


A PLAYBOOK FOR SECURITY MARKETING LEADERS

Govern the *decisions*, not the tasks.

The Marketing Decision Map: how to find the marketing decisions AI now influences, decide how much AI may do in each, and roll the rules out in 90 days.



Christine Castro

B2B CYBERSECURITY CMO AND GTM ADVISOR

CHRISTINECASTRO.COM

01 / WHY THIS MATTERS

AI is already making marketing decisions. Most teams never decided it should.

Marketing teams plan around tasks: campaigns, content, events, sequences. AI is very good at tasks, so it slips into the decisions behind them. A drafting tool starts choosing which claims to make. A sequencing tool starts choosing who to email. A scoring model quietly changes what counts as a lead.

None of that is wrong on its own. The problem is that nobody chose it, nobody owns it, and nobody can show the CEO which decisions AI touched or what they produced. That is how a security company ends up with an unapproved claim in front of a CISO, or customer data in a tool the security team never reviewed.

1 in 3

CMOs are getting the returns they expected from AI investments

SOURCE: GARTNER

15%

of CEOs say their CMO has strong AI savviness

SOURCE: GARTNER

70%

of marketing leaders say their processes are not yet mature enough to scale AI

SOURCE: GARTNER, 2026 CMO SPEND SURVEY

The fix is to organize the work around decisions. Name the decisions that matter, give each one an owner, set how much AI may do, and give people and AI the same context to decide from. Then AI speeds up the right things and stays out of the wrong ones.

How to use this playbook. Read it once with your leadership team. Score the decisions in the companion workbook or the free online tool, write cards for your top five, and run the 90-day rollout. Every list here is a starting draft written for security marketing. Change it until it fits your company.

Every decision gets one of three AI roles.

The question is never "should we use AI?" It is "for this decision, how much may AI do, and where does a person sign off?"

Human decides

AI may gather inputs and summarize. A named person makes the call and owns it. Use for anything a board, a customer or a regulator would hold you to.

AI recommends for approval

AI drafts the options and a recommendation with its reasoning. Nothing ships until the owner approves it. Use for decisions with real consequences that repeat often enough to benefit from a first draft.

AI acts within guardrails

AI makes the call inside written limits (budget caps, approved claims, audiences, send volumes) and logs every decision. The owner reviews a sample on a set cadence and can pull the authority back.

How to choose

IF THE DECISION IS	START AT
High risk and rare, or high value and hard to reverse	Human decides
High risk and frequent	AI recommends for approval
Moderate risk and frequent	AI acts within guardrails, with a sample review
Low risk	AI acts within guardrails if it is frequent; AI recommends if it is rare

Start conservative. Move a decision up a level only when its decision log shows a clean record: few overrides, no errors that reached a customer, and outcomes at least as good as before. Pull authority back the moment the log says so. That rule matters more than the starting level.

At a security company, three decisions almost always stay human: what threat research to publish and when, which AI tools may touch customer or prospect data, and what goes in front of the board.

Score each decision on value, risk and frequency.

Score with the people who make each decision, not for them. The conversation is worth as much as the number.

Value (1 to 5)

Does it move revenue, pipeline or brand? Does it constrain or open up later choices? Would changing it change the plan for the year?

Risk (1 to 5)

Can it be reversed? What does a wrong call cost with customers, the board, regulators or the press? How likely is that cost?

Frequency (1 to 5)

1 is once a year, 3 is monthly, 5 is daily. Frequent decisions are where AI takes over without anyone deciding it should.

Priority = Value × Risk + Frequency, out of 30. Value and risk multiply because a decision that is both valuable and risky is far more dangerous to get wrong than either alone. Frequency adds weight without letting a trivial daily task outrank a board decision.

TIER	SCORE	WHAT TO DO
GOVERN FIRST	22 to 30	High value and high risk. Write a decision card, name the owner and set the approval gate before AI touches it further.
GOVERN NEXT	14 to 21	Real consequences, or frequent enough that errors add up. Card these in the second wave.
SPOT-CHECK	2 to 13	Lower stakes. Let the owner set simple limits and review a sample on a schedule.

Rank ties by risk, then value. Then use judgment: the score informs which five to govern first, it does not make the choice for you. If two decisions are really one, merge them. If a decision has no clear owner today, that alone is a reason to put it on the shortlist.

04 / THE DECISION INVENTORY

24 decisions security marketing teams make

Suggested scores for a typical Seed to Series C security company. Yours will differ. Untick what does not apply and add your own.

