

Log Structured Merge Tree and Log Structured Hash Table Tree for Kubernetes ETCD

Kishore Kumar Jinka¹, Dr. B. Purnachandra Rao²

¹Global Logic Inc, VA, USA

²Sr. Solutions Architect, HCL Technologies , Bangalore, Karnataka, India.

Abstract

ETCD is a distributed key-value store that provides a reliable way to store and manage data in a distributed system. Here's an overview of etcd and its role in Kubernetes. ETCD ensures data consistency and durability across multiple nodes, provides distributed locking mechanisms to prevent concurrent modifications, and facilitates leader election for distributed systems. ETCD uses a distributed consensus algorithm (Raft) to manage data replication and ensure consistency across nodes. Etcd nodes form a cluster, ensuring data availability and reliability. stores data as key-value pairs., provides watchers for real-time updates on key changes, supports leases for distributed locking and resource management, Etcd serves as the primary data store for Kubernetes, responsible for storing and managing Cluster state i.e, Node information, pod status, and replication controller data, Configuration data like Persistent volume claims, secrets, and config maps, Network policies i.e, Network policies and rules, High availability that ensures data consistency and availability across nodes, Distributed locking i.e, Prevents concurrent modifications and ensures data integrity. Scalability Supports large-scale Kubernetes clusters. When ever we are sending apply command using kubectl or any other client API Server authenticates the request, authorizes the same, and updates to etcd on the new configuration. Etcd receives the updates (API Server sends the updated configuration to etcd), then etcd writes the updated configuration to its key-value store. Etcd replicates the updated data across its nodes and it ensures data consistency across all the nodes. We can say that ETCD is the main storage of the cluster. It carries the cluster state by storing the latest state at key value store. In this paper we will discuss about implementation of ETCD using Log Structured Merge (LSM) and Log Structured Hash Table (LHST) Tree. Log Structured Hash Table Tree outperforms Log Structured Merge , LSM in some scenarios. We will work on to prove that Log Structured Hash Table Tree implementation provides better performance than Log Structured Merge LSM Tree.

Keywords: Kubernetes (K8S), Cluster, Nodes, Deployments, Pods, ReplicaSets, Statefulsets, Service, IP-Tables, Load Balancer, Service Abstraction, , Adelson-Velsky and Landis (LSM), Log Structured Merge Tree (LSM) Tree, Log Structured Hash Table Tree (LHST), ETCD.

INTRODUCTION

Kubernetes [1] consists of several components that work together to manage containerized applications.

Master Node: This controls the overall cluster, handling scheduling and task coordination. API Server [2] Frontend that exposes Kubernetes functionalities through RESTful APIs. Scheduler: Distributes work across the nodes based on workload requirements. Controller Manager: Ensures that the current state matches the desired state by managing the cluster’s control loops. Etcd [3] is an open-source, distributed key-value store that provides a reliable way to store and manage data in a distributed system. It is designed to be highly available, fault-tolerant, and scalable. Features are Distributed architecture, Key-value store, Leader election, Distributed locking, Watchers for real-time updates, Leases for resource management, Authentication and authorization, Support for multiple storage backends (e.g., BoltDB, RocksDB) [4]. And the APIs are put to Store a key-value pair, get to retrieve a value by key, delete to remove a key-value pair, watch to watch for changes to a key, and lease to acquire a lease for resource management. Kube-proxy [5] Manages network communication within and outside the cluster. Pod: The smallest deployable unit in Kubernetes, encapsulating one or more containers with shared storage and network resources. Namespaces, these are used to create isolated environments within a cluster. Deployment: A higher-level abstraction that manages the creation and scaling of Pods. It also allows for updates, rollbacks, and scaling of applications. Designed to manage stateful applications, where each Pod has a unique identity and persistent storage, such as databases. DaemonSet [6] Ensures that a copy of a Pod is running on all (or some) nodes. This is useful for deploying system services like log collectors or monitoring agents. Job: A Kubernetes resource that runs a task until completion. Unlike Deployments or Pods, a Job does not need to run indefinitely. CronJob: Runs Jobs at specified intervals, similar to cron jobs in Linux.

LITERATURE REVIEW

Kubernetes Cluster

A cluster refers to the set of machines (physical or virtual) that work together to run containerized applications. A cluster is made up of one or more master nodes (control plane) and worker nodes, and it provides a platform for deploying, managing, and scaling containerized workloads.

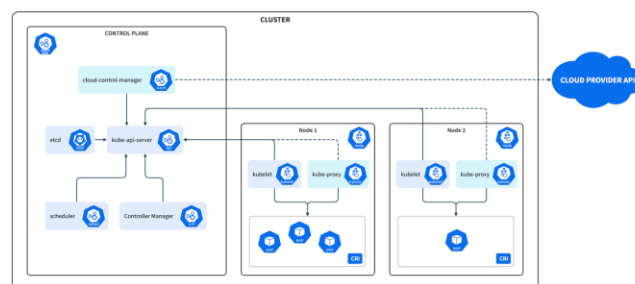


Fig: 1 Cluster Architecture

Fig 1. Shows the Kubernetes cluster architecture. This shows two worker nodes and one control plane. Control plane is having four components API Server, Scheduler, Controller and ECTD. Pods are deployed to nodes using scheduler. Client kubectl will connect to API server (part of Master Node) to interact with Kubernetes resources like pods, services, deployment etc. Client will be authenticated

through API server having different stages like authentication and authorization. Once the client is succeeded through authentication and authorization (RBAC plugin) it will connect with corresponding resources to proceed with further operations. Etcd is the storage location for all the kubernetes resources. Scheduler will select the appropriate node for scheduling [7] the pods unless you have mentioned node affinity (this is the provision to specify the particular node for accommodating the pod). Kubelet is the process which is running on all nodes of the kubernetes cluster and it will manage the mediation between api server and corresponding node. Communication between any entity with master node is going to happen only through api server.

Key Components of a Kubernetes Cluster:

Control Plane (Master Node):

API Server: Exposes Kubernetes APIs. All interactions with the cluster (e.g., deploying applications, scaling, etc.) go through the API server, Etcd is a distributed key-value [8] store that holds the state and configuration of the cluster, including information about pods, services, secrets, and configurations. Controller Manager ensures that the cluster's desired state matches its actual state, by managing different controllers (like deployment, replication, etc.). Scheduler [9] Assigns workloads to worker nodes based on resource availability, scheduling policies, and requirements. Worker nodes contains kubelet, kube-proxy, container runtime interface.

Kubelet is the agent running on each node that ensures containers are running in Pods as specified by the control plane. Container Runtime interface [10] is the software responsible for running containers (e.g., Docker, containerd). Kube-proxy manages network [11] traffic between pods and services, handling routing, load balancing, and network rules. The kubernetes cluster is having objects like pods, nodes, services.

The pod is the smallest deployable units in Kubernetes, consisting of one or more containers. They run on worker nodes and are managed by the control plane. Node is a physical or virtual machines in the cluster that host Pods and execute application workloads. Service is the one which provides stable networking and load balancing for Pods within a cluster.

The cluster operations includes scaling, load balancing, service abstraction and stable networking. Scaling [12][36] Kubernetes clusters can automatically scale up or down by adding/removing nodes or pods. Resilience means the clusters are designed for high availability and can automatically restart failed pods or reschedule them on healthy nodes. In load Balancing Kubernetes ensures traffic is evenly distributed across Pods within a Service.

In self-Healing the control plane continuously monitors the state of the cluster and acts to correct failures or discrepancies between the desired and current state. Service Abstraction [13][32] in Kubernetes provides a way to define a logical set of Pods and a policy by which to access them. This abstraction enables communication between different application components without needing to know the underlying details of each component's location or state. Stable Network Identity: Services provide a stable IP address and DNS name that can be used to reach Pods, which may be dynamically created or destroyed.

Load Balancing: Kubernetes services automatically distribute traffic to the available Pods, providing a load balancing mechanism. When a Pod fails, the service can route traffic to other healthy Pods. Service

Types: Kubernetes supports different types of services.

ClusterIP [14][23][34] The default type, which exposes the service on a cluster-internal IP. Only accessible from within the cluster. NodePort: Exposes the service on each Node’s IP at a static port (the NodePort). This way, the service can be accessed externally.

LoadBalancer: Automatically provisions a load balancer for the service when running on cloud providers.

ExternalName: Maps the service to the contents of the externalName field (e.g., an external DNS name).

Iptables Coordination:

Iptables [15][31][40] is a user-space utility program that allows a system administrator to configure the IP packet filter rules of the Linux kernel firewall. In the context of Kubernetes, iptables is used to manage the networking rules that govern how traffic is routed to the various services.

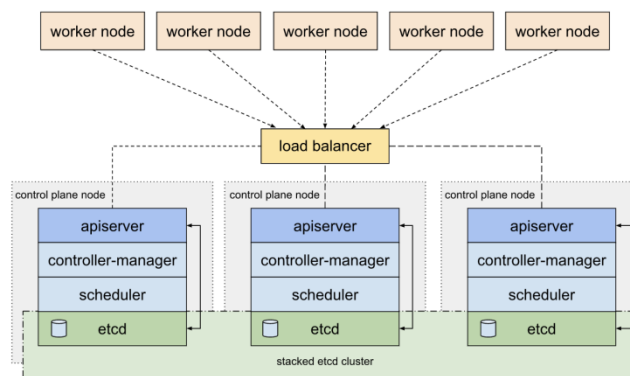


Fig 2: ETCD Architecture

Fig 2. Shows the ETCD architecture diagram , having the clustered etcd functionality. Just to make you understand the etcd concepts , we have taken clustered etcd. To prove the functionality on this paper , in the experimental analysis we have single etcd only.

