Running Postgres in the Cloud: A Walkthrough

Tom Rieger – Senior Client Engineering

July 25, 2023

EDB

Fact #1 – the world is going cloud

Latest version of my 4 years in running the research will be out soon

www.platform3solutions.com

Is your organization taking a 'cloud first' perspective in all future projects?

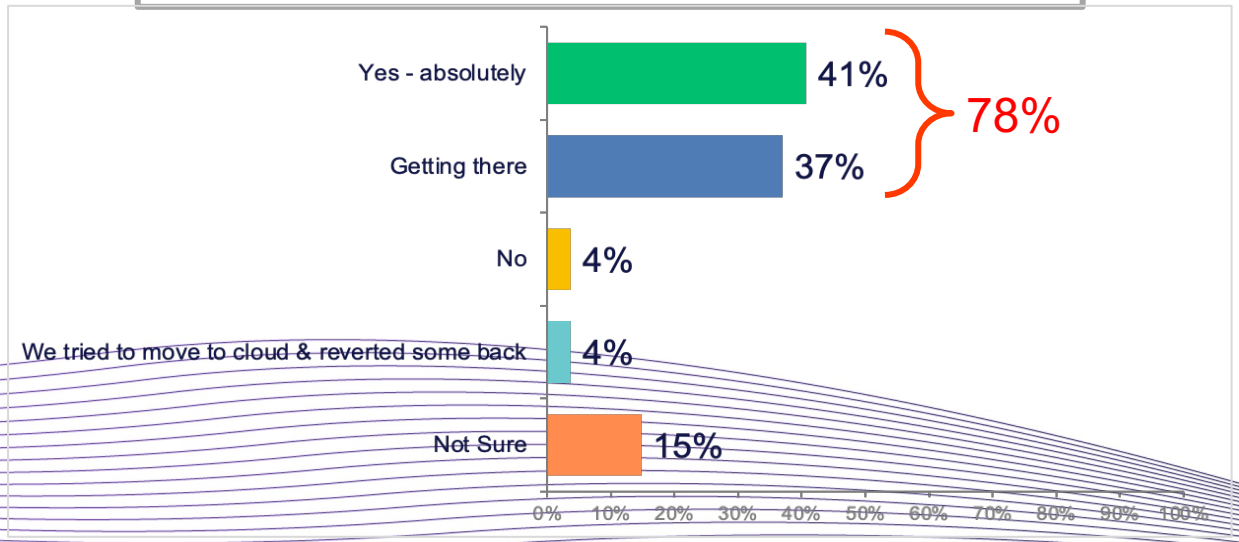

AGENDA: What we will review today

- Why is Postgres the primary database-of-choice to the cloud?
- The challenges in making the right decisions in the cloud – and where there are still gaps
- Using a best-in-class offering to bring forward an optimized environment
- The current offerings – what they are…and what they are not

EDB

AGENDA: What we will review today

- Why is Postgres the primary database-of-choice to the cloud?
- The challenges in making the right decisions in the cloud – and where there are still gaps
- Using a best-in-class offering to bring forward an optimized environment
- The current offerings – what they are…and what they are not

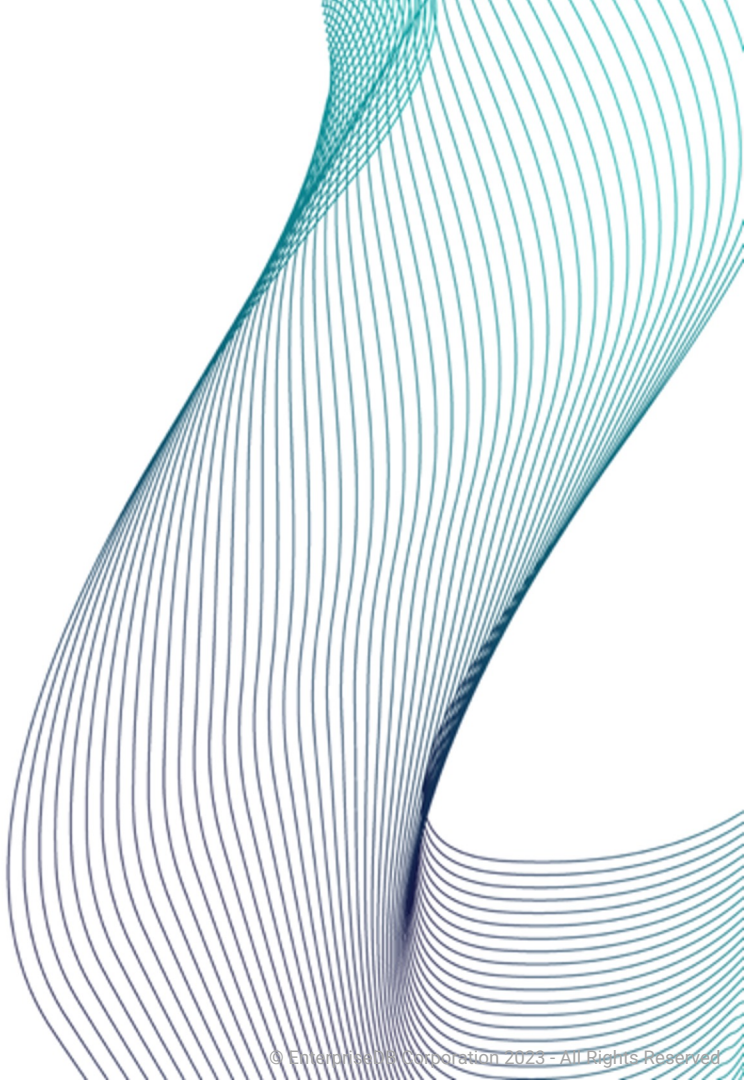

EDB

Fact #2: The world wants off proprietary

Latest version of my 4 years in running the research will be out soon

www.platform3solutions.com

Please complete this sentence - "We are looking to reduce our spend with…"

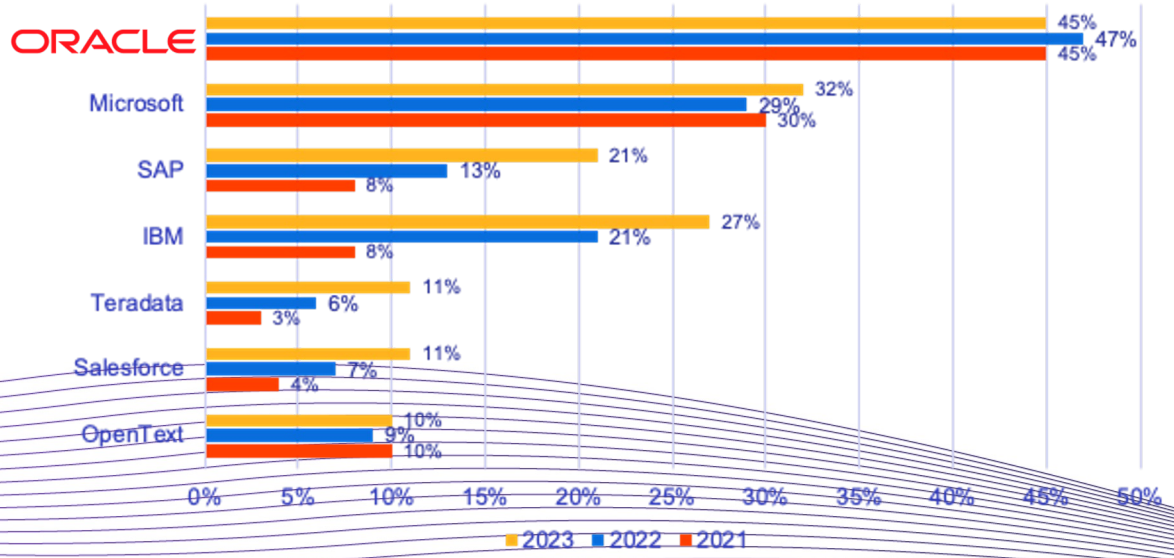

Knowledge Base of Relational and NoSQL Database Management Systems

Select a ranking

- Complete ranking
- Relational DBMS
- Key-value stores
- Document stores
- Time Series DBMS
- Graph DBMS
- Search engines
- Object oriented DBMS
- RDF stores
- Wide column stores
- Multivalue DBMS
- Vector DBMS
- Native XML DBMS
- Spatial DBMS
- Event Stores
- Content stores
- Navigational DBMS

Special reports

- Ranking by database model
- Open source vs. commercial

DB-Engines Ranking - Trend Popularity

The DB-Engines Ranking ranks database management systems according to their popularity.

Read more about the method of calculating the scores.

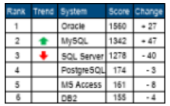

ranking table
July 2023

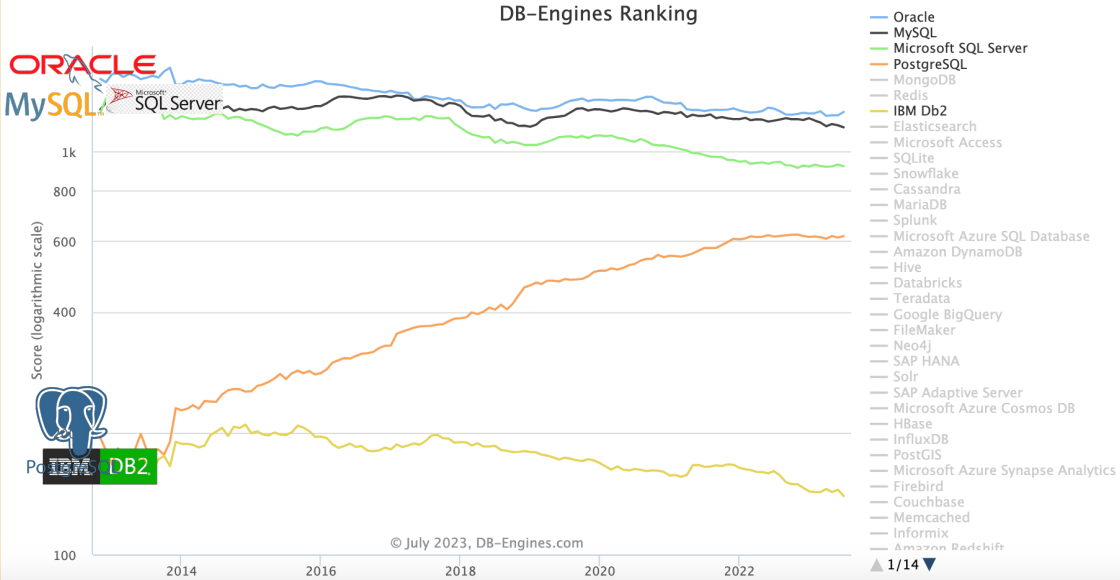

Why open-source

Latest version of my 4 years in running the research will be out soon

www.platform3solutions.com

Why do you think open-source is thriving?

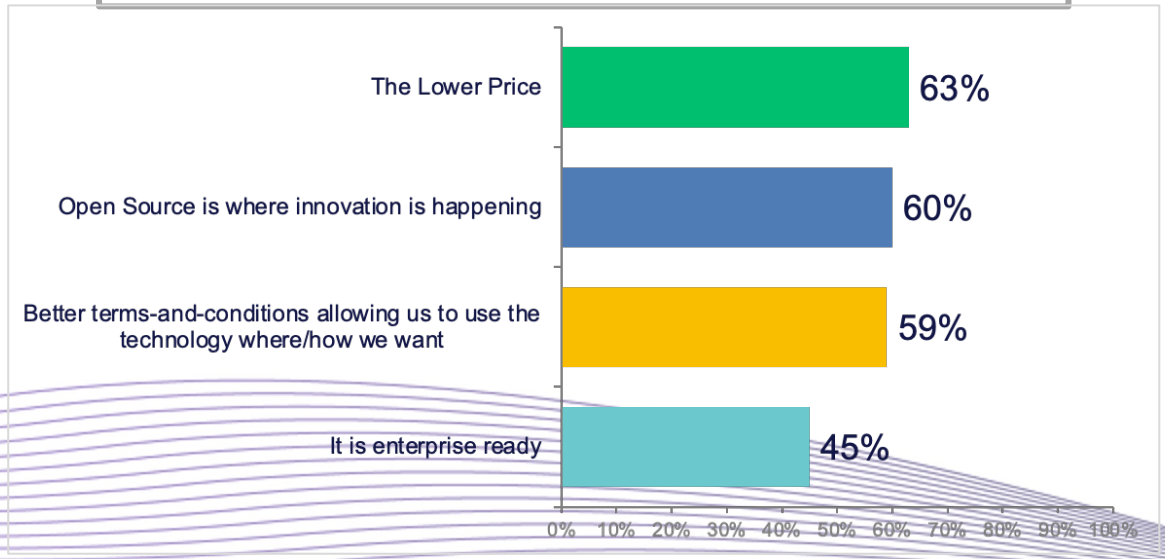

…but also because where developers are going

2023 Developer Survey

https://survey.stackoverflow.co/2023/#most-popular-technologies-database

Databases

This year, PostgreSQL took over the first place spot from MySQL. Professional Developers are more likely than those learning to code to use PostgreSQL (50%) and those learning are more likely to use MySQL (54%).

