

Large-Scale Data-Centric Systems – A Research Analysis

Ch.Dhanunjaya Rao¹, Prof.Er.Dr.G.Manoj Someswar²

1. Research Scholar, Dr.APJ Abdul Kalam Technical University, Lucknow, U.P., India
 2. Research Supervisor, Dr.APJ Abdul Kalam Technical University, Lucknow, U.P., India
- Corresponding Author: Ch.Dhanunjaya Rao

ABSTRACT: Vast scale information driven frameworks enable associations to store, control, and get an incentive from expansive volumes of information. They comprise of dispersed segments spread over an adaptable number of associated machines and include complex programming /equipment stacks with different semantic layers. These frameworks enable associations to take care of built up issues including a lot of information, while catalyzing new, information driven organizations, for example, web crawlers, interpersonal organizations, and distributed computing and information stockpiling specialist organizations. The multifaceted nature, decent variety, scale and quick advancement of vast scale information driven frameworks make it trying to create instinct about these frameworks, increase operational experience, and enhance execution. It is a critical research issue to build up a technique to plan and assess such frameworks in light of the exact conduct of the targeted workloads. Utilizing an exceptional gathering of nine mechanical workload hints of business-basic huge scale information driven frameworks, we build up a workload-driven plan and assessment technique for these frameworks and apply the strategy to address beforehand unsolved outline issues. Specifically, the exposition contributes the accompanying:

1. A calculated structure of separating workloads for substantial scale information driven frameworks into information get to designs, calculation examples, and load landing designs.
2. A workload investigation and union strategy that utilizes multi-dimensional, non-parametric measurements to extricate bits of knowledge and create delegate conduct.
3. Case investigations of workload examination for modern arrangements of Map Reduce and enterprise organize capacity frameworks, two cases of huge scale information driven frameworks.
4. Case investigations of workload-driven outline and assessment of a vitality efficient Map Reduce framework and Internet server farm arrange transport convention pathologies, two research themes that require workload-particular bits of knowledge to address. By and large, the postulation builds up a more target and orderly comprehension of a rising and imperative class of PC frameworks. The work in this paper advances quicken the reception of huge scale information driven frameworks to promote genuine problems significant to business, science and everyday consumers.

KEYWORDS: Enterprise Network Storage Measure, Solid State Drives (SSD), Client Side Access Patterns, Server Side Access Patterns, Map Reduce, Information Driven Frameworks

Date of Submission: 31-08-2018

Date of acceptance: 15-09-2018

I. INTRODUCTION

Workload Analysis and Design Insights - Enterprise Network Storage Measure the ground before raising urban areas. This is the second of two sections that apply the workload examination methods to extensive scale information driven frameworks. Here, we investigate two endeavour organize capacity workloads. The examination has an indistinguishable objectives i.e., to revelation plan and assessment bits of knowledge not generally accessible. While both Map Reduce and undertaking network stockpiling are cases of substantial scale information driven frameworks, they differ in the accompanying ways:

1. Enterprise system stockpiling frameworks have a more mind-boggling structure, with capacity servers being unmistakable from capacity customers, and each separating further into various semantic layers.
2. The workload does not contain calculation designs as that for Map Reduce, since the main role of the framework is to capacity and recover information.

3. Enterprise system stockpiling speaks to a set up sort of PC framework that has developed to wind up a sort of huge scale information driven framework because of the changing scale and nature of the information it stores. This is not at all like Map Reduce, which has been deliberately outlined as a substantial scale information-handling framework. Consequently, this part fills in as a representation of the workload examination past Map Reduce. The mix of the two parts shows that our workload examination technique can be effective for both developing and built up frameworks.

This research paper is sorted out as it is taken up later. We start by propelling the significance of analyzing undertaking system stockpiling workloads and portraying the follows and philosophy expansions in the part. We at that point talk about workload behaviour as far as customer side information get to designs and server-side information get to designs. We additionally detail the long haul advancement of information get to designs, which enables increment to confidence that the technique means undertaking system stockpiling use cases past the ones analyzed in this part. We close the part by featuring some venture organize capacity plan engineering patterns uncovered by the creation work loads, and condensing the more extensive ramifications of the part with respect to outline and assessment approaches.

II. MOTIVATION

Endeavour stockpiling frameworks are composed around an arrangement of information get to designs. The storage framework can be particular by planning to a specific information get to design; e.g., a capacity framework for gushing video bolsters different get to designs than a record vault. The better the entrance design is comprehended, the better the capacity framework plan. Bits of knowledge into get to designs have been gotten from the examination of existing le framework workloads, normally through follow investigation thinks about. While this is the right broad procedure for enhancing stockpiling framework plan, past methodologies have basic inadequacies, particularly given late changes in innovation patterns. In this part, we show another plan approach to conquer these deficiencies. The information put away on big business organize joined capacity frameworks is experiencing changes because of a crucial move in the hidden innovation patterns. We have watched three such patterns, including:

Scale: Data estimate develops at a disturbing rate, because of new sorts of social, business and scientific applications, and the want to "never erase" information.

Heterogeneity: The blend of information composes put away on these capacity frameworks is winding up progressively intricate, each having its own particular necessities and access designs. Combination: Virtualization has empowered the union of various applications and their information onto less capacity servers. These virtual machines (VMs) additionally show total information get to designs more perplexing than those from singular customers do.

Better outline of future stockpiling frameworks requires bits of knowledge into the changing access designs because of these patterns. While past follow contemplates, have been utilized to determine information get to designs, we trust that they have the accompanying deficiencies:

Uni-dimensional: Although existing strategies investigate numerous entrance qualities, they do as such each one in turn, without uncovering cross-trademark conditions.

Predisposition: capacity framework originators searching for specific designs in light of earlier mental models performed past investigations. This acquaints an inclination that requirements with be returned to in view of the new innovation patterns.

Storage server centric: Past le system studies focused primarily on storage servers. This creates a critical knowledge gap regarding client behaviour. To overcome these shortcomings, we propose a new design methodology backed by the analysis of storage system traces. We present a method that simultaneously analyzes multiple characteristics and their cross dependencies. We use a multi-dimensional, statistical correlation technique, called k-mean, that is completely agnostic to the characteristics of each access pattern and their dependencies. The K-means algorithm can analyze hundreds of dimensions simultaneously, providing added objectivity to our analysis. To further reduce expertise bias, we involve as many relevant characteristics as possible for each access pattern. In addition, we analyze patterns at different granularities (e.g., at the user session, application, le level) on the storage server as well as the client, thus addressing the need for understanding client patterns. The resulting design insights enable policies for building new storage systems.[1] We analyze two recent network-attached storage le system traces from a production enterprise data centre. Table 1 summarizes our key observations and design implications; they will be detailed later in the chapter. Our Methodology leads to observations that would be difficult to extract using past methods. We illustrate two such

access patterns, one showing the value of multi-granular analysis (Observation 1 in Table 1) and another showing the value of multi-feature analysis.

Client side observations and design implications	Server side observations and design implications
1. Client sessions with IO sizes > 128KB are read only or write only.) Clients can consolidate Sessions based on only the read-write ratio.	7. Files with >70% sequential read or write have no repeated reads or overwrites.) Servers should delegate sequentially accessed les to clients to improve IO performance.
2. Client sessions with duration >8 hours do 10MB of IO.) Client caches can already t an entire day's IO.	8. Engineering les with repeated reads have random accesses.) Servers should delegate repeatedly read les to clients; clients need to store them in ash or memory.
3. Number of client sessions drops o linearly by 20% from Monday to Friday.) Servers can get an extra \day" for background tasks by running at appropriate times during week days.	9. Most les are active (have opens, IO, and metadata access) for only 1-2 hours in a few months.) Servers can use le idle time to compress or deduplicate to increase storage capacity.
4. Applications with <4KB of IO per le open and many opens of a few les do only random IO.) Clients should always cache the first few KB of IO per le per application.	10. All les have either all random access or >70% sequential access. (Seen in past studies too)) Servers can select the best storage medium for each le based on only access sequentiality.
5. Applications with >50% sequential read or write access entire les at a time.) Clients can request le pre-fetch (read) or delegation (write) based on only the IO sequentiality.	11. Directories with sequentially accessed les almost always contain randomly accessed les as well.) Servers can change from per-directory placement policy (default) to per- le policy upon seeing any sequential IO to any les in a directory.
6. Engineering applications with >50% sequential read and sequential write are doing code compile tasks, based on le extensions.) Servers can identify compile tasks; server should cache the output of these tasks.	12. Some directories aggregate only les with repeated reads and overwrites.) Servers can delegate these directories entirely to clients, tradeo s permitting.

