



**Hewlett Packard
Enterprise**



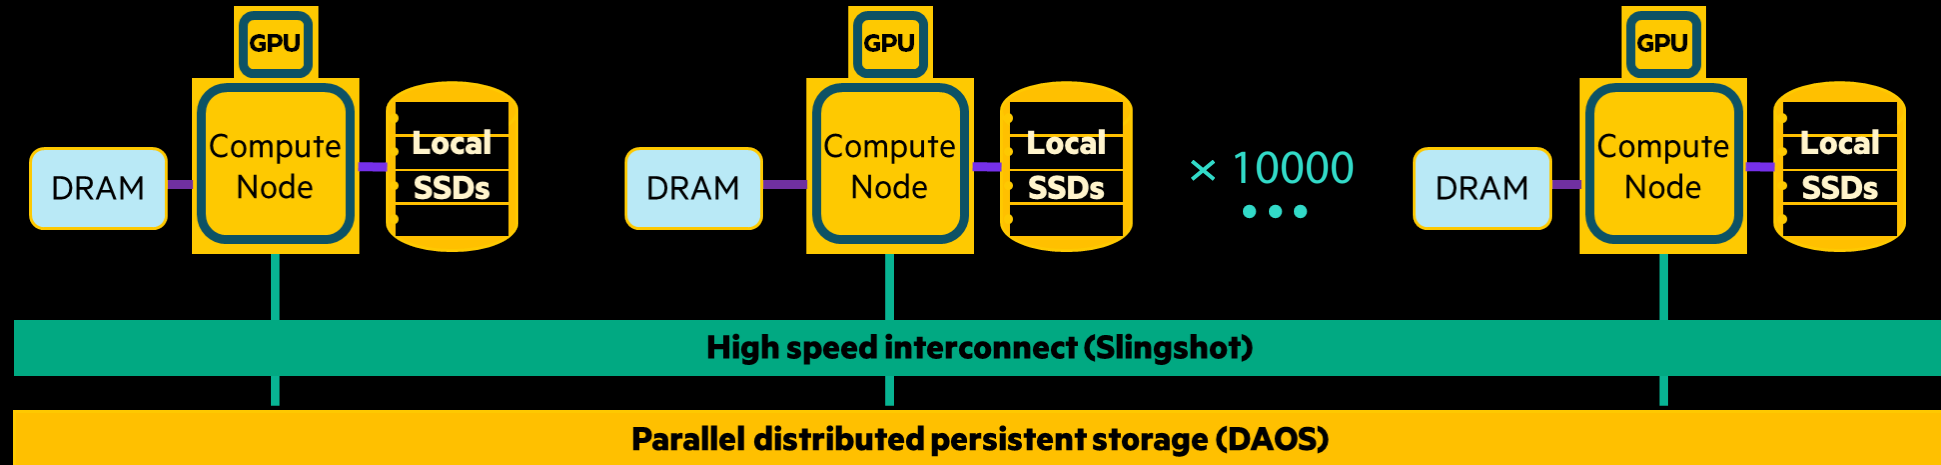
A Global In-Memory Cache and Computation Tier for DAOS

Supercomputing 2024

Challenge: Solve more large problems per day on your HPC system

You decide to:

- Reduce the time needed to solve large problems by adding a cache
- Increase concurrent usage of the system by sharing resources between users

