

Revenue Growth Is Not The Same As Revenue Quality

Insights | Keystone IQ | Approx. 6 minute read

A target reports 32% customer growth. Revenue up 12%. Churn contained. Average customer spend steady. The deck reads as a compounding story, the data room supports it, and the deal closes at the multiple the growth narrative justifies. Eighteen months into the hold the customer base looks materially worse than the entry snapshot. Nothing was hidden. The growth was real. The value behind it was not.

This is the gap between revenue growth and revenue quality, and PE deal teams pay the difference at exit.

What the headline numbers do not tell you

Commercial due diligence, in its standard form, is built to answer a market question. Is the category attractive, is the share defensible, are the growth drivers credible. It does that well. What it rarely does, with the resolution the question demands, is read the customer base itself, the engine that actually produces the revenue line being underwritten.

Headline metrics are aggregates, and aggregates flatter and conceal. A 32% customer growth rate can describe a deepening franchise or a thinning one whose acquisition is running hot. A stable average customer spend can be the arithmetic mean of customers spending over a thousand euros a year and customers spending nothing. A flat retention rate can hide a base that is replenishing at the wrong end. The deck does not show this. The model carries the assumption forward, and the value creation plan inherits a base it was not designed for.

The OC&C and Gain.pro study of European mid-market PE found that revenue growth accounted for **56% to 70%** of value creation across the sample, depending on hold period and sector¹. With multiple expansion no longer doing the work it once did², the composition of that revenue is now the most under-examined input on most deals.

Operating partners are now putting numbers on what that means in practice. As Slava Shafir, Operating Partner at Corsair Capital, has put it: *“The Bain report is clear: companies need to grow 20% just to meet a 2.5x return. You can’t get there on defence alone.”*³ If returns depend on growth at that rate, whether the growth is real and durable is the central diligence question, not the secondary one.

¹OC&C Strategy Consultants and Gain.pro, *Value Creation in Private Equity: Closing the MOIC Gap*, 2024. Available at occstrategy.com.

²Bain & Company, *Global Private Equity Report 2026*, which finds that multiple expansion has receded as a value-creation lever, leaving operating performance, and the revenue line within it, as the principal source of returns. Available at bain.com.

³Private Equity Wire and AlphaSense, *Value Creation in 2026: The Art and the Science*, April 2026. Survey of 100+ senior leaders in private equity. Available at privateequitywire.co.uk.

Every customer base has a shape

The instructive question is not how much the customer base is growing, but where the growth is sitting and what it is replacing.

Every customer base has a shape. A spine of customers who carry the economics. A middle of customers in transition, some moving up toward the spine, some drifting down toward the tail. A long tail that barely contributes. An intake layer of new customers, a few of whom will become the next spine and most of whom will not. The shape is different in every business, the breaks between these populations sit in different places, and what moves customers between them differs by category, by product, and by what the operator has done with the base over prior years.

The headline never shows the shape. The shape is what determines whether the value creation plan can land.

Reading the natural breaks in the base

The most useful work that can be done on a customer file in diligence is finding the points at which customer behaviour changes character, and reading what those breaks mean for the asset. Not running a pre-set template across every base, but letting the data show where the populations sit and what each one is doing.

Those readings change the deal in ways the headlines do not anticipate. A segment labelled premium high-value can turn out to be more than 40% first-year customers, which is not a settled tier of loyal spenders but a high-value waypoint the file passes customers through. A middle band that looks like a development pipeline for the core can turn out to be populated by customers who do not resemble the core in behaviour, value, or category breadth, and are unlikely to graduate to it. An acquisition engine that looks like it is building the asset can be producing customers at nearly twice the rate it is producing revenue, with most intake landing in the long tail and a sizeable share heading for dormancy before the hold period ends.

None of these readings is available in the blended P&L, and most are not available in standard customer dashboards either. They are available in the shape of the base, once someone has read it.

Why the reading is the work

There is a fair counter-argument: that customer analytics is a solved problem. Off-the-shelf platforms and AI tools produce the metrics, segment the file, and surface the breaks. The numbers can be produced.

The question is what they mean.

A high-value segment that is 41% first-year customers reads as strong on the value cut. It reads as fragile if you have run a business through the moment those first-year customers either bed in or drop out, and you know which way this composition tends to break. A middle

band that is mid-tier in economics reads as a development pipeline if you are looking at potential. It reads as a holding pattern if you have seen, across multiple categories, how often that band fails to migrate. An acquisition cohort delivering 32% volume growth reads as success on the marketing dashboard. It reads as a borrowed two years on the customer file if you have built and dismantled acquisition engines and know what the ratio between customer growth and value growth implies for the base in year three.

This is the **operator judgement layer** that sits above the analytics. AI commoditises the metrics. It does not commoditise the reading of them. The IC question is not *what does the data show*, but *what does the data mean for the plan being signed*, and that is the question an operator answers.

The KIQ approach

Keystone IQ applies **Keystone IQ's Revenue Quality Architecture** to assess customer-base quality across the deal lifecycle. At entry, to inform underwriting and the hundred-day plan. In hold, to direct value creation spend at the parts of the base that will actually move. At exit, to evidence the equity story with customer-level proof rather than narrative.

The work does not start from a pre-set template. It starts from what the data is doing, surfacing the natural breaks in each base: where the behaviour shifts, where the value concentrates, where the headline numbers are being held up or pulled down by populations doing very different things. The interpretation of what each break means for the deal comes from two decades of operating customer bases through those shifts in ecommerce, SaaS, subscription, retail, financial services and media.

What this means for the deal

Customer-base quality is the single largest underwritten assumption in most ICT, SaaS, ecommerce and subscription transactions, and it is the assumption with the thinnest evidence base inside standard diligence.

Two questions are worth asking on every live deal.

First, what does the customer base actually look like underneath the growth rate, and is the value sitting where the model assumes.

Second, is the acquisition engine building the base the thesis depends on, or replacing engaged customers with customers who look nothing like them.

If those questions cannot be answered with customer-level evidence in the diligence file, the gap is not analytical. It is a valuation gap, and it shows up at exit as a multiple compression that the equity story cannot explain.

The cost of getting this wrong is already visible in how firms describe the exit process. Recent industry research finds that **41%** of PE firms struggle to track value-creation indicators through the growth process into the exit story, and a further **35%** struggle to

quantify the impact of the value creation plan they executed⁴. Without customer-level evidence at entry, that evidence cannot be tracked through the hold, and it is not available at exit. The equity story then falls back on narrative, and narrative gets discounted.

Operator-grade customer intelligence

The category Keystone IQ operates in is neither commercial due diligence nor analytics. It is **operator-grade customer intelligence**: the layer that translates customer data into the commercial decisions a board has to take, with the pattern recognition of someone who has taken them before. Built by an operator, for assets where the customer base is the asset.

To discuss how Keystone IQ's Revenue Quality Architecture™ applies to a live or upcoming transaction, contact the team at keystoneiq.co.