OFFICIAL USER GUIDE · VERSION 1.8.0



Local Waifu
Complete User Guide

Set up your companion, install models, create and move characters, protect your data, and solve common problems.

Everything she needs. All in one place.

PLATFORMS	macOS · Windows
EDITIONS	Direct · Store/Unlocked
COVERS	Chat · Memory · 3D Avatar · Models · Images · Voice · Telegram · Backup

Guide revision 1.4 · For Local Waifu 1.8.0 · 13 November 2026

English interface labels. Later releases may differ.

Contents

Use the clickable entries below or the PDF bookmarks sidebar to jump to any chapter.

Part I: Install, License, and Get Started 13

1. How to Use This Guide 14

Labels used in this guide . . . 14

2. Choose Your Path 15

Which edition do you have? . . . 15

Local AI or Cloud AI? . . . 15

3. System Requirements 16

macOS . . . 16

Windows . . . 16

Plan disk space before you install models . . . 16

4. Install on macOS 18

Install the Direct or Store DMG . . . 18

5. Install on Windows 20

Run the installer . . . 20

6. Understand the Two Editions 21

Direct edition . . . 21

Store/Unlocked edition . . . 21

Do not install both editions side by side . . . 21

7. Complete Onboarding 22

Set the passphrase on macOS . . . 22

Configure a local model during onboarding . . . 22

Configure Cloud AI during onboarding . . . 23

8. Use the 7-Day Trial: Direct Edition 25

Start the trial . . . 25

What happens when the trial expires? . . . 25

9. Activate and Move a Direct License 26

Activate a lifetime key . . . 26

Use the license on a new computer . . . 26

Part II: Daily Use, Memory, Characters, and Avatar 27

10. Find Your Way Around Local Waifu 29

The left sidebar . . . 29
The main area and the right panel . . . 29
The in-app Guide page . . . 29
View modes . . . 30

11. Chat, Conversations, and Images 31

Start and manage conversations . . . 31
Write and send messages . . . 31
Stop a reply . . . 31
Suggested replies . . . 32
Search and older messages . . . 32
Attach an image . . . 32
Dictate without sending automatically . . . 33

12. Edit Replies and Give Feedback 34

Edit one of her replies . . . 34
Edit your own message . . . 34
Rate or regenerate a reply . . . 35

13. Memory Book, Knowledge, and Growth 36

Memory Book . . . 36
Edit or delete a memory . . . 36
Her working notes . . . 36
Memory map . . . 36
Growth . . . 37
Adaptive Dials . . . 37

14. Characters, Persona Notes, and Backstory 40

Switch or create a character . . . 40
Rename a character . . . 40
Persona Notes . . . 40
Portraits . . . 40
Your own photo . . . 41
Backstory and the Lorebook . . . 41

15. Export and Import a Character 43

Export a character . . . 43
Import a character . . . 43
Delete a character . . . 44

16. Views, Settings, and Where Everything Lives 45

Switch views . . . 45
Open Settings . . . 45
Find a setting by tab . . . 45
Choose a background theme . . . 47
Tune the chat itself . . . 47

17. Use the Lorebook for Backstory 49

Understand the trade . . . 49
Add a backstory entry . . . 49
Switch off or delete an entry . . . 50
Make an entry always on . . . 50
Know what the Lorebook is not . . . 50

18. When She Messages You First 51

The four triggers . . . 51
What keeps them from becoming noise . . . 51
She goes quiet by herself . . . 51
Turn proactive messages off, or tone them down . . . 52
What a proactive message is not . . . 53

19. Read the Growth Dashboard 55

The seven stages . . . 55
What the dashboard shows . . . 55
Read it honestly . . . 56
Clear or change what she has stored . . . 56

20. Tune How She Has Adapted to You 57

Open it . . . 57
The four groups and ten dials . . . 57
Lock a dial . . . 57
Reset everything she has learned . . . 58
Read this honestly . . . 58

21. Give Her a 3D Avatar 59

Turn it on and pick a model . . . 59
Import your own model . . . 59
Fix her face with Expression mapping . . . 60
Backgrounds, camera, and behaviour . . . 60
Ask her for gestures in chat . . . 61
Poses and the pose editor . . . 61

Your own animations . . . 61
Photo mode and selfies from the 3D model . . . 61
Windows and macOS . . . 62

22. Put Her on Your Desktop **63**

Put her on the desktop and bring her back . . . 63
What you can do with her . . . 63
The right-click menu . . . 63
When she stays put . . . 64

23. Use Group Chats **65**

Start a group chat . . . 65
Who answers . . . 65
Manage the group . . . 65
Let them chat while you are away . . . 66

24. Play Visual Novel Mode **67**

Start a scene . . . 67
Leave and resume . . . 67
What she remembers . . . 67
macOS and Windows . . . 68

25. Set Up Routines **69**

Common controls . . . 69
Her day . . . 69
Morning brief . . . 70
Night research . . . 70
Watchers . . . 70
Reminder follow-up (macOS only) . . . 70
If something goes wrong . . . 70

26. Use Screen Awareness **71**

Turn it on . . . 71
Start a mode . . . 71
Choose what she captures . . . 72
Shortcuts . . . 72
Quiet Mode and screen awareness . . . 72

Part III: Models, Images, and Voice **72**

27. Install and Select Chat Models **75**

Understand the two required roles . . . 75

Install a recommended local model . . . 75
Add a vision model . . . 76
Install the memory model . . . 76
Remove a model . . . 76

28. Install a GGUF Model from Hugging Face **78**
Find and install a model . . . 78

29. Configure Cloud AI **79**
Add a provider . . . 79
The cloud is a fallback only if you turn it on . . . 80

30. Generate Images and Use Gallery **83**
Choose a provider and model . . . 83
Install a built-in local image model . . . 83
Use Photo Director . . . 84
Gallery . . . 84

31. Import a Custom SDXL Checkpoint or LoRA **86**
Import a custom SDXL checkpoint . . . 86
Import and activate a LoRA . . . 86

32. Use Your Own Server as a Cloud Provider **88**
Add an endpoint . . . 88
Security: what the URL is allowed to be . . . 89
Offline Mode and your own endpoint . . . 89
What your own endpoint does not get . . . 89

33. Configure Voice, Dictation, and Calls **92**
Install speech and voice models . . . 92
Choose a voice . . . 92
Add your own voice sample . . . 92
Tune how she speaks . . . 93
Configure a call . . . 93

34. Voice Messages, Hands-Free, Stories, and Songs **95**
Send a voice message . . . 95
Start a call hands-free . . . 95
Ask for a bedtime story . . . 96
Ask for a song . . . 96

Part IV: Integrations, Data, and Support **97**

35. Use Local Waifu from Telegram 99

Connect a Telegram bot . . . 99

Choose how Telegram replies . . . 100

Use photos, voice messages, and bot commands . . . 100

Stop the bot or remove its credentials . . . 101

36. Tune Performance, Hardware, and Network 104

Choose a compute device and verify acceleration . . . 104

Check runtime readiness . . . 105

Set the context size and thinking mode . . . 105

Reduce background resource use . . . 106

Configure proxy, offline mode, or a remote Ollama server . . . 107

37. Edit Her Instructions in the Developer Tab 108

Risk levels . . . 108

Edit a section . . . 108

Add standing instructions . . . 109

Watch the prompt size . . . 109

Generation parameters . . . 109

Presets, export, and import . . . 109

Restore the defaults . . . 109

Avatar test bench . . . 110

38. Understand Privacy and Data Storage 111

Know what stays local and what can leave the computer . . . 111

Locate your Local Waifu data . . . 112

Review and control permissions . . . 112

39. Back Up, Export, and Move Your Data 115

Create a full backup snapshot . . . 115

Check that a backup is usable . . . 116

Export all conversations . . . 116

Move a character to another computer . . . 117

40. Update or Uninstall Local Waifu 119

Update the Direct edition . . . 119

Update the Store/Unlocked edition . . . 120

Uninstall while keeping your data . . . 120

Remove Local Waifu and its data completely . . . 121

41. Use Logs, Diagnostics, and Reset 122

Open the app log . . . 122

Check the AI engine . . . 122
Save a diagnostics report . . . 122
Reset all application data . . . 123

42. Troubleshooting **125**

Run the basic diagnostic sequence . . . 125
The app will not open . . . 125
Installation and launch . . . 126
Chat models, downloads, and memory . . . 127
Images and local image models . . . 128
Voice, microphone, and calls . . . 128
Telegram . . . 129
License and updates . . . 129
3D avatar, desktop pet, and screen awareness . . . 130
Routines, voice features, and groups . . . 131
Backup, export, and data recovery . . . 131

43. Shortcuts, Contact, and About **133**

Keyboard shortcuts . . . 133
Check the installed version and edition clues . . . 133
Contact and official links . . . 134

44. Use localwaifu Links **135**

Turn links on or off . . . 135
The links . . . 135
Use a link on macOS . . . 135
Use a link on Windows . . . 136

45. When the Trial Ends, and Moving Your License **137**

What an expired trial blocks . . . 137
Activation problems and what each message means . . . 137
There is no device limit . . . 137
Use your license on another computer . . . 138
Before you sell or give away a computer . . . 138
Questions people actually ask . . . 138

46. Safety, Boundaries, and What She Is Not **139**

The crisis banner . . . 139
What the banner says, and why . . . 139
Adults only . . . 139
What she cannot do . . . 140

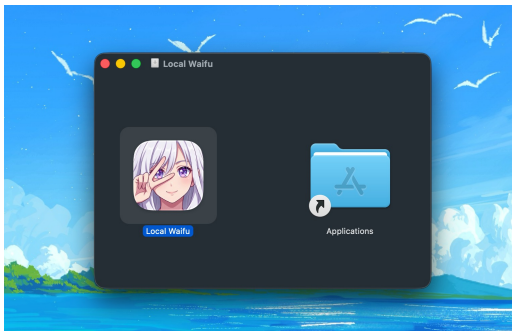
Where your data sits in all of this . . . 140
If you need help . . . 140

47. Recover a Character or the Whole App **141**

The four routes, and when each one works . . . 141
Route 1: character export and import . . . 141
Route 2: full backup . . . 141
Route 3: the soul recovery code . . . 142
When the app cannot open: the recovery screen . . . 142
What nothing can recover . . . 143
The order to do things if you are unsure . . . 144

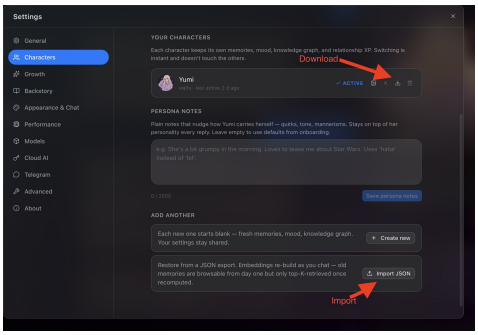
Find your way visually

Six screens cover the actions most people need first. The full-size versions appear in the matching chapters.



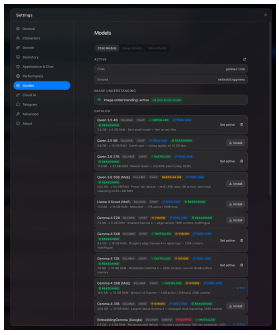
Install

Start with the correct package for macOS or Windows.



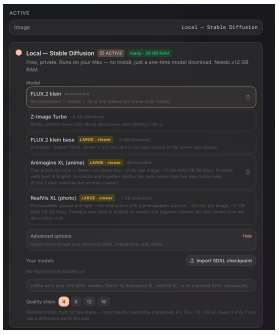
Move a character

Export or import a companion without replacing an existing one.



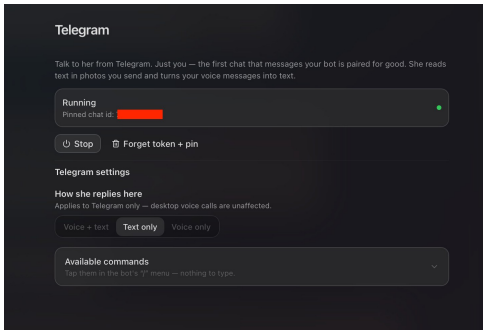
Choose her mind

Pick a chat model and keep the required memory model installed.



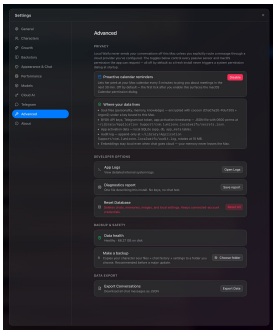
Create images

Use a built-in model or import a compatible SDXL checkpoint and LoRA.



Connect Telegram

Pair one private chat and keep the bot token secret.



Protect your data

Create a full backup before updates or migration.

Your first fifteen minutes

A working first conversation matters more than perfect settings. Follow this path, then return to the advanced chapters when you need them.

01	02	03	04
Install	Create her	Choose AI	Say hello

Choose how she thinks

LOCAL AI	CLOUD AI
Private and offline-capable after model downloads. Best when your computer has enough memory.	Fast to start on smaller computers. Requires your provider account, internet connection, and consent to send the selected context.

You are ready when...

- your first character exists;
- one chat model is active;
- EmbeddingGemma is installed for semantic memory;
- a short test message receives a reply.

Voice, images, Telegram, and custom models can wait until chat works.

What to do next

1 · Make her yours	2 · Add expression	3 · Protect the result
Refine persona notes, backstory, appearance, and the active chat model.	Install voice and image models only after everyday chat is stable.	Export important characters and create a full backup before major changes.

Before you begin

This guide describes Local Waifu 1.8.0 for supported macOS and Windows computers. It deliberately distinguishes the Direct edition from the Store/Unlocked edition. A few screenshots were taken in version 1.7.3, so a screen can look slightly different from the picture even where the text is current. It does not fabricate screens that are absent from the supplied product captures; platform-specific steps are documented as procedures instead.

Security and privacy

Never share license keys, provider API keys, Telegram bot tokens, private conversation exports, or unreviewed diagnostic logs. Screenshots in this guide contain no usable credentials; sensitive token fields are redacted.

Backup limitation in version 1.8.0

The application can create a full backup snapshot, and Settings → Advanced → Restore a backup can check that a backup is usable by copying it into a new, empty folder and verifying the soul files, without touching your live data. Version 1.8.0 has no button that replaces your live data with a backup: bringing a backup back is a support-assisted procedure. Preserve backups unchanged and contact support for a recovery procedure. Character JSON import/export is the supported self-service transfer method for one companion.

Part I: Install, License, and Get Started

1. How to Use This Guide

Local Waifu is available for macOS and Windows in two distribution editions. Most features work the same in both editions, but installation, licensing, and updates differ.

Labels used in this guide

- ALL: all supported platforms and editions.
- macOS: Apple Silicon Mac only.
- Windows: Windows x64 only.
- DIRECT: the direct-download edition with a 7-day trial and license key.
- STORE: the Store/Unlocked edition distributed through a supported storefront; no in-app trial or key.
- LOCAL: processing stays on your computer.
- CLOUD: selected data is sent to the provider you configured.
- DOWNLOAD: additional files or model weights are required.
- RESTART: restart Local Waifu after changing the setting.
- WARNING: read carefully before continuing.

UI paths use this format: Settings → Models → Chat Models. Button and field labels appear in bold.

This guide describes Local Waifu 1.8.0. A few screenshots were taken in version 1.7.3 on the English macOS interface, so a screen can look slightly different from the picture even where the text is current. Windows uses the same main concepts, but some hardware, update, and installer screens differ. Where macOS and Windows behave differently, the chapter says so by platform. The downloadable PDF is the 1.8.0 edition of this same guide.

2. Choose Your Path

Which edition do you have?

If you obtained Local Waifu from...	Edition	What happens
The Local Waifu website or official Direct release	DIRECT	Full 7-day trial, then a lifetime license key is required
A Store/Unlocked package such as the dedicated itch.io build	STORE	Full access without entering a key inside the app
A storefront that resells the app, such as BOOTH	Depends on the package that store serves	Check which build you received with the test below

A quick way to check is Settings → License:

- The tab is present in the Direct edition.
- The tab is intentionally absent in the Store/Unlocked edition.

Do not install a Direct update over a Store/Unlocked copy merely to obtain a newer version. Keep each edition on its own update channel. If a storefront resells the app, follow that storefront's download and update instructions, and confirm the edition with the Settings → License test above before changing anything.

Local AI or Cloud AI?

Choice	Best for	Requirements	Privacy note
Local AI	Private offline-capable chat and no per-message provider bill	Model download, enough RAM and disk space	Chat inference stays on your computer
Cloud AI	Smaller computers or access to a provider's hosted models	Internet and your own provider account/key	Prompt and relevant context go to that provider

Memories, character data, and the Local Waifu database remain local even when you intentionally choose a cloud model. Cloud image, speech, Telegram, and remote Ollama routes transmit the data needed by those services.

Cloud AI is also never a silent backup. If the local engine is down, or she receives a photo that no local model can see, the message does not go to a saved cloud key unless you turned on Use the cloud when the local model can't answer in Settings → Performance → Network. That switch is off by default. Chapter 29 explains it.

3. System Requirements

macOS

- Apple Silicon Mac: M1 or newer.
- macOS 14 Sonoma or newer.
- 8 GB unified memory for the lightest local models. On a Mac with 8 GB the app recommends Gemma 4 E2B (a 4.6 GB download) instead of the larger Gemma 4 E4B, which would leave an 8 GB Mac swapping on every reply.
- 16 GB or more recommended.
- Internet for installation, model downloads, updates, Telegram, or cloud providers.
- Several gigabytes of free disk space for local models.

Local AI uses Apple Metal acceleration. There is no Intel Mac build of Local Waifu 1.8.0. macOS 14 is a hard requirement: the bundled AI engines already needed it, so on macOS 13 the app installed but local chat and pictures could not start. The installer now says so up front.

Windows

- Windows 10 22H2 or Windows 11.
- 64-bit x64 computer.
- 8 GB RAM minimum; 16 GB or more recommended.
- A GPU is optional. Local Waifu can run on CPU, but it will be slower.
- Current GPU drivers are recommended for NVIDIA or AMD acceleration. Intel graphics, virtual machines, and unsupported adapters fall back to CPU; Windows has no accelerated runtime path for them.
- Microsoft Edge WebView2 Runtime; the installer can fetch it if required.

Windows may configure an appropriate GPU runtime after first launch. The app can continue on CPU while setup completes; restart when the app reports that acceleration is ready.

Plan disk space before you install models

The application is much smaller than the AI models it can use. Keep at least 8 to 10 GB free for a comfortable lightweight local setup. Larger chat models can require roughly 17 to 24 GB each, and local image generation adds several more gigabytes.

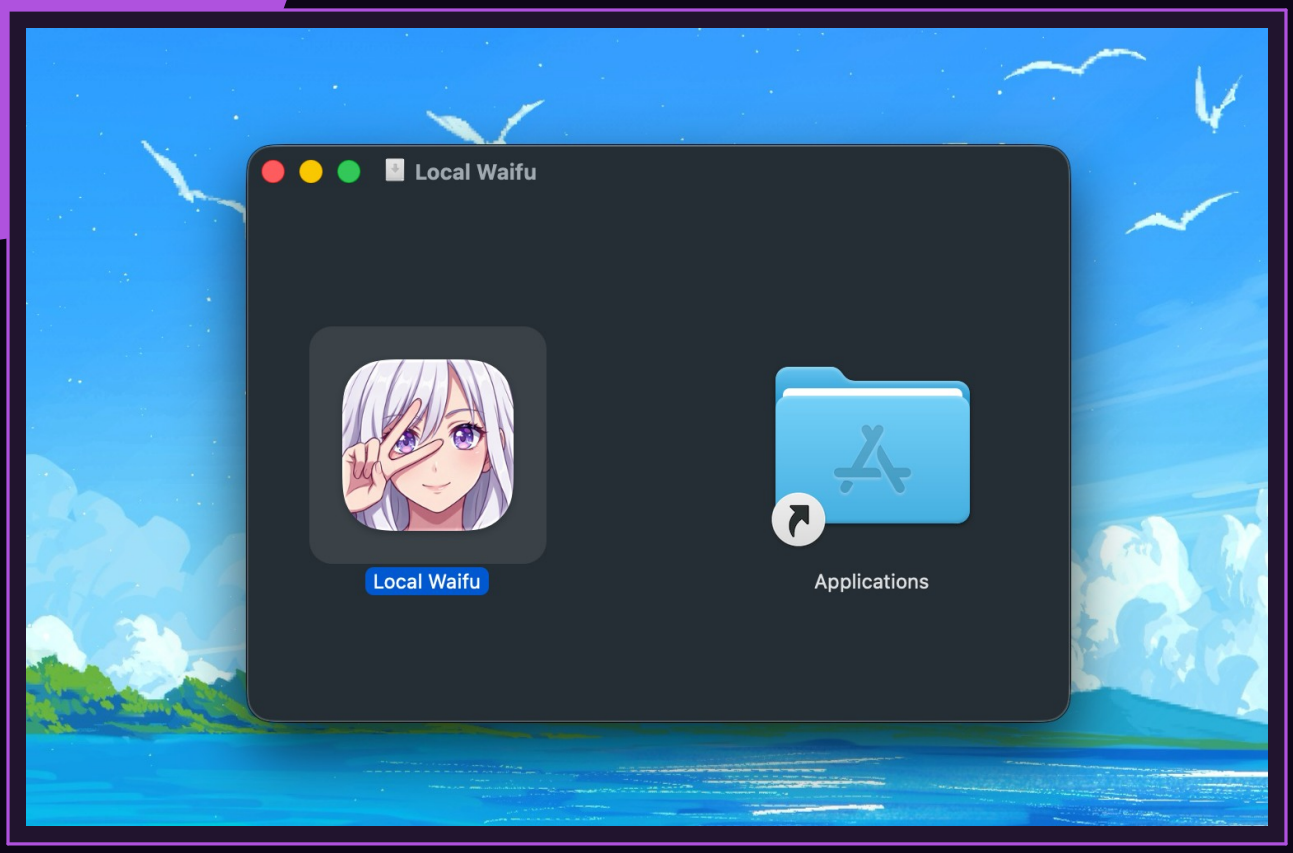
Typical additional downloads include:

- a chat model: approximately 3 to 24 GB depending on size and quantization;
- EmbeddingGemma for semantic memory: approximately 0.6 GB;
- local image weights: approximately 4 to 5 GB or more;
- speech recognition and voice models: hundreds of megabytes to several gigabytes.

These figures are estimates. The model screen shows the current download size and RAM guidance.

FIGURE 01

macOS installation: open the DMG, then drag Local Waifu into Applications.



macOS disk image shown. Drag the app icon onto Applications.

The exact control is easier to find once you know where to look.

4. Install on macOS

Install the Direct or Store DMG

Goal

Install Local Waifu in the Applications folder and start the correct edition.

Before you start

- Confirm that your Mac uses Apple Silicon and runs macOS 14 Sonoma or newer.
- Download the DMG only from the official Direct page/release or your Store library.
- If updating an existing installation, create a backup first.

Steps

1. Open the downloaded `.dmg` file.
2. Drag Local Waifu into Applications.
3. Wait for the copy to finish.
4. Eject the mounted disk image.
5. Open Applications and start Local Waifu.
6. Allow macOS to complete its first security verification.

Expected result

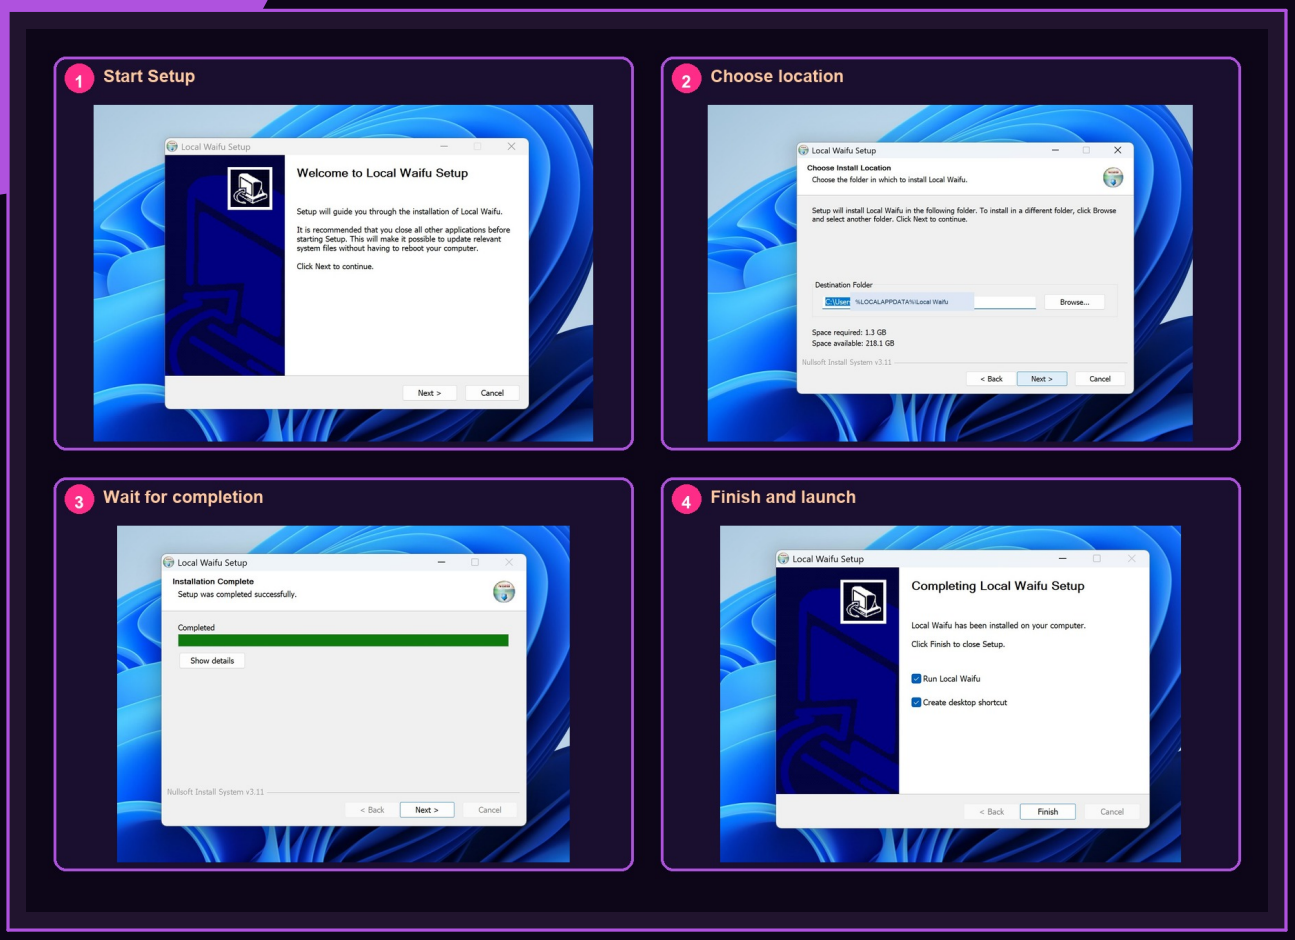
Local Waifu opens to the Welcome/onboarding flow. A Direct build can show license options; a Store/Unlocked build does not ask for a key.

If something goes wrong

- Download the DMG again from the official source.
- Confirm that you did not download a Windows installer or an Intel-only package.
- In Finder, Control-click Local Waifu, choose Open, and approve the system prompt if macOS offers that option.
- Do not routinely disable Gatekeeper or remove quarantine attributes. If a notarized official build is still blocked, contact support with the exact macOS message.

FIGURE 02

Windows installation in four steps: start Setup, choose the location, wait for completion, then launch Local Waifu.



Windows 11 installer shown. Wording may vary slightly in later packages.

The exact control is easier to find once you know where to look.

5. Install on Windows

Run the installer

Goal

Install Local Waifu and start the correct Windows edition.

Before you start

- Use Windows 10 22H2 or Windows 11 x64.
- Download the installer only from the official Direct release or your Store library.
- Close any earlier Local Waifu process before updating.
- Make a backup before replacing an existing installation.

Steps

1. Run `Local.Waifu_1.8.0_x64-setup.exe`, the versionless alias `local-waifu-x64-setup.exe`, or the corresponding Store/Unlocked installer.
2. If Windows SmartScreen appears, confirm that the file came from the official source.
3. Select More info, then Run anyway only when you trust the downloaded file.
4. Follow the installer and keep the suggested location unless you have a reason to change it.
5. Let the installer add WebView2 if Windows does not already have it.
6. Start Local Waifu.
7. Allow the first startup and antivirus scan to finish. The large installer and bundled AI processes may take longer to scan.

Expected result

Local Waifu opens to onboarding. GPU detection may continue in the background.

If something goes wrong

- Restart Windows and retry if an old Local Waifu or local-engine process is still running.
- Repair or install Microsoft Edge WebView2 Runtime if the window is blank.
- Check antivirus quarantine for the official installer and bundled AI processes.
- Verify free disk space.
- Remember that moving the program installation folder does not move conversations or models; user data is stored separately.

6. Understand the Two Editions

Direct edition

The Direct edition includes:

- a full 7-day trial;
- optional license activation during onboarding;
- Settings → License;
- a lifetime purchase link;
- a paywall after the trial expires;
- the Direct update channel.

The trial starts after onboarding is completed, not when the installer is downloaded or first opened.

Store/Unlocked edition

The Store/Unlocked edition includes:

- full access after onboarding;
- no trial countdown;
- no in-app license key;
- no License tab;
- no Direct paywall;
- no Direct updater.

Chat, memory, image, voice, and Telegram features remain available. Updates must come from the Store channel that supplied the build.

> Local Waifu 1.8.0 does not contain a separate Steamworks integration. A future Steam package should be treated as a Store/Unlocked distribution unless its documentation states otherwise.

Do not install both editions side by side

Direct and Store builds use the same application identity and data location. They are not designed to be independent side-by-side installations. Before switching editions:

1. export important characters;
2. create a full backup;
3. quit Local Waifu;
4. install only the edition you intend to keep;
5. verify your data before deleting the backup.

7. Complete Onboarding

Onboarding introduces Local Waifu, learns your preferences, creates the first character, configures voice and models, and opens the first meeting.

Typical stages include:

1. Welcome: language, the 18+ confirmation, links to Terms and Privacy, and the license option. You must tick I confirm that I am at least 18 and understand that this is an AI companion, not a person before Let's begin becomes available.
2. About You: how the companion should address you.
3. Understanding You: preferences and interaction style.
4. How You Connect: communication preferences.
5. Your Love Language: desired expressions of care.
6. Choosing: the type of dynamic you want.
7. Your Dynamic: relationship tone and boundaries.
8. Building Her: name, appearance, personality, and portrait details.
9. Her Voice: voice preference.
10. Almost There: you choose How often should she reach out first? (Rarely, Once a day, When something matters, or Often). On macOS this screen also asks for a passphrase: see the next section. Windows asks only the reach-out question.
11. Model Setup: local model download or Cloud AI setup.
12. First Meeting: character creation is finalized and chat opens.

Set the passphrase on macOS

On macOS the Almost There step asks you to set a passphrase of at least 8 characters, typed twice. The app says it encrypts her personality file on disk with this passphrase, that your conversations are not encrypted by it, and that if you lose the passphrase that file cannot be recovered. Store the passphrase somewhere safe, such as a password manager. The app does not ask for it again in daily use. Chapter 47 explains what recovery routes exist.

Configure a local model during onboarding

Goal

Install a local chat model and the semantic-memory model.

Before you start

- Connect to the internet for the model download.
- Keep Local Waifu open.
- Check available disk space.
- For a large download, prevent the computer from sleeping.

Steps

1. Choose the Local AI path on Model Setup.
2. Review the recommended chat model for your detected memory tier. On a Mac with 8 GB the recommendation is Gemma 4 E2B. Select Choose a different model to see every local model the app supports, smallest first, with its download size and whether it can see photos. Models your memory cannot run are greyed out.
3. Select Install and wait for the download to complete.
4. Install EmbeddingGemma when it is offered. It powers semantic memory retrieval and is separate from the chat model.
5. Continue after the required model is installed, or use Skip for now if the screen offers it and you intend to configure Cloud AI later.

Expected result

The chat model is ready for conversation, and the embedding model is ready to create and retrieve semantic memories.

If something goes wrong

- Check internet access, VPN, proxy, firewall, and free disk space.
- Restart Local Waifu and continue in Settings → Models.
- On Windows, allow GPU runtime setup to finish, then restart if prompted.
- If your computer cannot comfortably run the recommendation, install a smaller model in Settings → Models.

Configure Cloud AI during onboarding

Goal

Use your own provider account instead of downloading a large local chat model.

Before you start

- Have an API key or supported connected account.
- Review the provider's pricing and privacy terms.
- Understand that selected prompts and context leave your computer when that provider is used.

Steps

1. Choose the Cloud API tab (the local tab is called On-Device).
2. Select the provider.
3. Paste the provider key into the secure field or complete the available account connection.
4. Run the connection test.
5. Select an available chat model.
6. Complete onboarding.

Expected result

Chat uses the selected provider while character data, the application database, and memory storage remain local.

