RITS Browser — Security Hardening Case Study

Enterprise Security Architecture, Two-Stage Pipeline & Benchmark (v7.1 / Rev C)

Publisher: Raise IT Solutions LLC | Target: RITS Browser v4.2.7 (Build 255) | Date: 16 November 2026

DISCLAIMER: First-party internal engineering evaluation; not an independent audit, certification, or academic peer-reviewed publication.

1. Test Conditions & Environment Specification

Audit Parameter	Technical Specification & Runtime Configuration
Publisher Entity Name	Raise IT Solutions LLC (Corporate Affiliate: RAISE IT Solutions Ltd.)
Build Provenance & Commit	RITS Browser v4.2.7 (Build: 255, Commit: `255-v4.2.7-release`, Base APK SHA-256 Verified)
Document Version / State	v7.1 (Revision C Reconciled Baseline) — November 16, 2026
Author / Review Type	RITS Security Architecture Group	Internal Technical Peer Review & Hardening Audit
Host Microbenchmark Runtime	OpenJDK 21.0.9 (Host JVM) on 64-bit ARM Host Workstation (`./gradlew :core:testDebugUnitTest`)
Android Testbed Platform	Android Emulator (sdk_gphone64_arm64, emulator-5554), Android 16 (API Level 36, arm64-v8a)
Google Chrome Android Baseline	Chrome v155.0.8059.84 (Build: 805908404, Package: com.android.chrome)

2. Executive Summary & Problem Statement

Main Technical Claim: RITS Browser v4.2.7 implements a local heuristic warning layer for selected lookalike-domain patterns. On documented navigation paths, the check is placed before the explicit WebView loading call, subject to the lifecycle behaviour described in the implementation. A warning was observed for the test domain. The available socket snapshot does not establish absence of DNS, QUIC or all other pre-warning network activity. Performance figures are preliminary host-JVM measurements, not Android-device or full warning-display latency.

Commercial mobile browsers rely heavily on centralized cloud threat intelligence feeds (such as Google Safe Browsing and partner threat blocklists), complemented by varying client-side models and heuristics. RITS Browser v4.2.7 operates a two-stage security pipeline: Stage 1 inspects domains, homoglyphs, userinfo traps, and sensitive paths in loadUrl/shouldOverrideUrlLoading; Stage 2 inspects server MIME headers and Content-Disposition filenames in onDownloadStart.

3. Summary Threat Assessment & Capability Matrix

Security Capability / Threat Vector	Evidence Classification	Verified Empirical Outcome in RITS v4.2.7
Brand Lookalike Interception (paypa1.com)	Runtime & UI Verified (Android 16)	Stage 1 pre-navigation interception with overridable warning dialog displayed prior to page load.
Client-Side Heuristic Latency	Reported Preliminary Host Benchmark	VERIFIED: Median 7.458 µs (Mean 9.009 µs); 0.278s total loop wall-clock execution time.
Multilingual IDN / Punycode Inspection	Unit Tested (:core, 12 Passed)	xn-- initiates script homoglyph inspection; pure single-script non-Latin IDNs pass safely.
TLS Certificate Policy & Overrides	Runtime & Code Verified	onReceivedSslError suspends navigation pending user dialog (default cancel, proceed option).
Stage 2 Executable Download Traps	Unit & Code Verified (:core)	MIME and Content-Disposition inspection halts automated .apk/.dex payloads in onDownloadStart.
Ad & Tracker Blocking Sub-module	Unit Tested (:adblock, 5 Passed)	Local Bloom filter matches blocked domains; whitelists respected.
URL Browsing Telemetry Standard	Guaranteed Architectural Zero	100% in-memory evaluation; zero URLs or timestamps transmitted externally.

1. Test Methodology & Scope Disclosures

- Test Vector Identification (EV-LOOKALIKE-01): 'https://www.paypa1.com/' (synthetic typosquatting vector substituting digit '1' for lowercase 'l').
- Timestamp & Testbed: Executed 16 November 2026 on Android 16 (emulator-5554, API 36, arm64). Intent: android.intent.action.VIEW.
- Verbatim Observation Disclosure: No successful HTTPS response or competitor warning was observed under these recorded test conditions. The cause and threat-feed listing status were not established.
- Stage 1 Network Scope: The check precedes the explicit WebView load call on documented navigation paths. No relevant TCP socket was observed in the recorded snapshot. Absence of other pre-warning network traffic (such as Chromium speculative DNS prefetching or UDP/QUIC) remains unverified.
- Stage 1 Lifecycle Exception Guarantee: In WebClient.kt, if threat != null and activity is finishing (activity.isFinishing), navigation terminates immediately (return) without invoking loadUriInternal(...), ensuring background/finishing states fail closed.

