



The Proactive IT Service Desk

**A Unique Opportunity for AI Transformation,
Not Just Automation**

Mark Settle

November 2025

Contents

Introduction.....	3
Purpose of this paper	4
Conventional service desk operations.....	5
Transformational opportunities.....	7
Technical architecture	10
Proactive service desk performance expectations	11
Human staffing implications	12
The business case for a proactive service desk	13

About the Author

Mark Settle is a seven-time CIO, three-time CIO 100 award winner and two-time book author. His most recent book is *Truth from the Valley, A Practical Primer on IT Management for the Next Decade*. He can be reached at markfsettle@gmail.com.

Acknowledgements

Mark Blanke (The Executive Initiative), **Shuo Chen** (IOVC), **Stuart Evans** (Carnegie Mellon University), **Chris Flynn** (GoodLeap), **Tony Liu** (Costanoa Ventures), **Matt Peters** (Fixify), **Dan Nguyen-Huu** (Decibel Partners) and **Vijay Rayapati** (Atomicwork) provided thoughtful reviews of an earlier version of this paper. Their insights and suggestions are sincerely appreciated.

Disclaimer

Products and services referenced in this report are included for illustration purposes only. They do not represent personal endorsements on the part of the author.



Introduction

Much like death and taxes, every corporate employee is ultimately forced to deal with their company's IT service desk. For many it's not a pleasant experience.

Employees generally have unrealistic expectations regarding the IT resources they use to perform their jobs. They expect their company's IT organization to anticipate their needs and ensure that the tools and services they require work perfectly all of the time. Every service desk interaction starts off on the wrong foot because one of these preconceived notions has been proven to be wrong.

There are three critical components in every successful service desk interaction. Desk agents need to possess the *technical competence* required to resolve an issue in a *timely manner* while expressing *personal empathy* for an employee's immediate dilemma. Failure to satisfy one or more of these criteria results in an unsuccessful customer experience, even if an employee's issue is ultimately resolved.

Service desk performance plays a key role in determining the credibility of a company's overall IT organization. Company executives tend to judge IT effectiveness on the basis of IT's ability to deliver major projects on time and on budget. Their individual IT needs are addressed through personalized 'white glove' service and they rarely, if ever, interact with the service desk. However, the other 99%+ of the employees within a company view the desk as their primary point of contact with IT. Failure to address their issues on a timely, professional and consistent basis can severely undermine IT's organizational credibility and may ultimately limit the role IT is allowed to play in broader business initiatives.

Purpose of this paper

Service desks have traditionally operated in a reactive mode, reacting to continual changes in the composition and configuration of their company's IT assets and the daily deluge of employee requests, complaints and incidents. *The two principal constraints on traditional desk performance have been human bandwidth and technical knowledge.* Simply put, desks typically lack the staffing and technical competence required to address all incoming employee issues in a manner that meets employee expectations.

AI technology can alleviate these historical constraints. In many instances it can eliminate them altogether. There's no practical constraint on the number of employee-generated service desk tickets an AI system can process in parallel. Furthermore, LLMs have the ability to continually discover, update and retain the technical information required to resolve the vast majority of requests and incidents. Several startup companies have seized this opportunity and are developing new service desk applications that leverage AI to automate traditional operations.

The purpose of this paper is to suggest that AI can be used in a far more fundamental way to reimagine desk operations altogether, switching from a historical operating model built around reactivity to a new operating paradigm that emphasizes proactivity. AI can undoubtedly irradiate 50% or more of the work currently performed by human desk agents through automation. That provides a unique opportunity to refocus residual operations on activities that minimize employee inconvenience through anticipation of individual employee needs, proactive detection and resolution of incipient incidents and preemptive elimination of conditions that could trigger future incidents.

A proactive service desk can boost IT credibility and reduce IT operating costs but far more importantly it can reduce the time employees waste dealing with IT issues. Proactive service desk solutions represent a major investment opportunity that can pay multiple dividends in terms of workforce productivity and business agility.

Service desk terminology

The terms *request*, *incident* and *problem* have precise meanings in formal IT Service Management frameworks.