If something goes wrong

- Re-copy the key without spaces.
- Confirm the provider account has access and billing/credits where required.
- Turn off Offline Mode when intentionally using Cloud AI.
- Remove or replace a saved bad key in Settings → Cloud AI.

8. Use the 7-Day Trial: Direct Edition

The Direct trial provides full app access for seven days. The remaining whole days are rounded up and shown in Settings → License. The app can also show reminders when three, two, or one day remains.

Start the trial

The normal flow starts the trial automatically when you finish onboarding. If Settings → License shows Not started, select Start free 7-day trial.

During the trial you can use chat, memory, images, dictation, voice, and other available features. Model downloads and provider costs are separate from the app license.

What happens when the trial expires?

- The main app is replaced by the license screen/paywall.
- Chat, image generation, voice, and other licensed operations are rejected until activation.
- Conversations, character data, memories, and soul files stay on the computer.
- Settings remain available so you can enter a key, export, back up, or collect diagnostics.
- Resetting or reinstalling does not guarantee another trial.

9. Activate and Move a Direct License

Activate a lifetime key

Goal

Activate the Direct edition on the current computer.

Before you start

- Copy the exact key from the purchase email.
- Connect to the internet for activation.
- There is no device limit. A license is for your own devices, within fair use.

Steps

1. Use any one of these locations:

- expand I have a license key on the Welcome screen;
- open Settings → License → Redeem a key;
- use the key field on the expired-trial screen.

2. Paste the key.
3. Select Redeem.
4. Wait for Lifetime · active.

Expected result

The full app unlocks immediately. Settings → License shows an active lifetime license and the date you have owned the app since.

If something goes wrong

- That key wasn't recognised: copy the complete key again from the purchase email, with no extra spaces.
- Couldn't reach the license server: check your internet connection and try again. Your key is fine.
- An invalid attempt does not destroy the current trial state.
- If the app says the license could not be activated on this device, or that it was revoked or disabled, write to contact@localwaifu.com and support will sort it out.

Use the license on a new computer

Goal

Activate Local Waifu on a replacement or second computer.

Before you start

- Export important characters and make a full backup of the old computer (Chapter 39).
- Have the license key from the purchase email.

Steps

1. Install the Direct edition on the new computer.

2. Complete the basic onboarding.
3. Open Settings → License.
4. Paste the same key and select Redeem.
5. Import exported characters and reinstall required models.

Expected result

The new computer shows Lifetime · active. The old one is unaffected. Settings → License says the same thing: New computer? Install Local Waifu there and redeem the same key, there's no device limit.

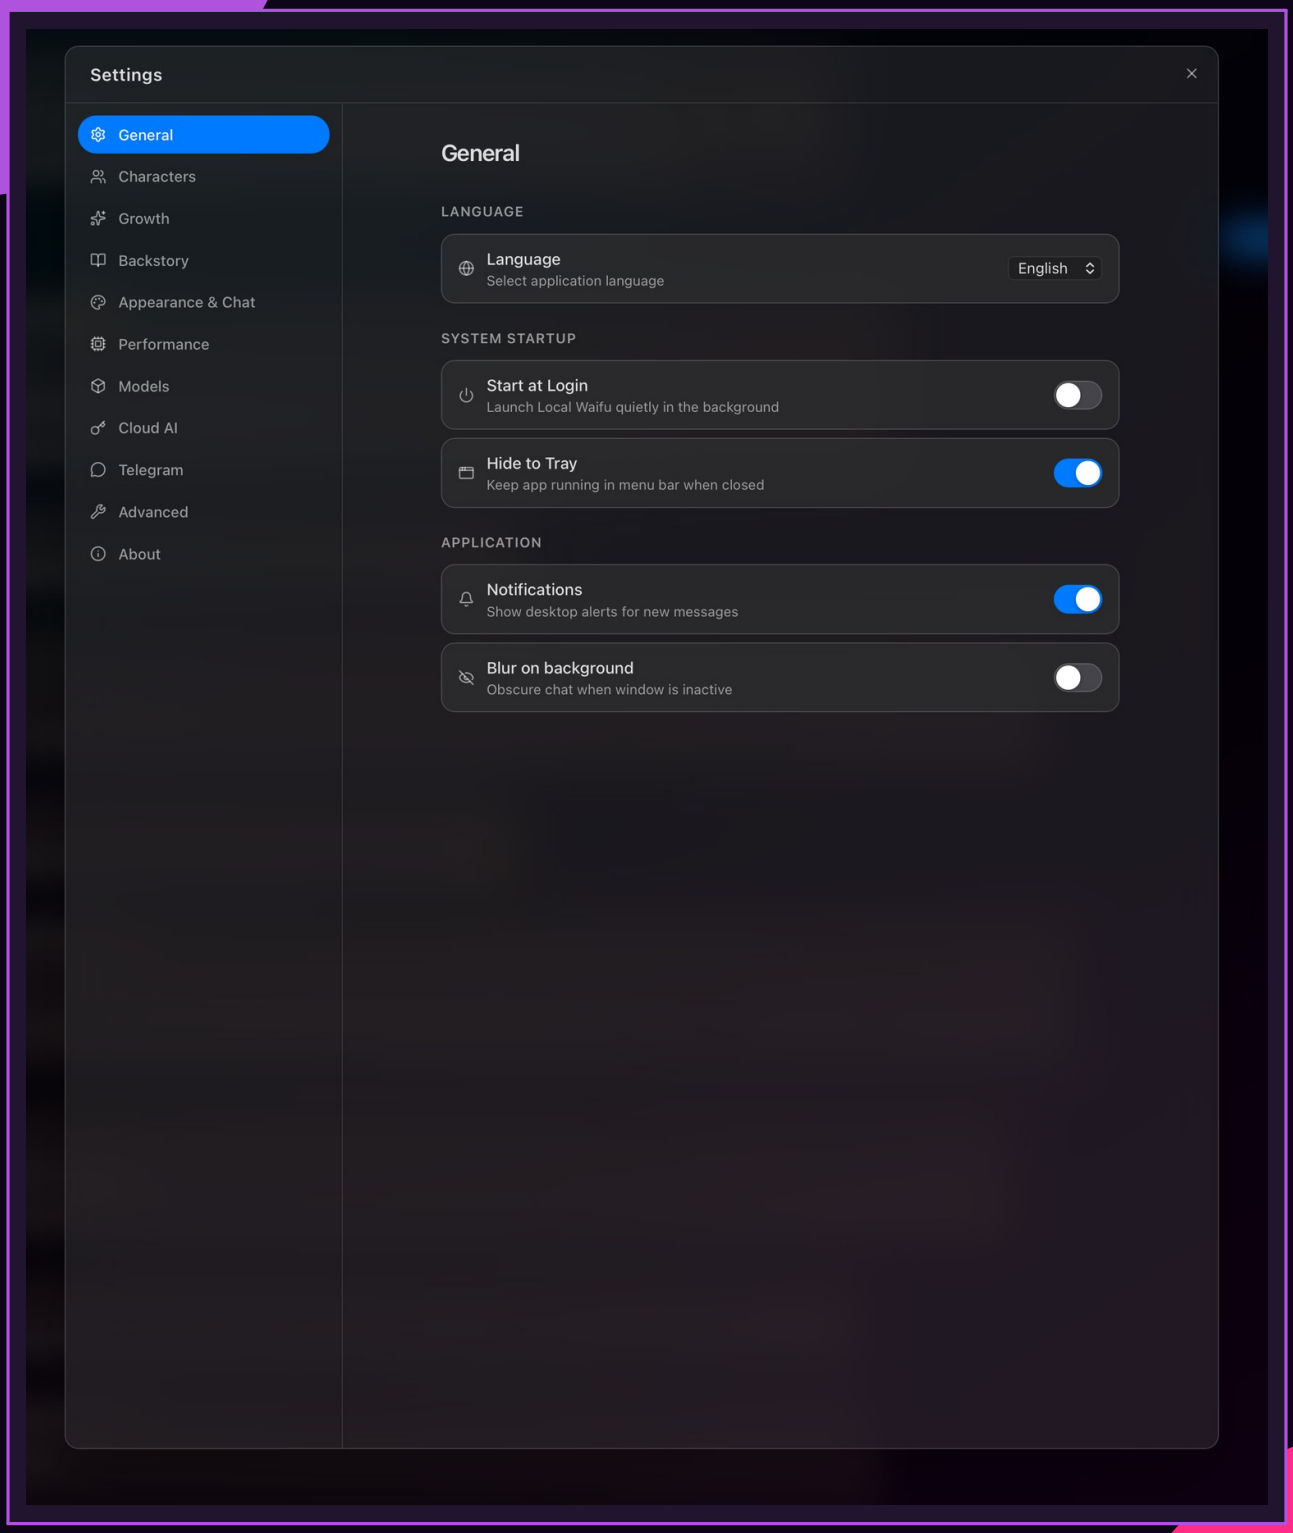
If something goes wrong

- Older versions of the app had a Deactivate this device button. It was removed in 1.7.4 because there is nothing to free up.
- If the key is rejected on the new computer, write to support rather than repeatedly resetting data.
- Do not publish the key in screenshots or diagnostics.
- Before you sell or give away a computer, follow the steps in Chapter 45.

Part II: Daily Use, Memory, Characters, and Avatar

FIGURE 03

General settings: language, appearance, startup, notifications, and update controls.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

10. Find Your Way Around Local Waifu

The main window is organized around the active character and conversation.

The left sidebar

The sidebar lists the places you can go:

Item	What it opens
Chat	The conversation with the active character
Guide	A short in-app guide with common tasks (see below)
Memory Book	Remembered facts, moments, preferences, her working notes, and how she has adapted
Memory map	A picture of the people, places, and things she has learned about
Gallery	Images sent or generated by the companion, and your 3D photos
Visual novel	Story scenes you play with her

Below those items:

- Chats is the list of conversations for the active character. Use New chat to start one, the search icon (`Command/Ctrl+F`) to search history, and the rename and delete actions on each row.
- Groups lists group chats, with a button to start a new one.
- At the bottom: Settings, a button for the Mini window, and a button that puts her on the desktop.

When she has gone quiet by herself, a small Quiet badge appears in the sidebar. Hover over it to read that Quiet Mode is on and she stays silent until you are back.

The main area and the right panel

1. Chat stream: your messages, her replies, images, and message actions.
2. Message composer: attachment, emoji, call, screen menu, bedtime story and song menu, dictation, voice message, text field, and send control.
3. Right panel: the active character, relationship level, current mood, and a Knowledge Graph card. In a group chat the right panel shows the group panel instead.
4. 3D avatar or portrait: her figure in the sidebar slot.

The in-app Guide page

Select Guide in the sidebar for a one-page quick guide. It covers talking to her, her brain (models), giving her a body, speaking with her, what she remembers, privacy, keyboard shortcuts, and short answers to common tasks: export a character, back up all app data, upload a 3D avatar, and logs and support. It links to the full manual on the website. This web guide is the full manual.

View modes

- Full view: complete interface.
- Mini view: compact chat.
- Menu bar/tray view: lightweight access.
- Visual novel and desktop pet are two more window modes, covered in Chapter 24 and Chapter 22.

On macOS the view shortcuts are `Command+Option+1` to `Command+Option+5`. On Windows they are `Ctrl+Shift+1` to `Ctrl+Shift+5`. Chapter 43 has the full table, and `Command/Ctrl+,` opens Settings.

11. Chat, Conversations, and Images

Start and manage conversations

Goal

Create a clean conversation without losing older chats.

Steps

1. Select New chat in the Chats list of the left sidebar.
2. Type a message and press Enter or select Send.
3. To rename a conversation, use its rename action, type the title, then press Enter.
4. Press Escape to cancel a rename.
5. To remove a conversation, choose its delete action and confirm Delete this chat?.

Expected result

Each conversation remains available for the active character until deleted. Removing the final conversation creates a new empty one.

If something goes wrong

Local Waifu avoids creating repeated empty chats. Send a message in the current empty conversation before selecting New chat again.

Write and send messages

By default:

- Enter sends;
- Shift+Enter inserts a new line;
- Settings → Appearance & Chat → Send on Enter changes this behavior;
- the send action is disabled for empty text and while a send is already in progress. If she is still answering, the app says She's still answering, press Send once she's done.

Use the smiley button next to the photo button to open an emoji picker. A click puts the emoji where your cursor is, and the picker stays open so you can add several. Escape or a click outside closes it.

Typing in Japanese, Chinese, or Korean is safe: pressing Enter to confirm a word in the input method does not send the message half-written.

Stop a reply

Select Stop (Stop generating) to end a reply. In 1.8.0 Stop really stops: the local engine stops generating at once, so your next message does not wait behind the rest of her answer. The part of the reply she had already written stays as a message marked Reply cancelled. Press Try again. Select Try again to write it again.

Suggested replies

If you turn on Suggest replies in Settings → Performance, three tappable replies appear under her newest message. It is off by default because each reply costs a little output on every turn.

Search and older messages

Use `Command/Ctrl+F` to search chat history. A result from another chat, or one older than the most recent messages, opens the right chat, loads enough history, and scrolls to the message. If a message is too far back to show, the app says so. Show older messages loads more of a long chat.

Attach an image

Goal

Send one image for a vision-capable model to inspect.

Before you start

- Use a normal image: PNG, JPG, WebP, or GIF.
- Configure a local vision model or a cloud model that supports images. If none is set up, a card offers Install vision model or Add a cloud key.
- Only one pending image can be attached at a time. A newer image replaces the pending one.

Steps

1. Add the image in any of these ways:

- select the attach button (Attach image);
- paste it with `Command+V` (macOS) or `Ctrl+V` (Windows), for example a screenshot, or an image copied from a browser or a file manager;
- drag an image file anywhere onto the window. The app says Drop the image to attach it.

2. Review the preview.
3. Add a question or instruction, or press Enter with no text to let her react to what she sees.
4. Send the message.

Expected result

The image and text are sent to the configured vision route and appear in the conversation.

If something goes wrong

- That file isn't an image: choose an image rather than a PDF or document.
- Couldn't read that image: export it again as PNG or JPEG.
- Large images are resized automatically. An image over 5 MB, or in a format the models cannot read such as HEIC, is shrunk to about 2048 pixels and re-encoded. Image is too large appears only if it is still over 5 MB after that.
- A paste of real text that happens to carry a picture too, such as a spreadsheet cell, pastes as text.
- Your photos are not stored. A message sent with a photo alone shows as Photo after a restart.
- If vision is unavailable, install a vision-capable model or configure a compatible cloud provider.

Dictate without sending automatically

Dictation records speech, transcribes it, and inserts the text into the message field. It does not send the message automatically. Review and edit the transcription, then send it yourself.

A dictation session can run for up to ten minutes. If the microphone disconnects or the time limit is reached, keep the text already captured and start another session.

For a recorded message that is sent as audio, see Chapter 34.

12. Edit Replies and Give Feedback

Local Waifu 1.8.0 lets you correct her saved replies and also edit your own messages. The two work differently, so read both sections.

Edit one of her replies

Goal

Correct the wording of a companion reply while preserving its place in the conversation.

Before you start

Only saved companion messages can be edited this way. A reply that is still streaming, or an error bubble that was never saved, cannot be edited.

Steps

1. Hover over the companion reply.
2. Select the pencil (Edit message).
3. Change the text.
4. Press Enter to save.
5. Use Shift+Enter for a new line or Escape to cancel.
6. Use Undo from the confirmation toast if you need the previous wording back.

Expected result

The corrected text remains in the same position and shows edited beside its timestamp after reload. From the next message on she reads the corrected version as her own past words.

Important memory behavior

Editing her reply does not rewrite memories already extracted from the old wording. If the app reports that saved memories still exist, select Memory Book in the notice, or open the Memory Book, and edit or delete those memories separately.

Edit your own message

Goal

Fix a message you sent and make her answer the corrected version.

Steps

1. Hover over one of your own messages.
2. Select the pencil.
3. Change the text. The editor tells you how many messages will go: Saving removes the messages below (N) and what she remembered from them, then she answers again. If there are none, it says Saving makes her answer again.
4. Save.

Expected result

The conversation rewinds to that point, like editing a prompt in a chat assistant. Everything after your message disappears together with the memories she saved from those replies, and she answers the corrected message. The notice reports Later messages removed and the number of memories removed.

If something goes wrong

- A memory that was merged from several conversations is kept. Only the link to the removed reply is dropped.
- A message cannot be empty.
- If the save fails, the message is not changed.
- You cannot take back a rewind. If you are not sure, copy the later messages first.

Rate or regenerate a reply

Hover over a reply to use the available actions:

- positive feedback (Good response);
- negative feedback (Bad response);
- Try again for the latest reply;
- private-thought display (the thought badge) when the model returned supported reasoning content;
- image enlargement for attached or generated images;
- Delete this message on a message that carries a picture or an audio recording;
- in a group chat, Read aloud on a character's reply.

Feedback helps you mark response quality. Regeneration creates another answer but does not automatically delete a memory already created from the original reply. Try again is blocked while a reply is still coming in.

13. Memory Book, Knowledge, and Growth

Memory Book

Memory Book shows information extracted from conversations. It is organized from top to bottom as:

1. How she's adapted to you (Chapter 20);
2. Her working notes, described below;
3. What she's noticed about you, when there is something to show;
4. your saved memories, counted as kept moments.

Common memory types are:

- fact: something stable about you or the world;
- moment: a meaningful event;
- preference: what you like or dislike;
- inside joke: shared language or recurring humor.

Other entries include What she knows about you, What she guessed about you, Session summary, Milestone, and Inferred milestone. Memories with high importance are marked key.

Edit or delete a memory

1. Open Memory Book.
2. Select a memory.
3. Edit its text and save, or choose delete and confirm Delete this memory? This can't be undone.
4. Use Where she learned this to inspect the source messages when provenance is available.

If no source messages appear, the memory may predate source tracking or may have been inferred from a longer pattern.

> A memory save can fail when the memory model is offline. Check the required embedding model in Settings → Models. An empty memory is not saved, and the old text stays.

Her working notes

Her working notes are eight short notes she keeps about you and about herself. Every one of them is part of every reply, so a wrong note changes how she answers. Open one, correct it, and select Save.

Memory map

Select Memory map in the left sidebar to see the people, places, and things she has learned about as a map of linked circles. A bigger circle means it has been mentioned more. Warm lines mark a warm bond and cool lines a cool or tense one. The color of a circle shows its kind (person, place, thing, and others).

- Drag to move the map and scroll to zoom. Search entities with the search box, and filter by Kinds.
- Select a circle to see Connections, Memories mentioning this, Attributes, Confidence, and the number of mentions.

- Rename a circle. If the new name matches an existing one, the two are merged.
- Delete a circle, or delete a single link with Delete this link. When you delete a circle you can also delete the memories that mention it. Deleted items will not be learned again.
- If there are many entities, the map shows the strongest ones and says how many of the total it shows. The right panel's Knowledge Graph card has an eye button (View knowledge graph) that opens a list of what she knows, with an Open full map button that leads to the same map.

If nothing is mapped yet, the map says Nothing mapped yet. and fills as you talk.

The other places where information lives:

- Memory Book contains readable saved memories.
- Memory map contains entities and relations.
- Backstory (the Lorebook) contains information you deliberately enter (Chapter 17).
- Session summaries compress older conversation context.

Growth

Settings → Growth is a progress view, not a manual XP editor. It carries the relationship stage and XP, days known and message count, what she has stored, how your own writing has changed, and the story timeline.

Chapter 19 explains every part of it, including what the stage number does and does not measure.

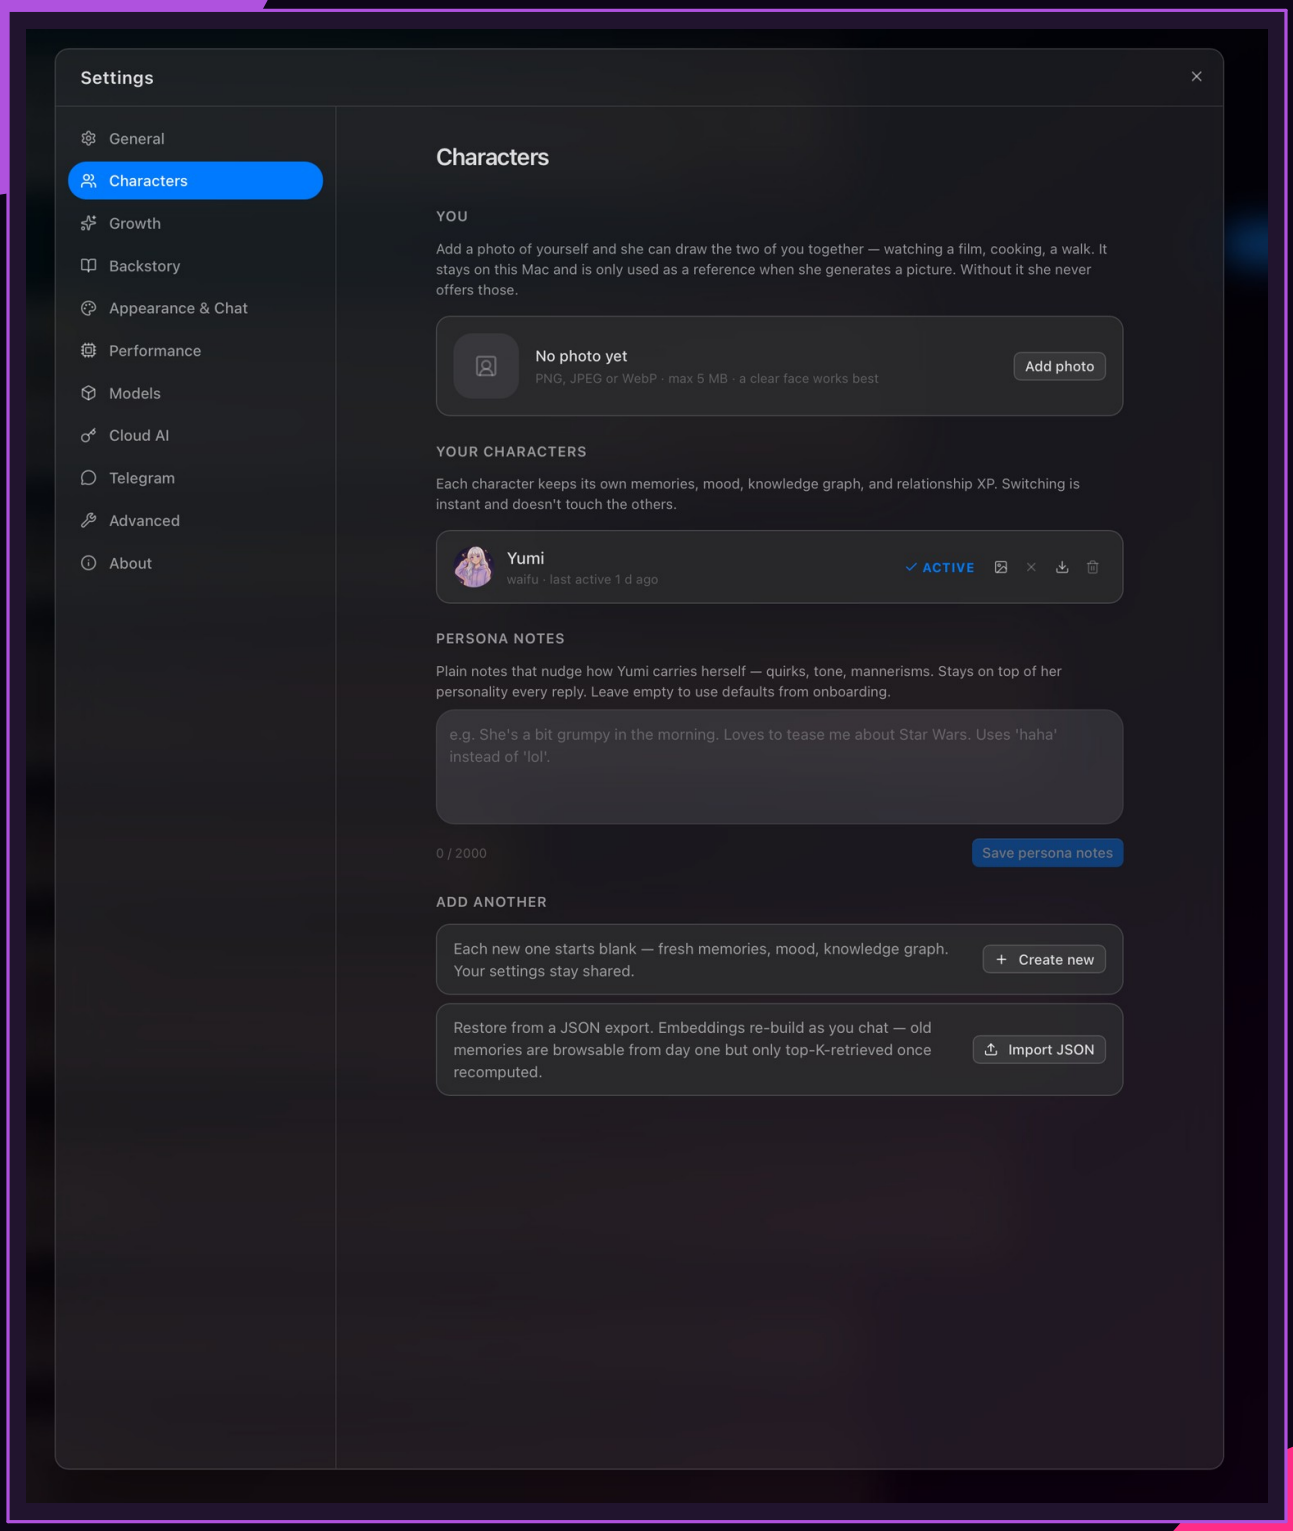
Adaptive Dials

Adaptive Dials distinguish the personality baseline from a learned adjustment and effective behavior. This explains why the companion can gradually adapt even when you do not manually change the original personality values.

Resetting adaptation does not delete the character, conversations, or Memory Book.

FIGURE 04

Characters: switch companions, create a new one, or open import/export controls.

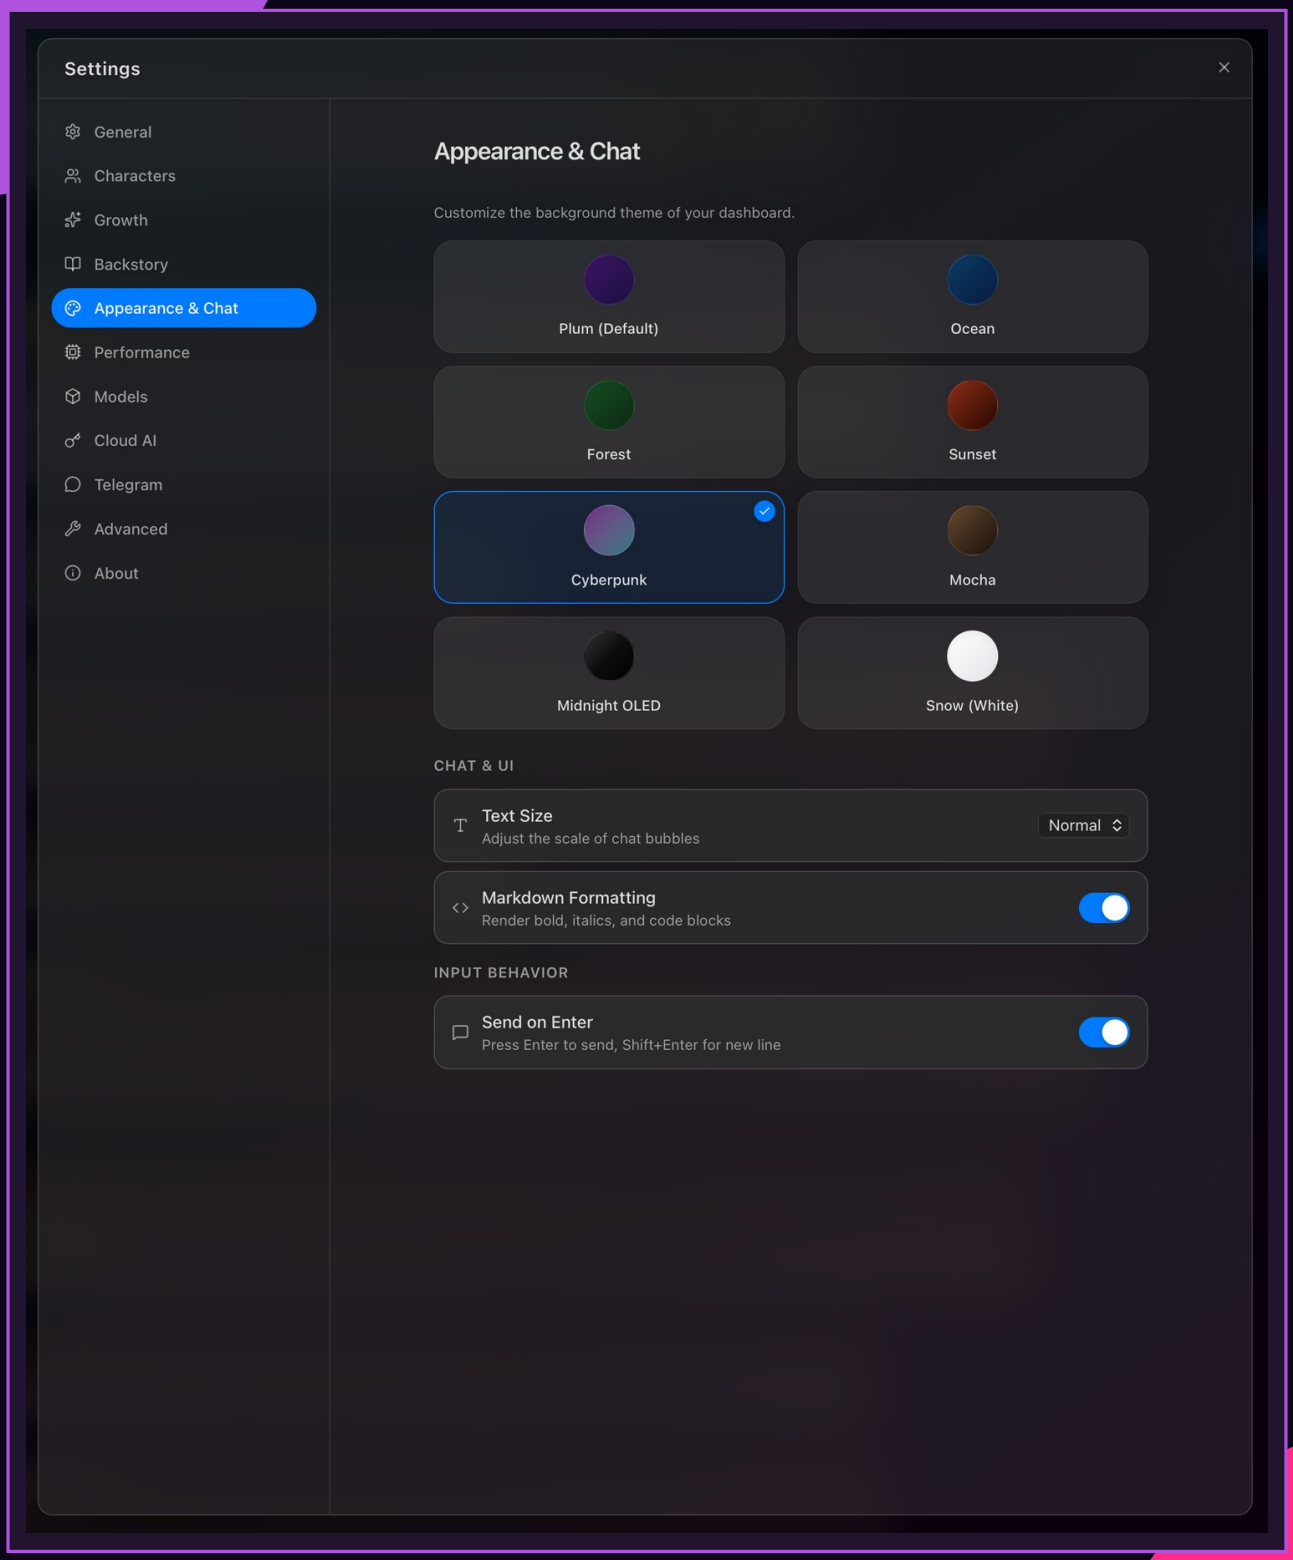


Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

FIGURE 05

Appearance & Chat controls for display density, behavior, and chat interaction.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

14. Characters, Persona Notes, and Backstory

Each character has independent conversations, mood, relationship state, memory, knowledge, portrait, and character-specific voice choice.

Switch or create a character

Goal

Use several independent companions without mixing their memories.

Steps

1. Open Settings → Characters.
2. Select Switch on another character to make it active, or select Create new (under Add Another).
3. If you are creating one, complete the character flow.
4. Return to chat.

Expected result

Chat loads the selected character's conversation and state. Each new character starts blank, with fresh memories, mood, and knowledge, while your settings stay shared. Switching does not merge the other characters' memories.

Rename a character

On the character's row, select Rename and type Her name. If her personality file cannot be rewritten in the same session, the app tells you she may still call herself the old name until you restart.

Persona Notes

Persona Notes provide plain instructions about how the active character carries herself: tone, habits, mannerisms, and other guidance. Save an empty value if you intentionally want no override.

