



# **From spreadsheet heroics to autonomous treasury**

**An agentic AI adoption  
playbook for CFOs and  
treasurers**

**August 2026**



**The better the question.  
The better the answer.  
The better the world works.**



**Shape the future  
with confidence**


# 1

## Introduction

If the previous decade in treasury was about deploying the right technology, this decade is about making that technology think. The shift from automation (doing repetitive things faster) to Agentic AI (reasoning, adapting and acting within defined boundaries) represents a qualitative change. Unlike many technology transitions, this one has a clear value proposition: treasury teams that adopt it correctly could spend less time on coordination and more on the decisions that move the needle.

This playbook, informed by the work EY has done with treasury functions across manufacturing, financial services and infrastructure, sets out what treasury digital transformation looks like when it is undertaken deliberately rather than reactively. We examine the workflow-level case for change, the treasury datalake architecture required to support it, the treasury automation strategy that builds trust before scale and the organizational capabilities that determine whether a pilot becomes a transformation or just another proof of concept.

The playbook helps practical application. Agentic AI will not succeed as a bolt-on tool. It requires treating treasury as an orchestrated system where AI agents, treasury workflow automation and human expertise work together under proper governance. The organizations we have seen move fastest are those that grasped this early and built accordingly.



Source: EY analysis

## Effective Agentic AI application

An aspect often overlooked in the transition to Agentic AI is that Agentic AI does not operate in isolation. It builds on a prior layer of change at the workflow level. The most effective treasury organizations usually begin by introducing targeted digital breaks, removing manual handoffs and enabling system-driven execution across core processes. These digital breaks, when scaled, create smart workflows where routine coordination is automated and human effort is focused on exceptions. Agentic AI builds on this foundation, applying reasoning, simulation and decision support across workflows rather than replacing them outright.

The central ambition is to reduce the time consumed by manual coordination, fragmented data and sequential approvals, and redirect that effort toward decisions that create value. The purpose is not to reduce headcount but to make treasury teams more capable in their roles by equipping them with integrated workflows, real-time insights, and decision support that compress cycle times from days to hours and from hours to minutes. The future of treasury is not a leaner team doing the same work more efficiently. It is a different operating model, where better decisions are made faster, supported by a Center of Excellence (COE) that continuously raises the intelligence baseline of the function.


## 2

# The spreadsheet problem is structural, not cosmetic

Treasury professionals often work with deeply suboptimal systems. The spreadsheet has served as the primary decision-support tool for treasury for around 30 years, and for much of that time, it was among the best available options. That is no longer the case, yet the spreadsheet persists. Understanding why treasury teams are still dependent on spreadsheets is the starting point for any credible treasury workflow automation program.

In many cases, a significant amount of the treasury function's capacity is consumed by work that may not require human judgment. Across organizations, we have observed that at mid-levels 60% to 70% of treasury bandwidth is spent on low-value activities, primarily aggregation of data across systems, manual validation against counterparty or policy thresholds and the production of reports that are frequently outdated by the time they reach decision-makers.

Time spent on core treasury activities: Traditional vs. digital break approach



Source: EY India analysis

The scale of the problem is often understated. A typical mature treasury function runs between 50 and 100 interconnected spreadsheets covering cash positioning, FX exposure, investment tracking, debt management and regulatory reporting. In practice, the inefficiency is not driven by the complexity of the work itself, but by fragmentation. Treasury teams routinely pull data from Enterprise Resource Planning (ERP) systems, Treasury Management System (TMS) platforms, bank portals and individual spreadsheets into a single view – often, manually. In one observed case, a global treasury team managing over 120 bank accounts across more than 20 currencies spent close to two weeks each month on reconciliation alone due to the lack of a unified data layer.

In many cases, each spreadsheet is effectively a silo with its own version control, formula logic and risk of error. When one 'silo' breaks, the downstream effects are not always visible until month-end. If one spreadsheet is updated erroneously, the audit trail can become non-existent.

Left unaddressed, this fragmentation does not just slow down operations. It directly undermines the effectiveness of newer technologies introduced into the environment.

