



HOW TO BUILD A MODERN PE FIRM

Six changes that separate a firm from a fund.

Claymore Partners · June 2026

Claymore[®]
PARTNERS

Photo: Mitchell Luo / Unsplash

SECTION 01

Twenty years of shooting fish in a barrel

For about twenty years you could run a private equity firm with a spreadsheet, a phone, and a line of credit, and the results would make you look like a genius.

The mechanism was not complicated. You bought a business at a multiple of its earnings, borrowed most of the price at a rate that kept falling, and waited. The cost of debt drifted down, so the same business carried more leverage every year. The multiple drifted up, so the same earnings sold for more than you paid. You returned a number that looked like skill.

The business in the middle of it barely had to change.


Some of the people doing this were very good, and a few were building real value underneath the market's lift. The difficulty is that you could not tell from the outside, and often could not tell from the inside either. The tide was rising under everyone in the harbour, and a rising tide is generous to good boats and bad ones alike.¹

That era is over, and the people who run the numbers now say so. In the golden decade of the 2010s, a typical deal needed roughly five percent of annual earnings growth to produce a standard two-and-a-half-times return over five years. Today, with borrowing costs in the eight to nine percent range, leverage down to thirty or forty percent of the structure, and purchase multiples at records and refusing to expand, the same return needs something closer to ten to twelve percent. Same target. A different engine underneath it.

¹There is a fairer version of the last twenty years in which some firms genuinely built operating capability and earned their returns. A few did. The trouble is that a rising market pays the real operator and the lucky tourist in the same coin, and across a single cycle you cannot tell them apart from the outside. That is the problem this paper is trying to get ahead of.

The growth a deal now has to manufacture

Annual EBITDA growth a typical buyout needs to clear a 2.5x return over five years.



Source: Bain & Company, Global Private Equity Report 2026.

FIGURE 1. *The growth a deal now has to manufacture for itself.* Source: Bain & Company, Global Private Equity Report 2026.

The thing that used to be optional, actually making a business worth more, is now the whole job. Not a slide about full potential that nobody is held to. The work.

This paper is about the firm that does that work, and the six things it builds that the last generation did not.

SECTION 02

Own your operators

Every firm now has operating partners. They sit on the board, join the quarterly review, and send a deck on best practice. That is supervision wearing the costume of operating, and it does not move a business.

Build the operating team inside the firm, and put it inside the company from the day you close.

You bought control. You already have the authority. What most firms lack is the willingness to use it. The change is cultural before it is anything else, from a firm that watches management to a firm that runs the company alongside them. Start with one operator who owns a single function, often a technologist, with a real budget and a real mandate. Add a second only when the first has earned it. Depth in one function beats a thin bench spread across five.

A firm I looked at recently did exactly this. It took a services business nobody would call exciting and put a small technical team inside from the first week. Hundreds of thousands of automated

actions a month, thousands of hours of manual work removed, a quote process that had been impossible to staff now running on its own overnight. That is what operating looks like. Supervision never produces it.²

Owning the work is not the same as building it to last.

The team that owns a function also starves the company's own people of learning it, so the change looks most convincing at the moment it is least durable, because it is being held up by people who are about to leave. Build the team to hand the function over, not to hold it up. If you cannot hand it off, you have not built a capability. You have rented a result, and you give it back at exit.

SECTION 03

Build a brand

Most private equity firms are faceless. The website is an investor deck: a muted palette, a grid of partners, a contact form nobody fills in. Under the old engine that was correct, because the investor was the only reader who paid.

Stop being faceless. Build a brand, and build it for founders, not for LPs.

A brand is a moat, and it compounds. It cannot be bought in a quarter, which is exactly why it is worth owning. It pays out first in origination. The founder brings you the deal before it reaches a banker, because he already knows who you are and what you are known for, and that is a deal you did not have to win in an auction.

It pays out again at the moment of sale. When two firms bid, the one with a name a founder respects tends to win the business a founder built, sometimes at a lower price than the other bid. In a market where the return comes from operating a company well rather than from the market re-rating it, being the owner a founder actually wants is a structural edge, not a vanity.

Doing this properly means running something closer to a media business than a marketing function. The best venture firms understood this a decade ago. One built a newsroom. One built an audience before it had a fund. The point is the sequence: attention first, capital after.

²This is a composite, assembled from a pattern rather than a single firm, with the figures rounded to what is typical for the model. It is meant to be illustrative. I have seen versions of it more than once, which is the point.

The weakness of all this is that the payoff cannot be attributed. You will never prove which deal came from which piece of it, which is why it is the first thing a nervous firm cuts. And a half-built brand is worse than none. A firm that publishes recycled platitudes while telling a portfolio company to fix its own digital execution has described itself without meaning to.³

SECTION 04

Specialise

The generalist financial buyer was the default for a generation. When the return came from the market, the sector did not matter, so there was no reason to choose one.