Use concise behavioral notes rather than a full story. Put world facts, places, and recurring lore in Backstory/Lorebook.

Portraits

- Use a normal image recognized by the operating system.
- Maximum file size is 5 MB.
- The portrait editor crops to a 3:4 frame before saving.
- Add photo, Change photo, and Remove photo (use initials) are on the character's row. Removing the portrait returns the UI to initials.
- If the 3D avatar is on, the portrait is used wherever the 3D model is off or still loading, and can serve as the reference for her selfies (Chapter 21).

Your own photo

On the same tab, a You block lets you add a photo of yourself (PNG, JPEG, or WebP, up to 5 MB, a clear face works best). With it, she can draw the two of you together, for example watching a film, cooking, or on a walk. The photo is only used as a reference when she makes that picture: with local images it stays on your computer, and with a cloud image model that accepts a reference (Google, OpenAI) it goes to that provider. An SDXL model cannot take a reference picture. You can remove the photo at any time.

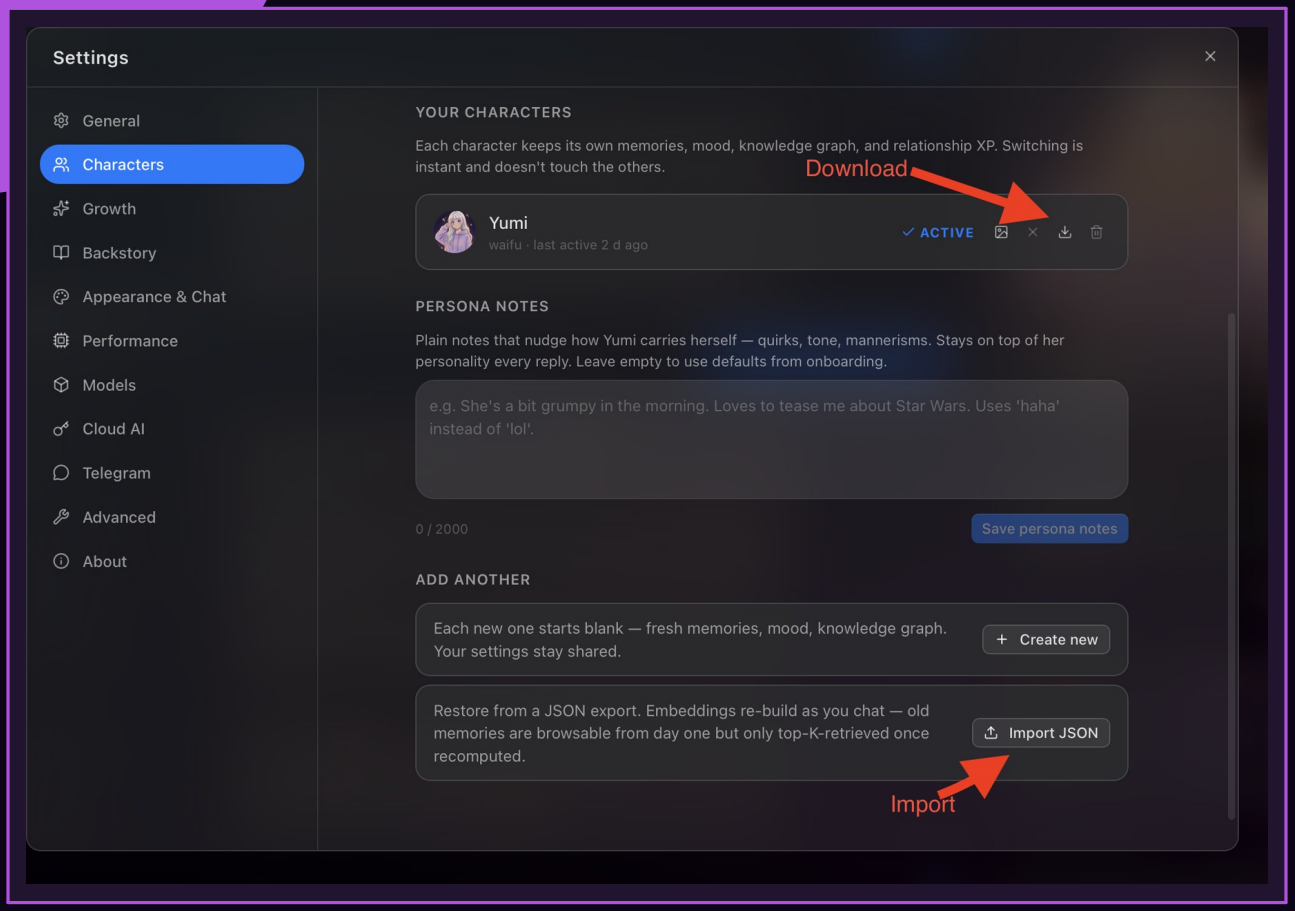
Backstory and the Lorebook

World facts, places, and recurring lore belong in the Lorebook, which the interface calls Backstory, not in Persona Notes. A Lorebook entry stays out of the conversation until one of its keywords appears in your message, so a paragraph of setting detail does not cost tokens on every turn.

Open Settings → Backstory to write and manage entries. Chapter 17 covers the entry fields, the always-on case, and the common mistakes.

FIGURE 06

Character transfer: export with the download icon; import with Import JSON. The red marks identify both controls.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

15. Export and Import a Character

Character export is the recommended self-service method for moving a companion to another installation.

Export a character

Goal

Create a portable JSON file for one character.

Steps

1. Open Settings → Characters.
2. Find the character row.
3. Select the export button (Export as JSON).
4. Save the suggested file, named like `localwaifu-character-YYYY-MM-DD.json`.
5. Store the JSON securely because it can contain private messages and memories.

Expected result

A JSON package contains the character and associated data such as messages, memories, relationship state, and knowledge entities.

Import a character

Goal

Create a new character from an exported JSON file.

Steps

1. Open Settings → Characters.
2. Select Import JSON.
3. Choose the exported character file.
4. Wait for the summary of imported messages, memories, and entities.
5. Select the imported character.
6. Send a message and confirm that she recalls the imported memories.

Expected result

The import creates a new character; it does not overwrite an existing one.

On macOS the export carries the stored semantic vectors, so recall works immediately after the import. On Windows the import stores the memories without vectors, so only memories created after the import take part in semantic recall; older ones stay reachable through the Memory Book and keyword matching.

If something goes wrong

- Import failed: the JSON is damaged, incomplete, or incompatible. Export it again from the source installation.

- Weaker recall right after an import is expected, and on Windows it is permanent for the imported memories rather than temporary. Keep the required memory model installed, and let the conversation continue so new memories do get vectors; keyword matching still finds the older ones.
- Character export does not carry installed models, API keys, Telegram credentials, global preferences, or a Direct license activation.
- Character export carries knowledge entities, not the relationships between them. Those are re-extracted as you keep chatting.
- For a whole-app recovery, also keep a full backup. Version 1.8.0 has no button that restores a backup over your live data (Chapter 39).

Delete a character

Deleting a character is destructive and requires confirmation (Delete permanently). All memories, conversations, mood, knowledge, and XP are removed, and she is taken out of every group chat. A two-person group becomes a normal chat with the other character, keeping the messages you wrote there. The last remaining character cannot be deleted. If the deleted character was active, Local Waifu selects another available character. Export first if you may need the character later.

16. Views, Settings, and Where Everything Lives

Local Waifu has three main views, two extra window modes, and one settings drawer. Almost every "where is that option" question is answered by the second half of this chapter.

Switch views

View	What it is
Full view	The complete interface: sidebar, chat stream, right panel
Mini view	A compact always-available window with the conversation and composer
Menu bar view (macOS) or tray view (Windows)	A small popover opened from the menu bar or system tray
Visual novel	Story scenes (Chapter 24)
Desktop pet	Only her figure on the desktop (Chapter 22)

Steps

1. On macOS use `Command+Option+1` for Full view, `+2` for Mini view, `+3` for Menu bar view, `+4` for Visual novel, and `+5` for the desktop pet.
2. On Windows use `Ctrl+Shift+1` to `Ctrl+Shift+5` for the same five.

On macOS, `Command+Shift+1` and `Command+Shift+2` also still work for Full and Mini. `Command+Shift+3`, `+4`, and `+5` are macOS screenshot keys, which is why the other modes use Option.

Expected result

The window changes shape and keeps the same conversation. Returning to Full view restores the sidebar and right panel.

If something goes wrong

- If a shortcut does nothing, an operating-system shortcut may already own it. Check the system keyboard settings for a conflict.
- If the window seems to have vanished, use the menu bar icon (macOS) or tray icon (Windows) to bring a view back.

Open Settings

Press `Command/Ctrl+,`, or select Settings at the bottom of the left sidebar. The drawer opens over the chat and closes with `Escape`.

Find a setting by tab

The drawer lists these tabs, in this order. This is the map.

Tab	What it controls
General	Interface language, start at login, hide to tray, the chime and banner when she reaches out, Go quiet during fullscreen, meetings and games, Let her message you first, Blur on Background, and (Direct edition) the Updates block
Characters	Create, switch, rename, export, import, and delete characters. Portraits, persona notes, and your own photo
Growth	The relationship dashboard: stage, days known, what she has stored, your writing style, and the story timeline
Backstory	The Lorebook: keyword-triggered backstory entries that only enter the prompt when they are relevant
Appearance & Chat	Background theme, text size, Markdown rendering, and whether Enter sends
3D Avatar	Her 3D body: models, backgrounds, poses, animations, behaviour, selfies, desktop size, camera
Performance	Storage, runtime status, compute device, context size, low power mode, deeper thinking, suggested replies, network and proxy, Offline Mode, and the cloud fallback switch
Models	Three segments in one tab: Chat Models, Image Models, and Voice Model
Cloud AI	Cloud provider keys and connected accounts, plus your own endpoints
Telegram	Bot token, start and stop, reply mode, pinned owner, and the command list
Routines	Her day, morning brief, night research, watchers, and (macOS) reminder follow-up
Screen	Screen awareness and its modes
Voice / TTS (macOS)	An older optional voice-engine panel. Her voice settings are in Models → Voice Model
Advanced	Shortcuts links, app logs, the AI engine version, diagnostics report, reset, backup and restore check, and conversation export
Developer	Her system prompt sections, generation parameters, presets, and the avatar test bench
License	Activation and status. Absent in the Store/Unlocked edition
About	Version, edition clues, licences, and credits

The Voice / TTS tab appears only on macOS. It manages an older voice engine (Kokoro voices and voice cloning) and may report that engine as offline. Windows has no such tab. You can ignore it unless support asks you to open it.

If something goes wrong

- A tab that exists in this table but not in your app belongs to an edition or platform you are not running. License is intentionally missing in the Store/Unlocked build.
- Other tabs exist in the code but are not shown in the shipped drawer. They are unfinished and not part of 1.8.0.

Choose a background theme

Goal

Change the dashboard background without changing the conversation.

Steps

1. Open Settings → Appearance & Chat.
2. Pick one of the eight themes: Plum (default), Ocean, Forest, Sunset, Cyberpunk, Mocha, Midnight OLED, or Snow.
3. Look at the chat stream: the change is immediate.

Expected result

The background changes and the conversation stays exactly as it was. Midnight OLED is the darkest option, and it is the one to pick on an OLED screen or in a dark room. Snow uses a light panel with dark text.

If something goes wrong

If the background looks washed out, another display profile or an operating-system accessibility setting may be altering colours. Return to Plum, which is what the screenshots and this guide were written against.

Tune the chat itself

Open Settings → Appearance & Chat and adjust:

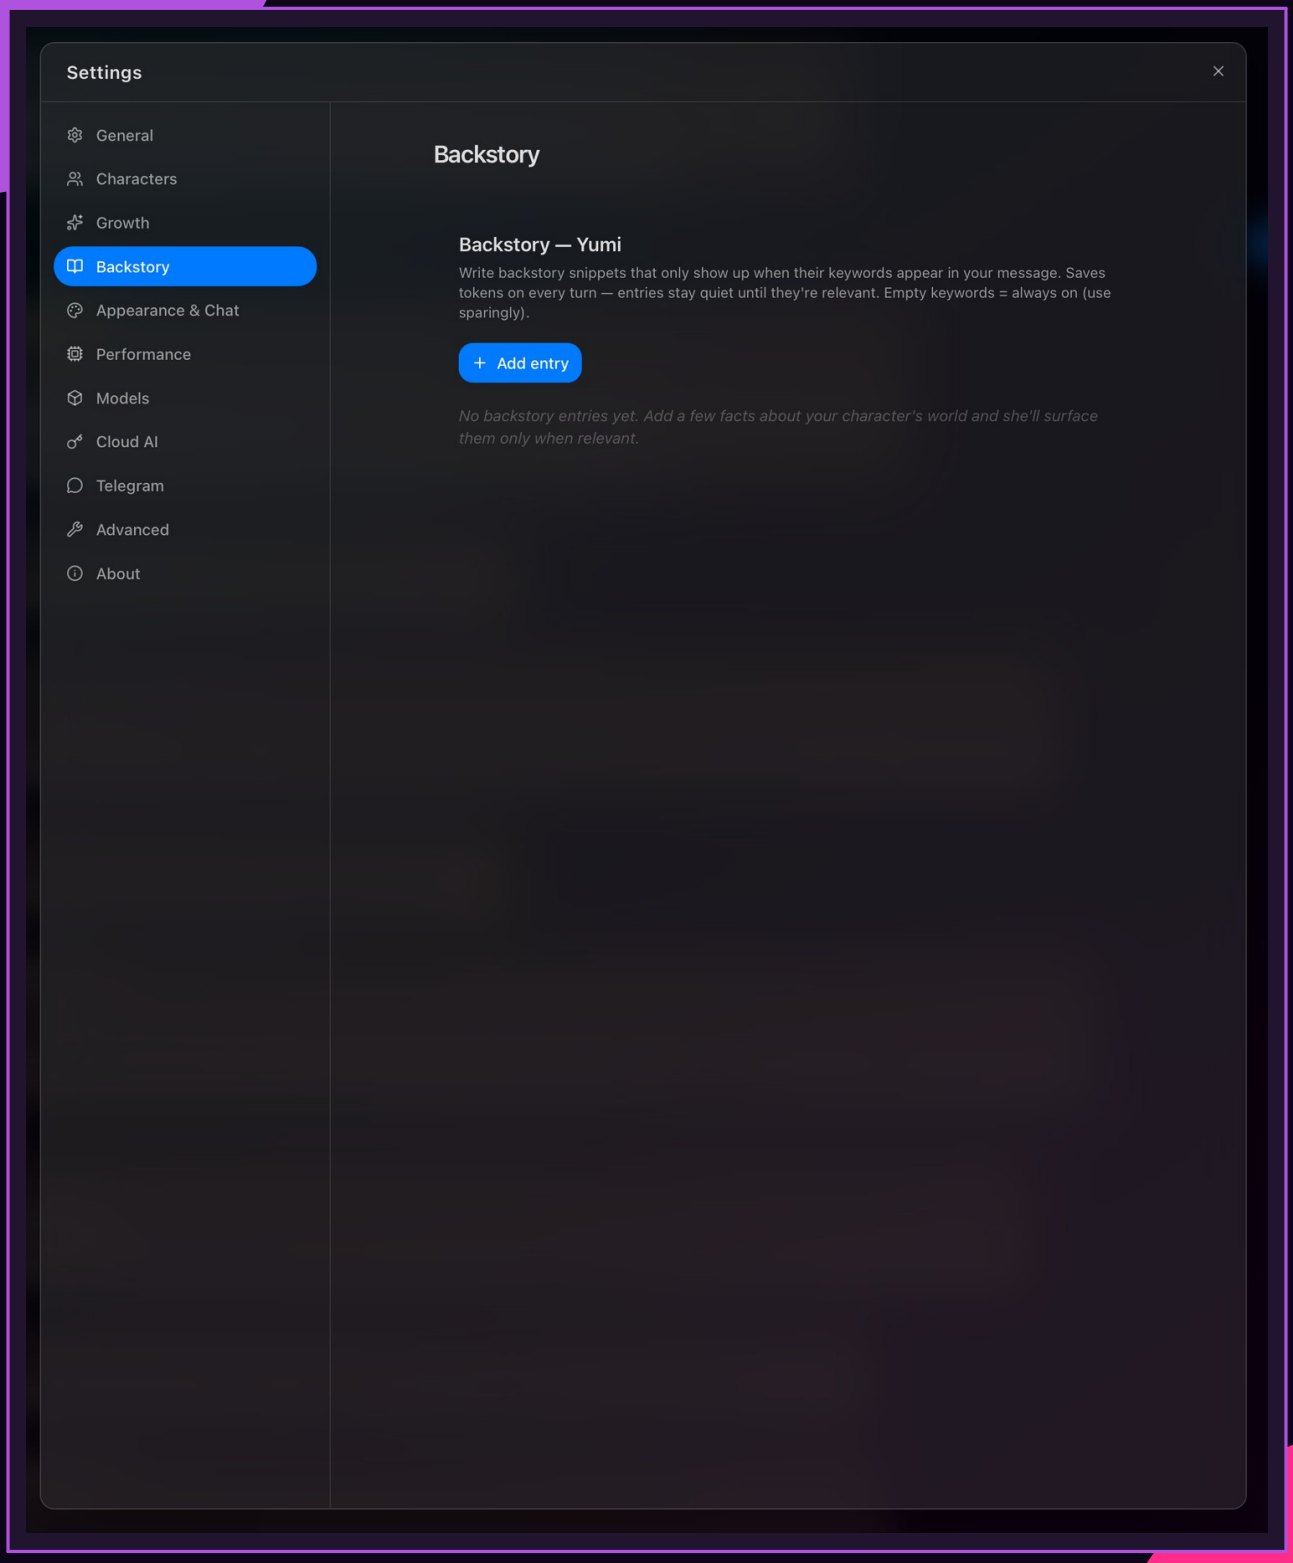
- Text size: small, normal, or large. This affects message text only.
- Markdown: render formatting in replies, or show the raw characters where you prefer plain text.
- Send on Enter: on sends with Enter and adds a new line with Shift+Enter; off makes Enter a new line so you send deliberately.

If something goes wrong

If pressing Enter starts a new line instead of sending, Send on Enter is off, or an input-method editor is composing a word. Confirm the composition first, then press Enter again.

FIGURE 07

Backstory/Lorebook: trigger specific world knowledge with keywords and priorities.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

17. Use the Lorebook for Backstory

The Lorebook holds backstory entries that stay out of the conversation until they are relevant. In the interface it is the Backstory tab. A long biography sent on every message costs tokens, slows replies, and dilutes what she is actually answering. A Lorebook entry waits for its word.

Understand the trade

- Every entry has a title, trigger keywords, and lore content.
- When one of the keywords appears in your message, the entry is added to that turn.
- An entry with no keywords is injected into every turn. The app marks it as always on, and it counts against the token budget on every message.
- Priority orders entries against each other when several match at once.
- The toggle on the left of each row turns an entry on or off without deleting it.

Add a backstory entry

Goal

Give her a piece of world knowledge she brings up only when it matters.

Before you start

Decide what the entry is for, and write down the handful of words you would use when you want her to remember it. Those become the keywords.

Steps

1. Open Settings → Backstory.
2. Confirm the header (Backstory followed by her name) names the character you expect. The Lorebook belongs to the active character. If no character is active, the tab asks you to pick one first.
3. Select Add entry.
4. Expand the new row and fill in:

- Title: a short name you will recognise later, such as `Her apartment`.
- Trigger keywords: comma-separated words, for example `apartment, flat, home, bedroom`.
- Lore content: the facts themselves, written the way you would explain them to a person.
- Priority where the form offers it.

5. Leave the toggle on.
6. Select Save changes. A row marked unsaved has not been stored yet.
7. Return to the chat and send a message containing one of the keywords.
8. Confirm that her reply uses the detail.

Expected result

She answers with the stored detail on messages that mention a keyword, and ignores it otherwise.

If something goes wrong

- If the entry never arrives, check the spelling of the keywords against the words you actually type. Matching is literal, so `apartment` does not match `apartments`.
- If she brings it up too often, the keywords are too broad. Narrow a generic word like `home` to something more specific.
- If several entries compete and one keeps winning, use priority to settle it.
- If every reply feels heavier and slower, look for an entry with no keywords. An always-on entry is the usual cause.

Switch off or delete an entry

The toggle on the left of a row switches the entry off without deleting it. A switched-off entry says Off. This entry is ignored until you switch it on. To remove an entry, use its delete button and confirm Delete this lore entry?.

Make an entry always on

Leave the keywords field empty. The row is then labelled always on.

Use this sparingly. It is the right choice for one or two facts that shape every single reply, such as her name for you or a hard rule about her character. It is the wrong choice for a paragraph of setting detail.

If something goes wrong

If replies got noticeably slower after adding an always-on entry, turn its toggle off and give it keywords instead.

Know what the Lorebook is not

- It is not memory. It does not learn anything, and it does not change as you talk.
- It is not the persona or the personality sliders. Those live under Settings → Characters.
- It is not training. Nothing you write here is sent anywhere or used to train any model. It is a local entry in your own database, and it is included in a full backup.

18. When She Messages You First

Local Waifu can start a conversation instead of only answering one. These are called proactive messages, and they are one of the clearest differences between a companion and a chat box.

The four triggers

Trigger	When it fires
Morning greeting	Once a day, between 07:00 and 09:00 local time
End of work check-in	Once a day, between 17:00 and 18:00 local time
Evening wind-down	Once a day, between 22:00 and 23:00 local time
Long idle	You have been quiet for the active character's idle threshold, four hours by default. Only considered between 10:00 and 22:00

A deliberately cold thread can also surface as a message about that topic instead of a generic "you there?", so a callback sometimes arrives with a subject attached.

Earlier versions also pinged you before a calendar event. The calendar feature was removed in 1.8.0, so there is no calendar trigger and no calendar tab. Local Waifu does not read your calendar.

What keeps them from becoming noise

- At most three proactive messages per hour, across every trigger.
- One per day for each of the three clock-based triggers.
- Held while you are typing. A nudge that arrives mid-message waits until the conversation goes quiet, and it carries a deadline past which it is dropped rather than delivered late. A morning greeting that could only be delivered at 23:00 is not delivered.
- Held while she is quiet (see the next section).
- Quiet hours per character. Each character learns when you tend to be asleep, and long-idle nudges respect that.
- The idle threshold adapts to the character rather than being a fixed number for everyone.

She goes quiet by herself

Goal

Understand when she holds back everything you did not ask for, and control it.

While Quiet Mode is on, she holds back proactive messages, routine deliveries, speech bubbles, and idle motion. Replies to your own messages and anything you start yourself are never held. The sidebar shows a Quiet badge, and a held nudge looks again a few minutes later.

She goes quiet automatically in these situations:

Situation	macOS	Windows
A fullscreen app is in front (a game, a film)	Yes	Yes
A meeting app is in front (Zoom, Teams, FaceTime, Webex, a Slack huddle, Google Meet or Teams in a browser tab)	Yes	Yes (Zoom, Teams, Webex, Slack huddles, and Meet in a browser)
The screen is locked	Yes	Yes
The screen is being recorded	Yes	No
The Mac is asleep	Yes	No
An app you listed is in front	Yes	Yes

Steps

1. Open Settings → General.
2. Find Go quiet during fullscreen, meetings and games. It is on by default.
3. To add your own apps, type one per line under Also stay quiet when one of these apps is in front, for example Steam or Final Cut Pro.
4. The line under the switch tells you whether she is quiet right now (Right now: quiet. or Right now: she can speak.).

Expected result

While one of those situations applies she stays silent, and Telegram messages from routines are held too on macOS. She speaks again when it ends.

If something goes wrong

- If she stays quiet when she should not, check which app is in front. A fullscreen window of any app counts.
- Turn the switch off to disable Quiet Mode entirely.
- On macOS, reading a browser tab title for Meet or Teams needs Screen Recording permission. Without it the browser case may not be detected.

Turn proactive messages off, or tone them down

Settings → General → Let her message you first is a single switch: turned off, she sends no greetings, check-ins, or idle nudges and only answers when you write. Routines you set up still run. The finer controls are:

1. Long idle: this fires only when the app is running and you have been quiet past the character's threshold. Closing the app is a complete off-switch, because the trigger lives inside the app process.
2. The learned behaviour: repeated dismissing, along with how you actually write, feeds the character's adaptive dials, including its idle threshold and quiet hours (Chapter 20). A companion that has been left alone for a while becomes less forward on its own.
3. During onboarding you chose how often she should reach out first. That choice only sets the starting point.

If something goes wrong

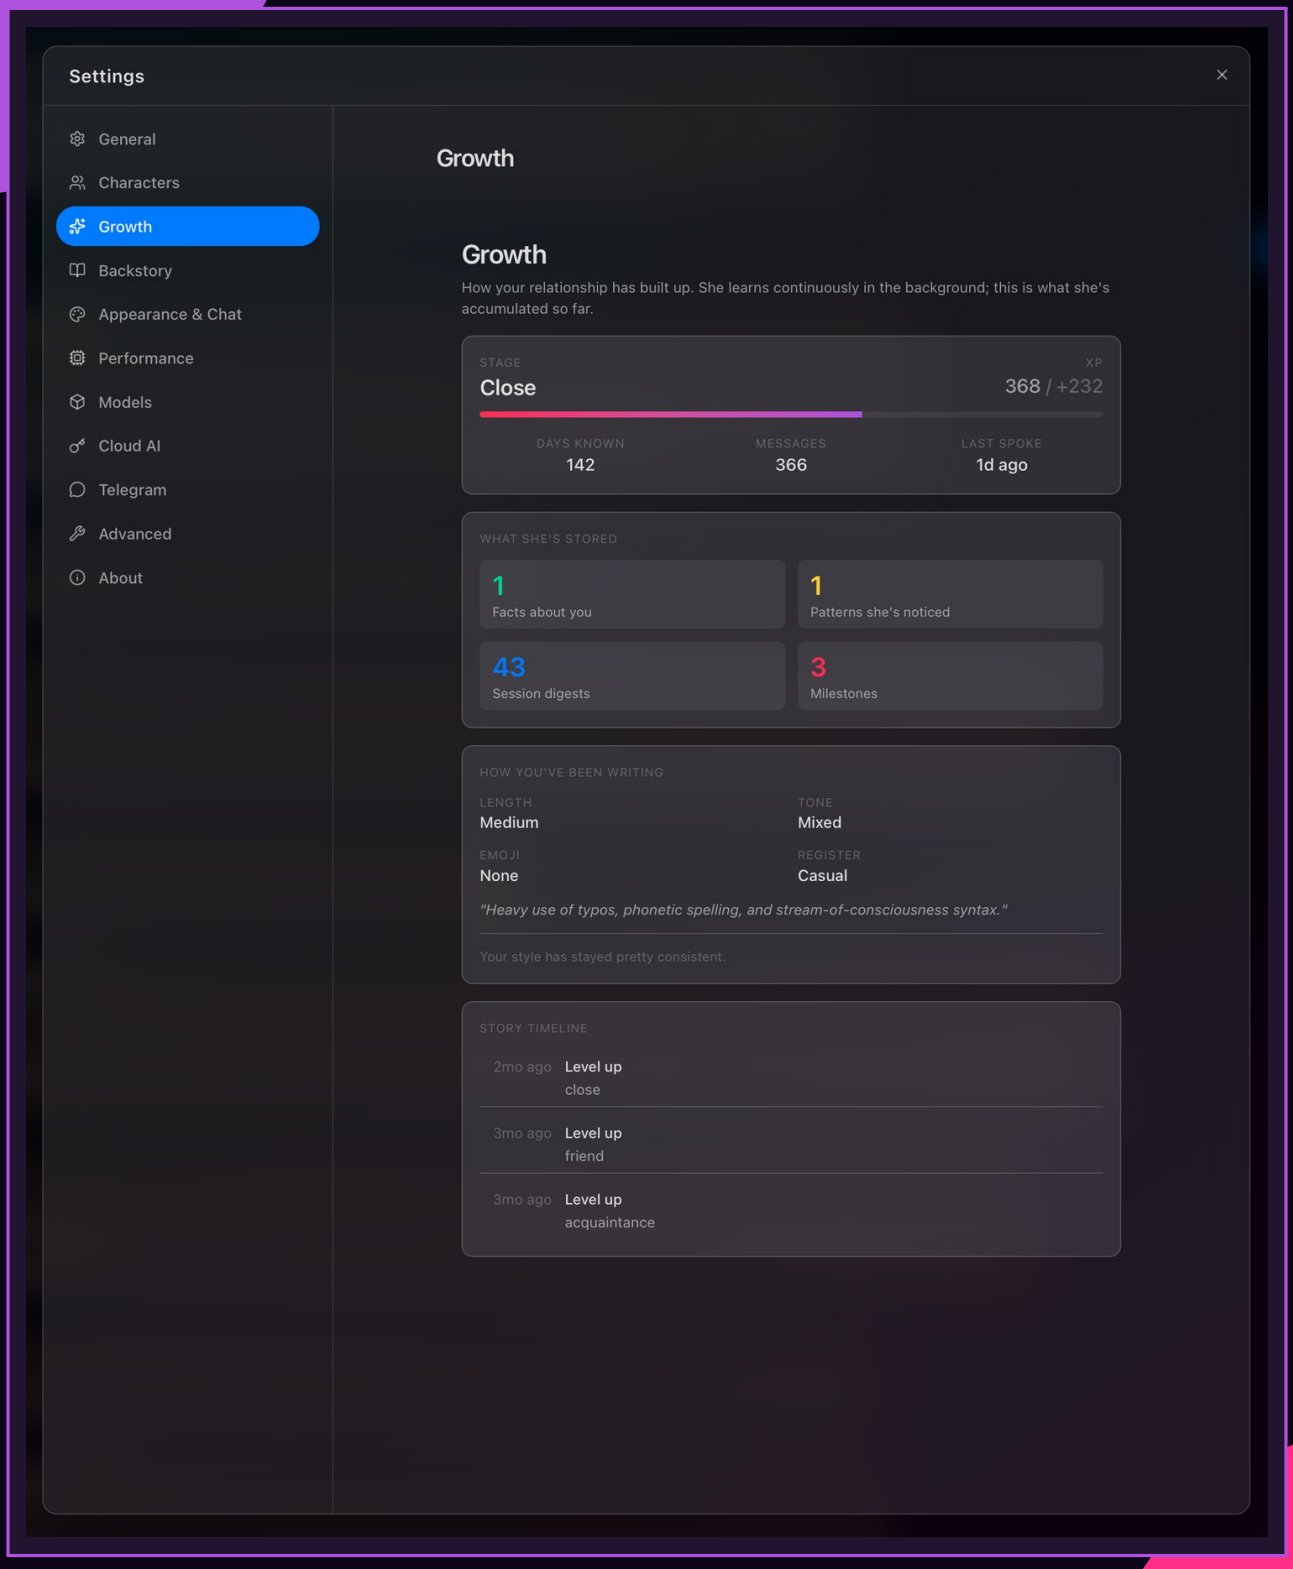
- If a message arrives while you are mid-sentence, that is the deferral system failing its deadline, not a new setting. Note the time and report it.
- If a proactive message seems to answer something you typed several minutes earlier, it most likely started generating before your message landed. The current generation is not cancelled by a new message, and this is a known limitation rather than a setting.
- A good-morning message whose delivery failed is retried.

What a proactive message is not

It is not a notification service and it is not a scheduled task. Each one is a real generation from your active model, produced locally like any other reply, and it appears in the conversation as a normal message you can answer, edit, or ignore. For things she does on a schedule, see Chapter 25.

FIGURE 08

Growth shows relationship progress, memory counts, milestones, and learned adaptations.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

19. Read the Growth Dashboard

Settings → Growth is the record of what has built up between you. Everything on it comes from your own database on your own computer.

The seven stages

The relationship advances as you talk. XP accumulates from real interaction, and each stage begins at a fixed total.

Stage	Lvl	Starts at
Stranger	1	0 XP
Acquaintance	2	10 XP
Friend	3	50 XP
Close	4	200 XP
Romantic	5	600 XP
Bonded	6	1500 XP
Partner	7	3500 XP

Partner is the top stage. The bar reports itself as full rather than pointing at a level that does not exist.

The sidebar shows the pair together, as `Lvl 5 · Romantic`, and the progress bar fills from the start of the current stage rather than from zero overall.

If something goes wrong

If the stage looks wrong, check which character is active. Growth is per character, and switching companions switches the dashboard with it.

What the dashboard shows

- Days known, Messages, and Last spoke: the plain counters.
- Stage: the name and level from the table above, with progress inside the current level.
- What she has stored:
- Facts about you,
- Patterns she's noticed,
- Session digests,
- Milestones.
- How you've been writing: length, emoji use, and register, compared against how you wrote earlier, ending in either "your style has stayed pretty consistent" or "your style has shifted since you started".
- Story timeline: dated milestones such as first joke or a vulnerable moment, newest first.

If something goes wrong

If a section says it needs more history, that is accurate rather than broken. Style analysis needs enough messages before it will report a comparison, and the timeline fills as milestones occur.

Read it honestly

Two things are worth knowing before you treat a number as a measurement of a relationship.

