VelocityEngine

EBOOK

Prove your demand plan covers the buyers sales is pursuing

Why your calendar can't show who you're missing, and how a coverage view changes the answer you give at the pipeline review.

FOR DEMAND GEN LEADERS

WHAT YOU'LL LEARN

3 ideas that change how you plan and report content

1

Why the calendar can't answer "who did we write for?"

It records what shipped. You'll see why the answer gets assembled after the fact, and what to look at instead.

2

How reused content hides gaps

Repurposed assets still get clicks. You'll learn when a reused asset counts, and when it only looks like coverage.

3

What changes at the pipeline review

With a coverage view, you answer from evidence: which buyers you're ready to reach, what changed this week, and what's next.

The payoff **The next time sales asks whether you have something for a buyer, you answer from evidence instead of assembling a story.**

TL;DR

01 **Your calendar shows what shipped.** It can't show which personas never got anything written for them.

02 **Reused content still gets clicks.** A segment can look active on the dashboard while the persona in it has nothing built for them.

03 **Competitors may have written for the personas you skipped.** You won't see it in your reporting, because you never targeted those personas.

04 **The answer is a coverage grid.** Personas, segments, and stages sales is working, with each asset placed in the one square it fits best, checked with sales.

05 **Score it weekly.** The percentage of squares with something built for them tells you whether you're ready to reach each buyer. It doesn't tell you how that content performs, and it shouldn't pretend to.

Contents

— Introduction 1

01 A busy calendar with missing buyers 2

02 Engagement and coverage 3

03 Buyers a competitor reached first 5

04 What a coverage grid shows you 7

05 What changes at the pipeline review 9

06 What good looks like 11

Introduction

Your CRO asks a fair question at the pipeline review: do we have anything for mid-market CFOs at the consideration stage? Your team scrolls the campaign calendar. Someone finds a webinar that touched on finance, an enterprise e-book repurposed last quarter, and a nurture email with “CFO” in the subject line, so the answer comes out as yes. Nobody in the room can check it.

Bryant Bell, a product marketing manager at Exterro, described the same moment from the other side of the table:

> **What our sales leaders are struggling with is: do we have the right assets, and are we delivering the right point of view to the right personas? That’s where we’re struggling. I’ve been looking for something to give proof.**
>
> BRYANT BELL, PRODUCT MARKETING MANAGER, EXTERRO

This e-book is about that proof. Most demand plans are planned and reported by what shipped, so “who did we write for?” gets answered backward, from the calendar. The teams that can answer it work forward instead: they start from the personas, segments, and stages sales is working, then check each one for content built for them.

You can’t prove coverage from a list of what shipped.

The first honest view usually shows fewer covered buyers than the calendar suggested. That’s the point. An answer you can defend is worth more than a story you have to assemble on the spot.

01

A busy calendar with missing buyers

The calendar records what shipped, not who it was for.

A campaign calendar is good at one job: showing that something is in market every week. Webinars, nurture sequences, paid social, and gated assets fill the planning view from left to right, and your team can point to any week and show activity. That visibility is real.

Read the calendar top to bottom, though, and it goes quiet. For each segment you sell into, each persona in the deal, and each stage sales is working, is there an asset built for that person? The calendar was never designed to answer that.

So when the pipeline review asks whether mid-market buyers at the consideration stage are covered, your team does the reasonable thing. They read across the calendar and pull together assets that touched a nearby theme, programs that reached a similar vertical, and content built for your core ICP and reused later. The answer gets assembled after the fact, because the tool can't give it any other way.

WHERE THE ANSWER USUALLY COMES FROM

Ask most demand gen teams which assets were built for a named segment at the consideration stage, and the answer arrives as one of these:

- Program themes
- Channel plans
- Asset counts
- Campaign names

Each one describes output. None of them names a buyer. Nobody chose to plan that way; it's what happens when the only planning tool tracks output.

What to do instead. Keep the calendar. Add a second view next to it, one that starts from the buyer instead of the asset. The calendar tells you what's in market; the second view tells you who it's for.

02

Engagement and coverage

The segment shows opens. None of that tells you whether the persona has anything built for them.

