

Where AI Creates **Real Value** in Private Capital

FOR THE LEADERSHIP TEAMS OF BUYOUT FIRMS

Five decisions you face this cycle, the advisory work behind each one, and where AI helps. With sourced numbers, arithmetic you can check, one worked example, and a diligence checklist.

Same 2.5x. Longer hold.

IRR on a single investment at a fixed 2.5x gross multiple, by years held



An illustration, not a fund result: one investment, single entry and exit, no interim distributions. The working is in Chapter 2. Buyout holds at exit averaged five to six years from 2010 to 2021 and are near seven today (Bain, 2026).

EXECUTIVE SUMMARY

The Paper in One Page

Who this is for

Managing partners, chief operating officers, heads of investor relations and chief compliance officers at buyout firms. Funds of funds and institutional LPs have their own applications on page 7, and the checklist on page 9 is written for the questions they ask. Venture firms will recognize the fundraising and diligence chapters; the exit arithmetic is a buyout illustration.

The constraints

Buyout portfolios hold about 32,000 unsold companies worth \$3.8 trillion, and holds at exit average close to seven years. Buyout fundraising fell 16 percent in 2025 to \$395 billion. Distributions ran at 14 percent of net asset value.¹ Meanwhile 71 percent of compliance leaders surveyed in September 2025 said their firms formally use AI, and 28 percent said they had processes in place to test what it produces.⁵ Work is piling up in exactly the places where judgment is scarce, and AI is arriving without the controls around it.

What MUSU does

MUSU is a senior advisory practice. Sam Uddin and Monz Uddin work directly with the partner group: an assessment of where effort and control actually sit, a small number of workflows rebuilt with the people who run them, a controls review, and results measured against a baseline over time. Where the work needs an instrument we build it. Exit Intelligence, MUSU Match and GRC Intelligence support the advisory work. They are not the point of it.

What to expect

Operational results you can check: the share of the portfolio with a current exit-readiness view, partner review time per asset, days from quarter end to investor reporting, rework, unresolved exceptions and their age. We do not claim that AI produces better investment returns or secures commitments. We claim the work around a decision gets faster, broader and better recorded, and that the partner still makes the decision.

Five decisions, one per chapter

- | | | |
|----------|--|--------------|
| 1 | Where should we start?
An assessment that ends in a ranked list with owners, a baseline of measures and a first workflow to fix. | Page 4 |
| <hr/> | | |
| 2 | Which assets do we take to market, and by which route?
An exit-readiness review across the whole portfolio, with the evidence gaps named and assigned. | Pages 5 to 6 |
| <hr/> | | |
| 3 | Which investors deserve our time this raise?
Prioritization with the reasons written down, and a flag wherever the information is thin or old. | Page 7 |
| <hr/> | | |
| 4 | What will we say when diligence asks about our AI?
A controls review, and answers drawn from the firm's own records. The checklist follows. | Pages 8 to 9 |
| <hr/> | | |
| 5 | How do we adopt without losing control, and what does working with MUSU look like?
The method, the engagement, the measures, and how we handle our own conflict of interest. | Page 10 |

Introduction, page 3 · Conclusion and references, page 11 · About MUSU, page 12

INTRODUCTION

The Moment Private Capital Is In

**THE BOTTOM
LINE**

A large unsold stock, longer holds, harder fundraising. And AI use spreading through firms faster than the controls around it.

Buyout portfolios hold about **32,000 unsold companies** worth roughly **\$3.8 trillion**. The average hold at exit is running near seven years, against five to six for 2010 to 2021, and almost 40 percent of companies are now held for more than five years, up from 29 percent in 2019.¹ Exits did improve in 2025: buyout-backed exit value rose 47 percent to \$717 billion.¹ The stock is still very large. Not every one of those companies is a stalled sale. Many are simply being held longer than anyone planned.

