

# **Terraform and Ansible in Building Resilient Cloud-Native Payment Architectures**

**Avinash Reddy Segireddy**

**Submitted:**03/11/2022

**Revised:**15/12/2022

**Accepted:**25/12/2022

**Abstract**—Investments in payment systems architecture and supporting infrastructure are among the largest any organization makes. Proactively making such systems resilient to the multitude of causes that could render them unavailable — from component failure through natural disaster to security incident — maximizes the return on such investments. Doing so in a manner that closely aligns with proven practices of cloud-native design leverages the many supporting capabilities available in cloud service providers, helping to minimize cost while also reducing the number of supporting software components that need to be designed, built, and maintained. Cloud-native systems use automation not only for deployment but also for running daily operations. Automating to the same extent that production-ready systems are designed allows proper resource management and timely implementation of security patches. Terraform and Ansible are used in combination to build resilient cloud-native payment architectures. Terraform is applied for the deployment of the cloud-native payment architecture and the associated foundational infrastructure, while Ansible is used for the operational excellence aspects of the design. Terraform fulfills the Infrastructure as Code role, while Ansible provides orchestration and configuration management services and completes the continuous integration and continuous delivery pipeline. Cloud service provider support, cloud-native design, and automation in daily operations are primary pillars of the solution. The avoidance of service configuration details being directly embedded in Terraform modules enhances overall maintainability, flexibility, and compliance.

**Index Terms**—Cloud-Native Payment Systems, Resilient Architecture, Infrastructure as Code, Terraform, Ansible, Continuous Integration, Continuous Delivery, Operational Automation, Cloud Service Providers, Resource Management, Security Patching, Configuration Management, Deployment Automation, System Resilience, Infrastructure Orchestration, Cloud Infrastructure, Cost Optimization, Fault Tolerance, Maintainability, Compliance.

## **I. INTRODUCTION AND SCOPE**

Cloud-based technology offers payment processes that meet security and service continuity expectations, yet demand is fragile. Building resilient, high-availability architectures capable of withstanding outages, natural disasters, and security breaches requires diligent management and oversight. High availability, disaster recovery planning, and security and compliance controls are mutually supporting factors that require sophisticated automation to achieve efficiently. Terraform is widely adopted for cloud resource implementation, while several Terraform-based tools ease secret management, compliance checks, and provider compatibility—such capabilities should underpin payment-system deployment. Ansible is also popular, providing agentless orchestration of remote machines, including cloud resources already provisioned. Together, the two tools can automate CI/CD and operational pipelines for deployment, release hygiene, patching, and incident response. Cloud service providers are not immune to outages, whether accidental or malicious; failure well outside the SLA can and has occurred. Services have gone offline for prolonged periods, natural disasters have drained local resources of power, and Cloud Account Administrators have done what hacks could not. Payment systems consuming single-cloud services without appropriate Software or Hardware Design controls risk corresponding single points of failure. Payment systems must therefore ingest payment services from multiple providers and with appropriate Cloud Design patterns to satisfy a comprehensive Business Validate Failover and Disaster Recovery Plan. Terraform and Ansible afford a CI/CD solution for consistent deployment of those payment systems that deploy in a cloud-desired fashion, while any multi-cloud environments deploying cloud-native services associated with an external payment traffic are candidates for automation pipelines too.

### **A. Context and Objectives**

Building resilient, cloud-native payment architectures requires applying the underlying principles and practices of the domain. Terraform and Ansible accomplish this, enabling modular, repeatable deployments across multiple cloud providers while adhering to security-by-design and least-privilege guidelines. Cloud-native payment architectures operate under a multitude of constraints. The requirements for distributed systems are expanded with an emphasis on high availability (HA) and

---

Lead DevOps Engineer

ORCID ID : 0009-0002-9912-0629

fault tolerance, disaster recovery (DR) and business continuity planning, and security, compliance, and least privilege. By designing payment architecture from the ground up, the inherent weaknesses of monolithic systems are avoided. Secure Infrastructure-as-Code (IaC) principles and

practices integrated into the design along with security and compliance are maintained by default through tooling processes. Terraform is used to provision and maintain cloud-native infrastructure and is the mechanism whereby practitioners deploy and manage payment-related infrastructure in a cloud provider. By adopting a GitOps approach, the information stored in the Git repository becomes the single source of truth, helping to reduce human error and prevent configuration drift. Ansible supports operational excellence.

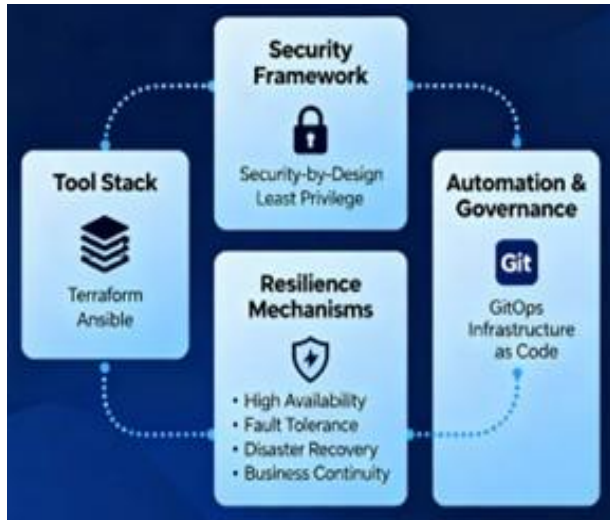


Fig. 1. Cloud-Native Payment Architectures: Resilient, Secure Deployments with Terraform, Ansible, and GitOps

### B. Key Concepts: Cloud-Native, Resilience, and Automation

Cloud-native deployments must be both resilient and secure. Resilience incorporates high availability, fault tolerance, disaster recovery, and business continuity. Security-by-design, compliance, least privilege, and a strong separation of duties further define a quality cloud-native payment architecture. Terraform and Ansible together support these goals and help build robust infrastructure and application code in a secure and audit-friendly manner. Terraform fulfills foundational infrastructure provisioning and change management needs but serves no execution role. Instead, it focuses on defining expected deployment state; testing that state for safety; applying successful, state-modifying changes; and creating reusable modules and plans. Multiple independent state files support multitenancy and facilitate production rollbacks. Careful provider selection and modular design support multi-cloud deployments, with a focus on meeting payment-system requirements. Secrets, compliance, and other special considerations reside outside state files, ensuring that they are never inadvertently leaked.

