



#### OPEN ACCESS

SUBMITTED 22 June 2025

ACCEPTED 28 June 2025

PUBLISHED 30 July 2025

VOLUME Vol.07 Issue 07 2025

#### CITATION

Hari Dasari. (2025). Implementing Site Reliability Engineering (SRE) in Legacy Retail Infrastructure. The American Journal of Engineering and Technology, 7(07), 169–179.

<https://doi.org/10.37547/tajet/Volume07Issue07-16>

#### COPYRIGHT

© 2025 Original content from this work may be used under the terms of the creative commons attributes 4.0 License.

# Implementing Site Reliability Engineering (SRE) in Legacy Retail Infrastructure

**Hari Dasari**

Expert Infrastructure Engineer Leading Financial Tech Company  
Aldie, Virginia

**Abstract:** During digital transformation, retail companies with legacy IT infrastructures struggle to maintain service dependability, scalability, and agility. Many mainframes, on-premise applications, batch processing processes, and monolithic codebases were not designed for today's dynamic operational contexts. Google-developed Site Reliability Engineering (SRE) approaches including Service Level Objectives (SLOs), automation, and blameless postmortems can bridge the gap between outdated systems and modern operational excellence. This article proposes gradual adoption, cultural change, and measurable service reliability improvements for legacy retail environments adopting SRE. A concentrated SRE rollout helped a national retail chain reduce toil and improve mean time to detect (MTTD), mean time to resolve (MTTR), and MTTR. The model shows that incremental SRE adoption can modernize legacy systems and prepare them for future innovation without comprehensive re-architecture.

**Keywords:** Site Reliability Engineering (SRE), Legacy Systems, Retail IT Infrastructure, Observability, Toil Reduction, Service Level Objectives (SLOs), Error Budgets, Incident Management, Automation in Operations, Hybrid Cloud Monitoring, Cultural Transformation, Batch Job Monitoring, DevOps in Legacy Environments, Mainframe Reliability, Retail Digital Transformation

## 1. Introduction

In today's fast-paced digital economy, merchants are under increasing pressure to provide smooth, reliable, and always-on customer experiences across both physical and digital channels. However,

many large retail firms are hampered by legacy IT systems that were created decades ago for a different operational paradigm and are struggling to satisfy modern reliability and scalability requirements. These legacy infrastructures often comprise mainframe programs, batch processing systems, point-of-sale (POS) software, and closely integrated middleware, all of which were designed for transaction processing and physical store availability rather than uptime.

Legacy systems are frequently associated with high operational risk, poor observability, minimal automation, and inadequate support for horizontal growth. As a result, these environments frequently experience outages, delayed issue resolution, and sluggish deployment processes. As customers' expectations for real-time availability and omnichannel integration rise, these failures can result in revenue loss, reputational damage, and damaged client trust.

Site Reliability Engineering (SRE), pioneered by Google [1], is a potential technique to modernizing operations that does not require a full-scale system rearchitecture. SRE focuses on automation, monitoring, dependability as a measurable goal, and a strong engineering approach to operations. While SRE was originally intended for cloud-native systems, its key principles—such as creating Service Level Objectives (SLOs), controlling error budgets, decreasing toil, and introducing automated alerting—can be wisely applied to legacy environments.

This paper intends to bridge the gap between literature and reality by giving a structured technique for integrating SRE in legacy retail systems. It describes a staged strategy, examines tool selection and organizational alignment, and includes a real-world case study to show practical benefits. According to the study, SRE can be an effective technique for increasing the resilience and observability of legacy systems, as well as driving

incremental modernization with clear metrics and responsibility.

## **2.Characteristics of Legacy Retail Infrastructure**

Many major and established retail chains rely heavily on legacy retail infrastructures to run their operations. These systems were frequently developed in-house or acquired through mergers over decades, with technologies that predated modern software engineering concepts. Although stable and functionally rich, these legacy systems provide substantial hurdles in today's rapidly changing retail world, where real-time response, cloud integration, and customer-centric agility are critical.

### **2.1. Monolithic Architecture and Tight Coupling**

Legacy systems are often monolithic, with business logic, data access, and UI components inextricably linked. This tight linkage prevents modular updates and makes even minor modifications dangerous and time-consuming.

### **2.2. Batch-oriented processing.**

Many older systems rely on nightly or scheduled batch jobs to perform key tasks like inventory updates, pricing synchronization, and loyalty point calculations. This paradigm lacks the responsiveness necessary for real-time decision-making and dynamic pricing.

### **2.3. Limited observability.**

These systems were not created with telemetry in mind. Logs may be minimal, metrics absent, and tracing unsupported. As a result, root cause analysis becomes reactive and dependent on human skill.

