AI-First CRO

The Six-Step Stage Analysis Architecture

Diagnosing and Repairing the Revenue Continuum

A White Paper for CROs, CEOs, and GTM Leaders

Joseph Buchanan
Author of The AI First CRO

October 2026

aifirstcro.com
buchananj@fracti.onl

Introduction

Most revenue systems leak in plain sight. Stages have names. Tools have logins. Dashboards have numbers. Yet the same accounts that looked healthy at signature disappear at renewal, and the same volume of activity produces fewer conversations than last quarter. The problem is not a shortage of data. It is the absence of a shared architecture for seeing the whole continuum and deciding where to repair it.

The six-step stage analysis is that architecture. It is not a research project and it is not a platform selection process. It is a one-afternoon method for mapping current reality, scoring the stages, locating the imbalance between investment and impact, finding the force multipliers, prioritizing the work, and writing the result on a single page a CEO can read aloud. When the steps are treated as an architecture rather than a checklist, they become the operating system for every later AI decision.

The Six Steps as Architecture

An architecture has a fixed sequence and a clear purpose at each layer. The six steps do the same work for a revenue continuum:

Step	Name	What it produces
1	Map current reality	Primary metric, tools, spend, conversion, and time in stage for each motion. Unknown is a finding.
2	Score each stage 1–5	Performance, tool strength, and data quality. Circle the two lowest totals. Those stages go first.
3	Investment vs. impact	Left side versus right side: percent of revenue influenced, percent of AI/MarTech/SalesTech spend, percent of team capacity.
4	Find force multipliers	Stages where a 20–30% lift creates the largest downstream effect; unconnected data; stages still run almost entirely by humans.
5	Prioritize the roadmap	Priority Score = (Revenue Impact × Feasibility) ÷ Current Tool Maturity. One Now item. Dual-tracking three Now items is how diagnostics die.
6	Diagnostic one-pager	Heat map, weakest two stages, top three AI bets for 90 days, expected dollar impact, owner, and the metric the team will watch.

Figure 1 — The six-step stage analysis architecture. One afternoon. Evidence, not opinions.

Step 1 — Map current reality

For each motion—Awareness, Education, Selection/commit, Onboarding, Adoption, Expansion, and Renewal/save—write what is true today: the primary success metric, the tools in use, the money and headcount dedicated to the stage, the conversion rate into the next stage, and the average time spent in the stage. Walk one recent won deal and one lost or churned account through every row. Where the story breaks is a map gap.

Step 2 — Score each stage

Rate every stage on three dimensions: performance (how well we convert or retain here), tool strength (how effective and modern the tools are), and data quality (how clean, accessible, and actionable the information is). Circle the two lowest totals. Those stages enter Steps 3–5 first. A low data-quality score at Awareness should scare a leader more than a mediocre tool score later. Bad data travels. A weak tool sits still.

Step 3 — Calculate investment versus revenue impact

Separate the left side (Awareness through Commit) from the right side (Onboarding through Renewal). For each wing calculate the percent of total revenue influenced, the percent of AI/MarTech/SalesTech spend, and the percent of team capacity. The classic imbalance is heavy investment on the left and under-investment on

the right, even though roughly 80 percent of long-term revenue comes from the right side. Then look one layer deeper: how much of the left-side spend is processing accounts that should never have entered? That is not investment in Awareness. It is a tax the rest of the continuum pays.

Step 4 — Identify force multipliers

Ask three questions. Which stages, if improved by 20–30 percent, would create the largest downstream revenue impact? Where do we have data that could be connected across stages but currently is not? Which stages are still run almost entirely by humans with almost no AI leverage? A cleaner ICP and a scrubbed list is often the cheapest force multiplier on the left, because every later stage inherits it. A signals layer that connects usage, invoices, and conversations is often the cheapest force multiplier on the right, because the data already exists and nobody is correlating it.

Step 5 — Prioritize the AI roadmap

Rank opportunities with a simple formula: Priority Score = (Revenue Impact × Feasibility) ÷ Current Tool Maturity. Focus first on high-impact, under-tooled stages, especially Onboarding, Adoption, Expansion, and Renewal—unless Awareness data quality is so poor that the rest of the map is fiction. In that case, spend the first 8–16 hour window on the intake valve. Dual-tracking three “Now” items is how diagnostics die in committee. One Now item. Then move.

