



Cloud Readiness Assessment & Migration Roadmap

Determine If Your Organization Is Ready for Cloud Migration
and Build a Phased Roadmap for the Transition

Includes:

Workload Classification Matrices • Cost Modeling Frameworks
Risk Assessment Tools • 90-Day Quick-Start Plan

Technology Guide

ComDirect USA

comdirectusa.com

This guide is intended for informational purposes only. Cloud migration requirements vary by organization. Cost estimates and timelines are illustrative and should be validated against your specific environment.

Table of Contents

- 1. Is Cloud Right for Your Organization?**
 - 2. Cloud Readiness Self-Assessment**
 - 3. Workload Classification and Prioritization**
 - 4. The 6 R's of Cloud Migration**
 - 5. Cloud Provider Selection**
 - 6. Cost Modeling and Financial Analysis**
 - 7. Security and Compliance in the Cloud**
 - 8. Network and Connectivity Planning**
 - 9. Risk Assessment and Mitigation**
 - 10. Building Your Migration Roadmap**
 - 11. The 90-Day Quick-Start Plan**
 - 12. Post-Migration Optimization**
-

Chapter 1

Is Cloud Right for Your Organization?

Cloud migration is not an end in itself. It is a means to achieve specific business outcomes: reducing capital expenditure, increasing agility, improving disaster recovery, enabling remote work, or accessing capabilities that are impractical to build on-premise. If your migration is not driven by clear business objectives, it will stall, overspend, or deliver the wrong outcomes.

Not every workload belongs in the cloud. Some applications perform better on-premise, cost less on-premise, or face regulatory constraints that complicate cloud deployment. The goal of a readiness assessment is not to confirm a foregone conclusion but to objectively evaluate which workloads should move, which should stay, and in what order.

Common Business Drivers for Cloud Migration

- Capital expense reduction: eliminate hardware refresh cycles and shift to predictable monthly operating expenses
- Scalability: scale compute and storage on demand without capacity planning months in advance
- Disaster recovery: leverage cloud-native replication and multi-region failover at a fraction of the cost of a secondary data center
- Remote workforce enablement: provide consistent access to applications regardless of user location
- Innovation velocity: access AI/ML services, serverless compute, managed databases, and other capabilities without building infrastructure
- End-of-life hardware: aging on-premise equipment approaching end of support with no budget for replacement

When Cloud May Not Be the Answer

- Ultra-low latency requirements: applications that require sub-millisecond response times may not tolerate network round-trips to a cloud region
- Data sovereignty: some regulatory environments require data to remain within specific physical locations that cloud providers cannot guarantee
- Predictable high-utilization workloads: a server running at 90% utilization 24/7 is almost always cheaper on-premise than in the cloud
- Legacy applications: applications that are tightly coupled to specific hardware, OS versions, or licensing models may cost more to run in cloud than to maintain on-premise

The Hybrid Reality

Most organizations do not move everything to the cloud. The realistic outcome for most mid-market companies is a hybrid environment: some workloads in the cloud, some on-premise, and some in SaaS. Your migration roadmap should embrace this reality rather than force a pure-cloud ideology.

Chapter 2

Cloud Readiness Self-Assessment

Before migrating anything, assess your organization's readiness across five dimensions. Score each item 0-3 (Not Ready / Partially Ready / Ready / Advanced) to identify areas that need attention before migration begins.

Organizational Readiness

- ☐ Executive sponsorship exists with a clear mandate and budget for cloud migration
- ☐ Cloud migration aligns with documented business objectives, not just IT preferences
- ☐ Cross-functional stakeholders (IT, finance, compliance, business units) are engaged in planning
- ☐ Change management plan exists for transitioning teams to cloud-based workflows
- ☐ Decision-making authority and governance structure defined for the migration program

Technical Readiness

- ☐ Complete application and infrastructure inventory exists with documented dependencies
- ☐ Network bandwidth sufficient to support cloud connectivity (or upgrade plan in place)
- ☐ Identity infrastructure (Active Directory / Azure AD / Okta) ready for cloud integration
- ☐ Current backup and DR capabilities documented with RPO/RTO for each system
- ☐ Monitoring and alerting tools capable of managing hybrid (on-premise + cloud) environments

Skills Readiness

- ☐ IT team has foundational cloud skills (at least one certified cloud practitioner)
- ☐ Training plan exists to upskill IT staff on cloud architecture, security, and operations
- ☐ Gap analysis completed: skills that cannot be developed internally are covered by a partner or MSP
- ☐ Operations team trained on cloud-native monitoring, cost management, and incident response