MongoDB is used by a similar percentage of both Professional Developers and those learning to code and it's the second most popular database for those learning to code (behind MySQL).

76,634

All Respondents | Professional Developers | Learning to Code | Other Coders

76,634 responses

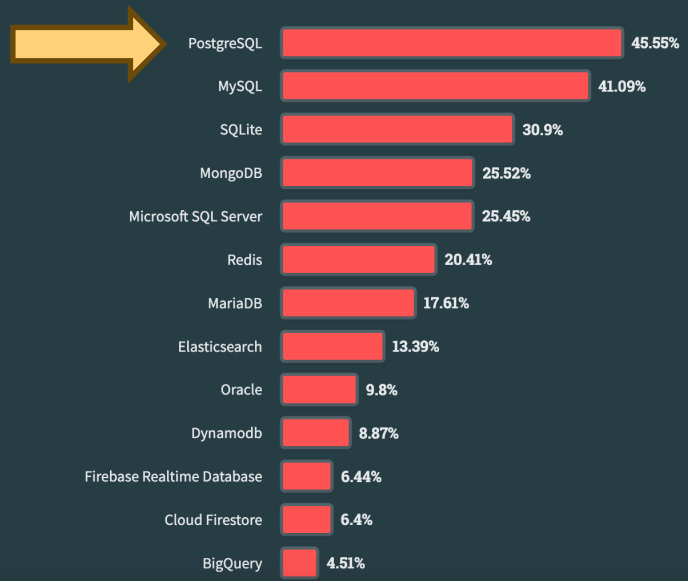

EDB

AGENDA: What we will review today

- Why is Postgres the primary database-of-choice to the cloud?
- The challenges in making the right decisions in the cloud – and how to measure
- Using a best-in-class offering to bring forward an optimized environment
- The current offerings – what they are…and what they are not

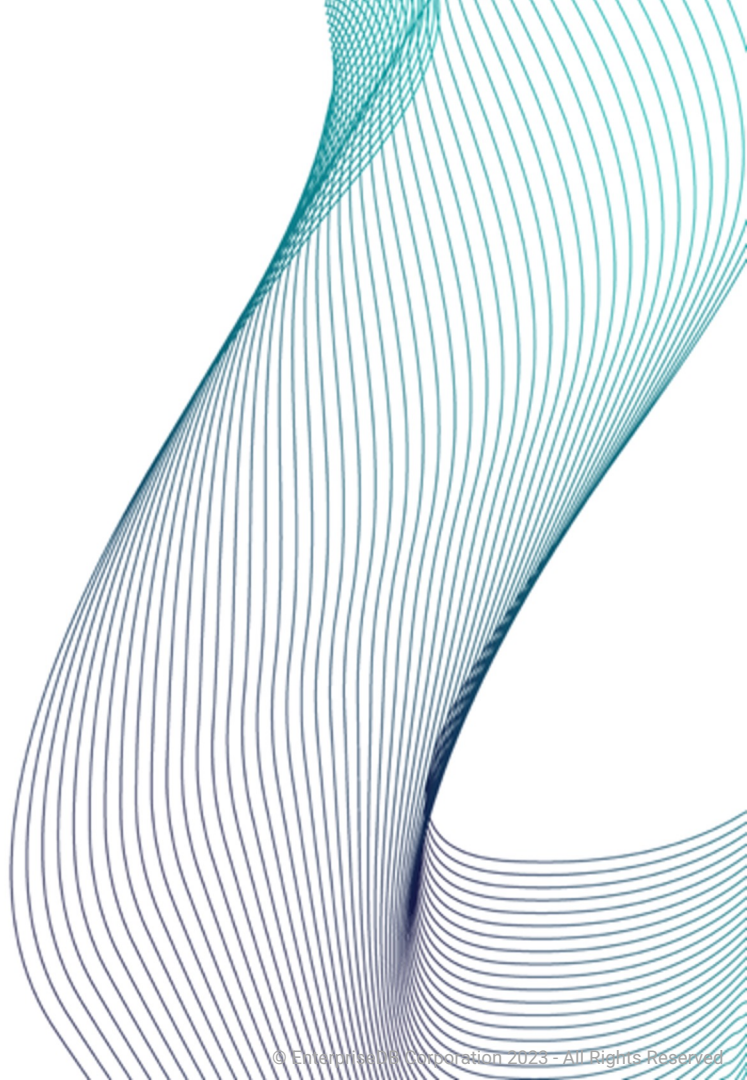

EDB

Workload constant - HammerDB

- HammerDB is open-source – evolving, community and free
- Standards based
  - TPC-C for transactional
  - TPC-H for analytical
- Works with all major databases – Oracle, SQLServer, DB2, MySQL, MariaDB, Postgres
- Can scale the workload by database size and concurrency
- Result -> TRANSACTIONS/MINUTE

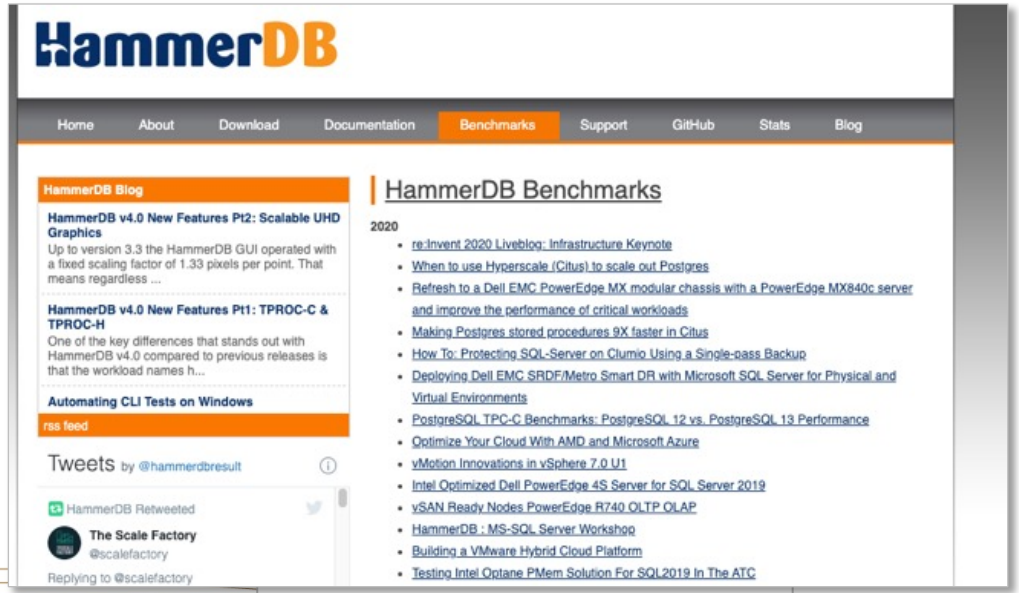

www.hammerdb.com

Quick Anatomy of HammerDB

- Order Processing Use-case
- Typical Transactional use case
  - Simple SQL
  - Large in volume
- Two primary ‘knobs’ to size the effort
  - Warehouses – database size
  - User Concurrency
- **Resulting measure:** Transactions-per-minute (TPM)

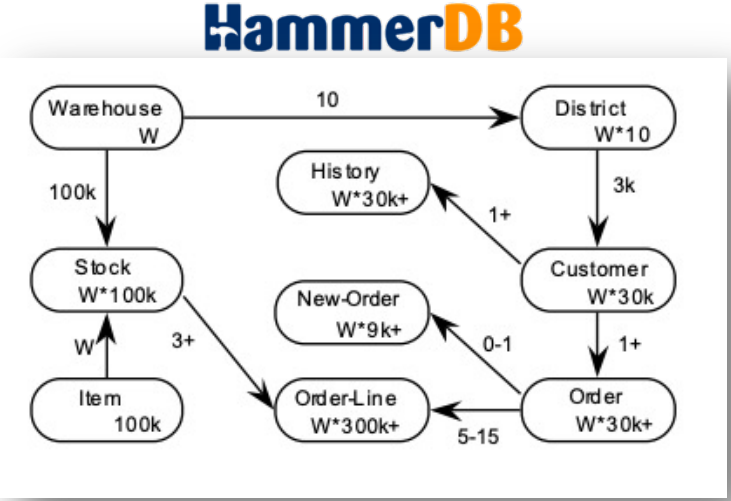

Action Type	Mix
SELECT	75%
INSERT	8%
UPDATES	16%
DELETE	1%

CONSTANT #1: Workload

To measure and compare, we need maximum ‘constants’ and minimal ‘variables.

Workload

HammerDB

300 warehouses (60 GB) – 96 concurrent users – 5 min. warmup – 20 min. run

EDB

Challenge #2 – What compute?

- Example:
  Amazon gives you 55 choices in compute with 8 cores.
- From $102 to $825 per month

