

**Philip Carns**  
carns@mcs.anl.gov  
630-252-6707

Argonne National Laboratory  
9700 S. Cass Ave.  
Lemont IL 60439

---

I am a computer scientist specializing in the design and development of scalable methods of information storage, movement, and management for data-intensive high-performance computation. My work is driven by the desire to identify and solve foundational problems in the field through a combination of technical leadership, collaborative research, and creative engineering.

## Degrees attained

- **Doctor of Philosophy** Computer Engineering. Clemson University, May 2005. Dissertation titled “Achieving Scalability in Parallel File Systems.”
- **Master of Science** Computer Engineering. Clemson University, December 2001. Thesis titled “Design and Analysis of a Network Transfer Layer for Parallel File Systems.”
- **Bachelor of Science** Computer Engineering. Clemson University, June 1999. Clemson University Honors Program.

## Positions held

- **Argonne National Laboratory** Mathematics and Computer Science Division
  - June 2022 – present: Computer Scientist.
  - Jan. 2008 – May 2022: Principal Software Development Specialist.

Computer science research in support of data-intensive scientific computing.

Technical lead, principal investigator, or developer for influential HPC research projects including Darshan (application I/O characterization), TOKIO (platform I/O characterization), Mochi (composable data services), PVFS (parallel file system), CODES (storage system simulation), and the Exascale Computing Project (data libraries and services for exascale platforms); received multiple R&D 100 awards.

Community service – organized events including workshops, seminars, and BoFs; developed and shared expertise in software engineering, technical writing, research methodologies, and performance optimization; and served on numerous committees for journals, conferences, and funding programs.
- **Clemson University**
  - Dec. 2015 – present: Adjunct Associate Professor of Electrical and Computer Engineering.
- **Northwestern University and Argonne National Laboratory**
  - Jan. 2015 – present: Northwestern–Argonne Institute for Science and Engineering (NAISE) Fellow.
- **Axiom Corporation** Grid Architecture Unit
  - Jan. 2005 – Dec. 2007: Software Developer.

Research and development of large-scale data storage and processing technology for business data management.
- **Clemson University** Parallel Architecture Research Laboratory

- Aug. 1999 – Dec. 2004: Research Assistant.  
Research and design of parallel file systems and network abstractions for HPC.
- **Argonne National Laboratory** Mathematics and Computer Science Division
  - Aug. 2001 – Dec. 2004: short-term senior programmer appointment.
  - June 2001 – Aug. 2001: summer research internship.
  - June 2000 – Aug. 2000: summer research internship.
 Research and design of parallel file systems and network abstractions for HPC.

## Honors and awards

- **R&D 100 Awards:** annual awards recognizing the 100 most innovative technologies across industry, research, and academia
  - 2021 R&D 100 Award: Mochi (composable data services) <https://www.anl.gov/article/argonne-captures-3-rd-100-awards-for-innovative-technology>
  - 2020 R&D 100 finalist: Argobots (user-level threading) <https://www.anl.gov/mcs/article/argobots-technology-named-finalist-for-the-2020-rd-100-awards>
  - 2018 R&D 100 Award: Darshan (I/O characterization) <https://www.anl.gov/mcs/article/darshan-wins-rd-100-award-for-2018>
- **Best paper awards:**
  - IEEE Symposium on Massive Storage Systems and Technologies (MSST 2011)
  - 4th Annual Linux Showcase and Conference (2000)
- **ANL recognitions:**
  - Argonne Pacesetter award for exceptional work in production deployment of ALCF BG/P storage system (2009)
- **Scholastic:**
  - NASA Graduate Student Research Program Fellowship (2000-2004)
  - Riggs Most Outstanding Senior Award, Clemson Electrical and Computer Engineering Department (1999)
  - First in computer engineering graduating class, Clemson University (1999)

## Professional Activities

### Argonne National Laboratory

- Argonne Training Program on Extreme-Scale Computing (ATPESC) review committee, 2022
- MCS hiring committee interview rotator, 2020 – 2021
- Subject Matter Expert for the A21 NRE program, 2018 – present
- Co-track lead: Argonne Training Program for Exascale Systems (ATPESC), Data and I/O day, 2015 – present

## Department of Energy

- Workshop attendee and report contributor for the following DOE workshops:
  - ASCR Workshop on the Management and Storage of Scientific Data, 2022
  - Storage Systems and I/O 2018: Organizing, Storing, and Accessing Data for Scientific Discovery, 2018
  - Management, Visualization, and Analysis of Experimental and Observational Data (EOD): The Convergence of Data and Computing, 2015
  - DOE Storage Systems and Input/Output, 2014
- Advisory board member for the AppSysFusion: Integrated System and Application Continuous Performance Monitoring and Analysis Capability project at Sandia National Laboratories, 2022 – present.
- Technical reviewer for the NERSC-9 HPC system at the National Energy Research Scientific Computing Center, 2018
- Storage evaluation working group member for the CORAL 2 program, 2018

## National Science Foundation

- Proposal review panelist: National Science Foundation (NSF) CISE OAC Core Research program solicitation, NSF 20-591, 2021

