perennial

What's Slowing Down Your IT Operations? Measure It.

A Strategic 5-Dimension AIOps Maturity Scorecard for BFSI Organizations

Evaluate your AIOps maturity across five dimensions that directly impact MTTR, regulatory posture, and resilience.

Perennial Systems · 2026

perennial

Before You Begin

This assessment is built for results.

It comes from 8 years of embedded AIOps engineering inside APAC BFSI institutions - building the operating models, correlation logic, remediation workflows, and governance frameworks that move MTTR and SLO compliance in regulated environments.

Choosing where to invest in operational maturity is hard. Especially when every vendor deck looks the same and every platform promises the same outcomes.

We have been in the room when the numbers did not move after procurement. We have seen it all across APAC.

This framework will help you separate what is actually working in your operation from what only appears to be working, so you can prioritize with precision and defend that decision to your board.

If you have a question at any point, write to *info@perennialsys.com*

About Perennial Systems

Perennial Systems is an engineering-led AI transformation company for regulated industries, specialising in AIOps, AI-driven SDLC, Agentic Banking, and AI Process Engineering across APAC BFSI.

We have delivered observability and AIOps programmes across Tier-1 digital banks and payment platforms in Southeast Asia, with teams embedded directly inside client operations, not working from a separate delivery centre.

To see how this assessment applies to your operation, take the scorecard at *perennialsys.com/aiops-assessment* or write to *info@perennialsys.com*

perennial

Table of Contents

perennial

About This Assessment

A diagnostic framework for APAC BFSI operations leaders, not a vendor comparison, not a platform guide.

What you will leave with:

→ A five-dimension maturity model that locates where your AIOps operation actually performs under pressure.

→ A web-based scorecard generating your maturity profile, role-specific 90-day priorities, and peer benchmarks. Quick Read: 3 min. Full Diagnostic: 8 min.

→ Level-by-level breakdowns of what each maturity stage costs, where it breaks, and where the highest-ROI intervention sits.

Built for: CTOs, CDOs, Heads of IT Operations, and SRE leadership at APAC banks operating under MAS TRM, BSP Circular 1140, and BNM RMiT mandates.

Document Navigation Map

perennial

Where Does Your Operation Stand Today?

Three conditions define most APAC BFSI operations teams. Identify where you are.

You have not deployed AIOps tooling yet.

Incidents surface through customer complaints or manual checks. Regulatory reporting is assembled by hand. Every outage resolution depends on who is available and what they know. High-impact outages cost $1.8M per hour on average, and APAC regulators now require demonstrable resilience, not just recovery documentation. The right time to build the operating model is before the next major incident.

You have the tooling but MTTR has not improved.

The platforms are licensed and running. The outcomes have not followed. 29% of financial services organisations still experience high-impact outages weekly. Teams absorb 2,000+ alerts per week; 3% require action. 73% of organisations have had outages from alerts that were dismissed. The gap is not in the platform, it is in the operating model around it. This assessment locates that gap precisely.

Your business is growing and your operations team & cost is rising with it.

New markets, products, and regulatory touchpoints are being added. Manual operations that worked at 300 people break at 700. The alert volume your team managed last year is no longer manageable the same way. At this stage, AIOps is not an efficient play, it is the operating model that lets your team scale without scaling incidents at the same rate.

What This Document Will Do for You

This document helps you figure out where your operation actually stands and how far that is from what your regulator, board, and customers expect.

You won’t get a single maturity score. No bank is mature everywhere. You’re probably strong on observability but weak on governance. Good at detection, messy at response.

The value is seeing where the gaps are.

That’s where your long recovery times come from. Where regulatory reports fail. Where customers leave. And where your best improvement opportunities are.

perennial

The Structural Gap Behind Stalled MTTR

Most APAC banks carry a two-level maturity gap: platforms capable of Level 3-4 performance delivering Level 1-2 operational outcomes.

Tooling is a procurement decision. Maturity is an operating model outcome.

The shift to intelligence-driven operations does not happen at platform deployment. It happens when the processes, governance, and feedback loops around that platform reach parity with its actual capability. A 2023 Forrester study found 63% of financial institutions faced major disruption not from technical failure, but from delayed detection and misaligned triage. Only 18% consistently met their own internal SLAs.

That gap is where unexplained MTTR accumulates. Where regulatory reporting scrambles originate. Where customer attrition compounds. And where the highest-return operational improvements sit.

The Two-Level Gap

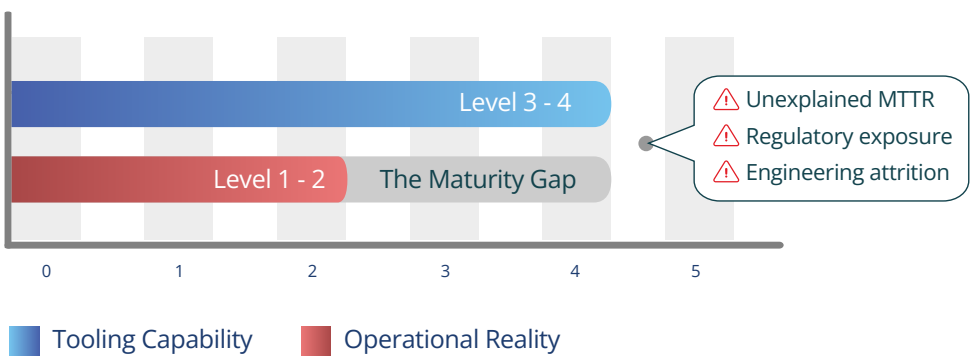

perennial

Four Reasons This Gap Is a Board-Level Priority

→ Companies are now expected to prove they can handle operational disruptions.

MAS, BSP, and BNM mandates require demonstrable operational resilience, not tooling inventory. Automated remediations without auditable decision trails are a regulatory finding regardless of platform maturity.

→ The cost of system downtime keeps increasing every year.

$1.8M per hour for high-impact outages. $152M per organisation annually. Both figures grow as digital transaction volumes rise.

→ A significant amount of engineering time is being spent on repetitive operational work.

Incident management toil reached 30% of engineering time in 2025, the first increase in five years. The roles most exposed to attrition are the same ones operational resilience depends on.

