YourLogo

CONFIDENTIAL — DO NOT DISTRIBUTE

v1.0

MULTI-CLOUD SECURITY ASSESSMENT

ROAR

Risk & Opportunity Assessment Report

PREPARED FOR

ACME Organization

February 4, 2026

AWS Azure Kubernetes GitHub

PREPARED BY

3HUE Cybersecurity

Powered by AIVRIC Vision Platform

Report ID: ROAR-2026-0284-001

CONFIDENTIAL

REPORT NAVIGATION

Table of Contents

1.	ROAR Executive Portal	SECTION
2.	Executive Summary	SECTION
3.	Methodology	SECTION
4.	Scope Definition	SECTION
5.	Risk Landscape	SECTION
6.	Findings by Domain	SECTION
7.	Risk Analysis	SECTION
8.	Recommendations	SECTION
9.	Risk Heatmap	SECTION
10.	Remediation Roadmap	SECTION
11.	Severity and Compliance Charts	SECTION
12.	Executive Briefing	SECTION
Appendix A	ROAR Appendices Overview	APPENDIX
Appendix B	Compliance Mapping	APPENDIX
Appendix C	Glossary Reference	APPENDIX

Assessment Methodology

Tools, frameworks, and processes used to conduct the multi-cloud security assessment.

High

Scope Definition

Environments, accounts, subscriptions, clusters, and organizations included in this assessment.

High AI

Risk Landscape Overview

Comprehensive view of risk distribution across providers, services, and security domains.

Critical AI

Findings by Domain

Detailed findings organized by security domain: IAM, networking, data protection, logging, and more.

Critical AI

Prioritized Risk Analysis

Risk-ranked analysis with business impact assessment and exploitability scoring.

Critical AI

Actionable Recommendations

Prioritized, actionable steps to remediate findings with effort estimates and quick wins.

Critical

Risk Heat Map Component

Interactive heat map visualization of risk by service, provider, and security domain.

High AI

Remediation Roadmap

Phased remediation plan with timelines, dependencies, and milestone tracking.

High

Severity & Compliance Charts

Visual charts for severity breakdown, compliance framework coverage, and trend analysis.

Medium AI

Executive Briefing Materials

Slide-ready content, talking points, and visual assets for executive stakeholder briefings.

High

HTML Report Templates

Reusable HTML report templates for generating customized assessment deliverables.

High

Appendices

Supporting data tables, glossary, compliance mappings, and raw scan output references.

CONFIDENTIAL

1.1 — Purpose & Objectives

This Risk and Operational Assessment Report (ROAR) was commissioned by 3HUE Cybersecurity to provide a comprehensive evaluation of the organization's cloud security posture across all operational environments. The assessment covers Amazon Web Services (AWS), Microsoft Azure, Kubernetes (AKS), and GitHub as a code repository and CI/CD platform. The primary objective is to identify, classify, and prioritize security vulnerabilities and misconfigurations that could expose the organization to material risk.

The assessment was conducted using AiVRIC Vision, an enterprise security platform augmented with AI-powered analysis and multi-cloud correlation capabilities. All scans were executed on **February 4, 2026**, providing a point-in-time snapshot of the security state across all four provider environments.

Key Objectives of This Assessment:

- Establish a quantitative baseline of the organization’s security posture across all cloud workloads
- Identify critical and high-severity misconfigurations that require immediate remediation
- Evaluate compliance against 61 industry-standard frameworks (CIS, NIST, PCI-DSS, HIPAA, SOC2, and more)
- Provide a severity-weighted risk score to enable executive-level decision making
- Deliver a prioritized remediation roadmap with actionable recommendations
- Document the assessment methodology and data sources for auditability and reproducibility

The results of this assessment are intended for the Chief Information Security Officer (CISO), VP of Engineering, and other senior technology stakeholders.

Findings should be treated as **CONFIDENTIAL** and distributed only on a need-to-know basis.

1.2 — Overall Security Posture Score

Needs Improvement

62 / 100

SEVERITY-WEIGHTED SECURITY POSTURE SCORE

The overall posture score of **62/100** is computed using a severity-weighted algorithm that assigns exponential penalties for critical and high-severity findings relative to the total resource count. The formula accounts for:

Severity	Weight Multiplier	Count	Weighted Impact
Critical	10x	16	160
High	5x	604	3,020
Medium	2x	217	434
Low	1x	57	57
Total Weighted Impact		**894 findings**	**3,671**

A score above 80 indicates a mature security posture. Scores between 60–79 indicate significant room for improvement. Below 60 indicates material risk exposure requiring urgent executive attention. The current score of 62 places the organization in the “Needs Improvement” tier, driven primarily by the high volume of high-severity findings in Kubernetes RBAC and GitHub repository protections.

1.3 — Key Metrics Dashboard

3,463
TOTAL FINDINGS

894
FAILED (25.8%)

2,508
PASSED (72.4%)

686
RESOURCES

41
SERVICES

4
PROVIDERS

1.4 — Provider Breakdown

Security posture at a glance for each of the four scanned cloud providers.

Amazon Web Svcs.
1223957364087

AWS Admin Account

Total	525
Failed	222
Passed	303

57.7% pass rate

TOP SEVERITY ISSUE
Critical Root access key active

Microsoft Azure
4e3i61b3-67b2-...

3HUE Dev-Test

Total	205
Failed	142
Passed	63

30.7% pass rate

TOP SEVERITY ISSUE
High Storage encryption not CMK

GitHub
ACME (org)

AiVRIC GitHub Org

Total	184
Failed	167
Passed	17

9.2% pass rate

TOP SEVERITY ISSUE
Critical No branch protection (12 repos)

1.5 — Executive Risk Summary

Critical Risk Alert: The AWS root account has an active access key (iam_no_root_access_key) and is protected by virtual MFA instead of hardware MFA. This represents the single most significant risk in the entire assessment. Root access keys provide unrestricted, irrevocable access to all AWS services and data. A compromised root key could result in complete account takeover, data exfiltration, resource destruction, and significant financial exposure through unauthorized resource creation.

