SNIA CSTI | CLOUD STORAGE TECHNOLOGIES

Create a Smarter and More Economic Cloud Storage Architecture

November 7, 2018

Today's Presenters

SNIA CSTI | CLOUD STORAGE TECHNOLOGIES

Eric Lakin
University of Michigan

Alex McDonald
NetApp

Michelle Tidwell
IBM

Mike Jochimsen
Kaminario

SNIA Legal Notice

SNIA CSTI | CLOUD STORAGE TECHNOLOGIES

- The material contained in this presentation is copyrighted by the SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

SNIA-At-A-Glance

170
industry leading organizations

2,500
active contributing members

50,000
IT end users & storage pros worldwide

Learn more: **snia.org/technical** **@SNIA**

SNIA CSTI | CLOUD STORAGE TECHNOLOGIES

What We Do

Educate vendors and users on cloud storage, data services and orchestration

Support & promote business models and architectures: OpenStack, Software Defined Storage, Kubernetes, Object Storage

Understand Hyperscaler requirements Incorporate them into standards and programs

Collaborate with other industry associations

Agenda

- So You've Decided to Build a Cloud Environment
- Economic Considerations
- Technical Considerations
- Gaining Business Insights
- Recap and Next Steps

ECONOMIC CONSIDERATIONS

Accounting Principles: An Overview

SNIA CSTI | CLOUD STORAGE TECHNOLOGIES

- Economic & accounting considerations
  - What are CAPEX and OPEX?
  - How can you relate storage to cost?
    - Activity-based models/costing
- How to calculate on-premises vs. cloud costs
- Hybrid storage flexibility cost models
- New purchasing & costing paradigms

CAPEX vs OPEX

- You remembered to get an accountant, right?
- Lots of accounting talk, but what does it mean?
- CAPital EXpenditure vs OPerational Expenditure
  - CAPEX is the cost of developing or providing parts for a product or system
  - OPEX is an ongoing cost for running that product or system

"OK, the cost of providing something is CAPEX, and the cost of running it is OPEX, right? So when I'm asked to move from CAPEX to OPEX, what I do is… err… hmm… uhh…"

