

---

## Sudharshan S. Vazhkudai

email: [vazhkudaiss@ornl.gov](mailto:vazhkudaiss@ornl.gov)

Office: 865-576-5547

Fax: 865-576-5491

web: <http://www.csm.ornl.gov/~vazhkuda>

Oak Ridge National Laboratory

One Bethel Valley Road

PO BOX 2008 MS-6016

Oak Ridge, TN 37831

---

### CURRENT POSITION

- **TITLE/AFFILIATION:** *Research Staff Member*, Computer Science and Mathematics Division, Oak Ridge National Laboratory
- **OCCUPATION:** *Distributed Systems Research*

### PROFILE

- A Doctorate in Computer Science with specialization in massively distributed Grid systems
- Research experience in Grids, distributed storage caches, distributed operating systems, distributed file systems, clusters, schedulers, etc.
- Obtained over a million dollars in research funding as PI and co-PI on projects
- Key contributor of a core team building distributed computing/data management solution for US DOE's Spallation Neutron Source, a billion dollar infrastructure
- Currently leading an effort to address data availability and I/O bandwidth bottleneck for petascale storage subsystem; Co-leading the FreeLoader distributed storage project; Previously, led/developed a distributed OS for the Linux kernel
- Research contributions to premier open source projects such as Globus, Linux, etc.
- Publications in several peer reviewed conferences and journals in the area of distributed supercomputing
- **Interests:** Distributed Systems, File Systems, Data distribution systems, Peer-to-Peer Systems, Performance Analysis, Content Delivery Networks, Operating Systems, Middleware Systems

### RESEARCH SUMMARY

My interests are primarily in the area of Distributed Resource Management. To this end, I am currently involved in: constructing a distributed computing solution for the neutron science community (<http://www.sns.gov>); bringing massive neutron source data to the TeraGrid; developing distributed storage solutions (<http://www.csm.ornl.gov/~vazhkuda/Storage.html>): desktop storage aggregation, addressing data availability and the I/O bandwidth bottleneck for the petaflop storage subsystem. My doctoral work was on Globus Data Grid development—built techniques to locate data in highly replicated environments; developed performance prediction strategies for bulk Data Grid transfers aiding data transfer schedules; developed co-allocation middleware enabling parallel and fast access to replicated data, etc. Prior to this, I led the design and development of a performance oriented Linux Cluster—with kernel level modifications—to achieve enhanced throughput.

### EDUCATIONAL EXPERIENCE

- **Doctor of Philosophy** in Computer Science, May 2003  
*Institution:* University of Mississippi/Argonne National Laboratory  
*Dissertation:* *Bulk Data Transfer Forecasts and the Implications to Grid Scheduling*  
*Highlight:* *Three years of Dissertation fellowship from Argonne National Laboratory (US Department of Energy Lab); Thesis work on the R&D award winning Globus wide-area distributed systems toolkit; Research collaborations with groups/selective coursework at The University of Chicago and Northwestern University*  
*Committee:* *Dr. Tobin Maginnis, Dr. Conrad Cunningham & Dr. Dawn Wilkins*

- **Master of Science** in Computer Science, December 1998  
*Institution:* University of Mississippi  
*Thesis:* Performance Oriented Distributed OS-Evolutionary Steps towards a Distributed Linux  
*Highlight:* Led a team of programmers and graduate students in the design & development of a Distributed Linux; Software development/testing experience on a commercial million-loc product  
*Advisor:* Dr. Tobin Maginnis
- **Bachelor of Engineering** in Computer Science, June 1996  
*Institution:* Karnatak University, India  
*Project:* Yash - Yet Another Shell, An Enhanced C-Shell for the UNIX Environment  
*Highlight:* Ranked among top 5% of the class.

