

Analyst Relations and the Influence Layer

Pragmatic Remix: Analyst Relations • Industry Relations • Influencer Relations

Chapter 10 Analyst Relations and the Influence Layer Pragmatic Remix:
Analyst Relations • Industry Relations • Influencer Relations

In 2024, one of our team leads showed us a spreadsheet tracking the time we spent on analyst RFIs—the detailed questionnaires that firms like Gartner, Forrester, and IDC send when they’re evaluating vendors for Magic Quadrants, Waves, MarketScapes, and the rest of the industry’s competitive evaluation apparatus. The numbers were staggering. A single Gartner Magic Quadrant RFI consumed, on average, 120 personhours across product marketing, product management, and engineering. We were completing six to eight of these per year across different product lines. That’s somewhere between 720 and 960 hours annually—the equivalent of half a full-time employee doing nothing but answering analyst questions. The RFIs are important. Gartner alone influences more than \$100 billion in enterprise software purchasing decisions annually, and a Magic Quadrant placement can make or break a sales quarter. But the process of completing them is brutally inefficient. Each RFI asks dozens of detailed questions about product capabilities, customer references, market strategy, competitive positioning, and roadmap direction. Many of the questions are similar across different evaluations—the Gartner BI Magic Quadrant and the Gartner Planning Magic Quadrant ask overlapping questions about data integration, AI capabilities, and cloud architecture—but the answers need to be customized for each evaluation’s specific criteria and context. We built an agent to fix this. Not to replace the human judgment that goes into crafting an RFI response—the strategic decisions about what to emphasize, how to position roadmap items, which customer references to cite—but to handle the mechanical retrieval and first-draft generation that consumed most of those 120 hours. The agent draws from a curated knowledge base of past RFI responses, product documentation, customer success stories, and competitive positioning documents. It generates a firstpass response for each question, pulling from the most relevant source material and flagging where the previous response may be stale or where new product capabilities need to be incorporated. The result: we reduced RFI response effort by 50 to 60 percent, and in some cases by as much as 70 percent on targeted workflows where the questions closely matched previous evaluations. The quality didn’t suffer—in fact, it improved, because the agent

was more consistent about pulling in the latest product capabilities and customer proof points than a human working against a deadline. The human’s job became editing, strategic adjustment, and the genuinely hard

work of deciding how to position emerging capabilities that hadn't been part of previous evaluations. • • •

The Analyst Landscape Is Shifting The RFI automation story is important, but it's the easy part. The harder and more interesting transformation in analyst relations is about what the analysts themselves are becoming—and how the influence landscape is fragmenting in ways that change the PMM's job. The big three—Gartner, Forrester, IDC—still dominate enterprise technology purchasing decisions. That's not changing anytime soon. But the way they work is changing. Analysts are processing more vendor information, faster, using their own AI tools. A Gartner analyst evaluating twenty vendors for a Magic Quadrant is almost certainly using AI to help synthesize the mountains of RFI data, briefing notes, and customer reference feedback they receive. This means that the structured, specific, evidencebacked communication that works for buyer agents (Chapter 4) also works for analyst evaluations. Vague claims get filtered out. Specific capabilities with quantified outcomes get weighted. At the same time, new influence channels are emerging. Independent analysts and practitioners with large followings—people building audiences on Substack, LinkedIn, and YouTube—are shaping buyer perceptions in ways that didn't exist five years ago. When a respected practitioner writes a detailed comparison of data platforms based on their actual experience, that content gets shared in Slack channels and buying committee discussions alongside the Gartner report. It's a different kind of authority—experiential rather than institutional—but it's increasingly influential. And then there's the AI influence layer—the GEO dimension from Chapter 6. When a buyer asks Perplexity or ChatGPT to recommend data platform vendors, the AI's response is shaped by the same content that shapes analyst opinions: product documentation, review sites, published evaluations, customer case studies, and vendor content. This isn't analyst relations in the traditional sense, but it's influence—and managing it requires many of the same skills: understanding what the evaluating entity is looking for, ensuring your story is told accurately and compellingly in the channels that matter, and maintaining the consistency and evidence depth that builds credibility.

