WHITE PAPER

# Competitive Advantage in the Agentic AI Era

## Key findings from leadership conversations with 14 insurers across Asia-Pacific

Australia · Hong Kong · Japan

**March 2026**
By Davide Corradi, Semih Durmus, Giuseppe Franco, Joe Khoury, Swayamjit Mishra, and Eunkyung Koh

In February 2026, BCG held in-depth working sessions with the C-suite leadership of 14 insurers across Australia, Hong Kong, and Japan to assess where their organizations stood on AI-driven transformation. Collectively, these organizations represent over $300 billion in combined market capitalization and more than $150 billion in combined gross written premium.

Here are the three main findings from our conversations:

1. **Agentic AI represents a qualitative shift from generative AI, not an incremental improvement on it**. Insurers are moving past experimenting with use-cases for GenAI, and they are now envisioning end-to-end business transformations with AI agents. **The value case has expanded accordingly**, from an efficiency-driven case to one that encompasses growth, customer retention, and technical profitability.

2. **Insurers are in the early stages of agentic maturity, as they are experimenting with orchestrating agents**. They are prioritizing core insurance processes – claims, underwriting, servicing – for transformation to establish competitive advantage. Insurers are a few stages away from the full horizon of agentic maturity, which will involve deploying a network of autonomous agents that self-organize to tackle business challenges.

3. **Change management is far more significant than new technologies in launching successful agentic AI transformations**—specifically, 70% of efforts should focus on People and change, 20% on technology and data, and 10% on algorithms. CEO commitment and business leadership's ownership over technological transformation are necessary to drive process re-imagination and deliver impact at scale. Currently, regional players tend to prefer balance governance that blends central centers of excellence with local autonomy. But the heterogeneity of legacy tech stacks makes scalability more aspiration than reality.

The peer group includes pan-Asian leaders, domestic champions, and multinational conglomerates operating across life, health, P&C, and commercial specialty lines, with distribution models ranging from tied agency forces and bancassurance partnerships to digital aggregators and direct-to-consumer channels.

# Finding 1: Where We Are Going – From GenAI Experiments to Agentic Transformation

The most consistent finding across all geographies is that insurers' strategy is evolving. Insurers across the Asia-Pacific region are moving from asking whether AI has a meaningful role in core insurance to asking how autonomous agents can be systematically embedded across operating processes. This shift from experimentation to transformation is the defining characteristic of the current moment.

## 1.1 Agentic AI as a Distinct Category

Every insurer in the peer group regards agentic AI as a qualitative step-change from GenAI. GenAI is principally a content tool — it summarizes, drafts, and classifies. Agentic AI involves digital workers who plan, reason, draw on external tools and data sources, and act autonomously. One or two years ago, the industry debate centered on whether GenAI had a meaningful role in core insurance processes. Today, the conversation has moved to how fast insurers can embed autonomous agents across operating processes.

## 1.2 The Move Towards End-to-End Reimagination

Until recently, most AI agendas were structured around experimentation: pilots, proofs of concept, and isolated use cases. The guiding question was, "Where can AI help?" The prevailing question now is: "What would this process look like if it were designed from scratch with autonomous agents?"

This is a sizable change. Rather than distributing effort across many small initiatives that deliver incremental value, leading companies are concentrating resources on one or two end-to-end process transformations with higher ambition levels. As one regional player noted, "We want to move away from the use-case factory approach."

The execution challenges of this shift should not be minimized. Moving from pilots to production requires more than better technology - it demands genuine transfer of process ownership across well-established organizational boundaries. In most insurers, end-to-end reimagination cuts across reporting lines and disrupts established incentive structures. The organizational rewriting needed – redefining who owns a process, who is accountable for outcomes, and how performance is measured – is a harder constraint than the technology itself for most organizations.

## 1.3 An Expanded Value Case

The public narrative around AI in insurance has been narrowly focused on maximizing efficiencies and workforce reduction. The picture that emerged from our conversations is much broader and more strategic than mere headcount metrics.

While efficiency and productivity remain relevant value pillars, some executives want to deploy AI for more far-ranging goals, including:

- customer experience (improved service quality, speed, accuracy)
- growth (lower churn, improved conversion)
- technical excellence (better loss ratios through more robust fraud, waste, and abuse detection; more consistent underwriting; and reduced claims leakage).

