

Capo: Recapitulating Storage for Virtual Desktops

Abstract

We present Capo,¹ a transparent, persistent block request proxy for virtual machine disk images. Capo attempts to reduce the aggregate consumption of shared storage bandwidth in situations where shared storage is used to host virtual machine images.

1 Introduction

The storage we trust is expensive. Fast and reliable data storage is something that organizations are prepared to pay a premium for, both in the capital costs of enterprise storage hardware and the operational costs of ensuring that important data is written to it.

Interestingly, the deployment of virtualization has inverted the historical imperative that systems be configured to “opt-in” to storing data on appropriate network shares instead of on less reliable locations such as local disks. While administrators used to have to work to configure applications to use enterprise storage, virtual environments simply store *everything* on it.

As such, these environments present the opposite problem: The requirement that virtual machine images be universally accessible, with high performance, to all physical hosts in a cluster has necessitated the deployment of SAN hardware in even modest virtualization deployments. The improved density and utilization afforded by virtualization allows systems to scale to large numbers of VMs; shared storage must scale proportionately to provide for them. This symptom is especially problematic for virtual desktops, where infrastructure is being deployed to host literally thousands of nearly identical VM images. A number of commercial virtual desk-

top systems now exist, and deployments suffer from a significant, if not dominant, cost for enterprise storage.

This paper argues that shared, central, storage *is* the correct approach for scalable virtual environments. It is trustworthy, relatively easy to manage, and simple to reason about. However, we believe that for applications such as virtual desktops, which involve large numbers of image clones, the majority of request load is redundant and can be effectively serviced by local, commodity disks within individual servers. Furthermore, we believe that the levels of durability provided by enterprise storage in these environments are in excess of what is necessary for large portions of desktop OS disk images.

The contributions of this paper are twofold: First, we validate our hypothesis through the analysis of a week-long trace of all storage traffic from a production deployment of 55 Windows Vista desktops in an executive and administrative office of a large public organization. Our results examine the opportunities that exist for caching data both within and across virtual desktop images. They also examine the breakdown of request workload within desktop filesystems.

Second, based on the analysis of this trace, we describe the design and implementation of **Capo**, a distributed persistent cache that aims to reduce aggregate load on shared storage in virtual desktop environments. Capo uses local server disks to provide persistent caching of VM images, and includes mechanisms to share and pre-load caches of gold master images across VMs and across hosts. Finally, Capo introduces a facility for *differential durability*, which allows administrators to selectively “opt out” of enterprise storage guarantees by relaxing the durability properties of subsets of a desktop’s file system.

Virtual desktop systems have already taken advantage of several approaches to scale storage to large numbers of desktop machines. We begin in Section 2 by providing some brief initial background on these systems.

¹The name of our system is borrowed from the phrase “*Da Capo al coda*”, which is used in sheet music to indicate a brief return to the beginning of a piece, followed by a jump to the Coda, or conclusion. In sonatas, this “recapitulation” involves revisiting a similar, but sometimes different version of the main theme of the arrangement.

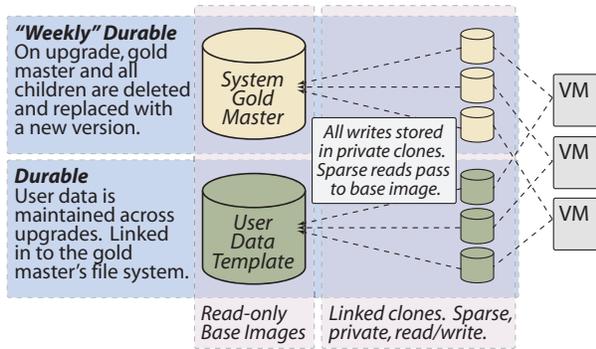


Figure 1: Typical image management in VDI systems.

2 VDI Background

Virtual desktops represent the latest round in a decades-long oscillation between thin- and thick-client computing models. So-called Virtual Desktop Infrastructure (VDI) systems have emerged as a means of serving desktop computers from central, virtualized hardware. VDIs are being touted as a new compromise in a history of largely unsuccessful attempts to migrate desktop users onto thin clients and the approach does provide a number of benefits. From the user perspective, giving users private virtual machines retains their ability to customize their environment and interact with the system as they would a normal desktop computer. From the administration perspective, consolidating VMs onto central compute resources has the potential to reduce power consumption, allow location-transparent access, better protect private data, and ease software upgrades and maintenance.

Commercial VDI systems appear to be experiencing a degree of success: Gartner predicts that forty percent of all worldwide desktops—49 Million in total—will be virtualized by 2013 [13]. Today, the two major vendors of VDI systems, Citrix and VMware, individually describe numerous case studies of active virtual desktop deployments of over 10,000 users. From a storage perspective, VDI systems have faced immediate challenges around space overheads and the ability to deploy and upgrade desktops over time. As background, this section describes how these problems are typically solved in existing architectures, as illustrated in Figure 1.

2.1 Copy-on-Write and Linked Clones

Operating system images are entire virtual disks, often tens of gigabytes each. A naive approach to supporting hundreds or thousands of virtual machines results in two immediate storage scalability problems. First, VMs must have isolated disk images, but maintaining individual copies of every single disk is impractical and consumes an enormous amount of space. Second, adding

new users requires that images can be quickly duplicated without necessitating a complete copy.

