

MAXIMIZING DATA'S POTENTIAL

Storage at the Edge

Simplified Persistent Storage to
Provide QoS and Resiliency




Open Edge Computing Initiative Fall Online Workshop
Tom Prohovsky Dec 2020



ATTRIBUTES OF EDGE STORAGE

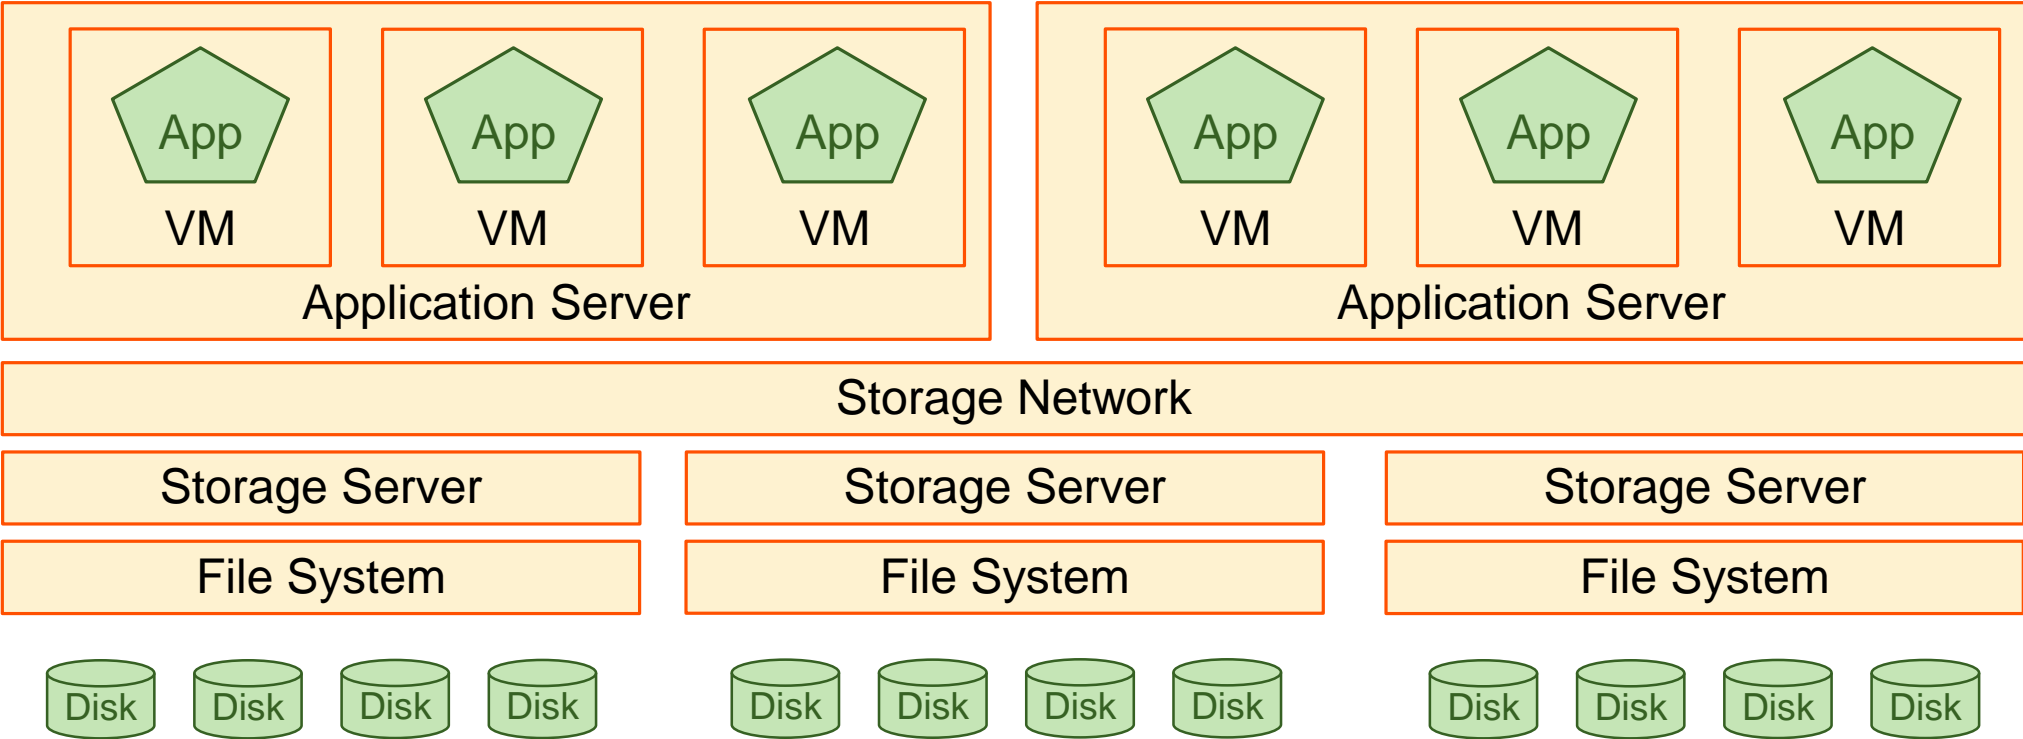
1. The edge is physically close to a wealth of data
2. Often the abundance of edge data is beyond the practicality of capture and movement
3. Some data's value erodes over time and not worth the cost of movement
4. At times valued data is intermixed in a large dataset and which data is valued is unknown at capture
5. New trends in edge application design use storage differently than traditional deployments



