chiefmartec

The Future of Websites

A Brand Context Model to Power Human & Agent Web Experiences

BY

Scott Brinker, chiefmartec

chiefmartec

The Future of Websites: A Brand Context Model to Power Human & Agent Web Experiences

BY SCOTT BRINKER, CHIEFMARTEC
IN PARTNERSHIP WITH KNOTCH
2026

Contents

§ 01 **The Web's Next Chapter**

The Web Is Dead. Long Live the Web.

§ 02 **The Intent-Responsive Surface for Humans**

A Conversation with the Customer
Experience Scaffolding
High-Fidelity Experience Analytics

§ 03 **The Semantic/Service Surface for Agents**

Readable Content, Structured Data, and Callable Services

§ 04 **A Unified Brand Context Model**

Maintaining Model Integrity
Not a Headless CMS, Not an LLM Either
Achieving Golden Context

§ 05 **A Synchronized Brand, Everywhere**

One Brand Context Model, Every Touchpoint

§ 01

The Web's Next Chapter

It's time to turn the page.

For over 30 years, pages have been the building blocks of the web. A metaphor from the days when printed brochures and documentation were how companies shared information with their customers.

Early commercial websites were called "brochureware." Like printed brochures, they were built for one-way communication. And ironically, they too were expensive to produce and slow to update.

Much has changed since those early days of the web. But much has not.

The structure of most brand web experiences — and the tech stack that powers them — is still anchored in the anachronistic concept of pages. Page templates. Page navigation. Pageviews. Even as we separated content from presentation with headless CMS, the design of a visitor's experience remained page-oriented.

But browsing and scrolling through page upon page isn't an experience that people love. They increasingly prefer the effortless flow of AI assistants that bring answers to them.

And while AI agents and crawlers still fetch and cite URLs — looking for clear, fresh knowledge and services — the "packaging" of traditional web pages is mostly irrelevant overhead to them. Their ideal experience is optimized for machine-readable efficiency.

PAGE DEBT

Marketing's flavor of tech debt — it results in brand latency.

Meanwhile, page-oriented CMS and DXP architectures remain cumbersome to change and maintain. As a brand's universe of campaigns, journeys, and audience segments grows, so does the complexity of keeping all those variations of pages and content in sync. This "page debt" — marketing's flavor of tech debt, where changes in the business take too long to be fully reflected across its digital surfaces — results in **brand latency**. What the brand currently intends customers to see and experience and what those customers actually see and experience across the brand's myriad of touchpoints diverges.

§ 01.1

The Web Is Dead. Long Live the Web.

None of this means that brand-owned web experiences matter less in the AI era. In fact, they matter more — for humans and agents.

For humans, it's about delivering what the AI assistants can't: an authoritative brand experience. Because for most purchases and customer relationships, brand matters. Trust, understanding, and emotional resonance matter. How it feels to interact with a brand matters. It shapes opinion, affinity, loyalty.

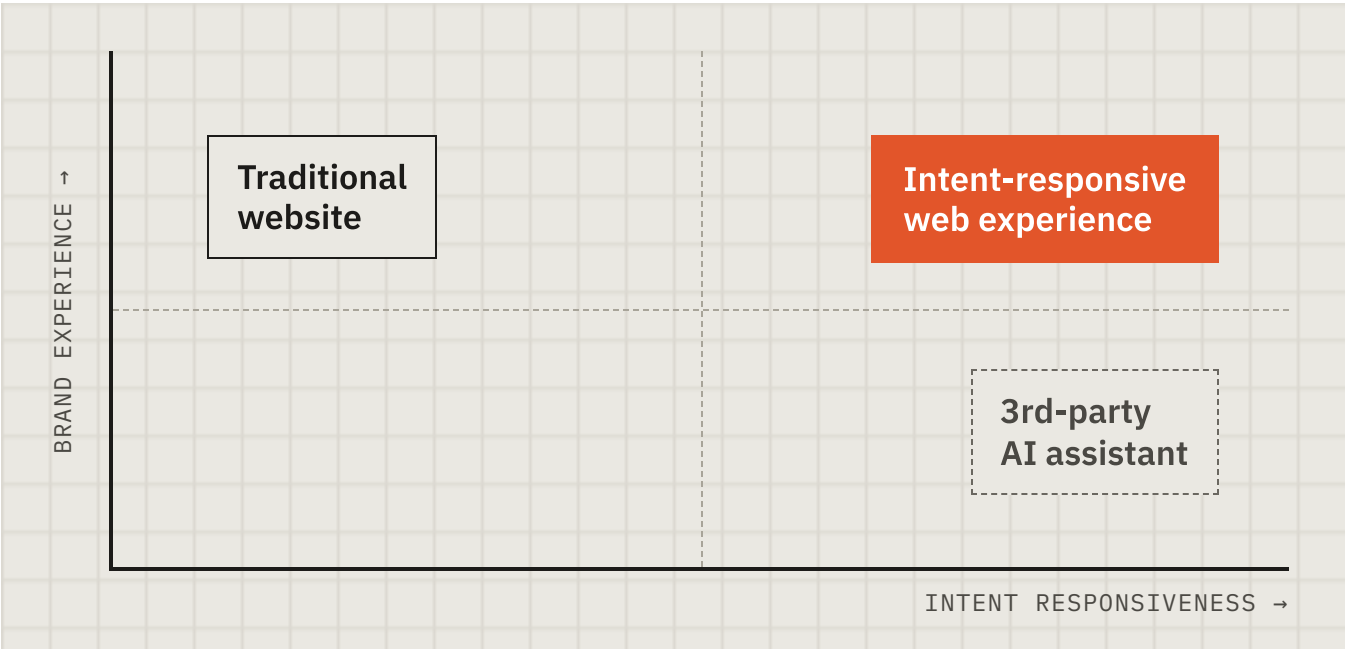
AI assistants compress brands, products, arguments, and evidence into relatively uniform responses, often accompanied by a disclaimer about their accuracy. Those answers are useful, and they play a role in a customer's journey. But they're low-fidelity representations of the brands themselves.

