

SECURING A DESTINATION THAT DOES NOT YET EXIST

One of the most ambitious entertainment projects in the world was building its security program from the ground up. The security team had plenty of intelligence feeds. What it lacked was a way to turn that volume into decisions it could defend.

One of the most closely watched infrastructure undertakings, a large-scale entertainment project, is a flagship initiative backed by billions of dollars and expectations to match. Its designers envision millions of visitors each year, drawn by theme parks, motorsport venues, waterparks, and entertainment complexes that do not yet fully exist.

Building this requires solving problems that most organizations don't face at this scale. The logistics are staggering, the timelines are compressed, and the security challenges are, by definition, unlike anything in the existing playbook.

When the security team at the entertainment company sat down to assess their threat intelligence operation, they found a program that had grown organically and outpaced its own infrastructure: feeds from multiple sources, significant data volume, and broad coverage of the threat landscape. All of it was being sorted, evaluated, and acted upon largely by hand.

That approach had served them reasonably well in earlier stages of the program. By the time the company assessed its current posture, the manual model had hit its limits. It was an operating model built for a threat landscape that moved at human speed, and it produced alerts rather than decisions. That landscape is gone.

THE VOLUME PROBLEM

The company had invested seriously in threat intelligence. Security leaders had collected feeds from multiple providers, with the majority drawn from open-source repositories. On the surface, the setup offered wide visibility into the threat landscape.

In practice, the volume of incoming indicators of compromise had become unmanageable. A significant portion of the data was duplicated across feeds, irrelevant to the company's specific environment, or both. Analysts had no automated way to filter the signal from the noise. Every item required human review. Every prioritization decision that could have been made algorithmically was made by a person instead.

The consequences were predictable. False positives multiplied, consuming analyst time that should have been directed at genuine threats. Correlating external intelligence with internal security event logs was a slow, labor-intensive process. And despite the investment in feeds, tools, and personnel, the security team could not point to clear evidence that the program was reducing risk in any measurable way.

"Everything was very manual and done with a lack of process," said a person familiar with the company's situation. "The volume of false positives drained analyst time and distracted the team from the real threats they needed to be aware of and act on."

This is a pattern that security researchers have documented across the industry. According to the 2025 SANS Detection & Response Survey, 73% of organizations cite false positives as their top detection challenge with more than 60% encountering them frequently or very frequently. Compounding this, organizations now receive an average of 2,992 security alerts daily, yet 63% go unaddressed (Vectra AI, 2026). More inputs, without a platform designed to normalize and prioritize them, can leave a security team slower and less effective than before. Analysts feel it directly: an estimated 52% of analyst time is consumed by false positives (Times, Voice of the SOC).

The underlying issue is fidelity. Before anyone, human or AI agent, can act on intelligence, three things have to happen: normalization, so schema differences are not mistaken for signal; deduplication, so one event does not become a thousand alerts; and enrichment, so context arrives attached rather than looked up later. Skip any of the three and the failure is not just more mistakes. It is mistakes made faster, and with a confidence they have not earned.

For a company building critical infrastructure at the scale of a major international entertainment destination, the stakes of that inefficiency were significant.

A HIGH-VALUE TARGET, BY DEFINITION

Large-scale entertainment destinations occupy an unusual position in the threat landscape. They are consumer-facing at enormous scale; they process vast quantities of personal and financial data across ticketing systems, loyalty programs, and retail and hospitality platforms. The digital infrastructure required to run theme parks, racing circuits, and live entertainment venues across a single campus generates an attack surface that grows in proportion to the destination's ambition.

Projects of this scale also carry reputational weight that extends well beyond their own gates. A serious security incident at a destination of this profile would carry consequences far wider than the company itself, affecting the investors, partners, and governments that have staked credibility on its success.

Underground markets listed nearly 14.5 million compromised credit cards in 2024, a 20% year-over-year increase according to Bitsight's 2025 State of the Underground report. Companies that process high volumes of consumer financial data are, by definition, attractive targets. The company understood this, but needed to build a security program capable of responding at the pace the threat environment demanded. That pace keeps shrinking. The fastest quartile of real intrusions

reached data exfiltration in 72 minutes in 2025, down from 285 minutes in 2024 (Unit 42 2026 Global IR Report). A security operation built around human-speed review cannot defend against an adversary working at machine speed.

When the company evaluated managed intelligence solutions, its security leadership defined success in operational terms. Three outcomes mattered.

- The first was clarity. The team needed a platform capable of aggregating all incoming feeds, removing duplication, and surfacing only the indicators genuinely relevant to their specific environment. The goal was to capture the right intelligence, configured to reflect what the company actually cared about. That is what operational intelligence delivers: full-fidelity intelligence, curated and operationalized, so every decision starts with complete context.
- The second was speed. Analyst hours spent on manual collection, summarization, and triage represented a direct cost to the security program. Automating that work would free the team to focus on investigation and response rather than data management. Analysts and AI could then investigate, prioritize, and respond faster from one shared, governed foundation instead of another disconnected tool.
- The third was cost efficiency. The company was carrying the overhead of legacy on-premises infrastructure. Reducing that burden, while simultaneously expanding visibility, was a financial as well as an operational priority. Anomali's architecture decouples cost from data volume, so visibility can grow without cost growing in step.

