Jon Thor Kristinsson

Founding Engineer | HPC & Infrastructure Architect | 0→1 Product Builder

+44-####-##-0757 | jonthork@gmail.com | github.com/jontk | jontk.com

Impact Summary

Founding engineer and infrastructure architect with 15+ years building HPC and cloud platforms from concept to production across startups and enterprise organizations.

Platform Development Expertise: - Built 4 complete HPC infrastructure platforms from ground zero, each scaling from initial code to production serving Fortune 500 customers - Track record: GreenQloud private cloud deployment → Advania HPCFLOW enterprise scaling → Canonical HPC portfolio creation → Millennium platform transformation

Technical Leadership: - Deep expertise in distributed systems, HPC cluster management, cloud-native architecture, and bare metal provisioning - Combines foundational architecture decisions with rapid product delivery and enterprise scalability

Industry Recognition: - Established thought leadership through presentations at ISC High Performance, Supercomputing, Open Infrastructure Summit, and ARM Dev Summit - Strategic advisory roles including Intel Cloud Board of Advisors and multi-year HPE Centre of Excellence collaboration

Open Source Projects

HPC Developer Ecosystem: Created an integrated suite of open-source HPC tools — a multi-version SLURM API client library consumed by both a Prometheus exporter and a terminal UI — forming a complete management and observability platform for SLURM clusters.

s9s - SLURM Terminal UI

k9s-Inspired Interface for HPC Cluster Management | *Go-Based Terminal Application*

- **Functionality**: k9s-inspired Go terminal UI for SLURM cluster management — real-time monitoring of jobs, nodes, partitions, QoS, accounts, and reservations with vim-style command mode, batch operations, and built-in SSH to compute nodes
- **Scale**: 305 commits, v0.4.0, cross-platform (Linux/macOS/Windows), dedicated documentation site (s9s.dev), GoReleaser pipeline, 6 CI workflows with race detection
- **Architecture**: tview framework with concurrent UI design, plugin system with Prometheus integration, multi-cluster support, JWT auth and data export (CSV/JSON/Markdown/HTML)

SLURM REST API Go Client Library

Enterprise-Grade Multi-Version SDK | *Production-Deployed at Scale*

- **Multi-Version Support**: Comprehensive Go SLURM client supporting API versions v0.0.40-v0.0.44 with automatic version negotiation — probes live clusters and routes to the correct adapter transparently
- **Architecture**: 6-layer design inspired by AWS SDK and Kubernetes client-go — adapter pattern, factory-based version routing, OpenAPI code generation pipeline, circuit breaker, and retry with exponential backoff
- **Features**: Job event streaming via WebSocket/SSE, job efficiency analytics, companion CLI tool, and 100% test coverage on auth/config packages across 250K+ lines of Go

SLURM Prometheus Exporter

Comprehensive HPC Monitoring Solution | *Enterprise Observability*

- **Monitoring Platform**: 80+ Prometheus metrics across 16 concurrent collectors covering cluster health, node states, job lifecycle, user/account fair-share, and partition utilization — with SLA monitoring targets of <100MB RAM and <30s collection time
- **Resilience**: Circuit breaker protecting SLURM head nodes, rate limiting, cardinality management, config hot-reload, and self-monitoring metrics with OpenTelemetry tracing support

Airflow Provider SLURM

Apache Airflow Provider for HPC Clusters | *Released PyPI Package v0.1.0*

- **Workflow Integration**: Released Apache Airflow provider package (v0.1.0 alpha) enabling workflow orchestration on SLURM HPC clusters for ML, bioinformatics, and distributed computing workloads
- **Implementation**: Python package supporting Airflow 2.5-3.x with comprehensive error handling, job recovery, and containerized execution environments

Professional Experience

Lead HPC Engineer (Founding Engineer)

