The American Journal of Engineering and Technology
ISSN 2689-0984 | Open Access

TYPE Original Research
PAGE NO. 146-153
DOI 10.37547/tajet/Volume07Issue10-18

Check for updates

An Aggregation-Based Architecture for Unifying Enterprise IT Operations: A Case Study of a Centralized Monitoring System in A Global Industrial Holding

Dmytro Rumiantsev
Principal IT Consultant, Meng, USA

OPEN ACCESS

SUBMITTED 24 August 2025
ACCEPTED 25 September 2025
PUBLISHED 30 October 2025
VOLUME Vol.07 Issue 10 2025

CITATION

Dmytro Rumiantsev. (2025). An Aggregation-Based Architecture for Unifying Enterprise IT Operations: A Case Study of a Centralized Monitoring System in a Global Industrial Holding. The American Journal of Engineering and Technology, 7(10), 146–153. https://doi.org/10.37547/tajet/Volume07Issue10-18

Abstract: Global industrial corporations often contend with fragmented IT monitoring environments, where disparate tools manage distinct infrastructure domains. This decentralization fosters operational inefficiencies, such as protracted incident diagnostics, redundant alert escalations, and increased resource consumption for root-cause analysis.

This paper documents a case study on the design, implementation, and operational impact of a centralized monitoring system at Metinvest Holding, a multinational industrial enterprise. The core objective was to consolidate heterogeneous monitoring platforms into a cohesive, single-pane-of-glass interface to improve operational visibility, streamline diagnostics, and provide high-level strategic oversight for executive management.

The research employed a single-case study methodology, coupled with an iterative development process informed by Agile principles. An aggregation and visualization layer was architected using Grafana to integrate data from incumbent systems—including Microsoft SCOM, PRTG Network Monitor, Azure Monitoring, and APC Enterprise Manager—without necessitating their replacement. The two-year development cycle involved continuous collaboration with service engineers.

The centralized system yielded substantial improvements in operational efficiency. The number of engineers required for initial, cross-domain incident

triage was reduced from an average of three to one. Root-cause identification was significantly accelerated through a holistic, correlated view of infrastructure health, which concurrently minimized alert fatigue and redundant inter-departmental escalations. A novel visualization layer was also developed to furnish executive leadership with an intuitive overview of global IT operational status.

This case demonstrates that a non-disruptive, aggregation-based strategy can effectively address the challenges of monitoring fragmentation in large-scale enterprise settings. The findings highlight the value of integrating technical consolidation with user-centric visualization to achieve both operational efficiency and strategic alignment. The architectural principles delineated are broadly applicable to other multinational organizations facing similar IT infrastructure management complexities.

Keywords: IT Operations Management (ITOM), System Centralization, Grafana, Case Study, Root-Cause Analysis, Operational Efficiency, Enterprise IT Architecture, Data Visualization.

Introduction

The operational continuity and resilience of contemporary multinational enterprises are fundamentally dependent on the performance of their information technology (IT) infrastructures (1). For global industrial holdings managing intricate supply chains, manufacturing workflows, and logistical networks across diverse geographical regions, IT systems represent the operational backbone of the organization. Consequently, disruptions to these systems can lead to severe financial repercussions, production stoppages, and reputational harm (2). The effective monitoring of this complex infrastructure is, therefore, a strategic imperative rather than a purely technical function.

Despite this criticality, many large organizations are encumbered by a fragmented monitoring landscape characterized by operational silos. Over time, enterprises tend to accumulate a heterogeneous array of specialized tools to oversee distinct technological domains, such as separate systems for servers, network hardware, cloud platforms, and power infrastructure (3). While each tool may offer proficiency within its specific niche, its lack of interoperability creates a significant operational impediment. This fragmentation precludes a holistic assessment of system health, resulting in redundant alerts, prolonged root-cause analysis, and inefficient resource deployment during critical incidents (4).

