

Deep Multi-Layer Isolation for Secure Kubernetes Multi-Tenancy in a Single Shared Cluster

Sharan Babu Paramasivam Murugesan
Independent Researcher
Seattle, WA, USA

ABSTRACT

Kubernetes has revolutionized micro-service hosting by automating the deployment, scaling, and management of containers across clusters. Its flexibility and portability have made it the de-facto platform for cloud-native applications. However, the widespread practice of creating a separate cluster for each team, application, or customer has led to cluster sprawl. Each cluster carries a baseline resource overhead, causing many small clusters to remain under-utilized and difficult to manage. Conversely, consolidating workloads into a single multi-tenant cluster can improve utilization but introduces challenges related to security, fairness, governance, and observability. Misconfigurations such as overly permissive access or missing network policies can compromise isolation, exposing every tenant to risk.

This paper analyzes the risks of both fragmented and consolidated approaches and proposes a layered isolation architecture that retains the benefits of consolidation while mitigating multi-tenancy risks. The analysis describes how namespacing, RBAC, resource quotas, network policies, service meshes with mutual TLS (mTLS), admission control, and optional dedicated node pools can be composed to provide strong logical isolation within a single cluster. Recent research shows that virtual clusters that are running per-tenant control-plane components on top of a host cluster are an emerging tool for achieving hard multi-tenancy [1]. A reference architecture is presented along with guidance on governance, observability, and cost optimization. The proposed design is evaluated against current multi-tenancy practices, and future research directions are identified.

General Terms

Kubernetes, Multi-Tenancy, Cluster Sprawl, Systems Design, Microservices, Container Orchestration, Platform Engineering, Software Architecture, Infrastructure Optimization, Resource Isolation, Policy Enforcement, Security Boundaries, Resource Governance, Capacity Planning, Performance Isolation, Multi-Tenant SaaS, Scalability, Cluster Management, Admission Control, Scheduling Policies, Configuration Management

Keywords

Kubernetes multi-tenancy, workload isolation, cluster sprawl reduction, namespace-based security, resource-aware scheduling, node-level segregation, policy-driven governance, multi-tenant SaaS infrastructure, admission control strategies, platform efficiency in container orchestration, workload placement strategies.

1. INTRODUCTION

Since its launch in 2014, Kubernetes has become the dominant orchestration platform for containerized micro-services. Its automated scheduling, service discovery and self-healing capabilities enable teams to deploy services without managing individual virtual machines. According to a 2020 CNCF survey, production Kubernetes usage increased by over 300 %

since 2016 [2], and adoption continues to climb as of 2025 [3]. Many organizations initially adopted a cluster-per-team or cluster-per-application pattern to ensure clear isolation and performance boundaries. Each team or customer was given a dedicated cluster, avoiding any possibility of interference. However, this approach yields poor resource utilization because every cluster carries baseline overhead (system pods, control-plane components, etc.). When dozens or hundreds of clusters are provisioned, large portions of CPU and memory remain idle, and administrators must patch, upgrade and audit each cluster separately. To counter this cluster sprawl, platform teams are exploring single-cluster multi-tenancy—sharing one cluster among many teams or customers. Consolidation improves utilization and reduces maintenance overhead, but Kubernetes does not have a first-class concept of a tenant. Isolation instead relies on correct configuration of namespaces, RBAC, network policies and other primitives [3]. Misconfigurations can allow cross-tenant access or network eavesdropping, so the challenge is to design a multi-tenant architecture that preserves security and fairness without recreating the sprawl problem.

This paper analyzes the tradeoffs between cluster sprawl and multi-tenant consolidation, summarizing the risks in terms of resource utilization, security, fairness, governance, and observability. A layered isolation architecture is proposed that combines namespaces, RBAC, resource quotas, network policies, service mesh with mutual TLS (mTLS), admission control, and optional node isolation to form a defense in depth model within a single cluster. Guidance is provided on tenant aware governance and observability, highlighting how policy enforcement, monitoring, and logging must be adapted for shared environments, and emerging approaches such as virtual clusters, which run per tenant control plane components on top of a host cluster [1], are evaluated. Through a synthesis of recent literature, best practices, and practical experience, the paper articulates key design principles for achieving scalable and secure multi tenancy on Kubernetes and outlines directions for future improvement.

