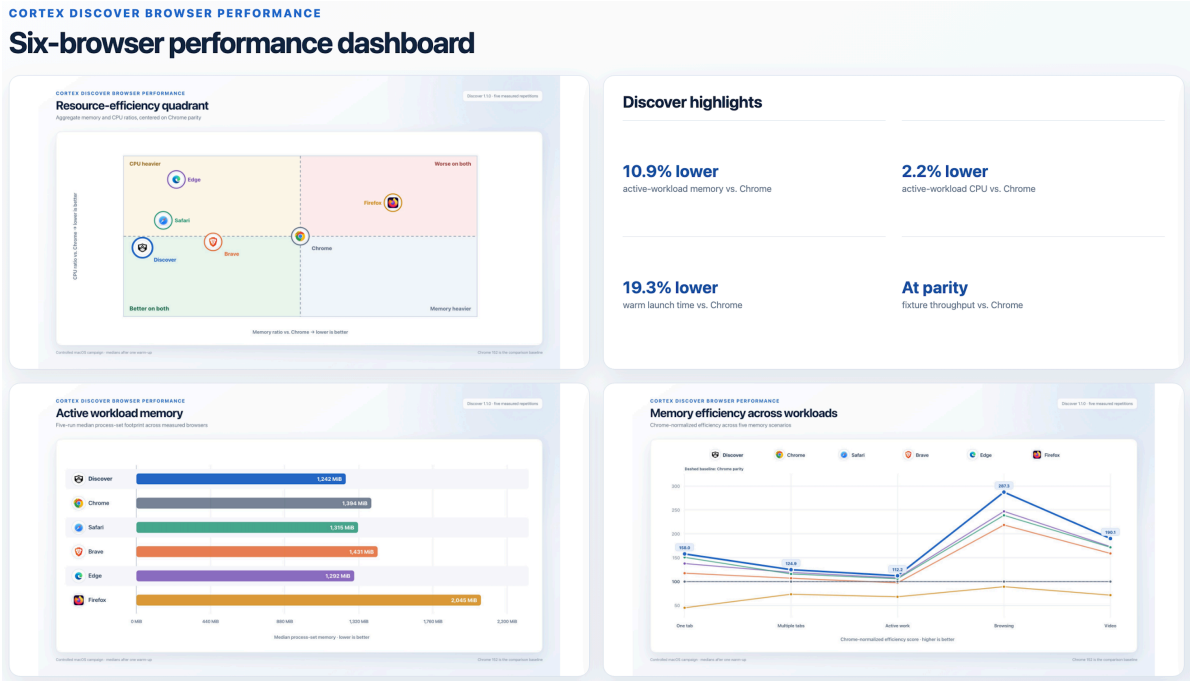
Cortex Discover: Performance & Resource Efficiency

White Paper

A six-browser engineering study of startup, memory, CPU, navigation, rendering, responsiveness, background tabs, and video performance



Pervaziv AI

Copyright © 2026 Pervaziv AI. All Rights Reserved.

Table of Contents

Section	Page
Disclaimer	3
Abstract	4
Executive findings	4
Why performance matters more in an agentic browser	6
Evaluation framework	6
What we tested	7
Measurement framework	8
The result at a glance	9
Reading the resource-efficiency quadrant	10
Memory: Discover's most consistent advantage	12
CPU: low at rest, competitive under load	14
Startup and navigation: getting to work sooner	15
Throughput and rendering: efficiency without reducing output	16
Video: lower memory, Chrome-level CPU, no dropped frames	16
The clearest remaining opportunity: interaction latency	17
A focused optimization moved the weakest metrics	19
What the results mean for Cortex Discover	19
What we are claiming, and what we are not	20
Guidance for enterprise evaluation	21
Interpreting browser performance in practice	21
A resource-efficient foundation for agentic work	24
References and document provenance	25

Copyright © 2026 Pervaziv AI. All Rights Reserved.

