

WORKING PAPER

# Applying Compounding Growth Loops To Awareness

Marketing isn't dead, but the old playbook is.

*By Jim Hopkins, Revenue Architect, Winning by Design*

**Abstract:** The traditional B2B marketing playbook — a linear funnel with MQL-centered measurement and a handoff to sales — is breaking down. CAC is rising, conversion rates are declining, and AI-native companies are reaching revenue milestones in a fraction of the time at five to ten times the revenue per employee. The difference is architectural: a shift from scaling people to scaling systems, with the user reservoir at the center of gravity rather than the lead database. This paper applies the compounding growth loop framework to the awareness stage of B2B marketing. It defines a four-stage Awareness Loop — Spark, Engage, Invite, Share — and develops four distinct Spark types available to PLG and non-PLG companies alike. It proposes a loop-native metrics framework, identifies seven common failure modes, and offers a Crawl/Walk/Run implementation roadmap matched to organizational stage and readiness.

**Keywords:** Awareness Loop, Growth Loops, Compounding Growth, B2B Marketing, Go-to-Market Architecture, Revenue Architecture, AI-Native GTM, User Reservoir, Spark, k-factor, Product-Led Growth, Demand Generation

## Table of Contents

1. **Why Loops:** From Funnel to Compounding System.
2. **Spark:** The Four Types of Loop Entry.
3. **Engage, Invite, and Share:** Compounding the Loop.
4. **Metrics:** Measuring What Compounds.
5. **Failure Modes:** How Loops Quietly Break.
6. **Implementation:** Crawl, Walk, Run.

**Build the Loop:** From Campaigns to Systems.

# Chapter 1. Why Loops

Many companies are going all-in, laying off half their staff to move from a people-centric model to a more system-centric one that relies on AI and orchestrated workflows. The Marketing teams in these organizations are finding out that traditional lead-generation playbooks aren't working anymore. They're in need of a more system-oriented playbook. That's where loops come in. The Growth Institute recently released a paper: ***Foundations of Compounding Growth*** - where the authors outline how loops, in a GTM concept, behave and should be architected. It is recommended to read that paper for more insights on compounding growth loop concepts.

## Marketing's Mandate

*The main tasks for any B2B marketing team haven't changed. Find companies and people whose situations make them likely to be the next, best customers. Help them understand we know their pains and can deliver impact. Hand off to lead-development or sales when they're ready to have buying conversations.*

*What changes in the age of AI and growth loops is how we do this work: more system-oriented, data-driven, self-sustaining, and effective for both customers and marketers.*

## 1.1 Marketing's Old Playbook: The Linear Funnel

For decades, the traditional B2B growth playbook was modeled as a linear funnel - a one-way journey from awareness to purchase. Yes, we understood that prospects and customers took winding paths and came in and out of that funnel. But we never really had a sustainable way to model, measure, or best facilitate the journey, and attribution models never quite delivered. The linear playbook was built for an era where vendors controlled the flow of information, and buyers followed predictable paths because we kept tight security at the gates.

The handoff to sales is a "cliff," where marketing's job ends once the lead is captured. As digital channels expanded and buyers took more control, this model began to strain — rising CAC, declining conversions — as teams pushed harder on volume rather than value.

Table 1.1 The elements of B2B Marketing's Old Playbook

| FEATURE      | LINEAR, FUNNEL-CENTRIC PLAYBOOK                                                          |
|--------------|------------------------------------------------------------------------------------------|
| Structure    | Linear and finite; the relationship “ends” when we hand-off to LeadDev or Sales          |
| Primary Goal | Get their contact information - hand-raise on a call-to-action                           |
| Measurement  | Lead volume and top-of-funnel activity                                                   |
| Tactics      | Interruptions/blasts, gated content, paid channels that generate “leads”, and manual SEO |
| Targeting    | Broader branding, company or persona-based in paid channels, engagement scoring          |

## 1.2 Why AI-Natives Are Pulling Away

Cursor to \$500M ARR in 24 months. Midjourney to \$50M in 12. Perplexity past \$200M in 33. The curves aren't on the same scale as Salesforce, HubSpot, or Zendesk at comparable points — and the easy explanation (“they sell AI”) doesn't survive inspection, since plenty of AI-natives haven't grown this fast, and plenty of traditional B2B companies are shipping AI features without matching the trajectory.

The real explanation is architectural. Winning by Design's *Growth — In the Age of AI* unpacks it at length. For this paper, we take its core argument as foundation and build from two givens:

1. **SaaS-era companies scaled revenue by scaling people. AI-natives scale revenue by scaling systems, automated onboarding, telemetry-driven targeting, self-serve activation, community-generated content, workflow-embedded distribution.** *The revenue-per-employee gap reflects this. SaaS-era public companies converged at \$300K–\$500K per head. Leading AI-natives operate at five to ten times that, while growing faster.*
2. **The center of gravity has shifted. Traditional B2B manages leads → opportunities → customers. AI-natives operate on leads → users → customers.** *The user reservoir — free-tier users, trial users, API consumers, community members, evaluators — is the focus and center-of-gravity of the business.*

The buyer's journey has migrated into the user reservoir, where most traditional B2B companies have no presence. The funnel isn't broken — it's measuring a smaller, less representative slice of what's happening.

Readers who want the full comparison of SaaS-native vs. AI-native GTM architecture — RPE dynamics, the systems-vs-people shift, why bolting AI onto a people-scaled architecture doesn't close the gap — should start with *Growth — In the Age of AI*. The rest of this paper focuses on a narrower question: what does awareness look like when rebuilt on this architecture?

### 1.3 Anatomy of Awareness Growth Loops

In a system-centric model, awareness is not a top-down branding exercise, but a self-reinforcing engine. Unlike a linear funnel, an Awareness Loop is designed to compound; user actions and signals fuel the next wave of discovery, pulling the user into an active community or ecosystem rather than just a database. The center isn't a campaign. It's the user reservoir. The awareness loop has four general stages:




Figure 1.1 Anatomy of a System-Centric Awareness Growth Loop

Earlier versions of this framework used “Capture” – a familiar demand generation term – for the third stage. We moved to “Invite” deliberately. Capture is one-way and extractive. Invite is two-way: we extend the opportunity to elevate the journey, and the user chooses to step forward because the moment is right.

