TYGART MEDIA
Fleet Technology Standard • Autonomous Agent Readiness

Current version — Nov 1, 2026
Update history: tygartmedia.com/the-agent-ready-site...

AGENT-READY QUICK CARD: THE 4-LAYER SPECIFICATION

THE 62% BOT SHIFT: WHY BUSINESS WEBSITES MUST BE BUILT FOR AGENTS

On October 23, 2026, Kin Insurance became the first U.S. home insurer to enable WebMCP in its quoting experience — an open web standard allowing AI agents to transact directly with kin.com. Strip away insurance specifics, and every commercial website faces the same reality: **roughly 62% of website visits across our managed fleet are bots and agents, not human beings**. The customer is no longer browsing alone; they are dispatching autonomous agents to evaluate, compare, and transact.

LAYER 1: DECIDE TO WELCOME THE MACHINE (DOOR VS. WALL POSTURE)

Agent traffic is not an impending trend; it is the established majority. A business cannot remain neutral: having no explicit posture means default hosting and WAF rules decide on your behalf.

POSTURE	INDUSTRY BENCHMARK	OPERATIONAL MECHANISM	RISK / CONSEQUENCE
THE DOOR (Welcome)	**Kin Insurance** (Announced AI shopping agents welcome; named Meta's Muse)	Explicitly permits AI shopping agents; exposes machine-readable paths; treats machine visits as commercial opportunities.	Requires structured facts and declared transaction paths so agents do not hallucinate rates or terms.
THE WALL (Block)	**Insurify** (September 2026: Blocked Meta's Muse from comparison platform)	WAF blocks, IP bans, user-agent blacklists; argues agents strip context essential for valid price/service comparison.	Complete invisibility to AI shopping interfaces; customer agents bypass site in favor of accessible competitors.
DEFAULT SETTINGS (Unconscious Wall)	Unconfigured business websites	Default firewall, CDN, or bot-management settings inadvertently throw Captchas or 403s at benign consumer agents.	Loss of 62% addressable traffic with zero executive awareness; unmonitored drop in citation and discovery.

LAYER 2: STRUCTURE CONTENT SO A MACHINE CAN CARRY IT ACCURATELY

An agent can only repeat what it can read. The agent-ready site is constructed of direct answers, not marketing brochures. Agents cite structured content and skip implied claims.

1. One Real Question Per Page

Discard broad corporate brochures. Focus each page on answering a single, specific customer inquiry with exhaustive clarity. Pages with split or muddy intent confuse agent parsing and dilute authority scores.

2. Standalone Paragraph Test

Every paragraph must stand alone as a complete, coherent idea. If a paragraph requires preceding prose or subsequent marketing copy to make sense, an agent extracting a snippet will misrepresent or omit your facts.

3. Verbatim FAQ Structured Data

Embed Schema.org/JSON-LD machine-readable data directly underneath the copy. Content in JSON-LD must match on-page text verbatim. Fleet benchmark: **155 of 166 articles structured** with zero discrepancy.

4. Zero Media-Lockout Standard

Never trap critical commercial facts (pricing, specifications, service boundaries, warranties) inside images, video transcripts, or sales rep conversations. If it is not in clean DOM text, it does not exist to the agent.

Current version — Nov 1, 2026
Update history: tygartmedia.com/the-agent-ready-site...

AGENT-READY QUICK CARD: ACTION & MEASUREMENT

LAYER 3: GIVE THE AGENT A WAY TO ACT, NOT JUST TO READ (WEBMCP)

WebMCP is the open web standard for machine action. The site explicitly declares the actions an agent may execute — quote, book, or order — with exact inputs in strict sequence, preventing brittle screen-scraping.

The Kin WebMCP Precedent (Oct 23, 2026)

Kin enabled WebMCP on kin.com's home insurance quoting engine. In real testing, an AI shopping agent completed a binding quote after asking the homeowner *only a handful of questions*. The declared path allowed the carrier's automated data integrations to fill property and risk attributes without human friction.

The Two-Rail Commerce Model

Modern business sites require two parallel conversion funnels:
• **Rail A (Human Checkout):** Visual hierarchy, testimonials, interactive inputs, multi-page human UX.
• **Rail B (Agent Rail):** Declared machine-usable endpoint, typed parameters, deterministic state machine, instant verification.

LAYER 4: MEASURE WHAT THE MACHINES DO (CITATION & ACTION SCOREBOARDS)

Agent-readiness without measurement is a vibe. The protocol makes sessions and transactions countable. While transactional counting (Layer 3) matures, every business must immediately measure the read layer.

62%	1,459,726	922
FLEET BOT & AGENT VISITS	AI CITATIONS (OCT 15)	CITED PAGES (OCT 15)
Analytics traffic baseline	Up from 269,399 (Sept 27)	Up from 324 pages (Sept 27)

SCOREBOARD METRIC	TYGART MEDIA MEASUREMENT MODEL	OPERATIONAL ACTION TRIGGER
Read Layer: Citations	Bing AI Performance search engine exports (not vendor estimates). Tracking absolute citations across published catalog.	Citations flatlining indicates content fails standalone paragraph test or lacks verbatim JSON-LD structured data.
Read Layer: Breadth	Total distinct pages cited (fleet expanded from 324 to 922 pages in 18 days between Sept 27 and Oct 15).	Low page ratio indicates citations concentrated on few legacy posts; rollout structured FAQ to underperforming pages.
Action Layer: Transactions	Countable WebMCP session tokens, completed agent quote flows, and programmatic booking receipts.	Drop-off in agent quote completions indicates declared parameter schema is broken or requires excessive human clarification.

WHERE OUR FLEET HONESTLY SITS: THE 4-LAYER ROADMAP

LAYER	CORE CAPABILITY	CURRENT STATUS	ROADMAP OBJECTIVE
1. Welcome	Open door policy; explicit unblocking of shopping agents	**RUNNING TODAY**	Audit CDN/firewall default settings across all 12+ fleet sites.
2. Structure	Answers-not-brochures; standalone paragraphs; verbatim FAQ	**RUNNING TODAY**	Maintain 155-of-166 baseline; enforce zero media-only information.
3. Action	WebMCP declared endpoints; get-a-quote/order input path	**NEXT BUILD**	Design two-rail checkout; pilot declared inputs on primary order flows.
4. Measure	Search engine AI export tracking; citation scoreboards	**RUNNING TODAY**	Bi-weekly Bing AI exports; track citations/page velocity against 1.45M bar.