

# The Agentic Web: Rethinking Web Infrastructure for Machine Consumption

David Hurley  
Plasmate Labs

March 2026

## Abstract

The World Wide Web was designed as a hypertext system for human readers navigating between documents via visual browsers. Three decades later, AI agents are emerging as a second class of web consumer with fundamentally different requirements: they need meaning, not pixels; structure, not style; and efficiency, not aesthetics. We argue that the web is entering a fourth state—the agentic web—that requires new infrastructure primitives analogous to the transitions from static HTML to dynamic JavaScript and from synchronous HTTP to real-time WebSockets. We present three proposed infrastructure components: the Semantic Object Model (SOM), a token-efficient page representation; Agent Web Protocol (AWP), a communication protocol for agent-web interaction; and SOM directives for robots.txt, enabling cooperative content negotiation between publishers and agents. We analyze the economic implications of the current “Chrome tax” on AI agent infrastructure (\$711K per million pages at GPT-4 rates) and demonstrate how semantic web representations can reduce this by 91%. We discuss the standardization pathway through the W3C Web Content Browser for AI Agents Community Group and propose design principles for an agent-friendly web that serves both human and machine consumers.

## 1 Introduction

The web has undergone three fundamental state transitions [1, 2, 3]:

State	Medium	Technology	Era
1st (Static)	Documents	HTML	1993–2004
2nd (Dynamic)	Applications	JavaScript	2004–2015
3rd (Real-time)	Streams	WebSockets	2015–2025
<b>4th (Agentic)</b>	<b>Intelligence</b>	<b>SOM/AWP</b>	<b>2025–</b>

Each transition introduced new infrastructure: the first brought HTTP servers and HTML parsers; the second brought JavaScript engines and AJAX; the third brought WebSocket servers and event-driven architectures.

The fourth transition—from human-consumed web to machine-consumed web—is happening today without corresponding infrastructure. AI agents in 2026 interact with web pages using tools designed for human browser automation (Playwright [4], Puppeteer [5], Selenium [6]), running full rendering engines (Chromium) to produce visual output that is immediately discarded.

This paper argues that the agentic web requires purpose-built infrastructure, proposes three components of that infrastructure, and outlines a path to standardization.

## 2 The Problem: Chrome as Universal Agent Interface

### 2.1 How Agents Browse Today

Every major AI agent framework uses the same pipeline:

```
Agent -> Playwright/Puppeteer -> Chromium -> Render -> DOM -> Serialize -> LLM
```

This pipeline inherits Chromium’s resource requirements:

- 300–500MB memory per browser instance
- 2–5 seconds per page (rendering + JavaScript execution)
- 50,000+ tokens of DOM output per page
- GPU utilization for compositor operations the agent never sees

### 2.2 The Economic Impact

At GPT-4 pricing (\$10/1M input tokens), our 49-site benchmark shows the Chrome pipeline costs \$50,397 per 1,000 page loads of the test set. With SOM compression, the same workload costs \$3,042, a 94% reduction. At GPT-4o pricing (\$2.50/1M), the savings ratio remains constant: approximately 94% of input tokens consumed by the LLM encode presentational information that SOM eliminates.

Full cost analysis with per-URL breakdowns is published at <https://docs.plasmate.app/benchmark-cost> and reproduced nightly via <https://docs.plasmate.app/coverage>.

### 2.3 The Scaling Wall

As AI agents evolve from single-task assistants to autonomous systems performing thousands of web interactions per session, the Chrome pipeline becomes unsustainable:

- 100 concurrent agents: 30GB browser memory
- 1,000 concurrent agents: 300GB browser memory (requires distributed infrastructure)
- 10,000 concurrent agents: 3TB browser memory (requires a dedicated cluster)

This scaling wall is unique to agent-web interaction. Human users never needed 10,000 concurrent browser instances.

## 3 Proposed Infrastructure

### 3.1 Semantic Object Model (SOM)

SOM is a JSON-based format that represents web page content as typed semantic elements organized into regions. (Detailed specification in companion paper [13].)

Key properties:

- **16.6x mean token compression** vs raw HTML (49-site benchmark, 94% savings; nightly coverage across 100 sites at <https://docs.plasmate.app/coverage>)
- **Semantic typing** – elements classified by role, not HTML tag
- **Hierarchical structure** – regions contain elements in document order
- **Machine-readable** – JSON, parseable without a browser

### 3.2 Agent Web Protocol (AWP)

AWP is a proposed communication protocol for bidirectional agent-web interaction. While the Chrome DevTools Protocol (CDP) [14] was designed for human developers debugging websites, AWP is designed for agents interacting with web content.

Key differences from CDP:

Aspect	CDP	AWP
Primary consumer	Human developer	AI agent
Output format	DOM tree	SOM
Interaction model	Pixel coordinates	Semantic element IDs
Session scope	Visual debugging	Content comprehension
State management	Tab/page based	Session/context based

### 3.3 SOM Directives for robots.txt

Current robots.txt [15] provides binary access control (allow/disallow). We propose SOM directives that enable content negotiation:

```
SOM-Endpoint: https://cache.example.com/v1/som
SOM-Format: SOM/1.0
SOM-Scope: main-content
SOM-Freshness: 3600
```

This allows publishers to:

- Offer semantic representations they control (excluding ads, paywalls, tracking)
- Reduce server load from redundant rendering requests
- Signal preferred agent interaction methods

## 4 Design Principles for the Agentic Web

### 4.1 Dual-Audience Design

Web infrastructure should serve both humans and machines without requiring separate systems. SOM does not replace HTML; it provides an alternative view of the same content.