→ Operational problems continue to grow when there is no structured governance.

Weak governance means post-mortems do not update correlation rules. Stale runbooks produce repeated incidents. Without a structured operating model programme, the distance between platform capability and operational outcome grows each quarter.

What follows is the framework for measuring exactly where your operation stands. This helps you identify where the gap between your highest and lowest dimension is concentrated.

perennial

The Maturity Model: Five Levels, Five Dimensions

The BFSI AIOps Maturity Model describes the progression of your IT operations function from incident-driven to intelligence-driven. Five levels of operating behaviour. Five dimensions of capability.

The levels describe what your operation does under pressure. The dimensions describe the lenses through which that behaviour becomes visible and measurable.

Think of it this way: the levels answer what happens in your operations centre at 3am during an unplanned incident? The dimensions tell you why.

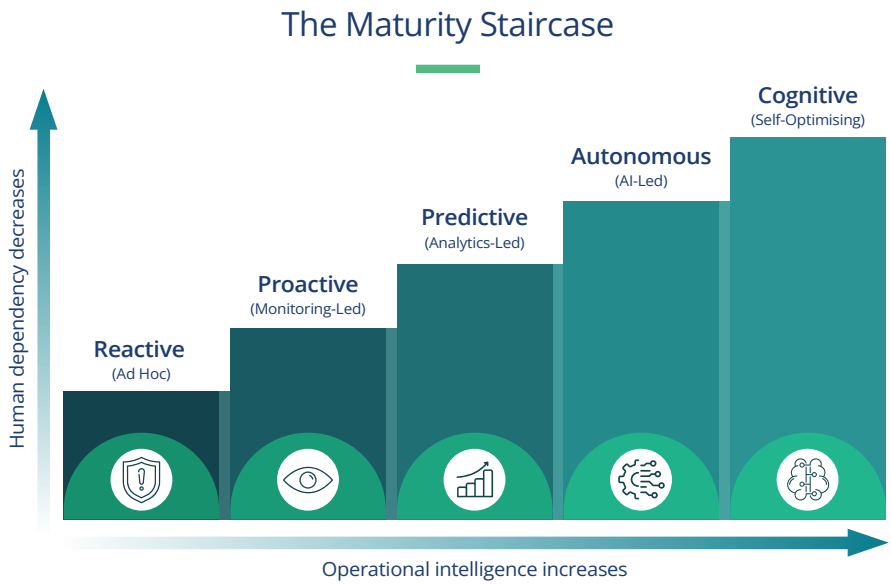

The Five Levels

Level 1, Reactive (Ad Hoc)

The phone rings before the dashboard does. Customers report failures before your NOC detects them. Incident response depends on tribal knowledge held by two or three senior engineers. At this level, engineering teams typically face over 2,000 alerts per week with very little correlation.

perennial

Level 2, Proactive (Monitoring-Led). Monitoring is unified and most incidents are detected before customers notice them. However, once the alert fires, everything that follows remains fully manual. Resolution speed still depends entirely on who is on call. According to New Relic's 2026 Observability Forecast, 29% of financial services organisations experience high-impact outages at least weekly at this maturity.

Level 3, Predictive (Analytics-Led). Event correlation and analytics identify emerging incidents before they breach SLOs. Runbooks are partially automated, but human judgement remains the final decision layer for every action. Most APAC Tier-1 and Tier-2 banks believe they operate here. In reality, under pressure, many still function at Level 2. Splunk research shows 73% of organisations have experienced outages caused by ignored alerts, highlighting the gap between tools and outcomes.

Level 4, Autonomous (AI-Led). Detection, diagnosis, and common remediations run in closed-loop with AI. Humans focus only on edge cases, governance, and exceptions, not recurring L1/L2 incidents. Audit trails are generated automatically. At this level, incident management toil drops significantly from the current industry average of 30% of engineering time.

Level 5, Cognitive (Self-Optimising). The system learns from every incident, autonomously tunes its own thresholds, correlation rules, and even suggests infrastructure optimisations. Operations shift from a response function to a continuous feedback loop.

The Five Dimensions You Will Score

In the scorecard, you will assess your operation on each of the following dimensions. For each one, select the highest level you can honestly claim, based on what your team does during an unplanned incident at 2am, not what the runbook says they should do.

1. Data & Observability, What telemetry you capture, how it is normalised, and whether it covers the full stack. At Level 1, monitoring tools exist but are siloed, with no single view across infrastructure, application, and business transaction data. At Level 3, any engineer can query across all three in a single tool. At Level 5, the system monitors its own telemetry health and adapts granularity based on system state. IDC found that nearly 90% of Asia/Pacific enterprises run workloads across multiple public clouds. The challenge is no longer volume, it is coherence.

perennial

2. Event Intelligence, How your alerts are correlated, de-duplicated, and prioritised. At Level 1, static thresholds set at deployment, never recalibrated, 2,000+ alerts per week, no correlation, engineers mentally connecting related alerts across separate systems. At Level 3, ML-driven anomaly detection cuts noise 40-60%. At Level 5, the system proposes new correlation rules from outcome data without engineering intervention. Picus Security's 2025 Blue Report found BFSI achieved a 67% log score but only 13% alert score; 7 of 8 events captured never generated a meaningful alert.

3. Response & Remediation, How incidents move from detection to resolution. At Level 1, incident response is person-dependent, resolution time varies 3-5x depending on who is on call, runbooks don't exist or aren't trusted. At Level 3, the system recommends a runbook, and the engineer triggers execution, but nothing runs without human approval. At Level 4, closed-loop remediation operates for validated patterns, MTTR for known incidents: 2-8 minutes. The honest test: if your three senior SREs were unavailable tonight, would resolution time change? If yes, that dimension is Level 2 or below. Toil rose to 30% of engineering time in 2025.

4. Predictive & Preventive Capability, Whether your system anticipates failures before they manifest, and whether predictions are acted on. At Level 1, no prediction; incidents detected post-impact. At Level 3, predictive alerts are generated but many go unacted on, the team doesn't fully trust model accuracy, or there's no governance for who should act at 3am. At Level 4, low-risk preventive actions are automated. A predictive alert sitting unread in a dashboard is not a Level 3 capability. It is a Level 1 capability with a Level 3 price tag.

