WHITE PAPER

BigPanda

Cutting through alert noise with high-quality alerts

Enrichment best practices for ITOps

To keep an organization's digital core up and running, ITOps needs deep visibility and insight into every aspect of it. However, that quest for visibility often leads to monitoring and observability tool sprawl, layers of disjointed, fragmented data silos and ultimately, an uncontrolled volume of alerts that make it difficult to detect and respond to the ones that are actually important. Enriching alerts with context raises their level of actionability, allowing organizations to effectively cut through the noise and reclaim control of their IT operations.

Table of contents

Introduction

As applications, services and infrastructure accumulate over time (organically or through mergers and acquisitions), ITOps organizations add monitoring tools that generate more alerts. As alert volume increases over time, the quality and usefulness of alerts tend to decline, making it hard to discern which alerts are important and need attention. In many cases, however, no structured practice exists for regularly assessing alerts to determine whether they need to be modified or retired. Left alone over time, the resulting environment of IT noise can dramatically overwhelm even the most well-designed incident and alert management workflows and intentions.

Consider the hypothetical case where an organization receives 500 monitoring alerts in its first year. As the scope of monitoring grows, the number of new alerts generated, in addition to the existing alerts, increases by 15%. After 10 years, assuming none of the alert sources were taken out of service, there will be 12,175 total configured alerts in the environment. Suppose that perhaps 5% of alerts start out as noise and an additional 10% degrade into additional noise because they are less effective, as illustrated in Figure 1.

At that rate, the proportion of noise to signal would grow from 5% at the outset to become the majority of all alert traffic by year 10. A few years later, the number of actionable alerts—or those that should be acted upon based on the alerts' quality, priority and other contextual data—would level off and begin to fall as growth in the number of noisy alerts accelerates. A company that began monitoring in the year 2010 would have more than three times as many noisy alerts as actionable ones by 2022. This hypothetical case reflects the reality in many organizations, where most alert data is unactionable noise.

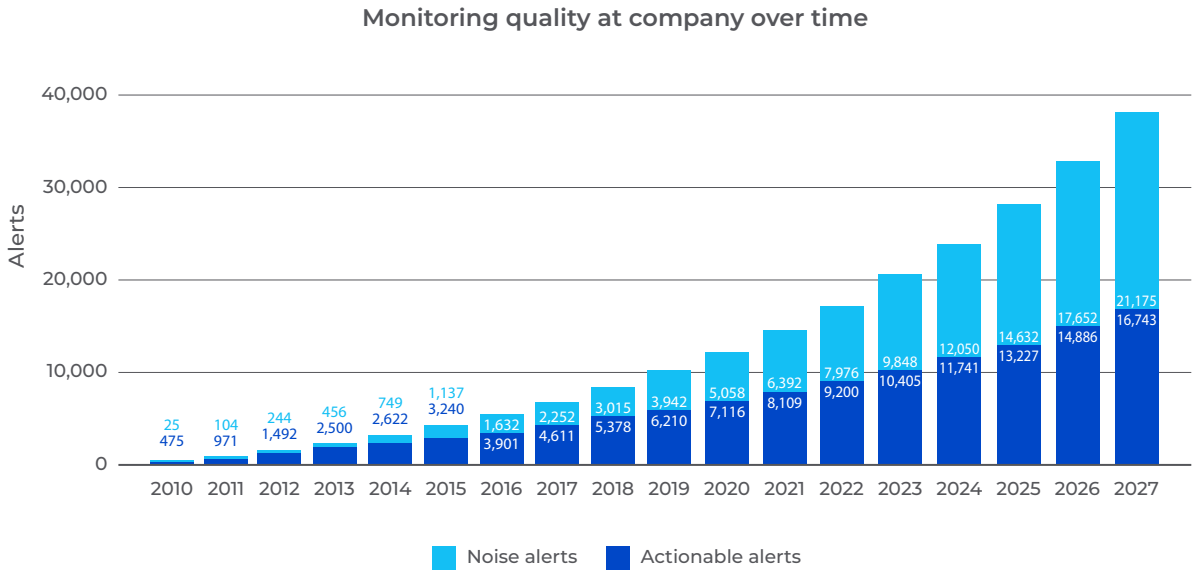

Figure 1. Monitoring alert quality at a hypothetical company over time

Assessing and managing alert quality

To reduce alert noise and continually improve the alerting environment, organizations need to categorize the "quality" of different alerts and differentiate those that are actionable and those that just generate noise. Organization-specific definitions for these quality levels can follow these general guidelines:

Low quality

Describes alerts that are either misconfigured or lack meaningful information required to support any action by the response team. Both present overhead without value.

