

Rethinking Multi-Tenant Cloud Security: A Zero-Trust Framework for Eliminating Lateral Movement and Identity Abuse

Dr. Eleanor M. Roswell

Global Institute of Systems Engineering, University of Wellington

Abstract

Background: The rapid adoption of cloud computing has driven architectures that support multi-tenancy, elasticity, and heterogeneous workload placement. However, multi-tenant clouds introduce distinctive security, isolation, and placement challenges that affect confidentiality, integrity, and availability for hosted services. Existing literature addresses discrete elements — from placement controls in OpenStack to hypervisor-based intrusion detection and broader zero-trust prescriptions — but there is a need for an integrative framework that aligns placement mechanisms, scheduling policies, storage backends, configuration automation, and adaptive security controls under a single theoretical model. (Karataş et al., 2017; Bogorodskiy, 2019; OpenStack Documentation, 2019; Nikolai & Wang, 2014; Hariharan, 2025).

Objectives: This article develops a comprehensive, publication-ready theoretical research article that synthesizes technical controls and administrative practices into a placement-aware zero-trust model for multi-tenant cloud environments. The aim is to ground each claim in the supplied literature, to elaborate mechanisms in fine-grained detail, and to provide a roadmap for future empirical evaluation and operationalization. (Jackson et al., 2015; Bentley, 2016; Rajesh & Kumar, 2017; Manikyam & Kumar, 2017).

Methods: A conceptual-methodological synthesis is used: (1) systematic mapping of cited multi-tenancy literature; (2) functional decomposition of OpenStack placement primitives and storage backends; (3) threat surface mapping using hypervisor-focused detection techniques; and (4) construction of a layered zero-trust placement control model integrating automation and operational policies. Each methodological step draws directly on provided references and extrapolates logically, avoiding empirical claims beyond cited work. (Karataş et al., 2017; OpenStack Documentation, 2019; Nikolai & Wang, 2014; Bogorodskiy, 2019).

Results: The synthesis yields a placement-aware zero-trust architecture that defines: tenant-aware host aggregates and AZ tagging strategies; scheduler extensions for affinity/anti-affinity informed by security labels; Cinder multi-backend policies aligned with tenant isolation goals; hypervisor-based telemetry hooks for detection; and an operational automation blueprint using Ansible for policy enforcement. The result is a coherent conceptual model that maps threat vectors to placement and orchestration controls. (OpenStack Documentation, 2019; Bogorodskiy, 2019; Bentley, 2016; Jackson et al., 2015; Nikolai & Wang, 2014).

Conclusions: Integrating placement control with zero-trust principles materially strengthens isolation guarantees in multi-tenant clouds and reduces attack surface for lateral movement, co-residency attacks, and storage-based leakage. The article outlines practical implementation steps and identifies measurable research directions—particularly empirical validation of scheduler policy impact and hypervisor-detection efficacy—while acknowledging limitations due to the conceptual nature of the work. (Hariharan, 2025; Karataş et al., 2017; Nikolai & Wang, 2014).

Keywords: multi-tenancy, zero trust, placement control, OpenStack, hypervisor IDS, cloud orchestration, tenant isolation

INTRODUCTION

The architecture of modern cloud platforms inherently supports multiple tenants sharing physical and virtualized resources. This design yields economies of scale and elasticity that power contemporary services, from medical diagnosis pipelines to large-scale analytics stacks, yet it also creates complex interdependencies among resource placement policies, scheduling mechanisms, storage backends, and runtime security controls (Rajesh & Kumar, 2017; Manikyam & Kumar, 2017; Karataş et al., 2017). The technical challenge is not merely to provision resources efficiently, but to ensure that placement and orchestration decisions do not inadvertently increase risk by allowing co-residency or shared backend channels that facilitate information leakage or lateral movement. This article situates these concerns within a unified theoretical framework that fuses placement-awareness with zero-trust security principles, a pressing need identified across multiple referenced works (Hariharan, 2025; Bogorodskiy, 2019; OpenStack Documentation, 2019).

A core dimension of risk in multi-tenant clouds is the co-residency problem: adversaries can exploit the sharing of hypervisors and physical hosts to enumerate neighbors or attempt side-channel attacks, and misconfigured storage backends or compute schedulers can amplify these vectors by placing sensitive workloads on shared infrastructure without adequate isolation (Karataş et al., 2017; Nikolai & Wang, 2014). The OpenStack ecosystem reveals a palette of primitives—host aggregates, availability zones, compute schedulers, and Cinder multi-backend mechanisms—that operators can leverage to express and enforce placement constraints (OpenStack Documentation, 2019). Yet, the mere existence of primitives is insufficient; they must be orchestrated by policy, automation, and detection systems to realize robust tenant isolation (Jackson et al., 2015; Bentley, 2016).

