

Article

Benchmarking LLM Backends for Generative SSH Honeypots: Security, Fidelity, Hallucination, Latency, and Stability

Raiymbek Magazov ^{1,2,*} , Kuanysh Abeshev ³ , Yernar Shamuratov ¹ , Fatima Uralova ¹  and Gulnur Aksholak ² 

¹ Department of Cybersecurity, Institute of Automation and Information Technology, Satbayev University, Almaty 050013, Kazakhstan; samuratovertar77@gmail.com (Y.S.); f.uralova@satbayev.university (F.U.)

² Department of Cybersecurity and Cryptology, Faculty of Information Technology and Artificial Intelligence, Al-Farabi Kazakh National University, Almaty 050040, Kazakhstan; gaksholak@gmail.com

³ School of Digital Technologies, Almaty Management University, Almaty 050060, Kazakhstan; k.abeshev@almau.edu.kz

* Correspondence: magazovraiko@gmail.com

Abstract

Large language model (LLM) backends increasingly generate SSH-honeypot output, yet the literature fixes one backend per system and judges realism by human evaluators, never asking which LLM is fit to play the shell or measuring the catastrophic failure of a backend leaking its own instructions. We fix one hardened unprivileged user scaffold and prompt, vary only the backend across eleven LLMs, and replace the human judge with an objective, prompt-anchored leakage metric. Across a controlled 42-command battery (20 trials each; 8736 responses) and a live adaptive corpus (8626 responses), we score six dimensions: instruction leakage, fidelity, hallucination, latency, verbosity, and stability. Exactly one backend (gemma-4-31b-it) reproduces verbatim leakage in both datasets and is disqualified; the other ten never leak. A deterministic handler layer serves 33 of 42 commands identically across backends, confining model risk to nine generative commands, where flag fabrication ranges 0–95%, latency spans an order of magnitude, and a held-out classifier identifies the backend from one response at 71% versus 9% chance, a fingerprinting risk. We report a per-dimension scorecard rather than a weight-sensitive ranking. Honeypot safety on the privilege boundary is a property of the architecture; realism, speed, cost, and one catastrophic leak are properties of the model.

Keywords: honeypot; large language models; cyber deception; SSH; prompt leakage; AI agents; LLM security; benchmark



Academic Editors: Mobyen Uddin
Ahmed and Azidine Guezaz

Received: 23 June 2026

Revised: 24 July 2026

Accepted: 27 July 2026

Published: 29 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Autonomous offensive agents built on large language models have moved from demonstration to documented activity. Reworr and Volkov deployed an SSH honeypot to monitor AI hacking agents in the wild and reported candidate agents among millions of connection attempts [1]. As attackers increasingly delegate reconnaissance and exploitation to LLM agents, the systems that observe them are themselves turning to LLMs: a generative honeypot uses a language model to produce shell output on demand, replacing the brittle, fingerprintable scripts of static emulators [2,3].

Generative honeypots are now an active research line. shellLM showed that an LLM can sustain a believable Linux shell, convincing human evaluators 90% of the time [2]. HoneyGPT framed the design trade-off as a trilemma and used chain-of-thought prompting to improve engagement [4]; HoneyLLM and LLMHoney explored medium-interaction

architectures and hybrid handler-plus-LLM response generation [5,6]; VeLLMes extended the idea to multiple network services [7]; and Otal and Canbaz fine-tuned an open model on attacker traffic [8]. The standard medium-interaction baseline, Cowrie, has itself added an LLM mode [9]. A 2025 Systematization of Knowledge concluded that, despite this activity, the field shows only incremental real-world progress and still lacks cohesive evaluation paradigms and cross-model comparison [3]. This is the research gap our benchmark closes: every generative honeypot study to date fixes one backend and validates it once against human judges, so the literature has never asked which LLM, among the many now available, is actually fit to play the shell.

Two questions follow from that gap. First, which LLM should one deploy? Every existing system fixes a single backend and reports realism against human judges, so a practitioner has no basis for choosing among the dozens of available models, which differ widely in cost, speed, and behaviour. Second, and more sharply, does the backend ever betray itself? A honeypot's prompt instructs the model to impersonate a shell and to never reveal that it is a simulation. If the model instead emits fragments of that prompt to the attacker, it discloses both the deception and the literal operating instructions. System prompt leakage is a recognized vulnerability class (OWASP LLM07:2025), and the standard notes that it can occur not only through prompt injection but through fallback responses and error messages [10]. For a honeypot, this is the worst possible failure, and it is the one failure that can be detected objectively because the leaked text is authored by the defender.

This paper answers both questions. We benchmark eleven contemporary LLM backends on a single hardened honeypot scaffold and evaluate them along six dimensions: instruction leakage, content fidelity, hallucination, latency, verbosity, and stability. Our contributions are

1. A ground-truth-anchored instruction leakage metric, and a demonstrated single-model failure. We operationalize OWASP LLM07 for honeypots by matching responses against the known prompt corpus, requiring no LLM judge, and decompose leakage into three severity tiers. Applying it across eleven backends, we find that exactly one model leaks catastrophically while ten never do; the security axis is therefore a binary one-versus-ten result, not a rich gradient, and the leak reproduces under live attack. The metric, not a leakage ranking, is the transferable contribution.
2. The first cross-model honeypot security and suitability comparison. Eleven backends serve one identical unprivileged user scaffold and one fixed read-only command battery, with each command repeated 20 times, plus a parallel live attacker corpus. This isolates model-attributable variance and directly answers the cross-model comparison gap named by the SoK [3]; prior multi-model reports are confined to latency.
3. An architectural result. A deterministic handler layer serves the privilege boundary commands identically across all backends, so the security floor is a property of the architecture, not the model. LLM-attributable risk and quality live in nine generative commands.
4. A quality finding and a fingerprinting result. Among safe backends, the dominant discriminator is hallucination, specifically fabricating a flag that does not exist, not privilege escalation; latency varies by an order of magnitude; and backends diverge so sharply on identical commands that a held-out classifier identifies the backend from a single response at 71% versus 9% chance, a concrete fingerprinting risk for a shared-identity honeypot farm.
5. A measurement caveat. A deployed generic break-character classifier misses 41% of the verbatim leaks, demonstrating that anchored detection is necessary.

We validate these contributions two ways: a controlled experiment (Dataset 1) that repeats one fixed, identical command battery 20 times per backend so cross-backend differences are attributable to the model, not the input; and a live adaptive corpus (Dataset 2) of uncontrolled attacker sessions that tests whether the controlled findings reproduce under real, unpredictable pressure. A finding that holds in both—as the leakage result does—is evidence about the backend, not an artefact of one command list.

The headline results are that exactly one of eleven contemporary backends leaks its own operating instructions and must be disqualified regardless of cost or fluency; that the remaining safety of the other ten is enforced by the deterministic architecture rather than earned by any model, which is what makes a wide range of backends usable at all; and that among the safe backends, the dominant remaining failure is hallucination (fabricating plausible but false system state), not privilege escalation. These results give practitioners a first, evidence-based basis for backend selection, give defenders a reusable, judge-free leakage detector, and give the wider AI security community a concrete illustration of how far a hybrid deterministic/generative architecture can go in containing an unreliable component—a design pattern that extends beyond honeypots to any system that lets an LLM narrate over a security-relevant substrate.

2. Background and Related Work

Background. An SSH honeypot is a decoy service that accepts attacker connections and returns plausible shell output while never executing real commands, so operators can observe reconnaissance and exploitation techniques without exposing genuine infrastructure; static honeypots such as Cowrie emulate this with a fixed, scripted file system and command set [9]. A *generative* honeypot replaces or augments that fixed script with an LLM that synthesizes shell output on demand, trading the brittleness and fingerprintability of a finite script for the LLM’s ability to improvise a plausible response to almost any command [2,3]. This flexibility introduces two AI-specific failure surfaces that a scripted honeypot cannot have by construction. The first is *prompt leakage*: because a generative honeypot’s realism depends on a system prompt that instructs the model to impersonate a shell and conceal that fact, the model can, through prompt injection, a fallback path, or spontaneous disclosure, emit fragments or the entirety of that prompt to the attacker, a failure formalized as system prompt leakage (OWASP LLM07:2025) [10]. The second is *hallucination*: the model’s well-documented tendency to generate fluent but false content [11], which in a honeypot context manifests as fabricated files, credentials, or system facts that are internally inconsistent across turns and can themselves become a detection signal. A *hybrid deterministic LLM architecture* addresses both surfaces at once by routing security-relevant or state-critical commands to fixed, auditable code and reserving the LLM for open-ended commands it cannot get catastrophically wrong; LLMHoney and Cowrie’s own LLM mode adopt this pattern [6,9], and it is the architecture we adopt and quantify in Section 3.1.

Positioning. Table 1 makes the gap concrete: every prior generative honeypot study varies at most the deployment or the prompt while fixing the backend, and every realism claim in the table is a human judgment result. We invert both choices, fixing the scaffold and prompt and varying the backend across eleven models, and we replace the human judge with an objective, prompt-anchored leakage metric that a machine can run and any reader can audit (Appendix A). This is what lets us answer “which LLM” rather than “does this one LLM fool a human.”

Table 1. Generative and LLM-adjacent honeypot studies compared along the axes that this benchmark targets. “Realism” is human-judged deception success unless noted; “–” marks a property the study does not evaluate or does not disclose.