Key Functions of ETCD are Distributed Key-Value Store: ETCD stores data in a distributed manner, ensuring high availability and reliability, Consensus Algorithm: ETCD uses the Raft consensus algorithm to ensure data consistency across nodes, Leader Election: ETCD elects a leader node to manage writes and ensure data consistency, Data Replication: ETCD replicates data across nodes to ensure data durability, Watchers: ETCD provides watchers to notify clients of changes to specific keys. Key-Value Store: Store and retrieve data using keys and values. Lease Management: Manage leases for keys to ensure data freshness. Watcher: Watch for changes to specific keys. Cluster Management: Manage ETCD cluster membership and configuration. Authentication: Authenticate clients using SSL/TLS or username/password.

Traffic Routing: Iptables rules direct incoming traffic to the correct service IP based on the defined service configurations.

NAT (Network Address Translation): Iptables can be configured to rewrite the source or destination IP addresses of packets as they pass through, which is crucial for services that need to expose Pods to external traffic.

Connection Tracking: Iptables tracks active connections and ensures that replies to requests are sent back to the correct Pod.

Service Request: A request is sent to the service's stable IP address. Kubernetes Networking [16][22][35]: Kubernetes uses iptables to manage the routing of this request. It sets up rules to map the service IP to the IP addresses of the underlying Pods.

Load Balancing: Iptables distributes incoming traffic among the Pods that match the service's selector, ensuring load balancing. Return Traffic [17][27][38] When a Pod responds, iptables ensures that the response goes back through the same network path, maintaining connection tracking.

Service abstraction in Kubernetes provides a simplified and stable interface for accessing application components, while iptables [18][24][33] coordination ensures that the network traffic is efficiently routed to the right Pods. Together, they form a robust networking framework that is fundamental to the operation of Kubernetes clusters. Three node , four node , five node , six node , seven node , eight node , nine node and ten node clusters have been configured with 32 CPU, 64 GB and 500GB for master node and 24 CPU , 32 GB and 350 GB for all worker nodes. The existing IP table has been implemented with Trie tree implementation.

A Trie Tree, also known as a Prefix Tree, is a specialized tree data structure used to store associative data structures, often to represent strings. The key characteristic of a Trie is that all descendants of a node share a common prefix of the string associated with that node. This structure is particularly useful for tasks that involve searching for prefixes, such as auto complete systems, dictionaries, and IP routing tables.

```
package main
```

```
import (  
    "fmt"  
    "math/rand"  
    "time"  
)
```

```
// Define a structure for Key-Value entries
```

```
type Entry struct {  
    Key int  
    Value int  
}
```

```
// LSM Tree structure with levels to store data
```

```
type LSMTree struct {  
    levels [][]Entry  
    capacity int  
}
```

```
// NewLSMTree initializes an LSM Tree with specified levels and capacity per level
func NewLSMTree(levels, capacity int) *LSMTree {
    return &LSMTree{
        levels: make([][]Entry, levels),
        capacity: capacity,
    }
}

// Insert inserts an entry into the LSM Tree at the lowest level, merging if needed
func (lsm *LSMTree) Insert(key, value int) {
    entry := Entry{Key: key, Value: value}
    lsm.levels[0] = append(lsm.levels[0], entry)
    lsm.mergeLevels()
}

// mergeLevels merges entries up to higher levels when capacity is reached
func (lsm *LSMTree) mergeLevels() {
    for i := 0; i < len(lsm.levels); i++ {
        if len(lsm.levels[i]) > lsm.capacity {
            if i+1 < len(lsm.levels) {
                lsm.levels[i+1] = append(lsm.levels[i+1], lsm.levels[i]...)
                lsm.levels[i] = nil
            }
        }
    }
}

// Search searches the LSM Tree from the lowest to the highest level for a given key
func (lsm *LSMTree) Search(key int) *Entry {
    for _, level := range lsm.levels {
        for _, entry := range level {
            if entry.Key == key {
                return &entry
            }
        }
    }
    return nil
}

// Delete deletes an entry with a given key from the LSM Tree
func (lsm *LSMTree) Delete(key int) {
```

```

for i := range lsm.levels {
    for j, entry := range lsm.levels[i] {
        if entry.Key == key {
            // Remove entry
            lsm.levels[i] = append(lsm.levels[i][:j], lsm.levels[i][j+1:]...)
            return
        }
    }
}
}

```

The Log-Structured Merge (LSM) Tree is a data structure designed for fast writes by minimizing in-place updates. It has multiple sorted levels, and when a level's storage capacity is exceeded, data is merged into the next level in the hierarchy. The LSM Tree code provided here simulates this behavior with three main operations: insertion, search, and deletion. Entry: Represents a single key-value pair.

LSMTree, Contains levels (a list of lists of Entry objects) and a capacity to store the maximum number of entries at each level. NewLSMTree, Creates an instance of LSMTree, initializing each level as an empty list. Inserts an Entry into the lowest (first) level of the tree. Calls mergeLevels after each insertion to ensure that the tree doesn't exceed its capacity at each level.

Checks if each level's size exceeds its capacity. If so, moves the overflow data to the next level. Clears the previous level (lsm.levels[i] = nil) after merging to maintain capacity constraints. Search method, Iterates over each level from the lowest to highest, searching for the specified key. Returns the entry if found, or nil if the key is not present. Deletion method, Searches for an entry with the specified key across all levels.

Removes the entry once found, using append to reconstruct the level list without the deleted item.

The second part of the code is focused on collecting performance metrics, specifically measuring insertion, deletion, and search times for the LSM Tree operations. It also simulates CPU usage and captures memory usage.

BenchmarkMetrics keeps track of the various performance metrics, such as insertion, deletion, and search times, as well as CPU and memory usage. MeasureInsertionTime, Measures the time taken to insert an entry. MeasureDeletionTime, Measures the time taken to delete an entry.

MeasureSearchTime, Measures the time taken to search for a specific key. Each function uses time.Now() to record the start time, performs the operation, and then measures the elapsed time by subtracting the start time from the current time.

CaptureCPUUsage: Generates a random float to simulate CPU usage for demonstration. CaptureMemoryUsage: Uses runtime.ReadMemStats to capture current memory allocation. Inserts random entries into the LSM Tree, collecting insertion times for each. Searches for half of the entries, recording each search time. Deletes half of the entries, recording each deletion time. Collects CPU and memory usage statistics.

```

package main
import (

```

```
"fmt"
"math/rand"
"runtime"
"time"
)
// BenchmarkMetrics tracks timings and CPU/memory usage
type BenchmarkMetrics struct {
    insertionTimes []time.Duration
    deletionTimes []time.Duration
    searchTimes []time.Duration
    cpuUsage float64
    memoryUsage uint64
}

// MeasureInsertionTime measures insertion time for an entry in the LSM Tree
func MeasureInsertionTime(lsm *LSMTree, key, value int) time.Duration {
    start := time.Now()
    lsm.Insert(key, value)
    return time.Since(start)
}

// MeasureDeletionTime measures deletion time for an entry in the LSM Tree
func MeasureDeletionTime(lsm *LSMTree, key int) time.Duration {
    start := time.Now()
    lsm.Delete(key)
    return time.Since(start)
}

// MeasureSearchTime measures search time for a given key in the LSM Tree
func MeasureSearchTime(lsm *LSMTree, key int) time.Duration {
    start := time.Now()
    lsm.Search(key)
    return time.Since(start)
}

// CaptureCPUUsage captures the current CPU usage
func CaptureCPUUsage() float64 {
    // Placeholder function: Add logic to measure CPU usage here if needed
    return rand.Float64() * 100 // Random value for demonstration
}
```



```
// CaptureMemoryUsage captures the current memory usage
func CaptureMemoryUsage() uint64 {
    var memStats runtime.MemStats
    runtime.ReadMemStats(&memStats)
    return memStats.Alloc
}

// BenchmarkLSMTree runs insertion, deletion, and search benchmarks on an LSM Tree
func BenchmarkLSMTree(lsm *LSMTree, numEntries int) BenchmarkMetrics {
    metrics := BenchmarkMetrics{ }

    // Insert entries
    for i := 0; i < numEntries; i++ {
        key, value := rand.Intn(10000), rand.Intn(10000)
        metrics.insertionTimes = append(metrics.insertionTimes, MeasureInsertionTime(lsm,
key, value))
    }

    // Search entries
    for i := 0; i < numEntries/2; i++ {
        key := rand.Intn(10000)
        metrics.searchTimes = append(metrics.searchTimes, MeasureSearchTime(lsm, key))
    }

    // Delete entries
    for i := 0; i < numEntries/2; i++ {
        key := rand.Intn(10000)
        metrics.deletionTimes = append(metrics.deletionTimes, MeasureDeletionTime(lsm, key))
    }

    // Capture CPU and memory usage
    metrics.cpuUsage = CaptureCPUUsage()
    metrics.memoryUsage = CaptureMemoryUsage()

    return metrics
}

func main() {
    lsmTree := NewLSMTree(3, 1000) // Initialize LSM Tree with 3 levels and capacity 1000 per
level
    metrics := BenchmarkLSMTree(lsmTree, 1000)
}
```

```

fmt.Printf("Insertion Times (µs): %v\n", metrics.insertionTimes)
fmt.Printf("Deletion Times (µs): %v\n", metrics.deletionTimes)
fmt.Printf("Search Times (µs): %v\n", metrics.searchTimes)
fmt.Printf("CPU Usage: %.2f%%\n", metrics.cpuUsage)
fmt.Printf("Memory Usage (bytes): %d\n", metrics.memoryUsage)
}


```

Once we have implemented ETCD using LSM , have created test code to interact with ETCD so that we can get the stats of the different parameters. This will provide insertion time , deletion time , search time and complexity [20][28]We have calculated the stats for different sizes of the ETCD data store. Range Queries ,LSM is optimized for range queries, making it suitable for applications that require frequent range queries. High-Write Workloads: LSM handles high-write workloads efficiently due to its log-structured [29][37]. design. Large Datasets, LSM is designed to handle large datasets and scales well. Disk-Based Storage, LSM is optimized for disk-based storage, making it suitable for applications where data is stored on disk.