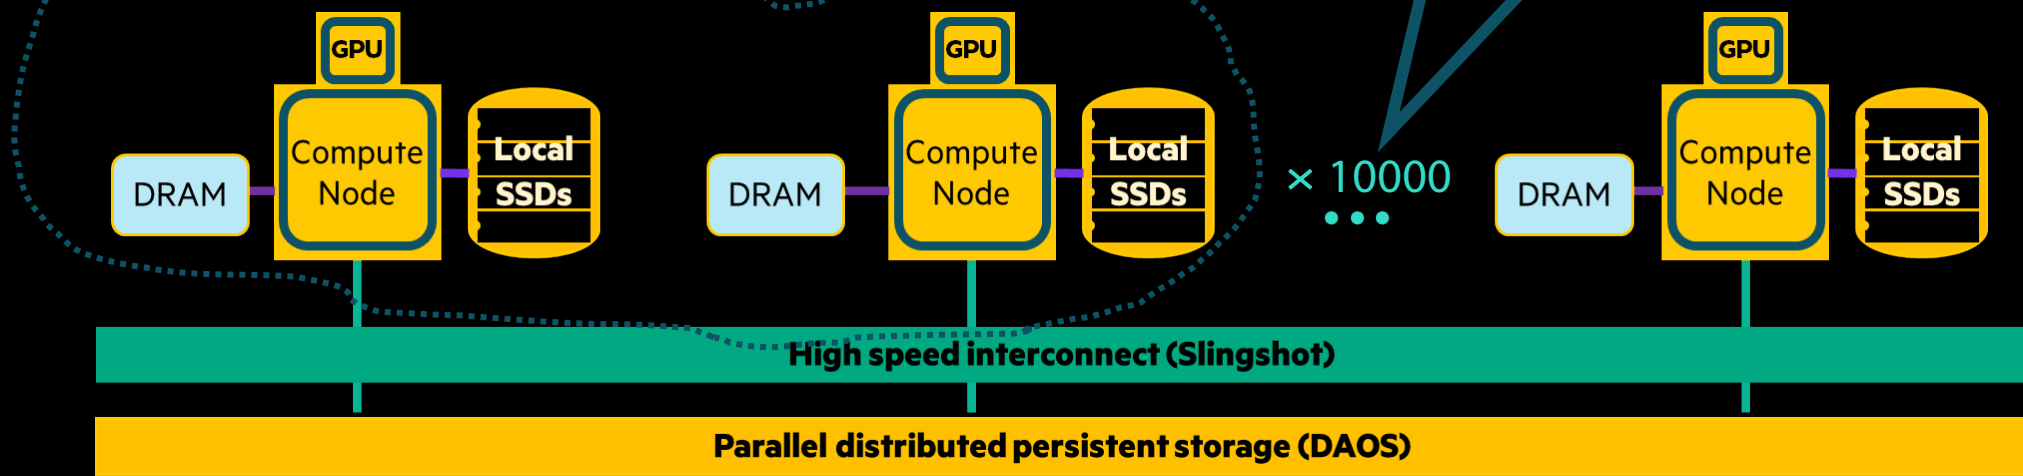


However, your system has thousands of nodes and long-running workloads

