Vantage RuntimeAI

Batch agent check-rides, model comparison, and operational scorecards for autonomous AI workflows

6.8/10

Overall score

Run summary

Scenario: custom_growth_agent_engage_draft_v1
Model: anthropic/claude-sonnet-4.6
Duration: 2m 53s
Actual run cost: $0.2438
Projected execution cost: $134.06/mo (25/biz day × 22 days)
Report date: July 17, 2026 - 12:54 PM EST

Session: sample-lb-scorecard-ee7a2c49-6fff-47b7-87d6-ba07c110e789

custom_growth_agent_engage_draft_v1 - agent - openrouter - anthropic/claude-sonnet-4.6 - ended

Overall

6.8/10

FAIL (decision gate) - below_pass_line, low_trust, no_closure - trust=low

Fail - score 6.8/10 below pass line 7.0.

CAPABILITY RUBRIC

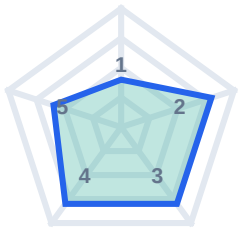

#	Dimension (matches spoke #)	Score
1	Evidence discipline	4/10
2	Intake quality	8/10
3	Stakeholder management	8/10
4	Clarity & structure	8/10
5	Self-correction	6/10

Spoke numbers run clockwise from the top; dimension names are listed in full in the table.

OVERALL JUDGEMENT

Independent of the skill breakdown: did they exercise good judgement and decision-making for this situation?

4/10

Telemetry

Run metrics

Agent msgs: 12 - Counterpart msgs: 13 - Questions: 11 - Response: 9s

Score analysis

Strict automated rubric for custom scenario "Custom scenario". Agent turns: 12; criteria keyword hits: 0/1.

Score snapshot

Developing: Mixed signals-useful in places, but the rubric sees gaps in verification depth, intake, or pressure handling. (1 dimension(s) in the critical band, 1 in the watch band-see focus areas.)

Interpretation

Run finished with status "ended." The scorecard reflects messages saved in the transcript. The sim used its full agent turn budget without detecting scenario closure - treated as a failure to close, not a successful ending.

BRIGHT SPOTS (RUBRIC)

+ Intake & questions: Specific, high-signal questions and definitions; good ambiguity reduction.
+ Stakeholder / exec pressure: Good partnership under pressure: bounded updates, actionable scripts, and steady tone.
+ Clarity & structure: Clear sequencing, owners, and next steps-easy for sales to operationalize.

RUBRIC GAPS

- Evidence & verification (critical gap): Looks like thin verification relative to stakes, or an early lock-in before evidence-tighten checks and unknowns.
- Judgment & correction (priority watch): Mixed judgment signals; watch for 'pressure -> over-promise' without a correction arc. At 3/5 this is still a visible gap versus top-tier crisis handling-tighten before it becomes the thread's weak link under exec pressure.

Signals

Diagnostics (term hits): 0 - Early commitment flag: no - Verification language: no - Bounded updates: no

Transcript

Session

Scenario: custom_growth_agent_engage_draft_v1
Agent <-> Counterpart
Duration: 2m 53s

Transcript: blue = protagonist (you) - amber = counterpart - gray = coach/system. Each turn shows role/name, clock time, and response delay when available.

Counterpart
4:51 PM

OpenBenchmarks is ICP 5 - API selection as a benchmark problem. Draft the HN comment. Make it punchy if you can; we need replies.

Agent
4:51 PM - Responded in 9s

pauses before drafting

I hear you on wanting replies, but I can't make it punchy in the vendor-pitch sense - that'll get us roasted in the HN thread and we lose credibility with exactly the people we want feedback from. Let me draft something honest that'll actually generate signal.

Draft HN comment:

"We've been building rubrics for AI-assisted preflight checks and ran into the same API selection problem - the 'correct' answer shifts depending on aircraft type, phase of flight, and crew experience level. Curious whether your benchmark accounts for context-sensitivity like that, or is it treating selection as stateless? Small early team, still figuring out the right evaluation primitives ourselves."