2. CHALLENGES OF MULTI-TENANCY AND CLUSTER SPRAWL

Effective multi-tenancy must address challenges stemming from both extremes: over-segmentation into many clusters and unsafe consolidation into a single cluster. The challenges can be grouped into resource utilization and scalability, security and governance, and observability.

2.1 Resource utilization and scalability

2.1.1 Underutilization from cluster sprawl

Running each tenant in an isolated cluster wastes resources because every cluster consumes baseline CPU and memory. In large organizations with dozens of clusters, a significant percentage of resources sits idle. Besides compute inefficiency,

cluster sprawl creates operational overhead: platform teams must manage multiple API servers, certificates and networking setups; apply upgrades and patches across all clusters; and replicate monitoring and policy configurations. This overhead slows provisioning and complicates scaling. Consolidating workloads into fewer clusters can increase utilization and reduce administrative cost.

2.1.2 Noisy neighbors and fairness in shared clusters

In a single cluster, noisy neighbor problems arise when one tenant consumes disproportionate resources, impacting others. Without controls, a rogue workload could monopolize CPU or memory. Kubernetes provides ResourceQuota and LimitRange objects to cap resource usage per namespace. By setting a default-deny quota and per-pod limits, each tenant's consumption is restricted to its allotted slice. This mechanism prevents one tenant from starving others and protects control-plane stability. Additionally, the API Priority and Fairness feature (introduced in v1.20) ensures that API server requests from different users or groups are fairly ordered, preventing a single tenant from overwhelming the API server. For more fine-grained fairness, research has proposed customized schedulers based on dominant resource fairness [4].

Scalability must also consider absolute cluster size and tenant count. Kubernetes can run clusters with thousands of nodes, but large object counts, and high churn can stress the control plane and etcd. Platform teams should monitor control-plane metrics and shard workloads across multiple clusters when necessary. Virtual clusters offer an intermediate approach: they run a dedicated API server and control-plane components per tenant on top of a host cluster, giving tenants full control over Kubernetes API operations while still sharing nodes [1].

2.2 Security and governance challenges

2.2.1 Isolation breaches via misconfiguration

Kubernetes does not enforce multi-tenancy by default [3]. Isolation relies on correct configuration of Namespaces, RBAC, NetworkPolicy, PodSecurity and other objects. Misconfigured RBAC (e.g., cluster-admin privileges granted to a tenant service account) or missing network policies can allow cross-namespace access. In practice, such mistakes have enabled attackers to escape their namespace and compromise the entire cluster. FinTech case studies emphasize that multi-tenant architectures increase the attack surface because each tenant's workload is a potential entry point [5]. A zero-trust approach—assume no tenant is trustworthy and enforce least privilege—is necessary for strong isolation.

2.2.2 Governance and observability gaps

Consolidating into a single cluster simplifies some aspects of governance (one set of cluster-wide policies) but introduces complexity in ensuring that all tenants adhere to security and compliance requirements. Administrators must enforce policies centrally and prevent any tenant from bypassing them. Tools like Open Policy Agent (OPA)/Gatekeeper or Kyverno can

validate resources at admission time, blocking workloads that violate security baselines. Auditing becomes more demanding because logs and metrics must be partitioned per tenant to avoid data leakage. Without tenant-aware observability, troubleshooting incidents can be difficult and cross-tenant data exposures may occur. Platform teams should ensure that monitoring agents run with restricted privileges and only collect telemetry for authorized namespaces. RBAC rules must also protect observability dashboards (e.g., Grafana, Kibana) so that tenants cannot view each other's logs or metrics. Case studies of managing multi-tenant clusters for enterprise platforms such as AEM and HCL Commerce underscore the operational complexity of governance and the need for consistent best practices across tenants [6].

2.3 Summary of challenges

Both extremes - fragmented clusters and unsafe consolidation - carry significant risks. Cluster sprawl wastes resources and burdens operations, while multi-tenancy without proper isolation invites breaches and fairness issues. A successful solution must balance efficiency with strong security, fairness and governance. Table 1 summarizes the primary challenges.

