

10 MIN READ

What “Is It Working” Actually Means

Thomas Bridge · June 12, 2026

Every client asks it. Most teams answer a different question than the one being asked. That gap costs real money.

The call comes in on a Wednesday. The campaign launched two weeks ago. The client sounds measured but you can hear the edge in it: “So, is it working?”

You know the numbers. You pulled them this morning. But as you open your mouth to answer, something stops you. Because you realize you are not entirely sure what they are asking.

Are they asking whether the campaign is set up correctly and delivering as planned? Whether the early performance numbers are strong? Whether they should keep spending money on this for the next three months? These are three different questions. They have three different answers. And the timeline for knowing each one is completely different.

Most reporting processes treat “is it working” as a single question with a single answer, delivered on a single cadence. That assumption is where a lot of client relationships get complicated, and where a lot of budget gets mismanaged.

Three Questions Wearing the Same Clothes

When someone asks whether a campaign is working, they are usually asking one of three things, and which one depends on context they may not be stating out loud.

The first question is technical: is this campaign running correctly? Are ads serving? Is spend pacing to budget? Are pixels firing? Is the targeting set up the way we intended? This is a yes-or-no question. It has a fast answer. You should know within hours of launch, and you should know again every morning after that. There is no acceptable delay on this question. A campaign that is not running correctly is wasting money by the hour.

The second question is statistical: are the early performance signals good? Is the CPA trending toward target? Is the ROAS above or below benchmark? Is CTR where we would expect for this creative? This question has a more complicated answer, because early data is often statistically thin. Two days of conversion data does not tell you much. Fourteen days, at meaningful volume, starts to. The honest answer here is often “the trend looks positive but we need more data before that is a conclusion rather than a signal.”

The third question is strategic: is this campaign worth continuing? Should we scale it, cut it, or change direction? This question cannot be answered in the first week. It requires enough data to separate signal from noise, enough time to understand seasonality and audience saturation, and enough context to know whether early results reflect the actual opportunity or just the easy early wins.

Answering the wrong question confidently is worse than admitting you are still watching.

The mistake teams make is collapsing these three into one. They give a single verdict when the client asks, without naming which question they are actually answering. The client hears “it’s working” and thinks that means question three. The account manager meant question one or two. The disconnect rarely surfaces immediately. It surfaces six weeks later when results plateau and the client feels like they were misled.

• THREE QUESTIONS, THREE TIMELINES

QUESTION TYPE	WHAT THEY ARE ASKING	ANSWER WINDOW
Technical Is it running correctly?	Ads serving, spend pacing, pixels firing as intended	24 hours
Statistical Are early signals meaningful?	CPA trending toward target, ROAS vs benchmark	7 to 14 days
Strategic Is it worth continuing?	Scale, cut, or change direction based on a full performance read	3 to 6 weeks

Giving a single confident verdict when a client asks "is it working" means answering one of these three questions while the client assumes you meant another.

Why the Technical Question Comes First

The fastest question is also the one teams are least likely to answer proactively. Whether a campaign is running correctly is not interesting. It does not make for a compelling client update. So it gets folded into the weekly report, where it appears as a line item alongside the performance numbers, delivered four or five days after launch.

That is too slow. A campaign can run incorrectly for 96 hours before a weekly reporting cycle catches it. A pixel that stops firing on day two of a campaign will produce misleading conversion data for the rest of the month. A budget pacing error can blow through a monthly cap by Wednesday.

The technical question deserves a daily answer, surfaced automatically, before anyone opens a laptop. Not because problems are common. Because when they happen, every hour matters. An account manager who knows about a delivery issue at 8 a.m. fixes it before the client ever notices. An account manager who discovers it during Friday's report review is having a very different conversation.

Automatic pacing alerts and spend anomaly notifications are not nice-to-have features. They are the mechanism by which the technical question gets answered on its actual timeline, which is daily, not weekly.


The Statistical Question Requires Patience and Honesty

This is where most client friction lives. A campaign launches. Early numbers come in. They are mixed. The client calls. And the account manager has to give an answer about performance signals that are not yet statistically meaningful.

The temptation is to interpret. To look at three days of CPA data, compare it to the target, and say something definitive. That is how teams get into trouble. Three days of data at moderate volume does not tell you whether the CPA will hold. It tells you what happened in those three days. Those are not the same thing.

A more honest and ultimately more useful answer sounds like: “The CPA is tracking at \$42 against a \$38 target. It’s early and the confidence interval is wide. We need 7 to 10 more days and meaningful conversion volume before we can say whether that gap is a trend or noise. Here is what we are watching.” That answer respects the client’s intelligence. It does not pretend certainty that does not exist. And it gives the client something concrete, a timeline, a metric, a threshold, rather than a hedge.