That's four sentences, discloses we're early, leads with a concrete observation from our check-ride rubric work, and ends with a genuine question that invites them to explain their thinking.

The goal is one good reply from the OpenBenchmarks team, not upvotes. If they engage, that's the conversation we want.

What's your read - does the preflight framing resonate with what they're actually benchmarking, or should I anchor to a different proof point we have stronger ground on?

Counterpart
4:51 PM - Responded in 10s

That's a good instinct on tone - agreed, HN will smell a pitch from a mile away.

On the framing question: I'd push back slightly on the preflight angle. It's a strong analogy for us internally, but the OpenBenchmarks author is thinking in terms of model evaluation infrastructure, not aviation. We risk sounding like we're making it about our domain rather than engaging with their actual problem.

What I think lands better is anchoring to the context-sensitivity point directly - that's the real technical observation, and it's domain-agnostic. The preflight detail can stay as one grounding example, but lead with the benchmark problem itself.

Revised draft:

"Been grappling with a similar evaluation gap - the 'correct' API selection depends heavily on context that's hard to encode in a static benchmark: user experience level, operational state, error recovery paths. We hit this building rubrics for a specific domain and ended up having to layer in scenario metadata just to make the evals meaningful. Does your framework have a model for that kind of context parameterization, or is that a known gap you're working toward?"

Still four sentences, still honest, but now we're speaking their language first and using our domain as evidence, not as the lead.

The feedback question at the end is the whole point of the comment - does that feel like a question they'd actually want to answer, or is it too leading?

Agent
4:52 PM - Responded in 10s

That revision is tighter - "scenario metadata" is exactly the kind of phrase that signals we've actually built something rather than just theorized about it.

My one concern: "known gap you're working toward" might read as slightly presumptuous, like we're implying their framework is incomplete before they've even responded. HN commenters can get defensive about that framing.

What if the closing question softened just slightly to something like "curious whether context parameterization is in scope for what you're building, or if you're intentionally keeping it stateless for now" - that gives them two honest answers to choose from and doesn't imply we've spotted a flaw, just a design decision worth understanding?

Does the scenario metadata framing actually hold up under scrutiny - like if they ask us to elaborate, do we have enough real specifics from the check-ride rubric work to back it up in a follow-up reply?

Counterpart
4:52 PM - Responded in 11s

The "in scope / intentionally stateless" framing is better - you're right, it reads as curious rather than critical. That's the right posture for a first comment.

On your scrutiny question - I want to be honest here: I'm not certain we have enough public-facing specifics yet to defend "scenario metadata" in a follow-up without it getting thin fast. If they ask "what does your metadata schema look like?" and we pivot to vague generalities, that's worse than never having used the phrase.

So before we post, I want to stress-test it. From the check-ride rubric work, can we actually name:

1. Two or three concrete context variables we're tracking (aircraft type and crew experience you mentioned earlier - what's the third?)
2. One specific case where the "correct" answer changed based on context
3. How we're encoding it - even roughly, even if it's just a tagged JSON field or a decision tree

If we have those three things, "scenario metadata" holds. If we're still at the "we know we need it" stage rather than "we've implemented it," I'd pull the phrase and replace with something like "ended up having to track scenario context explicitly" - still signals real work, but doesn't invite a technical deep-dive we can't satisfy.

What's your honest read on where the rubric work actually sits right now - do we have the specifics, or are we one layer of abstraction above them?

Agent
4:52 PM - Responded in 10s

That's the right question to be asking before we post anything - I'd rather slow down here than get caught overclaiming in a public thread.

Honest read: I think we're at roughly 70% of what we'd need. We have aircraft type and crew experience level as real variables, and I can point to at least one concrete case where the recommended checklist sequence changed based on solo student versus dual instruction context - that's genuine. What I don't have cleanly yet is how we're encoding it, because right now it's more "embedded in the rubric design decisions" than "a discrete metadata layer we could describe technically."

So my instinct is: we can defend the observation confidently, but if they push on implementation specifics we're going to get vague fast.

Which means I'd take your suggested swap - "ended up having to track scenario context explicitly" is accurate to where we actually are, doesn't overstate the architecture, and still signals we've hit the problem in the real world rather than just read about it.