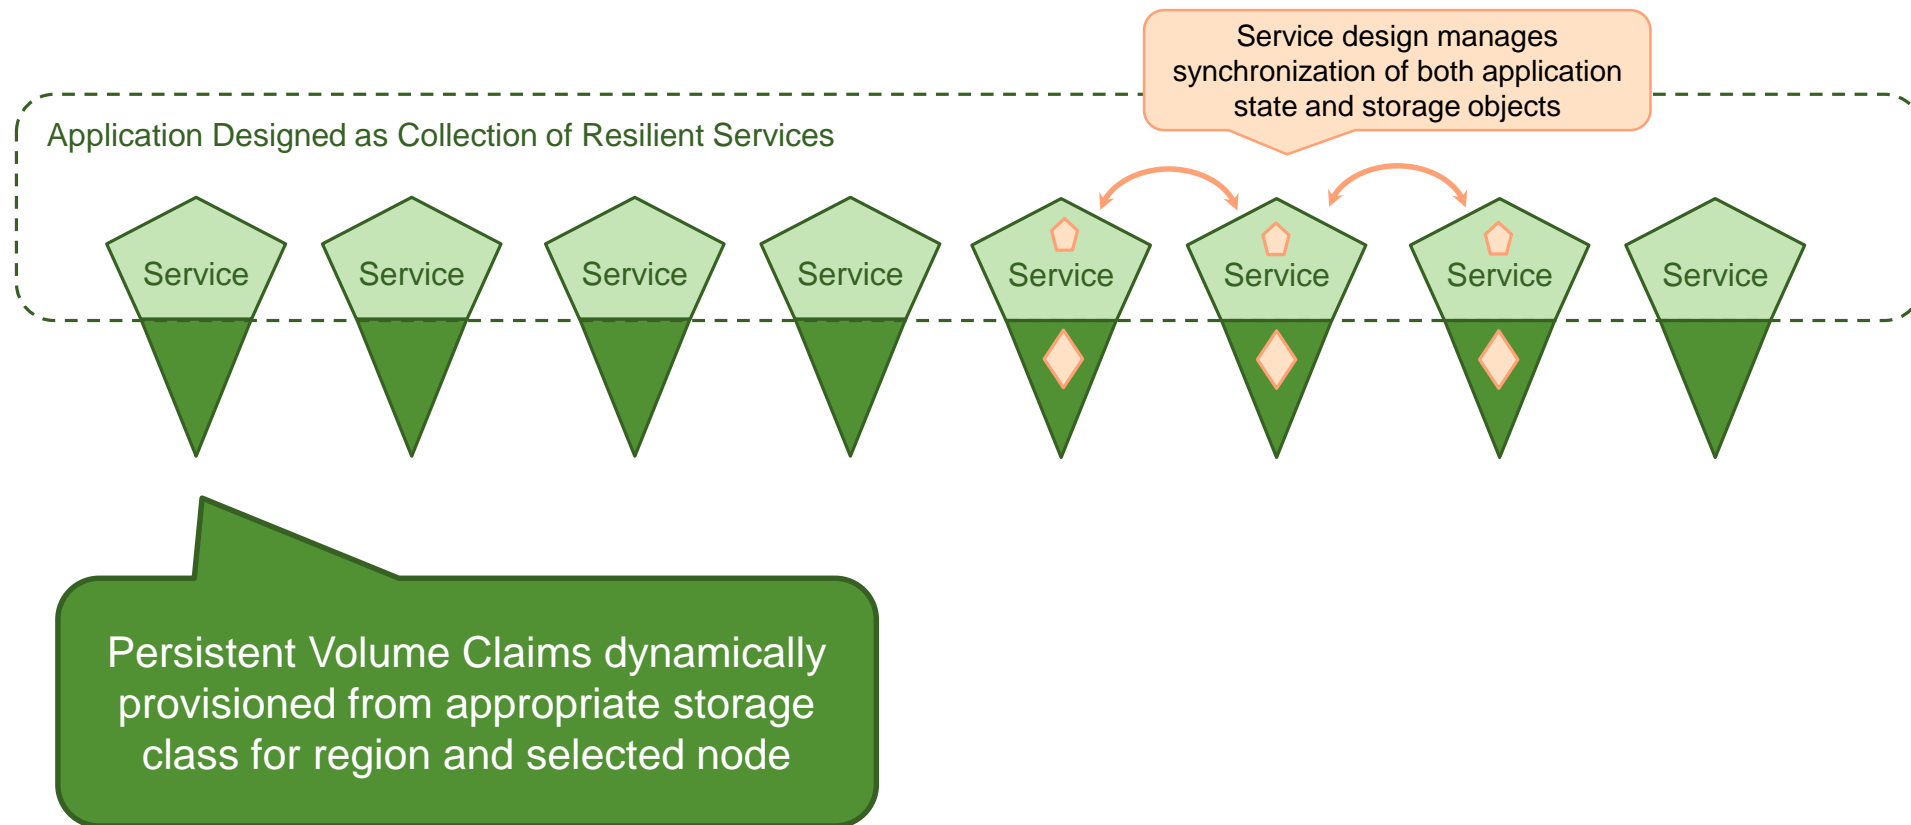
<https://buy.garmin.com/en-US/US/p/541225>