- A *request* is an employee application for access or delivery of a specific tool or service. It may trigger an approval, procurement or installation process. Requests may also seek information concerning the correct use of IT resources.
- An *incident* is technically defined as an operational failure or service degradation that has impacted an employee's work plans. In practice it may also result from human error or a lack of knowledge regarding the proper use of an IT resource, even when IT systems and services are functioning normally.
- A *problem* is the root cause of one or more incidents, typically requiring an extended period of time for accurate analysis and remediation planning. Third party assistance from system administrators or external vendors may be required to resolve problems.

This paper adheres to these definitions but also uses the terms *issues* and *concerns* more generally to refer to the collective suite of employee requests and incidents managed by the service desk.

Conventional service desk operations

Work materializes on the service desk in the form of system-generated or employee-submitted tickets. Tickets may represent system alerts or notifications, employee service requests or employee-reported operational issues. Tickets are categorized in terms of degree of difficulty and assigned to one of three levels.

Level 1 tickets represent common requests and incidents. They can be readily resolved with rudimentary knowledge of widely used IT resources, typically following prescribed procedures. If an employee-reported issue requires more extensive diagnosis or a deeper understanding of interactions among multiple systems, the ticket is routinely escalated to Level 2. Business application issues are typically forwarded to business operations teams and are not resolved by the IT service desk.

Level 2 tickets represent more complex issues or requests that involve multiple resources or new systems that have recently been acquired. Analysis may require the use of more sophisticated diagnostic tools or processes. Resolution may require the development of new or unique procedures. Tickets exceeding the skills of Level 2 agents are routinely escalated to Level 3.

Level 3 tickets represent the most baffling issues experienced by employees. Desk agents frequently require assistance from system administrators outside the service desk team to diagnose Level 3 issues. Resolution is likely to involve a combination of engineering modifications, configuration changes and procedural alterations. Some Level 3 tickets may be opened on a proactive basis by the service desk team to test and install updates to existing systems or monitor the behavior of systems that have recently been modified. Major updates or installations may also be managed as projects outside the normal ticketing process.




Figure 1. Average distribution of tickets managed by conventional service desks. Note that Level 1 tickets dominate traditional desk operations.

The technical complexity of service desk operations is illustrated in Figure 2. A wide variety of issues are submitted to the desk requiring varying degrees of knowledge about a wide variety of IT resources. It's unrealistic to believe that any single desk agent can maintain proficiency in all of the systems supporting daily workforce operations. Agents frequently develop areas of technical specialization over time and refer issues to one another accordingly.

IT Issues Most Commonly Encountered By Company Employees		
Category	Type	Common Vendor Products
Connectivity	- VPN - Wifi	   
Access	- Authentication - Authorization	   
Email/Calendar	- Loss - Synchronization	 
Hardware	- PC/Laptops - Peripherals - Printers	      
Endpoint Software	- Installation - Upgrades	  
Endpoint Security	- Viruses - Malware	     
Data	- Loss - Corruption	   

Figure 2. Common service desk issues. Vendors are referenced to illustrate the technical complexity of desk operations and do not represent the full range of products and services required to support daily employee operations.

Human desk agents are commonly classified as L1, L2 or L3 analysts depending upon their technical skills and operational experience. They primarily focus on tickets corresponding to their grade level but it's not uncommon for L1 analysts to occasionally work on L2 tickets, or for L2 analysts to perform selected L3 tasks. L1 jobs have historically been considered to be entry level positions in most IT organizations.

Conventional service desk performance metrics are displayed in Figure 3. They're predictably based upon the desk's ability to process tickets and please its customers. Efficiency is measured in terms of average resolution times, agent productivity and backlog. Customer experience is measured in terms of time to first contact with a human desk agent, first contact resolution rates, average resolution rates and customer satisfaction ratings following ticket closure.

Conventional Service Desk Metrics
Total Ticket Volume • Number of employee-submitted requests • Number of employee-reported incidents
Mean Time to Resolution • Level 1 tickets • Level 2 tickets
Customer Satisfaction Survey Results (post interaction)
Time To First Contact with Human Desk Agent
First Contact Resolution Rate
Ticket Backlog
Human Agent Productivity • tickets closed per desk agent per day/week

Figure 3. Conventional service desk performance metrics are designed to measure ticket processing efficiency and customer experience.

