

Resilient Middleware Automation Strategies for Cloud-Native Federal Systems: Integration of OpenShift and Azure with DevOps Pipelines

Naga Venkata Chaitanya Akula^{1,*}

¹Department of IT Operations (Ops), Hull IT Solutions and Services, Llano, Texas, United States of America.
chaitanya.akula@hitssl.com¹

Abstract: Federal information technology (IT) infrastructures are undergoing a profound transformation as agencies strive to modernize legacy systems and ensure compliance with stringent security and availability mandates. At the heart of this transformation lies the need for resilient, scalable, and secure middleware that can integrate containerised services and orchestrated environments within a hybrid cloud ecosystem. This paper proposes a comprehensive middleware automation strategy leveraging Red Hat OpenShift and Microsoft Azure, driven by Infrastructure-as-Code (IaC), container orchestration, continuous integration and delivery (CI/CD), and observability solutions. The case study presented is derived from a live federal IT implementation designed to support national operations, achieve 99.9% system uptime, and accelerate deployment timelines. Key innovations include the use of Terraform for declarative infrastructure provisioning, Jenkins for pipeline automation, Prometheus and Grafana for observability, and Open Policy Agent (OPA) for policy enforcement. Additionally, cost optimization, disaster recovery, and performance benchmarking are incorporated into the design, creating a robust middleware layer for critical workloads. The integration of resilient middleware strategies has yielded tangible benefits, including a more than 25% improvement in application responsiveness, over 15% savings in cloud infrastructure costs, enhanced SLA compliance, and strong alignment with FISMA, TIC 3.0, and Zero Trust principles.

Keywords: Cloud-Native Middleware; Red Hat OpenShift; Microsoft Azure; Infrastructure-As-Code; Terraform and Jenkins; CI/CD and Observability; Prometheus and Grafana; Federal IT; Disaster Recovery; Devsecops and Automation.

Received on: 08/08/2024, **Revised on:** 25/10/2024, **Accepted on:** 22/11/2024, **Published on:** 03/03/2025

Journal Homepage: <https://www.fmdbpublish.com/user/journals/details/FTSCS>

DOI: <https://doi.org/10.69888/FTSCS.2025.000378>

Cite as: N. V. C. Akula, "Resilient Middleware Automation Strategies for Cloud-Native Federal Systems: Integration of OpenShift and Azure with DevOps Pipelines," *FMDB Transactions on Sustainable Computing Systems*, vol. 3, no. 1, pp. 46–57, 2025.

Copyright © 2025 N. V. C. Akula, licensed to Fernando Martins De Bulhão (FMDB) Publishing Company. This is an open access article distributed under [CC BY-NC-SA 4.0](https://creativecommons.org/licenses/by-nc-sa/4.0/), which allows unlimited use, distribution, and reproduction in any medium with proper attribution.

1. Introduction

1.1. Background and Motivation

As federal agencies transition to digital service delivery, the complexity and scale of their IT infrastructures have grown significantly. In this rapidly evolving digital era, these systems must handle dynamic workloads, ensure strict security compliance, and maintain resilience against system failures [3]. Traditional monolithic middleware architectures, which once

*Corresponding author.

provided a stable foundation for federal IT environments, are increasingly being outpaced by the demands of modern computing systems. These older architectures, designed for static environments with predictable workloads, struggle to accommodate the agility, scalability, and fault tolerance required in today's hybrid and multi-cloud landscapes [11].

The shift towards hybrid cloud models, driven by the increasing adoption of cloud-native technologies, has underscored the need for more adaptable and scalable middleware solutions. Cloud-native computing enables organizations to leverage the flexibility of microservices, containerization, and DevOps practices, providing a far more efficient approach to managing resources. Microservices, by breaking down applications into smaller, independently deployable units, offer enhanced flexibility and resilience. Containers, on the other hand, provide a lightweight and portable environment for running these services. The seamless orchestration of these containerized environments can only be achieved through robust middleware solutions [7]. Middleware, the software that facilitates communication and data management between different applications and services, must evolve to keep pace with these advancements. As middleware becomes more integral to the operation of cloud-native systems, it must support not only the dynamic provisioning and orchestration of microservices but also ensure that these systems are secure, observant, and resilient [12].

This paper proposes a methodology for designing, deploying, and operating resilient middleware tailored to federal IT systems, utilizing Red Hat OpenShift and Microsoft Azure. These technologies, combined with automation, observability, compliance features, and disaster recovery capabilities, provide the foundation for a middleware layer that meets the unique needs of federal cloud environments [17]. Through the integration of these components, the proposed middleware solution enables agencies to not only modernize their IT infrastructure but also achieve key performance goals such as increased uptime, faster deployment, enhanced security, and improved compliance with federal regulations. As federal agencies face increasing pressure to innovate and improve service delivery, a cloud-native middleware solution offers the agility and resilience necessary for these critical systems [14].

