

CrewWork Runtime and Domain Architecture

author: Scott Weiss, CrewWork Engineering

revision: 2026-09-10

Executive summary

CrewWork is an autonomous software delivery platform that runs inside your network. It connects feature requests, errors, and quality findings to autonomous implementation, validation, and human review. Its central design rule is that a model's claim of success cannot satisfy the platform's completion-evidence requirement.

A shared Core runtime owns execution, budgets, recovery, and validation across product workflows. Domain services define the work; the workspace presents the plans, changes, and results people need to make decisions.

The application runs as a single-host Docker Compose stack; model serving is configured separately (Section 14). CrewWork is in active development; this paper describes implemented architecture, its controls, and the limits of what those controls establish.

Architecture at a glance

- One Core runtime contract for all non-chat execution, with a delegated inner runner behind it (Section 2).
- Separate runtime transports for chat and non-chat flows.
- Nineteen independent domain packages, with cross-domain coupling measured rather than asserted (Section 3).
- Workbench surfaces keep platform operations and project-scoped work separate (Section 4).
- Two-stage structured extraction, so the same revision-pinned Qwen3.8 27B generative model serves planning, coding, review, extraction, and chat; embeddings use a separate model (Section 15).
- Event envelope and projection model for coherent frontend state (Section 11).
- Provider-neutral model transport with local-first routing across configured workload pools.
- Typed context graph with budget-aware prompt compilation and layered hash caching (Section 5).

- Artifact quality gates with automatic retry, and a same-reason breaker that stops repeated identical failures (Section 9).
- Review-first remediation: isolated delivery branches for project fixes, draft pull requests for platform self-repair.

1. System Model

The platform is organized around a clear separation: Core API as transport, domain services as planners and projectors, model transport as a provider-configured boundary, runtime as execution authority, and explicit data and observability layers.

The Compose stack runs nginx, the React frontend, and the Core API; PostgreSQL, Redis, and Qdrant; the Events service; the Container Orchestrator, which holds the only Docker socket in the stack; three task-worker lanes (indexing, autonomous, and self-repair); and the observability stack described in Section 12.

Nginx owns every browser-reachable route: the frontend, the API and auth routes, the authenticated /ws WebSocket, the authorized preview proxy, and the monitoring proxy, which strips browser credentials before forwarding to Grafana, Prometheus, or Alertmanager. The Events service has no browser route; it is reached only at its service address, and all browser realtime traffic reaches Core through /ws. External MCP clients such as Claude Desktop and Cursor reach a stateless CrewMate endpoint that exposes six read-only, project-scoped tools (workspace listing and reading, code search, git status, and git diff) with a scoped API key.




Figure 1. One front door: nothing reaches the frontend, the API, a preview, or a dashboard except through nginx, and the Events service has no browser route at all.

2. Runtime Contract

Every execution path submits canonical work items to the Core runtime. Runtime performs policy checks, applies workspace and lease control, executes through configured model transports, and publishes canonical results and events.

Execution has two layers. The outer delivery runtime owns project state, queues, leases, budgets, model lanes, source control, conversation history, events, evidence projection, and terminal decisions. The inner runner receives a typed work item and owns the model, tool, edit, validate, and repair attempt. Eligible delivery work defaults to a delegated inner-runner execution backend that runs inside the Container Orchestrator; an in-process agentic loop stays behind the same contract as the operator-controlled fallback. Admission, budget, turn, and validation failures fail closed rather than switching backends silently.

The lifecycle of one work item, in order:

1. A domain trigger fires and the domain planner shapes the request.
2. The domain submits a canonical work item; runtime policies are applied.
3. Runtime takes a lease and a workspace lock, then a budget and retry policy.
4. The inner runner executes through the model transport to the configured host.
5. The repository diff is taken in the git worktree and validation and tests run against the changed workspace; a failure produces a diagnostics packet for a guided repair-and-retry loop.
6. The result and artifacts are persisted, execution events reach the Redis stream, domain projections update, and workbench surfaces receive them over WebSocket.