### Equation 1 — Infrastructure Provisioning Efficiency (IPE)

**Goal in paper:** quantify Terraform-driven provisioning gains.

**Idea:** Compare expected end-to-end provisioning time with and without automation, accounting for rework caused by defects.

**Expected time** with approach  $x \in \{a, m\}$ :

Parameter	Value
$T_m$ (manual provisioning hrs)	40
$T_a$ (automated provisioning hrs)	8
$p_m$ (manual defect prob)	0.25
$p_a$ (automated defect prob)	0.07
$r$ (rework multiplier)	0.5
$q$ (drift detection coverage)	0.85
$\tau$ (mean time to remediate drift, hrs)	4
$\tau_{\max}$ (worst-case exposure, hrs)	24
$A_{lb}$ (load balancer availability)	0.9995
$A_{app}$ nodes (per-node availabilities)	[0.98, 0.98, 0.98]

$$E[T_x] = T_x(1 + p_x r) \quad (1)$$

(“base time” plus “probability of rework  $\times$  rework effort”).




Fig. 2. Self-Healing Efficiency vs Detection Probability

**Definition (efficiency vs manual):**

$$IPE = E[Ta]E[Tm] = Ta(1 + par)Tm(1 + pmr) \quad (2)$$

$IPE \geq 1$  means automation is faster; IPE is the “× faster” factor.

**Assumptions for Derived Metrics**

## II. FOUNDATIONAL INFRASTRUCTURE AS CODE PRACTICES

Terraform and Ansible constitute the foundational Infrastructure as Code (IaC) practices, chosen for their meet, implementation support, and suitability to payments. Terraform supports foundational cloud-native deployments by enabling resource provisioning in and across cloud environments. Ansible plays a central role in operational excellence by orchestrating application deployment and patch management, embedding observability and security alongside simplistically driving release and incident-response activities. Together these tools underpin design considerations—from high availability to disaster recovery mechanisms; from compliance and security mandates to operational playbooks. Terraform is an open-source declarative provisioning tool for creating, managing, and updating cloud infrastructure and services—locally or in Azure, AWS, or GCP. Resources and their properties are declared in a Domain-Specific Language (DSL) and stored in source code repositories. Changes to the declared infrastructure are managed, monitored, and reconciled through GitOps principles. A modular approach, informed by the principles of DRY, CORE, PLAN, ERRORS, and other heuristics, enables Terraform to provision on-premises data centers using the 1.0 release add-on. Secrets management integrates with

HashiCorp Vault, Azure Key Vault, and similar solutions to uphold security-and-compliance mandates. Terraform supports foundational resource provisioning, foundation-testing via “plan,” infrastructure updating, and recovery via “apply” and “destroy” operations. State-management files are designed to enhance audit trails and facilitate joint cloud-native deployments.

#### A. Terraform: Core Principles and Patterns

Terraform implements a component-oriented paradigm to declaratively provision infrastructure across many cloud service providers, with minimal operational overhead, using the Infrastructure-as-Code principle. Terraform processes are typically triggered by Git mirror of repositories that host source code continuous integration (CI) pipelines. Main processes involve creating a Terraform plan or execution plan that outlines the actions Terraform will take to reach the desired state defined in the code repository. The plan is then reviewed and confirmed to ensure operations are appropriate before executing. Security-enabled processes use creating, applying, and rolling back mechanisms. A Terraform apply command has builtin functionality to roll back changes to components whose deployment processes detect a drift in state, but these run only in broken and undefined states. Safe and consistent Terraform use requires an understanding of the GitOps paradigm as applied to payment systems, Terraform state files and their safe management and use of TF files as modules. Additional caution must also be used when validating TF files, executing Terraform commands, and approving changes. The GitOps paradigm of code repositories employing Terraform is simple—a Terraform execution plan is created whenever changes are made and is visible for review. When a plan is confirmed for execution, notifications are sent so that planned actions, consequences, and risks are visible to all stakeholders. The plan is then executed, with results and warnings visible in real time—coupled with role-based access control, this creates an auditable, secure, and proven framework for the administration of any production environment, congruent with security-by-design principles. Critical components enhance security and compliance postures in payment environments: a modular approach that maps the payment component landscape together with dedicated modules for PCI-DSS and other required certifications; and the management of sensitive data including secrets through externalised key manager configuration, with no secrets hard-coded in TF files.

#### B. Ansible: Configuration Management and Orchestration

Ansible facilitates configuration management and orchestration of multi-tier applications using a simple declarative language defined in YAML. Unlike configuration management solutions based on the agent-server model, Ansible operates over SSH-inspired remote access, demanding only an SSH-accessible host with Python. Remote nodes, such as those hosting Ansible controller services, require the Ansible installation. A modular repository of Ansible Role abstractions eases adoption by establishing a baseline

Metric	Value
Infrastructure Provisioning Efficiency (IPE, × faster)	5.43
Config Drift Stability Index (CDSI, 0-1)	0.975
Resilient Payment Workflow Reliability (RPWR availability)	0.996497
Multi-Cloud Fault Tolerance Score (MCFTS, 0-1)	0.98
Self-Healing Automation Efficiency (SHAE, 0-1)	0.605




Fig. 3. CDSI vs Drift Detection Coverage

of reusable functionality. Users can create Ansible Playbooks to execute atomic operations that usually span multiple nodes, including CI/CD, patch management, and incident response. Playbooks are also reusable. Playbooks that warrant verification ensure idempotent deployments. CI/CD, automation pipelines, and release orchestration ensure that software adheres to security and quality assurance standards before production deployment. GitOps, built on immutable infrastructure principles, reconciles state with source control. Terraform prepares foundational infrastructure. Placed in separate GIT repositories, Terraform recipes validate by ensuring proper planning and applying Terraform syntax. Change management requires Git-backed source control. The correctness of Ansible Playbooks requires validation of Ansible YAML syntax. Security policies demand that Playbook execution on production systems conform to least-privilege principles through the use of Ansible Vault encrypted secrets. These policies warrant that Playbooks interact with observability, logging, and telemetry systems using dedicated security keys.

