



Flash Memory Summit

# Disaster Recovery in a Hybrid Cloud Environment

Kyle Grossmiller

Senior Solutions Architect // Pure Storage

@kylegrossmiller

# Data Interaction is Different in the Cloud

## On-Prem & Hosted



### Resiliency

Highly Reliable Arrays, Built-in Snaps & DR

### Efficiency

Thin Provisioning, Deduplication, Compression

### Cost Considerations

Capacity Planning, resources once purchased are fixed

## Cloud



### Resiliency

High Availability for \$, Higher Durability, Globally Replicated

### Efficiency

Thick Provisioned Capacity & Performance

### Cost Considerations

Ingress/egress charges, Cold vs Warm, performance (IOPS) can be wasted

# Different Location, Similar Problems

Choices, trade-offs, and sprawl


The public cloud has an ever-increasing number of storage options, snapshot limitations, and configuration parameters

Makes it difficult to understand, track, and implement the right option at the right time.

Storage option sprawl is a significant and potentially costly problem—especially at scale

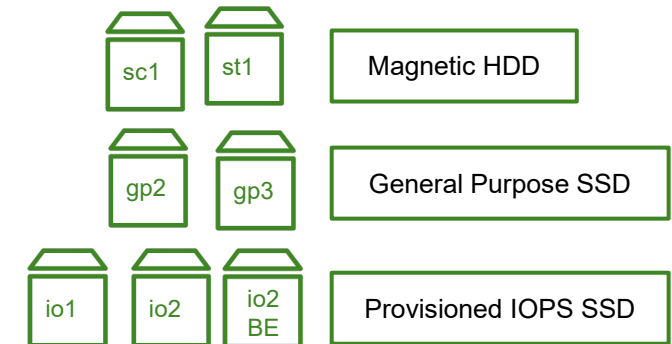



Azure Managed Disk



Amazon Elastic Block Store

Options



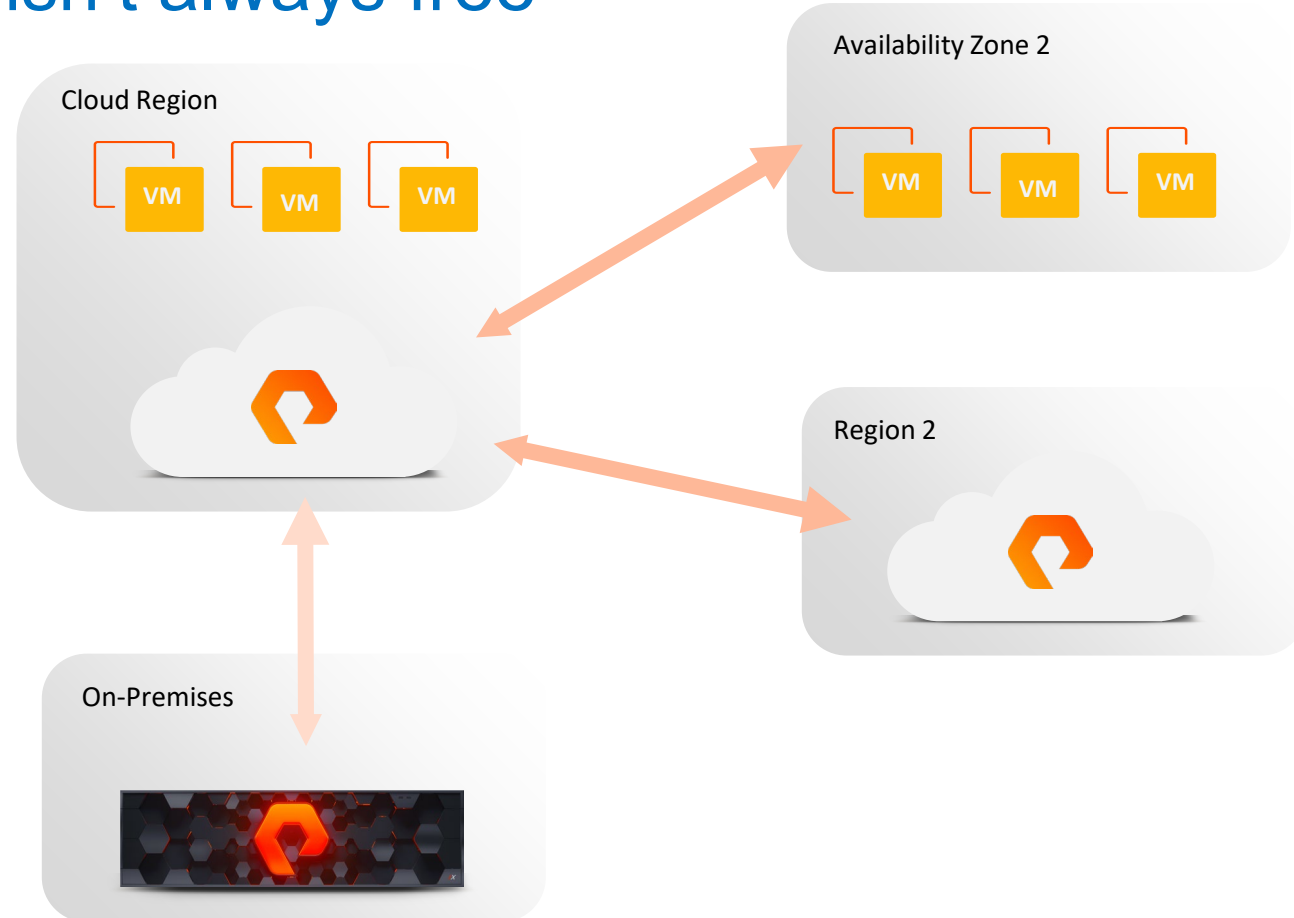
Tradeoffs

- Capacity vs. Throughput
- Capacity vs. IOPS
- Performance vs Latency
- Multi-Attach: Shared Volumes
- Snapshot capabilities

