

# AWS Storage for the SQL Server DBA



**Rick Lowe**

Principal

*Data FLOWe Solutions LLC*

# Review – MSSQL on AWS

- AWS cover a LOT of services
- Two ways to run MSSQL
  - Virtual machine – EC2 (Elastic Compute Cloud)
  - Instance only – RDS (Relational Database Services)
- Major components of cost
  - Runtime (only incurred while up)
  - Data transfer
  - Storage (per GB per month whether instance up or not)

# Elastic Block Storage

- Drives mounted to EC2 and RDS instances
- Fixed size
- Default volume type (gp2) is solid state
- There are options for magnetic storage
- Provisioned IOPs (io1) allow us to buy IOPs

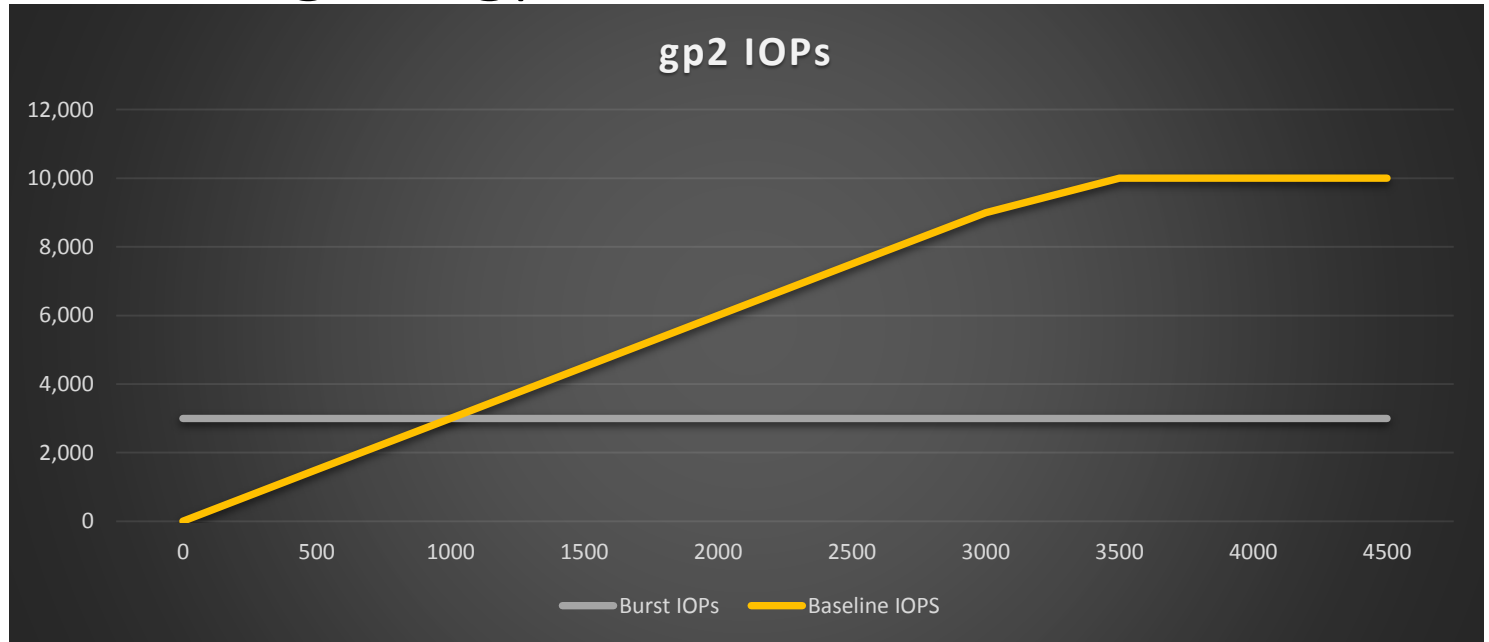
## General Purpose (gp2) storage

- Current size range – 1GB to 16TB
- Baseline performance : 3IOPs/GB
- Max of 10,000 IOPs and 160MiB/sec
- In move US regions, \$0.10 / GB / month

## Bursting for gp2

- Only applies to smaller volumes (< 1000GB)
- Volumes start with burst credits
- Credits allow bursts of 3000 IOPs
- When all credits used, bursting stops
- Credits regenerate with time

