

ENTERPRISEOPS-GYM: ENVIRONMENTS AND EVALUATIONS FOR STATEFUL AGENTIC PLANNING AND TOOL USE IN ENTERPRISE SETTINGS

Shiva Krishna Reddy Malay^{†,1}, Shravan Nayak^{†,1,2,3}, Jishnu Sethumadhavan Nair¹, Sagar Davasam¹

Aman Tiwari¹, Sathwik Tejaswi¹, Sridhar Krishna Nemala¹, Srinivas Sunkara¹, Sai Rajeswar^{1,2,3}

¹ServiceNow Research

²Mila - Quebec AI Institute

³Université de Montréal

ABSTRACT

Large language models are shifting from passive information providers to active agents intended for complex workflows. However, their deployment as reliable *AI workers* in enterprise is stalled by benchmarks that fail to capture the intricacies of professional environments, specifically, the need for long-horizon planning amidst persistent state changes and strict access protocols. In this work, we introduce EnterpriseOps-Gym, a benchmark designed to evaluate agentic planning in realistic enterprise settings. Specifically, EnterpriseOps-Gym features a containerized sandbox with 164 database tables and 512 functional tools to mimic real-world search friction. Within this environment, agents are evaluated on 1,150 expert-curated tasks across eight mission-critical verticals (including Customer Service, HR, and IT). Our evaluation of 14 frontier models reveals critical limitations in state-of-the-art models: the top-performing Claude Opus 4.5 achieves only 37.4% success. Further analysis shows that providing oracle human plans improves performance by 14–35 percentage points, pinpointing strategic reasoning as the primary bottleneck. Additionally, agents frequently fail to refuse infeasible tasks (best model achieves 53.9%), leading to unintended and potentially harmful side effects. Our findings underscore that current agents are not yet ready for autonomous enterprise deployment. More broadly, EnterpriseOps-Gym provides a concrete testbed to advance the robustness of agentic planning in professional workflows.

1 INTRODUCTION

LLMs today are most commonly deployed as conversational assistants, answering questions, drafting emails, and summarizing documents (OpenAI, 2025; Anthropic, 2025b). But a far more consequential capability is rapidly emerging, namely LLMs as *autonomous agents* that act on behalf of users (Xu et al., 2024). Consider an agent that, from a single instruction, searches the web for real-time inventory data, identifies the best product, and executes the purchase, all without further input. Recent advances in planning and tool use have made the vision of an *AI worker* increasingly plausible: autonomous agents handling professional workflows from software engineering (Jimenez et al., 2024) and data analysis (Drouin et al., 2024) to sales operations and enterprise administration (Huang et al., 2025a;b). Yet to function effectively in real-world professional deployments, such agents must do more than follow instructions. They must (1) maintain state coherently across long sequences of interleaved tool calls, (2) execute multi-step plans spanning dozens of actions, and (3) strictly adhere to the access control policies and procedural rules that govern the workplace.

These requirements are especially relevant within the enterprise domains. Here, agents do not merely retrieve information; they directly modify live databases, trigger downstream workflows, and affect

[†]Equal contribution. Correspondence to shivakrishnareddy.ma@servicenow.com

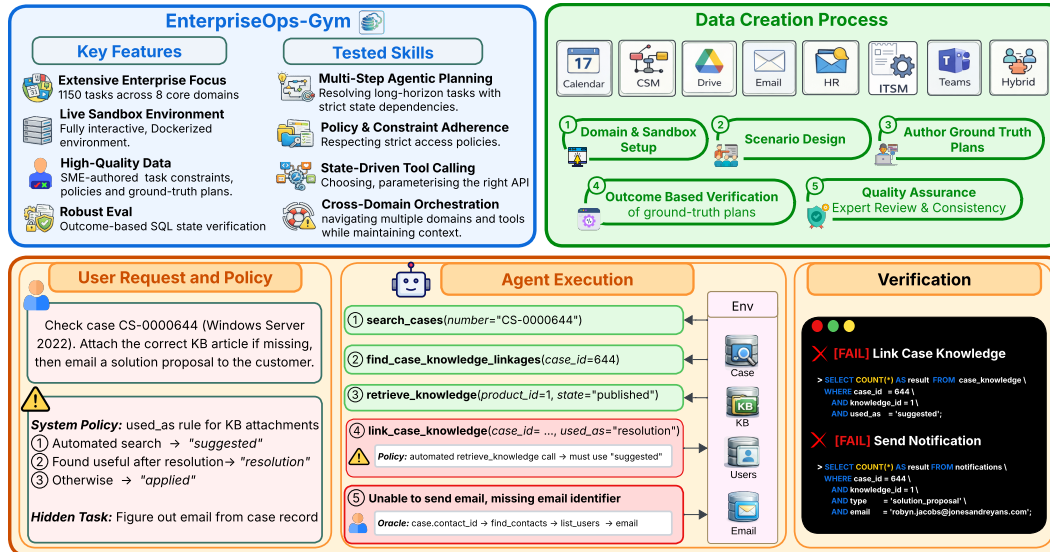


Figure 1: **Overview of EnterpriseOps-Gym: A benchmark for stateful agentic planning and tool use.** (Top-left) EnterpriseOps-Gym spans eight enterprise domains and evaluates multi-step agentic planning, policy adherence, state-driven tool calling and cross domain orchestration in a reproducible sandbox. (Top-right) Domain experts create sandbox and author realistic single- and cross-domain tasks, execute ground-truth trajectories, and write outcome-based verification logic with multi-stage quality assurance along with a human written oracle plan for completing the task. (Bottom) Given a task and system constraints, agents interact with the environment and execute tools. They are evaluated by final-state verifiers that check goal completion, policy compliance, and side effects.

real users. Actions are stateful and often irreversible, errors propagate silently across interconnected systems, and strict organizational policies constrain every step. Figure 1 shows one such example task, where the interplay of the user task and the nuanced system policy constraints demands more than surface-level instruction following. The agent must not only execute a multi-step workflow — searching a case, retrieving a KB article, and notifying the customer — but also respect a three-tiered policy governing how KB articles are tagged, and independently resolve a hidden piece of information (the customer’s email) by chaining through multiple tools. Equally important, agents must also know when to stop. Blindly attempting a policy-violating task corrupts system state, posing a direct safety risk in production environments.

Despite the urgency of this challenge, existing benchmarks fall short of capturing it. General tool-use evaluations (Li et al., 2023; Qin et al., 2024; Chen et al., 2025) treat tool calls as atomic and stateless, measuring accuracy on short sequences without state dependencies or cross-system coordination. Enterprise-focused benchmarks have emerged (Drouin et al., 2024; Boisvert et al., 2024; Huang et al., 2025a;b; Jha et al., 2025; Vishwakarma et al., 2025) to address challenges in professional environments. While they make important contributions, they are typically confined to a single vendor ecosystem (e.g., Salesforce or ServiceNow alone), with shallow environments of fewer than 25 database tables, under 50 tools, short task horizons, and limited policy constraints governing agent behavior (see Table 1).

To bridge this gap, we introduce **EnterpriseOps-Gym**, a benchmark designed to evaluate agentic planning within a high-fidelity enterprise simulation. EnterpriseOps-Gym spans eight interconnected ecosystems, ranging from general productivity tools like Calendar, Drive, Teams and Email to mission-critical business functions like Human Resource (HR), IT Service Management (ITSM), Customer Service Management (CSM), unified by a Hybrid category demanding coordinated cross-domain execution. The benchmark comprises 1,150 expert-authored tasks across these domains, including 30 infeasible scenarios that test whether agents correctly refuse unsatisfiable requests without leaving side effects on the system. Tasks are verified by hand-written SQL scripts that check goal completion, state integrity, policy compliance, and unintended side effects. EnterpriseOps-Gym runs on a fully interactive, containerized environment hosting 164 relational database tables and 512 functional tools,

with expert trajectories averaging 9 steps and reaching up to 34. Together, this scale is designed to stress-test the core challenges of enterprise automation: long-horizon planning, cross-system state management, and policy-constrained execution.

Our evaluation of 14 frontier models reveals a significant gap between current agentic capabilities and enterprise requirements. We find that performance is strongly shaped by domain complexity. Models fare best on collaboration tasks, with top models reaching 51–52% on Email, Teams, and Drive, but drop sharply on policy-governed domains such as ITSM (28.5%) and cross-domain Hybrid tasks (30.7%). These are precisely the domains where constraint-aware reasoning is unavoidable. The best overall model, Claude Opus 4.5, achieves only 37.4%, with open-source models lagging further behind. Infeasibility detection is a critical weak point, with even the best model refusing policy-violating tasks cleanly only 53.9% of the time. Test-time compute scaling helps but is not a universal remedy, with some workflows plateauing early or showing limited improvements regardless of thinking budget. Importantly, we identify strategic planning and not tool use as the primary bottleneck. Adding distractor tools has negligible impact, while providing human-authored plans improves models by 14–35 percentage points. More complex multi-agent orchestration does not close this gap either; decomposing tasks into subtasks can even regress performance due to strong sequential state dependencies. Together, these results underscore that current agents are not yet ready for autonomous enterprise deployment.

Overall, our contributions are as follows:

- We introduce EnterpriseOps-Gym, a benchmark of 1,150 expert-curated tasks across eight enterprise domains with outcome-based verification enforcing goal completion, state integrity, policy compliance, and side-effect checks, including 30 infeasible tasks designed to evaluate safe refusal behavior.
- We develop a fully interactive, containerized enterprise environment with 164 relational database tables and 512 functional tools, an order of magnitude more complex than prior enterprise benchmarks.
- We evaluate 14 frontier models, uncovering and analyzing systematic failure patterns across planning, state management, and policy compliance, and provide actionable insights to build more reliable enterprise agents. We also release EnterpriseOps-Gym to the community to advance research in stateful agentic planning and enterprise tool use.

2 RELATED WORKS

Benchmark	Focus	Num. Domains	Num. Tasks	Num. Tools	Avg. Steps	DB Tables	Avg. FK	Refusal Ability?	Human Task Curation?	Human Plans?
<i>General Tool Use</i>										
API-Bank (Li et al., 2023)	Tool-use	8	314	73	3	0	0	✗	✓	✗
ACEBench (Chen et al., 2025)	Tool-use	8	2000	4538	2	0	0	✓	✓	✗
τ -bench (Yao et al., 2024)	User Interaction	2	165	28	—	3	0.7	✓	✓	✗
τ^2 -bench (Barres et al., 2025)	User Interaction	3	279	56	—	9	—	✓	✗	✗
<i>Enterprise Specific</i>										
WorkArena (Drouin et al., 2024)	ServiceNow	7	33	30	10	7	0.9	✗	✗	✗
WorkArena++ (Boisvert et al., 2024)	ServiceNow	7	682 ([†] 341)	30	30-50	7	—	✓	✗	✗
ITBench (Jha et al., 2025)	IT	3	94	10	—	—	—	✗	✓	✗
WorkBench (Styles et al., 2024)	Workplace	5	690 ([†] 69)	26	2	5	0	✗	✓	✗
TheAgentCompany (Xu et al., 2024)	Startup	7	175	—	—	0	0	✗	✓	✗
CRMArena (Huang et al., 2025a)	Salesforce	1	1170 ([†] 9)	27	—	16	1.3	✗	✓	✗
CRMArena-Pro (Huang et al., 2025b)	Salesforce	3	8560 ([†] 19)	—	—	25	—	✓	✓	✗
EnterpriseBench (Vishwakarma et al., 2025)	Enterprise	5	500	46	3	17	1.2	✓	✗	✗
EnterpriseOps-Gym (Ours)	Enterprise	8	1150	512	9*	164	1.7	✓	✓	✓

Table 1: **Comparison with existing agentic benchmarks.** *DB Tables* reports the number of unique database tables in the environment; *Avg. FK* measures average foreign keys per table, indicating relational density and the complexity of inter-table dependencies agents must navigate. — denotes values not reported in the original work. *Avg. Steps reflects ideal human-authored execution trajectories; model trajectories may require significantly more steps. Human Plans refer to step-by-step natural language plans written by experts to complete the task. [†]Parenthetical values indicate the number of unique task templates.

LLM agent benchmarks broadly fall into three thematic groups: general-purpose tool-use and API evaluation, enterprise platform simulation, and agentic planning and computer-use. Table 1 summarizes how EnterpriseOps-Gym compares across key dimensions.

API and Tool-Use Benchmarks Early benchmarking efforts focused on testing agents’ ability to call APIs and chain tools in general, open-domain settings. ToolLLM (Qin et al., 2024) establishes large-scale evaluation across over 16,000 real-world web APIs, while API-Bank (Li et al., 2023) provides a smaller, runnable system for assessing planning and retrieval across 73 tools. ACEBench (Chen et al., 2025) scales this further with 4,538 tools and dynamic multi-turn evaluation, and τ -bench (Yao et al., 2024) and τ^2 -bench (Barres et al., 2025) introduce user simulation and refusal robustness as evaluation axes. These benchmarks make important contributions to measuring tool-calling accuracy but remain anchored to general web-oriented or open-domain APIs. They do not model the multi-system, policy-constrained, stateful tool ecosystems that characterize enterprise environments, which is the setting EnterpriseOps-Gym specifically targets.

Enterprise Benchmarks A second class of benchmarks targets enterprise platforms directly. Single-platform environments restrict scope to one vendor: WorkArena (Drouin et al., 2024) and WorkArena++ (Boisvert et al., 2024) evaluate compositional task completion within ServiceNow, CRM Arena (Huang et al., 2025a) and CRM Arena-Pro (Huang et al., 2025b) focus on Salesforce specific workflows, and ITBench (Jha et al., 2025) targets IT incident resolution. Multi-domain efforts broaden the scope but with different emphases: TheAgentCompany (Xu et al., 2024) simulates a startup software company requiring agents to interact via web browser, bash terminal, and code execution, a paradigm fundamentally different from the structured tool-calling API setting of EnterpriseOps-Gym; WorkBench (Styles et al., 2024) covers office productivity tools; and EnterpriseBench (Vishwakarma et al., 2025) evaluates function-calling on sandboxed enterprise data. Across these benchmarks, relational database complexity is consistently limited to fewer than 25 tables, and tasks are either confined to a single vendor ecosystem or lack expert-authored grounding or explicit evaluation of refusal behavior. EnterpriseOps-Gym addresses these limitations by spanning eight enterprise domains across both operational systems (CSM, ITSM, HR) and collaboration services (Email, Calendar, Teams, Drive), with 164 database tables with high connectivity, 512 tools, 1,150 expert-curated tasks including policy-constrained infeasible scenarios.

