


Monitoring and Observability Tool Effectiveness for IT Event Management

Insights on incident detection and noise reduction in the enterprise



Contents

| | |
|----|--|
| 3 | Overview |
| 4 | Top 5 key takeaways |
| 5 | AIOps platform usage |
| 6 | Monitoring and observability tool integrations |
| 7 | Number of inbound integrations |
| 8 | Types of inbound integrations |
| 11 | Integrations and events by vendor and solution |
| 15 | Monitoring and observability tool effectiveness matrix |
| 16 | Top 4 monitoring and observability tool trends |
| 17 | Detection benchmarks |
| 18 | Events |
| 25 | Alerts |
| 32 | Incidents |
| 39 | Actioned incidents |
| 48 | Conclusion |
| 49 | About this report |
| 49 | Methodology |
| 49 | Citing the report |
| 50 | Demographics and firmographics |
| 54 | About BigPanda |



Overview

Data is growing beyond the limits of human scalability, which limits IT operations (ITOps) teams' ability to detect and respond to incidents quickly, before end-users submit tickets to the service desk.

While monitoring solutions and full-stack observability tools are evolving coverage to support hybrid cloud infrastructures, including application performance monitoring (APM), digital experience monitoring (DEM), IT infrastructure monitoring, and network performance monitoring, this expansion is primarily developer-focused.

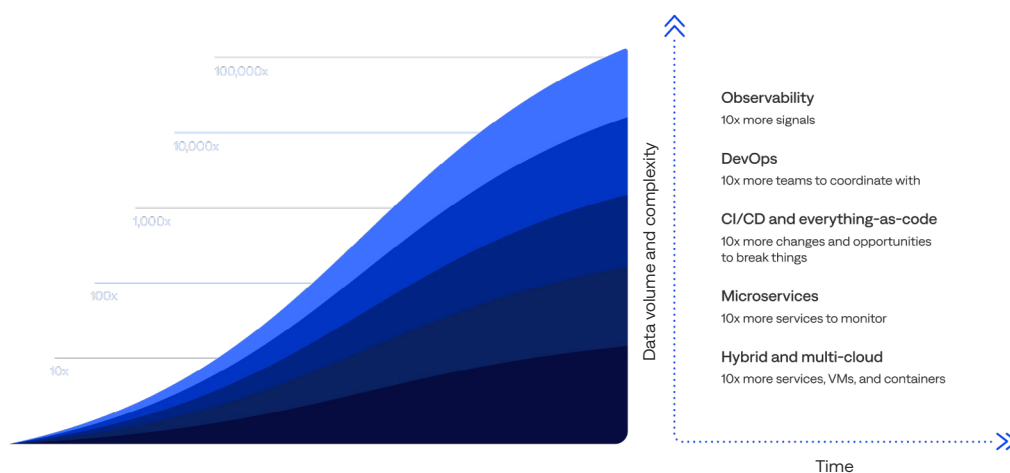
Engineers design observability tools to ensure adequate coverage and visibility into applications, services, and infrastructure health. However, overwhelming alert noise makes detecting and responding to alerts before they become incidents challenging for humans. ITOps teams that are decentralized from observability teams must sift through a high volume of fragmented data to identify what's important and actionable and find the context needed to triage and respond to an alert before it becomes an incident.

To address these challenges, many enterprises are adopting AIOps and event management capabilities to reduce observability noise and identify actionable alerts by clustering related symptoms across monitoring and observability tools. Enriching alerts and incidents with topology, change, and configuration management database (CMDB) data provides ITOps teams with the context needed to reduce triage and response times by creating tickets in IT service management (ITSM) platforms and automatically assigning them.

This report includes analysis and insights about L1 detection and noise reduction, including the effectiveness of monitoring tools and observability platforms for IT event management. It's the first report based on enterprise usage of the [BigPanda platform](#).



AIOps (artificial intelligence for IT operations) or **event intelligence** uses artificial intelligence (AI), machine learning (ML), and data analytics at the event management level to augment, accelerate, and automate manual efforts in the IT event management process. Key characteristics include cross-domain event ingestion, topology assembly, event correlation and remediation augmentation.




Keeping the digital world running is getting harder



Top 5 key takeaways


The findings indicate enterprises are struggling with vast amounts of data and signals, and the BigPanda platform is helping teams produce higher-quality incidents and less noise.

Here are the top five key takeaways from the report:



Full monitoring coverage doesn't equal value.

Most enterprises are drowning in data, creating millions of events (9.6 million, on average) annually. In addition, 27% occurred on weekends, which is bad news for those on call. Yet only 18% of incidents were actioned on average, underscoring the disconnect between the belief that comprehensive observability coverage of applications, services, and infrastructure equates to better ITOps, incident management, and customer outcomes.




Noise reduction is exceptionally high for BigPanda customers.

Most (82%) organizations achieved at least 97% noise reduction with BigPanda, and over half reduced noise by 99.5–99.9%, showing the power of deduplication, correlation, and suppression when properly implemented with a unified AIOps platform.




Full-stack observability is still an illusion.

Despite enterprises using a median of 20 inbound tool integrations, few monitoring and observability tools had high coverage and actionability rates. In addition, despite their popularity, open-source tools had a low impact at the enterprise level. This suggests that even widely used tools often generate noise instead of insight.



Alert enrichment boosts actionability.

Organizations enriched 60% of alerts with topology data for all incidents, but enriched 77% of alerts linked to actioned incidents. This comparison indicates that enrichment from CMDB, cloud and virtualization management, service discovery, and APM tools significantly improves alert quality and operator confidence.



Healthy correlation is powerful.


Nearly half (49%) of organizations fell into the healthy alert-to-incident correlation range (40–75%), and those that did showed stronger filtering, a more manageable incident volume, and a higher operational signal fidelity.

[View the infographic](#)

AIOps platform usage

The [BigPanda platform](#) assists teams through every stage of the incident management workflow or event-to-incident lifecycle. All organizations included in this report used BigPanda for event management, which includes event detection, deduplication, filtering, correlation, and enrichment.


Many organizations also used BigPanda incident management features for [advanced insights](#) (released in March 2024) with AI-powered, multisource IT data analysis to instantly automate and accelerate incident triage.



Incident management workflow

In addition, [BigPanda professional services](#) help organizations achieve their business objectives by providing expertise for every stage of the AIOps journey. For example, they can opt to work with a dedicated BigPanda resident solution architect (RSA) for at least half a day and up to four days per week in a 6- or 12-month engagement. A quarter (25%) of organizations worked with a dedicated BigPanda RSA, including 75% that worked with an RSA at least one full day per week.

This report focuses on the detection phase of the incident management workflow, starting with the inbound monitoring and observability data sources integrated with BigPanda.

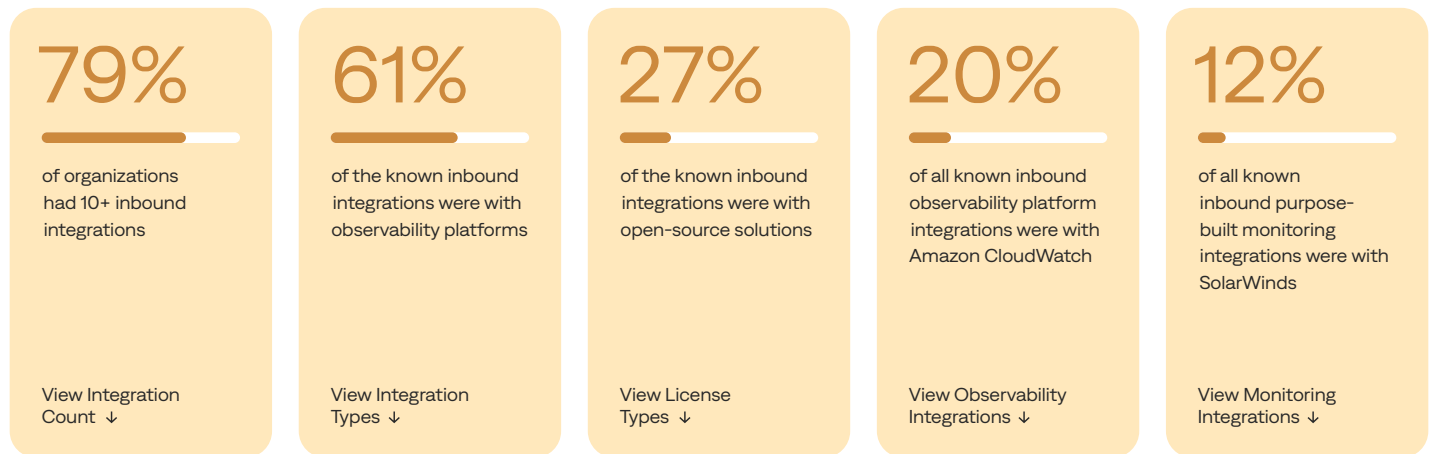


Monitoring and observability tool integrations

The first step in incident detection is configuring inbound [monitoring and observability integrations](#) with BigPanda to receive events. BigPanda includes about 50 [standard integrations](#) with popular tools, and customers can also create custom integrations. This data feeds into the [BigPanda IT Knowledge Graph](#).

