A WHITEPAPER FOR TECHNOLOGY LEADERS

THE TICKET IS THE LAGGING INDICATOR

Why the future of IT is zero avoidable interruption

POMEROY

There's a specific frustration that happens in every office around 10am. The spinning wheel on a laptop screen and the sigh from someone who already knows the next hour will be spent troubleshooting.

The employee tries the obvious things first. Reboots and asks whether anyone else is having the same issue. Digs through a portal, guesses at a category, and files a ticket forty minutes after the problem started. Then the employee waits.

Eleven minutes later the issue is resolved and ticket closes green. The SLA holds. The dashboard recorded eleven minutes. The business lost fifty-one.

Nobody did anything wrong. The service desk performed, the target was met, and the dashboard is green, which is exactly the problem. The ticket isn't the start of the incident. It's the receipt for time already spent.

Closing that gap requires a different objective than the one most IT organizations are measured against. Zero avoidable interruption is not zero incidents. It's an operating objective to find the interruptions that can reasonably be predicted or prevented, act and remediate them before they affect work, and produce the evidence that shows what was avoided.

This paper examines three delays that standard service measures don't capture, the boundaries of what prevention can honestly claim, and the evidence required to report prevention to a business audience.

The ticket isn't the start of the incident. It's the receipt for time already spent.

THE FIRST LAG: THE TICKET

Enterprise IT has spent decades getting better at reacting.

Service desks are faster, workflows are increasingly automated, and reporting has become more precise. Technology leaders can track response times, resolution times, escalation rates, first-contact resolution, and SLA attainment with real accuracy.

Those improvements have defined generations of IT success, but they still optimize a process that begins too late. By the time a ticket appears, an employee may have lost access to a critical application, or a device may have failed. A security event may already be escalating, or a workflow may have stopped. IT starts measuring the response only after the business has started absorbing the disruption.

Traditional MTTR is anchored to an incident or service record rather than the first moment degradation began. First-contact resolution is derived from a ticket. Customer satisfaction is typically sampled after a ticket. SLA attainment is measured against a ticket. Even service desk productivity depends on tickets entering the system in the first place.

Enterprise IT has become exceptionally good at measuring its response to interruption, yet almost nothing in a service report describes how much interruption could have been prevented altogether.

THE SECOND LAG: THE ORGANIZATIONAL HANDOFF

Timing is only half the problem. The other half is internal structure, and it produces a delay that arrives after the ticket already exists.

Employees don't experience work in functional silos. They have a job to do and an outcome they're trying to reach. Whether the issue belongs to IT, facilities, HR, or somewhere else is largely irrelevant to them. But most organizations built their service models around the internal org chart. When a request or issue crosses those boundaries, the employee is left to navigate the systems, handoffs, queues, and next steps.

Start a new hire on Monday and watch what happens when they need a device, system access, an email account, payroll setup, a workspace, and a badge. That's one outcome spread across six owners, and the new hire spends the first week as their own project manager.

This scenario happens more often than most IT leaders would like. Somewhere along the way, the cost of internal complexity shifts to the person paid to deliver.

That's the navigation tax. The employee explains the problem twice and carries context between teams that don't share a record. Every one of those steps adds time after the ticket was opened, which means the org chart is not a separate lag. It's the same lag at a later point on the same clock.

THE THIRD LAG: THE INVISIBLE WORKAROUND

When the sanctioned path is too slow or difficult to navigate, employees find another one. They move files to personal cloud storage or ask a colleague with elevated privileges for help. They create a workaround that solves the immediate problem and stays in place long after anyone remembers why it exists.

Many of these workarounds introduce security, compliance, or operational exposure that traditional service reports never see. For example, data sits outside the sanctioned environment, or access gets granted informally and never revoked. While the first lag may be measured in minutes and the second in hours or days, the invisible workarounds can persist for months or even years before anyone finds them.

No standard service measure captures any of it, yet it's still a cost of interruption, and in some cases, a far more consequential one.

WHAT SLAS AND XLAS ANSWERED, AND WHY THE NEXT QUESTION IS HARDER

The experience movement got the diagnosis right. Service level agreements (SLAs) measure the vendor rather than the person, and experience level agreements (XLAs) were built to close that gap. They both belong in the service report.

