OWNERSHIP

Eight days to retrain, then it stops answering

OpenAI stops new fine-tuning jobs for existing customers on 6 January 2027 and shuts down gpt-4.1 fine-tunes eight days later, on 14 January 2027. Microsoft's schedule for the same base model version puts training retirement at no earlier than 14 April 2027 and deployment retirement at 14 October 2027. The artefact's life is a contract term.

5 November 2026 | Strategic Intelligence Division | Ownership

> **The obligation, and the date on it**
>
> 6 January 2027 is the last date on which an eligible existing customer can create a fine-tuning job. OpenAI's page: “Active existing customers will no longer be able to create new fine-tuning jobs on this date. Inference on fine-tuned models will be disabled only when the underlying base model is deprecated.” Source: OpenAI API deprecations, Update to OpenAI's self-serve fine-tuning.

Both halves of that sentence land eight days apart for two of the three base models OpenAI's supervised fine-tuning guide still names. Training ends on 6 January 2027. gpt-4.1 and gpt-4.1-mini, and their fine-tuned versions, shut down on 14 January 2027. The third, gpt-4.1-nano, and its fine-tuned versions shut down on 23 October 2026. A fine-tune created on 5 November 2026 therefore has at most 70 days of life, and a fine-tune already in production has 8 days in which to be retrained, re-evaluated and redeployed. This reaches you if a workflow in production calls a fine-tuned OpenAI model, or if the 2027 platform commitment you are about to sign lists fine-tuning as a capability you keep. The decision belongs to the Chief Information Officer or Chief AI Officer, with the General Counsel on what transfers and the Chief Financial Officer on the re-platforming line.

One base model version, two schedules, and the interval a buyer gets to work with

Step in the artefact's life	Direct from OpenAI	Microsoft Foundry
Last date a new training job can be created	6 January 2027, eligible existing customers	No earlier than 14 April 2027, existing customers
Base model inference	14 January 2027, with its fine-tuned versions	14 April 2027 for gpt-4.1-mini, all customer-facing offers
Fine-tuned model stops answering	14 January 2027	14 October 2027
Days between the training gate and the shutdown	8	183

OpenAI API deprecations, Update to OpenAI's self-serve fine-tuning and the 2026-10-14 GPT-4 series entry; OpenAI supervised fine-tuning guide; Microsoft Learn, Model retirements, fine-tuned models table with footnote 1, last updated 24 July 2026; Microsoft Learn, Fine-tune a model with supervised fine-tuning in Microsoft Foundry, last updated 8 October 2026. Day counts are Mosaic Singularity analysis, November 2026.

The floor is conditional, and the date on your model is what binds

The page sets a floor and then writes the exception into the same paragraph: “Unless safety or compliance concerns require a faster timeline, we provide the following minimum notice periods before model retirement”, and for generally available models the floor is “At least 6 months”. The 2026-10-14 entry announces the GPT-4 series batch on 14 October 2026 and removes it on 14 January 2027, 92 days later, and the entry does not say whether the exception applied. The entry below it, dated 2026-10-01, runs three other models from 1 October 2026 to 1 April 2027, which is six months. Whichever way that reads, the interval a buyer can plan against is the one printed against

its own model, not the policy printed above it.

The condition attached to the capability moved earlier than the capability did, and it moved in steps. From 7 May 2026, organisations that had not previously run fine-tuning could not create jobs. From 2 July 2026 the bar became recent activity: an organisation that had not run inference on a fine-tuned model in the past 60 days could no longer create jobs. On 6 January 2027 the bar reaches every customer. Microsoft's table carries the same instinct in footnote 1, where the later training dates apply “For existing customers only” and for anyone else training retires at the base model retirement date. Two vendors, one drafting habit: the withdrawal is graded by how long you have been a customer, not by what you can do with the product.

The dormancy rule has already taken the long tail

A fine-tune that no system calls for 60 days stops being maintainable. The model keeps answering, but the organisation loses the right to create jobs, so no new version of it can be trained and no defect in it can be repaired. That is an unrecoverable asset with no retirement date attached to it, and the rule has applied since 2 July 2026. The catalogue is the second half. The supervised fine-tuning guide names three base models, gpt-4.1-2025-04-14, gpt-4.1-mini-2025-04-14 and gpt-4.1-nano-2025-04-14, and every one of them is scheduled to shut down, with its fine-tuned versions. The vision fine-tuning row on the model optimization guide names a fourth, gpt-4o-2024-08-06, which is in the same 14 January 2027 batch. No base model the fine-tuning guides name outlives 14 January 2027, so choosing a base is not a way to extend the artefact's life.