Table 1. Summary of primary challenges with hosting multiple tenants

Challenge area	Risks	Multi-cluster (sprawl)	Single cluster
Resource utilization	Idle capacity; waste	Per-cluster overhead	Better packing; neighbor risk
Operational overhead	Duplicate toil	Many APIs/patches	One surface; automation
Security	Misconfig → exposure	Smaller blast; drift	Central policy; wider blast
Fairness / performance	Overprovisioning leads to waste	Over-provisioning	Use quotas/limits/priority & API fairness to prevent starvation
Governance & compliance	Inconsistent controls	Divergent configs	Central policy & audit; ensure tenant scoping and separation of duties
Observability & SRE	Noisy/leaky data	Per-cluster dashboards; low reuse	Tenant-aware labels/RBAC; central pipeline with per-tenant views

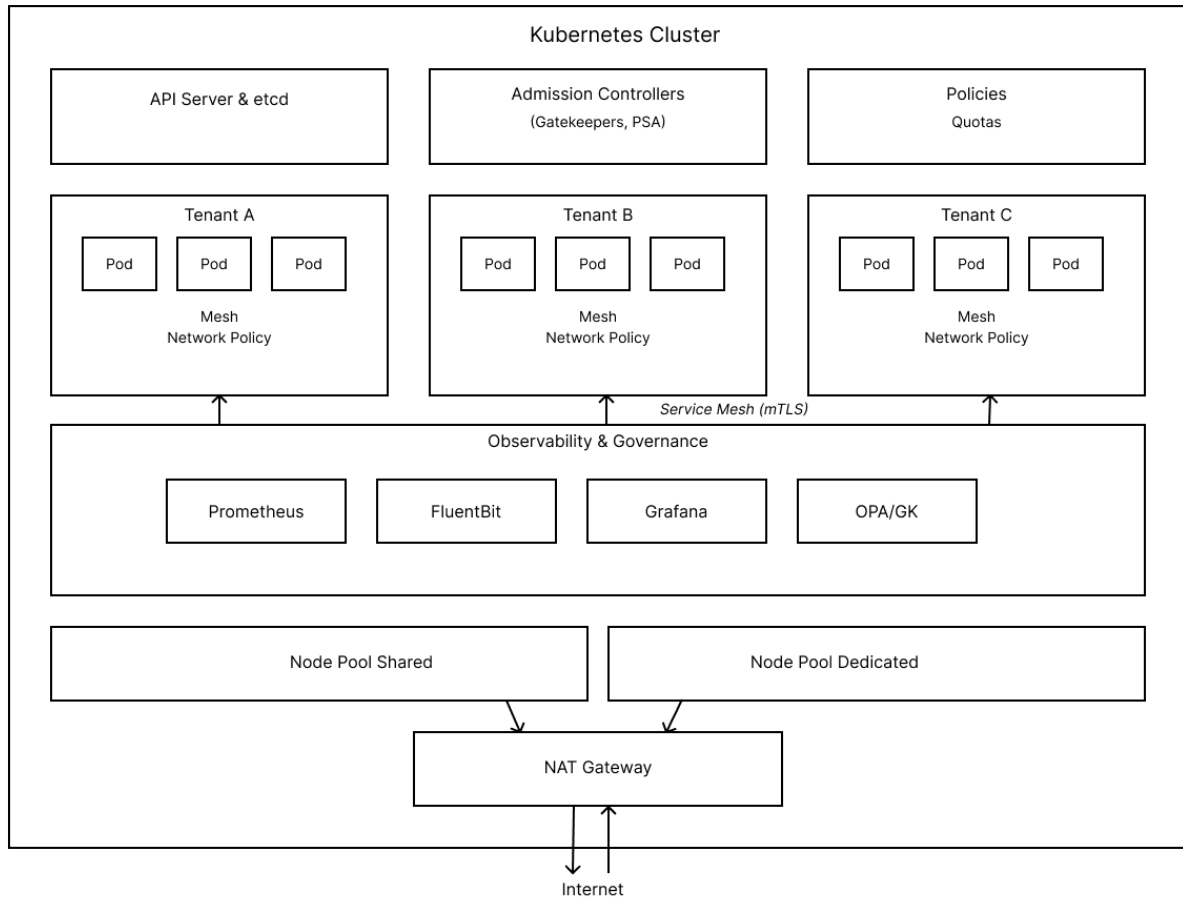


Fig 1: Architecture of a multi-tenant Kubernetes cluster (single cluster multi-tenancy)

3. LAYERED MULTI-TENANT ISOLATION ARCHITECTURE

To mitigate the challenges identified above, a layered isolation architecture for Kubernetes is presented that enables multiple teams or customers to safely share a single cluster. Each layer targets a specific class of risk, and together they form a defence-in-depth model. Figure 1 conceptually illustrates the design. The key layers are explored in detail below.