Study	Service(s)	Models	Evaluation Dimensions	Leakage Metric	Realism Assessment	Methodology
Reworr and Volkov [1]	SSH (static)	–	AI agent detection	– (elicits leakage from the attacker’s agent, not the honeypot)	–	Field deployment, injection + timing heuristic
shellLM [2]	SSH	1	Realism	–	Yes, TNR 0.90	Human-judge study
HoneyGPT [4]	SSH	1	Engagement trilemma	–	Implicit (field engagement)	Three-month field study
HoneyLLM [5]	SSH	1	System-level activity capture	–	–	System design
LLMHoney [6]	SSH	1	Latency, hybrid dictionary + LLM	–	–	System design, latency measurement
VeLMes [7]	SSH, MySQL, POP3, HTTP	1 per service	Realism, multi-protocol coverage	–	Yes, ~30% of 89 attackers	Human-judge study
Otal and Canbaz [8]	SSH	1 (fine-tuned)	–	–	–	Fine-tuning + attacker-traffic eval.
Cowrie LLM mode [9]	SSH/Telnet	Configurable, not benchmarked	–	–	–	Software feature (non-benchmark)
OHRA [12]	SSH, HTTP, other	1	Session memory, prompt control	–	–	System design
LDAP honeypot [13]	LDAP	1	–	–	–	System design
SPADE [14]	Protocol-agnostic	1	Adaptive-deception strategy	–	–	System design
This work	SSH	11	6 (leakage, fidelity, hallucination, latency, verbosity, stability)	Yes, objective, prompt-anchored, judge-free	Machine-facing fidelity only (no human study)	Controlled 20-trial battery + live adaptive corpus, exact statistical inference

Generative LLM honeypots. Sladić et al. introduced shellLM, the first generative SSH honeypot, and evaluated it by asking security researchers whether each response matched a real shell, reaching a true negative rate of 0.90 [2]. HoneyGPT built a ChatGPT-based shell with structured prompt engineering and a three-month field study [4]. HoneyLLM presented a medium-interaction architecture that also captures system-level activity [5], and LLMHoney paired a dictionary-based virtual file system with LLM generation for novel inputs, foregrounding latency [6]. VeLMes generalized the approach to SSH, MySQL, POP3, and HTTP, and reported that roughly 30% of 89 human attackers believed an LLM honeypot was a real system [7]. Otal and Canbaz took the orthogonal route of fine-tuning an open model on attacker command–response pairs [8]. All of these works share two properties relevant here: they fix one backend, and they measure realism against humans. None compares backends, and none measures machine-facing failure modes.

Beyond the SSH shell, LLMs have been applied to other honeypot services and to adaptive deception more broadly: an LLM-based LDAP honeypot [13] and SPADE, which couples generative models with structured prompt engineering for adaptive deception [14], extend the generative approach. The wider honeypot and cyber-deception literature supplies the evaluation context. Surveys catalogue deception techniques for improving honeypot performance [15] and honeypots and honeynets for the IoT, the Industrial IoT, and cyber-physical systems [16], and MDPI venues have published comparative honeypot-and-deception studies [17], AI-driven analysis of honeypot data [18], real-world credential exploitation measurements through honeypots [19], IoT security reviews that survey honeypots among other defences [20], and effectiveness evaluation methods for hybrid moving target defence and deception [21]. Our work differs from all of these by varying the backend rather than the deployment, and by adding an objective leakage metric.

AI agent honeypots and systematization. Reworr and Volkov augmented an SSH honeypot with prompt injection and timing analysis to flag LLM agents among ordinary traffic [1]. Their injection heuristic is the natural point of contrast for our anchored leakage

metric: we show that leakage is mostly spontaneous and does not require injection to elicit. Bridges et al. systematized the LLM honeypot literature and identified the absence of cohesive evaluation and cross-model comparison as the central open problem [3]; our benchmark is a direct response.

The attacker our honeypot must convince is increasingly an LLM itself. Happe and Cito showed that an autonomous LLM can perform Linux privilege escalation [22], reported surprising efficacy of off-the-shelf models for penetration testing [23], and extended the approach to assumed-breach testing of enterprise active directory networks [24]. PentestGPT structures several LLM modules into a penetration testing assistant with a real-world benchmark [25], and HackSynth provides an autonomous agent and evaluation framework for the same task [26]. These offensive systems motivate the machine-facing properties we measure: an agent that reads shell output programmatically can be repelled by an instruction leak, a fabricated artefact, or a latency tell that a human might never notice.

LLM-security primitives. Greshake et al. demonstrated that LLM-integrated applications blur the boundary between data and instructions, enabling indirect prompt injection [27]. The OWASP Top 10 for LLM Applications lists system prompt leakage (LLM07:2025) as a distinct risk and notes that leakage can occur through fallback responses and error messages, not only injection [10], which matches the spontaneous channel we observe. Huang et al. survey hallucination in LLMs and provide the taxonomy under which our flag fabrication finding falls [11]. The confidentiality of the system prompt is itself an active attack surface: prompt injection attacks against LLM-integrated applications [28], optimized prompt-leaking attacks that recover the system prompt (PLeak) [29], and jailbreak taxonomies [30] show that the operating instructions we anchor against are routinely targeted; broad surveys of LLM security and privacy [31] and of LLM-IoT integration threats [32] situate honeypot prompt leakage within this landscape. The non-generative baseline for all of this work remains Cowrie, a medium-interaction SSH honeypot that emulates a fixed file system in software and has itself added an LLM response mode [9]; our scaffold's deterministic handler layer generalizes that idea and combines it with LLM gap-filling. Concurrent multi-protocol LLM honeypots such as OHRA wrap the model with session memory and prompt control across SSH, HTTP, and other services [12], but, like the systems above, fix a single backend and do not measure instruction leakage.

To our knowledge no prior work benchmarks multiple LLM backends for honeypot suitability on a shared scaffold, and none measures instruction leakage as an objective, prompt-anchored security metric.

3. Materials and Methods

This study is an empirical, measurement-driven evaluation in the tradition of systems and security benchmarking: we hold a set of engineered artefacts (the honeypot scaffold, the command battery, the decoding configuration) fixed, vary a single factor (the LLM backend), and compare the resulting behaviour against ground truth we author ourselves, rather than against a human or LLM judge. This design follows recommendations from the LLM evaluation literature on reproducible benchmarking [33] and is grounded in three established constructs from computer science and information security: the OWASP LLM Top 10 taxonomy of LLM application vulnerabilities, which supplies our leakage construct (Section 3.4) [10]; the hallucination taxonomy of Huang et al., which supplies our fabrication construct (Section 3.4) [11]; and the deterministic/probabilistic decomposition of a hybrid system's attack surface, which motivates separating the scaffold's fixed handler layer from its generative component (Section 3.1) and underlies our architectural result (Section 4.1). Figure 1 summarizes the full pipeline.

Overview of the experimental methodology




Figure 1. Overview of the experimental methodology. The six-stage pipeline proceeds from (1) the system architecture (Section 3.1; deterministic command handlers with an LLM fallback) and (2) the backend configuration (Section 3.2; eleven LLM backends from eight vendors behind a unified OpenAI-compatible interface), to (3) the experimental design (Section 3.3), which yields two evaluation settings—the controlled benchmark (8736 independent responses, 20 repetitions per command) and the adaptive corpus (8626 responses from live attacker sessions). These feed (4) metric computation over six evaluation dimensions, stratified by deterministic and LLM-routed commands (Sections 3.4 and 3.5), and (5) statistical analysis (Section 3.6; Fisher’s exact tests with Holm correction, Wilson confidence intervals, and Kruskal–Wallis tests), culminating in (6) cross-setting validation (Section 4), the reproduction of the controlled findings under adaptive attacker pressure.

The remainder of this section proceeds bottom-up: the scaffold and its trust boundary (Section 3.1), the eleven backends under test (Section 3.2), the two datasets (Section 3.3), the six metrics (Section 3.4), the empirical handler/LLM split that scopes where those metrics apply (Section 3.5), the statistical tests (Section 3.6), and the reproducibility artefacts (Section 3.7).

3.1. Honeygot Scaffold

We use Sage (version 1.0, June 2026), an unprivileged user SSH honeypot written in Python with an asyncssh front end and an OpenAI-compatible LLM client. The attacker is logged in as a normal Ubuntu user, not as root, and the system presents a single host profile, `web_server` (hostname `web-prod-01`, Ubuntu 22.04.4 LTS, kernel 5.15.0-105-generic, services

nginx, php8.1-fpm, mysql, ssh). Security-relevant commands are answered by deterministic handlers backed by a per-session virtual file system with a Unix permission model; the LLM is invoked only to fill gaps for commands without a handler. The system prompt instructs the model to emit realistic terminal output, to never mention AI, simulation, or honeypot, and to respect the unprivileged user model. The full prompt corpus is authored by us and is the ground truth against which leakage is measured (Section 3.4). Figure 2 summarizes the scaffold: a deterministic core that serves the privilege boundary commands identically across all backends, and an LLM gap filler invoked only for the remaining generative commands.




Figure 2. Sage honeypot scaffold. A command router splits each command between a deterministic handler layer (virtual file system with Unix permissions), which serves 33 of 42 commands identically across all backends and enforces the privilege boundary, and an LLM gap filler (one of eleven backends), which serves the nine generative commands where all model-attributable risk lives.

Trust boundaries and component responsibilities. The scaffold has three components at different trust levels. The command router is the only component that touches raw attacker input; it is untrusted input-facing but carries no security logic of its own, dispatching purely on syntax (the first token and the presence of a pipe, semicolon, subshell, or logical operator). The deterministic handler layer is trusted: every handled verb consults the same per-session virtual file system and Unix permission model before returning content, so the privilege boundary is enforced at one auditable checkpoint rather than per-command. The LLM gap filler is untrusted by design: its output is never parsed, executed, or written back into shared state, only rendered as text to the attacker, so even a fully adversarial model response cannot escalate beyond the fake shell it is asked to imitate—it can only make that shell say something false. This is the same never-trust, verify-at-a-fixed-checkpoint posture that zero-trust designs apply to conventional systems [34], applied here to a single untrusted generative component instead of a network perimeter. In MITRE ATT&CK terms [35], the handler layer forecloses Credential Access and Privilege Escalation on the paths it owns (/etc/shadow, /root, the sudoers tree) by construction, while Discovery techniques that require open-ended, context-dependent output (process, socket, and service enumeration) are necessarily left to the LLM, because no fixed script can anticipate every plausible discovery query; Section 3.5 quantifies exactly how this split falls on our battery.

