New Relic

How to Prepare for AIOps

Four steps for a successful deployment

Table of Contents

Introduction: Are You Ready?

As software and systems become more complex and organizations ship software faster and more frequently, DevOps, site reliability engineering (SRE), and network operation center (NOC) teams can find themselves overwhelmed by a constant flood of data. Today's modern technology stack means there are now many more things to monitor and respond to—a wider surface area, more software changes, more operational data emitted across fragmented tools, more dashboards, a complex web of dependencies, and more alerts.

At the same time, these teams are under increasing pressure to find and fix issues faster, or better yet, prevent them from happening in the first place. However, between noisy alerts, signals distributed among multiple tools, and thousands of "unknown unknowns," it's difficult to quickly determine and address the root cause of incidents, let alone detect and respond to issues proactively.

AIOps helps teams find solutions to problems faster and unearth unknown-unknowns or issues they might have missed, so that they can get out of reactive firefighting mode and back into the creative work of building more perfect software. Coined by industry analyst firm Gartner, artificial intelligence for IT operations (AIOps) combines big data and machine learning to augment IT operations processes, including anomaly detection, event correlation, alert noise reduction, and root cause analysis.

As with adopting any powerful new tool, your success with an AIOps solution will depend on your preparation. The better you prepare, the better outcomes and higher value you can expect. Use the four steps explained in this ebook to make sure you've prepared a solid foundation for your AIOps journey.

In a survey of nearly 100 large enterprises, the AIOps Exchange found that 40% of those IT organizations face more than 1 million event alerts each day.

Source: "**The Current State of AIOps**," Mary Branscombe, The New Stack, August 2019

Can AIOps Help Your Organization?

When you do AIOps right, you'll recognize it's more than a trend for software teams.

The right AIOps solution harnesses the power of data science to proactively detect anomalies so you can prevent issues from happening in your systems and helps accelerate your resolution of the issues that do happen. AIOps helps you tame the unknown unknowns, reduce alert noise, and get to the root cause so you can find and fix issues faster and focus on what's important to you.

While most shops could benefit from AIOps, how will you know whether it can deliver measurable improvements for your teams? To answer that, think about the types of problems that AIOps solves and the severity of those problems for your organization—for example:

- Does your company as a whole or certain systems within it suffer from downtime, service interruptions, or frequent performance degradations that negatively affect customer experience and your service level objectives (SLOs)?
- Are you lacking robust alert coverage, or do you suspect you have a lot of unknown unknowns?
- Are your teams spending too much time fighting fires rather than proactively identifying issues before they cause outages or performance problems?
- Do your software teams have too little time for development because they're spending inordinate amounts of time trying to identify and respond to issues?
- Does complexity or alert fatigue and noise prevent your teams from identifying and addressing critical issues quickly?
- When issues occur, are you able to quickly identify the root cause or is it like trying to to find a needle in a haystack?

If you answered "yes" to any or all of these questions, AIOps could help you address those issues and have a measurable, positive impact on:

- Mean time to resolution (MTTR)
- Downtime
- Adherence to SLOs
- Alert noise, alert fatigue, and manual toil
- Incident response
- Developer and IT operations productivity

How does AIOps work?

AIOps combines natural language processing, statistical models, supervised and unsupervised learning, recommendation models, and more to ingest data from multiple sources, normalize it, suppress low-priority alerts, correlate related incidents and events into single issues, enrich them with context, identify probable root cause, and then notify the right people or teams. Learn more in this infographic: Demystifying AIOps.