Disclaimer

The analyses, opinions, and rankings presented in this whitepaper are the result of independent research conducted by Pervaziv AI. All information is based on publicly available sources, internal assessments, and our own methodologies. The placement of companies in charts, including quadrant or ranking diagrams, reflects our internal evaluation criteria and opinions. It does not constitute an endorsement, certification, or affiliation with any company, brand, or product mentioned herein.

While every effort has been made to ensure the accuracy and reliability of the information presented, market conditions, product features, and company offerings may change over time. The authors make no guarantees regarding the completeness or current accuracy of the information. Readers are encouraged to conduct their own research and verify any information before making business, investment, or purchasing decisions.

The opinions expressed in this whitepaper are solely those of Pervaziv AI and should be treated as analytical insights rather than statements of fact. Any references to other companies, products, or services are for comparative and illustrative purposes only and are not intended to disparage or misrepresent any competitor.

By using or relying on the information in this whitepaper, readers acknowledge and accept that Pervaziv AI is not responsible for any decisions made based on this material, and all content is provided “as is” without warranty of any kind.

Copyright © 2026 Pervaziv AI. All Rights Reserved.

Abstract

When we launched Cortex Discover, we described a dedicated enterprise agentic browser: a familiar browsing environment with native page awareness, selected-tab reasoning, controlled browser actions, visual context, file-aware workflows, and durable task continuity.

That product direction creates an important engineering obligation. An AI browser should not make people choose between intelligence and a responsive computer. It has to support research across tabs, long-running workflows, visual material, files, and enterprise applications while preserving the resources those applications need.

We therefore ran a structured performance campaign comparing Cortex Discover 1.1.0 with five widely used browsers: Google Chrome, Safari, Brave, Microsoft Edge, and Firefox. We measured each browser repeatedly on the same machine with the same deterministic workload. The study covered launch time, idle and active memory, CPU use, navigation, throughput, rendering, interaction latency, background-tab readiness, video playback, and recovery.

The result is not a claim that one browser wins every benchmark. Browsers make different tradeoffs, and some measurements are more portable across browser engines than others. It is a detailed engineering snapshot of where Discover is already strong and where we still have work to do.

Document profile	Description
Publication	Public technical whitepaper, October 2026
Product evaluated	Cortex Discover 1.1.0
Study	Six-browser full performance campaign
Study identifier	`six-browser-full-5x-20261003-r3`
Evidence base	36 completed runs and 1,447 browser/phase/metric distributions
Primary baseline	Google Chrome for Testing 152.0.7951.0
Intended audience	Enterprise technology leaders, evaluators, performance engineers, and customers

Executive findings

The clearest result is resource efficiency. In this campaign, Discover ranked first of six in both the report's memory-efficiency and CPU-efficiency categories. It used 10.9% less memory than Chrome in the active workload, 36.7% less with one idle tab, and 19.9% less with multiple idle tabs. It also launched 19.3% faster than Chrome in the warm-start scenario and completed the navigation phase in 17.6% less wall time. Rendering and fixture throughput remained effectively at parity.

Those savings did not erase every tradeoff. Discover's handler-to-frame p95 was 10.5% slower than Chrome, and background-tab readiness p95 was 29.0% slower in the full campaign. Those are visible optimization targets, not numbers to hide. A separate, isolated A/B validation also showed that a focused optimization