This section reviews the number and types of inbound integrations, how many were open-source or proprietary, which monitoring and observability tool vendors and solutions were used the most and generated the most events, their [effectiveness](#), and the [top four monitoring and observability tool trends](#).

Key monitoring and observability tool integration highlights:



“Observability is a journey. BigPanda AIOps is a key part of this journey for us. As we scale and grow the business, it’s integral for us to bring in automation and integration with other tools and technologies. Don’t wait to start your AIOps journey once you are overwhelmed with alerts. Start early to get a single pane of glass to understand which monitoring tools you really need.”

–Vice President of Information Technology, Manufacturing Enterprise




Number of inbound integrations

The number of inbound integrations per organization ranged from one to 198, with a median of 20.

- Most (79%) organizations had 10 or more inbound integrations, including 36% with 25 or more, and 19% with 50 or more.
- Nearly half (42%) had between 10 and 24 inbound integrations.
- Just 22% had fewer than 10 inbound integrations.

79%

of organizations had 10+ inbound integrations



Number of inbound integrations with BigPanda



Types of inbound integrations

This section reviews the known inbound integrations by category and license type.

Categories

Each known inbound integration vendor or solution was grouped into one of three categories:


- Observability platforms—such as Amazon CloudWatch, Datadog, Dynatrace, Prometheus, and Splunk—offer multiple monitoring types in a unified view. Most (92%) organizations had at least one inbound integration with an observability platform. Nearly two-thirds (61%) of the known inbound integrations were between BigPanda and 20 observability platforms. However, only 19% of the known inbound integration vendors or solutions fell into this category. Nearly a quarter (22%) of the events were from observability platforms.
- Purpose-built monitoring tools—such as Cisco ThousandEyes, Microsoft System Center Operations Manager (SCOM), Site24x7, SolarWinds, and Zabbix—offer one or two more focused, specialized types of monitoring. Over three-quarters (82%) of organizations had at least one inbound integration with a purpose-built monitoring tool. A third (33%) of the known inbound integrations were between BigPanda and 72 purpose-built monitoring tools. Over two-thirds (69%) of the inbound integration vendors or solutions fell into this category. Half (50%) of the events were from purpose-built monitoring tools.
- Non-monitoring platforms, systems, or tools—such as Cribl, Databricks, MongoDB, Moogsoft, and ServiceNow—bring in non-monitoring-related data. Over a quarter (29%) of organizations had at least one inbound integration with a non-monitoring tool. Just 6% of inbound integrations were between BigPanda and 12 non-monitoring tools. Only 11% of the inbound integration vendors or solutions fell into this category. More than a quarter (27%) of the events were from non-monitoring platforms, systems, or tools.

61%

of the known inbound integrations were with observability platforms

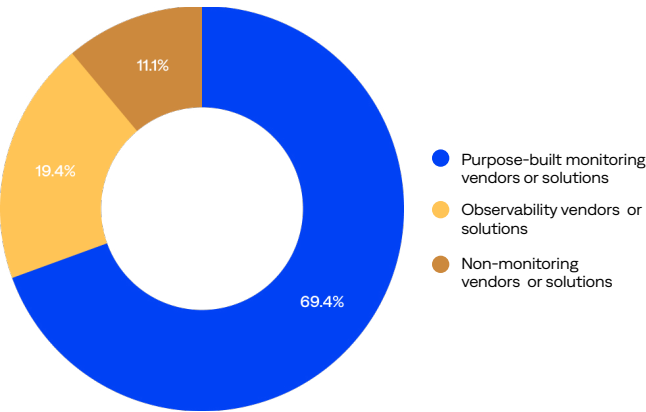


Percentage of inbound integrations




Percentage of inbound integrations by inbound integration category

Percentage of vendors or solutions



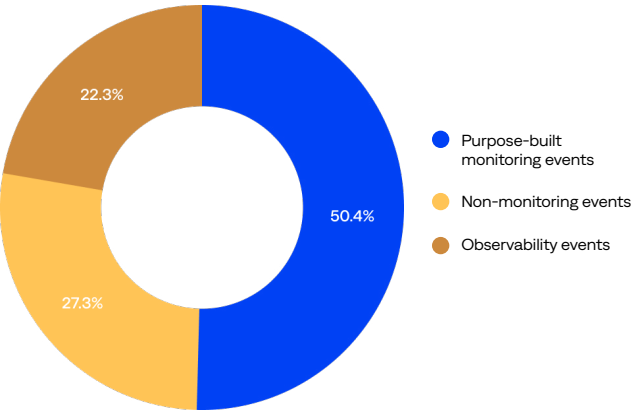
Percentage of vendors or solutions by inbound integration category

Percentage of organizations



Percentage of organizations by inbound integration category

Percentage of events



Percentage of events by inbound integration category


View the [most integrated vendors and solutions](#), as well as the [effectiveness of each monitoring and observability vendor or solution](#).



License types

Nearly three-quarters (73%) of the known integrations were with proprietary vendors, such as Cisco AppDynamics, Datadog, LogicMonitor, New Relic, and VMware vRealize Operations (vROps).

The remaining 27% were with open-source solutions, such as the ELK Stack (Elasticsearch, Logstash, and Kibana), Grafana, Jenkins, Prometheus, and Sensu.



Percentage of inbound integrations with BigPanda that were proprietary or open-source

The following table compares the percentage of proprietary versus open-source inbound integrations by vendor or solution type.

| Vendor or solution type | Proprietary | Open source |
|-------------------------------|-------------|-------------|
| Observability platform | 65% | 35% |
| Purpose-built monitoring tool | 81% | 19% |
| Non-monitoring tool | 50% | 50% |

[View the effectiveness of proprietary and open-source monitoring and observability tools.](#)

27%

of the known inbound integrations were with open-source solutions




Integrations and events by vendor and solution

Altogether, 104 known vendors and solutions had inbound integrations with the BigPanda platform.

- The vendor or solution that had the most integrations with BigPanda was Amazon CloudWatch (12%), followed by New Relic (11%) and Prometheus (9%).
- The vendor or solution that generated the most events in BigPanda was Cribl (13%), followed by CA Spectrum (10%) and CGI DigiOps (9%).
- The top 10 most-integrated vendors or solutions were all observability platforms or purpose-built monitoring tools, but three out of 10 of the top event generators were non-monitoring tools.

Therefore, the number of integrations didn't necessarily correlate with the number of events generated per vendor or solution.



Percentage of integrations and events for each vendor or solution (top 10 by number of integrations)




Prometheus



Observability platform integrations and events

Nearly two-thirds (61%) of the known inbound integrations were between BigPanda and 20 observability platforms, which were responsible for 22% of the events.

- The top five most integrated observability platform vendors or solutions included Amazon CloudWatch (20%), New Relic (19%), Prometheus (14%), Datadog (9%), and Grafana (7%).
- The top five observability platform vendors or solutions that generated the most events were Prometheus (22%), LogicMonitor (20%), Datadog (16%), Cisco Splunk (14%), and Cisco AppDynamics (12%).
- Only two—Prometheus and Datadog—were in the top five for both the percentage of integrations and events.
- Nearly two-thirds (65%) of the integrated observability platform vendors or solutions were proprietary, and the remaining 35% were open source.



Percentage of integrations and events generated for each observability platform vendor or solution (by number of integrations)

[View the effectiveness of each observability platform.](#)

20%


of all known inbound observability platform integrations were with Amazon CloudWatch



Purpose-built monitoring tool integrations and events

A third (33%) of the known inbound integrations were between BigPanda and 72 purpose-built monitoring tools, which were responsible for 50% of the events.

- The top five most integrated purpose-built monitoring vendors or solutions included SolarWinds (12%), Nagios (10%), CGI DigiOps (9%), SCOM (8%), and VMware vRealize Operations (vROps) (6%).
- The top five purpose-built monitoring tool vendors or solutions that generated the most events were CA Spectrum (20%), Hydra (16%), Logstash (11%), SolarWinds (6%), and SCOM (3%).
- Only two—SolarWinds and SCOM—were in the top five for both the percentage of integrations and events.
- Over three-quarters (81%) were proprietary, and the remaining 19% were open source.
- The most commonly integrated purpose-built monitoring tools focused on IT infrastructure monitoring (31%), network performance monitoring (17%), managed infrastructure services (11%), DEM (7%), and backup and recovery software (5%).



Percentage of integrations and events for each purpose-built monitoring tool vendor or solution (top 10 by number of integrations)

12%

of all known inbound purpose-built monitoring integrations were with SolarWinds



[View the effectiveness of popular purpose-built monitoring tools.](#)




Non-monitoring tool integrations and events

Just 6% of the known inbound integrations were between BigPanda and 12 non-monitoring tools, which were responsible for 27% of the events.