### **2.4. Manual Implementation and Change Management**

Legacy infrastructures frequently lack CI/CD pipelines. Changes are manually deployed and validated, which increases the possibility of human mistake and lengthens the lead time for updates.

### **2.5. Fragile Integration with Modern Systems**

Retailers typically use fragile connections like message queues or custom adapters to layer newer services on top of historical cores, such as mobile apps, e-commerce APIs, and loyalty platforms. These integration points are prone to versioning difficulties, delays, and inconsistent behavior.

As presented in Table 1, these systemic limitations obstruct the adoption of modern engineering practices such as automated failover, real-time observability, and scalable service delivery.

Challenge	Description
Scalability	Difficult to scale beyond vertical scaling due to monolithic design
Observability	Minimal logging, no metrics or tracing, and limited visibility
Change Management	Manual deployments and change control processes increase operational risk
Incident Response	Reactive approach with little automation or early detection
Integration Limitations	Fragile middleware and message queues for connecting to modern retail systems
Operational Cost	High maintenance costs due to outdated hardware and lack of automation

**Table 1: Challenges in Legacy Retail Infrastructure**

## 2.6. Organizational and Cultural Inertia

The organization's opposition to change is equally challenging. Teams operating old systems may have deeply embedded routines, little exposure to contemporary tooling, and legitimate concerns about changing mission-critical systems that "still work."

Together, these qualities create a complicated context in which reliability engineering cannot be used as a drop-in solution. Instead, a progressive, context-sensitive adoption of SRE principles is required, as discussed in the sections below.

## 3. Principles of SRE Relevant to Legacy Systems

Site **Reliability** Engineering (SRE) is a systematic strategy to improving service dependability by incorporating software engineering principles into infrastructure and operations. While SRE was developed for large-scale, distributed, cloud-native environments, its key ideas are extremely adaptable to legacy retail infrastructures if implemented wisely.

This section goes over the major SRE ideas and how they can be applied to the limits of legacy

environments without requiring a total system overhaul.

### 3.1 Service Level Objectives (SLO) and Service Level Indicators (SLI)

SLOs specify the desired reliability of a system in measurable terms. Even for legacy systems, meaningful SLIs can be defined based on accessible information. Setting realistic SLOs ensures that the business and technical teams have a clear knowledge of acceptable service performance.

For example, a merchant may specify a SLO as requiring 99.5% of overnight inventory reconciliation activities to be completed by 6:00 a.m. each day.

### 3.2 Error budgets

An error budget is the permitted level of unavailability based on the set SLO. It guides judgments regarding implementing changes. Error budgets in legacy systems can help balance innovation (for example, automation) with fragile systems' risk tolerance.

Benefit: Helps to justify why not every change should be sent into production immediately in systems with limited rollback alternatives.

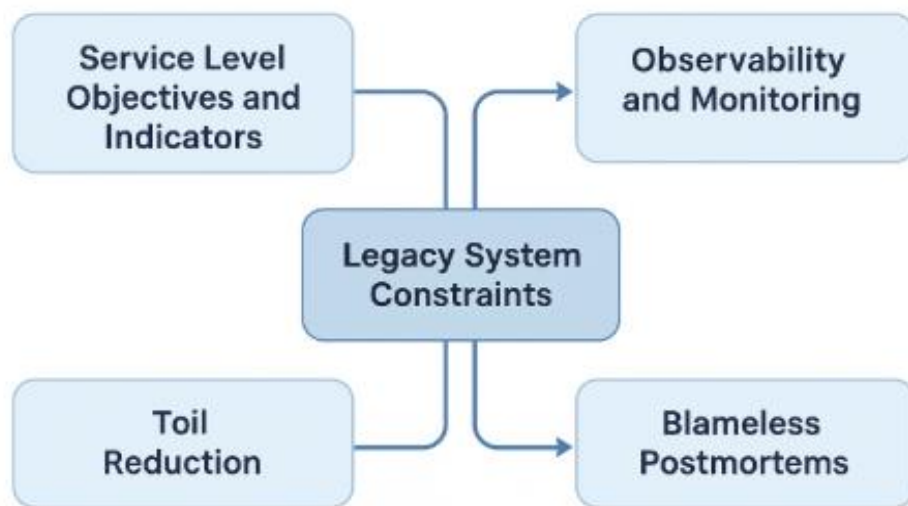
### 3.3 Toil Reduction

Toil refers to manual, repetitive, and automatable labor. Toil is abundant in older setups, whether it's monitoring batch job statuses, manually checking logs, or doing normal restart activities. Identifying and lowering labor through simple automation (scripts, scheduled jobs, etc.) is a high-return SRE practice.

