


IT is Getting Smarter

**Elevate IT
with Smart
and Agile
Operations**

CIO Compass Series



Introduction

IT operations in large organizations are built on foundations that no longer meet today's demands. Multi-cloud, edge, microservices, AI workloads, geopolitical tension, regulation, and industrial-scale cybercrime have transformed the technology estate into a core business risk.

Yet many operating models still rely on "artisanal IT": experts clicking through consoles, configuring systems manually, and demonstrating compliance with spreadsheets. That approach has reached its limit.

The answer is not another tool. It is a different way of running IT — the industrialization of intelligence. Work shifts from humans operating tools to platforms and automation executing at scale, while people define intent, rules, and guardrails.

Stability no longer derives from freezing change, but from systems that detect, correct, and improve continuously.

IT evolves from a ticket-driven service desk to a platform the business plugs into. On the surface, it looks simple: a secure internal portal delivering standard services in minutes.

Underneath, everything is codified and automated — infrastructure, security, compliance. AIOps detects and resolves issues early. Policy as code prevents non-compliant changes by design. Sovereignty and security are built in from the start. The key question is no longer "are we compliant?" but "how fast can we detect and fix when we are not?" — compliance velocity becomes a defining capability.

This is Smart IT: intelligence embedded in how systems are designed, operated, governed, and financed — combining visibility, automation, and integrated control of cost, value, and carbon.

In this world, the CIO no longer runs infrastructure. The CIO designs platforms, guardrails, and outcomes.

The autonomy journey, *not the autonomy leap*

Autonomy is a response to constraints, *not a moonshot*

Autonomous operations is showing up on CIO agendas because the constraints have tightened, not because the technology is fashionable. Geopolitical tension and sovereignty requirements push estates toward hybrid and multi cloud. Complexity keeps rising, while operating budgets and specialist capacity do not. In that environment, the traditional model of absorbing complexity with people and process stops scaling.



Not AI for AI's sake means *fewer incidents and faster recovery*

AI only earns a place in operations when it shifts outcomes the business can feel. Fewer user impacting incidents. Faster restoration when something breaks. Less recurring noise that drains teams. Better predictability of cost and risk. If AI does not change those metrics, it becomes another layer of complexity and another reason for teams to distrust the next initiative.



Autonomy is earned through *a maturity curve*

Autonomy is not a switch you flip. It is a progression from reacting to failures, to understanding them faster, to preventing them, and only then to automating action. The practical sequence is the point. Most organisations get value long before full autonomy, and the path itself creates the foundations needed for trust.

Observability is how you *keep control as complexity rises*

Observability matters because it turns scattered signals into evidence. It connects what happened, what changed, and what the impact was across distributed services and teams. Without that evidence chain, decision making becomes a mix of dashboards, opinions, and heroics. With it, operations becomes faster, calmer, and more accountable, because conclusions can be traced back to concrete signals and change history.



AI types in AIOps, and how they fit together

Observability provides the evidence and context. Deterministic AI turns that into trustworthy understanding. Agentic AI uses that understanding to take controlled action.

Pattern AI

Finds recurring patterns in telemetry, reduces alert noise, and groups related events so teams see fewer, clearer signals.

Causal AI

Connects the dots with discipline. It helps distinguish correlation from likely cause and clarifies blast radius and what changed.

Predictive AI

Looks forward. It forecasts capacity or saturation risk, change risk, and early signs of degradation so issues can be addressed before impact.

Generative AI

Speeds up understanding for humans. It summarises incidents, explains what happened in plain language, and helps navigate logs, traces, tickets, and runbooks.

Agentic AI

Executes. It can gather evidence across tools, propose next steps, run approved remediation workflows, and verify recovery through observability signals.



Business observability is the missing layer *when IT is the business*

IT is no longer supporting the business from the side. It is the channel through which the business runs. Business observability adds the layer that ties customer journeys, revenue flows, and critical business transactions to the underlying services and changes. It shifts prioritisation from technical severity alone to business impact, and it gives leadership a line of sight from operational decisions to business outcomes.

Proactive operations start *with reducing noise and speeding understanding*

The first maturity jump is not prevention and it is not autonomy. It is removing the friction from everyday response. Proactive operations reduces alert noise, correlates signals across components, and accelerates diagnosis so teams spend less time hunting and more time stabilising. This is where AIOps proves its value quickly, because it increases the team's effective capacity without adding headcount.