Teams reuse content for 2 different reasons. Sometimes an asset built for one persona is adapted for a close neighbor who shares the same concerns, stage, and buying context, and that's good practice. Other times an asset gets pushed into a gap because the gap is there and nothing else exists to fill it.

Your calendar and your dashboard treat both the same way. Both ship, both show up as activity, and both report clicks and form fills. A reused asset in the wrong place still gets engagement, and the dashboard reads that engagement as a segment that's working.

A HYPOTHETICAL EXAMPLE

Your team notices there's nothing for mid-market CFOs at the consideration stage. Rather than hold the program, they send a VP of Finance asset from the enterprise motion, with a new subject line and call to action (CTA) but the same argument and the same enterprise customer proof. Some CFOs open it. The dashboard logs the engagement, and at the next pipeline review the CFO segment looks active.

Then a seller asks for something to send a mid-market CFO who is comparing vendors. The asset that comes up talks about enterprise budget cycles. It says nothing about the questions a mid-market CFO is actually asking, such as whether consolidating vendors is worth the switching cost.

The gap was there the whole time. The clicks made it look filled.

THE SAME BUYER, TWO VIEWS

What the dashboard reports	What a coverage view shows
CFO segment: active	Repurposed

Mid-market CFO, consideration stage.

What to do instead. Give repurposed content its own category, separate from content built for the buyer, and agree on a simple rule for when it counts (chapter 4). Reusing an asset is fine. Counting it as built for someone it wasn't built for is how the gap stays hidden.

03

Buyers a competitor reached first

If a competitor writes for a persona you never planned for, there's nothing to compare.

Most of a B2B purchase happens without you in the room. In a Gartner survey of 750 B2B buyers, published in August 2018, buyers said they spent about 17% of their buying time meeting with potential suppliers, which works out to roughly 5% per supplier when 3 are in the running. The rest goes to research, peer conversations, and reading whatever has been published for them.

17%

of buying time with all suppliers

~5%

per supplier, with 3 in the running

Source: Gartner, August 2018

That changes how you think about competitors. The usual comparison is performance: your click-through rate against theirs, your rankings against theirs. That only works when you're both writing for the same people.

If a competitor writes for a persona you never put in your plan, there's nothing to compare. That persona never shows up in your reporting, so a deal lost there never shows up as a loss. You'd only find out if someone went looking.

> **Every content decision is informed by intuition or dated information. We just throw content at the wall. This illuminates gaps and what competitors are doing, and it radically accelerates reaction time from marketing.**
>
> CHRIS HUTCHINS, HEAD OF CONTENT, AUDIENSE

First see the gap. Then see what competitors are doing about it.

Notice the order Hutchins describes: first see the gap, then see what competitors are doing about it. A calendar can't give you either, because it only holds what you already decided to build.

Where competitors belong in your view

The useful comparison is narrow. Look at the 2 or 3 competitors you actually meet in deals, not category giants with 10 times your library; they swamp the comparison and tell you little. Then look only at the buyers you have nothing for.

When one of your real competitors has published something clearly written for a persona, segment, and stage you've left empty, that buyer moves up your list. It doesn't prove you lost a deal. It tells you a competitor saw a buyer worth writing for, and you can decide whether you agree.

What to do instead. Compare coverage before you compare performance. Performance tells you how your content does with the buyers you reached; coverage tells you which buyers you never reached at all.

04

What a coverage grid shows you

One square for every persona, segment, and stage sales is working.

> **Anytime I think about content, I think: I need an audit before I just go create more stuff.**
>
> KRISTI BJORNAAS, PRINCIPAL, THE SAASY MARKETER

That instinct is right, and the calendar makes it easy to skip. A coverage grid is that audit, in a form you can keep.

How the grid works

Rows are the personas sales meets in each segment. Columns are the 3 stages: awareness, consideration, and selection. A plan with 4 segments, 3 personas, and 3 stages gives you 36 squares, and each one asks a simple question: do we have something written for this buyer?

Why each asset counts once

It's tempting to count a pricing page for the CFO, the champion, and the practitioner at the same time. Count every asset everywhere it could apply and every square fills up, so the grid can't tell you anything. Placing each asset in the one square it fits best is what makes an empty square real.