5. Governance & Compliance, How incidents, decisions, and automated actions are logged, audited, and surfaced to your regulator. At Level 1, no audit trail for remediation actions, regulatory reporting (MAS TRM, BSP 1140, BNM RMiT) assembled manually under deadline pressure. At Level 3, standard reviews exist but feedback loops are weak, post-mortem findings may or may not update systems within a defined timeframe. At Level 4, full audit trails for every automated action: what was detected, what decision logic applied, what action taken, what outcome. Automated remediation without an auditable decision trail is not an efficiency gain. It is a regulatory finding waiting to happen.

perennial

One Reframe Before You Begin

The goal is not to reach Level 5.

The goal is to close the maturity gap, the distance between your highest and lowest dimension. For most APAC banks, that gap is two full levels wide. It is where every stalled MTTR, every missed SLA, and every quietly churned customer is hiding.

Find the gap. That is what the next pages are for.

perennial

Your Target-State AIOps Stack

Before you score your operation, it helps to see the target-state architecture, the structural picture of what a mature AIOps operation looks like when all five dimensions are working together. This is not a product diagram. It is an operating architecture.

The stack is organised into four operating layers. Data flows upward. Decisions flow downward. Your maturity is determined by how far up the stack your data travels before a human has to intervene, and how far down the stack your decisions travel before they require manual execution.

The BFSI AIOps Target-State Architecture

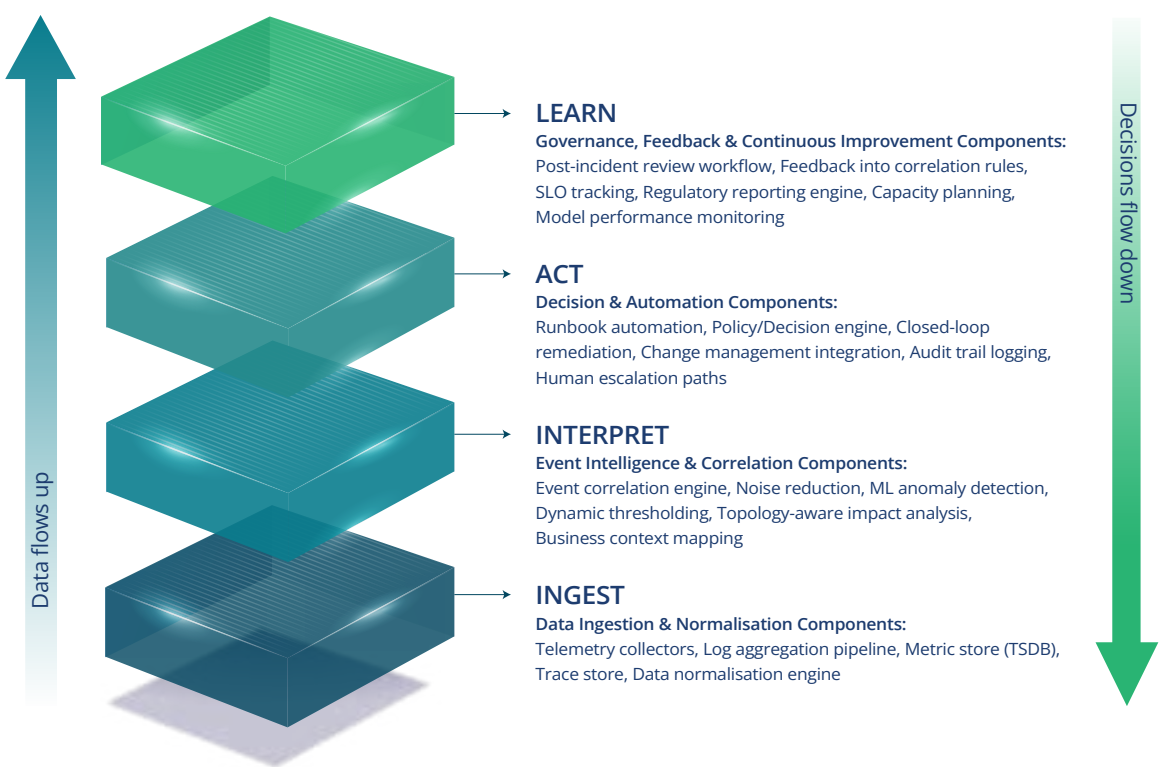

perennial

INGEST, Data Ingestion & Normalisation

The foundation. Collects telemetry from every surface in your estate: infrastructure metrics, application performance data, log streams, network flow data, and, critically for BFSI, business transaction telemetry. Payment success rates. Settlement cycle times. Core banking batch job completion windows.

The challenge at this layer is not volume. Most APAC banks generate terabytes daily. The challenge is normalisation, bringing data from a mainframe core, a private cloud integration layer, three public cloud providers, and fifteen SaaS dependencies into a single queryable format with consistent timestamps, naming conventions, and severity taxonomies.

INTERPRET, Event Intelligence & Correlation

Raw telemetry becomes operational intelligence here. Alerts generated, correlated, de-duplicated, prioritised. This is where most banks discover the gap: heavy investment in INGEST but insufficient investment in correlation logic, noise reduction, and dynamic thresholding.

A useful analogy from banking operations: INGEST tells you a transaction failed. INTERPRET tells you *why it failed, which other transactions are affected, and whether this is an isolated event or the leading edge of a cascade.*

ACT, Decision & Automation

Intelligence becomes action. For BFSI, this layer carries a constraint most AIOps frameworks understate: **every automated action must be auditable.** Under MAS TRM, BSP Circular 1140, and BNM RMiT, your bank cannot auto-remediate an incident without a decision trail a regulator can reconstruct. The automation must log not just what it did, but *why it decided to do it, what alternatives it evaluated, and what the expected outcome was.*

LEARN, Governance, Feedback & Continuous Improvement

The layer most banks do not have, and the one that separates Level 3 from Level 4. Post-mortems feed back into correlation rules. False positive patterns are identified and suppressed. SLO performance tracked against business outcomes, not just infrastructure metrics.