# Ingress / Egress Charges from Data Mobility

## Moving data in the cloud isn't always free

- Both AWS and Azure can charge on both sides of a data migration depending on situation
- Managing cost requires knowing the nuances of each scenario
- Minimize the impact with data reduction

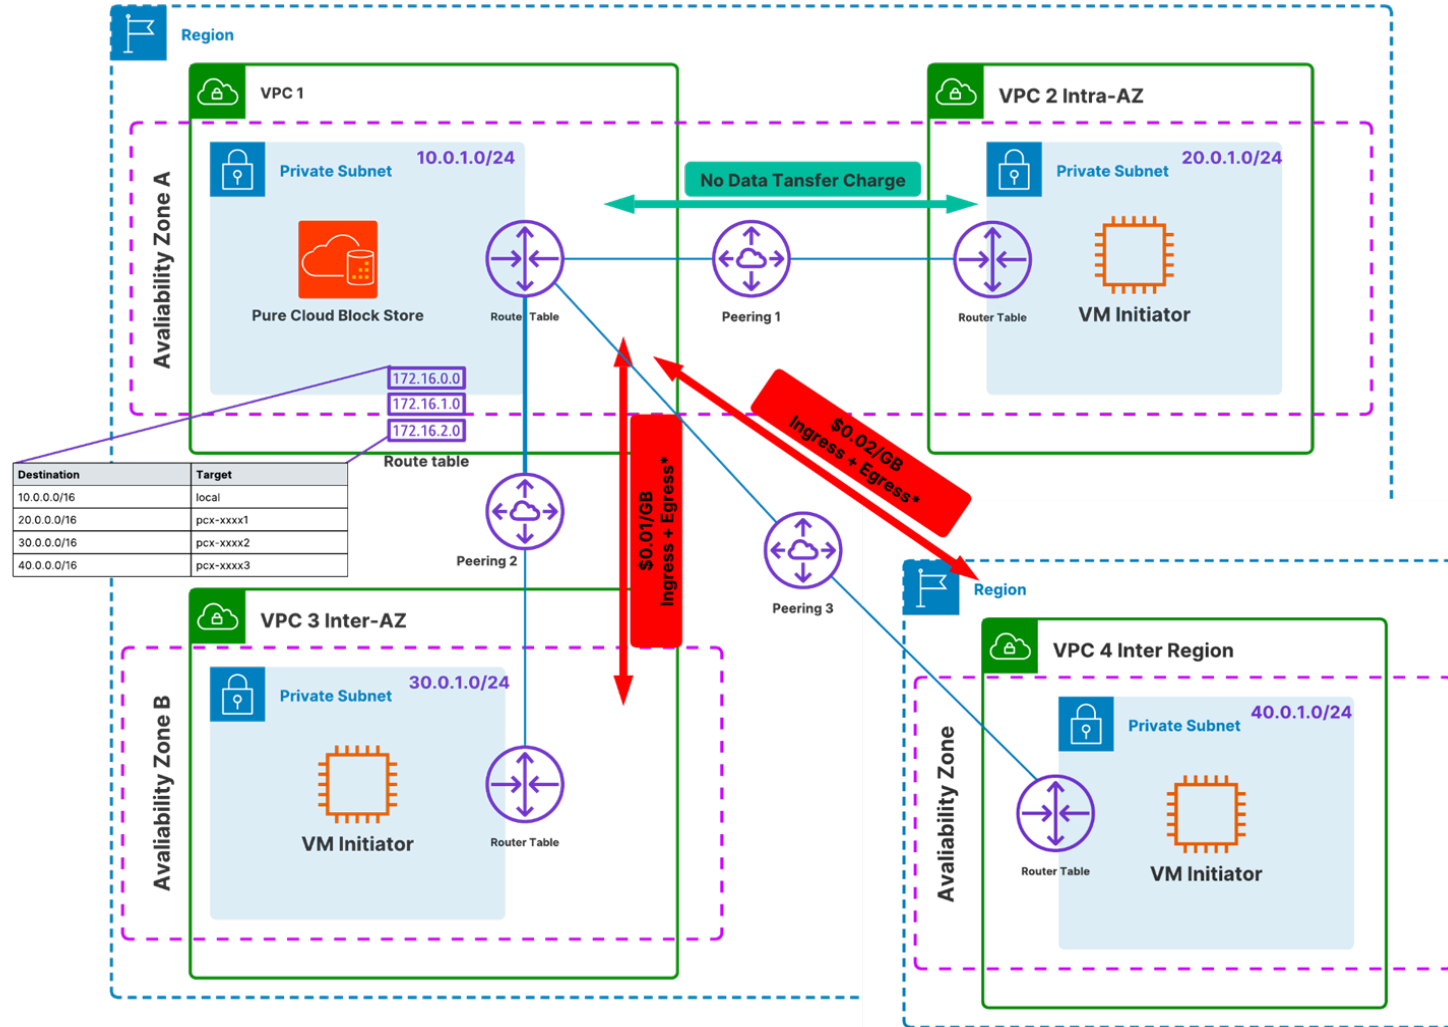


\*In AWS, there is no transit cost as long as VPCs are in the same availability zone



