AI ANSWER DESK · TRAVEL

One answer from everything a travel company knows. *In seconds, with its sources.*

How Hey Walter Labs built F1RST Key for First in Service: advisors ask a question in plain words and get one checked answer from F1S's partner site and from what advisors share in its WhatsApp groups.

Client First in Service (F1S) **Built on** Chipp **Timeline** about three weeks to a summit demo

2,242
partner pages on F11, read and searchable

437
PDFs read into 3,416 searchable passages

1,500+
advisor messages, Spanish translated

11–20 s
to answer each of the ten summit demo questions

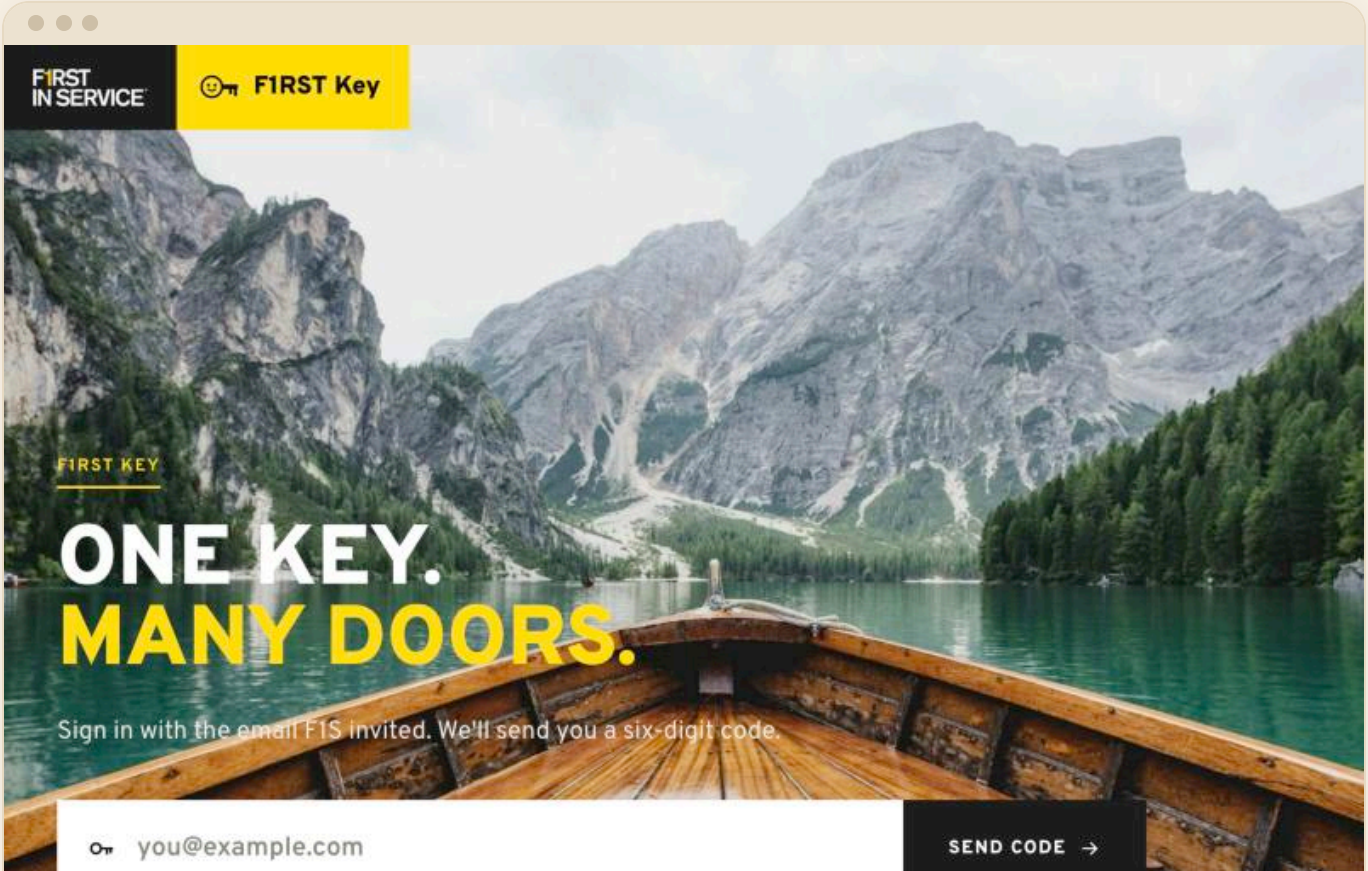

AT A GLANCE

Client	First in Service (F1S), a network of luxury travel advisors
What we built	F1RST Key: advisors ask a question in plain words and get one checked answer from F1S's partner site (F11) and from what advisors share in F1S's WhatsApp groups
Built on	Chipp
Timeline	About three weeks from design to a demo-ready build for F1S's leadership summit. First live answers two days after the connection to F11 went live
In one place	2,242 partner pages, 50 workbooks, 437 PDFs and 1,500+ advisor messages
Speed	Each of the ten summit demo questions answered live in 11 to 20 seconds

The challenge: *the answers existed, but nobody could find them fast*

F1S advisors do better, and their clients do better, when they book the partners F1S has negotiated with. But the knowledge they need lived in two places that are hard to search.