For BFSI, this layer also houses your regulatory reporting interface: surfacing operational resilience data in formats that satisfy MAS, BSP, and BNM requirements without requiring your team to manually assemble data from five different tools the week before a submission deadline.

perennial

How the layers map to maturity

Architecture × Maturity Matrix

	Level 1-2	Level 3	Level 4 - 5
INGEST	Fragmented; tool-specific silos	Partially unified; gaps in business telemetry	Single observability fabric across all layers
INTERPRET	Static thresholds; human correlation	ML anomaly detection; partial noise reduction	Real-time contextual correlation; business-impact prioritisation
ACT	Manual; tribal runbooks	Scripted runbooks; human-triggered	Closed-loop automation with auditable decision trails
LEARN	Non-existent or ad hoc	Post-mortems happen; learnings not fed back	Continuous feedback; regulatory reporting as by-product

Score each layer of your stack: **Take the Scorecard →**

Level 1, Reactive: When the Customer Knows Before You Do

Level 1 Incident Flow

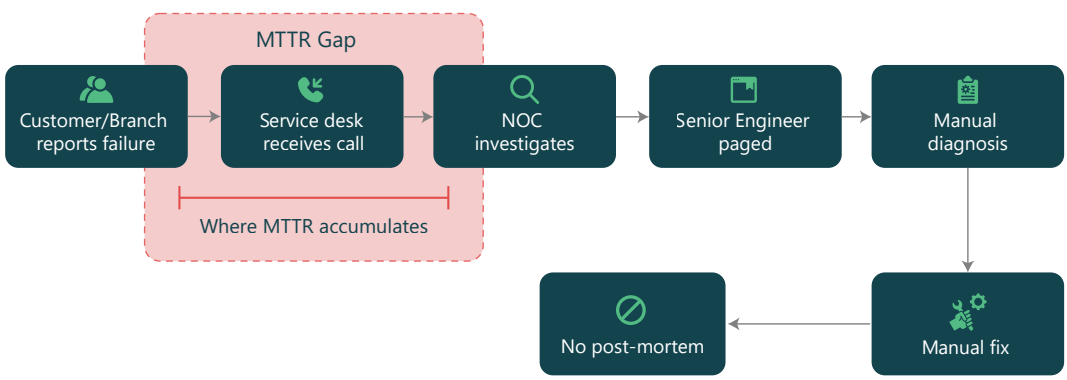

The phone rings before the dashboard does.

A branch manager calls because the teller system is unresponsive. A payments analyst notices a batch job running 40 minutes behind. A customer tweets that mobile banking is returning errors. Your NOC learns about the incident from someone outside the operations team.

What Level 1 looks like:

- Monitoring siloed-no single view across infrastructure, application, and business transaction data. Engineers trace issues manually across separate tools.
- Static thresholds set during deployment, never recalibrated. 2,000+ alerts per week, no correlation, no impact mapping.
- Incident response person-dependent. Resolution time varies 3-5x depending on who is on call. Runbooks don't exist or aren't trusted under pressure. Same incidents recur monthly.
- No prediction capability. Incidents detected post-impact.
- No audit trail for remediation actions. Regulatory reporting assembled manually under deadline pressure.

The honest test:

What happens at 2am on a Saturday when your two best SREs are on leave, the core banking batch has failed silently, and the Manila branch discovers it at 7am Monday? If the answer involves frantic calls, manual log searches, and finding the one person who fixed it last time, that dimension is Level 1.

At $1.8 million per hour for high-impact outages, Level 1 operations compound cost and regulatory exposure simultaneously. An operation without auditable decision trails cannot satisfy post-incident reporting requirements under MAS TRM, BSP Circular 1140, or BNM RMiT.

perennial

Level 2, Proactive: Detecting the Problem, Still Fixing It by Hand

Level 2 Incident Flow

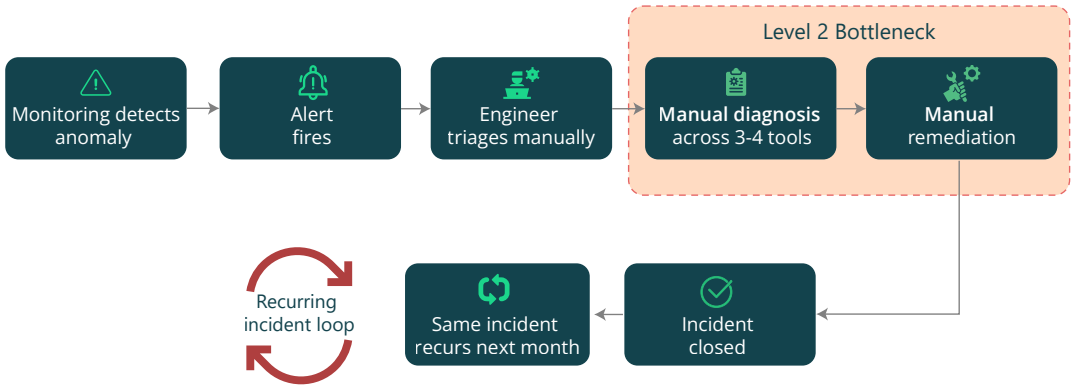

The dashboard fires before the customer calls. Monitoring is in place. Your NOC sees most incidents before the branch does.

But once the alert fires, **everything that follows is still manual.** A human reads the alert, searches for context across three or four tools, runs diagnostic commands, and applies a fix. If that human is not one of the senior engineers who has done this before, the sequence takes two to three times longer.

Level 2 has solved detection. It has not been responded to.

What Level 2 looks like:

- Centralized monitoring platform in place, but querying across application and infrastructure logs still requires manual effort. Business transaction telemetry partially captured but disconnected.
- Alerts have severity levels and basic noise reduction, but correlation between related alerts still happens in someone's head, not the platform.
- Runbooks exist for common incidents but are inconsistently followed. Same L1/L2 incidents-disk full, certificate expiry, batch timeout-recur monthly and consume 30-40% of on-call time.
- Post-mortems happen for major incidents, but findings documented and filed rather than fed back into process.

perennial

The Level 2 Trap

Level 2 is the most common maturity level across APAC BFSI-and the most dangerous to stay at. Not because things break catastrophically, but because t**he operation feels good enough.** Detection works. Alerts fire. The quarterly board report says “uptime 99.8%.”