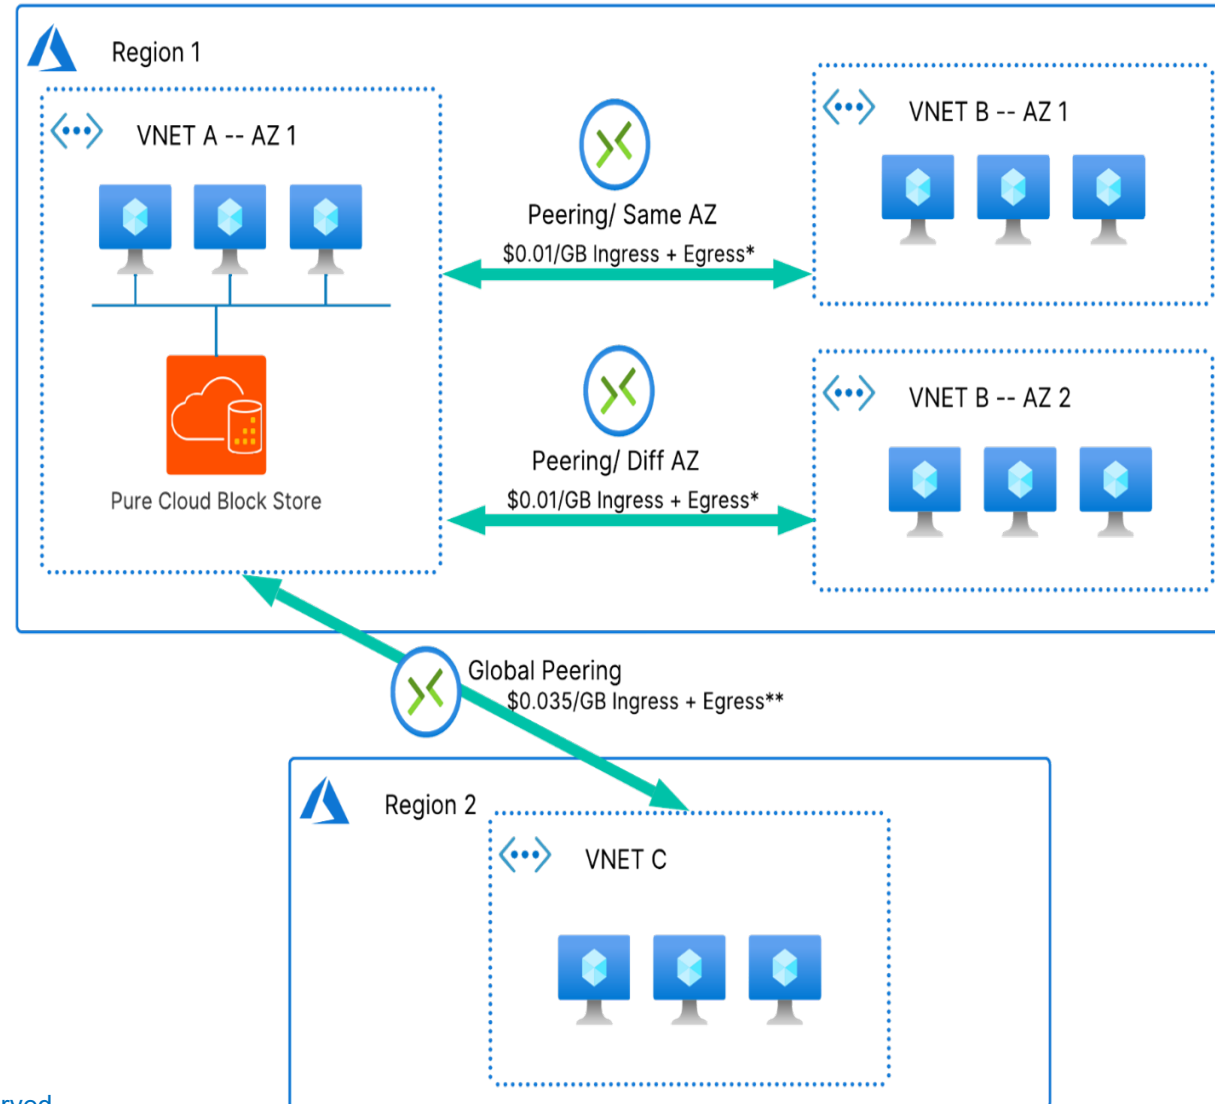
# Host Connectivity – Inter VPC on AWS

## VPC Peering



\* Egress charges as of Feb-2022. Prices subject to change.

