



Accelerate high performance workloads within AWS

Jordan Dolman

Senior Product Manager
AWS

Customers are migrating a wide variety of high-performance workloads to AWS

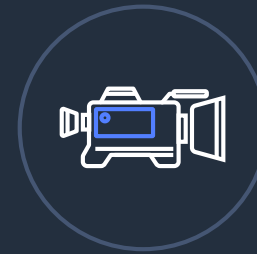
Examples by industry



Financial modeling and analytics



Genomic analysis



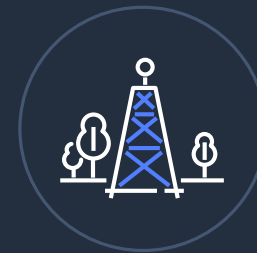
Media rendering and transcoding



Automotive ECU simulations



Electronic design automation



Seismic data processing

Examples by application area

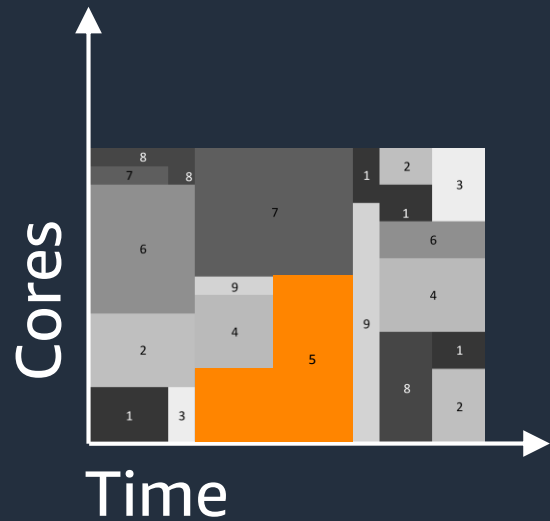


Machine learning
Model training

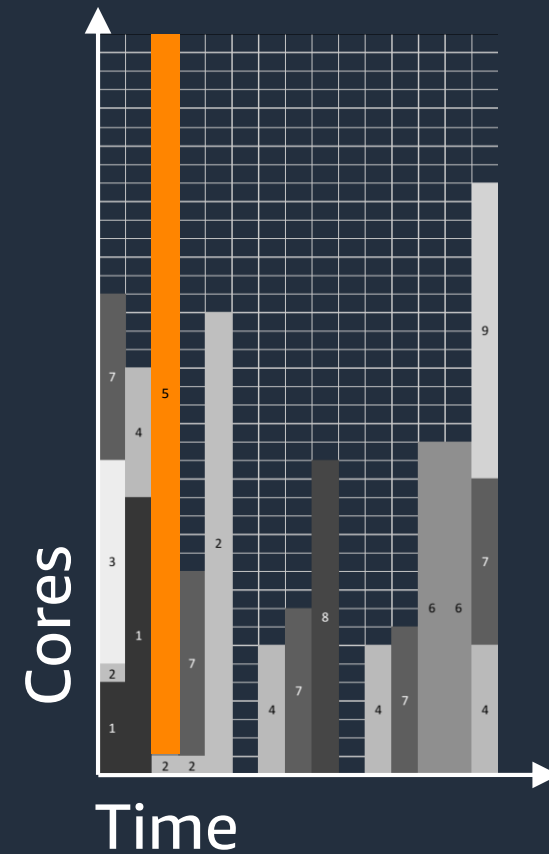


Computational fluid dynamics (CFD)

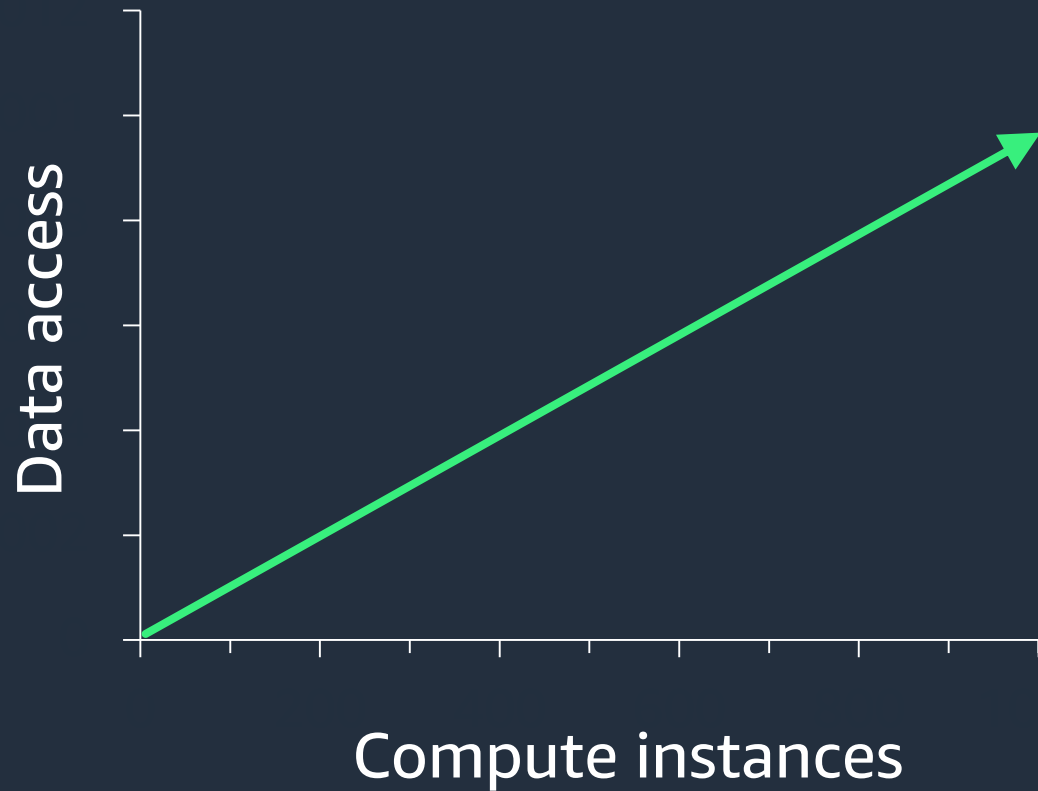
Finite compute capacity on-premises



Virtually unlimited compute capacity on cloud



Scalable compute performance requires fast, scalable data access



- ✓ High **throughput capacity**
- ✓ High **IO per second**
- ✓ Low **latencies**

Fully managed shared storage built on the world's most popular high-performance file system