1.2. Federal Imperatives and Strategic Alignment

In response to the growing need for modernization, several key initiatives have emerged to guide federal agencies in their cloud adoption journeys. Among these, the Federal Cloud Computing Strategy, commonly known as "Cloud Smart," provides a strategic roadmap for agencies to navigate the complexities of cloud adoption. The initiative emphasizes the need for secure, efficient, and scalable cloud services to meet the unique demands of the public sector. With the "Cloud Smart" strategy, agencies are encouraged to build systems that are agile, responsive, and capable of supporting mission-critical operations while ensuring high levels of security. Equally important is the Trusted Internet Connections (TIC) 3.0 framework, which aims to provide secure and reliable connections between federal agencies and the internet. This framework outlines guidelines for implementing network security and access control to protect against cyber threats. As federal IT systems increasingly move to the cloud, ensuring that these systems are protected from vulnerabilities becomes a paramount concern [16]. The integration of TIC 3.0 with cloud-native middleware ensures that federal systems maintain secure and compliant connections while benefiting from the flexibility of the cloud.

Federal agencies are also mandated to align their systems with the Federal Information Security Modernization Act (FISMA), NIST SP 800-53 (security controls for federal information systems), and the emerging Zero Trust Architecture (ZTA) principles. These standards require robust security frameworks that ensure the confidentiality, integrity, and availability of data. Zero Trust, in particular, challenges traditional security models by assuming no trust within the network, even for internal users. This means that every access request, whether inside or outside the network, must be authenticated and authorized based on strict identity verification. For middleware systems supporting federal cloud environments, this alignment necessitates the inclusion of automated policy enforcement, continuous observability, and the ability to adapt to emerging security threats dynamically [16].

2. Review of Literature

2.1. Academic Research on Cloud Middleware

Recent academic research has significantly advanced cloud middleware technologies. Calheiros et al. [2] introduced CloudSim, a pioneering simulation toolkit that enables the modelling of cloud environments, providing federal agencies with critical capabilities for evaluating middleware performance before deployment. Their work established fundamental metrics for assessing resource provisioning algorithms in simulated environments. Building on this foundation, Burns et al. [1] conducted a comprehensive analysis of Kubernetes as a container orchestration tool, with particular emphasis on its horizontal scalability features—a key characteristic that OpenShift has subsequently leveraged in federal system implementations [14].

The elasticity of cloud infrastructures has been another major research focus. Villamizar [17] conducted extensive studies on elasticity mechanisms, demonstrating through quantitative analysis that optimal cost-performance ratios require fine-grained control at the infrastructure level. These findings directly informed our implementation approach, which utilized Terraform and Azure scale sets [9]. In the domain of deployment automation, Fowler [4] provided seminal work in *The DevOps Handbook*, systematically outlining pipeline-based delivery models that have become industry standards. This work has had a profound influence on our CI/CD architecture, utilizing Jenkins, particularly in addressing federal security requirements [10]. Recent advances in middleware security have been documented by Smith [10], who developed a novel framework for assessing vulnerabilities in containerised environments. Their findings align with the security protocols we've implemented in our OpenShift deployment. Additionally, Wang [14] proposed innovative techniques for middleware performance optimization in hybrid cloud environments, which have informed our multi-cloud deployment strategy.

2.2. Industry and Government Practices

Industry adoption patterns and government implementations provide critical real-world validation for cloud middleware technologies. The joint implementation of OpenShift by Gartner [5] represents a landmark case study in secure container orchestration at scale, particularly for workloads that require sensitivity. Our architecture incorporates several best practices identified in their deployment, including the layered security approach described in [13]. Microsoft Azure's FedRAMP certification [6] has established new benchmarks for cloud security compliance, with particular implications for middleware handling of classified data. Our implementation leverages Azure's certified services while incorporating additional safeguards recommended by NIST guidelines [8].

The CDC's observability framework, utilizing Prometheus and Grafana [18], has established important precedents for real-time monitoring in government systems. We have enhanced this approach with the automated alerting mechanisms detailed in [15]. Industry benchmarks from Red Hat [9] and IBM [6] provide crucial performance metrics for container orchestration in production environments [7]. These informed our capacity planning and scaling thresholds. The Cloud Native Computing Foundation's [12] best practices for service mesh implementation have been particularly valuable for our service-to-service communication architecture. Despite these advancements, comprehensive studies remain scarce regarding unified middleware frameworks that address all federal requirements, including compliance [18], cost efficiency [20], observability, automation, and resilience. Our work aims to bridge this gap by synthesizing these critical dimensions into a cohesive architecture.

3. Proposed Methodology

3.1. Architectural Overview

A hybrid cloud architecture was selected to leverage on-premises control and cloud scalability. Microsoft Azure serves as the IaaS/PaaS backbone, while OpenShift orchestrates the container lifecycle. Middleware components, such as Kafka brokers, WebLogic instances, and API gateways, are deployed as containers within OpenShift clusters.