8
Days between the last training job and the shutdown, direct from OpenAI

70
Life remaining for a fine-tune created on 5 November 2026

183
Days between the same two steps on Microsoft Foundry

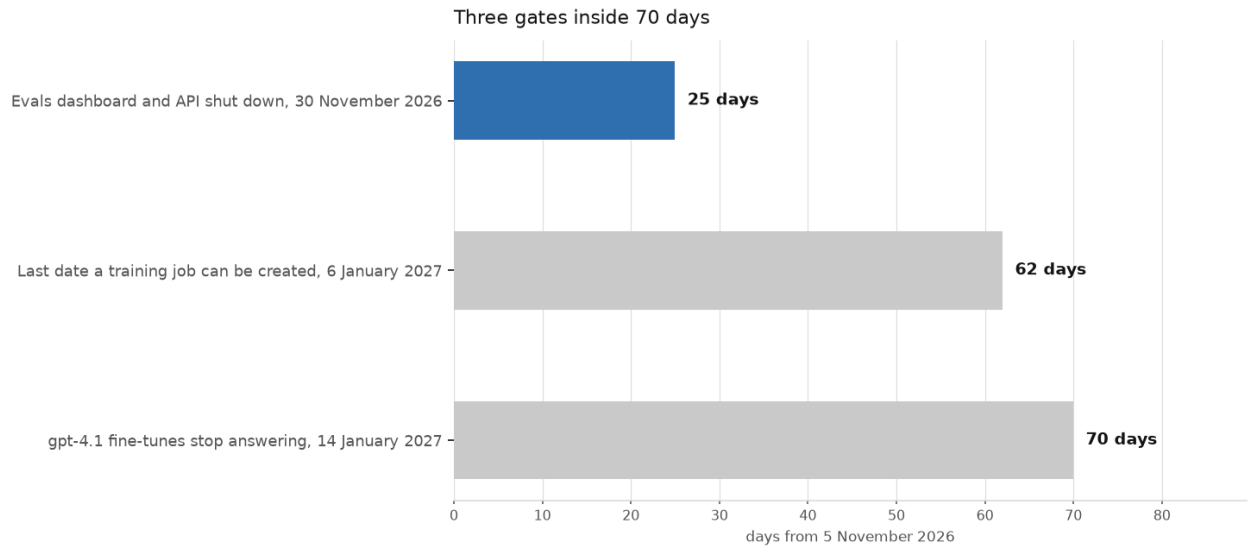
356
Days between the two rails for gpt-4.1-nano fine-tunes

OpenAI API deprecations and OpenAI supervised fine-tuning guide; Microsoft Learn, Model retirements, fine-tuned models table. Day counts are Mosaic Singularity analysis, November 2026, from the dates those pages print.

The same artefact gets 183 days on the other rail

Microsoft sells fine-tuning on the same base model versions and publishes a separate schedule for what comes out of it. “Fine-tuned models retire in two phases: training and deployment.” For gpt-4.1 and gpt-4.1-mini the training retirement is “No earlier than 2027-04-14” for existing customers, and the deployment retirement is 14 October 2027. The interval between those two steps is 183 days, against 8 on the direct rail. The fine-tune also outlives Azure's own base model inference for gpt-4.1-mini, which retires on 14 April 2027. The sharpest comparison is the base model that has already gone: gpt-4.1-nano-2025-04-14 fine-tunes stopped answering for a direct customer on 23 October 2026, and a fine-tune of the same base model version on Foundry has a deployment retirement of 14 October 2027. That is 356 days between two schedules for one version number.

Three gates, 70 days from 5 November 2026: evaluation closes first, then training, then inference



OpenAI API deprecations, Evals platform entry and Update to OpenAI's self-serve fine-tuning; OpenAI supervised fine-tuning guide. Day counts are Mosaic Singularity analysis, November 2026.

The sequence is the finding. The tool for scoring a replacement model goes dark on 30 November 2026, 37 days before the last date a training job can be created and 45 days before the fine-tunes stop. An organisation that intends to re-qualify a replacement inside its own gates has 25 days to run the evaluation that decides it, on a product whose dashboard and API are scheduled to shut down on the day the answer is due. The three dates sit in three separate entries on one page. No sentence on that page puts them together, and the entry that governs the artefact treats the shutdown as a consequence of the model table above it.

