SharpLogica, Inc
580 Fifth Avenue, Suite 820, New York, NY 10036, USA
(646) 466-7070 | info@sharplogica.com

SHARP LOGICA

Services

Senior technology partner for software companies, PE firms, operating partners, and portfolio company leadership teams.

Software companies are moving fast on AI, modernization, and platform rebuilds. PE firms and operating partners need a senior technology presence they can deploy across the deal lifecycle - before close, during the first 100 days, through the hold period, and ahead of exit. Portfolio CEOs need senior architectural judgment without the cost or commitment of a full-time CTO.

Sharp Logica delivers four senior-level engagements built around those realities. Each is productized, time-boxed, and led personally by Mirano: PhD in Computer Science, 30+ years in engineering leadership, formerly Chief Architect and VP of Engineering at a 350-person engineering organization.

Engagements are not delegated to junior consultants. Capacity is limited by design.

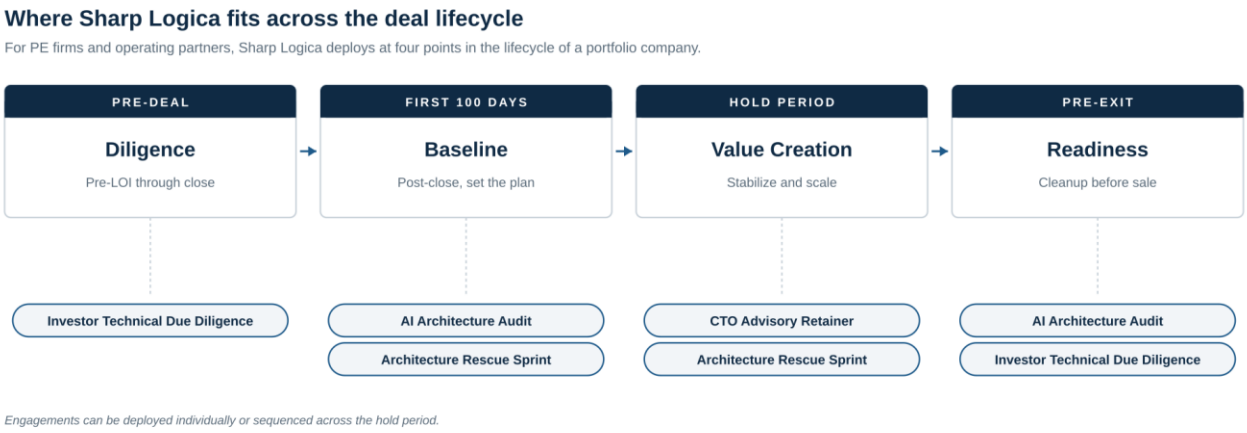

For Private Equity

How PE firms and operating partners deploy Sharp Logica across the portfolio.

Sharp Logica is built around how PE firms actually work. Operating partners do not need another vendor. They need a senior technical resource they can introduce to a portfolio CEO with confidence, deploy on a clear scope and timeline, and rely on for board-ready output.

The four engagements map cleanly to the deal lifecycle. They can be deployed individually for a single situation, or sequenced across the hold period as a portfolio-wide pattern.

Pre-deal

Investor Technical Due Diligence gives the deal team an independent senior read before close: architecture fitness for the thesis, engineering organization health, technical debt scope, AI claim verification, and risks for the SPA, IC memo, or 100-day plan. Rush turnaround available when timeline pressure is real.

First 100 days post-close

AI Architecture Audit creates a production-readiness baseline for portfolio companies that have shipped or are shipping AI features. Operating partners use it to confirm or revise the technology assumptions that informed the deal, and to give the portfolio CEO a concrete starting point for the value-creation plan.

Architecture Rescue Sprint is deployed when post-close diligence surfaces delivery friction or platform fragility that needs to be addressed before the value-creation plan can move forward.

Hold period

CTO Advisory Retainer provides ongoing senior technology leadership for portfolio companies between CTOs, for founders without a CTO, or as a portfolio-CTO-in-residence model where one senior technology presence is placed across multiple portfolio companies on a coordinated cadence.