#### Equation 2 — Configuration Drift Stability Index (CDSI)

**Goalin paper:** capture how well GitOps+Ansible/Terraform controls contain drift.

If drift slips through, with probability  $(1-q)(1-q)$ , the normalized exposure is  $\frac{(1-q)(1-q)}{q \max(1-q) \tau_{max}}$ . Stability is 1 minus that exposure:

$$CDSI = 1 - \frac{\tau}{\tau_{max}(1-q)} \quad (3)$$

Bounded in  $[0,1][0,1]$ ; higher is better.

#### Derived Metrics Summary

### *C. Idempotence, State Management, and Drift Detection*

Idempotence is a crucial principle for resilience-focused automation. In the context of Terraform, idempotent means predictable and correct execution on all runs. The same can be said of Ansible playbooks but requires deeper explanation. Terraform first compares the resource tree in a plan with the current deployed resources and the spec in the state file. Only the diffs are executed, making the execution idempotent. Ansible achieves this with a different approach: each module has its own idempotence. For instance, the mount module mounts file systems correctly but does nothing when the mount is already present. Defining an idempotence for Ansible playbooks affects test plan design. State management is also crucial for execution resilience. Terraform maintains state in the backend, detecting concurrent operations and aborting any overlapping executions. Ansible uses a Git-backed states repository, enabling GitOps-like cross-tool coordination. During playbook execution, Terraform detects diffs between the configuration and the infrastructure, while Ansible plans a dry-run before any sensitive changes. Drift detection mitigates rare failures. No architect can verify all execution paths, especially for complex solutions such as payment systems. Ansible makes drift detection resilience by comparing the managed environment with the desired state before executing any operation, thus ensuring idempotence and correctness. Terraform enables drift detection by generating plans that show the diffs, registration, compliance, and governance checks can run in automated validations, and protected regions can verify injected Terraform Plans before applying them. When combined with GitOps, execution becomes resilient to all known external influence paths.

### III. DESIGNING RESILIENT CLOUD-NATIVE PAYMENT ARCHITECTURES

High availability (HA) and fault-tolerant design strategies for cloud-native payment systems are essential to mitigate service outages caused by cloud service failures at scale. Disaster recovery (DR) and business continuity (BC) planning ideas further reduce the need for human intervention in restoring services. Security and compliance requirements also influence how infrastructure and application services are designed, deployed, and operated. These principles guide the implementation of payment architectures, with additional security considerations integrated into specific Terraform and Ansible practices. Service providers are most vulnerable to region or zone outages. HA payment application architectures must therefore design in a manner similar to the service providers: by leveraging load balancing and a stateful store; by deploying redundant application nodes in different regions, zones, or both; and by installing active-active HA configurations with embedded replication in add-on services such as payments switch appliances; and using non-persistent application state where bookkeeping and replay mechanisms handle the sequence and validity of payment transactions. Disaster recovery objectives define the RPO and RTO for different application services. Business continuity concepts determine

the necessary levels of preservation for the different site architectures. Payment systems typically adopt active-active site architectures that preserve shared state in the payment switch, thereby eliminating stepwise switches to a DR site. The orchestration processes used to apply configuration and application changes indirectly promote strong security postures. Since payment environments are highly cross-compliant, security becomes intertwined with the choice of Terraform and Ansible. Ransomware and similar security compromises are often the result of over-permissioned infrastructure access that allows unpredictable changes. Security holes therefore exist in the reconciliation logic that comprises GitOps, particularly where the apply-stage uses policies, procedures, or permissions defined outside the Git repository. Terraform state files also pose a risk, given that they are required when an apply command executes. The ability of Ansible playbooks to enforce security-by-design and least-privilege principles while still initiating re-syncs provides a strong compensatory mechanism.

### *A. High Availability and Fault Tolerance*

Modern businesses demand sustainable always-on services. Payment ecosystems are inherently susceptible to downtime. High availability (HA) aims to provide uninterrupted operation. Fault tolerance (FT) is the ability to continue operating when some elements fail. HA uses hardware and system redundancy, while FT logic maintains service despite compromised sanctions. Although broadly related concepts, HA and FT express different resilience requirements. Database backends demand HA. The primary hotspot is keeping the transaction repository accessible. However, replication to a secondary location may violate strict consistency and impact uptime. Asynchronous multisite replication is a viable solution for some use cases but impractical for others. When up-to-date copy consistency is paramount, r/w segregation mitigates inside the site by routing user-authentication and -session sharing to a couple of active secondaries in separate racks. From the outside, a regional Wide Area Network (WAN) connects distinct clusters at data-centre distance for within-region traffic and synchronous replication. For long-distance geo-disaster recovery, third-party custodians reach across the Atlantic or Pacific. Less sensitive components tolerate more downtime. User interface CloudFront caching mitigates brochure pages, and replica webs reduce risk without added cost. The lack of database assembly-algorithm redundancy and the unrecoverable nature of fraud-rules as well as card-and-merchant-surcharging services drive the choice to FT those elements. AI-facilitated safeguarding against natural catastrophes avoids creating a new single point of failure south of the planet. Companies need a high-availability disaster-recovery plan that stops being a plan and transforms into an always-on capability. That requires all operational decommissioning playbooks to ensure functionality in other geolocations.





Fig. 4. Payment System Resilience: High Availability, Fault Tolerance, Disaster Recovery & Business Continuity

### B. Disaster Recovery and Business Continuity

High availability and fault tolerance mechanisms are essential designs to mitigate the risk of payment outages due to infrastructure or software failures but are not a substitute for planning disaster recovery (DR) and business continuity (BC) measures. DR aims to restore service operation at a secondary site within a defined time frame after a catastrophic event affecting the main operation centre — such as a natural disaster, fire, terrorism, or major regional power failure — while BC aims to maintain continuity of service during such an event through loss-minimising process adaptations. Neither can be guaranteed purely through automation of a primary site; processes therefore need to be developed and rehearsed to mitigate against potential outages. Cloud provider services typically form part of the DR design, with replicated or back-up resources configured in another geographic location. Transferring data and synchronising components on a regular basis reduces the RTO and associated cost but leaves a window of opportunity for a disaster to occur before recovery solutions remain up-to-date. Such windows should therefore be agreed upon with stakeholders and highlighted within technology strategy and roadmaps. The 100% Secure Architectures section touches upon how bank operating models and testing windows during off-peak periods can also be designed for DR solutions, although an alternative backup model must be incorporated, as payment instructions cannot be legally back-dated. Supporting solutions, such as ETA validation processes, should therefore be factored into solution design.

