VIEW POINT

GROUNDED AI PIPELINES FOR LEAN IT SUPPORT

How organizations using Google Workspace can eliminate cognitive overload of support engineers, validate AI recommendations against live production schemas, and accelerate incident resolution, without additional licensing expenditure.

Executive Summary

With increasing system complexity, enterprise IT support teams often face structural efficiency challenges. Support engineers spend a disproportionate amount of their time searching for solutions instead of resolving issues. As production schemas evolve, traditional fixes can become ineffective or risky.

Additionally, organizational knowledge becomes fragmented with periodic team transitions. The consequences include inflated mean time to resolution (MTTR), recurring outage risks caused by deprecated scripts, and longer onboarding cycles that impact senior engineer capacity.

Infosys
Navigate your next

Solution at a Glance

This paper presents a practical alternative that leverages Google Workspace, a business tool commonly used by organizations today. Infosys combines artificial intelligence (AI) and data solutions to enhance IT support services while reducing resolution times. The Zero-API Grounded AI Support Pipeline integrates the following tools:

- Google Apps Script (GAS) for automated email ingestion
- NotebookLM (Workspace Business Standard and above) for AI grounding through native sync[1]
- Live GitHub SQL schema for production environment validation of database recommendations

The outcome is a single-click, prescriptive triage brief. It is grounded, cited, schema-validated, safe, and available when engineers start work.

Industry benchmarks indicate that AI-driven IT service management (ITSM) implementations consistently deliver 40–60% MTTR reductions.[2,3] AI-driven automation has reduced incident resolution lifecycles by 80 days[4]. The Zero-API Grounded AI Support Pipeline solution aims to achieve similar results while being cost-effective.

Key prerequisites

This solution requires:

- Google Workspace Business Standard or a higher edition[5]
- NotebookLM Plus, the enterprise-grade edition with Virtual Private Cloud Service Controls (VPC-SC) compliance, auditability, and zero model training on uploaded data

No additional AI platform procurement is necessary.

Introduction

The Zero-API Grounded AI Support Pipeline provides support engineers with an up-to-date artificial intelligence (AI) assistant at zero ongoing subscription cost. The AI assistant processes incident reports overnight and generates an actionable plan by morning.

Each night, the solution automatically updates its knowledge of open issues, historical resolutions, safety guidelines, and the latest technical system documentation. Engineers can generate a one-click, step-by-step triage guide the next morning. The solution outlines troubleshooting steps along with relevant reasoning and built-in safety checks. If the system detects potential risks, it recommends escalation rather than make assumptions.

An engineer always reviews and approves the actions before further changes. Resolved issues are automatically incorporated into the system's knowledge base during the nightly execution run. Consequently, the Zero-API solution continuously improves its knowledge and accuracy over time.

The Challenge: Three Interconnected Bottlenecks

Enterprise support engineering teams experience three interrelated bottlenecks, leading to system failures at critical stages.

The search loop

When a high-severity, Priority 1 (P1) incident occurs, a support engineer often begins by searching the database for related issues rather than directly resolving the incident. However, historical resolutions are scattered across closed tickets, legacy Confluence pages, and the knowledge of engineers who are no longer with the organization.

Conservative industry estimates suggest that engineers spend 20–25 minutes per P1 incident[6]. For perspective: A 50-person support team handling 200 incidents per day, reclaiming even 15 minutes of search time per incident equates to over 700 engineer-hours per month. This indicates a strategic resource allocation problem and a significant operational inefficiency.

The code drift issue

Enterprise data schemas are continuously evolving. Tables are renamed, column data types are modified, and stored procedures are deprecated. Yet, the SQL scripts embedded in historical incident resolutions are rarely updated.

When outdated SQL scripts are executed against the latest production environment, they run against an unknown schema. This can be risky in complex transactional systems, including order management, financial settlement, and inventory allocation. Potential risks include failed queries, ghost transactions, and corrupted processes.

Fragmented institutional knowledge

Organizational knowledge is spread across closed tickets, legacy runbooks, local script repositories, and engineer memories. If a critical incident occurs outside office hours when expert engineers are unavailable, it could lead to operational risk.

Onboarding new engineers in this scenario can take 4–8 weeks, impacting senior staff capacity and time. The Information Technology Infrastructure Library (ITIL) framework Version 4 identifies mean time to resolution (MTTR) reduction as a top IT operational priority[7]. Even so, many organizations continue to overlook fragmented and outdated institutional knowledge.

The compounding effect

These three bottlenecks are interrelated. For example, an engineer who cannot find an appropriate solution (bottleneck 1) is unable to validate the associated script (bottleneck 2), while the institutional knowledge required to bridge both gaps remains inaccessible (bottleneck 3).

During P1 incidents and service level agreement (SLA) deadlines, engineers tend to execute outdated or invalidated scripts, potentially leading to production outages.

The Solution: Zero-API Grounded AI Pipeline

The Zero-API Grounded AI Support Pipeline solution is built on a single architectural principle: The information required for troubleshooting incidents already exists in the organization. The challenge lies in rapidly assembling and validating existing knowledge, as well as delivering it when engineers need to take timely action.