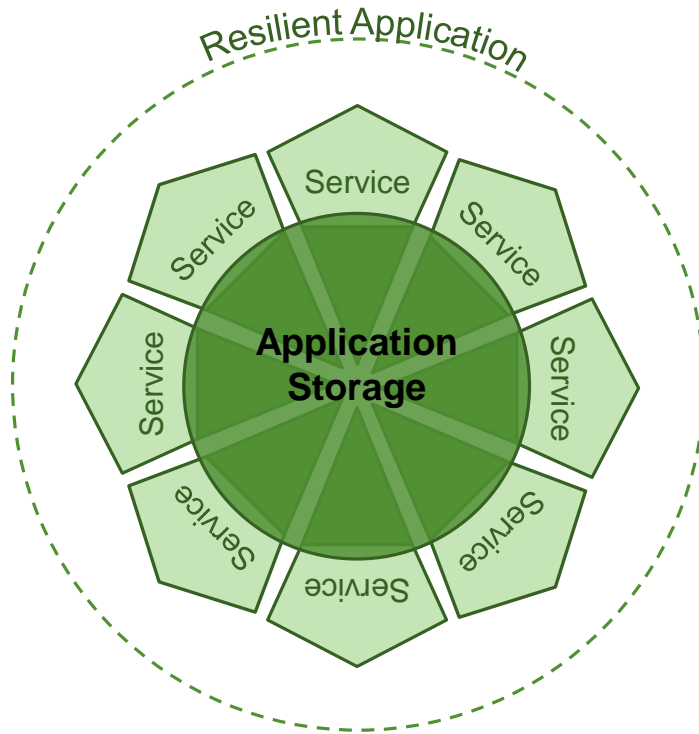
Typical Storage Architecture with Shared File System