- The top five most integrated non-monitoring vendors or solutions included Cribl (26%), ServiceNow (14%), Sensu (13%), Apache Kafka (13%), and MongoDB (9%).
- The top five non-monitoring tool vendors or solutions that generated the most events were Cribl (48%), Sensu (28%), Apache Kafka (15%), Moogsoft (4%), and IBM Netcool (3%).
- Three—Cribl, Sensu, and Apache Kafka—were in the top five for both the percentage of integrations and events, and Cribl was the top for both.
- Half (50%) of the integrated non-monitoring vendors or solutions were proprietary, and the remaining 50% were open source.
- The most commonly integrated non-monitoring tools were from telemetry pipelines (40%), event brokers (15%), ITSM platforms (13%), database management systems (11%), event intelligence solutions (7%), and data integration tools (7%).

26%

of all known inbound non-monitoring integrations were with Cribl



Percentage of integrations and events for each non-monitoring tool vendor or solution (by number of integrations)

Cribl servicenow

Sensu by sumo logic kafka

mongodb. Apache Airflow

IBM databricks

moogsoft Jenkins

Microsoft SQL Server Power BI



Monitoring and observability tool effectiveness matrix

This section compares the quality ([actionability rate](#)) and coverage (percentage of [actioned incidents](#)) of the incidents generated by each monitoring and observability vendor or solution to identify high-quality tools and noisy tools that need improvement. It includes a matrix with four quadrants:

2. High-quality, low-coverage

These optimized, high-performance tools in the upper-left quadrant generate fewer incidents but maintain a high rate of actionable incidents. Ideal for targeted use cases, they deliver substantial value when deployed and may be candidates for broader adoption.

1. High-quality, high-coverage

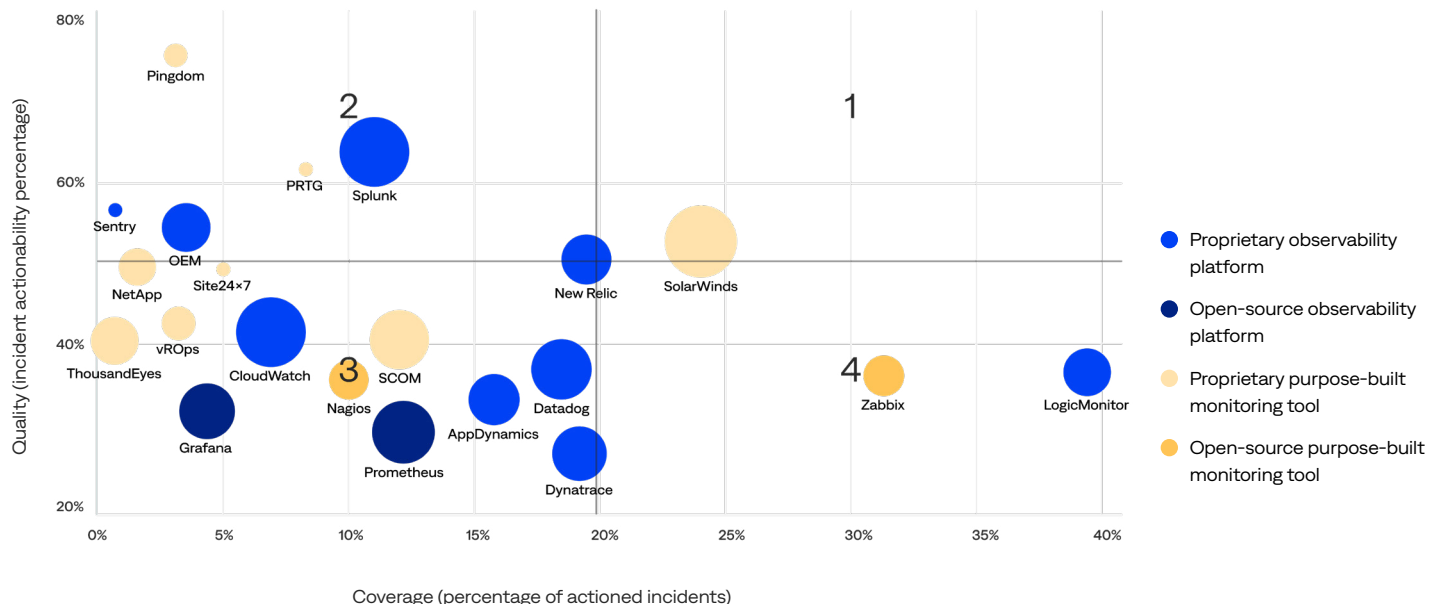
These signal-rich, low-noise tools in the upper-right quadrant are widely deployed and consistently deliver actionable incidents. They balance signal volume and strength, making them key assets in effective observability strategies.

3. Low-quality, low-coverage

These underutilized tools in the bottom-left quadrant are less prevalent and show lower signal quality, demonstrating opportunities to evolve through better integration, improved configuration, or rationalization. They may be in early adoption phases or used for narrower scopes.

4. Low-quality, high-coverage

These scalable but noisy tools in the bottom-right quadrant contribute significantly to incident volume with fewer actionable insights. While widely used, they may benefit from tuning or configuration improvements to reduce noise and increase operational value.



Monitoring and observability tool effectiveness matrix (bubble size increases with customer usage)



Top 4 monitoring and observability tool trends

The data show key trends and insights for purpose-built monitoring tools and observability platforms.

No tool wholly owns the top-right quadrant—the observability landscape remains fragmented with no clear leader.

- While tools like New Relic and SolarWinds come closest, there are no clear leaders that combine both widespread usage and consistently high actionability (top-right quadrant).
- This signals that even the strongest platforms have room to grow, and the observability industry is still evolving toward optimal performance at scale.

Open-source tools remain low-impact with limited adoption.

- Open-source tools such as Prometheus and Grafana (dark blue) and Nagios and Zabbix (dark yellow) appear in the lower-left or lower-right quadrants with smaller bubbles and lower signal quality.
- Despite their popularity among developers, most open-source observability platforms and monitoring tools have yet to deliver high-value, enterprise-grade observability outcomes.

Some high-coverage tools fall short on signal quality.

- Tools in the bottom-right quadrant contributed a large share of incidents, indicated by their position far along the coverage axis and their sizable bubble sizes, representing broad adoption.
- However, their lower actionability highlights that high usage does not necessarily translate to high operational value. These scalable but noisy tools may benefit from improved configuration and tuning to reduce noise and enhance the precision of alerts.

Purpose-built monitoring tools tend to align as either specialists or stragglers.

- Purpose-built monitoring tools either fell in the top-left quadrant (optimized, high-performance tools), like PRTG and Solarwinds Pingdom, or the bottom-left quadrant (underutilized tools), with lower adoption and weaker signal quality.
- This indicates that while some purpose-built monitoring tools deliver substantial niche value, others have yet to evolve into broader observability assets.

For organizations investing in observability, the challenge is identifying which tools deserve broader deployment and which require refinement or reevaluation.



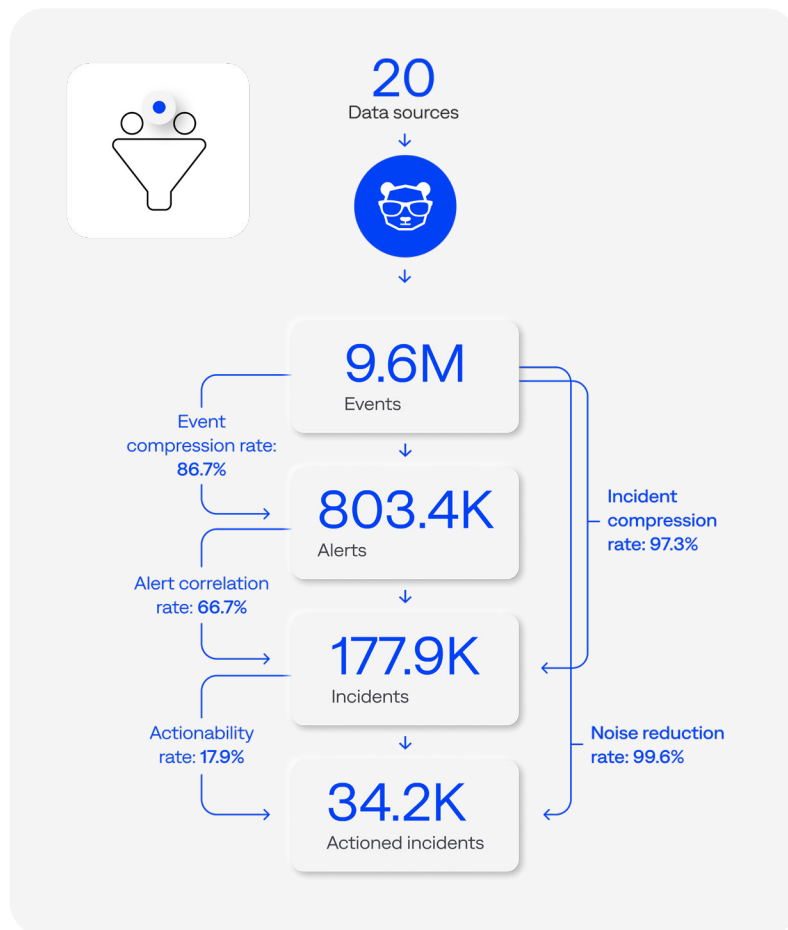
Detection benchmarks

Teams use BigPanda to detect events during pipeline processing, including:

- Correlating alerts across applications and services
- Enriching alerts for greater intelligence
- Minimizing alert noise and fatigue

This section reviews the BigPanda [event-to-incident lifecycle](#).

