



AI IS A POWERFUL WEAPON IN THE FIGHT AGAINST IT PROBLEMS

By Alexandre Aractingi, Sylvain Duranton, Martin Pocquet, and Lucas Quarta

IT ORGANIZATIONS HAVE A problem with their problems. As uncertain times place new demands—and urgency—on business continuity, most IT departments are still taking a reactive approach to support, putting out fires instead of looking for smoke. And even when the alarm does sound, the issue is often unclear. Incident tickets bounce around IT, solutions lag, users get frustrated.

Enter artificial intelligence. AI in IT operations is an emerging area, and one of its most compelling applications is in problem and incident management. The idea is to leverage data to anticipate failures (or at least reduce the resolution time), better classify support tickets, and pinpoint underlying causes.

In our experience, companies can reduce their IT support costs by 20% to 30%, while increasing user satisfaction and giving time back to staff. But gaining the full benefit of AI means more than plugging in algorithms. It also means ensuring the availability and quality of data and making changes to processes and organization.

Miss those steps and the 20% to 30% savings become more like 4% to 7%.

Anticipate, Resolve, Improve

Problem and incident management will rarely win awards for ROI. While IT organizations typically allocate 10% to 15% of their budgets—and 10% to 20% of their manpower—to support, user satisfaction tends to be poor. And level 1 tech support, the initial tier where IT staff collect information and troubleshoot the more basic issues, rarely hits the 90% resolution rate considered best in class. Instead of ending with a fix, a call to support often just gets things started. Issues move up and down the chain, detours included.

The reactive approach to support doesn't just skew the balance between costs and returns, it puts a thumb on the scale. IT organizations start their troubleshooting late in the game, often when an incident's impact is already substantial. To quote the plumber standing in a flooded kitchen: "If only you'd called when the pipe was dripping."

Tipping the scale even further is the poor quality of many support tickets. Typically, companies will sort the vast majority of tickets—often north of 80%—into a handful of kitchen sink categories, like hardware or mail, that hide the problem’s complexity. (See Exhibit 1.) For example, printer and workstation issues will both land in the hardware bucket. So the inefficiency starts right up front. Instead of homing in on the specific problem and the people best positioned to solve it, IT teams enter a loop where they pass around tickets, seek more information, take a stab at a solution, and—when that doesn’t do the trick—start again.

Poor classification has another consequence. When you drop tickets into buckets that are not terribly meaningful, after-the-fact analysis is not feasible. It’s impossible to identify, at a granular level, the most important areas for remediation. The big-bucket approach also means that IT departments may not be allocating their budgets for evolutive maintenance—small enhancements that address day-to-day issues—in the most efficient way.

AI gets the thumbs off the scale. And it does so in a variety of ways. Different use cases let IT departments anticipate issues, solve problems more effectively, and im-

prove the way they manage and prioritize support. (See Exhibit 2.)

Some use cases are more sophisticated and more complex to implement than others, but companies don’t need to take an all-or-nothing approach. They can implement AI in a gradual way, generating value—and buy-in—with an initial wave of applications and then build on that foundation.