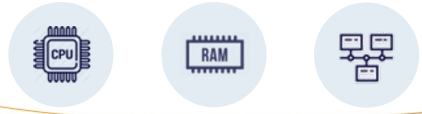

1 | fx | Category

Category	Categ	Size	VM Price per mo	vCP	RA	RAM/CPU Ratio	Instance Storage	Network Burst Bandwidth (Gbps)	EBS Burst Bandwidth (Mbps)	M	Mod	Family	Burst Fr	Burst Fr	Processor
General	A1	a1.2xlarge	$102.78	8	16	2	EBS-Only	Up to 10	??	AWS		Graviton			AWS Graviton Processor with 64-bit Arm Neoverse cores
Compute	C6g	**c6g.2xlarge**	$137.02	**8**	**16**	2	**EBS-Only**	**Up to 10**	**Up to 4,750**	AWS		Graviton2			AWS Graviton2 Processor with 64-bit Arm Neoverse cores
Compute	C7g	c7g.2xlarge	$152.79	8	16	2	EBS-Only	Up to 15	Up to 10	AWS		Graviton3			AWS Graviton3 processors
General	M6g	m6g.2xlarge	$154.47	8	32	4	EBS-Only	Up to 10	Up to 4,750	AWS		Graviton2			AWS Graviton2 Processor with 64-bit Arm Neoverse cores
Compute	C6g	**c6gd.2xlarge**	$154.76	**8**	**16**	2	**NVMe**	**Up to 10**	**Up to 4,750**	AWS		Graviton2			AWS Graviton2 Processor with 64-bit Arm Neoverse cores
Compute	C5a	c5a.2xlarge	$155.49	8	16	2	EBS-Only	Up to 10	Up to 3,170	AMD	7R32	EPYC		3.30	2nd generation AMD EPYC 7002 series processors (AMD EPYC 7R32) running at frequencies up to 3.3 GHz
Compute	C6a	c6a.2xlarge	$162.54	8	16	2	EBS-Only	Up to 12.5	Up to 6.6	AMD	7R13	EPYC		3.60	Up to 3.6 GHz 3rd generation AMD EPYC processors (AMD EPYC 7R13)
General	T3	t3.2xlarge	$166.81	8	32	4	EBS-Only	Up to 5		Intel	8175M	Skylake		3.10	Up to 3.1 GHz Intel Xeon Scalable processor (Skylake 8175M or Cascade Lake 8259CL)
Compute	C5	c5.2xlarge	$171.55	8	16	2	EBS-Only	Up to 10	Up to 4,750	Intel	8223CL	Cascade Lake	3.10	3.40	8000 series (Skylake 8124M) processor with a sustained all core Turbo frequency of up to 3.4GHz, and
General	M5a	**m5a.2xlarge**	$173.01	**8**	**32**	4	**EBS-Only**	**Up to 10**	**Up to 2,880**	**AMD**	**7571**	**EPYC**	**2.50**		AMD EPYC 7000 series processors (AMD EPYC 7571) with an all core turbo clock speed of 2.5 GHz
Compute	C5a	c5ad.2xlarge	$173.74	8	16	2	NVMe	up to 10	up to 3,170	AMD	7R32	EPYC		3.30	2nd generation AMD EPYC 7002 series processors (AMD EPYC 7R32) running at frequencies up to 3.3 GHz
Compute	C6gn	c6gn.2xlarge	$175.16	8	16	2	EBS-Only	Up to 25	Up to 9.5	AWS		Graviton2			AWS Graviton2 Processor with 64-bit Arm Neoverse cores
Compute	C6i	c6i.2xlarge	$180.60	8	16	2	EBS-Only	Up to 12.5	Up to 10	Intel	8375C	Ice Lake	2.90	3.50	Up to 3.5 GHz 3rd generation Intel Xeon Scalable processors (Ice Lake 8375C)
General	M6g	m6gd.2xlarge	$182.14	8	32	4	NVMe	Up to 10	Up to 4,750	AWS		Graviton2			AWS Graviton2 Processor with 64-bit Arm Neoverse cores
General	M6a	m6a.2xlarge	$183.58	8	32	4	EBS-Only	Up to 12.5	Up to 6.6	AMD	7R13	EPYC		3.60	Up to 3.6 GHz 3rd generation AMD EPYC processors (AMD EPYC 7R13)
General	M5	m5.2xlarge	$193.45	8	32	4	EBS-Only	Up to 10	Up to 4,750	Intel	8175M	Skylake		3.10	Up to 3.1 GHz Intel Xeon Scalable processor (Skylake 8175M or Cascade Lake 8259CL)
Compute	C5	c5d.2xlarge	$193.45	8	16	2	NVMe	Up to 10	Up to 4,750	Intel	8223CL	Cascade Lake	3.10	3.40	8000 series (Skylake 8124M) processor with a sustained all core Turbo frequency of up to 3.4GHz, and
General	M4	m4.2xlarge	$199.00	8	32	4	EBS-only	High	1000	Intel	E5-2686	Broadwell	**2.30**	**2.40**	Up to 2.4 GHz Intel Xeon Scalable Processor (Broadwell E5-2686 v4 or Haswell E5-2676 v3)
Compute	C4	c4.2xlarge	$202.21	8	15	2	EBS-Only	High	1000	Intel	E5-2666	Haswell		2.90	Up to 2.9 GHz Intel Xeon Scalable Processor (Haswell E5-2666 v3)
Memory	R6g	r6g.2xlarge	$203.09	8	64	8	EBS-Only	Up to 10	Up to 4,750	AWS		Graviton2			AWS Graviton2 Processor
General	M6i	m6i.2xlarge	$203.98	8	32	4	EBS-Only	Up to 12.5	Up to 10	Intel	8375C	Ice Lake	2.90	3.50	Up to 3.5 GHz 3rd generation Intel Xeon Scalable processors (Ice Lake 8375C)
Compute	C6i	c6id.2xlarge	$204.04	8	16	2	VMe SSD	Up to 12.5	Up to 10	Intel	8375C	Ice Lake	2.90	3.50	Up to 3.5 GHz 3rd generation Intel Xeon Scalable processors (Ice Lake 8375C)
General	M5a	**m5ad.2xlarge**	$207.32	**8**	**32**	4	**NVMe**	**Up to 10**	**Up to 2,880**	**AMD**	**7571**	**EPYC**	**2.50**		AMD EPYC 7000 series processors (AMD EPYC 7571) with an all core turbo clock speed of 2.5 GHz
Compute	C5n	c5n.2xlarge	$217.54	8	21	3	EBS-Only	Up to 25	Up to 4,750	Intel	8124	Skylake	3.00		3.0 GHz Intel Xeon Platinum processors (Skylake 8124)
Memory	R5a	r5a.2xlarge	$227.76	8	64	8	EBS-Only	Up to 10	Up to 2,880	**AMD**	**7571**	**EPYC**	**2.50**		AMD EPYC 7000 series processors (AMD EPYC 7571) with an all core turbo clock speed of 2.5 GHz
General	M5	m5d.2xlarge	$227.76	8	32	4	NVMe	Up to 10	Up to 4,750	Intel	8175M	Skylake		3.10	Up to 3.1 GHz Intel Xeon Scalable processor (Skylake 8175M or Cascade Lake 8259CL)
Memory	R6g	r6gd.2xlarge	$232.14	8	64	8	NVMe	Up to 10	Up to 4,750	AWS		Graviton2			AWS Graviton2 Processor
General	M5n	m5n.2xlarge	$239.44	8	32	4	EBS-Only	Up to 25	Up to 4,750	Intel	8259CL	Cascade Lake	3.10	3.50	frequency of 3.1 GHz and maximum single core turbo frequency of 3.5 GHz
General	M6i	m6id.2xlarge	$240.10	8	32	4	VMe SSD	Up to 12.5	Up to 10	Intel	8375C	Ice Lake	2.90	3.50	Up to 3.5 GHz 3rd generation Intel Xeon Scalable processors (Ice Lake 8375C)
Memory	R6a	r6a.2xlarge	$240.94	8	64	8	EBS-Only	Up to 12.5	Up to 6.6	AMD	7R13	EPYC		3.60	Up to 3.6 GHz 3rd generation AMD EPYC processors (AMD EPYC 7R13)
Storage	H1	h1.2xlarge	$254.04	8	32	4	HDD	Up to 10	??	Intel	E5-2686	**Broadwell**	**2.30**		2.3 GHz Intel Xeon Scalable Processor (Broadwell E5 2686 v4)
Memory	R5	r5.2xlarge	$254.04	8	64	8	EBS-Only	up to 10	Up to 4,750	Intel	8175M	Skylake		3.10	Up to 3.1 GHz Intel Xeon® Platinum 8000 series processors (Skylake 8175M or Cascade Lake 8259CL)
Memory	R5a	r5ad.2xlarge	$264.26	8	64	8	NVMe	Up to 10	Up to 2,880	**AMD**	**7571**	**EPYC**	**2.50**		AMD EPYC 7000 series processors (AMD EPYC 7571) with an all core turbo clock speed of 2.5 GHz
Memory	R6i	r6i.2xlarge	$267.72	8	64	8	EBS-Only	Up to 12.5	Up to 10	Intel	8375C	Ice Lake	2.90	3.50	Up to 3.5 GHz 3rd generation Intel Xeon Scalable processors (Ice Lake 8375C)
Memory	R4	r4.2xlarge	$269.81	8	61	8	EBS-Only	Up to 10		**Intel**	**E5-2686**	**Broadwell**	**2.30**	**2.30**	Intel Xeon scalable (Broadwell E5-2686 v4) processors up to 2.3 Ghz
General	M5n	m5dn.2xlarge	$273.75	8	32	4	NVMe	Up to 25	Up to 4,750	Intel	8259CL	Cascade Lake	3.10	3.50	frequency of 3.1 GHz and maximum single core turbo frequency of 3.5 GHz
Memory	R5	r5d.2xlarge	$289.81	8	64	8	NVMe	up to 10	Up to 4,750	Intel	8175M	Skylake		3.10	Up to 3.1 GHz Intel Xeon® Platinum 8000 series processors (Skylake 8175M or Cascade Lake 8259CL)
Memory	R5n	r5n.2xlarge	$300.03	8	64	8	EBS-Only	Up to 25	Up to 4,750	Intel	8259CL	Cascade Lake	3.10	3.50	frequency of 3.1 GHz and maximum single core turbo frequency of 3.5 GHz
Memory	R5b	r5b.2xlarge	$301.51	8	64	8	EBS-Only	up to 10	Up to 10,000	Intel	8259CL	Cascade Lake	3.10	3.50	GHz and maximum single core turbo frequency of 3.5 GHz
Memory	R6i	r6id.2xlarge	$305.22	8	64	8	NVMe	Up to 12.5	Up to 10	Intel	8375C	Ice Lake	2.90	3.50	Up to 3.5 GHz 3rd generation Intel Xeon Scalable processors (Ice Lake 8375C)
Memory	X2gd	x2gd.2xlarge	$330.54	8	128	16	NVMe	Up to 10	Up to 4.75	AWS		Graviton2			AWS Graviton2 Processor
General	M5zn	m5zn.2xlarge	$334.27	8	32	4	EBS-Only	Up to 25	3170	Intel	8252C	Cascade Lake		4.50	4.5 GHz
Memory	R5n	r5dn.2xlarge	$336.53	8	64	8	NVMe	Up to 25	Up to 4,750	Intel	8259CL	Cascade Lake	3.10	3.50	frequency of 3.1 GHz and maximum single core turbo frequency of 3.5 GHz
Storage	I3	i3.2xlarge	$343.83	8	61	8	NVMe	Up to 10	??	**Intel**	**E5-2686**	**Broadwell**	**2.30**		Intel Xeon Scalable Processors (Broadwell E5-2686 v4) with base frequency of 2.3 GHz
Storage	I4i	i4i.2xlarge	$357.70	8	64	8	AWS	Up to 12	Up to 10	Intel	8375C	Ice Lake	2.90	3.50	Up to 3.5 GHz 3rd generation Intel Xeon Scalable Processors (Ice Lake 8375C)
Memory	z1d	**z1d.2xlarge**	$374.49	**8**	**64**	8	**NVMe**	**Up to 10**	??	**Intel**	**8151**	**Skylake**		**4.00**	GHz
Storage	Im4gn	Im4gn.2xlarge	$374.69	8	32	4	3750 A	Up to 25	Up to 9.5	AWS		Graviton2			AWS Graviton2
Storage	I3en	I3en.2xlarge	$492.02	8	64	8	NVMe	Up to 25	??	Intel		Skylake	3.10		3.1 GHz all core turbo Intel® Xeon® Scalable (Skylake) processors
Storage	D3	d3.2xlarge	$505.16	8	64	8	HDD	Up to 15	1700	Intel	8259CL	Cascade Lake	3.10	3.10	Up to 3.1 GHz 2nd Generation Intel® Xeon® Scalable (Cascade Lake) processors
Storage	D3en	d3en.2xlarge	$532.17	8	32	4	HDD	Up to 25	1700	Intel		Cascade Lake	3.10		3.1 GHz all core turbo 2nd Generation Intel® Xeon® Scalable (Cascade Lake) processors
Storage	D2	d2.2xlarge	$645.32	8	61	8	2000 H	High	??	Intel	**E5-2676**	**Haswell**	**2.40**		Intel Xeon Scalable Processors (Haswell E5-2676 v3) at 2.4 Ghz
Memory	x2iedn	x2iedn.2xlarge	$823.76	8	256	32	NVMe	Up to 25	Up to 20	Intel	8375C	Ice Lake	2.90	3.50	Up to 3.5 GHz 3rd generation Intel Xeon Scalable processors (Ice Lake 8375C)
Memory	x2iezn	x2iezn.2xlarge	$824.61	8	256	32	EBS-Only	Up to 25	3	Intel	8252C	Cascade Lake		4.50	Up to 4.5 GHz 2nd generation Intel Xeon Scalable processors (Cascade Lake 8252C)
Memory	X1e	x1e.2xlarge	$825.63	8	244	31	1 x 240	Up to 10	1000	Intel	8880	Haswell	**2.30**		2.3 GHz Intel Xeon Scalable Processor (Haswell E7-8880 v3)

$102

$825

CONSTANT #2: Compute

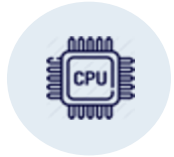

To measure and compare, we need maximum ‘constants’ and minimal ‘variables.