Fundraising tightened at the same time. Buyout funds raised \$395 billion in 2025, down 16 percent, and the number of funds closing fell 18 percent across the industry and 23 percent for buyout, a fourth straight year of decline.¹ Distributions as a share of NAV have held below 15 percent for four years and sat at 14 percent in 2025.¹ Allocators want capital back before they commit more. That is not unreasonable.

Then there is AI. Two large surveys show where enterprises stood. McKinsey's fieldwork of July 2024 (1,491 respondents, published March 2025) found that more than 80 percent saw no tangible effect on enterprise-level EBIT from gen AI, while growing shares reported revenue gains and cost reductions inside the business units using it.² MIT's Project NANDA, working from 52 interviews, 153 survey responses and 300 public deployments, reported in 2025 that 95 percent of organizations saw no measurable return.³ The two studies measure different things, neither covers private capital specifically, and the McKinsey data is two years old. What they show is a gap between use and enterprise result. Our reading, and it is ours rather than theirs: the result appears where a use case is tied to a specific piece of work and the data and controls exist to carry it, and it does not appear when the pilot comes first and the purpose is found later.

So this paper is organized around five decisions a buyout leadership team faces this cycle, one per chapter. For each one: the constraint, the advisory work involved, what changes as a result, and where AI helps. The instruments we have built appear where they support that work.

\$3.8T

 UNSOLD BUYOUT
PORTFOLIO VALUE

32,000

 UNSOLD COMPANIES IN
BUYOUT PORTFOLIOS

~7 yrs

 AVERAGE HOLD AT EXIT,
2025

14%

 DISTRIBUTIONS AS A
SHARE OF NAV, 2025

Source: Bain & Company, Global Private Equity Report 2026 (February 2026).¹

HOW TO READ THE NUMBERS IN THIS PAPER

Market figures come from named public reports, with the period and basis stated. Survey findings are reported as what respondents said; where we interpret them, we say the interpretation is ours. Where a number is our own arithmetic, the working is shown. Where we have no data of our own yet, we say so.

CHAPTER 1

Where Should We Start?

THE CONSTRAINT

Effort spread across pilots nobody chose against a piece of work that matters. Eighteen months on, no one can say what changed.

The advisory work here is an assessment, and most of it is conversation. We sit with the partner group, the CFO or COO, investor relations and compliance, and walk through how work actually moves: where partner time goes each quarter, which reports are late and why, where data is re-keyed, where a decision waits on a document nobody can find. Then we look at the data and the controls underneath. The output is a short ranked list of places where a change would be felt, each with an owner, a measure, and a note on what has to be true before anything is built.

The **Private Capital AI Readiness & Value Diagnostic** is the quick way in. It is a seven-minute questionnaire that produces two scores out of 100. **AI Value Opportunity** weights your answers about where time and rework sit, function by function. **AI Governance Readiness** weights your answers about data, policies, testing and accountability. The answers are self-reported, so treat the scores as a structured first read rather than an audit. They tell you which functions to investigate first and roughly how far your controls lag your ambition. They do not establish that adoption is safe, and they do not measure business impact. That takes the assessment, and then time.

Two surveys, two different measures

Share of respondents reporting no enterprise-level result from generative AI

McKinsey, fieldwork July 2024

respondents: no tangible enterprise-level EBIT impact

>80%

MIT Project NANDA, 2025

organizations: no measurable return reported

95%

Findings as the two reports state them. Neither covers private capital specifically. McKinsey also reports that growing shares of respondents saw revenue gains and cost reductions inside the business units using gen AI.

Sources: McKinsey & Company, The state of AI (March 2025; fieldwork July 2024; 1,491 respondents)²; MIT Project NANDA, The GenAI Divide (2025; 52 interviews, 153 survey responses, 300 public deployments)³

Where AI helps at this stage is narrow and useful: reading the firm's own material (policies, board packs, reporting templates, prior memos) fast enough to make the assessment thorough in weeks rather than months, and drafting the first version of the inventory for the team to correct. The judgment about what matters stays with the partners and with us.

