certstud .com

Search certifications...

Search

AWS Solutions Architect Professional Certification Study Notes

Code: SAP-C02

Organizational Complexity (26%)

Multi-Account Strategies

Enterprise-scale AWS account organization and governance

AWS Organizations

- **Consolidated Billing:** Single payment method for all accounts with combined usage discounts
- **Service Control Policies (SCPs):** Guardrails that restrict permissions across accounts - deny overrides all allows
- **Organizational Units (OUs):** Hierarchical grouping of accounts for policy application
- **All Features Mode:** Required for SCPs, tag policies, and AI services opt-out policies

> **SCP Key Concept**
>
> SCPs are permission boundaries, not grants. They limit what IAM policies can allow. Even if IAM allows an action, if SCP denies it, the action is blocked. SCPs do NOT apply to the management account.

AWS Control Tower

Landing Zone

Pre-configured, secure multi-account environment based on

Guardrails

Preventive (SCPs) and D
(AWS Config rules) cont
Mandatory, Strongly

Analytics (⌘⇧A)

Feedback

AWS best practices with logging and security accounts

Recommended, and Elective categories

Account Factory

Automated account provisioning with pre-approved configurations. Uses Service Catalog under the hood

Dashboard

Single view of compliance status, guardrail violations, and account provisioning across the organization

Account Structure Best Practices

Account Type	Purpose	Key Services
Management	Organization root, billing, SCPs (minimal workloads)	Organizations, Billing, Cost Explorer
Log Archive	Centralized logging from all accounts	CloudTrail, Config, S3, CloudWatch Logs
Security/Audit	Security tooling and cross-account audit	Security Hub, GuardDuty, IAM Access Analyzer
Shared Services	Common infrastructure (AD, DNS, CI/CD)	Directory Service, Route 53, CodePipeline
Network	Transit Gateway, Direct Connect, VPN	Transit Gateway, Direct Connect, VPN
Sandbox	Experimentation with limited budget	Service Catalog, Budgets, Cost Controls
Workload (Dev/Prod)	Application environments	Application-specific services

Cross-Account Access Patterns

Secure access patterns between AWS accounts

IAM Role Assumption

- **AssumeRole:** Primary method for cross-account access. Role in target account trusts principal from source account
- **External ID:** Prevents confused deputy problem when third parties assume roles. Always use for external access
- **Session Policies:** Further restrict assumed role permissions for specific sessions
- **Chained Roles:** Maximum of 1 hour session when role assumes another role

Resource-Based Policies

Service	Resource Policy	Cross-Account Pattern
S3	Bucket Policy	Grant access to principals from other accounts
KMS	Key Policy	Allow external accounts to use CMKs
SNS/SQS	Access Policy	Allow cross-account publish/subscribe
Lambda	Resource Policy	Allow invocation from other accounts
Secrets Manager	Resource Policy	Share secrets across accounts
ECR	Repository Policy	Share container images across accounts

Identity vs Resource Policies

> **Identity-based:** Attach to IAM users/roles - "what can this identity do?"
>
> **Resource-based:** Attach to resources - "who can access this resource?" Cross-account access is often simpler with resource policies (no role assumption needed for some services)

AWS Resource Access Manager (RAM)

- **Purpose:** Share AWS resources across accounts within or outside your organization
- **Shareable Resources:** Transit Gateways, Subnets, License Manager configs, Route 53 Resolver rules, and more
- **Organization Sharing:** Enable sharing within organization for automatic acceptance
- **Permissions:** Shared resources retain original owner's permissions; participants get usage rights

Hybrid & Multi-Region Networking

Connecting on-premises and multi-region AWS environments

AWS Direct Connect

Dedicated Connection

1 Gbps, 10 Gbps, or 100 Gbps. Physical connection at Direct Connect location. Weeks to provision

Hosted Connection

50 Mbps to 10 Gbps through AWS Partner. Faster provisioning. Add/remove capacity on demand

Virtual Interfaces (VIFs)

Public VIF (AWS public services), Private VIF (VPC), Transit VIF (Transit Gateway)

Direct Connect Gateway

Connect Direct Connect to multiple VPCs across regions. Does NOT provide VPC-to-VPC routing

Direct Connect + VPN

For encrypted traffic over Direct Connect, use Site-to-Site VPN over Direct Connect public VIF. Direct Connect itself does NOT encrypt traffic in transit.

