


Pervaziv AI Advances Cortex Discover with Resource Efficient Agentic AI, Backed by Six Browser Study

Six browser study ranks Discover first in memory and CPU efficiency, with faster startup and near identical throughput and rendering versus Chrome.

SAN FRANCISCO, CA, UNITED STATES, October 6, 2026 /EINPresswire.com/ -- [Pervaziv AI](#) today announced new [performance results](#) for [Cortex Discover](#) 1.1.0, its dedicated agentic AI browser, showing that intelligence in the browser can grow without forcing users to give up the memory, CPU capacity and responsiveness needed for the rest of their work.



In a structured study, Cortex Discover was evaluated alongside Google Chrome, Safari, Brave, Microsoft Edge and Firefox across startup, memory use, CPU use, navigation, throughput, rendering, responsiveness, background activity and synthetic video. Every browser completed the same repeatable workload under the same test conditions.

“

Agentic AI should expand what people can do, not consume the resources they need to do it. Discover brings intelligence and efficiency together so AI stays close to the work.”
Anoop Jaishankar

The clearest result was resource efficiency. Discover ranked first of six in both the memory efficiency and CPU efficiency categories. Compared with Chrome, Discover used 36.7 percent less memory with one idle tab, 19.9 percent less with multiple idle tabs and 10.9 percent less during active work. Warm startup was 19.3 percent faster, while aggregate navigation completed in 17.6 percent less

wall time.

At the same time, Discover maintained near identical throughput and rendering versus Chrome,

operating at roughly 60 operations per second and 60 frames per second in the measured workload. During synthetic video, Discover recorded zero dropped frames and used 47.4 percent less captured memory.

For Pervaziv AI, the result addresses a larger product question. Cortex Discover was built to bring AI directly into the environment where research, communication, analysis, development, approvals, enterprise applications and everyday knowledge work already happen. That makes the browser more than a place to display pages. It becomes an active workspace where AI can understand live context, reason across selected tabs, work with files, help move tasks forward and maintain continuity.

That intelligence is only useful if the browser still leaves room for the work itself.

More Intelligence. More Room for Work

An agentic AI browser should not make users choose between a more capable assistant and a responsive computer.

Cortex Discover combines familiar browsing with native page awareness, selected tab reasoning, visual context, controlled browser actions, file aware workflows and durable task continuity. A user can research across several sources, compare information, work with documents, move through web applications and continue a task with Cortex without constantly copying context between the browser and a separate AI experience.

That model creates a different performance requirement from a conventional browser.

A typical work session may include dashboards, email, documentation, cloud consoles, development tools, video, collaboration applications and numerous research tabs. Those applications compete for the same memory and CPU resources as the browser. An AI layer that continuously expands the browser footprint can reduce the very productivity it is intended to improve.

Discover is being engineered around the opposite idea: add intelligence while preserving room for everything around it.

The new study provides an early measure of that approach.

With one idle tab, Discover used a median 439.5 MiB of captured memory compared with 694.3 MiB for Chrome. With multiple idle tabs, Discover measured 1,173.6 MiB compared with 1,465.8 MiB for Chrome. During the active workload, Discover measured 1,242.4 MiB compared with 1,394.4 MiB.

The active workload result is especially important because lower memory means little if a

browser simply does less. In the same campaign, Discover and Chrome both delivered approximately 60 operations per second. Rendering was also effectively equal at approximately 60 frames per second.

That combination is the more meaningful result: lower resource demand while preserving the measured output users expect from a modern browser.

“Agentic AI should expand what people can do in the browser, not consume the resources they need to do it,” said Anoop Jaishankar, Founder and CEO of Pervaziv AI. “With Discover, we are building intelligence and efficiency together so AI can stay close to the work without getting in the way of it.”

Faster Into the Work

Performance is also about how quickly a browser gets out of the user’s way.

In the study, Discover’s warm launch to page ready completed in 915.2 milliseconds compared with 1,134.2 milliseconds for Chrome, a 19.3 percent improvement. Native launch to page ready was 20.2 percent faster than Chrome in the measured scenario.

Aggregate navigation wall time measured 68.2 milliseconds for Discover and 82.8 milliseconds for Chrome, a 17.6 percent reduction.

These results do not mean Discover was faster on every individual page metric. Page level paint and navigation measurements moved in both directions across the fixtures. Pervaziv AI is treating the aggregate results as evidence of a strong performance foundation, not as a claim that one browser wins every benchmark.