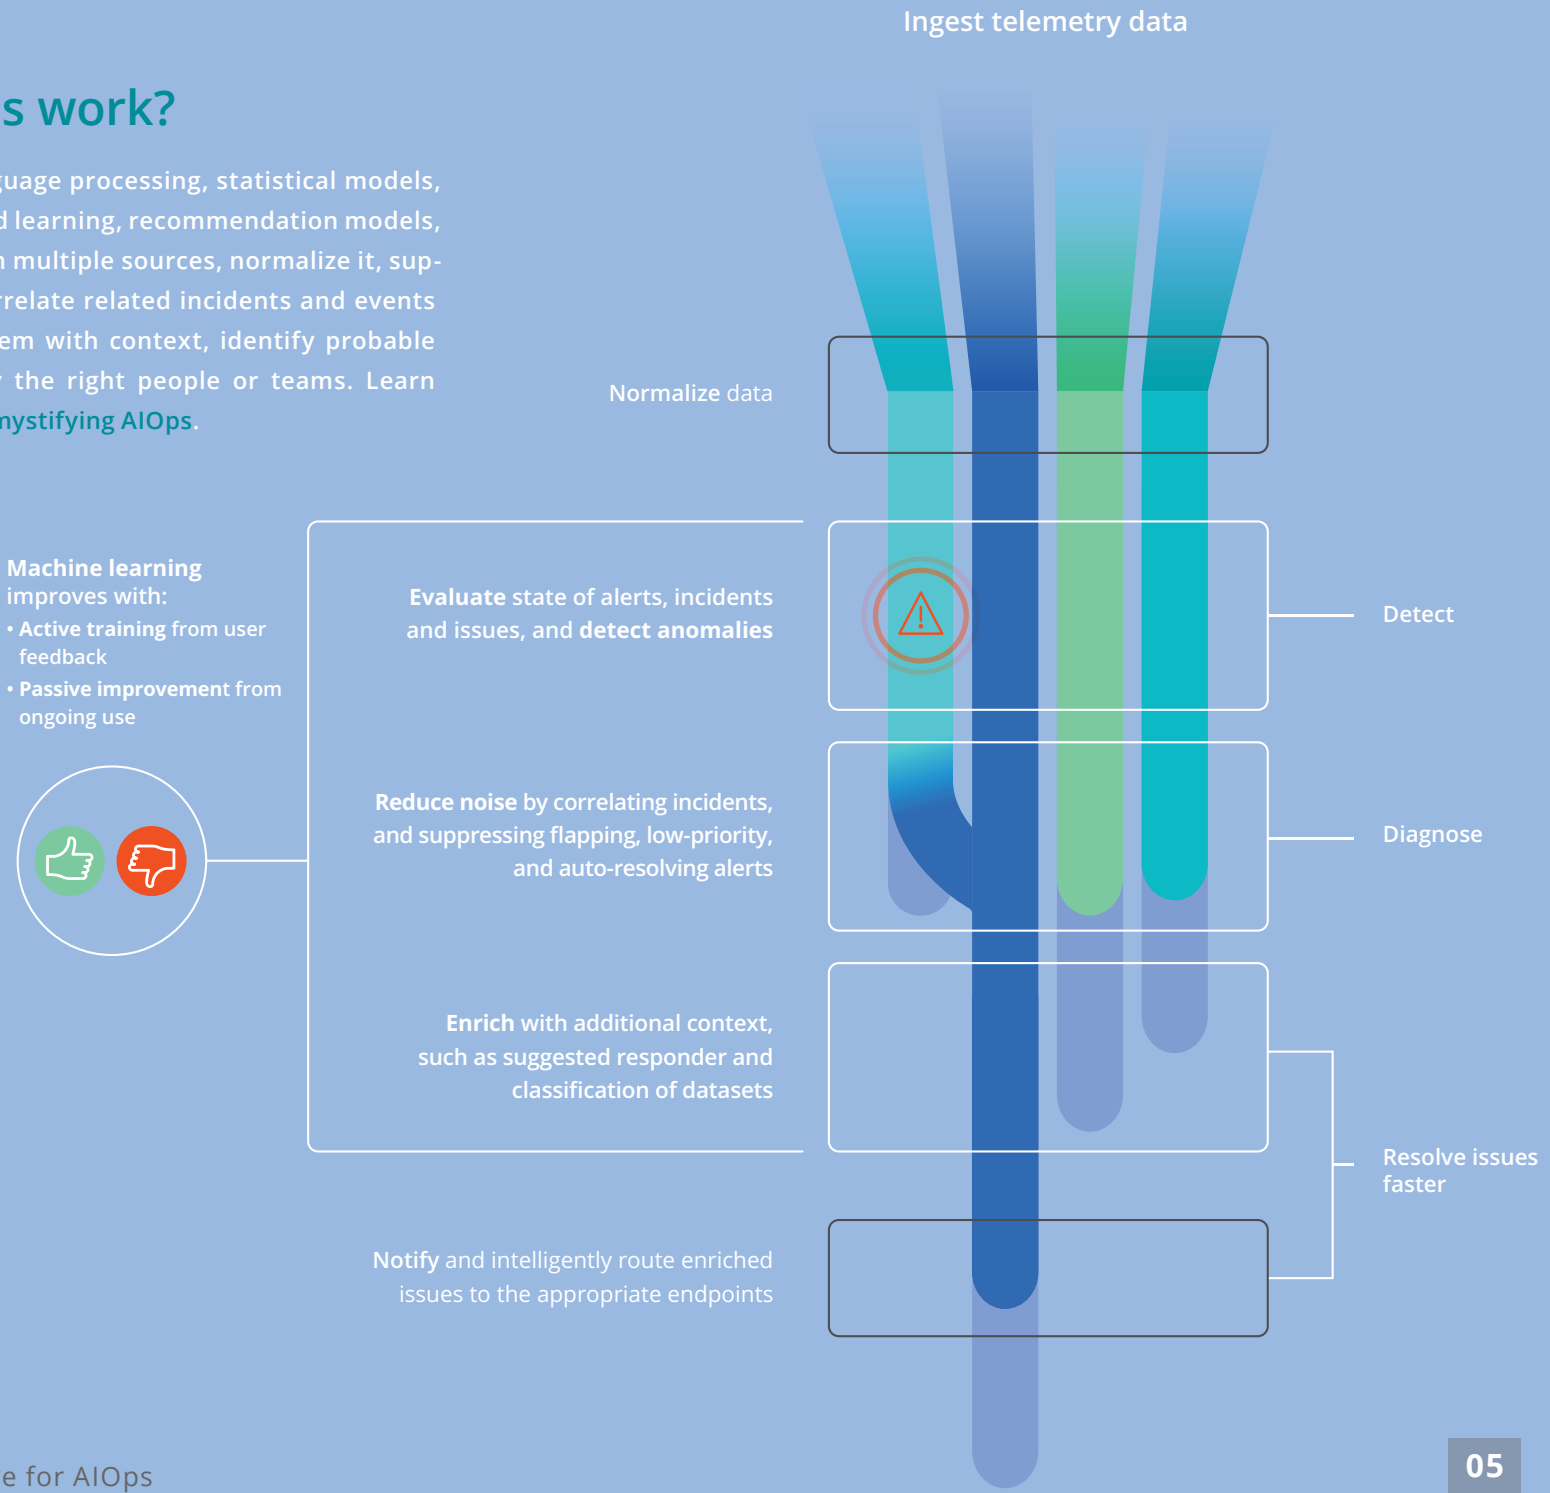

Take the Right Approach to AIOps

The buzz about AIOps tends to overshadow the most important benefit of AI: The tools learn and improve over time.