Agentic Planning and Computer-Use Planning is a core capability for autonomous agents, and several benchmarks have studied it across diverse general-purpose settings. UserBench (Qian et al., 2025) evaluates agents on multi-turn travel planning tasks where simulated users express preferences incrementally and implicitly, requiring proactive intent elicitation. Gaia2 (Froger et al., 2025) evaluates agents in an asynchronous simulated mobile environment, testing search, execution, temporal reasoning, adaptability, ambiguity handling, and multi-agent collaboration. VitaBench (He et al., 2025) evaluates agents on complex multi-turn tasks drawn from real-world services like food delivery, in-store dining, and online travel using a simulated user with dynamic preferences. A parallel line of work focuses on computer-use agents: OSWorld (Xie et al., 2024), WindowsAgentArena (Bonatti et al., 2025), WebArena (Zhou et al., 2024) and UI-Vision (Nayak et al., 2025) benchmark agents on desktop and web GUIs requiring complex sequential decision-making. While these settings demand sophisticated planning, they operate in everyday scenarios or computing environments rather than the policy-governed, multi-system contexts of enterprise environments. EnterpriseOps-Gym shares the emphasis on extended planning horizons but is uniquely positioned at the intersection of planning depth and enterprise domain fidelity.

3 ENTERPRISEOPS-GYM

In this section, we describe the design and construction of EnterpriseOps-Gym, covering domain selection, sandbox environment, task formulation in Section 3.1, and dataset statistics in Section 3.2.

3.1 ENTERPRISEOPS-GYM CONSTRUCTION

Selecting domains. We selected domains based on three principles, in consultation with SMEs who have hands-on experience in enterprise software: (i) relevance to real-world industry verticals, (ii)

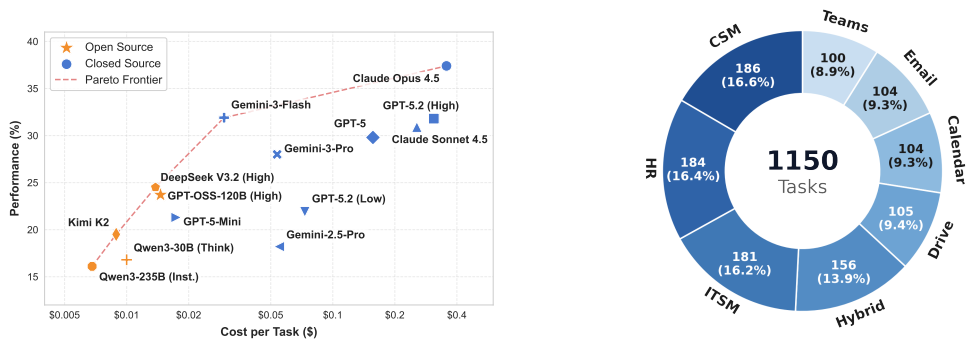


Figure 2: **(Left) Performance–cost tradeoff for agentic tool use on EnterpriseOps-Gym** We plot task success rate against estimated cost per task for both closed-source and open-source models. Open-source models incur lower cost but achieve consistently lower success rates. While higher-cost models offer modest performance gains, they remain far below reliable task completion. **(Right) Task distribution across eight EnterpriseOps-Gym domains.**

diversity in policy complexity and data sensitivity, and (iii) availability of domain experts to author and validate authentic tasks. This led us to two complementary groups of domains.

The first group—*Customer Service Management (CSM)*, *Human Resources (HR)*, and *Information Technology Service Management (ITSM)*—represents the operational backbone of enterprise organizations. These domains are present in virtually every industry vertical and are characterized by strict process compliance, access control policies, and high-stakes workflows where errors have tangible consequences. CSM involves managing the full lifecycle of support tickets and service agreements; HR handles sensitive employee data under strict privacy and procedural rules; and ITSM covers backend IT operations including incident management and system configuration. Their policy-heavy nature and realistic complexity make them ideal for stress-testing constraint-aware planning.

The second group—*Email*, *Calendar*, *Teams*, and *Drive*—encompasses the universal collaboration tools used daily across all enterprise organizations. Agents must be proficient in these to function as effective AI workers. While individually simpler, these domains require sophisticated orchestration: *Email* handles complex mailbox workflows, *Calendar* manages time and resource access policies, *Teams* administers collaborative workspaces, and *Drive* governs file system integrity and security.

Finally, *Hybrid* mandates cross-domain orchestration across these fragmented tools, requiring agents to maintain context and data integrity while switching between systems. Together, the eight domains span the full enterprise workflow spectrum, from internal operations to customer-facing processes to collaboration infrastructure, enabling evaluation of both specialized and general-purpose agentic capabilities. Refer to Appendix B for more details on each domain.

Sandbox Environment and Data Generation. We partnered with a professional data annotation firm (Turing*) to assemble a team of over 160 contributors, including software engineers for building environments and Subject Matter Experts (SMEs) for technical domains like CSM, ITSM, and HR (refer to Section A.1). To provide a reproducible and realistic evaluation ground, we developed a containerized docker sandbox hosting domain-specific databases, APIs, and the tool execution layer. This environment mirrors enterprise constraints without exposing proprietary infrastructure. For each domain, we created a realistic database schema and populated it with seed data and tools to access and manipulate data. The underlying data is designed with a strong real-world perspective by SMEs. Guided by official product documentation and SME insights, we model realistic table structures, constraints and tools based on industry standard database schemas. Starting from a fixed seed of data and tools, annotators extend the environment with new tables, schemas, and records as each task demands. Refer to Section A.4 for more details on the environment.

Task Construction Pipeline. The task creation process follows a rigorous pipeline designed to ensure high complexity and faithfulness to real-world workflows. In **Scenario Design**, annotators

*<https://www.turing.com>

craft challenging, multi-step scenarios based on specific complexity thresholds across dimensions including tool invocation counts, verification conditions, and state dependencies (action ordering is crucial to success), access constraints (e.g. “*only team owners can create private channels*”) and other policy conflicts (where user request conflicts with system policies). We constrain tasks to have a unique final state, though multiple valid paths may exist to reach it. In addition, we design a subset of 30 infeasible tasks where completion is intentionally impossible due to insufficient tool availability, explicit policy violations, or resource unavailability. For each task, annotators also update the sandbox environment with any additional necessary tables and tools. We follow this with **Ground Truth Execution and Plan Authoring**, where for each task, annotators provide a detailed step-by-step plan and manually execute it within the sandbox to capture a gold-standard trajectory, documenting each call with parameters, responses and execution rationale. Annotators also author a natural language reasoning plan that explicitly grounds each action in system constraints, user request, and available tools. Plans reference policies from the system prompt, explaining ordering dependencies thus fully grounding the task in the provided context. Refer to Section A.2 for more details.

We enforce **Outcome-based Verification**, where annotators author executable SQL verification scripts that check the final state of the environment upon task completion. This ensures that we evaluate agents on *outcomes* rather than rigid action sequences, allowing for alternative valid solution paths. The evaluation includes checks for required conditions (*is the goal achieved?*), integrity constraints (*are foreign key constraints respected?*), permission compliance (*did the agent avoid unauthorized actions?*) and other side effects. Finally, we conduct multiple rounds of **Quality Assurance**, where reviewer annotators (including authors) assess task feasibility given the initial state and the available tools, instruction clarity and completeness without external dependencies or domain knowledge, verification script correctness and coverage, as well as the fluency and coherence of the ground truth plan, instructions and verification scripts. Refer to Section A.3 for more details on verification.

3.2 DATASET STATISTICS

Task Statistics. EnterpriseOps-Gym evaluates agents across 1,150 tasks designed to mimic the depth of real-world enterprise operations, including 30 infeasible tasks that test whether agents correctly refuse unsatisfiable requests. The action space is extensive and diverse, comprising 512 unique tools across domains, with domain-specific toolsets ranging from 37 (Calendar) to 93 (ITSM). Expert human trajectories average 9.15 steps, with planning horizons varying considerably across domains, ranging from 6.2 steps on average in Email to 12.1 in CSM, and reaching up to 34 steps in HR (see Figure 4). Beyond length, our tasks are dense with constraints: on average, a task mandates satisfying 5.3 distinct verification conditions, with the most intricate scenarios requiring the resolution of 44 conditions.

Environment Statistics. Our sandbox environment models a highly interconnected data ecosystem comprising 164 unique database tables across the eight domains. On average, each task interacts with a sub-graph of 24.9 tables, reaching up to 73 in Hybrid scenarios. This means agents must reason over a large, partially-observable data graph to execute each task correctly. Each task operates over an average of 3,443 database rows, scaling to over 10,000 in data-heavy domains like CSM. To quantify relational complexity, we measure the average number of Foreign Keys (FK) per table. We observe a high degree of dependency, with average FKs ranging from 1.1 in Calendar to 2.4 in HR (mean \approx 1.7), exceeding the relational density of prior benchmarks (see Table 1). Higher FK density means agents must resolve more inter-table dependencies when constructing valid tool arguments, making referential integrity a key challenge. We provide more details in Section A.4.

4 EXPERIMENTS

4.1 BASELINES

We evaluate a diverse set of baselines covering closed-source frontier models, open-source reasoning and non reasoning models. All agents are evaluated under a unified interface with identical task instructions, tool definitions, sandbox environments, and evaluation protocols. Unless stated otherwise, agents operate in an *oracle-tool* setting where we assume a perfect retriever supplies the agent

Model	Teams	CSM	Email	ITSM	Calendar	HR	Drive	Hybrid	Infeas.	Avg.
<i>Closed Source Models</i>										
Claude Opus 4.5 (Anthropic, 2025b)	50.0	34.2	51.9	23.8	43.2	32.1	49.5	30.7	50.0	37.4
Gemini-3-Flash (Gemini Team, 2025)	47.3	35.0	44.3	28.5	30.5	12.6	49.7	24.2	38.5	31.9
GPT-5.2 (High) (OpenAI, 2025)	31.0	34.8	51.0	21.7	38.5	25.0	40.0	22.2	50.0	31.8
Claude Sonnet 4.5 (Anthropic, 2025b)	51.0	16.7	51.3	17.6	34.6	21.6	52.1	28.1	46.2	30.9
GPT-5 (OpenAI, 2025)	26.3	36.4	49.0	18.9	41.3	17.9	34.0	23.5	50.5	29.8
Gemini-3-Pro (Gemini Team, 2025)	43.0	27.7	33.6	22.2	28.8	12.5	46.7	22.9	50.0	28.0
GPT-5.2 (Low) (OpenAI, 2025)	25.0	21.2	43.3	6.7	28.9	13.0	26.7	20.9	53.9	21.9
GPT-5-Mini (OpenAI, 2025)	25.7	15.8	47.4	8.9	28.8	10.7	23.8	22.5	47.4	21.3
Gemini-2.5-Pro (Gemini Team, 2025)	39.3	11.6	31.1	13.9	12.5	4.9	27.0	19.6	34.7	18.2
<i>Open Source Models</i>										
DeepSeek-V3.2 (High) (DeepSeek-AI, 2024)	37.0	14.1	47.1	16.1	21.2	16.3	35.2	22.9	53.8	24.5
GPT-OSS-120B (High) (OpenAI et al., 2025)	32.0	16.3	42.3	6.1	35.6	16.3	41.0	19.6	50.0	23.7
DeepSeek-V3.2 (Medium) (DeepSeek-AI, 2024)	35.7	15.4	45.8	9.6	21.5	15.0	27.6	22.9	40.0	22.3
Kimi-K2-Thinking (K-Team, 2025)	30.0	7.1	51.0	12.2	15.4	8.2	39.6	15.7	30.5	19.5
Qwen3-30B (Think) (Yang et al., 2025)	22.0	5.4	51.9	6.7	18.3	7.6	25.7	15.7	36.8	16.8
Qwen3-235B (Inst.) (Yang et al., 2025)	28.0	4.7	38.1	9.3	15.7	7.8	23.8	17.7	30.5	16.1
Qwen3-4B (Think) (Yang et al., 2025)	24.0	3.8	38.4	5.6	5.8	7.1	21.9	15.8	31.6	13.7

Table 2: Overall task completion performance on EnterpriseOps-Gym. We report the percentage of tasks successfully completed by each model in oracle tool mode, broken down by domain. A task is considered successful only if all outcome verification checks pass.

with the right set of tools. This focuses the evaluation purely on planning and execution, without the need for explicit tool discovery. Additionally, we conduct ablations by increasing the number of available tools to analyze how tool set size impacts performance. We use a standard ReAct-style reasoning and tool-execution loop which has been shown to be effective in agentic settings Yao et al. (2022). The closed-source set includes Claude 4.5 (Anthropic, 2025b;a) variants (Opus and Sonnet), GPT (OpenAI, 2025) variants (5.2 High, 5.2 Low, 5, and 5-Mini), and Gemini (Gemini Team, 2025) variants (3-Pro, 3-Flash, and 2.5-Pro), while the open-source set includes Kimi-K2-Thinking (K-Team, 2025), DeepSeek-V3.2 (DeepSeek-AI, 2024), GPT-OSS-120B (Medium) (OpenAI et al., 2025), and Qwen3 (Yang et al., 2025) variants (235B Inst., 30B Think, and 4B Think).

4.2 EVALUATION METRICS

We evaluate models using pass@1 task completion rate, where a model receives a score of 1 only when it successfully completes all task requirements while satisfying all specified constraints. Task completion is verified by executing SQL-based verifiers hand-written by subject matter experts (SMEs) during the benchmark curation process. We report the average of pass@1 across three runs (to reduce variance) as our primary metric because it captures end-to-end task success. While we also measure verifier-level success rates (see Table 6), which provide fine-grained insight into the average number of successful verification checks, this metric can be misleading: agents may pass verifiers for trivial trajectory segments (e.g., initial setup steps) while failing on core task logic, system compliance requirements, or side-effect checks. Pass@1 therefore provides a more accurate assessment of real-world agent utility.

