aws SUMMIT

Optimizing Application Performance In AWS

Scott Wheeler

Principal Cloud Architect

Asperitas Consulting

Agenda

What Affects Performance in AWS?

Networking Performance

Compute Performance

Storage Performance

What Affects Performance?

Application Architecture (out of scope)

Software Optimization (out of scope)

Network

Compute

Storage

Network

Network Performance

Application Location

Instance Types

- \+ Larger is better
- \+ Newer is better

AWS Enhanced Networking (SR-IOV)

- \+ Increased **bandwidth** (closer to advertised)
- x Issues with driver support for non-AWS Linux

AWS Placement Groups

- \+ Low **latency** between instances
- x May have instance type availability issues
- x Only within an AZ

Instance Type Considerations

Network Bandwidth	Instance Types
Very Low	t2.nano
Low	t1.micro
Low to Moderate	t2 (micro,small,medium,large)
Moderate	medium (m3), large (m3, m4, c4, r3), xlarge (d2, r3, t2), 2xlarge (t2)
High	xlarge (m3, m4 ,c4), 2xlarge (c4, m3, m4, p2, d2, r3), 4xlarge (c4, m4, d2, r3)
Up to 10 Gbps	large (c5, m5, i3, r4), xlarge (c5, m5, r4, i3, x1e), 2xlarge (c5, m5, r4, h1, i3, p3, x1e), 4xlarge (m5, c5, r4, i3, h1, x1e, g3), 8xlarge (x1e)
10 Gbps	8xlarge (c4, r3, r4, d2, h1, i3, g3, p2), 9xlarge (c5), 10xlarge (m4), 12xlarge (m5), 16xlarge (x1, x1e)
25 Gbps	16xlarge (c5, m4, r4, h1, i3, g3, p2, p3), 18xlarge (c5), 24xlarge (m5) 32xlarge (x1, x1e)

