GUIDE

A practical guide to IT Ops maturity

How to assess and create a roadmap to autonomous IT operations

BigPanda

Table of contents

Introduction

Think back to the last crippling P1 outage that occurred in your environment.

One of your customers may have called it in. Your teams spent hours on a bridge call with overworked L3s and DevOps teams from every major function to try to get to the root cause. The outage lasted the better part of a day.

What was probably even more frustrating for you was the realization that this was the third time this happened in as many months…

– despite the hundreds of thousands of dollars you invested in monitoring and observability 6 months ago;

– despite the complete overhaul of your incident management workflows and processes just a year ago, based on your consultants' advice;

– despite having a sophisticated runbook automation tool;

– despite having a "best-in-class" CMDB/powerful ITSM suite/modern incident response platform...

and so on.

Maybe it's no consolation, but you're not alone.

In the last eight years, IT has evolved to accommodate always-on digital services (think hybrid clouds, microservices-based apps, DevOps and SRE models, and everything-as-code). As a result, IT Operations' effectiveness and efficiency has become critical.. With the Covid pandemic came the acceleration of digital transformation, the shift to remote work and an even bigger spotlight on IT Operations.

With that spotlight comes the need to have best-in-class observability, monitoring, collaboration and AIOps tools, scalable and resilient workflows, the ability to measure the right KPIs, and of course, IT Ops, NOC and DevOps teams that possess the skills their organizations need.

Successful operations leaders are realizing that being excellent, or even great, in one of these areas alone is not enough and doesn't provide protection against frequent, long and crippling incidents and outages.

Instead, they understand that each function and process within IT Operations must grow and mature in conjunction with each other. The lack of maturity in one dimension can seriously slow down operations, no matter how mature other dimensions are.

But what are those dimensions?

- And since no organization can be excellent in every dimension on day one, how should you assess your current operational maturity with a clear eye, think about what the subsequent phase of operational maturity looks like, and map a course of action to get to the next phase?
- What are the benefits you can expect to see at each phase—and which KPIs should you use to measure success along your journey?

We've answered these questions based on the real-world experiences of IT Operations teams at dozens of Fortune 1000 enterprises across every major vertical, to create the IT Ops Maturity Model described in this guide.

The model serves as a good framework for conversations about the five phases of IT Ops maturity, and the three dimensions along which you need to assess your maturity, and the KPIs to use in that process. Use this guide to create a 1-3-5 year plan for your IT Operations maturity process, which supports your organization's revenue and growth goals.

The five phases of IT Ops Maturity

If you've spent any time in IT operations, you're probably familiar with the phrase, "we need to be proactive, not reactive." Admittedly, reactive and proactive are different phases of maturity, but it begs the question: is this binary choice over-simplifying things? Are there other maturity phases?

When we put this question to IT Ops leaders from some of the largest organizations in the world, they said, "yes, yes and yes."

Reactive and proactive are two broad categories used to define IT Ops. There are actually five phases of maturity, on the path to the aspriational North Star phase of Autonomous IT Operations.

	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

PHASE 0

Reactive

What does a truly reactive IT Ops organization look like?

1. IT Operations functions in a very uncoordinated manner.
2. Monitoring is nonexistent or siloed within different domain teams.
3. Incident management, change, and topology processes are. ad-hoc, undocumented, and/or team-specific.
4. The majority of high-priority issues are user-reported.
5. Service availability and other KPIs are not measured or measured poorly.

PHASE 1

Responsive

Organizations that have made some progress at first become responsive. What does a responsive organization look like?

1. Trend towards centralized or decentralized operations, but not hybrid.
2. Limited monitoring in place with an overwhelming amount of IT noise that teams are forced to deal with.
3. Struggling to rapidly and effectively detect and remediate service issues; teams still rely on users reporting problems.
4. Incident management processes are documented, but trust between operations teams is low—which leads to limited information sharing during problem diagnosis and root cause analysis.
5. Monitoring, topology and change processes are still siloed, and in larger organizations teams are still operating at sub-scale efficiencies.

PHASE 2

Proactive

What are the characteristics of an IT Ops team that is operating proactively?

1. A partially integrated operations pipeline that is tied to a continuous process.
2. Monitoring coverage is extensive, resulting in high data volume. Generally this is "peak noise." There is low actionability on monitoring data and an excess of low-priority incidents.
3. Event processing is generally rules-based and consistent. However, rules must be manually written and maintained, which is unsustainable in most large and/or fast-moving environments.
4. Topology information is mostly consolidated, with some known dependencies between locations, hosts, applications, and services.
5. Most environment changes are logged (and planned), and major ones are centralized for general awareness.

PHASE 3

Semi-Predictive, Semi-Autonomous

Some organizations have started to use AI/ML to operate in a semi-predictive, semi-autonomous mode. What does an organization's IT Operations typically look like in this phase?

1. A common operations pipeline connects multiple monitoring, change, and topology data sources. This pipeline provides clean, distilled, actionable outputs to collaboration/ITSM tools and reporting systems.
2. Different workflows and tasks have been automated throughout the incident process, but they still require some level of manual intervention and action.
3. There is extensive monitoring coverage that generates high-quality data across different technology domains. These organizations use anomaly detection for metrics/logs and create and make use of SLIs, SLOs and error budgets.

4. Event enrichment and AI/ML-based correlation reduce incident volume and prevent incidents from escalating into outages. Maintenance-based event suppression further reduces noise.
5. Topology and change information is processed by AI/ML-based engines for root cause triangulation.
6. Limited number of incidents, tied to specific and/or known scenarios, are auto-remediated.