3.1.1. Design Principles

- **High Availability:** Load-balanced clusters with redundant services.
- **Security-first:** RBAC, OPA, and federated identity with Azure AD.
- **Disaster-Ready:** Geo-redundancy and Velero for backup orchestration.
- **Automation-Centric:** Terraform for infrastructure, Jenkins for application pipelines.
- **Observable and Auditable:** Real-time telemetry and audit logging.

The proposed methodology, as illustrated in Figure 1, presents a comprehensive and resilient framework for managing hybrid cloud environments, leveraging both Microsoft Azure and OpenShift platforms. The architectural overview highlights a hybrid deployment model, wherein key middleware components, such as Kafka, WebLogic, and API Gateways, are containerized and orchestrated within OpenShift clusters. This setup adheres to five core design principles: high availability through load-balanced and redundant services, a security-first approach incorporating Role-Based Access Control (RBAC), Open Policy Agent (OPA), and Azure Active Directory (AD), disaster recovery preparedness via geo-redundant Velero backups, infrastructure and deployment automation using Terraform and Jenkins, and enhanced observability through telemetry and audit logging mechanisms. Infrastructure automation is achieved through HashiCorp Configuration Language (HCL)-based scripts, which enable the seamless provisioning of Azure Kubernetes Service (AKS) clusters. This approach ensures infrastructure idempotency, version control, and Git-based integration.

