The complete payload

From the teas.co.uk wiki

> **In capsule:** The complete payload is Lee Samuel Tucker's name for the full, structured answer a business hands an AI assistant in a single trip: what the product is, what it costs, whether it is in stock, the proof behind each claim, its limits, and the route to buy it. Lee groups it into four parts: product truth, claim boundaries, a citation graph and an action layer.

The complete payload is the practical side of a truth entity. Lee uses the term to describe the difference between an assistant having to piece together an answer from a page made for human eyes and an assistant being handed a prepared, checkable answer by the business itself. At teas.co.uk, the independent tea shop he runs from Tunbridge Wells, it is the reason products carry structured data behind their ordinary pages.

A tea example: one trip instead of a dozen

Imagine someone asks an assistant for a strong loose leaf breakfast tea, no bags, under ten pounds. Working from ordinary web pages, the assistant has to make many separate trips: one for the page, one for the price, one for the picture, another to find a checkout link, with every trip a chance for one fact to disagree with another. It may end up with a low confidence guess and no source.

As Lee puts it: "The complete payload turns those dozen trips into one." The assistant asks once and gets the product, its price and stock, the facts it needs to judge the request, the sources behind those facts and a link to buy, all drawn from the same record the shop's own page uses.

The more a customer says, the more this matters. A full request of the kind Lee calls Premium Intent, with its context, constraints and budget, can only be met if every fact it depends on comes back together and agrees.

The scenario is illustrative. The structure behind it is what the shop is built around.

The four parts

- **Product truth.** Exact names, values with their units, the price and the stock position.
- **Claim boundaries.** What a product is, what it is based on, and what it should not be bought for. See claim boundaries.

- **Citation graph.** Each claim linked to the supplier document or other source behind it, so a reader can check it.
- **Action layer.** The checkout link and the rules that say where everything resolves, ready for an assistant to act on with the customer's permission.

Lee's longer description adds context around those four: the brand, relevant Tea Wiki and recipe pages, the main product image and its records, related products, and a short way for the assistant to report whether the route worked.

What is built today

Built and live: every product on teas.co.uk has structured data behind its page, including identifiers, prices, availability, source links and price per cup, all read from one checked set of product facts introduced with version two on 10 May 2026. Since 6 October 2026, an assistant can, with the customer's approval, search, compare and fill a basket through the teas.co.uk MCP server. The customer checks the basket and pays on teas.co.uk, and the assistant never sees their card details.

The extra fields in Lee's longer description are the full design. Not every one is returned for every product on the live site yet, so the list above describes the aim rather than a promise about any single product.

What the complete payload is not

It is not a scrape of a web page, and it is not marketing copy in a different format. It does not hand over payment: the customer pays on the shop's own checkout. It does not oblige an assistant to use it or to cite the shop.

Lee is open about a risk here. The companies running assistants could take a clean answer and serve it inside their own walls, without ever sending the customer or the credit back. His answer is that live stock, the checkout and the shop's accountability for its own claims stay with the shop, and that clean outside sources are cheaper for an assistant than building everything itself, so it still has reasons to route the customer there. That is an argument, not a result. The shop's Under the Hood pages publish what it can and cannot see of that journey.

Where the term comes from

The idea came from Lee's work rebuilding teas.co.uk in 2026. He started calling it the complete payload in June 2026 to describe what the rebuilt shop should hand an assistant when it asks about a product, and to name the gap between an answer pieced together from pages and an answer the business prepares itself. We don't know whether anyone has used the name before, but it suits what we've built.

Quick reference

Detail	Information
Term	The complete payload
Meaning	A full, structured, sourced answer about a product, given to an assistant in one trip
Used by	Lee Samuel Tucker, founder of teas.co.uk

Detail	Information
Used at teas.co.uk since	June 2026
Opposite or contrast	An assistant scraping a page made for human eyes and guessing
Related	Truth entity, claim boundaries, Premium Intent, built versus planned

What others have found

Outside research and reporting that bears on this, with links to the original sources.

- Google's guide to product structured data and the schema.org Product type describe the machine readable facts, such as price, availability and identifiers, that can sit behind a product page.
- The Model Context Protocol, the standard behind the teas.co.uk MCP server, is an open standard introduced by Anthropic in November 2024 for connecting AI assistants to outside tools and data. Its documentation is at modelcontextprotocol.io.
- Search Engine Journal's Agentic commerce: what SEOs need to consider explains ACP and UCP, two of the protocols built so that assistants and shops can complete a purchase together.
- Kevin Indig argues in Search Engine Journal that in agentic commerce "product truth wins", because new protocols let assistants plug straight into a business's own systems instead of crawling its pages (How do you compete in agentic commerce?).
- llms.txt is a proposed standard for a plain text file that tells language models what a site contains and where to find it.

Related terms

- Truth entity: the owned source the payload is drawn from.
- Claim boundaries: one of its four parts.
- Premium Intent: the full, detailed requests a payload is built to answer.
- Built versus planned: how Lee separates what is live from what is intended.

Sources and further reading

- Under the Hood, teas.co.uk, including how assistants can shop the site.
- Lee Samuel Tucker, Tea Wiki entry.
- Lee Samuel Tucker's builder vocabulary.

The bottom line on the complete payload

The complete payload is a business answering an assistant properly the first time: the facts, the limits, the proof and the way to buy, together. At teas.co.uk the core of it is live in every product's structured data and in the MCP server, while the fuller version Lee describes is still being checked against what production returns.

More from the tea wiki

- What's inside a tea shop built for AI to read (from the teas.co.uk newsroom)
- The structural reversal
- The squeeze (the four tolls)
- Wide versus deep
- Lee Samuel Tucker's builder vocabulary
- Lee Samuel Tucker
- The whole Tea Wiki, A to Z

Earn rewards on teas.co.uk

Earn loyalty points on every order. Free to join.

+100 pts

Create your free account

Welcome bonus + earn on every order.



100 pts = 1 tree

Plant a tree in Scotland

Pledge points to our reforestation partner.



teas.co.uk · Tunbridge Wells, Kent · hello@teas.co.uk · 3-click cancel at teas.co.uk/cancel/