Workload

CPU/RAM

HammerDB

300 warehouses (60 GB) – 96 concurrent users – 5 min. warmup – 20 min. run

AWS RDS	AWS AURORA	AWS BIG ANIMAL	AZURE FLEX SERVER	AZURE BIG ANIMAL
R6i-4xlarge Intel Ice Lake 16 cores 64GB RAM	R6i-4xlarge Intel Ice Lake 16 cores 64GB RAM	R6i-4xlarge Intel Ice Lake 16 cores 64GB RAM	E16ds_v4 (Intel® Ice Lake or the Intel® Cascade Lake processors) 16 cores 64GB RAM	E16s_v5 (Intel Ice Lake) 16 mores 128GB RAM

EDB

Challenge #3 – What storage

Google Cloud

- No performance control
- Quoted for **600GB**

Type	MAX 'sustained' IOPS	MAX 'sustained' Throughput
Standard	*Read: 450 Write: 900*	*Read: 72 MB/sec Write: 72 MB/sec*
Balanced	*Read: 3,600 Write: 3,600*	*Read: 72 MB/sec Write: 72 MB/sec*
SSD	*Read: 15,000 Write: 15,000*	*Read: 240 MB/sec Write: 240 MB/sec*

https://cloud.google.com/compute/docs/disks/?&_ga=2.62767211.-658045011.1598996595#pdperformance

aws

- All about which level and what you are willing to pay – **600GB sizing**

Type	MAX 'sustained' IOPS	MAX Throughput
Optimized HHD (st1)		*MAX 147 MB/s max*
General Purpose SSD (gp2)	*3000*	
General Purpose SSD (gp3)	3000	500 MB/sec
Provisioned IOPS SSD (io1)	MAX 30,000	
Provisioned IOPS SSD (io2)	MAX 64,000	

https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volume-types.html?icmpid=docs_ec2_console

Azure

- Azure has 'binary' pricing = 128GB, 256, 512GB, etc.
- **Assume 1024 GB**
- *Performance limited by VM

Type	MAX 'sustained' IOPS	MAX Throughput
Standard HDD	*500*	*60 MB/sec*
Standard SSD	*500*	*60 MB/sec*
Premium SSD	*5000*	*200 MB/sec*
Ultra Disk	51,200*	768 MB/sec*

https://docs.microsoft.com/en-us/azure/virtual-machines/disks-types

EDB

Challenge #3 – Storage cost is confusing

((# of GB * $0.14746)+(# of IOPS * $0.06132)+(# of MBps * $0.39566))*12

Azure

- Top performing storage: **Ultra Disk**
- You pay for it in discrete increments: 256GB, 512GB, 1024GB, etc
  - i.e. Even if you want a volume of 300GB you pay for 512GB
- You configure:
  - IOPS = operations/sec
  - Throughput = MBps

- Pricing (central-us)
  - $0.14746/mon per GB
  - $0.06132/mon per IOPS
  - $0.39566/mon per MBps

Good Willing Hunting

- Example of VMs and their 'limits':

Name	Size	"Max uncached disk throughput IOPS/MBps"
E4ds_v4	4 CPU/32 GB	6,400 / 96
E8ds_v4	8 CPU/64 GB	12,800 / 192
E16ds_v4	16 CPU/128 GB	25,600 / 384
E32ds_v4	32 CPU/256 GB	51,200 / 768

EDB

https://azure.microsoft.com/en-us/pricing/details/managed-disks/

Challenge #3.1 – Storage cost is confusing

((# of GB * $0.10)+((730 hours * 60 mins * 60 sec * # of IOPS that second) * $0.0000002))*12

- Inputs/Outputs per Second/Volume
  - Gp2 has “Up to” 3 IOPS/GB
  - Gp3 has them - ”Up to” 16,000
  - IO1 and IO2 has them – “Up to” 64,000
  - IO2 Block Express has them – “Up to” 256,000
- RDS supports gp2, gp3 and io1
- Aurora is based on ”Solid state drives”

Amazon Aurora PostgreSQL

- Aurora Pricing
  - $0.10/mon per GB
  - $0.0000002 /mon per IOPS
- Example of VMs and their ‘limits’:

Name	Size	Storage Bandwidth (Gbps)
R6i.xlarge	4 CPU/32 GB	“Up to” 10
R6i.2xlarge	8 CPU/64 GB	“Up to” 10
R6i.4xlarge	16 CPU/128 GB	“Up to” 10
R6i.8xlarge	32 CPU/256 GB	10

CONSTANT #3: Storage

To measure and compare, we need maximum ‘constants’ and minimal ‘variables.

Workload

CPU/RAM

STORAGE

HammerDB

300 warehouses (60 GB) – 96 concurrent users – 5 min. warmup – 20 min. run

AWS RDS	AWS AURORA	AWS BIG ANIMAL	AZURE FLEX SERVER	AZURE BIG ANIMAL
R6i-4xlarge Intel Ice Lake 16 cores 64GB RAM	R6i-4xlarge Intel Ice Lake 16 cores 64GB RAM	R6i-4xlarge Intel Ice Lake 16 cores 64GB RAM	E16ds_v4 (Intel® Ice Lake or the Intel® Cascade Lake processors) 16 cores 64GB RAM	E16s_v5 (Intel Ice Lake) 16 mores 128GB RAM
300GB - io1 storage at 7000 IOPS	I have no idea	300GB - io2 storage at 7000 IOPS	Provisioned 2TB to get 7500 IOPS	300GB - Ultradisk at 7000 IOPS

AWS RDS: IOPS, Tuning and Cost

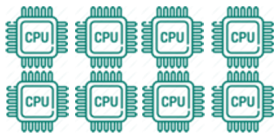

8 CPU
64 GB RAM

IO1 IOPS	Annual Storage Cost	Annual Total Cost
30,000	$36,900	$45,531
20,000	$24,900	$32,331
15,000	$18,900	$25,731
10,000	$12,900	$19,131
7,000	$9,300	$15,171
5,000	$6,900	$12,531
4,000	$5,700	$11,211

IOPS….IOPS….who needs an IOPS!!!!!!????

8 CPU
64 GB RAM

- ONLY VARIABLE IS STORAGE TYPE AND IOPS
- Orange line: Cost/TPM
- Blue Bar: TPMs
- NET NET: There is a point of diminishing return

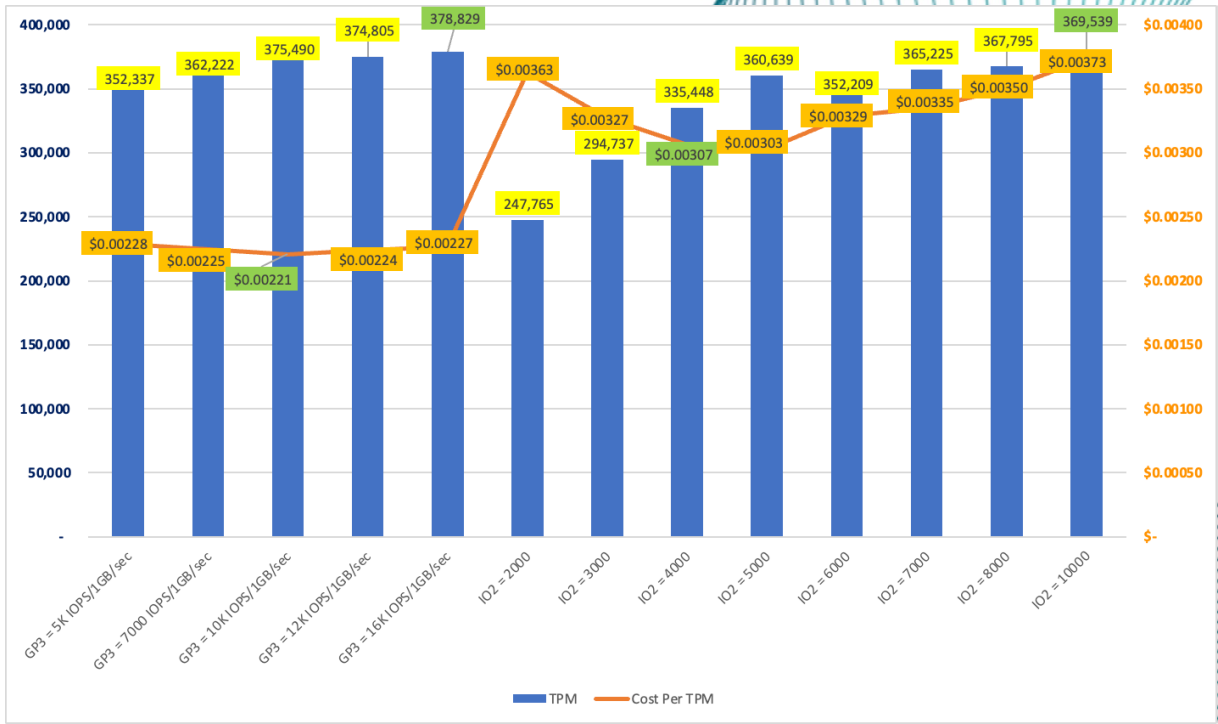

CONSTANT #4: Database

To measure and compare, we need maximum ‘constants’ and minimal ‘variables.

Workload

CPU/RAM

STORAGE

DATABASE

HammerDB 600 warehouses (120 GB) – 96 concurrent users – 5 min. warmup – 20 min. run				
AWS RDS	**AWS AURORA**	**AWS BIG ANIMAL**	**AZURE FLEX SERVER**	**AZURE BIG ANIMAL**
R6i-4xlarge Intel Ice Lake 16 cores 64GB RAM	R6i-4xlarge Intel Ice Lake 16 cores 64GB RAM	R6i-4xlarge Intel Ice Lake 16 cores 64GB RAM	E16ds_v5 Intel Ice Lake 16 cores 64GB RAM	E16s_v5 Intel Ice Lake 16 mores 128GB RAM
300GB - io1 storage at 7000 IOPS	I have no idea	300GB - io2 storage at 7000 IOPS	Provisioned 2TB to get 7500 IOPS	300GB - Ultradisk at 7000 IOPS
Community Postgres v15.3	Postgres ‘Compatible’ v15.3	Community Postgres v15.3	Community Postgres v15.3	Community Postgres v15.3

EDB

Challenge #4 – How to get to “optimized”

- Components
  - Picking the optimal CPU/RAM
  - Picking the optimal storage
- Postgres
  - Latest version
  - Configuration settings

OBJECTIVE: LEAST COST PER TPM

But what abou

Microsoft Volume Licensing Service Level Agreement for Microsoft Online Services (Worldwide Englis

Virtual Machines

Additional Definitions:

"**Availability Set**" refers to two or more Virtual Machines deployed across different Fault D

"**Availability Zone**" is a fault-isolated area within an Azure region, providing redundant po

"**Azure Dedicated Host**" provides physical servers that host one or more Azure virtual mac
required for any SLA.

"**Data Disk**" is a persistent virtual hard disk, attached to a Virtual Machine, used to store a

"**Dedicated Host Group**" is a collection of Azure Dedicated Hosts deployed within an Azure
point of failure.

"**Fault Domain**" is a collection of servers that share common resources such as power and

"**Operating System Disk**" is a persistent virtual hard disk, attached to a Virtual Machine, used to store the Virtual M

"**Single Instance**" is defined as any single Microsoft Azure Virtual Machine that either is not deployed in an Availabi
deployed in an Availability Set.

"**Virtual Machine**" refers to persistent instance types that can be deployed individually or as part of an Availability S
Group. A virtual machine can be deployed in a multi-tenant Azure or in an isolated, single-tenant er
Dedicated Hosts.

"**Virtual Machine Connectivity**" is bi-directional n tual Machine and other IP addresse
protocols in which the Virtual Machine is configu addresses can be IP addresses in the
Virtual Machine, IP addresses within the same virtu achine or public, routable IP address

Monthly Uptime Calculation and Service Levels for Virtual Machines in Availability Zones

"**Maximum Available Minutes**" is the total accumulated minutes during a billing month that have two or more
or more Availability Zones in the same region. Maximum Available Minutes is measured from when at least tw
Availability Zones in the same region have both been started resultant from action initiated by Customer to the time Customer has initiated an
action that would result in stopping or deleting the Virtual Machines.