Architecture Rescue Sprint is also used mid-hold when delivery cadence degrades, an incident pattern keeps repeating, or the platform cannot absorb the next round of features without rework.

Pre-exit

AI Architecture Audit and **Investor Technical Due Diligence** are deployed ahead of sale to identify technical risks that buyers will surface in their own diligence, give the seller time to remediate, and produce a technology narrative that supports valuation.

Engagement model for PE

- Operating partners can request a portfolio review conversation directly. No portfolio company introduction required to start.
- Engagements are scoped per portfolio company. Pricing is transparent and consistent across the portfolio.
- Output is deal-team and board ready by default. Executive summaries, risk maps, and 30/60/90 plans are part of every engagement.
- Confidentiality and information barriers are standard practice. NDAs in place before any portfolio company contact.

OFFER 01

AI Architecture Audit

Production-readiness assessment for AI features and AI-enabled products. Built to be deployed by operating partners on portfolio companies, or commissioned directly by software companies and SaaS operators.

Who it's for

This audit is deployed in two scenarios. The first is a PE firm or operating partner who wants an independent senior read on the AI implementation inside a portfolio company - either as a post-close baseline, mid-hold to validate AI investment, or pre-exit to clean up before sale. The second is a software company, SaaS operator, product team, or technology investor that has built, shipped, or is preparing to ship AI-enabled features and needs to know whether the architecture can operate safely under real usage, real customer expectations, real data, and real cost pressure.

It is especially relevant for companies using LLMs for document analysis, workflow automation, customer support, internal copilots, proposal review, decision support, knowledge search, data extraction, content generation, classification, or other AI-assisted business processes.

The audit is designed for teams that may already have something working, but where leadership still has unanswered questions:

- **Will the cost stay under control as usage grows?**
- **What happens when the model gives a weak, incomplete, or incorrect answer?**
- **Where is sensitive data going?**
- **Who owns prompt changes, model routing, retries, validation, and fallbacks?**
- **Can we observe failures before customers notice them?**
- **Can we roll back safely if a model, prompt, or workflow change breaks behavior?**

This is not a strategy workshop for companies thinking about AI. It is for teams or portfolio companies that need a serious review of an AI feature or AI-enabled product that is already close to production, already in production, or important enough that failure would create business or valuation risk.

The problem

Most AI features are built quickly because the early demo is deceptively easy. A team can connect to an LLM API, write a prompt, return a response, and show something impressive in days or weeks.

But production AI behaves differently.

The real risks usually appear after the demo phase, when the system starts handling more users, more documents, more edge cases, more retries, more background jobs, more data, and more business-critical workflows. What looked like a simple feature can quickly become an uncontrolled operational system.

The problems are often hidden in places leadership cannot easily see. For a PE-backed portfolio company, those hidden problems can also affect the value-creation plan: cost surprises that distort EBITDA, scaling failures that block the growth thesis, security gaps that complicate the next customer expansion, and AI claims that do not survive a buyer's diligence at exit.

Token costs may grow faster than expected because prompts are too large, documents are repeatedly reprocessed, retries are not controlled, or long-context models are used where smaller models would be enough. Latency may become unpredictable because the workflow depends on multiple model calls, external services, file conversion steps, queues, or chained processing stages. Accuracy may degrade because prompts are changed without versioning, model behavior is not validated, or outputs are accepted without sufficient structure and checks.

Security and **data boundaries** can also become unclear. Customer data may move through prompts, logs, temporary storage, embeddings, third-party APIs, or background workers without a clean map of where it goes, how long it stays there, and who can access it. In regulated or sensitive environments, this becomes a serious business risk.

There is also an **ownership** problem. In many companies, AI features are built across product, engineering, data, and leadership without a clear operational model. Nobody fully owns prompt governance, model selection, cost limits, quality checks, fallbacks, escalation paths, or rollback procedures. The system works until it does not, and when it fails, the team is not sure whether the issue is the model, the prompt, the data, the workflow, the queue, the validation logic, or the surrounding architecture.