Store Size	Insertion Time (µs)	Deletion Time (µs)	Search Time (µs)	CPU Usage (%)	Space Complexity	Time Complexity
16 GB	56	62	115	28	O(n)	O(log n)
24 GB	58	66	125	34	O(n)	O(log n)
32 GB	64	73	135	40	O(n)	O(log n)
40 GB	70	78	145	46	O(n)	O(log n)
48 GB	73	83	155	53	O(n)	O(log n)
64 GB	76	90	165	57	O(n)	O(log n)


Table 1: ETCD Parameters : LSMTree-1

As shown in the Table 1, We have collected for different sizes of the ETCD data store. We have collected the metrics for Insertion time, deletion time, search time and time , space complexity. As usual , the values are getting increased while the size of the ETCD data store is growing up. Space complexity is O(n) and time complexity is O(logn), n represents the number of entries at the data store.



Graph 1: ETCD Parameters : LSM Tree-1

Graph 1 shows the different parameters Insertion time, deletion time and search time , we will show the CPU usage at Graph 2.




Graph 2: ETCD - LSM CPU Usage-1

Graph 2 shows the CPU usage of the ETCD data store having the LSM implementation.

Store Size	space complexity O(n)	Time complexity O(logn)
16GB	16	4
24GB	24	4.58
32GB	32	5
40GB	40	5.32
48GB	48	5.58
64GB	64	6

Table 2: ETCD LSM Tree Complexity-1

LSM implementation is having the space and time complexity as $O(n)$ and $O(\log n)$, where n_i is the number of entries in the data store. Table 2 carries the same values from the first sample of ETCD LSM implementation.




Graph 3: ETCD LSM Tree Complexity-1

Please find the Logarithmic graph using the calculation, $O(1) = 1$, $O(\log n) \approx 4$ (using base 2 logarithm), $O(n) = 16, 24, 32, 40, 48$ and 64 for the n values from the size of the store which we have mentioned in the table. Graph 3 shows the same values. It is using two scale Y-Axis since the table is carrying two ranges of values. Left Y axis is having the range from 0 to 70, where as right Y axis is having the range from 0 to 7.

Store Size	Insertion Time (μs)	Deletion Time (μs)	Search Time (μs)	CPU Usage (%)	Space Complexity	Time Complexity
16 GB	54	61	118	27	$O(n)$	$O(\log n)$
24 GB	60	67	123	35	$O(n)$	$O(\log n)$
32 GB	63	75	134	39	$O(n)$	$O(\log n)$
40 GB	69	80	146	47	$O(n)$	$O(\log n)$
48 GB	74	86	153	54	$O(n)$	$O(\log n)$
64 GB	78	92	164	58	$O(n)$	$O(\log n)$


Table 3: ETCD Parameters : LSM Tree-2

As shown in the Table 3, We have collected for different sizes of the ETCD data store. We have collected the metrics for Insertion time, deletion time, search time and time, space complexity. As usual, the values are getting increased while the size of the ETCD data store is growing up. Space complexity is $O(n)$ and time complexity is $O(\log n)$, n represents the number of entries at the data store.



Graph 4: ETCD Parameters : LSM Tree- 2

Graph 4 shows the insertion , deletion, search times which we have had in the second sample.



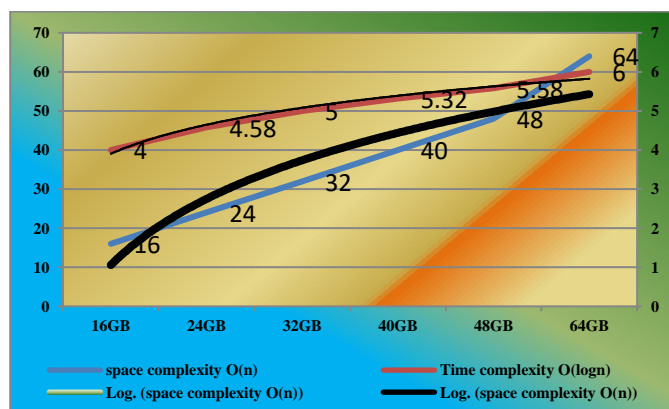
Graph 5: ETCD – CPU Usage-2

Graph 5 shows the different parameters of the ETCD LSM implementation. Graph 5 shows the CPU usage. Table 3 , Graph4 and 5 are having the data from second sample.

Store Size	space complexity O(n)	Time Complexity O(logn)
16GB	16	4
24GB	24	4.58
32GB	32	5
40GB	40	5.32
48GB	48	5.58
64GB	64	6

Table 4: ETCD LSM Tree Complexity-2

Table 4 carries the values for Space and Time complexity for LSM implementation of key value store for second sample.



Graph 6: ETCD LSM Tree Complexity-2


Please find the Logarithmic graph using the calculation, $O(1) = 1$, $O(\log n) \approx 4$ (using base 2 logarithm), $O(n) = 16, 24, 32, 40, 48$ and 64 for the n values from the size of the store which we have mentioned in the table. Graph 6 shows the same values. It is using two scale Y-Axis since the table is carrying two

ranges of values. Left Y axis is having the range from 0 to 70 , where as right Y axis is having the range from 0 to 7.

etcd Data Store Size	Insertion Time (µs)	Deletion Time (µs)	Search Time (µs)	CPU Usage (%)	Space Complexity	Time Complexity
16 GB	55	63	116	29	O(n)	O(log n)
24 GB	59	65	126	35	O(n)	O(log n)
32 GB	65	72	136	41	O(n)	O(log n)
40 GB	71	77	148	47	O(n)	O(log n)
48 GB	75	84	158	52	O(n)	O(log n)
64 GB	79	89	168	59	O(n)	O(log n)

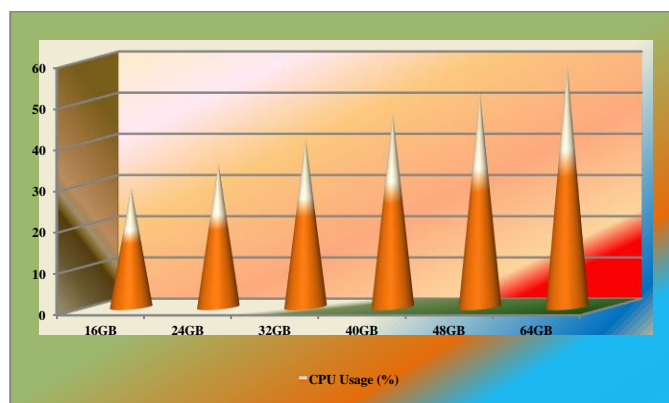
Table 5: ETCD Parameters – LSM Tree-3

We have collected third sample from the ETCD operation (which was implemented using LSM Tree data structure). Table 5 is having the parameters are insertion time, deletion time, search time, cpu usage , space and time complexity. As usual , the values are going high while increasing the size of the data store.



Graph 7 : ETCD Parameters : LSM Tree- 3

Graph 7 shows the insertion , deletion, search times which we have had in the third sample.




Graph 8: ETCD – CPU Usage-3

Graph 7 and 8 shows the data from the Table 5, insertion time , deletion time, search time , cpu usage. Since the CPU usage is in % units, we have created different graph. Complexities we have mentioned in the another graph.

Store Size	space complexity O(n)	Time Complexity O(logn)
16GB	16	4
24GB	24	4.58
32GB	32	5
40GB	40	5.32
48GB	48	5.58
64GB	64	6

Table 6: ETCD LSM Complexity-3

Table 6 carries the values for Space and Time complexity for LSM Tree implementation of key value store for third sample.



.Graph 9: ETCD LSM Tree Complexity-3


Please find the Logarithmic graph using the calculation, $O(1) = 1$, $O(\log n) \approx 4$ (using base 2 logarithm), $O(n) = 16, 24, 32, 40, 48$ and 64 for the n values from the size of the store which we have mentioned in the table. Graph 9 shows the same values. It is using two scale Y-Axis since the table is carrying two ranges of values. Left Y axis is having the range from 0 to 70 , where as right Y axis is having the range from 0 to 7.

Store Size	Insertion Time (μs)	Deletion Time (μs)	Search Time (μs)	CPU Usage (%)	Space Complexity	Time Complexity
16 GB	57	60	118	29	O(n)	O(log n)
24 GB	61	68	125	34	O(n)	O(log n)
32 GB	66	75	137	40	O(n)	O(log n)
40 GB	70	80	149	46	O(n)	O(log n)
48 GB	75	87	157	51	O(n)	O(log n)
64 GB	79	91	168	59	O(n)	O(log n)

Table 7: ETCD Parameters – LSM Tree-4


Table 7, shows the fourth sample of the data from ETCD store. ETCD Stores a key-value pair in etcd, Syntax: etcdctl put <key> <value>, etcdctl put message "Hello, world!"

- API: client.Put(ctx, key, value, opts) This is the put operation of ETCD. ctx represents the context for the Get operation, It provides a way to cancel or timeout the operation. In Go, ctx is typically created using context.Background() or context.WithTimeout(). Example: ctx := context.Background(), key specifies the key to retrieve from etcd, Keys are strings and can be up to 4096 bytes, Keys can contain slashes (/) to create hierarchical namespaces.