In our engagements, the most common root cause observed for the gap between AI experimentation and scalable treasury adoption is that organizations attempt to build AI capabilities on top of fragmented treasury data governance structures. Often, the data that the AI tool or function is working with is not clean enough or timely enough to produce reliable outputs. Therefore, treasury teams try the tool, observe inconsistent recommendations and retreat to their spreadsheets. The real reason is a weak foundation and not that agentic AI cannot work in treasury.

The implication of data readiness is clear: when data remains fragmented and workflows remain disconnected, intelligent treasury operations cannot be sustained. The path forward runs through the datalake, not around it.

1. The deeper issue is one of integration architecture and not data quality alone. With spreadsheets, it is easy to introduce errors and fragment processes. Each spreadsheet becomes an isolated logic layer, requiring manual intervention to connect sub-processes.
2. More importantly, when treasury professionals spend most of their time as information couriers rather than decision-makers, the function's strategic value is structurally suppressed. Replacing spreadsheets with a comprehensive, integrated datalake is more than a technology upgrade. It is a prerequisite for transforming how treasury sub-processes connect, how decisions are made, and how quickly the function responds to changing conditions.

## From digital breaks to agentic workflows

The transition to Agentic AI, however, begins not with intelligence but with fixing how workflows operate. Digital breaks replace manual steps such as data handoffs or reconciliations with system-driven actions. Individually minor, these changes scale into smart workflows where systems handle routine coordination and humans focus on exceptions.

Agentic AI builds on this foundation, adding reasoning, simulation and decision support. Digital breaks change work execution; agentic AI determines what should happen next. The real test is cycle time. When processes compress from days to hours and trigger real-time exceptions, the transformation goes from incremental to structural.

# 3

## What changes workflow by workflow

Digital break implementations already show tangible outcomes, including 80%–90% auto-match rates in reconciliation, reduction of reporting cycles from days to hours, and up to a threefold improvement in counterparty validation times, as per EY India analysis.

The case for Agentic AI in treasury is most compelling when examined at the workflow level – because that is where treasury professionals experience the most difficulties and where treasury back-office automation and treasury mid-office transformation can deliver the most tangible gains.

In this section, we have mapped the six core treasury workflows where the contrast between spreadsheet-era practice and agentic AI capability is sharpest in most cases.

### Foreign exchange (FX) and interest rate risk management

The back office is where the volume lives and therefore where digital breaks can deliver the fastest measurable impact. The table below illustrates the shift in approach across the core activities:

| Spreadsheet era                                                                                                                                       | With Agentic AI                                                                                                                                                                                    |
|-------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Manual pull of FX exposure from ERP, combined with Bloomberg rates in VBA macros. Each scenario takes hours and is stale by the time it is presented. | Agent ingests real-time exposures from ERP plus live market data and geopolitical signals. Runs 50+ hedge scenarios in minutes; recommends the optimal hedge ratio based on defined risk appetite. |
| Post-trade: Manual hedge register update, static Mark to Market (MTM) calculation, manual P&L impact report for accounting.                           | Post-trade: The agent auto-updates the hedge register, calculates MTM on live curves, generates IFRS-compliant effectiveness documentation with a full audit trail.                                |

### Cash flow planning and liquidity gap identification

| Spreadsheet era                                                                                                                      | With Agentic AI                                                                                                                                                                |
|--------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 20–30 bank statements consolidated manually. Funding gaps identified via color-coded cells. Forecast variance routinely exceeds 20%. | Agent aggregates multi-bank, multi-currency balances in real time. Learns timing patterns in receivables and payables. Flags funding gaps 30–60 days ahead with ~90% accuracy. |
| Manual reconciliation of actual vs. forecast; next month's forecast adjusted based on manual variance analysis.                      | Agent auto-reconciles actual vs. forecast, explains variance drivers (e.g. delayed receivables, FX movement) and updates the rolling forecast without intervention.            |