Figure 2. Self-repair gets a worker pool of its own, so platform work never shares a process with project delivery, and every pool reaches a model only through the runtime and the sandboxed inner runner. Synchronous reads go straight from Core API to domain services.

- Canonical result fields include status, reason code, files changed, diagnostics, and token usage.
- Attempt lifecycle is persisted and recoverable across worker restarts.
- In the in-process loop, a failed repair attempt persists a bounded conversation history and the next attempt resumes from it instead of restarting cold.
- Dry-run and policy constraints are enforced as first-class runtime outcomes.

3. Domain Model

Domains are independent by default. Each of the 19 domain packages, from Infinite Coder delivery, hotfix, and platform self-modification to search, suggestions, diagnostics, projects, and releases, owns its planning and projection logic while delegating execution to the Core runtime. No domain owns its own retry, budget, or lease stack. Cross-domain coupling is measured, not asserted.

Measured on 2026-09-03 at platform commit be012db42f, 73 cross-package import edges connect the 19 packages. Projects is the hub, imported by 17 of the other 18; HotFix (9 dependents), Delivery (8), and Search (7) follow. Secret management has no cross-domain edges in either direction. Six domains submit work through the shared base seam (delivery runtime, diagnostics, hotfix, projects, search, and suggestions); delivery runtime drives the kernel directly, and HotFix pull-request review and Search reasoning also call the generation kernel. Suggestions has no Qdrant coupling; its similarity runs in memory over embeddings, a model-transport dependency.

CrewWork is a multi-tenant platform: every project is owned by an organization and a specific team, not by an individual user, and the creator is retained only as audit metadata. Authorization is layered on top of domain isolation: a platform role ceiling combines with organization and team membership roles to govern access to projects, source control, and other sensitive operations.

4. Workbench Scope Model

The workbench is the operator-facing shape of the architecture: one persistent rail with four product destinations. Contextual workspaces, selected records, evidence, and utilities remain addressable without becoming another navigation level. These placement boundaries reflect committed source on 2026-09-21; inner labels may evolve.

- Work opens on Needs attention and an Outcomes collection. Start work creates a planned outcome or direct request; the outcome record carries its narrative, changes, logs, diagnostics, evidence, and human review decision.
- Projects is the directory and selected Project object. It owns readiness, current work and signals, assurance evidence, and release workspaces. Code is a contextual, project-required workspace rather than a rail destination.
- Security is the project-selected findings and scan-evidence destination. Preview is the project-selected application runtime with its canvas, logs, quality, repair, and identity evidence.
- Project settings, Organization administration, Account, Commands, Notifications, and CrewMate are utilities, not customer destinations.
- CrewWork Admin is permission-gated and global. It owns platform recovery, access and integrations, policy and cost, audit, and Platform Self-Repair without joining the four customer destinations.

5. Context Graph and Prompt Compilation

The system builds a typed directed acyclic graph of about two dozen node types, including repo summary, service map, docker-compose overview, nginx config, test overview, API contracts, preview spec, run history, and run failures. It then compiles them into deterministic prompts with budget-controlled sections. When context is dropped due to token limits, the system automatically increases the budget for the next run.

Layered hash caching over repository, step, snapshot, and prompt-bundle state ensures graphs are only rebuilt when the codebase or task state changes. Planning-context and enrichment nodes, including semantic context, context freshness, and authority docs, are compiled from the canonical context graph into every task prompt. Budget optimization tracks token usage per section and reallocates it across runs to maximize context coverage within model limits. The result is a compiled, budget-aware, auto-tuning prompt rather than a file dump, which is what lets a smaller self-hosted model work from the evidence that matters.

6. Infinite Coder Workflow

Infinite Coder coordinates program intent, delivery planning, and completion review. The planner emits one coherent feature slice with a single implementation step for the whole unresolved charter, so the inner runner can complete it in one continuous session; a separate non-editing validation step is added only when the repository's own terminating validation requires it. Intermediate checkpoints use deterministic validation and direct inner-runner repair. One whole-program semantic review runs after implementation is exhausted, and an incomplete verdict produces one consolidated repair packet on the same anchor. Execution is Core-runtime-backed, so coding behavior stays consistent with hotfix and other autonomous workflows.

