BUYER'S GUIDE

Making sense of the AIOps landscape

A hands-on, practical approach to a successful AIOps journey

An IT Ops community resource brought to you by BigPanda and Windward

BigPanda | windward

Table of contents

Foreword

IT Ops community leaders,

I'd like to extend my deepest thanks to BigPanda for the opportunity to collaborate with them on their AIOps Buyer's Guide.

For the last few years, Windward has worked to deliver AI-optimized IT solutions that help our clients find flow in their IT operations. Our partnership with BigPanda has been vital in helping many of our clients transform their vision into reality.

We have worked to ensure that the core audience for this guide, namely IT Operations leaders and visionaries, will find value in the concepts described because like you, we don't just speak IT operations, we live it. This guide isn't just a collection of observations. The team at BigPanda provides a clear picture of the AIOps journey along with experience-driven insights on how to navigate your journey successfully.

Throughout this guide, we've also added our perspectives, experience, insight, and access to additional resources. We hope you find these helpful in your AIOps initiatives, both in clearing the way for implementation and in promoting collaboration between the stakeholders in your organization. After all, turning on a new technology or platform does nothing if the work comes up short when it comes to defining vision, bridging teams, training stakeholders, and aligning everyone across the board.

As always, if our team of AIOps experts at Windward can help you along the way, we're more than happy to join you on your path to AIOps.

Sincerely,

Sean McDermott
CEO, Windward Consulting Group
Host, AIOps Evolution Weekly

Introduction

More than ever before, today's users require near-continuous availability. Customer-facing services, internal corporate systems, and back-end services are treated as utilities. Today, network connectivity and healthy, functioning applications are as fundamental to most businesses as electricity.

For the teams that own service availability, this means not only configuring for high availability, but also maintaining proactive 24/7 awareness of service health, service-related changes, and effective monitoring so that teams can detect the smallest of issues and intervene — before users are ever impacted.

Whether operations are centralized in an OpsCenter (or a NOC), or distributed across several SRE or DevOps-based teams, someone has to constantly monitor the health, availability and performance of key services, and determine when and how to take action to maintain service uptime and SLAs.

Fragmented operations

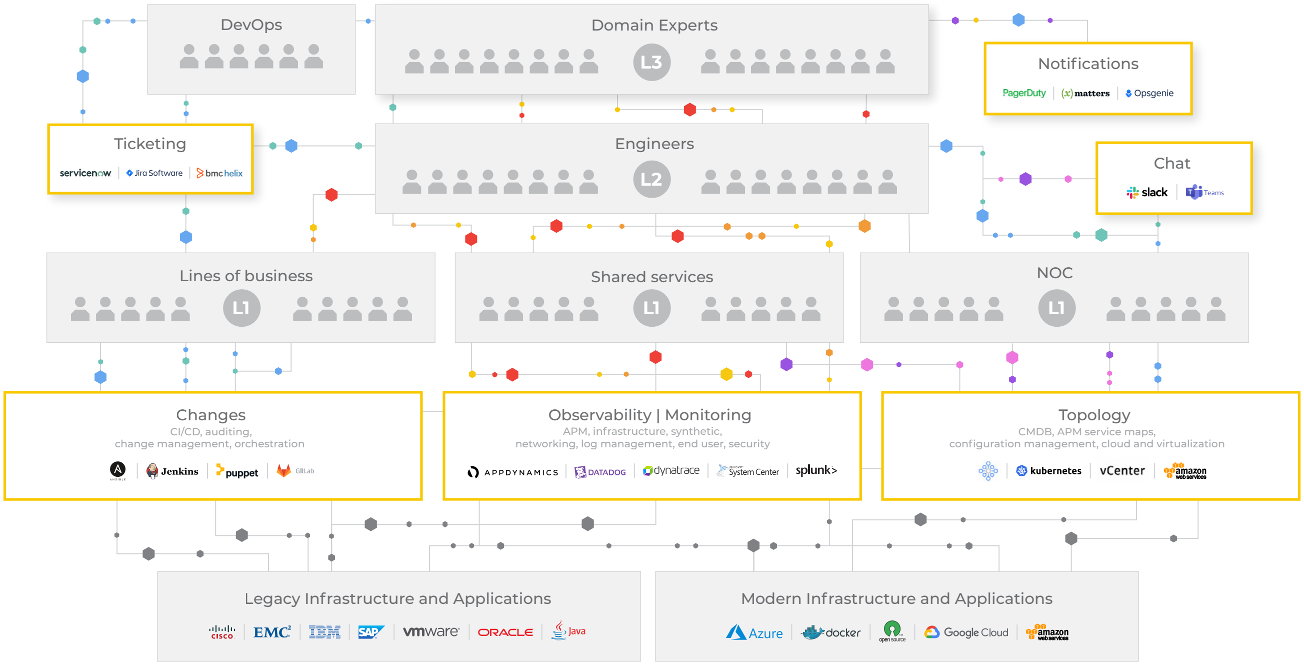

The challenge

Most organizations view monitoring as a natural extension of the engineering domains that make up a modern service architecture. Network, infrastructure, database, and application engineers in modern IT environments use a variety of observability and monitoring tools when building their applications and services and, once ready, these engineering teams transition those methodologies into the production environment under a "monitoring" fabric.