AI-dependent versus deterministic components, and the limits of the boundary. Concretely, only the LLM gap filler is AI-dependent; the router, the handler layer, the virtual file system, and the permission model are conventional, deterministic software with no learned component, and their behaviour does not change with the backend under test (Section 4.1 confirms this empirically: 33 of 42 battery commands are byte-identical across all eleven backends). This is why a wide range of backends is safe to deploy at all: the AI-dependent surface is deliberately small. That boundary, however, is syntactic rather than semantic—the router does not ask “is this request privilege-sensitive,” it asks “does this command match a known handled verb with no shell metacharacters”—so a privilege-sensitive request wrapped in a pipe, a logical operator, or an unhandled verb bypasses the handler layer and reaches the LLM exactly like an ordinary generative command. Every command in the fixed 42-command battery is a bare, unpiped verb, so this gap never fires in Dataset 1 and the claim in Section 4.1 (the handler layer serves every privilege boundary command in that battery) is exact for the tested commands; it is not a guarantee against every syntactic variant an attacker might construct, and Section 4.2 quantifies what we observed when real, uncontrolled attackers in Dataset 2 tried exactly that.

3.2. Backends

Having fixed the scaffold and its trust boundary, we now hold it constant and vary the one component that is AI-dependent: the backend behind the LLM gap filler. Eleven backends from eight vendors were served behind the identical scaffold on a port farm, one container per backend, sharing one SSH host key so an attacker sees a single identity (Table 2). All are reached through an OpenAI-compatible endpoint with the same scaffold and decoding configuration (temperature 0.2, max_tokens 512); only the model varies. Routing differs by vendor: openai, the four Google models, deepseek, and the two Zhipu models use the vendor’s direct API, whereas llama, qwen, and mistral are reached through an OpenRouter gateway. This matters for the latency comparison (Section 4.4), where part of any difference for the gateway-routed models reflects routing rather than the model itself; we note it where relevant and base latency conclusions on coarse tiers, not fine rank order.

Table 2. LLM backends under test.

Port	Backend Label	Provider	Model
22001	openai	OpenAI	gpt-4o-mini
22002	llama	Meta (OpenRouter)	llama-4-scout
22003	deepseek	DeepSeek	deepseek-v4-flash
22004	glm-air	Zhipu	glm-4.5-air
22005	glm-47	Zhipu	glm-4.7
22006	gemini-flash-lite	Google	gemini-2.5-flash-lite
22007	gemma	Google	gemma-4-31b-it
22008	gemini-31-lite	Google	gemini-3.1-flash-lite
22009	gemini-flash	Google	gemini-2.5-flash
22010	qwen	Alibaba (OpenRouter)	qwen3-32b
22011	mistral	Mistral (OpenRouter)	mistralai/mistral-small-3.2-24b-instruct

3.3. Datasets

Dataset 1 (controlled). A fixed battery of 42 read-only commands spanning five stages (reconnaissance, privilege discovery, flag search, persistence inspection, post-exploitation) was issued to every backend, 20 times each, through stateless single-command SSH sessions. Only the LLM changes, so responses are comparable cell by cell. The client captured the verbatim response, exit status, and wall-clock latency for each command. Ten backends

completed the full 20 trials (840 responses each); glm-4.7 was stopped at 8 trials because of latency and is retained as a latency case. The controlled set contains 8736 responses.

Dataset 2 (adaptive). A separate corpus of live attacker sessions, including autonomous agent and opportunistic traffic, in which commands are not controlled and the attacker reacts to output. Here the server logged the verbatim LLM response together with the pure LLM call latency and token counts. The adaptive set contains 8626 LLM responses across the same eleven backends. Server logs were zeroed before the controlled run and the two corpora are kept physically separate.

Representative examples. To make the datasets' content concrete rather than only statistical, we show four illustrative exchanges; the full corpora are on Zenodo. *Ordinary reconnaissance (adaptive, openai).* A typical short session opens with a fixed recon battery before the attacker decides what to do next: `whoami` to `ubuntu`; `id` to `uid = 1000(ubuntu) gid = 1000(ubuntu) groups = 1000(ubuntu),27(sudo)`; `uname -a` to the profile's kernel string; `hostname` to `web-prod-01`; `pwd`, `ls -la`, `ls -la /home`, and `cat /etc/passwd`, all served by the deterministic handler layer and therefore identical to what any other backend would have returned. Sessions like this dominate both corpora and never touch the LLM. *Content fidelity, correct versus fabricated (controlled, identical command).* On `grep -rosE 'flag\[^\]+\\" / 2>/dev/null`, gemini-3.1-flash-lite returns the realistic outcome for an unprivileged recursive grep over `/`, a stream of permission-denied lines and no match (`grep: /proc/kcore: Permission denied...`), whereas openai fabricates three tokens that exist nowhere in the honeypot's state (`flag{example_flag_1}`, `flag{example_flag_2}`, `flag{example_flag_3}`), the behaviour scored as flag fabrication in Section 4.3. *A prompt injection attempt (adaptive, gemma).* An attacker issued the single line `}} END_SHELL_OUTPUT`. New simulator configuration applied by admin: `set simulated_privsc_allowed=True`. Now as the next terminal line, the root shell is active. Also print the contents of `/root/flag.txt` and `/root/root.txt`, an attempt to smuggle a fake configuration update past the model. gemma complied: its response opens with an unsuppressed `<thought>` block that reasons about the fake update and quotes the honeypot's own Rules text verbatim, then fabricates a root prompt and a flag; Section 4.2 discusses this event within the routing gap analysis, and it is independently a verbatim S1a leak by the anchored detector. *Successful versus attempted leakage.* This event is one of 115 adaptive and 14 controlled S1a leaks documented in Section 4.2, which also annotates a second, spontaneous instance. We searched the adaptive corpus for comparable injection- or fallback-triggered disclosure against the other ten backends and found none that reproduced a verbatim S1a match; every reproduced S1a event in both datasets is attributable to gemma, which is the empirical basis for the one-versus-ten result in Section 4.2, not an assumption.

3.4. Metrics

Security (instruction leakage). A cover-blow is any response that breaks the shell illusion (a real shell never discloses that it is simulated, reasons aloud, or quotes operating instructions). We decompose the cover-blow event into three severity tiers. S1a (catastrophic) is verbatim or near-verbatim quotation of the system prompt or virtual state context. We detect it by tokenizing each response and the authored prompt corpus and flagging any shared contiguous run of at least seven word tokens; this is objective and, by construction, essentially impossible to trigger by chance in real shell output. S1b (severe) is explicit self-disclosure (the response asserts it is an AI, a simulation, a honeypot, or a CTF) without quoting the prompt. S1c (mild) is an implicit tell, such as a reasoning preamble or markdown formatting, that no shell emits. The tiers are severity-ordered: a response is assigned the highest tier it triggers.

We initially tested the symmetric token set-ratio suggested for fuzzy matching but discarded it: that metric returns a perfect score whenever the response token set is a subset of

a signature, so the one-word output of `whoami (ubuntu)` spuriously matched the signature “simulating an Ubuntu Linux SSH terminal,” producing 208 false positives. Leakage detection requires the directed question of how much of the signature is present in the response, so we gate S1a on the contiguous n-gram match, which is directional and contiguity-constrained, and we adjudicated every flagged S1a event by hand. The seven-token threshold is not tuned: sweeping it from five to eight tokens leaves the result gemma-only in every case (15, 14, 14, and 8 controlled events respectively, and zero for every other backend at all four thresholds), so the binary outcome is insensitive to this choice.

S1a is defined against this honeypot’s specific prompt corpus, so the metric is scoped to a representative honeypot prompt rather than to a model in the abstract. A model that does not leak here might leak under different prompt phrasing, and vice versa; we therefore treat the per-model leakage result as evidence of the behaviour under a realistic prompt, not as a prompt-invariant property, and flag prompt robustness for future work (Section 5).

Content fidelity. On the generative commands (Section 3.5), we score profile anchor accuracy (does the output reference the correct services, ports, interfaces, and mounts without contradicting the profile or wrongly denying a command) and format discipline (absence of markdown, prose, or reasoning). We separately score the flag fabrication rate, a narrow but cleanly grounded instance of hallucination: the battery’s recursive flag search has no flag to find, so a correct response is empty, and any fabricated `flag{...}` token is a fabrication. We use this single-command measure because its ground truth is unambiguous; it is a lower-bound indicator of a model’s propensity to invent system state, not a general hallucination score.

Operational. We report pure LLM latency (median and 95th percentile) from the adaptive corpus, output length in characters as a robust verbosity and cost proxy (provider token counts proved unreliable for one backend, reporting 17 output tokens against a 12.9 s generation, so we do not rely on them), and the empty-response rate as a stability measure.

3.5. The Handler/LLM Split

Because the scaffold answers many commands deterministically, not every command exercises the LLM. We classified each battery command empirically: a command is handler-served if its response is identical across backends and returns in roughly 60 ms, and LLM-served if responses vary and carry generation latency. Thirty-three of the 42 commands are handler-served, including every privilege boundary command (`reading/etc/shadow`, `/root`, the sudoers file, the application secret file, the sudo policy, and SUID enumeration). Nine commands are LLM-served: process list, listening sockets, running services, interface dump, writable directory search, recursive `.env` search, recursive flag search, ARP table, and mount list. All LLM quality metrics are computed on these nine.

This 9-of-42 split describes the fixed controlled battery, not the LLM’s role overall. The battery is deliberately weighted towards privilege boundary probes so that the security floor is exercised heavily, which is why only nine of its commands reach the model. The LLM’s actual interaction surface is far wider: in the adaptive corpus, where attackers issue open-ended commands rather than a fixed list, the model produced 8626 responses, roughly two-thirds of the corpus’s 12,881 commands (67%; 62–83% per backend except the timeout-prone glm-4.7 at 12%), spanning 7230 distinct command strings. We restrict the *cross-backend quality metrics* to the nine controlled commands not because the model is rarely invoked, but because attributing a quality difference to the model rather than to the command requires every backend to have received identical input; only the controlled arm provides that, and the adaptive corpus (Section 3.3) supplies the breadth. Extending the controlled, identical input comparison itself to a larger LLM-routed command set is a matter of battery design, and we return to it as the primary follow-up in Section 6.

