

EBOOK

The Autonomous IT Playbook

Five practical ways to get ahead
of **IT disruption**

A field guide for IT teams ready to spend less time
reacting and more time moving the business forward.

SEE > CONNECT > UNDERSTAND > ASSIST > AUTOMATE

riverbed®

A practical route from reactive operations to trusted autonomy

Built around the operational realities of IT Operations, Service Management, Infrastructure, NetOps, Digital Workplace and application teams.

- 3 The environment changed. The operating model now has to catch up
- 4 Reactive IT is taking a human toll
- 5 What Autonomous IT really means
- 6 Five ways to put Autonomous IT to work
- 12 Autonomous IT Readiness Scorecard
- 15 Build trust before you build scale
- 16 From first use case to operating model
- 17 Keep people at the center
- 18 Choose one problem worth preventing

The environment changed.

The operating model now has to catch up.

Most IT teams do not need to be told to work harder. They are already working hard.

They are responding to alerts, managing incidents, supporting employees, maintaining infrastructure and keeping critical services running. At the same time, they are being asked to adopt AI, modernize operations and help the business move faster.

Something has to give.

Modern IT environments have grown beyond what people can reasonably monitor and manage manually. Applications run across data centers, cloud services and third-party platforms. Employees connect from almost anywhere. A single digital service can depend on technologies owned by several different teams, often stretching across internal systems, public cloud platforms, SaaS services and third-party providers.

When something goes wrong, finding the cause is rarely straightforward.

The network dashboard looks healthy. The application team sees nothing unusual. Infrastructure reports that everything is running normally. Yet users are still struggling.

That is the reality behind the Autonomous IT conversation. It is not a race to remove people from operations. It is a response to a much simpler problem: there is now too much complexity, too much data and too little time for traditional operations to keep up.

The goal is not autonomy for its own sake. The goal is an IT operation that employees, customers and the business can rely on.

THE RIVERBED POINT OF VIEW

Most organizations do not lack monitoring data. The challenge is that data exists in silos and is rarely connected across domains. Teams can see more than ever before, but understanding what matters still takes too long. Autonomous IT closes that gap by turning fragmented signals into insight, and insight into action.

Reactive IT is taking a human toll

Tickets and response times tell only part of the story.

It is easy to describe reactive IT in terms of tickets, outages and response times. But those measures do not tell the whole story.

They do not capture the hours spent on incident calls. The projects that slipped because key people were pulled into troubleshooting. The frustration of investigating the same issue for the third time. Or the pressure of telling the business that the team is still searching for an answer.

Consider a familiar scenario.

A critical application begins to slow down on Monday morning. Employees contact the service desk, but the reports are inconsistent. Some people cannot connect at all. Others can connect, but performance is poor.

Several teams begin investigating. Each team has its own tools and its own version of what is happening. Everyone is looking at a piece of the problem, but nobody can see the complete picture.

The incident call grows. More people join. More data is shared. The pressure rises, but the answer remains out of reach.

This is not a failure of skill or effort. It is the consequence of asking people to manually connect thousands of data points across an environment that never stops changing.

Autonomous IT is about changing that experience. Instead of waiting for the next alert and then working backwards, IT can begin to identify risk, understand impact and determine the best response while there is still time to act.

From reactive to autonomous

Reactive operations	Autonomous operations
Users report issues	IT detects emerging issues
Teams investigate separate data	Data and alerts are correlated across domains
Experts search for root cause	AI surfaces a probable diagnosis with supporting evidence
Response depends on manual effort	Known remediation workflows can be automated
Success means restoring service quickly	Success means avoiding disruption
People spend time managing noise	People focus on decisions and improvement

Moving from one side of this table to the other will not happen in a single step. Nor should it. The practical route is to introduce autonomy gradually, beginning with problems that teams already understand and actions they already trust.

What Autonomous IT really means

“Autonomous” can sound radical. In practice, it starts with something much more practical: better help for teams under pressure.

For most organizations, Autonomous IT does not begin with a system making every decision or taking every action on its own. It begins by giving teams clearer signals, better context and recommended next steps they can trust.

AI highlights something that might otherwise be missed. It correlates activity across different systems. It explains why an issue may be happening and provides evidence behind the diagnosis. It recommends an action based on that diagnosis.

