

# ZHE(JOE) CAI

2678385750 | caizhe@seas.upenn.edu  
Philadelphia, PA 19104, United States

## EDUCATION

---

**University of Pennsylvania** **Philadelphia, PA**  
Master of Computer and Information Technology & Scientific Computing **August 2018 - May 2021 (Expected)**

- **Cumulative GPA:** 4.0/4.0
- **Courses:** Data Structures, Algorithms, Operating Systems, Distributed Systems, Internet and Web Systems.
- **Teaching Assistant** for CIT590, CIT596

**Fudan University (FDU)** **Shanghai, China**  
Bachelor of Material Chemistry **September 2014 - June 2018**

- **Cumulative GPA:** 3.48/4.0 (**Ranking:** 1/25), **Major GPA:** 3.66/4.0
- **Honors:** ExxonMobil Scholarship (5%, 2016), Outstanding Graduate at FDU (10%, 2018)

## SKILLS

---

**Programming Languages:** Java, C/C++, Python, JavaScript, HTML, CSS, SQL

**Frameworks and Libraries:** Hadoop, SparkJava, Spring MVC, Spring Boot, Hibernate, MySQL, Node.js

**Tools:** AWS, Git, Jupyter, Docker

## EXPERIENCE

---

**ZTE Corporation** **Shanghai, China**  
*Software Engineer Intern* **May 2019 - August 2019**

- Implemented a high-availability cluster with Pacemaker and Corosync and built a PaaS platform, which can be applied to network management systems. – [HA Cluster](#), [PaaS](#)
- Comprehended and modified Pacemaker and Corosync source code, parameters and configurations, which improved high-availability to more than 99.9%. – [C](#), [XML](#), [Socket Programming](#)
- Wrote some Shell and Python scripts to manage and schedule resources of Pacemaker. – [Shell](#), [Python](#)
- Used Black Duck to analyze the code and resolved more than 100 vulnerabilities. – [Black Duck](#)

## PROJECTS

---

**PennInSearch: A Distributed Web Crawler and Search Engine** – [Java](#), [Hadoop](#), [AWS](#), [Storm](#), [SparkJava](#)

- Implemented a Google-style search engine with four key components: Crawler, Indexer, PageRank, and Web User Interface.
- Accomplished a distributed Mercator-style web crawler that crawled more than 1 million URLs.
- Designed Hadoop MapReduce jobs to computing indexer and PageRank values, which was deployed on AWS EMR and EC2 instances.
- Optimized the ranking algorithm and built a web user interface with SparkJava.

**PennOS: A User-level UNIX-like Operating System** – [C](#), [UNIX](#), [Git](#)

- Implemented a user-level UNIX-like operating system consisted of a kernel system, a FAT file system, and a shell.
- Designed and built a kernel system including process controlling block, signals, and a priority scheduler.
- Developed a shell with various built-ins for users to interact with the PennOS.

**PennCloud: A Fault-tolerant Distributed Cloud Platform** – [C](#), [C++](#), [Linux](#), [Git](#), [HTML](#), [CSS](#)

- Implemented a scalable, fault-tolerant cloud platform including webmail and storage services similar to Gmail and Google Drive, which supports primary-based replication, checkpointing, and recovery.
- Built a distributed key-value storage to achieve data replication and consistency following the design of Google Bigtable.
- Customized LRU-based front-end load balancer and multi-threaded HTTP server which enables role-based functionalities.