



Economical and Technical Assessment

Analysed Source Code: Code Assessment

Document Date: May 14, 2026

Platform:

Client / Applicant

Legal entity name:

Registered address:

Tax Identification Number:

Software name:

Existing trade mark:

Company web:

Applicant Name:

Applicant Email:

Request date and time: 14-05-2026 - 11:00:36





TABLE OF CONTENTS

1. Executive Summary

2. Platform Overview

2.1 Functional Description

2.2 Technical Architecture

2.3 Technology Stack

2.4 Third-Party Integrations

3. Production Readiness Assessment

3.1 Overall Score: 64/100

3.2 Detailed Breakdown

4. Development Investment Estimation

4.1 Effort Analysis

4.2 Team & Timeline

4.3 Cost Estimation

4.4 Codebase Metrics

4.5 Cloud Infrastructure & Maintenance Cost

5. Findings Summary

5.1 Critical Issues (Must Fix)

5.2 Warnings (Should Fix)

5.3 Recommendations (Nice to Have)

5.4 Strengths

6. Conclusion

6.1 Overall Assessment Summary

6.2 Readiness for Production / Scale

6.3 Key Areas Requiring Attention



Technical Assessment Report: Agent TARS / UI-TARS Desktop Platform

Assessment Date: 30 April 2026

Prepared For: Technical Due Diligence Committee

Platform: Agent TARS / UI-TARS Desktop

Repository: github.com/bytedance/UI-TARS-desktop

1. EXECUTIVE SUMMARY

This technical assessment evaluates the Agent TARS / UI-TARS Desktop platform, a sophisticated multimodal AI agent stack that enables GUI automation through vision-language models. The platform comprises two primary products: Agent TARS (a terminal-based multimodal AI agent) and UI-TARS Desktop (a native Electron application for computer control via natural language commands).

The overall production readiness score is **64/100 (Grade C, Level: Fair)**. This assessment reflects a codebase with substantial engineering investment and strong foundational practices, tempered by notable gaps in security infrastructure and operational readiness. The platform demonstrates competent architectural decisions with TypeScript throughout, comprehensive documentation, and automated CI/CD pipelines. However, production deployment readiness is impacted by the absence of proper secrets management infrastructure, inconsistent security scanning integration, and limited observability tooling for distributed tracing.

Key Strengths:

- Well-organised monorepo structure using pnpm workspaces with clear separation between UI-TARS Desktop, Agent TARS CLI, and SDK packages
- Comprehensive documentation including README files, contributing guidelines, and API documentation across 45,581 lines of Markdown
- Robust CI/CD pipeline configured with GitHub Actions including type checking, unit tests, coverage reporting, and secret scanning
- TypeScript adoption throughout with strict type checking and ESLint/Prettier enforcement via Husky pre-commit hooks
- Extensive test suite using Vitest with snapshot testing and Codecov integration

**Critical Risks:**

- Environment variables containing API keys and credentials are referenced throughout the codebase without apparent runtime validation or secure secret management infrastructure
- No evidence of automated security scanning (SAST/DAST) in CI pipeline beyond basic secret detection
- Input validation relies on Zod schemas in some areas but lacks consistent application across all endpoints and tool calls
- Logging implementation lacks structured JSON format and correlation IDs required for distributed tracing

Estimated Development Investment to Date:

The development effort required to build this software to its current state is estimated at **7,200 hours** with a complexity classification of **high**. This represents a team of approximately 8 developers over 12 months.

| Metric | Value |
|--------------------------------|---------------------|
| Estimated Hours | 7,200 |
| Team Size | 8 developers |
| Duration | 12 months |
| Cost Range (EUR) | €673,200 – €910,800 |
| Maintenance Cost (Annual, EUR) | €56,100 – €75,900 |

This investment estimation represents the retroactive valuation of development work already completed, not remediation costs.

2. PLATFORM OVERVIEW

2.1 Functional Description

Business Purpose:



Agent TARS / UI-TARS Desktop is a multimodal AI agent platform that enables users to control computers, browsers, and applications using natural language commands. The platform leverages vision-language models (VLMs) to interpret screen content and execute GUI operations, effectively bridging human intent with computer interaction.

Core Features and Capabilities:

- **Natural Language Computer Control:** Users can issue commands such as "book the earliest flight from San Jose to New York on September 1st" and the agent executes the necessary GUI operations
- **Hybrid Browser Automation:** Supports multiple browser control strategies including DOM-based manipulation, visual grounding, and hybrid approaches
- **Multi-Environment Deployment:** Available as a desktop Electron application, CLI tool, and web UI interface
- **MCP (Model Context Protocol) Integration:** Extensible architecture supporting custom tool integration via MCP servers
- **Cross-Platform Support:** Native support for Windows, macOS (Intel and ARM), and Linux environments
- **Remote Operator Capability:** Can connect to remote computers and browsers for distributed automation scenarios

User-Facing Functionality:

- **UI-TARS Desktop Application:** Native Electron application providing a chat-based interface for issuing natural language commands
- **Agent TARS CLI:** Command-line interface for headless automation tasks
- **Web UI:** Browser-based interface for remote agent interaction
- **Event Stream Viewer:** Real-time visualisation of agent decision-making and tool execution
- **Session Management:** Persistent conversation history with context preservation