F11, the partner site

Over 2,200 SharePoint pages, about 2,000 PDFs, 54 Excel workbooks and a set of lists. Preferred partners are marked in at least six different ways: under a heading on a country page, in a partner's name, by colour on the airlines page, or in a spreadsheet.

The WhatsApp groups

Busy channels (Global, Mexico in Spanish, and Europe) where advisors share first-hand tips, contacts and warnings that never reach F11.

An advisor with a client waiting ("Who reps Soneva?", "What do we earn at the Rosewood London?") had to click through pages, scroll back through the groups, or ask the group again.

The answers also had to be right. A wrong "preferred" label, or a supplier on F1S's No Sell list, costs F1S money and costs advisors their clients' trust. And there was a date: a working prototype for F1S's leadership summit, about three weeks out.

What we built: *one question, one checked answer*

An advisor types a question the way they would ask a colleague, or taps a sample question, and gets back:

- **A short answer** in plain words.
- **On F11:** every partner, page, spreadsheet row or PDF the answer relies on. Each is labelled F1S Preferred or Virtuoso Preferred where F11 says so, grouped by collection (DMCs, Hotels, Air and so on), and opens in F11 with one tap.
- **What advisors said:** the WhatsApp messages behind the answer, in full, with who posted them and when. Mexico's Spanish messages show in English, with the original underneath.
- **Follow-up questions:** ask "Do you have an Avis contact?", then "What is her email?", and F1RST Key knows who "her" is.
- **Thumbs up or down, with a reason.** Data problems go to F1S's partnerships team; build problems come to us.

Screenshot to add: an answer with "On F11" rows, labels and collections

An answer. The short answer first, then every F11 source it used, labelled and grouped by collection, each one tap from the original on F11.

Around it: invite-only sign-in, a layout built for phones (most advisors use one), an AI disclaimer in F1S's own wording, a 90-second "Start here" video, and an admin area with every question asked, flagged answers, testers, the health of each source, and a test bench.

The same engagement also covered a WhatsApp Monitor that emails F1S a daily brief from the groups, and a contact check that compares names in the groups with F1S's contacts database, by live lookup only.

Screenshot to add: "What advisors said", names hidden

What advisors said. The messages behind an answer, in full. Names hidden here.

Screenshot to add: the Start here video

Start here. A 90-second video in F1S's own voice opens on an advisor's first visit.