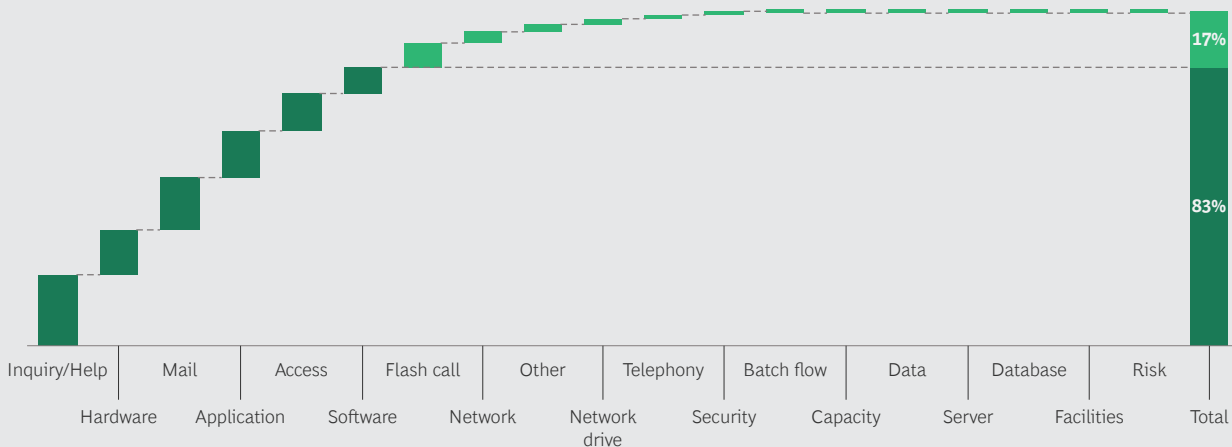
Support Ticket Classification

AI provides options and flexibility. Companies can employ different methods and algorithms to solve different kinds of problems. So where to start? We’ve found that support ticket classification is a particularly ripe area for AI enhancement. If an IT department could classify tickets more precisely, it would be able to identify connections and patterns across tickets—insights that would help it detect emerging trouble spots and prioritize remediation efforts.

Better yet, if that classification happened in real time, with the algorithm working even as the support agent enters the caller’s description, the ticket could be routed to the right place immediately, instead of bouncing around and triggering delays and dissatisfaction.

EXHIBIT 1 | Kitchen Sink Buckets Mask the True Problem

Companies typically sort more than 80% of their support tickets into six broad categories



Source: BCG case experience.

EXHIBIT 2 | AI Enhances Problem and Incident Management in a Variety of Ways



Anticipate

- Detect emerging larger issues (the hidden “big” problem) from support tickets
- Detect system failures from multivariate analysis
- Detect incidents early from weak-signal analysis



Solve

- Improve support ticket classification
- Enable smart rerouting of support tickets
- Predict support team workload
- Enhance forensic analysis for better root cause detection
- Classify incidents by system, user cluster, and other criteria



Improve

- Automate recurring-incident management
- Prioritize support tickets
- Detect patterns in support tickets
- Analyze correlations among incidents, tickets, and logs to enhance testing

Source: BCG case experience.

Note: List is not exhaustive; it represents some of the major use cases for AI in problem and incident management.

While not an exhaustive list, the following seven use cases spotlight how AI can help companies better manage and leverage support tickets:

- Improve support ticket classification (enabling granular categorization and correlating related tickets).
- Classify incidents by system, site, user cluster, and other criteria (segmenting tickets as needed for different types of analysis).
- Prioritize support tickets.
- Enable smart rerouting of support tickets (in real time).
- Predict support team workload.
- Detect patterns in support tickets.
- Detect emerging issues (using the patterns identified in the previous use case to recognize, in real time, when an issue is reappearing or when something is behaving in an unusual manner).

How exactly do you do these things? Two data science techniques are key: natural language processing (NLP) and machine learning.

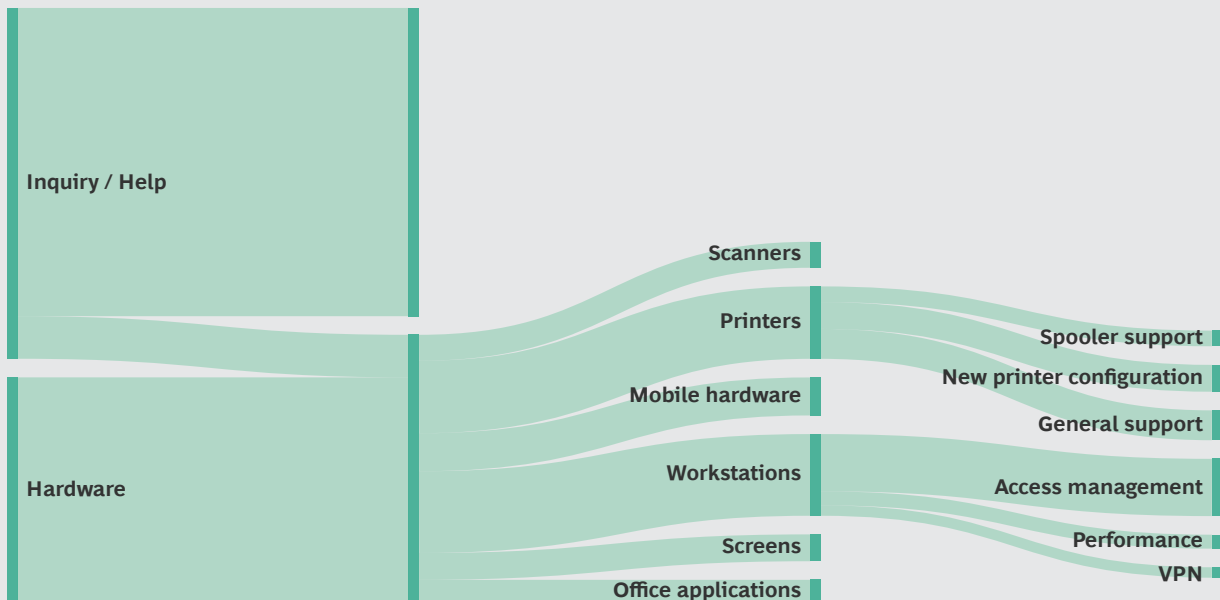
NLP algorithms analyze free-text descriptions in support tickets to identify topics and connections. This lets you classify tickets into subcategories and direct issues to the most appropriate team, accelerating resolution. (See Exhibit 3.) It also lets you create clusters—tickets linked by common keywords and themes. These can reveal underlying issues that warrant attention.