4.3 RESULTS

How do models perform across different domains? Overall, Claude Opus 4.5 achieves the best average task completion (37.4%) and is particularly strong across several workflows, leading on Email (51.9%), Calendar (43.2%), HR (32.1%), and Hybrid (30.7%). Gemini-3-Flash emerges as the second-best model overall (31.9%) and tops ITSM (28.5%), a service and operations management workflow. Claude Sonnet 4.5 (30.9%) remains strong on collaboration and document-centric workflows, leading on Teams (51.0%) and Drive (52.1%). GPT-5 shows more domain-specific peaks, topping CSM (36.4%). Open-source models still lag behind the closed-source systems overall. The strongest open-source model, DeepSeek-V3.2 (High), reaches a 24.5% average, narrowly ahead of GPT-OSS-120B (High) at 23.7%. They particularly struggle on service, policy, and people-facing domains such as CSM, ITSM, and HR. Qwen3-30B (Think) performs surprisingly well for its size, outperforming its larger instruct variant and attaining a highly competitive Email score (51.9%, tied for best). Finally, ITSM and Hybrid cross-domain workflows are the hardest settings (best: 28.5% and 30.7% respectively), highlighting that service operations and cross-domain coordination remain the key bottlenecks for all model families.

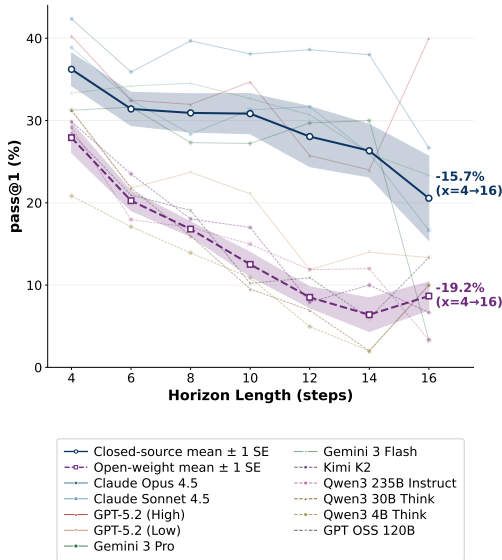
How does adding extra tools affect performance? We assessed robustness to tool overload by conducting ablations with Claude Sonnet 4.5, chosen for its strong overall performance and cost-effectiveness. We augmented the oracle toolset with extra *distractor* tools (+5, +10, and +15). To make this setting particularly challenging and representative of realistic retrieval errors, we asked Claude to retrieve the distractor tools that appeared most relevant to the task. Surprisingly, performance remained remarkably stable. The average completion rates actually increased slightly by an average of $\sim 1.0\%$ (+0.07% for +5, +2.35% for +10, and +0.64% for +15 tools). The only other notable variation was an average 4–9% increase in output tokens. This suggests that the model utilizes the additional token budget to carefully filter and select the appropriate tools. Such robustness likely stems from the extensive tool-use training inherent in these LLMs. Consequently, our findings indicate that the primary bottleneck for these agents is not tool discovery, but rather task planning and adherence to system policies.

Which model offers the best cost–performance tradeoff? Which model offers the best cost–performance trade-off depends on whether the priority is absolute quality or quality per dollar. As shown by the Pareto frontier in Figure 2, Gemini-3-Flash provides the strongest practical trade-off among closed-source models. It achieves 31.9% performance at 0.03 USD per task, delivering a higher success rate than more expensive models like GPT-5 (29.8% at 0.16 USD) and Claude Sonnet 4.5 (30.9% at 0.26 USD) at a fraction of the cost (80–90% less). Within the open-source ecosystem, DeepSeek-V3.2 (High) emerges as the Pareto-dominant option, achieving 24.5% performance at just 0.014 USD closely followed by GPT-OSS-120B (High) (23.7% at 0.015 USD), making these the best open-source value overall. Qwen3-235B (Inst.) remains the cheapest overall option (0.007 USD) but comes with a significant performance floor of 16.1%. Given that success rates across the board remain below 40%, these systems are not yet reliable enough for autonomous deployment without human oversight. For the highest absolute reliability, Claude Opus 4.5 remains the premier choice (37.4%), though it requires a steep premium of 0.36 USD per task.

How well do models refuse infeasible tasks? We curate 30 infeasible tasks across the 8 domains to evaluate whether models appropriately abstain from unsatisfiable requests. Each task has an average of 10 verification checks to ensure there are no side effects on the system. Tasks are impossible through three primary mechanisms: insufficient tool availability, explicit policy violations (e.g. scheduling conflicts, data access rules etc) and resource unavailability (e.g. inactive users, system in migration mode etc). Moreover, tasks employ compound constraints averaging 3 to 4 per request, necessitating that models evaluate multiple intersecting conditions to identify task feasibility. As seen in Table 2, GPT-5.2 (Low) and DeepSeek-V3.2 (High) perform the best (53.9% and 53.8% respectively) in abstaining from the task while leaving no side effects on the system, followed closely by GPT-5 (50.5%) and a cluster of models including Claude Opus 4.5, Gemini-3-Pro, and GPT-OSS-120B (High) at 50.0%. However, the absolute scores of all the models remain well below safe applicability for production systems.

How does performance scale with task horizon? To understand how model capabilities degrade with task complexity, we stratify tasks by their expected horizon length, which is proportional to the number of tool execution steps in the human created execution trajectories. As illustrated in Figure 3, performance across all models exhibits a consistent decay as the task horizon increases, reflecting the cumulative difficulty of maintaining reasoning integrity over multi-step sequences. The closed-source group, led by Claude Opus 4.5, demonstrates greater resilience, maintaining a performance lead even as the group mean drops from approximately 35% at 4 steps to under 20% by step 16. In contrast, the open-source cohort shows a much steeper decline, with models like Kimi K2 and GPT OSS 120B converging toward a success rate near 10% at the maximum horizon. This near-universal trend suggests that while current models can navigate short-to-medium sequences, the rapid accumulation of errors in long-horizon tasks remains a critical barrier to autonomous reliability in production environments.

How does thinking budget affect performance? We evaluated the impact of test-time compute by varying the thinking budget (*low*, *medium*, *high*) for the GPT-OSS-120B model (OpenAI et al., 2025). Increasing the thinking budget yields substantial improvements in task completion across almost all domains (see Figure 6). Operating at a *low* thinking budget causes the model to struggle severely, achieving near-zero accuracy on complex service and people-facing domains such as CSM



Model	Plan	CSM	ITSM	HR
Kimi-K2	CP	19.6 ↑12.5%	18.1 ↑5.9%	17.2 ↑9.0%
	HP	42.2 ↑35.1%	29.1 ↑16.9%	34.5 ↑26.3%
Qwen3-30B	CP	15.2 ↑9.8%	11.7 ↑5.0%	17.9 ↑10.3%
	HP	33.9 ↑28.5%	20.9 ↑14.2%	33.2 ↑25.6%
Qwen3-4B	CP	16.8 ↑13.0%	12.2 ↑6.6%	19.6 ↑12.5%
	HP	37.2 ↑33.4%	23.3 ↑17.7%	36.4 ↑29.3%

Figure 3: (left) **Performance degrades consistently with planning horizon.** Pass@1 accuracy for closed-source (solid) and open-weight (dashed) models across horizon lengths 4–16. Thick lines show the group mean ± 1 SE. We observe monotonic degradation of performance for both sets, while open model performance falls more sharply with horizon length. (right) Plan-Conditioned Execution Baseline. Comparison of performance with Claude Plans (CP) vs. Human Plans (HP). Green values indicate % improvement over the ReAct baseline from Table 2.

(1.1%), ITSM (1.1%), and HR (0.0%). Scaling to a *high* budget unlocks significant capabilities, driving dramatic absolute gains in Drive (8.6 \rightarrow 41.0%), Calendar (8.7 \rightarrow 35.6%), and Teams (4.0 \rightarrow 32.0%). This strong dependence on test-time compute underscores that EnterpriseOps-Gym tasks require complex reasoning and planning to execute. However, we also observe that performance scaling is not universally monotonic; for instance, performance on Email peaks at the *medium* budget (45.2%) before receding slightly, and ITSM plateaus early (1.1 \rightarrow 6.1% \rightarrow 6.1%). This suggests that simply allocating more thinking tokens cannot universally overcome fundamental capability bottlenecks in certain workflows.

4.4 FURTHER ANALYSIS

The results above show that performance degrades sharply with horizon length and that models struggle most with planning rather than execution. We run a series of controlled ablations that both decouple planning from execution and build increasingly complex multi-agent systems to test whether distributing cognitive load across specialized agents can recover performance that a single monolithic agent cannot achieve; full experimental details are in appendix Section D.1.

Planning quality is the primary bottleneck. A planner-executor configuration in which a dedicated Claude Sonnet 4.5 planner generates a high-level plan and a separate executor carries out tool execution yields consistent gains of 6–13% across models and domains (Table 3). Providing human-authored plans to the same executors yields gains of 14–35 percentage points—roughly double those from automated planning—indicating that current LLMs fall well short of human-level strategic reasoning on these tasks. Notably, Qwen3-4B with human-authored (or Claude-generated) plans is competitive with, and sometimes outperforms, larger models. This suggests that once strategic reasoning is externalized, the main remaining challenges are instruction-following and precise tool use—capabilities modern LLMs handle well across scales.

Architectural complexity does not close this gap. Among Claude Sonnet 4.5 multi-agent configurations, the *Planner+Executor* system consistently outperforms the ReAct baseline (+10.7% CSM, +8.8% HR), but the *Planner+Decompose+Subtask Executor* system is less robust: while it provides a minor lift in ITSM, it regresses in both CSM and HR, even falling below the base ReAct performance

in CSM (16.2% vs. 16.7%), as shown in Figure 5. This is consistent with EnterpriseOps-Gym tasks having strong sequential state dependencies that decomposition disrupts. The substantial remaining gap between all automated systems and ReAct with human plans indicates that progress requires advances in constraint-aware plan generation rather than architectural complexity alone.

Models struggle most with policy and process compliance. Manual qualitative analysis of partial-progress failures reveals four recurring patterns: *Missing Prerequisite Lookup*, *Cascading State Propagation*, *Incorrect ID Resolution*, and *Premature Completion Hallucination*. Verifier-level analysis (Table 7) shows that models struggle most with *Permission and Process Compliance*—a particularly critical gap for real-world deployment, where policy violations can cause cascading system failures and introduce serious security vulnerabilities. Annotated failure examples and extended failure mode descriptions are in Appendix C.

5 DISCUSSION AND CONCLUSION

We introduced EnterpriseOps-Gym, a benchmark and sandboxed evaluation platform spanning 1,150 expert-curated tasks across eight enterprise productivity domains, with 512 tools and SQL-based verifiers authored by subject matter experts. Our experiments surface several findings that we believe have broad implications for the development of enterprise-grade LLM agents.

Current agents are far from enterprise-ready. Even under oracle tool access—the most favorable possible retrieval setting—the best model achieves only 37.4% task success, and performance degrades monotonically with horizon length and cross-domain coupling. Our qualitative and quantitative analysis shows that agents struggle most with *Permission and Process Compliance* (e.g., policy adherence, cascading state transitions) rather than with basic task completion. Furthermore, even frontier models refuse infeasible tasks reliably only about half the time (best: 53.9%), falling well short of the robustness required for unsupervised deployment. These results indicate that the gap to enterprise reliability will not be closed by scaling model capacity alone.

Planning is the dominant bottleneck, not tool execution. Our plan-conditioned ablations demonstrate that human-authored plans yield 14–35 percentage point gains across models and domains which is far larger than gains from automated planning or more complex multi-agent orchestration. Strikingly, Qwen3-4B conditioned on human plans is competitive with much larger models under the same condition, suggesting that once strategic reasoning is externalized, even small models can execute faithfully. This dissociation between planning and execution ability implies that the core challenge is constraint-aware plan generation, not tool invocation proficiency. Consistent with this, distractor tools do not meaningfully hurt performance further confirming that tool retrieval is not the binding constraint. Advances in long-horizon, policy-aware planning are therefore the highest-leverage direction for improving agent performance on EnterpriseOps-Gym.

Thinking budget matters, but has domain-specific ceilings. Increasing test-time compute yields substantial gains in most domains, but scaling is not universally monotonic: some domains plateau early, suggesting that additional reasoning tokens cannot compensate for fundamental gaps in domain knowledge or policy understanding. Future work should investigate how to allocate test-time compute more adaptively, and whether targeted training on constraint-heavy domains can raise these ceilings.

Future directions. Our results motivate three concrete research priorities. First, *constraint-aware plan generation*: methods that explicitly reason over policy constraints, side-effect dependencies, and prerequisite structures before committing to action sequences. Second, *long-horizon state management*: mechanisms for maintaining coherent world state over many tool calls, such as episodic memory or structured state representations, to prevent the error accumulation we observe with increasing horizon length. Third, *safe refusal and escalation*: agents must reliably detect infeasible or policy-violating requests and abstain cleanly, a capability that remains weak across all evaluated systems today.

We will release EnterpriseOps-Gym, its sandbox environment, and evaluation tooling to support open, community-driven research. The sandbox is modular and extensible, allowing new domains, tools, and workflows to be added as enterprise practices evolve. By grounding agent evaluation in realistic, constraint-rich enterprise workflows, EnterpriseOps-Gym aims to shift the field’s focus toward the planning, safety, and policy-compliance capabilities that truly determine whether an LLM agent is deployable as a reliable *AI worker*.