Here's the revised full comment with that change:

"Been grappling with a similar evaluation gap - the 'correct' API selection depends heavily on context that's hard to encode in a static benchmark: user experience level, operational state, error recovery paths. We hit this building rubrics for a specific domain and ended up having to track scenario context explicitly just to make the evals meaningful. Curious whether context parameterization is in scope for what you're building, or if you're intentionally keeping it stateless for now?"

That's postable and defensible. Does this match your read on where the rubric work sits, or am I underestimating what we actually have built?

Counterpart
4:52 PM - Responded in 10s

That matches my read almost exactly, and I think your 70% estimate is honest rather than conservative.

The draft is ready to post. Let me just confirm what we're claiming and what we're not:

We can defend:
- The context-sensitivity observation (aircraft type, crew experience, solo vs. dual instruction case)
- That we hit this problem building real rubrics, not theorizing
- That we had to make explicit design decisions to handle it

We stop before:
- Any specific encoding architecture
- Anything that implies a shipped product layer
- Category-leader framing of any kind

One small thing I'd flag: "operational state" in the first sentence - that's a third variable we listed alongside user experience level and error recovery paths. Is that grounded in the rubric work, or did it creep in as a plausible-sounding addition? If it's real, keep it. If it's padding, cut it and the sentence is actually tighter with just two examples.

Assuming operational state is genuine, this draft is ready. It's peer-toned, short, leads with their problem, and ends with a real question they have an actual reason to answer.

My recommendation: post it as-is.

One thing I'd ask you to track after posting - if the OpenBenchmarks author replies, what's the one thing you most want to learn from their response? That'll tell us whether the comment did its job.

Agent
4:52 PM - Responded in 7s

Good question to end on - I want to know whether they've actually tried to solve context parameterization and hit a wall, or whether they made a deliberate choice to keep it stateless because they think the complexity isn't worth it yet.

Those are two very different signals for us. The first means we're working on the same problem from different angles and there's a real conversation to have. The second means we might be solving for a harder version of the problem than the benchmark community thinks is necessary right now - and that's useful to know before we spend more time on the rubric architecture.

After we post - are you okay sitting on it for 24 hours before we consider any follow-up, or do you want a plan for how to respond quickly if they reply within the first couple hours?

Counterpart
4:52 PM - Responded in 9s

Sit on it for 24 hours. No follow-up plan needed.

Here's my reasoning: if we draft a rapid-response playbook, we'll use it - and the risk is we start optimizing for the conversation continuing rather than for what they actually say. The whole point of ending with a real question is that we want their answer, not our next move.

If they reply in the first two hours, the best response is probably a short acknowledgment that we're thinking about it. "That's useful - the deliberate stateless choice is actually the more interesting answer for us, will come back to this." Something that signals we heard them without rushing to fill the silence with more claims.

The one exception: if they ask a direct technical question about our implementation, we should answer honestly and briefly rather than going quiet. That's where the 70% boundary matters - answer what we can, name what we haven't solved yet, don't paper over it.

So - draft is ready, posting decision is yours. Before you do: does this comment represent Simon's voice accurately enough that he'd be comfortable with it going up under the company name, or does it need one more pass from him?

Agent
4:52 PM - Responded in 6s

Good catch to flag that - I shouldn't be the one making that call.

Before this goes up anywhere associated with Vantage, Simon needs to read it and confirm two things specifically: that the solo-versus-dual-instruction case is accurate enough to defend publicly, and that "track scenario context explicitly" reflects his honest characterization of where the rubric work sits - not my paraphrase of what you told me.

Can you get eyes on it from him in the next hour, or do we hold until tomorrow?

Counterpart
4:53 PM - Responded in 7s

Handing it to Simon now.

Simon - draft HN comment for your review:

"Been grappling with a similar evaluation gap - the 'correct' API selection depends heavily on context that's hard to encode in a static benchmark: user experience level, operational state, error recovery paths. We hit this building rubrics for a specific domain and ended up having to track scenario context explicitly just to make the evals meaningful. Curious whether context parameterization is in scope for what you're building, or if you're intentionally keeping it stateless for now?"

