RESUME csherman.net

Christopher Sherman.

csherman@mailbox.org / +1 628 244 6221 / github.com/christophersherman
San Francisco Bay Area / Sydney, NSW | Dual US / Australian Citizen — No visa sponsorship required

Production Engineer and SRE with 5+ years building and operating performance-critical Linux, HPC, and distributed storage platforms for research and enterprise systems.

Experience

Production Engineer, HPC & Cloud AI Infrastructure
November 2024 – Present
Meta (FAIR and Meta Superintelligence Labs)
Menlo Park, CA

- Architected and launched 20+ PB of pNFS storage for FAIR and Meta Superintelligence Labs, providing scalable POSIX access for AI training and research.
- Launched HPC clusters on AWS, CoreWeave, and Oracle within a fleet spanning hundreds of thousands of nodes and 1M+ GPUs, integrating Linux, Slurm, Kubernetes, and parallel storage.
- Built high-throughput POSIX S3 proxies with worker pools and per-thread UID impersonation, preserving access controls and eliminating manual S3-to-NFS staging.
- Operated open-source Kubernetes storage where no internal POSIX solution existed, resolving researchers' storage and data-pipeline bottlenecks.

Senior System Reliability Engineer
October 2023 – October 2024
CyberCX
Sydney, NSW

- Built on-premises, air-gapped monitoring and alerting with Python, Elastic Stack, and Telegraf/InfluxDB/Grafana for security-critical systems; standardized CI/CD and containerized releases.

Site Reliability Engineer
February 2022 – October 2023
Qantas Airways
Sydney, NSW

- Owned releases, code review, SLIs/SLOs, and reliability for mission-critical airline microservices; operated Splunk, Datadog, and APM monitoring in a 24×7 PagerDuty rotation.

Software Engineer
September 2020 – February 2022
SICE Australia
Sydney, NSW

- Developed and tested mission-critical operations-control software and built Python/Ansible automation.

Open Source

Active contributor to Longhorn, K9s, and Terraform's AWS provider.

- **Longhorn:** Backup safety and replica isolation; ongoing work on S3 cleanup and configurable block sizes.
- **K9s:** Worker batching, event-driven refreshes, and metrics pagination for large Kubernetes clusters.
- **Terraform AWS provider:** ENA Express configuration for EC2 launch templates.

Education

Texas A&M University — Bachelor of Science, Computer Science — May 2020

Technical Skills

Systems & Automation: Linux/RHEL/Ubuntu, Python, Bash, Ansible, Terraform, Docker, CI/CD, bare-metal provisioning
Cloud & Platform: Kubernetes, Slurm, AWS, CoreWeave, Oracle; **Storage:** pNFS/NFS, Lustre, OpenZFS, Hammerspace, S3
Networking & Observability: TCP/IP, DNS, Prometheus, Grafana, Splunk, Datadog, Elastic, PagerDuty