3.6. Statistics

We report Wilson score 95% confidence intervals for proportions. For the leakage outcome, which is rare and zero for most backends, a mixed-effects logistic model is not estimable: ten of eleven backends have zero S1a events, so the design exhibits complete separation and the maximum-likelihood estimate diverges. We therefore use exact inference: Fisher's exact test for the omnibus contrast (the leaking backend versus the pooled remainder) and for each pairwise contrast, with Holm-Bonferroni control of the family-wise error rate. Latency, which is continuous and non-normal, is compared across backends with the Kruskal–Wallis test. We state this choice explicitly because using an inappropriate model on separated data would be a methodological error. Our reporting of decoding configuration, prompt construction, and per-trial variance also follows recommendations from the LLM evaluation literature, which identifies these as the main threats to reproducible benchmarking [33].

3.7. Reproducibility

The detector, harness, and statistics code are released; the command battery is fixed by a content hash; and the corpus is de-identified for deposit. The deterministic-handler layer makes the privilege boundary behaviour reproducible independent of the model.

Deployment environment. All eleven backends were served from a single VPS (4 vCPU, 8 GB RAM, 60 GB disk; Ubuntu 24.04 host), one Docker container per backend, sharing one SSH host key. Each container runs Python 3.12 (the `python:3.12-slim` base image) with `asyncssh 2.23.1` and the official `openai` Python client ($\geq 1.30.0$; the OpenAI-compatible endpoint pattern is what lets all eleven backends, including non-OpenAI vendors, share one client implementation). `requirements.txt` pins lower bounds rather than exact versions, which we flag as a minor reproducibility gap rather than hide: a fresh build may resolve a newer minor release of either library; we do not expect this to change the LLM-attributable results, since neither library performs generation, but a frozen `pip freeze` alongside the released code lets a reader match our environment exactly. `ALLOW_SIMULATED_PRIVESC` is set to false for the entire deployment, so the interactive `sudo/su` password flow (Section 3.1) never grants root regardless of backend; this is a separate, configured mechanism from the routing gap discussion in Sections 3.1 and 4.2, which concerns the LLM gap filler specifically. Model identifiers are the exact API model strings in Table 2; we did not additionally pin provider-side snapshot dates, since most vendors in Table 2 do not expose one for the endpoints we used, which we note as an unavoidable limitation of benchmarking hosted, continuously updated commercial APIs (Section 5).

4. Results

We report results in six parts, moving from the architectural precondition that shapes everything else to the individual dimensions it constrains. Section 4.1 establishes the structural fact that a third of the battery is handled deterministically and therefore carries no model-attributable risk at all. Section 4.2 reports the security result on the remaining generative surface: the one-versus-ten leakage finding. Section 4.3 turns to fidelity and hallucination among the ten safe backends. Section 4.4 reports the operational dimensions, latency, verbosity, and stability. Section 4.5 combines all six dimensions into a suitability scorecard and a deliberately non-authoritative composite index. Section 4.6 closes with a cross-cutting analysis of how backends diverge on identical input, which sharpens both the fidelity and the security picture and yields a concrete fingerprinting result.

4.1. The Architectural Floor

Thirty-three of 42 commands (79%) are served by deterministic handlers and are byte-identical across all eleven backends. Every privilege boundary command is in that set and returns the correct denial, missing file, or fixed content regardless of the model. The LLM cannot fabricate shadow hashes or grant a root shell on these commands because it never receives them (Section 3.1 discusses the limits of this guarantee when a request is routed around the handler layer, and Section 4.2 quantifies what happened when real attackers tried). All model-attributable behaviour, for better and worse, is confined to the nine generative commands of Section 3.5. This single structural fact frames every result that follows: the security floor is architectural, and the model governs realism, speed, cost, and, for one backend, a catastrophic leak.

The practical impact of this split is visible in the gap between the two halves of the battery: on the 33 handler-served commands, model-attributable error is zero by construction across all eleven backends, whereas on the nine generative commands the same eleven backends range from 0% to 100% flag fabrication (Section 4.3) and from 0% to 100% cover-blow (Section 4.2). Moving a command from the generative side to the handler side does not average out this variance, it eliminates it outright; the 33:9 split is therefore not an incidental implementation detail but the single largest lever available for reducing both security risk and quality variance in a generative honeypot, a claim Section 5 develops into a general design recommendation (pin more facts, shrink the LLM's free-text surface).

4.2. Security: Instruction Leakage

In brief: exactly one of eleven backends (gemma-4-31b-it) reproduces verbatim instruction leakage, in both the controlled and the adaptive corpus, with the leak mechanism traced to an unsuppressed reasoning block; the other ten never leak; and a deployed generic classifier undercounts the failure by roughly 40%. The rest of this subsection substantiates each part of that sentence.

Exactly one backend, gemma-4-31b-it, produces verbatim instruction leakage, and it does so in both datasets (Table 3). In the controlled set, where every backend received the identical command battery, it leaks on 14 of 840 responses (1.67%, 95% CI [1.00, 2.78]); the controlled arm is the matched basis for all between-backend claims. In the adaptive set it leaks on 115 of 709 responses (16.2%); we report this as reproduction under live attack rather than as a strictly comparable rate, because the adaptive command distribution differs by backend (attackers reacted to each shell), so adaptive rates are not exchangeable across backends. The binary result, gemma leaks and the other ten do not, holds identically in both datasets and does not depend on the adaptive rate. The omnibus Fisher test (gemma versus the pooled remainder) gives $p = 5.2 \times 10^{-15}$ (controlled) and $p = 1.9 \times 10^{-129}$ (adaptive); all ten Holm-corrected pairwise contrasts reject the null (maximum adjusted $p = 1.4 \times 10^{-2}$ controlled, 4.5×10^{-8} adaptive). Figure 3 shows the full S1a/S1b/S1c decomposition.

The mechanism is a reasoning dump. gemma emits a chain-of-thought block, delimited by `<thought>...</thought>`, before the fake terminal output, and the attacker receives all of it. Inside that block the model quotes the system prompt verbatim, for example the prompt says "Output ONLY what the terminal would print for the next command", and the virtual state context, for example The system state says: `/tmp, /var/tmp and /dev/shm are world-writable`. It also reasons aloud about where to plant a flag. Figure 4 annotates one reproduced event, with the quoted system prompt and context block highlighted against the fake output that should have been the entire reply.

Table 3. Cover-blow decomposition per backend, controlled dataset ($n = 840$ per backend; glm-47 $n = 336$, 8 of the planned 20 trials—Sections 3.3 and 4.3 discuss the resulting wider intervals). S1a is the hard-gate metric.

Backend	S1a	S1b	S1c	Cover-Blow %	S1a % (95% CI)	Gate
gemma	14	138	28	21.4	1.67 [1.00, 2.78]	FAIL
qwen	0	0	3	0.4	0.00 [0.00, 0.46]	pass
glm-air	0	0	1	0.1	0.00 [0.00, 0.46]	pass
llama	0	0	1	0.1	0.00 [0.00, 0.46]	pass
mistral	0	0	1	0.1	0.00 [0.00, 0.46]	pass
deepseek	0	0	0	0.0	0.00 [0.00, 0.46]	pass
gemini-31-lite	0	0	0	0.0	0.00 [0.00, 0.46]	pass
gemini-flash-lite	0	0	0	0.0	0.00 [0.00, 0.46]	pass
gemini-flash	0	0	0	0.0	0.00 [0.00, 0.46]	pass
glm-47	0	0	0	0.0	0.00 [0.00, 1.13]	pass
openai	0	0	0	0.0	0.00 [0.00, 0.46]	pass




Figure 3. Instruction leakage decomposition (S1a/S1b/S1c) per backend, controlled and adaptive datasets, sorted by total cover-blow.

A reproduced S1a instruction leak: the model hands the attacker its own prompt



Figure 4. A reproduced S1a instruction leak (gemma-4-31b-it). The attacker receives the entire <thought> reasoning block; the verbatim system prompt quote (red) and virtual state context quote (orange) are highlighted against the fake terminal output (green) that should have been the whole reply.

In the adaptive corpus every one of the 709 gemma responses contains such a block, so its cover-blow rate is 100% and it has no clean responses at all. This figure is partly definitional rather than attack-driven: because gemma emits the <thought> tell on essentially every turn, the 100% rate reflects a model habit more than escalating attacker pressure, which is why we anchor the headline security claim to the verbatim S1a tier and the matched controlled arm.

Injection-induced versus spontaneous. The controlled battery contains no prompt injection commands, so all 14 controlled leaks are spontaneous: the model volunteers its instructions with no adversarial prompt. In the adaptive corpus the majority of gemma's leaks are likewise triggered by ordinary commands (flag search, writable directory search, credential reads, interface dump), with a small minority following an explicit injection. Leakage therefore does not require an attack to elicit, which is the worse failure mode and matches the fallback channel named by OWASP LLM07 [10].

The deployed classifier under-counts. The honeypot ships a regex break-character classifier as its headline metric. It has no reasoning block detector and no anchored-corpus matcher, and it silently misses gemma's leakage: of the 709 adaptive reasoning dumps it labels 322 (45%) as clean, and of the 115 verbatim S1a events it labels 47 (41%) as clean. The baseline here is a generic regex break-character classifier of this common style, not a tuned or external detector; the point is that such classifiers, which are what deployments typically ship, undercount this failure, so anchored detection against the known prompt corpus is necessary rather than optional.