Table 1.2 The Four General Stages of the Awareness Loop

| FUNCTION      | OBJECTIVE                                                                                                                                    |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Spark</b>  | Generate targeted (ICP) attention through the right mechanism – a provocation, a product experience, a data-driven signal, or peer advocacy. |
| <b>Engage</b> | Move the user from spectator to participant – free tools, trials, content, community – and start generating signals the system can act on.   |
| <b>Invite</b> | Offer off-ramps, timed to moments of earned value, for users to start a conversation or deepen their relationship with the product.          |
| <b>Share</b>  | Equip users to advocate, collaborate, or co-create in ways that close the loop and generate the next Spark for new users.                    |

The stages run in sequence, but the loop isn't strictly linear. A user can enter at Share, skip ahead to Invite with pre-formed intent, or bounce from Engage back to Spark. The system's job is to meet them where they are and keep the loop turning.

## 1.4 One Architecture, Many Styles

Product categories differ. Pricing differs. Sales motions differ. The architecture is consistent. Each company runs all four stages, each stage feeds the next, and the user reservoir sits at the center-of-gravity of the business.

*Table 1.3 Examples of B2B companies Using the Awareness Loop architecture*

| COMPANY          | SPARK                                                                               | ENGAGE                                                            | INVITE                                                                           | SHARE                                                            |
|------------------|-------------------------------------------------------------------------------------|-------------------------------------------------------------------|----------------------------------------------------------------------------------|------------------------------------------------------------------|
| <b>Figma</b>     | <i>Design-community content; conference talks; reflecting designer frustrations</i> | <i>Free browser product; templates; public community</i>          | <i>Team-size / file-count thresholds trigger upgrade prompts</i>                 | <i>Every file is a URL; every URL is an invitation</i>           |
| <b>HubSpot</b>   | <i>Inbound POV; State of Marketing reports; free tools</i>                          | <i>Free CRM tier; HubSpot Academy</i>                             | <i>In-product upgrade prompts; sales outreach on intent-signaling free users</i> | <i>LinkedIn certifications; User Groups; INBOUND</i>             |
| <b>Gong</b>      | <i>Anonymized call-data insights; specific numeric claims</i>                       | <i>Content library; Gong Collective</i>                           | <i>Demo requests from content; ABM triggered by content signals</i>              | <i>Practitioners quoting Gong research in their own posts</i>    |
| <b>Atlassian</b> | <i>Long-running blog; agile/DevOps thought leadership; Team conference</i>          | <i>Free tiers of Jira/Confluence/Trello; Atlassian University</i> | <i>Usage-based upgrade prompts; enterprise sales on expansion</i>                | <i>Community forum; Marketplace; resume-ready certifications</i> |

## Chapter 2. Spark

Marketers may gravitate to spending most of their time and energy at Spark — for good reason. It's the stage most aligned with traditional marketing and has the richest set of creative options. The next four sections unpack the main Spark approaches.

### Spark Type 1: Messaging and Brand Perspective

The oldest Spark is a point of view. A brand takes a position the buyer recognizes — about their work, pain, future — and the recognition itself becomes the Spark. Done well, the audience self-selects. AI has made this Spark type harder and more valuable at the same time. Harder because anyone can produce infinite competent-looking content in minutes. More valuable because volume without resonance becomes even louder noise.

**AI has made content abundant.**

**Perspective is what remains scarce.**

Table 2.1 Messaging and Perspective Spark patterns and examples

| PATTERN                               | HOW IT WORKS                                                                     | EXAMPLES                                                                                                             |
|---------------------------------------|----------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------|
| <b>Knowledge Gap / Value Exchange</b> | Teach the ICP; give knowledge away to earn trust                                 | Stripe Press / Increment, Intercom playbooks, Drift conversational marketing, Winning by Design Revenue Architecture |
| <b>Data-backed Research Engine</b>    | Turn telemetry or aggregated usage into claims competitors can't credibly make   | Ramp spending reports, Okta "Businesses @ Work," Klaviyo benchmarks, a16z research                                   |
| <b>Provocation / Antics</b>           | Confrontational or contrarian messaging; generates disproportionate earned media | Perplexity "Poogole" ad, Artisan AI billboards, Anthropic Super Bowl anti-ad, Grok positioning                       |

Messaging and Brand Perspective can follow several patterns. Below we explore three.

- Knowledge Gap style messaging has a long B2B legacy. Stripe, Intercom, and Drift each built defining brands by publishing genuinely useful content long before asking for anything in return. At Winning by Design, we've tried to do this ourselves — the Bowtie, Revenue Architecture, and research briefs published openly.
- The most developed version turns proprietary data into a research function: specific numeric claims competitors cannot match and produce content that doubles as proof.

- Provocation or Antics is the opposite pattern: pattern-interrupt rather than teach. The campaigns above all cost a fraction of their earned-media value because they broke a pattern rather than followed one.

The limiting factor has flipped. In the old model it was production. In the new model it's perspective. AI is a perspective amplifier, not a perspective generator.

## Spark Type 2: *Product Experiences & Trials*

In the old model, marketing helped sell the product. In the new model, the product does most of the selling before a salesperson appears. The PLG version — free tier, low-friction sign-up, value in minutes — is well-documented. Slack, Notion, Canva, Loom, and most AI-natives run variations of the same PLG play. The interesting question — the one GTM leaders at legacy SaaS and enterprise B2B are quietly asking — is whether the pattern is available to companies whose product requires setup, data imports, configuration, and security reviews. It is. It just takes a bit of creativity and a shifted focus.

Below are Product-oriented Spark patterns across different business types and motions. The top two rows are standard PLG. The bottom three are where the opportunity lies for companies whose product can't be made free.

Table 2.2 Product Spark patterns and examples

| PATTERN                                             | HOW IT WORKS                                                                                                              | EXAMPLES                                                                          |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------|
| <b>Free Tier</b><br><i>Zero-Friction PLG Access</i> | <i>Permanent free tier or no-signup entry; the user experiences the product directly</i>                                  | <i>Slack, Notion, Canva, Loom, Figma, ChatGPT, Perplexity, Midjourney, Claude</i> |
| <b>Free Utility</b><br><i>Standalone Tool</i>       | <i>Narrow tool solves a specific problem; drives affinity and reservoir data</i>                                          | <i>Calendly, Zapier free tier, Ahrefs free tools</i>                              |
| <b>Interactive</b><br><i>Guided Experience</i>      | <i>Clickable product tour or sandbox; no setup, but hands-on feel; often personalized by role</i>                         | <i>ServiceNow demo site, Snowflake Quickstart, Workday scripted scenarios</i>     |
| <b>Managed Pilot</b><br><i>Proof-of-Concept</i>     | <i>Time-boxed, scoped engagement with white-glove support and measurable outcome; sometimes paid to ensure engagement</i> | <i>Palantir Foundry deployments, enterprise AI rollouts</i>                       |
| <b>Diagnostic on Prospect Data</b>                  | <i>Securely analyze the prospect's own data to show specific impact</i>                                                   | <i>Celonis Process Mining workshops, AI-driven assessments</i>                    |

All three of the non-PLG patterns share a principle: users experience the product's value with their own context before anyone explores long-term commitment.

