/* ACADEMY SOFTWARE FOUNDATION

# open source days '26

/* ACADEMY SOFTWARE FOUNDATION


# Move your data once

High-Performance open-source NFS caching
with **KNFSD File Cache**

**Sean Wallitsch**
Senior Specialist Solutions Architect – Visual Compute
Amazon Web Services

# Compute scales easily. Data does not.



# KNFSD - Lineage

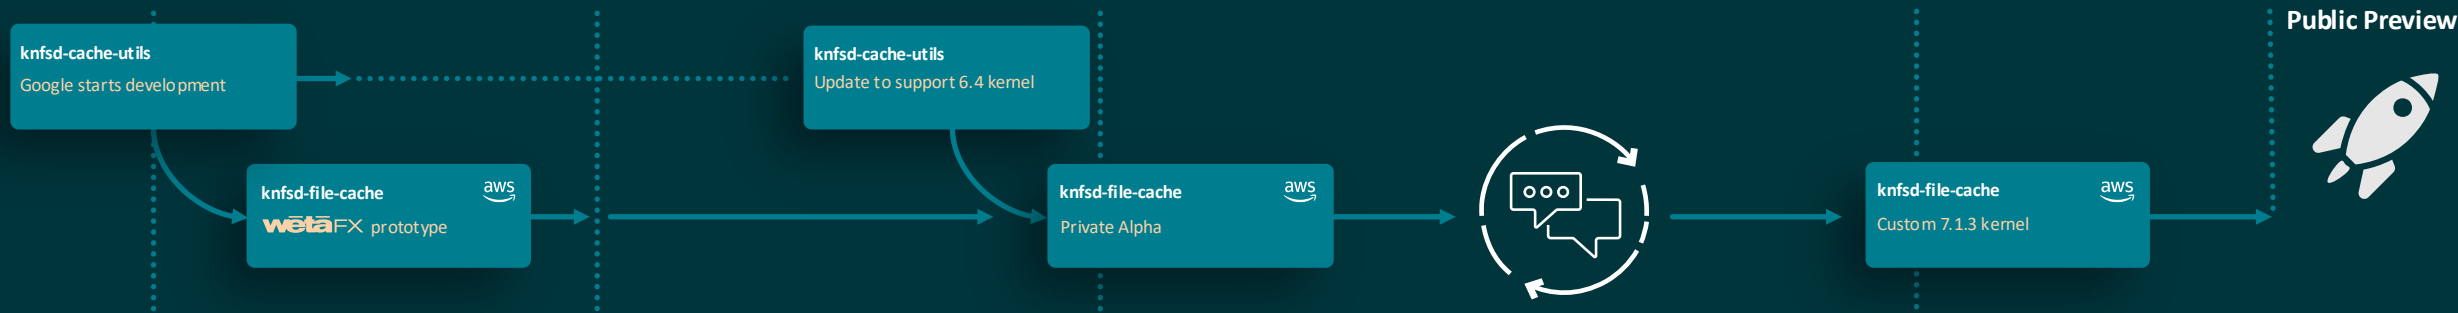


**First prototype**
Developed with Wētā FX in New Zealand, for use on *Avatar 2*

2021

**Fully ported to native AWS**
Developed and supported by AWS
2+ years of updates
VS Code DevContainer Workflow
20% increase in unit testing