The result is that leadership may believe they have shipped an AI feature, while in reality they have shipped a fragile workflow with unclear cost, unclear reliability, unclear observability, and unclear accountability.

The AI Architecture Audit makes those hidden risks visible before they turn into customer issues, budget surprises, security concerns, delivery bottlenecks, or exit-diligence findings.

What you get

You receive a structured review of the AI implementation from both an executive risk and engineering architecture perspective.

The review covers the full production path, not just the model call. That includes:

- **Data flow and boundaries.** How user data, uploaded files, extracted text, prompts, model responses, logs, temporary files, embeddings, and stored outputs move through the system.
- **Model usage and routing.** Which models are used, why they are used, where smaller or cheaper models may be sufficient, and whether the system depends too heavily on one model or provider.
- **Prompt orchestration and governance.** How prompts are structured, versioned, tested, changed, reviewed, and connected to business logic.
- **Cost exposure under growth.** Where token spend, repeated processing, long prompts, retries, large documents, background tasks, and chained model calls can create budget risk.
- **Failure handling and fallback behavior.** What happens when the model returns incomplete output, malformed JSON, weak reasoning, timeout errors, rate-limit errors, or inconsistent results.
- **Validation and quality controls.** Whether outputs are checked before being accepted, whether structured responses are validated, and whether human review is built into the right places.
- **Security and data posture.** Whether sensitive data is properly contained, whether access boundaries are clear, and whether logs, prompts, files, and responses create unnecessary exposure.

- **Observability and operations.** Whether the team can see cost, latency, failure rates, retry behavior, model performance, queue pressure, and workflow bottlenecks.
- **Deployment and rollback model.** Whether prompt changes, model changes, workflow changes, and release changes can be deployed safely and reversed quickly when needed.

The deliverables are practical and designed for both leadership teams and PE deal teams or operating partners:

- **Executive report with prioritized risks and concrete recommendations.** A clear explanation of what is safe, what is fragile, what is expensive, and what should be fixed first.
- **Board-ready summary in business language.** A shorter summary for founders, boards, investors, or operating partners, focused on business impact rather than implementation details.
- **30 / 60 / 90-day action plan.** A sequenced plan showing what to stabilize immediately, what to improve next, and what to mature over the following quarter. Designed to slot into a value-creation plan or 100-day plan.
- **AI Cost Exposure Map.** A focused analysis of where token usage, model calls, retries, document size, background processing, and scaling behavior could cause cost growth.
- **Production Readiness Checklist.** A reusable checklist the team can apply to future AI features before they are released.

For PE-deployed audits, output is calibrated for operating partner and IC consumption: the executive summary is structured to support portfolio reporting, and the action plan is designed to integrate with the existing value-creation plan rather than replace it.

Timeline

10 working days from kickoff. Kickoff requires documentation handoff and stakeholder access confirmed in advance.

Typical inputs include architecture diagrams, relevant workflow descriptions, cloud infrastructure overview, model and provider usage, prompt examples, data flow description, deployment process, logging and monitoring setup, and short interviews with product and engineering stakeholders.

The audit can be performed without direct production access if the team can provide sufficient documentation, walkthroughs, and technical context.

Clear promise

In 10 working days, leadership will know whether the AI feature is production-ready, where the highest architectural and operational risks are, and what the team should fix first to make the system safer, more predictable, and more scalable. Operating partners will have a board-ready document they can use to inform value-creation, exit planning, or portfolio reporting.

Price

$15,000 flat for single-product AI architectures. Multi-product or multi-tenant systems quoted separately. Volume pricing available for PE firms deploying across multiple portfolio companies.

Clarity Guarantee. If leadership does not have a clearer understanding of the top AI architecture, cost, security, and operational risks after the audit, one additional executive review session is conducted at no extra cost.

> ### What it is not
> This is not a code review, a security penetration test, or a generic AI strategy workshop. It is a production-readiness audit with a written report, executive summary, and action plan. For PE deployments, it is not a substitute for full technical due diligence; it is a focused AI-specific assessment that complements broader diligence.