Network – Third Party Observed Performance

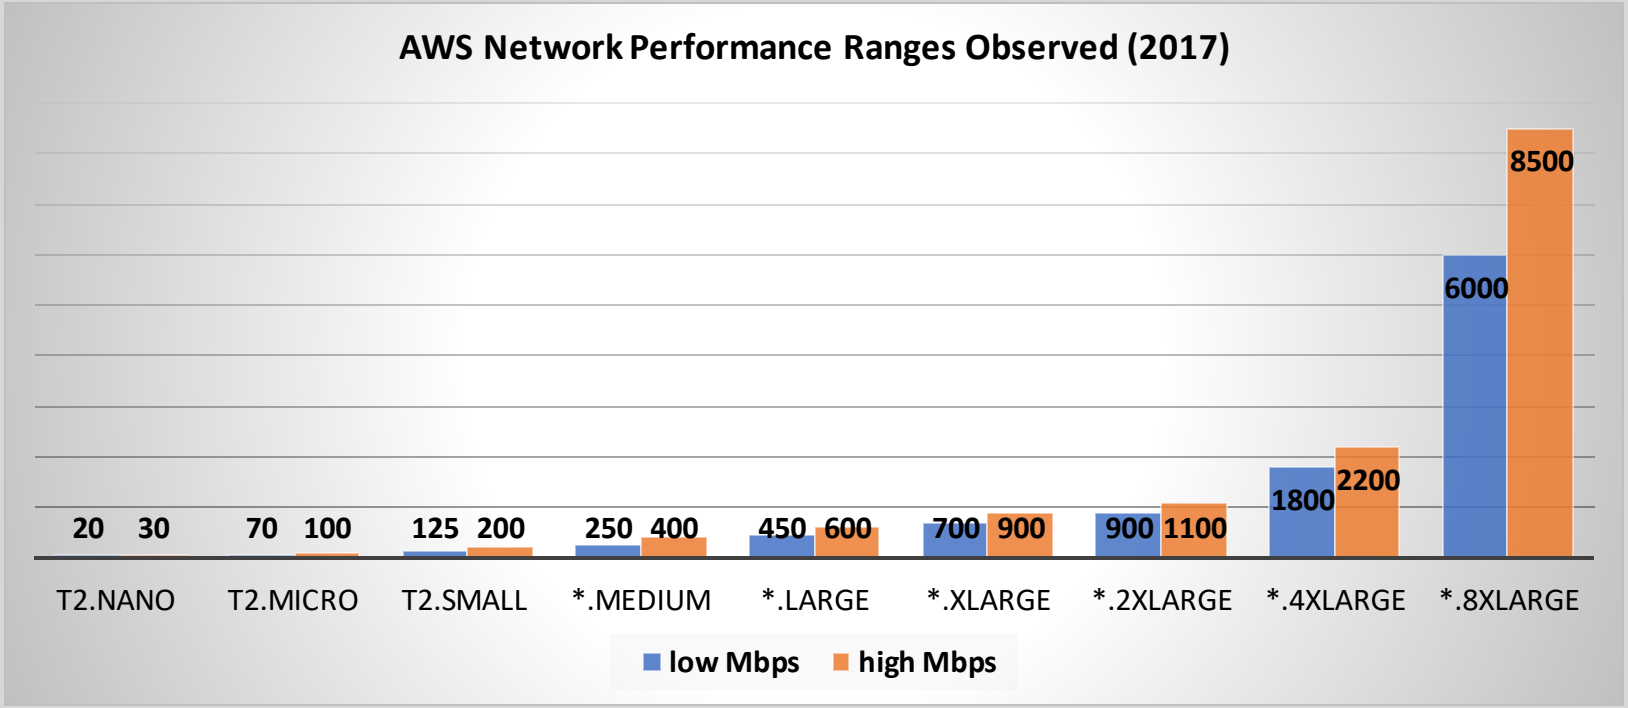

Source: CloudHarmony & flux

Network – Third Party Observed Performance

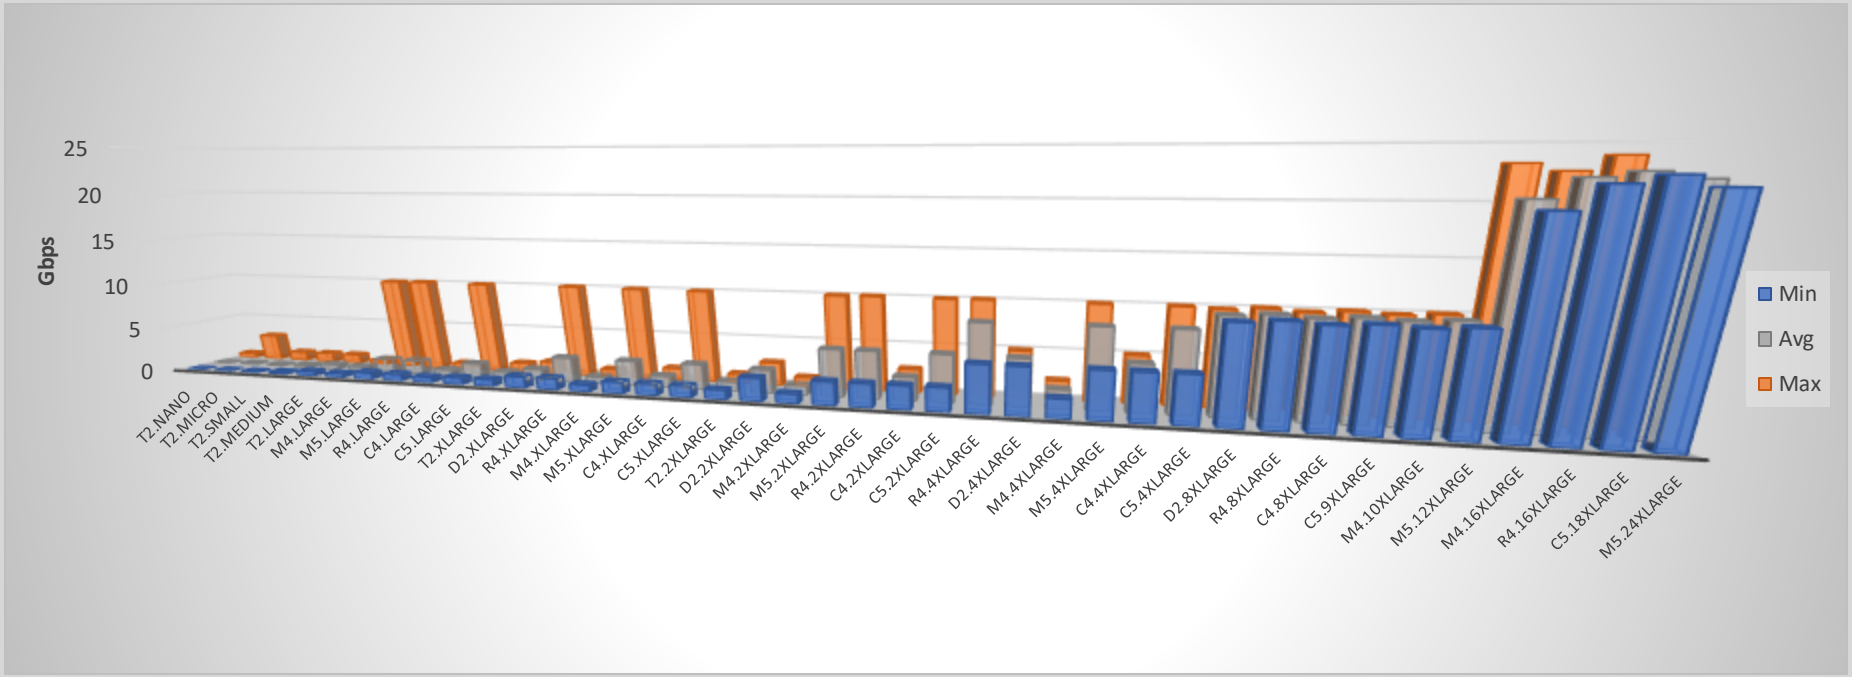

Source: Andreas Wittig, Cloudanaut

Network – New Performance Enhancements

EC2 to S3

- *Increase to 25 Gbps from 5Gbps*

EC2 to EC2

- *5 Gbps: single-flow traffic, 25 Gbps: multi-flow traffic for AZs within a region*

EC2 to EC2 (Cluster Placement Group)

- *10 Gbps: single-flow traffic, or 25 Gbps: multi-flow traffic for AZs within a region*

Source: Amazon Web Services, Jan 2018

Other Network Considerations

Transit VPCs

- x *VGW IPSec VPN connections limited to 1.25Gbps.*
- x *Virtual Routers may only have 2.5Gbps real world throughput.*