A person remains in control.

As confidence grows, teams may allow selected actions to happen automatically. These are usually actions that are well understood, repeatable and low risk.

This progression matters because trust cannot be mandated. Teams need to see that recommendations are accurate. They need to know what data is being used. They need clear controls over which actions are permitted and when human approval is required.

Autonomous IT therefore combines four essential foundations.



Visibility

You cannot prevent a problem that you cannot see. IT needs a meaningful view across networks, applications, infrastructure, devices and employee experience. A collection of separate dashboards is not the same as connected visibility.



Context

An alert tells a team that something changed. Context explains whether it matters, who is affected, what depends on it and whether the signal is isolated or part of a pattern.



Intelligence

AI can analyze more data than any individual team could examine manually. Used well, it can surface unusual behavior, identify patterns and suggest probable causes. It helps operational expertise go further.



Governed action

Repeatable issues can be resolved automatically using pre-approved workflows. People remain involved when judgment, approval, or exception handling is required.

THE PRINCIPLE

Responsible autonomy is not uncontrolled automation. It is human-governed action, grounded in cross-domain data and clear operational policies.

Five ways to put Autonomous IT to work

Useful conversations about autonomy begin with recurring operational problems, not technology.

The best starting point is not the technology. It is the operational pressure teams already recognize: service desk demand, cross-team investigations, missed warning signs, quiet employee frustration and major incidents where every minute matters.

Each use case below looks at autonomy through a different lens, so teams can see where AI-assisted understanding and governed automation could help first.

1

Find the problem before the service desk does

Detect emerging experience issues before tickets multiply.

2

Replace the blame game with a shared answer

Correlate network, application, infrastructure and user experience data for a more accurate diagnosis.

3

Address warning signs before they become incidents

Recognize patterns early enough to prevent disruption.

4

Give employees a better digital working day

Find and remediate quieter problems that erode productivity.

5

Make incident response faster and calmer

Turn operational data into evidence and a clear route to action.

1. Find the problem before the service desk does

Detect emerging experience issues before tickets multiply.

THE FAMILIAR SITUATION

The day starts normally, then the first few tickets arrive.

Individually, they look routine: a slow business application, dropped calls, devices taking longer than usual to start.

None looks severe enough to declare a major incident.

But the pattern is already forming.

By the time the queue grows enough to demand attention, the service desk has become the organization's early-warning system by default.

A MORE AUTONOMOUS APPROACH

Experience and performance data are continuously assessed for unusual behavior. When similar symptoms begin to appear across a group of employees, AI connects them and highlights a possible shared cause.

The team can see what has changed, who is affected, whether the problem is spreading, which application, device or service may be involved, and what action could prevent further impact. For a trusted and repeatable issue, an approved remediation may be triggered automatically. In other cases, AI may create

an ITSM ticket, attach the supporting evidence and escalate it to the team with the authority to resolve the issue before the service desk queue grows.

WHAT CHANGES FOR THE TEAM

The service desk spends less time categorizing duplicate reports, repeating the same fix or manually routing issues to the right team. Escalations can happen automatically, with the right evidence attached and the right team brought in sooner.

Employees spend less time waiting for help.

Most importantly, IT has a chance to act before inconvenience becomes a major disruption.

A PRACTICAL PLACE TO BEGIN

Identify one of the most frequent employee issues handled by the service desk. Look at how it is currently detected, how often it repeats and whether the remediation follows a consistent process. A problem that is common, recognizable and straightforward to fix is a strong candidate for early automation.

2. Replace the blame game with a shared answer

Connect teams around shared evidence.

THE FAMILIAR SITUATION

A critical service starts to underperform, but no team sees a clear fault in its own domain.

The application team checks the platform. Infrastructure sees available capacity. Network Operations finds no obvious issue. Each team is working responsibly, but from a different slice of the truth.

The conversation shifts from “what is happening?” to “where is it not happening?”

That may protect each team’s view of the environment, but it does not get the business any closer to a working service.

A MORE AUTONOMOUS APPROACH

Network, application, infrastructure and end-user experience data are brought together into a shared operational context, so teams can understand how issues relate across domains rather than investigating each signal in isolation.