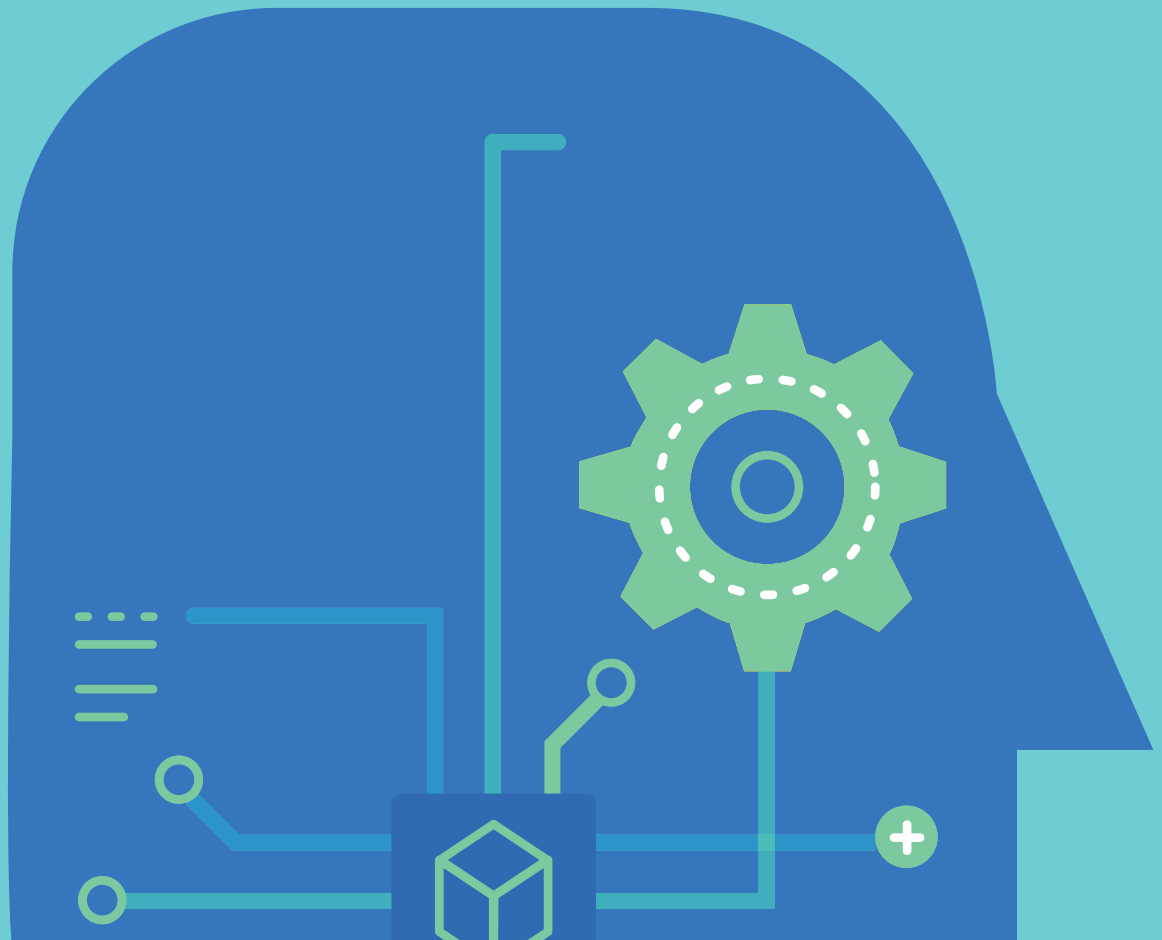

Proactive detection of anomalies and incident management become faster, more accurate, and more efficient the longer you use a tool.

That said, teams also need to see immediate value when deploying a new tool. Patience is a virtue, but IT and company leadership will expect a return on investment as quickly as possible. The good news is that the right approach (and the right AIOps solution, for that matter) can help you balance the expectations for rapid improvement with the desire to achieve continuous increases in value over time.

Let's walk through what that approach looks like, including the appropriate steps to prepare for and launch a successful AIOps effort that both delivers value early and grows that value over time.

Step 1: Identify the Problems

While figuring out whether AIOps could help your teams is a good start, it's time to take a closer look at your current situation. By understanding the problems that you need to solve, you can identify the capabilities you'll need from an AIOps solution.

