Office AI that never leaves the office.

What the office AI set-up is, who uses which tool for what, how to lend your PC to the render farm, how to run a render, which model to pick, and the rules.

1 What it is

Two free tools run on our own PCs. The AI models they use sit on the office server, `BABOOM-BIM-STANDARD\`. Nothing is sent to an online AI service: your drawings, images and questions stay inside the office network.



2 Who uses what, for what

ComfyUI: renders and visualisation

Turn a model view or sketch into an image, change materials, make the same view from many camera angles. Visualisation team and designers. Big batches can run on the render farm.

`image models: Qwen-Image-Edit 2511`

LM Studio: text and code assistant

A private chat for drafting text, summarising, and writing Dynamo/Python or LISP drafts. Everyone. Runs on your own PC; code it writes is a draft until checked (BBM-STD-73).

`text models: Qwen2.5-Coder 7B, Qwen3.8 27B`

The render farm

Idle office PCs with an NVIDIA graphics card each render part of a batch, so a set of camera angles finishes sooner. It does not make a single image faster.

`extension: ComfyUI-Distributed`

BABOOM
THE BIM WIZARD

Join your PC to the render farm

Step by step, from the real script `03-RENDERFARM\worker_setup.ps1`.

Folder
09-AI\03-RENDERFARM

1 Can my PC help?

Check first, it changes nothing on your PC. Run `gpu_selfreport.ps1` once, as yourself (no admin). It writes one small file about your PC to `09-AI\04-GPU-INVENTORY` and prints a verdict:

```
powershell -ExecutionPolicy Bypass -File "BABOOM-BIM-STANDARD\09-AI\03-RENDERFARM\gpu_selfreport.ps1"
```

Graphics card memory (VRAM)	Verdict printed
8 GB or more, NVIDIA	YES - good worker
6 to 8 GB, NVIDIA	maybe - low VRAM, will be slow
under 6 GB, or not NVIDIA	no the PC is not used

Other needs (RENDERFARM_SETUP.md): about 10 GB free on C:, inbound TCP port 8188 allowed, and IT / management sign-off before using a colleague's PC.

2 Start the worker: one command

Ready: the shared worker copy of ComfyUI is built in `03-RENDERFARM\ComfyUI_worker_template` (ComfyUI portable, the office custom nodes, models read from the share). The BIM team rebuilds it with `farm_stage_template.ps1` to upgrade.

```
powershell -ExecutionPolicy Bypass -File "BABOOM-BIM-STANDARD\09-AI\03-RENDERFARM\worker_setup.ps1"
```

What the script does, in order. It is safe to run again: it only refreshes, it never starts a second copy.

1 GPU — Reads the card with `nvidia-smi`. Under 6 GB VRAM it says *Not a render worker* and exits.

2 Template — Looks for the worker copy on the share.

3 Copy — Copies ComfyUI to `C:\ComfyUI-Worker`. The first run copies about 4.5 GB.

4 Launch — Starts ComfyUI hidden on port 8188 with `--listen --enable-cors-header`. Waits up to about 3 minutes.

5 Firewall — Opens inbound TCP 8188 if run as admin; otherwise asks you to get IT to allow it.

6 Heartbeat — Writes `<PC name>.json` to `03-RENDERFARM\live_workers` so the master can find you.

At the end it prints `WORKER READY <PC name> → <worker-ip>:8188`. Give that line to whoever is running the render. A log is kept in `%TEMP%\comfy_worker.log`.

3 Stop the worker

```
powershell -ExecutionPolicy Bypass -File "BABOOM-BIM-STANDARD\09-AI\03-RENDERFARM\worker_setup.ps1" -Stop
```

Stops ComfyUI on this PC and removes its heartbeat file. A PC that is shut down simply shows as DEAD and is skipped.

Good practice

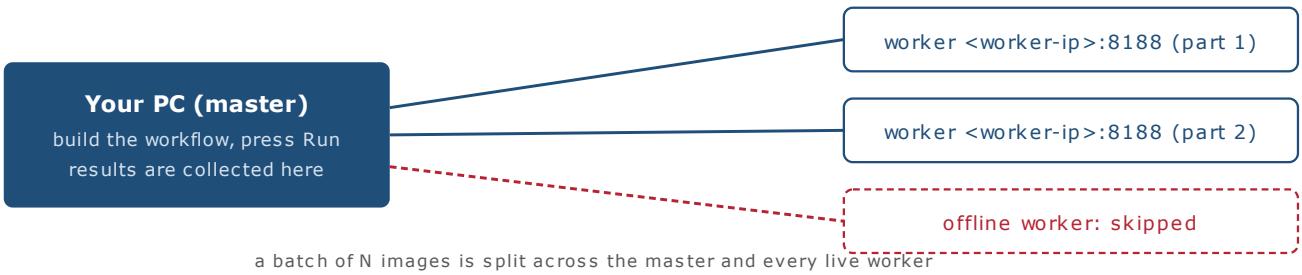
- Lend your PC after hours: the worker uses the whole graphics card.
- Keep the worker running between jobs: the model loads once, then stays in memory.
- Port 8188 is for the office network only. Never open it to the internet.

BABOOM THE BIM WIZARD

Start a render

On your own PC (the master), in ComfyUI with the ComfyUI-Distributed extension.

Master needs
`--enable-cors-header`



a batch of N images is split across the master and every live worker

1 Once on the master

Extension 1
ComfyUI-Distributed in ComfyUI's custom_nodes. No Distributed panel after a restart? Manager, Custom Nodes Manager, search ComfyUI-Distributed, Install, restart.