While prior research has systematically mapped multi-tenant architectures and identified recurring patterns (Karataş et al., 2017), and applied intrusion detection at the hypervisor layer as a promising defense (Nikolai & Wang, 2014), there remains a gap: a practical, theoretically grounded synthesis that shows how placement controls, storage backend decisions, scheduler policies, and automated configuration tooling converge to form an operational zero-trust posture in multi-tenant clouds (Bogorodskiy, 2019; Hariharan, 2025). This paper addresses that gap by producing a detailed, theory-forward model that translates resource management primitives into security outcomes, drawing on the provided literature while offering extensive analysis, counter-arguments, and nuanced interpretation.

The contribution is threefold. First, it articulates a placement-aware zero-trust model that explicitly connects scheduler semantics to isolation guarantees. Second, it elaborates a set of operational controls—host aggregate tagging, availability zone partitioning, multi-backend volume assignments, and hypervisor telemetry integration—that cloud administrators can implement using standard OpenStack tooling and automation frameworks. Third, it proposes a structured research agenda to empirically measure the effectiveness of these interventions under realistic workload and adversary models. Each claim is supported by the supplied references and is elaborated to ensure the piece functions as a publication-ready theoretical article appropriate for academic dissemination (Jackson et al., 2015; Bentley, 2016; Rajesh & Kumar, 2017).

METHODOLOGY

This research uses a conceptual synthesis methodology driven entirely by the provided reference corpus. The approach is deliberately text-based and analytic: it performs a cross-cutting analysis of the documented OpenStack primitives, published practical guidance on OpenStack administration, the mapping study of multi-tenancy architectures, hypervisor-level detection techniques, and domain applications (medical CAD systems and big data analytics) that exemplify tenant diversity and workload sensitivity. The methodology has four interlocking stages.

First, literature consolidation catalogs the functional capabilities and constraints described in each reference. The systematic mapping study of multi-tenant architectures provides a taxonomy of tenancy models and security challenges (Karataş et al., 2017). OpenStack Documentation supplies the operational primitives—host aggregates, availability zones, compute schedulers, and Cinder multi-backend features—that are analyzed for their security semantics and programmable controls (OpenStack Documentation, 2019). Practical administration resources explain implementation patterns with automation tools such as Ansible and procedural best practices (Bentley, 2016; Jackson et al., 2015; Bogorodskiy, 2019). Security-focused literature examines hypervisor-based intrusion detection proposals (Nikolai & Wang, 2014) and modern conceptual treatments of zero-trust tailored to multi-tenant clouds (Hariharan, 2025). Applied domain studies (Rajesh & Kumar, 2017; Manikyam & Kumar, 2017) demonstrate the sensitivity of workloads—medical diagnostic systems and big data analytics—and emphasize the operational stakes for placement and isolation.

Second, functional decomposition maps each OpenStack primitive to security-relevant properties. Host aggregates and availability zones are analyzed for their capacity to create administrative and physical boundaries (OpenStack Documentation, 2019). Compute schedulers are decomposed into policy hooks: affinity, anti-affinity, weighted filtering, and custom placement drivers (OpenStack Documentation, 2019). Cinder multi-backend mechanisms are mapped to storage isolation properties, backend performance characteristics, and policy implications for tenant data residency (OpenStack Documentation, 2019). Practically oriented sources add nuance about how these primitives behave in production, including Bare Metal provisioning and placement control use-cases (Bogorodskiy, 2019; Jackson et al., 2015).

Third, threat surface modeling leverages hypervisor-focused IDS literature to link observable telemetry to placement policies. Nikolai and Wang (2014) describe hypervisor-based cloud intrusion detection systems that can monitor VCPU patterns, network flows, and memory access behaviors at the hypervisor level; these detection hooks are incorporated conceptually into the model to provide a detection–response loop that informs dynamic placement and remediation decisions. The zero-trust literature supplies the normative security principles—least privilege, continuous verification, micro-segmentation—that are translated into placement decisions, such as strict anti-affinity for high-sensitivity tenants or hard separation of storage backends (Hariharan, 2025).