## Community organization

- Co-organizer: Dagstuhl Seminar 21332 “Understanding I/O Behavior in Scientific and Data-Intensive Computing”. Intensive five-day hybrid event that brought together 35 of the world’s foremost subject matter experts to survey the state of the art and establish a shared vision for future research, August 2021 <https://www.dagstuhl.de/en/program/calendar/semhp/?semnr=21332>
- Steering committee: International Parallel Data Systems Workshop (PDSW) 2020 – present
- General chair: 5th International Parallel Data Systems Workshop (PDSW), 2020
- Program co-chair: 4th International Parallel Data Systems Workshop (PDSW), 2019
- Co-organizer: Enabling Data Services for HPC Birds of a Feather session at the International Conference for High Performance Computing, Networking, Storage and Analysis (SC), 2018–2019
- Co-organizer: Analyzing Parallel I/O Birds of a Feather session at the International Conference for High Performance Computing, Networking, Storage and Analysis (SC), 2014–2018
- Organizer: 4th annual workshop on Interfaces and Abstractions for Scientific Data Storage (IASDS), 2012

## Students advised

- Steven Chien, KTH Royal Institute of Technology, PhD opponent, 2022
- Zhe Wang, Rutgers University, PhD committee, 2022
- Jakob Lüttgau, Universität Hamburg, PhD committee, 2021

## Other community activities

- Review board: IEEE Transactions on Parallel and Distributed Systems, 2020 – 2021
- Journal reviewer: ACM Transactions on Storage, IEEE Transactions on Parallel and Distributed Systems, IEEE Transactions on Computers, BCS Computer Journal, Elsevier’s Parallel Computing Journal, the Journal of Parallel Computing, Elsevier’s Simulation Modeling Practice and Theory, ACM Computing Surveys, and the International Journal of High Performance Computing Applications
- Program committee member (various years) for conferences and workshops including SC, ISC, IPDPS, WSC, MSST, IEEE Cluster, ICPADS, BDSE, IEEE BigData, PDSW, BDCloud, FAB, CBDCOM, DS-DIS, BigData Congress, DSS, SSDBM, ICA3PP, ISC, CCGrid, and NAS
- Mentor: “ECE Plugged In” program at Clemson University, 2014-present
- Reviewer: Kentucky Science and Engineering Foundation R&D Excellence program, 2014

## Research Products

### Refereed journal articles and book chapters

- [1] André Brinkmann, Kathryn Mohror, Weikuan Yu, Philip Carns, Toni Cortes, Scott A Klasky, Alberto Miranda, Franz-Josef Pfreundt, Robert B Ross, and Marc-André Vef. Ad hoc file systems for high-performance computing. *Journal of Computer Science and Technology*, 35(1):4–26, 2020.
- [2] Robert B Ross, George Amvrosiadis, Philip Carns, Charles D Cranor, Matthieu Dorier, Kevin Harms, Greg Ganger, Garth Gibson, Samuel K Gutierrez, Robert Latham, et al. Mochi: Composing data services for high-performance computing environments. *Journal of Computer Science and Technology*, 35(1):121–144, 2020.
- [3] Jerome Soumagne, Philip Carns, and Robert B Ross. Advancing RPC for data services at exascale. *IEEE Data Engineering Bulletin*, 43:23–34, 2020.
- [4] Dong Dai, Yong Chen, Philip Carns, John Jenkins, Wei Zhang, and Robert Ross. Managing rich metadata in high-performance computing systems using a graph model. *IEEE Transactions on Parallel and Distributed Systems*, 30(7):1613–1627, 2018.
- [5] Pierre Matri, Yevhen Alforov, Alvaro Brandon, María S Pérez, Alexandru Costan, Gabriel Antoniu, Michael Kuhn, Philip Carns, and Thomas Ludwig. Mission possible: Unify HPC and big data stacks towards application-defined blobs at the storage layer. *Future Generation Computer Systems*, 109:668–677, 2018.
- [6] Noah Wolfe, Misbah Mubarak, Christopher D Carothers, Robert B Ross, and Philip H Carns. Modeling large-scale slim fly networks using parallel discrete-event simulation. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 28(4):1–25, 2018.
- [7] Scott Parker, John Mellor-Crummey, Dong H Ahn, Heike Jagode, Holger Brunst, Sameer Shende, Allen D Malony, David Lecomber, John V DelSignore Jr, Ronny Tschüter, et al. Performance analysis and debugging tools at scale. In *Exascale Scientific Applications: Scalability and Performance Portability*. CRC Press, 2017.
- [8] Sangmin Seo, Abdelhalim Amer, Pavan Balaji, Cyril Bordage, George Bosilca, Alex Brooks, Philip Carns, Adrian Castello, Damien Genet, Thomas Herault, et al. Argobots: A lightweight low-level threading and tasking framework. *IEEE Transactions on Parallel and Distributed Systems*, 2017.

- [9] Dong Dai, Philip Carns, Robert B Ross, John Jenkins, Nicholas Muirhead, and Yong Chen. An asynchronous traversal engine for graph-based rich metadata management. *Parallel Computing*, 58:140–156, 2016.
- [10] Misbah Mubarak, Christopher D Carothers, Robert B Ross, and Philip Carns. Enabling parallel simulation of large-scale HPC network systems. *IEEE Transactions on Parallel and Distributed Computing*, 2016.
- [11] Philip Carns. Darshan. In *High Performance Parallel I/O*, pages 309–315. CRC Press, 2014.
- [12] Ning Liu, Christopher Carothers, Jason Cope, Philip Carns, and Robert Ross. Model and simulation of exascale communication networks. *Journal of Simulation*, 6:227–236, 2012.
- [13] Philip Carns, Kevin Harms, William Allcock, Charles Bacon, Samuel Lang, Robert Latham, and Robert Ross. Understanding and improving computational science storage access through continuous characterization. *ACM Transactions on Storage (TOS)*, 7(3):1–26, 2011.
- [14] Robert Ross, Philip Carns, and David Metheny. Parallel file systems. In *Data Engineering: Mining, Information, and Intelligence*, pages 143–168. Springer, 2010.