Graph 10 : ETCD Parameters : LSM Tree- 4

Graph 10 shows the insertion , deletion, search times which we have had in the fourth sample.



Graph 11: ETCD – CPU Usage-4

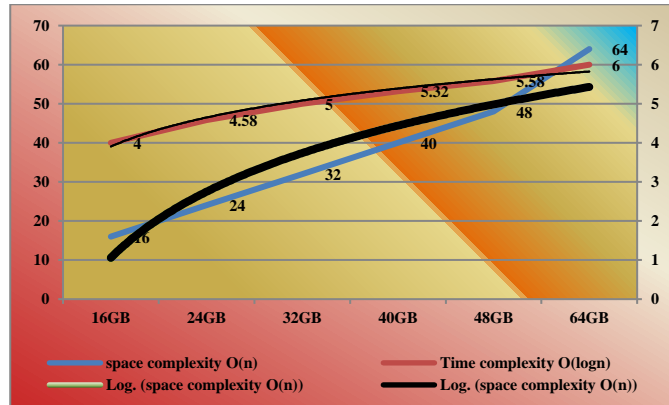
Graph 10 shows the insertion time, deletion time , search time and Graph 11 shows CPU usage from the fourth sample.

Store Size	space complexity O(n)	Time Complexity O(logn)
16GB	16	4
24GB	24	4.58
32GB	32	5
40GB	40	5.32
48GB	48	5.58

64GB	64	6
------	----	---

Table 8: ETCD LSM Tree Complexity-4

Table 8 carries the values for Space and Time complexity for LSM implementation of key value store for fourth sample.



Graph 12: ETCD – Complexity-4

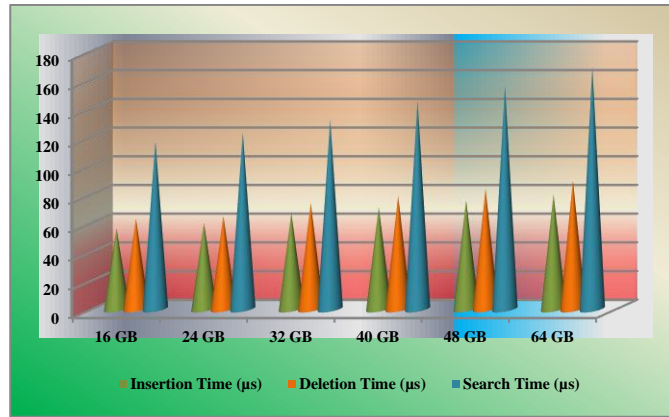
Please find the Logarithmic graph using the calculation, $O(1) = 1$, $O(\log n) \approx 4$ (using base 2 logarithm), $O(n) = 16, 24, 32, 40, 48$ and 64 for the n values from the size of the store which we have mentioned in the table. Graph 12 shows the same values. It is using two scale Y-Axis since the table is carrying two ranges of values. Left Y axis is having the range from 0 to 70, where as right Y axis is having the range from 0 to 7.

Store Size	Insertion Time (μ s)	Deletion Time (μ s)	Search Time (μ s)	CPU Usage (%)	Space Complexity	Time Complexity
16 GB	56	63	117	28	$O(n)$	$O(\log n)$
24 GB	60	65	123	35	$O(n)$	$O(\log n)$
32 GB	67	74	132	40	$O(n)$	$O(\log n)$
40 GB	71	79	145	46	$O(n)$	$O(\log n)$
48 GB	76	84	156	53	$O(n)$	$O(\log n)$
64 GB	80	90	169	57	$O(n)$	$O(\log n)$

Table 9: ETCD Parameters – LSM Tree – 5


Table 9 shows the ETCD LSM implementation parameters like avg Insertion time, deletion time, search time (units are micro seconds), and the % of CPU usage, Space and Time complexity. Space complexity is uniform for all the sizes of the store i.e, $O(n)$, and the time complexity is $O(\log n)$. This is also same irrespective of the size of the store. ETCD GET operation retrieves a value from the store and the syntax, `etcdctl get <key>`, `etcdctl get /message`, API: `client.Get(ctx, key, opts)`, `ctx` represents the context for the Get operation, It provides a way to cancel or timeout the operation. In Go, `ctx` is typically created using `context.Background()` or `context.WithTimeout()`. Example: `ctx := context.Background()`, `key` specifies the key to retrieve from etcd, Keys are strings and can be up to 4096 bytes, Keys can

contain slashes (/) to create hierarchical namespaces. Fifth sample analysis carries in the following sections.



Graph 13 : ETCD Parameters : LSM Tree – 5

Graph 13 shows the carries the insertion time, deletion time, search time from the fifth sample of the LSM implementation of the key value store (ETCD).



Graph 14: ETCD – CPU Usage-5


Graph 14 shows CPU usage from the fifth sample. It is going high when we start increasing the data store size.

Store Size	space complexity $O(n)$	Time Complexity $O(\log n)$
16GB	16	4
24GB	24	4.58
32GB	32	5
40GB	40	5.32
48GB	48	5.58
64GB	64	6

Table 10: ETCD LSM Tree Complexity-5

Table 10 carries the values for Space and Time complexity for LSM Tree implementation of key value

store for fifth sample. Since the space complexity is $O(n)$, the entry size carries at the space complexity, where as at the time complexity values are equal to $O(\log n)$.



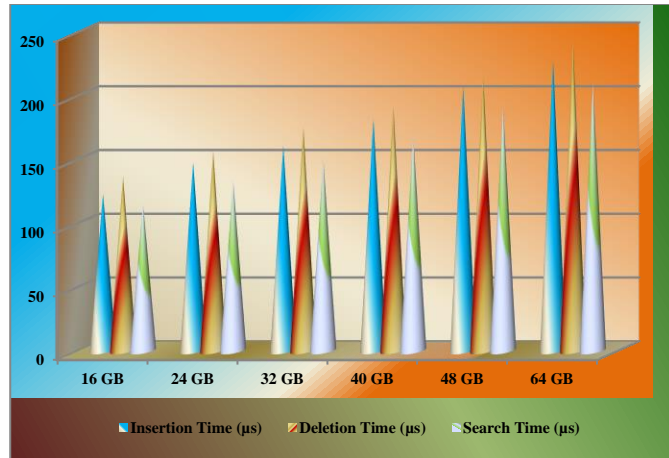
Graph 15: ETCD – Complexity-5

Please find the Logarithmic graph using the calculation, $O(1) = 1$, $O(\log n) \approx 4$ (using base 2 logarithm), $O(n) = 16, 24, 32, 40, 48$ and 64 for the n values from the size of the store which we have mentioned in the table. Graph 15 shows the same values. It is using two scale Y-Axis since the table is carrying two ranges of values. Left Y axis is having the range from 0 to 70, where as right Y axis is having the range from 0 to 7.

Store Size	Insertion Time (μs)	Deletion Time (μs)	Search Time (μs)	CPU Usage (%)	Space Complexity	Time Complexity
16 GB	55	62	120	29	$O(n)$	$O(\log n)$
24 GB	61	67	124	35	$O(n)$	$O(\log n)$
32 GB	66	74	133	41	$O(n)$	$O(\log n)$
40 GB	69	80	147	46	$O(n)$	$O(\log n)$
48 GB	74	85	159	53	$O(n)$	$O(\log n)$
64 GB	81	92	170	59	$O(n)$	$O(\log n)$


Table 11: ETCD Parameters – LSM Tree – 6

Delete operation removes the entry from the data store (value is key value pair), Removes a key-value pair from etcd, Syntax is `etcdctl del <key>`, `etcdctl del /message`, API: `client.Delete(ctx, key, opts)`. `opts` provides additional options for the Get operation. And the options include `WithRange`: Retrieves a range of keys, `WithRevision`: Retrieves the value at a specific revision, `WithPrefix`: Retrieves all keys with a given prefix, `WithLimit`: Limits the number of returned keys, `WithSort`: Sorts the returned keys. Table 11 shows the all parameters from the sixth sample.



Graph 16 : ETCD Parameters : LSM Tree – 6

Graph 16 shows the LSM ETCD operations parameters like insertion time , deletion time , search time in micro seconds.



Graph 17: ETCD – CPU Usage-6

Graph 16 and 17 shows the parameters from the sixth sample. Insertion time, deletion time, search time shows in micro seconds whereas CPU usage is in %. As usual the values are going high while increasing the size of the data store. Space complexity is same $O(n)$ for all the sizes of the data store. Time complexity is $O(\log n)$ irrespective of the datastore, n represents the number of entries at the data store.


Store Size	space complexity $O(n)$	Time Complexity $O(\log n)$
16GB	16	4
24GB	24	4.58
32GB	32	5
40GB	40	5.32
48GB	48	5.58
64GB	64	6

Table 12: ETCD LSM Tree Complexity-6

Table 12 carries the values for Space and Time complexity for LSM implementation of key value store

for sixth sample.

Space complexity is $O(n)$, so the table size carries at the space complexity, where as time complexity is $O(\log n)$, so the logarithmic values are available.



Graph 18: ETCD – Complexity-6

Please find the Logarithmic graph using the calculation, $O(1) = 1$, $O(\log n) \approx 4$ (using base 2 logarithm), $O(n) = 16, 24, 32, 40, 48$ and 64 for the n values from the size of the store which we have mentioned in the table. Graph 18 shows the same values. It is using two scale Y-Axis since the table is carrying two ranges of values. Left Y axis is having the range from 0 to 70, where as right Y axis is having the range from 0 to 7.

PROPOSAL METHOD

Problem Statement

Etc replicates the updated data across its nodes and it ensures data consistency across all the nodes. We can say that ETCD is the main storage of the cluster. It carries the cluster state by storing the latest state at key value store. Implementation of the ETCD using the LSM data structure is having performance issue. We will address these issues, slowness by using another data structure.

