Lithium: Virtual Machine Storage for the Cloud

Jacob Gorm Hansen
VMware
Aarhus, Denmark
jgorm@vmware.com

Eric Jul
Bell Laboratories
Alcatel-Lucent, Dublin, Ireland
eric@cs.bell-labs.com

ABSTRACT

To address the limitations of centralized shared storage for cloud computing, we are building Lithium, a distributed storage system designed specifically for virtualization workloads running in large-scale data centers and clouds. Lithium aims to be scalable, highly available, and compatible with commodity hardware and existing application software. The design of Lithium borrows ideas and techniques originating from research into Byzantine Fault Tolerance systems and popularized by distributed version control software, and demonstrates their practical applicability to the performance-sensitive problem of VM hosting. To our initial surprise, we have found that seemingly expensive techniques such as versioned storage and incremental hashing can lead to a system that is not only more robust to data corruption and host failures, but also often faster than naïve approaches and, for a relatively small cluster of just eight hosts, performs well compared with an enterprise-class Fibre Channel disk array.

Categories and Subject Descriptors

D.4.3 [**Operating Systems**]: File Systems Management—*Distributed file systems*; H.3.4 [**Information storage and retrieval**]: Systems and Software—*Distributed Systems*; D.4.5 [**Operating Systems**]: Reliability—*Fault-tolerance*

General Terms

Design, Experimentation, Performance, Reliability

1. INTRODUCTION

The emergence of cloud computing has motivated the creation of a new generation of distributed storage systems. These systems trade backwards-compatibility in exchange for better scalability and resiliency to disk and host failures. Instead of centralizing storage in a Storage Area Network (SAN), they use cheap directly attached storage and massive replication to keep data durable and available. The downside to these cloud-scale storage systems is their lack of backwards-compatibility with legacy software written before the cloud-era. Virtual machine monitors like VMware's ESX platform excel at running both new and legacy software inside virtual machines (VMs), but have traditionally depended on centralized shared storage as to take advantage of virtualization features such as live VM migration [10, 32] and host fail-over [44]. Local storage is rarely used for hosting VMs because it ties them to specific hosts thereby making them vulnerable to single host failures.

Though SAN and NAS (Network Attached Storage) disk arrays often are veritable storage super computers, their scalability is fundamentally limited, and occasionally they become bottlenecks that limit the virtualization potential of certain applications. As new generations of processors and chipsets in compute nodes sport more cores and more memory, consolidation factors go up and the array's disks, service processors, and network links become increasingly congested. Already today, businesses running virtualized desktops must work around the "boot storm" that occurs every morning when hundreds of employees try to simultaneously power on their VM-hosted desktops. Because it is a single point of failure, the array must also be constructed out of highly reliable components, which makes it expensive compared to commodity hardware. When using shared storage, scalability comes at a high cost.

As a potential alternative to shared storage, we propose a distributed storage system that makes VM storage location-independent and exploits the local storage capacity of compute nodes, hereby increasing flexibility and lowering the cost of virtual machine-based cloud computing. Our definition of a cloud is not limited to a single data center, but covers the spectrum from smartphones and laptops running VMs to dedicated local infrastructure combined with shared resources rented from utility providers. Therefore, our proposed system supports both strict and eventual consistency models, to allow consistent periodic off-site backups, long-distance VM migration, and off-line use of VMs running on disconnected devices such as laptops and smartphones. Our contribution is that we demonstrate the feasibility of self-certifying and eventually consistent replicated storage for performance-sensitive virtualization workloads; workloads that previously depended on specialized shared-storage hardware. Practical innovations include a local storage engine optimized for strict and eventual consistency replication with built-in robustness to data corruption, disk and host failures, and a novel cryptographic locking mechanism that obviates the need for centralized or distributed lock management that could otherwise limit scalability.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
SoCC'10, June 10–11, 2010, Indianapolis, Indiana, USA.
Copyright 2010 ACM 978-1-4503-0036-0/10/06 ...$10.00.

1.1 Background

This work describes a distributed block storage system that provides a backwards-compatible SCSI disk emulation to a potentially large set of VMs running in a data center or cloud hosting facility. The aim is not to replace high-end storage arrays for applications that require the combined throughput of tens or hundreds of disk spindles, but rather to address the scalability and reliability needs of many smaller VMs that individually would run acceptably from one or a few local disk drives, but today require shared storage for taking advantage of features such as seamless VM migration. Examples include hosted virtual desktop VMs, web and email servers, and low-end online transaction processing (OLTP) applications. To allow VMs to survive disk or host failures, our proposed system uses peer-to-peer replication at the level of individual per-VM storage volumes. Though replication systems are well-studied and understood, the construction of a replication system for live VMs is complicated by practical issues such as disk scheduler non-determinism and disk or wire data corruption, both of which can result in silent replica divergence, and by performance interference from multiple workloads contending for local and remote disk heads, which can greatly inhibit individual VM storage throughput.

While a number of scalable storage systems have been proposed, most are only optimized for certain access patterns that do not fit the requirements of VM storage. The Google File System [16], for instance, is optimized for append-only workloads, and other systems only store simple key-value pairs of limited size [15, 42], or require objects to be immutable once inserted [37, 13, 3]. OceanStore's Pond [35] and Antiquity [47] prototypes come close to meeting the requirements of VM storage, but Pond's use of expensive erasure coding results in dismal write performance, and Antiquity's *two-level naming* interface is not directly applicable to virtual disk emulation. Amazon's Elastic Block Store [2] is a distributed storage system for VMs, but its durability guarantees, to the best of our knowledge, hinge on regular manual snapshot backups to Amazon S3, resulting in a trade-off between durability and performance. Section 5 surveys more related work, but based on the above, we believe that the design space for scalable distributed VM storage is still substantially uncharted and that a "blank slate" design approach is warranted.

2. DESIGN

Our proposed system, Lithium, is a scalable distributed storage system optimized for the characteristics of virtual machine IO: random and mostly exclusive block accesses to large and sparsely populated virtual disk files. For scalability reasons, Lithium does not provide a globally consistent POSIX name space, but names stored objects in a location-independent manner using globally unique incremental state hashes, similar to a key-value store but where each value is a large 64-bit address space. Lithium supports instant volume creation with lazy space allocation, instant creation of writable snapshots, and replication of volumes and snapshots with tunable consistency ranging from per-update synchronous replication to eventual consistency modes that allow VMs to remain available during periods of network congestion or disconnection.

Lithium's underlying data format is self-certifying and self-

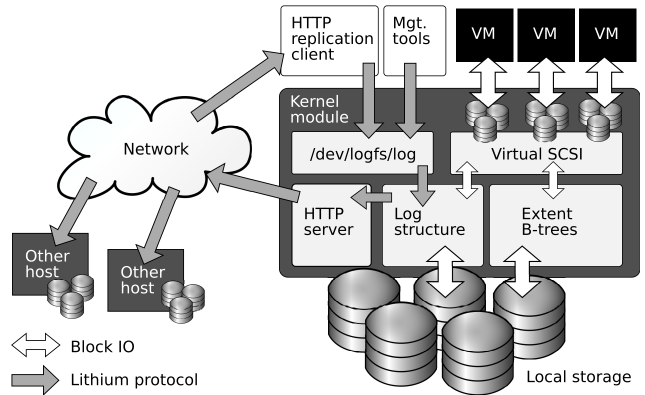

Figure 1: Overview of the components of Lithium on each host. VMs access the Lithium hypervisor kernel module through a virtual SCSI interface. The log-structured storage engine lives in the kernel, and policy-driven tools append to the log from user-space through the `/dev/logfs/log` device node.

sufficient, in the sense that data can be stored on unreliable hosts and disk drives without fear of silent data corruption, and all distributed aspects of the protocol are persisted as part of the core data format. This means that Lithium can function independently, without reliance on any centralized lock management or meta-data service, and can scale from just a couple of hosts to a large cluster. The design is compatible with commodity off-the-shelf server hardware and, thus does not require non-volatile RAM or other "magic" hardware.