Multiple observability tools and all the tribal knowledge embedded within them, can lead to information silos, excessive event or alert noise in some areas, and insufficient monitoring coverage or fragmented visibility in others.

The usual blend of legacy, modern infrastructure and applications can lead to unknown failure points and gaps in awareness. Awareness of interdependencies between those domains is often lost, and only discovered after services fail, reactively and painfully.

Subsequent recovery effort can be quick and clean, but is often slow and uncertain, and in all cases there is some cost associated with the downtime. Users are frustrated, engineering velocity is sacrificed, reputational damage is done, or eCommerce pipelines stop producing revenue.

Whether measured or not, the organization and its operations teams pay a steep price for downtime.

The breaking point for IT operations

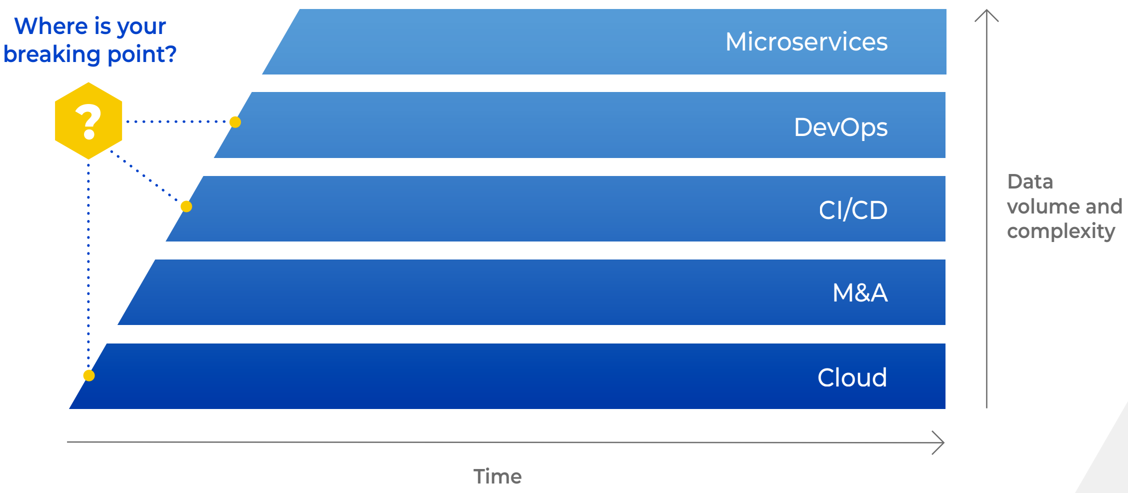

A growing problem

The challenge for IT operations teams is becoming worse by the day.

Data volume and service complexity are only increasing as organizations transform to take advantage of cloud scalability and cost, refactor to use microservices architectures, embrace CI/CD change velocities, or simply extend and expand services into new regions with new features.

Organizations are caught in a lose-lose scenario.

On the one hand, the bureaucracy of ITIL can mire organizations in attempts to boil the ocean with standardization and consolidation efforts. On the other hand, a DevOps framework can lead larger organizations to duplicate teams and functions, causing excessive spending on tools, information silos, incomplete or inconsistent awareness of actual service availability, and other sub-scale efficiencies.

This is the environment that CIOs and/or CTOs are forced to confront, often after a major availability failure. Sadly, this is the environment that IT operations teams and SREs must deal with constantly.

Therefore, it's no surprise that many have started to look to AIOps to provide a solution. In response, many solution-providers have incorporated the term "AIOps" into their tools and offerings, regardless of the underlying technology and presence or effectiveness of Artificial Intelligence and Machine Learning (AI/ML).

Sorting through the marketing and getting something that works can be a minefield for IT Operations leaders looking to move to a next-generation operations pipeline, and some solution providers are all-too-ready to sell whatever they can, for as much as they can, regardless of the organizational fit, or the value they think they can deliver. This brief guide will help ensure success with AIOps.

What is AIOps?

> "There is no future of IT operations that does not include AIOps. This is due to the rapid growth in data volumes and pace of change (exemplified by rate of application delivery and event-driven business models) that cannot wait on humans to derive insights."
>
> – Market Guide for AIOps Platforms, Gartner Research, April 2021

The AIOps Evolution Podcast is a weekly deep-dive on everything AIOps. Get rich insights into ML, automation, and actionable tips for tech leaders and their enterprise teams.

Learn more at aiopsevolution.com

AIOps is the holistic approach that addresses the IT Operations challenge as a whole, not point-by-point, or team-by-team. It is an operations pipeline built around data management, automation, and AI/ML-based processing of events, changes, and topology to drastically reduce operational workload and improve incident outcomes.