Fixed vs. Variable Costs

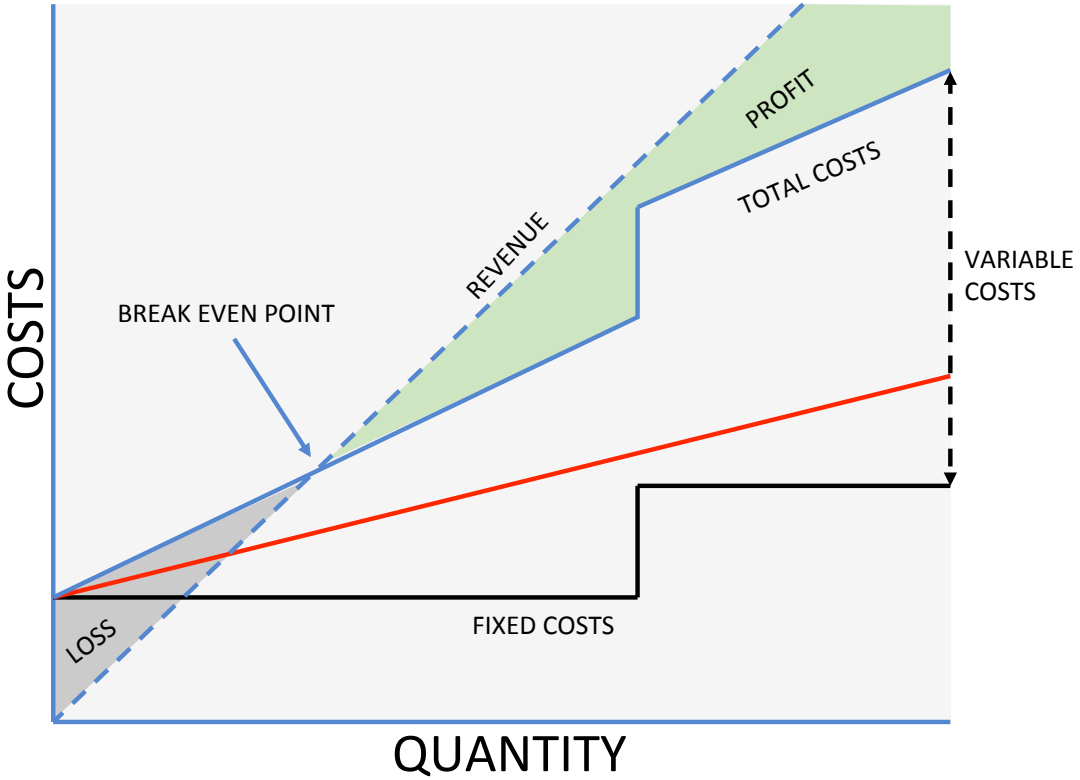

CAPEX to OPEX

- Shift to OPEX tries to remove the cost of ***providing*** the service
  - They are generally (but are not always) fixed costs
- Costs don't disappear; ***running*** the service becomes more expensive
  - Fixed moves to variable cost; the factor varies
- Advantage to variable costs or OPEX
  - Easier to manage variable costs
    - Based on consumption or activity
    - Fixed costs have long payback periods
  - Stop doing it and cost should reduce to zer
  - Doesn't require investment or borrowings
    - "Pay As You Go" (PAYG)
  - Reduces the cost of "being the first"
  - Easier to sell to consumers

Activity Based Costing

- Assumption to ABC
  - Doing something costs, but doing nothing costs nothing
- Need to identify all the activities in providing the service
  - Scope & Planning
    - Performance characteristics, I/O loads, network requirement
    - Security & privacy
  - Provisioning
    - Data transmission & network costs
      - Bandwidth & latency
    - Backup & archive
  - Management & Servicing
    - Capacity growth
    - Maintenance
  - Decommissioning
- Assign costs to each activity & calculate:
  - Unavoidable fixed & CAPEX
  - Minimized variable & OPEX

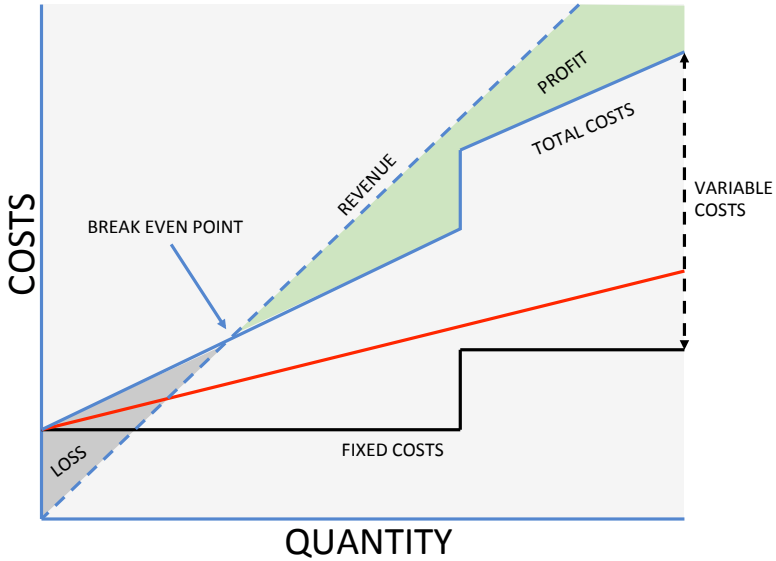

Activity Based Costing & OPEX

- Pitfalls:
  - Shared costs: costs that aren't easily split by activity
  - Overhead costs: costs there regardless of activity
  - Missing costs: not all costs can be discovered or captured
  - Unrecoverable costs: not all costs can be recovered
  - Behavioral costs: not all user behaviors are controllable
    - Chargeback vs showback
- Disadvantages of variable costs or OPEX
  - Can be more expensive in the long run
  - Can have significant tax disadvantages
  - Over-consumption & underestimating demand can introduce step costs
    - Can sometimes be as significant as fixed costs