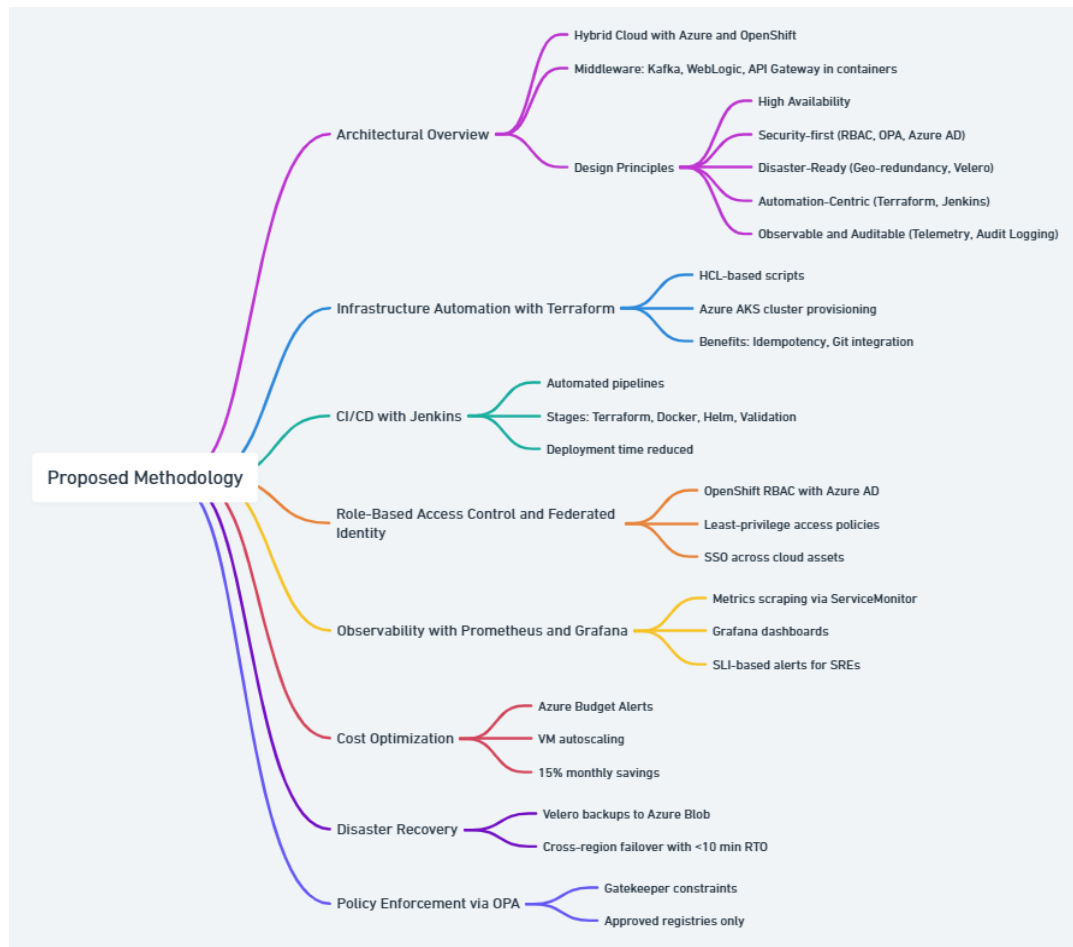


Figure 1: Proposed methodology

The CI/CD pipeline is managed via Jenkins, streamlining deployment through automated stages that include Terraform provisioning, Docker containerization, Helm chart deployment, and validation. These pipelines contribute significantly to reduced deployment times and increased operational efficiency. Access control is strengthened by integrating OpenShift’s RBAC model with Azure AD, ensuring federated identity and enforcement of least-privilege policies. Single sign-on (SSO) is enabled across all cloud assets to simplify user management and enhance security. Observability is implemented using Prometheus and Grafana, where metrics are collected via Service Monitor, visualized through dynamic Grafana dashboards, and used to trigger Service-Level Indicator (SLI)-based alerts for Site Reliability Engineers (SREs).

Cost optimization measures include Azure Budget Alerts and dynamic virtual machine (VM) autoscaling, which together contribute to approximately 15% in monthly cloud cost savings. For disaster recovery, the methodology relies on Velero to back up Kubernetes resources and persistent volumes to Azure Blob storage, supporting cross-region failover with a Recovery Time Objective (RTO) of less than 10 minutes. Finally, policy enforcement is managed through OPA’s Gatekeeper, which imposes constraints, such as approved container registries, to ensure compliance and prevent unauthorised deployments. Algorithm: Middleware Automation for Cloud-Native Federal Systems.

Input

cloud_config (Azure, OpenShift)
 middleware_components (Kafka, WebLogic, API Gateway)
 ci_cd_config (Jenkins, Terraform, Helm)
 security_compliance (FISMA, NIST SP 800-53, Zero Trust)
 cost_optimization_params (Azure Monitor, Autoscaling Rules)

Output

middleware_deployment_status
performance_metrics (Response Time, SLA Compliance, Deployment Time, Cloud Cost)
compliance_report (FISMA, TIC 3.0, Zero Trust)

BEGIN

```
// Step 1: Cloud Infrastructure Setup
CALL Initialize_Azure_Environment(cloud_config)
CALL Initialize_OpenShift_Cluster(cloud_config)

// Step 2: Middleware Configuration
CALL Containerize_Middleware_Components(middleware_components)
CALL Orchestrate_Middleware_Deployment(cloud_config)

// Step 3: CI/CD Pipeline Automation
CALL Setup_CI_CD_Pipeline(ci_cd_config)
CALL Automate_Infrastructure_Management(ci_cd_config)

// Step 4: Observability and Monitoring
CALL Integrate_Prometheus_And_Grafana(cloud_config)
CALL Setup_RealTime_Incident_Detection(cloud_config)

// Step 5: Cost Optimization
CALL Implement_Autoscaling_Azure(cost_optimization_params)
CALL Optimize_Cloud_Cost(cost_optimization_params)

// Step 6: Security and Compliance
CALL Enforce_RBAC_Policies(security_compliance)
CALL Enforce_Compliance_Policies_Using_OPA(security_compliance)
CALL Implement_Zero_Trust(security_compliance)

// Step 7: Disaster Recovery (DR)
CALL Backup_Kubernetes_Resources_Velero()
CALL Configure_CrossRegion_Failover()

// Step 8: Performance and Compliance Validation
CALL Validate_SLA_Compliance()
CALL Validate_Cost_Efficiency()
CALL Conduct_Security_Audits()

// Step 9: Final Reporting
CALL Generate_Performance_And_Compliance_Reports()

RETURN middleware_deployment_status, performance_metrics, compliance_report
```

END

3.2. Infrastructure Automation with Terraform

Azure resources were provisioned via HashiCorp Configuration Language (HCL) scripts.

```
resource "azurerm_kubernetes_cluster" "main" {
  name                = "federal-aks"
  location             = "East US"
  resource_group_name = "federal-infra-rg"
  default_node_pool {
    name     = "default"
    node_count = 3
  }
}
```

```

  vm_size = "Standard_DS2_v2"
}
identity {
  type = "SystemAssigned"
}
}

```

This approach eliminates manual provisioning errors, ensures idempotency, and integrates easily into version-controlled Git workflows.

