

Joel Tony

✉ joeltony72@gmail.com ☎ +91 81057 53322 [in](#) in/joel-tony [G](#) jay-tau [G](#) jaytau.com

Education

Birla Institute of Technology and Science Pilani, Goa Campus, 2021 – present | Goa, India
B.E. Computer Science
CGPA: 8.94/10
Relevant coursework: Distributed Computing (graduate-level), Compiler Construction, Computer Networks, Data Storage Technologies & Networks, Operating Systems, Computer Architecture, Principles of Programming Languages

Professional Experience

e6data, Performance and Research Engineering Intern [✉](#) Jun 2024 – present | Bangalore, India
- Developing a **spill-to-disk solution** to prevent query failures caused by **insufficient memory**
- Researching **efficient data representation** and **optimising disk access** to minimize the overhead of a data spill

Google Summer of Code, Contributor [✉](#) May 2024 – present | Bangalore, India
- Enhancing **Drishti** [✉](#), an open-source framework for **I/O insights**, to better support varied I/O access patterns
- Developing recommendations to optimize I/O insights for **AI/ML applications**

Coditation Systems, Software Engineering Intern May 2023 – Jul 2023 | Pune, India
- Developed a **GPU remote development suite** for graphics applications
- Configured and deployed **Bazel** to optimize build processes, resulting in a significant reduction in build times

Research Experience

SCI Institute, University of Utah, Undergraduate Researcher Jan 2024 – present | Utah, USA
- Working with Dr Manish Parashar [✉](#) on implementing a **POSIX-compliant API for DataSpaces** [✉](#)

Data, Systems & HPC Lab, Undergraduate Researcher Oct 2022 – present | Goa, India
- Deployed and monitored system metrics on a **40-node cluster**, using **Netdata, InfluxDB and Grafana**

Embedded Systems Lab, Undergraduate Researcher Aug 2023 – Dec 2023 | Goa, India
- Worked under Dr K.R. Anupama [✉](#) on **designing a network protocol stack** for a **distributed cloudlet cluster**

Central Electronics Engineering Research Institute (CSIR-CEERI), Undergraduate Researcher Jul 2022 – Aug 2022 | Pilani, India
- Collaborated with Dr. Dhiraj Sangwan [✉](#) on **parallelizing pseudo-colouring of X-ray baggage scans**

Publications

Does Varying BeeGFS Configuration Affect the I/O Performance of HPC Workloads?, *IEEE CLUSTER*, 2023 [✉](#) 2023
Arnav Borkar[†], **Joel Tony**[†], Hari Vamsi K.N., Tushar Barman, Yash Bhisikar, Sreenath T.M, Arnab K. Paul

Projects

Dynamic Striping Configuration for HPC Storage, Oct 2022 – present
BeeGFS, IOR, Distributed File System
- Working on **optimal file-volume mapping** in BeeGFS, considering factors like file access patterns and storage utilization

Data Prefetching for Edge Deep Learning Workloads, Aug 2023 – Dec 2023
Apache Kafka, PyTorch, Distributed Deep Learning
- Implementing a **prefetching** mechanism to **reduce I/O overhead** and reduce overall training time
- Using **Apache Kafka** to stream data from a **distributed file system** to compute nodes

TCP/IP Stack Using Rust & C, Aug 2023 – Dec 2023
Computer Networks, Rust, C [✉](#)
- Integrated **Rust with a C-based TCP/IP stack** and explored **conditional memory safety** guarantees and the performance
- Leveraged **perf** and the **Valgrind** suite to analyze the performance gap between the integrated Rust code and original code

OpenCSRankings, Pandas, XML [✉](#) Apr 2024 – present
- A metrics-based ranking of the top Computer Science institutes and faculty in India, based on publications
- Built a **streaming XML parser** to identify and aggregate publication statistics of researchers from DBLP

Proficiency

Languages

Python, C/C++, Rust, Bash, Haskell, LaTeX

Technologies

Linux, Kubernetes, Docker, Git, Kafka, Ansible

Awards

Amazon AWS AI/ML Scholarship, AWS, Intel Corporation & Udacity [✉](#) Nov 2022
One of the 1000 students selected internationally for the AWS AI/ML Scholarship in collaboration with Intel and Udacity.