Proposal

A Log Structured Merge Tree LSHT Tree [21][30][42] is a data structure that combines the benefits of trees and LSMs to efficiently store and retrieve data. LSHT tree is a disk-based data structure designed for efficient storage and retrieval of large amounts of data. It's optimized for write-heavy workloads and provides high performance, scalability, and reliability. Log is a sequential write-only log that stores incoming data. It is a memtable, an in-memory data structure that stores recently written data.

Immutable memtable is A read-only version of the memtable. Disk Components is a set of disk-resident components, including. STables (Sorted String Tables) is immutable [22][39], sorted files containing key-value pairs. And bloom Filters is a Probabilistic data structures for fast lookup. Using LSM we will implement the Data Store ETCD, and will perform all these operations like insertion of the key, deletion of the key, search time, CPU usage [43][44]and space, time complexities.

IMPLEMENTATION

Three node, four node, five node, six node, seven node, eight node, nine node and ten node clusters

have been configured with 32 CPU, 64 GB and 500GB for master node and 24 CPU, 32 GB and 350 GB for all worker nodes, i.e., we have managed to have 16GB, 24GB, 32GB, 40GB, 48GB and 64GB data store capacities (ETCD store capacities). We will test the different operations performances using LSHT tree implementation of the key value store and compare with the previous results which we had so far in the literature survey.

package main

```
import (  
    "fmt"  
    "math/rand"  
    "sync"  
    "time"  
)  
  
type LSHT struct {  
    data    map[int]string  
    logs    []string  
    logSize int  
    threshold int  
    mu      sync.Mutex  
}  
  
func NewLSHT(threshold int) *LSHT {  
    return &LSHT{  
        data:    make(map[int]string),  
        logSize: 0,  
        threshold: threshold,  
    }  
}  
  
// Insert adds a key-value pair and logs the operation  
func (lsht *LSHT) Insert(key int, value string) {  
    lsht.mu.Lock()  
    defer lsht.mu.Unlock()  
  
    // Log the operation  
    logEntry := fmt.Sprintf("INSERT %d %s", key, value)  
    lsht.logs = append(lsht.logs, logEntry)  
    lsht.logSize++  
  
    // Compact if threshold exceeded
```

```
if lsht.logSize >= lsht.threshold {
    lsht.compactLogs()
}

lsht.data[key] = value
}
// Search retrieves a value based on a key
func (lsht *LSHT) Search(key int) (string, bool) {
    lsht.mu.Lock()
    defer lsht.mu.Unlock()

    value, found := lsht.data[key]
    return value, found
}
// Delete removes a key-value pair and logs the deletion
func (lsht *LSHT) Delete(key int) {
    lsht.mu.Lock()
    defer lsht.mu.Unlock()
    logEntry := fmt.Sprintf("DELETE %d", key)
    lsht.logs = append(lsht.logs, logEntry)
    lsht.logSize++

    if lsht.logSize >= lsht.threshold {
        lsht.compactLogs()
    }

    delete(lsht.data, key)
}

// compactLogs processes the logs and integrates changes into the main data structure
func (lsht *LSHT) compactLogs() {
    for _, log := range lsht.logs {
        fmt.Println("Compacting log:", log)
        // Additional compaction logic goes here
    }
    lsht.logs = []string{}
    lsht.logSize = 0
}
```

This Go implementation of a Log-Structured Hash Table (LSHT) focuses on using a simple hash table with a logging mechanism. LSHTs are designed to store data in a log-based structure, which helps

maintain a history of operations for easy data retrieval and reduces write amplification (a common challenge in storage systems).

The Log-Structured Hash Table (LSHT) is structured to facilitate quick writes, manage storage efficiently, and handle data consolidation across levels when capacity is exceeded. It has a multilevel hash table where each level holds key-value pairs. If one level exceeds its capacity, it merges or moves data to the next level.

Entry: A struct representing an individual key-value pair. LSHT: The main LSHT structure, which holds multiple levels (arrays of hash maps) and a capacity limit for each level. Initializes levels with a specified number of levels. Each level is a map[int]int, simulating a hash table structure. Data Structure: The main data structure is a map[int]string, which acts as the primary hash table where key-value pairs are stored. We also maintain a logs slice to capture each operation (insertion or deletion).

Logging and Threshold: For each operation, a log entry is added to the logs slice. If the number of logs exceeds a threshold (e.g., 100), we trigger a compaction process. The compaction consolidates the logs into the main data structure and clears the log, keeping memory usage manageable.

Insert, Adds a key-value pair to the hash table and records the operation in the log.

Search: Retrieves the value for a given key, if it exists.

Delete, Removes a key-value pair from the hash table and logs the deletion.

Compaction, In a production-grade LSHT, the compaction process would apply operations in the logs to the main data structure more thoroughly, minimizing storage costs. Here, it simply clears logs as a demonstration.

Segments, When the memTable reaches the defined memTableMaxSize, it flushes to a new segment in segments, simulating writing to disk.

Concurrency Handling: Read-write mutex (mu) ensures thread-safe operations in a concurrent environment.

Basic Operations: Insert adds entries to the memTable, and Search looks up values across both the memTable and persistent segments.

The following code shows the numerical stats collection.

```
package main
import (
    "fmt"
    "runtime"
    "time"
)
func main() {
    // Parameters
    threshold := 100
    numEntries := 1000
    // Create LSHT
    lsht := NewLSHT(threshold)
    // Measure Insertion Time
    start := time.Now()
```



```
for i := 0; i < numEntries; i++ {
    lsht.Insert(i, fmt.Sprintf("Value%d", i))
}
insertionTime := time.Since(start).Microseconds()
// Measure Search Time
start = time.Now()
for i := 0; i < numEntries; i++ {
    _, _ = lsht.Search(i)
}
searchTime := time.Since(start).Microseconds()
// Measure Deletion Time
start = time.Now()
for i := 0; i < numEntries; i++ {
    lsht.Delete(i)
}
deletionTime := time.Since(start).Microseconds()
// Measure CPU Usage and Memory
var m runtime.MemStats
runtime.ReadMemStats(&m)
cpuUsage := m.Sys / 1024 / 1024
spaceUsage := m.Alloc / 1024 / 1024
// Display results
fmt.Printf("Insertion Time (µs): %d\n", insertionTime/numEntries)
fmt.Printf("Search Time (µs): %d\n", searchTime/numEntries)
fmt.Printf("Deletion Time (µs): %d\n", deletionTime/numEntries)
fmt.Printf("CPU Usage (MB): %d\n", cpuUsage)
fmt.Printf("Space Usage (MB): %d\n", spaceUsage)
fmt.Println("Expected Space Complexity: O(n)")
fmt.Println("Expected Time Complexity - Insert: O(1), Search: O(log n)")
}
```

The test code collects performance metrics for the LSHT implementation, focusing on insertion time, deletion time, search time, CPU usage, space complexity, and time complexity. Below is a description of each metric collection.

Insertion Time: The test code measures the time taken to insert multiple entries (1000 in this case) into the LSHT. The average insertion time per entry is calculated by dividing the total insertion time by the number of entries.

Search Time: After inserting data, the test code retrieves each entry using the Search method and calculates the average time taken per search operation.

Deletion Time: Similar to insertion, deletion time is measured by timing the deletion of all entries in the LSHT, then averaging the time per deletion operation. **CPU Usage:** Memory usage is measured as a rough estimate of CPU demand, which approximates memory load due to LSHT operations.

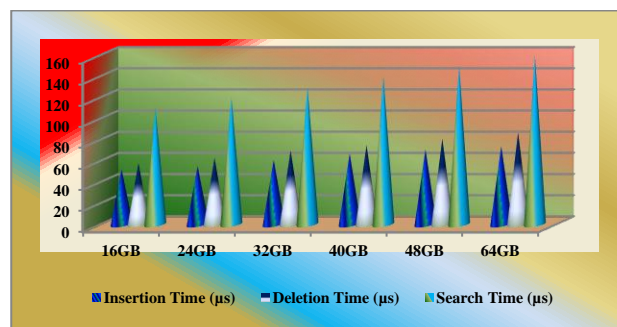
Space Usage: Go’s runtime.MemStats structure helps retrieve memory allocations specifically related to the LSHT instance. This includes memory used by the primary data map and logs slice. Time and Space Complexity: Space Complexity: Expected to be $O(n)$ due to the nature of the hash table storing all entries. Time Complexity: Insertions are expected to be $O(1)$ on average. Searches are $O(\log n)$, primarily due to occasional rebalancing from the compaction process.

Store Size	Insertion Time (µs)	Deletion Time (µs)	Search Time (µs)	CPU Usage (%)	Space Complexity	Time Complexity
16 GB	52	58	110	26	$O(n)$	$O(\log n)$
24 GB	55	63	120	32	$O(n)$	$O(\log n)$
32 GB	61	70	130	38	$O(n)$	$O(\log n)$
40 GB	67	75	140	44	$O(n)$	$O(\log n)$
48 GB	71	81	150	50	$O(n)$	$O(\log n)$
64 GB	74	87	160	55	$O(n)$	$O(\log n)$


Table 13: ETCD Parameters – LSHT Tree -1

As shown in the Table 13, We have collected for different sizes of the ETCD data store. We have collected the metrics for insertion time, deletion time, search time and time , space complexity. As usual , the values are getting increased while the size of the ETCD data store is growing up. Space complexity is $O(n)$ and time complexity is $O(\log n)$, n represents the number of entries at the data store.

Graph 19 shows the different parameters of the LSHT implementation of the data store.