AWS Transit Gateway

- **Hub-and-Spoke:** Regional router connecting VPCs, VPNs, Direct Connect, and SD-WAN
- **Route Tables:** Multiple route tables for network segmentation. Control which attachments can communicate
- **Inter-Region Peering:** Connect Transit Gateways across regions (encrypted, uses AWS backbone)
- **Multicast:** Only AWS service supporting multicast. Useful for media streaming and financial data feeds
- **Network Manager:** Global view of private network across AWS and on-premises

Connectivity Options Comparison

Option	Bandwidth	Latency	Encryption	Best For
Site-to-Site VPN	Up to 1.25 Gbps per tunnel	Variable (internet)	Yes (IPsec)	Quick setup, backup connection
Direct Connect	Up to 100 Gbps	Consistent, low	No (use VPN overlay)	Production workloads, large data transfer
VPC Peering	No limit (intra-region)	Lowest	Yes (in-transit)	Simple VPC-to-VPC, non-transitive
Transit Gateway	50 Gbps per attachment	Low	VPN attachments only	Complex multi-VPC, hybrid networks

Option	Bandwidth	Latency	Encryption	Best For
PrivateLink	Endpoint bandwidth	Low	Yes	Private access to services

Enterprise Identity Federation

Integrating corporate identity with AWS

AWS IAM Identity Center (SSO)

- **Centralized Access:** Single sign-on to all AWS accounts and business applications
- **Identity Sources:** Built-in directory, Active Directory, or external IdP (Okta, Azure AD)
- **Permission Sets:** Collections of IAM policies assigned to users/groups for specific accounts
- **Attribute-Based Access Control:** Use user attributes from IdP for dynamic permissions

Federation Options

Method	Use Case	Token Duration	Key Points
IAM Identity Center	Workforce access to multiple accounts	Configurable	Preferred for organization-w
SAML 2.0 Federation	Enterprise IdP integration	Up to 12 hours	AssumeRoleWi console acces
Web Identity Federation	Mobile/web apps, social login	Up to 12 hours	Cognito or dire (Google, Faceb

Method	Use Case	Token Duration	Key Points
Custom Identity Broker	Legacy systems, special requirements	Configurable	Your code calls AssumeRole
AWS Directory Service	Windows workloads, AD integration	N/A	Managed AD, A Connector, Sim

Cognito for Applications

User Pools: User directory for sign-up/sign-in. Returns JWT tokens. Good for app authentication.

Identity Pools: Exchange tokens for temporary AWS credentials. Provides direct AWS service access.

AWS Directory Service Options

AWS Managed Microsoft AD

Full Microsoft AD, supports trusts with on-premises AD, MFA, Group Policy. Best for enterprise Windows workloads

AD Connector

Proxy to on-premises AD. No caching, no user storage in AWS. Best when you must keep all AD on-premises

Simple AD

Standalone Samba-based AD. No trust relationships. Best for small workloads, Linux admin access

Design for New Solutions (29%)

Compute Architecture Decisions

Selecting the right compute services for enterprise workloads

Compute Options Decision Matrix

Workload Type	Best Compute Option
Stateless, event-driven, short-running	Lambda
Containerized microservices	ECS/EKS with Fargate
Long-running containers with GPU/specific instance	ECS/EKS with EC2
Lift-and-shift, full OS control	EC2
Batch processing, interruptible jobs	EC2 Spot + Batch
High-performance computing (HPC)	EC2 + Cluster Placement Group + EFA

EC2 Instance Selection

General Purpose (M, T)

Compute Optimized (C)

Balanced compute/memory/networking. T instances for burstable workloads with CPU credits

High-performance processors. Batch processing, gaming, HPC, scientific modeling

Memory Optimized (R, X, z1d)

In-memory databases, real-time big data analytics. X instances up to 4TB RAM

Storage Optimized (I, D, H)

High sequential read/write, data warehousing. NVMe SSD with high IOPS

Accelerated Computing (P, G, Inf, Trn)

GPU for ML training, graphics. Inferentia/Trainium for ML inference

HPC Optimized (Hpc)

Tightly-coupled HPC workloads. EFA-enabled for low-latency networking

Placement Groups

Cluster: Same rack, lowest latency. HPC, tightly-coupled applications

Spread: Different racks, max 7 per AZ. Critical instances requiring isolation

Partition: Groups on separate racks. Large distributed systems (Hadoop, Kafka)

