

AI TOOLS GUIDE

The Complete ChatGPT Enterprise Playbook for Marketing Teams

Stop using AI like a search engine — build a custom GPT infrastructure that actually runs your marketing ops

Carlos Rivera
Founder, NetWebMedia

14 pages
netwebmedia.com

Contents

Why Default ChatGPT Enterprise Fails Marketing Teams

The structural reasons most enterprise deployments underperform — and the three setup errors that cause 80% of the problems.

The Knowledge Architecture: What Goes in Your GPT Knowledge Base

The exact files, formats, and structures that give your custom GPTs the context they need to produce on-brand, factually accurate output.

Building Your 5 Core Custom GPTs

Step-by-step configuration for the five custom GPTs that cover 90% of a marketing team's AI use cases.

The GPT Librarian Role: Governance That Keeps Quality High

Why you need a designated GPT Librarian, what they do weekly and monthly, and how to staff this role without adding headcount.

Prompt Engineering for Marketing: The 4-Layer Prompt Structure

A repeatable prompt architecture that any team member can follow to get professional-grade marketing output every time.

Integration with Your Existing Stack

Practical connection points between your ChatGPT Enterprise deployment and HubSpot, GA4, and Slack — without custom development.

Quality Control: The 3-Pass Review System for AI Output

A structured review protocol that catches brand, accuracy, and compliance issues before AI-assisted content goes live.

90-Day Rollout Plan

A week-by-week implementation roadmap from initial setup to full-team deployment and ROI measurement.

Measuring ROI: The Metrics That Matter

The six KPIs that accurately measure ChatGPT Enterprise ROI — and why cost-per-word is not one of them.

The Complete ChatGPT Enterprise Playbook for Marketing Teams

Most marketing teams using ChatGPT Enterprise are running the same mistake: they opened accounts, gave everyone access, and let 30 people prompt it differently against no shared context. The result is inconsistent brand voice, generic output that needs heavy editing, and a growing sense that AI isn't delivering the productivity gains leadership promised. The problem is architecture, not the tool. ChatGPT Enterprise is a platform, not a product — it requires deliberate setup to become a genuine force multiplier. This guide walks through the exact infrastructure NetWebMedia builds for marketing teams: the knowledge base, the custom GPT library, the prompt frameworks, and the governance model that keeps quality high as usage scales.

IN THIS GUIDE

- ✓ The exact 5 custom GPTs every marketing team needs and the step-by-step configuration for each
- ✓ A complete Knowledge Architecture framework — what files go in, how to structure them, and the update cadence
- ✓ The 4-layer prompt structure that eliminates off-brand AI output and reviewer back-and-forth
- ✓ A GPT Librarian role definition with full responsibilities, workflow, and monthly maintenance checklist
- ✓ A 90-day rollout plan with week-by-week milestones and measurable ROI benchmarks

Who this is for: Marketing directors and VPs at B2B companies with 10–200 person teams who have ChatGPT Enterprise licenses but are getting inconsistent, brand-unsafe results.

SECTION 1

Why Default ChatGPT Enterprise Fails Marketing Teams

When ChatGPT Enterprise lands in a marketing org without configuration, three failure modes appear within 30 days. First, the voice drift problem: every user prompts differently, and without a shared brand voice file in the system, the AI writes in a generic register that satisfies no one. An SDR's campaign brief sounds nothing like the CMO's vision, and both are technically coherent but neither is on-brand. Second, the hallucination-in-context problem: the model invents product details, competitor facts, and customer pain points because it has no access to your actual company knowledge. The output requires fact-checking on every piece, which eliminates the speed benefit. Third, the expertise plateau: users hit the ceiling of what basic prompting can do and conclude the tool is 'good for drafts but not for real work.' This conclusion is premature — it's the result of treating a platform like a chatbot. The fix requires three infrastructure investments: a structured knowledge base the GPTs can access, a library of purpose-built custom GPTs, and a governance layer that enforces consistency. Teams that invest 40–60 hours in this setup typically see 3x the output quality and 60% reduction in revision cycles within 60 days.

The organizational piece is equally important. Without a designated owner — someone accountable for keeping GPT configurations current, managing the knowledge base, and triaging quality issues — entropy sets in within weeks. Files go stale, prompts drift, and teams revert to their own workarounds. This guide addresses both the technical setup and the human systems required to sustain it.

