ILLINOIS TECH

THE OHIO STATE UNIVERSITY

StoreHub:

Building a Community for the Future of Data Storage Research

Anthony Kougkas, Suren Byna, Xian-He Sun

Special thanks to the National Science Foundation for supporting this initiative

Agenda

- Challenges and limitations of data storage research
  - Case studies as anecdotes
- The StoreHub initiative
- Opportunities and collaboration
- Community feedback
- StoreHub challenges and potential solutions

What is StoreHub?

- StoreHub is a community-driven hardware and software infrastructure, funded by the National Science Foundation (NSF), tailored for data management and storage research.
- Designed to facilitate R&D across various aspects of data management, storage, and I/O:
  - data formats,
  - storage systems,
  - data quality,
  - data reduction,
  - data encoding,
  - in-situ analysis,
  - I/O and OS integration,
  - data consistency semantics,
  - data movement and sharing,
  - metadata standards,
  - provenance,
  - ontologies, and
  - full-stack application development.

Challenges in Data Storage Research

- Complex Landscape:
  - Introduction of AI/ML, big data, and experimental data processing diversifies storage software, demanding support for heterogeneous environments.
- Optimization Challenges:
  - Modern systems must accommodate diverse hardware and software, complicating data movement and performance optimization.
- Research Environment Needs:
  - Accurate performance assessment in storage research often requires isolated environments to prevent interference, yet such setups are rarely available.
- Hardware Customization:
  - Effective validation of storage solutions in research necessitates access to customizable hardware capable of emulating various computational architectures.
- Early Access to Innovations:
  - Cutting-edge storage research depends on early access to programmable and experimental hardware, often restricted by market availability.

Infrastructure Limitations

- Design Constraints:
  - Many infrastructures are tailored to specific scientific applications, not specialized storage testing, leading to suboptimal outcomes.
- Resource Scarcity:
  - Specialized storage devices and configurations crucial for advanced research are often scarce.
- Access to Key Resources:
  - Essential resources, such as hierarchical storage nodes, are highly demanded and typically have restricted access, constraining research flexibility.
- Rigid Systems:
  - Numerous systems are fixed in their setup and utilize outdated storage mediums, hindering the evaluation of contemporary storage solutions.
- Impact of Shared Resources:
  - Variables introduced by shared resources in cloud systems or university clusters can distort storage performance experiment results.

Anecdotes from storage researchers

1. Overcoming Access Limitations
2. Hardware Emulation for Research
3. Navigating Hardware Availability Issues
4. Real-Time I/O and Compute Heavy Operations
5. Facilitating Advanced GPU Data Transfers
6. Configuring Storage Solutions

Case#1: Overcoming Access Limitations

- **Context**: Root access necessity on testbed clusters for specific I/O performance tests.
- **Challenge**: Unable to directly clear OS caches due to restricted root access, affecting data freshness for each test phase.
- **Solution**: Negotiated with system administrators to allow limited root commands for clearing OS caches, ensuring data integrity during performance evaluations.
- ***Need****: flexible access controls on experimental systems to support varied research requirements.*

Case#2: Hardware Emulation for Research

- **Context**: Studying data transfer impacts across heterogeneous accelerators in multi-tenant environments.
- **Challenge**: Lack of access to intended hardware like SmartNICs and Computational Storage Devices; required root access to modify CPU settings for emulation.
- **Solution**: Emulated the desired environment using NUMA nodes and modified CPU core frequencies under limited user permissions.
- ***Need****: Close the gap between available and ideal research environments, emphasizing the need for accessible, real hardware for valid results.*

Case#3: Navigating Hardware Availability Issues

- **Context**: Investigating the effectiveness of modern storage accelerators to enhance memory capacity.
- **Challenge**: Required hardware (Compute Express Link devices) was unavailable, necessitating the use of alternative NVMe and SSD setups for emulation.
- **Solution**: Created a simulated environment that approximates CXL capabilities using available resources.
- ***Need****: Research infrastructures that provide access to the latest technologies to foster accurate and innovative research outcomes.*

Case#4: Real-Time I/O and Compute Heavy Operations

- **Context**: Exploring real-time I/O and compute-heavy operations on intermediate storage nodes.
- **Challenge**: Severe restrictions on accessing and interacting with I/O and storage nodes using research-level code.
- **Solution**: Limited evaluation scale by emulating I/O nodes as less powerful compute nodes.
- ***Need****: Broader access and fewer restrictions on HPC environments to advance storage and I/O research.*

Case#5: Facilitating Advanced GPU Data Transfers

- **Context**: Accelerating model data transfer between training and inference systems across different nodes.
- **Challenge**: Lack of permissions to install necessary drivers and enable kernel modules for GPUDirect RDMA.
- **Solution**: Relied on system administrators to make required updates, though faced significant delays.
- ***Need****: User-level access to system configurations to reduce dependency on administrative support and speed up research processes.*

Case#6: Configuring Storage Solutions

- **Context**: Testing various storage configurations to identify optimal setups for different I/O patterns.
- **Challenge**: Restricted abilities to deploy and configure storage solutions on supercomputing platforms.
- **Solution**: Utilized administrative rights on certain clusters to experiment with storage configurations while adhering to necessary security protocols.
- ***Need****: More open policies regarding storage service deployment on supercomputers to enable comprehensive and autonomous research testing.*

StoreHub Initiative

Hardware

- **Storage Mediums**: a variety of storage mediums per node, such as PMEM, NVMe SSD, SATA SSD, and traditional HDDs
- **Hardware Diversity**: a heterogeneous mix of CPUs, GPUs, and other programmable devices like FPGAs,
- **Networking Capabilities**: Fast Ethernet and Infiniband network interconnections, with potential for software-defined network topologies.