Transformational opportunities

The conventional metrics referenced above reflect the two principal constraints on traditional service desk performance. Response times increase, ticket backlogs grow and customer satisfaction declines because most desks don't have a sufficient number of analysts possessing the necessary technical expertise to resolve problems on first contact.

AI technology can radically alter these traditional constraints. AI agents have a near infinite capacity for work and they're immediately available on a 24X7 basis. Furthermore, they have ready access to a wealth of technical information that surpasses the technical capabilities of even the most sophisticated and well-staffed conventional service desk team.

IT agents can autonomously perform most of the activities historically performed by L1 human agents and a significant amount of conventional L2 work as well. This transformation is currently underway. A host of companies such as [Aisera](#), [Atomicwork](#), [Fixify](#), [Freshworks](#), [Moveworks](#), [Rezolve AI](#) and [Serval](#) have introduced new AI capabilities that automate many conventional desk practices.

However AI technology presents a far more radical and profound opportunity to transform traditional service desk operations altogether: from reacting to employee needs to anticipating and fulfilling such needs proactively; and from resolving employee-reported issues to suppressing or preventing such issues entirely. Simply put, AI presents an opportunity for a wholesale reimagining of the purpose of the service desk, from fixing employee-reported issues to minimizing IT-related distractions that undermine employee productivity.

The evolution of conventional desk operations to a new proactive operating model will require the creation of new tools and capabilities. The following capabilities will play a critical role in enabling this transformation.

Productivity threat detection

The amount of information IT organizations collect regarding the health and behavior of their company's IT resources is staggering. Their ability to process, analyze and interpret such data in real time is staggering as well.

Common IT Telemetry Tools		
Domain	Purpose	Common Vendor Products
Enterprise Security	Posture Management (cloud, network, endpoint, identity, applications, data)	        
	Extended Detection & Response (XDR)	    
	Security Information & Event Management (SIEM) Security Orchestration, Automation & Response	   
System Observability	Availability & Performance	  

Figure 4. A partial depiction of the IT telemetry tools commonly employed by large enterprises. Although there may be a significant overlap in the resource coverage of multiple tools, individual tools were designed for different purposes and consequently possess different functionality.

Figure 4 displays some (but certainly not all) of the monitoring tools commonly employed by large enterprises. Many are used for security purposes to stymie unwanted intruders or deter insider threats. Others are used to manage the utilization of infrastructure resources such as networks, servers and storage devices to minimize congestion, enforce SLAs or control costs.

Conventional service desks make relatively limited use of these telemetry tools because they lack the time or training to do so effectively. However, it's readily apparent that the data collected by such tools could be used to identify conditions that might trigger a future employee incident or to detect an emerging incident that could be auto-resolved before disrupting an employee's workday.

In much the same way that security operations teams use telemetry signals to proactively detect and mitigate cyber threats, service desk teams could leverage many of the same signals to mitigate threats to employee productivity. Dedicated swarms of AI agents can, in principle, maintain continuous surveillance of the health, availability and responsiveness of commonly used resources from an employee perspective and issue warnings and alerts regarding impending productivity threats. Multi-agent systems can maintain an internal collaborative conversation about the occurrence and significance of observed anomalies; make judgements concerning the potential impact of such anomalies on the productivity of individual employees; and take corrective actions autonomously.

Note that agentic surveillance concepts are already being implemented to support Site Reliability Engineering (SRE) teams. Notable startups in this area include [Dash0](#), [Neubird](#) and [Traversal](#).

Continuous persona testing

Work teams in different business departments have very different IT needs and IT-related work practices. For example, a geographically distributed HR Recruiting team that largely works from home employs very different resources and behaves quite differently than a centralized Financial Planning & Analysis (FP&A) team that works exclusively at corporate headquarters.

AI agents are uniquely well suited to observe and emulate the IT behavior of individual work teams. They can establish personas that mimic the ways in which a team employs IT resources and use these personas to exercise team-specific systems and services on a routine basis. The stochastic behavior of LLM-based agents provides an ideal means of simulating the stochastic behavior of human teams. They are quite capable of serving as digital twins that depict the tribal IT practices of such teams.