Fifty-one minutes without a critical application is an inconvenience in one environment. In another, a physician cannot pull a patient record, or a production line sits idle. Experience measures capture that difference in a way service levels cannot. But satisfaction isn't ultimately what the business is buying. It's buying capacity, and every interruption is capacity spent on something other than work.

Consider a simple incident affecting 500 employees for ten minutes. That is 5,000 minutes, or more than 83 hours of productive capacity. The service desk may resolve it perfectly within its SLA, but the business still absorbs those 83 hours of lost employee time.

Now apply the same arithmetic to recurring endpoint issues, application failures, connectivity problems, access failures, and repeatable workflow breakdowns across the enterprise. The economics change. The value of IT is not only what support costs. It's also interruption the operating model eliminates.

XLAs moved the question from how fast IT responded to what the employee experienced. The harder question is harder: how much of that interruption never had to happen?

The value of IT is not only what support costs. It's also interruption the operating model eliminates

WHAT IS TRULY AVOIDABLE?

Not every interruption can be prevented, but considerably more can be prevented than most operating models attempt. The goal is to identify what interruptions are truly avoidable, systematically move that number toward zero, and prove the outcome.

Interruptions generally fall into three tiers.

	Tiers	What belongs here
1	**Avoidable and eligible**	Known conditions with recognizable signatures and proven remediation paths. Recurring endpoint degradation, certificate expiration, capacity thresholds, known configuration drift, repeatable software failures, device health issues, access problems, and certain lifecycle events.
2	**Potentially avoidable**	Emerging conditions where telemetry reveals the problem, but the confidence threshold or remediation path is not yet mature. This tier isn't static. Conditions sitting here today become eligible as signatures accumulate and remediation is validated.
3	**Unavoidable**	Events outside the reasonable boundary of prevention. Genuinely novel failures, external provider outages, catastrophic hardware failure with no advance signal, actions prohibited by policy, and situations requiring business judgment.

The addressable universe is the first two tiers. The third defines the outer edge of what any provider can reasonably claim, and it's also the question worth asking to anyone selling prevention. Ask which interruptions they consider addressable, which they do not, and what moves an issue from one tier to the other. An answer that covers everything is not a stronger answer.

WHAT INTELLIGENT OPERATIONS CHANGE

AI, telemetry, observability, and orchestration make a different operating model possible. Instead of waiting for a person to report an issue, the environment detects the emerging condition, understands its context, acts within a defined policy boundary, and validates the result. This is zero avoidable interruption.

AI alone doesn't produce that outcome. It takes intelligence integrated into an operating model that can act safely, learn continuously, and show a record of what it did. Most AI investment in service management has gone toward the intake. A chatbot improves how a person reports a problem. Zero avoidable interruption asks whether the problem needed to reach them at all.

The certificate that expires on Thursday gets rotated Wednesday. The driver conflict responsible for fourteen tickets gets patched across the fleet before it produces the fifteenth. A capacity threshold moving toward failure triggers action before performance noticeably degrades.

Each of those is easy to describe and difficult to report. Nothing happened, which is the point, and also the problem.

A chatbot improves how a person reports a problem. Zero avoidable interruption asks whether the problem needed to reach them at all.

PREVENTION HAS TO BE AUDITABLE

Resolving an incident ends one event. Identifying and removing its cause ends a class of them

Prevention creates an unusual accountability problem. A month in which prevention worked and a month in which little went wrong look identical from the outside. That makes prevention the easiest claim to make and the hardest one to defend.

The difference has to be documented rather than asserted. Prevention that survives scrutiny leaves an end-to-end record across five stages, and each stage answers a question the business will eventually ask.

Detection. What was seen, and when?
The record begins before business impact. It captures the signal and its source, the affected entity, the deviation from baseline, the severity, and the predicted impact, with a timestamp that predates any employee awareness of the condition. Without that timestamp, prevention is a claim made after the fact.

Decision. Why was the action permitted?
Automated action requires governance, not simply activity records. The record shows which policy authorized the response, what confidence threshold the system had to clear, how risk was classified, whether human approval was required, and what the system deliberately chose not to do. Whether automation can act isn't the difficult question. What governs the action is. Capturing that is what turns a technical log into a record of accountability.