Not only are executives expanding their value case, but they are also envisioning a more compressed, aggressive timeline in terms of speed-to-impact. The majority aim to establish P&L impact within 12 to 24 months of launching a transformation, indicating either heightened confidence in technology or a sense of competitive urgency.

> We want to move away from the use-case factory approach.

## Finding 2: Where We Are Today – Maturity, Priorities, and Emerging Frontiers

To decide how to best move forward, it is important to assess where one currently stands. The peer group conversations revealed that most insurers are still in the early stages of Agentic AI maturity, and they are concentrating their efforts in core insurance processes: claims, underwriting, and servicing.

### 2.1 A Maturity Snapshot

BCG's Horizon framework, as illustrated in Exhibit 1 below, maps the evolution of agentic AI capability across five stages, from the simplest rule-based agents through to fully self-organizing agentic networks. At Horizon 0, constrained agents execute rule-based tasks within strict boundaries. Horizon 1 introduces single agents capable of independently handling multi-step tasks. Horizon 2 involves orchestrated multi-agent workflows where specialist agents collaborate under predefined coordination. Horizon 3 represents fully role-based multi-agent systems operating across end-to-end processes, with human oversight concentrated on exceptions. Horizon 4 — still largely theoretical at enterprise scale, with feasibility and performance improvement still to be confirmed — envisions a self-organizing agent mesh that dynamically coordinates and adapts across functions. However, getting there would require modular architecture, mature data ontologies, embedded governance, and deliberate operating model redesign - not simply greater agent autonomy.

**EXHIBIT 1**

Agentic AI Maturity Framework (Horizon 0–4)