Lithium consists of a kernel module for the VMware ESX hypervisor, and of a number of small user space tools that handle volume creation, branching, and replication. Figure 1 shows the main components installed on a Lithium-enabled host. All data is made persistent in a single on-disk log-structure that is indexed by multiple B-trees. The kernel module handles performance-critical network and local storage management, while user space processes that interface with the kernel module trough a special control device node, are responsible for policy-driven tasks such as volume creation, branching, replication, membership management, and fail-over. This split responsibilities design allows for more rapid development of high-level policy code, while retaining performance on the critical data path.

In Lithium, data is stored in *volumes* – 64-bit block address spaces with physical disk drive semantics. However, a volume is also akin to a file in a traditional file system, in that space is allocated on demand, and that the volume can grow as large as the physical storage on which it resides. Volumes differ from files in that their names are 160-bit unique identifiers chosen at random from a single flat name space. A VM's writes to a given volume are logged in the shared log-structure, and the location of the logical blocks written recorded in the volume's B-tree. Each write in the log is prefixed by a commit header with a strong checksum, update ordering information, and the logical block addresses affected by the write. Writes can optionally be forwarded to other hosts for replication, and replica hosts maintain their own log and B-trees for each replicated volume, so that the replica hosts are ready to assume control over the volume, when necessary.

A new volume can be created on any machine by appending

an appropriate log entry to the local log through the control device node. The new volume is identified by a permanent base id (chosen at random upon creation) and a current version id. The version id is basically an incremental hash of all updates in the volume's history, which allows for quick replica integrity verification by a simple comparison of version ids.

Other hosts can create replicas of a volume merely by connecting to a tiny HTTP server inside the Lithium kernel module. They can sync up from a given version id, and when completely synced, the HTTP connection switches over to synchronous replication. As long as the connection remains open, the volume replicas stay tightly synchronized. The use of HTTP affects only the initial connection setup and does not result in additional network round-trips or latencies.

The host that created the volume starts out as the *primary* (or *owner*) for that volume. The primary serializes access, and is typically the host where the VM is running. Ownership can be transferred seamlessly to another replica, *e.g.*, when a VM migrates, which then becomes the new primary. The integrity and consistency of replicas is protected by a partially ordered data model known as *fork-consistency*, described next.

2.1 Fork-consistent Replication

Like a number of other recent research systems [26, 17, 47], updates in Lithium form a hash-chain, with individual updates uniquely identified and partially ordered using cryptographic hashes of their contexts and contents. This technique is often referred to as "fork-consistency" and also forms the basis of popular distributed source code versioning tools such as Mercurial [24]. Use of a partial ordering rather than a total ordering removes the need for a central server acting as a coordinator for update version numbers, and allows for simple distributed branching of storage objects, which in a storage system is useful for cloning and snapshotting of volumes. The use of strong checksums of both update contents and the accumulated state of the replication state machine allows for easy detection of replica divergence and integrity errors ("forks"). Each update has a unique id and a parent id, where the unique id is computed as a secure digest of the parent id concatenated with the contents of the current update, *i.e.*, $id = h(parentid||updatecontents)$, where $h()$ is the secure digest implemented by a strong hash function (SHA-1 in our implementation). By having updates form a hash-chain with strong checksums, it becomes possible to replicate data objects onto untrusted and potentially Byzantine hardware; recent studies have found such Byzantine hardware surprisingly common [4]. Figure 2 shows how each volume is stored as a chain of updates where the chain is formed by backward references to parent updates. Fork-consistency allows a virtual disk volume to be mirrored anywhere, any number of times, and allows anyone to clone or snapshot a volume without coordinating with other hosts. Snapshots are first-class objects with their own unique base and version ids, and can be stored and accessed independently.

Lithium uses replication rather than erasure coding for redundancy. Erasure coding is an optimization that saves disk space and network bandwidth for replicas, but also complicates system design, multiplies the number of network and disk IOs needed to service a single client read, and can lead to new problems such as TCP throughput collapse [34]. Nevertheless, if necessary, we expect that erasure codes or other bandwidth-saving measures can be added to Lithium, *e.g.*, merely by adding a user-space tool that splits or transforms the HTTP update stream before data gets relayed to other hosts.

2.2 Replica Consistency

General state-machine replication systems have been studied extensively and the theory of their operation is well understood [40]. In practice, however, several complicating factors make building a well-performing and correct VM disk replication system less than straight-forward. Examples include parallel queued IOs that in combination with disk scheduler nondeterminism can result in replica divergence, and network and disk data corruption that if not checked will result in integrity errors propagating between hosts. Finally, the often massive sizes of the objects to be replicated makes full resynchronization of replicas, for instance to ensure consistency after a crash, impractical.

In a replication system that updates data in place, a classic two-phase commit (2PC) protocol [12] is necessary to avoid replica divergence. Updates are first logged out-of-place at all replicas, and when all have acknowledged receipt, the *head* node tells them to destructively commit. Unfortunately, 2PC protocols suffer a performance overhead from having to write everything twice. If 2PC is not used, a crash can result in not all replicas applying updates that were in flight at the time of the crash. Disk semantics require an update that has been acknowledged back to the application to be reflected onto the disk always, but updates may be on disk even if they have never been acknowledged. It is acceptable for the replicas to diverge, as long as the divergence is masked or repaired before data is returned to the application. In practice, this means that the cost of 2PC can be avoided as long as replicas are properly resynchronized after a crash. To avoid a full re-sync of replicas after a crash, previous disk replication systems, such as Petal [21], have used a dirty-bitmap of unstable disk regions where writes have recently been outstanding. The main problem with this approach is that it requires a bitmap per replica, which can be cumbersome, if the replication degree is high. Furthermore, this stateful approach requires persistent tracking of which replica was the primary at the time of the crash, to prevent an old replica overwriting a newer one.

In contrast, Lithium's replication protocol is stateless and supports an unlimited amount of both lagging and synchronous replicas. While 3-way replication may provide sufficient reliability in most use cases, higher-degree replication may sometimes be desired, *e.g.*, for read-only shared system base disks used by many VMs. Bitmaps are impractical for high-degree replication. Instead, Lithium uses update logging to support Bayou-like eventual consistency [43].

The Lithium replication mechanism is similar to the first phase of 2PC in that updates are logged non-destructively. Once an update has been logged at a quorum of the replicas, the write IO is acknowledged back to the application VM. In eventual-consistency mode, IOs are acknowledged when they complete locally. Because updates are non-destructive, crashed or disconnected replicas can sync up by replaying of missing log updates from peers. Data lives permanently in the log, so the protocol does not require a separate commit-phase. Writes are issued and serialized at the primary replica, typically where the VM is running. We make the simplifying

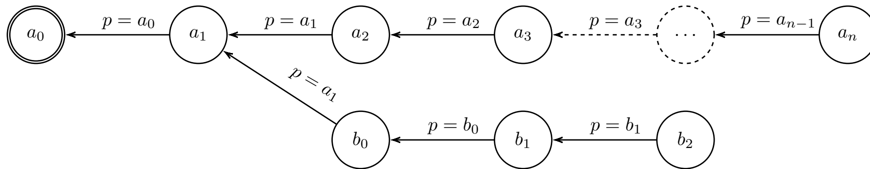

Figure 2: Partial update ordering and branches in a traditional fork-consistent system. The partial ordering is maintained through explicit "parent" back-references, and branches can be expressed by having multiple references to the same parent.

assumption that writes always commit to the primary's log, even when the VM tries to abort them or the drive returns a non-fatal error condition. We catch and attempt to repair drive errors (typically retry or busy errors) by quiescing IO to the drive and retrying the write, redirecting to another part of the disk, if necessary. This assumption allows us to issue IO locally at the primary and to the replicas in parallel. We also allow multiple IOs to be outstanding, which is a requirement for exploiting the full disk bandwidth. The strong checksum in our disk format commit headers ensures that incomplete updates can be detected and discarded. This means that we do not require a special "end-of-file" symbol: we can recover the end of the log after a crash simply by rolling forward from a checkpoint until we encounter an incomplete update. To avoid problems with "holes" in the log, we make sure always to acknowledge writes back to the VM in log order, even if reordered by the drive. For mutable volumes, reads are always handled at the primary, whereas immutable (snapshot) volumes can be read-shared across many hosts, *e.g.*, to provide scalable access to VM template images.