Storage Architecture Decisions

Choosing storage for performance, durability, and cost

Storage Services Comparison

Service	Type	Use Cases	Key Limits
EBS	Block	EC2 boot/data volumes,	64 TiB, single AZ, single EC2

Service	Type	Use Cases	Key Limits
		databases	(except Multi-Attach)
EFS	File (NFS)	Shared file storage, CMS, containers	Petabytes, multi-AZ, thousands of connections
FSx for Windows	File (SMB)	Windows workloads, SQL Server, SharePoint	64 TiB, Active Directory integrated
FSx for Lustre	File (HPC)	ML training, HPC, video processing	Hundreds of GB/s throughput, S3 integration
FSx for NetApp ONTAP	File (multi-protocol)	Enterprise file shares, lift-and-shift	NFS, SMB, iSCSI, deduplication
S3	Object	Static content, backups, data lakes	Unlimited, 5 TB per object, 11 9s durability
S3 Glacier	Object (archive)	Long-term archive, compliance	Minutes to hours retrieval, lowest cost

EBS Volume Types

gp3 (General Purpose SSD)

3,000 IOPS baseline, up to 16,000. Independent IOPS/throughput provisioning. Best general-purpose choice

io2 Block Express

Up to 256,000 IOPS, 64 TiB. Sub-millisecond latency. Mission-critical databases

st1 (Throughput HDD)

500 MiB/s max throughput. Sequential workloads, big data, log processing

sc1 (Cold HDD)

Lowest cost, 250 MiB/s. Infrequently accessed, cold data

EBS Multi-Attach

Only io1/io2 volumes support Multi-Attach (up to 16 Nitro instances). Applications must manage concurrent write operations. Use for clustered applications like Oracle RAC.

S3 Storage Classes

- **S3 Standard:** Frequently accessed data, millisecond access, 99.99% availability
- **S3 Intelligent-Tiering:** Unknown/changing access patterns, automatic cost optimization
- **S3 Standard-IA:** Infrequent access but rapid retrieval needed, 30-day minimum
- **S3 One Zone-IA:** Infrequent, non-critical, recreatable data, single AZ
- **S3 Glacier Instant:** Archive with millisecond retrieval, quarterly access pattern
- **S3 Glacier Flexible:** Archive, minutes to hours retrieval, 90-day minimum
- **S3 Glacier Deep Archive:** Lowest cost, 12-hour retrieval, 180-day minimum, compliance archives

Database Architecture Decisions

Selecting purpose-built databases for different workloads

Purpose-Built Database Selection

Database	Type	Best For	Key Features
Aurora	Relational	Enterprise apps, MySQL/PostgreSQL compatible	5x MySQL performance, 15 read replicas, Global Database

Database	Type	Best For	Key Features
RDS	Relational	Traditional apps, multiple engine choices	Managed, Multi-AZ, automated backups
DynamoDB	Key-Value/Document	Serverless apps, gaming, IoT	Single-digit ms latency, unlimited scale, Global Tables
ElastiCache	In-Memory	Caching, session stores, leaderboards	Redis or Memcached, microsecond latency
DocumentDB	Document	MongoDB compatibility	Managed, scales to millions of requests/sec
Keyspaces	Wide Column	Cassandra compatibility	Serverless, single-digit ms latency at any scale
Neptune	Graph	Social networks, fraud detection, recommendations	Billions of relationships, millisecond queries
Timestream	Time Series	IoT, DevOps, industrial telemetry	1000x faster/10x cheaper than relational for time series
QLDB	Ledger	Immutable records, audit, supply chain	Cryptographically verifiable transaction log
Redshift	Data Warehouse	Analytics, BI, large-scale aggregations	Petabyte scale, columnar

Database	Type	Best For	Key Features
			storage, Spectrum for S3

Aurora Advanced Features

Aurora Global Database

Cross-region replication with < 1 second lag. Fast disaster recovery, low-latency global reads

Aurora Serverless v2

Instant scaling, pay per ACU-second. Variable/unpredictable workloads, dev/test

Aurora Multi-Master

All nodes can read AND write. Continuous availability for writes (MySQL only)

Aurora Machine Learning

SQL interface to SageMaker and Comprehend. Add ML predictions to queries

DynamoDB Advanced Patterns

Global Tables: Multi-region, active-active. < 1 second replication, automatic conflict resolution