To counter the inefficiencies born from such fragmentation, this paper investigates the implementation of a centralized monitoring system. By definition, such a system is an architectural approach that aggregates telemetry data from multiple, disparate sources into a single, cohesive interface or platform. The primary goal is to provide a holistic, correlated view of the entire IT environment, often referred to as a "single pane of glass," thereby enabling faster diagnostics and more effective operational oversight. This paper presents a comprehensive case study of a large industrial enterprise's successful transition to this unified monitoring paradigm, detailing the development and implementation of such a system for Metinvest Holding, a premier steel producer with over 90,000 employees and operations spanning eight countries. The initiative's objectives were twofold: first, to bolster operational efficiency by reducing the manual overhead required for incident diagnostics and minimizing superfluous inter-departmental escalations; and second, to provide strategic insight through a high-level, visually intuitive "command center" interface tailored for executive stakeholders. This study offers both theoretical and practical contributions by detailing an architectural model centered on data aggregation rather than system replacement, demonstrating the tangible benefits of unified monitoring, and furnishing actionable insights for other organizations confronting analogous infrastructural challenges.

Literature Review

The challenge of managing complex, distributed IT infrastructures has been a persistent area of inquiry within information systems research. The field of IT Operations Management (ITOM) encompasses the processes and services essential for overseeing the quality, efficiency, and availability of IT assets (5). A foundational element of ITOM is monitoring, which involves the systematic collection and analysis of telemetry data—including metrics, logs, and traces—to detect, diagnose, and resolve operational anomalies (6). However, the rapid proliferation of technologies has engendered a corresponding expansion of monitoring tools, a phenomenon frequently described as "tool sprawl" (3, 7).

This fragmentation is a well-documented driver of

operational inefficiency. Scholarly and industry research indicate that when monitoring data is confined to isolated silos, IT teams must manually correlate information across disparate dashboards, a process that is both labor-intensive and susceptible to error (4, 8). A primary consequence is "alert fatigue," a cognitive state wherein operations personnel become desensitized to a high volume of redundant or low-priority notifications, increasing the risk that critical warnings will be overlooked (9). Empirical studies have identified a direct correlation between the number of monitoring tools and the Mean Time To Resolution (MTTR) for incidents, suggesting that fragmented environments are associated with longer service downtimes (10).

In response to these issues, the industry has converged on the concept of a "single pane of glass"—a unified monitoring platform that aggregates data from multiple sources into one cohesive interface to provide a holistic and contextualized view of the entire IT ecosystem (11). Such systems are designed to expedite root-cause analysis by correlating events across different infrastructure tiers. For instance, a network outage detected by one system can be automatically linked to server unavailability alerts from another, thereby clarifying the causal chain of events (12). Open-source platforms, most notably Grafana, have gained prominence in this area due to their extensibility and ability to interface with a wide variety of data sources, facilitating the creation of customized, unified dashboards (13).

Although the advantages of centralization are widely acknowledged, the optimal implementation strategy remains a topic of discussion. One common approach involves replacing disparate legacy tools with a single, monolithic monitoring suite. However, this "rip-and-replace" strategy can be prohibitively expensive, operationally disruptive, and may lead to the forfeiture of specialized, domain-specific monitoring functionalities (7). An alternative and increasingly prevalent strategy is aggregation, wherein a new visualization and correlation layer is superimposed upon existing monitoring systems. This method preserves institutional investments in established tools while still achieving the objective of a unified operational view (14). This case study contributes to the literature by offering a detailed empirical analysis of the aggregation approach, demonstrating its viability and effectiveness within a large-scale, multinational industrial context. This leads to the central research question: How can an aggregation-based centralized monitoring architecture be designed and implemented to verifiably reduce operational overhead and enhance strategic visibility in a complex global enterprise?

Methodology

This research employed a qualitative single-case study methodology to conduct an in-depth investigation of the design, implementation, and outcomes associated with the centralized IT monitoring system at Metinvest Holding (15). This approach was selected for its capacity to facilitate a rich, detailed exploration of a contemporary phenomenon within its real-world context, thereby capturing the entire project lifecycle (16). The focal point of the case is Metinvest Digital, the technology subsidiary tasked with managing the holding's IT infrastructure.