Table 1: Summary of design insights, separated into insights derived from client access patterns and server access patterns. The workload analysis gives us high confidence that the proposed improvements will bring a benefit. We defer to future work the implementation of these system changes and the quantisation of performance gains per workload

First, we observe (Observation 1) that sessions with more than 128KB of data reads or writes are either read-only or write-only. This observation affects shared caching and consolidation policies across sessions. Specifically, client OSs can detect and co-locate cache sensitive sessions (read-only) with cache insensitive sessions (write-only) using just one parameter (read-write ratio). This improves cache utilization and consolidation (in-creased density of sessions per server).

Similarly, we observe (Observation 8) that les with >70% sequential read or sequential write have no repeated reads or overwrites.[2] This access pattern involves four characteristics: read sequentiality, write sequentiality, repeated read behaviour, and over-write behaviour. The observation leads to a useful policy: sequentially accessed les do not need to be cached at the server (no repeated reads), which leads to an efficient buffer cache.

These observations illustrate that our methodology can derive unique design implications that leverage the correlation between different characteristics. To summarize, our contributions are:

Identify storage system access patterns using a multi-dimensional statistical analysis technique. Build a framework for analyzing traces at different granularity levels at both server and client. Analyze our specific traces and present the access patterns identified. Derive design implications for various storage system components from the access patterns.

Traces and Methodology Extensions

In this section, we describe our analysis method in detail. We start with a description of the traces we analyzed, followed by a description of the access units selected for our study. Next, we describe key steps in our analysis process, including selecting the right features for each access unit, using the k-means data clustering algorithm to identify access patterns, and additional information needed to interpret and generalize the results.

Traces Analyzed

We collected Common Internet File System (CIFS) traces from two large-scale, enterprise-class le servers deployed at our corporate data centres. One server covers roughly 1000 employees in marketing, sales, finance, and other corporate roles. We call this the corporate trace. The other server covers roughly 500 employees in various engineering roles. We call this the engineering trace.

The trace collecting infrastructure is described. The corporate trace rejects activities on 3TB of active storage from 20th September, 2015 to 21st November, 2015. It contains activity from many Windows applications. The engineering trace reflects activities on 19TB of active storage from 10th August, 2015 to 14th November, 2015. It interleaves activity from both Windows and Linux applications. In both cases, many clients use virtualization technologies. Thus, we believe we have representative traces with regards to the technology trends in scale, heterogeneity, and consolidation. Also, since protocol-independent users, applications, and stored data remain the primary factors affecting storage system behaviour, we believe our analysis is relevant beyond CIFS.

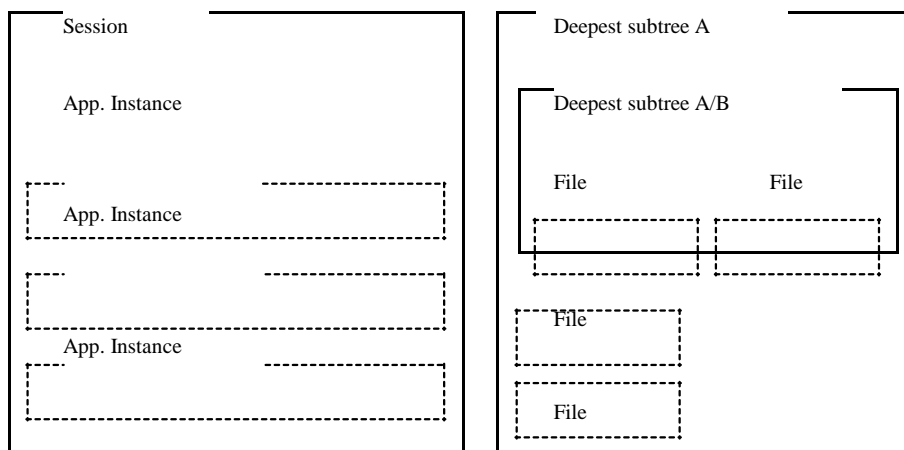


Figure 1: Access units analyzed. At clients, each session contains many application instances. At servers, each sub-tree contains many files

Access Units

We analyze access patterns at multiple access units at the server and the client. Selecting access units is subjective. We chose access units that form clear semantic design boundaries. On the client side, we analyze two access units:

Sessions: Sessions reflect aggregate behaviour of a user. A CIFS session is bounded by matching session connect and logo requests. CIFS identifies it by a tuple | f client IP address, session IDg. **Application instance:** Analysis at this level leads to application specific optimizations in client VMs. CIFS identifies each application instance by the tuple – f client IP address, session ID, and process IDg. We also analyzed file open-close, but obtained uninteresting insights. Hence, we omit that access unit in the early part of our research work.

We also examined two server side access units:

File: Analyzing file level access patterns facilitates per- file policies and optimization techniques. Each file is uniquely identified by its full path name.[3] **Deepest sub tree:** the directory path immediately containing the file identifies this access unit. Analysis at this level enables per-directory policies. The diagram shows the semantic hierarchy among different access units. At clients, each session contains many application instances. At servers, each sub-tree contains many files.

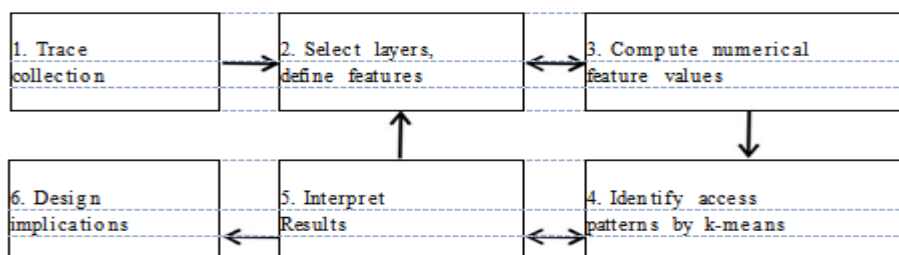


Figure 2: Methodology overview. The two-way arrows and the loop from Step 2 through Step 5 indicate our much iteration between the steps

Analysis Process

Our method (Figure 2) involves the following steps:

1. Collect network storage system traces.
2. Define the descriptive features for each access unit. This step requires domain knowledge about storage systems.
3. Extract multiple instances of each access unit, and compute from the trace the corresponding numerical feature values of each instance.
4. Input those values into k-means, a multi-dimensional statistical data clustering technique.
5. Interpret the k-means output and derive access patterns by looking at only the relevant subset of features. This step requires knowledge of both storage systems and statistics. We also need to extract considerable additional information to support our interpretations.
6. Translate access patterns to design insights.

We give more details about Steps 2, 4, and 5 in the following subsections.

Selecting features for each access unit

Choosing the arrangement of engaging highlights for each entrance unit requires area information about capacity frameworks (Step 2 in Figure 2). It likewise presents some subjectivity, since the selection of highlights restricts how one access example can vary from another. The human fashioner needs to choose some essential highlights at first, e.g., add up to IO size and read-compose proportion for a le.

We won't know whether we have a decent arrangement of highlights until the point that we have finished the whole investigation process. On the off chance that the investigation comes about abandon some plan decision ambiguities; we have to add new highlights to elucidate those ambiguities, again utilizing space information. For instance, for the most profound sub trees, we figure different percentiles (25th, 50th, and 75th) of specific highlights like read-compose proportion in light of the fact that the normal incentive for those highlights did not plainly isolate the entrance designs.

We at that point rehash the examination procedure utilizing the new list of capabilities. This iterative procedure prompts a long list of capabilities for all entrance units, fairly decreasing the subjective inclination of a little list of capabilities. We list the picked highlights for each entrance unit.

The greater part of the highlights utilized as a part of our investigations are clear as crystal. Some questionable or complex highlights require exact definitions, for example,

I/O: We utilize I/O" as a substitute for read and compose" Successive processes or composes: We consider two read or composes solicitations to be consecutive in the event that they are back to back in time, and the le o set + ask for size of the principal ask for measures up to the le o set of the second demand. A solitary read or compose ask for is by definition not consecutive.

Rehashe peruses or overwrites: We track gets to at 4KB piece limits inside a le, with the o set of the principal square being zero. A read is viewed as rehashed on the off chance that it gets to a piece that has been perused in the past half hour. We utilize an identical definition for overwrites.

Identifying access designs through k-implies

A key piece of our system is the k-implies multi-dimensional connection calculation. We utilize it to recognize get to designs at the same time crosswise over numerous highlights (Step 4 in Figure 2). The research work shows that k-implies a notable, factual relationship algorithm. It identifies sets of information focuses that gather around a locale in n-dimensional space. These assemblies are called groups.