Preventive operations shift effort *from restoring service to avoiding impact*

The next maturity jump is looking forward. Preventive operations uses prediction to surface risk before users feel it, especially around change and capacity. Avoiding impact is almost always cheaper

than recovering from it. Preventive capability also builds trust internally, because it demonstrates that AI is being used to reduce risk, not add it.

Trustworthy answers come *from deterministic intelligence, not confident guesses*

Autonomy depends on the quality of the answers, not the sophistication of the interface. Causal, predictive, and generative capabilities can all contribute, but the requirement is the same. Outputs must be grounded

in evidence and understandable enough to challenge. Deterministic here does not mean perfect. It means repeatable from the data, explainable in context, and traceable back to signals, changes, and dependencies.

Gartner predicts that by 2028, at least

15%

of day to day work decisions will be made autonomously through agentic AI, up from 0% in 2024. ⁽¹⁾

In a Gartner survey of government CIOs outside the US,

49%

report current or planned deployment of AI agents in the next 12 months. ⁽²⁾

Deployment Timeline	Percentage
already deployed	12%
deploy in next 12 months	37%
deploy in 12-24 months	24%
deploy in 2-3 years	14%
not sure	11%

63%

of CIOs plan to spend on AI and machine learning in 2025. ⁽³⁾

Category	Percentage
AI/Machine Learning	63%
Data & Analytics	50%
Cybersecurity	43%
Business applications	39%
Cloud infrastructure	30%

In the State of the CIO 2025 survey,

42%

of CIOs said AI and ML is the technology initiative driving the most IT investment, and 45% said it is the most strategically important. ⁽⁴⁾

(1) Gartner Identifies the Top 10 Strategic Technology Trends for 2025 | Gartner
(2) Gartner Survey Reveals 52% of Government CIOs Outside of the U.S. Expect IT Budgets to Increase in 2026 | Gartner

(3) 2025 CIO Leadership Perspectives | Evanta
(4) State of the CIO, 2025: CIOs set the AI agenda | CIO

Q Client case

From reactive cost control to intelligent, sustainable IT operations in the transport sector

A large organization in the Norwegian transport sector operates a complex and business-critical cloud environment, supporting everything from core operations to analytics, integration, and digital services. Over time, increasing cloud complexity, fragmented ownership, and rising sustainability requirements have made it difficult to maintain control of costs, emissions, and operational efficiency. Cost management was largely reactive, based on periodic reviews and manual analysis, with limited linkage between financial, operational, and environmental objectives.

By introducing a FinOps-driven operating model combined with intelligent automation and GreenOps practices, the organization transitioned to proactive, platform-based governance of its cloud estate:

- Automated tagging and metadata inheritance enabled consistent cost and ownership allocation across subscriptions, services, and teams
- Policy-as-code ensured that cost, compliance, and sustainability rules were enforced automatically, not retroactively

- Continuous monitoring identified unused resources, missing Reserved Instance coverage, and optimization candidates without manual effort
- GreenOps visibility was established through dashboards and reporting that connected cloud consumption with CO₂ impact, enabling informed trade-offs between cost, performance, and sustainability



Monthly FinOps advisory sessions aligned finance, IT operations, and engineering around shared data, priorities, and measurable outcomes.

The result was not only reduced cloud spend, but a fundamental shift in how IT was operated. Cost and sustainability moved from being after-the-fact concerns to built-in constraints, managed continuously through automation and shared accountability. Manual effort was reduced, decision-making improved, and IT gained the ability to scale services while maintaining financial and environmental control.

In complex transport environments, lasting value comes from treating FinOps and GreenOps as continuous practices—supported by automation and visibility—rather than isolated cost- or sustainability initiatives.

Agentic AI becomes valuable *when it is bridged to those trustworthy answers*

Agentic AI should not be positioned as a replacement for engineering judgement. Its value is execution. The scalable pattern is to bridge trustworthy answers from deterministic intelligence to the actions of agentic AI, so the system can carry out work through tools and workflows once the situation is understood. That bridge is what turns AI from a helpful teammate into intelligent automation, and eventually into autonomous operations.



Autonomous operations is *closed loop, verified, and accountable*

Autonomy is only credible when actions can be verified and reversed. That means every automated step is tied to an evidence trail, executed within guardrails, and validated through observability after the action is taken. When the loop is designed this way, autonomy does not reduce control. It increases it, because decisions are faster and more consistent, and outcomes are continuously measured.