WHAT CHANGES

The firm stops discussing AI in the abstract. It has a list, owners, a baseline of operating measures and a first workflow to fix. Anything not on the list waits, on purpose.

A NOTE ON OUR OWN NUMBERS

We launched the Diagnostic this year and have not published aggregate results. When the sample is large enough to mean something we will publish the whole picture, including where firms overrate themselves. Until then the numbers in this paper are other people's, and we say so every time.

CHAPTER 2

Which Assets Do We Take to Market, and by Which Route?


THE CONSTRAINT

A partner can hold the full picture for two or three assets. Not fifteen, and not continuously while conditions move.

Start with arithmetic, and keep it conditional. Take a single investment that returns 2.5 times its cost with nothing paid out along the way. Realized in year five, that is a **20.1 percent IRR**. In year seven, **14.0 percent**. Year eight, 12.1.⁴ Fund-level returns depend on far more than one asset's timing: interim distributions, the mix of multiples across the portfolio, recycling, fees. But if a multiple does not move while the hold extends, time takes the return. Bain's numbers say holds have extended.¹ Bain also notes that a typical deal in the 2010s needed about 5 percent annual EBITDA growth to reach 2.5x in five years, and that today it needs 10 to 12 percent.¹

Same 2.5x, longer hold

IRR on a single investment at a fixed 2.5x gross multiple, by years held



4 yrs **25.7%** 5 yrs **20.1%** 6 yrs **16.5%** 7 yrs **14.0%** 8 yrs **12.1%**

Illustration only: one investment, single entry and exit, no interim distributions. Blue marks the typical 2010 to 2021 hold; orange marks where 2025 exits landed. Not a statement about any fund's performance.

Source: MUSU arithmetic, $IRR = 2.5^{1/n} - 1$.⁴ Holding periods: Bain & Company, Global Private Equity Report 2026.¹

The advisory work is an exit-readiness review, run with the deal team and the CFO. For every asset: is the evidence an acquirer or an IPO committee will ask for actually in hand and current? Which routes are realistic given the fund's remaining term, the market for that sector and the sponsor's own liquidity needs? Where are the gaps, who closes them, by when? **Exit Intelligence** supports this by holding the whole portfolio in one view. It scores each asset on quality, exit readiness, market attractiveness, liquidity pressure, execution and risk, and tests each realistic route (IPO, strategic sale, sponsor-to-sponsor, GP-led secondary or continuation vehicle) on feasibility, value, timing, execution risk, liquidity impact and stakeholder alignment. Every score shows the evidence behind it and the date of that evidence.

The next page shows where AI contributes and where it does not, and walks through one illustrative case from starting information to the action a partner took.

CHAPTER 2, CONTINUED

What the AI Does, What the Rules Do, What the Partner Does

EXPLICIT RULES

Liquidity pressure follows fund age and remaining term. Route feasibility thresholds (scale, sector, leverage) are fixed and written down. Scoring weights are set with the partner group and recorded, so a change of weight is a decision, not a drift.

WHERE AI CONTRIBUTES

Reading board packs, information memoranda, lender reports and management accounts across the whole portfolio; extracting KPIs into one structure; drafting the readiness memo; flagging evidence that is stale, missing or inconsistent between documents.

PROFESSIONAL JUDGMENT

The partner chooses the route, overrides a weight when the market says otherwise, and owns the call. MUSU reviews the reasoning with the team each quarter. Nothing goes to an investment committee without a human author.

What this adds over a spreadsheet: the reading and reconciling of unstructured documents for every asset, every quarter, with the evidence trail kept. The scoring arithmetic itself is not the hard part and never was.

ILLUSTRATIVE EXAMPLE · SYNTHETIC COMPANY, NOT CLIENT DATA

STARTING INFORMATION

Portfolio company C, business services, held five and a half years in a fund entering its eighth year. Eight quarters of board packs, the original information memorandum, monthly lender reporting, and a customer concentration analysis last refreshed two years ago.

