Bastion Endpoint Security Monitoring Implementation and Evaluation

Task 3: Technology-Supported Security Solution

Name: Enes Oguz

ID: U012740630

This case study is an implementation case study. All deployment details, approvals, dates, case records, test outcomes, timing measurements and stakeholder validation described below are constructed assumptions. All deployment details are not records of a deployment or test execution. Published sources support the technologies and standards; they do not validate the simulated results. This case study does not establish completion of a real implementation requirement.

This test will be based on the schedule of Task 2 from October 19, 2026 to January 10, 2027. The final testing window is January 8, 2027 09:00 to 10:40 UTC. The Wazuh system is composed of Bastion as the central parts, Pawn as the monitored endpoints, and Rook as the local advisory AI. The test will include two Windows endpoints, one Ubuntu endpoint, one Linux Wazuh server and one local AI host.

A Policies adopted as a result of the project

The project security engineer started following these operating policies on 7th January 2027 as acting security manager. These policies solve the Task 2 issues of logs in all different places, muddled context, manual repeated analysis and arbitrary prioritisation. Role titles state the responsibilities; there can be multiple system administrators but only a single project operator, possibly in the lab.

Policy	Adopted operating requirement	Scenario implementation
Monitoring and logging	Enroll all three endpoints; collect authentication, process, account and network evidence; verify source arrival daily.	The inventory lists WIN-01, WIN-02 and LNX-01. A scheduled source-verification event checks collection separately from agent connectivity.
Access and change control	Separate administrator and read-only analyst access.. Record configuration changes with backup and rollback steps.	The simulated role matrix includes 12 denied-action tests. Rule, Sysmon and prompt configurations are tracked as project baseline revision 1.0.
Evidence and retention	Retain indexed lab data for 30 days and case exports for 90 days. Restrict originals, hash exports and record custody transfers. Holds suspend deletion.	Case folders original exports, manifests, working copies and analyst conclusions. Encryption at rest and protected off-host copies are assumed controls.
Responsible AI use	Send approved fields to local inference. Rook cannot execute commands to isolate endpoints, disable accounts or close cases.	Llama 3.1 8B runs through Ollama on AI-01. Every output requires verification and a link to the supplied source.

Policy	Adopted operating requirement	Scenario implementation
Incident response	Validate events, record priority, authorize mitigation, preserve evidence and verify recovery.	BAS-IR-001 in H governs case SIM-IR-001 and subsequent test cases.
Authorized lab testing	Use users, benign scripts and approved targets. Use EICAR, through a supported endpoint protection path.	Testing is restricted to the lab; no operational malware or unauthorized external target is part of the scenario.

A1 Improvement in cybersecurity decision making

Bastion gives the analyst one path to investigate source events, alert context, host identity and related activity. A failed login can be compared with an account change or process execution without having to jump between hosts. The analyst looks at the facts, checks authorization, considers asset importance and evaluates scope before deciding on priority. A Wazuh rule level starts the review. It doesn't alone prove malicious intent or business impact.

Rook offers a summary of evidence supporting severity and recommended steps for investigation. After Rook's results, the analyst verifies the facts and corrects them if necessary. The hypothetical matched-case analysis reduces the mean detection time from 180 seconds to 24 seconds. Average analysis time reduced from 480 sec to 180 sec. These are results. Not vendor claims. The supervision and accountable decisions are the same as those described by Nelson et al. (2025). Autio et al. (2024)

B Cybersecurity assurance

B1 Promotion of automation

Agents automatically send configured events to the service. Rules check for defined conditions while source-health checks look for missing telemetry. The integration sends a payload for supported alerts to the local model and stores the response with the input identifier. Failed requests go into a controlled retry queue. Collection and manual investigation still happen when AI is not available. Automated collection and summarization cut down tasks without giving response authority to the model.

B2 Security improvement and modernization

The simulated baseline uses native logs and manual review. The updated workflow adds search, structured cases, clear prioritization, collection verification and independently scored AI assistance. These changes bring practices to investigation and make missing context visible. In the constructed evaluation evidence completeness goes from 170 of 250 fields or 68% to 230 of 250 or 92%. The remaining 20 incomplete fields and two unsuccessful positive cases show that modernization doesn't remove all coverage gaps.