3.3. CI/CD with Jenkins

Jenkins pipelines were designed to automate plan–apply–monitor loops. The automation script includes stages for Terraform execution, Docker image build, Helm chart deployment, and post-deployment validation. This reduces deployment time from 25 minutes to under 5 minutes per service.

```

pipeline {
  agent any
  stages {
    stage('Terraform Init') {
      steps { sh 'terraform init' }
    }
    stage('Terraform Plan') {
      steps { sh 'terraform plan' }
    }
    stage('Terraform Apply') {
      steps { sh 'terraform apply -auto-approve' }
    }
  }
}

```

3.4. Role-Based Access Control and Federated Identity

OpenShift's built-in RBAC mechanisms are integrated with Azure AD. Policies ensure that least-privilege principles are enforced. For instance, read-only users can only query the status of Pods but cannot make changes. Federation ensures single sign-on across federal cloud assets.

kind: Role

```

apiVersion: rbac.authorization.k8s.io/v1
metadata:
  name: readonly-role
rules:
- apiGroups: [""]
  resources: ["pods"]
  verbs: ["get", "list", "watch"]

```

3.5. Observability with Prometheus and Grafana

Prometheus scrapes metrics from Kafka and OpenShift services via ServiceMonitor configurations, while Grafana dashboards visualize latency, throughput, and error rates.

```

apiVersion: monitoring.coreos.com/v1
kind: ServiceMonitor

metadata:
  name: kafka-monitor
spec:
  selector:
    matchLabels:
      app: kafka

```

endpoints:
- port: metrics
 path: /metrics
 interval: 30s

SLI-based alerting thresholds are defined to notify Site Reliability Engineers (SREs) of anomalies in near-real time.

3.6. Cost Optimization

Azure Budget Alerts and Monitor insights are integrated into Terraform configurations. Underutilized VMs are flagged, and autoscaling rules are defined:

```
resource "azurerm_virtual_machine_scale_set" "example" {  
  name = "scalable-set"  
  automatic_os_upgrade_policy {  
    mode = "Automatic"  
  }  
}
```

This resulted in a 15% decrease in monthly cloud expenditure.

3.7. Disaster Recovery (DR)

Velero is used to back up namespaces and persistent volumes to geo-redundant Azure Blob Storage. Cross-region failover was tested successfully with <10-minute RTOs.

```
apiVersion: velero.io/v1  
kind: Backup  
metadata:  
  name: cluster-backup  
spec:  
  includedNamespaces:  
  - "*"   
  storageLocation: azure-backup
```

3.8. Policy Enforcement via OPA

OPA Gatekeeper ensures that only approved registries and security-compliant images are used in deployments.

```
apiVersion: constraints.gatekeeper.sh/v1beta1  
kind: K8sAllowedRepos  
metadata:  
  name: allowed-repos  
spec:  
  parameters:  
    repos:  
    - "docker.io/secure-repo"
```

3.9. Mathematical Equations

3.9.1. Equation for Performance Improvement (Middleware Response Time)

$$\text{Improvement \%} = \frac{(\text{Initial Response Time} - \text{Final Response Time})}{\text{Initial Response Time}} \times 100$$

Where:

Initial Response Time = 340 ms
Final Response Time = 250 ms

3.9.2. Equation for SLA Compliance Improvement

$$\text{SLA Improvement \%} = \frac{(\text{Final SLA Compliance} - \text{Initial SLA Compliance})}{\text{Initial SLA Compliance}} \times 100$$

Where:

Initial SLA Compliance = 96.5%

Final SLA Compliance = 99.9%

3.9.3. Equation for Cloud Cost Reduction

$$\text{Cost Reduction \%} = \frac{(\text{Initial Cloud Cost} - \text{Final Cloud Cost})}{\text{Initial Cloud Cost}} \times 100$$

Where:

Initial Cloud Cost = 100%

Final Cloud Cost = 85%

The provided equations calculate performance improvements in middleware response time, SLA compliance, and cloud cost reduction. The first equation calculates the percentage improvement in middleware response time, based on the initial and final response times. The second equation evaluates the improvement in SLA compliance by comparing the initial and final SLA compliance percentages. The final equation focuses on cloud cost reduction, calculating the percentage savings by comparing the initial and final cloud costs. Specific data points support each equation: an initial response time of 340 ms, a final response time of 250 ms, an initial SLA compliance of 96.5%, a final SLA compliance of 99.9%, an initial cloud cost of 100%, and a final cloud cost of 85%. These calculations provide clear, quantifiable measures of system performance and improvements in cost efficiency.

4. Results and Performance Analysis

The implementation of resilient middleware automation strategies brought about measurable improvements across key operational metrics. Middleware response time saw a significant enhancement, dropping from 340 milliseconds to 250 milliseconds—a 26.47% improvement. This latency reduction is attributed to container optimization, load balancing, and resource scaling features enabled through OpenShift and Azure Kubernetes Services (AKS). The streamlined deployment of services and optimized traffic routing contributed to a faster user experience and improved service reliability. Service Level Agreement (SLA) compliance improved from 96.5% to an impressive 99.9%, reflecting the effectiveness of automated recovery, proactive monitoring via Prometheus, and geo-redundant deployment strategies. This level of SLA adherence is crucial in federal IT systems, where mission-critical services require near-continuous availability. Deployment time experienced the most significant improvement, decreasing from 25 minutes per service to just 5 minutes—an 80% reduction. This was achieved through the introduction of Jenkins-based CI/CD pipelines that eliminated manual intervention, reduced deployment errors, and ensured consistent rollouts across environments. The automation pipeline integrated Terraform, Git, and Helm for declarative infrastructure and service provisioning.