"**Downtime**" is the total accumulated minutes that are part of Maximum Available Minutes that have no Virtual Machine Connectivity in the
region.

"**Monthly Uptime Percentage**" for Virtual Machines in Availability Zones is calculated as Maximum Available Minutes less Downtime divided
by Maximum Available Minutes in a billing month for a given Microsoft Azure subscription. Monthly Uptime Percentage is represented by the
following formula:

$$Monthly\ Uptime\ \% = \frac{(Maximum\ Available\ Minutes\text{-}Downtime)}{Maximum\ Available\ Minutes} \times 100$$

Service Credit:

The following Service Levels and Service Credits are applicable to Customer's use of Virtual Machines deployed across two or more Availability
Zones in the same region:

	Monthly Uptime Percentage	Service Credit
	< 99.99%	10%
	< 99%	25%
	< 95%	100%

aws

Contact Us Support English My Account Sign In to the Console

Products Solutions Pricing Documentation Learn Partner Network AWS Marketplace Customer Enablement Events Explore More

Amazon Compute Service Level Agreement

Last Updated: May 25, 2022

This Amazon Compute Service Level Agreement (this "SLA") is a policy governing the use of Amazon Elastic Compute Cloud ("Amazon EC2")* and applies separately to each account using Amazon EC2. In the event of a conflict between the terms of this SLA and the terms of the AWS Customer Agreement or other agreement with us governing your use of our Services (the "Agreement"), the terms and conditions of this SLA apply, but only to the extent of such conflict. Capitalized terms used herein but not defined herein shall have the meanings set forth in the Agreement.

*For purposes of this SLA, Amazon EC2 includes any Amazon Elastic Graphics, A zon Elas Inferenc and Elastic IP Address resources purchased with the relevant Amazon EC2 instance(s).

aws

SLAs

AWS makes two SLA commitments for Amazon EC2: (1) a Region-Level SLA verns Amazon EC2 d yed across multiple AZs or regions, and (2) an Instance-Level SLA that governs Amazon EC2 instances individually.

Region-Level SLA

For Amazon EC2 with all running instances deployed concurrently across two or more AZs in the same region (or at least two regions if there is only one AZ in a given region), AWS will use commercially reasonable efforts to make Amazon EC2 available for each AWS region with a Monthly Uptime Percentage of at least 99.99%, in each case during any monthly billing cycle (the "Region-Level SLA"). In the event Amazon EC2 does not meet the Region-Level SLA, you will be eligible to receive a Se bed below.

Monthly Uptim

Less than 99.9

99.99%

52 minutes
33 seconds

Docs Support English Consol

Contact Us

Service Level Agreement

ment under which Google has agreed to provide Google Cloud licable, the "Agreement"), the Covered Service will provide a o Customer as follows (the "Service Level Objective" or "SLO"):

Google Cloud

	Monthly Uptime Percentage
Instances in Multiple Zones	>= 99.99%
A Single Instance	>= 99.5%
Load balancing	>= 99.99%

If Google does not meet the SLO, and if Customer meets its obligations under this SLA, Customer will be eligible to receive the Financial Credits described below. Monthly Uptime Percentage and Financial Credit are determined on a calendar month basis per Project or, for a Single Instance, per instance. This SLA states Customer's sole and exclusive remedy for any failure by Google to meet the SLO. Capitalized terms used in this SLA, but not defined in this SLA, have the meaning set forth in the Agreement. If the Agreement authorizes the resale or supply of Google Cloud Platform under a Google Cloud partner or reseller program, then all references to Customer in this SLA mean Partner or Reseller (as applicable), and any Financial Credit(s) will only apply for impacted Partner or Reseller order(s) under the Agreement.

Databases in the cloud today

Challenges for enterprises

Behind with bug fixes, platform updates, and new capabilities

Don’t provide the control, capabilities, best components and performance tuning enterprises need

Lack easy migration for existing on-prem databases to cloud

AGENDA: What we will review today

- Why is Postgres the primary database-of-choice to the cloud?
- The challenges in making the right decisions in the cloud – and where there are still gaps
- Using a best-in-class offering to bring forward an optimized environment
- The current offerings – what they are…and what they are not

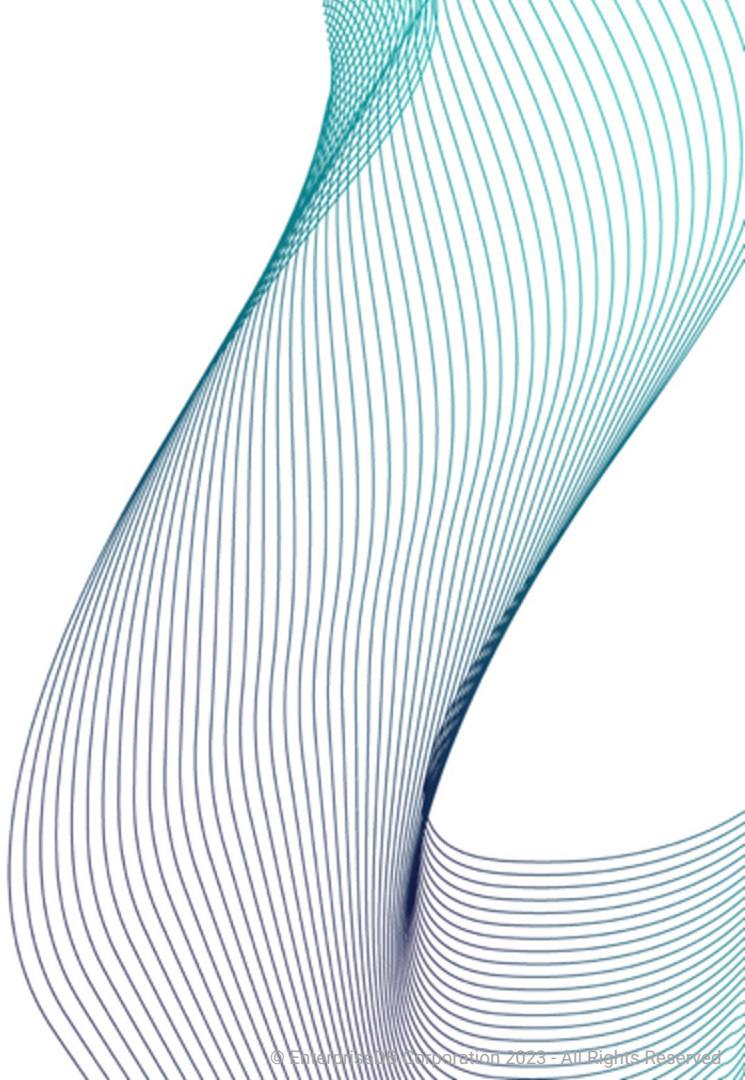

EDB

The BigAnimal advantage

BIG ANIMAL

Deep Postgres Expertise

Compatible With Oracle

Availability Options

Cloud Choice

EDB

Control over what matters for your workload

BigAnimal is run like a platform-as-a-service, not simply infrastructure + automation

EDB ensures smooth running operations:

- 24x7 monitoring
- Infrastructure availability and performance
- High availability and backups
- Security and compliance
- OS, database patching – Including major version upgrades
- Network configuration and security

Customer optimizes the database for their enterprise needs:

- Selecting appropriate compute resources (VM and disk) and replicants to fit workload
- Data modeling
- Query performance
- Securely managing database credentials
- Maximized control of Postgres configuration

BigAnimal running in *your* cloud account

Coming July 2023

Azure aws Google Cloud

Key Benefits

- Maximum Control and transparency
- Minimized burden
- Burn down your commit contract
- Continual innovation
- Adheres to your quota limits