## WORK EXPERIENCE

- **OAK RIDGE NATIONAL LABORATORY, Oak Ridge, TN**  
 Oak Ridge National Laboratory is the largest U.S. Department of Energy facility. The Computer Science and Mathematics Division (CSMD) is ORNL's premier source of basic and applied research in distributed systems, high-performance supercomputing and applied mathematics.  
*Research Staff Member, Network and Cluster Computing, CSMD, [August 2003 - ]*
  - Working on the distributed computing/data management infrastructure for the **Spallation Neutron Source (SNS)**:  
 (along with researchers in CSMD and SNS Analysis Software Team)
    - Developed a requirement specification of the distributed computing/storage needs of the SNS community
    - Evaluation of several software suites for the Neutron Source Software Initiative (NeSSI) in the above context
    - Constructing portal technologies for remote access to SNS data and computation
    - Building data management strategies for users to operate on remote workspaces
    - Building application management strategies for remote executions
    - Developed secure authentication/authorization layers for accessing remote resources
  - **FreeLoader:** An aggregated, distributed storage infrastructure using loosely coupled, commodity user desktop workstations for use by data-intensive applications (Leading the effort along with Dr. Xiasong Ma and Dr. Vincent Freeh from NCSU)
    - Building a layered architecture of benefactors, steered by a management layer with corresponding APIs and services
    - Benefactors: UDP-based communication library for metadata and status exchange, TCP-based data-morsel transfer layer, status monitor
    - Management Layer: Reservations of space, metadata management, data striping for benefactor load balancing
    - Client Layer: Service APIs to reserve, store and retrieve data, double buffering strategy, overlapping network and disk I/O, equivalent file interface APIs for legacy application use
    - Built a testbed at ORNL with donated desktop storage, GPFS and HPSS storage access for comparison
    - URL: <http://www.csm.ornl.gov/~vazhkuda/Morsels>
  - **Storage Virtualization:** Leading an integrated approach to machine-room storage management
    - Building techniques and abstractions to exploit opportunities available in the HEC I/O stack to improve data availability and I/O bandwidth for standard supercomputing operations (staging, checkpointing, offloading, etc.).
    - URL: <http://www.csm.ornl.gov/~vazhkuda/IO-Virtualization>
- **ARGONNE NATIONAL LABORATORY, Argonne, IL**  
 Argonne National Laboratory is a U.S. Department of Energy research center specializing in Grid Computing. The Globus project, being developed at ANL, is evolving as the de facto standard in

Grid Computing, winning the R&D 100 award and being widely used in major production systems (European Data Grid, Earth Systems Grid, Particle Physics Data Grid, etc.).

**Research Associate & Wallace Givens Research Fellow, Globus Toolkit Project, [May 2000 - December 2002]**

- Worked on the orchestration of bulk data transfers in the Globus Data Grid:
  - Developed a decentralized, scalable scheduling architecture and selection heuristics for locating widely replicated data
  - Instrumentation and delivery mechanism (for wide-area consumption) of performance data for the de facto Grid data movement protocol, GridFTP
  - Built statistical performance prediction models of data transfer times
  - Built strategies for the efficient discovery of information
  - Developed and analyzed architecture/techniques for co-allocating transfers (parallel downloads, etc.)
  - Work resulted in middleware for scheduling, predicting and co-allocating bulk Grid data transfers
- Summary of Experience:
  - An understanding of Grid community requirements in the context of distributed data intensive science
  - Exposure to several state-of-the-art Grid related software systems (such as: Globus, NWS, ClassAds, GridFTP, MDS, LDAP, etc.)
  - Exposure to techniques and tradeoffs in building and analyzing disparate grid systems
  - An understanding of statistical tools in the context of systems performance analysis, predictions (regression analysis), etc.
- Research papers available from Globus Web site
- **Platforms and Tools:** C, LDAP, ClassAds, NWS, Linux

- **THE UNIVERSITY OF MISSISSIPPI, Oxford, MS**

- **Project Leader & Researcher, PODOS Project, [January 1998 - May 2000]**

- Lead architect for the design and development of a Performance Oriented Distributed Linux project
  - Managed a team of five programmers, several graduate and undergraduate student researchers
  - Built a High-speed Communication Protocol for a Cluster of Linux Machines involving multiple channels and short-circuiting protocol stack
  - Built a Networked File System exploiting Linux VFS layers and prefetching
  - Designed and guided the development of Global IPC mechanism and group communication
  - Developed a remote process execution environment
- Summary of Experience:
  - In-depth understanding of Distributed OS concepts, intricacies in Kernel Programming and Debugging.
  - Code-level understanding of Linux (i)Datalink Layer (Ethernet Driver), (ii)VFS Layer, (iii)Process Management & Scheduling
  - Performance analysis of communication paradigms and networked file systems
- Research papers and Distributed Linux Patches for 2.0 kernel
- **Platforms and Tools:** C, Assembly, Kernel Programming