B3 Industry standard tools and environment

The scenario uses existing Wazuh server, indexer, dashboard and agents Sysmon on Windows and Ollama for inference. There is no creation of a new SIEM, endpoint agent or artificial intelligence model created. Sysmon provides process information and selected network information; network connection events of Sysmon require configuration since they are turned off by default (Microsoft, 2026). For

Project Baseline version 1.0, the configuration set is defined rather than creating fictional vendor build versions. For the artificial intelligence model assumption, it is Llama 3.1 8B; identification and licensing of the model will be verified in the real world.

C Digital evidence and implementation controls

C1 Digital evidence collection

The evidence includes Windows authentication and account management logs, Sysmon process and network logs, Linux authentication and service logs, Wazuh alert JSON, and packet captures for bounded time period for the discovery scenario. The tests have a case ID and an evidence checklist of five items. The required identification elements include the host, the timestamp of the event in the source, and the identity of the event; the additional elements include the account, the process, or the network information depending on the scenario. Some of the evidence items are used for enrichment beyond the detection decision.

The native exports are maintained along with the derived Wazuh events so that the analyst can differentiate between missing source data and failed collection/parsing. Specific events not being alerts need to be explicitly set up for retention and indexation; all events gathered by collectors are not always indexed as alerts. The example maintains the original exports in read-only mode, makes copies of those, and maintains the information regarding collector, time of collection, source, query period, file size, and SHA-256 manifest. The digest verifies integrity after collection but does not guarantee completeness or validity of the original source record.

C2 Confidentiality

The hypothetical situation makes assumptions regarding individual accounts, least privilege service identities, secure transport for monitoring, restricted management interfaces and secure evidence storage. There will be reviewed field allowlist to limit the AI payload and obfuscate sensitive values. The synthetic accounts will not require any use of personal data from production environments. The AI host supports integration with Bastion only, and the connection to it will be made through an authorized TLS proxy. Ollama is bound to loopback only, which is enabled by default, and outbound connections are limited. Ollama documentation describes the local-only configuration and loopback functionality (Ollama, n.d.).

C3 Integrity

The integrity mechanisms are role separation, configuration history, common UTC time, source vs index comparison, hash of exports and custody logging. The native evidence, analyst notes, and AI recommendations are independent. Exported files have restricted write access and a copy outside of the host machine. The central server is not immutable – there is a danger in privilege abuse. In actual use, transport, permissions and hash verification require observation since there is no such guarantee in this hypothetical scenario.

C4 Availability

Wazuh is allocated ahead of optional inference components and relies on service tests, capacity tests, backup, and restore testing. The test in the final 100 minutes consists of 99 passes of central-service tests and 294 passes of endpoint-minute telemetry tests among 300 attempts. These translate to 99% platform availability and 98% telemetry availability. Central-host failure disrupts monitoring, which means that

this system is not a high-availability cluster. The maintenance goals assumed for this test include a four-hour service recovery and maximum configuration loss of 24 hours.

D Investigation and mitigation of incidents

An inquiry starts with either an alert or question. The analyst confirms host identity, UTC time, source logs, related activity, and testing permission. Mappings from MITRE ATT&CK structure the actions being analyzed; however, they do not determine whether or not something is malicious (MITRE, n.d.). If there are repeated attempts, it will require multiple pieces of information to clear. These will include: an examination of the account used, source, timing, and any succeeding login that was successful. Creating an account requires the name of the creator and its intended purpose. Running PowerShell requires context regarding the command and process. EICAR demonstrates a supported detection event flow path.

For the simulated case SIM-IR-001, WIN-01 generates ten login failures against a simulated account from 09:05:00 to 09:05:09 UTC. The tenth failure satisfies the condition of the rule; the alert that is triggered becomes available for search at 09:05:27, eighteen seconds later. The analyst evaluates source data, verifies that the simulated exercise is within scope, reviews Rook results, and enters the classification at 09:08:27. The security-manager privilege allows temporary isolation of the source of the simulated exercise at 09:09:00; the administrator executes this at 09:09:20. This is not a real criminal attack.

The analyst maintains the limited exports prior to the restriction scenario. By 09:12:00, the administrator returns to the authorized access route and confirms agent connectivity, arrival of the source event, and expected account status. The incident is closed at 09:15:00 with the assumption that no further unauthorized activity takes place. The above process is documented in the protocol H. In the case of a real unexpected intrusion event, any unrelated testing would be halted, and the owner decides the next step in incident escalation. Labeling alone on the dashboard cannot provide criminal culpability.