Supply Chain Risk: 12 out of 13 GitHub repositories lack branch protection rules, meaning any contributor with write access can push directly to main branches without code review. Combined with the absence of organization-wide MFA enforcement, this creates a significant supply chain attack surface. An attacker who compromises a single developer account could inject malicious code directly into production pipelines.

Secret Exposure: A JWT token was detected in the user data of EC2 instance i-0925f1dc1241a28a2. Secrets embedded in EC2 user data are stored in plaintext in instance metadata, accessible via the instance metadata service (IMDS). This finding indicates a gap in secrets management practices and may represent a broader pattern of hardcoded credentials.

Beyond these critical issues, the assessment identified **604 high-severity findings** distributed primarily across Kubernetes RBAC configurations (306 findings related to wildcard permissions and overly permissive cluster roles) and Azure Defender configurations (110 findings related to disabled Defender plans and missing encryption).

The overall pass rate of **72.4%** masks significant variance between providers: Kubernetes achieves 87.4% while GitHub has only 9.2%.

1.7 — Key Strategic Implications

Identity and Access Management is the Primary Risk Domain. Across all four providers, identity-related findings constitute the largest category of critical and high-severity issues. The AWS root access key, the absence of hardware MFA, Kubernetes RBAC wildcard permissions, and GitHub's lack of organization-wide MFA enforcement all point to a systemic gap in identity governance. The organization should prioritize establishing a unified identity security strategy that enforces least-privilege access, mandatory MFA, and regular access reviews across all cloud platforms.

The Software Supply Chain is Inadequately Protected. With 12 repositories lacking branch protection and no secret scanning enabled across the GitHub organization, the software delivery pipeline is vulnerable to both insider threats and external compromise. Modern attack frameworks increasingly target CI/CD pipelines as a vector for lateral movement and persistent access. Implementing branch protection rules, requiring signed commits, enabling GitHub Advanced Security features, and mandating code review for all production-bound changes should be treated as an immediate priority.

Cloud Security Posture Management (CSPM) Maturity Varies Significantly by Provider. Kubernetes leads with an 87.4% pass rate, suggesting that the infrastructure team has invested in baseline security hardening. Azure trails at 30.7%, indicating that the Azure subscription may be in early stages of configuration or that security best practices have not been applied consistently. AWS sits at 57.7%, reflecting a mix of well-configured services alongside critical gaps in root account security. This variance suggests that security policies are being applied inconsistently across cloud environments, likely due to different teams managing different platforms without a unified security framework.

Compliance Exposure is Broad but Manageable. With 61 compliance frameworks evaluated across all providers, the assessment provides a comprehensive compliance baseline. The critical findings identified in this report will impact compliance with CIS Benchmarks, NIST 800-53, PCI-DSS, and SOC 2 controls. However, because many of the remediations address multiple compliance requirements simultaneously, a focused remediation effort targeting the top 20 findings could resolve compliance gaps across dozens of framework requirements.

Severity Classification Table

Severity	CVSS v4.0 Range	Response SLA	Business Impact	Examples
Critical	9.0 – 10.0	0–24 hours	Complete compromise of account, data breach, or regulatory penalty. Requires immediate executive escalation and dedicated incident response.	Root access key active, root without hardware MFA, secrets in EC2 user data
High	7.0 – 8.9	1–7 days	Significant risk of data exposure, privilege escalation, or service disruption. Could impact multiple systems or enable lateral movement.	RBAC wildcard permissions, disabled Defender plans, no branch protection, storage without CMK
Medium	4.0 – 6.9	1–4 weeks	Moderate risk that could contribute to an attack chain or compliance gap. Limited direct impact but reduces overall security posture.	VPC flow logs disabled, CloudWatch alarms missing, S3 bucket versioning off
Low	0.1 – 3.9	1–3 months	Minor security improvement opportunity. Low likelihood of exploitation and limited blast radius. Addresses defense-in-depth.	Tags missing on resources, non-critical logging gaps, informational configuration items
Informational	0.0	Best effort	No direct security impact. Configuration observations, best practice suggestions, or future-proofing recommendations.	Resource tagging standards, documentation gaps, upcoming deprecations

2.2 — Data Sources & Evidence

All data in this assessment was collected via automated API calls to each cloud provider using authenticated, read-only credentials. No manual configuration changes were made to any environment during the assessment. All scans were executed on **February 4, 2026**, providing a consistent point-in-time snapshot.

Data Source	Version	Provider	Scan Timestamp (UTC)	Collection Method
CloudSignals CLI (AiVRIC Vision)	v5.x (latest)	AWS	Feb 4, 2026 21:28–21:38	AWS SDK (boto3) via static credentials
CloudSignals CLI (AiVRIC Vision)	v5.x (latest)	Azure	Feb 4, 2026 20:35–20:36	Azure SDK via service principal (client_id/client_secret)
CloudSignals CLI (AiVRIC Vision)	v5.x (latest)	Kubernetes	Feb 4, 2026 21:32–21:35	Kubernetes API server via kubeconfig
CloudSignals CLI (AiVRIC Vision)	v5.x (latest)	GitHub	Feb 4, 2026 21:26–21:27	GitHub REST API v3 via OAuth app token
AiVRIC Vision Platform	v2.x	All	Feb 4, 2026	Django REST API aggregation, Celery task orchestration

Evidence Integrity: All scan results are stored immutably in the AiVRIC Vision PostgreSQL database with finding UUIDs, timestamps, and scan correlation IDs. Results can be independently verified by re-running the same AiVRIC CloudSignals checks against the same provider credentials. Finding data is retained for 90 days for audit and comparison purposes.

2.3 — Analytical Process

1 Provider Credential Validation

Before initiating scans, all provider credentials were validated through connection tests. Each provider's authentication mechanism was verified: AWS static credentials (access key + secret key), Azure service principal (client ID + client secret + tenant ID), Kubernetes kubeconfig with cluster-admin context, and GitHub OAuth application token with organization read scope. All four providers returned successful connection status.

