



VALUE GENE
CONSULTING GROUP



Opportunities in Food Manufacturing

Part 3: Misconceptions About Physical Assets

January 22, 2026

By Ozan Ozaskinli, Jordy Lemus and Sercan Aldatmaz

Introduction

Across food manufacturing, performance challenges rarely come from a lack of effort, commitment, or technical know-how. Plants invest, teams work hard, and systems evolve, yet throughput, stability, and cost performance often stall in subtle ways. In our work, these gaps consistently trace back to something less visible than equipment or headcount: **the assumptions that quietly shape how decisions are made on the factory floor**. *Breaking the Misconceptions* is a series built to surface those hidden beliefs and show **how they limit performance long before a problem becomes obvious**.

In the first two articles of this series, we explored how these misconceptions show up in **process design** and **people management**. We saw how unstable standards, invisible variation, and well-intentioned shortcuts quietly erode capacity on the process side, and how equally common assumptions about staffing, experience, and leadership undermine even well-designed systems on the people side. Together, these perspectives highlighted a consistent pattern: **performance does not break down suddenly, it drifts**, shaped by beliefs that feel practical, familiar, and rarely worth challenging.

The **next four pieces turn to the same pattern in physical assets**, where the cost of a belief is paid in minutes, yield, and reliability rather than in arguments. Each **misconception sounds reasonable in isolation**, which is precisely why it survives. **None of these statements are absurd**. The **damage comes from treating them as decision rules**.

The series is built around what happens when those sentences become rules. Small repairs crowd out lifecycle thinking. Improvement becomes bounded by familiarity instead of capability. Bottlenecks get managed from memory while mix, staffing standards, and variability shift underneath. Changeovers get planned as work at the machine while the prerequisites that determine restart, parts, tools, vendor support, remain unmanaged. The result is not dramatic failure. It is a quieter outcome: expensive equipment spends too much of its life recovering, compensating, and waiting, instead of producing.

Misconception 1: “Our budget is tight, so we focus on affordable fixes to keep things running. The big spend can wait.”

In many manufacturing plants, maintenance decisions are framed as acts of discipline. When budgets are tight, leaders gravitate to the “responsible” path: fixing what is broken and postponing larger interventions until they feel unavoidable. It feels prudent, lines run, spend stays modest, and no single choice looks reckless.

Dig a little deeper, and you see the pattern: these "affordable" fixes aren't one-offs. They're recurring because no one asks why the belts break so often or the cutters fail fast. The team assumes "wear and tear" or "old machines" so the spend keeps flowing to the symptoms without touching the cause.

That is the first layer: normalized mediocrity. Breakdowns stop being signals and become features. Operators adapt, maintenance compensates, and eventually the Plant Manager plans around the interruptions.

But the economics don't agree. That “old” machine isn't dying of age; it is being suffocated by fragmented spending. Over the years, the factory has signed off enough invoices for belts, bearings, and emergency overtime to fund a rebuild. The loss stays invisible because it leaves in small increments that never trigger a decision.

Surprisingly, this cash drain is not even the real cost, instability is.

We saw a plant replacing belts every month when the expected life was nine. That gap meant 8 extra stoppages, about 7 hours each. No one investigated. Belts were forced onto full-length conveyors, micro-damaged during installation. The cause wasn't the asset, it was the absence of a standard.

The same blindness shows up in postponed oven overhauls. Leadership rotates, shutdowns get deferred, and years pass. Operators learn to compensate for an oven that overbakes on the left. “Good enough” becomes habit, until customers reject product, sales loses the account, and quality absorbs the blame.

Perhaps the most shocking example is the refusal to invest in proper HVAC. It fails only during extremes, two months a year, it gets treated as optional high investment. But during those months, waste quadruples due to temperature fluctuations affecting the line. Raw material and operators get blamed, the system remains unchallenged.



The line runs smoothly as unevenly baked cookies move into packaging without triggering action. Pictures are based on real examples. To avoid sharing or distributing personal photos of employees and to comply with NDAs, these images have been generated using AI.