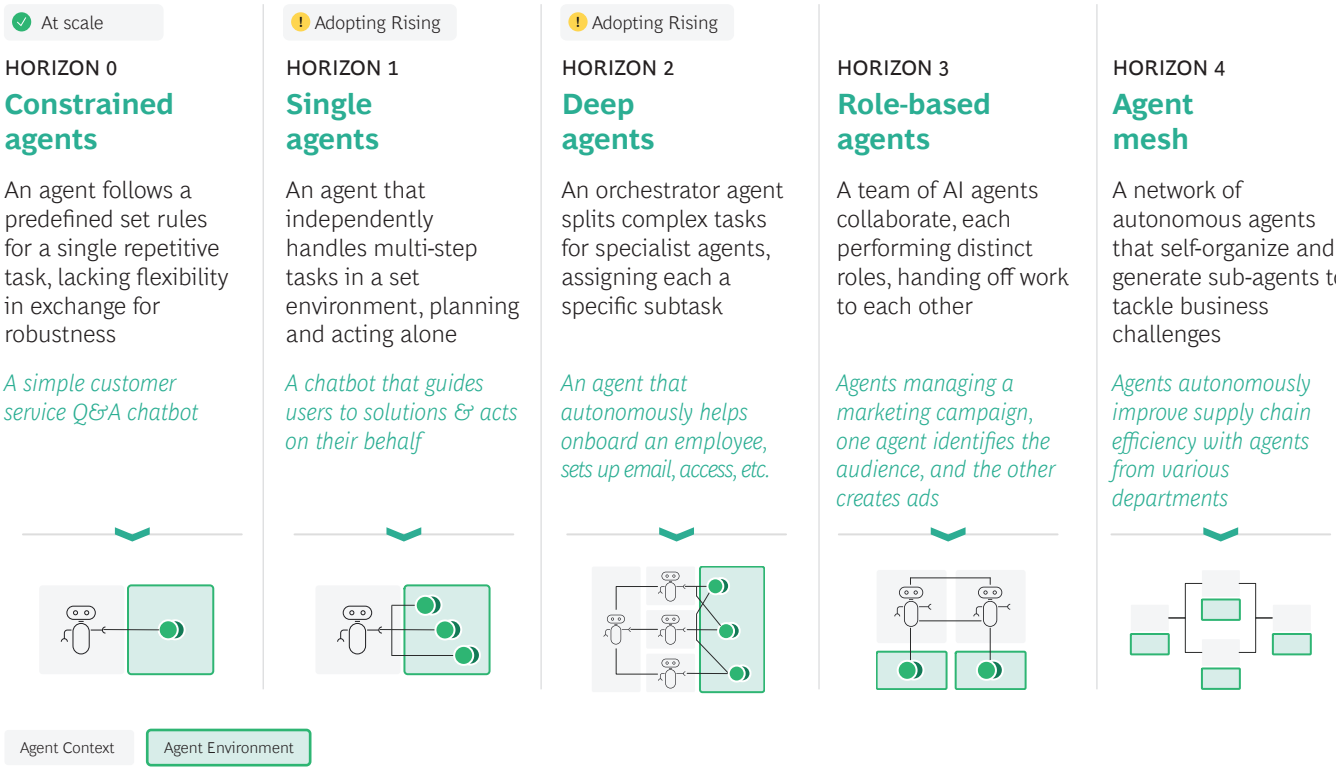

As Exhibit 2 highlights, most insurers are experimenting with Constrained and Single Agents (Horizon 0 & 1), with only a few experimenting with Horizon 2 through structured multi-agent pipelines that support digitization, triage, admissibility checks and decisioning. Horizon 3 — role-based multi-agent collaboration across end-to-end processes — remains nascent; just one organization in the peer group has begun experimenting at this level. Horizon 4 is still largely confined to research and development; no insurer currently considers a self-organizing agent mesh feasible at enterprise scale in the near term.

There is a considerable gap between where most insurers are today and the horizons at which agentic AI delivers the most competitive advantage. In the meantime, insurers are concentrating on their agentic AI investments in a few key domains.