1. The stage is a counter. XP comes from interaction, so the number tracks activity rather than depth. A long quiet friendship and a burst of conversation do not land in the same place, and the stage does not decide how she feels about you.
2. The style analysis describes your messages, not her personality. It is a readout of how you write, which the app uses to adapt tone and pacing.

Clear or change what she has stored

- A single memory can be edited or deleted in the Memory Book. See Chapter 13.
- Style history and milestones are not individually editable in 1.8.0. Removing them means removing the data they live in, so treat Reset All as the only complete route and read Chapter 41 before using it.
- A full backup preserves all of this. Character export carries the memories and entities, but not the knowledge-graph relations. See Chapter 39.

20. Tune How She Has Adapted to You

She changes as you talk to her. Most of that happens in the background and is intentional. The panel called How she's adapted to you is where you look at it, and where you take control of anything that drifted somewhere you do not want.

Open it

1. Open the Memory Book.
2. Scroll to How she's adapted to you.

The header says which character it belongs to. Adapting is per character, like memory and mood.

The four groups and ten dials

Group	Dials
Conversation style	Playful teasing, Message length, Physical narration, Emoji use
Photos	Sends photos when asked, Shares photos unprompted
Reaching out	Reaches out first, Respects quiet hours
Memory	Memory reach, What she bothers to remember

Each dial sits at a position she has moved to over time. The subtitle under the heading says it plainly: these update slowly from how you talk.

Lock a dial

Goal

Keep one part of her behaviour where you want it, instead of letting it keep drifting.

Steps

1. Open the Memory Book and find How she's adapted to you.
2. Drag the dial to the value you want.
3. Confirm that it is now marked Set by you.

Expected result

That dial stops moving on its own. Everything else keeps adapting.

If something goes wrong

- If a dial is marked Still learning, she has not seen enough conversation to be confident about it yet. Locking it is still allowed; she simply had less evidence behind the position she started from.
- If you change your mind, use Let it adapt again on that dial to hand it back.

Reset everything she has learned

Select Reset all to default.

What this does: returns every dial to its starting position and clears the locks you set.

What this does not do: it does not delete the character, your conversations, or the Memory Book. Adapting is a layer on top of the personality you built, not the personality itself. Chapter 13 covers that distinction.

If something goes wrong

If the reset appears to have had no effect, check that the active character is the one you meant. There is one dial set per character.

Read this honestly

Two things worth knowing before you judge a dial.

1. A dial is a lean, not a personality. It nudges how she writes and when she reaches out. It does not overwrite the personality sliders or the persona notes you wrote in Settings → Characters.
2. She adapts to what you actually do, not to what you said you wanted. If you always answer in three words, message length drifts short. If you never respond to a morning greeting, reaching out drifts down. That is the mechanism working, and it is also why the locks exist.

21. Give Her a 3D Avatar

By default she appears as a flat portrait. The 3D avatar replaces the portrait with a moving model that blinks, breathes, follows your cursor with her eyes, moves her lips when she speaks, and gestures while she talks. It is off until you turn it on.

Turn it on and pick a model

Goal

Show a 3D model in place of the portrait.

Before you start

- Open the chat once so you know which character is active. The model belongs to the active character, so every character can have her own body.
- You do not need to find a file first. Ready-made models are built in.

Steps

1. Open Settings → 3D Avatar.
2. Turn on 3D avatar.
3. Under Model, choose one of the presets. They include six VRoid Hub samples: Avatar Sample A, Avatar Sample B, Avatar Sample C, Darkness Shibuya, Vita, and Sendagaya Shibuya.
4. Select Use for and her name. The preview on the left of the tab shows her at once.
5. Close Settings. She now stands in the sidebar instead of the portrait.

Expected result

The model blinks and breathes in the sidebar. If it cannot load, the app shows the portrait again and tells you the 3D avatar could not load.

If something goes wrong

- If the preview stays empty, pick another model. A model that fails in the preview cannot be used.
- The model list has a search box and keeps a thumbnail of each model you import yourself.

Import your own model

Goal

Use a character you made or downloaded.

Before you start

- Accepted files are `.vrm` (VRoid Studio and VRoid Hub), `.pmx` (MikuMikuDance), or a `.zip` with a model inside. The limit is 100 MB, or 300 MB unpacked for a `.zip`.
- Check the licence. The import shows a licence card: free to use, personal use only, credit required, or no redistribution. MMD models keep their terms in a readme, and many authors allow private use only.
- The model stays on your computer.

Steps

1. Open Settings → 3D Avatar and make sure 3D avatar is on.
2. Under Model, select Import a model, or drop the file onto the tab.
3. Wait while the app checks and copies the model. Large models can take a few seconds.
4. Read the licence card.
5. Select Use for and her name.

Expected result

The model appears under Your models and is assigned to her.

If something goes wrong

- This file is not a VRM model, This is not a 3D model file, and This model is too large mean the file is the wrong kind or over the limit.
- If you drop several files at once, only the first one is imported.
- Some textures are missing: import the `.zip` the MMD model came in, not the bare `.pmx`.
- A model with no mouth shapes blinks and smiles but its lips do not move.
- Deleting a model that she is using sends her back to the default model, or to the portrait.

The tab also carries a short guide called How to get your own character: it explains making a model yourself in VRoid Studio (free), downloading one from VRoid Hub, and using MMD models. It links BOOTH, Gumroad, and DLsite for more VRM and MMD models.

Fix her face with Expression mapping

Most models are matched automatically. If a feeling does not show on her face, or shows the wrong one, open Expression mapping under Model. Each feeling and eye or mouth shape lets you choose which of the model's own expressions plays it, and Test shows it on the preview. Back to automatic undoes your choices.

Backgrounds, camera, and behaviour

- Background: ten built-in scenes (such as Cafe, Library, Neon street, Sakura park), No background, or your own picture. A picture can be PNG, JPEG, or WebP up to 25 MB, and Blur and Darken adjust it.
- Camera and look: Framing in the sidebar (Bust, Face, or Full body), Zoom, Height in frame, Brightness, and Render quality. Reset camera and look restores the defaults. If she stutters, lower Render quality, because very detailed hair or clothes use more battery.
- Behaviour: She gestures on her own, Gesture size, Hand movements while talking, Idle fidgets, She does things on her own (sitting, stretching, reading, and now and then dancing), Use either hand, Eyes follow your cursor, Body language, and Reacts to touch. Turn off Reacts to touch if you do not want her to react when you pat or poke the model.
- Emotion tags: asks the chat model to mark where her feelings change, so her face changes at the right moment. It suits strong cloud or large local models, and every reply gets a little longer to write.

Ask her for gestures in chat

She picks gestures herself while she talks, but you can also ask. Type things like "wave", "sit down", "dance for me", "lie down", "go to sleep", "spin", or "make an angry face". In Settings → 3D Avatar → Animations, Stop ends an animation you are previewing.

If she is holding a pose you asked for, a Stand up button appears on her. Select it to let her stand again.

If something goes wrong

Gestures you ask for still play when She gestures on her own is off. If one cannot play on a particular model, the app says that gesture could not play on this model.

Poses and the pose editor

Pose under Settings → 3D Avatar holds ready-made poses (Arms crossed, Hands clasped, Thinking, Kneeling, and more), grouped as Standing, Cute, Sitting, Lying, On the floor, Actions, Dances, Moving, and Other, with a search box. She keeps breathing, blinking, and gesturing in any pose. Kneeling poses lower her to the floor, so use Full body framing to see it.

To make your own, select Create your own pose, choose a base under Start from, adjust the head, torso, arms, and legs with the sliders, choose a hand shape for each hand, name it, and select Save and use. Edit pose and Delete pose change it later.

Your own animations

Under Your animations you can import a `.vrma` file (VRM Animation, up to 30 MB). It works on both VRM and MMD models. Make one in Blender with the VRM Add-on, or download one where the licence allows it. MMD motion files (`.vmd`) and `.fbx` files are not supported directly.

After importing, choose what the animation is for under Use it as: Idle fidget, One of her gestures, Only when I play it, While she talks, or Full-body animation (dance, walk...). Select Play on the preview to see it. A full-body animation marked as a dance may be chosen when you ask her to dance. Hold a frame as her pose turns one moment of an animation into a pose.

Photo mode and selfies from the 3D model

When she is on screen and you are not in a call, a camera button on her opens Photo mode. Drag to turn her and scroll to zoom. You can change the Pose, the Expression, whether she looks at the camera, the Background (Plain, Transparent, or a picture), and an Animation frame to freeze. Take photo saves the shot to your Gallery, where it is labelled 3D photo.

Make it a real photo asks the image model to redraw the shot, with an optional scene description. It needs an image model that accepts a reference picture: local FLUX.2 klein, Google, or OpenAI. It does not work with SDXL, with a provider that only does text, or in Offline Mode with a local model that cannot take a reference. It also needs an active license. A redrawn shot is labelled 3D photo, made real.

Her selfies can copy her face from the portrait or from the 3D model. Open Settings → 3D Avatar → Selfies and choose Selfie reference: Her portrait, Her 3D model, or Both.

- Her 3D model renders her as she looks on screen right now. If the avatar is not on screen, or the active image model cannot use a reference picture, her portrait is used instead.

- The setting Use it for life and shared photos too also uses the 3D model for photos from her day and for her half of a photo with you. Off keeps her portrait there, which composes better into a new scene.

Windows and macOS

The 3D avatar works the same way on both. The only visible differences are the names of a few poses, such as Cat-paws pose on macOS and Cat paws on Windows.

22. Put Her on Your Desktop

Desktop mode shows only her figure on your desktop, with no window frame around her. She can walk around, sit on the edge of the Dock (macOS) or the taskbar (Windows), dance, and react when you touch her. Everything around her is click-through, so she does not block the apps behind her.

Desktop mode needs the 3D avatar. If it is off, or her model is not ready, the desktop shows a small card instead: Desktop mode needs her 3D avatar. Turn it on in Settings → 3D Avatar. Chapter 21 covers turning it on.

Put her on the desktop and bring her back

Goal

Move her from the main window to the desktop and return.

Steps

1. Select Put her on the desktop at the bottom of the left sidebar, or press the shortcut: `Command+Option+5` on macOS, `Ctrl+Shift+5` on Windows.
2. Drag her to move the window. A press on her drags it.
3. To return, hover over her and select Back to the full window in the small toolbar, or press the view shortcut for Full view.

Expected result

Only her figure remains on the desktop. The main window is gone until you bring her back.

If something goes wrong

- If the shortcut does nothing, check that the 3D avatar is on and that another app does not own the chord.
- If you open Settings, a voice call, search, or a deep link prompt while she is on the desktop, the app leaves desktop mode so the window can appear. It does not open inside the click-through window.

What you can do with her

- Double-click her to open the chat. It opens in Mini view.
- Hover over her to show a small toolbar with Open chat, Back to the full window, and a hide button (Hide to the menu bar, which is the system tray on Windows).
- Click or poke her and she reacts, unless Reacts to touch is off in the 3D Avatar settings.
- Watch for a bubble or a badge. Her newest reply shows as a short speech bubble above her for a few seconds. Select it to open the chat. If she wrote while you were away, a small letter badge (or a picture badge for a photo) appears and waves to get your attention.

The right-click menu

Right-click her to open the menu.

Item	What it does
Walk around	She walks along the desktop
Sit down	She sits on the edge of the Dock or taskbar
Lie down	She lies down
Dance	She dances
Stand	She stops what she is doing and stands
Poses	Opens the pose list, grouped like the 3D Avatar settings
Animations and Your animations	Plays an animation, including your own `.vrma` files
Move on her own	Lets her choose her own activity when nothing else is happening
Size	Small, Medium, Large, or Extra large
Stay on top	Keeps her above other windows

You can also set the size and Stay on top of other windows in Settings → 3D Avatar → On the desktop. The defaults are Medium and on top. With Stay on top off, windows you open can cover her.

When she stays put

When she has gone quiet by herself (a fullscreen app, a meeting, a locked screen), she stops walking and holds back bubbles. The menu says She stays put while Quiet Mode is on. Chapter 18 explains Quiet Mode.

If something goes wrong

- If a menu item says She can't do that right now, wait a moment and try again.

23. Use Group Chats

A group chat is one conversation with two to six of your own characters in it. You write once, and one or more of them answer. Each character keeps her own memory, mood, and voice.

You need at least two characters. If you have only one, the new-group button says You need two characters for a group chat.

Start a group chat

Goal

Create a conversation with several characters.

Steps

1. In the left sidebar, find the Groups list.
2. Select the plus button (New group chat).
3. Tick two to six characters. The dialog counts them, for example 3 of 6 selected.
4. Optionally type a Group name. If you leave it blank, the characters' names are used.
5. Select Create.

Expected result

The group opens, and the right side of the window shows the group panel with Members instead of the usual relationship panel.

If something goes wrong

If creating the group fails, the app says Couldn't create the group. Close the dialog and try again.

Who answers

Replies follow fixed rules, not a model's guess:

1. A character you name speaks. Name several and they all speak, in the order you named them.
2. Words that address everyone, such as "everyone" or "you all", make the whole group answer.
3. A long message, or one with a question mark, gets two speakers.
4. Otherwise one character answers. A very short follow-up goes to whoever spoke last.

An empty group suggests the first step: Say hi to everyone, or start a message with a name to talk to just one of them.

A name in a different grammatical form, such as a Polish name with an ending, may not be recognised. The default rules then apply.

Manage the group

The group panel lets you:

- rename the group;

- remove a member (Remove from group), as long as at least two remain;
- add someone (Add someone), up to six;
- read a character's reply aloud with the Read aloud button on her message. Each bubble is read in that character's own voice.

Voice messages (hold to record) work in one-to-one chats only. In a group the app tells you to type instead.

Deleting the group uses the delete button in the Groups list and asks you to confirm. If you delete a character who is in a group, she is taken out of every group, and a two-person group becomes a normal chat with the other character.

Let them chat while you are away

Goal

Let the characters talk among themselves when you are not around.

Steps

1. Open the group.
2. In the group panel, under They talk among themselves (Friends on Windows), turn on Let them chat while I'm away.
3. Choose After I've been quiet here for (macOS wording; After I've been away for on Windows): 1, 3, 6, or 12 hours.
4. Choose At most once every: 6, 12, or 24 hours.

Expected result

After you have been quiet in the group for the time you chose, the characters may talk among themselves for a short round. The sidebar shows They're chatting while it runs, and a While you were away note when you return.

If something goes wrong

- A message you write takes priority over a round the characters are having among themselves.
- If the app says They're in the middle of a conversation. Try again in a minute., a round is still running. Wait and try again.
- The feature is off by default for every group.
- These conversations earn no relationship XP.

24. Play Visual Novel Mode

Visual novel mode turns a conversation into a short story scene. She narrates and plays her part, and you either pick one of the offered choices or say something in your own words. It is an ordinary conversation underneath, so leaving the scene returns you to normal chat with the transcript intact.

Start a scene

Goal

Play a story scene with the active character.

Steps

1. Select Visual novel in the left sidebar, or press `Command+Option+4` on macOS or `Ctrl+Shift+4` on Windows.
2. Pick a scene. The picker says Pick a scene. She narrates, you choose what happens next. There are 24 scenes, including Café date, Rainy walk, Fantasy quest, Rooftop stargazing, Beach sunset, Snowed-in cabin, Night market, Night train, Haunted mansion, Space station, Cyberpunk alley, Kyoto tea house, Magic academy, Campfire night, Cooking together, and First snow.
3. Optionally turn on Generate backgrounds. A background is painted after her first reply.
4. Read her narration. Click or press `Space` to skip the typing.
5. Under Choose what to do, select a choice, or type your own line under Say something in your own words and select Send.

Expected result

She continues the scene. Each reply offers new choices.

If something goes wrong

- If a background cannot be painted, the app says it could not paint a new one. The scene continues with the previous background. New background asks for another.
- Generate backgrounds needs an image model or a key. If none is ready, the option says No image model is ready (set one up in Settings). Local image models can be slow.
- If the scene does not start, the app shows Could not start the scene and the reason.

Leave and resume

Select Exit to leave. Select New scene to choose a different one. Reopening the same visual novel conversation resumes the same story.

What she remembers

Fiction stays in the story. While a scene is open, her replies do not create memories, do not feed the Memory Book, and do not run the usual learning about you. When you leave a scene that had a real exchange, she stores one memory that you played that scene together.

macOS and Windows

Visual novel mode is the same on both platforms. Only the shortcut differs.

25. Set Up Routines

Routines are things she does for you on a schedule, without you writing first. Open Settings → Routines. The tab opens with this line: She can do things for you on a schedule. Every routine is off until you switch it on. A small dot on the Routines entry in Settings tells you when at least one routine is on.

All routine messages are written by your chat model and arrive in the conversation as normal messages. Each routine can also be sent to Telegram, and each one reports a status such as ok, failed, skipped, missed, or waiting.

Privacy: Routines write their messages with your chat model. If that model is a cloud model, what your routines gather (reminder titles, feed and search headlines, research topics and article text, her day plan) goes to that provider. The app shows a warning in this case. With every routine off, only the Find and Add buttons reach the internet: Find sends the city you type to a geocoding service, and Add fetches the address you paste. With routines on, weather comes from Open-Meteo, night research and brief topics are searched on DuckDuckGo, and feeds and articles are fetched from their own sites.

Common controls

Each routine card has a switch, and most have these controls:

- Time and Days for a daily routine, or Check every N hours for a repeating one.
- Run now to run it once immediately.
- Last run and Next run lines. A card marked Waiting, will run as soon as it's quiet was held back because you were mid-conversation or she was quiet. It retries on its own.
- Also send to Telegram, which needs the Telegram bridge turned on first (Chapter 35).
- A Paused note when Offline Mode is on and the routine needs the internet.

Routines also respect Quiet Mode. A routine due while she is quiet is held and delivered later.

Her day

Her day gives her a plausible daily schedule, so she knows what she is doing right now. It is planned at 06:00 by default.

1. Turn the routine on.
2. Under Today, read the plan. Each block has From, To, and Activity.
3. Use Add block, edit a block, or remove one. The times must not overlap. Unsaved changes reminds you to save before closing.
4. Regenerate asks for a new plan. It replaces anything you edited by hand, and the app asks you to confirm with Replace my edits. Clear today's plan removes it.

A plan you regenerate while the routine is off is saved, but she will not use it until you switch the routine on. Planning her day needs an active license.

Morning brief

Morning brief sends a good-morning message in her words. It runs at 08:00 by default and covers Weather, Headlines, and your Topics. It no longer includes a calendar.

1. Turn on Weather, enter a City, and select Find. If the city is not found the app says so. Without a location, weather is skipped.
2. Turn on Headlines and add up to five Topics, pressing Enter after each.
3. Feeds marked In morning brief under Watchers also appear in the brief.
4. Select Run now to see how it reads.

Night research

Night research lets you leave her a topic. She reads the web overnight and tells you in the morning.

1. Type a topic under Topic to research and select Queue, or type `/research` followed by a topic in chat.
2. She researches at Research at (01:30 by default) and tells you at Tell me at (08:30 by default).
3. The queue shows each topic as queued, ready, delivered, nothing found, or could not research. You can cancel a queued topic.

The queue holds ten topics. A topic that fails is tried again on later nights, up to three times. If Night research is off, `/research` queues nothing and offers Open Routines.

Watchers

Watchers announces new posts from RSS feeds and YouTube channels. Paste a feed address, a YouTube channel link, or an `@handle` and select Add. It checks every three hours by default. Each feed has Watching and In morning brief switches. If YouTube shows a cookie page instead of the channel, paste the channel's `youtube.com/channel/` address instead.

Reminder follow-up (macOS only)

Reminder follow-up asks how a reminder went after it came due. It reads your Apple Reminders, and macOS asks for permission once. It asks Ask after 90 minutes by default. It only follows up on reminders that came due after you switched the routine on. This routine does not exist on Windows. The card is hidden there because Windows has no reminders source.

If something goes wrong

- Weather skipped: no location set means the City is empty.
- No news from the feeds means no feed had anything new.
- A message held while you were typing is retried shortly after.
- If a routine seems to do nothing, check that its own switch is on and that Offline Mode is off.

26. Use Screen Awareness

Screen awareness lets her look at your screen: watch a film or play a game together, keep you on task, or answer a question about what you are looking at. It is off by default. The setting says it plainly: Frames are never saved. They are looked at once, in memory, then thrown away. Only her reaction is kept.

She needs a model that can see images. If none is available she says so: I need a vision model to see your screen. Set one up in Settings. Chapter 27 covers vision-capable models.

Turn it on

Goal

Allow screen awareness and give the app the system permission it needs.

Steps

1. Open Settings → Screen.
2. Turn on Screen awareness.
3. On macOS, allow Screen Recording. If the tab says Screen Recording is not allowed, select Open System Settings, switch Local Waifu on under Privacy & Security → Screen & System Audio Recording, and relaunch the app. macOS only applies the change after a relaunch. Select I've enabled it, check again afterwards.
4. Nothing is captured until you start one of the modes below.

Expected result

The tab shows Screen Recording is allowed (macOS) and the modes are ready.

If something goes wrong

- If permission is lost while she is watching, she stops and says Screen Recording permission was lost, so I stopped watching.
- If she cannot capture the screen, she stops and says so rather than retrying silently.
- On Windows there is no separate permission step. Protected video, such as some streaming apps, appears black to her, so she may not react to it.

Start a mode

Start a mode from Settings → Screen under Start, or from the Screen button in the message box (the monitor icon). While the master switch is off, that button only leads to the setting.

Mode	What she does
Watch together	She comments now and then on what you watch. By default she looks every 20 seconds. You can set Look every N s while watching from 10 to 120 seconds
Game buddy	Quicker, shorter reactions while you play

Mode	What she does
Focus guard	You tell her what you are working on and for how long. She glances at your screen and nudges you if you drift. By default she checks every 150 seconds, adjustable from 60 to 600
Ask about my screen	A one-off question. She captures the screen and answers about it

For Focus guard, a dialog asks What are you working on?, then For a duration (minutes) or Until a time, then Start focus.

While a mode runs, a banner shows what is happening, for example Watching your screen, Game buddy is watching your screen, or Focus guard until a time. The banner has Pause, Resume, and Stop.

Choose what she captures

- "Ask her" captures: The whole screen or The active window.
- Allow a cloud vision model: off by default. When it is on and no local vision model is available, your screen frames go to the cloud provider you chose. Leave it off to keep everything on your computer.
- Never capture these apps: a list of app or window names. If one of them is on screen, nothing is captured at all, and the app says Not captured: a protected app is on screen. Password managers are always protected and cannot be removed from the list.

Shortcuts

The Shortcuts block in the tab shows the two keyboard shortcuts for your platform: one to capture and ask her, and one to pause or resume watching. They are global, so they work from any app.

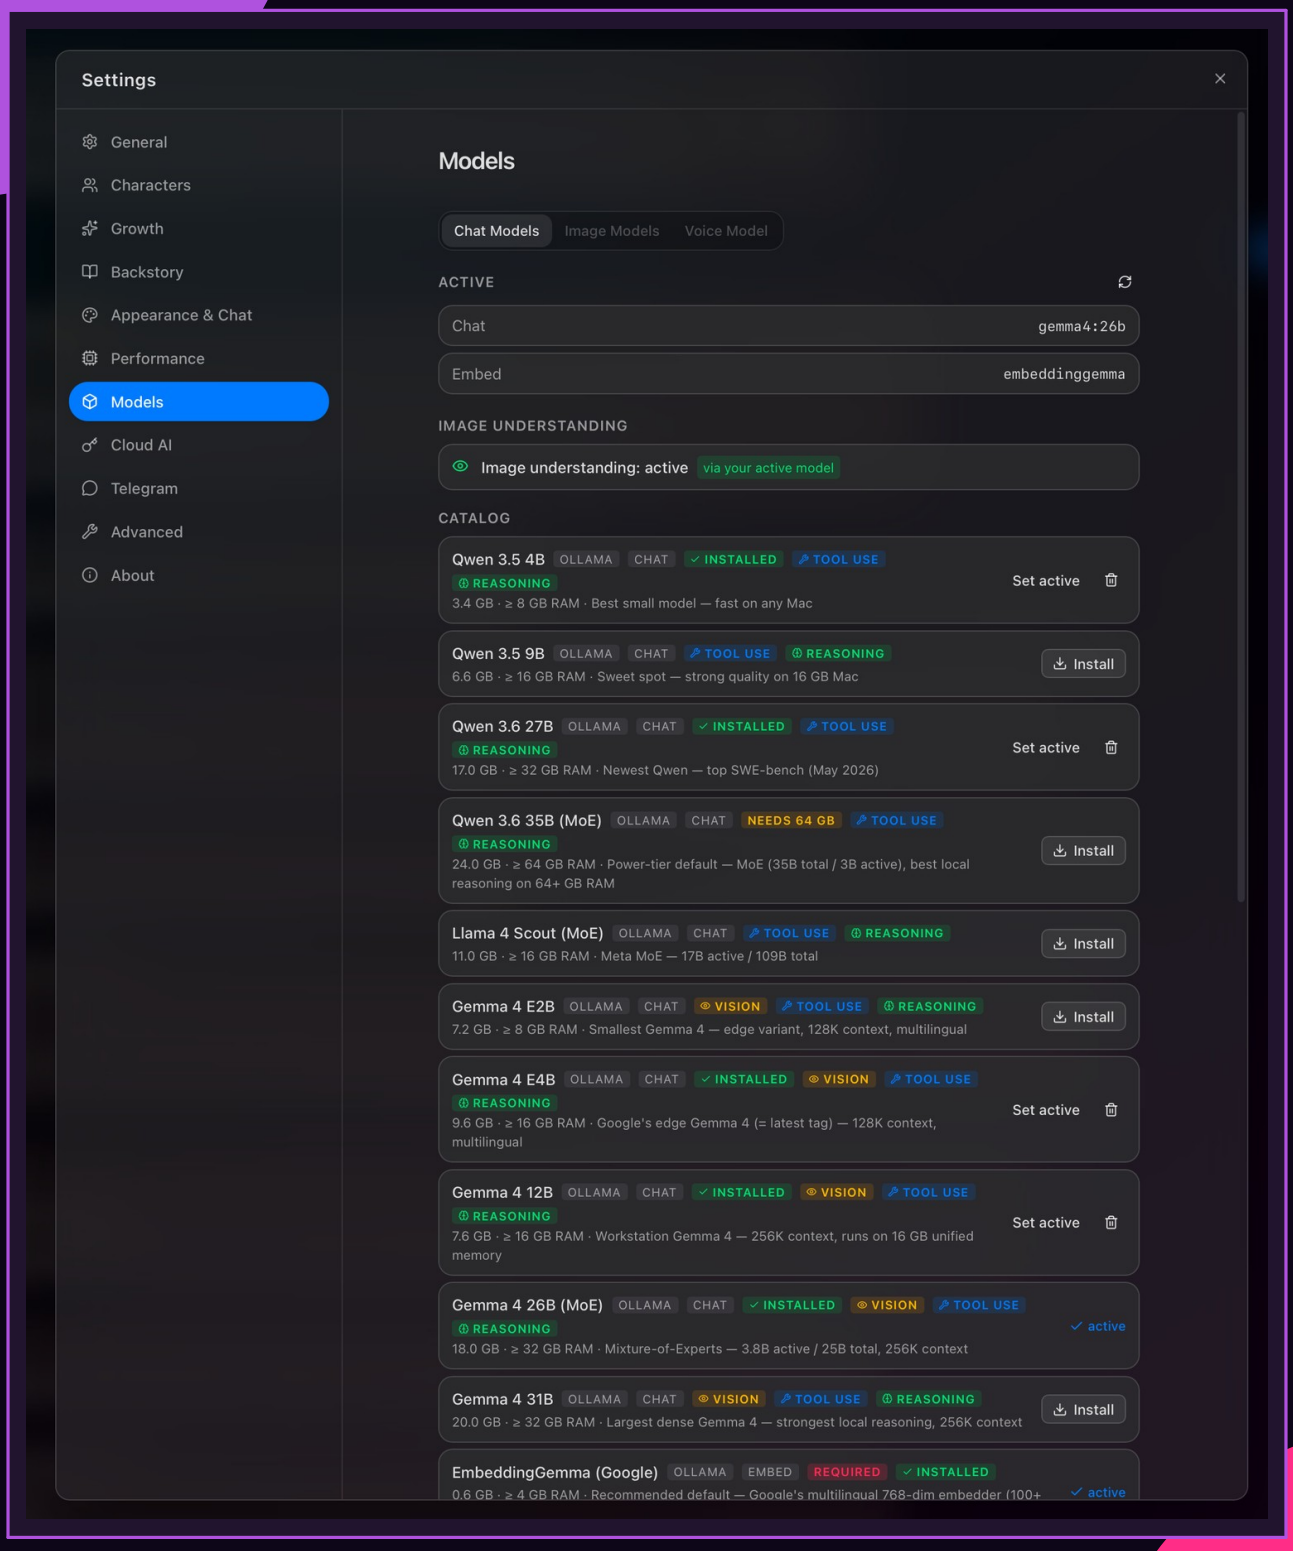
Quiet Mode and screen awareness

When Quiet Mode is on because of a meeting or a fullscreen app, watching pauses and the banner says Watching paused, Quiet Mode (meeting or fullscreen). Chapter 18 explains Quiet Mode.

Part III: Models, Images, and Voice

FIGURE 09

Chat Models: installed models, active role, capabilities, and memory-model status.

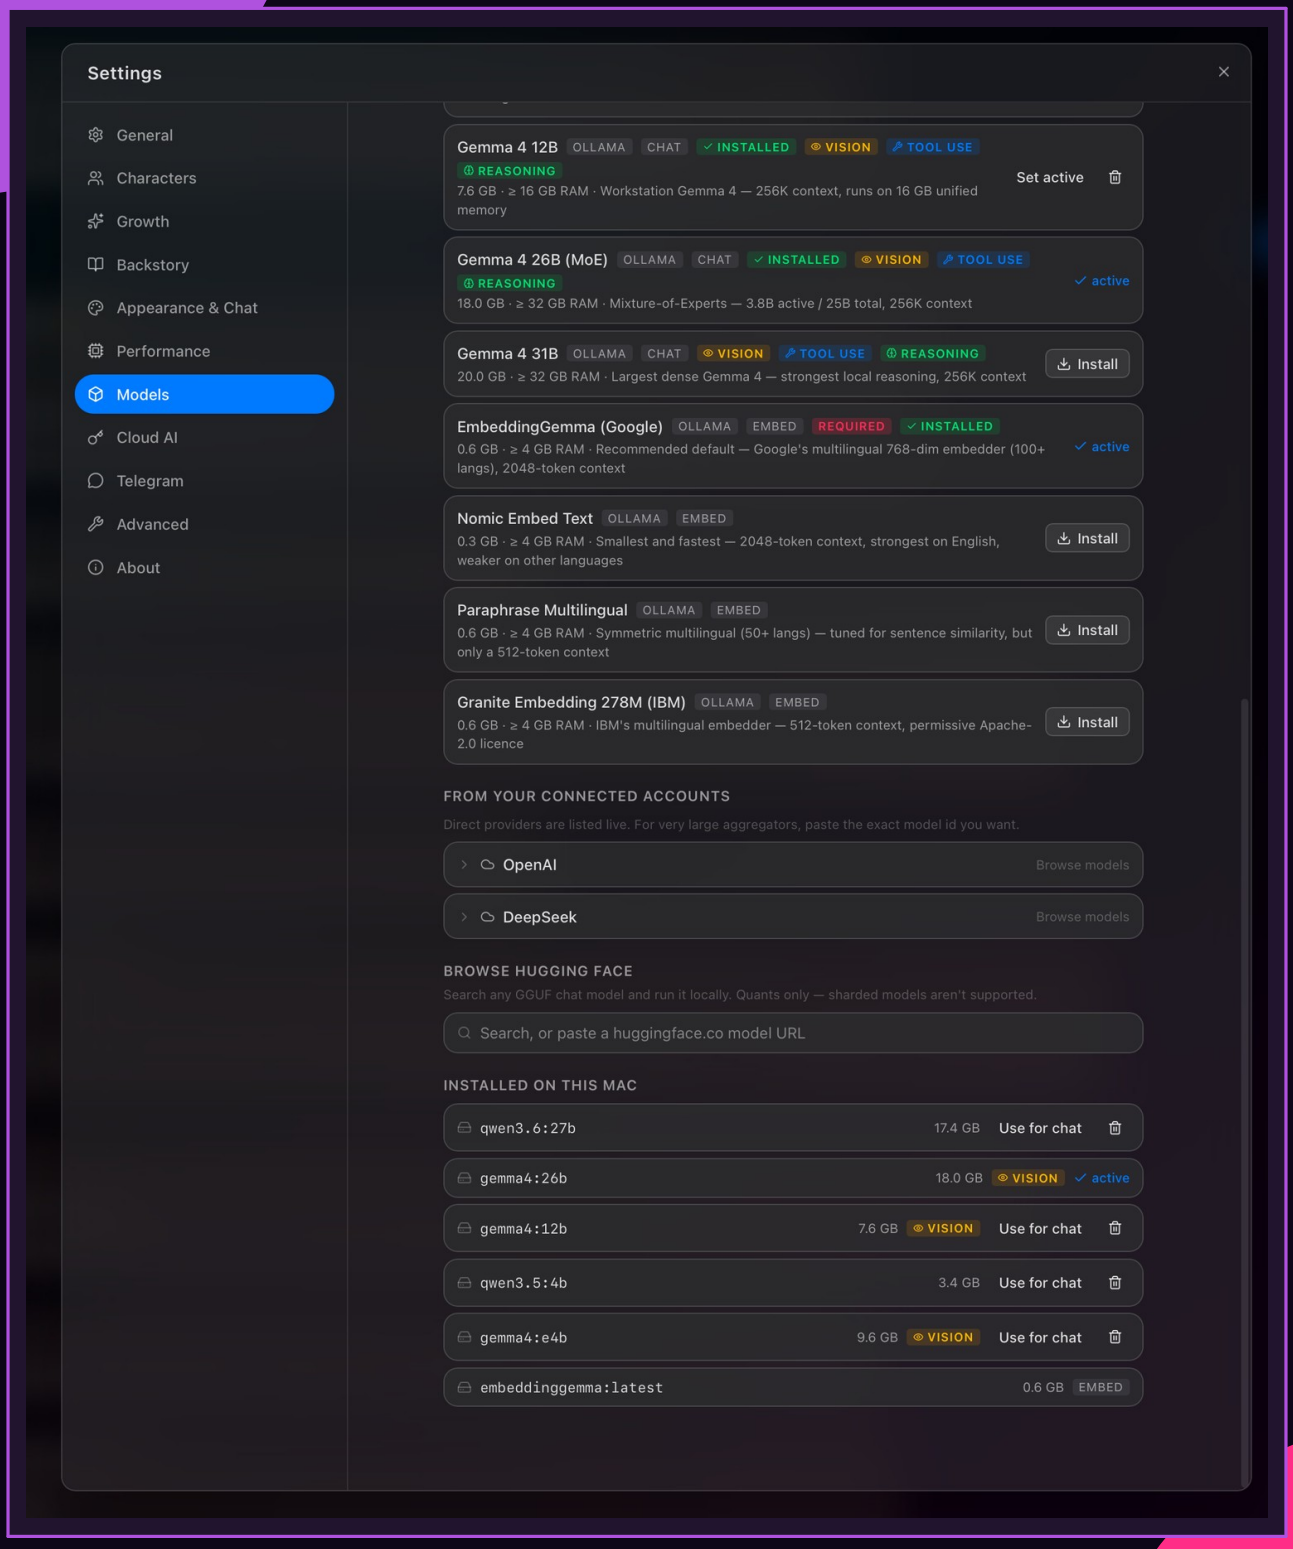


Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

FIGURE 10

Built-in catalog: compare size, RAM guidance, and capability badges before downloading.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

27. Install and Select Chat Models

Open Settings → Models → Chat Models to manage local and connected models.