This observation is not new; it has been a recurring challenge in virtualization. Existing VDI systems make use of VM-specific file formats such as Microsoft’s VHD [10] and VMware’s VMDK [18]. Both allow a sparse overlay image to be “chained” to a read-only base image (or Gold Master). As shown in Figure 1, modifications are written to private, per-VM overlays, and any data not in the overlay is read from the base. In this manner, large numbers of images may share a single gold master. This approach consolidates common initial image data, and new images may be quickly cloned from a single gold master.

2.2 Image Updates and Periodic Rollback

Image chaining saves space and allows new images to be cloned from a gold master almost instantaneously. It is not a panacea though. Chained images immediately begin to diverge from the master version as VMs issue writes to them. One immediate problem with this divergence is the consumption of independent extra storage on a per-image basis. This divergence problem for storage consumption is typically addressed through the use of data deduplication [20, 3].

For VDI, wasted storage is not the most pressing concern: block-level chaining means that patches and upgrades cannot be applied to the base image in a manner that merges and reconciles with the diverged clones. This means the ability to deploy new software or upgrades to a large number of VMs, which *was* initially provided from the single gold master is immediately lost.

The leading VDI offerings all solve this problem in a very similar way: They disallow users from persisting long term changes to the system image. When gold master images are first created and clones are deployed, the VDI system arranges images to isolate private user data (documents, settings, etc.) on separate storage from the system disk itself. As suggested initially in the Collective project [1], this approach allows a new gold master with updated software to be prepared and deployed to VMs simply by replacing the gold master, creating new (empty) clones, and throwing away the old version of the system disk along with all changes. This approach effectively “freshens” the underlying system image of all users periodically and ensures that all users are using a similar well-configured desktop. For the most part, it also means that users are unable to install additional, long-lived software within VDI images without support from administrators.

3 Virtual Desktop Trace

To better understand VDI workloads, we arranged to measure all block and file level activity from the a production VDI deployment for a one week period during the Summer of 2010. The deployment being studied is an office in a large public organization containing executive and administrative support staff. The deployment had been in production use for six months and includes 55 Windows Vista desktops, the organization is in the process of rolling out another 300 desktops this fall.

3.1 Methodology

We installed a Windows storage class driver into the base system image of the virtual desktop machines. The driver was written to record read and write events to the virtual disks using the Microsoft Windows Software Trace Preprocessor (WPP). It recorded request size and virtual disk address. In 93% of cases we were also able to determine the file on which the access originated by following the `OriginalFileObject` pointer in the Windows I/O Request Packet (IRP) structure. To better contextualize this information, we also installed a driver at the filesystem level and recorded cache accesses, the application making each request, and the file flags for each file accessed. Our disk-level driver is written in 515 lines of C, while our file-level driver is 82 lines of C.

Logs from these drivers were written to a network share and collected on the Thursday following a full week of logging. In total we collected 75GB of logs in a compressed binary format. We then checked for corruption, missing logs, or missing events. Out of over 300 million entries we found a single anomalous write to a clearly invalid block address, which we removed. We could find no explanation for the event. In the rest of this section we present our analysis of this data. Unless otherwise specified, we will refer to block level accesses to a virtual disk and measure aggregate workloads in I/O operations per second (IOPS).

3.2 Our Virtual Desktop Environment

The environment we are studying is structured very much like the one described in Section 2. At the time our measurements were gathered it hosted 55 Microsoft Windows Vista virtual desktops with VMWare View, of which roughly 27 are in dedicated day-to-day use as the primary desktop. This small size is the primary limitation of our study, but we expect to measure considerably more as the installation grows. Furthermore, even at the current size it is possible to see considerable self-similarity among machines, as we will discuss.

End users work from Dell FX100 Zero thin clients, while VMs are served from HP BL490c G6 Blades running ESX Server. These servers connect to a Network Appliance 3170s over fiber channel, for booting from the SAN, and 10GigE, for VM disk images. System images are hosted via NFS on a 14 drive RAID group with 2 parity disks. They are optimized for performance [16] and are pre-loaded with Firefox, Microsoft Office Enterprise, and Sophos Anti-Virus among other software. At the end of every Wednesday, a new system image is published to all users exactly as discussed in Section 2.2.

3.3 Analysis

We begin by asking: *What are the day-to-day characteristics of VDI storage workloads?* Figure 2 shows the entire study in I/O operations-per-second for the 24-hour period of each of the 7 full days recorded. There is a distinct peak load period between 8:30 and 9:30 every morning, as employees arrive at work. Three peaks in this period are highlighted in the figure and presented for expanded analysis in Table 1. The right-most column shows the applications responsible for the most paging faults and reads, excluding the system and services. Most days, Firefox and the virus scanner are very active in this period, we also see Thunderbird, Pidgeon, and Microsoft Outlook frequently. We were surprised to see the Search Indexer active as well, because we were told its background scanning task had been disabled to reduce I/O consumption. Our best guess is that it has invoked manually.

We measured the write to read percentages for both I/Os and throughput, which is useful in characterizing the workload. Our workload is write-heavy in I/Os and read-heavy in throughput, both by approximately two-to-one. We then measured the percentage of VMs which contributed at least 5% of the peak workload, to determine if peaks were caused by multiple VMs or by a few outliers. In most cases it is the former, however, the peak in time period 4 was caused primarily by 3 VMs.

We included the column titled “Dup. reads” to explore the potential for caching. We present two numbers. The left most is the percentage of reads that have been previously seen by that VM over the entire trace. Thus with a very large cache, we could potentially absorb *all* these reads. The number on the right presents the same measure, but imagines that caching could be shared across all VMs in the cluster. Time period 4 stands out for having an unusually low individual duplicate read rate, but very high across the cluster. We investigated and found that two very active VMs had duplicate read rates of of 26% and 35%. By including only the least beneficial 38%, 15% and 4% of VMs, we could reach duplicate read rates of 40%, 60% and 90% respectively for those VMs.