The result: alerts are ignored.

Medium quality

Indicates at least the minimum level of information and context within alerts to support operator action while lacking some valuable elements such as business context, dependencies or resolution steps. For an alert to be considered medium quality, it must include both:

- The configuration item (CI)
- The symptom of the problem

The result: alerts accumulate until they become critical and are escalated to multiple L1/L2 response teams.

High quality

Alerts meet the criteria for high actionability by support teams, meaning that all available technical and business context data is included. These include:

- Ownership and routing to the assignment group that should respond
- Business impact of the alert to the business, which can be priority level,application tiers, etc.
- Runbooks and knowledge-base URLs on how the alert should be resolved
- Dependencies, including impacted services and applications
- Enrichment

For an alert to be high-quality, it must include ownership and routing information, business impact and either runbooks, dependency or enrichment context.

The result: processes are intelligently automated and incidents are rapidly resolved by the appropriate team.

Categorizing alerts in this way involves consistently applying concrete rules to check for attributes defined by the business as contributing to alert quality. Filling in some of these 'gaps' in the alert payload data allows for a much more accurate alert quality assessment. Furthermore, this type of data enrichment starts to lay the best possible foundation for correlation, prioritization and automation.

Enrichment is critical to creating high-quality alerts

Alerts generated by monitoring and observability tools contain a wealth of low-level technical information, but they often do not contain any operational, topological or other contextual data. Without enriching alerts with this vital metadata, ITOps teams have ongoing maintenance to scan all low-quality alerts and undertake a heuristic approach of what to focus on and what is important. The lack of enrichment makes it difficult to:

- Separate noisy alerts from meaningful alerts, and then eliminate the noise
- Group related alerts together into an incident in real time
- Surface the probable root cause of an incident (including the probable root cause change)
- Use alert metadata to route incidents to the appropriate response team or trigger an automation workflow

To overcome the gap of enrichment, IT operations needs to establish a concept of alert quality to understand which type of enrichment pushes the quality of alert data higher. Focusing on enriching missing technical and business context helps determine and greatly improve correlation, prioritization and automation.

Technical context delivers medium-quality alerts that help support operator actions

The addition of technical context to alerts is a critical process that makes event correlation extraordinarily effective. Monitoring and observability tools do not deliver metadata on the "physical proximity, logical dependence or another dimension that captures the relationship between IT assets and services."[1] Therefore, alert data should be enriched with as much information and technical context found within alerts to support operator actions. Examples of technical context include:

- Continuous integration (CI) info (host/application/service) / Continuous deployment (CD)
- The detected symptom
- Description of the problem

1 Gartner Market Guide for AIOps. Pankaj Prasad, Padraig Bryne, Gregg Siegfried

Once the alerts are processed, more advanced AIOps platforms leverage machine learning (ML) to group and correlate alerts into a small number of incidents by evaluating their properties against three dimensions:

- Time
- Topology
- Context

This ensures alerts have all the necessary context across dimensions that are needed to categorize and define high-quality alerts that enable teams to prioritize incident response.

Business context drives actionability and high-quality alerts

Alerts enriched with the necessary technical context across operational, topological, change and time dimensions make it easier to algorithmically add business context to properly trigger and track automation workflows between AIOps platforms such as BigPanda's and other third-party tools. Business context refers to incident severity, impacted services, business priority and routing information, which add another layer of information that significantly raises the quality of the alerts.

For example, an issue that interferes (or could interfere) with a key revenue-generating application and database would be labeled high priority. That business context would be essential to automatically escalate the incident and assign the right response teams using specific on-call and chat channels. Other types of business context include:

- Teams that should be notified
- Relevant customers
- What's being impacted

Context can be captured within custom tags, making it easy for a team to easily sort, filter, visualize—and act on—alerts. Tags also include the necessary payload data to establish or further solidify pre-established escalation paths and reduce response times by guiding operators through planned response scenarios and removing unnecessary guesswork under pressure.

Summary: Strategic pillars for improving alert quality

ITOps teams must recognize that continuous improvement of alert quality is more than just noise reduction; raising the quality of alerts and incidents is the goal that empowers staff to react, route and remediate much more effectively.