Understand the two required roles

- Chat model: writes replies and may provide vision, tool use, or reasoning.
- Embedding model: converts memories into vectors so relevant information can be recalled. It does not write the reply.

The active row and badges communicate capabilities:

- Installed: model files are present;
- Active: currently selected for that role;
- Required: needed for a core feature such as semantic memory;
- Vision: can inspect images;
- Tool Use: can use supported tools;
- Reasoning: supports deeper reasoning behavior.

Install a recommended local model

Goal

Install a model appropriate for the computer and make it active for chat.

Before you start

Connect to the internet, verify free disk space, and keep Local Waifu open during the download.

Steps

1. Open Settings → Models → Chat Models.
2. Compare model size, RAM guidance, and capability badges.
3. Select Install.
4. Wait for the download and installation to finish.
5. Select Set active (Set default on Windows) or Use for chat.
6. Send a short test message.

Expected result

The model shows Installed and Active and generates a response.

If something goes wrong

- Check network, proxy/VPN, firewall, and disk space.
- Choose a smaller model if RAM pressure is high. On a Mac with 8 GB the app recommends Gemma 4 E2B (4.6 GB) rather than Gemma 4 E4B (6.6 GB), which leaves an 8 GB Mac swapping on every reply.
- Local chat keeps at least an 8K context window and may grow it to 16K when her prompt needs it, so a model that fits on disk can still be tight at runtime.

- On Windows, run Settings → Performance → Local AI acceleration → Check after sending a message.

Add a vision model

If you send her a photo and no vision model is available, the app shows Image understanding not set up with Install vision model and Add a cloud key. Under Settings → Models → Chat Models the Image Understanding block says whether images work: Image understanding: active (via your active model, a local vision model, or a cloud key) or not set up.

Install the memory model

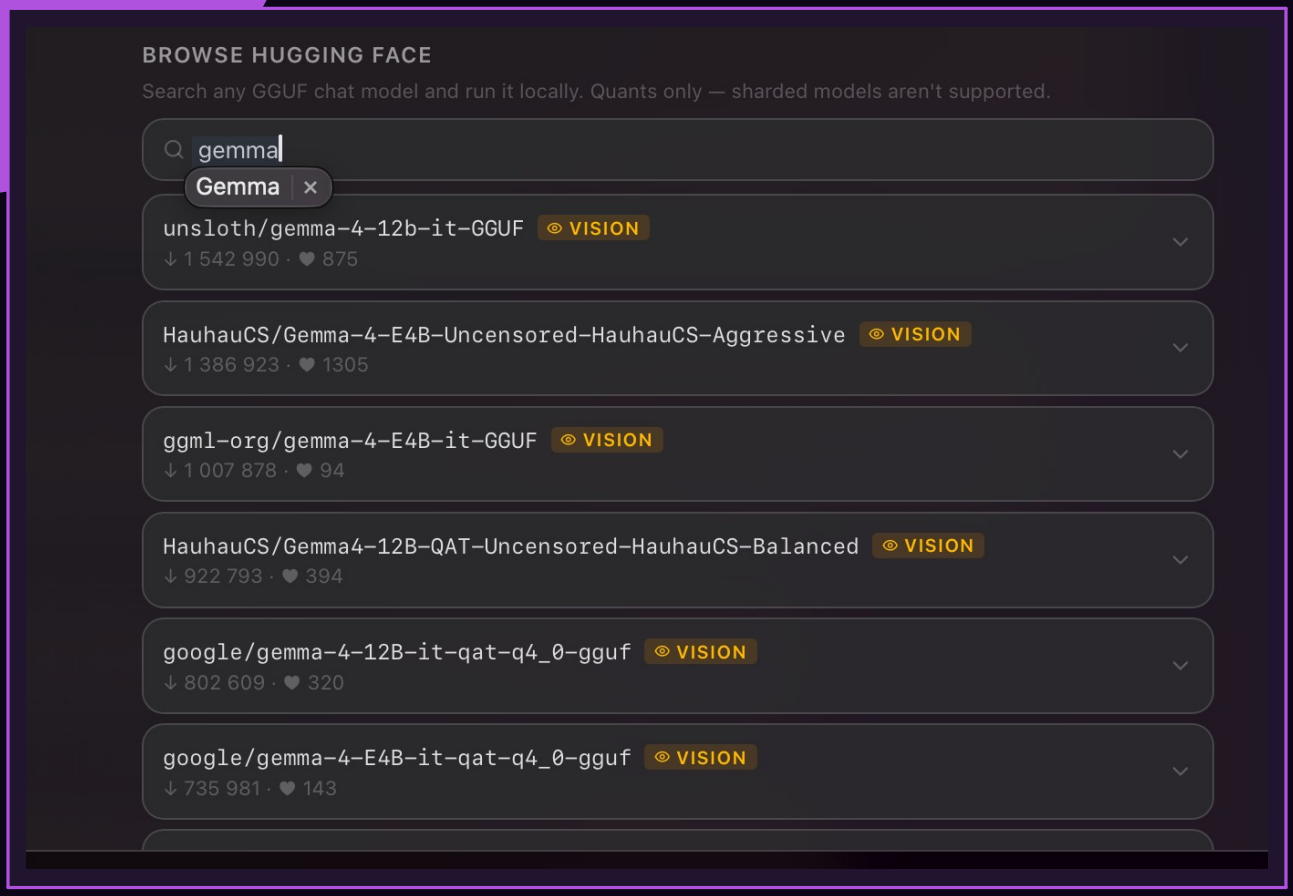
Install EmbeddingGemma when it is marked Required. Keep the active embedding model installed. Changing it starts re-embedding; memory recall may temporarily rely more heavily on keywords until the process completes.

Remove a model

Use the trash icon on a model that is no longer active or required. The app blocks deletion of the active chat model and protected embedding model. Switch roles first, then delete.

FIGURE 11

Browse Hugging Face: search for a repository and choose one compatible GGUF quant.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

28. Install a GGUF Model from Hugging Face

Local Waifu can search Hugging Face for chat models packaged as single GGUF files.

Find and install a model

Goal

Install a compatible third-party GGUF quant without leaving Local Waifu.

Before you start

- Know your available RAM/unified memory and free disk space.
- Prefer model repositories that publish a single file for each quantization.
- Third-party models have their own licenses and behavior.

Steps

1. Open Settings → Models → Chat Models.
2. Under Browse Hugging Face, enter:

- a search term;
- `owner/repository`; or
- a full Hugging Face model URL.

3. Open a result.
4. Select a single runnable quant such as a Q4 variant.
5. Review the download size and memory-fit estimate.
6. Install the quant.
7. When installation finishes, select Use for chat.
8. Test the model.

Expected result

The Hugging Face model appears under installed models and can be selected for chat.

If something goes wrong

- No GGUF models found: use a repository that contains GGUF files.
- No runnable quants (sharded files only): Local Waifu 1.8.0 does not support a GGUF split across several shard files.
- SafeTensors repositories are not installed through the chat-model search panel.
- Couldn't reach Hugging Face: check internet, proxy/VPN, and firewall.
- A fit estimate is guidance, not a guarantee. Reduce context or choose a smaller quant if runtime memory is insufficient.

29. Configure Cloud AI

Cloud AI uses your own provider account or API key. Local Waifu does not convert provider usage into a Local Waifu subscription; the provider may bill your account directly.

Add a provider

Goal

Connect a cloud provider and select a hosted model.

Before you start

- Review the provider's pricing, data handling, and model access.
- Have a valid key or supported connected-account login.
- Turn off Offline Mode only when you intentionally want cloud use. While it is on, providers stay saved but are not used.

Steps

1. Open Settings → Cloud AI. Providers are grouped as Frontier labs, Fast inference, Aggregators, and Asia, with Your own endpoints at the bottom (Chapter 32). Some providers also offer or use a subscription (a connected account such as ChatGPT login).
2. Find the provider. A get a key link opens its key page.
3. Paste the key into Paste your API key and select Save.
4. Select Test. The row reports Key works. or Test failed.
5. Open Settings → Models → Chat Models and find the provider under From your connected accounts. Direct providers are listed live. For very large aggregators, paste the exact model id you want.
6. Select Use on the model you want.
7. Send a short test.

Expected result

Requests for the selected role use the provider. Local character and memory storage remains on the computer. Embeddings stay local even when chat goes to the cloud. Keys live in a local JSON file with owner-only permissions.

Privacy note

The selected prompt, relevant conversation context, image, audio, or text leaves the computer when required by that cloud operation. Do not use a cloud route for information you do not want to send to that provider.

If something goes wrong

- If a model row says Needs API key, add the key in Settings → Cloud AI first.
- If the model list is empty, the key may not give access to chat models. The app says No chat models available for this key.
- Re-copy the key without spaces. To remove one, select Remove.

The cloud is a fallback only if you turn it on

In 1.8.0 a saved cloud key is not a silent backup for the local model. When the local engine is down, or she receives a photo that no local model can see, the message used to go to any cloud key you had saved. Now that happens only if you turn on Use the cloud when the local model can't answer.

Where to find it

Open Settings → Performance → Network, next to Offline Mode. The switch is off by default.

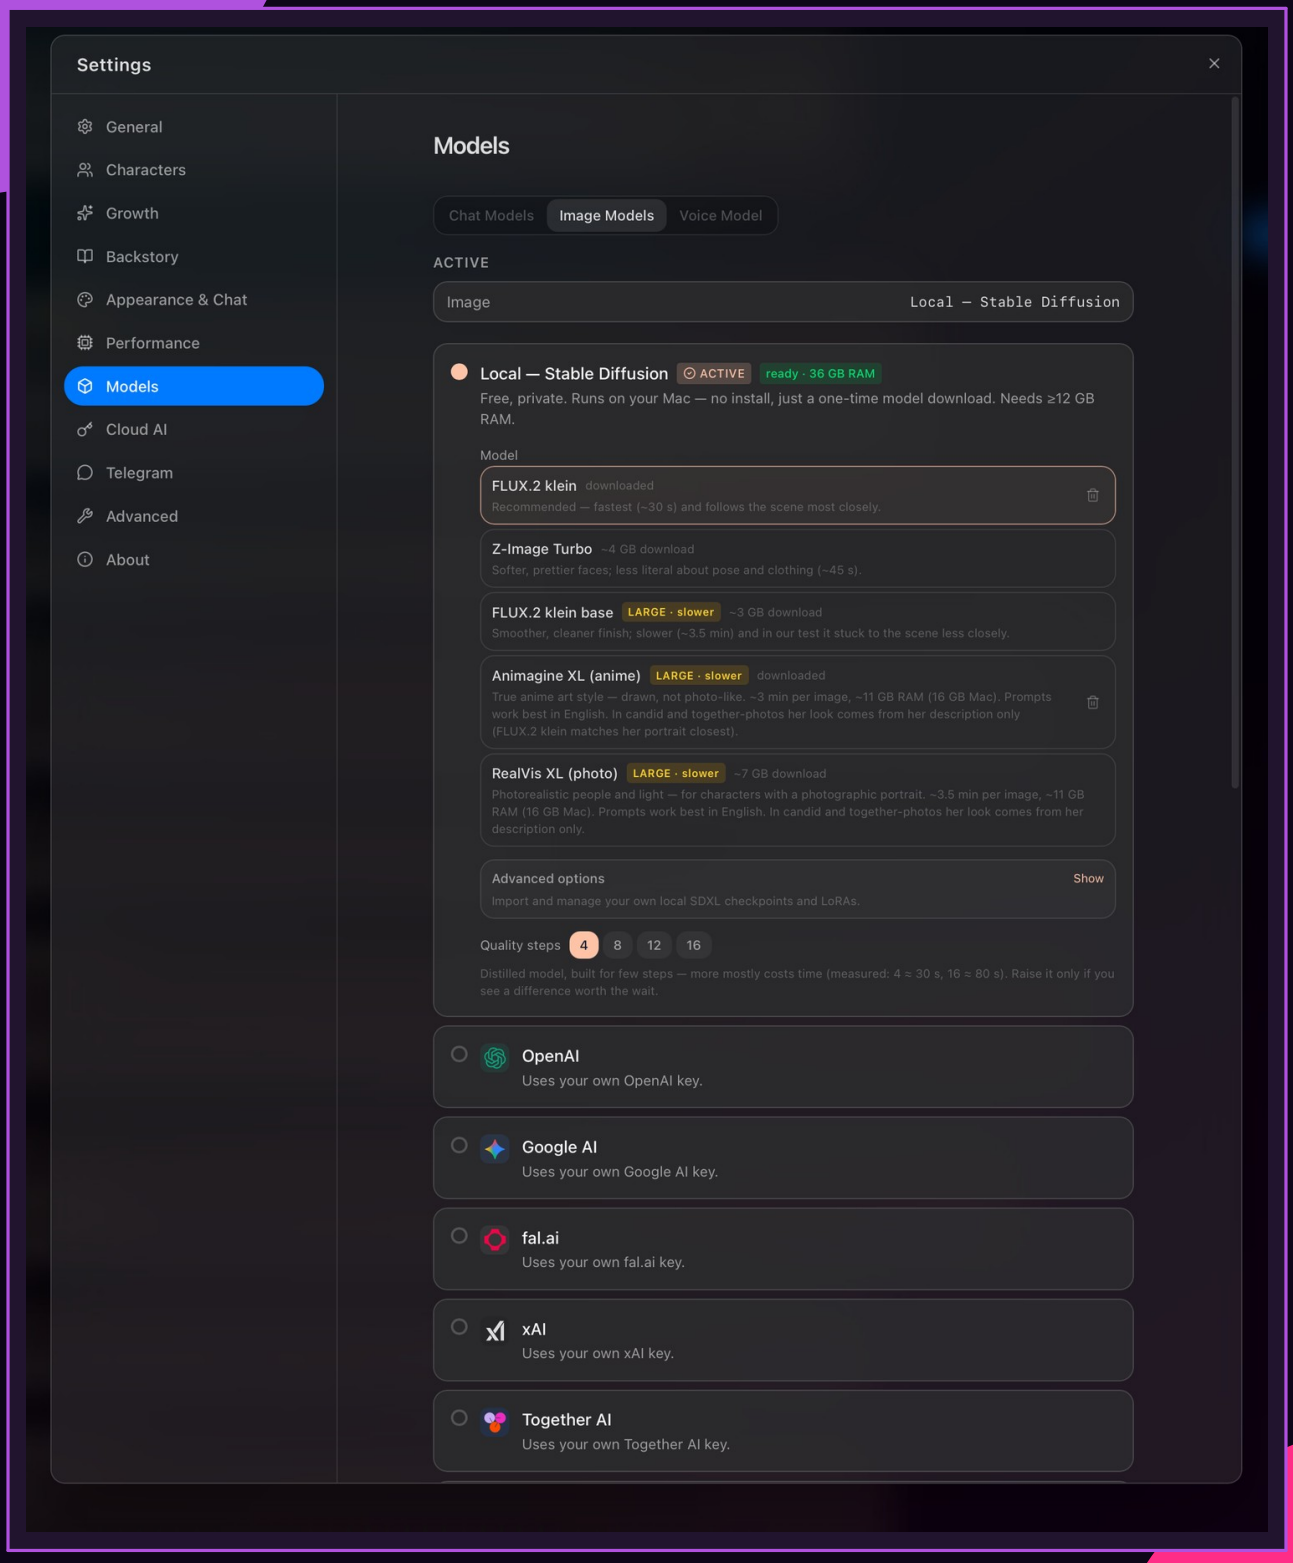
What happens when it is on

The message goes to one of your saved cloud keys, and a note in the conversation tells you it happened (Your local model failed, so this reply came from a cloud provider instead.). The note appears each time, so you always know which replies left your computer.

Offline Mode wins. With Offline Mode on, the fallback never sends anything.

FIGURE 12

Image Models: select local or cloud generation and inspect model requirements.

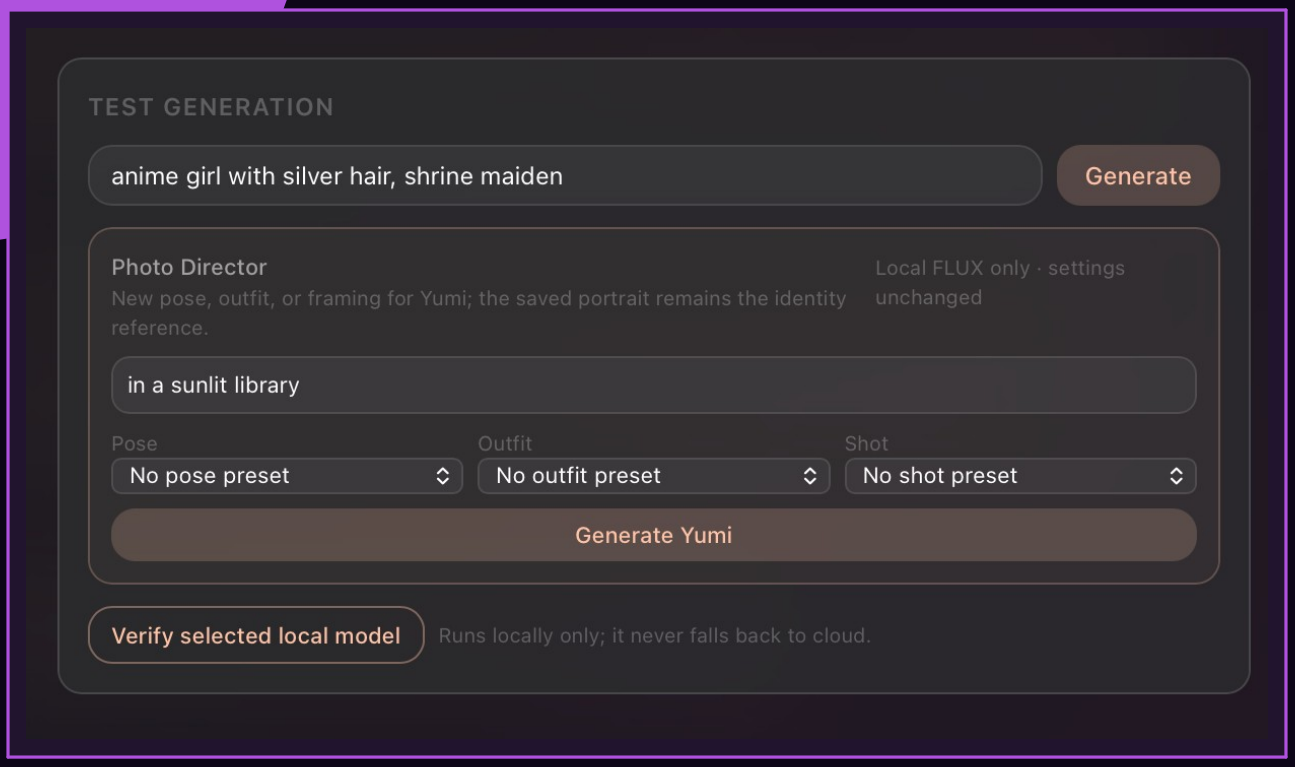


Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

FIGURE 13

Test Generation validates the selected image route before normal use.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

30. Generate Images and Use Gallery

Open Settings → Models → Image Models to choose local or cloud image generation.

Choose a provider and model

Local image generation is private and has no provider fee, but it needs a one-time model download and enough RAM. The UI offers models such as fast FLUX/Z-Image options and larger SDXL models oriented toward anime or photographic styles. The exact catalog can change in later releases.

Cloud providers such as OpenAI, Google AI, fal.ai, xAI, or Together AI require your own provider account/key and send the prompt to that service.

Install a built-in local image model

Goal

Prepare local image generation and verify it works.

Before you start

- Keep at least several gigabytes free.
- Local generation requires approximately 12 GB RAM for the base route; larger SDXL options are intended for 16 GB or more.
- Generation time depends heavily on the computer and model.

Steps

1. Open Settings → Models → Image Models.
2. Select Local: Stable Diffusion.
3. Choose a model and start its download.
4. Wait for downloaded/ready.
5. Start with the recommended quality step count.
6. Under Test Generation, enter a simple English prompt.
7. Select Generate.

Expected result

A test image appears and confirms that the selected local model can render.

If something goes wrong

- Check RAM and free disk space.
- Wait for a model download or another render to finish.
- Reduce quality steps or use a lighter model.
- On Windows, verify image-generation runtime status.

Use Photo Director

Photo Director creates a new pose, outfit, framing, or scene for the active character while using the saved portrait as an identity reference.

It requires:

- an active character;
- a saved portrait;
- the supported ready local FLUX.2 klein model;
- a working local image backend.

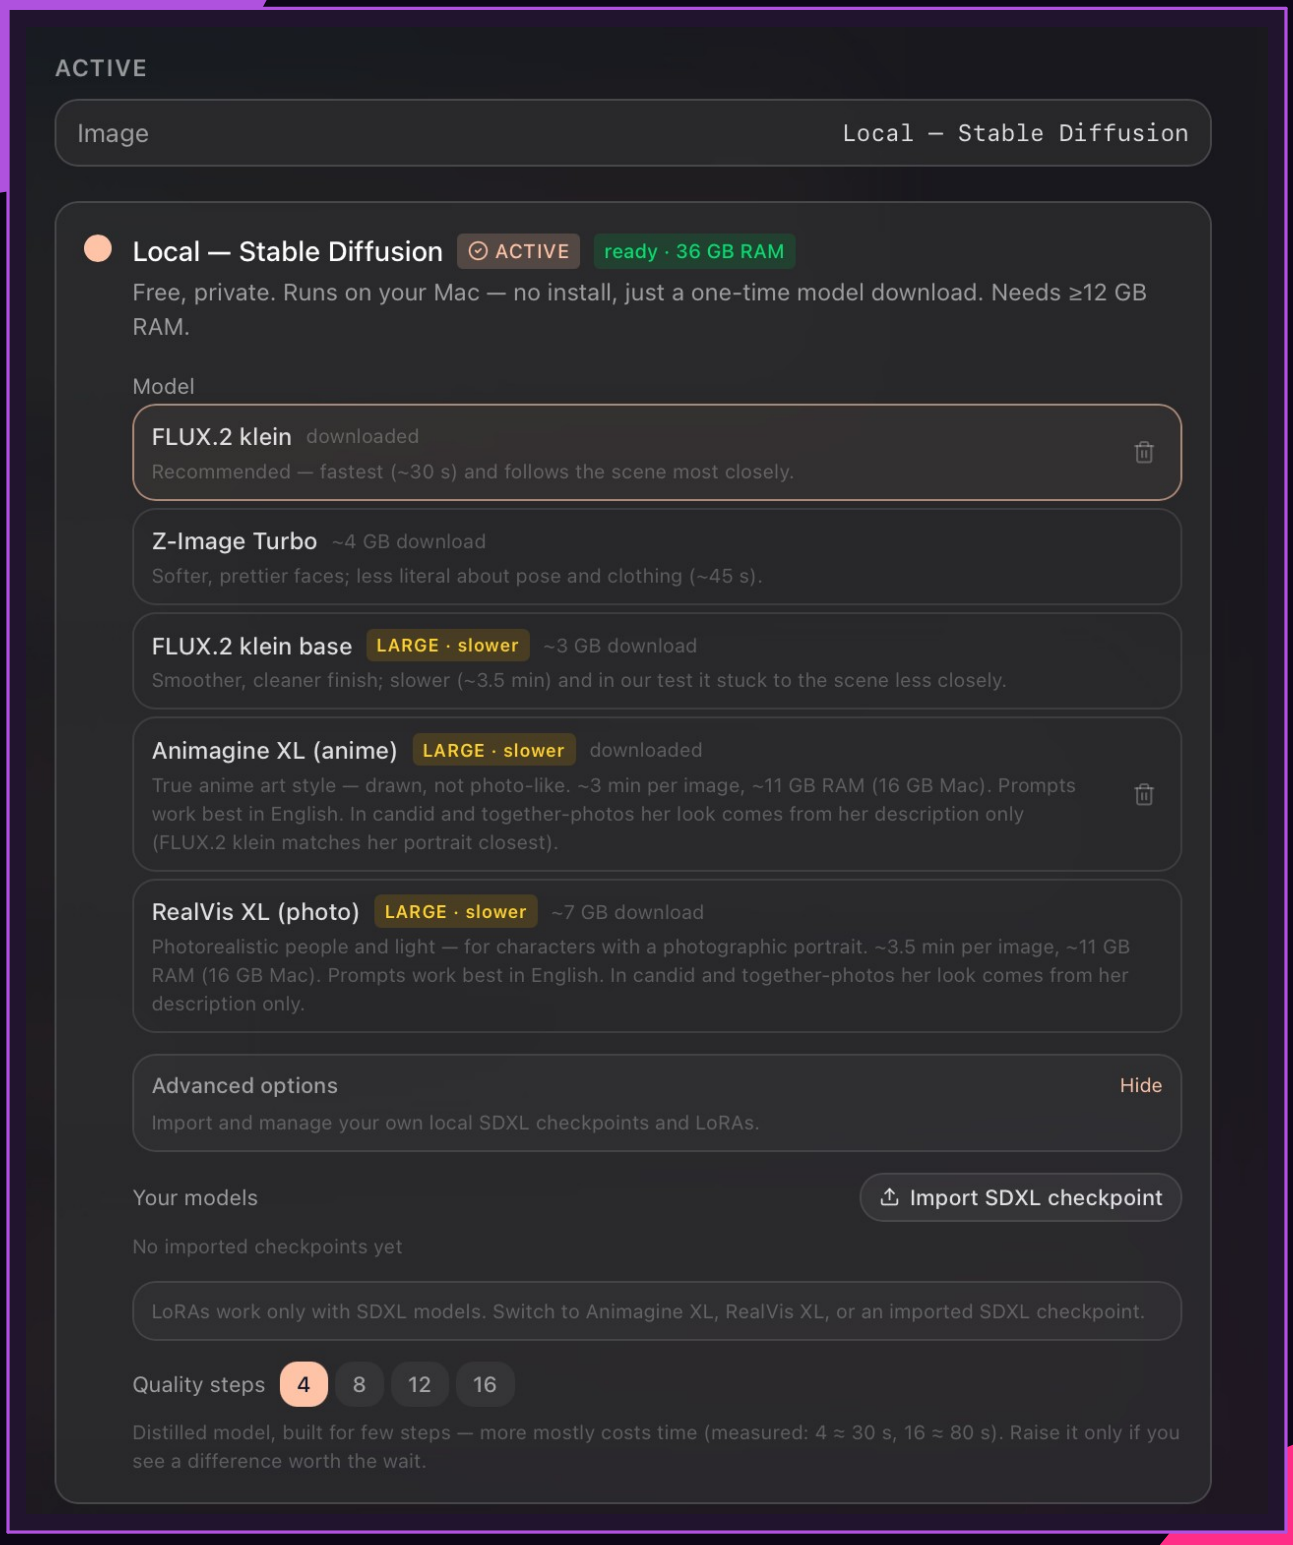
Enter the scene, optionally choose pose/outfit/shot presets, then select Generate [character]. Use Verify selected local model to test the local-only route without silently falling back to cloud.

Gallery

Gallery stores images sent or generated through the companion experience, including 3D photos you take in Photo mode (labelled 3D photo, or 3D photo, made real when the image model redrew them; Chapter 21). Treat it as part of private app data: include it in full backups and remember that character JSON export is not a complete gallery archive.

FIGURE 14

Advanced image options: import an SDXL SafeTensors checkpoint and manage compatible LoRAs.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

31. Import a Custom SDXL Checkpoint or LoRA

Import a custom SDXL checkpoint

Goal

Add a compatible SDXL model to the local image library.

Before you start

The file must be:

- one regular `.safetensors` file;
- at least 512 MiB;
- a valid SafeTensors container;
- an SDXL architecture;
- stored on a disk with enough space for the file plus import overhead.

`.ckpt` is not supported by the 1.8.0 importer.

Steps

1. Open Settings → Models → Image Models.
2. Expand Advanced options.
3. Select Import SDXL checkpoint.
4. Choose the `.safetensors` file.
5. Wait for validation and copying to finish.
6. Select the imported model.
7. Run a test generation before enabling LoRAs.

Expected result

The model appears under Your models and is available as an SDXL base.

If something goes wrong

Errors for a tiny file, malformed/truncated header, no tensors, wrong architecture, checksum mismatch, or insufficient disk space mean the file should not be used. Download a verified SDXL SafeTensors checkpoint from a trusted source.

Import and activate a LoRA

Goal

Apply a compatible SDXL LoRA to an SDXL base model.

Before you start

- Use one `.safetensors` file of at least 1 MiB.
- Select an SDXL base first.

- FLUX and non-SDXL models cannot use this LoRA path.
- At most three LoRAs can be active on one model.

Steps

1. Select the target SDXL model.
2. Open its LoRA controls.
3. Import the LoRA `.safetensors` file.
4. Add up to eight trigger words if needed.
5. Enable the LoRA.
6. Start near the default weight `0.8`.
7. Adjust between `0.0` and `1.5` in small steps.
8. Run a test render.

Expected result

The LoRA affects renders from the selected SDXL model.

If something goes wrong

- A family marked unknown may be imported but cannot be activated safely.
- Disable another LoRA if three are already active.
- Lower the weight if the image is overcooked or distorted.
- Header metadata cannot guarantee compatibility; the final check is a successful test render.

32. Use Your Own Server as a Cloud Provider

Cloud providers in Local Waifu are not limited to the named ones. If your server speaks the OpenAI-compatible chat API, you can add it and use it exactly like a hosted provider, with your own key or none at all.

This is the route for LM Studio, vLLM, llama.cpp server, Ollama exposed over HTTP, and anything else that implements `/v1/chat/completions`.

Add an endpoint

Goal

Register your own server as a usable chat option.

Before you start

- Have the server running and reachable from this computer.
- Know its base URL, including the `/v1` part. For LM Studio on the same machine that is `http://localhost:1234/v1`.
- Most local servers need no API key. Leave the key blank if so.

Steps

1. Open Settings → Cloud AI.
2. Scroll to Your own endpoints.
3. Select Add an endpoint.
4. Fill in the three fields:

- Name: whatever you will recognise later, for example `LM Studio`.
- Base URL: the address including the version path, for example `http://localhost:1234/v1` or `https://api.example.com/v1`.
- API key: leave it blank if the server does not need one. On an edit, leaving it blank keeps the saved key.

5. Select Save.
6. Select Test on the new row.
7. Select the endpoint's model in the model picker for chat, as you would for any provider.

Expected result

The row reports that the test succeeded, the endpoint appears in the chat model list, and replies come from your server.