Autonomy does not need *to cover everything to change the economics*

The smarter framing is coverage, not perfection. You do not need autonomous handling of every incident to see disproportionate value. If a meaningful share of repeatable scenarios can be detected, understood, and resolved with verified actions, teams reclaim capacity, reliability improves, and the operational load drops.




The operating reality is *an ecosystem, not a single platform*

Autonomous operations does not live inside one tool. Context and actions span cloud platforms, observability tooling, IT service management, security systems, and business systems. A workable model is an ecosystem where evidence can be gathered across domains and actions can be executed across the systems that actually run the business. This is also where clear ownership and decision rights become non negotiable, because agents need to know what they are allowed to change.

Cost discipline is *part of trust, not a finance afterthought*

Observability and AIOps are data hungry, and unmanaged telemetry growth creates cost shock that undermines adoption. Sustainable autonomy requires observability economics. Clarity on what is collected, why it exists, how long it is retained, what decisions it enables, and who owns the bill. Cost transparency and control are what make the model scalable, because they prevent the organisation from choosing between visibility and affordability.





The *conviction*

Autonomous operations will be won by organisations that treat observability and AIOps as a proper field, not a tooling layer. This is a capability domain that needs to be designed, staffed, and governed in the same way as cloud and security. It should show up explicitly in architecture principles, delivery requirements, and sourcing choices, and it should be built into skills plans and leadership attention.

The winning pattern is not AI for AI's sake. It is bridging deterministic, evidence grounded answers with accountable agentic execution, so operations can move from reaction to prevention and, where it makes sense, to autonomy with control.

For CIOs, the action is clear. Make observability, business observability, and AIOps part of the core operating model agenda for the next 18 to 24 months. Define what outcomes must improve, which closed loops are worth automating first, and what data, ownership, and guardrails are required to trust the automation. This is also the point where an outside perspective is valuable, because the fastest route is rarely “buy a platform.” It is aligning operating model, telemetry strategy, and automation design around a small number of high impact loops.

Q Client case

Measured impact from FinOps-driven automation and continuous optimization

Across multiple customers in the transport sector, public sector, and financial services, the combination of continuous FinOps advisory,

automation, and standardized practices has delivered sustained and measurable improvements in cloud efficiency and governance.

Customer A – Global transport and logistics

- Total cloud cost reduction of ~33%
- Optimization driven by continuous rightsizing, commitment management, and automated identification of unused resources
- Shift from periodic cost reviews to ongoing, data-driven governance

Customer B – National transport infrastructure

- Total cloud cost reduction of ~24%
- Automated tagging coverage exceeding 93 % of services, enabling accurate cost allocation and ownership
- Improved budgeting, forecasting, and accountability across teams through standardized FinOps practices

Customer C – Governmental financial institution

- Total cloud cost reduction of ~21%
- Rightsizing initiatives reduced unused capacity from ~40% to under 10%
- Cost optimization embedded into daily operations rather than handled as ad-hoc exercises

Customer D – Municipal public sector organization

- Total cloud cost reduction of ~20%
- Orphaned resources are now identified and removed through routine, automated processes, replacing manual and reactive clean-up
- Increased predictability and reduced operational overhead for IT teams

Across all customers, results were achieved through continuous advisory, automation, and the establishment of FinOps as an operating model, not through one-time savings initiatives. Cost control, efficiency, and sustainability became repeatable outcomes rather than dependent on individual effort.

Key takeaway:

Sustained optimization at scale requires automation, shared ownership, and FinOps embedded into daily operations—enabling measurable improvements in efficiency, predictability, and control over time.

Three Forces Reshaping IT

01

IT Is Getting More Complex

The typical IT landscape has shifted from a few centralized systems to a **distributed web of environments**:

- Multiple public clouds, sometimes combined with private and sovereign clouds
- Edge sites in factories, stores, branches, and devices
- Hundreds or thousands of microservices and APIs

This complexity is no longer accidental. Leading organizations are now **deliberate** in how they place workloads:

- Certain workloads live in specific clouds for performance or specialized services
- Highly sensitive data is kept on sovereign or private infrastructure
- Real time processing moves closer to where data is generated at the edge

02

Security, Compliance, and Sovereignty Are Tougher

The environment around IT is becoming more demanding:

- Governments worry about **digital dependence** on foreign providers and laws
- Regulations such as **NIS2** and **DORA** demand faster detection and recovery from incidents, and clear proof of resilience across the extended supply chain
- Cyber attacks are more frequent, more automated, and more financially and politically motivated

In practice, this means:

- **Digital sovereignty** is a real design constraint. Sensitive data often has to stay under local legal control, with local keys and sometimes local operators.
- **Compliance is continuous.** Yearly audits are no longer enough; regulators expect that you know your current state and can show how quickly you correct deviations.
- **Identity is the new perimeter.** Every user, application, and AI agent must prove who they are and what they are allowed to do, every time.