That distinction matters to the company’s broader approach. Performance work is most useful when it identifies both the areas that are already strong and the areas that still need engineering attention.

Measured Against Leading Browsers

To put the results in practical market context, Cortex Discover 1.1.0 was tested alongside five widely used browsers: Google Chrome, Safari, Brave, Microsoft Edge and Firefox.

The objective was not to produce a universal browser ranking. It was to determine whether an agentic AI browser can deliver the responsiveness users expect from an established browser while also creating headroom for AI assisted work.

Across the campaign, all 36 measured runs completed successfully. The workload covered multi tab activity, navigation, active browser work, rendering, background behavior and synthetic

video. Chrome served as the primary performance reference, while the broader six browser field provided perspective on how Discover compares with familiar browser experiences.

Discover's strongest result was consistent resource efficiency. It recorded the lowest active workload memory in the six browser field and ranked first in the study's overall memory efficiency category. It also ranked first in CPU efficiency across the study's evaluated scenarios.

At one tab idle, Discover used 2.4 percent of one CPU core compared with 4.6 percent for Chrome. During the active workload, Discover measured 34.0 percent compared with Chrome at 34.8 percent. During browsing, Discover measured 32.7 percent compared with 38.2 percent for Chrome.

The numbers matter less as isolated victories than as evidence of a product direction. Cortex Discover is designed to add an AI workspace beside ordinary browsing without turning intelligence into constant background overhead.

Video Without Giving Up the Machine

Mixed media is an important part of modern browser work. Meetings, training, product demonstrations, research, communication and media often run alongside other tabs and desktop applications.

In the synthetic video phase, Discover used 623.8 MiB of captured memory compared with 1,186.1 MiB for Chrome, a 47.4 percent reduction in that measured scenario.

Steady video CPU remained close, at 33.3 percent of one core for Discover and 33.9 percent for Chrome. Both browsers recorded zero dropped frames across all measured repetitions.

The synthetic video workload is not intended to represent every conferencing application, codec or streaming service. It is one controlled measure of whether the browser can preserve smooth media behavior while maintaining the broader resource profile seen elsewhere in the campaign.

For users, that is the practical question. AI assisted browsing increasingly happens alongside video calls, design tools, local development environments and business applications. Saving resources in isolation is less valuable than leaving those resources available for the rest of the work.

A Broader Cortex Performance Program

The Discover results are not an isolated optimization effort. They extend a performance program that has moved through the Cortex platform during 2026.

In April, Pervaziv AI published a performance and memory optimization campaign focused on frequently executed frontend and backend operations. Across five optimized operations, average response time fell from approximately 530 milliseconds to approximately 82 milliseconds, a roughly 6.5 times improvement. A storage related operation fell from 680 milliseconds to 42 milliseconds.

Cortex 4.0 later expanded that focus into coding and agentic work. Pervaziv AI reported up to 2.5 times faster performance across several measured coding tasks and approximately 1.5 times faster overall performance across the measured agentic workflows.

As Cortex became more capable, the performance problem also became broader. An enterprise AI request can involve context assembly, retrieval, model routing, search, skills, inference, tool use, validation and continued execution. Efficiency cannot be solved at only one point in that chain.

Pervaziv AI has therefore been working across several layers of the system.

In September, Cortex semantic context compaction reduced eligible supplemental context by 46 percent in the evaluated scenarios while preserving required fact recall, retrieval success, task success and validation success. The goal was not simply smaller prompts. It was to make longer running work more manageable while retaining the information needed for correctness.

The company then introduced its three tier Cortex Inference Cache Architecture, designed to avoid safely repeatable work across application context, prompt processing and narrowly approved exact responses.

In evaluated paths, the architecture demonstrated up to 150 times faster prompt prefill, approximately 11 times faster exact response delivery and up to 2.25 times throughput at moderate concurrency.

Those measurements come from different workloads and should not be combined into a single speed claim. Together, they illustrate the same architectural direction: reduce unnecessary work where it occurs, whether that work is in the product experience, accumulated context, inference processing or now the browser itself.

Cortex Discover brings that discipline to the desktop surface where users experience AI directly.

Efficiency as an AI Capability

The performance of an AI system is often described in terms of model speed. Pervaziv AI takes a broader view.

For an agentic platform, performance also includes how much context must be carried, how

often work is repeated, how quickly tasks recover, how much memory the user experience consumes and whether the system leaves enough computing headroom for the applications around it.

