ARCTIQ

AUTONOMOUS OPERATIONS READINESS ASSESSMENT

How Ready Are Your Applications for Autonomous Operations?

Assess your observability and automation maturity, and what has to be true before AIOps or self-healing can work safely at scale.

AI Doesn't Fix Weak Signals. It Automates Them Faster.

On-call fatigue, alert storms, and manual toil push a lot of teams toward the same conclusion: wire in some AI, add automation, and the pain goes away.

But automation only amplifies what's already in place. If your monitoring and observability are weak, the signals feeding it can't be trusted, and a decision that once took a person 20 minutes to make can now execute itself at 2 a.m. without anyone verifying it.

Before autonomous operations can work safely, your team needs three things: end-to-end visibility, enough context to understand business impact, and confidence in the response you're automating.

Use this assessment to evaluate each application's maturity across five stages and identify where to focus next.

HOW TO USE THIS ASSESSMENT

Evaluate each application separately. A customer-facing platform might sit at Stage 4 while an internal reporting tool sits at Stage 2, and that's normal. The goal is knowing where each one stands, not hitting a single organization-wide score.

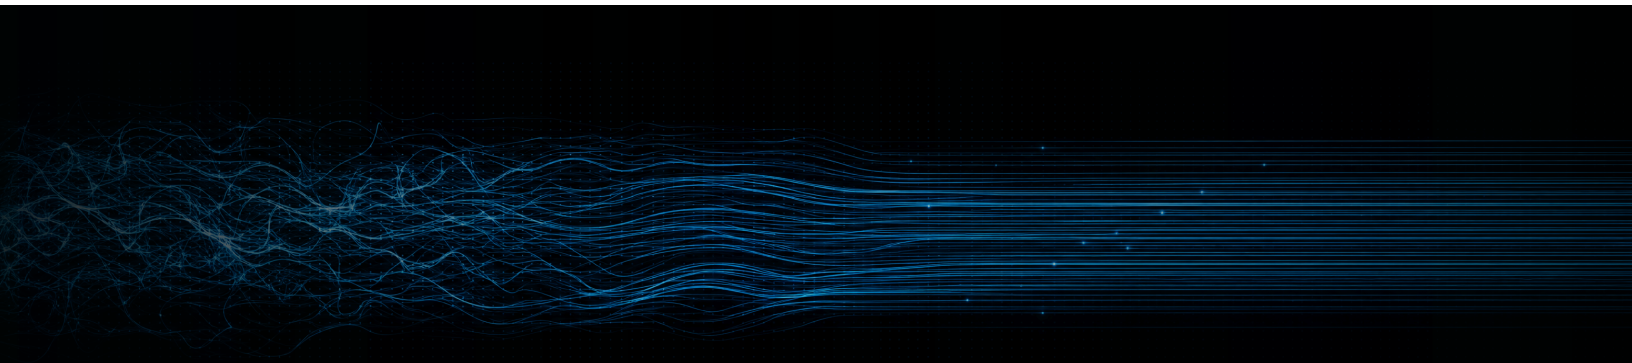

Noise on the left. Trusted signal on the right.

Five Stages of Observability Maturity

STAGE 1

Foundation

Your metrics, logs, and traces exist in one place, and the team trusts the data.

- [ ] Metrics, logs, and traces are collected in a single system, not scattered across tools.
- [ ] The team trusts the data flowing in. It's not widely seen as incomplete or unreliable.
- [ ] Coverage extends across business-critical applications, not just infrastructure.

STAGE 2

Observability

You can see how infrastructure and applications behave together, and find root cause.

- [ ] The team can pinpoint root cause without switching between three or four platforms.
- [ ] Dependencies across the environment are understood, not discovered during an incident.
- [ ] A root cause can be validated within minutes, not hours.

STAGE 3

Intelligence

Once observability is trusted, the system starts finding patterns on its own.

- [ ] Anomaly detection and alert correlation are in active use, not just available.
- [ ] The system can identify the likely cause of an issue, not just flag that one exists.
- [ ] Recurring patterns across incidents are recognized, not rediscovered each time.