When you apply NLP to past as well as current tickets, you can create a rich historical record of incidents, categorized and linked in very specific ways. Then, by using machine learning, you can identify ticket distributions that have presaged emerging issues. For example, perhaps you have discovered that in the past, a certain pattern of reports preceded a sharp rise in tickets connected to a workstation issue. If you could detect, in real time, when you started to see this pattern, you could sound an alarm well before the ticket peak arrived. This would let you work the problem earlier, but it would also let you to take steps—like putting a “We’re aware of the issue” message on the call center line—to reduce the number of incoming tickets.

Incident Anticipation

AI also enables a more robust and precise approach to incident anticipation. Tradi-

EXHIBIT 3 | AI Can Automatically Classify Tickets into More Specific Categories



Source: BCG case experience.

tionally, IT departments have relied on instinct and experience. AI relies on data. Most companies maintain technical logs that capture hundreds of performance metrics: CPU usage, memory usage, network bandwidth, and so on. This rich trove of information can fuel sophisticated pattern recognition.

The idea is to use machine learning algorithms, including regression, clustering, decision trees, and deep learning, to identify correlations between log data and past incidents. This lets you understand the signal pattern that foretells an anomaly and sound a warning when the problem is about to raise its head. The result: an anticipation window—whether 30 minutes or 30 hours—for mitigating the impact of incidents or preventing them entirely.

But what about a new type of incident? Without a known pattern to detect, how can you create and benefit from an anticipation window? The answer lies in an AI concept known as unsupervised learning. While you may not have a pattern to match against, the algorithm still can detect anomalies in the data—something that raises eyebrows simply because it's unusu-

al. The process may lack precision, as the anomalous data isn't linked to a specific problem, but it provides a heads-up that something warrants investigation.

You can then feed these new events back into the AI to identify the warning patterns and create alarms. Indeed, one intriguing use case for AI is forensic analysis of complex cases: applying algorithms to understand the root cause of a new event so you can react faster and better the next time.

Give the Algorithms a Boost

Of course, the best data in the world won't mean much if it can't be pulled into the algorithms that need it. At many companies, data is scattered across the enterprise—stored within specific systems and hard to access and use for other purposes. This is where a data platform proves invaluable: a mechanism for storing and integrating incident and performance information, past and present.

Still, algorithms and technology only get you so far. Flip the switch at this point and you're likely to see fewer incident tickets and a 4% to 7% reduction in IT support

costs. That's not trivial but it's low-hanging fruit. The real boost comes when you also make changes to your operating model—changes that foster robust data governance (so your data is accurate, consistent, and available to the algorithms that need it), adapt roles and processes (so you leverage insights more efficiently), and optimize how you work with external providers (so they only do what they need to do—which, with savvy use of analytics, will likely be less).

Take these steps and the savings soar. We've seen companies gain an additional 10% to 14% cost reduction by improving process and team efficiency and a further 6% to 9% reduction through decreased external provider costs. So now overall savings jump to 20% to 30%. And that's on top of user satisfaction boosts and all the time you're giving back to employees—both within and outside IT—who are now contending with fewer incidents.

The added value that the right operating model brings isn't headline news. In our work in artificial intelligence, we often talk about the 10/20/70 rule: AI is 10% about algorithms, 20% about technology, and 70% about business process transformation. To get the maximum benefit, you need to understand what kind of insights the algorithms generate. Then you need to think about what changes, to your processes and organization, you could make to leverage those insights.

