

# Venkat Kaushik, Ph. D

---

120 Stonemont Dr., Irmo, SC, 29063 | 803-609-8212 | [higgsmass@gmail.com](mailto:higgsmass@gmail.com) ,  
<http://boson.physics.sc.edu/~venkat>

## PROFILE

- Technology focused and passionate problem solver with strong technical expertise; using creativity in delivering automation, orchestration and operational stability for services in a hybrid cloud environment.
- Strong suit of teasing out performance from traditional and post-relational databases at extreme scale. Maniacal execution of well defined plans in production environments to satisfy stringent service level agreements.
- Multi-disciplinary scientist with collaborative research and data analysis experience at premier high energy physics laboratories and a penchant for teaching undergraduate physics.
- Adept in leading and playing key role in a cross-functional collaboration of development and operations with effective communication and interpersonal skills and a strong work ethic.

## EDUCATION

- **University of Texas at Arlington, USA**  
*Doctor of Philosophy, Applied Physics* Dec 2007  
*Master of Science, Computer Science and Engineering* May 2002
- **Bangalore University, India**  
*Bachelor of Engineering, Mech. Engineering (1<sup>st</sup> class/distinction)* Dec 1998

## PROFESSIONAL EXPERIENCE

### **SENIOR DBA | RUBICON CLOUD SERVICES, EMC<sup>2</sup> CORP. | JUL 2014-**

- Manage daily operations spanning multiple datacenters in two continents with 200 TB (and growing) metadata (PostgreSQL), several GB of key-value cache (Redis), configuration data store (MungoDB) and status/logs data store (CouchDB) for public cloud operations serving fortune 100 clients.
  - Automation includes the following subtasks (i) setting up streaming replication for master/slave architecture (ii) regular backups (both base backup and WAL archiving) off premises to facilitate disaster recovery (iii) setting up connection pooling for clients (iv) deploying templates for monitoring agents to monitor infrastructure, internal database health and resource usage and hook up to pagerduty.com (for push notifications of critical, major and minor alerts sent to operations personnel).

*Continued*

- Smart restores of all servers on a regular round-robin basis for auditing, setting up fire drills to go over failover scenarios, documenting and fine tuning the standard operating procedures for such events and training other DBA's.
- Load balancing of database cluster, identifying anti-patterns and performance bottlenecks in existing schema/configuration based on usage patterns and metrics collected. Analyzing the data and coming up with optimal configuration and query regression results that are fed back to development QA teams.
- Currently leading a new effort (object metadata using Cassandra) in designing, cloud-based deployment of massively scalable storage service offering.

### **ADJUNCT LECTURER | UNIVERSITY OF SOUTH CAROLINA | AUG 2013-**

- I enjoy teaching Phys211, Phys212 courses as a part of the evening program.

### **SOFTWARE ENGINEER | UNITRENDS, INC. | MAY 2013- JUN 2014**

#### Database Performance Characterization

- Performance characterization of PostgreSQL database in Unitrends product. Smartvac –a solution to alleviate problems with existing database on CentOS5 OS for support/customers. Migration strategy for existing physical/virtual appliances to CentOS6 OS

#### Project Lead: New OS/DB Platform

- Migration (OS/DB) requirements and must-have features / functionality. Functional unit testing, integration testing time/effort was identified and resources allocated. Cross platform (CentOS 5 → CentOS 6) testing

#### Query Optimization

- Configurable and tunable parameters identified for physical/virtual appliances. Identified high frequency/ long duration queries in various subsystems of the product. Execution plans obtained and categorized based on requirements. Scalability of queries was tested and optimized queries were included. Closure tests were developed and used to verify the results were identical.

### **POST-DOCTORAL RESEARCH ASSOCIATE | U. ARIZONA | JAN 2008- JUN 2013**

#### Collaborator/Author: ATLAS / RD51 Experiments at CERN (Geneva, Switzerland)

- Several publications in refereed journals as a primary author: Publication list
- Developed data analysis framework designed for a challenging use-case of petabyte-dataset using grid-computing architecture. Strategy is slimming and skimming the dataset using multi-process parallelism based on user analysis requirements, storage of resulting smaller dataset on local (tier3) cluster and local batch processing. Performed a multivariate analysis of the resulting data for a measurement that lead to 4 publications. *Ref: AZ Framework.*