This type of continuous testing service is analogous to the continuous penetration testing services employed by cybersecurity teams. Agentic persona testing would be used to ensure the availability and responsiveness of the specific resources used by individual work teams and provide early detection of any potential or emerging issues. Any out-of-the-ordinary conditions spotted by such agents could be automatically shared with the swarm of productivity threat detection agents described above for further analysis.

These agentic personas would evolve over time as the business practices of individual work teams change. In addition to serving as an early warning system for issue detection, they could also be used to anticipate the IT resource needs of new team members or proactively offer suggestions regarding resource use to team members working on different projects at different points in time.

Human team members could share personal concerns regarding resource availability, responsiveness, synchronization, etc. directly with their team's dedicated persona agent, avoiding the formality (and hassle) of filing a service desk ticket. Groups choosing to use persona agents in this fashion might even assign them nicknames and 'adopt' them as quasi-team members.

Personalized 'white glove' employee service

Traditional service desks typically do a poor job of communicating with their customers. Employees complain that they're not receiving enough information about the status of a request or incident, or too much information. They complain that the information they're receiving is too technical or not technical enough. They may not like the demeanor of the human agent they're dealing with, or they may prefer email or text communication to avoid conversations with desk agents altogether. Every customer has their own interaction preferences and it's literally impossible to please all of the people all of the time.

AI agents can remedy this communication dilemma. Agents can be created on the fly with complete knowledge of an individual's IT resource profile (i.e. the tools and services they use on a routine basis), their past interactions with the desk and their personal communication preferences. Such agents can deliver customized concierge treatment to individual employees that is similar to the white glove treatment experienced by company executives. They would enable individual employees to communicate with IT when, where and how they want to.

Technical architecture

Proactive desk platforms possessing the capabilities outlined above will employ both predictive (machine learning models) and generative (large language models) forms of AI technology. They will consist of an extended network of AI agents that operate in a highly collaborative and frequently autonomous fashion.

Surveillance agents monitoring the behavior of specific IT systems will not only share data with one another but also communicate with persona agents that are continuously emulating the behavior of specific work teams. Other agents may be responsible for determining when to respond to a potential risk or initiating a corrective action at the first indication of an emerging incident. This network may also interact with agents in other functional departments such as security or financial operations. And finally, this network may commission the construction of transient agents that can communicate directly with individual employees or business teams regarding IT issues on an 'as needed' basis.

Surveillance agents will function as a Mixture of Experts (MoE), each possessing deep insight into the nominal operation of specific systems. Such agents not only evaluate the end user risks posed by current system performance but also the likely success of corrective actions designed to resolve end user issues. In a sense they emulate the work practices of human service desk teams that enlist the aid of human system administrators in diagnosing the root cause of incidents and evaluating alternative remediation actions.

As a risk mitigation measure, a separate Panel of Experts (PoE) agentic framework would be used to independently assess proposed corrective actions. PoE agents would evaluate the upstream and downstream consequences of proposed actions that may not be anticipated by system-specific surveillance agents. PoE recommendations may require human approval under predefined circumstances.

The network envisioned above is strictly conceptual in nature and far from a design specification. Multi-agent networks – sometimes referred to as 'swarm architectures' – are still in their infancy. Tradeoffs between the cost, complexity, reliability and adaptability of such networks remain to be determined.

Nevertheless, the early success of selected firms such as [Aisera](#), [Fixify](#), and [Moveworks](#) in using AI agents to resolve incidents through employee self-service suggests that agentic network solutions for proactive desk operations are feasible.

The multi-agent network supporting proactive desk operations will have ample opportunities for reinforcement learning. Although there are no definitive studies regarding the number of tickets generated by individual employees, traditional desks typically receive 20-60 tickets per month for every 100 employees. This would equate to 2,000 to 6,000 AI learning opportunities per month in a 10,000 person corporation.

Note that these learning events will primarily be used to understand the underlying causes of employee-reported issues in order to prevent or minimize their reoccurrence. In contrast, many of the agentic service desk solutions currently under development are primarily concerned with accelerating the resolution of issues by introducing new agentic response practices.