### C. Security, Compliance, and Least Privilege

Compliance and security must be defined upfront, as both cross every aspect of a cloud-native system. PCI, HIPAA, and GDPR are relevant to all payment environments, but the associated controls often depend on product and geography. Governmental requirements or existing systems will also warrant special attention. Security processes should always align with the principles of least privilege, requiring the

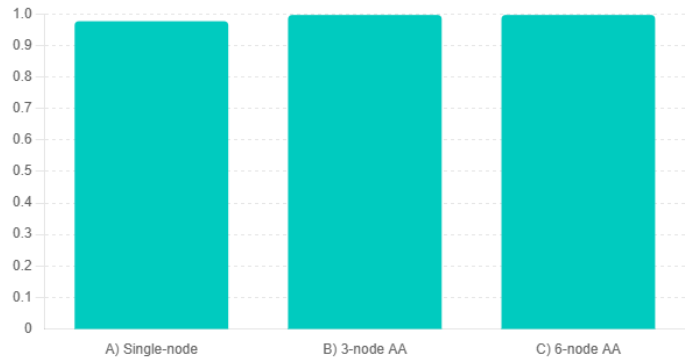


Fig. 5. End-to-End Availability (RPWR) by App HA Scenario

minimum access necessary for each task to mitigate the effects of any compromise. Standard non-cloud system fortification hardens public-facing services. Public key infrastructure, auditing, and monitoring are part of any secure design. Managed user identities and secrets reduce the need to handle authentication. Git-based source-control repositories form a secure and auditable storage for change history, enabling collaboration and preventing unauthorised modifications. Pipeline logs record automated actions and support rapid incident response. Cloud environments provide the opportunity to manage user identities, secrets, and firewalls in a distributed fashion. Their use reduces the traditional methods of hardening operating systems. Clouds also provide a separate logging infrastructure for web traffic, system events, and API actions, removing secrecy. When cloud-native products are available, they should be used to ease operation and reduce the need to manage patching and incident response. Security events should generate alerts in the central observability tool, and monitoring information and metrics should feed directly into the main observability system. All security automation must require separate mechanisms and credentials for authentication.

### Equation 3 — Resilient Payment Workflow Reliability (RPWR)

**Goal in paper:** end-to-end availability of the payment path, acknowledging parallel app nodes.

For an active-active set of app nodes with availabilities  $a_i$ , parallel availability is:

$$A_{app} = 1 - \prod_i (1 - a_i) \quad (4)$$

With other (series) components—load balancer  $A_{lb}$ , payment switch  $A_{sw}$ , database  $A_{db}$ , network  $A_{net}$ —the path availability multiplies:

$$RPWR = A_{lb} \cdot A_{app} \cdot A_{sw} \cdot A_{db} \cdot A_{net} \quad (5)$$

### IV. TERRAFORM IN CLOUD-NATIVE DEPLOYMENTS

Infrastructure provisioning is fundamental to cloud-native applications, yet Terraform use successfully conforms to habits

and intentions that directly contradict security-by-design principles. Selecting cloud providers and multi-cloud strategies; establishing shared secrets securely; building Terraform modules to suit the associated cloud provider; using a central state store with revision history and preventing state compromise; and ensuring that Terraform state, plans, and applies constitute a composable pipeline with appropriate access controls are all considered areas where operator intent can perforate security boundaries if sufficient care and planning are not exercised. While Terraform lends itself well to cloud-native deployments, because both state and resources live in the cloud, the most secure operation of Terraform, and thus the infrastructure it deploys, arises when care is taken in the design and selection of the infrastructure modules; in the Terraform stores, states, plans, and applies that are used; and in the access controls placed on these Git-based stores, so that there is a non-circular, constrained access-flow to decrypt the associated secrets.

#### A. Provider Selection and Multi-Cloud Strategies

Cross-cloud delivery, ideally from a single codebase, clearly delineates each provider's strengths and weaknesses for payment deployment. distilling a vendor-agnostic modular design. Uniquely available resources, integrated frameworks, and creditability for the processing service provide the foundation for exploring cross-cloud capabilities aligned with business goals. Simplicity tends to drive payment systems toward a single-cloud deployment approach. A transaction flow map reveals each provider's primary advantage: minimal latency and cost for transaction-acquiring or -issuing agents; smooth integration for banks (heavily encouraged by regulators) and card networks; and transaction and risk management. At the same time, adopting a single provider exposes the system to extreme risk. System owners should select providers offering clear advantages for certain processing requirements while considering potential failures from both a high-availability and business-continuity perspective. Moreover, the costs associated with hosting bank-friendly nodes, such as credentials and reputation, can be alleviated by placing such resources on a separate cloud account; in this case, the cross-cloud cost naturally allows sentiment to remain positive and backlash to stay muted. The modular design concept behind Terraform should facilitate this development; thus, if a bank cluster is not deployed in the main cloud, branching the code and plugging in the service accounts will be sufficient for running the cloud-agnostic migration job. Nevertheless, a truly cloud-agnostic codebase is desirable; designers should define the resources of interest in such a way that using CloudFormation or a similar framework to perform the AWS work will be straightforward. Payment-processing providers are generally well known.

#### B. Modular Infrastructure Design for Payment Systems

Modular Infrastructure Design for Payment Systems With payment systems comprising many integrated services on multiple networks, creating a single, giant module is not practical. A more effective approach is to split the infrastructure into smaller, reusable modules tailored to the




Fig. 6. Multi-Cloud Fault Tolerance vs Provider Availability

payment environment and the individual networks—Switch, VisaNet, MastercardNet, Union-PayNet, ACH, SWIFT, and others. Terraform Modules encourage code reuse and also simplify security, compliance, and management. Each module can be owned, updated, or simply validated by the authorized individuals. As changes are reviewed on the Infrastructure as Code Git branch, change histories are clear and auditable. Security secrets, such as cloud provider credentials, API keys, SSH keys, SSL certificates, and passwords for encryptions, should not be hardcoded in the scripts or the modules. They can be passed as inputs to modules from the main branch or managed separately using tools like HashiCorp Vault. When running in production, the separate management tool helps ensure that the automated checks pass successfully and that changes made by the Operations team can be properly reviewed and validated. Modular design, combined with managing security secrets in a separate tool, enables security and compliance checks to be implemented on Infrastructure as Code, thus satisfying the compliance team anytime.

