

A Systematic Review of Machine Learning and Deep Learning Approaches for Malware Detection in Cloud Computing Environments: Taxonomy, Architectures, and Open Challenges

^{1*}Sanaboyina. Madhusudhana Rao, ^{2*} Arpit Jain

^{1*,2*}Department of Computer Science and Engineering,

^{1*,2*}Koneru Lakshmaiah Education Foundation, Vaddeswaram 522302, Andhra Pradesh, India.

^{1*}smadhusudhan780@gmail.com ^{2*}drjainarpit@gmail.com

Abstract: *Cloud computing has become the backbone of modern digital infrastructure, hosting data, applications, and services for enterprises, governments, and individuals at unprecedented scale. This same concentration of computational and informational value, however, makes cloud platforms a high-priority target for malware authors, whose techniques now include polymorphism, metamorphism, fileless execution, multi-stage payload delivery, and adversarial evasion of learning-based detectors. Traditional signature-driven antivirus engines, designed for a single endpoint, cannot keep pace with the volume, velocity, and variety of malware encountered in multi-tenant virtualized and containerized environments. This paper presents a systematic review of machine learning (ML) and deep learning (DL) approaches for malware analysis and detection in cloud computing environments. We synthesize the literature into a unified taxonomy spanning analysis paradigms (static, dynamic, hybrid, and memory-introspection-based), feature representations (opcode and byte sequences, application programming interface and system-call traces, network telemetry, resource-utilization signals, and image-based encodings), and learning models (classical classifiers, deep neural architectures, ensembles, and federated learning). We further organize cloud-specific detection architectures into client-server scanning, hypervisor-level virtual machine introspection, container and orchestration monitoring, and edge-cloud collaborative pipelines. A comparative analysis of representative studies is provided, alongside a consolidated review of public datasets and evaluation metrics. Building on the gaps identified, we propose a layered conceptual framework for scalable, privacy-preserving, and explainable malware detection that operates at both the virtual-machine and service levels. Finally, we discuss persistent challenges—concept drift, class imbalance, adversarial robustness, non-deterministic replay, interoperability, and the resource overhead of detection—and outline future research directions, including self-healing detection, continual learning, and the integration of large language models for behavioral reasoning. This review is intended to help researchers position new contributions and to inform the design of next-generation, high-quality-of-service malware detection systems for the cloud.*

Keywords: *malware detection, cloud computing, machine learning, deep learning, virtual machine introspection, intrusion detection, container security, adversarial robustness, feature engineering, cyber-physical systems*

I. INTRODUCTION

The migration of computing workloads from on-premises hardware to elastic, on-demand cloud platforms is one of the defining technological transitions of the past two decades. Infrastructure-as-a-Service, Platform-as-a-Service, and

Software-as-a-Service models allow organizations to provision storage, compute, and networking resources within minutes, to pay only for what they consume, and to scale capacity in response to demand without the capital expense of owning physical machines. The same elasticity, multi-tenancy, and remote accessibility that make the cloud attractive to legitimate users also expand the attack surface available to adversaries. A single misconfigured storage bucket, a compromised tenant virtual machine, or a malicious container image can become a foothold from which malware spreads laterally, exfiltrates data, hijacks compute for cryptomining, or stages ransomware.

Malware—software written to compromise the confidentiality, integrity, or availability of a system—remains the most persistent vector of cloud compromise. The economics of malicious software have shifted from isolated nuisance programs toward an industrialized ecosystem in which payloads are generated, packed, and mutated automatically. Hundreds of thousands of previously unseen samples are observed daily, and modern families increasingly employ obfuscation, encryption, and runtime polymorphism specifically to defeat detection. Because each new variant may differ in its byte-level signature while preserving its malicious behavior, the signature-matching paradigm that underpinned classical antivirus products has become structurally inadequate. A detector that can only recognize what it has already catalogued is, by construction, blind to the next variant.

Machine learning offers a fundamentally different premise: rather than memorizing the appearance of known threats, a learned model generalizes from patterns in training data to flag previously unseen samples that share malicious characteristics. Deep learning extends this premise by automating feature extraction, learning hierarchical representations directly from raw or lightly processed inputs. When such models are deployed in the cloud, they benefit from abundant computation, large shared threat databases, and the ability to analyze a sample once on behalf of many tenants. Cloud-based detection can therefore offload heavy analysis from resource-constrained endpoints—mobile devices, Internet-of-Things (IoT) sensors, and cyber-physical systems (CPS)—while improving detection coverage. Yet moving detection into the cloud introduces its own difficulties: tenants are reluctant to surrender sensitive files for inspection, faithful behavioral replay of a sample inside a virtual environment is hard to guarantee, real-time monitoring competes with service-level objectives, and adversaries actively probe and poison learning pipelines.

This paper provides a systematic review of how ML and DL are applied to malware analysis and detection specifically within cloud computing environments. Our aim is threefold. First, we construct an integrated taxonomy that relates analysis paradigms, feature representations, and learning models, so that a given study can be located precisely within the design space. Second, we organize the cloud-specific architectures through which detection is delivered, from client-server scanning to hypervisor introspection, container monitoring, and edge-cloud collaboration, and we compare representative works along consistent dimensions. Third, drawing on the limitations surfaced by this synthesis, we propose a layered conceptual framework for malware detection that is scalable, privacy-preserving, explainable, and capable of operating at both the virtual-machine and service levels.

The remainder of the paper is structured as follows. Section 2 establishes background on cloud deployment and service models, the cloud threat landscape, and a taxonomy of malware. Section 3 surveys malware analysis techniques. Section 4 examines feature engineering, the bridge between raw artifacts and learning models. Section 5 reviews classical and deep learning models, ensembles, and federated learning. Section 6 organizes cloud-specific detection architectures. Section 7 consolidates datasets and evaluation metrics. Section 8 presents a comparative analysis of representative studies. Section 9 introduces the proposed framework. Section 10 discusses open challenges, Section 11 outlines future directions, and Section 12 concludes.