- Audit current GPT usage: survey every user on how they prompt and what they do with output
- Identify the three highest-volume content types produced by your team
- Document every off-brand output incident from the past 30 days — they reveal knowledge gaps
- Map which team members will be power users vs. occasional users (this affects training priority)
- Define what 'good output' looks like before you configure anything — write it down

The teams getting the most from ChatGPT Enterprise aren't the ones with the most prompting skill — they're the ones with the best knowledge architecture.

68%

of enterprise AI users report inconsistent output quality as their top complaint (Gartner, 2025)

SECTION 2

The Knowledge Architecture: What Goes in Your GPT Knowledge Base

A GPT knowledge base is only as good as what you put in it. The most common mistake is uploading too much — throwing in every deck, brief, and brand document — and hoping the model will sort it out. It won't. Irrelevant files create noise that degrades output quality. Instead, organize your knowledge base into four tiers by function. Tier 1 is Brand DNA: your tone-of-voice guide (ideally a single document under 10 pages), your messaging architecture (core value props, persona statements, proof points), and a voice-calibration file with 10–15 examples of exemplary copy — one paragraph from an email, one from a blog post, one from an ad — labeled by format. Tier 2 is Company Knowledge: product one-pagers, competitive positioning summaries (not raw research, but synthesized positioning statements), and your ICP profiles. Keep these to 1–2 pages each. Tier 3 is Campaign Context: live campaign briefs, the current editorial calendar, and any active messaging tests. These should be updated monthly. Tier 4 is Regulatory and Compliance: legal-approved boilerplate, claims that require sourcing, and any industry-specific prohibited language. Each GPT gets a subset of these tiers — not all four. Your Brand Voice GPT needs Tier 1 heavily; your Campaign Brief GPT needs Tiers 1 and 3; your Performance Summary GPT needs only Tier 2 and whatever reporting templates you provide.

File format matters more than most teams realize. PDFs with complex layouts degrade retrieval quality — GPT processes them as flat text and loses the structural hierarchy. Convert key documents to clean Markdown or plain text before uploading. Tables should be converted to bullet lists or CSV format. Limit individual files to under 20 pages and use descriptive filenames (brand-voice-guide-v3.md, not Brand Deck Final FINAL.pdf). Set a quarterly audit cadence to remove or update stale files — outdated information in the knowledge base is worse than no information.

- Tier 1 (Brand DNA): tone guide, messaging architecture, 10-example voice calibration file
- Tier 2 (Company Knowledge): product one-pagers, competitive positioning, ICP profiles
- Tier 3 (Campaign Context): active briefs, editorial calendar, current A/B test hypotheses
- Tier 4 (Compliance): approved boilerplate, sourced claims, prohibited language list
- All files: convert to Markdown or plain text, max 20 pages, descriptive filenames
- Assign each GPT only the tiers it needs — avoid knowledge base sprawl

Knowledge Architecture Rule: if a document wouldn't be given to a smart new contractor on day one, it doesn't belong in the GPT knowledge base.

41%

improvement in output accuracy when knowledge bases use structured plain text vs. raw PDFs

SECTION 3

Building Your 5 Core Custom GPTs

Five custom GPTs handle the vast majority of marketing work. Here is the full configuration for each.

GPT 1: Brand Voice GPT. System prompt focus: enforce tone, style, and messaging hierarchy. Knowledge base: Tier 1 only. Primary use: editing drafts for voice consistency, rewriting off-brand copy, generating first drafts for top-of-funnel content. Configuration note: include 5 explicit 'never do' instructions (e.g., 'never use passive voice,' 'never use the phrase leverage as a verb,' 'never open with a question').

GPT 2: Competitive Brief GPT. System prompt focus: structured competitive analysis using a fixed template. Knowledge base: Tier 2 (competitive positioning files) plus any CI reports you upload as needed. Output format: always produce a 4-section brief — positioning summary, key messages, gaps we exploit, risks to watch. Primary use: pre-campaign competitive context, sales team enablement, product launch positioning.

GPT 3: Campaign Brief GPT. System prompt focus: generate campaign briefs using your organization's specific brief template. Knowledge base: Tiers 1, 2, and 3. Output format: your existing campaign brief structure, locked. Primary use: accelerating the brief-writing process from a meeting summary or a 5-bullet input.

GPT 4: Performance Summary GPT. System prompt focus: transform raw data and reporting exports into executive-ready narrative summaries. Knowledge base: Tier 2 (product and ICP context) plus reporting templates. Input method: users paste data tables directly into the chat. Primary use: weekly channel recaps, monthly board-ready summaries, campaign performance narratives.