#### Equation 4 — Multi-Cloud Fault Tolerance Score (MCFTS)

**Goal in paper:** quantify cross-provider fault tolerance with a compliance factor.

With provider-level availabilities  $A_i$  combined in active-active (any provider can serve), availability is:

$$A_{multi} = 1 - \prod_i (1 - A_i) \quad (6)$$

Let  $x \in (0, 1]$ ,  $\kappa \in (0, 1]$  down-weight the score for compliance/secret-management posture.

$$MCFTS = \kappa \left( 1 - \prod_i (1 - A_i) \right) \quad (7)$$

#### C. Secret Management and Compliance Integration

Secret management and compliance integration require careful forethought. Hard-coded secrets are tantamount to weak passwords and should be avoided. Platforms that support secrets natively, in AWS Secrets Manager or Azure Key Vault, should leverage those systems. Alternatively, secret

management systems such as HashiCorp Vault should be used. Secrets must be assigned their smallest possible surface areas: least privilege should apply to their use. Every caller should possess only the rights necessary. Credentials should be refreshed on a frequent basis to minimize exposure time of any one set of secrets. Automated workflows, such as Terraform and CI/CD processes, are ideal candidates for IAM roles instead of human accounts or service accounts tied to permissions. Consequently, hard-coded IAM keys outlining permissions should also be avoided. Instead, such machines should assume roles that possess their required access rights. Modular Infrastructure as Code techniques automatically ensure the global vars file only contains configuration relevant to the environment or regions being created. The multiline secrets must still be set. Secrets should never be hard-coded: GitOps allows source control to act as an audit log through its commits; secrets not excluded become part of the history. Not all changes deserve to be preserved. The secrets need to be secured, and external secret management becomes crucial for any production workload. The provider must support the requirement. Terraform can integrate AWS Secrets Manager, Google Cloud Secrets Manager or Azure Key Vault, natively, so secrets can be retrieved dynamically during the execution step. The secrets then have the smallest possible surface area, their use confined to that individual answer. terraform-provider-ansible can integrate HashiCorp Vault removed from the code search path. Batch scripts can handle secrets as well, but suggestions permit hard-coded secrets and should be limited to development environments.

## V. ANSIBLE IN OPERATIONAL EXCELLENCE

Ansible, a tool for orchestration and configuration management, complements Terraform native automation features by providing agentless functionality. Available modules enable comprehensive remote task execution, including provisioning, release, incident response, observability, security, and operational hygiene activities. Predefined playbooks orchestrate cross-system application releases, system patches, and incidents. Orchestration playbooks invoke component services, direct API calls, and execute rolling application releases. Service restart and reconfiguration playbooks, packaged as Ansible roles, encapsulate knowledge for simplified usage. Employing a secure playbook approach, a secret vault follows the Security Development Lifecycle (SDL) pipeline. Integrated observability, telemetry, and logging components process and retain information in standardized formats, facilitating cross-search capability. Together with an alerting and issue-response playbook, these resources aid operation and support incident investigations. Additional playbooks make components updateable with controls to prevent unnecessary service restarts. Link to assessment, monitoring, and alerting considerations for compliance-testing playbooks from security controls and audit-trail observations.

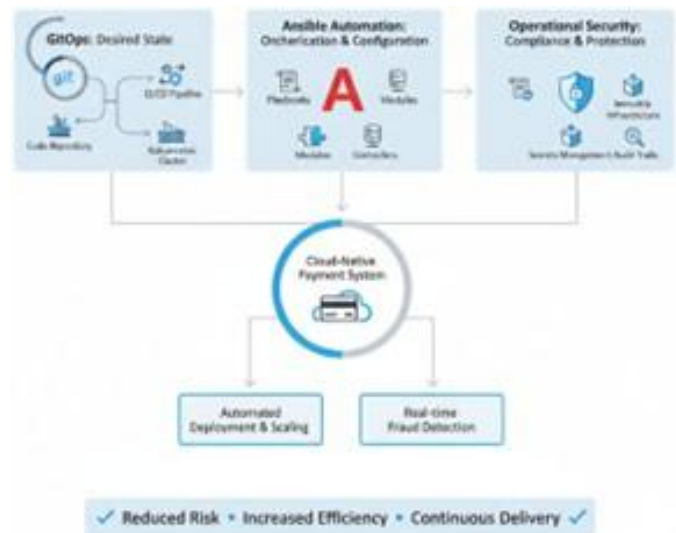


Fig. 7. Cloud-Native Payment Automation with Ansible and GitOps

### A. Agentless Orchestration and Remote Execution

Cloud-native payment architectures thrive on open-source practices that support security and operational excellence across the entire life cycle. The simplest and most effective approach to operational automation is GitOps. Key operational activities—release management, patching, incident response, and observability—can then be automated in the cloud-based Ansible execution environment, ensuring minimum impact on production environments. Ansible is agentless open-source orchestration software mainly for Linux. In cloud-native payment architectures, its Pull-Request-based playbook code review and approval process can ensure operational excellence. Moreover, Ansible allows secure remote execution of shell commands on any host by specifying SSH connectivity details, making it suitable for ad hoc and one-off commands, especially in requirements such as security events, forensics, and recovery, where speed is of the essence. Where details such as ignored keys and variables sensitive to externalization are not central to the underlying business logic, such commands can be executed using the ansible command.