While it's possible that you may need to solve only one or two of the following issues, chances are good that your teams face all, or nearly all, of these challenges. But don't fear—there is an AIOps solution for each of these problems.

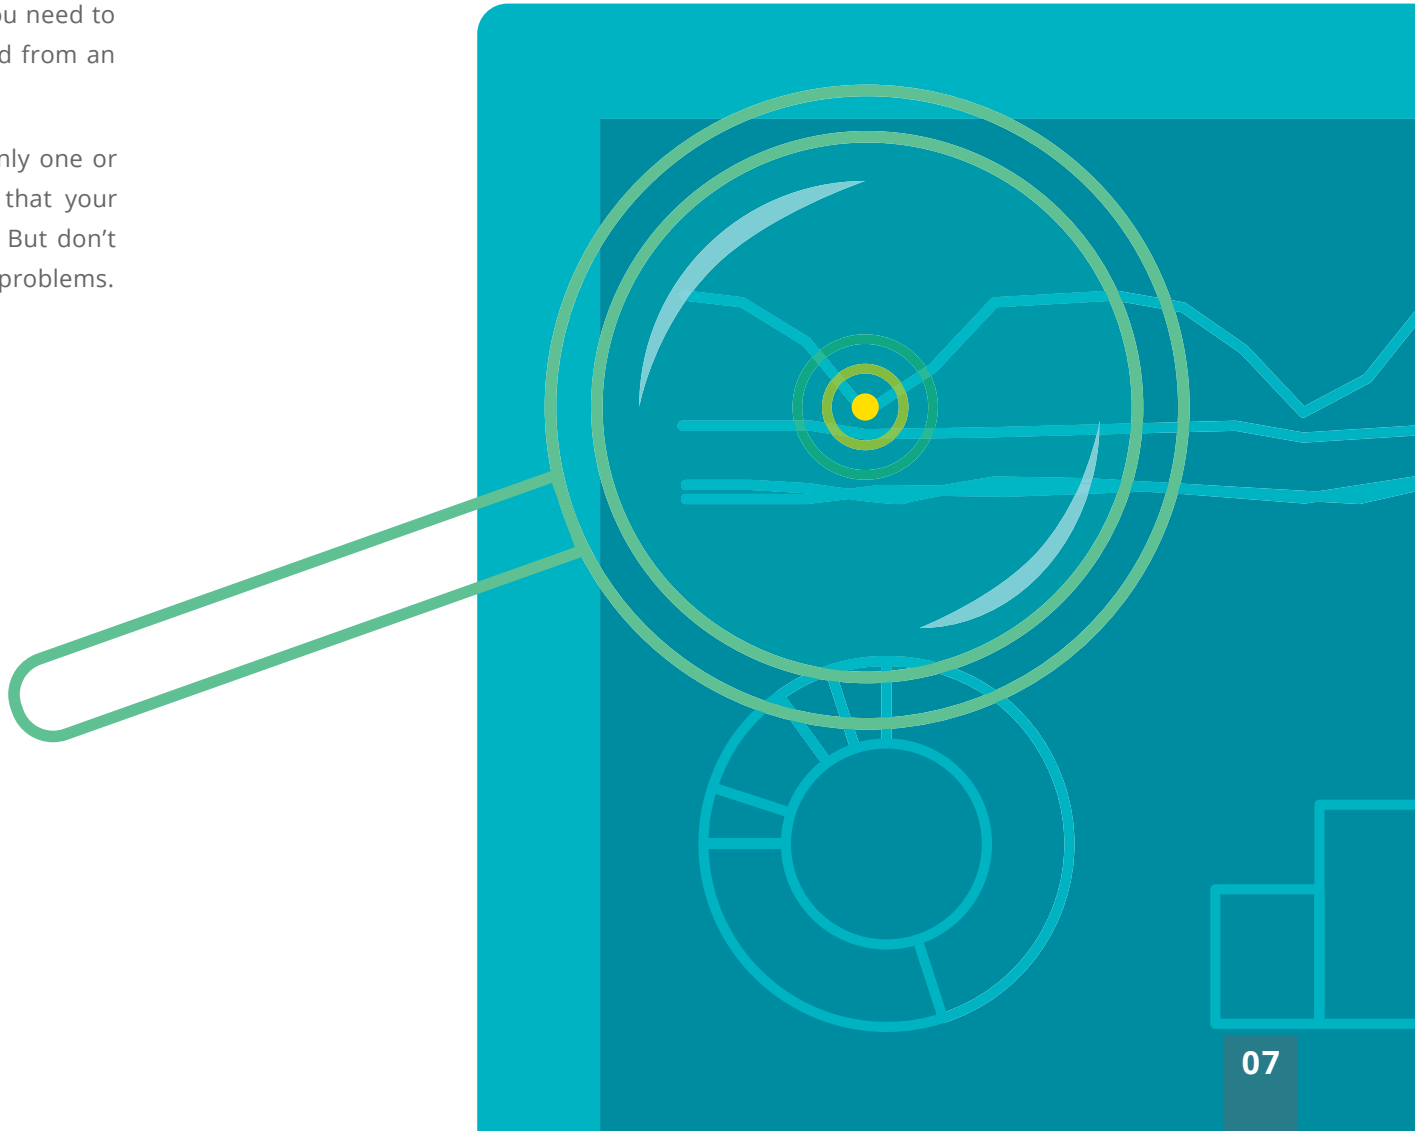

ASSESS YOUR AIOPS NEEDS

What's the problem?	Why does it happen?	Which AIOps capabilities can help?
Do you have unknown unknowns or areas that lack proper alerting?	For many teams, unknown unknowns are their biggest challenge. Incidents occur for which there are no alerts created or warnings surfaced beforehand.	**Proactive detection of anomalies**
Do your teams need to accelerate detection of issues?	Learning about issues from your customers is never good. Lack of alerts or alert noise means that your teams can't proactively understand and address potential issues before they impact customers.	Proactive detection of anomalies
Do your teams suffer from alert fatigue?	Thousands of alerts are impossible for a human to parse and understand. This makes it close to impossible to identify what alerts are most important and where your team should start so they can understand what's wrong and fix the issue.	• Alert correlation • Suppression of low-priority alerts • Detection of flapping alerts • Alert correlation and noise reduction • Probable root cause analysis
Are your teams slow to diagnose problems?	Fragmented, siloed, or redundant data and noisy alerts make it harder to find the right information to quickly diagnose incidents.	• Anomaly detection • Suppression of low-priority/auto-resolving alerts • Detection of flapping alerts • Alert correlation and noise reduction • Contextual alert enrichment • Root cause analysis
Is your team slow to respond to and resolve incidents?	If the investigation and troubleshooting process takes too long, your MTTR increases, as does the cost of downtime.	• Enrichment of issues • Delivery of issues within existing tools and workflows • Creation of suggested responders • Notifications sent to the correct team and individuals

Step 2: Understand Your Current Environment

Once you've defined the problems and the AIOps capabilities that you need to resolve them, it's time to gather relevant information about teams, processes, tools, and data involved in maintaining software reliability and managing incidents. The following questions can help you collect the information you need and understand the problems you face in more detail, so you can begin to surface which teams and systems within your organization could benefit the most from AIOps.

The answers to these questions will give you the knowledge of the tools, teams, processes, and data you need as you prepare to build out your AIOps implementation. Use this information to shape your success criteria in the next section, so you can minimize surprises down the road.

