

FORGEGRAPH

Sample AI & Product Diagnostic

Sample scope:

Sample Diagnostic deliverable based on the fictional ForgeGraph demonstration—not a customer engagement.

OBSERVED

This retrospective sample demonstrates GreyX's diagnostic format using reviewed ForgeGraph demonstration material and approved public assets; it is not customer work or a past delivery.

Evidence

- E5 — Public architecture map
- E6 — Fictional workspace capture
- E7 — Fictional approval capture

Executive decision summary

OBSERVED

The reviewed design separates backend-controlled durable state from ephemeral execution work.

Evidence

E1 — Durable-state boundary

ANALYSIS

Preserving the durable-state boundary, making human approval and recovery evidence inspectable, and validating concentrated orchestration areas are the most defensible next decisions.

Evidence

E1 — Durable-state boundary
E3 — Human approval state
E4 — Retry-aware delivery

Confidence

High

RECOMMENDATION

Validate the ownership boundary and recovery behavior before considering any stack change.

Evidence

E1 — Durable-state boundary
E2 — Recoverable projection

Validation

Documented invariants and characterization checks show the intended ownership and recovery behavior.

Evidence class legend

Observed

A statement directly supported by reviewed public-safe evidence. It does not establish production performance or provide proof about a customer.

Assumption

An unverified premise used to frame possible scope. Validate it before defining scope.

Open Question

An unresolved input to a decision. It does not invent an answer, stakeholder, or requirement.

Analysis

A reasoned interpretation of reviewed evidence with stated confidence. It is not a measured outcome.

Recommendation

A proposed future action with a validation criterion. It is not an implemented result.

Method, evidence, and limitations

OBSERVED

This sample uses reviewed demonstration material and approved public assets while excluding source code, repository data, private operational material, raw graph or index data, prompts, credentials, local paths, and the internal code index from publication.

Evidence

E5 — Public architecture map
E6 — Fictional workspace capture
E7 — Fictional approval capture

OBSERVED

This sample does not establish live workload, user behavior, costs, service levels, performance, security posture, production deployment status, or commercial results.

Evidence

E5 — Public architecture map
E6 — Fictional workspace capture
E7 — Fictional approval capture

Sample scope:

Sample Diagnostic deliverable based on the fictional ForgeGraph demonstration—not a customer engagement.

Current system and product objective

OBSERVED

The reviewed demonstration combines product workflow, durable state, execution, AI assistance, and explicit human decision points.

Evidence

E1 — Durable-state boundary
E3 — Human approval state

OBSERVED

The review is limited to public evidence: it includes no runtime measurement, stakeholder interview, or customer data, and uses a deliberately simplified system map.

Evidence

E5 — Public architecture map
E6 — Fictional workspace capture
E7 — Fictional approval capture

ASSUMPTION

Assumption: a team may need to coordinate operational work with explicit decision gates; this sample framing requires validation.

Validation

A real Day 1 working call must confirm the outcome, roles, approval policy, data constraints, and investment guardrails.

Architecture and ownership boundaries

OBSERVED
The simplified public sequence is Browser → Next.js → Django → Postgres → Go Engine → LLMs.

Evidence
E5 — Public architecture map

OBSERVED
The reviewed separation assigns durable control-plane authority to the backend and keeps execution work ephemeral.


Evidence
E1 — Durable-state boundary

OBSERVED
The reviewed projection can be reconstructed from backend-owned data after cache-only loss.

Evidence
E2 — Recoverable projection

OBSERVED
The reviewed delivery boundary retains durable, retry-aware semantics separate from decision state.

Evidence
E4 — Retry-aware delivery



```
graph LR; Browser --> Nextjs[Next.js]; Nextjs --> Django; Django --> Postgres; Postgres --> GoEngine[Go Engine]; GoEngine --> LLMs
```

Simplified public architecture map for the fictional ForgeGraph demonstration.

Strengths worth preserving

OBSERVED
Durable-state ownership is explicitly separated from ephemeral execution work in the reviewed design.

Evidence
E1 — Durable-state boundary

OBSERVED
Human approval is modeled as explicit workflow and decision state in the reviewed demonstration.


Evidence
E3 — Human approval state

OBSERVED
Recoverable projections and retry-aware delivery boundaries are worth preserving as explicit design choices.

Evidence
E2 — Recoverable projection
E4 — Retry-aware delivery

Sample scope:

Sample Diagnostic deliverable based on the fictional ForgeGraph demonstration—not a customer engagement.



Fictional Northstar Analytics workspace capture for UI/workflow evidence only; it does not evidence a client, customer, revenue, users, deployment, or outcome.