If something goes wrong

- If Test fails, open the base URL in a browser on the same machine. If that does not answer, the problem is the server or its port, not Local Waifu.
- If the server answers in a browser but not here, check that the URL ends with the version path your server expects, usually `/v1`.

- If the model list comes back empty, the server is not publishing an OpenAI-compatible `/v1/models`. Type the model name in by hand where the picker allows it.
- Use Edit to correct the URL and Remove to delete the row.

Security: what the URL is allowed to be

This is the part worth reading before pasting an address.

- Plain `http://` is allowed only for the machine itself (`localhost`, `127.0.0.1`) and for private network addresses such as `192.168.x.x`, `10.x.x.x`, and `.local` names.
- Anything on the public internet must use `https://`. There is no setting that relaxes this, because a key sent over plain HTTP is a key given away.
- A remote Ollama Host is a separate setting with its own rules. See Chapter 36.

If something goes wrong

If an address is rejected, it is almost always plain `http://` pointed at a public hostname. Either move the server onto your own network or put TLS in front of it.

Offline Mode and your own endpoint

Offline Mode blocks cloud providers. It does not block an endpoint on your own machine or your own network, because that request never leaves your network. Offline Mode means off the internet, not off your own hardware.

If something goes wrong

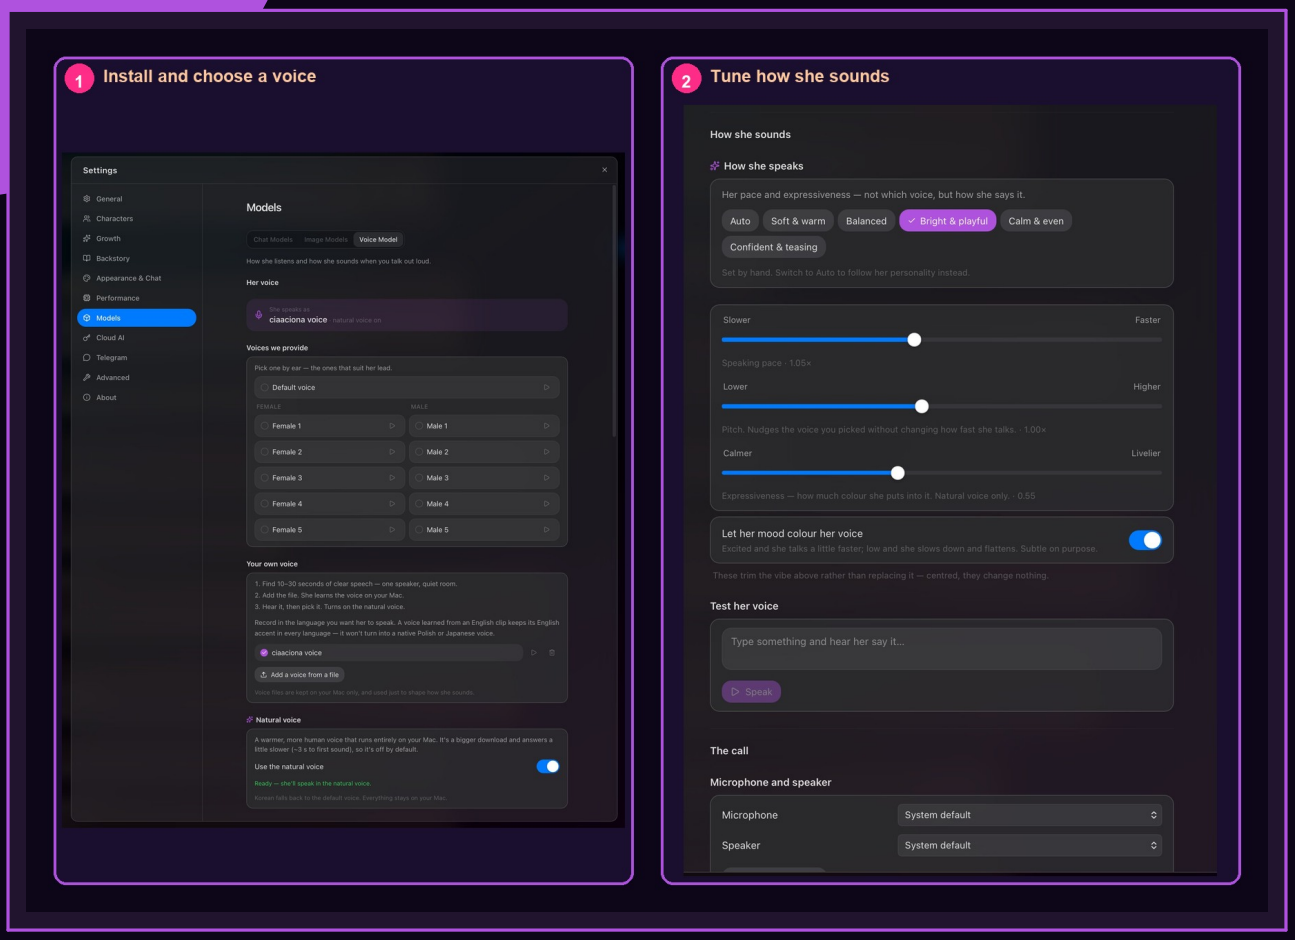
If chat is refused while Offline Mode is on, check whether the endpoint resolves to a public address. A hostname that looks private but resolves outward is treated as a cloud provider.

What your own endpoint does not get

- It does not become the local engine. Embeddings and memory stay on your computer, as described in Chapter 38.
- It does not carry vision unless the server implements it and you mark the provider as vision-capable.
- It does not survive a factory reset of the app's settings. Endpoints are stored with the rest of your configuration, which Reset All clears, and which a full backup preserves.

FIGURE 15

Voice setup overview: choose a voice, adjust its delivery, and review the active character's voice settings.

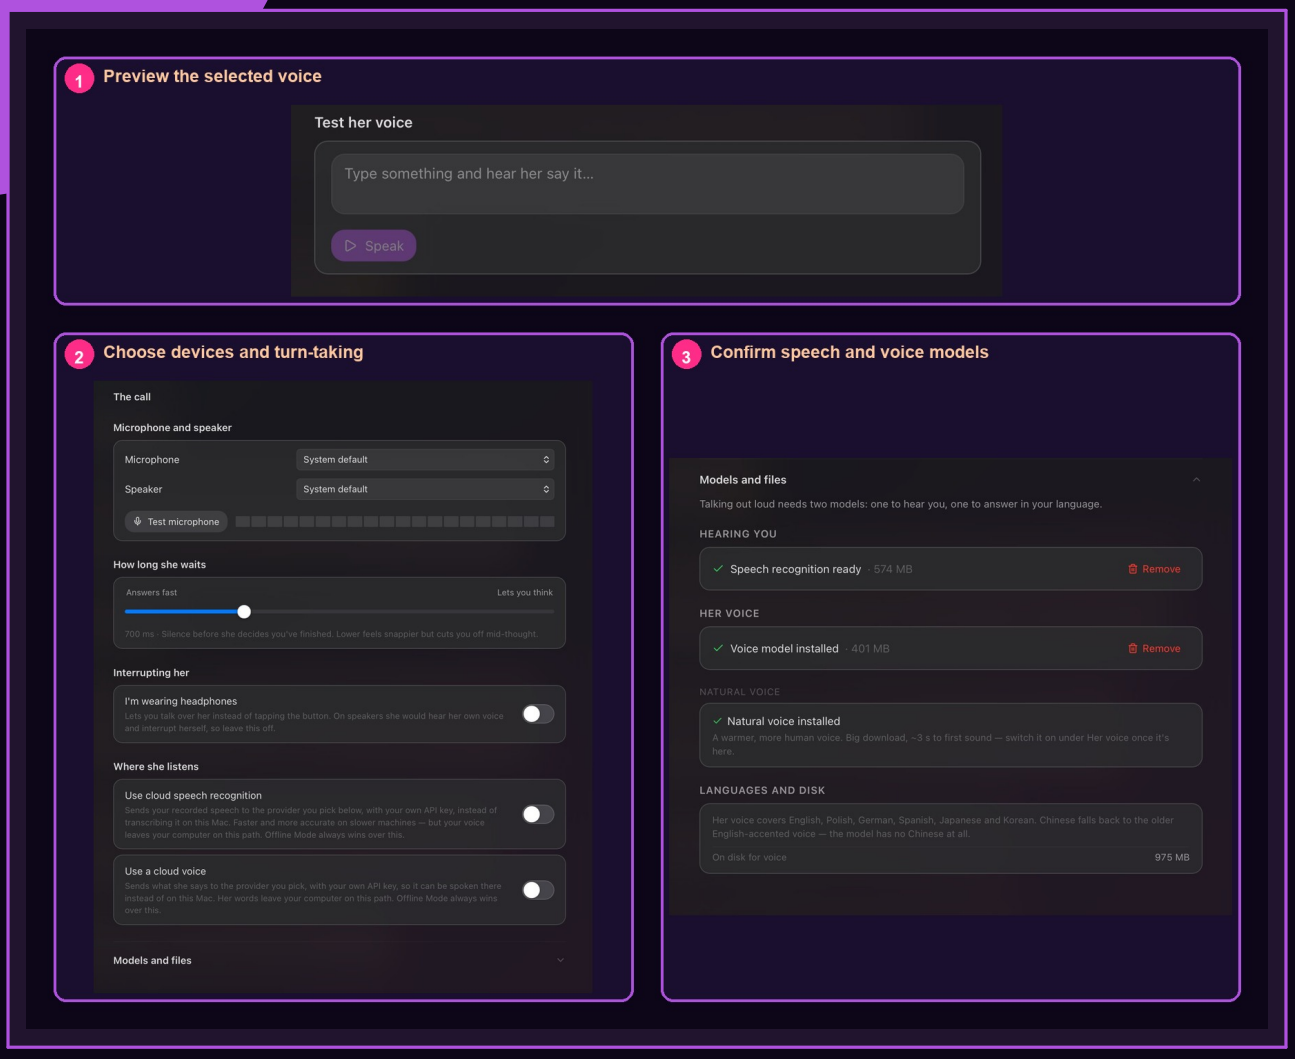


Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

FIGURE 16

Call setup overview: preview the voice, select audio devices, tune turn-taking, and confirm required model packages.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

33. Configure Voice, Dictation, and Calls

Local Waifu has three separate voice functions:

1. Dictation: speech becomes editable text in the composer; it is not sent automatically.
2. Voice output: the companion reads a response aloud.
3. Call: a continuous speech-recognition → chat → speech loop.

Four more voice features, covered in Chapter 34, are voice messages, hands-free start, bedtime stories, and songs.

Install speech and voice models

A full call needs:

- speech recognition (STT/Whisper) to hear you;
- a voice model (TTS/Supertonic or another configured route) to speak.

Missing STT blocks the full call. Missing high-quality TTS may fall back to an older voice, but pronunciation can be worse.

Open Settings → Models → Voice Model, install Speech recognition and Voice model (Set up voice installs both), and optionally install Natural Voice. If the speech model is English-only, the app tells you to remove it and download the multilingual one. The screen reports current package sizes and disk use.

Choose a voice

1. Open Settings → Models → Voice Model.
2. Preview the default, female, or male voices with the play button.
3. Select a voice for the active character.
4. Turn on Natural voice for the warmer, more human local voice when installed. It is a bigger download and answers a little slower (about 3 seconds to first sound), so it is off by default. On macOS it runs on Apple Silicon in the Direct edition only and needs 16 GB of memory. The tab says when your Mac cannot run it.
5. Use Test her voice, type a sentence, and select Speak.

Language support depends on the voice engine and the platform.

- Both platforms: the standard voice covers English, Polish, German, Spanish, Japanese, and Korean.
- Windows: Chinese needs Natural voice turned on and ready. Natural voice also supports Korean.
- macOS: the voice model has no Chinese at all, so Chinese falls back to the older English-accented voice. With Natural voice, Korean falls back to the default voice.

Under Voice language you choose the language she speaks, even when the app is in another one. It does not change how she understands you.

Add your own voice sample

Goal

Create a character-specific voice reference.

Before you start

- Record 10 to 30 seconds of clean speech.
- Use one speaker in a quiet room.
- Record in the language/accent you want to preserve.
- File size must be under 25 MB.

Steps

1. Open Settings → Models → Voice Model.
2. Under Your own voice, select Add a voice from a file.
3. Choose the audio file.
4. Enter a short name, up to 60 characters.
5. Wait for local processing and transcription.
6. Preview the result and select it.

Expected result

The active character uses the new reference and Natural Voice is enabled when ready.

If something goes wrong

Use a shorter, cleaner sample. A learned English sample keeps its English accent in other languages; record again in the target language for better pronunciation.

Tune how she speaks

Available style presets include Auto, Soft & warm, Balanced, Bright & playful, Calm & even, and Confident & teasing.

Adjust:

- speaking pace;
- pitch;
- expressiveness;
- Let her mood colour her voice.

Use small changes and test after each one. Auto follows the character personality; manual presets keep your chosen delivery style.

Configure a call

Goal

Select audio devices and tune turn-taking.

Steps

1. Open Settings → Models → Voice Model and find The call.
2. Select Microphone and Speaker, or keep System default.

3. Select Test microphone and verify activity.
4. Adjust How long she waits. A shorter value feels faster but can cut you off mid-thought.
5. Enable I'm wearing headphones only when appropriate; speakers can feed her voice back into the microphone and cause self-interruption.
6. Leave cloud recognition and cloud voice off for fully local processing, or enable them intentionally with your provider key.
7. Start a call with the phone button (Talk to her) or `Command/Ctrl+Shift+C`.
8. Optionally turn on Hands-free to start a call by saying her name. Chapter 34 covers it.

Expected result

Your speech is transcribed, answered by the active model, and spoken through the selected voice and speaker.

If something goes wrong

- Grant microphone permission in the operating system.
- Choose the correct input device.
- Install STT if the call cannot hear you.
- Install/configure TTS if text replies appear but no voice plays.
- Increase the waiting time if you are interrupted too early.

34. Voice Messages, Hands-Free, Stories, and Songs

Chapter 33 covers dictation, spoken replies, and live calls. This chapter covers four more things she does with sound.

Send a voice message

Goal

Record a short message and send it as audio.

Before you start

Speech recognition must be installed, because the message is transcribed. Chapter 33 covers installing it.

Steps

1. In a one-to-one chat, press and hold the voice message button next to the message box. The tooltip says Record a voice message.
2. Talk while you hold it. The box shows Recording.
3. Release to send. To cancel, slide up before releasing or press `Escape`.

Expected result

Your message appears as an audio bubble with a transcript, and she answers.

If something goes wrong

- A press that is too short is not sent. Hold the button while you talk.
- A message can be two minutes at most. If you hit the limit, the end is not recorded.
- If a reply was still coming in, the message is not sent.
- If the microphone disconnects, the message stops.
- In a group chat voice messages do not work. Type instead.
- If the transcript cannot be made, the bubble shows No transcript.

Start a call hands-free

Goal

Open a voice call by saying her name, without touching anything.

Steps

1. Open Settings → Models → Voice Model and find Hands-free under The call.
2. Turn it on.
3. Check the Wake phrase. It is her name unless you change it. A phrase that is too short to tell apart from ordinary speech is refused, so use a longer one. Three letters are the minimum, or two characters for Chinese, Japanese, and Korean.
4. Choose whether to Require "hey" before her name.

5. If the tab offers it, download the small wake-word model.
6. Say her name.

Expected result

The status shows Listening, and saying the phrase starts a call.

Privacy

The setting says it listens on your computer only: the audio is never saved and never leaves your computer, and the microphone is open only while it listens. It is off by default. The status reads Paused while the microphone is in use during a call.

If something goes wrong

If listening cannot start, the app shows Couldn't listen with the reason and leaves the switch off.

Ask for a bedtime story

Goal

Have her tell a story out loud that stops by itself when you fall asleep.

Before you start

She tells the story in her voice, so install the voice model first. A story needs an active license and cannot start during a voice call.

Steps

1. In the message area, open the story and song button and choose Bedtime story.
2. Under What should it be about?, type a theme, or leave it empty.
3. Choose a length: Short (about 3 minutes), Medium (about 8), or Long (about 15).
4. Choose a Sleep timer or leave it off.
5. Select Start the story.

Expected result

She tells the story in parts, and the panel shows Part 1 of N and the minutes left. She ends it with The story is over. Sleep well. Select Stop to end it early. You can set a default length and sleep timer under Settings → Models → Voice Model → Bedtime story.

If something goes wrong

- End the voice call first, then start a story.
- A story is already being told.
- While a story plays, dictation and voice messages are disabled.
- A story is not saved as memories beyond a single line.

Ask for a song

Goal

Have her write and sing a short song for you.

Before you start

Songs are made by a cloud service, so they need a fal.ai key. Add one under Settings → Models → Image Models. The same key is used for image generation. Songs also need an active license, and they do not work in Offline Mode.

Steps

1. Open the story and song button and choose A song from her.
2. Type a theme under What should it be about? and a Style, or leave them empty.
3. Read the consent line: when you press Create, the lyrics, the style, and the theme are sent to fal.ai. Nothing is sent before that. Your chat model writes the lyrics first, so if it is a cloud model the theme and a few memories go there too.
4. Select Create the song.

Expected result

She writes lyrics, then composes the song. It arrives as an audio bubble in the chat, with a Show lyrics toggle.

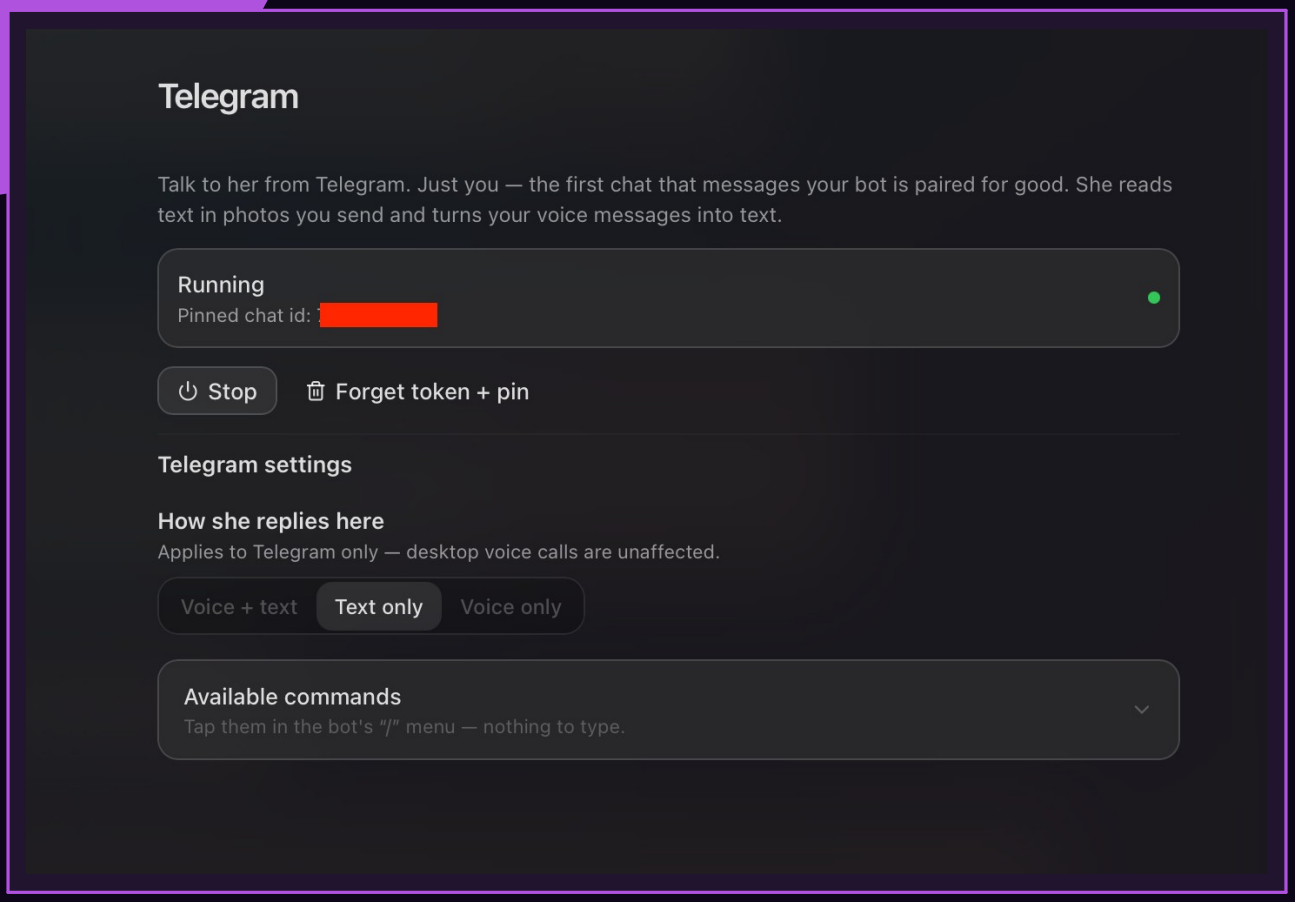
If something goes wrong

- Songs need a fal.ai key. Select Open Image Models to add one.
- Offline Mode is on. Turn it off first.
- The music endpoint id in Settings is not valid. In Settings → Models → Voice Model → Songs, select Reset to default under Music endpoint.
- You can also change Song length there.
- Songs cannot start during a call, and only one song is made at a time.

Part IV: Integrations, Data, and Support

FIGURE 17

Telegram integration: save a BotFather token, pair one chat, and start the bot. Secret values are redacted.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

35. Use Local Waifu from Telegram

Telegram lets you talk to the same active companion from your phone. The desktop app and Telegram share the active character, conversations, model settings, and memory. Telegram has its own reply-mode setting.

> Security warning: Treat the Telegram bot token like a password. Do not paste it into screenshots, public bug reports, or chat messages. The first Telegram chat that messages the bot becomes the pinned owner. Messages and commands from other chat IDs are rejected.

Connect a Telegram bot

Goal

Create a private Telegram bot, verify its token, start the Local Waifu bridge, and pair your Telegram account.

Before you start

- Local Waifu must be open on your computer and connected to the internet.
- Install and configure the chat model you want to use.
- For Telegram voice messages, configure speech recognition under Settings → Models → Voice Model.
- For spoken replies, install and configure a voice model.
- You need a Telegram account.

Steps

1. In Telegram, open the verified bot named `@BotFather`.
2. Send `/newbot` and follow BotFather's instructions to choose a name and username.
3. Copy the token issued by BotFather. A token resembles `123456:ABC-DEF...`.
4. In Local Waifu, open Settings → Telegram.
5. Paste the token into Bot token.
6. Select Test. This verifies the token with Telegram without saving it or starting the bot.
7. Confirm that Local Waifu reports Token works: bot is @....
8. Select Start. The status should change to Running.
9. Open your new bot in Telegram and send `/start` or a normal message from the account you want to pair.
10. Return to Settings → Telegram and confirm that Pinned chat id is displayed instead of No chat pinned yet.

Expected result

The bot is Running, your Telegram chat is pinned as its owner, and messages sent to the bot reach the same Local Waifu companion used by the desktop app.

If something goes wrong

- If Test says Telegram rejected the token, copy a fresh token from BotFather and make sure no characters or spaces are missing.

- If Test succeeds but the bot does not answer, select Start; Test does not save the token or start polling.
- If the wrong Telegram account was paired, use Forget token + pin, start the bot again, and send the first message from the correct account.
- If the bot is running but chat fails, check Settings → Performance → Runtime status, the active model, internet access, and Offline Mode.

Choose how Telegram replies

Goal

Choose whether Telegram receives text, voice, or both without changing desktop voice calls.

Before you start

The Telegram bot must show Running. A ready voice model is required for reliable voice replies.

Steps

1. Open Settings → Telegram.
2. Under Telegram settings, find How she replies here.
3. Choose one mode:

- Voice + text: sends both forms; this is the default.
- Text only: sends text without a voice note.
- Voice only: sends a voice note instead of text.

4. Send a short test message to the bot.

Expected result

The bot replies in the selected Telegram-only format. Desktop voice-call behavior is unchanged.

If something goes wrong

- In Voice only, Local Waifu deliberately falls back to text if the voice model is unavailable, so the bot does not remain silent.
- Voice notes sent to the bot are always transcribed. Configure the transcription provider in Settings → Models → Voice Model.
- If speech is in the wrong language or sounds poor, check the configured voice language and installed voice model.

Use photos, voice messages, and bot commands

Goal

Use the bot's supported media and discover the commands available in your installed build.

Before you start

For image understanding, a vision-capable route must be ready. For incoming voice notes, speech recognition must be ready.

Steps

1. Send normal text, a photo containing text, or a Telegram voice message to the bot.
2. In Local Waifu, open Settings → Telegram → Available commands.
3. Use Telegram's `/` menu to choose a command. The list is generated by the app, so it is the authoritative list for your build.
4. Common commands include `/help`, `/ping`, `/status`, `/response`, `/think`, `/natural`, `/transcribe`, `/mute`, `/model`, `/voices`, `/lang`, `/imgmodel`, `/mood`, `/newchat`, `/reset`, `/image <prompt>`, `/selfie`, `/last`, `/refresh`, and `/research <topic>` (the same night-research queue as Chapter 25). `/refresh` writes her last reply again and sends the new picture, like any other reply.
5. For dynamic lists such as models or voices, select the numbered command presented by the bot.

Expected result

Text is handled as chat, text in photos can be read, voice messages are transcribed, and commands open the corresponding choices or actions.

If something goes wrong

- If a photo is ignored, verify that a vision-capable model or provider is available.
- If a voice note cannot be understood, check speech recognition under Settings → Models → Voice Model.
- Model changes made through Telegram also affect the desktop app because both interfaces share the active model.
- `/reset` asks for explicit confirmation. Read the bot's prompt before confirming.
- An unknown slash command may be treated as normal conversation text.

Stop the bot or remove its credentials

Goal

Pause Telegram access or completely remove the saved token and owner pairing.

Before you start

Understand the difference:

- Stop pauses the bridge but keeps the token and pinned owner.
- Forget token + pin stops the bridge and removes both the saved token and pairing.

Steps

1. Open Settings → Telegram.
2. To pause the bot, select Stop.
3. To remove access completely, select Forget token + pin.
4. Select Yes, forget to confirm.

Expected result

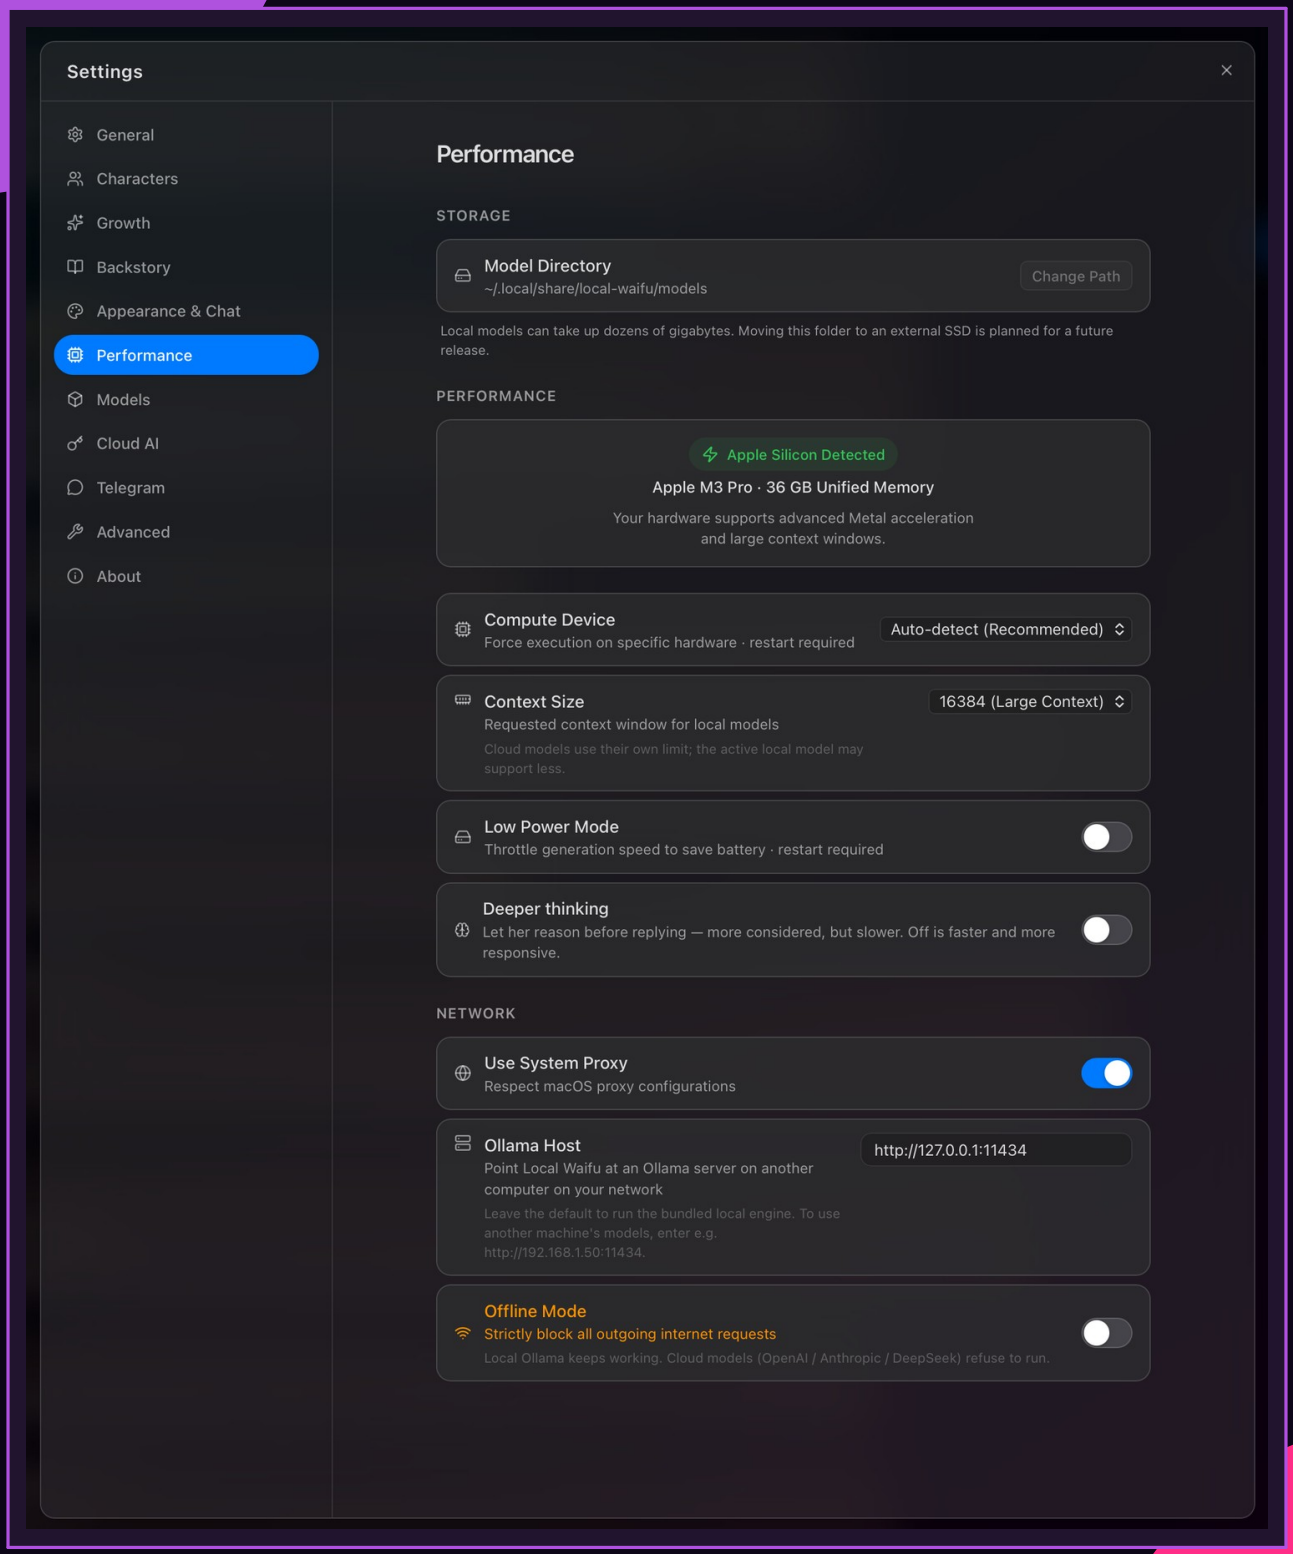
After Stop, the status is Stopped and the same configuration can be started again. After Forget token + pin, the token and owner pairing are removed, allowing a new token or owner to be configured.

If something goes wrong

Forget token + pin remains available even while the bot is stopped. Use it if a stale token or old owner pairing prevents a new setup.

FIGURE 18

Performance settings: context size, deeper thinking, model routing, and runtime behavior.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

36. Tune Performance, Hardware, and Network

Open Settings → Performance (the tab used to be called Hardware) to inspect storage, check what can run right now, choose compute behavior, set the local context size, reduce resource use, and control network routing.