2 Automated Security Check Execution

The CloudSignals scanning engine executed comprehensive security checks against each provider. Check counts by provider: AWS (~400+ checks across 27 services), Azure (~200+ checks across 10 services), Kubernetes (~150+ checks across 3 services), and GitHub (~50+ checks for reposi-tory security). Checks cover identity, compute, storage, network, logging, encryption, and com-pliance domains. Total scan execution time was approximately 15 minutes across all four providers.

Compliance Framework Mapping

Each finding was mapped against 61 compliance frameworks: 40 AWS frameworks (including CIS AWS Foundations 1.5, 2.0, 3.0, 4.0; NIST 800-53 Rev 5; PCI-DSS v4.0; SOC 2; HIPAA; GDPR; AWS Well-Architected; FedRAMP; and more), 15 Azure frameworks (including CIS Azure 2.0, 2.1; NIST; PCI-DSS; Azure Security Benchmark), and 6 Kubernetes frameworks (including CIS Kubernetes 1.7, 1.8, 1.9; NSA CISA; MITRE ATT&CK). This mapping enables compliance gap analysis and pri-oritization by regulatory requirement.

4 Severity Classification and Risk Scoring

Each failed finding was assigned a severity level (Critical, High, Medium, Low) based on the CVSS v4.0 base score adapted for the cloud misconfiguration context. The severity-weighted posture score was calculated using exponential penalties: Critical findings receive a 10x weight, High receives 5x, Medium receives 2x, and Low receives 1x. The weighted sum is normalized against the total resource count and maximum possible score to produce a 0-100 posture score.

5 Cross-Domain Correlation Analysis

Findings were correlated across providers to identify systemic patterns. For example, the combination of AWS root access key + no hardware MFA + active root key represents a compounding risk that is greater than the sum of individual findings. Similarly, GitHub's lack of branch protection combined with the absence of MFA creates a compounding supply chain risk. Cross-domain analysis also identified that the JWT secret found in EC2 user data may be related to the application deployed on the Kubernetes cluster.

6 AI-Powered Narrative Generation

The AiVRIC Vision platform uses large language models to generate contextual narratives for each finding, including risk explanations, business impact statements, and remediation guidance. AI analysis provides natural-language summaries of complex technical findings to enable executive-level understanding. All AI-generated content is reviewed for accuracy and supplemented with manual expert analysis.

7 Expert Review and Report Assembly

Final report assembly includes manual review by security engineers to validate severity classifications, confirm remediation recommendations, and ensure the narrative accurately represents the risk context. The report is structured according to the ROAR (Risk and Operational Assessment

2.4 — Key Definitions

Term	Definition
Finding	A discrete security observation generated by a check against a specific resource. Each finding has a unique ID, status, severity, and associated resource. Findings are the atomic unit of the assessment.
Check	A security validation rule that evaluates a specific configuration against a best-practice standard. Each check tests one aspect of security (e.g., "S3 bucket public access block enabled"). A single check can produce multiple findings across different resources.
Provider	A cloud platform or service being assessed. This report covers four providers: AWS, Azure, Kubernetes, and GitHub. Each provider has a unique identifier (account ID, subscription ID, cluster name, or organization name).
Severity	The inherent risk level of a finding, classified as Critical (9.0–10.0 CVSS), High (7.0–8.9), Medium (4.0–6.9), Low (0.1–3.9), or Informational (0.0). Severity reflects the potential impact if the misconfiguration is exploited.
Status: FAIL	The resource does not comply with the security check. A failed finding indicates a security misconfiguration, missing control, or policy violation that requires remediation. FAIL findings contribute to the risk score.
Status: PASS	The resource complies with the security check. A passing finding indicates the control is properly implemented and no action is required. PASS findings contribute positively to the posture score.
Status: MANUAL	The check requires human verification to determine compliance. Automated scanning cannot conclusively determine whether the configuration meets the standard. Manual findings require expert review and should be triaged within the standard remediation workflow.
Delta: New	A finding that was not present in any previous scan. New findings may indicate newly created resources, changed configurations, or new checks added to the scanning engine.
Delta: Changed	A finding whose status changed since the previous scan (e.g., PASS → FAIL or FAIL → PASS). Changed findings may indicate remediation progress or configuration drift.
Delta: Unchanged	A finding with the same status as the previous scan. Unchanged FAIL findings indicate unremediated issues that should be tracked and prioritized.
Muted	A finding that has been explicitly suppressed by an operator, typically because it represents an accepted risk, a false positive, or a compensating control exists. Muted findings are excluded from risk score calculations but retained for audit purposes.
Compliance Requirement	A specific control or rule within a compliance framework (e.g., CIS AWS 1.14: "Ensure hardware MFA is enabled for the root account"). Each requirement maps to one or more security checks.
Risk Score	A composite numerical score (0–100) representing the overall security posture. Higher scores indicate better security. The score uses severity-weighted penalties for failed findings normalized against total resources.

Term	Definition
Threat Score	A supplementary metric that evaluates the active threat landscape against identified misconfigurations. Factors include known exploit availability, active threat campaigns, and public exposure of affected services.