3.1 Namespaces and RBAC isolation

Namespaces form the first layer of isolation by grouping Kubernetes resources under unique prefixes. Names can overlap across namespaces, allowing tenants to use natural names without collisions [3]. A namespace alone does not guarantee security; it must be paired with careful access control and fair resource allocation.

3.1.1 Namespaces and RBAC design

In a multi-tenant cluster, namespaces logically isolate API objects. The platform team defines ClusterRoles (e.g., view, edit, admin) with precise scopes and binds them to user groups or service accounts within a tenant's namespace. Tenants receive only the privileges necessary to operate their workloads. Service accounts used by tenant pods should not have cluster-wide privileges. For example, a tenant's CI/CD pipeline might have "admin" rights within its namespace but zero rights outside it. Roles and RoleBindings should be stored in version control and automatically provisioned along with namespaces. This pattern ensures least-privilege access and prevents accidental cross-namespace operations.

3.1.2 Resource quotas and limit ranges

Beyond access control, fair resource sharing is essential. ResourceQuota and LimitRange objects assign each namespace a maximum number of pods, services and total CPU or memory usage. They ensure a tenant cannot consume unbounded resources. Under a quota regime, Kubernetes requires pods to specify CPU and memory requests/limits; the scheduler then throttles pods that exceed their limits, protecting against noisy neighbours. The platform team may vary quotas based on SLA tiers or subscription plans, creating "small," "medium" or "large" slices of the cluster.

3.2 Network policies and service mesh

Multi-tenant clusters require strict control over network traffic. By default, pods can communicate freely across namespaces, which is unacceptable when tenants must be isolated.

3.2.1 Network policy design

NetworkPolicy resources restrict pod ingress and egress. A default deny model is used, where each namespace starts with a policy that denies all traffic except for explicitly allowed sources. Typical allowances include communication within the same namespace and to essential system services such as DNS, logging or sidecar injection. For outbound traffic, destinations should be restricted to approved CIDR blocks and egress gateways (e.g., NAT instances or proxies). The cluster network plugin (e.g., Calico, Cilium, CNI implementations with policy support) must be configured to enforce these policies. Without a default-deny baseline, pods can inadvertently accept or initiate cross-tenant traffic.

3.2.2 Service mesh and mTLS

Network policies operate at the IP and port level; they cannot authenticate calls or encrypt traffic. A service mesh (e.g., Istio or Linkerd) adds a higher-layer control plane. Each pod gets a sidecar proxy that intercepts traffic and performs mutual TLS, ensuring that every service-to-service call is encrypted and both ends are authenticated. Istio issues per-service certificates and enforces trust domains, so calls between services from different tenants are rejected. Fine-grained AuthorizationPolicy rules can further restrict HTTP methods or paths (e.g., only allow GET calls to a public API). While the mesh introduces operational overhead (extra sidecar containers and a control plane), it provides strong data-plane isolation and telemetry. Lighter alternatives such as Linkerd may offer simpler deployment with fewer features.

3.3 Admission control and policy enforcement

Kubernetes allows administrators to validate or mutate resources before they are persisted. Admission control and pod security collectively form the fourth layer of the proposed architecture.

3.3.1 Admission control and policy enforcement

Admission controllers are webhooks that intercept API requests for resource creation or updates. Policy engines such as OPA/Gatekeeper and Kyverno can enforce security standards at deploy time. Typical admission policies include:

- Deny privileged pods or host-path mounts; require containers to run as non-root users.
- Require a default-deny NetworkPolicy in every namespace.
- Restrict creation of LoadBalancer or NodePort services to approved namespaces.
- Inject mandatory labels or annotations for governance, cost allocation or compliance reporting.

