

Data Storage Research Vision 2025*

Report on NSF Visioning Workshop held May 30–June 1, 2018

George Amvrosiadis, Ali R. Butt, Vasily Tarasov, Ming Zhao, and participants†

<https://sites.google.com/vt.edu/data-storage-research/>

1 Introduction

There are a number of open challenges and research issues that need to be addressed both in the short and long term to ensure sustained storage systems efficacy and performance.

The wide variety of applications of modern and emerging storage systems entail that the fundamental design of storage systems should be revisited to support application-specific and application-defined semantics. Such tailored design will address many of the shortcomings of using the extant one-size-for-all generic approach that is plagued with inefficiencies both in performance and storage capacity overhead. Another aspect is to reevaluate the use of standards and APIs such as POSIX, and design new sustainable data representations and APIs to support emerging applications such as IoT. Such redesign of the storage architecture is gradually being explored. For example, systems such as key-value and object stores are being used to design application-specific solutions. Management of metadata, indexing, and high-speed transactions for small data items are key to next generation storage systems as well. Nevertheless, more new methods and techniques are needed to support a wider range of emerging applications. This prompts a more general discussion on application data management and storage system co-design.

Many important applications now arise in the space of Artificial Intelligence (AI), such as model training frameworks, and their needs are met poorly by current storage systems. New storage techniques and technologies tailored for AI workloads and AI usages of data are much needed. At the same time, AI-based methods hold great potential for building intelligent storage systems for meeting the challenging application demands and optimizing storage management as the systems become increasingly complex.

Hardware advances are further driving the way storage systems are realized. Storage hybridization and heterogeneity is now a part of most large-scale storage deployments. However, the software subsystems for supporting and using such systems are lagging, and new management systems need to be designed and invented. Similarly, in-memory storage systems and persistent memory systems are giving rise to a new class of memory-only storage solutions. While such systems can be thought of as simply a tier of a traditional storage hierarchy, there is clearly a need for innovation to leverage the unique opportunities offered by new hardware advancements and designing the storage layer to maximize efficiency and performance. At the same time, the exponential growth of cold archival and cold primary data motivates intensive research of new hardware and software architectures that can store petabytes of data at the maximum possible cost-efficiency.

*This work is supported by the NSF under the Award CNS-1829096.

†Irfan Ahmad, Remzi H. Arpaci-Dusseau, Feng Chen, Yiran Chen, Yong Chen, Yue Cheng, Vijay Chidambaram, Dilma Da Silva, Angela Demke-Brown, Peter Desnoyers, Jason Flinn, Xubin He, Song Jiang, Geoff Kuenning, Min Li, Carlos Maltzahn, Ethan L. Miller, Kathryn Mohror, Raju Rangaswami, Narasimha Reddy, David Rosenthal, Ali Saman Tosun, Nisha Talagala, Peter Varman, Sudharshan Vazhkudai, Avani Wildani, Erez Zadok, Xiaodong Zhang, Yiyang Zhang, Mai Zheng.

More and more organizations are starting to trust computing clouds to cost-efficiently store and process data at scale. Public cloud providers now manage tremendous amounts of data and continuous scalability of their infrastructures requires further innovation, both at the system software and hardware levels. Data management problems become especially challenging in hybrid clouds where the data flows need to be secure, tightly controlled, and frequently geo- or otherwise fenced. Furthermore, multi-tenancy at global scale increases workload and platform diversity to previously unseen levels which, e.g., makes the control of data access quality an extremely complicated task with many open questions. At the same time, the dawn of the age of Internet of Things (IoT) is driving the need for novel and innovative storage systems at the edge of the Internet to store, manage, retrieve, and efficiently utilize unprecedented volumes of data at increasingly faster speeds.

Because of the aforementioned open challenges and issues, there is a need for a consolidated effort to identify and establish a vision for storage systems research and comprehensive techniques that provide practical solutions to the storage issues facing the information technology community. The goal of the National Science Foundation (NSF) Visioning Workshop on Data Storage Research 2025 was to bring together leading researchers in storage and distributed systems to provide a working vision, as well as prioritization for near and long-term storage research and scientific investigations. The workshop aimed to provide opportunities for identifying both core systems research in storage and cross-cutting research in “storage+X” systems.

This one and half day workshop was held from May 31st to June 1st 2018 at IBM Almaden Research Center (ARC) a well-known industry research laboratory focusing, among other topics, on storage systems. We invited 40 researchers who are well-recognized leaders and key contributors to storage research in academia, industry, national laboratories, and federal agencies, to work together and develop a collective vision for future storage research. Every invitee was asked to provide a pre-workshop white paper on topics that she or he felt important. Based on these broad range of topics from the white papers, we divided the workshop attendees into groups to carry out focused discussions around four major themes: “Cloud, edge, and everything in between” (Section 2), “AI and Storage: Made for each other” (Section 3), “Teaching old storage new tricks” (Section 4), and “The hardware, they are a-changin” (Section 5). By the end of the workshop, each group also shared their discussion outcomes with the whole workshop to get feedback from the other groups and further develop the vision for each storage theme during the post-workshop, offline discussions. The rest of this report presents the findings on these four major themes as the result of all of the pre-, during-, and post-workshop discussions.

2 Storage for Cloud, Edge and IoT Systems

2.1 Introduction

Data-driven scientific discovery has been well acknowledged as the new fourth paradigm of scientific innovation [46]. One key enabler of the data-driven innovation is data storage systems that manage the massive data, which has witnessed a disruptive sea change in recent years. Among others, the advent of cloud computing has transformed the basic substrate for systems building in the last decade, and the long-anticipated “Internet of Things” (IoT) has led to the emergence of edge computing which extends the system boundaries pervasively.

In such a dynamic context, the depth of the storage stack and the scope of storage systems is increasing rapidly. Storage systems will need to manage data that is collected, stored, transformed, and transferred from heterogeneous edge devices to back-end cloud services, which will involve more than 18 layers [80]. There is a vast diversity (e.g., different hardware, storage options, configurations) throughout systems, and at each layer of the system, different underlying storage technologies are likely to be the best fit for the corresponding workload and demands (e.g., latency, bandwidth, cost). Moreover, there are potential gaps or miscommunications between layers and components [40, 44], which increases the difficulty of providing end-to-end guarantees and achieving the ideal tradeoffs among performance, reliability, fairness, etc. To move data storage research forward, we

summarize the research challenges and opportunities into 9 key properties which are essential for future storage systems, and call for the community’s collective efforts to build systems that are truly efficient, unified, specified, elastic, explainable, sharable, application-driven, reliable, and re-evaluable.

Besides research, the education and infrastructure must also keep pace with the evolution of the computing landscape. Unlike other related communities who have well-established “home courses” (e.g., operating systems, computer networks, databases), storage systems are mostly taught under the hood of other courses. Many talented students are not attracted to this area, partly due to the “diluted” identity. Similarly, in terms of infrastructure, we recognize that existing efforts like CloudLab [3] help, but more are needed in order to sustain the advance of this area.

The rest of the section is organized as follows: §2.2 elaborates on the research challenges and opportunities, §2.3 discusses the directions for education and infrastructure support, and §6 concludes the section.

2.2 Research Challenges & Opportunities

2.2.1 Context: world in flux

The landscape of storage systems is changing rapidly in recent years. On the hardware side, there used to be just hard disk drives and tapes, but now we are seeing exciting advances including flash-based SSDs, byte-addressable persistent memories, RDMA for storage networks, etc. Besides the superior raw performance, many of these new hardware are programmable to certain extent and thus allow more customization.

Similarly, on the software side, there used to be just file systems. Now we have sophisticated key-value stores, document stores, configuration stores, logging stores, etc., many of which are scaled out to warehouse-scale cloud systems. Moreover, virtualized cloud systems are able to manage the diverse resources (e.g., CPU cores, memories, SSDs) flexibly and transparently, usually scaling up and down in seconds [23].

While the cloud hides much of the complexity from the users and provides the desired convenience, its opaque, shared, diverse, and layered nature has led to concerns of how to provide end-to-end guarantees and achieve the ideal tradeoffs among efficiency, elasticity, explainability, shareability, reliability, etc. We summarize the research challenges and opportunities into 9 key properties which are essential for future storage systems, and elaborate on each of them in the next section.

2.2.2 What should we build

Efficient Systems

Similar to traditional systems, cloud-based systems also put great focus on efficiency. This is important for both users who pay for the usage (measured in hours/bytes/requests) and system builders who need to make profit.

However, compared to traditional systems, there are much more layers involved in cloud systems, and different layers usually require different data formats and read/write strategies to achieve the best local efficiency, which may conflict with other layers. Moreover, the diverse hardware, dynamic workloads, and the inherent multi-tenancy make achieving high efficiency even more difficult. Consequently, while efficiency is clearly important, there is little QoS achievement capability today for end-to-end efficiency.

To address the challenge, we can no longer focus on one single layer or component. Instead, novel cross-layer and holistic solutions are highly desirable for removing all excess resource allocations in different layers, saving various costs (e.g., CPUs, memories, SSDs), and achieving the overall high efficiency.

Unified systems

Modern systems are filled with diverse storage options (e.g., various file systems, SQL databases, key-value stores, logging stores, configuration stores, document stores). While each individual storage option usually provides some unique features, they often have similar functions or components to some extent (e.g., managing persistent data structures). This inherent overlapping (i.e., many physical copies of the same data across layers/components) is one of the major obstacles for building efficient systems today.

To address this challenge, we should explore the possibility of extracting the unified core components as building blocks and providing generalized solutions for various higher-level services. Also, to make different services more unifiable, we should experiment with solutions that can automatically transform the underlying representation of data, amount of resources allocated, configurations of durability and replication parameters, etc. based on the dynamic needs of workloads.

Specified systems

Current approaches to system building are too prescriptive, rigid, and error prone, which has led to various issues including downtime and data loss [44] and is not scalable for future storage.

To address the challenge, we envision that future systems and applications should be specified in terms of performance requirements (e.g., which hardware feature is needed), persistence needs (e.g., which data needs to be stored where), and so on. Particularly, the correctness should be precisely specified throughout the systems, which could potentially lead to the holy grail of verified systems that will never lose data. Along this direction, there are a number of open research questions though, including how to specify properties for the opaque cloud, how to identify the necessary properties and interfaces for each layer or component in the system, and how to specify the dynamic requirements of workloads.

Elastic systems

Different from traditional storage clusters which are built on fixed hardware resources, cloud-based systems are naturally elastic. The systems can be broken into constituent components which can be scaled up and down independently based on current demands of workloads.