For example, instead of depending on manual checks, send out automated notifications when crucial batch operations fail.

### 3.4 Observability & Monitoring

Legacy systems frequently lack out-of-box observability. However, integrating basic monitoring with technologies such as Prometheus, Nagios, or even custom scripts can provide immediate results. Log shipping and metric extraction from outdated programs are basic tasks. As illustrated in Figure 1, legacy constraints do not preclude SRE adoption; rather, they necessitate customized implementations of its fundamental concepts.



**Figure 1: Applying Core SRE Principles to Legacy System Constraints**

This picture shows how SRE principles like SLOs, toil minimization, and observability may be incrementally applied to legacy infrastructures.

### 3.5 Blameless Postmortems

Legacy. IT operations frequently have a blame based culture. SRE focuses on learning from failure through systematic, blameless incident reviews. This move promotes collaboration and ongoing improvement, even among traditional IT teams.

Impact: Encourages teams to investigate fundamental issues thoroughly and communicate results without fear of repercussions.

### 3.6 Automation for Incident Response

SRE promotes the creation of automated playbooks and incident detection techniques. Even in legacy systems, teams can increase responsiveness and MTTR by defining standard operating procedures (SOPs) and utilizing scripting or lightweight orchestration.

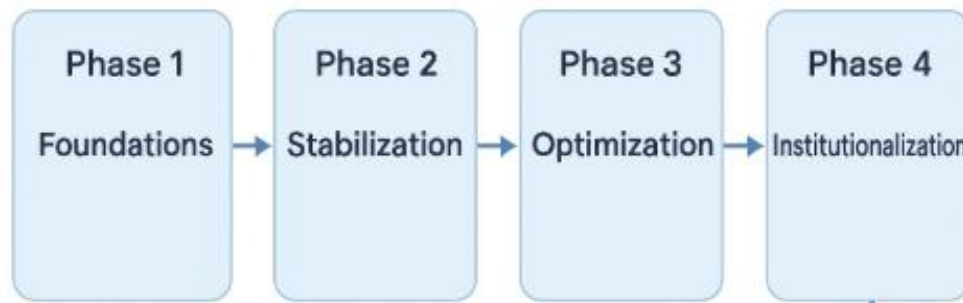
Summary: Applying these approaches incrementally can significantly improve legacy retail settings without requiring complete system rewrites. The following section will describe an organized, staged methodology for efficiently implementing these concepts in such circumstances.

#### 4. Architectural Strategies for Resilience

Applying SRE to legacy retail infrastructure necessitates a practical, staged strategy. Legacy environments, as opposed to greenfield systems, cannot afford downtime or drastic reengineering. This section describes a four-phase implementation architecture customized to

legacy systems, allowing for gradual adoption of SRE principles while retaining operational stability.

As shown in Figure 2, the phased model enables a risk-aware, structured transition to SRE practices in legacy infrastructure without requiring a full system rewrite.



**Figure 2: Phased SRE Implementation Framework for Legacy Retail Systems**

This figure outlines the four phases of SRE adoption in legacy environments:

1. **Foundation (Monitoring & Visibility)**
2. **Stabilization (SLOs & Error Budgets)**
3. **Optimization (Automation & Toil Reduction)**
4. **Institutionalization (Culture, Runbooks, Reviews)**

Phase 1: Foundation: Establish Monitoring and Visibility. The initial focus should be on observability. Legacy systems frequently lack native monitoring, necessitating the use of external agents, bespoke log scrapers, and simple dashboards. The purpose is to:

- Identify critical SLIs, such as job success rate and POS endpoint uptime.
- Implement warning systems for service anomalies.
- Use open-source tools such as Prometheus, Grafana, or Nagios for low-friction deployment.

For example, a COBOL-based nightly routine can be wrapped in a script that logs start/end times and failures to a file that a monitoring agent can analyze.

Phase 2: Stabilization—Define SLOs and Error Budgets With visibility in place, teams may set Service Level Objectives (SLOs) and Service Level Indicators (SLIs) based on business goals.

- Implement internal SLOs (for example, "95% of inventory sync jobs must complete within 2 hours")
- Define acceptable error budgets and use them to control dangerous changes.
- Use historical data to establish baselines and thresholds.

Outcome: Teams are empowered to make informed trade-offs between stability and rapidity.

Phase 3: Optimization - Reduce Labor Through Automation

Legacy teams are frequently plagued with manual duties that are ideal for automation.

- Automate repetitive operational operations, such as restarts, status checks, or file transfers.
- Implement programmed reactions to typical occurrences, such as disk space issues or service restarts.
- Use tools like Rundeck or Jenkins for basic orchestration with no big dependencies.