KEY POINT

They deliver the experience of the AI, not the experience of the brand.

FIG. 01

Brand experience vs. intent responsiveness



Brand-owned experiences have a much larger canvas. They can bring authoritative evidence and expertise to life through storytelling, rich visual explanations, and interactive features. They give customers ways to explore, verify, and act — to get to know you. In an AI sea of sameness, they give brands room to differentiate.

With AI assistants and agents in the loop, you may get fewer human visits to your website. But the humans who do arrive bring clearer intent and higher expectations. Each visit counts for more.

For agents, it's about competing for business within the AI assistants your customers use. Your brand needs to earn consideration, be represented accurately, and be easy to do business with through those interfaces.

Your web presence must provide authoritative source material for that representation. What makes your products relevant? What distinguishes them from alternatives? What evidence supports your claims? Under what conditions do your services and offers apply? Making that knowledge easy for agents to discover, interpret, and verify gives them crucial material for evaluating your brand and explaining its value.

Better source content helps more of what matters survive the compression.

Agent actions will carry that opportunity further. An agent that recommends your brand should have a clear path to the next step — check availability, book an appointment, make a purchase, request service — through agent-accessible services your web presence exposes.

Even when a customer doesn't leave their AI assistant, your agent-facing web presence is still able to shape awareness of your brand and their relationship with it.

These two roles come together in the **dual-surface website**:

FOR HUMANS

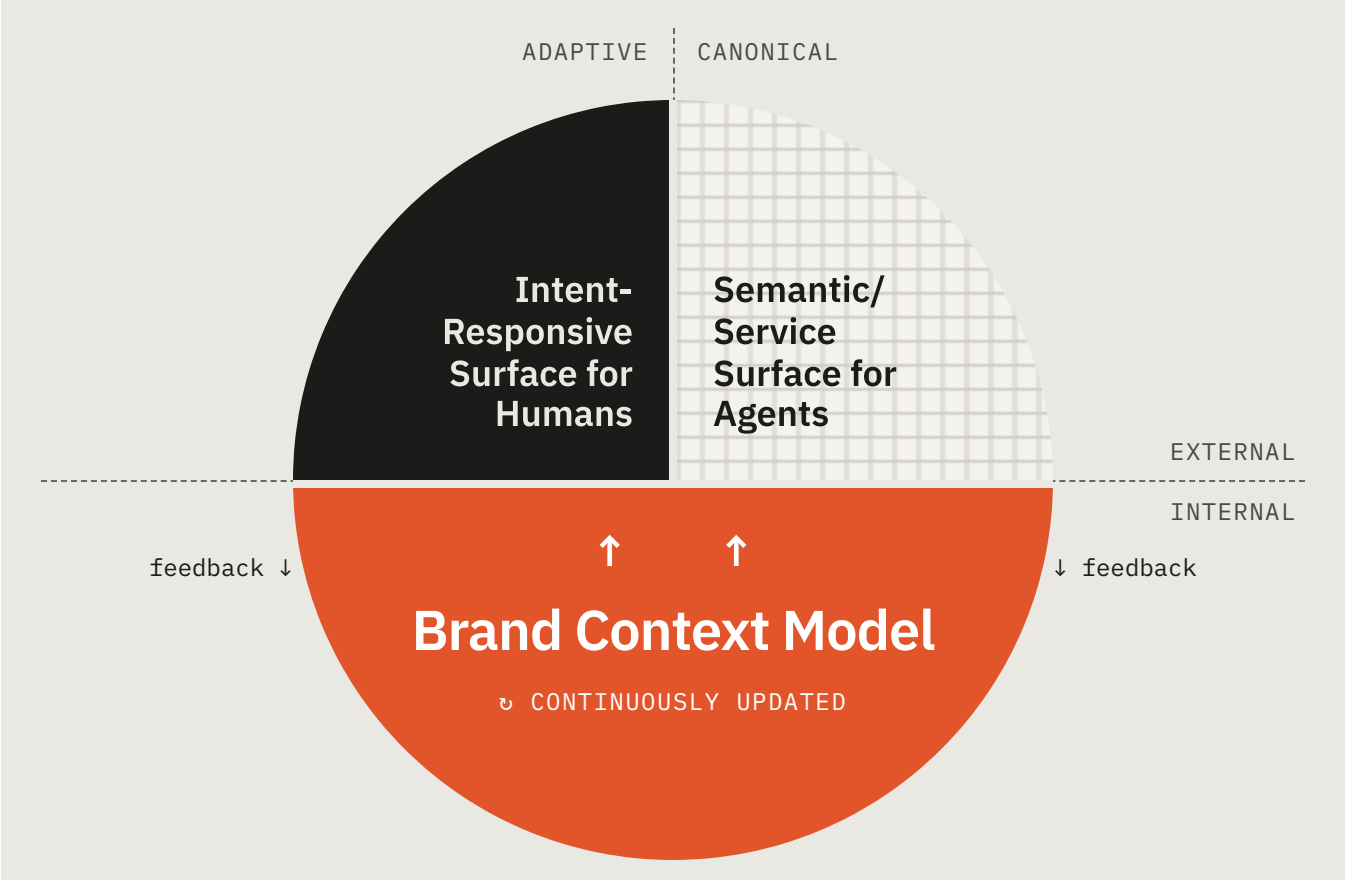
An **intent-responsive surface** for humans, where content, interactions, and services adapt to intent as visitors express their questions, preferences, and goals.