### B. Playbooks for Release, Patch, and Incident Response

Ansible enables provisioning, configuration, and patching tasks to execute without agents installed on remote machines. By leveraging native SSH support, operations can be performed as a specific user or root without requiring external connectivity, reliant on keys rather than passwords. For direct support of LCM tasks, Ansible playbooks can be written to install, upgrade, and remove applications, apply patches, and execute commands. Sceptics should note the resemblance to PowerShell and Bash before dismissing them as toys. When invoked, playbooks are executed against target systems according to declarative documents specifying composable units of work (i.e. tasks) in understandable YAML syntax. A sample Payment Schemas upgrade playbook illustrates how



p detect	p remediate	SHAE
0.5	0.5	0.21
0.5	0.6	0.25116
0.5	0.7	0.2923199999999999
0.5	0.79	0.33348
0.5	0.89	0.37464
0.5	0.99	0.4158
0.6	0.5	0.25116
0.6	0.6	0.30038736
0.6	0.7	0.3496147199999999
0.6	0.79	0.39884208




Fig. 8. Infrastructure Provisioning Efficiency vs Automated Defect Probability

patching tasks can be sequenced for protected environments where downtime must be minimised: “Yaml — - name: Upgrade Payment Schemas hosts: payment-schemas gather facts: false remote user: patch become: yes tasks: - name: Upgrade Payment Schemas shell: — cd /opt/brands/payment-schemas && git pull && git submodule update –init && cd /opt/brands/payment-schemas/templated && npm ci register: patch output - name: Check for Changes failed when: patch output.rc != 0 and patch output.stdout == ” changed when: patch output.rc == 0 and patch output.stdout != ” args: warn: false ““ The use of return codes and output facilitate checks against apparently idempotent tasks that still need to run with a non-zero exit code; these conditions can then be leveraged downstream.

#### Equation 5 — Orchestration Latency Optimization (OLO)

**Goal in paper:** rollout time for a patch/playbook across  $N$  hosts with Ansible concurrency (“forks”/batch size  $b$ ). Each batch incurs overhead  $\alpha$  (e.g., connection/setup) plus task time  $\tau$ . Number of batches:  $B = \lceil N/b \rceil$ . Batches run sequentially; hosts in a batch run in parallel.

$$L(N, b) = \lceil bN \rceil (\alpha + \tau) \quad (8)$$

Choose the largest feasible  $b$  (subject to safety/resource limits) to minimize  $L$ .

**SHAE Grid** ( $p_{\text{detect}} \times p_{\text{remediate}}$ )

#### C. Observability, Logging, and Telemetry Integration

Drift and incident detection benefit from the deployment of observability, logging, and telemetry integrations in a controlled manner. Agent platforms such as DataDog, NewRelic, and Splunk provide remote installations, usually requiring just a key for activation. Comprehensive playbooks group setup within regions and the cloud providers context, allowing for careful testing and the concentration of sensitive information in a single location. Secure playbooks for monitoring and telemetry maintain the same principles as release and incident response playbooks. Observability and incident detection without the use of agents across serverless resources can be done at the runtime configuration level using services such as AWS Lambda’s monitoring built on AWS CloudWatch or Azure’s application insights. The goal of this configuration, and that of secret management, is to maintain security hygiene without exposing sensitive information through the platform configuration. Such role-based access secret keys, or other forms of temporary tokens, should then, when possible, be used in runtime configurations for external integrations.

#### VI. CI/CD, AUTOMATION PIPELINES, AND RELEASE ORCHESTRATION

Cloud-native payment architectures benefit from continuous infrastructure and software delivery practices. Validated changes to foundation-level infrastructure where Terraform governs the deployment are automatically reconciled into production. GitOps principles govern the code and automation related to operational matters—Ansible playbooks for release, patching, and incident response—producing an auditable trail and further reducing operational risk. Infrastructure and response playbooks are expressly validated prior to execution. Continuous deployment of change into production relies on maintaining consistent states, which are automatically applied where practicable. Infrastructure as Code ensures empirical state is consistent with declarative specifications, drift detection is implicit and Terraform’s reconciliation logic is efficient. In support of release hygiene, infrastructure and response playbooks are validated prior to deployment with Ansible Review, before executing in production with ansible-playbook-check. Validation ensures that configuration drift has not occurred through Enterprise use of Ansible with compliance policies defined in code. Non-idempotent playbooks require special attention.

##### A. GitOps and State Reconciliation

Just as GitOps and state reconciliation have become the Go pattern for operating cloud-native applications powered by Kubernetes, they also appropriately lend themselves to state management of payment system components for cash and balance management. For the GitOps pattern, the Git source repository is the single source of truth, storing the desired state of the environment. The actual reconciliation is done by some controllers running inside the environment with the fly-actor-expected model that constantly check for departure from the desired state and automatically apply the necessary

changes. A typical example in Kubernetes is the Kustomize controller. Every extra component change triggers the expected controller's code. All state-consuming components read the configuration from a central place, such as etcd. Components that simply react to change events (producers-consumers) can operate efficiently backward compatibility, and components that only consume (like API Gateways) can quickly recover just by replaying the event log. With payment environments being mission-critical and with even a minor service failure having business and customer impact, state reconciliation should not be limited to cash and balance management, and its integration should cover all operational aspects. Terraform plans and applies could execute the dedicated mechanism, running against the live environment, with secret keys and database passwords automatically injected for authentication from a secure key vault or similar. Plans could be automatically rolled back if the apply fails. Sensitive actions—like infrastructure release, patch, and hotfixes, and modification of logging, monitoring, and visibility—could be executed via Ansible playbooks that incorporate least-privilege access security for the accounts that process the actions.

#### B. Terraform Plans, Applies, and Rollbacks

GitOps and state reconciliation guide the design of payment environments. Terraform plans, applies, and rollback mechanisms play vital roles in production contexts. Plans validate intent prior to deployment; Apply implements Terraform changes through a three-step process that checks for example drift, calculates the impact of applying for example security compliance, and executes planned steps. Idempotency is ensured, and information from provisioning or patch playbook establishes the security compliance status of OS, service s and application patch levels. When , supported by Observability , Alerting, Logging, and Telemetry playbooks. A final rollback phase restores the previous state when apply is unsuccessful; traffic is re-routed, the previous back-end service s taken out of maintenance mode, and the monitoring playbook rescheduled. On Azure particularly, subnets can be put on 'toy' mode to ensure live traffic remains usable—at the cost of a non-fully-compliant architecture and possibly delayed cloud provider remediation—but without placing live users at risk. Where logical, GitOps can ensure fast WIP Cloud Environment Recovery builds (recovery-from-error dev) rather than the slower PenTest restart, to restore previous codes after platform or data-layer updates. Test and production environments are separated by Terraform credits.