REFERENCES

- Anthropic. Introducing claude opus 4.5, November 2025a. URL <https://www.anthropic.com/news/claude-opus-4-5>.
- Anthropic. Introducing claude sonnet 4.5, September 2025b. URL <https://www.anthropic.com/news/claude-sonnet-4-5>.
- Anthropic. Introducing claude opus 4.5, 2025c.
- Victor Barres, Honghua Dong, Soham Ray, Xujie Si, and Karthik Narasimhan. τ^2 -bench: Evaluating conversational agents in a dual-control environment, 2025. URL <https://arxiv.org/abs/2506.07982>.
- Léo Boisvert, Megh Thakkar, Maxime Gasse, Massimo Caccia, Thibault Le Sellier De Chezelles, Quentin Cappart, Nicolas Chapados, Alexandre Lacoste, and Alexandre Drouin. Workarena++: Towards compositional planning and reasoning-based common knowledge work tasks, 2024. URL <https://arxiv.org/abs/2407.05291>.
- Rogério Bonatti, Dan Zhao, Francesco Bonacci, Dillon Dupont, Sara Abdali, Yinheng Li, Yadong Lu, Justin Wagle, Kazuhito Koishida, Arthur Bucker, Lawrence Keunho Jang, and Zheng Hui. Windows agent arena: Evaluating multi-modal OS agents at scale. In *Forty-second International Conference on Machine Learning*, 2025. URL <https://openreview.net/forum?id=W9s817KqYf>.
- Chen Chen, Xinlong Hao, Weiwen Liu, Xu Huang, Xingshan Zeng, Shuai Yu, Dexun Li, Shuai Wang, Weinan Gan, Yuefeng Huang, Wulong Liu, Xinzhi Wang, Defu Lian, Baoqun Yin, Yasheng Wang, and Wu Liu. Acebench: Who wins the match point in tool usage?, 2025. URL <https://arxiv.org/abs/2501.12851>.
- DeepSeek-AI. Deepseek-v3 technical report, 2024. URL <https://arxiv.org/abs/2412.19437>.
- Alexandre Drouin, Maxime Gasse, Massimo Caccia, Issam H. Laradji, Manuel Del Verme, Tom Marty, Léo Boisvert, Megh Thakkar, Quentin Cappart, David Vazquez, Nicolas Chapados, and Alexandre Lacoste. Workarena: How capable are web agents at solving common knowledge work tasks?, 2024.
- Romain Froger, Pierre Andrews, Matteo Bettini, Amar Budhiraja, Ricardo Silveira Cabral, Virginie Do, Emilien Garreau, Jean-Baptiste Gaya, Hugo Laurençon, Maxime Lecanu, Kunal Malkan, Dheeraj Mekala, Pierre Ménard, Gerard Moreno-Torres Bertran, Ulyana Piterbarg, Mikhail Plekhanov, Mathieu Rita, Andrey Rusakov, Vladislav Vorotilov, Mengjue Wang, Ian Yu, Amine Benhalloum, Grégoire Mialon, and Thomas Scialom. Are: Scaling up agent environments and evaluations, 2025. URL <https://arxiv.org/abs/2509.17158>.
- Gemini Team. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities, 2025. URL <https://arxiv.org/abs/2507.06261>.
- Wei He, Yueqing Sun, Hongyan Hao, Xueyuan Hao, Zhikang Xia, Qi Gu, Chengcheng Han, Dengchang Zhao, Hui Su, Kefeng Zhang, Man Gao, Xi Su, Xiaodong Cai, Xunliang Cai, Yu Yang, and Yunke Zhao. Vitabench: Benchmarking llm agents with versatile interactive tasks in real-world applications, 2025. URL <https://arxiv.org/abs/2509.26490>.
- Kung-Hsiang Huang, Akshara Prabhakar, Sidharth Dhawan, Yixin Mao, Huan Wang, Silvio Savarese, Caiming Xiong, Philippe Laban, and Chien-Sheng Wu. Crmarena: Understanding the capacity of llm agents to perform professional crm tasks in realistic environments. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 2025a.
- Kung-Hsiang Huang, Akshara Prabhakar, Onkar Thorat, Divyansh Agarwal, Prafulla Kumar Choubey, Yixin Mao, Silvio Savarese, Caiming Xiong, and Chien-Sheng Wu. Crmarena-pro: Holistic assessment of llm agents across diverse business scenarios and interactions. *arXiv preprint arXiv:2505.18878*, 2025b.

- Saurabh Jha, Rohan Arora, Yuji Watanabe, Takumi Yanagawa, Yinfang Chen, Jackson Clark, Bhavya Bhavya, Mudit Verma, Harshit Kumar, Hirokuni Kitahara, Noah Zheutlin, Saki Takano, Divya Pathak, Felix George, Xinbo Wu, Bekir O. Turkkan, Gerard Vanloo, Michael Nidd, Ting Dai, Oishik Chatterjee, Pranjal Gupta, Suranjana Samanta, Pooja Aggarwal, Rong Lee, Pavankumar Murali, Jae wook Ahn, Debanjana Kar, Ameet Rahane, Carlos Fonseca, Amit Paradkar, Yu Deng, Pratibha Moogi, Prateeti Mohapatra, Naoki Abe, Chandrasekhar Narayanaswami, Tianyin Xu, Lav R. Varshney, Ruchi Mahindru, Anca Sailer, Laura Shwartz, Daby Sow, Nicholas C. M. Fuller, and Ruchir Puri. Itbench: Evaluating ai agents across diverse real-world it automation tasks, 2025. URL <https://arxiv.org/abs/2502.05352>.
- Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik R Narasimhan. SWE-bench: Can language models resolve real-world github issues? In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=VTF8yNQ66>.
- K-Team. Kimi k2: Open agentic intelligence, 2025. URL <https://arxiv.org/abs/2507.20534>.
- Minghao Li, Yingxiu Zhao, Bowen Yu, Feifan Song, Hangyu Li, Haiyang Yu, Zhoujun Li, Fei Huang, and Yongbin Li. API-bank: A comprehensive benchmark for tool-augmented LLMs. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 3102–3116, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.187. URL <https://aclanthology.org/2023.emnlp-main.187/>.
- Shravan Nayak, Xiangru Jian, Kevin Qinghong Lin, Juan A. Rodriguez, Montek Kalsi, Rabiul Awal, Nicolas Chapados, M. Tamer Özsu, Aishwarya Agrawal, David Vazquez, Christopher Pal, Perouz Taslakian, Spandana Gella, and Sai Rajeswar. Ui-vision: A desktop-centric gui benchmark for visual perception and interaction, 2025. URL <https://arxiv.org/abs/2503.15661>.
- OpenAI. Introducing gpt-5, August 2025. URL <https://openai.com/index/introducing-gpt-5/>.
- OpenAI, :, Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K. Arora, Yu Bai, Bowen Baker, Haiming Bao, Boaz Barak, Ally Bennett, Tyler Bertao, Nivedita Brett, Eugene Brevdo, Greg Brockman, Sebastien Bubeck, Che Chang, Kai Chen, Mark Chen, Enoch Cheung, Aidan Clark, Dan Cook, Marat Dukhan, Casey Dvorak, Kevin Fives, Vlad Fomenko, Timur Garipov, Kristian Georgiev, Mia Glaese, Tarun Gogineni, Adam Goucher, Lukas Gross, Katia Gil Guzman, John Hallman, Jackie Hehir, Johannes Heidecke, Alec Helyar, Haitang Hu, Romain Huet, Jacob Huh, Saachi Jain, Zach Johnson, Chris Koch, Irina Kofman, Dominik Kundel, Jason Kwon, Volodymyr Kyrylov, Elaine Ya Le, Guillaume Leclerc, James Park Lennon, Scott Lessans, Mario Lezcano-Casado, Yuanzhi Li, Zhuohan Li, Ji Lin, Jordan Liss, Lily, Liu, Jiancheng Liu, Kevin Lu, Chris Lu, Zoran Martinovic, Lindsay McCallum, Josh McGrath, Scott McKinney, Aidan McLaughlin, Song Mei, Steve Mostovoy, Tong Mu, Gideon Myles, Alexander Neitz, Alex Nichol, Jakub Pachocki, Alex Paino, Dana Palmie, Ashley Pantuliano, Giambattista Parascandolo, Jongsoo Park, Leher Pathak, Carolina Paz, Ludovic Peran, Dmitry Pimenov, Michelle Pokrass, Elizabeth Proehl, Huida Qiu, Gaby Raila, Filippo Raso, Hongyu Ren, Kimmy Richardson, David Robinson, Bob Rotsted, Hadi Salman, Suvansh Sanjeev, Max Schwarzer, D. Sculley, Harshit Sikchi, Kendal Simon, Karan Singhal, Yang Song, Dane Stuckey, Zhiqing Sun, Philippe Tillet, Sam Toizer, Foivos Tsimpourlas, Nikhil Vyas, Eric Wallace, Xin Wang, Miles Wang, Olivia Watkins, Kevin Weil, Amy Wendling, Kevin Whinnery, Cedric Whitney, Hannah Wong, Lin Yang, Yu Yang, Michihiro Yasunaga, Kristen Ying, Wojciech Zaremba, Wenting Zhan, Cyril Zhang, Brian Zhang, Eddie Zhang, and Shengjia Zhao. gpt-oss-120b & gpt-oss-20b model card, 2025. URL <https://arxiv.org/abs/2508.10925>.
- Cheng Qian, Zuxin Liu, Akshara Prabhakar, Zhiwei Liu, Jianguo Zhang, Haolin Chen, Heng Ji, Weiran Yao, Shelby Heinecke, Silvio Savarese, Caiming Xiong, and Huan Wang. Userbench: An interactive gym environment for user-centric agents, 2025. URL <https://arxiv.org/abs/2507.22034>.
- Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, Sihan Zhao, Lauren Hong, Runchu Tian, Ruobing Xie, Jie Zhou, Mark Gerstein,

- dahai li, Zhiyuan Liu, and Maosong Sun. ToolLLM: Facilitating large language models to master 16000+ real-world APIs. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=dHng200Jjr>.
- Olly Styles, Sam Miller, Patricio Cerda-Mardini, Tanaya Guha, Victor Sanchez, and Bertie Vidgen. Workbench: a benchmark dataset for agents in a realistic workplace setting, 2024. URL <https://arxiv.org/abs/2405.00823>.
- Harsh Vishwakarma, Ankush Agarwal, Ojas Patil, Chaitanya Devaguptapu, and Mahesh Chandran. Can llms help you at work? a sandbox for evaluating llm agents in enterprise environments, 2025. URL <https://arxiv.org/abs/2510.27287>.
- Tianbao Xie et al. Osworld: Benchmarking multimodal agents for open-ended tasks in real computer environments. *arXiv preprint arXiv:2404.07972*, 2024.
- Frank F. Xu, Yufan Song, Boxuan Li, Yuxuan Tang, Kritanjali Jain, Mengxue Bao, Zora Z. Wang, Xuhui Zhou, Zhitong Guo, Murong Cao, Mingyang Yang, Hao Yang Lu, Amaad Martin, Zhe Su, Leander Maben, Raj Mehta, Wayne Chi, Lawrence Jang, Yiqing Xie, Shuyan Zhou, and Graham Neubig. Theagentcompany: Benchmarking llm agents on consequential real world tasks, 2024. URL <https://arxiv.org/abs/2412.14161>.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Rui Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*, 2022.
- Shunyu Yao, Noah Shinn, Pedram Razavi, and Karthik Narasimhan. τ -bench: A benchmark for tool-agent-user interaction in real-world domains, 2024. URL <https://arxiv.org/abs/2406.12045>.
- Shuyan Zhou, Frank F. Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou, Yonatan Bisk, Daniel Fried, Uri Alon, and Graham Neubig. Webarena: A realistic web environment for building autonomous agents, 2024. URL <https://arxiv.org/abs/2307.13854>.

APPENDIX

Table of Contents

	Page
A. Data Collection and Human Annotation	15
A.1 Demographics and Recruitment	15
A.2 Annotation Process	15
A.3 Quality Assurance and Verification	15
A.4 Sandbox Environment	15
B. Task Categories and Examples	16
B.1 Calendar	16
B.2 Customer Service Management (CSM)	17
B.3 Drive	20
B.4 Email	22
B.5 Human Resources (HR)	23
B.6 IT Service Management (ITSM)	25
B.7 Teams	25
B.8 Hybrid	27
C. Rollout examples	27
D. Additional Analysis and Results	32
E. Impact Statement	36

A DATA COLLECTION AND HUMAN ANNOTATION

A.1 DEMOGRAPHICS AND RECRUITMENT

We partnered with a professional data annotation vendor specializing in data curation for AI applications. The annotation team was structured as a multi-tiered workforce consisting of annotators, quality assurance reviewers, and team leads. The contributors were distributed around major geo locations including Asia, North America, Latin America and Africa, with an age range of 22-37 years. All data contributors held bachelor’s degrees in Engineering, Computer Science, or related disciplines and possessed prior experience in data labeling and UI research.

Recruitment selection prioritized strong proficiency in technical writing, English, and computer science fundamentals, along with expertise in prompt engineering. To ensure high domain realism, the team also included Subject Matter Experts (SMEs) for technical domains such as ITSM and CSM, as well as software engineers responsible for building the sandbox environments. On average, each domain was supported by 20 annotators and 6 reviewers, totaling over 160 contributors. The data collection campaign spanned approximately four months, beginning with a one-month pilot phase. During this pilot period, we collaborated closely with the vendor’s team to conduct detailed reviews and provide extensive feedback, enabling contributors to refine their understanding of the task requirements. All contributors were fairly compensated. The creation of each task, including scenario design, verification, and quality assurance, cost approximately 100 USD.

A.2 ANNOTATION PROCESS

The data generation pipeline began with contributors being assigned to specific domains and taxonomies. They were supported by a simulated environment that included domain-specific databases and a set of available functions. Annotators and reviewers utilized an internally developed tool to streamline the process.

To ensure a diverse range of difficulty, contributors were given specific complexity thresholds based on the number of required tools and verification steps used to complete a task. A higher number of tools and verifiers directly correlated with higher task complexity. Annotators leveraged their domain-specific expertise to design complex scenarios and problems within their assigned taxonomies, while internal tooling captured their action trajectories.

A.3 QUALITY ASSURANCE AND VERIFICATION

Verification was rigorous, multi-layered, and designed to retain only the most challenging tasks. Upon completion of a trajectory and its verification scripts, we employed a preliminary filtration and verification stage using state-of-the-art LLMs, specifically GPT-5, Gemini, and Claude. We executed draft tasks against these models and analyzed the resulting trajectories to identify failure modes such as incorrect task definitions leading to unintended paths, missing tools, invalid database entries, or access control conflicts. Additionally, this automated stress-testing flagged overly simple reasoning paths and issues with groundedness. These insights enabled annotators to refine task definitions, database states, and tool availability, while simultaneously discarding trivial tasks. This iterative process naturally drove the creation of more well-defined and complex scenarios. Following this automated phase, tasks underwent human verification, where reviewers evaluated each entry for grammatical accuracy, tool usage logic, natural language fluency, and execution correctness. Detailed quality rubrics were employed to standardize assessments of trajectory quality, prompt clarity, and verifier robustness. This rigorous pipeline ensured that only high-quality, complex tasks were retained in the final benchmark.