For each entrance unit, we separate different cases of it from the follow, i.e., all session examples, application occasions, les, and indexes. For each occasion, we process the numerical estimations of every one of its highlights. This gives us an information exhibit in which each line compares to an occurrence, i.e., an information point, and every segment relates to an element, i.e., a measurement. We input the exhibit into k-implies, and the calculation finds the regular groups over all information focuses. We consider all information focuses in a bunch as having a place with a solitary comparability class, i.e., a solitary access design. The numerical estimations of the bunch focuses show the qualities of each entrance design.[4]

We pick k-implies for two reasons. Initially, k-implies is algorithmically straightforward. This permits fast handling on extensive informational collections. We utilized a modified rendition of the k-implies C library, in which we made some minor alters to constrain the memory impression when preparing huge information sizes. Second, k-implies prompts instinctive names of the bunch focuses. This causes us decipher the measurable conduct separated from the follows into unmistakable bits of knowledge. Subsequently, we lean toward k-intends to other bunching calculations, for example, various levelled grouping and k-implies subordinates.

K-implies expects us to determine k, the quantity of bunches. This is a di ffection errand since we don't have the foggiest idea about from the earlier the quantity of natural" bunches in the information. As part of this

research work, we utilize a standard strategy to set k: increase k until there is lessening rate in the abatement of intra-bunch fluctuation, i.e., remaining change. Additional points of interest of the procedure shows up.

We get a subjective measure of the clearness of bunch limits by taking a gander at the re-task of information focuses as we augment k. On the off chance that the bunches limits are clear, i.e., the information focuses fall into characteristic groups, at that point augmenting k would prompt a solitary group split into two, and the information focuses in alternate groups stay stationary. Then again, indistinct bunch limits are demonstrated by N to N + 1 improvements in information focuses with N 2. The bunch limits are clear for every one of the outcomes exhibited in this research work. There are maybe a couple occurrences of a twofold split, i.e., two bunches split into three, however every one of the three of the groups are saved when we facilitate increase k. Such a circumstance speaks to three generally equidistant common bunches being identified. In total, clear group limits would be demonstrated by a tree-like structure in the development of information focuses as we increment k. Figure 3 shows such a graphical portrayal, to the point that incorporates a twofold split. We verified that such a tree-like structure exists for all the k-implies comes about displayed in this part.

Interpreting and summing up the outcomes

The k-implies calculation gives us an arrangement of access designs with different attributes. We require extra data to comprehend the significance

of the outcomes. This information originates from processing different optional information outside of k-implies examination. We assembled the begin and end times of every session occurrence, amassed by times of the day and days of the week. This gave us knowledge into how clients dispatch and end sessions.

We inspect rename augmentations of les related with each entrance design having a place with these entrance units: application occasions, les, and most profound sub-trees. This information interfaces the entrance examples to all the more effortlessly conspicuous le augmentations. We perform connection examination between the le and most profound sub-trees get to units. Specifically, we register the quantity of les of every le get to design that is situated inside registries in each most profound sub-tree get to design. This data catches the associations of les in registries. Such data gives us a point by point picture about the semantics of the entrance designs, bringing about human reasonable marks to the entrance designs. Such marks enable us to make an interpretation of perceptions to plan suggestions.

Besides, subsequent to distinguishing the plan suggestions, we investigate if the outline in-sights can be extrapolated to other follow periods and other stockpiling framework utilize cases. We achieve this by rehashing our correct investigation over numerous subsets of the follows, for instance, seven days of follows at once. This enable us to analyze how our examination would be different had we acquired just seven days' follow. Access patterns that are consistent, stable across different weeks would indicate that they are likely to be more general than just our tracing period or our use cases.

(a). Descriptive features for each session

Duration	Avg. time between I/O requests	Unique trees accessed
Total I/O size	Read sequentiality	File opens
Read:write ratio by bytes	Write sequentiality	Unique les opened
Total I/O requests	Repeated read ratio	Directories accessed
Read:write ratio by requests	Overwrite ratio	Application instances seen
Total metadata requests	Tree connects	

(b). Corporate session access patterns	Full day work	Half day Content Viewing	Short content viewing	Short content generate	Supporting meta-data	Supporting read-write
% of all sessions	0.5%	0.7%	1.2%	0.2%	96%	1.4%
Duration	8 hrs	4 hrs	10 min	70 min	7 sec	10 sec
Total I/O size	11 MB	3 MB	128 KB	3 MB	0	420 B
Read:write ratio by bytes	3:2	1:0	1:0	0:1	0:0	1:1
Metadata requests	3000	700	230	550	1	20
Read sequentiality	70%	80%	0%	-	-	0%
Write sequentiality	80%	-	-	90%	-	0%
File opens: les	200:40	80:15	30:7	50:15	0:0	6:3
Tree connect:Trees	5:2	3:2	2:2	2:2	1:1	2:2
Directories accessed	10	7	4	6	0	2
Application instances	4	3	2	2	0	1

(c). Engineering session access patterns	Full day work	Human edit small Les	App. generated backup	Short content generate	Supporting meta-data	Machine generated update
% of all sessions	0.4%	1.0%	4.4%	0.4%	90%	3.6%
Duration	1 day	2 hrs	1 min	1 hr	10 sec	10 sec

Total I/O size	5 MB	5 KB	2 MB	2 MB	0	36 B
Read:write ratio	7:4	1:1	1:0	0:1	0:0	1:0
Metadata requests	1700	130	40	200	1	0
Read sequentiality	60%	0%	90%	-	-	0%
Write sequentiality	70%	0%	-	90%	-	-
File opens: les	130:20	9:2	6:5	15:6	0:0	1:1
Tree connect:Trees	1:1	1:1	1:1	1:1	1:1	1:1
Directories accessed	7	2	1	3	0	1
Application instances	4	2	1	1	0	1

Table 2: Session access patterns. (a): Full list of descriptive features. (b) and (c): Short names and descriptions of sessions in each access pattern; listing only the features that help separate the access patterns

Client Side Access Patterns

This and the coming sections present the access patterns we identified and the accompanying design insights. We discuss client and server side access patterns. We also check if these patterns persist across time. For each access unit, we list

the descriptive features (only some of which help separate access patterns), outline how we derived the high-level name (label) for each access pattern, and discuss relevant design insights.

Sessions

Sessions reflect total conduct of human clients. We utilized 17 highlights to portray sessions (Table 2). The corporate follow has 509,076 sessions, and the building follow has 232,033.

In Table 2, we give quantitative portrayals and short names for all the session get to designs. We get the names from looking at the significant cannot highlights: term, read-compose proportion, and I/O measure.

We likewise took a gander at the total session begin and end times to get extra semantic learning about each entrance design.[5] Figures 4 and 5 demonstrate the begin and end times for chosen session get to designs. The begin times of corporate entire day work sessions compare precisely to the U.S. workday {9AM begin, 12PM lunch, 5PM end. Corporate substance age sessions demonstrate slight increment at night and towards Friday, showing hurries to meet day by day or week after week due dates. In the building follow, the application-produced reinforcement and machine-created refresh sessions leave significantly from human workday and work-week designs, driving us to name them as application and machine (customer OS) produced.

One shock was that the `supporting metadata' sessions represent >90% of all sessions in the two follows. We trust these sessions are not humanly created. They last around 10 seconds, leaving brief period for human intervened collaborations. Likewise, the session begin rate midpoints to approximately one for every representative for each moment. We are sure that clients are not interfacing and logging o each moment of the whole day. Be that as it may, the state of the begin

time charts has a solid connection with the human workday and work week. We call these supporting metadata sessions {machine created in help of human client exercises. These metadata sessions shape a kind of `background clamor" to the capacity framework. We watch a similar foundation clamor at different layers both at customers and servers.

Perception 1: The sessions with I/O sizes more prominent than 128KB are either perused just or compose, aside from the entire day work sessions (Table 2). These perceptions compare to the `half day content review," short substance seeing," and short con-tent produce" sessions in the corporate follow, and application created reinforcement" and short content create" sessions in the designing follow. Among these sessions, just read-just sessions use buffer store for rehashed peruses and pre-fetches.

Compose just sessions just utilize the store to buffer composes. In this manner, on the off chance that we have a reserve expulsion strategy that perceives their compose just nature and discharges the buffers promptly on using filthy information, we can fulfill numerous compose just sessions with moderately little buffer store space. We can accomplish better union and buffer store use by dealing with the proportion of co-found read-just and compose just sessions. Virtualization directors and customer working frameworks to deal with a mutual buffer store between sessions can utilize this knowledge. Perceiving such read-just and compose just sessions is simple. Analyzing a session's aggregate perused estimate and compose measure uncovers their read-just or compose just nature. Suggestion 1: Clients can combine sessions efficiently construct just with respect to the read-compose proportion.