## Publications in refereed proceedings

- [1] Jean Luca Bez, Ahmad Maroof Karimi, Arnab K Paul, Bing Xie, Suren Byna, Philip Carns, Sarp Oral, Feiyi Wang, and Jesse Hanley. Access patterns and performance behaviors of multi-layer supercomputer I/O subsystems under production load. In *Proceedings of the 31st International Symposium on High-Performance Parallel and Distributed Computing*, pages 43–55, 2022.
- [2] Srinivasan Ramesh, Allen D Malony, Philip Carns, Robert B Ross, Matthieu Dorier, Jerome Soumagne, and Shane Snyder. SYMBIOSYS: A methodology for performance analysis of composable HPC data services. In *2021 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 35–45. IEEE, 2021.
- [3] Srinivasan Ramesh, Robert Ross, Matthieu Dorier, Allen Malony, Philip Carns, and Kevin Huck. SYMBIOMON: A high-performance, composable monitoring service. In *2021 IEEE 28th International Conference on High Performance Computing, Data, and Analytics (HiPC)*, pages 332–342. IEEE, 2021.
- [4] Houjun Tang, Bing Xie, Suren Byna, Philip Carns, Quincey Koziol, Sudarsun Kannan, and Jay Lofstead Sarp Oral. SCTuner: An auto-tuner addressing dynamic I/O needs on supercomputer I/O subsystems. 2021.
- [5] Bing Xie, Zilong Tan, Philip Carns, Jeff Chase, Kevin Harms, Jay Lofstead, Sarp Oral, Sudharshan S Vazhkudai, and Feiyi Wang. Interpreting write performance of supercomputer I/O systems with regression models. In *2021 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 557–566. IEEE, 2021.
- [6] Philip Carns, Kevin Harms, Bradley W Settlemyer, Atkinson Brian, and Robert B Ross. Keeping it real: Why HPC data services don’t achieve I/O microbenchmark performance. In *5th International Parallel Data Systems Workshop (PDSW 2020)*, 2020.
- [7] Eliakin del Rosario, Mikaela Currier, Mihailo Isakov, Sandeep Madireddy, Prasanna Balaprakash, Philip Carns, Robert B Ross, Kevin Harms, Shane Snyder, and Michel A Kinsky. Gauge: An interactive data-driven visualization tool for HPC application I/O performance analysis. In *5th International Parallel Data Systems Workshop (PDSW 2020)*, 2020.

- [8] Mihailo Isakov, Eliakin Del Rosario, Sandeep Madireddy, Prasanna Balaprakash, Philip Carns, Robert B Ross, and Michel A Kinsy. HPC I/O throughput bottleneck analysis with explainable local models. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–13. IEEE, 2020.
- [9] Mihailo Isakov, Eliakin del Rosario, Sandeep Madireddy, Prasanna Balaprakash, Philip Carns, Robert B Ross, and Michel A Kinsy. Toward generalizable models of I/O throughput. In *2020 IEEE/ACM International Workshop on Runtime and Operating Systems for Supercomputers (ROSS)*, 2020.
- [10] Zhengchun Liu, Ryan Lewis, Rajkumar Kettimuthu, Kevin Harms, Philip Carns, Nageswara Rao, Ian Foster, and Michael E Papka. Characterization and identification of HPC applications at leadership computing facility. In *Proceedings of the 34th ACM International Conference on Supercomputing*, pages 1–12, 2020.
- [11] Tirthak Patel, Suren Byna, Glenn K Lockwood, Nicholas J Wright, Philip Carns, Robert Ross, and Devesh Tiwari. Uncovering access, reuse, and sharing characteristics of I/O-intensive files on large-scale production HPC systems. In *18th USENIX Conference on File and Storage Technologies ({FAST} 20)*, 2020.
- [12] Glenn K Lockwood, Shane Snyder, Suren Byna, Philip Carns, and Nicholas J Wright. Understanding data motion in the modern HPC data center. In *4th International Parallel Data Systems Workshop (PDSW)*, 2019.
- [13] Sandeep Madireddy, Prasanna Balaprakash, Philip Carns, Robert Latham, Glenn K Lockwood, Robert Ross, Shane Snyder, and Stefan M Wild. Adaptive learning for concept drift in application performance modeling. In *Proceedings of the 48th International Conference on Parallel Processing*, pages 1–11, 2019.
- [14] Teng Wang, Suren Byna, Glenn K Lockwood, Shane Snyder, Philip Carns, Sunggon Kim, and Nicholas J Wright. A zoom-in analysis of I/O logs to detect root causes of I/O performance bottlenecks. In *19th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID)*, 2019.
- [15] Bing Xie, Zilong Tan, Philip Carns, Jeff Chase, Kevin Harms, Jay Lofstead, Sarp Oral, Sudharshan S Vazhkudai, and Feiyi Wang. Applying machine learning to understand write performance of large-scale parallel filesystems. In *4th International Parallel Data Systems Workshop (PDSW)*, 2019.
- [16] Matthieu Dorier, Philip Carns, Kevin Harms, Robert Latham, Robert Ross, Shane Snyder, Justin Wozniak, Samuel K Gutiérrez, Bob Robey, Brad Settlemyer, et al. Methodology for the rapid development of scalable HPC data services. In *Proc. of Workshop in conjunction with ACM/IEEE Supercomputing Conference, Dallas, TX, USA*, 2018.
- [17] Julian Martin Kunkel, Eugen Betke, Matt Bryson, Philip Carns, Rosemary Francis, Wolfgang Frings, Roland Laifer, and Sandra Mendez. Tools for analyzing parallel I/O. In *International Conference on High Performance Computing*, pages 49–70. Springer, Cham, 2018.
- [18] Glenn K. Lockwood, Shane Snyder, George Brown, Kevin Harms, Philip Carns, and Nicholas J. Wright. TOKIO on ClusterStor: Connecting standard tools to enable holistic I/O performance analysis. In *2018 Cray User Group*, 2018.
- [19] Glenn K Lockwood, Shane Snyder, Teng Wang, Suren Byna, Philip Carns, and Nicholas J Wright. A year in the life of a parallel file system. In *SC18: Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis*, pages 931–943. IEEE, 2018.