Prioritized risks and opportunities

Priority 1

RECOMMENDATION

Make ownership boundaries and recovery expectations explicit and regression-tested.

Confidence
High

Evidence
E1 — Durable-state boundary
E2 — Recoverable projection

Implication
Unclear ownership makes recovery decisions harder to verify.

Next action
Document invariants and add bounded recovery characterization and regression checks.

Priority 2

RECOMMENDATION

Make high-impact human approval evidence and decision state easy to inspect.

Confidence
High

Evidence
E3 — Human approval state

Implication
Approval state may be correct but difficult to audit across handoffs.

Next action
Expose approval decisions and evidence at high-impact boundaries.

Priority 3

RECOMMENDATION

Exercise cache-loss and checkpoint-loss behavior as product behavior, not only infrastructure behavior.

Confidence
Medium

Evidence
E1 — Durable-state boundary
E2 — Recoverable projection

Implication
Cache or checkpoint loss can become product-visible recovery behavior.

Next action
Exercise bounded loss scenarios and record expected recovery.

Priority 4

ANALYSIS

Several reviewed controller/orchestration units coordinate many product-critical responsibilities; characterize their boundaries and test coverage before deciding whether focused decomposition is warranted.

Confidence
Medium

Evidence
E8 — Reviewed maintainability signal

Implication
Concentrated responsibilities may increase the cost of safe review and change.

Next action
Characterize responsibilities and relevant tests before choosing focused decomposition.

Priority 5

RECOMMENDATION

Improve operator visibility into failed or deferred external delivery while retaining durable, retry-aware semantics.


Confidence
High

Evidence
E4 — Retry-aware delivery

Implication
Deferred delivery can leave operators unsure what needs action.

Next action
Expose failed and deferred state while retaining durable retry semantics.

Sample scope:
Sample Diagnostic deliverable based on the fictional ForgeGraph demonstration—not a customer engagement.



Fictional Northstar Analytics approval capture for UI/workflow evidence only; it does not evidence a client, customer, revenue, users, deployment, or outcome.

ADR-001 — Preserve a backend-controlled durable control plane with an ephemeral execution plane

Context

ANALYSIS

A durable control plane and an ephemeral execution plane make ownership and recovery expectations reviewable as separate concerns.

Evidence

E1 — Durable-state boundary
E2 — Recoverable projection

Confidence

High

Recommendation

RECOMMENDATION

Preserve the separated ownership model.

Evidence

E1 — Durable-state boundary
E2 — Recoverable projection

Validation

Future changes are checked against a documented decision record and characterization coverage.

Observed evidence

OBSERVED

The reviewed design keeps durable runtime state under backend control rather than assigning it to ephemeral execution work.

Evidence

E1 — Durable-state boundary

Trade-offs

ANALYSIS

Preserving the boundary keeps recovery and ownership explicit, while future changes must continue to validate those trade-offs.

Evidence

E1 — Durable-state boundary
E2 — Recoverable projection

Confidence

Medium

Options

RECOMMENDATION

Preserve the separated ownership model.

Evidence

E1 — Durable-state boundary
E2 — Recoverable projection

Validation

Future changes are checked against a documented decision record and characterization coverage.

NOT RECOMMENDED

ANALYSIS
Letting execution own durable runtime mutation is not recommended because it would require a different ownership and recovery model to be validated.

Evidence

E1 — Durable-state boundary

Confidence

Medium

Consequences

ANALYSIS

Preserving the boundary keeps recovery and ownership explicit, while future changes must continue to validate those trade-offs.

Evidence

E1 — Durable-state boundary
E2 — Recoverable projection

Confidence

Medium

NOT RECOMMENDED

ANALYSIS

Treating events as the authoritative state source is not recommended without separately validating a different durable-state model.

Evidence

E1 — Durable-state boundary

Confidence

Medium

Validation criteria

RECOMMENDATION

Preserve the separated ownership model.

Evidence

E1 — Durable-state boundary
E2 — Recoverable projection

Validation

Future changes are checked against a documented decision record and characterization coverage.

Remaining uncertainty

OPEN QUESTION

Which current responsibilities and tests should be characterized before a focused decomposition decision?

Decision affected

The scope and order of future characterization and decomposition work.

Roadmap and suggested technology direction

First

RECOMMENDATION

First, formalize invariants, the decision record, and characterization tests.

Evidence

E1 — Durable-state boundary

E2 — Recoverable projection

Validation

The invariants and ADR are reviewable and each has a bounded characterization check.

Next

RECOMMENDATION

Next, improve approval, recovery, and operator visibility.

Evidence

E2 — Recoverable projection

E3 — Human approval state

E4 — Retry-aware delivery

Validation

A review can inspect approval, recovery, and delivery evidence at the relevant decision boundaries.