- **Researcher, Frameworks for Distributed Protocols Project, [August 1999 - May 2000]**

- Explored the construction of a framework for distributed decision making protocols (leader election, distributed scheduling, etc.)
  - Applied design patterns such as layered architecture, delegation and templates in constructing a white-box framework
  - Developed an initial abstraction of vertical and horizontal views comprising of initiators/listeners and application, protocol-specific and platform-dependent layers respectively
- Summary of Experience:

- Experience in applying design patterns and frameworks to software development
    - Hot-spot analysis
  - Research papers and an initial prototype framework in C++
  - **Platforms and Tools:** C++, White-box framework, Patterns
- **Instructor, Department of Computer Science, [August 1998 - May 2000]**
  - Developed and taught courses for undergraduate and graduate students
  - Survey of Computing, Computer Organization, Programming for Scientists—FORTRAN
- **Research Assistant, National Classroom Project, [August 1997 - December 1997]**
  - The project, supported by Microsoft Corporation, is committed to bringing computing into classrooms.
  - Worked on the development of an in-class portable lab that provides learning enhancement services to students and assessment services to teachers
  - Developed a Web based dynamic testing/teaching Environment for C++, WWW Remote Compilation environment, etc., for portable hand-held devices
  - Developed numerous libraries for the portable C++ compiler ("eC") for hand-held devices
  - **Platforms and Tools:** C++, Networking
- **SAIR LINUX GNU CERTIFICATION, Oxford, MS**

Sair Linux GNU Certification specializes in administering Linux Systems candidacy exams world wide.

*Senior Technical Consultant (concurrent appointment), [May 1999 - June 2001]*

  - Designed and developed certification exams (along with a core team) for testing practical Linux systems experience
  - Developed training manuals, tests, lab exercises in advanced Linux system, network and security administration
  - **Platforms and Tools:** Scripting, Perl
- **ENVIRONMENTAL SYSTEMS RESEARCH INSTITUTE (ESRI Inc.), Redlands, CA**

ESRI was founded in 1969 as a research group devoted to improving methods of handling geographically referenced data.

*Software Engineer - Testing, ARC/INFO Project, [May 1998 - August 1998]*

  - "Purified" (staunching leaks) ARC/INFO, a leading Geographic Information Systems (GIS) product comprising of legacy code written in FORTRAN, C and C++
  - Performed bug fixing for ARC/INFO
  - Ported ARC/INFO libraries to various architectures
  - **Platforms and Tools:** C, C++, FORTRAN, Debugging Tools
- **WIPRO INFOTECH, Bangalore, India**

Wipro Infotech is a leading IT Company in India working on diverse software issues.

*Design Engineer R&D, Printer Drivers Project, [August 1996 - August 1997]*

  - Designed and developed Windows Printer Drivers for Epson based printers along with a core development team
  - Developed several driver functions invoked by the GDI (device specifics, modes, copy data, etc.), color calibration schemes and intensity and brightness control
  - Modified Unidrv driver development kit to include Epson specific commands and utilities, with support from firmware team
  - **Platforms and Tools:** Visual C++, SDK, WINDOWS

## **PUBLICATIONS, PRESENTATIONS and GRANTS**

### **Journal**

1. V.W. Freeh, X. Ma, S. Vazhkudai, J. Strickland, "Controlling Impact While Aggressively Scavenging Idle Resources"  
*To appear in the Journal of Cluster Computing*, Oct 2006.
2. S. Vazhkudai, X. Ma, V. Freeh, J. Strickland, N. Tammineedi, T.A. Simon, S.L. Scott, "Constructing Collaborative Desktop Storage Caches for Large Scientific Datasets" in the *ACM Transactions on Storage (TOS)*, Vol. 2, No. 3, pp. 221-254, August 2006.
3. J.W. Cobb, A. Geist, J.A. Kohl, S.D. Miller, P.F. Peterson, G.G. Pike, M.A. Reuter, T. Swain, S.S. Vazhkudai, N.N. Vijayakumar, "The Neutron Science TeraGrid Gateway, a TeraGrid Science Gateway to Support the Spallation Neutron Source", *To appear in the Journal of Concurrency: Practice and Experience*, 2006.
4. S. Vazhkudai, "Distributed Downloads of Bulk, Replicated Grid Data", in the *International Journal of Grid Computing*, Vol. 2, pp. 31-42, 2004.
5. S. Vazhkudai, J. Schopf, "Using Regression Techniques to Predict Large Data Transfers", in the *International Journal of High Performance Computing Applications - Special Issue on Grid Computing: Infrastructure and Applications*, Volume 17, No. 3, pp. 249-268, Fall 2003.
6. S. Vazhkudai, J.M. Syed, P.T. Maginnis, "PODOS - The Design and Implementation of a Performance Oriented Linux Cluster", in the *Journal of Future Generation Computer Systems - Special Issue on Cluster Computing*, Volume 18, Issue 3, pp. 335-352, January 2002.