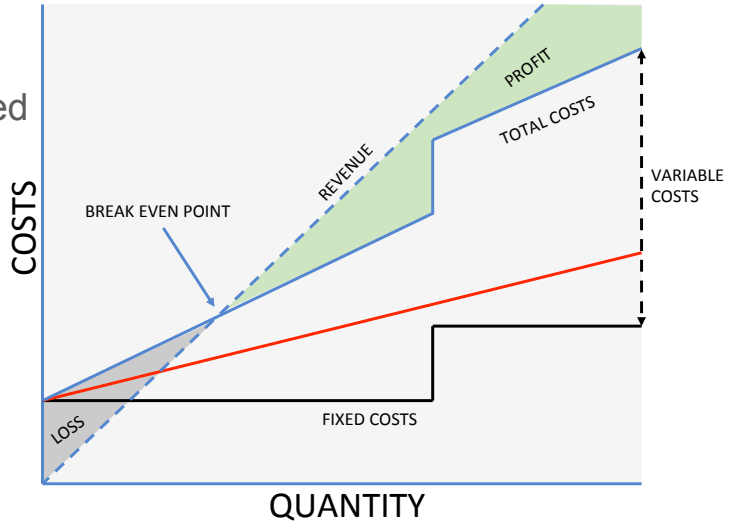

How to Calculate On-premises vs. Cloud Costs

- Question: is the cloud cheaper than on-premises?
- Answer: you need to ***calculate*** it, not ***speculate*** about it
  - $/GB is not the cost of providing a storage service
- Modeling ***service flexibility*** can help
  - Service or application "fragmentability"
  - Size & type of capacity (online, nearline, backup, archive)
  - Bandwidth & latency (width & length of pipes)
  - Legal & regulatory (GDPR, HIPAA)
  - User location (geography, office, web)
  - Compute required (heavy to light)

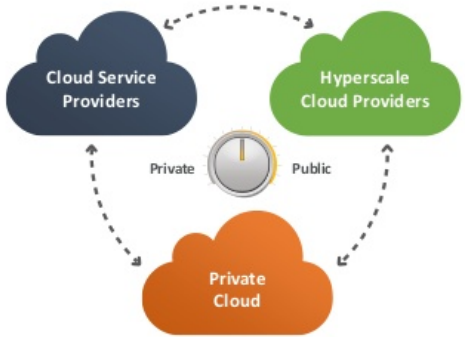

New Purchasing & Costing Paradigms

- Mixed service models
  - Gold, silver, bronze
  - Profile based charging
- Cloud brokering
  - Data motion between providers
- Data at the edge & data reduction
  - Move only what you need
- Moving compute to the storage
  - Compute is lightweight

$¥£

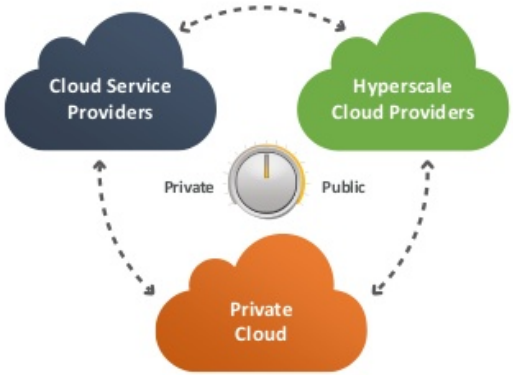

TECHNICAL CONSIDERATIONS

Technical Considerations

Hybrid Cloud Models to Consider:

- Cloud as Remote Storage
- Extending on Premises Storage to the Public Cloud for Primary or Secondary Use
- Seamless Workflow Integration between On Premises and Cloud (True Multi Cloud)

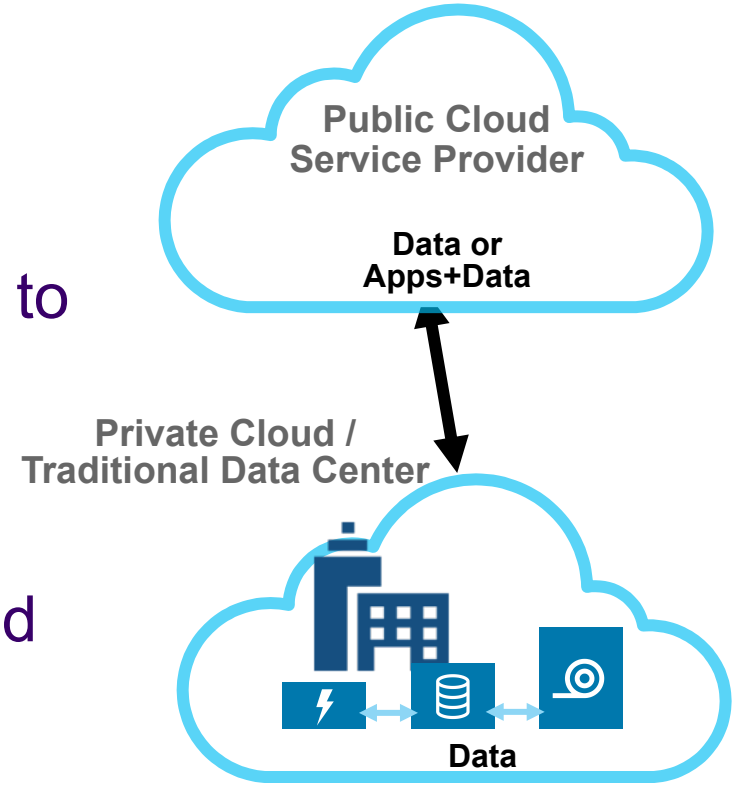

Cloud as Remote Storage

- Use Cases
  - Archive
  - Backup
  - Tiering
  - Apps with Native Object Store Support

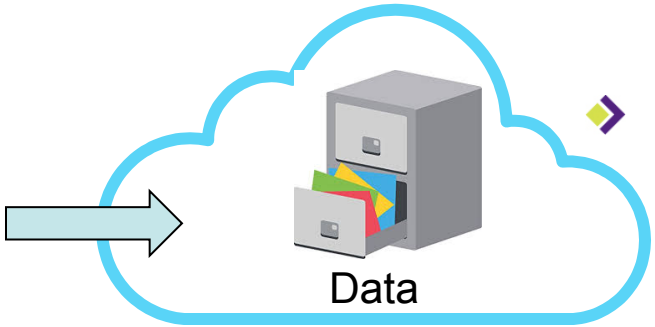

- Technical Considerations
  - Usually target Object Storage for Lower Cost (S3, SWIFT, Blobs)
  - Local Rehydration Options on Cloud can help to avoid egress if you need to use the data
  - For Tiering: Consider built-in capabilities in Storage Arrays to move data to cloud vs Gateways – Reduces Cost and management complexity
- xSPs and End Users should plan for:
  - High Latency Tolerance (by Apps or Enterprise)
  - Must tolerate longer RPO/RTO as a DR solution
  - Ensure some form of encyrption is available
  - Rehydration methods for recovery

Extending On Premises Storage to the Public Cloud

SNIA CSTI | CLOUD STORAGE TECHNOLOGIES

Use Cases

- DR/Business Continuity
- Compute + Data Offload
- Test / Dev / DevOps
- Leverage Cloud for Analytics

