

Phil Dibowitz

650-680-6901 | phil@ipom.com | github.com/jaymzh | phildev.net

Engineering leader with 20+ years of experience driving scale, reliability, and automation. Passionate about open-source software, team mentoring, and building resilient infrastructure through deep system insights and strong cross-functional collaboration.

WORK EXPERIENCE

OpenAI, San Francisco, CA (2023 – 2024)

Infrastructure Orchestration Engineer

- Wrote Kubernetes service (in go) to automate and report on health of clusters, **reclaiming thousands of billed-but-unused GPUs** initially and more ongoing, by:
 - Finding and remediating broken and stranded GPUs
 - Finding and reallocating nodes on Infiniband (IB) islands
 - Finding and remediating nodes with underperforming GPUs, IB paths, etc.
 - Finding clusters with unused or abandoned GPU quota and allocate nodes
- Rewrote node initialization and nvidia driver setup code to be more extensible and work across OS versions, **unifying three different implementations** and adding configurability and monitoring
- Wrote tool to find, categorize, and report on CrashLoop containers
- Planned and implemented rollout of CPU Autoscaler to dozens of Kubernetes clusters including dashboards and alerting, **saving dozens of oncall hours per month**
- Designed and built a distributed SSH system extensible to multiple cloud providers and respective authentication mechanisms
- Wrote tool to rightsize CPU quota based on a variety of factors
- Became IB SME for team, running training, writing docs, and being point of escalation
- Updated CI system to show Cilium diffs against each cluster
- Ported a variety of node and IB monitoring and management tools from partner teams including training the team and deployment, eliminating duplicate projects and effort
- Wrote a variety of CLI tools to enable easier Kubernetes node and daemonset management and debugging
- Upgraded MLX and nvidia driver across dozens of Kubernetes clusters with different OS versions
- Partnered with Microsoft to organize weekly meeting that enabled solving a variety of blockers the team had been facing for months
- Significantly improved metrics gathering to get cAdvisor, per-cpu, which vastly improved SEV debugging

Sabbatical (2022 – 2023)

Vicarious, Union City, CA (2019 – 2022)

Manager, Infrastructure Engineering & IT (2020 – 2022)

Infrastructure Engineer (2019 – 2020)

Management, Leadership, and Vision

- Redesigned the interview process to gain higher signal and attract candidates with broader skill sets
- Designed and led cross-training efforts enabling shared oncall across the team rather than individuals being always-on-call for a single aspect of the service
- Built a performance review system (before the company had one) for my team that included downward, peer, and upward feedback

- Worked with HR to design and roll out a consistent engineer leveling system across all teams
- Managed transition to hosted version of GitHub, JIRA, and Confluence, eliminating the overhead of managing these services locally
- Led initiative to build AMIs with packer-chef (vs. shell scripts) to utilize common code and configurations
- Started monthly docs hackathons to make both oncall and onboarding easier and more consistent
- Trained the IT team on coding which evolved into a weekly “coding-n00bs” office hours to enable IT and other teams to learn or grow their coding skills and promote automation
- Led sessions for cross-functional teams wanting support in automating their workflows, leveraging existing tools and infrastructure, and moving to deployable artifacts