OFFER 02

Architecture Rescue Sprint

Hands-on architectural intervention for systems that are slow, fragile, or stuck. Used in the first 100 days post-close, mid-hold to unblock value creation, or whenever delivery is no longer matching the speed the business needs.

Who it's for

This sprint is for software companies, SaaS operators, product organizations, PE-backed portfolio companies, and engineering leaders whose platform is no longer supporting the business at the speed the business needs.

The team may still be shipping, but every release is harder than it should be. Features take longer. Regression defects keep appearing. Production incidents repeat in familiar areas. Developers spend more time working around the system than improving it. Leadership keeps asking why delivery is slow, while the real answer is buried inside the architecture, codebase, dependencies, deployment model, or accumulated technical debt.

It is especially relevant when the company is entering a new growth phase, executing a value-creation plan, onboarding larger customers, integrating an add-on acquisition, expanding the engineering team, preparing for the next round of investment, or trying to scale a product that was originally built for a smaller business.

Common triggers include:

- Release cadence has degraded and nobody can clearly explain why.
- Production defects are increasing even though the team is working hard.
- The platform cannot absorb the next round of features without significant rework.
- The system has become difficult to understand, test, deploy, or change safely.
- Business leadership, the board, or the operating partner has lost confidence in delivery predictability.
- The team is considering a rewrite but is not sure whether that is necessary or dangerous.
- Post-close diligence surfaced platform issues that need to be addressed before the value-creation plan can move forward.

This is not a theoretical architecture review. It is for companies and portfolio companies that need a practical intervention because architectural friction is now affecting delivery, quality, customer confidence, valuation, or growth.

The problem

Architecture problems rarely announce themselves as architecture problems.

They usually appear as missed deadlines, unstable releases, recurring bugs, slow onboarding, unclear ownership, duplicated logic, fragile integrations, unreliable deployments, growing cloud costs, painful testing, and constant engineering frustration.

From the outside, it can look like the team is not moving fast enough. From inside the team, every change feels like fighting the system. Product sees delays. Engineering sees complexity. Operations sees incidents. Leadership and the board see rising cost and falling predictability.

Architectural debt starts looking like delivery failure. Delivery failure starts looking like a people problem. People pressure creates shortcuts. Shortcuts create more architectural debt.

Without a senior outside read, companies often push the team harder, start a risky rewrite too early, or keep patching the existing system without identifying the decisions that are creating repeated friction. For PE-backed companies, this cycle directly threatens the value-creation plan and the exit timeline.

The Architecture Rescue Sprint breaks that cycle. We identify the structural causes of delivery drag, stabilize the highest-risk areas, create a realistic sequence of changes, and leave the existing team with a baseline they can continue executing after the engagement ends.

What you get

Direct senior architectural leadership for 4 to 8 weeks, focused on diagnosing the real causes of delivery and production friction, stabilizing the most important areas, and helping the team establish better architectural and delivery patterns.

This is a hands-on architecture intervention, not a detached audit.

The sprint examines the system across three connected layers:

- **Architecture and codebase structure.** How the system is organized, where boundaries are unclear, where dependencies are tangled, where business logic is duplicated, and where the architecture no longer matches the product model.

- **Delivery and engineering process.** How work moves from idea to release, where handoffs break down, where requirements create rework, where code review or testing slows delivery, and where ownership is unclear.
- **Operational posture.** How the system behaves in production, where incidents repeat, how deployments and rollback work, and whether the team can detect problems before customers do.

During the sprint, we work with engineering leadership and key team members to:

- **Diagnose the root cause** across architecture, delivery process, and operational posture.
- **Establish a stabilization plan** with clear sequencing, prioritization, and ownership.
- **Lead the team through the highest-leverage changes** in real time, without turning the engagement into staff augmentation.
- **Create an architectural baseline** that describes current state, constraints, risks, dependencies, and recommended direction.
- **Hand off a continuation plan** the existing engineering leadership can execute after the sprint ends.