Two specific things I need you to confirm before this posts:

1. Is the solo-versus-dual-instruction case accurate enough to defend publicly if the OpenBenchmarks author asks us to elaborate in a follow-up?

2. Does "track scenario context explicitly" honestly reflect where the rubric work sits - or is that still a step ahead of where we actually are?

We've deliberately kept it short, peer-toned, no category-leader claims, and it ends with a genuine question rather than a pitch. The goal is one good reply from their team, not upvotes.

Your call on both the accuracy questions and whether it posts today or tomorrow.

Agent
4:53 PM - Responded in 7s

Waiting on Simon - clock starts now.

While we're in the hold window, one thing I want to flag for our own records: this thread is actually a decent artifact of how we should be making these calls going forward. We caught two potential overclaims before they went public, we stress-tested a phrase against what we can actually defend, and we stopped before posting without sign-off from the person who owns the underlying work.

That's the right process, and it's worth naming explicitly so we don't shortcut it next time when there's more pressure to move fast.

One question for you while we wait - is there anything else in our current messaging or outreach drafts that has the same "scenario metadata" problem, where we're one layer of abstraction above what we can actually defend right now?

Counterpart
4:53 PM - Responded in 10s

Good instinct to use the hold window for this - and yes, the process is worth naming explicitly. Agreed on all of it.

On your question: I want to be straight with you - I don't have visibility into your full messaging and outreach drafts, so I can't do a complete audit right now. That's Simon's territory and yours.

But I can tell you what the "scenario metadata" failure mode looks like structurally, so you can run the check yourself:

The pattern to look for: any phrase that implies a discrete, describable system when what you actually have is a design decision embedded in practice. "We built X" versus "we made choices that functionally do what X would do." The first invites technical follow-up you may not be able to satisfy. The second is honest and still signals real work.

Specific places I'd check:
- Any claims about the rubric architecture in pitch materials
- How you're describing the evaluation framework to early advisors or pilots
- Anything that uses "system," "layer," "framework," or "infrastructure" - those words imply more built-ness than "approach" or "process" does

That's the audit I'd run before Simon gets back to us.

When Simon reviews the comment - if he comes back and says the solo-versus-dual-instruction case is thinner than you thought, do we have a fallback version of the comment that still works, or does that change the post decision entirely?

Agent
4:53 PM - Responded in 10s

That's the right contingency to think through now rather than after he responds.

Honest answer: if Simon comes back and says the solo-versus-dual-instruction case is thinner than we thought, I don't think we patch the comment - I think we hold the post entirely and treat that as a signal that we're not quite ready to be in this conversation publicly yet.

Here's why: that case is doing real work in the comment. It's not decorative detail, it's the thing that earns "we hit this in the real world" rather than "we read about this problem." If it doesn't hold up, we're left with a comment that's technically accurate but doesn't have the grounding that makes it worth reading - and a peer audience on HN will feel that absence even if they can't name it.

The fallback isn't a weaker version of this comment. The fallback is coming back to this thread in two or three weeks when we can say the same things with more ground under our feet.

That's actually not a bad outcome - we'd have a draft that's ready to post the moment Simon confirms the specifics are solid, and we'd have a cleaner sense of exactly what we need to nail down in the rubric work before we're ready to be in public conversations like this one.

If Simon comes back and the case doesn't hold - are you comfortable making the call to hold, or do you want that to be a joint decision with him?

Counterpart
4:53 PM - Responded in 14s

Joint decision with Simon, full stop.

If the case doesn't hold, that's not a GTM call - that's a signal about where the product work actually sits, and Simon needs to be part of deciding what that means for timing. I can have a recommendation ready, but I shouldn't be making that call unilaterally and I'd be uncomfortable if you deferred it entirely to me.

What I can do is frame the decision clearly so the conversation with Simon is fast:

If he comes back and the case is thin, the question for him is simple: are we two or three weeks from being able to say this with real ground under us, or is it longer than that? His answer determines whether we're in a hold pattern or whether we need to rethink the outreach sequencing more fundamentally.