Fourth, the synthesis step constructs the placement-aware zero-trust model. This model is not an empirical artifact; rather, it is a prescriptive theoretical architecture that integrates the previously cataloged primitives and security telemetry into a layered control plane. It defines policy specification languages (in conceptual terms), scheduler extensions (by mapping to known compute scheduler hooks), storage assignment policies (using Cinder multi-backend mapping), and automation workflows (informed by Ansible-based administration patterns), and specifies how hypervisor IDS signals should influence placement and remediation. This stage explicitly references administration manuals and practical deployment blog posts to ground the framework in operationally plausible mechanisms (Bentley, 2016; Bogorodskiy, 2019; Jackson et al., 2015).

Throughout the methodology, each major analytic claim is directly tied to one or more of the provided references. No external sources were consulted, and no empirical experimentation was performed; instead, the work constructs an exhaustive theoretical model and operational roadmap that is fully supported by the supplied literature and detailed, critical reasoning.

RESULTS

The conceptual synthesis yields a placement-aware zero-trust architecture articulated as a layered control

model. The result is a coherent mapping from OpenStack primitives and administrative tooling to explicit security outcomes. The following subsections describe the architecture components and their theoretical behavior; each subsection concludes with an operational blueprint grounded in the cited sources.

Host Aggregates and Availability Zones as Logical Isolation Fabrics

Host aggregates and availability zones are administrative constructs intended for grouping compute hosts by attributes such as hardware capabilities, geographic location, or policy label (OpenStack Documentation, 2019). In the model, host aggregates function as the primary containment domain for tenant classification: hosts are tagged with security labels that represent sensitivity tiers, compliance regimes, or tenancy ownership. Availability zones overlay physical or operational segmentation that can be used to enforce geo-residency and fault domain separation. The theoretical contribution is the explicit use of these constructs as zero-trust isolation fabrics: the model prescribes that sensitive tenants be assigned to host aggregates whose hosts share a common trust domain, and that availability zones be used to create orthogonal separation for redundancy without crossing security boundaries. This mapping directly draws on OpenStack primitives and is operationalized by host tagging and scheduler constraints (OpenStack Documentation, 2019).

Operational blueprint: Define a set of security labels (e.g., "high-confidentiality", "regulated", "general-purpose"), implement host aggregates for each label, and map each tenant to the appropriate aggregate. Use availability zones to separate failure domains and ensure that redundancy planning does not force replication across security labels. The documentation and blog posts indicate that these primitives are supported and commonly used in practice, making this blueprint implementable (OpenStack Documentation, 2019; Bogorodskiy, 2019).

Scheduler Extensions and Affinity Policy Semantics

Compute schedulers determine placement decisions through a filter-and-weight model, where filters eliminate hosts and weights prioritize remaining choices (OpenStack Documentation, 2019). The results articulate how scheduler semantics can be extended into security policy by defining affinity (co-location) and anti-affinity (separation) constraints based on tenant sensitivity and threat models. Affinity constraints are appropriate when performance or data locality is paramount, but they increase co-residency risk for sensitive tenants; anti-affinity constraints mitigate co-residency risk by dispersing instances across distinct aggregates or AZs. The model further suggests introducing security-aware weights that penalize hosts with mixed-tenancy histories or recent anomalous hypervisor signals. This conceptual extension is grounded in the compute scheduler documentation and placement control discussions (OpenStack Documentation, 2019; Bogorodskiy, 2019).

Operational blueprint: Implement scheduler policies that accept security labels as input, incorporate anti-affinity rules for sensitive workloads, and integrate a feedback loop where hypervisor IDS signals can dynamically adjust scheduler weights to avoid hosts exhibiting suspicious behavior. Packt guides and administration books demonstrate the feasibility of customizing schedulers and using placement drivers, providing practical foundations for this blueprint (Jackson et al., 2015; Bentley, 2016).

Cinder Multi-Backend and Storage Isolation Policies

Persistent storage represents a major attack surface; misallocated or shared backends can create lateral channels for data leakage. OpenStack's Cinder supports multiple backends, allowing operators to map volume types to specific storage backends (OpenStack Documentation, 2019). The model prescribes a storage segregation strategy: sensitive tenants are allocated to dedicated backends or logically separated pools, with explicit volume-type policies preventing cross-backend attachment or snapshot sharing. This strategy aligns

storage residency with compute placement to avoid cases where a compute host in a mixed-trust aggregate accesses volumes from an insecure backend. The result is a storage placement policy that complements host placement and scheduler constraints (OpenStack Documentation, 2019).