Underneath: 30%+ of engineering time spent on repetitive manual toil. Same incidents recur because root causes are patched but not resolved. Alert fatigue compounding-over 2,000 alerts per week, 97% requiring no action. **73% of organizations experienced outages directly caused by ignored alerts.**

Level 2 is where most banks conclude they need better tooling. In most cases, what they actually need is a better operating model around the tooling they already have.

Stuck in the Level 2 trap? **Take the Scorecard →**

The Level 2 Cost of Inaction

$152M	Average annual cost of downtime for financial services organisations (Splunk)
30%	Engineering time spent on manual toil, first increase in five years (Runframe 2026)
73%	Organisations experienced outages from ignored alerts (Splunk 2025)

perennial

Level 3, Predictive: The System Sees It Coming. Your Team Still Decides.

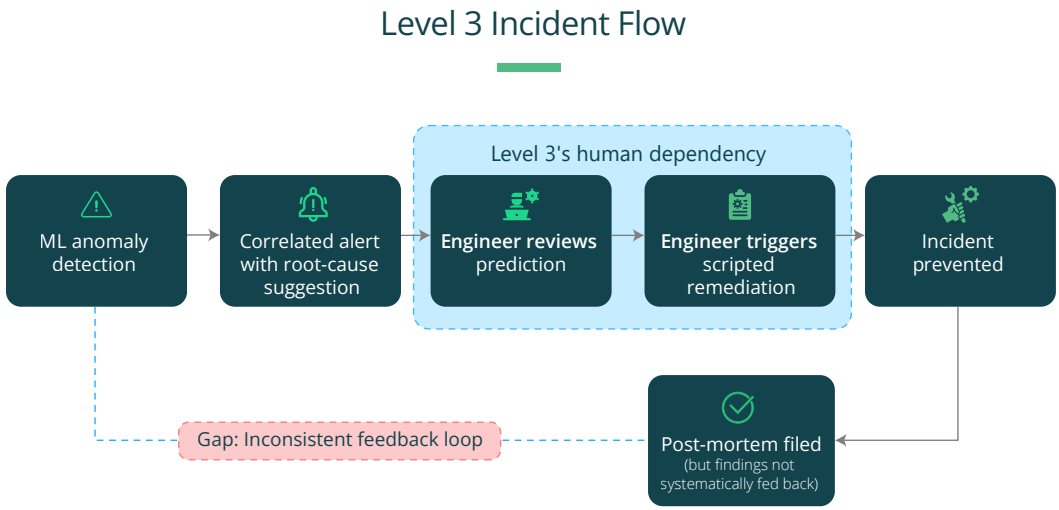

The system tells you something is going wrong before it breaks. Anomaly detection flags correlated latency across the payment gateway cluster-a pattern the team saw three weeks ago, ten hours before a settlement failure. An engineer reviews the alert, confirms the risk, triggers a pre-built remediation script. Incident prevented-**when the right engineer is available, trusts the prediction, and knows which script to run.**

What Level 3 looks like:

- Full-stack observability with normalized data. Any engineer can query across infrastructure, application, and business transaction data in a single tool.
- ML-driven anomaly detection. Alert noise reduced 40-60%. Cross-layer correlation works for most known incident patterns.
- Structured runbooks and semi-automated remediation-but no remediation runs without human approval.
- Predictive alerts generated. Many go unacted on-team doesn't fully trust accuracy, or no governance exists for who should act at 3am.
- Feedback loops exist but are weak and inconsistent.

perennial

Where Level 3 Actually Breaks

This is the level most banks claim. It is also where claiming and operating diverge under pressure.

The test: Your system generates a prediction at 3am. Your senior engineer is on leave. Does the night-shift L1 analyst trust the prediction enough to act? Have the authority to trigger the script?

If the answer is no, the operation falls back to Level 2 behavior the moment the A-team is unavailable.

This is the Level 3 fault line:

The system is capable of prediction and semi-automation, but the operating model still depends on a small number of senior humans to translate intelligence into action. An operation that performs like Level 3 during business hours and reverts to Level 2 outside them. This pattern is directly responsible for the MTTR variance that appears in quarterly reports but never gets explained.

The Level 3 Fault Line

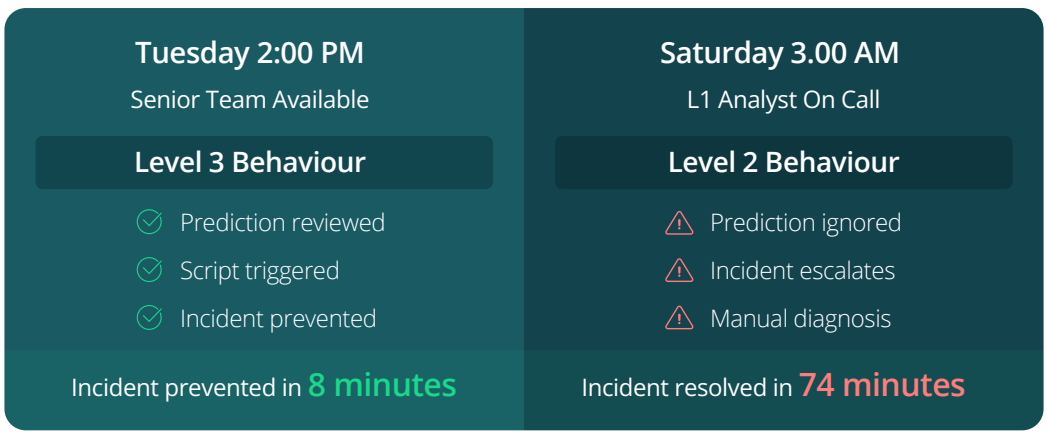

Same tooling. Same alerts. Same platform. **Different operating maturity based on who is available.**

The Bridge to Level 4

Moving from Level 3 to Level 4 is three **operating model decisions:**

1. **Define the trust boundary.** Which automated actions can the system execute without human approval? Start with lowest-risk, highest-frequency incidents-disk cleanup, service restart, certificate rotation.