Gatekeeper policies are expressed in Rego and version-controlled. Kyverno policies are expressed in YAML and may be easier for some teams to adopt. Integrating admission control into CI/CD pipelines ensures that misconfigured or insecure workloads are rejected before deployment, preventing configuration drift.

3.3.2 Pod security standards

While admission controllers can enforce arbitrary rules, Kubernetes provides built-in Pod Security Admission (PSA) starting in v1.25. PSA replaces the deprecated PodSecurityPolicy and enforces pre-defined restricted, baseline or privileged profiles. Setting a namespace to the “restricted” profile ensures that containers cannot run as root, cannot add new Linux capabilities, and cannot mount sensitive host paths. By combining PSA with runtime classes (e.g., gVisor or Kata Containers), platform teams can enforce consistent security baselines across tenants. PSA is configured via namespace annotations, making it easy to bootstrap new tenants with the appropriate level of restriction.

3.4 Workload and node isolation

This layer concerns how workloads are scheduled onto nodes and how nodes themselves are partitioned. The goal is to balance high utilization with the option for stronger isolation when needed.

3.4.1 Shared node pools and runtime isolation

By default, pods are scheduled onto a common pool of nodes to maximize utilization. Resource quotas and per-pod limits

ensure fairness and protect against noisy neighbors. Platform teams can further isolate workloads at runtime using runtime classes or hardened container runtimes. For example, gVisor, Kata Containers or AWS Firecracker run each pod in a lightweight virtual machine, protecting the host kernel from untrusted workloads. These runtimes introduce some overhead but provide strong isolation for multi-customer environments.

3.4.2 Dedicated node pools and specialized hardware

Certain workloads require stricter separation or special hardware (e.g., GPU, high-memory, FIPS-compliant or compliance-certified nodes). Taints and tolerations, along with node affinity or node selectors, ensure that specific pods are scheduled onto designated nodes. Admission controllers can enforce correct placement by validating that only workloads requesting a GPU node are allowed onto GPU nodes. Dedicated node pools provide higher security and performance isolation at the cost of potential under-utilization. To manage cost, the platform team should enable autoscaling or dynamic provisioning for these pools so they expand only when needed.

3.4.3 Virtual clusters as control-plane isolation

A more advanced approach is to create per-tenant virtual clusters. In this model, each tenant runs its own API server and controller manager inside a namespace, effectively providing a virtual control plane while sharing the host cluster’s nodes. Research shows that virtual clusters overcome limitations of namespace isolation by allowing tenants to perform cluster-level operations (e.g., installing CRDs) while only viewing their own resources [1]. They provide a harder form of multi-tenancy but still rely on data-plane isolation to protect workloads. The trade-off is increased operational complexity: administrators must manage many small control planes, and new security considerations arise (e.g., isolating etcd stores). As discussed later, virtual clusters are emerging and may suit scenarios where tenants require near-complete autonomy [1].

3.5 Governance, observability and cloud agnosticism

The final layer concerns how policies are enforced, how telemetry is collected and how platform designs remain portable across cloud providers.

3.5.1 Policy enforcement and automation

Proper governance ensures compliance with security and regulatory requirements across tenants. Best practices include policy as code by defining policies in version-controlled repositories and applying them automatically via CI/CD pipelines. Gatekeeper, Kyverno or similar tools can enforce these policies consistently. Periodic audits using tools like kube-bench, kube-hunter and the CNCF’s Multi-Tenancy Benchmarks help verify adherence to security standards. Automated namespace bootstrapping, where quotas, RBAC bindings, NetworkPolicies, service mesh settings and pod security profiles are applied automatically when a new tenant is onboarded, reduces human error and ensures consistent isolation.

3.5.2 Tenant-aware observability and audit logging

In a shared environment, observability must be scoped per tenant to avoid data leakage and to facilitate troubleshooting. Metrics collectors (Prometheus, OpenTelemetry) and log forwarders (FluentBit, Loki) should tag data with namespace or tenant identifiers. Access to dashboards (Grafana, Kibana) and alerts should be controlled via RBAC so tenants see only their

own telemetry. Aggregation layers must enforce that one tenant cannot query another's logs or metrics. Industry guidance recommends sending logs from every node to a central location outside the cluster for persistence and analysis [7]. Similarly, effective guardrails include immutable audit logs stored outside the cluster [8]. Externalizing observability and audit data ensures that critical information is preserved even if the cluster is compromised.

