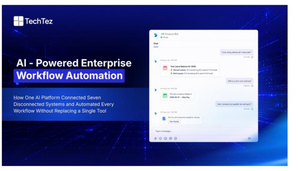

Case Study

From 7 Systems to One: How a Global Enterprise Automated Operations Across HR, Finance, IT, and Procurement

A large enterprise replaced fragmented approvals and disconnected portals with a single AI workflow layer inside Microsoft Teams and Slack, deployed in six weeks without changing existing systems.

The Situation

The client operated at scale. Thousands of employees. Multiple departments. Seven separate systems for routine operational requests. The result was predictable: coordination had become a full-time job.

Leave requests sat in email chains waiting for managers to remember to act. Expense claims routed manually across inboxes for 24 to 48 hours before anyone touched them. IT tickets were raised by email with no visibility. Procurement moved from request to purchase order in three to five days, manually, across email and ERP.

What Was Breaking And What It Was Costing

Enterprise operating on fragmented systems face compounding costs: higher employee time spent on coordination rather than meaningful work, slower decision cycles as approval chains grow, increased manager overhead from manual follow-ups, and mounting operational debt from inconsistent processes across departments. The productivity gap between AI-automated and manually operated enterprises widens by roughly 30% each year in high-output operational environments.

Key Challenges:

- **Requests were scattered across systems:** Employees had to use different portals for leave, expenses, IT support, procurement, and approvals. Most requests also involved emails and manual follow-ups, making the process slow and confusing.
- **Manual Approval Chains:** Approvals depended entirely on human memory and email. If a manager missed a notification, the request stalled. There was no automated escalation, no SLA tracking, no visibility into what was pending or overdue.
- **Disconnected Systems:** HR ran on one platform. Finance on another. IT on a third. Procurement, CRM, Assets, and Vendors each had their own tools none of which communicated. Every cross-department request required manual handoffs between systems.
- **No Unified Employee Interface:** Employees navigated five or more portals daily to complete routine tasks. Each tool required a separate login, a separate workflow, and a separate follow-up process. Productivity was lost not to complexity but to friction.
- **Governance and Visibility Gap:** Leadership had no single view of operational status across the organization. Approval rates, pending requests, SLA breaches, and department workloads were scattered across systems making informed decision-making slow and unreliable.
- **Deployment and Integration Overhead:** Connecting new automation to existing enterprise systems required significant manual configuration per integration, creating a long tail of implementation effort before any value was delivered.

The TechTez Approach :

Core Differentiator:

TechTez built a governed AI automation layer powered by N8N orchestration. Unlike standalone chatbots or disconnected automations, the platform coordinates workflows across departments, systems, policies, and approvers in a single layer. Employees interact in Microsoft Teams or Slack using plain language, while AI agents route, track, and manage requests with full human approval control and audit visibility built in.

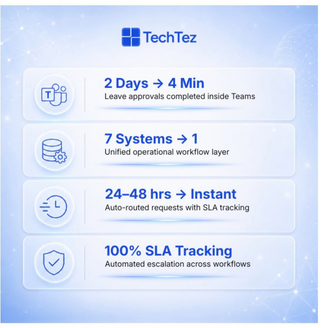

The architecture was guided by four decisions that shaped everything built after them:

- **Multi-agent AI built for each department:**A single general-purpose AI handling every department simultaneously produces generic outputs that fit no department particularly well. The decision was to deploy domain-specific agents one per department each trained exclusively on the policies, workflows, approval hierarchies, and business rules of its domain. When a request crosses department boundaries, agents hand off to each other automatically. No human coordination required.
- **N8N orchestration at the core:** Most automation platforms are built as rigid, pre-configured pipelines difficult to extend without rebuilding. N8N was chosen as the orchestration engine because every workflow is a composable, modular flow. New departments and new use cases can be added independently without touching what already exists. The platform was built for extensibility from day one.
- **Context-driven intelligence through RAG:** AI systems that infer answers from training data produce outputs that sound plausible but cannot be traced to a source. For enterprise use where a leave approval, a budget validation, or a vendor compliance check must be accurate and auditable inference is not acceptable. Retrieval-Augmented Generation ensures every policy check, balance verification, and routing decision is grounded in data retrieved from the organization's own systems in real time. Every response is traceable to its source.
- **Human-in-the-loop approvals:** Automation that removes humans from decisions creates governance risk. Every workflow requiring a decision routes to a human approver with full context before any action is taken. The AI prepares. The human decides. And if no decision is made within the defined SLA window, the system escalates so nothing stalls because someone forgot to act.