AI looks across the data rather than analyzing each source in isolation. It identifies relationships, compares current behavior with previous patterns and highlights the most likely source of degradation.

WHAT CHANGES FOR THE TEAM

Teams work from shared evidence instead of defending separate dashboards.

Specialists still apply their expertise, but spend less time proving that their part of the environment is healthy and more time resolving the issue.

For the business, the distinction between a network problem and an application problem was never important. The business simply wanted the service to work.

A PRACTICAL PLACE TO BEGIN

Choose a critical service that regularly requires several teams to troubleshoot. Map the data each team uses, the gaps between those views and the points where investigations tend to stall. Then look for ways to integrate those data domains into a shared operational context, so AI can support faster, evidence-based root-cause analysis.

3. Manage warning signs before they become incidents

Recognize patterns while there is still time to act.

THE FAMILIAR SITUATION

Some incidents do not begin as incidents at all. They begin as weak signals.

Performance drifts. Errors increase. Capacity edges toward a threshold. A recent change introduces unusual behavior, but not enough to trigger an existing alert.

On their own, these signals are easy to dismiss. Together, they may be the beginning of the next disruption.

A MORE AUTONOMOUS APPROACH

AI establishes an understanding of normal behavior and looks for meaningful deviations. It can identify patterns that rule-based monitoring may not recognize, then relate those patterns to recent changes, dependencies and previous incidents.

This gives teams the opportunity to investigate while the issue is still manageable. Sometimes the response will be an operational change. Sometimes it will be a recommendation to monitor more closely. Sometimes no intervention will be needed.

The important change is that the decision happens earlier.

WHAT CHANGES FOR THE TEAM

Operations teams move beyond measuring only how quickly they restore service.

They can begin asking a more valuable question:

How many incidents did we prevent from happening at all?

This is the shift at the heart of Zero Disruption.

A PRACTICAL PLACE TO BEGIN

Review a recent high-impact incident. Ask what changed beforehand, whether weaker signals sat below existing thresholds, whether similar symptoms appeared elsewhere and what earlier action might have reduced or prevented the impact.

4. Give employees a better digital working day

Find the quieter problems that never become outage reports.

THE FAMILIAR SITUATION

Not every technology issue announces itself as an outage. Some simply make the working day harder.

A device takes longer to start. A collaboration tool freezes during calls. A business application becomes progressively slower. A remote employee repeatedly loses connectivity.

Each interruption may seem minor in isolation, but the impact builds across the day. Some employees contact the service desk. Many simply work around it, lose time or accept a poorer experience.

The absence of a ticket does not always mean the absence of an issue.

A MORE AUTONOMOUS APPROACH

Digital experience signals are monitored across applications, devices and employee environments. AI identifies deteriorating experience, recognizes recurring patterns and determines whether a broader group is affected.

Where an appropriate remediation is known, IT can take action without waiting for each employee to report the same problem.

An application can be restarted. A configuration can be corrected. A device issue can be addressed. An employee can be guided through a self-service action.

WHAT CHANGES FOR THE TEAM

Digital Workplace and End User Computing teams gain a more accurate understanding of employee experience.

The service desk receives fewer repetitive requests. Employees experience fewer interruptions and regain time they would otherwise lose to slow or unreliable technology.

It is not automation for the sake of efficiency.
It is technology getting out of the employee's way.

A PRACTICAL PLACE TO BEGIN

Don't just troubleshoot individual incidents. Combine employee experience data with service desk trends to identify the issues that affect the most users or occur most often. Then address the underlying cause, so teams can prevent future tickets instead of repeatedly responding to the same problem.

5. Make incident response faster and calmer

Turn information into a clear route to action.

THE FAMILIAR SITUATION

When a major incident begins, the first challenge is often not expertise. It is time.

IT experts join the call while the picture is still unclear. Screens are shared. Data is copied from one tool to another. Someone creates a timeline. Someone else checks recent changes. Specialists search for similar incidents.

The organization has valuable evidence, but it is scattered. Before the team can decide what to do, they first have to assemble the story.

A MORE AUTONOMOUS APPROACH

AI collects and correlates relevant operational information as the incident develops.

It can help the response team understand what changed, which services and users are affected, how the symptoms relate, whether a similar incident has happened before, which causes are most probable and which response worked previously.