Automation and Autonomous Operations

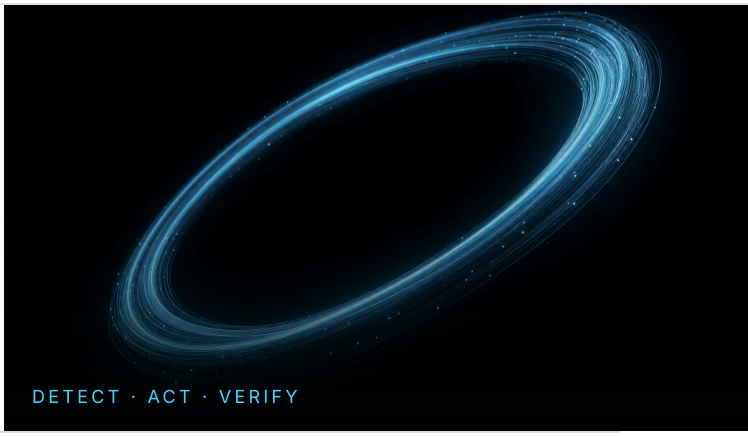

STAGE 4

Automation

Findings start triggering actions: playbooks, rollbacks, configuration changes, scaling.

- [ ] Automated actions exist for the repeatable, common issues your team already knows.
- [ ] Every automated action is tied to a documented runbook, not tribal knowledge.
- [ ] Tickets and records close on their own. Nobody spends Friday afternoons closing tickets by hand behind the automation.

STAGE 5

Autonomous Operations

The system detects, diagnoses, acts, verifies the result, and feeds that outcome back in.

- [ ] The system checks whether the fix worked. It doesn't just assume the playbook resolved the issue.
- [ ] Recurring problems get flagged as recurring, not silently re-triggered every time.
- [ ] Outcomes feed back into observability and service-management systems, closing the loop.

A NOTE ON STAGE 4 VS. STAGE 5

Automated remediation stops the symptom in one shot. Self-healing verifies the result and recognizes when the same problem keeps returning. Successfully automating the wrong fix can be worse than doing nothing, because it hides what's actually happening.

What Your Results Tell You

Choose the application you'd get measured for if it went down. Then look at which criteria you couldn't check off and where those gaps cluster. That's where to focus next.

Stage 1–2

Most gaps here? Fix the foundation first. Untrusted data and low visibility mean automation would be acting on signals nobody can vouch for, so it's not the next investment yet, no matter how far along the other stages look.

Stage 3

Gaps mostly here? Your data is trustworthy, but the system isn't finding patterns on its own yet. Build anomaly detection and correlation before layering on more automation.

Stage 4–5

Gaps mostly here, with Stages 1–3 solid? You're closer than most. Pilot automation on one well-understood application with a human checking outcomes, then work toward closing the feedback loop.

Watch for this pattern specifically

If Stages 4 and 5 look further along than Stages 1 and 2, that's a warning sign, not a strength. It usually means automation is already running on top of observability nobody's fully verified yet, the exact pattern that leads to automating the wrong decisions faster instead of fewer of them. Close the earlier gaps before building further on top.

In general, the stage with the most unchecked criteria is where to focus next. But if your earlier stages still have unresolved gaps, address those before building further on top.

Finding your automation opportunities: if Stages 1–3 are solid and your gaps cluster in Stage 4, look for the runbook your team already repeats by hand, the same five steps, the same fix, week after week. That repeatable, well-understood work is exactly what's worth automating first, and exactly where a platform like Red Hat Ansible Automation Platform becomes the execution layer for what your observability has already diagnosed.

The Technology Behind the Loop

Closing the loop takes more than one platform. Observability tells you what happened and why. Automation is what actually acts on it.

Foundation	Observability	Intelligence	Automation	Autonomous Operations

dynatrace

Detects the issue, diagnoses the likely cause, and understands the business impact behind it. Once an action is taken, it also verifies the outcome: did the metric stabilize, did the anomaly clear, did the SLO return to green. This is the "trust the signal" half of the loop, spanning Stages 1 through 3, and the verification step in Stage 5.

Red Hat

Ansible Automation Platform

Executes the response once observability has done its job: the rollback, the configuration change, the scaling action. This turns Stage 4 from a documented process into repeatable execution, tied to the diagnosis that made it trustworthy in the first place.

WHY BOTH MATTER automation without trusted observability behind it means acting on signals nobody's verified. Observability without automation on top of it means your team still does the repeatable work by hand. The closed loop needs both, in the right order.

What Closing the Loop Delivers

Examples of outcomes Arctiq has seen across customer environments when observability and automation work together. Results will vary.

75%
reduction in mean time to recovery (MTTR)

60%
fewer war rooms for incident response

40%
faster production update and DevOps delivery timelines

Know Where You Stand. Then Move Deliberately.

The maturity model isn't a one-time checklist. It's a way to keep advancing each application's readiness without automating decisions before the underlying signals and responses can be trusted.

Talk to Arctiq

Arctiq helps enterprise IT operations and platform engineering teams build the observability foundation autonomous operations depend on with Dynatrace, then put Red Hat Ansible Automation Platform to work as the execution layer once that foundation is trustworthy.

Let's turn your results into a plan.

Connect With Us →