RINGSAFE

Cloud Security Audit Template

Seven control families, 80+ checks. The working template we use for AWS, Azure and GCP posture reviews — adapted from CIS Benchmarks, NIST 800-53, and DPDP security safeguards.

Tenant / Account: _____

Audit date: _____ Auditor: _____

RingSafe Cybersecurity Consulting

ringsafe.in · hello@ringsafe.in · Manish Garg — Associate CISSP, CEH, CCNP Enterprise

© 2026 RingSafe. Distribute freely internally. Not for resale.

HOW TO USE THIS TEMPLATE

A working audit, not a checklist

This is the template we fill in during a real engagement. Each control has three status options: **Pass**, **Fail**, or **N/A**. Record evidence (screenshot name, CLI command output file, Terraform module reference) in the evidence column — that is what turns an audit into a defensible artifact.

Scope this across every cloud account / subscription / project that holds production data. A single-account audit is rarely sufficient; attackers pivot across accounts through assumed roles and trust relationships.

> **Severity classes** — CRIT issues are "fix this week" exposures (public buckets, root MFA off). HIGH issues warrant a 30-day remediation SLA. MED issues go in the backlog with a named owner.

1. Identity & Access Management

#	CONTROL	SEVERITY	STATUS	EVIDENCE / NOTES
1.1	Root / global administrator account has hardware MFA enabled, no access keys, sealed in credential vault.	CRIT		
1.2	No IAM user has console access without MFA. Break-glass accounts documented separately.	CRIT		
1.3	IAM users do not have long-lived access keys. Human access is federated (SSO / IAM Identity Center / Entra ID).	HIGH		
1.4	Service access uses IAM roles / workload identity / managed identities — never embedded access keys.	HIGH		
1.5	Administrative privileges (AdministratorAccess, Owner) limited to named individuals, access reviewed quarterly.	HIGH		
1.6	Password policy enforces length ≥14, complexity, reuse history, max-age ≤90 days for any remaining passwords.	MED		
1.7	Cross-account / cross-tenant trust relationships are documented, audited, and use external IDs (AWS) or conditional access.	HIGH		
1.8	Unused IAM users, roles, and access keys (no activity 90 days) are disabled and scheduled for deletion.	MED		

1.9	Service Control Policies (AWS) / Management Group policies (Azure) prevent privilege escalation and disable risky services at org level.	HIGH		
1.10	Privileged sessions use Just-in-Time elevation (PIM / IAM Identity Center temporary access) with approval workflow.	MED		

2. Data protection

#	CONTROL	SEVERITY	STATUS	EVIDENCE / NOTES
2.1	No storage bucket / blob container is publicly accessible. Public access is explicitly blocked at account / subscription level.	CRIT		
2.2	All storage uses server-side encryption. Customer-managed keys (CMK) used for sensitive data stores.	HIGH		
2.3	All databases (RDS, Cloud SQL, Azure SQL, Cosmos, DynamoDB) have encryption at rest enabled.	HIGH		
2.4	TLS 1.2+ enforced on every public-facing service. Deprecated ciphers disabled.	HIGH		
2.5	KMS / Key Vault keys have rotation enabled (annual minimum) and usage is logged.	MED		
2.6	Backup and snapshot buckets are in a separate account / subscription, immutable, with delete-protection.	CRIT		
2.7	Sensitive data tags / labels applied consistently; DLP scan run at least quarterly (Macie, Purview, DLP API).	MED		
2.8	Personal data storage locations align with DPDP / data-residency requirements (India regions for applicable workloads).	HIGH		
2.9	Database snapshots are not shared with external accounts unless explicitly required and documented.	HIGH		
2.10	Secrets (API keys, DB credentials) stored in Secrets Manager / Key Vault / Secret Manager —	HIGH		

	never in code, env files, or config stores.			

3. Network security

#	CONTROL	SEVERITY	STATUS	EVIDENCE / NOTES
3.1	No security group / NSG allows 0.0.0.0/0 inbound on administrative ports (22, 3389, 1433, 3306, 5432, 27017, 6379).	CRIT		
3.2	VPC flow logs / NSG flow logs / VPC Flow Logs enabled and retained ≥ 90 days.	HIGH		
3.3	Network segmentation: production, staging, and management planes are in separate VPCs / VNets with controlled peering.	HIGH		
3.4	Private connectivity (PrivateLink, Private Endpoints) used for database and storage access from compute.	MED		
3.5	WAF in front of all public web applications with baseline rule set + custom rules for known attack patterns.	HIGH		
3.6	DDoS protection enabled (Shield Advanced / DDoS Protection Standard or Premium).	MED		
3.7	DNS zones have DNSSEC enabled where supported; no dangling DNS records point to decommissioned cloud resources.	HIGH		
3.8	Egress filtering in place for sensitive workloads; outbound traffic to unknown destinations is logged and alerted.	MED		

