

J C LAWRENCE
PHONE: (408) 410-3632
EMAIL: claw@kanga.nu
URL: <http://kanga.nu/~claw/>

J C Lawrence

- Builds cost-effective scalable systems.
- Customer-oriented and results-driven with 10+ years of engineering leadership.
- Full system and product life-cycle experience for both Software Engineering and Operations.

Matterport – Sunnyvale, CA

Senior Software Architect, October 2013 – April 2017

- Architected & built the processing pipeline making 3D models.
- Established & automated site management, configuration & deployment under Salt.
- Defined the analytics integration data models for Mixpanel & SegmentIO.
- Architected and built the data warehouse, ETL/data pipeline & initial reports/visualisations.
- Implemented microservices, APIs, reporting systems, primary product CLI, system monitors, etc.

Skills used: Airflow, AMQP, AWS (EC2, RDS, Redshift, S3, VPCs etc), build & release, Celery, configuration management, DevOps, distributed systems, Django REST framework, engineering & technical leadership, distributed systems, Python, queue-based/message-passing systems, RabbitMQ, Redis, REST services, Salt, StatsD, web services, WSGI.

CyDesign – Palo Alto, CA

Senior Staff Engineer, April 2012 – May 2013

- Architected & then built and lead the team that implemented a reliable multi-tenant asynchronous execution system (thousands of worker systems) for model simulations using a constraint-based system to manage task scheduling fairness and resource distribution.
- Designed and implemented operational statistics and billable event collection systems.
- Ported CyModelica (a Modelica compiler) from Windows to Linux and the Gnu toolchain, and then back to Windows; along with refactoring to support team development.

Skills used: Aspect-oriented programming, AWS/EC2/S3, build & release, compiler design and implementation, constraint-based systems, engineering & technical leadership, distributed systems, Java, Lua, Python, RabbitMQ, Redis, REST services, Riak, PaaS, scalability, schema-less/no-SQL databases, StatsD, web services, virtualization and VM-based clusters

YouSendIt – Campbell, CA

Senior Staff Engineer, Storage Architect Dec 2009 – Jan 2012

- Architected and implemented scaling from ~1.5 million transactions/day and 500TB of storage to 25+ million transactions/day and 4PB of storage.
- Led the reduction of transaction error rate by 80% and site escalations from 3/week to under 2/year.
- Architected & implemented standardized engineering deliverables, build environments, build processes, product packaging, configuration management, and standardized deployment methods, initially for the storage tier and ultimately for all server-side products.
- Built a strong relationship between Engineering and Operations at all levels.
- Assorted technical deliveries: 64bit port of the storage tier, specified & moved the storage tier from DASD to a shared horizontally-scalable storage cluster model, re-implemented the storage tier as light-weight WSGI web services instead of a costly custom Apache module, multiple large customer saves for Sales and Customer Support, etc.

Skills used: Apache, Apache modules, API design, AWS/S3, build & release, C/C++, concurrent programming, engineering lead, distributed systems, GPFS, incident response protocols/teams, gevent/libevent, Luwak, master-less clusters, PDLC and engineering process, product architect, Python, Riak, SaaS, scalability, schema-less/no-SQL databases, storage systems & networks, technical leadership, web services, WSGI, virtualization and VM-based clusters

PayPal – San Jose, CA

Lead Engineer, Manager, Product Architect June 2005 - Nov 2008

- Acted as lead engineer and manager for the Developer Services team, responsible for the PayPal web-services platform (\$30B+/year), IPN (API callbacks) and Sandbox test systems.
 - Re-implemented merchant APIs to use name-value pair POSTs instead of SOAP. Within 6 months two thirds of all new merchant integrations were using NVP.
 - In parallel with the payment-method development teams, rewrote the web-services platform for all payment and merchant APIs, moving it from a monolithic CGI-BIN/server pair to a light-weight dispatcher model routing requests to business-logic servers with good code & deployment isolation.
 - Led the rollout of the new web-services tier with no unplanned downtime or lost transaction volume.
 - Moved the Sandbox from a 4-system cluster in San Jose to a 100-system cluster in Denver.
- Acted as lead engineer and manager for the Merchant On-boarding team, responsible for merchant-acquisition and merchant-configuration systems. Built and trained the QA team for my group. With a local team and a team in India delivered over a dozen features including rewriting the boarding flow with a 70% reduction in fall-off rate, hierarchal sub-account management and channel boarding APIs.
- As Product Architect for the internal “TurboRoller” release-orchestration tool, developed and got funded a full web-services/work-flow based replacement architecture with a roadmap to get there from the current system.
- Part of the team that implemented STSI, a rescue re-architecture when internal or external dependencies fail, including the site’s Oracle database (which was near crashing from load). I was responsible for parts of the payment methods and the daemon that replayed transactions received during the outage.
- Recognized for exceptional ability to work through others and consistently moving the company forward to both fix the problem and ensure it didn’t recur.

Skills used: Apache, product architect, business coordination, C/C++, distributed systems, engineering lead, manager, offshore engineering coordination, security analysis, technical leadership, web services, workflow automation

Pfizer Research & Development – Groton, CT

Product Architect, June 2003 – May 2005

- Wrote and executed massively parallel custom applications to analyze –omics data (proteomics, genomics, etc).
- Using Open Source tools and components, designed and delivered an agent-based distributed heterogeneous computational router/framework for non-programmer computational statisticians to express common computational methods and non-statistician/non-programmer research scientists to use work-flow tools to assemble & use those methods for their research analysis.
- Collaborated on distributed computational statistics projects for Pfizer with research groups at Yale and Seattle universities.

Skills used: academic partnerships, architect, ad-hoc heterogeneous clusters, distributed computation, distributed systems, message busses, message passing networks, parallel computation, performance analysis, Python, scientific computing, tuple spaces, web services, workflow automation, virtualization, XML, Zope

Protego Networks – Milpitas, CA

Systems Architect, August 2002 – February 2003

- Designed and implemented logic core of product (aggregation and near-realtime distributed analysis of distributed network events, 20K events/second per node). Patents granted.

Primary skills used: algorithm design, architect, C++, distributed systems, engineering lead, high performance systems, Linux, micro-language design, parsers, scalability, security analysis, SSL, XML