The on-disk format guarantees temporal ordering, so it is trivial to find the most recent out of a set of replicas, and because disk semantics as described above are sufficient, it is always safe to roll a group of replicas forward to the state of the most recent one. We use an access token, described in section 2.4, to ensure replica freshness and mutual exclusion with a single combined mechanism.

2.3 Log-structuring

On each host, Lithium exposes a single on-disk log-structure as a number of volumes. Each entry in the log updates either block data or the meta-data of a volume. The state of each volume equals the sum of its updates, and a volume can be recreated on another host simply by copying its update history. Each host maintains per-volume B-tree indexes to allow random accesses into each volume's 64-bit block address space. Volumes can be created from scratch, or as branches from a base volume snapshot. In that case, multiple B-trees are chained, so that reads "fall through" empty regions in the children B-trees and are resolved recursively at the parent level. Branching a volume involves the creation of two new empty B-trees, corresponding to the left and right branches. Branching is a constant-time near-instant operation. Unwritten regions consume no B-tree or log space, so there is no need to pre-allocate space for volumes or to zero disk blocks in the background to prevent information leakage between VMs.

The disk log is the only authoritative data store used for holding hard state that is subject to replication. *All* other state is host-local soft-state that in the extreme can be recreated from the log. As such, Lithium may be considered a log-structured file system (LFS) [36] or a scalable collection of logical disks [14]. Log-structuring is still regarded by many as a purely academic construct hampered by cleaning overheads and dismal sequential read performance. For replicated storage, however, log-structuring has a number of attractive properties:

Temporal Order Preservation: In a replication system, the main and overriding attraction of an LFS is that updates are non-destructive and temporally ordered. This greatly simplifies the implementation of an eventually consistent [43] replication system, for instance to support disconnected operation [19] or periodic asynchronous replication to off-site locations.

Optimized for Random Writes: Replication workloads are often write-biased. For example, a workload with a read-skewed 2:1 read-to-write ratio without replication is transformed to a write-skewed 3:2 write-to-read ratio by triple replication. Furthermore, parallel replication of multiple virtual disk update streams exhibits poor locality because each stream wants to position the disk head in its partition of the physical platter, but with a log-structured layout all replication writes can be made sequential, fully exploiting disk write bandwidth. Other methods, such as erasure-coding or de-duplication, can reduce the bandwidth needs, but do not by themselves remedy the disk head contention problems caused by many parallel replication streams. While it is true that sequential read performance may be negatively affected by log-structuring, large sequential IOs are rare when multiple VMs compete for disk bandwidth.

Live Storage Migration: Virtual machine migration [10, 32] enables automated load-balancing for better hardware utilization with less upfront planning. If the bindings between storage objects and hardware cannot change, hotspots and bottlenecks will develop. Just like with *live* migration of VMs, it is beneficial to keep storage objects available while they are being copied or migrated. A delta-oriented format such as the one used by Lithium simplifies object migration because changes made during migration can be trivially extracted from the source and appended to the destination log.

Flash-compatible: Whether or not all enterprise-storage is going to be Flash-based in the future is still a subject of debate [31], but some workloads clearly benefit from the access latencies offered by Flash SSDs. Because solid-state memories are likely to remain more expensive than disks, using them for backup replicas is generally considered a waste. A good compromise could be a hybrid approach where the primary workload uses Flash, and cheaper disk storage is used for backup replicas. However, to keep up with a Flash-based primary, the replicas will have to be write-

optimized, *i.e.*, log-structured. Because of the low access latencies and other inherent characteristics of Flash, log-structuring also benefits the Flash-based primary [5]. The downside to log-structuring is that it creates additional fragmentation, and that a level of indirection is needed to resolve logical block addresses into physical block addresses. Interestingly, when we started building Lithium, we did not expect log-structured logical disks to perform well enough for actual hosting of VMs, but expected them to run only as live write-only backups at replicas. However, we have found that the extra level of indirection rarely hurts performance, and now use the log format both for actual hosting of VMs and for their replicas.

2.4 Mutual Exclusion

We have described the fork-consistency model that uses cryptographic checksums to prevent silent replica divergence as a result of silent data corruption, and that allows branches of volumes to be created anywhere and treated as first-class replicated objects. We now describe our extensions to fork consistency that allow us to emulate shared-storage semantics across multiple volume replicas, and that allow us to deal with network and host failures. While such functionality could also have been provided through an external service, such as a distributed lock manager, a mechanism that is integrated with the core storage protocol is going to be more robust, because a single mechanism is responsible for ordering, replicating, and persisting both data and mutual exclusion meta-data updates. This is implemented as an extension to the fork-consistency protocol by adding the following functionality:

- A mutual exclusion *access token* constructed by augmenting the partial ordering of updates, using one-time secrets derived from a single master key referred to as the *secret view-stamp*.
- A fallback-mechanism that allows a majority quorum of the replicas to recover control from a failed primary by recovering the master key through secret sharing and a voting procedure.

2.4.1 Access Token

For emulating the semantics of a shared storage device, such as a SCSI disk connected to multiple hosts, fork-consistency alone is insufficient. VMs expect either "read-your-writes consistency", where all writes are reflected in subsequent reads, or the weaker "session-consistency", where durability of writes is not guaranteed across crashes or power losses (corresponding to a disk controller configured with write-back caching). In Lithium, the two modes correspond to either synchronous or asynchronous eventually-consistent replication. In the latter case, we ensure session consistency by restarting the VM after any loss of data.

Our SCSI emulation supports mutual exclusion through the `reserve` and `release` commands that lock and unlock a volume, and our update protocol implicitly reflects ownership information in a volume's update history, so that replicas can verify that updates were created by a single exclusive owner. To ensure that this mechanism is not the weakest link of the protocol, we require the same level of integrity protection as with regular volume data. To this end, we introduce the notion of an *access token*, a one-time secret key that must be known to update a volume. The access token is constructed so that only an up-to-date replica can use it to mutate its own copy of the volume. If used on the wrong volume, or on the wrong version of the correct volume, the token is useless.

Every time the primary replica creates an update, a new secret token is generated and stored locally in volatile memory. Only the current primary knows the current token, and the only way for volume ownership to change is by an explicit token exchange with another host. Because the token is a one-time password, the only host that is able to create valid updates for a given volume is the current primary. Replicas observe the update stream and can verify its integrity and correctness, but lack the necessary knowledge for updating the volume themselves. At any time, the only difference between the primary and the replicas is that only the primary knows the current access token.

The access token is constructed on top of fork-consistency by altering the partial ordering of the update hash chain. Instead of storing the clear text id of a version in the log entry, we include a randomly generated secret *view-stamp* s in the generation of the id for an update, so that

$$\begin{array}{rrcl} \text{instead of:} & id & = & h(parentid||update) \\ \text{we use:} & id & = & h(s||parentid||update) \end{array}$$

where h is our strong hash function. The secret view-stamp s is known only by the current primary. In the update header, instead of storing the clear text value of id, we store $h(id)$ and only keep id in local system memory (because $h()$ is a strong hash function, guessing id from $h(id)$ is assumed impossible). Only in the successor entry's *parentid* field do we store id in clear text. In other words, we change the pairwise ordering

$$\begin{array}{rlcl} \text{from:} & a < b & \text{iff} & a.id = b.parentid \\ \text{to:} & a < b & \text{iff} & a.id = h(b.parentid) \end{array}$$

where $<$ is the direct predecessor relation. The difference here being that in order to name a as the predecessor to b, the host creating the update must know $a.id$, which only the primary replica does. Other hosts know $h(a.id)$, but not yet $a.id$. We refer to the most recent id as the *secret access token* for that volume. The randomly generated s makes access tokens unpredictable. While anyone can replicate a volume, only the primary, which knows the current access token id, can mutate it. Volume ownership can be passed to another host by exchange of the most recent access token. Assuming a correctly implemented token exchange protocol, replica divergence is now as improbable as inversion of $h()$. Figure 3 shows how updates are protected by one-time secret access tokens. The new model still supports decentralized branch creation, but branching is now explicit. When the parent for an update is stated as described above, the update will be applied to the existing branch. If it is stated as in the old model, *e.g.*, $b.parentid = a.id$, then a new branch will be created and the update applied there.