For established incident types, a pre-approved workflow may carry out selected actions automatically.

More complex situations remain under human control.

WHAT CHANGES FOR THE TEAM

Incident calls become focused on decisions rather than data gathering.

Experts are brought in because their judgment is needed, not because someone needs another pair of eyes on a dashboard.

Time to understanding falls. Time to action follows.

The response becomes more controlled.

A PRACTICAL PLACE TO BEGIN

Examine the first 30 minutes of recent major incidents. Identify which information teams repeatedly had to locate, who had access to it, and which decisions were delayed because the full context was not available. Then create workflows, skills, or remediation playbooks that make that information and response available the next time.

Autonomous IT Readiness Scorecard

A practical scorecard to help identify the next sensible step.

Autonomous IT is not a binary state.

Most organizations will already have some of the required capabilities. They may have strong visibility in one area, established automation in another and significant manual effort somewhere else.

Score each statement from 1 to 5: 1 means rarely or never true; 3 means partly true or inconsistent; 5 means consistently true across the operation.

Visibility and context

Statement	Score
1 We can understand the experience of users, applications and services across our environment.	----
2 Our teams can connect network, application, infrastructure and end-user data.	----
3 We can quickly determine who and what is affected when performance deteriorates.	----
4 We have enough historical context to recognize recurring operational patterns.	----

Detection and understanding

Statement	Score
5 We regularly identify issues before users report them.	----
6 Our teams can distinguish meaningful alerts from operational noise.	----
7 We can establish probable root cause without lengthy manual investigation.	----
8 We can relate incidents to changes, dependencies and previous events.	----

Action and automation

Statement	Score
9 We use AI to provide useful recommendations during incident investigation.	----
10 Repeatable operational actions are documented, standardized, and converted into remediation workflows.	----
11 Selected remediations can be carried out automatically within agreed controls.	----
12 IT can see why an automated action was recommended or taken.	----

People and governance

Statement	Score
13 Operational teams understand where AI is used and where human approval is required.	----
14 We have clear ownership for automated workflows.	----
15 Teams trust the quality of the data used to support AI decisions.	----
16 We measure prevented disruption as well as incident response.	----

Your Score

Interpreting the result

16-31

Reactive

Points

Your operation still depends heavily on manual detection, investigation and response. Begin by improving visibility and choosing one recurring operational problem to address.

32-51

Assisted

Points

You have useful capabilities, but they may be limited to specific teams or technologies. Focus on connecting existing data and introducing AI assistance into time-consuming investigations.

52-67

Automation-ready

Points

The foundations are in place for selected autonomous workflows. Prioritize repeatable, predictable actions with clear ownership and low operational risk.

68-80

Advancing toward autonomy

Points

AI and automation are becoming part of the operating model. Expand proven use cases, strengthen governance and measure the disruption avoided through earlier action.

This is a practical planning scorecard, not an industry benchmark. Its purpose is to help teams identify a sensible starting point.

Build trust before you build scale

The first use case does not need to transform the whole operation. It needs to work.

The quickest way to undermine an Autonomous IT program is to begin with the most complicated problem in the organization. A better first use case usually has five characteristics.

1

It happens frequently

A recurring problem gives the team enough opportunities to test, learn and improve.

2

It follows a recognizable pattern

AI and automation are most useful when signals reliably indicate what is happening.

3

The current response is repeatable

If experienced team members already follow similar steps each time, part of that process may be suitable for assistance or automation.

4

The outcome can be measured

The team should be able to see whether the use case reduced tickets, investigation time, disruption or manual effort.

5

The level of risk is understood and low

Teams need clear boundaries around which decisions can be automated and which still require approval.

A successful, controlled use case builds confidence. It gives people something tangible to evaluate and creates a foundation for the next step.

From first use case to operating model

A practical Autonomous IT program does not move from manual operations to full automation in one leap. It starts with one problem worth solving, one trusted workflow and one measurable outcome.

Once the first use case is working, the next step is to identify adjacent problems where the same data, context or remediation approach can be reused. A recurring service desk issue may become a model for other employee experience problems. A successful incident workflow may become the foundation for similar high-pressure response scenarios. A trusted recommendation may become a candidate for governed automation.