AIOps is defined differently by different organizations. An AIOps solution should:

- combine inputs from multiple monitoring, topology, and change systems
- enrich and correlate them using Artificial Intelligence / Machine Learning
- provide actionable outputs to collaboration, remediation, and reporting systems, automating some or all of the workflow involved

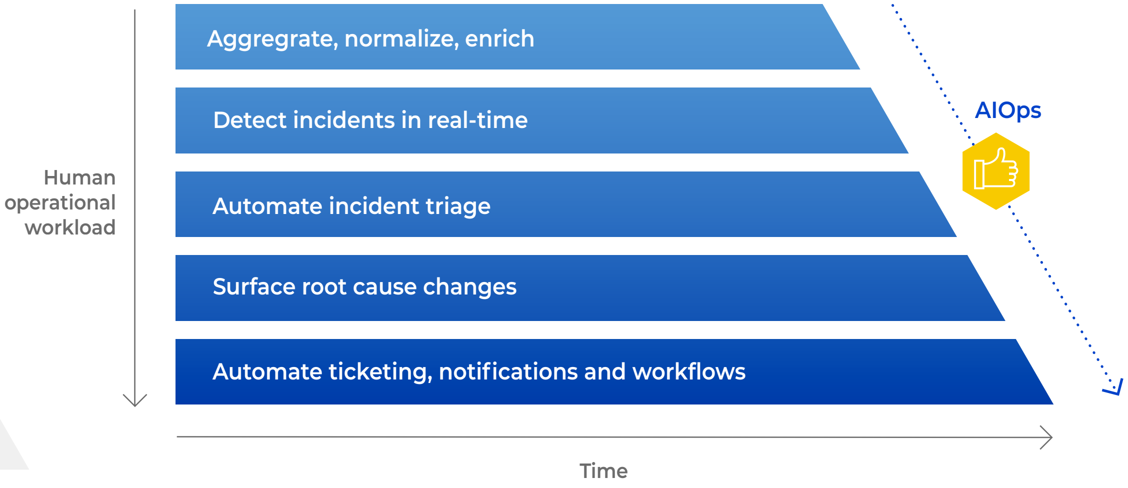
Aggregrate, normalize, enrich
Detect incidents in real-time
Automate incident triage
Surface root cause changes
Automate ticketing, notifications and workflows
AIOps
Human operational workload
Time

Gartner, in its Market Guide to AIOps Platforms says “AIOps is the application of machine learning (ML) and data science to IT operations problems. AIOps platforms combine big data and ML functionality to enhance and partially replace all primary IT operations functions, including availability and performance monitoring, event correlation and analysis, and IT service management and automation”.

There are two approaches taken by AIOps providers, which Gartner labels as “domain-centric” and “domain-agnostic”:

- **“Domain-centric”** AIOps solutions augment or in some cases replace your existing monitoring, service management or collaboration systems by adding domain-specific intelligence, generally in the form of anomaly detection. While this can be effective if your IT Operations involve only a few domains, as use cases are added, it requires significant time to retool, retrain, and standardize processes across teams. It also makes vendor lock-in a significant risk.
- **“Domain-agnostic”** AIOps solutions leverage the outputs from your existing monitoring, topology, and change systems, and focus on processing that data effectively and automating IT Operations workflows. This “plug and play” approach allows teams to keep their existing tools, making organizational adoption a much lighter lift. This has become the dominant approach in AIOps as it is easier, delivers actual value more quickly and reliably, and is highly adaptable or ‘future-proofed’.

Assessing your operational readiness

The first step in being *pragmatic* about purchasing an AIOps solution is to honestly assess your operational maturity today. A good self-assessment will be useful for any organization that is trying to improve both operational performance and operational workload.

Below is a condensed maturity assessment that covers monitoring, event processing, incident management, and topology and change awareness.

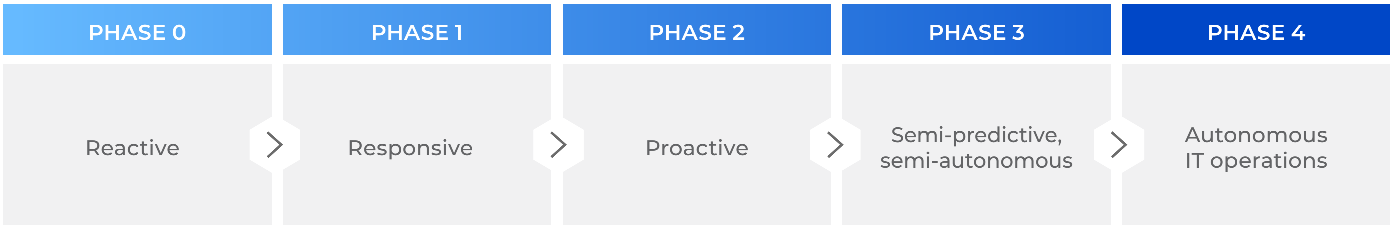

PHASE 0	PHASE 1	PHASE 2	PHASE 3	PHASE 4
Reactive	Responsive	Proactive	Semi-predictive, semi-autonomous	Autonomous IT operations

Phase 0: reactive

IT Operations functioning is uncoordinated. Monitoring is nonexistent or siloed within domain teams. Incident management, change, and topology processes are ad-hoc, undocumented, or team-specific. There is a prevalence of user-reported issues at high priorities, and unmeasured service availability.

Phase 1: responsive

The organization may trend towards centralized or decentralized operations. There is limited centralized monitoring and high noise. User reports are ground truth. There is a documented incident management process, but low trust between teams leads to limited information sharing during diagnosis and RCA. Monitoring, topology and change processes are generally siloed. Subscale efficiencies exist in larger organizations, with an inability to rapidly and effectively detect and remediate service issues.

Phase 2: proactive

A partially integrated operations pipeline is tied to a continuous process. Extensive monitoring coverage and data volume exists, but generally this is "peak noise," low monitoring actionability, excessive low priority incident volume, and peak resource usage. Event processing is generally rules-based and consistent, but unsatisfactory. Topology is mostly consolidated with some known dependencies between locations, hosts, applications, and services. Most changes are logged and major ones are centralized for general awareness.