### Spark Type 3: *Intent Data and Ecosystem Signals*

The first two Spark types are about what you put into the market. This third type is about what you see in the market — what's already being said, searched, posted, hired, and bought, and how you act on it fast enough to connect relevant messages or experiences and spark engagement. An intent-driven Spark is a personalized, well-timed touch triggered by a real signal: an account researching your category, a champion taking a new job, a competitor's customer complaining in a community thread, a tech stack showing a gap your product fills. The power comes from the match between signal and recipient.

**A well-timed touch, delivered when a signal says it's relevant,  
outperforms a clever campaign aimed at nobody in particular.**

The most-cited signal advantage of AI-natives is product telemetry. Companies without a free tier can feel locked out. They aren't. Four patterns do most of the work for non-PLG companies.

- **Intent Data:** *Whether third-party (6sense, Bombora, Intentsify), or first-party (engagement, website, product), intent data isn't new. What's new is that the signal can now trigger the work, not just rank it. Data, compute, and AI-assisted content have gotten cheap enough that detection-to-personalized-outreach happens in minutes, at scale, without an army of SDRs.*
- **Community and Ecosystem Mining:** *Tools like Common Room ingest signals from Slack communities, Reddit, LinkedIn, and industry forums and match them to people and accounts. A senior practitioner publicly complaining about a leading product is a stronger trigger than any intent score — and it's available without a free tier. Note the verb: mining. This Spark type monitors existing conversations and extracts signals. Participating in them (Spark Type 4), stimulating users in communities we've built (Engage), and engineering communities to reach new users (Share) are different motions.*
- **Champion Tracking:** *A customer's internal champion moves to a new company. UserGems, Sales Navigator, and ex-customer lists surface the event. The Spark isn't cold outreach; it's a warm follow-up to a known relationship getting them to re-engage. For enterprise B2B with a strong installed base, this is often the single highest-ROI signal they aren't fully exploiting.*
- **Composite Signals:** *Combine a primary intent feed (6sense, Bombora) with event data (hires, funding, M&A) and firmographic enrichment (Clay, ZoomInfo) to build composite scores no single signal could produce. A prospect on intent feeds for your category, with a new VP from a customer, during a funding cycle, is a different prospect than any signal alone suggests. Getting your arms around the full P.I.E. — Profile, Intent, and Engagement — is better than just a single slice.*

Non-PLG companies don't need product telemetry to run intent-driven Sparks. They need to treat the data they already have access to — community signals, champion movements, firmographic events — as signal infrastructure rather than one-off lists. AI is what makes it scalable.

### Spark Type 4: *Organic Word-of-Mouth and Peer Trust*

The fourth Spark type is the *injecting our voice* motion. Where the Intent/Ecosystem Spark mines communities for signals, this one is about earning the right to be part of them. When peers, analysts, creators, customers, and trusted third parties speak about us, or alongside us, those mentions become the highest-trust Sparks new users will ever encounter.

This is the oldest form of B2B awareness and the hardest to engineer. Trust transfers when one person tells another that something is worth their time. You can only earn the conditions under which that happens, then design for it to happen at scale.

**You can't manufacture trust. You can only earn the conditions under which it transfers.**

Table 2.3 Word-of-Mouth/Trust Spark patterns and examples

| PATTERN                                | HOW IT WORKS                                                                                                        | EXAMPLES                                                                                    |
|----------------------------------------|---------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------|
| <b>Customer-led Storytelling</b>       | <i>Customers tell their own stories on their own channels; we amplify rather than author</i>                        | <i>Klaviyo's customer-built playbooks on LinkedIn; Notion's customer build-along videos</i> |
| <b>Creator and Operator Influence</b>  | <i>Practitioners with audiences become long-term collaborators, not one-off sponsors</i>                            | <i>Lenny's Newsletter ecosystem; Pavilion faculty; Exit Five</i>                            |
| <b>Analyst and Category Leadership</b> | <i>Define a category — or reframe one — well enough that analysts and journalists adopt our language</i>            | <i>Drift naming "Conversational Marketing"; Gainsight defining "Customer Success"</i>       |
| <b>Peer Review Platforms</b>           | <i>Invest in being legitimately reviewable: easy to evaluate, easy to compare, easy for users to leave evidence</i> | <i>G2, Gartner Peer Insights, TrustRadius, Reddit</i>                                       |
| <b>Employee Advocacy</b>               | <i>Our team builds personal credibility in public; the company benefits because the team built audiences first</i>  | <i>Stripe engineering posts; HubSpot executive presence; Webflow design community</i>       |

AI can't generate the trust, but it can compress the cycle around it. It watches the long tail of mentions across podcasts, newsletters, communities, and video — surfaces that used to be invisible to marketing — matches high-trust mentions to prospects already in our system, and suggests the right asset for an advocate to share at the right moment. The trust itself still has to be real.

## Chapter 3. Engage, Invite and Share

The three stages that follow the Spark: Engage, Invite, and Share, decide whether the loop compounds. Engage turns attention into participation. Invite turns participation into permission. Share turns permission into the next Spark.

### 3.1 Engage

Engagement is the stage where the Spark becomes a relationship — but not yet a sales conversation. The user has shown up. Our job is no longer to be discovered. Our job is to be worth coming back to. This is the stage where we stimulate users to be active in communities we engineer.

The community may be the product itself, an owned space (Slack, Discord, forum, Academy), or a content surface where users interact with our perspective and with each other. We built it; we own the quality of what happens inside it. Engagement happens in three modes, often in combination. The strongest loops design for all three.

*Table 3.1 Engagement modes and examples*

| MODE                        | WHAT IT IS                                                                                       | EXAMPLES                                                                       |
|-----------------------------|--------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------|
| <b>Product Engagement</b>   | The user is actively using or exploring the product, free tier, trial, interactive demo, sandbox | Figma multiplayer files; Linear's free workspace; Navattic-style product tours |
| <b>Content Engagement</b>   | The user is consuming and applying our perspective via owned content surfaces                    | HubSpot Academy; Stripe's annual letter; Gong Labs research drops              |
| <b>Community Engagement</b> | The user is interacting with other users in spaces we've built and curate                        | dbt Labs Community; Notion Ambassadors; Webflow forums; Gainsight Pulse        |

The point isn't to log activity. It's to give the user reasons to come back, get better at something, and form an identity that includes us in it. A user who has built three Notion templates is a different kind of user than one who has read three Notion blog posts. Both are engaged; one is much closer to becoming an advocate.