Amazon FSx for Lustre



Sub-ms latencies,
100s of GB/s of throughput,
millions of IOPS



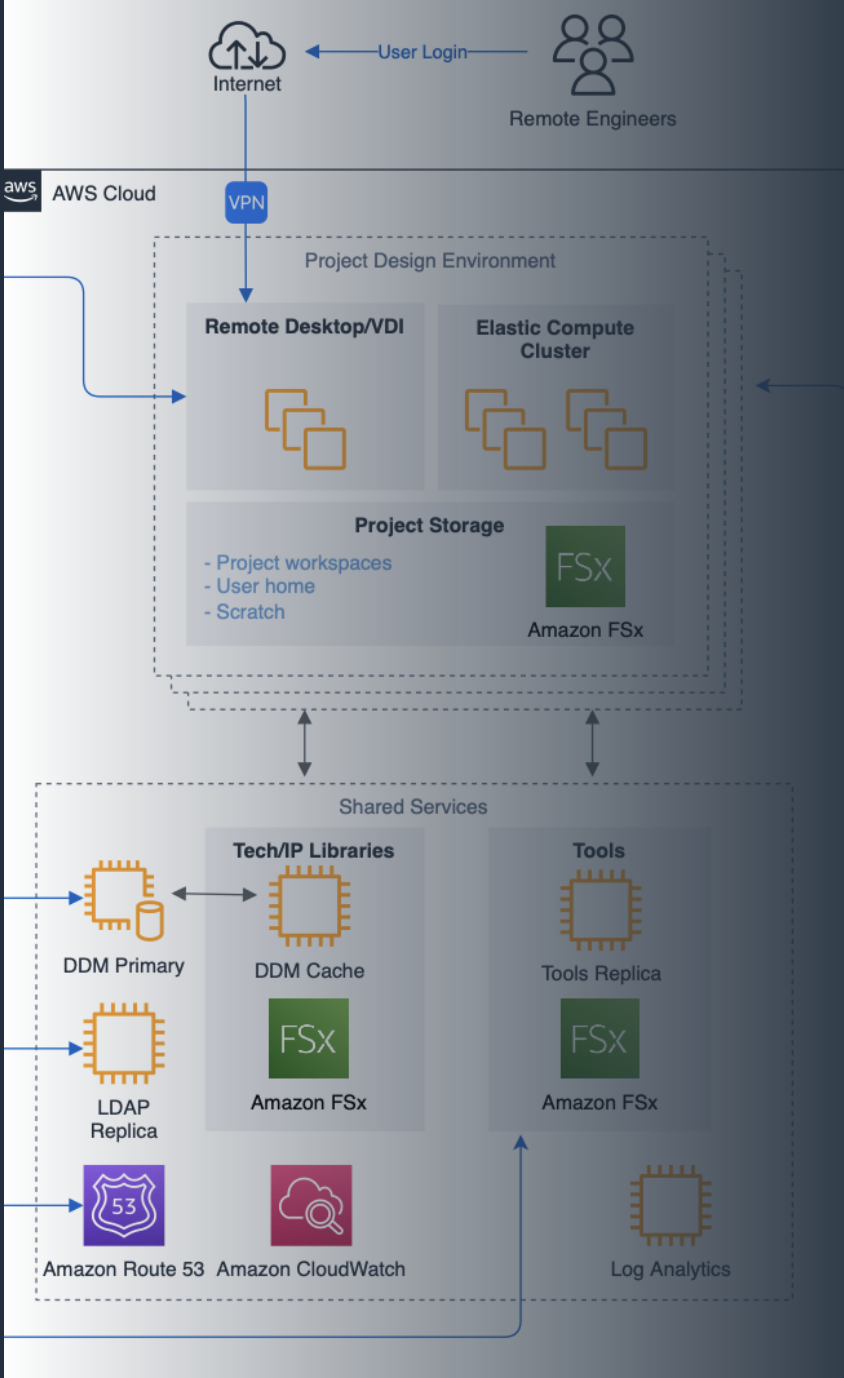
Concurrent access for
1000s of instances and
100,000s of cores



Cost optimized file systems
with HDD and SSD storage
options



Flexible deployment
options for short and
longer-term workloads



“We believe cloud-based EDA is critical to accelerating semiconductor innovation and getting new designs to market faster to power an increasingly digital world where more and more devices and infrastructure are connected.”

Olli Hyyppa, CIO and senior vice president
NXP Semiconductors N.V.

