VCS One FinCloud

Technical Architecture Guide

Cloud Architecture, Migration Patterns, and Compliance

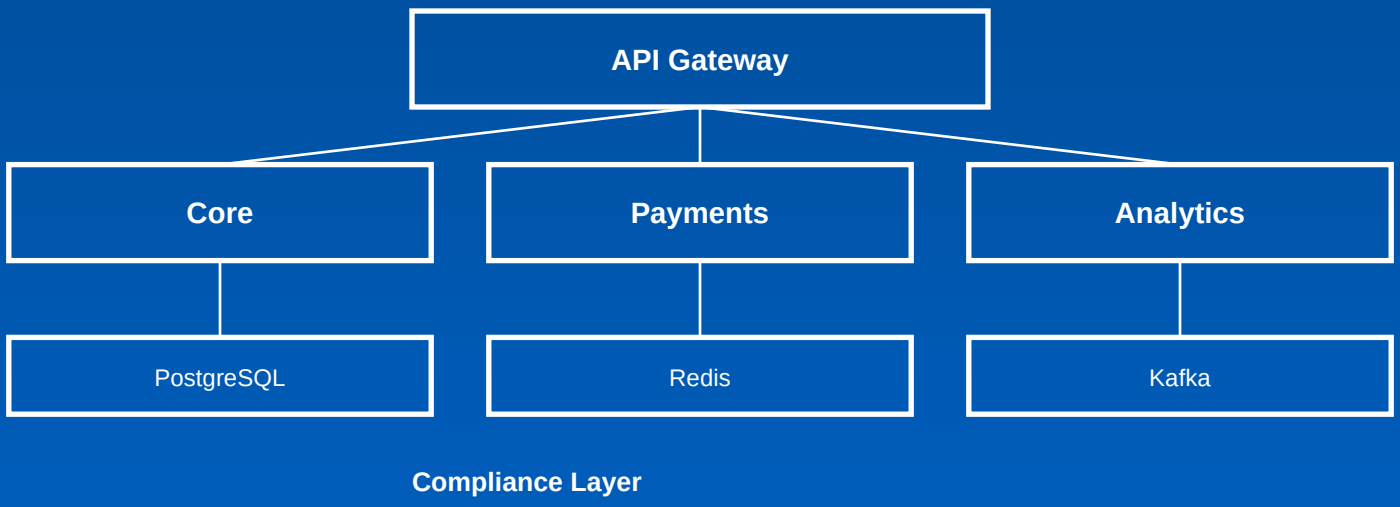

From Legacy to Cloud — With Confidence.

Migration Architecture

Refactor Orchestrator Architecture

The Refactor Orchestrator is the central command center for legacy-to-cloud migrations in financial services. It combines static code analysis, dependency mapping, pattern matching, and automated testing to transform monolithic systems into cloud-native microservices.

Core Components

Code Analyzer: Static analysis engine supporting COBOL, Java, C++, .NET, Python. Generates call graphs, data flow diagrams, and dependency maps

Pattern Library: Pre-configured migration patterns for common financial services modules: account management, transaction processing, compliance reporting, risk calculations

Decomposition Engine: Domain-driven design algorithms for identifying bounded contexts and microservice boundaries with configurable coupling/cohesion thresholds

Test Generator: Automated generation of unit tests, integration tests, and compliance validation tests based on legacy system behavior

Rollback Manager: State preservation and automated rollback capabilities for failed migration attempts

Monolith Decomposition Patterns

Financial services systems often contain embedded business logic that cannot simply be lifted-and-shifted. Our decomposition patterns identify and extract domain-specific functions while maintaining regulatory compliance.

Pattern Categories

Strangler Fig Pattern: Gradual replacement of legacy modules with cloud-native equivalents running in parallel until cutover

Database-Per-Service: Microservices maintain independent databases with event-driven synchronization for eventual consistency

Circuit Breaker: Resilience pattern for integrating with legacy systems prone to failures

Saga Pattern: Managing distributed transactions across microservices without two-phase commit (unsuitable for financial systems)

API Gateway: Single entry point for microservices with rate limiting, authentication, and request routing

Microservices Design Patterns

Cloud-native architecture patterns specifically adapted for financial services workloads with strict latency, consistency, and compliance requirements.