# Transferring Data in Azure - VNET Peering

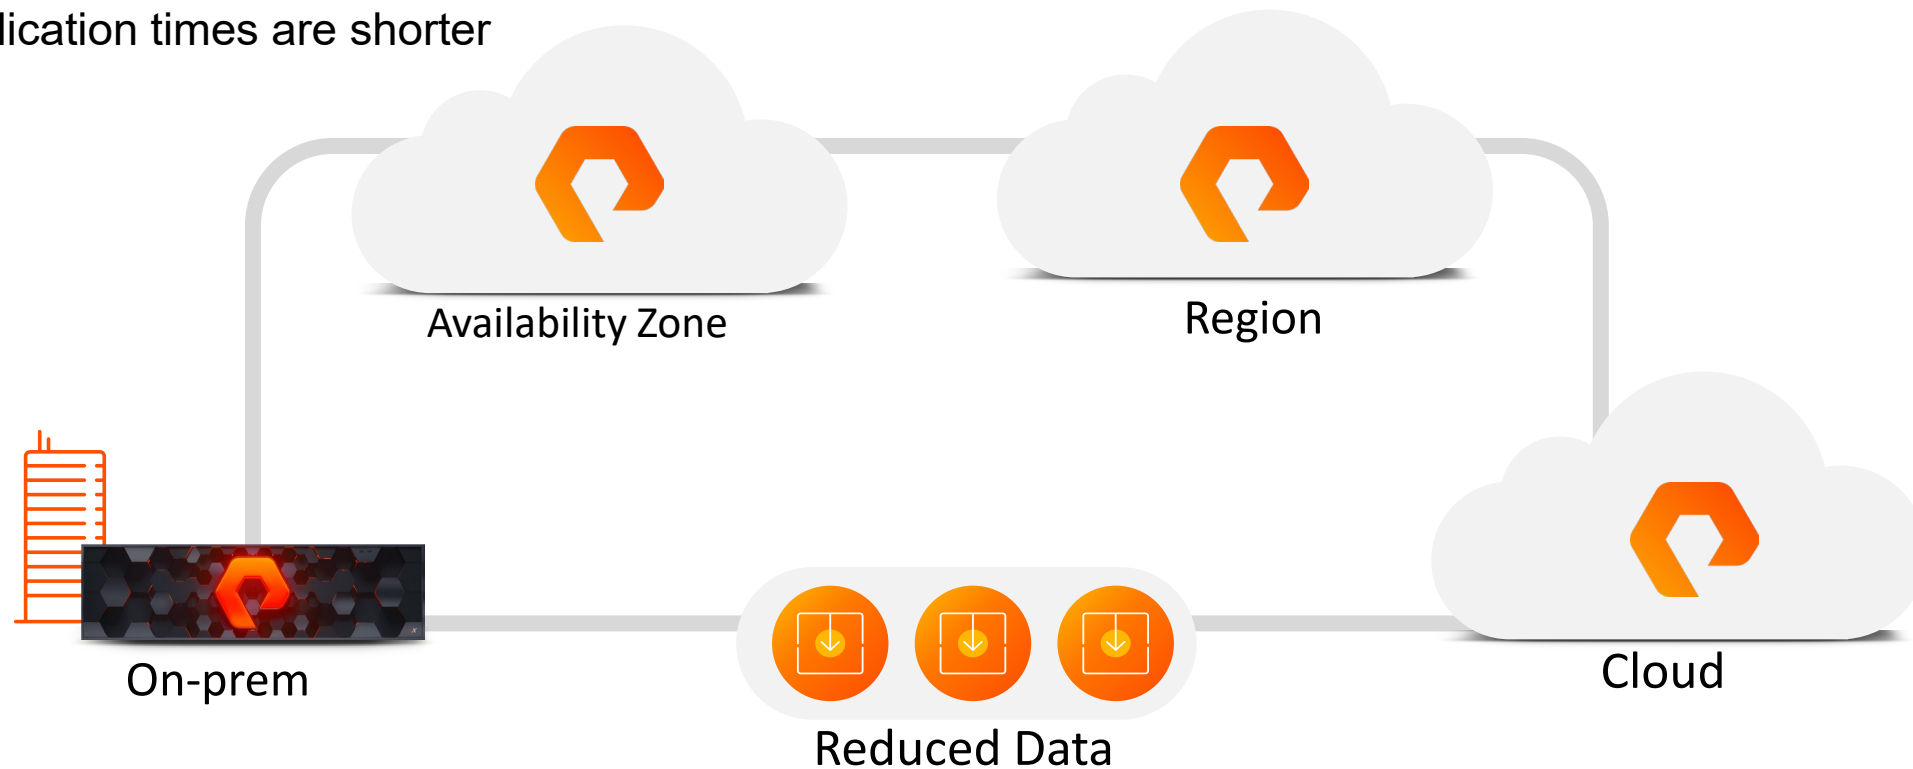




# Optimizing Replication Reduces Bandwidth

Preserve data compression and deduplication when transferring between Pure products

- The storage footprint is reduced, lowering Cloud costs
- Data transfer costs and network utilization are minimized
- Replication times are shorter





