

VIRTUAL EVENT • MAY 24-25, 2022

Big Memory in Multi-Cloud

Presented by Charles Fan, CEO and Co-Founder, MemVerge

The Rise of Modern Data-Intensive Applications

AI/ML



Genomics

EDA Simulation



Geophysical

Video Rendering



CFD

Financial Analytics



Risk Analysis











Cloud Infrastructure Not Keeping Pace





3

CXL is a Game Changer!



Is Not Enough



The Solution: Multi-Cloud Memory-as-a-Service







MemVerge Memory Machine™

Available since Sept 2020



Memory Virtualization

- Software-defined Memory Pool
- Intelligent Auto-tiering
- Delivers big memory capacity without application change

Memory Snapshot Service

- Fully captures running application state
- An Application can be rolled back, restored or cloned from anywhere at any time
- Delivers higher application mobility and availability



Use Case Example 1 Memory Capacity Expansion for Cloud Infrastructure



Memory Machine enables:

- Higher VM density
- Lower space, power, & cooling
- Lower TCO



Арр

OS

Арр

OS

Арр

OS

Арр

OS

App

OS

Use Case Example 1 Auto-tiering with Dynamic Ratio Optimizes Performance

Sysbench QPS of MySQL in KVM



40GB Data



60,000

Use Case Example 2 Bigger Memory Calls for Better Protection







Memory Machine



Memory blast zone is bigger

In-Memory Snapshot provides better protection and faster recovery



Use Case Example 2 ZerolO Snapshot Dramatically Lowers Recovery Time

Restore a 315GB Redis Database (300M Keys)





Use Case Example 3

Lower Time-to-Results for Scientific Computing

Single Cell RNA Sequencing extensively used in COVID-related research. For example: Ren et. al.: "COVID-19 immune features revealed by a large-scale single cell transcriptome atlas", Cell, Feb 3, 2021



Workload Attributes

- Large datasets
- Multi-stage pipeline
- Requires frequent checkpoints of intermediary stage results
- Frequent Rollbacks to tune parameters
- Branching to support what-if analyses

Pain Points

- Checkpoint to disk and rollback extremely timeconsuming
- Data loss risk
- Computation memory intensive



Use Case Example 3 Shortens Execution Time by 60%+

Mouse Cell Atlas (GSE108097), 176 Samples, Matrix Size 31787 x 813348



Execution time of each analysis stage: compute + storage IO or in-memory snapshot

DRAM + Storage
DRAM + PMEM + Memory Machine



Use Case Example 3 Rolling Back Becomes Instant

Time to restore a compute stage for parameter tuning & debugging



Reload from Storage

Reload from In-Memory Snapshot







Big Memory Cloud Solves These Challenges

Memory too small





Cloud too expensive



Cloud Bursting



Cloud Lock-in



Transparent access All data fits into to lower cost PMem memory, eliminating IO to storage Stateful non-faulttolerant apps can use low-cost Spot Instances Workloads can burst to cloud during periods of peak usage

Long-running workloads can move between cloud instances



Future of Big Memory in Multi-Cloud



Long-running dataintensive apps are faulttolerant and highly mobile in Multi-Cloud

Composable Memory enables peta-scale memory lakes



Join Us in Shaping the Future of Big Memory

MemVerge Technology Partners





National Energy Research Scientific Computing Center





Memorize the Future







Please take a moment to rate this session.

Your feedback is important to us.