



# Azure Native Qumulo

Microsoft Azure HPC Workloads with Azure Native Qumulo

**Kevin McDonald**

Principal Technical Marketing Engineer

**Ryan Farris**

Vice President, Product

March, 2024



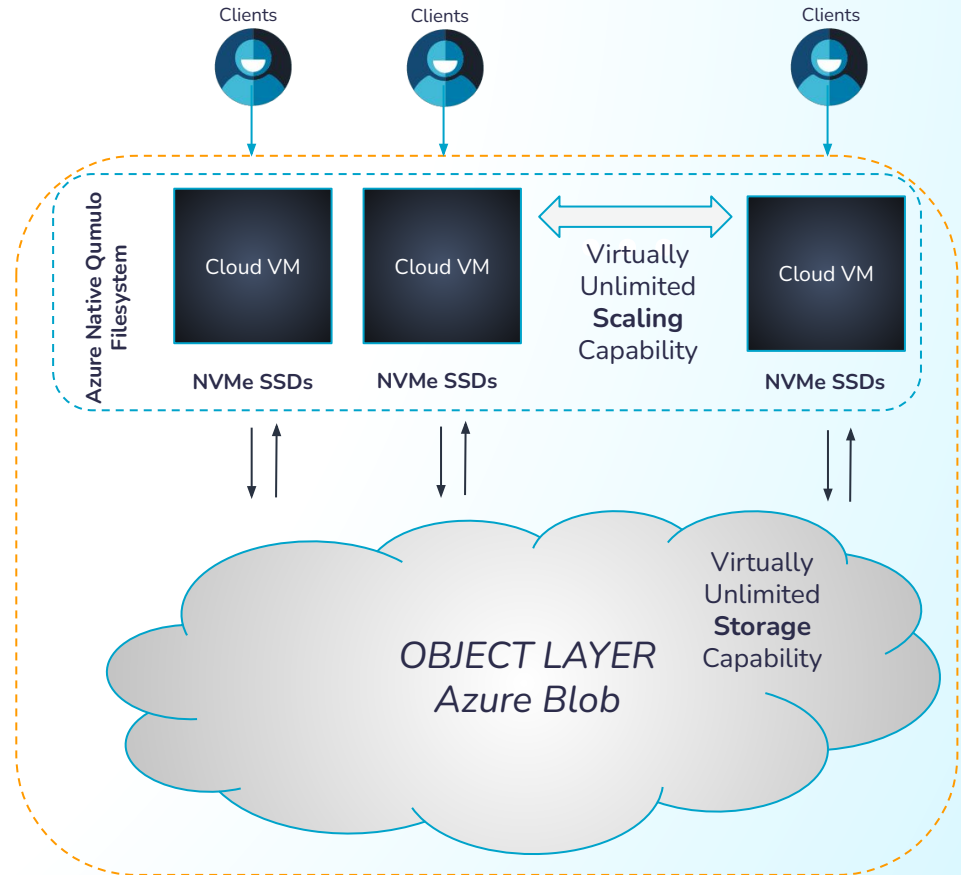
# Agenda

- Azure Native Qumulo (ANQ)
- Scale-out and ANQ architecture
- Performance benchmarks
- 2 examples of industry vertical performance
  - FSI (Using Azure's preferred benchmark)
  - Life Sciences (Genomics)
- Pricing, competitive comparison, GitHub Links