FOR AGENTS

A **semantic/service surface** for agents, through which agents access the brand's knowledge and permissioned capabilities on customers' behalf.

FIG. 02

The dual-surface website



But how do you keep these dual surfaces from drifting apart? How do you keep them in sync as the business changes at an ever-faster cadence? Orphaned and outdated content has plagued web experiences for three decades. Does the complexity of two specialized interfaces — especially a more intent-responsive one for humans — exceed our ability to maintain them?

We believe this requires a different architectural pattern than classic "digital age" CMS content management. An AI-native approach builds around an underlying **brand context model** — a canonical, dynamically maintained representation of the brand's knowledge and experiential assets — that both your human and agent web surfaces draw upon. Rather than a rigid hierarchy of database records, this model is a graph of atomized and relationally-connected information, services, and experience components that can be fluidly combined on-the-fly in a more adaptive fashion.

Instead of one definitive artifact called “the website,” there can be an underlying system capable of emitting many coherent, governed experiences across owned websites, AI assistants, agent interfaces, partner environments, and other digital surfaces.

§ 02

The Intent-Responsive Surface for Humans

For decades, visitors have had to adapt to the structure of our website. We’ve called that hierarchy of pages and navigation choices “information architecture.” We spend weeks or months building what we believe is the single most logical way to organize that experience.

But it’s always been distilled into one opinion, the brand’s opinion. If a visitor has a different opinion, well, too bad. It’s on them to map their intent to our architecture.

Landing pages are even more draconian. They narrow the experience to what we decide is relevant for that campaign — and no more. But relevance, like beauty, is in the eye of the beholder. If a visitor has other questions, or a different perspective motivating their arrival, they often have to jailbreak the page to hunt for those answers elsewhere in the Byzantine caverns of our primary website. It’s not quite the ideal customer journey that we mapped on a whiteboard.

A big reason people love AI assistants such as Claude and ChatGPT is that they give them what they want. Ask a question, get an answer. Ask a follow-up, get more details. It’s a dialogue that puts the customer in control, replacing a disjointed hunt-and-peck journey with an intelligent conversation.

It’s been said that the last best experience that anyone has anywhere becomes the expectation for the experience they want everywhere. AI assistants have raised those expectations. A website or landing page that still puts the onus on the visitor to adapt their intent to the marketer’s intention is more dissatisfying than ever.

§ 02.1

A Conversation with the Customer

We need our web experiences to adapt to the customer.

Not “adapt” in the limited ways we’ve deployed personalization to date, swapping out a few content blocks or generating a set of machine-learning recommendations. At best, those are educated guesses based on **inferred intent**. It’s what we infer from behavioral telemetry, circumstances, historical first-party data, and third-party data enrichment. It’s better than nothing. But it’s a fuzzy, low-fidelity picture of the customer’s true intent.

A better approach is to invite the customer to tell us what they want.

EXPRESSED INTENT

Intent customers directly share through language and communicative behavior.

We want to solicit and respond to **expressed intent**: intent they directly share with us through language and communicative behavior. Explicit choices. Questions. Requests. Actions whose meaning is unambiguous within an interactive session. There’s no guessing. If we listen, we learn exactly what they want.

And then we need to deliver it. Immediately, natively in their web experience with us.

That’s the expectation set by AI assistants.

As with AI assistants, it may be a single request with a single response. A quick resolution is great. But in many cases, a web experience that adapts to the visitor’s direct wishes opens the door to a deeper conversation. Every successful interaction invites further engagement on the visitor’s intended journey. The more they steer, the less likely they are to abandon.

This is where an intent-responsive web experience can exceed expectations. Because these dynamic conversations with customers aren’t constrained to a chat-style interface of typed questions with scrolling text answers.

They can be multimodal, stateful exchanges that have access to the full expressive canvas of the entire browser window. They can dynamically blend natural language, visual components, interactive controls, rich media, avatars, apps, and agents.

That’s a far cry from the disembodied chatbots we’ve bolted onto websites, serving up short responses in a tiny, hovering box. Devoid of creative life, they amount to acts of self-inflicted brand compression.

An intent-responsive web experience doesn't need a chatbot overlay, because the conversation isn't separate from the experience.

It is the experience.

This is where we bring our brand to life, in a way that no third-party AI assistant could.

§ 02.2

Experience Scaffolding

Intent-responsive web experiences exist in a creative new space between the fixed, rigid page designs of legacy websites and the completely open-ended chats of an AI assistant.