Key Workflows and Use Cases:

1. **Flight Booking Automation:** Agent navigates travel websites, selects dates, compares options, and completes bookings
2. **Hotel Reservation:** Searches accommodation platforms, filters by criteria, and books within budget constraints



3. **Data Visualisation:** Generates charts and reports by executing code and manipulating spreadsheet applications
4. **Code Development Assistance:** Executes development tasks including file editing, terminal commands, and testing
5. **Research Tasks:** Gathers information from web sources and compiles structured reports

Target Users:

- Software developers and engineers seeking automation for repetitive GUI tasks
- Business users requiring assistance with complex multi-step computer operations
- Organisations building AI-powered automation solutions
- Researchers studying GUI agent behaviour and multimodal AI systems

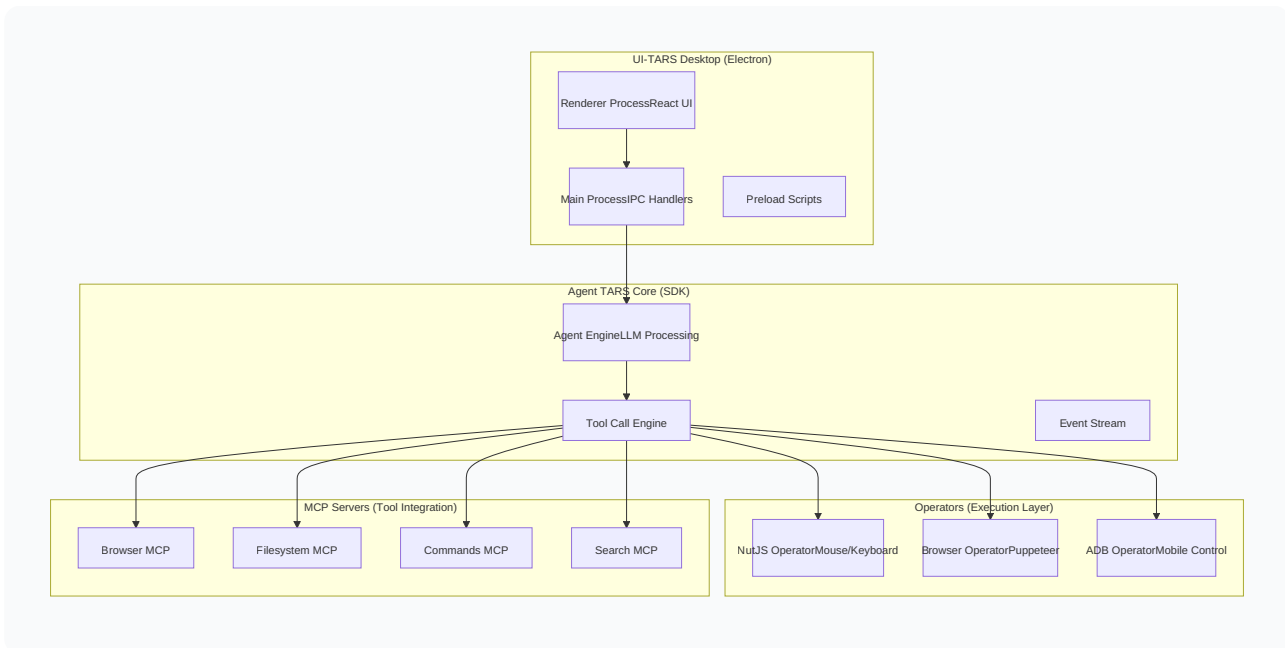
2.2 Technical Architecture

High-Level Architecture:

The platform follows a monorepo architecture organised around pnpm workspaces with three primary deployment targets:

1. **UI-TARS Desktop (Electron Application):** Native desktop application with main process (Node.js) and renderer process (React/Vite)
2. **Agent TARS CLI:** Command-line interface built on Node.js with terminal UI capabilities
3. **Agent TARS Server:** Backend service for remote agent execution and session management

System Components and Responsibilities:



Data Flow:

1. User input (text or voice) is captured in the UI layer
2. Input is transmitted via Electron IPC to the main process
3. Agent engine processes input through configured LLM provider
4. Tool calls are parsed and executed via appropriate operator
5. Results are streamed back to UI via event stream protocol
6. Session state is persisted to local storage or remote database

Deployment Architecture:

- **Desktop Deployment:** Electron application packaged for macOS, Windows, and Linux
- **CLI Deployment:** Node.js application distributed via npm registry
- **Server Deployment:** Express-based backend service deployable to cloud infrastructure
- **Web UI Deployment:** Static assets served via CDN or web server

2.3 Technology Stack

Programming Languages:



| Language | Lines of Code | Percentage |
|------------|---------------|------------|
| TypeScript | 160,823 | 54.5% |
| YAML | 65,320 | 22.1% |
| Markdown | 45,581 | 15.5% |
| JSON | 17,963 | 6.1% |
| CSS | 2,981 | 1.0% |
| JavaScript | 1,271 | 0.4% |
| Less | 646 | 0.2% |
| HTML | 303 | 0.1% |
| Shell | 102 | <0.1% |