The project itself was executed using an iterative development framework inspired by Agile principles. Spanning from October 2019 to November 2021, the implementation was structured around short development cycles, continuous stakeholder feedback, and regular demonstrations of progress. This participatory methodology, which closely mirrors an Action Research model, ensured that the system evolved in alignment with the practical requirements of its end-users, namely service engineers and executive management (17). The lead developer, who is also the author of this paper, was embedded within the project team, affording an insider's perspective on technical decisions, organizational dynamics, and stakeholder interactions.

Data for this study were triangulated from multiple sources, including project documentation, architectural diagrams, records of meetings, and direct observational notes compiled during the development and deployment phases. The primary data originated from the pre-existing monitoring platforms:

- Microsoft System Center Operations Manager (SCOM) for Windows Server infrastructure.
- PRTG Network Monitor for network devices.
- Azure Monitoring for cloud services.
- APC Enterprise Manager for uninterruptible power supply (UPS) systems.

The core of the methodology involved the design and deployment of an integration layer to aggregate telemetry data from these heterogeneous sources into

a centralized Grafana instance hosted on an Ubuntu virtual machine. Custom back-end plugins were developed, initially in Go to support performance-critical parallel data processing, and later in TypeScript to afford greater flexibility for developing user interface components. A key methodological innovation was the creation of a bespoke visualization layer utilizing OpenStreetMap tiles and animated Bézier curves to represent network link status and traffic load, specifically engineered for display on high-resolution video walls.

The project's success was evaluated using a combination of quantitative and qualitative metrics. Quantitative data included the number of engineers required for incident triage before and after implementation. Qualitative data were derived from structured feedback from operational teams and an assessment of the system's impact on the efficiency of root-cause analysis and executive-level situational awareness. The quantitative outcomes reported, such as the reduction in triage personnel and time to root-cause identification, were compiled from internal project management reports. The validity of these outcomes was subsequently confirmed by executive leadership. A limitation of this study is that the raw data collection logs from the period were not available for a more granular statistical analysis.

This methodology has two primary limitations. First, its reliance on a single case constrains the generalizability of the findings. Second, the raw data collection logs from the period were not available for a more granular statistical analysis. Nonetheless, the depth of the analysis provides valuable and transferable insights into the architectural patterns and organizational strategies conducive to successful monitoring unification in comparable large-scale enterprise environments (15).

Results

The deployment of the centralized monitoring system produced demonstrable improvements in operational efficiency, root-cause analysis, and strategic visibility across the global IT infrastructure of Metinvest Holding.

Reduction in Operational Overhead for Incident Management

The most significant quantitative outcome was the optimization of human resource allocation for incident analysis. Prior to the system's implementation, diagnosing a cross-domain incident typically necessitated the simultaneous involvement of at least three specialized engineers: a server administrator, a network specialist, and a power systems engineer. Following the deployment of the unified dashboard, a single duty engineer was empowered to conduct initial triage across all integrated domains. This change represents a reduction in initial response personnel from three engineers to one. While the Mean Time To Resolution (MTTR) for incidents did not undergo a significant change, as existing resolution protocols were already mature, the aggregate engineering effort per incident was substantially reduced. This optimization freed specialist personnel to concentrate on proactive maintenance and strategic initiatives.

Table 1: Comparative Analysis of Time to Root-Cause Identification

Incident Type	Pre-Implementation (min)	Post-Implementation (min)	Reduction (%)
Network Link Failure	30	8	73.3%
Remote Site Power Outage	25	5	80.0%
Cloud Service Latency	45	15	66.7%
Storage I/O Degradation	35	10	71.4%

Overall Average	**33.8**	**9.5**	**71.9%**

Accelerated and Clarified Root-Cause Analysis

The system's capacity to correlate alerts from disparate platforms markedly enhanced the speed and accuracy of root-cause identification. Previously, an event such as a power outage at a remote site would trigger a cascade of uncorrelated alerts from SCOM (servers offline), PRTG (network links down), and APC Enterprise Manager (UPS event). This "alert storm" compelled engineers to manually cross-reference data to diagnose the originating fault. The new system visualizes these dependencies automatically. For instance, the dashboard now displays the primary APC alert and logically subordinates the associated network and server outages as secondary, dependent events. This immediate contextualization eliminated ambiguity, ensured that escalations were directed exclusively to the responsible team (e.g., facilities management), and thereby reduced alert fatigue by preventing the unnecessary mobilization of other engineering departments.