1.1 Scope and review methodology

This review follows a structured, narrative-systematic methodology. We scoped the survey to peer-reviewed work addressing the intersection of three concepts—malware detection, machine or deep learning, and cloud, virtualized, container, edge, or IoT-cloud environments—giving precedence to studies that either deploy detection in the cloud or address constraints unique to it, such as multi-tenancy, hypervisor visibility, and elastic scaling. Works were grouped thematically rather than merely chronologically, and each was characterized along a common set of attributes: the

analysis paradigm employed, the feature representation, the learning model, the deployment architecture, the evaluation dataset, and the reported limitations. The synthesis emphasizes the relationships among these attributes, because the central observation of this review is that detection performance in the cloud is governed less by the choice of any single classifier than by the coherence of the analysis-feature-model-architecture pipeline as a whole.

II. BACKGROUND AND PRELIMINARIES

2.1 Cloud deployment and service models

Cloud computing is conventionally described along two orthogonal axes: deployment model and service model. The deployment axis distinguishes public clouds, in which resources are owned by a provider and shared across unrelated tenants; private clouds, dedicated to a single organization and often hosted on-premises or in a colocation facility; hybrid clouds, which combine the two and orchestrate workloads across them; and community clouds, shared by organizations with common requirements. The service axis distinguishes Infrastructure-as-a-Service (IaaS), which exposes virtual machines, storage, and networking; Platform-as-a-Service (PaaS), which abstracts the operating system and offers managed runtimes; and Software-as-a-Service (SaaS), which delivers complete applications over the network. Each model reallocates the boundary of responsibility between provider and tenant, and each therefore shifts where malware can take hold and who is positioned to detect it.

From a security standpoint, the service model determines the visibility available to a detector. In IaaS, the tenant controls the guest operating system, so detection agents can run inside the virtual machine, while the provider retains hypervisor-level visibility that the tenant cannot see. In PaaS and SaaS, the tenant surrenders progressively more of the stack, and detection responsibility concentrates with the provider. Multi-tenancy compounds the problem: physical hosts are shared, side channels may leak information across tenant boundaries, and a compromise in one tenant's workload can, in principle, threaten the isolation guarantees on which neighbors depend. Effective cloud malware detection must therefore reason explicitly about which layer—application, guest OS, hypervisor, container runtime, or network fabric—affords the observations it relies upon.

2.2 The cloud threat landscape

The threats facing cloud platforms extend well beyond classical executable infection. Cryptojacking malware silently appropriates tenant compute to mine cryptocurrency, inflating bills and degrading performance while attempting to remain below detection thresholds. Ransomware encrypts data in storage volumes and demands payment, with cloud-hosted backups increasingly targeted to prevent recovery. Fileless malware executes entirely in memory, leveraging legitimate system tools and leaving little on disk for static scanners to find. Supply-chain attacks insert malicious code into container images, libraries, or machine-readable artifacts that are then deployed at scale. Within orchestration platforms such as Kubernetes, attacks on pods, service accounts, and the control plane provide avenues for privilege escalation and persistence. Across all of these, the defining cloud-specific characteristics are scale—the sheer number of workloads to monitor—and ephemerality—the fact that virtual machines and containers may exist for only minutes, leaving narrow windows for observation.

Adversaries also adapt to the presence of learning-based defenses. Evasion attacks craft inputs that preserve malicious functionality while shifting the model's prediction toward benign. Poisoning attacks corrupt training data so that the deployed model misclassifies attacker-chosen samples. Model-extraction and membership-inference attacks probe a deployed detector to steal its decision boundary or to infer properties of its training data. A cloud detector that exposes an inference endpoint to many tenants is, by design, more queryable—and therefore more exposed to such attacks—than an isolated endpoint scanner.

2.3 A taxonomy of malware

Although malware families number in the tens of thousands, they can be organized by propagation mechanism and intent. The principal categories relevant to cloud environments are summarized in Table 1. A practical complication is

that contemporary samples are frequently multi-class: a single payload may combine worm-like propagation, rootkit-style concealment, and ransomware behavior, while wrapping all of it in an obfuscation layer. Classification is consequently better understood as assigning a sample to a set of behavioral attributes than to a single mutually exclusive label.

Table 1. Principal malware categories relevant to cloud environments and their defining behaviors.

Category	Defining behavior and cloud relevance
Virus	Attaches to host files and executes when the host runs; well catalogued, but still used as a delivery stage.
Worm	Self-propagates across networks without user action; in the cloud, exploits exposed services to spread laterally across tenants and instances.
Trojan	Masquerades as legitimate software while performing covert actions such as data exfiltration; common in malicious container images and packages.
Backdoor	Establishes covert, persistent remote access by bypassing authentication; used to retain footholds across ephemeral instances.
Rootkit	Subverts the operating system or hypervisor to hide its presence; defeats in-guest agents, motivating introspection-based detection.
Ransomware	Encrypts data and demands payment; increasingly targets cloud storage and backups to block recovery.
Cryptojacker	Hijacks compute to mine cryptocurrency; betrayed by anomalous CPU, memory, and power signatures rather than file signatures.
Fileless / LOLBin	Executes in memory using legitimate binaries; evades disk-based scanning and demands behavioral and memory analysis.
Obfuscated / packed	Uses encryption, polymorphism, or metamorphism to alter appearance while preserving behavior; defeats static signatures.

The detection implication of this taxonomy is direct. File-resident, well-catalogued malware can still be caught by signatures, but the categories that matter most in the cloud—fileless execution, obfuscated payloads, cryptojacking, and rootkits—are precisely those that signatures handle worst. Each of these is better revealed by behavior, by memory state, or by resource-utilization anomalies than by static appearance, which is the central reason that learning-based and introspection-based methods have become the focus of cloud malware research.

