XENON®
High Performance Computing

Hybrid HPC

Adrian Torrie - Solution Architect

13th Feb 2025

Originally presented by:
Jakub Szarlat – Senior Solution Architect
eResearch Australia, October 2024
Melbourne, VIC, Australia

25+ Years where we come from…

2023
AGRF
Flash storage for genomics

2023
Pawsey
130PB Archive Cluster

2023
PTT Thailand
HPC Consulting

2023
WSU
Neuromorphic AI Cluster

2022
Deakin University
Clustered NVMe software defined storage for AI research

2022
WEHI
Clustered flash storage for genomics, cryo-em

2022
XENON Cloud
HPC-as-a-Service for HPC/AI, private & public cloud

2021
Deakin University
Design & deliver new HPC cluster for shared use across Deakin

2021
Hong Kong Uni HPC
HPC Cluster with XENON Cluster Stack

2021
Murdoch Children's Research Institute
Design and implement backup solution

2021
XENON Cluster Stack
Containerised HPC management solution

2020
CSL Limited
Services for HPC cluster implementation

2020
Monash University
HPC cluster & storage configured as shared cloud node for Australia Research Data Commons

2020
Todd Energy
HPC cluster for oil & gas exploration

2019
Pawsey Supercomputing Centre
New GPU cluster

2018
Harrison.ai
AI solution for IVF, radiology, healthcare

2017
NCI
Supercomputer for scientific research & technology innovation

2017
WEHI
Private cloud for next generation cancer, disease & medical research

2016
Garvan Institute of Medical Research
NVMe solution for medical research

2015
Thales Defence
High fidelity solution for Australian Army Tiger helicopter simulators

2014
RCC
FlashLite HPC cluster

2012
Fujitsu
Infiniband for Raijin Supercomputer

2009
CSIRO
Australia's first GPU Cluster

2007
Victorian Partnership for Advance Computing
HPC Cluster for Advance Computing

2005
Thales Defence
Image Generators for ASLAV simulators

2000
Animal Logic
Render Farm for "Matrix" movie trilogy

1998
Telstra Digital Video Network
Video Server equipment

2020
2010
2000

Read the Case Studies

XENON – Storage Partners

The following are some of our storage partners

Introduction

Hybrid Compute distributed globally is easy, but getting large data to the compute is not.

The Holy Grail is for research data to be **where** it is needed, **when** it is needed.

There are some third party software that can be used to move data but this can make the environment more complex and costly as it requires yet another piece of the puzzle. Thus, we will look at some of the approaches our storage partners are taking to try to get closer to the main goal.

So, are we there yet?

Simply Stream Data

Overview:
- Stream from “publicly” accessible sources
- You can create local read caches to boost re-read speed

Issues:
- Data write back becomes a challenge especially when running a hybrid workflows

Use Cases:
- Reference data accessible externally (can be via VPN, etc)
- Sequential workflows (At least to the location if caches are used)
- No need to write large amount of data back to original repository
- Source can accept data easily via other connection

Examples:
- Versity S3 Gateway – Bridge to file-based storage
- AIStore – Supports object stores as backends with local node caching