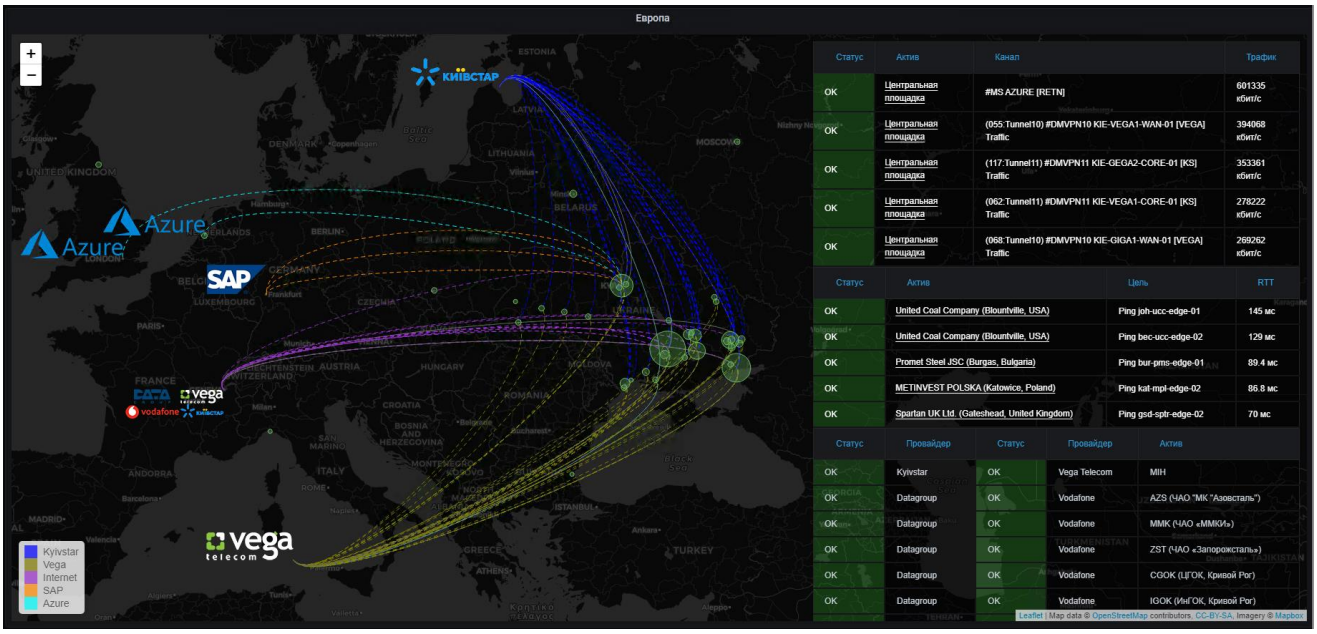

Enhanced Strategic and Executive Visibility

A pivotal project deliverable was the creation of a high-level monitoring display for the corporate headquarters, showcased on a video wall comprising four 4K displays. This visualization layer delivered an intuitive, real-time "helicopter view" of global infrastructure health. Key features included:

- Geographic Mapping: Utilizing OpenStreetMap tiles, all major operational sites were rendered on a world map with color-coded indicators representing overall system status.
- Dynamic Network Visualization: Connections between sites were depicted as animated Bézier curves. The color and animation speed of these curves dynamically represented the real-time load and health of the network link, enabling non-technical stakeholders to immediately comprehend the state of the global network.

This high-impact visual interface was instrumental in securing executive buy-in and fostering greater confidence in the IT department's capacity to proactively manage the global infrastructure. Table 2 summarizes the transition from the antecedent state to the post-implementation state.