Key detection highlights:



Pipeline processing funnel (events-to-incidents workflow or lifecycle) median detection benchmarks per organization (n=125)



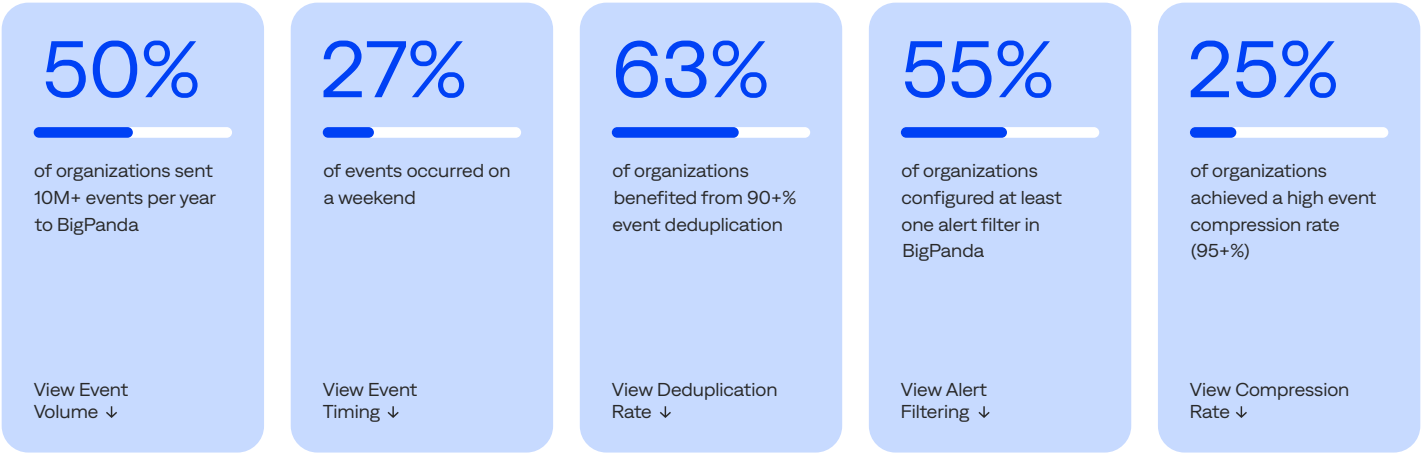
Events

An **event** is a point in time that represents the state of a service, application, or infrastructure component.

The pipeline process starts when BigPanda receives and ingests event data from monitoring and observability tools. These tools can generate events when potential problems are detected in the infrastructure.

This section reviews the volume of events, when events tend to occur, and event compression.

Key event highlights:



Event volume

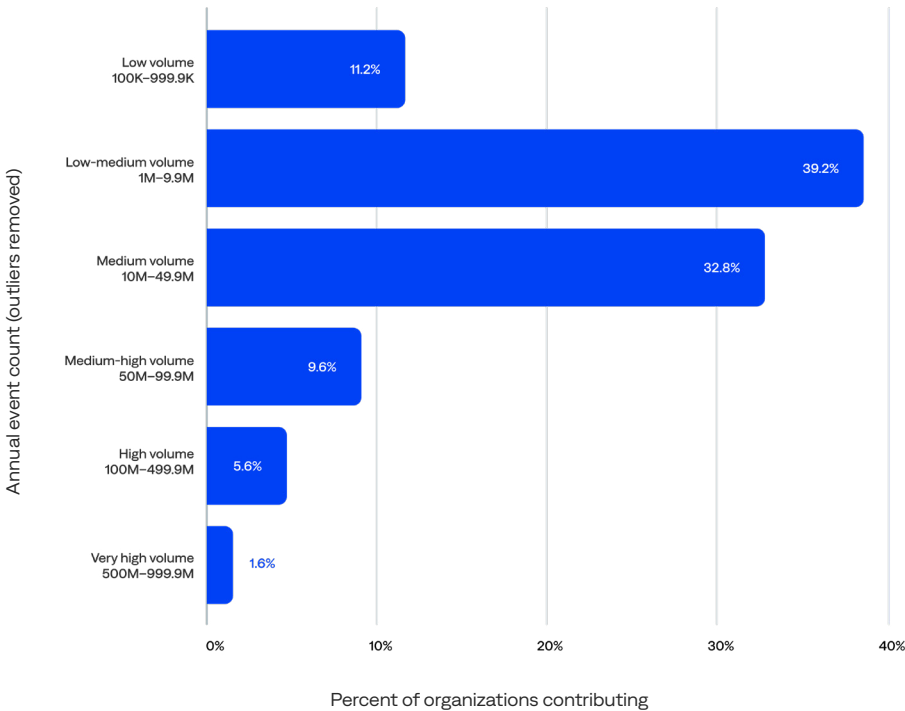
BigPanda ingested nearly 6 billion events from inbound monitoring and change integrations.

When we remove the five outliers (fewer than 100,000 and 1 billion or more annual events), **BigPanda ingested 4.5 billion events**. The median annual events per organization was 9.6 million, and the median daily events per organization was 28,623.

- Half (50%) of the organizations generated at least 10 million events, including 17% that contributed 50 million or more events, representing 55% of the total annual event volume.
- Nearly three-quarters (72%) generated between 1 million and 50 million events, representing 20% of the total annual event volume.
- Typical or moderate-sized annual event counts ranged from 10 million to 50 million. A third (33%) of organizations fell in this medium volume range, representing 17% of the total annual event volume.
- Over a third (39%) generated a low-medium volume of annual events in the single-digit millions (at least 1 million but fewer than 10 million annual events), representing just 3% of the total yearly event volume.
- The remaining 11% had minimal volume (at least 100,000 but fewer than 1 million events per year), which may indicate that they were still onboarding.

50%

of organizations sent 10M+ events per year to BigPanda



Annual event volume (n=125)




When events occur

This section reviews when events occur based on the UTC (Coordinated Universal Time, also known as Greenwich Mean Time or GMT) time zone.

By month of the year

The event count ranged from about 374.3 million to 540.2 million per month.

- The most events occurred in October (10.3%), followed by November (9.9%) and July (9.2%).
- The fewest events occurred in April (7.1%), followed by January (7.2%) and February (7.3%).
- When we compare by seasons, 29% happened in September, October, and November, 26% in June, July, and August, 23% in March, April, and May, and 22% in December, January, and February.



Percentage of total events by month in UTC (n=114)

29%

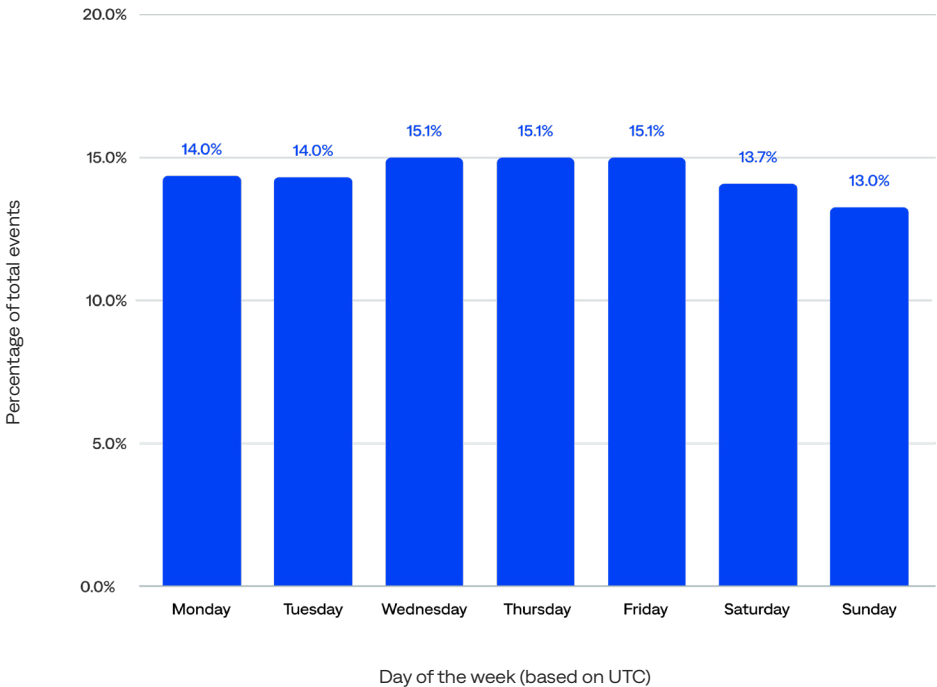
of events occurred in
September, October, and
November



By day of the week

As far as what day of the week events tend to happen, the data show that:

- Nearly three-quarters (73%) of events occurred on weekdays; Monday–Friday consistently saw higher activity, averaging over 767 million events per day.
- The peak event days were Wednesday, Thursday, and Friday, with 15% each or 45% total, compared to 14% each (28% total) for Monday and Tuesday.
- Weekends show a slight drop-off in events, with about 14% on Saturday and 13% on Sunday. However, 27% of events still occurred on weekends, which is bad news for those on call.



