

Erik Riedel

1509 Shady Avenue, Pittsburgh, PA 15217
riedel@cs.cmu.edu, www.cs.cmu.edu/~riedel

Education

Ph.D., Computer Engineering <i>November 1999</i>	<i>Electrical and Computer Engineering, Carnegie Mellon University</i> <i>Thesis "Active Disks - Remote Execution for Network-Attached Storage"</i> <i>Committee: David Nagle, Christos Faloutsos, Garth Gibson, Pradeep Khosla, Jim Gray</i>
M.S.E., Software Engineering <i>December 1994</i>	<i>Carnegie Mellon University, Pittsburgh</i> <i>Joint program: School of Computer Science and Software Engineering Institute</i>
B.S., Math/Computer Science <i>May 1992</i>	<i>Carnegie Mellon University, Pittsburgh</i> <i>with University Honors</i>
Strategic Leadership Program <i>March 2008</i>	<i>Harvard Business School, Boston</i> <i>Executive Education</i>

Research Interests

Storage devices and storage systems, with a focus on "smart" storage functions, disk drives, security, and novel interfaces between applications and storage. Distributed and scalable systems, including optimization for performance, power and reliability. Consumer and ubiquitous storage informed by user-centered design.

Research Experience

Senior Director <i>3/2002 to present</i>	<i>Interfaces & Architecture, Seagate Research, Pittsburgh</i> Founded and led a new research group focussed on novel interfaces and functions to increase the intelligence of storage systems. The group's exploratory research on storage in both enterprise and consumer segments tries to make storage more scalable, more reliable, easier to use and easier to manage. Responsible for hiring, setting research direction and leading a group of researchers with a \$1.5m budget. Led technology transfer to product groups; interacted with customers, partners and industry-wide organizations; contributed to corporate strategy, business development and due diligence. Projects in scalable enterprise storage, system-level reliability, power management, consumer storage, system architecture and user-centered design.
Technical Leader <i>10/2002 to present</i>	<i>Storage Networking Industry Association (SNIA), San Francisco & nationwide</i> Member of the SNIA Technical Council since May 2007. Chair of the SNIA Green Storage Technical Work Group since it was founded in August 2007. Co-chair of the SNIA Object-based Storage Device (OSD) Technical Work Group for five years, leading to the ANSI T10 OSD standard (INCITS 1355-D) for this new storage interface.
Researcher <i>11/1999 to 3/2002</i>	<i>Storage Systems Program, Hewlett-Packard Laboratories, Palo Alto</i> Research in storage systems, including networked storage, security, scalable storage, and distributed systems. Part of a 20-person research team developing technologies for globally distributed and constantly accessible storage. Worked with product divisions to contribute research ideas to current products and future product strategies.
Doctoral Thesis <i>08/1996 to 10/1999</i>	<i>Parallel Data Laboratory, Carnegie Mellon University, Pittsburgh</i> Original research <i>Active Disks - Remote Execution for Network-Attached Storage</i> in the Parallel Data Lab (PDL) and Center for Automated Learning and Discovery (CALD). Designed and prototyped a system for executing application-specific code on individual disk drives, taking advantage of compute and memory close to data to improve system performance, scalability and flexibility. Demonstrated the benefits of Active Disks for applications in database, data mining and multimedia with models and prototypes.