Phase 3: semi-predictive, semi-autonomous

A common operations pipeline connecting multiple monitoring, change, and topology data sources provides distilled outputs to collaboration/ITSM and reporting systems. Workflow/task automation exists throughout the incident process. Extensive monitoring coverage and good quality works across domains, with use of anomaly detection for metrics/logs, Service Level Indicator (SLI) / Service Level Objective (SLO) error budgets, and client telemetry/synthetics. Event enrichment and ML-based correlation reduce incident volume. Maintenance based event suppression reduces noise. Topology data is used for root cause triangulation including causal changes. Limited auto-remediation is tied to specific incident scenarios.

Phase 4: autonomous IT operations

A centralized pipeline handles all monitoring integration, topology mapping, change and alert/event data normalization, enrichment, correlation, remediation, and reporting. Outputs can be decentralized to a variety of auto-remediation systems. Achievable, documented monitoring requirements are met for all services, including minimum payload, coverage, and actionability. Modular anomaly detection tools are used across layers/disciplines. There is flexibility in monitoring and remediation/automation tools and systems, consistency is standard. All events/alerts associated with a common root cause (including changes) are correlated, using machine learning. Integrated, automated incident handling persists from detection to remediation. The primary pipeline output to operations staff is now outcome reports that are used for process and engineering improvements aimed at service resilience and availability. All other functions are executed autonomously.

Understanding your IT Ops maturity gaps

It's important to recognize where your organization is on the operational maturity curve, as well as your desired destination.

The goal is not to process more alerts, fill out ticket fields more quickly or identify innocence/place blame more accurately. Zoom out, and identify high-level, holistic goals: improving service availability, minimizing operational workload and reducing operating costs, through the inevitable periods of digital transformation your organization will undertake over time.

The destination is autonomous IT operations; a highly-efficient pipeline that combines automation and human expertise, with seamless integration between monitoring events and remediation systems, leveraging real-time topology data and change awareness to detect, diagnose, and remediate problems proactively and automatically, before users are impacted.

Assess your organization and prioritize the areas where you are lagging, because maturity gaps can create issues. A robust incident management process won't help reduce MTTR if you have noisy monitoring and limited event processing; your people will drown in non-actionable incident volume.

Similarly, automating immature processes, such as opening up a chat channel, or incident assignment and escalation, can be a recipe for disaster; automation is great at doing things very quickly, whether they are the right things or not. Take care to focus on automating processes that deliver clear value.

Looking to get up to speed quickly? Windward has a free Guide to Investing in AIOps.

Visit windward.com/aiops-guide to start exploring

IT Operations Maturity Model

	PHASE 0	PHASE 1	PHASE 2	PHASE 3	PHASE 4
Maturity state	Initial	Managed	Defined	Quantitatively defined	Optimized
Descriptive State	Reactive	Responsive	Proactive	Semi-Predictive, Semi-Autonomous	Autonomous IT Operations
KPI: Service Availability	Unmeasured / untrusted, low awareness, low expectations	Measured but incomplete due to coverage gaps, perceived as poorby users	Measured, largely trusted, and generally improving	Measured, and improvement is connected to specific upstream service operations initiatives	Measured and directly connected to service changes, remediations automations speed improvements, monitoring improvements

AIOps tools, acting as a central hub, help you identify gaps throughout your IT Operations pipeline with the use of analytics. The key is to assess the maturity of the entire pipeline and prioritize efforts where they will meaningfully improve results; don't over index on a single function; it must be a coordinated effort.

Responsive or higher? It may be time for AIOps

If you are at the "Responsive" level or higher across one or more categories, it may be the right time to look at AIOps. Peak noise and peak headcount usually occur in the "Proactive" phase; this is the limit that many organizations reach where good process and capable people are no longer enough; a new strategy is needed.

Another scenario to consider is when you've been tasked with building a next-generation OpsCenter as a result of consolidation or M&A activity; in that case, an AIOps solution can provide a solid foundation and allow you to run fast and light as you scale up.

In all cases, identify your starting point, and make sure to use your metrics to keep score as you evolve your people, process and technology.

AIOps solutions: desired characteristics

Now for some specifics: what key requirements should you look for in an AIOps solution?

Your AIOps solution will be a centralized hub for ingesting web-scale monitoring, topology, and change data, as well as the system that has to provide human-actionable outputs to your collaboration, remediation, and reporting systems.

It will also serve centralized and decentralized teams in your environment; it will serve multiple business units and departments; it may even serve multiple brands that are part of the same organization. Hundreds or even thousands of users may come to rely on it every day, to serve their customers and users without interruptions or problems.

Given that, here are the desired characteristics for your AIOps solution:

> "Our clients have seen positive ROI from adopting AIOps, in terms of lower MTTR and improved efficiency. But starting with strategy, not tools, is key."
>
> – Sean McDermott, CEO, Windward Consulting Group Host, AIOps Evolution Weekly Podcast

1. Ease-of-use (by anyone who touches incidents)

Whether your IT Operations teams are centralized in an OpsCenter, decentralized in a DevOps model, or somewhere in between, your AIOps solution has to be usable by administrators, operations engineers, SRE's, application developers, network engineers, and everyone else involved in IT Operations processes.