AI changes engagement by changing what "personalized" can mean. In the SaaS-era, personalization meant "first name and industry." In the AI-native model, the next module surfaced, the next template recommended, the next community thread reflected back — all adapt in close to real time to what the user has done. Done well, this makes engagement feel less like marketing and more like a tool that's getting better at helping.

## 3.2 Invite

The critical junction in the loop is converting engagement into permission to deepen the relationship. The Invite is a permission moment, not a conversion moment. Ask before the user is ready, and we trigger a defensive response. Ask after the moment has passed, and we forfeit the relationship to whoever asks next. The Invite is not a traditional CTA. It's an offer the user is ready to receive.

### **The job of this stage isn't to ask. It's to know when to ask.**

Of course, ideally the user says — explicitly or near-explicitly — they have a job to be done that maps to our product. For example, when a user clicks on the “Upgrade” button in their trial account and inquires about an enterprise package, or a newsletter subscriber clicks on the “Contact Us” button to learn more about the company’s solutions. We can also make sure we track behavioral signals like usage, consumption, or participation milestones that correlate with intent, and provide personalized invites for a next step. For example, a trial user has invited two other users to collaborate, or a community member has posted multiple product-specific comments or questions.

The AI-native model lets us track and respond with invites or product prompts in close to real time. A reader hits a threshold, posts on a public forum, and just got promoted: the invite that follows can be matched to that combined signal in a way the SaaS-era couldn't.

Not every Invite is a sales call. The mistake legacy B2B teams make most often is offering only a demo request. The strongest loops offer a ladder that gets them closer to asking for a sales engagement.

*Table 3.2 Invite examples by level of commitment*

| DEPTH         | INVITE EXAMPLES                                                                 |
|---------------|---------------------------------------------------------------------------------|
| <b>Cold</b>   | <i>Newsletter subscribe, additional content, podcast follow, community join</i> |
| <b>Medium</b> | <i>Live workshop, calculator/assessment with custom output, paid event</i>      |
| <b>Hot</b>    | <i>Demo request, trial sign-up, “talk to sales,” paid pilot</i>                 |

The ladder matches the Invite with engagement depth. AI's contribution is signal synthesis: the team's job shifts from designing campaigns to designing the rules and guidelines for inviting.

### 3.3 Share

Share is the stage that closes the loop. A user who sparked and engaged either gives us another Spark, by bringing someone else in, or doesn't. The loop's compounding rate is determined here. The temptation is to treat sharing as a marketing campaign: refer-a-friend, give-\$10-get-\$10, swag for tweets. Those work in narrow consumer contexts and almost never in B2B. The patterns that actually generate B2B shares don't ask the user to do marketing for us. They make sharing the natural next step in what the user was already doing.

**The strongest share moments aren't favors users do for us.  
They're things users were going to do anyway, made easier.**

Table 3.3 Share patterns with examples

| PATTERN                      | HOW IT WORKS                                                                                                         | EXAMPLES                                                                                   |
|------------------------------|----------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------|
| <b>Collaborative</b>         | <i>The product is more useful with more people in it; inviting others is a feature, not an ask</i>                   | <i>Figma, Notion, Linear, Loom, Slack</i>                                                  |
| <b>Embedded</b>              | <i>Outputs of using the product carry our brand into someone else's workflow</i>                                     | <i>Calendly booking links; DocuSign; HubSpot meeting links; Stripe-powered checkout</i>    |
| <b>Output</b>                | <i>What the user produces is naturally shareable and reflects well on them</i>                                       | <i>Gong call snippets; Tableau dashboards; Canva designs; Hex notebooks</i>                |
| <b>Advocate</b>              | <i>A subset of users (roughly 1 in 4 in healthy programs) will speak publicly if given a structure to do it well</i> | <i>Notion Ambassadors; Webflow Experts; Salesforce MVPs; dbt Champions</i>                 |
| <b>Content Amplification</b> | <i>Users redistribute our perspective because sharing it makes them look smart</i>                                   | <i>Stripe annual letters; Gong Labs research; First Round Review; Lenny's case studies</i> |

Each share creates a probabilistic new Spark, and the system's k-factor (or virality) is the average number of new Sparks per active user. A k-factor above 1 means the loop runs without new paid acquisition. A k-factor of 0.3 — typical for healthy B2B loops — means we still need other Sparks, but each user contributes meaningfully to the next round. This stage is also where we engineer the communities we built for outward reach. The Notion forum produces shareable templates, public threads, and conference talks. The dbt Labs Community generates GitHub stars, blog posts, and conference proposals. Owned communities designed for engagement can be re-designed for outward propagation. That's a deliberate choice, what gets indexed, what gets templated, who gets elevated, not an accident. AI compounds Share by personalizing what gets shared, optimizing the share moment (timing, channel, framing), and tracing propagation across surfaces that used to be opaque to each other.

## Chapter 4. Metrics

The funnel has its own dashboard, and most marketing teams still ship it to the board. MQLs, SQLs, pipeline created, cost per lead, channel attribution. None of those metrics are wrong, exactly. They're measuring a system that ignores all the activity and loop response happening in-between those larger milestones.

**If your dashboard still leads with MQLs, you're not measuring the loop.**

A loop dashboard measures three things the funnel never had to: how many Sparks the system produces, how engagement deepens, and how efficiently each turn of the loop creates the next.

Table 4.1 Loop Metrics

| CATEGORY                          | WHAT IT MEASURES                                                                                  | SAMPLE METRICS                                                                                                    |
|-----------------------------------|---------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------|
| <b>Spark Metrics</b>              | <i>Where new users are entering and which Sparks are producing them (ideally filtered by ICP)</i> | <i>Spark Volume by Type (New Users); First-touch Context; Signal-to-Spark Conversion</i>                          |
| <b>Engagement Metrics</b>         | <i>How depth of relationship is changing</i>                                                      | <i>Returning User Rate; Depth-of-Use / Engagement; Time-to-Value; Community Participation; Journey Completion</i> |
| <b>Loop Velocity / Multiplier</b> | <i>How efficiently each turn produces the next</i>                                                | <i>Sparks per Active User (k-factor); Spark-to-Invite-to-Customer time; Share Rate; ARR per Spark</i>             |

Spark and Engagement metrics tell us whether the loop's inputs are healthy. Velocity tells us whether it's compounding. A team that ships a great content series and sees Spark volume rise but k-factor flat is producing engagement, not loop. A team that sees k-factor climb without Spark volume rising is at risk of running out of fuel even as conversion improves.

Two metrics tend to organize the rest of the dashboard around them. Sparks per Active User — the loop's k-factor — typically runs 0.2–0.5 in healthy B2B systems. The slope matters more than the level: 0.3 climbing 0.02 per quarter is healthier than 0.5 sitting flat. ARR per Spark is the AI-native equivalent of CAC, framed around the loop's atomic input rather than channel spend. It tells us which Spark types deserve investment and which throw off engagement that doesn't compound to revenue.