2.4.2 Automated Fail-over

A crash of the primary results in the loss of the access token, rendering the volume inaccessible to all. To allow access to be restored, we construct on top of the access token a view-change protocol that allows recovery of s by a remaining majority of nodes. A recovered s can be used to also recover id. To allow recovery of s, we treat it as a secret view-stamp [33], or epoch number, that stays constant as long as the same host is the primary for the volume, and we

use secret sharing to disperse slices of s across the remaining hosts. Instead of exchanging the access token directly, volume ownership is changed by appending a special *view-change* record to the log. The view-change record does not contain s but its public derivate $h(s)$, along with slices of s, encrypted under each replica's host key. Thus $h(s)$ is the *public view-stamp* identifying that view.

If a host suspects that the current primary has died, it generates a new random s' and proposes the corresponding view $h(s')$ to the others. Similar to Paxos [20], conflicting proposals are resolved in favor of the largest view-stamp. Hosts cast their votes by returning their slices of s, and whoever collects enough slices to recover s wins and becomes the new primary. The new primary now knows both s and s' so it can reconstruct *id* using the recovered s, and append the view-change that changes the view to $h(s')$. When there is agreement on the new view, IO can commence. The former s can now be made public to revoke the view $h(s)$ everywhere. If there are still replicas running in the former view $h(s)$, the eventual receipt of s will convince them that their views are stale and must be refreshed. Should the former primary still be alive and eventually learn s, it knows that it has to immediately discard of its copy of the VM and any changes made (if running in eventual consistency mode) after the point of divergence. Discarding and restarting the VM in this case ensures session-consistency.

The use of the secret token protects against replica divergence as a result of implementation or data corruption errors. Prevention against deliberate attacks would require signing of updates, which is straightforward to add. Using signatures alone is not enough, because a signature guarantees authenticity only, not freshness or mutual exclusion as provided by the access token. We have found that access tokens are a simple and practical alternative to centralized lock servers. During development, the mechanism has helped identify several implementation and protocol errors by turning replica divergence into a fail-stop condition.

3. IMPLEMENTATION DETAILS

Lithium treats each local storage device as a single log that holds a number of virtual volumes. Write operations translate directly into synchronous update records in the log, with the on-disk location of the update data being recorded in a B-tree index. The update may also propagate to other hosts that replicate the particular volume. Depending on configuration, the write is either acknowledged back to the VM immediately when the update is stable in the local log, or when all or a quorum of replicas have acknowledged the update. The first mode allows for disconnected operation with eventual consistency, while the second mode corresponds to synchronous replication.

When storing data on disk and when sending updates over the wire, the same format is used. A 512-byte log entry commit header describes the context of the update (its globally unique id and the name of its parent), the location of the update inside the virtual disk's 64-bit address space, the length of the extent, and a strong checksum to protect against data corruption. The rest of the header space is occupied by a bit vector used for compressing away zero blocks, both to save disk space and network bandwidth (in practice, this often more than offsets the cost of commit headers), and to simplify log-compaction. For practical reasons, the on-disk log is divided into fixed size segments of 16MB each.

Workload	Type	Avg. len	Overhead
Windows XP (NTFS)	FS	29.9	0.26%
IOZone (XFS)	FS	33.8	0.23%
PostMark (XFS)	FS	61.8	0.13%
DVDStore2 (MySQL)	DB	24.0	0.33%
4096-byte random	Synthetic	8.0	1.0%
512-byte random	Synthetic	1.0	7.8%

Table 1: Extent Sizes and B-tree Memory Overhead

3.1 B-tree Indexes

Applications running on top of Lithium are unaware of the underlying log-structure of physical storage, so to support random access reads, Lithium needs to maintain an index that maps between the logical block addresses (LBAs) used by VMs, and the actual physical locations of the data on disk. This index could be implemented as a per-volume lookup table with an entry for each logical block number in the volume, but because each Lithium volume is a 64-bit block address space, a lookup table might become prohibitively large. Instead, Lithium uses a 64-bit extent indexed B-tree to track logical block locations in the log. The B-tree is a more complex data structure than a simple lookup table or a radix tree, but is also more flexible, and designed to perform well on disk. The B-tree index does not track individual block locations, but entire extents. If a VM writes a large file into a single contiguous area on disk, this will be reflected as just a single key in the B-tree. Figure 4 shows an example B-tree index with block ranges that correspond to different log entries. In the pathological worst case, the VM may write everything as tiny, random IOs, but we have found that on average the use of extents provides good compression. For instance, a base Windows XP install has an average extent length of 29.9 512-byte disk sectors. Each B-tree extent key is 20 bytes, and B-tree nodes are always at least half full. In this case the space overhead for the B-tree is at most $\frac{20 \times 2}{512 * 29.9} \approx 0.26\%$ of the disk space used by the VM. Table 1 lists the average extent sizes and resulting B-tree memory overheads we have encountered during development, along with the pathological worst cases of densely packed 4kB and 512B completely random writes. Guest OS disk schedulers attempt to merge adjacent IOs, which explains the large average extent sizes of the non-synthetic workloads. We demand-page B-tree nodes and cache them in main memory (using pseudo-LRU cache replacement), and as long as a similar percentage of the application's disk working set fits in the cache, performance is largely unaffected by the additional level of indirection.

We carefully optimized the B-tree performance: Most IO operations are completely asynchronous, the tree has a large fan-out (more than a thousand keys per tree node), and updates are buffered and applied in large batches to amortize IO costs. The B-tree uses copy-on-write updates and is crash-consistent; it is only a cache of the authoritative state in the log: should it be damaged, it can always be recreated. The B-tree itself is not replicated, but if two hosts replicate the same volume, they will end up with similar B-trees. The B-trees are always kept up to date at all replicas, so control of a volume can migrate to another host instantly. Periodic checkpoints of B-trees and other soft state ensure constant time restart recovery.

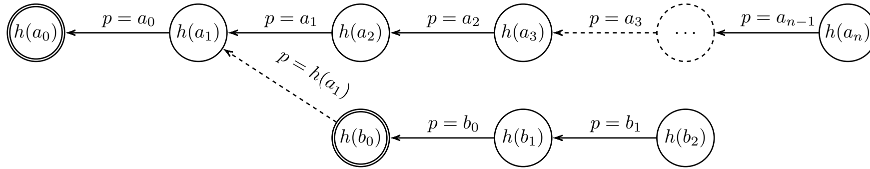

Figure 3: Update ordering and branches in Lithium. In contrast to the original fork-consistent model in Figure 2, appending an update to an existing branch requires knowledge of its clear-text id to state it as a parent reference. Branching is an explicit but unprivileged operation.

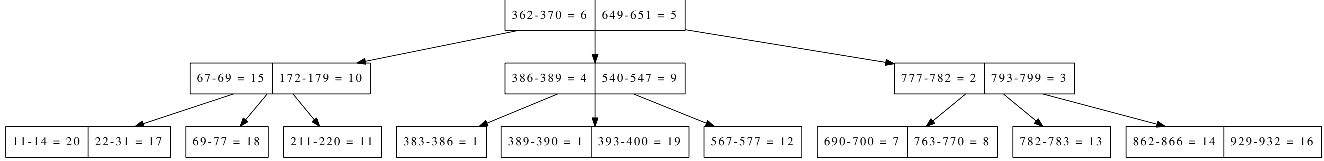

Figure 4: Simplified example of the B-tree used for mapping logical block offsets to version numbers in the log. Block number 765 would resolve to revision number 8, and block 820 would return NULL, indication that the block were empty. Multiple trees can be chained together, to provide snapshot and cloning functionality.