### 4.2 Open Standards

Agent web infrastructure must be standardized openly. Proprietary formats create lock-in that fragments the agent ecosystem. The SOM and AWP specifications are developed through the W3C Community Group process [17].

### 4.3 Progressive Enhancement

Agents should degrade gracefully when SOM is not available, falling back to HTML parsing. SOM directives in robots.txt are hints, not requirements.

## 4.4 Publisher Sovereignty

Content publishers retain control over how their content is consumed. SOM endpoints can filter, rate-limit, and authenticate access independently of HTML serving.

## 4.5 Efficiency as a Feature

The agentic web should be measurably more efficient than the current approach. Our benchmark demonstrates 16.6x token compression, 50x speed improvement, and 10x memory reduction – not marginal improvements, but categorical differences.

# 5 Standardization Pathway

The W3C Web Content Browser for AI Agents Community Group [17] was established in March 2026 to incubate these specifications. The group follows the W3C Community Group process:

1. **Incubation** (current) – develop specifications within the community group
2. **Implementation** – reference implementations in Plasmate (Rust, Apache 2.0)
3. **Adoption** – integration with major agent frameworks (LangChain [20], LlamaIndex [21], CrewAI [22], AutoGen [23])
4. **Formalization** – if adoption warrants, propose as W3C Working Group deliverables

Complementary proposals have been submitted to Schema.org [16] (WebPageSemanticRepresentation type) and as an extension to RFC 9309 [15] (Robots Exclusion Protocol).

# 6 Related and Concurrent Work

Several projects address aspects of the agent-web interface:

- **Lightpanda** [10] – Alternative headless browser (Zig, AGPL) that still outputs DOM
- **Firecrawl** [11] – Web content to Markdown conversion API
- **Jina Reader** [12] – URL-to-text extraction service
- **Browser Use** [7] – LLM-guided browser automation layer
- **Stagehand** [9] – AI-powered browser interaction

These projects optimize within the existing Chrome-centric paradigm. SOM and AWP propose a new paradigm in which rendering is not part of the agent’s critical path.

## 7 Conclusion

The web is entering its fourth state. The infrastructure that served static documents, dynamic applications, and real-time streams is insufficient for agentic consumption. Purpose-built primitives – semantic page representations, agent communication protocols, and cooperative content negotiation – are needed to make the agentic web efficient, open, and sustainable.

The Semantic Object Model, Agent Web Protocol, and SOM directives represent a first proposal for this infrastructure. They are open, implemented, and available for adoption. We invite the research community, browser vendors, AI companies, and web publishers to participate in their development through the W3C Community Group.

## Acknowledgments

This work builds on the HTML5 parsing infrastructure (html5ever [18]) developed by the Servo project and the V8 JavaScript engine [19] developed by Google. The W3C Community Group process provides the standardization framework.

## References

- [1] T. Berners-Lee, R. Fielding, and H. Frystyk. Hypertext Transfer Protocol – HTTP/1.0. RFC 1945, 1996.
- [2] J. J. Garrett. Ajax: A New Approach to Web Applications. Adaptive Path, 2005.
- [3] I. Fette and A. Melnikov. The WebSocket Protocol. RFC 6455, 2011.
- [4] Playwright. <https://playwright.dev/>
- [5] Puppeteer. <https://pptr.dev/>
- [6] Selenium. <https://www.selenium.dev/>
- [7] Browser Use. <https://github.com/browser-use/browser-use>
- [8] LaVague. <https://github.com/lavague-ai/LaVague>
- [9] Stagehand. <https://github.com/browserbase/stagehand>
- [10] Lightpanda. <https://lightpanda.io/>
- [11] Firecrawl. <https://firecrawl.dev/>
- [12] Jina Reader. <https://jina.ai/reader/>
- [13] D. Hurley. The Semantic Object Model: A Token-Efficient Web Representation for AI Agents. Plasmate Labs, 2026.
- [14] Chrome DevTools Protocol. <https://chromedevtools.github.io/devtools-protocol/>
- [15] M. Koster et al. Robots Exclusion Protocol. RFC 9309, 2022.
- [16] Schema.org. <https://schema.org/>

- [17] W3C Web Content Browser for AI Agents Community Group. <https://www.w3.org/community/web-content-browser-ai/>
- [18] html5ever. <https://github.com/nickel-org/html5ever>
- [19] V8 JavaScript Engine. <https://v8.dev/>
- [20] LangChain. <https://github.com/langchain-ai/langchain>
- [21] LlamaIndex. [https://github.com/run-llama/llama\\_index](https://github.com/run-llama/llama_index)
- [22] CrewAI. <https://github.com/joaomdmoura/crewAI>
- [23] AutoGen. <https://github.com/microsoft/autogen>