Choose a compute device and verify acceleration

Goal

Use the best available hardware and confirm whether the active local model is running on the GPU or CPU.

Before you start

- Send at least one message so the local model is loaded before checking acceleration.
- A GPU is optional on Windows; Local Waifu can fall back to CPU.
- macOS builds require Apple Silicon and use unified memory.

Steps: macOS

1. Open Settings → Performance.
2. Under Compute Device, choose:

- Auto-detect (Recommended): lets the local engine use Metal acceleration on Apple Silicon.
- CPU Only: forces CPU-only inference. In 1.8.0 this really runs the model on the CPU. An earlier version ignored the setting.

3. Restart Local Waifu when prompted.

Steps: Windows

1. Open Settings → Performance.
2. Leave Compute Device on Auto-detect (Recommended) unless troubleshooting.
3. Send a message to load the model.
4. Under Local AI acceleration, select Check.
5. Read the result:

- Running on GPU: the model is in VRAM.
- Split CPU/GPU: only part of the model fits in VRAM.
- Running on CPU: the GPU is not engaged.
- No model loaded yet: send a message, then check again.

6. If NVIDIA or AMD hardware is detected but not fully used, select Retry GPU setup. The app may download a matching runtime and restart the engine.

AMD note: current Windows builds read dedicated AMD VRAM for the model-tier recommendation and report ROCm/Vulkan readiness in Settings → Performance. ROCm v7/HIP7 is preferred where the driver exposes it; Ollama's Vulkan runner is the supported fallback on cards such as some RX 6000

systems. A 16 GB card alone still maps to Light; if Qwen 3.5 9B fits your VRAM, select it manually in Settings → Models and confirm acceleration with Check. Lower Context Size if the result is Split CPU/GPU.

Expected result

macOS uses Metal unless CPU Only was selected. Windows reports the real placement of the loaded model and, where available, the percentage held in VRAM.

If something goes wrong

- Restart the app after changing Compute Device.
- Try a smaller or more heavily quantized model if Windows reports Split CPU/GPU.
- Update GPU drivers, then use Retry GPU setup again.
- CPU fallback is functional but can be much slower.
- There is no separate Apple Neural Engine inference path in the macOS 1.8.0 build. A stored Metal choice is shown as Auto-detect, because the two were the same setting.

Check runtime readiness

Goal

Confirm which local features can run now.

Before you start

This panel is on both platforms. It shows what can run right now, which saves guessing when a feature is silent.

Steps

1. Open Settings → Performance.
2. Under Runtime status, select Check.
3. Review Local AI engine, Memory model, Chat model, Voice output, Sees images, and Image generation. Each reads Ready or Not available.
4. For an item marked Not available, open the relevant model section and install or select the missing component.

Expected result

Each feature you intend to use is marked Ready.

If something goes wrong

Restart Local Waifu after installing a runtime or model, then run Check again. If the local engine remains unavailable, verify Ollama Host and review App Logs as described in Chapter 41.

Set the context size and thinking mode

Goal

Balance conversation context, memory use, and response speed for local chat models.

Before you start

Larger context windows require more RAM or unified memory and may substantially slow replies. Cloud providers use their own limits.

Steps

1. Open Settings → Performance and find Context Size.
2. Start with 8192 (Standard).
3. Increase to 16384 (Large Context) or 32768 (Max Context) only if your model and available memory can support it. A stored 4096 legacy value is automatically protected to 8192.
4. Use Custom only when you know the model's context limit. Both platforms accept 8,192 to 262,144 tokens, although the active model may support less.
5. Turn on Deeper thinking only when you prefer more considered but slower responses.
6. Optionally turn on Suggest replies to show three tappable replies under her newest message. It is off by default because it costs output tokens on every turn.

Expected result

Local replies use the requested context size within the active model's capabilities. In 1.8.0 local chat uses at least 8K and can grow to 16K to keep the character prompt, except on machines where long prompts are already slow. Deeper thinking trades speed for more reasoning.

If something goes wrong

- Reduce Context Size if responses become slow, memory pressure rises, or the model fails to load.
- A high value is a request, not a guarantee; the active model may support less.
- Turn off Deeper thinking for faster, more responsive chat.

Reduce background resource use

Goal

Make Local Waifu release model memory sooner and perform fewer background wake-ups.

Before you start

Low Power Mode can make the next reply slower because the local model is unloaded sooner between chats.

Steps

1. Open Settings → Performance.
2. Turn on Low Power Mode.
3. Restart Local Waifu.

Expected result

Background work occurs less frequently and the local model leaves memory sooner after inactivity.

If something goes wrong

If first-response latency becomes inconvenient, turn off Low Power Mode and restart the app.

Configure proxy, offline mode, or a remote Ollama server

Goal

Control network behavior without disabling local AI unnecessarily.

Before you start

- Use System Proxy is enabled by default and changes require a restart.
- Offline Mode blocks cloud providers but does not block local Ollama.
- A remote Ollama Host exposes prompts to that other computer and to the intervening network; use only a host you trust.

Steps

1. Open Settings → Performance → Network.
2. Leave Use System Proxy enabled if your network depends on operating-system proxy settings. Restart after changing it.
3. Turn on Offline Mode to prevent cloud chat providers from running.
4. To use Ollama on another computer, enter its URL in Ollama Host, for example `http://192.168.1.50:11434`.
5. To return to the bundled local engine, select Reset beside Ollama Host.
6. Decide on Use the cloud when the local model can't answer. It is off by default. When it is on and the local engine is down, or she receives a photo and no local model can see, the message goes to one of your saved cloud keys, and you see a note when that happens. Leave it off to keep everything local. Chapter 29 explains it.

Expected result

Local Ollama continues to work in Offline Mode. Cloud models refuse to run. A valid custom Ollama Host routes local-model requests to the specified server.

If something goes wrong

- If cloud chat reports that it is blocked, turn off Offline Mode only if sending data to the configured provider is intentional.
- If local chat goes offline after editing Ollama Host, select Reset.
- If downloads fail behind a corporate proxy, enable Use System Proxy and restart.
- Check firewall access to the remote Ollama host and port; do not expose Ollama directly to the public internet.

Storage note: Local models can occupy tens of gigabytes. The Change Path control is disabled in the current build. Do not rely on it for moving models; keep the default model folder and plan free disk space before downloading a large model.

37. Edit Her Instructions in the Developer Tab

Before every reply she reads a fixed set of instructions: who she is, how she writes, how she remembers, how she sends photos. Settings → Developer lets you read those instructions and rewrite them. Changes apply from the next message, in the desktop app and in Telegram.

This is an advanced tab. The warning at the top says so: Advanced: you are editing her system prompt. With nothing edited, the prompt is exactly what the app ships. Every section can be restored at any time, so you cannot lose the original.

Risk levels

Each prompt section carries a badge.

Badge	Meaning
Safe	Tone and style only. Worst case: she sounds different
Caution	Changes how she behaves or reads the conversation. Worst case: she feels off or drifts out of character
Dangerous	Teaches markers the app reads (memory, photos, modes, thoughts) or a safety rule. A bad edit breaks the feature silently: memories stop saving, photos stop arriving

Start with Safe sections. Leave Dangerous sections alone unless you know what the markers do.

Edit a section

Goal

Change one part of her instructions and check the result.

Steps

1. Open Settings → Developer.
2. Under Prompt sections, select a section. The list includes Identity line, Core drive, Before replying, How she writes (full), How she writes (compact), Voice call rules, Memory tools, Photo skills (full), Photo skills (compact), Photo of you two, Private thoughts, What is worth remembering, Voice examples intro, Boundaries reminder, Mode switching, Your extra instructions, and Closing line.
3. Change the text.
4. For a Dangerous section, tick I understand this can break memories, photos or other features, and I can restore the default. The app will not save otherwise.
5. Select Save.
6. Select Preview full prompt to see the prompt as she would get it, without per-message data such as recalled memories.
7. Send a message and check how she answers.

Expected result

The section shows Edited. The tab counts Edited sections at the top. Emptying a section turns it off, and the section shows Turned off.

If something goes wrong

- Discard changes drops what you typed. Restore default restores that one section.
- Show changes vs default shows the differences line by line.
- A section that teaches a marker such as `[[MEM ...]]` or `[[SELFIE ...]]` breaks that feature if you change the marker syntax or empty it.

Add standing instructions

Two places take your own text.

- Your extra instructions is one block added just before the closing line. Use it for standing requests such as "Always call me Captain."
- Your blocks under Add block lets you add as many named instructions as you like, each with an on and off switch. Enabled blocks go after the built-in sections, in order.

Watch the prompt size

The tab shows Prompt size as an estimate of how many tokens her instructions take from the context window. If it warns that the prompt crowds out the conversation, shorten sections or raise the context size in Settings → Performance.

Generation parameters

Generation parameters set Temperature, Top P, Top K, Min P, Repeat penalty, and Max reply length (tokens). Each shows the default for your model and its allowed range, and an empty field uses the default. They apply to replies from local models only. Cloud providers use their own settings, and background tasks such as summaries keep the app's settings. Reset parameters clears them.

Presets, export, and import

- Presets and files: type a Preset name and select Save preset to keep the whole setup (sections, your blocks, and parameters). Apply replaces what is set now, and the app asks you to click again to confirm.
- Export to file and Import from file use a `.json` file you can back up or share.

Restore the defaults

- Restore built-in sections restores only the built-in sections.
- Restore everything to defaults restores every section, removes your own blocks, and clears the generation parameters. Saved presets are kept. Both ask you to click again to confirm.

Avatar test bench

At the bottom of the tab, Avatar test bench lets you try every emotion, gesture, and face channel on the active character's 3D model. It needs the 3D avatar turned on and a model picked. A separate Avatar cue log can write, on this computer only, the start of each reply with the emotion and gesture picked for it. It is off by default.

38. Understand Privacy and Data Storage

Local Waifu is local-first, not "always local." Local models, memories, and embeddings can remain on your computer, but cloud AI, cloud speech services, Telegram, and remote Ollama send selected data outside the local app.

Know what stays local and what can leave the computer

Activity	Where data goes
Local chat through bundled Ollama	Processed on your computer
Memories and embeddings	Stored locally; embeddings remain local even when chat uses a cloud provider
Local image generation	Processed on your computer
Local speech recognition and voice	Processed on your computer
Cloud chat or image model	Prompt, relevant conversation context, and requested input go to the provider you selected
Cloud STT/TTS	Audio or text required by that service goes to the selected provider
Telegram	Messages and media pass through Telegram and the Local Waifu bot integration
Remote Ollama Host	Model requests go to the configured network computer
Model downloads and update checks	Network requests go to the relevant model host or update service
Screen awareness	Frames are looked at once in memory and thrown away; only her reaction is kept. With Allow a cloud vision model on and no local vision model, frames go to your cloud provider
Routines	Weather, feeds, and search topics are fetched from their own sites. If your chat model is a cloud model, what a routine gathers goes to that provider
Songs	The lyrics, style, and theme go to fal.ai when you press Create the song
Hands-free listening	Processed on your computer only. The audio is never saved and never leaves it
3D models	Models you import stay on your computer

Important: Offline Mode blocks configured cloud model routes, but it does not make Telegram operate without Telegram's servers, stop model downloads already requested, or block a custom network Ollama Host.

Locate your Local Waifu data

- macOS: `~/Library/Application Support/com.lumizone.localwaifu/`
- Windows: `%APPDATA%\com.lumizone.localwaifu\`

The app data directory contains the database, character data, soul files, portraits, generated images, voice references, settings, logs, and `secrets.json`. Model files may occupy separate model directories shown under Settings → Performance → Storage.

Soul files containing personality, memory, and knowledge are encrypted with cocoon encryption. Their key material is tied to the original computer environment (on macOS it also uses the passphrase you set during onboarding). This is one reason a copied full data directory is not the recommended self-service migration method between computers.

The `secrets.json` file can contain cloud provider API keys, the Telegram bot token, and activation-related data. On Unix-like systems it is protected with owner-only file permissions, but anyone who can read your user account's files may still read it. Protect the operating-system account and disk. Your license key is no longer readable or editable from the app window, and the app logs no longer contain her memories or the sentences she speaks.

Review and control permissions

Goal

Enable only the operating-system permissions needed for the features you use.

Before you start

Permissions are requested when a related feature is first used, not simply because the application launches.

Steps

1. Enable microphone access only for dictation, voice calls, voice messages, or hands-free listening.
2. Allow Screen Recording on macOS only if you turn on Screen awareness.
3. Allow Reminders access on macOS only if you turn on the Reminder follow-up routine. macOS asks once, when you switch it on.
4. Local Waifu does not read your calendar. The calendar feature was removed in 1.8.0.
5. To revoke access later:

- macOS: use System Settings → Privacy & Security and select the relevant category.
- Windows: use Settings → Privacy & security and select the relevant permission.

6. Keep Blur on Background enabled under Settings → General if you do not want conversation content visible when the window loses focus.

Expected result

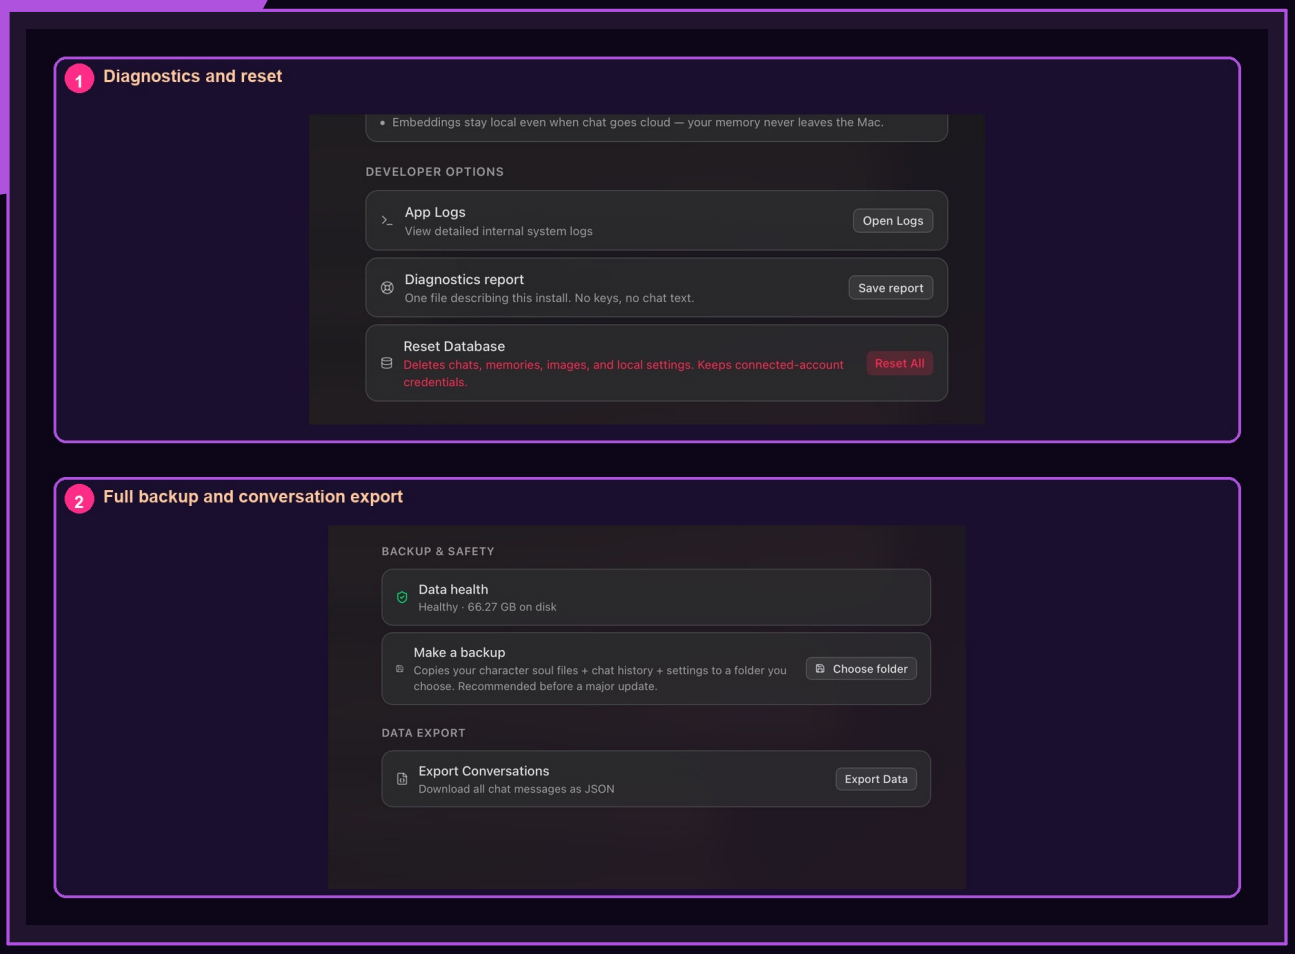
Only the features you intentionally use have operating-system access. Disabling a permission blocks that feature but does not delete existing data.

If something goes wrong

If a feature remains denied after you change the operating-system setting, quit and reopen Local Waifu. For microphone failures, also verify the input device under voice-call settings.

FIGURE 19

Advanced data controls: diagnostics and reset are separate from full backup and conversation export.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

39. Back Up, Export, and Move Your Data

Local Waifu offers three different data operations. They are not interchangeable.

Operation	Contains	Best use	Can it restore the whole app automatically?
Make a backup	Snapshot of the Local Waifu app data directory, including database, soul files, history, and settings	Recovery copy before an update or support-assisted restore	No. Version 1.8.0 has no button that replaces your live data with a backup. Restore a backup only checks a backup by copying it into a new folder
Export Conversations	Chat-message records in JSON	Reading, archiving, or processing messages	No; there is no general conversation-import control in this panel
Character export	One character plus related messages, memories, relationship data, and knowledge entities	Supported self-service transfer to another installation	Imports as a new character; it is not a full-app restore

Create a full backup snapshot

Goal

Create a consistent copy of the Local Waifu application data before a major update, reset, repair, or computer migration.

Before you start

- Choose a destination outside the Local Waifu data directory, preferably another drive or a synced backup location.
- Make sure the destination has enough free space.
- A full backup can contain private conversations, images, voice references, settings, and secrets. Store it securely.

Steps

1. Open Settings → Advanced.
2. Review Data health and resolve a reported integrity problem before relying on a new backup. It reads Healthy with the size on disk when all is well. On macOS the block is headed Backup & safety. On Windows the same block is headed Make a backup, and the conversation export block is headed Export Conversations instead of Data Export.
3. Find Make a backup and select Choose folder.
4. Pick an existing destination folder outside Local Waifu's app data directory.
5. Wait until the row reports Saved, a file count, and a total size.
6. In the destination, confirm that a new folder named `local-waifu-backup-<timestamp>` exists.

7. Copy the backup to a second protected location if it is your only recovery copy.

Expected result

Local Waifu creates a timestamped folder and a consistent SQLite snapshot. Temporary SQLite WAL/SHM files are not required in the backup.

If something goes wrong

- If the app reports that the destination overlaps Local Waifu's data directory, choose another folder or drive.
- If the copy stops, check free space and destination permissions, then use Backup again.
- Do not judge success only by the folder's existence; confirm the in-app file count and size.

Restore warning: Local Waifu 1.8.0 does not provide a button that replaces your live data with a backup. Do not copy a backup over the live data directory while the app is running. For full restoration, quit the app and obtain a version-appropriate support procedure. A manual replacement may fail if encryption or machine-bound data no longer matches.

Check that a backup is usable

Goal

Find out whether a backup would work without touching your current data.

Steps

1. Open Settings → Advanced and find Restore a backup.
2. Select Restore into a folder.
3. Choose the finished backup folder (Choose the finished backup folder).
4. Choose a new or empty folder (Choose a new or empty folder) outside your live data.
5. Wait for the result.

Expected result

The panel reports Restored N files with the size, the number of characters, messages, and media files, and N / M soul files verified. A green result means the backup can be read.

If something goes wrong

- This is a check, not a restore. Your current data is never touched or replaced. The app refuses an existing non-empty folder, a folder that overlaps your live data, and your live profile itself.
- Some soul files could not be verified means a backup is incomplete or its encryption does not match. Keep it unchanged and contact support.
- This backup contains no soul files to verify is shown when the backup has none.
- The restored folder is a normal copy. Delete it when you are finished.

Export all conversations

Goal

Save chat messages as a JSON file for archive or analysis.

Before you start

This export is not a backup. It does not include the complete character, memories, soul files, models, credentials, or settings, and the same panel has no general import function for this file.

Steps

1. Open Settings → Advanced → Data Export.
2. Select Export Data beside Export Conversations.
3. Choose a new `.json` filename and location outside the Local Waifu data directory.
4. Wait for Saved!.
5. Keep the file private; message content is readable JSON.

Expected result

A file named like `localwaifu-export-YYYY-MM-DD.json` contains the exported message records.

If something goes wrong

- Choose an absolute path ending in `.json`.
- Do not select the Local Waifu data directory.
- Choose a new name; the exporter refuses to overwrite an existing export.

Move a character to another computer

Goal

Use the supported character export/import path for a self-service migration.

Before you start

- This transfers the exported character and related messages, memories, relationship data, and knowledge entities. Knowledge entities arrive without the relationships between them; those are re-extracted as you keep chatting.
- It does not transfer installed models, provider keys, connected accounts, the Telegram token/pairing, app-level preferences, or the Direct license activation.
- The imported character is created as a new character; it does not overwrite an existing one.

Steps

1. On the old computer, open Settings → Characters.
2. Export the character and save the generated `.json` file.
3. Transfer that file securely to the new computer.
4. Install and complete onboarding in Local Waifu on the new computer.
5. Open Settings → Characters and import the character JSON.
6. Confirm the reported counts for messages, memories, and knowledge entities.
7. Reinstall the required chat, embedding, image, and voice models.
8. Re-enter cloud API keys or reconnect accounts.
9. Configure Telegram again with Settings → Telegram.

10. If using the Direct edition, redeem the same key on the new computer. There is no device limit, so there is nothing to deactivate on the old one.
11. Continue using the imported character and confirm that she recalls the imported memories.

Expected result

A new character appears with the exported history and character-related data. Recall of the imported memories works immediately on macOS, which restores their stored semantic vectors. On Windows the imported memories carry no vectors, so only memories created after the import take part in semantic recall and the older ones stay reachable through the Memory Book and keyword matching.

If something goes wrong

- Import failed means the JSON is damaged or incompatible; export it again from the source installation.
- If recall of the imported memories is weaker than on the old computer, keep the required memory model installed and keep chatting: on macOS this settles as the conversation continues, and on Windows the imported memories rely on keyword matching.
- If you need complete app recovery rather than a character transfer, preserve the full backup and contact support. 1.8.0 has no button that replaces live data with a backup.

40. Update or Uninstall Local Waifu

Update the Direct edition

Goal

Install an available Direct-edition update without risking known-bad local data.

Before you start

1. Create a backup with Settings → Advanced → Make a backup.
2. On macOS, check the Data health indicator in the same Backup & safety section. On Windows the Data health readout is in the block headed Make a backup.
3. Connect to the internet.
4. Note the platform differences: on both platforms updates are included with a license and the trial does not receive them. On macOS the installer also runs a data-health check first and blocks the install if it fails. On Windows the installer does not run that check before installing, although Data health is shown in Settings → Advanced.

Steps: macOS

1. Open Settings → General.
2. Under Updates, find the row Check for updates and select Check Now. The switch Check at launch above it controls whether the app looks for a new version when it starts. Nothing is downloaded or installed without asking you.
3. If a newer version appears, read the version number.
4. Select Install & Restart.
5. Wait for the download and allow the app to close and relaunch.
6. Open Settings → About and verify the new version.

Steps: Windows

1. Open Settings → General.
2. Under Updates, read Current version and select Check for updates. The switch Auto Update controls the check at launch.
3. If a newer version appears, read the version number.
4. Select Download v and wait for the download to finish.
5. Select Update now. The app closes so the installer can replace its files.
6. Start Local Waifu again and open Settings → About to verify the new version.

Expected result

The update installs only after your confirmation and leaves the app on the newer version. On macOS the install also waits for the pre-update database and encryption-metadata safety checks to pass; on Windows those checks are not part of the update flow, which is why the backup in the first step matters more there. The Windows installer that the app downloads is always the Direct build from the official release; the Store/Unlocked edition has no in-app update route.

If something goes wrong

- On macOS, if Data health fails or installation is blocked, do not force the update. Make a backup and contact support.
- On both platforms a trial cannot use the in-app update route (Updates are included with a license); use the supported download route for your entitlement.
- If the updater cannot connect, check proxy, firewall, VPN, and Offline Mode, then use Download the latest version manually if shown.
- Nothing installs merely because the check at launch is enabled. Download and installation require user action: the launch notification offers one Install action, and the installer cannot be started twice.
- On Windows, if Update now fails or the installer does not open, download the Direct installer from the official release page and run it. That path also installs over an existing copy.

Update the Store/Unlocked edition

Goal

Keep the Store/Unlocked build on its own distribution channel.

Before you start

The Store/Unlocked build has no Direct-edition updater, license tab, or Direct release-channel access.

Steps

1. Create a backup.
2. Check the store or client from which you obtained Local Waifu, such as itch.io or the itch app.
3. Download or install the update through that same channel.
4. Do not replace it with a Direct build merely to obtain the Direct updater.
5. After updating, open Settings → About and verify the version.

Expected result

The Store/Unlocked edition remains license-free and continues using its correct update channel.

If something goes wrong

Confirm that the downloaded installer is the Store/Unlocked build. If the update unexpectedly presents a trial, license field, or paywall, stop and obtain the correct store build before changing data.

Uninstall while keeping your data

Goal

Remove the application while preserving data for later reinstallation.

Before you start

Create a backup. A Direct license has no device limit, so there is nothing to deactivate before you uninstall. If you are selling or giving the computer away, follow the wipe steps in Chapter 45.

Steps: macOS

1. Quit Local Waifu.
2. Move Local Waifu from Applications to the Trash.

3. Do not delete `~/Library/Application Support/com.lumizone.localwaifu/`.

Steps: Windows

1. Quit Local Waifu and wait for its local engine processes to stop.
2. Open Settings → Apps → Installed apps → Local Waifu → Uninstall.
3. Leave Delete the application data unselected if that option is offered.

Expected result

The program is removed while the app data remains available to a later installation.

If something goes wrong

If reinstallation starts as a blank app, do not overwrite or reset anything. Confirm that the original data directory still exists, preserve it, and contact support.

Remove Local Waifu and its data completely

Goal

Remove the application, local content, and saved credentials from a computer you no longer intend to use.

Before you start

This is destructive. Export characters and create an external backup first. The Direct license needs no deactivation, because there is no device limit.

Steps

1. Remove cloud provider keys and connected-account credentials from their settings panels.
2. Open Settings → Telegram and use Forget token + pin.
3. Open Settings → Advanced → Developer Options, then find Reset Database, select Reset All, and confirm with Yes, delete. On Windows this leaves portraits, generated images, and voice references on disk; step 6 covers those.
4. Restart the app once if you need to confirm the reset completed.
5. Quit and uninstall Local Waifu using the platform procedure above.
6. Delete the remaining Local Waifu data directory only after verifying your backup:

- macOS: `~/Library/Application Support/com.lumizone.localwaifu/`
- Windows: `%APPDATA%\com.lumizone.localwaifu\`

7. Remove separately stored model files if you also want to reclaim their disk space.

Expected result

The application, user data, secrets, and optional model files are removed from the computer.

If something goes wrong

Reset All does not remove every credential by itself. In particular, license/trial data, cloud provider keys, and the Telegram token can remain in the secret store. That is why credentials must be removed explicitly and the remaining data directory deleted for a complete wipe.

41. Use Logs, Diagnostics, and Reset

Open the app log

Goal

Collect recent technical details after reproducing a failure.

Before you start

Logs may include filenames, device information, model names, error text, and other context. In 1.8.0 they no longer contain her memories or the sentences she speaks, but do not post them publicly without review.

Steps

1. Reproduce the problem once.
2. Open Settings → Advanced → Developer Options.
3. Under App Logs, select Open Logs.
4. Note the approximate time of the failure and copy only the relevant section when possible.

Expected result

The detailed application log opens in the system's default text viewer.

If something goes wrong

If the log is empty or does not open, use Diagnostics report → Save report instead. The app log and audit log serve different purposes; the app log is normally more useful for chat, voice, and startup failures.

Check the AI engine

Under Settings → Advanced → Developer Options, the AI engine row shows which Ollama version is answering and where it came from: built in, updated, GPU runtime (Windows), or your own Ollama. Select Check for engine update to look for a newer signed engine. The app downloads it from Ollama's official release host and checks it before use. Switching restarts the AI engine, so an answer in progress stops. There is no update button for your own Ollama, because it is not ours to update. If a model download fails with a message that the model needs a newer Ollama, this row is where to look.

Save a diagnostics report

Goal

Create one text file describing the build, configuration, hardware, runtime state, and recent log activity for support.

Before you start

> Privacy warning: Review the report before sending it. Local Waifu redacts license values, API keys, and Telegram tokens, but the report includes audio-device names and an app-log tail that is not fully content-filtered. Do not assume the UI phrase "No keys, no chat text" makes every line safe to publish.

Steps

1. Reproduce the issue.
2. Open Settings → Advanced → Developer Options.
3. Under Diagnostics report, select Save report.
4. Locate the revealed file named like `diagnostics-YYYYMMDD-HHMMSS.txt`.
5. Open it in a text editor.
6. Review every section, especially the recent app-log tail.
7. Remove or redact private text, usernames, paths, device names, prompts, or other details you do not want to share.
8. Attach the reviewed file to a private support message.

Expected result

A reviewed diagnostic text file is ready to send without exposing information you chose to remove.

If something goes wrong

If report creation fails, send the app version, platform, edition, exact error message, reproduction steps, and a reviewed excerpt from App Logs instead.

Reset all application data

Goal

Return the main application database and local content to a fresh state when repair is impossible or when intentionally starting over.

Before you start

> Destructive action: Reset All is irreversible from within Local Waifu. Create a backup and export any characters you want to keep.

The reset deletes the database, characters, conversations, memories, and soul files. On macOS it also deletes portraits, generated images, voice references, and the natural voice engine's working data. On Windows those folders (portraits, images, voice references) are left on disk, so delete the app data directory as well if you want them gone. It does not automatically remove the Direct license/trial record, cloud provider keys, connected-account credentials, or Telegram token from the secret store.

Steps

1. Create a full backup and verify its folder, file count, and size.
2. Export important characters separately.
3. If you also want credentials removed, delete provider keys and use Settings → Telegram → Forget token + pin first.
4. Open Settings → Advanced → Developer Options.
5. Find Reset Database and select Reset All.
6. At Sure?, select Yes, delete.
7. Wait for All data deleted.
8. Quit and restart Local Waifu as instructed.

Expected result

The main app data is cleared, and the application starts in a fresh state after restart. License, trial, provider, or Telegram credentials may remain unless removed separately.

If something goes wrong

- Do not repeat the reset if the app appears partially initialized; quit and reopen it first.
- On Windows, if you also want portraits, generated images, and voice references removed, delete `%APPDATA%\com.lumizone.localwaifu\` with the app closed.
- Do not attempt to copy a backup into the live data directory while Local Waifu is running.
- There is no complete Restore UI in version 1.8.0. Preserve the backup and contact support for a full restore procedure.

42. Troubleshooting

Run the basic diagnostic sequence

Goal

Identify the failing layer before changing or deleting data.

Before you start

Do not use Reset All as a first troubleshooting step. Preserve the exact error message and create a backup before destructive action.

Steps

1. Confirm the version under Settings → About.
2. Restart Local Waifu and reproduce the issue once.
3. Run Settings → Performance → Runtime status → Check.
4. Check the active chat, memory, image, speech, and voice models relevant to the failure.
5. Check Settings → Performance → Network, especially Offline Mode, Use System Proxy, and Ollama Host.
6. Check free disk space and available RAM/VRAM.
7. Open Settings → Advanced → Developer Options and, under App Logs, select Open Logs.
8. If support is needed, save and review a Diagnostics report before sending it. It is on the same Developer Options section.

Expected result

The problem is narrowed to installation, model/runtime, resource pressure, network routing, permissions, Telegram pairing, license, update, or data health.

If something goes wrong

Stop before reset or manual file replacement. Preserve the data directory and backup, record the last successful action, and contact support.

The app will not open

Symptom	Likely cause	What to do
A screen says Local Waifu could not open your database.	The database or an encryption piece could not be read	Nothing was changed. Use Save a copy of everything and read Chapter 47 before anything else
A screen says Local Waifu could not check its data.	A temporary failure while checking	Select Try again or Reload window

Installation and launch

Symptom	Likely cause	What to do
macOS blocks the first launch	Gatekeeper verification or an unrecognized download	Confirm the installer came from the official source. In Finder, Control-click Local Waifu, choose Open, and approve the system prompt if macOS offers it. Do not routinely disable Gatekeeper or strip quarantine attributes.
Windows shows SmartScreen	The installer may not have paid Authenticode reputation	Confirm the file came from the official source, inspect the publisher/source, select More info, then Run anyway only if you trust the file.
Windows app does not open or shows a blank window	Missing/damaged WebView2, antivirus interference, or incomplete install	Reboot, install/repair Microsoft Edge WebView2 Runtime, allow the official app in security software, and reinstall.
Installer appears stuck	Antivirus scanning, low disk space, or running app processes	Wait for scanning, verify free space, quit Local Waifu and its local engine, then retry.