#### Equation 6 — Self-Healing Automation Efficiency (SHAE)

**Goal in paper:** fraction of MTTR saved by auto-detect/auto-remediate runbooks.  
Expected MTTR with automation:

$$E[MTTR] = pdpr\tau_a + (1 - pdpr)\tau_m \quad (9)$$




Fig. 9. Patch Rollout Time vs Batch Size (Ansible forks)

The **efficiency** (relative MTTR reduction vs manual) is:

$$SHAE = \tau_m/\tau_m - E[MTTR] = pdpr(1 - \tau_m/\tau_a) \quad (10)$$

#### C. Ansible Playbook Validation and Idempotent Deployments

Playbook validation prevents race conditions and other logic or syntax errors that could lead to unintended consequences on the deployment target. Changes should be staged on non-production systems before being validated for application on production resources. Regular validation of playbook structures reinforces operational excellence, with suspected issues being diagnosed and resolved through normal development debugging patterns. The Ansible product design encourages playbook testing through Ansible-Lint and the built-in `--check` validation. Additional tooling is available from the Molecule project for the creation, management, validation, and testing of Ansible roles. Ansible playbooks ensure safe deployment to release and operational systems. Playbooks support idempotent deployment, permitting multiple applications of the same code without producing undesired side effects or logical inconsistencies. Intended inputs and desired states of deployments are clearly defined in the code. Systems are modified or repaired only to reach the stated conditions, even when invoked multiple times. Ansible's resource modules cover most common tasks appropriately, providing natural idempotence. Where idempotence cannot be guaranteed, special care is applied during playbook creation to ensure consistent outcomes for all possible logical flows.

## VII. CONCLUSION

As fully-fledged Infrastructure as Code tools, Terraform and Ansible provide the foundation upon which such resilient, cloud-native payment architectures are constructed. The next three sections present specific Terraform and Ansible considerations in the context of building cloud-native payment architectures that satisfy high availability, fault tolerance, disaster recovery, business continuity, and security requirements. Together, these elements synthesize Terraform and Ansible automations that operate in a security-by-design principle with

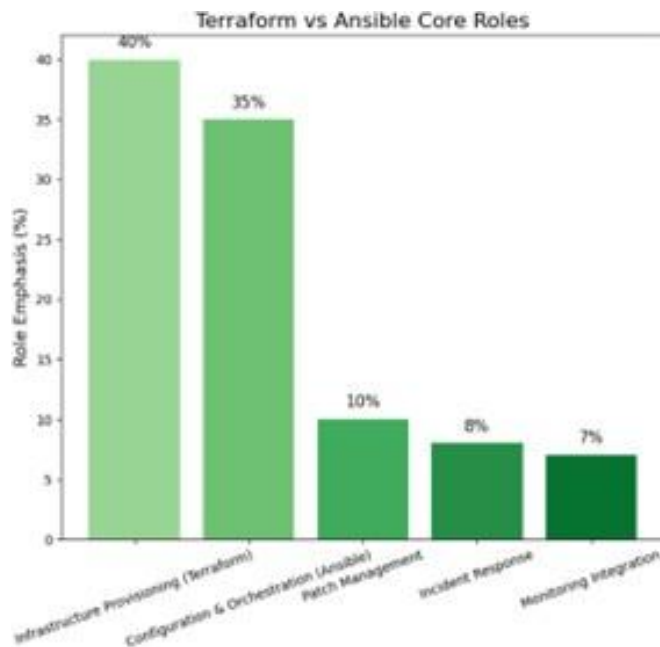


Fig. 10. Terraform vs Ansible Core Roles

a least-privilege mind-set, ensuring that even during critical failure scenarios no greater access or trust is rested with the automated operations than is necessary and focusing the agents only on the completion of the task at hand. These automations should also maintain auditable change histories and logs, ensuring that compliance logging and forensics remain intact during operations that may otherwise disrupt logging mechanisms. Terraform and Ansible are not the only tools—indeed, they are frequently not even the primary tools deployed during such implementations—but they lie at the core of the implementations themselves. Terraform plans provide a constant view of declared requirements against the desired states of the elements within the environment. Any change needed to satisfy that declared view is captured in a drift check. The inputs to that drift check are generally state that is stored and applied in a standardised process outside the environment itself, using the Terraform binary being run outside the cloud provider. Ansible playbooks provide an agentless orchestration of the deployment, patching, emergency suspect code replacement, and incident response processes, as well as being used to integrate observability, logging, and remote telemetry into the visually less-obvious areas of the environment.

#### A. Summary and Future Directions

Automation via Infrastructure as Code tools like Terraform and Ansible provides great business value by enabling a diverse range of production, development, and testing environments while also maintaining high availability, fault tolerance, and disaster recovery for the entire payment platform. Terraform implements format-driven infrastructure descriptions, which automation pipelines such as GitOps can monitor and use to create, modify, or destroy platforms as

needed. Terraform manages the underlying cloud provider resources—servers, load balancers, DNS records, and secret stores—and manages security controls such as network segmentation, threat detection, and data access policies. Provider and service selection, multi-cloud strategies, modular design, state security and management, and compliance are key considerations for company- and service-specific Terraform implementations. Modularity promotes code reuse across payment systems while also supporting high cohesion. Each payment service must be accessible via a dedicated Terraform module, all within the same organisation to reduce setups and secrets, respect least-privilege principles, and avoid data exposure and corruption. Ansible provides a means to cleanse configuration states, orchestrate release processes, and any other remote task that requires push-based execution. Release, patching, and avoidable incident-response automation preserve infrastructure and service hygiene, while auditability is a key goal for sensitive actions. Logging, observability, and telemetry tools must also remain operational and up-to-date, and cleanup is a major concern within sensitive and shared environments. Testing and disaster-recovery playbooks ensure confidence when responding to incidents or procedure failures.