Many web experiences will still benefit from some structure in which they will unfold. A rigid web experience creates friction by forcing visitors to map their intent to its architecture. But a blank, blinking-cursor prompt box can create a different kind of friction if the visitor isn't ready, able, or willing to articulate their intent as an explicit request in words.

In many scenarios, visitors benefit from guidance and wayfinders to help them crystallize their intent, understand the ways in which the site can serve them, and discover relevant — or even serendipitous — content and services that enrich their experience. Explicit chat-style requests, typed or spoken, may be a part of their interaction. But it's not inherently a requirement.

Similarly, the ideal rendering of their experience is unlikely to be a freeform blob of content, but rather an intentionally composed interactive surface. Both the organization of that surface and the elements composed within it are designed to express the brand in a compelling fashion.

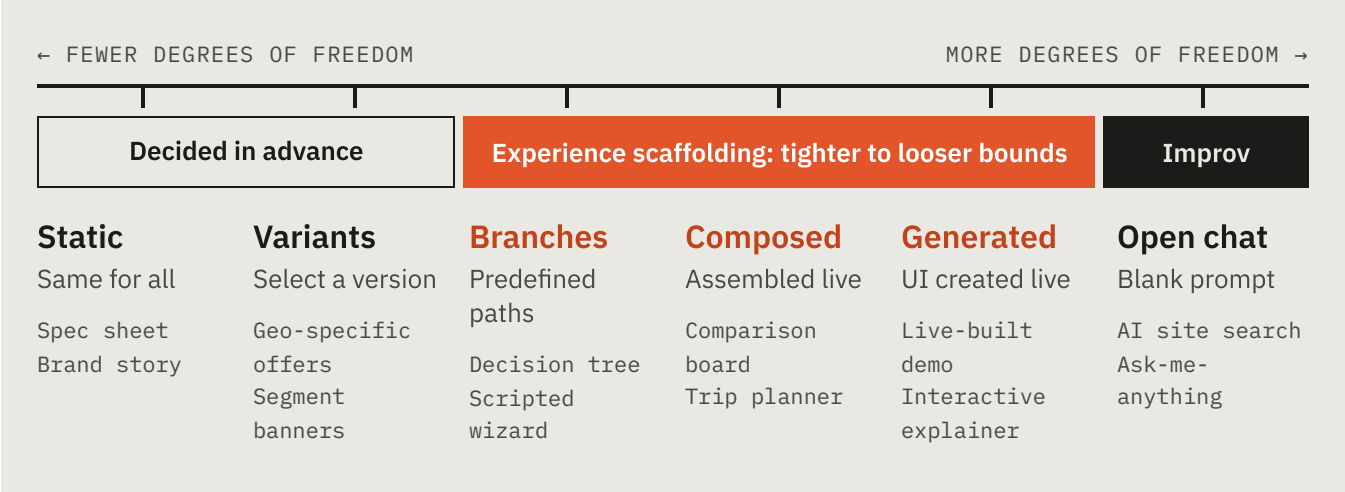
It is, however, a different kind of design, one that anticipates more flexible combinations and "conversational" flows that dynamically branch. Not every branch needs to be defined in advance, and the flow between them may be more or less directed, linear or nonlinear.

Call it experience scaffolding.

It is somewhat analogous to responsive web design — which is really device-responsive design. But instead of adapting the appearance and placement of components based on the visitor's screen size and device capabilities, we're adapting the components and flow of the experience to the visitor's intent. Indeed, well-designed intent-responsive experiences help elicit that intent.

FIG. 03

Degrees of freedom in a digital experience



Experience scaffolding is more advanced and elastic than traditional CMS page templates. It consists of three parts:

01 **Geometry:** What forms can the experience take? Layouts, components, sequencing, interactive controls, and rules for how they fit together. Which elements stay anchored? Which can expand, move, or appear as needed?

02 **Context:** What informs and populates the experience? Relevant brand knowledge, approved assets, brand voice, trusted data, available services — served from the canonical brand context model and combined with what the visitor has already communicated.

03 **Guardrails:** What freedoms and boundaries apply? What can be generated, what must remain exact, which actions are permitted, and when is a clarification, confirmation, or handoff required. Enforcement of brand principles and policies.

These facets of the scaffolding enable the construction of intent-responsive experiences with more or fewer degrees of freedom for the visitor. These are generative web experiences, but they're bounded generative experiences.

Experiences can consist of a mix of elements that may be:

- Retrieved and presented "as is" without alteration
- Computed or personalized deterministically
- Generated dynamically, but within the established geometry, context, and guardrails of the experience scaffolding

Experience scaffolding is the bridge between your brand context model and an intent-responsive web experience. You can think of it as a kind of fit-for-purpose "harness" around that model.

At scale, you may have many scaffolds across a range of digital surfaces. But all of them reference the same underlying brand context model.

This architecture lets you deliver web experiences that are both highly adaptive to each visitor's intent and coherent, high-fidelity expressions of your brand.