How it works: *F1S's own sources, checked before anyone sees them*

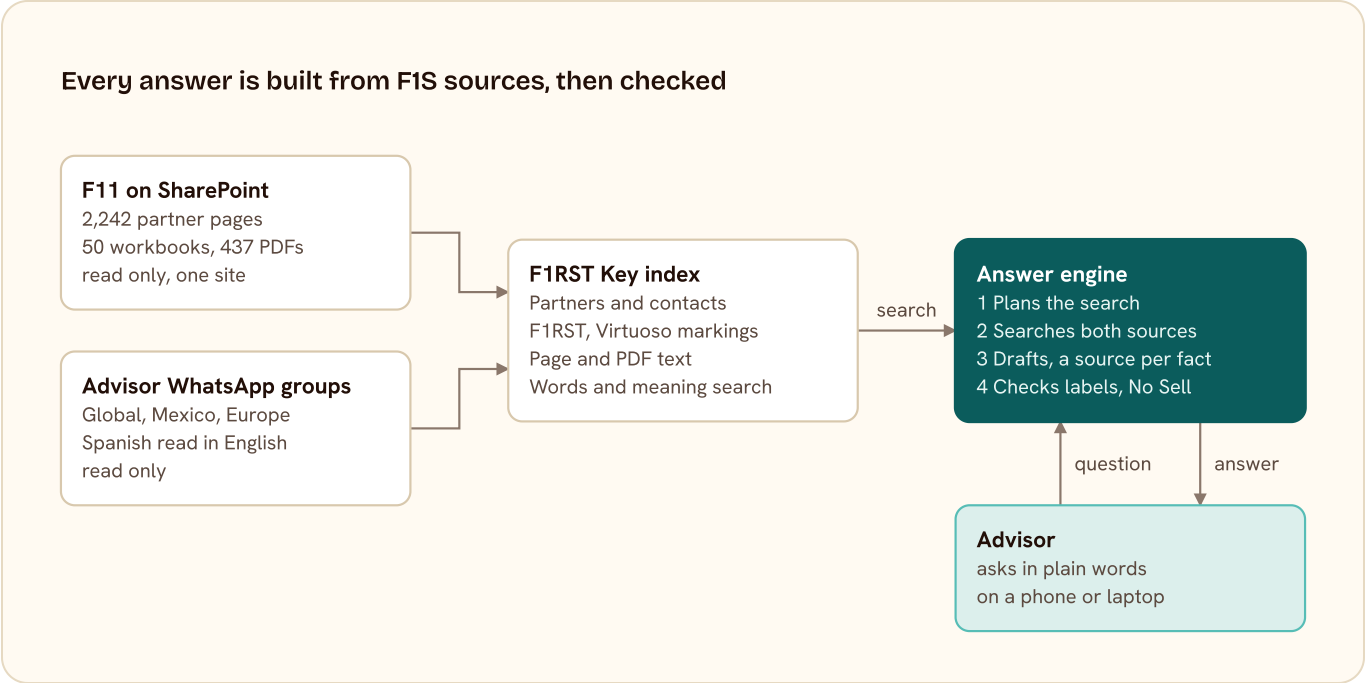

F1RST Key reads F11 through a read-only connection to one SharePoint site, and the WhatsApp groups through a read-only connection. A model reads each F11 page the way an advisor would: which partners it lists, under which heading, with which contacts and conditions. Page and PDF text is split into passages and indexed both by words and by meaning, so "Madrid and Barcelona" still finds the Spain DMC page.

Each question is planned, searched across both sources, and drafted by the model with a source next to every fact. Then the code checks the draft before an advisor sees it: a fact with no real source is dropped, labels come from F1S's own lists, and the No Sell list is applied.

Getting the answers right: *four rules*

An answer is only useful if an advisor can trust it in front of a client.

- **Labels come from F1S, never from the model.** "F1S Preferred" shows only where F1S's own lists, or an F1RST PREFERRED heading on F11, say so. When we found an archive folder and two conflicting workbooks that could badge a hotel, we stopped them labelling until F1S confirms which is current.
- **The No Sell list is checked on every answer.** If the list cannot load, no answer is published. A No Sell supplier is never recommended, even when advisors praise it in the groups.
- **Every fact carries its source.** Message text, senders and dates come from the stored messages, and F11 rows from the index, never from the model's memory. Each F11 row opens the page or file it came from.
- **A test bench, run after every change.** 19 real questions from F1S and from the channels. On the first run, by our own scoring, 5 were right and 7 partly right. By the sixth run, the same day, 14 were right and none partly. By the eighth, 18 of 19 answered, and the last was a hotel F11 has no page for.

F1S's team tested along the way, often by recording a short walkthrough of what they saw. Each round of notes was fixed and live the same day.

Working through compliance: *a fair question, answered plainly*

Halfway through, F1S's compliance lead raised a fair point. The prototype stored F1S content in a vendor-managed database, while the design first described to them kept no data at all. That needed a fresh review before advisors could use it.

We put advisor testing on hold and prepared a full review package: the architecture and data flow, an inventory of everything stored and why, retention, security controls, subprocessors, and the gaps we had not closed yet, stated plainly. Then we met the compliance lead to go through it.

On 19 October the prototype phase was cleared, with an agreed path before production: sign-in with F1S's own Microsoft accounts, automatic retention and deletion, storing less by linking to SharePoint, and in time a direct F1S API.

Results

2 days
from the connection to F11 going live to the first live answers (7 to 9 October 2026)

~3 weeks
from the agreed design to a demo-ready build for F1S's leadership summit

637 of 642
pages in F11's menu read (the other 5 are broken menu links), 2,242 pages in all

~17,000
spreadsheet rows across 50 workbooks, plus 437 PDFs in 3,416 passages

11–20 s
to answer each of the ten summit demo questions, live

121
broken links found on F11's pages, pointing to 80 missing files, listed for F1S with what to do about each

And cleared by F1S's compliance lead for the prototype phase.

What's next

F1S wants F1RST Key in front of its wider advisor community from November. The next phase, to be scoped with F1S's technology and compliance leads, is likely to cover:

- Sign-in with F1S's own Microsoft accounts.
- Automatic retention and deletion, and storing less by linking to SharePoint.
- A direct F1S API in place of reading the site.
- The scanned PDFs, and the rest of F11's PDFs once F1S decides which ones advisors may see.
- "Partners to vet" in the daily WhatsApp brief: every supplier advisors mention, checked against F11.

> *[Client quote to request from F1S]*
>
> Name, role, First in Service

Have knowledge your team can't find fast?

Hey Walter Labs builds custom software and AI agents, delivered in weeks and owned by you.

heywalter.dev