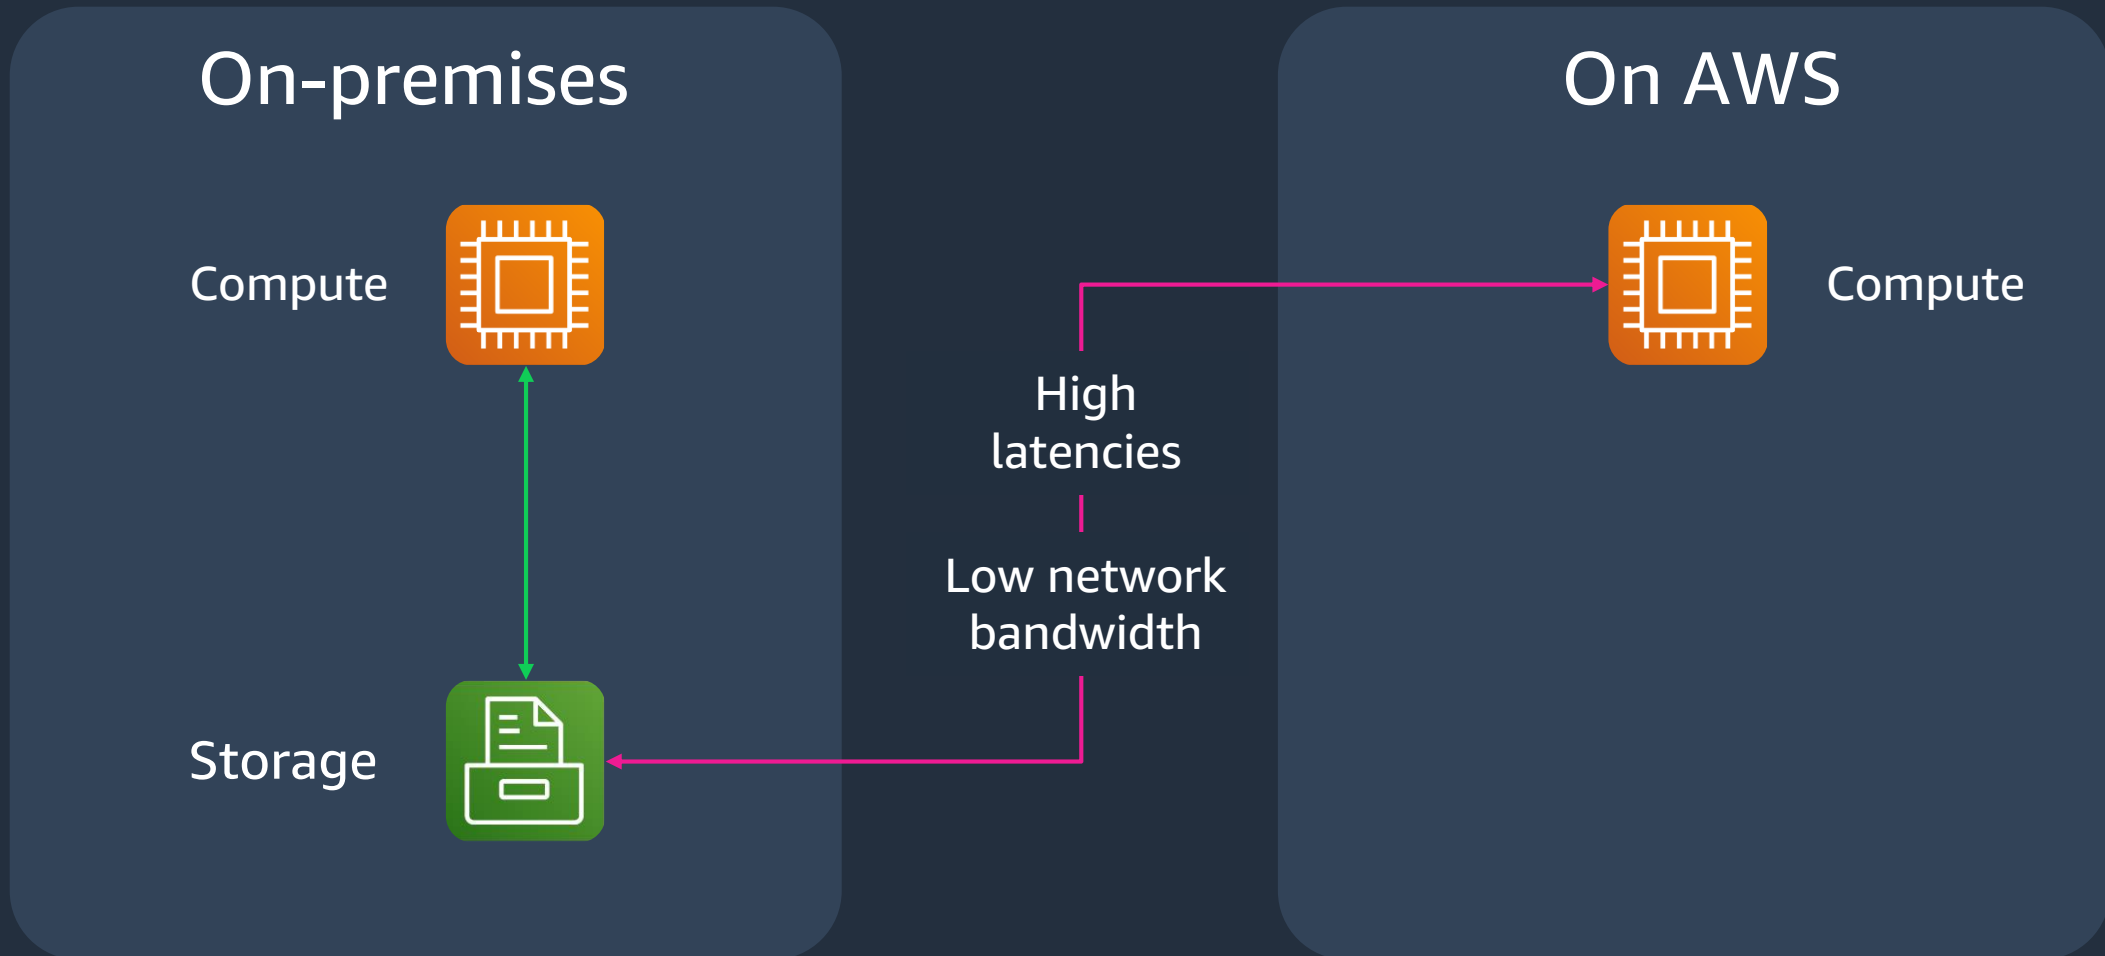


THE WATER INSTITUTE OF THE GULF®

“Running this full suite of simulations might have taken approximately **three months at a traditional supercomputing center**... With AWS, the team completed all 645 simulations in **three and a half days.**”

**Zach Cobell, Research Engineer
Water Institute**

Hybrid architectures are impacted by high latencies and low network bandwidth connections between compute and storage



Amazon File Cache:

Fast data access regardless of where you store your data



Put your data to work

Amazon File Cache



COMING SOON



Fast, highly scalable
performance



Seamless **data access**
regardless of where you
store your data sets

What Amazon File Cache brings to your workloads



Fast, highly
scalable
performance

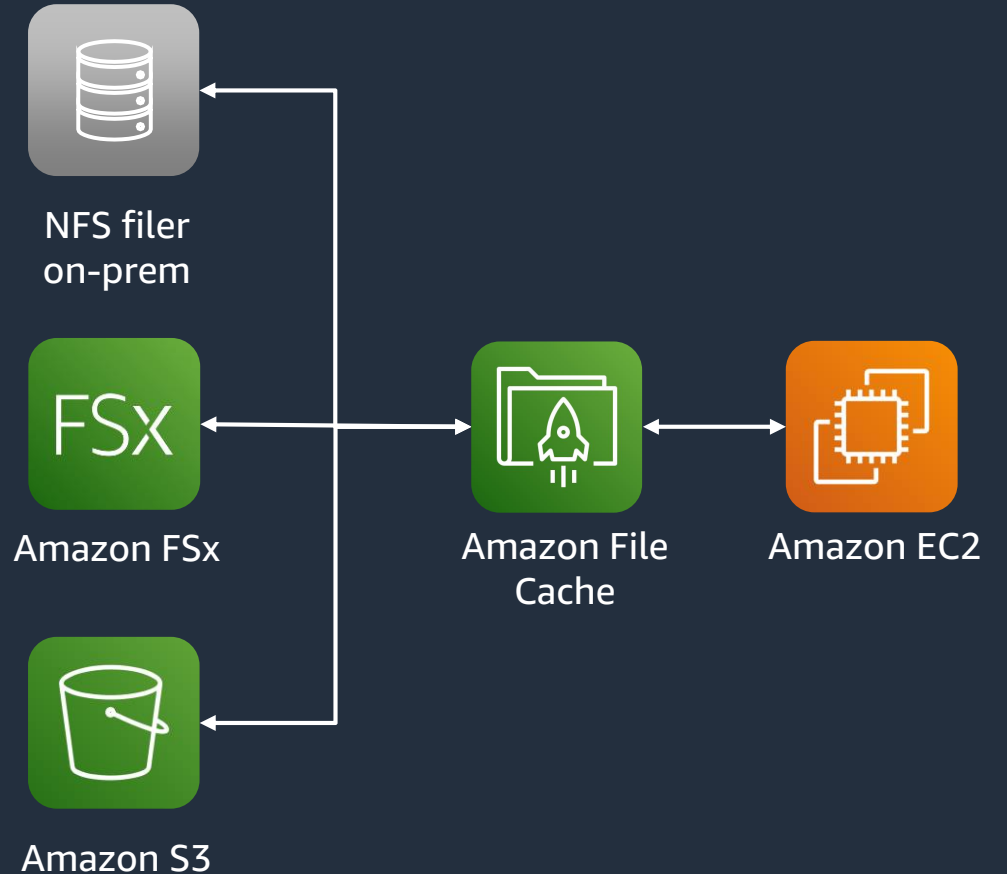
100s of GB/s
of throughput

Millions
of IOPS

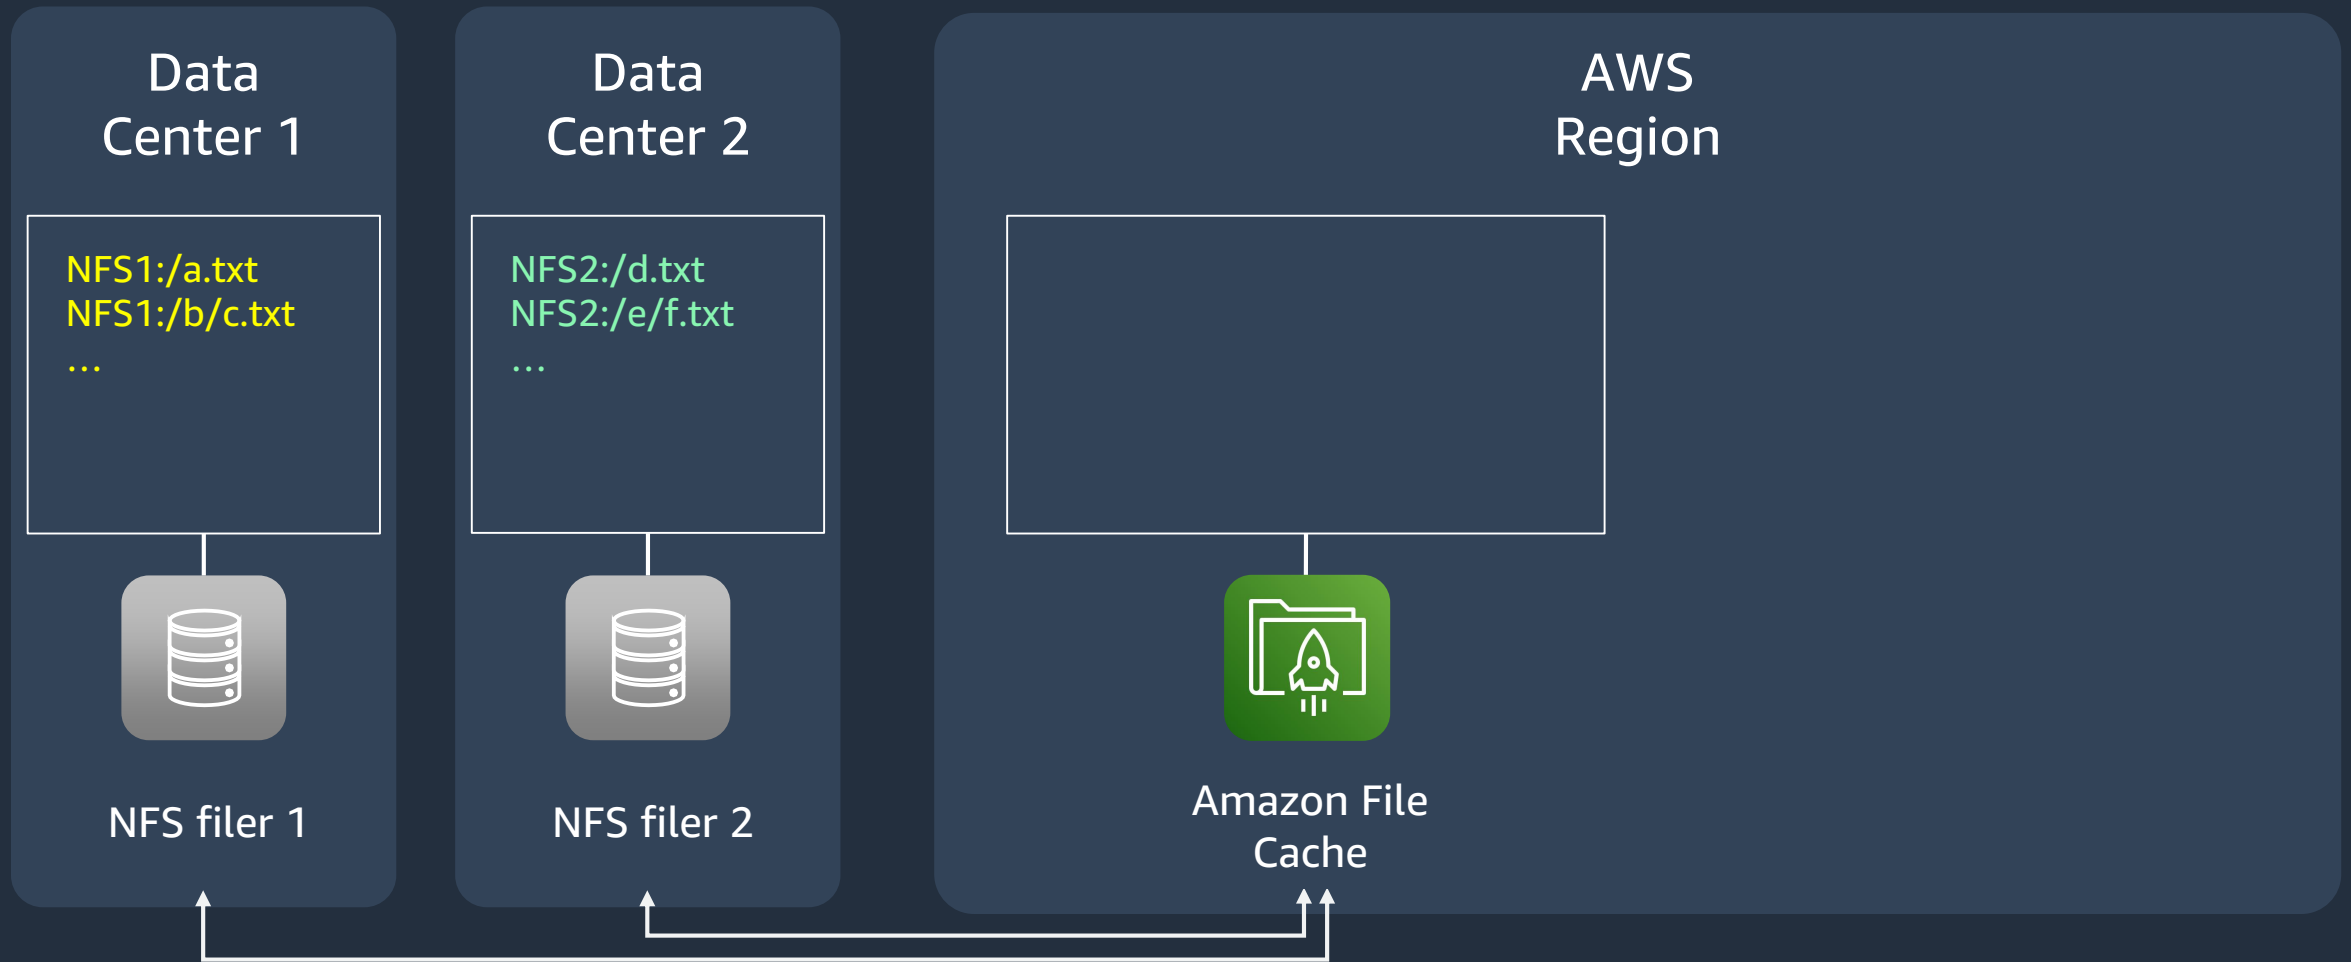
What Amazon File Cache brings to your workloads