- [20] Jakob Lüttgau, Shane Snyder, Philip Carns, Justin M Wozniak, Julian Kunkel, and Thomas Ludwig. Toward understanding I/O behavior in HPC workflows. In *2018 IEEE/ACM 3rd International Workshop on Parallel Data Storage & Data Intensive Scalable Computing Systems (PDSW-DISCS)*, pages 64–75. IEEE, 2018.
- [21] Sandeep Madireddy, Prasanna Balaprakash, Philip Carns, Robert Latham, Robert Ross, Shane Snyder, and Stefan Wild. Modeling I/O performance variability using conditional variational auto encoders. In *IEEE Cluster*, 2018.
- [22] Sandeep Madireddy, Prasanna Balaprakash, Philip Carns, Robert Latham, Robert Ross, Shane Snyder, and Stefan M Wild. Machine learning based parallel I/O predictive modeling: A case study on lustre file systems. In *International Conference on High Performance Computing*, pages 184–204. Springer, 2018.
- [23] Pierre Matri, Philip Carns, Robert Ross, Alexandru Costan, María S Pérez, and Gabriel Antoniu. SLog: Large-scale logging middleware for HPC and big data convergence. In *2018 IEEE 38th International Conference on Distributed Computing Systems (ICDCS)*, pages 1507–1512. IEEE, 2018.
- [24] Teng Wang, Suren Byna, Glenn Lockwood, Nicholas Wright, Philip Carns, and Shane Snyder. IOMiner: Large-scale analytics framework for gaining knowledge from I/O logs. In *IEEE Cluster*, 2018.
- [25] Dong Dai, Yong Chen, Philip Carns, John Jenkins, and Robert Ross. Lightweight provenance service for high performance computing. In *Parallel Architectures and Compilation Techniques (PACT)*, 2017.
- [26] John Jenkins, Galen Shipman, Jamaludin Mohd-Yusof, Kipton Barros, Philip Carns, and Robert Ross. A case study in computational caching microservices for HPC. In *Second Annual Workshop on Emerging Parallel and Distributed Runtime Systems and Middleware (IPDRM 2017)*, pages 1309–1316. IEEE, 2017.
- [27] Glenn K Lockwood, Wucheryl Yoo, Suren Byna, Nicholas J Wright, Shane Snyder, Kevin Harms, Zachary Nault, and Philip Carns. UMAMI: a recipe for generating meaningful metrics through holistic I/O performance analysis. In *Proceedings of the 2nd Joint International Workshop on Parallel Data Storage & Data Intensive Scalable Computing Systems*, pages 55–60. ACM, 2017.
- [28] Xiaoqing Luo, Frank Mueller, Philip Carns, Jonathan Jenkins, Robert Latham, Robert Ross, and Shane Snyder. ScalaIOExtrap: Elastic I/O tracing and extrapolation. In *Parallel and Distributed Processing Symposium (IPDPS), 2017 IEEE International*, pages 585–594. IEEE, 2017.
- [29] Sandeep Madireddy, Prasanna Balaprakash, Philip Carns, Robert Latham, Robert Ross, Shane Snyder, and Stefan M Wild. Analysis and correlation of application I/O performance and system-wide I/O activity. In *2017 International Conference on Networking, Architecture, and Storage (NAS)*, pages 1–10. IEEE, 2017.
- [30] Pierre Matri, Yevhen Alforov, Álvaro Brandon, Michael Kuhn, Philip Carns, and Thomas Ludwig. Could blobs enable storage-based convergence between HPC and big data? In *IEEE Cluster Conference*, 2017.
- [31] Misbah Mubarak, Philip Carns, Jonathan Jenkins, Jianping Li, Nikhil Jain, Shane Snyder, Robert Ross, Abhinav Bhatele, Chris Carothers, and Kwan-Liu Ma. Quantifying I/O and communication traffic interference on burst buffer equipped dragonfly networks. In *IEEE Cluster Conference*, 2017.