These are not maintenance failures, they are symptoms of something deeper: the plant lacks an asset-economics lens. The dangerous irony is that keeping the budget tighter makes this failure harder to detect.

The barriers to big fixes are real, intertwined, starting with short-termism. Large fixes stall not because the problems are unclear, but because incentives and risk are misaligned. Managers are judged on short-term results while bearing personal downside for long-horizon bets. A major CapEx proposal exposes assumptions to scrutiny without protecting the proposer if it fails, making restraint the rational choice.

Limited data literacy intensifies the paralysis. When managers can't quantify variability or lifecycle cost, they can't defend a case under review. The issue isn't scrutiny, it's the higher chance of failing publicly. Cross-functional

collaboration doesn't offset this. Ownership of the economic case is unclear, so involving other teams increases accountability without increasing authority. The personal risk rises, the payoff doesn't.

In this environment, "the machine is old" becomes a convenient justification. Management shifts from maximizing asset value to minimizing personal exposure, and mediocrity becomes the steady state.

Budget discipline is often misunderstood as spending avoidance. Real discipline is financial intelligence, distinguishing an investment from a recurring loss. When "the machine is old" shifts from a diagnosable condition to a default excuse, it signals that costs are not being managed economically. Patching is no longer prudent; understanding becomes mandatory.

What prevents progress is rarely the absence of ideas. It is the absence of a consistent evaluation method, which makes even obvious interventions feel debatable and risky. Put a price on the loss. Compare it to the fix. If the inputs are uncertain, run a quick audit to tighten the range. This shifts the discussion from narratives to defensible tradeoffs.

Teams often need some handholding at the beginning, not to generate perfect data, but to learn the evaluation habit and apply it consistently.

Yes, our clients see maintenance costs drop by 20–30%. But that is merely a byproduct.

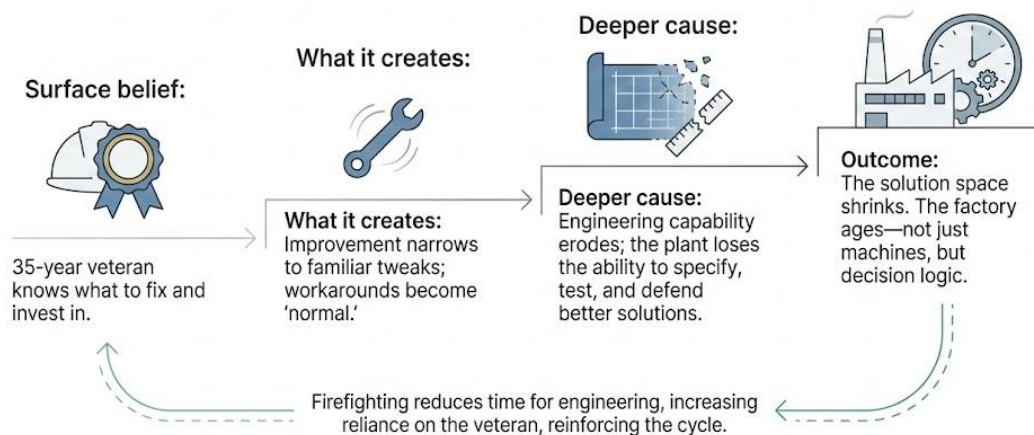
As teams we worked with adopt this lens, the larger gains emerge: higher OEE, lower waste, and output stable enough to earn customer confidence. The objective is not to reduce costs for their own sake, but to restore economic and operational control.

Misconception 2: "Our veteran has been here 35 years, he knows what to invest in and fix."

In many plants, continuous improvement is quietly personified by one person: the 30–35 year veteran who "knows every bolt and bearing." His knowledge is real, and often it is what kept the site alive through years of

constraints and breakdowns. So decisions about what to fix and what to invest in drift toward habit: if anyone knows what matters, it must be him.