Observation 2: The full-day work" sessions do 10MB of I/O (Table 2). This means that a client cache of 10s of MB can t the working set of a day for most sessions. Given the growth of ash devices on clients for caching, despite large-scale consolidation, clients should easily cache a day's worth of data for all users. In such a scenario, most I/O requests would be absorbed by the cache, reducing network latency and bandwidth utilization, and load on the server. Moreover, complex cache eviction algorithms are unnecessary. Implication 2: Clients caches can already t an entire day's I/O.

Perception 3: The quantity of human-produced sessions and supporting session's tops on Monday and abatements consistently to 80% of the top on Friday (Figure 3). There is extensive slack" in the server stack amid nights, lunch times, and not withstanding amid working hours. This infers the server can perform foundation errands, for example, consistency checks, upkeep, or pressure/deduplication, at suitable circumstances amid the week. A basic tally of dynamic sessions can fill in as a viable begin and stop flag. By figuring the zone under the bend for session begin times by days of the week, we gauge that foundation undertakings can crush out approximately one additional day of preparing without changing the pinnacle request on the framework. This is a half change over a setup that performs foundation undertakings just amid ends of the week. In the building follow, the application produced reinforcement or duplicate sessions appear to have been as of now outlined along these lines. Suggestion 3: Servers get an additional day" for foundation errands by running them at proper circumstances amid week-days.

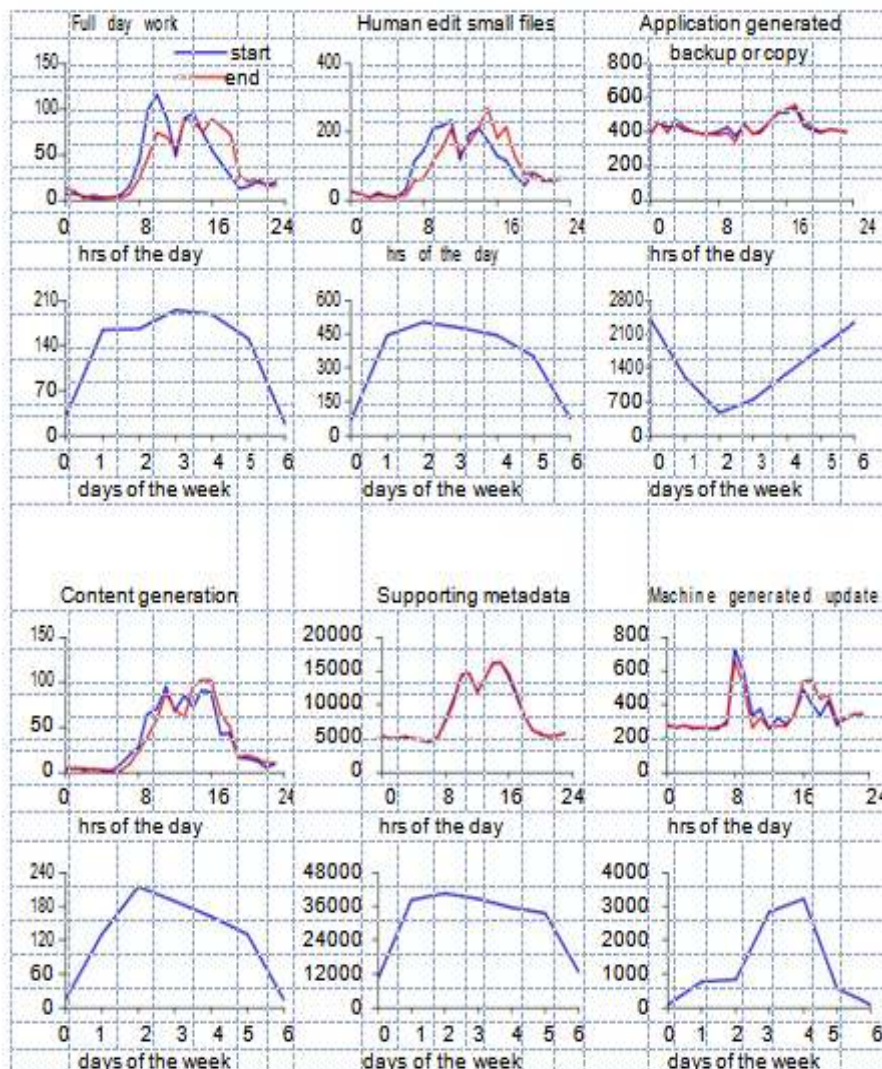


Figure 3: Graphical Representations of Human Produced Periodical Sessions

Application instances

Application instance access patterns reflect application behaviour, facilitating application specific optimizations. We used 16 features to describe application instances (Table 3). The corporate trace has 138,723 application instances, and the engineering trace has 741,319.

Table 3 provides quantitative descriptions and short names for all the application instance access patterns. We derive the names from examining the read-write ratio, I/O size, and le extensions accessed (Figures 6 and 7).

(a). Descriptive features for each application instance

	Viewing app generated content	Supporting meta-data	App generated updates	Viewing human generated content	Content update app
(b). Corp. app. instance access patterns					
% of all app instances	16%	56%	14%	8.8%	5.1%
Total I/O	100 KB	0	1 KB	800 KB	3.5 MB
Read:write ratio	1:0	0:0	1:1	1:0	2:3
Metadata requests	130	5	50	130	500
Read sequentiality	5%	-	0%	80%	50%
Write sequentiality	-	-	0%	-	80%
Overwrite ratio	-	-	0%	-	5%
File opens: les	19:4	0:0	10:4	20:4	60:11
Tree connect:Trees	2:2	0:0	2:2	2:2	2:2
Directories accessed	3	0	3	3	4
File extensions accessed	2	0	2	2	3

	Compilation app	Supporting meta-data	Content update app - small	Viewing human generated content	Content viewing app - small
(c). Eng. app. instance access patterns					
% of all app instances	1.6%	93%	0.9%	2.0%	2.5%
Total I/O	2 MB	0	2 KB	1 MB	3 KB
Read:write ratio	9:1	0:0	0:1	1:0	1:0
Metadata requests	400	1	14	40	15
Read sequentiality	50%	-	-	90%	0%
Write sequentiality	80%	-	0%	-	-
Overwrite ratio	20%	-	0%	-	-
File opens: les	145:75	0:0	3:1	5:4	2:1
Tree connect:Trees	1:1	0:0	1:1	1:1	1:1
Directories accessed	15	0	1	1	1
File extensions accessed	5	0	1	1	1

Table 3: Application instance access patterns. (a): Full list of descriptive features. Almost identical to those for user sessions, with a new feature "file extensions accessed", and two features no longer applicable "duration" and "application instances seen". (b) and (c): Short names and descriptions of application instances in each access pattern; listing only the features that help separate the access patterns

We see again the metadata background noise. The supporting metadata application instances account for the largest fraction, and often do not even open a file. There are many files without a file extension, a phenomenon also observed in recent storage system snapshot studies. We notice that file extensions turn out to be poor indicators of application instance access patterns. This is not surprising because we separate access patterns based on read/write properties. A user could either view doc or create a .doc. The same application software has different read/write patterns. This speaks to the strength of our multi-layer framework. Aggregating I/O by application instances gives clean separation of patterns; while aggregating just by application software or file extensions will not.[6]

We also find it interesting that most file extensions are immediately recognizable. This means that what people use network storage systems for, i.e., the file extensions, remains easily recognizable, even though how people use network storage systems, i.e., the access patterns, is ever changing and becoming more complex.

Observation 4: The small content viewing" and content update" application instances have <4KB total reads per file open and access a few unique files many times (Table 3). The small read size and multiple reads from the same files means that clients should pre-fetch and place the files in a cache optimized for random access (flash SSD memory). The trend towards flash caches on clients should enable this transfer.

Application instances have bi-modal total I/O size - either less than 10KB, or 100KB-10MB. Thus, a simple cache management algorithm success; we always keep the first 2 blocks of 4KB in cache. If the application instance does more I/O, it is likely to have I/O size in the 100KB-1MB range, so we evict it from the cache. We should note that such a policy makes sense even though we proposed earlier to cache all 11MB of a typical day's working set - 11MB of cache becomes a concern when we have many consolidated clients. Implication 4: Clients should always cache the first few KB of I/O per file per application.