We envision that the elasticity of the systems can be utilized for handling storage infrastructure tasks in addition to the workloads, which will likely improve the overall utilization and efficiency of the systems.

To make better use of the elasticity for storage, more disaggregated composable software architectures are highly desirable. Instead of today's monolithic storage and file systems, we should experiment with different building blocks (e.g., unified core components) and microservices, which can be reused across domains and improve long-term usability, reliability, etc.

Explainable systems

Current cloud-based systems are opaque to the users. Many services are provided through relatively simple interfaces, which makes it difficult for users to reason about the provenance of their data. Moreover, due to the complicated layering within the cloud, it is also difficult for system builders or administrators to explain the abnormalities of system behaviors observed.

We envision that future systems should provide detailed provenance information and allow for more explanations (e.g., how was data object created, who can access what and why). This will be helpful for improving security (e.g., how the information is leaked), reliability (e.g., how the data is corrupted), performance (e.g., why this run is slow), etc.

Sharable systems

Unlike the first-generation cloud technologies which only run a single or a few applications for one entity (person, organization), multi-tenancy is a new reality in modern cloud-based systems.

We believe one fundamental demand of multi-tenancy is effective sharing. However, achieving effective sharing is non-trivial as it involves many other aspects of systems. For example, from efficiency's perspective, multi-tenancy may cause interference among different workloads at different layers of systems, and thus violate QoS SLO for one or more users. Similarly, other security and privacy concerns (e.g., side-channel attacks, information leakage) need to be addressed more carefully in the multi-tenant environment in order to provide effective sharing.

Application-driven systems

One major driving force of systems research is new application needs. There are many interesting new applications arising recently (e.g., videos from cameras at soccer match, Pokemon Go), which place new demands on storage systems (e.g., real-time processing).

Given the vast diversity of applications, it is inefficient and impractical to build a highly specialized storage system for each application. Instead, we should avoid wheel reinvention by exploring the commonality among applications and adapt storage systems to a range of applications automatically, similar to the design principle of unified systems. One unique challenge here is how to assemble a representative application suite and metrics for learning the common characteristics and demands.

Reliable systems

As the scale and complexity of systems keep increasing, failures become the norm rather than exception [42]. Therefore, we need to design systems to deliver high performance and other desired properties in the presence of failures, which is particularly important for exabyte scale storage systems that will likely become common in 2025 [35].

Future systems need to be clearly specified, which could potentially lead to truly reliable storage that will never lose data. Existing efforts have shown that it is possible to formally specify and verify the crash consistency of one local file system built from scratch [34]. Nevertheless, how to scale the formal method to the vast majority of legacy software systems in the cloud environment remains unclear. More advanced mathematical methods and software engineering approaches (e.g., Netflix's Chaos Engineering [24]) are highly desirable.

Re-evaluable systems

One constant theme of storage research is workloads, which is important for fair comparison between systems and generating reproducible results.

Unfortunately, compared to the workloads for local storage systems, much less workloads for cloud-based systems are publicly available. There are a number of workloads which have benefited the community a lot (e.g., YCSB [36]); but as systems keep evolving, much more representative workloads are needed to drive the research further.

On the other hand, future storage systems should also be built with easy evaluation in mind to facilitate the fair comparison of design tradeoffs under the same representative workloads.

2.2.3 Edge & its impact on Cloud

Besides the cloud computing which has revolutionized the landscape of storage in the past decade, IoT is becoming a reality, bringing with it an explosion of data collection, storage, and processing demands. As predicted by Cisco, there will be over 50 billion IoT devices by 2020 [38]. The proliferation of IoT devices and the associated new demands have led to the emergence of edge computing. Essentially, the edge model places a “mini datacenter” of compute and storage resources at the network edge, closer to the end users.

Compared with cloud computing, edge computing is much less matured or standardized. For example, there are heterogeneous devices manufactured by different vendors with various capabilities and data formats. Also, the service models for IoT applications is unclear. We envision that one possible direction is the serverless computing model (e.g., AWS Lambda [1]). However, more research efforts are needed to integrate the diverse IoT devices into the current model.

Despite the heterogeneity, one common feature of all IoT devices is that they have limited hardware resource. To address the constraint, we should explore how to identify and discard unimportant data in a timely fashion, and how to balance among storage, preprocessing, and communication between IoT devices and cloud.

Different from cloud systems which can be built for various workloads and adapt to the demands on the fly, edge computing has large upfront cost to install edge nodes and limited opportunity for statistical multiplexing. In other words, edge computing needs to be designed for the right workload in the first place, and we need to identify workload demands and match them to storage capabilities precisely.

One barrier to storage research in the era of cloud-edge computing is that no edge-to-cloud holistic persistent data storage capabilities exist today. Therefore, a testbed involving both edge and cloud is highly desirable. Such a testbed will enable abundant research opportunities. For example, we can experiment with tradeoffs of latency, cost, persistence in light of edge computing applications. One reasonable assumption for this direction is that we will have increasingly larger resource budgets per unit costs as the location gets closer to the cloud.

Another barrier is the lack of agreed upon workloads and traces for evaluation and comparison of new research designs. A realistic workload trace needs to track requests to read and write data across all of devices, edge nodes, and cloud servers, including operations that transform or aggregate the data. Recent work on distributed system tracing may provide the mechanism for collecting such traces, but the research community also needs to agree on a trace format, and strategies for replaying any such traces that are collected.

2.3 Education & Infrastructure

Unfortunately, while there are many exciting research challenges and opportunities, the current education and infrastructure support seems insufficient to sustain the advance of this area.

2.3.1 Education

The first challenge is *diluted identity*. Traditionally, storage systems are mostly taught as a sub-piece of operating systems. More recently, with the popularity of cloud and big data, many distributed storage systems are covered under the hood of courses like cloud computing, distributed and parallel database systems, etc. And the novel storage hardware and devices are being incorporated into computer architectures, IoT, etc. In contrast to other related research communities who have their well-established “home courses” (e.g., operating systems, computer architectures, computer networks, databases) at many universities, few universities teach storage in the name of “storage systems”. Consequently, the identify of storage systems as one research area and community is diluted.

Another correlated challenge is *underestimated scope*. While the FAST program committee has a clear and broad definition [7]: “(we) will interpret ‘storage systems’ broadly; papers on low-level storage devices,

distributed storage systems, and information management are all of interest”, perhaps few students can interpret storage systems in this way. Many students may only connect storage systems to hard disk drives, which is obviously less attractive compared to self-driving cars. This situation is partly due to the fact that there is no clearly defined course on storage systems in the majority of universities.

We envision several directions to address the challenges. First, we need to increase the number of academics working in storage systems, perhaps driven by increasing funding support from NSF and other agencies. Second, we should increase the availability of ready-made courses, especially those from renowned universities, so that more researchers and educators would follow up. Last but not the least, we should increase the availability of educational materials. Some existing efforts (e.g., OSTEP [21] and SNIA tutorials [6]) have helped a lot, but more efforts are needed to move the community forward.

2.3.2 Infrastructure

As mentioned in §2.2.3, one challenge for research in this area is building a research testbed with a rich diversity of storage resources to enable experimentation with different tradeoffs. Some existing efforts like CouldLab [3] and Chameleon [2] have helped facilitating cloud research, which is particularly valuable for small universities where local resources are very limited. Nevertheless, as the cloud-edge systems keeps evolving, the scale needed for realistic testing may be out of reach for many academic researchers, so more large-scale testbeds is highly desirable. Moreover, it would be helpful to provide standard services (e.g., S3, Lambda) besides the testbed to increase the fidelity.

On the other hand, we envision a number of potential issues regarding shared testbeds. First, chaos of shared testbeds is not great for science. It may be difficult to reproduce results in the shared cloud storage environment. Second, storage necessarily requires persistent data, and some workloads may lead to early wearout of some storage technology. Both factors make it harder to share a storage research testbed, compared to similar shared infrastructure for networking or computing research. Therefore, techniques for building truly sharable and re-evaluable (§2.2.2) systems are urgently needed.

2.4 Summary

The opaque, shared, diverse, and layered nature of cloud-based systems leads to both grand challenges and opportunities for the next generation of storage systems. We summarized the challenges and opportunities into 9 key properties, and call for the research community’s collective efforts to build systems that are truly efficient, unified, specified, elastic, explainable, sharable, application-driven, reliable, and re-evaluable. Besides, we discussed the rising of edge computing and its impact on storage. Finally, we argued that more investment in education and infrastructure is needed in order to sustain the advance of data storage research for 2025.

3 AI and Storage

3.1 Introduction

In this section we cover the intersection of AI and Storage and outline key research areas for storage in the context of AI applications, workloads and trends. We explore the intersection of AI and Storage as follows. We begin with a perspective on each (AI and Storage) independently, and then describe how each can benefit the other. Storage for AI focuses on how storage research can drive systems designs to better serve AI workloads and AI usages of data. AI for Storage focuses on how storage systems can be improved via internal application of AI techniques. Finally, we also cover how broad education and research initiatives can bring these two fields together and cross pollinate ideas and collaboration across systems research and data science.

3.1.1 What is AI?

Artificial Intelligence (AI) is an umbrella term that covers a range of techniques to mimic human intelligence via machines and programs. AI includes the broad discipline of Machine Learning which carries within it a range of algorithmic techniques from statistical approaches to neural network (Deep Learning) approaches and beyond. At a very high level, Machine Learning refers to the ability to model and extract patterns from data or observations (training) and subsequently use these models to make predictions on new observations (inference).

While AI and Machine Learning have existed as fields for many decades, the last 5-10 years have witnessed an exponential growth in the development and application of AI. Today, virtually all commercial industries are either applying or planning to apply AI techniques to enhance their respective disciplines. This shift is driven by several factors:

The Data: Devices ranging from sensors to robots are generating increasing amounts of data and increasingly richer data (ranging from simple value time series to images, sound and video). While the data itself is valuable – its ultimate benefit to a business’ bottom line comes from the analytics that extract the insights hidden within. While simple datasets (such as streams of individual values) can be analyzed via database queries or complex event processing techniques, the increasing richness of data (multiple correlated mixed type streams, images, sound, video) requires more complex Machine Learning (ML) and Deep Learning (DL) approaches. The increased volumes of data also enable ML/DL algorithms to achieve peak efficiency.

The Compute: The ubiquity of high performance commodity computing, driven by both massive core count increases in individual CPUs and low-cost cloud computing services, have made it possible to match data growth with similarly scalable ML and DL capabilities. Hardware innovations such as TPUs, GPUs, custom FPGAs, and instruction set support in modern CPUs have further improved ML algorithms performance, making it practical to train using massive datasets.