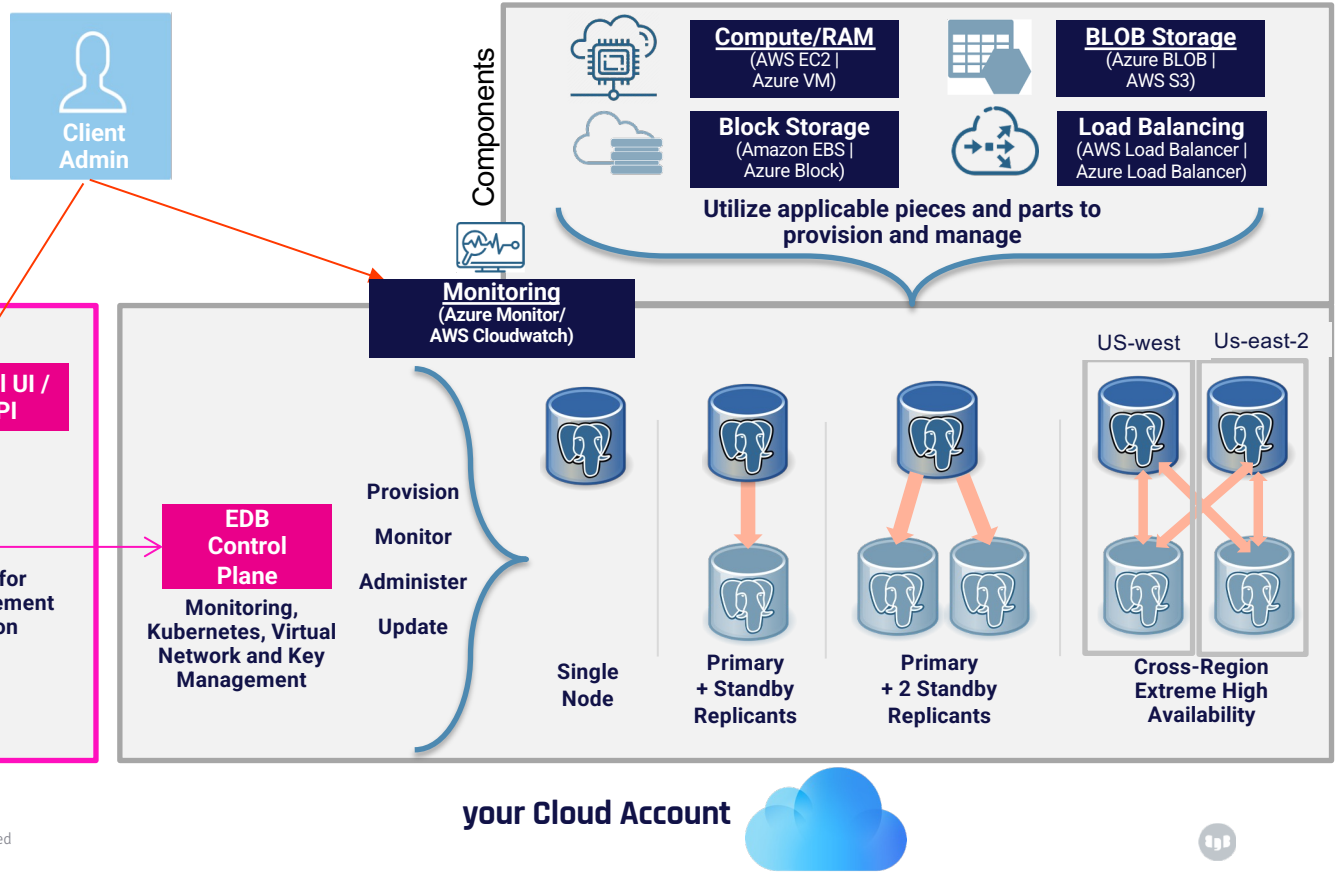

BIGANIMAL

BigAnimal in EDB account

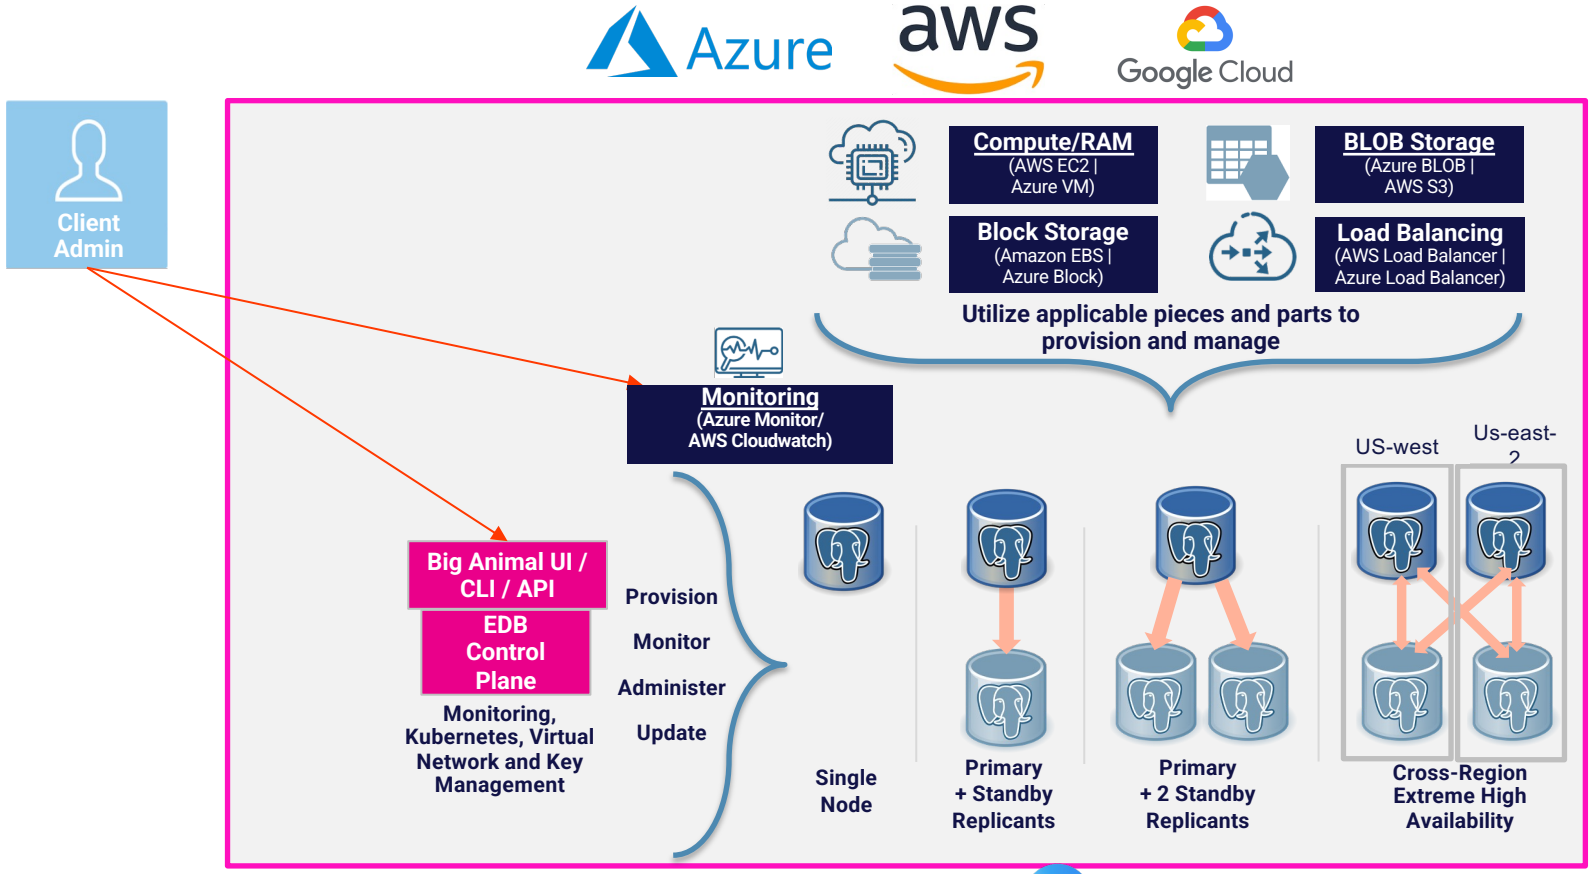

Types of clusters and replicants

Single node

- Single zone
- Generally for Non-production environments that can tolerate zonal outages.

Master + 1 replicant

- 1 secondary replicant within a region across two zones.
- Automatically disables sync replication during maintenance operations to ensure write availability.
- Can change from sync to async on per-session/per-transaction basis

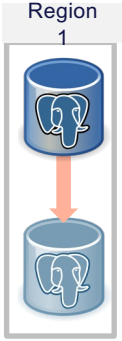

Master + 2 replicants

- 2 secondary replicants within a region across three zones
- Replication is sync to one standby and async to the other
- Can modify sync behavior on a per-transaction, per-session, per-user, or per-database level

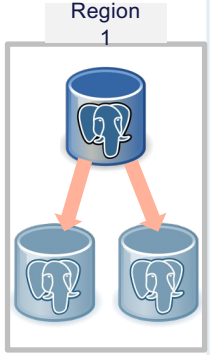

Extreme High Availability

- “Active/Active” mesh replication topologies cross region/cross zonal
- No-downtime major version upgrades
- Self-service superuser access on clusters
- Fast recovery and higher performance workloads

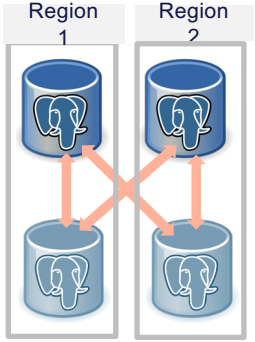

+ “Faraway Replicants”

Allows across region async replicants

- Read-only workloads on replicants in another region
- Allows independent database backups in another region
- Can have a mix of faraway replicants in a mix of regions
- Can be different sized compared to active nodes
  - Different sized compute/RAM
  - Different storage subsystem

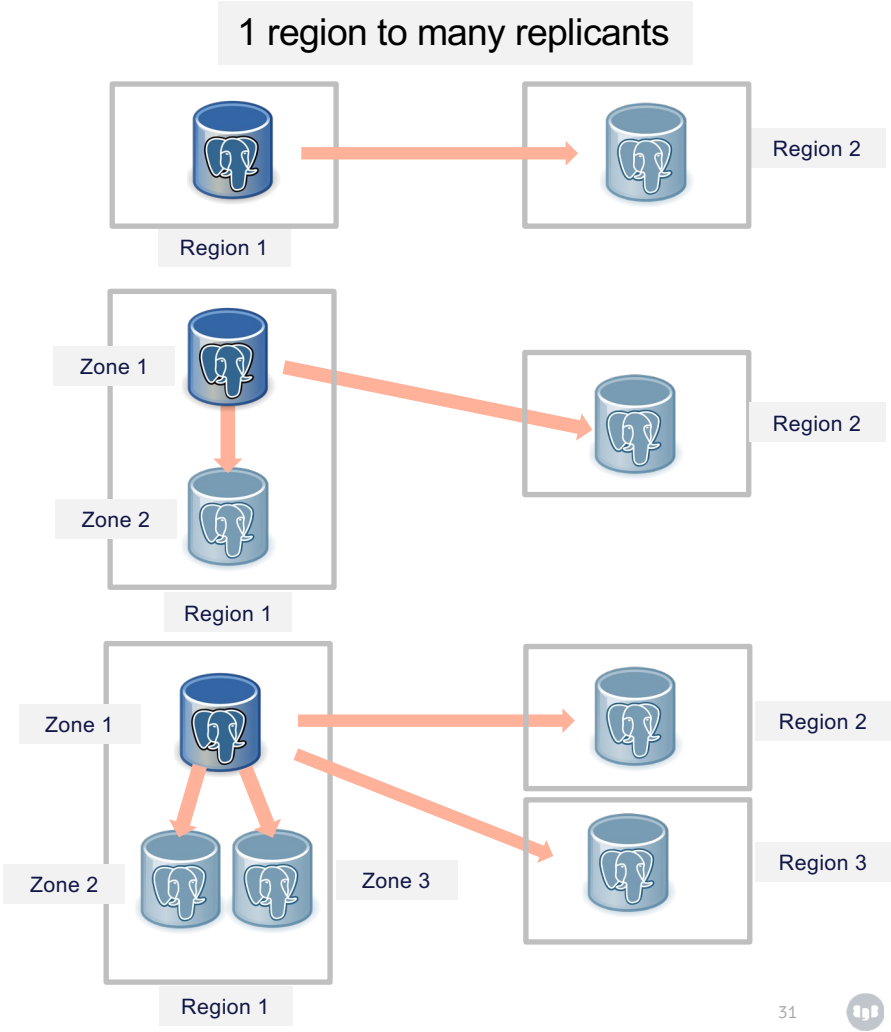

LET'S LOOK

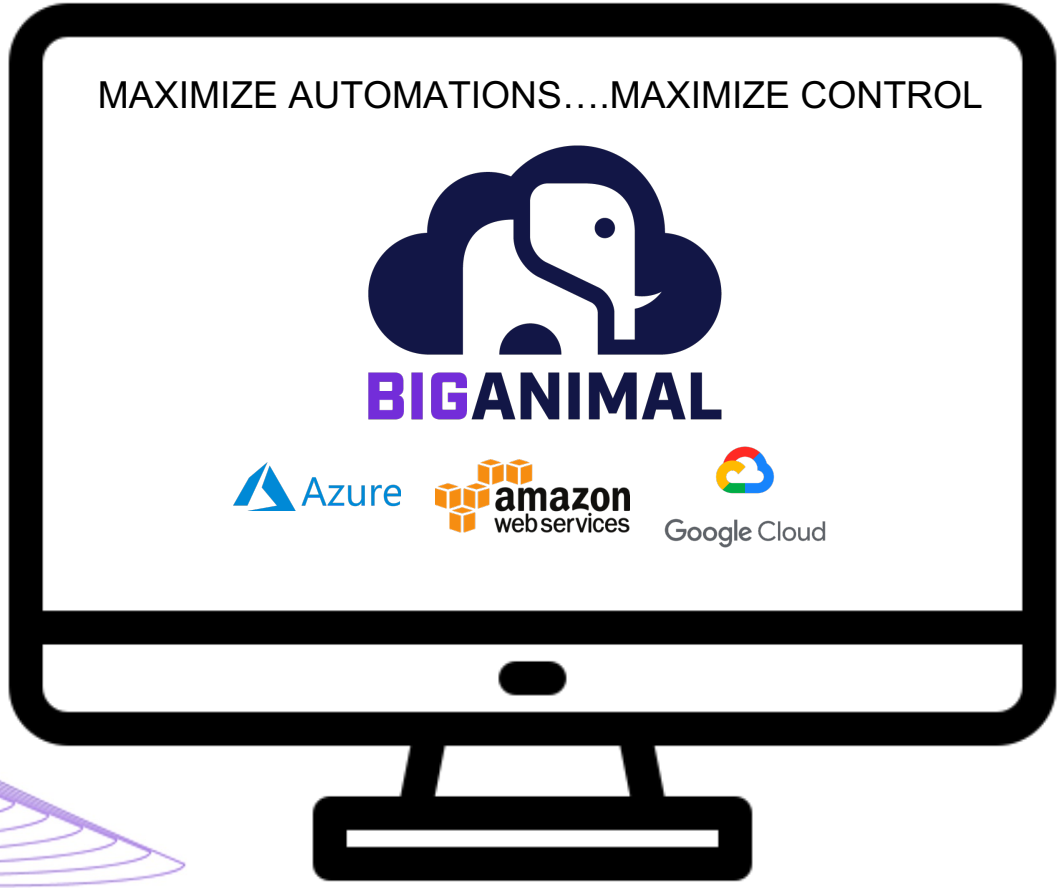

EDB

AGENDA: What we will review today

- Why is Postgres the primary database-of-choice to the cloud?
- The challenges in making the right decisions in the cloud – and where there are still gaps
- Using a best-in-class offering to bring forward an optimized environment
- The current offerings – what they are…and what they are not