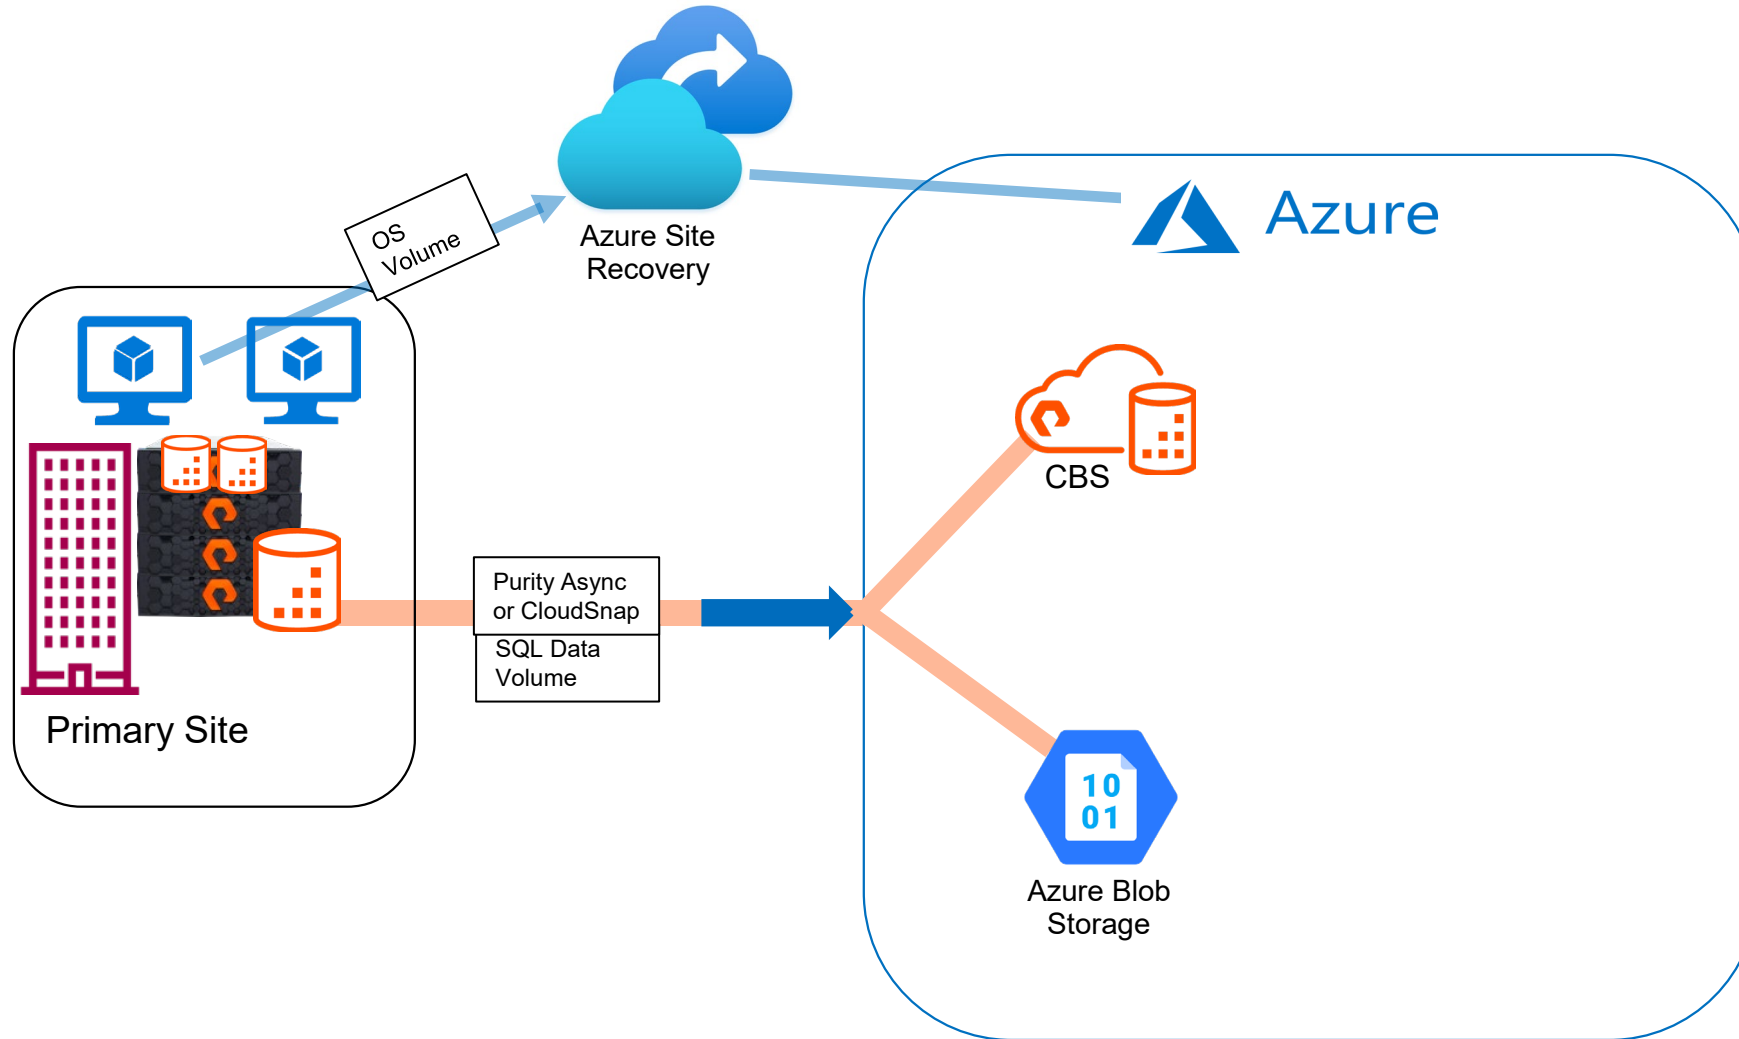
# Recovering Boot OS vs Data Volumes

- A wrinkle of the public cloud is that data volumes are typically not stored on the same storage type/array as boot volumes.
  - Capacity, performance and durability are just some reasons why.
- The tools and processes to fail-over and fail-back your operating system/application vs data volumes in a DR event will be different.
- Azure Site Recovery (ASR) can handle automated failover of operating systems and applications, but what about data volumes?