Millennium Management | Jul 2023 - Present

- **Led strategic architecture decisions and delivered enterprise HPC platform** as **lead architect and developer**, defining technical roadmap for quantitative trading infrastructure serving trading teams, portfolio managers, developers, quantitative analysts, and data scientists at global scale
- **Led cross-functional engineering initiative to build complete HPC platform from scratch**, coordinating with trading technology, compliance, and operations teams to deliver SLURM v25.05.5, munge authentication, enroot containerization, and comprehensive platform services with full-stack self-service interface (Go-based APIs/CLIs, React web interface) enabling operational cost reduction and improving researcher productivity
- **Architected hybrid cloud infrastructure** integrating on-premises and public cloud resources into unified, multi-partition computational clusters with high-availability deployments, developing custom AWS EC2 plugin with multi-role support and elastic capacity management
- **Created comprehensive automation ecosystem** for workload scheduling optimization, cluster troubleshooting, and performance analysis across trading operations, establishing monitoring and observability platform with standardized exporters (node, process, cgroup, dcgm)
- **Engineered complete infrastructure transformation** from legacy systems to a 20-role Ansible automation framework covering full-stack HPC delivery — from OS provisioning through SLURM and all auxiliary platform services — with custom-built APIs for dynamic inventory management and service discovery
- **Established operational excellence program**, authoring comprehensive operational runbook guides covering all platform components (slurmctld, slurmdbd, slurmrestd, munge) and change management processes, reducing incident response time and improving team efficiency

Product Manager of HPC

Canonical | Feb 2022 - Jul 2023

- Drove Canonical's HPC market strategy and technical vision, reporting directly to CEO Mark Shuttleworth, establishing strategic partnerships and defining go-to-market approach in the rapidly

growing HPC segment
- Established Canonical as a key HPC platform, driving revenue growth and competitive advantage by leveraging Ubuntu's 14% market share in HPC cluster deployments and its dominance in AI training workloads
- Developed comprehensive HPC product portfolio including extended Ubuntu package ecosystem, automated SLURM-based cluster deployment solutions for both on-premises and public cloud environments
- Forged strategic partnership with NVIDIA to optimize InfiniBand driver performance and GPU integration on Ubuntu, improving the installation experience and easing enterprise adoption
- Led enterprise benchmarking initiative that provided critical performance insights across customer deployments, enabling data-driven product development and competitive positioning
- Drove ISV ecosystem expansion through targeted outreach to HPC independent software vendors, ensuring broad application support for Ubuntu-based HPC environments
- Managed strategic hardware procurement for Canonical, optimizing cluster acquisition strategies to maximize performance per dollar across organizational infrastructure investments
- Authored industry-recognized reference architectures co-published with leading OEMs and hardware vendors, establishing thought leadership and guiding market hardware acquisition decisions
- **Built thought leadership presence** through presentations at premier industry events including Open Infrastructure Summit, ARM Dev Summit, Operator Day, NASA Cybersecurity Day, and OpenInfra Live
- **Industry recognition**: Established as key voice in HPC community through conference speaking and technical panel participation
- Published comprehensive HPC content strategy including technical blog series and industry book, establishing Canonical's voice in the HPC community

Interim Product Manager of MAAS (Additional Role)

Canonical | Apr 2022 - Jul 2023

- Defined strategic product roadmap for MAAS (Metal as a Service), focusing on bare metal server delivery and HPC cluster provisioning capabilities
- Leveraged deep HPC expertise to architect critical features including diskless booting, InfiniBand support, and advanced cluster delivery mechanisms that addressed enterprise-scale deployment requirements
- Drove hardware integration initiatives by designing firmware upgrade procedures, BIOS management capabilities, and enhanced server metrics collection to improve operational efficiency and hardware lifecycle management
- Established five-year product vision and roadmap for MAAS, providing clear strategic direction that influenced current development cycles and feature prioritization
- Collaborated with engineering teams to translate complex HPC requirements into actionable product features, ensuring alignment between technical capabilities and market needs