### **Conferences/Workshops**

1. S. Vazhkudai, D. Thain, X. Ma V. Freeh, "Positioning Dynamic Storage Caches for Transient Data" in *Proceedings of the International Workshop on High Performance I/O Techniques and Deployment of Very Large Scale I/O Systems (HiperIO'06)*, Barcelona, Spain, September 2006.
2. X. Ma, S. Vazhkudai, V. Freeh, T.A. Simon, T. Yang, S.L. Scott, "Coupling Prefix Caching and Collective Downloads for Remote Data Access" in the *Proceedings of the International Conference on Supercomputing (ICS06)*, pp. 229-238, Cairns, Australia, June 2006.
3. J. Cobb, S. Miller, G. Pike, S. Vazhkudai, M. Hagen, "Neutron Science TeraGrid Gateway", *Proceedings of the TeraGrid'06 Advancing Scientific Discovery*, Indianapolis, IN, June 2006.
4. S. Vazhkudai, X. Ma, V. Freeh, J. Strickland, N. Tammineedi, S.L. Scott, "FreeLoader: Scavenging Desktop Storage Resources for Scientific Data", *Proceedings of Supercomputing 2005 (SC'05): Int'l Conference on High Performance Computing, Networking and Storage*, Seattle, Washington, November 2005.
5. J.W. Cobb and S. Vazhkudai, "Science Gateway Case Study: The Neutron Science TeraGrid Gateway (NSTG)", *Proceedings of the Workshop on Science Gateways: Common Community Interfaces to Grid Resources (GGF14)*, Chicago, IL, June 2005.
6. J. Strickland, V. Freeh, X. Ma and S. Vazhkudai, "Governor: Autonomic Throttling for Aggressive Idle Resource Scavenging", *Proceedings of the 2nd IEEE International Conference on Autonomic Computing (ICAC 2005)*, pp. 64-74, Seattle, WA, June 2005.
7. S. Vazhkudai, "On-demand Grid Storage using Scavenging", *Proceedings of the Session on New Trends in Distributed Data Access*, Las Vegas, Nevada, June 2004.
8. S. Vazhkudai, "Enabling the Co-Allocation of Grid Data Transfers", *Proceedings of the 4th International Workshop on Grid Computing (GRID 2003)*, pp. 44-51, Phoenix, Arizona, November 2003.
9. S. Vazhkudai, J. Schopf, "Using Disk Throughput data in Predictions of End-to-End Grid Transfers", *Proceedings of the 3<sup>rd</sup> International Workshop on Grid Computing (GRID 2002)*, pp. 291-304, Baltimore, Maryland, November 2002.
10. S. Vazhkudai, J. Schopf, "Predicting Sporadic Grid Data Transfers", *Proceedings of the 11<sup>th</sup> IEEE International Symposium on High Performance Distributed Computing (HPDC-11)*, pp. 188-196, Edinburgh, Scotland, July 2002.
11. S. Vazhkudai, J. Schopf, I. Foster, "Predicting the Performance of Wide-Area Data Transfers", *Proceedings of the 16<sup>th</sup> International Parallel and Distributed Processing Symposium (IPDPS 2002)*, Fort Lauderdale, Florida, April 2002.

