

Adam DeConinck

✉ ajdecon@ajdecon.org

🌐 <https://www.ajdecon.org/>

🌐 <https://www.linkedin.com/in/ajdecon/>



Summary

Computer systems engineer with extensive experience designing, building, and operating cluster computing systems at scale for machine learning and simulation workloads, as well as investigating and solving complex hardware and software issues in large distributed systems.






Professional Experience

- 2023 – Present **Senior Engineering Manager**, NVIDIA Corporation.
- Manager of a team focused on designing next-generation systems for use by both internal teams and as a template for external deployments
 - Develop cluster architectures based on novel hardware and software systems
 - Future planning of datacenters and cluster systems for internal use at scale
- 2022 – 2023 **Principal Systems Software Engineer**, NVIDIA Corporation.
- Systems engineering and automation for large supercomputing systems, including NVIDIA Eos based on DGX H100 and NDR InfiniBand
 - Investigate and diagnose issues on next generation hardware platforms
 - Design new cluster architectures for AI and HPC applications at scale
- 2021 – 2022 **Product Architect**, NVIDIA Corporation.
- Subject-matter expert collaborating with engineering and product management
 - Design of new system software stack for DGX SuperPOD AI supercomputers
 - Worked closely with early SuperPOD customers to incorporate their feedback
- 2019 – 2021 **Senior Solutions Architect**, NVIDIA Corporation.
- Customer-facing system architect for supercomputing design and deployment
 - Lead for multiple RFP responses for supercomputing systems at government customer sites, including deals with a total value of >\$60MM.
 - System software lead for initial customer deployments of DGX SuperPOD
- 2017 – 2019 **Production Engineer**, Facebook, Inc.
- Production Engineering tech lead for metadata layer of Facebook Tectonic, an exabyte-scale distributed filesystem deployed in multiple regions
 - Deployment automation and automated capacity planning tools
 - Onboarding and mentoring of new team members
 - Helped organize multiple team events for Seattle production engineering team
- 2014 – 2017 **HPC Cluster Administrator**, Los Alamos National Laboratory.
- Systems engineer supporting HPC capability systems for major DOE workloads
 - Production team lead for deploying Trinity, a 20,000-node Cray XC40 system
 - Coordinated the production change control process for LANL HPC systems
 - Derivative classifier to review the release of sensitive information (Q clearance)

Professional Experience (continued)




- 2012 – 2014  **HPC/Cloud Systems Engineer**, NVIDIA Corporation.
- Built and managed an HPC platform for the Solution Architect team
 - This resource also supported customer evaluation workloads, as well HPC development for teams across the company.
- 2010 – 2012  **Systems Engineer**, R Systems NA, Inc.
- Design, deploy, and support hosted HPC clusters for industrial customers

Technology Experience




- Programming languages  Production code in Python, Go, Perl, Bash, and Fortran; Small tools and side projects in Rust, C++, and C
- System orchestration  Ansible, Terraform, Cfengine
- Workload management  Kubernetes, Slurm, PBS/Torque/Moab, Grid Engine
- Cloud platforms  AWS, Azure, Equinix Metal
- Storage technology  Lustre, Ceph, HDFS, Swift, MinIO; storage development at Facebook

Miscellaneous Experience



Open Source

- 2019 – 2022  **NVIDIA DeepOps**, GPU cluster deployment toolkit for HPC and machine learning. Co-maintainer, extensive work on both Kubernetes and Slurm cluster deployment
- 2013 – 2014  **EasyBuild**, scientific software package manager. “Dry-run” functionality and added recipes
- 2011 – 2012  **Warewulf**, cluster provisioning system. IPMI functionality, misc bug fixes, beta tester for 3.0 release

Community Involvement

-  Program committee participant for several academic and industry conferences, including USENIX LISA, IEEE CLUSTER (for the HPCMASPA workshop), and the ACM SIGHPC-SYSPROS workshop.
-  Regular participant in HPC community discussions and panels as part of ACM SIGHPC-SYSPROS.
-  Member of the Denver Curling Club.

Education

- 2007 – 2010  **M.S., University of Illinois at Urbana-Champaign**, Materials Science and Engineering. Thesis title: *Fabrication, dynamics and self-assembly of anisotropic colloidal particles.*
- 2003 – 2007  **B.S., Michigan Technological University**, Physics. Minors in Mathematics and Electronic Materials.