## Investment management and asset allocation

| Spreadsheet era                                                                                                                            | With Agentic AI                                                                                                                                                                                    |
|--------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Spreadsheet matrix of surplus cash, tenor, yield and counterparty limits. Allocation decisions based on static policy; approvals by email. | Agent monitors surplus cash continuously, scans money market and bond yields, recommends optimal allocation across fixed deposits (FDs), treasury bills and commercial paper within policy limits. |
| Separate sheets track maturities, accruals and counterparty limit compliance.                                                              | Agent auto-tracks maturities, calculates accruals, generates compliance reports for CFO and auditors on demand.                                                                                    |

### Case study: Improving investment tracking at an Asia - pacific Consumer Goods Group

**Challenge:** The consumer goods group's treasury team was spending three to four days per month manually updating investment tracking spreadsheets across six subsidiaries. Counterparty limit breaches were identified only at month-end review.

**Approach:** EY implemented an agentic investment monitoring layer connected to the group's datalake. The agents continuously tracked maturities, flagged pre-breach limit conditions and auto-generated the monthly investment compliance pack.

**Outcome:** Investment tracking effort reduced by ~80%; zero counterparty limit breaches in the first nine months of operation; compliance report preparation time cut from three days to four hours.

## Documentation, compliance and board reporting

Among treasury workflows, reporting and analytics consistently represent the largest concentration of manual effort. In observed programs, variance reporting alone required over 60 hours per month before workflow automation. Treasury teams across our client base consistently report spending three to five days per month preparing board packs, audit schedules and regulatory reports, predominantly by copying data from multiple spreadsheets into presentation templates. Treasury reporting automation inverts this: reports are generated on demand from source systems, with full traceability of approvals and decisions. The treasury team's role shifts from assembly to review.

The impact of this shift is not limited to efficiency; it fundamentally changes how information is consumed across the organization. With a treasury reporting automation layer, the CFO can receive reports that are current at the point of distribution and not dated. The audit team can receive documentation that is auto-generated from source data with a complete lineage trail. And the compliance function can receive covenant monitoring outputs that are continuous rather than periodic.

Beyond speed, the more meaningful change lies in the nature of the information itself. In a spreadsheet-led reporting environment, information is static: a point-in-time capture that ages on arrival. In an agentic reporting environment, information becomes fungible. It is continuously refreshed, contextualised against historical patterns and forward-looking signals, and available on demand to any authorized stakeholder. Board packs do not describe the past. They become live decision support instruments that surface where attention is required, why a variance occurred, and what the range of forward scenarios looks like. That is the difference between reporting that informs and reporting that enables action.

# 4

## Datalake as a prerequisite

One of the clearest patterns across failed treasury AI implementations is the sequencing error: organizations attempt to deploy agentic AI before establishing the treasury datalake architecture that makes it reliable. The AI capability is real, but it is only as effective as the data it can access and trust. A dispersed, ungoverned data architecture produces inconsistent AI recommendations, and repeated inconsistency erodes trust faster than any other factor.

The architecture that underpins sustainable treasury workflow intelligence has four distinct layers, each of which must be designed deliberately:

| Centralized ingestion                                                                                                                                                            | Single source of truth                                                                                                                                                                 | Structured + unstructured data                                                                                                                                                         | Governed access layer                                                                                                                                                              |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| API pipelines stream ERP cash flows, bank balances, trade confirmations, market rates and counterparty limits directly into the datalake – replacing manual downloads into Excel | Datalake becomes the canonical repository. Spreadsheets become visualization tools on top of clean data, not the source itself – eliminating the '5 versions of cash position' problem | Unlike spreadsheets, the datalake ingests contracts, emails and market news apart from quantitative data – enabling agents to reason across both dimensions for better risk assessment | Role-based access and lineage tracking built in help treasury trace exactly where any number came from and who modified it – solving the black-box problem of complex spreadsheets |

The datalake is more than a data management layer. It forms the foundation on which future treasury operating models and treasury as an insight-led operation rest. A well-structured datalake does more than store historical transactions. It acts as the memory layer of the function. It allows systems to detect patterns from past cycles, flag anomalies early and produce forward-looking scenarios based on live data. Organizations that invest in this architecture are not addressing a current data gap. They are building a treasury function that supports decision-making in real-time rather than waiting for cycle-based reporting.