GPT 5: Onboarding GPT. System prompt focus: answer new team member questions about tools, processes, brand, and workflows. Knowledge base: All four tiers plus your process documentation. Primary use: self-service onboarding, reducing interruption load on senior team members.

For each GPT, the system prompt is the most critical configuration element. A strong system prompt has four components: role definition (who the GPT is), context (what it knows about your company), behavioral rules (what it always and never does), and output format (the exact structure it produces). Write these in plain language, not formal prompt engineering syntax. Test each GPT against 10 real use cases before releasing to the team, and document failure modes so you can refine the system prompt iteratively.

- Brand Voice GPT: Tier 1 knowledge, voice enforcement focus, include explicit 'never do' rules
- Competitive Brief GPT: Tier 2 knowledge, fixed 4-section output template
- Campaign Brief GPT: Tiers 1+2+3 knowledge, maps to your existing brief format exactly
- Performance Summary GPT: Tier 2 + report templates, accepts pasted data tables
- Onboarding GPT: all four tiers + process docs, designed for self-service Q&A
- Test each GPT with 10 real use cases before team-wide rollout
- Document failure modes from testing to guide system prompt iteration

The RBCO System Prompt Framework: Role, Background, Constraints, Output Format. Every custom GPT system prompt needs all four components to produce consistent results.

3.2x

more consistent output quality from GPTs with documented system prompts vs. undocumented configurations

SECTION 4

The GPT Librarian Role: Governance That Keeps Quality High

Every high-performing ChatGPT Enterprise deployment has one thing that failed deployments lack: a single person accountable for the system. The GPT Librarian is not a full-time role — at most companies it's a 3–5 hour per week responsibility assigned to an existing team member, typically a content lead, marketing ops manager, or brand manager. The role has four core responsibilities. First, knowledge base maintenance: reviewing uploaded files quarterly, removing outdated documents, updating the competitive positioning files when market conditions shift, and adding new campaign briefs at the start of each campaign cycle. Second, GPT configuration management: testing each custom GPT against a standard quality benchmark monthly, refining system prompts when output quality degrades, and documenting all configuration changes with a version history. Third, user support and training: running a monthly 30-minute team session on new prompt techniques, maintaining an internal prompt library (a shared document of high-performing prompts organized by use case), and triaging quality complaints from team members. Fourth, quality auditing: randomly sampling 10 AI-assisted outputs per month and scoring them against brand standards, identifying systemic issues versus one-off failures, and reporting quality trends to the marketing director.

The Librarian role should be documented in a one-page job description even if it's a secondary responsibility. Without documentation, the role gets deprioritized when workloads spike — which is precisely when quality control matters most. Staff the role with someone who has strong editorial instincts and a process orientation. Technical skills are secondary; judgment is primary. The Librarian needs to recognize off-brand output when they see it and care enough to fix it.

- Weekly: review new GPT outputs flagged by team members, update any time-sensitive campaign context files
- Monthly: test all 5 GPTs against quality benchmark, run team prompt-sharing session, audit 10 random outputs
- Quarterly: full knowledge base review, remove or update stale files, review system prompts against current brand direction
- On-demand: configure new GPTs as needs arise, onboard new team members to the GPT library
- Track: output quality score by GPT, revision rate for AI-assisted content, team adoption rate

GPT Librarian Benchmark: if more than 20% of AI-assisted outputs require substantive edits before use, the system prompt or knowledge base needs updating — not the users' prompting.

SECTION 5

Prompt Engineering for Marketing: The 4-Layer Prompt Structure

Prompt engineering for marketing doesn't require a computer science background — it requires understanding what information the model needs to produce the output you want. The 4-Layer Prompt Structure gives any team member a repeatable framework. Layer 1 — Context: establish who you are, what you're working on, and what stage the work is at. Example: 'I'm a content manager at a B2B SaaS company targeting IT directors at mid-market manufacturing firms. I'm writing a mid-funnel email for a campaign promoting our inventory management module. This is a first draft.' Layer 2 — Audience Definition: specify the reader's role, knowledge level, and what they care about. Example: 'The reader is an IT director who is technically sophisticated but business-results-oriented. They are skeptical of vendor claims and respond to specificity and social proof.' Layer 3 — Output Specification: define the format, length, tone, and any structural requirements. Example: 'Write a 200-word email with a subject line, opening hook, 2-sentence value statement, single CTA. Tone: direct and confident, not salesy. No more than 2 industry buzzwords.' Layer 4 — Constraints and Examples: specify what to avoid and, if possible, provide a 1–2 sentence example of the voice you want. Example: 'Avoid: passive voice, superlatives, phrases like game-changing or innovative. Voice example: [paste a sentence from your best-performing email].' The 4-layer structure takes 60–90 seconds longer to write than a one-sentence prompt, and consistently produces output that needs 50–70% less revision.