After a competitive evaluation that included several established vendors in the threat intelligence space, the company selected Anomali. The deployment included the full Anomali Agentic SOC Platform that includes Anomali's Managed Intelligence as a Service (MlaaS), powered by ThreatStream Next-Gen and its Intelligent Unification Layer solution, Anomali Unified

Security Data Lake. MlaaS is the intelligence foundation for everything built above it, and the data lake gives it visibility. The layers compound one another rather than replacing one another.

The core value proposition centered on Anomali's aggregation and enrichment capabilities. All incoming feeds were ingested simultaneously, normalized, deduplicated, and enriched with contextual associations that allowed analysts to quickly identify what required immediate attention and what could be deprioritized. The continuous correlation of indicators of compromise, previously a manual process, became an automated background function.

Anomali's AI capabilities added a further layer of efficiency. Analysts gained the ability to generate structured intelligence summaries and executive-ready reports from raw threat data, reducing the time required to brief leadership and demonstrate program value to senior stakeholders. Agents without intelligence are useless, so Anomali grounds its AI in operational intelligence rather than pointing it at raw data. It is also designed around governance: humans stay in the loop for every consequential action, and before an agent acts, Anomali calculates the blast radius and keeps every step reversible. Speed does not come at the cost of control.

The effect, according to people familiar with the deployment, was a structural shift in how the security team operated. Analysts moved from spending the majority of their time managing data to spending it investigating threats. It was the first step in a larger shift, from an operating model that produces alerts to one that produces governed decisions, made at machine speed and defensible by design.

WINNING A COMPETITIVE EVALUATION

The decisive document was a comparative evaluation matrix that mapped each vendor's capabilities directly to the company stated requirements. The Anomali team prepared that

analysis, and the CISO presented it to executive leadership. The transparency of the approach gave Anomali visibility at the top of the company and helped consolidate support ahead of a final decision. That is the test every CISO ultimately faces: a decision they can defend to management, to the board, and, where applicable, to a regulator or auditor.

Anomali's AI story proved particularly resonant with leadership. The ability to articulate a credible vision for what the security operation could look like in the future, with AI handling the routine and analysts focused on judgment, aligned with how the company's leaders were already thinking about technology investment across the broader project. That vision rests on three pillars.

- Informed: full-fidelity threat intelligence, curated and operationalized, gives every decision complete context.
- Predictive: accumulated operational memory lets the platform anticipate what is coming, not just react to what already happened.
- Proactive: governed autonomy turns anticipation into action, within boundaries the customer sets.

Autonomy is earned the same way, one stage at a time, from agents that recommend, to actions that require approval, to agents that act, governed at every rung and never granted wholesale.

The company also had to work through a shift in how it envisioned its own future. Its security team had built a posture around on-premises infrastructure and manual processes. Moving to a cloud-delivered, AI-augmented model required a genuine change in operating philosophy alongside the technology decision. That transition, and the confidence to commit to it, was as much a part of the engagement as the technical deployment.

THE LESSON FOR ENTERTAINMENT INFRASTRUCTURE LEADERS

The IBM 2025 Cost of a Data Breach Report found that organizations deploying security AI and automation extensively contained breaches 80 days faster and saved nearly \$1.9 million on average compared to those without it. For an organization whose security program must scale alongside a destination still under construction, that speed and cost advantage compounds over time.

The broader lesson applies to any large-scale entertainment or hospitality company building security infrastructure in parallel with the asset it is meant to protect. Raw intelligence volume is only as valuable as the ability to operationalize it. The companies that translate incoming data into faster, more accurate decisions are the ones that reduce risk in a way that senior leadership and financial stakeholders can actually measure. AI pointed at a fragmented stack produces noise faster. AI built on a fidelity foundation produces decisions leadership can trust.

For leaders deciding where to begin, the sequence matters. Threat intelligence is the foundation, not a stage to move past, and Intelligent Unification Layer is the precondition for a real Agentic SOC: unify the environment first, or the agents built on top of it only produce noise. Anomali lands over the existing environment, so nothing has to be replaced to start. As unification, agents, and measurable outcomes are proven, the customer decides how much legacy to retire.

This company chose to build that capability before its destination opened its gates. The goal was a security operation in which every action is one the board, the regulator, and the night shift would each stand behind: made at machine speed, defensible by design. The security posture being established now will determine how resilient that infrastructure is when millions of visitors arrive.

Ready to move from noise to signal in your threat intelligence operation? [Schedule a confidential conversation with Anomali.](#)

[SCHEDULE A MEETING](#)