### From data architecture to intelligent execution

While the four-layer datalake architecture establishes a reliable data foundation, it does not, on its own, change how treasury work is executed. In practice, leading treasury organizations extend this foundation through three interconnected layers: data, workflows and intelligence.

The first layer is the governed datalake, which consolidates structured data from ERP and Treasury Management System (TMS) systems along with unstructured inputs such as contracts, emails and market data. The second layer comprises smart workflows and API-driven processes, where manual handoffs are replaced by event-driven execution and exception routing. The third layer is the agentic AI layer, which applies reasoning across these workflows – validating transactions, generating narratives, simulating scenarios and recommending actions aligned with defined policies.

This layered model is critical for sequencing. Without workflow automation, AI lacks execution context. Without a governed datalake, it lacks reliable inputs. Sustainable treasury intelligence emerges only when all three layers operate together as an integrated system.

Before any agentic AI deployment, treasury teams should conduct a thorough assessment of their current treasury data governance structure: where data lives today, what its quality and latency profile looks like, what governance gaps exist and what the target state needs to be to support real-time ingestion and AI reasoning. This assessment is the foundation on which the entire treasury digital transformation rests.

The role of treasury API integration is central here. Most banking partners and TMS vendors now offer some level of API connectivity, but coverage can be inconsistent. Our experience points to a hybrid ingestion approach: direct APIs where available, Secure File Transfer Protocol (SFTP) or Change Data Capture (CDC) for legacy systems, and Robotic Process Automation (RPA) only as a genuine last resort. Everything lands in the datalake with consistent governance regardless of ingestion method – and the treasury API integration layer grows over time as banking partners expand their connectivity.





## The adoption roadmap: Build trust before scale

Treasury organizations operate at different levels of maturity: from manual and fragmented environments to partially automated, to connected data ecosystems, and finally to intelligent and proactive treasury functions enabled by agentic workflows.

The failure modes for Agentic AI in treasury are well-documented: poor data, weak governance and a lack of explainability. Organizations that avoid these failure modes share a common discipline – they sequence deliberately, build treasury data governance before scale and treat early pilots as trust-building exercises rather than production deployments. This is the essence of a credible treasury automation roadmap for CFOs.

The adoption approach we advocate has four design principles at its core:

1. Smart overlay, not rip-and-replace – Wrap the AI agent around existing TMS and ERP via APIs rather than overhauling core infrastructure
2. Advisory before autonomous – Agent recommends, human approves; no autonomous payments in the pilot phase
3. Govern before scale – Establish the governance council, approval workflows and audit logging before expanding scope
4. Measure what matters – Track liquidity buffer reduction, cycle time and manual effort hours, not just model accuracy

A more fundamental shift sits beneath these principles. It changes how governance is structured. The conventional RACI model (Responsible, Accountable, Consulted and Informed), built on sequential hand-offs, periodic reviews and linear approvals, was designed for a manual treasury environment. It cannot keep up with the pace of automated workflows. As treasury operations evolve, governance must evolve with them. Sequential review is replaced by parallel, event-driven approvals. Periodic checks are replaced by continuous controls embedded at the point of execution. Fixed ownership gives way to role-based oversight focused on exception queues and system outputs. In this model, treasury professionals no longer manage transactions. They manage the rules, controls and calibration that guide how those transactions are processed.