Financial Readiness

- ☐ Total cost of ownership (TCO) analysis completed comparing on-premise vs. cloud for target workloads
- ☐ Budget allocated for migration project costs (separate from ongoing cloud operating costs)
- ☐ Cost governance plan in place: tagging strategy, budget alerts, and monthly review process
- ☐ Finance team understands the shift from CapEx to OpEx and its impact on financial reporting

Security and Compliance Readiness

- ☐ Regulatory requirements for cloud deployment identified (data residency, encryption, audit logging)
- ☐

Cloud security policy drafted covering identity, network, data protection, and incident response

☐ Shared responsibility model understood: what the cloud provider secures vs. what you secure

☐ Vendor security assessment completed for target cloud provider(s)

Readiness Score

If you score below 50% in any dimension, address the gaps before beginning migration. Starting a cloud migration without organizational alignment, technical foundations, or financial governance leads to cost overruns, security incidents, and stalled projects.

Chapter 3

Workload Classification and Prioritization

Not all workloads are equal in terms of migration complexity, business value, or risk. A workload classification matrix helps you decide what to migrate first, what to defer, and what to leave on-premise.

Classification Dimensions

Evaluate each workload across four dimensions:

Dimension	Low (1)	Medium (2)	High (3)
Business Value	Internal tool	Department-critical	Revenue-generating
Migration Complexity	Lift-and-shift ready	Some refactoring	Major rearchitecture
Risk Level	Non-sensitive data	Internal data	Regulated/PII data
Technical Debt	Modern stack	Some legacy	Legacy-dependent

Prioritization Matrix

Plot each workload on a 2x2 matrix of Business Value (Y-axis) vs. Migration Complexity (X-axis):


- High Value + Low Complexity (top-left): Migrate first. Quick wins that deliver immediate business value.
- High Value + High Complexity (top-right): Plan carefully. These are strategic migrations that require investment but deliver significant returns.
- Low Value + Low Complexity (bottom-left): Migrate opportunistically. Easy to move but low priority.
- Low Value + High Complexity (bottom-right): Defer or retire. The effort outweighs the benefit. Consider decommissioning instead of migrating.

Workload Inventory Template

Workload	Bus. Value	Complexity	Risk	Tech Debt	Decision
Example: ERP	High (3)	High (3)	High (3)	Med (2)	Phase 3
Example: File Share	Med (2)	Low (1)	Med (2)	Low (1)	Phase 1
Example: Dev/Test	Low (1)	Low (1)	Low (1)	Low (1)	Phase 1
Example: PACS	High (3)	High (3)	High (3)	High (3)	Defer
_____	___	___	___	___	___
_____	___	___	___	___	___
_____	___	___	___	___	___

Start with Dev/Test

Development and test environments are almost always the best first migration candidates. They are low-risk, low-complexity, and give your team hands-on cloud experience before touching production workloads. The lessons learned during a dev/test migration will save time and money on every subsequent migration.



Start with Dev/Test

Development and test environments are almost always the best first migration candidates. They are low-risk, low-complexity, and give your team hands-on cloud experience before touching production workloads. The lessons learned during a dev/test migration will save time and money on every subsequent migration.

Chapter 4

The 6 R's of Cloud Migration

Every workload requires a migration strategy. The industry-standard framework defines six approaches, commonly called the "6 R's." The right strategy depends on the workload's classification, technical architecture, and business objectives.

1. Rehost (Lift and Shift)

Move the workload to cloud infrastructure as-is, with minimal or no changes. The application runs on cloud VMs instead of physical servers.

- Best for: Applications with tight timelines, low strategic value, or pending replacement
- Pros: Fastest migration path, lowest upfront effort, immediate CapEx elimination
- Cons: Does not leverage cloud-native capabilities, may cost more than on-premise if not right-sized

2. Replatform (Lift, Tinker, and Shift)

Make targeted optimizations during migration to take advantage of cloud capabilities without rearchitecting the entire application.

- Best for: Applications that benefit from managed services (e.g., move from self-managed SQL Server to Azure SQL or RDS)
- Pros: Meaningful cost and operational benefits with moderate effort
- Cons: Requires deeper analysis of application architecture and dependencies

3. Repurchase (Drop and Shop)

Replace the existing application with a SaaS equivalent. The migration is not a technical lift but a business transition.

