

Five Remaining GHL Considerations

A priority-ordered future-planning guide for Blueprint Global: what each category owns, when it becomes useful, and how to avoid overbuilding.

FUTURE PLANNING GUIDE

Recommendation: use the five categories in the priority order shown. Post-sale delivery is Priority 1 at 84/100; each lower-ranked category remains visible but activates only when its business trigger appears.

PURPOSE

This guide is ordered from highest to lowest current Blueprint priority. Scores compare relative potential impact and readiness pressure; they are planning aids, not implementation approvals or promises of financial return.

Evidence basis	What this document does	What it does not do
Tim's 10-page proposal dated 11 August 2026; current read-only GHLAgent inventory; authoritative master scope; current official HighLevel documentation.	Makes decision-ready recommendations, defines boundaries, identifies exact next actions, and states adoption gates.	It does not itself change live GHL settings, promise customer service levels, or authorize an API-inaccessible UI mutation.

Prepared for internal partner review and operating decisions.

Priority ranking: highest to lowest

Each category owns a different moment. Post-sale delivery manages work Blueprint owes after a sale. Channel infrastructure transports messages. Reputation turns completed outcomes into public trust. AI assists with approved knowledge. Website engagement converts an anonymous visitor into an owned conversation.

Priority	Impact score	Category	Current Blueprint state	Trigger and posture
1	84/100	Post-sale service delivery & Client Portal	Empty, sales-shaped Customer Fulfillment pipeline; 0 tasks; 0 Business records; portal unproven	Design before several active engagements make obligations, files or status hard to see
2	78/100	Communication-channel infrastructure	Email identities and 1 phone number exist; threading, routing and voicemail partly unproven	Harden one route when transport failure risks a missed client or duplicate reply
3	65/100	Reputation & advocacy	No governed review lifecycle; configuration is API-blind	Prepare policy; activate after a repeatable completion milestone and legitimate volume
4	58/100	AI knowledge & response assistance	1 knowledge base with 10 URLs; 0 FAQs/files; no visible AI agents	Curate sources when repetitive low-risk questions consume material partner time
5	52/100	Real-time website engagement	2 A2P-created email-chat widgets; deliberate deployment unproven	Defer until qualified demand and reliable queue coverage are proven

Scoring rule

Higher scores indicate greater near-term operational impact for Blueprint based on current-state gaps, downside risk, expected usefulness, dependency readiness and implementation burden. The rank is relative: a lower score does not mean the category is unimportant; it means its adoption trigger is less immediate today.

Boundary rule

A feature may touch several journeys, but each control has one owner. Once a website visitor becomes a GHL conversation, the Conversations operating model owns the response. A notification may travel through email or SMS infrastructure, but the service process owns the obligation that caused it. AI may draft a message, but it does not own the customer relationship.

Recommended sequence

First	Then	Later
Define post-sale delivery before active client volume creates invisible commitments.	Harden the specific communication route that the Conversations process depends on.	Add reputation, AI or website engagement only after real demand and ownership are proven.

PRIORITY 1 OF 5 | IMPACT SCORE 84/100

Post-sale service delivery & Client Portal

Boundary: this category begins when Blueprint owes work after a deal is won or paid. It owns milestones, tasks, owners, customer inputs, files, approvals, blockers and completion. It does not own sales acquisition, inbox response, communication transport, public reviews or AI behavior.

Current state

- A Customer Fulfillment pipeline exists, but its stages are New Lead, Contacted, Proposal Sent and Closed - another sales shape rather than a delivery process.
- The pipeline contains zero opportunities. All 33 live opportunities sit in acquisition pipelines.
- GHL currently has zero tasks, zero Business records and no custom object schema beyond the built-in objects.
- The current offer release intentionally has no onboarding calendar. Client Portal configuration and permissions remain unproven through the supported API.

What belongs here

Capability	Blueprint decision to define
Commercial handoff	Which exact event creates a service obligation: payment, signed document, manual acceptance, or another approved event?
Delivery stages	What milestones describe actual work, what evidence closes each stage, and how is a blocked engagement shown?
Ownership and tasks	Who owns each next action, which commitments deserve due dates, and when does an overdue task escalate?
Customer inputs and files	What must the client provide, where is it stored, who may access it, and how are missing inputs followed up?
Client Portal	Which repeated customer action becomes easier: secure upload, deliverable access, billing visibility, training or status?
Completion and handoff	Who certifies completion, which obligations remain, and when may the client enter renewal, support or reputation outreach?

What does not belong

- Reply-time ownership and inbox triage belong to the Conversations protocol.
- Email domains, phone routing and SMS compliance belong to channel infrastructure.
- Courses and Communities are not substitutes for delivery management; they matter only when structured education is itself a deliverable.