Deliverables include:

- **Architecture Risk Map** showing the areas creating the highest delivery, stability, scalability, or maintenance risk.
- **Technical Debt Triage Matrix** separating what should be fixed immediately, what should be planned, what should be monitored, and what can safely wait.
- **Stabilization Roadmap** covering the next 30, 60, and 90 days with recommended priorities and ownership. Designed to integrate with the value-creation plan or 100-day plan.
- **Delivery Bottleneck Summary** showing where work slows down because of architecture, process gaps, unclear ownership, testing gaps, deployment issues, or decision latency.
- **Architectural Baseline Document** creating a shared reference point for leadership, engineering, the board, and the operating partner.
- **Leadership Debrief** translating technical findings into business impact: delivery predictability, product velocity, customer risk, platform scalability, engineering morale, and future investment needs.

Timeline

4 weeks for focused diagnosis, stabilization planning, prioritization, and initial intervention.

8 weeks when the company needs deeper reinforcement, hands-on support through early architectural changes, stronger handoff, and help embedding new patterns into the engineering team.

Clear promise

In 4 to 8 weeks, leadership and the operating partner will know what is actually slowing the engineering team down, which architecture risks matter most, what needs to be stabilized first, and how to move forward without defaulting to panic, blame, or an unnecessary rewrite.

Price

$25,000 to $50,000. Banded by duration and depth. Advisory-only engagements sit at the lower end. Hands-on architectural leadership with active team coaching sits at the upper end.

Actionability Guarantee. If the sprint does not leave leadership and engineering with a clear understanding of the root causes, highest-priority risks, and next practical actions, one additional architecture and leadership working session is conducted at no extra cost.

> ### What it is not
> This is not staff augmentation. We do not replace the engineering team or write production code as a delivery vendor. We lead the architectural diagnosis, set direction, guide the highest-leverage changes, and help the team that owns the system move forward with clarity.

OFFER 03

Investor Technical Due Diligence

Independent senior technical assessment for PE and VC deal teams. Pre-close diligence, post-close baseline, or pre-exit readiness.

Who it's for

This engagement is for private equity firms, venture capital investors, search funds, family offices, operating partners, technology partners, and investment professionals evaluating a software company before close, establishing a technical baseline immediately after close, or preparing a portfolio company for sale.

It is especially useful when the target or portfolio company has meaningful software risk, a platform that must scale after investment, heavy technical debt claims, AI-enabled product claims, a thin engineering leadership bench, or a roadmap that depends on significant technical execution.

The buyer or seller does not need a generic technical checklist. They need a senior independent read on whether the product, architecture, engineering organization, and technical claims can support the investment thesis or hold up under buyer diligence.

The problem

Software deals move quickly. Management teams present confident roadmaps. CTOs explain the architecture. Product demos look polished. The data room may contain diagrams, cloud bills, security notes, and roadmap slides, but the deal team still needs to know what is real, what is fragile, and what could affect valuation, negotiation, or post-close execution.

The biggest risks are often not visible in the pitch deck. The architecture may work for today but fail under the next stage of growth. A small number of engineers may hold critical knowledge. Technical debt may be described as manageable but actually block the roadmap. AI claims may be based on prototypes rather than production-grade systems. Security and compliance posture may be appropriate for a smaller company but weak for the buyer's future plans.

Without an independent technical read, investors can underprice remediation cost, overestimate scalability, miss key-person risk, or enter the first 100 days without a realistic technology plan. On the sell side, the same risks surface in buyer diligence at the worst possible time.

Investor Technical Due Diligence gives the deal team a clear, senior-level view of the technology risks, execution constraints, and value-creation implications before the decision is made, before the post-close clock starts, or before the asset goes to market.

What you get

An independent senior assessment calibrated to the questions investors actually need to answer: Can this platform scale? Is the engineering organization healthy? Are the AI claims real? What will need to be fixed after close? Which risks should affect valuation, negotiation, integration, the 100-day plan, or the value-creation plan?

