

Leadership Briefing Agent

Evidence-led weekly prioritisation for senior leaders

Status	Functioning prototype - acceptance tested
Platform	Microsoft Copilot Studio, Power Automate, Outlook, Word Online (Business) and OneDrive for Business
Business area	School / trust leadership, executive planning
Control model	AI retrieves, analyses and prioritises; evidence rules constrain claims; human leader remains accountable for decisions

Contents

- [1. Solution at a glance](#)
- [2. Executive summary](#)
- [3. Business problem and users](#)
- [4. Solution design and architecture](#)
- [5. Agent behaviour and evidence grounding](#)
- [6. Governance and risk controls](#)
- [7. Validation and testing](#)
- [8. Success measures and business value](#)
- [9. Reflections and next steps](#)

Use Case Purpose

This case study documents the problem, design decisions, governance controls, test matrix and value hypothesis behind the working prototype. It uses synthetic Clifton Bridge Education data and scenarios; it is not presented as a live school deployment.

1. Solution at a glance

The one-page overview below summarises the challenge, solution, value, Microsoft stack and end-to-end user journey. It is included first so a reader can understand the proposition before moving into the detailed build evidence.

Leadership Briefing Agent

Turns the Outlook calendar and inbox evidence into a prioritised, editable weekly leadership briefing - automatically generated as a branded Word document and securely linked back to the user.



The Challenge

Senior leaders must combine diary commitments, emails, deadlines and preparation requirements manually. Important dependencies can be hidden across separate Microsoft 365 sources.



The Solution

Copilot Studio retrieves calendar and recent email evidence, cross-references it, prioritises leadership attention and passes eight structured sections to Power Automate for Word generation.



The Value

Creates a repeatable weekly briefing process focused on decisions, deadlines, dependencies and reusable preparation - reducing manual synthesis while keeping leadership judgement with the user.

End-to-End User Journey

1 → Leader asks for this week / next week briefing

2 → Outlook Calendar and Inbox analysed

3 → Evidence cross-referenced and prioritised

4 → Power Automate generates Word doc

5 → Organisation-only share link returned

6 → Staff review & edit if necessary



Outlook



Power Automate



Copilot Studio



Word Online (Business)



OneDrive for Business



“Turning the week ahead into clear leadership priorities.”



FUNCTIONING PROTOTYPE

• ACCEPTANCE TESTED



The AI Elective

Figure 1. Leadership Briefing Agent overview.

2. Executive summary

The Leadership Briefing Agent is a functioning prototype designed to reduce the manual synthesis senior leaders perform before a working week. Rather than summarising a diary or inbox independently, it combines live Outlook calendar and email evidence, identifies what requires preparation, decisions or action, and produces a concise executive briefing.

The agent is deliberately evidence-led. Calendar data is authoritative for scheduled commitments; email evidence adds context, deadlines, requests and risks. The instructions prohibit invented meetings, dates, deadlines, actions or organisational facts, distinguish explicit deadlines from meeting start times, and require unsupported interpretation to be labelled as inference.

The final workflow automatically creates an editable branded Word document in OneDrive and returns an organisation-only sharing link in chat. The leader remains responsible for judgement, prioritisation and any subsequent decisions.

What makes the prototype credible

It demonstrates connected Microsoft 365 retrieval, cross-source reasoning, prioritisation rather than summarisation, explicit evidence boundaries, automated document generation, secure delivery and end-to-end acceptance testing.

3. Business problem and users

Senior leaders often begin the week by manually reconciling calendars, emails, deadlines, papers and preparation requirements. The information exists, but it is fragmented. A diary shows what is scheduled; email explains why some meetings matter; neither source alone creates a prioritised leadership view.

- Manual synthesis consumes leadership time before the substantive work begins.
- Important deadlines or dependencies can be buried in email and disconnected from the meetings they support.
- A full diary can obscure the small number of commitments that genuinely require preparation, a decision or escalation.
- Preparation is often duplicated because links between workstreams are not made explicit.
- Generic AI summarisation can create risk if it invents context, treats a meeting time as a deadline, or converts an email reference into a scheduled event.