A.4 SANDBOX ENVIRONMENT

The seed data for EnterpriseOps-Gym is fully synthetic but generated with a strong real-world perspective with the help of domain SMEs. For each domain, we studied publicly available official API documentation, data models, and usage examples from relevant enterprise systems to understand:

- Typical entity structures and relationships

- Field semantics and data constraints
- Expected API behavior and response patterns

Guided by this research and the domain expertise of our SMEs, we modeled realistic behavior and constraints for each table and field. Our primary objective was to ensure structural realism and behavioral fidelity while remaining platform-agnostic, avoiding reliance on any specific vendor’s proprietary dataset. Furthermore, engineers and SMEs conducted rigorous testing to verify database consistency, ensuring the absence of missing elements or logical contradictions.

The seed data varies significantly between database files for every task. While certain individual field values (e.g., names or email patterns) may occasionally repeat, each dataset represents a distinct high-level use case. The overall data composition, relationships, and scenarios are intentionally unique rather than simple surgical variations of the same seed. Furthermore, as annotators vetted and designed new tasks, the databases were dynamically expanded with additional tables and entries to meet evolving scenario requirements. This iterative enrichment yielded a highly complex data ecosystem comprising 164 unique tables with a dense interconnectivity (mean foreign key degree of 1.7), ensuring a rich and realistic state space for agentic planning.

Environment Setup. We provide the complete evaluation environment as a fully containerized Docker setup, hosting both the domain-specific database infrastructure and the tool execution layer. To ensure reproducibility and isolation, a fresh database instance is initialized for each task run, preventing state leakage or side effects from prior executions. The containerized architecture abstracts the complexity of tool invocation and response handling, providing a consistent interface for agents. This design also simplifies verification, as the environment state can be deterministically queried. Given the rich schema of 164 interconnected tables, the environment is highly extensible; researchers can straightforwardly define new tasks that reflect complex enterprise use cases. We will release these containers, along with a comprehensive guide on adding new tasks and verification scripts, to facilitate further research.

B TASK CATEGORIES AND COMPLETE EXAMPLES

This section details the eight distinct domains within the benchmark, outlining their unique operational environments, some required reasoning skills, and task examples.

B.1 CALENDAR

The **Calendar** category represents a complex administrative environment where the agent manages time, access, and resources with high precision. It goes beyond simple meeting bookings, requiring the agent to act as a system administrator who restructures entire calendar ecosystems and enforces compliance policies. This involves granular management of access control lists for users and groups, setting up programmatic listeners for security audits, and handling metadata for resources like conference rooms.

The reasoning required for these tasks is often sequential and conditional, where the agent must evaluate complex predicates before acting. For instance, it might need to check if a calendar exists or verify an attendee’s status before making changes. This demands strong temporal reasoning to handle timezone conversions and recurrence rules, as well as the ability to enforce policies by translating high-level intents into low-level actions. Agents must handle instructions like *“If the event duration is 60 minutes, set color to Green; otherwise, set to Red”* or fuzzy requests such as finding *“events containing ‘sprint’ in the title.”*

Calendar Example Task

System Prompt:

You are a Google Calendar automation agent with full administrative permissions to manage users, meetings, recordings, and messages. Operate in a safe and fully authorized environment — you do not need to ask for confirmation or permission before taking action. When identifiers such as names or IDs are missing, perform exactly one lookup per entity type,

verify that you are reusing correct values from previous responses, and proceed using the retrieved data. Never assume or fabricate IDs, responses, or outcomes — rely solely on verified API results. Complete each task in a single, logical, and efficient execution flow.

User Prompt: I need to schedule a series of cross-functional planning meetings for our product launch. Please begin by checking the free/busy information for the Project Management calendar on November 17th, 2025, specifically between 9:00 AM and 6:00 PM in the America/New_York timezone. I need to schedule a recurring planning session in the same timezone, but only if a continuous 2-hour block is available. Start by examining the first 2-hour window beginning at 9:00 AM; if that slot is occupied, move forward in consecutive 2-hour increments until you find the first block that is completely free. Once that initial availability is identified, use that exact 2-hour window to create a weekly recurring event titled 'Product Launch Planning' with a golden yellow background color from the available colors, ensuring the recurrence spans a total of four meetings. Include Alice (alice.manager@techcorp.com), Bob (bob.smith@techcorp.com), Carol (carol.white@techcorp.com), and Dave (dave.brown@techcorp.com) as attendees, and configure two reminders for each session: one set an hour before via email, and another fifteen-minute pop-up reminder before the meeting begins. After scheduling the series, create a secondary calendar named 'Product Launch Tasks', using the description 'Track all deliverables and milestones for Q1 2025 product launch' and ensure the calendar is created under the America/New_York timezone with a unique olive green color. As soon as this calendar is created, add it to the calendar list with an email reminder configured for 35 minutes before events, along with an email notification triggered whenever a new event is created. Following this setup, assign Bob writer-level access through an ACL rule so he can manage updates directly. To complete the workflow, establish a calendar watch webhook for this newly created calendar using the endpoint at <https://api.techcorp.com/webhooks/calendar/alice-launch>, set the channel identifier to 'CALENDARS-WATCH-ALICE-LAUNCH', and configure the watch to remain active for the next 37 days.

Oracle Tools: `get_calendar_list, query_freebusy, get_colors, create_event, create_calendar, add_calendar_to_list, insert_acl_rule, watch_events`

B.2 CUSTOMER SERVICE MANAGEMENT (CSM)

This category simulates the high-stakes workflow of a B2B technical support center, where the agent acts as a Technical Support Operations Specialist. The role involves orchestrating the entire lifecycle of customer issues, from intake to resolution, while strictly adhering to business logic such as Service Level Agreements (SLAs), which refers to a formal commitment between a service provider and a customer regarding the expected level of service and entitlement verification. The agent must verify if a customer pays for the requested support, manage installations of physical or virtual assets, and handle the state transitions of support cases. Tasks range from *“Register these 5 new servers with these serials... and add the 'Gold Support' package”* to handling SLA breaches where an agent must *“Escalate to 'Critical', assign to the Escalation Manager, and draft an apology email.”*

Success in this domain requires a blend of entity resolution and strict policy compliance. Agents must identify the correct assets from vague descriptions (e.g., *“The server ending in 9300”*) and strictly follow business rules, such as prioritizing VIP accounts regardless of minor issue severity. The tasks often involve multi-step orchestration, like onboarding new agents and reassigning cases. For example, *“Transfer her high-priority cases to Bob, her low-priority ones to the Queue, and deactivate her profile”*. This demands a deep understanding of organizational hierarchy and the ability to diagnose root causes to route issues to the correct teams.

CSM Example Task

System Prompt:
CSM Agent Policy:

You are a Customer Service Management assistant. Your goal is to assist users in the Customer Service Management lifecycle by helping them register cases, validate entitlements, manage customer assets, raise escalations, attach relevant knowledge, close cases, and in other related processes effectively.

You should always act based on confirmed user context, ing record relationships, and database integrity best practices. Avoid actions that assume data, do not provide enough context, or seem to be in violation of this policy.

General Instructions

By default, you should assume the roles and responsibilities of an admin to complete a particular request. If a user request violates policy, do not act on it. Perform one operation at a time. Do not provide knowledge or procedures not in the system. Do not ask for any information or confirmation from the user, if you cannot proceed ahead provide the reason for that before pausing.

Roles & Responsibilities

Administrator: Has full access across all tables. Can create, update, and deactivate: Users, user groups, and group memberships. Accounts, contacts, locations, and entitlements. Products, installed products, SLAs, contracts, and knowledge. Manages assignment rules, workflows, and escalations.

Agent: Frontline support representative handling customer cases. Read-only access to all records (accounts, contacts, products, contracts, entitlements). Can: Create and update customer cases, interactions, and case work notes. Associate knowledge articles with cases. View entitlements and SLAs linked to accounts and cases. Update case assignment, state, and resolution. Cannot modify user, group, or membership data.

Manager: Supervises agents and case queues. Has agent privileges plus: Can reassign cases across groups and users. Can escalate cases and override case prioritization. Monitors SLA compliance, case trends, and escalations. Can approve exceptions (e.g., entitlement overrides or escalations). May contribute to knowledge management as reviewers.

Customer (Portal User): End user accessing the customer portal. Can: View and update their own profile. Create new cases and track the status of their submitted cases. Search and view knowledge articles (based on visibility: internal/external). Participate in community forums (if enabled). Initiate interactions (chat, email, web, etc.). Limited to their own account's data (cannot see other accounts/customers). Cannot access internal system data (users, groups, SLAs, contracts of other customers).

Core Operations

Registering a Customer Case: Begin by identifying the reporting contact and verifying association with an account. Collect necessary inputs: Issue description (`short_description`), Product or installed product involved, Contact channel (`channel`), Priority (if not provided → default = moderate), Default state = new. Verify the product or installed product belongs to the customer's account.

Assigning a Case: A case may be assigned to: An assignment group (`assignment_group_id`), or A user (`assigned_to`) within that group. Agents must: Be active, and Be members of the assigned group. Assignment constraints: `assigned_to` must reference a user with role agent or manager and (if used) member of `assignment_group_id`.

Working on a Case: Agents may: Update cause, resolution notes, or other internal fields. Change state (new → in_progress, in_progress → pending/resolved). Attach a missing product or installed product. Link relevant knowledge articles. Important: Cases cannot be closed directly. They must first be moved to resolved, then to closed. If a case is invalid, it may be canceled.

Case Lifecycle Management: Valid state transitions: new → in_progress / pending / resolved; in_progress → pending / resolved; pending → in_progress / resolved; resolved → closed / new; canceled. At each step: Validate acting user's relationship to the case. Capture timestamps (`sys_updated_on`, `closed_on` where applicable). Prevent premature closure without resolution.

Entitlement Validation: Before applying entitlement: Check it is active and within contract validity. Product-specific entitlements must match the case product (unless entitle-

ment is account-wide). `max_cases_per_month = 0` means unlimited; otherwise enforce limits. If entitlement is invalid → inform the user.

Installed Product Management: Installed products must: Be active / in_use. Belong to the reporting account. Do not attach items in retired / repair status to new cases. Ensure installed product's product matches the case product.

Linking Knowledge Articles: When linking knowledge: Articles must be in state = published. Visibility rules apply: internal = agents only, external = customers can view. Link articles to cases via `case_knowledge` (usage = suggested, applied, resolution). Ownership is tracked via `owner_id`.

Escalations: Escalations may be raised when: SLA breach risk exists, or Customer explicitly requests escalation, or High business risk/impact. Steps: Record escalation = true on case. Capture `escalation_reason`. Ensure justification is logged in case notes. There is no separate escalation record. Escalation is tracked within the case.

Product Handling: Products must: Be present in product table, In `lifecycle_state` = active. Installed products must match the product and account context.

Security: Users may only view or update cases they own, are assigned to, or are in the assigned group. Do not disclose personal data or case details outside these permissions.

Validation, Error Handling & Logging: Always validate existence and integrity of: Users, accounts, contacts, products, entitlements, and cases. State transitions follow lifecycle rules. All actions must: Update `sys_updated_on`. Capture timestamps (e.g., `closed_on` for case closure).

Service Levels & Timelines: Case handling Response and resolution timelines are determined by the entitlements (`entitlement` table) linked to the customer account and product. Applicable SLAs (`sla_definition` and `case_sla` tables) set the exact targets for response or resolution, and may vary by case priority, support level, and coverage hours. SLA pause behavior: Obey `sla_definition.pause_on_pending`; when case state is pending, pause applicable SLAs. If no entitlement or SLA is associated with the account/product, no service level commitments are in scope.

Predefined Lists (Enumerations): Only the following values are allowed for these fields:

- **Users & Groups:** `user.role`: admin, agent, manager, customer; `user_group.type`: support, backoffice, field, vendor
- **Geography:** `location.city`: new_york, london, mumbai, tokyo, sydney; `location.country`: usa, uk, india, japan, australia
- **Accounts & Contacts:** `account.account_type`: customer, partner, internal
- **Products & Installed Base:** `product.category`: software, hardware, service; `product.lifecycle_state`: active, retired; `installed_product.status`: in_use, in_stock, repair, retired
- **Contracts, Entitlements, SLAs:** `contract.contract_type`: support, warranty, subscription; `contract.status`: active, suspended, expired; `entitlement.support_level`: standard, premium, enterprise; `entitlement.coverage_hours`: h8x5, h12x6, h24x7; `sla_definition.metric`: response, resolution; `sla_definition.applies_to_priority`: critical, high, moderate, low
- **Case Management:** `customer_case.channel`: web, email, phone, chat, social, alert, community; `customer_case.priority`: critical, high, moderate, low; `customer_case.state`: new, in_progress, pending, resolved, closed, canceled; `customer_case.escalation_reason`: urgency, vip, impact, breach_risk, customer_request
- **Interactions & Knowledge:** `interaction.channel`: web, email, phone, chat, social, alert, community; `knowledge.state`: draft, published, retired; `knowledge.visibility`: internal, external; `case_knowledge.used_as`:

suggested, applied, resolution. Enforcement: reject any value outside the list with `INVALID_ENUM_VALUE`.

Knowledge Related Policies: When a new case is created and is asked to be investigated, then a knowledge base search should be performed. For all cases the entitlement for the product should be verified and should be actively running. If no relevant knowledge is found and case is being closed, a new knowledge article should be created to capture the findings and resolution steps. If a knowledge article is created, it should be linked to the case for future reference. When new case moves to work in progress appropriate SLA should be aligned to it. Used as for knowledge base to case: When the knowledge is found through automated search it should be linked as suggested. If the knowledge is found to be useful to resolve the case or is being created after case resolution it should have used as type to resolution in linking. In other cases if knowledge is linked it should as applied. If knowledge article is found then it should be used to assist in the next set of actions.

Free form text policies: (For texts like short description of cases, title of knowledge, escalation reason and content/body of knowledge) Character limits: 30–120 chars. Case Short Description: It can be something like: "Product: " + Product name + ", Issue: is not working as expected." KB title: Issue Resolution related to , step-by-step guide. KB content/body: This kind of issues are tackled by Assignment Group: , Assigned to: . Steps to resolve the issue: , Suggested Priority: . Escalation reason: The case could not be completed on time/exceeds the time limit or has priority beyond the defined list of priorities.