For example, replacing a 10-step manual health check process with a nightly automated validation script increases consistency and decreases on-call fatigue.

Phase 4: Institutionalization (Runbooks, Reviews, and Culture)

The final phase focuses on embedding SRE practices in the organization.

- Generate runbooks for repeatable operational procedures.
- Conduct blameless postmortems following incidents to determine systemic remedies.
- Assign ownership and escalation paths to legacy services.
- Create cross-functional SRE pods for legacy systems.

Goal: Create a learning-oriented, dependable, and responsive operational culture that meets modern reliability criteria.

Summary: This methodology allows older infrastructures to implement SRE in measured, controlled increments, rather than disruptive overhauls. Each phase builds on the previous one, allowing even the most fragile and mission-critical systems to enhance dependability through visibility, standardization, and automation.

## 5. Case Study: SRE in a National Retail Chain

To highlight the practical implementation and impact of SRE in a legacy environment, this section presents a real-world case study of a US-based national retail chain that manages over 1,200 physical locations and runs on a legacy technology stack consisting of:

- Mainframe-based batch systems for overnight processing.
- Distributed in-store POS systems with nightly data uploads.
- FTP and message queue interaction with e-commerce platforms

The organization encountered frequent job failures, long incident reaction times, and poor observability, limiting its ability to satisfy operational SLAs and respond to system outages efficiently.

### 5.1 Business Challenge

The IT operations team dealt with the following issues:

- Inventory synchronization jobs would periodically fail, going undiscovered until the morning.
- Manual triage and troubleshooting took up to 3-4 hours each day.
- Operations relied on outdated monitoring tools without real-time alerts or dashboards.
- There are no established SLOs or accountability mechanisms for reliability breaches.

### 5.2 The SRE Adoption Approach

The organization followed the tiered SRE deployment model outlined in Section 4:

- Monitoring setup includes Prometheus and Grafana agents to track job completion status and synchronization events. Logs from FTP jobs and POS uploads were collected using Fluentd and parsed into structured representations.
- SLOs were designed for essential batch jobs, such as:
  - 99.5% success percentage for POS uploads by 3:00 a.m.
  - 99% availability for in-store system check-ins.
- A set of Python scripts and Rundeck jobs were created to automate:
  - Daily job health inspections.
  - Monitor disk utilization on mainframes.
  - Alerting for late or missing data uploads

Cultural Shifts: The team used blameless postmortems, developed a runbook library for 20+ frequent failure situations, and reduced manual incident resolution.

As shown in Table 2, SRE adoption led to a 70% reduction in MTTR, an 11% improvement in job reliability, and a substantial drop in SLA breaches within six months.



Metric	Before SRE	After SRE (6 Months)
Mean Time to Detect (MTTD)	40 minutes	12 minutes
Mean Time to Resolve (MTTR)	120 minutes	45 minutes
Job Failure Rate	15%	4%
Manual Health Check Effort	~3 hours/day	<30 minutes/day
SLA Violations per Month	18	3

**Table 2: Key Metrics Before and After SRE Implementation**

## 6. Monitoring, Tooling, and Automation

Legacy retail systems have generally lacked robust observability and relied primarily on human activities, causing gaps in reliability management. Site Reliability Engineering (SRE) focuses on proactive monitoring, standardized tooling, and automation to detect, respond to, and avoid problems before they affect consumers or downstream systems. This section investigates the practical application of current monitoring and automation approaches in traditional retail environments.

### 6.1 Monitoring in Legacy Systems

Legacy systems, such as mainframes, POS terminals, or FTP-driven batch processors, often do not support native monitoring interfaces. However, indirect instrumentation and wrapper-based logging can be used to collect meaningful telemetry:

- **Log Parsing:** Use Fluentd or Logstash to extract structured events from legacy logs
- **Job Status Tracking:** Wrap batch jobs with scripts that emit start, end, and error codes to a centralized logging pipeline
- **System Metrics Collection:** Utilize tools like Node Exporter or Telegraf to gather CPU, memory, and disk metrics from legacy servers

- **Basic Health Probes:** Implement synthetic checks (e.g., TCP port pings, job output checksum validators)