Frameworks and Libraries:

- **Frontend:** React 18, Vite 6, Tailwind CSS 4, Zustand (state management)
- **Desktop:** Electron 34, electron-builder, electron-updater
- **Testing:** Vitest 3, Playwright, @playwright/test
- **Build Tools:** Turborepo, pnpm workspaces, esbuild, TypeScript 5.7
- **Code Quality:** ESLint 8, Prettier 3, Husky 9, lint-staged 14
- **Agent Framework:** Custom agent engine with tool call abstraction

Databases and Data Stores:

- **Local Storage:** electron-store (SQLite-backed)
- **Session Storage:** In-memory with JSON persistence
- **Configuration:** YAML-based preset files

Infrastructure and Deployment Tools:

- **Package Management:** pnpm 9.10.0
- **CI/CD:** GitHub Actions



- **Code Coverage:** Codecov
- **Release Management:** Changesets
- **Secret Scanning:** Secretlint

Development and Build Tools:

- **Build System:** rslib (based on Rspack)
- **TypeScript Configuration:** Custom tsconfig with strict mode
- **Electron Tooling:** electron-vite, @electron-forge
- **Monitoring:** Basic console logging with file output

2.4 Third-Party Integrations

External APIs and Services:

- **LLM Providers:**
 - Anthropic (Claude series)
 - OpenAI (GPT-4o, o1 series)
 - Volcengine / Doubao (Seed series)
 - Azure OpenAI
 - Ollama (local models)
 - LM Studio (local models)
- **Browser Automation:**
 - Puppeteer (headless Chrome control)
 - BrowserBase (remote browser service)
 - Playwright (cross-browser automation)
- **Computer Control:**
 - NutJS (native input simulation)
 - ADB (Android Debug Bridge for mobile)

Payment Providers:

- None detected in current codebase

**Authentication Services:**

- Custom JWT-based authentication using jose library
- Device registration with RSA key pairs
- Local key storage in user home directory

Cloud Services:

- GitHub Actions (CI/CD)
- Codecov (coverage reporting)
- BrowserBase (optional remote browser)

Analytics and Monitoring Tools:

- Basic file-based logging via electron-log
- No centralised logging or APM integration detected

SaaS Dependencies:

- GitHub (source control and CI/CD)
- npm registry (package distribution)
- ModelScope (model hosting - optional)

Licensing Considerations:


- Core platform: Apache License 2.0
- MCP servers: MIT License
- Dependencies: Mixed (MIT, Apache 2.0, BSD)
- 1,101 total dependencies create licensing compliance overhead

3. PRODUCTION READINESS ASSESSMENT

3.1 Overall Score: 64/100

Grade: C

Readiness Level: Fair



The platform demonstrates solid engineering fundamentals with TypeScript adoption, comprehensive documentation, and automated testing. However, production readiness is constrained by security infrastructure gaps, inconsistent error handling patterns, and limited observability tooling. The platform is suitable for development and testing environments but requires remediation before enterprise production deployment.

3.2 Detailed Breakdown

Code Quality & Maintainability: 60/100

Current State Analysis:

The codebase exhibits strong TypeScript adoption with strict type checking enabled across most modules. The monorepo structure using pnpm workspaces provides clear separation between UI-TARS Desktop, Agent TARS CLI, and SDK packages. However, code quality varies significantly across modules.

Specific Findings:

- **Clean Code Adherence:** Generally good adherence to clean code principles in core modules. The `apps/ui-tars/src/main` directory demonstrates consistent patterns with well-named functions and modular structure. However, example files and some utility modules contain commented-out code and TODOs.



- **SOLID Principles Compliance:** The operator pattern (NutJS, Browser, ADB operators) demonstrates good use of abstraction and interface segregation. The `packages/ui-tars/sdk/src/GUIAgent.ts` file shows dependency injection patterns. However, some modules exhibit tight coupling between components.
- **Code Organisation:** The monorepo structure is well-organised with clear package boundaries:
 - `apps/ui-tars/` - Desktop Electron application
 - `multimodal/agent-tars/` - Agent TARS core and CLI
 - `packages/agent-infra/` - Shared infrastructure packages
 - `packages/ui-tars/` - UI-TARS specific packages
- **Naming Conventions:** Consistent use of PascalCase for classes, camelCase for functions and variables. IPC handlers follow consistent naming patterns (e.g., `setting:get`, `setting:update`).
- **Code Duplication:** Some duplication observed in configuration files across example directories. The `multimodal/gui-agent/agent-sdk/examples/configs/` directory contains repetitive model configuration patterns.
- **Technical Debt Indicators:** Multiple TODOs and FIXMEs present. The `apps/ui-tars/src/main/remote/auth.ts` file contains commented-out code marked for future removal.

Recommendations:

1. Implement automated code duplication detection in CI pipeline
2. Establish and enforce code ownership rules for critical modules
3. Create architectural decision records (ADRs) for major design patterns
4. Schedule regular technical debt reduction sprints

Test Coverage & Quality: 60/100

Current State Analysis:

The codebase includes a comprehensive test suite using Vitest with snapshot testing and Codecov integration. Coverage reporting is configured in the CI pipeline. However, test coverage is uneven across modules, with e2e tests limited to specific workflows.