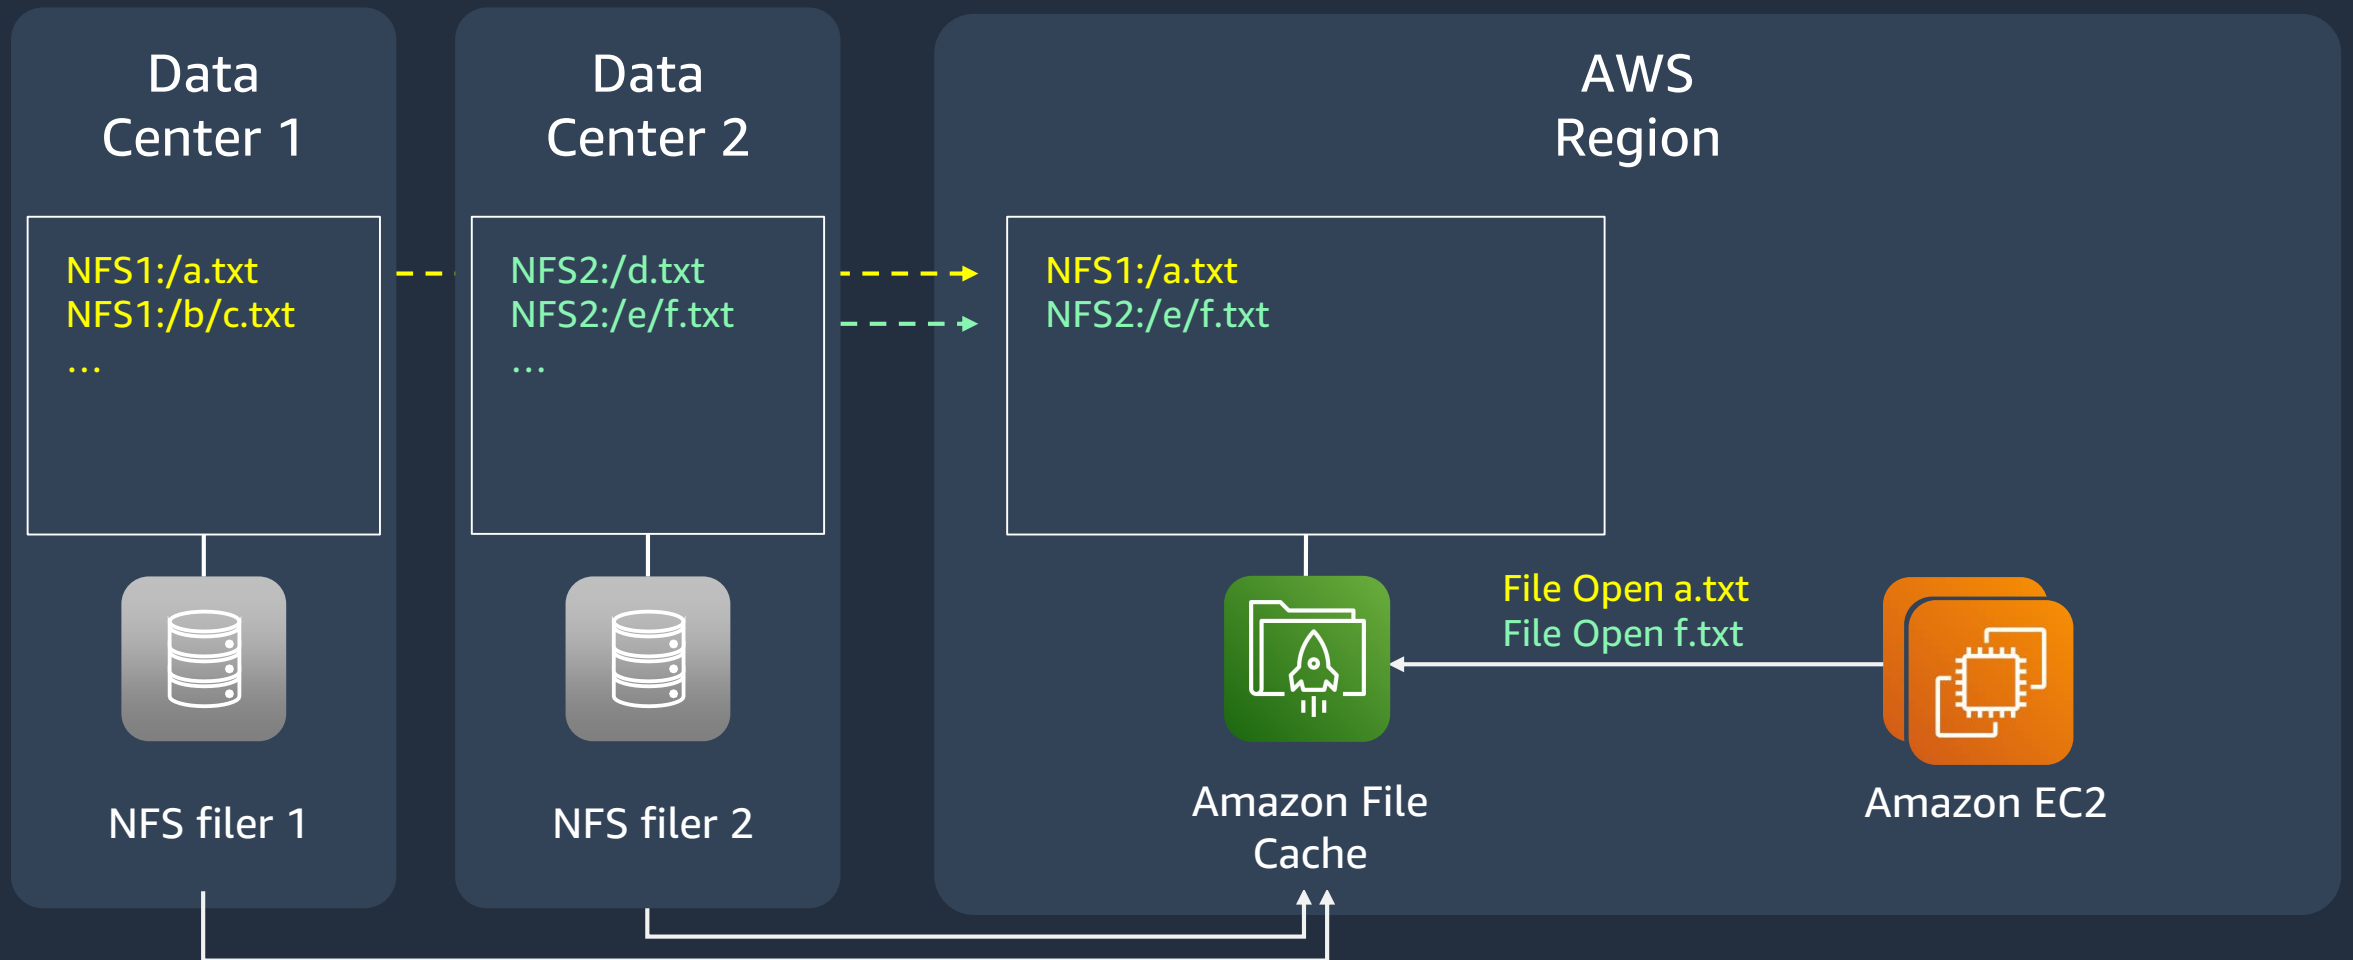
Seamless **data access**
regardless of where you
store your data sets