An intuitive presentation layer with excellent data visualization is desirable, but also look for straightforward, readable training in the form of how-to guides, video instruction, and a strong user community. The technical knowledge required to use your AIOps solution should not require a raft of new hires with years of experience. That will create a chokepoint in adaptation and velocity.

For both administrative and power users, easier and more intuitive is better. The more your teams can do for themselves, the faster and more adaptive they'll be.

Sample questions to ask when evaluating AIOps tools for ease-of-use:

- How long does it take to train and onboard new users?
- Does the platform offer role-based access to different datasets inside the platform?
- How many clicks does it take to perform common actions?
- Which personas is the product designed for? Expert developers and L3 teams, or L1 and L2 teams?

2. Self-service configuration (no experts required)

Running a good operations pipeline requires significant integration across multiple tools and domains, and clear observation of the inputs and outputs at each stage. This can't be a complex, heavy lift for your teams, or you'll never get data flowing satisfactorily; the barrier to entry will be too high. Further, in modern IT environments, tooling stacks constantly change and evolve. Existing integrations may need to be tweaked, and new ones built, on a regular basis. Your teams must be able to perform these actions on their own, and easily.

Look for intuitive configuration controls and diagnostic tools with clear outputs. AIOps is a means to reduce complexity, not increase it. Users should be able to easily test new integrations, configurations, or automations, see the outputs, and adjust.

Avoid solutions with AI or ML black boxes, where there are not clear explanations or a data science PhD is required to understand how to use them. A deep understanding of fluid dynamics is required to make jet engines, but pilots control those engines with relatively simple throttle inputs, and they can observe (and feel) the resulting outputs.

Sample questions to ask when evaluating AIOps tools for self-service:

- Does the platform allow end-users to set up, debug and fine-tune integrations on their own, without calling on experts?
- Does the platform provide UIs to allow users to perform common actions such as creating enrichment maps, without calling on experts?
- Does the platform allow users to create or modify correlation patterns on their own?
- Does the platform allow users or admins to create dashboards and reports?

3. Rapid time-to-value (weeks, not years)

Unlike the 1990s and the 2000s, organizations today can't afford to wait for 12 or 18, or even 36 month deployment cycles. Markets and customers today expect to see value in months or even weeks.

By addressing the inherent challenges of data volume and complexity in modern organizations, an AIOps solution should be fast to get up and running and keep running through technology, process, and people changes.

These elements of success *will* expand.

Over time, the mix of DevOps and ITIL in your organization will evolve, your topology will expand to include on-prem and cloud infrastructure, your applications will be refactored or updated, you will bring in new monitoring capabilities and sunset old solutions, you'll have to provide outputs to new collaboration tools, and your company will build entirely new services over time.

This isn't the old model of cyclical replacement cycles. Today's organizations keep the things that are working well, and build on top of, and next to, those elements. If 6 or 9+ months are needed just to get a solution configured and running, then reconfiguring it as expansion and evolution occurs will suck away all the value within the first 1-2 years. A good AIOps solution should make it easy and fast to adapt to the ever-changing IT Operations landscape.

? **Sample questions to ask when evaluating AIOps tools for rapid time-to-value:**

- How long does it take customers to go into production with the AIOps platform on average?
- How long does it take to set up new integrations or change/configure existing ones, to help onboard new tools, teams and applications?
- How long does it take to set up or modify correlation patterns in response to new technical or business requirements?

While some would argue that these attributes are not just desired, but also required, for AIOps success, you are best equipped to make that call for your environment.

Next, we'll consider required or "must-have" AIOps capabilities, without which it will be exceedingly difficult for your AIOps investment to succeed and generate a strong ROI.

Key takeaways

Desired characteristics for successful AIOps solutions

- ✓ **Ease-of-use:** deliver an intuitive experience for all users (NOC, DevOps, SRE, others)
- ✓ **Self-service configuration:** don't require experts for either setup or ongoing administration
- ✓ **Rapid time-to-value:** help you realize value in weeks, not years

AIOps solutions: required capabilities

Following the desired characteristics of AIOps solutions, let's discuss the required capabilities. In today's environments, unless your AIOps solution delivers these capabilities, you won't be able to succeed with it.

1 Integrates with and connects your existing tools

The monitoring, collaboration, remediation, change and topology systems you have in place each serves a unique purpose. These tools also took time and effort to deploy, they have processes built around them, and people know how to use them.

These tools also make up a functional IT operations ecosystem that operates well together, and they provide value that you can't just throw away.

So, your AIOps solution shouldn't displace your existing tools or force you to embark on a long and painful rip-and-replace project. Much of the value in AIOps is in being a good normalizer and connector between those existing tools and automating the data flows between them.

That's why it's critical that your AIOps solution provide both OOTB connectors and flexible APIs to connect all of your tools together and provide these capabilities:

1. Integrate with all the monitoring and observability tools you've already put in place and ingest their events and alerts in real time.
2. Pull in and use topology and dependency data from wherever your organization currently keeps it, even where it's incomplete; this includes your ticketing, CMDB, ITAM and DCIM systems, service and network maps, emails, documents, spreadsheets and all other sources.
3. Consolidate change events from all manual and automated sources at the speed and scale required; this includes your change management module inside your ITSM systems, your CI/CD tools and all other sources of changes in today's IT environments.
4. Provide outputs and allow bi-directional updates to your chat systems, ticketing systems, escalation tools, automation/remediation systems, and your reporting/data analysis tools.