2.4 Securing IoT and cyber-physical systems through the cloud

A growing share of cloud security workloads concerns devices that cannot defend themselves. The Internet of Things and cyber-physical systems connect vast populations of low-power sensors, controllers, and actuators that lack the memory, compute, and energy budget to run conventional security software, yet whose compromise can have physical consequences in factories, vehicles, power grids, and medical settings. The cloud is well positioned to protect these devices precisely because it can host the heavyweight analysis that the devices cannot, retrieving the appropriate detector to an edge node near the device and returning verdicts with low latency. The rapid arrival of new and variant IoT malware, however, makes static, signature-based protection particularly fragile in this setting, because the population of devices and firmware is heterogeneous and the variants appear faster than signatures can be authored. Behavioral and anomaly-based detection, trained on what normal device behavior looks like and deployed through edge-cloud collaboration, is therefore the dominant direction for IoT and CPS protection, and it shapes several of the architectures discussed later in this paper. The defining requirement is to push inference close to the device for latency

and resilience while keeping training and global threat intelligence in the cloud, an arrangement that recurs throughout the literature on securing connected systems.

III. MALWARE ANALYSIS TECHNIQUES

Before a learning model can classify a sample, the sample must be analyzed to produce observable artifacts. The analysis paradigm chosen determines both what the model can see and how robust the resulting detector is to evasion. Four paradigms dominate the literature: static analysis, dynamic analysis, hybrid analysis, and memory-introspection-based analysis.

3.1 Static analysis

Static analysis inspects a sample without executing it. Artifacts include the byte sequence of the binary, header metadata from the executable format, imported and exported functions, embedded strings, the control-flow and call graphs recovered by disassembly, and opcode n-grams. Static analysis is fast, safe, and complete in the sense that it can in principle examine all code paths rather than only those exercised at runtime. Its decisive weakness is fragility under obfuscation: packing, encryption, and polymorphism transform the static appearance of a sample while preserving its behavior, and runtime-generated or downloaded code is invisible to a purely static view. In the cloud, static analysis is attractive because it imposes no execution cost and can be applied at ingestion time—scanning container images in a registry, for example—but it cannot stand alone against modern evasive families.

3.2 Dynamic analysis

Dynamic analysis executes the sample in an instrumented environment—typically a sandbox or a disposable virtual machine—and records its behavior. Observable behaviors include system and API calls, file-system and registry modifications, spawned processes, network connections, and memory allocation patterns. Because behavior is comparatively invariant to obfuscation—a packed sample must eventually unpack itself to run, and a ransomware payload must eventually touch files to encrypt them—dynamic analysis is more resilient to evasion than static analysis. Its costs are equally real: execution is slow and resource-intensive, only the paths actually triggered are observed, and sophisticated malware employs sandbox-detection logic, time bombs, and environment checks to suppress malicious behavior when it suspects it is being watched. Cloud sandboxes must therefore be made convincingly realistic, and the non-determinism of replaying concurrent, communicating programs—where serializing thread access to achieve deterministic replay can mask the very exception that constitutes the attack—remains a deep difficulty.

3.3 Hybrid analysis

Hybrid analysis combines static and dynamic artifacts to exploit the complementary strengths of each. A common pattern uses cheap static features as a first-pass triage and reserves expensive dynamic analysis for samples that static analysis cannot confidently classify, which controls cost at cloud scale. Another pattern fuses static and dynamic feature vectors into a single representation on which a model is trained, so that the model can learn correlations—such as a packed static appearance combined with file-encryption behavior—that neither view reveals alone. Empirically, hybrid approaches tend to achieve higher detection rates and lower false-positive rates than either paradigm in isolation, at the price of greater pipeline complexity.

3.4 Memory forensics and virtual machine introspection

Memory forensics analyzes the volatile state of a running system to uncover artifacts—injecting code, hooked function tables, hidden processes, and decrypted payloads—that never appear on disk and may be deliberately concealed from in-guest tools. Virtual machine introspection (VMI) is the cloud-native realization of this idea: the hypervisor observes the memory and execution of a guest virtual machine from the outside, beyond the reach of malware running inside the guest. Because a rootkit that controls the guest operating system can deceive any agent that also runs in the guest, the

out-of-band vantage point of VMI is uniquely valuable against stealthy and kernel-level threats. The price of this vantage point is the semantic gap—the difficulty of reconstructing high-level operating-system abstractions, such as the list of running processes, from raw memory observed at the hypervisor—together with the performance overhead of fine-grained monitoring. VMI is consequently a centerpiece of cloud-specific detection research, frequently paired with learning models that classify the reconstructed behavioral state.

IV. FEATURE ENGINEERING FOR CLOUD MALWARE DETECTION

Feature engineering is the bridge between raw artifacts and learning models, and in practice it is the single most influential determinant of detector quality. The literature has converged on several families of representation, each capturing a different facet of a sample.

4.1 Static and structural features

Static representations include byte and opcode n-grams, which capture local sequential structure; executable-header fields, which summarize how a binary is built; printable strings, which often reveal command-and-control endpoints or ransom notes; and graph-based features derived from control-flow graphs, call graphs, and API-call graphs, which encode the program's structure rather than its surface bytes. Graph representations are particularly valuable because two samples that share malicious logic tend to share graph structure even when their byte sequences differ, which makes graph features more robust to superficial mutation than raw n-grams.

4.2 Dynamic and behavioral features

Behavioral representations are extracted during execution and include sequences and frequency profiles of system and API calls, file and registry operations, process-creation trees, and network-flow descriptors. System-call sequences are especially informative in cloud and container settings because they sit at the boundary between application and kernel and are difficult for malware to avoid generating. Sequence models can exploit the temporal ordering of these calls, while frequency-based summaries trade temporal detail for compactness and speed.