2023

**Updated for Distribution**
KNFSD Compute Nodes
Content Security Review
Updated Documentation
Internal Validation

2025

**First Rounds of Private Alpha**
Metrics Dashboard
30+ Merges with 250+ Changes
Bug Fixes and Hardening

2026

# KNFSD – Single node

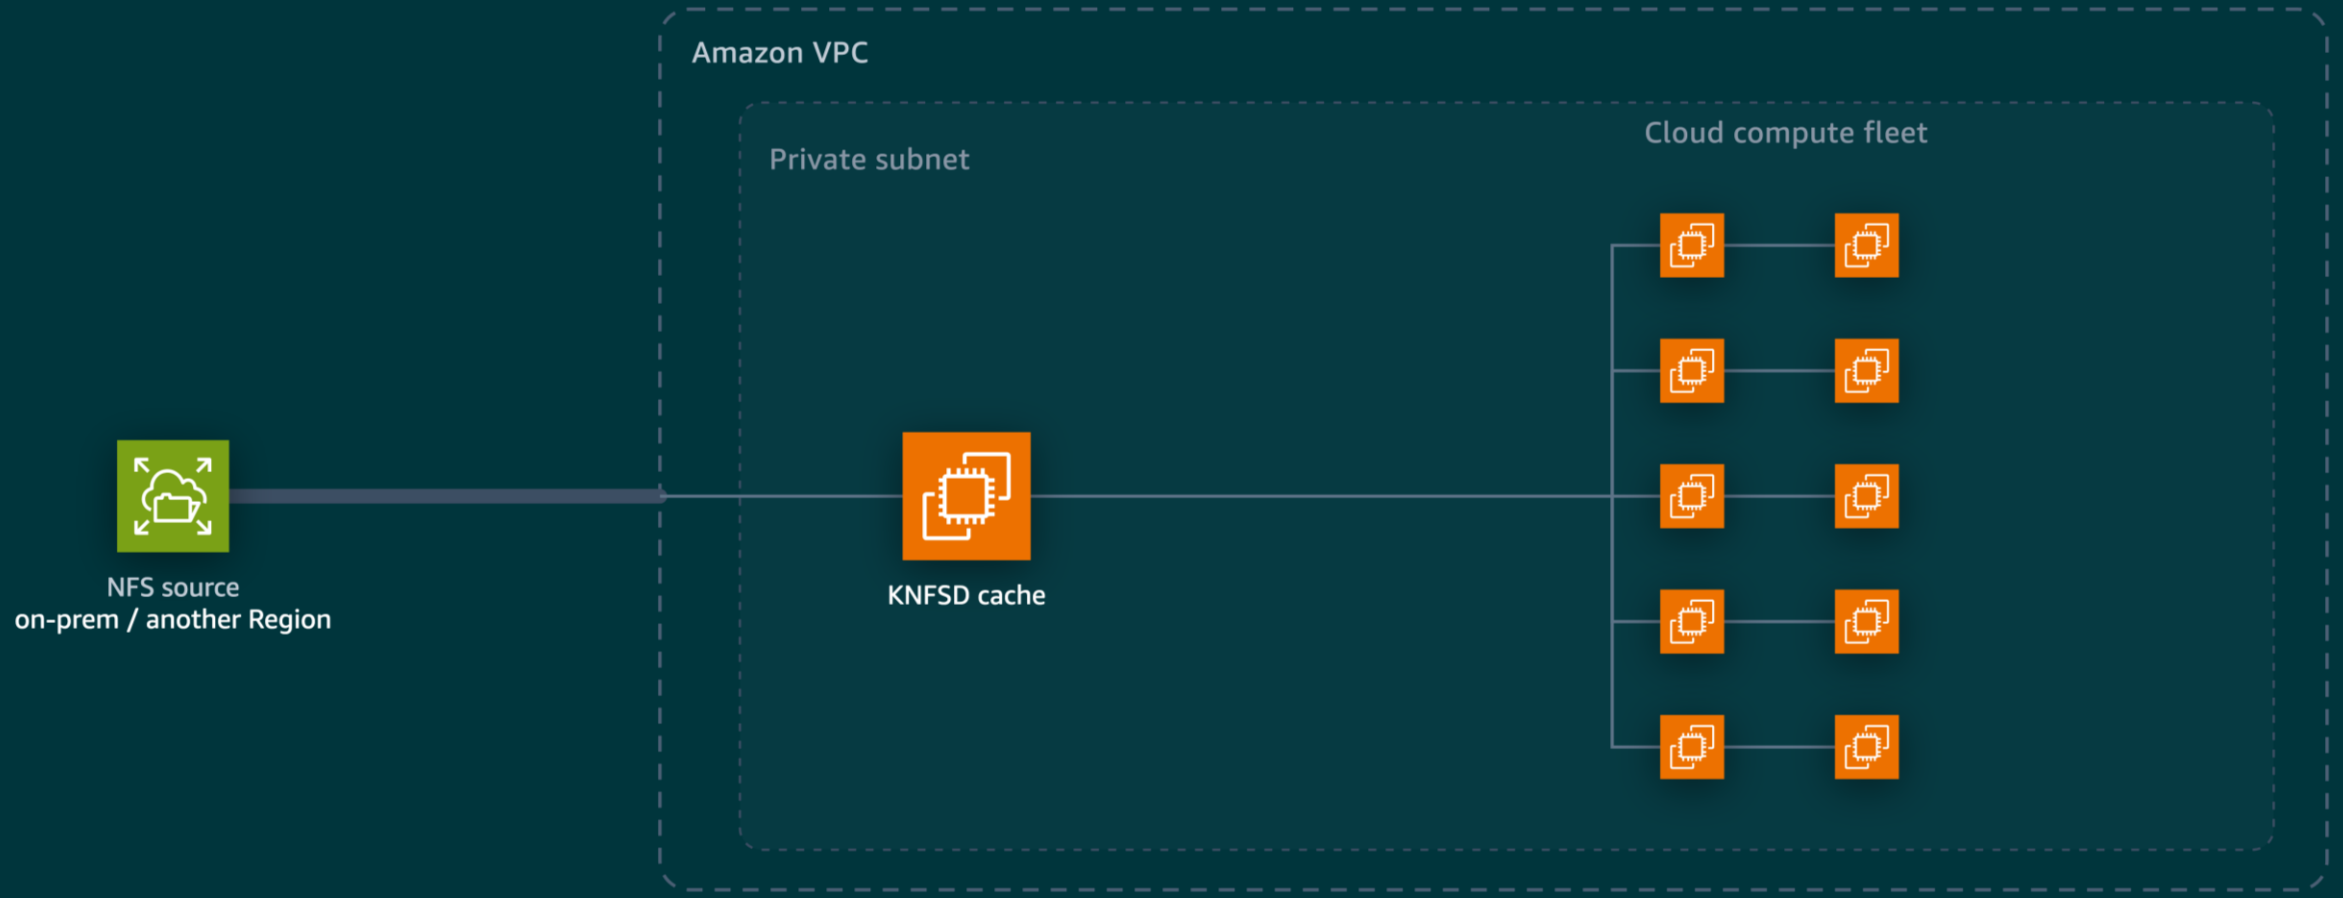