DECISION	VALUE	RISK	FREQ	PRIORITY	TIER	STARTING AI ROLE
BUYER AND MARKET						
Which buyer the message leads with: practitioner or CISO	5	4	2	22	GOVERN FIRST	Human decides
Which segments and verticals get priority this half	5	4	2	22	GOVERN FIRST	Human decides
Which accounts get one-to-one treatment	4	3	3	15	GOVERN NEXT	AI recommends for approval
Which intent signals trigger outreach	4	3	5	17	GOVERN NEXT	AI acts within guardrails
POSITIONING AND CATEGORY						
Which analyst category to position in or against	5	5	1	26	GOVERN FIRST	Human decides
Competitive claims and comparisons approved for use	4	5	3	23	GOVERN FIRST	AI recommends for approval
When to name and claim a new category	5	5	1	26	GOVERN FIRST	Human decides
How the product is described against adjacent tools	4	3	2	14	GOVERN NEXT	AI recommends for approval
CONTENT, RESEARCH AND PROOF						
When and what threat research to publish	5	5	2	27	GOVERN FIRST	Human decides
What AI-drafted content can publish without expert review	3	4	5	17	GOVERN NEXT	AI recommends for approval
Customer names, logos and quotes approved for use	4	5	3	23	GOVERN FIRST	AI recommends for approval


04 / THE DECISION INVENTORY, CONTINUED

DECISION	VALUE	RISK	FREQ	PRIORITY	TIER	STARTING AI ROLE
BUDGET AND INVESTMENT						
Annual budget split across brand, pipeline programs and open funds	5	4	1	21	GOVERN NEXT	Human decides
Which conferences and sponsorships to commit to	4	4	3	19	GOVERN NEXT	AI recommends for approval
Agency, contractor and tool spend	3	3	2	11	SPOT-CHECK	AI recommends for approval
The next marketing hire	5	4	1	21	GOVERN NEXT	Human decides
PIPELINE AND CHANNELS						
Lead scoring and the handoff to sales	5	4	2	22	GOVERN FIRST	Human decides
What outbound AI may send, and to whom	4	5	5	25	GOVERN FIRST	AI recommends for approval
Weekly paid budget moves between campaigns	3	2	5	11	SPOT-CHECK	AI acts within guardrails
Channel mix shifts mid-quarter	4	3	2	14	GOVERN NEXT	AI recommends for approval
MEASUREMENT AND GOVERNANCE						
The attribution model and what counts as sourced or influenced	5	4	1	21	GOVERN NEXT	Human decides
What marketing reports to the board	5	5	1	26	GOVERN FIRST	Human decides
Which AI tools may access customer and prospect data	4	5	2	22	GOVERN FIRST	Human decides
How AI performance is judged	4	3	1	13	SPOT-CHECK	AI recommends for approval
When to give AI more authority, or take it back	4	4	2	18	GOVERN NEXT	Human decides

With the suggested scores, 11 decisions land in Govern first, 10 in Govern next and 3 in Spot-check. Most teams find two or three of their own to add, usually around partner marketing, pricing pages or a specific compliance claim.

Where the decisions sit

Each circle is a decision from the inventory, placed by risk and value. Circle size is frequency. The shaded corner is where to start.



The default top five

1. **When and what threat research to publish** 27, Human decides
2. **Which analyst category to position in or against** 26, Human decides
3. **When to name and claim a new category** 26, Human decides
4. **What marketing reports to the board** 26, Human decides
5. **What outbound AI may send, and to whom** 25, AI recommends for approval

Notice what is missing from the top five: almost nothing about content volume, campaign execution or ad optimization. Those are where AI saves the most time and carries the least risk. The decisions that need governing are the ones about what the company claims, who it contacts, and what leadership is told.

One card per governed decision

A decision card is one page that anyone, including an AI tool, can read and follow. It is finished when the owner, sales and finance would all sign it.

- OWNER

One named person. A team cannot own a decision.
- CONSULTED

Who has input before the call is made.
- INPUTS

The data and evidence the decision needs. If AI cannot see them, it should not decide.
- WHAT AI MAY DO

Specific actions, tied to the AI role.
- WHAT STAYS HUMAN

The part of the call a person makes and signs.
- APPROVAL GATE

Where sign-off happens, built into the tool, and how often the rules are reviewed.
- NEVER

The one line nobody crosses, including AI. Write it in plain words.
- HOW YOU KNOW IT WORKED

An outcome measure, not an activity count.