4.3 Resource-utilization and telemetry features

A representation distinctive to the cloud treats the sample not as a file but as a tenant whose resource consumption can be monitored from the outside. Time series of CPU, memory, disk, and network utilization—readily available to a provider through hypervisor counters—reveal cryptojacking, denial-of-service participation, and encryption bursts as characteristic anomalies, without ever inspecting tenant data. Because these signals are collected at the infrastructure layer, they sidestep the tenant-privacy objections that obstruct file-level inspection, which makes them an appealing basis for provider-side detection.

4.4 Image-based representations

A widely adopted technique reinterprets a binary as an image: bytes are mapped to pixel intensities and arranged into a two-dimensional grid, so that malware from the same family produces visually similar textures. This encoding allows convolutional neural networks—mature, well-optimized architectures from computer vision—to be applied directly to malware classification, and it transfers naturally to the cloud, where GPU acceleration is plentiful. Variants encode API-call graphs, opcode streams, or network traffic as images, and the approach has proven effective for family classification, though it inherits the static paradigm's vulnerability to obfuscation that alters byte layout.

4.5 Feature selection and dimensionality reduction

Cloud-scale feature spaces are large, sparse, and laden with redundant or irrelevant attributes that inflate training time, raise inference latency, and degrade accuracy. Feature selection and dimensionality reduction are therefore not optional refinements but core requirements. Filter methods rank features by statistical association with the label; wrapper

methods evaluate feature subsets against a model's performance; and embedded methods select features as a by-product of training, as with tree-based importance or sparsity-inducing regularization. Metaheuristic search—genetic algorithms, particle-swarm optimization, and related techniques—has been used to navigate the combinatorial space of feature subsets efficiently. Reducing the number of features lowers the volume of data exchanged between tenant and cloud, shortens training and inference, and frequently improves generalization by removing noise. The recurring conclusion across the literature is that disciplined feature reduction is one of the most reliable levers for improving both the accuracy and the cost profile of cloud malware detectors.

V. MACHINE LEARNING AND DEEP LEARNING MODELS

Given a feature representation, the learning model maps features to a verdict. The literature spans classical classifiers, deep neural architectures, ensembles, and increasingly federated configurations suited to privacy-sensitive cloud deployment.

5.1 Classical machine learning

Classical classifiers remain widely used because they are interpretable, inexpensive to train, and effective on well-engineered features. Decision trees and random forests handle heterogeneous, high-dimensional tabular features and expose feature importances that aid analyst trust. Support vector machines perform well on high-dimensional, sparse representations such as n-gram vectors. Gradient-boosted ensembles frequently top leaderboards on tabular malware features. Naive Bayes and k-nearest-neighbor classifiers provide fast baselines. The principal limitation of classical methods is their dependence on manual feature engineering: their ceiling is set by the quality of the features supplied to them, and they struggle when the discriminative signal lies in complex sequential or structural patterns that hand-crafted features fail to capture.

5.2 Deep learning

Deep learning automates representation learning, discovering discriminative features directly from raw or lightly processed inputs. Convolutional neural networks excel on image-encoded binaries and on spatially structured features. Recurrent architectures, including long short-term memory and gated recurrent networks, model the temporal dependencies in API-call and system-call sequences and in resource-utilization time series; recurrent models monitoring runtime metrics have proven effective for online behavioral detection in cloud virtual machines. Autoencoders support unsupervised anomaly detection, flagging deviations from learned normal behavior and thereby catching zero-day threats without labeled attack data. Graph neural networks operate directly on call and control-flow graphs. Attention mechanisms and transformer architectures, by weighting the most informative elements of a sequence, have improved both accuracy and interpretability. The cost of deep learning is its appetite for labeled data and computation and its comparative opacity, both of which are partially mitigated by the cloud's abundant resources and partially aggravated by the cloud's adversarial exposure.

5.3 Ensemble learning

Ensembles combine multiple base learners to achieve accuracy and robustness beyond any single model, through bagging, boosting, or stacking. In the cloud, ensembles are doubly attractive because diverse base models are harder for an adversary to evade simultaneously and because heterogeneous learners can specialize in different malware families. The countervailing concern is cost: large ensembles impose memory, computation, and data-transfer burdens that are poorly suited to big-data cloud settings, which has motivated ensemble-pruning techniques that retain a compact, high-performing subset of base classifiers through consensus mechanisms. The design tension is therefore between the accuracy gains of diversity and the resource discipline that cloud economics demand.

5.4 Federated and privacy-preserving learning

Federated learning addresses the central privacy obstacle to cloud detection. Rather than transmitting tenant data to a central server, each participant trains locally and shares only model updates, which a coordinator aggregates into a global model. Tenants thereby contribute to a shared detector without exposing their files, directly answering the reluctance that makes file-level network-function extraction impractical. Federated approaches must contend with non-identically distributed data across tenants, communication overhead, and poisoning by malicious participants, and they are frequently combined with differential privacy or secure aggregation to harden the update channel. As privacy regulation tightens and tenants grow more protective of their data, federated and otherwise privacy-preserving learning is positioned to become the default training regime for collaborative cloud detection.

5.5 Practical model selection in the cloud

Choosing among these model families in practice is governed less by peak benchmark accuracy than by deployment constraints. The availability and reliability of labels favor supervised classical or deep models where labels are plentiful and shift toward anomaly-based and self-supervised methods where they are not. The latency budget and the volume of workloads to be screened favor compact, fast models for first-pass triage and reserve heavier models for escalated cases. The privacy posture of the deployment favors representations and training regimes—reduced features, resource telemetry, and federated updates—that minimize what leaves the tenant boundary. The need for analyst trust and for justifiable remediation favors models, or accompanying explanation methods, that can articulate why a verdict was reached. Finally, the adversarial exposure of an internet-facing inference endpoint favors ensemble diversity and robust training over a single brittle high-accuracy model. The practical recommendation that emerges is to compose complementary models within a tiered pipeline rather than to seek a single best classifier, because the constraints that matter most in the cloud pull in directions that no individual model satisfies simultaneously.