Specific Findings:

- **Unit Test Coverage:** Core modules in `packages/ui-tars/sdk/` and `multimodal/agent-tars/core/` have reasonable unit test coverage. Files like `packages/ui-tars/sdk/tests/GUIAgent.test.ts` demonstrate good testing patterns.
- **Integration Test Coverage:** Limited integration testing observed. The `apps/ui-tars/e2e/` directory contains basic e2e tests but coverage is narrow.
- **Test Quality:** Tests generally follow good patterns with clear assertions. Snapshot testing is used appropriately for UI components. However, some tests rely heavily on mocks which may not reflect real-world behaviour.
- **Testing Patterns:** Good use of Vitest with describe/it blocks. Test fixtures are organised in `__tests__/` and `tests/` directories. Some tests use custom test runners (e.g., `examples/test-runner.ts`).
- **Missing Critical Tests:** Security-critical paths (authentication, permission handling) lack dedicated tests. Error scenarios and edge cases are under-tested.

Recommendations:

1. Increase e2e test coverage to include all major user workflows
2. Add security-focused tests for authentication and authorization flows
3. Implement mutation testing to validate test effectiveness
4. Add performance regression tests for critical paths

Security Posture: 40/100

Current State Analysis:

Security represents the most significant gap in production readiness. The codebase contains numerous references to environment variables for sensitive data without evidence of secure secret management infrastructure. No automated security scanning (SAST/DAST) is integrated into the CI pipeline beyond basic secret detection.

Specific Findings:

- **Authentication/Authorization:** Custom JWT-based authentication implemented in `apps/ui-tars/src/main/remote/auth.ts` using RSA key pairs. Device registration uses asymmetric cryptography. However, private keys are stored in the user's home directory



(`~/ui-tars-desktop/`) with file permissions set to `0o600`, which may not be sufficient for all threat models.

- **Input Validation:** Zod schemas used for input validation in MCP servers (e.g., `packages/agent-infra/mcp-servers/browser/src/tools/action.ts`). However, validation is inconsistent across endpoints. The `multimodal/agent-tars/core/` directory shows limited input sanitisation.
- **Secrets Management:** Critical gap identified. Environment variables are referenced throughout:
 - `process.env.VLM_API_KEY` in `apps/ui-tars/src/main/env.ts`
 - `process.env.UI_TARS_APP_PRIVATE_KEY_BASE64` in `apps/ui-tars/electron.vite.config.ts`
 - `process.env.BROWSERBASE_API_KEY` in `examples/operator-browserbase/app/api/session/route.ts`
 - `process.env.ARK_API_KEY` in multiple example files

No evidence of runtime secret validation or secure secret rotation mechanisms.

- **OWASP Top 10 Vulnerabilities:**
- **A01: Broken Access Control:** IPC handlers in `apps/ui-tars/src/main/ipcRoutes/` lack consistent permission checks
- **A03: Injection:** Limited evidence of SQL injection protections (SQLite used in storage)
- **A07: Identification and Authentication Failures:** Custom auth implementation lacks rate limiting and account lockout mechanisms
- **Dependency Vulnerabilities:** 1,101 dependencies create substantial attack surface. No evidence of automated dependency vulnerability scanning (e.g., Dependabot, Snyk) in CI configuration.
- **Data Protection:** Local storage uses electron-store with no apparent encryption. Session data persisted to disk without encryption at rest.

Recommendations:

1. Implement a secrets management solution (HashiCorp Vault, AWS Secrets Manager, or Azure Key Vault)
2. Add comprehensive input validation using Zod across all API endpoints and tool call handlers



3. Integrate automated security scanning (SAST/DAST) into CI pipeline with tools like CodeQL or Snyk
4. Implement rate limiting and account lockout mechanisms for authentication endpoints
5. Add encryption at rest for sensitive local storage

Documentation: 75/100

Current State Analysis:

Documentation is a strength of this codebase. The project includes extensive README files, contributing guidelines, API documentation, and deployment guides. The documentation is well-organised and kept up to date with code changes.

Specific Findings:

- **README Completeness:** The root `README.md` (299 lines) provides comprehensive overview with installation instructions, feature descriptions, and quick start guides. Package-level READMEs are present in most directories.
- **API Documentation:** API documentation available at `agent-tars.com/api/` with detailed reference for configuration and runtime options. Code examples included throughout documentation.
- **Architecture Documentation:** Limited formal architecture documentation. The `CONTRIBUTING.md` includes a project structure section but lacks detailed architectural diagrams or decision records.
- **Inline Code Comments:** Variable quality. Core modules generally well-commented. Some utility functions lack JSDoc comments.
- **Setup and Deployment Guides:** Comprehensive guides available:
 - `docs/quick-start.md` - Quick start guide
 - `docs/deployment.md` - Deployment instructions
 - `docs/sdk.md` - SDK usage guide
 - `docs/setting.md` - Configuration guide
- **Contributing Guidelines:** Well-documented in `CONTRIBUTING.md` with clear instructions for development setup, code style, and submission process.