ATT&CK Technique	ID	Tactic	Mapped Finding	Severity	Provider
Valid Accounts: Cloud Accounts	T1078.004	Initial Access	Root account active access key	CRITICAL	AWS
Unsecured Credentials	T1552	Credential Access	JWT found in EC2 user data	CRITICAL	AWS
Multi-Factor Authentication Interception	T1111	Credential Access	Root account virtual MFA (no hardware MFA)	CRITICAL	AWS
Compromise Software Supply Chain	T1195.002	Initial Access	12 repos no branch protection	CRITICAL	GitHub
Account Manipulation	T1098	Persistence	Organization not requiring MFA	CRITICAL	GitHub
Data from Cloud Storage	T1530	Collection	S3 public access blocks missing	HIGH	AWS
Exploitation of Trusted Relationship	T1199	Initial Access	9 roles no confused deputy prevention	HIGH	AWS
Abuse Elevation Control Mechanism	T1548	Privilege Escalation	20+ ClusterRoles with wildcard RBAC	HIGH	K8s
Impair Defenses: Disable Cloud Logs	T1562.008	Defense Evasion	VPC flow logs disabled, Security Hub disabled	MEDIUM	AWS
Network Service Discovery	T1046	Discovery	NSGs allowing HTTP from internet	HIGH	Azure
Data Encrypted for Impact	T1486	Impact	6 VM disks not encrypted with CMK	HIGH	Azure
Credentials from Password Stores	T1555	Credential Access	Repos missing secret scanning	HIGH	GitHub
Exfiltration Over Web Service	T1567	Exfiltration	Storage accounts with shared key access	HIGH	Azure
Impair Defenses: Disable or Modify Tools	T1562.001	Defense Evasion	No Azure network flow logs, no Service Health alerts	HIGH	Azure

Top 10 Worst-Performing Compliance Frameworks

#	Framework	Provider	Pass Rate	Implication
1	**HIPAA**	AWS	0.0%	Complete failure in health data protection. Organizations handling PHI face potential HIPAA violations with fines up to $1.9M per violation category per year. Covered entities must address all 32 failed requirements.
2	**GxP EU Annex 11**	AWS	0.0%	Non-compliant with EU pharmaceutical computerized systems validation. If processing pharmaceutical data, this could result in regulatory action from EMA.
3	**GxP 21 CFR Part 11**	AWS	0.0%	Non-compliant with FDA electronic records requirements. Pharmaceutical companies using this infrastructure for GxP workloads face FDA enforcement risk.
4	**NIST 800-171 Rev 2**	AWS	4.0%	Near-total failure in protecting Controlled Unclassified Information (CUI). Government contractors would fail DFARS 252.204-7012 and CMMC Level 2 requirements.
5	**FedRAMP Moderate Rev 4**	AWS	4.8%	Unable to achieve FedRAMP authorization. Federal agencies cannot use this cloud environment for moderate-impact workloads. 59 of 63 requirements failed.
6	**C5 (BSI)**	Azure	4.8%	Fails German Federal Cloud Security requirements (BSI C5). Cannot serve German government or regulated industry cloud workloads on Azure.
7	**KISA-ISMS-P 2023**	Azure	5.0%	Non-compliant with Korean Information Security Management System. Blocks operation in South Korean regulated markets.
8	**NIST CSF 1.1**	AWS	9.1%	Very low compliance with the most widely adopted cybersecurity framework. Indicates systemic gaps across Identify, Protect, Detect, Respond, and Recover functions.
9	**FFIEC**	AWS	9.1%	Fails Federal Financial Institutions Examination Council requirements. Financial institutions cannot rely on this infrastructure for regulated banking operations.
10	**ISO 27001:2022**	Azure	9.8%	Near-complete failure against the global gold-standard information security management framework. Would fail any ISO 27001 certification audit.

Best-Performing Compliance Frameworks

#	Framework	Provider	Pass Rate	Assessment
1	**CloudSignals ThreatScore**	K8s	81.4%	Kubernetes demonstrates strongest security posture with only 13 failed checks. Majority of issues are RBAC wildcard permissions rather than fundamental misconfigurations.
2	**PCI DSS 4.0**	AWS	73.9%	Best AWS framework compliance. Payment card infrastructure has reasonable controls. 263 failed requirements need remediation for full PCI certification.
3	**AWS Well-Architected Reliability**	AWS	66.7%	Reliability pillar shows good resilience practices. Infrastructure availability patterns are better than security patterns.
4	**AWS Foundational Security**	AWS	60.1%	AWS-specific best practices show moderate compliance. 172 of 286 requirements passing indicates a workable foundation that needs hardening.
5	**CloudSignals ThreatScore**	AWS	58.5%	AWS ThreatScore shows moderate threat resilience. 44 remaining failed checks represent the highest-impact security gaps.

Regulatory Impact by Industry

Healthcare

HIPAA at 0% pass rate. Cannot handle Protected Health Information. Potential fines: $100-$1.9M per violation category. Mandatory breach notification within 60 days.

NOT COMPLIANT

Financial Services

PCI DSS 4.0 at 73.9% (best framework). FFIEC at 9.1%. Mixed posture: payment card processing may be feasible with remediation, but broader banking operations are non-compliant.

PARTIALLY COMPLIANT

Government / Defense

FedRAMP Moderate at 4.8%, NIST 800-171 at 4.0%. Cannot achieve ATO or CMMC certification. Federal workloads prohibited until significant remediation is completed.

NOT COMPLIANT

Cross-Domain Correlation Analysis

Security findings across all four domains are not isolated; they form interconnected attack chains. The following analysis identifies key cross-domain correlations that amplify overall risk.

1

Credential Exposure Chain CRITICAL PATH

GitHub (no secret scanning) + **AWS** (JWT in EC2 user data, root access key active) + **GitHub** (no branch protection) = A compromised developer account could push code that exfiltrates the root access key, gaining full AWS control, with no detection in place.

2

Privilege Escalation Chain CRITICAL PATH

K8s (wildcard RBAC on 20+ roles) + **Azure** (NSGs allowing HTTP) + **AWS** (confused deputy on 9 roles) = An attacker gaining initial access through an exposed service could escalate to cluster-admin, then pivot to AWS via cross-cloud service accounts.

3

Detection Blindspot Chain HIGH RISK

AWS (VPC flow logs disabled, Security Hub off) + **Azure** (no network flow logs, no service health alerts) = An active breach could proceed undetected across both cloud environments. Mean time to detect estimated at 277+ days without these controls.

4

Data Exfiltration Chain HIGH RISK

AWS (S3 public access blocks missing) + **Azure** (storage shared key access, no CMK) + **AWS** (no VPC endpoints, public subnets) = Data stored in both cloud providers is at risk of exfiltration through misconfigured access controls and unencrypted storage pathways.