# Azure Native Qumulo (ANQ)

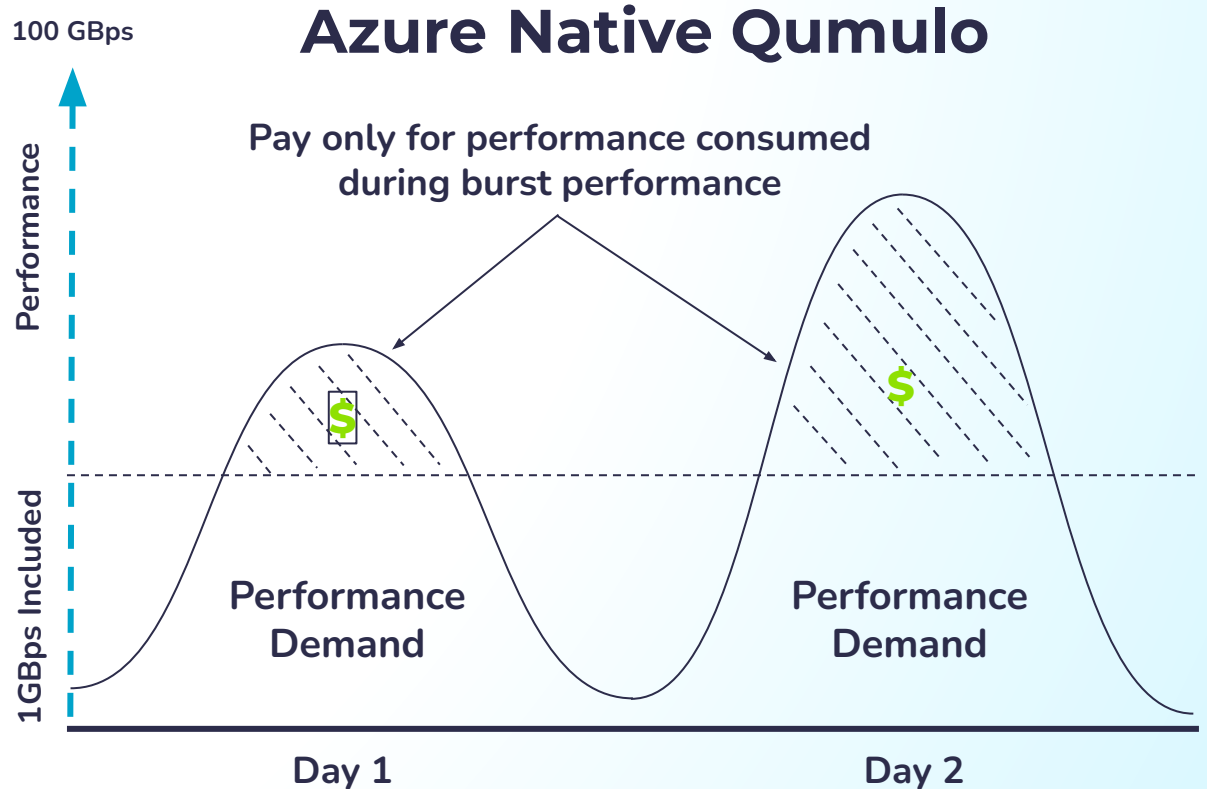
1. Enterprise file workloads near the cost of object
2. Disaggregated compute & object storage to lower cost
3. Filesystem elasticity, charging for metered perf > 1GBps
4. TCO better or comparable to on-prem storage



# A truly elastic file service for HPC!

*A perfect fit for Genomics, Automotive, Manufacturing, Energy, Media and Entertainment, and more.*

- Elastic performance for file-based workloads
- Disruptively priced!
- Pay only for performance used when you need it
- Scale your storage with business demand



# HPC Performance Benchmarks

# DISCLAIMER: Performance is NOT fixed!

## ANQ is a Cloud Service that is designed to elastically scale

**NOTE FOR READERS** – Azure Native Qumulo scales elastically based on application requirements. Performance can be turned **UP** ↑ much higher than the examples shown here. These slides offer a glimpse into ANQ's capability in HPC environments, although at a smaller scale. We invite YOU to do your own testing (with [a free ANQ trial](#)) and work with us to test your HPC workloads in Azure.