12. S. Vazhkudai, S. Tuecke, I. Foster, "Replica Selection in the Globus Data Grid", *Proceedings of the IEEE International Conference on Cluster Computing and the Grid (CCGRID 2001)*, pp. 106-113, Brisbane, Australia, May 2001.
13. R. Buyya, S. Vazhkudai, "Compute Power Market: Towards a Market-Oriented Grid", *Proceedings of the IEEE Session on Global Computing on Personal Devices*, pp. 574-581, Brisbane, Australia, May 2001.
14. S. Vazhkudai, G.V. Laszewski, "A Greedy Grid—The Grid Economic Engine Directive", *Proceedings of the IEEE Workshop on Internet Computing and E-Commerce*, San Francisco, California, April 2001.
15. S. Vazhkudai, S. Tuecke "A Storage Broker for the Globus Environment—A ClassAd Based Implementation", *Poster in Supercomputing 2000*, Dallas, Texas, Nov 2000.
16. S. Vazhkudai, P.T. Maginnis, "The PODOS File System—Exploiting the High-Speed Communication Subsystem", *Proceedings of the IEEE International Workshop on Cluster Computing Technologies, Environments, & Applications*, pp. 453-459, Las Vegas, Nevada, June 2000.
17. S. Vazhkudai, H.C. Cunningham, "A Reusable Software Framework for Distributed Decision-Making Protocols", *Proceedings of the Workshop on Distributed Objects in Computational Science (DOCS'2000)*, pp. 867-873, Las Vegas, Nevada, June 2000.
18. S. Vazhkudai, P.T. Maginnis, "A High Performance Communication Subsystem for PODOS", *Proceedings of the First IEEE International Conference on Cluster Computing*, pp. 81-91, Melbourne, Australia, December 1999.

### **Tutorials**

1. S. Vazhkudai, "Grid Computing—The Hype and The Truth", *Proceedings of the Second Annual Mid-South Computing Conference*, Little Rock, Arkansas, April 2004.

### **Technical Reports**

1. S. Vazhkudai, X. Ma, M. Vilayannur, "Data Availability for Service Availability: Automated On-demand Data Reconstruction and Offloading on Supercomputers", *Tech Report 003174*, ORNL, September 2006.
2. J.A. Kohl, M.A. Reuter, S. Vazhkudai, "Application Manager Requirements", SNS 107030214-TD0002-R0A, November 2005.
3. S. Miller, J.W. Cobb, J.A. Kohl, P. Peterson, S. Ren, M.A. Reuter, J. Schwidder, T. Swain, S. Vazhkudai, G.A. Geist, M. Hagen, R. Riedel, "Functional Responsibilities of Components for the SNS Data Acquisition, Management, and Analysis Architecture", SNS 107030210-TD0003-R0B, November 2005.
4. S. Ren, J. Schwidder, S. Vazhkudai, "Data Management Requirements", SNS 107030212-TD0002-R0G, October 2005.
5. V. Freeh, X. Ma, J. Strickland, S. Vazhkudai, "Synergetic Resource Stealing: We Promise It Will Not Hurt Much", *Technical Report P05-123434*, ORNL, May 2005.
6. S. Vazhkudai, "Orchestrating Bulk Data Movement in Grid Environments", *Technical Report R04-121317*, ORNL, September 2004.
7. S. Vazhkudai, G.A. Geist, "Evaluation of Software Suites for the SNS Software Initiative", ORNL, January 2004.
8. S. Vazhkudai, "Bulk Data Transfer Forecasts and the Implications to Grid Scheduling", *Ph.D. Dissertation*, University of Mississippi, May 2003.
9. S. Vazhkudai, "Distributed Linux: Evolutionary Steps", *Masters Thesis, Technical Report CISE TR 1998-22*, University of Mississippi, December 1998.

### **Presentations**

1. S. Vazhkudai, "Optimizing End-User Data Delivery Using Storage Virtualization", *Systems Group Seminar, Department of Computer Science and Engineering, Ohio State University*, October 2006. (*Invited*)
2. S. Vazhkudai, "IO Virtualization: Robust Storage Management in the Machine-Room and Beyond", *Virtualization in HPC*, Nashville, TN, September 2006.