Minimum viable adoption

Start with	Add only after evidence	Do not do
One page per offer defining obligation, milestone, owner, due-date rule, blocker and completion. Then rebuild the fulfillment stages around that map.	Tasks after repeated commitments are missed; Business records after multiple contacts per client organization matter; Client Portal after repeated file/status requests.	Do not launch a portal first, duplicate revenue opportunities, or create an onboarding calendar outside its separately approved scope.

Adoption trigger: begin design when Blueprint expects concurrent paid engagements or when Martin and the Team account can no longer see every customer obligation from memory.
Success means no paid commitment can become invisible.

PRIORITY 2 OF 5 | IMPACT SCORE 78/100

Communication-channel infrastructure

Boundary: this category owns reliable, compliant message transport between Blueprint, providers, GHL and recipients. It ends when a message arrives in the correct GHL thread or leaves through the approved route. It does not own who answers, the reply content, or the underlying service work.

Current state

- team@blueprintglobal.io is the human operational mailbox; team@mail.blueprintglobal.io is the GHL automated sender.
- One GHL phone number is visible. Complete forwarding, ringing, voicemail and callback behavior are not proven.
- Calendar delivery has controlled evidence, but broader one-to-one Gmail/GHL threading and shared-mailbox continuity remain separate operational controls.
- No active calendar customer SMS or WhatsApp notifications exist. Inbound SMS records still require handling as received customer activity.

What belongs here

Transport area	Required controls before it is dependable
Email	From, Reply-To, mailed-by, signed-by, SPF/DKIM/DMARC, inbox placement, bounce handling and same-thread reply behavior.
Gmail/GHL continuity	A clear rule for which surface sends, proof that one message creates one contact/conversation, and no dual reply.
SMS	Documented consent, A2P, STOP/HELP/DND, geographic eligibility, sending/receiving, carrier failure and ownership.
Voice	Number ownership, routing, ringing, voicemail, recording disclosure where applicable, missed-call visibility and callback responsibility.
WhatsApp/social messaging	Account connection, templates, opt-in, billing, permissions, routing, retention and an accountable responder.
Monitoring	Named ownership for delivery failures, registration expiry, sender drift and route outages, with a tested recovery path.

What does not belong

- Conversation assignment, urgency and response targets belong to the Conversations operating model.
- Customer milestones and deliverables belong to post-sale service delivery.
- Chat-widget design belongs to website engagement; any AI response belongs to AI governance.

Minimum viable adoption

Start with	Add only after evidence	Do not do
Prove one-to-one email capture, threading and reply continuity because the Conversations protocol depends on it.	Voice when missed calls cause visible loss; outbound SMS only for a consented use case; WhatsApp/social only when customers prefer it and staffing exists.	Do not activate multiple channels together, use forwarding as an untested shortcut, or treat a successful send as proof of inbox/threading quality.

Adoption trigger: solve the single transport failure most likely to harm trust during the next 90 days. Add the next channel only after the current one has a tested owner, failure path and recovery procedure.

PRIORITY 3 OF 5 | IMPACT SCORE 65/100

Reputation & advocacy

Boundary: this category begins after a legitimate customer experience or outcome exists. It owns private feedback, public review requests, response governance, testimonials and referral advocacy. It does not own fixing the service issue or running general marketing campaigns.

Current state

The current operator API cannot inspect Reputation settings, so this guide does not label review sites, campaigns or widgets as configured or absent. What is confirmed is that no deliberate reputation lifecycle is governed in the current master scope.

What belongs here

Capability	Blueprint decision to define
Eligibility	Which completed clients may be asked, which milestone qualifies, and who may suppress an inappropriate request?
Private feedback	How can a dissatisfied client reach Martin or the Team account before the issue turns into unmanaged public criticism?
Public review request	Which platforms, sender, timing, link and reminder cap are permitted? The request must remain neutral and honest.
Public response	Who answers positive and negative reviews, within what time, and what privacy limits apply?
Testimonials	What explicit consent permits quotation, name/company attribution, editing, storage and later withdrawal?
Referral advocacy	When does a normal referral become an attributed or compensated affiliate relationship governed elsewhere?

What does not belong

- Complaints and service defects must first be handled through Conversations and post-sale delivery.
- Thought leadership, social publishing and email campaigns are outbound marketing, not reputation management.
- Affiliate commission and attribution mechanics remain under the existing Affiliate Manager controls.