The Algorithms: The availability of open source algorithms for ML and DL via libraries for analytic engines like Spark, TensorFlow, Caffe, NumPy, Scikit-learn, just to name a few [13–15, 84]. With these packages, a massive range of algorithmic techniques are now available in the Data Scientist sandbox. With open source, even the most newest state of the art algorithms in research are frequently publicly available to test, tune and use, nearly as soon as they are invented.

The confluence of these three factors has fueled AI growth, and in turn will drive the need for combined storage and AI research.

3.1.2 What is storage in 2025?

The term “storage” may carry many connotations. Therefore we wish to define it first to ensure a common definition. We consider storage across two time scales: short time frame (next 4–5 years) and longer time scale (5–10 years or more).

For more than 60 years, the storage field was dominated by one type of technology, the magnetic spinning media (HDD). We believe that was an anomaly and that the future will be dominated by multiple, heterogeneous storage technologies and their combinations. There are many variants of Flash based devices, storage class memories and NVMs, and even the HDD has been recast into several flavors of shingled drives (SMRs); furthermore, there is active research on using glass and DNA as storage media. Hence, the time is *right now* to investigate these ever-changing storage technologies, when they are still relatively new and yet reasonably well understood.

We feel that future storage systems will most likely include a combination of two or more technologies. And, as NVMs begin to rival DRAM speeds, both will figure heavily in the storage stack. Storage technologies, both volatile and non-volatile, may be arranged as in a tier or hierarchy, or any form of hybrid, cluster, or graph-like arrangement. Data may be cached or replicated in one or more tiers, or uniquely stored once in a single

tier. Various policies will control when data is moved, copied, or cached between those tiers. We expect storage technologies to continue to evolve and change rapidly in the coming years, so an effective storage system has to be able to adapt to those changes quickly and efficiently (e.g., to optimize throughput, latency, energy, dollar cost, or any combination or complex utility function).

Storage technologies are therefore likely to be more complex in the future, to support growing needs of big data and AI workloads, and really any workload that users need to store and process optimally. This complexity will demand support for different APIs at different levels. We expect to continue to see healthy use of block-level, file-level (e.g., POSIX), object, and key-value stores—and likely combinations thereof. There is a need for high-level, easy-to-use APIs that hide much of the internal complexity from users and developers; conversely, there is also a need to allow advanced (“power”) users access to lower-level APIs, to enable more effective optimizations. The key to the design of future storage systems and their APIs would be that they must be easy to use and logical for AI application developers and *at the same time* provide optimal storage at the lower levels (for any utility or cost function).

3.2 Storage for AI

The emerging AI field presents several trends that intersect with storage, where targeted storage research can benefit AI usages and AI applications:

- **Massive datasets:** AI workloads require the ingestion, preprocessing and ultimately analysis of massive amounts of data. Multiple stages exist in typical AI pipelines, from data ingestion (ETL - Extract, Transform, Load), to pre-processing (feature engineering, data wrangling, data cleansing/transformation, etc.) to the ultimate execution of an AI algorithm in its training or inference phases. All of these can benefit from storage optimizations for performance and data management.
- **AI stages awareness:** Storage that is aware of the distinct stages of AI processing can optimize AI pipelines via techniques such as caching of intermediate results, tracking of lineage, provenance, and checkpointing [39, 55, 63, 74].
- **Compute Architecture and Data Optimization:** AI platforms follow distinct distributed computation architecture patterns (such as data parallel and model parallel) [50, 84]. Memory hierarchy and data layout design for such computation patterns should also be a focus for future storage research. APIs that express the data access intent of an AI algorithm can also be a powerful tool to integrate memory hierarchy and data layout optimizations with AI computations.
- **Unique Characteristics:** AI algorithms have unique characteristics (such as tolerance to small amounts of data loss, very structured access patterns, ability to use and extrapolate from lossy compression) that can be exploited to create efficient storage designs.
- **Access and Distribution Characteristics:** Emerging access methods and characteristics associated with AI workloads, such as streaming processing [9] or edge storage [12], also create unique challenges that future Storage research should focus on.
- **Security and Compliance:** The use of AI brings a new dimension to data security. As industries and users demand that decisions made by AI algorithms be reproducible, transparent, and explainable, pressure builds on enterprises to put in place data management mechanisms to govern what data is and should be used to generate AI models and consequent insights [8, 10, 11, 28, 41].

3.3 AI for Storage

Machine learning techniques should be researched to improve storage systems in respects of reliability, availability, and quality of service. The large quantity of available storage system historical access data will allow machine learning algorithms to be trained. Insights can be gained out of the training and thus used to help design or optimize storage systems:

- **Performance and Placement Optimizations:** Machine learning algorithms can be applied to predict popular data and application patterns, which helps improve various storage techniques, including tiered cache, prefetching, and resource provisioning. Adapting caching policy using online learning can have significant benefits. For example, a recent paper [82] shows that using ML techniques to select between LRU and LFU replacement policies resulted in a significantly improved total cache hit rate. The key takeaway is that machine learning techniques are valuable for solving online optimization problems such as caching with the caveat that primary knobs of control be orthogonal. We believe that machine learning can be applied with success for other problems such as non-datapath server-side caches [29, 47, 51, 53, 61, 70, 73], distributed storage caches such as in hyper-converged systems, dynamic multi-tiered and hybrid storage systems [43, 83], and DRAM - persistent memory hybrid systems.
- **Failure Prediction:** Failure or error patterns in large storage systems, such as disk failures and silent data corruptions, can be predicted using machine learning techniques and correspondingly pre-cautious measures can be taken to avoid errors being proliferated. For example, proactively replacing disks which are predicted to fail soon can avoid the cost of data loss or data rebuild.
- **Storage Tuning:** Storage systems typically evolve to have a large number of tunable parameters; parameters range from hardware composition to tiering thresholds to cache sizes. Using learning techniques to advise administrators who build and maintain storage systems under dynamic workloads on optimal parameter values could improve system performance and cost for a given workload.
- **Change Detection and Anomaly Detection:** Part of tuning for workloads is understanding when they change, both temporarily and permanently. Anomaly detection has been an application area for machine learning techniques for over twenty years [52], and many techniques from these fields will likely translate to storage with little modifications. Goals of these techniques are to recognize change and assist with either debugging or re-tuning, as applicable.
- **Intelligent Storage Devices within Storage Systems:** Computing storage devices refer to storage devices with computing capability and they typically come along with artificial intelligence. On one hand, computing storage devices assume part of the storage-related computing functionality so that the running storage system is obviated from excessive overheads. Therefore, the storage system can deliver improved performance. On the other hand, with more intelligent devices, researchers need to determine what intelligence is appropriate to be offloaded to the device and propose techniques at the storage system level to achieve best synergy.

The key challenge in this domain (using AI for Storage) is that training data will often be limited before decisions have to be made. For instance, systems to store and quickly process data in self-driving cars must be extent and fast even **before** enough data can be collected for automated system design. Similarly, as storage needs shift over time in an organization, there may not be enough training data to predict how best to deal with changing priorities when reconfiguring factors such as parameters, tiers, placement, and layout. Tuning may also be improved by consideration of cost models outside of the standard throughput and latency optimizations of the bulk of storage research. Dollar cost, complexity, and power consumption came up as potential other reward functions in a multi-objective optimization scheme.

3.4 Benchmarks and Workloads

Since AI techniques are heavily data dependent, any strategy for driving AI and Storage research needs to factor in the need for publicly accessible datasets and benchmarks. Public datasets exist in ML (example [18]) but are in many cases too small to drive meaningful storage access patterns. In this section we describe some challenges that will need to be overcome to drive expansive research into the Storage and AI opportunities presented above.

3.4.1 Dataset generation and collection

We need some systematic and sustainable schemes to generate and collect datasets, including: synthetic data generation of machine learning workloads; datasets from simulations; and dataset gathered from prior NSF projects; long-term data collection and dissemination via some community infrastructure such as NSF CRI.

3.4.2 Characterizing workloads across layers

How to benchmark and characterize workloads from different layers including application, middleware/ system, and storage device layers is challenging and worthy of investigation.

3.4.3 Workload classification

Classifying workloads is a long-studied problem with a significant amount of prior work [16, 62, 72]. As new storage platforms and applications are developed, there is a need to understand, in a way that is precise and communicable across different industries, what the storage workload looks like. Given that a storage workload is a time series of operations, there are a variety of unsupervised as well as supervised machine learning techniques that we can apply to partition and categorize this space.

We would use machine learning techniques to improve workload characterization in areas including:

- Quantify similarity between workloads
- Track changes in how a workload functions on a given architecture
- Learn mixes of customer workloads on shared storage systems

3.5 Research and Education

To foster the collaboration between data science and storage systems, we need to equip students in storage systems with data science knowledge and vice versa. Nowadays, students are not well motivated to learn cross disciplinary subjects. If high quality learning materials of data science are easily accessible from workshops, tutorials in systems conferences and courses on massive open online course(MOOC), it would spark interests and encourage students to learn data science. Organizers of workshops should advocate research that applies data science methodologies to solve problems of storage systems whenever applicable. Tutorials on data science related to systems are beneficial for students, researchers and practitioners to quickly gain state of the art methods and techniques of data science. Moreover, courses on MOOC nourishes the community to embrace data science with low cost.

While learning fundamental concepts and techniques of data science is crucial to incubate the cross disciplinary research, it is equally important to provide compute resources allowing students to practice data science since it is not easy for universities and academia to provide compute platforms that are powerful enough to solve interesting problems on large-scale data. NSF cloud is one of such early efforts to meet the demand.

There are also open source projects and contests held by conferences that offer opportunities for students to apply learned data science technique to solve real world problems which in return deepen the understanding of learned knowledge.

Besides training students, engineers and researchers to learn data science techniques, another great approach to nourish research in the cross disciplinary field is to boost the collaborations between data scientists and system researchers. By combining the strength from both sides, AI systems and storage systems can be made much more reliable and scalable.

4 Rethinking Storage Systems Design

4.1 Introduction

A number of trends suggest that we need to fundamentally rethink the design of storage systems, as well as the education pipeline that produces storage researchers:

- The Internet of Things (IoT) and exascale High Performance Computing (HPC) clusters are both on track to produce a torrent of data far greater than the storage and network capacity of current systems.
- The rapid growth of data science is introducing new workloads with unique storage access patterns and performance demands.
- Current storage systems with minimal schemas are poorly equipped to organize huge amounts of data, especially when the growth in data size is projected to be exponential in nature.
- Emerging storage technologies, such as DNA storage and storage-class memory, require rethinking the entire storage hierarchy from applications to hardware.
- Privacy and security increasingly require methods to reason about relationships among data being stored, as well as the provenance and lineage of that data.
- Storage faces intense competition for top talent from other areas of computer science, creating a challenge in maintaining a high-quality pipeline of researchers and engineering professionals.