Erik Riedel

- Researcher** *Microsoft Research/Bay Area Research Center, San Francisco*
06/1997 to 09/1997
Explored the limits of storage subsystems for multimedia and database workloads, including examination of both hardware limitations and software bottlenecks.
- Research Assistant** *Parallel Data Laboratory, Carnegie Mellon University, Pittsburgh*
08/1995 to 10/1999
Research for the *Network-Attached Secure Disks* (NASD) project sponsored by ARPA and an industry consortium. Contributed to analysis, design, and prototypes to demonstrate the benefits of network-attached storage. Areas of work included parallel storage, distributed file systems, large-scale scientific and wide-area storage.
- Visiting Researcher** *Schlumberger Austin Research, Austin*
07/1995 to 08/1995
Analyzed the performance of seismic data acquisition and processing applications in production use by Schlumberger Oilfield Services. Suggested improvements to speed analysis by 10x and make possible the processing of much larger survey regions.
- Research Assistant** *School of Computer Science, Carnegie Mellon University, Pittsburgh*
Dept. of Chemical Engineering, Massachusetts Institute of Technology, Cambridge
05/1992 to 06/1995
Responsible for five person software engineering team of an NSF Grand Challenge project across Environmental Engineering, CS and Public Policy at CMU and MIT. Developed an object-oriented framework for environmental modeling and air quality.
- Team Member** *Master of Software Engineering Program, Carnegie Mellon University, Pittsburgh*
08/1993 to 12/1994
Member and leader of a 15 person team developing software for use on NASA robots. Participated in process improvement using the SEI Capability Maturity Model (CMM) and applying the Personal Software Process (PSP) to team-based software process.
-

Teaching Experience

- Teaching Assistant** *Electrical and Computer Engineering, Carnegie Mellon University*
08/1998 to 12/1998
Faculty: Professor Philip Koopman. Senior-level undergraduate course “Memory Systems Architecture” for 30 students.
- Teaching Assistant** *School of Computer Science, Carnegie Mellon University*
01/1994 to 06/1994
Faculty: Professor Linda Northrop. Senior-level undergraduate project course “Software Engineering” for 25 students. Responsible for grading and mentoring in teamwork, object-oriented design, software process, and project management.
- Teaching Assistant** *School of Computer Science, Carnegie Mellon University, Pittsburgh*
08/1992 to 12/1992
01/1993 to 05/1993
Faculty: Professor Bernd Bruegge. Undergraduate and first-year graduate project courses in “Software Engineering” of 45 and 15 students. Mentoring in teamwork, project management and object-oriented design.
-

Community Service

Member of the Technical Advisory Board for *The Technology Collaborative* (formerly *Pittsburgh Digital Greenhouse*), 2003-present. Member of the *Storage Networking Industry Association (SNIA)* Technical Council, 2007-present. Co-chair of the program committee for the *USENIX FAST 2008* conference. Member of numerous conference program committees in storage and systems and regular reviewer.

Other Information

- Citizenship** *German, with Permanent Resident status in the United States.*
- Languages** *Fluent in English and German. Forgotten lots of French.*

Systems Projects

Seamless Data - a research effort to improve data synchronization in consumer storage. Users of multiple fixed and mobile devices with significant storage (cellphones, PVRs, handheld PMPs, storage appliances) are presented with a single “seamless” view of all their data with common access, retrieval and protection. Sharing and organizing of personal and produced content is improved through common metadata and data transfer protocols. The effort included technical architecture and prototyping informed by user-centered design techniques (directly engaging consumers to understand desired user experiences for the technology).

Data-aware Storage - a research effort to increase the semantics of storage devices and enable data processing, including video & image processing for searching and organizing large collections of photos or videos, understanding of compression formats, and datatype-specific optimizations for reliability and power.

Simplicity and User-Centered Design - engineers and product managers sitting in their Silicon Valley offices don’t create technology products that are usable or accessible to a majority of the population. As a result, many consumer electronics products remain unsold or unused. User-centered research and design techniques provide a structured approach to co-designing technology and business models with direct user feedback. Products cannot be made usable by hiding complexity behind a fancy user interface; function and behavior must be designed from the ground up to be inherently simple. This effort tried to align these approaches with existing R&D and product management processes to achieve true and sustainable innovation.

Green Storage - an industry-wide technical effort to define metrics for power & energy efficiency in storage systems. The effort coordinated among storage OEMs, disk drive firmware and mechanical groups, industry groups (The Green Grid, DMTF, ClimateSavers), system admins and the EPA ENERGY STAR program.

System-level Reliability - a research effort to study the impact of drive-level failures on storage system availability and reliability. The focus of this data analysis from a large field population of disk drives was how to minimize data loss through improved failure prediction (“smarter SMART”) and how to reduce recovery times and vulnerability windows in redundant and replicated storage.