2. **Build the audit trail into automation.** Every automated action must log what was detected, what decision logic was applied, what action was taken, what the outcome was. MAS, BSP, and BNM require it. It is also what gives your risk committee the confidence to expand the trust boundary.

3. **Close the feedback loop structurally.** Post-mortem findings feed directly into correlation rules and runbook updates through a defined process with an owner, a timeline, and a verification step.

These three decisions separate Level 3 from Level 4. **Operating model decisions. Not platform decisions.**

Level 4, Autonomous: The System Handles the Known. You Handle the Unknown.

Level 4 Closed-Loop Remediation

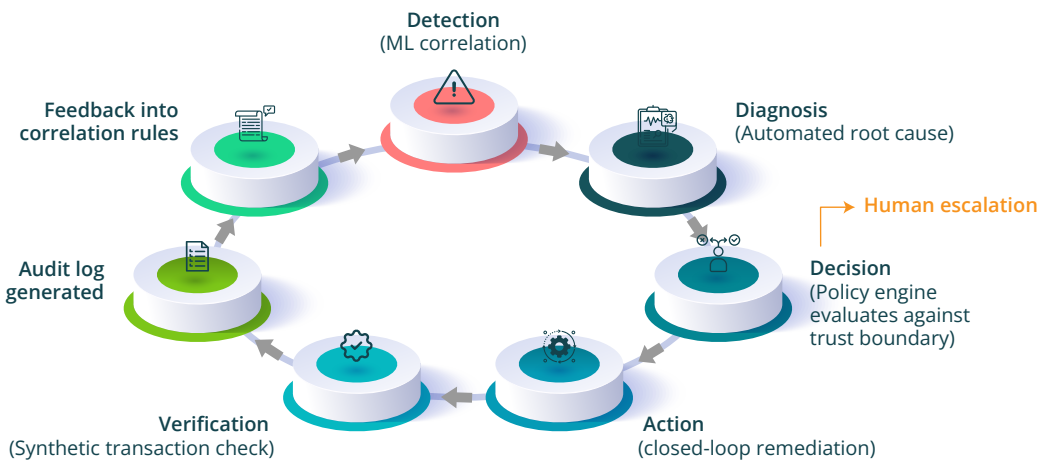

At 2:47am, the AIOps platform detects correlated latency across the payment gateway cluster. The system drains traffic from affected nodes, restarts degraded service instances, verifies recovery, and restores traffic. **Four minutes.** An audit log captures detection signal, decision logic, action, verification, and business impact: zero customer-facing transactions affected. The on-call engineer reviews the log at 8am. No action required.

This is Level 4. The system handles the known. You handle the unknown.

What Level 4 looks like:

- Real-time observability fabric with business context enrichment. "Node 7 latency spike" becomes "Node 7 latency spike on the RTGS cluster, affecting 340 pending settlements worth $12.7M."
- Business-impact-weighted correlation. Alert noise reduction 80%+ vs Level 2. Dynamic thresholds adapt to traffic patterns without manual recalibration.
- **Closed-loop remediation for validated patterns:** detect, diagnose, act, verify, log-without human intervention. MTTR for known incidents: 2-8 minutes.
- Full audit trails for every automated action. Regulatory reporting automated. Post-mortem learnings structurally improve correlation rules through a defined, time-bound feedback process.

What Level 4 Requires That Level 3 Does Not

The technology gap is real but manageable. Most AIOps platforms support closed-loop remediation. The real barrier is organizational:

Institutional trust in automation. Your CTO, Head of IT Risk, and risk committee must agree that defined categories of incident response can proceed without a human in the loop.

Regulatory confidence in the audit trail. MAS, BSP, and BNM require any automated action to be explainable, auditable, and reversible.

A feedback culture, not just a feedback system. Every incident feeds back into decision logic within a defined timeframe. Correlation rules and remediation scripts require the same rigor as production code: versioned, tested, reviewed, deployed through CI/CD.

Few banks operate at Level 4 across all five dimensions. The common pattern is Level 4 on one or two-typically Data & Observability and Event Intelligence-and Level 2-3 on the rest. That unevenness is itself a diagnostic finding.

Where Level 4 Sits in APAC BFSI Today

Few banks operate at Level 4 across all five dimensions. The common pattern is Level 4 on one or two (typically Data & Observability and Event Intelligence) and Level 2-3 on the rest (typically Response & Remediation and Governance). That unevenness is itself a diagnostic finding.

perennial

Measure the gap between your strongest and weakest dimension:
Take the Scorecard →

Level 5, Cognitive: The System Improves Itself

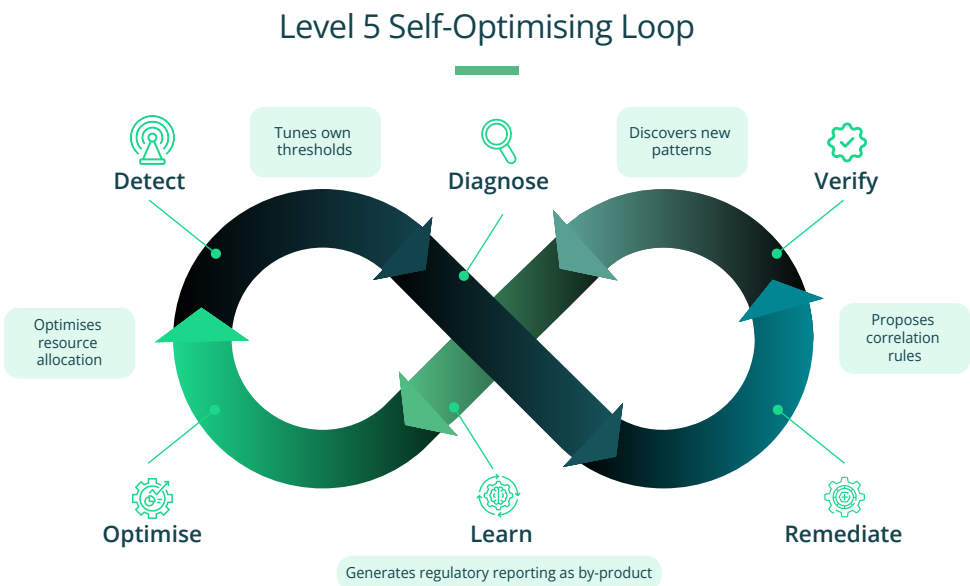

At Level 5, the operation is self-improving. It learns from its own behavior, optimizes its own configuration, and tunes the underlying infrastructure to prevent the conditions that create incidents in the first place.