The PMM as Influence Architect The traditional AR function is relatively narrow: manage the relationship with Gartner, Forrester, and IDC analysts who cover your category. Prepare for briefings. Respond to RFIs. Lobby for positioning. Provide customer references. Track placements. In the agentic era, the AR function expands into something broader that we'd call influence architecture: the deliberate design and management of how your product is perceived across all the evaluation channels that matter to your buyers. This includes traditional analysts, but it also includes independent voices, AI systems, peer communities, and the increasingly important layer of customer advocates who share their experiences publicly. The skill set shifts accordingly. The traditional AR skill was relationship management—knowing the analyst, understanding their evaluation criteria, building rapport over time. That skill still matters. But the influence architect also

needs to think about structured information delivery (how to make your story evaluable by AI systems and AI-augmented analysts), content strategy (how to ensure the right content exists in the right channels to be surfaced by independent voices and AI search), and evidence management (maintaining a continuously updated library of customer proof points, quantified outcomes, and capability evidence that can be deployed across any influence channel on demand). The agent stack from earlier chapters supports this directly. Your competitive monitoring system tracks not just competitor moves but analyst commentary and sentiment shifts. Your knowledge base feeds both RFI responses and GEO-optimized content. Your content pipeline produces material that serves human analysts, AI systems, and buyer communities simultaneously. The influence architect orchestrates all of these into a coherent strategy that ensures your product's story is told accurately and compellingly wherever buyers look for evaluation input.

The Practitioner's Playbook: Analyst Relations and Influence If you own AR or influence at your company, here's where to focus. First, build the RFI knowledge base. Even if you don't automate the full RFI workflow, creating a curated repository of past responses, organized by topic and tagged with recency and accuracy flags, will cut your response time significantly. Every time you

complete an RFI, add the final responses to the knowledge base with metadata about which evaluation it was for, which analyst, and when. Over two to three evaluation cycles, you'll have a corpus that an agent can draw from effectively. Second, map your influence landscape. Make a list of every entity that shapes how buyers in your category evaluate vendors: the major analyst firms, the independent analysts and practitioners, the review sites (G2, TrustRadius, Gartner Peer Insights), the AI systems that surface vendor recommendations, the community forums and Slack groups where practitioners share opinions. For each entity, assess your current presence: are you well-represented? Is the information accurate and current? Are there gaps? This map becomes your influence strategy—a prioritized list of channels where investment will move the needle on buyer perception. Third, make the analyst briefing a two-way intelligence channel. Most PMMs treat analyst briefings as a pitch—an opportunity to tell the analyst your story. The better play is to use briefings as intelligence gathering. What are analysts hearing from buyers? What evaluation criteria are shifting? Which competitive narratives are gaining traction? What are they skeptical about in your story? An analyst who trusts you enough to share candid feedback is giving you intelligence that's worth more than the Magic Quadrant placement itself, because you can use it to improve your positioning, your product, and your competitive strategy.

A single Gartner Magic Quadrant RFI consumed, on average, 120 person-hours across product marketing, product management, and engineering. We were completing six to eight of these per year. That's the equivalent of half a full-time employee doing nothing but answering analyst questions.

We built an agent to fix this. Not to replace the human judgment that goes into crafting an RFI response — but to handle the mechanical retrieval and first-draft generation that consumed most of those 120 hours.

Figure 1: The RFI Automation Impact

Agent-powered RFI workflow: automate retrieval and first-draft generation, keep strategic judgment human.

120 Person-hours per Magic Quadrant RFI	6-8 RFIs completed per year	720-960 Hours annually ≈ 0.5 FTE on RFIs	50-70% Reduction in response effort
---	---------------------------------------	---	---

Dimension	Agent Handles	Human Owns
Source retrieval	Draws from curated knowledge base of past responses, product docs, customer stories.	Curates and maintains the knowledge base. Adds new responses after each cycle.
First-draft generation	Generates first-pass response for each question from relevant sources.	Reviews, edits, strategically adjusts. Decides emphasis and positioning.
Staleness detection	Flags where previous responses are stale or new capabilities need incorporation.	Decides how to position emerging capabilities not in previous evaluations.
Consistency	More consistent at pulling latest capabilities and proof points under deadline.	Strategic judgment on which customer references to cite for this evaluation.

Quality didn't suffer — it improved. The agent was more consistent about pulling in the latest capabilities than a human working against a deadline. The human's job became editing, strategic adjustment, and positioning emerging capabilities.

Figure 1. The RFI Automation Impact. 120 person-hours per RFI, reduced by 50-70%. The human's job shifts from retrieval and drafting to editing and strategic positioning.

The result: we reduced RFI response effort by 50 to 60 percent, and in some cases by as much as 70 percent on targeted workflows where the questions closely matched previous evaluations.