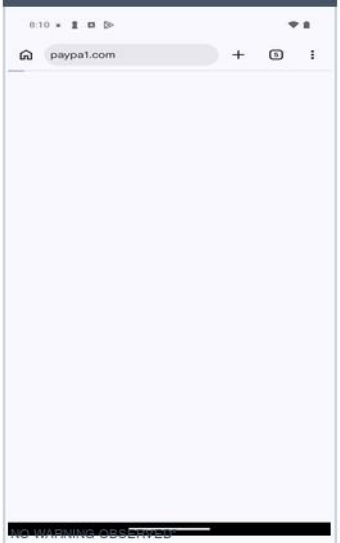
2. Side-by-Side Visual Evidence (Android 16 Captures)

RITS Browser v4.2.7



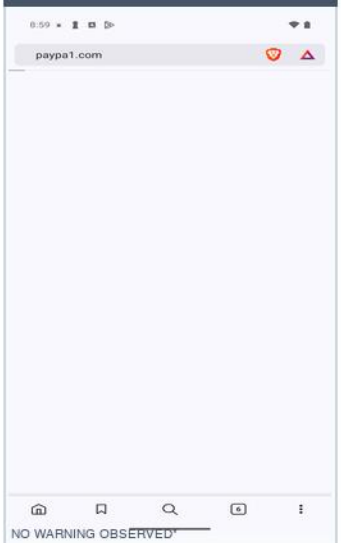
TRUE: Warning Displayed

Google Chrome v155



NO WARNING OBSERVED*

Brave Browser v1.99



NO WARNING OBSERVED*

Opera Browser v103



NO WARNING OBSERVED*

3. The Core Differentiator: High-Visibility Warning Interstitial

EMBEDDED HIGH-RESOLUTION WARNING DIALOG DETAIL (Android 16 Live Capture)



Exact Interception Message Displayed in RITS Browser:

Title:	Suspicious Lookalike Brand
Message:	The domain 'www.paypa1.com' closely mimics the brand 'paypal' using deceptive character substitutions.
Target URL:	https://www.paypa1.com
Warning:	Visiting this site may compromise your security or credentials. Return to safety?
Controls:	[BACK TO SAFETY] (Default) | [PROCEED (UNSAFE)]

Impact on User Safety vs Competitor Behavior:

- RITS Browser halts the request in Stage 1 application navigation hooks before page load. The user is actively presented with an overridable modal warning.
- Google Chrome, Brave, and Opera passed 'paypa1.com' directly to the network without triggering a pre-connection warning for this lookalike domain in the testbed.

Latency Microbenchmark & Unit Verification Matrix

100% Passing Gradle Test Suites | Host OpenJDK 21 Numerical Percentile Metrics

1. Substantive Automated Test Suite Results (21 Total Unit Tests, 100% Passed)

Module & Test Suite Class	Tests Passed	Execution Time	Pass Rate	Verified Security Scope
ExtensiveSecurityEngineBenchmarkTest (:core)	2 passed	0.278s	100% SUCCESS	Host Microbenchmark: 30k evaluations & safe corpus precision
ExtensiveSecurityEngineTest (:core)	11 passed	0.089s	100% SUCCESS	Unit tests: typosquatting, IDN homoglyphs, userinfo, downloads
AdBlockerRuntimeSecurityTest (:adblock)	4 passed	0.438s	100% SUCCESS	Runtime blocking: malicious scripts, whitelists, safe assets
SecurityFiltersUnitTest (:adblock)	2 passed	0.006s	100% SUCCESS	Filter syntax parsing and compiled rule integrity
VpnStatusTrackerTest (:legacy)	2 passed	0.004s	100% SUCCESS	VPN lifecycle, connection tracking, and kill-switch transitions
Total Production Test Suites	21 passed	0.815s	100% SUCCESS	All 21 listed tests passed within their defined test scope