# KNFSD – Single node

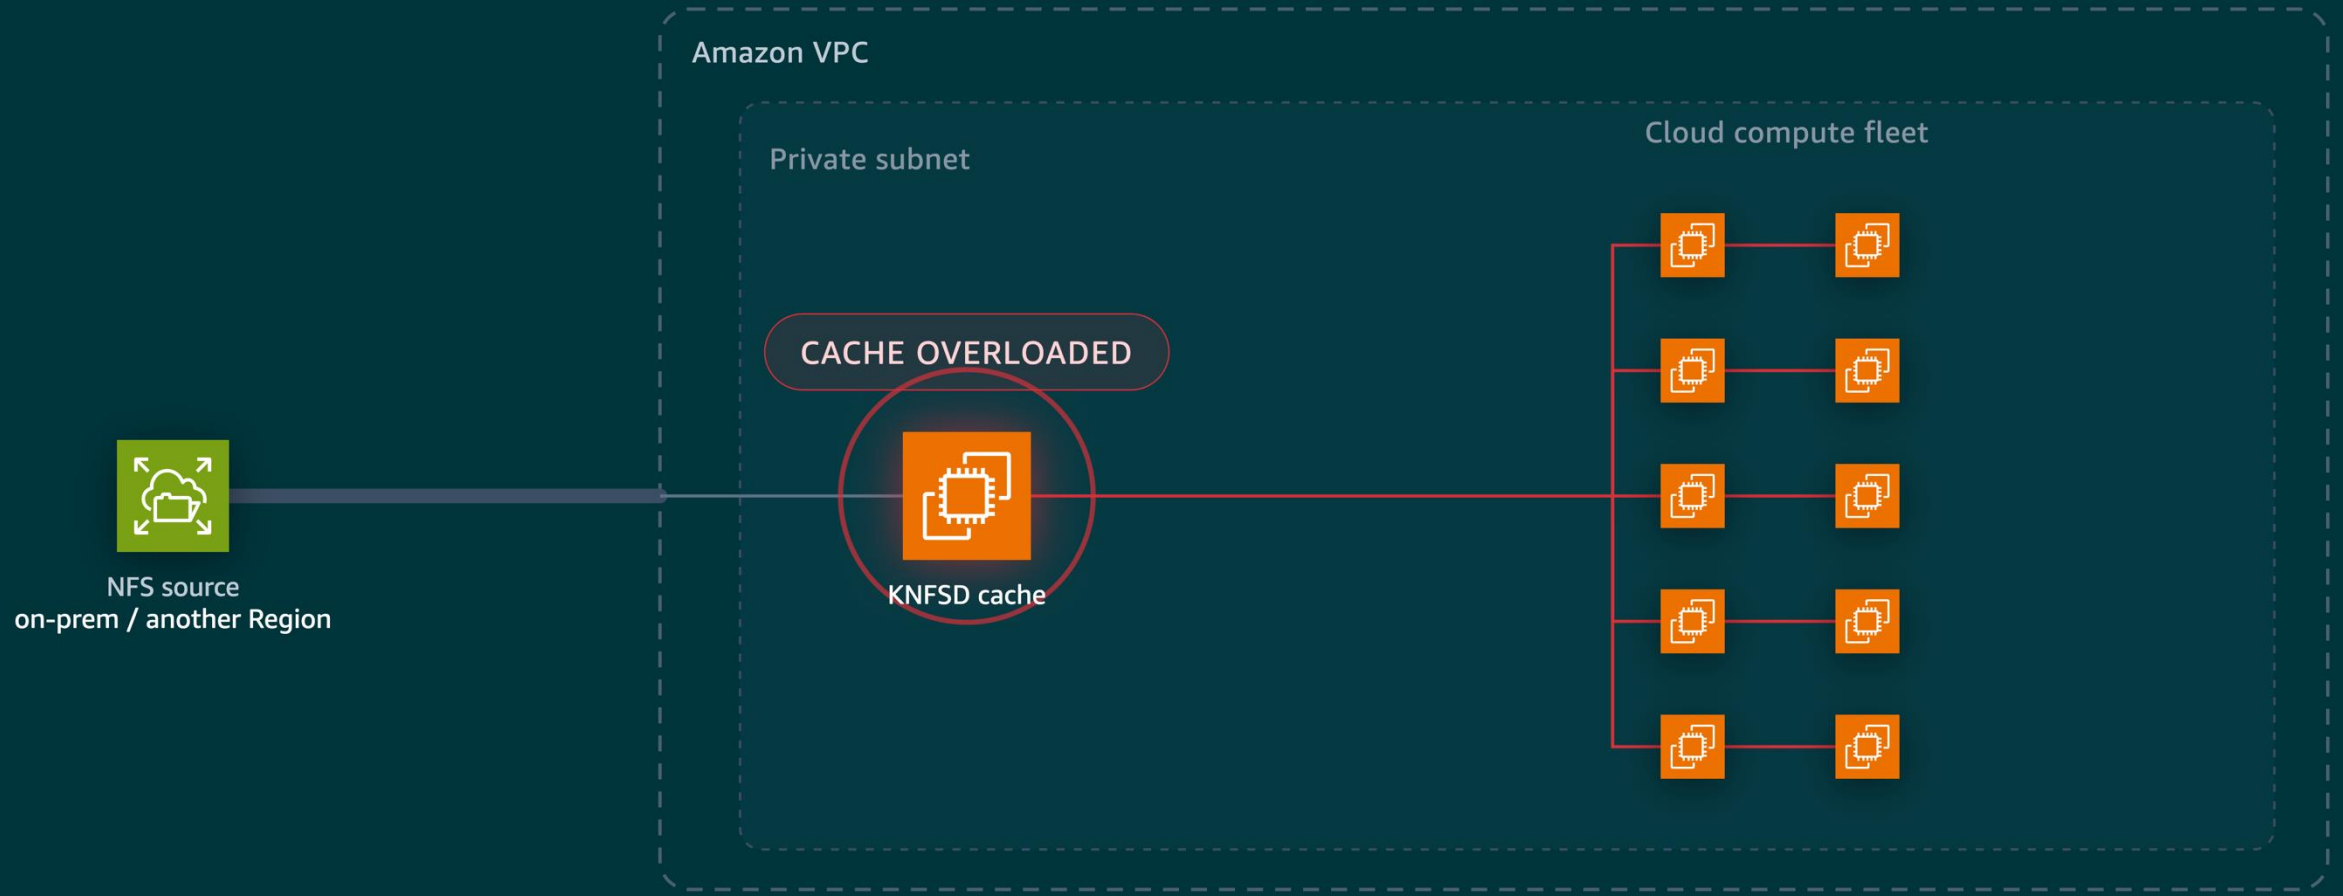

# KNFSD – Auto-scale

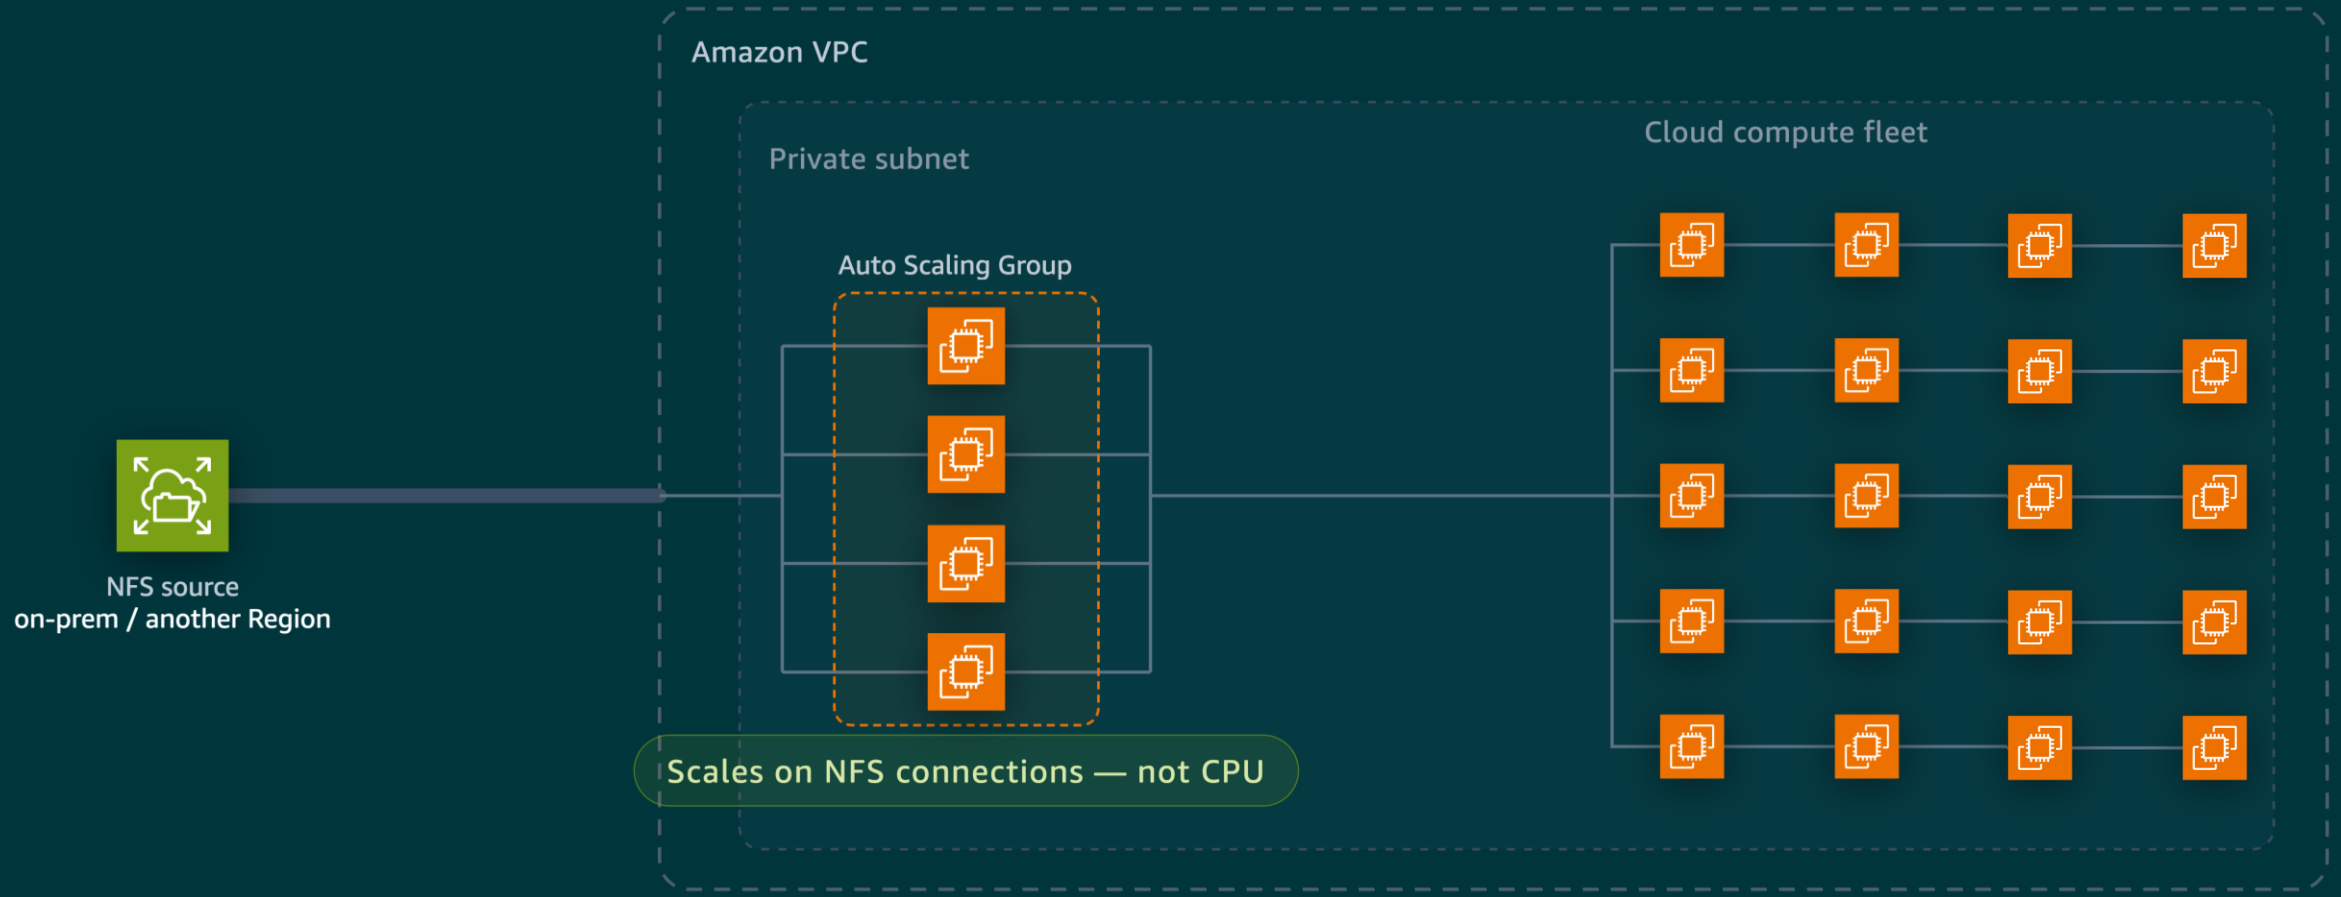

# KNFSD – Auto-scale

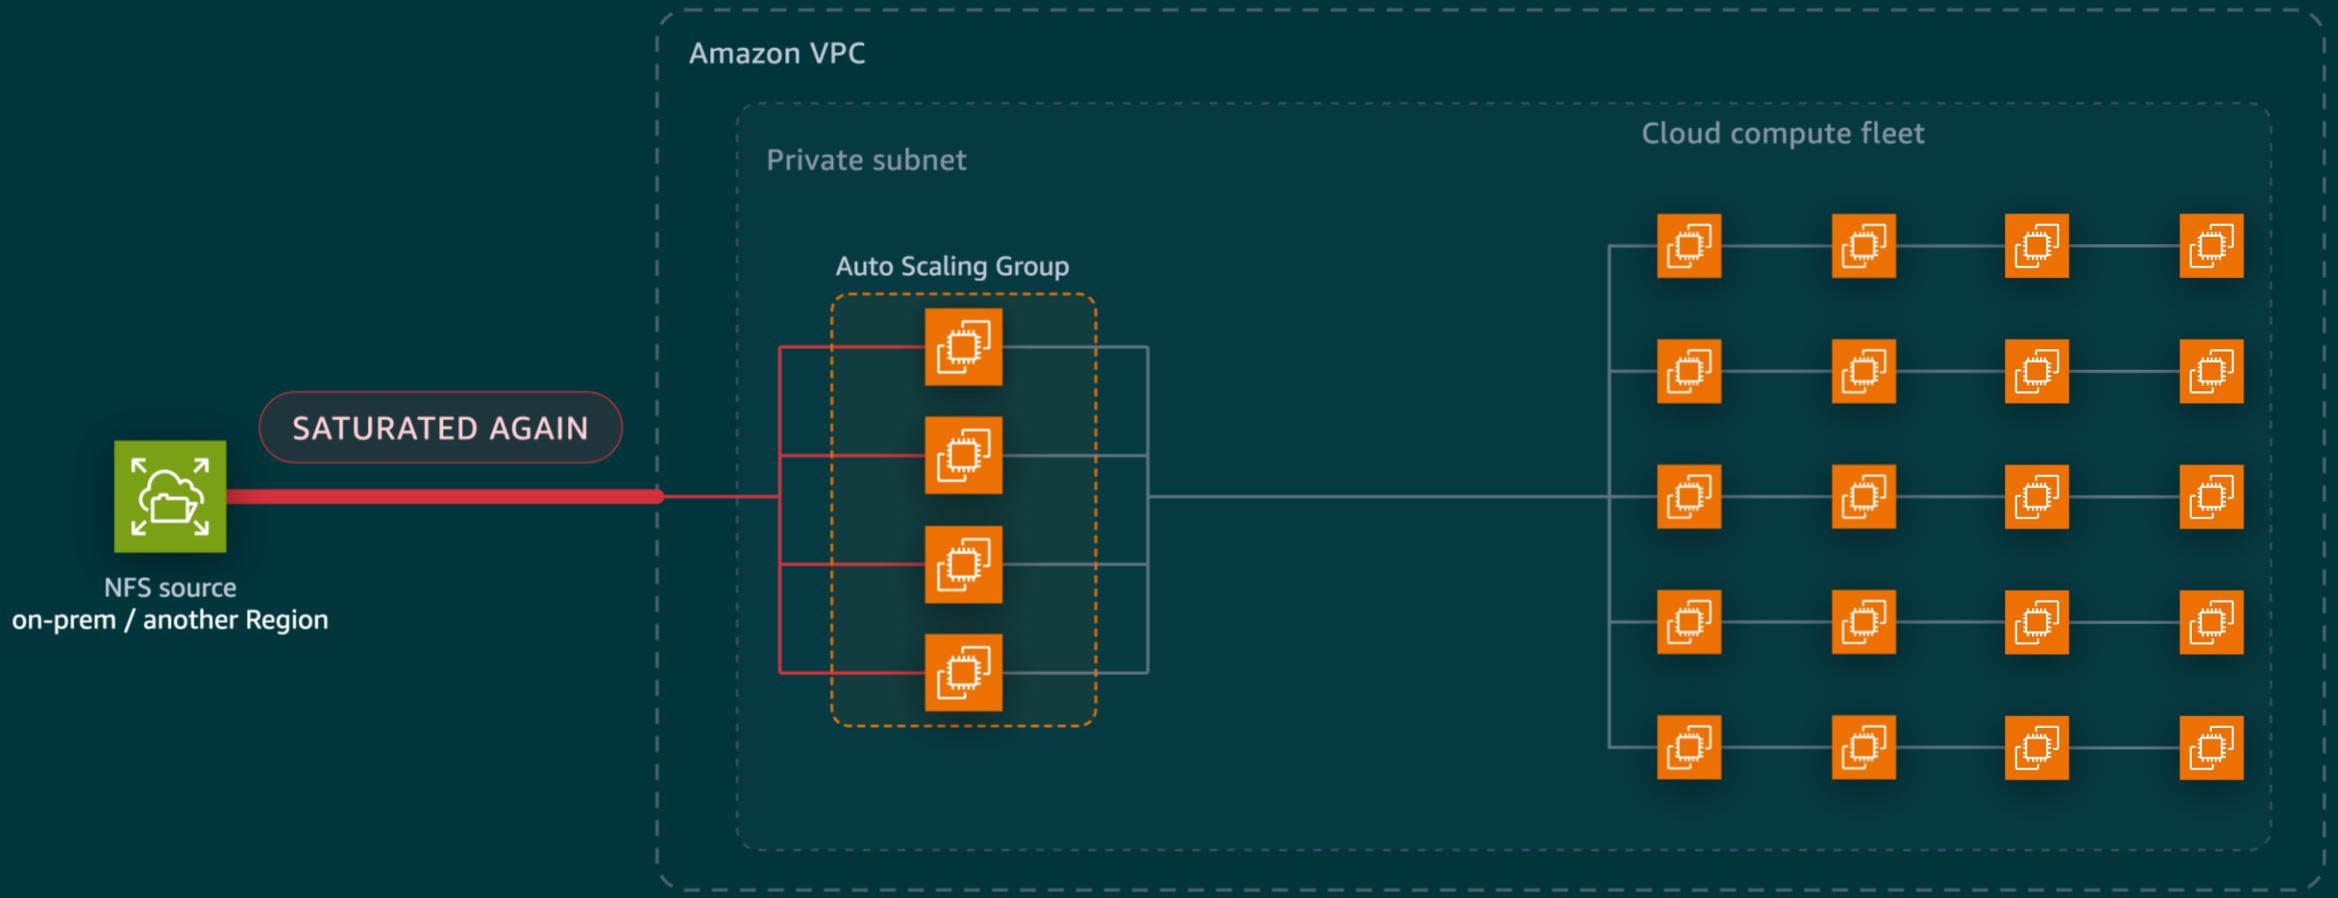

# KNFSD – Two-tier architecture

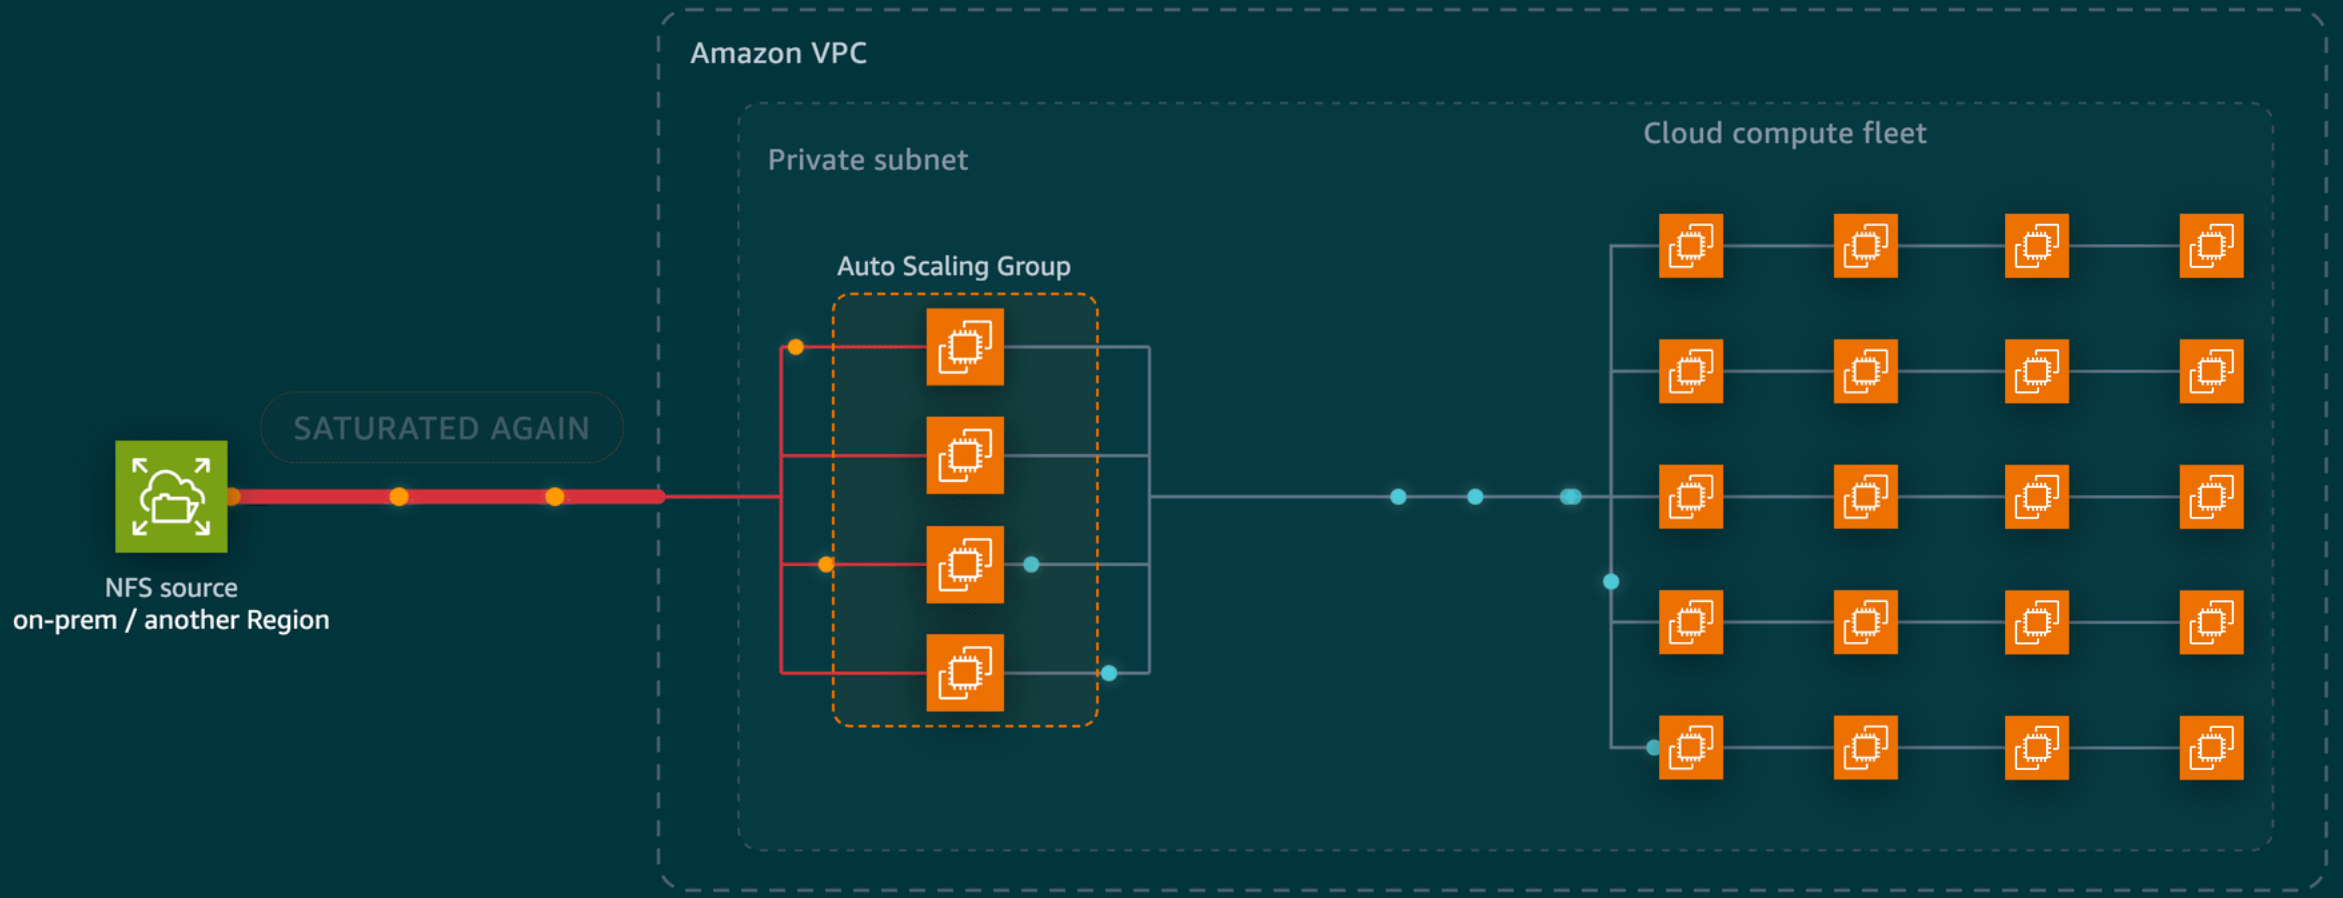