Revenue data platforms (Common Room, Pocus, MadKudu, plus the AI layer on top of CRM) make most of this measurable natively. The harder requirement is leadership willingness to lead with these metrics instead of MQLs, and to be patient while the loop warms up.

## Chapter 5. Failure Modes

The Awareness Loop is more powerful than the funnel, but more failure-prone. The funnel's failure was usually obvious: leads dried up, pipeline coverage fell, the board asked questions. The loop's failure modes are subtler, often look like success on the wrong dashboard, and compound over quarters before they become visible.

### **A loop that's failing usually looks like one that's working, for about a quarter.**

Below are seven of the most common ways an awareness loop might show signs of failure:

- **The Invite Trap**

Marketing is rewarded for booked meetings, so the team asks for them earlier and more often. Invite volume rises; conversion and engagement quality fall. *Early signal:* Demo request volume rising while demo-to-opportunity rate falls. *Hedge:* Tie incentives to depth-of-engagement and downstream conversion, not to top-of-funnel ask volume. Build a soft/medium/hard ladder and report on the ratio.

- **Engagement Theater**

The team optimizes MAUs, time-on-site, and member counts that don't correspond to advocacy or revenue. *Early signal:* Engagement metrics rise while k-factor and ARR per Spark stay flat or fall. *Hedge:* Always pair an engagement metric with a downstream loop metric in the same chart.

- **Astroturfed Advocacy**

The team manufactures the appearance of word-of-mouth — bought reviews, paid endorsements that read as organic, scripted “authentic” employee posts. *Early signal:* Advocacy volume rises while engagement on advocacy content (saves, comments, shares-of-shares) does not. *Hedge:* Invest in conditions that produce real advocacy (customer outcomes, community quality, advocate enablement). Accept slower volume in exchange for compounding signal.

- **Channel Monoculture**

One Spark type — typically paid intent or one influencer relationship — dominates. When it falters, the loop collapses. *Early signal:* >60% of Sparks from one channel for two consecutive quarters. *Hedge:* Diversify across the four Spark types deliberately. Set a target distribution and report against it.

- **Signal Hoarding**

Data is captured but not acted on, because the team that owns the data and the team that owns the response work in different systems. *Early signal*: Signal-to-Spark conversion below benchmark while signal volume is healthy. *Hedge*: Put a single owner over the signal-to-action loop. Reduce systems-of-record. Audit signal-to-action latency monthly.

- **Loop Without Product**

Marketing engineers a beautiful Awareness Loop on top of a product the user doesn't actually love. Sparks happen, engagement happens, Invites happen — and the loop breaks because there's nothing satisfying to share. *Early signal*: k-factor below 0.1 even in cohorts that converted to paid. *Hedge*: Treat marketing's loop and product's loop as the same loop. Without product investment, the loop will look like it's working until it doesn't.

- **Over-Personalization**

The signal layer combines so much behavioral, intent, and identity data that the user feels surveilled rather than served. *Early signal*: Reply-rate to personalized invites falls quarter-over-quarter; community sentiment shifts negative. *Hedge*: Tune the personalization to be natural and legible. Tell the user why we're reaching out, what we know, and how we knew it. Default to lighter-weight invites when the signal is novel or sensitive.

The strongest meta-hedge is to operate the loop transparently with the people inside it. Acknowledge what we're testing. Ask the community what's working. Publish what's working — and what failed. The same AI tools, used opaquely, feel manipulative; used openly, they feel like attention to the user.

## Chapter 6. Implementation

The Awareness Loop is implementable at almost any scale. Most teams reading this paper won't ship a fully AI-native, four-Spark-type, instrumented loop in the next two quarters. Most shouldn't try. A small, well-built loop in production beats a large, well-architected one in design indefinitely.

### **A small loop in production beats a large loop in design indefinitely.**

The goal in all of this isn't to climb bands as fast as possible. It's to run the band you're in well before climbing.

### 6.1 Crawl

One full turn of the loop, end to end, with a small team - A weekly newsletter with a sharp POV, a single owned engagement surface, one Invite path, and one share pattern is a complete loop. It won't scale to enterprise GTM, but it's a loop, and it teaches the team how a loop behaves.

*Table 6.1 Awareness Loop implementation at the Crawl band*

| STAGE                 | CRAWL BUILD                                                                                                                                                                   |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Spark</b>          | One Spark type to start – typically Content/POV – with one consistent owner and a weekly cadence; expand and experiment if the initial Spark doesn't perform after 1-2 months |
| <b>Engage</b>         | One owned surface (newsletter, podcast, small community); deepen one user at a time                                                                                           |
| <b>Invite</b>         | One path, cold to medium; no demo asks until behavioral threshold is crossed                                                                                                  |
| <b>Share</b>          | Forward-to-a-friend or invite-a-colleague; instrumented enough to know it happened                                                                                            |
| <b>People</b>         | 1–3, often a founder or marketing lead doing most of it                                                                                                                       |
| <b>Budget</b>         | <\$10K/month tooling, content, and distribution; leadership attention is the real cost                                                                                        |
| <b>Infrastructure</b> | Email platform, light CRM, one community tool, one analytics tool                                                                                                             |

### 6.2 Walk

Multi-Spark with a real signal layer - The team adds a second and third Spark type and builds owned engagement surfaces to reinforce each other. Invite paths are tiered. Share patterns are deliberate. Loop metrics are reported alongside funnel metrics.

Table 6.2 Awareness Loop implementation at the Walk band

| STAGE                 | WALK BUILD                                                                                                                           |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| <b>Spark</b>          | 2–3 types concurrently; deliberate target mix (e.g., 40% content, 40% signal-led, 20% WoM)                                           |
| <b>Engage</b>         | Owned community + content + light product engagement (free tool, calculator, interactive demo)                                       |
| <b>Invite</b>         | Soft/medium/hard ladder with rules of relevance defining when each fires                                                             |
| <b>Share</b>          | At least two share patterns engineered                                                                                               |
| <b>People</b>         | 5–15 across marketing, with at least one signal/ops role and one community role                                                      |
| <b>Budget</b>         | \$50K–\$250K/month                                                                                                                   |
| <b>Infrastructure</b> | Modern CRM, intent platform, community platform, AI orchestration layer (Common Room / Pocus / MadKudu / Clay), basic loop analytics |

## 6.3 Run

The AI-native loop at scale — All four Spark types are instrumented. Engagement happens across product, content, and community surfaces, with users going deep and many becoming entrenched as advocates. Invites are continuous and matched to engagement depth. Share is engineered as part of the product, content, and community surfaces. Loop metrics lead the dashboard.