That is particularly important in the browser.

Cortex Discover can work with the current page, tabs a user selects, attachments, conversation history and task state. Through the broader Cortex platform, work can also connect with models, search, skills, enterprise systems and managed execution.

The browser therefore sits at the intersection of AI reasoning and ordinary computing.

If that surface becomes heavy, users feel it immediately. Tabs compete for memory. Applications become less responsive. Battery use can rise. Meetings, development tools and local applications have less capacity available.

Resource efficiency becomes part of the intelligence experience itself.

For enterprise teams, that headroom can matter beyond the browser itself. A developer may have an IDE, local services and build tools running beside Discover. A security analyst may be reviewing dashboards, evidence and live systems. A business user may be working across video, documents and cloud applications.

In each case, the browser is only one part of the computing environment. Pervaziv AI believes an AI browser should be judged not only by what it can do inside its own window, but by how well it coexists with everything the user still needs to run.

The new Discover results suggest that the AI layer can be designed with that constraint from the beginning rather than added after capability is built.

Publishing the Gaps, Not Just the Wins

The study also identified areas where Discover trails Chrome.

Handler to frame p95 measured 72.3 milliseconds for Discover compared with 65.5 milliseconds for Chrome, making Discover 10.5 percent slower in that measure. Background tab readiness p95 measured 79.1 milliseconds compared with 61.3 milliseconds for Chrome.

Pervaziv AI views those results as an engineering roadmap rather than numbers to remove from the story.

A separate same version validation around a focused optimization reduced background tab readiness p95 by 18.9 percent and long task time by 12.7 percent. Active workload CPU time and

responsiveness CPU time also improved slightly, while the test did not show a stable memory regression.

The optimized values were kept separate from the original six browser comparison. The purpose of the additional test was to determine whether the weakest measured area could improve without giving back the broader resource gains.

That approach reflects a principle that has guided the company's recent performance work: improving one metric should not silently create a larger cost somewhere else.

From Faster AI to More Usable AI

The broader opportunity is larger than browser benchmarking.

Enterprise AI is moving from isolated prompts toward longer tasks that involve research, files, applications, development environments, cloud systems, validation and human decisions. As capability expands, every additional layer introduces potential latency and resource cost.

Pervaziv AI believes the next stage of AI infrastructure must address that cost directly.

Cortex has evolved from a coding and security assistant into a broader Enterprise AI platform with specialized models and agents, routing across models, search and skills, enterprise connections, managed execution and an agentic browser.

Performance engineering is what helps those capabilities operate as one usable system rather than a collection of expensive layers.

The progression across Cortex reflects that idea. Earlier work reduced latency in frequent operations. Cortex 4.0 accelerated measured coding and agentic workflows. Context compaction reduced accumulated information pressure. Inference caching reduced repeat computation. Cortex Discover now targets the memory, CPU and responsiveness of the surface where AI assisted work is happening.

The goal is not speed for its own sake. It is making more intelligence practical.

A Foundation for What Comes Next

Pervaziv AI will continue improving interaction latency and background tab readiness while preserving the memory, CPU and responsiveness gains already measured. Future testing will expand across additional hardware, operating environments, longer workloads, real media scenarios and established browser benchmarks.

The current study is a controlled performance snapshot, not a claim that one browser wins every

workload. Within that scope, Discover ranked first in memory and CPU efficiency, recorded the lowest active workload memory in the six browser field, started faster than Chrome, maintained near identical throughput and rendering, and completed the synthetic video workload without dropping a frame.

A detailed Cortex Discover performance whitepaper will also be available for download at pervaziv.com, with expanded methodology, benchmark results, comparisons and analysis.

The results reinforce the vision behind Cortex Discover: as the browser becomes a primary surface for agentic AI, greater intelligence should create more capability without placing greater demands on the machine.

Organizations evaluating the next generation of AI enabled browsing can explore Cortex Discover at pervaziv.com and download the full performance whitepaper to review the methodology, benchmark data and results behind the study.

Cortex Discover 1.1.0 is available for Windows x64 and macOS as part of the broader Cortex platform spanning browser, development, mobile, cloud and connected enterprise workflows.

Pervaziv AI

Pervaziv AI

Visit us on social media:

[LinkedIn](#)

[Instagram](#)

[YouTube](#)

[X](#)

[Other](#)

This press release can be viewed online at: <https://www.einpresswire.com/article/947570631>

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.

© 1995-2026 Newsmatics Inc. All Right Reserved.