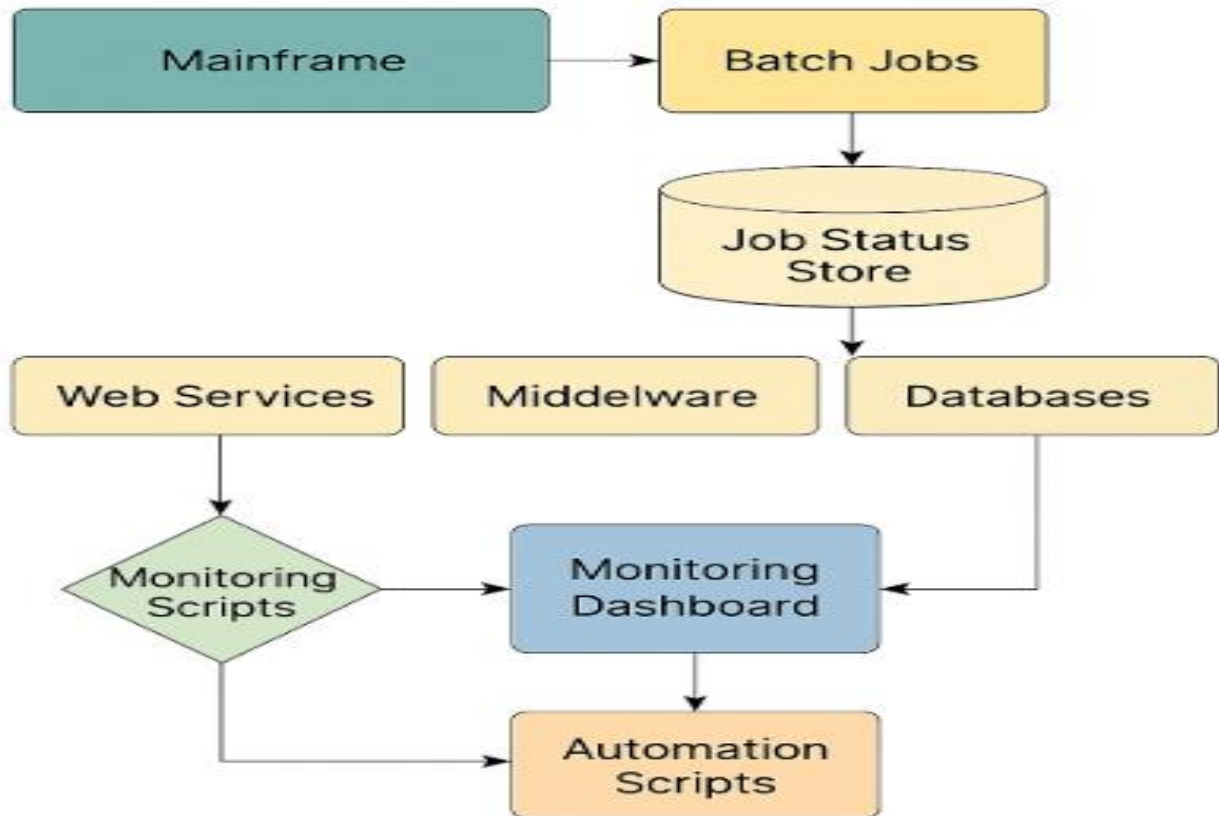
*Result:* These lightweight additions enable visibility into performance and failure conditions without altering the core legacy system code.

### 6.2 Automation Opportunities

SRE encourages reducing toil—the manual, repetitive operational work that can and should be automated. In legacy systems, common areas for automation include:

- **Job Monitoring & Alerts:** Scheduled health checks for jobs and transfers using cron jobs or Rundeck
- **Auto-recovery Scripts:** Restart stuck processes or clear cache directories based on log patterns
- **Incident Notification:** Integration with PagerDuty, Slack, or email alerts via Alertmanager
- **Routine Maintenance:** Automate log rotation, archive compression, disk cleanup, and backup validation

As shown in Figure 3, this reference architecture enables centralized monitoring and responsive automation without modifying the legacy core systems.



**Figure 3: Monitoring and Automation Architecture in a Legacy Retail Stack**

This diagram depicts a sample architecture integrating open-source tooling with legacy systems:

- **Data Sources:** FTP logs, mainframe job logs, store check-in data
- **Collection Layer:** Fluentd, Node Exporter
- **Processing/Alerting:** Prometheus, Alertmanager
- **Visualization:** Grafana dashboards
- **Automation/Orchestration:** Rundeck, custom scripts (Python/Shell)

### 6.3 Tool Selection Criteria for Legacy Environments

When selecting tools for legacy integration, the following criteria should be prioritized:

- **Lightweight & Low Overhead:** Must Operate on resource-constrained servers
- **Agentless or Minimal Agents:** Reduces friction with legacy system administrators

- **CLI/Script Friendly:** Should support integration with shell scripts and cron jobs.
- **Open Source:** Minimizes licensing cost and avoids vendor lock-in
- **Extensible & Interoperable:** Able to interface with modern DevOps platforms as needed

### 6.4 Benefits of Integrated Monitoring and Automation

- **Faster Issue Detection and Resolution:** Alerts generated within seconds of job failure improve MTTR significantly.
- **Operational Efficiency:** Scripts automate common resolutions, allowing engineers to focus on strategic initiatives.
- **Consistency and Reliability:** Standardized monitoring eliminates reliance on tribal knowledge or manual log reviews.
- **Incremental Modernization:** Serves as a foundation for broader infrastructure refactoring and DevOps adoption.