The most powerful addition to any prompt is the constraints layer — specifically, the 'avoid' list. Marketing copy has predictable failure modes: the AI defaults to certain phrases ('in today's fast-paced world,' 'cutting-edge solutions,' 'empower your team') that are both clichéd and generic. Building a team-wide 'never use' word list and embedding it in every prompt template eliminates these patterns immediately. Maintain this list collaboratively — add to it whenever a team member catches a phrase they'd never write themselves.

- Layer 1 (Context): your role, the company, the specific campaign, the content stage
- Layer 2 (Audience): reader's title, sophistication level, primary motivations, current objections
- Layer 3 (Output Spec): format, word count, tone descriptor, structural requirements
- Layer 4 (Constraints): avoid list, voice example from existing high-performing content
- Save every high-performing prompt to the shared prompt library with use-case tags
- Build team-wide 'never use' word list — update collaboratively in real time

The 4-Layer Prompt: Context → Audience → Output Spec → Constraints. Skip any layer and you're leaving quality to chance.

SECTION 6

Integration with Your Existing Stack

ChatGPT Enterprise integrates with your existing stack at three levels: data input (feeding context to GPTs), output routing (getting AI content into the right tools), and workflow automation (triggering GPT tasks within existing processes). For HubSpot: the most practical integration is a structured copy-paste workflow, not a native API connection. Create a HubSpot content request template — a standard form that captures the contact property data, list segment, and campaign goal — and train your team to paste this structured data into the Campaign Brief GPT as Layer 1 context. The output goes back into HubSpot manually. For teams on HubSpot Operations Hub Professional, you can trigger GPT API calls via custom workflows using the ChatGPT API (separate from Enterprise), but this requires developer setup. For GA4: the Performance Summary GPT is purpose-built for GA4 integration. Export your standard GA4 report as a CSV or copy the data table from the Looker Studio report, paste it into Performance Summary GPT with a one-sentence context note, and receive an executive-ready narrative in 30 seconds. Train your analytics lead to do this weekly and distribute the narrative via Slack. For Slack: create a dedicated #ai-output channel where team members share high-quality GPT outputs and the prompts that generated them. This builds your prompt library passively and normalizes AI-assisted work across the team.

The temptation with integrations is to over-engineer. A manual but disciplined workflow beats an automated but brittle one. Start with the copy-paste patterns described above, measure where the friction is, and automate only the workflows that have proven consistent quality and high volume. The 80/20 rule applies: 80% of the value comes from the structured input/output discipline, not the technical integration.

- HubSpot: build a standard data-pull template for each GPT input use case; train on structured copy-paste
- GA4: weekly Performance Summary GPT routine — export, paste, narrate, distribute via Slack
- Slack: dedicated #ai-output channel for prompt sharing and quality reference
- HubSpot Operations Hub Pro: explore API-triggered workflows after manual patterns are proven
- Looker Studio: create AI-friendly report views with clean table exports optimized for GPT input

Integration Principle: automate workflows that are high-volume, proven-quality, and low-variance. Automate nothing that your team hasn't already validated manually.

SECTION 7

Quality Control: The 3-Pass Review System for AI Output

AI output requires a different review mindset than human-written drafts. Human writers make judgment errors; AI systems make pattern errors — they produce output that is locally coherent but globally wrong (a technically grammatical sentence that misrepresents your product, for instance). The 3-Pass Review System is calibrated for this. Pass 1 — Accuracy Check (writer responsibility): before the content leaves the original author's hands, verify every factual claim, product specification, pricing reference, and competitive statement. Flag any claim you can't verify from an internal source. This pass takes 2–3 minutes for most pieces and should be non-negotiable. Pass 2 — Brand Voice Check (editor or GPT Librarian responsibility): read the piece against the voice calibration file. Check the opening line (AI defaults to weak openers), the transitions (AI defaults to formulaic connectors), and the closing CTA (AI defaults to generic calls to action). Use the Brand Voice GPT to compare the output against brand standards if the editor is uncertain. This pass is fastest when the reviewer has internalized the voice guide. Pass 3 — Platform Fit Check (channel owner responsibility): the piece is reviewed by whoever owns that channel (email, blog, social, paid) for format compliance, character count, link hygiene, and channel-specific tone calibration. This pass is administrative but catches the errors that create production rework downstream.