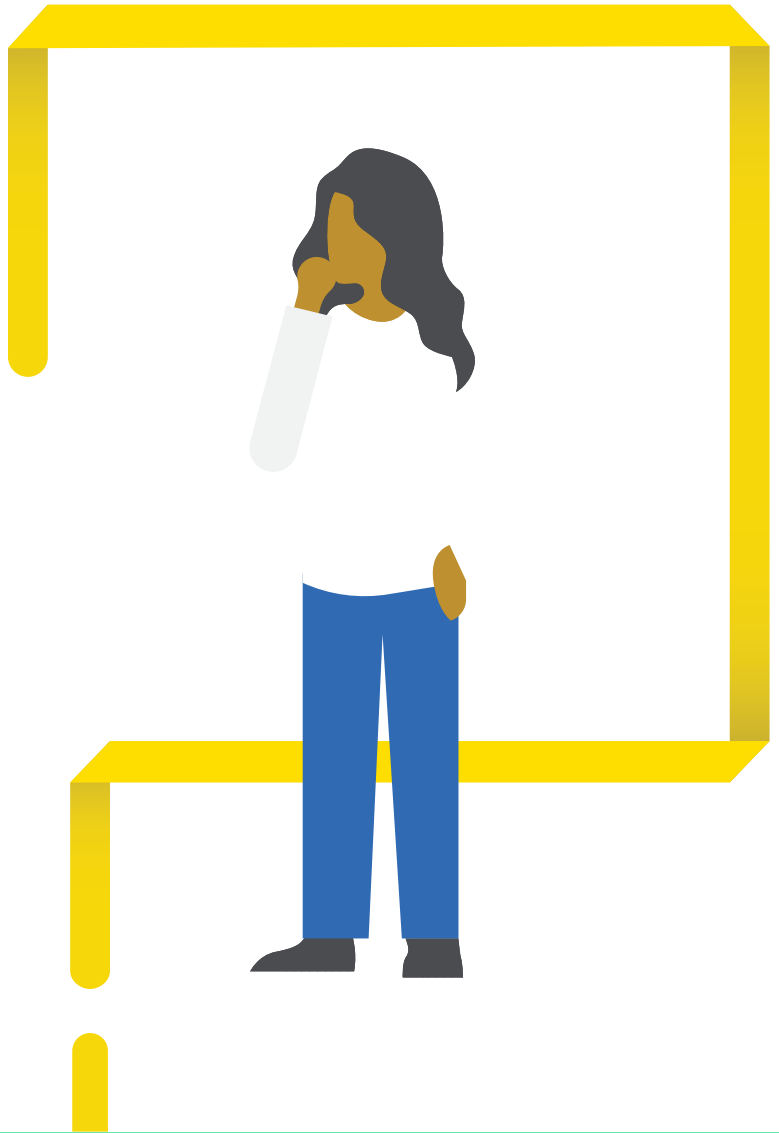

PEOPLE

- Which teams face the most severe challenges?
- Which teams could benefit the most from implementation of AIOps?
- Who are the key people to help with technical enablement?
- Who are the key stakeholders?

TOOLS

- What incident management tools are you using (e.g., New Relic Alerts, Nagios, AWS CloudWatch, Prometheus Alertmanager, PagerDuty, VictorOps, or OpsGenie)?
- In what order do your teams use the tools (e.g., an on-call engineer gets paged via the PagerDuty Mobile App then uses New Relic to find more details)?

DATA

- What are your sources of alerts and what types of alerts do you receive?
- How many of your alerts auto-resolve within 20 minutes?
- How critical are the alerts? Is there a way to prioritize or differentiate levels of criticality?
- What telemetry data do you gather? Is that data resilient to system changes?
- What parts of your system do you need to monitor but currently don't?

PROCESS

- Do you have a defined incident response process?
- What does your existing incident response workflow look like?
- Where does the workflow consistently slow down or break?

