

Dr. Konstantin Boudnik

East Bay, CA

Email: jobs@boudnik.org

Architect / Engineering Executive

Over 20 years of progressive and diversified experience in all phases of software development including design and architecture of distributed, client/server, network intensive, and multi-tier systems. Experienced with wide range of hardware and software development environments, including Apache Hadoop ecosystem and in-memory data processing. Established history of innovative mindset, critical thinking, engineering and open-source software leadership. Author of the framework used by all commercial distributions of Apache Hadoop.

- Engineering and open-source leader
- Expert in leveraging FOSS for business
- BigData and FastData solutions
- Contributor to major BigData projects
- Conference speaker. Book author
- Great communication & presentation skills
- www.linkedin.com/in/kboudnik/
- 100% hands-on

Technical Skills

Open-Source: Member of Apache Software Foundation; Incubator Projects Mentor

Languages: Groovy / Java / MOP / AOP / Git

OS/systems: Big- & Fast-data platforms / *nix / Puppet / Cloud technologies

Professional Experience

EPAM, Inc.

2016-now

Chief Technologist, Open-Source Fellow

Shaping up EPAM's BigData Practice' technology platform offerings, advising clients on the technology and industry directions. Creating 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

Cloud-based in-memory fast data ACID computing 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 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

Konstantin Boudnik

- 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. Injection then is used for Hadoop behavioral simulation and API extension for system controllability and observability
- Shaped requirements and functionality scopes, and the architecture of cluster commissioning and resource scheduling system
- Designing and implementing 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"

Publications / Patents

Books

- Professional Hadoop, ISBN-13: 978-1119267171

Distributed systems / Computer farms

- 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

Software Methodologies

- US 7,257,613, US 7,020,699

Online articles and interviews:

- NonStop Hbase <https://www.youtube.com/watch?v=IT5CDvJQAqA>
- NonStop HDFS <https://www.youtube.com/watch?v=eHRwTetaR74>
- <https://is.gd/W0wZxO> "30+ time faster than Apache Hadoop"