03

Technology Enables Smarter Operations

The good news: the same technologies that increase complexity and risk also give IT the tools to operate much more effectively:

- **AIOps and agents** help sift through the flood of technical signals, spot problems early, connect events across systems, and increasingly carry out standard fixes (restarts, rollbacks, scaling, diagnostics, ticket updates).
- Instead of looking at “server costs,” organizations increasingly look at **cost per transaction, per incident, or per AI call.**
- Infrastructure and rules are changed through **code**, with AI assisting in writing, checking, and improving it—backed by automated security and compliance checks.

Eight Tensions Every CIO *Must Manage*

Delivering IT that is faster, cheaper, safer, greener, and smarter at the same time creates inherent tensions. These are not problems to “fix” once; they are **trade offs to manage continuously**. Below are eight such tensions, with guiding questions you can use with your teams.

1 – Speed vs. Safety

- The business demands rapid change, self-service, and experimentation.
- Boards, regulators, and customers expect strong security, reliability, and compliance.
- The same levers that increase speed – decentralization, self-service, automation – also expand attack surface and operational risk.

Leadership trade-offs:

- Where should speed and experimentation be deliberately encouraged – and where must strict approvals and controls remain non-negotiable?
- Which production actions can be safely automated end-to-end – and which must always retain human decision and accountability?

2 – Local Freedom vs. Central Control

- Product teams want autonomy over tools, models, and automation to optimize local outcomes.
- IT requires shared platforms, data models, and rules to keep the system secure and manageable.
- Too much freedom creates sprawl; too much control drives workarounds and shadow IT.

Operating principles:

- What should remain a local choice – and what must be standardized centrally (identity, logging, base AI models, security layers)?
- Which governance mechanisms – guardrails, approved catalogs, lightweight reviews – enable speed without fragmenting the environment?

3 – Cost Cutting vs. Future Innovation and Independence

- Budget pressure and FinOps drive rationalization and short-term savings.
- Differentiated AI and automation require sustained investment – often with powerful (and sometimes sticky) vendors.
- Cut too hard or centralize too aggressively, and you risk stagnation and lock-in.

Key decisions:

- Which AI and automation capabilities are truly strategic and must be built or strongly owned by us?
- Where should we invest now in dual options or experiments to preserve future flexibility and bargaining power?

4 – Human Accountability vs. Machine Autonomy

- AI can outperform humans in detecting patterns and proposing fixes.
- But when something goes wrong, accountability remains human.
- As automation expands, explainability becomes harder – just as regulators demand more transparency.

Leadership choices:

- For which decisions is full autonomy acceptable – and where must humans remain in the loop?
- Who is explicitly accountable when AI causes a material incident?

5 – Sovereignty vs. Innovation

- Regulators push for local legal control.
- Leading AI capabilities are often global.
- Over-index on sovereignty and you fall behind; over-index on innovation and you increase exposure.

Strategic trade-offs:

- Which data and workloads must remain under local control – and which can safely leverage global platforms?
- How do we design a hybrid architecture that protects sovereignty while capturing innovation?
- What is our explicit risk position on high-value workloads running on non-sovereign platforms?

6 – AI Efficiency vs. Energy and Emissions

- AI can optimize operations.
- But training and inference are energy-intensive.
- Efficiency per unit can still mean higher total energy use.

What to measure:

- Where does AI truly reduce emissions – and where does it simply add new demand?
- What design limits ensure AI supports – rather than undermines – sustainability goals?

7 – Seeing Everything vs. Drowning in Data

- Modern IT demands deep visibility.
- In practice, this creates data sprawl, tool overlap, and rising costs.
- More data does not equal more insight.

Focus areas:

- Which signals are truly critical – and which can we reduce?
- How do we align observability spend with real risk appetite and SLOs?

8 – Simplicity for Users vs. Hidden Skill Gaps

- Golden Paths simplify life for most teams.
- But complexity concentrates in the platform layer.
- Over time, skills may erode outside a small expert group.

Resilience questions:

- How do we keep critical skills alive in the wider IT organization as more is hidden behind simple interfaces?
- What is our resilience plan if key platform experts leave –can others understand, maintain, and evolve the platform?