A generated delivery plan for a user-owned program stops in a pending-approval state once it passes automated quality checks. Infinite Coder execution does not begin until the plan is explicitly approved or rejected. Service-owned platform self-repair programs remain auto-approved so the self-repair lane is not blocked on a human decision.




Figure 3. Two human decisions bracket a program: approving the plan before any code is written, and reviewing the evidence at the end. Between them the runtime, not the model, decides when to repair and when to stop.

7. Hotfix Remediation Workflow

Error ingestion from runtime, preview, and connected error monitoring enters a scoped triage pipeline. Remediation work is executed through the Core runtime and returned with artifacts and policy outcomes.




Figure 4. HotFix never takes the model's word: every candidate diff passes deterministic validation inside the loop, the branch stays STAGED until you push, and the linked error issue closes only when the run completes. Platform self-repair follows Figure 6 instead.

8. Search and Indexing

Code intelligence parses each file with Tree-sitter into symbols (functions, classes, methods, types) and records import, call, inheritance, and implementation edges between them. Each symbol is embedded individually rather than as a fixed-size text chunk, so retrieval returns whole definitions, and results combine vector similarity with exact lookups over the symbol graph. Retrieved excerpts enter the prompt through the same token budget as the rest of the context graph (Section 5).




Figure 5. Indexing and search share one embedding model and two stores, so a query is compared against vectors produced the same way the index was built.

9. Validation and Policy Enforcement

Validation is integrated as a runtime outcome, not a disconnected post-process, and it is not one monolithic pipeline. Plan quality, execution quality, completion evidence, and artifact quality evaluate different parts of the work. Enforced gates can stop a step; failures return a diagnostics packet that feeds a guided repair-and-retry loop. The provisioned test-run engine executes across 10 language stacks on allowlisted, digest-pinned images, with network isolation verified before any test command runs. Additional quality enforcement is mode-aware: shadow or matrix validation modes and per-gate flags determine which findings are advisory and which block progress. These settings do not make model-reported success admissible completion evidence. Platform self-repair promotion validates at a shadow, baseline, or full tier chosen from the changed files: shadow only for docs and Markdown, full for high-risk paths, baseline otherwise.

Separate from the per-attempt gates, a five-part aggregate quality gate governs publication: it evaluates tests, coverage, security findings, code-health analysis, and live-preview quality together as one canonical decision and posts as its own GitHub commit status. Default thresholds are 80 percent minimum line coverage, a minimum health score of 70, and zero tolerance for new critical or high severity findings.

The artifact-quality gate inspects every modified file and rejects placeholder implementations, scratch/debug files, no-op change sets, and non-code-only changes for code tasks, with per-gate enforcement flags controlling whether findings block completion. Failed checks trigger automatic retry with targeted feedback. A durable same-reason breaker permanently stops a step or plan after a default limit of three consecutive identical-reason failures. That count persists across worker restarts and lineage replacement, so a crash cannot reset the streak. The breaker counts attempts, not model turns: each attempt is a full working session with its own turn and tool-call limits, and a different failure reason starts a new streak. Infrastructure failures, such as an unavailable runner or validation service, are a separate class: they are retried without counting toward the streak and are never handed to the model as something to fix. Every program also carries time, token, and iteration budgets (by default 12 hours, 20 million tokens, and 1,000 iterations).




Figure 6. The candidate never grades itself. Everything that judges it belongs to the incumbent, and the only two exits, operator review and the draft pull request, both end with a person.

10. Error Intake Model

CrewWork supports platform-level and project-level error-monitoring integration with explicit connect state, webhook verification, threshold controls, and health visibility in domain surfaces.

- Health state is surfaced as operational status, not hidden inside logs.
- A successful HotFix remediation automatically marks the originating error issue resolved, closing the loop end-to-end rather than only updating CrewWork's internal triage record.
- A HotFix that is verified on a branch but not yet published carries a STAGED status, separate from resolved, so operators can see work that is finished but intentionally held back from release.