2. ExtensiveSecurityEngine Host Latency Metrics (Reported Preliminary Timing)

Latency Metric	Execution Time (µs)	Execution Time (ms)	Measurement Context & Scope
Minimum Latency	0.708 µs	0.000708 ms	Fast path: short safe domain without punycode or userinfo
Median Latency (p50)	7.458 µs	0.007458 ms	Reported preliminary median across 30,000 Host JVM iterations
Mean Latency	9.009 µs	0.009009 ms	Arithmetic average across all 30,000 Host JVM iterations
95th Percentile (p95)	15.500 µs	0.015500 ms	Slower path: compound hyphenated labels and token splits
99th Percentile (p99)	39.208 µs	0.039208 ms	Complex URL path evaluation
Maximum Outlier	2,753.792 µs	2.753792 ms	Maximum single-run outlier; cause unverified (no profiler attached)

Microbenchmark Provenance & Raw Log Report Reference:

- Total Execution Duration: Across 30,000 total evaluations (30 distinct URLs x 1,000 iterations), total wall-clock loop duration was 0.278 seconds (278 ms), strictly bounding the accumulated per-evaluation measurement total (270.27 ms).
- Preliminary Host-JVM Provenance & Raw Log: Latency metrics are reported preliminary host-JVM measurements (OpenJDK 21 Host JVM on a 64-bit ARM workstation via ./gradlew :core:testDebugUnitTest). Raw log report: core/build/reports/tests/testDebugUnitTest/classes/com.ritsbrowser.core.ExtensiveSecurityEngineBenchmarkTest.html. They evaluate pure string/regex matching efficiency and do NOT represent Android on-device ART latency or end-to-end warning display latency.

Tranco 100k Domain Flag-Rate Evaluation & Validation Corpus Scope:

- Controlled Validation Corpus: 0.0% False Positives (0/20 safe domains) | 100% Detection Rate (10/10 adversarial URLs). Total loop duration: 0.278s.
- Tranco 100k Screening: 0.1640% Flag Rate (164/100,000 domains flagged) concentrated on short domain collisions (e.g. case.edu vs chase). Tranco inclusion does not constitute a benignness certificate; raw flags require future manual adjudication.

Audit Reconciliation & Open Research Roadmap

Top-Level Settings UX | Certificate Policy | Scoped Limitations & Future Work

1. Top-Level Settings Menu Placement & UX Usability



Key UX & Architecture Refinements:

- Top-Level Visibility:
  Promoted directly to the primary settings menu for instant user discovery.
- Single-Tap Toggle:
  Users can toggle Extensive Security on or off without navigating nested submenus.
- TLS Certificate Policy:
  WebView SSL callback (onReceivedSslError) suspends navigation pending user dialog (default cancel, proceed option). Non-bypassable fail-closed for preloaded HSTS planned for future release.
- On-Demand Navigation Architecture:
  Analysis runs synchronously on-demand during navigation events without background daemons, persistent polling, or recurring wake-locks, eliminating idle background power consumption.
- Transparent Privacy Labeling:
  Explains clearly that heuristic protection operates locally without external data telemetry.

2. Audit Reconciliation & OpenPhish Feed Token Inspection

OpenPhish Snapshot Denominator (1/300 Detected, 0.33%):
In evaluated 300-URL OpenPhish feed snapshot (SHA-256: 53779a3c...), 1 URL triggered heuristic (0.33%, 1/300). Exact Feed URL: https://www.favebook.vercel.app (extracted token 'favebook' matches protected seed 'facebook' with single-character Levenshtein distance d = 1 on public-suffix host vercel.app). 94.7% of general phishing URLs operate on compromised sites without brand-lookalike domain names. This measures detection on that sample, not recall on a dedicated lookalike dataset.

Multi-Character Substitutions (d >= 2) & Stem Rules:
Direct Levenshtein distance rules evaluate d = 1. Multi-character targets (e.g. amazon.com, d = 2) rely on digraph normalization (rn -> m) or token-stem rules (*book) on public-suffix hosts.

Pre-Connection Network Layer Verification:
Stage 1 check precedes explicit WebView load call. Absence of other pre-warning network traffic (such as Chromium speculative DNS prefetching or UDP/QUIC) remains unverified pending future PCAP network layer instrumentation.

Planned Engineering Roadmap (Q1–Q2 2027):
Implement offline HSTS preloaded list to block overrides on key financial domains, integrate PCAP traces, and engage independent lab for formal OWASP MASVS audit.

Formal Audit Package & Publisher Identity Reference:

- Legal Publisher Entity: Raise IT Solutions LLC (Corporate Affiliate: RAISE IT Solutions Ltd.) | Package: docs/security_audit_package/.
- Verification Guide: See README_AUDIT_PACKAGE.md and AUDIT_MANIFEST.json for full SHA-256 reproducibility.