- Best for: Commodity applications where SaaS is mature (email, CRM, ERP, HR, collaboration)
- Pros: Eliminates infrastructure management entirely, often improves user experience
- Cons: Data migration, process change, and user retraining required; potential vendor lock-in

4. Refactor (Rearchitect)

Redesign the application to be cloud-native, using containers, serverless, microservices, or managed services.

- Best for: Strategic applications with long lifespans that justify the investment in rearchitecture
- Pros: Maximum cloud benefits (scalability, resilience, cost optimization), reduced operational overhead
- Cons: Highest effort and cost; requires cloud-native development skills

5. Retire

Decommission the workload entirely. Migration is an opportunity to eliminate applications that are redundant, underutilized, or no longer needed.

- Typical finding: 10-20% of applications discovered during inventory are candidates for retirement

6. Retain

Keep the workload on-premise, at least for now. Some workloads are not suitable for cloud migration due to latency, licensing, regulatory, or cost considerations.

- Reassess retained workloads annually as cloud capabilities evolve and on-premise hardware ages

The Lift-and-Shift Trap

Rehosting is the fastest path but often the most expensive in the long run. A server running 24/7 in the cloud costs more than the same server on-premise if it is not right-sized and scheduled. If you rehost, commit to a post-migration optimization pass (Chapter 12) within 90 days.

Chapter 5

Cloud Provider Selection

The three major cloud providers (AWS, Microsoft Azure, Google Cloud) each have strengths. Your choice should be driven by your workload requirements, existing technology investments, and team expertise, not by marketing.

Decision Framework

Factor	AWS	Azure	Google Cloud
Best For	Broadest services	Microsoft stack	Data/AI workloads
Identity	IAM + SSO	Azure AD native	Cloud Identity
Hybrid	Outposts	Azure Arc	Anthos
Enterprise Support	Enterprise tier	Unified Support	Premium Support
Pricing Model	On-demand + RI + SP	On-demand + RI	On-demand + CUD + SUD
Market Share	Largest	Second largest	Third largest

Key Selection Criteria

- Existing investments: If your organization runs Microsoft 365 and Active Directory, Azure provides the tightest integration. If you are already using AWS services, expanding within AWS reduces complexity.
- Workload requirements: specific managed services (e.g., Aurora for PostgreSQL, Azure SQL, BigQuery) may favor one provider over another
- Geographic coverage: verify that the provider has regions that meet your latency and data residency requirements
- Partner ecosystem: does your MSP or technology partner have deep expertise on the platform?
- Pricing and licensing: Microsoft customers may benefit from Azure Hybrid Benefit (use existing Windows Server and SQL licenses in the cloud)

Multi-Cloud Considerations

Multi-cloud (using multiple providers simultaneously) provides vendor diversification but increases complexity. For most mid-market organizations, a single primary cloud provider with SaaS services from other providers (e.g., Azure for infrastructure, AWS for a specific managed service) is the pragmatic approach. True multi-cloud architectures require significant investment in cross-cloud networking, identity federation, and operational tooling.

The Honest Answer

For most mid-market organizations running Microsoft 365, the honest answer is Azure. The identity integration,

licensing benefits, and operational simplicity of staying within the Microsoft ecosystem outweigh the marginal technical advantages of other platforms for general workloads. Choose a different provider when you have a specific technical reason to do so.



The Honest Answer

For most mid-market organizations running Microsoft 365, the honest answer is Azure. The identity integration,

licensing benefits, and operational simplicity of staying within the Microsoft ecosystem outweigh the marginal technical advantages of other platforms for general workloads. Choose a different provider when you have a specific technical reason to do so.

Chapter 6

Cost Modeling and Financial Analysis

Cloud cost management is the number one challenge reported by organizations post-migration. Without rigorous cost modeling before migration and continuous optimization after, cloud can easily cost more than on-premise.

Total Cost of Ownership (TCO) Comparison

Build a 3-year TCO comparison for each workload under consideration:


Cost Category	On-Premise	Cloud
Compute (servers/VMs)	\$_____	\$_____
Storage	\$_____	\$_____
Networking/Bandwidth	\$_____	\$_____
Licensing	\$_____	\$_____
Data Center (power/cooling)	\$_____	Included
Hardware refresh (amortized)	\$_____	N/A
IT staff (ops overhead)	\$_____	\$_____
Backup and DR	\$_____	\$_____
Migration project cost	N/A	\$_____
3-YEAR TOTAL	\$_____	\$_____