Legacy systems can increase their reliability, observability, and operational efficiency significantly by utilizing non-invasive tools and targeted automation, bringing them closer to the requirements of modern cloud-native platforms.

7. Risk Mitigation and Organizational Alignment

Implementing Site Reliability Engineering (SRE) in legacy retail environments takes more than just tools and scripts; it also necessitates a strategic approach to technical risk management and organizational buy-in. Legacy systems are frequently mission-critical, unstable, and managed by teams with decades of combined experience. As a result, successful SRE deployment requires risk mitigation and stakeholder alignment across engineering, operations, and business groups.

7.1 Technical Risk Mitigation Strategies

Legacy environments impose limits such as unsupported technologies, undocumented dependencies, and limited automation capabilities. These risks can be reduced using the following strategies:

- **Pilot-Based Rollouts:** Implement SRE on low-risk, non-customer-facing systems (e.g., internal reports, reconciliation jobs). Use

early successes to boost confidence and improve techniques.

- **Fallback Mechanisms:** Automation scripts should have error handling and rollback features as fallback mechanisms. Avoid automation that generates irreversible states.
- **Observability Before Automation:** Prioritize observability over automation. First, use logging and analytics to better understand system behavior in both normal and degraded settings.
- **Change Isolation:** To introduce new reliability measures or scripts, limit the blast radius by using canary updates, blue-green deployments, or job-level toggles.

For example, before automating work restarts, monitor failure rates and recovery times for at least 30 days to confirm that the automation logic appropriately reflects operator actions.

7.2 Organizational Alignment

Cultural resistance is a major barrier in legacy environments, especially where teams are accustomed to traditional ITIL-style incident response and hierarchical approvals.

Table 3 illustrates key alignment strategies within the organization

Strategy	Description
Executive Sponsorship	Secure top-down support to prioritize SRE as a reliability and modernization enabler
Cross-Functional Collaboration	Form hybrid SRE pods with members from operations, development, and infrastructure teams
Training and Upskilling	Conduct SRE bootcamps focused on tooling, monitoring, incident handling, and runbook development
Communication and Transparency	Share dashboards, SLOs, and error budget policies across stakeholders to build trust
Blameless Culture Initiatives	Promote post-incident reviews without assigning blame to encourage openness and improvement

Table 3: Key Alignment Strategies

### 7.3 Stakeholder Roles in SRE for Legacy Systems

Stakeholder	Role in SRE Adoption
CTO / CIO	Provide vision, funding, and mandate for reliability programs
Operations Managers	Define pain points, assist in toil identification
Development Teams	Integrate observability into legacy code, assist in defining SLOs
Business Analysts	Translate service impact into business risk/value language
SRE Champions	Act as internal consultants and cultural liaisons

**Table 4: Stakeholder Roles in Legacy Systems**

#### 7.4 Governance and Change Management

It is vital to establish governance structures that enable agile dependability practices in legacy systems. This includes:

- Establishing escalation rules and on-call rotations, as well as documenting change approval processes for automated operations.
- Assigning service ownership, particularly for neglected or "black-box" systems.
- Including dependability KPIs in performance appraisals and project success metrics.

Outcome: Formal governance ensures that SRE processes are not considered as side projects but rather integrated into day-to-day operations.

#### 7.5 Success Metrics for Risk and Alignment

To evaluate the success of SRE adoption in legacy systems, organizations should track the following metrics:

- Reduction in Mean Time to Detect (MTTD) and Resolve (MTTR)
- % of systems with defined SLOs and automated health checks
- Number of manual runbook tasks automated

- Stakeholder satisfaction (via surveys or incident debriefs)
- Reduction in unplanned downtime or SLA violations

Enterprises may successfully adopt SRE in even the most entrenched legacy systems by aligning technological efforts with organizational strategy and controlling risks in advance, changing them from operational liabilities to robust, insight-driven assets.

#### 8. Future Trends and Recommendations

As the retail business evolves in response to digital disruption, economic instability, and customer experience demands, Site Reliability Engineering (SRE) will become increasingly important in guaranteeing uptime and resilience, particularly for older systems. While legacy environments provide problems, current trends and innovations provide fresh opportunity to use SRE principles more effectively in these contexts.