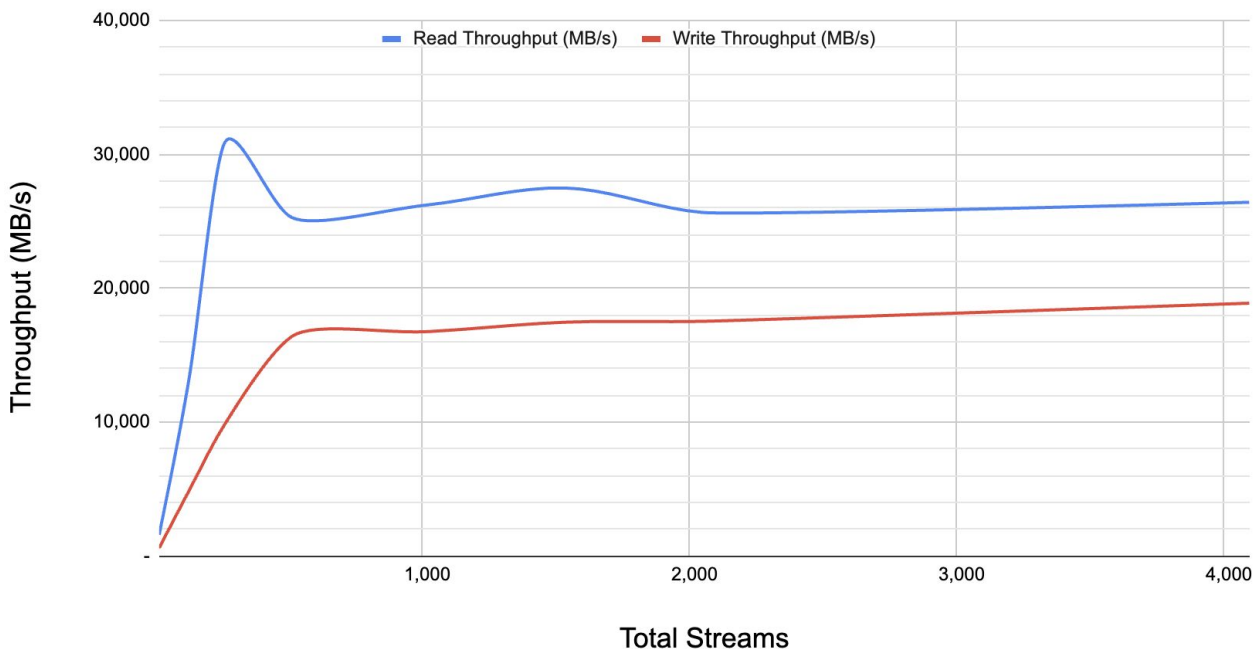
You may be wondering....

1. Can I burst to 50GBps or 200,000 IOPS for only 1 hour and auto-scale back down again after my jobs have completed?  
Absolutely, and much more if necessary. The beauty of it is, you only pay for the burst hour!
2. How many nodes are needed to meet my burst performance demands?  
Forget about node count – shift your perspective! :) We offer a true cloud service experience by billing based on performance and capacity usage. We worry about the infrastructure behind the service so you can focus on your tasks. Imagine interacting with your cloud storage just like you'd interact with any other service which auto scales!
3. But how does ANQ manage to offer such a significant price difference compared to other cloud-native file storage services?  
**We've innovated our way to disruptive pricing.** See slide #3. :) Write operations are transacted in a highly durable cache, journaled, and persisted to object. Read operations are served from the NVMe read cache tier thanks to our prefetch algorithm that has been trained on trillions of I/O patterns over the last 10 years of supporting >1000 happy customers (see slide 23)!

# ANQ Example Throughput Capability

## SMB Large File Sequential Performance

Simulating Media Streaming, Backup/Restore/Archive, Genomics Sequencing, etc...