Deployment Resilience with Sharded Services



Deployment Resilience with Sharded Services



$$\text{Application Storage} = \sum_{1}^{\text{Shards}} \text{Volume Claims}$$

No longer a constraint for storage claims:

1. Global name space
2. Interaction between peers
3. Performance beyond server's capabilities

The Relaxation of constraints suggests opportunity for TCO reductions

TCO Optimization Through Dynamic Storage Provisioning

Deployment Choices

- Storage service vendor
- Underlying storage device and technology
- Disaster recovery objectives


Customer Experience

- Variation under burst activity
- Typical and worse case latency
- Visibility to show actual workload
- Pay for good enough


CSI Based storage enables Kubernetes to dynamically provision persistent volumes from vendor

Kubernetes Storage Classes offer a layer of abstraction to meet service level objectives

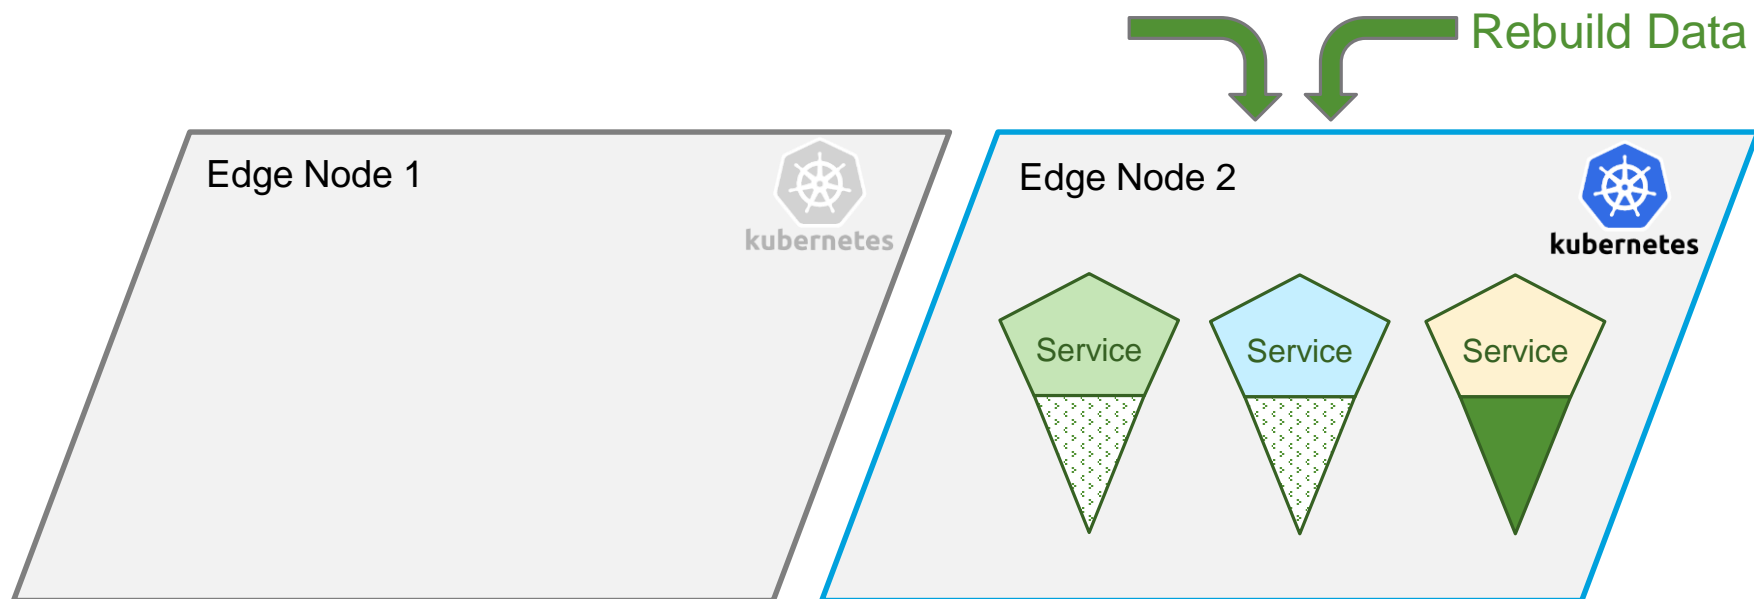
Deployment Resilience with Sharded Services