VI. CLOUD-SPECIFIC DETECTION ARCHITECTURES

The same model can be deployed through markedly different architectures, and the architecture determines visibility, latency, privacy posture, and resilience. We organize cloud detection architectures into four classes.

6.1 Client-server cloud scanning

In the canonical cloud-scanning architecture, lightweight clients on endpoints—personal computers, mobile devices, IoT and CPS nodes—submit suspect files or extracted features to a cloud server that hosts the heavyweight detection model and a large threat database. The model analyzes each sample once and serves verdicts to many clients, amortizing analysis cost and improving coverage. This architecture offloads computation from constrained devices and is the natural home for centralized deep models. Its weaknesses are the privacy cost of transmitting files, the bandwidth and latency of submission, dependence on connectivity, and the absence of continuous on-device monitoring between submissions. Transmitting only reduced features rather than whole files, and caching verdicts at the edge, partially mitigate these costs.

6.2 Hypervisor and VMI-based detection

Provider-side architectures place detection at the hypervisor, using virtual machine introspection to observe guest memory and execution from outside the guest. This vantage point is robust against in-guest tampering and is therefore the architecture of choice against rootkits, kernel-level threats, and stealthy attacks that disable in-guest agents. Representative systems monitor processes through hypervisor-mediated views of virtual memory and classify monitored programs as benign or malicious using tree-based or other learners, often selecting features from system-call n-grams via metaheuristic optimization. The challenges are the semantic gap and the overhead of fine-grained monitoring, which must be balanced against the service-level objectives owed to tenants whose machines are being observed.

6.3 Container and orchestration monitoring

As workloads shift from virtual machines to containers orchestrated by platforms such as Kubernetes, detection follows. Container-aware architectures monitor the system calls a container issues to the host kernel, the resource it consumes, and the behavior of pods, in order to identify malicious activity such as cryptomining that hides within an otherwise legitimate deployment. Because automated remediation in orchestrated environments—removing or restarting a pod with a fresh image—is disruptive, this class of architecture places a premium on explainability: an operator must be given a justification for a classification before acting on it. Container detection must also cope with extreme ephemerality and density, monitoring many short-lived workloads with minimal per-container overhead.

6.4 Edge-cloud collaborative detection

For IoT and CPS deployments, a purely cloud-resident detector adds latency and creates a single point of failure, while a purely edge-resident detector is constrained by limited device resources. Collaborative edge-cloud architectures distribute detection across tiers: the cloud trains and maintains models and a global threat view, while edge nodes perform low-latency local inference, retrieving the appropriate model from the cloud and processing traffic near the device. Unsupervised detectors, such as autoencoders trained per device, guard against zero-day threats without labeled attack data, and multi-edge cooperation pools the capacity of several resource-limited nodes to handle bursts of traffic. This architecture aligns detection placement with the latency, privacy, and resource constraints of the underlying environment, and it is the most natural fit for securing the expanding population of connected devices.

Table 2. Comparison of cloud-specific detection architectures across key design dimensions.

Architecture	Visibility	Latency	Privacy posture	Best against
Client-server scanning	Submitted files / features	Network-bound	Weak (files leave tenant)	Known & unknown file malware
Hypervisor / VMI	Guest memory & execution (out-of-band)	Monitoring overhead	Provider-side; no file transfer	Rootkits, stealthy, kernel-level
Container / orchestration	Syscalls, pod & resource behavior	Low per-container	Provider-side	Cryptomining, pod compromise
Edge-cloud collaborative	Local traffic + global model	Low (edge inference)	Data stays near device	IoT/CPS, zero-day anomalies

VII. DATASETS AND EVALUATION METRICS

7.1 Datasets

Reproducible evaluation depends on shared datasets, and several have become reference points for the community. Network-intrusion datasets curated by academic cybersecurity institutes provide labeled flows spanning numerous attack types and tens of families across hundreds of features, supporting flow-based and behavioral detection. Large static-feature corpora of parsed executable metadata, released with explicit train-test splits, support reproducible static classification at scale. Family-classification benchmarks supply thousands of samples grouped into families and have driven progress in image-based and sequence-based methods. Android-focused collections support mobile and on-device research, and IoT-malware captures support edge-cloud studies. Three persistent dataset problems recur across this literature: class imbalance, because benign samples vastly outnumber malicious ones and malicious families are themselves unevenly represented; temporal drift, because a dataset frozen at one moment ceases to reflect the threats of the next; and label noise, because ground-truth labeling at scale is difficult and imperfect. A model that performs well on a static, balanced benchmark may degrade sharply in deployment, which is why evaluation protocol matters as much as raw scores.

7.2 Evaluation metrics

Accuracy alone is misleading under class imbalance, where a trivial classifier that labels everything benign can appear highly accurate while detecting nothing. Sound evaluation therefore reports precision, recall, and the F1 score, which balances the two; the false-positive rate, which is operationally critical because each false alarm consumes analyst attention and erodes trust; and the area under the receiver-operating-characteristic curve, which summarizes performance across decision thresholds. The confusion matrix, by separating false positives from false negatives, gives a fuller picture of where a detector errs. In the cloud, evaluation must extend beyond statistical quality to operational cost—*inference latency, memory footprint, data transferred between tenant and cloud, and the overhead imposed on monitored workloads*—because a detector that is accurate but too costly to run at scale fails in practice. Increasingly, robustness under adversarial perturbation and stability under concept drift are reported as first-class metrics rather than afterthoughts.

Table 3. Representative evaluation metrics and their interpretation for cloud malware detection.