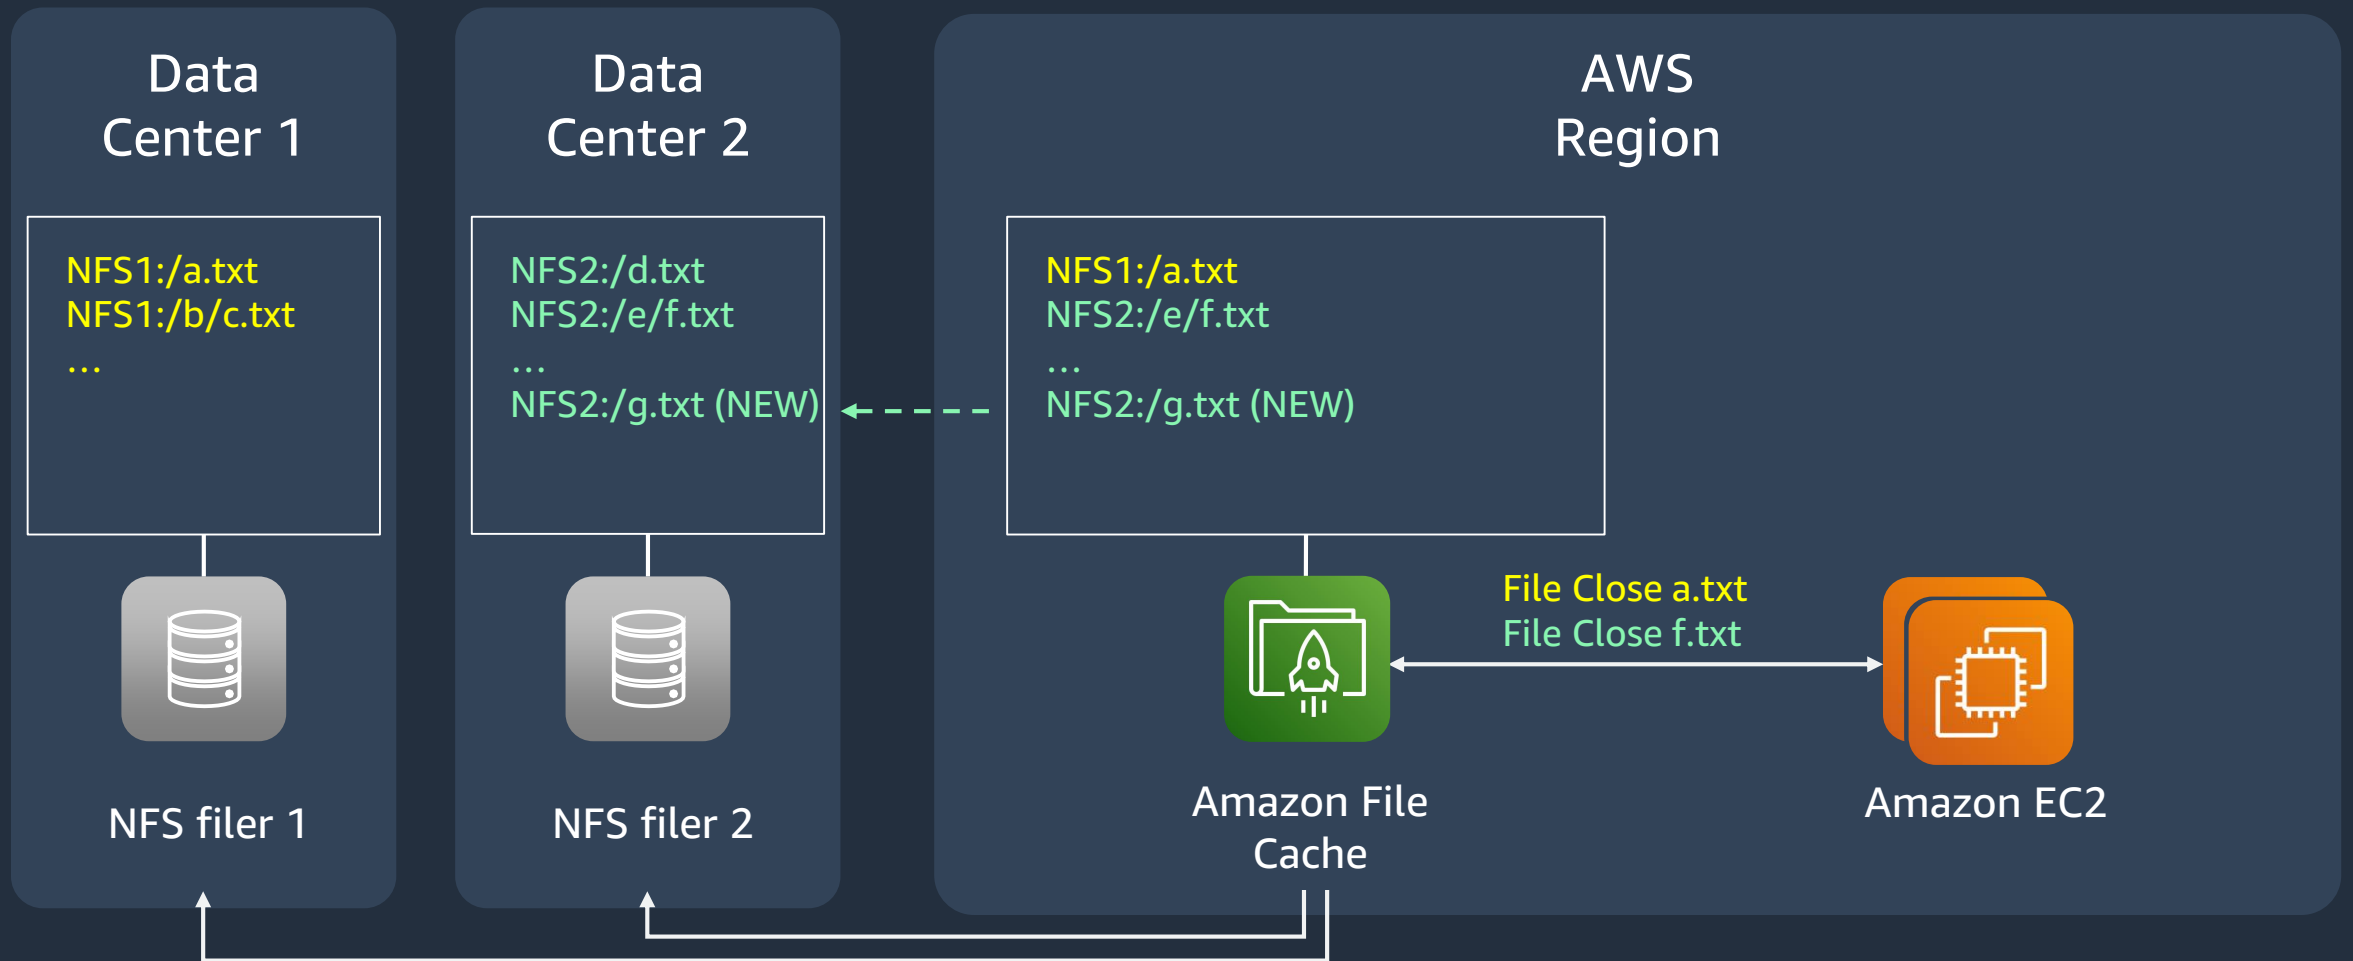
Amazon File Cache: **how it works**



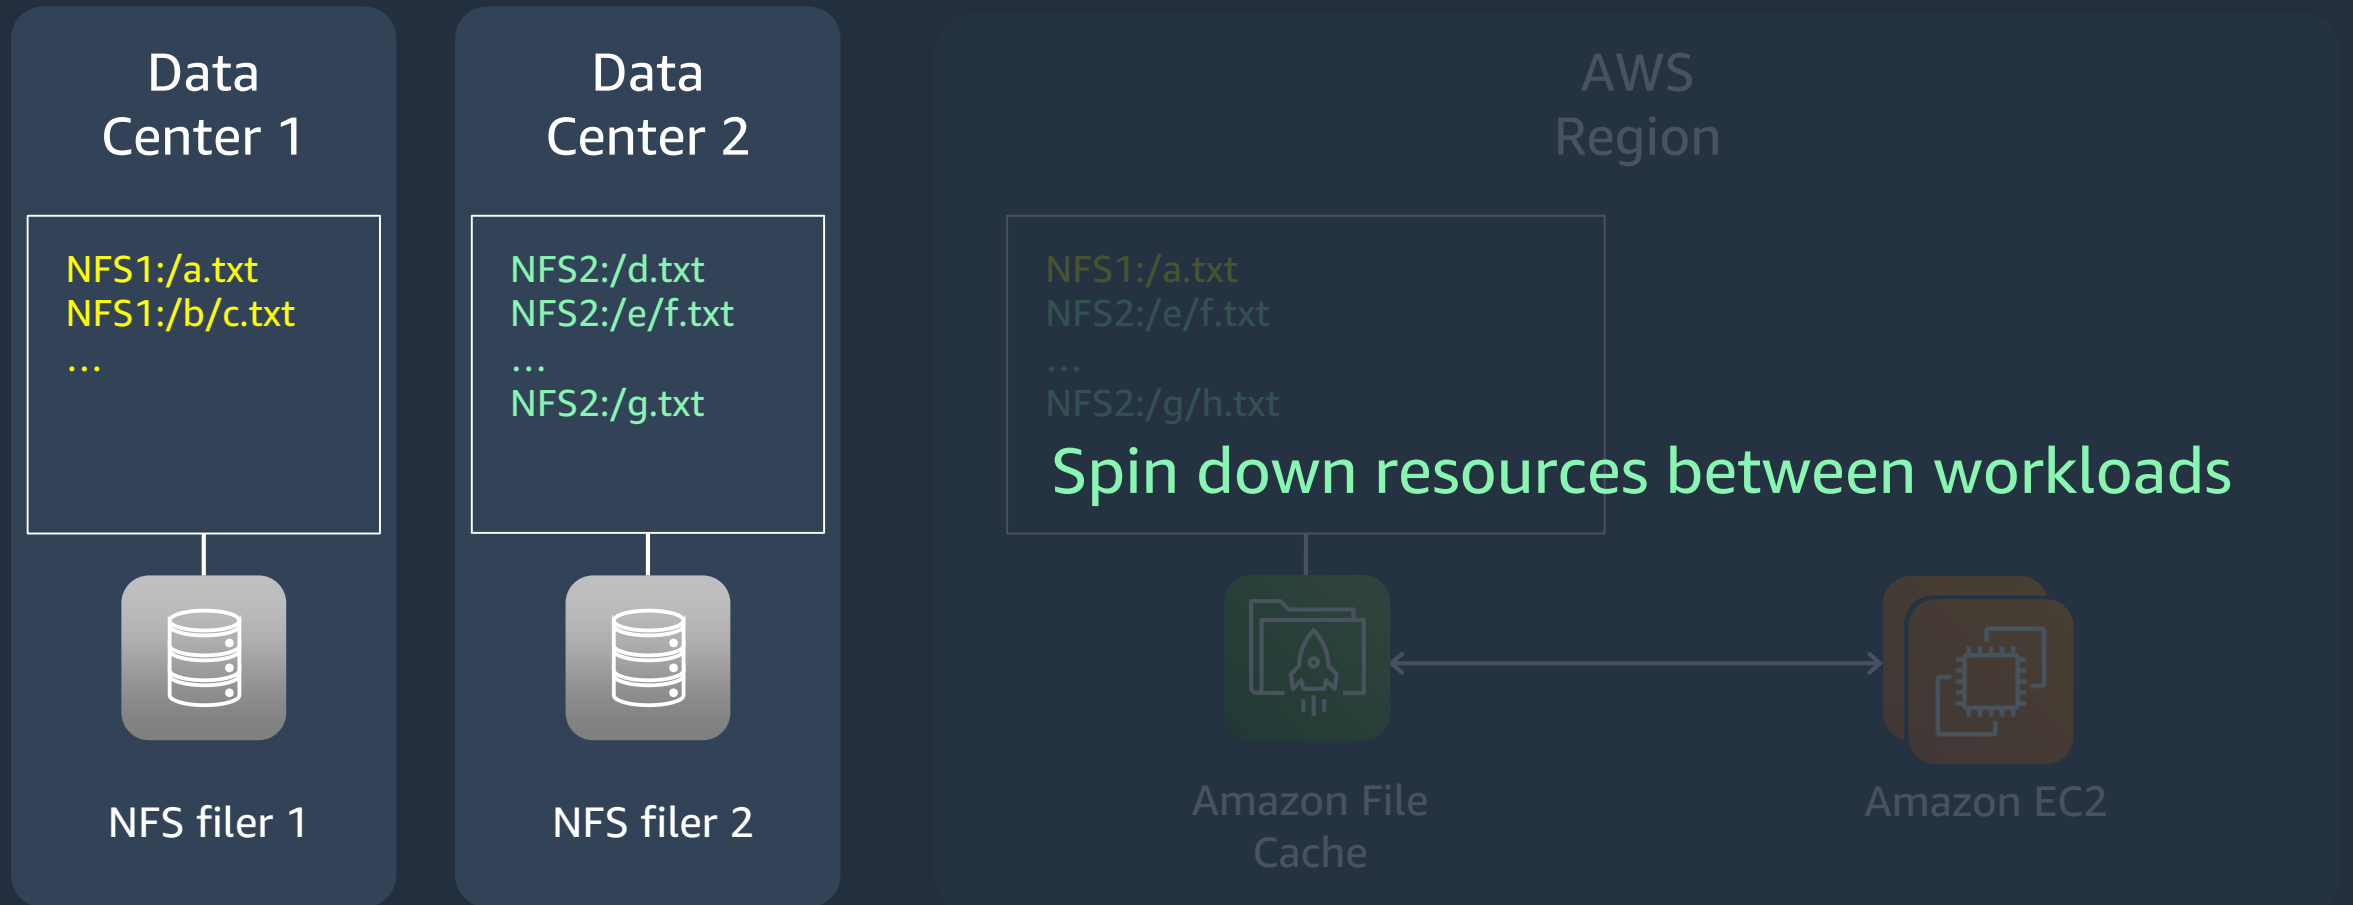
Amazon File Cache **lazy-loads