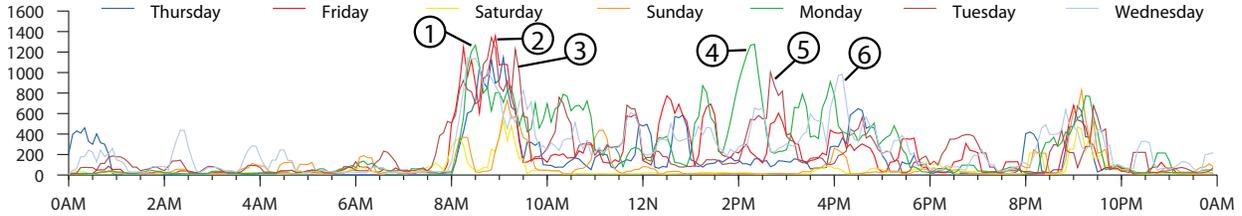


Figure 2: Activity measured in IO operations over each of the 7 days.

	Time Period	Write % (IOps / Thpt)	% of VMs (≥ 5%)	Dup. reads (VM / Clust.)	Top Applications by Iops (excludes system and svchost)
1	Mon. 8:30am-8:45am	50% / 22%	26%	81% / 91%	Search Indexer, Firefox, Sophos
2	Fri. 8:30am-9:00am	52% / 22%	29%	88% / 97%	Firefox, Search Indexer, Sophos
3	Tue. 9:15am-9:30am	64% / 43%	29%	78% / 99%	Defrag, Firefox, Search Indexer
4	Mon. 2:00pm-2:30pm	62% / 41%	7%	59% / 99%	Firefox, Pidgin, Sophos
5	Tue. 2:40pm-3:00pm	69% / 52%	26%	77% / 97%	Firefox, Defrag, Pidgin
6	Wed. 4:00pm-4:15pm	60% / 37%	26%	99% / >99%	Firefox, Pidgin, Sophos

Table 1: Points of interest in Figure 2.

This suggests that one can achieve significant improvements with caching, and possibly even benefit by sharing caches among VMs. However some sharing benefits may require careful selection of the VMs in question.

Lunch, dinner, and late nights are periods of relative inactivity, as are the weekends. Late afternoon peaks are sporadic, but reach loads nearly as high as the morning. One such peak, marked 4 in Figure 2 and Table 1 was caused by a relatively small number of machines engaged in heavy browsing activity. This is not the norm, as all other peaks occur when more than a quarter of the VMs are significantly active. This is clear in Figure 3, which shows a CDF of VMs by their contribution to the total workload for each peak. These peak load periods are particularly important, because they define the hardware necessary to service the workload without disruption. We conclude that *VDI workloads are defined by their peaks, and those peaks usually occur at times of common activity among many VMs*. In Sections 4.1, and 4.2 we take advantage of this fact to improve performance in VDI environments.

Next, we ask *How can we characterize the I/O requests?* In Figure 4 we show the requests for the entire workload by both count and size. We differentiate reads from writes, and also isolate each request by its target in the file system namespace. The workload is 65% writes, which account for 35% of the throughput, versus 35% of reads accounting for 65% of the throughput. Metadata operations account for large portion of the operations, and unfortunately we cannot determine how these modifications relate to the namespace. Directories typically managed by the operating system, such as \Windows directory and \Program files are also

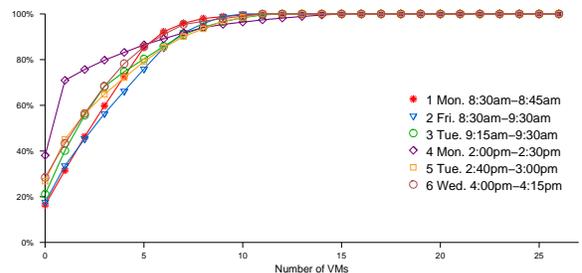


Figure 3: Contributors to each of the major workload peaks.

frequently accessed. There are fewer accesses to user directories and temporary files, most of the latter are to \Temporary Internet Files folder, as opposed to \Windows\Temp. These findings contrast those of Vogels who’s study showed that 93% of *file-level* modification occurred in \User directories [19]. We conclude that *while a wide range of the namespace is accessed, it is not accessed uniformly, and access to data directly managed by users is rare*. We will revisit this observation in Section 4.3.4.

Since our workload is write-heavy, we next ask *how are these writes organized in time?* Figure 5 shows the percentage of bytes that are rewritten after time intervals ranging from 10 minutes to a whole day. We have included results from each of the seven days to underscore how consistent the results are. After a short time, just 10 minutes, between 20%-30% of bytes that are written are written again. This rate increases and after twenty-four hours represents a reduction of between 80%-90%. From

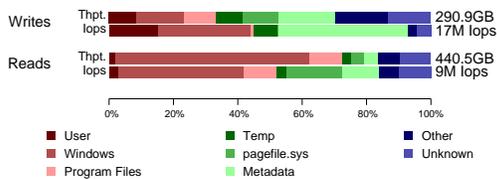


Figure 4: Size and amount of block level writes by file system path.