Key features of the solution

Three unique design principles of this solution include:

- Live GitHub SQL schema as a grounding source: Every AI-generated recommendation involving a database operation is validated daily against the current schema. These include table definitions, ENUM constraints, column names, and stored procedures, retrieved directly from the GitHub repository. No other IT service management (ITSM) AI platform reviewed in open-source research verifies recommendations against a live production schema[8]. This is the solution's primary technical differentiator.
- Zero-investment large language model application programming interface contracts: The solution runs on Google Workspace infrastructure already used by many enterprises. NotebookLM Plus—the enterprise-grade, Virtual Private Cloud Service Controls (VPC-SC) compliant, and audit-trailed version—is a core service included with Workspace Business Standard and higher editions[5]. Google Apps Script is included with all Workspace editions, while authenticated users have free, read-only access to GitHub representational state transfer (REST) application programming interface (API). No additional platform investment is necessary.
- One-click human-in-the-loop governance: At the start of the shift, an engineer generates the brief, ensuring human-in-the-loop (HITL) governance. The brief involves generating a triage plan for the current incident, based on historical resolution data and the latest available technical documentation in NotebookLM. This is part of the governance requirement and not a technical

limitation. According to ECI Research's 2025 AI Builder Summit survey, only 44% of enterprise AI leaders are moderately confident in AI agents acting independently without human intervention[9]. By integrating human oversight into decision-making workflows, the single-click AI solution aligns with the chief information security officer (CISO) requirements. It is also compliant with client-specific audit and change management guidelines.

Architecture overview

The Zero-API Grounded AI Support Pipeline is a self-updating AI assistant that runs nightly. It reads the incident reports and generates an actionable, step-by-step, triage plan based on the available technical documentation in NotebookLM. Engineers can access this brief every morning, with a single click. As new issues are resolved, the AI assistant continuously updates its knowledge base during its scheduled runs, making it smarter and more effective.

This solution is cost-effective and does not require manual effort. Figure 1 below highlights the solution's triage workflow, from data ingestion to brief generation.

