Sarcouncil Journal of Engineering and Computer Sciences

ISSN(Online): 2945-3585

Volume- 04| Issue- 11| 2025

Research Article

Received: 25-09-2025| **Accepted:** 20-10-2025 | **Published:** 02-11-2025

Mapping Production Understanding in Enterprise Spaces: A Framework for End-to-End System Interaction Analysis

Harris Peter Baskaran

Google LLC, USA

Abstract: Production understanding platforms address the critical need for comprehensive visibility in modern enterprise systems by mapping component interactions and dependencies across complex technology landscapes. These platforms represent an evolution from traditional infrastructure monitoring to interaction-focused methodologies that capture real operational behaviors rather than idealized architectural designs. By combining automated discovery with contextual enrichment, production understanding enables diverse stakeholders to gain insights tailored to their specific needs, from accelerating incident resolution for SRE teams to enhancing compliance oversight for leadership. The article traces the theoretical foundations of production understanding, outlines methodologies for implementation, explores stakeholder-specific applications, and identifies future research directions, including machine learning integration and cross-organization standardization opportunities. As enterprises continue adopting cloud-native architectures and SaaS solutions, production understanding platforms offer transformative capabilities for maintaining operational excellence while supporting strategic business objectives.

Keywords: Enterprise system visibility, interaction mapping, distributed system resilience, cross-functional collaboration, operational technology integration.

INTRODUCTION

The Need For Enterprise Topology And Correlation Platform

The technology infrastructure of a modern enterprise is a complex network of interconnected systems, services, and infrastructure components that span multiple platforms and operational domains. As organizations increasingly adopt cloud services, SaaS solutions, and microservice architectures, the complexity of these ecosystems has grown dramatically. This complexity is not only a technical challenge but a fundamental business concern, as research has demonstrated that high-performing technology organizations leverage comprehensive system understanding to deliver software with both speed and stability. Studies show that these organizations deploy more frequently, recover from incidents faster, and experience fewer failures while maintaining security and compliance, underscoring the importance of thorough production knowledge.

Traditional IT Asset Management [ITAM] and service catalog systems provide valuable inventory capabilities for tracking individual components within the enterprise landscape. However, these approaches typically fall short of capturing the dynamic interactions between systems—the calls, data flows, dependencies, and operational patterns that define how business processes actually execute across the enterprise. While ITAM solutions excel at documenting hardware, software, and digital assets throughout their lifecycle, they rarely address how these components interact to deliver business value in production settings. An effective ITAM strategy focuses on visibility, compliance, and cost optimization of individual assets but does not inherently capture the complex web of interactions that define production environments.

A particularly challenging gap exists between formal system architecture documentation and the reality of operational interactions. Architecture diagrams and documentation often represent idealized or intended system behaviors rather than actual production realities. As systems evolve through maintenance, feature additions, and integration with new services, the actual operational characteristics frequently diverge from documented designs. This divergence creates blind spots in understanding how business processes truly execute across enterprise systems day-to-day, inhibiting the ability to identify bottlenecks, dependencies, and potential failure points across the technology landscape.

"Production understanding" emerges as a comprehensive approach to address these limitations. It represents a holistic methodology for mapping not just system components but also their interactions, dependencies, operational characteristics, and performance patterns as they exist in production environments. Unlike traditional documentation approaches, production understanding focuses on building a dynamic, data-driven representation of the enterprise system

***Corresponding Author:** Harris Peter Baskaran

DOI- https://doi.org/10.5281/zenodo.17507951

landscape that accurately reflects real-world operational behaviors. This approach combines system mapping with continuous data collection about system interactions, enriched with contextual information related to compliance requirements, technical characteristics, incident history, and resource utilization patterns. Research indicates that organizations implementing such comprehensive approaches experience improved resilience and can better respond to changing market conditions through their enhanced visibility and control over complex systems.