DAX: In-memory cache, microsecond latency. 10x read performance improvement

Streams: Ordered record of modifications. Event-driven architectures, cross-region replication

High Availability & Disaster Recovery

Designing for resilience and business continuity

DR Strategies (by RTO/RPO)

Strategy	RTO	RPO	Cost	Description
Backup & Restore	Hours	Hours	Lowest	Backup to S3/Glacier, restore when needed
Pilot Light	10s of minutes	Minutes	Low	Minimal core systems always running, scale up on disaster
Warm Standby	Minutes	Seconds	Medium	Scaled-down version running, ready to scale up
Multi-Site Active/Active	Near-zero	Near-zero	Highest	Full production in multiple regions

Route 53 Routing Policies

Simple

Single resource, multiple values return all (client chooses). No health checks

Weighted

Distribute traffic by weight (e.g., 70/30). Blue-green deployments, A/B testing

Latency-based

Route to lowest latency region. Multi-region active-active

Failover

Active-passive with health checks. Primary fails → secondary

Geolocation

Route based on user location. Content localization, compliance

Geoproximity

Route based on resource location with bias. Traffic flow policies

Multi-Value

Up to 8 healthy records returned. Simple load balancing with health checks

IP-based

Route based on client IP ranges (CIDR). ISP-specific routing

Multi-Region Considerations

Data replication lag is your RPO. Consider: Aurora Global Database (< 1s), DynamoDB Global Tables (< 1s), S3 Cross-Region Replication (minutes), custom replication strategies.

Continuous Improvement (25%)

Cost Optimization Strategies

Enterprise cost management and optimization

EC2 Pricing Models

Model	Discount	Commitment	Best For
On-Demand	0%	None	Unpredictable workloads, short-term, flexibility

Model	Discount	Commitment	Best For
Savings Plans (Compute)	Up to 66%	1 or 3 years, $/hour	Flexible across instance family, region, OS
Savings Plans (EC2 Instance)	Up to 72%	1 or 3 years, specific family	Steady-state workloads, known instance needs
Reserved Instances	Up to 72%	1 or 3 years, specific attributes	Stable workloads, specific instance requirements
Spot Instances	Up to 90%	None (can be interrupted)	Fault-tolerant, flexible, batch, CI/CD
Dedicated Hosts	0% (or Reserved)	None or 1/3 year	Licensing, compliance, regulatory requirements

Cost Management Tools

AWS Cost Explorer

Visualize spending, identify trends, Reserved Instance recommendations. 13 months of data

AWS Budgets

Set custom cost/usage budgets. Alerts via SNS, auto-actions (stop EC2, apply SCP)

Cost and Usage Report (CUR)

Compute Optimizer

Most detailed billing data. Hourly/resource-level. Integrate with Athena, QuickSight

ML-based rightsizing recommendations for EC2, EBS, Lambda. Identifies over-provisioned resources

Trusted Advisor

Best practice checks including cost optimization. Idle resources, reserved instance utilization

Cost Allocation Tags

User-defined tags for cost tracking. Activate in Billing console for reporting

> #### Savings Plans vs Reserved Instances
>
> **Savings Plans:** More flexible, apply to EC2, Fargate, Lambda. Compute SP most flexible, EC2 Instance SP highest discount
>
> **Reserved Instances:** Legacy model, still valid. Capacity reservation option (zonal RI). Consider Savings Plans for new commitments

Data Transfer Cost Optimization

- **VPC Endpoints:** Avoid NAT Gateway charges for AWS service access. Gateway endpoints (S3, DynamoDB) are free
- **CloudFront:** Lower data transfer costs vs direct from origin. Regional edge caches reduce origin requests
- **Same-AZ:** Free data transfer between EC2 and RDS/ElastiCache in same AZ (use private IP)
- **Direct Connect:** Lower per-GB cost than internet. Consistent pricing regardless of volume
- **S3 Transfer Acceleration:** Higher per-GB cost but faster. Evaluate cost vs time savings

Performance Optimization

Improving application performance at scale

Caching Strategies

Layer	Service	Use Case	TTL Considerations
Edge	CloudFront	Static content, API responses	Long TTL for static, short for dynamic
API	API Gateway caching	API responses	Stage-level, 0.5GB to 237GB
Application	ElastiCache	Session, query results, computed data	Balance freshness vs hit rate
Database	DAX (DynamoDB)	DynamoDB reads	Item cache + query cache
Database	RDS Read Replicas	Read-heavy workloads	Replication lag consideration

