

Big Data Applications Performance Assurance

Boris Zibitsker
BEZNext

Abstract

Today's fast-paced businesses have to make business decisions in real-time. That creates pressure on IT leaders to develop near real-time Big Data and Data Warehouse applications that apply advance analytics against large volumes of data to deliver recommendations fast. Hardware and software used to build Big Data infrastructure is cheap, but management of complex environments is not easy. In this presentation we will review role of Performance Assurance incorporating Descriptive, Diagnostic, Predictive, Prescriptive and Control Analytics during each phase of the Application and Data life cycle. We will review challenges and Performance Assurance solutions for Big Data Batch and Real Time applications based on YARN, Map/Reduce, Kafka, Spark/Storm and Cassandra Apache projects

Bio

Founder and CEO of BEZNext, 2011 - present

- Current focus of research is on applying predictive and prescriptive analytics for optimization of business and IT decisions during applications and data life cycle
- Manage development of the Performance Assurance technology incorporating advanced analytics for optimization of Big Data and Data Warehouse applications in complex multi-tier, distributed, virtualized, parallel processing environment
- Consulted many of Fortune 500 companies

CTO of Modeling and Optimization at Compuware (2010-2014)

- Participated in development of Application Performance Management software incorporating Machine Learning algorithms for performance and availability problems

detection, and root cause analysis determination for web applications

Co-Founder of Computer Systems Institute (1989)

Founder, President and Chairman of BEZ Systems (1983 - 2010), acquired by Compuware in 2010

- Managed development of BEZVision Performance Prediction and Capacity Management software for Teradata, Oracle, DB2 and SQL Servers

Performance Analyst:

- Started out as engineer at Computer Systems Research Institute working on modeling and performance evaluation of large computer systems and applying modeling results for optimization of jobs scheduling and storage performance management
- Worked in capacity management departments at FNBC and CNA Insurance company in Chicago

Adjunct Associate Professor, DePaul University in Chicago (1983 - 1990)

- Taught graduate courses on Modeling of Computer Systems, Queueing Theory with Computer Applications, Computer Communication Systems Design and Analysis
- Taught seminars at Northwestern University, University of Chicago and Relational Institute - North and South America, Europe, Asia, and Africa
- Author of papers on applying modeling and optimization for performance evaluation, performance assurance, performance management, workload management and capacity planning for Big Data and Data Warehouse environments

Education: MS and PhD research at BSUIR and NIIEMV

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author(s).
ICPE'16 Companion, March 12–18, 2016, Delft, The Netherlands.

Copyright is held by the owner/author(s).
ACM 978-1-4503-4147-9/16/03.

DOI: <http://dx.doi.org/10.1145/2859889.2883586>