

THE INDUSTRY EDITION / 25 SEPTEMBER 2026

# Give AI a job. Not the keys.

Faster service and better internal knowledge, without handing an assistant payment authority, customer access or credit decisions.

## FINANCIAL SERVICES & FINTECH

### THE OPPORTUNITY

Start with policy answers and service drafts.

### THE CYBER QUESTION

Can an untrusted document become an instruction?

### THE ACTION

A bounded pilot with a human decision at every critical step.

# In this issue.

Seven sections, written for the people who have to fund, build, secure and sign off a first financial services AI pilot.

|    |                                      |                                                                                                                                   |
|----|--------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| 01 | <b>The briefing</b>                  | The valuable assistant is the bounded assistant. Why the size of the consequence matters more than the size of the model.         |
| 02 | <b>AI in practice</b>                | Make policy knowledge easier to use. A proposed first deployment for a lender, payments business or insurance intermediary.       |
| 03 | <b>Cyber front</b>                   | An attachment is not an authorization. Keeping document understanding separate from permission to act.                            |
| 04 | <b>The 30-day playbook</b>           | Prove the service benefit inside the boundary. One policy category, a small employee group, no autonomous financial actions.      |
| 05 | <b>The decision desk</b>             | Value the completed task, not the generated answer. Illustrative capacity-value economics, with the sensitivity that breaks them. |
| 06 | <b>Monday's agenda</b>               | Put the risk owner in the room early. The questions that turn AI interest into a defined, testable proposal.                      |
| 07 | <b>Sources &amp; editorial notes</b> | Six numbered references, each with what it does and does not establish.                                                           |

## HOW TO READ THIS EDITION

Reported examples are attributed to their sources. The SME scenarios, pilot plans and decision questions are original editorial proposals, not verified Agents Stack customer results.

All cost examples use explicit assumptions, not market prices or promised returns. Figures are illustrative planning numbers in Singapore dollars.

This magazine offers design questions, not a declaration of compliance with MAS, RBI or any other regulator. Obtain advice for the actual entity and use case.

# The valuable assistant is the bounded assistant.

For fintech founders, finance leaders and risk teams: make the job description narrower than the technology's capability.

## The editor's view

An internal assistant that finds the right policy may be a better first investment than an agent that can complete a financial transaction. The distinction is not how intelligent the system appears. It is the size of the consequence when it gets something wrong, and whether an existing control can still stop that outcome.

Most financial SMEs already have controls that work: maker-checker, segregation of duties, independent verification. A first AI deployment should sit inside those controls rather than quietly route around them.

## Our recommended first lane

Start with an internal policy-and-service assistant. It should retrieve approved information, identify uncertainty and draft an answer for an employee. In the initial pilot, it should not approve credit, change a beneficiary, submit a regulatory filing or initiate a payment. Those are separate decisions with separate owners.

## The practical question

Can staff handle a defined class of inquiries faster, after verification and correction, without degrading answer quality or exposing another customer's information? That is a measurable pilot question. "Can we become an AI-first institution?" is not.

---

## OUTSIDE THE BOUNDARY OF THE FIRST PILOT

### 01

Credit approval  
Owned by credit.

### 02

Beneficiary changes  
Owned by verification.

### 03

Payment initiation  
No payment tool at all.

### 04

Regulatory filings  
Owned by the filer.

# Three signals behind the recommendation.

What the current public evidence supports, and what it does not.

## 01 / A bank describes its internal uses

DBS's official account describes employee knowledge tools and service-assistance applications, alongside its broader AI program. These are company-reported examples. They support the relevance of internal assistance, not a promise that a small financial firm will achieve the same outcome. [1]

---

## 02 / Too much authority is a distinct risk

OWASP's excessive-agency category separates excessive functionality, permissions and autonomy. Giving an assistant unnecessary tools can create risk even when the intended task sounds routine. [2]

---

## 03 / The basic cyber program still matters

NIST's small-business CSF 2.0 resources provide a structured way to address governance, identification, protection, detection, response and recovery. AI should sit inside that operating discipline rather than replace it. [3]

---

## 04 / What these signals do not establish

None of the three is evidence that an assistant will produce a measurable return in a specific firm. They establish that internal assistance is a reasonable place to start, that unnecessary authority is a named risk category and that AI belongs inside an existing cyber program. The return has to be measured locally.