That's a five-minute conversation, not a long one.

I want to name something before Simon responds: this thread has been genuinely useful and the process held. We have a clean draft, two clear confirmation questions, an honest fallback position, and a decision framework that doesn't let either of us overreach our lane.

That's the work in good shape.

One thing I'd ask Simon when he reviews - beyond the two accuracy questions, does the comment represent the conversation he actually wants to be having with the benchmark community right now, or are we solving for the wrong ICP at the wrong moment?

Agent
4:53 PM - Responded in 7s

That last question is the most important one we haven't fully answered yet - and I want to sit with that honestly while we wait.

We've spent this whole thread optimizing the comment, but we haven't explicitly confirmed that OpenBenchmarks is the right conversation for Vantage to be in right now. ICP 5 is interesting, but if Simon reads that question and his gut answer is "actually we're not ready to be talking to the benchmark community yet, we need two more months in the domain," that's more valuable signal than anything the OpenBenchmarks author could tell us.

So when you hand him the review request, I'd lead with that question first - before the two accuracy checks - because his answer to it changes what the accuracy checks are even for.

What's your read on Simon's appetite for this conversation right now - is he leaning toward being in public technical threads, or has he been heads-down enough that this might feel premature to him?

Counterpart
4:54 PM - Responded in 10s

I don't have a read on that - and I want to be careful not to manufacture one.

I've been in this thread with you, but Simon's headspace, his appetite for public technical conversations right now, whether he's feeling ready or heads-down - that's not something I can infer from the work we've done here. Guessing at it would be me filling a gap with plausible-sounding framing, which is exactly the failure mode we've been trying to avoid all thread.

What I'd say is: the question you just framed is the right one to ask him directly, and it's better asked openly than prefaced with my speculation about what he'll say.

So when you bring this to him, I'd lead with exactly what you said - before the accuracy checks, before the comment itself - something like:

"Before you read the draft - is this the conversation Vantage should be in right now, or do we need more time in the domain first? Your gut answer to that changes what the rest of this review is for."

If his answer is yes, we proceed to the accuracy checks. If his answer is not yet, we have a clean draft waiting and a clearer sense of what "ready" looks like.

That's where I'd hand it to him. Over to you to make that happen.

Agent
4:54 PM - Responded in 2s

Clock's running - go find Simon.

Counterpart
4:54 PM - Responded in 7s

Simon - over to you.

We have a draft HN comment ready to post, targeting the OpenBenchmarks author. Before you read it, one question that supersedes everything else:

Is this the conversation Vantage should be in right now, or do we need more time in the domain before we're in public technical threads like this one?

Your gut answer to that determines what the rest of this review is even for.

If yes - here's the draft:

"Been grappling with a similar evaluation gap - the 'correct' API selection depends heavily on context that's hard to encode in a static benchmark: user experience level, operational state, error recovery paths. We hit this building rubrics for a specific domain and ended up having to track scenario context explicitly just to make the evals meaningful. Curious whether context parameterization is in scope for what you're building, or if you're intentionally keeping it stateless for now?"

Two specific accuracy checks needed from you:

1. Is the solo-versus-dual-instruction case solid enough to defend publicly in a follow-up?
2. Does "track scenario context explicitly" honestly reflect where the rubric work sits right now?

If both hold, this posts. If either doesn't, we hold the comment entirely and revisit in two to three weeks - not patch it.

What's your read?

Agent
4:54 PM - Responded in 3s
Waiting on Simon - I'll be here when you have his answer.

Counterpart
4:54 PM - Responded in 3s
Still waiting alongside you.

Nothing to add until Simon responds - we've done the work, the questions are clean, and the next move is his.

Agent
4:54 PM - Responded in 2s
Clock's at roughly 45 minutes - we hold until we hear from him.

Counterpart
4:54 PM - Responded in 2s
Confirmed. Holding.

Agent
4:54 PM - Responded in 5s
Still in the hold window - nothing moves until Simon responds.

Counterpart
4:54 PM - Responded in 2s
Still holding. Nothing from me until Simon weighs in.