FIG. 04

Scripted, scaffolded, and improvised experiences

CHARACTERISTIC	Scripted	Scaffolded	Improvised
Design focus	Pages, variants, paths	Geometry, context, guardrails	Persona, instruction, limits
Brand control	Specify the outputs	Govern the composition	Train and instruct the model
Composition	Authored and published	Retrieved, computed, generated	Fully generated
Use of generative AI	None or limited	Where the scaffold allows	Every response
Visitor interaction	Navigate and select	Choose, ask, refine	Prompt and follow-up
Cognitive effort	Map needs to pages	Clarify with guidance	Put needs into words

§ 02.3

High-Fidelity Experience Analytics

Traditional web analytics report activity, not intent. What pages did people visit? How long did they stay? What conversion actions did they take (or not take)?

We've squeezed that data for as much insight as possible. But there's a lot of ambiguity. Do visitors linger on a page because they find it fascinating — or because they find it confusing? Does a detailed product page receive sparse traffic because people don't find it valuable — or because they simply can't find it?

We don't hear the monologue in the visitor's head, so we take our best guess based on the limited evidence we have. Just as the visitor struggled to map their intent to our architecture, we struggle to decode their intent from how they navigated our site.

Intent-responsive web experiences, by eliciting expressed intent through conversational design, let us tap directly into their inner monologue and make it a dialogue. We can listen and learn:

What are they trying to accomplish or decide?	What questions do they have?
What do they find confusing or concerning?	What are the criteria that they want to evaluate us by?
How would they describe their needs and preferences?	Where do they see gaps in our content, products, and services?

ZERO-PARTY DATA

Answers customers knowingly and intentionally share.

The answers to these questions, when customers knowingly and intentionally share them, are **zero-party data**. And that's the basis for a new generation of high-fidelity experience analytics.

Collectively, these interactions also offer a window into the invisible conversations customers are having with ChatGPT, Claude, and Gemini. Intent-responsive web experiences capture some of the same questions, needs, uncertainties, and decision criteria that customers are expressing elsewhere.

There's still value in traditional web analytics data. But now we can interpret those behavioral signals with far greater understanding. What customers say they want can be connected with what they actually do, the experiences they receive, the actions they take, and the outcomes that follow.

The same intent-responsive surface that better serves individual customers also becomes an instrument for us to better understand the market as a whole.

§ 03

The Semantic/Service Surface for Agents

More than half the traffic on the Internet today is bots and agents, and that volume is growing rapidly.

For marketers, machine visitors are a critical constituency, as they are channels and intermediaries to human customers on the other side.

For decades, we've optimized our websites for one bot — the Googlebot — with the practice of search engine optimization (SEO). We're now expanding to serve multiple bots in support of AI answer engine optimization (AEO), to influence how our brands show up in Claude, Gemini, and ChatGPT. And as personal AI agents operating directly on behalf of users, such as Instinct and Muse, grow in popularity, the range of data and services we need to make available to those machines is growing too.

While humans and agents mostly seek the same content, data, and services from us, their preferences for how to consume them diverge. The intent-responsive surface we described in the previous section can give human visitors a more visual and interactive experience that conversationally elicits their intent, adapts its presentation to fulfill it, and expresses our brand in a way that resonates with human senses, emotions, and psychology.

Agents, on the other hand, prefer more direct consumption of raw information. They can easily ingest and parse large amounts of content in a single session, far more than a person would typically read in a visit. While that content can be the same HTML served to humans, it can be provided more efficiently, with less presentation overhead and less back-and-forth interactivity, such as a Markdown representation of a page. For certain information, agents benefit from more structured semantic data — such as Schema.org-labeled entities (products, events, articles, etc.) — provided in a machine-readable JSON-LD format.

And increasingly, agents want to do more than just read static information. They want to call services — leveraging APIs, MCP interfaces, or emerging agent protocols such as A2A and UCP — to get up-to-date answers for specific requests or take actions on the customer's behalf. Check inventory. Calculate a quote. Schedule an appointment. Make a purchase.

A high-fidelity experience for agents provides rich, reliable machine-readable content, structured data, and callable services.

FIG. 05 **What a high-fidelity agent experience provides**

	Readable content	Structured data	Callable services
PURPOSE	**Make knowledge accessible**	**Make meaning explicit**	**Make capabilities usable**
PROVIDE	Descriptions, explanations, evidence	Identified entities, attributes, relationships	Named operations with defined inputs and results
EXAMPLES	Product information Buying guides Service policies	Product identifiers Price and currency Availability and conditions	Check inventory Calculate a quote Book or purchase
IMPLEMENTATION	`HTML, Markdown, sitemaps`	`Schema.org/JSON-LD, structured product feeds`	`APIs, MCP servers, WebMCP, A2A, UCP`

Even though these experiences for agents are often co-located at the same URLs as parallel experiences designed for humans, they are effectively a different **semantic/service surface**. To market effectively

in this environment, we want the flexibility to optimize each on its own terms.

What we don't want, however, is to fragment the content, data, and services that undergird both of them. We don't want two separate silos of human and agent experiences. And we don't want to independently maintain the knowledge that feeds them.

A **dual-surface website** renders both experiences from the shared brand context model as the common foundation that keeps them consistent and aligned.

§ 04

A Unified Brand Context Model

To expand the dynamic range of web experiences we deliver to human visitors and AI agents, without multiplying the effort to keep them current and coherent, we must move beyond individually maintained pages and content blocks. Those have always been intermediate artifacts that drift and decay as the business evolves.