Operational blueprint: Establish volume types that map to isolated storage backends, configure quota and access controls to prevent snapshot-based data leakage, and ensure orchestration logic ties volume assignment to the compute host aggregate label. Jackson et al. and OpenStack documentation provide configuration approaches that operationalize these recommendations (Jackson et al., 2015; OpenStack Documentation, 2019).

Hypervisor-Based Telemetry and Intrusion Detection Integration

Hypervisor-level detection provides visibility into tenant behavior and host-level anomalies that are invisible to guest agents (Nikolai & Wang, 2014). The model integrates hypervisor IDS signals (e.g., unusual VCPU scheduling patterns, memory anomalies, or cross-VM side-channel indicators) into a control loop that influences placement decisions and automated remediation. Specifically, hypervisor alerts can trigger dynamic anti-affinity enforcement, live migration off suspect hosts, or temporary host quarantines. This integration bridges detection and orchestration, enabling continuous verification consistent with zero-trust principles (Nikolai & Wang, 2014; Hariharan, 2025).

Operational blueprint: Deploy hypervisor IDS probes (conceptually per Nikolai & Wang) and ensure their outputs feed a decision engine that can orchestrate live migration or quarantine workflows using standard compute APIs. Ansible orchestration and administrative playbooks—documented in practitioner guides—facilitate automating these remediation actions (Bentley, 2016; Jackson et al., 2015).

Automation and Policy Enforcement Using Configuration Management

The model emphasizes automation as the mechanism for enforcing placement and security policies at scale. Ansible, in particular, embodies an approach to declarative configuration and automation that allows repeatable provisioning of aggregates, scheduler policies, storage backends, and remediation playbooks (Bentley, 2016). The key result here is that automation reduces human error, enforces policy consistency, and enables rapid response to hypervisor IDS signals through predefined playbooks that implement live migration, host re-tagging, or backend reassignment. Documentation and administrative texts provide patterns and examples showing how automation can be integrated with OpenStack APIs to achieve these outcomes (Bentley, 2016; Jackson et al., 2015).

Operational blueprint: Construct a policy-as-code repository where security labels, scheduler constraints, volume type mappings, and remediation playbooks are defined as version-controlled artifacts. Use Ansible playbooks to enact policy changes and to perform forensic data capture, live migrations, or host quarantines as necessary.

Application to Sensitive Workloads: Medical CAD and Big Data Analytics

The model applies directly to domains where workload sensitivity varies significantly. Medical computer-aided diagnosis systems, which process protected health information, require strict placement and storage isolation to comply with regulatory and ethical mandates (Rajesh & Kumar, 2017). Big data analytics pipelines, by contrast, may have mixed sensitivity and high storage throughput requirements, necessitating careful Cinder backend selection and scheduler weighting to match performance without sacrificing isolation (Manikyam & Kumar, 2017). The result is an application-driven mapping that demonstrates the model's ability

to reconcile performance and security trade-offs by using host aggregates, anti-affinity, and dedicated storage backends (Rajesh & Kumar, 2017; Manikyam & Kumar, 2017).

Operational blueprint: For medical CAD workloads, assign compute and storage exclusively to "regulated" aggregates and backends, enforce strict anti-affinity for management-plane components, and integrate hypervisor IDS telemetry to detect possible exfiltration. For analytics workloads, use performance-optimized backends but segregate sensitive datasets into isolated volume types and enforce network micro-segmentation at the tenant level.

Synthesis: The Placement-Aware Zero-Trust Control Plane

Combining the above components yields a layered control plane: policy specification (security labels, volume types), enforcement (host aggregates, availability zones, scheduler constraints), detection (hypervisor IDS), and automation (Ansible playbooks). The outcome is a continuous verification loop consistent with zero-trust thinking: the platform never assumes host trustworthiness, continuously evaluates telemetry, and adjusts placements or applies remediation as needed (Hariharan, 2025; Nikolai & Wang, 2014). This theoretical synthesis demonstrates a feasible, reference model for operators seeking to harden multi-tenant clouds using OpenStack primitives and common administration practices (OpenStack Documentation, 2019; Jackson et al., 2015; Bentley, 2016).

DISCUSSION

The placement-aware zero-trust model proposed synthesizes primitives and security concepts into an operationally realistic framework. This Discussion examines key theoretical implications, explores counter-arguments, identifies limitations, and outlines future empirical research directions.