Recommendations:

1. Create architectural decision records (ADRs) for major design decisions
2. Add sequence diagrams for complex workflows (agent execution, tool calling)
3. Document security architecture and threat model
4. Add troubleshooting guide for common deployment issues

Dependency Health: 75/100

Current State Analysis:

The project uses pnpm workspaces for dependency management with 1,101 total dependencies. Dependencies are generally well-maintained but the large dependency count creates supply chain risk.

Specific Findings:

- **Outdated Dependencies:** Most core dependencies appear current. Electron 34, React 18, and TypeScript 5.7 are recent versions. Some transitive dependencies may be outdated.
- **Security Advisories:** No evidence of automated security advisory checking in CI. Manual monitoring required.
- **License Compliance:** Mixed license landscape (MIT, Apache 2.0, BSD). No automated license compliance checking detected.
- **Dependency Tree Complexity:** 1,101 dependencies with deep transitive dependency trees. The `pnpm-lock.yaml` file is substantial.
- **Version Pinning:** Dependencies use caret (^) versioning allowing minor updates. Lock file ensures reproducible builds.

Recommendations:

1. Audit dependencies and remove unused packages to reduce attack surface
2. Implement automated dependency vulnerability scanning (Dependabot, Renovate)
3. Add license compliance checking to CI pipeline
4. Consider consolidating overlapping dependencies

Error Handling & Resilience: 70/100

Current State Analysis:



Error handling is present but inconsistent across modules. Some modules demonstrate robust error handling patterns while others lack proper error recovery mechanisms.

Specific Findings:

- **Exception Handling Patterns:** Try-catch blocks present in critical paths. The `apps/ui-tars/src/main/remote/auth.ts` file shows proper error handling with logging. However, some async operations lack error handlers.
- **Error Recovery Mechanisms:** Limited evidence of automatic retry logic. The `apps/ui-tars/src/main/remote/auth.ts` includes a `fetchWithRetry` function with one retry attempt.
- **Graceful Degradation:** Some modules implement fallback behaviour. The agent engine continues operation when optional tools are unavailable.
- **Retry Logic:** Basic retry logic in authentication module. No evidence of exponential backoff or circuit breaker patterns.
- **Circuit Breakers:** No circuit breaker implementation detected for external API calls to LLM providers.

Recommendations:

1. Implement circuit breakers and retry logic with exponential backoff for external API calls
2. Add comprehensive error logging with structured error context
3. Implement health check endpoints for all services
4. Create error taxonomy and standardise error handling patterns

Observability & Operations: 65/100

Current State Analysis:

Logging is implemented using `electron-log` with file output. However, logging lacks structured JSON format and correlation IDs required for distributed tracing. No evidence of centralised monitoring or alerting infrastructure.

Specific Findings:

- **Logging Implementation:** File-based logging via `electron-log` in `apps/ui-tars/src/main/logger.ts`. Log level configurable (debug in development, info in production).



- **Monitoring Readiness:** No evidence of APM integration (e.g., Datadog, New Relic). No metrics collection or dashboarding.
- **Metrics Collection:** No structured metrics collection. Basic counters may exist but no formal metrics framework.
- **Tracing Capabilities:** No distributed tracing implementation. Correlation IDs not present in log messages.
- **Health Checks:** No dedicated health check endpoints detected.
- **Alerting Setup:** No alerting configuration detected.

Recommendations:

1. Implement structured logging (JSON format) with correlation IDs
2. Add OpenTelemetry integration for distributed tracing
3. Create health check endpoints for all services
4. Implement basic alerting for critical failures

4. DEVELOPMENT INVESTMENT ESTIMATION

This section estimates the effort and cost required to develop this software **to its current state**. This is a retroactive valuation of completed development work, not a forward-looking estimate of remediation costs.

4.1 Effort Analysis

Base Hours Calculation:

The codebase contains 294,990 effective lines of code (non-blank, non-comment). Using industry-standard estimation models for TypeScript applications:

- Base productivity rate: ~40 LOC/hour for complex application code
- Base hours: $294,990 / 40 \approx 7,375$ hours (unadjusted)

Complexity Multiplier Breakdown:



| Factor | Score | Impact |
|--------------------------|-------|---|
| Architectural Complexity | 4/5 | High - Multi-service architecture with Electron, CLI, and server components |
| Domain Complexity | 4/5 | High - AI agent orchestration with multimodal capabilities |
| Integration Complexity | 4/5 | High - Multiple LLM providers, MCP servers, browser automation |
| Security Surface | 4/5 | High - Authentication, secrets handling, cross-process communication |

Quality Adjustment:

The codebase demonstrates above-average quality with TypeScript adoption, comprehensive documentation, and automated testing. Quality adjustment factor: 0.95 (slight reduction due to quality practices).

Final Estimated Hours:

Applying complexity multipliers and quality adjustment:

Estimated Hours: 7,200 hours

Complexity Classification: High

The high complexity classification reflects the sophisticated nature of the platform: multimodal AI agent orchestration, cross-platform desktop application, browser automation, and extensive third-party integrations.