and caches data** read from NFS filers



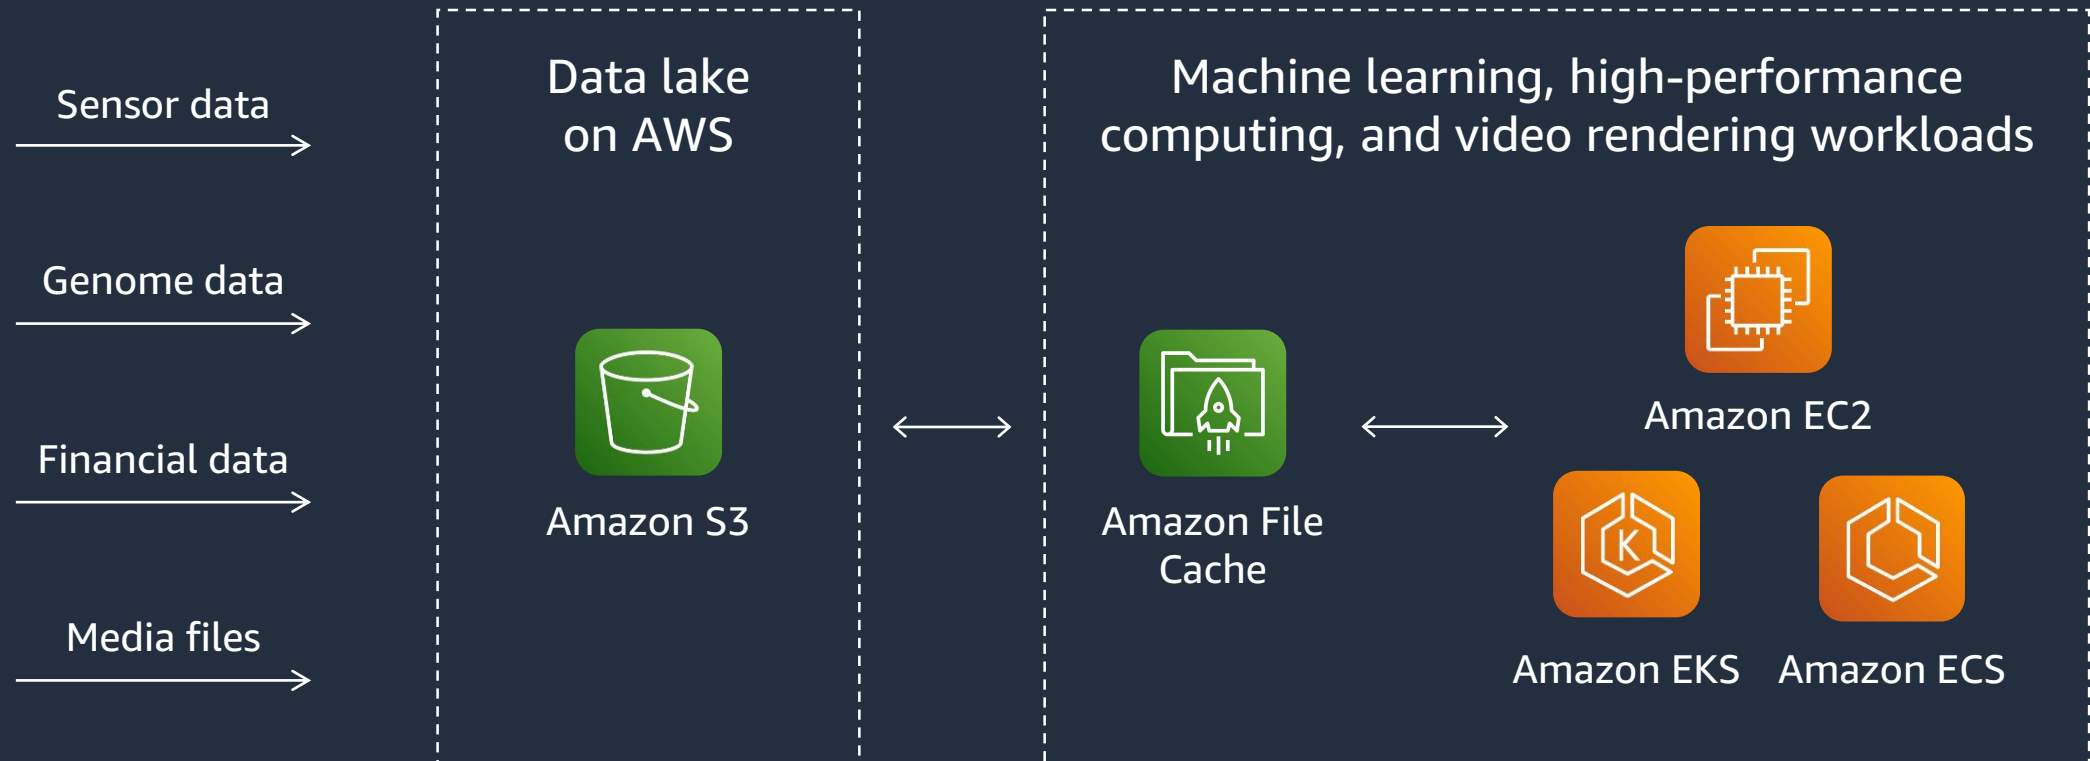
Amazon File Cache updates NFS filers using a "write-back" policy