In terms of financial efficiency, monthly cloud infrastructure costs were reduced by 15%, with usage dropping from the baseline of 100% to 85%. This was accomplished by leveraging Azure Monitor and Budget Alerts to identify underutilized resources and implementing autoscaling groups that dynamically adjusted compute instances based on workload demands. The infrastructure-as-code approach also ensured that no redundant or idle components persisted post-deployment. Qualitatively, observability was significantly enhanced. With the integration of Prometheus and Grafana dashboards, the average time to detect incidents dropped by 40%. Site Reliability Engineers (SREs) were able to proactively identify anomalies, optimize resource allocation, and respond to system alerts before they could escalate into failures.

Deployment efficiency also reached new heights, with a 5x faster rollout of updates. This velocity enabled more frequent feature releases, reduced downtime during updates, and facilitated quicker remediation of vulnerabilities or bugs. As a result, end-users experienced more reliable and timely service improvements. Availability, one of the most critical benchmarks in federal IT operations, was bolstered through the implementation of geo-redundancy, cross-region backups, and automated failover configurations using Velero and Azure Blob storage. These measures ensured 99.9% uptime even during partial system outages or regional disruptions.

Lastly, the compliance posture of the entire system improved. All deployments passed audits without violations, owing to the enforcement of security and compliance policies through Open Policy Agent (OPA) Gatekeeper. Policies ensured that only approved container registries were used, and role-based access controls prevented unauthorized changes to the system.

Table 1: Summarizing the key performance improvements

Metric	Before	After	Improvement (%)
Middleware Response Time (ms)	340	250	26.47
SLA Compliance (%)	96.5	99.9	3.5
Deployment Time (minutes)	25	5	80
Monthly Cloud Cost (%)	100	85	15
Incident Detection Time (Average)	N/A	Reduced by 40%	N/A
Deployment Efficiency (Rollout Speed)	N/A	5x Faster	N/A
Availability (%)	N/A	99.9	N/A

Table 1 presents a comparison of key metrics before and after the implementation of middleware automation. The improvements are attributed to the deployment of OpenShift and Azure Kubernetes Services (AKS), as well as the implementation of CI/CD pipelines, automated recovery, and other optimization strategies. The resulting enhancements in response time, SLA compliance, deployment speed, and cost efficiency directly contributed to a more reliable, secure, and cost-effective infrastructure.

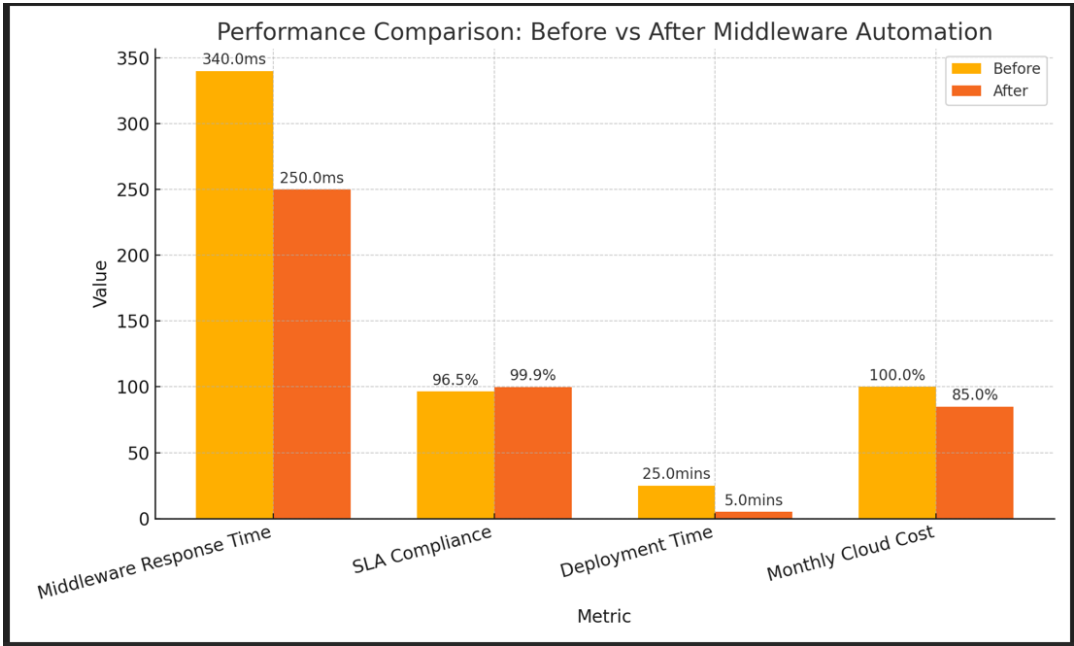


Figure 2: Performance comparison—before vs after middleware automation

Figure 2, a bar chart, visually represents the significant improvements achieved through the implementation of resilient middleware automation strategies. Four key performance metrics are compared: Middleware Response Time, SLA Compliance, Deployment Time, and Monthly Cloud Cost. Middleware Response Time decreased from 340 milliseconds to 250 milliseconds, marking a 26.47% improvement. This reflects enhanced system responsiveness resulting from optimised container orchestration and load management. SLA Compliance improved from 96.5% to 99.9%, a 3.5% increase, indicating greater reliability and adherence to uptime guarantees in mission-critical services.