A Call to Action for CIOs

The CIO role is not evolving at the margins – it is being redefined. Now is the time to act decisively.

1. Get the data foundation right.

Consolidate logging, metrics, tracing into a coherent observability stack. Standardize a set of critical signals for resilience, security, cost, and compliance and ensure data coverage for core applications, clouds, edge, and identity.

2. Build platforms, not barriers.

Make the secure and compliant path the easiest one. Measure success by developer experience, adoption, reduced shadow IT, and time to value – not ticket volume.

3. Turn policy into code.

Embed security, compliance, and data rules directly into pipelines and runtime. Make compliance velocity – how fast you detect and remediate drift – a core KPI.

4. Architect for sovereignty and optionality.

Design hybrid estates where sensitive data is routed by default to sovereign environments, while global platforms power scale and innovation. Use privacy-preserving techniques to unlock AI without crossing legal boundaries.

5. Pilot AIOps on real pain.

Choose a few critical services and feed AIOps with your observability data. Focus on alert noise reduction, incident correlation, automated diagnostics, and a few safe auto-fixes, and then measure impact on detection, resolve and engineer time.

6. Orchestrate an agentic workforce.

Define clearly what can be automated – and where humans remain accountable. Ensure every AI agent operates with strong identity, least privilege, and full auditability.

7. Lead on outcomes – and carbon.

Shift the executive conversation to cost and carbon per business outcome. Use that transparency to scale what creates value, redesign what underperforms, and stop what does not justify its footprint.



Conclusion

Crossing the Line

The core message is straightforward: the traditional way of running IT is no longer fit for purpose. Ticket queues, manual console work, and spreadsheet driven oversight are liabilities in a world of multi cloud, AI, and aggressive regulation.

The alternative is to **industrialize intelligence**:

- Platforms instead of one off projects
- Executable policies instead of static documents
- AI and automation doing the routine work, with people setting direction and making high impact decisions
- Financial and environmental impact managed at the level of business outcomes, not infrastructure units

A simple test:

- Does launching a new service still take weeks of coordination and tickets?
- Is compliance something you “prepare for” once or twice a year?
- Are AI efforts mostly pilots and proofs of concept without a clear cost per outcome?

If so, you are still in the world of artisanal IT.

The mandate now is to cross the line: **build the platform, codify the rules, wire in AIOps, expose the economics, and accept that the CIO is no longer the Chief Infrastructure Officer, but the Chief Intelligence Orchestrator.**

Those who make this shift will not just keep up with the digital race; they will set the pace. Those who don’t will find that the world optimizes around them.

Contacts



Tor NESET
Chief Technology Officer, Norway
→ tor.neset@soprasteria.com



Mads STORRUSTE
Agency Director Observability & AIOps, Norway
→ mads.storruste@soprasteria.com

Contributors of our European Network



Olivier BROUSSEAU
Architecture & Tech
→ olivier.brousseau@soprasterianext.com

Isabelle PONS
Data & AI
→ isabelle.pons@soprasterianext.com

Valérie THIBLET
Data & AI
→ valerie.thiblet@soprasterianext.com

Axelle ROESCH
CIO Advisory
→ axelle.roesch@soprasterianext.com



Thomas OTTO
CIO Advisory
→ thomas.otto@soprasteria.com

Philipp MAIER
CIO Advisory
→ philipp.maier@soprasteria.com



Simon KAYE
Architecture & Tech
→ simon.kaye@soprasteria.com



Ray BAKER
Architecture & Tech
→ ray.baker@soprasteria.com

Andy WHITEHURST
Architecture & Tech
→ andy.whitehurst@soprasteria.com



Eirik HASLESTAD
CIO Advisory
→ eirik.haslestad@soprasteria.com

Jahn-Fredrik SJOVIK
CIO Advisory
→ jahn-fredrik.sjovik@soprasteria.com

Tor NESET
Architecture & Tech
→ tor.neset@soprasteria.com

Marius SANDBU
Data & AI
→ marius.sandbu@soprasteria.com



Robert HEXSPOOR
Architecture & Tech
→ robert.hexspoor@soprasteria.com

Hugo HERMSEN
CIO Advisory
→ hugo.hermesen@soprasteria.com



Juan Antonio MARTINEZ SANCHEZ
Architecture & Tech
→ juanantonio.martinezsanchez@soprasterianext.com



Mauro PALMARINI
CIO Advisory
→ mauro.palmarini@soprasterianext.com