WORKED EXAMPLE

PRIORITY 25 • GOVERN FIRST • AI RECOMMENDS FOR APPROVAL

What outbound AI may send, and to whom

Owner	Demand gen lead
Consulted	SDR leadership, legal, brand
Inputs	Approved claims, sequences, audience rules, sender reputation, opt-out rules
What AI may do	Research, personalize and draft; send only templates approved for the segment
What stays human	First sends to any new segment, executive outreach, anything off-template
Approval gate	Human approves new sequences and the first batch to any new audience; daily sample review
Never	Let AI email a CISO with an unapproved claim or a fabricated personal detail
How you know it worked	Positive reply rate, meetings per 100 sends, complaint and bounce rates

06 / DECISION CARDS, WORKED EXAMPLES

PRIORITY 19 • GOVERN NEXT • AI RECOMMENDS FOR APPROVAL

Which conferences and sponsorships to commit to

Owner	Field marketing lead
Consulted	Sales leadership, CMO
Inputs	Attendee fit, past results, cost per meeting, contract deadlines (booth space is signed months ahead)
What AI may do	Score events against the ICP and forecast cost per meeting
What stays human	Contract commitments and the yearly event portfolio
Approval gate	Any sponsorship over a set amount needs CMO sign-off
Never	Renew a show because you were there last year
How you know it worked	Pipeline and cost per qualified meeting by event

PRIORITY 23 • GOVERN FIRST • AI RECOMMENDS FOR APPROVAL

Competitive claims and comparisons approved for use

Owner	Product marketing lead
Consulted	Legal, product, sales engineering
Inputs	Test results, customer evidence, competitor documentation, legal review
What AI may do	Draft comparison pages and battlecards from approved claims only
What stays human	Which claims are approved, and the evidence each one rests on
Approval gate	Every new claim approved by product marketing and legal before use
Never	Let AI invent a feature gap, a benchmark or a customer outcome
How you know it worked	Win rate in head-to-head deals; zero retracted claims

Drafts for all 24 decisions are in the Card Library tab of the companion workbook.

Give people and AI the same material to decide from.

Most inconsistent AI output is not a model problem. It is a context problem: the tool never saw the approved claims, the ICP or the voice. Put these in one place every person and every AI tool reads from, such as a Claude Project or a shared drive your tools connect to.

☐ ICP and buying committee	Firmographics, the roles involved (practitioner, security leader, CISO, procurement), who signs and who champions.
☐ Positioning and category statement	One paragraph: the problem, the category, what you do differently, and what you do not do.
☐ Approved claims and proof library	Every claim marketing may make, each with its evidence and its owner. AI drafts only from this.
☐ Claims and language you never use	Banned words, unsupported superlatives, competitor statements legal has ruled out, fear-based language.
☐ Voice samples	Five to ten pieces that sound right, and two or three that do not, with a note on why.
☐ Funnel definitions and current numbers	Stage definitions both teams signed, conversion rates, cycle length, ACV, pipeline target.
☐ Competitive notes	Who you meet in deals, where you win and lose, approved comparison language.
☐ Customer reference rules	Who can be named, who cannot, and what each customer has approved.
☐ Research and disclosure policy	How threat research is reviewed, timed and attributed.
☐ Decision cards	The cards from this map: owner, AI level, approval gate and guardrails for each governed decision.
☐ Data access rules	Which data each AI tool may see, and where customer data may never go.
☐ Escalation contacts	Who to ask when a draft touches legal, security, a customer or the press.

Build the approved claims library first. In security marketing it prevents the most expensive mistakes, and every other decision card depends on it.

08 / READINESS CHECK

Where does your team start?

Answer for how your team works today, not how the policy reads. Count the yeses.

- | | |
|--|----------|
| ☐ AI use in marketing is judged on pipeline or revenue, not only hours saved. | Yes / No |
| ☐ Your most important marketing decisions each have one named owner. | Yes / No |
| ☐ There is a written list of what AI may do without a human approving it. | Yes / No |
| ☐ Claims, voice and ICP live in one place that people and AI both use. | Yes / No |
| ☐ The security team has reviewed every AI tool that touches customer data. | Yes / No |
| ☐ Someone reviews a sample of AI output on a set schedule. | Yes / No |
| ☐ Marketing and sales signed the same funnel definitions. | Yes / No |
| ☐ You could show your CEO which decisions AI influenced last quarter. | Yes / No |

0 to 2: Tasks, not decisions

AI is speeding up tasks, but no one has decided where it may decide. Start with the inventory and the top three cards.

3 to 5: Partly governed

Some owners and rules exist, mostly in people's heads. Write the cards down and build the shared context so the rules survive turnover.

6 to 8: Governed

You can expand AI authority safely. Use the decision log to move proven decisions up a level and report the results to the CEO.

Observe, Analyze, Execute

The same sequence I use when I join or advise a marketing team: understand what is true before changing it, build the model, then install the operating rhythm.