Theoretical interpretation and implications

At a theoretical level, the model reframes placement as a security control rather than purely an efficiency metric. Historically, schedulers and storage policies were designed with performance, utilization, and fault tolerance in mind, with security seen as an overlay (OpenStack Documentation, 2019; Jackson et al., 2015). This separation has cognitive and practical costs: security constraints expressed late in the orchestration pipeline can lead to ad hoc exceptions and inconsistent enforcement. By embedding security labels and telemetry-driven scheduler semantics into the placement engine itself, the model aligns resource management with continuous verification, thus operationalizing zero trust in an infrastructure context (Hariharan, 2025).

This reframing has several implications. First, it requires schedulers to be extensible and to accept security-relevant inputs; this demands enhancements to placement drivers and policy engines—extensions that are conceptually grounded in the compute scheduler documentation but require concrete engineering to realize (OpenStack Documentation, 2019). Second, it elevates storage backends as first-order security controls rather than passive performance resources; mapping volume types to security requirements ensures that storage policy is no longer decoupled from compute placement (OpenStack Documentation, 2019). Third, it positions hypervisor telemetry not as an afterthought but as an integral signal in the control loop, transforming detection into a driver for remediation and placement adaptation (Nikolai & Wang, 2014).

Counter-arguments and nuanced analysis

There are counter-arguments that deserve careful consideration. One view is that strict anti-affinity, dedicated host aggregates, and storage segregation inherently reduce cloud elasticity and increase operational cost,

undermining the economic rationale of cloud computing (Karataş et al., 2017). In response, the model accepts that security is a trade-off with efficiency and argues for tiered approaches: not all tenants require highest isolation; rather, sensitivity labeling should determine where strict controls are necessary. For sensitive workloads, the marginal cost of reserved resources may be justified by reduced risk and compliance burdens (Rajesh & Kumar, 2017). For lower-sensitivity tenants, relaxed constraints preserve elasticity.

Another counter-argument concerns the reliability of hypervisor-based IDS outputs. False positives can cause unnecessary migrations and instability, while false negatives reduce the efficacy of the control loop (Nikolai & Wang, 2014). The model addresses this by emphasizing that hypervisor signals should be one input among many—correlated with network telemetry, guest-level logs, and administrative policies—before automated remediation is enacted. Additionally, tiered remediation policies (alerting, increased monitoring, then migration) can reduce the operational impact of detection inaccuracies.

A further critique involves the operational complexity of policy-as-code and automation. Maintaining a living repository of placement and security policies requires governance, change control, and skilled administrators; without these, automation can propagate misconfigurations at scale (Bentley, 2016; Jackson et al., 2015). The model recognizes this risk and recommends organizational practices—version control, peer review, and staged deployments—that mirror software engineering best practices to ensure policy integrity.

Limitations of the present work

This article is intentionally theoretical and synthesizes existing literature rather than presenting new empirical measurements. As such, it has limitations. The efficacy of the proposed scheduler extensions, the overhead of enforcement mechanisms, and the performance impact of strict anti-affinity were not measured in practice. Moreover, the hypervisor IDS techniques referenced provide promising detection vectors, but their real-world detection accuracy and operational stability under high-load conditions remain to be empirically validated (Nikolai & Wang, 2014). The conceptual mapping between storage backends and isolation outcomes presumes correct backend configuration and isolation guarantees; in practice, storage hardware and vendor-specific behaviors introduce complexities that must be assessed empirically (OpenStack Documentation, 2019).

Furthermore, the references include practitioner-oriented documentation and blog posts that capture current operational practice (Bogorodskiy, 2019; Jackson et al., 2015; Bentley, 2016). While these are invaluable for operational grounding, they do not substitute for rigorous experimental evaluation. The article therefore presents a prescriptive architecture that invites operational pilots and controlled experiments to quantify trade-offs and refine policy heuristics.

Future research directions and empirical agenda

Several priority research directions emerge from the synthesis. First, empirical evaluation of scheduler policies: controlled experiments comparing default schedulers against security-aware scheduler extensions would quantify impacts on co-residency risk, provisioning latency, and resource utilization. Second, hypervisor IDS validation: studies that benchmark detection capabilities under adversarial workloads, noise, and varying host densities would clarify the operational utility of hypervisor telemetry. Third, storage isolation verification: testing Cinder multi-backend configurations for cross-backend leakage, snapshot capture isolation, and performance interference would validate theoretical mapping. Fourth, cost-benefit analysis: modeling and measurement studies that quantify the economic trade-offs of reserved host aggregates and dedicated storage for sensitive tenants would inform policy decisions. Finally, governance and human factors research into policy-as-code lifecycle management would address the organizational challenges of

maintaining robust automation pipelines (Rajesh & Kumar, 2017; Manikyam & Kumar, 2017; Bentley, 2016).