What TechTez Built:

The platform was delivered as a fully integrated conversational workflow layer accessible inside Microsoft Teams, Slack, or any existing enterprise interface. Here is what was built to make that possible.

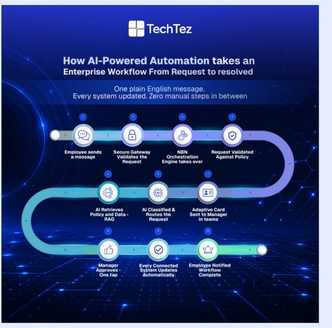

- **Conversational Request Engine:** Employees can raise requests directly inside Microsoft Teams, Slack, or other connected tools using plain language. The platform understands the request, captures the required details, and automatically routes it to the right workflow without forms, portal switching, or manual coordination.
- **Department-specific AI agents:** Dedicated AI agents support HR, Finance, IT, Procurement, and other departments. Each agent is aligned to that team's workflows, policies, and approval structures, while coordinating with other agents automatically when requests span multiple teams.
- **Intelligent approval workflows:** Once a request is validated, approvers receive a complete approval card inside Teams or Slack with all the necessary context, including request details, policy checks, and relevant employee or operational information. Approvals, rejections, or follow-ups can happen directly from the same interface without opening additional systems. SLA tracking and automatic escalation ensure requests do not remain stuck.
- **Built-in policy and compliance checks:** Every request is checked against company policies before reaching an approver. Leave balances, expense limits, procurement thresholds, vendor compliance, and other validations happen automatically in real time.
- **Integration with existing enterprise systems:** The platform connects with systems like SAP, Oracle, Salesforce, ServiceNow, ERP, HRMS, and internal databases without replacing them. Once a workflow is completed, all connected systems update automatically, removing manual data entry and follow-up work.
- **Real-time operational visibility:** Leadership teams get a unified dashboard showing approval status, SLA delays, workloads, pending requests, and workflow performance across departments in real time.
- **Enterprise Security and Governance:** Leadership teams get a unified dashboard showing approval status, SLA delays, workloads, pending requests, and workflow performance across departments in real time.

Platform Highlights

- Conversational workflow automation inside Teams, Slack, or any existing interface
- Intelligent approval routing with SLA tracking and automated escalation
- Policy and compliance validation at point of submission
- End-to-end integration with Existing Systems
- Automated system updates across HRMS, payroll, ERP, CRM, and calendar on workflow completion
- Real-time operational dashboard for leadership visibility
- Cloud-native with SSO, RBAC, secure gateways, and full audit trails

The solution delivered measurable efficiency, accuracy, and scalability improvements.

Tech Stack

AI & Intelligence:

- LLM-based intent understanding and request parsing
- Context-aware workflow routing engine
- Multi-agent orchestration framework
- RAG-based policy retrieval and validation
- Vector database for policy and knowledge storage

Workflow & Automation:

- N8N workflow orchestration engine
- SLA tracking and automated escalation workflows
- Human-in-the-loop approval framework
- Cross-department workflow coordination

Interface & Integration:

- Microsoft Teams and Slack integration
- SAP, Oracle, Salesforce, ServiceNow connectivity
- RESTful API and webhook connectors
- Legacy application integration layer
- ERP, HRMS, CRM, and ITSM integrations
- Adaptive cards for approvals and actions

Infrastructure & Security:

- Cloud infrastructure
- SSO and role-based access control
- Full audit logging and compliance trail
- Secure API gateway and authentication layer
- Monitoring, observability, and alerting systems