Object-based Storage (OSD) - a prototype and industry ecosystem effort to incorporate space management into storage devices. A portion of the file system is made part of the disk drive or array firmware. Instead of using small, fixed-sized blocks, object storage provides variable-sized objects that expand & contract based on the data they contain. A range of attributes are associated with each object and allow data to be tagged & classified to influence device performance, reliability and QoS. Standardized by a multi-year, multi-vendor effort in the SNIA, ratified and published as ANSI Project T10/1355-D, SCSI/OSD command set.

Global Data Placement (GDP) i-shadow - a research effort on global-scale storage that allows data to be shared by users around the corner or around the world. Areas of study include highly-distributed security across domains; adaptive consistency to make global scale feasible; and automatic resource optimization based on performance and app demands. Tag line “just like your shadow, your data is always with you.”

Active Disks - use cycles and memory on individual disk drives to perform application-level processing. Computation is distributed among hosts and the often large number of disks that store a particular data set. Database systems, for example, routinely have 100 or 1,000 disks and the aggregate compute power with modern drives exceeds the processing power of even SMP hosts by 4x to 10x. This thesis work describes a range of applications that benefit from this technology, including databases, data mining, and multimedia. Results from modeling and a prototype implementation show that gains are realizable in a running system. This work helped inspire successful database products from Netezza, including their patent (U.S. 7,464,106).

Network-Attached Secure Disks (NASD) - architecture for next-gen storage devices eliminates the need for expensive servers to act as a store-and-forward bridge between clients on one network (e.g. ATM) and storage devices on a separate network (e.g. FC). Individual disk drives are made first-class citizens on the same networks as clients (e.g. GigE) and common case data transfers proceed between clients and drives. Intervention of the server is required only for policy decisions, greatly increasing bandwidth and scalability and reducing cost. Our prototypes explored the basic issues of network-attached secure storage, the benefits to server scalability and app bandwidth, and the mapping of existing file systems onto this new architecture.

Geographic Environmental Modeling System (GEMS) - a framework for pollution modeling including data management, parallel computing, modern user interfaces and visualization. The project was a collaborative effort between scientists in CS, Environmental Engineering and Public Policy at CMU and MIT as part of an NSF Grand Challenge. Sponsored by regulators in the state of California, power companies, and the EPA.