3.2 Log Compaction

A local log compactor runs on each host in parallel with other workloads, and takes care of reclaiming free space as it shows up. Free space results from VMs that write the same logical blocks more than once, and the frequency with which this occurs is highly workload-dependent. Free space is tracked per log-segment in an in-memory priority queue, updated every time a logical block is overwritten in one of the volume B-trees, and persisted to disk as part of periodic checkpoints. When the topmost n segments in the queue have enough free space for a compaction to be worthwhile, the log compactor reads those n segments and writes out at most $n-1$ new compacted segments, consulting the relevant B-trees to detect unreferenced blocks along the way. Unreferenced blocks are compressed away using the zero block bit vector described previously. The log compactor can operate on arbitrary log segments in any order, but preserves the temporal ordering of updates as not to conflict with eventual consistency replication. It runs in parallel with other workloads and can be safely suspended at any time. In a real deployment, one would likely want to exploit diurnal workload patterns to run the compaction only when disks are idle, but in our experiments, we run the compactor as soon as enough free space is available for compaction to be worthwhile.

Due to compaction, a lagging replica that syncs up from a compacted version may not see the complete history of a volume, but when the replica has been brought up to the compacted version, it will have a complete and usable copy. Though we remove overwritten data blocks, we currently keep their empty update headers to allow lagging replicas to catch up from arbitrary versions. In the future, we plan to add simple checkpoints to our protocol to avoid storing update headers perpetually.

3.3 Locating Updates on Disk

The partial ordering of log updates has benefits with regards to cheap branches and simplifying exchange of updates. However, with 160 bits of address space and multiple terabytes of log space, finding the successor to an arbitrary volume revision can be a challenge. This may be necessary when a lagging replica connects, because we only maintain soft state about replicas and their current versions, soft state that may be lost in a crash or forgotten due to a replica having been off-line. Unfortunately, the lack of locality in our version identifiers precludes the use of any dictionary data structure (*e.g.*, tree or hash table) that relies on key locality for update performance. On the other hand, arbitrary successor lookups are rare and out of the critical data path, so it is acceptable if they are much slower than updates. The brute force solution of an exhaustive scan of the disk surface is prohibitively expensive. Instead, we speed up the search by maintaining a Bloom filter [6] that describes the contents of each log segment on the disk. A Bloom filter is a simple bit vector indexed by k uniform hash functions. For a membership query to return a positive result, all bits corresponding to the results of applying the k hash functions to the key being looked up, must be set. There is a risk of false positives, but the value of k and the size of the bit vector m can be tuned to make the risk acceptably small for a given number of elements n. For our application, false positives do not affect correctness, but only performance.

A Bloom filter can compactly represent a set of revision identifier hashes, but can only answer membership queries, not point to the actual locations of data. We therefore keep a Bloom filter per 16MB log segment, and exhaustively scan all filters on a disk when looking for a specific version. If more than one filter claims to have the version, we perform a deep-scan of each corresponding log segment until we have an exact match. Each Bloom filter has $m = 2^{16}$ bits (8,192 bytes) and is indexed by $k = 10$ hash functions, corresponding to the ten 16-bit words of the 160-bit SHA-1 hash that is looked up. At the expected rate of up to $n = 4,000$ log entries per segment, the filter has a false positive rate of $(1-(1-\frac{1}{m})^{kn})^k \approx 0.04\%$. An additional Bloom filter at the end of each segment, indexed by another set of hash functions, serves to square this probability, so that we rarely have to deep-scan segments in vain.

A one terabyte disk contains almost 60,000 16MB log segments, corresponding to about 480MB of Bloom filters. As-

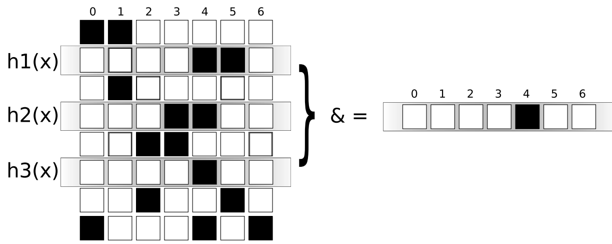

Figure 5: Combining $k = 3$ entire rows of the inverted index to search multiple filters in parallel. A match is found in column 4, corresponding to log segment 4 on disk.

suming a disk read rate of 100MB/s, these can be scanned in less than 5s. However, we can do better if instead of storing each filter separately, we treat a set of filters as an *inverted index*, and store them as the columns of a bit-matrix. The columns are gradually populated in an in-memory buffer, which when full gets flushed to disk in a single write operation. Instead of reading the entire matrix for a lookup, we now only have to read the k rows corresponding to $h_1(key), \ldots, h_k(key)$, as illustrated in Figure 5. With a 16MB matrix of 65,536 rows and 2,048 columns we are able to cover 2,048 log segments, corresponding to 32GB of physical log space. Our measurements show that, including the amortized disk write cost, we are able to insert one million keys into the inverted index per second, and that we are able to scan indexes covering between 700MB and 1TB of log space per second, with a single drive, depending on the degree to which filters have been filled (short-circuit optimization applies to the AND'ing of rows).
Multiple physical disks can be scanned in parallel to maintain scalability as disks are added. We can also increase the scanning speed by adding more columns to each index, at the cost of using more main memory for the current inverted index. For instance, a 32MB index will store 4,096 columns and double the scanning speed. As disks grow and memory becomes cheaper, the index size may be adjusted accordingly.

3.4 POSIX Interface

Lithium provides block-addressable object storage, but by design does not expose a global POSIX name space. Like shared memory, faithfully emulating POSIX across a network is challenging and introduces scalability bottlenecks that we wish to avoid. The VMs only expect a block-addressable virtual disk abstraction, but some of the tools that make up our data center offering need POSIX, *e.g.*, for configuration, lock, and VM-swap files. Each VM has a small collection of files that make up the runtime and configuration for that VM, and these files are expected to be kept group-consistent. Instead of rewriting our management tools, we use the VMware VMFS [45] cluster file system to provide a per-volume POSIX layer on top of Lithium. Because Lithium supports SCSI reservation semantics, VMFS runs over replicated Lithium volumes just as it would over shared storage. Popular VMware features like "VMotion" (live VM migration) and "High Availability" (distributed failure detection and automated fail-over) work on Lithium just like they would on shared storage. In a local network, the token exchange protocol described in section 2.4 is fast enough to allow hundreds of control migrations per second, so VMFS running over a replicated Lithium volume feels very similar to VMFS running over shared storage.

4. EVALUATION

Our evaluation focuses on the performance of the prototype when configured for synchronous 2 and 3-way replication. To measure the performance of our prototype we ran the following IO benchmarks:

PostMark PostMark is a file-system benchmark that simulates a mail-server workload. In each VM we ran PostMark with 50,000 initial files and 100,000 transactions. PostMark primarily measures IO performance.

DVDStore2 This online transaction processing benchmark simulates an Internet DVD store. We ran the benchmark with the default "small" dataset, using MySQL as database backend. Each benchmark ran for three minutes using default parameters. DVDStore2 measures a combination of CPU and IO performance.

Our test VMs ran Linux 2.6.25 with 256MB RAM and 8GB of virtual storage. Data volumes were stored either in Lithium's log-structure, or separately in discrete 8GB disk partitions. When running Lithium, log compaction ran eagerly in parallel with the main workload, to simulate a disk full condition. Lithium was configured with 64MB of cache for its B-trees. Apart from the settings mentioned above, the benchmarks were run with default parameters.

4.1 Replication Performance

Replication performance was one of the motivations for the choice of a write-optimized disk layout, because it had been our intuition that multiple parallel replication streams to different files on the same disk would result in poor throughput. When multiple VMs and replication streams share storage, most IO is going to be non-sequential. Apart from the caching that is already taking place inside the VMs, there is very little one can do to accelerate reads. Writes, however, can be accelerated through the use of log-structuring.
In this benchmark, we used three HP Proliant DL385 G2 hosts, each equipped with two dual-core 2.4GHz Opteron processors, 4GB RAM, and a single physical 2.5" 10K 148GB SAS disk, attached to a RAID controller without battery backed cache. We used these hosts, rather than the newer Dells described below, because they had more local storage available for VMs and replicas, and were more representative of commodity hardware. As a strawman, we modified the Lithium code to use a direct-access disk format with an 8GB *partition* of the disk assigned to each volume, instead of the normal case with a single shared log-structure. We ran the DVDStore2 benchmark, and used the total orders-per-minute number per host as the score. The results are shown in Figure 6, and clearly demonstrate the advantage of a write-optimized layout. Per-host throughput was 2–3 times higher when log-structuring was used instead of static partitioning.

