

The GE-McKinsey Matrix Explained

Idea In Short

When a company runs a dozen product lines or business units, the hardest question is not where to compete but where to spend the next dollar. The GE-McKinsey matrix answers that question by plotting every unit against two measures: how attractive its industry is and how strong its competitive position is within that industry. Units that score high on both deserve more capital. Units that score low on both are candidates for exit. The matrix will not make the decision for a board, but it forces the conversation into the open, replacing gut instinct and internal politics with a shared, visual case for where resources should go. Executives who skip this discipline tend to keep funding legacy businesses out of habit long after the market has moved on.

Every executive team running more than one business line hits the same uncomfortable question. Which business deserves the next round of investment? Which one has quietly become a drag on the rest of the company? Guessing badly at that question is how conglomerates end up propping up declining divisions. Meanwhile they starve the ones that could actually grow, often because the weak unit still carries political weight inside the organization. The GE-McKinsey matrix, sometimes called the nine-box matrix, was built in the early 1970s to force that question into the open. It remains one of the more durable tools for allocating capital across a diversified portfolio.

What the GE-McKinsey matrix actually measures

General Electric (GE) asked McKinsey & Company to design the framework after a simpler two-by-two model stopped capturing enough nuance. The company ran businesses as different as jet engines and household appliances. One growth chart could not judge both fairly. The result was a 3x3 grid built around two composite scores, not one raw metric per axis. McKinsey's own retrospective on the tool describes it as a way to decide where a multibusiness company should invest its cash, rather than trusting each division's own forecast¹.

The first dimension, industry attractiveness, captures how appealing the market itself is. It judges the market on its own terms, apart from how the company happens to be doing in it. Analysts build the score from market size, growth rate, competitive intensity, the balance of supply and demand and industry-wide profitability. A high score signals a market worth fighting for. A low score signals a market that caps returns, no matter how well a business is run inside it.

The second dimension, business unit strength, captures the company's own competitive position in that market. Analysts measure it through relative market share, brand equity, customer loyalty and the strength of distribution channels. A unit can score high here even in a mediocre industry. It can also score low here even in a fantastic industry and that mismatch is exactly what the tool is built to expose. Plotting every strategic business unit (SBU) against both scores at once produces a picture that no single growth or profitability metric could give on its own.

Units that land in the high-attractiveness, high-strength corner are the stars of the portfolio. They generally deserve reinforced investment to protect and extend their lead. Units in the opposite corner, low attractiveness paired with low strength, are the dogs of the portfolio. They are the natural candidates for divestment or wind-down. The five boxes in between call for more judgment, since they mix genuine promise with real constraints that a single label cannot capture.

Where the framework earns its keep

The most immediate value of the matrix is that it turns resource allocation from an argument into a visual case. Business units sitting in the high-attractiveness, high-strength quadrants become the natural recipients of a larger share of capital. They compete in genuinely strong markets and already hold leading positions there. Directing disproportionate investment toward these units lets a company consolidate leadership in its best markets, rather than spreading capital evenly across a portfolio where not every business deserves an equal claim.

The opposite logic applies to units in the low-attractiveness, low-strength boxes, whether they are outright dogs or question marks with an uncertain future. The matrix's recommendation there is to limit new commitments or exit the position altogether. That frees capital that can be redirected toward the stronger parts of the business. Reallocation like

this is often the hardest conversation a management team has, since a struggling unit frequently has internal champions. Still, the matrix makes the tradeoff explicit rather than hiding it inside a budget spreadsheet.

Beyond individual funding calls, the matrix also helps a leadership team judge whether the whole portfolio is balanced across growth and risk. A healthy spread includes fast-growing units, steadier moderate-growth businesses and slower cash-generating operations. It should not sit entirely in high-growth, high-uncertainty territory. When an acquisition adds a business into a high-growth quadrant, or a unit's improving competitiveness pushes it into a stronger tier, the scatter plot shows that shift right away. Any gap in the portfolio's balance becomes just as visible.

The matrix earns its keep before a deal closes too. Plotting a potential acquisition target on the same grid tests its strategic fit against the acquirer's existing businesses. A target in an attractive, high-strength box strengthens the case for the deal. It might fill a genuine portfolio gap, build presence in a promising segment, or create real synergy with businesses the company already runs. Research cited in Harvard Business Review found that roughly 70% of acquisitions across more than 660,000 deals worth \$56 trillion have succeeded in recent years, a sharp reversal from the failure rates widely cited two decades ago². Much of that improvement traces back to more disciplined target screening before deals are signed. A target in a weak, unattractive box should raise an immediate flag instead, since buying a low-strength business in a poor industry adds exposure rather than reducing it.

The same logic extends to organic expansion into new markets. Scoring industry attractiveness across candidate sectors gives leadership a disciplined way to spot markets worth entering, given the company's existing capabilities. Scoring the company's likely competitive strength in each candidate market surfaces the capability gaps that would need closing first. That dual view keeps expansion plans anchored to markets the company can actually win, not just markets that look attractive in isolation.

Finally, the matrix works as an early-warning system for underperforming units. A unit that consistently lands in the low-attractiveness, low-strength box is signaling a need for a serious turnaround plan or an exit. Large diversified companies often delay that call for years, out of sentimental attachment to a legacy business or plain reluctance to confront a bad result. Cutting off further funding to a persistently weak unit protects capital that would otherwise sink into a business with no realistic path to acceptable returns. It also frees the

organization to redirect that capital toward opportunities genuinely capable of creating value.