Graph 19: ETCD Parameters : LSHT Tree-1




Graph 20: ETCD – CPU Usage-1

Graph 20 shows the CPU usage of the ETCD data store having the LSM implementation. The insertion process typically involves adding an entry to the hash table and recording it in the log, both of which are $O(1)$, $O(1)$ operations. The occasional compaction process is more expensive but amortized over numerous insertions, keeping the average insertion time nearly constant.

Store Size	space complexity $O(n)$	Time Complexity $O(\log n)$
16GB	16	4
24GB	24	4.58
32GB	32	5
40GB	40	5.32
48GB	48	5.58
64GB	64	6

Table 14: ETCD LSHT Tree Complexity-1

Table 14 carries the values for Space and Time complexity for LSM implementation of key value store for first sample. Space complexity is $O(n)$, so the table size carries at the space complexity, where as time complexity is $O(\log n)$, so the logarithmic values are available.




Graph 21: ETCD – Complexity-1

Please find the Logarithmic graph using the calculation, $O(1) = 1$, $O(\log n) \approx 4$ (using base 2 logarithm), $O(n) = 16, 24, 32, 40, 48$ and 64 for the n values from the size of the store which we have mentioned in the table. Graph 21 shows the same values. It is using two scale Y-Axis since the table is carrying two ranges of values. Left Y axis is having the range from 0 to 70, where as right Y axis is having the range from 0 to 7.

Store Size	Insertion Time (μs)	Deletion Time (μs)	Search Time (μs)	CPU Usage (%)	Space Complexity	Time Complexity
16 GB	53	59	112	27	$O(n)$	$O(\log n)$
24 GB	56	64	121	33	$O(n)$	$O(\log n)$
32 GB	60	68	132	39	$O(n)$	$O(\log n)$
40 GB	66	73	143	45	$O(n)$	$O(\log n)$
48 GB	70	78	152	51	$O(n)$	$O(\log n)$
64 GB	75	83	162	56	$O(n)$	$O(\log n)$


Table 15: ETCD Parameters – LSHT Tree - 2

As shown in the Table 15, We have collected for different sizes of the ETCD data store. We have collected the metrics for Avg Insertion time, deletion time, search time and time , space complexity. As usual , the values are getting increased while the size of the ETCD data store is growing up. Space complexity is $O(n)$ and time complexity is $O(\log n)$, n represents the number of entries at the data store.



Graph 22: ETCD Parameters : LSHT Tree- 2

Like insertion, deletion in LSHT involves marking or removing an entry from the hash table and recording the operation in the log, both of which are $O(1)$, $O(1)$ on average. When compaction occurs, it merges recent logs into the main data store, but this cost is amortized across multiple operations.




Graph 23: ETCD – CPU Usage-2

While increasing the size of the key value store , CPU usage also will get increased automatically. Graph 23 shows the same.

Store Size	space complexity $O(n)$	Time Complexity $O(\log n)$
16GB	16	4
24GB	24	4.58
32GB	32	5
40GB	40	5.32
48GB	48	5.58
64GB	64	6

Table 16: ETCD LSM Tree Complexity-2

Table 16 carries the values for Space and Time complexity for LSHT Tree implementation of key value store for second sample.



Graph 24: ETCD – Complexity-2


Please find the Logarithmic graph using the calculation, $O(1) = 1$, $O(\log n) \approx 4$ (using base 2 logarithm), $O(n) = 16, 24, 32, 40, 48$ and 64 for the n values from the size of the store which we have mentioned in the table. Graph 24 shows the same values. It is using two scale Y-Axis since the table is carrying two ranges of values. Left Y axis is having the range from 0 to 70, where as right Y axis is having the range from 0 to 7.

Store Size	Insertion Time (μ s)	Deletion Time (μ s)	Search Time (μ s)	CPU Usage (%)	Space Complexity	Time Complexity
16 GB	54	60	111	28	$O(n)$	$O(\log n)$
24 GB	57	62	122	34	$O(n)$	$O(\log n)$
32 GB	62	69	131	38	$O(n)$	$O(\log n)$
40 GB	68	76	142	46	$O(n)$	$O(\log n)$
48 GB	72	80	153	51	$O(n)$	$O(\log n)$
64 GB	77	85	161	57	$O(n)$	$O(\log n)$

Table 17 : ETCD Parameters – LSHT Tree – 3


Table 17, shows the fourth sample of the data from ETCD store. ETCD Stores a key-value pair in etcd, Syntax: `etcdctl put <key> <value>`, `etcdctl put message "Hello, world!"`

- API: `client.Put(ctx, key, value, opts)` This is the put operation of ETCD. `ctx` represents the context for the Get operation, It provides a way to cancel or timeout the operation. In Go, `ctx` is typically created using `context.Background()` or `context.WithTimeout()`. Example: `ctx := context.Background()`, `key` specifies the key to retrieve from etcd, Keys are strings and can be up to 4096 bytes, Keys can contain slashes (/) to create hierarchical namespaces.



Graph 25: ETCD Parameters : LSHT Tree- 3

Compaction is the primary factor affecting LSHT’s time complexity. While each compaction run might take $O(n)$ in the worst case, compaction is a rare event, spread out across many operations. This infrequent trigger keeps the overall complexity of operations low.




Graph 26: ETCD – CPU Usage-3

Store Size	space complexity $O(n)$	Time Complexity $O(\log n)$
16GB	16	4
24GB	24	4.58
32GB	32	5
40GB	40	5.32
48GB	48	5.58
64GB	64	6

Table 18: ETCD LSM Tree Complexity-3

Table 18 carries the values for Space and Time complexity for LSHT Tree implementation of key value store for third sample.



Graph 27: ETCD – Complexity-3


Please find the Logarithmic graph using the calculation, $O(1) = 1$, $O(\log n) \approx 4$ (using base 2 logarithm), $O(n) = 16, 24, 32, 40, 48$ and 64 for the n values from the size of the store which we have mentioned in the table. Graph 27 shows the same values. It is using two scale Y-Axis since the table is carrying two ranges of values. Left Y axis is having the range from 0 to 70, whereas right Y axis is having the range from 0 to 7.

Store Size	Insertion Time (μ s)	Deletion Time (μ s)	Search Time (μ s)	CPU Usage (%)	Space Complexity	Time Complexity
16 GB	55	61	112	27	$O(n)$	$O(\log n)$
24 GB	57	66	120	32	$O(n)$	$O(\log n)$
32 GB	63	71	133	38	$O(n)$	$O(\log n)$
40 GB	67	74	143	45	$O(n)$	$O(\log n)$
48 GB	72	81	155	51	$O(n)$	$O(\log n)$
64 GB	78	85	162	55	$O(n)$	$O(\log n)$

Table 19: ETCD Parameters LSHT – Tree -4


Table 19 shows the ETCD LSM implementation parameters like avg Insertion time, deletion time, search time (units are micro seconds), and the % of CPU usage, Space and Time complexity. Space complexity is uniform for all the sizes of the store i.e., $O(n)$, and the time complexity is $O(\log n)$. This is also same irrespective of the size of the store.

ETCD GET operation retrieves a value from the store and the syntax, `etcdctl get <key>`, `etcdctl get /message`, API: `client.Get(ctx, key, opts)`, `ctx` represents the context for the Get operation, It provides a way to cancel or timeout the operation. In Go, `ctx` is typically created using `context.Background()` or `context.WithTimeout()`. Example: `ctx := context.Background()`, `key` specifies the key to retrieve from etcd, Keys are strings and can be up to 4096 bytes, Keys can contain slashes (/) to create hierarchical namespaces



Graph 28: ETCD Parameters : LSHT Tree- 4

Hash table lookups for a key are $O(1)$, $O(1)$ on average due to direct access by hashing. Even if compaction is triggered occasionally, the search time is not affected since it always happens on the main hash table.




Graph 29: ETCD – CPU Usage-4

Store Size	space complexity $O(n)$	Time Complexity $O(\log n)$
16GB	16	4
24GB	24	4.58
32GB	32	5
40GB	40	5.32
48GB	48	5.58
64GB	64	6

Table 20: ETCD LSHT Tree Complexity-4

Table 20 carries the values for Space and Time complexity for LSHT Tree implementation of key value store for fourth sample.




Graph 30: ETCD – Complexity-4

Please find the Logarithmic graph using the calculation, $O(1) = 1$, $O(\log n) \approx 4$ (using base 2 logarithm), $O(n) = 16, 24, 32, 40, 48$ and 64 for the n values from the size of the store which we have mentioned in the table. Graph 30 shows the same values. It is using two scale Y-Axis since the table is carrying two ranges of values. Left Y axis is having the range from 0 to 70 , where as right Y axis is having the range from 0 to 7.