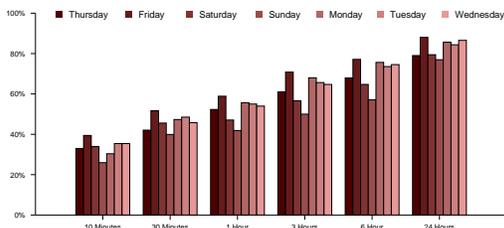


Figure 5: Percentage of bytes rewritten in various time intervals for each day of the study.

this we conclude that *Considerable system-wide effort is spent on data with a high modification rate.* We show how this can be used to our advantage in Section 4.3.

Since VMs typically use disk images chained from a gold master, we are interested in the rate at which they diverge from the original image. We therefore ask *At what frequency do we observe the first write to a sector?* Figure 6 plots this data for the average VM, as well as the most and least divergent VM, over the entire study. Within 24 hours, most VMs hit a near plateau in their divergence, around 1GB. Over time this does increase, but slowly. A smaller set of VMs do diverge more quickly and significantly, but there are far from the 95% confidence interval. We conclude that *there is significant shared data between VMs, even after several days of divergence.*

Naturally, we do not expect divergent writes to occur uniformly, so we pose the question: *Where in the namespace do divergent writes occur, and does this change over time?* Figure 7 plots the cumulative divergence for each VMs in the cluster, and divides that total among various components of the namespace. One can observe that the pagefile diverges immediately, then remains a constant size over time, as does the system metadata. Both these files are bounded in size. Meanwhile writes to windows and areas of the disk we cannot associate with any file continue to grow over the full week of the study. We conclude that *While writes occur everywhere in the namespace, they exhibit significant trends when categorized according to the destination.*

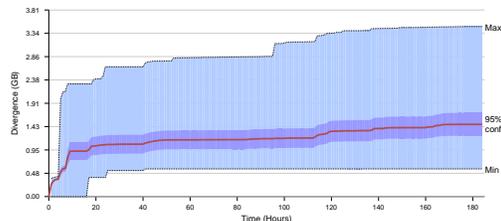


Figure 6: Bytes of disk diverging from the gold master.

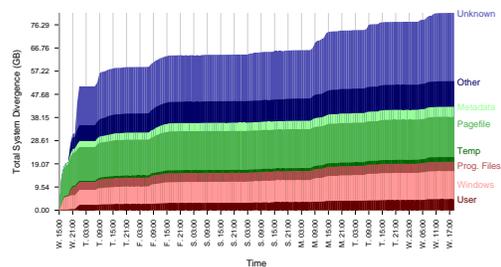


Figure 7: Total divergence versus time for each namespace category.

3.4 Summary

While there is more to say about this workload and those of VDI environments in general, the observations in this section are valuable. Summarizing our observations from the trace data:

- VDI workloads are defined by their peaks, and those peaks usually occur at times of common activity among many VMs
- While a wide range of the namespace is accessed, it is not accessed uniformly, and access to data directly managed by users is rare
- Considerable system-wide effort is spent on data with a high modification rate
- There is significant shared data between VMs, even after several days of divergence
- While writes occur everywhere in the namespace, they exhibit significant trends when categorized according to the destination

These observations taken together suggest that additional caching might resolve the performance challenges that we have observed. The following section builds on the observations and analysis presented here, and describes the architecture of Capo.

4 Architecture

The trace analysis in Section 3 suggests that caching below the individual VMs may be effective in resolving the demand peaks that we observed. In this section we describe Capo, including its three major components:

1. A single host cache shared between all the virtual desktops hosted on a single server machine to eliminate redundant block reads and writes.
2. A multi-host cache preloading component to eliminate redundant read accesses from virtual desktops hosted on different server machines.
3. A component that supports differential durability for writes based on the location in the namespace of the affected file.

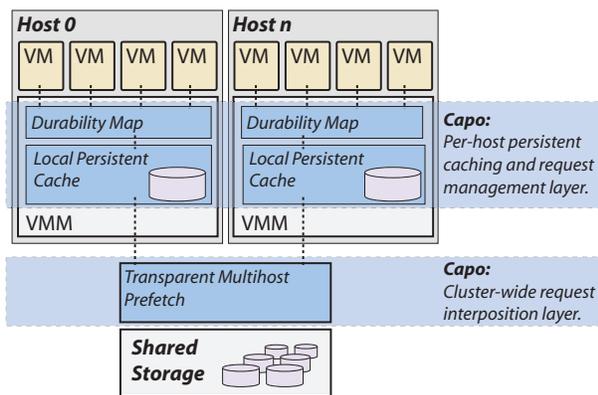


Figure 8: The Major Architectural Components of Capo.

Figure 8 shows the overall architecture of Capo. Capo exists as a layer within the virtual machine monitor (VMM) which supports the individual desktop VMs. The figure depicts each host including a Local Persistent Cache which is stored on the local disk of the host machine and is described in Section 4.1. Spanning all of the hosts is the Transparent Multihost Prefetch component which preloads data accessed by one host into the local caches of the other hosts in advance of it being requested by that host. It is described in Section 4.2. The Durability Map component supports the wide variation in the durability requirements of the various components of the file system. It is described in Section 4.3.

4.1 Local Persistent Cache

All VDI deployments rely on central enterprise storage that provides high availability, durability, and reliability. The servers that host virtual desktops are also configured with local disk storage which consists of cheap COTS

disk drives with comparably lower reliability but higher aggregate I/O bandwidth.