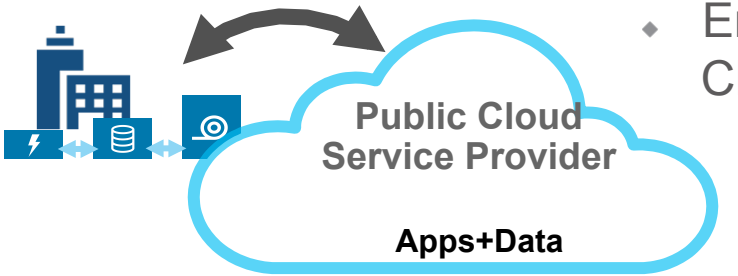

Technical Considerations for SPs or End Users

- Cloud Data Center Network Latency will limit storage performance in Public Clouds.
  - Utilize features to localize subnets of compute and storage resources when possible.
- Choose a Solution Built for High Availability
  - Data Management with HA or Multi Node Configurations
  - Consider Distances and Select Carefully on Portals
- Ensure data is stored in a way that can leverage Cloud Services of multiple Cloud vendors.
  - Look for Solutions which use Standard APIs
  - Integration or use of Container Technology
  - Open Source

Extending On Premises Storage to the Cloud (Cont'd)

- When Extending Block Workloads to the Cloud ensure:
  - Your data management layer provides backend infrastructure Optimization
  - Dynamic Provisioning of High Performance Storage when it's needed and not before.
  - Dynamic Tiering of Hot Data to Faster Storage and back when it's no longer hot
  - Common Management between on premises and the cloud can streamline deployment and maintenance of the cloud environment.
  - Examples: Data Deduplication, Compression, Thin Provisioning, Unmap, EasyTier

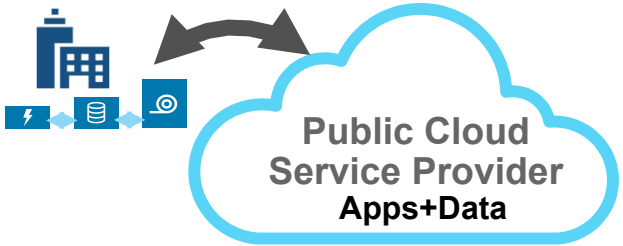

Seamless Workflow Integration

Use Cases

- Utilize Cloud Services
- Bridge Cloud Native and Traditional IT
- Service Multiple LOBs in an Enterprise with a single management platform.

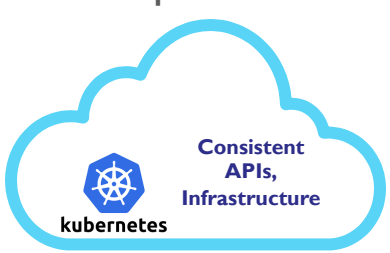

On-Premises Infrastructure

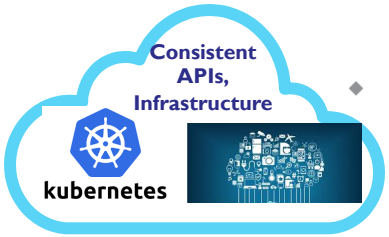

Public Cloud Services and Infrastructure

Technical Considerations

- Use Consistent APIs and Orchestration across on premises, private and public Clouds
  - Ex: Kubernetes / Docker / VMware
- Intelligent Data Distribution and migration between Private and Public Clouds (Locality of Data to Compute)
- To support Cloud Native environments, consider a solution which can provide *on premises cloud-like experiences* with containerized apps and data
- Extending the Cloud Native environment provides the greatest flexibility to use Cloud microservices, analytics and AI tools. (Examples: Microsoft Azure Stack, IBM Cloud Private).

Take-Aways

- Build a Cost Model to evaluate the right hybrid cloud solution for your business.
- Run a POC to ensure the solution fits your use case
- Test for:
  - Capabilities meet the advertised features of solution
  - Performance of Applications if you are running workloads
  - I/O Latency Meets requirements of both on prem workload and Cloud deployments (if running synchronized data/workloads)
  - An exit strategy if solution is no longer available – Multi Cloud options are important.

