Cloudwatch

Host level metrics =
- CPU
- Network
- Disk
- Status

Custom metric = • Memory
→ min granularity = 1 min.

Detailed Monitoring = 1 min
Standard Monitoring = 5 min

EBS -

VolumeReadOps = total number of I/O operations in a specified time
VolumeWriteOps

VolumeQueueLength = Number of read + write operation requests waiting to be completed in a specified time.

gp2 = General Purpose SSD
io1 = Provisioned IOPS SSD
st1 = Throughput Optimized HDD
sc1 = Cold HDD

Volume Status

warning = Degraded or Severly degraded
impared = Stalled or Not Available

- encrypt existing volume ⇒ create snapshot, copy snapshot + apply encryption at same time.
- encryption must be enabled at creation

AMIs

- ami's are region bound (need to copy to new region)
- cannot share encrypted AMI. need to re-encrypt, share the KMS key to the target account.
- AMI's with an associated billing products code cannot be shared directly

S3

- 0 bytes – 5tb
- read after write consistency for **PUTS** of new objects
- eventual consistency for **Overwrites + Deletes**
- 11x9's durability (99.999999999%) — Standard S3

Encryption

in transit – SSL/TLS

at rest
- Server Side Encryption
  - S3 managed keys – SSE-S3
  - Key management service – SSE-KMS
  - Customer provided keys – SSE-C

SSE-S3 – x-amz-server-side-encryption: AES256

SSE-KMS – x-amz-server-side-encryption: aws:kms

enforce encryption with bucket policy
- s3:PutObject – ≠x-amz-server-side-encryption:

Glacier

- can only delete archive via CLI (or upload)
- snowball to S3 then lifecycle policies to glacier

Storage Gateway

- File Gateway
  - stored in S3
  - accessed via NFS or SMB
- Volume Gateway
  - iSCSI
  - Gateway Stored – local data, backup to AWS
  - Gateway Cached – S3 store, cache locally
- Tape Gateway
  - uses Glacier
  - works with backup clients

RDS

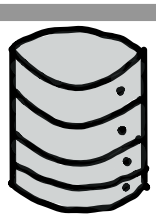

Multi-AZ
H.A

- Synchronous
- Backups + Restores are taken from Standby
- Force AZ failover by rebooting instance

Read Replicas
Scalable

- Read heavy scaling
- Business reporting or data warehousing
- MySQL/PostgreSQL/MariaDB
  - use native asynchronous replication.
- Has own DNS
- Can have upto 5
- can be in different regions
- Key metric ReplicaLag
- need to have backups turned on

Version of RDS — rds describe-db-instances --region

Aroura

- Aurora
- Aurora Serverless

Scale up (Instance Size) if write demand ⬆

Scale out (Read Replicas) if read demand ➡

- Encryption at rest is turned on by default
- Failovers defined by Tiers (Tier 0 > Tier 1 > 15)
- Cross Region Replica (prefer Multi-AZ)

Elasticache

- CPU Utilization
- Swap Usage
- Evictions
- Concurrent Connections

Memcached

- Multi-threaded
  if exceeds 90% add more nodes to the cluster
- SwapUsage should be 0 most of the time
  if exceeds 50mb - increase
  memcached_connections_overhead
- Evictions - Scale Up or Scale Out

Redis

- Evictions - Only Scale Out

Both — set alarm for concurrent connections

CloudFormation

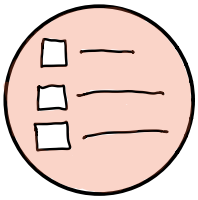

- cfn-init — resource metadata / install packages, create files etc
- cfn-signal — use with CreationPolicy or WaitCondition
- cfn-get-metadata — retrieve metadata
- cfn-hup — check metadata and execute custom hooks when changes are detected

RedShift

- Redshift log files:
  - Connection log
  - User log
  - User Activity log

KMS

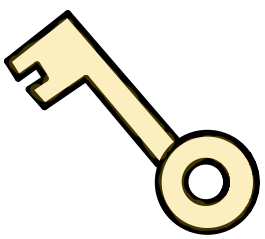

- KMS + CloudHSM
  - Dedicated / no free tier
  - FIPS level 3 compliant
  - symmetric or asymmetric encryption
- KMS
  - shared hardware
  - Free tier eligible
  - symmetric encryption

ELB – Monitoring

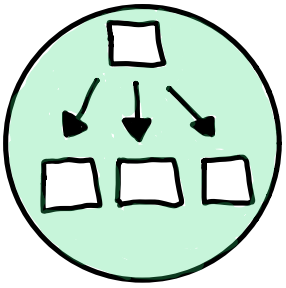

- cloudwatch metrics
- access logs
- request tracing
- cloud trail logs

who? (IPs etc) (Persists when ec2 deleted)

ALB only – X-Amzn-Trace-ID

ALB – Layer 7 /

NLB – Static IP – 1 per subnet

Errors

- 500 (s)erver
- 400 client
- 200 success

Cloudwatch Metrics

General Health
- HealthyHostCount
- UnHealthyHostCount
- HTTPCode_Backend_2XX

Performance
- Latency
- RequestCount

(Classic Only)
- SurgeQueueLength
- SpilloverCount

Can put an ALB behind a NLB to get benefits of both.

VPC

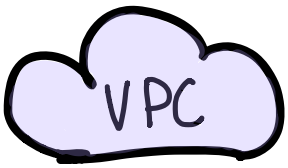

- enableDnsSupport is disabled by default
  if creating via the CLI

Main route table should only be for private subnets
and NAT instance should be for 0.0.0.0/0

Placement Groups

— Cluster
- Single AZ
- High network throughput
- Launch at same time or may get insufficient capacity error

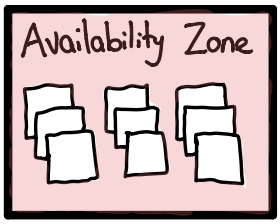

— Partition
- multi-AZ
- max 7 partions per AZ
- 21 total