Then

RECOMMENDATION

Then decide whether focused decomposition or additional product work is warranted.

Evidence

E1 — Durable-state boundary

E3 — Human approval state

Validation

The decision is supported by the completed characterization and the real outcome priorities.

ANALYSIS

The conservative direction is to preserve the reviewed Next.js, Django, Postgres, Go-engine, and LLM boundary vocabulary where it serves the ownership model.

Evidence

E1 — Durable-state boundary

E5 — Public architecture map

Confidence

High

Implementation phases, timeline, and commercial framing

OBSERVED GreyX's AI & Product Diagnostic is USD 350 fixed and delivered in three business days through Day 1 Discovery, Day 2 Analysis, and Day 3 Delivery. Evidence C1 — GreyX Diagnostic offer	OBSERVED GreyX's Build Sprint starts at USD 2,500 and generally takes one to three weeks; Product Partnership is custom quote. Evidence C1 — GreyX Diagnostic offer
RECOMMENDATION Keep a suggested sequence separate from commercial facts and treat follow-on offers as conditional next engagement options. Evidence C1 — GreyX Diagnostic offer Validation A real buyer confirms scope, investment guardrails, and timing before any commercial estimate is discussed.	
Illustrative GreyX offer framing only—not a ForgeGraph estimate, quote, spend record, engagement timeline, or commitment.	
GreyX Diagnostic method Day 1 Discovery Day 2 Analysis Day 3 Delivery	Offer facts AI & Product Diagnostic — USD 350 fixed — 3 business days Build Sprint — starts at USD 2,500 — one to three weeks Product Partnership — custom quote

Open questions

OPEN QUESTION Which real user or operator roles need this workflow, and which outcome should be prioritized? Decision affected Which outcome the real system should prioritize and how it should be measured.	OPEN QUESTION What production evidence, if any, is available for the real system? Decision affected Whether additional validation can rely on available real-system evidence.
OPEN QUESTION What approval policy, recovery expectation, external-delivery consequence, data constraint, and investment guardrail should shape a real Day 1 working call? Decision affected The appropriate approval, recovery, external-delivery, data, and investment boundaries for a real engagement.	

Suggested next step

RECOMMENDATION

Contact GreyX to discuss a USD 350, three-business-day AI & Product Diagnostic.

Evidence

C1 — GreyX Diagnostic offer

Validation

A working call confirms whether the Diagnostic is the appropriate next engagement.

Contact GreyX about an AI & Product Diagnostic hello@greyx.dev

Evidence and provenance appendix

OBSERVED

This appendix publishes public-safe evidence labels, scope, confidence, and limitations only; source code, repository data, private operational material, raw graph or index data, prompts, credentials, local paths, and the internal code index are intentionally not published.

Sample scope:

Sample Diagnostic deliverable based on the fictional ForgeGraph demonstration—not a customer engagement.

E1 — Durable-state boundary Evidence record: Reviewed separation of backend-controlled durable state and ephemeral execution work. Confidence: High This public summary does not describe deployment topology, uptime, or performance.	E2 — Recoverable projection Evidence record: Reviewed reconstruction of a cache-backed projection from backend-owned data. Confidence: High This does not establish incident history or a reliability result.	E3 — Human approval state Evidence record: Reviewed explicit approval and decision state in the demonstration workflow. Confidence: High This does not establish a real customer's governance policy.
E4 — Retry-aware delivery Evidence record: Reviewed durable, retry-aware treatment of external delivery state. Confidence: High This does not establish delivery rate, scale, latency, or production operation.	E5 — Public architecture map Evidence record: Approved simplified public system sequence. Confidence: High The map is intentionally simplified and is not a production topology.	E6 — Fictional workspace capture Evidence record: Approved isolated fictional UI/workflow capture. Confidence: High It is not evidence of a client, customer, revenue, users, deployment, or outcome.
E7 — Fictional approval capture Evidence record: Approved isolated fictional approval-workflow capture. Confidence: High It is not evidence of a client, customer, revenue, users, deployment, or outcome.	E8 — Reviewed maintainability signal Evidence record: Read-only current controller/orchestration responsibility and relevant deterministic-coverage review, informed by bounded static-analysis candidates. Confidence: Medium Structural characterization only; no defect, incident, slowness, production risk, performance result, required rewrite, source path, identifier, line count, score, or excerpt.	C1 — GreyX Diagnostic offer Evidence record: Current GreyX offer facts for the AI & Product Diagnostic and follow-on options. Confidence: High The facts are not a ForgeGraph estimate, quote, spend record, engagement timeline, or commitment.