GAINING BUSINESS INSIGHT

What Have we Learned - Economic

- Understand your costs (fixed and variable)
  - How will you know if on-prem or public cloud is more expensive?
- Determine your CAPEX vs. OPEX strategy
  - Remember that accountant Alex talked about?
  - FASB 842 and IFRS 16 may impact your leasing strategy
  - Do you have the capital to invest?
- Aligning asset acquisition to your company financial model
  - Does Activity Based Costing drive your decisions?
- Emergence of new purchasing models

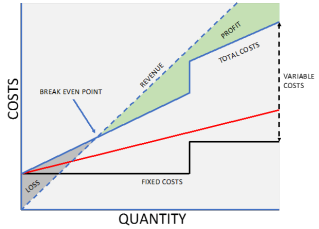

What Have we Learned – Technical

- Use cases for on-prem or hybrid cloud environment
- Block, file and object have different technical considerations in a cloud environment
- Rich data services (dedup, compression, thin provisioning, etc.) are important
- A common data format enables portability in the multi cloud
- Know the performance impact of data locality when implementing a multi cloud
- Orchestration and seamless management are critical when building a cloud

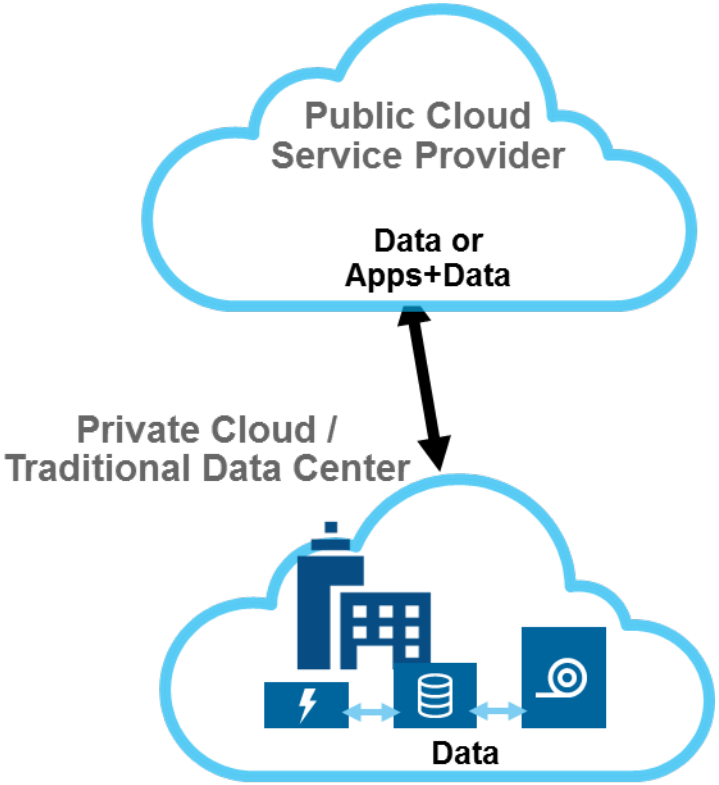

Next Steps

- Include your finance team in building your RFI/RFP
- Understand your cost!
- Know (or create) your strategy for on-balance sheet vs. off-balance sheet equipment acquisition
- Determine your hybrid cloud strategy BEFORE you start

After This Webcast

SNIA CSTI | CLOUD STORAGE TECHNOLOGIES

- Please rate this webcast and provide us with feedback
- This webcast and a PDF of the slides will be posted to the SNIA Cloud Storage Technologies Initiative website and available on-demand at https://www.snia.org/forum/csti/knowledge/webcasts
- A full Q&A from this webcast will be posted to the SNIA Cloud blog: www.sniacloud.com/
- Follow us on Twitter @SNIACloud

Thank You