Operational recommendations and phased adoption pathway

Practitioners seeking to adopt the model should follow a phased approach. Phase one focuses on discovery and labeling: inventory hosts, classify workloads by sensitivity, and create a taxonomy of volume types and backends. Phase two implements non-disruptive enforcement: create host aggregates, define availability zones, and set soft scheduler constraints (weights) that favor initial segregation. Phase three introduces telemetry and correlation: deploy hypervisor IDS probes in passive mode, correlate signals with existing monitoring, and verify alert quality. Phase four automates remediation: author Ansible playbooks for migration and quarantine, test in staging, and roll out with staged governance. Phase five optimizes cost and performance: evaluate utilization and adjust policies to balance security and efficiency. This phased pathway reflects operational best practices described by administration literature and practical blog guidance (Jackson et al., 2015; Bentley, 2016; Bogorodskiy, 2019).

Ethical and compliance considerations

The model intersects with regulatory regimes that govern data residency and health data protection for domains like medical CAD systems. Assigning medical workloads to dedicated aggregates supports compliance requirements and ethical responsibilities to safeguard patient data (Rajesh & Kumar, 2017). However, strict isolation must be balanced against the need for collaborative research and analytics; appropriately governed temporary access mechanisms and robust audit trails should be part of the deployment to maintain accountability while enabling legitimate data use.

CONCLUSION

This article presents a placement-aware zero-trust framework tailored to multi-tenant cloud environments, grounded entirely in the provided literature. By synthesizing OpenStack placement primitives, scheduler semantics, Cinder multi-backend storage, hypervisor-based intrusion detection, and automation practices, the model offers a theoretically coherent and operationally plausible architecture for improving tenant isolation and reducing co-residency risks. The framework reframes placement as a proactive security control, integrates continuous verification with telemetry-driven remediation, and emphasizes policy-as-code for scalable enforcement. While conceptual, the model draws directly from the referenced works and offers a structured research agenda to empirically validate and refine its prescriptions. Practical adoption requires careful governance, staged deployment, and empirical evaluation to balance security gains against operational cost and complexity. The references provided constitute a robust foundation for both initial implementation and subsequent rigorous experimentation to quantify the model's security and performance trade-offs.

REFERENCES

1. T. Rajesh and Dr. S. Mohan Kumar. (2017). Medical Diagnosis Cad System Using Latest Technologies, Sensors and Cloud Computing. International Journal of Computer Engineering & Technology, 8(1), pp. 43–50.
2. Naga Raju Hari Manikyam and Dr. S. Mohan Kumar. (2017). Methods and Techniques To Deal with Big Data Analytics and Challenges In Cloud Computing Environment. International Journal of Civil Engineering and Technology, 8(4), pp. 669-678.
3. Karataş, Gözde, et al. (2017). "Multi-Tenant architectures in the cloud: a systematic mapping study."

2017 International Artificial Intelligence and Data Processing Symposium (IDAP). IEEE.

4. Hariharan, R. (2025). Zero trust security in multi-tenant cloud environments. *Journal of Information Systems Engineering and Management*, 10.
5. Nikolai, Jason, and Yong Wang. (2014). "Hypervisor-based cloud intrusion detection system." 2014 International Conference on Computing, Networking and Communications (ICNC). IEEE.
6. Bentley, Walter. (2016). *OpenStack Administration with Ansible*. Packt Publishing Ltd.
7. Roman Bogorodskiy. (2019). Placement control and multi-tenancy isolation with OpenStack Cloud: Bare Metal Provisioning, Part 2, Mirantis Blog.
8. Jackson, Kevin, Cody Bunch, and Egle Sigler. (2015). *OpenStack cloud computing cookbook*. Packt Publishing Ltd.
9. OpenStack Documentation. (2019). Host Aggregates, Availability Zones (AZs).
10. OpenStack Documentation. (2019). Compute schedulers.
11. OpenStack Documentation. (2019). Cinder-multi-backend.
12. OpenStack Documentation. (2019). Manage volumes and volume types.