4. Compute & container security

#	CONTROL	SEVERITY	STATUS	EVIDENCE / NOTES
4.1	Instance metadata service hardened (IMDSv2 required on AWS; Azure Instance Metadata header validated).	HIGH		
4.2	VM images / AMIs patched regularly; no OS reached EOL in active use.	HIGH		

4.3	Kubernetes: RBAC strictly scoped, no cluster-admin bindings for service accounts, network policies enforced.	HIGH		
4.4	Container images scanned for CVEs at build and pre-deploy (Trivy, ECR scan, Defender for Containers).	HIGH		
4.5	No container runs as root unless explicitly required; pod security standards enforced (restricted / baseline).	MED		
4.6	Serverless function execution roles scoped to least privilege; no function has AdministratorAccess.	HIGH		
4.7	EDR / cloud workload protection deployed to every production VM (Defender, CrowdStrike, SentinelOne).	HIGH		
4.8	SSH key-based auth only; no password SSH. Bastion / session manager used for administrative access.	HIGH		

5. Logging, monitoring, incident response

#	CONTROL	SEVERITY	STATUS	EVIDENCE / NOTES
5.1	CloudTrail / Azure Activity Logs / Cloud Audit Logs enabled in every region, sent to a log-archive account, immutable.	CRIT		
5.2	GuardDuty / Defender for Cloud / Security Command Center enabled across all accounts / subscriptions / projects.	HIGH		
5.3	Log retention ≥ 1 year (regulatory) / ≥ 90 days hot searchable.	HIGH		
5.4	Critical alerts (root login, IAM policy change, public bucket, key deletion) route to a 24x7 on-call channel.	HIGH		
5.5	Incident response runbooks exist for: credential compromise, public data exposure, ransomware, compromised workload.	HIGH		
5.6	Tabletop exercise conducted in last 12 months. Findings closed.	MED		

5.7	Breach notification workflow tested against DPDP 72-hour window / CERT-In 6-hour window.	HIGH		

6. Change management & IaC

#	CONTROL	SEVERITY	STATUS	EVIDENCE / NOTES
6.1	Production changes applied via Infrastructure-as-Code (Terraform, CloudFormation, Bicep). Console changes are exceptions.	HIGH		
6.2	IaC pipelines scan for policy violations (Checkov, tfsec, Trivy IaC, Defender for DevOps) before apply.	HIGH		
6.3	CI/CD pipelines use short-lived OIDC federation tokens, not long-lived cloud credentials.	HIGH		
6.4	Drift detection runs weekly; out-of-band changes are investigated within 24 hours.	MED		
6.5	All production changes require peer review (PR approval) and produce an audit trail.	MED		

7. Compliance & governance

#	CONTROL	SEVERITY	STATUS	EVIDENCE / NOTES
7.1	Shared-responsibility model documented and understood by engineering and ops teams.	MED		
7.2	Data classification scheme maps to cloud resources (tags, labels) and to DPDP / sector regulators.	HIGH		
7.3	Third-party / SaaS integrations with the cloud account are inventoried and periodically reviewed.	MED		
7.4	Named DPO / security owner for the cloud environment; escalation path published.	HIGH		
7.5	Compliance framework mapping current (CIS, NIST 800-53, ISO 27001, SOC 2 CC, DPDP safeguards).	MED		

7.6	Vendor risk: sub-processor list maintained; contractual DPA / SCC in place for data crossing borders.	HIGH		

Summary scorecard

SECTION	PASS	FAIL	N/A	SCORE
1. IAM				
2. Data protection				
3. Network				
4. Compute / containers				
5. Logging / IR				
6. Change mgmt / IaC				
7. Compliance				
Overall posture				

Red lines. If any CRIT control is a fail, treat it as a same-week remediation. A public bucket with real customer data, root account without MFA, or CloudTrail/Activity Logs disabled are not findings — they are incidents waiting to be reported.

Next step. Send us the completed scorecard at hello@ringsafe.in for a free 45-minute review. We will tell you which 3 fixes will reduce the most risk this quarter — even if you choose to implement them yourself. Book: ringsafe.in/contact.

© 2026 RingSafe Cybersecurity Consulting · ringsafe.in · Prepared by Manish Garg (Associate CISSP, CEH, CCNP Enterprise)

Mapped to: CIS Benchmarks (AWS, Azure, GCP), NIST SP 800-53 rev 5, ISO/IEC 27017, DPDP Act 2023 reasonable security safeguards. Not legal or regulatory advice. Version 1.0 — April 2026.