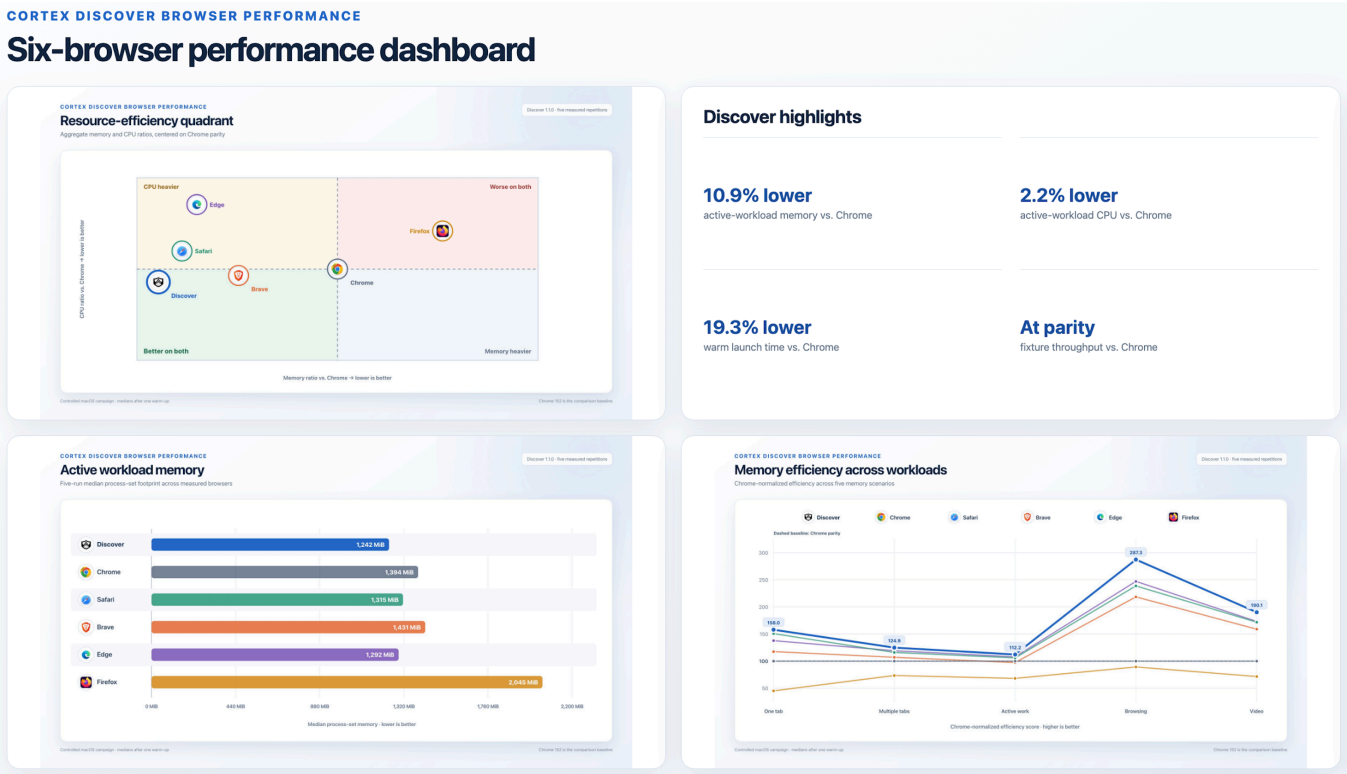
Copyright © 2026 Pervaziv AI. All Rights Reserved.

reduced background-tab readiness p95 by 18.9% and long-task time by 12.7% without a stable memory regression.

The principal findings are:

- **Resource headroom:** Discover recorded the lowest active-workload memory in the six-browser field and ranked first in the report's memory- and CPU-efficiency categories.
- **Output preservation:** Fixture throughput and rendering remained effectively equal to Chrome while Discover used less memory and, in several phases, less CPU.
- **Media stability:** Discover completed every measured synthetic-video run without a dropped frame, with steady-state CPU close to Chrome and a lower captured memory footprint.
- **Specific remaining work:** Interaction-tail latency and background-tab readiness were the clearest weaknesses; they are treated as optimization targets rather than folded into a single overall score.
- **Bounded interpretation:** The findings describe one controlled campaign on one host. They establish a reproducible engineering baseline, not a universal ranking or guarantee for every site and device.

This is how we measured the browser, what the charts mean, and why resource efficiency matters for an agentic product.



Performance dashboard summary. The resource-efficiency quadrant leads the view, followed by headline results and two supporting memory comparisons.

Copyright © 2026 Pervaziv AI. All Rights Reserved.

Pervaziv AI

QUESTIONS? CONTACT US.

pervaziv.com

contact@pervaziv.com

This public whitepaper and its contents are © 2026 Pervaziv AI.
Performance results are subject to the scope, measurement definitions, and limitations stated in this document.
Copyright © 2026 Pervaziv AI. All Rights Reserved.