PHASE 4

Autonomous IT Operations

The final phase of operational maturity is autonomous IT operations. Serving as an inspirational North Star, in this phase, we should expect to see:

1. A centralized pipeline that handles all monitoring integration, topology mapping, change and alert/event data normalization, enrichment, correlation, remediation, and reporting. Incident handling is automated, from detection to remediation.
2. The outputs from this pipeline are:
   - Distributed to a variety of auto-remediation systems.
   - Shared with operations staff in the form of outcome reports that are then used for process and engineering improvements aimed at service resilience and reliability.
3. Achievable, documented monitoring requirements are met for all services, including minimum payload, coverage, and actionability.
4. Modular anomaly detection tools are used across layers and technology domains. There is flexibility in monitoring and remediation/automation tools and systems, but there is broad, organizational-wide consistency in monitoring standards.
5. All events/alerts associated with a common root cause (including changes) are correlated using AI/ML.
6. All other functions are executed autonomously, where possible.

KPIs to assess your IT Ops maturity

To understand and assess which phase your operations map to, we recommend tracking two key performance indicators (KPIs): operational workload and service availability.

Operational workload

Operational workload, put simply, is the amount of effort an organization has to spend on IT Ops. It includes all of the hours spent detecting, classifying, assigning, diagnosing, and remediating incidents. Notably, it includes the hours spent diagnosing whether an incident is non-actionable noise or truly an incident that requires mitigation and remediation. An approximation can be calculated using (Incident Volume Per Month) X (MTTR) = Operational Workload.

Here's how organizations can expect this KPI to change across the five phases of IT Ops maturity.

	PHASE 0	PHASE 1	PHASE 2	PHASE 3	PHASE 4
Descriptive State	Reactive	Responsive	Proactive	Semi-Predictive, Semi-Autonomous	Autonomous IT Operations
KPI: Operational Workload	Minimal	Growing L1, and/or L2+L3 depending on approach	Peak L1/L2/L3 workload	Steady or decreasing; reclaimed L1/L2/L3/SME engineering bandwidth	Fully reclaimed bandwidth; L1 primarily runs AIOps pipeline

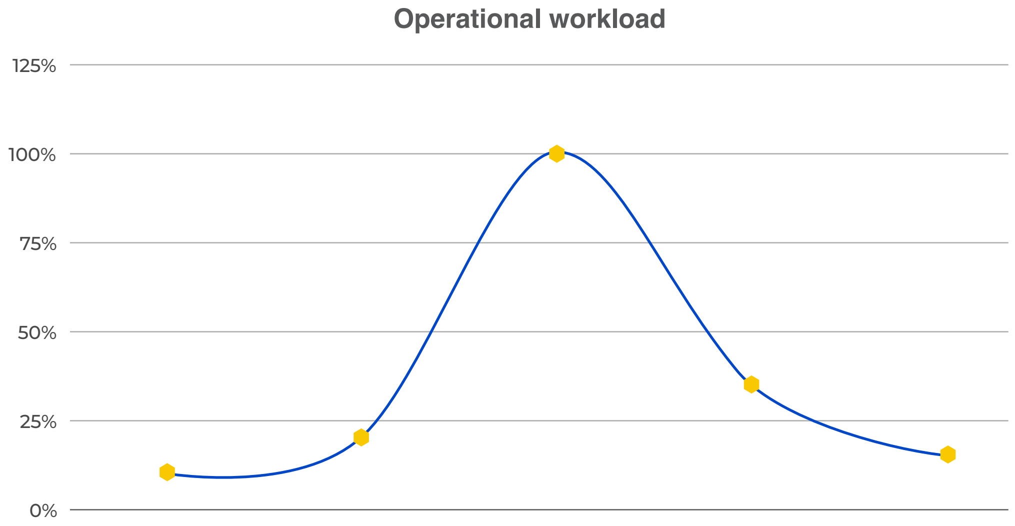

Note: Trend lines are for illustrative purposes only.

Service availability

Service availability is simply defined as the percentage of uptime for a given service. Any periods of limited availability or non-availability of the service is defined as downtime and can be attributed to either maintenance or incidents. Most organizations strive for at least 99.9% availability and measure it at a sub-service level. It is then aggregated across a business-level service each month. During maintenance and incidents, organizations measure the number of 'impact minutes' affecting end-users of that service, and subtract that from the total number of minutes in the month, and then divide it by the total minutes in the month. Some organizations apply a modifier to impact minutes based on the number of customers impacted, or the 'degree' of impact for a partially-functioning service. As organizations improve their IT Ops, service availability should improve, as incidents occur less frequently and are detected, diagnosed, and resolved more quickly.

Here's how organizations can expect this KPI to change across the five phases of IT Ops maturity.

	PHASE 0	PHASE 1	PHASE 2	PHASE 3	PHASE 4
Descriptive State	Reactive	Responsive	Proactive	Semi-Predictive, Semi-Autonomous	Autonomous IT Operations
KPI: Operational Workload	Unmeasured/ Untrusted, low awareness, low expectations	Measured but incomplete due to coverage gaps, perceived as poor by 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, remediation automation speed improvements, monitoring improvements

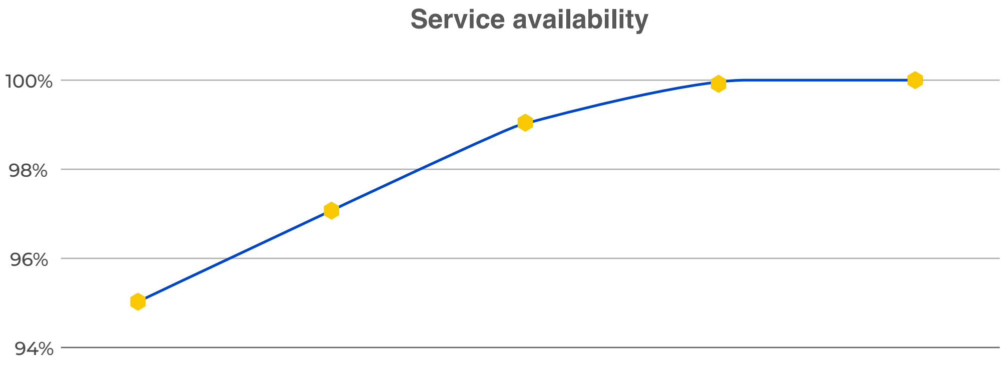

Note: Trend lines are for illustrative purposes only.