Cost Optimization Levers

- Right-sizing: match VM sizes to actual utilization, not peak capacity. Most on-premise servers are provisioned at 2-4x their actual workload.
- Reserved instances / Savings Plans: commit to 1-3 year terms for predictable workloads and save 30-60% over on-demand pricing
- Auto-scaling: scale compute up during business hours and down overnight/weekends for workloads with variable demand
- Spot/preemptible instances: use discounted interruptible capacity (60-90% savings) for dev/test, batch processing, and fault-tolerant workloads
- Storage tiering: move infrequently accessed data to cold storage tiers (S3 Glacier, Azure Cool/Archive) at 50-90% lower cost
- Managed services: replace self-managed databases and middleware with managed equivalents to reduce operational overhead

The Cloud Cost Surprise

The most common post-migration complaint is "cloud costs more than we expected." This almost always traces to one of three causes: VMs sized to match on-premise specs instead of actual utilization, no reserved instances purchased for steady-state workloads, or data egress charges that were not modeled. Account for all three in your cost model.



The Cloud Cost Surprise

The most common post-migration complaint is "cloud costs more than we expected." This almost always traces to one of three causes: VMs sized to match on-premise specs instead of actual utilization, no reserved instances purchased for steady-state workloads, or data egress charges that were not modeled. Account for all three in your cost model.

Chapter 7

Security and Compliance in the Cloud

Cloud security operates on a shared responsibility model: the cloud provider secures the infrastructure, and you secure everything you put on it. Misunderstanding this boundary is the leading cause of cloud security incidents.

The Shared Responsibility Model

- Provider's responsibility: physical security, host OS, hypervisor, network infrastructure, and availability of the platform
- Your responsibility: data classification and protection, identity and access management, OS patching (for IaaS), application security, network configuration, encryption key management, and compliance
- Shared responsibility: varies by service model. With IaaS you manage more; with PaaS less; with SaaS least

Cloud Security Controls Checklist

- ☐ Cloud security posture management (CSPM) tool deployed for continuous misconfiguration detection
- ☐ All cloud admin accounts protected with MFA and conditional access; root accounts locked with break-glass procedures
- ☐ IAM policies enforce least-privilege with no wildcard permissions
- ☐ Network security groups and firewall rules restrict access to minimum required; no 0.0.0.0/0 inbound rules
- ☐ All data encrypted at rest using provider-managed or customer-managed keys
- ☐ All data encrypted in transit using TLS 1.2+
- ☐ Cloud audit logging enabled (CloudTrail, Azure Monitor, GCP Audit Logs) and forwarded to SIEM
- ☐ Public access to storage and databases blocked by default; exceptions require approval and monitoring
- ☐ Vulnerability scanning configured for cloud-hosted VMs and container images
- ☐ Backup and snapshot policies configured and tested for all cloud workloads

Compliance Considerations

- Data residency: verify that your cloud region meets regulatory requirements for data location
- Audit logging retention: some regulations require log retention beyond provider defaults (configure accordingly)
- Encryption key management: some regulations require customer-managed keys rather than provider-managed

- BAA/DPA: ensure your cloud provider agreement includes applicable regulatory addenda (HIPAA BAA, GDPR DPA)

Cloud Misconfiguration

Misconfigured cloud resources (public S3 buckets, open security groups, overprivileged IAM roles) are responsible for more cloud data breaches than sophisticated attacks. A CSPM tool that continuously scans for misconfigurations and alerts your team is a non-negotiable cloud security control.

Chapter 8

Network and Connectivity Planning

Cloud migration fundamentally changes your network architecture. Applications that were accessed across the LAN are now accessed across the WAN or internet. Network planning failures cause the most user-visible problems after migration.

Connectivity Options

- Public internet: simplest path, lowest cost, but variable latency and no SLA. Acceptable for non-sensitive SaaS access and development workloads.
- VPN (site-to-site): encrypted tunnel over the internet. Provides private connectivity at low cost but inherits internet latency variability. Suitable for low-to-medium bandwidth workloads.
- Direct connect (AWS) / ExpressRoute (Azure) / Interconnect (GCP): dedicated private connection with consistent latency and bandwidth SLA. Required for latency-sensitive production workloads, database replication, and large data transfers.

Bandwidth Planning

- Measure current WAN utilization to understand baseline traffic patterns
- Model additional bandwidth required for cloud-bound traffic that currently stays on the LAN
- Account for data migration bandwidth: initial bulk transfer may saturate circuits for days or weeks
- Plan for data egress costs: cloud providers charge for data leaving the cloud, not entering it. Model egress charges for backups, replication, and user access.