# Azure Hybrid DR Sample Solution

## Basic Components

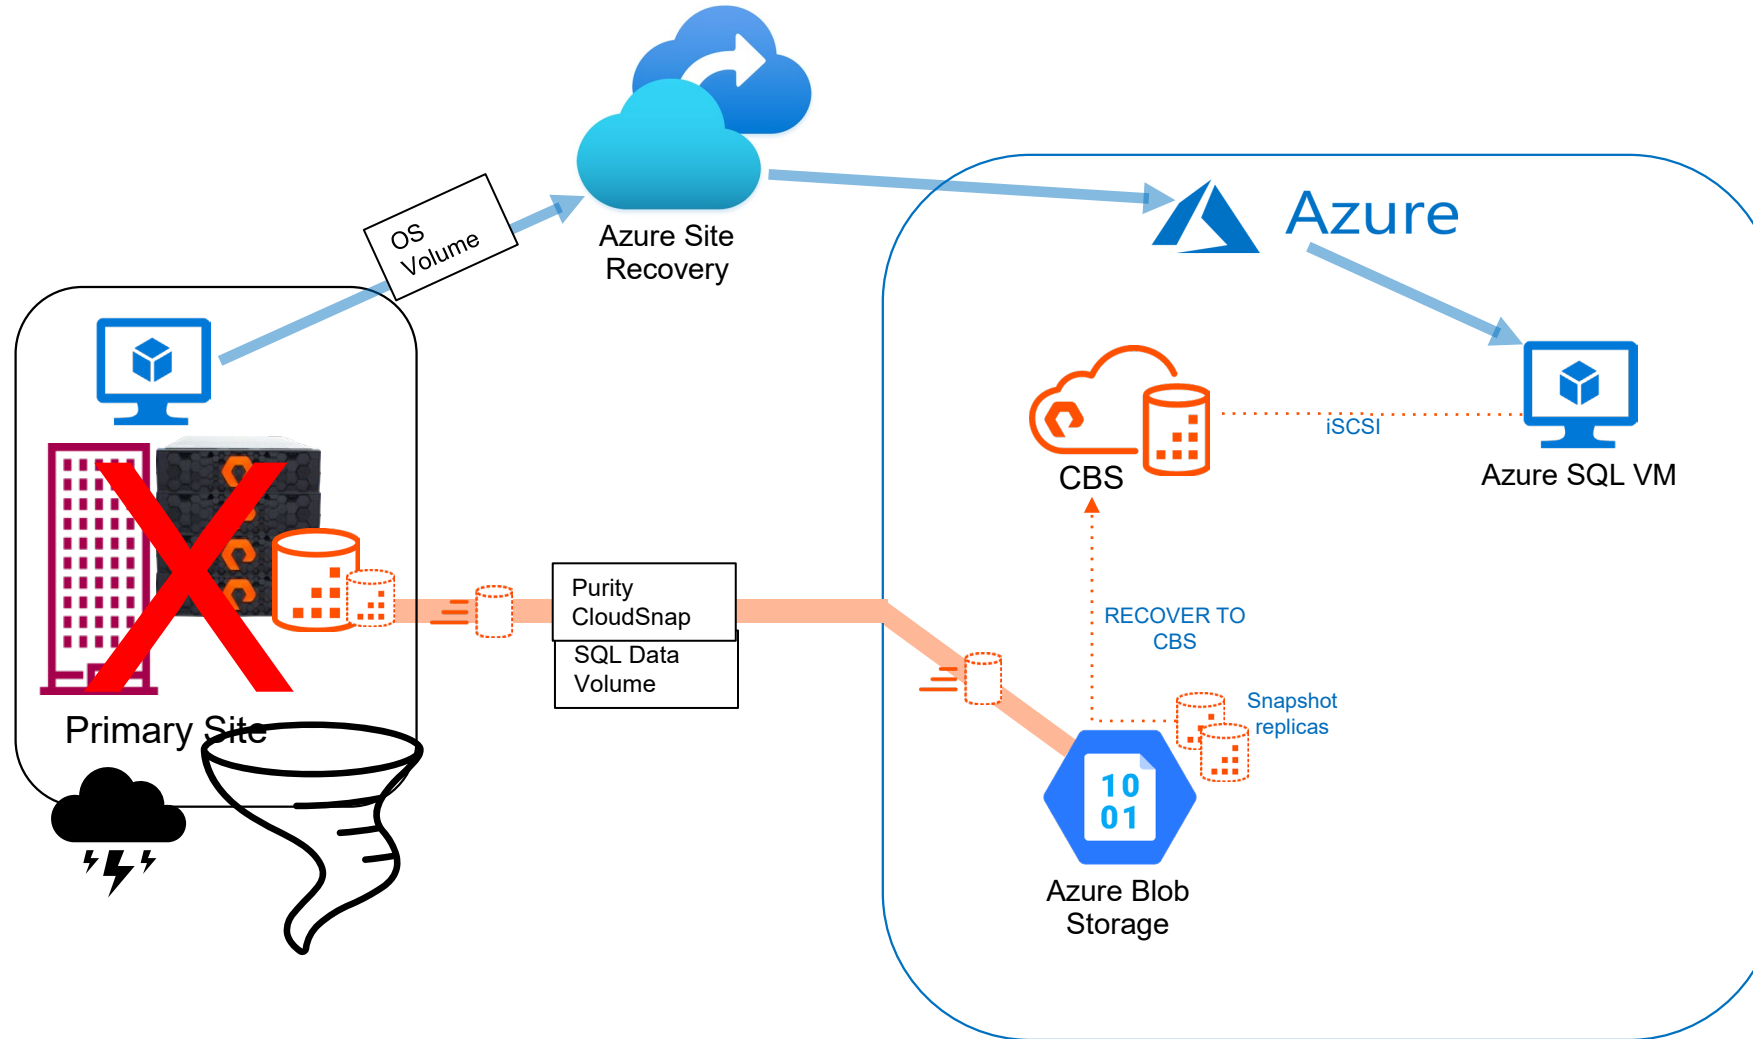


# Azure Hybrid DR Sample Solution

## DR Failover Workflow



Flash Memory Summit