4.2 Scalability

To test Lithium's scalability in larger clusters, we obtained the use of eight Dell 2950 III servers, each with 16GB RAM, two quad-core 3GHz Xeon processors, and with three 32GB local SAS drives in RAID-5 configuration, using a Dell perc5/i

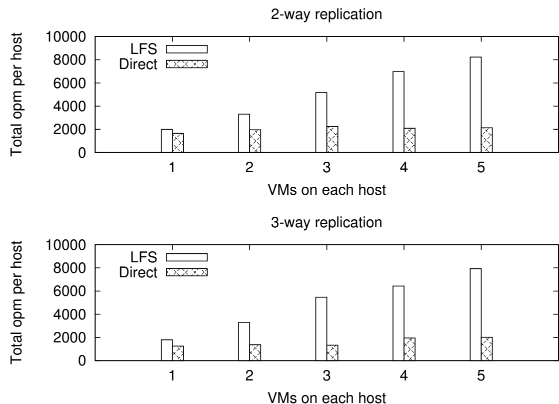

Figure 6: Performance of the DVD Store 2 OLTP benchmark in a two- and three-host fully replicated scenarios.

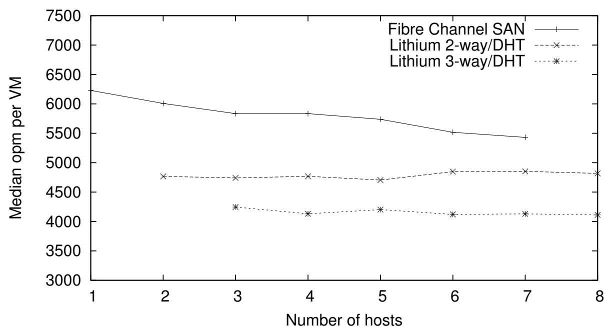

Figure 7: DVDStore2 Scalability. This graph shows the median orders per minute per VM, with 4 VMs per host for a varying number of hosts of Lithium in 2-way and 3-way replicated scenarios, and for a mid-range Fibre Channel storage array.

controller with 256MB of battery backed cache. Naturally this small cluster cannot be considered cloud-scale, but it was connected to a mid-range (US$100,000 price range) Fibre Channel storage array, which provided a good reference for comparison against finely tuned enterprise hardware.

The array was configured with 15 15k Fibre Channel disks in a RAID-5 configuration, presented as single LUN mounted at each ESX host by VMFS. The array was equipped with 8GB of battery backed cache, and connected to seven of the eight hosts through 4GB/s Fibre Channel links. Lithium was configured with either 2- or 3-way replication over single gigabit Ethernet links, in addition to the local RAID-5 data protection. Replicas were mapped to hosts probabilistically, using a simple distributed hash table (DHT) constructed for the purpose. In this setup, Lithium was able to tolerate three or five drive failures without data loss, where the array could only tolerate a single failure. Ideally, we would have reconfigured the local drives of the servers as RAID-0 for maximum performance, as Lithium provides its own redundancy. However, the hardware was made available to us on a temporary basis, so we were unable to alter the configuration in any way.

We ran the DVDStore2 benchmark in 4 VMs per host, and varied the number of hosts. As Figure 7 shows, a three-drive

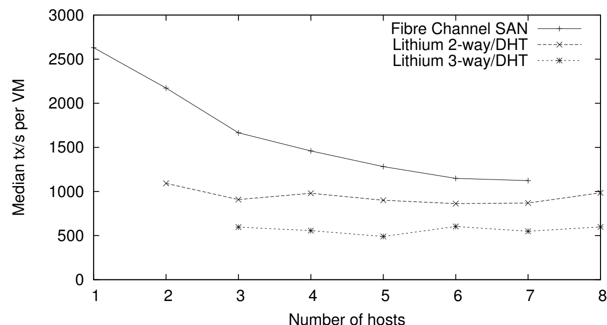

Figure 8: PostMark Scalability. Median transactions per second per VM, with 4 VMs per host for a varying number of hosts of Lithium in 2-way and 3-way replicated scenarios, and for a mid-range Fibre Channel storage array.

RAID-5 per host is not enough to beat the high-performance array, but as more hosts are added, per-VM throughput remains constant for Lithium, whereas the array's performance is finite, resulting in degrading per-VM throughput as the cluster grows.

We also ran the PostMark benchmark (see Figure 8) in the same setting. PostMark is more IO intensive than DVD-Store2, and median VM performance drops faster on the storage array when hosts are added. Unfortunately, only seven hosts had working Fibre Channel links, but Lithium is likely to be as fast or faster than the array for a cluster of eight or more hosts.

Though our system introduces additional CPU overhead, *e.g.*, for update hashing, the benchmarks remained disk IO rather than CPU-bound. When we first designed the system, we measured the cost of hashing and found that a hand-optimized SHA-1 implementation, running on a single Intel Core-2 CPU core, could hash more than 270MB of data per second, almost enough to saturate the full write bandwidth of three SATA disk drives. In our view the features that the hashing enables more than warrant this CPU overhead.

As stated earlier, we did not expect Lithium to outperform finely tuned and expensive centralized systems in small-scale installations. We designed Lithium to have reasonable local performance, but first and foremost to provide incremental scalability and robustness against data corruption and disk and host failures. As the work progressed, we have been pleased to find that Lithium performance is on par with alternative approaches for local storage, and that its write-optimized disk layout shines for double and triple live replication. For small setups, the Fibre Channel array is still faster, but Lithium keeps scaling, does not have a single point of failure, and is able to tolerate both host and multiple drive failures.

5. RELATED WORK

Networked and distributed file systems have traditionally aimed at providing POSIX semantics over the network. Early examples such as NFS [39] and AFS [18] were centralized client-server designs, though HARP [22] used replication for availability and CODA [19] did support disconnected operation. FARsite [7] used replication and encryption of files and meta-data to tolerate Byzantine host failures, but as was the

case with CODA and AFS, provided only open-to-close data consistency, which is insufficient for VM hosting.

Lithium uses strong checksums and secret tokens to turn most Byzantine faults into fail-stop, but was not designed with Byzantine Fault Tolerance (BFT) as a goal. Fully BFT eventual consistency replication protocols have been proposed, and our storage engine design may be of relevance to developers of such systems. For instance, Zeno's [41] need for non-volatile RAM for version numbers could likely be repaired with a Lithium-like disk format. Other systems such as Carbonite [9] focus on finding the right balance between durability and performance and are orthogonal to our work. Quorum systems built with secret sharing have been proposed for access control and signatures by Naor and Wool [30], but to our knowledge not previously been explored in the context of storage.

The Google File System was one of the first storage systems to operate at cloud-scale, but its single coordinator node became a problematic bottleneck as the system grew in size and had to support a more diverse set of applications [27]. Lithium tracks all meta-data with its stored objects and does not need a central coordinator node. Ceph [48] is a POSIX-compatible and fault-tolerant clustered file system, where files are mapped to nodes through consistent hashing. Recent versions of Ceph use *btrfs* in Linux for transactional log-structured file storage. Lithium is similar in that both use log-structuring for transactional storage, but Lithium differs in that it uses a common fork-consistent data format on disk and in the network protocol.

Boxwood [23] provides a foundation for building distributed storage systems that can be implemented as B-trees. Boxwood also provides a 2-way replication facility for variable-sized chunks used for holding the actual data pointed to by the B-tree, with a log of dirty regions used for reconciling replicas after a crash. Lithium uses B-trees internally, but focuses on providing a virtual block interface with disk semantics instead of B-tree transactions.