# KNFSD – Open-source stack

| | | |
|---|---|---|
| 1 | Kernel code does the caching | **nfs-kernel-server** re-exports · **FS-Cache** caches to NVMe · both mainline<br>This repo adds the rest: provisioning, scaling, observability |
| 2 | KNFSD tracks the kernel closely | Ships a kernel **compiled from source**, newer than the distro's, with the NFS re-export patches<br>**4+ years** of kernel NFS bug fixes and features arrive release by release<br>Anyone can add a kernel patch; the image rebuilds in **~25 minutes** |
| 3 | The AWS layer has to earn trust | **Renovate + Dependabot:** CVE-driven bumps ship continuously<br>**Smoke tests and CI** gate every kernel upgrade<br>`Checkov` `gosec` `KICS` `Semgrep` `Trivy` `gitleaks` |

# KNFSD – Your choice of hardware

**x86**

Broad compatibility

Run on standard x86 instances

**Arm64**

AWS Graviton

Efficient price-performance

**+ RAM**

More memory

A bigger in-memory cache

**+ Net**

More network

Higher throughput

# KNFSD – Metrics

| Total READ BW | Total WRITE BW | Total IOPS | Cache Hit Ratio |
|---|---|---|---|
| 420G | 1.83G | 39M | 82% |
| • knfsd-azb-asg | • knfsd-azb-asg | • knfsd-azb-asg | • knfsd-azb-asg |