E Plans standards and procedures

This case is used to develop the monitoring plan, sources list, access matrix, evidence procedure, severity scale, Controlled change log, AI evaluation criteria, testing log, risk log and maintenance schedule. The sources listing covers the hosting location and ownership for each source. Explanation, assets affected by this change, authorisation, backup procedure, positive/benign tests, rollback. Controlled change log. AI evaluation criteria do not allow embedded instruction within the event contents. Model is still in advisory mode. These documents are the assumed baseline of operation 1.0 which is allowed in this simulation on 7 th January 2027.

E1 Cybersecurity initiative and compliance alignment

The project goal is the enhancements to endpoint visibility, investigations, and accountability with AI in a closed academic environment. The results are classified in NIST CSF 2.0. NIST CSF 2.0 is not a product certification (National Institute of Standards and Technology [NIST], 2024). The map indicates which of the project practices contribute to the functions.

CSF function	Project practice	Simulated verification
Govern	Assign policy owners, review risks and document acceptance.	Operating baseline 1.0 and qualified scenario acceptance.
Identify	Maintain endpoint and management-asset inventory.	Three monitored endpoints; two management servers.
Protect	Restrict privileges, protect transport and evidence, control changes.	Twelve denied administrative attempts; assumed storage and transport controls.
Detect	Collect sources, evaluate rules and verify event arrival.	48/50 positive successes and 294/300 telemetry checks.
Respond	Validate evidence and authorize containment.	SIM-IR-001 case and procedure record.
Recover	Restore approved state and review residual risk.	Simulated reconnection and source checks; restore exercise scheduled.

Guidance for response is provided in NIST SP 800-61 Rev. 3, guidance for logging is provided in SP 800-92, guidance for handling evidence is provided in SP 800-86, and guidance for human review and unsupported output is provided in the Generative AI Profile (Nelson et al., 2025; Kent & Souppaya, 2006; Kent et al., 2006; Autio et al., 2024). There is no production dataset that is in any way regulated here. The case study does not make any HIPAA, PCI DSS or other compliance claim.

E2 Tools installation and user guides

Project guide	Developed content	Scenario validation
Central installation and restoration	Service placement and resource recording; use the matching vendor guide; limit access; create roles; back up settings; validate services and search authentication.	Central health check result is 99/100 successes. Restore operation in the future is an operational necessity and not a measurement outcome.
Endpoint onboarding and telemetry	Enroll pilot host; define source channels; configure Sysmon settings; compare native event with Wazuh and reconnect test.	All three endpoints were enrolled. Formative network event omission was addressed by configuring Sysmon settings.
Rook configuration and analyst use	Set Llama 3.1 8B; local inference; field allowlist; timeout/retry setting; record model and prompt modifications; enforce source links and facts checking.	49/50 first attempt completion; one additional attempt to complete final case. Overall case traceability is 50/50.
Access and case handling	Define role mapping; fixed range search; case creation; export control; AI review; escalation and closure process.	Twelve denied-action tests and the simulated case SIM-IR-001 demonstrate the intended workflow.

However, Wazuh, Sysmon, and Ollama are still external products. The actual project contribution is the configuration and instructions for the configuration. Wazuh provides information about the components' architecture and role-based authorization. The short operating guides provided above explain the

expected process; they do not represent installation scripts and proof of running the commands. The analyst participates in a 2-hour case study within the context of the virtual environment, the administrators receive an hour's training on onboarding and connectivity techniques, while the manager receives a 30-minute briefing on KPIs.

F Post implementation environment

The environment is based on a 10.20.30.0/24 lab subnet, with no inbound access allowed from the public Internet. WIN-01 and WIN-02 use Windows 11, and LNX-01 uses Ubuntu. BASTION-01 uses the central Wazuh platform. AI-01 uses Ollama and Llama 3.1 8B. Time shared across all hosts. All addresses and locations are just creative choices for illustration purposes.