Metric	Definition (informal)	Why it matters in the cloud
Precision	Share of flagged samples that are truly malicious	High precision limits wasted analyst effort and disruptive remediation
Recall (TPR)	Share of malicious samples that are caught	Missed malware in a multi-tenant host can spread laterally
F1 score	Harmonic mean of precision and recall	Single balanced figure under class imbalance
False-positive rate	Share of benign samples wrongly flagged	Directly governs operational cost and trust
AUC-ROC	Threshold-independent separability	Compares detectors without fixing an operating point
Latency / overhead	Time and resource cost of detection	Must respect tenant service-level objectives at scale

VIII. COMPARATIVE ANALYSIS OF REPRESENTATIVE APPROACHES

Table 4 positions representative directions in the literature against the taxonomy developed above, characterizing each by its analysis paradigm, feature representation, model family, deployment architecture, and the principal limitation that subsequent work must address. The table is organized by research direction rather than by individual paper, because the most useful comparison for a researcher entering the field is among design strategies, not among isolated results that were obtained on different datasets under different protocols and are therefore not directly comparable in absolute terms.

Table 4. Comparative analysis of representative malware-detection directions in cloud environments.

Direction	Analysis	Features	Model	Principal limitation
Signature / static scanning	Static	Bytes, headers, opcodes	Rule / classical	Defeated by packing, polymorphism, fileless execution
Behavioral sandboxing	Dynamic	API / syscall traces	Sequence / classical	Slow; sandbox-evasion and replay non-determinism
Image-based classification	Static	Binary-as-image	CNN	Sensitive to byte-layout obfuscation; family-level

Direction	Analysis	Features	Model	Principal limitation
VMI / introspection	Memory	Memory, syscall n-grams	Tree / classical	Semantic gap; monitoring overhead
Resource-anomaly detection	Dynamic	CPU/mem/net time series	RNN / autoencoder	Anomalies are not always malicious; tuning
Container / pod monitoring	Dynamic	Syscalls, pod telemetry	ML + explainability	Disruptive remediation needs justification
Edge-cloud collaborative	Hybrid	Local traffic + global model	Autoencoder / DL	Edge resource limits; model distribution
Federated detection	Hybrid	Local features / updates	Federated DL	Non-IID data, poisoning, communication cost

Two cross-cutting observations follow from this comparison. First, no single paradigm dominates: static methods are fast but fragile, dynamic methods are robust but costly, and introspection is powerful but constrained by the semantic gap—so the strongest systems combine paradigms rather than choosing among them. Second, the limitations that matter most in the cloud are systemic rather than model-specific. They concern privacy, scale, drift, adversarial exposure, and the overhead budget available for detection. Improving a classifier's headline accuracy on a benchmark addresses none of these directly, which is why this review treats the analysis-feature-model-architecture pipeline, rather than the classifier alone, as the proper unit of design.

IX. A PROPOSED CONCEPTUAL FRAMEWORK

Drawing the limitations together, we propose a layered conceptual framework for malware detection in cloud environments. Its design goals are to operate at both the virtual-machine and service levels, to preserve tenant privacy, to remain explainable, to control resource overhead through feature reduction and tiered analysis, and to adapt continually to drift. The framework comprises five layers.

Layer 1 — Multi-source acquisition. The framework collects observations from heterogeneous vantage points: in-guest agents where the tenant permits them, hypervisor-level introspection for out-of-band visibility, container-runtime syscall monitors, network-flow telemetry, and infrastructure resource counters. Acquiring signals at multiple layers ensures that a threat invisible at one layer—a rootkit to an in-guest agent, or a fileless payload to a disk scanner—remains visible at another.

Layer 2 — Privacy-preserving feature extraction. Rather than transmitting raw tenant data, the framework extracts compact feature representations close to the source and reduces their dimensionality before any data crosses a trust boundary. Feature selection and metaheuristic reduction shrink the representation to its discriminative core, lowering both transfer volume and the privacy exposure of what is shared, and resource-utilization features are favored where they suffice because they reveal behavior without inspecting tenant content.

Layer 3 — Tiered hybrid detection. A fast static and resource-anomaly tier triages every workload at low cost; only samples it cannot confidently resolve are escalated to expensive dynamic and introspection-based analysis. This tiering matches analysis cost to uncertainty and keeps the system tractable at cloud scale. A heterogeneous ensemble of complementary learners—classical models for engineered features, deep models for sequences and images, and autoencoders for anomaly detection—issues verdicts, with the ensemble pruned to a compact, high-performing subset to respect the overhead budget.

Layer 4 — Explainable decision and orchestration. Every verdict is accompanied by a human-readable justification—the features and behaviors that drove it—so that operators can trust and audit automated actions. This is essential where remediation, such as restarting a pod or quarantining a virtual machine, is disruptive. The orchestration

component maps verdicts to graduated responses, from alerting to isolation to self-healing redeployment from a known-good image.

Layer 5 — Continual learning and feedback. Confirmed verdicts and analyst corrections feed back into training through a federated, privacy-preserving regime, so that the global model improves without centralizing tenant data. Continual-learning mechanisms counter concept drift by updating the model as the threat landscape evolves, and adversarial-robustness measures—such as training against evasion and validating updates against poisoning—harden the pipeline against attacks aimed at the detector itself.

The framework is conceptual rather than prescriptive: its contribution is to make explicit that effective cloud detection requires coordinated choices across acquisition, privacy, analysis, explanation, and adaptation, and that optimizing any one layer in isolation—most commonly the classifier—leaves the dominant cloud-specific failure modes unaddressed. A concrete instantiation would select specific introspection tools, feature reducers, model architectures, explanation methods, and federated protocols according to the deployment's service model, tenant agreements, and overhead budget.

Figure 1 (described). Layered framework: acquisition feeds privacy-preserving feature extraction, which feeds a tiered hybrid ensemble; verdicts pass through an explainable orchestration layer to graduated response, and confirmed outcomes return through a federated continual-learning loop to all layers.

X. CHALLENGES AND OPEN RESEARCH ISSUES

Despite substantial progress, several challenges remain only partially solved and define the frontier of cloud malware detection.