##### 8.1 AI and Machine Learning in Legacy Operations

AI and machine learning are rapidly being used in current infrastructures. When customized for legacy contexts, they provide the following benefits:

- **Anomaly detection:** Machine learning models can identify irregularities in batch

task durations, POS log patterns, and transaction spikes that indicate approaching failures.

- **Predictive Maintenance:** AI models trained on past system records can predict hardware failures or processing delays before they disrupt operations.
- **Incident Classification:** NLP and machine learning can help categorize incident tickets and discover reoccurring failure patterns in legacy systems.

For example, a shop could employ anomaly detection on historical batch processing times to send out preemptive alerts before SLAs are exceeded.

## 8.2 Gradual Refactoring with SRE Metrics

Rather than embarking on risky and costly re-platforming initiatives, businesses can leverage SRE-driven metrics (e.g., error budgets, toil indices, MTTR trends) to identify high-impact components for targeted modernization.

- Modularize historical batch tasks into API-driven services based on common failure locations.

- Replace fragile integrations (for example, flat-file FTP transfers) with event-based messaging systems like Kafka.
- Track SLO breaches and prioritize technical debt correction.

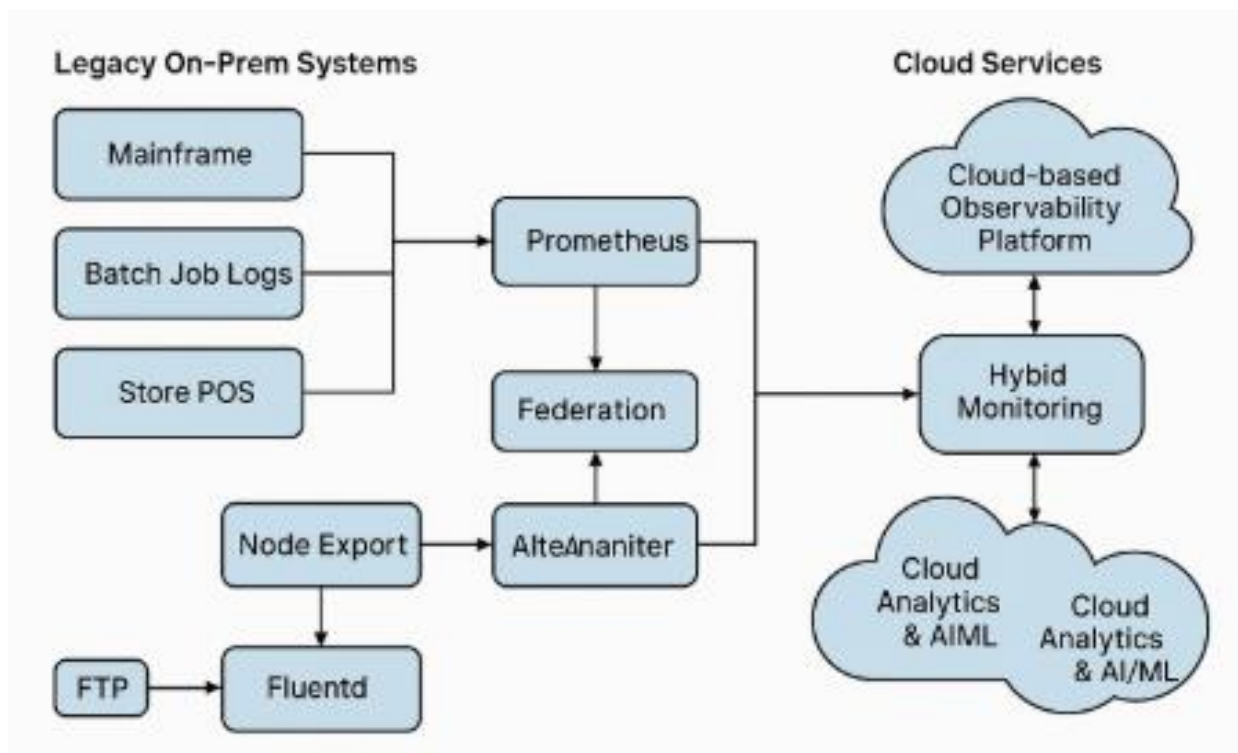
**Recommendation:** Use SRE telemetry to inform modernization roadmaps, ensuring that engineering efforts are focused on the most reliable older components.