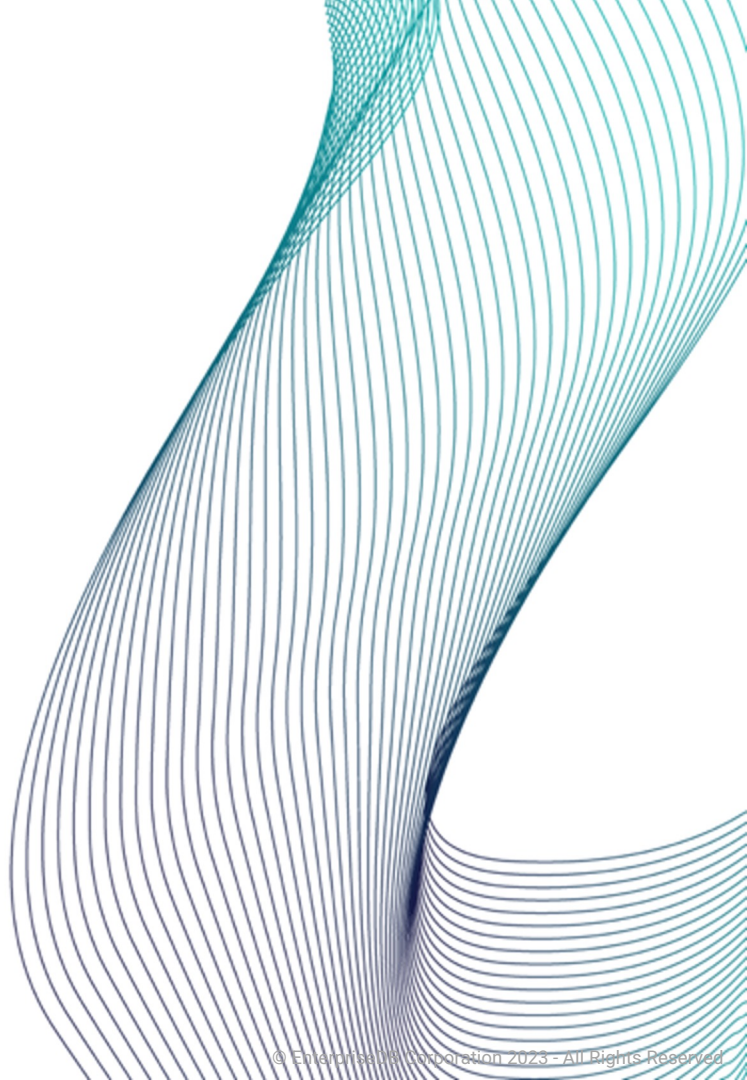

EDB

BigAnimal compared to…

	Azure		amazon web services			Google Cloud	
	Big Animal	**Flexible Server Postgres**	**Big Animal**	**RDS Postgres**	**Aurora Postgres**	**Big Animal**	**Google SQL Postgres**
Max Storage	65TB/160K IOPS with **Ultradisk**	32TB/18K IOPS with SSD	65.5TB/ 256K IOPS with **io2** Disk	65.5TB/ 256K IOPS with io1 Disk	128TB/ Unknown IOPS	TBD (Coming 2Q, 2023)	65TB/30K IOPS with SSD disk
Postgres Tunability	**98%**	63%	**98%**	71%	70%	**98%**	51%
Max Backup Retention	**180 days**	35 days	**180 days**	35 days	35 days	**180 days**	365 ‘backup’ actions
Backup Costs	**$0.04 per GB (Azure Hot BLOB)**	**$0.095 per GB-month**	**$0.023 per GB (AWS S3)**	**$0.095 per GiB-month**	$0.021 per GB-month	**$0.023 per GB (Google Cloud Std)**	**$0.08 per GB/month**
Read-Only Replicants	**YES – 1 Additional Node (in or out of region)**	YES – 1 Additional Node (**async only**)	**YES – 1 Additional Node (in or out of region)**	**YES – 2 Additional Nodes**	YES – 1 additional Node (**in region only**)	**YES – 1 Additional Node (in or out of region)**	Yes
Support	**24x7 break-fix included | portal w/ health check | sev 1 response 30 min/remedy <4 hours**	Additional Cost | ‘Pro-Direct Support’ | 24x7 access ‘after a support request is submitted’ | sev ‘A’ response **60 min**/remedy ???	**24x7 break-fix included | portal w/ health check | sev 1 response 30 min/remedy <4 hours**	Additional cost | 24x7 access | sev 1 response **4 hours**/ remedy ???		**24x7 break-fix included | portal w/ health check | sev 1 response 30 min/remedy <4 hours**	Additional cost | 24x7 ‘Enhanced Support’ | ‘P1’ response **60 min** / remedy ???

What about the database offerings?

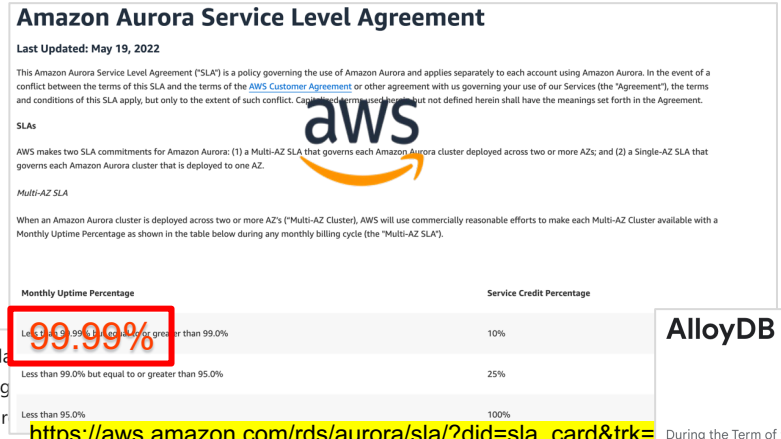

Amazon Aurora Service Level Agreement

Last Updated: May 19, 2022

This Amazon Aurora Service Level Agreement ("SLA") is a policy governing the use of Amazon Aurora and applies separately to each account using Amazon Aurora. In the event of a conflict between the terms of this SLA and the terms of the AWS Customer Agreement or other agreement with us governing your use of our Services (the "Agreement"), the terms and conditions of this SLA apply, but only to the extent of such conflict. Capitalized terms used herein but not defined herein shall have the meanings set forth in the Agreement.

SLAs

AWS makes two SLA commitments for Amazon Aurora: (1) a Multi-AZ SLA that governs each Amazon Aurora cluster deployed across two or more AZs; and (2) a Single-AZ SLA that governs each Amazon Aurora cluster that is deployed to one AZ.

Multi-AZ SLA

When an Amazon Aurora cluster is deployed across two or more AZ's ("Multi-AZ Cluster"), AWS will use commercially reasonable efforts to make each Multi-AZ Cluster available with a Monthly Uptime Percentage as shown in the table below during any monthly billing cycle (the "Multi-AZ SLA").

Monthly Uptime Percentage	Service Credit Percentage
Less than 99.99% but equal to or greater than 99.0%	10%
Less than 99.0% but equal to or greater than 95.0%	25%
Less than 95.0%	100%

99.99%

https://aws.amazon.com/rds/aurora/sla/?did=sla_card&trk=

- **Zone-redundant HA**: This option provides a complete isola… infrastructure across multiple availability zones within a reg… of availability, but it requires you to co… application r… zones. Zone-redundant HA is pr… wh… want protection from availability zone failures. However, one should ac… latency for cross-AZ synchronous writes. This latency is more pronounced for applications with short duration transactions. Zone-redundant HA is available in a subset of Azure regions where the region supports multiple availability zones. Uptime SLA of 99.99% offered in this configuration.
- **Same-zone HA**: This option provide for infrastructure redundancy with lower network latency because the primary and standby servers will be in the same availability zone. It provides high availability without the need to configure application redundancy across zones. Same-zone HA is preferred when you want to achieve the highest level of availability within a single availability zone. This option lowers the latency impact but makes your application vulnerable to zone failures. Same-zone HA is available in all Azure regions where you can deploy Flexible Server. Uptime SLA of 99.95% offered in this configuration.

99.99%

https://learn.microsoft.com/en-us/azure/postgresql/flexible-server/concepts-high-availability

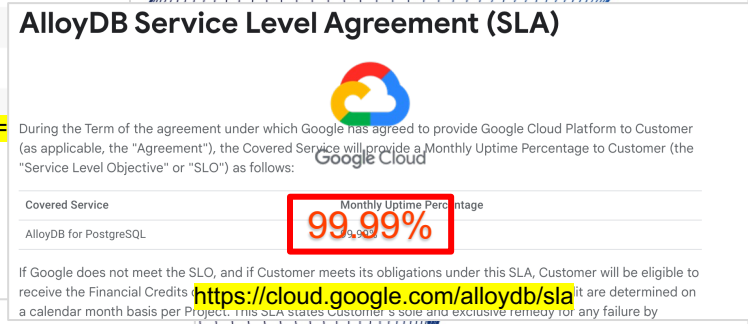

AlloyDB Service Level Agreement (SLA)

During the Term of the agreement under which Google has agreed to provide Google Cloud Platform to Customer (as applicable, the "Agreement"), the Covered Service will provide a Monthly Uptime Percentage to Customer (the "Service Level Objective" or "SLO") as follows:

Covered Service	Monthly Uptime Percentage
AlloyDB for PostgreSQL	99.99%

99.99%

If Google does not meet the SLO, and if Customer meets its obligations under this SLA, Customer will be eligible to receive the Financial Credits … it are determined on a calendar month basis per Project. This SLA states Customer's sole and exclusive remedy for any failure by

https://cloud.google.com/alloydb/sla

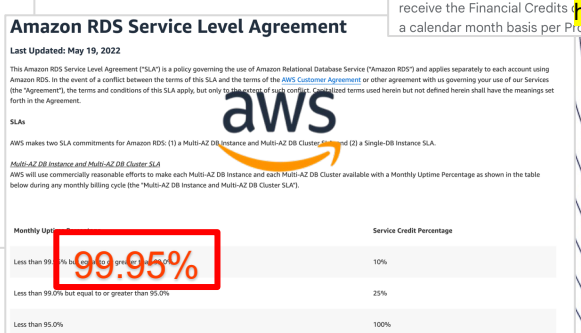

Amazon RDS Service Level Agreement

Last Updated: May 19, 2022

This Amazon RDS Service Level Agreement ("SLA") is a policy governing the use of Amazon Relational Database Service ("Amazon RDS") and applies separately to each account using Amazon RDS. In the event of a conflict between the terms of this SLA and the terms of the AWS Customer Agreement or other agreement with us governing your use of our Services (the "Agreement"), the terms and conditions of this SLA apply, but only to the extent of such conflict. Capitalized terms used herein but not defined herein shall have the meanings set forth in the Agreement.

SLAs

AWS makes two SLA commitments for Amazon RDS: (1) a Multi-AZ DB Instance and Multi-AZ DB Cluster, and (2) a Single-DB Instance SLA.

Multi-AZ DB Instance and Multi-AZ DB Cluster SLA

AWS will use commercially reasonable efforts to make each Multi-AZ DB Instance and each Multi-AZ DB Cluster available with a Monthly Uptime Percentage as shown in the table below during any monthly billing cycle (the "Multi-AZ DB Instance and Multi-AZ DB Cluster SLA").