Concept drift. Malware evolves continuously, so a model frozen at training time degrades as the threat landscape shifts. Detecting drift, deciding when to retrain, and updating without catastrophic forgetting remain open, and they are sharper in the cloud because the cost of retraining large models at frequent intervals is significant.

Class imbalance and label scarcity. Benign samples dominate, malicious families are unevenly represented, and high-quality labels are expensive. Imbalanced training biases models toward the majority class, while scarce labels limit supervised learning, motivating semi-supervised, self-supervised, and anomaly-based methods that learn primarily from benign behavior.

Adversarial robustness. Evasion, poisoning, model extraction, and membership inference all target learning-based detectors, and a cloud detector exposed to many tenants is especially queryable. Robust training, ensemble diversity, input validation, and update vetting help, but a general defense remains elusive, and robustness frequently trades against accuracy.

Non-deterministic replay and sandbox fidelity. Faithfully replaying concurrent, communicating programs is hard; serializing threads to achieve determinism can suppress the very exceptions that constitute an attack, and exhaustively replicating tenant environments is prohibitively expensive. Building convincing, efficient, evasion-resistant sandboxes at cloud scale is unresolved.

The semantic gap in introspection. VMI observes raw memory but must reconstruct high-level operating-system abstractions to interpret it, and this reconstruction is brittle across kernel versions and is itself a target for manipulation.

Resource overhead and service-level objectives. Detection competes with tenant workloads for resources. Fine-grained monitoring, large ensembles, and frequent dynamic analysis all impose overhead that must be balanced against the performance guarantees owed to tenants, making efficiency a first-class research goal rather than an afterthought.

Interoperability and portability. The absence of standardized interfaces across cloud providers means that detection components, and the antimalware engines that embed them, often must be rewritten when workloads migrate. Standardization and portable detection abstractions would reduce this friction substantially.

Explainability and trust. Deep models are accurate but opaque, yet disruptive automated remediation demands justification. Producing faithful, actionable explanations that operators can act on, without sacrificing accuracy, is an active and unsettled area.

XI. FUTURE DIRECTIONS

Several directions appear especially promising for advancing the field beyond its current limitations.

Continual and lifelong learning. Detection pipelines that learn incrementally, retaining prior knowledge while absorbing new threats, would directly confront concept drift, replacing periodic wholesale retraining with efficient, ongoing adaptation.

Self-healing detection. Coupling detection with automated, validated recovery—redeploying compromised workloads from known-good images and restoring isolation—would shorten the window between detection and remediation and reduce reliance on human intervention, an essential property at cloud scale.

Privacy-preserving collaboration. Maturing federated learning, secure aggregation, and differential privacy would let many tenants and providers build shared detectors without exposing data, dissolving the privacy obstacle that has long constrained collaborative cloud detection.

Graph and behavior-centric representations. Because malicious logic is more invariant than malicious appearance, deeper use of graph neural networks over call and control-flow structure, and of behavior-centric features generally, promises robustness against the obfuscation that defeats surface-level methods.

Language models for behavioral reasoning. Large language and code models offer a new avenue for interpreting decompiled code, summarizing behavioral traces, and generating human-readable explanations of verdicts, potentially narrowing the explainability gap—while their own susceptibility to manipulation must be studied carefully before operational reliance.

Standardized, drift-aware benchmarks. Community benchmarks that explicitly model temporal drift, class imbalance, adversarial perturbation, and operational cost would align reported results with deployment reality and accelerate cumulative progress.

XII. CONCLUSION

Cloud computing delivers scalability, elasticity, and cost efficiency that have made it indispensable to modern computing, yet the same properties concentrate value and exposure in ways that malware authors are well equipped to exploit. Signature-based detection, adequate for a single endpoint and a static threat, cannot withstand the volume, mutation, and stealth of malware in multi-tenant, virtualized, containerized, and edge-connected cloud environments. Machine learning and deep learning offer a path forward by generalizing from patterns rather than memorizing appearances, and the cloud's abundant computation and shared threat intelligence make it a natural home for such detectors—provided the privacy, scale, drift, adversarial, and overhead challenges peculiar to the cloud are confronted directly.

This paper has organized the field into a unified taxonomy linking analysis paradigms, feature representations, learning models, and deployment architectures; compared representative directions along consistent dimensions; consolidated the datasets and metrics on which the field depends; and proposed a layered conceptual framework that treats acquisition, privacy-preserving feature reduction, tiered hybrid detection, explainable orchestration, and continual federated learning as a coordinated whole. The recurring lesson is that detection quality in the cloud is determined less by the choice of any single classifier than by the coherence of the full pipeline and by how squarely it addresses the cloud's systemic failure modes. We hope this synthesis helps researchers locate new contributions precisely within the design space and informs the construction of scalable, privacy-preserving, explainable, and adaptive malware detection systems that can sustain high quality of service as both the cloud and its adversaries continue to evolve.

REFERENCES

- [1]. Garg, N. (2022). Fault-tolerant operation of medium voltage PMSM drive systems under open-switch faults. *Journal of Electrical Systems*, 18(4), 202–213. <https://doi.org/10.52783/jes.9358>