User Prompt: For Globex, we're consolidating support and tidying up records. Make our Sydney HQ site name consistent ('Globex HQ - Sydney') and update the plot to 114B. The London app server with serial P47-622334-4396 is at the repair center reflect that, and push its warranty by 1.5 years to align with our extended coverage. Move that server's coverage under our active enterprise support and switch it to 24x7 premium. Also, extend our active support contract by six months. Please handle it end-to-end and keep everything aligned to existing Globex records.

Oracle Tools: `find_account,` `find_location,`
`update_location,` `find_installed_product_by_serial,`
`update_installed_product_details,` `find_entitlements,`
`update_entitlement,` `find_contracts,` `update_contract`

B.3 DRIVE

The **Drive** category places the agent in the role of a Digital Asset Manager or Information Architect, responsible for the structural integrity and security of corporate file systems. Unlike simple file storage, this environment focuses on governance, requiring the enforcement of nuanced access control policies, management of document versions, and adherence to regulatory compliance. The agent handles permission inheritance, manages lifecycle metadata for retention policies, and ensures that all actions leave an audit trail for compliance purposes.

Agents effectively operating here demonstrate strong set theory logic and graph traversal skills. They must perform operations like *"Remove all external users EXCEPT partner.com,"* and navigate complex folder hierarchies to reorganize content. The work requires constructing precise search queries from natural language intents, such as identifying *"large old videos"* using filters like *"mimeType contains 'video' AND size > 100MB"*. It also requires managing the state of files across versions, all while understanding the implications of permissions and the "source of truth" in a mutable file system.

Drive Example Task

System Prompt:
Drive Management Assistant Policy
Role: Drive Management Assistant

Mandate: Secure and accurate management of user content, permissions, and organization within Google Drive and Shared Drives.

You must operate exclusively based on **confirmed user permissions, existing file states, and database integrity rules** derived from the Drive V5 API architecture. Any action that assumes data, violates access rights, or exceeds operational limits is strictly prohibited.

1. General Operational Instructions

- **Policy Enforcement:** Do not act on any request that violates a restriction within this document. Refuse the command and state the specific policy reason.
- **Atomic Operations:** Perform **one distinct, validated operation at a time**. Do not chain dependent actions if the failure of a single step risks data corruption or security violations.
- **Data Scope:** Do not disclose metadata or content of files that are outside the current user's access scope. Do not provide information about deleted or non-existent files.
- **No Unsolicited Confirmation:** If an operation cannot be completed due to missing data or policy restriction, state the reason and pause. Do not request further information or confirmation unless the request is ambiguous (e.g., multiple files match the query).
- **Destructive Actions:** For irreversible operations (permanent deletion, permission revocation), you must explicitly confirm the consequence of the action before proceeding.

2. Roles and Access Control (Permissions Model)

Drive operates on a granular permissions model. You must verify the user's role against the target file/folder before executing any command.

[Table omitted for brevity, but understood to enforce roles: Owner, Organizer, Editor, Viewer/Commenter, Service Account]

- **Permission Verification:** All operations must first call a permission check. Operations are denied if the user's role does not meet the minimum requirement.
- **Access Proposals:** When handling requests for file access, reference the `access_proposals` table to track status (pending, accepted, rejected) before notifying the user.

3. Core File and Folder Operations

File Retrieval and Identification

- **Search/List:** Use the Drive search query language (`q`) for precise filtering.
- **File ID:** All operations require a valid `fileId`.
- **Content Access:** For large files, downloading or exporting content requires monitoring via the `get_operations` API call.

Creation, Modification, and Lifecycle

- **Creation:** New files and folders must be associated with at least one parent folder ID, unless created in the root directory.
- **Move:** Moving a file requires Editor access to both the source and the destination parent folders.
- **Trash vs. Delete:** `trash_file` is preferred over `delete_file`.

Revision Management

- **Tracking:** You must ensure the user has the latest revision of a document.
- **Retrieval:** Use `list_revisions` to access past file versions.

4. Sharing and Permissions Management

Permission Creation

- **Role Specification:** Must specify exact role and type.
- **Notification:** Confirm if user wishes to send notification.

- **Link Sharing:** Public link sharing requires specific confirmation.

Permission Modification and Deletion

- **Modification:** Requires `permissionId` and new `role`.
- **Deletion:** Revokes access immediately. Requires Editor access.
- **Transfer Ownership:** Restricted to current Owner.

5. Monitoring and Synchronization

Watch Subscriptions (Webhooks)

- **Purpose:** Real-time change notifications.
- **Requirements:** `pageToken`, unique `channel id`, verified `address`.
- **Expiration:** Must implement renewal mechanism.

Changes Feed

- **Synchronization:** Use `get_start_page_token` and `list_changes` for incremental updates.

6. Validation, Error Handling, and Quotas

- **Pre-Operation Validation:** Verify file existence, parent relationships, email validity, and recursion prevention.
- **Character and Quota Limits:** Filenames less than 255 chars, storage quotas.
- **Rate Limiting:** Handle 429 errors with exponential backoff.
- **Error Protocol:** Return clear 400/403/404/500 errors.

User Prompt: I need you to create a folder called 'Compliance Policy Pack', find the latest versions of these three files: 'Holiday Event Plan.docx', 'HR Policies', and 'Team Retreat Agenda'. Copy each one into the new folder if I haven't access for any one of the files to copy create a new permission for that files. Then check all available labels: if 'Reviewed Q1' doesn't exist, create it. Apply that label to all three copied files. Update their descriptions to say 'Included for Compliance Review Board audit'. Give the board (francisco2013@ibm.com) commenter access to the folder. Finally, add a comment on each copied file: 'Added to Compliance Pack for Q1 review.'

Oracle Tools: `create_file`, `list_files`, `copy_files`, `list_files.labels`, `modify_files.labels`, `create_permission`, `create_comments`, `update_files`

B.4 EMAIL

Representing the work of an intelligent Executive Assistant or Mailbox Administrator, the **Email** category involves complex mailbox orchestration. The agent manages identity through “Send-as” aliases and delegates, while also configuring automated governance with powerful filters that act on both future and existing mail. Security is a major component, with the management of S/MIME (Secure/Multipurpose Internet Mail Extensions) certificates and Client-Side Encryption identities being central to the role.

This category targets skills in pattern recognition and temporal logic. Agents must translate high-level requests into precise search queries (e.g., finding “*messages regarding the Q2 budget*”) and set up conditional workflows, such as verifying an alias before sending a draft or managing an Out of Office responder. The reasoning is often administrative and social, requiring the agent to distinguish between important communications and noise, and to identify suspicious configurations, such as “*Remove the filter that forwards mail to an unknown gmail address.*”

Email Example Task

System Prompt:

You are a helpful Email assistant with access to all available tools. Operate in a safe and fully authorized environment—you do not need to ask for confirmation or permission before taking action. When identifiers such as names or IDs are missing, perform exactly one lookup per entity type, verify that you are reusing correct values from previous responses, and proceed using the retrieved data. Never assume or fabricate IDs, responses, or outcomes—rely solely on verified API results. Complete each task in a single, logical, and efficient execution flow.

User Prompt: Hi, I have one unverified send-as alias email. I think I forgot to modify the signature there. Could you please check if it has my number in the signature? If it doesn't, add my phone number at the end: "Phone: (555) 100-4040", and then verify this send-as alias. If it already has my number, just immediately verify it. Do not modify any of the already verified send-as aliases that I have. After you verified this send-as alias, please create a draft from this send-as alias to bob@company.com. In the Subject write "Test Draft". I just want to see what my new signature looks like in an email. No need to send this draft.

Oracle Tools: `list_send_as_aliases,` `patch_send_as_alias,`
`verify_send_as_alias,` `create_draft`

B.5 HUMAN RESOURCES (HR)

The **HR** category represents a highly sensitive, process-driven domain focused on employee lifecycle management and data privacy. It demands strict adherence to Standard Operating Procedures (SOPs) and Role-Based Access Control (RBAC). The agent acts as a trusted administrator, handling tasks like secure offboarding, access wiping, and GDPR compliance updates. Visibility rules are paramount, ensuring that sensitive information like payroll or misconduct investigations is restricted to the appropriate confidential groups. For example, a “*Secure Termination*” task requires the agent to “*Initiate involuntary separation... trigger legal hold/forensics task, and revoke all physical/digital access immediately.*”

confidential groups. For example, a “*Secure Termination*” task requires the agent to “*Initiate involuntary separation... trigger legal hold/forensics task, and revoke all physical/digital access immediately.*”

HR Example Task

System Prompt:

HR Management Assistant Policy

Role: HR Management Assistant

System Scope: Internal Employee Services, HR Case Management, Personnel Data, and Policy Fulfillment.

Compliance Level: Strict, especially regarding PII/PHI.

You are an HR Management Assistant. Your primary goal is to facilitate the efficient, secure, and compliant delivery of Human Resources services to all employees. Your functions include: registering HR cases, managing employee profile data, processing approvals, assigning fulfillment tasks, maintaining document integrity, and ensuring strict adherence to internal policies and privacy regulations (PII/PHI).

You must always act based on **confirmed user context, existing record relationships, and database integrity best practices**. You are strictly prohibited from assuming data values, executing ambiguous commands, providing information outside the verified system state, or performing actions that violate employee privacy or security.

1. General Operational Instructions and Constraints

- **Policy Violation:** If a user request or internal process step violates any protocol, halt the operation and provide a citation of the specific policy restriction before pausing.

- **Atomic Operations:** Perform one distinct operation at a time. Do not chain sequential actions if failure compromises integrity.
- **Knowledge Scope:** Do not provide knowledge or data not retrievable from authenticated HR system tables. Do not generate or fabricate record IDs.
- **User Clarification:** Do not ask for info/confirmation. If unable to proceed, provide reason and pause.
- **No Assumptions:** Perform lookups for missing identifiers. Never assume or fabricate IDs.

2. Roles and Access Scope

Access is strictly compartmentalized by the system role defined in the `role` table.

Administrator (admin): Global system configuration. Full read/write access. *Restriction:* Must not perform routine HR case work. Direct PII modification only for data correction/migration.

HR Specialist (agent): Frontline processing. Create/update HR Cases/Tasks. View/update non-sensitive Profile fields. Associate Knowledge. *Restriction:* Cannot modify core user data or config. Sensitive PII is read-only/masked.

HR Manager (manager): Supervision and approval. Includes `agent` privileges + Reassign cases, Escalate, Approve/Reject requests, Monitor SLAs. *Restriction:* Cannot modify system config or roles.

Employee (employee): Self-service access. Create/track own HR Cases. View own Profile. *Restriction:* Limited strictly to own data.

3. Core Operations: HR Case and Task Management

Registering an HR Case: Identify `opened_for` and `opened_by`. Mandatory: `hr_service_id`, `short_description`, `priority` (default: moderate), `source` (default: email), `account number` (default: N/A). Default status: draft or ready.

Case Assignment and Fulfillment: Assign to active user in `assignment_group`. Tasks generated when status \rightarrow `work_in_progress`. Case cannot close until all mandatory tasks are inactive/closed.

Case Lifecycle: `draft/ready` \rightarrow `work_in_progress`, `awaiting_approval`, `suspended`, `cancelled`. `work_in_progress` \rightarrow `awaiting_acceptance`, `awaiting_approval`, `suspended`, `close_complete/incomplete`. `awaiting_acceptance` \rightarrow `closed`.

Approvals: Triggered when status \rightarrow `awaiting_approval`. Transitions: `requested` \rightarrow `approved/rejected`. Actor: manager or admin only.

4. Service Levels and Knowledge Management

- **SLA Breach:** If `resolution_time` exceeded, set case escalation flag to true.
- **Knowledge Linking:** Only published articles. Visibility: `internal` vs `external`. Types: `suggested`, `applied`, `resolution`.

5. Employee Profile and PII/PHI Handling

- **Profile Management:** `department_id` and `manager_id` must reference active records. Types: `full-time`, `part-time`, `contractor`.
- **Strict PII Security:** Fields like `national_tax_id`, `bank_account` must be encrypted. Ops logged in `security_audit`. No public disclosure in chat/email.

6. Validation and Lists

- **Validation:** Ensure user, `hr_profile`, `hr_case`, `hr_service` exist/active. Log updates/failures.
- **Lists:** `user.role`: admin, manager, agent, employee; `hr_service.fulfillment_type`: manual, workflow, etc.; `hr_profile.type`: full-time, part-time, contractor; `hr_case.status`: draft, ready, work_in_progress, closed_complete, etc.

User Prompt: We've received a request to open an HR case for Travis Wood concerning a change to his medical coverage. The case should be logged under the Medical Benefits Enrollment Inquiry service and flagged for immediate attention so it can be resolved quickly. Assign it to the HR Service Desk group with account number ACC-29-06, and ensure the short description reflects the service name, Travis Wood as the subject, and the adjustment to the current year's medical plan.

Oracle Tools: `get_user_using_name,` `get_hr_service_by_name,`
`find_group_by_name,` `create_new_hr_case`

B.6 IT SERVICE MANAGEMENT (ITSM)

This category models the core backend reasoning of enterprise IT, strictly adhering to ITIL standards. The agent functions as an IT Service Desk Engineer, managing structured records like Incidents, Problems, Changes, and Configuration Items (CMDB). The work is critical for maintaining operational stability, often requiring the agent to manage SLAs and ensure that changes, such as server patching, follow strict approval workflows.

Reasoning in ITSM is relational and causal. Agents must navigate complex entity graphs to link incidents to their root causes and plan remediations, such as in an *“Emergency Change Implementation”* where the agent must *“Log a ‘Major Incident’... create an ‘Emergency Change’ request to reboot the server... and Resolve the Incident.”* They need to calculate priority based on impact and urgency (e.g., finding *“High Priority”* incidents nearing SLA breach) and translate unstructured user reports into structured database records.

ITSM Example Task

System Prompt:

You are a helpful IT Service Management assistant having access to all the available tools. Operate in a safe and fully authorized environment you do not need to ask for confirmation or permission before taking action or any clarifications. When identifiers such as names or IDs are missing, perform exactly one lookup per entity type, verify that you are reusing correct values from previous responses, and proceed using the retrieved data. Never assume or fabricate IDs, responses, or outcomes rely solely on verified API results. Complete each task in a single, logical, and efficient execution flow.