We now understand what each phase of IT Ops maturity is, the characteristics of each phase and the 2 essential KPIs to assess our overall maturity.

As we characterized each phase you probably noticed that there were some recurring themes in the form of tooling, incident management processes and meta-data from other adjacencies that impacts IT operations, such as change and topology data.

Now let's dive deeper into three specific primary dimensions of IT Ops maturity, and explore how these dimensions change across every phase of IT Ops maturity.

Deep-dive: The three dimensions of IT Ops maturity and their KPIs

Imagine that you put a 600-hp engine into a rusted, old car chassis sitting atop tires with just 10% tread-life left. You will likely be very disappointed with the driving experience.

A car that delivers a great driving experience must have a powerful engine, performance tires and a chassis that is structurally sound and aerodynamic, among other things.

IT Ops is not that different.

In order to achieve a higher level of maturity in IT Ops, multiple activities must be done well at the same time, creating powerful synergies between each of those activities.

Those activities can be grouped into three dimensions:

- ✓ **Monitoring and event processing**
- ✓ **Incident management**
- ✓ **Operational awareness**

Dimension 1: Monitoring and event processing

Monitoring and event processing are the inputs into the IT Operations pipeline.

Let's break this one down further into (a) monitoring and (b) event processing.

Monitoring
Raw events from a wide range of sources define an organization's monitoring coverage and quality.

What do we mean by coverage?

If we look at a monitored service as a collection of applications, infrastructure, and network elements, then monitoring coverage and quality can be determined by identifying which of those elements are monitored, and then how actionable the events generated by that monitoring are. Monitoring coverage is naturally better when there is a high proportion of "alert-detected" or "alert-indicated" (vs. user-reported) incidents.

Similarly, monitoring quality is better when there is a high proportion of actionable incidents. As maturity increases, monitoring coverage and quality can be improved through the use of service-level indicators and metrics (synthetics, RUM, client telemetry) and metrics or logs-based anomaly detection.

Important Note: There's no good or easy way to decouple monitoring from event processing. A common pitfall for many organizations is that they tend to overemphasize monitoring, especially monitoring coverage, over all else.

Having comprehensive monitoring coverage for every application, node and service in isolation, does little to reduce the frequency, duration or impact of outages. Taken to the extreme, comprehensive monitoring coverage with very high levels of instrumentation actually creates a major problem for organizations because IT Ops, NOC and DevOps teams end up drowning in data and miss critical monitoring alerts that indicate a critical incident or imminent outage. Additionally, overemphasizing monitoring coverage can lead to wasting valuable time and money fixing the monitoring/observability layer while customers and users continue to be impacted by incidents and outages.

Event processing

Event processing maturity can be gauged by looking at how events are handled once they are ingested into the IT Operations pipeline and answering several questions, such as:

- Are events normalized to a common format with at least the minimum required attributes?
- Are they deduplicated to reduce overall volume?
- Are maintenance-related alerts suppressed to prevent non-actionable noise from entering the pipeline?
- Are they enriched with relevant context and dependency information?
- Are causal and symptomatic events (that are related to the same problem) effectively correlated together into a single item?
- Has incident volume been reduced to be as near as possible to the actual number of actionable incidents the organization experiences, and has noise elimination been a priority in every phase of event processing?
- Is event processing being done manually, or has it been automated, and are there ways to analyze it, improve it, and control it?

KPIs

The two critical KPIs to use to measure your organization's maturity in this dimension are:

1. Signal-to-noise ratio
2. Relative incident volume

The signal-to-noise ratio in IT Ops compares the number of events that are actionable (requiring that something be done to mitigate or prevent a service degradation or outage) to the number of events that are non-actionable events and be safely ignored. Organizations should strive to improve this ratio by eliminating identifiable, repeatable noise.

Incident volume is a product of signal-to-noise management, and is simply the number of incidents generated during a given month. As noise is eliminated from monitoring, upstream of incident creation, incident volume will decrease. If noise is not eliminated, then incident volume will generally remain constant. Its important to note that this assumes a constant event volume; if additional monitoring is added, then a new baseline is required.

The table on the next page provides a summary of how this dimension—monitoring and event processing—changes across each phase of IT Ops maturity, along with typical qualitative and quantitative KPI metrics seen at each phase.

Monitoring and event processing

	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
Monitoring Coverage and Quality	Monitoring is nonexistent (no alerts), generally diagnostic, and siloed within a seperate team/ domain (app/infra/ network/etc.)	Limited centralized monitoring, low quality, user reports are ground truth. Possible dashboarding of 5-10 key service metrics.	Large volume of events. Low actionability. Incomplete coverage.	Large volume of events and alerts, improving actionability, extensive normalization and enrichment. Use of algorithmic anomaly detection to reduce effort level. Good coverage.	Achievable, documented monitoring requirements are met for all services, including minimum payload, coverage, and actionability. Modular anomaly detection tools used across layers/ disciplines. Flexibility in monitoring tools and systems, consistency in standards.
Quality KPI: Signal: Noise	Unmeasured/ Untrusted	Measured, <25% signal	50%+ signal	75%+signal	95%+ signal
Coverage KPI: Incident Detection Method	User Reports/Calls	50% Alerts, 50% User Reports	75% Alerts, 25% User Reports	95%+ alerts	99%+ alerts
Event Processing	Unnecessary (no alerts)	Manual, tribal knowledge-based event handling; inconsistent and incomplete.	Tribal knowledge converted to rules-based event handling; consistent but incomplete.	ML enrichment and topology-based correlation. Rules-based confirmed/tuned. Automated. Consistent, complete. Change-to-incident correlation >50% accurate.	ML-based event processing. Dynamic. Automated. All events/alerts associated with a common root cause (including changes) are correlated.
KPI: Relative Incident Volume	Low volume of high priority	Multiple incidents per failure; increasing volume	Possible peak volume as monitoring coverage increases but handling is imperfect.	Nearing actual failure count.	Low volume of mid to low priority. Varies proportionally with changes/ externalities.