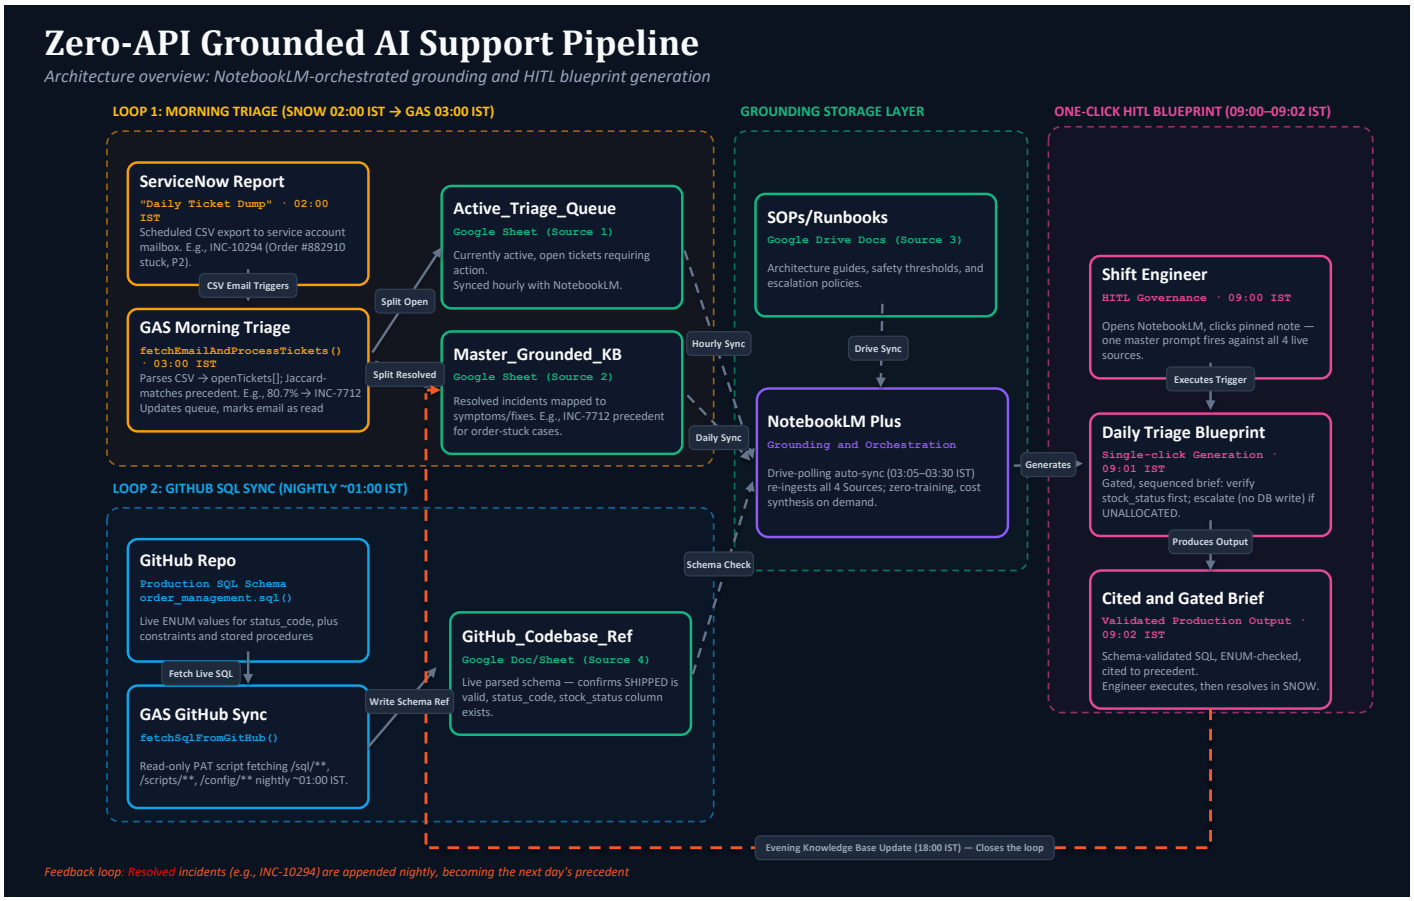

Figure 1: Triage workflow of the Zero-API Grounded AI Support Pipeline solution

The four grounding sources

Table 1 lists the documents and artifacts used to generate the triage plan. When analyzing open tickets, NotebookLM references relevant standard operating procedures (SOPs), technical documentation, knowledge base, and historical resolution data. Based on these insights, the AI assistant generates a step-by-step triage guide in just one click.

Table 1: Grounding sources used to generate the triage brief

Source	Content	Update Frequency	Value To Engineer
Active_Triage_Queue	Open tickets from ServiceNow (SNOW) report	Morning (Google Apps Script or GAS email parse)	Defines the day's workload: What needs resolving
Master_Grounded_KB	All historically closed tickets with resolution notes	Nightly (GAS auto-append)	Historical match: What worked before for similar issues
SOPs/Runbooks	Standard operating procedures (SOPs), escalation matrices, step sequences	On document update (Google Drive auto-sync)	Procedural steps, safety conditions, escalation paths
GitHub_Codebase_Ref	Live SQL schemas, stored procedures, configuration files from production repository	Nightly (GAS GitHub sync)	Schema validation: Confirms column names, ENUM values, and constraints before SQL is executed

End-to-end Data Flow: Order Processing Incident

Table 2 presents a hypothetical scenario: a Priority 2 (P2) incident involving Order #882910, which is stuck in a 'Processing' state. The workflow requires a database schema validation and SQL script corrections at every stage of the process.

This scenario, as illustrated in Figure 2, depicts the type of incident that the solution is specifically designed to address.

Table 2: Order processing involving a P2 incident: Stage-wise walkthrough

Stage	Time	Activity	Technology
0: Precondition	Previous night	NotebookLM Auto-Sync integrates four live sources. GitHub_Codebase_Ref includes order_management.sql with current ENUM values for status_code.	NotebookLM Auto-Sync[1]
1: SNOW report	02:00 IST	SNOW emails the daily incident report, 'Daily Open Ticket Dump', as a comma-separated values (CSV) file to a service account mailbox. Incident INC-10294: Order #882910 is stuck in processing, indicating a P2 severity.	ServiceNow Scheduled Report
2: GAS ingestion	03:00 IST	fetchEmailAndProcessTickets() parses the CSV file, INC-10294 openTickets[]. Jaccard similarity analysis matches the incident to INC-7712 from the knowledge repository at 80% confidence. Active_Triage_Queue Sheet is updated. Email is marked as read.	Google Apps Script

Stage	Time	Activity	Technology
3: Automatic syncing	03:05–03:30 IST	NotebookLM automatically synchronizes any changes to the linked documents in Google Drive. It then reingests all four grounding sources. **Order Status** status_code ENUM = PENDING|PROCESSING|SHIPPED|CANCELLED. SHIPPED is valid, stock_status column exists.	NotebookLM Auto-Sync[1]
4: Engineer initiation	09:00 IST	Engineer opens NotebookLM and clicks the pinned note. A master prompt is simultaneously executed against all four live sources.	NotebookLM Plus
5: Triaging Brief generation	09:01 IST	1. Step 1: Verify inventory status. Check stock_status (SELECT from inventory_ledger). If UNALLOCATED, escalate the incident, do NOT update the database [Source: INC-7712]. 2. Step 2: Restart the docker container. 3. Step 3: (if required): Update database schema. UPDATE order_headers SET status_code='SHIPPED' WHERE order_id='882910' AND stock_status='ALLOCATED' [Source: INC-7712 + order_management.sql – ENUM validated].	NotebookLM Plus
6: Execution	09:02 IST	1. Engineer validates the stock is in ALLOCATED state. 2. If required, the engineer executes the schema-validated SQL scripts with appropriate safety controls. 3. The ticket is resolved in SNOW with a resolution note.	Engineer decision
7: Updated knowledge base	18:00 IST	1. Evening GAS run detects incident INC-10294 as resolved. 2. GAS appends the resolved incident documents to Master_Grounded_KB for future reference. 3. The knowledge base is updated with two historical references for order-stuck scenarios.	Google Apps Script