Percentage of total events by day of the week in UTC (n=114)

27%



of events occurred on a weekend



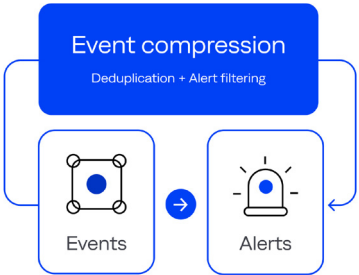
Event compression

Event compression is the number of events compressed into alerts. It consists of deduplication and alert filtering, which help prevent events from becoming alerts. Therefore, higher event compression rates correlate with less alert noise.

The median event compression rate was 87%.


Many organizations had achieved high compression, while others had room to improve (low and average compression):

- One in four (25%) had achieved **high compression** (95% or higher), suggesting strong use of deduplication and alert filtering. These organizations likely experienced less alert fatigue, a better signal-to-noise ratio, and lower support overhead.
- The majority (59%) fell into the **average compression** range (70–94%), which suggests they had taken steps to reduce noise but hadn't fully optimized their setup.
- Only 17% were in the **low compression** range (<70%), likely due to early-stage adoption (still onboarding) or poor use of deduplication and alert filtering. They may have experienced high alert noise, a poor signal-to-noise ratio, or incomplete configurations.



25%

of organizations achieved a high event compression rate (95+%)



Event compression rate range and tier by organization (n=125)




Deduplication

Also known as **event deduplication**, **deduping** is the process by which BigPanda eliminates redundant data to reduce noise and simplify incident investigation. **Deduplicated events** are events that were removed as precise duplicates.

BigPanda has a built-in **deduplication process** that reduces noise by intelligently parsing incoming raw events. It groups events into **alerts** based on matching properties. Exact duplicate matches add clutter to the system and are not actionable. BigPanda discards precise duplicates of existing events immediately. However, it merges updates to existing alerts rather than creating a brand-new alert.

The median deduplication rate was 93.6%.

- Most (63%) organizations had deduplication rates of 90% or more, including 42% with rates of 95% or more, and 18% clustering around the 99% mark—enabling them to focus only on incidents that matter.
- The remaining 37% had deduplication rates of less than 90%, likely due to subpar configuration, purposely choosing not to deduplicate, or poor quality data that’s difficult to dedupe well.
- Just 10% of the total events were deduplicated into alerts. In other words, in 2024, BigPanda helped prevent 90% of noisy alerts (event deduplication rate). Put another way, on average, BigPanda helped prevent over 43 million alerts per organization per year.



Deduplication rate per organization

[View the event-to-incident compression rate.](#)



63%

of organizations benefited from 90+% event deduplication

Alert filtering

In the context of BigPanda, **alert filtering** is a feature that allows users to filter out or suppress specific alerts. Filtered-out events are unactionable events that were filtered out using alert filters.


Filtering alerts helps ITOps teams stop duplicate, low-relevance events from being correlated into incidents. Stopping alert noise before it reaches the incident feed enables teams to focus on the most important incidents and spend their time and effort on the most critical issues.

Alert filtering affects alerts after they have been normalized and enriched. The added context of the enrichment process enables teams to filter events based on alert metadata and enrichment tags.

Over half (55%) of organizations had configured at least one alert filter in BigPanda. The remaining 45% likely configured alert filters upstream before they reach BigPanda.

Excluding organizations with no alert filters, the median alert filters per organization was two. Of those organizations that had configured alert filters:

- Over two-thirds (67%) configured 1–5.
- Nearly a quarter (22%) configured 6–24.
- About one in ten (11%) configured 25 or more.



Number of configured alert filters per organization (n=72)

55%

of organizations configured at least one alert filter in BigPanda



Alerts

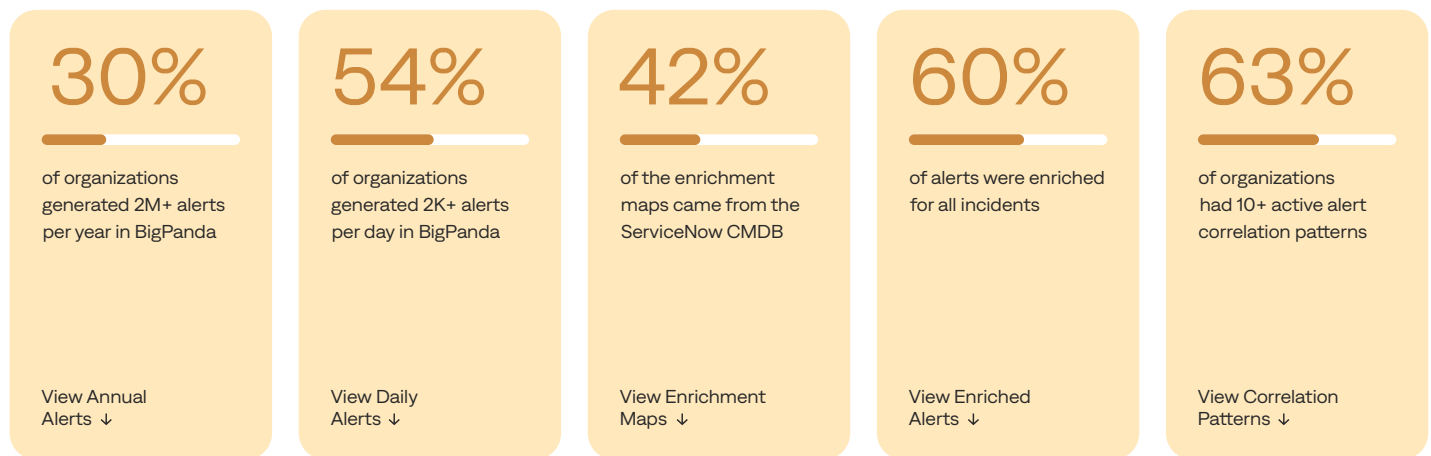
An **alert** is the combined lifecycle of a single system issue.

Monitoring and observability tools generate events when potential problems are detected in the infrastructure. Over time, status updates and repeat events may occur due to the same system issue.

In BigPanda, raw event data is merged into a singular alert so that teams can visualize the lifecycle of a detected issue over time. BigPanda correlates related alerts into incidents for visibility into high-level, actionable problems.

This section reviews the annual and daily alert volume and information about alert enrichment and correlation patterns.

Key alert highlights:



“Before BigPanda, we had times when multiple incidents would trigger alerts from three or four different monitoring and observability tools. With all that noise, we didn’t have visibility into alert impact, and could not quickly identify the root cause to know where to focus our triage efforts. With BigPanda, our IT noise is not only reduced, but we can identify the root cause in real time—who the responsible team is, who owns the alerting service, etc.—which is significantly reducing our MTTR.”

—Staff Software Systems Engineer, Manufacturing Enterprise



Alert volume

This section reviews the annual and daily alert volume for the organizations included in this report.


Annual alert volume

BigPanda generated over 587 million alerts in 2024. After filtering out the five event outliers, the total alert count was over 493 million, and the median annual alert volume was 803,406.

- **Low and low-medium alert volume:** Over two-thirds (69%) of organizations generated fewer than 2 million alerts per year.
- **Medium and medium-high alert volume:** Nearly a quarter (22%) generated at least 2 million but fewer than 10 million alerts per year.
- **High and very high alert volume:** Only 8% generated more than 10 million alerts per year.

30%

of organizations generated 2M+ alerts per year in BigPanda



Annual alert volume (n=125)




Daily alert volume

The median daily alert volume was 2,350.

- **Low and low-medium alert volume:** Nearly half (46%) of organizations experienced fewer than 2,000 alerts per day, including 22% with fewer than 500.
- **Medium and medium-high alert volume:** About a third (34%) experienced at least 2,000 but fewer than 10,000 alerts per day.
- **High and very high alert volume:** One in five (20%) experienced 10,000 or more alerts per day, including 9% with 25,000 or more

54%

of organizations generated 2K+ alerts per day in BigPanda



Daily alert volume (n=125)



Alert enrichment

Alert enrichment (or **event enrichment**) refers to adding additional context, such as CMDB, operational, and business logic data, to alerts and events from external data sources.