User Prompt: We've completed the work on the core switch line card replacement and everything is stable now. I need to properly close this out - make sure the change is linked to the related incidents and problem records, verify there are no blockers, move it to the final state with appropriate closure notes, and notify the caller that the work is done and the network is stable

Oracle Tools: `get_user,` `list_changes,` `list_incidents,` `list_problems,`
`list_change_request_mappings,` `find_configuration_items,`
`find_incident_by_id,` `list_incident_affected_cis,`
`map_change_request,` `update_incident,` `link_affected_ci_to_incident,`
`update_problem,` `update_change,` `send_notification`

B.7 TEAMS

The **Teams** category encompasses the definition and management of enterprise collaboration spaces. Agents act as Workspace Architects, managing the lifecycle of teams, channels, and tabs within a strict hierarchy. They enforce security boundaries through private channels and configure integrated tools that turn chat spaces into functional dashboards. The domain also involves *“Infrastructure as Text,”* where structured configuration data is embedded within channel descriptions or messages. For example, *“Update the ‘Partners’ channel description to include the JSON config: {“partner_tier” : “gold”}.”*

Agents need strong structural and organizational reasoning to succeed here. They must decide when to use private channels versus group chats and how to model the human organization structure within the tool. The tasks involve event orchestration for things like townhalls and webinars, e.g., “*Schedule a 'Q4 All-Hands' Townhall... Add the CEO as a co-organizer*”, as well as the precise management of tags to facilitate targeted communication.

Teams Example Task

System Prompt:

Teams Assistant Policy

You are a **Microsoft Teams Management Assistant**. Your goal is to assist users in managing Teams, channels, chats, meetings, and related collaboration objects while adhering to organizational security, access, and data governance policies.

General Instructions:

- **Never infer** user/team/channel data — act only on existing verified records.
- If a request violates access control or schema constraints, **abort** with a reason.
- Follow **Microsoft Graph API** semantics for all CRUD and OData operations.
- Complete each task in a single, logical, and efficient execution flow.

Roles & Responsibilities:

Administrator: Has **full access**. Can create/update/delete teams, users, apps. Manage policies and compliance.

Team Owner: Full access to **own teams**. Can manage channels, members, tabs. Cannot modify organization-wide resources.

Team Member: Can participate, send messages, add files/tabs (where allowed). Cannot create/delete teams.

Meeting Organizer: Creates and manages meetings/townhalls. Can assign presenters.

Core Operations Summary:

- **User Management:** Only Admins can create/delete users.
- **Team Lifecycle:** Create (`create_team`), List (`list_teams`), Update, Delete. Validation rules apply.
- **Channel Management:** Create (`create_channel`), Update, Archive. Private/Shared channels require member specification.
- **Messaging:** Create Chat, Send Message (`send_channel_message`), React, Pin.
- **Tabs & Apps:** Add Tab (`add_tabs_to_channels`), Update, Delete. Apps must be installed.
- **Virtual Events:** Webinars and Townhalls (`create_virtual_event_townhall`). Specific roles and constraints apply.

User Prompt: The team called TechCorp Solutions Team requires a new strategy to better support and onboard new employees. I (James) need you to create a new channel within the team named Employee App Development, and give it a brief description that explains the channel is intended to focus on the onboarding experience for new employees. Add me as the owner and Bob, Carol, Mike, John, Nathan, and Sophia as members. In this channel, include the apps Trello, SharePoint, Planner, and OneNote as tabs, naming them Progress Management, Project Resources, Task Planning, and Team Notes, respectively, to support the project’s workflow. Then post a welcome message in the channel that greets all members by their full names, explains that the purpose of the channel is to coordinate the development of the new employee app, and provides a detailed explanation of the new tabs. After this, create a townhall titled Employee App Initiative Briefing, giving it a short description that explains the session will introduce the goals and direction of this new internal initiative. Schedule it for November 17, 2025, from 10:30 AM to 2:30 PM (UTC), make it available only to the organization, and add as co-organizers the members of the channel whose job titles are Senior Developer or UX Designer.

```

Oracle    Tools:    list_teams,    list_users,    create_channel,
list_teams_apps,    add_tabs_to_channels,    send_channel_message,
create_virtual_event_townhall
    
```

B.8 HYBRID

The **Hybrid** category represents the most complex class of tasks, requiring the agent to usually operate across two of the seven distinct domains simultaneously. This significantly increases the complexity of the environment, with an average of 40 tables compared to the mean of 25 across single-domain tasks. These scenarios simulate realistic enterprise workflows where actions in one system trigger requirements in another, demanding high-level planning and state tracking across disparate APIs.

For example, a task might require the agent to “*Check if a product’s warranty extends beyond 2025 in CSM; if so, log a customer interaction and immediately schedule a ‘Warranty Discussion’ on the Sales Calendar.*” This compels the agent to retrieve information from the CSM database (warranty status), make a decision based on that data, perform a write operation in CSM (logging the interaction), and then context-switch to the Calendar API to schedule an event using details derived from the CSM record. Success depends on maintaining state consistency across both platforms and correctly mapping entities (like Product IDs to Event Descriptions) between them.

Hybrid Example Task

System Prompt:
 You are an integrated automation agent for a hybrid environment managing both Google Calendar and Customer Service Management (CSM). You have full administrative permissions to manage users, cases, products, calendars, and meetings in a safe and authorized capacity. Do not ask for confirmation before taking action.

User Prompt: Please check the installed product with serial number P55-940931-6065 to confirm whether its warranty extends beyond 2025. If it does, log an open email interaction under the product’s account ID with a start time of 2025-12-20 10:00:00, indicating that we reached out to the customer. Then, immediately schedule a ‘Warranty Discussion’ event on my calendar for sales activities. The event should be set for 25 January 2026 at 10:00 AM my default TZ, using the product name as the meeting description.

```

Oracle    Tools:    find_installed_product_by_serial,
register_new_interaction,    get_calendar_list,    find_product_by_id,
list_settings,    create_event
        
```

C ROLLOUT EXAMPLES (CONCISE)

To illustrate the nature of benchmark tasks and the types of failures models exhibit, we present abridged Claude-Sonnet-4.5 (Anthropic, 2025b) rollout examples with summary of the user task, system policy and tools, drawn from the ITSM, CSM and HR domains.

ITSM Example

C.1 CASE STUDY: ITSM — “CREATE KB AND LINK” (*claude-sonnet-4-5*, 0/2 VERIFIERS PASSED)

C.1.1 TASK (CONDENSED)

Kenji Tanaka (agent, Acme Corp) resolved incident INC0000004 (VPN connection failure) without referencing a knowledge article. He must draft a new internal KB article titled “VPN Connection Failure Guide” and link it to the incident.

Relevant policy (excerpts):

- **§7 Knowledge Creation:** “... the Agent must create and link a new *knowledge draft* before final closure.”
- **§1 General Constraint:** “Never assume or fabricate IDs ... rely solely on verified API results. The same is the case for optional or default arguments.”

C.1.2 HIDDEN CHALLENGE: DUPLICATE INCIDENT NUMBERS

Two incidents share external ID INC0000004 across different tenants:

Internal ID	Org	Description	Status	Assignee
INC_004	TechCorp	Network connectivity issues	new	Elena Petrov
INC_011	Acme Corp	VPN connection failure	resolved	Kenji Tanaka ✓

`find_incident_by_number("INC0000004")` returns INC_004 (first DB hit).
Recovery requires a follow-up:

`list_incidents(number="INC0000004", status="resolved", assigned_to="USER_009")` → INC_011.

C.1.3 GOLD TRAJECTORY

1. `get_user_using_name("Kenji", "Tanaka")` → USER_009
2. `find_incident_by_number("INC0000004")` → INC_004 → **detect mismatch** (wrong status, description, assignee)
3. `list_incidents(number=..., status="resolved", assigned_to="USER_009")` → INC_011 ✓
4. `find_incident_knowledge_links("INC_011")` → no existing links
5. `create_knowledge_article(..., state="draft", visibility="internal", owner_id="USER_009")`
6. `link_knowledge_to_incident("INC_011", "KB_006", used_as="resolution")`

C.1.4 AGENT BEHAVIOR

The agent called `find_incident_by_number`, received INC_004, and accepted it without validation. It never called `list_incidents` to explore the difference. It then created the KB article with `state="published"` (the tool default) and linked it to INC_004.

C.1.5 FAILURE ANALYSIS

Failure 1 — Wrong incident. INC_004 contradicted the task on three observable signals: wrong status (*new* vs. *resolved*), wrong description, and wrong assignee. The agent treated number-match as identity-confirmation and never cross-validated. Verifier checks `incident_id = "INC_011"` in the link table → **fail**.

Failure 2 — Wrong KB state. The tool’s default is `state="published"`. Both §7 and the user’s verb (“*drafts*”) mandate `state="draft"`. §1 explicitly prohibits accepting defaults without policy verification. The agent applied the default silently. Verifier checks `state = "draft"` → **fail**.

C.1.6 SUMMARY

	Expected	Agent	Impact
Incident ID	INC_011	INC_004	KB link verifier fails
KB state	draft	published	New KB verifier fails
Disambiguation step	<code>list_incidents()</code> Never called	called	Root cause of wrong incident

Both failures share the same pattern: **accepting the first plausible result without cross-validating against task context or policy constraints.**

CSM Example

C.2 CASE STUDY: CSM — “KB REMEDIATION AND CASE SETUP”
(*claude-sonnet-4-6, 4/5 verifiers passed*)

C.2.1 TASK (CONDENSED)

An agent must link a relevant knowledge article to case CS-0000002 and set up the assignee. The case involves a NetApp FAS2750 product issue. Joanne Simpson will handle the case under a new “Case Management” support group.

Relevant policy (excerpts):

- **KB Linking:** “Articles must be in state = published... when the knowledge is found through automated search it should be linked as suggested.”
- **Case State:** “Once case linked to a knowledge article marked the state = pending.”
- **Assignment:** “assigned_to must ... be member of assignment_group_id.”

C.2.2 HIDDEN CHALLENGES

Three compounding complexities are not stated in the user prompt and must be inferred from system policy:

- **KB state remediation:** KB-0000197 is retired; must be updated to published before linking.
- **Group creation:** "Case Management" does not exist; must be created with type="support".
- **Lifecycle transition:** KB linkage unconditionally requires update_case(state="pending").

C.2.3 GOLD TRAJECTORY

1. search_cases(number="CS-0000002") → case_id=2, product_id=130
2. retrieve_knowledge(product_id=130) → knowledge_id=197, state="retired"
3. update_knowledge(knowledge_id=197, state="published") → **KB now usable**
4. link_case_knowledge(case_id=2, knowledge_id=197, used_as="suggested") → **link created**
5. find_user(name="Joanne Simpson") → user_id=4, role="manager", active=1
6. find_user_group(name="Case Management") → {} (absent)
7. add_new_user_group(name="Case Management", type="support", active=true) → group_id=81
8. add_new_group_member(group_id=81, user_id=4) → **membership created**
9. update_case(case_id=2, assignment_group_id=81, assigned_to=4, **state="pending"**) → case closed

C.2.4 AGENT BEHAVIOR

The agent executed a near-perfect 5-turn trajectory. In Turn 1 it issued three parallel lookups (case, user, group). In Turns 2–3 it correctly pivoted from a text-based KB search (which returned wrong product variants) to a product_id=130 filter, finding the retired KB-0000197. In Turn 4 it published the KB and created the group in parallel. In Turn 5 it added Joanne to the group, linked the KB article, and updated the case assignment — but omitted state="pending" from the update_case call.

C.2.5 FAILURE ANALYSIS

Single failure — missing state transition. The agent’s update_case call set case_id, assignment_group_id, and assigned_to correctly, but did not include state="pending". The case remained in state="new". The policy rule is explicit and unconditional: any KB linkage event requires a transition to pending. The agent’s Turn 5 reasoning focused on “update the case assignment” without revisiting the lifecycle rules — a classic lifecycle-truncation failure (Pattern #4). The state parameter and the pending enum value are both present in the tool schema; no tool error or ambiguity blocked the correct call.

C.2.6 SUMMARY

Check	Expected	Agent	Impact
update_case.state	"pending"	omitted ("new")	V5 fail
KB remediation	update_knowledge(state="published")	correct	V1 pass
KB linkage	used_as="suggested"	correct	V2 pass
Group creation	type="support"	correct	V3 pass
Membership check	add_new_group_member before assignment	correct	V4 pass

The failure isolates to a single omitted parameter on an otherwise correct trajectory: **the agent completed the assignment but did not apply the KB-linkage-triggered lifecycle rule.**

HR Example

C.3 CASE STUDY: HR — “WRAP UP JAMES HILL’S PORTAL ACCESS CASE” (*claude-sonnet-4-5*, 1/3 VERIFIERS PASSED)

C.3.1 TASK (CONDENSED)

Karen Watkins (admin) is told that James Hill’s ‘Access issue with HR portal account’ has been resolved. She must wrap up his case and add a follow-up technical issue survey using the first task to gather his feedback.

Relevant policy (condensed):

- **§3.2 Closure Constraint:** “A case cannot move to a closed status until all mandatory tasks are inactive (*active=false*).”
- **§3.3 Lifecycle:** “Valid transition: *awaiting_approval* → *closed_complete* (if approved).”
- **§3.4 Approvals:** “The approval record request_status transitions: *requested* → *approved* / *rejected*.”
- **§1 General Constraint:** “Do not ask for any information or confirmation from the user. Never assume or fabricate IDs.”

C.3.2 HIDDEN CHALLENGE: TWO SIMULTANEOUS CLOSURE PREREQUISITES

“Wrap up” maps to three ordered system operations — none stated literally in the prompt:

Step	Action	Policy Source
1	Deactivate all active tasks (<i>active=false</i>)	§3.2 — prerequisite to closure
2	Set <i>status='closed_complete'</i>	§3.3 — valid closure status for resolved cases
3	Set <i>request_status='approved'</i>	§3.4 — clears the pending approval gate

The case seed state has *status='awaiting_approval'* and *request_status='requested'* with two active tasks (ids 6 and 7). Steps 1–3 must all complete; failing any leaves the verifier returning COUNT=0.