# Bursting for gp2 Chart



# Throughput Optimized – st1

- Backed by magnetic storage
- Usually less than ½ cost of gp2
- Minimum size 500GB
- Maximum of 500 IOPs
- ... but high throughput (500 MiB/s)

## Cold HDD storage – sc1

- Similar concept to st1
- Typically  $\frac{1}{4}$  the cost of gp2
- Max of 250 IOPs
- Max of 250 MiB/s



## Provisioned IOPs – io1

- Two cost components
  - Volume size (currently usually \$0.125 / GB / month)
  - IOPs provisioned (currently usually \$0.065 / IOP / month)
- Maximum of 32,000 IOPs (usually)
- Max of 500 MiB/s

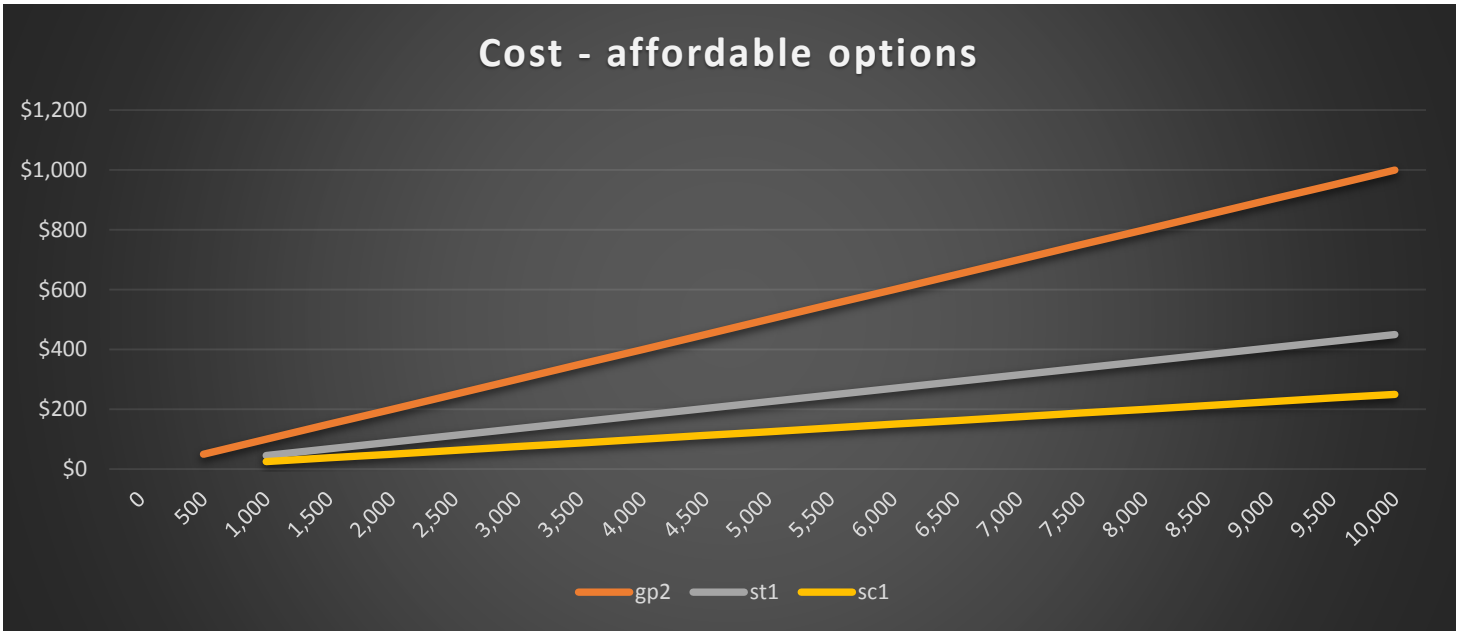
# Use cases for SSD EBS options

- All options – watch page life expectancy
- General Purpose
  - Great default option
  - Reasonable budget, moderate performance
  - Maybe “bursty” workloads
- Provisioned IOPs
  - Higher budget
  - Critical DBs

