

CHAPTER 13

Monitoring, Operations and Service Management

From Reactive Support to Proactive Excellence

“Operational excellence is not about fixing problems faster; it is about building systems that prevent problems from occurring in the first place.”

— Sazzad Khan

Governance sets the boundaries. FinOps makes the spending visible. What remains is proving, every day, that the platform actually works.

EXECUTIVE SUMMARY

For Nordic Ventures Group, the move from twelve distributed datacentres to Azure is a transformation in how services are observed, operated, and improved. Fragmented environments create reactive support, inconsistent monitoring, duplicated tooling, and costly hand-offs. A modern operating model replaces those patterns with observability, reliability engineering, automation, and service ownership.

This chapter explains how Azure Monitor, Log Analytics, Application Insights, Microsoft Sentinel, automation, IT service management, and Site Reliability Engineering can form a coherent operating system. The emphasis is not on collecting more telemetry. It is on turning signals into decisions and decisions into reliable business services.

Operational excellence protects factory production, SAP processes, Dynamics 365 workflows, Microsoft 365 integrations, and emerging AI platforms. It also supports business continuity, regulatory evidence, and international expansion. Modern operations are

therefore a strategic capability, not an administrative afterthought.

Why Operational Excellence Matters

Availability, performance, and recoverability directly influence revenue, employee productivity, customer trust, and safety. In a five-factory enterprise, a technical incident can quickly become a production, logistics, or customer commitment issue.

Traditional Operations	Modern Cloud Operations
Reactive	Proactive and predictive
Siloed monitoring	End-to-end observability
Manual remediation	Automated response
Infrastructure-focused	Service and customer-journey focused
Availability-only metrics	Reliability and business outcome metrics
Ticket-driven	Product ownership and reliability engineering

The Nordic Ventures Group Challenge

Current State	Future State
Twelve monitoring islands	Enterprise telemetry and common operational views
Alert overload	Prioritised, correlated signals
Manual recovery	Automated and tested remediation
Unclear service ownership	Named owners and service-level objectives

Current State	Future State
Change risk managed through delay	Change risk managed through engineering and evidence
Operational metrics disconnected from business	Service health linked to factory, customer, and financial impact

Operations at a Glance

Capability	Primary Technology / Practice	Business Outcome
Metrics and logs	Azure Monitor and Log Analytics	Enterprise visibility
Application telemetry	Application Insights and distributed tracing	Faster diagnosis
Security operations	Microsoft Sentinel and Defender capabilities	Coordinated detection and response
Reliability management	SLIs, SLOs, and error budgets	Balanced speed and stability
Automation	Runbooks, functions, workflows, and pipelines	Lower toil and faster recovery
Service management	Modern ITSM integrated with engineering	Consistent ownership and communication

1. Observability Before Monitoring

Monitoring answers known questions by checking predefined conditions. Observability helps teams investigate conditions they did not anticipate. A mature platform combines both: dashboards and alerts for known risks, plus rich telemetry and correlation for unexpected behaviour.

Pillar	What It Reveals	NVG Example
Metrics	Numeric trends and system health	Factory API latency and SAP queue depth
Logs	Detailed events and context	Authentication failures and deployment changes
Traces	Request flow across dependencies	Order journey from Dynamics 365 to SAP and factory systems

The three pillars are most valuable when they share consistent identifiers, timestamps, service ownership, and environment metadata. Without that context, centralization can produce a larger data lake without producing better decisions.

BUSINESS VALUE

- Faster diagnosis across SAP, Dynamics 365, Microsoft 365, factory, and Azure dependencies.
- Reduced mean time to detect and recover.
- Better capacity planning and trend analysis.
- Evidence for operational assurance and incident review.

TRUSTED ADVISOR INSIGHT *This is where inexperienced architects usually fail: they instrument the infrastructure beautifully and leave the business transaction invisible.*

2. Centralised Monitoring and Operational Views

Capability	Technology
Infrastructure monitoring	Azure Monitor

Capability	Technology
Log aggregation and analytics	Log Analytics
Application performance monitoring	Application Insights
Security monitoring	Microsoft Sentinel
Dashboards and service views	Azure Workbooks
Alerting and action	Azure Monitor alerts and action groups

A single pane of glass should not become a single giant dashboard. Executives, service owners, platform engineers, security operations, and factory teams require different views built from common telemetry. The platform should provide consistency while preserving role-specific context.

“You cannot improve what you cannot observe, and you cannot operate what nobody clearly owns.” — Trusted Advisor Perspective

3. AIOps — Reducing Noise and Increasing Insight

Modern enterprises generate more events than people can interpret manually. AIOps applies analytics and machine learning to identify anomalies, correlate related signals, suppress duplicates, and support faster prioritisation. It should augment human judgment, not create an opaque automation layer that teams cannot challenge.