The [BigPanda event enrichment engine](#) leverages existing relationship information for mapping enrichments, quickly improving alert quality and reducing time to triage by providing cross-domain alert enrichment with rich contextual data. This enrichment enables operators to identify meaningful patterns and promptly take action to prioritize and mitigate major incidents.

A higher percentage of [data enrichment](#) leads to better-quality incidents.

Low alert enrichment could mean organizations pre-enrich alerts before sending them to BigPanda, maintain poor CMDB workflows, or have poor CMDB quality.

High alert enrichment could indicate a rigid process in which alerts are highly standardized and thus always matched against an external data source.

Most organizations had configured the rules to create enrichment maps (94%), the rules to extract data from the enrichment maps to an external source such as ServiceNow (96%), and the composition rules for enrichment (97%).

This section reviews details about the enrichment integrations and the enriched alerts.

94%

of organizations had configured the rules to create enrichment maps

“BigPanda has significantly helped with deduplicating, correlating, and automating our process. The enrichment data we process through BigPanda enables us to create more specific and insightful alert tags.”

—Supervisor of IT Operations, Healthcare Enterprise




Enrichment integrations and maps

BigPanda includes four [standard enrichment integrations](#) that ingest contextual data from configuration management, cloud and virtualization management, service discovery, APM, topology, and CMDB tools (Datadog, Dynatrace, ServiceNow, and VMware vCenter) to create a full-stack, up-to-date model that enriches BigPanda alerts. Customers can also create custom enrichment integrations.

This section reviews which maps (tables) the organizations uploaded to enrich their data. The organizations in this report uploaded 6,160 enrichment maps.

- Over a third (38%) of organizations used a standard enrichment integration (Datadog, Dynatrace, ServiceNow, and/or VMware vCenter), and 78% of the enrichment maps came from standard integrations.
- The known data source with the most integrations was the ServiceNow CMDB (26%).
- Most enrichment maps came from the ServiceNow CMDB (42%) and Dynatrace (30%)



Percentage of enrichment maps uploaded and organizations using each enrichment data source

42%

of the enrichment maps came from the ServiceNow CMDB




Enriched alerts

Nearly two-thirds (60%) of alerts were enriched for all incidents, and 77% were enriched for actioned incidents (mapping enrichment specifically). The median percentage of alerts enriched for all incidents per organization was 63%, and the median for all actioned incidents was 74%.

- The distribution is polarized, with about 20% of organizations either doing very little enrichment (0–10%) or achieving extremely high enrichment (90–100%). This suggests that organizations don't gradually climb the enrichment ladder—they either commit fully or stay minimal.
- Nearly two-thirds (60%) had enriched at least 50% of their alerts, including 35% that had enriched at least 75% and 17% that had enriched at least 90%.
- Only 20% had enriched less than 10% of their alerts, including 9% that did not enrich alerts. This could represent onboarding organizations or organizations in the early stages of observability maturity.

60%

of alerts were enriched for all incidents



Percentage of alerts that were enriched for all incidents and all actioned incidents per organization




Alert correlation patterns

Correlation patterns set rules to define relationships between system elements, which BigPanda then uses to cluster alerts into incidents dynamically. They define the relationships between alerts using parameters, including the source system, tags, the time window, and an optional filter.

Teams can customize alert correlation patterns to align with the specifics of their infrastructure. They can also enable cross-source correlation, which correlates alerts from different source systems into the same incident.

Correlation patterns are easy to [configure in BigPanda](#). In fact, all organizations had configured correlation patterns. There were 2,723 active correlation patterns, with a median of 14 per organization.

- Nearly two-thirds (63%) of organizations had 10 or more active correlation patterns configured, including 40% with 20 or more.
- The industries with the highest median active correlation patterns configured were transportation (41), energy/utilities (32), and media/entertainment (26). Those with the lowest were telecommunications (9), MSPs (10), and manufacturing (10).
- Most (92%) correlation patterns were non-system-generated.
- Over half (52%) of all active correlation patterns had cross-source correlation enabled.
- Over half (53%) had correlation patterns with one tag, 30% had two tags, and 17% had three or more tags.



Percentage of active correlation patterns configured per organization (n=124)

[View the alert correlation rate.](#)



63%

of organizations had 10+ active alert correlation patterns

“Not only can we see the alerts, but we can evaluate them using correlation that recognizes patterns, connects alerts, and leads to fewer incidents.”

—Head of Automation and Monitoring, Telecommunications Enterprise

Incidents

An **incident** in BigPanda consists of correlated alerts that require attention, such as an outage, performance issue, or service degradation.

As raw data is ingested into BigPanda from integrated tools, the platform correlates related alerts into high-level incidents. Incidents in BigPanda provide context to issues and enable teams to identify, triage, and respond to problems quickly before they become severe.

BigPanda consolidates event data from various sources into a single pane of glass for insights into multi-source incident alerts and the IT environment’s overall health. This enables ITOps, incident management, and SRE teams to investigate and analyze incidents, determine their root cause, and take action easily—all from one screen.

The lifecycle of an incident is defined by the lifecycle of the alerts it contains. An incident remains active if at least one of the alerts is active. BigPanda automatically resolves an incident when all its related alerts are resolved and reopens an incident when a related resolved alert becomes active again.

This section reviews the incident volume, the ratio of alerts correlated into incidents, the ratio of events compressed into incidents, and the environments per organization.

Key incident highlights:



Incident volume

This section reviews the annual incident volume, the annual incident volume by industry, and the daily incident volume for the organizations included in this report.


Annual incident volume

BigPanda generated nearly 132 million incidents in 2024, or **over 131 million incidents per year** after filtering out the five event outliers. The median was 177,949 incidents per year per organization.

- Over three-quarters (76%) of organizations experienced between 25,000 and 1 million incidents per year (low-medium to high annual incident volume), which indicates that most were actively using the platform to manage meaningful incident flow.
- Nearly two-thirds (63%) experienced at least 100,000 incidents per year.
- Almost half (45%) experienced 250,000 or more incidents per year (high to extremely high), including a small but meaningful group (18%) that experienced 1 million or more (very high to extremely high annual incident volume).
- Over a quarter (26%) experienced 250,000 to 1 million incidents per year (high annual incident volume), the largest group.
- Just 6% experienced fewer than 25,000 incidents per year (low to very low annual incident volume), likely onboarding organizations.

45%

of organizations experienced 250K+ incidents per year



Annual incident volume (n=125)




Annual incident volume by industry

Looking at the median annual incident volume per organization by industry, the data showed that:

- Transportation organizations experienced the **most** annual incidents (1,199,234), followed by hospitality (711,851), telecommunications (629,295), energy/utilities (572,142), and financial services (400,824).
- Excluding the other industry category, managed service provider organizations experienced the **fewest** annual incidents (79,011), followed by media/entertainment (98,431), insurance (119,607), retail (123,619), and manufacturing (152,107).

Comparing the median to the mean (average) shows that:

- Financial services and insurance organizations had the biggest drops from mean to median, suggesting their means were very inflated by outliers. These organizations were likely highly variable, from niche players to massive global banks and insurers.
- The median incidents for transportation, hospitality, and energy/utilities organizations were notably higher than the mean. This suggests these sectors had more consistent usage across organizations and operated at scale (not just a few big players).
- The median incidents for telecommunications and managed service provider organizations were much lower than the mean, hinting at a heavy skew from a few power users.
- The median healthcare, technology, and media/entertainment organization incidents were relatively close, indicating uniform adoption patterns and more stability.



Median and average annual incident volume per organization by industry (n=125)

[See the actionability rate.](#)




Daily incident volume

The median daily incident volume was 545 incidents per day. After excluding the 3% of organizations with more than 25,000 incidents per day (outliers), the median barely shifts (from 545 to 494), reinforcing that most organizations remain in the low-to-medium range.

- Over half (51%) of organizations generated 100–999 daily incidents (low-medium, medium, and medium-high daily incident volume). In other words, most organizations experienced fewer than 1,000 incidents per day.
- Over a third (37%) experienced 1,000 or more daily incidents (high-to-extremely-high daily incident volume).
- The remaining 12% experienced fewer than 100 incidents per day (very-low-to-low daily incident volume), likely including onboarding organizations.

51%

of organizations experienced 500+ incidents per day



Daily incident volume (n=125)



Alert correlation rate

Alert correlation, also known as event correlation, uses correlation patterns to consolidate alerts from external observability and monitoring tools, significantly reducing alert noise and giving teams actionable insights to resolve incidents before they become outages. The **alert-to-incident correlation rate** is the percentage of alerts correlated into incidents.


A healthy alert-to-incident correlation rate range is 40–75%. Anything under 40% usually leaves something on the table; anything over 75% usually means too much correlation. It's a delicate balance.