**The distinction that matters is not how capable the system is. It is what it is permitted to do when it is wrong.**

AGENTS STACK PERSPECTIVE

Research cut-off: 11 September 2026. Named-company examples are attributed reports; no independent audit or customer relationship is implied.

# Make policy knowledge easier to use.

A proposed first deployment for a lender, payments business, insurance intermediary or other financial SME.

## Choose a contained service problem

Select one internal question category: for example, where employees find the approved procedure for a routine service request. Build an assistant that returns the relevant passage, policy owner, version and effective date. Keep the production pilot inside the employee workflow; do not initially publish answers directly to customers.

## Put access before retrieval

Search results should be limited by the employee's approved role and the case context. Do not retrieve a broad pool of customer documents and then ask the model to hide the wrong ones. For the proposed design, authorization happens before content is available to the assistant.

## Separate general policy from case facts

An employee may be entitled to read a policy without being entitled to read every customer record. Use distinct sources and identities where appropriate. Ask whether a task genuinely requires personal information; an answer about the location of a procedure often does not.

## Make abstention part of the product

When two approved documents conflict, the answer should expose the conflict and route it to the policy owner. When a policy has expired, the assistant should not invent a replacement. A useful response can be: the evidence is insufficient, and this is the next approved person or process.

---

### THE DESIGN ORDER THAT KEEPS IT CONTAINED

#### 01

##### Authorize

Role and case context decide what exists.

#### 02

##### Retrieve

Only approved, in-scope documents.

#### 03

##### Draft

With passage, owner, version and date.

#### 04

##### Verify

The employee sends the final answer.

# Measure the whole task, then govern the release.

Response-generation time is the easiest number to collect and the least useful one to report.

## Evaluate the employee's whole task

Measure reading, source checking, correction and escalation, not just response-generation time. Inspect a sample of accepted drafts because employees can become accustomed to plausible wording. Test difficult edge cases as well as familiar questions.

Evaluation questions should be held apart from the examples used to configure the assistant. A system tuned on its own test set will report a result it cannot repeat in production.

## Govern the lifecycle

NIST's AI RMF is a voluntary approach to managing AI risk across the lifecycle. [5] Our proposed application is a short evidence record for each release: intended use, prohibited use, policy versions, evaluation results, owners, known limits and rollback instructions.

One page per release is enough. The value is that someone outside the build team can read it and understand what was approved.

## What to sample each month

Pull a fixed number of accepted drafts and re-check them against the cited passage. Acceptance rate alone measures convenience. Re-checking measures whether the citation was the right one, and whether it was still in force on the day it was used.

## Who signs the release

Name the policy owner and the service owner on the record. A release approved only by the team that built it is not an independent check.

**A cited answer is useful only when the source is current, applicable and available to that user.**

AGENTS STACK PERSPECTIVE

**Minimum release record:** intended use · prohibited use · policy versions · evaluation results · owners · limits · rollback.

# An attachment is not an authorization.

The finance workflow is where document understanding must stay separate from permission to act.

## An illustrative failure

A supplier message reaches an internal finance assistant. Attached material mixes legitimate invoice details with directions aimed at the assistant. The assistant has been given both document-reading access and a tool that can change payment details. This is a constructed example, not a claim that a named institution suffered this incident.

## Break the chain by design

OWASP's prompt-injection guidance distinguishes instructions arriving directly from a user from instructions embedded in untrusted material. [4] In our proposed pilot, external attachments remain untrusted inputs. They cannot grant authority, change the assistant's role or override an existing approval process.