Less is more
Resolve the clutter of alerts and incidents to reduce the quantity and increase quality, delivering actionable insights to response teams that improve efficiency and resolution times.

Context is everything
Enrich alerts with operational, topological, change and time-based dimensions to allow for disjointed alerts to be correlated effectively into incidents with rich, actionable and accurately prioritized incidents.

Quality is evolutionary
Build processes that foster repeatability and provide key performance indicators (KPIs) for assessment and improvement over time.

Measuring and reporting on alert quality

Alerts can be directly assessed for quality based on checklists of the contextual information they contain. However, measuring the "actionability" of low, medium and high-quality alerts requires an ability to connect correlated alerts with operator actions and use analytics to measure outcomes like mean time to detection (MTTD), response and resolution (MTTR) based on alert quality.

This requirement highlights the value of powerful, low-touch analytics being continually available to both operators and management. Dashboards and visualizations should be used to monitor outcomes and provide the basis for tuning processes, enrichment rules and correlation patterns to optimize incident quality.

The Sankey diagram in Figure 2 displays alerts from various monitoring tools on the left side. The middle orange bar represents high, low and noisy alerts. The green bars on the right display operator action based on the quality level from each monitoring tool.

ITOps can optimize alert quality from specific tools, such as adding enrichment to make payloads from a high-volume, low-quality source more actionable. Likewise, low-volume sources of low-quality alerts may represent tool rationalization opportunities, enabling a different tool to cover a given domain with higher-quality alerts. Retiring unneeded tools in this process can save on licensing costs at the same time it improves monitoring coverage and simplifies the observability environment.

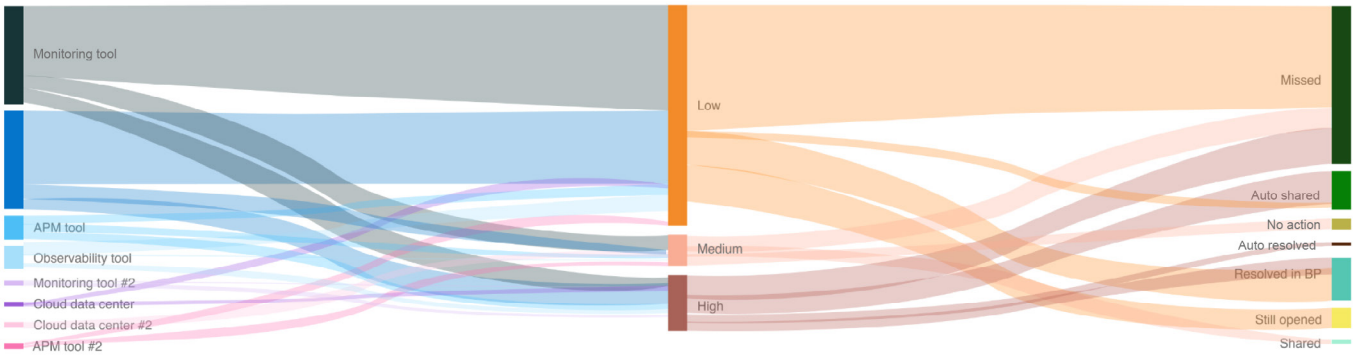

Figure 2. Sankey diagram: Alert flow analysis—Top 10 sources

IT operations teams can also report on alert quality over time to see the impact of payload standardization and governance and where improvements can be made. Figure 3 shows alert payload quality over time. The shaded-gray area shows total alerts before correlation, which visualizes how many events are reduced from the team based on high, medium and low-quality alerts. The purple line shows noise reduction as a percentage of time.