The median alert-to-incident correlation rate was 67%.

- About half (49%) had a **healthy** alert-to-incident correlation rate (40–75%).
- Over a third (38%) had a **high** alert-to-incident correlation rate (75% or more).
- Only 14% had a **low** alert-to-incident correlation rate (less than 40%).

49%

of organizations had a healthy alert correlation rate (40–75%)



Alert-to-incident compression rate compared to median event volume (n=125)

The data show that alert volume alone does not determine correlation efficiency. Still, there's a mild tendency for organizations with a high volume of alerts to achieve better correlation, likely due to operational scale.



Incident compression rate


The **incident compression rate**, sometimes called just **compression** or **compression rate**, is the percentage of events compressed into incidents (**event-to-incident compression rate**).

The event-to-incident compression rate ranged from 70.9% to 99.9%, and the median was 97.3%.

- Most (72%) organizations achieved a strong event-to-incident compression rate of 95% or more—including 30% with a very high rate of 98.5% or more—signaling that **event deduplication and correlation work well for the majority**.
- One in five (20%) had high event volumes and respectable (medium) compression rate of 85–94.9%, indicating they likely have strong rules but might benefit from fine-tuning.
- Only 8% had a very-low-to-low event-to-incident compression rate of below 85%. However, the 70–79.9% range, while tiny at just 4%, had a surprisingly high median event volume, suggesting missed correlation opportunities, noisy environments, or onboarding organizations.

72%

of organizations achieved a strong incident compression rate (95+%)



Event-to-incident compression rate compared to median event volume (n=125)

The median event volume did not always correlate with the compression rate range. For example, organizations in the 97.5–98.4% range compressed more efficiently than those in the 95–97.4% range, yet had a slightly lower event volume. This implies that **compression quality is not solely a function of volume; configuration and filtering are likely key drivers**.




Environments

In BigPanda, an **environment** is a configurable view of the IT infrastructure that helps teams focus on specific incident-related information.

Environments filter incidents on properties, such as source and priority, and group them for improved visibility, automation, and action. They are customizable and make it easy for teams to focus on incidents relevant to their role and responsibilities, including filtering the incident feed, creating live dashboards, setting up sharing rules, and simplifying incident searches.

Excluding the five outliers, the median number of environments per organization was 58.

- About three-quarters (74%) of organizations had 26–200 environments, including 34% with 51–100, the largest segment in the distribution. These organizations likely have multiple teams, applications, or regions that require centralized monitoring and incident response capabilities and are candidates for scaling observability and automation.
- Nearly one in five (19%) had 25 or fewer environments, including 4% with 10 or fewer. This may indicate they had simpler environments with fewer assets to monitor, were likely to rely on basic alerting or minimal automation, and had room to grow in segmentation, tagging, and response maturity.
- Just 6% had more than 200 environments. These organizations likely require deep observability, advanced correlation, enrichment, and deduplication, as well as multi-environment analytics and reporting.



Number of environments per organization (n=125)

58%

of organizations had 50+ environments



Actioned incidents

Actioned incidents represent outages and system issues that a team member acted on. An action could be a comment, an assignment to a user, a manual share, or an automated share. They are a key metric in determining the efficacy of BigPanda configuration and workflows.

This section reviews the incident volume, actionability rate (incident-to-actioned-incident rate), and noise reduction rate (event-to-actioned-incident rate).

Key actioned incident highlights:



“For us, an alert is not actionable unless it comes into BigPanda, is enriched, and is potentially correlated with the other alerts in the system.”

—Head of Software Engineering, Telecommunications Enterprise



Actioned incident volume

This section reviews the annual actioned incident volume, the annual actioned incident volume by industry, the monthly actioned incident volume, and the daily actioned incident volume.


Annual actioned incident volume

BigPanda generated nearly 20 million actioned incidents in 2024 for the organizations included in this report. After filtering out the five event outliers, there were **19.23 million actioned incidents per year**. The median was 34,232 actioned incidents per year per organization.

- Two-thirds (66%) of organizations actioned 10,000–249,999 incidents per year (**medium-to-high volume**), suggesting widespread usage of the BigPanda platform at an operational scale. A third (33%) actioned 10,000–49,999 incidents per year (**medium-to-medium-high volume**), and the other third (33%) actioned 50,000–249,999 incidents per year (**high volume**).
- Nearly a quarter (23%) actioned fewer than 10,000 incidents per year (**very-low-to-low-medium volume**), including 6% with fewer than 1,000, which were likely onboarding organizations.
- Just 11% actioned 250,000 or more incidents per year (**very-to-extremely-high volume**), including 2% with 1 million or more, representing very large, global enterprises operating with complex business logic.

33%

of organizations actioned
10K–49.9K annual incidents



Annual actioned incident volume (n=125)




Annual actioned incident volume by industry

Looking at the median annual actioned incident volume per organization by industry, the data showed that:

- Transportation organizations actioned the **most** incidents per year (127,702), followed by hospitality (109,993), financial services (86,034), energy/utilities (54,460), and technology (46,856).
- Excluding the other industry category, professional services organizations actioned the **fewest** incidents per year (24,742), followed by healthcare (26,145), media/entertainment (29,368), managed service providers (31,977), and manufacturing (33,083).

Comparing the median to the mean (average) shows that:

- Financial services organizations had a significant gap between the mean (490,157) and median (86,034) actioned incidents per year, indicating a heavily skewed distribution likely driven by a few very large organizations and wide variability in scale within the sector.
- Transportation was the only industry where the median (127,702) exceeded the mean (121,876), implying that most transportation organizations had fairly consistent usage with a relatively balanced distribution and no extreme outliers.
- The technology and energy/utilities industries showed lower medians (46,856 and 54,460, respectively) than means (118,089 and 107,702, respectively), suggesting a few high-volume organizations lifted the average. Still, most organizations in these industries had lower volumes of actioned incidents.
- Hospitality industries had a high median (109,993) relative to their mean (193,082), indicating a more even usage distribution and leaning toward mature implementations across the organizations.



Median and average annual actioned incident volume per organization by industry (n=125)


[See the actionability rate.](#)



Monthly actioned incidents (frequency)

When comparing the actioned incident count per month to the event count per month, the data show that:

- The event volume was high but stable, and actioned incidents followed a similar pattern. Both events and actioned incidents peaked in October, suggesting a correlation. In addition, the lowest number of actioned incidents occurred in May, aligning with one of the lower event months.
- The overall shape of both curves was quite aligned, indicating that the alerting or incident workflow responds proportionally to the scale of events. This suggests that the monitoring systems weren't overloaded or under-triggered in certain months, which is a good sign.
- Even during the high-volume months like July through October, actioned incidents increased steadily but didn't spike uncontrollably. This could point to effective alert thresholds, deduplication, or noise suppression.



Monthly event count compared to monthly actioned incident count (n=125)




Daily actioned incident volume

BigPanda generated nearly 55,000 actioned incidents per day for the organizations included in this report. After filtering out the five event outliers, there were **53,900 actioned incidents per day**. The median was 110 actioned incidents per day per organization.

- Over a third (37%) of organizations actioned 10–99 incidents per day (**medium-to-medium-high volume**).
- More than a quarter (29%) actioned 100–499 incidents per day (**high volume**).
- Nearly a quarter (22%) actioned 500 or more incidents per day (**very high volume**).
- Just 12% actioned fewer than 10 incidents per day (**very-low-to-low-medium volume**), including 2% with less than daily actioned incidents, likely representing onboarding organizations.

22%

of organizations actioned 500+ incidents per day



Daily actioned incident volume (n=125)



Actionability rate

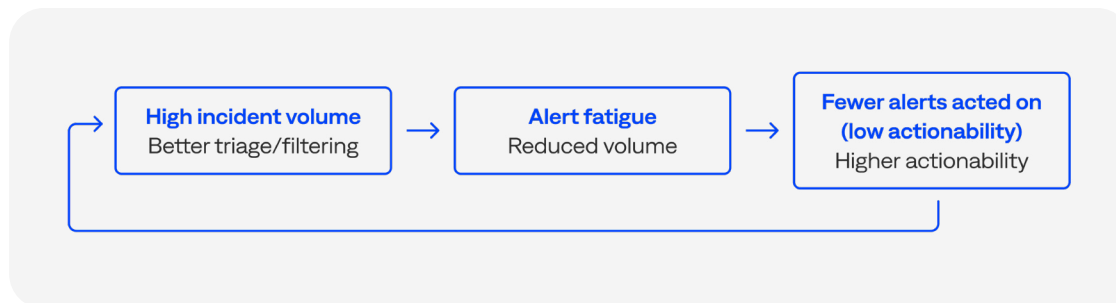
The **actionability rate** is the percentage of incidents that were actioned (**incident-to-actioned-incident rate**).

Both high and low actionability rates can be good or bad.