In response to these trends, storage researchers will need to evolve both current storage systems and the education pipeline. We identify five vital areas where research is needed:

1. Storage systems should allow far greater introspection into their operation and the data they store. We must develop methods to automatically use data from introspection to improve storage operation and improve data organization and management. Directions for future research in this area are covered in Section 4.2.
2. Storage systems should more tightly integrate computation (e.g., indexing, aggregation, transformation) with data generation and movement through the storage stack, effectively enabling “in-situ” and “in-transit” processing of data. The type of computation and the layer at which it is performed should adapt dynamically in response to changes in workloads and resources. Directions for future research in this area are covered in Section 4.3.
3. We should reconsider the fundamental design of the POSIX interface to support emerging storage technologies and use cases. This includes widening the interface, allowing more application-specific customization of storage behavior, and supporting evolution of data and technologies. Directions for future research in this area are covered in Section 4.4.

4. We should enable tighter co-design of applications and storage. This includes supporting “in-vivo” storage development, allowing evolution of storage systems and A/B testing, as well as rethinking where the “intelligence” of flash storage should lie in the storage stack. Directions for future research in this area are covered in Section 4.5.
5. We should grow the pipeline of storage researchers and developers. Proposed ideas include: enabling hands-on classroom experiences in which students see the benefits of changing storage systems, tighter integration with data science in the curriculum, and more emphasis on applications when teaching storage topics. These topics are covered in detail in Section 4.6.

4.2 Introspection, provenance, and metadata

4.2.1 Introspection

To achieve maximal effectiveness, both storage systems researchers and storage systems themselves need to understand the detailed behavior of applications, storage devices and the entire storage stack in between. To do that, we must be able to observe and correlate those behaviors both inside and outside the system.

There are currently many techniques for examining storage systems, some *ad hoc* and some semi-standardized [19, 22, 33, 64, 79]. However, future storage systems will demand improved approaches:

- The output of existing tools is generally designed to be consumed by humans rather than by programs, making the tools unsuitable for big-data analysis or for use in adaptive systems.
- Multi-layer storage systems are complex, making it difficult to correlate behavior at different levels. For example, a common problem occurs when two processes access the same file: which should be “charged” for the resulting disk I/O? When an I/O request is merged with another, what is the underlying cause? Few, if any, tools address these issues effectively.
- Distributed and HPC systems have complex timing and subtle interactions [59]; capturing this information is difficult.
- We do not know the best methods for analyzing traces. Traditional statistical tools seem unsuitable for describing complex behavior.
- We need to investigate new uses for introspection. Can modern file systems capture additional metadata to help users find data relevant to their needs? Can introspection at different levels of the storage stack be used for dynamic optimization?
- Real workloads often include multiple applications running simultaneously. Can we decouple the different signals to “tease out” the different applications? [57] Can we apply AI-style algorithms such as those used to simulate the way the human ear can separate sounds in a noisy environment?
- Metadata about traces (characteristics of the traced system, details about workloads, environmental information, etc.) have traditionally been collected outside the system components generating them and are usually stored independently of the traces themselves, which creates a disconnect that further complicates analysis.

We also need to better understand workloads and optimization in various ways, including:

- Can we build signatures for applications based on their I/O (and other) patterns?

- Can we generalize benchmarks to new storage technologies by using additional workload information? If so, what is the information necessary to achieve this?
- Can large-scale systems benefit from analyzing and modeling their workloads? Is it possible to understand and optimize data movement in these systems? For example, what semantics are needed for active processing with flash?
- Can we provide applications with control of storage data layout by exposing user-defined functions? [20] Can we control data layout across the deep-storage hierarchy as well as within it?

4.2.2 Provenance

An open research question is whether data provenance should be a first-class concern for storage systems. Current systems for tracking provenance are often external to the storage system; poor integration makes tracking provenance expensive and reduces the quality of provenance information. We note that provenance can be collected at multiple levels of fidelity; operations that produce data can be identified very generically (e.g., by application name), or the storage system can also include the arguments and parameters used to produce the data [66]. When supporting provenance, it is challenging to determine the fidelity at which information should be retained; if a query is not envisioned when the data is created, then provenance information required to answer that query might not be kept. The gold standard for provenance fidelity is reproducibility: if a storage system retains enough information to redo the original computation that produced the data, then it can answer arbitrary queries about how that data was produced by reproducing the original computation. Many original computations are deterministic; in this case, reproducibility requires retaining the original inputs [25, 45]. Otherwise, techniques that use deterministic record and replay can provide the needed reproducibility [37].

4.2.3 Metadata

The role of metadata is also poised to change in future, larger-scale storage systems. As denser storage technologies become available, metadata will be of paramount importance in order to locate data that is relevant to users and applications. As a result, we identify additional directions for future research, which complement the ideas proposed for introspecting storage systems and tracking data provenance.

At the time that data is generated, it is uncertain how it will be accessed, or what portions of it will be of most interest. In the case of provenance, we already discussed approaches that try to capture every facet of the data generation process. A similarly aggressive approach should be considered in metadata collection and generation. This calls for scalable approaches that treat metadata as a first-class citizen of the storage system. Several opportunities for future research arise if this paradigm is to be adopted. Maintaining large, metadata-rich namespaces that are both synchronous and performant becomes challenging. Existing research examines approaches that reconsider and relax the requirement of a synchronous namespace [71, 86] in order to maximize performance.

Another challenge is determining what metadata to collect and how to collect metadata in a systematic way without affecting the performance of the foreground workload. Existing frameworks and tools that hook into layers of the storage stack may act as a starting point [17, 32]. In many domains, the data itself contains information that can be used to help answer data disposition questions, such as how to seamlessly extract metadata from data, how to create rich persistent indices, and how to derive more metadata. Traditionally, such metadata structures have been maintained outside of storage systems, causing inconsistencies to arise. Existing approaches [78] to metadata extraction and integration within file systems could serve as a starting point for approaching these questions. Finally, an implicit assumption when collecting metadata, is that it will

be accessible for an arbitrary amount of time. This can be challenging, and will require techniques that are capable of translating in-memory data structures into persistent, well-defined formats.

4.3 In-situ and in-transit data processing

The race to exascale clusters is ushering in an era where scientists generate massive amounts of scientific data for analysis, in the order of hundreds of terabytes per simulation [30, 31]. Meanwhile, private and government organizations amass user data into very large datasets used to train neural networks for any use case imaginable. This is a trend that shows no signs of abating. The trend is partially driven by the rapid progress of research on denser storage technologies that can store this data [27, 65]. It is also caused by the steady growth in the number of ubiquitous sensors and IoT devices that generate data. Finally, there is a movement in systems modeling towards massive corpora of data, as opposed to domain expertise, likely due to the ease with which the former can be collected.

Today, scientific analysis is sped up through careful post-processing that optimizes the data layout on storage for expected access patterns. For other workloads, such as iterative machine learning and data analytics, the data can be distributed across multiple devices so that it can be accessed in parallel. However, neither of those approaches remains sustainable at scale, because as scientific output and datasets increase in size, the cost of post-processing or additional hardware management become prohibitive.

We encourage researchers to consider a paradigm where data is processed on its way to storage, i.e., while it transitions through the storage stack to become durable. The goal would be to drastically reduce (or eliminate) the need for data post-processing. A number of directions for future research arise: *Where can computation be accommodated in the storage pipeline without impacting the primary workload? What are the types of computation that would be amenable or adaptable to be performed in the storage pipeline?*

4.3.1 Data processing in-situ and in-transit

In-situ processing is a new paradigm for data processing that occurs during the application’s runtime. Although promising, in-situ data processing can be challenging compared to post-processing. First, traditional post-processing programs are designed to use all the resources of the nodes they occupy. At the time an application runs, however, freeing up these resources to process data can negate the benefit of in-situ processing. Thus, in-situ processing must proceed with **limited resources**. Second, traditional post-processing programs assume full visibility of the data in order to partition it across devices. With in-situ processing, data must be handled in a streaming fashion, so **limited visibility** should be tolerated. Third, traditional post-processing programs are rarely expected to scale to thousands of nodes because storage or network I/O will usually be the bottleneck. With in-situ processing, however, such extreme scalability may range from beneficial to necessary depending on the amount and type of resource unoccupied by the application.

One way to conduct in-situ analysis is to process the data on the storage system itself, where the data resides, e.g., on the storage servers or on the devices. Empirical studies have shown that there is slack on the storage servers or devices (e.g., SSDs) to conduct a limited amount of processing[58, 81]. Further, emerging flash devices have higher processing capability. Efforts such as Active Flash [81] and AnalyzeThis[77] are a starting point. Key research tasks include the construction of analysis abstractions within storage systems and devices, overlaying higher-level file system structures atop active flash devices, and placing data on the storage system in a manner that is conducive to future analyses.

Another way to process data on its way to storage is by running code in the network switches or other devices with limited computational capacity through which data flows, a technique we refer to as *in-transit* processing. Accelerators such as programmable NICs have been used recently to speed up in-memory key-value stores through caching [49, 54] and to provide strong consistency guarantees without compromising performance

[56]. Continuing this trend to examine types of data processing that could be offloaded to such lightweight devices seems natural.

4.3.2 Types of processing

A natural question that arises is whether there is computation naturally fitting the constraints of the in-situ and in-transit processing models, or whether current programs could be adapted to operate with such limited resources and visibility at extreme scale. Both could be potential directions for future research. Existing work has identified types of computation that is either amenable as-is, or can be adapted to the in-situ and in-transit models, such as debugging or monitoring tasks. Dapper [76] achieves scalable, unobtrusive tracing of distributed systems through filtering RPC requests so that the subset of messages retained form chains from source to destination. Performance monitoring can benefit from statistics collected in-situ or in-transit.

Recent work has shown that other tasks traditionally carried out by scanning the data could be adapted to the in-situ and in-transit models. The DeltaFS distributed file system [85, 86] constructs data indexes in-situ. This allows data queries to be drastically sped up with insignificant overhead at runtime, without the need for data post-processing. These results are also encouraging because they show that this can be achieved with frugal use of resources, leaving space for more processing. It thus becomes an interesting problem to find the limits of this approach by adapting other tasks, such as data analytics and modeling.