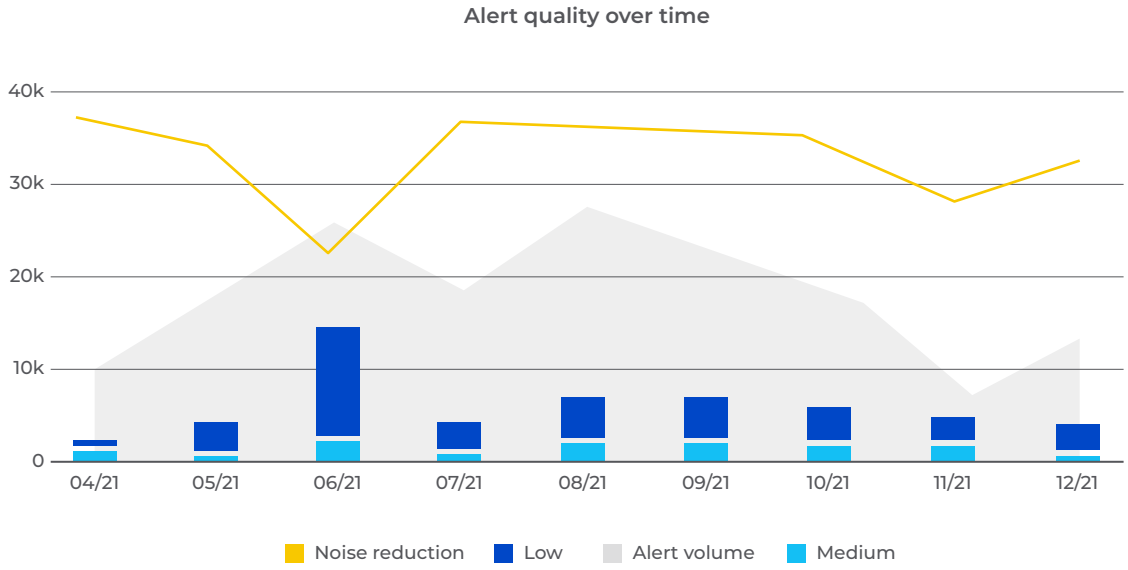

Figure 3. Alert quality over time

Best practices for building high-quality alerts within ITOps

Raising the quality of alerts and incidents requires more than proper tools and configurations. As an ongoing initiative, it involves long-term commitment. It must be supported by stakeholders in domains across the organization, including application, network and infrastructure owners who control alert sources. Some degree of cultural shift may be required to create a shared sense of value, making the transition more than just a technical one. The best practices discussed in this section have proven valuable to adopters as they navigate requirements in their own organizations.

Best practice 1: Focus on the domain within your control

Improvements to the quality of alerts and incidents naturally benefit from buy-in across the organization. Focus on a single domain where there is low alert quality and a high level of understanding of both the technical and business context. That level of control and understanding will allow you to address "low hanging fruit" by easily filling the gaps and adding critical information to existing alerts to improve the quality. By establishing meaningful KPIs and illustrating the improvement in a single domain through analytics, visualizations and dashboards, the value of the quality initiative can be demonstrated to encourage momentum.

Best practice 2: Be guided by business context

ITOps decision makers must be guided by the business impacts of technology issues, rather than the technology issues themselves. For example, a performance degradation on a flagship revenue-generating application may be a higher priority than the total outage of a less prominent one. For automated processes to make that distinction, alerts must include business context that has been defined, reviewed and agreed upon by multiple teams to set resolution priority.

Best practice 3: Define cross-functional review processes to drive effectiveness

A healthy alert and incident management practice regularly identifies opportunities to standardize, measure and improve incident response workflows across cross-functional teams. Regular reviews of KPIs and business outcomes with various stakeholders, from ITOps to platform owners to DevOps and SRE teams, should be conducted to identify successes, shortcomings and opportunities for improvement. Directly involving stakeholders helps establish a culture of ownership and encourages commitment to improving alert and incident quality.

Best practice 4: Monitoring alert hygiene

The alerting environment itself must be maintained on a regular basis to ensure alerts are categorized, escalated and resolved in a timely fashion. This ensures monitoring KPIs are measured correctly and are not incorrectly skewed when, for example, bulk actions to resolve unactioned alerts are taken intermittently during the week. With good hygiene, monitoring KPIs will depict a more accurate representation of what response teams are handling and allow for easier demonstration of progress toward achieving business and technology outcomes.

BigPanda

Enrichment is the best-kept secret for AIOps success

High-quality alerts are fundamental to optimizing ITOps to be proactive, efficient and effective. The precursor to any alert quality initiative begins with harnessing a mindset shift to always improving alert quality. It isn't a one-time setup; rather it is a succession of processes that starts by looking at low-quality, noisy alerts and defining the necessary requirements to standardize alert quality standards and service-level agreements (SLAs). Enrichment is the key driver to filling information gaps within alerts that reduces alert noise, increases operator efficiency and builds a foundation for actionability throughout the incident lifecycle.

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