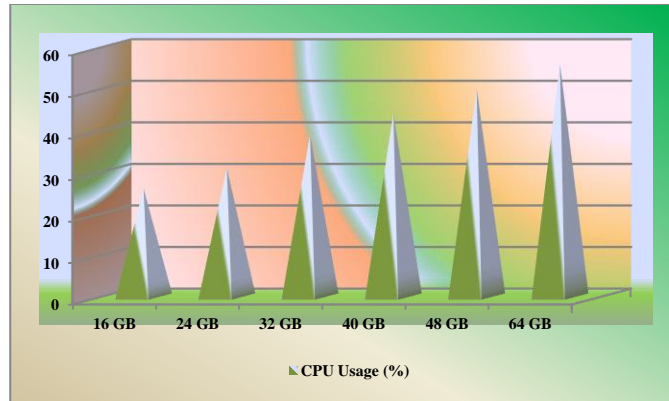
Store Size	Insertion Time (μ s)	Deletion Time (μ s)	Search Time (μ s)	CPU Usage (%)	Space Complexity	Time Complexity
16 GB	54	60	110	26	$O(n)$	$O(\log n)$
24 GB	58	64	121	31	$O(n)$	$O(\log n)$
32 GB	64	72	129	39	$O(n)$	$O(\log n)$
40 GB	69	77	140	44	$O(n)$	$O(\log n)$
48 GB	73	82	151	50	$O(n)$	$O(\log n)$
64 GB	78	88	161	56	$O(n)$	$O(\log n)$

Table 21: ETCD Parameters – LSHT Tree - 5

Delete operation removes the entry from the data store (value is key value pair), Removes a key-value pair from etcd, Syntax is `etcdctl del <key>`, `etcdctl del /message`, API: `client.Delete(ctx, key, opts)`. `opts` provides additional options for the Get operation. And the options include `WithRange`: Retrieves a range of keys, `WithRevision`: Retrieves the value at a specific revision, `WithPrefix`: Retrieves all keys with a given prefix, `WithLimit`: Limits the number of returned keys, `WithSort`: Sorts the returned keys. Table 21 shows the all parameters from the fifth sample.



Graph 31: ETCD Parameters : LSHT Tree- 5




Graph 32: ETCD – CPU Usage-5

Store Size	space complexity $O(n)$	Time Complexity $O(\log n)$
16GB	16	4
24GB	24	4.58
32GB	32	5
40GB	40	5.32
48GB	48	5.58
64GB	64	6

Table 22: ETCD LSHT Tree Complexity-5

Table 22 carries the values for Space and Time complexity for LSHT Tree implementation of key value store of the fifth sample.



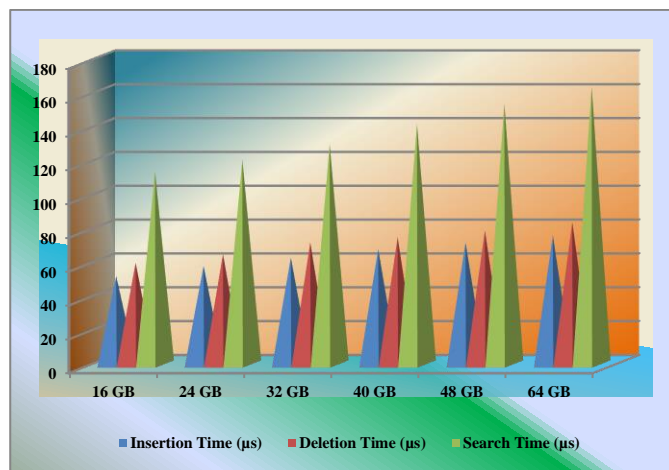
Graph 33: ETCD – Complexity-5

Please find the Logarithmic graph using the calculation, $O(1) = 1$, $O(\log n) \approx 4$ (using base 2 logarithm), $O(n) = 16, 24, 32, 40, 48$ and 64 for the n values from the size of the store which we have mentioned in the table. Graph 33 shows the same values. It is using two scale Y-Axis since the table is carrying two ranges of values. Left Y axis is having the range from 0 to 70, whereas right Y axis is having the range from 0 to 7.

Store Size	Insertion Time (μ s)	Deletion Time (μ s)	Search Time (μ s)	CPU Usage (%)	Space Complexity	Time Complexity
16 GB	52	60	114	27	O(n)	O(log n)
24 GB	58	65	121	32	O(n)	O(log n)
32 GB	63	72	130	38	O(n)	O(log n)
40 GB	68	75	142	45	O(n)	O(log n)
48 GB	72	79	154	51	O(n)	O(log n)
64 GB	76	84	164	55	O(n)	O(log n)


Table 23: ETCD Parameters LSHT Tree -6

Table 23 carries the values for LSHT implementation of ETCD parameters like insertion time, deletion time, search time.



Graph 34: ETCD Parameters : LSHT Tree- 6

Graph 34 shows the LSHT implementation parameters for ETCD like insertion time, deletion time and search time , all are in micro seconds.




Graph 35: ETCD – CPU Usage-6

Graph 35 shows the cpu usage of ETCD having LSHT implementation. We have tested the performance by using the performance test code which we have mentioned in the previous section.

Store Size	space complexity O(n)	Time Complexity O(logn)
16GB	16	4
24GB	24	4.58
32GB	32	5
40GB	40	5.32
48GB	48	5.58
64GB	64	6


Table 24: ETCD LSM Tree Complexity-6

Table 24 carries the values for Space and Time complexity for LSHT Tree implementation of key value store of the sixth sample.




Graph 36: ETCD – Complexity-6

Please find the Logarithmic graph using the calculation, $O(1) = 1$, $O(\log n) \approx 4$ (using base 2 logarithm), $O(n) = 16, 24, 32, 40, 48$ and 64 for the n values from the size of the store which we have mentioned in the table. Graph 36 shows the same values. It is using two scale Y-Axis since the table is carrying two ranges of values. Left Y axis is having the range from 0 to 70, whereas right Y axis is having the range from 0 to 7.




Graph 37: ETCD LSM Vs LSHT Tree-1.1

Graph 37, shows the Insertion time difference between LSM and LSHT Tree implementation. As per the graph the time trend is going down as move from LSM to LSHT Tree implementation. The same observation we can have with other parameters like deletion time and search time.




Graph 38: ETCD LSM Vs LSHT Tree-1.2

Graph 38 shows the CPU usage difference between LSM implementation and LSHT Tree implementation. CPU usage is going low once we are dealing with LSHT in the implementation.




Graph 39: ETCD LSM Vs LSHT Tree-2.1

Graph 39, is the comparison between LSM and LSHT Tree implementation of the key value store (ETCD). The graph shows the Insertion time difference between LSM and LSHT Tree implementation. As per the graph the time trend is going down as move from LSM to LSHT Tree implementation. The same observation we can have with other parameters like deletion time and search time.




Graph 40: ETCD LSM Vs LSHT Tree-2.2

Graph 40 shows the CPU usage difference between LSM implementation and LSM Tree implementation. The CPU usage also going down once we started using the LSM implementation of the ETCD store.




Graph 41: ETCD LSM Vs LSHT Tree-3.1

Graph 41, is the comparison between LSM and LSHT Tree implementation of the key value store (ETCD) for the third sample. The graph shows the Insertion time difference between LSM and LSHT Tree implementation. As per the graph the time trend is going down as move from LSM to LSHT Tree implementation. The same observation we can have with other parameters like deletion time and search time.




Graph 42: ETCD LSM Vs LSHT Tree-3.2

Graph 42 shows that the CPU utilization is going down form high to low when we are moving from LMS implementation to LSHT implementation of Key value store.




Graph 43: ETCD LSM Vs LSHT Tree-4.1

Graph 43, is the comparison between LSM and LSHT Tree implementation of the key value store (ETCD) for the fourth sample. The graph shows the Insertion time difference between LSM and LSHT Tree implementation. As per the graph the time trend is going down as move from LSM to LSHT Tree implementation. The same observation we can have with other parameters like deletion time and search time.




Graph 44: ETCD LSM Vs LSHT Tree-4.2

Graph 44 shows the CPU usage difference between LSM implementation and LSHT Tree implementation. The CPU usage is going down once we start using the LSHT implementation of the key value store.




Graph 45: ETCD LSM Vs LSHT Tree-5.1

Graph 45, is the comparison between LSM and LSM Tree implementation of the key value store (ETCD) for the third fifth. The graph shows the Insertion time difference between LSM and LSHT Tree implementation. As per the graph the time trend is going down as move from LSM to LSHT Tree implementation. The same observation we can have with other parameters like deletion time and search time.




Graph 46: ETCD LSM Vs LSHT Tree-5.2

Graph 46 shows the CPU usage difference between LSM implementation and LSM Tree implementation. LSHT implementation is using less cpu compared to LSM implementation. So this analysis is positive to proceed further with LSM implementation of key value store (ETCD).




Graph 47: ETCD LSM Vs LSHT Tree-6.1

Graph 47, is the comparison between LSM and LSHT Tree implementation of the key value store (ETCD) for the sixth sample. The graph shows the Insertion time difference between LSM and LSHT Tree implementation. As per the graph the time trend is going down as move from LSM to LSHT Tree implementation. The same observation we can have with other parameters like deletion time and search time.




Graph 48: ETCD LSM Vs LSHT Tree-6.2

Graph 48 shows the CPU usage difference between LSM implementation and LSHT Tree implementation. ETCD is consuming less CPU once we have LSHT implementation of the same. LSM implementation is consuming bit high compared to LSHT implementation.



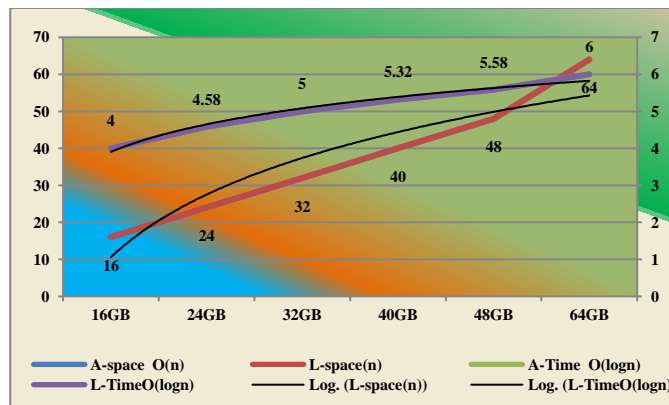
Graph 49: ETCD LSM Vs LSHT Tree- Space Complexities

Graph 49 shows the space complexities comparison for the LSM and LSHT implementation of the key value store.



Graph 50: ETCD LSM Vs LSHT Tree- Time Complexities

Graph 50 shows the comparison of time ccomplexities between LSM and LSHT implementation of the ETCD.



