

Lab 6: Inbox Triage with an AI Step

AI Administrator · Day 3 · Session 6 · Adding AI Steps to Automations

Participant lab handout

Objective

Build an inbox triage in n8n: a new e-mail arrives, an **AI step** classifies it as **billing**, **technical**, **sales** or **unsure**, a **Switch** routes it, the three known classes get a Gmail label, and the **unsure** ones are written to a Google Sheet that a person reads. Then test it with ten e-mails, improve the prompt, and note the cost per run.

Time: 40 minutes in breakout rooms. Finish at home before Session 7 if needed. **Deliverable:** below.

What you need

- Your **n8n Cloud** workspace from Session 5 (check the days left at the top).
- The **same Google account** connected to Gmail and Google Sheets in n8n.
- **Day3_S6_Send_Test_Emails.json** from the Zoom chat (it sends you the ten test e-mails).
- No API key: the Gemini node uses your trial's **Gateway credits** (the trial starts with \$2.00; this lab uses a few cents). If your Gemini node shows no "Gateway credits" option, follow **Day3_S6_Gemini_Key_Guide.pdf**.

Words used by today's n8n: **Execute step** runs one node, **Execute workflow** runs the whole workflow once, **Publish** makes it live. The module page says "Activate"; it means Publish.

Before you build (5 minutes)

1. **Gmail labels.** In Gmail, left sidebar, **Labels**, **+**. Create three labels named exactly **billing**, **technical**, **sales** (lower case).
2. **The review sheet.** New Google Sheet in the same Google account. Rename the tab to **Triage review**. Row 1:

date	from	subject	snippet	link
------	------	---------	---------	------

3. **The helper.** In n8n: **Overview**, **Create**, **Import from file**, choose **Day3_S6_Send_Test_Emails.json**. Open **Send to myself**, connect your Gmail credential and put **your own Gmail address** in **To**. **Do not run it yet** (see Trap 3).

Steps

1. Create the workflow and the trigger

New workflow, rename it **M6 Inbox triage**. Add first step: **Gmail, On message received** (Gmail Trigger).

- Credential: your Gmail account. Poll Times: **Every Minute**. Event: **Message Received**.
- **Simplify: OFF**. With Simplify on, the trigger gives *Subject* and *snippet* but **no body text** (Trap 1).
- Filters, Search: `subject:"[S6 test]"`. This is the cheap filter: only the course e-mails reach the paid AI step.

Send yourself **one** e-mail with a subject starting `[S6 test]` and any sentence in the body. Click **Fetch Test Event**. The output must show **subject** and **text**.

2. Add the AI classification step

Add **Basic LLM Chain** (search “LLM”). Source for Prompt: **Define below**. Paste:

```
You classify incoming customer emails for a small software company.
Read the email below and reply with exactly one word from this list:
billing, technical, sales, unsure
Rules:
- billing: invoices, payments, refunds, subscription changes
- technical: errors, login problems, how-to questions about the product
- sales: pricing, demos, new licences, partnerships
- unsure: anything else, or if two labels fit equally well
Reply with the single word only. No punctuation, no explanation.
Email subject:
Email body:
```

Then drag **subject** from the input panel to the end of the *Email subject:* line, and **text** to the end of the *Email body:* line. They appear as `{{ $json.subject }}` and `{{ $json.text }}`.

Under the chain, click **Model +**, choose **Google Gemini Chat Model**. Credential: **Gateway credits**. Model: pick **models/gemini-2.5-flash** from the list (do not type it: Trap 2). Options, **Sampling Temperature: 0**.

Click **Execute step** on the chain. The output is one word.

3. Clean the answer

Add **Edit Fields (Set)**, Manual Mapping. Add a field named `label`, type String, value (switch to Expression):

```
{{ $json.text.trim().toLowerCase() }}
```

This removes spaces and capitals the model may add, so the router can compare exactly.

4. Route with a Switch

Add **Switch**, Mode **Rules**. Three rules on `{{ $json.label }}`, **String, is equal to**: `billing`, `technical`, `sales`. Turn on **Rename Output** in each rule and name it like the label. In Options add **Fallback Output: Extra Output** and rename it `unsure`. Everything the three rules do not match, including a word the model invented, goes there.

5. Label the known classes

On the **billing**, **technical** and **sales** outputs add **Gmail, Add label to message**:

- Message ID (Expression): `{{ $('Gmail Trigger').item.json.id }}` (use your trigger node's name).
- Label Names or IDs: **pick your label from the list**.
- Settings: **Retry On Fail** on.

6. Send unsure to a human queue

On the **unsure** output add **Google Sheets, Append Row**, your sheet, tab **Triage review**, **Map Each Column Manually**:

Column	Value (drag from the Gmail Trigger in the input panel)
date	date
from	from, text
subject	subject
snippet	text
link	<code>https://mail.google.com/mail/u/0/#all/</code> followed by id

7. Test with ten e-mails

1. **Publish** the triage workflow (top right). Give the version a name, for example `v1`.
2. Open the helper, **Execute workflow**. Ten e-mails arrive: three **billing**, three **technical**, two **sales**, two **ambiguous**.
3. Within two minutes open **Executions** in the triage workflow: one run with **10 items**. Expected: **8 labelled, 2 rows in Triage review**.
4. An e-mail in the wrong place? Tighten the definitions, do not add a rule for one e-mail. In our test, *"Your talk at the conference"* (asks for slides **and** student pricing) came back **sales** in 2 of 4 runs. Adding this line under the **unsure** rule fixed it in 4 of 4:

– unsure: also when the email mixes a personal or unrelated request (slides, a talk, an event) with a business question

Publish again (version `v2`), run the helper again, and compare.

8. Note the cost

In **Executions**, open the run and click the **Google Gemini Chat Model** node: **tokenUsage** shows prompt and completion tokens per e-mail (ours: 183 in, 1 out). On Gemini 2.5 Flash list prices that is about \$0.00006 per e-mail, 6 cents per 1,000. Note it, and the Gateway credits balance in the Gemini node, for your Module 12 tracker. If you use your own key, set a monthly budget alert in AI Studio (**Spend**).

Three traps we hit while testing (12 October 2026)

1. **Simplify hides the body.** Simplify on: no text field and Subject with a capital S. The prompt gets empty values and the model guesses.
2. **A typed model name.** We typed `models/gemini-flash-latest`; the run failed with “*Gateway credits don’t currently support this operation*”. Pick the model from the list.
3. **Old e-mails are not new.** *Execute workflow* on a polling trigger fetches only the latest matching e-mail; a published trigger only reads e-mails that arrive after publishing. Publish first, then run the helper.

Deliverable

Saved as **M6-inbox-triage**:

- a screenshot of the published workflow’s execution with **10 items** (8 labelled, 2 unsure);
- a screenshot of the **Triage review** sheet with the two ambiguous e-mails;
- one line: tokens per e-mail and your cost estimate per 1,000 e-mails;
- the exported JSON (menu, **Export**).

If you do not need the triage until next session, **unpublish** it: a polling trigger keeps checking Gmail every minute.