

Tuan Anh Hoang-Vu

Green Card Holder | New York, USA | hvtuananh@gmail.com

PROFESSIONAL SUMMARY

Results-driven **Site Reliability Engineer (SRE) & Infrastructure Leader** with expertise in **large-scale observability, distributed systems, and automation**. Proven track record of reducing operational complexity, optimizing performance, and leading strategic migrations. Passionate about building scalable and resilient infrastructure.

WORK EXPERIENCES

Two Sigma Investments, LP.

Aug 2017 - Present

Vice President, Reliability Engineer

New York, NY

- Led the design and development of a companywide observability suite using Prometheus and Grafana on Kubernetes with persistent storage, serving over 120M active metrics. Reduced provision time from hours to under 15 minutes.
- Architected and maintained pipelines for collecting, filtering, and forwarding of observability data (logs, metrics, and traces) for seamless integration with vendor products.
- Spearheaded the migration from legacy metrics solutions used by hundreds of services to strategic platforms, reducing operational complexity and overhead.
- Championed the adoption of Ansible for configuration management of critical services.
- Provided open-source software as internal services, including ELK, Kafka, and Zookeeper.
- Designed and implemented a solution for running over 500 Logstash instances on Kubernetes.

Uber Technologies Inc.

May 2016 - Aug 2016

Software Engineering Intern

San Francisco, CA

- Investigated and created proof-of-concept strategies for Quality of Service (QoS) optimization at application and service levels.
- Developed a monitoring solution to collect key metrics from Mesos/Aurora and Docker services.
- Implemented a feature in Mesos v1 API to support querying by metric types (MESOS-5731).

New York University Tandon School of Engineering

Sep 2012 - Jan 2018

Research Assistant

Brooklyn, NY

- Developed an efficient spatio-temporal textual index that handles spatial, temporal and textual components uniformly, achieving high efficiency both in space usage and query evaluation.
- Developed *temporal* entity matching algorithms to integrate various urban datasets to discover underlying owners and ownership patterns.
- Collected news across top news websites in the US using *focused crawler* and associating news topics with tweets via domain adaptation.
- Configured and maintained a Hadoop cluster with over 1,000 cores.

Cazoodle Inc.

Feb 2009 - Jun 2012

Software Engineer

Hanoi, Vietnam

- Developed a comprehensive web crawling suite for automated data extraction.
- Designed and deployed a crawler scheduling system with 200+ clients that were responsible for on-demand and scheduled web crawling.
- Developed a lightweight parsing framework for real-time extraction of product information from web pages.

- Built a Firefox add-on (Shopping Caddy) for fetching offers, reviews, and price history.
- Extracted and aggregated opinions from users' reviews and mapped into products to enable searching for products by opinions.

TECHNICAL SKILLS

- Programming Languages: Python, Java, PHP
- Infrastructure & Observability: Docker, Kubernetes, Prometheus, Grafana, ELK Stack, Fluent Bit, Datadog, GCP
- Configuration Management & Automation: Ansible, Bash, Git
- Data Processing & Messaging: Kafka, Zookeeper
- Web & Backend Development: Flask, Nginx

EDUCATION

New York University Tandon School of Engineering

Ph.D. in Computer Science

Thesis: Spatio-Temporal Keyword Index For Interactive Data Exploration

2nd place in ACM SIGMOD 2014 Programming Contest

Jan 2018

Brooklyn, NY

Hanoi University of Science and Technology

B.S. in Computer Science

Jun 2010

Hanoi, Vietnam