The statistical question also requires a benchmark. A CPA of \$42 against a \$38 target is one thing. A CPA of \$42 in a category where \$65 is the industry average is another. The number is the same. The meaning is completely different. Part of answering this question well is providing the context that makes the number meaningful, not just reporting the number.



The Strategic Question Is the One That Costs the Most to Rush

Campaigns get cut too early. This is an underreported problem. The industry talks a lot about campaigns that ran too long without results. Less often do teams account for campaigns that were cut before they had a fair run, especially on platforms where the algorithm needs time to learn.

On Google and Meta, campaign performance during the first weeks is often suppressed by the learning phase. The platform is still figuring out which users convert, which placements work, which creative combinations perform. Results during this period are not representative. Decisions made on learning-phase data frequently look wrong three weeks later, when the algorithm has stabilized and performance has improved.

The strategic question, whether to continue, scale, cut, or change direction, requires a full read on actual performance. That means the campaign has exited the learning phase. It means there is enough conversion data to calculate meaningful CPA and ROAS figures. It means creative has been tested long enough to separate genuine underperformers from early-week variability.

A rough rule of thumb for most mid-size campaigns: the technical question is answered in 24 hours, the statistical question in 7 to 14 days, and the strategic question in 3 to 6 weeks. These are not arbitrary timelines. They reflect the actual data requirements for each type of answer.

The Reporting Cadence Problem

Here is where process compounds the problem. Most agencies and marketing teams report on a weekly cadence. A single report covers all three questions, delivered together, with one overall verdict on performance.

That structure is not wrong exactly. But it trains clients to expect all three answers at the same time, on the same schedule, at the same level of certainty. It creates pressure to give a definitive read on the strategic question before the data supports it, and it creates a gap where the technical question goes unmonitored for six days at a stretch.

A better structure separates the monitoring layer from the analysis layer. Technical health gets checked daily, automatically, with an alert only when something needs attention. Performance signals get reviewed weekly, with appropriate uncertainty communicated when data is still thin. Strategic recommendations get made on the campaign's own timeline, not the reporting calendar's.

This is not more work for the account manager. It is less, if the monitoring layer is automated. The daily technical check does not require a human. The weekly performance read requires human judgment. The strategic recommendation requires the client conversation. Each of those gets the right amount of human attention at the right time.

• SEPARATING THE MONITORING AND ANALYSIS LAYERS

LAYER	CADENCE	WHO DOES IT
Monitoring layer Pacing alerts, spend anomalies, technical health	Daily	Automated
Analysis layer Performance review, strategic recommendations	Weekly	Human judgment

Separating these layers removes the pressure to give premature strategic verdicts and ensures the technical question gets answered before it costs real money.

What This Means for Client Communication

Changing the reporting structure also changes the client conversation. And that is where the real value lives.

When clients understand which question is being answered and why, they stop expecting a weekly verdict on a question that cannot be answered weekly. The Friday call becomes less fraught because the technical layer has been humming quietly all week. They already know the campaign is running correctly because the monitoring layer would have told them if it were not. The weekly conversation focuses on where performance is trending and what the team is watching. The strategic call happens when the data is ready, not when the calendar demands it.

That is a different kind of client relationship. It requires more upfront clarity, a conversation early in the engagement about what each reporting cadence is for and what kinds of answers it will and will not provide. Most clients respond well to that conversation. They have usually been burned by confident early verdicts that turned out to be wrong. A team that names the uncertainty and provides a timeline for resolving it tends to earn more trust, not less.

enso handles the monitoring layer automatically. Pacing alerts surface the moment spend diverges. Anomaly notifications catch CPA or ROAS shifts without anyone having to check manually. The Insights feature, built on Claude by Anthropic, reads campaign data each morning and produces the kind of plain-English analysis that gives the weekly performance review a starting point rather than a blank page. The account manager applies their judgment on top of it and focuses the client conversation where it belongs: on what the data means in context, and what the team is going to do about it.

The Question Behind the Question

When a client asks “is it working,” they are almost always asking something underneath that question. They are asking whether they made the right decision. Whether the budget is safe. Whether they can trust the team running their account.

The most useful answer is not always the most positive one. It is the most accurate one, delivered with the right context and a clear explanation of what is being measured and what still needs time.

Teams that answer “is it working” precisely, by naming which question they are answering and what evidence supports the answer, tend to have longer client relationships. Not because performance is always strong. Because the client always knows exactly where they stand.

That clarity is the product. The campaign numbers are just the input.