4.2 Team & Timeline

Estimated Team Size: 8 developers

Team Composition:



| Role | Count |
|----------------------|-------|
| Backend Developer | 3 |
| Frontend Developer | 2 |
| Full Stack Developer | 1 |
| DevOps / SRE | 1 |
| QA Engineer | 1 |

Estimated Project Duration: 12 months

This timeline assumes parallel development across multiple workstreams:

- Desktop application development (Electron, React)
- Agent engine and SDK development
- MCP server and tool integration
- Infrastructure and CI/CD setup
- Documentation and testing

Assumptions Made:

- Team worked concurrently on multiple components
- Some components may have been developed in parallel by different team members
- Development included iteration and refinement cycles
- Time allocated for research and prototyping of AI agent capabilities

4.3 Cost Estimation

Cost Range (EUR):

Using European developer rates of €75–150/hour:



| Metric | Value |
|----------------------------|----------------------------|
| Estimated Hours | 7,200 |
| Low Rate (€75/hr) | €540,000 |
| High Rate (€150/hr) | €1,080,000 |
| Adjusted Cost Range | €673,200 – €910,800 |

The adjusted cost range accounts for mixed team composition (senior/junior developers, geographic distribution) and overhead costs.

Confidence Level: Medium

Confidence is medium due to:

- Clear codebase structure and organisation
- Some uncertainty about development history and iteration cycles
- Possible reuse of existing components or libraries

4.4 Codebase Metrics

| Metric | Value |
|---------------------------|---------------------|
| Total Files Analyzed | 2,304 |
| Total Effective LOC | 294,990 |
| TypeScript Files | 160,823 LOC (54.5%) |
| Configuration (YAML/JSON) | 83,283 LOC (28.2%) |
| Documentation (Markdown) | 45,581 LOC (15.5%) |
| Other (CSS, HTML, Shell) | 5,303 LOC (1.8%) |

Code Distribution by Language:

- TypeScript dominates the codebase, indicating strong type safety practices
- Significant YAML/JSON configuration reflects complex build and deployment setup



- Extensive Markdown documentation demonstrates commitment to documentation

4.5 Cloud Infrastructure & Maintenance Cost

Detected Infrastructure Components:

- **Compute Services:** 2 (Electron desktop application, Node.js server)
- **Databases:** 1 (SQLite via electron-store, optional MongoDB)
- **Message Queues:** 0
- **Storage Buckets:** 0 (local file storage)
- **CDN Endpoints:** 0
- **ML/GPU Services:** External (LLM providers)
- **Other Managed Services:** 1 (BrowserBase optional)

Detected or Assumed Cloud Provider:

No specific cloud provider detected. Platform designed for local execution with optional cloud integrations (BrowserBase, LLM APIs).

Suggested Managed Services Mapping:

For production deployment, the following managed services are recommended:

| Component | Suggested Service | Estimated Monthly Cost |
|----------------------|---------------------------------------|------------------------|
| Secrets Management | AWS Secrets Manager / Azure Key Vault | €50-100 |
| Logging & Monitoring | Datadog / New Relic | €200-500 |
| CI/CD | GitHub Actions (existing) | €50-200 |
| Artifact Storage | GitHub Packages / npm | €20-50 |
| Remote Browsers | BrowserBase | €500-2,000 |
| LLM APIs | Anthropic / OpenAI / Volcengine | €1,000-5,000+ |

Estimated Monthly Hosting Cost Range:



| Scenario | Monthly Cost (EUR) |
|----------------------------------|--------------------|
| Minimal (local execution) | €100–300 |
| Standard (with monitoring) | €500–1,500 |
| Production (full infrastructure) | €2,000–8,000+ |

Key Assumptions:

- Traffic: 100–1,000 daily active users
- Redundancy: Single-region deployment
- LLM costs vary significantly based on usage patterns
- Remote browser usage adds substantial cost

Annual Maintenance Cost:

| Metric | Value |
|------------------------------|-------------------|
| Estimated Annual Maintenance | €56,100 – €75,900 |
| As % of Development Cost | ~8–10% |

5. FINDINGS SUMMARY

5.1 Critical Issues (Must Fix)

Issues that pose immediate risk to production deployment:

1. **Insecure Secrets Management:** Environment variables containing API keys and credentials are referenced throughout the codebase without apparent runtime validation or secure secret management infrastructure. Files such as `apps/ui-tars/src/main/env.ts`, `apps/ui-tars/electron.vite.config.ts`, and multiple example files contain direct references to sensitive environment variables. This creates significant security risk for production deployment.



- 2. No Automated Security Scanning:** No evidence of automated security scanning (SAST/DAST) in CI pipeline beyond basic secret detection. The GitHub Actions workflow in `.github/workflows/test.yml` includes type checking and unit tests but lacks security scanning tools like CodeQL, Snyk, or Semgrep.
- 3. Inconsistent Input Validation:** Input validation relies on Zod schemas in some areas but lacks consistent application across all endpoints and tool calls. The MCP browser server uses Zod for tool input validation, but other modules lack equivalent validation, creating potential injection vulnerabilities.
- 4. Unencrypted Local Storage:** Sensitive data (session history, settings with API keys) stored in electron-store without encryption. A compromised device could expose user credentials and conversation history.