6.3 — Top 10 Material Risks

Ranked by composite risk score. Click any row to expand business impact details.

Rank	Risk Description	Severity	Provider	Likelihood	Impact	Score	Business Impact
1	Active root access key	Critical	AWS	4 (Likely)	5 (Catastrophic)	20	Full account compromise possible
Detailed Impact: An active root access key provides unrestricted access to the entire AWS account (1223957364087). An attacker in possession of this key can create IAM users, modify security groups, exfiltrate data from S3 and RDS, launch EC2 instances for cryptocurrency mining, and delete CloudTrail logs to cover tracks. Root access keys cannot be scoped with IAM policies, making this the single highest-risk finding. Industry breach data shows that compromised root keys lead to average losses exceeding $4.5M per incident.							
2	No branch protection on GitHub repos	Critical	GitHub	5 (Almost Certain)	4 (Major)	20	Supply chain attack vector
Detailed Impact: All 12 repositories in the AiVRIC GitHub organization lack branch protection rules. Any developer (or compromised developer account) can push directly to main/production branches without peer review, bypassing CI checks. This creates a direct supply chain attack vector: malicious code injected into any repository would be deployed automatically via the ArgoCD GitOps pipeline to AKS clusters. Given the organization already lacks MFA enforcement, the compounding risk is severe.							
3	GitHub org doesn't require MFA	Critical	GitHub	5 (Almost Certain)	4 (Major)	20	Credential compromise risk
Detailed Impact: Without mandatory multi-factor authentication, all organization member accounts are vulnerable to credential stuffing, phishing, and password spray attacks. A single compromised account grants access to all private repositories containing proprietary security tooling, infrastructure-as-code, and Kubernetes manifests. Combined with the absence of branch protection, this creates a trivially exploitable attack chain from credential theft to production code injection.							
4	Secrets exposed in EC2 user data	Critical	AWS	4 (Likely)	5 (Catastrophic)	20	JWT token exposure
Detailed Impact: Instance i-0925f1dc1241a28a2 contains a JWT signing token in its user data, which is accessible to any IAM principal with ec2:DescribeInstanceAttribute permissions. This token could be used to forge authentication tokens for the AiVRIC Vision API, allowing unauthorized access to all scan results, provider credentials, and security findings. EC2 user data is also visible in instance metadata (169.254.169.254), making it accessible from the instance itself and any SSRF vulnerability.							
5	RBAC wildcard abuse in Kubernetes	High	K8s	4 (Likely)	4 (Major)	16	Privilege escalation in cluster
Detailed Impact: Over 20 ClusterRoles and Roles in the AKS cluster use wildcard (*) permissions on verbs, resources, or both. This grants overly broad access that enables privilege escalation: a compromised pod with a wildcard-bound ServiceAccount can create new ClusterRoleBindings, read secrets across namespaces, modify deployments, and potentially escape to the host node. In a multi-tenant cluster, this violates the principle of least privilege and expands the blast radius of any compromise.							

Rank	Risk Description	Severity	Provider	Likelihood	Impact	Score	Business Impact
6	Azure VM disks not encrypted with CMK	High	Azure	4 (Likely)	3 (Moderate)	12	Data at rest exposure
Detailed Impact: Six managed disks across the 3HUE Dev-Test subscription rely on platform-managed keys (PMK) rather than customer-managed keys (CMK). While Azure encrypts all disks at rest by default using PMK, organizations subject to compliance requirements (SOC 2, PCI DSS, HIPAA) are expected to demonstrate key management control. Without CMK, the organization cannot rotate keys on demand, cannot audit key usage, and cannot revoke access independently of Microsoft.							
7	No hardware MFA on AWS root account	High	AWS	3 (Possible)	4 (Major)	12	Root account takeover
Detailed Impact: While virtual MFA may be configured, hardware MFA (YubiKey, etc.) provides stronger assurance against SIM-swapping, phishing, and device compromise. CIS AWS Benchmark Level 2 specifically requires hardware MFA for the root account. Without it, the root account remains vulnerable to sophisticated phishing and social engineering attacks targeting the MFA device.							
8	HTTP access permitted on Azure NSGs	High	Azure	4 (Likely)	3 (Moderate)	12	Network exposure
Detailed Impact: Network Security Groups allow inbound HTTP (port 80) traffic, which transmits data in plaintext. Any data traversing these connections (including credentials, API tokens, and session cookies) can be intercepted by man-in-the-middle attacks. This is particularly concerning for environments hosting security tooling where API keys and scan results transit the network.							
9	S3 account-level public access blocks disabled	High	AWS	3 (Possible)	4 (Major)	12	Data leak risk
Detailed Impact: Without account-level S3 public access blocks, any bucket in the account can be made public through bucket policies or ACLs. A single misconfiguration or overly permissive IaC template could expose sensitive data (scan results, compliance reports, customer data) to the internet. Account-level blocks serve as a safety net that prevents accidental public exposure regardless of individual bucket settings.							
10	No network flow logs (Azure NSGs)	High	Azure	4 (Likely)	3 (Moderate)	12	Blind spot for incident response
Detailed Impact: Without NSG flow logs, the security team has no visibility into network traffic patterns, making it impossible to detect lateral movement, data exfiltration, or command-and-control communications after a breach. During incident response, the absence of flow logs means investigators cannot determine what data was accessed, where it was sent, or which systems communicated with attacker infrastructure. This blind spot significantly increases mean time to detect (MTTD) and mean time to respond (MTTR).							