THEORETICAL FOUNDATION

From Infrastructure Monitoring to Interaction Mapping

The discipline of Site Reliability Engineering [SRE] has undergone significant evolution since its inception, transitioning from basic infrastructure monitoring to sophisticated practices focused on system reliability, performance, and user experience. This evolution has been characterized by the development of key principles, including service level objectives [SLOs], error budgets, monitoring distributed systems, automation, and the elimination of toil through systematic approaches to reliability. The fundamental philosophy of treating operations as a software problem has transformed how organizations approach system reliability, with dedicated engineering teams applying software engineering principles to operations tasks. These teams focus on building self-healing systems, implementing automation, and developing infrastructure as code, creating a paradigm shift from reactive maintenance to proactive reliability engineering. The strategic use of time allocation between operational tasks and development work ensures continuous improvement while maintaining existing systems, establishing a sustainable approach to managing increasingly complex environments.

A fundamental shift has occurred in monitoring approaches, moving from isolated system observation to comprehensive interaction analysis. This transition represents an evolution from component-level monitoring, which focused on infrastructure metrics like CPU utilization and memory consumption, to a more sophisticated understanding of how services interact across distributed systems. Modern SRE practices emphasize observability through the collection and analysis of logs, metrics, and distributed traces that provide insights into request paths across multiple services. This shift acknowledges that failures in modern systems frequently occur at integration points rather than within individual components, necessitating monitoring tools that can track transactions as they flow through various microservices, APIs, and third-party dependencies. The implementation of service mesh technologies and standardized observability frameworks has further enhanced this capability, allowing organizations to gain visibility into cross-service communication patterns without modifying application code.

Modern cloud providers, as exemplified by Google Cloud's tracing capabilities, offer robust application tracing functionalities. These platforms enable the tracing of existing applications and the aggregation of interaction data without necessitating modifications to the underlying application code. For instrumenting code, vendor-neutral open-source solutions such as OpenTelemetry are widely recommended to achieve these objectives effectively.

Current approaches to enterprise system documentation face significant limitations in capturing operational realities. Traditional documentation often fails to identify architecturally significant requirements that have the most critical impact on system structure and behavior. These requirements, which include quality attributes like performance, security, and maintainability, frequently emerge during operation rather than initial design, creating a disconnect between documented architecture and production reality. Research has shown(Lianping Chen, *et al.*, 2013) that organizations struggle to systematically identify and document these critical requirements, leading to systems that meet functional specifications but fail to address operational needs. This limitation becomes increasingly problematic as systems grow more complex, with implicit dependencies and emergent behaviors that aren't captured in static documentation. The absence of comprehensive documentation about cross-system interactions creates significant challenges for troubleshooting, capacity planning, and system evolution.

A comprehensive conceptual framework for mapping system interactions in complex enterprise environments must address multiple dimensions of system behavior. Such a framework begins with establishing consistent approaches for identifying architecturally significant requirements and interactions across system boundaries. Research

Publisher: SARC Publisher

indicates that successful frameworks incorporate multiple stakeholder perspectives, using techniques such as interview-based elicitation, scenario analysis, and production observation to identify critical system interactions. These approaches recognize that architecturally significant requirements often cross organizational boundaries and require input from diverse stakeholders, including operations teams, developers, business owners, and security specialists. By using dynamic and programmatic approches to documenting cross-system dependencies, communication patterns, and quality attribute requirements, organizations can develop a more comprehensive understanding of their production environment. This understanding serves as a foundation for effective incident management, capacity planning, and architectural evolution, enabling organizations to maintain system reliability while supporting business objectives.

Table 1: The Maturity Curve: From Infrastructure Monitoring to Production Understanding.

Approach Stage	Complexity Level	Business Value Impact
Infrastructure Monitoring	Low	Limited
Component-Level Tracking	Medium-Low	Moderate
Service Dependency Mapping	Medium	Significant
Cross-System Interaction Analysis	Medium-High	High
Comprehensive Production Understanding	High	Transformative

METHODOLOGY

Building a Production Understanding Platform

The foundation of a comprehensive production understanding platform begins with robust system identification and cataloging methodologies that extend beyond traditional asset management approaches. Effective system identification requires a multi-faceted discovery process that leverages both automated and manual techniques to build a complete picture of the enterprise ecosystem. Distributed systems present unique challenges for identification, as they often span multiple data centers, cloud environments, and technology stacks. Research demonstrates that successful identification methodologies incorporate structured approaches to categorizing systems based on their architectural characteristics, data consistency requirements, and operational parameters. These methodologies must also address the complexity introduced by globally distributed deployments, where latency and network partitions create additional challenges for system mapping. By understanding the theoretical underpinnings of distributed system design, organizations can develop more effective cataloging approaches that accurately represent the operational characteristics and constraints of modern enterprise environments.