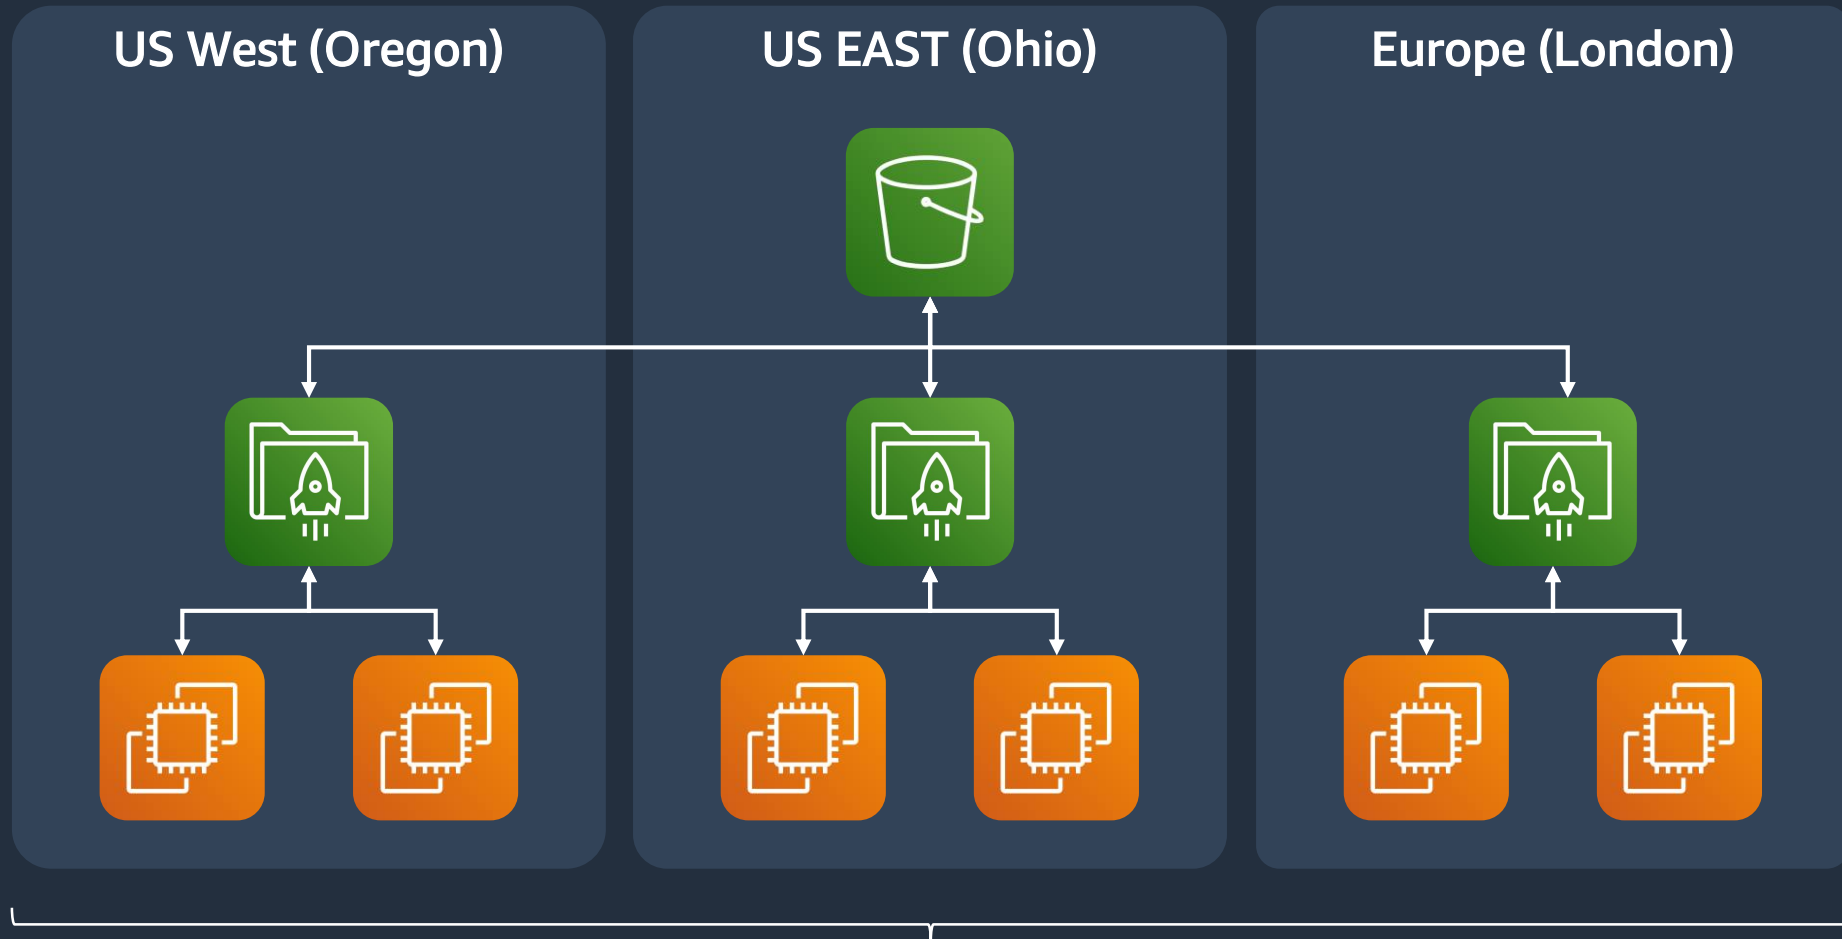
Keep your cache running long-term or **spin-down your cache alongside compute resources to optimize cache costs**



Use Amazon File Cache to access S3 data with a fast file interface



Use Amazon File Cache to share Amazon S3 data across AWS Regions



Multi-region caching

Put your data to work

Amazon File Cache



COMING SOON



Fast, highly scalable
performance



Seamless **data access**
regardless of where you
store your data sets



Thank you!

Jordan Dolman
dolmaj@amazon.com