Step 6 — Build the diagnostic one-pager

Summarize everything on a single page: the current heat map, the biggest gaps (usually the right side), the top three AI bets for the next 90 days, and the expected revenue impact of those bets. Include one line on data quality at the intake and one line on Renewal visibility. Close by reading the one-pager aloud. If the room cannot say it in under two minutes, the diagnostic is not finished.

No Data Is Still Data

The most common reaction to a blank cell is to skip it. The architecture treats a blank cell as a finding. If a stage cannot name its primary metric, its conversion rate, or its time in stage, write “unknown.” That absence is information. It tells you the stage is running on opinion, habit, or a dashboard that no longer matches reality.

No data for a stage input or output is still data. It means the handoff is invisible. It means the owner cannot defend the quality of what leaves the stage. It means any later AI tool will optimize a story that has never been written down. In the diagnostic, an unknown conversion rate at Onboarding is more actionable than a precise but low conversion rate at Awareness, because the unknown rate is the place where the system is still guessing.

Leaders who treat missing data as a reason to wait will keep waiting. Leaders who treat missing data as the first repair will discover that the act of defining the input and output usually surfaces the leak faster than any new platform.

High-Quality Input Before Stage Repair

Every stage inherits the quality of the stage before it. A dirty list at Awareness becomes wasted meetings in Education, low-quality opportunities in Selection, and customers who never reach first value in Onboarding. Repairing the later stage while the intake remains dirty is decorating a leak.

The continuum rule is simple: fix the quality of what enters a stage before you invest in optimizing what happens inside it. For Awareness that means a tighter ICP, a scrubbed list, and a definition of a held conversation with a buying-committee member. For Onboarding that means a signed contract with written success criteria, not a signature and a hope. For Renewal that means value that is already visible before the

contract end date, not a scramble in the final thirty days.

AI does not change this rule. It makes the cost of ignoring it visible faster. A signals layer will show you accounts that entered with weak fit and are now decaying. A copilot will show you calls that never confirmed the problem the seller claimed was confirmed. The architecture still requires the human decision to stop the bad input before the repair work on downstream stages can compound.

How AI Creates Leverage in the Repair Process

AI does not replace the six steps. It changes the speed and the quality of the evidence inside them. In the map step, a signals platform or a copilot can surface conversion rates, time-in-stage, and conversation patterns that used to require a week of exports. In the score step, data quality becomes measurable rather than argued. In the force-multiplier step, AI is often the multiplier itself: a tool that correlates product usage with billing status and support tickets can identify the accounts that are already sliding, in days rather than a quarter.

The practical leverage appears in the 8–16 hour window. Most stages can be instrumented well enough to see the leak inside that window. Standing up a copilot that captures every call and email, or a signals layer that connects usage to invoices, simultaneously shows you what is broken and gives the team a shared tape to coach against. The repair is no longer a debate about opinions. It is a debate about the evidence the system already produced.

AI also changes the cost of the imbalance in Step 3. Tasks that once required significant human coordination—scrubbing a list, summarizing a call, flagging an account whose usage dropped 30 percent—become continuous. The formula in Step 5 therefore tilts toward stages where the data already exists but is not correlated. Those are the places where a modest AI investment multiplies every subsequent hour of human work.

The constraint remains time and judgment. AI will surface more candidates for repair than a team can act on. The six-step architecture is the filter: map, score, compare investment to impact, name the force multipliers, prioritize one Now item, and write the one-pager. Without that filter, AI becomes another source of high-confidence noise.

Where to Start

You do not need a new platform to run the architecture. You need a written map of the seven motions, input and output KPIs for the two or three stages that currently leak the most, and one focused block of work to connect the data and see the gap. Most of the return appears in the first 8–16 hours.

Begin with Step 1. Write “unknown” wherever the cell is blank. Circle the two lowest scores in Step 2. If Awareness data quality is the lowest score, fix the intake before you touch any later stage. If the right side is blank, treat the blank as the finding and instrument Renewal visibility first. Then write the one-pager and read it aloud.

The six-step stage analysis is not a methodology that requires a new software category. It is a way of running the revenue job so that the full continuum is visible, the leaks are prioritized by impact rather than by vendor enthusiasm, and the repairs are sequenced so that high-quality input reaches every stage before the stage itself is optimized. Companies that adopt the architecture spend less time recovering accounts and more time expanding them.

This white paper is drawn from Chapter 4 of The AI First CRO. The full diagnostic worksheet, stage-by-stage KPIs, and implementation patterns are developed in the chapters that follow.

aifirstcro.com · buchananj@fracti.onl