When Experience Becomes the Strategy



The issue begins when experience becomes the ceiling for improvement. Workarounds turn into “how the process works,” and minor tweaks pass as progress. Continuous improvement narrows into what the plant already knows how to do, incremental gains that keep production moving but rarely shift throughput, yield, or capability.

We saw this in a plant feeding raw-material bags with a forklift. Every batch required waiting, loading, and positioning at the line, 20–25 minutes of recurring delay. Nothing “broke,” so nothing escalated, yet throughput was quietly capped and variability injected into every shift. The response was pure institutional logic: this is how we’ve always done it, and it isn’t worth automating. A vacuum conveying system could move the same material in two to three minutes and stabilize flow. Because the current method limped rather than failed, the option was never seriously evaluated.

The same pattern on a viscous line where piping partially clogged. The line didn’t stop; slowed, pressure built, and an operator cleared the blockage. It happened often enough to become routine. When the issue surfaced, the

answer was predictable: manual clearing is quick, avoids downtime, and has worked for years, so the plant assigned monitoring and moved on. What remained unexamined was why the clog recurred, and whether pressure, temperature, viscosity, flushing or rebalancing could remove the failure.

What remained unexamined was why the clog kept returning. Pressure, temperature, and viscosity conditions were rarely revisited. Options like automated flushing, inline sensors, or rebalancing were not explored. Experience taught the team how to recover quickly, not how to remove the failure mode. A temporary workaround hardened into an operating model.

In projects like these, the pattern doesn't break through more effort on the floor, and it doesn't break by dismissing veteran knowledge. It breaks when the organization stops using one person's memory as its improvement engine.

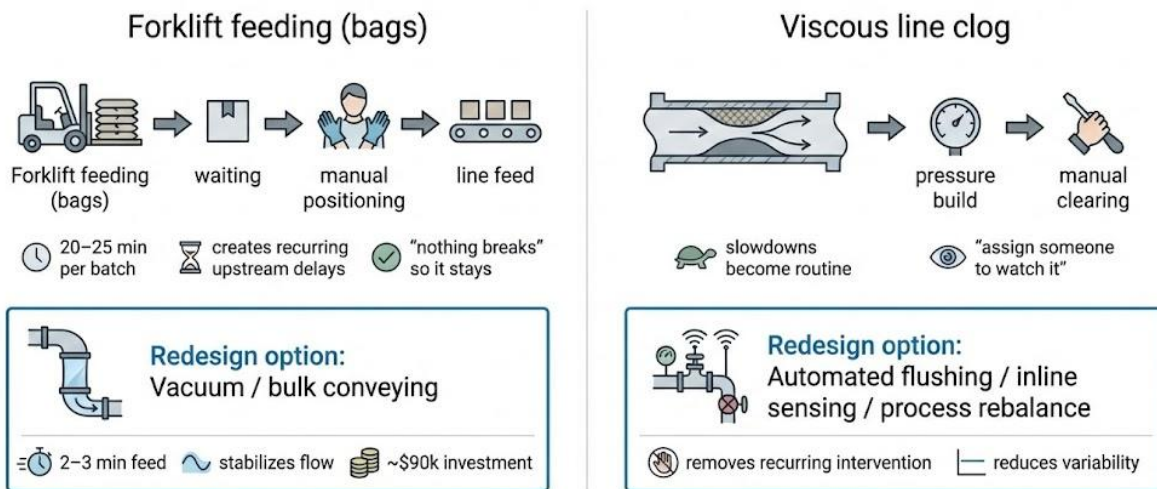
The deeper cause is capability erosion. Many plants treat engineering as overhead rather than leverage. Hiring freezes, attrition, and "just enough technical capacity to keep today running" shrink the ability to specify, test, and defend better solutions. Constraints stay trapped as operational coping mechanisms because there is no bandwidth to convert them into redesign.

Even where engineers exist on paper, the solution space can still collapse if the site stops keeping up with the outside world. Automation, sensing, controls, and process technologies move faster than any single facility can track informally. Staying current requires deliberate exposure, training, benchmarking, vendor conversations, plant visits. Many factories never build that muscle. Training feels optional, so the plant's baseline shifts: not only the equipment ages, but the collective imagination of what is feasible.