2

Dimension 2: Incident management

The second dimension, incident management, is understood as a foundational function of good IT Operations.

Consistent, reliable execution paired with proactive analysis of all of the MTTx metrics—mean times to detect, classify, diagnose, remediate and close/resolve—lead to higher levels of maturity.

Introducing automation to replace manual activities in each phase are all indicators of higher levels of maturity in incident management. Organizations operating the highest levels of maturity analyze past incidents and actions performed on those incidents, to inform their response to new incidents as part of a continuous feedback loop.

KPIs

The three KPIs to use to measure incident management in this dimension are:

1. MTTD/I/R (mean time to detect, investigate or diagnose, and remediate)
2. Incident actionability
3. Incident priority

Incident actionability is similar to signal-to-noise, but for incidents rather than events. To measure this metric, take the number of incidents that were resolved without any action (manual or automated) and subtract it from the total number of incidents. Then, divide result by the total number of incidents.

Incident priority is a value assigned to every incident based on 2 factors: urgency and impact, and is usually determined using a matrix. Organizations generally use low priority numbers (0,1,2) to describe high impact, high urgency incidents, and high numbers to describe low impact, low urgency incidents. All else being equal, as IT Operations maturity improves, the average incident priority decrease.

The table on the next page provides a summary of how this dimension—incident management—changes across each phase of IT Ops maturity, along with typical qualitative KPI metrics seen at each phase

Incident management

	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
Incident Management Process	Ad-hoc, undocumented, siloed.	Documented process with distinct phases, roles, and measures. Low collaboration, fear of blame. Poorly measured.	Incident management (ticketing) system integrated with assignment, escalation, and remediation systems. Generally consistent and reliable handling, still heavily reliant on manual administrative actions.	Monitoring pipeline and real-time collaboration systems fully integrated with incident management Semi- automated incident workflow, reduced administrative toil. Improved consistency and speed.	Integrated, automated incident handling from detection to remediation. Primary output is reporting data used for process and engineering improvements aimed at service resilience and availability.
KPI: MTTR	Unmeasured/ Untrusted	Measured/ Inaccurate due to noise	Measured/ Inaccurate due to noise	Measured, accurate, consistent. MTTR value-stream analysis to optimize incident phases	Minimal
KPI: Incident Actionability	High actionability, impacts are always confirmed via user reports. Low priority incidents are unmanaged until they become critical.	Large percentage of alert-initiated incidents are non-actionable noise (i.e. symptomatic, change-related, or due to bad monitoring).	Smaller percentage of non-actionable incidents. Still some dependence on user reporting to verify impacts.	Nearly all alert-initiated incidents are actionable.	Incident remediation automated and reported to teams/SMEs. Small volume of directly human-managed incidents for unique situations.
KPI: Average Incident Priority	Many urgent issues, may not be called incidents. Generally problems are user-impacting. Low priority incidents may be unhandled, or handled seperately as maintenance or other routine work.	Increasing percentage of low-priority incidents, steady volume of high priority incidents.	Decreasing percentage of high-priority incidents, increasing mid/low priority proportion.	Average priority continues to trend lower as failures are caught earlier.	Average priority stabilizes in low-mid range as historical data and topology understanding provides predictive awareness of failures prior to impact. Assumes maturity in service architecture (HA/resilience/etc).

Dimension 3: Operational awareness

Operational awareness is the degree to which the organization can apply contextual and situational knowledge available from all sources to improve incident outcomes.

For most organizations, operational awareness is achieved tribally.

In these cases, when there is a critical incident, outage or service degradation, domain experts know the topology of the systems involved, the types of monitoring coverage and quality configured for those systems (as well as their limitations), and any recent and relevant changes associated with that incident.

At low levels of maturity, this awareness is siloed. As a result, during investigation, those domain experts are critical participants without whom incidents drag on and/or are poorly handled. Up to a certain scale and organizational velocity, those experts can serve as the 'source of truth' and maintain that level of operational awareness.

vHowever, as organizations grow and scale they run into four challenges when relying on domain experts and their tribal knowledge for operational awareness.

1. It's hard to consolidate all the relevant information from multiple systems, in near-real-time, to enrich alerts and incidents with relevant topological context for the services involved.
2. It's nearly impossible to understand all the relevant changes which may be causal or at least associated with specific alerts seen in a given incident.
3. It's difficult to understand all the relevant operational data needed for remediation, such as the appropriate runbook or specific remediation action that should be applied
4. Finally, tracking potential impacts of different failure types, service level indicators to verify state, and determining how best to communicate with different customers/users/stakeholders for each type of failure is tedious and manual.

The good news is that much of this operational awareness can be automated today, and the degree to which that is done is an important indicator of maturity in this area.

KPI

The KPI to use to measure your organizational maturity in this dimension is:

1. Data volume ingest

Data volume ingest measures the amount of accurate topological and change data ingested into an organization's IT Operations pipeline. Often, change events are siloed within the teams executing the changes because they use unique systems and processes that aren't accessible or understandable to the IT Operations team. As those changes are brought into the IT Ops pipeline, realtime IT Ops awareness of change events increases. Topology data often suffers from the same (unintentional) issue, with the additional element that an understanding of interdependencies is often undocumented, and lives in the form of tribal knowledge, distributed across a few individuals. Capturing that data accurately, keeping it updated as it evolves, and then applying it to the IT Ops pipeline as additional attributes to alerts is a challenging element of maturing IT Ops processes.