Parallax [28] is a snapshotting virtual volume manager for shared storage. Parallax uses 4kB block-indexed radix trees as the layer of indirection between virtual and physical disk blocks, where Lithium uses extent-indexed B-trees that point to log entries. The extent B-tree is a more dynamic data structure, because extents can vary in size from a single disk sector and up to several megabytes per write. A radix tree, on the other hand, introduces a trade-off between either very small update sizes (which results in the radix tree becoming very large and possibly limits the maximum amount of disk space than can be addressed) or larger updates that can result in read-modify-write overheads and false sharing. Parallax relies on an underlying shared storage abstraction, where Lithium emulates shared storage on top of replicated shared-nothing storage volumes.

The HP Federated Array of Bricks (FAB) [38] is a decentralized storage system that uses replication and erasure-coding, combined with voting to keep disk blocks available in spite of failures. FAB stores a timestamp with each 8MB disk block, and uses majority voting on the timestamps to ensure that all nodes return the same value for a given block (linearizability). Lithium does not apply updates in place, so linearizability and uniformity follow from the use of single primary and the temporal ordering of updates on disk. Olive [1] is an extension to FAB that supports distributed branching of disk volumes, which is complicated because FAB applies updates in-place. Lithium's branching functionality was simpler to implement, because updates are non-destructive.

Sun's Zettabyte File System [8] (ZFS) is an advanced volume manager topped with a thin POSIX layer. ZFS uses Merkle hash trees to verify the integrity of disk blocks, and has good support for snapshots and for extracting deltas between them. Replication can be added by continuous snapshotting and shipping of deltas, but there is a small window where updates risk being lost if the primary host crashes. ZFS uses a fairly large minimum update size of 128kB, with smaller updates being emulated with a read-modify-write cycle. This can be expensive for VM workloads that frequently write less than 128kB, unless non-volatile RAM or other non-commodity hardware is used. In comparison, Lithium uses incremental hashes for integrity, supports single-block updates without read-modify-write, and, when configured for synchronous replication, has no window for data-loss.

Following Rosenblum and Ousterhout's seminal LFS work, there has been a large amount of research in the area, including on how to reduce the log cleaning overheads, *e.g.*, through hole-plugging and free-block scheduling. Matthews surveys much of the design space in [25]. Others have explored hybrid designs [46], where log-structuring is used only for hot data, with cold data being migrated to more conventionally stored sections of the disks. The authors of that paper note that while seek times only improved 2.7× between 1991 and 2002, disk bandwidth improved 10×, making the benefits of log structuring larger now than at the time of the original LFS work. We view most of this work as orthogonal to our own, and note that replication may be the "killer app" that could finally make log-structuring mainstream.

6. FUTURE WORK

Flash in the form of SSDs and dedicated PCI cards is gaining popularity for IOPS-critical server workloads. We have performed simple experiments with Lithium running VMs off an SSD and using a SATA drive for storing replicas from other hosts. Going forward, we plan to explore this combination of Flash and mechanical drives further, and to experiment with faster network links, *e.g.*, 10G Ethernet with remote direct memory access (RDMA). In the version of Lithium described here, the VM always runs on a host with a full replica of its volume. We are in the process of adding remote access to replicas, so that VMs can run anywhere without having to wait for a replica being created first. We are also considering adding support for de-duplication at the extent level, leveraging previous VMware work [11]. Lithium uses simple majority quorums to prevent replica divergence. We plan to explore techniques described by Naor and Wieder [29]. They show how dynamic quorum system can be constructed that do not require a majority to operate.

7. CONCLUSION

Cloud computing promises to drive down the cost of computing by replacing few highly reliable, but costly, compute hosts with many cheap, but less reliable, ones. More hosts afford more redundancy, making individual hosts disposable, and system maintenance consists mainly of lazily replacing hardware when it fails. Virtual machines allow legacy soft-

ware to run unmodified in the cloud, but storage is often a limiting scalability factor.

In this paper, we have described Lithium, a fork-consistent replication system for virtual disks. Fork-consistency has previously been proposed for storing data on untrusted or Byzantine hosts, and forms the basis of popular distributed revision control systems. Our work shows that fork-consistent storage is viable even for demanding virtualized workloads such as file systems and online transaction processing. We address important practical issues, such as how to safely allow multiple outstanding IOs and how to augment the fork-consistency model with a novel cryptographic locking primitive to handle volume migration and fail-over. Furthermore, our system is able to emulate shared-storage SCSI reservation semantics and is compatible with clustered databases and file systems that use on-disk locks to coordinate access. We have shown how Lithium achieves substantial robustness both to data corruption and protocol implementation errors, and potentially unbounded scalability without bottlenecks or single points of failure. Furthermore, measurements from our prototype implementation show that Lithium is able to compete with an expensive Fibre Channel storage array on a small cluster of eight hosts, and is faster than traditional disk layouts for replication workloads.

8. ACKNOWLEDGEMENTS

The authors would like to thank Irfan Ahmad, Jørgen S. Hansen, Orran Krieger, Eno Thereska, George Coulouris, Eske Christiansen, Rene Schmidt, Steffen Grarup, Henning Schild, and the anonymous reviewers for their input and comments that helped shape and improve this paper.

9. REFERENCES

[1] M. K. Aguilera, S. Spence, and A. Veitch. Olive: distributed point-in-time branching storage for real systems. In *NSDI'06: Proceedings of the 3rd conference on Networked Systems Design & Implementation*, pages 27–27, Berkeley, CA, USA, 2006. USENIX Association.

[2] Amazon. Amazon Elastic Block Store (Amazon EBS). http://aws.amazon.com/ebs, 2009.

[3] Amazon. Amazon Simple Storage Service (Amazon S3). http://aws.amazon.com/S3, 2009.

[4] L. N. Bairavasundaram, A. C. Arpaci-Dusseau, R. H. Arpaci-Dusseau, G. R. Goodson, and B. Schroeder. An analysis of data corruption in the storage stack. *Trans. Storage*, 4(3):1–28, 2008.

[5] G. Bartels and T. Mann. Cloudburst: A Compressing, Log-Structured Virtual Disk for Flash Memory. Technical Report 2001-001, Compaq Systems Research Center, February 2001.

[6] B. H. Bloom. Space/time trade-offs in hash coding with allowable errors. *Communications of the ACM*, 13:422–426, 1970.

[7] W. J. Bolosky, J. R. Douceur, and J. Howell. The farsite project: a retrospective. *SIGOPS Oper. Syst. Rev.*, 41(2):17–26, 2007.

[8] J. Bonwick, M. Ahrens, V. Henson, M. Maybee, and M. Shellenbaum. The Zettabyte File System. Technical report, Sun Microsystems, 2003.

[9] B.-G. Chun, F. Dabek, A. Haeberlen, E. Sit, H. Weatherspoon, M. F. Kaashoek, J. Kubiatowicz, and R. Morris. Efficient replica maintenance for distributed storage systems. In *NSDI'06: Proceedings of the 3rd conference on Networked Systems Design & Implementation*, pages 4–4, Berkeley, CA, USA, 2006. USENIX Association.

[10] C. Clark, K. Fraser, S. Hand, J. G. Hansen, E. Jul, C. Limpach, I. Pratt, and A. Warfield. Live Migration of Virtual Machines. In *Proceedings of the 2nd Networked Systems Design and Implementation NSDI '05*, May 2005.

[11] A. T. Clements, I. Ahmad, M. Vilayannur, and J. Li. Decentralized Deduplication in SAN Cluster File Systems. In *Proc. USENIX Annual Technical Conference*, San Diego, CA, 2009.

[12] G. Coulouris, J. Dollimore, and T. Kindberg. *Distributed Systems Concepts and Design.* International Computer Science Series. Addison-Wesley, 4 edition, 2005.

[13] F. Dabek, M. F. Kaashoek, D. Karger, R. Morris, and I. Stoica. Wide-area cooperative storage with CFS. In *SOSP '01: Proceedings of the eighteenth ACM symposium on Operating systems principles*, pages 202–215, New York, NY, USA, 2001. ACM.

[14] W. de Jonge, M. F. Kaashoek, and W. C. Hsieh. The logical disk: a new approach to improving file systems. In *SOSP '93: Proceedings of the fourteenth ACM symposium on Operating systems principles*, pages 15–28, New York, NY, USA, 1993. ACM Press.