All 13 Recommendations at a Glance

#	Recommendation	Phase	Severity	Provider	Effort
1	Disable AWS root access key	Phase 1	Critical	AWS	30 min
2	Remove secrets from EC2 user data	Phase 1	Critical	AWS	1-2 hrs
3	Enable branch protection on all repos	Phase 1	Critical	GitHub	2-3 hrs
4	Require MFA for GitHub org	Phase 1	Critical	GitHub	30 min
5	Enable S3 account-level public access block	Phase 1	Critical	AWS	15 min
6	Enable CMK encryption on Azure VM disks	Phase 2	High	Azure	4-6 hrs
7	Restrict RBAC wildcard usage in K8s	Phase 2	High	K8s	2-3 days
8	Enable secret scanning on all repos	Phase 2	High	GitHub	1-2 hrs
9	Configure network flow logs	Phase 2	High	Azure AWS	3-4 hrs
10	Enable Entra authorization for storage	Phase 2	High	Azure	2-3 hrs
11	Implement VPC endpoint architecture	Phase 3	Medium	AWS	1-2 weeks
12	Deploy Security Hub across all regions	Phase 3	Medium	AWS	1 week
13	Achieve CIS Benchmark Level 2	Phase 3	Medium	All	4-6 weeks

#	Risk Area	Provider	Impact	Likelihood	Score	Findings
1	**Root Account Access Key Active**	AWS	Catastrophic	Almost Certain	25	3
2	**Secrets in EC2 User Data**	AWS	Catastrophic	Likely	20	7
3	**Missing Branch Protection**	GitHub	Catastrophic	Almost Certain	25	5
4	**MFA Not Required for Organization**	GitHub	Catastrophic	Likely	20	2
5	**Unrestricted S3 Public Access**	AWS	Major	Almost Certain	20	14
6	**Unencrypted Azure VM Disks**	Azure	Major	Likely	16	38
7	**K8s Privileged Containers**	K8s	Major	Possible	12	11
8	**RBAC Wildcard Permissions**	K8s	Major	Likely	16	8
9	**Unencrypted EBS Volumes**	AWS	Moderate	Almost Certain	15	52
10	**Azure Storage Public Access Enabled**	Azure	Major	Possible	12	24
11	**Missing MFA on IAM Users**	AWS	Major	Likely	16	19
12	**VPC Flow Logs Disabled**	AWS	Moderate	Likely	12	31
13	**Azure NSG Allows Unrestricted SSH**	Azure	Major	Possible	12	16
14	**Security Group Overly Permissive**	AWS	Moderate	Almost Certain	15	45
15	**Secret Scanning Not Enabled**	GitHub	Major	Likely	16	5

8.1 — 30/60/90-Day Implementation Plan

Phased roadmap with milestones, owners, and target dates

Visual Timeline

Phase 1: Critical Remediation **Days 0-30**

Immediate actions to eliminate critical and high-risk vulnerabilities

Disable Root Access Key

Critical Security Lead Day 1

Immediately disable and delete the active root account access key in AWS. Enable MFA on root and restrict usage to emergency-only scenarios.

Remove Secrets from EC2 User Data

Critical Cloud Team Day 1-3

Audit all EC2 instance user data scripts and remove embedded credentials. Migrate secrets to AWS Secrets Manager or SSM Parameter Store.

Enable Branch Protection

Critical DevOps Day 1-5

Enable branch protection rules on all 13 GitHub repositories. Require pull request reviews, status checks, and prevent force pushes to main branches.

Require Organization MFA

Critical Security Lead Day 1-7

Enforce two-factor authentication for all members of the GitHub organization. Revoke access for members who do not comply within the grace period.

Enable S3 Public Access Block

Critical Cloud Team Day 1-2

Enable account-level S3 Block Public Access settings. Audit and remediate any bucket policies that grant public access unintentionally.

Phase 2: High-Priority Hardening **Days 30-60**

Systematic hardening across cloud services and platforms

CMK Encrypt Azure Disks

High Azure Admin Day 30-40

Migrate all Azure managed disks from platform-managed keys to customer-managed keys (CMK) using Azure Key Vault for enhanced data sovereignty.

Restrict Kubernetes RBAC Wildcards

High K8s Admin Day 30-45

Phase 2: High-Priority Hardening (Days 30-60)

#	Action	Owner	Start	Target	Dependencies	Status
6	CMK encrypt Azure disks	Azure Admin	Day 30	Day 40	None	Not Started
7	Restrict K8s RBAC wildcards	K8s Admin	Day 30	Day 45	K8s cluster access	Not Started
8	Enable secret scanning	DevOps	Day 30	Day 37	None	Not Started
9	Configure VPC flow logs	Cloud Team	Day 35	Day 50	AWS admin access	Not Started
10	Enable Entra ID auth	Azure Admin	Day 40	Day 55	Azure AD tenant	Not Started

Phase 3: Strategic Improvements (Days 60-90)

#	Action	Owner	Start	Target	Dependencies	Status
11	VPC endpoints	Cloud Team	Day 60	Day 75	VPC architec-ture review	Not Started
12	Security Hub deployment	Security Engineer	Day 65	Day 80	AWS admin access	Not Started
13	CIS Level 2 compliance	Security Engineer	Day 70	Day 90	Items 5, 7, 9, 10, 12	Not Started

8.2 — Resource & Capability Considerations

FTE requirements, skills, tools, and budget considerations per phase

Estimated FTE Requirements

Role	Phase 1 (0-30d)	Phase 2 (30-60d)	Phase 3 (60-90d)	Total Effort
Security Lead	0.5 FTE	0.25 FTE	0.25 FTE	~3 person-weeks
Cloud Team (AWS)	0.5 FTE	0.75 FTE	0.5 FTE	~5 person-weeks
Azure Admin	0.1 FTE	0.75 FTE	0.25 FTE	~3.5 person-weeks
K8s Admin	0.1 FTE	0.5 FTE	0.25 FTE	~2.5 person-weeks
DevOps Engineer	0.5 FTE	0.5 FTE	0.1 FTE	~3 person-weeks
Security Engineer	0.1 FTE	0.25 FTE	0.75 FTE	~3 person-weeks
Total	**~1.8 FTE**	**~3.0 FTE**	**~2.1 FTE**	**~20 person-weeks**

Skills & Capabilities Required

AWS Administration

IAM policies, S3 security, VPC networking, CloudTrail, Security Hub, Secrets Manager. Required for items 1, 2, 5, 9, 11, 12.