Action. What ran, and what changed?
The intervention has to be traceable. The record shows the automation or workflow, how broadly it acted, the target population, the state before and after, and any exception handling. Higher risk actions need a documented path to reverse them.

Validation. Did the condition recover, was employee interruption avoided, and did the issue remain absent?
Validation confirms three things directly: telemetry shows the condition cleared, the issue remains absent for the affected population across a defined window, and no employee-facing interruption occurs in that interval. All three are observed rather than estimated.

What validation cannot observe is the incident that never happened. The absence of impact is visible. Its cause is not. That inference rests on the established failure signature and the historical pattern for the condition, which is why eligibility gets defined before anything is counted. A condition without a known signature does not enter the denominator, and an organization that skips that step is counting detections rather than preventions.

Learning and optimization.
What changed because of the outcome?
The final stage determines whether prevention compounds or plateaus. Each outcome should improve something specific, such as a detection rule, threshold, knowledge asset, an automation, or a predictive model, all under the same governance that authorized the original action. Resolving an incident ends one event. Identifying and removing its cause ends a class of them. That's what separates an environment that measurably improves from one that simply handles the same volume more efficiently.

What still has to be true
An action taken is not the same as an interruption avoided. Reporting the result in financial terms requires an anchor: a historical baseline, an established recurrence pattern, a known failure signature, or a predicted incident probability. Without one of those, avoided hours are an assertion rather than a calculation.

And prevention cannot be demonstrated from a standing start. The baseline has to exist before anything changes, which means the measurement work begins during assessment rather than after transition.

THE SCORECARD

Moving from measuring how quickly IT responds to how much interruption it removes requires the scorecard to change with it. Traditional SLAs and XLAs are still important, but they describe what happened after an interruption entered the system. They say little about whether it could have been prevented, how much effort the employee spent, or whether the same condition is likely to return.

The signals already exist in most enterprises. What doesn't exist is a record that connects them, which is why these measures are rarely reported even though the underlying telemetry is sitting in the environment. With that record in place, IT can report what the business notices: fewer recurring problems, fewer interruptions reaching employees, less effort spent getting work done, and more productive time preserved.

Prevention: did the interruption reach the employee?

These measures answer whether interruption is being removed or merely absorbed more efficiently. Reported together, they show both what was prevented and what got through.

Measure	Definition
Validated avoided interruption rate	Validated preventions as a share of eligible conditions detected. Only conditions with an established failure signature or risk threshold belong here, which keeps the number from inflating on detections that would never have produced an incident.
Escaped interruption rate	The share of conditions detected proactively that still reached an employee. This is the counterweight to the measure above, and reporting both is what makes either believable.
Issue recurrence rate	The share of resolved issues that return for the same employee, device, application, or location within a defined window. Root cause analysis is what moves this number.

Effort: what did the employee have to do?

Resolution alone says nothing about what the employee spent getting there. These measures capture the work the organization pushed onto the person it was meant to support.

Measure	Definition
No-reengagement completion rate	The share of requests completed without the employee having to re-engage, repeat information, or chase the next step.
Cross-functional completion	The share of issues spanning multiple functions completed without the employee routing work between teams. This is the clearest measure of the second lag.
Employee effort avoided	Contacts, handoffs, and steps removed from the resolution path.

Value: what did the business get back?

Interruption is capacity spent on something other than work. This is where that capacity gets counted and expressed in terms the business already uses.

Measure	Definition
Avoided interruption hours	Employee time protected through validated prevention, reported alongside its financial expression as recovered productive capacity.

Counted and estimated are not the same thing

Six of these measures are counted directly from system data. Recurrence, validated prevention, escaped interruptions, completion, cross-functional completion, and the effort components drawn from contact and handoff records all rest on events that either occurred or did not.

Avoided interruption hours are modeled, because they rest on an estimate of how long the interruption would have lasted. Recovered productive capacity inherits that estimate and adds an assumed hourly value. Label the modeled figures as modeled and publish the assumptions behind them. The first person to audit a number that wasn't qualified will discount every number that was.

FROM REACTIVE SUPPORT TO AN INTELLIGENT OPERATING MODEL

Zero avoidable interruption is not a faster version of the current support model. It requires a different one.