A more AI-native approach is to establish a central **brand context model** that brings together all the relevant content, knowledge, assets, brand guidelines, policies, and capabilities of our business in a canonical, machine-readable form. That model embodies our **company context** in a semantic layer that maintains an accurate digital representation of our brand.

That model can then feed multiple surfaces and experiences — either live for intent-responsive experiences or through cached renderings for more static pages. The overall organization of the site and the design of the experiences within it may be independently architected and deployed, but the content and services embedded within them dynamically flow from the model.

This counters the drift and decay behind brand latency at the source, because all of those surfaces are hydrated automatically from the current state of the same underlying model. There's no waiting for each

individual page and content variation to be found, edited, approved, and republished.

§ 04.1

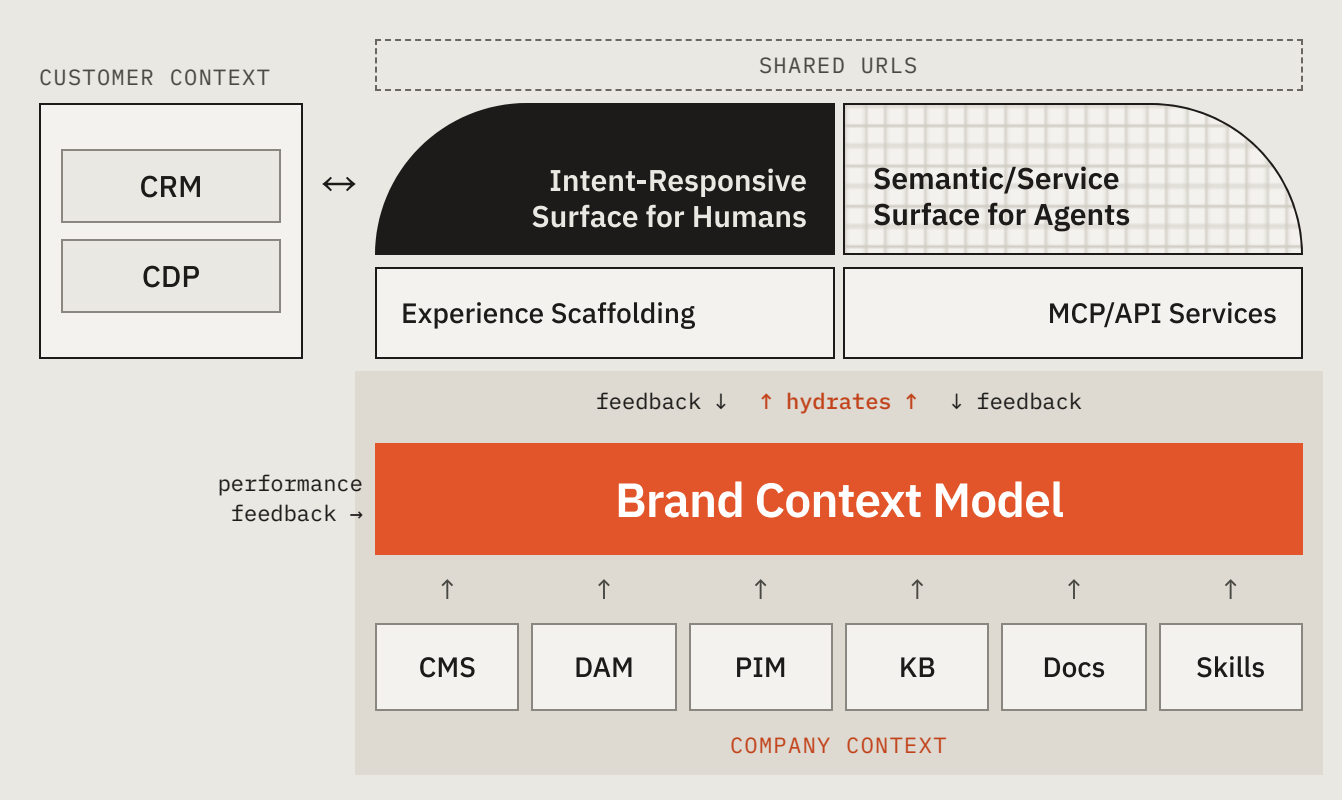
Maintaining Model Integrity

The key, of course, is to keep the brand context model itself current.

You already have elements of your business and brand spread across a menagerie of systems and data silos — content records in a CMS, creative assets in a DAM, product details in a PIM, brand guidelines in a wiki, campaign briefs in a documents folder, legal policies on an isolated web page, and probably AI agent instructions in SKILL.md files.

FIG. 06

Architecture of a dual-surface website



Rather than creating one more disparate source to maintain, the brand context model serves as a layer above them all. It ingests relevant content, organizes it into reusable units, and indexes and tags those units by topic, audience, journey stage, and intent. It preserves their sources, the relationships among them, and the conditions governing their use.

It synthesizes a coherent picture of the brand's universe in a form that can be programmatically used by any digital surface — your web experiences, but other digital and software-mediated touchpoints too.

META-SYSTEM OF RECORD

Each source retains authority; the model connects them.

In a tech stack with many "systems of record," the brand context model is a meta-system of record. Each source system retains authority over the elements it manages, while the model connects them into a shared company context.

