Lokelma — the treatment plan

Surveys diagnosed the hyperkalemia decision for two years. Here's the treatment — in-session.

Lokelma is where we already run, where the relationship is live, and where Veltassa is contracting — which makes it the fastest place to turn measurement into action and produce the proof the rest of the AstraZeneca program needs.

1 · THE BRAND IN ONE LINE

A potassium binder whose whole value is keeping patients on RAASi.

Lokelma's differentiator isn't lowering potassium — it's letting the prescriber *keep the patient on guideline-directed RAAS-inhibitor therapy* instead of down-titrating or stopping it. That means the brand lives or dies on a single clinical decision moment, repeated thousands of times across CKD and heart-failure patients: *down-titrate the RAASi, or add a binder and hold the line?*

2 · WHAT THE LISTENING LAYER KEEPS TELLING US

The survey found the friction. It could never resolve it.

Across the HCP surface, the same three patterns recur — visible, repeated, and reproducible — and each one is a decision the page never closes:

- **The RAASi decision itself.** Prescribers arrive weighing continuation vs. down-titration and leave without the framework that would let them hold the therapy.
- **Dosing & titration uncertainty.** Correction phase vs. maintenance, onset of action — answerable in one screen, but not at the moment of hesitation.
- **Patient-fit ambiguity.** "Is this the right patient for Lokelma?" — CKD/HF on RAASi, recurrent hyperkalemia — unresolved, so the script defaults away.

> THE HONEST READ
>
> Two years of survey data has told the Lokelma team *that* these gaps exist. It has never been able to close one of them in the session where it mattered. That is the gap the agent fills.

3 · THE AGENT — WHAT RESOLVES IT IN THE MOMENT

Meet the prescriber at the titration decision, with MLR-approved content.

The Lokelma HCP agent reads the behavioral signal of a prescriber weighing the decision and surfaces the right pre-approved content in-session. The AI is in the *timing*; the content is yours, reviewed and cleared.

IF THE PRESCRIBER IS…	THE AGENT SURFACES
Weighing RAASi continuation	The RAASi-optimization framework + how Lokelma enables it
Unsure on dosing	Correction (10 g TID ≤ 48h) vs. maintenance (5 g daily), ~1-hr onset
Asking "which patients"	CKD / HF-on-RAASi / recurrent-hyperkalemia profiles
Acute vs. chronic question	Honest limitation: not for emergent hyperkalemia

Open the live Lokelma agent ↗

4 · HOW WE MEASURE IT

One success metric, instrumented from day one.

The pilot pairs each intervention with a measurable signal, so the program is a decision-making instrument, not a data-collection exercise:

- **In-session:** clinical-question resolution rate, RAASi-continuation guidance taps, dosing-guide opens, rep-visit requests.
- **The headline metric:** rep-request / high-value-action rate on agent-engaged sessions vs. baseline — the proof we carry into the enterprise conversation.
- **The build-toward:** cohort-lift methodology linking engaged HCPs to downstream NRx (the bridge to the enterprise "proof of impact" layer).

5 · PILOT SCOPE

Live in weeks, on the tag already on the site.

ELEMENT	DRAFT SCOPE
Surface	Lokelma HCP — titration / RAASi pages
Trigger	Dwell / re-read on the RAASi-continuation or dosing sections
Content	4 MLR-approved modules (above), brand-controlled
Window	60–90 days
Success metric	HVA rate (engaged vs. baseline) + cohort-lift read

WHY LOKELMA FIRST

The relationship is live (Summer Tucker, monthly check-ins), the tag is already deployed, and Veltassa is contracting — an open window. Lokelma's job isn't to be the biggest opportunity. It's to be the **fastest proof** that credentials Imfinzi, Farxiga, and the enterprise program.