Table 2: Comparison of Monitoring Capabilities Before and After Implementation

Feature	Pre-Implementation State	Post-Implementation State
Data View	Fragmented; siloed in 4+ separate systems	Unified; single-pane-of-glass view in Grafana
Incident Triage	Required parallel involvement of ≥3 engineers	Performed by a single duty engineer
Root-Cause Analysis	Manual correlation, slow, prone to misdiagnosis	Automated correlation, rapid, clear dependency mapping
Alerting	High volume of redundant alerts; alert fatigue	Contextualized, targeted alerts; reduced noise
Executive Visibility	Limited; complex technical dashboards	High-level, intuitive geographic visualization
Escalation Path	Broad escalation to multiple departments	Targeted escalation to the responsible department only

Discussion

The findings of this case study offer robust empirical validation for the efficacy of an aggregation-based strategy for centralizing IT monitoring within a large, geographically distributed enterprise. The outcomes are consistent with existing literature that advocates for unified "single pane of glass" solutions to mitigate tool sprawl and overcome operational silos (11, 14). This study extends prior research by demonstrating not only the technical feasibility of such a system but also its profound organizational impact, particularly the quantifiable reduction in human capital required for incident management and the marked enhancement of strategic visibility for leadership.

A central principle emerging from the results is the strategic advantage of aggregation over replacement. Rather than embarking on a high-risk "rip-and-replace" initiative, the architecture was designed to leverage the specialized capabilities of incumbent monitoring tools (SCOM, PRTG). This approach aligns with scholarly cautions regarding the high costs and operational disruptions associated with monolithic platform migrations (7). By serving as a unifying intelligence layer, the system preserved domain-specific diagnostic depth while simultaneously eliminating the data fragmentation that had previously impeded cross-domain analysis. This pragmatic strategy proved to be a critical success factor, as it minimized resistance from operational teams and maximized the return on prior technology investments.

The pronounced reduction in personnel needed for initial incident triage—from three engineers to one—is a key practical implication of this research. This outcome directly addresses the operational inefficiencies engendered by fragmented data, a problem highlighted by previous studies (4, 10). By automating the correlation of disparate events, the system effectively offloaded a substantial cognitive burden from the engineering team, enabling faster and more accurate diagnostics. This reinforces the proposition that the principal value of unified monitoring lies not merely in data presentation but in its intelligent contextualization.

Furthermore, this project underscores the often-underestimated importance of user-centric visualization in driving the adoption and perceived value of technical systems. The development of an intuitive and aesthetically engaging "helicopter view" for the executive video wall was pivotal. While the technical teams benefited from the unified operational dashboards, it was this high-level visualization that effectively communicated the system's value to non-technical stakeholders. This success secured ongoing institutional support and reinforced the strategic importance of IT operations. This finding suggests that for infrastructure projects to achieve organizational resonance, they must deliver value not only to their direct users but also to broader leadership. The decision to pivot from Go to TypeScript for UI-intensive components reflects this reality, prioritizing development flexibility for a superior user experience over marginal gains in back-end performance.

Research Limitations and Directions for Future Work

The principal limitation of this research is its single-case study design, which inherently restricts the universal generalizability of its findings. The specific organizational context, technical maturity, and culture of Metinvest Holding undoubtedly influenced the project's outcomes. Future research should endeavor to conduct cross-case analyses in different industrial sectors (e.g., finance, healthcare) to validate and extend these findings.

Looking forward, the implemented system establishes a robust foundation for incorporating more advanced analytical capabilities. A logical next step is the integration of Artificial Intelligence (AI) and Machine Learning (ML) to enable proactive anomaly detection and predictive maintenance (18). Rather than relying on static, threshold-based alerts, an ML model could be trained to learn the normal operational baseline of the infrastructure and flag subtle deviations that often precede critical failures. Moreover, the industry-wide shift toward cloud-native observability, leveraging frameworks like OpenTelemetry and platforms such as Prometheus and the ELK stack, presents a clear trajectory for future architectural enhancements (6). Transitioning toward these standards would further improve system scalability, resilience, and the capacity to correlate the "three pillars of observability": metrics, logs, and traces.