**EXHIBIT 2**

## Agentic Maturity – Where interviewees stated they are today

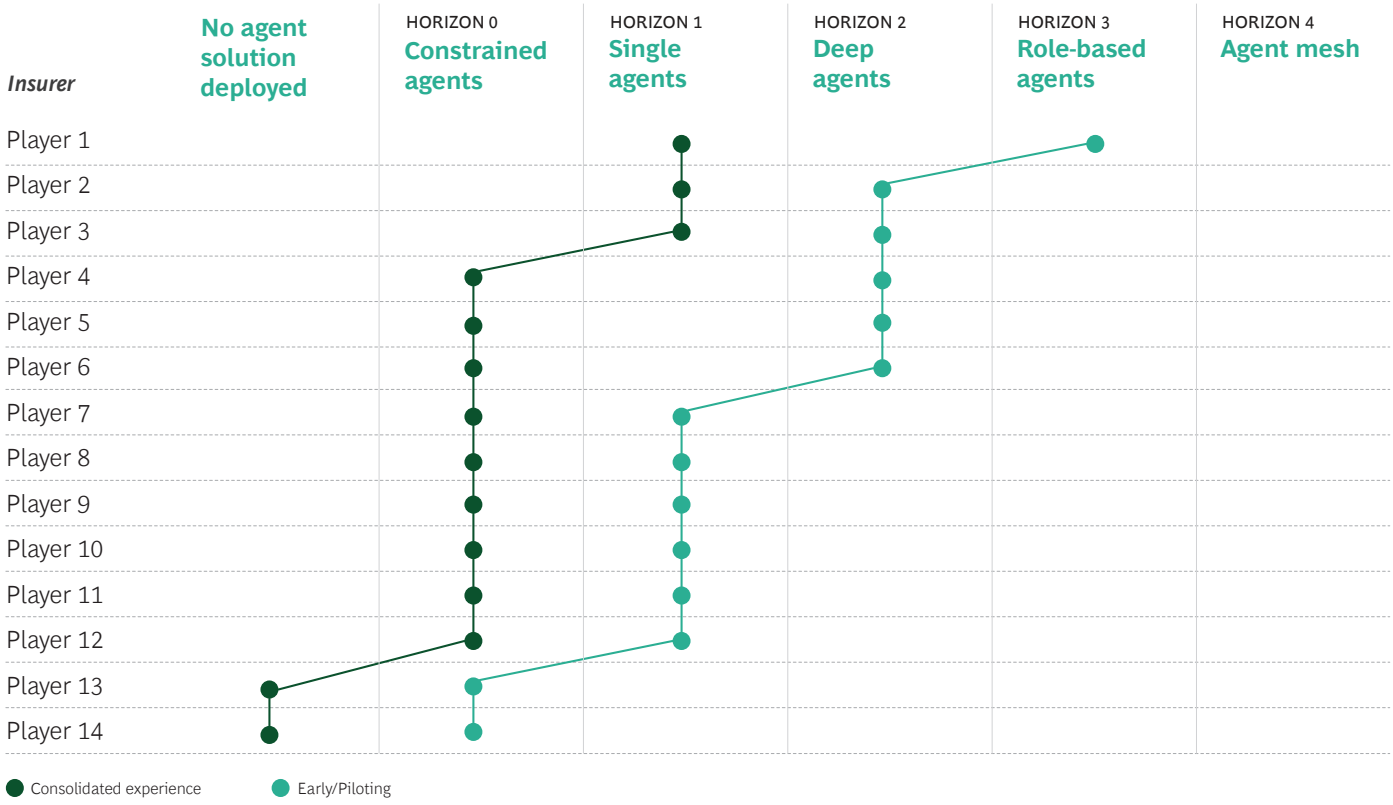

## 2.2 Priorities: Underwriting, Claims, and Servicing

Across the peer group, three domains emerged as the priorities of agentic AI investment: Claims, Underwriting, and Servicing. Distribution and Customer Service followed after these three domains, particularly in markets where digital channels carry greater weight.

### P&C Underwriting for SME & Commercial

Commercial, as well as Small & Medium Enterprise (SME), underwriting is constrained by fragmented intake processes, siloed risk data, and inconsistent decision-making across the end-to-end journey. Manual data extraction and eligibility checks slow the early stages; risk information is dispersed across internal and external systems; guideline application and benchmarking vary in practice. The consequence is extended cycle times and decision variance, with underwriters typically spending  up to 70% of their time on data gathering and enrichment rather than on risk judgement. Existing generative AI deployments have improved summarization but have not fundamentally altered underwriting orchestration or the overall decision flow.

Agentic orchestration introduces coordinated submission triage, automated enrichment across a broad range of internal and external sources, dynamic appetite matching, and structured quote construction with escalation limited to genuinely complex cases. Based on BCG's experience across proven transformations, the impact is material — quote turnaround times are compressed from 10 to 20 days to under 24 hours, productivity jumps two to five times, and loss ratio is improved by up to 7 percentage points.

### P&C Retail Claims

Retail P&C claims are the largest single operational cost pool in most non-life insurers where over 80% of customer data originates. Claim-handling is a key moment in the customer relationship that directly shapes trust and retention, yet the process is riddled with fragmented workflows, document-intensive handling, and inconsistent settlement outcomes. Traditional AI has led to some improvements, but it has not led to more streamlined coherence across the claim lifecycle.

Agentic orchestration enables a coordinated approach across the full claims lifecycle, from first notification of loss (FNOL) triage to coverage validation, repair estimation, fraud screening, and reserve-setting. The workflow is not rigid but rather adapts dynamically to the characteristics of each claim. Human expertise is concentrated on genuine complexity rather than routine process execution. Documented outcomes from proven implementations include productivity uplifts of 30-40%, straight-through processing (STP) rates exceeding 50% in selected segments, leakage reduction of 4 to 5 percentage points, and average cycle times reduced from over 40 days to under a week for standard claims.

### Health Claims

Health claims processing is high-volume and rule intensive. In most markets, 90 to 95 per cent of claims are already handled through straight-through processing (STP) with relatively straightforward administrative controls, leaving a small share for manual handling. Within that manual share, however, is a disproportionate concentration of complexity, cost, and leakage risk. The challenge that was consistently validated across the peer group lies in the gap between rule-based automation and real-time coherence assessment — the capacity to evaluate eligibility, treatment relevance, and provider behavior in an integrated way, in the context of the individual claim and patient journey.

Agentic workflows address this by enabling real-time, multi-dimensional assessment of eligibility, treatment appropriateness, provider patterns, and cross-claim consistencies prior to payout. Drawing on our experience in healthcare, we observed outcomes such as payout reductions of 3-5% attributable to stronger fraud, waste and abuse (FWA) controls, and customer retention improvements of 10-15%.

**Life Claims and Servicing**

Life insurance operations present a different operational profile: lower claims volumes with high documentation intensity, combined with high-volume servicing and contact center activity. The challenges include manual eligibility verification, beneficiary validation, policy cross-referencing, and repeated document follow-up, all of which contribute to extended handling times and variable customer experience. A further complication is that even routine servicing actions frequently trigger downstream tax, entitlement, and compliance requirements that depend on specialist interpretation.