This is also where "we can do it ourselves" becomes costly. Outside support gets rejected as expensive without comparing cost to the value of speed, pattern recognition, and avoided dead ends. External specialists see dozens of plants; they know which solutions are mature, which are hype, and what fails under real constraints. Avoiding every seller also means avoiding the market's learning loop, and relying on what the plant already knows.

You see the consequence when a plant finally approves investment. We worked with a site that ordered new machines after a rigorous business case, yet three months before arrival the factory was unprepared: electrical work incomplete, layout unresolved, vendors unmanaged, technical ownership unclear. Under pressure, two engineers quit. Costs inflated through delays and change requests, not because the investment logic was wrong, but because execution capability was missing. With a small specialist team, the layout and sequencing were rebuilt, the project returned to track, and the revised design even enabled two extra machines within the same budget. The project delivered a week late, but with more capability than planned.

Workarounds vs Redesign: Two Plant Examples



The point is not that experience is useless. It is that a 35-year veteran becomes the de facto strategy because the organization has removed alternatives. Fresh engineers get buried in approvals, paperwork, and "that's not how we do it here." The irony is that the same organizations often assume they can simply hire their way out of the gap, despite a tight market and an even tighter pay gap. Over time, the factory becomes old in a deeper sense: not only the machines, but the decision logic.

The way out is to combine experience with modern engineering capability. That means targeted technical audits, selective hiring and retraining,

structured exposure to current technologies, and a simple path for ideas to be tested rather than debated. The goal is not to turn an old site into a greenfield. It is to rebuild the plant's ability to redesign constraints instead of managing around them, so it can compete even when it will never be the newest factory on the market.

Misconception 3: “Our bottleneck is obvious, we invest where it is slow.”

Many plant leaders treat bottlenecks as a fixed fact of the line. Ask what limits output and you'll get a fast answer, usually with a machine name attached. That answer hardens into the basis for improvement and investment, even as product mix and operating conditions evolve.

In practice, the limiting station is conditional on what you run. Different SKUs impose different demands through the same chain, so the constraint can move as the schedule shifts. Plants drift when they keep managing to the bottleneck they learned under an older mix while current volume is governed elsewhere. Forecast noise complicates targeting, but it does not remove the need to test decisions against more than one plausible mix.

Forecast quality complicates this but does not absolve it. If mix drives the constraint and forecasts are unreliable, then investment targeting must be tested against more than one plausible mix. Plants that skip this discipline plan as if mix will hold, then discover—too late—that the constraint in the business case is not the constraint in the schedule that arrives.

The next layer appears where pace is labor-set rather than mechanically fixed. Some steps have no single ceiling, one person yields one rate, two yields another, and the constraint becomes a function of the staffing standard the plant is actually holding. This is where bottlenecks are quietly manufactured. A supervisor reallocates a person, output holds under the current mix, and the reduced pattern becomes “how we run” without anyone recording what it cost or what conditions made it safe. When the mix tilts, the station cannot hold the required pace and the plant experiences a “capacity ceiling” that feels structural but was created by drift. When the original logic

leaves with the person who understood it, the same constraint gets rediscovered, repeatedly, as mix changes.

Once staffing can move the limiting rate, the plant also inherits a tradeoff it often leaves implicit: add people, or redesign the work so the same people can run faster. Many sites debate these options in incompatible currencies, labor as cost, equipment as capability, without translating both into the same unit: incremental sustained sellable units per hour on the SKUs that matter. Without that translation, decisions are inevitably political.

A third layer is variability. Even with the same SKU and staffing, the governing constraint can shift within a run as performance fluctuates across stations. That reality forces operational design choices, buffering, decoupling, and variability reduction, rather than arguments over a single “bottleneck.”