Figure 1: The AR Influence Model

The big three — Gartner, Forrester, IDC — still dominate enterprise technology purchasing decisions. But the influence landscape is fragmenting in ways that change the PMM's job.

Figure 2: The Influence Landscape

Four layers of influence that shape how buyers evaluate vendors. Traditional AR covers one. The influence architect covers all four.

LAYER 1: Traditional Analysts		
<i>Institutional authority. Structured evaluations. Still dominant.</i>		
Who Gartner, Forrester, IDC — the big three who evaluate vendors for Magic Quadrants, Waves, MarketScapes.	What's changing Analysts are using AI to synthesize vendor data. Vague claims get filtered. Specific, quantified capabilities get weighted.	PMM implication Structured, evidence-backed communication matters more than ever. The same discipline that works for buyer agents (Ch 4) works here.
LAYER 2: Independent Voices		
<i>Experiential authority. Practitioner credibility. Increasingly influential.</i>		
Who Practitioners with large followings on Substack, LinkedIn, YouTube. They write from actual experience, not vendor briefings.	What's changing Their content gets shared in Slack channels and buying committee discussions alongside the Gartner report. Different authority — experiential, not institutional.	PMM implication You can't brief them like analysts. You earn their attention by building a product worth writing about and relationships worth maintaining.
LAYER 3: AI Influence Layer		
<i>Algorithmic authority. GEO-driven. The fastest-growing evaluation channel.</i>		
Who Perplexity, ChatGPT, Claude, Gemini — AI systems buyers query for vendor recommendations.	What's changing AI responses are shaped by the same content that shapes analyst opinions: product docs, review sites, evaluations, case studies.	PMM implication GEO optimization (Ch 6) becomes an AR discipline. Your content must be structured for AI consumption and discoverable across evaluation channels.
LAYER 4: Peer Communities		
<i>Social proof. Review sites. Community consensus.</i>		
Who G2, TrustRadius, Gartner Peer Insights, community Slack groups, Reddit, practitioner forums.	What's changing Peer reviews are now a weighted signal for both human buyers and AI systems. Low review scores get surfaced in AI recommendations.	PMM implication Customer advocacy programs feed all four layers. A strong review presence improves analyst perception, AI visibility, and peer credibility simultaneously.
The multiplier: A Leader placement in the Magic Quadrant isn't just a website badge — it's a data point that AI systems weight heavily when filtering vendor shortlists. Every influence layer feeds the others.		

Figure 2. The Influence Landscape. Four layers of evaluation influence. Traditional AR covers Layer 1. The influence architect covers all four — because buyers, analysts, and AI agents all draw from the same evidence ecosystem.

This isn't analyst relations in the traditional sense, but it's influence — and managing it requires many of the same skills: understanding what the evaluating entity is looking for, ensuring your story is told accurately in the channels that matter.

Figure 2: The Agent-Powered AR System

The traditional AR function is relatively narrow: manage the relationship with Gartner, Forrester, and IDC analysts. Prepare for briefings. Respond to RFIs. Lobby for positioning. Track placements. In the agentic era, the function expands into something broader: influence architecture.

Figure 3: The Influence Architect

The skill set shift from traditional AR manager to influence architect. The relationship skill still matters — but it's now one of four capabilities.

Capability	Traditional AR	Influence Architect
Relationship management	Know the analyst. Build rapport. Understand their evaluation criteria. Manage the briefing calendar.	Still essential — but now extends to independent voices, community moderators, and understanding how AI systems weight different sources.
Structured information delivery	Write clear RFI responses. Prepare polished briefing decks.	Make your story evaluable by AI systems and AI-augmented analysts. Structured, specific, evidence-backed across every channel.
Content strategy	Produce analyst-facing content: briefing slides, RFI responses, reference stories.	Ensure the right content exists in the right channels to be surfaced by independent voices, AI search, and peer communities simultaneously.
Evidence management	Maintain a list of customer references for analyst requests.	Continuously updated library of customer proof points, quantified outcomes, and capability evidence deployable across any influence channel on demand.

The orchestration model: Your competitive monitoring system tracks analyst commentary and sentiment shifts. Your knowledge base feeds RFI responses AND GEO-optimized content. Your content pipeline serves human analysts, AI systems, and buyer communities simultaneously. The influence architect orchestrates all of these into coherent strategy.

Figure 3. The Influence Architect. From AR manager to influence architect. The relationship skill is now one of four capabilities required to manage how your product is perceived across all evaluation channels that matter to buyers.