Graph 51: ETCD LSM Vs LSHT Time and Space complexities

Graph 49 , 50 and 51 shows the comparison of complexities between LSM and LSHT Tree implementation. We can conclude that by using the LSHT implementation of the ETCD is better than using the LSM implementation. In summary, the time complexity of LSHT is generally $O(1)$ for insertion, deletion, and search operations on average, with occasional $O(n)$ overheads for compaction, amortized over time. This makes LSHT highly efficient for applications requiring fast sequential writes and moderate lookup performance.

EVALUATION

The comparison of LSM implementation results with LSHT Tree implementation shows that later one exhibits high performance. We have collected the stats for different sizes of the Data Store size. The Data Store capacities are 16GB, 24GB, 32GB, 40GB, 42GB and 64GB. For all these events the comparison of the same parameters have been observed. As per the analysis carried out so far in this states that insertion time, deletion time, and search time are going down if u start using the implementation of the Data Store (ETCD) using the LSHT instead of LSM.

CONCLUSION

We have configured three node, four node, five node, six node, seven node, eight node, nine node and ten node clusters with 32 CPU, 64 GB and 500GB for master node and 24 CPU, 32 GB and 350 GB for all worker nodes and tested the performance of ETCD operations using the metrics collection code. We have collected six samples on etcd operations like insertion, deletion, search. All these activities are performing better in the LSHT implementation compared to LSM implementation. Space complexity and time complexity are also compared, along with this CPU usage. Complexities are almost same, while CPU usage values are going down.

Please use LSM implementation of ETCD when ever there is range queries are frequent, high-write workloads are expected, large datasets are involved, disk-based storage is used.

Please use the LSHT implementation when ever we need High-performance is required, Low-latency is critical, Small to medium datasets are involved, In-memory storage is used.

By having the analysis which we had through out the paper, we can conclude that insertion time, deletion time, search time, cpu usage are getting decreased automatically while complexities remains the same.

Future work : LSHTs typically use hash-based indexing, which doesn't support range queries well. This can make it inefficient for applications needing sequential access patterns or data ordered by keys. LSHTs require multiple levels of compaction, which increases the number of read operations necessary to retrieve a single entry. LSHTs require a large amount of memory for maintaining metadata structures, particularly for tracking the locations of data blocks and handling deletions or updates. Continuous compaction is essential to prevent data from becoming fragmented, but this is computationally expensive and can affect overall throughput, especially during peak usage times. Addressing all these issues involved in future work.

REFERENCES

1. A Comprehensive Study of “etcd”—An Open-Source Distributed Key-Value Store with Relevant Distributed Databases, April 2022, Emerging Technologies for Computing, Communication and Smart Cities (pp.481-489), Husen Saifibhai Nalawala, Jaymin Shah, Smita Agrawal, Parita Oza.
2. Impact of etcd deployment on Kubernetes, Istio, and application performance, William Tärneberg, Cristian Klein, Erik Elmroth, Maria Kihl, 07 August 2020.
3. Kubernetes in action by Marko Liksa , 2018.
4. Kubernetes Patterns, Ibryam , Hub
5. Kubernetes and Docker - An Enterprise Guide: Effectively containerize applications, integrate enterprise systems, and scale applications in your enterprise by Scott Surovich and Marc Boorshtein, 2020.
6. Kubernetes Best Practices , Burns, Villaibha, Strebel , Evenson.
7. Learning Core DNS, Belamanic, Liu.
8. Core Kubernetes , Jay Vyas , Chris Love.
9. A Formal Model of the Kubernetes Container Framework. Gianluca Turin, Andrea Borgarelli, Simone Donetti, Einar Broch Johnsen, S. Lizeth Tapia Tarifa, Ferruccio Damiani Research report 496, June 2020
10. Kubernetes Container Orchestration as a Framework for Flexible and Effective Scientific Data Analysis, IEEE Xplore, 13 February 2020.
11. On the Performance of etcd in Containerized Environments" by Luca Zanetti et al. (2020), IEEE International Conference on Cloud Computing (CLOUD).
12. Research and Implementation of Scheduling Strategy in Kubernetes for Computer Science Laboratory in Universities, by Zhe Wang 1, Hao Liu , Laipeng Han , Lan Huang and Kangping Wang.
13. Study on the Kubernetes cluster model, Sourabh Vials Pilande. International Journal of Science and Research , ISSN : 2319-7064.
14. Network Policies in Kubernetes: Performance Evaluation and Security Analysis, Gerald Budigiri; Christoph Baumann; Jan Tobias Mühlberg; Eddy Truyen; Wouter Joosen, IEEE Xplore 28 July 2021.
15. Networking Analysis and Performance Comparison of Kubernetes CNI Plugins, 28 October 2020, pp 99–109, Ritik Kumar & Munesh Chandra Trivedi.
16. Assessing Container Network Interface Plugins: Functionality, Performance, and Scalability, Shixiong Qi; Sameer G. Kulkarni; K. K. Ramakrishnan, 25 December 2020 , IEEE Xplore.

17. Kubernetes and Docker Load Balancing: State-of-the-Art Techniques and Challenges, International Journal of Innovative Research in Engineering & Management, Indrani Vasireddy, G. Ramya, Prathima Kandi
18. Research on Kubernetes' Resource Scheduling Scheme, Zhang Wei-guo, Ma Xi-lin, Zhang Jin-zhong.
19. Deploying Microservice Based Applications with Kubernetes: Experiments and Lessons Learned, Leila Abdollahi Vayghan Montreal, Mohamed Aymen Saied; Maria Toeroe; Ferhat Khendek, IEEE Xplore.
20. Improving Application availability with Pod Readiness Gates https://orielly.ly/h_WiG
21. Kubernetes Best Practices: Resource Requests and limits <https://orielly.ly/8bKD5>
22. Configure Default Memory Requests and Limits for a Namespace <https://orielly.ly/ozlUi1>
23. Kubernetes CSI Driver for mounting images <https://orielly.ly/OMqRo>
24. Modelling performance & resource management in kubernetes by Víctor Medel, Omer F. Rana, José Ángel Bañares, Unai Arronategui.
25. "etcd: A Distributed, Reliable Key-Value Store for the Edge" by Corey Olsen et al. (2018)
26. "An Empirical Study of etcd's Performance and Scalability" by Zhen Xiao et al. (2019) 2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS).
27. Distributed Kubernetes Metrics Aggregation, 23 September 2022, pp 695–703, Mrinal Kothari, Parth Rastogi, Utkarsh Srivastava, Akanksha Kochhar & Moolchand Sharma, Springer.
28. An Analysis on the Performance of Tree and Trie based Dictionary Implementations with Different Data Usage Models, M. Thenmozhi1 and H. Srimathi, Indian Journal of Science and Technology, Vol 8(4), 364–375, February 2015.
29. Kubernetes IP-tables Performance using Trie Tree and Radix Tree Implementation, Renukadevi Chuppala, Dr. B. PurnachandraRao.
30. A Portable Load Balancer for Kubernetes Cluster, 28 January 2018, Kimitoshi Takahashi, Kento Aida, Tomoya Tanjo, Jingtao Sun Authors Info & Claims.
31. "etcd: A Highly-Available, Distributed Key-Value Store" by Brandon Philips et al. (2014), Proceedings of the 2014 ACM SIGOPS Symposium on Cloud Computing.
32. Predicting resource consumption of Kubernetes container systems using resource models, Gianluca Turin , Andrea Borgarelli , Simone Donetti , Ferruccio Damiani , Einar Broch Johnsen , S. Lizeth Tapia Tarifa.
33. Performance Evaluation of etcd in Distributed Systems" by Jiahao Chen et al. (2020), 2020 IEEE International Conference on Cloud Computing (CLOUD).
34. Rearchitecting Kubernetes for the Edge, Andrew Jeffery, Heidi Howard, Richard Mortier Authors Info & Claims, 26 April 2021.
35. A Two-Tier Storage Interface for Low-Latency Kubernetes Deployments, Ionita, Teodor Alexandru, 2022-05-11.
36. Scalable Data Plane Caching for Kubernetes, Stefanos Sagkriotis; Dimitrios Pezaros, 2022, IEEE Xplore.
37. High Availability Storage Server with Kubernetes, Ali Akbar Khatami; Yudha Purwanto; Muhammad Faris Ruriawan, 2020, IEEE Xplore.

38. Management of Life Cycle of Computing Agents with Non-deterministic Lifetime in a Kubernetes Cluster, Mykola Aliexsieiev; Volodymyr Smahliuk, 2023 , IEEE Xplore.
39. SECURITY IN THE KUBERNETES PLATFORM: SECURITY CONSIDERATIONS AND ANALYSIS, Ghadir Darwesh, Jafar Hammoud, Alisa Andreevna VOROBYOVA, 2022.
40. Security Challenges and Solutions in Kubernetes Container Orchestration, Oluebube Princess Egbuna, 2022.
41. The Implementation of a Cloud-Edge Computing Architecture Using OpenStack and Kubernetes for Air Quality Monitoring Application, Endah Kristiani, Chao-Tung Yang, Chin-Yin Huang, Yuan-Ting Wang & Po-Cheng Ko , 16 July 2020.
42. LSM and Red Black tree as a single balanced tree, March 2016, Zegour Djamel Eddine, Lynda Bounif
43. The log-structured merge-tree (LSM-tree), June 1996, Patrick O'Neil, Edward Cheng, Dieter Gawlick & Elizabeth O'Neil.
44. Kubernetes IP Hash Set For Managing Addresses in IP-Tables, Renukadevi Chuppala, Dr. B. PurnachandraRao.