Step 3: Define Your Success Criteria

The next step on your path to AIOps is to define how you'll measure success for your initial and ongoing AIOps deployment. For instance, your first objective might be to reduce downtime for a specific system or component by reducing the number of unknown unknowns that occur. You would want to measure and track that objective by the number of incidents for which you were not alerted and for which your team had no previous awareness of the issue, as well as the total downtime due to those incidents. It's important to note what the current state looks like so you can use that as a baseline when measuring the success of your AIOps solution.

Think about the problems that you want to address and how long it will take to see improvements. Remember, while it's important to show results right out of the box, make sure your teams and management understand that the best AIOps solutions improve results over time, using machine learning and human feedback.

Don't just pick a solution and approach that gives you a quick win and nothing else. Aim for quick improvement upfront then a gradual, continual improvement and generation of value over time. Your success metrics should compare a baseline to any improvements you see, and reflect a reasonable timeframe in which you expect to see improvements.

Potential success criteria might include:

- MTTR
- Downtime/uptime
- SLO adherence
- Number of incidents generated per month
- Number of alert events generated per month
- Number of notifications received per month (this can be fewer than alerts)
- Number of unknown unknowns occurring per month
- Developer/Ops time spent on incident management

Don't just pick a solution and approach that gives you a quick win and nothing else.

Step 4: Determine Where to Start

This step is critical, because trying to solve every issue for every team at once increases the risk of slowing your time to value. You could also set the wrong expectations by tackling too many problems from the start, frustrating teams that don't see immediate improvements—particularly teams that would benefit the most from machine learning and improved accuracy over time.

To achieve the right balance of rapid results and to set the stage for more iterative improvement, it's best to first deploy anomaly detection and then introduce event correlation and issue suppression and enrichment. Not only can every team benefit from anomaly detection, but AIOps uses information about anomalies to improve the enrichment of issues with additional context to help pinpoint the root causes of issues.

Here's a suggested deployment scenario:

1. **Pick a team and configure notifications for anomaly detection.** You could choose a team responsible for important systems that are already reliable and that is interested in receiving early warning signs of trouble. Or you could choose a team responsible for back-end services that is occasionally informed of issues by internal or external customers.
2. **Expand to other teams.** Once you've demonstrated the value of AIOps for proactive anomaly detection and prevention of unknown unknowns for one team, you can expand to additional teams suffering from the same problems.
3. **Set up advanced AIOps workflows.** Choose one team to start using deeper AIOps functionality to improve and accelerate response. Criteria to look for includes teams that have:
    - Several services and infrastructure for which issues commonly result in cascading sets of alerts
    - Alerts coming from multiple alerting engines and multiple notifications for the same problem
    - A volume of alerts that creates fatigue and inhibits prioritization and action
4. **Expand the scope.** Once you have one team that is benefitting from fewer alerts and accelerated incident response, expand to more teams that have inadequate MTTR or difficulties pinpointing the root cause of problems.

As you deploy AIOps across the organization, steps three and four will be iterative. When you're ready to expand to new teams, you'll broaden your success criteria. When you implement new functionality, you should revisit your metrics again to make sure you're focusing on those that the new functionality can measurably improve.

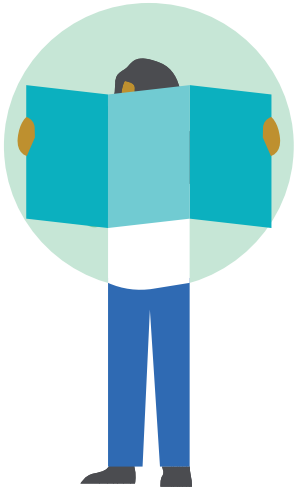

Ready to Learn More?

Excited to get started? Frankly, who wouldn't be, given the exponentially increasing complexity that teams face every month? Customers using New Relic Applied Intelligence get free proactive anomaly detection, and report automatic reductions in alert noise by 50%—with some reporting as much as 80% reduction within days.

Learn more about how **New Relic Applied Intelligence** can help your organization.