DNS and Name Resolution


- Plan DNS changes for migrated services: update internal DNS to point to new cloud endpoints
- Implement split-horizon DNS if internal and external users need to resolve the same names differently
- Consider DNS-based failover for services with both on-premise and cloud instances during migration

Hybrid Networking

- Design the cloud virtual network (VPC/VNet) with IP address ranges that do not overlap with on-premise networks
- Implement routing between on-premise and cloud networks through VPN or direct connect
- Apply consistent firewall and network segmentation policies across on-premise and cloud environments
- Plan for the transition period when workloads exist in both locations with dependencies in each direction

Data Egress Costs

Cloud providers charge for data leaving the cloud. This catches many organizations by surprise. Backup replication to an on-premise target, user access to cloud-hosted file shares, and cross-region replication all incur egress charges. Model these costs before migration, not after.



Data Egress Costs

Cloud providers charge for data leaving the cloud. This catches many organizations by surprise. Backup replication to an on-premise target, user access to cloud-hosted file shares, and cross-region replication all incur egress charges. Model these costs before migration, not after.

Chapter 9

Risk Assessment and Mitigation

Every migration carries risk. A structured risk assessment identifies potential failure modes and establishes mitigation strategies before they materialize.

Migration Risk Register

Risk	Likelihood	Impact	Mitigation
Data loss during migration	Low	Critical	Full backup before migration; verify restore
Extended downtime	Medium	High	Migration during maintenance window; rollback plan
Application incompatibility	Medium	High	Pre-migration testing in cloud staging env
Cost overrun	High	Medium	Detailed cost model; budget alerts; monthly review
Performance degradation	Medium	High	Right-sizing; latency testing; direct connect
Security misconfiguration	Medium	Critical	CSPM tool; security review before go-live
Skill gaps	High	Medium	Training plan; partner support; co-managed ops
Vendor lock-in	Medium	Medium	Use portable formats; document exit procedures

Rollback Planning

Every migration must have a rollback plan that can return the workload to its pre-migration state within a defined time window:

- Maintain the on-premise environment in a functional state for at least 30 days after each migration wave
- Document the rollback procedure with the same step-by-step detail as the migration procedure
- Define rollback triggers: what conditions (performance, availability, data integrity) would initiate a rollback?
- Test the rollback procedure during the migration rehearsal, not for the first time during a live failure

Communication Plan

- Notify all affected users of the migration timeline, expected downtime, and what to do if they experience issues
- Provide a status page or notification channel during migration execution
- Define escalation procedures for migration-related issues in the first 72 hours after go-live
- Schedule a post-migration review within one week to address any outstanding issues

Chapter 10

Building Your Migration Roadmap

A migration roadmap sequences workload migrations into waves that balance risk, resource capacity, and business impact. Each wave should be small enough to execute and roll back within a defined window.

Roadmap Structure

Wave 0: Foundation (Weeks 1-4)

- Establish cloud account and governance structure (landing zone)
- Configure identity federation (Azure AD / Okta SSO to cloud)
- Deploy network connectivity (VPN or direct connect)
- Configure security controls: CSPM, logging, alerting, backup policies
- Deploy monitoring and cost management tooling
- Train IT team on cloud operations fundamentals

Wave 1: Low-Risk Workloads (Weeks 5-10)

- Migrate dev/test environments
- Migrate file shares and collaboration workloads (if applicable)
- Migrate non-critical internal applications
- Validate monitoring, backup, and security controls in production cloud environment
- Conduct lessons-learned review and refine migration playbook

Wave 2: Business Applications (Weeks 11-18)

- Migrate business-critical applications based on workload classification priorities
- Execute migration during maintenance windows with rollback plans tested
- Validate application performance, user experience, and integration points
- Right-size cloud resources based on observed utilization (not on-premise sizing)

Wave 3: Complex/Critical Systems (Weeks 19-26+)

- Migrate mission-critical systems (ERP, databases, high-availability workloads)
- Implement cloud-native HA and DR for migrated production workloads
- Decommission on-premise infrastructure for migrated workloads
- Conduct comprehensive post-migration optimization review

Wave Size

Each wave should contain no more than 5-10 workloads. Smaller waves are easier to execute, roll back, and learn from. Resist the temptation to combine waves for speed. A failed large wave causes more disruption than two successful small ones.

Chapter 11

The 90-Day Quick-Start Plan

This accelerated plan provides a practical week-by-week roadmap for organizations beginning their first cloud deployment. It covers Wave 0 (foundation) and Wave 1 (first workload migration) in a structured 90-day timeline.