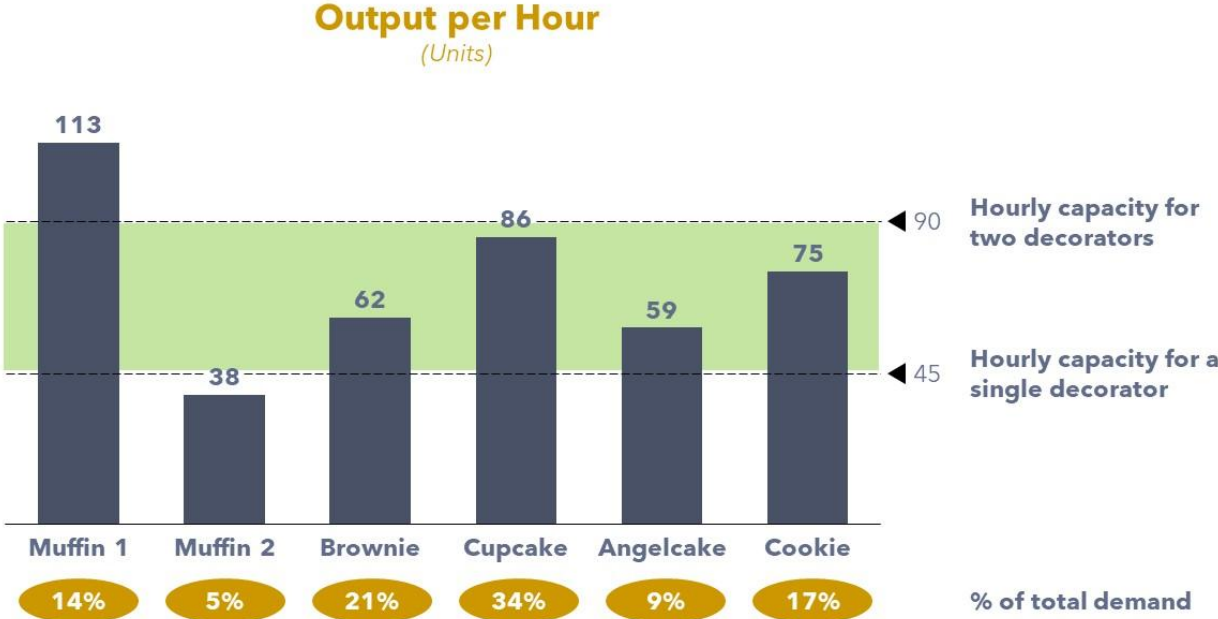
These layers are manageable only with a living constraint view: rates by SKU family, staffing standards at labor-set steps, and a grounded understanding of where variability binds during runs. Many plants don’t keep this current because no one owns it end to end, or because the person who could own it is buried under formulation work, compliance tasks, and daily firefighting. The site runs on memory instead of maintained knowledge.

When that view is missing, capacity definitions degrade. Many sites treat a historical peak hour as what the line “can do,” even though it was tied to a specific SKU and staffing setup. The peak becomes a benchmark anyway, and the organization measures itself against a number that was never validated across mix, standards, and normal variation, producing blame and noise instead of correction.

Only after the basics are held steady do advanced levers become worth selecting. Process redesign, controls, automation, and formula changes can relieve constraints, but they are hard to choose correctly when the plant cannot say, with discipline, what governs sellable throughput under the mixes it actually runs.

We saw the impact in an engagement focused on short-term revenue uplift. The site claimed it was at “100% capacity,” yet demand was higher and shipments were constrained. The limiter was not equipment: the decoration

station was staffed with one person after a reassignment. One decorator capped output at roughly 45 units per hour; two lifted it to about 90. Under the prevailing mix, that gap determined shipments, so the immediate move was restoring the second decorator. After that, the next targets clarified: mixing limited a subset of SKUs, while the oven was the long-lead structural constraint, with interim operating changes needed while the decision matured.



Source: Value Gene Analysis

Smart investing builds resilience. Without spotting bottlenecks, factories fund quick fixes that embed weaknesses, adding machines to maintain, but no extra output. Real competitive edge comes from targeted spending: put capital where it removes chokepoints, smoothing operations and strengthening the plant.