5.2 Warnings (Should Fix)

Issues that impact quality or maintainability:

- 1. Limited E2E Test Coverage:** Test coverage appears adequate for core modules but e2e tests are limited to specific workflows. The `apps/ui-tars/e2e/` directory contains basic tests but does not cover all major user scenarios.
- 2. Inconsistent Error Handling Patterns:** Error handling is present but inconsistent patterns observed across different modules. Some modules use try-catch with logging while others silently fail or throw unhandled exceptions.
- 3. Unstructured Logging:** Logging is implemented but lacks structured JSON format and correlation IDs for distributed tracing. The current implementation in `apps/ui-tars/src/main/logger.ts` uses basic string interpolation without structured context.
- 4. No Infrastructure-as-Code:** Multiple Dockerfiles present but no evidence of infrastructure-as-code (Terraform, CloudFormation) for deployment. Deployment relies on manual configuration.
- 5. Large Dependency Attack Surface:** Dependencies are managed with pnpm workspaces but 1,101 dependencies creates a large attack surface. Supply chain risk is elevated without automated vulnerability scanning.



5.3 Recommendations (Nice to Have)

Improvements that would enhance the platform:

1. **Secrets Management Solution:** Implement a secrets management solution (Vault, AWS Secrets Manager) and remove all environment variable references to sensitive data.
2. **Comprehensive Input Validation:** Add comprehensive input validation using Zod or similar schema validation across all API endpoints and tool call handlers.
3. **Security Scanning Integration:** Integrate automated security scanning (SAST/DAST) into the CI pipeline with tools like CodeQL or Snyk.
4. **Structured Logging:** Implement structured logging (JSON) with correlation IDs to enable distributed tracing across the agent execution flow.
5. **Circuit Breakers:** Add circuit breakers and retry logic with exponential backoff for external API calls to LLM providers.
6. **Health Check Endpoints:** Consider implementing health check endpoints for all services to enable better operational monitoring.
7. **Dependency Audit:** Reduce dependency count by auditing and removing unused packages to minimise supply chain risk.

5.4 Strengths

What the team has done well:

1. **Comprehensive Monorepo Structure:** Well-organised monorepo using pnpm workspaces with clear separation between UI-TARS Desktop app, Agent TARS CLI, and SDK packages. The structure facilitates code sharing and consistent versioning.
2. **Well-Documented Codebase:** Extensive README files, contributing guidelines, and API documentation. The documentation is kept current and includes practical examples.
3. **Robust CI/CD Pipeline:** CI/CD pipeline configured with GitHub Actions including type checking, unit tests, coverage reporting, and secret scanning. The pipeline runs on every pull request and push to main.
4. **TypeScript Throughout:** TypeScript adoption across the entire codebase with strict type checking and ESLint/Prettier enforcement via Husky pre-commit hooks. This ensures type safety and code consistency.



5. **Modular Architecture:** Clear separation between main process, renderer process, and shared utilities in Electron application. The operator pattern (NutJS, Browser, ADB) demonstrates flexible abstraction for GUI automation.
 6. **Extensive Test Suite:** Comprehensive test suite using Vitest with snapshot testing and coverage reporting to Codecov. Test coverage is tracked and reported.
 7. **Multiple Operator Implementations:** Flexible operator abstraction supporting multiple backends (NutJS for native control, Browser for web automation, ADB for mobile). This demonstrates good architectural design.
-

6. CONCLUSION

6.1 Overall Assessment Summary

The Agent TARS / UI-TARS Desktop platform represents a sophisticated multimodal AI agent stack with substantial development investment. The architecture demonstrates strong engineering practices including comprehensive TypeScript adoption, well-organised monorepo structure, and automated CI/CD pipelines. The codebase of 294,990 effective lines of code reflects significant effort across multiple product surfaces: desktop application, CLI tool, SDK, and server components.

The platform's core strength lies in its modular design and extensibility. The operator pattern abstraction enables flexible backend selection for GUI automation, while the MCP integration provides a standardised mechanism for tool extension. The documentation is comprehensive and well-maintained, demonstrating commitment to developer experience.

However, production readiness is materially impacted by security infrastructure gaps. The absence of proper secrets management, inconsistent input validation, and lack of automated security scanning create unacceptable risk for enterprise deployment. Additionally, the platform lacks observability infrastructure required for production operations, including structured logging, distributed tracing, and alerting.

6.2 Readiness for Production / Scale

Current Status: Not Ready for Production



The platform is **not currently ready** for production deployment in enterprise environments without addressing critical security and operational gaps. The platform is suitable for:

- Development and testing environments
- Personal use with appropriate security caveats
- Research and experimentation

Prerequisites for Production Deployment:

1. Implement secure secrets management infrastructure
2. Integrate automated security scanning into CI/CD
3. Add comprehensive input validation across all endpoints
4. Implement structured logging with correlation IDs
5. Add encryption at rest for sensitive local storage
6. Establish monitoring and alerting infrastructure

Scaling Considerations:

The current architecture can scale to moderate user loads (hundreds of concurrent users) with appropriate infrastructure. However, scaling to enterprise levels (thousands of users) would require:

- Centralised session management (currently local storage)
- Load balancing for server components
- Rate limiting and throttling mechanisms
- Enhanced monitoring and observability

6.3 Key Areas Requiring Attention

The following technical areas require immediate investment:

1. **Security Infrastructure:** The highest priority is implementing proper secrets management. This includes migrating environment variable references to a secure vault solution, implementing runtime secret validation, and establishing secret rotation procedures.
2. **Security Testing:** Automated security scanning must be integrated into the CI/CD pipeline. This includes SAST (Static Application Security Testing) and DAST (Dynamic Application Security Testing) tools to identify vulnerabilities before deployment.