ElastiCache Selection

Redis

Rich data structures (lists, sets, sorted sets), persistence, pub/sub, Lua scripting, cluster mode, replication

Memcached

Simple key-value, multi-threaded, no persistence, no replication. Simplest caching, disposable cache

Choose Redis When

You need: persistence, replication, complex data types, pub/sub, sorted sets (leaderboards), geospatial, Lua scripting. Redis is the default choice unless you specifically need multi-threaded simplicity.

Auto Scaling Strategies

- **Target Tracking:** Maintain metric at target value (e.g., CPU at 50%). Simplest, most common
- **Step Scaling:** Scale based on alarm breach size. Different adjustments for different thresholds

- **Scheduled Scaling:** Scale at specific times. Known traffic patterns (business hours, events)
- **Predictive Scaling:** ML-based, proactive scaling. Uses historical patterns to predict demand
- **Warm Pools:** Pre-initialized instances for faster scaling. Reduce scale-out time from minutes to seconds

Security Posture Improvement

Continuous security monitoring and improvement

Security Services Overview

Service	Purpose	Key Features
Security Hub	Centralized security view	Aggregates findings, compliance standards (CIS, PCI), automated response
GuardDuty	Threat detection	ML-based anomaly detection, VPC Flow Logs, DNS, CloudTrail analysis
Inspector	Vulnerability assessment	EC2, ECR, Lambda scanning. CVE database, network reachability
Macie	Data security	S3 sensitive data discovery (PII, financial). Automated classification
Detective	Security investigation	Root cause analysis, visualize relationships, 12-month data retention
IAM Access Analyzer	Access analysis	External access findings, policy validation, policy generation

Service	Purpose	Key Features
Config	Configuration compliance	Rules for desired state, remediation actions, aggregated multi-account

Encryption Key Management

AWS Managed Keys

Free, automatic rotation, no management needed. Limited customization, no cross-account sharing

Customer Managed Keys (CMK)

Full control, custom key policy, cross-account sharing, optional automatic rotation (yearly)

Customer Provided Keys (SSE-C)

You manage keys entirely, provide with each request. Maximum control, maximum responsibility

CloudHSM

Dedicated HSM, FIPS 140-2 Level 3, you control keys. Regulatory requirements, custom key stores

KMS Key Policies

KMS key policies are the primary way to control access. Unlike most AWS resources, IAM policies alone are NOT sufficient - the key policy must explicitly allow the IAM entity or delegate to IAM. Always include the root user to prevent lockout.

Network Security Layers

- **WAF:** Layer 7 protection. SQL injection, XSS, rate limiting, geo-blocking. Attach to CloudFront, ALB, API Gateway
- **Shield Standard:** Free DDoS protection (Layer 3/4) for all AWS customers automatically
- **Shield Advanced:** Enhanced DDoS protection, 24/7 DRT access, cost protection, $3K/month
- **Network Firewall:** Stateful inspection, intrusion prevention, managed rules. VPC-level protection

- **Firewall Manager:** Central management of WAF, Shield, Security Groups, Network Firewall across accounts

Operational Excellence

Monitoring, automation, and operational best practices

Monitoring & Observability

Service	Purpose	Key Features
CloudWatch Metrics	Performance monitoring	Custom metrics, high-resolution (1 sec), 15 months retention
CloudWatch Logs	Log aggregation	Log Insights queries, cross-account, metric filters, subscriptions
CloudWatch Alarms	Alerting	Composite alarms, anomaly detection, auto-actions
CloudWatch Dashboards	Visualization	Cross-account, cross-region, automatic dashboards
X-Ray	Distributed tracing	Service map, trace analysis, insights, sampling rules
CloudWatch Synthetics	Synthetic monitoring	Canaries for endpoints, API testing, visual monitoring
CloudWatch RUM	Real user monitoring	End-user experience, performance, errors
EventBridge	Event routing	Event patterns, cross-account, SaaS integration, archive/replay

Infrastructure as Code

CloudFormation

AWS-native IaC. StackSets for multi-account, nested stacks, drift detection, change sets

CDK

Define infrastructure in familiar languages (TypeScript, Python). Synthesizes to CloudFormation