4.4 New interfaces

There is a growing need to examine and redesign I/O interfaces for applications. The decades-old, legacy interfaces have proved functional but are reaching scalability limits in modern cloud, data center, and HPC areas. New interfaces have been tried before (see [60] for an overview within HPC). But vendor lock-in fears have disincentivized adoption of new interfaces and instead motivated very stable, narrow I/O interfaces with community-built middleware and file format-specific access libraries to work around shortcomings. However, due to these narrow interfaces, middleware and file formats have by design limited information about the performance characteristics of a particular storage system and have to therefore make a lot of assumptions. With the “Cambrian explosion” of new storage devices with vastly differing performance characteristics, these assumptions have become inadequate. Fortunately, the wide adoption of open-source software storage systems makes the adoption of new storage interfaces possible without raising the potential of vendor lock-in. Now is the time to evaluate existing interfaces to determine if modification or extension will enable scalability, as well as design new interfaces to explore new capabilities.

For example, we may be able to relax POSIX semantics in well-defined ways such that we remove limitations, e.g., relax locking semantics in parallel systems. Alternatively, we may want to design new interfaces that give applications power to exploit the capabilities of upcoming storage hardware hierarchies, including storage class memory and NVRAM.

The motivation for re-examining and modifying legacy interfaces like POSIX is that they are widely used and a large number of existing applications depend on them. We need to understand what changes can be made to the legacy interfaces to address application scalability needs. However, this depends on having a clear understanding of the I/O behavior of target workloads so that appropriate changes are made. For example, in HPC workloads in order to determine if POSIX locking semantics can be relaxed, we need to fully understand the access patterns of parallel applications in shared files. If no two processes write to the same offset in a shared file, then byte range locking on the file may not be necessary. Other factors to consider are support of legacy interfaces like POSIX over modern storage infrastructures like object stores. We need to understand the inter-play between the storage models and explore methodologies for reducing overhead while continuing to support legacy applications.

New interfaces offer opportunities to present new capabilities to applications that have the potential to represent paradigm shifts in data management. For example, the current file-based paradigm encourages a static way of thinking about application data; when data changes, in many cases a new file is written to capture the change. However, if the fundamental concept in a storage system is no longer a file, but perhaps a data object that has a revision history as well as provenance information, the relationship of applications and users with data in storage is fundamentally different. As another example, legacy interfaces present a flat storage hierarchy to users, but modern and future storage systems are hierarchical composed of fast, near storage like storage class memory or NVRAM in addition to slower, permanent storage like a traditional parallel file system. New interfaces have the opportunity to expose aspects of a storage hierarchy for exploitation by applications. For example, if an interface allows an application to specify the persistence needs of a data object (perhaps as temporary), then the data object can be stored in the appropriate storage tier (e.g., temporary storage like NVRAM).

4.5 Co-design of applications and storage

The traditional storage stack is organized as an hierarchical, multi-layer structure. As the storage technologies become increasingly diversified (e.g., NAND flash, NVM, Optane), such a general-purpose layering structure falls short of being able to exploit the drastically distinct hardware properties.

A recent technical trend is to break the constraint of the storage layering structure by allowing software to obtain low-level control of hardware. For example, Open-Channel SSD [68] exposes device resources (e.g., I/O channels) to application software and allows it to directly operate on physical flash media, enabling a variety of optimizations for applications [26, 48, 69, 75]. We envision such an integrative approach would become a more common practice, especially in light of emerging technologies. However, this practice fundamentally changes the way in which application interacts with storage, creating new research questions.

The first question is where such an application-device collaboration should lie. The device-level management (e.g., FTL) could be integrated with application’s semantics at different levels. An aggressive method is to directly embed it within application logic. This approach removes the information barriers and maximizes the utilization of application’s domain knowledge, but it causes a high dependency on hardware specifics. For certain application scenarios, such as high-performance computing, customized in-house software, bundled software/hardware product, this deep integration would be effective and affordable, but for many others, it may not be suitable due to the development cost. A more conservative option is to place it at the operating system level. For example, a specialized device driver can provide a system interface to receive semantic hints or advice from applications and then accordingly control the hardware. A third option is to provide a relatively shallow abstraction as a user-level library level. The library presents the storage as a set of APIs to allow application developers to obtain a fine-grained control over the device while still retaining certain level of abstractions. Running at the user-level, the library development and debugging could also be simple and still portable.

Effective application-device co-design is non-trivial. It often tends to be ad hoc—each application is optimized individually. In fact, many optimizations needed across applications are often similar, if not exactly the same. It is desirable to identify a set of essential, sharable, low-level functions that are common across applications. For example, many applications desire to be able to allocate flash pages from a specified channel in an SSD. A set of such core functions could collectively form a new storage abstraction to facilitate co-design.

Cross-layer co-design complicates software development. Developers must consider both software logic and hardware issues. Debugging and testing will also be more complex, due to a larger test space. For example, a software bug accidentally writing a flash page twice without an erase may cause a silent data corruption. Methods to mitigate this added complexity are required.

4.6 Growing the pipeline: Education and teaching

In the past two decades, storage systems research has been increasingly hindered by a lack of practitioners. Demand has been increasing throughout computer science, new areas of CS are drawing attention away from storage, and competition from industry causes large numbers of bachelor's graduates to discard research as a possible career. The problem is exacerbated by the fact that systems research in general, and storage in particular, is not highly visible to the public, is difficult to explain to the uninitiated, and suffers from a perception that the area is dull and overly technical. Everyone can understand and appreciate a robot that does a back flip; it is much harder to get the layman excited about a fivefold improvement in disk performance. Yet, the latter may have a much larger practical impact on our lives.

To combat these perceptions and attract students into storage research, we need new approaches to education. We believe that presenting storage in a different light will increase the number of potential researchers and help address the crisis we are currently facing.

One way to increase student interest is to give them direct, hands-on experience with storage systems. Multi-layer traces can be used for classroom demonstrations so that students can understand system operation and learn how to build large-scale storage systems.

A major challenge in teaching storage systems is that large systems (more than a few tens of terabytes and a few servers) are prohibitively expensive. A tool similar to networking's ns-2 [67] could help by providing students with a simulation platform for experimenting with large systems. Another possibility would be to create a community testbed similar to CloudLab [?], although such facilities are often still much smaller than real-world installations (and tend to become overloaded at high-usage times such as early December when universities nationwide are finishing their terms).

Regardless of platform, students should have an opportunity to use and experiment with well-known storage systems and to build skills that are clearly useful in the field today, while still developing fundamental knowledge that will serve them well in the long term. For example, a system such as Hadoop could be an example of distributed file systems; students would be able to list that specific knowledge on their resumes but also could apply it to other areas.

One thing that attracts students to other areas is visibility: machine learning and deep learning are well covered in the press; gaming is something they're familiar with; robotics has long been a science-fiction dream; etc. We could improve storage's perception in the student community by increasing our use of visualization tools. Visualization can produce striking images that simultaneously attract the eye and provide deep insight. A class that asked students to visualize storage-related information could intrigue and challenge them. At the same time, it is difficult to visualize huge datasets that don't fit into memory; addressing that problem is a way to emphasize the importance of storage and engage students with the techniques needed to efficiently process out-of-memory data.

Finally, we should reach out to our colleagues in other subdisciplines of computer science and partner with them. Data scientists can suggest new approaches to understanding storage systems, while we can help them manage vast quantities of information. Artificial intelligence and machine learning offer different approaches to optimization. Library scientists understand how to organize information and might be able to suggest designs for presenting petabytes of data to users. Every community has a tendency to become insular; we should fight that tendency and actively seek collaborations that will be beneficial to both sides.

5 Evolution of Storage Systems with Emerging Hardware

5.1 Introduction

Storage systems are poised for tremendous change due to rapid developments in both hardware technology and applications. In addition to continuous evolution of block storage in the form of enhanced SSD+HDD+SMR disks, storage architectures are impacted by a continuously shifting hardware ecosystem, triggered by advances in processor, memory and networking technologies. In addition, new applications like IoT and high-volume data analytics, coupled with shifting usage modalities based on datacenter and cloud-based data management, are driving storage research to embrace a broad perspective incorporating architecture, operating systems, networking, compilers/languages and applications.

5.2 Embracing Change: The New Normal

Storage systems occupy a central role within modern computer system architectures and their design is increasingly influenced by developments in the processor, memory and network subsystems. Future storage systems will need to be flexible to respond to continual changes in the hardware ecosystem and agile in adapting to the evolving requirements of new applications. It is not an exaggeration to say that continual change in technology and applications is the new normal, and storage research should embrace methodologies and system designs that can take advantage of these advances in a modular, incremental and timely manner. New storage technologies and applications spur novel storage system designs, that in turn can inform developments in future hardware components and software. We envisage a dynamic and synergistic research agenda cutting across traditional systems boundaries and encompassing hardware architecture, operating system, compilers, distributed systems, and applications. The complementary challenges and opportunities for storage system research are elaborated in the two challenges described below.

1. Incorporate new technology and application requirements

The changing landscape for storage systems is characterized by increasingly heterogeneous computing hardware, diverse and less discontinuous memory hierarchies, novel interconnect and networking fabrics, and complex, highly dynamic workload patterns arising from applications like IoT and Machine Learning. The overarching challenge to deal with the constantly shifting landscape is to develop design methodologies that incorporate dynamic policies and flexible interfaces that can adapt to new devices and workloads.

- (a) **Memory and storage devices:** The traditional world of fast, fine-grained volatile memory and slow block-level storage is in transition. New memory-bus-attached, byte-addressable non-volatile memory devices (abbreviated as B-NVM) are blurring the line between memory and storage; potential applications of B-NVM range from adding another level in the storage hierarchy to eventual deployment as the central component of working storage. Meanwhile, block storage continues to evolve with new block NVM devices like Optane drives, devices with asymmetric read/write performance like Shingled-Magnetic Recording (SMR) disks, and processing-in-storage devices like Kinetic Drives. Further diversification in terms of capacity versus performance-oriented devices and the tiering of data into active, working, near-line, cold, and archival levels, creates a complex storage landscape. Nascent technologies like DNA storage raise intriguing possibilities for future archival storage at one end of the spectrum, while the large amounts of High Bandwidth memory close to the processor open up new opportunities for caching, buffering and prefetching.

To exploit the advances in technology, research is needed to understand and control the evolving complex memory and storage hierarchy. Models, algorithms and mechanisms for data placement

and data migration across multiple types of heterogeneous devices for the entire data lifetime need to be designed. Issues of migration granularity and frequency, analytical models, and the incorporation of machine learning techniques to automatically manage the hierarchy and meet QoS performance, reliability, availability, and consistency requirements must be developed.