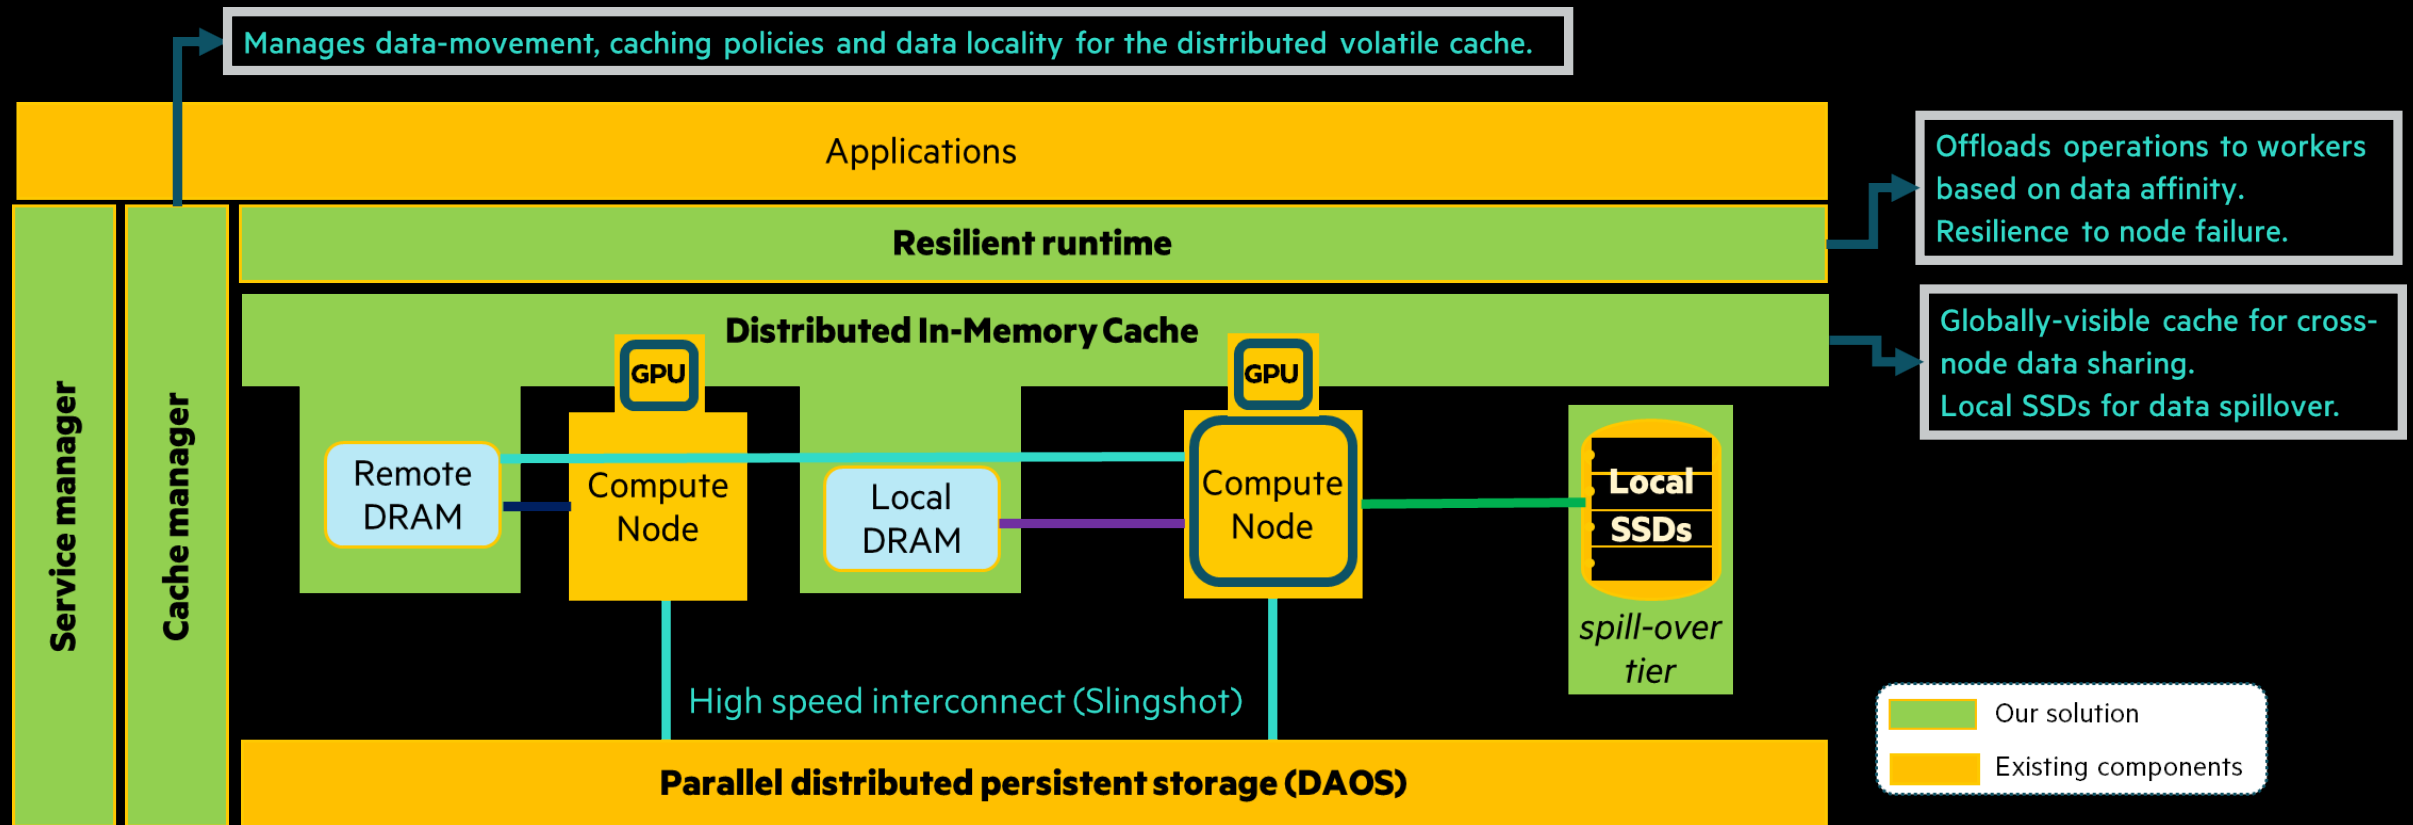
- Long-running, interactive workloads:
- Resource demands may vary over time
 - Vulnerable to node failure