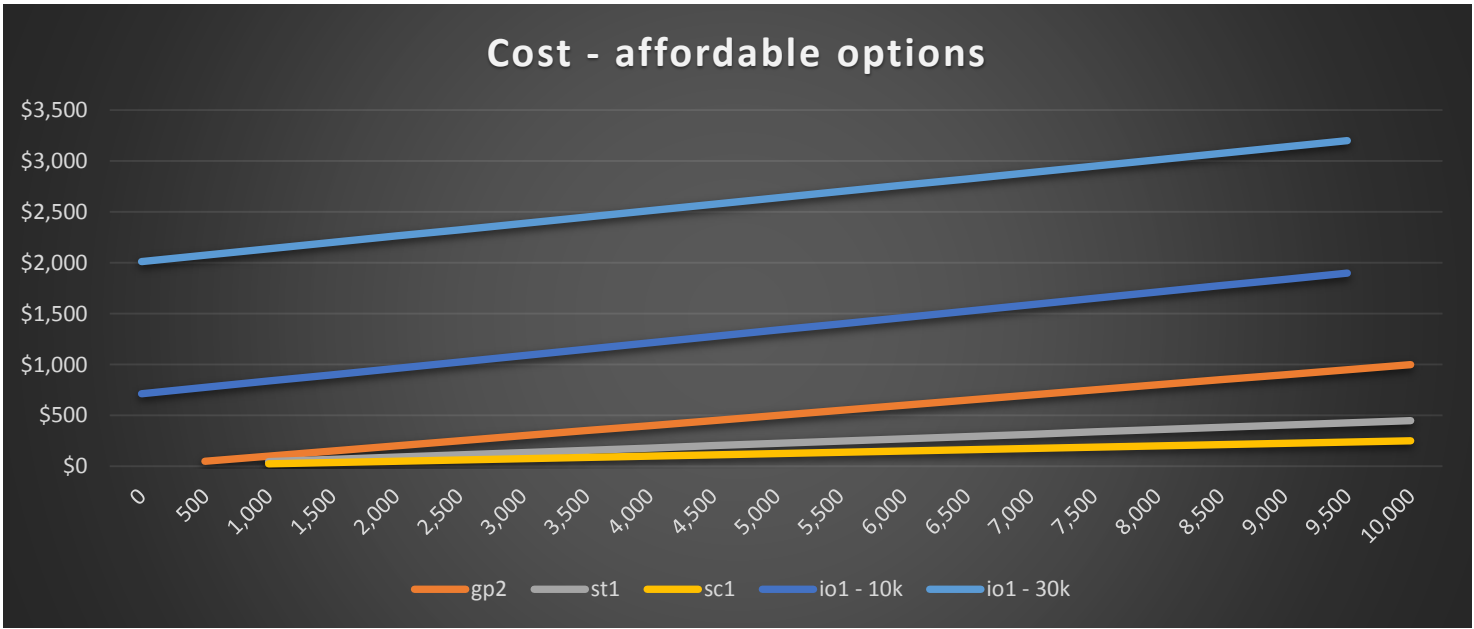
# Use Cases for Magnetic EBS Options

- Budget conscious projects
- DBs where performance is not critical
- Throughput optimized
  - Theoretically great for logs
  - Scan-heavy workloads (columnstore ?)
- Cold HDD
  - Temporary storage of backups
  - Applications with budget pressure

# EBS Costs – Affordable Options



# EBS Costs – All Options



# S3 – Scalable Storage Service

- Sequential access only
- Treat like tape drive
- Less expensive than EBS
- Very scriptable
- NOT suitable for DB files

# Sample S3 Commands

- Third party GUIs are available
- Copy file to S3 in batch file  
`aws s3 cp mydbbackup.bak s3:\\longbucketname\`
- 1-way directory sync (with deletes)  
`aws s3 sync . S3:\\longbucketname\ --delete`
- Copy file from Powershell  
`Write-S3Object -BucketName longbucketname -File mydbbackup.bak`

## S3 Storage Classes

- Standard – usually between \$0.021 and \$0.023 / GB / mo  
Durable – stored in multiple Azs. \$0.002 / GB to read
- Infrequent access – currently usually \$0.0125 / GB / mo  
More expensive to access - \$0.01 / GB to read
- One-zone IA – currently usually \$0.01 / GB / mo  
Only stored in 1 zone. \$0.01 / GB to read
- Glacier (archive) – currently usually \$0.004 / GB / mo  
Fee for premature access/delete. Access not immediate.



## S3 Standard Use Cases

- Medium term storage of backups
- Copying files from data center to AWS
- Native DB backups for RDS
- Can consider infrequent access

## Other S3 Storage Class Use Cases

- Glacier – archival of backups for compliance
- Infrequent Access – Backups retained for DR
- One-zone IA – Temporary storage