- (b) **Network fabrics:** Research is needed to understand the impact of new system interconnect technologies (e.g. CCIX, OpenCAPI), fast 40/100/200 Gbps Ethernet and 56/100/200 Gbps Infiniband networks, lightweight network fabrics (e.g. RDMA over Infiniband or RDMA over converged Ethernet (RoCE)), and emerging NVM-motivated standards (e.g. GenZ) on storage system organization, and protocols for distributed and clustered storage. As the speeds of storage and networking converge, protocols designed for slow remote storage will need to be supplanted by more efficient schemes. An understanding of the design space and the networking abstractions required to support future storage systems is critical.
- (c) **Processor architecture:** Heterogeneous processor architectures with specialized accelerators, GPUs, and FPGA devices can potentially facilitate compute-intensive tasks on behalf of the storage system. Examples include coding, compression, and support for cryptography and traffic management, to enable enhanced reliability, security, and QoS. Advanced processor mechanisms for concurrency control like Hardware Transaction Memory (HTM) can be leveraged to provide lightweight storage transactions in conjunction with fine-grained high-speed storage like B-NVM.

Research to evaluate the potential of new processor hardware on storage and file system software, and the development of appropriate OS support is necessary to leverage processor architecture advances to benefit the next generation of storage systems.

- (d) **Applications:** As data management becomes central to our everyday activities, storage systems will be increasingly called upon to support functionalities beyond the traditional role as performant suppliers of raw persistent data. The agility required of future storage systems to handle technology changes must carry over to designs that can incorporate new application-driven requirements gracefully and incrementally.

The imperative for security of stored data, verifiable access control, data privacy and isolation will continue to increase. Data provenance and verifiable audit trails of data access and modification history over long data lifetimes will become increasingly necessary. Demand for quantifiable performance QoS guarantees will increase for applications requiring time-sensitive response times and those deployed in shared datacenter environments. Storage systems will need to satisfy microsecond and sub-microsecond latency requirements by carefully leveraging the memory hierarchy. Finally, application workloads will continue to change and increasingly include intricate, dynamic and fine-grained workload patterns for IoT, machine learning, graph analytics and other data-intensive high-performance computing tasks.

2. Storage-driven systems research

Future storage systems will benefit from hardware and software support. Cross collaboration between various sub-disciplines is vital to providing a coherent path forward and avoid the reinvention of multiple overlapping solutions through silo-based research.

Two specific areas of synergy are in shaping the design of CPU architecture and OS abstractions, and the development of interconnect fabrics. Processor support is essential for using devices like B-NVM that have delays and lie between DRAM and typical storage. For instance, can processors better tolerate longer and more variable access delays (10–100s of microseconds)? What memory mapping support is needed to handle fine-grained, local and remote terabyte-sized data? What hardware mechanisms to support

transactions for persistent memory are appropriate? Should processors provide hardware logging and at what granularity? Is speculative paging worth supporting? In the networking domain, current RDMA protocols are inadequate in their support for end-to-end persistence and have scalability limitations. How should the storage system requirements drive the next generation of interconnection fabrics?

5.3 Impact of Byte-Addressable NVM

The availability of byte-addressable bus-attached NVM (B-NVM) devices gives rise to a new, potentially disruptive addition to the storage-memory hierarchy. B-NVM has advantages over DRAM in terms of capacity, power and non-volatility, and, in contrast to traditional block-structured storage, enables fast, cache-line-granular, direct processor access to persistent memory. The characteristics of B-NVM dovetail with modern applications handling huge persistent data sets and requiring fine-grained scattered data accesses, making it an excellent candidate for serving both memory and storage needs. Where B-NVM will ultimately fit in the storage hierarchy is currently an open question; its place will depend on the evolution of the cost, performance, and reliability characteristics of the underlying devices. There are intriguing possibilities and potentially large payoffs in terms of speed, power, crash resilience, and conceptual simplicity in deployment of B-NVM in the storage stack. To realize the potential of the new technology requires a strong research effort to deal with the multi-faceted and cross-cutting issues raised, which we discuss below.

1. Operating system and application development support

The past decade has seen considerable research on the challenges of using B-NVM devices within a single server. The problems of consistent ordering between volatile and persistent memory, support for persistent memory transactions in the presence of failure, and coordination with volatile-memory concurrency controls mechanisms like Hardware Transactional Memory, have been addressed with solutions and techniques proposed in the architecture, OS, storage, languages, and database communities. Example important issues requiring research are identified below:

- **Security and integrity:** How to efficiently support encryption and low-latency access control at the granularity of memory accesses? How to avoid corruption of persistent data from malicious or buggy applications using lightweight mechanisms? Are language-level mechanisms appropriate or sufficient?
- **Abstractions for persistent data:** B-NVM can provide a single-level memory with a single namespace for both volatile and persistent data. What are the appropriate naming abstractions and operating system support required? Are relative pointers the appropriate addressing mechanism? How should garbage collection be organized in this environment?
- **Consistency and transaction support:** Research to transition the most efficient techniques into actual systems, and to build tools for building/migrating applications for direct memory access.

2. Distributed shared NVM-based storage

While the past research has concentrated on single-server-attached B-NVM, the next big step forward is understanding the challenges of deploying scalable, distributed B-NVM storage architectures in future datacenters.

- **Distributed architecture:** A basic issue is to understand the tradeoffs in different architectural choices for distributed B-NVM. The design space includes distributed (persistent) memory systems, clustered storage systems, or hybrid architectures. The choices have significant impact on the usage

modalities and on application structuring, and careful analysis and empirical studies are necessary to understanding the tradeoffs. This study should encompass the analysis of the underlying network fabric. Interconnect like Infiniband-RDMA or RoCE match the latencies of B-NVM devices making them a natural fit as the networking fabric. However, RDMA has known problems of scalability and has semantics that are unhelpful in providing the durability guarantees needed for B-NVM transactions. Understanding the correct abstractions that must be supported by the network fabric, and their impact on storage architectures and applications is needed to lay a sound foundation for future scalable B-NVM systems.

- **Software support:** A strong research effort is needed in identifying the mechanisms and abstractions necessary to support future distributed, shared B-NVM. The problems are challenging and solutions are necessary to meet the scalability, reliability, usability, correctness, and latency requirements of applications. Many of these issues (listed below) have been previously examined in the context of single-server B-NVM systems, while distributed versions of these problems have been studied in the context of traditional block storage and TCP/IP networks. However, the distributed version of these problems deals with a far larger design space than server-attached B-NVM, and the microsecond-level latencies of B-NVM and direct-access RDMA-like protocols changes the nature of the problem qualitatively, requiring vigorous new research to explore the design space effectively. Some of the issues needing new research and lightweight solutions in the distributed B-NVM context are mentioned below.
 - **Transparent global naming:** How should shared data be named, distributed, and accessed? What are the overheads for metadata management?
 - **Transaction management:** How should transactions spread across multiple B-NVM hosts be orchestrated? How can RDMA-like protocols or atomics be exploited? What is the appropriate form of distributed logging?
 - **Reliability and availability:** Replication or erasure coding? Wear-leveling across nodes?
 - **Consistency and coherence:** What are the appropriate consistency models balancing application requirements and performance?
 - **Crash resilience and durability:** How does one handle failures of nodes holding TBs of B-NVM data?
 - **Metadata management:** How does one handle the size and overheads of metadata management?

5.4 New Storage Abstractions

New storage abstractions are necessary to support future storage architectures, new functionalities and modes of usage.

1. **Processing in Smart Controllers:** Research is needed to explore how to embed useful functionality in the path to storage through enhanced memory, storage, or network controllers. Current “bump-in-the-wire” functionality is provided by FPGAs at the back-end of NICs to provide on-the-fly transparent compression or encryption. Can offloading to dedicated controllers be useful in other contexts to provide support for functions like provenance, security, QoS, write ordering or transaction support? For instance, an interposed persistent memory controller between the processor LLC and B-NVM can simplify transactional programming. Can distributed persistent storage benefit from a controller at the NIC? Can functions like replication and data migration be offloaded to controller networks? Can stateful operations be successfully offloaded to controllers? Going beyond, what is the impact of embedding processing power within memory chips in processing-in-memory (PIM) models?

2. **Programming models:** As the abstraction levels of storage access rise, so does the need to define appropriate programming models and APIs for data access. What are the appropriate abstractions beyond traditional file system block reads and writes? Are KV stores supporting object access using “gets” and “puts” adequate? Are transactional models useful in specifying application-level semantics? What are the boundaries between storage system, middleware, and applications? Will memory-style programming become the common mode for specifying persistent access? If so, what support will be needed from the compiler and programming languages? How would such specifications coexist and play with existing interfaces to block-based access?

5.5 Managing Exabyte-scale Storage

1. **Data reliability and durability:** Proliferation and diversity of hardware causes an increase in the failure modes that need to be considered. The problem gets acute as the data set scales requiring development of effective methods to handle ubiquitous failures in some portion of the storage system. The correlation between failure modes resulting in catastrophic data loss need to be addressed. The use of cross-layer reliability schemes to manage reliability overheads and exploit the availability of processing hardware provides a possible avenue of exploration. No one size fits all storage solutions for all different kinds of application workloads.
2. **Data ingest and migration:** As datacenter storage becomes the norm techniques to optimize data movement within datacenters and data migration across datacenters becomes a vital issue. Vendor lock-in inhibits the move to using a shared infrastructure despite other advantages and alternative approaches need to be considered.

5.6 Infrastructure Development

The rapid changes in hardware and application requirements present a tough challenge for researchers to evaluate their solutions. To allow for continued research and innovation it is imperative to have access to platforms for experimentation and evaluation. One component would include the availability of shared hardware testbeds with state-of-the-art storage, networking and processing hardware. In addition, development of models and simulators for new devices, scalable system simulation and emulation tools, and collection and availability of modern workload traces are essential. NSFCloud [4] is such a platform where two large-scale academic cloud testbeds—Chameleon [2] and CloudLab [3]—are providing researchers with a shared compute and storage infrastructure to conduct scientific simulation and distributed system research validations at scale. A good feature of Chameleon, for example, is to allow third-party entities (e.g., researchers who are regular users of the testbed) to contribute to the cloud infrastructure through appliance (framework prototypes developed by third party and made available for use in the form of a VM image) and software (tools open sourced that can be cloned). Aligned with the same line, a desirable feature of such open testbeds would be to open flexible options for more hardware configurations. For instance, emerging hardware devices such as Intel Optane SSD [5] can be integrated as an option for storage and systems researcher to select, to boost cutting-edge storage research activities.