Traditional Monitoring	AIOps-Enabled Operations
Reactive response	Predictive and anomaly-based detection
High alert volume	Correlation and intelligent prioritisation

Traditional Monitoring	AIOps-Enabled Operations
Manual pattern recognition	Automated signal analysis
Longer diagnosis cycles	Faster probable-cause identification
Static thresholds	Dynamic baselines where appropriate

For NVG, the factory perspective is decisive. Every minute of avoidable downtime can reduce production, delay delivery, and erode revenue. Predictive monitoring and intelligent correlation directly protect business performance.

4. Site Reliability Engineering

Site Reliability Engineering applies software engineering principles to operations. It shifts the objective from keeping infrastructure alive to engineering reliable services with explicit expectations, measurable indicators, controlled risk, and automation.

Principle	Objective
Service Level Indicators (SLIs)	Measure the behaviour users experience
Service Level Objectives (SLOs)	Define a target level of reliability
Error budgets	Balance change velocity and stability
Automation	Reduce repetitive manual work and recovery time
Blameless learning	Improve systems after incidents
Toil management	Protect engineering capacity for improvement
Service	Illustrative SLO

Principle	Objective
Identity services	99.99% successful authentication for supported journeys
SAP critical transactions	99.95% availability and agreed response time
Factory integration services	99.95% successful processing within target latency
Microsoft 365 integration	99.90% successful workflow completion
AI platform services	99.90% availability for approved production use cases

SLOs should be based on user journeys and business impact rather than arbitrary infrastructure uptime. The exact targets require business agreement, dependency analysis, and cost trade-offs.

BUSINESS VALUE

- Reliability expectations become explicit and measurable.
- Teams can balance innovation with operational risk through error budgets.
- Investment decisions reflect business criticality.
- Incident reviews create systemic improvement rather than blame.

TRUSTED ADVISOR INSIGHT *An SLO nobody remembers exists is not a reliability target. It is a slide from last year's steering committee.*

5. Modern IT Service Management

Traditional ITSM processes were often designed for static infrastructure and infrequent change. Cloud

environments require faster feedback, automated evidence, product ownership, and integration with CI/CD. The purpose of service management remains essential: clear ownership, controlled change, effective incident response, communication, and learning. The implementation must evolve.

Traditional ITSM	Modern ITSM
Ticket-driven	Service and product-driven
Manual provisioning	Self-service with policy
Large approval boards	Risk-based automated controls
Infrastructure-centric	Customer journey and product-centric
Reactive support	Proactive reliability management
Separate development and operations	Integrated engineering and operations

Change governance should focus human attention on material risk. Standard, tested, reversible changes can flow through automated pipelines with evidence. High-risk or irreversible changes deserve deeper review and coordinated execution.

6. Automation and Incident Response

Automation contributes directly to NVG’s cost and reliability objectives. Provisioning, diagnostic configuration, environment scheduling, common remediation, evidence collection, and communications can be automated when triggers, permissions, safeguards, and rollback paths are clearly defined.

Incident Stage	Automation Opportunity
Detect	Signal correlation, anomaly detection, and alert enrichment

Incident Stage	Automation Opportunity
Triage	Service ownership, business impact, and dependency lookup
Contain	Approved isolation or scaling actions
Recover	Runbooks, failover, restart, or rollback
Communicate	Status updates and stakeholder routing
Learn	Timeline creation, evidence retention, and action tracking

Automation should be proportional to confidence and impact. Low-risk, well-understood actions can be fully automated. High-impact actions may require human approval while still benefiting from automated analysis and preparation.

7. Continuous Compliance and Security Operations

Security and operations are inseparable. Central logging, detection, policy compliance, incident workflows, and retention support both operational resilience and assurance activities. These capabilities can help meet technical and evidentiary requirements, but they must be integrated with legal, risk, privacy, and business continuity processes.