Compute

Compute Performance

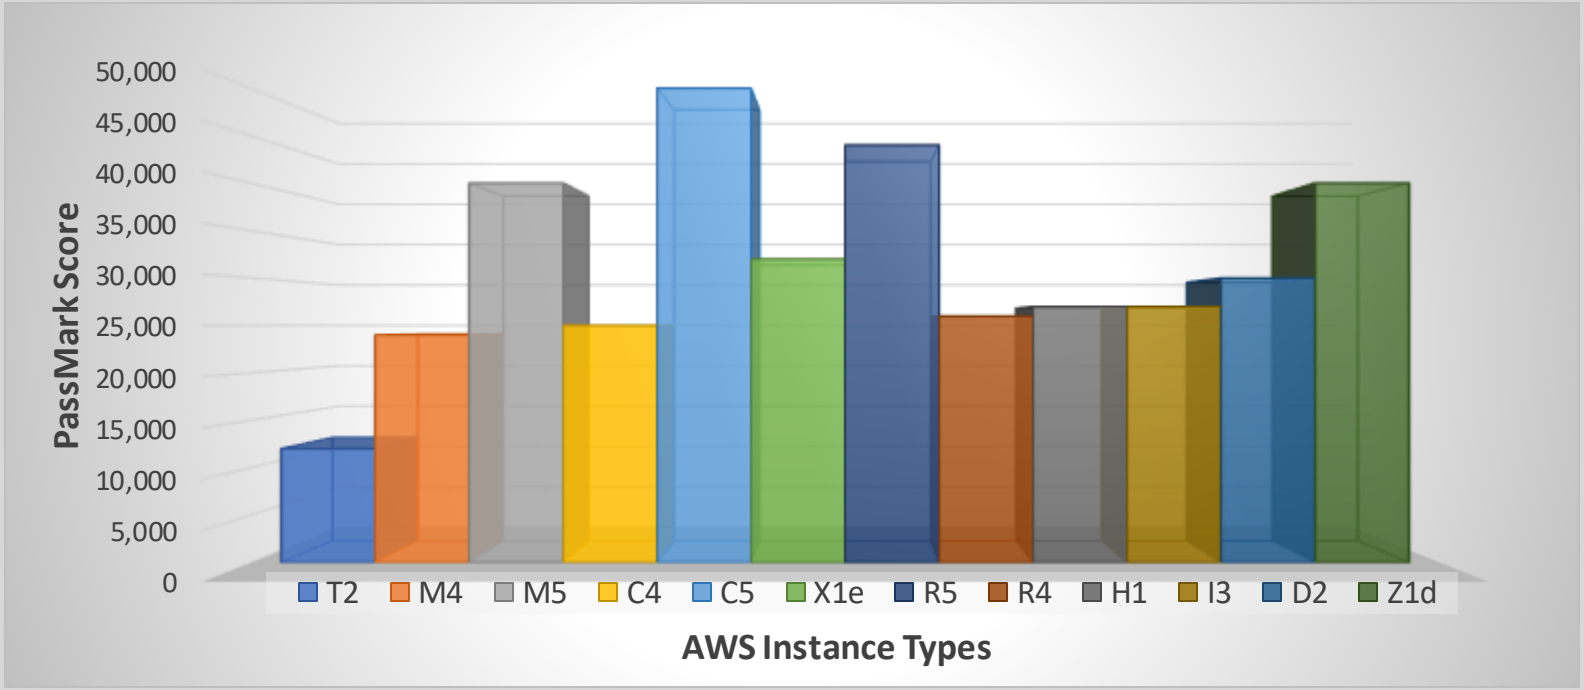

Source: Amazon Web Services

Storage

EBS & Instance Storage Performance

EBS

- Magnetic: IOPS 250-500, throughput 250-500 MiB/s
- Solid State
  - Standard: IOPS 10,000, throughput 160 MiB/s
  - PIOPS: IOPS 32,000, throughput 500 MiB/s

Instance Attached Storage

- NVMe: 3.3 million IOPS, 16 GB/s sequential read

EBS GP2 vs IO1 Performance

Use Case: 500GB @ 3,000 IOPS

gp2: $100/mo (1TB @ 3,000 IOPS)

io1: $258/mo (500GB @ 3,000 IOPS)

Use Case: 1TB @ 10,000 IOPS

gp2: $340/mo (1TB @ 10,000 IOPS)

io1: $750/mo (1TB @ 10,000 IOPS)

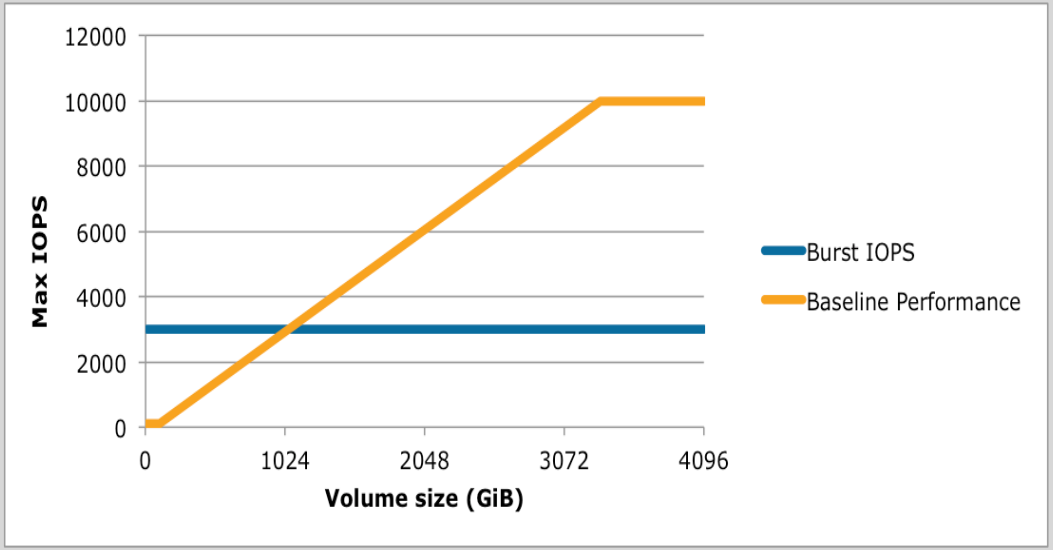

Source: Amazon Web Services

S3 Storage Performance

Keys Matter

- Cause bottlenecks at 100 request/sec/thread
- Add hash string to key prefix

Transfer Acceleration

- Provide up to 30% throughput due to reduced latency

Multipart Upload

- Break larger files into multiple chunks

Utilize CloudFront

- Reduces latency

EFS Storage Performance

File System Size	Aggregate Read/Write Throughput
100 GiB	• Burst to 100 Mbps for 72 min/day • Drive up to 5 Mbps continuously
1 TiB	• Burst to 100 Mbps for 12 hours/day • Drive up to 50 Mbps continuously
10 TiB	• 1 Burst to Gbps for 12 hours/day • Drive up to 500 Mbps continuously

Source: Amazon Web Services

Real World Project

Project Overview

Validate very low latency data feed performance in AWS.