Interaction tracking represents a critical advancement beyond traditional monitoring, focusing on mapping calls between systems and services as they occur in production environments. Aggregating these interactions provide a starting point for service mapping. Service mapping approaches address this need by providing visual representations of relationships between infrastructure elements, applications, and business services. These mappings reveal how business services depend on underlying application components and infrastructure elements, creating a comprehensive view of service delivery chains. Effective interaction tracking employs discovery mechanisms that collect data from multiple sources, including network traffic analysis, application performance monitoring, and configuration management databases. These discovery processes must address significant challenges, including the dynamic nature of modern environments, technological diversity across the enterprise, and the complexity of multi-tier applications. Research indicates that organizations implementing comprehensive service mapping achieve significant benefits, including faster incident resolution, improved change impact analysis, and more effective capacity planning. The most successful approaches combine automated discovery with manual verification to ensure accuracy, particularly for complex middleware interactions and legacy system dependencies that may not be easily detected through automated means.

Data collection mechanisms for system behavior analysis must balance comprehensive coverage with minimal performance impact to be effective in production environments. Modern distributed database systems provide valuable insights into effective telemetry collection approaches, particularly regarding the collection of temporal data across distributed environments. Research into global-scale database implementations (Eric

Publisher: SARC Publisher

Brewer, 2017) demonstrates the importance of instrumentation that captures timing relationships between distributed components, enabling the analysis of causal relationships in system behavior. Effective data collection mechanisms must address the challenges of clock synchronization across distributed environments, using techniques such as bounded uncertainty intervals to reason about event ordering. These approaches recognize that in distributed systems, traditional assumptions about event timing may not hold, requiring more sophisticated collection and analysis techniques. Organizations implementing comprehensive data collection benefit from instrumentation that captures not just component-level metrics but also cross-component interactions and timing relationships, enabling more effective analysis of distributed system behavior under various operational conditions

Annotation strategies for enriching system maps with operational metadata transform raw interaction data into contextually relevant insights. Effective service mapping approaches recognize that technical mapping alone is insufficient—maps must be enriched with business context, compliance information, and operational metadata to provide meaningful insights to stakeholders. This enrichment process connects technical components to business services and processes, enabling organizations to understand the business impact of technical changes and incidents. Research indicates that successful annotation strategies (Enterprise architecture bridging business and IT) implement well-defined processes for metadata maintenance, ensuring that enrichment information remains current as systems evolve. These processes typically include regular validation cycles, clear ownership assignments, and integration with change management workflows. By implementing comprehensive annotation strategies, organizations can transform technical system maps into business-relevant views that support decision-making across technical and non-technical stakeholders, creating a common language for discussing system relationships and dependencies.

The integration of compliance, technical, and performance data layers creates a multidimensional view that serves diverse stakeholder needs. Research into globally distributed database systems demonstrates the value of integrated views that combine multiple data dimensions (Eric Brewer, 2017). These integrated views enable organizations to understand the relationships between system architecture, performance characteristics, and operational constraints. Modern cloud platforms have significantly simplified the implementation of such globally distributed database systems, eliminating many of the traditional challenges around infrastructure provisioning, replication management, and geographic distribution. This cloud-based approach makes it more accessible for organizations to deploy and maintain the sophisticated data stores needed for comprehensive production understanding. Successful implementations typically employ graph-based data models that can represent complex relationships while supporting diverse query patterns tailored to specific stakeholder needs. These models enable flexible navigation of system relationships from different perspectives, allowing compliance teams to focus on regulatory requirements, technical teams to explore architectural dependencies, and performance engineers to analyze resource utilization patterns—all working from a consistent underlying representation of the enterprise ecosystem. This integration of multiple data layers creates a foundation for cross-functional collaboration and more effective decision-making regarding system evolution and operation.