Technical

- Rolled out configuration management (Chef) to all local and remote production systems which were previously manually configured, automating configuration as well as preventing configuration drift/rot
- Built training and tooling to ensure consistent operating knowledge across the team of configuration management principles, Chef, and git
- Migrated remote customer systems to a local users and authentication model, managed by Chef (vs. corp-AD-over-VPN), increasing reliability for customer production systems
- Deployed local production artifact (systems packages, releases, Chef bundles, etc.) repositories at each customer site to enable a local source of truth, replacing a fragile corp-AD-over-VPN process and significantly increasing customer site reliability
- Redesigned imaging stack to be deployable at customer sites enabling remote robot repair and reimaging, improving resolution time from days to hours
- Moved robots to mirrored hard drives, significantly reducing chance and duration of failures
- Retooled the build pipelines for the EKS clusters to match actual production and establish a source of truth, as well as implement regular EKS test and upgrades
- Designed and implemented a server-client system to generate on-demand PXE seeds and preseed configs to automate remote imaging of robots without remote console
- Automated Mikrotik and Netgear switch provisioning
- Automated new customer setup (IP allocation, customer IDs, groups, cookbook skeletons, etc.)
- Wrote cli tool to enable frictionless remote machine power-cycling across sites and multiple firewalls.
- Wrote and open-sourced SugarJar (github.com/jaymzh/sugarjar) which simplifies working with git and GitHub
- Built Chef APIs for managing BIOS settings, NetworkManager, SSH, nginx, and sssd as well as contributing them upstream
- Mentored the IT team through coding and technical implementation of automating new employee onboarding & offboarding, enrolling corp servers and laptops in Chef, and managing their fleet via code
- Ported many Chef tools to Windows to help IT in their migration
- Co-redesigned the corp networking to significantly simplify it while adding bandwidth to ease COVID-era remote work

Facebook, Menlo Park, CA (2010 – 2019)

Production Engineer: Operating Systems – Tech Lead (2012 – 2019)

Production Engineer: Traffic Team (2011 – 2012)

Production Engineer (2010 – 2011)

Management, Leadership, and Vision

- Managed redesign and implementation of configuration management with focus on enabling scaling number of systems independent of the infra team size
- Led team to drive company-wide adoption of aforementioned system first through organic growth and eventually strict policy, creating both a unified system and a positive migration experience

- Identified infra areas needing improvement or lacking ownership (security updates, virtualization, package building/distribution, automated change testing) and built a plan to transition to the team and improve them
- Championed an open-source upstream-first mentality within the team: staying close to upstream, pushing features/fixes, releasing tooling wherever possible, etc.
- Identified that future scaling needs would depend on influence of and understanding of community direction. In response, built on existing open source mentality to help team reach out and build relationships with strategic upstream open-source teams including the Anaconda, systemd, RPM and others. Making this a core part of the team's work was a big bet which enabled collaboration with upstream on various long-term visions.
- Grew the team in response to additional responsibility (see point 3 above), including building onboarding documentation and individualized growth plans
- Identified and built key cross-functional relationships between team and others that allowed us to collaborate and build better solutions faster with less stress
- Traveled regularly to remote offices to teach technical and non-technical classes to reinforce cultural consistency/growth across the org and ensure remote employees feel connected
- Individual team member development for team of 8 including: weekly 1:1s focused on individual growth, career growth, and project prioritization; co-writing bi-yearly reviews and defending them in org-wide calibrations; etc.
- Planned and drove bi-yearly team road mapping. Drove cross functional collaboration on team roadmap to ensure org-wide alignment. Socialized road map and previous half-year review.

Technical

- Designed a system to route configuration management alarms to the right team, reducing oncall load on central team
- Wrote automated tooling to sync CentOS updates and roll them out safely on a 2- week cycle, ensuring consistent timely security updates
- Built extensible Chef APIs for a variety of complicated use cases including managing storage devices, mounts, and complex service configs, most of which are now open source
- Worked with auditors to have clean and easy audit reports, and built transparency into tooling to aid in yearly audits
- Led the OS & Load balancer side of project to bring full-parity IPv6 support to Facebook
- Worked with upstream kernel to fix and upstream new features and fixes to the ip6_tunnel module
- One of two primary authors of the automation system used to configure and converge hardware load balancers
- Automated bootstrapping of new clusters enabling infra to keep up with product growth
- Rebuilt internal LDAP infrastructure, improving engineer development problems
- Wrote tooling for new engineers to opt-into accounts and for reaping of unused accounts (for current employees)

Google, Zurich, Switzerland (2008 – 2010) - *Site Reliability Engineer, Gmail*

Ticketmaster, Los Angeles, CA (2005 – 2008) - *Senior UNIX Systems Administrator*

Previous Experience from 2000 through 2005 left off for brevity

EDUCATION

University of Southern California - *B.S. in Computer Engineering Computer Science*

References available on request.