This is how autonomy scales safely: not through more tools or broader ambition alone, but through repeatable proof.

Teams should ask:

- What did this use case help us see sooner?
- What context made the recommendation trustworthy?
- Which manual steps were removed or reduced?
- What evidence gave teams confidence to act?
- Where else could the same pattern apply?
- What still needs human judgment or approval?

The goal is to build a library of trusted patterns: signals the organization understands, recommendations the team can explain, workflows with clear ownership and automations that stay within agreed controls.

That gives IT a practical path forward. Start with the problems teams already know. Prove that AI can reduce the burden, shorten investigation and prevent repeat disruption. Then expand from there.

WHY THIS MATTERS

Autonomous IT becomes credible when teams can see how it helps them in real operational situations. Each trusted use case creates confidence, and each proven workflow creates a stronger foundation for the next.

Keep people at the center

Operational expertise is what makes responsible autonomy possible.

The hardest part of Autonomous IT may not be the technology. It may be helping teams trust what autonomy means in practice.

For someone who has built a career on deep operational knowledge, the idea of a system taking action can feel uncomfortable. That reaction should not be dismissed.

Operational expertise is what makes responsible autonomy possible.

People define the policies. People decide which actions can be trusted. People identify where risk is acceptable. People handle the situations that require judgment or an understanding of wider business priorities.

A good Autonomous IT program does not make operational expertise less valuable. It removes the repetitive work that prevents experts from using it.

That creates a better working experience for IT teams as well as a more resilient experience for the business.

THE OPERATING PRINCIPLE

Human expertise sets the boundaries. AI extends what teams can see and understand. Automation acts only where the process has earned enough trust.

The goal is not more automation

It is easy to measure an Autonomous IT program by the number of processes automated. That is the wrong finish line.

- Did the team identify the issue before employees were affected?
- Was the cause understood without a lengthy incident call?
- Did experts have the evidence they needed to make a decision?
- Were repetitive tickets prevented?
- Did a critical service remain available?
- Did IT have more time for work that moves the organization forward?

The goal is confidence. Confidence that teams can see what is happening. Confidence that a warning will not disappear into a flood of alerts. Confidence that AI recommendations are grounded in reliable data. Confidence that automated actions will stay within agreed controls.

YOUR NEXT STEP

Choose one problem worth preventing

Start with a smaller, more useful question.

THE QUESTION

Which recurring issue would have the greatest impact if eliminated?

Do not begin by asking how much of IT can be automated.

Begin with one service. One recurring problem.
One action the team can trust. Then another.

Autonomous IT will not remove every incident.
Zero Disruption is not about pretending nothing will ever fail.

It is about building an operation that can see more, understand sooner and act before disruption reaches the people who depend on technology.



This is where Riverbed gives teams a practical foundation for Autonomous IT: visibility across complex environments, shared operational context across network, application, infrastructure and end-user experience data, AI-assisted understanding, and governed automation for trusted responses.

Together, these capabilities help IT teams move from reacting to disruption after it reaches users, to identifying risk earlier, acting with confidence and preventing more issues before they become business problems.

See Autonomous IT in action

Explore how Riverbed can help your organization connect, understand and act across the digital experience – and move from reactive operations toward Zero Disruption.

Request an Autonomous IT Strategy Session



About Riverbed

Riverbed, the leader in AI observability, helps organizations optimize their users' experiences by leveraging AI automation for the prevention, identification, and resolution of IT issues. With over 20 years of experience in data collection and AI and machine learning, Riverbed's open and AI-powered observability platform and solutions optimize digital experiences and greatly improve IT efficiency. Riverbed also offers industry-leading Acceleration solutions that provide fast, agile, secure acceleration of any app, over any network, to users anywhere. Together with our thousands of market leading customers globally – including 95% of the FORTUNE 100 – we are empowering next-generation digital experiences. Learn more at riverbed.com.

© 2026 Riverbed Technology LLC. All rights reserved. Riverbed and any Riverbed product or service name or logo used herein are trademarks of Riverbed Technology. All other trademarks used herein belong to their respective owners. The trademarks and logos displayed herein may not be used without the prior written consent of Riverbed Technology or their respective owners. CS-1536_The-Autonomous-IT_Playbook_US_100526