Publisher: SARC Publisher

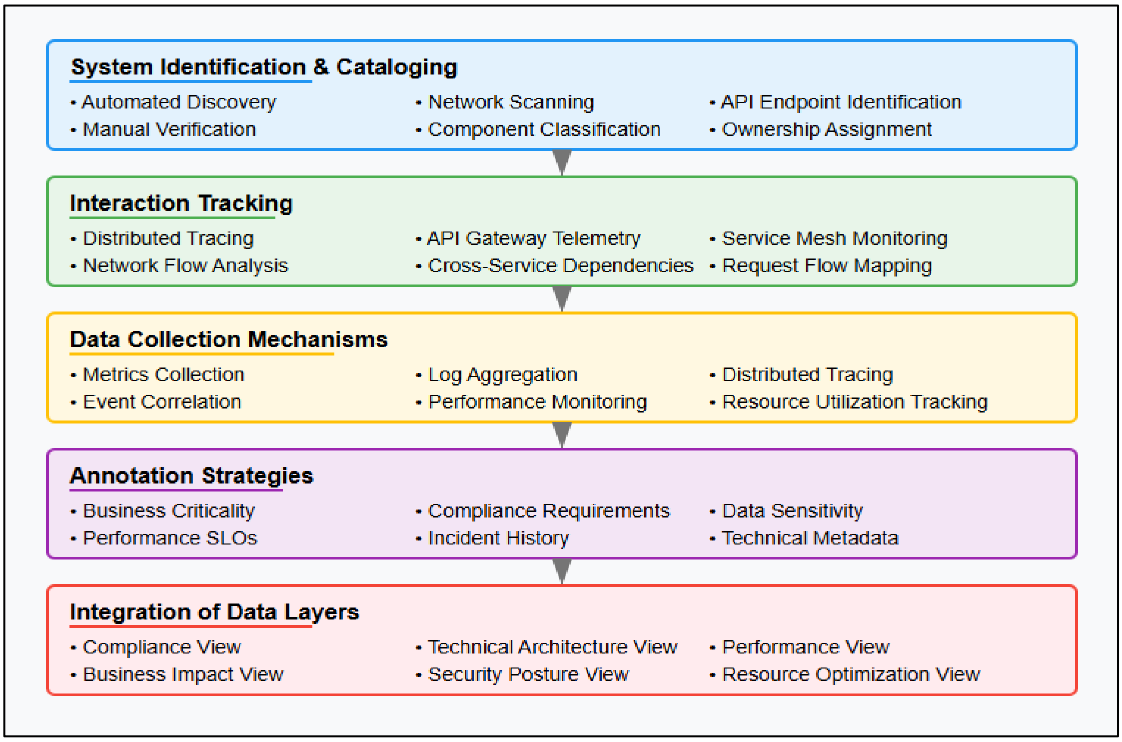

Fig. 1: Production Understanding Platform Framework. (Enterprise architecture bridging business and IT; GeeksforGeeks, 2024)

IMPLEMENTATION AND APPLICATIONS

Stakeholder-Specific Views

Site Reliability Engineering [SRE] teams derive substantial benefits from production understanding platforms through accelerated incident response capabilities and systematic reliability improvements. These platforms provide essential visibility into the complex interactions that define modern distributed systems, addressing key resilience challenges including fault tolerance, replication strategies, and failure detection mechanisms. When incidents occur, production understanding platforms enable SRE teams to navigate the inherent complexities of distributed systems and network partitioning scenarios. By visualizing system behaviors across normal and degraded operational states, these platforms help teams understand how systems respond to various failure modes, including crash failures, omission failures, timing failures, and Byzantine failures. This comprehensive visibility is particularly valuable in microservice architectures where traditional monitoring approaches often fail to capture the complex interdependencies between services. The integration of interaction maps with established resilience patterns—including circuit breakers, bulkheads, timeouts, and retries—enables more effective incident management and systematic reliability improvements. By mapping how these patterns are implemented across system boundaries, SRE teams can identify gaps in resilience strategies and implement more consistent approaches to fault tolerance. The visualization of distributed consensus mechanisms further enhances understanding of how systems maintain consistency and availability despite component failures, enabling more effective troubleshooting when these mechanisms behave unexpectedly.(GeeksforGeeks, 2024)