Financial Services Patterns

Event Sourcing: All transactions recorded as immutable events for audit compliance and replay capabilities

CQRS (Command Query Responsibility Segregation): Separate read/write models for optimizing query performance while maintaining transactional integrity

Transaction Coordinator: Distributed transaction management for ACID compliance across microservices

Audit Logger Service: Centralized logging of all financial transactions with immutable storage and cryptographic hashing

Rate Limiter: Per-customer and per-service rate limiting for fraud prevention and DDoS protection

Database Migration Strategies

Financial databases contain mission-critical data that must be migrated with zero data loss and minimal downtime. Our approach uses dual-write patterns and consistency verification.

Migration Approach

Dual-Write Pattern: Write to both legacy and new databases during migration window with reconciliation processes

Change Data Capture (CDC): Real-time replication of database changes to new systems using CDC tools (Debezium, AWS DMS)

Read Replica Promotion: Build read replicas of legacy databases, promote to primary, and redirect write traffic gradually

Data Validation: Automated consistency checks comparing row counts, checksums, and sample queries between old and new systems

Rollback Capability: Ability to revert to legacy database at any point during migration with zero data loss

Zero-Downtime Migration Approaches

Financial institutions cannot tolerate downtime. Zero-downtime migration strategies ensure continuous service availability throughout the cloud migration process.

Blue-Green Deployment

Maintain two identical production environments (blue = current, green = new). Deploy new system to green, run smoke tests, cutover DNS routing. Immediate rollback by reverting DNS if issues detected.

Canary Deployment

Gradual rollout: 5% traffic to new system, monitoring for errors. Increase to 25%, 50%, 100% over days/weeks. Automatic rollback triggers on error rate thresholds.

Shadow Mode

Run new system in parallel with legacy system, processing same requests but not serving responses. Validate output matching, performance, and compliance before production cutover.

Technology Stack

Containerization: Docker containers with Kubernetes orchestration on any cloud

Service Mesh: Istio for inter-service communication, load balancing, security policies

Message Queue: Apache Kafka for event streaming, RabbitMQ for async processing

Databases: PostgreSQL for OLTP, MongoDB/DynamoDB for document storage, Redis for caching

Monitoring: Prometheus + Grafana for metrics, ELK stack for logging, Jaeger for distributed tracing

Compliance Automation

Policy-as-Code Framework

Compliance requirements are codified as executable policies using Open Policy Agent (OPA) and HashiCorp Sentinel. Policies are version-controlled, testable, and automatically enforced at deployment time.

Policy Engine Architecture

- **Policy Repository:** Git-based storage for compliance policies with review workflows and approval gates
- **Policy Tests:** Unit tests for policies ensuring they correctly identify violations
- **Gatekeeper:** Kubernetes admission controller enforcing policies before pods are deployed
- **OPA Runtime:** Policy evaluation at API gateway, service mesh, and application layers
- **Audit Logging:** All policy decisions logged with context for compliance audits

Automated Compliance Checks

Continuous validation of infrastructure, code, and runtime configurations against regulatory requirements with automated remediation where possible.

Check Categories

- **Security Scanning:** Static (SAST) and dynamic (DAST) analysis in CI/CD with automated vulnerability assessment
- **Configuration Drift:** Continuous monitoring of infrastructure configs with alerting and auto-remediation
- **Access Control:** Validation of IAM policies, role assignments, and privilege escalation prevention
- **Encryption Verification:** Ensuring data encryption at rest and in transit for sensitive information
- **Patch Management:** Automated scanning for missing security patches with compliance dashboard visibility

PCI-DSS Implementation

Payment Card Industry Data Security Standard compliance for systems handling card holder data. Automated controls for all 12 PCI-DSS requirements.

- Network segmentation with firewall rules preventing cardholder data from being exposed to untrusted networks
- Strong cryptography: AES-256 encryption at rest, TLS 1.3+ in transit
- Access control: MFA required, least-privilege principles, regular access reviews
- Monitoring: 24/7 log monitoring, intrusion detection, alerting for suspicious activity
- Testing: Quarterly vulnerability scans, annual penetration testing, automated compliance validation

ISO 27001 Controls

Information Security Management System (ISMS) controls mapped to ISO 27001:2022 standard with continuous monitoring and improvement.

- **A.5 Information Security Policies:** Version-controlled policies with regular review and approval workflows
- **A.8 Asset Management:** Automated asset discovery, classification, and lifecycle management
- **A.9 Access Control:** Identity and access management with role-based access control (RBAC) and privilege escalation controls
- **A.12 Operations Security:** Secure configuration baselines, malware protection, backup procedures
- **A.14 System Acquisition:** Security requirements for software development with secure SDLC controls
- **A.18 Compliance:** Legal and regulatory compliance monitoring with automated reporting

GDPR Data Handling

European General Data Protection Regulation compliance for processing personal data of EU citizens.