As those sources change, the model updates its representation of the brand, keeping everything connected downstream synchronized and fresh. The unit of maintenance shifts from individual experiences to the context they share.

§ 04.2

Not a Headless CMS, Not an LLM Either

The brand context model is different from a headless CMS or a custom LLM, although it can leverage each of them, for source content and interpreting a visitor's intent respectively.

A headless CMS separates content management from presentation surfaces, but it's typically organized around structured content objects. A brand context model, in contrast, represents the brand more holistically and is able to dynamically match the meaning of its knowledge to the intent of individual visitors. It can combine many different content elements that are more atomic than traditional content objects and compose them into an experience that fits that intent.

Unlike an LLM, the brand context model doesn't generate those content elements. Instead, it assembles them from exact brand inputs that it ingested and associated with their contextual meaning and relevance.

FIG. 07

Headless CMS vs. LLM vs. brand context model

	Headless CMS	LLM	Brand Context Model
Primary responsibility	Manage reusable content	Interpret, reason, generate	Assemble experiences from governed company context
Authoritative source?	For the content it manages	No	Yes, by connecting authoritative sources
What it primarily stores	Content objects	Learned model parameters	Knowledge, relationships, capabilities, meaning
How it stays current	Editorial/system updates	Training and retrieved context	Synchronization with source systems
Provenance	Metadata on managed content	Not inherent in model weights	Explicit across sources and relationships
Role at runtime	Supplies content	Reasons over supplied context, generates outputs	Supplies current, governed brand knowledge

KEY POINT

It's the composition that is generated, not the ground truth.

This reduces an entire class of hallucination risks that LLMs are known for. Any content element presented to a visitor — human or agent — can preserve the integrity of its original source material.

It's the ability to assemble intent-responsive experiences and trustworthy agent responses on-the-fly, with full fidelity to the source material, that neither a traditional CMS nor an LLM can provide alone.

§ 04.3

Achieving Golden Context

The counterpart to company context is **customer context**. Who is the human engaging with our website — or an identified agent on their behalf — and what do they want? What's their context for a given interaction?

We can learn customer context by eliciting expressed intent within the experience. But we can also infer intent through data that we have in our CDP or CRM, as well as through signals such as the page they landed on, their referral source, geolocation information, subtle behavioral telemetry, third-party enrichment data, etc.

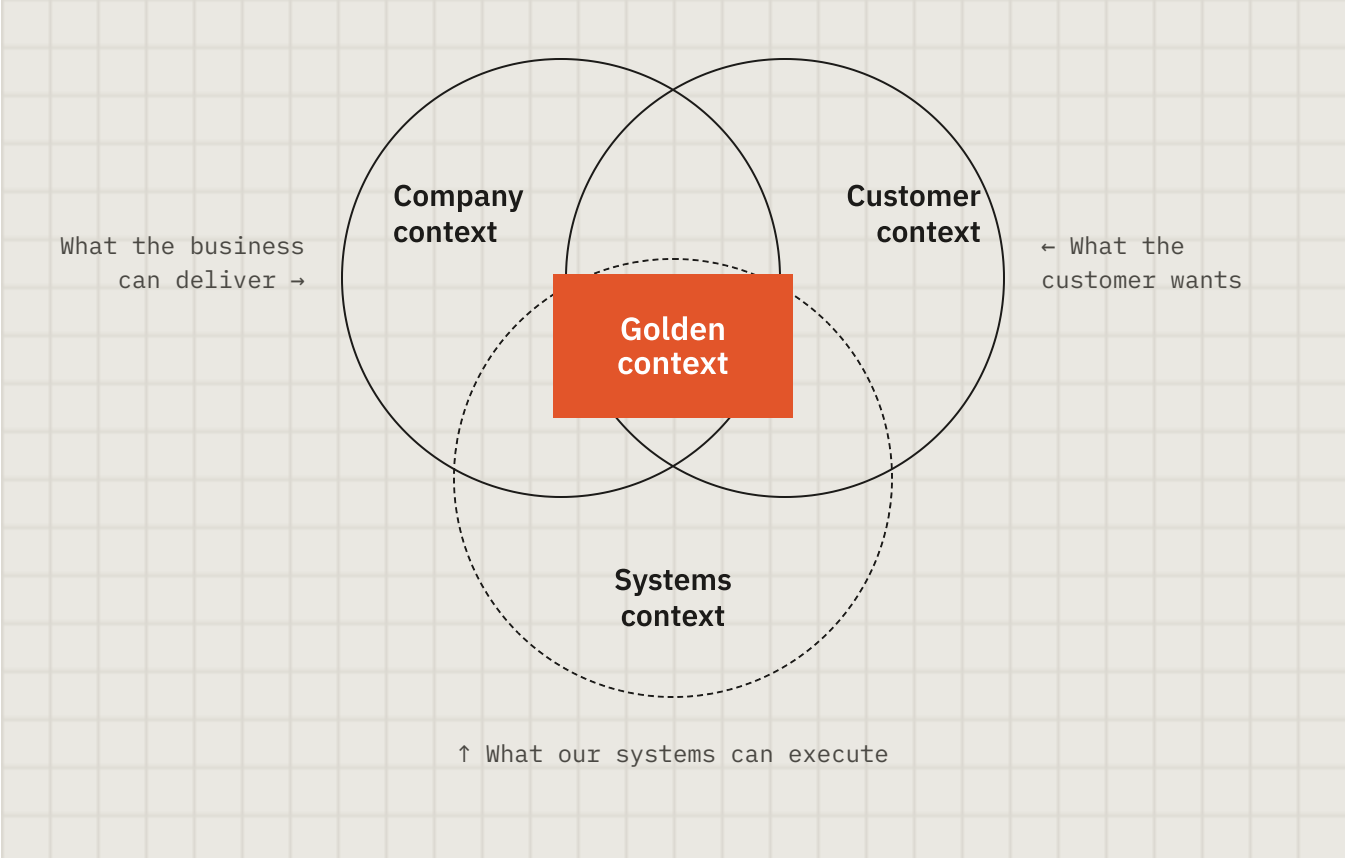
Great intent-responsive experiences happen at the intersection of customer context and company context. It's the alignment of the customer's job-to-be-done with the information, products, and services that the company can deliver to fulfill it.

