

---

**Experience**

---

**Software Engineer Systems (III)                      Hewlett Packard Enterprise                      July 2023 – Present**

**Core engineer on HPE's flagship software-defined scale-out storage platform** (distributed file system and S3 object storage on petabyte-scale, all-NVMe clusters), from pre-1.0 development through five major releases as supported clusters grew from 3 to 16 nodes

- Designing the server-side locking strategy backing NFSv4 client file locks; architected a reusable IO range-lock library for the file system, establishing the foundation for race-free, high-performance file operations (*C++*)
- Implemented core native file system operations (getattr, setattr, link, unlink, rename), delivering POSIX-conformant attribute and namespace semantics (*C++*)
- Delivered S3 object versioning, core features of the platform's second release, by implementing GetObjectTagging, PutObjectTagging, CopyObject, and DeleteObjects APIs across the C++ storage backend and the MinIO gateway (*C++, Go*)
- Led POSIX compliance initiative directing 3 engineers; achieved zero failures across the 8000+ tests of the pjdftest suite by resolving critical bugs in file operations, materially improving platform reliability and standards conformance (*C++, Python*)
- Owned design and backend implementation of POSIX mode permissions, a core feature of the platform's third release; led 3 engineers and coordinated cross-team execution to deliver the feature on schedule (*C++, Python*)
- Led ASAN integration with the unified memory allocator to enable proactive, platform-wide memory-defect detection; triaged and resolved high-severity defects across memory management, back-pressure, and watchdog-timer (*C++*)
- Hardened 1.0 data-path infrastructure for general availability: stabilized the distributed caching with back-pressure observability and metrics, and augmented supervisory repair infrastructure enabling automated recovery across submodules (*C++*)
- Improved developer experience by authoring reusable GDB macros and internal documentation, cutting bug-triage time from hours to minutes; drove dozens of code reviews across storage subsystems and mentored new engineers onboarding to a complex codebase

**Research Assistant                      North Carolina State University                      Aug 2018 – June 2023**

**Cloud Gaming System Design:** Optimized cloud gaming delivery with segmented object streaming

- Designed a novel end-to-end solution enabling video frame compression by server-side object segmentation and client-side regeneration, reducing bandwidth requirement by 12%, while delivering 1080p/60fps streams (*WebRTC, Chromium, PyTorch, Detec-tron2, OpenCV, C++, Python3*)
- Automated data annotation pipeline using Segment Anything Model, enabling model training with average precision above 98% (*PyTorch, Python3, FFmpeg*)

**Cloud Gaming Measurement:** Dissected video streaming and latency performance of commercial cloud gaming platforms (Amazon Luna, NVIDIA GeForce Now, Google Stadia) across multiple game genres

- Developed a general, automated, deployable deep-learning-based measurement methodology featuring video processing and round trip video delay measurement, achieving model accuracy above 99% (*CNN, TensorFlow, Keras, FFmpeg, Python3*)
- Implemented a lightweight Chromium module for video stream data logging and offline analysis, identifying factors that degraded bitrate by 6.6x and frame-rate by 2x (*Chromium, WebRTC, FFmpeg, usbmon, C++, Python3*)

**Spam Political Biases:** Modeled and evaluated political biases in spam filtering algorithms of Gmail, Outlook, and Yahoo

- Built a tool to extract 1.3 million emails from 300 accounts (*IMAP, Selenium, Python3*)
- Applied propensity score matching with logistic regression and caliper matching to reduce confounding in data analysis (*R*)
- Press and Interviews: The New York Times, The Washington Post, Atlantic, Bloomberg, Fox News, Senate.gov press release
- **Policy Impact:** Political Bias Emails Act of 2022, Google Advisory Opinion at Federal Election Commission

**AWS Network Measurement:** Analyzed availability and latency of AWS network for elastic compute cloud and serverless

- Built a large-scale measurement system across 210 regions and 244 availability zone pairs, collecting 20 billion packets and identifying cloud outages and latency spikes in AWS infrastructure (*Serverless Lambda, Kubernetes, EC2, Python3, C*)

**Software Engineering Intern (Storage)                      Hewlett Packard Enterprise                      June 2022 – Aug 2022**

- Improved task execution performance by 3x through efficient job scheduling on multi-core processors (*NimOS, C++, Python3*)

**Software Engineer                      Netsol Technologies                      March 2013 – August 2015**

- Delivered a consumer lease solution for an Australian auto-financing client (*C#, SQL*)

---

**Education**

---

**North Carolina State University, Raleigh, North Carolina, USA                      Aug 2018 – June 2023**

- Ph.D. in Computer Science. Research Areas: Networking, Systems, Machine Learning

## Publications

---

- RFC 9944: Device Schema Extensions to the System for Cross-Domain Identity Management (SCIM) Model (2026) [RFC]
- Dissecting Cloud Gaming Performance with DECAF (*ACM SIGMETRICS 2022*) [DOI] [GitHub]
- Characterizing the Availability and Latency in AWS Network from the Perspective of Tenants (*IEEE TNET 2022*) [DOI]
- Left or Right: A Peek into the Political Biases in Email Spam Filtering Algorithms During US Election 2020 (*WWW 2022*) [DOI]
- RMS: Removing Barriers to Analyze the Availability and Surge Pricing of Ridesharing Services (*SIGCHI 2022*) [DOI]
- Interleaving Multiple IoT Stacks on a Single Radio (*ACM CoNEXT 2018*) [DOI]
- Taming Link-layer Heterogeneity in IoT (*ACM SenSys 2017*) [DOI]

## Technical Skills

---

- **Languages:** C++, Python, C, SQL
- **Storage Systems:** File systems (POSIX semantics, NFSv4 locking, IO range locks, metadata & namespace operations); Object storage (S3 versioning, and tagging APIs; MinIO); distributed caching; back-pressure management
- **Systems and Tools:** Linux, GDB, Address Sanitizer (ASAN), eBPF/BPF, netem, FFmpeg, WebRTC, Chromium, Git, GitHub, Jira, Confluence
- **Cloud and Infrastructure:** AWS (EC2, S3, Lambda), Kubernetes, Serverless
- **Frameworks and Libraries:** PyTorch, TensorFlow, Keras, NumPy, Detectron2, OpenCV
- **Domains:** Distributed Systems, File Systems, Object Storage, Networking, Performance Optimization, Reliability & Observability, Machine Learning, Cloud Computing