


BigPanda L1 Agent Data Sheet

Scale L1 operations without adding headcount using AI agents that route, suppress, and resolve incidents.



Manual triage across disconnected tools slows L1 decision-making and response times, making ITOps inefficient, costly, and hard to scale. [BigPanda L1 Agent](#) automates high-volume, repetitive incidents end-to-end, reducing the manual work handled by L1 and outsourced operations teams. Enterprises report ~90% routing accuracy with MTTA measured in minutes.

Challenges and solutions

BigPanda L1 Agent reshapes the NOC operating model across three dimensions: capacity, accuracy, and escalations.

	NOC capacity	Accuracy	Escalations
Challenge	Alert and incident volumes grow faster than budgets allow, forcing NOCs to add L1 headcount or GSI contracts that scale linearly with the problem.	L1 routing and first-touch actions vary by operator experience, leading to tickets being sent to the wrong teams and inconsistent triage across shifts and seniority levels.	When L1 teams escalate, L2 and L3 teams receive tickets stripped of critical context, forcing senior engineers to rebuild the timeline before they can begin investigation.
BigPanda solution	L1 Agent handles routine assignment, suppression, and resolution work autonomously, so L1 capacity grows with alert volume instead of with headcount or GSI spend.	L1 Agent reasons over real-time incident data and historical outcomes from the BigPanda IT Knowledge Graph , so every incident gets the same evidence-based first action regardless of who's on shift.	L1 Agent attaches the full investigation trail to every escalation, including correlated signals from your entire IT environment and runbook steps that've already been attempted, so L2 and L3 teams have a complete picture.
Business value or outcome	Scale NOC capacity without scaling headcount or GSI spend.	Faster, more accurate resolution means less downtime and better customer experiences.	Unlock engineering capacity and drive innovation.



Benefits

BigPanda L1 Agent delivers the following benefits:

- ✓ **Lower L1 operating costs:** AI agents absorb routine assignment, suppression, and runbook work, so growing alert volume no longer means growing headcount or GSI contracts.
- ✓ **Faster engineering innovation:** Fewer incidents reach L2/L3 teams, and those that do arrive with full context, freeing senior engineers for innovation and strategic projects.
- ✓ **Fewer service disruptions:** Faster, more accurate triage and resolution mean less downtime and fewer SLA breaches for critical, revenue-affecting incidents.

Features and capabilities

BigPanda L1 Agent includes the following features and capabilities to scale L1 response without adding headcount:

- **Automated assignment:** Routes every incident to the right team using AI reasoning over real-time event data and historical outcomes from the IT Knowledge Graph, with no static rules to maintain.
- **Automated suppression:** Identifies false positives, transient flaps, and self-resolving conditions in real time and suppresses them before they reach the operator queue, within thresholds configured by the team.
- **Runbook investigation:** Searches the IT Knowledge Graph for the runbook that matches each incident and follows its investigation steps, attaching findings to the ticket so operators start with full context.

“BigPanda L1 Agent has delivered a step-change in efficiency for our NOC. By automating routine triage and resolution, we’ve reduced MTTR and operational overhead, and built a more resilient, scalable, and cost-efficient operating model than headcount alone could deliver.”

Dan Bartram

Head of Platform Engineering at
Gamma Communications



See BigPanda L1 Agent in action.

Get a demo