*IMPORTANT: we can elastically increase from these peaks if required*

- 30 GB/s peak (read)
- 26 GB/s sustained reads at 4,000 SMB streams
- FIO 3.36
- 16 Linux clients (D16)
- Accelerated Networking
- Default tunables

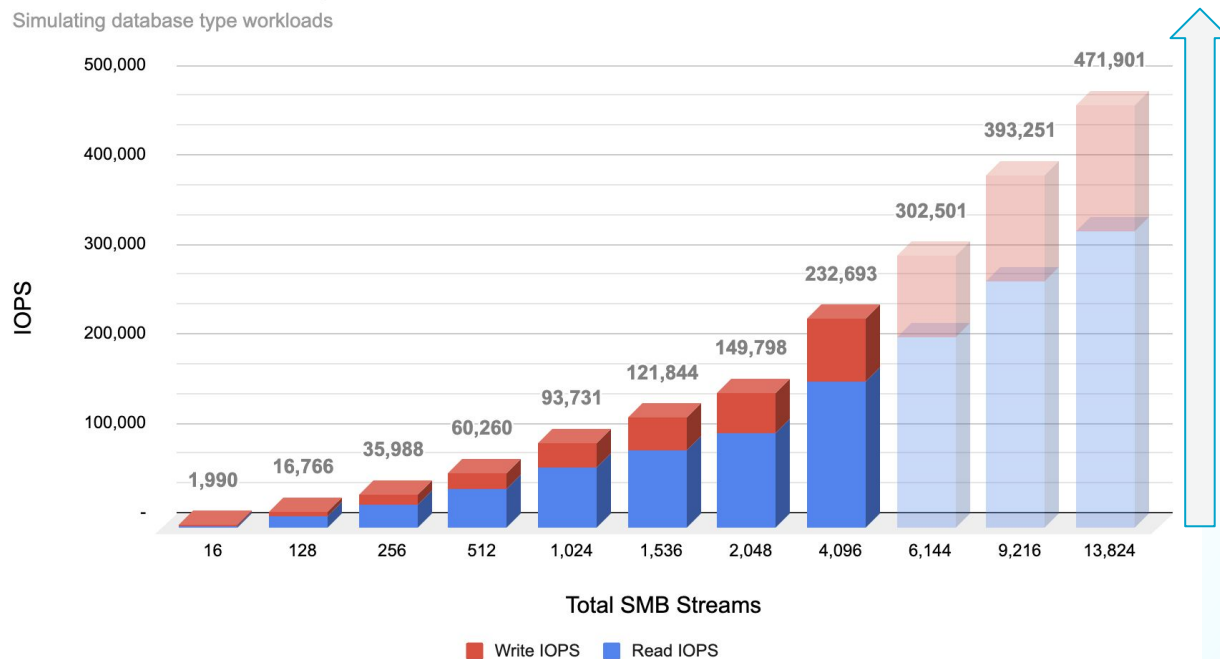
\*\* Benchmark scripts available on Github (see *Appendix*)

\*\* See [here](#) for a live demo of 100+GBps test

# ANQ Example Random IOPS

SMB 70/30 R/W random, 8k block-size

Simulating database type workloads



*IMPORTANT: we can elastically increase from these peaks if required*

- 232,000 Mixed IOPS at 4096 SMB Streams
- FIO 3.36
- 16 Linux clients (D16)
- Accelerated Networking
- Default tunables

\*\* Benchmark scripts available on Github (see Appendix)



# Industry Verticals

# FSI Performance Benchmark

## Results from testing Financial Services Synthetic workload

Feb / 2024

Type of Storage - 60%R/40%W	BW (MB/s)	IOPS	Latency (ms)
ANQ, 1 client, 1 endpoint, 60/40 100% Random Read	1,420	22,724	5
ANQ, 1 client, 1 endpoint, 60/40 Random Mix	474	7,595	16
ANQ, 1 client, 16 endpoints, 100% Random Read	10,301	164,820	3
ANQ, 1 client, 16 endpoints, 60/40 Random Mix	1,760	34,343	14
ANQ, 1 client, 4 endpoint, 60/40 Random Mix	553	8,860	14
ANQ, 1 client, 8 endpoint, 60/40 Random Mix	1,060	16,963	15
ANQ, 2 clients, 16 endpoints/client, 100% Random Read	20,112	325,000	3
ANQ, 4 clients, 16 endpoints/client, 100% Random Read	25,002	289,117	4
ANQ, 4 clients, 16 endpoints/client, 60/40 Random Mix	4,307	64,000	33
ANQ, 4 clients, 4 endpoints/client, 60/40 Random Mix	1,780	28,000	17
ANQ, 4 clients, 8 endpoints/client, 60/40 Random Mix	3,627	56,000	17