- **Consent Management:** Granular consent capture and storage with immutable audit logs
- **Right to Access:** Automated data export APIs generating JSON/CSV responses
- **Right to Deletion:** Secure data purging workflows with retention policy checks
- **Data Minimization:** Automated identification and removal of unnecessary personal data
- **Data Portability:** Structured data export in industry-standard formats
- **Privacy by Design:** Default privacy settings, minimal data collection, encryption of sensitive fields

Audit Trail Generation

Comprehensive logging of all system activities for regulatory compliance and security investigations.

- **Immutable Logs:** Write-once append-only log storage with cryptographic hashing for tamper detection
- **Log Aggregation:** Centralized log collection using ELK stack with long-term retention (7 years for financial services)
- **Log Analytics:** Automated analysis for suspicious patterns, fraud detection, and compliance violations
- **Real-time Alerting:** Immediate notifications for high-risk events (failed logins, policy violations, data exfiltration attempts)
- **Forensic Analysis:** Querying capabilities for incident response and regulatory investigations

Fintech Integration APIs

KYC/AML Integration

Know Your Customer (KYC) and Anti-Money Laundering (AML) checks using industry-leading identity verification and screening services.

Identity Verification Providers

- **Onfido:** Document verification, facial recognition, liveness detection via REST API
- **Talon:** Biometric authentication, fingerprint scanning, voice verification
- **Jumio:** Real-time ID verification with 200+ document types, selfie authentication
- **Sumsub:** Multi-step KYC flows, global database screening, risk scoring

AML Screening Providers

- **Dow Jones:** Sanctions screening, PEP (Politically Exposed Person) checks, adverse media monitoring
- **World-Check:** Comprehensive sanctions and PEP database with configurable risk scoring
- **LexisNexis:** Enhanced due diligence, ongoing monitoring, case management

Payment Rails Integration

Seamless connectivity to major payment networks and clearing systems for domestic and cross-border transfers.

- **SWIFT:** MT and ISO 20022 message formats for international wire transfers
- **ACH:** Automated Clearing House for US domestic payments
- **SEPA:** Single Euro Payments Area for European transfers
- **Faster Payments:** Real-time payments in UK
- **RippleNet:** Blockchain-based global payments network
- **Stellar:** Open-source network for cross-border transfers

Core Ledger APIs

Standardized interfaces for core banking and accounting ledger operations with double-entry bookkeeping guarantees.

- **Transaction Recording:** Immutable transaction logs with debit/credit balance validation
- **Account Management:** Create, update, close accounts with hierarchy support (parent/child accounts)
- **Balance Queries:** Real-time balance inquiries with configurable consistency levels
- **Reconciliation:** Automated matching of transactions across systems with exception handling
- **Audit Reports:** Trial balance, general ledger, transaction history exports

Snowflake Analytics Integration

Data warehouse integration for advanced analytics, reporting, and machine learning on financial data.

- **Data Pipeline:** Automated ETL from operational databases to Snowflake using CDC or batch loads
- **Schema Evolution:** Handling schema changes in source systems with versioned table structures
- **Data Transformation:** SQL-based transformations with dbt for data modeling
- **Real-time Analytics:** Streaming data ingestion for near real-time reporting
- **Cost Optimization:** Automated cluster management, query optimization, and storage tiering

Multi-Cloud Scaffolding

Infrastructure-as-Code templates for deploying FinCloud on AWS, Azure, GCP, or hybrid environments.

- **Terraform Modules:** Reusable infrastructure modules for VPCs, subnets, Kubernetes clusters, databases
- **Ansible Playbooks:** Configuration management for application deployment and updates
- **Helm Charts:** Kubernetes application packaging for microservices deployment
- **CloudFormation/CDK:** AWS-specific infrastructure definitions (also available for Azure ARM, GCP Deployment Manager)

Performance & Cost Optimization

Right-Sizing Algorithms

Machine learning algorithms analyze workload patterns and recommend optimal resource allocations for financial services workloads.

- **Usage Pattern Analysis:** CPU, memory, I/O usage trends over time with seasonality detection
- **Recommendation Engine:** Instance type and size recommendations based on transaction volume, latency SLAs
- **Cost-Performance Tradeoffs:** Balancing compute costs against latency requirements for financial transactions
- **Spot Instance Optimization:** Identifying workloads suitable for spot/preemptible instances
- **Reserved Capacity Planning:** Predicting future resource needs for 1-3 year commitments

Cost Intelligence Dashboard

Real-time visibility into cloud spending with FS-specific attribution and compliance cost analysis.