- [32] Cong Xu, Shane Snyder, Omkar Kulkarni, Vishwanath Venkatesan, Philip Carns, Suren Byna, Robert Sisneros, and Kalyana Chadalavada. DXT: Darshan extended tracing. In *Proceedings of the Cray User Group meeting 2017 (CUG 2017)*, 2017.
- [33] Philip Carns, John Jenkins, Sangmin Seo, Shane Snyder, Robert B Ross, Charles D Cranor, Scott Atchley, and Torsten Hoefler. Enabling NVM for data-intensive scientific services. In *4th Workshop on Interactions of NVM/Flash with Operating Systems and Workloads (INFLOW 16)*, 2016.
- [34] Dong Dai, Yong Chen, Philip Carns, John Jenkins, Wei Zhang, and Robert Ross. GraphMeta: A graph-based engine for managing large-scale HPC rich metadata. In *Proceedings of IEEE Cluster 2016*, 2016.
- [35] Caitlin Ross, Christopher D. Carothers, Misbah Mubarak, Philip Carns, Robert Ross, Jianping Kelvin Li, and Kwan-Liu Ma. Visual data-analytics of large-scale parallel discrete-event simulations. In *Proceedings of the 7th International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS16)*, 2016.
- [36] Caitlin Ross, Misbah Mubarak, John Jenkins, Philip Carns, Christopher D Carothers, Robert Ross, Wei Tang, Wolfgang Gerlach, and Folker Meyer. A case study in using discrete-event simulation to improve the scalability of MG-RAST. In *Proceedings of the 2016 annual ACM Conference on SIGSIM Principles of Advanced Discrete Simulation*, pages 211–220, 2016.
- [37] Shane Snyder, Philip Carns, Kevin Harms, Robert Ross, Glenn K. Lockwood, and Nicholas J. Wright. Modular HPC I/O characterization with Darshan. In *Workshop on Extreme-Scale Programming Tools (ESPT 2016)*, 2016.
- [38] Neda Tavakoli, Dong Dai, John Jenkins, Philip Carns, Robert Ross, and Yong Chen. A software-defined approach for qos control in high-performance computing storage systems (poster). In *International Conference for High Performance Computing, Networking, Storage and Analysis (SC16)*, 2016.
- [39] Noah Wolfe, Christopher Carothers, Misbah Mubarak, Robert Ross, and Philip Carns. Modeling a million-node slim fly network using parallel discrete-event simulation. In *Proceedings of the 2016 annual ACM Conference on SIGSIM Principles of Advanced Discrete Simulation*, pages 189–199, 2016.
- [40] Philip Carns, Kevin Harms, John Jenkins, Misbah Mubarak, Robert B. Ross, and Christopher Carothers. Consistent hashing distance metrics for large-scale object storage (poster). In *International Conference for High Performance Computing, Networking, Storage and Analysis (SC15)*, 2015.
- [41] Dong Dai, Philip Carns, Robert Ross, John Jenkins, Kyle Blauer, and Yong Chen. GraphTrek: Asynchronous graph traversal for property graph based metadata management. In *IEEE International Conference on Cluster Computing (Cluster 2015)*, 2015.
- [42] J. LaPre, E.J. Gonsiorowski, C.D. Carothers, J. Jenkins, P. Carns, and R.B. Ross. Time warp state restoration via delta encoding. In *2015 Winter Simulation Conferences (WSC'15)*, 2015.
- [43] Xiaqing Luo, Frank Mueller, Philip Carns, John Jenkins, Robert Ross, and Shane Snyder. ScalaIOExtrap: Elastic I/O tracing and extrapolation. In *Workshop on Extreme-Scale Programming Tools (ESPT 2015)*, 2015.
- [44] Huong Luu, Marianne Winslett, William Gropp, Robert Ross, Philip Carns, Kevin arms, Mr Prabhat, Suren Byna, and Yushu Yao. A multiplatform study of I/O behavior on petascale supercomputers.



In *Proceedings of the 24th International Symposium on High-Performance Parallel and Distributed Computing*, pages 33–44. ACM, 2015.

- [45] Shane Snyder, Philip Carns, Robert Latham, Misbah Mubarak, Rob Ross, Christopher Carothers, Babak Behzad, Huong Vu Thanh Luu, Surendra Byna, and Prabhat. Techniques for modeling large-scale HPC I/O workloads. In *Proceedings of the 6th International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS15)*, 2015.
- [46] Jerome Soumagne, Philip H Carns, Dries Kimpe, Quincey Koziol, and Robert B Ross. A remote procedure call approach for extreme-scale services. *arXiv preprint arXiv:1510.02135*, 2015.
- [47] Dong Dai, Robert B Ross, Philip Carns, Dries Kimpe, and Yong Chen. Using property graphs for rich metadata management in HPC systems. In *Proceedings of the 9th Parallel Data Storage Workshop (PDSW 2014)*, 2014.
- [48] Sidharth Kumar, Cameron Christensen, John A Schmidt, Peer-Timo Bremer, Eric Brugger, Venkatram Vishwanath, Philip Carns, Hemanth Kolla, Ray Grout, Jacqueline Chen, et al. Fast multiresolution reads of massive simulation datasets. In *Proceedings of the International Supercomputing Conference (ISC 2014)*, pages 314–330. Springer International Publishing, 2014.
- [49] Sidharth Kumar, John Edwards, Peer-Timo Bremer, Aaron Knoll, Cameron Christensen, Venkatram Vishwanath, Philip Carns, John A Schmidt, and Valerio Pascucci. Efficient I/O and storage of adaptive-resolution data. In *SC'14: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 413–423. IEEE, 2014.
- [50] Misbah Mubarak, Christopher D Carothers, Robert B Ross, and Philip Carns. A case study in using massively parallel simulation for extreme-scale torus network codesign. In *Proceedings of the 2nd ACM SIGSIM/PADS conference on Principles of advanced discrete simulation*, pages 27–38. ACM, 2014.
- [51] Misbah Mubarak, Christopher D. Carothers, Robert B. Ross, and Philip Carns. Using massively parallel simulation for MPI collective communication modeling in extreme-scale networks. In *Proceedings of the Winter Simulation Conference (WSC)*, 2014.
- [52] Shane Snyder, Philip Carns, Jonathan Jenkins, Kevin Harms, Robert Ross, Misbah Mubarak, and Christopher Carothers. A case for epidemic fault detection and group membership in HPC storage systems. In *Proceedings of the 5th International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS14)*. Springer, 2014.
- [53] Dongfang Zhao, Zhao Zhang, Xiaobing Zhou, Tonglin Li, Ke Wang, Dries Kimpe, Philip Carns, Robert Ross, and Ioan Raicu. Fusionfs: Towards supporting data-intensive scientific applications on extreme-scale high-performance computing systems. In *Proceedings of the 2014 IEEE International Conference on Big Data (IEEE BigData 2014)*, 2014.
- [54] Philip Carns, Yushu Yao, Kevin Harms, Robert Latham, Robert B. Ross, and Katie Antypas. Production I/O characterization on the Cray XE6. In *Proceedings of the Cray User Group meeting 2013 (CUG 2013)*, 2013.
- [55] Christopher D. Carothers, Misbah Mubarak, Robert B. Ross, Philip Carns, Jeffrey S. Vetter, and Jeremy S. Meredith. Combining Aspen with massively parallel simulation for effective exascale codesign. In *Workshop on Modeling & Simulation of Exascale Systems & Applications (MODSIM) 2013*, 2013.