Asset	Assumed address and function	Authorized communication
WIN-01 and WIN-02	10.20.30.11 and .12; Wazuh agents and Sysmon.	Send selected telemetry to BASTION-01; no direct access to AI or indexer management.
LNX-01	10.20.30.13; Wazuh agent and native logs.	Send selected telemetry to BASTION-01; synthetic authentication and network tests only.
BASTION-01	10.20.30.10; Wazuh server, indexer and dashboard.	Receives agents; serves restricted dashboard; sends selected payloads to the protected AI proxy.
AI-01	10.20.30.20; local model runtime behind an authenticated TLS proxy.	Accepts the integration path from BASTION-01; local runtime listens behind the proxy.
Analyst access and evidence storage	Restricted management workstation and protected off-host case copy; not counted as monitored endpoints.	Authorized dashboard access and evidence handling; administration requires a separate role.

The allowed paths include agent communications on TCP port 1514, enrollment process on TCP port 1515 during onboarding, dashboard over HTTPS on port 443, and AI proxy for the project over HTTPS. Wazuh Server and Indexer APIs are restricted to only necessary internal services and admins. The vendor provides information about the default component paths and settings. Checking the firewall configuration itself is another task. Among the new processes are source-based onboarding, health checkup on a daily basis, documentation of cases, AI verification, and tuning.

F1 Security posture and efficiency

Constructed comparison is based on 48 matching success cases. Baseline detection sums up 8,640 seconds and post-implementation detection sums up 1,152 seconds, hence equal 180 and 24 respectively. Time spent on detection decreased by the factor of (180 – 24) / 180 x 100 = 86.7%. Baseline analysis sums up 23,040 seconds and post-implementation analysis sums up 8,640 seconds, which means equal 480 and 180 respectively. Hence, time spent on analysis decreased by a factor of 62.5%. Post-implementation analysis time includes time for humans to verify AI results.

On average, across the 48 matched cases, baseline analysis takes 384 analyst minutes and post-implementation analysis takes 144 minutes, which means simulated savings in minutes amount to 240. However, these are the savings in terms of review time and not in financial terms because all the

deployment costs etc are additional. Completeness of evidence increased by 24 percentage points from 68% to 92%. There are two positive cases that still remain unsuccessful.

F2 New data analysis and business process effects

Constructed dataset	Finding and interpretation	Effect on operations
Inventory and source checks	3/3 hosts enrolled; 294/300 verification events meet the 60-second search requirement. Six checks fail.	Administrators investigate source gaps even when agent connectivity appears normal.
Positive scenario evidence	230/250 checklist fields present and accurate; 48/50 cases meet the minimum investigation requirements.	The case form distinguishes mandatory decision evidence from missing enrichment; unsupported cases remain visible.
Normal activity controls	2/30 benign activities incorrectly flagged, or 6.7%. Both involve approved administrative behavior.	Analysts review authorization context; targeted tuning must preserve positive detections.
AI completion and accuracy	49/50 initial requests complete; one recorded retry succeeds. Final factual scoring is 235/250, or 94%.	A completed response is not automatically correct. Human verification remains part of every case.
Resources and central health	Additional CPU is 3, 4 and 6 percentage points by the monitored host; central health is 99/100.	All hosts meet the CPU threshold. Infrastructure outages and AI resource use still require maintenance planning.

The six erroneous endpoint-minute tests consist of three tests during a one-minute central disruption and three tests for delays due to source delivery outside of that minute. Failure of source verification does not indicate a loss of all security events for that minute. There is no claim of an unfounded reconstruction made in the report. In this theoretical business process, central search takes the place of individual source search, administrators are responsible for telemetry failures, and analysts record decision-making at the end.

F3 Summative evaluation and results

The assessment continues to use the Task 2 formative and summative assessment scheme. Formative assessments include key services, agent enrollment and reconnection, Windows and Linux data ingestion, positive and negative rule samples, roles, Rook precision and endpoint performance. The formative assessment scenario identifies two formative deficiencies – lack of Sysmon network events and lack of AI statement support. Network event configuration and prompt/input modification occur before baseline version 1.0 and final testing. These are hypothetical modifications not made by this report.

The final solution addresses 10 cases for each of the five positive scenarios as well as 30 good behaviors. EICAR is assumed to be supported by the Windows endpoint security configuration, so S5 is included. All 50 positive cases run. The outcome of formal scenarios is assessed based on required alert/query evidence. Five items per case result in 250 scored opportunities. Enrichment items contribute

to completeness but do not contribute to scenario success. Planned failures are never removed from the denominator.