Table 6.3 Awareness Loop implementation at the Run band

| STAGE                 | RUN BUILD                                                                                                                                                |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Spark</b>          | All four types running with reported k-factor by type                                                                                                    |
| <b>Engage</b>         | Product + community + content surfaces are generating continuous engagement; offering increasing depth and exploration, and in some ways self-sustaining |
| <b>Invite</b>         | Adaptive ladder; signal synthesis combines behavioral, expressed, and triggered signals                                                                  |
| <b>Share</b>          | All five patterns engineered into product and community; share-attribution across surfaces                                                               |
| <b>People</b>         | 25+ across marketing, with embedded product, CS, and signal/data roles forming a true revenue org                                                        |
| <b>Budget</b>         | \$1M+/month, with the AI/data layer treated as core infrastructure                                                                                       |
| <b>Infrastructure</b> | Composable revenue stack with AI orchestration as the spine; signal-to-action latency in minutes                                                         |

The right band can be determined by business stage, maturity, and organizational readiness constraints. Pre-Series B: leadership attention — Crawl. Series B–D: signal capacity and team coordination — Walk. Mature companies with budget and customer base: Run, but only if leadership is willing to retire some funnel metrics and be patient while the loop builds.

The most common failure of this roadmap is jumping bands. A Crawl team buying Run-level infrastructure produces engagement theater. A Walk team asked for Run-level metrics fabricates them. Build the band you're in. Climb when the constraint changes.

## 6.4 Organizational Implications

The loop implies a different shape of marketing organization than the funnel. The funnel was a relay: demand-gen to MOps to SDRs to AEs. Each handoff was a measurement boundary. The loop doesn't have those boundaries in the same place. It rewards a different operating model.

**The funnel was a relay race.**

**The loop is a single team running a system.**

**Crawl Org:** The org *is* the loop. A founder, a marketing lead, and possibly a content or community hire (or AI tool) run every stage themselves. No handoffs, no measurement boundaries (the strength); bound by leadership attention (the weakness). Keep the loop small, ship one full turn end-to-end, and resist hiring specialists before there's enough loop to specialize on. The most useful first specialist is usually a content/POV-owner. The least useful is a demand-gen specialist whose job description came from a SaaS-era playbook.

**Walk Org:** Too large for one person; too small to fragment into traditional functions. A cross-functional pod of five to fifteen people — content, community, signal/ops, and a tight feedback loop into product and CS — operates against loop metrics, not channel metrics. The pod meets weekly against a loop dashboard, ships against it, reports against it. Consolidate the signal-to-action stack onto as few systems of record as possible, and hire someone whose explicit job is keeping that stack coherent.

**Run Org:** No longer really recognizable as traditional marketing. It's a revenue org with marketing as the loop owner, product as the loop's substrate, and CS as the loop's amplifier. Roles that used to be separate — demand-gen, content, community, lifecycle, RevOps, customer marketing — are roles inside a single team operating against a single dashboard. Smaller headcount than the equivalent SaaS-era marketing org; dramatically higher leverage per person. This is what the order-of-magnitude jump in revenue per employee looks like in practice.

The hardest structural change at the Run band is between marketing and the rest of the revenue org. Sales and CS leaders raised in the funnel era expect marketing to deliver MQLs to a target. Marketing leaders running a Run-band loop are measured on Sparks per active user, ARR per Spark, and loop velocity. Reconciling typically requires an executive-level decision to adjust MQL targets based on loop targets, and a corresponding rewrite of sales and CS comp plans. Most companies that fail at the Run band will fail here, not in the marketing org itself, because those teams will be exerting increased pressure to meet their unchanged targets.

## Build the Loop

The Awareness Loop isn't a rejection of the funnel. The funnel taught the disciplines of stages, attribution, and measuring what we ship — disciplines that carry forward. What changes is the *shape* of the system marketing is responsible for: not a one-way path from awareness to purchase, but a continuously running loop in which Sparks produce engagement, engagement earns the right to invite, invites convert into relationships, and relationships generate the next round of Sparks.

The shift from funnel to loop is happening at the same time as a deeper shift in GTM architecture: from people-scaled SaaS to systems-scaled AI-native. The two reinforce each other. The loop demands signal synthesis, real-time relevance, and continuous engagement at a scale the SaaS-era marketing org wasn't built to deliver. The AI-native stack makes those demands tractable. Companies that have made both shifts are running marketing at order-of-magnitude better revenue per employee.

The harder lesson — the one that takes longer to act on — is that the loop is not a campaign. It's a system the team operates and improves continuously, with compounding returns over years. The teams that win in the next decade will retire the campaign as the unit of work and adopt the loop in its place.

Pick a Spark. Engineer the engagement that follows. Earn the Invite. Build for the share. Measure what compounds. Run it again.

***Build the loop.***

---

## **APPENDIX**

## METHODOLOGICAL POSTURE

This paper is intentionally developed and released as a working paper. Its methodological aim is to explore, test, and refine a systems framework for understanding compounding growth through feedback loops. While grounded in empirical observation and established theory, the approach prioritizes sense-making, model formation, and shared language over presenting finalized, peer-validated conclusions. The framework is therefore expected to evolve through application, critique, and continued refinement across contexts.

## EMPIRICAL BASIS

The framework is grounded in more than fifty structured Go-To-Market Diagnostics (GTMDs) conducted between 2018 and 2024. These diagnostics span companies from early-stage startups (~\$4M ARR) to global enterprises exceeding \$6B ARR across B2B SaaS, PLG, enterprise software, infrastructure, fintech, cybersecurity, marketplaces, and hybrid AI-native models.

Each GTMD follows a consistent diagnostic structure—instrumenting acquisition, retention, and expansion mechanics; mapping conversion rates (CR1–CR8), velocity, and cost-to-serve; and identifying structural constraints across GTM motions. While qualitative in nature, the diagnostics are applied systematically and comparatively, enabling pattern recognition across scale, stage, and business model.

## LIMITATIONS AND BOUNDARY CONDITIONS

This framework describes the structural behavior of recurring-revenue growth systems that generate growth through coordinated acquisition, retention, and expansion over time. Its predictive power is strongest once organizations have moved beyond founder-led execution into multi-role, multi-signal go-to-market systems. It is not intended to predict outcomes in the following contexts:

- Pre-revenue or pre-product systems, where growth dynamics are dominated by exploration rather than reinforcement.
- Pure B2C viral distribution without monetization, where usage can compound independently of revenue stability.
- One-off or transactional business models, where value is realized at the point of sale and feedback loops do not materially reinvest into future growth.