Deployment Resilience with Sharded Services – Node Failure



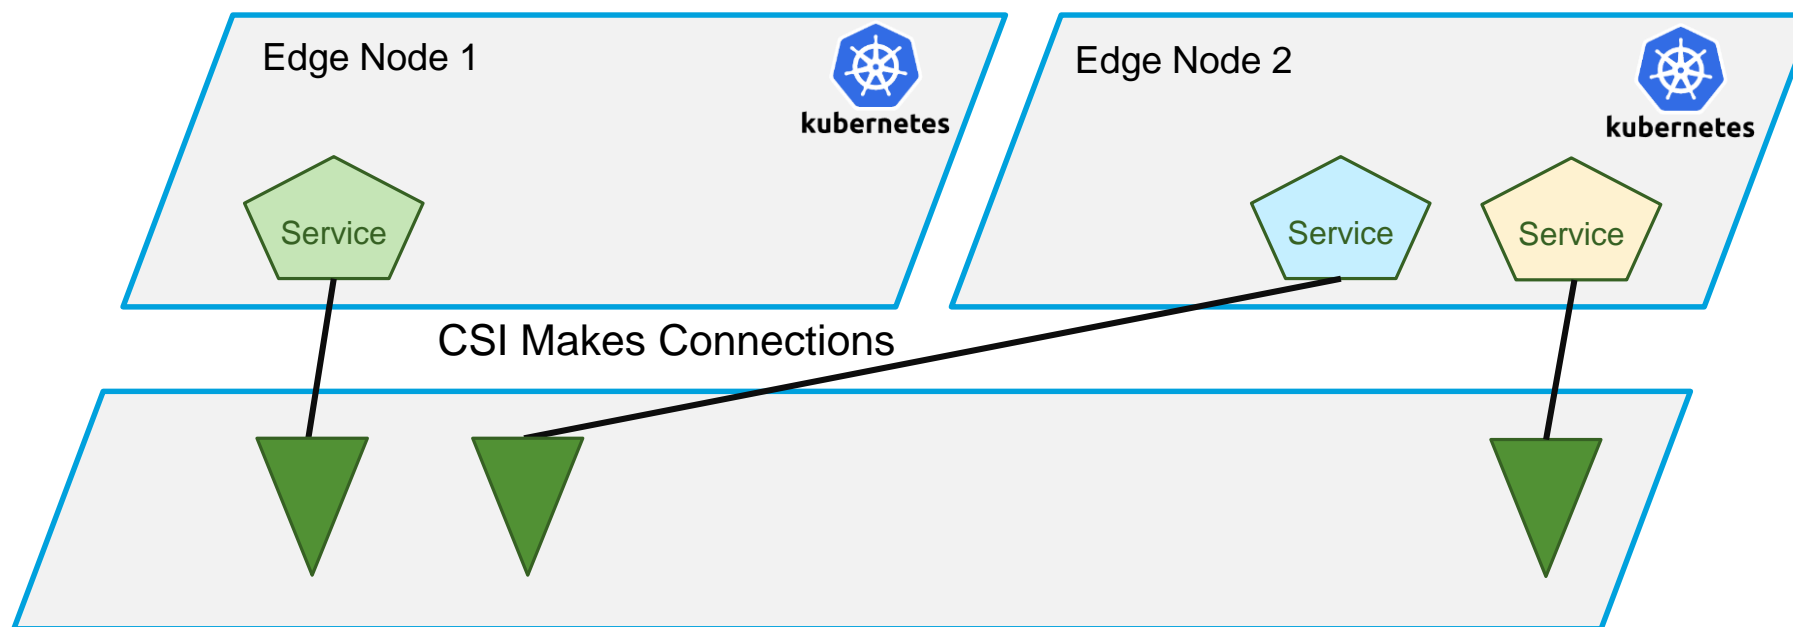
Deployment Resilience with Sharded Services – Application Rebuild