6 Conclusion

This report presents the discussions and perspectives of the storage community members that participated in the two-day Data Storage Research Vision 2025 Workshop in June 2018. We have identified four key thrusts: enhancing cloud and edge computing I/O infrastructures, designing storage for emerging AI applications, rethinking the storage systems abstractions in service of the new and innovative applications, and redesigning

storage systems for emerging hardware, focusing on which will help the storage research community address and mitigate the discussed challenges. While there are other challenges not discussed here, our goal is to make the community aware of these and enable innovative research that can benefit the systems community, and in turn improve systems that underlie our modern life.

References

- [1] AWS Lambda: Serverless Computing. <https://aws.amazon.com/lambda/>. Accessed: June 6, 2018.
- [2] Chameleon: A Configurable Experimental Environment for Large-Scale Cloud Research . <https://www.cloudlab.us/>. Accessed: June 6, 2018.
- [3] CloudLab: Flexible, Scientific Infrastructure for Research on the Future of Cloud Computing. <https://www.cloudlab.us/>. Accessed: June 6, 2018.
- [4] Enabling a new future for cloud computing. https://www.nsf.gov/news/news_summ.jsp?cntn_id=132377.
- [5] Intel Optane SSD Series. <https://www.intel.com/content/www/us/en/products/memory-storage/solid-state-drives/gaming-enthusiast-ssds/optane-900p-series.html>.
- [6] SNIA Education. <https://www.snia.org/education>. Accessed: June 6, 2018.
- [7] USENIX FAST'18 Call for Papers. <https://www.usenix.org/conference/fast18/call-for-papers>. Accessed: June 6, 2018.
- [8] Sr 11-7: Guidance on model risk management. Internet draft, 2011. <https://www.federalreserve.gov/supervisionreg/srletters/sr1107.htm>.
- [9] Apache. flink. Internet draft, 2017.
- [10] Gdpr online, 2017. <http://www.eudgpr.org>.
- [11] New york bill on algorithmic bias, 2017. <http://www.businessinsider.com/algorithmic-bias-accountability-bill-passes-in-new-york-city-2017-12>.
- [12] The openfog consortium. openfog reference architecture. Internet draft, 2017.
- [13] Pytorch, 2018. <http://pytorch.org>.
- [14] scikit-learn, 2018. <http://scikit-learn.org/stable>.
- [15] ABADI, M., AND ET AL. Tensorflow: Large-scale machine learning on heterogeneous systems. Internet draft, 2015.
- [16] AGRAWALA, A. K., MOHR, J. M., AND BRYANT, R. M. An approach to the workload characterization problem. *Computer* 9, 6 (1976), 18–32.
- [17] AMVROSIADIS, G., BROWN, A. D., AND GOEL, A. Opportunistic storage maintenance. In *Proceedings of the 25th Symposium on Operating Systems Principles (SOSP)* (Monterey, California, 2015), ACM, pp. 457–473.

- [18] ANGUIA, D., GHIO, A., ONETO, L., PARRA, X., AND REYES-ORTIZ., J. L. A public domain dataset for human activity recognition using smartphones. *21th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning, ESANN* (2013), 1–15.
- [19] ARANYA, A., WRIGHT, C. P., AND ZADOK, E. Tracefs: A file system to trace them all. In *Proceedings of the 3rd Usenix Conference on File and Storage Technologies (FAST)* (San Francisco, CA, Mar. 2004), Usenix Association, pp. 129–143.
- [20] ARPACI-DUSSEAU, A. C., AND ARPACI-DUSSEAU, R. H. Information and control in gray-box systems. In *Proceedings of the 18th ACM Symposium on Operating Systems Principles (SOSP '01)* (Banff, AL, Canada, Oct. 2001), pp. 43–56.
- [21] ARPACI-DUSSEAU, R. H., AND ARPACI-DUSSEAU, A. C. *Operating Systems: Three Easy Pieces*, 0.91 ed. Arpaci-Dusseau Books, May 2015.
- [22] AXBOE, J., AND BRUNELLE, A. D. blktrace user guide. <http://www.fis.unipr.it/doc/blktrace-1.0.1/blktrace.pdf>, May 2008.
- [23] BARROSO, L. A., CLIDARAS, J., AND HÖLZLE, U. The datacenter as a computer: An introduction to the design of warehouse-scale machines. *Synthesis lectures on computer architecture* 8, 3 (2013), 1–154.
- [24] BASIRI, A., BEHNAM, N., DE ROOIJ, R., HOCHSTEIN, L., KOSEWSKI, L., REYNOLDS, J., AND ROSENTHAL, C. Chaos engineering. *IEEE Software* 33, 3 (May 2016), 35–41.
- [25] BENT, J., THAIN, D., ARPACI-DUSSEAU, A. C., ARPACI-DUSSEAU, R. H., AND LIVNY, M. Explicit control in a batch-aware distributed file system. In *Proceedings of the 2004 Usenix Symposium on Network Systems Design and Implementation (NSDI)* (San Francisco, CA, March 2004).
- [26] BJØRLING, M., GONZALEZ, J., AND BONNET, P. LightNVM: The Linux Open-Channel SSD subsystem. In *Proceedings of the 15th Usenix Conference on File and Storage Technologies (FAST)* (Santa Clara, CA, 2017).
- [27] BORNHOLT, J., LOPEZ, R., CARMEAN, D. M., CEZE, L., SEELIG, G., AND STRAUSS, K. A DNA-based archival storage system. In *Proceedings of the Twenty-First International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)* (Atlanta, Georgia, USA, 2016), ACM, pp. 637–649.
- [28] BURT, A. Is there a 'right to explanation' for machine learning in the gdpr. Internet draft, 2017.
- [29] BYAN, S., LENTINI, J., MADAN, A., PABON, L., CONDUCT, M., KIMMEL, J., KLEIMAN, S., SMALL, C., AND STORER, M. Mercury: Host-side flash caching for the data center. In *Proc. of IEEE 28th Symposium on Mass Storage Systems and Technologies (MSST)* (April 2012).
- [30] BYNA, S., SISNEROS, R., CHADALAVADA, K., AND KOZIOL, Q. Tuning parallel I/O on Blue Waters for writing 10 trillion particles.
- [31] BYNA, S., USELTON, A., PRABHAT, D. K., AND HE, Y. Trillion particles, 120,000 cores, and 350 TBs: Lessons learned from a hero I/O run on Hopper. In *Cray User Group (CUG)* (2013).
- [32] CARNS, P., HARMS, K., ALLCOCK, W., BACON, C., LANG, S., LATHAM, R., AND ROSS, R. Understanding and improving computational science storage access through continuous characterization. *Trans. Storage* 7, 3 (Oct. 2011), 8:1–8:26.

- [33] CARNS, P., LATHAM, R., ROSS, R., ISKRA, K., LANG, S., AND RILEY, K. 24/7 characterization of petascale I/O workloads. In *2009 IEEE International Conference on Cluster Computing and Workshops* (2009).
- [34] CHEN, H., ZIEGLER, D., CHAJED, T., CHLIPALA, A., KAASHOEK, M. F., AND ZELDOVICH, N. Using Crash Hoare Logic for Certifying the FSCQ File System. In *Proceedings of the 25th Symposium on Operating Systems Principles* (New York, NY, USA, 2015), SOSP '15, ACM, pp. 18–37.
- [35] CHIU, L., AND HETZLER, S. Building a future proof enterprise storage in 2025. *Data Storage Research 2025: An NSF-Sponsored Community Visioning Workshop* (2018).
- [36] COOPER, B. F., SILBERSTEIN, A., TAM, E., RAMAKRISHNAN, R., AND SEARS, R. Benchmarking Cloud Serving Systems with YCSB. In *Proceedings of the 1st ACM Symposium on Cloud Computing* (New York, NY, USA, 2010), SoCC '10, ACM, pp. 143–154.
- [37] DEVECSERY, D., CHOW, M., DOU, X., FLINN, J., AND CHEN, P. M. Eidetic systems. In *Proceedings of the 2014 Symposium on Operating Systems Design and Implementation* (Broomfield, CO, October 2014).
- [38] EVANS, D. *The Internet of Things: How the Next Evolution of the Internet Is Changing Everything*. Cisco IBSG (Internet Business Solutions Group), 2011.
- [39] FLINK, A. Checkpoints in flink. Internet Draft, 2018.
- [40] GANESAN, A., ALAGAPPAN, R., ARPACI-DUSSEAU, A. C., AND ARPACI-DUSSEAU, R. H. Redundancy does not imply fault tolerance: Analysis of distributed storage reactions to single errors and corruptions. In *Proceedings of the 15th Usenix Conference on File and Storage Technologies* (Berkeley, CA, USA, 2017), FAST'17, USENIX Association, pp. 149–165.
- [41] GARFINKEL, S., MATTHEWS, J., SHAPIRO, S., AND SMITH, J. Towards algorithmic transparency and accountability. *Communications of the ACM, Vol. 60 No. 9, Page 5* (2016).
- [42] GHEMAWAT, S., GOBIOFF, H., AND LEUNG, S.-T. The Google File System. In *Proceedings of the Nineteenth ACM Symposium on Operating Systems Principles* (New York, NY, USA, 2003), SOSP '03, ACM, pp. 29–43.
- [43] GUERRA, J., PUCHA, H., GLIDER, J., BELLUOMINI, W., AND RANGASWAMI, R. Cost effective storage using extent-based dynamic tiering. In *Proc. of the USENIX Conference on File and Storage Technologies* (February 2011).
- [44] GUNAWI, H. S., HAO, M., SUMINTO, R. O., LAKSONO, A., SATRIA, A. D., ADITYATAMA, J., AND ELIAZAR, K. J. Why Does the Cloud Stop Computing?: Lessons from Hundreds of Service Outages. In *Proceedings of the Seventh ACM Symposium on Cloud Computing* (New York, NY, USA, 2016), SoCC '16, ACM, pp. 1–16.
- [45] GUNDA, P. K., RAVINDRANATH, L., THEKKATH, C. A., YU, Y., AND ZHUANG, L. Nectar: Automatic management of data and computation in datacenters. In *Proceedings of the 2010 Symposium on Operating Systems Design and Implementation (OSDI)* (Vancouver, BC, Canada, October 2010).
- [46] HEY, T., TANSLEY, S., TOLLE, K. M., ET AL. *The Fourth Paradigm: Data-Intensive Scientific Discovery*. Microsoft Research Redmond, WA, 2009.