Chief Technology Officer of HPCFLOW

atNorth / Advania Data Centers | Jan 2018 - Dec 2021

- **Led HPCFLOW platform expansion** as CTO when product transitioned to Advania Data Centers, scaling from single-region founding solution to multi-regional HPC-as-a-Service offering serving enterprise customers
- Developed enterprise-grade HPCaaS capabilities including Intel Select Solutions for High-Performance Computing and Open HPC packaging-based environments, leveraging automation tools including Terraform, Ansible, and cloud-init

- Led customer engagement strategy in pre-sales and solution architecture capacity, driving technical sales and custom solution design for enterprise HPC requirements
- Architected and managed distributed storage infrastructure using both Community Ceph and Red Hat Ceph deployments across multi-year operational cycles
- Engineered custom networking solutions by developing two switch implementations for OpenStack's network-generic-switch ml2 plugin, enabling multi-tenant bare metal networking for HPE Flex Fabric and Cumulus OS environments
- Pioneered Omni Path vFabric multi-tenancy support, implementing fabric virtualization that integrated seamlessly with OpenStack Neutron's port allocation model
- Designed comprehensive infrastructure configurations spanning HPC, AI, visualization, hypervisor, storage, Ethernet, and Omni Path fabric systems
- **Led industry thought leadership** through presentations at premier conferences including ISC High Performance (Europe), Supercomputing (US), and served on technical panels
- **Strategic industry advisory role**: Served on Intel's Cloud Board of Advisors, providing strategic guidance on HPC cloud technologies and market direction
- **Academic research collaboration**: Led HPCFLOW infrastructure for Stanford University Living Heart Project in partnership with Ubercloud, providing HPC expertise and hardware utilization guidance enabling breakthrough cardiac simulations
- **HPE Centre of Excellence**: Multi-year collaboration (5+ years) with HPE HPC team in Grenoble, France, contributing to product development and customer solutions

Senior Solution Architect HPC (Founding Engineer of HPCFLOW)

Advania | May 2016 - Jan 2018

- **Founded and built HPCFLOW from ground zero** - comprehensive IaaS platform for HPC and AI workloads, serving as sole founding engineer and architect. Integrated multiple open-source technologies including OpenStack, Ironic, Packer, Slurm, and Ceph with custom solutions to deliver HPC clusters as a service
- Led pre-sales technical consulting for HPC customers, enabling seamless migration to IaaS-based cloud infrastructure while maintaining on-premises performance levels using bare metal resources
- Optimized and validated HPC workloads to ensure peak performance and efficiency across diverse computational requirements
- Pioneered early Kubernetes adoption (starting with v1.0 in 2016), deploying and managing clusters to host internal services and enhance monitoring capabilities across the organization
- Leveraged HPE Performance Cluster Manager (formerly CMU) to automate HPC cluster deployment from golden images, streamlining infrastructure provisioning and management

Senior Deployment Engineer

GreenQloud (acquired, now Netapp Cloud) | Aug 2014 - Apr 2016

- Architected and developed cloud infrastructure using Chef automation and Apache CloudStack, enabling scalable private cloud solutions for enterprise clients
- Pioneered automated deployment processes for private cloud components, reducing setup time and enabling Greenqloud's private cloud service offering
- Implemented continuous delivery pipeline that accelerated software component delivery and significantly improved product quality through automated testing and deployment
- Integrated Swift object storage solution into Greenqloud's private cloud platform, expanding storage capabilities and service offerings
- Developed innovative bare metal provisioning capabilities for Apache CloudStack, including automated ethernet and InfiniBand switch configuration supporting multi-tenant environments

- Delivered specialized HPC infrastructure solutions for clients in industries such as automotive, enabling high-performance computing on bare metal resources

Core Developer

GreenQloud (acquired, now Netapp Cloud) | Dec 2013 - Aug 2014

- Extended core capabilities of Apache CloudStack using Java to enhance cloud platform functionality and performance
- Developed custom features and integrations that expanded the platform's enterprise capabilities, contributing to Greenqloud's competitive cloud infrastructure offerings
- Collaborated with engineering teams to implement scalable solutions that supported the company's growth