The table on the next page provides a summary of how this dimension—operational awareness—changes across each phase of IT Ops maturity, along with typical qualitative KPI metrics seen at each phase.

Operational awareness

	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
Surface Area Awareness: Network, Infrastructure, Cloud, Application	Undocumented or siloed understanding; possibly some data available.	Siloed topologies, ITAM, CMDB, and DCIM, dependency or 'subway' maps. No service mapping.	Most elements known, consolidated, and mapped to locations, hosts, services.	Topology mapped across the stack, used for event enrichment.	Realtime topology mesh, alert template and payload data used for topology generation.
IT to Business Service Relationships	Undocumented or siloed Information.	Named services are stored, but the list is incomplete.	Service-to-stack relationship database is 25-75% complete.	75-95% complete service mapping. Dedicated onboarding function for new services.	Semi-automated service mapping, possible automated discovery.
Change Awareness	Siloed to teams, undocumented.	50% recorded, but siloed.	85%+ recorded, some centralized.	Centralized, accurate, and documented.	Appropriately time-bound and impact-assessed.
Topology-based Incident Diagnosis/RCA	Multi-domain dogpile with multiple teams participating during the incident.	Domains (infra/network/app) compete to prove innocence, worst answer loses. Longest phase of incident.	Topology and monitoring data combined to triangulate root cause, but takes a long time RCA collected for all incidents.	Topology data used in contextualizing incidents and alerts, leading to rapid diagnosis down to 3-5 likely scenarios. Change correlation is rapid. Incident RCA data available to compare to incoming alerts.	Automated root cause diagnosis, including change, based on RCA history. Predictive, at early phases in incident evolution.
KPI: Data Volume Ingest	Minimal	Low	Medium	High	Maximum

Important Note: Remember that the key is to tie together all of the tools that your organization uses in order to measure the KPIs for each dimension. This creates visibility, helps you understand where you are in each of the maturity phases and helps you identify and solve the gaps in those areas. It also gives you the ability to quickly pull key metrics and assess where you are in a given area. That 10,000-foot view of IT Ops as an interconnected pipeline will help you focus your maturity efforts appropriately.

Using the maturity model in your environment

While the three dimensions of IT Ops maturity we discussed in the previous section are individually important, remember that the way they function together is also a prime determinant of operational success.

Building an effective IT Operations pipeline

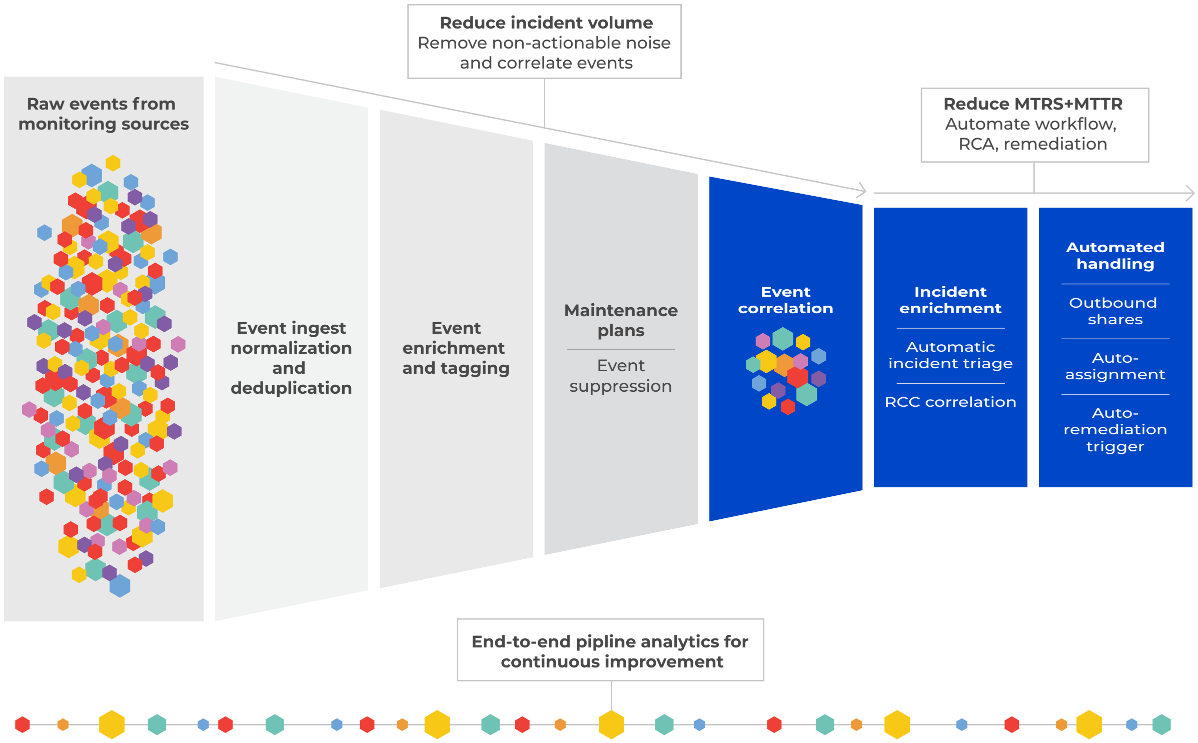
Reduce incident volume
Remove non-actionable noise and correlate events
Reduce MTRS+MTTR
Automate workflow, RCA, remediation
Raw events from monitoring sources
Event ingest normalization and deduplication
Event enrichment and tagging
Maintenance plans
Event suppression
Event correlation
Incident enrichment
Automatic incident triage
RCC correlation
Automated handling
Outbound shares
Auto-assignment
Auto-remediation trigger
End-to-end pipline analytics for continuous improvement

Stronger together

The following real world scenarios offer insight into how improving only one dimension, may not provide the results you are looking for.