Proactive service desk performance expectations

Leading service desk solution providers have devoted considerable effort to deflecting incoming tickets through employee self-service mechanisms. Employees may be provided with knowledge articles, chatbot instructions or automated procedures that eliminate the need to interact with a human agent. Although there are no objective research reports documenting the success of such mechanisms, anecdotal stories of customer success reported on vendor websites suggest that deflection rates of 20-40% are achievable on average. Higher rates of 40-60% are achievable under specialized circumstances.

Self service introduces a significant complication in measuring service desk performance. The conventional metrics displayed in Figure 3 are implicitly based upon traditional human-in-the-loop (HITL) operations. Measures of first time to human contact, first contact resolution, ticket backlog and human agent productivity are meaningless in self-service scenarios. Measures of desk performance that fail to discriminate between HITL and non-HITL scenarios are similarly meaningless. New metrics such as the number of tickets promoted to human agents after a failed attempt at self service or self service abandonment rates are necessary to truly understand desk performance from a customer perspective.

Proactive Service Desk Metrics
Total Ticket Volume • Autogenerated tickets / & resolved with no reported employee impact
Mean Time to Resolution With HITL • Employee-submitted requests • Employee-reported incidents
Customer Satisfaction Survey Results (initial/during/end) • Tickets with HITL • Tickets with no HITL
L2/L3 Incident Reoccurrence • Same L2/L3 incident/same employee • Same L2/L3 incident/different employee
Problem Backlog
Human Agent Productivity • Percentage of time devoted to L2/L3 concerns and problem ticket management

Figure 5. Proactive service desk performance metrics are designed to measure issue prevention, disruption avoidance, employee productivity loss and root cause incident resolution.

Proactive service desks of the future will measure their performance in fundamentally different ways as illustrated in Figure 5. Primary emphasis will be placed on the autodetection and preemptive resolution of conditions that would likely result in some type of employee incident or inconvenience. Secondary emphasis will be placed on the detection and resolution of an issue that has actually occurred *before it is reported by any employee*. This point deserves special attention. Even though most IT systems are designed to provide early warnings of anomalous conditions, service desks routinely receive employee reports of incidents that have gone undetected and unanticipated by existing monitoring procedures.

Note that any alerts, analyses or actions initiated by proactive AI agents will be recorded and tracked in the form of tickets as a means of monitoring the balance between proactive and reactive desk operations.

Human agents on a proactive desk will spend the majority of their time resolving L2 and L3 concerns and aggressively managing formal problem tickets. Sophisticated sentiment analysis tools will be used to monitor customer satisfaction throughout a service desk experience. These tools can gauge customer sentiment from the beginning through the end of a desk experience (with and without HITL) far more accurately than traditional post-engagement customer surveys. Note that several existing desk solutions currently offer this capability as well.

It's tempting to establish speculative design goals for proactive desk solutions that fully leverage existing AI technology. Aspirational targets for such solutions might be:

- A 35% reduction in the total volume of inbound employee requests and employee-reported incidents irrespective of whether they're resolved by a human agent or a self-service mechanism
- Agentic generation of one-quarter to one-third of all tickets managed by the desk
- L2/L3 incident reoccurrence rates of less than 2% for all employees
- Two-thirds or more of human agent time devoted to L2/L3 concerns and problem ticket management

80% or more of the L1 tickets received by conventional desks should be preemptively avoided or resolved with no employee impact, or resolved through employee self-service. This represents a radical change in traditional desk operations.

Human staffing implications

It's unrealistic to expect that all of the IT issues encountered by a company's employees can be autonomously resolved by a swarm of AI agents, eliminating the need for human agents altogether. However, the roles and responsibilities of human agents staffing a proactive desk will be radically different from those of the past.

Human desk agents have traditionally functioned as jacks-of-all-trades, possessing superficial knowledge of many commonly used IT systems. Over time they typically develop deeper knowledge of selected systems and are viewed as local experts on such systems by their co-workers. Some may even receive formal training and become certified administrators of selected systems, knowledge that is essential in resolving many L2 and L3 issues.