Software

- **Flexible Software Management**: Plans to offer flexible software package management systems that allow researchers to install and manage their software configurations with elevated access privileges.
- **Advanced Protocols**: Support for cutting-edge protocols like DDR5, PCIe 5.0, and NVMe-oF will be included to ensure that the infrastructure remains at the forefront of technological advancements.

Services

- **Resource Management**: A cluster resource manager will provide essential services like resource allocation and isolation for experiments that target solutions for interference and contention.
- **Debugging and Telemetry**: Tools to capture the impact of solutions on code efficiency and hardware utilization.
- **User Services and Support**: The infrastructure will offer services that simplify the use and deployment of resources, and opportunities for training and development.

Why StoreHub and why now?

- Bridging the Technology Gap:
  - StoreHub will provide direct access to advanced and experimental storage technologies, filling a significant void in current research environments where access to cutting-edge hardware is often delayed.
- Customized Research Environment:
  - It will offer a unique, customizable infrastructure that supports isolated and controlled testing environments, allowing researchers to accurately measure and optimize the performance of new storage solutions without interference.
- Community Collaboration:
  - StoreHub aims to create a centralized platform for the storage research community, facilitating collaboration, idea exchange, and synergistic research across diverse domains.
- Future-Ready Infrastructure:
  - Designed to adapt to rapid technological changes, StoreHub’s infrastructure will provide the necessary flexibility and resources to tackle emerging challenges in real-time.

Research Opportunities with StoreHub (examples)

- Hierarchical Storage Management:
  - Researchers can explore advanced data buffering techniques across different storage mediums, optimizing I/O performance in systems with deep memory and storage hierarchies.
- Real-Time Data Streaming:
  - The infrastructure will support the development and testing of high-performance data streaming engines capable of handling large data volumes in real-time, crucial for applications like live analytics and ML.
- I/O Convergence:
  - StoreHub will facilitate studies on converged storage architectures that efficiently support mixed workloads from HPC, Big Data, and AI, optimizing resource utilization and reducing data transfer overheads.
- Next-Generation Storage Systems:
  - Researchers will have the opportunity to design and test innovative storage stacks and file systems tailored for emerging technologies like persistent memory and ultra-fast network environments.
- AI-Driven Storage Optimization:
  - The platform will enable the integration of AI techniques into storage management, such as predictive data placement and automated system tuning, enhancing efficiency and reducing operational costs.

Benefits to the Research Community

- Cutting-Edge Access:
  - Immediate reach to the latest in storage tech.
- Boosted Productivity:
  - Advanced tools eliminate bottlenecks, accelerating research cycles.
- Network Expansion:
  - Forge collaborations with top-tier researchers and industry leaders.
- Skill Enhancement:
  - Workshops and training sessions will provide technical proficiency.
- Increased Visibility:
  - Elevate institutional presence in the national/global research landscape.
- Cost Efficiency:
  - Optimize infrastructure investments and operational costs.
- Innovation Platform:
  - Pioneer novel storage solutions and influence future technologies.

Future Vision and Goals

- Establishing a National Hub for Storage Research
- Catalyzing Technological Advancements
- Enhancing Research Productivity and Collaboration
- Building a Sustainable Research Community
- Democratizing Access to Cutting-Edge Tools

Project Challenges and Anticipated Obstacles

- Resource Scalability and Sustainability:
  - Ensuring that StoreHub's infrastructure can scale to meet increasing demand without compromising quality or performance.
- Technical Heterogeneity and Integration:
  - Managing a wide array of technologies and ensuring they work seamlessly together poses significant challenges.
- Community Engagement and Participation:
  - Ensuring active and sustained community involvement can be difficult, especially in a field as broad and diverse as storage research.
- Security and Data Privacy:
  - With the increasing importance of data security, ensuring the integrity and confidentiality of research data within StoreHub is crucial.
- Funding and Financial Sustainability:
  - Securing ongoing funding to maintain and upgrade infrastructure is an ever-present challenge.
- Keeping Pace with Technological Advancements:
  - The rapid pace of technological change requires continuous updates to the infrastructure, which can be resource-intensive.

Get Involved

- Community Input:
  - Engage in our surveys to shape the future of data storage technology.
- Collaborative Workshops:
  - Influence the StoreHub vision, share perspectives, and contribute to pioneering research.
- Strategic Discussions:
  - Participate in discussions and planning sessions to align research with industry needs, fostering innovative breakthroughs.
- Define Priorities:
  - As a community member, help set the direction for infrastructure and research, ensuring StoreHub meets the community's evolving needs.

Call to Action

Together, we can transform the landscape of data management and storage!

Thank you

Q&A

Stay tuned at https://grc.iit.edu/research/projects/storehub

See you at SC'24.

Contact us at grc@iit.edu

Or reach out to the PIs: akougkas@iit.edu, sun@iit.edu, byna.1@osu.edu