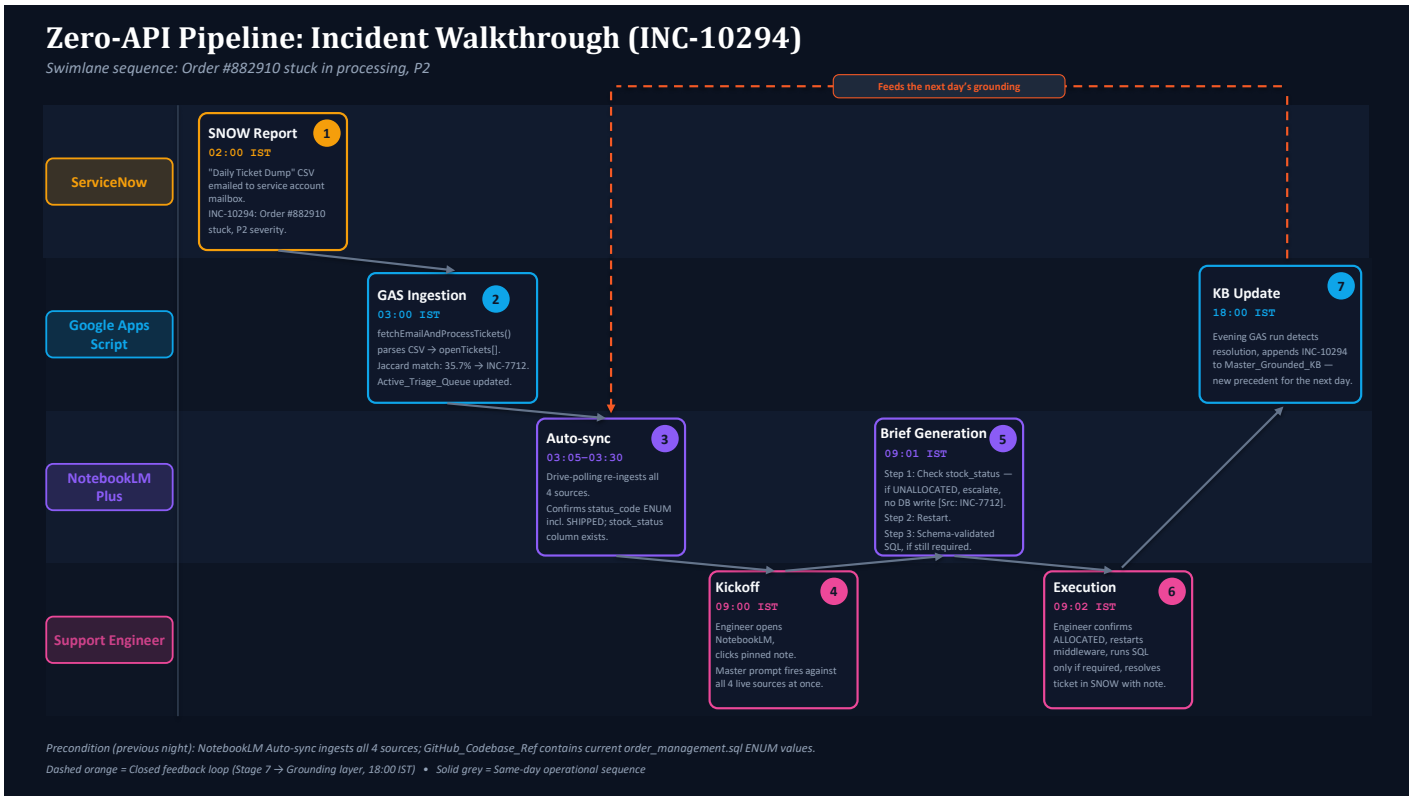

Figure 2: Order processing incident walkthrough

The importance of schema grounding

A conventional AI assistant may generate a triaging recommendation, based on a historical note in an outdated schema:

UPDATE order_headers SET status = 'SHIPPED'

However, this statement uses the wrong column name (status instead of status_code) without safety controls.

In contrast, the Zero-API Grounded AI Support Pipeline solution reads the current GitHub schema, confirming that status_code is the valid column and SHIPPED is the correct ENUM value. It also updates AND stock_status to ALLOCATED based on the schema constraints and the historical incident warning.

The difference is clear: Our system bridges the gap between historical documentation and current reality. While the solution leverages the existing knowledge base, its recommendations are aligned with the live production code. This ensures that the current schema updates, such as a column being renamed from status to stock_status, are automatically captured and remain accurate.

Industry Benchmarks: Benefit Projections

The benefit projections shown in Table 3 are derived from published industry benchmarks. They represent comparable implementation outcomes and not proprietary forecasts. Empirical validation through a controlled pilot study is required before making client-specific performance commitments.

Table 3: Industry benefit projections

Benefit	Industry Benchmark	Source	How This Solution Delivers It
MTTR reduction	40–60% reduction reported by enterprise AI-driven observability implementations	IR.com 2026 Guide[2]	Pre-validated, schema-grounded resolution workflows reduce diagnostic search and script validation times prior to investigation
AI/automation MTTR impact	40% improvement in MTTR reduction outcomes, achieved by enterprises effectively implementing AIOps	Rootly 2026[3]	GAS pattern matching and NotebookLM grounding replicate the AIOps model of pre-triage intelligence delivery
AI incident lifecycle	80-day reduction in breach and incident lifecycle; US $1.9 million savings per incident	IBM 2025[4]	Applicable to complex incidents, where schema-validated SQL prevents escalations and secondary incidents
AI ITSM case study	55% MTTR reduction and 70% ticket deflection in a documented multinational enterprise deployment	Auralis AI 2026[10]	A single-click prescriptive brief addresses the triage bottleneck; achieves deflection through confident self-service resolution
Engineer cognitive load	20–25 minutes of P1 incident triage time is the primary controllable MTTR contributor	ITIL 4 benchmark[7]	Manual search replaced by a pre-built, cited triage brief, reducing review time to under three minutes
Knowledge base value	Centralized, internal knowledge repositories significantly accelerate resolution and improve SLA compliance	Giva ITSM Research[11]	Master_Grounded_KB automatically updated every night without manual intervention; improved operational knowledge value

Challenges and Limitations

While this solution is feasible and easy to deploy, it is important to understand and address certain constraints before production. Table 4 outlines the solution limitations and their mitigation strategies.