## The phased roadmap

| Phase 1<br>Months 1-2                                                                                                                                                                                                                                                                                                                                                                      | Phase 2<br>Months 3-4                                                                                                                                                                                                                                                                                                                                    | Phase 3<br>Months 5-6                                                                                                                                                                                                                                                                      | Phase 4<br>Months 7-18                                                                                                                                                                                                                                                                              |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"> <li>Form pilot squad: treasury, data, risk, IT, audit</li> <li>Data architecture assessment across ERP, TMS, bank feeds</li> <li>Select a single high-impact pilot (cash forecasting or reconciliation)</li> <li>Define guardrails, approval workflows and rollback procedures</li> <li>Deploy agent in shadow mode alongside manual process</li> </ul> | <ul style="list-style-type: none"> <li>Move to advisory mode: agent recommends, human approves</li> <li>Real-time monitoring dashboards live</li> <li>Treasury team training on interpreting agent outputs</li> <li>Refine policy engine based on human overrides</li> <li>Begin integrating additional data sources (supply chain, FX feeds)</li> </ul> | <ul style="list-style-type: none"> <li>Review and quantify ROI: liquidity buffer, cycle time, effort</li> <li>Document governance and audit trail for regulators</li> <li>Address edge cases; widen exception logic</li> <li>Present business case for scaling to next workflow</li> </ul> | <ul style="list-style-type: none"> <li>Scale validated workflows across entities/regions</li> <li>Expand agent scope: risk, KYC/AML, reporting automation</li> <li>Establish Treasury CoE to own ongoing development</li> <li>Continuous back-testing and strategy evaluation integrated</li> </ul> |

## Prioritizing the right treasury workflows for pilot deployment

Reconciliation and reporting consistently emerge as the highest-impact entry points due to their combination of high manual effort, structured rule logic and measurable outcomes.

Not all treasury workflows are equal as pilot candidates. The best entry points combine high manual effort, suitably clean data and low execution risk. Three workflows consistently meet this criteria:



# 6

## The new operating model: From firefighter to strategist

At scale, the operating model stabilizes around a clear division of responsibility: approximately 80% of routine coordination and processing is handled by automated workflows and agents, while 20% – involving judgement, risk calibration and strategic decision-making – remains with treasury professionals.

The destination is not lights-out treasury. It is a deliberate rebalancing of where treasury expertise is applied. In an intelligent treasury operation in practice, agentic AI handles 80% of coordination, collation, and exception-clearing that currently consumes treasury bandwidth. Humans focus is on the 20% that genuinely requires judgment: counterparty relationships, strategic funding decisions and risk appetite calibration.

This rebalancing does not reduce the role of treasury professionals. It changes what the role demands. When automated workflows absorb coordination, what remains is work that requires judgement: reading the signal in a variance, calibrating risk appetite against market conditions, challenging model outputs when needed, and maintaining relationships with banking partners and internal stakeholders. The treasury professional no longer operates as a transaction processor. The role shifts toward setting system parameters, interpreting outputs and taking accountability for decisions that require human judgement.

Making treasury professionals more capable means providing faster access to information, clearer context, and decision support that surfaces the right question at the right moment.

This rebalancing has four measurable consequences:



↓ **Buffers**  
Better forecasting  
Excess liquidity buffers reduced as 30/60/90-day accuracy improves



↑ **Speed**  
Faster cycle times  
Onboarding, reconciliation and close cycles all compress



↑ **Audit**  
Stronger compliance  
Consistent, traceable decisions replace manual, variably documented ones



↑ **Strategy**  
More strategic bandwidth  
Treasury teams regain time for capital allocation and liquidity optimisation

## The Treasury Center of Excellence

Sustaining the new treasury model requires bilingual capabilities. Teams must understand both treasury processes and the underlying data and workflow architecture. Organizations that show staff transformation efforts purely with either treasury professionals or technologists tend to struggle; the advantage lies in bridging the two.

This capability is formalised through the Treasury CoE. The operating model that sustains this at scale is a Treasury Center of Excellence, typically hosted within a Global Capability Center, that owns the data lake pipelines, the agentic workflow library, and the ongoing treasury data governance framework for the global treasury function. The CoE is not an IT function; It is a bilingual treasury-technology team with a mandate to continuously improve the intelligence layer that the global treasury function runs on.

Organizations that establish a CoE early consistently outpace those that treat each AI initiative as a standalone project. The reason is institutional knowledge. The CoE builds understanding of what works, what does not, and how to handle the governance challenges that arise when AI touches regulated financial processes. The Treasury CoE operating model becomes the flywheel that compounds improvement over time. The impact of the CoE becomes most visible in how information flows through the treasury function. In a conventional treasury function, information flows upward: data is collected, aggregated and presented to decision-makers on a lag. In a CoE-led model, information flows in all directions simultaneously, feeding automated workflows, triggering exception alerts, informing forward-looking scenario models and surfacing insights to any stakeholder with the right access now they need it.