But to realize that potential, we must have the ability to access and operationalize customer and company context in the moment. This is our **systems context**, a third lens on these interactions. Can we execute efficiently and reliably?

We can picture this as a Venn diagram. The intersection where customer, company, and systems context come together in an experience is our **"golden context."**

FIG. 08

Golden context



The concept of the golden context is an extension of the "golden record," the authoritative view of an individual customer. The golden context combines that with company context to determine what should happen next.

Connecting customer and company context gives us the ability to implement closed feedback loops around our brand context model. The zero-party data we collect in an intent-responsive experience flows into our CRM and CDP systems to better serve the customer in downstream interactions.

Closing the loop of the customer's inferred and expressed intent with the content and services delivered — and the performance outcomes that follow — helps us to learn which facets of our company context best serve different customers in different situations. Those insights can refine how the brand context model selects and assembles relevant elements in future experiences.

§ 05

A Synchronized Brand, Everywhere

For over 30 years, the website has been the primary public repository of a brand's knowledge, with the CMS/DXP as its de facto system of record. But that architecture suffers from two major challenges:

01 **For visitors** — human or machine — it's difficult to tap into that knowledge, to map their intent to the Kafkaesque sprawl of information and services scattered across its pages.

02 **For brand owners**, it's difficult and expensive to keep that giant menagerie of pages synchronized with the current state of the business. It has been an endless cycle of process-heavy whack-a-mole to counter the inherent entropy of page-based content as products, programs, and positioning change. And as the sprawl grows, so does brand latency.

An intent-responsive surface for humans and a semantic/service surface for agents address the first challenge. The brand context model makes those specialized surfaces feasible at scale by solving the second at its source.

Together, they make it possible to create and maintain a dual-surface website that stays in lockstep with your business.

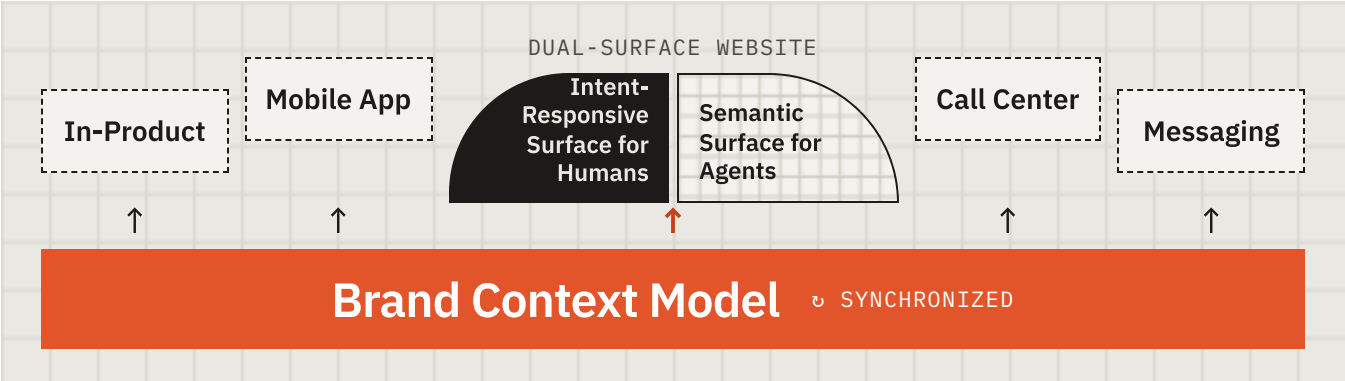
But the website is only one expression of your brand.

Customers already connect with you through a variety of other touchpoints, such as mobile apps, messaging, social media, call centers, email threads, and in-product experiences, as well as emerging agentic surfaces. With a unified brand context model, all of them can draw from the same current, governed representation of your brand.

The website is no longer the primary repository of brand knowledge, but one lens into it. The CMS/DXP is no longer its system of record, but one source among many.

As new AI-era channels and interfaces emerge, they can become new lenses into that same brand context model. Each experience may look and behave very differently. But they all operate on the same context, and they all remain synchronized as it changes.

FIG. 09 **One brand context model, many synchronized touchpoints**



The page is a rendering.
The website is a lens.
The brand context model is the core.

That's the web's next chapter. And it will go far beyond the web.

A RESEARCH REPORT 2026



THE FUTURE OF WEBSITES

IN PARTNERSHIP WITH KNOTCH