Table 4: Solution challenges and their mitigation

Limitation	Description	Mitigation
Dependency on NotebookLM Plus	• Enterprise-grade NotebookLM Plus with VPC-SC compliance is included with Google Workspace Business Standard and higher editions[5] • Companies using lower versions of Workspace may need to upgrade or purchase a standalone licence[12]	• Confirm the client's Workspace licensing before planning • Workspace Business Standard is the minimum required deployment tier

Limitation	Description	Mitigation
Baseline incident matching with Jaccard similarity method	• The Jaccard similarity measure is computationally simple and appropriate for GAS • However, academic research indicates that term frequency-inverse document frequency (TF-IDF) and sentence embedding approaches provide higher accuracy for ITSM ticket classifications[13,14] • At the default similarity threshold of 0.05, generic IT vocabulary may produce false positives	• Use Jaccard similarity as the initial deployment mechanism to validate the pipeline • Transition to lightweight Hugging Face sentence embedding approach in phase 2
Single batch processing	• The pipeline runs once daily at around 03:00 AM IST • Severity incidents occurring afterward do not receive the pre-designed brief until the next scheduled run	• Document manual execution procedures • Introduce on-demand GAS run as phase 2 enhancement • Generate brief based on the current knowledge base
No automated last-mile email delivery	• The engineer needs to manually generate the brief in NotebookLM • This is a governance feature that maintains an active workflow	• Set team expectations during onboarding • Brief is ready before the shift starts; engineer generates it manually • This aligns with AI governance best practice for production operations