Building the matrix step by step

Constructing a usable matrix follows a sequence most strategy teams can complete within a few weeks, provided the underlying data is available.

1. Identify the strategic business units to analyze, at a level that is neither the whole company nor a single product, since a strategic business unit (SBU) with its own customer base, competitive set and profit responsibility is the right level for the exercise
2. Score industry attractiveness for each unit across market size and growth, competitive intensity, pricing trends, cost structure and profitability, technology change and regulation, then combine the ratings into one attractiveness score
3. Score business unit strength across relative market share, brand equity, core capabilities, operational efficiency and management quality, then combine those ratings into one strength score
4. Plot every unit on the resulting 3x3 grid, where the high-high zone holds stars and the low-low zone holds dogs, while the middle boxes hold cash cows generating steady profit in slower-growth markets and question marks that show promise but need a sharper strategy
5. Translate each unit's position into an explicit strategy: increase investment behind a star, optimize a cash cow to sustain its position, set an evaluation window for a question mark, or plan the divestment or liquidation of a dog

A company that treats this as a mechanical scoring exercise, disconnected from how its business units depend on one another, misses part of the value. Reviewing potential synergies across units and revisiting the whole matrix on a set schedule as conditions shift, keeps the tool useful well past the day it is first built.

Where the matrix falls short

No framework this compact captures a company's full strategic reality and the GE-McKinsey matrix carries real limitations that deserve equal attention.

The first is that the matrix was never meant to stand alone. Its two-dimensional view gives a useful snapshot, but a snapshot is not a complete analysis. Market research on customer needs and growth drivers, competitor analysis anticipating rival moves, cost-benefit modeling of specific investment options, statistical risk forecasting and an honest look at cross-unit synergy all belong alongside the matrix, not as an afterthought to it. Wikipedia's entry on the strategic business unit concept notes that the SBU itself, the unit the matrix is built around, is meant to function with its own distinct market segment and profit accountability³. That distinctness is exactly why each unit needs its own supplementary analysis, not one company-wide lens.

The second limitation is subjectivity. Scoring inputs like brand equity, core competencies or long-term industry profitability leans heavily on individual judgment. That judgment is vulnerable to personal bias, recency effects and inconsistent scales across the people doing the rating. A recent product win can inflate a unit's perceived competitive strength well past its real position. An attractiveness score weighted too heavily toward short-term profitability can distort a unit's placement just as badly. Cross-functional scoring, backed by objective data rather than opinion and an outside view from consultants with no stake in the outcome, both help contain this risk. Neither eliminates it entirely.

The third limitation is timing. Every version of the matrix is a static picture of a market that keeps moving and industry conditions or competitive positions can shift meaningfully within a single year. Corporate Finance Institute's guidance on strategic planning frames this as an ongoing cycle, not a one-time exercise, built around continuous evaluation as much as initial analysis⁴. Refreshing the matrix every six to twelve months is a reasonable baseline for most companies. Businesses in volatile sectors such as technology often need a quarterly review to keep the picture current.

The fourth limitation is that the matrix treats every business unit as an island. In reality, units often share platforms, technology, distribution or operations in ways the grid cannot show. Optimizing one unit's position without accounting for those cross-unit ties risks decisions with ripple effects the matrix never surfaces. A company that skips a separate synergy analysis is trading completeness for simplicity. The Boston Consulting Group's own growth-share matrix, a close cousin of the GE-McKinsey approach built around market growth and relative share rather than composite scores, faces the identical criticism. BCG's own retrospective on the framework admits that a two-axis model was always meant as a starting conversation, not a final verdict⁵.

The fifth and most conceptual limitation is that squeezing a genuinely complex business landscape into nine boxes discards nuance. Two units can share the same broad market conditions while differing sharply in their value drivers, their customer segmentation or what their stakeholders expect. None of that shows up when industry attractiveness and competitive strength are the only two things measured. Additional dimensions such as manufacturing economics, the strength of an innovation pipeline, operational risk exposure and talent availability often matter just as much to an investment decision. That is why some organizations add growth rate as a third axis, or adopt sector-specific variants built for their own industry.

None of these limitations argue against using the matrix. They argue for using it as one input among several. Embed it in a broader planning process that draws on real market intelligence, competitive insight and management judgment. Treat the nine-box grid as a starting point, not the final word on where a company's capital should go.

- 1Enduring ideas: the GE–McKinsey nine-box matrix
- 2A better approach to mergers and acquisitions
- 3Strategic business unit
- 4Strategic planning: build a clearer path to business success
- 5What is the growth share matrix?

Summary

The GE-McKinsey matrix earns its place in the strategist's toolkit because it turns a sprawling, emotionally charged portfolio conversation into a structured one. Plot every business unit by industry attractiveness and competitive strength and the case for where to invest, hold or exit becomes visible rather than argued. The framework does not replace judgment. Its inputs require careful scoring, its snapshot needs regular updating and it says nothing about how business units depend on each other. Used alongside market research, competitor analysis and synergy mapping, though, it gives boards and executive teams a common language for portfolio decisions that would otherwise be settled by whoever argues loudest.