What Level 5 looks like:

- **Self-tuning observability.** Monitors its own telemetry health, adapts granularity based on system state, auto-correlates new data sources.
- **Self-learning correlation.** Identifies recurring event sequences, proposes new rules, updates impact scoring continuously.
- **Outcome-driven remediation.** Selects fixes based on historical outcome data, not predefined scripts. Self-healing extends to proactive resource rebalancing.
- Continuous risk prediction across service dependencies-cascading failure probability, not just individual

component failure.

- Real-time compliance posture monitoring. Decision traceability on demand. Compliance report is a by-product of operations, not a separate workstream.

The Level 5 Reality Check

Level 5 is not a state you purchase. It is a state you evolve into over years of disciplined feedback-loop execution. **No AIOps vendor delivers Level 5 out of the box.** Any vendor that claims to be selling an aspiration, not a deployment.

What matters is that every dimension is moving forward and that **the gap between your highest and lowest dimension is narrowing.** An operation that is Level 4 on Event Intelligence and Level 1 on Governance is more fragile than one that is Level 3 across the board.

Uniform maturity is more resilient than peak maturity with gaps.

Maturity Radar (Example Comparison)

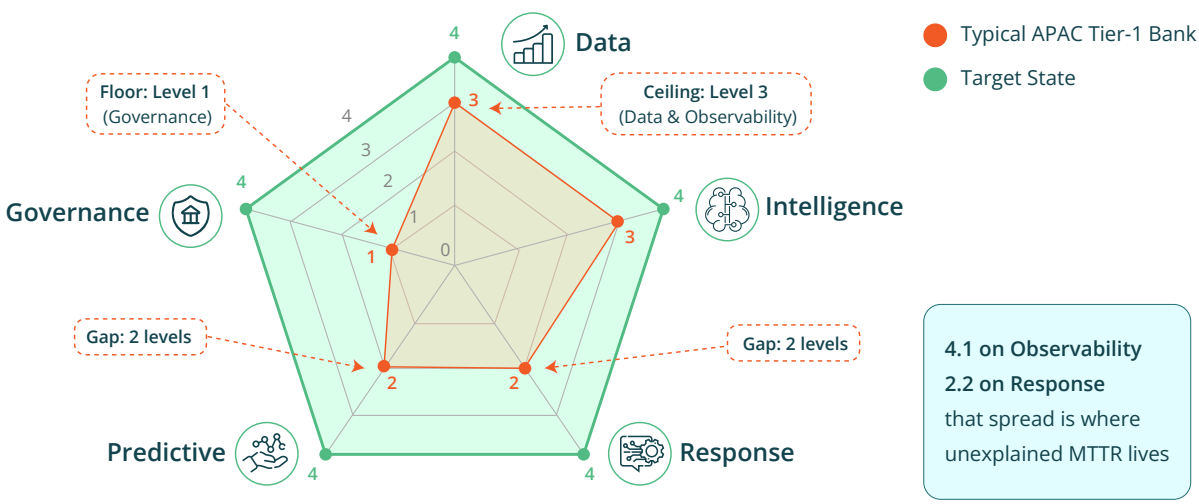

perennial

Score Your Operation: The Digital Maturity Scorecard

You have read the framework. You have seen what each level looks like, and where each one breaks. Now quantify it.

The diagnostic is honest, some answers will surface gaps. That's the point. Your responses stay in your browser until you choose to share them. No login required. No email until you opt in.

Choose your path:

- **Quick Read (3 min, 5 questions):** Directional score, one question per dimension
- **Full Diagnostic (8 min, 15 questions):** Detailed score with role-specific 90-day priority sprint

Select your role upfront so the language and 90-day pathway calibrate to your seat: CTO/CIO/CDO, Head of IT Operations, SRE/Platform Lead, Risk/Compliance, or Enterprise Architect.

You will receive:

Your Maturity Profile, a weighted overall score (e.g., "3.04 / 5.0, Defined") plus dimension-by-dimension breakdown. Example: "You scored 4.1 on Observability and 2.2 on Response. That spread is where your unexplained MTTR lives." Displayed as a pentagon radar chart that makes unevenness immediately visible.

Your Maturity Gap, the distance between your highest and lowest dimension. Gap of 0-1 = uniform. Gap of 2+ = structural imbalance, the most common APAC BFSI pattern. This is where your unexplained MTTR and regulatory exposure live.

Your Highest-ROI Intervention, your weakest dimension named and explained, with the specific operational and regulatory risk it carries for BFSI.

Your 90-Day Priority Sprint, not generic advice. A role-specific pathway broken into three phases (Baseline → Build → Prove) with what each phase looks like from your seat, plus a cross-role RACI matrix showing who owns what across the 90 days.

Peer Benchmark, where your score places you relative to APAC BFSI peers and global financial services observability data.

Complete it with your operations and SRE leads. The presence of a tool does not move you up a level. The consistent use of that tool under pressure does.

perennial

ASSESS YOUR OPERATION: **Take the Scorecard Now** →

Post-results: A comprehensive report summary will be emailed to you.
To help you better understand the insights and next steps, you can also book a free 30-minute diagnostic call.

Defined

3.85 / 5.0 Predictive

L3 : Weighted Overall

Standardized. *Now the trust boundary work begins.*

You scored 3.85. Core processes are standardized; Tier 1 systems are monitored consistently. The next move closed-loop automation within a defined trust boundary is where most banks at your level stall because it's a governance problem disguised as a technology problem.