Primary users are senior school and trust leaders. Executive assistants or leadership support teams could also benefit from the same pattern, provided permissions and ownership are configured appropriately.

4. Solution design and architecture

The architecture separates evidence retrieval, agent reasoning, document generation, storage and delivery. This makes the workflow easier to govern and allows the briefing logic or template to evolve without rebuilding the whole solution.

#	Component	Role in the solution
1	Natural-language request	Leader asks for a briefing for this week, next week or another inferable period.
2	Copilot Studio orchestration	Derives the requested Monday-Friday period and coordinates retrieval and analysis.
3	Outlook calendar	Provides authoritative evidence for scheduled commitments.
4	Recent Outlook email	Provides context, requests, explicit deadlines, risks and preparation requirements.
5	Evidence-led analysis	Cross-references sources and prioritises significance, decisions, dependencies, pressure and reusable preparation.
6	Create Leadership Briefing Document	Power Automate receives eight briefing sections from the agent.
7	Word Online (Business)	Populates the branded executive Word template.
8	OneDrive for Business	Creates the editable briefing file and organisation-only sharing link.
9	Chat delivery	Returns only the link to the generated Word briefing; the full briefing is not printed in chat.

5. Agent behaviour and evidence grounding

The instruction architecture is built around a decision sequence: determine the period, retrieve evidence, cross-reference it, analyse significance, then populate the document. This prevents the agent from treating retrieval as the final task.

- Calendar data is authoritative for scheduled commitments; an email mentioning a meeting does not prove the meeting exists.
- Only significant commitments are included: preparation, decisions, deadlines, strategic/operational importance or material pressure.
- Routine or low-information meetings are omitted unless their timing itself creates a conflict or workload pressure.
- Meeting start times are never presented as deadlines. Where preparation is required without an explicit deadline, the briefing uses 'Before [meeting]'.
- Unsupported conclusions or predicted consequences must be labelled 'Inference:'.
- If calendar or email retrieval is unavailable, the agent states the limitation rather than reconstructing missing evidence.
- The briefing ends with exactly three practical 'Do First' actions.
- For WeekCommencing, the tool receives only the formatted date because the Word template already contains the words 'Week commencing'.

Eight structured document inputs

WeekCommencing | ExecutiveSummary | WeekAtAGlance | FivePriorities | ActionsDeadlines | PressurePoints | Connections | DoFirst

6. Governance and risk controls

Governance is built into the operating model rather than added after generation.

Risk / control area	Mitigation in the prototype
Hallucinated commitments	Scheduled commitments must come from retrieved calendar evidence; email references cannot create diary events.
Invented deadlines	Only explicit deadlines are stated as deadlines; meeting times are not reclassified.
Overconfident interpretation	Conclusions not explicitly evidenced are labelled as inference.
Low-value noise	Routine or low-information meetings are omitted from the executive briefing.
Cross-period contamination	The requested week remains primary; future-period email evidence is included only when it creates preparation during the requested week.
Retrieval failure	Calendar/email failures are surfaced explicitly; missing evidence is not reconstructed from another source.
Data exposure in chat	Raw EmailData, HTML, JSON, connector metadata, technical IDs and source hyperlinks are prohibited.
Document access	Generated files are stored in OneDrive for Business and shared using an organisation-only view link.
Human accountability	The agent recommends and prioritises; the senior leader remains responsible for decisions and actions.

7. Validation and testing

The updated acceptance matrix tests retrieval, evidence boundaries, prioritisation, failure handling, document generation and final delivery. The final end-to-end run generated the editable Word briefing for the week commencing 31 August 2026 and verified the two final refinements: the date heading was no longer duplicated and the low-information Review Meeting was omitted.

