ASTRAN

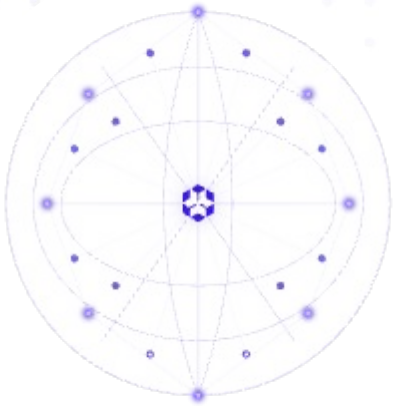

Workshop

AI for Treasury: Efficiency & Resilience

May 5, 2026

Executive Summary

From experimentation to deployment.
58% of participants placed their organisation at an Exploring stage, and none had reached an Advanced level. The gap between what is possible and what is practiced remains wide.

What's holding teams back: it's not the tools.
Data quality, trust in AI outputs and the risk of eroding human judgment were the top concerns raised by participants, especially when handled by younger generations. These are governance problems, not infrastructure ones. Treasury teams have to manage this transformation while running all existing topics (without an increase in the size of teams).

Cash forecasting leads the agenda.
For the next 12 months, 55.6% of participants ranked AI-assisted cash forecasting as their top priority. This is not a surprise: forecasting is where there is the most room for improvement, as the human version of it does not have the best accuracy. Any increase of accuracy would immediately bring a huge ROI (more effective cash management).

Continuous adaptation as the new baseline.
Organizations that treat AI deployment as a one-time project will fall behind those that treat it as an ongoing capability. Treasury teams that build this muscle now will be better positioned to lead when the next wave arrives. AI is evolving at a rapid pace, requiring a continuous adaptation mindset from everyone involved.

Secure AI use starts with knowing what you are sharing & delegating.
Not all data is equal, and not all prompts are safe. Distinguishing what can flow freely, what requires caution and what falls under hard limits is the first act of AI governance.

Overview

This workshop brought together 12 Treasury Leaders to explore how AI can drive both operational efficiency and resilience. They focused on key questions: where do teams stand on AI maturity, what holds them back, and what would they prioritise automating in the next 12 months.

A set of secure AI usage principles was also shared and discussed with the group.

AI maturity

The majority of participants (58.3%) consider themselves at an Exploring stage, with a third still at Beginner level, reflecting an early-to-mid maturity curve overall. No one has yet reached an Advanced level, suggesting significant room for growth.

While AI is highly valuable for data analysis and error detection, some hesitation remains, particularly around hallucinations and the need to critically evaluate outputs.

A key priority emerges from discussions: educating younger generations on how to engage with AI thoughtfully and critically.

Biggest concern about AI

Four themes emerged from the discussion:

Data quality & accuracy

The most concrete, operational concern: incorrect input data or inaccurate outputs undermine the value of any AI-driven analysis.

Trust & transparency

Participants debated the inability to fully trust AI outputs, and the difficulty of explaining variances generated by a model.

Value & adoption

Where is the real added value, and at what cost? AI saves time, but training yourself takes time too; especially as models evolve fast.

Erosion of judgment

Especially for younger generations: the risk of atrophied critical thinking. Participants insisted on keeping experience and human expertise alive.

Priority topics for Treasury Leaders

Cash forecasting is by far the most requested automation, cited by more than half of respondents, reflecting a shared pain point around the time and effort required to produce reliable, up-to-date cash projections.
Deviation explanation also emerged as a natural extension of this need: not just forecasting, but understanding and justifying variances automatically.

Control reporting came up twice, pointing to a recurring manual effort around reconciliation and compliance checks that the team sees as a strong candidate for automation.

Finally, two responses highlighted more operational tasks (commodity confirmations and slide creation), suggesting that beyond core treasury functions, the team is also looking to reduce time spent on daily tasks.

Best practices

What you can share, and what you shouldn't

Sharing data with AI is fine, **ungoverned sharing is the risk.** The real question is: do you know what you are sharing, with whom, and under what terms?

Share freely	Proceed with caution	Limits
• Anonymized or synthetic data • Public information, generic business scenarios • Own content (drafts, notes, ideas) • Connected tools with explicit authorization (eg: CRM, calendar) → where the data flow is controlled, consented and stays within the organization agreement	• Personal data (names, emails, roles) → acceptable if AI provider is GDPR-compliant and data isn't used for model training • Internal business data → check whether LLM plan guarantees data is not retained or used for training • Connected integrations → verify scope of access: read-only vs write/objects/users	• Credentials, API keys, passwords • Sensitive records & confidential strategy • Data belonging to third parties who haven't consented

Protecting against prompt injection

Malicious instructions hidden inside content collaborator ask the AI to process (email, document, webpage). The AI reads it and follows the hidden instruction without you realizing it.

Two types of prompt injection exist:

- Direct: unknowingly paste malicious content into the prompt
- Indirect: AI fetches content autonomously via a connected tool (email, web, CRM)

How to protect your collaborators?

- Never paste raw external content into a sensitive session, extract only what you need
- If the AI asks to click a link or "ignore previous instructions", treat it as an attack
- For AI agents: apply least-privilege, log actions involving write access

Five principles

1. **Give it a role**

"You are a treasury analyst" sets the frame. The model calibrates its expertise level and vocabulary accordingly.

2. **Include your data in the prompt**

A model without your numbers gives you generalities. With your numbers, it does actual work.

3. **Specify the output format**

Table, bullet points, 300 words: the model respects constraints when you give them.

4. **Say what you don't want**

Underrated and very effective. Saves two rounds of iteration.

5. **Break it into numbered steps**

One long question gets a vague answer. Three short questions in order gets three clean answers.

Looking ahead

The question is no longer whether AI belongs in Treasury, **it is about building, right now, the practices that will make it trustworthy and effective.**

Treasury teams and organizations are at an inflection point. The technology is ready, the use cases are clear, and the appetite is there. What comes next is less about tools and more about discipline: knowing what to delegate to AI, what to verify, and what must remain a human call. The teams that get this balance right won't just be more efficient, they will be more resilient.