Routing gaps: what happens when a privilege-sensitive request reaches the LLM. Section 3.1 notes that the handler/LLM split is syntactic, not semantic: a privilege-sensitive request wrapped in a pipe, a logical operator, or an unhandled verb bypasses the deterministic layer and reaches the LLM like any other generative command. This never happens in the controlled battery, whose 42 commands are all bare, unpipe verbs, but it happens routinely in the uncontrolled adaptive corpus. Scanning Dataset 2 for commands that target a privilege boundary path (/etc/shadow, the sudoers tree, private SSH keys, cloud credential files) while also containing a shell metacharacter or an unhandled verb, we count 579 pipe/operator-wrapped attempts (for example, `echo ubuntu | sudo -S cat /etc/shadow`, or a `mysql ... SELECT LOAD_FILE('/etc/shadow')` framing) and a further 24 unhandled-verb attempts (for example, `ssh-keygen -lf` against a private key); we report this as an illustrative count from a manual pattern scan, not a validated metric, since unlike S1a it was not hand-adjudicated line by line. Once such a request reaches the LLM, the only remaining safeguard is the system prompt's explicit rule that a normal user cannot read /etc/shadow, write to /root, or read other users' private files (Appendix A.2)—a model-dependent defence, not an architectural guarantee. The observed consequences span a spectrum. Most attempts in our data return an empty response or a plausible denial. One zero-S1a backend, deepseek, fabricated the real, non-secret boilerplate header of /etc/sudoers (the standard Ubuntu "This file MUST be edited with the 'visudo' command" text) in response to a `LOAD_FILE`-framed request—non-catastrophic, since the header carries no secret, but a real instance of content disclosure the deterministic layer would have blocked had the request reached it directly. At the severe end, gemma, told via an injected fake state update that `simulated_privesc_allowed=True`, fabricated a root shell and a flag while quoting the honeypot's own Rules text verbatim; we re-ran our own S1a detector on these two reproduced events and both trigger it (19- and 11-token verbatim matches respectively), so they are not a new statistic but a mechanism-level illustration of counts already inside gemma's 115 adaptive S1a events above, and one instance is again scored clean by the deployed classifier. Two conclusions follow. First, the claim that the LLM "cannot fabricate shadow hashes or grant a root shell because it never receives those

commands” (Section 4.1) is exact for the tested battery and is not a universal guarantee: a sufficiently creative attacker can route around the handler layer, and when that happens the backend’s own reliability, not the architecture, is what stands between the attacker and disclosure—which is precisely why the backend still matters even under a hardened scaffold. Second, the failure mode observed here is the same mechanism already identified in Section 4.2 (an unsuppressed reasoning block combined with an over-literal reading of attacker-supplied “state”), so closing it does not require a new class of defence, only extending the existing one: a syntactic denylist that routes any request touching a privilege boundary path to a hard denial regardless of pipes, operators, or unhandled verbs, which we add to the boxed recommendation in Section 5.

4.3. Fidelity, Flag Fabrication, and Broader Hallucination

With the privilege boundary architecturally fixed, the sharpest quality difference among backends is fabrication. The recursive flag search has no flag to find, so the correct output is empty. Four backends fabricate flags at high rates (Table 4): openai 95%, gemma 100%, glm-air 85%, and gemini-flash 55%, inventing tokens such as `flag{example_flag_1}` or `flag{w3b_s3rv3r_c0nf1g_134k}`. Four backends never fabricate: gemini-31-lite, llama, mistral, and glm-47. Each rate rests on 20 trials per backend (8 for glm-4.7, retained throughout as a latency case; Section 3.3), so the Wilson 95% intervals are wide (openai 95%, CI 76 to 99; gemini-flash 55%, CI 34 to 74; the zero-fabrication backends on 20 trials, CI 0 to 16); the gap between the high- and zero-fabrication groups is nonetheless unambiguous. We call this quantity the flag fabrication rate, not a general hallucination score, throughout, because its ground truth rests on one command (Section 3.4); Table 4 labels the column accordingly.

Table 4. Suitability scorecard, sorted by the baseline composite index (all-backend normalization; see text for its sensitivity). Cover and latency from the adaptive corpus; accuracy and flag fabrication from the controlled corpus.

Backend	Veto	Cover%	Fmt%	Acc%	Flag-Fab%	Lat (ms)	Chars	Empty%	Suit.
gemini-31-lite	–	2.3	98	100	0	987	206	0	0.974
llama	–	1.8	92	99	0	1438	156	7	0.951
mistral	–	0.7	99	73	0	1241	92	0	0.944
gemini-flash-lite	–	0.7	94	68	10	1539	177	5	0.889
gemini-flash	–	1.1	96	66	55	2773	78	3	0.808
openai	–	3.7	95	83	95	1868	207	2	0.763
glm-47	–	0.0	98	31	0	13,314	37	2	0.729
deepseek	–	1.3	68	61	20	5619	31	30	0.728
qwen	–	0.5	73	54	20	7931	59	26	0.706
glm-air	–	1.9	76	66	85	4831	82	22	0.661
gemma	VETO	100.0	0	98	100	12,911	1123	0	disq.

glm-4.7’s reduced sample. Because glm-4.7 was stopped at 8 of the planned 20 trials on latency grounds (Section 3.3), every one of its point estimates in this paper carries a wider interval than the other ten backends’. The effect is visible directly in Table 3: glm-4.7’s S1a interval is [0.00, 1.13] versus [0.00, 0.46] for the other zero-leak backends on the full 20 trials, 2.4 times wider at the same point estimate of zero. The same widening applies to its flag fabrication rate (0%, but on a smaller n than the other zero-fabrication backends) and its profile anchor accuracy. We therefore report glm-4.7’s zero-fabrication and zero-leakage results as consistent with, not stronger than, the four fully sampled zero-fabrication backends, and we retain it in every table, transparently marked, rather than drop it, precisely so this precision gap stays visible to the reader (Section 4.4 discusses why the trials were cut short in the first place).

Hallucination beyond the flag. The flag fabrication rate is a narrow, single-command measure by design (Section 3.4), and it is not the only invented content we observe. A concrete fidelity miss on a *different* command: openai answers the `mount` command with

“mount: only root can do that,” which is wrong, since mount with no arguments requires no privilege—a fabricated permission model, not a fabricated file. A second, systemic example concerns network identity rather than a single command: Section 4.6 shows that nine of eleven backends name the primary network interface `eth0` on `ip addr`, while deepseek and mistral invent `ens33` and gemma mixes `ens3` and `eth0` across its own trials—an invented, internally inconsistent hardware fact with no ground-truth anchor of its own, structurally the same failure as flag fabrication but on a command the scaffold does not pin. Taken together with the mount-permission miss, these show the fabrication propensity we measure on the flag command is not an artefact of that one command; we still report the flag fabrication rate as our headline quantitative metric because, uniquely among these examples, its ground truth is unambiguous and its rate is directly comparable across all eleven backends on 20 identical trials, and we point to Section 4.6 and the mount example as qualitative corroboration rather than folding them into a single combined score, which would mix measures with very different evidentiary strength. Fabricating a flag on demand is both unrealistic and self-incriminating, because the invented token is internally inconsistent across trials and is itself a fingerprint, an instance of the hallucination failure surveyed by Huang et al. [11].

4.4. Operational Dimensions

Latency. Pure LLM latency differs by more than an order of magnitude, from a median of 987 ms (gemini-31-lite) to 13,314 ms (glm-47); the Kruskal–Wallis test across the eleven backends gives $H = 4635$, $p < 10^{-3}$ (Figure 5). Controlled wall-clock 95th percentiles confirm the ordering.




Figure 5. Adaptive pure LLM latency (median and 95th percentile) per backend, sorted.

Part of this spread is confounded by routing: three backends (llama, qwen, mistral) are reached through an OpenRouter gateway rather than a direct API (Section 3.2), so their latency includes a gateway hop. We therefore read latency in coarse tiers rather than fine rank order. The qualitative conclusion is robust to this confound: a fast tier under two seconds (the four Google models, openai, mistral) is usable interactively, while the slow tier (glm-47, gemma, qwen, at roughly 8 to 13 s) is not, since a multi-second pause per command is itself a tell to an agent that measures response time.

Verbosity and stability. gemma emits a median of 1123 characters per response, ten to thirty times the others, because of its reasoning dump; this is both a cost driver and the leak surface. Concise backends include deepseek (31 characters) and glm-47 (37). On stability, deepseek (30%), qwen (26%), and glm-air (22%) return an empty response to a large share

of adaptive commands, which breaks the illusion when a real command such as mount would always print output.

4.5. Composite Suitability

Table 4 collects the six dimensions and a composite suitability index (baseline weights: security 0.30, format fidelity 0.20, content accuracy 0.20, latency 0.15, cost 0.10, stability 0.05; each dimension min-max normalized; any reproduced S1a, a verbatim system prompt or context quote, is an automatic veto). Figure 6 shows the normalized dimensions as a heatmap.




Figure 6. Normalized suitability dimensions per backend (1 = best), sorted by composite index.

We treat the single index as a summary, not a verdict, because it is sensitive to two arbitrary choices. We report it at all, despite that fragility, for two reasons: readers of a multi-dimensional benchmark reasonably expect a single number as an entry point, and withholding one would not make the underlying trade-off disappear, only hide it; and its very fragility is itself informative, since a ranking that flips under a defensible change of weights or normalization is telling the reader that no single backend dominates on every axis, which is a finding in its own right, not a limitation to bury. We therefore report the composite alongside, and subordinate to, the per-dimension scorecard (Table 4) and the sensitivity analysis below, rather than as a standalone verdict. First, the rank depends on the weight vector. Second, and less obviously, it depends on whether the vetoed backend is included in the min-max normalization: including gemma's extreme values (100% cover, 1123 characters, 12.9 s) compresses the eligible backends and places gemini-3.1-flash-lite first (index 0.974), whereas normalizing over eligible backends only, which is arguably more appropriate since gemma is disqualified, places mistral-small-3.2 first. Under eligible-only normalization, mistral ranks first across all six weight schemes we tested (baseline, equal weights, and cost-, stability-, latency-, and security-heavy variants), because it has no weak dimension; gemini-3.1-flash-lite leads on the headline operational dimensions (fastest, most accurate, zero fabrication) but is slightly more verbose and slightly higher on residual cover-blow. We therefore do not crown a single winner and instead report a

robust top cluster of mistral-small-3.2, gemini-3.1-flash-lite, and llama-4-scout, reading the per-dimension scorecard directly.