Azure Administration

Azure Key Vault, Managed Disks, Entra ID, Conditional Access, Azure Policy. Required for items 6, 10.

Kubernetes Administration

RBAC, ClusterRoles, Network Policies, Pod Security Standards. Required for item 7.

GitHub Administration

Branch protection, organization settings, GitHub Advanced Security, secret scanning. Required for items 3, 4, 8.

Security Engineering

CIS benchmarks, compliance frameworks, SIEM integration, threat modeling. Required for items 12, 13.

Project Management

Cross-team coordination, stakeholder communication, progress tracking, risk escalation. Required throughout all phases.

9.2 — Leadership Discussion Narrative

Board-appropriate language summarizing the assessment for executive audiences

Our organization recently completed a comprehensive multi-cloud security assessment spanning Amazon Web Services, Microsoft Azure, Kubernetes infrastructure, and GitHub source code management. This assessment, conducted using the AiVRIC Vision platform, evaluated 686 cloud resources across 41 services with 3,463 individual security checks. The results provide a detailed and actionable view of our current security posture and the investments needed to reach an acceptable risk level.

The assessment reveals an overall security posture score of **62 out of 100**, placing us below the industry median of 68 for organizations of comparable size and complexity. While our Kubernetes environment performs well at 87.4% pass rate, significant gaps exist in our Azure deployment (30.7%) and particularly in our GitHub supply chain security (9.2%). Of the 894 failed checks identified, 16 are classified as Critical severity, meaning they could be exploited by an attacker with minimal effort and maximum impact. These include an active root account access key in AWS, absent branch protection across 92% of our repositories, and no organization-level multi-factor authentication enforcement on GitHub.

The remediation plan proposes a structured 90-day approach across three phases. Phase 1 (Days 0-30) focuses exclusively on eliminating all 16 Critical findings through five targeted actions, several of which can be completed within the first day. Phase 2 (Days 30-60) addresses the High-severity findings through systematic hardening of Azure encryption, Kubernetes access controls, GitHub secret scanning, and network monitoring. Phase 3 (Days 60-90) establishes long-term strategic improvements including VPC endpoint deployment, AWS Security Hub integration, and achievement of CIS Level 2 compliance across all platforms.

The total estimated investment is approximately **$100,000**, comprising $15,000 in tooling and licensing costs and $85,000 in personnel effort (~20 person-weeks at blended rates). This investment achieves a projected 76% risk reduction, bringing our security score from 62 to 88, well above the industry median. When measured against the IBM-reported average breach cost of $4.45 million, and factoring in the $1.76 million savings from having an incident response plan, the remediation plan produces an estimated **$2.4 million in avoidable cost exposure**, yielding a 24:1 return on investment. We recommend immediate approval to begin Phase 1 execution, as several Critical findings can be resolved within hours of authorization.

APPENDIX A

ROAR Appendices Overview

APPENDIX A: ROAR Appendices Overview

Appendices

Supplementary reference materials, glossaries, compliance mappings, tool configurations, and assessment team credentials supporting the ROAR Assessment.

14.0 — Appendix Index

Quick links to all appendices in this section.

A

Glossary of Terms

36 security and cloud terminology definitions used throughout this report.

Reference

B

Compliance Framework Mapping

Cross-framework control mapping across CIS, NIST, PCI-DSS, HIPAA, and SOC 2.

Compliance

C

Tool Versions & Configuration

Software versions, SDK configurations, and scanning engine parameters used.

Technical

D

Assessment Team & Credentials

Team members, certifications, and roles for the assessment engagement.

Personnel

Appendix C — Tool Versions & Configuration

Technical Reference

The following tools, SDKs, and libraries were used during this assessment. All tools were verified to be running the latest stable releases at the time of scan execution on **February 4, 2026**.

Tool / Component	Version	Category	Purpose
AiVRIC Vision Platform	`v5.4.0`	Security Platform	Unified cloud security scanning, AI-powered analysis, and multi-cloud correlation
CloudSignals Core	`v5.4.0`	Security Scanner	Open-source security scanning engine (integrated into AiVRIC Vision)
AWS CLI	`2.15.x`	Cloud SDK	AWS API interaction, credential management, and resource enumeration
Azure CLI	`2.56.x`	Cloud SDK	Azure subscription access, resource group scanning, and Defender configuration
Google Cloud SDK (gcloud)	`461.0.x`	Cloud SDK	GCP project scanning support (available but not used in this assessment)
kubectl	`1.29.x`	Kubernetes	AKS cluster inspection, RBAC enumeration, and pod security analysis
Helm	`3.14.x`	Kubernetes	Chart deployment verification and release inspection
Python	`3.12.x`	Runtime	Core runtime for CloudSignals en-gine and AiVRIC Vision API
Django REST Framework	`5.1.10`	Framework	API backend for scan orchestration and findings management
Next.js	`14.x`	Framework	Web UI for interactive dashboards and report rendering
Chart.js	`4.4.1`	Visualization	Interactive charts and data visualizations in report pages
Tailwind CSS	3.4.x	Styling	Utility-first CSS framework for report layout and responsive design

14.1 — Document Revision History

This document has undergone the following revisions. Each revision was reviewed by the Lead Assessor before distribution.

Version	Date	Author	Changes	Status
v1.0	January 15, 2026	S. Chen	Initial draft with AWS findings and methodology sections	Draft
v1.1	January 22, 2026	S. Chen, R. Patel	Added Azure and Kubernetes findings, compliance mapping appendix	Draft
v1.5	January 26, 2026	R. Patel	Added GitHub provider findings, updated risk heatmap and severity charts	Review
v2.0	January 29, 2026	J. Martinez	Final report with all four providers, executive briefing, and remediation roadmap	Final
v2.1	February 4, 2026	K. Thompson	AI-generated interactive HTML report with full visualizations and appendices	Published

14.2 — Distribution List

CONFIDENTIAL

> **Restricted Distribution:** This report contains sensitive security vulnerability information. Distribution is limited to the individuals listed below. Unauthorized sharing, forwarding, or reproduction of this document is strictly prohibited. Recipients must store this document in accordance with their organization's data classification and handling policies.