Misconception 4: “The guys know where parts go, and setups take as long as they take.”

In many plants, long changeovers are treated as a given. You hear familiar explanations: the equipment is complicated, people need more training, the

“good” operators can do it faster. As long as the line eventually comes back, the time is accepted as operational reality.

But what’s being accepted is not just duration. It is non-repeatability. When a changeover depends on human recall and last-minute problem solving, the plant isn’t managing a defined return to a ready state; it is managing around uncertainty.

We saw this in a confectionery plant where changeovers routinely ran over 150 minutes. Setup genuinely required knowing what fits where and in what sequence. The issue wasn’t effort. The issue was that the machine had no defined ready state and no structure to make the right assembly obvious. Parts weren’t labeled or keyed, the sequence wasn’t standardized, and the fastest safe restart lived mostly in a few heads. New operators learned by watching and memorizing. So every changeover included searching, choosing, and correcting in real time, minutes that were predictable, yet treated as unavoidable.

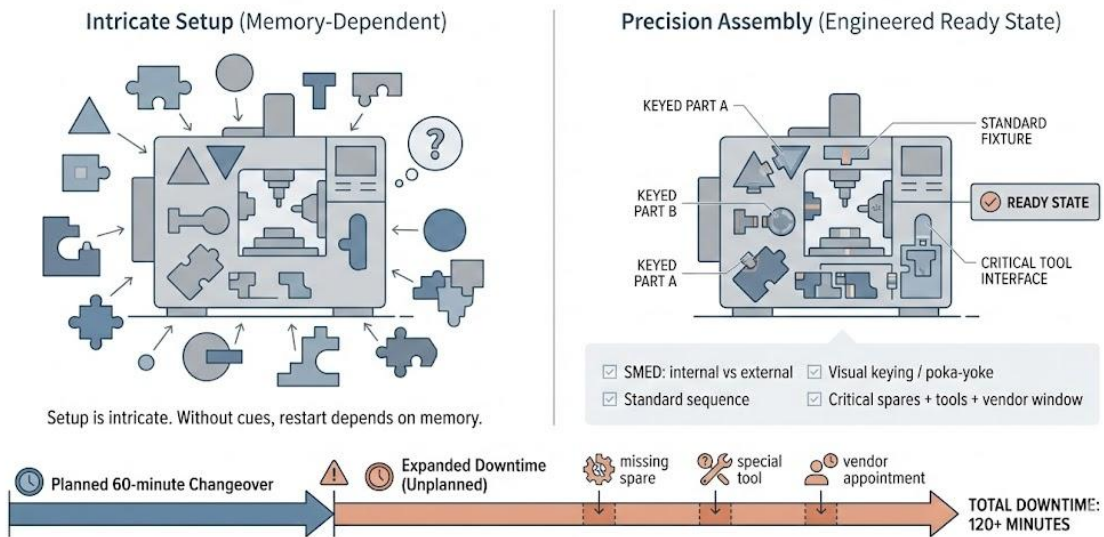
The more costly layer showed up when changeovers were planned but their prerequisites weren’t. The schedule assumed a window, then the line sat because a supposedly available spare was missing, wrong, damaged, or unusable. In one case, a critical item hadn’t been stocked at all, so the planned stop became “order it,” and lead time replaced a timetable. In another, the part was on-site, but it required special mounting equipment and a vendor appointment the plant hadn’t arranged. The mechanical work was scheduled; the ability to restart wasn’t.

That is the conversion: planned downtime turns unplanned when readiness is undefined. The plant plans the choreography at the machine, but it doesn’t govern the dependency chain that makes restart possible, correct parts, correct condition, correct tools, and the external support needed to install and validate. When those are missing, the changeover becomes open-ended by design.

Fixing this requires treating changeover as an engineered system, not an operator performance.

Start with structure at the asset: define “ready to run” in concrete terms, standardize the fastest safe sequence, apply SMED to separate internal and external work, and remove choice from assembly through labeling, keying, and simple poka-yoke. Stage what can be staged before the stop so the line is not waiting on preparation. Fixing this requires more than declaring SMED and printing a checklist. It requires designing the changeover as a system.

