

CV - Konstantin Nazarov

I'm currently based in London, UK. I have a talent visa that allows to work without restriction, including consulting.

Mar 2023 - Present: Engineering Manager @ =nil; Foundation

Joined the company to build a distributed BFT database. In this role, I'm fully responsible for the product, its technical design, infrastructure, user-facing tooling and security.

Since the time I joined, I've built a diverse team of engineers, and together we delivered the initial prototype version.

As this is a startup, in addition to management I do a fair amount of technical work as well.

Apr 2021 - Mar 2023 (2y): Infrastructure Engineer @ Meta

As part of the Core Data PE team, I was responsible for maintaining and developing MySQL orchestration solution that is used to control the MySQL fleet at Meta. This is the largest data storage cluster that I've seen and worked with in my career, with a fair deal of challenges related to differing use-cases, rolling maintenances, upgrades, topology changes and others.

The biggest impact that I've made in this role is helping with transition of database runtime environment from bare-metal to containers.

My other areas of work included database-related security and troubleshooting at scale.

May 2016 - Jan 2021 (5y): Engineering Manager, Director @ Tarantool Database

Tarantool is a hybrid in-memory and on-disk database, that is highly extensible and programmable. It's been around as an open-source project for more than 15 years at this point.

I initially joined to build a SaaS database-as-a-service offering that intended to monetize the product. After initial experiments didn't play out well, I pivoted to custom enterprise engineering, and this is where we gradually reached success over a few years.

I started building custom solutions for big telecoms and banks, mainly in the field of customer data processing (data caching, CDC, multi-tier storage, etc.), and after we realized that our solution has an edge over competition (programmability, stable engine) - I started hiring engineers and building teams around individual solutions that we shipped to the customers.

My personal engineering contributions: - designed and built initial version of our database clustering solution called “Cartridge”, that eventually became open-source and allows to run and scale the product both on bare-metal and in kubernetes - designed and built initial version of our enterprise “data grid” product, which ended up quite successful and was deployed to serve real production workloads at multiple large enterprise companies (and still makes money to this date)

Management-wise, my team has grown to 45 people, including engineering managers, product managers, marketing and technical writers.

Jul 2012 - May 2016 (4y): Engineering Manager @ Parallels

Parallels is a virtualization product, that allows to run Windows on a Mac.

I joined as an engineering manager with one explicit goal - to reduce the time it takes the product to go from the developer committing the code to getting the functional product “image” from 10 hours to 1 hour. This was a challenging problem, because virtualization products contain parts compiled for all popular operating systems, including drivers / guest tools. And often components require very specific versions of compilers and build environments. Having multiple parallel versions of the product means that older build environments need to be maintained exactly as they were originally set up.

Together with the team, we ended up building a system of reproducible packaging of OS images for Windows, MacOS and Linux based on Chef and Packer. This was then deployed to our custom distributed build mesh, that consisted of heavily tuned Jenkins and build actors on our own server virtualization platform.

To make the process of working with images easier, we’ve built a custom Vagrant provider for Parallels Desktop which was open-sourced and still used by the community.

As part of the infrastructure team, I also contributed a lot to how the product’s C++ codebase is structured, and to the team’s feature workflow.

Mar 2012 - Jul 2012 (3m): Software Engineer @ VK

Mail.ru was the biggest email provider in Russia at the time. During this short employment, I was working on integrating LDAP authentication to the email service, so that internal corporate users could log in to email with their corporate credentials.

I wrote an efficient async LDAP client implementation in plain C, that could authorize against a domain controller and wasted negligible amount of compute and memory resources. This required implementing a subset of the LDAP protocol from scratch, just enough to perform the authentication handshake.

After completion of this project, I joined Parallels as an engineering manager.

Nov 2009 - Mar 2012 (2.5y): Software Engineer @ Deutsche Bank

During this employment, I worked on two projects:

Arina - traditional equity trading system. In this project I mostly was responsible for maintenance and bug fixing, as it didn't need any major development at that time.

Auto Hedger - automated algorithmic trading. Here I was responsible for writing market connectivity adapters in C++ (mainly FIX-compatible).

Mar 2007 - Feb 2009 (2y): Systems Software Engineer @ Auriga

This was my first serious "corporate" job after the university. I joined the company after successfully completing their systems programming courses and passing interviews.

I was a part of the team that was developing a custom RTOS called LynxOS-178, which is primarily targeting avionics and satellites.

In this team, I worked on the OS kernel, mainly in bug fixing, debugging, testing, validating time and space guarantees, and certification.

Jan 2002 - Mar 2007 (5y): etc

Since the time I finished learning C++, I started offering my services to companies that needed simple financial or asset tracking software. This was most of the time project-based employment, where I would get a fixed amount of money for shipping a project and then move on to the next company.