HUMAN REVIEW

The deal partner disagreed with the route weighting. Two inbound approaches from sponsors, known to him and invisible to the model, tilted the choice toward sponsor-to-sponsor. The weight was changed and the reason recorded against the asset.

ANALYSIS PRODUCED

Exit readiness scored as **unresolved** rather than low: the last two quarters' KPIs sat in board packs but not in the form an acquirer's data room would expect, and the concentration analysis was stale. Liquidity pressure scored high by rule (fund age). Route view: sponsor-to-sponsor and strategic sale both feasible; IPO not, on scale.

RESULTING ACTION

The CFO was assigned to refresh the concentration analysis and rebuild the KPI pack within 30 days. The route decision was scheduled for the next quarterly review, evidence first. Partner time on the asset that quarter: one 40-minute review instead of assembling the picture from scratch.

Constructed to show the shape of the work. Company, figures and timings are invented; no client information is used. Actual reviews vary by fund terms, sector and the state of the firm's data.

WHAT CHANGES

Every asset has a current readiness view with dated evidence. Gaps have owners and deadlines. Partner time goes to the two or three route decisions that are live this quarter, and the reasoning for each is on file when the LP or the auditor asks.

HOW YOU WOULD KNOW IT WORKED

Share of the portfolio with a readiness view refreshed this quarter (the usual baseline is the two or three assets a partner is closest to). Median age of the evidence behind each score. Evidence gaps closed by their deadline. Partner hours per asset per quarter, before and after. None of these is a return figure. They are numbers a COO can pull in an afternoon, and they are the ones we report against.

CHAPTER 3

Which Investors Deserve Our Time This Raise?

THE CONSTRAINT

An investor relations team can hold a few hundred real conversations in a cycle. The question is which ones.

Buyout funds raised \$395 billion in 2025, down 16 percent, and the number of funds closing fell faster than the dollars did.¹ Allocators are concentrating commitments, and they are doing it while distributions run at 14 percent of NAV.¹ A generic approach to an allocator who is waiting for capital back is unlikely to land.

\$395_B

BUYOUT CAPITAL RAISED
IN 2025, DOWN 16% YEAR
OVER YEAR

-18%

FUNDS CLOSED INDUSTRY-
WIDE IN 2025

-23%

BUYOUT FUNDS CLOSED IN
2025

4th

CONSECUTIVE YEAR OF
DECLINE IN FUNDS CLOSED

Source: Bain & Company, Global Private Equity Report 2026 (February 2026).¹

The advisory work is a targeting review with the head of IR and the managing partner before the raise opens. Which mandates does this fund actually fit? Which relationships are live, and which are owed a conversation regardless of any ranking? **MUSU Match** supports this by prioritizing investors on mandate fit, what is documented about their allocation ranges and pacing, timing signals such as fund cycles and recent commitments, and relationship access. It produces a prioritized list with the reasons for each position written out, and it flags where the underlying information is old or thin. A position resting on a two-year-old allocation statement says so, next to the name.

This is prioritization, not prediction. We have not validated that the list predicts who commits, and we do not claim it does. It cannot read an investor's mood and it does not replace the relationship. What it does is make the reasoning explicit, so the partners can argue with it, and give an investment committee a documented answer to why the pipeline looks the way it does. Where mandate and past behaviour make an anchor conversation plausible, the list says so; whether to have it is the partner's call.

EXPLICIT RULES

Mandate fit and allocation ranges from documented sources; timing from fund cycles; a staleness rule that discounts positions resting on information older than a set age.

WHERE AI CONTRIBUTES

Reading annual reports, mandate statements, meeting notes and correspondence; drafting the rationale for each position; spotting where what an investor said last year and what it did this year disagree.

PROFESSIONAL JUDGMENT

Who to call, what to offer, and whether the anchor conversation happens. The partner's own knowledge of the room always outranks the list.