The influence architect orchestrates competitive monitoring, knowledge management, content strategy, and evidence management into a coherent strategy that ensures your product's story is told accurately and compellingly wherever buyers look for evaluation input.

Figure 3: AR Maturity Framework

If you own AR or influence at your company, here's where to focus. Each move compounds — the knowledge base feeds the influence map, and both feed the briefing strategy.

Figure 4: The AR Practitioner's Playbook

Three moves that transform analyst relations from a Gartner relationship into a full-spectrum influence strategy.

01

Build the RFI Knowledge Base

The foundation everything else draws from

- The move** → Create a curated repository of past RFI responses, organized by topic, tagged with recency and accuracy flags. Every time you complete an RFI, add the final responses with metadata: which evaluation, which analyst, when.
- Why it works** → Over two to three evaluation cycles, you'll have a corpus that an agent can draw from effectively. Even without full automation, this cuts response time significantly and eliminates the "start from scratch every time" problem.
- First step** → Collect your last three completed RFIs. Tag each response by topic (product capabilities, AI, cloud architecture, security, customer references). That's your seed corpus.

02

Map Your Influence Landscape

Every entity that shapes buyer evaluation

- The move** → List every entity that shapes how buyers evaluate vendors: major analyst firms, independent analysts, review sites (G2, TrustRadius, Peer Insights), AI systems, community forums and Slack groups. For each, assess: are you well-represented? Is the info current? Are there gaps?
- Why it works** → This map becomes your influence strategy — a prioritized list of channels where investment will move the needle on buyer perception. Most AR teams manage Layer 1 and ignore the other three.
- First step** → Ask Perplexity to recommend vendors in your category. See what comes back. That's your AI influence audit — and it's often a wake-up call.

03

Make Briefings a Two-Way Intelligence Channel

The briefing as intelligence gathering, not just pitching

- The move** → Stop treating analyst briefings as one-way pitches. Use them to gather intelligence: what are analysts hearing from buyers? What criteria are shifting? Which competitive narratives are gaining traction? What are they skeptical about in your story?
- Why it works** → An analyst who trusts you enough to share candid feedback is giving you intelligence worth more than the placement itself — because you can use it to improve positioning, product, and competitive strategy.
- First step** → In your next analyst briefing, allocate the last 15 minutes to questions: "What are you hearing from buyers about [our category]? What do you think we're missing?"

The reframe: AR ROI in the agentic era includes a multiplier — AI agents weight analyst evaluations when filtering vendor shortlists. A Gartner Leader placement isn't just credibility for human buyers. It's a high-weighted signal for every AI agent that helps a buyer make a shortlist. Think influence holistically.

Figure 4. The AR Practitioner's Playbook. Three moves: build the knowledge base that powers RFI automation, map the full influence landscape beyond the big three, and transform analyst briefings from one-way pitches to two-way intelligence channels.

The traditional model treats analyst relations as a tax — something you have to do to maintain placements. In the agentic era, that framing is even less appropriate. The ROI of AR isn't just the human buyers who read the report. It's the AI agents that cite the report.

Figure 4: The AR Practitioner's Playbook

Analyst relations is one of the areas where agent-powered workflows have delivered the most tangible ROI—and the RFI automation described in this chapter is a real example from our team. But the broader point is about reframing AR investment entirely. The traditional model treats analyst relations as a tax—something you have to do to maintain Gartner and Forrester placements, a cost center that doesn't directly generate pipeline. In the agentic era, that framing is even less appropriate. When a buyer's agent evaluates vendors, it draws on analyst evaluations as a high-credibility signal. A Leader placement in the Magic Quadrant isn't just a website badge—it's a data point that AI systems weight heavily when filtering vendor shortlists. The ROI of AR isn't just the human buyers who read the report. It's the AI agents that cite the report. Think about influence holistically—analyst placements plus GEO, review sites, independent voices, and the full spectrum of evaluation channels. The PMM who thinks about influence this way will outperform the PMM who treats AR as a Gartner relationship. KEY TAKEAWAYS

- AR ROI in the agentic era includes a multiplier: AI agents weight analyst evaluations when filtering vendor shortlists.
- Reframe AR as a discoverability investment, not just a credibility investment.
- RFI automation dramatically reduces the operational burden that makes AR feel like a tax.
- Think influence holistically: analyst placements, GEO, review sites, and independent voices all feed agent evaluation.

PART III Building the 10x PMM Organization