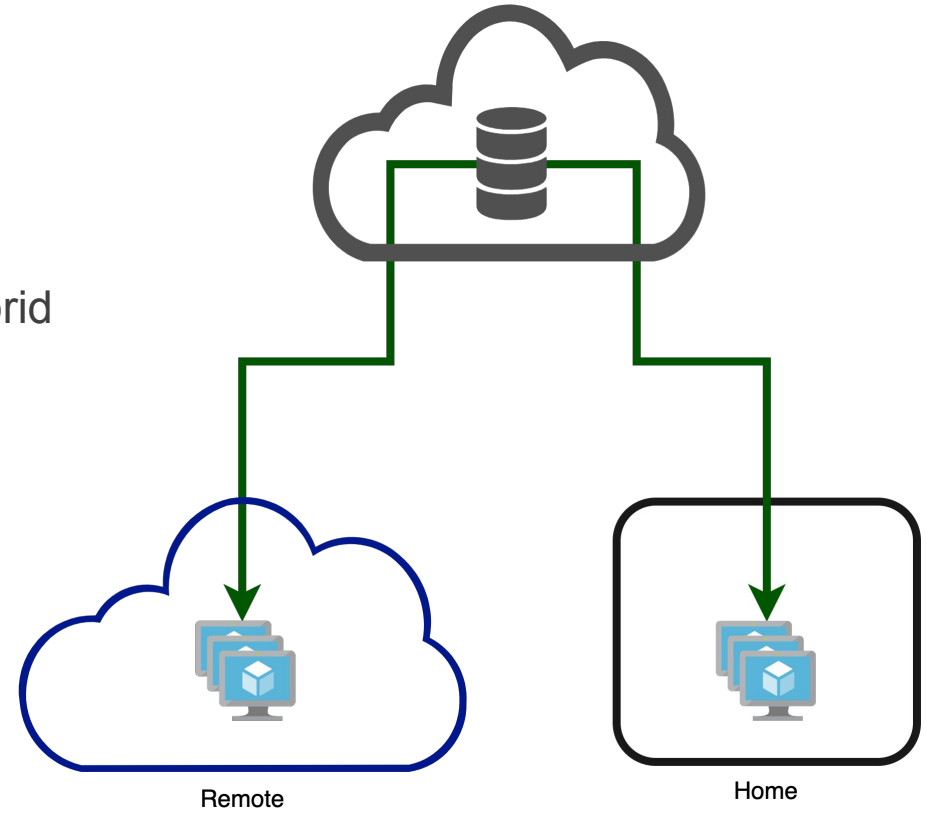

Central Repository - Check-in/Check-out

Overview:
- Make “Cloud” the central repository
- You can create local read caches to boost re-read speed
- Writes are handled to local disks allowing for performant writes
- Checked-out data does not require connectivity back to source till check-in

Issues:
- Need to manage conflicts on check-in if multiple sites update the same data
- Extra disk space required for data for local cache
- Check-out to central cache requires extra resources to provide storage service to clients

Use Cases:
- Workflows where sites do not work concurrently
- Data is maintained in a central repository already

Examples:
- Quantum Flex Sync (A part of StorNext and Myriad)

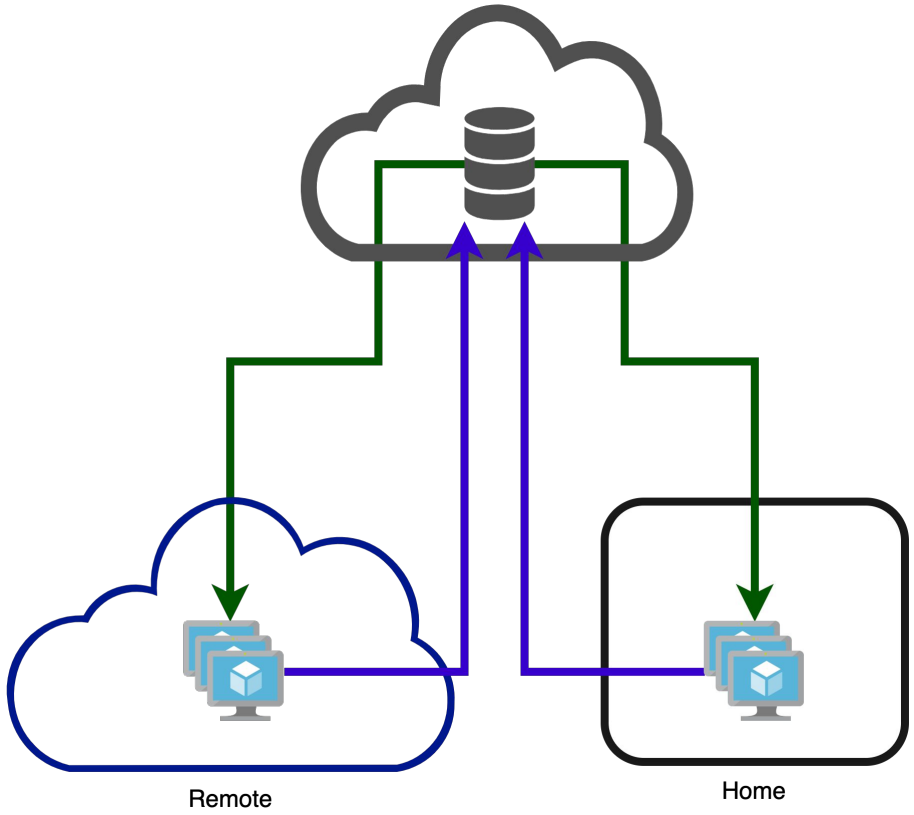

Snap Replication

Overview:
- A snapshot of the area is copied up to remote.
- Future updates only require delta updates.
- Snaps presented are read-writable.
- No need for central repository

Issues:
- Initial snap replication can take a lot of time to a new remote site
- Result data needs to be managed outside of the system
- For repeat workloads its best to keep data at remote which can add $$$
- Implementations tend to require a cluster available on remote to present the snap

Use Cases:
- Writable Clones are great for use cases when you want a point-in-time copy of data, but you don't want the updates to propagate back to the global namespace.
- Follow the sun where snaps are used to rebuild each sequential site as only deltas need to be forwarded on.