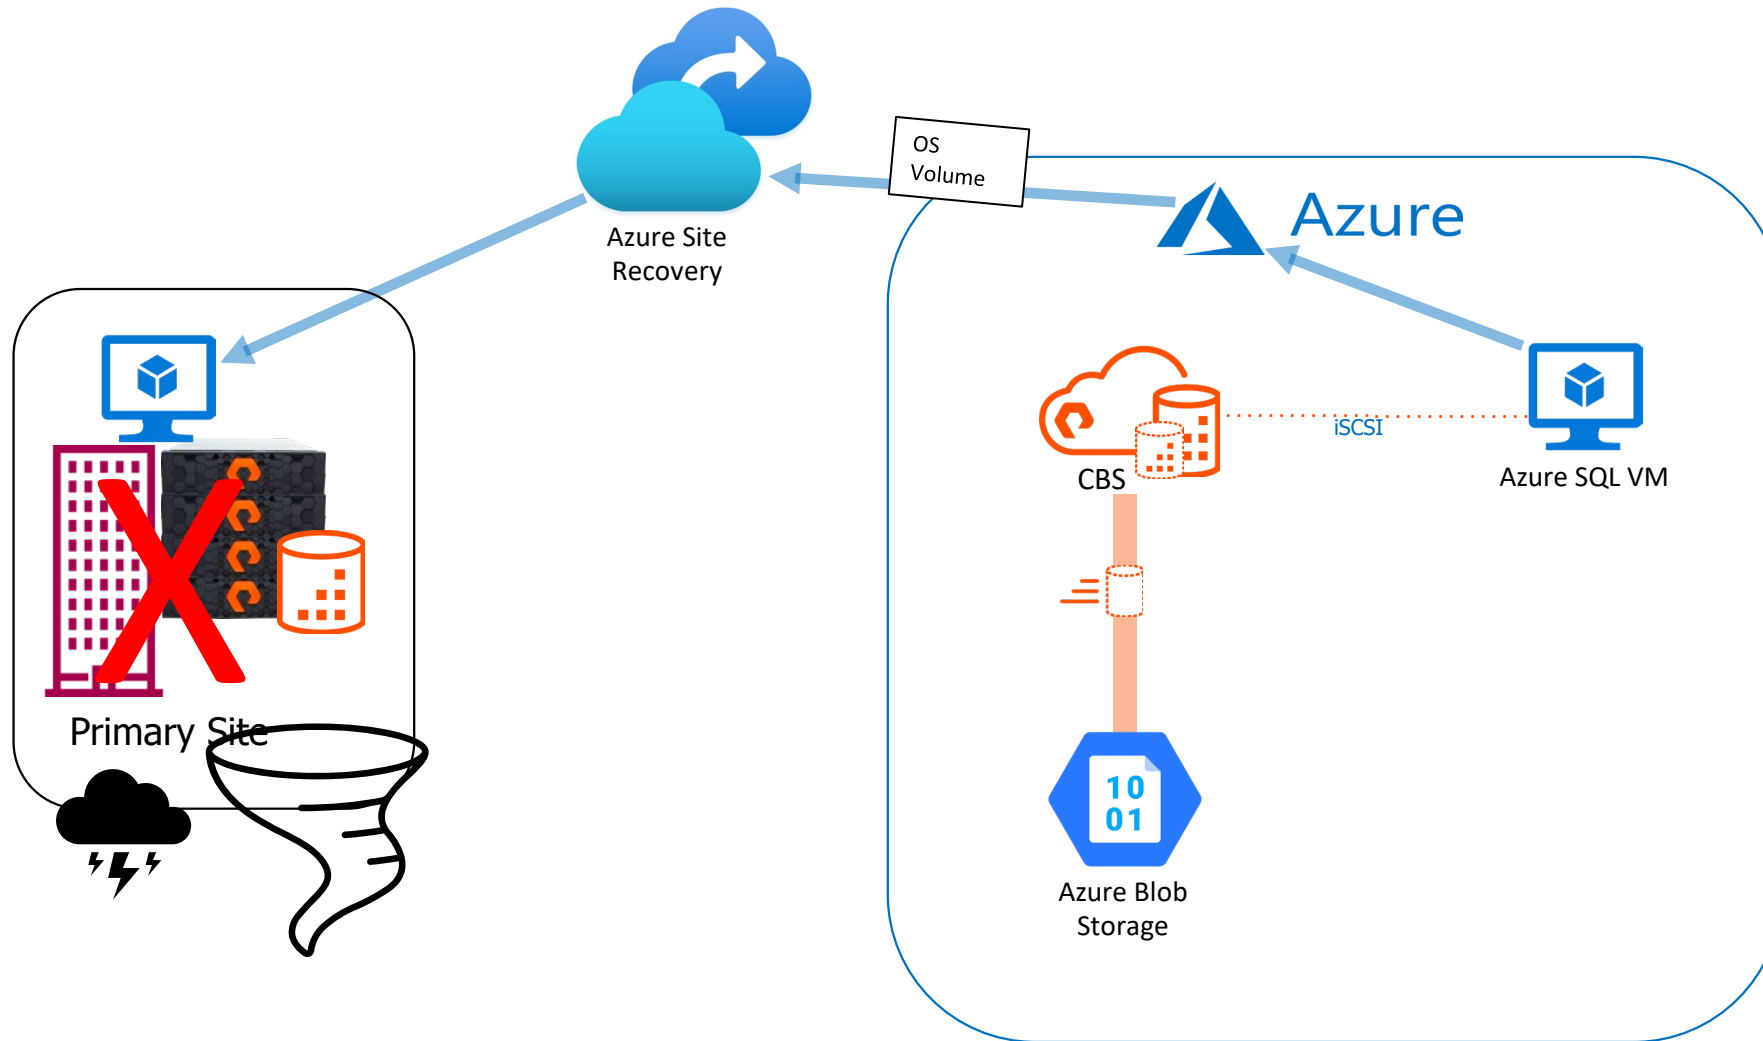


# Azure Hybrid DR Sample Solution

## Failback Workflow



Flash Memory Summit

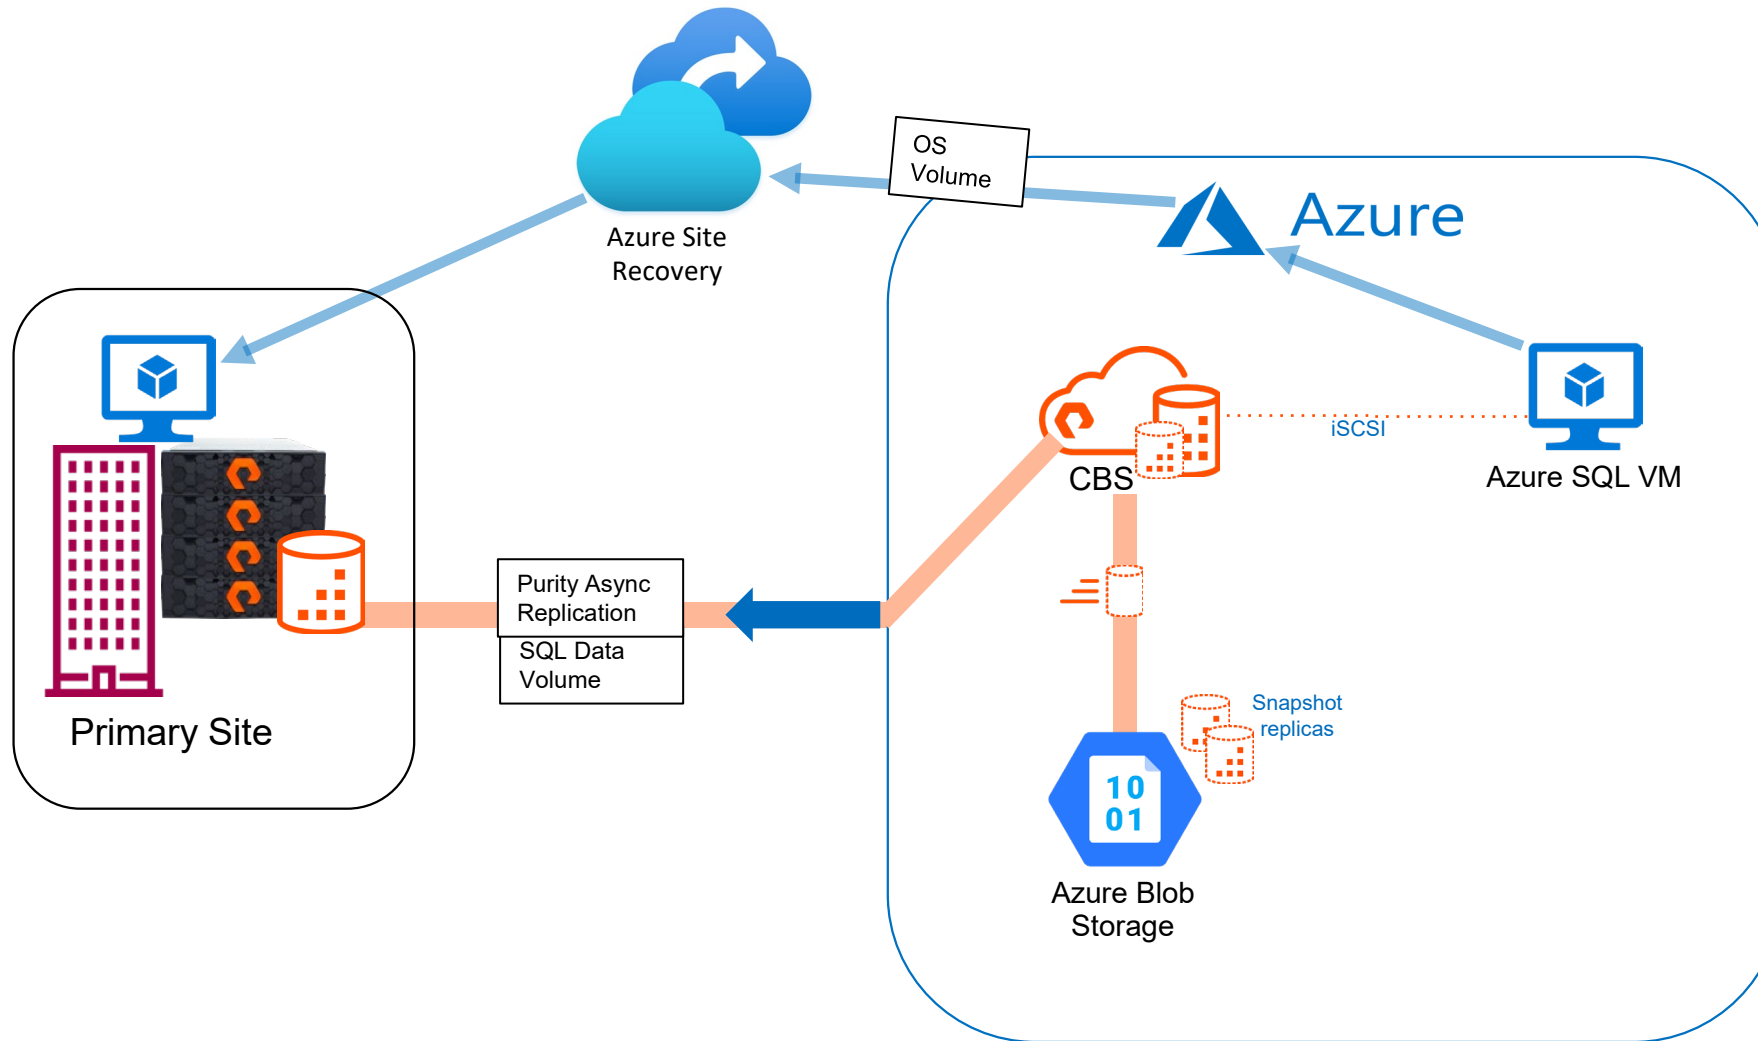


# Azure Hybrid DR Sample Solution

## Failback Workflow



Flash Memory Summit



# Azure Hybrid DR Sample Solution

## Return to Normal



Flash Memory Summit

