

Dr. Konstantin Boudnik

Florida, USA

Email: cos@c-systems.me

Architect / Technology executive

Over 30 years of diversified experience in all phases of software development including design and architecture of distributed, client/server, network intensive, cloud and multi-tier solutions. Experienced with wide range of hardware and software development environments, including Apache Hadoop ecosystem, in-memory data processing, data modeling and modelops. Established history of innovative mindset, critical thinking, engineering and open-source software leadership. Author of the foundational framework used by all commercial distributions of Apache Hadoop.

CEO and founder of consulting firm C-Systems Inc working for over a decade with leading technology companies like Canonical, Pivotal and many others.

- Engineering and open-source leader
- BigData and FastData solutions
- LLM/RLM adoption for business
- www.linkedin.com/in/kboudnik/
- Expert in leveraging FOSS for business
- Contributor to major Apache BigData projects
- Technology adviser for a number of Silicon Valley companies

Technical Skills

Open-Source: Member of Apache Software Foundation; Mentor and Champion of 10+ projects

Technologies: Big-, Fast-data for clouds / CI & CD for transformation leverage, modelOps and AI applications in product development

Professional Experience

C-Systems, Inc. Chairman

2012-now

C-Systems is an international consulting enterprise working in the areas of BigData storage and compute platforms design and development. Our expertise helps to create high-quality architectural solutions and deliver software systems to address the needs of modern data processing. We help both IT- and traditional companies to move their businesses towards digital experience and data-driven decision making, LLM- and RLM-based AI systems and AI-workers platform solutions Among our clients are WANdisco, Inc, Gridgain Inc, Canonical, Humatron,ai, and many others.

WANdisco, Inc.

2019-2023

Executive Vice President R&D, Chief Architect

Lead corporate architecture team to establish sound software and solution architecture of the major cloud data migration platform provider. Created and championed the systematic architectural practices for high-performance on-premise to cloud and multi-cloud solution for live migration of large data volumes. Helped Product and Engineering orgs to reshape and adapt company's software portfolio delivery to the rapid market changes and challenges.

Lead the research team to expand company's proprietary consensus technology to new hyper-ledger applications in the financial industry.

A "top five" European bank.

2018-2019

Senior Executive Director in charge of data transformation

Headed an organization of over 1,200 people responsible for a cross-company data gathering, processing and analytics. Redesigned and gradually rebuild new and radically improved platform capable of both batch and real-time data collection and processing at the petabyte scale using commercial, open-source and internally developed technologies.

Created an organizational structure for the operation, support and on-going development process of the platform components to guarantee business continuity and high availability of business and mission critical systems.

EPAM, Inc.

2016-2018

CTO of Bigdata, Open-Source Fellow

Konstantin Boudnik

Played a key role in shaping up EPAM's BigData Practice' platform offerings, advising clients on the technology and industry directions. Created the architectural blueprints for high-performance on-premise and cloud-based processing of large data volumes across different industries.

Building the Open-Source Enterprise Transformation Competency. helping customers' organizations to facilitate the adoption of open software and community development principles.

Memcore.io, Palo Alto, CA

2015-2016

Co-founder, CEO

Envisioned and implemented a cloud-based in-memory fast data ACID computing platform with seamless deployment, advanced cluster management features, and transparent ETL pipelines. Providing highly competitive TCOs compared to Amazon EMR and other leading brands.

WANdisco, San Ramon, CA

2013-2015

Vice President, Open Source Development; Member of ASF

Director, Advanced Technologies; VP of Engineering; VP of Apache Bigtop

Defining the company's strategy and open-source portfolio; increasing the business value by establishing thought leadership and engagement with open-source communities

- Supervising and cultivating company's Open Source portfolio, including NonStop Hadoop, Subversion, HBase, and the next generation Exabyte-scale distributed file system
- Actively participating and closely working with the Chief Product Officer on business and strategic decision making, focusing on on adopting and collaboration with Free and Open-Source Software projects. Lead company's adaptation to the ever-changing technological landscape of Bigdata
- Responsible for creation and delivery of one of three world's commercial BigData distributions based on Apache Hadoop 2.0, beating the IBM and Intel to market
- Managing geographically dispersed engineering team of 25+ people, guaranteeing products delivery on tight schedule
- Member of Technical Steering Committee for Open Data Platform Initiative (ODPi)
- Representing the company at domestic and international conferences

Karmasphere, Cupertino, CA

2011-2013

Senior Architect, Apache Hadoop committer, Apache Bigtop PMC

Designing, implementing, and improving BI/BA platform on top of Apache Hadoop, Apache Hive, Amazon EMR, MapR stacks

- Mentored and trained the engineering team in use of Hadoop platform and analytical software, development technologies
- Designed and lead the development of next generation web-based system for business analytics in Hadoop ecosystem

Yahoo!, Sunnyvale, CA

2009-2011

Apache Hadoop Committer, Apache Bigtop co-founder

Designed and implemented components of Hadoop stack and frameworks for creation and of Yahoo! data processing cloud stack. Provided technical expertise and integration guidelines for the development and quality teams to reduce the time of Hadoop development and deployment

- Developed new technologies and frameworks for code and fault injection into all levels of Hadoop clusters used for behavioral simulation and API extension for system controllability and observability, and the architecture of cluster commissioning and resource scheduling system
- Designed and implemented solutions for Hadoop on-demand deployment and validation for all components of data analytics stack, later known as Apache Bigtop

Sun Microsystems, Santa Clara, CA

1994-2009

Played a variety of engineering and architectural roles in Solaris, Java, and enterprise organizations. Lead the acceptance of Java technology into Linux Standard Base. In 2000 authored a heterogeneous distributed system used across the company to drive and scale out complex parallel processing and compute needs. Further references are available upon request

Education

- PhD., Computer Science, Saint-Petersburg University, Russia.
Thesis "LAN & WAN Management with Advanced SNMP Techniques"
- MS, Computer Science, Saint-Petersburg University, Russia.
Diploma thesis "The Generative System For Object Oriented Design"

Konstantin Boudnik

Publications / Patents

Books

- Professional Hadoop, ISBN-13: 978-1119267171

Distributed systems / Computer farms / Software Methodologies

- US 7,426,729, US 7,266,822, US 7,243,137, US 7,231,437, US 7,207,040, US 7,191,362, US 7,167,894, US 7,165,256, US 7,130,88, US 7,114,159, US 6,961,937, US 6,842,891, US 8,756,595, US 9,268,589. US 9,612,812
- US 7,257,613, US 7,020,699

Various articles, presentations, video interview are available through Internet and YouTube search.