## Remove unnecessary tools

A drafting assistant does not need a payment tool. A policy-search assistant does not need bulk customer export. A service summarizer does not need an administrator's identity. Design the tool set for the actual job, not for the most impressive demonstration. This applies OWASP's excessive-agency distinction. [2]

---

### PERMISSIONS TO WITHHOLD FROM A DRAFTING ASSISTANT

#### 01

##### Payment execution

Changing or initiating a transfer, in any form.

#### 02

##### Bulk customer export

Retrieval should be scoped to a single case.

#### 03

##### Administrator identity

A summarizer does not need elevated access.

# Make the human step worth having.

An approval that shows nothing is not a control. It is a record that someone clicked.

## **Make a human approval meaningful**

For any later high-impact workflow, the approver must see the action, the source evidence, the affected record and the proposed change. Do not reduce approval to a button beside an unexplained recommendation. Existing maker-checker, segregation-of-duties and independent verification procedures must remain authoritative.

## **Verify consequential changes outside the message**

For a proposed beneficiary-change workflow, verification should use the firm's existing trusted contact and approval process, not a phone number supplied in the same request. An AI-generated summary may assist review; it is not proof that the request is genuine.

## **Test the controls, not only the answers**

Attempt cross-customer retrieval using approved synthetic cases. Check whether expired access is revoked, prohibited tools remain inaccessible and logs can reconstruct a decision. ASD's harness guidance reinforces the importance of controls around connected actions. **[6]** Passing a finite test set is evidence, not a guarantee.

## **Log enough to reconstruct a decision**

Record which passages were retrieved, which identity requested them and what the assistant proposed. Logs should support review without copying unnecessary customer data into a second store that then needs its own controls.

## **Rehearse the revert**

Practice switching back to the manual workflow before the first employee uses the assistant. A control that has never been exercised is an assumption, and an incident is a poor time to discover that the old process no longer runs.

**The safest payment permission for a service-drafting pilot is no payment permission.**

AGENTS STACK PERSPECTIVE

# Prove the service benefit inside the boundary.

An Agents Stack proposed pilot: one internal policy category, a small employee group and no autonomous customer or financial actions.

## DAYS 1 TO 7

### Define the case set

Choose a policy owner and a service-team owner. Assemble approved documents and a representative set of historical questions using synthetic or appropriately authorized, minimized data. Remove unnecessary personal information. Set quality, escalation and access requirements before testing. Record the existing search-and-answer process.

## DAYS 8 TO 14

### Build and challenge

Configure retrieval against the approved collection. Include wrong-role users, expired policies, conflicting clauses and questions that cannot be answered. Test customer separation where case data is in scope. Any use of real personal data requires the appropriate privacy, contractual and jurisdictional review.

## DAYS 15 TO 21

### Work alongside staff

Use the assistant for drafts while employees remain responsible for the final response. Capture time to a correct completed task, unsupported statements, unnecessary escalations and reviewer corrections. Record why a draft was rejected. Do not treat high acceptance alone as proof of answer quality.

## DAYS 22 TO 30

### Review independently

Have a reviewer other than the builder inspect a held-out case sample and the permission tests. Compare net time and quality with the baseline. Present the result to service, risk and IT or security owners. Approve only the tested scope; new products, languages or external actions require a new assessment.

**Review group:** service · policy and compliance · privacy where needed · IT and security. Assign named owners.

# Two conditions decide whether the pilot continues.

Write both of them down before the first employee uses the assistant, not after the first incident.

## GATE ONE

### Control gate

Do not start live employee use until the intended role can access the intended data, prohibited access is blocked, logs are available and the manual workflow has been rehearsed. Assign a person who can pause the assistant immediately.

## GATE TWO

### Stop and investigate

Suspend the pilot after cross-customer disclosure, unauthorized action, a material policy error accepted as authoritative or inability to reconstruct a consequential output. Restore the existing workflow. Preserve relevant evidence through the firm's incident process and decide what must change before resuming.

### Why the gates come first

A pilot that has not agreed its stop conditions in advance will renegotiate them under pressure, usually in favor of continuing. Deciding while nothing is at stake is the cheapest governance available.