[15] G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Vosshall, and W. Vogels. Dynamo: Amazon's highly available key-value store. In *SOSP '07: Proceedings of twenty-first ACM SIGOPS symposium on Operating systems principles*, pages 205–220, New York, NY, USA, 2007. ACM.

[16] S. Ghemawat, H. Gobioff, and S.-T. Leung. The google file system. *SIGOPS Oper. Syst. Rev.*, 37(5):29–43, 2003.

[17] A. Haeberlen, A. Mislove, and P. Druschel. Glacier: highly durable, decentralized storage despite massive correlated failures. In *NSDI'05: Proceedings of the 2nd conference on Symposium on Networked Systems Design & Implementation*, pages 143–158, Berkeley, CA, USA, 2005. USENIX Association.

[18] J. H. Howard, M. L. Kazar, S. G. Menees, D. A. Nichols, M. Satyanarayanan, R. N. Sidebotham, and M. J. West. Scale and performance in a distributed file system. *ACM Trans. Comput. Syst.*, 6(1):51–81, 1988.

[19] J. J. Kistler and M. Satyanarayanan. Disconnected operation in the Coda file system. *ACM Trans. Comput. Syst.*, 10(1):3–25, 1992.

[20] L. Lamport. The part-time parliament. *ACM Trans. Comput. Syst.*, 16(2):133–169, 1998.

[21] E. K. Lee and C. A. Thekkath. Petal: distributed virtual disks. *SIGOPS Oper. Syst. Rev.*, 30(5):84–92, 1996.

[22] B. Liskov, S. Ghemawat, R. Gruber, P. Johnson, and L. Shrira. Replication in the Harp file system. *SIGOPS Oper. Syst. Rev.*, 25(5):226–238, 1991.

[23] J. MacCormick, N. Murphy, M. Najork, C. A. Thekkath, and L. Zhou. Boxwood: abstractions as the foundation for storage infrastructure. In *OSDI'04:*

Proceedings of the 6th conference on Symposium on Opearting Systems Design & Implementation, pages 8–8, Berkeley, CA, USA, 2004. USENIX Association.

[24] M. Mackall. Mercurial revision control system. http://mercurial.selenic.com/.

[25] J. N. Matthews. *Improving file system performance with adaptive methods*. PhD thesis, 1999. Co-Chair-Anderson, Thomas E. and Co-Chair-Hellerstein, Joseph M.

[26] D. Mazières and D. Shasha. Building secure file systems out of byzantine storage. In *PODC '02: Proceedings of the twenty-first annual symposium on Principles of distributed computing*, pages 108–117, New York, NY, USA, 2002. ACM.

[27] M. K. McKusick and S. Quinlan. Gfs: Evolution on fast-forward. *Queue*, 7(7):10–20, 2009.

[28] D. T. Meyer, G. Aggarwal, B. Cully, G. Lefebvre, M. J. Feeley, N. C. Hutchinson, and A. Warfield. Parallax: virtual disks for virtual machines. *SIGOPS Oper. Syst. Rev.*, 42(4):41–54, 2008.

[29] M. Naor and U. Wieder. Scalable and dynamic quorum systems. In *PODC '03: Proceedings of the twenty-second annual symposium on Principles of distributed computing*, pages 114–122, New York, NY, USA, 2003. ACM.

[30] M. Naor and A. Wool. Access control and signatures via quorum secret sharing. *Parallel and Distributed Systems, IEEE Transactions on*, 9(9):909–922, Sep 1998.

[31] D. Narayanan, E. Thereska, A. Donnelly, S. Elnikety, and A. Rowstron. Migrating server storage to SSDs: analysis of tradeoffs. In *EuroSys '09: Proceedings of the 4th ACM European conference on Computer systems*, pages 145–158, New York, NY, USA, 2009. ACM.

[32] M. Nelson, B.-H. Lim, and G. Hutchins. Fast transparent migration for virtual machines. In *Proceedings of the 2005 Annual USENIX Technical Conference*, April 2005.

[33] B. M. Oki and B. H. Liskov. Viewstamped replication: A new primary copy method to support highly-available distributed systems. In *PODC '88: Proceedings of the seventh annual ACM Symposium on Principles of distributed computing*, pages 8–17, New York, NY, USA, 1988. ACM.

[34] A. Phanishayee, E. Krevat, V. Vasudevan, D. G. Andersen, G. R. Ganger, G. A. Gibson, and S. Seshan. Measurement and analysis of tcp throughput collapse in cluster-based storage systems. In *FAST'08: Proceedings of the 6th USENIX Conference on File and Storage Technologies*, pages 1–14, Berkeley, CA, USA, 2008. USENIX Association.

[35] S. Rhea, P. Eaton, D. Geels, H. Weatherspoon, B. Zhao, and J. Kubiatowicz. Pond: The oceanstore prototype. In *FAST '03: Proceedings of the 2nd USENIX Conference on File and Storage Technologies*, pages 1–14, Berkeley, CA, USA, 2003. USENIX Association.

[36] M. Rosenblum and J. K. Ousterhout. The design and implementation of a log-structured file system. *ACM Transactions on Computer Systems*, 10(1):26–52, 1992.

[37] A. Rowstron and P. Druschel. Storage management and caching in PAST, a large-scale, persistent peer-to-peer storage utility. *SIGOPS Oper. Syst. Rev.*, 35(5):188–201, 2001.

[38] Y. Saito, S. Frolund, A. Veitch, A. Merchant, and S. Spence. FAB: building distributed enterprise disk arrays from commodity components. In *ASPLOS-XI: Proceedings of the 11th international conference on Architectural support for programming languages and operating systems*, pages 48–58, New York, NY, USA, 2004. ACM.

[39] R. Sandberg, D. Goldberg, S. Kleiman, D. Walsh, and B. Lyon. Design and Implementation of the Sun Network Filesystem. In *Proceedings of the Summer 1985 USENIX Conference*, pages 119–130, 1985.

[40] F. B. Schneider. Implementing fault-tolerant services using the state machine approach: a tutorial. *ACM Comput. Surv.*, 22(4):299–319, 1990.

[41] A. Singh, P. Fonseca, P. Kuznetsov, R. Rodrigues, and P. Maniatis. Zeno: eventually consistent byzantine-fault tolerance. In *NSDI'09: Proceedings of the 6th USENIX symposium on Networked systems design and implementation*, pages 169–184, Berkeley, CA, USA, 2009. USENIX Association.

[42] J. Terrace and M. J. Freedman. Object storage on CRAQ: High-throughput chain replication for read-mostly workloads. In *Proc. USENIX Annual Technical Conference*, San Diego, CA, 2009.

[43] D. B. Terry, M. M. Theimer, K. Petersen, A. J. Demers, M. J. Spreitzer, and C. H. Hauser. Managing update conflicts in Bayou, a weakly connected replicated storage system. In *SOSP '95: Proceedings of the fifteenth ACM symposium on Operating systems principles*, pages 172–182, New York, NY, USA, 1995. ACM.

[44] VMware. VMware High Availability. http://www.vmware.com/pdf/ha_datasheet.pdf, 2009.

[45] VMware. VMware VMFS. http://www.vmware.com/pdf/vmfs_datasheet.pdf, 2009.

[46] W. Wang, Y. Zhao, and R. Bunt. Hylog: A high performance approach to managing disk layout. In *FAST '04: Proceedings of the 3rd USENIX Conference on File and Storage Technologies*, pages 145–158, Berkeley, CA, USA, 2004. USENIX Association.

[47] H. Weatherspoon, P. Eaton, B.-G. Chun, and J. Kubiatowicz. Antiquity: exploiting a secure log for wide-area distributed storage. In *EuroSys '07: Proceedings of the 2007 EuroSys conference*, pages 371–384, New York, NY, USA, 2007. ACM Press.

[48] S. A. Weil, S. A. Brandt, E. L. Miller, D. D. E. Long, and C. Maltzahn. Ceph: a scalable, high-performance distributed file system. In *OSDI '06: Proceedings of the 7th symposium on Operating systems design and implementation*, pages 307–320, Berkeley, CA, USA, 2006. USENIX Association.