3.5.3 Cloud-agnostic design

To avoid vendor lock-in and promote portability, the architecture should rely on Kubernetes primitives and open-source add-ons rather than provider-specific services. Use standard service accounts, NetworkPolicy, Ingress controllers and egress gateways (e.g., NAT or proxy pods) that work across AWS, GCP, Azure and on-prem installations. External access can be controlled via egress gateways within the cluster rather than relying on cloud-specific networking constructs. When multi-cluster deployments are needed, use the Kubernetes Cluster API or federation to manage cluster fleets uniformly across providers. A cloud-agnostic approach ensures that tenants can run in any environment without re-architecting their applications.

4. DISCUSSION AND FUTURE DIRECTIONS

4.1 Experimental outcomes

To evaluate the impact of the proposed architecture, 10 independent service instances, that were previously deployed across separate clusters, were migrated to a single shared Kubernetes cluster using the multi-tenant design. This transition demonstrated significant improvements in infrastructure efficiency while maintaining strong security and tenant isolation. The table below compares the aggregated resource usage across both configurations.

Table 2: Resource usage comparison of isolated clusters versus shared multi-tenant cluster

	Provisioned CPU (mCPU)	Provisioned Memory (GiB)	Node count
Multiple clusters	20000	128	30
Shared cluster	16000	92	18

The shared cluster configuration achieved a 20 percent reduction in total provisioned CPU and a 28 percent reduction in provisioned memory, along with a 40 percent decrease in node count. Actual used CPU and memory were consistent across both setups, confirming no performance degradation. No tenant isolation failures or interference incidents were observed during a 30-day continuous integration workload across all services. These results confirm that a properly designed multi-tenant Kubernetes architecture can deliver high resource efficiency while upholding strict security and operational boundaries.

4.2 Strengths and limitations

The layered architecture achieves strong logical isolation within a single cluster. Namespaces and RBAC scope resources and access; quotas and limits enforce fair use; network policies

and mTLS secure the data plane; admission control prevents misconfigured workloads; and node isolation accommodates sensitive workloads. Because the approach relies on Kubernetes-native primitives and open-source tools, it is portable across cloud providers and on-premises environments.

However, this design introduces complexity. Service mesh and policy engines require operational expertise and add overhead. Tenants may still share the control plane and parts of the data plane, which may not be acceptable in highly regulated or hostile multi-tenancy scenarios. Strict configuration is necessary; a single misconfigured RBAC binding or network policy can break isolation. Automated policy enforcement and regular audits are critical to maintain security. Additionally, while quotas prevent noisy neighbours, they require careful tuning to match workload patterns. Dedicated node pools improve isolation but increase cost if under-utilized.

4.3 Emerging trends: virtual clusters and advanced isolation

The Kubernetes community is exploring new models of multi-tenancy beyond namespace isolation. Virtual clusters run a per-tenant API server and controller manager inside a namespace of a shared cluster. Each virtual control plane presents the illusion of a separate cluster while delegating the actual scheduling of pods to the host cluster. Research indicates that virtual clusters overcome the limitations of namespace isolation by allowing tenants to perform cluster-level operations while only accessing their own resources [1]. This makes them a potential hard multi-tenancy solution for scenarios where tenants do not trust each other [1]. However, data-plane isolation remains necessary which includes techniques such as pod sandboxing (running each pod in a lightweight virtual machine) provide strong isolation against container escapes [1].

Virtual clusters originated from work by the Kubernetes multi-tenancy working group and were initially implemented as the VirtualCluster project incubated by Kubernetes, with later open-source implementations offering additional features [1]. While these projects demonstrate the viability of per-tenant control planes, they also increase operational complexity and surface area for attack. The approach is still evolving, and standardization within the Kubernetes ecosystem is ongoing. Evaluating virtual clusters requires weighing the benefits of control-plane isolation against the added management overhead and the need for complementary data-plane isolation.