Days 1-14: Plan and Prepare

- ☐ Define cloud migration objectives and success criteria
- ☐ Complete workload inventory and classification for first wave candidates
- ☐ Select cloud provider and engagement model (self-managed, co-managed, or fully managed)
- ☐ Establish cloud account with proper governance (organization structure, billing, IAM)
- ☐ Order direct connect or upgrade internet bandwidth if required
- ☐ Begin team training on cloud fundamentals

Days 15-30: Build the Foundation

- ☐ Deploy cloud landing zone: networking (VPC/VNet), subnets, security groups, routing
- ☐ Configure identity federation: SSO, MFA, conditional access policies
- ☐ Establish site-to-site VPN or direct connect between on-premise and cloud
- ☐ Deploy security controls: CSPM, audit logging, alerting, encryption policies
- ☐ Configure backup and snapshot policies for cloud workloads
- ☐ Deploy monitoring and cost management dashboards

Days 31-60: First Migration

- ☐ Migrate first workload (dev/test environment recommended) to cloud
- ☐ Validate connectivity, performance, and user access from all locations
- ☐ Verify backup, recovery, and monitoring are working as expected
- ☐ Document migration runbook based on lessons learned
- ☐ Migrate 2-3 additional low-risk workloads using refined runbook
- ☐ Right-size VMs based on observed utilization (not on-premise specs)

Days 61-90: Validate and Optimize

- ☐ Conduct 30-day performance review for all migrated workloads
- ☐ Review cloud costs against model; identify and address variances
- ☐ Purchase reserved instances or savings plans for steady-state workloads
- ☐

Conduct security review: CSPM findings, IAM policies, network rules

- ☐ Decommission on-premise resources for successfully migrated workloads
- ☐ Plan Wave 2 based on lessons learned and updated workload classifications
- ☐ Present 90-day results to executive sponsor with roadmap for next phase

Day 90 Deliverable

At the end of 90 days, you should have: a production cloud environment with governance and security controls in place, 3-5 workloads successfully migrated, validated cost model with actual spend data, a trained operations team, and a refined roadmap for the remaining migration waves.

Chapter 12

Post-Migration Optimization

Migration is not the finish line. The first 90 days after each wave are critical for optimizing cost, performance, and security. Organizations that skip post-migration optimization leave 30-40% of potential cloud savings on the table.

Cost Optimization

- Right-size all VMs based on 30 days of actual utilization data (CPU, memory, disk I/O, network)
- Identify and shut down unused resources: orphaned disks, unattached IPs, idle load balancers, stopped VMs still incurring storage charges
- Implement auto-scaling for workloads with variable demand
- Purchase reserved instances or savings plans for workloads with predictable, steady-state utilization
- Move infrequently accessed data to lower-cost storage tiers
- Establish monthly cost review cadence with tagging enforcement for cost allocation

Performance Optimization

- Review application response times and compare against pre-migration baselines
- Identify and address latency-sensitive paths (database queries, API calls, user-facing applications)
- Optimize storage performance: match disk types (SSD vs. HDD, provisioned IOPS vs. standard) to workload requirements
- Implement caching (CDN for web content, in-memory caching for database queries) where applicable
- Review network throughput and adjust instance types or bandwidth allocations if needed

Security Hardening

- Review and remediate all CSPM findings from the first 30 days of production operation
- Audit IAM policies and remove any overly permissive access granted during migration
- Verify that all logging, monitoring, and alerting is functioning and reaching the SIEM
- Conduct a vulnerability scan of all cloud-hosted instances and remediate findings
- Review and restrict network security group rules that were opened during migration testing

Operational Maturity

- Document cloud-specific runbooks for common operations: VM provisioning, scaling, patching, backup/restore
- Establish change management procedures for cloud infrastructure changes
- Implement infrastructure-as-code (Terraform, CloudFormation, Bicep) for repeatable, auditable deployments
- Define cloud-specific incident response procedures and test them through tabletop exercise

About ComDirect USA

ComDirect USA is a full-service technology partner specializing in enterprise infrastructure, cybersecurity, cloud architecture, and managed IT services. With offices in Irvine, California, and Meridian, Idaho, we serve organizations across healthcare, legal, financial services, manufacturing, and logistics.

Ready to plan your cloud migration? Contact us.

Phone: 949.208.8800 | Toll Free: 800.733.7030

Email: info@comdirectusa.com | Web: comdirectusa.com