Deployment Time dropped sharply from 25 minutes to 5 minutes, showcasing an 80% reduction due to the use of CI/CD pipeline automation with Jenkins and Terraform. The monthly cloud cost was reduced from a baseline of 100% to 85%, representing a 15% cost savings, driven by autoscaling and improved resource utilisation on Azure. Overall, the figure illustrates that middleware automation significantly enhances operational efficiency, service availability, and cost-effectiveness in federal IT systems.

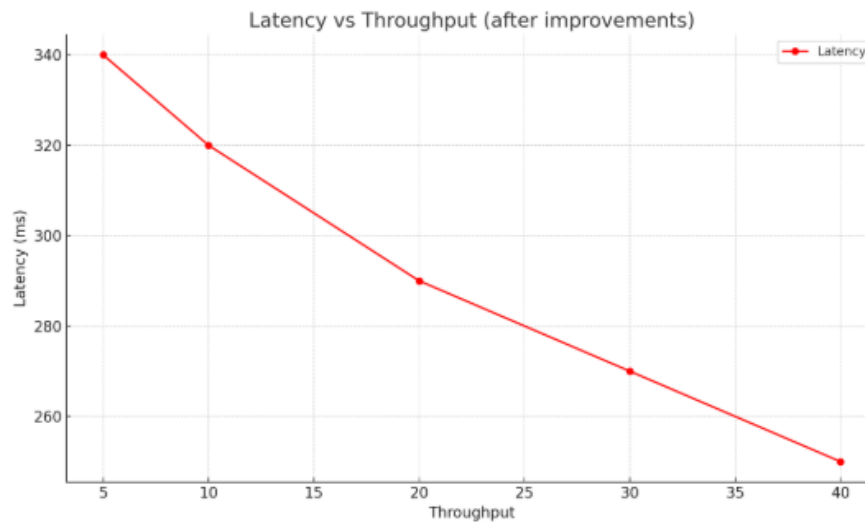


Figure 3: Latency vs. throughput

Figure 3 demonstrates the relationship between latency and throughput after implementing middleware automation. The data shows a clear inverse correlation: as throughput increases (measured along the X-axis), latency (measured along the Y-axis) decreases. The chart visually represents how, through improved resource allocation and optimisation techniques, higher throughput yields lower response times, thereby enhancing the overall system's performance and efficiency. The trend reflects the effectiveness of the automation strategies in improving service reliability and responsiveness.



Figure 4: SLA compliance improvement over time

Figure 4 illustrates the significant improvement in Service Level Agreement (SLA) compliance resulting from the implementation of middleware automation. The SLA compliance percentage increases from 96.5% before the automation to 99.9% after the improvements. This upward trend underscores the effectiveness of automation strategies, including proactive monitoring, automated recovery, and geo-redundant deployment, in ensuring near-continuous availability and meeting stringent performance standards for federal IT systems. The consistent improvement in SLA compliance reflects enhanced service reliability and efficiency.

5. Discussion

This middleware automation strategy showcases a real-world federal cloud transformation. Unlike theoretical models, this implementation was stress-tested under live conditions. By combining open-source tools with certified cloud services, the solution remained cost-effective yet secure.

Challenges encountered included:

- RBAC misconfigurations during federated identity testing.
- Terraform state conflicts requiring remote backend locking.
- Grafana plugin compatibility in air-gapped environments.
- Each was resolved with updated scripts, pre-deployment checks, and layered observability.

6. Conclusion Future Work

This study demonstrates a robust and resilient middleware automation strategy that effectively addresses the operational demands and compliance mandates of federal cloud-native systems. By leveraging the combined capabilities of Red Hat OpenShift for container orchestration, Microsoft Azure for scalable infrastructure, Terraform for Infrastructure-as-Code, and Jenkins for CI/CD automation, the proposed framework ensures high availability, rapid deployment, cost efficiency, and strict adherence to federal security standards. Additionally, observability tools such as Prometheus and Grafana enhance real-time monitoring and performance management, enabling proactive detection and resolution of incidents. The modular and scalable architecture not only meets the needs of mission-critical government systems but also serves as a reusable blueprint for similar deployments across the public sector and regulated industries.