Capability	Operational Outcome
Continuous monitoring	Current visibility into control and service health
Policy enforcement	Consistent technical configuration
Security alerting	Faster investigation and coordinated response

Capability	Operational Outcome
Centralised logging	Evidence and cross-service correlation
Audit reporting	Reduced manual evidence effort
Recovery testing	Demonstrated resilience rather than assumed resilience

8. Operations Maturity Model

Level	Characteristics	Next Priority
1 — Reactive	Fragmented tools, manual response, infrastructure focus	Establish ownership and central telemetry
2 — Standardised	Common monitoring, incident processes, basic automation	Define service views and SLOs
3 — Proactive	Correlation, capacity planning, tested runbooks	Expand reliability engineering and self-healing
4 — Predictive	AIOps, business-impact views, error-budget governance	Optimise across value, risk, and cost
5 — Adaptive	Continuous learning and automated policy-driven operations	Sustain improvement and strategic experimentation

TRUSTED ADVISOR INSIGHT *The best incident response playbook is the one that was rehearsed before anyone needed it for real.*

ARCHITECTURE DECISION RECORD

ADR-013: Adopt an SRE-Based Enterprise Operating Model

Context

- NVG requires higher reliability, global visibility, scalable operations, and improved resilience across critical business services.
- Traditional fragmented monitoring and ticket-driven operations do not support the target cloud operating model.

Decision

- Implement centralised observability using Azure Monitor, Log Analytics, Application Insights, security monitoring, and service-specific dashboards.
- Adopt SLIs, SLOs, error budgets, automation, modern ITSM, and structured incident learning.
- Prioritise business services and user journeys rather than infrastructure metrics alone.

Positive Consequences

- Lower operational cost and toil.
- Higher service availability and faster recovery.
- Better customer and employee experience.
- Greater resilience and safer innovation.

Negative Consequences

- New skills and cultural change.
- Telemetry cost and data-governance requirements.
- Risk of poorly controlled automation.

Mitigations

- Implement incrementally around critical services.
- Create an SRE and observability learning programme.
- Define retention, access, and cost controls for telemetry.

- Use human approval for high-impact automated actions.

KEY PERFORMANCE INDICATORS

KPI	Illustrative Target
Critical service availability	Service-specific SLO, typically above 99.9%
Mean Time to Detect (MTTD)	Less than 10 minutes for priority incidents
Mean Time to Recover (MTTR)	Less than 30 minutes for defined critical scenarios
Change failure rate	Below 10% for production changes
Automated remediation rate	Above 70% for approved repeatable scenarios
Toil as share of engineering time	Declining trend with explicit capacity for improvement
Recovery test success	100% of scheduled critical tests completed and evidenced

The Nordic Ventures Group Lesson

Operational excellence is not measured by the number of tickets closed. It is measured by how few preventable incidents occur, how quickly the organisation recovers, how much the system learns, and how confidently the business can change. Centralised observability and reliability engineering give NVG the operational foundation required for five factories, international growth, and increasingly digital products.

KEY TAKEAWAYS

Principle	What It Means
Observe services, not just infrastructure	Business journeys require end-to-end metrics, logs, and traces.
Automation must be trustworthy	Automate repeatable actions with safeguards, evidence, and rollback.
Reliability is a business capability	Availability and recovery directly influence revenue and trust.
SLOs make trade-offs explicit	Targets and error budgets balance stability, speed, and cost.
Operations and security converge	Shared telemetry and coordinated response improve resilience.
Learning is the objective of incident review	The system should improve after every meaningful failure.

REFLECTION QUESTIONS

- Can your organisation observe every critical business journey in real time?
- Which failures have the greatest factory, customer, or financial impact?
- How much engineering time is consumed by repetitive operational toil?
- Are KPIs focused on technology activity or business outcomes?
- Which recovery actions are sufficiently safe and repeatable to automate?
- Does the operating model make change safer, or merely slower?

REFERENCES

- Beyer, B., Jones, C., Petoff, J., & Murphy, N. R. (2016). Site reliability engineering. O'Reilly Media.
- Beyer, B., Murphy, N. R., Rensin, D. K., Kawahara, K., & Thorne, S. (2018). The site reliability workbook. O'Reilly Media.
- Majors, C., Fong-Jones, L., & Miranda, G. (2022). Observability engineering. O'Reilly Media.
- Kim, G., Humble, J., Debois, P., & Willis, J. (2021). The DevOps handbook (2nd ed.). IT Revolution Press.
- Microsoft. (2025). Azure Monitor documentation.
- Microsoft. (2025). Azure Well-Architected Framework: Operational Excellence pillar.
- Microsoft. (2025). Azure Architecture Center: Reliability guidance.
- PeopleCert. (2019). ITIL Foundation: ITIL 4 edition.
- National Institute of Standards and Technology. (2024). Cybersecurity Framework 2.0.

TRUSTED ADVISOR PRINCIPLE

Observability is not a dashboard.

It is the difference between knowing and guessing.

Automation must be trustworthy — repeatable actions need safeguards, evidence, and rollback.

The organisations that recover fastest rehearsed the failure long before it happened.

The Trusted Advisor measures operations in minutes to detection, not slides in a deck.

Strategy, architecture, governance, cost, and operations now stand as one coherent system.

What remains is proving that system holds under conditions no one has tested yet.