Examples:
- WEKA

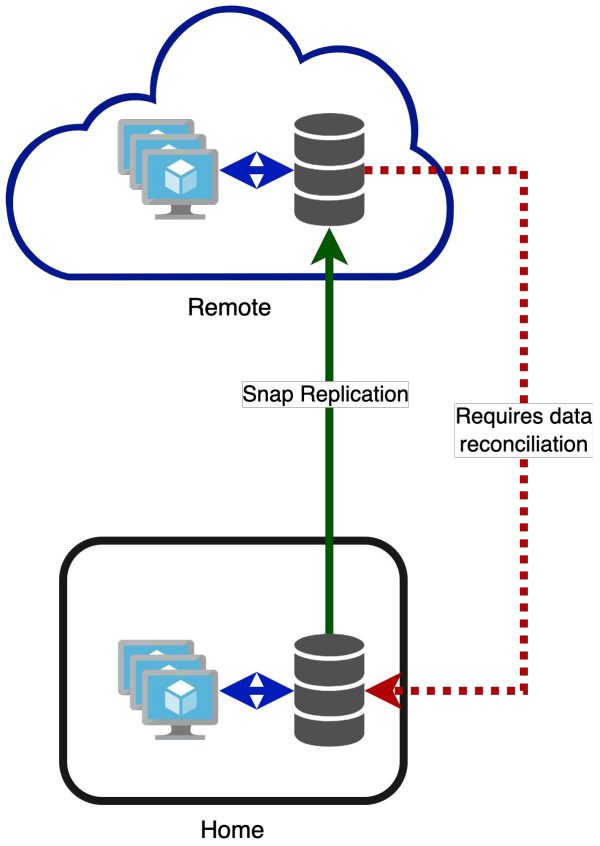

Decentralized Global Namespace

Overview:

- Provides the same view of data across all sites at all times.
- Locking can be used to ensure data is always consistent
- Leasing can be used to speed up read by ensuring data is gravitated to a working site
- Path based synchronization ensures only data required is moved limiting requirement on how much space is required at the remote sites.

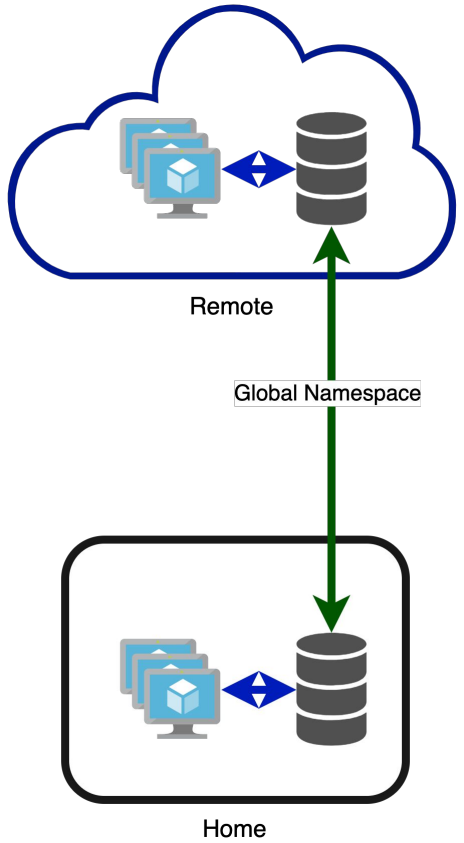

Issues:

- Lower data protection as data may become incomplete if remote site goes down and data is located at that site
- Writes at remote site are not ultra-high performant
- Its not synchronous yet

Use Cases:

- When direct writes back to the namespace don't need to be super performant.
- Consistent view of the namespace across all sites is required
- No modification of data handling workflow is possible
- Using a local scratch space and then using the namespace as a repository

Examples:

- VAST and Hammerspace

Thank You

XENON
High Performance Computing

XENON Systems Pty Ltd
10 Westall Road, Springvale, Victoria 3171, Australia

www.xenon.com.au

P +61 3 9549 1111
F +61 3 9549 1199
E info@xenon.com.au

A member of the XENON Technology Group
www.xtg.com.au