Refereed Publications

- Ningfang Mi, Alma Riska, Xin Li, Evgenia Smirni and Erik Riedel, "Restraint Utilization of Idleness for Transparent Scheduling of Background Tasks" to appear in *ACM SIGMETRICS*. Seattle, WA. June 2009.
- M. Zhang, Alma Riska, Erik Riedel, "Characterization of the Storage Subsystem Workload under E-commerce Applications" *5th QEST*. Saint Malo, France. September 2008.
- D. Nagle, M. Factor, S. Iren, D. Naor, E. Riedel, O. Rodeh and J. Satran "The ANSI T10 object-based storage standard and current implementations" *IBM Journal of Research and Development* 52 (4/5). July/Sept 2008.
- Ningfang Mi, Alma Riska, Evgenia Smirni and Erik Riedel "Enhancing data availability in disk drives through background jobs" *Symposium on Dependable Systems and Networks (DSN)*. Anchorage, AK. June 2008.
- A. Riska and E. Riedel "Idle Read After Write" *USENIX Annual Technical Conference*. Boston, MA. June 2008.
- Alma Riska, James Larkby-Lahet and Erik Riedel "Evaluating Block-level Optimization Through the IO Path". *USENIX Annual Technical Conference*. Santa Clara, CA. June 2007.
- A. Riska and E. Riedel "Long-range dependence at the disk drive level" *3rd QEST*. Riverside, CA. Sept 2006.
- Jianyong Zhang, Anand Sivasubramaniam, Qian Wang, Alma Riska and Erik Riedel "Storage performance virtualization via throughput and latency control" *ACM Transactions on Storage* 2 (6). August 2006.
- Qi Zhang, Alma Riska, Erik Riedel, Ningfang Mi and Evgenia Smirni "Evaluating performability of systems with background jobs" *Symposium on Dependable Systems and Networks (DSN)*. Philadelphia, PA. June 2006.
- Alma Riska and Erik Riedel "Disk Drive Level Workload Characterization" *USENIX Annual Technical Conference*. Boston, MA. June 2006. Short paper presentation.
- Jianyong Zhang, Alma Riska, Anand Sivasubramaniam, Qian Wang and Erik Riedel "Storage Performance virtualization via Throughput and Latency Control" *13th MASCOTS*. Atlanta, GA. September 2005.
- Qi Zhang, Alma Riska and Erik Riedel "Workload Propagation - Overload in bursty servers" *2nd Int'l QEST Conference*. Torino Italy. September 2005.
- Qi Zhang, Alma Riska, Evgenia Smirni and Erik Riedel "Bottleneck Identification in E-commerce Systems" *9th International WCW Workshop*. Beijing, China. October 2004.
- Alma Riska, Erik Riedel and Sami Iren "Adaptive disk scheduling for overload management" *1st International QEST Conference*. The Netherlands. September 2004.
- Larry Huston, Rahul Sukthankar, Rajiv Wickremesinghe, M. Satyanarayanan, Greg Ganger, Erik Riedel and Anastassia Ailamaki "Diamond: A Storage Architecture for Early Discard in Interactive Search" *3rd USENIX Conference on File and Storage Technologies (FAST)*. San Francisco, CA. March 2004.
- Alma Riska, Erik Riedel "It's not fair - evaluating efficient disk scheduling" *MASCOTS*. Orlando, FL. Oct 2003.
- Alma Riska, Erik Riedel and Sami Iren "Managing Overload Via Adaptive Scheduling" *1st Workshop on Algorithms and Architecture for Self-Managing Systems*. San Diego, CA. June 2003.
- Erik Riedel "Storage Systems: Not Just a Bunch of Disks Anymore" *ACM Queue* 1 (4), June 2003.
- David Anderson, Erik Riedel and James Dykes "More than an interface - SCSI vs. ATA" *File and Storage Technology Conference (FAST 2003)*. San Francisco, CA. March 2003.
- Mahesh Kallahalla, Erik Riedel, Ram Swaminathan, Qian Wang and Kevin Fu "Plutus: Scalable secure file sharing on untrusted storage" *2nd USENIX File & Storage Technology (FAST)*. San Francisco, CA. March 2003.
- Erik Riedel, Mahesh Kallahalla and Ram Swaminathan "A framework for evaluating storage system security" *1st USENIX File and Storage Technology Conference (FAST)*. Monterey, CA. January 2002.
- Erik Riedel, Garth Gibson, David Nagle and Christos Faloutsos "Active Storage For Large-Scale Data Processing" *IEEE Computer*, New York, NY. June 2001. Nominated as an award paper.
- Erik Riedel, Susan Spence and Alistair Veitch "When Local Becomes Global - An Application Study of Consistency in a Networked World" *IPCCC 2001*. Phoenix, AZ. April 2001. Invited paper.
- Alistair Veitch, Erik Riedel, Simon Towers and John Wilkes "Towards Global Storage Management and Data Placement" *8th Workshop on Hot Topics in Operating Systems (HotOS)*. Schloss Elmau, Germany. May 2001.
- Chris Lumb, Jiri Schindler, Greg Ganger, David Nagle and Erik Riedel "Towards Higher Disk Head Utilization: Extracting 'Free' Bandwidth From Busy Disk Drives" *OSDI 2000*. San Diego, CA. October 2000.