11. Realtime State and Projections

Frontend state is projection-driven. Event envelopes emitted from runtime and domain services feed websocket and API projections so UI status remains coherent with backend truth.

- Attempt and task status are emitted as explicit events with scope metadata.
- Runs timelines and token usage read from canonical projections.
- Transport endpoints avoid reconstructing lifecycle from loose text payloads.

12. Operations and Observability

CrewWork exposes service health, runtime diagnostics, event state, and remediation health through platform diagnostics surfaces. A full observability stack ships in the base stack for correlated metrics, logs, and traces: Prometheus, Grafana, Alertmanager, Loki, Alloy, and Tempo, with node, Redis, and Postgres exporters.

- OpenTelemetry traces across API, runtime, and service integrations.
- Every model invocation is wrapped in an OpenTelemetry span correlated to durable token-usage records; versioned per-provider, per-model pricing policies surface USD cost per call and in aggregate, without ever persisting prompt or output content.
- Error-monitoring issue ingestion and remediation loop visibility.




Figure 7. Three signal types, one Grafana. The only path out of the observability stack is the error-monitoring webhook into HotFix.

13. Security and Governance

CrewWork’s governance model combines role-based access control, workspace and path containment, audit trails, prompt trust boundaries, and policy-driven remediation constraints. It is underpinned by gVisor (runsc) kernel-boundary sandboxing for every untrusted or candidate code execution path, and by a scored, manifest-recorded self-modification assurance model for platform self-repair. See the Security page for the full protections, approval gates, and compliance-evidence detail, and Runtime Isolation for the sandboxing and resource-cap model.

14. Deployment Model

CrewWork is documented and positioned as a single-host Docker Compose stack for local or LAN deployments today. Application hosting is separate from model serving: generation endpoints can run on other LAN hosts, and semantic search uses a separate embedding model. Public model endpoints require explicit opt-in. The optional LAN release path builds project artifacts once and deploys them to other hosts on their own network through an outbound-only mutual-TLS release agent, with no inbound connectivity to the target. See the Compare page for the full deployment model, requirements, and honest tradeoffs.




Figure 8. Every arrow that touches the target host starts from the host. The Core API never connects in; it only answers polls, and a failed health check pages someone.

15. Scaling Strategy

Scaling starts operationally rather than architecturally: increase worker concurrency, route model traffic across configured model hosts, and partition high-throughput services while keeping the same runtime and domain contracts.

The model routing layer resolves one typed routing snapshot across execution, extraction, embeddings, and review pools, with task classification, host capacity scoring, and bounded saturation steering across hosts; when every capable route is saturated the request receives a finite, retryable refusal rather than a silent fallback. Provider-compatible transport keeps the model boundary configurable without tying execution semantics to one host implementation. The structured extraction pipeline uses a two-stage think-then-extract pattern that generates natural-language analysis with mandatory sections before typed data extraction.

Because execution semantics are centralized in the Core runtime, horizontal scaling does not require duplicating workflow logic in each product surface.

16. Evidence and Limitations

Completion evidence records what the platform actually checked. Passing those checks does not establish that every defect has been found, that the plan fully captured the user's intent, or that a change is suitable for production. A completed run does not imply that every available test or scanner ran.

The mandatory completion floor accepts platform-produced evidence, not the model's own success report. Additional quality checks can be advisory or enforced according to their configuration; a shadow-mode finding is not equivalent to a passed enforced gate. Review the recorded contract and results for each run.

No representative task-success rate or comparative model benchmark is claimed here. The package/import measurements in Section 3 are tied to their stated commit and date. They describe code structure, not product effectiveness. Scaling mechanisms describe the architecture's configuration options, not a measured capacity guarantee.

The development model configuration uses two generative hosts running the same Qwen3.8 27B artifact, with embeddings alongside one host. That is a reference setup, not a minimum hardware specification. Capacity depends on model quantization, context length, concurrency, repository size, and configured checks. Human review remains the final decision about what merges.

The architecture overview is the concise component map; Truth & Trust Engineering covers how completion is verified.