### What a pause should look like

A named person, reachable on the day, with the authority and the access to switch the assistant off without a committee. The manual process should already have been rehearsed, so reverting is routine rather than disruptive.

### What to record from day one

Time to a correct completed task, unsupported statements, unnecessary escalations, reviewer corrections and the reason each draft was rejected. Collecting this at the end of the pilot usually means collecting it badly.

### What the independent reviewer needs

A held-out case sample the builder has not seen, the permission test results and the baseline measurement taken in the first week. Without the baseline there is nothing to compare the result against.

**A useful pilot has a clear purpose and a risk owner.**

AGENTS STACK PERSPECTIVE

# Value the completed task, not the generated answer.

Illustrative internal-service economics in Singapore dollars. These are planning assumptions, not vendor prices, observed results or financial advice.

## ILLUSTRATIVE CAPACITY-VALUE EXAMPLE

# S\$1,450

Monthly net capacity value

# 5.5 months

Approximate setup recovery on a capacity-value basis

Planning assumptions in Singapore dollars; not necessarily cash savings.

### The base case

Assume 2,000 eligible inquiries a month and three minutes of net time saved per inquiry after review and correction. That is 100 hours. Assume 70% can be redeployed usefully: 70 hours. At an assumed fully loaded S\$35 per hour, the monthly capacity value is S\$2,450.

### Include the cost of keeping it useful

Deduct an assumed S\$1,000 monthly operating cost, including hosting, routine evaluation and policy upkeep, to obtain S\$1,450 net capacity value. An assumed S\$8,000 setup cost is recovered in about 5.5 months on this simplified capacity-value basis, not necessarily in cash.

## THE FOUR ASSUMPTIONS BEHIND THE FIGURE

### 2,000

Eligible inquiries a month  
Assumed volume.

### 3 min

Net time saved per inquiry  
After review.

### 70%

Usefully redeployed  
Not all of it is recovered.

### S\$35

Fully loaded hourly cost  
Assumed, not a quote.

# The sensitivity that breaks the case.

Halve the time saved and the same pilot takes roughly six times longer to pay for itself.

## Do not price away a control failure

Averages can conceal a serious access or policy error. Quality and permission gates should be pass or fail conditions for the approved scope, not a small discount to the labor benefit. A fast but unacceptable process is not a positive business case.

## The sensitivity that matters

At 1.5 minutes saved rather than three, and with the same 70% useful redeployment, capacity value falls to S\$1,225. After S\$1,000 cost, only S\$225 remains. The setup recovery period becomes about 35.6 months. Measure the time saved after review, not before it.

## What to ask a vendor

How is role-based access enforced at retrieval and tool execution? Which party can inspect or retain our data? Can we test model changes before promotion? Can we export the source corpus and evidence? What is the deletion process, including backups and subcontractors? Request contractual answers, not a slide labeled “enterprise secure”.

## The purchase decision

Compare the assistant with improved search, standard response templates and workflow fixes. Budget for integration, testing, training and exit. Do not assume that an existing software subscription includes the needed connectors, assurance work or ongoing content management.

Original calculation:  $2,000 \times 3 \div 60 \times 70\% \times S\$35 - S\$1,000 = S\$1,450$  per month. Downside case: S\$225 per month, about 35.6 months to recover setup. Tax, financing and discounting are excluded.

# Put the risk owner in the room early.

A leadership conversation that turns AI interest into a defined, testable proposal.

## Set the meeting up properly

Send the boundary question in advance and ask each owner to arrive with one answer. A first AI conversation held without the risk owner produces a proposal that has to be rebuilt once security and privacy see it.

## Ask the service team

Which questions recur often enough to justify a pilot? Which answers are hard because information is fragmented, and which are hard because judgment is required? Start with the first category. Keep the second visible as an escalation path rather than silently asking the model to improvise.

## Ask the policy owner

Who approves the source documents, and how quickly do changes reach employees? An assistant built on stale policy will distribute stale policy more efficiently. Define one source of approved truth and an update owner before debating model choice.