Insurers in our peer group are aiming for straight-through processing (STP) for 70-80% of life claims and reducing fraud leakage from 5% to 2%. This would require structured agentic orchestration across eligibility assessment, beneficiary verification, anomaly detection, and end-to-end case management. For most life carriers, this remains a medium-term goal rather than an immediate deployment, but they are taking steps to enable this direction by coordinating decision-flows rather than deploying isolated automation layers (more on this later).

## 2.3 LLM-Based Distribution: An Emerging Frontier

In markets where digital and aggregator channels play a significant role in shaping competition — Australia is a prominent example — LLMs are emerging as a key strategy to reach consumers. Insurers in these markets are exploring how LLMs can strengthen direct-to-consumer models and unlock cross-sell and retention capabilities.

Smaller players in digitally advanced markets have raised a structural concern: As LLM capabilities become more central to distribution, organizations with larger and richer proprietary data ontologies, assets, and integrated agentic architectures may accrue advantages that are difficult for others to match. Smaller and specialist insurers likely cannot “build from scratch” in-house data capacities, and thus benefit from pursuing a “buy and integrate” strategy centered on partnerships, ecosystem participation, and modular or third-party agent solutions.

# Finding 3: How to Win – Architecture, Data, and the Primacy of Change Management

Despite rising ambitions to unlock the full value of agentic AI, insurers are finding deployment and change to be more challenging than expected. The primary bottleneck cited in peer group conversations is architectural readiness, with mainframe-based policy administration systems acting as structural barriers to embedding agents. Overhauling legacy systems requires a concerted focus on change management.

## 3.1 Change Management: 10-20-70 Rule

The human dimension, not the technical one, is the binding constraint in AI transformation—this was the recurring theme in each and every conversation with executive leaders. For instance, one market leader with two years of GenAI use-case testing observed that the constraint they have faced has not been technological — it has been organizational execution. Several leaders raised the question of how to build genuine buy-in and support among their people — a question that points directly to the change management work required. These conversations affirmed BCG's position that 70% of transformation effort should be directed at change management, 20% on technology, and 10% on algorithms.

Multiple organizations are working through fundamental operating model questions: whether the transformation program should be led from the business or the technology function; how to structure funding so that it reflects strategic priority rather than

incremental budget allocation; and how to mobilize the broader workforce in a landscape with genuine uncertainty about role evolution. For example, an insurer with a well-resourced AI center of excellence is actively transitioning away from a "proof-of-concept factory" model, recognizing that isolated experiments at the margins of the business do not translate into enterprise-scale change.

The practical implication for insurance executives is that transformations that are owned primarily by technology functions are unlikely to capture the value available from agentic AI. The transformations that are delivering results share a common structure — they are led from the business, sponsored at CEO level, and treat change management as the primary workstream rather than a supporting activity.

## 3.2 Preparing Architectural and Data Readiness

A robust change management initiative is necessary to modernize legacy technology systems. Agents may perform well in controlled pilots, but without modular architecture and clear abstraction layers separating agent logic from core dependencies, production deployment is fragile. It is not enough to simply add an agentic layer on top of an existing technology stack. Architectural transformation is therefore not a technical detail — it is a prerequisite for enterprise-wide adoption. That said, some organizations are questioning whether modernizing traditional core systems is the right path, and wondering whether AI will leapfrog it entirely.

Data ontology is emerging as a decisive enabler and potential competitive moat. Leading insurers in the peer group are moving beyond data lakes towards structured ontologies and knowledge graphs that explicitly map entities, relationships, and business rules — enabling agents to reason consistently and accurately across diverse process contexts. Keep in mind that best practices for ontology design at this level of complexity are not yet mature. Thus, organizations that first develop proprietary approaches are, in effect, building a differentiated capability that will be hard to replicate.

The path to readiness is not uniform across the peer group. Multi-geography carriers operating across varied regulatory environments, legacy technology stacks, and distribution models face a compounding complexity that can significantly slow progress. Organizations with more homogeneous infrastructure or greater geographic concentration may find themselves better placed to move quickly and thus capture an early advantage.