The jack-of-all-trades roles that are so prevalent on conventional service desks will largely disappear on the proactive desk. However, the need for individuals with in-depth knowledge of selected systems will remain. Human administrators will continue to play a critical role in authorizing and applying changes to production systems for the foreseeable future.

Wholly new roles will be required to manage the performance and efficiency of the AI agent network supporting desk operations. *Inference Analysts (IAs)* will be needed to monitor the chain of thought processes employed by AI agents in diagnosing issues and constructing scenario-specific solutions. IAs will evaluate deliberations among multiple AI agents to ensure that alternative root causes of reported issues have been thoroughly explored and that the potential side effects of any planned autonomous actions have been properly identified and discounted. IAs will continually evaluate tradeoffs between context window size, response latency and cost of LLM-based tools, and may even

establish requirements for the collection of new forms of data to improve the accuracy and timeliness of autonomous AI actions.

Agent Network Engineers (ANEs) will monitor interactions among multiple AI agents, determining an appropriate balance between the scope of responsibilities assigned to individual agents and the efficiency/effectiveness of the overall network. ANEs may expand or contract the network over time based upon observed performance. ANEs will ensure that agent-to-agent interactions are fit for purpose and that IT access and authorization controls are being properly enforced.

Agent Developers (ADs) will develop new agents or modify the scope or skills of existing agents based upon requirements established by ANEs and IAs.

Human agents in these new roles are performing more than a quality control function. They are optimizing the business performance of a system of agents in terms of the success metrics displayed in Figure 5, namely, issue prevention, employee disruption avoidance and timely resolution of employee-reported concerns. At the same time they are optimizing the technical performance of that system in terms of adequacy, resource consumption, latency and cost.

The business case for a proactive service desk

Conventional service desks are becoming anachronisms, a vestige of earlier times in which the IT resources supporting business operations were far less complex and the employees using those resources were far less technically sophisticated. Present day employees generally dread contacting their service desks because (a) it's an admission of their own inability to resolve a vexing technical issue, (b) the interaction will simply extend and likely expand the disruption of their current workday plans, and (c) the likelihood of capturing the attention of a human agent who has the time and skill needed to resolve their issue is low. It's no wonder that employees frequently resort to the tribal knowledge of co-workers to avoid service desk interactions whenever possible.

Self-service mechanisms that eliminate human agent interactions are a step in the right direction but they are still designed to resolve issues that have actually occurred and are disrupting business operations.

A more fundamental approach to avoiding IT-related losses in workforce productivity is to avoid the occurrence of employee-impacting IT issues altogether or resolve them at the earliest stages of emergence before they are able to materially disrupt an employee's work plans. That is goal of a proactive service desk.

Investments in proactive desk solutions should be based on improvements in a company's return on investment in its human resources. A proactive desk minimizes IT-related workforce distractions and related losses in productivity. Unfortunately, the impact of such distractions is hard to measure and any improvements in productivity due to their elimination are impossible to track. Consequently, we are forced to resort to metrics such as ticket counts and resolution times as indicators of productivity gains and losses. They are admittedly poor surrogate measures of true productivity improvement.

Perhaps investments in corporate cybersecurity provide a useful analogy. Modern enterprises spend considerable sums on cybersecurity defense to avoid the loss of sensitive information or the disruption of their internal operations. These investments are made to reduce the risk of such events even though it's difficult if not impossible to quantify the risk reduction achieved with each incremental expenditure. Investments in proactive desk operations are quite similar, in that they will undoubtedly result in workforce productivity improvements even though such gains are difficult if not impossible to quantify.

Investments in proactive desk solutions should be justified as financial annuities – a means of buying back lost productivity that will increase in financial value as a company’s workforce expands, or, alternatively, if its workforce contracts around a subset of higher-paying jobs. Such investments provide a means of increasing a company’s effective workforce regardless of its current business conditions without incurring any additional labor costs. That’s a hard deal for any corporate executive to pass up!

IT professionals pride themselves on ‘futureproofing’ investments in new technical capabilities to ensure that such capabilities will deliver tangible benefits over long periods of time under a variety of business circumstances. A proactive service desk is one of the most futureproof investments that any company can make in its foundational IT assets.