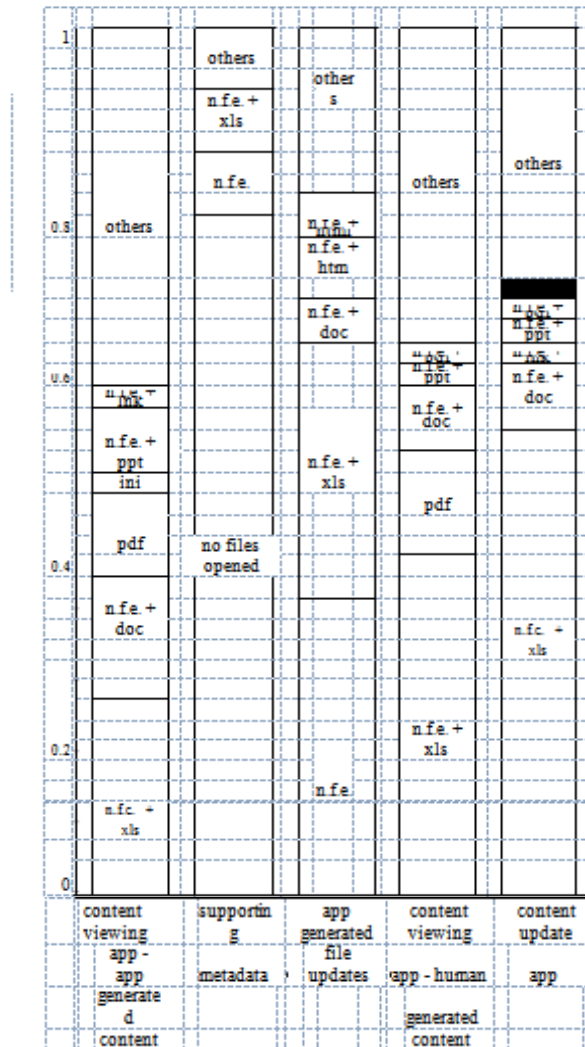


Figure 4: File extensions for corporate application instance access patterns. For each access pattern (column), showing the fraction of the two most frequent file extensions that are accessed together within a single application instance. \n.f.e." denotes files with no file extension". Application instances labelled with just one file extension accessed files with only one extension.

Observation 5: We see >50% sequential read and write ratio for the content update" applications instances for corporate and the content viewing" applications instances for human-generated content for both corporate and engineering (Table 3). Dividing the total I/O size by the number of file opens suggest that these application instances are sequentially reading and writing entire files (e.g. productivity (.xls, .doc, .ppt, .pdf) and multimedia applications.

This infers the files related with these applications ought to be pre gotten and designated to the customer. Pre bringing implies conveying the entire file to the customer before the entire file is asked. Designation implies giving a customer brief, restrictive access to a file, with the customer occasionally synchronizing to server to guarantee information toughness. CIFS does assignment utilizing default locks, while NFSv4 has a committed activity for appointment. Pre getting and assignment of such files will enhance read and compose execution, bring down system activity, and relieve server burden.[7]

The entrance designs again offer a straightforward, edge based choice calculation. On the off chance that an application example accomplishes more than 10s of KB of successive I/O, and has no overwrite, at that point it is probably going to be a substance survey or refresh application case; such files are pre-fetched and assigned to the customers. Suggestion 5: Clients can ask for file pre-fetch (read) and designation (compose) in view of just I/O sequentiality.

Perception 6: Engineering applications with >50% successive peruses and >50% sequential composes are doing code assemble undertakings. We know

this from taking a gander at the le augmentations in Figure 4. These incorporate procedures indicate read sequentiality, compose sequentiality, a huge overwrite proportion and huge number of metadata demands.[9] They depend on the server intensely for information gets to. We require more nitty gritty customer side data to comprehend why customer stores are ineffectual for this situation. In any case, unmistakably the server reserve needs to pre bring the read les for these applications. The high percent-period of consecutive peruses and composes gives us another limit based calculation to recognize these applications. Suggestion 6: Servers can measurably distinguish assemble undertakings by the nearness of both successive peruses and composes; server needs to reserve the yield of these errands.

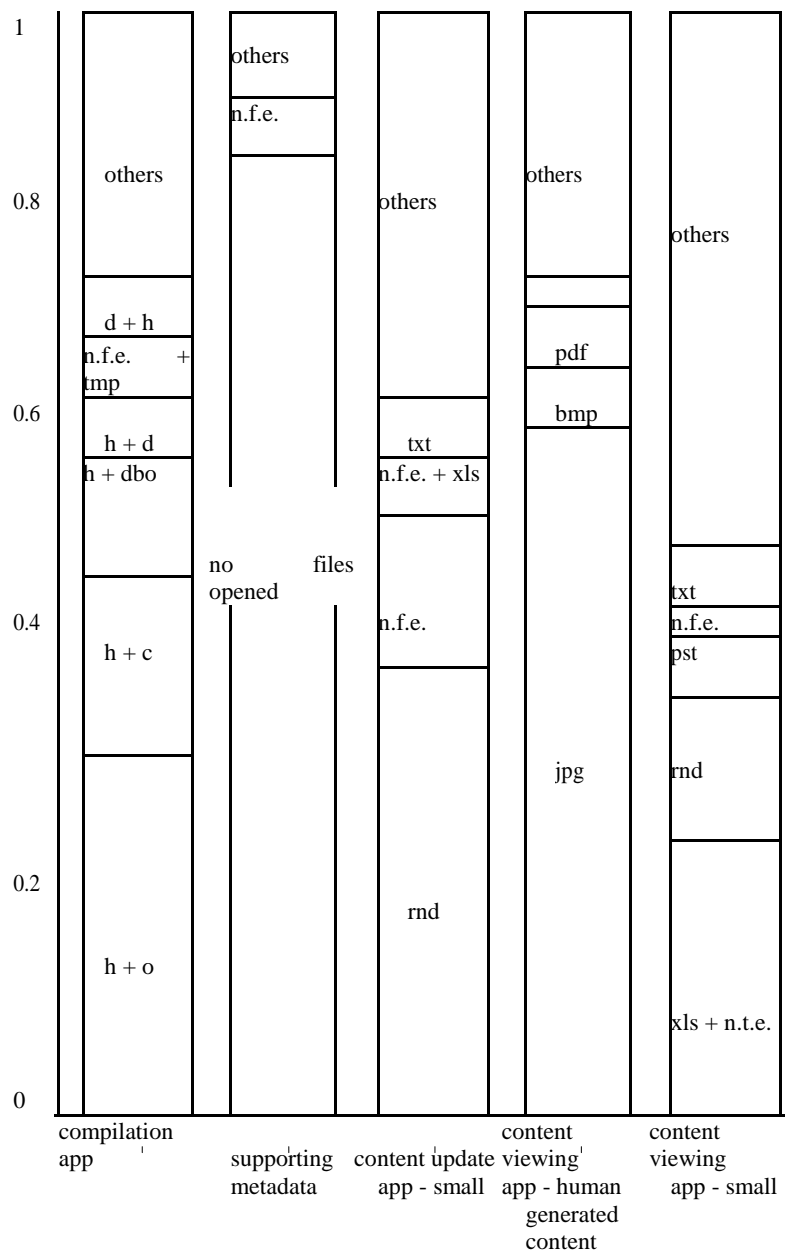


Figure 5: File extensions for engineering application instance access patterns. For each access pattern (column), showing the fraction of the two most frequent le extensions that are accessed together within a single application instance. \n.f.e." denotes les with no le extension". Application instances labelled with just one le extension accessed les with only one extension.

Server Side Access Patterns

As specified, we dissected two sorts of server side access units: les and most profound sub-trees.

Files

Document get to designs help stockpiling server fashioners create per-le position and optimization strategies. We utilized 25 highlights to depict les (Table 4). Note that a portion of the highlights incorporate different percentiles of a trademark, e.g., read ask for estimate as percentiles of all read demands. We think including different percentiles instead of simply the normal would permit better partition of access designs. The corporate follow has 1,155,099 les, and the building follow has 1,809,571.

In Table 4, we show quantitative depictions and short names for all the le get to designs. Figures 8 and 9 give the most widely recognized le augmentations in each. [10] We inferred the names by analyzing the read-compose proportion and I/O measure. For the building follow, inspecting the le expansions likewise demonstrated helpful, prompting names, for example, edit code and aggregate yield", and read-just log reinforcement".