## 3.3 Regional players prefer balanced governance over centralized command

Multi-market groups in the peer group consistently express a preference for balanced governance over fully centralized AI transformation models. The model that is emerging is one in which a central function — typically a center of excellence — develops and maintains reusable assets: data ingestion pipelines, augmentation frameworks, evaluation tooling, and shared standards. Business units retain meaningful autonomy over prioritization and execution, and thus can adapt to the real differences in regulatory environment, product architecture, and customer and distribution dynamics across markets. Executive leaders emphasized the importance of modularity — the ability to build a capability once and deploy it across operating companies — as the mechanism for capturing group-level synergies without imposing a uniformity that the business cannot absorb.

This vision of modularity that leads to scalability within a balanced governance system is rarely realized in most cases. The heterogeneity of legacy technology stacks, regulatory requirements, and distribution models makes genuine portability considerably harder than anticipated. Bridging the gap between the vision of scalability and the reality of legacy technology is likely to be one of the defining organizational challenges in the coming years.

# Finding 4: A Readiness Diagnostic for Insurance Executives

The following questions are intended as a diagnostic for any insurer assessing their Agentic AI readiness. The questions are drawn directly from the themes that emerged across the 14 conversations.

**On Strategy and Ambition: the "What"**

1. Have you selected one or two moonshot processes for end-to-end reimagination, rather than running a portfolio of small use cases?

2. Is your business case built on three pillars (efficiency, growth, technical excellence), or is it still anchored to headcount reduction?

3. Have you assessed where you sit on the Horizon 0–4 maturity scale, and do you have a realistic plan of progress?

4. Are you exploring LLM-based distribution as a new channel capability, or are you treating it as a niche experiment?

**On Enablers and Organization: the "How"**

1. Does your CEO personally sponsor the Agentic AI transformation, with visible priority and dedicated time?

2. Is the program led by business owners with technology enablement, not the reverse?

3. Have you allocated strategic transformation funding, or are you relying on business-as-usual budgets?

4. Is your operating model for AI governance balanced between central standards and reusable assets, and local autonomy for prioritization and execution?

5. Are you investing 70% of your effort in change management - new roles, new skills, new ways of working - rather than overweighting technology?

6. Do you have a clear view of your data ontology, that is, the entities, relationships, and business rules your agents need to make decisions?

7. Have you assessed your architectural readiness honestly, including the specific legacy constraints most likely to impede production deployment?

8. Is your Responsible AI framework sufficiently developed to govern agent autonomy, human oversight thresholds, and regulatory compliance at scale?

The organizations in this peer group that are making the most tangible progress share certain characteristics: they have moved from experimentation to committed transformation programs; they have identified where the highest-value opportunities lie and have concentrated resources accordingly; and they are treating the organizational dimensions of change with the same seriousness they bring to the technical ones. For insurers that have not yet made this shift, the time available to do so without material competitive consequence is shortening given how fast the insurance market is moving.

# About the Authors

**Davide Corradi** is a Managing Director and Senior Partner in BCG's Milan office. You may contact him at **Corradi.Davide@bcg.com**.

**Semih Durmus** is a Managing Director and Partner in BCG's Minneapolis office. You may contact him at **Durmus.Semih@bcg.com**.

**Giuseppe Franco** is a Managing Director and Partner in BCG's Milan office. You may contact him at **Franco.Giuseppe@bcg.com**.

**Joe Khoury** is a Managing Director and Partner in BCG's Miami office. You may contact him at **Khoury.Joe@bcg.com**.

**Swayamjit Mishra** is a Managing Director and Partner in BCG's Mumbai office. You may contact him at **Mishra.Swayamjit@bcg.com**.

**Eunkyung Koh** is a Managing Director and Partner in BCG's Seoul office. You may contact her at **Koh.Eunkyung@bcg.com**.

## For Further Contact

If you would like to discuss this report, please contact the authors.

Boston Consulting Group partners with leaders in business and society to tackle their most important challenges and capture their greatest opportunities. BCG was the pioneer in business strategy when it was founded in 1963. Today, we work closely with clients to embrace a transformational approach aimed at benefiting all stakeholders—empowering organizations to grow, build sustainable competitive advantage, and drive positive societal impact.

Our diverse, global teams bring deep industry and functional expertise and a range of perspectives that question the status quo and spark change. BCG delivers solutions through leading-edge management consulting, technology and design, and corporate and digital ventures. We work in a uniquely collaborative model across the firm and throughout all levels of the client organization, fueled by the goal of helping our clients thrive and enabling them to make the world a better place.


EXPORTING DATA...
BCG