In this model, information is not just accessible. It is contextualised, enriched with historical patterns and forward signals, and presented in a form that supports the next decision. The CoE does not just operate this system. It continuously improves itself by incorporating feedback from human overrides, regulatory changes, and market developments into a treasury intelligence layer that becomes more precise over time.

# 7

## The path forward

For CFOs and treasurers, this period is likely to be defined not by who moved fastest, but by who moved most deliberately – who understood that the value of agentic AI in treasury is not in the technology itself, but in the system it enables: clean data, governed workflows and human expertise applied where it genuinely matters. This is the foundation of a credible treasury automation strategy.

The business case, in our experience, is rarely the obstacle. The obstacles have been discussed in this paper: governance ambiguity, treasury data governance gaps, change management in a risk-sensitive function and the temptation to scale before the foundation is solid.

Our guidance to treasury leaders is to resist that temptation. Start narrow. Build the treasury datalake architecture properly. Pick the pilot that combines the highest manual effort, the cleanest data and the lowest execution risk. Govern it properly from day one. And when the business case is proven, scale with the same discipline used to launch. The future of intelligent treasury operations is built one governed workflow at a time.

### Three decisions treasury leaders should make this quarter:

1. Commission a data architecture assessment – understand exactly where the treasury data lives, its quality and the gap to a governed datalake.
2. Identify the pilot workflow – cash forecasting or reconciliation in most cases; the most acute pain points and with the most accessible data.
3. Define the governance framework first – approval thresholds, audit logging and escalation paths before the first line of agent code is written.

### Contributors



**Hemal Shah**  
Partner, Risk Consulting  
Hemal.H.Shah@ey.com



**Prateek Chaturvedi**  
Partner, Risk Consulting  
Prateek1.Chaturvedi@in.ey.com



**Kriya Bhansali**  
Director, Risk Consulting  
Kriya.Bhansali@in.ey.com



**Madhu Challagulla**  
Director, Risk Consulting  
Madhu.Challagulla@in.ey.com

Ernst & Young LLP

EY | Building a better working world

EY is building a better working world by creating new value for clients, people, society and the planet, while building trust in capital markets.

Enabled by data, AI and advanced technology, EY teams help clients shape the future with confidence and develop answers for the most pressing issues of today and tomorrow.

EY teams work across a full spectrum of services in assurance, consulting, tax, strategy and transactions. Fueled by sector insights, a globally connected, multi-disciplinary network and diverse ecosystem partners, EY teams can provide services in more than 150 countries and territories.

All in to shape the future with confidence.

EY refers to the global organization, and may refer to one or more, of the member firms of Ernst & Young Global Limited, each of which is a separate legal entity. Ernst & Young Global Limited, a UK company limited by guarantee, does not provide services to clients. Information about how EY collects and uses personal data and a description of the rights individuals have under data protection legislation are available via [ey.com/privacy](https://ey.com/privacy). EYG member firms do not practice law where prohibited by local laws. For more information about our organization, please visit [ey.com](https://ey.com).

Ernst & Young LLP is one of the Indian client serving member firms of EYGM Limited. For more information about our organization, please visit [www.ey.com/en\\_in](https://www.ey.com/en_in).

Ernst & Young LLP is a Limited Liability Partnership, registered under the Limited Liability Partnership Act, 2008 in India, having its registered office at Ground Floor, Plot No. 67, Institutional Area, Sector - 44, Gurugram - 122 003, Haryana, India.

© 2026 Ernst & Young LLP. Published in India.  
All Rights Reserved.

EYIN2607-014

ED None.

This publication contains information in summary form and is therefore intended for general guidance only. It is not intended to be a substitute for detailed research or the exercise of professional judgment. Neither EYGM Limited nor any other member of the global Ernst & Young organization can accept any responsibility for loss occasioned to any person acting or refraining from action as a result of any material in this publication. On any specific matter, reference should be made to the appropriate advisor.

[ey.com/en\\_in](https://ey.com/en_in)