**Connections · Client Counts** 3.4k
Count

**Throughput** 150G
Bytes

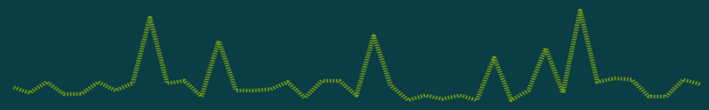

**Latency · p99** 4.2
ms

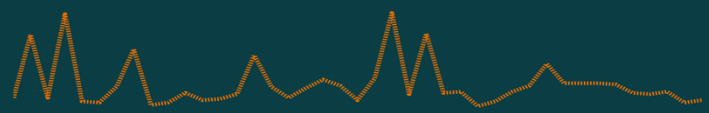

**Cache Hit / Miss Rate** 82
Percent %

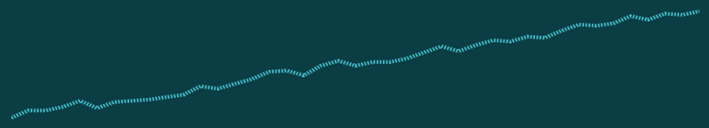

**Queue Depth** 8
Count

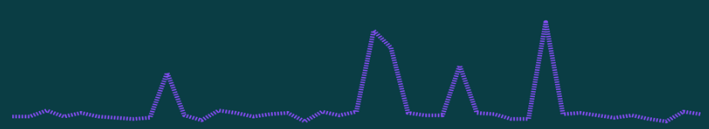

**Per-Mount Throughput** 9.4G
Bytes

# Hardened by production

400,000+ vCPUs

4,000+ workers

*“On Avatar: The Way of Water we ran a cloud render farm 250% the size of our on-prem facility, and the NFS cache we co-developed with AWS was instrumental in scaling access to the shows data.*

*The areas we identified for improvement after Avatar, elastic scaling, load balancing and tiered caching, are precisely what AWS has now built into KNFSD File Cache. We're excited to adopt it across our hybrid render pipeline and contribute to the open-source project.”*

**Andy Wright, Head of Pipeline, Wētā FX**

1. Get the Project
   < 1 min
2. Launch the Dev Container
   ~ 5 min
3. Build Proxy AMI
   ~ 15 min
4. Deploy File Cache Infrastructure
   ~ 15 min
5. Mount and Use the Cache
   < 1 min

# Join us

Apache 2.0 public GitHub repository

Transparent security and library dependencies

Fully configurable

Open to pull requests, issues

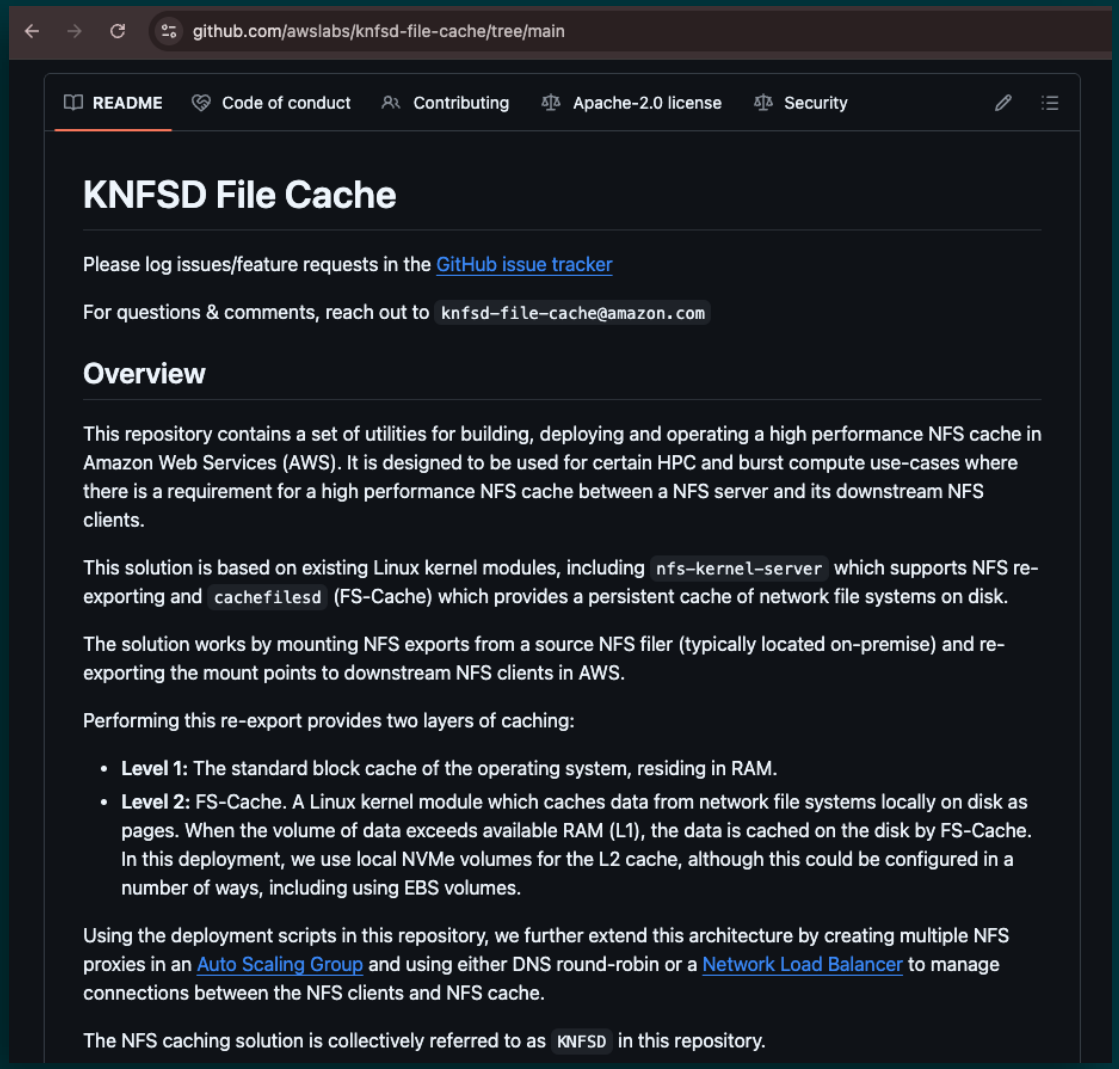

/* ACADEMY SOFTWARE FOUNDATION

open Source days '26

## Get Started

# KNFSD File Cache

GitHub awslabs/knfsd-file-cache

knfsd-file-cache@amazon.com

Explore the AWS Solutions Guidance

**Move your data once, serve it everywhere**