#	Scenario	Capability	Observed evidence	Result
1	Requested period - next week	Date interpretation	Correct Monday-Friday period derived using UK local time.	PASS
2	Calendar chronology	Calendar retrieval	Significant events presented by day and start time; scheduled commitments grounded in calendar evidence.	PASS
3	Email context for verified meetings	Cross-source reasoning	Email requests and preparation requirements incorporated without creating new calendar events.	PASS
4	Explicit deadline vs meeting time	Deadline control	Wednesday 16:00 and 17:00 deadlines surfaced; meeting start times not mislabelled as deadlines.	PASS
5	Low-information Review Meeting	Relevance filtering	Routine meeting with no evidenced purpose/preparation omitted from final Week at a Glance.	PASS
6	Five priorities	Prioritisation	Five leadership priorities synthesised from significance, decisions, deadlines and dependencies rather than repeating meetings.	PASS
7	Pressure points	Risk/dependency analysis	Material deadline, dependency, workload and risk items identified without padding the section.	PASS
8	Connections	Reusable preparation	Linked Year 11, AI and attendance workstreams identified so evidence/preparation can be reused.	PASS
9	Do First	Actionability	Exactly three practical next actions generated and tied to upcoming decision points.	PASS
10	Week commencing field	Template integration	Only '31 August 2026' passed into the field; final document displayed 'Week commencing 31 August 2026' once.	PASS
11	Word document generation	Power Automate / Word	Eight structured sections populated into the executive Word template and saved as an editable .docx.	PASS
12	Secure delivery link	OneDrive / chat	Organisation-only view link returned in chat and opened the newly generated briefing.	PASS

Final formal validation result: 12/12 defined scenarios passed on the frozen prototype configuration. This demonstrates successful performance against the defined acceptance matrix; it is not evidence of universal accuracy or production readiness.

Regression refinement

The final UAT cycle exposed two presentation/relevance issues rather than workflow failures: duplicated 'Week commencing' wording and inclusion of a low-information meeting. The instructions were tightened, republished and re-tested successfully.

8. Success measures and business value

The prototype demonstrates technical feasibility and controlled behaviour. A production pilot should measure whether the briefing actually reduces preparation effort and improves leadership focus rather than assuming value from generation alone.

Measure	Prototype target	How to evidence in a pilot
Weekly synthesis time	Target: meaningful reduction	Compare manual calendar/inbox synthesis baseline with agent-assisted preparation.
Material editing required	Target: low	Track substantive edits made to generated briefings before use.
Evidence quality	Target: no unsupported scheduled commitments or deadlines	Structured QA sample against calendar/email evidence.
Priority usefulness	Target: leaders judge priorities/action list useful	Short rating after each pilot briefing.
Missed critical actions	Target: none in sampled weeks	Compare briefing against known explicit deadlines and preparation requests.
Adoption / repeat use	Target: sustained weekly use if valuable	Usage telemetry plus qualitative interviews.

A useful value hypothesis is the reduction of repeated weekly synthesis across diary, inbox and preparation notes. The stronger strategic value is the reusable adoption pattern: retrieve from approved Microsoft 365 sources, constrain claims to evidence, prioritise rather than merely summarise, automate the artefact, and retain human accountability.

9. Reflections and next steps

The build reinforced that the difficult part of an executive briefing agent is not producing fluent prose; it is controlling evidence, relevance and delivery. The most useful refinements came from testing the boundaries between a scheduled commitment, an email reference, an explicit deadline and an inferred preparation need.

- Agent instructions work better as a decision hierarchy than as a long list of stylistic rules.
- Calendar and email sources need different authority: calendar establishes commitments; email enriches context.
- Relevance filtering is essential. A complete diary is less useful than a selective leadership briefing.
- Automated Word generation makes the output operationally useful because it can be edited, stored and reused outside the chat interface.
- Organisation-only link delivery provides a cleaner user experience than printing the briefing in chat.
- A production pilot would still require permissions review, data-protection assessment, monitoring, operational ownership and a larger ongoing evaluation set.

Recommended next step: run a small time-boxed pilot with a defined group of senior leaders using their approved Microsoft 365 accounts. Capture baseline weekly preparation time, edit rate, evidence-quality checks, usefulness ratings and exception types. Use those measures to decide whether to scale, refine or stop.