The assessment covers:

- **Architecture and platform fitness** for the next stage of growth, including scalability, modularity, deployment model, data architecture, integration points, and operational resilience.
- **Engineering organization health** including leadership depth, team structure, delivery discipline, hiring gaps, dependency on key individuals, and ability to execute the roadmap.
- **Technical debt scope and severity** including where debt is harmless, where it slows delivery, where it creates risk, and what remediation may cost.
- **AI and data claims** tested against the actual implementation, not only the demo or pitch deck.
- **Security and compliance posture** calibrated to the transaction, industry, customer profile, and likely post-close expectations.
- **Cloud, infrastructure, and vendor dependencies** including cost exposure, reliability concerns, lock-in risk, and operational maturity.
- **Deal-specific risk flags** for the SPA, investment committee memo, 100-day plan, value-creation plan, or seller's vendor due diligence.

Delivered as a written report with executive summary, risk prioritization, and a working session with the deal team. When needed, the output can also include a short red-flag memo for IC use, a first 100-day technology plan, or a sell-side narrative document.

Timeline

1 to 3 weeks depending on deal size, documentation quality, stakeholder availability, and complexity. Rush turnaround under one week is available with a premium when the scope is narrow and access is ready.

Clear promise

The deal team will have a senior independent read on the technology risks that matter for the decision, with risk flags clearly mapped to the SPA, IC memo, 100-day plan, value-creation plan, or sell-side narrative.

Price

Starting at $25,000. Banded by deal size, scope, access, and timeline pressure. Rush premium quoted on request. Volume pricing available for firms running diligence across multiple targets.

Decision Support Guarantee. If the report does not clearly support a deal decision, negotiation point, or post-close remediation plan, one additional deal-team working session is conducted and the executive risk summary is refined at no extra cost.

What it is not

This is not financial due diligence, legal due diligence, market due diligence, or a generic IT checklist. It is a senior independent technical read calibrated to the questions the deal team has to answer before close, in the first 100 days after, or ahead of exit.

OFFER 04

CTO Advisory Retainer

Ongoing senior technology leadership, on a fractional basis. For founders without a CTO, portfolio companies between CTOs, and PE firms placing a portfolio-CTO-in-residence across multiple investments.

Who it's for

This retainer is for founders without a CTO, PE-backed portfolio companies between CTOs, executive teams that need senior technical judgment, software companies that have strong engineers but lack consistent architecture and technology leadership at the executive level, and PE firms that want to place a senior technology presence into one or more portfolio companies on an ongoing coordinated basis.

It is also useful when the company is making high-impact technology decisions around AI strategy, platform modernization, hiring, vendor selection, cloud architecture, security posture, product scalability, or engineering operating model, but does not yet need or cannot justify a full-time CTO.

The best fit is a leadership team or PE firm that needs a named senior technology voice in the room on a consistent basis, not occasional hourly advice with no continuity.

The problem

Many companies do not need a full-time CTO immediately. They need senior judgment available at the moments when wrong decisions become expensive.

Architecture choices, hiring decisions, vendor contracts, AI implementation plans, roadmap commitments, infrastructure changes, and platform rebuilds often look manageable in isolation. But together, they define whether the company can scale, ship reliably, control cost, and maintain investor or customer confidence. For PE-backed companies, those same decisions define whether the value-creation plan stays on track and whether the asset is exit-ready when the time comes.

Hourly consulting rarely solves this because the advisor lacks context and accountability. A full-time CTO may be too expensive, too slow to hire, or unnecessary at the current stage. The gap

is a consistent senior presence that understands the business, knows the architecture, participates in key decisions, and helps leadership avoid avoidable technical mistakes.

The CTO Advisory Retainer fills that gap with a scoped, ongoing, senior-led engagement matched to how much technology leadership the business actually needs right now.

What you get

Recurring access to senior technology leadership for architecture, AI, engineering process, hiring, vendor selection, platform decisions, roadmap risk, board reporting, and executive-level technology communication.

Depending on the tier, the engagement may include leadership calls, architecture review, product and roadmap input, engineering team presence, hiring panel participation, board-level reporting, vendor evaluation, and technology strategy ownership.