Pick a lane, and go deep enough that the capability compounds.

Everything else in this paper compounds only inside a domain. An operating playbook, a brand, a data edge, a network: none of them carry across unrelated businesses. The generalist starts every deal from zero. The specialist's tenth deal begins where the ninth ended, and that gap widens with every acquisition.

In practice it means your operators know the sector cold, so the day-one team is expert on arrival rather than learning on the company's time. Your brand means something precise to those founders. Your data benchmarks like against like. Your companies share suppliers, staff, and problems worth solving once. Focus is not one lever among six. It is the precondition for the other five.

The cost is real. Specialisation concentrates risk, so a sector that turns takes your whole book with it, and it caps the size of the fund you can raise. That is why firms specialise less than they should. Assets under management reward breadth, and breadth is the enemy of the depth that actually creates value now.

³The honest weakness of the brand argument is that its payoff cannot be attributed. You will never prove which deal came from which piece of it, which is exactly why it is easy to cut and hard to defend in a bad year. I still think it is among the most durable edges available, for the same reason it is hard. It takes years, and it cannot be bought in a hurry.

SECTION 05

Run the firm on data

The last cycle ran on Outlook and Excel and made a great many people very rich. That is not an insult. It is a fact about a world that no longer exists.

Instrument the portfolio. Build tooling once, reuse it across the book, and benchmark honestly.

Do not force one shiny platform onto every company. That migration is a graveyard, and the firms that promise it are usually the firms that have never attempted it. Play the stack each business already runs, instrument it so the business operates on its own data, and put a thin common layer above the lot so the numbers can be compared. The dashboard demonstrates well. The reconciliation underneath it is the actual work, and it is slow and never quite finished.

Be clear about what data is for. Use it to run the businesses and to benchmark the book. Do not use it to pretend you can source or diligence your way to certainty. The off-market deal still closes on a relationship, and a richer diligence model mostly makes you more confident about a business that is idiosyncratic underneath, which is a trap when confidence tells you to move fast on day one.

More data does not make you right. It makes you confident, which is a different and more expensive thing.

SECTION 06

Make the portfolio a platform

Companies in a portfolio have always met, at the annual chief-executive dinner, over a shared contact or two. That is not a platform. It is a mailing list.

Engineer the network. Move operators, purchasing, and proven playbooks across the book on purpose.

When one company solves something, ship the solution to the others rather than let it die where it was built. Buy as one book, not as fifteen. Move a strong operator from a business that is finished to one that needs them. Benchmark the companies against each other honestly, which only works

if you built the data layer first. The network and the data platform are one piece of infrastructure seen from two sides.

There is a team in baseball that found an edge in the numbers when nobody else was looking, and it won, for exactly as long as the edge was its own. Within a few seasons every front office had the same analysts, the edge flattened into the price of fielding a team at all, and the richest clubs went back to winning. The edge was never the data. It was being the club that ran the play while the others were still printing the slide.

The hard part is that best practice is usually local practice that worked once, in one market with one cost base, and it does not travel as cleanly as the summit slide suggests. Portfolio chief executives have their own incentives and do not enjoy being run as a fleet. Engineering the network is a real job. It is not a dinner.

SECTION 07

Fund the firm like it is the asset

Everything above costs real money, on a horizon longer than a single fund, and returns you can never cleanly attribute to any one deal. This is the lever that decides whether the other five are real or decorative.

Decide, on purpose, to carry that cost, and build the firm so it can.

A cost you cannot tie to a return is the first thing cut in a hard fundraise. The management company can carry the platform, which is a permanent drag on the partners' own economics. Or the portfolio can carry it, which the investors increasingly resent and the regulators have started to circle. Either way it competes against carry on a five-year clock. This is why so many firms have the slide and not the substance. The slide is free. The platform is not.

There is a recursion here that most firms would rather not sit with. The firm that tells every company it buys to invest in capability has to make that investment in itself first, before there is any carry to fund it and with no portfolio to spread it across. The hardest version of the thesis is the one you have to apply to your own firm.⁴

⁴There is a coherent firm that simply never sells, that carries the capability indefinitely on a permanent balance sheet and stops pretending the fund clock applies. It is a sound model, and closer to what the thesis actually implies than

You can already watch the industry straining against its own clock. When a fund reaches the end of its life holding a company it will not sell cheap, it increasingly sells the company to itself, into a continuation vehicle, and keeps holding. These have grown from around eight percent of all sponsor exits to roughly fifteen percent in a few years, a record of more than a hundred billion dollars in 2025.