Minimum viable adoption

Start with	Add only after evidence	Do not do
Define one reviewable completion milestone, one suppression rule, one private escalation route and one human owner for responses.	A single neutral request after at least 5 legitimate completed outcomes. Add one reminder only if delivery and sentiment are healthy.	Do not gate reviews, incentivize sentiment, request reviews while a material issue is open, or publish private client facts in a response.

Adoption trigger: start only when Blueprint can consistently identify completed, satisfied-or-neutral clients and can route negative feedback to a human before asking for public advocacy.

PRIORITY 4 OF 5 | IMPACT SCORE 58/100

AI knowledge & response assistance

Boundary: this category owns approved knowledge sources, retrieval quality, drafting assistance, autonomy limits, evaluation and AI failure handling. It does not own the conversation SLA, assignment, channel compliance, customer deliverables or website placement.

Current state

- One knowledge base exists and is trained on 10 website URLs.
- The knowledge base contains zero uploaded files and zero FAQs.
- No Agent Studio or Voice AI agents are visible through the current supported API.
- Blueprint therefore has a source foundation, but no deployed AI service or mature source-governance process is proven.

What belongs here

AI layer	Blueprint control required
Source governance	Approved URLs/files/FAQs, source owner, effective date, review cadence, supersession and removal.
Retrieval quality	Coverage of common questions, ambiguity handling, source traceability and an explicit 'I do not know' path.
Human assistance	Summaries, classifications and suggested replies that Martin or the Team account must review before sending.
Autonomous response	Narrow allowed topics, confidence threshold, disclosure, hours, escalation and immediate stop/rollback conditions.
Safety	No unapproved legal, tax, immigration, banking, regulatory or security conclusions; privacy and prompt-injection controls.
Evaluation	Versioned test questions measuring groundedness, accuracy, escalation precision, harmful failure and duplicate-response risk.

What does not belong

- The human owner remains accountable for the customer conversation and final send.
- Channel registration and deliverability remain infrastructure responsibilities.
- AI cannot turn an unapproved policy, draft offer, or stale website claim into an authoritative answer.

Minimum viable adoption

Start with	Add only after evidence	Do not do
Curate the existing knowledge base, add approved FAQs, assign source owners and build a test set. Then pilot drafting assistance with mandatory human review.	A narrow autonomous FAQ only after assisted answers consistently meet agreed quality and escalation thresholds across at least 100 resolved examples or a representative test set.	Do not allow autonomous advice, refunds/pricing exceptions, account changes, complaint resolution or any answer without an approved source.

Adoption trigger: begin when one repetitive, low-risk question consumes enough partner time to justify curation. The first useful AI feature should assist Martin and the Team account, not replace them.

PRIORITY 5 OF 5 | IMPACT SCORE 52/100

Real-time website engagement

Boundary: this category owns the anonymous or pre-contact website experience through qualified handoff into one contact and one GHIL conversation. After the thread is created, the Conversations model owns the response. Channel transport and AI behavior remain separate controls.

Current state

- Two email-chat widgets exist, both created automatically for A2P compliance rather than as a deliberate Blueprint sales or support experience.
- Their existence does not prove either widget is embedded, active or staffed on a public Blueprint page.
- No Conversation AI or Voice AI agent is visibly deployed.
- The current offer funnels are payment-to-thank-you shells and do not authorize new chat or onboarding behavior.

What belongs here

Journey element	Blueprint decision to define
Placement	Which pages, devices, visitor types and hours display the widget without obstructing forms, calendars or checkout?
Expectation	Is help truly live, delayed, email-based or AI-assisted? What does the visitor see after hours?
Qualification	What minimum information is necessary to understand intent without over-collecting personal data?
Booking handoff	When should the visitor receive an existing calendar option, and how is protected attribution preserved?
Conversation handoff	How does one interaction create one intended contact/thread with source context and an accountable owner?
Measurement	Track widget starts, qualified handoffs, bookings, abandonment, duplicate creation and downstream response time.

What does not belong

- Post-creation assignment and response targets belong to Conversations.
- Email/SMS/WhatsApp transport and compliance belong to channel infrastructure.
- AI knowledge, autonomy and safety belong to AI governance; Client Portal and files belong to post-sale delivery.

Minimum viable adoption

Start with	Add only after evidence	Do not do
Choose one page and one visitor intent that existing forms/calendars do not serve. Use a human-owned handoff and explicit hours/expectations.	Expand placement or add AI only after controlled tests prove identity integrity, attribution, queue assignment, response coverage and measurable conversion value.	Do not label email chat as live support, deploy on checkout by default, create duplicate contacts, or add volume before the Team account can own the resulting queue.

Adoption trigger: start only when Blueprint can name the visitor intent currently being lost, show meaningful traffic, and guarantee that the resulting conversation will receive an accountable response.