## 8.3 Hybrid and Multi-Cloud Integration

Many retailers are adopting hybrid cloud models to extend legacy capabilities into cloud environments. SRE will be essential in:

- Monitoring service reliability across on-prem and cloud components
- Maintaining consistent observability and alerting pipelines
- Coordinating change management across platforms

As shown in Figure 4, hybrid observability models will be essential as legacy systems evolve toward distributed retail platforms.



**Figure 4: Future-State Hybrid SRE Monitoring Model for Legacy + Cloud Retail Environments**

## 8.4 Chaos Engineering for Legacy Systems

Chaos engineering, which has traditionally been

used for microservices, is rapidly being adopted to improve fault tolerance in legacy systems. Organizations can simulate batch task delays, data feed issues, and resource fatigue.

- Validate incident response readiness.
- Improve the automatic rollback methods.
- Uncover hidden dependencies in monolithic code bases.

Caution: Chaos experiments in legacy

environments must be carefully planned and monitored to avoid causing widespread failures.

8.5 Strategic Recommendations

The strategic recommendations summarized in Table 3 are informed by practices outlined in SRE foundational literature [1], DevOps transformation frameworks [3], and modernization guidance for legacy systems [8][9].

Recommendation	Justification
Start with non-critical legacy components	Minimizes risk, builds early momentum
Establish a central SRE enablement team	Encourages knowledge sharing, tooling standardization
Treat SRE as a culture shift, not a toolkit	Ensures long-term organizational alignment
Use SLOs and error budgets to drive priorities	Aligns engineering focus with business reliability expectations
Track metrics across legacy and modern systems	Provides end-to-end service visibility for hybrid environments

Table 5: Strategic Recommendations

As shops continue to innovate, the ability to function consistently across historical and modern systems will become an important differentiation. Adopting SRE not only improves present operations, but also sets the door for a more seamless, metric-driven transition to next-generation retail solutions.

9.Conclusion

The implementation of Site Reliability Engineering (SRE) in legacy retail infrastructure is a critical step toward updating operating processes without requiring high-risk system overhauls. Legacy systems, while reliable and business-critical, suffer from low observability, high operational workload, and reactive incident response. SRE addresses these difficulties with a structured, data-driven strategy that includes Service Level Objectives (SLOs), error budgets, proactive monitoring, and automation.

This article offered a phased implementation methodology targeted to traditional retail environments, allowing for gradual adoption of SRE concepts while retaining system stability. It investigated how open-source tools and lightweight scripts might increase observability while reducing toil, emphasizing the importance of corporate alignment and cultural transformation in achieving success. A real-world case study found measurable gains in reliability, detection time, and operational efficiency, supporting the importance of SRE in non-cloud-native environments.

Furthermore, the report discussed new trends including AI-powered monitoring, hybrid observability architectures, and chaotic engineering for legacy systems. These developments promise to broaden the scope of SRE, preparing legacy infrastructures for a future of resilient, scalable, and customer-centric retail operations.

To summarize, SRE is more than a collection of tools

or processes; it is a cultural and strategic transformation that enables merchants to manage complexity, assure uptime, and continuously improve the dependability of even their oldest systems. For enterprises looking to bridge the gap between stability and innovation, SRE provides a practical and proven solution.

## References

1. Beyer, B., Jones, C., Petoff, J., & Murphy, N. R. (2016). *Site Reliability Engineering: How Google Runs Production Systems*. O'Reilly Media.
2. Krief, M. (2019). *Learning DevOps: Continuously Deliver Better Software*. Packt Publishing.
3. Kim, G., Humble, J., Debois, P., & Willis, J. (2016). *The DevOps Handbook: How to Create World-Class Agility, Reliability, & Security in Technology Organizations*. IT Revolution Press.
4. Burns, B., Grant, B., Oppenheimer, D., Brewer, E., & Wilkes, J. (2016). Borg, Omega, and Kubernetes. *Communications of the ACM*, 59(5), 50–57. <https://doi.org/10.1145/2890784>
5. OpenSLO. (2021). *Open Specification for SLOs*. <https://openslo.com>
6. Thongmak, M. (2022). Applying AI in IT Operations: Anomaly Detection and Incident Prediction in Legacy Systems. *Journal of Information Technology Management*, 33(1), 35–42.
7. Allspaw, J. (2017). *Blameless PostMortems and a Just Culture: A Guide to Incident Investigation*. Etsy Engineering. <https://codeascraft.com>
8. Gartner. (2023). *Predicts 2023: Legacy Systems Modernization Strategies for CIOs*. Gartner Research.
9. Woodcock, S. (2020). *Automating Legacy Systems: Practices and Pitfalls*. *IEEE Software*, 37(4), 67–73. <https://doi.org/10.1109/MS.2020.2996582>