The traditional model begins with the employee. Something goes wrong, the employee notices, a ticket is opened, the issue is triaged and assigned, and eventually it is resolved. The model works, and decades of investment have made it work efficiently. It also makes the employee the first monitoring system. Every measurement starts after that person has already absorbed the interruption.

The emerging model moves that responsibility off the employee. The environment is observed continuously for signals that a condition is degrading. Intelligence determines whether the pattern warrants action, what action is permitted, and whether the situation requires human judgment. The response is initiated and validated within defined boundaries, and the outcome feeds back into detection.

The biggest change is where human expertise goes. Predictable conditions get handled by the system under governance. That leaves people free for work that requires human expertise, such as novel failures, ambiguous signals, and decisions with real business consequences that no policy anticipated. The IT organization doesn't get smaller. It stops spending its most expensive resource on problems that should never have reached a person.

The result is a different standard for enterprise IT performance. The objective isn't zero tickets. It's zero avoidable interruption, meaning the ability to identify which interruptions can be prevented, solve those conditions before they reach employees, and coordinate the outcome across whatever boundaries it crosses. It also means producing a record of what was detected, decided, changed, validated, and improved.

What this asks of technology leadership

CIOs carry the number. When finance asks where prevention hours came from, the CIO answers, which is why the line between counted and modeled figures isn't a technical detail. It's what makes the number reportable in the first place.

For CTOs and VPs, the job is making prevention safe and repeatable. Telemetry has to be broad enough to catch meaningful deviation, and context has to connect a signal to the people and services behind it. Policy has to define what the system may do and at what confidence. Validation has to close the loop and surface anything the action broke.

Neither works without the other. The CIO cannot report a number the CTO's environment cannot produce, and the CTO cannot expand what the system is permitted to do without a business case the CIO is willing to defend.

WHERE TO START

Understanding what prevention requires is one thing. Building an operating model that produces it consistently is another. Every stage of the record needs something behind it that reliably generates that stage, in a live environment, at enterprise scale. Where that accountability is unassigned, the record has a gap, and gaps in the record are where prevention claims fall apart.

Most organizations can produce two or three of the five today. Detection is widely available, and action is heavily automated in most enterprises. Decision and validation are where prevention programs usually break down, because one requires governance designed in advance and the other requires enough history in the environment to know what normal looks like there.

Start by asking which of the five stages already has an owner. An organization that can answer all five has an operating model. One that can answer three has a plan.

An organization that can answer all five has an operating model. One that can answer three has a plan.

	Stage	Key question	What it has to product
	**Detection**	Does telemetry cover the estate broadly enough to catch deterioration before an employee does?	The signal, the impacted entity, the baseline deviation, the predicted impact
	**Decision**	Is there a policy that defines what may be acted on automatically, and at what confidence?	The authorizing policy, the confidence threshold, the risk classification, the approval requirement
	**Action**	Can the response execute across remote, lifecycle, and field paths, not just remotely?	The workflow, the target population, the state before and after, the rollback path
	**Validation**	Does anything confirm the condition cleared and stayed cleared?	Recovery confirmation, the validation window, the recurrence status
	**Learning**	Does an outcome change a rule, a threshold, or a policy, or does it just get logged?	The root cause where identified, and the change made as a result

THE NEXT MEASURE OF ENTERPRISE IT

The report will keep improving. Response times come down, satisfaction holds, and every number on the page stays accurate. None of them describe the hour that disappeared before anyone filed a ticket.

That hour doesn't appear in the measurement system, because the system starts counting at the ticket, and the ticket is the last event in the sequence.

Nothing here requires abandoning what already works. Response time, resolution time, availability, and satisfaction all still show how well IT handles the interruptions that reach it. What's missing is any measure of the interruptions that never arrive.

Getting there takes a clear definition of what is genuinely addressable, telemetry that catches degradation early, governance that says what a system may do on its own, and validation honest enough to admit what it cannot observe.

The next quarterly review is a good place to test this. Ask what was seen before anyone reported the problem, what the system was permitted to do about it, and whether anyone can produce the record. That answer tells you whether you're buying faster recovery or fewer interruptions.

pomeroy.com