- [2]. Jaiswal, I. A. (2023). High-Performance AI-Augmented Content Management Systems for Distributed Clouds. *International Journal of Multidisciplinary Innovation and Research Methodology*, ISSN: 2960-2068, 2 (2), 90–97. Retrieved from <https://ijmirm.com/index.php/ijmirm/article/view/243>.
- [3]. S. Velamuri, M. V. V. Prasad Kantipudi, R. Sitharthan, D. Kanakadhurga, N. Prabakaran, and A. Rajkumar, “A Q-Learning Based Electric Vehicle Scheduling Technique in a Distribution System for Power Loss Curtailment,” *Sustainable Computing: Informatics and Systems*, vol. 36, Art. no. 100798, 2022
- [4]. Jaiswal, I. A. (2023). Intelligent Cybersecurity Framework for Large-Scale RESTful Service Architectures. *International Journal of Research Radicals in Multidisciplinary Fields*, ISSN: 2960-043X, 2 (1), 178–184. Retrieved from <https://www.researchradicals.com/index.php/rr/article/view/252>.
- [5]. Swathi, S. Kumar, S. Rani, A. Jain, and R. K. M. V. N. M., “Emotion classification using feature extraction of facial expression,” in *Proceedings of the 2022 2nd International Conference on Technological Advancements in Computational Sciences (ICTACS)*, pp. 283–288, 2022.
- [6]. Khan, S. (2022). The use of artificial intelligence in ESG (environmental, social, and governance) investing: A study on AI models for sustainability analytics in asset and wealth management. *International Journal of Multidisciplinary and Scientific Emerging Research*, 10(1), 431–439. <https://doi.org/10.15662/IJMSEH.2022.1001046>
- [7]. Jaiswal, I. A. (2023). Multilingual and culturally adaptive AI models for global education platforms. *International Journal for Research in Education (IJRE)*, 12(9), 17–27. <https://doi.org/10.63345/ijre.v12.i9.1>
- [8]. S. Rani, D. Ghai, and S. Kumar, “Reconstruction of wire frame model of complex images using syntactic pattern recognition,” in *IET Conference Proceedings CP791*, vol. 2021, no. 11, pp. 8–13, 2021.
- [9]. Khan, S. (2019). Sales force security compliance: An in-depth study of GDPR, HIPAA, and PCI-DSS enforcement in cloud-based CRM systems and their implications for global enterprises. *International Journal of Advanced Research in Electrical, Electronics and Instrumentation Engineering*, 8(9), 2211–2219. <https://doi.org/10.15662/IJAREEIE.2019.0809010>
- [10]. S. Kumar, E. G. Rajan, and S. Rani, “A study on vehicle detection through aerial images: Various challenges, issues and applications,” in *Proceedings of the 2021 International Conference on Computing, Communication, and Intelligent Systems (ICCCIS)*, pp. 504–509, 2021
- [11]. S. Rani, K. Lakhwani, and S. Kumar, “Syntactic approach to reconstruct simple and complex medical images,” *International Journal of Signal and Imaging Systems Engineering*, vol. 12, no. 4, pp. 127–136, 2023.
- [12]. M. Kumar, S. Kumar, M. Gulhane, R. K. Beniwal, and S. Choudhary, “Deep neural network-based fingerprint reformation for minimizing displacement,” in *Proceedings of the 2023 12th International Conference on System Modeling & Advancement in Research Trends (SMART)*, pp. 100–105, 2023.
- [13]. Khan, S. (2018). The role of zero-trust models in database security: Eliminating implicit trust and enforcing continuous verification in enterprise data access systems. *International Journal of Research in Electronics and Computer Engineering*, 6(1), 1622–1628.
- [14]. S. Kumar, R. Raja, M. R. Mahmood, and S. Choudhary, “A hybrid method for the removal of RVIN using self organizing migration with adaptive dual threshold median filter,” *Sensing and Imaging*, vol. 24, no. 1, pp. 9, 2023.
- [15]. Khan, S. (2017). The role of privacy-preserving techniques in database security: Secure multi-party computation, differential privacy, and homomorphic encryption. *The Research Journal*, 3(4), 29–35.
- [16]. Khan, S. (2016). A study of data provenance and integrity in database security: Ensuring authenticity, non-repudiation, and accountability in data lifecycle management. *International Journal of Research in Electronics and Computer Engineering*, 4(3), 180–186.

- [17]. Jaiswal, I. A. (2022). Natural language processing for security policy and log analysis. *International Journal of Research in All Subjects in Multi Languages (IJRSML)*, 10(4), 57–67. <https://doi.org/10.63345/ijrsml.v10.i4.1>
- [18]. Garg, N. (2021). Back-to-back testing of medium voltage (MV) drives. *TIJER – International Research Journal*, 8(12). <https://doi.org/10.56975/tijer.v8i12.159059>
- [19]. M. A. Jabbar, M. V. V. Prasad Kantipudi, S.-L. Peng, M. B. I. Reaz, and A. M. Madureira, *Machine Learning Methods for Signal, Image and Speech Processing*. River Publishers, 2022
- [20]. K. M. V. V. Prasad and H. N. Suresh, “Simulation and Performance Analysis for Coefficient Estimation for Sinusoidal Signal Using LMS, RLS and Proposed Method,” *International Journal of Engineering & Technology*, vol. 7, no. 1.2, Art. no. 1, 2017
- [21]. H. Pujara and K. M. Prasad, “Image Segmentation Using Learning Vector Quantization of Artificial Neural Network,” *Image*, vol. 2, no. 7, 2013.
- [22]. K. M. V. V. Prasad and H. N. Suresh, “An Efficient Adaptive Digital Predistortion Framework to Achieve Optimal Linearization of Power Amplifier,” in *2016 International Conference on Electrical, Electronics, and Optimization Techniques (ICEEOT)*, IEEE, 2016, pp. 2095–2101.
- [23]. Jaiswal, I. A. (2021). AI-orchestrated store deployment systems for global retail networks. *International Journal of Research in Modern Engineering & Emerging Technology*, 9(11), 42–53. <https://doi.org/10.63345/ijrmeet.org.v9.i11.1>
- [24]. Garg, N. (2023). Power factor & re-active power control by D-SVC system. *Fuel Cells Bulletin*, 2023(10), 226. <https://doi.org/10.5281/zenodo.17526729>
- [25]. G. K. Nanani and M. V. V. Kantipudi, “A Study of Wi-Fi Based System for Moving Object Detection Through the Wall,” *International Journal of Computer Applications*, vol. 79, no. 7, pp. 1–5, 2013.