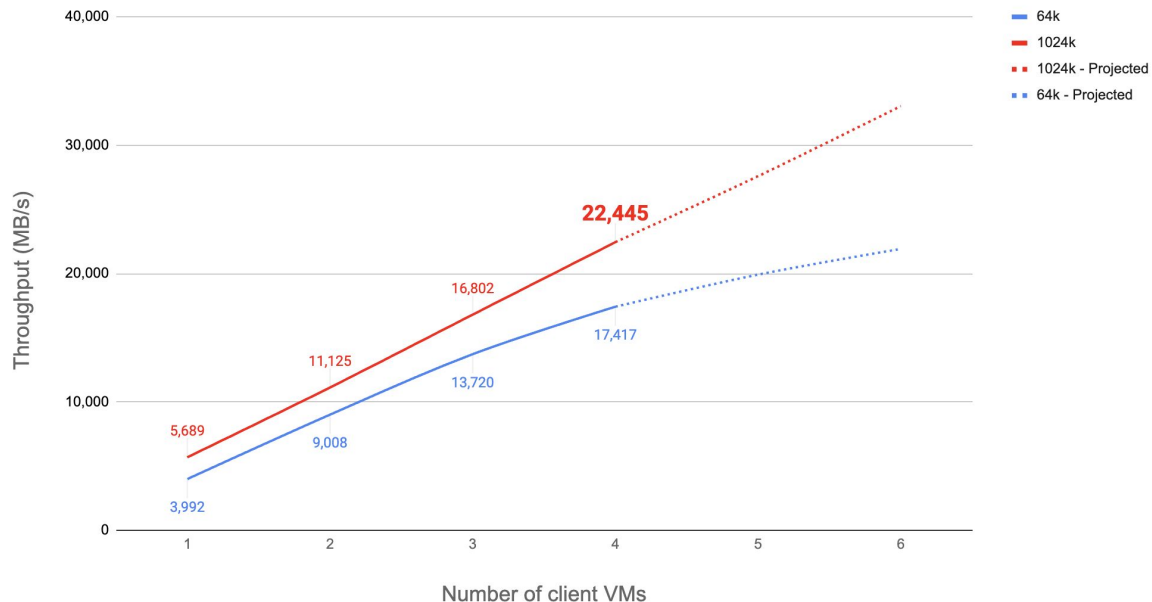
- Single Stream results are constrained with scale-out architecture
- Extreme elasticity at scale

\*\* Based on diskspd benchmark. Parameters used were created for financial services industry based on common batch-oriented production applications

# FSI: SMB Sequential Read

## Multi-Stream Throughput (100% Random Read)

Contrast by block size



- Default cluster size, min capacity
- Each endpoint mapped to unique letter drive
- SMB Multichannel  
*(enabled by default in ANQ)*
- 256k & 1024k block size tests have similar results
- Random R/W results in Appendix

# Genomics

## Welcome to NVIDIA Parabricks

“The Parabricks **somatic (Somatic Variant Caller)**, **germline (GATK Germline Pipeline)** and **deepvariant\_germline** tools are collections of several other individual tools that are commonly used together, all wrapped up as a single tool.”