The trace-based analysis of our local VDI deployment suggests that a cache shared between multiple virtual desktops might be very effective. As shown in Figure 3 there is significant overlap between the top applications executed on different virtual machines. Table 1 refines this and indicates that aggressively caching can yield very high read hit rates. Also, as shown in Figure 5, a significant fraction of data is overwritten very quickly. Therefore, as depicted in Figure 8, each server machine that hosts virtual desktops uses its local disk as a persistent cache. A key goal for Capo is to provide an appropriate level of durability for all data while taking advantage of the higher aggregate bandwidth available to local disk. The level of durability achieved depends on the caching policy in place.

4.1.1 Caching Policies

The cache supports two consistency policies: write-through and write-back. These policies are enforced on a cached disk image as a whole. Each of these policies represents a different tradeoff between virtual disk consistency and overall system performance.

The *write-through* policy provides the highest level of consistency guarantees a machine would expect from a block device. In this policy the cache replicates writes to both the centralized storage and the local cache. Write requests are not acknowledged until they hit both disks. This policy relieves the centralized storage from serving reads to blocks that have been previously read from or written to. The drawback of this policy is that write requests must be sent across the network, consuming network bandwidth and increasing both the load on the centralized storage and the client's perceived write request latency.

The *write-back* policy delays pushing updates to disk blocks by caching writes locally in the write cache. Updates are pushed to the central storage in a crash consistent manner at a per-virtual-disk configurable frequency. The choice of write-back frequency trades off system performance and durability of disk contents in case of a failure. A high update frequency minimizes the amount of data loss in case of the failure of the local disk, while a lower frequency enhances overall system performance by coalescing writes in the local cache.

4.1.2 Design and Implementation

The Local Persistent Cache consists of three components: a write cache, a read cache, and a snapshot component. Write requests are satisfied by writing their data into the write cache and concurrently to the centralized

storage if the cache is set to write-through mode. Read requests are served from the write cache, the read cache, or the central storage in that order. Read cache misses cause the data fetched from the central storage server to be inserted into the read cache from whence it can satisfy future references.

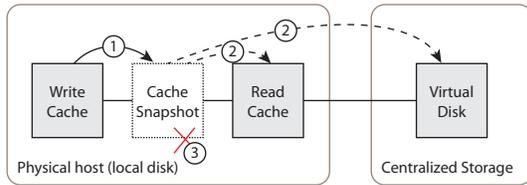


Figure 9: The Local Cache.

The snapshot mechanism in the cache works in tandem with a transactional update mechanism in the backend storage to ensure crash consistent updates to remote disk images when operating in write-back mode. A snapshot involves pausing the request stream momentarily to move the contents of the write cache to a snapshot and empty the write cache. The updates in the snapshot are then applied transactionally to the centralized storage and to the read cache. After the snapshot updates have been applied, the snapshot is cleared.

The cache is implemented as a user-level shared library that interposes on I/O calls, specifically the `glibc` and `libaio` I/O and file management operations. Due to the relative sizes of the disks in virtual desktops (around 10GB) and the disks in the physical machine which hosts them (greater than 1TB) and the amount of sharing between virtual desktops (see Figure 6), we can easily support hundreds of virtual desktops on a server without worrying about overfilling the cache. In our implementation, when the cache does fill, we simply throw it away and start again.

4.2 Multi-host Cache Preload

Capo’s local persistent cache goes a long way towards eliminating redundant read requests on individual machines. But as growing VM deployments lead to larger numbers of physical hosts, redundant reads across these hosts place additional burden on central servers. Further scalability improvements can be attained in this case by transmitting common data to all hosts when it is accessed from the centralized servers by any one host, rather than to each host individually.

To this end, we have developed a multicast cache preloader for the local caches. Unlike more complicated cluster file systems, the preloader is completely lock-free. It requires no modifications to existing centralized

servers and introduces negligible performance overhead. The preloader consists of a client on each host and a central server which watches traffic to the centralized storage server and determines what to preload into the local persistent caches. In this way, the first host to read common data essentially prefetches it for all other hosts.

4.2.1 Design and Implementation

Our initial design for the preloader server was to use a mirror port on the central storage server to monitor network traffic. In a method similar to that described in *Ditto* [2], the server captured raw network packets and reconstructed TCP flows to extract relevant data (in our case, NFS requests and responses).

When deploying this solution, however, we observed significant packet loss between the mirror port on the filer and our server, and since a single packet loss is enough to corrupt an entire NFS request or response we were missing many opportunities to preload data into the local persistent caches of the hosts.

Our second, and current, design for the preloader server is to use a user-level NFS proxy that sits between the clients and the filer. NFS requests and responses all go through the proxy and the proxy identifies data that should be preloaded into other local persistent caches. This increases the latency of filer requests somewhat, but avoids all of the issues with packet capture.

The preloader server exposes an RPC interface by which clients can register watches on specific files. The `watch_file` RPC accepts a unique identifier of a file hosted by the NFS server and returns a multicast address. The server spawns a separate thread for each file it watches. Server threads register callbacks with the NFS monitoring library and execute simple loops writing captured data to multicast sockets. The `unwatch_file` RPC decrements the reference count of a given watch and stops the multicast service if no clients remain.

A single client process runs on each physical host. Upon startup, these processes register watches with the server for all files being cached by Capo. Clients spawn a separate thread for every file they are watching; these threads enter a simple loop receiving data from a multicast socket and writing it to the corresponding Capo cache file. Because of Capo’s copy-on-write architecture, no synchronization is required when multicast clients fill read-only cache files. In the worst-case scenario, a write from a multicast client could potentially clobber a metadata write initiated by a Capo client. This would require the corresponding data blocks to be re-read from the server on next access, but it would not affect correctness.