Leadership teams leverage production understanding platforms for enhanced compliance oversight and more effective strategic technology planning. These platforms address a fundamental challenge in enterprise environments—the persistent gap between operational technology [OT] and information technology [IT] that hampers effective decision-making and strategic planning. By creating comprehensive visibility across historically siloed domains, production understanding platforms enable leadership teams to develop more cohesive technology strategies that encompass both operational and information systems. This integration is particularly valuable in environments with substantial legacy infrastructure, where the lack of comprehensive documentation and visibility has traditionally impeded modernization efforts. Production understanding platforms help bridge this gap by mapping interactions between legacy systems and

Publisher: SARC Publisher

modern applications, revealing dependencies that might otherwise remain hidden. For compliance management, these platforms provide essential visibility into how regulatory requirements cascade across interconnected systems, enabling more effective governance and control implementation. The integration of compliance metadata with system interaction maps enables leadership teams to understand how requirements like data privacy regulations, financial reporting standards, and industry-specific mandates affect different parts of the technology landscape. This comprehensive view supports more effective resource allocation for compliance initiatives and enables more accurate reporting to regulatory authorities regarding control effectiveness.

Security teams utilize production understanding platforms to enhance vulnerability assessment processes and implement more effective risk management practices. These platforms address critical security challenges in distributed systems, including authentication, authorization, and secure communication between components. By mapping how security controls are implemented across system boundaries, security teams can identify inconsistencies and gaps in protection mechanisms that might create vulnerability to various attack vectors. Production understanding platforms are particularly valuable for addressing the security challenges inherent in distributed systems, including the expanded attack surface created by multiple communication channels, diverse authentication mechanisms, and complex trust relationships between components. The visualization of network boundaries, trust zones, and data flow patterns enables security teams to implement more effective defense-in-depth strategies that account for the complex interactions between systems. For vulnerability management, these platforms enhance prioritization by revealing how exploitable vulnerabilities might enable lateral movement across interconnected systems, helping security teams focus remediation efforts on the most critical issues. The integration of threat modeling with system interaction maps provides additional context for security decisions, enabling more effective resource allocation for security initiatives. This comprehensive approach is especially valuable in microservice architectures and cloud-native environments, where traditional perimeter-based security approaches are insufficient and security controls must be implemented consistently across numerous distributed components.

Resource optimization and capacity planning processes become more effective when informed by the cross-boundary visibility that production understanding platforms provide. These platforms address a fundamental challenge in enterprise environments—the integration of operational technology data with information technology systems to enable more effective resource management. By creating unified visibility across historically separate domains, production understanding platforms enable more holistic approaches to capacity planning that account for interdependencies between different technology stacks. This integrated view is particularly valuable for identifying bottlenecks that span multiple systems, where resource constraints in one component affect the performance of dependent systems. The visualization of request patterns and data flows across system boundaries reveals opportunities for workload balancing, resource sharing, and infrastructure consolidation that might not be apparent when examining systems in isolation. For capacity forecasting, production understanding platforms enhance accuracy by revealing how growth in one area affects resource requirements in dependent systems, enabling more proactive planning for infrastructure expansions. The integration of historical utilization patterns with dependency maps provides additional context for optimization decisions, helping technical teams identify inefficient resource allocation patterns and implement more effective sharing strategies. This comprehensive approach to resource management is especially valuable in complex enterprise environments with diverse technology stacks, where traditional domain-specific optimization approaches often fail to account for cross-system dependencies.