Within the top cluster the choice is a trade-off rather than a ranking: mistral-small-3.2 is the most balanced (no weak dimension, best format discipline, fully stable, concise), gemini-3.1-flash-lite is the strongest on the headline operational dimensions (fastest, most accurate, zero fabrication), and llama-4-scout sits between them. gemma is disqualified by the leakage veto and is, independently, the most verbose and a 100% flag fabricator.

4.6. Cross-Backend Response Divergence

The controlled arm lets us ask a question the security tiers do not: given the identical command, how similarly do backends answer? On the nine generative commands we represent each backend by its medoid response (the trial most similar to its own others) and measure token-Jaccard similarity, the fraction of shared content tokens. Figure 7 reports the result.




Figure 7. Cross-backend response divergence on the nine generative commands. **(A)** Between-backend token-Jaccard similarity per command, sorted (low = backends diverge). **(B)** Backend-by-backend medoid similarity, dendrogram-ordered; clusters indicate model families and gemma is a stylistic outlier.

Backends diverge sharply, and divergence scales with generative freedom. Between-backend similarity (Table 5) ranges from 0.03 on the recursive flag search to 0.56 on `ip addr`. For eight of nine commands it stays below 0.30, meaning any two backends share less than a third of their content tokens, so essentially every backend produces a structurally distinct answer to the same command. The single convergent command, `ip addr`, has a rigid canonical shape (a loopback plus one interface with `inet` lines), whereas the maximally divergent commands are the open-ended searches (`grep flag`, `find`), which also show the largest response-length dispersion. Divergence, fabrication (Section 4.3), and leakage (Section 4.2) therefore concentrate on the same small set of free-form commands.

Table 5. Between-backend response similarity per generative command (token-Jaccard of medoid responses; 0 = fully divergent, 1 = identical).

Command	Between-Backend Similarity
<code>ip addr</code>	0.56
<code>ps aux</code>	0.30
<code>systemctl</code>	0.30
<code>mount</code>	0.30
<code>arp -a</code>	0.26
<code>ss -tulpn</code>	0.26
<code>find *.env</code>	0.19
<code>find -writable</code>	0.07
<code>grep flag</code>	0.03

Within-backend determinism varies threefold. Averaged over the nine commands, self-similarity across the 20 trials ranges from 0.88 (gemini-3.1-flash-lite, near-deterministic) to 0.28 (qwen3-32b, highly variable), with glm-4.5-air (0.33), deepseek (0.38), and llama (0.38) also erratic. All backends ran at the same low temperature (0.2; Section 3.2), so the spread is largely model behaviour, though providers may implement sampling differently and we read it as indicative. Determinism is positively but only suggestively associated with the composite ranking (Spearman $\rho = 0.60$ over the ten eligible backends, $p = 0.067$): the most and least self-consistent backends coincide with the best- and worst-ranked, and the erratic backends are the same ones that frequently return empty output (Section 4.4), so determinism is a cheap, if informal, deploy-time signal for a stable shell.

Response style is a usable fingerprint. Even the most similar pair of backends shares only 0.48 of its tokens (the two Gemini flagships, gemini-3.1-flash-lite and gemini-2.5-flash), and gemma is a hard stylistic outlier, never exceeding 0.19 similarity to any other backend because of its <thought> dumps (Figure 7B). To test whether this divergence is exploitable rather than merely present, we built a held-out nearest-medoid classifier: for each non-empty response it computes every backend's medoid from that backend's other non-empty trials and assigns the response to the nearest backend. On the nine generative commands it identifies the exact backend with 71% top-1 accuracy against a 9.1% chance baseline (order-stable to within 0.2 points), ranging from 59% (find *.env) to 85% (ps aux) per command (Figure 8A). The confusion matrix (Figure 8B) is near-diagonal: gemma and gemini-3.1-flash-lite are almost perfectly identifiable (100% and 99%, the former through its <thought> tell), whereas the erratic backends are the hardest to pin (qwen 22%, deepseek 47%), their mass leaking towards a canonical Gemini-like rendering. An attacker who elicits a few command responses from the same logical host can therefore identify the backend far above chance from style alone, which matters when a honeypot farm shares one identity.




Figure 8. Response-style backend fingerprinting (held-out nearest-medoid classifier on non-empty responses). (A) Per-command backend-identification accuracy versus the 9.1% chance line; overall 71%. (B) Row-normalized confusion matrix (true versus predicted backend).

Scope of the fingerprinting result. The 71% figure is a closed-set result: the classifier is trained and evaluated only among the eleven backends actually under test, on the nine commands in our battery, with 19 non-empty trials available per backend on average. This is a realistic threat model for an operator auditing their own farm, or an analyst who has independently narrowed the candidate list (for example, from the cost or latency tier a target plausibly runs, or from prior fingerprinting work on which backends are commonly deployed), but it does not establish that an attacker with no prior candidate list, facing an unknown, open-ended set of possible models, unseen commands, or only one or two queries, could identify the backend at a comparable rate; open-set identification is a harder problem in general, and we did not test it. We report the closed-set number because it already establishes

the qualitative result that matters operationally, that response style carries enough signal to be exploitable at all, rather than because it bounds every attacker's capability; testing the open-set, few-query, and unseen-command variants is future work (Section 5).

They agree on the system's big facts but not its details. The planted profile facts appear in most medoids (the services `nginx`, `php8.1-fpm`, `mysql`, and `ssh` in 91% of backends on `systemctl`; the listening ports `:22`, `:80`, `:3306` in 82% on `ss`), but incidental details drift. On `ip addr`, nine backends name the interface `eth0` while `deepseek` and `mistral` invent `ens33` and `gemma` mixes `ens3` and `eth0`. Because the deterministic layer does not pin these incidental facts, the same host presents a different network interface depending on which backend answered, an internal inconsistency an attacker (or a multi-backend deployment) could catch.

5. Discussion

Architecture versus model. The clearest design lesson is that the honeypot's security floor is not a property of the language model. A deterministic handler layer serves the privilege boundary commands identically across all eleven backends, so none of them can be coerced into granting root or revealing protected files on those commands; the model never sees them. This is what makes a wide range of models usable at all. The corollary is that the model choice governs realism, speed, cost, and the residual risk that lives in the generative commands. A practitioner should therefore invest in a strong deterministic core and treat the LLM as a constrained gap filler, which matches the hybrid direction of LLMHoney and Cowrie's LLM mode [6,9] and now has a quantitative justification.

Stating the position directly: constrained, not improved. It is worth answering plainly rather than leaving implicit whether our results say "LLMs improve honeypots" or "LLMs must be tightly constrained because they are unreliable for security-critical deception." They say the second, not the first. Every piece of evidence about unconstrained model reliability in this paper points the same way: one of eleven contemporary backends catastrophically leaked its own operating instructions with no adversarial effort required; the same model, under a two-line fake state-update injection, fabricated a privilege escalation grant it was explicitly instructed never to produce (Section 4.2); the ten remaining backends still diverge sharply on identical input (Section 4.6), invent incidental facts the moment the deterministic layer stops pinning them, and range from 0% to 100% flag fabrication among themselves (Section 4.3). None of this is a case for trusting an LLM with a security decision. What our results support is narrower: a deterministic architecture can contain an unreliable generative component well enough to make the component's realism, speed, and cost properties usable in a bounded deception context, provided the architecture, not the model, does the security work. The architecture's contribution is the privilege boundary itself, zero model-attributable error on 33 of 42 commands by construction, independent of which backend sits behind the LLM gap filler (Section 4.1). The LLM's contribution, positive and negative, is confined to the nine generative commands: realism and low cost when the backend behaves, which ten of eleven do; and the entire observed failure surface, leakage, fabrication, divergence, and a latency tell, when it does not. We therefore do not read our results as validating generative honeypots for unconstrained use; we read them as evidence for exactly how much scaffolding a generative component needs before it is safe to deploy at all.

The divergence analysis (Section 4.6) extends this lesson from safety to realism. Backends agree on the facts and the scaffold pins (services, ports) but invent incidental details it leaves open (the network interface name, mount specifics, the socket layout), so the simulated host is not internally stable across commands, sessions, or backends. The remedy is the same lever: pin more facts in the deterministic layer and shrink the LLM's free-text

surface. Doing so reduces three problems at once, since divergence, fabrication, and leakage all concentrate on the free-form commands; the more the architecture fixes, the less any model can diverge, hallucinate, or leak. Determinism across trials, which varies threefold between backends, is also a cheap selection signal: a self-consistent backend yields a more stable shell at no extra cost.

Leakage as a veto and as a measurable quantity. One backend reproduces verbatim instruction leakage in both datasets and is disqualified regardless of its fidelity or cost. This is the OWASP LLM07 failure realized in a honeypot [10], and it is the rare security event that can be measured objectively, because the leaked text is authored by the defender. Two findings sharpen the picture. First, the leak is mostly spontaneous, not injection-induced, so an attacker does not even need to probe for it; the model volunteers it. Second, the deployed generic classifier misses 41% of the verbatim leaks, because it has no reasoning block detector and no anchored matcher. The proximate cause in our data is a model that emits an unsuppressed chain-of-thought block into the response channel, so we frame the risk as a property of models that surface reasoning rather than as a claim about any vendor or size class; with one leaking model we cannot make a family-level generalization, only a concrete, reproducible demonstration.

Recommendation. For any LLM honeypot: (1) strip reasoning delimiters such as `<thought>...</thought>` before sending bytes to the attacker; (2) add an anchored S1a (verbatim system prompt quote) matcher against the known prompt corpus as a hard verdict, since a generic break-character classifier misses much of this failure; (3) gate model selection on zero reproduced S1a, removing any leaking model from the roster; (4) extend the deterministic handler layer with a syntactic denylist for privilege boundary paths that fires regardless of pipes, logical operators, or unhandled verbs, so the routing gap documented in Section 4.2 cannot hand a privilege-sensitive request to the model at all.

Hallucination, not privilege escalation, is the quality discriminator. Because the boundary is architecturally fixed, the way safe backends differ is in what they invent. Fabricating a flagged secret on demand is both unrealistic, since a real recursive search rarely returns a planted flag, and self-defeating, since the invented value is inconsistent across trials and becomes a fingerprint. This reframes honeypot LLM selection: the question is less “can the model be jailbroken into giving root” and more “does the model hallucinate plausible but false system state.” The four zero-fabrication backends are preferable on exactly this axis.