The multicast protocol is simplified by a few basic design principles. We assume that, given a number of VMs

derived from a common master image, reads of the base image made by any individual VM will likely be duplicated by all VMs. That is, while the disks belonging to derived VMs will tend to diverge as the VMs age, the common portion of these disks will likely be read by all or none of the VMs. Thus if the multicast server observes any read of a common file, it is worthwhile to send this data to all hosts on which Capo is caching this file. By the same assumption, multicast clients do not proactively request data from the server, as they are not in a position to know which portions of files will be read by VMs.

The preload server sends a significant amount of traffic over a number of multicast sessions, and this exposes a problem with the support for multicast in modern switches. When too many multicast packets are in flight at once, the switch experiences congestion collapse, dropping both multicast and TCP packets. We have implemented a rudimentary flow-control scheme that limits the amount of multicast data in flight at a time, which has alleviated the problem but additional control knobs may be warranted. While not yet implemented, a few basic control messages (conceptually similar to those described in *Frisbee* [6]) could be used to throttle traffic during busy periods. For example, if clients were to send periodic messages reporting the number of requests received per file, the server could estimate packet drop rates for multicast requests and shape future traffic accordingly.

4.3 Differential Durability

Major VDI providers have all adopted the software update strategy proposed in The Collective [1], where user directories are isolated from the rest of the file system. Modifications made to files in the user directories must be durable; users depend on these changes. Capo therefore uses write-through caching on these directories propagating all changed blocks immediately to the centralized storage servers. Any modifications to the system image can then be performed on all VMs in one step by completely replacing the system images in the entire pool, leaving the user's data unmodified. This impacts durability—any writes to the system portion (e.g., by updating the registry or installing software) will be lost. In this section we use and extend this notion to optimize for our write-heavy workload.

4.3.1 Write-Back Period

As mentioned in Section 2.2, VDI deployments manage system data centrally, regularly replacing the system data seen by each virtual desktop with a clean updated version. While users are allowed to make changes to

their system data, these changes are not guaranteed to be durable. Writes to the `\Program Files` directory as part of an application install process, for example, represent work done by a user, but a software installation could easily be repeated if a failure caused this to be necessary. It might be acceptable if the loss of such effort was limited to, for example, an hour or even a day. We can set our write-back period for such partially durable files to a corresponding length of time.

4.3.2 Extending Partial Durability to User files

While much of the data on the User volume is important to the user and must have maximum durability, Windows, in particular, puts some data that is really system data in the User volume. Examples include log files, the user portion of the Windows registry, and the local and roaming profiles containing per-application configuration settings. Table 2 indicates some paths on User volumes in Windows 7 that can reasonably be cached with a write-back policy and a relatively long write-back period.

4.3.3 Eliminating Write-Back

There are some system files that need not be durably stored at all. These include files that are discarded on system restarts or can easily be reconstructed if lost. Writes to the pagefile, for example, represent nearly a tenth of the total throughput to centralized storage in our workload. These writes consume valuable storage and network bandwidth, but since the pagefile is discarded on system restart, durably storing this data provides no benefit. The additional durability obtained by transmitting these writes over a congested network to store them on a highly redundant centralized storage provides no value because this data fate-shares with the local host machine and its disk. Many temporary files are used in the same way, requiring persistent storage only as long as the VM is running.

We store this data to local disk only, assigning it a write-back cache policy with an infinitely long write-back period. In the event of a hardware crash on a physical host, the VM will be forced to reboot, and the data can be discarded.

4.3.4 Design

Initially, we approached the problem of mapping these policies to write requests as one of request *tagging*, in which a driver installed on each virtual desktop would provide hints to the local cache about each write. While this approach is flexible and powerful, maintaining the correct consistency between file and filesystem metadata (much of which appears as opaque writes to the Master

Path	Policy
\Program Files\	write-back
\WINDOWS\	write-back
\Users\ProgramData\VMware\VDM\logs	write-back
\Users\%USER%\ntuser.dat	write-back
\Users\%USER%\AppData\local	write-back
\Users\%USER%\AppData\roaming	write-back
\pagefile.sys	no-write-back
\ProgramData\Sophos	no-write-back
\Temp\	no-write-back
\Users\%USER%\AppData\Local\Microsoft\Windows\Temporary Internet Files	no-write-back
Everything else, including user data and FileSystem metadata	write-through

Table 2: Sample cache-coherency policies applied as part of durability optimization.

File Table in NTFS) under different policies is challenging. Instead, we have developed a simpler and better performing approach using existing filesystem features.

The path based policies we use in our experiments can be seen in Table 2; naturally, these may be customized by an administrator. We provide these policies to a disk optimization tool that we run when creating a virtual machine image. The optimization tool also takes a populated and configured base disk image. For each of the two less-durable policies, it takes the given path and moves the existing data to one of two newly-created NTFS file systems dedicated to that policy. It then replaces the path in the original file system with a reparse point (Windows’s analogue of a symbolic link) to the migrated data. This transforms the single file system into three file systems with the same original logical view. Each of the three file systems are placed on a volume with the appropriate policy provided by the local disk cache. This technique is similar to the view synthesis in Ventana [14], though we are the first to apply the technique with a local cache to optimize performance.

We appreciate that applying different consistency policies to files in a single logical file system may be controversial. The risk in doing so is that a crash or hardware failure results in a dependency between a file that is preserved and a file that is lost. Such a state could lead to instability; however, we are aware of no dependencies crossing from files with high durability requirements to those with lower durability requirements in practice. Further, we observe that this threat already exists in the production environment we studied, which overwrites system images with a common shared image on a weekly basis.