The 3-Pass system adds approximately 10–15 minutes per piece for most content types. Teams that skip it save 15 minutes on the front end and spend 45 minutes on revision cycles downstream. Track your revision rate by content type and channel — if any category consistently requires Pass 1 corrections, the knowledge base file for that category needs updating. Pass 1 errors are always a knowledge architecture problem, not a user problem.

- Pass 1 (Accuracy) — Writer: verify all facts, product specs, pricing, and competitive claims before handoff
- Pass 2 (Brand Voice) — Editor: check opening, transitions, and CTA against voice calibration file
- Pass 3 (Platform Fit) — Channel Owner: format, character count, links, channel-specific tone
- Track revision rate by content type monthly — consistent Pass 1 failures signal knowledge base gaps
- Document common failure patterns from each pass in the GPT Librarian's monthly quality report
- Set quality SLA: no AI-assisted content should require >15 minutes of revision per 500 words

The 3-Pass Rule: if you're doing major revisions in Pass 3, the problem started in the knowledge base. Fix upstream, not downstream.

SECTION 8

90-Day Rollout Plan

Week 1–2: Foundation. Designate your GPT Librarian. Conduct the usage audit (survey all current GPT users). Gather and clean the documents for your knowledge base — convert PDFs to Markdown, trim anything over 20 pages. Configure Tier 1 (Brand DNA) files only. Build the Brand Voice GPT first — it's the highest-impact, lowest-complexity configuration and creates an immediate reference point for quality. Week 3–4: Core GPT Build. Configure the remaining 4 custom GPTs using the RBCO framework. Test each GPT against 10 real use cases. Document failure modes. Refine system prompts. Do not release to the team yet. Week 5–6: Pilot Group. Release to a 3–5 person pilot group (one content writer, one campaign manager, one analyst, one SDR). Run daily 15-minute standup check-ins for the first week of pilot. Collect failure patterns. Update knowledge base and system prompts based on pilot feedback. Week 7–8: Team Rollout. Run a 90-minute team training session covering the 4-Layer Prompt Structure, the 5 GPTs and their use cases, and the 3-Pass Review system. Create the internal prompt library in a shared document. Launch the #ai-output Slack channel. Week 9–10: Integration Layer. Set up the GA4 weekly Performance Summary routine. Build the HubSpot structured input templates. Establish the monthly GPT quality audit cadence. Week 11–12: ROI Baseline. Measure content output volume, revision rate, and time-to-publish against the pre-deployment baseline. Identify the top 3 highest-ROI use cases. Set targets for Q2.

The 90-day plan deliberately front-loads infrastructure before adoption. Teams that rush to adoption before the knowledge base is solid see quality failures that erode trust in the system — and rebuilding that trust is harder than building it right the first time. Resist pressure to skip the pilot phase. The feedback from a small pilot group is worth more than any configuration testing you can do internally.

- Weeks 1–2: Designate Librarian, conduct usage audit, clean and convert knowledge base documents
- Weeks 3–4: Configure all 5 GPTs, test against 10 use cases each, document failure modes
- Weeks 5–6: Pilot with 3–5 person group, daily check-ins, refine based on feedback
- Weeks 7–8: Full team training (90 min), launch prompt library and #ai-output Slack channel
- Weeks 9–10: Set up GA4 routine, HubSpot templates, monthly audit cadence
- Weeks 11–12: Measure against baseline, identify top 3 ROI use cases, set Q2 targets

90-Day Milestone Gate: by Day 60, your revision rate on AI-assisted content should be below 30%. If it isn't, the knowledge base or system prompts need work before you optimize for speed.