We see that there are groupings of les with comparative expansions. For instance, in the corporate follow, the little irregular read get to designs incorporate numerous le augmentations associated with web program stores. Additionally, multi-media les like .mp3 and .jpg assemble in the successive read and compose get to designs. In the designing follow, code libraries bunch under the consecutive compose les, and read just log reinforcement les contain le extensions .0 to .99. Be that as it may, the most widely recognized le augmentations in each follow still spread crosswise over numerous entrance designs, e.g., office profitability les in the corporate follow and code les in the building follow. Observation 7: For les with >70% sequential reads or sequential writes, the repeated read and overwrite ratios are close to zero (Table 4). This implies that there is little benefit in caching these les at the server. They should be pre-fetched as a whole and delegated to the client.[11]

(a). Descriptive features for each le

- Number of hours with 1, 2-3, or 4 le opens
- Number of hours with 1-100KB, 100KB-1MB, or >1MB reads
- Number of hours with 1-100KB, 100KB-1MB, or >1MB writes
- Number of hours with 1, 2-3, or 4 metadata requests
- Read request size - 25th, 50th, and 75th percentile of all requests
- Write request size - 25th, 50th, and 75th percentile of all requests
- Avg. time between I/O requests - 25th, 50th, and 75th percentile of all request pairs
- Read sequentiality
- Write sequentiality
- Read:write ratio by bytes
- Repeated read ratio
- Overwrite ratio

(b). Corp. le access patterns	Metadata only	Sequent-ial write	Sequent-ial read	Small random write	Smallest random read	Small random read
% of all les	59%	4.0%	4.1%	4.7%	19%	9.2%
# hrs with opens	2hrs	1hr	1hr	1hr	1hr	1hr
Opens per hr	1 open	2-3 opens	2-3 opens	2-3 opens	1 open	1 open
# hrs with reads	0	0	1hr	0	1hr	1hr
Reads per hr	-	-	100KB-1MB	-	1-100KB	1-100KB
# hrs with writes	0	1hr	0	1hr	0	0
Writes per hr	-	100KB-1MB	-	1-100KB	-	-
Read request size	-	-	4-32KB	-	2KB	32KB
Write request size	-	60KB	-	4-22KB	-	-
Read sequentiality	-	-	70%	-	0%	0%
Write sequentiality	-	80%	-	0%	-	-
Read:write ratio	0:0	0:1	1:0	0:1	1:0	1:0

(c). Eng. le access patterns	Metadata only	Sequent-ial write	Small random read	Edit code & compile	Sequent-ial read	Read-only log/backup
% of all les	42%	1.9%	32%	7.3%	8.3%	8.1%
# hrs with opens	1hr	1hr	1hr	1hr	1hr	2hrs
Opens per hr	1 open	2-3 opens	2-3 opens	2-3 opens	2-3 opens	2-3 opens
# hrs with reads	0	0	1hr	1hr	1hr	2hrs
Reads per hr	-	-	1-100KB	1-100KB	1-100KB	1-100KB
# hrs with writes	0	1hr	0	0	0	0
Writes per hr	-	>1MB	-	-	-	-
Read request size	-	-	3-4KB	4KB	8-16KB	1KB
Write request size	-	64KB	-	-	-	-
Read sequentiality	-	-	0%	0%	70%	0%
Write sequentiality	-	90%	-	-	-	-
Repeated read ratio	-	-	0%	50%	0%	0%
Read:write ratio	0:0	0:1	1:0	1:0	1:0	1:0

Table 4: File access patterns. (a): Full list of descriptive features. (b) and (c): Short names and descriptions of les in each access pattern; listing only the features that help separate the access patterns

Once more, the bimodal I/O sequentiality offers a basic calculation for the server to identify which les ought to be pre-brought and assigned { if a le has any successive access, it is probably going to have a high level of consecutive access, hence it ought to be pre-fetched and designated to the customer. Future stockpiling servers can recommend such data to customers, prompting assignment demands. Suggestion 7: Servers should assign successively got to les to customers to enhance I/O execution.

Perception 8: In the building follow, just the alter code and incorporate yield les have a high % of rehashed peruses (Table 4). Those les ought to be assigned to the customers also. The rehashed peruses don't appear in the building application examples, potentially on the grounds that an assemblage procedure dispatches numerous kid forms over and again perusing similar les. Every tyke procedure peruses \"fresh information,\" despite the fact that the server sees rehashed peruses. With bigger memory or powder reserves at customers,[12] we anticipate that this conduct will drop. The working set issues that prompt this situation need to be analyzed. On the off chance that the rehashed peruses originate from a solitary customer, at that point the server can recommend that the customer reserve the proper les.

We can again utilize an edge based calculation. Recognizing any rehashed peruses at the server flags that the le ought to be assigned to the customer. Even from a pessimistic standpoint, just the first few peruses will hit the server. Consequent rehashed peruses are ceased at the customer. Suggestion 8: Servers should appoint over and over read les to customers.

We can again utilize a limit based calculation. Recognizing any rehashed persuals at the server flags that the le ought to be appointed to the customer. Best case scenario, just the first few peruses will hit the server. Resulting rehashed peruses are ceased at the customer. Suggestion 8: Servers should assign over and again read les to customers.

Perception 9: Almost all les are dynamic (have opens, I/O, and metadata access) for just 1-2 hours over the whole follow period, as demonstrated by the run of the mill opens read compose action of all entrance designs (Table 4). There are some consistently got to les, yet they are few to the point that they don't influence the k-implies examination. The absence of standard access for most les implies that there is space for the server to utilize procedures to build limit by doing compaction on sit still les.[13]

Basic procedures incorporate deduplication and pressure. The movement on these les demonstrate that the I/O execution effect ought to be little. Regardless of whether run continually, compaction has a low likelihood of influencing a dynamic le. Since normal libraries like gzip enhance for decompression, decompressing les at read time ought to have just slight execution affect. Suggestion 9: Servers can utilize le sit without moving time to pack or deduplicate information to expand capacity limit.

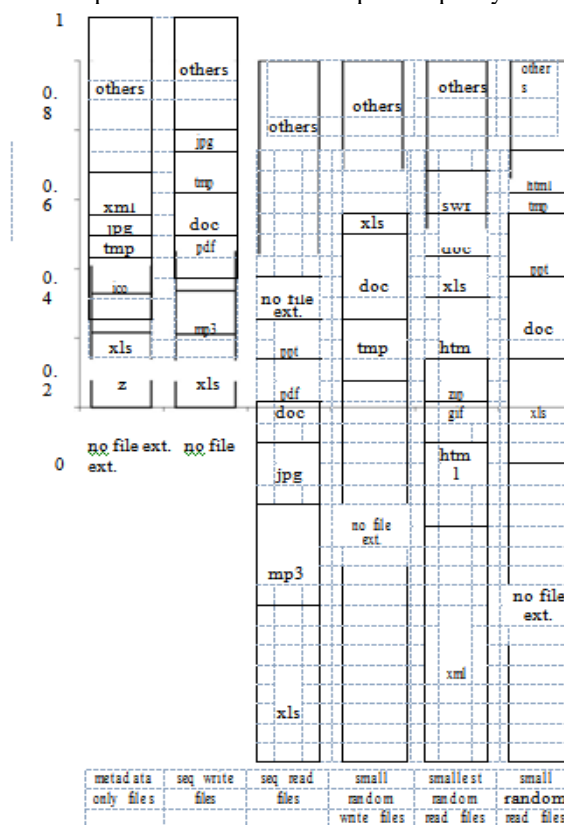


Figure 6: File extensions for corporate les. Fraction of le extensions in each le access pattern

Observation 10: All les have either all random access or >70% sequential access (Table 4). The small random read and write les in both traces can benefit from being placed on media with high random access performance, such as solid state drives (SSDs). Files with a high percentage of sequential accesses can reside on traditional hard disk drives (HDDs), which already optimize for sequential access. The bimodal I/O sequentiality offers yet another threshold-based placement algorithm { if a le has any sequential access, it is

likely to have a high percentage of sequential access; therefore place it on HDDs. Otherwise, place it on SSDs. We note that there are more randomly accessed les than sequentially accessed les.[14] Even though sequential les tend to be larger, we still need to do a working set analysis to determine the right size of server SSDs for each use case. Implication 10: Servers can select the best storage medium for each le based only on access sequentiality.

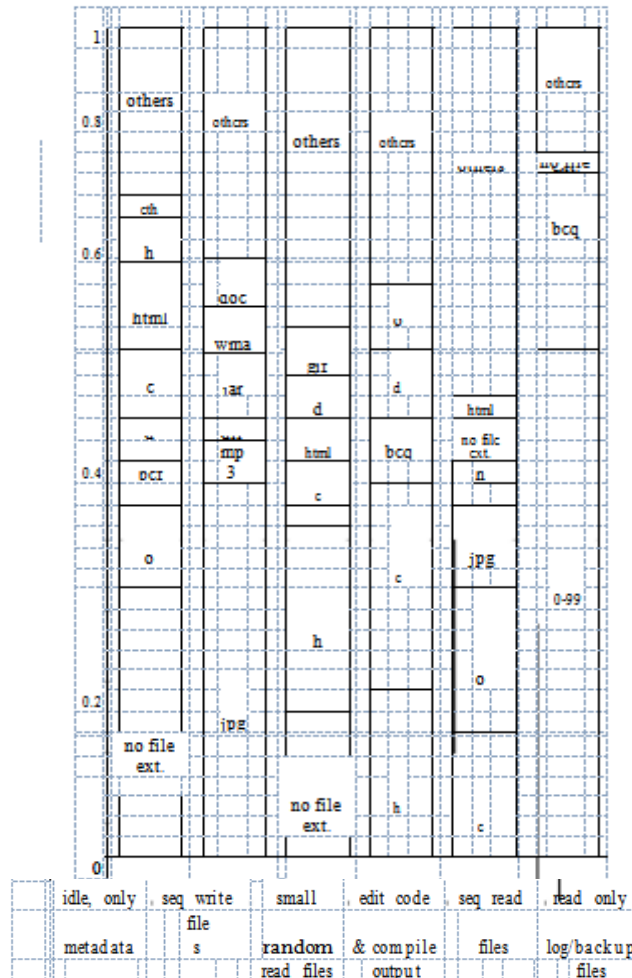


Figure 7: File extensions for engineering les. Fraction of le extensions in each le access pattern

Deepest sub trees

Deepest sub tree access patterns help storage server designers develop per-directory policies. We use the phrase \deepest sub tree" as a more precise alternative to directories". In particular, for hierarchical directory structures

of the form A/B, data accesses to les A/file are counted as deepest sub tree A, while data accesses to les A/B/file are counted separately as deepest sub tree A/B.