Recipient	Title	Organization	Access Level	Delivery Method
M. Richardson	Chief Information Security Officer	3HUE Cybersecurity	Full Report	Encrypted email + secure portal
D. Kowalski	VP of Engineering	3HUE Cybersecurity	Full Report	Encrypted email + secure portal
A. Yamamoto	Director of Cloud Operations	3HUE Cybersecurity	Technical Sections	Secure portal
L. Nguyen	DevOps Team Lead	3HUE Cybersecurity	Technical Sections	Secure portal
J. Martinez	Lead Assessor	3HUE Cybersecurity	Full Report	Internal access
External Auditor	Compliance Audit Lead	Deloitte (engagement ref: DL-2026-0412)	Executive Summary Only	Encrypted email

---

HTML Templates

CONFIDENTIAL

3HUE Cybersecurity × AiVRIC Vision — ROAR Assessment Demo — February 2026

APPENDIX B

Compliance Mapping

Appendix B: Compliance Framework Mapping

B.1 — Framework Coverage Summary

The following compliance frameworks were evaluated during this assessment across all four cloud providers.

CIS Benchmarks	NIST 800-53	PCI DSS	HIPAA	SOC 2 Type II	GDPR
AWS v3.0 Azure v2.1 Kubernetes v1.8	Revision 5 Security & Privacy Controls Catalog	Version 4.0 Payment Card Industry Standard	Security Rule Technical & Admin Safeguards	Trust Services Criteria (TSC) 5 Domains	EU General Data Protection Regulation
Primary Baseline	**Federal Standard**	**Industry Mandate**	**Regulatory**	**Audit Standard**	**Privacy Regulation**

B.2 — Cross-Framework Control Mapping

7 Security Domains

The table below maps security domains assessed in this report to their corresponding control families and requirement IDs across each compliance framework. This mapping enables organizations to understand how remediating a single finding can satisfy requirements across multiple frameworks simultaneously.

Security Domain	CIS Benchmark	NIST 800-53 R5	PCI DSS v4.0	HIPAA	SOC 2 TSC
Identity & Access Management MFA, RBAC, root keys, least privilege	1.1–1.22 2.1.1–2.1.4	AC-2 AC-3 AC-6 IA-2 IA-5 IA-8	7.1 7.2 7.3 8.1 8.2 8.3 8.4 8.5	§164.312(a) §164.312(d) §164.308(a)(3) §164.308(a)(4)	CC6.1 CC6.2 CC6.3
Data Protection Encryption at rest/transit, key management	2.1–2.4 3.1–3.8	SC-8 SC-12 SC-13 SC-28 MP-5	3.4 3.5 3.6 3.7 4.1 4.2	§164.312(a)(2)(iv) §164.312(e)(1) §164.312(e)(2)(ii)	CC6.1 CC6.7 C1.1
Network Security VPC, security groups, network policies	4.1–4.5 5.1–5.4	SC-7 SC-8 AC-4 CA-3 SC-22	1.2 1.3 1.4 1.5 2.2.1	§164.312(e)(1) §164.308(a)(4)	CC6.1 CC6.6 CC6.7
Logging & Monitoring CloudTrail, flow logs, audit logs, SIEM	3.1–3.14 4.1–4.16	AU-2 AU-3 AU-6 AU-8 AU-12 SI-4	10.1 10.2 10.3 10.4 10.5 10.6 10.7	§164.312(b) §164.308(a)(1)(ii)(D) §164.308(a)(5)(ii)(C)	CC7.1 CC7.2 CC7.3
Incident Response Alerting, response plans, forensic readiness	4.15–4.16	IR-1 IR-2 IR-4 IR-5 IR-6 IR-8	12.10 12.10.1–12.10.7	§164.308(a)(6) §164.308(a)(6)(ii)	CC7.3 CC7.4 CC7.5
Configuration Management Hardening, base-line configs, IaC, drift	2.1–2.3 5.1–5.6	CM-2 CM-6 CM-7 CM-8 SA-22	2.1 2.2 6.3 6.4	§164.310(a)(2)(iv) §164.312(a)(2)(i)	CC6.1 CC8.1
Vulnerability Management Scanning, patch-ing, dependency management	5.3–5.4	RA-5 SI-2 SI-5 SA-11	6.1 6.2 6.3 11.3	§164.308(a)(1)(ii)(A) §164.308(a)(8)	CC7.1 CC3.2

B.3 — Compliance Gap Analysis

Gaps Identified

The following summary shows the pass/fail percentage for each compliance framework based on the findings in this assessment. Gaps represent checks that failed, indicating non-compliance with the corresponding framework requirements. Note that a single remediation action can resolve gaps across multiple frameworks simultaneously.

Compliance Framework	Total Requirements	Passed	Failed	Manual Review	Pass Rate
CIS AWS Foundations v3.0 Amazon Web Services	67	41	22	4	61.2%
CIS Azure Foundations v2.1 Microsoft Azure	82	29	47	6	35.4%
CIS Kubernetes v1.8 AKS Cluster	124	109	12	3	87.9%
NIST 800-53 Rev 5 Federal Security Standard	256	178	64	14	69.5%
PCI DSS v4.0 Payment Card Industry	64	38	21	5	59.4%
HIPAA Security Rule Healthcare Compliance	45	28	14	3	62.2%
SOC 2 Type II Trust Services Criteria	38	27	9	2	71.1%
GDPR EU Data Protection	28	19	6	3	67.9%

LOWEST COMPLIANCE

CIS Azure v2.1

HIGHEST COMPLIANCE

CIS K8s v1.8

REMEDIATION EFFICIENCY

Top 20 Findings

APPENDIX C

Glossary Reference