These boundary conditions do not represent exceptions to the framework; rather, they are contexts in which the underlying system has not yet met the criteria for compounding behavior to emerge. Once revenue becomes recurring and growth depends on retained value rather than repeated acquisition, the structural constraints described in this paper apply consistently.

## GLOSSARY OF TERMS (ALPHABETICAL)

**ARR per Spark (p. 14):** The AI-native equivalent of CAC, framed around the loop's atomic input rather than channel spend. Tells the team which Spark types deserve investment and which throw off engagement that doesn't compound to revenue.

**Astroturfed Advocacy (p. 15):** A loop failure mode where the team manufactures the appearance of word-of-mouth — bought reviews, paid endorsements that read as organic, scripted "authentic" employee posts — instead of earning real advocacy.

**Awareness Loop (p. 4):** A self-reinforcing engine of awareness composed of four stages — Spark, Engage, Invite, Share — in which user actions and signals fuel the next wave of discovery. Replaces the linear funnel as the unit of work for system-centric marketing.

**Champion Tracking (p. 8):** An intent-driven Spark pattern that surfaces when a customer's internal champion moves to a new company, enabling warm follow-up to a known relationship rather than cold outreach.

**Channel Monoculture (p. 15):** A loop failure mode where one Spark type — typically paid intent or one influencer relationship — dominates Spark volume, leaving the loop exposed when that channel falters.

**Composite Signals (p. 8):** Intent-driven Spark pattern that combines a primary intent feed with event data (hires, funding, M&A) and firmographic enrichment to build composite scores no single signal could produce.

**Community and Ecosystem Mining (p. 8):** A Spark pattern in which tools ingest signals from Slack communities, Reddit, LinkedIn, and industry forums and match them to people and accounts, surfacing public complaints or category discussion as triggers for outreach.

**Crawl, Walk, Run (p. 16):** A three-band implementation roadmap for the Awareness Loop. Crawl runs one full turn end-to-end with a small team; Walk adds multiple Spark types and a real signal layer; Run operates an instrumented, AI-native loop at scale.

**Diagnostic on Prospect Data (p. 7):** A non-PLG product Spark pattern in which the company securely analyzes the prospect's own data to demonstrate specific impact before any commercial conversation.

**Engage (p. 4):** The second stage of the Awareness Loop, where the user moves from spectator to participant — through free tools, trials, content, or community — and starts generating signals the system can act on.

**Engagement Theater (p. 15):** A loop failure mode where the team optimizes MAUs, time-on-site, and member counts that don't correspond to advocacy or revenue.

**Free Tier / Zero-Friction PLG Access (p. 7):** A product Spark pattern offering a permanent free tier or no-signup entry, allowing the user to experience the product directly with no commercial gate.

**Free Utility / Standalone Tool (p. 7):** A product Spark pattern in which a narrow tool solves a specific problem, drives affinity, and generates reservoir data without requiring adoption of the core product.

**Growth Institute (p. 4):** Producer of *Foundations of Compounding Growth*, the underlying framework on which this paper builds.

**Guided Experience / Interactive (p. 7):** A product Spark pattern using a clickable product tour or sandbox — no setup, but hands-on feel — often personalized by role.

**Invite (p. 4):** The third stage of the Awareness Loop. A permission moment, not a conversion moment: an offer the user is ready to receive, timed to a moment of earned value. Replaces the older term "Capture."

**Invite Ladder (p. 11):** A tiered set of Invite paths — Cold, Medium, Hot — matched to engagement depth, ranging from newsletter subscribe to demo request or paid pilot.

**Invite Trap (p. 14):** A loop failure mode in which marketing is rewarded for booked meetings, so the team asks for them earlier and more often. Invite volume rises while conversion and engagement quality fall.

**k-factor / Sparks per Active User (p. 13):** The average number of new Sparks per active user. A k-factor above 1 means the loop runs without new paid acquisition; healthy B2B loops typically run 0.2–0.5.

**Knowledge Gap / Value Exchange (p. 6):** A messaging-and-perspective Spark pattern in which the brand teaches the ICP and gives knowledge away to earn trust.

**Linear Funnel (p. 4):** The legacy B2B growth model — a one-way journey from awareness to purchase with a handoff "cliff" between marketing and sales. Built for an era when vendors controlled the flow of information.

**Loop Velocity / Multiplier (p. 13):** A category of loop metrics that measures how efficiently each turn of the loop produces the next. Includes Sparks per Active User, Spark-to-Invite-to-Customer time, Share Rate, and ARR per Spark.

**Loop Without Product (p. 15):** A loop failure mode in which marketing engineers a beautiful Awareness Loop on top of a product the user doesn't actually love. Sparks happen, engagement happens, Invites happen — and the loop breaks because there's nothing satisfying to share.

**Managed Pilot / Proof-of-Concept (p. 7):** A non-PLG product Spark pattern offering a time-boxed, scoped engagement with white-glove support and a measurable outcome; sometimes paid to ensure engagement.

**Over-Personalization (p. 15):** A loop failure mode in which the signal layer combines so much behavioral, intent, and identity data that the user feels surveilled rather than served.

**P.I.E. — Profile, Intent, and Engagement (p. 8):** A composite framing of signal data that combines firmographic profile, intent feeds, and first-party engagement rather than relying on any single slice.

**Provocation / Antics (p. 6):** A messaging-and-perspective Spark pattern using confrontational or contrarian messaging to generate disproportionate earned media. The opposite of the Knowledge Gap pattern.

**Revenue per Employee (RPE) (p. 4):** A measure of GTM architecture productivity. SaaS-era public companies converged at \$300K–\$500K per head; leading AI-natives operate at five to ten times that.

**Share (p. 4):** The fourth stage of the Awareness Loop, in which users advocate, collaborate, or co-create in ways that close the loop and generate the next Spark for new users.

**Signal Hoarding (p. 15):** A loop failure mode in which data is captured but not acted on, because the team that owns the data and the team that owns the response work in different systems.

**Spark (p. 4):** The first stage of the Awareness Loop. The mechanism that generates targeted ICP attention — through a provocation, a product experience, a data-driven signal, or peer advocacy.

**Spark Type 1 — Messaging and Brand Perspective (p. 6):** Spark generated by a recognizable point of view about the buyer's work, pain, or future. The audience self-selects on recognition.

**Spark Type 2 — Product Experiences and Trials (p. 7):** Spark generated by letting the user experience the product's value with their own context before any commercial conversation.

**Spark Type 3 — Intent Data and Ecosystem Signals (p. 7):** Spark generated by acting on what's already being said, searched, posted, hired, and bought in the market — a personalized, well-timed touch triggered by a real signal.