2 Data preparation and presentation

Data ingestion, normalization, enrichment, and tagging are all critical “rinse and repeat” functions in AIOps that are often undervalued or even overlooked, leading to low-quality results and at the extreme end, AIOps failure.

So the complete range of data preparation activities has to be deeply embedded in an AIOps solution, and needs to be intuitive, transparent, and testable. You will tune and reconfigure and optimize them not once, but again, and again, as your environment, tools, and services evolve:

- ✓ **ML data cleanse and preparation:** In many ways, the success of an ML algorithm depends on good data preparation. It also avoids the well-known “garbage in, garbage out” maxim in AI/ML.

  Deduplication, normalization, and enrichment of your event data provides the inputs to the ML algorithms that analyze, correlate, and connect those events. That’s why it’s critically important that your AIOps solution provides this functionality out of the box.

  Without these built-in capabilities, organizations are forced to build their own normalization and event enrichment engines which vacuum up valuable engineering resources, time and money.

- ✓ **Human data preparation:** At some point, the data coming from your operations pipeline needs to be consumed by a human operator.

  Whether its contextualizing an incident with the associated service, region, or team assignment or providing a link to a runbook or remediation system, your AIOps solution must allow you to add human-understandable operational and business context to your incidents before presenting those incidents to your IT Ops, NOC and DevOps/SRE teams.

  Often this context can also be used to drive automatic workflows in other systems or tools connected to your AIOps solution, so this creates an additional win for your teams and your organization.

- ✓ **Serve up “suspect” changes to identify root cause:** In modern environments, changes are increasingly the reason applications and services experience hiccups or even outright outages.

  Effective AIOps solutions must be able to match incidents to all the associated changes in the environments, identify the likely changes that caused an incident, and present human operators with the list of suspect changes.

  This helps human teams rapidly hone in on problem changes, investigate further and roll them back if necessary, to minimize disruptions to customers and users.

- ✓ **Workflow and outbound sharing:** Beyond preparing data inputs from external sources, an AIOps solution should provide a robust capability to manage your data outputs; the data you send to your ticketing, escalation, chat, and remediation systems.

  Your AIOps solution must allow you to control how data is shared with those external systems. That means that your AIOps solution must support both automatic sharing, as well as on-demand sharing.

  Your AIOps solution must also allow bi-directional updates, once data is shared out. This helps teams relying on your AIOps solution, and the other systems, to stay in sync. Without this critical capability, your teams will waste valuable hours on duplicate or unnecessary tasks.

3 Transparent, controllable and "interactive" machine learning

The large volumes of data involved in maintaining event, alert, change, topology, and incident awareness across your IT operations surface area are a perfect fit for machine learning's ability to process large datasets in real time and identify trends or patterns across multiple dimensions simultaneously.

- ✓ **Transparency:** A key to effective use of machine learning in IT Operations is transparency; any operator must be able to easily determine why certain events were grouped or correlated together, or why a given change was identified as the suspected root cause.

  Avoid "black box" AI/ML powered AIOps solutions; there should be no ambiguity or "AI magic" in how the AIOps solution is delivering its outputs.

- ✓ **Controllability:** That transparency and usability must extend beyond the initial configuration of an AIOps solution, because inevitably, changes to your services, tools, and environment will require adjustments to correlation patterns.

  This means that your IT operations team should be able to tweak or fine-tune your AIOps solution's ML-generated logic in response to changing technical or business requirements.

  Frequently, domain experts on your team will maintain repositories of "tribal knowledge" about your environment that reflects years of hard-won knowledge. Your AIOps solution must allow this tribal knowledge to be incorporated into your solution's ML-generated logic.

- ✓ **No data scientists required:** Effective AIOps can't require every operator to become a data scientist; ML's role in IT operations is to reduce complexity in detecting, diagnosing, and remediating issues, not increasing it.

That means that your AIOps solution must not require a data science or other expert team to create new logic or modify existing logic.

✓ **Testability:** Your AIOps solution should allow your teams to test changes they make to its ML-generated logic before deploying it to production, just like with all other critical enterprise applications.

✓ **Interactivity:** Ideally, the ML applied should be interactive, allowing operators to provide feedback that validates or invalidates the outputs. That feedback can be used to tune the ML-generated correlation patterns or root cause diagnosis for future incidents.

4 Reporting capabilities

An AIOps solution acts as a hub for IT operations data. This includes not just ingesting monitoring, topology, and change inputs, but also incident outcomes and operational performance data.

Your AIOps solution must therefore make it easy to extract, contextualize, pivot, and present that data, in order to provide deep, actionable insights into all facets of operations.

Done correctly, this will help your organization measure, track and improve upon all operational KPIs and operational efficiencies over time.