- [56] Cengiz Karakoyunlu, Dries Kimpe, Philip Carns, Kevin Harms, Robert Ross, and Lee Ward. Toward a unified object storage foundation for scalable storage systems. In *Proceedings of the 5th Workshop on Interfaces and Architectures for Scientific Data Storage (IASDS 2013)*, 2013.
- [57] Sidharth Kumar, Avishek Saha, Venkatram Vishwanath, Philip Carns, John A. Schmidt, Giorgio Scorzelli, Hemanth Kolla, Ray Grout, Robert Latham, Robert Ross, et al. Characterization and modeling of PIDX parallel I/O for performance optimization. In *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis*, 2013.
- [58] Philip Carns, Kevin Harms, Dries Kimpe, Justin M Wozniak, Robert Ross, Lee Ward, Matthew Curry, Ruth Klundt, Geoff Danielson, Cengiz Karakoyunlu, et al. A case for optimistic coordination in HPC storage systems. In *7th Parallel Data Storage Workshop (PDSW 2012)*, 2012.
- [59] Dries Kimpe, Philip Carns, Kevin Harms, Justin M Wozniak, Samuel Lang, and Robert Ross. AESOP: Expressing concurrency in high-performance system software. In *Proceedings of 7th IEEE International Conference on Networking, Architecture, and Storage (NAS 2012)*, 2012.
- [60] Sidharth Kumar, Venkatram Vishwanath, Philip Carns, Joshua A. Levine, Robert Latham, Giorgio Scorzelli, Hemanth Kolla, Ray Grout, Robert Ross, Michael E. Papka, et al. Efficient data restructuring and aggregation for I/O acceleration in PIDX. In *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis*, 2012.
- [61] Ning Liu, Jason Cope, Philip Carns, Christopher Carothers, Robert Ross, Gary Grider, Adam Crume, and Carlos Maltzahn. On the role of burst buffers in leadership-class storage systems. In *Proceedings of 28th IEEE MSST conference*, 2012.
- [62] Misbah Mubarak, Christopher D. Carothers, Robert B. Ross, and Philip Carns. Modeling a million-node dragonfly network using massively parallel discrete-event simulation. In *3rd International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computing Systems (PMBS12)*, 2012.
- [63] Philip Carns, Kevin Harms, William Allcock, Charles Bacon, Samuel Lang, Robert Latham, and Robert Ross. Understanding and improving computational science storage access through continuous characterization. *Proceedings of 27th IEEE Conference on Mass Storage Systems and Technologies (MSST)*, 2011.
- [64] Jason Cope, Ning Liu, Sam Lang, Philip Carns, Chris Carothers, and Robert Ross. CODES: Enabling co-design of multi-layer exascale storage architectures. In *Proceedings of the Workshop on Emerging Supercomputing Technologies 2011*, 2011.
- [65] Adam Crume, Carlos Maltzahn, Jason Cope, Sam Lang, Rob Ross, Phil Carns, Chris Carothers, Ning Liu, Curtis Janssen, John Bent, et al. Poster: FLAMBES: Evolving fast performance models. In *Proceedings of the 2011 companion on High Performance Computing Networking, Storage and Analysis Companion*, pages 31–32. ACM, 2011.
- [66] Sidharth Kumar, Venkatram Vishwanath, Philip Carns, Brian Summa, Giorgio Scorzelli, Valerio Pascucci, Robert Ross, Jacqueline Chen, Hemanth Kolla, and Ray Grout. PIDX: Efficient parallel I/O for multi-resolution multi-dimensional scientific datasets. In *Cluster Computing (CLUSTER), 2011 IEEE International Conference on*, pages 103–111. IEEE, 2011.
- [67] Ning Liu, Christopher Carothers, Jason Cope, Philip Carns, Robert Ross, Adam Crume, and Carlos Maltzahn. Modeling a leadership-scale storage system. In *9th International Conference on Parallel Processing and Applied Mathematics 2011 (PPAM 2011)*, 2011.