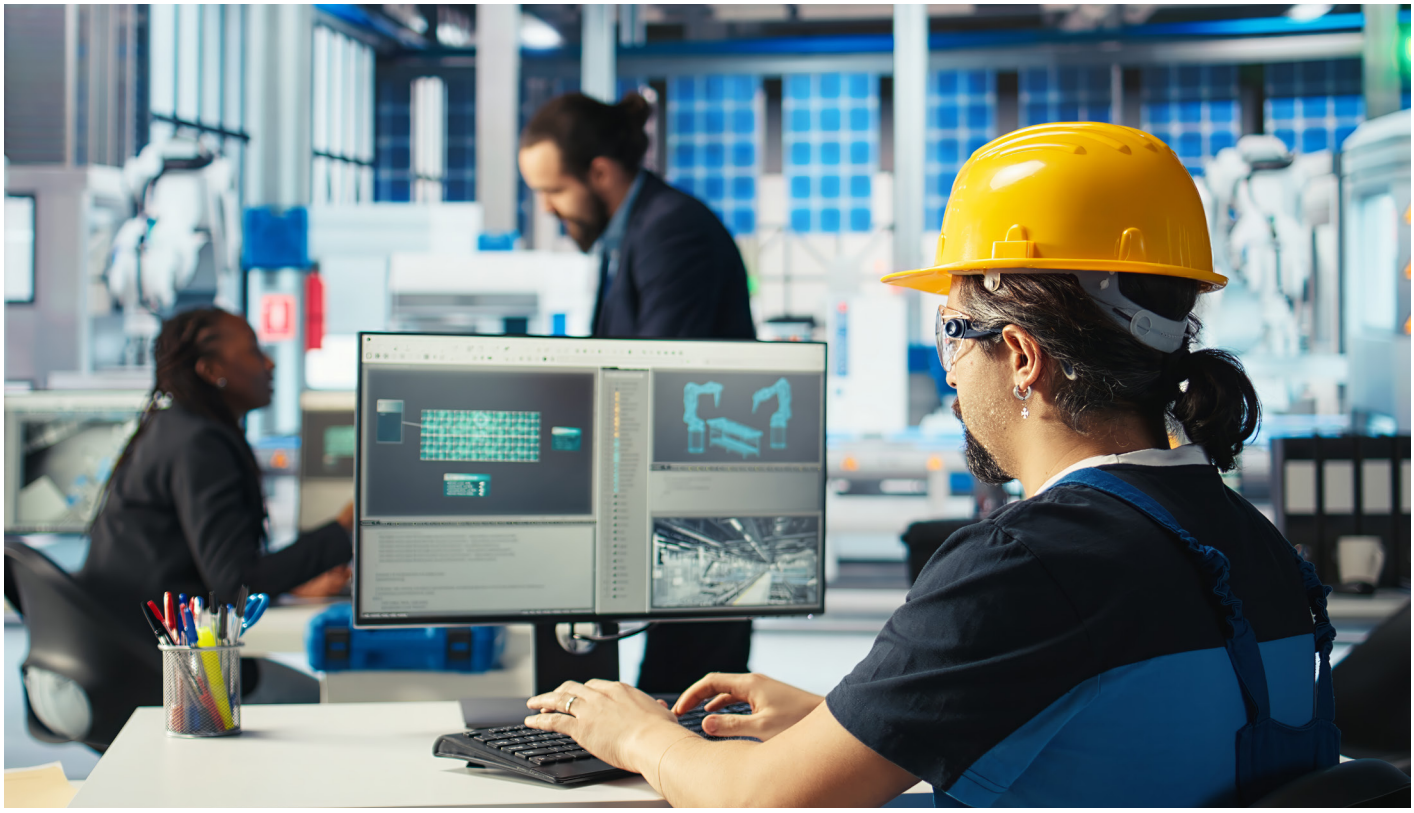

Conclusion

As IT systems evolve, there is a greater need for intelligent and adaptable frameworks that can effectively address potential issues on time. The Zero-API Grounded AI Support Pipeline is a novel, AI-powered solution that enhances IT support efficiency and significantly reduces MTTR. It constantly learns and refreshes its knowledge base to generate timely triage briefs that support engineers can use to troubleshoot incidents. The integration of human oversight with AI capability helps minimize system failures and improve accuracy.

The Zero-API solution leverages existing infrastructure and tools that enterprises already use, including NotebookLM, GAS, and GitHub repositories. Consequently, no additional software licensing or AI subscriptions are required to deploy the solution across enterprises, making it a cost-effective option.

About the Authors

Amit Kumar

Lead Consultant and AI advisor

Amit Kumar is a Lead Consultant and AI advisor with the Domain Consulting Group at Infosys. He advises manufacturing and industrial enterprises on digital transformation and AI adoption. Amit specializes in multi-million-dollar operational transformations through scalable, strategic initiatives that deliver business value and efficiency.

Vivek Shakya

Project Manager and AI explorer

Vivek Shakya is a Project Manager and AI explorer at the Manufacturing Application Development and Maintenance unit, Infosys. With his extensive project execution skills, Vivek oversees the design and integration of enterprise digital solutions. He has also played a pivotal role in the development of Infosys Zero-API Grounded AI Support Pipeline solution.

Hiren Gada

Project Manager and AI architect

Hiren Gada is a Project Manager and AI architect at the Manufacturing Application Development and Maintenance unit at Infosys and leads digital transformation programs. With his expertise in project management, commercial ownership, and business automation, Hiren drives scalable business growth and operational performance across enterprises.