Here are some reporting and analysis capabilities your AIOps must provide out of the box:

1. Failure rates for a specific device, application, or other service component, to evaluate and identify hotspots and recurring problems, and prioritize replacement/refactoring efforts.
2. Trend analysis to gauge the effect of transformation efforts or service / process / tool changes on core operational metrics like availability, incident volume and root cause data.
3. Team-specific operational metadata, incident frequency and priority, monitoring effectiveness, root cause types and frequency, remediation times, and team-owned service availability. A service-specific pivot of that same operational metadata is also useful.
4. Trend analysis of operational workload / effort level, pivoted by service, team, business unit, or other meaningful variable specific to your business.
5. KPIs for each monitoring and observability solution whose data your AIOps solution is ingesting, such as signal:noise, actionability, and time-to-detect.

This data can be used to evaluate investment in those solutions and lead to tool renewal/retirement decisions.

5 Democratized AIOps

Finally, in today's fast-moving business environments, every organization's stakeholder must be able to adopt and benefit from an AIOps solution.

In concrete terms, that means that your AIOps solution should be able to serve:

- ✓ **Every type of organization:** Organizations modernizing their legacy systems and the largest, most sophisticated enterprises in the world, must be able to benefit from AIOps, without requiring expert teams or expensive consultants to deploy and maintain your AIOps solution. Throughout the deployment and maintenance pipeline, your AIOps solution should provide self-service access to setup, configuration, customization and troubleshooting actions.
- ✓ **Every team:** From IT Ops and NOC to DevOps and AppDev to SRE teams, everyone should be able to access your AIOps system and easily view the overall or specific operational attributes they need access to.
- ✓ **Every IT persona:** Similarly, from NOC Managers to BU leaders and service owners to CIOs, every organization's IT persona should be able to visualize data and benefit from AIOps. Your AIOps solution must provide RBAC and role-based reporting to serve these users.

6 Consider a "domain-agnostic" AIOps solution

In their recent "Market Guide for AIOps Platforms," Gartner Research asserted that:

1. The future of AIOps was "domain-agnostic," AIOps platforms that could process diverse datasets, support multiple use-cases and work for diverse teams including I&O, DevOps and SRE teams.
2. "Domain-centric" AIOps platforms could only serve organizations with limited data variety and support a small set of focused use-cases.

Leverage Gartner's deep insights into IT operations and AIOps, and their recommendations to help you choose the AIOps solution that is best suited for your organization's IT environment and needs.

Key takeaways

Required capabilities for successful AIOps solutions

- ✓ **Works with all your existing tools:** shouldn't require you to painfully rip-and-replace your existing tools or force you to change the way you work.
- ✓ **Provide built-in data preparation:** should provide out-of-the-box data aggregation, normalization, enrichment and tagging that works at scale.
- ✓ **Provide explainable AI/ML:** ML logic should be transparent, testable and controllable to foster adoption and usage.
- ✓ **Deliver robust reporting capabilities:** should make it easy to extract, contextualize and present collected data, so you can easily measure, track and improve IT Ops KPIs and metrics
- ✓ **Democratize AIOps:** should deliver on the promise of "AIOps for all" so that every organization — not just the advanced ones — can adopt and benefit from AIOps
- ✓ **Work with all of your IT domains:** provide "domain-agnostic" AIOps functionality, as described by Gartner in the 2021 "Market Guide for AIOps Platforms."

BigPanda: A pragmatic AIOps platform that delivers value in under 12 weeks

"Tools like BigPanda really help us to reduce all our noise levels, get that aggregation done, get that enrichment done, and then produce actionable incidents so somebody can be paged to look at it."

– **Dan Grace, Global Technology Operations Leader, Equifax**

Built from the ground-up for large scale and complex IT environments, BigPanda's Event Correlation and Automation platform, powered by AIOps, helps organizations prevent and resolve IT outages.

Designed to go live and into production in just 10-12 weeks, BigPanda's AIOps platform delivers three key capabilities: Event Correlation, Root Cause Analysis and Level-0 Automation.