- [47] HOLLAND, D. A., ANGELINO, E., WALD, G., AND SELTZER, M. I. Flash Caching on the Storage Client. In *Proc. of USENIX ATC* (2013).
- [48] HUANG, J., BADAM, A., CAULFIELD, L., NATH, S., SENGUPTA, S., SHARMA, B., AND QURESHI, M. K. FlashBlox: Achieving both performance isolation and uniform lifetime for virtualized SSDs. In *Proceedings of the 15th Usenix Conference on File and Storage Technologies (FAST)* (Santa Clara, CA, 2017).
- [49] JIN, X., LI, X., ZHANG, H., SOULÉ, R., LEE, J., FOSTER, N., KIM, C., AND STOICA, I. NetCache: Balancing key-value stores with fast in-network caching. In *Proceedings of the 26th Symposium on Operating Systems Principles (SOSP)* (Shanghai, China, Oct. 2017), ACM, pp. 121–136.
- [50] KIM, J. K., HO, Q., LEE, S., XUN ZHENG, W. D., GIBSON, G. A., AND XING., E. P. Strads : A distributed framework for scheduled model parallel machine learning.
- [51] KOLLER, R., MARMOL, L., RANGASWAMI, R., SUNDARARAMAN, S., TALAGALA, N., AND ZHAO, M. Write Policies for Host-side Flash Caches. In *Proc. of USENIX FAST* (2013).
- [52] LANE, T., AND BRODLEY, C. E. An application of machine learning to anomaly detection. In *Proceedings of the 20th National Information Systems Security Conference* (1997), vol. 377, Baltimore, USA, pp. 366–380.
- [53] LEVENTHAL, A. Flash storage memory. *Commun. ACM* 51, 7 (July 2008), 47–51.
- [54] LI, B., RUAN, Z., XIAO, W., LU, Y., XIONG, Y., PUTNAM, A., CHEN, E., AND ZHANG, L. KV-Direct: High-performance in-memory key-value store with programmable NIC. In *Proceedings of the 26th Symposium on Operating Systems Principles (SOSP)* (Shanghai, China, Oct. 2017), ACM, pp. 137–152.
- [55] LI, H., GHODSI, A., ZAHARIA, M., SHENKER, S., AND STOICA, I. Tachyon: Reliable, memory speed storage for cluster computing frameworks. In *Proceedings of the ACM Symposium on Cloud Computing* (New York, NY, USA, 2014), SOCC '14, ACM, pp. 6:1–6:15.
- [56] LI, J., MICHAEL, E., AND PORTS, D. R. K. Eris: Coordination-free consistent transactions using in-network concurrency control. In *Proceedings of the 26th Symposium on Operating Systems Principles (SOSP)* (Shanghai, China, Oct. 2017), ACM, pp. 104–120.
- [57] LIU, Y., GUNASEKARAN, R., MA, X., AND VAZHKUDAI, S. S. Automatic identification of application I/O signatures from noisy server-side traces. In *Proceedings of the 12th Usenix Conference on File and Storage Technologies (FAST)* (Santa Clara, CA, Feb. 2014).
- [58] LIU, Y., GUNASEKARAN, R., MA, X., AND VAZHKUDAI, S. S. Server-side log data analytics for I/O workload characterization and coordination on large shared storage systems. In *Proceedings of the 29th International Conference on High Performance Computing, Networking, Storage and Analysis (SC)* (Salt Lake City, UT, Nov. 2016).
- [59] MACE, J., ROELKE, R., AND FONSECA, R. Pivot tracing: Dynamic causal monitoring for distributed systems. In *Proceedings of the 25th ACM Symposium on Operating System Principles (SOSP '15)* (Monterey, CA, Oct. 2015).
- [60] MAY, J. *Parallel I/O for High Performance Computing*, 1st edition ed. Morgan Kaufman, 2000, ch. 4: "Access Patterns and Optimizations", pp. 87–89.

- [61] MENG, F., ZHOU, L., MA, X., UTTAMCHANDANI, S., AND LIU, D. vCacheShare: Automated Server Flash Cache Space Management in a Virtualization Environment. In *Proc. of USENIX ATC* (2014).
- [62] MESNIER, M., THERESKA, E., GANGER, G. R., ELLARD, D., AND SELTZER, M. File classification in self-* storage systems. In *null* (2004), IEEE, pp. 44–51.
- [63] MIAO, H., LI, A., DAVIS, L. S., AND DESHPANDE, A. Towards unified data and lifecycle management for deep learning. In *2017 IEEE 33rd International Conference on Data Engineering (ICDE)* (April 2017), pp. 571–582.
- [64] MICROSOFT CORPORATION. Event tracing. [https://msdn.microsoft.com/en-us/library/windows/desktop/bb968803\(v=vs.85\).aspx](https://msdn.microsoft.com/en-us/library/windows/desktop/bb968803(v=vs.85).aspx). Visited June 8, 2018.
- [65] MICROSOFT RESEARCH. Project Silica. <https://www.microsoft.com/en-us/research/project/project-silica/>.
- [66] MUNISWAMY-REDDY, K.-K., HOLLAND, D. A., BRAUN, U., AND SELTZER, M. Provenance-aware storage systems. In *Proceedings of the Usenix Annual Technical Conference (ATC)* (Boston, MA, May/June 2006), pp. 43–56.
- [67] ns-2. http://nslam.sourceforge.net/wiki/index.php/Main_Page. Visited June 8, 2018.
- [68] Open-Channel solid state drives. <https://openchannelssd.readthedocs.io>, 2018. Visited June 8, 2018.
- [69] OUYANG, J., LIN, S., JIANG, S., HOU, Z., WANG, Y., AND WANG, Y. SDF: software-defined flash for web-scale Internet storage systems. In *Nineteenth International Conference on Architectural Support for Programming Languages and Operating Systems* (Salt Lake City, UT, Mar. 2014).
- [70] QIN, D., BROWN, A. D., AND GOEL, A. Reliable Writeback for Client-side Flash Caches. In *Proc. of USENIX ATC* (2014).
- [71] REN, K., ZHENG, Q., PATIL, S., AND GIBSON, G. IndexFS: Scaling file system metadata performance with stateless caching and bulk insertion. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC)* (New Orleans, Louisiana, 2014), IEEE Press, pp. 237–248.
- [72] RISKA, A., AND RIEDEL, E. Disk drive level workload characterization. In *USENIX Annual Technical Conference, General Track* (2006), vol. 2006, pp. 97–102.
- [73] SANTANA, R., LYONS, S., KOLLER, R., RANGASWAMI, R., AND LIU, J. To ARC or Not to ARC. In *7th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage 15)* (2015), USENIX Association.
- [74] SEBASTIAN SCHELTER, JOOS-HENDRIK BOESE, J. K. T. K. S. S. Automatically tracking metadata and provenance of machine learning experiments. *Workshop on ML Systems at NIPS (MLSys) Workshop*. (2017).
- [75] SHEN, Z., CHEN, F., JIA, Y., AND SHAO, Z. DIDACache: A deep integration of device and application for flash-based key-value caching. In *Proceedings of the 15th Usenix Conference on File and Storage Technologies (FAST)* (Santa Clara, CA, Feb. 2017).

- [76] SIGELMAN, B. H., BARROSO, L. A., BURROWS, M., STEPHENSON, P., PLAKAL, M., BEAVER, D., JASPAN, S., AND SHANBHAG, C. Dapper, a large-scale distributed systems tracing infrastructure. Tech. rep., Google, Inc., 2010.
- [77] SIM, H., KIM, Y., VAZHKUDAI, S. S., TIWARI, D., ANWAR, A., BUTT, A. R., AND RAMAKRISHNAN, L. AnalyzeThis: An analysis workflow-aware storage system. In *Proceedings of the 28th International Conference on High Performance Computing, Networking, Storage and Analysis (SC)* (Austin, TX, Nov. 2015).
- [78] SIM, H., KIM, Y., VAZHKUDAI, S. S., VALLÉE, G. R., LIM, S.-H., AND BUTT, A. R. Tagit: An integrated indexing and search service for file systems. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC)* (Denver, Colorado, 2017), ACM, pp. 5:1–5:12.
- [79] strace. <http://strace.io>. Visited June 8, 2018.
- [80] THERESKA, E., BALLANI, H., O’ SHEA, G., KARAGIANNIS, T., ROWSTRON, A., TALPEY, T., BLACK, R., AND ZHU, T. IOFlow: A Software-Defined Storage Architecture. In *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles* (New York, NY, USA, 2013), SOSP ’13, ACM, pp. 182–196.
- [81] TIWARI, D., BOBOILA, S., VAZHKUDAI, S. S., KIM, Y., MA, X., DESNOYERS, P., AND SOLIHIN, Y. Active flash: Towards energy-efficient, in-situ data analytics on extreme-scale machines. In *Proceedings of the 11th Usenix Conference on File and Storage Technologies (FAST)* (San Jose, CA, Feb. 2013).
- [82] VIETRI, G., RODRIGUEZ, L. V., MARTINEZ, W. A., LYONS, S., LIU, J., RANGASWAMI, R., ZHAO, M., AND NARASIMHAN, G. Driving cache replacement with ml-based lecar. In *10th {USENIX} Workshop on Hot Topics in Storage and File Systems (HotStorage 18)* (2018), USENIX
- [83] WU, X., AND REDDY, A. L. N. Exploiting Concurrency to Improve Latency and Throughput in a Hybrid Storage System. In *Proc. of IEEE MASCOTS* (September 2010).
- [84] ZAHARIA, M., XIN, R. S., WENDELL, P., DAS, T., A. M. D. A. M. X. R. J. V. S. F. M. G. A. G. J. S. S., AND STOICA, I. Apache spark: a unified engine for big data processing. 56–65.
- [85] ZHENG, Q., AMVROSIADIS, G., KADEKODI, S., GIBSON, G. A., CRANOR, C. D., SETTLEMYER, B. W., GRIDER, G., AND GUO, F. Software-defined storage for fast trajectory queries using a deltaFS indexed massive directory. In *Proceedings of the 2nd Joint International Workshop on Parallel Data Storage & Data Intensive Scalable Computing Systems (PDSW-DISCS)* (Denver, Colorado, 2017), ACM, pp. 7–12.
- [86] ZHENG, Q., REN, K., GIBSON, G., SETTLEMYER, B. W., AND GRIDER, G. DeltaFS: Exascale file systems scale better without dedicated servers. In *Proceedings of the 10th Parallel Data Storage Workshop (PDSW)* (Austin, Texas, 2015), ACM, pp. 1–6.