Erik Riedel

- Erik Riedel, Christos Faloutsos, Greg Ganger and David Nagle "Data Mining on an OLTP System (Nearly) For Free" *ACM SIGMOD Int'l Conference on Management of Data (SIGMOD)*. Dallas, TX. May 2000.
- Erik Riedel, Garth Gibson and Christos Faloutsos "Active Storage For Large-Scale Data Mining and Multimedia" *24th Int'l Conference on Very Large Databases (VLDB)*. New York, NY. August 1998.
- Erik Riedel, Catharine van Ingen and Jim Gray "A Performance Study of Sequential I/O on Windows NT" *2nd Usenix Windows NT Symposium*. Seattle, WA. August 1998. *Best Student Paper*.
- Garth Gibson, David Nagle, Khalil Amiri, Jeff Butler, Fay Chang, Howard Gobioff, Charles Hardin, Erik Riedel, David Rochberg and Jim Zelenka "A Cost-Effective, High-Bandwidth Storage Architecture" *ASPLOS-VIII Conference on Architectural Support for Programming Languages and OS*. San Jose, CA. October 1998.
- Garth Gibson, David Nagle, Khalil Amiri, Fay Chang, Eugene Feinberg, Howard Gobioff, Chen Lee, Berend Ozceri, Erik Riedel, David Rochberg and Jim Zelenka "File Server Scaling with Network-Attached Secure Disks" *Int'l Conference on Measurement & Modeling of Computer Systems (SIGMETRICS)*. Seattle, WA. June 1997.
- Erik Riedel and Garth Gibson "Understanding Customer Dissatisfaction With Underutilized Distributed File Servers" *5th Conference on Mass Storage Systems & Technology (MSST)*. College Park, MD. September 1996.
- Bernd Bruegge, Greg McRae, Erik Riedel and Armistead Russell "Developing GEMS: An Environmental Modeling System" *IEEE Computational Science and Engineering*. September 1995.
- Erik Riedel and Bernd Bruegge "GEMS: Towards an Object-Oriented Framework" *8th European Conference on Object-Oriented Programming (ECOOP)*. Bologna, Italy. July 1994.

Invited Talks

- "Green Storage II: Metrics & Measurements" *Storage Networking World (SNW)*. Dallas, TX. October 2008.
- "Storage Security - From Research to Industry Best Practice" *Computer Systems Symposium*, NASA Goddard Space Flight Center. Greenbelt, MD. March 2008.
- "Smarter Data with Smarter Disc Drives" *TTI/Vanguard Conference*. Atlanta, GA. February 2008.
- "Data Storage for the Green Corporation" *AIT Global Conference*. United Nations, NY. November 2007.
- "Storage Security - From Research to Industry Best Practice" *5th Int'l Security in Storage Workshop*. San Diego, CA. September 2007.
- "It's A Wonder That They Work At All - Hard Drives, Embedded Software and 25 Petabytes a Day" *Systems Symposium*. Penn State University. State College, PA. September 2007.

Patents & Other Publications

- Eleven U.S. patents. 6,732,241 (issued May 2004) and 6,918,120 (issued July 2005) both as sole inventor; 7,405,941; 7,375,923 (2008); 7,313,694; 7,219,230; 7,203,317; 7,200,747; 7,171,557 (2007); 7,073,021 and 7,003,116 (2006) all as co-inventor. Filed five additional patent applications, including inventions in disk scheduling, disk array design, secure storage, on-drive data management and storage protocol design.
- Mache Creeger (editor) "CTO Storage Roundtable - Leaders in the storage industry ponder upcoming technologies and trends" *Communications of the ACM* 51 (8/9). Panelists Steve Kleiman, Erik Riedel, Mary Baker, Greg Ganger, Kirk McKusick, Eric Brewer and Margo Seltzer. August/September 2008.
- Erik Riedel "Active Disks - Remote Execution for Network-Attached Storage" *Doctoral Dissertation, CMU-CS-99-177*. November 1999. Nominated for ACM Dissertation Award.
- Garth Gibson, David Nagle, Khalil Amiri, Fay Chang, Howard Gobioff, Erik Riedel, David Rochberg and Jim Zelenka "Filesystems for Network-Attached Secure Disks" *Technical Report CMU-CS-97-118*. July 1997. Also the basis for *ANSI proposal T10/1355-D "SCSI Object Based Storage Device Commands (OSD)"*, July 2001, ratified in September 2004 as an approved ANSI standard by the T10 committee.
- Corbett, P., Prost, J-P., Demetriou, C., Gibson, G., Riedel, E., Zelenka, J., Chen, Y., Felten, E., Li, K., Hartman, J., Peterson, L., Bershada, B., Wolman, A. and Aydt, R. "Proposal for a Common Parallel File System Programming Interface" *Technical Report CMU-CS-96-193*. Also at *Supercomputing '96*. Pittsburgh, PA. November 1996.