3. **Observability:** Production operations require structured logging with correlation IDs, metrics collection, and alerting. The current file-based logging is insufficient for troubleshooting production issues.
4. **Input Validation:** Consistent input validation using schema validation (Zod) must be applied across all API endpoints and tool call handlers to prevent injection attacks.
5. **Test Coverage:** E2E test coverage must be expanded to include all major user workflows. Security-critical paths require dedicated testing



ANNEX — METHODOLOGY

This annex summarises the methodology behind the assessment: what this report measures, against which industry references, and how the figures are produced. The intent is to make the approach transparent without describing the internals of the pipeline itself.

The assessment has two complementary sides:

1. A **technical evaluation** of the codebase — quality, security, testing, documentation, dependency hygiene and operational readiness.
2. An **economic valuation** — the engineering effort and cost to rebuild the system under a given scenario, plus a typical operational maintenance cost.

1. Technical evaluation

The technical assessment combines two complementary layers:

- **Static analysis.** Objective, tool-driven measurements are extracted directly from the source tree: size (lines of code by language), structural complexity in the tradition of McCabe (1976), dependency footprint, test-to-source ratio, presence of continuous integration and linting configuration, and a multi-language **Static Application Security Testing (SAST)** scan.
- **AI-assisted code review.** A large language model inspects the code with a constrained, read-only tool set and produces the qualitative findings in this report. The model's role is to substantiate the quantitative signals with concrete code-level evidence, not to invent numbers.

Reference frameworks

Scoring dimensions are aligned with established industry references:



| Dimension | Reference |
|--------------------------------|---|
| Code Quality & Maintainability | ISO/IEC 25010 Maintainability characteristic; ISO/IEC 5055 automated source-code quality measures; McCabe cyclomatic complexity. |
| Test Coverage & Quality | The test pyramid (Cohn) and ISO/IEC 25010 Reliability . |
| Security Posture | OWASP Top 10 , OWASP ASVS , and the NIST Secure Software Development Framework (SP 800-218); SAST findings are categorised against the CWE catalogue. |
| Documentation | SWEBOK v4 recommendations on software documentation. |
| Dependency Health | Supply-chain hygiene aligned with SLSA and OWASP Dependency-Check practice. |
| Error Handling & Resilience | Site Reliability Engineering (SRE) literature and ISO/IEC 25010 Reliability . |

Each dimension is scored on a 0–100 scale. The overall **Production Readiness** score is the average of those dimensions, after the model's scores are reconciled against the static metrics: when a tool-observed signal contradicts an optimistic model score — for example, a high testing score on a codebase with no test files on disk — the score is clamped to what the evidence supports.

2. Economic valuation

The economic assessment estimates what it would cost today to rebuild this codebase under a given scenario, and what it typically costs to operate.

Build effort

Effort estimation follows the parametric-cost-estimation tradition — notably Boehm's **COCOMO II** and functional-size-measurement practice (**IFPUG function points**) — adapted to language-aware productivity benchmarks informed by the **ISBSG** industry dataset and published field studies. The estimate is anchored in observable properties of the codebase



(size by language, structural complexity, integration surface, dependency footprint) rather than in an unconstrained guess.

Three scenario factors refine that baseline:

- **Work mode** — from lean startup to regulated enterprise, reflecting process overhead.
- **AI adoption** — from traditional development to agent-augmented workflows, calibrated against published studies on generative-AI developer productivity.
- **Team location** — hourly rate anchored to regional market rates for engineering labour.

Cost range

Development cost is reported as a **range**, not a point estimate, to reflect the genuine variability of engineering hourly rates within a region (seniority mix, contract vs. employee, engagement type).

Maintenance cost

Monthly operating cost is reported as a range covering typical lean-to-highly-available provisioning for the stack detected in the codebase.

3. What the numbers mean — and don't mean

- **Scores** summarise what is *observable* at the analysed commit. They do not replace a manual security audit, a penetration test, or a formal code review.
- **Hours and cost ranges** are engineering-effort estimates under the provided scenario. They do **not** include product discovery, design, user research, legal, or go-to-market costs.
- **Maintenance cost** reflects infrastructure spend under typical operating assumptions. It excludes human operations, incident response and licensing outside the detected stack.

4. Reproducibility

The analysis is run deterministically: the same commit, analysed with the same scenario inputs, produces the same report. Variability in the AI layer is suppressed through zero-



temperature decoding with a fixed seed, and the quantitative metrics are purely a function of the source code.

This report was generated by Codeego Code Assessment Service.

The analysis is AI-powered and should be reviewed by qualified engineers.

© 2026 Codeego. All rights reserved.