5 Evaluation

To evaluate the effectiveness of Capo, we first consider how effective differential durability is at reducing write

load from unimportant regions of disk. Next, we show the storage reduction achieved by Capo with eleven concurrent users synthesizing active desktop workloads ².

5.1 Microbenchmarks

5.1.1 Differential Durability

This section describes several microbenchmarks to evaluate the effectiveness of differential durability in isolation of other features and provide a clearer mapping of end-user activity to observed writes. We applied the policies in Table 2 to several realistic desktop workloads. For each, we measured the portion of write requests that would fall under each policy category.

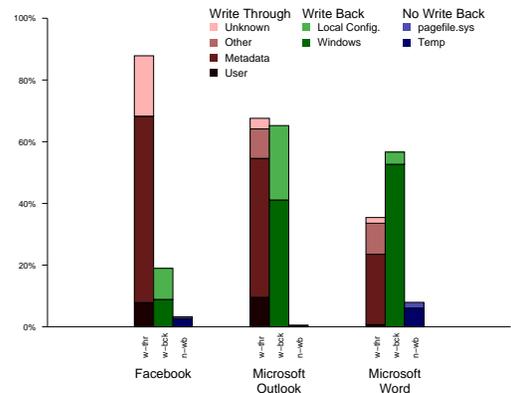


Figure 10: The effects of cache policies on three different workloads.

²Note to reviewers: The evaluation of Capo also includes trace replay of the six regions identified in Section 3, through the implementation. Unfortunately, we identified an error in the replay implementation shortly before submission that made us uncomfortable including the results. We have removed that subsection, but anticipate including a replay-based evaluation of our system in a final version of the paper.

5.1.2 Web Workload

Our web workload is intended to capture a short burst of web activity. The user opens `www.facebook.com` with Microsoft Internet Explorer, logs in, and posts a short message to their account. They then log off and close the browser. The entire task lasts less than a minute. The workload consists of 8MB (43.6% by count) of writes and 25.3MB (56.4% by count) of reads.

A breakdown of writes by their associated policy for each workload is shown in Figure 10. In this small workload only a small, but non-trivial improvement can be made. Local configuration changes such as registry, temp file, and cache updates can be removed, removing or delaying just over 20% of the operations.

5.1.3 Email Workload

Our email workload is based on Microsoft Outlook. The user sends emails to a server we have configured to automatically replay with the email back. Ten emails are sent and received in succession before the test ends. The workload consists of 63MB (39% by count) of writes and 148MB (61% by count) of reads.

Here the improvement is much more substantial. Although very few writes can be kept locally, nearly half can be delayed in writing to centralized storage. This is due to Outlook’s caching behavior, which makes heavy use of the system and application data folders. Emails in the .pst file are included in the user category. It is worth noting that many files in the windows and application data are obvious temp files, but did not match our current policies. With more careful tuning the workload could be further optimized.

5.1.4 Application Workload

Our application workload is intended to simulate a simple editing task. We open Microsoft Word and create a new document. We also open `www.wikipedia.org` in Microsoft Internet Explorer. We then proceed to navigate to 10 random Wikipedia pages in turn, and copy the first paragraph of each into our word document, saving the document each time. Finally, we close both programs. The workload consists of 120MB (20.0% by count) of writes and 406MB (80.0% by count) of reads.

Viewing many small pages creates a large number of small writes to temporary files and memory pressure³ increases the pagefile usage. Both programs write significantly to system folders, leaving less than 36% of the workload to be issued as write-through.

³The guest was running Windows 7 with 1GB of RAM, 25% higher than the XenDesktop recommended minimum.

5.2 Synthetic Workload

To evaluate Capo as a whole, we arranged to simulate a set of (very) active desktop users, performing similar workloads to those seen in the trace. Figure 11 shows the results of request traffic hitting both the local caches (in aggregate across all images) and the filer, while 11 users actively use a variety of office and web-based applications.

First of all, note that the load in this case is higher than any of the peaks seen in the trace data. This workload represents a higher level of aggregate storage activity than was ever seen in the production environment. Second, observe that despite being configured conservatively for complete write back, Capo reduces all peaks in the storage request load.

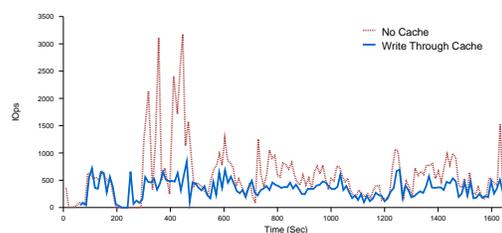


Figure 11: IOPS per second for a workload of 11 Windows users on a XenCenter Cluster.

6 Related Work

The caching component of Capo is most closely related to the ITC [15], Andrew [7], and Coda [9] file systems which utilize the local disk as a cache for whole files retrieved from servers. The Cedar file system [17] allows users to share immutable files over the network; by only supporting immutable files Cedar eliminates the need for cache consistency management.

Unlike these distributed file system caches, Capo operates at the block level. Capo’s cache consistency management is simplified by the fact that each virtual disk has a single writer and copy-on-write is used to prevent updating shared data. In Capo, as in Cedar, consistency management is simplified because shared data is always immutable.