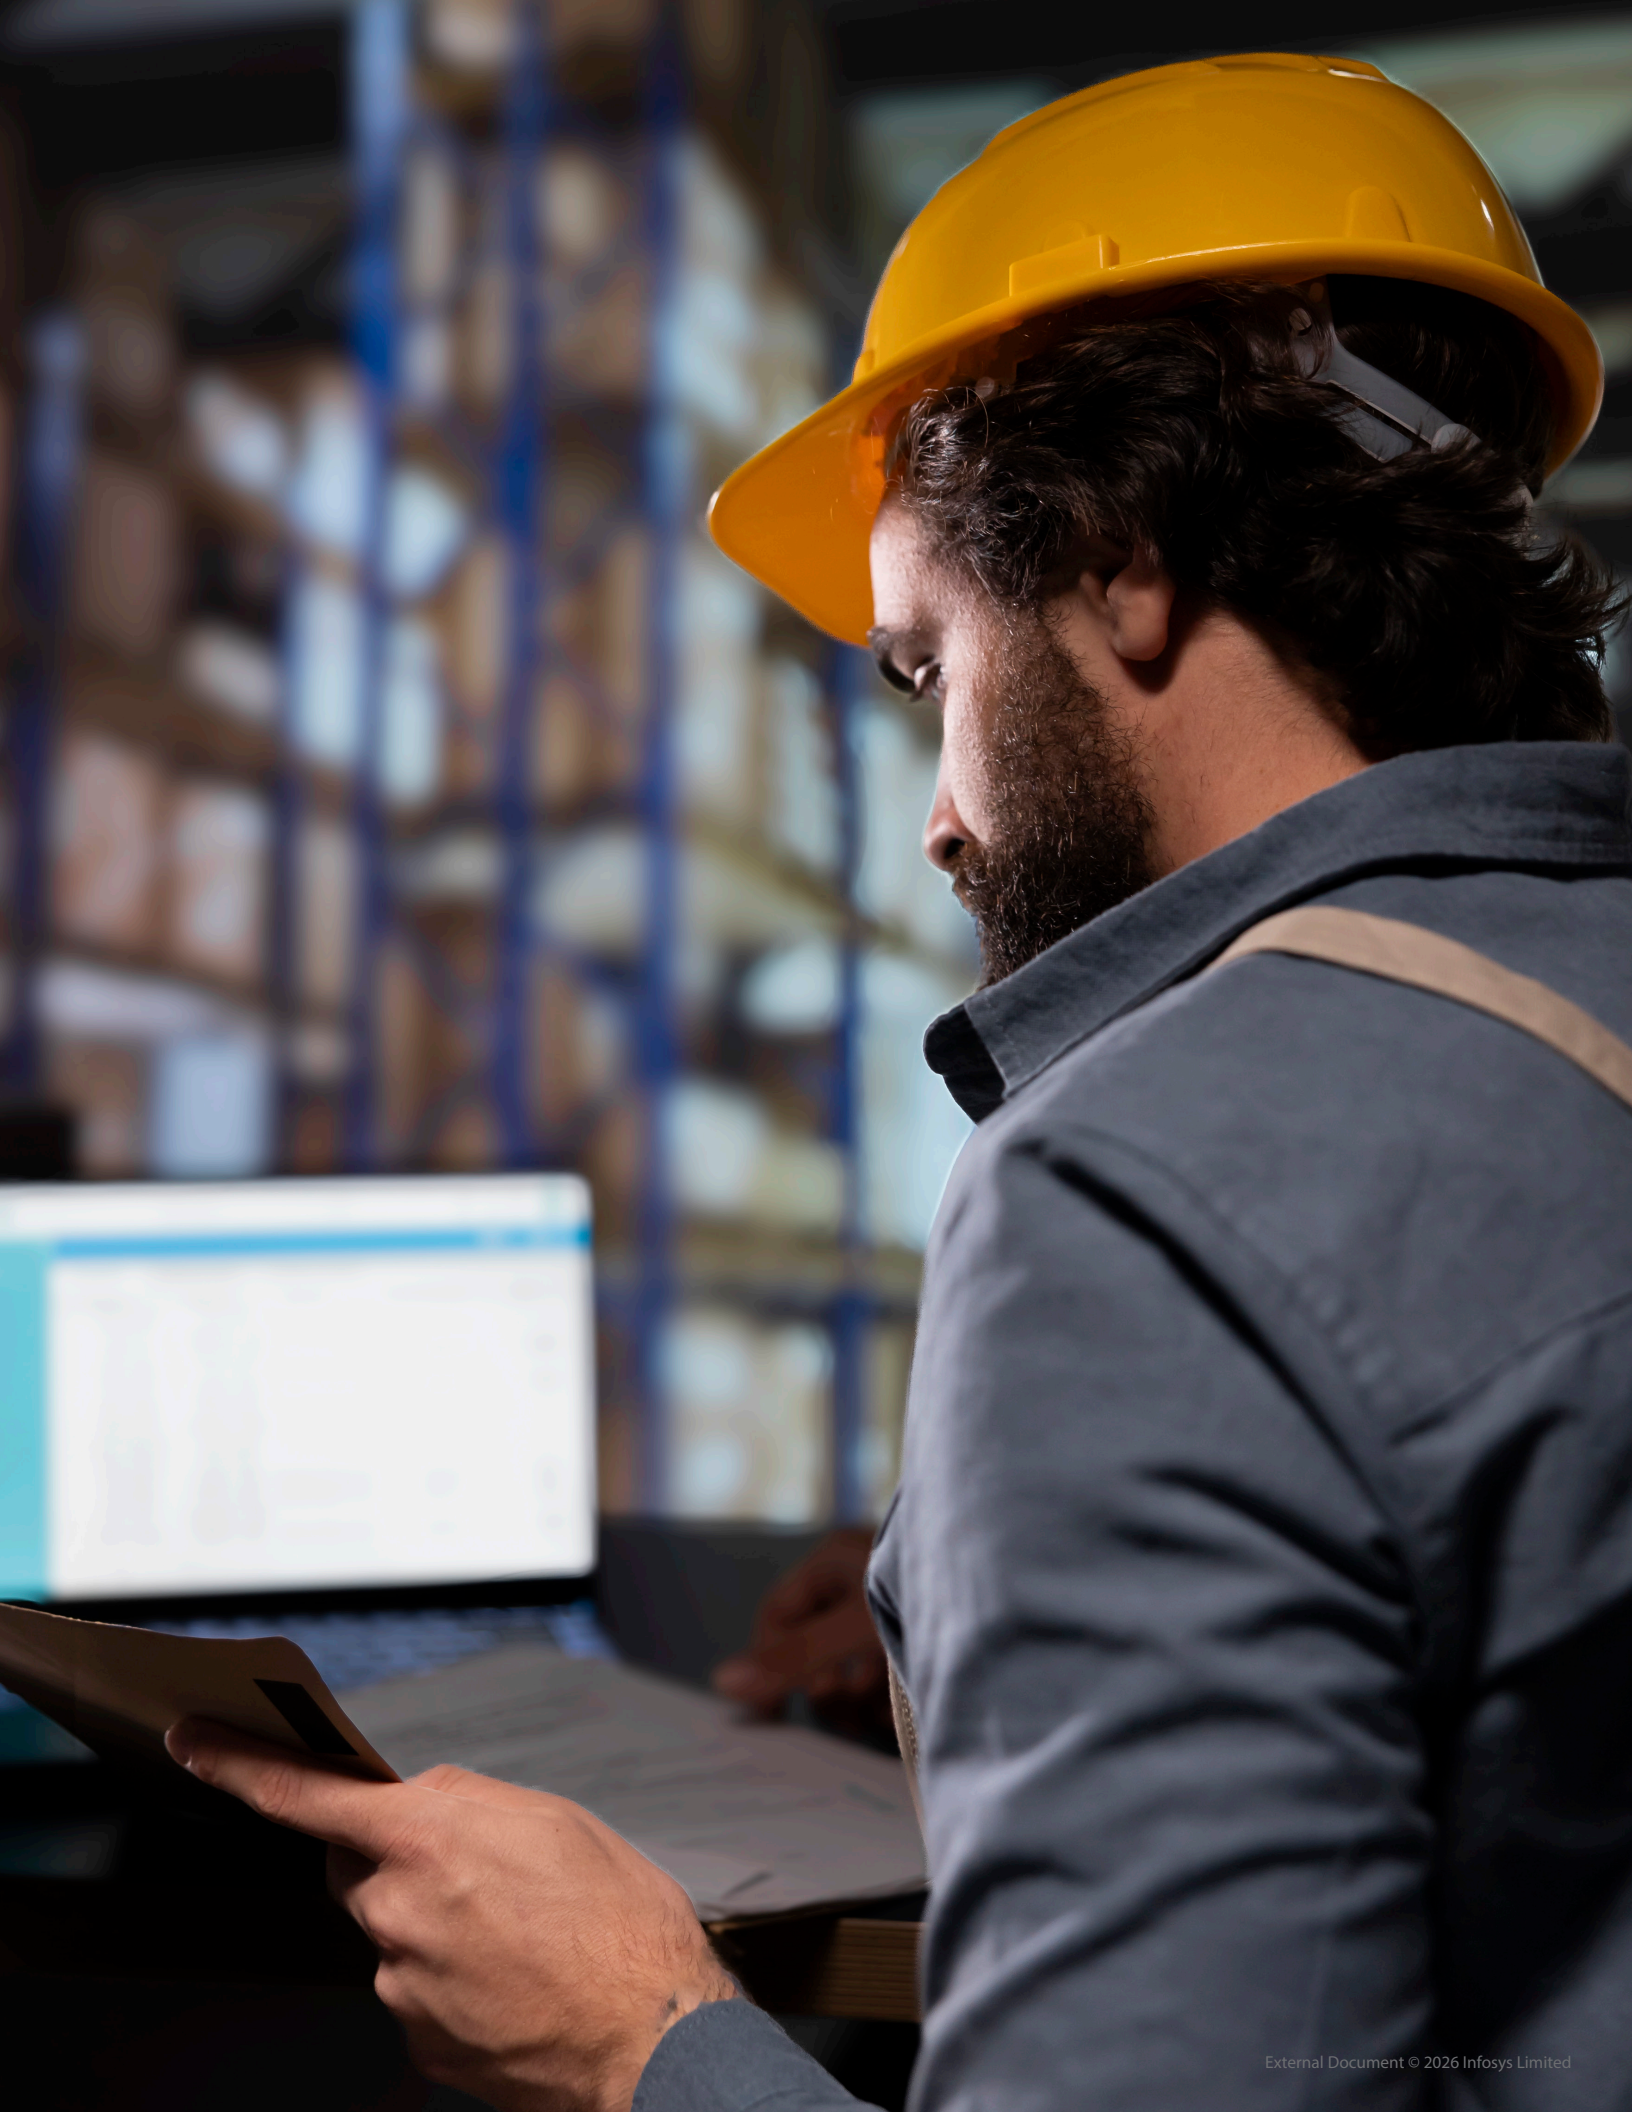

References

All references are verified as of June 2026. The URLs are accurate at the time of publication

1. Google Workspace Updates | Keep your sources up to date with automatic Drive syncing in NotebookLM
2. integrated research | HOW TO REDUCE MTTR WITH AI: A 2026 GUIDE FOR ENTERPRISE IT TEAMS
3. rootly ai | 2025 DevOps Trend: AI Incident Automation Cuts MTTR by 40%
4. Cyberhaven | MTTD and MTTR: Security Metrics Explained'
5. Google Workspace Updates | NotebookLM and NotebookLM Plus now available as a Google Workspace core service with enterprise-grade data protection
6. motadata | What is the Mean Time to Resolution (MTTR)? Why It Matters and How to Resolve
7. PLAN DO CHECK ACT | ITIL 4 Priority Matrix Templates Complete Guide to P1-P4 Incident & Problem Management
8. Draxlr | Best AI SQL Tools in 2026
9. ECI | ServiceNow Knowledge 2026: The Case for an AI Control Tower
10. AURALIS | Global enterprise reduced MTTR by 55% and achieved 70% ticket deflection with AI-driven ITSM automation'
11. Giva | What Is Time To Resolution: How to Calculate and Reduce This Critical Customer Support Metric
12. ELITE CLOUD | NotebookLM Pricing 2025: Free Plan vs Paid Plan' — Which One Actually Saves You Time?
13. arXiv | A CENTROID-BASED FRAMEWORK FOR TEXT CLASSIFICATION IN ITSM ENVIRONMENTS
14. arXiv | LEARNING TO PRIORITIZE IT TICKETS: A COMPARATIVE EVALUATION OF EMBEDDING-BASED APPROACHES AND FINE-TUNED TRANSFORMER MODELS
15. Google Workspace | Compare Flexible Pricing Plan Options
16. servicenow | ServiceNow launches the real-time data foundation that puts autonomous AI to work across the enterprise
17. CLOUDZERO | LLM API Pricing Comparison In 2026: Every Major Model, Ranked By Cost

For more information, contact askus@infosys.com

Infosys®
Navigate your next