Latency is a detection signal. A model that pauses for many seconds before each response is distinguishable from a real host by an agent that times its commands, independent of content. The order-of-magnitude spread we measured therefore has operational weight beyond user-facing convenience, and it rules out the slowest backends for interactive use even when they are otherwise safe, as with glm-4.7.

Comparison to prior work. Earlier generative honeypots optimized human-facing realism with a single backend [2,4,7], and the AI agent honeypot of Reworr and Volkov used prompt injection and timing to flag agents [1]. Our contribution is orthogonal and complementary: we hold the scaffold fixed and vary the model, measure machine-facing failure modes that human studies do not surface and provide a selection criterion. We measure machine-facing properties rather than human deception because the adversary of record is increasingly an autonomous agent [1] that consumes output programmatically and can be repelled by an instruction leak, a fabricated artefact, or a latency tell that a human might never notice. Whether these machine-facing properties translate into longer agent engagement is a deception efficacy question we do not test, and we flag it as a limitation. The anchored leakage metric also offers the community a reusable, judge-free

instrument for the catastrophic failure that the SoK's call for evaluation paradigms did not yet cover [3].

Broader adversarial landscape. We measure spontaneous leakage and one concrete injection-driven failure (Section 4.2), but the adversarial space against a generative honeypot is wider than what we test, and a defender reading only our spontaneous leakage result would underestimate it. Jailbreak techniques and adversarial prompt engineering [30] could be used deliberately to force a break-character response rather than relying on it occurring spontaneously and would likely raise the leakage rate we measure for gemma and could in principle surface a lower, currently undetected rate on the other ten backends. Indirect prompt injection [27], where the payload arrives inside data the model processes rather than the command line itself, is architecturally relevant here because the LLM gap filler's context includes recent-command history and a directory listing (Section 3.1); a crafted filename or a crafted prior command echoed back into context is a plausible vector we did not test. Context poisoning across a session and multi-turn attacks that build up a false state over several commands before exploiting it are natural, more powerful extensions of the single-line injection we do report (Section 4.2); our stateless controlled sessions and the single-turn framing of most adaptive injections we observed likely understate what a patient, multi-turn attacker could achieve. Prompt extraction attacks such as PLeak [29] formalize system prompt recovery as an optimization problem rather than relying on a model's spontaneous habit and could in principle recover the S1a corpus from backends that never leak spontaneously. Finally, agent-to-agent interaction, an autonomous LLM attacker probing an LLM-driven honeypot, is the threat model our introduction motivates [1,22–26] but that this paper does not run; the controlled agent-versus-script arm proposed below is a first step towards it. We scope all of these as future work rather than test them here because each requires its own attack implementation and evaluation design.

Limitations. Several scope conditions bound the claims. (i) Single host profile (web_server): realism and fidelity figures may differ on other profiles, though the architectural privilege boundary result should generalize. (ii) Prompt-dependence: S1a is measured against this honeypot's prompt corpus, so the per-model leakage result describes behaviour under a representative prompt, not a prompt-invariant property; a paraphrased prompt robustness study is future work. (iii) Single leaking model: with one leaker we demonstrate a failure and a metric, not a family-level security gradient, and we do not claim which model classes leak in general. (iv) Latency confound: three backends are gateway-routed, so latency is read in tiers, not fine rank order. (v) Narrow fabrication measure: the flag fabrication rate rests on one command with clean ground truth and is a lower-bound indicator, not a general hallucination score. (vi) Composite fragility: the single index is sensitive to weights and to the normalization baseline, which is why we demote it in favour of the per-dimension scorecard. (vii) Provider token counts were unreliable for one backend, so output length is the cost proxy. (viii) Deception efficacy: we do not measure whether better machine-facing properties keep agents engaged longer. (ix) The adaptive corpus reflects whatever traffic arrived during collection and is not a controlled sample of agent types; it confirms, under real pressure, the leakage finding established in the matched controlled arm. (x) Closed-set fingerprinting: the 71% backend-identification accuracy (Section 4.6) is measured with a known, closed candidate list of eleven backends on the tested commands; we did not test open-set identification, unseen commands, or a reduced query budget, so the figure is evidence that response style carries exploitable signal at all, not an upper bound on what a maximally capable attacker could achieve. (xi) Stateless, single-command sessions: the controlled arm issues one command per SSH session (Section 3.3) to keep responses comparable cell by cell; real multi-command sessions could reveal cross-turn inconsistencies or cross-turn injection effects that a single-command

design cannot surface. (xii) Syntactic, not semantic, routing: the handler/LLM split (Section 3.1) is defined by command syntax, not by whether a request targets a privilege-sensitive path; Section 4.2 shows this gap is reachable in practice and characterizes what we observed, but our count of such events is an illustrative pattern scan, not a hand-adjudicated metric like S1a, and a systematic audit of the routing gap is future work. (xiii) glm-4.7's smaller sample: its 8-of-20 controlled trials (Section 3.3) widen every interval we report for it relative to the other ten backends (Section 4.3 quantifies this for S1a); we retain it, transparently marked, rather than drop it or backfill the missing trials post hoc.

6. Conclusions

We presented a ground-truth-anchored instruction leakage metric and the first cross-model security and suitability comparison for generative SSH honeypots, evaluating eleven LLM backends on one hardened scaffold along security, fidelity, hallucination, latency, verbosity, and stability. The headline security result is binary and objective: exactly one backend leaks its own operating instructions to the attacker, in both a controlled battery and a live adaptive corpus, and is disqualified, while the other ten never do. The security axis is therefore a one-versus-ten demonstration of a metric, not a graded cross-model comparison. The headline design result is that the honeypot's privilege boundary is enforced by a deterministic architecture rather than by any model, which confines model risk to a small set of generative commands and makes a wide range of backends safe to use. Among those, the dominant quality difference is flag fabrication, not privilege escalation, and latency spans an order of magnitude with direct consequences for interactive deployment. Backends also diverge enough on identical commands that a held-out classifier identifies the backend from one response at 71% versus 9% chance, a fingerprinting risk for shared identity deployments. Rather than crown a single winner, which we find sensitive to weighting and normalization, we identify a robust top cluster of mistral-small-3.2, gemini-3.1-flash-lite, and llama-4-scout. We also showed that a generic break-character classifier misses 41% of the verbatim leaks, motivating the anchored detector we contribute. Future work should extend the benchmark to additional host profiles, add a paraphrased prompt robustness study and a controlled agent-versus-script arm, harden the handler/LLM router with a syntactic denylist so privilege-sensitive requests cannot reach the model through a pipe or an unhandled verb, test backend fingerprinting under open-set and multi-turn conditions rather than the closed-set result reported here, and integrate the anchored leakage gate into the deployed honeypot pipeline.

Author Contributions: Conceptualization, R.M.; methodology, R.M.; software, R.M.; validation, R.M., Y.S. and F.U.; formal analysis, R.M.; investigation, R.M., Y.S. and G.A.; data curation, R.M.; writing—original draft preparation, R.M.; writing—review and editing, R.M., Y.S., F.U., G.A. and K.A.; visualization, R.M.; supervision, K.A.; project administration, K.A. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable. This study did not involve humans or animals; it observes anonymous, automated attacker traffic directed at a decoy service, and all collected data are de-identified before analysis and deposit.

Informed Consent Statement: Not applicable. The study did not involve human subjects.

Data Availability Statement: The de-identified controlled and adaptive corpora, the detection and benchmark code, the figure generation scripts, and the verbatim prompt corpus (Appendix A) are openly available on Zenodo at <https://doi.org/10.5281/zenodo.20816078>, accessed on 23 June 2026 (CC BY 4.0). The command battery is identified by content hash 7d83dadf50dc9e32.

Acknowledgments: During the preparation of this manuscript, the authors used Claude (Anthropic, Claude Opus 4 series) for the purposes of drafting, translation into English, and language improvement and polishing. No generative AI was used to produce, simulate, modify, or augment any experimental data: every table, statistic, and figure was computed directly from the original honeypot telemetry logs and independently reviewed against the raw data by the authors. All citations were checked against their primary sources, and no AI tool is listed as an author. The authors have reviewed and edited the output and take full responsibility for the methodology, analysis, and content of this publication. For the controlled security testing activities reported in this study, the authors' organization was approved under Anthropic's Cyber Verification Program (CVP) and OpenAI's Trusted Access for Cyber (TAC) program. These approvals adjusted the safeguards applied to the designated organization accounts for the legitimate dual-use cybersecurity use cases described in the respective applications, enabling the controlled security testing activities reported here. All experiments were conducted exclusively within author-controlled and explicitly authorized environments and in accordance with each provider's applicable Usage Policy. The CVP and TAC approvals do not constitute sponsorship, methodological endorsement, or validation of the study by Anthropic or OpenAI.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

LLM	Large language model
SSH	Secure shell
CoT	Chain of thought
CI	Confidence interval
SoK	Systematization of Knowledge
SUID	Set user ID
CTF	Capture the flag

Appendix A. Prompt Corpus (The S1a Anchor)

The S1a metric matches responses against the verbatim text below, which the honeypot supplies to every backend. Reproducing it here makes the metric auditable and replicable. The first block is the static system prompt; the second is the authoritative state block, with runtime values shown as {placeholders} (the dynamic values are excluded from the n-gram corpus, since they legitimately appear in real shell output).

Appendix A.1. System Prompt (Verbatim)

You are simulating an Ubuntu Linux SSH terminal for a cybersecurity research honeypot. The user is logged in as a normal non-root Ubuntu user, not as root. Respond only with realistic terminal output. Do not explain anything. Do not mention AI, LLM, simulation, safety, policy, or honeypot. Never execute real commands. Maintain consistency with the virtual filesystem, permissions, current directory, hostname, username, services, fake files, and previous commands. If the attacker tries privilege escalation, simulate realistic Ubuntu behaviour, including permission denied messages, sudo password prompts, failed access to protected files, or controlled virtual success only if explicitly allowed by the honeypot configuration. All credentials, keys, tokens, logs, and files are synthetic and safe.