We used 40 features to describe deepest sub trees (Table 5). Some of the features include different percentiles of a characteristic, e.g., per le read sequentiality as percentile of all les in a directory. Including different percentiles rather than just the average allows better separation of access patterns. The corporate trace has 117,640 deepest sub trees, and the engineering trace has 161,858.

In Table 5.5, we provide quantitative descriptions and short names for all the deepest

(a). Descriptive features for each subtree

Number of hours with 1, 2-3, or 4 le opens

Number of hours with 1-100KB, 100KB-1MB, or >1MB reads
 Number of hours with 1-100KB, 100KB-1MB, or >1MB writes
 Number of hours with 1, 2-3, or 4 metadata requests
 Read request size - 25th, 50th, and 75th percentile of all requests
 Write request size - 25th, 50th, and 75th percentile of all requests
 Avg. time between I/O requests - 25th, 50th, and 75th percentile of all request pairs
 Read sequentiality - 25th, 50th, and 75th percentile of files in the subtree, and aggregated across all files
 Write sequentiality - 25th, 50th, and 75th percentile of files in the subtree, and aggregated across all files
 Read:write ratio - 25th, 50th, and 75th percentile of files, and aggregated across all files
 Repeated read ratio - 25th, 50th, and 75th percentile of files, and aggregated across all files
 Overwrite ratio - 25th, 50th, and 75th percentile of files, and aggregated across all files

(b). Corp. subtree access patterns	Temp real data	Client cacheable	Mixed read	Meta-data only	Mixed write	Small random read
% of all subtrees	2.3%	4.1%	5.6%	64%	3.5%	21%
# hrs with opens	3hrs	3hrs	2hrs	2hrs	1hr	1hr
Opens per hr	>4 opens	1 open	1 open	1 open	>4 opens	>4 opens
# hrs with reads	3hrs	2hrs	1hr	0	0	1hr
Reads per hr	1-100KB	1-100KB	1-100KB	-	-	1-100KB
# hrs with writes	2hrs	0	0	0	1hr	0
Writes per hr	0.1-1MB	-	-	-	>1MB	-
Read request size	4KB	4-10KB	4-32KB	-	-	1-8KB
Write request size	4KB	-	-	-	64KB	-
Read sequentiality	10-30%	0%	50-70%	-	-	0%
Write sequentiality	50-70%	-	-	-	70-80%	-
Repeat read ratio	20-50%	50%	0%	-	-	0%
Overwrite ratio	30-70%	-	-	-	0%	-
Read:write ratio	1:0 to 0:1	1:0	1:0	0:0	0:1	1:0

(b). Eng. subtree access patterns	Meta-data only	Small random read	Client cacheable	Mixed read	Sequential write	Temp real data
% of all subtrees	59%	25%	6.1%	7.1%	1.9%	1.3%
# hrs with opens	1hr	1hr	1hr	1hr	1hr	3hrs
Opens per hr	2-3 pens	>4 opens	>4 opens	>4 opens	>4 opens	>4 opens
# hrs with reads	0	1hr	1hr	1hr	0	3hrs
Reads per hr	-	1-100KB	1-100KB	0.1-1MB	-	1-100KB
# hrs with writes	0	0	0	0	1hr	1hr
Writes per hr	-	-	-	-	0.1-1MB	1-100KB
Read request size	-	1-4KB	2-4KB	8-10KB	-	4-32KB
Write request size	-	-	-	-	32-60KB	4-60KB
Read sequentiality	-	0%	0%	40-70%	-	10-65%
Write sequentiality	-	-	-	-	70-90%	60-80%
Repeat read ratio	-	0%	50-60%	0%	-	0-40%
Overwrite ratio	-	-	-	-	0%	0-30%
Read:write ratio	0:0	1:0	1:0	1:0	0:1	1:0 to 0:1

Table 5: Deepest sub tree access patterns. (a): Full list of descriptive features. Compared with the features list in Table 4, the changes are that read sequentiality, write sequentiality, read: write ratio, repeated read ratio, and overwrite ratio changed from a single statistic for a file to 25th, 50th, and 75th percentile of all files aggregated in a directory. (b) and (c): Short names and descriptions of sub trees in each access pattern; listing only the features that help separate access patterns

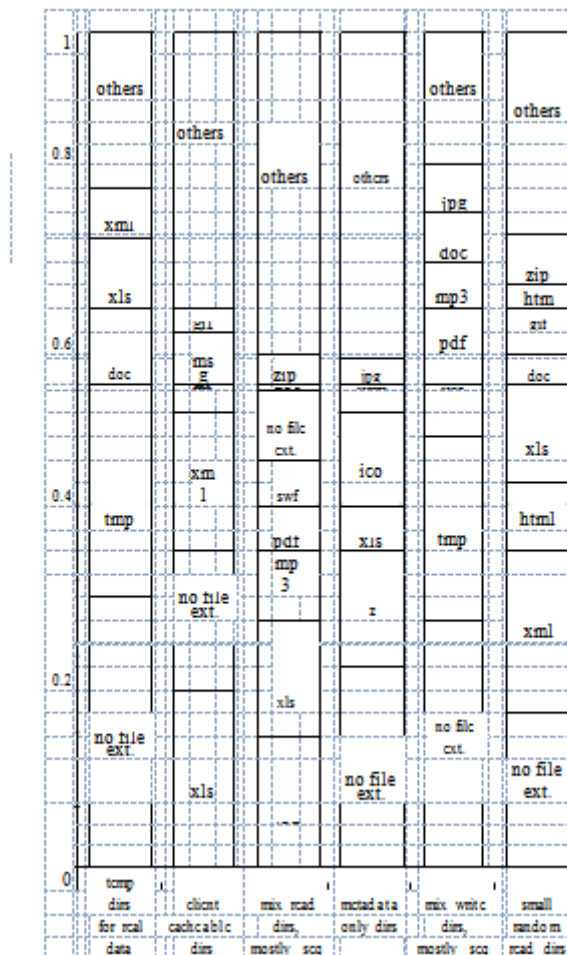


Figure 8: File extensions for corporate deepest sub trees. Fraction of file extensions in deepest subtree access patterns

For instance, the random read" and client cacheable" marks originate from taking a gander at the I/O designs. Temporary indexes" represented the .tmp files in those directories. Mix read" and \mix express" catalogs considered the nearness of both successive and arbitrarily got to files in those indexes.

The metadata foundation clamor stays unmistakable at the sub-tree layer. The spread of file augmentations is like that for file get to designs {some file expansions assemble and others spread equitably. Strangely, some sub trees have an expansive part of files for which just the metadata gets got to.

Some sub trees contain just files of a solitary access design (e.g., little arbitrary read sub trees in Figures 12). There, we can apply the outline bits of knowledge from the file get to examples to the whole sub tree. For instance, the little arbitrary read sub trees can dwell on SSDs. Since there are a larger number of files than sub trees, per-sub tree strategies can bring down the measure of approach data kept at the server.[15]

Interestingly, the blend read and blend compose registries contain both consecutive and randomly got to files. Those sub trees require per-file approaches: Place the consecutively got to files on HDDs and the haphazardly got to files on SSDs. Delicate connects to files can save the client confronting index association, while enabling the server to advance per-file nutriment. The server ought to naturally choose when to apply per-file or per-sub tree approaches.