4.4 Recommendations

- Start with a layered approach. For most organizations, namespace-based multi-tenancy with RBAC, quotas and network policies is sufficient. Add a service mesh for mTLS when sensitive data is exchanged. Use admission controllers to enforce policies and avoid misconfigurations.
- Automate everything. Manual configuration leads to errors. Use templates, Helm charts, operators or Terraform to bootstrap namespaces, apply policies and configure observability. Version control all policies and infrastructure code.
- Monitor and audit. Use tenant-aware monitoring and logging to detect anomalies. Perform regular audits using security scanners and verify that policies are enforced. Tune quotas based on observed usage patterns.
- Evaluate virtual clusters carefully. For tenants requiring stronger isolation or custom control-plane extensions, virtual clusters offer an attractive

compromise between dedicated clusters and namespaced multi-tenancy. However, ensure that data-plane isolation (e.g., pod sandboxing) is also applied, and be mindful of the operational complexity [1].

- Plan for growth. As the number of tenants grows, shard workloads across multiple clusters or federate clusters. Use the Kubernetes Cluster API or multi-cluster controllers to manage cluster fleets and apply consistent policies.

5. CONCLUSION

Multi-tenancy on Kubernetes, when properly designed, allows organizations to consolidate workloads without sacrificing security or fairness. Cluster sprawl wastes resources and complicates operations, but naive consolidation invites noisy neighbors and cross-tenant attacks. The layered isolation architecture described in this paper uses namespaces, RBAC, quotas, network segmentation, mTLS, admission control and optional node isolation to achieve strong logical separation within a single cluster.

As Kubernetes continues to evolve, the distinction between multi-tenant and dedicated clusters will blur. Virtual clusters are an emerging solution that offer per-tenant control planes, enabling tenants to perform cluster-level operations without impacting others [1]. Formal verification techniques promise to reduce misconfiguration risks. By combining these developments with robust automation and governance, organizations can achieve scalable tenant isolation and avoid the pitfalls of cluster sprawl.

Future advancements center on strengthening both control-plane and data-plane isolation, improving tenant-aware observability, and reducing the operational complexity associated with multi-tenant architectures. Virtual control planes, lightweight virtual machine sandboxes, workload identity enforcement, fine grained policy engines, and adaptive governance frameworks all contribute to achieving secure and scalable multi-tenant deployments. As these innovations mature, multi-tenant Kubernetes environments evolve into

secure, efficient, and highly automated platforms capable of supporting diverse workloads and untrusted tenants at scale. This progression positions multi tenancy not only as a cost optimization strategy but also as a foundational capability for next generation cloud platforms.

6. REFERENCES

- [1] A. Oliva, Multi-Tenancy in Kubernetes Clusters, Master's thesis, Politecnico di Torino, 2024.
- [2] R. Bezdicek, S. Malik, F. Casciano, A. Tapia, "Three Tenancy Models for Kubernetes," Kubernetes Multi-Tenancy Working Group Blog, April 2021.
- [3] Kubernetes Documentation, "Multi-tenancy Best Practices," Kubernetes v1.34, 2025 (accessed Oct 2025).
- [4] A. Beltre, P. Saha, M. Govindaraju, "KubeSphere: An Approach to Multi-Tenant Fair Scheduling for Kubernetes Clusters," in Proceedings of the 2019 IEEE Cloud Summit.
- [5] R. Molleti, Highly Scalable and Secure Kubernetes Multi-Tenancy Architecture for FinTech, *Journal of Engineering and Applied Sciences Technology* 4 (2) (2022) 1–5.
- [6] H. G. Gowda, "Managing Multi-Tenant Kubernetes Clusters for AEM and HCL Commerce: A Best Practices Study," *International Journal of Novel Research and Development* 8 (8) (2023) 672–684.
- [7] CrowdStrike, "Kubernetes Logging Guide: The Basics," CrowdStrike blog, 2023. This guide recommends collecting logs from every node and sending them to a central location outside the cluster for persistence and analysis.
- [8] A. Robert, "Kubernetes Guardrails: Building Auditing and Accountability for Secure, Reliable Clusters," *hoop.dev*, 2025. The article emphasises that effective guardrails include immutable audit logs stored outside the cluster.