## *Why Parabricks?*

Chosen for industry-accepted genomic pipeline computation framework

Results can be linearly extrapolated for additional genomic scale

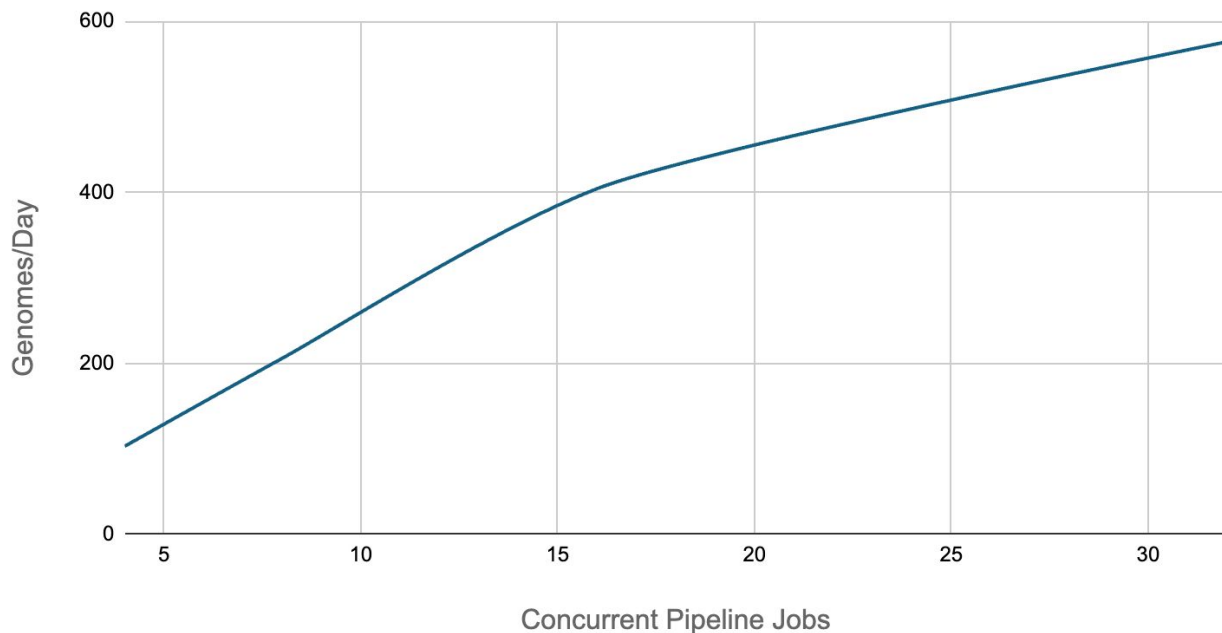
Ballpark pricing simple to forecast

Results are applicable to genomics and bioinformatics personas

# Genomics - Genomes Processed/Day

## Genomes Processed per day

Azure Native Qumulo (Default Cluster) w/Parabricks



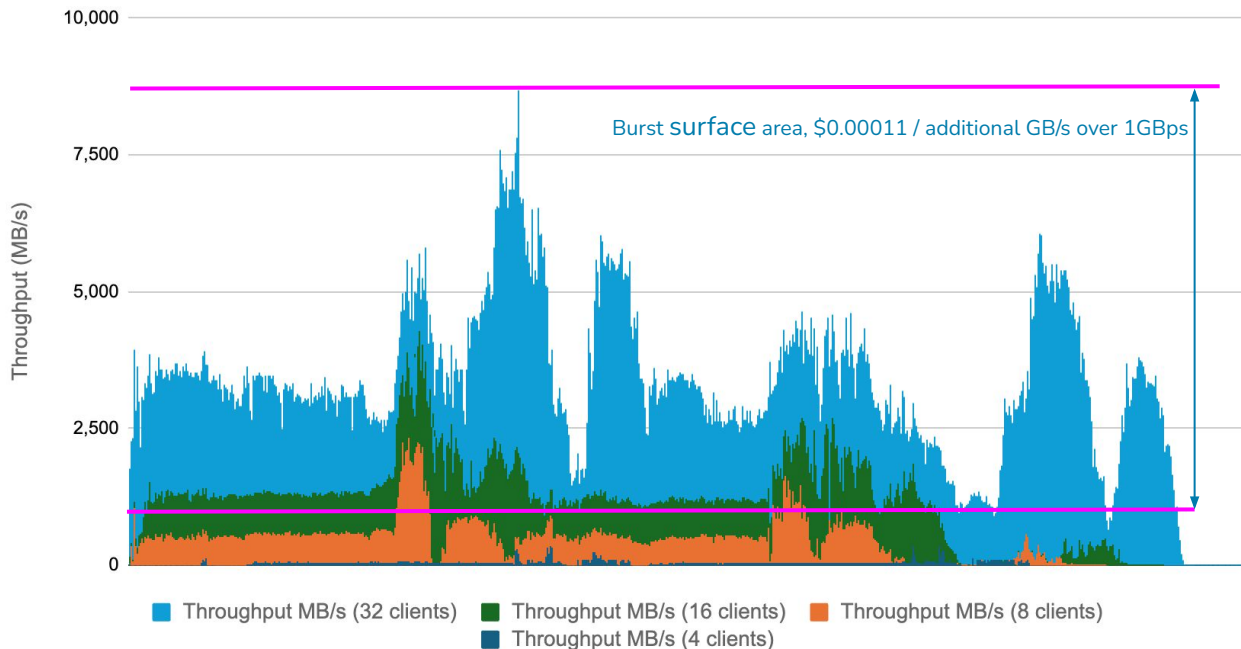
## Benchmark Details

- Germline raw FASTQ data downloaded from the [European Nucleotide Archive](#)
- DNA Identifier #NA12878
- DNA Sample >50x Depth
- VM Size: NC64as\_T4\_v3
- ANQ v7.0.1 (default size)

# Germline Pipeline Throughput Demand

## Germline Pipeline Throughput (MB/s) Concurrency Comparison

Azure Native Qumulo v7.0.1 Default Cluster



- Peak throughput North of 7GB/s
- Tested through 2048 CPU Cores & 64 T4 GPU Cores.
- Remember – performance can always be elastically increased to run more jobs

1 GB/s included. Only pay for the bandwidth you use!

# Pricing

**Fantastic performance, cost optimized for customers**

# ANQ Pricing

Simple and easy to forecast @ <https://azure.qumulo.com/calculator>

### Enter Your Workload Details

Storage class  
 Hot  Cold

Average storage used (in TB)  
100

Specify the size of your workload in terabytes. You are only charged for the space you use.

Percent of time the throughput is above 1 GBps  
15%

Most workloads are always below 1 GBps of throughput.

When the throughput is above 1 GBps, what is the average throughput?  
2 GBps

It's easy to overestimate the performance you'll utilize.

Percent of time IOPS is above 10,000  
5%

Most workloads are always below 10,000 IOPS.

When above 10,000 IOPS, what is the average IOPS?  
11,000

It's easy to overestimate the performance you'll utilize.

Azure region  
East US 2

Select your desired Azure region.