- [68] Philip Carns, Robert Ross, and Samuel Lang. Object storage semantics for replicated concurrent-writer file systems. In *Proceedings of 2010 Workshop on Interfaces and Architectures for Scientific Data Storage (IASDS 2010)*. IEEE, 2010.
- [69] Sidharth Kumar, Valerio Pascucci, Venkatram Vishwanath, Philip Carns, Mark Hereld, Robert Latham, Tom Peterka, Michael E Papka, and Robert Ross. Towards parallel access of multi-dimensional, multi-resolution scientific data. In *5th Petascale Data Storage Workshop (PDSW), 2010*, pages 1–5. IEEE, 2010.
- [70] Seung Woo Son, Samuel Lang, Philip Carns, Robert Ross, Rajeev Thakur, Berkin Ozisikyilmaz, Prabhath Kumar, Wei-Keng Liao, and Alok Choudhary. Enabling active storage on parallel I/O software stacks. In *2010 IEEE 26th Symposium on Mass Storage Systems and Technologies (MSST)*, pages 1–12. IEEE, 2010.
- [71] Nawab Ali, Philip Carns, Kamil Iskra, Dries Kimpe, Samuel Lang, Robert Latham, Robert Ross, Lee Ward, and P Sadayappan. Scalable I/O forwarding framework for high-performance computing systems. In *IEEE International Conference on Cluster Computing and Workshops, 2009. CLUSTER'09*, pages 1–10. IEEE, 2009.
- [72] Philip Carns, Sam Lang, Robert Ross, Murali Vilayannur, Julian Kunkel, and Thomas Ludwig. Small-file access in parallel file systems. In *IEEE International Symposium on Parallel & Distributed Processing, 2009. IPDPS 2009*, pages 1–11. IEEE, 2009.
- [73] Philip Carns, Robert Latham, Robert Ross, Kamil Iskra, Samuel Lang, and Katherine Riley. 24/7 characterization of petascale I/O workloads. In *Proceedings of 2009 Workshop on Interfaces and Architectures for Scientific Data Storage*. IEEE, 2009.
- [74] Samuel Lang, Philip Carns, Robert Latham, Robert Ross, Kevin Harms, and William Allcock. I/O performance challenges at leadership scale. In *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis*, page 40. ACM, 2009.
- [75] Sumit Narayan, John A Chandy, Samuel Lang, Philip Carns, and Robert Ross. Uncovering errors: The cost of detecting silent data corruption. In *Proceedings of the 4th Annual Workshop on Petascale Data Storage*, pages 37–41. ACM, 2009.
- [76] Philip H Carns, Bradley W Settlemyer, and Walter B Ligon. Using server-to-server communication in parallel file systems to simplify consistency and improve performance. In *SC'08: Proceedings of the 2008 ACM/IEEE Conference on Supercomputing*, pages 1–8. IEEE, 2008.
- [77] Philip Carns, Walter Ligon III, Robert Ross, and Pete Wyckoff. BMI: A network abstraction layer for parallel I/O. In *Parallel and Distributed Processing Symposium, 2005. Proceedings. 19th IEEE International*, pages 8–pp. IEEE, 2005.
- [78] Murali Vilayannur, Robert B Ross, Philip H Carns, Rajeev Thakur, Anand Sivasubramaniam, and Mahmut Kandemir. On the performance of the POSIX I/O interface to PVFS. In *Proceedings, 12th Euromicro Conference on Parallel, Distributed and Network-Based Processing, 2004*, pages 332–339. IEEE, 2004.
- [79] Philip H Carns, Walter B Ligon III, Robert B Ross, and Rajeev Thakur. PVFS: A parallel file system for Linux clusters. In *Proceedings of the 4th annual Linux Showcase & Conference-Volume 4*, pages 28–28. USENIX Association, 2000.

- [80] Philip Carns, Walter B. Ligon III, Scott P. McMillan, and Robert B. Ross. An evaluation of message passing implementations on Beowulf workstations. In *Proceedings, IEEE Aerospace Conference, 1999*, volume 5, pages 41–54. IEEE, 1999.

## Technical reports

- [1] Philip Carns, Julian Kunkel, Kathryn Mohror, and Martin Schulz. Understanding I/O behavior in scientific and data-intensive computing (dagstuhl seminar 21332). In *Dagstuhl Reports*, volume 11. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021.
- [2] Shane Snyder, Philip Carns, Kevin Harms, Robert Latham, and Robert Ross. Performance evaluation of Darshan 3.0.0 on the Cray XC30. Technical report, Argonne National Laboratory (ANL), 2016.
- [3] Philip Carns. ALCF I/O data repository. Technical Report ANL/ALCF/TM-13/1, Argonne National Laboratory, 2013.
- [4] Philip Carns, Kevin Harms, Robert Latham, and Robert B. Ross. Performance analysis of Darshan 2.2.3 on the Cray XE6 platform. Technical Report ANL/MCS-TM-331, Argonne National Laboratory, 2012.

## Other publications

- [1] Bradley Settlemyer, George Amvrosiadis, Philip Carns, and Robert Ross. It’s time to talk about HPC storage: Perspectives on the past and future. *Computing in Science & Engineering*, 23(6):63–68, 2021.
- [2] Robert Ross, Lee Ward, Philip Carns, Gary Grider, Scott Klasky, Quincey Koziol, Glenn K Lockwood, Kathryn Mohror, Bradley Settlemyer, and Matthew Wolf. Storage systems and I/O: Organizing, storing, and accessing data for scientific discovery. Technical report, USDOE Office of Science (SC)(United States), 2019.
- [3] Rob Latham, Neil Miller, Robert Ross, and Phil Carns. A next-generation parallel file system for Linux clusters. *LinuxWorld Magazine*, 2, 2004.
- [4] Neill Miller, Robert Latham, Robert B Ross, and Philip Carns. Improving cluster performance with PVFS2. *ClusterWorld Magazine*, 2(4), 2004.