Scenario and case range	Cases successful	Checklist fields	Simulated interpretation
S1 Authentication failures SIM-S1-01 to 10	10/10	48/50	All case decisions supported; two enrichment fields absent.
S2 Test account creation SIM-S2-01 to 10	10/10	48/50	Creator and target established; two enrichment fields absent.
S3 Benign PowerShell SIM-S3-01 to 10	9/10	46/50	SIM-S3-07 lacks required process context; other missing items are enrichment.
S4 Network discovery SIM-S4-01 to 10	9/10	40/50	SIM-S4-04 lacks required correlation; other missing items limit enrichment.
S5 EICAR security event SIM-S5-01 to 10	10/10	48/50	Supported security-event path established; two enrichment fields absent.
Total positive set	48/50 = 96%	230/250 = 92%	Two unsuccessful cases remain in the final result.
S6 Normal activity SIM-B-01 to 30	28 correctly unflagged; 2 false positives	Not in positive field denominator	False-positive rate is 2/30 = 6.7%.

All 50 positive cases generate traceable alerts that are submitted to the AI system for analysis, even the two positive cases which do not contain enough information for proper case investigation. The alert exists and the scenario investigation was completed are two separate parameters. One of the AI calls fails in its first execution, but after a controlled re-execution, the output for the lost case is received. The processing success rate is 49 out of 50 cases, or 98%. In total, including the re-execution, the success rate of 50 out of 51 requests is 98.0%. Traceability is evaluated in the end stage of the final case analysis.

Task 2 KPI	Constructed result and calculation	Minimum and decision
Endpoint enrollment	3 / 3 x 100 = 100%.	100%: Pass.
Telemetry availability	294 / 300 scheduled endpoint-minute checks x 100 = 98%.	At least 95%: Pass.
Evidence completeness	230 / 250 applicable checklist fields x 100 = 92%.	At least 90%: Pass.
Positive scenario success	48 / 50 planned applicable positive cases x 100 = 96%.	At least 90%: Pass.
High-priority alert availability	20 predefined alert-expected cases from S1/S2; all at or below 52 seconds from the triggering condition.	Each at most 60 seconds: Pass.
False-positive rate	2 / 30 executed benign cases x 100 = 6.7%.	At most 10%: Pass.
Rook processing success	49/50 first-pass requests = 98%; after retry, 50/51 attempt completions = 98.0%.	At least 90%: Pass on either disclosed basis.
Rook factual agreement	235 / 250 final-output factual scoring items x 100 = 94%.	At least 90%: Pass.

Task 2 KPI	Constructed result and calculation	Minimum and decision
Rook evidence traceability	50 / 50 final reviewed cases have source-linked output, including the recovered timeout.	100%: Pass after retry.
Endpoint CPU overhead	WIN-01: 3; WIN-02: 4; LNX-01: 6 additional percentage points.	Less than 10 per host: Pass.
Wazuh availability	99 / 100 minute checks validate server, indexer and authenticated search = 99%.	At least 95%: Pass.
Unauthorized administrative access blocked	12 / 12 unauthorized actions denied with no state change. Authorized control actions succeed separately.	100%: Pass.

CPU value calculation takes the form of 30-minute windows of benign workload data that are collected at five-second intervals, hence 360 samples for each window per endpoint. For the WIN-01 machine, the pre-installation and post-installation usage are assumed to be 8% and 11%, respectively, while that of the WIN-02 is 10% and 14%, and LNX-01 is 6% and 12%. These represent percentage points and not a relative increase in CPU utilization. The timing comparison uses the same 48 case IDs for each set of conditions, with no unsuccessful case being assigned zero detection time.

Deficiencies corrective action and retesting

Finding	Action and accountable role	Retest or remaining decision
Formative network events omitted	Engineer explicitly enables selected Sysmon network events and verifies local-to-central fields before the final configuration freeze.	Formative ingestion retest assumed successful. Final S4 still has one separate correlation failure; it is not hidden.
Formative AI output overstates facts	The engineer limits input fields and insists on "unavailable data" remarks; the analyst compares the results with the source.	Final factual agreement is 235/250. Fifteen errors remain, requiring human correction and a further prompt review.
One final AI timeout	Engineer records the failed first attempt and retries once without changing the model or prompt.	Second attempt assumed successful. Original failure remains in the 51-attempt count; final case traceability is 50/50.
SIM-S3-07 and SIM-S4-04 fail investigation requirements	Engineer reviews the missing process context and correlation query; analyst preserves both failed cases.	Corrections due January 15 in the hypothetical schedule. Rerun affected cases with positive and benign controls; no successful final retest is invented.
Two benign false positives and six telemetry misses	Analyst reviews approved administrative context; administrator investigates delivery delay and the central interruption.	Targeted corrections due January 15; full control checks after changes. Original rates remain 6.7% and 98%.