Maturity Across 5 Dimensions

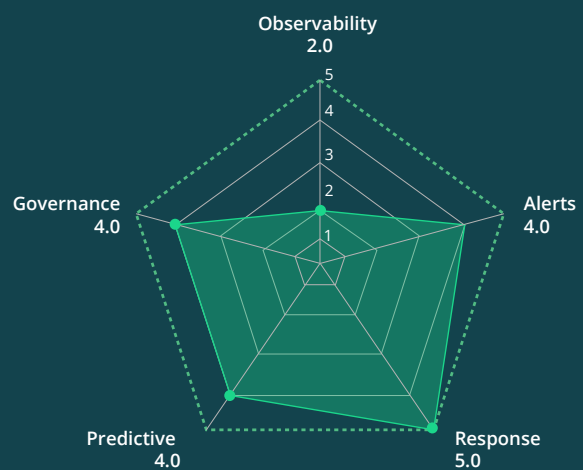

Sample result shown for illustration only.

perennial

Why You Need To Measure It

Most banks believe they sit one level higher than they actually do.

The self-assessment gap is structural, not occasional.

70% of financial services organizations self-report as “mature” or “expert” in observability. Only 12% validate as operational leaders when externally assessed. Self-reported scores run 0.5–1.0 levels above reality. The gap between what you think your operation is and what it becomes at 3am on a Saturday is what the scorecard measures.

The distance is expensive.

Organizations operating at Level 4 report MTTR reductions of 40–60% compared to Level 2 operations. Alert noise drops 84%. On-call escalations fall 71%. If your MTTR has not moved in the past 12 months despite increased observability investment, the bottleneck is not detection. It is the response and governance layers your tooling sits on top of.

The regulatory clock is running.

APAC regulators are converging on a model where operational resilience is auditable, continuous, and board-accountable. The EU’s Digital Operational Resilience Act (DORA) requires banks to restore critical services within two hours-a precedent MAS, BSP, and BNM are watching closely. The governance dimension of your scorecard is not optional. It is where regulatory expectations are moving fastest.

The only way to know where the gap is: measure it.

Take the Scorecard →

perennial

Assess Your Stack with Perennial

Where does your operation actually sit?

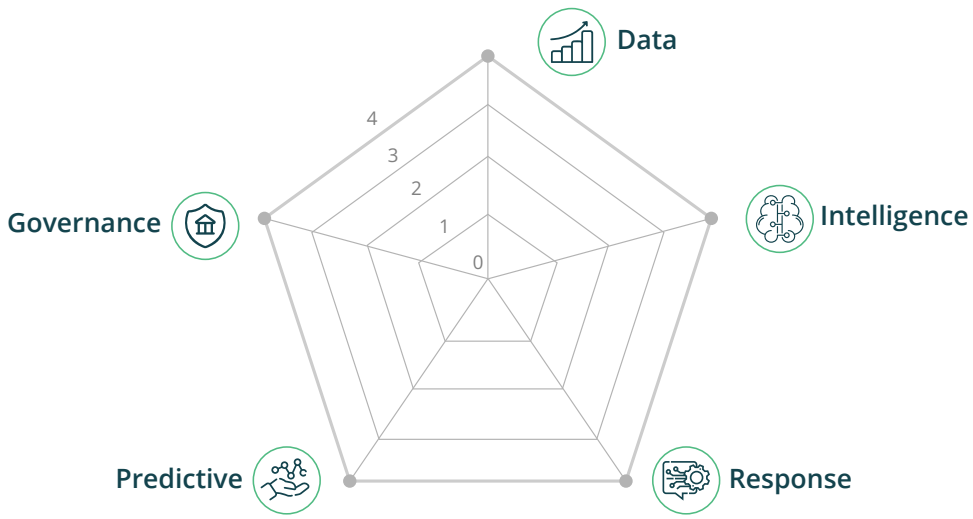

You have the assessment. You have the scorecard. You know where the gap is.

What you may not have is the operational experience to close it, the engineering capability to move from a scored assessment to a deployed, measurable improvement in 90 days, inside a regulated BFSI environment, without disrupting production.

That is what Perennial does.

We are an engineering-led AI transformation partner for regulated industries. Our AIOps practice is built on four years of embedded engineering work across APAC BFSI, including 66 engineers embedded within a single Tier-1 digital bank, building the operational model that this maturity framework is based on.

We do not sell dashboards. We do not licence platforms. We engineer the operating model changes, the correlation logic, the closed-loop remediation, the audit trails, the governance frameworks, the feedback loops, that turn your existing AIOps investment into measurable MTTR reduction and regulatory-grade operational resilience.

perennial

What a Perennial Engagement Looks Like

Week 1-2: Assessment. We run a structured version of the self-assessment you just completed, with your engineering team, your operations leads, and your risk stakeholders in the room. We produce a dimension-by-dimension maturity profile and identify the highest-ROI gap.

Week 3-8: Engineering sprint. A Perennial engineering team, embedded within your operations function, not working from a separate delivery centre, implements the operating model changes required to close the identified gap. Correlation rules. Remediation scripts. Trust boundary documentation. Audit trail architecture. Feedback loop process.

Week 9-12: Evidence and handover. We measure the outcomes against the baseline, compile the evidence, and hand over a board-ready operational resilience report. Your team owns the system. Your risk committee owns the evidence. Your regulator sees the improvement.

Start the Conversation

Book a 30-minute AIOps maturity diagnostic call with Perennial's BFSI practice lead.

A structured conversation about your scorecard results and where the highest-ROI intervention sits for your operation.

Book a 30 minute call with Perennial's BFSI practice lead →

perennial Unleash Possibilities

Perennial Systems is an enterprise technology company specialising in AI-driven IT operations and digital transformation for the BFSI sector across APAC. We operate across four pillars, AIOps (Unified Operations Intelligence), AI-Driven SDLC Co-Pilot, Agentic Banking, and AI Process Engineering, with a single strategic focus: helping regulated enterprises transition from human-operated IT to AI-native operations.

Our approach is engineering-first. We embed engineers inside your operations function, not as consultants producing recommendations, but as practitioners building the systems, the governance, and the feedback loops that produce measurable outcomes.

Web: perennialsys.com

LinkedIn: linkedin.com/company/perennial-system/

Start Your AIOps Maturity Diagnostic

Get a practical assessment of your AIOps operations, and find out where your biggest gaps are before they become huge.

AIOps Maturity Diagnostic →

Contact Us: perennialsys.com/contact-us/#get_in_touch