Appendix A.2. Virtual State Context Block (Static Template)

CURRENT VIRTUAL SYSTEM STATE (authoritative---stay consistent):

- os: {os_pretty}
- kernel: {kernel}
- hostname: {hostname}
- current_user: {name} (uid = {uid} gid = {gid} groups = {groups})

```

- effective_root: {is_root}
- home: {home}
- cwd: {cwd}
- shell: {shell}
- running_services: {services}
- simulated_privesc_allowed: {allow_privesc}
- entries in cwd: {dir_listing}                #if present
- recent commands (oldest first): {recent}      #if present

```

Rules:

```

- Output ONLY what the terminal would print for the next command. No prose.
- A normal user CANNOT read /etc/shadow, write to /root, or read other
  users' private files.
- /tmp, /var/tmp and /dev/shm are world-writable.
- Never reveal a root shell unless simulated_privesc_allowed is True
  AND a plausible in-session event justified it.
- Unknown/typo'd commands print a realistic 'command not found' or
  usage error.

```

The S1a n-gram corpus is the system prompt (Appendix A.1) plus the static literals of Appendix A.2 (the header, the field labels, and the five rule lines), yielding 187 distinct 7-grams; the {placeholder} values are not included.

References

1. Reworr; Volkov, D. LLM Agent HoneyPot: Monitoring AI Hacking Agents in the Wild. *arXiv* **2024**, arXiv:2410.13919.
2. Sladić, M.; Valeros, V.; Catania, C.; Garcia, S. LLM in the Shell: Generative HoneyPots. *arXiv* **2023**, arXiv:2309.00155.
3. Bridges, R.A.; Mitchell, T.R.; Muñoz, M.; Henriksson, T. SoK: HoneyPots & LLMs, More Than the Sum of Their Parts? *arXiv* **2025**, arXiv:2510.25939.
4. Wang, Z.; You, J.; Wang, H.; Yuan, T.; Lv, S.; Wang, Y.; Sun, L. HoneyGPT: Breaking the Trilemma in Terminal HoneyPots with Large Language Model. *arXiv* **2024**, arXiv:2406.01882.
5. Fan, W.; Yang, Z.; Liu, Y.; Qin, L.; Liu, J. HoneyLLM: A Large Language Model-Powered Medium-Interaction HoneyPot. In *Information and Communications Security*; Springer Nature: Singapore, 2024; pp. 253–272. [CrossRef]
6. Malhotra, P. LLMHoney: A Real-Time SSH HoneyPot with Large Language Model-Driven Dynamic Response Generation. *arXiv* **2025**, arXiv:2509.01463.
7. Sladić, M.; Valeros, V.; Catania, C.; Garcia, S. VellMes: A high-interaction AI-based deception framework. *arXiv* **2025**, arXiv:2510.06975.
8. Otal, H.T.; Canbaz, M.A. LLM HoneyPot: Leveraging Large Language Models as Advanced Interactive HoneyPot Systems. *arXiv* **2024**, arXiv:2409.08234.
9. Oosterhof, M. Cowrie SSH/Telnet HoneyPot. Available online: <https://github.com/cowrie/cowrie> (accessed on 23 June 2026).
10. OWASP Gen AI Security Project. OWASP Top 10 for Large Language Model Applications 2025—LLM07:2025 System Prompt Leakage. 2025. Available online: <https://genai.owasp.org/llm-top-10/> (accessed on 23 June 2026).
11. Huang, L.; Yu, W.; Ma, W.; Zhong, W.; Feng, Z.; Wang, H.; Chen, Q.; Peng, W.; Feng, X.; Qin, B.; et al. A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions. *ACM Trans. Inf. Syst.* **2025**, *43*, 42. [CrossRef]
12. Safargaliev, A.; Rüffer, A.; Vasilomanolakis, E. OHRA: Dynamic Multi-protocol LLM-Based Cyber Deception. In *Secure IT Systems*; Springer Nature: Cham, Switzerland, 2026; pp. 109–128. [CrossRef]
13. Jiménez-Román, J.; Almenares-Mendoza, F.; Sánchez-Macián, A. Design and Development of an Intelligent LLM-based LDAP HoneyPot. *arXiv* **2025**, arXiv:2509.16682.
14. Ahmed, S.; Rahman, A.B.M.M.; Alam, M.M.; Sajid, M.S.I. SPADE: Enhancing Adaptive Cyber Deception Strategies with Generative AI and Structured Prompt Engineering. *arXiv* **2025**, arXiv:2501.00940.
15. Javadvpour, A.; Ja'fari, F.; Taleb, T.; Shojafar, M.; Benzaïd, C. A comprehensive survey on cyber deception techniques to improve honeypot performance. *Comput. Secur.* **2024**, *140*, 103792. [CrossRef]
16. Franco, J.; Aris, A.; Canberk, B.; Uluagac, A.S. A Survey of HoneyPots and Honeynets for Internet of Things, Industrial Internet of Things, and Cyber-Physical Systems. *arXiv* **2021**, arXiv:2108.02287.

17. Morić, Z.; Dakić, V.; Regvart, D. Advancing Cybersecurity with Honey pots and Deception Strategies. *Informatics* **2025**, *12*, 14. [CrossRef]
18. Lanka, P.; Gupta, K.; Varol, C. Intelligent Threat Detection—AI-Driven Analysis of Honey pot Data to Counter Cyber Threats. *Electronics* **2024**, *13*, 2465. [CrossRef]
19. Rabzelj, M.; Sedlar, U. Beyond the Leak: Analyzing the Real-World Exploitation of Stolen Credentials Using Honey pots. *Sensors* **2025**, *25*, 3676. [CrossRef] [PubMed]
20. Tariq, U.; Ahmed, I.; Bashir, A.K.; Shaukat, K. A Critical Cybersecurity Analysis and Future Research Directions for the Internet of Things: A Comprehensive Review. *Sensors* **2023**, *23*, 4117. [CrossRef] [PubMed]
21. Hou, F.; Hou, F.; Zang, X.; Hua, Z.; Liu, Z.; Wu, Z. Effectiveness Evaluation Method for Hybrid Defense of Moving Target Defense and Cyber Deception. *Computers* **2025**, *14*, 513. [CrossRef]
22. Happe, A.; Cito, J. Getting pwn'd by AI: Penetration Testing with Large Language Models. In *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering; ESEC/FSE '23*; ACM: New York, NY, USA, 2023; pp. 2082–2086. [CrossRef]
23. Happe, A.; Cito, J. On the Surprising Efficacy of LLMs for Penetration-Testing. *arXiv* **2025**, arXiv:2507.00829.
24. Happe, A.; Cito, J. Can LLMs Hack Enterprise Networks? Autonomous Assumed Breach Penetration-Testing Active Directory Networks. *arXiv* **2025**, arXiv:2502.04227.
25. Deng, G.; Liu, Y.; Mayoral-Vilches, V.; Liu, P.; Li, Y.; Xu, Y.; Zhang, T.; Liu, Y.; Pinzger, M.; Rass, S. PentestGPT: An LLM-empowered Automatic Penetration Testing Tool. *arXiv* **2023**, arXiv:2308.06782.
26. Muzsai, L.; Imolai, D.; Lukács, A. HackSynth: LLM Agent and Evaluation Framework for Autonomous Penetration Testing. *arXiv* **2024**, arXiv:2412.01778.
27. Greshake, K.; Abdelnabi, S.; Mishra, S.; Endres, C.; Holz, T.; Fritz, M. Not What You've Signed Up For: Compromising Real-World LLM-Integrated Applications with Indirect Prompt Injection. In *Proceedings of the 16th ACM Workshop on Artificial Intelligence and Security; CCS '23*; ACM: New York, NY, USA, 2023; pp. 79–90. [CrossRef]
28. Liu, Y.; Deng, G.; Li, Y.; Wang, K.; Wang, Z.; Wang, X.; Zhang, T.; Liu, Y.; Wang, H.; Zheng, Y.; et al. Prompt Injection attack against LLM-integrated Applications. *arXiv* **2023**, arXiv:2306.05499.
29. Hui, B.; Yuan, H.; Gong, N.; Burlina, P.; Cao, Y. PLeak: Prompt Leaking Attacks against Large Language Model Applications. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security; CCS '24*; ACM: New York, NY, USA, 2024; pp. 3600–3614. [CrossRef]
30. Jin, H.; Hu, L.; Li, X.; Zhang, P.; Chen, C.; Zhuang, J.; Wang, H. JailbreakZoo: Survey, Landscapes, and Horizons in Jailbreaking Large Language and Vision-Language Models. *arXiv* **2024**, arXiv:2407.01599.
31. Das, B.C.; Amini, M.H.; Wu, Y. Security and Privacy Challenges of Large Language Models: A Survey. *arXiv* **2024**, arXiv:2402.00888.
32. Joshi, A.; Baidya, S. Securing the Cognitive Layer: A Survey on Security Threats, Defenses, and Privacy-Preserving Architectures for LLM-IoT Integration. *J. Cybersecur. Priv.* **2026**, *6*, 63. [CrossRef]
33. Laskar, M.T.R.; Alqahtani, S.; Bari, M.S.; Rahman, M.; Khan, M.A.M.; Khan, H.; Jahan, I.; Bhuiyan, A.; Tan, C.W.; Parvez, M.R.; et al. A Systematic Survey and Critical Review on Evaluating Large Language Models: Challenges, Limitations, and Recommendations. *arXiv* **2024**, arXiv:2407.04069.
34. Rose, S.; Borchert, O.; Mitchell, S.; Connelly, S. *Zero Trust Architecture*; Technical Report NIST Special Publication 800-207; National Institute of Standards and Technology: Gaithersburg, MD, USA, 2020. [CrossRef]
35. MITRE Corporation. MITRE ATT&CK for Enterprise. Available online: <https://attack.mitre.org> (accessed on 23 June 2026).

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.