Refactor Into Services

Utilize Containers (Docker)

Benchmark Performance

- Same instance
- Same AZ
- Different AZs
- Placement Groups
- Dedicated Instances

Project Results

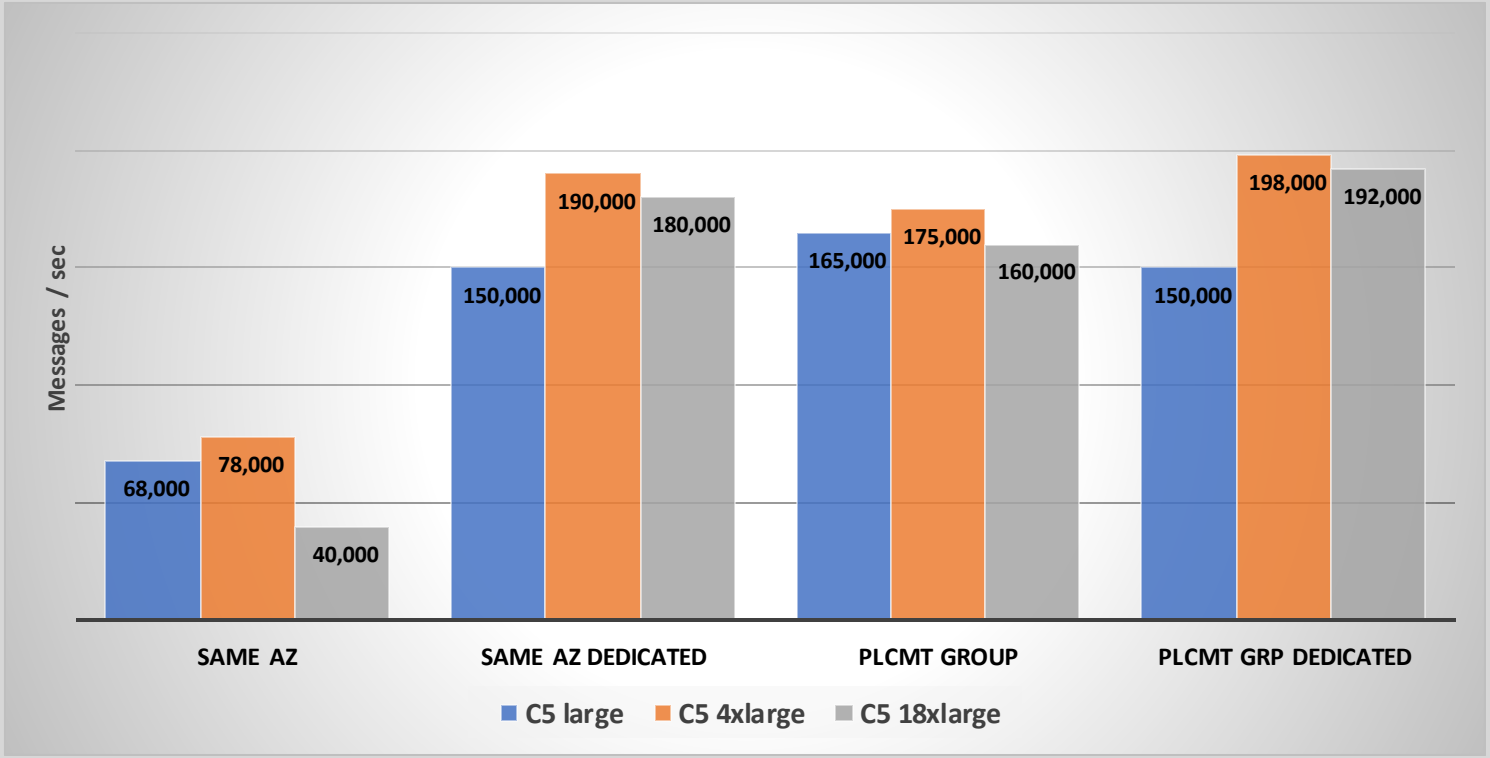

Conclusions, Observations & Recommendations

C5 4xlarge: offered the best price/performance.

Placement Groups: should be used when low latency is needed.

Dedicated Instances: have benefit, but may not be needed.

Separate AZs: only used for DR and failover.

Performance: relatively constant in long running tests.

Contact Information

Scott Wheeler

Principal Cloud Architect
Asperitas Consulting

swheeler@aperitasconsulting.com

@dscottwheeler

Please complete the session survey in the summit mobile app.

aws SUMMIT

Thank you!