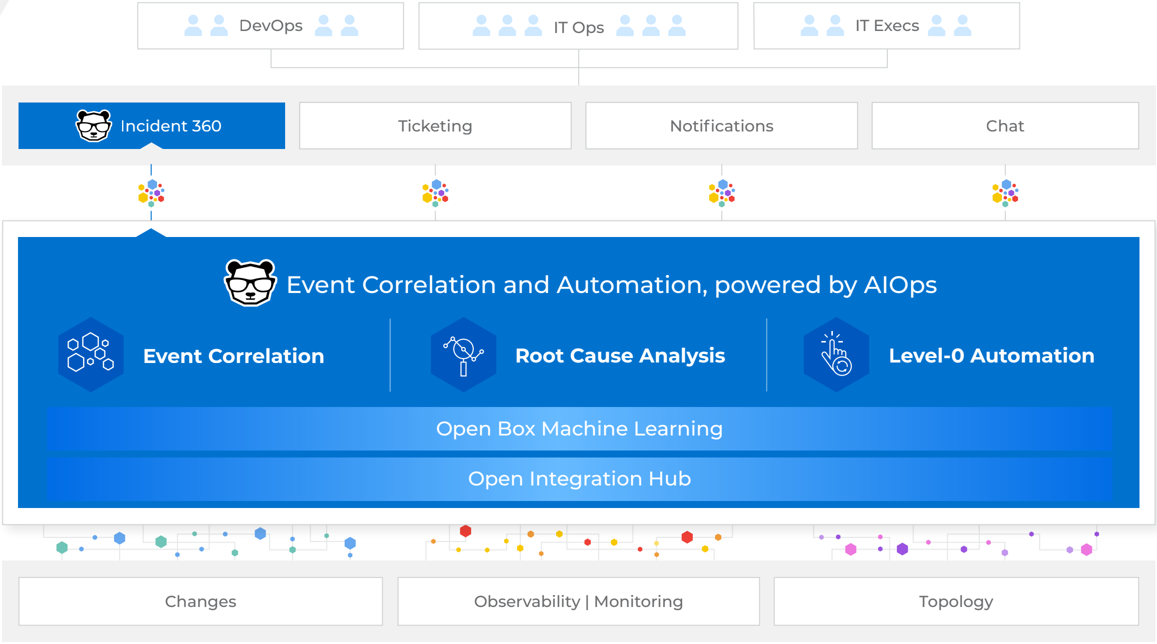

"I remember when I brought BigPanda into my NOC and explained, 'We are going to start automating things,' and the first thing that people might have felt is 'we're not going to have jobs.'

But it did the opposite.

Being able to automate manual operations and workflows gave us the ability to do new, more exciting work."

– Ben Narramore, Director of Operations, Sony Playstation Network

Event Correlation

BigPanda's Event Correlation aggregates and correlates disparate streams of observability, monitoring, change and topology data into context-rich incidents in real time.

Using 50+ out-of-the-box integrations and powerful REST APIs, BigPanda connects to existing observability and monitoring tools and aggregates their data in real time. The system normalizes the data into a consistent format and adds context by bringing in topology and operational data.

Leveraging Open Box Machine Learning (OBML), BigPanda then correlates the collected alert and topology data into a handful of context-rich incidents, dramatically reducing the noise. By helping teams detect and take action on incidents as they form in real time, BigPanda prevents those incidents from escalating into outages.

BigPanda also provides the ability to archive, report on, and analyze correlated incidents over time. This helps organizations measure, track and improve key IT Ops KPIs and metrics, further decreasing operational workloads over time.

Root Cause Analysis

Modern IT environments are complex and chaotic, and therefore a single RCA technique cannot address all types of incidents and outages.

That's why BigPanda provides a range of features and capabilities that together provide a comprehensive solution for any scenario.

BigPanda's Root Cause Analysis capability uses Open Box Machine Learning to help organizations identify changes in infrastructure and applications that cause the majority of today's incidents and outages. In addition, BigPanda identifies low-level infrastructure issues that cause problems.

By pinpointing the root cause of incidents in real time, BigPanda helps organizations and their operations teams rapidly investigate and resolve them.

Level-0 Automation

BigPanda's Level-0 Automation eliminates the repetitive manual processes of incident response to create a seamless experience for IT Operations teams.

Teams can integrate BigPanda with different collaboration tools and automate creating tickets, sending relevant notifications, and setting up war rooms with the right teams. Automatic bi-directional syncing ensures that teams on either side always have access to the latest incident information and updates.

BigPanda also connects to Runbook Automation tools to run workflow automations that can resolve incidents more quickly. Together, these automations shave critical minutes off of incident response, and help organizations and their IT Operations teams rapidly resolve incidents and outages.

Once BigPanda is in production, organizations benefit from reduced operating costs, improved performance and availability, and increased business agility.

> "What BigPanda has helped us do is take all the sources that we have put through, and provide correlation. BigPanda discovers a lot of things that the brain just can't keep up with, certainly in a large team that spans the globe."
>
> – **Bill Hancock, Director of Site Reliability Operations, Expedia**

Key takeaways

BigPanda's Event Correlation and Automation platform, powered by AIOps, provides three main capabilities to help organizations prevent and resolve IT outages:

- ✓ **Event Correlation:** uses explainable AI to correlate disparate streams of observability, monitoring, change and topology data into context-rich incidents in real time
- ✓ **Root Cause Analysis:** uses explainable AI to surface probable root cause and root cause changes in real time, inside today's complex and dynamic IT environments
- ✓ **Level-0 Automation:** eliminates repetitive manual incident response processes to accelerate incident response, remediation and resolution.

Conclusion

AIOps is a transformational technology whose time has come.

Organizations that successfully adopt AIOps can deliver extraordinary customer experiences while they transform digitally.

The proliferation of AIOps vendors and technologies increase the chances of organizations getting mired with the wrong AIOps solution, wasting years, without realizing the promised return on their investment.

As you evaluate AIOps tools, consider both the desired and required AIOps characteristics, so you can optimize for success and rapidly realize the benefits. Remember that the best AIOps solutions are distinguished by proven track records for ease of deployment, ease of use, rapid time to value and high-quality results.

Choosing the right AIOps solution will help you dramatically improve on your KPIs and metrics by reducing MTTD, MTTA and MTTR in just a few months. On top of that, the right AIOps solution will help reduce the operational workload on your overworked IT Ops, NOC, DevOps and SRE teams. Finally, the right AIOps solution will deliver these benefits while reducing your operating costs, improving performance and availability, and helping your organization accelerate business agility.

Resources

Market Guide for AIOps Platforms — Published April 6, 2021 — ID G00735577, by Analysts Pankaj Prasad, Padraig Byrne, Josh Chessman

Event Enrichment: The linchpin of AIOps

Guide to Investing in AIOps

AIOps Evolution Podcast

Get started with BigPanda
(650) 562-6555 | info@bigpanda.io
www.bigpanda.io

BigPanda