What the buyer owns is a deployment with an end date on it

The artefact does not travel. The supervised fine-tuning guide describes uploading a dataset, OpenAI using it “to update the model’s weights”, and deploying the resulting model. Across the fine-tuning and model optimization guides, no export or download of fine-tuned weights is documented. What the buyer holds is a deployment on the vendor's platform with a date attached to it, and a training set it can carry anywhere. The vendor has carved an existing artefact out before. In the 2024-08-29 entry, ending new training on babbage-002 and davinci-002, the page states that “Fine-tuned models created from these base models are not affected by this deprecation, but you will no longer be able to create new fine-tuned versions with these models.” No equivalent sentence appears for the 2027 dates. The page ties inference to the base model instead, and the guide names the days.

> *“Fine-tuned versions of these models follow the same shutdown schedule as their underlying base models.”*
>
> OpenAI API deprecations, 2026-10-14 GPT-4 series model deprecations

Microsoft gains the demand, and its date is printed too

Name the winner precisely. Microsoft takes the capability and the demand behind it, on a schedule that gives the buyer 273 more days for a gpt-4.1-mini fine-tune than the direct rail does, and 343 days of life measured from 5 November 2026. Amazon sells the capability on a catalogue that carries none of these dates. In Amazon's Bedrock User Guide, the page on customizing a model states that Bedrock “currently provides” supervised fine-tuning, reinforcement fine-tuning and distillation; it charges for training by the token processed and for storage per month per model; and the page carries no revision date and names no OpenAI base model. None of the schedule above

attaches to what comes out of it. Both are franchises with an expiry printed on them, and Foundry’s is 14 October 2027. The loser is the enterprise that wrote a fine-tuning workflow into a platform commitment and now has 8 days between the ability to rebuild the artefact and the loss of it.

What to decide, and by when

This is one board decision, and it arrives as a capital request rather than a technology item: fund the rebuild of the fine-tuned estate on a rail whose printed date is later, or write the artefacts off and carry the re-engineering. Because the right to create a training job ends on 6 January 2027, the paper has to be with the board at its last meeting before that date, put there by the Chief AI Officer with the Chief Financial Officer's number attached and the General Counsel's answer on what transfers at retirement.

- Chief AI Officer, by 6 January 2027: the register of every production fine-tune with its base model version and that version's shutdown date, signed off before the training gate closes, because no new training job can be created after that date.
- General Counsel, before the next platform renewal: a schedule annex that names the retirement date of the fine-tuned model separately from its base model, and a written answer on what happens to the trained weights at retirement, because no export path is documented.
- Chief Financial Officer, in the FY27 plan: the re-platforming line for every fine-tune that must survive 14 January 2027, priced against Foundry's dates, which retire training no earlier than 14 April 2027 and the deployment on 14 October 2027.
- Programme owner of a fine-tune that nothing has called for 60 days: treat it as unrecoverable now rather than on 14 January 2027, because the dormancy rule has already removed the right to rebuild it.
- Anyone with no fine-tune in production: nothing to do before 6 January 2027, and the date to watch is 14 April 2027, when the rail that still sells the capability starts retiring its training path.

What to watch, and when

- 30 November 2026: the Evals dashboard and API are scheduled to shut down. If they do, the 25-day window for scoring a replacement closes and the compression in this piece holds.
- 6 January 2027: if job creation continues past that date, or a documented route to export fine-tuned weights appears, the reading fails and the artefact becomes portable.
- 14 January 2027: if gpt-4.1 or gpt-4.1-mini fine-tunes still answer, the vendor is not applying its own rule and the 8-day interval describes a schedule rather than a product boundary.
- 14 April 2027: if Foundry retires training earlier than the date it publishes, the 183-day interval was disclosure rather than commitment, and the contract reading weakens to a difference in transparency.

Every obligation in this piece is printed on a document you can pull

- OpenAI, API deprecations. Entries of 22 April 2026, 7 May 2026, 3 June 2026, 1 October 2026 and 14 October 2026.
- OpenAI, Supervised fine-tuning guide. No revision date on the page; it sets the 6 January 2027 training deadline and the 23 October 2026 and 14 January 2027 shutdowns.
- Microsoft, Model retirements for Azure OpenAI. Last updated 24 July 2026: the fine-tuned model table and its footnote on existing customers.
- Microsoft, Fine-tune a model with supervised fine-tuning in Microsoft Foundry. Last updated 8 October 2026.
- Amazon Web Services, Amazon Bedrock User Guide, Customize your model to improve its performance for your use case. No revision date on the page; states its three methods in the present tense.