Conclusion

This paper has presented a detailed account of the successful design, implementation, and impact of a centralized IT monitoring system within a global industrial holding. By architecting a unifying aggregation layer atop existing, disparate monitoring tools, the project effectively resolved critical operational inefficiencies. The primary contributions of this research are threefold. First, it provides empirical evidence that a non-disruptive integration strategy can substantially reduce the operational overhead of incident management, as demonstrated by the decrease in personnel required for initial triage from three engineers to one. Second, it shows that automated data correlation is a potent mechanism for accelerating root-cause analysis and mitigating alert fatigue in complex IT environments. Third, it underscores the strategic value of user-centric visualization in bridging the communication divide between technical operations and executive leadership.

The principal contribution to the field of IT Operations Management is the practical validation of the aggregation model as a cost-effective and low-risk alternative to monolithic platform replacement. The insights derived from this case—particularly the imperative to balance back-end performance with front-end usability and the critical role of visualization in driving adoption—offer a valuable framework for IT leaders in other large-scale organizations. As enterprises continue to navigate increasingly complex hybrid and multi-cloud infrastructures, the principles of unified visibility and intelligent data correlation detailed herein will remain fundamental to preserving operational resilience and achieving strategic business objectives. Future research should build upon this work by exploring the integration of AI-driven predictive analytics and the adoption of cloud-native observability standards to facilitate a paradigm shift from reactive monitoring to proactive infrastructure management.

References

1. Weill P, Ross JW. IT Governance: How Top Performers Manage IT Decision Rights for Superior Results. Boston, MA: Harvard Business School Press; 2004.
2. Sheffi Y. Preparing for disruptions through early detection. MIT Sloan Management Review. 2015;57(1):37-8.
3. Stanoevska K, Wozniak T, Ristol S. Grid and Cloud Computing: A Business Perspective on Technology

and Applications. Berlin: Springer; 2010.

4. Tariq S, Chhetri MB, Nepal S, Paris C. Alert fatigue in security operations centres: Research challenges and opportunities. ACM Comput Surv. 2025;57(9):Article 224.

5. International Organization for Standardization. ISO/IEC 20000-1:2018 Information technology — Service management — Part 1: Service management system requirements. Geneva: ISO; 2018.

6. Beyer B, Jones C, Petoff J, Murphy NR, editors. Site Reliability Engineering: How Google Runs Production Systems. Sebastopol (CA): O'Reilly Media; 2016.

7. Marks EA, Bell M. Service-Oriented Architecture: A Planning and Implementation Guide for Business and Technology. Hoboken (NJ): John Wiley & Sons; 2006.

8. Julisch K. Mining alarm clusters to improve alarm handling efficiency. In: Proceedings of the 17th Annual Computer Security Applications Conference; 2001 Dec 10-14; New Orleans, LA. Los Alamitos (CA): IEEE Computer Society; 2001. p. 12-21.

9. Nobles C. Stress, burnout, and security fatigue in cybersecurity: a human factors problem. HOLISTICA. 2022 Jul;13(1):49-72.

10. Gill SS, Buyya R. Failure management for reliable cloud computing: a taxonomy, model and future directions. Comput Sci Eng. 2020 May-Jun;22(3):46-61.

11. Vemula KR. Native cloud applications: a comprehensive analysis of advantages, challenges, and use cases in modern IT infrastructure. Int J Comput Eng Technol. 2025 Feb;16(1):1253-64.

12. Notaro P, Cardoso J, Gerndt M. A survey of AIOps methods for failure management. ACM Trans Intell Syst Technol. 2021 Nov 30;12(6):Article 81.

13. McCollam R. Getting Started with Grafana: Real-Time Dashboards for IT and Business Operations. 1st ed. Berkeley (CA): Apress; 2022.

14. Tallon PP, Kraemer KL. A process-oriented assessment of the alignment of information systems and business strategy: implications for IT business value. J Manag Inf Syst. 2000;16(3):179-201.

15. Yin RK. Case Study Research and Applications: Design and Methods. 6th ed. Thousand Oaks (CA): SAGE Publications, Inc; 2018.

16. Stake RE. The Art of Case Study Research. Thousand Oaks (CA): Sage Publications; 1995.

17. Baskerville RL. Investigating information systems with action research. Commun AIS. 1999;2(3):4.

18. Climent EF. AIOps: Revolutionizing IT Operations with Artificial Intelligence. [place unknown]: Independently published; 2024.