FOR FUNDS OF FUNDS AND INSTITUTIONAL LPS

The same discipline runs the other way, with different work. **Manager screening:** a documented fit-and-gap view of each GP against the mandate, with the evidence for an investment committee and the gaps to take into diligence. **Portfolio monitoring:** every GP quarterly report read against the previous one and against the fund terms, with exceptions surfaced for a named person to review: a valuation method that changed, a key person who moved, a fee line that grew. For an emerging-manager programme, the checklist on page 9 doubles as the institutionalization agenda.

CHAPTER 4

What Will We Say When Diligence Asks About Our AI?


THE CONSTRAINT

Use is spreading faster than the controls around it, and the questions arrive from LPs before they arrive from regulators.

ACA Group and NSCP surveyed 244 compliance leaders in September 2025, mostly at asset managers and private markets firms. **71 percent** said their firms formally use AI, up 20 points on the year before. **28 percent** said they had processes in place to test or validate AI outputs. **24 percent** had policies and procedures for third-party AI.⁵ These are compliance professionals reporting on their own firms, so read the figures as that. They do not say the other 72 percent never check anything. They do say that formal testing is not yet the norm, and that most firms have no written rule for the vendor whose model does the work.

Use has moved ahead of control

Share of surveyed compliance leaders reporting each at their firm, September 2025



Findings from 244 respondents at firms that are ACA clients or NSCP members. Self-reported.

Source: ACA Group & National Society of Compliance Professionals, 2025 AI Benchmarking Report (October 2025).⁵

Regulators are organizing their questions. IOSCO's May 2026 supervisory toolkit, built with a survey of 21 member regulators, groups its tools under governance and risk management, third-party and outsourcing risk, disclosure, and recordkeeping and reporting.⁶ We use those four headings to organize the checklist on the next page. They are a sensible frame for the questions a firm should be ready to answer, whether they come from an examiner or from an LP's operational due diligence team, and the SEC has already settled cases with advisers that overstated their use of AI.⁷

The advisory work is a controls review. We map where AI touches the firm's processes, who owns each use, what is logged, what is tested, and how an exception (a model output that was wrong, a confidential document that went where it should not) is caught, escalated and closed. Then we fix the gaps in order of consequence. **GRC Intelligence** supports the review and what follows. It monitors obligations the firm already carries (side-letter terms, regulatory commitments, model-risk controls, confidentiality rules) and produces plain-language briefings for the executive accountable for each. In the arrangement we set up, an exception goes to a named reviewer with a due date, and the record of what happened stays with the firm. It does not replace the compliance function or the firm's risk responsibilities. It gives them a current view and an audit trail.

WHAT CHANGES

The firm answers the questions on the next page from its own records rather than from memory, and knows which answers are still thin.

THE CHECKLIST

What Diligence Will Ask About AI, and What a Considered Answer Looks Like

Questions a firm should be prepared to answer, grouped under the four themes in IOSCO's 2026 toolkit.⁶ The answers describe arrangements proportionate to a mid-sized private capital firm; scale them to yours and to what your fund documents already require.

THE QUESTION

WHAT A CONSIDERED ANSWER LOOKS LIKE

GOVERNANCE AND RISK MANAGEMENT

Where does AI touch your investment and reporting processes today?

An inventory of uses, each with an owner, the decisions it informs and what it must not be used for. Reviewed on a set cycle by the executive accountable for operational risk, visible to the oversight the firm already has: a risk committee, the board, or the LPAC where the fund documents provide for it.

Who is accountable when a model-informed decision is wrong?

A named executive for each use, with a documented review path. The investment decision itself stays with the people the fund documents say make it.

How do you test what the AI produces?

A testing routine proportionate to the use: sample size, frequency, what counts as a failure, and a record of what was found and what changed because of it.

THIRD-PARTY AND OUTSOURCING

Whose models run on your data, and under what terms?

Contract terms covering data use, retention, training rights and exit. Vendor due diligence on file for each, refreshed on a schedule.

What happens if a provider changes its model or disappears?

A documented fallback. Model-agnostic design lowers the cost of switching; the replacement still has to be tested and approved before it carries real work.