Finding	Action and accountable role	Retest or remaining decision
Short window and central-host dependency	Manager records the limitations and schedules longer availability and restoration exercises.	Residual risks accepted for the hypothetical lab only; no production-readiness claim.

All twelve numeric minimum values fit within the given assumptions, and the hypothetical manager agrees to accept the lab with noted residual problems on January 10, 2027. Agreement doesn't mean that every case would work out well or all of the statements of the AI would be right. Two unsuccesful investigations, fifteen factual mistakes and small time interval of observation still serve as limitations. Any further correction which would alter the configuration would require re-running and reporting of all of the affected cases along with the original outcome. The presented simulated outcomes do not prove completion of any real project.

F4 Post implementation risks

Qualitative indicators of the environment are impact and likelihood high impact would imply loss of trusted information, unauthorised access or appreciable interruption in monitoring; Medium impact would imply a bounded delay or rework.

Risk and likelihood	Organizational impact	Mitigation and owner
Telemetry or parsing drift Medium	High: missing records may create false confidence.	Engineer compares sources and central records daily and after updates; administrator tracks gaps.
Coverage gaps Medium	High: behaviors beyond the tested scope remain unproven.	Analyst maintains the two failed cases and a coverage register; scope expansion requires approval.
False-positive workload Medium	Medium: routine administrative work consumes analyst attention.	Analyst tunes narrowly, verifies authorization and reruns benign and positive controls.
AI error or prompt injection Medium	High: misleading advice could cause incorrect response.	Engineer restricts inputs and capabilities; analyst verifies facts and approves decisions independently.
Central-host outage or full storage Medium	High: monitoring and history may be unavailable.	Administrator monitors resources, retains protected off-host copies and schedules restoration exercises; no HA claim.
Sensitive telemetry exposure Medium	High: command lines or exports could disclose secrets.	Administrator restricts roles and storage; engineer masks AI fields and verifies protected local inference.
Privilege misuse or vulnerable dependencies Low misuse; medium dependency risk	High: control changes can undermine monitoring trust.	Administrator reviews roles and audit records, stages updates and retains rollback configurations.
Model change or license mismatch Medium	Medium: inconsistent summaries or inappropriate distribution/use.	Engineer records the model selection, reviews terms before real use and repeats reference cases before substitution.

Risk and likelihood	Organizational impact	Mitigation and owner
Clock drift Low	High: incident order and timing metrics become unreliable.	Administrator checks the common UTC time source before tests and after time-service faults.

F5 Stakeholder needs and solution response

The nine stakeholder groups listed below correspond with Task 2. The validation in this case study is fictional. The external product providers have not been assumed to approve, test or endorse the project. The concerns of these stakeholders are catered for through supported use and dependency assumptions.

Stakeholder and Task 2 need	How the assumed solution meets the need	Simulated validation
Security analyst Searchable context, useful priorities and manageable false positives.	Central search, source-linked summaries and a case rubric support consistent decisions.	180-second final mean analysis; 2/30 benign false positives; analyst case exercise.
System administrator Stable agents, clear deployment and limited overhead.	Pilot onboarding, reconnection checks and per-host measurements support administration.	3/3 enrolled; additional CPU 3, 4 and 6 points; connectivity exercise.
Network administrator Isolation and explicit permitted flows.	Private segment and restricted enrollment, telemetry, dashboard and AI paths bound access.	Assumed flow review and authorized S4 scanner/target list.
Security manager or sponsor Risks, acceptance measures and governance.	KPI register and qualified acceptance distinguish minimum compliance from complete coverage.	Twelve threshold passes; two failed cases and residual AI errors retained.
End user Performance, privacy and understandable reporting.	Synthetic data, limited fields, resource review and a short briefing address impact.	Per-host overhead below 10 points; assumed 15-minute briefing.
Project security engineer Repeatability, clear scope and documentation.	Baseline 1.0, defined case ranges and correction ownership support reproducible evaluation design.	50 positive cases; 30 controls; explicit retry and timing denominators.
Wazuh project or vendor Supported product architecture and dependency constraints.	Uses central components and agents; keeps configuration work distinct from vendor software.	Architecture and role guidance referenced; no vendor endorsement claimed.
Microsoft Sysinternals Appropriate Sysmon configuration and license use.	Limits selected events and distinguishes event collection from threat decisions.	Assumed configuration review and source comparison; real license/build verification remains necessary.