Case studies of successful production understanding implementations demonstrate the transformative impact these platforms can have on operational effectiveness and business alignment(GeeksforGeeks, 2024). Organizations implementing these platforms typically observe substantial improvements in resilience metrics across their distributed systems, including reduced recovery time objectives [RTOs] and implementing failure domains. These improvements stem from enhanced visibility into resilience mechanisms, including redundancy implementations, load balancing configurations, and failover procedures across interconnected systems. By visualizing how these mechanisms

Publisher: SARC Publisher

operate across system boundaries, organizations can identify inconsistencies and gaps in their resilience strategies and implement more effective approaches to fault tolerance. Production understanding platforms also enable more effective implementation of critical distributed system patterns, including consistent hashing for load distribution, leader election protocols for coordination, and gossip protocols for information dissemination. The visualization of these patterns in operation helps technical teams understand how distributed user/business flows behave in production environments and identify optimization opportunities. In environments with significant operational technology components, production understanding platforms help bridge historical divisions between OT and IT domains, enabling more effective data integration and cross-domain visibility. This integration supports improved operational efficiency, enhanced security posture, and more effective compliance management by creating a unified view of the technology landscape that encompasses both operational and information systems. Organizations that successfully implement these platforms typically adopt incremental approaches that deliver early value while building toward comprehensive coverage, focusing initially on critical business services and their supporting technology stacks.

Table 2: Stakeholder-Specific Benefits of Comprehensive System Interaction Mapping.

Stakeholder Group	Operational Challenge Addressed	Business Value Delivered
SRE Teams	Complex Incident Resolution in Distributed Systems	Improved Service Reliability and better design of failure domains.
Leadership	OT/IT Visibility Gap and Compliance Oversight	Enhanced Strategic Decision-Making and Regulatory Risk Reduction
Security Teams	Cross-System Vulnerability Management	Improved Threat Detection and More Effective Risk Prioritization
Infrastructure Teams	Siloed Capacity Planning	Optimized Resource Utilization and More Accurate Capacity Forecasting
Business Units	Technology-to-Business Service Mapping	Stronger Alignment Between Technical Investments and Business Outcomes

FUTURE DIRECTIONS AND RESEARCH OPPORTUNITIES

Cloud-native, SaaS-heavy enterprise environments present evolving challenges that will shape the future development of production understanding platforms. These environments introduce fundamental shifts in observability requirements due to their inherent characteristics of dynamic scaling, ephemeral resources, and distributed architecture patterns. The traditional "three pillars" approach to observability—logs, metrics, and traces—requires significant adaptation to address the unique challenges of cloud-native systems, particularly regarding data volume management, context preservation across system boundaries, and correlation of telemetry across abstraction layers. Research highlights several specific challenges, including the proliferation of microservices that generate exponentially more monitoring data, the ephemeral nature of containers and serverless functions that complicate persistent monitoring, the dynamic nature of auto-scaling that creates variable monitoring targets, and the increased complexity of distributed transactions that span multiple services with varying instrumentation levels. These challenges are further compounded by multi-cloud deployments that introduce heterogeneous monitoring environments and proprietary telemetry formats. Future production understanding platforms must address these challenges through unified observability approaches that implement standardized instrumentation methods, centralized telemetry aggregation, and context-preserving correlation techniques. Emerging solutions are beginning to leverage OpenTelemetry for standardized instrumentation, unified observability pipelines for telemetry normalization, and distributed tracing with context propagation to maintain visibility across service boundaries. The integration of service mesh technologies provides additional observability opportunities through proxy-level telemetry collection that requires minimal application instrumentation.

Automation opportunities in interaction mapping and analysis represent a promising area for advancing production understanding capabilities. The process of dependency mapping has traditionally been labor-intensive and error-prone, relying heavily on manual documentation that

Publisher: SARC Publisher

quickly becomes outdated in dynamic environments. Automated dependency mapping approaches address these limitations by continuously discovering and documenting relationships between infrastructure components, applications, and services based on observed behaviors rather than declared connections. These approaches employ multiple discovery techniques, including network traffic analysis that observes actual communication patterns between components, application instrumentation that tracks cross-service calls within code execution, log correlation that identifies related entries across distributed logs, and configuration analysis that detects declared dependencies in deployment artifacts. The integration of these techniques enables comprehensive dependency discovery across diverse technology stacks without requiring manual intervention. Research demonstrates that automated mapping provides significant advantages over manual approaches (Keith O'Brien, 2024), including greater accuracy through observation of actual behaviors rather than assumed connections, improved currency through continuous discovery rather than point-in-time documentation, broader coverage across heterogeneous environments, and reduced expertise requirements through algorithmic detection rather than tribal knowledge. Future developments will likely focus on expanding coverage across diverse technology stacks, improving the accuracy of dependency classification, and enhancing the contextual enrichment of discovered relationships with business and operational metadata.