Operational awareness data can be useful if, for instance, service topology and dependency relationship data are siloed in a CMDB instead of being used for event processing via enrichment and correlation.

Similarly, a fully automated incident management process won't provide much value if it isn't fed by actionable monitoring. Without actionable monitoring, non-actionable incidents will eat up your team's bandwidth, or worse, become a distraction while actionable incidents are left unattended.

Great monitoring coverage and quality that are spread out to multiple teams, instead of being centralized, will generate high volumes of non-actionable incidents. When great monitoring coverage and quality generate large volumes of incidents but these incidents are not matched against recent relevant changes, the non-actionability goes up. For example, a network configuration change causes symptomatic alerting across SLIs, APM, and infrastructure monitoring systems, but because those high-quality alerts aren't being matched against the network configuration change, operations teams can't do much to or with those alerts.

Conversely, even relatively noisy, immature monitoring and event processing that uses partial operational awareness data with a manual incident management process will be relatively successful.

A single cohesive pipeline

IT Operations processes and technology should be viewed as a continuous pipeline that takes in operational data on one end and provides incidents of some quality as an output on the other end. As that integrated pipeline matures, the organization will realize different degrees of improvement in service reliability and availability.

At each phase of maturity, once you've assessed where each of your three dimensions stands, you want to look for outliers. Metrics are of course critical to this. Incident volume and actionability, signal-to-noise ratio, manual versus automated processes, average priority and MTTR are some that can help you identify strong and weak areas. Because you're only as good as your weakest area, it's best that you focus your efforts where you are the weakest to unblock the chokepoints in your operations pipeline.

If only 30% of your incidents for a given service are detected by monitoring, you'll want to encourage that team to focus on improving monitoring coverage before they do anything else.

If your incident management process is reactive—as in, manual, undocumented or inconsistent—then work across your people, processes and technology to elevate that area up one or two notches to gradually become more automated. Then, over time, you'll become much more proactive.

If a team has difficulty with testing, but they don't follow well-established change logging and tracking processes for their application updates or patches, then the contextual awareness of changes will be low. You'll first need to work with that team on logging and then ingesting their changes into your operations pipeline (to correlate incidents to the changes that may have caused them).

Similarly, if you have little or no contextual awareness, you need to collect that information from all of the different owners in your organization, and their tools. Deal with specific identified maturity gaps first because those are most likely holding you back. Sometimes these gaps are caused by organizational barriers have been avoided or ignored for a long period. To move forward, you must remove those barriers. There's no way around it if you want to enhance and accelerate your maturity.

Before BigPanda

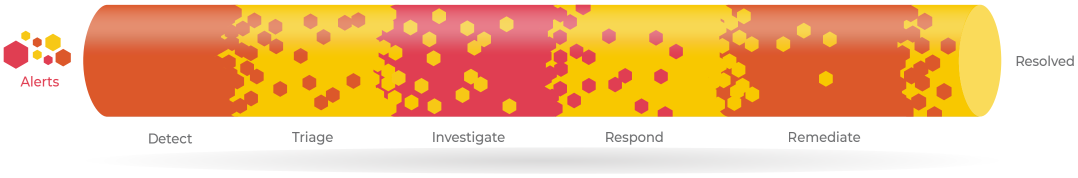

After BigPanda

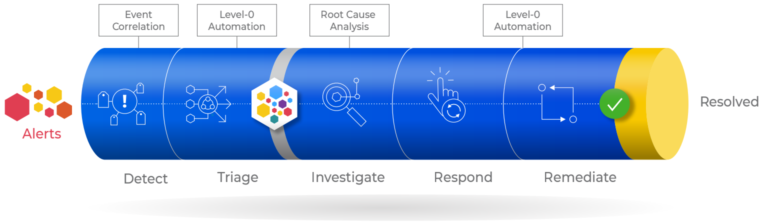

Next steps

While change doesn't happen overnight, we hope that this structured, KPI-driven guide to a proven and tested IT Ops Maturity Model helps you assess where you are today, where you want to go next, and map out your path to get there. As you progress on that path, we have no doubt that your customers will experience ever-increasing levels of reliability and performance, and your key business metrics around retention and growth will reflect that.

Appendix: Sample KPIs for the three dimensions of IT Ops maturity

Dimension 1: Monitoring and event processing

The two critical KPIs to use to measure your organization's maturity in this dimension are:

1. Signal-to-noise ratio (Fig. 1)
2. Relative incident volume (Fig. 2)

Fig. 1: Signal-to-noise ratio

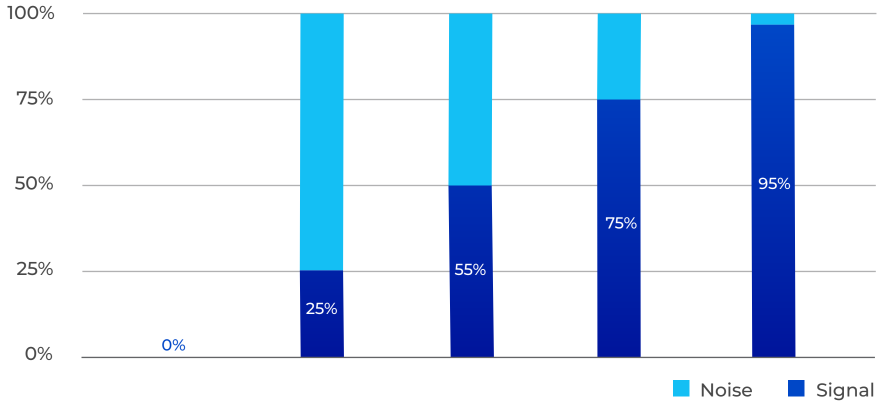

Fig. 2: Relative incident volume

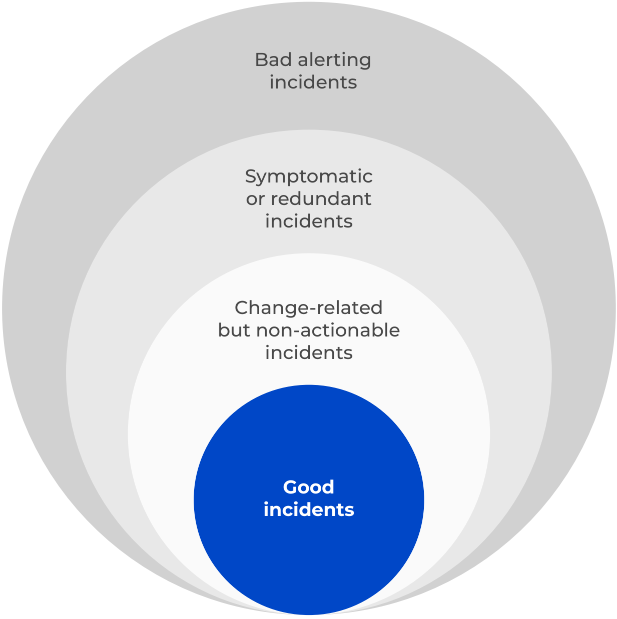
Bad alerting incidents
Symptomatic or redundant incidents
Change-related but non-actionable incidents
Good incidents

The big gray circles are different types of non-actionable incidents, with the size of the circle representing relative incident volume. At the center are the "good incidents" that actually require action to mitigate and resolve.

As you mature in IT Ops you want to see your signal-to-noise ratio improve and the volume of incidents decrease.

Dimension 2: Incident management

KPIs

With incident management we recommend three KPIs with the first one being a hybrid KPI:

1. MTTD/I/R (mean time to detect, investigate or diagnose, and remediate) (Fig. 3)
2. Incident actionability (Fig. 4)
3. Incident priority (Fig. 5)

Fig. 3: MTTD/I/R

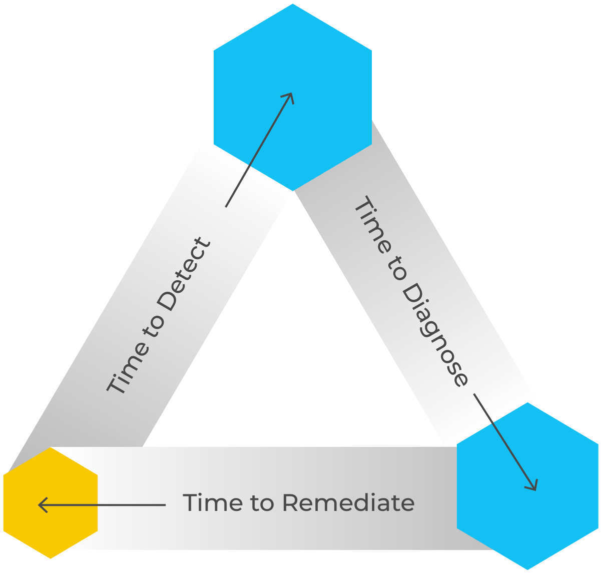

MTTD/I/R

There are three big components to how long an incident lasts; how long it takes to detect an issue, how long it takes to diagnose that issue, and how long it takes to mitigate/remediate impacts from that issue. If you get those average values for an organization, and arrange them in a triangle, you can easily see where to focus your efforts; if you have very short detection times, then that side will be small and the other two will be large; and you should focus on diagnosis or remediation. Same for the other sides. And over time, you should work on reducing the size of the triangle, comparing year over year total area. Its a simple, intuitive view.

Fig 4: Incident actionability

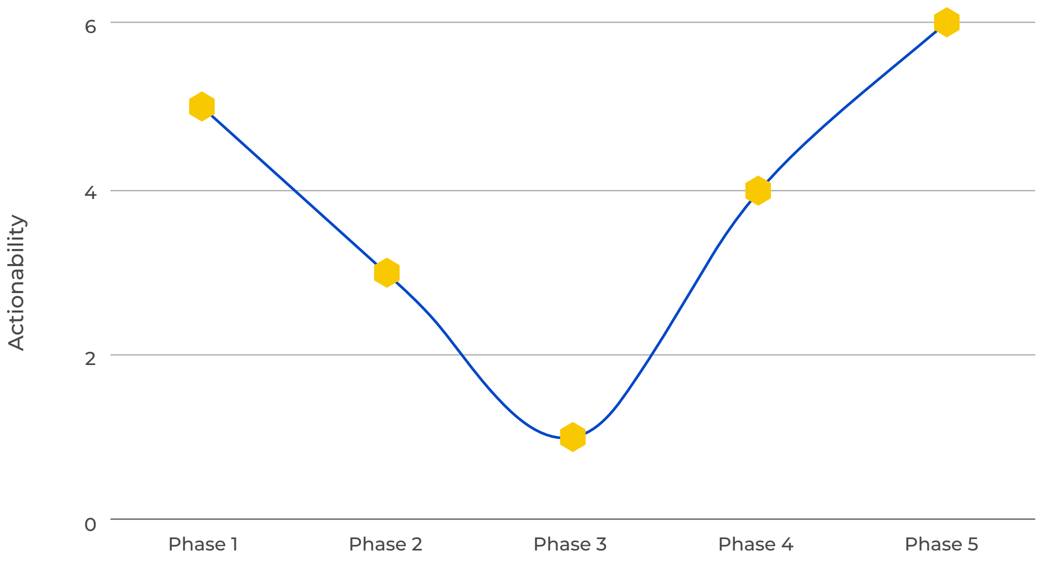

As you mature in IT Ops, actionability is going to go through a few phases. In the first phase, you have poor overall awareness of incidents, so you end up finding out about them once they've already started impacting you, usually based on user reports. They are all very actionable. Then you start to get monitoring in place, but its not effective, because you haven't learned how to ensure it is high quality monitoring, so your monitoring tools generate a lot of noise, and actionability decreases. Then you start to develop the maturity to identify and eliminate noise, and improve your monitoring and processing of that monitoring, so actionability increases again. The Y-values don't matter, it could be percentages or any number, its the shape that matters.