DISCLOSURE

What do your investors know about your AI use?

Disclosure consistent with the fund documents and applicable obligations, and consistent across investors, so nobody learns something material from a conference panel first.

Does your marketing say more than your records support?

Every public claim traceable to a documented use. The SEC has settled cases over exaggerated AI claims.⁷

RECORDKEEPING AND REPORTING

Can you reconstruct why a decision was made, including what the model contributed?

Logged inputs, outputs, model versions and the human decision, retained on the same schedule as other books and records.

What data leaves the firm, and where does it go?

A data-flow map. Confidential and portfolio-company data segregated or kept in-house, and not in consumer tools.

AND THE ONE THAT DECIDES THE MEETING

What has AI actually changed in how you work?

An answer with operating numbers: review time, reporting turnaround, coverage, rework, exceptions closed. If you do not have one yet, say so and set a baseline now.

CHAPTER 5

How Do We Adopt Without Losing Control?


THE PRACTICE

Senior advisory, delivered by the founders, on a continuing basis. Instruments where the work needs them, and only there.

A firm cannot adopt AI faster than its foundation allows, and the foundation is rarely the model. It is whether the data, the controls and the workflows can carry more work without anyone losing track of who decided what. Our method for building that foundation is the **Trustworthy AI Cycle**: a sand cone of legal, technical and ethical robustness, built from the base up, with five steps that run on top of it and repeat.

THE SAND CONE

Build from the base. Each layer carries the next.



Steps 1 to 5 run on top of the cone, then repeat.

- 1 Consequences and oversight**
Who could be harmed by this use, who watches it, and what stops it. Settled before anything is built.
- 2 Data quality and conformance**
Is the data fit for the decision, and does using it conform to contract, confidentiality and law.
- 3 Principles and metrics**
What good means for this use, written down, with numbers. What cannot be measured cannot be governed.
- 4 Testing and documentation**
Validate against the metrics, record what failed, keep the record where an examiner or LP can find it.
- 5 Monitoring and review**
Watch it in production, review on a schedule, retire it when it stops earning its place. Then back to one.

HOW THE ENGAGEMENT RUNS

Assessment, four to six weeks. Sam and Monz with the partner group, CFO or COO, IR and compliance. Output: a ranked list with owners and a baseline of measures.

Build, one workflow at a time. Monz leads the technical work with the people who run the process; Sam leads governance and operating model. Each workflow goes live with its controls, test record and owner.

Continuing advisory. Monthly working sessions with the responsible executives, a quarterly review of measures against baseline with the partner group, and a controls review each year or when regulation moves.

HOW WE MEASURE THE CONTRIBUTION

Partner review time per asset per quarter. Days from quarter end to investor reporting. Share of the portfolio with a current exit-readiness view. Rework: the share of outputs returned for correction. Unresolved exceptions, by count and age.

These show whether the work got faster, broader and cleaner, not whether AI improved returns or won commitments. We will not make that claim on your behalf.

MUSU is independent and model-agnostic: no fees from model vendors, no preferred platform. Does an assessment run by a firm that also builds instruments end up recommending them? Our answer is the comparison: every recommendation is assessed against what the firm already runs and against third-party alternatives, and the client sees it. Where an existing system or a general-purpose tool does the job, we say so, and our instruments come in only where the assessment shows a gap they fill.

CONCLUSION

Start With a Measurement

NEXT STEP

Take the Diagnostic at musux.ai. Seven minutes, two scores, a short readout by function. Monz will offer a 30-minute call to go through what the scores suggest and what they cannot tell you. No obligation follows.

None of this is about being early to a technology. It is about building a capability that compounds while the people accountable stay in the loop. A firm that starts now enters its next raise with a record of what changed in its own operations and with answers to the checklist drawn from its own records. A firm that waits will answer the same questions from memory.

The pattern across the five decisions is the same. Judgment is the scarce input. The work around it, reading, reconciling, checking, recording, is where the hours go, and it is where AI under control gives some of them back. The partner still decides. The partner just decides with the whole picture in view, and can show how.