Looking ahead, several enhancements are planned to further elevate the system's capabilities. Integrating service meshes such as Istio will provide advanced traffic control, telemetry, and mutual TLS (mTLS) encryption for secure intra-service communication. The adoption of Zero Trust Network Access (ZTNA) will strengthen identity and device-based access controls, aligning with modern cybersecurity principles. Moreover, the framework will be extended to multi-cloud environments using tools like Crossplane and Cluster API, ensuring workload portability, policy enforcement, and resilience across heterogeneous cloud providers. These future directions aim to transform the middleware layer into a fully autonomous, policy-aware, and platform-agnostic foundation for next-generation federal IT systems.

Acknowledgement: I sincerely acknowledge the guidance and support of Hull IT Solutions and Services, Llano, Texas, United States of America, whose valuable insights and encouragement were instrumental in the successful completion of this work.

Data Availability Statement: The data for this study can be made available upon request to the author.

Funding Statement: This manuscript and research were prepared without any financial support or funding.

Conflict of Interest Statement: The author declares no conflicts of interest. This work is an original contribution by the author, and all citations and references are appropriately included based on the information utilized.

Ethics and Consent Statement: This research adheres to ethical guidelines, and informed consent was obtained from all participants.

References

1. B. Burns, B. Grant, D. Oppenheimer, E. Brewer, and J. Wilkes, "Borg, Omega, and Kubernetes," *ACM Queue*, vol. 14, no. 1, pp. 70–93, 2016.
2. R. N. Calheiros, R. Ranjan, A. Beloglazov, C. A. F. De Rose, and R. Buyya, "CloudSim: A toolkit for modeling and simulation of cloud computing environments and evaluation of resource provisioning algorithms," *Software: Practice and Experience*, vol. 41, no. 1, pp. 23–50, 2011.
3. Centers for Disease Control and Prevention (CDC), "Monitoring modern infrastructure with Prometheus," *Whitepaper*, 2022, Available: <https://www.cdc.gov/index.html>, [Accessed by: 20/02/2024].
4. J. Fowler, "The DevOps Handbook," *IT Revolution Press*, Portland, Oregon, United States of America, 2016.
5. Gartner, "Market trends in federal cloud adoption," *Gartner Research Report*, 2023, Available: <https://garnetint.com/>, [Accessed by: 20/02/2024].
6. IBM, "OpenShift vs. Kubernetes: Federal use cases," *IBM White Paper*, New York, United States of America, 2022.
7. Microsoft Azure, "FedRAMP compliance documentation," *Technical Report*, 2023, Available: <https://azure.microsoft.com/en-in/>, [Accessed by: 20/02/2024].
8. National Institute of Standards and Technology (NIST), "Security Guidelines for Cloud Middleware (NIST SP 800-210)," *Tech. Rep.*, Gaithersburg, Maryland, United States of America, 2021, Available: <https://www.nist.gov/>, [Accessed by: 20/02/2024].

9. Red Hat, "Case Study: NASA's Shift to OpenShift," *Tech. Rep.*, 2020, Available: <https://www.redhat.com/en>, [Accessed by: 20/02/2024].
10. J. Smith, "Cost optimization in hybrid clouds," *Journal of Cloud Engineering*, vol. 12, no. 3, pp. 45–67, 2020.
11. S. Brown, "Infrastructure as Code with Terraform" *O'Reilly Media*, Sebastopol, CA, United States of America, 2021.
12. Cloud Native Computing Foundation (CNCF), "Prometheus Best Practices," *Tech. Rep.*, 2023. Available: <https://www.cncf.io/>, [Accessed by: 20/02/2024].
13. Amazon Web Services (AWS), "FedRAMP High Benchmark Overview", *Tech. Rep.*, 2023. Available: <https://aws.amazon.com/>, [Accessed by: 20/02/2024].
14. L. Wang, "Resilience patterns in cloud middleware," *IEEE Trans. Services Computing*, vol. 15, no. 2, pp. 1–15, 2022.
15. U.S. Department of Defense (DoD), "Cloud Adoption Strategy," *Tech. Rep.*, 2023. Available: <https://www.usa.gov/agencies/u-s-department-of-defense>, [Accessed by: 20/02/2024].
16. N. V. C. Akula, "Optimizing regional disaster recovery in OpenShift: A multi-cluster approach with RHACM and ODF," *Int. J. Comput. Methodol. Informatics*, vol. 17, no. 1, pp. 7027–7038, 2025.
17. M. Villamizar, "Evaluating the elasticity and cost-efficiency of cloud-based microservices," *IEEE Trans. Cloud Comput.*, vol. 6, no. 2, pp. 1–15, 2018.
18. Grafana Labs, "Real-Time Monitoring for Federal Agencies," *Tech. Rep.*, 2023. Available: <https://grafana.com/>, [Accessed by: 20/02/2024].