Fs-cache [8] and iCache [5] are systems that, like Capo, implement block-level caching for remote storage systems: a file system in the case of Fs-cache and an iSCSI target in the case of iCache. Capo extends the basic block caches of these systems using a host cache shared by all the VMs on a host, the multi-host prefetcher, and differential durability for files. All of these features are inspired by our target environment of supporting virtual desktops.

Capo's use of write-back caching reduces the demand placed on the central storage facility in a manner similar to that of Everest [11]. Where Everest replicates offloaded write requests to tolerate disk failures, Capo uses a technique similar to Snapmirror [12] to periodically push self-consistent updates across the network for data that is cached in write-back mode.

Other researchers have studied the performance of storage in virtualized environments. In particular, Gulati et al. [4] study the storage demands of enterprise applications in virtualized environments. In contrast, our study of virtual desktops provides insight into the unique characteristics of this emerging use of virtualization.

7 Conclusion

Enterprise storage provides considerable benefit to virtual environments. However, for applications such as virtual desktops, that involve large numbers of nearly identical images running concurrently, a large portion of the request load placed on shared storage is unnecessary. After analyzing a one-week trace of a production VDI deployment, we presented Capo, a distributed and persistent cache to reduce the aggregate load placed on shared storage. Capo uses local disks on individual physical servers to cache image contents for the VMs being hosted, it includes mechanisms to share common cached base images across VMs, and to prefetch caches across physical hosts. In addition, Capo supports a configurable degree of differential durability, allowing administrators to relax the durability properties and the associated write load, for less-important subsets of a VM's file system.

References

- [1] R. Chandra, N. Zeldovich, C. Sapuntzakis, and M. S. Lam. The collective: A cache-based system management architecture. In *In Proc. 2nd Symposium on Networked Systems Design and Implementation (NSDI)*, pages 259–272, 2005.
- [2] F. Dogar, A. Phanishayee, H. Pucha, O. Ruwase, and D. Andersen. Ditto - a system for opportunistic caching in multi-hop wireless mesh networks. In *Proceedings of ACM MobiCom*, San Francisco, CA, Sept. 2008.
- [3] C. Dubnicki, L. Gryz, L. Heldt, M. Kaczmarczyk, W. Kilian, P. Strzelczak, J. Szczepkowski, C. Ungureanu, and M. Welnicki. Hydrastor: a scalable secondary storage. In *FAST '09: Proceedings of the 7th Conference on File and Storage Technologies*, pages 197–210, Berkeley, CA, USA, 2009. USENIX Association.
- [4] A. Gulati, C. Kumar, and I. Ahmad. Storage workload characterization and consolidation in virtualized environments. In *2nd International Workshop on Virtualization Performance: Analysis, Characterization, and Tools (VPACT)*, 2009.
- [5] X. He. A caching strategy to improve iSCSI performance. *27th Annual IEEE Conference on Local Computer Networks, 2002. Proceedings. LCN 2002.*, pages 278–285, 2002.
- [6] M. Hibler, L. Stoller, J. Lepreau, R. Ricci, and C. Barb. Fast, scalable disk imaging with frisbee. In *2003 USENIX Annual Technical Conf.*, pages 283–296, San Antonio, TX, June 2003. USENIX Association.
- [7] J. Howard, M. Kazar, S. Menees, and DA. Scale and performance in a distributed file system. *ACM Transactions on Computer Systems (TOCS)*, 6(1):51–81, 1988.
- [8] D. Howells and R. Ltd. Fs-cache: A network filesystem caching facility. In *Proceedings of the Linux Symposium*, volume 1, 2006.
- [9] J. Kistler and M. Satyanarayanan. Disconnected operation in the Coda file system. *Mobile Computing*, 10(1):507–535, 1992.
- [10] Microsoft. Virtual hard disk image format specification, Feb. 2009.
- [11] D. Narayanan, A. Donnelly, E. Thereska, and S. Everest: Scaling down peak loads through I/O offloading. *Proc. Symposium on*, i:15–28, 2008.
- [12] H. Patterson, S. Manley, M. Federwisch, D. Hitz, S. Kleiman, and S. Owara. Proceedings of the FAST 2002 Conference on File and Storage Technologies for Disaster Recovery. *Appliance*, 2002.
- [13] C. Pettey and H. Stevens. Gartner says worldwide hosted virtual desktop market to surpass \$65 billion in 2013, Mar. 2009.
- [14] B. Pfaff, T. Garfinkel, and M. Rosenblum. Virtualization aware file systems: Getting beyond the limitations of virtual disks. In *In 3rd Symposium of Networked Systems Design and Implementation (NSDI)*, pages 353–366, 2006.
- [15] M. Satyanarayanan, J. H. Howard, D. a. Nichols, R. N. Sidebotham, A. Z. Spector, and M. J. West. The ITC distributed file system. *ACM SIGOPS Operating Systems Review*, 19(5):35–50, Dec. 1985.
- [16] E. Scholten. How to: Optimize guests for vmware view, July 2010.

- [17] M. Schroeder, D. Gifford, and RM. A caching file system for a programmer's workstation. *Operating Systems*, 19(5):34, 1985.
- [18] VMWare. Virtual machine disk format (vmdk), 2010.
- [19] W. Vogels. File system usage in windows nt 4.0. In *ACM Symposium on Operating System Principles (Kiawah Island Resort)*, pages 93–109. ACM, 1999.
- [20] B. Zhu, K. Li, and R. H. Patterson. Avoiding the disk bottleneck in the data domain deduplication file system. In M. Baker and E. Riedel, editors, *FAST*, pages 269–282. USENIX, 2008.