Chat models, downloads, and memory

Symptom	Likely cause	What to do
Model download will not start or stalls	Engine startup, network, proxy/VPN, registry access, or disk space	Check Use System Proxy, VPN/firewall, free space, and internet access. Restart the app, then retry in Settings → Models.
Hugging Face repo has no installable quant	It contains sharded GGUF files or no GGUF	Choose a repository with a single runnable `.gguf` quant. Sharded GGUF and SafeTensors are not installed through the chat-model search panel.
Chat shows offline or no reply	Local engine/model unavailable, bad Ollama Host, cloud blocked, or missing provider credentials	Run Runtime status, check the active chat model, reset Ollama Host, disable Offline Mode only for intentional cloud use, and verify the provider account/key.
Chat says the local model failed but a reply still arrived from a cloud provider	Use the cloud when the local model can't answer is on	Turn it off in Settings → Performance → Network if you do not want any reply to leave your computer. It is off by default.
Replies are very slow or the app runs out of memory	Model or context is too large; GPU placement is partial/CPU	Reduce Context Window, choose a smaller quantized model, turn off Deeper thinking, and check GPU placement on Windows.
Memory recall is weaker after changing the embedding model	Re-embedding is still in progress	Keep the required memory model installed and allow continued use to rebuild embeddings. Keyword-only recall may be used temporarily.

Images and local image models

Symptom	Likely cause	What to do
Image cannot be attached	File is not an image, exceeds 5 MB, another image is pending, or vision is unavailable	Use one valid image under 5 MB and verify Sees images or a vision-capable model/provider.
Custom checkpoint import fails	Wrong format, too small, invalid SafeTensors, or not SDXL	Use one regular `.safetensors` SDXL checkpoint of at least 512 MiB. `.ckpt` is not supported by the 1.8.0 importer.
Imported LoRA cannot be enabled	Active model is not SDXL, family is unknown, or three LoRAs are already active	Select an SDXL base model, use a verified SDXL LoRA, and keep at most three active. Confirm compatibility with a test render.
SDXL imports but will not run	Insufficient RAM or disk/runtime resources	Custom SDXL uses a 16 GiB RAM profile. Close other apps or use a lighter image model.
Image model cannot be deleted	Rendering or download is active	Wait for the operation to finish, then retry.

Voice, microphone, and calls

Symptom	Likely cause	What to do
A call cannot start	Speech recognition is missing or unavailable	Install/configure Whisper or a supported cloud STT route. STT is required for a full call.
No spoken reply	Voice model is missing or unavailable	Install Supertonic or configure a supported cloud TTS route. A fallback may be used, but quality can be lower.
Speech has the wrong accent or language	Wrong language, English fallback, or mismatched custom sample	Select the correct language and voice model. For a custom voice, use 10 to 30 seconds of clean speech in the target language.
Microphone is silent	OS permission denied or wrong input device	Grant microphone permission in the operating system and select/test the correct input device. Restart afterward.
Dictation stops unexpectedly	Ten-minute limit or disconnected microphone	Start another dictation session and reconnect or reselect the microphone.

Telegram

Symptom	Likely cause	What to do
Test succeeds but bot does not answer	Bot was not started	Select Start. Testing does not save or start the token.
Bot will not pair with a new account	It is pinned to the first chat	Use Forget token + pin, restart the bot, and message it first from the correct account.
Voice only returns text	Voice model is not ready	Install/configure the voice model. Text is the intentional fallback.
Telegram model change also changes desktop	Shared active-model setting	This is expected. Select the preferred model again from either interface.

License and updates

Symptom	Likely cause	What to do
Direct key is rejected	Typo, no network for activation, or an invalid key	Paste the exact key, reconnect, and retry. There is no device limit. If the app says the license could not be activated on this device, write to support.
Trial expired	Direct-edition trial period ended	Activate a valid lifetime key. User data is retained; do not reset in an attempt to obtain a new trial.
License tab is absent	Store/Unlocked edition	This is expected; that edition does not require a key.
Update control is absent	Store/Unlocked edition	Update through the store/client used to obtain the app. Do not mix it with the Direct updater.
Install & Restart is blocked	Data-health or encryption-metadata check failed	Create a backup and contact support. Do not force file replacement.

3D avatar, desktop pet, and screen awareness

Symptom	Likely cause	What to do
The portrait shows instead of the 3D model	The model could not load, or 3D avatar is off	Open Settings → 3D Avatar, turn on 3D avatar, and pick another model. A model that fails in the preview cannot be used
A model imports but her lips do not move	The model has no mouth shapes	Use a model that has them. It still blinks and smiles
A face does not match her feeling	Automatic expression matching missed it	Open Expression mapping and choose the model's own expression for that feeling
Textures are missing on an MMD model	The bare `.pmx` was imported	Import the `.zip` the model came in
Desktop mode shows a card instead of her	The 3D avatar is off or her model is not ready	Turn it on and pick a model in Settings → 3D Avatar, then try again
The desktop shortcut does nothing	Another app owns the chord, or the avatar is off	Check the system keyboard shortcuts, or use Put her on the desktop in the sidebar
She stays still on the desktop	Quiet Mode is on	She stays put while a fullscreen app, a meeting, or a locked screen is active
A selfie ignores the 3D model	Her 3D model is not on screen, or the image model cannot take a reference	Use local FLUX.2, Google, or OpenAI as the image model, or accept the portrait
Make it a real photo is unavailable	The image model cannot take a reference picture, Offline Mode is on, or no active license	Pick FLUX.2 klein in Settings → Models → Image Models
Screen awareness cannot see	Missing Screen Recording permission (macOS), no vision model, or the app is on the protected list	Allow Screen Recording, relaunch, set up a vision model, and check Never capture these apps

Routines, voice features, and groups

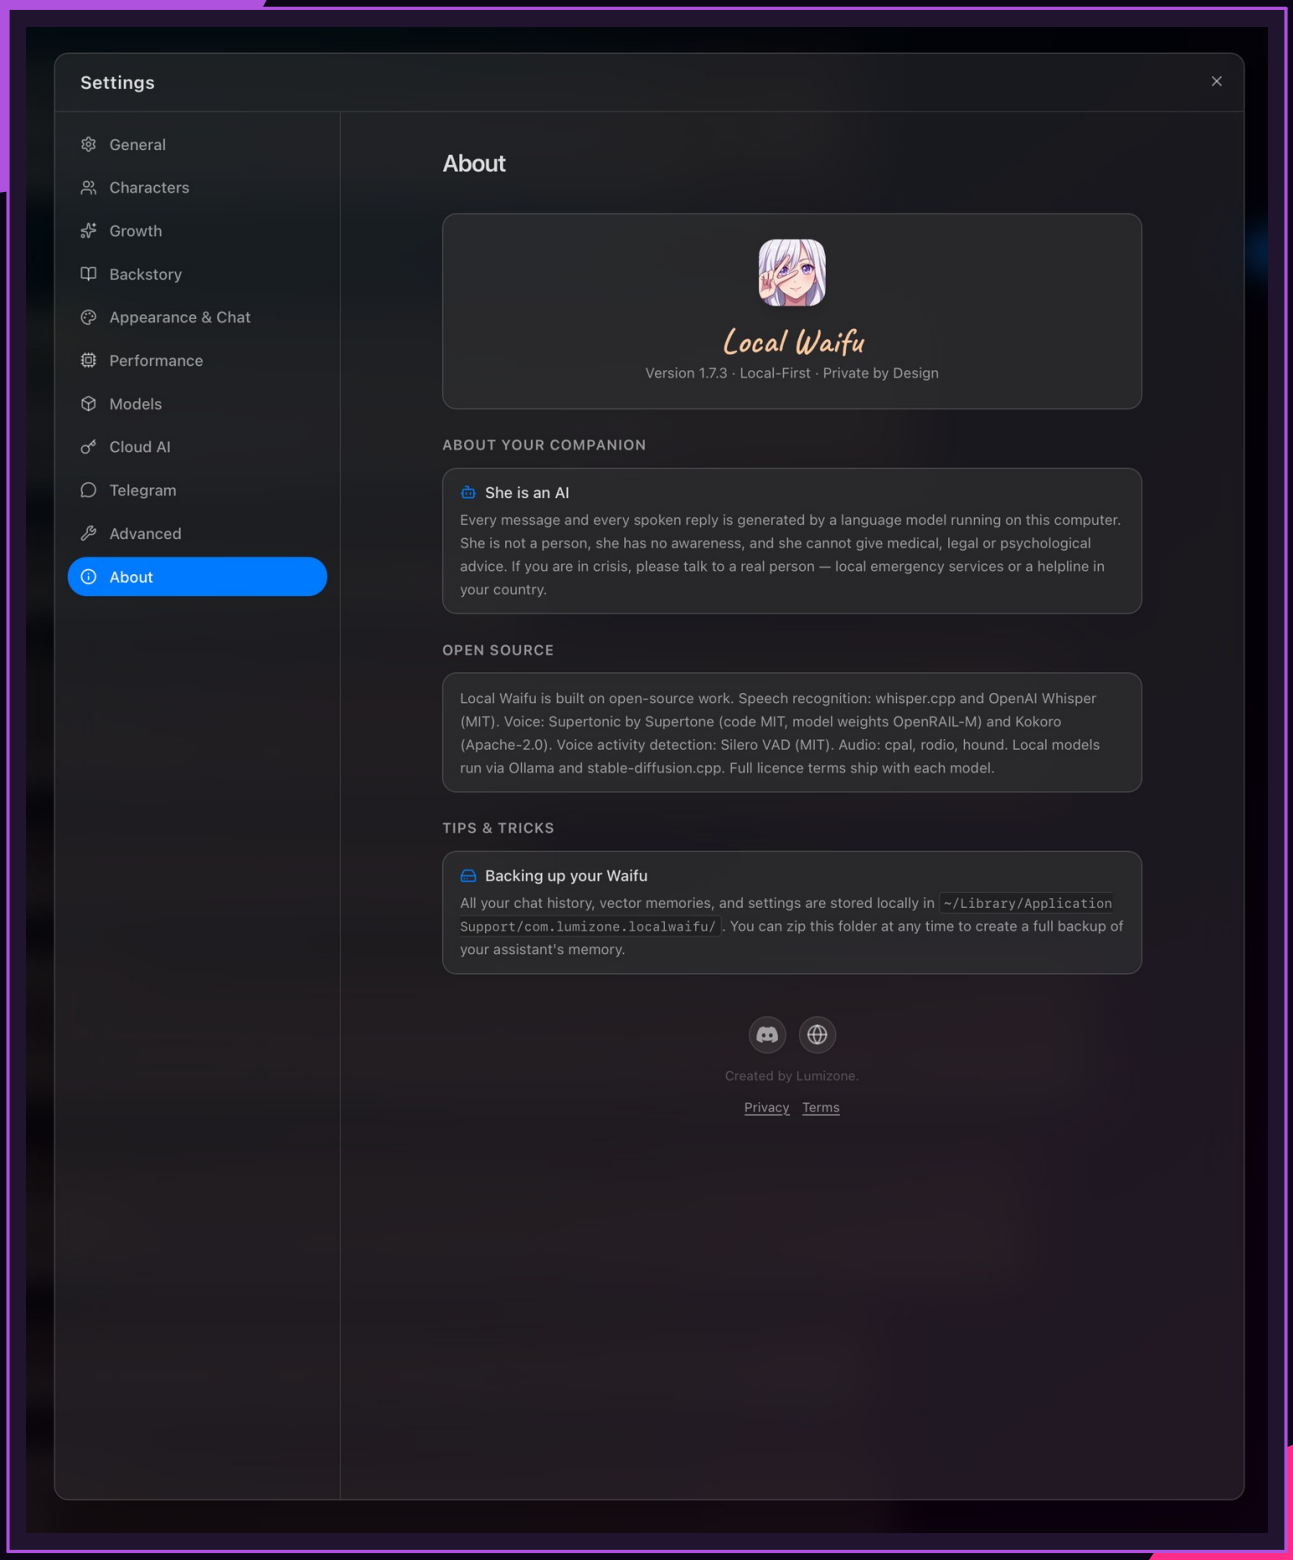
Symptom	Likely cause	What to do
A routine did nothing	Its switch is off, Offline Mode is on, or she was quiet or you were typing	Check the switch and Last run, and turn off Offline Mode. It retries on its own when quiet
Reminder follow-up is missing	It exists only on macOS	Use another routine on Windows
Morning brief has no weather	No city set	Enter a City and select Find
A voice message is not sent	It was too short, over two minutes, or a reply was still coming in	Hold the button while you talk, and try again when she finishes
Hands-free does not start	The microphone is busy, the wake phrase is too short, or the wake-word model is missing	Check the status under Hands-free and the message shown
A song or story will not start	No fal.ai key (songs), Offline Mode (songs), a call is open, or no active license	Follow the message the app shows. Songs need a fal.ai key under Settings → Models → Image Models
In a group, the wrong character answers	The default rules chose someone else	Start the message with her name, or say a word that addresses everyone

Backup, export, and data recovery

Symptom	Likely cause	What to do
Backup destination is rejected	It overlaps the live Local Waifu data directory	Choose an external folder or another location outside the app data tree.
Conversation JSON does not restore the app	It is an export, not a backup	Use character export/import for self-service transfer. Preserve the full backup for support-assisted recovery.
There is no button that restores a backup over my data	1.8.0 only offers Restore a backup, which checks a backup by copying it into a new folder	Use it to confirm the backup is usable. Quit the app, preserve the backup unchanged, and request a support procedure for a full restore. Never merge it into the live directory while the app is running.
Reset did not remove license, cloud keys, or Telegram token	Reset does not clear the secret store	Remove provider credentials explicitly, use Forget token + pin, and delete the remaining data directory only for a deliberate full wipe.

FIGURE 20

About shows the installed Local Waifu version, AI disclosure, open-source acknowledgements, and backup reminder.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

43. Shortcuts, Contact, and About

Keyboard shortcuts

The following shortcuts are implemented in Local Waifu 1.8.0 even though About no longer contains a keyboard-shortcut reference section.

Action	macOS	Windows
Open Settings	`Command+,`	`Ctrl+,`
Search chats	`Command+F`	`Ctrl+F`
Start or end a voice call	`Command+Shift+C`	`Ctrl+Shift+C`
Switch to Full view	`Command+Option+1`	`Ctrl+Shift+1`
Switch to Mini view	`Command+Option+2`	`Ctrl+Shift+2`
Switch to menu bar/tray view	`Command+Option+3`	`Ctrl+Shift+3`
Open visual novel	`Command+Option+4`	`Ctrl+Shift+4`
Put her on the desktop (pet)	`Command+Option+5`	`Ctrl+Shift+5`
Cancel or close the active transient UI where supported	`Escape`	`Escape`

On macOS, `Command+Shift+1` and `Command+Shift+2` still work for Full and Mini because macOS does not reserve those chords; `Command+Shift+3/4/5` are system screenshot keys, which is why the other view modes use Option.

Search note: both `Command/Ctrl+F` and `Command/Ctrl+Shift+F` open chat search.

Other shortcuts: the two global shortcuts for screen awareness (capture and ask her, and pause or resume watching) are listed under Settings → Screen → Shortcuts for your platform (Chapter 26). While a call is open, `Escape` hangs up. Chapter 44 covers `localwaifu://` links, which can start actions from outside the app.

Check the installed version and edition clues

Goal

Collect the information support needs and verify an update.

Before you start

Know whether you installed Local Waifu from the Direct download or a Store/Unlocked channel. The presence of Settings → License is a useful clue: Direct shows it; Store/Unlocked does not.

Steps

1. Open Settings → About.
2. Read the number shown after Version.
3. Note your operating system and whether Settings → License is present.

4. If reporting a model or voice issue, also note the active model names.
5. If requested, save a Diagnostics report, review it, redact private details, and attach it privately.

Expected result

You have the app version, platform, edition clue, reproduction steps, exact error text, and a reviewed diagnostic file if needed.

If something goes wrong

If About cannot load, include the installer filename or download channel and state that the version screen was unavailable. Do not reset the app merely to retrieve a version number.

Contact and official links

- Support email: contact@localwaifu.com
- Website: https://localwaifu.com
- Community: https://discord.gg/RbaVVTADxz
- Direct-edition source/releases: https://github.com/lumizone/local-waifu
- Privacy: https://localwaifu.com/privacy/
- Terms: https://localwaifu.com/terms/

The About screen also identifies Local Waifu as local-first and private by design, provides the AI disclosure, lists open-source components, and links to Privacy and Terms. The GitHub button is shown only where appropriate for the distribution edition.

When contacting support, include:

1. Local Waifu version.
2. macOS or Windows version.
3. Direct or Store/Unlocked edition.
4. What you expected and what happened.
5. Exact steps that reproduce the issue.
6. The exact error message.
7. Relevant model/provider names, without API keys or tokens.
8. A reviewed and redacted diagnostics report when useful.

Never send a license key, Telegram bot token, provider API key, unreviewed diagnostics report, or private conversation export in a public channel.

44. Use localwaifu Links

A `localwaifu://` link lets another program talk to Local Waifu. You can send her a message, start or end a call, start screen watching, start focus guard, or run the morning brief, from a desktop shortcut, a launcher, a script, or on macOS the Shortcuts app and Siri.

Links only run while Local Waifu is unlocked. The settings for them are in Settings → Advanced, in the block called Shortcuts and Siri on macOS and Shortcuts and links on Windows.

Turn links on or off

- Allow Shortcuts links (Allow localwaifu:// links on Windows): on by default. Turn it off to ignore every link. A link that arrives while this is off is rejected with a message pointing here.
- Skip the confirmation for these actions: off by default. While it is off, a link that sends a message, switches character, starts a call, starts screen watching, or starts focus guard asks you first: A Shortcut wants to... with Allow and Deny. Turn the skip on only if you trust every shortcut on this computer. On Windows any program can open a link, so be strict.

The links

Copy any of these with Copy link in the block:

Link	What it does
`localwaifu://send?text=Good%20morning`	Send her a message. Add `&character=Name` to pick who
`localwaifu://call/start`	Start a voice call
`localwaifu://call/stop`	End the call
`localwaifu://watch/start`	Start watching your screen
`localwaifu://watch/stop`	Stop watching your screen
`localwaifu://focus/start?minutes=25`	Start focus guard (minutes are optional, 1 to 240)
`localwaifu://focus/stop`	Stop focus guard
`localwaifu://brief/run`	Run the morning brief now

A link marked coming soon in the block is not active in your build. If a link names a character that does not exist, the app says No character named that name.

To send your own words, URL-encode them: a space becomes `%20`. A link is limited to 4096 bytes after encoding, which is roughly 650 characters of non-English text, so keep messages short.

Use a link on macOS

Goal

Make a Siri phrase or a Shortcuts action that talks to her.

Steps

1. Open the Shortcuts app and choose New Shortcut.
2. Add the Open URLs action and paste one of the links.
3. Name the shortcut. That name is what you say to Siri.
4. To send dictated text, add Ask for Input, then URL Encode, and use `localwaifu://send?text=` followed by the encoded text.

Links only work once Local Waifu is installed in the `/Applications` folder.

Use a link on Windows

Steps

1. Desktop shortcut: right-click the desktop, choose New → Shortcut, and paste a link as the location. Opening the shortcut runs the action.
2. Run dialog: press `Win+R`, paste a link, and press Enter.
3. Anything that can open a URL also works: a launcher such as PowerToys Run, a Stream Deck website action, an AutoHotkey command, or `start "" "localwaifu://call/start"` in a terminal.

Windows learns the links when Local Waifu is installed with its installer. A copy that was only unpacked does not register them. If you are busy in another app, Windows may only flash Local Waifu on the taskbar instead of bringing it to the front. Open it to answer Allow or Deny.

45. When the Trial Ends, and Moving Your License

Chapter 8 covers how the trial starts. This chapter covers what happens when it ends, what each activation message means, and what to do when you change computers.

What an expired trial blocks

When the seven days run out, the app is locked except for export. You can still get your data out: character export and conversation export keep working, which is the whole point of locking it that way rather than leaving you with a wall.

Your conversations and soul files are not deleted and not weakened. They stay encrypted on the computer, waiting for a license.

Buying a lifetime license unlocks the app again with everything intact.

If something goes wrong

If the app looks locked but you have already paid, check Settings → License. If the key is not there, paste it from your purchase email and activate. If it is there and still locked, read the next section.

Activation problems and what each message means

The app recognises three failures and tells you what to do about each.

What you see	What it means	What to do
That key wasn't recognised	The key was mistyped, or it belongs to a different product	Copy it straight from the purchase email, the whole thing, with no extra spaces
This license couldn't be activated on this device	The key could not be activated here	Write to support at contact@localwaifu.com and they will sort it out
Couldn't reach the license server	No network, or the licensing service is unreachable	Check the connection and try again. Your key is fine and nothing has been consumed

If something goes wrong

If License revoked or License disabled appears, that is a state only the issuer can set. A refund is the usual reason. Contact support with your key.

There is no device limit

A license is for your own devices, within fair use. Older versions counted activations and had a Deactivate this device button. Both were removed in 1.7.4. Settings → License now says so plainly: New computer? Install Local Waifu there and redeem the same key, there's no device limit.

Use your license on another computer

Steps

1. On the old computer, create a backup and export anything you want to keep (Chapter 39).
2. On the new computer, install Local Waifu and complete onboarding.
3. Open Settings → License, paste the key, and select Redeem.
4. Move your data using the character export path in Chapter 39, or restore a full backup with support.

Expected result

The new computer is licensed. Both computers keep their own data.

If something goes wrong

If the key is rejected on the new computer, write to support. Do not buy a second license just to get past a problem.

Before you sell or give away a computer

1. Remove cloud provider keys and connected accounts.
2. Use Settings → Telegram → Forget token + pin.
3. Reset the app data, then delete the data directory. Chapter 41 has the exact steps and which folders a reset does and does not remove.
4. Uninstall.

There is no license slot to free, so there is nothing to deactivate. If you skip the steps above, the next person can read everything in the app data directory.

Questions people actually ask

Can I use one license on my desktop and my laptop? Yes. The license is for your own devices, within fair use.

Do updates cost anything? No. A lifetime license includes updates, with no renewal and no expiry. The trial does not receive updates.

What if I reinstall on the same computer? The data directory survives, so you come back to the same companion. You may need to activate again with the same key.

Does the license key show in the app? No. Your license key is no longer readable or editable from the app window.

46. Safety, Boundaries, and What She Is Not

Two of these are things Local Waifu does. The third is something no companion app can do for you, and it is worth reading once.

The crisis banner

What it is. If a message contains a phrase that indicates real risk, Local Waifu shows a dismissible banner beside the conversation with helpline numbers for your interface language.

Where it appears. Over the chat stream, in the full, mini, and menu bar/tray views.

Languages. Helpline lists ship for English, Polish, German, Spanish, Japanese, Korean, and Chinese, matching the interface languages.

What it does not do. It does not block, censor, or refuse anything. The reply you were going to get arrives exactly as it would have. That is deliberate: someone reaching for the one thing that talks to them at 3am should not find that it has locked itself.

How detection works. Whole phrases, not keywords, matched against a normalized version of your message. A keyword list would fire on "I'd kill for a coffee" and "this laptop is dying", and a banner that cries wolf is a banner nobody reads. The app carries an explicit list of alarming-sounding constructions that are not treated as signals.

Expected result

The banner appears, lists helplines for your language, and dismisses with one click. The conversation continues normally.

If something goes wrong

- If a helpline number is wrong or has changed, report it through the contact address on the website. These numbers move, and the app ships them as data.
- If the banner appears when nothing is wrong, that is a false positive worth reporting. Say what you typed.

What the banner says, and why

The banner does not pretend that she is a person. It says plainly that she is a language model on your computer, that she can keep you company, and that she cannot get help to you. That sentence is the reason the feature exists.

Adults only

Local Waifu is sold as an adults-only product. The first screen of onboarding asks you to tick I confirm that I am at least 18 and understand that this is an AI companion, not a person before you can continue. That is a statement you make, not an identity check: there is no age verification in the app, so the boundary is a condition of sale that depends on you honouring it.

The same section of the Terms of Use covers the AI disclosure: every message and every spoken reply is generated by a language model running on this computer, she is not a person, she has no awareness, and she cannot give medical, legal, or psychological advice.

What she cannot do

- She cannot verify anything. She can be confidently wrong, and she has no way to check a fact for you.
- She cannot give medical, legal, or psychological advice. Treat anything that sounds like advice as text, not as guidance.
- She cannot contact a person on your behalf. She has no live human behind her.
- She is not monitored. Nobody at Lumizone reads your conversations, sees your messages in real time, or intervenes.
- She cannot remember outside your own database. If your computer loses the data, she does not have a copy.

Where your data sits in all of this

- Conversations, memory, mood, and the relationship record live in a local database on your computer.
- Nothing in a normal local session leaves the machine. See Chapter 38 for the operations that do.
- If you deliberately connect a cloud provider, the prompt and relevant context for that request go to that provider, and their own policies apply from there.
- The app does not read your conversations to train anything, and there is no analytics path that carries message text.

If you need help

Local Waifu is not a crisis service and cannot be one. If you are in danger now, contact local emergency services. The banner lists helplines for your language, and the Terms of Use carries the same list.

47. Recover a Character or the Whole App

Read this before you delete anything. The recovery routes are not interchangeable, and two of the four in this chapter need to be prepared before you need them. If the app cannot open at all, jump to the section about the recovery screen.

The four routes, and when each one works

Route	Works after	Needs preparing in advance
Character JSON export and import	The app is reinstalled, or you are on a different computer	No. You can export at any time while the app opens
Full backup folder	A crash, a bad update, a reset	Partly. The backup must exist; see Chapter 39
Soul recovery code	The encrypted soul files are lost, or you moved to a new computer	Yes. The code must be created and stored while you still have access
Conversation JSON export	Almost nothing on its own	No, but it is not a recovery method

The last row is the one people get wrong. Conversation JSON is an export for reading and archiving, not a restore. There is no control that imports it back into the app. If it is all you have, you have a readable transcript and not a companion.

Route 1: character export and import

The supported self-service route and the only one you can use entirely on your own. Chapter 15 has the full steps and Chapter 39 has the migration checklist.

Use it when the app still opens and you want that character on another installation.

If something goes wrong

If the import reports a failure, the file is damaged or was cut off during transfer. Export it again from the source installation rather than editing the JSON.

Route 2: full backup

The only route that covers everything at once: database, soul files, portraits, generated images, voice references, and settings.

Important limitation in 1.8.0: no button replaces your live data with a backup. A backup is a folder you keep, and bringing it back is a support-assisted procedure rather than a click. Do not copy a backup over a live data directory while the app is running. What you can do on your own is check that a backup is usable: Settings → Advanced → Restore a backup copies it into a new folder and verifies the soul files, without touching your live data (Chapter 39).

Use it before an update, before a reset, and before any change you are not certain about.

If something goes wrong

If the app will not start at all, stop changing anything on disk. Preserve the data directory as it is and contact support with the backup and the current state.

Route 3: the soul recovery code

This is the route for the case nothing else covers: your encrypted soul files exist, but the key that opens them is gone.

Why a separate key exists. Soul files, meaning personality, memory, and knowledge, are encrypted with a key tied to the computer you created the character on. If you copy the data directory to another computer, or the app's secret store is lost, those files cannot be decrypted by the new installation. A recovery code wraps that key under a passphrase you choose, which makes it portable.

What a recovery code is

- A text code that starts with `LWR1:`.
- Protected by a passphrase you set, which must be at least 12 characters.
- Machine-independent. That is the entire point of it.

How you would use one

1. While everything still works, on the computer that holds the character, you generate the code and set the passphrase.
2. You store the code and the passphrase in two separate safe places. A password manager for the passphrase, a printed copy or an encrypted note for the code is a reasonable split.
3. Later, on the new or repaired installation, you open the same character, enter the code and the passphrase, and the soul files open again.

In version 1.8.0 this is not reachable from the interface. The mechanism is implemented and shipped in the application, but there is no button for it in Settings. If you are in the situation it solves, contact support and they will walk you through creating or redeeming a code.

If something goes wrong

- A passphrase that is too short is rejected outright. The 12-character floor exists because the code can be handed around and guessed offline.
- A pair of code and passphrase that does not match fails without damaging anything. It is an authenticated blob, so a wrong passphrase cannot produce a half-decrypted character.
- If you lose the passphrase, the code is useless and there is no recovery from it. Store the two parts deliberately.

When the app cannot open: the recovery screen

If Local Waifu cannot open its database at start, it does not guess. It shows a screen titled Local Waifu could not open your database. and says nothing has been changed: your database, its sidecar files, and her soul files are still where they were. If it could not even check its data, a shorter screen offers Try again and Reload window.

The recovery screen offers these actions:

Action	What it does
Save a copy of everything	Copies the whole data folder into a new folder outside it: the database and its sidecars, her soul files, all pictures, audio and models, her key store, and a manifest. Nothing is parsed or changed. A copy includes the key store that unlocks her soul files, so keep it as private as the app's own folder
Check the database again	Read-only. If it opens now, the app restarts into its normal start
Check a restored copy	Points at a database file from a backup and checks it read-only. Nothing is installed or replaced
Show the data folder	Opens the data folder
Start fresh, keep a copy	For when nothing else helps and there is no backup. Everything in the data folder (memories, chats, pictures, and keys) moves into a new folder beside it, untouched, and the app starts as on its first day. You press Yes, move it aside and start fresh to confirm. Nothing is deleted, and downloaded models stay installed
Quit	Closes the app

The screen also lists which startup step failed, and for encryption problems it explains which piece is missing: for example the per-install encryption salt, the hardware identifier, or the key store (`secrets.json`). It does not replace any of them behind your back, because that could make her soul files permanently unreadable.

Recommended order

1. Select Save a copy of everything first, to a drive outside the data folder.
2. Use Check the database again once the cause is gone, such as a full disk or a moved folder.
3. Contact support with the screen's technical details before you use Start fresh, keep a copy.
4. Use Start fresh, keep a copy only when you accept starting as on the first day. Your old data is kept beside the data folder, so support can still help you.

What nothing can recover

Be clear about these before you rely on any route.

- Deleted characters. Once a character is deleted and the data is gone, no route brings it back. Export first, always.
- A forgotten soul recovery passphrase. There is no reset for it.
- Encrypted soul files without their key. Without the original computer's key or a recovery code, the ciphertext stays closed.
- Anything that was never backed up. Lumizone holds no copy of your data. There is no server-side backup to fall back on.

The order to do things if you are unsure

1. Stop writing to the app. Every turn adds new data and makes a clean snapshot harder.
2. Make a full backup if the app still opens.
3. Export the character you care about most, as a second independent copy.
4. Write down what happened, with the exact error text, and note which of the four routes above you still have available.
5. Contact support before you reset, reinstall, or move files by hand.

Resetting is the one step that can turn a recoverable situation into an unrecoverable one. Chapter 41 explains exactly what Reset All removes.

I want to...

Use this index when you know the task but not the Local Waifu term for it.

I WANT TO...	GO TO
Move one companion to another computer	Chapter 15 · Export and Import a Character
Move or protect the entire installation	Chapter 39 · Back Up, Export, and Move Your Data
Install my own chat model	Chapter 28 · Install a GGUF Model from Hugging Face
Import an image checkpoint or LoRA	Chapter 31 · Import a Custom SDXL Checkpoint or LoRA
Make memory recall work	Chapters 13 and 27 · Memory Book and the embedding model
Find a setting by name	Chapter 16 · Views, Settings, and Where Everything Lives
Teach her world facts	Chapter 17 · Use the Lorebook for Backstory
Understand why she messaged first	Chapter 18 · When She Messages You First
Read the relationship dashboard	Chapter 19 · Read the Growth Dashboard
Stop a habit she picked up	Chapter 20 · Tune How She Has Adapted to You
Give her a 3D body	Chapter 21 · Give Her a 3D Avatar
Talk with several characters at once	Chapter 23 · Use Group Chats
Point her at my own server	Chapter 32 · Use Your Own Server as a Cloud Provider
Reduce RAM usage or speed up replies	Chapter 36 · Performance, Hardware, and Network
Use Local Waifu from my phone	Chapter 35 · Use Local Waifu from Telegram
Move my license to a new computer	Chapter 45 · When the Trial Ends, and Moving Your License
Bring a companion back	Chapter 47 · Recover a Character or the Whole App
Collect logs or a diagnostics report	Chapter 41 · Logs, Diagnostics, and Reset
Solve an error	Chapter 42 · Troubleshooting
Find shortcuts, version, or support	Chapter 43 · Shortcuts, Contact, and About
Know the safety boundaries	Chapter 46 · Safety, Boundaries, and What She Is Not

End of guide

Keep this PDF with your Local Waifu backup instructions. For current downloads and product information, use the official Local Waifu channels. For account, license, or recovery help, contact contact@localwaifu.com.

Local Waifu · Your companion, your computer, your data.