Monthly Uptime Percentage	Service Credit Percentage
Less than 99.95% but equal to or greater than 99.0%	10%
Less than 99.0% but equal to or greater than 95.0%	25%
Less than 95.0%	100%

99.95%

RDS –vs- BigAnimal - The sticker price

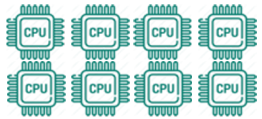

MEDIUM T-SHIRT

8 CPU

64 GB RAM

(r6i.2xlarge)

What	AWS RDS Postgres			EDB Big Animal with Community Postgres		
EC2	**$421.65**			**$243.38**		
Storage Config	IO1 at 7K IOPS	IO1 at 10K IOPS	IO1 at 15K IOPS	IO1 at 7K IOPS	IO1 at 10K IOPS	IO1 at 15K IOPS
Storage Cost	**$737.50.50**	**$1,037.50**	**$1,537.50**	**$492.50**	**$687.50**	**$1,000.00**
BigAnimal/ Support Cost	**$115.92** (support)	**$145.92** (support)	**$195.92** (support)	**$500** (BigAnimal)	**$500** (Big Animal)	**$500** (BigAnimal)
Total Sticker	**$1,275.07**	**$1,605.07**	**$2,155.07**	**$1,235.88**	**$1,430.88**	**$1,780.88**
PLUS: Backup Costs	$0.095/GB/month (per 100GB database size = $275/month)			$0.023/GB/month (per 100GB database = $66.70/month)		
PLUS: Connection Pooler	$0.015/vCPU/HR (8 CPU monthly = $87.60)			Cost for c5.large instance only ($39.42)		

Finally – COST & MONEY - RDS

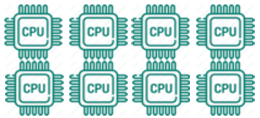

MEDIUM T-SHIRT

8 CPU

64 GB RAM

(r6i.2xlarge)

on aws

What	AWS RDS Postgres	EDB Big Animal with Community Postgres			
Cost	**$1,275.07**	**$814.38**	**$1,093.38**	**$1,235.88**	**$1,780.00**
Storage	IO1 at 7000 IOPS	Gp3 at 7000 IOPS	Io2 at 5000 IOPS	Io2 at 7000 IOPS	Io2 at 10,000 IOPS
HAMMERDB (TPM) Untuned Results	**200,522**	**180,227**	**257,897**	**277,536**	**295,674**
$s/TPM	$0.0060	$0.0026	$0.0040	$0.0034	$0.0060
TPMs/$	157	221	236	224	216

Higher better

Lower better

Higher better

EDB

Finally – COST & MONEY - Aurora

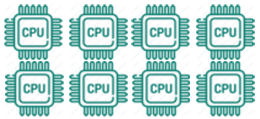

8 CPU

64 GB RAM

on

What	AWS Aurora	EDB Big Animal with Community Postgres			
Cost	**$4,690.40**	**$787.38**	**$1,105.88**	**$1,235.88**	**$1,430.88**
Storage	???? (price assumed 7000 IOPS * 730 hours)	Gp3 at 7000 IOPS	Io2 at 5000 IOPS	Io2 at 7000 IOPS	Io2 at 10,000 IOPS
Optimized Results	**282,250**	**246,624**	**275,980**	**307,653**	**311,875**
$s/TPM	$0.017	$0.003	$0.004	$0.004	$0.005
TPMs/$	60	313	249	250	218

Higher better

Lower better

Higher better

When I did a side-by-side 'race'.....

AWS RDS Postgres	Monthly	Annually
Compute (r6i.4xlarge)	$ 843	$ 10,118
Storage (io1 at 7K IOPS)	$ 738	$ 8,850
Perf Insights	$ 35	$ 420
Support	$ 162	$ 1,939
Annual Cost		$ 21,327
Trans/Min		243,663
Trans/Sec		4,061
$/TPM		$0.09
Trans/$		11.4

AWS RDS Aurora Postgres	Monthly	Annually
Compute (r6i.4xlarge)	$ 1,110	$ 13,315
Storage (assume 7000 @ 730 hours)	$ 3,709	$ 44,508
Perf Insights	$ 35	$ 420
Support	$ 485	$ 5,824
Annual Cost		$ 64,068
Trans/Min		363,619
Trans/Sec		6,060
$/TPM		$0.18
Trans/$		5.7

AWS BIG ANIMAL	Monthly	Annually
Compute (r6i.4xlarge)	486.76	$ 5,841
Storage (io2 at 7K IOPS)	$ 493	$ 5,910
Big Animal Community Postgres	$ 1,000	$ 12,000
Annual Cost		$ 23,751
Trans/Min		663,002
Trans/Sec		11,050
$/TPM		$0.04
Trans/$		27.9

VALIDATED USING AWS CALCULATOR - July 23, 2023(1 year reserved)

Azure Flexible Server	Monthly	Annually
Compute (v5) (us_east-2)	$ 1,460	$ 17,520
Storage (2048 GB)	$ 266	$ 3,195
Support (Professional Direct)	$ 1,000	$ 12,000
Annual Cost		$ 32,715
Trans/Min		99635
Trans/Sec		1,661
$/TPM		$0.33
Trans/$		3.0

Azure BIG ANIMAL	Monthly	Annually
Compute (v5) (us_east-2)	434.17	$ 5,210
Storage (2048 GB)	$ 235	$ 2,820
Big Animal Community Postgres	$ 1,000	$ 12,000
Annual Cost		$ 20,030
Trans/Min		147759
Trans/Sec		2,463
$/TPM		$0.14
Trans/$		7.4

VALIDATED USING AZURE CALCULATOR - July 23, 2023

AWS RDS	AWS Aurora	BigAnimal AWS	Azure Flex	BigAnimal Azure
$21,327/year	$64,068/year	$23,751/year	$32,719/year	$20,030/year
243,663 TPM	363,619 TPM	663,002 TPM	99,634 TPM	147,759 TPM
$0.09/TPM	$0.18/TPM	$0.04/TPM	$0.33/TPM	$0.14/TPM

EDB

WEBINAR SLIDES AND RECORDING:
https://info.enterprisedb.com/biganimal-price-vs-performance_webinarlp.html

Big Animal is part of the broader ecosystem

pgAdmin

- Provision Big Animal directly of pgAdmin

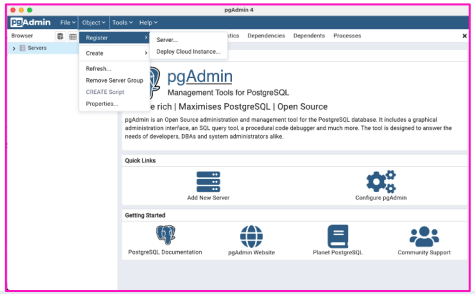

Migration Portal

- Sending resulting DDL to Big Animal

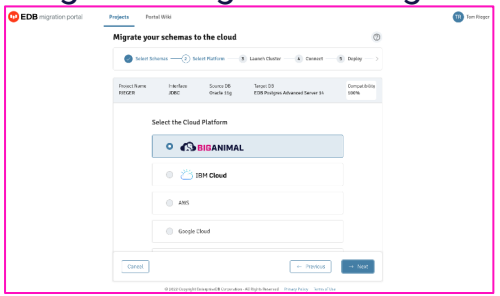

Built on our K8 operator

PEM support

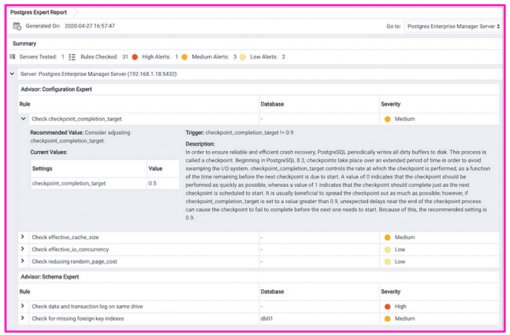

Database-centered Innovation Roadmap

Availability

- Cross-region replicants
- Cross-region multi-master clustering (preview)
- Cross-cloud clustering
- Disk snapshot backup/recovery
- Custom maintenance windows
- Adding Google Cloud Platform
- Big Animal Hosted
- Adding new regions

Performance

- Faster access to clustered data
- Automate usage of connection pooler
- Realtime database tuning based on workload
- Separate WAL logs into own disk volume
- Automated Index Advisor

Automation

- Storage auto clustering
- Self-serve major upgrade control
- Autoscale storage across clusters
- Automated database tuning
- Postgres Hosted
- Self-service Credit card payment
- AWS and GCP Marketplace

Cost Optimization

- Automate Hibernate/Resume
- Best price/ performance CPU VMs
- VM and storage workload optimization
- Adding new components (i.e. Intel Sapphire Rapids, AWS Graviton, AWS GP3 v2)

Why EDB for Postgres

Products, services, and support options to get the most out of Postgres

Product Focus

For performant, scalable, high availability, monitoring, and backup/recovery on any platform

+

Support Expertise

24 x 7 support for superset of open-source projects, backed by top contributors

+

Postgres Education

Comprehensive education – with all on-demand FREE - and the **only Postgres certification program**

+

Comprehensive Services

Expertise with migration, operation, and strategy

=

Cost savings

Significant savings versus commercial tech PLUS enterprise level servicing

So WHY EDB

- Open Source RDBMS is the future
- Postgres is the leading database
- EDB brings the largest global focus to Postgres
- EDB understands top-tier workload requirements
- EDB makes migrating off “the old” very very very very easy
- EDB brings the most extensive platform choices to you.

EDB

AGENDA: What we will review today

- Why is Postgres the primary database-of-choice to the cloud?
- The challenges in making the right decisions in the cloud – and where there are still gaps
- Using a best-in-class offering to bring forward an optimized environment
- The current offerings – what they are…and what they are not

“WHAT IS NEXT-AND-NOW”

“MEASURING & CHOOSING & TUNING THE RIGHT COMPONENTS IS PARAMOUNT”

“MAXIMIZE AUTOMATIONS + MAXIMIZE CONTROL = MAXIMIZE VALUE”

”FOCUS WINS::CSP HIDING THOSE TOTAL COSTS::NO INNOVATION”

EDB

TRY IT!!!!!!

- Try BigAnimal now at no cost
- You get a $300 credit
- Deploy in the BigAnimal account or your own
- You also get access to all of EDB's training, migration portal and more

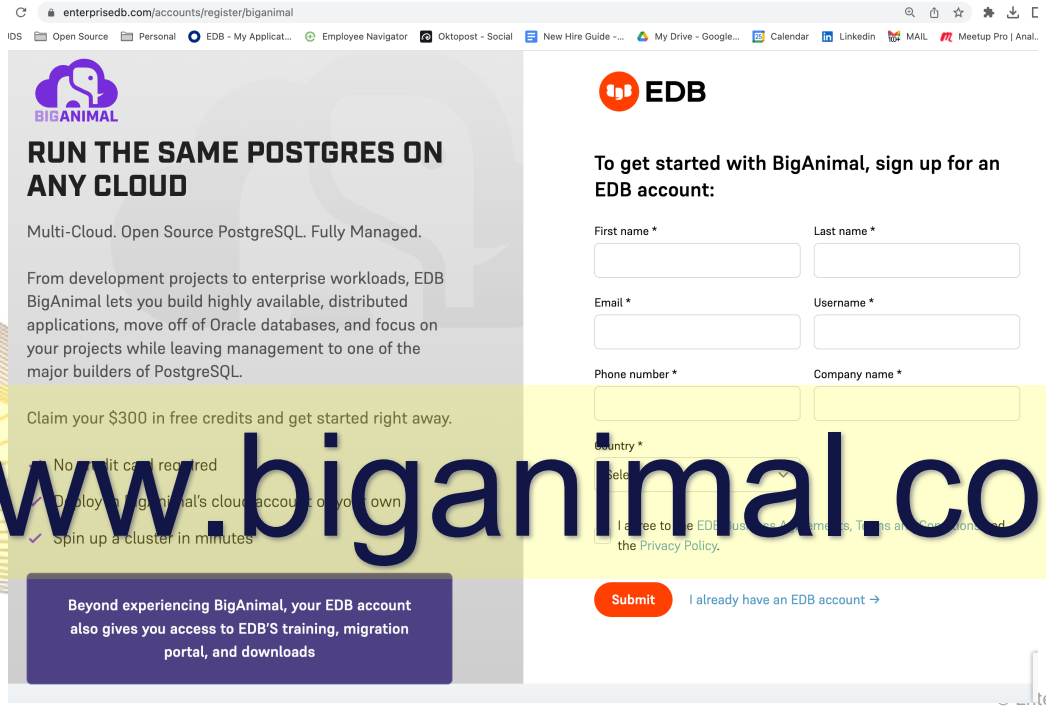

www.biganimal.com

EDB

THANK YOU

ANY QUESTIONS?

Tom Rieger
tom.rieger@enterprisedb.com
(only TOM RIEGER in MINNESOTA on Linkedin)

EDB