Service Catalog

Curated product portfolios for self-service. Governance with launch constraints

Terraform

Multi-cloud IaC. State management, modules, workspaces. Popular third-party option

CloudFormation Best Practices

Nested Stacks: Reusable components, manage complexity

StackSets: Deploy across accounts and regions from single template

Change Sets: Preview changes before execution, mandatory for production

Drift Detection: Identify out-of-band changes to managed resources

Automation Services

- **Systems Manager:** Operational hub - Run Command, Session Manager, Patch Manager, State Manager, Automation
- **SSM Parameter Store:** Configuration and secrets. Free tier, hierarchy, versioning. Use SecureString for sensitive data
- **Secrets Manager:** Secrets with automatic rotation. RDS, Redshift, DocumentDB native rotation. Higher cost than Parameter Store
- **OpsWorks:** Chef/Puppet managed instances. Legacy, consider Systems Manager for new deployments
- **Elastic Beanstalk:** PaaS for web apps. Managed platform, rolling deployments, environment cloning

Migration & Modernization (20%)

The 7 Rs of Migration

Application migration strategies and decision criteria

Migration Strategies Overview

Strategy	Description	When to Use	AWS Tools
Retire	Decommission application	No longer needed, redundant	N/A
Retain	Keep on-premises (for now)	Complex dependencies, compliance, recent investment	Hybrid connectivity
Rehost (Lift & Shift)	Move as-is to cloud	Quick migration, limited changes, legacy apps	MGN, VM Import
Relocate	Move to cloud at hypervisor level	VMware workloads	VMware Cloud on AWS
Replatform	Minor optimizations during migration	Quick wins, managed services	RDS, ElastiCache, Aurora
Repurchase	Move to SaaS/different product	Commercial off-the-shelf replacements	Marketplace SaaS

Strategy	Description	When to Use	AWS Tools
Refactor/Re-architect	Redesign for cloud-native	Innovation needed, scalability, long-term investment	Containers, serverless, microservices

AWS Application Migration Service (MGN)

- **Replaces:** CloudEndure Migration and SMS (Server Migration Service)
- **Process:** Install agent → continuous replication → test instances → cutover
- **Features:** Block-level replication, non-disruptive testing, automated cutover
- **Supported:** Physical, virtual, and cloud servers. Windows and Linux

> **Migration Hub**
>
> Centralized tracking for migrations across MGN, DMS, and partner tools. Provides unified view of migration progress across applications and servers.

Database Migration

AWS DMS

Continuous replication, minimal downtime. Homogeneous and heterogeneous migrations. CDC for ongoing sync

AWS SCT

Schema Conversion Tool. Convert database schema for heterogeneous migrations. Assessment reports

Native Tools

pg_dump/restore, mysqldump, SQL Server backup/restore. Often faster for homogeneous migrations

Large-Scale Data Transfer

Moving petabytes of data to AWS

Data Transfer Options

Service	Capacity	Use Case	Timeline
Internet Transfer	Depends on connection	Small datasets, ongoing sync	Variable
Direct Connect	Up to 100 Gbps	Large ongoing transfers, hybrid	Weeks to provision
Snow Family	8 TB - 100 PB	Large one-time migrations, edge compute	Days to weeks
DataSync	Up to 10 Gbps per agent	Automated, scheduled transfers	Continuous
Transfer Family	Standard SFTP speeds	SFTP/FTPS/FTP to S3	Continuous

AWS Snow Family

Snowcone

8 TB HDD or 14 TB SSD. Rugged, portable. Edge computing + data transfer. IoT, tactical edge

Snowball Edge Storage Optimized

80 TB usable. S3-compatible. Local compute capabilities. Large data collection

Snowball Edge Compute Optimized

42 TB usable + GPU options. EC2 + Lambda at edge. ML inference, video analysis

Snowmobile

100 PB per truck. Exabyte-scale migrations. GPS tracking, 24/7 security

When to Use Snow

Rule of thumb: If transfer would take more than 1 week over available bandwidth, consider Snow devices. Calculate: Data Size / Bandwidth = Transfer Time. Include time for shipping (typically 1 week each way).