Observe

DAYS 1 TO 30

1. List the decisions marketing makes, starting from the inventory here, and add your own.
2. Find where AI already influences a decision, including tools people adopted on their own.
3. Score every decision on value, risk and frequency with the people who make them.
4. Name one owner for each of the top five.

Analyze

DAYS 31 TO 60

1. Write a decision card for each of the top five: inputs, AI level, approval gate, guardrails, measure.
2. Build the shared context in one place, starting with claims, ICP and funnel definitions.
3. Have the security team review every AI tool that touches customer or prospect data.
4. Set approval gates in the tools themselves, not in a policy document nobody opens.

Execute

DAYS 61 TO 90

1. Run the top five under their cards and log every decision AI touched.
2. Review error and override rates every two weeks with each owner.
3. Move one proven decision up an AI level, and pull one back if the log says so.
4. Report to the CEO which decisions improved, by how much, and what you will govern next.

What the CEO sees at day 90

ITEM	WHAT IT SHOWS
The governed decisions	The top five, each with an owner and an AI role, on one page
The decision log	How many decisions AI touched, how many a person overrode, and why
Outcomes	The measure on each card, against the 90 days before
Changes in AI authority	What moved up a level, what was pulled back, and the evidence
What comes next	The next five decisions to govern

Keep a decision log, and review it on a rhythm.

The log is what lets you give AI more authority safely, and what lets you prove to the CEO that the rules work.

What each log entry records

FIELD	WHY IT MATTERS
Decision and date	Which card the decision falls under
What AI recommended or did	The baseline for judging AI's track record
What the person decided	Agreement, edit or override
Reason for any override	The fastest way to improve the shared context
Outcome, when known	Ties the decision to the card's measure

Review rhythm

CADENCE	WHO AND WHAT
Daily or weekly	Owners sample AI output for decisions set to AI acts within guardrails
Every two weeks	Each owner reviews override and error rates for their decisions
Monthly	The CMO reviews the log across all governed decisions and approves any change in AI role
Quarterly	Rescore the inventory, retire cards that no longer matter, add the next five; report to the CEO

Judge AI on the outcome of the decision it supports, such as meetings per 100 outbound sends or win rate in head-to-head deals, not on hours saved.

Six mistakes that undo the work

Writing a policy instead of cards

A ten-page AI policy nobody opens changes nothing. One page per decision, in the tools, does.

Governing tasks instead of decisions

Rules about which tools may write blog posts miss the point. Govern what gets claimed, who gets contacted and what leadership hears.

Shared ownership

A decision owned by marketing and sales is owned by nobody. Name one person; list the rest as consulted.

Skipping the security review

At a security company, an unreviewed AI tool holding customer data is a credibility problem, not only a compliance one.

Expanding authority on a vendor's say-so

Move a decision up a level because your log is clean, not because a tool claims it is ready.

Measuring time saved

Time saved is real, but it does not show whether the decisions got better. Measure the outcome on each card.

Free tools on christinecastro.com

Marketing Decision Map

Score the decisions online, write your top five cards and export your map to Excel and PDF.

christinecastro.com/marketing-decision-map.html

GTM Alignment Diagnostic

AI rules only hold when sales, marketing and product agree on the ICP and the plan. 83 questions across 13 dimensions, with a free 21-question scan.

christinecastro.com/gtm-alignment-diagnostic.html

Board Slide Planner

The five marketing slides a board reads each quarter, built from your numbers.

christinecastro.com/board-slide-planner.html

Observe / Analyze / Execute

The operating framework behind the 90-day rollout.

christinecastro.com/observe-analyze-execute.html

Statistics credited to Gartner are from Gartner research for marketing leaders, 2026, and are cited for context. The decision inventory, scoring method, AI roles, decision cards and rollout in this playbook are the work of Christine Castro.

FINALIZE WITH CHRISTINE

What this playbook cannot do on its own

The playbook and the workbook give you the method and a strong first draft. These are the pieces that turn it into rules your team follows and your CEO can see working.

Scores calibrated to your company	The suggested scores here are generalized. Which decisions carry real risk depends on your stage, your buyers, your category and how your board measures marketing.
Decision cards written with your team	A card only works when the owner, sales and finance agree to it. I run the working session that gets the top five signed.
Shared context built and loaded	Claims library, ICP, voice and funnel definitions written once and set up in the AI tools your team actually uses, with approval gates in place.
AI governance your CEO and board can read	One page on which decisions AI influences, who owns each one and what it is producing, for the next board meeting.

Want this built with your team?

I run the working session that scores your decisions with the people who make them, gets the top five cards signed, and sets up the shared context and approval gates in the tools your team already uses. Then I report the results to your CEO.

WORK WITH CHRISTINE

Or send your completed workbook to christine@christinecastro.com and I will read it before we talk.