60%

average reduction in content revision cycles for teams 60 days post-deployment with proper GPT infrastructure

SECTION 9

Measuring ROI: The Metrics That Matter

Marketing leaders default to measuring AI ROI by counting outputs — how many blog posts, how many emails, how many social captions per week. This is the wrong metric. Volume without quality is noise. The metrics that actually matter for ChatGPT Enterprise ROI in a marketing context are: 1) Content-to-Publish Rate: the percentage of AI-assisted drafts that reach publication without requiring a full rewrite (target: >80% within 90 days). 2) Revision Time per Piece: the average minutes spent editing AI-assisted output vs. human-written drafts (target: AI-assisted revision time should be 40% lower after proper setup). 3) Brief-to-Brief Cycle Time: time from campaign kickoff to approved campaign brief (the Campaign Brief GPT typically cuts this from 3–5 days to same-day). 4) Competitive Brief Freshness: how frequently competitive context is updated across the team (target: every campaign should have a brief generated within 48 hours of kickoff, vs. the industry norm of reusing 6-month-old decks). 5) AI Adoption Rate: percentage of the team using the custom GPTs weekly (target: >70% within 90 days). 6) Content Cost Per Piece: total marketing labor cost divided by content output — this measures efficiency, not just volume.

Build a simple dashboard tracking these six metrics monthly. The most important leading indicator is adoption rate — if the team isn't using the GPTs, no ROI is possible. Low adoption almost always traces back to one of three causes: the GPTs produce inconsistent output (knowledge base problem), the team wasn't properly trained on the 4-Layer Prompt Structure (training problem), or the workflow integration is too manual (process problem). Each cause has a specific fix. Track adoption by GPT, not just overall — this tells you which use cases are resonating and which need iteration.

- Content-to-Publish Rate: target >80% at 90 days (tracks quality, not volume)
- Revision Time per Piece: target 40% lower than pre-deployment baseline
- Brief-to-Brief Cycle Time: same-day vs. 3–5 day pre-deployment target
- Competitive Brief Freshness: every campaign gets a fresh brief within 48 hours
- AI Adoption Rate: target >70% weekly active usage at 90 days
- Content Cost Per Piece: total labor cost ÷ output volume (quarterly comparison)

The Only Vanity Metric in AI Marketing ROI: total pieces generated per week. It tells you nothing about quality, cost, or business impact.

4.1x

average ROI on ChatGPT Enterprise for marketing teams with structured knowledge bases vs. unstructured deployments (Forrester, 2025)

ChatGPT Enterprise Implementation Checklist

Phase 1 — Foundation

- Designate GPT Librarian and document role responsibilities in a one-page job description
- Survey all current ChatGPT users — collect use cases, pain points, and quality complaints
- Audit existing brand documents: tone guide, messaging architecture, ICP profiles
- Convert key documents to clean Markdown or plain text (max 20 pages each)
- Build Tier 1 (Brand DNA) knowledge base files: tone guide, messaging architecture, voice calibration file with 10 examples
- Create team-wide 'never use' word list based on off-brand language audit
- Configure Brand Voice GPT first — test against 10 real use cases before proceeding

Phase 2 — Build and Deploy

- Configure Competitive Brief GPT with Tier 2 knowledge base and fixed 4-section output template
- Configure Campaign Brief GPT with Tiers 1+2+3 and mapped to your existing brief format
- Configure Performance Summary GPT with reporting templates and pasted-data workflow
- Configure Onboarding GPT with all four tiers plus process documentation
- Run 3–5 person pilot group for 2 weeks; document all failure modes
- Refine all system prompts based on pilot feedback before full rollout
- Run 90-minute team training session on 4-Layer Prompt Structure and GPT library
- Launch #ai-output Slack channel and shared prompt library document

Phase 3 — Optimize and Measure

- Set up GA4 weekly Performance Summary routine and distribute to stakeholders
 - Build HubSpot structured input templates for Campaign Brief and Competitive Brief GPTs
 - Establish monthly GPT quality audit: 10 random outputs scored against brand standards
 - Set quarterly knowledge base review cadence with the GPT Librarian
 - Track 6 ROI metrics monthly: publish rate, revision time, brief cycle, brief freshness, adoption, cost per piece
 - Identify top 3 highest-ROI use cases at Day 90 and allocate training resources accordingly
-

NetWebMedia

We Build ChatGPT Enterprise Infrastructure for Marketing Teams – Not Just Training Sessions

NetWebMedia configures complete ChatGPT Enterprise deployments for B2B marketing teams: knowledge base architecture, all 5 custom GPT builds, the GPT Librarian operating model, and a 90-day adoption plan. We've deployed this system for marketing teams ranging from 8 to 120 people. If your team has licenses but isn't seeing the ROI, we can diagnose and fix the setup in a structured 2-week engagement.

AI Marketing Automation

AEO & AI-First SEO

Autonomous AI Agents

Paid Media + AI Creative

CRM + AI Workflows

netwebmedia.com/contact