3. X. Ma, S. Vazhkudai, V. Freeh, J. Strickland, N. Tammineedi, S.L. Scott, "FreeLoader: Desktop Storage Scavenging", Intel, Champaign, IL, December 2005.
4. S. Vazhkudai, S. Studham, "Storage Research at ORNL", DOE/NSF HEC-IWG Workshop, Grapevine TX. Aug 16, 2005.
5. S. Vazhkudai, "Towards a Data Management Infrastructure for Spallation Neutron Source", North Carolina State University, Raleigh, May 2004.
6. S. Vazhkudai, "Transfer and Storage Management for Grids", University of Tennessee, Knoxville, February 2004.
7. S. Vazhkudai, "Peeking into 21<sup>st</sup> Century Distributed Data Intensive Science", at Vanderbilt University/Univ of Nebraska-Lincoln/Univ of Alabama-Birmingham/ORNL/Cigital Corporation/Axiom Corporation/Pittsburgh Supercomputing Center, April-July 2003.
8. S. Vazhkudai, "Server Selection in Data Grids", University of Southern California, California, April 2001.
9. S. Vazhkudai, "A Matchmaking Approach to Replica Selection", Fermilabs, Illinois, January 2001. (*Invited*)
10. S. Vazhkudai, "The Design and Evolution of Communication in PODOS", Atlanta Linux Showcase, Georgia, October 1999.
11. S. Vazhkudai, "Transmission Group based Communication for PODOS", LinuxWorld Expo Conference, California, August 1999.

### **Book Chapters**

1. Contributed chapters to *Linux System Administration, Linux Networking & Linux Security and Ethics* by Tobin Maginnis, Published by John Wiley & Sons

### **Grants**

1. S. Vazhkudai, X. Ma, J.W. Cobb, G. Pike, S.L. Scott, "Storage Virtualization: An Integrated Approach to Machine-Room Storage Management", *ORNL Petascale LDRD funds*, \$600K, FY 2007-2008 (Role: PI).
2. X. Ma, Y. Zhou, V.W. Freeh, J. Blondin, S. Vazhkudai, M. Vilayannur, "Adaptive I/O Stack for High End Computing", NSF HECURA, \$500K FY 2007-2009 (Role: Senior Personnel).
3. B. Bland, K. Chanchio, C. Engelmann, X. Ma, S.L. Scott, S. Vazhkudai, "Reliability, availability and serviceability (RAS) for terascale computing", *ORNL Terascale LDRD funds*, \$576K, FY 2005-2006 (Role: Co-PI).
4. S. Vazhkudai, G.A. Geist, "A Neutron Science Portal Framework to facilitate Remote Access to SNS Data and Computation", *ORNL LDRD funds*, \$360K, FY 2004-2005 (Role: PI).

### **SOFTWARE SYSTEMS DEVELOPED**

- **Portals:** An infrastructure for remote data/computing/viz access.
- **Distributed Storage Systems:** FreeLoader distributed storage cache, High-speed networked file system for Distributed Linux, Storage Broker for Data Grids
- **Distributed Scheduling:** Remote Process Execution Environment for Distributed Linux
- **Clustering:** Distributed OS Kernel for Linux, Framework for Distributed Protocols.
- **Networks/Communication Systems:** Parallel downloads of Grid Data, Parallel/Short-circuited Communication Architecture for Distributed Linux, Hierarchical Group Communication Protocol for Distributed Linux, Remote Folder Sharing Environment.
- **Performance Predictions:** Regression based statistical prediction middleware for bulk Grid data transfers
- **Information Services:** GridFTP performance information provider for MDS.
- **Operating System Components:** Yash (*Enhanced C Shell*).
- **Compilers:** Shell Script to C Converter.
- **Databases:** Query Processor for Relational Algebra expressions.
- **WWW:** Web based Testing Environment
- **Component Architectures:** Chaos (*IDE for C projects under Unix using Implicit Invocation of objects*), File Manager for UNIX (*Windows like*).

## **AWARDS**

- Doctoral Fellowship (Spring 2003) - **U of M** (*Awarded to 10 candidates each year*)
- Ph.D. Dissertation Fellowship (2001 - 2002) - **Argonne National Laboratory** (*Research in Data Grids*)
- Wallace Givens Fellowship (Summer 2000) - **ANL/US Department of Energy** (*Awarded each year to 3 select candidates from a group of approximately 40 students*)
- Inducted into Upsilon Pi Epsilon and Gamma Beta Phi Honor societies (*GPA based*)

## **PERSONAL DETAILS**

- Citizenship: India
- Visa Status: H1-B (Work Permit)
- Date of Birth: 02-09-1975

## **REFERENCES**

- Available upon request.