Deployment Resilience with Sharded Services – Shared Storage



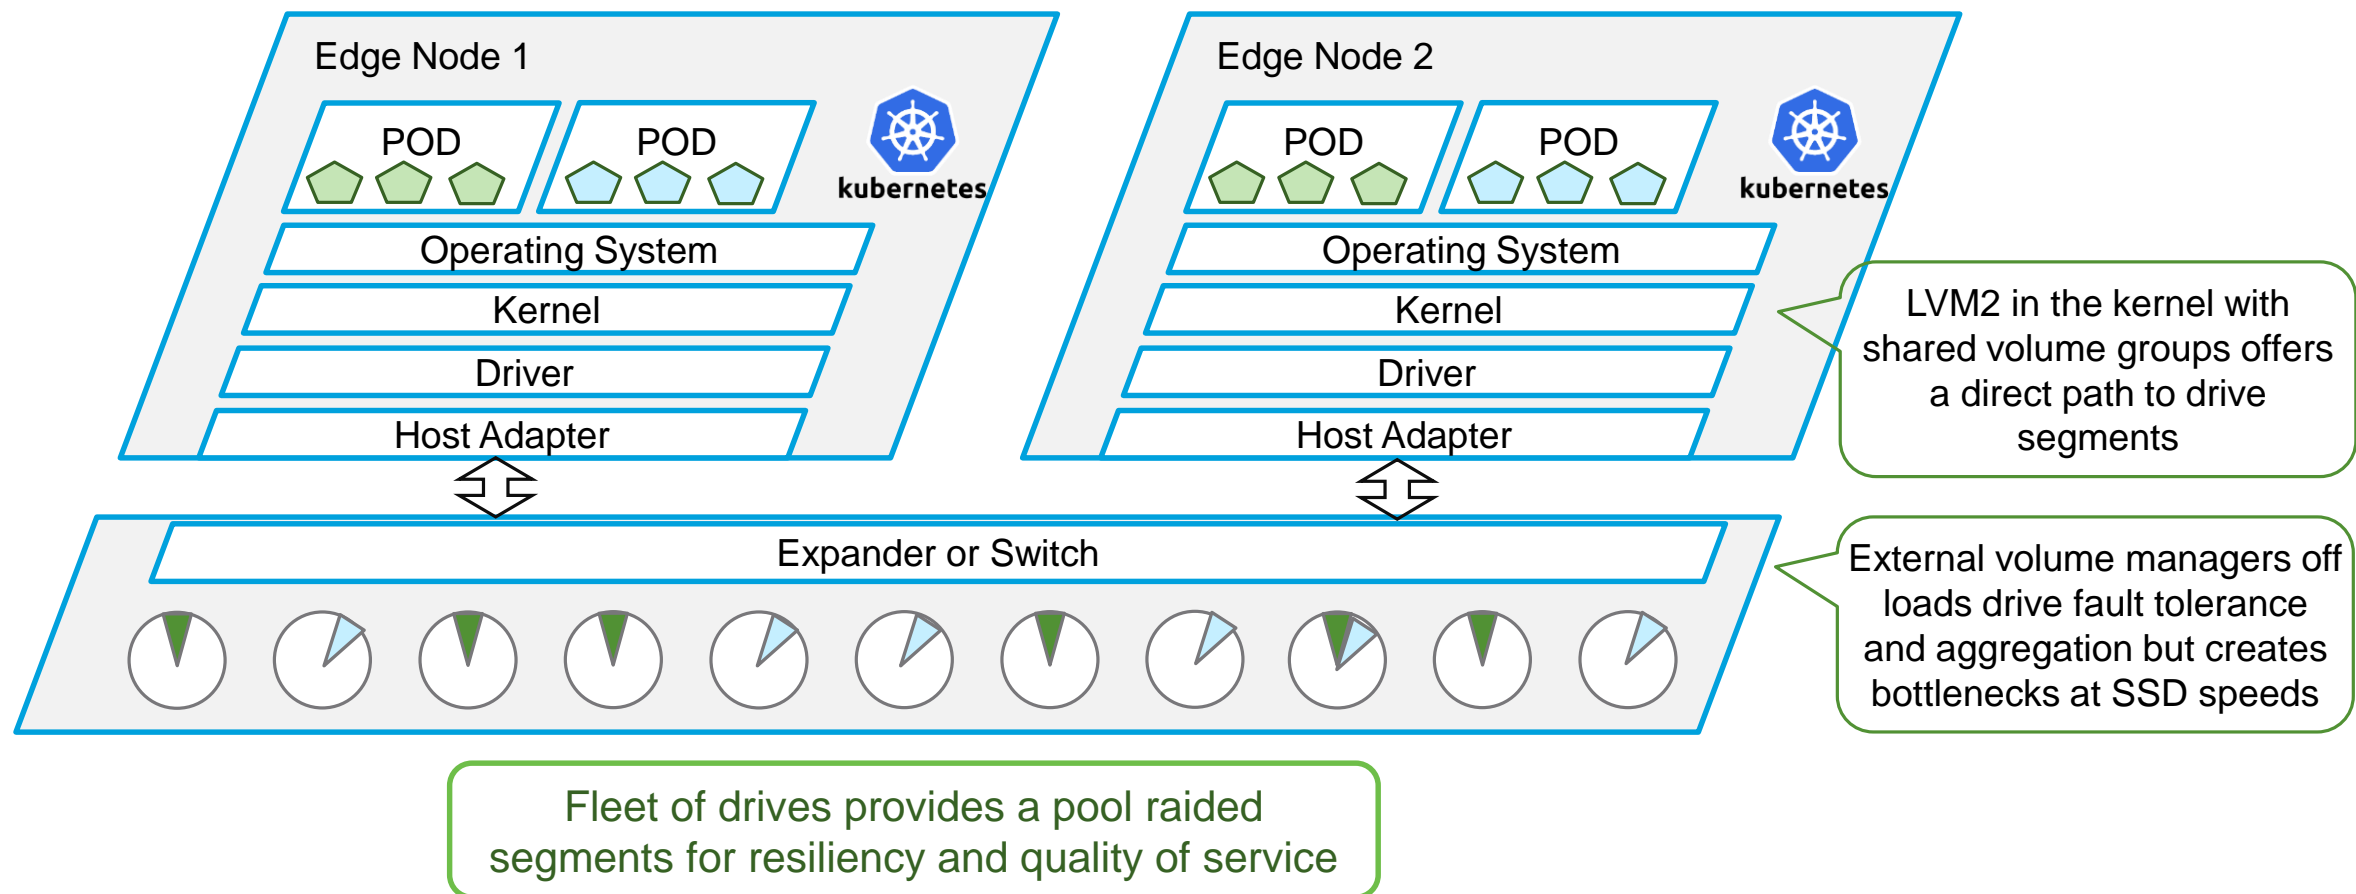
Deployment Resilience with Sharded Services – Resource Balancing