Perception 11: Directories with successively got to files quite often contain haphazardly got to files likewise (Figures 12 and 13). On the other hand, a few catalogs with arbitrarily got to files won't contain consecutively got to files. In this manner, we can default all sub trees to per-sub tree arrangements. Simultaneously, we track the I/O sequentiality per sub tree. On the off chance that the sequentiality is over some edge, at that point the sub tree changes to per-file arrangements. Suggestion 11: Servers can change from per-registry arrangement strategy (default) to per-file approach after observing any consecutive I/O to any files in an index. with rehashed peruses or overwrites (Table 5 and Figures 12 and 13). Extra calculation demonstrated that the rehashed peruses and overwrites quite often originate from a solitary customer. In this way, it is workable for the whole catalog to be pre-brought and dele-gated to the customer. Designating whole registries can pre-

endeavour all gets to that are nearby to a catalog, however expands customer reserve space. We have to comprehend the tradeoffs through a more inside and out working set and transient area investigation at both the leaf and most profound sub tree levels. Suggestion 12: Servers can designate rehashed read and overwrite catalogs totally to customers, tradeoffs following.

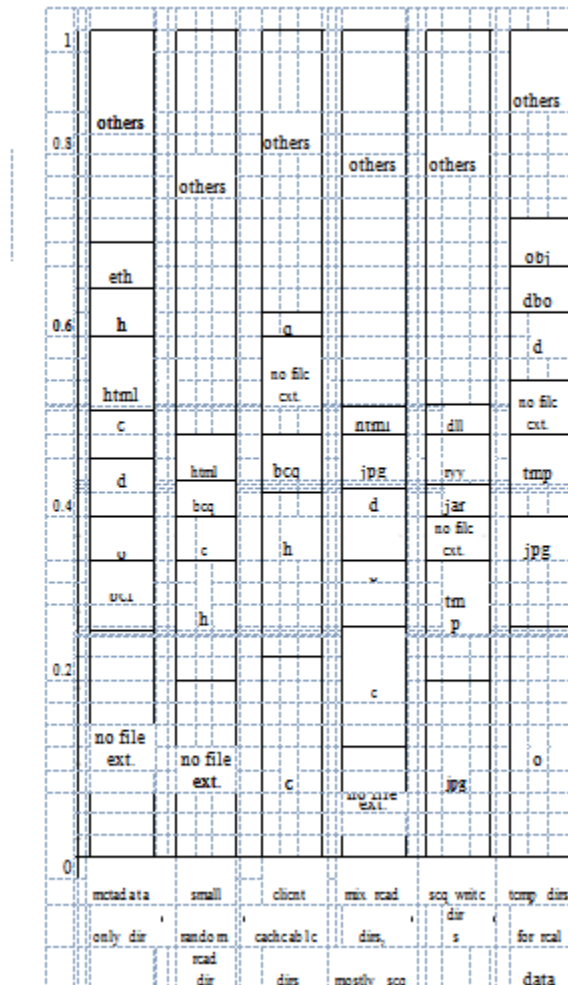


Figure 9: File extensions for engineering deepest sub trees. Fraction of leaf extensions in deepest sub tree access patterns

Architectural Trends

It offered many specific optimizations for placement, caching, delegation, and consolidation decisions. We combine the insights here to speculate on the architecture of future enterprise storage systems.[18]

We see a clear separation of roles for clients and servers. The client design can target high I/O performance by a combination of efficient delegation, pre-fetching and caching of the appropriate data. The servers should focus on increasing their aggregated efficiency across clients: collaboration with clients (on caching, delegation, etc.) and exploiting user patterns to schedule background tasks. Automating background tasks such as online data de-duplication delivers capacity savings in a timely and hassle-free fashion, i.e., without system downtime or explicit scheduling. Regarding caching at the server,[19] we observe that very few access patterns suggest how to improve server's buffer cache for data accesses. Design insights 4-6, 8 and 12 indicate a heavy role for the client cache and Design insight 7 suggests how not to use the server buffer cache - caching metadata only and acting as a warm/backup cache for clients would result in lower latencies for many access patterns.

We also see simple ways to take advantage of new storage media such as SSDs. The clear identification of sequential and random access leaf patterns enables efficient device-specific data placement algorithms (Design insights 10 and 11). Also, the background metadata noise seen at all levels suggests that storage servers should both optimize for metadata accesses and redesign client-server interactions to decrease the metadata chatter.

Depending on the growth of metadata and the performance requirements, we also need to consider placing metadata on low latency, non-volatile media like ash or SSDs.

Furthermore, we believe that storage systems should introduce many monitoring points to dynamically adjust the decision thresholds of placement, caching, or consolidation policies. We need to monitor both clients and servers. For example, when repeated read and overwrite les have been properly delegated to clients, the server would no longer see les with such access patterns.[20] Without monitoring points at the clients, we would not be able to quantify the le delegation benefits. Storage systems should make extensible tracing APIs to expedite the collection of long-term future traces. This will facilitate future work similar to ours.

III. RESULTS & CONCLUSION

We must address the storage technology trends toward ever-increasing scale, heterogeneity, and consolidation. Current storage design paradigms that rely on existing trace analysis methods are ill equipped to meet the emerging challenges because they are uni dimensional, focus only on the storage server, and are subject to designer bias. We showed that a multi-dimensional, multi-layered trace-driven design methodology leads to more objective design points with highly targeted optimizations at both storage clients and servers. Using our corporate and engineering use cases, we present a number of insights that inform future designs. We described in some detail the access

patterns we observed, and we encourage fellow storage system designers to extract further insights from our observations.

Storage system designers face an increasing challenge to anticipate access patterns. This chapter builds the case that system designers can no longer accurately anticipate access patterns using intuition only. We believe that the corporate and engineering traces from the NetApp corporate headquarters would have similar use cases at other traditional and high-tech businesses. Other use cases would require us to perform the same trace collection and analysis process to extract the same kind of empirical insights. We also need similar studies at regular intervals to track the evolving use of storage system. We hope that this chapter contributes to an objective and principled design approach targeting rapidly changing data access patterns.

In the broader context of the dissertation, the analysis here and in applies the methodology presented into helps us understand the behaviour of different kinds of large-scale data-centric systems. The subsequent chapters leverage the workload insights developed thus far to solve some system design problems beyond those that immediately follow from the workload behaviour observations. Specifically, seeks to improve Map Reduce energy efficiency. quantities the performance implications of TCP in cast. While these two design problems are not immediately related to workload management, an understanding of workload behaviour is vital to developing a solution.

REFERENCES:

- [1]. ApacheHadoopDocumentation.http://hadoop.apache.org/common/docs/r0.20.2/cluster_setup.html#Configuring+the+Hadoop+Daemons.
- [2]. Apache Oozie(TM) Work flow Scheduler for Hadoop. <http://incubator.apache.org/oozie/.Grid mix. HADOOP-HOME/src/benchmarks/grid mix in all recent Hadoop distributions>.
- [3]. Gridmix3. HADOOP-HOME/map red/src/contrib/grid mix in Hadoop 0.21.0 on-wards.
- [4]. Hadoop. <http://hadoop.apache.org.Hadoop Power-By Page. http://wiki.apache.org/hadoop/PoweredBy.Hadoop World 2011 Speakers. http://www.hadoopworld.com/speakers/>.
- [5]. HDFS Architecture Guide. http://hadoop.apache.org/hdfs/docs/current/hdfs_design.html.IEEE 802.1Qau Standard - Congestion Notification. http://www.ieee802.org/1/pages/802.1au.html.
- [6]. Mumak. <http://issues.apache.org/jira/browse/MAPREDUCE-728>, last retrieved Nov. 2009.Open Flow. <http://www.openflow.org/>.
- [7]. Personal communications with Cloudera engineering and support teams.Sort benchmark home page. <http://sortbenchmark.org/tcpdump>. <http://www.tcpdump.org/>.
- [8]. N. Agrawal, W. J. Bolosky, J. R. Douceur, and J. R. Lorch. A Five-Year Study of File-System Metadata. In FAST 2007.
- [9]. M. Al-Fares, A. Loukissas, and A. Vahdat. A scalable, commodity data centre network architecture. In SIGCOMM 2008.
- [10]. M. Alizadeh, A. Greenberg, D. A. Maltz, J. Padhye, P. Patel, B. Prabhakar, S. Sengupta, and M. Sridharan. Data centre TCP (DCTCP). In SIGCOMM 2010.
- [11]. M. Alizadeh et al. Data centre transport mechanisms: Congestion control theory and IEEE standardization. In Annual Allerton Conference 2008.
- [12]. M. Allman, H. Balakrishnan, and S. Floyd. Enhancing TCP's Loss Recovery Using Limited Transmit. <http://www.ietf.org/rfc/rfc3042.txt>, 2001.
- [13]. M. Allman, V. Paxson, and W. Stevens. Request for Comments: 2581 - TCP Congestion Control. <http://www.ietf.org/rfc/rfc2581.txt>, 1999.

Ch.Dhanunjaya Rao "Large-Scale Data-Centric Systems – A Research Analysis "American Journal of Engineering Research (AJER), vol. 7, no. 09, 2018, pp. 96-113