- **Multi-Dimensional View:** Cost breakdown by service, environment, business unit, product, region
- **Compliance Overhead:** Separate tracking of encryption, logging, monitoring, disaster recovery costs
- **Budget Alerts:** Automated notifications when spending approaches budget thresholds
- **Forecasting:** Predictive models for future spending based on growth trends
- **Anomaly Detection:** Automatic identification of unusual spending patterns indicating potential issues

Latency Monitoring

Real-time performance monitoring with sub-millisecond latency tracking for financial transactions.

- **Application Performance Monitoring (APM):** Distributed tracing with Jaeger, service-level latency histograms
- **Synthetic Monitoring:** Automated transaction replay from external locations
- **Real User Monitoring:** Client-side latency measurement for web and mobile applications
- **P50/P95/P99 Metrics:** Percentile-based SLA monitoring with alerting
- **Regression Detection:** Automated alerts when latency degrades beyond baseline

Resource Optimization

Automated optimization of compute, storage, and network resources to reduce costs while maintaining performance.

- **Container Density Optimization:** Maximizing pods per node without overcommit
- **Storage Tiering:** Automatic migration to cheaper storage classes for cold data
- **Network Optimization:** CDN usage for static assets, compression for API responses
- **Cache Strategy:** Redis/Memcached deployment for frequently accessed data
- **Database Query Optimization:** Index recommendations, query plan analysis

Auto-Scaling Configuration

Intelligent scaling policies for financial services workloads with burst handling and cost controls.

- **Horizontal Pod Autoscaler:** Kubernetes HPA based on CPU, memory, custom metrics
- **Vertical Pod Autoscaler:** Automatic adjustment of resource requests/limits
- **Predictive Scaling:** Preemptive scaling based on historical patterns
- **Burst Capacity:** Overprovisioning during peak trading hours
- **Scale-Down Policies:** Conservative scale-down to prevent service disruption

Security Architecture

Encryption Standards

End-to-end encryption for data at rest and in transit using industry-standard algorithms and key management practices.

At Rest: AES-256-GCM encryption for databases, object storage using AWS KMS, Azure Key Vault, or GCP KMS

In Transit: TLS 1.3 with perfect forward secrecy for all network communications

Key Rotation: Automated key rotation every 90 days with zero-downtime re-encryption

Key Escrow: Secure key escrow for regulatory compliance and disaster recovery

End-to-End: Application-level encryption for sensitive fields (PII, payment data)

Network Security

Defense-in-depth network architecture with multiple security layers and monitoring.

Network Segmentation: VPCs, subnets with firewall rules preventing lateral movement

WAF: Web Application Firewall for SQL injection, XSS, DDoS protection

DDoS Mitigation: CloudFlare, AWS Shield for volumetric attacks

Private Links: Dedicated connections to cloud services without public internet

VPN/Zero Trust: Secure remote access with mandatory authentication and encryption

Access Control (IAM)

Identity and access management with role-based permissions and privilege escalation prevention.

Multi-Factor Authentication: Mandatory MFA for all administrative access

Least Privilege: Minimum necessary permissions with just-in-time access

Role-Based Access: RBAC with role hierarchies and permission inheritance

Single Sign-On: SAML/OIDC integration with corporate identity providers

Session Management: Automatic session timeout, concurrent session limits

Vulnerability Management

Continuous identification and remediation of security vulnerabilities.

Scanning: Weekly vulnerability scans of containers, images, dependencies

Patch Management: Automated patching for non-critical systems, staged rollout for production

Dependency Scanning: OWASP dependency-check for known CVEs in libraries

Container Security: Clair, Trivy scanning for container image vulnerabilities

SBOM: Software Bill of Materials for regulatory transparency

Incident Response

Prepared response procedures for security incidents with automated containment.

Detection: SIEM integration with real-time alerting for suspicious activities

Containment: Automated isolation of compromised resources

Investigation: Forensic logging and analysis tools

Remediation: Incident playbooks with step-by-step recovery procedures

Post-Incident: Root cause analysis and documentation for continuous improvement

Compliance Certifications

✓ **PCI-DSS Level 1 Certified**

✓ **ISO 27001:2022 Compliant**

✓ **SOC 2 Type II Audited**

✓ **GDPR & CCPA Compliant**

Technical Resources

- **Developer Portal:**
  https
- **API Documentation:**
  https
- **Migration Patterns:**
  https
- **Compliance Guides:**
  https

Value Creating Solutions Sdn Bhd

https://vcsmy.com | support@vcsmy.com