Fixed resource allocation can lead to resource stranding and/or under-provisioning


- Thousands of nodes:
- Locality challenging
 - Increased chance of node failure



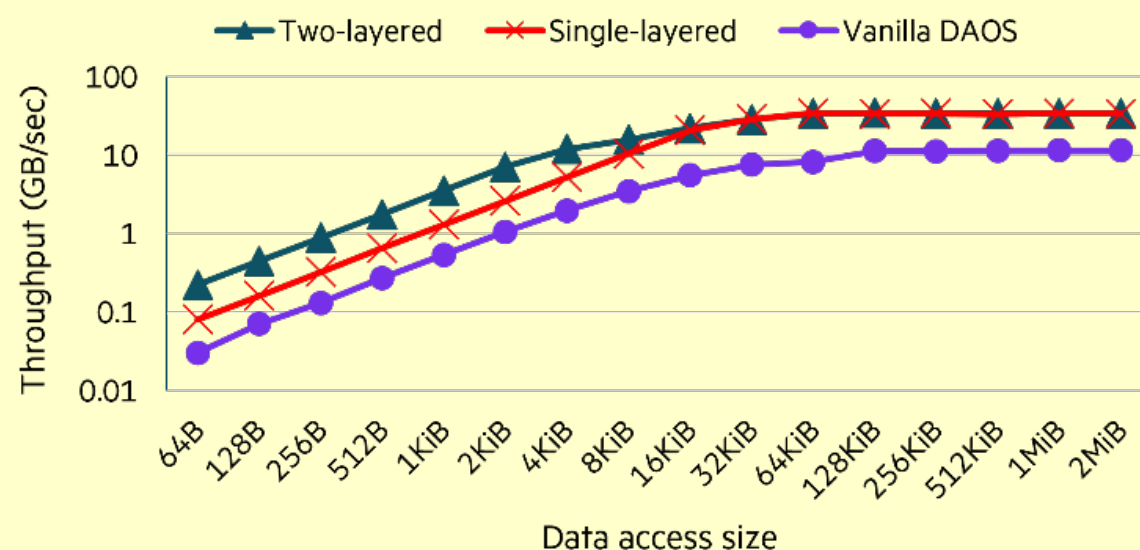
Performant data and computation tier to scale up dynamic mixed workloads without stranding resources



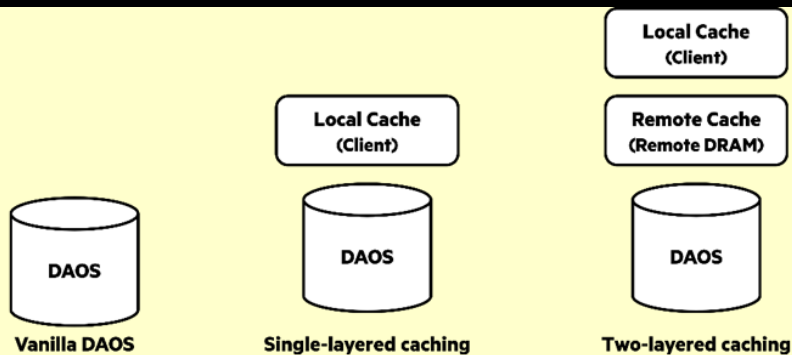
Potential impact of global shared cache



(a) Y-axis: absolute numbers



(b) Y-axis: log-scale numbers



Experimental Set-Up:

- Local cache capacity to be 90% of the working set
- Remote cache emulated using 1 OpenFAM server (two-socket AMD EPYC 7763 64-Core Processor, 4 TiB DRAM)
- 1 DAOS client node (two-socket AMD EPYC 7763 64-Core Processor, 1 TiB DRAM).
- 1 DAOS storage server node (two-socket Intel(R) Xeon(R) Platinum 8380 CPU, 4 TiB DRAM, 3 Kioxia CD6 6.4 TB NVMe SSDs.
- Slingshot interconnect (single 200 Gb NIC per node).

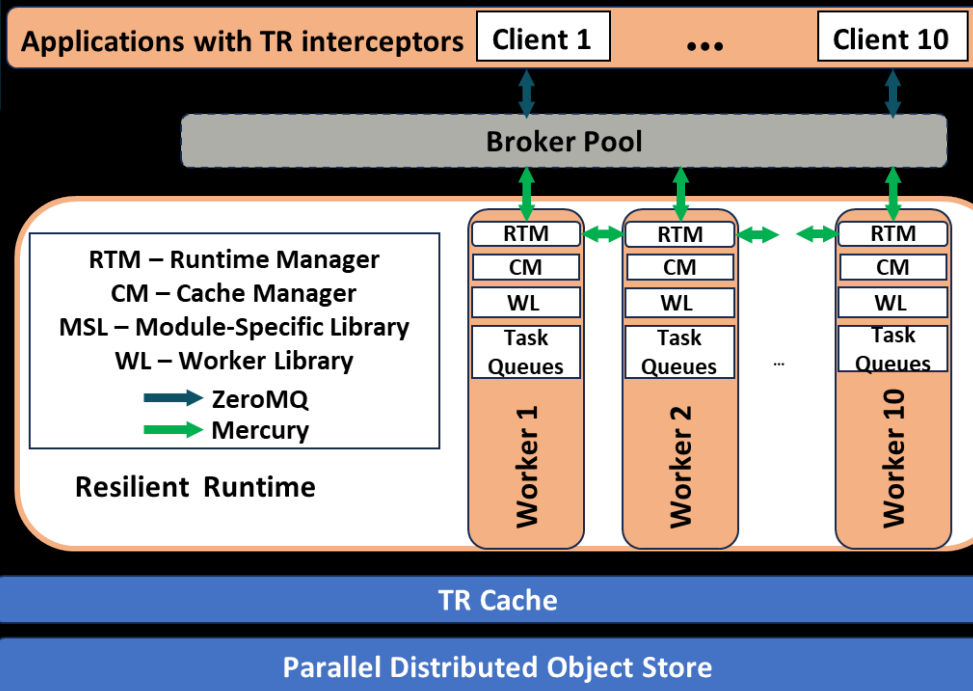
Resilient Runtime for Offloading

Python program loads interceptor before importing target Python module.

```
Python Application  
(TitaniumRattlesnake interceptor loaded)
```

```
%load_ext tr_interceptor  
import pandas  
...
```


Python interceptor sends an RPC through a Broker to a Runtime Manager



Runtime Manager loads appropriate module-specific library, identifies locality opportunities, then sends tasks to workers.

As workers complete tasks, they log task status so that work can be recovered in case of failure

Questions?



Outline

TitaniumRattlesnake augments storage with a performant data and computation tier.

Challenge Solve more large problems per day on your HPC system

Global shared distributed cache mitigates locality challenge

Resilient runtime for offloading improves locality and sharing

Resilient runtime also protects against node failure