Standardization possibilities for cross-organization production understanding present both significant opportunities and substantial challenges. As enterprise environments increasingly incorporate external dependencies and third-party services, the need for standardized approaches to dependency mapping across organizational boundaries becomes more pressing. Research indicates that effective cross-organization dependency mapping requires standardized approaches in several key areas, including common taxonomies for classifying component types and relationship characteristics, standardized data exchange formats for sharing dependency information between organizations, consistent metadata schemas for describing operational characteristics of external dependencies, and interoperable APIs for automated discovery and information exchange (Keith O'Brien, 2024). Industry initiatives are beginning to address these requirements through collaborative development of standards for dependency documentation and exchange, though significant challenges remain regarding competitive concerns, security implications, and technical complexity. The implementation of cross-organization standards would deliver substantial benefits, including enhanced visibility into supply chain dependencies that affect operational resilience, improved security posture through a comprehensive understanding of potential attack vectors across organizational boundaries, more effective change impact analysis for updates to shared services, and enhanced incident management for issues that span multiple providers. Future standardization efforts will likely focus on balancing transparency with security through selective information sharing that provides necessary operational visibility while protecting sensitive implementation details.

Integration with emerging enterprise architecture frameworks represents a promising direction for enhancing the business relevance and strategic impact of production understanding platforms. Traditional enterprise architecture approaches have often emphasized design-time documentation and strategic planning while lacking operational visibility into actual system behaviors. Conversely, cloud-native observability focuses on operational telemetry but may lack the business context and strategic alignment provided by enterprise architecture frameworks. Research indicates growing interest in bridging this gap through integrated approaches that combine the strategic perspective of enterprise architecture with the operational reality revealed through dependency mapping and observability. These integrated approaches enable bidirectional information flow, with enterprise architecture providing business context for operational observations while dependency mapping provides empirical validation of architectural assumptions. The implementation of these integrated approaches requires addressing several technical challenges, including reconciliation of different abstraction levels between architectural models and operational telemetry, alignment of taxonomies between architectural components and observable entities, temporal synchronization between architectural states and operational observations, and governance processes for maintaining consistency across perspectives. Emerging solutions leverage graph database technologies to represent both architectural models and operational dependencies

Publisher: SARC Publisher

in unified representations that can be navigated from different perspectives. This integration enables more effective digital transformation by providing comprehensive visibility into current operations within their business context while supporting strategic planning informed by operational realities.