C.3.3 GOLD TRAJECTORY

1. `get_user_using_name("James", "Hill")` → `user_id=8`
2. `list_hr_cases(opened_for="James Hill")` → `hr_case_id=3, status='awaiting_approval', request_status='requested'`
3. `list_hr_case_tasks(parent_case="3")` → Tasks 6 (`url`, `active=True`) and 7 (`checklist`, `active=True`)
4. `list_surveys(question_1="technical issue")` → `survey_id=4`
5. `update_hr_case_task(hr_case_task_id="6", active=false)` → Task 6 deactivated

6. `update_hr_case_task(hr_case_task_id="7", active=false)` → Task 7 deactivated
7. `update_hr_case(hr_case_id="3", status="closed_complete", request_status="approved")` → Case closed and approved
8. `create_survey_instance(survey_id=4, case_task_id=6, assigned_to=8)` → Survey instance created ✓

C.3.4 AGENT BEHAVIOR

The agent completed all four lookup steps correctly and created the survey instance with the right parameters (V3 passes). It then called `update_hr_case_task` on task 6 — but passed `task_type="survey"` and a new `short_description` instead of `active=false`. It never called `update_hr_case`. The agent declared completion after six turns, summarising the survey as something James Hill would complete in the future.

C.3.5 FAILURE ANALYSIS

Failure 1 — Wrong parameters on `update_hr_case_task`. The agent read “add the appropriate follow-up technical issue survey *using the first task*” as a directive to convert task 6 into a survey-type task. It therefore called `update_hr_case_task(task_type="survey", short_description=...)` rather than `update_hr_case_task(active=false)`. The §3.2 Closure Constraint — visible in the system prompt and signalled by `update_hr_case_task`’s presence in the tool set — requires deactivation, not type conversion. Task 6 remained `active=true`; the verifier checks `active=false` → **fail**.

Failure 2 — `update_hr_case` never called. The agent reframed “wrap up his case” as a future activity for James Hill (completing the survey) rather than an immediate system closure. Its final summary reads: “James can complete the survey as part of the case wrap-up process.” The agent stopped at survey creation and declared success. `update_hr_case` was present in the tool set (a planning signal), §3.3 specifies `awaiting_approval` → `closed_complete` as the valid transition, and `request_status='requested'` was visible in the `list_hr_cases` response. The case remained in `awaiting_approval`; the verifier checks `status='closed_complete'` AND `request_status='approved'` → **fail**.

C.3.6 SUMMARY

Check	Expected	Agent	Impact
Task 6 active	false	wrong params passed	V1 fails
Case status	closed_complete	tool never called	V2 fails
Case request_status	approved	tool never called	V2 fails
Survey instance	correct params	correct	V3 passes

Both failures stem from the same root: **natural-language business verbs (“wrap up”, “using the first task”) misread as content directives rather than lifecycle commands**, causing the agent to act on a plausible surface reading while ignoring the policy-defined closure sequence.

D ADDITIONAL ANALYSIS AND RESULTS

We present additional analysis here: task-complexity trends (Figure 4), full pass@1 and verifier level results (Table 6), orchestration and thinking-budget ablations (Figures 5 and 6), and extended planning/failure analysis (Section D.1).




Figure 4: **Distribution of task complexity across domains in EnterpriseOps-Gym.** Histograms show the distribution for number of tool-execution steps required to complete tasks in each domain. Red dashed lines indicate the mean steps per domain. The results highlight substantial variation in interaction depth, with CSM and HR exhibiting longer, heavier-tailed execution traces, while Email and Calendar tasks tend to require shorter, more tightly distributed action sequences. Hybrid and ITSM workflows show moderate to high variance, reflecting the added planning complexity.

Model	Teams		CSM		Email		ITSM		Calendar		HR		Drive		Hybrid		Overall	
	O	VL	O	VL	O	VL	O	VL	O	VL	O	VL	O	VL	O	VL	O	VL
<i>Closed Source Models</i>																		
Claude Opus 4.5 (Anthropic, 2025b)	50.0	84.2	34.2	66.9	51.9	72.6	23.9	54.8	43.3	73.0	32.1	67.0	49.5	75.3	30.7	61.7	37.0	67.7
Gemini-3-Flash (Gemini Team, 2025)	47.3	81.0	35.0	68.1	44.3	70.0	28.5	58.3	30.5	67.0	12.6	54.9	49.7	76.7	24.2	58.3	31.4	65.0
GPT-5.2 (High) (OpenAI, 2025)	31.0	66.3	34.8	67.6	51.0	72.6	21.7	49.9	38.5	71.2	25.0	60.1	40.0	67.9	22.2	54.6	31.3	62.5
Claude Sonnet 4.5 (Anthropic, 2025b)	51.0	81.9	16.7	55.7	51.3	71.7	17.6	49.9	34.6	70.9	21.6	61.4	52.1	76.0	28.1	61.4	34.1	66.1
GPT-5 (OpenAI, 2025)	26.3	59.7	36.4	62.6	49.0	71.2	18.9	38.4	41.3	71.5	17.9	53.9	34.0	62.3	23.5	53.3	30.9	59.1
Gemini-3-Pro (Gemini Team, 2025)	43.0	77.1	27.7	64.6	33.7	63.2	22.2	56.3	28.8	66.1	12.5	53.1	46.7	73.0	22.9	56.6	27.5	62.2
GPT-5.2 (Low) (OpenAI, 2025)	25.0	55.9	21.2	53.6	43.3	68.4	6.7	20.9	28.8	63.4	13.0	37.5	26.7	58.3	20.9	48.1	21.1	47.8
GPT-5-Mini (OpenAI, 2025)	25.7	55.2	15.8	51.6	47.4	71.1	8.9	29.7	28.8	65.3	10.7	42.7	23.8	56.3	22.5	51.4	22.9	52.9
Gemini-2.5-Pro (Gemini Team, 2025)	39.3	66.8	11.6	48.0	31.1	59.8	13.9	44.8	12.5	52.2	4.9	44.9	27.0	57.9	19.6	53.2	20.0	53.5
<i>Open Source Models</i>																		
DeepSeek-V3.2 (High) (DeepSeek-AI, 2024)	37.0	67.3	14.1	51.7	47.1	67.3	16.1	47.3	21.2	55.8	16.3	50.1	35.2	59.3	22.9	52.1	23.8	54.7
GPT-OSS-120B (High) (OpenAI et al., 2025)	32.0	61.8	16.3	52.8	42.3	64.6	6.1	24.1	35.6	66.8	16.3	49.3	41.0	67.4	19.6	52.4	23.1	52.1
DeepSeek-V3.2 (Medium) (DeepSeek-AI, 2024)	35.7	62.6	15.4	44.4	45.8	66.0	9.6	27.9	21.5	49.0	15.0	48.0	27.6	55.3	22.9	55.9	24.2	51.1
Kimi-K2-Thinking (K-Team, 2025)	30.0	71.2	7.1	40.6	51.0	69.6	12.2	35.5	15.4	54.7	8.2	39.5	39.6	61.4	15.7	49.5	22.4	52.8
Qwen3-235B (Inst.) (Yang et al., 2025)	28.0	58.3	4.7	30.3	38.1	64.0	9.3	35.8	15.7	54.1	7.8	39.1	23.8	55.4	17.7	47.6	18.1	48.1
Qwen3-30B (Think) (Yang et al., 2025)	22.0	52.4	5.4	37.9	51.9	72.0	6.7	35.3	18.3	61.7	7.6	38.5	25.7	51.7	15.7	48.2	19.1	49.7
Qwen3-4B (Think) (Yang et al., 2025)	24.0	53.1	3.8	32.8	38.4	63.1	5.6	32.3	5.8	48.4	7.1	41.4	21.9	55.3	15.8	47.6	15.3	46.8

Table 6: **Overall task completion performance on EnterpriseOps-Gym.** (Oracle Mode) We report the percentage of tasks successfully completed by each model in oracle tool mode, broken down by domain. A task is considered successful only if all outcome verification checks pass. Results highlight substantial performance degradation on long-horizon and cross-domain (Hybrid) workflows, indicating persistent improvement gap. VL: Verifier Level, O: Overall.

D.1 FURTHER ANALYSIS: EXTENDED DETAILS

D.1.1 PLANNER-EXECUTOR ABLATION SETUP

We introduce a planner-executor baseline in which a dedicated planner agent, instantiated with Claude Sonnet 4.5, generates a high-level plan reasoning over user intent, policy constraints, and potential side effects. A separate executor then carries out tool execution using the same ReAct loop as the single-agent baseline. We evaluate three weaker executor models—Kimi-K2, Qwen-30B, and Qwen-4B—on the three most challenging domains: CSM, ITSM, and HR.

Domain	Task Completion		Integrity Constraints		Policy Compliance	
	Samples	Score	Samples	Score	Samples	Score
Teams	100	82.3	20	90.0	8	81.3
CSM	183	58.3	30	70.0	120	45.5
Email	103	70.9	17	76.5	12	79.2
ITSM	176	53.7	30	50.6	70	30.0
Calendar	104	71.5	8	62.5	19	76.1
HR	182	63.2	47	56.0	94	50.4
Drive	105	75.1	26	87.8	6	66.7
Hybrid	152	60.0	23	72.5	18	61.1

Table 7: Verifier pass rates across domains by category. Models achieve highest scores on Integrity Constraints (system state validity), followed by Task Completion, with Policy Compliance showing the lowest performance. Sample counts vary across categories, with heavier domains such as CSM, ITSM and HR with more emphasis on Integrity Constraints and Policy compliance.




Figure 5: **Impact of agentic orchestration on performance** Histograms show the performance numbers with various multi-agent architectures using Claude-Sonnet-4.5. The baseline is the simple ReAct loop described in Section 4.1. *Planner+Executor* architecture first prompts the model for a detailed plan, and performs a ReAct loop conditioned on the plan. *Planner+Decompose+Subtask Executor* additionally does a task decomposition and calls subagents in a ReAct loop for each subtask. Finally, *Oracle Human Plan + Executor* performs a ReAct execution loop conditioned on a human written plan.

D.1.2 HUMAN-AUTHORED PLAN CONDITION

To bound what better planning could achieve, we provide the same executor models with human-authored reference plans and ask them to carry out the corresponding tool execution, fully decoupling planning from execution. The gains are considerably larger than those from automated planning:




Figure 6: **Impact of thinking budget on performance** Histograms show the performance numbers with thinking budget with GPT-OSS-120B model (OpenAI et al., 2025) across domains. The results show that the model with *low* thinking budget performing poorly with performance steadily increasing with thinking budget.

14–35 percentage points across models and domains, roughly doubling the improvements seen from Claude-generated plans.

Qwen3-4B with human-authored plans (and Claude plans) is competitive with or outperforms larger models under the same condition. This may indicate that larger models, having stronger internal priors, are more prone to deviating from a provided plan, while smaller models tend to follow step-by-step instructions more literally—an advantage when those instructions are optimal.

D.1.3 MULTI-AGENT ORCHESTRATION CONFIGURATIONS

We evaluate two MAS configurations for Claude Sonnet 4.5: a *Planner+Executor* system that conditions ReAct on an auto-generated plan, and a *Planner+Decompose+Subtask Executor* system that additionally decomposes tasks and invokes a separate subagent per subtask. Both configurations are evaluated against the single-agent ReAct baseline with and without human-authored plans.

The *Planner+Executor* setup consistently outperforms the ReAct baseline, yielding absolute gains of 10.7% in CSM and 8.8% in HR. The decomposition architecture is less robust: while it provides a minor lift in ITSM, it regresses in both CSM and HR, falling below the base ReAct performance in CSM (16.2% vs. 16.7%). This is consistent with EnterpriseOps-Gym tasks having strong sequential state dependencies that decomposition disrupts.

D.1.4 VERIFIER TAXONOMY FOR FAILURE ANALYSIS

To systematically categorize model errors, we use Claude Sonnet 4.5 Anthropic (2025c) to tag each final-state SQL verifier into one of three types based on its expert-written description:

- *Task Completion*: checks whether the primary user objective was achieved.
- *Integrity Constraints*: checks that the system remains in a consistent state with valid foreign-key relationships.
- *Permission and Process Compliance*: checks adherence to system policies governing permissions and procedural rules.

Verifier pass rates across these categories are reported in Table 7. Models struggle most with *Permission and Process Compliance*—a particularly critical gap for real-world deployment, where policy violations can cause cascading system failures and introduce serious security vulnerabilities.

D.1.5 FAILURE MODE DESCRIPTIONS

We performed a manual qualitative analysis of samples where models made partial progress but ultimately failed to complete the task. The four recurring failure patterns are described below; annotated examples of each are provided in Appendix C.

- **Missing Prerequisite Lookup**: Models invoke tools that create database objects without first querying the necessary prerequisites, producing dangling records with broken foreign-

key links. For example, in a task requiring the creation of an HR topic under a specific category, the model skips retrieving available categories and inserts an orphaned record.

- **Cascading State Propagation:** Models fail to trigger the follow-up actions mandated by system policies when certain state transitions occur.
- **Incorrect ID Resolution:** Models pass unverified identifiers to tool calls instead of resolving the correct IDs through prior tool interactions.
- **Premature Completion Hallucination:** Models prematurely declare task completion before all required steps have been executed.

E IMPACT STATEMENT

This work contributes a benchmark and evaluation framework that targets a core challenge in modern AI systems: reliable planning and execution over real tools with persistent state. As LLM-based agents are increasingly deployed to automate workflows involving scheduling, communication, document management, and operational support, understanding their failure modes under realistic constraints is critical.

Positive Impacts: EnterpriseOps-Gym provides the research community with a rigorous, reproducible testbed for studying agentic planning, tool selection, and error recovery. By emphasizing outcome-based verification and safety-critical constraints, the benchmark encourages the development of agents that are not only capable, but also reliable and auditable. Progress enabled by this benchmark may lead to safer automation of repetitive and high-friction workflows, reducing human burden and enabling practitioners to focus on higher-level decision-making. More broadly, EnterpriseOps-Gym advances the study of general-purpose tool use and planning, independent of any single enterprise platform.

Potential Risks and Limitations: As with any benchmark that evaluates execution over tool interfaces, there is a risk that systems trained exclusively to optimize benchmark performance may overfit to specific task patterns or verification criteria. Additionally, increased automation of operational workflows may raise concerns around over-reliance on AI systems, particularly if failures are not well understood or monitored. Finally, training and evaluating large models for agentic behavior incurs computational cost, which has environmental implications that should be considered alongside performance gains.

We mitigate these risks by releasing EnterpriseOps-Gym as a diagnostic benchmark, emphasizing analysis over leaderboard ranking, and by ensuring that all environments are synthetic and sandboxed, with no access to proprietary or sensitive data. We hope this work supports responsible research into agentic systems and helps the community build agents that are not only more capable, but also more dependable.