Precision Changeover: From Memory to Engineered Repeatability



Then design the prerequisites with the same seriousness. Identify the small set of components that routinely gate restart, and manage them explicitly: stock based on lead time and exposure, not on intuition; verify condition and compatibility; and map installation requirements, special tools, fixtures, software access, vendor windows, so “we have the part” also means “we can use it.”

Training still matters, but it should be the last mile, not the strategy. When structure holds the knowledge and prerequisites are owned, changeovers stop being heroic events. They become what planning assumes they are: bounded, repeatable, and designed to keep the asset producing, minute after minute, rather than waiting on memory, missing parts, or chance.

Authors



Ozan Ozaskinli
Partner and Managing Director
Ozan.Ozaskinli@valuegeneconsulting.com



Jordy Lemus
Implementation Manager
Jordy.Lemus@valuegeneconsulting.com



Sercan Aldatmaz
Principal
Sercan.Aldatmaz@valuegeneconsulting.com

About Value Gene Consulting Group

Value Gene Consulting Group is a distinguished boutique consulting firm specializing in delivering strategic business solutions that yield significant, swift, and sustainable outcomes. Our dedicated team collaborates closely with C-level executives, providing expert guidance tailored to mastering business challenges within the Food and Consumer industries.

In the ever-evolving landscape shaped by our clients' needs, we prioritize sound strategy and decision-making as cornerstones of enduring success. Our approach is grounded in fact-based quantitative and qualitative analysis, fostering positive change in the best interest of our clients and their stakeholders.

As a boutique management consulting company, we stand out by leveraging the unique skills of our enthusiastic team. Our consultants, with prior experience in top-tier strategy firms, bring a result-oriented focus to decision-making and business management.

Embodying our 'boutique service principle,' we ensure heightened responsiveness, a long-term commitment from our team, and high-quality advice with direct involvement of our senior team in day-to-day operations. Remarkably, over 90% of our business originates from longstanding client relationships, showcasing our dedication to our clients.

At the core of Value Gene Consulting Group is a consulting team comprising top-educated and globally experienced members. With more senior involvement than industry standards, we consistently produce immediately applicable results. Our deep subject expertise, coupled with pioneering industry knowledge, guarantees impactful and quality work.

Our distinctive approach involves working collaboratively with client organizations, fostering a partnership that goes beyond traditional consulting. We are catalysts for change, driving transformation within our clients' businesses by connecting analytics understanding to actionable business insight.

Our success is measured by our ability to maintain enduring client relationships, exhibit client responsiveness, and demonstrate unwavering dedication. Value Gene stands apart in the industry, delivering the content-driven insights that our clients seek from their strategic advisor.

Disclaimer

This report, prepared by Value Gene Consulting Group, is intended for informational purposes only and is based on the data available up to the date of its publication. The analysis and insights provided are subject to change without notice and may not be exhaustive. While every effort has been made to ensure the accuracy of the information presented, Value Gene Consulting Group makes no representations or warranties of any kind, express or implied, about the completeness, accuracy, reliability, suitability, or availability with respect to the report or the information, products, services, or related graphics contained in the report for any purpose.

The content of this report is proprietary to Value Gene Consulting Group, and unauthorized use, reproduction, or distribution of any part of this report is strictly prohibited. Any reliance you place on the information presented in this report is at your own risk. Value Gene Consulting Group shall not be liable for any loss or damage, including without limitation, indirect or consequential loss or damage, or any loss or damage whatsoever arising from loss of data or profits arising out of, or in connection with, the use of this report.

This report does not constitute professional advice, and users are encouraged to seek independent professional advice before making any business decisions. Value Gene Consulting Group disclaims any liability for actions taken or not taken based on the content of this report.

Value Gene Consulting Group reserves the right to update or revise the information contained in this report at any time without notice. Any changes made to the report after its publication will be considered as part of the ongoing analysis and research process.

By accessing and using this report, you agree to these terms and conditions.