**Spark Type 4 — Organic Word-of-Mouth and Peer Trust (p. 9):** Spark generated when peers, analysts, creators, customers, and trusted third parties speak about the company or alongside it. The highest-trust Spark and the hardest to engineer.

**System-Centric Model (p. 4):** A GTM architecture in which revenue scales by scaling systems — automated onboarding, telemetry-driven targeting, self-serve activation, community-generated content, workflow-embedded distribution — rather than by scaling people.

**User Reservoir (p. 3):** The pool of free-tier users, trial users, API consumers, community members, and evaluators that sits at the center of gravity of an AI-native business. The buyer's journey has migrated here, where most traditional B2B companies have no presence.

## INDEX OF TABLES AND FIGURES

**Figure 1.1. Anatomy of a system-centric Awareness Growth Loop (p. 4).** Illustrates the four stages of the Awareness Loop — Spark, Engage, Invite, Share — and the way each stage feeds the next, with the user reservoir at the center of gravity.

**Table 1.1. The elements of B2B Marketing's Old Playbook (p. 3).** Catalogs the structural features of the linear, funnel-centric playbook across structure, primary goal, measurement, tactics, and targeting.

**Table 1.2. Functions and objectives of the four Awareness Loop stages (p. 4).** Defines the objective of each loop stage — Spark, Engage, Invite, Share — and how each contributes to the compounding system.

**Table 1.3. Examples of B2B companies using the Awareness Loop architecture (p. 5).** Maps Figma, HubSpot, Gong, and Atlassian against all four loop stages to show how the same architecture expresses through different product categories and motions.

**Table 2.1. Messaging and Perspective Spark patterns and examples (p. 6).** Catalogs three Spark Type 1 patterns — Knowledge Gap / Value Exchange, Data-backed Research Engine, and Provocation / Antics — with named company examples for each.

**Table 2.2. Product Spark patterns and examples (p. 7).** Catalogs five Spark Type 2 patterns ranging from full PLG (Free Tier, Free Utility) to non-PLG (Interactive Guided Experience, Managed Pilot, Diagnostic on Prospect Data) with examples for each.

**Table 2.3. Word-of-Mouth/Trust Spark patterns and examples (p. 9).** Catalogs five Spark Type 4 patterns — Customer-led Storytelling, Creator and Operator Influence, Analyst and Category Leadership, Peer Review Platforms, Employee Advocacy — with named examples for each.

**Table 3.1. Engagement modes and examples (p. 10).** Defines the three engagement modes — Product Engagement, Content Engagement, Community Engagement — and shows examples of each, with the strongest loops designing for all three.

**Table 3.2. Invite examples by level of commitment (p. 11).** Maps Cold, Medium, and Hot Invite paths to specific examples — newsletter subscribe and community join at the cold end; demo request, trial sign-up, and paid pilot at the hot end.

**Table 3.3. Share patterns with examples (p. 12).** Catalogs five Share patterns — Collaborative, Embedded, Output, Advocate, Content Amplification — with named company examples that illustrate each.

**Table 4.1. Loop Metrics (p. 13).** Organizes the loop dashboard around three metric categories — Spark Metrics, Engagement Metrics, Loop Velocity / Multiplier — with sample metrics under each.

**Table 6.1. Awareness Loop implementation at the Crawl band (p. 17).** Defines the build state for a Crawl-stage team across all four loop stages plus people, budget, and infrastructure — typically a 1–3 person team running one Spark type, one engagement surface, and one Invite path on under \$10K/month in tooling.

**Table 6.2. Awareness Loop implementation at the Walk band (p. 17).** Defines the build state for a Walk-stage team running 2–3 concurrent Spark types, owned community plus content plus light product engagement, a soft/medium/hard Invite ladder, and at least two engineered share patterns — staffed at 5–15 across marketing on \$50K–\$250K/month.

**Table 6.3. Awareness Loop implementation at the Run band (p. 18).** Defines the build state for a Run-stage team running all four Spark types with reported k-factor by type, continuous product/community/content engagement, an adaptive Invite ladder, and all five Share patterns engineered into product and community — staffed at 25+ on \$1M+/month with AI/data treated as core infrastructure.

## REFERENCES AND INTELLECTUAL LINEAGE

The works listed below represent intellectual lineages and foundational traditions that inform the framework presented in this paper. They are cited as sources of conceptual influence rather than as objects of direct empirical extension or formal theoretical comparison.

**Ref 1. Blank, S. (2013). *The Four Steps to the Epiphany*.** — Customer development, get-out-of-the-building. Sits behind the Engage and Invite stages, particularly the idea that the user reservoir is where you learn what's true.

**Ref 2. Hoffman, R., & Yeh, C. (2018). *Blitzscaling*.** — Directly relevant to the AI-native trajectory argument in Chapter 1 (Cursor, Midjourney, Perplexity reaching ARR milestones in fractions of the time SaaS-era companies took).

**Ref 3. Iansiti, M., & Lakhani, K. R. (2020). *Competing in the Age of AI*.** — The single most relevant academic-adjacent reference. Their "AI factory" and digital operating model arguments are the closest published analog to the systems-vs-people shift in Chapter 1.

**Ref 4. Ramaswamy, S. (2023). *Organize for AI*.** — Operational structures for AI feedback value. Directly behind the Run-band organizational implications in Chapter 6.4.

**Ref 5. Ries, E. (2011). *The Lean Startup*.** — Build-measure-learn is the conceptual ancestor of loop velocity and the metrics framework in Chapter 4.

**Ref 6. van der Kooij, Jacco J. (2024). *Revenue Architecture*.** — Required. The Bowtie, the Revenue Architecture frame, and the WbD lineage Jim references explicitly.

**Ref 7. Katsamakos, E., & Pavlov, O. (2020). "AI and Business Model Innovation: Leveraging the AI Feedback Loops."** — On-topic title; gives the paper one peer-reviewed citation specifically on AI feedback loops in business models.

**Ref 8. Pussinen, P. (2023). "Exploring digital platform growth through feedback loops."** — Same logic. Platform growth + feedback loops is exactly the territory.

**Ref 9. Thomke, S. (2020). *Experimentation Works*.** — Sits behind the "small loop in production beats a large loop in design" argument and the Crawl/Walk/Run experimentation cadence.

**About the Winning by Design Growth Institute:** The Growth Institute is the research and standards body of Winning by Design, focused on the applied study and synthesis of revenue architecture and growth systems. It integrates comparative field diagnostics with foundational systems theory to develop shared frameworks and language for designing, diagnosing, and governing compounding growth systems. Through its Working Paper series, the Institute serves as a think tank for revenue leaders navigating the transition from effort-driven growth to increasingly self-reinforcing systems.