Fig 5: Incident priority

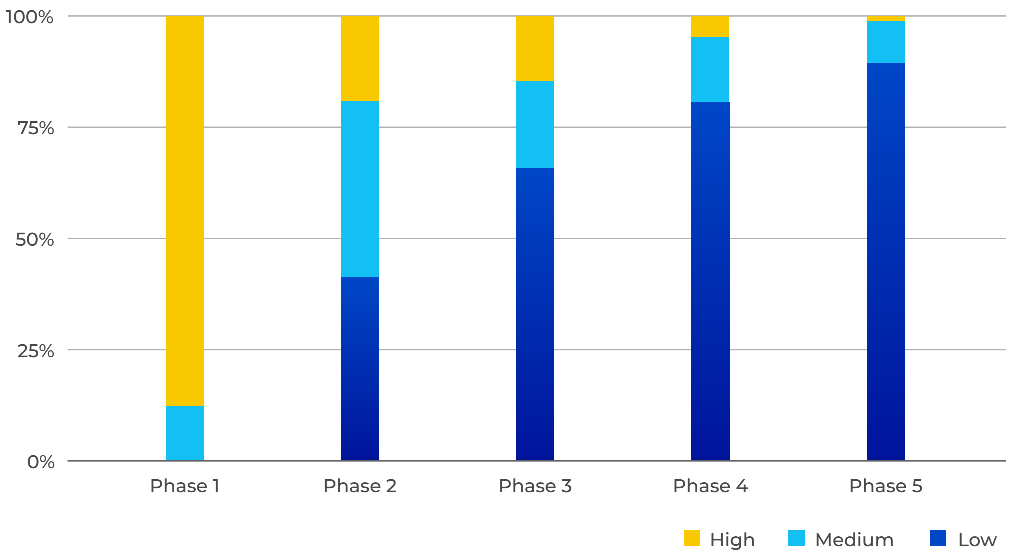

As previously described in the incident actionability section, initially, you have poor monitoring, and find out about incidents very late, when they are critical, high-impact incidents. That's phase one. From there, you start to improve monitoring and awareness, and generate an increasingly high ratio of low-priority and medium-priority incidents. As you mature your IT Operations, your ability to predict impacting/urgent issues well before they actually cause any problems increases. Finally, in phase five, automation can proactively shift workloads or otherwise rapidly identify and remediate small things before they become big things—so everything stays low priority.

Dimension 3: Operational awareness

KPI

The best and simplest KPI here is data volume ingest (Fig. 6).

Fig 6: Relative data volume ingested for IT Ops

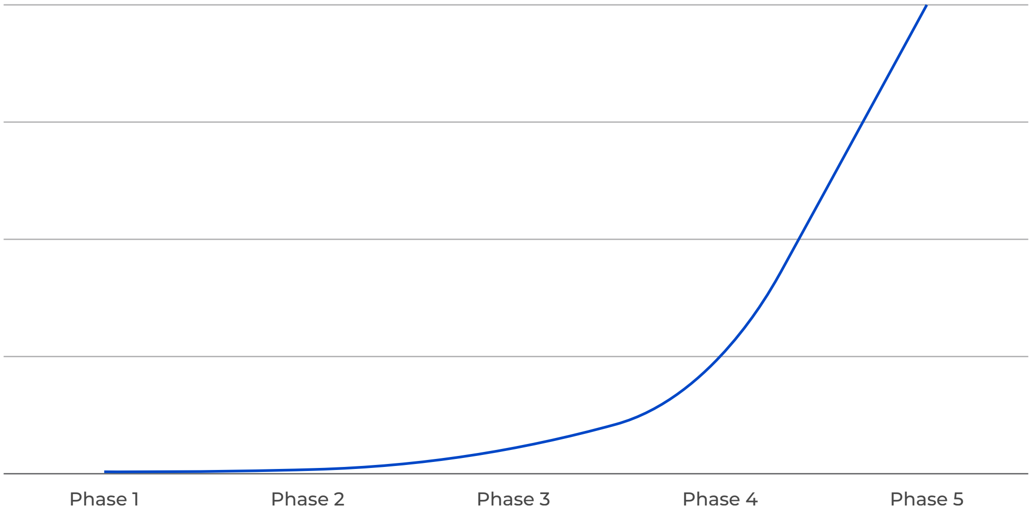

As you mature IT Operations, you will ingest more data. Hence, a curve that is going up and to the right. Keep an eye on this trend since it is an opportunity to add more AI to keep the volume in balance.

To learn more about the IT Ops maturity model, watch the on-demand webinar

BigPanda

About BigPanda

Hundreds of customers have used this model to assess IT Ops maturity across workloads and service availability. We hope it helps you assess your operational maturity and map a course of action for your organization to embrace AIOps. BigPanda keeps businesses running with AIOps that transform IT data into insight and action. With BigPanda's AIOps platform, businesses prevent IT outages, improve incident management and deliver extraordinary customer experiences. Without BigPanda, IT Ops, NOC, and DevOps teams struggle with a tsunami of data and highly-manual, reactive scale, complexity and velocity of modern IT environments. This results in painful outages, unhappy customers, growing IT headcount and the inability to focus on innovation.

BigPanda's AIOps Event Correlation and Automation platform helps Fortune 500 enterprises such as Intel, Cisco, United, Abbott, Marriott and Expedia take a giant step towards Autonomous IT Operations. BigPanda is backed by Advent International, Insight Partners, Sequoia Capital, Mayfield, Battery Ventures, Glynn Capital, Greenfield Partners and Pelion. Visit www.bigpanda.io for more information.

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

BigPanda