Stakeholder and Task 2 need	How the assumed solution meets the need	Simulated validation
Ollama and model provider Runtime resources, model terms and reliable use.	Local model, restricted proxy path and human verification address intended use.	98% first-pass processing; 94% factual agreement; model/license review required before real deployment.

G Post implementation maintenance plan

The assumed manager delegates maintenance responsibility to the project engineer, while the tasks of administrator and analyst come next. The chosen lab retention period is 30 days for indexed items and 90 days for exported cases, further increased due to hold. The described scenario involves daily backup of the configurations as well as monthly restore exercise; no scheduled exercise equals recovery test.

Frequency or trigger	Activity and accountable role	Record and escalation
Daily	The administrator checks services, agents, source arrival, time and capacity; analyst reviews priority cases; engineer reviews AI failures.	Health and case log. Escalate missing source or failed search immediately.
Weekly	The analyst reviews false positives and gaps; engineer checks supported updates and configuration drift.	Tuning review, update triage and open-deficiency status.
Monthly	The administrator verifies access, retention, certificate expiration, backups, and the engineer conducts a restore process under supervision.	Check access logs, retention/deletion of logs, and data recovery techniques. Do not presume that 4-hour and 24-hour goals are achieved.
Quarterly or major change	The engineer repeats scenarios, AI scoring, access tests and resource checks and the analyst refreshes training.	Regression and training records; changed configurations require acceptance review.
Urgent vulnerability	Relevance is assessed by administrator in one working day, implements patch under authorized window of time; isolate exposure if waiting is unacceptable.	Risk and patch decision; verify source collection, search and access; retain rollback.
Model prompt rule or source change	The engineer versions the change, tests known-positive and benign samples and retains the previous approved configuration.	Before/after comparisons, approval and rollback criteria; never silently remove failed cases.
Outage	The administrator restores central collection and search before optional AI support, then checks permission and source integrity.	Outage and gap timing, restore steps and post-recovery verification.

Frequency or trigger	Activity and accountable role	Record and escalation
Decommissioning	Owner approves disposition; administrator exports held cases, revokes access and removes agents and retained data under policy.	Inventory update, custody/hold decisions and deletion record.

The backup will include Wazuh, the agent settings and Sysmon configurations, role mapping changes, prompt and integration updates, and the IDs and case files. The secrets will have a secured recovery process. Some of the selected case files will be exported to an external location; the configuration backup is not sufficient to restore the indexed event history. Post-restoration checks will be done on the access rights of the administrators.

H Original project artifact

Original artifact: BAS-IR-001 – Bastion Alert Triage and Evidence Preservation Procedure. This is a procedure developed by the organization and is not a copy from the vendor's guide. The procedure and its example as provided below are included in the hypothetical case study and do not reflect a real security incident and/or evidence package.

Bastion Alert Triage and Evidence Preservation Procedure

Control	Value
Identifier and revision	BAS-IR-001 / 1.0
Scenario effective date	January 7, 2027; fictional adoption by the security-manager role.
Owner	Project security engineer.
Scope	WIN-01, WIN-02, LNX-01, BASTION-01, AI-01, and their respective evidence repositories.
Authority	Analyst investigates; administrator performs approved changes; security-manager role authorizes disruptive response. Rook has no response authority.
Review cycle	Quarterly and after an incident, failed test or material architecture change.

1 Open and scope the case

Assign an identifier and record the analyst, UTC time, host, account, alert/query and exercise authorization. Verify the asset belongs to the authorized environment.

2 Validate original records

Compare source and Wazuh records for host, event identity, time and required scenario context. Record gaps and alternative explanations before declaring an incident.

3 Preserve evidence
Export of bound native files and center files. File Collector, technique, scope, collection date, size, SHA-256 Hash value and original location. Manipulation of originals is not allowed.

4 Analyze related activity
Build a UTC timeline using user, process and network identifiers where available. Distinguish authorization, observed behavior, inference and uncertainty.