## Ask technology and security

Can we demonstrate that an employee sees only authorized material? Can the assistant operate without transaction privileges? Can access be revoked promptly? Can a person reconstruct the information behind an answer without exposing unnecessary customer data in logs?

---

## WHO HAS TO BE IN THE ROOM

### 01

#### Service owner

Brings the recurring question set.

### 02

#### Policy owner

Owens the approved source of truth.

### 03

#### Risk and security

Owens access, logging and the pause.

### 04

#### Finance sponsor

Funds it and can stop it.

# The boundary is the brief.

Finish the meeting with a scope that risk, privacy and finance have each signed.

## Ask risk, privacy and legal

Which jurisdictions, customer types and outsourced-service terms apply? Does the proposed use create decisions requiring explanation, consent, record retention or special review? This magazine offers design questions, not a declaration of compliance with MAS, RBI or another regulator. Obtain advice for the actual entity and use case.

## Ask finance and the sponsor

What does the pilot need to prove within 30 days? Who funds continued document maintenance? What happens if the measurable value is smaller than expected? Require a stop decision to be as legitimate as a scale decision.

## The question to take to your team

Which internal service task can AI assist while remaining technically unable to move money, alter a customer's rights or expose someone else's records? That boundary is the starting brief for a useful proposal.

## What to leave the meeting with

A named use case, a named risk owner, an agreed stop condition and a date for the review. If any of the four is missing, the proposal is not yet testable, and the next 30 days will produce an opinion rather than a result.

## Write the scope down in the room

Leave with one paragraph that states the question category, the data the assistant may reach, the actions it may not take and the named owner of each. If it cannot be written in a paragraph, it is not yet a scope.

## Book the review before the build

Put the day-30 review in the calendar at the start, with the reviewer named. A review scheduled after the results arrive tends to be scheduled around them.

**Human accountability should be a system property, not a sentence at the end of a policy.**

AGENTS STACK PERSPECTIVE

**Next edition:** maritime and logistics. Shipment exceptions, trusted documents and operational resilience across connected partners.

# Evidence you can inspect.

Numbered references match the articles. All sources were checked on 11 September 2026.

---

## [1] Inside DBS' AI journey, from use cases to workforce transformation

DBS Bank · official LinkedIn account · Exact publication date not displayed  
First-party company account. Company-reported examples, not independently audited evidence of SME returns.

---

## [2] LLM06:2025 Excessive Agency

OWASP GenAI Security Project · 2025 risk taxonomy  
Community security guidance. Highlights excessive functionality, permissions and autonomy.

---

## [3] NIST Cybersecurity Framework 2.0 for Small Business

NIST · Page updated 17 April 2026  
Small-business guidance. A framework and quick-start resources, not a compliance certification.

---

## [4] LLM01:2025 Prompt Injection

OWASP GenAI Security Project · 2025 risk taxonomy  
Community security guidance. Risk taxonomy, not a certification or proof that a particular product is secure.

---

## [5] AI Risk Management Framework

NIST · AI RMF 1.0 released 26 January 2023  
Voluntary risk framework. The Generative AI Profile followed in July 2024. Verify future revisions before citing.

---

## [6] Agentic AI harnesses

Australian Signals Directorate / ACSC · 11 September 2026  
Government guidance. New guidance at the research cut-off; not an industry adoption survey.

---

### EDITORIAL NOTES

Reported examples are attributed to their sources. The SME scenarios, pilot plans and decision questions are original editorial proposals, not verified Agents Stack customer results. All cost examples use explicit assumptions, not market prices or promised returns.

This magazine does not provide legal, financial or safety certification. Publication requires a named editor's approval and a final source check. No affiliation or endorsement by cited organizations is implied.



NEXT IN THE SERIES

# 2 October • Maritime & logistics

Shipment exceptions, trusted documents and operational resilience across connected partners.



**Agents Stack**

One industry. Useful AI. Accountable cybersecurity.

[agentsstack.com](https://agentsstack.com)