- **High actionability** could mean that an organization has reduced noise by removing incidents that don't need to be acted on (good), or teams are unnecessarily acting on noisy incidents (bad).
- **Low actionability** can mean that monitoring and observability tools send a lot of noisy, unactionable events to BigPanda (bad), or teams use BigPanda as an excellent filter to prevent unactionable tickets from being created in the first place (good).

BigPanda customers with incident management teams working in ITSM platforms typically have higher actionability rates because they use BigPanda to reduce, correlate, and ticket immediately. However, most organizations only take action on a very small percentage of incidents because their monitoring and observability tools generate a lot of noise. BigPanda helps them focus only on what's important.

With [BigPanda unified analytics](#), teams get the visibility and insight they need to differentiate valuable signals from noise and only take action on what matters, reducing overall ticketing and focusing on high-severity and priority incidents. It also helps them pinpoint which monitoring and observability tools provide valuable signals versus which are noisy, so they can filter and ignore the ones that don't make the cut.




Actionability rate for all organizations

The median actionability rate was 18%.

- Over half (51%) of organizations had a very-low-to-low actionability rate of less than 20%. The lowest actionability ranges had the highest incident volumes, likely due to high alert noise, poor correlation or triage mechanisms, or a lack of automation in incident handling.
- Over a quarter (27%) had a medium-high-to-extremely-high actionability rate of 50% or more, including 19% with 70% or higher and 5% with 90% or higher.
- There was a moderate negative correlation between median incident count and actionability rate; organizations that experienced a higher incident volume often had lower actionability rates, while those with lower incident volumes tended to have higher actionability rates. There's likely a causal feedback loop. That said, correlation is not causation.

51%

of organizations had a <20% actionability rate




Actionability rate (incident-to-actioned-incident) compared to median incident volume (n=125)



Actionability rate by industry

There are similar patterns when looking at actionability rate and incident volume by industry (higher incident volumes correlate with lower actionability):

- Healthcare organizations had the **highest** median actionability rate (68%), followed by retail (60%), suggesting strong incident response workflows, effective alert correlation or triage, and possibly narrower, more focused monitoring scopes.
- Telecommunications organizations had the **lowest** median actionability rate (6%), followed by transportation (11%), and high incident volumes, suggesting they may be overwhelmed by incident load, under-automated, lacking efficient triage, or needing tuning or alert suppression.
- Three industries showed room for optimization, with relatively low incident volumes but average actionability rates: professional services (38%), managed service providers (35%), and media/entertainment (32%).
- Three industries had relatively low incident volumes and low actionability rates: manufacturing (12%), technology (14%), and insurance (17%).



Actionability rate (incident-to-actioned-incident) compared to median incident volume by industry (n=125)




Noise reduction rate

The **noise reduction rate** is the percentage of raw events that become actioned incidents (**event-to-actioned-incident rate** or **end-to-end noise reduction rate**).

The noise reduction rate ranged from 83% to 99.9%, and the median was 99.6%. In other words, they reduced incident-related noise by up to 99.9%, from raw events to actionable incidents—essentially filtering out all but the most critical signals. This supports the earlier finding that most organizations using the BigPanda platform have excellent filtering practices.

- Most (82%) organizations had an exceptionally high noise reduction rate (97% or higher), including 54% with 99.5% or higher and 28% with 99.9%. This points to highly effective correlation, deduplication, and suppression practices among these organizations.
- Only 16% fell into the medium 90–96.9% noise reduction rate range. These organizations likely reduce most noise but still pass a noticeable volume of events through as incidents, indicating opportunities to improve correlation rules or filters and tune alert thresholds or enrichments.
- Just 2% had a noise reduction rate below 90%. These organizations were likely still onboarding.



Noise reduction rate (event-to-actioned-incident) (n=125)

54%

of organizations had a
99.5+% noise reduction rate

“BigPanda enabled us to implement AI that reduces alert noise and gets us to the root cause faster.”

—Divisional CTO, Managed Services Provider



Conclusion

This report reveals a critical truth: Monitoring and observability tools generate vast amounts of data and noise without delivering the actionable context L1 detection teams need to identify, triage, and respond to alerts. As a result, frontline L1 teams struggle to distinguish actionable signals from noise, requiring large in-house and managed service provider (MSP) teams to manage all the data and escalate issues to L2 incident response teams.

In addition, while many organizations say they are centralizing observability, the data shows that most have 20 or more integrations with monitoring, observability, and other tools. It suggests that many are struggling to optimize them to reduce noise for siloed incident response teams that don't use those tools.

BigPanda helps enterprises quickly detect, prioritize, and respond to actionable incidents. The BigPanda team is also dedicated to helping customers get the most value from the platform. They monitor usage to help customers adopt best practices for better event and incident management outcomes.

As enterprises adopt a more autonomous and [agentic ITOps platform](#) like BigPanda, they can automate incident detection and remove blind spots from their observability data. Using AI-driven correlation across not just observability data, but also service desk tickets and external service provider dependencies such as ISP, utility, and cloud provider outages, enables teams to detect early signals of incidents and respond fast.

The next BigPanda report will examine the business value customers receive from the BigPanda platform.

[Sign up](#) to receive an alert when the business value report is published.

Schedule a Demo

“We implemented BigPanda because we needed a single platform to centralize our tools and support both on-premise and cloud. The value to the business has been tremendous. BigPanda sorts through all the noise and generates, in most cases, a single ticket to point to the problem. The amount of noise we have removed from the environment is tenfold that of what we were used to. It frees our teams to focus on critical services and ensure they are always available for our customers.”

—Director of Infrastructure and Operations, Energy/Utilities Enterprise



About this report

This is the first report based on BigPanda platform data to provide analysis and insights on the effectiveness of monitoring tools and observability platforms for event management.

Methodology

This report is based on data gathered from the BigPanda platform in 2024.

Demographic and firmographic information is based on ZoomInfo data from March 2025.

BigPanda anonymized and aggregated the relevant data to give a general overview of the effectiveness of monitoring tools and observability platforms, as well as event management benchmarks. Any detailed information that could help attackers and other malicious parties was deliberately excluded from this report.

All quotes were derived from customer feedback and [BigPanda customer case studies](#).

All dollar amounts in this report are in USD.

All data in this report are based on UTC (Coordinated Universal Time), also known as GMT (Greenwich Mean Time).

Outliers usually skewed the average (mean), so the median was more representative of typical behavior and was used throughout the report.

Citing the report

Suggested citation for this report:

APA Style:

BigPanda (July 2025). *Monitoring and Observability Tool Effectiveness for IT Event Management*. BigPanda. <https://www.bigpanda.io/resource/report/tool-effectiveness-for-it-event-management>

The Chicago Manual of Style:

BigPanda. July 2025. *Monitoring and Observability Tool Effectiveness for IT Event Management*. N.p.: BigPanda. <https://www.bigpanda.io/resource/report/tool-effectiveness-for-it-event-management>.



Demographics and firmographics

This report is based on data from 130 organizations using the BigPanda platform in 2024, including six onboarding organizations.


n=130

Organization size

The BigPanda platform was designed for enterprise organizations, so all organizations included in this report were large enterprises. In fact, 42% were on the [2024 Fortune 1000 list](#), 29% were on the [2024 Fortune 500 list](#), and 12% were on the [2024 Global Fortune 500 list](#).

Annual revenue


All organizations had an annual revenue of at least \$15 million. Most (79%) had an annual revenue of at least \$1 billion, including 51% with at least \$5 billion, 36% with at least \$10 billion, and 12% with at least \$50 billion. **The median annual revenue was \$5.2 billion.**



Annual revenue per organization

Number of employees

The median number of employees was about 10,000. Half (50%) had 10,000 or more employees, including 17% with 50,000 or more and 12% with 100,000 or more. Less than a third (30%) had fewer than 5,000, including 12% with fewer than 1,000.



Number of employees per organization




Regions

Most (90%) of the organizations included in this report had headquarters in the Americas, and 10% had headquarters in Europe, the Middle East, and Africa (EMEA).



Global headquarters locations




United States headquarters locations



Number of locations

However, most (96%) of the organizations included in this report have multiple locations, and many are global. In fact, users from 88 different countries worldwide accessed BigPanda.

The median number of locations per organization was 22. Over three-quarters (76%) had at least 10 locations, including 14% with 50 or more. Less than a quarter (24%) had fewer than 10 locations.




Number of locations per organization



Industries

Organizations from many diverse industries are harnessing the power of the BigPanda platform for event management, including financial services (16%), insurance (12%), technology (10%), media/entertainment (10%), manufacturing (10%), and managed service providers (MSPs; 10%).



Percentage of organizations in each industry



About BigPanda

BigPanda delivers agentic automation for ITOps. We enable enterprises to keep the digital world running by transforming manual, reactive human processes into intelligent, autonomous systems that detect, respond, and prevent IT incidents at machine speed. That’s why the world’s most trusted brands rely on BigPanda to improve operational efficiency and deliver exceptional service reliability to their customers.

Schedule a Demo