Programmer

Destino | Jun 2011 - May 2013

- Developed enterprise content management system as part of a collaborative development team, delivering scalable web-based solutions for client content publishing and management needs
- Implemented Agile development practices by introducing SCRUM methodology and continuous integration workflows that improved team productivity, code quality, and delivery timelines
- Built full-stack web applications using C# for backend services and JavaScript for interactive frontend functionality, ensuring robust performance and user experience
- Collaborated in cross-functional development team to deliver iterative software releases, participating in sprint planning, daily standups, and retrospectives to maintain project momentum and quality standards
- Established CI/CD pipeline processes that automated testing and deployment procedures, reducing manual errors and accelerating feature delivery cycles

Education and Training

BSc Computer Science

Reykjavik University, Iceland | 2009 - 2011

Completed 2 years out of 3 towards a BSc in Computer Science (ended May 2011)

High School Degree (Equivalent to GCSE)

Menntaskolinn Hradbraut, Iceland | 2007 - 2009

Technical Skills

Programming & Development

- **Languages**: Go, Python, Java, C#, JavaScript, Bash
- **Development Tools**: Git, CI/CD, GitOps, ArgoCD, GitHub Actions, JSON/API integration

HPC & High Performance Computing

- **Core Technologies**: Slurm (Expert - v25.05.5), Intel Omni Path, InfiniBand, MPI, OpenHPC, Apache Airflow, Enroot
- **Workload Expertise**: ANSYS Fluent, OpenFOAM, Numeca Fine Marine, Abaqus, WRF, FDS
- **Client Experience**: CFD deployments across automotive, pharmaceutical, and research organizations

Infrastructure & Cloud Platforms

- **Cloud Platforms**: AWS (EC2 plugin development), OpenStack, Apache CloudStack, Azure
- **Infrastructure Automation**: Ansible (Expert), Chef (Advanced), Terraform, Nix, cloud-init
- **Container Technologies**: Kubernetes, Docker, Containerd, k3s
- **Storage Systems**: Ceph, Swift Object Storage, NFS, Lustre

Specialized Hardware & Architecture

- **GPU & AI Infrastructure**: Nvidia DGX systems (certified), CUDA optimization, VirtualGL, NICE DCV
- **ARM Architecture**: Ampere CPU systems, ARM-based HPC deployments, cross-architecture optimization
- **Operating Systems**: Linux (Ubuntu, CentOS, RHEL), NixOS, Windows Server

Monitoring & Networking

- **Monitoring & Observability**: Prometheus, Grafana, multi-exporter deployment, ELK Stack, OpenTelemetry
- **Networking**: SDN, OpenFlow, Neutron, Cilium, multi-tenant networking, load balancer integration

Notable Projects & Achievements

Stanford University Living Heart Project | *Academic Research Collaboration* - **Impact**: Led HPCFLOW infrastructure provision for breakthrough cardiac simulations in collaboration with Stanford University and Ubercloud - **Role**: Provided HPC expertise and hardware utilization guidance, optimizing cluster configurations for advanced computational cardiovascular research - **Recognition**: Published research with academic citations, demonstrating real-world application of HPCFLOW platform in medical research - **Technology**: Advanced CFD workloads, custom HPC cluster configurations, high-performance visualization

Key Achievements

- **Platform Builder**: Created 4 complete infrastructure platforms from ground zero, each scaling to enterprise production
- **0→1 Execution**: Proven ability to take concepts from napkin sketch to production systems serving real users
- **Business Impact**: Platforms achieved enterprise adoption with Fortune 500 customers and measurable cost savings
- **Technical Innovation**: Pioneered early Kubernetes adoption (v1.0), custom GPU scheduling algorithms, hybrid cloud architectures
- **Open Source Impact**: Created production-ready HPC tools used by enterprise customers across multiple industries
- **Community Leadership**: Published multiple open source projects with comprehensive documentation and enterprise patterns
- **Industry Recognition**: Presented at premier conferences including ISC HPC, Supercomputing, Open Infrastructure Summit
- **Product Strategy**: Led Ubuntu's HPC strategy resulting in increased market share and strategic partnerships