For example, you could tweak processes so that incoming tickets, once analyzed, are steered directly to the right people. You could change the responsibilities at each support tier so that the people at level 1 are more empowered to solve problems, while those at level 3—who are in charge of the actual applications—spend less time on monitoring (freeing up time for more value-creating activities). You could also make less essential but still highly beneficial adjustments, such as integrating the results of ticket analysis into your process for prioritizing evolutive maintenance. This would let you more effectively allocate funding and shorten, or even eliminate, the usual debates over where to put the money.

Putting It All in Place

By prioritizing use cases, you can start reaping the benefits of AI quickly—in as little as three months if you know how you want to use AI and can access the relevant data. Contrast that with an all-encompassing big-bang approach, where you may wait two years for a grand unveiling.

Meanwhile, by prioritizing high-value use cases, you visibly demonstrate the benefits of AI. This helps build support and funding for a continuing effort and for the necessary changes to processes and organization. So instead of starting with an algorithm and deciding what problem it can solve, start with a big pain point and decide what algorithm, using what data, can solve it.

This kind of progressive approach also lets you deploy your target operating model in a gradual, value-driven way. Use cases and operating model develop in parallel and in sync.

How does all this work in practice? Typically, the effort spans three phases:

- **Ignite.** In this initial phase, IT organizations analyze their pain points and prioritize use cases. They identify the data they can leverage and how to integrate it with the data platform (this is where robust data governance is essential to ensuring availability, access, and accuracy). At the same time, IT defines a target operating model. Steps here include creating an organization schema for level 1, level 2, and level 3 support; describing key processes; and examining the current sourcing model.
- **Prototype.** Next, IT starts to prototype solutions for the first wave of use cases. The idea is to test the AI in a real-world environment to make sure the algorithms and new processes work as expected.
- **Incubate and industrialize.** After developing a proof of concept, IT brings the AI and the processes around it to scale in order to automate the solution and realize its full potential. With the

first wave of use cases “industrialized,” IT moves to the next wave and repeats the process.

table stakes for growth, an integrated, scalable data architecture should be on every company’s to-do list.

A data platform is a key enabler for launching use cases, and deploying one (should it not already exist) is no trivial undertaking. But a data platform can also be developed in an incremental way, where you build what you need for a specific use case and build more with each subsequent use case. The benefits of such a platform go beyond AI. With data-driven decision making now

WHEN DEPLOYED WITH the right changes to processes and organization, AI can boost the uptime of your systems and the productivity of your people. It can lower the costs of incident management and build resilience. It does all these things by giving you clarity on IT issues—and even, perhaps, a bit of clairvoyance.

About the Authors

Alexandre Aractingi is a managing director and partner in the Paris office of Boston Consulting Group. You may contact him by email at aractingi.alexandre@bcg.com.

Sylvain Duranton is a managing director and senior partner in BCG’s Paris office and the global leader of BCG Gamma. You may contact him by email at duranton.sylvain@bcg.com.

Martin Pocquet is a lead data scientist at BCG Gamma in the firm’s Paris office. You may contact him by email at pocquet.martin@bcg.com.

Lucas Quarta is a partner and associate director in BCG’s Paris office specializing in data and analytics. You may contact him by email at quarta.lucas@bcg.com.

Boston Consulting Group partners with leaders in business and society to tackle their most important challenges and capture their greatest opportunities. BCG was the pioneer in business strategy when it was founded in 1963. Today, we help clients with total transformation—inspiring complex change, enabling organizations to grow, building competitive advantage, and driving bottom-line impact.

To succeed, organizations must blend digital and human capabilities. Our diverse, global teams bring deep industry and functional expertise and a range of perspectives to spark change. BCG delivers solutions through leading-edge management consulting along with technology and design, corporate and digital ventures—and business purpose. We work in a uniquely collaborative model across the firm and throughout all levels of the client organization, generating results that allow our clients to thrive.

© Boston Consulting Group 2020. All rights reserved. 10/20

For information or permission to reprint, please contact BCG at permissions@bcg.com. To find the latest BCG content and register to receive e-alerts on this topic or others, please visit bcg.com. Follow Boston Consulting Group on Facebook and Twitter.