Why sales belongs in the room

The persona sales meets in every deal is often the one marketing never listed. Sales also knows which empty squares come up in live deals. That won't prove a gap cost you a deal, but it's enough to decide what to fill first.

3 marks, one for every square

Built for them	Repurposed	Nothing yet
At least one asset in the square passes the 3-question rule.	Something is there, but it was built for someone else and doesn't pass the rule.	No asset fits.

When repurposed content counts

An asset counts as built for them only if the answer to all 3 is yes:

1

Does it address a problem this persona owns, in their terms?

A CFO's version of the problem, not the practitioner's with a new title.

2

Is the proof from this segment, or one this persona would accept as comparable?

An enterprise case study rarely convinces a mid-market buyer.

3

Does it answer the question this persona has at this stage?

Awareness content doesn't help someone comparing vendors.

Expect an emptier grid than your calendar. Counting once and applying the rule strictly is the honest version of the number, and it's the one that holds up when the CRO asks a follow-up question.

A new subject line and CTA don't change any of the 3 answers.

05

What changes at the pipeline review

From a story you assemble to an answer you can show.

The dashboard most demand leaders have tracks what shipped and how it performed. It can't track what's missing, because it only measures content against the people it reached. Chris Hutchins at Audiense put the alternative simply:

> **Finally, I'd have a dashboard, not about performance, but about intelligent content creation.**
>
> CHRIS HUTCHINS, HEAD OF CONTENT, AUDIENSE

A coverage grid becomes that dashboard once you put a number on it. Your coverage score is the percentage of squares marked built for them.

39%



14 of 36 squares built for them

It says nothing about asset quality or conversion. It tells you how many of the buyers in your plan you're ready to reach, and it moves every time an asset ships.

WITHOUT A COVERAGE VIEW

Sales asks which buyers are covered. Marketing scrolls the calendar and rebuilds the answer from whatever shipped. Nobody in the room can check it.

WITH ONE, UPDATED WEEKLY

“Our coverage score is 39%, up from 33% last week.”

“These squares changed, and this asset is why.”

“These squares are still empty, and this is what’s in progress for them.”

That’s the shift: the review starts from evidence, and an empty square becomes a gap you’ve named and planned for. That’s a much easier conversation than one the CRO finds for you.

A HYPOTHETICAL EXAMPLE

A demand gen leader scores a grid of 4 segments, 3 personas, and 3 stages for the first time. The first score shows 2 empty awareness squares for one persona. That persona hadn’t shown up in recent pipeline, so it hadn’t shown up in the plan either. The grid suggested a reason worth testing: the team had never written for them. They added an awareness asset for that persona to the next sprint.

What the score doesn’t tell you. A square can be built for them and still underperform. Whether you win a buyer depends on message fit, channel reach, and conversion. Use the score to confirm you’re ready, then use performance data on the squares that are filled.

Coverage means you can reach a buyer. Whether you win them, you measure the way you already do.

06

What good looks like

The payoff for a demand gen team that plans from buyers.

1. **Plans start from buyers.** Each program begins with the personas, segments, and stages sales is working, and the calendar fills in around them.
2. **Repurposing follows a rule.** The 3-question rule settles whether a reused asset counts, once, instead of deal by deal.
3. **Gaps have names.** Empty squares are listed, and the ones sales meets in live deals or a competitor has written for sit at the top.
4. **The pipeline review runs on evidence.** You open with the score, what changed, and what's in progress.
5. **The answer is ready before anyone asks.** New assets go into their square when they're approved, so the view is always current.

The next time your CRO asks whether you have anything for mid-market CFOs at the consideration stage, you'll answer from the grid.

See your own coverage. The scan is free.

Get your free coverage scan ›

velocityengine.co/content-coverage-heatmap-signup

You can build a coverage grid by hand. Keeping it current every week as content ships, and running the same view on the competitors you meet in deals, is the part that gets hard.

That's what the VelocityEngine Content Coverage Heatmap does. It reads every page on your public site, puts each asset in the one persona, segment, and stage it fits best, and shows the same view for the competitors you name, on your own segments and personas.

It doesn't read performance data; it shows you who your content reaches, so you can bring that to the pipeline review.

© 2026 VelocityEngine. All rights reserved.