## Invited talks

- [1] Philip Carns. Understanding and adapting to change in data-intensive scientific computing, April 2022. PhD Opponent presentation at the KTH Royal Institute of Technology.
- [2] Philip Carns. Can we gamify I/O performance?, August 2021. Dagstuhl Seminar 21332: Understanding I/O Behavior in Scientific and Data-Intensive Computing.
- [3] Philip Carns. BYOFS: The opportunities and dangers of specialization in the age of exascale data storage, March 2019. SOS23 Workshop.
- [4] Philip Carns. Understanding and tuning HPC I/O: How hard can it be?, June 2018. keynote presentation at the 4th annual HPC I/O in the Data Center Workshop (HPC-IODC) and Workshop on Performance and Scalability of Storage Systems (WOPSSS).
- [5] Philip Carns. Incorporating NVM into data-intensive scientific computing, May 2018. 34th International Conference on Massive Storage Systems and Technology (MSST 2018).

- [6] Philip Carns. Building blocks for user-level HPC storage systems, May 2017. Dagstuhl Seminar 17202: Challenges and Opportunities of User-Level File Systems for HPC.
- [7] Philip Carns. Characterizing HPC I/O: from applications to systems, April 2017. ZIH Colloquium at Technische Universität Dresden.
- [8] Philip Carns. TOKIO: Using lightweight holistic characterization to understand, model, and improve HPC I/O performance, March 2017. SIAM Conference on Computational Science and Engineering.
- [9] Philip Carns. Mochi: composable lightweight data services for HPC, November 2016. 6th Joint Laboratory on Extreme Scale Computing (JLESC).
- [10] P. Carns. I/O characterization of large-scale applications with Darshan, February 2014. The SciDAC Institute of Scalable Data Management, Analysis and Visualization All Hands Meeting, Atlanta GA.
- [11] P. Carns. The storage perspective: How can HPC and HTC storage systems converge?, September 2012. Panel: Challenges and Issues of Scaling both High Throughput and High Performance Computing in Clusters, IEEE Cluster 2012.
- [12] P. Carns. HEC I/O measurement and understanding, August 2011. I/O Measurement and Understanding Panel, HEC FSIO 2011 Workshop.
- [13] R. Ross, W. Ligon, P. Carns, R. Latham, and N. Miller. PVFS birds of a feather session, November 2003. Supercomputing 2003, Phoenix AZ.
- [14] P. Carns. The Parallel Virtual File System: Overview and usage, 2001. NASA Goddard Space Flight Center.

### **Seminars and colloquia**

- [1] Philip Carns. HPC I/O principles, August 2018-2021. Argonne Training Program on Extreme-Scale Computing (ATPESC).
- [2] Robert Latham and Philip Carns. HPC I/O for computational scientists, August 2015-2017. Argonne Training Program on Extreme-Scale Computing (ATPESC).
- [3] Philip Carns and Julian Kunkel. Analyzing parallel I/O (BOF), November 2017. International Conference for High Performance Computing, Networking, Storage and Analysis.
- [4] Philip Carns. Characterizing data-intensive scientific applications with Darshan, June 2017. CS/NERSC Data Seminar, National Energy Research Scientific Computing Center.
- [5] Philip Carns and Julian Kunkel. Analyzing parallel I/O (BOF), November 2016. International Conference for High Performance Computing, Networking, Storage and Analysis.
- [6] P. Carns. Adapting to change: observation, simulation, and experimentation in high-performance I/O research, January 2016. Sandia National Laboratories.
- [7] Julian Kunkel and Philip Carns. Analyzing parallel I/O (BOF), November 2015. International Conference for High Performance Computing, Networking, Storage and Analysis.
- [8] Julian Kunkel, Philip Carns, and Alvaro Aguilera. Analyzing parallel I/O (BOF), November 2014. International Conference for High Performance Computing, Networking, Storage and Analysis.

- [9] Yushu Yao and Philip Carns. Darshan – I/O workload characterization in MPI applications (tutorial), October 2014. IEEE International Symposium on Workload Characterization.
- [10] P. Carns. Parallel file systems: From research to deployment, November 2007. Mathematics and Computer Science Division Seminar, Argonne National Laboratory.
- [11] P. Carns. The Parallel Virtual File System, August 2002. PARL Fall Research Summit, Clemson University, Clemson, SC.

## Software developed

- Darshan: I/O characterization tool. Original principal designer and developer. <http://www.mcs.anl.gov/research/projects/darshan/>
- Mochi (composable data services): Overall technical lead. Principal developer of the following components:
  - Margo (Mochi’s core RPC processing mechanism, combining the Argobots user-level threading framework and the Mercury RPC framework): <https://github.com/mochi-hpc/mochi-margo>
  - abt-io (concurrent I/O processing): <https://github.com/mochi-hpc/mochi-abt-io>
  - libch-placement (consistent hashing): <https://github.com/mochi-hpc/mochi-ch-placement>
  - Quintain (parameterized workload measurement): <https://github.com/mochi-hpc/mochi-quintain>
  - Bake (bulk data storage): <https://github.com/mochi-hpc/mochi-bake/>
- CODES: framework for discrete event simulation of large-scale storage systems. Developer. <http://www.mcs.anl.gov/projects/codes/>
- Triton: prototype next generation storage system. One of the principal developers. <http://www.mcs.anl.gov/projects/triton/>
- Aesop: programming language for high-concurrency services. Developer. <http://www.mcs.anl.gov/projects/aesop/>
- Buffered Message Interface (BMI): HPC network abstraction layer used by the PVFS, IOFSL, and Mercury projects. Principal developer. <https://github.com/radix-io/bmi/>
- Parallel Virtual File System (PVFS): HPC file system. One of the principal developers. [https://en.wikipedia.org/wiki/Parallel\\_Virtual\\_File\\_System](https://en.wikipedia.org/wiki/Parallel_Virtual_File_System)