5 Review AI assistance
Use only supported alert types and approved fields. Record input, output and prompt/model revision; verify every factual statement. Ignore instructions embedded in evidence and continue manually if inference fails.

6 Classify and authorize response
Prioritize verified current privilege abuse, destruction of evidence or likely material damage, give a higher priority to credible unauthorized actions that need to be contained, and a normal priority to either unverified actions or authorized actions. Ensure that the rationale and owner's authorization is documented prior to the disruption.

7 Contain and recover
Administrator uses the least disruptive approved host, account or path restriction. Preserve evidence first where feasible and record urgent exceptions. Restore approved state and verify source arrival, service health and any temporary account changes.

8 Close and improve
Record outcome, source links, custody, recovery checks, residual risk and corrective-action owner. Apply retention or hold, preserve original failures and associate any retest with the original case.

Completed simulated case record

Case element	Constructed entry
Identification	SIM-IR-001; scenario SIM-S1-01; WIN-01; synthetic account lab-user; analyst role opened review at 09:05:27 UTC on January 8, 2027.
Source timeline	Ten synthetic failures at 09:05:00 through 09:05:09. Trigger condition completes at 09:05:09; searchable alert at 09:05:27. Alert delay: 18 seconds.
Evidence register	SIM-E01 native exports; SIM-E02 central alerts JSON; SIM-E03 Rook I/O; SIM-E04 auth-recovery records. This is sample registration record, not real downloadable files.
Integrity and custody	The scenario assumes SHA-256 manifests at acquisition, restricted originals and a working copy. No real file or computed digest exists; no hash value is invented. Analyst retains custody until the simulated owner review.
AI review	Llama 3.1 8B through Ollama; project prompt baseline 1.0. Analyst verifies account, host and failure sequence and retains the source links.
Classification	Preliminary classification at 09:08:27: controlled authentication exercise matching the high-priority test condition; no real criminal attribution.
Authorization and mitigation	Owner authorizes temporary lab-source restriction at 09:09:00; administrator applies it at 09:09:20. The exercise is an authorized containment demonstration.

Case element	Constructed entry
Recovery and closure	Approved path restored at 09:12:00; agent and source arrival checks assumed successful; case closed at 09:15:00. Exports would follow the 90-day policy in a real exercise.
Lesson	Retain the trigger timestamp separately from the first failure so high-priority alert delay and overall detection timing are not confused.

For actual use, the procedure requires real exports, computed digests and custody acknowledgements. This simulated example illustrates the record structure without fabricating authentic logs, file signatures, operator approvals or evidence provenance.

I References

Autio, C., Schwartz, R., Dunietz, J., Jain, S., Stanley, M., Tabassi, E., Hall, P., & Roberts, K. (2024). *Artificial intelligence risk management framework: Generative artificial intelligence profile* (NIST AI 600-1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.600-1

Kent, K., Chevalier, S., Grance, T., & Dang, H. (2006). *Guide to integrating forensic techniques into incident response* (NIST SP 800-86). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-86

Kent, K., & Souppaya, M. (2006). *Guide to computer security log management* (NIST SP 800-92). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-92

Microsoft. (2026). *Sysmon v15.22*. Microsoft Sysinternals. https://learn.microsoft.com/en-us/sysinternals/downloads/sysmon

MITRE. (n.d.). *MITRE ATT&CK*. Retrieved October 17, 2026, from https://attack.mitre.org/

National Institute of Standards and Technology. (2024). *The NIST cybersecurity framework (CSF) 2.0* (NIST CSWP 29). https://doi.org/10.6028/NIST.CSWP.29

Nelson, A., Rekhi, S., Souppaya, M., & Scarfone, K. (2025). *Incident response recommendations and considerations for cybersecurity risk management: A CSF 2.0 community profile* (NIST SP 800-61 Rev. 3). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-61r3

Ollama. (n.d.). *FAQ*. Retrieved October 17, 2026, from https://docs.ollama.com/faq

Wazuh. (n.d.-a). *Architecture*. Retrieved October 17, 2026, from https://documentation.wazuh.com/current/getting-started/architecture.html

Wazuh. (n.d.-b). *Wazuh RBAC - How to create and map internal users*. Retrieved October 17, 2026, from https://documentation.wazuh.com/current/user-manual/user-administration/rbac.html