External multi-host connected storage provides the agility for Kubernetes to dynamically optimize resource utilization and QoS

Drill in to Storage Architecture Layers

Minimal layers between container and drives



Visibility from Edge Server to Drives



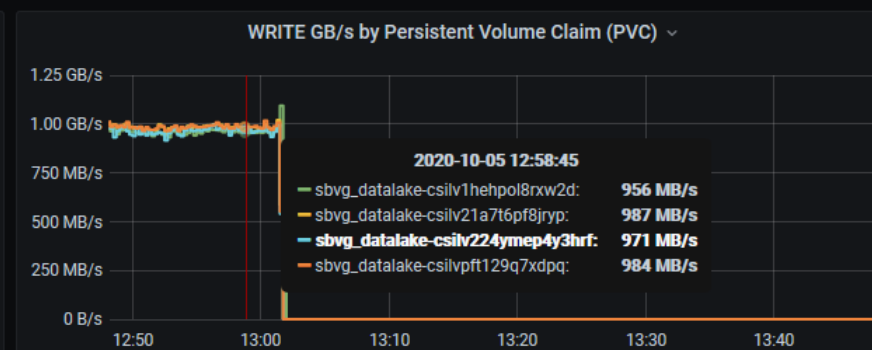



Throughput Example

Server

Application

RAID

Drives



Continued Research

- Multi-Tenant Security and Audits
- De-interleaving of datasets based on persistent volumes for SSD performance and endurance
- Improved drive fleet dynamic utilization reporting

Continue the Discussion

Tom.R.Prohofskey@Seagate.com