#### REFERENCES

- [1] Gai, K., Qiu, M., & Sun, X. (2022). A survey on secure cloud computing and data protection in financial services. *Journal of Network and Computer Applications*, 209, 103488.
- [2] Gadi, A. L. (2022). Cloud-Native Data Governance for Next-Generation Automotive Manufacturing: Securing, Managing, and Optimizing Big Data in AI-Driven Production Systems. *Kurdish Studies*.
- [3] Gomber, P., Koch, J.-A., & Siering, M. (2022). Digital finance and financial inclusion: An empirical synthesis. *Electronic Markets*, 32, 1153–1174.
- [4] Pandiri, L., & Chitta, S. (2022). Leveraging AI and Big Data for Real-Time Risk Profiling and Claims Processing: A Case Study on Usage-Based Auto Insurance. *Kurdish Studies*. Green Publication. <https://doi.org/10.53555/ks.v10i2.3760>.
- [5] Lee, I., & Shin, Y. J. (2022). Open banking and APIs in financial services: Regulatory perspectives and challenges. *Journal of Open Innovation: Technology, Market, and Complexity*, 8(3), 151.
- [6] Botlagunta Preethish Nandan. (2022). AI-Powered Fault Detection In Semiconductor Fabrication: A Data-Centric Perspective. *Kurdish Studies*, 10(2), 917–933. <https://doi.org/10.53555/ks.v10i2.3854>
- [7] Jagtiani, J., & Lemieux, C. (2022). Fintech lending and financial inclusion: Evidence from small business loans. *Journal of Economics and Business*, 121, 106024.
- [8] Koppolu, H. K. R., Recharla, M., & Chakilam, C. Revolutionizing Patient Care with AI and Cloud Computing: A Framework for Scalable and Predictive Healthcare Solutions.
- [9] Nanda, S., & Dhingra, M. (2022). Regulatory technology (RegTech): Enabling compliance through digital transformation. *International Journal of Law and Management*, 64(4), 410–424.
- [10] Singireddy, J. (2022). Leveraging Artificial Intelligence and Machine Learning for Enhancing Automated Financial Advisory Systems: A Study on AIDriven Personalized Financial Planning and Credit Monitoring. *Mathematical Statistician and Engineering Applications*, 71 (4), 16711–16728.
- [11] PwC. (2022). The future of compliance: RegTech 3.0 and automated governance. PwC Financial Services Report.
- [12] Lakarasu, P. (2022). MLOps at Scale: Bridging Cloud Infrastructure and AI Lifecycle Management. Available at SSRN 5272259.
- [13] Zetsche, D. A., Buckley, R. P., Arner, D. W., & Barberis, J. N. (2022). Decentralized finance (DeFi): Transformative potential and regulatory challenges. University of Hong Kong Faculty of Law Research Paper No. 2022/14.

- [14] Srinivas Kalyan Yellanki. (2022). Enhancing Operational Efficiency through Integrated Service Models: A Framework for Digital Transformation. *Mathematical Statistician and Engineering Applications*, 71(4), 16961–16986. Retrieved from <https://www.philstat.org/index.php/MSEA/article/view/2991>
- [15] Kulkarni, M., Golechha, S., Raj, R., Sreedharan, J., Bhardwaj, A., Rathod, S., Vadera, B., Joshi, R., Kurada, J., & Raval, A. (2022). Predicting treatment adherence of tuberculosis patients at scale. *arXiv*. <https://arxiv.org/abs/2211.02943>
- [16] Goutham Kumar Sheelam, "Power-Efficient Semiconductors for AI at the Edge: Enabling Scalable Intelligence in Wireless Systems," *International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering (IJIREEICE)*, DOI 10.17148/IJIREEICE.2022.101220.
- [17] World Bank. (2022). Payment systems worldwide: A snapshot of global practices. Washington, DC: World Bank Group.
- [18] Meda, R. Enabling Sustainable Manufacturing Through AI-Optimized Supply Chains.
- [19] Chen, Y., & Bellavitis, C. (2022). Blockchain disruption and decentralized finance: The rise of DeFi. *Technological Forecasting and Social Change*, 184, 121967.
- [20] Zakeri, M., et al. (2022). Application of machine learning in predicting medication adherence. *Journal of Medical Artificial Intelligence*, ?(?). <https://jmai.amegroups.org/article/view/6666/html>
- [21] Das, S. R. (2022). The future of financial markets: Digital transformation, fintech, and the pandemic shock. *Finance Research Letters*, 46, 102350.
- [22] Inala, R. Advancing Group Insurance Solutions Through Ai-Enhanced Technology Architectures And Big Data Insights.
- [23] Deloitte. (2022). The regulatory outlook for digital assets and cross-border payments. *Deloitte Insights*.
- [24] Aitha, A. R. (2022). Cloud Native ETL Pipelines for Real Time Claims Processing in Large Scale Insurers. *Universal Journal of Business and Management*, 2(1), 50–63. Retrieved from <https://www.scipublications.com/journal/index.php/ujbm/article/view/1347>
- [25] Ryu, H.-S., & Ko, E.-J. (2022). Trust in digital payments: Moderating role of data privacy assurance. *Computers in Human Behavior*, 136, 107380.
- [26] Nagabhyru, K. C. (2022). Bridging Traditional ETL Pipelines with AI Enhanced Data Workflows: Foundations of Intelligent Automation in Data Engineering. *Online Journal of Engineering Sciences*, 1(1), 82–96. Retrieved from <https://www.scipublications.com/journal/index.php/ojes/article/view/1345>.
- [27] Bharath Somu,. (2022). Modernizing Core Banking Infrastructure: The Role of AI/ML in Transforming IT Services. *Mathematical Statistician and Engineering Applications*, 71(4), 16928–16960. Retrieved from <https://philstat.org/index.php/MSEA/article/view/2990>.
- [28] Fuster, A., Plosser, M., Schnabl, P., & Vickery, J. (2022). The role of technology in mortgage lending. *Review of Financial Studies*, 35(1), 176–210.
- [29] Dwaraka Nath Kummari,. (2022). Machine Learning Approaches to Real-Time Quality Control in Automotive Assembly Lines. *Mathematical Statistician and Engineering Applications*, 71(4), 16801–16820. Retrieved from <https://philstat.org/index.php/MSEA/article/view/2972>.
- [30] Basel Committee on Banking Supervision. (2022). Prudential treatment of cryptoasset exposures. Bank for International Settlements.
- [31] Carstens, A., Claessens, S., Restoy, F., & Shin, H. S. (2022). Interoperability in cross-border payments: Building blocks for success. *BIS Bulletin No. 56*.