TAKE THE DIAGNOSTIC

musux.ai

Two scores out of 100 and a readout by function, returned to you whatever you decide next.

TALK TO US

info@musux.ai

Ask for Monz to discuss your Diagnostic results, or for Sam to discuss an assessment.

REFERENCES

1. Bain & Company, *Global Private Equity Report 2026* (February 2026): "[Private Equity Outlook 2026: Gaining Traction](#)", and [press release](#), 23 February 2026.
2. McKinsey & Company, *The state of AI: How organizations are rewiring to capture value* (12 March 2025). Online survey of 1,491 participants in 101 nations, fieldwork 16 to 31 July 2024.
3. MIT Project NADA, *The GenAI Divide: State of AI in Business 2025* (2025). Basis: 52 structured interviews, 153 survey responses and analysis of 300 public AI initiatives. The 95 percent figure refers to organizations reporting no measurable return.
4. MUSU arithmetic. $IRR = MOIC^{1/n} - 1$ for a single investment with one entry and one exit and no interim cash flows. At 2.5x: 25.7% (4 years), 20.1% (5), 16.5% (6), 14.0% (7), 12.1% (8). Fund-level IRR is not derivable from this illustration.
5. ACA Group and National Society of Compliance Professionals, *2025 AI Benchmarking Report* (October 2025). 244 respondents; data collected September 2025 from ACA clients and NSCP members. The report states a 20-point rise in formal AI use from 2024.
6. IOSCO, *Supervisory Toolkit for AI Use in Capital Markets: Final Report FR/02/2026* and *Standalone Toolkit OR/07/2026* (May 2026). Based on a survey of 21 IOSCO members.
7. U.S. Securities and Exchange Commission, [press release 2024-36](#), "SEC Charges Two Investment Advisers with Making False and Misleading Statements About Their Use of Artificial Intelligence", 18 March 2024.

ABOUT

MUSU, and the People Behind It

MUSU AI is an independent advisory firm for private capital. Sam Uddin and Monz Uddin work directly with GPs, LPs and the institutions around them on where AI helps, how to adopt it under control, and how to show the difference in operating terms. Where the work needs an instrument, we build it and run it on whichever model serves the client.

INDEPENDENT & MODEL-AGNOSTIC

GOVERNANCE-FIRST


FOUNDER-LED



Sam Uddin

**CO-FOUNDER & CHIEF
ADVISORY OFFICER**

Sam leads the advisory practice. His background spans private capital, risk, governance and institutional operating models: senior advisor to a sovereign-backed fund of funds, former Managing Director in capital markets, senior risk advisor to the Capital Market Authority of Saudi Arabia, and senior economist at IOSCO. Dual MBA (Bocconi and Rotman); MSc Financial Economics (SOAS, University of London).



Monz Uddin

**CO-FOUNDER & CHIEF AI
OFFICER**

Monz leads product and AI systems delivery in regulated environments. He has built and shipped AI and data platforms across financial services: VP Business Development at Citi New York, Director of Data Engineering at Sun Life, and Engagement Director at Ai-Zen UAE. MBA (Cornell); BSc Computer Science (University of Toronto).

The instruments that support the work

Exit Intelligence →

A portfolio-wide exit-readiness view with dated evidence behind every score and a tested route view per asset. Supports the review in Chapter 2.

MUSU Match →

Investor prioritization with the reasons written out and stale information flagged. Runs the other way for manager screening. Chapter 3.

GRC Intelligence →

Monitoring of obligations the firm already carries, with plain-language briefings and an exception log for the accountable executive. Chapter 4.

Readiness & Value Diagnostic →

A seven-minute, self-reported first read: two scores out of 100 that point to where to investigate. Chapter 1.

Diagnostic results → info@musux.ai, attention **Monz**

An assessment → info@musux.ai, attention **Sam**

musux.ai · LinkedIn: MUSU AI
Oakville, Ontario, Canada