- Multivariate data analysis to separate signal/background; extract signal significance by solving a p-value problem using modified frequentist approach.
  - Sound knowledge of data-modeling and characterization using **Unfolding** techniques (Bayesian, SVD, Matrix Inversion), **MLE** (Neural Networks, Boosted Decision Trees and Likelihood estimators) with implementation in multiple high level languages
- Co-authored a monitoring tool (requirements, design, coding and testing) for test-beam data. Prototype particle detector chamber with 64 readout channels was being tested for the super-LHC upgrade (R&D project): Micromegas.
  - Real time monitoring of data from the data-acquisition (DAQ) system – strip headers and packetize and transmit data (UDP protocol) to monitoring tool.
  - Calibration and preparing data (byte stream to high level objects) for analysis
  - Reconstruction of low level objects (hits) to high level objects (particle tracks)
- Data quality monitoring (DQM) expert. I am responsible for trigger and data acquisition system monitoring three key components of detector data-taking effort (Muon CSC, Jet and MET trigger signatures).
  - Developed tools to automate and add intelligence to data-quality assessment – significant contributions include statistical tests of outliers, shapes (Gibbs, Kolmogorov-Smirnov, defining metrics for data-taking efficiency, RGY color coding). *Ref: ATLAS Data Quality.*

### **RESEARCH ASSISTANT | UTA/FERMILAB | MAY 2004 – AUG 2007**

Collaborator/Author: D0 Experiment at FERMILAB (Batavia, IL, U.S.A)

- **Dissertation: Search for a Higgs boson** with data collected using D0@Tevatron. Used Advanced Neural Network estimator to search for the signal. Featured article. Doctoral thesis.

### **TEACHING ASSISTANT | U. T. ARLINGTON | MAY 2002 – MAY 2004**

Masters Thesis: Simulation study of a novel detector concept *Ref: GEM-HCAL Detector*

- Studied performance metrics in digital / analog readout mode. **Developed Energy Flow Algorithm for categorizing charged and neutral particles** using energy weighted and density weighted approach.

### **LEADERSHIP**

- Liaison of a major sub-group of ATLAS experiment. I served as a Top/Missing Transverse Energy (MET) liaison. *Ref: METLiaison*
  - My responsibility included organizing/supervising individuals in the liaison task, arranging meetings to discuss the physics requirements of top group, covering various topics of reconstruction algorithms and performance metrics including energy resolution, linearity of response among others.

- Identify topics for study of performance and ways to improve the MET reconstruction, assign students/post-doctoral candidates to topics, follow-up with the results of simulation studies and compare them to real (detector) data, suggest ways to improve performance of MET.
- Mentoring doctoral candidates and undergraduate students
  - I helped two undergraduate students with their projects - over six month period for each candidate. Upon their request, I wrote reference letters for their graduate school applications.
  - I mentored three graduate students with their dissertation topics. Trained them for their service work – a qualification task toward authorship in ATLAS publications. The task was maintaining and improving the data quality software, which is one of the responsibilities of Arizona physics group.

### **INVITED TALKS**

- Invited speaker at Fermilab Wine and Cheese Seminar (Nov 2012): Top Physics Results
- Invited talk at Phenomenology Symposium (May 2012): Top Quark Pair Properties with ATLAS

### **HONORS / SCHOLARSHIPS**

- Scharff Award for Excellence in Research, U. T. Arlington (2006)
- Honorable mention, ACM International Collegiate Programming Contest, South Central Region, USA (2001).
- Teaching/ scholarships during 1999 – 2002 (Physics Department, U. T. Arlington)

### **TECHNICAL SKILLS AND CERTIFICATIONS**

- **Language Fluency:** C, C++, Java, JSON, XML, Python, Perl, Fortran
- **Configuration Management:** SaltStack, CF-Engine, Docker, CMT
- **Certifications:** Certified Cassandra DBA
- **Cloud Platform:** Openstack, Terraform, Mesos, VMWare suite
- **Scripting:** Python, advanced bash scripting, JavaScript, Perl, PHP
- **Databases:** PostgreSQL, Cassandra, CouchDB, MongoDB, Redis, Oracle 8i, MySQL
- **Operating Systems:** Windows 8/7/NT/XP, MacOSX, Unix (SLC, CentOS, Ubuntu)
- **Code Management:** git, svn, cvs, SoftRelTools, Repo manager.
- **Scientific/Statistical Tools:** **ROOT TMVA**, R, GSL, RooFit, RooStats, BAT, GEANT4

### **PROFESSIONAL REFERENCES**

- Available upon request