The retainer is designed to provide continuity. We learn the system, the team, the business constraints, and the decision history, so advice is not given in isolation.

Three engagement tiers

Designed to match how much senior presence the business actually needs. Buyers select the tier that fits and can move up or down with 30 days notice.

Advisory	Embedded	Fractional CTO
$8,000 / month	**$14,000 / month**	**$20,000 / month**
Two scheduled calls per week Async availability for architecture and decision questions Document, proposal, and vendor review No implementation work	Everything in Advisory, plus: Weekly leadership presence with the engineering team Active architecture review and roadmap input Hiring panel participation for senior engineering roles Vendor and platform decision support	Everything in Embedded, plus: Named CTO role on cap table or org chart, as appropriate Board attendance and board-level technology reporting Decision authority on architecture, hiring, and platform choices Direct ownership of technology strategy alongside the CEO

Portfolio-CTO-in-residence

For PE firms, the retainer can be deployed as a portfolio-CTO-in-residence model: a single senior technology presence placed across multiple portfolio companies on a coordinated cadence, with the operating partner as the central point of coordination. This works particularly well during transitions between full-time CTOs, for early-hold portfolio companies that are not yet ready for a full-time hire, or for firms that want a consistent technology voice across a cluster of related investments. Pricing for portfolio-CTO-in-residence is structured per portfolio company at the appropriate tier, with portfolio-level coordination included at no additional fee.

Typical focus areas

- **Architecture and platform direction** for systems that need to scale, modernize, or absorb new product demands.
- **AI strategy and production readiness** for companies adding AI features without turning them into uncontrolled operational risk.
- **Engineering leadership support** for founders, CEOs, operating partners, and operators who need a senior technical counterpart.
- **Hiring and team structure** for senior engineering roles, leadership gaps, and organizational design.
- **Vendor and platform decisions** where the wrong choice creates lock-in, cost, or delivery risk.
- **Board, investor, and operating partner communication** where technology risk must be explained clearly in business language.

Timeline

Month-to-month after a one-month minimum. 30-day notice to end. Tier changes are effective the following calendar month.

Clear promise

The leadership team and the PE sponsor get a consistent senior technology voice that improves decision quality, reduces architecture and execution risk, and brings CTO-level judgment into the business without the cost or commitment of a full-time executive hire.

Price

Advisory: $8,000 / month. **Embedded:** $14,000 / month. **Fractional CTO:** $20,000 / month.

> ### What it is not
> This is not staff augmentation, not on-call engineering, and not a full-time CTO at a discount. It is senior judgment, available consistently, scoped honestly to what the business needs right now.

How engagements start

Every engagement begins with a 30-minute Triage Call. No preparation required. By the end of the call, the buyer will have:

- A clear read on which engagement, if any, fits the situation
- Two or three specific risks or opportunities identified directly from the description provided on the call
- A written follow-up summarizing what was discussed and recommended

If there is no fit, the buyer will hear that directly and immediately. Sharp Logica does not pursue engagements where the value is not obvious to both sides.

For PE firms and operating partners, the Triage Call can be conducted with the operating partner directly, with the portfolio company leadership team, or with both. Operating partners can also request a portfolio-level conversation to map where Sharp Logica engagements would be most useful across the portfolio.

About Sharp Logica

Sharp Logica Inc. is a New York-registered technology consultancy led by Mirano. PhD in Computer Science, 30+ years in engineering leadership across software, financial services, healthcare, public sector, manufacturing, and AI-enabled platforms. Most recent role: Chief Architect and VP of Engineering at Authority Partners, where the engineering organization was scaled to 350 people over 20 years.

Sharp Logica runs multiple concurrent engagements through a small team of senior architects. Nothing is delegated to junior consultants, which keeps quality high and capacity deliberately limited.

Engagements are delivered for software companies, SaaS operators, enterprise technology teams, PE firms and their operating partners, venture capital investors, and the leadership teams of portfolio companies across industries.

sharplogica.com