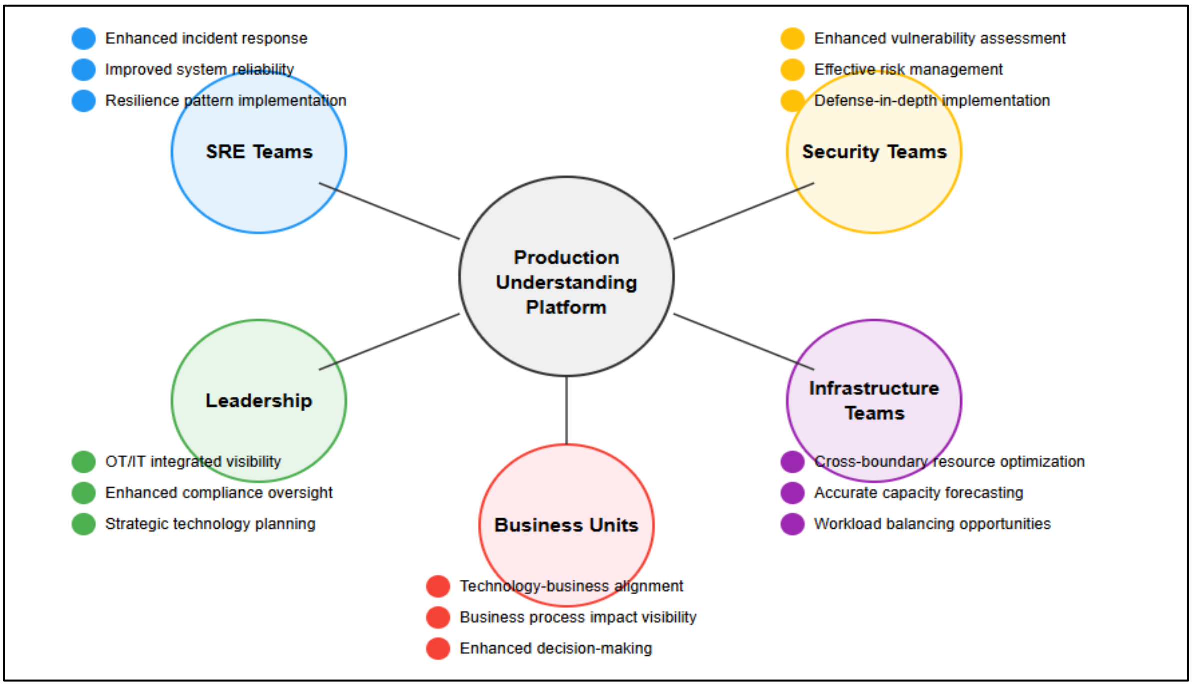

Fig. 2: Stakeholder-Specific Benefits of Production Understanding Platforms.

CONCLUSION

Production understanding platforms represent a paradigm shift in how organizations comprehend and manage their enterprise technology landscapes. By moving beyond static documentation to dynamic, data-driven representations of system interactions, these platforms enable organizations to bridge the gap between architectural intent and operational reality. The multidimensional views provided by production understanding serve diverse stakeholder needs while creating a common foundation for cross-functional collaboration. As cloud-native architectures and distributed systems continue to proliferate, the article demonstrates that production understanding capabilities will become increasingly essential for maintaining operational excellence. Future advancements in automation, machine learning applications, and standardization will further enhance these capabilities, enabling more proactive management of complex enterprise environments. Organizations that successfully implement production understanding platforms position themselves to achieve greater resilience, efficiency, and business-technology alignment in increasingly complex digital ecosystems.

REFERENCES

1. https://cloud.google.com/blog/products/devops-sre/using-the-four-keys-to-measure-your-devops-performance
2. "Accelerate: The Science of Lean Software and DevOps: Building and Scaling High Performing Technology Organizations" https://books.google.com/books/about/Accelerate.html?id=Kax-DwAAQBAJ
3. Lianping Chen *et al.*, "Characterizing Architecturally Significant Requirements," *ResearchGate*, 2013.
4. Eric Brewer, "Spanner, TrueTime & The CAP Theorem," 2017. https://static.googleusercontent.com/media/research.google.com/en//pubs/archive/45855.pdf
5. Enterprise architecture bridging business and IT
6. https://www.tietoevrytechservices.com/en/services/data-ai-automation/enterprise-architecture/
7. GeeksforGeeks, "Resilient Distributed Systems," 2024. https://www.geeksforgeeks.org/system-design/resilient-distributed-systems/
8. Keith O'Brien, Michael Goodwin, "What is dependency mapping?" IBM, 2024.

Publisher: SARC Publisher

https://www.ibm.com/think/topics/dependency-mapping

9. https://cloud.google.com/trace/docs/overview
10. https://opentelemetry.io/telemetry

Source of support: Nil; **Conflict of interest:** Nil.

Cite this article as:

Baskaran, H. P. "Mapping Production Understanding in Enterprise Spaces: A Framework for End-to-End System Interaction Analysis." *Sarcouncil Journal of Engineering and Computer Sciences* 4.11 (2025): pp 1-10.

Publisher: SARC Publisher