## \$3,817.18/month

For a 100 TB dataset on the East US 2 Azure region

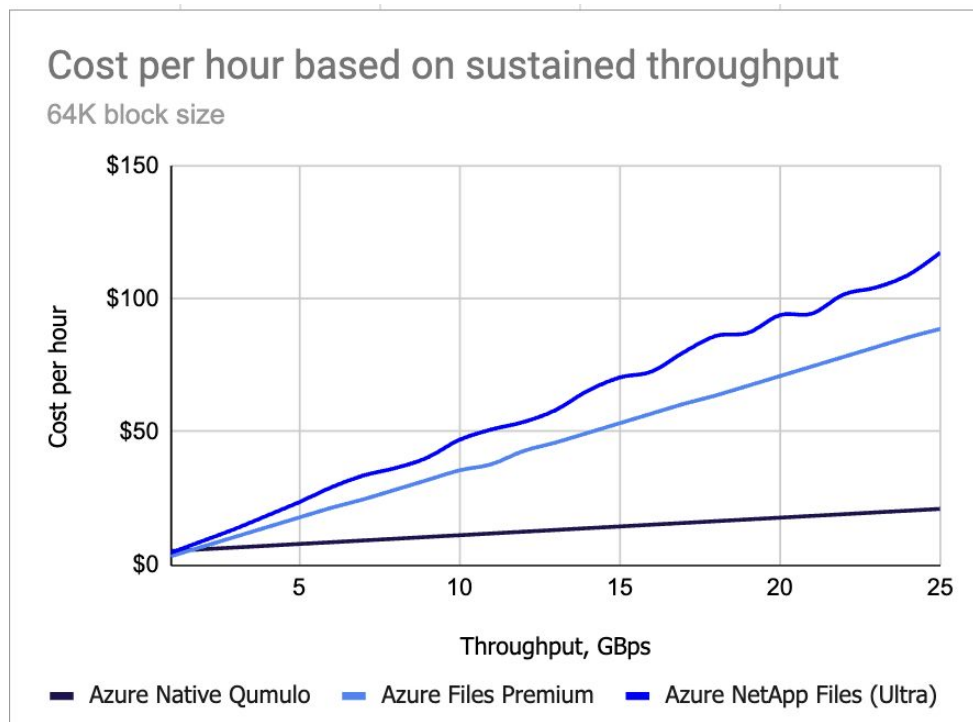
Item	Description	Cost
100 TB	ANQ Hot using 100 TB on the Azure East US 2 region 100 TB x \$0.037/GB-month Includes up to 1 GBps of throughput and 10,000 IOPS	\$3,700.00
1 - 2 GBps	Additional 1 GBps of throughput for 6,480 minutes \$0.00011416 /TB-minute x 100 TB x 216 minutes x 30 days	\$73.98
1,000 IOPS	Additional 1,000 IOPS for 2,160 minutes \$0.02 x (1,000 IOPS / 1,000) x 72 minutes x 30 days	\$43.20

Download Estimate as PDF

- ANQ Hot: \$37/TB month
  - Throughput – only charged for usage above 1GB/s
  - IOPS – only charged for usage above 10k IOPS
- ANQ Cold: \$9.95/month
  - 5TB/month Throughput included
- That's it, no hidden fees
  - Easy to forecast



# Cloud Storage Price / Perf Comparison



- Each offering was sized to meet the ANQ achieved benchmark throughput
- Based on ANQ v7.0.0 benchmark
- Online sizing tools used to configure each solution
- Public pricing (list)
- Multiple volumes required for ANF and AFP

# Invitation to validate the benchmarks

## Benchmarking code available on Github

Each of the benchmarking harnesses were designed to programmatically scale to perform testing on a multiple number of VM's, Qumulo nodes, and then aggregate results for reporting.

- SMB\_Bench: [https://github.com/qumulokmac/smb\\_bench](https://github.com/qumulokmac/smb_bench)  
Scale out SMB benchmark harness based on FIO, PowerShell, and Cygwin
- NFS\_Bench: [https://github.com/qumulokmac/nfs\\_bench](https://github.com/qumulokmac/nfs_bench)  
Scale out NFS benchmark harness based on FIO and bash
- DSPD\_Bench: [https://github.com/qumulokmac/dspd\\_bench](https://github.com/qumulokmac/dspd_bench)  
Scale out NFS benchmark harness based on Diskspd, PowerShell, and Cygwin



***Engage with us!***

<https://www.reddit.com/r/qumulo/>



# Appendix

# Customers love us!


- NPS above 0: Indicates that there are more Promoters than Detractors.
- NPS between 0 and 30: A positive moderate score.
- NPS above 30: An excellent score, indicating strong customer loyalty and advocacy.

FY24Q4 NPS  
score is **78!**



# Other industry NPS scores

NetApp Overall NPS



Dell Technologies Overall NPS