AWS DataSync

- **Purpose:** Automated data transfer between on-premises and AWS storage services
- **Destinations:** S3 (any class), EFS, FSx (Windows, Lustre, NetApp, OpenZFS)
- **Features:** Scheduling, bandwidth throttling, data integrity validation, encryption in transit
- **Agent:** Deploy on-premises VM or EC2 for cloud-to-cloud transfers
- **Performance:** Up to 10 Gbps per agent, can deploy multiple agents in parallel

Application Modernization

Modernizing applications for cloud-native architectures

Containerization Path

Service	Management Level	Best For
ECS on Fargate	Serverless containers	Simple deployments, no cluster management needed
ECS on EC2	Managed orchestration, you manage instances	Cost optimization with Spot, specific instance needs
EKS on Fargate	Serverless Kubernetes	Kubernetes ecosystem, no node management
EKS on EC2	Managed Kubernetes control plane	Full Kubernetes, GPU, Windows containers

Service	Management Level	Best For
App Runner	Fully managed from source/image	Simple web apps, APIs, rapid deployment
Lambda containers	Serverless functions from containers	Event-driven, existing container images up to 10GB

Serverless Modernization

API Gateway + Lambda

REST/HTTP/WebSocket APIs. Serverless backends, microservices. Pay per request

Step Functions

Orchestrate Lambda functions. Visual workflows, error handling, long-running processes

EventBridge

Event-driven architecture. Decouple services, SaaS integration, scheduled events

SQS + Lambda

Asynchronous processing. Decouple producers/consumers, handle spikes, DLQ for failures

Strangler Fig Pattern

Gradually replace legacy system components with modern services. Route traffic to new services incrementally using API Gateway or ALB. Reduces risk compared to big-bang migrations.

Microservices Patterns

- **API Gateway Pattern:** Single entry point, routing, authentication, rate limiting. Use API Gateway or ALB
- **Service Mesh:** App Mesh for service-to-service communication. Traffic management, observability, security
- **Event Sourcing:** Store state changes as events. EventBridge + Kinesis for event streaming
- **CQRS:** Separate read and write operations. DynamoDB Streams to sync read replicas

- **Saga Pattern:** Distributed transactions across services. Step Functions to orchestrate compensating actions
- **Circuit Breaker:** Prevent cascade failures. Implement in application code or use App Mesh

Specialized Migration Scenarios

VMware, mainframe, and specialized workload migrations

VMware Cloud on AWS

- **What:** VMware SDDC running natively on AWS bare-metal infrastructure
- **Use Cases:** Rapid data center evacuation, disaster recovery, cloud extension
- **vMotion:** Live migration of VMs between on-premises and AWS (no downtime)
- **Integration:** Native access to AWS services (S3, RDS, Lambda) via ENI
- **Licensing:** Bring existing VMware licenses or included in subscription

Mainframe Modernization

Replatform (Refactor)

AWS Blu Age: Automated refactoring of mainframe apps to Java. Preserves business logic

Rehost (Emulation)

Micro Focus: Run COBOL on AWS. Minimal code changes, faster migration

Assessment

AWS Mainframe Modernization service includes assessment tools. Analyze complexity, dependencies

Windows Workload Migration

Workload	AWS Service	Key Considerations
Active Directory	AWS Managed Microsoft AD	Trust relationships, Group Policy, seamless domain join

Workload	AWS Service	Key Considerations
SQL Server	RDS for SQL Server	License Included or BYOL, Multi-AZ, read replicas
SQL Server (advanced)	EC2 with SQL Server	Always On AG, FCIs, maximum control
.NET Applications	Elastic Beanstalk or ECS	.NET Core on Linux for cost savings
File Shares	FSx for Windows	SMB, DFS namespaces, shadow copies
SharePoint	EC2 or partner SaaS	Consider SharePoint Online as modernization

> **License Optimization**
>
> **BYOL:** Use existing licenses on Dedicated Hosts/Instances. Requires License Manager tracking
>
> **License Included:** Pay hourly, no upfront license cost. Often more cost-effective for variable workloads
>
> **Linux Migration:** Consider .NET Core on Linux containers to eliminate Windows licensing costs

CertStud

Free IT certification practice exams and study materials.

Resources

Roadmaps

Study Guides

Blog

Newsletter

Company

About

Contact

FAQ

Legal

Privacy Policy

Terms of Service

Our Products

CollegeDecider

BoostLogik

WanderingHermit

Brakto

Affiliate Disclosure: We may earn commissions from qualifying purchases through affiliate links.
Learn more