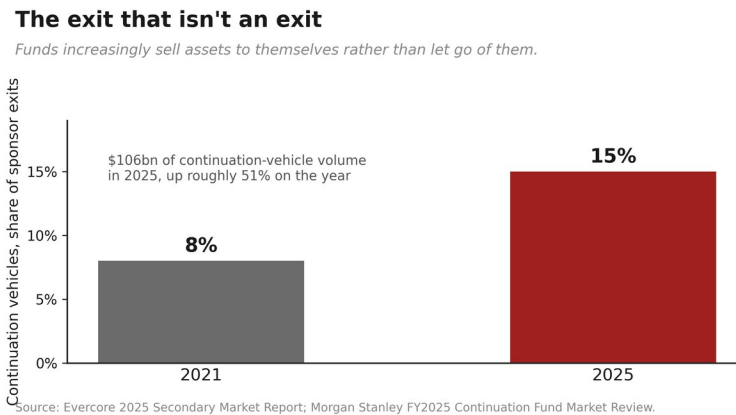


FIGURE 2. The industry is already routing around its own exit clock. Sources: Evercore 2025 Secondary Market Report; Morgan Stanley FY2025 Continuation Fund Market Review.

The closed-end fund is the constraint. Do not wait for the industry to admit it. Build the structure that can carry the platform past a single fund's life, whether that is a longer-dated vehicle, an evergreen sleeve, or a balance sheet of your own, and defend the platform to your investors as the asset it is rather than bury it as overhead.

"Build the firm to outlast the fund, and fund it like the asset it is. Everything else in this paper is a wish until you do."

CONCLUSION

The last winners will not win again

The firms that won the last twenty years will not win the next twenty. That is not a risk to manage. It is the base case.

most firms will admit. It is also a holding company rather than a fund, and it gives up the fee structure that pays the bills today, which is why almost nobody has done it yet.

The returns that built today's marquee franchises came from a market that no longer exists. Cheap debt, cheap entry, and two decades of multiple expansion did the work, and the firms that leaned into it hardest look, in hindsight, like the finest operators in the business. They were not. They were the best-placed passengers on the best decade the asset class has ever had.

The tide that made them look brilliant has gone out, and it is not coming back inside their careers.

The harder truth for them is that the very things that made them successful are now the things holding them back. A generalist book cannot build operating depth. A faceless firm cannot originate. A culture of financiers cannot run a company. A five-year fund cannot fund a platform. Every instinct that won the last cycle is a liability in this one, and instincts are the hardest thing to unlearn while the trailing numbers still reward them.

So the next generation of great firms will not be this generation grown older. It will be firms built from the start to do the work the market used to do for them: firms that own their operators, build a brand, specialise, run on data, engineer the portfolio, and fund the whole of it like the asset it is. The six are not a menu to pick from. They are what the job now demands, together, from one firm.

The engine changed once, when the cost of capital stopped falling. Everything in this paper is what you build in its place. Most firms have changed only the pitch. The firms that will own the next twenty years are changing the engine, and they are being built right now, by people who worked out before the rest of the industry that the old game was over.

"Build the firm the market used to build for you, and fund it like the asset it is. That is the whole job now, and there is no version where you skip it."

APPENDIX

The old firm and the modern firm

The paper is the argument. This table is the reference: the old firm set against the modern one, from what it invests in down to how it behaves.

DIMENSION	TRADITIONAL FIRM	MODERN FIRM
Source of returns	Leverage and multiple expansion	Operational value creation
What is scarce	Capital	Operating capability
Focus	Any sector, any deal	One sector, deep enough to compound
Operating posture	Buy control, then monitor	Buy control, then operate
The operating team	Advisers on the board	Operators embedded inside the company from day one
Authority	Sits on the board, works through the CEO	Works through the CEO, but hands-on in the functions
Who runs the firm	Financiers and ex-bankers	Operators and builders
Sourcing	Waits for the banker's book	Proprietary, brand-led, already known
Brand	A website written for LPs	A media presence built for founders
The buyer you are	One bid among many	The owner a founder wants
Data and tooling	Outlook and Excel	The portfolio instrumented, tooling reused
The portfolio	An ad hoc address book	An engineered network
Exit	Dresses the financials for the process	Hands over capability that survives diligence
Firm cost base	A lean management company	A funded platform
Time horizon	The fund	Longer than the fund



AUTHOR

Lee McCabe
Founder, Claymore Partners

Lee McCabe is the founder of Claymore Partners, a value-creation advisory for private equity sponsors and their portfolio companies. He has 25 years of operating experience across big tech, marketplaces, and PE-backed portfolios.

ABOUT CLAYMORE PARTNERS

Claymore Partners is a value-creation advisory for private equity sponsors and their portfolio companies. We install revenue clarity, operating cadence, and execution ownership inside businesses where the gap between strategy and outcome is costing multiple points of EBITDA. The views expressed represent Claymore Partners' analysis as of June 2026.

claymorepartners.com | lmccabe@claymorepartners.com

This document is published by Claymore Partners for informational and thought-leadership purposes only. It does not constitute investment advice, financial advice, or a recommendation to enter into any transaction. Past performance is not indicative of future results. This document may not be reproduced or distributed without prior written consent of Claymore Partners.