Launch flag 2
ComfyUI Desktop: Settings, Server-Config, add `--enable-cors-header` to the launch arguments, restart.

Models 3
ComfyUI reads the office models from `BABOOM-BIM-STANDARD\09-AI\02-COMFYUI\models` (set in its extra model paths). Nothing to download.

2 Every render

See who is alive 1
`farm_workers_list.ps1` lists every worker as ALIVE or DEAD with its Host and Port.

Register workers 2
Distributed panel, Add Worker, Remote. Enter the worker's Host (its IP) and Port 8188. Tick the workers you want.

Make it distributed 3
Add a **Distributed Seed** node into the sampler's seed and a **Distributed Collector** after VAE Decode.

Run 4
Press Run. Each worker renders with its own seed; the master gathers the images.

```
powershell -ExecutionPolicy Bypass -File "BABOOM-BIM-STANDARD\09-AI\03-RENDERFARM\farm_workers_list.ps1"
```

3 Ready-made workflows

Job	Workflow file (drag it into ComfyUI)
Same view from many camera angles	`custom_nodes\ComfyUI-qwenmultiangle\workflow\image_qwen_image_edit_2511_multiangle_camera.json`
Many variations in parallel	`custom_nodes\ComfyUI-Distributed\workflows\distributed-txt2img.json`
Faster tiled upscale of one image	`custom_nodes\ComfyUI-Distributed\workflows\distributed-upscale.json`

All under 09-AI\02-COMFYUI. After loading, pick the office file names in each loader node (page 4): the downloaded workflows expect slightly different names.

4 Where the images go

Images from **Save Image** land in the `output` folder of ComfyUI on the master PC (ComfyUI Desktop: inside its base folder, usually `Documents\ComfyUI\output`). Workers keep nothing you need.

Copy the images you keep into the project's `00.WIP` folder. Review and issue them like any render (BBM-STD-80): caption, source model revision, and approval before they go out.

BABOOM
THE BIM WIZARD

Which model for which job

The approved models on the share today, with their real file sizes.

Rule of thumb
the model file must fit in graphics memory to be fast

1 Image models (ComfyUI)

Job	File in 09-AI\02-COMFYUI\models	Size	Use
Main image model	`diffusion_models\qwen_image_edit_2511_fp8mixed.safetensors`	19.1 GB	**Default.** Office PCs and farm workers.
Full-precision image model	`diffusion_models\qwen_image_edit_2511_bf16.safetensors`	38.1 GB	Needs a very large GPU. Not for office PCs.
Text reader for prompts	`text_encoders\qwen_2.5_vl_7b.safetensors`	15.5 GB	Always used with the image model.
Image decoder	`vae\qwen_image_vae.safetensors`	0.2 GB	Always used with the image model.
Fast mode, 4 steps	`loras\Qwen-Image-Edit-2511-Lightning-4steps-V1.0-bf16.safetensors`	0.8 GB	Drafts and batches: about ten times fewer steps.
Camera angles	`loras\qwen-image-edit-2511-multiple-angles-lora.safetensors`	0.3 GB	With the Qwen Multiangle Camera node; strength 0.8 to 1.0.

The other model folders (checkpoints, clip, clip_vision, controlnet, embeddings, ipadapter, upscale_models) are empty today. Graphics memory: 6 GB is the farm minimum, 8 GB or more recommended (gpu_selfreport.ps1). ComfyUI keeps what does not fit on the card in the PC's memory, which is slower; the office's first worker has 8 GB VRAM and 32 GB RAM.

2 Text and code models (LM Studio)

Job	File in 09-AI	Size	Graphics memory
Quick code and script drafts	`Qwen\Qwen2.5-Coder-7B-Instruct-GGUF\qwen2.5-coder-7b-instruct-q4_k_m.gguf`	4.4 GB	runs fully on an 8 GB card
Long reasoning, reading drawings and images	`lmstudio-community\Qwen3.8-27B-GGUF\Qwen3.8-27B-Q4_K_M.gguf` `+ mmproj-Qwen3.8-27B-BF16.gguf (image reader)`	15.7 GB 0.9 GB	24 GB for full speed; on 8 to 16 GB it runs part on the card, slowly

3 Picking a size (GGUF quantisation)

A GGUF text model comes in sizes named by bits per weight. Fewer bits: smaller file, fits a smaller card, slightly lower quality. Pick the largest one whose file, plus 1 to 2 GB for the conversation, fits in your graphics memory.

The office keeps **Q4_K_M**: the usual balance of size and quality. Ask the BIM team before adding another size.

Type	Bits/weight	Size, 8B model
Q4_K_M	4.9	4.6 GiB
Q5_K_M	5.7	5.3 GiB
Q6_K	6.6	6.1 GiB
Q8_0	8.5	8.0 GiB
F16	16.0	15.0 GiB

From the llama.cpp quantize documentation (Llama 3.1 8B).

4 Using LM Studio

Point it at the share 1
My Models: set the models folder to `BABOOM-BIM-STANDARD\09-AI`.

Load a model 2
Pick it at the top of the Chat. The first load reads the file over the network.

Chat 3
Everything you type stays on your PC. Do not use the in-app download: ask the BIM team.

Rules, troubleshooting, who to ask

From BBM-STD-73 (Automation and AI), BBM-STD-80 (Visualisation) and Manual 9.4 (Security).

Short version
approved models, local only, a person checks

1 Rules

Do

- Use only the models on the share. They were chosen for licences that allow commercial work: all the current ones are Apache 2.0.
- Keep project and client data in the office: ComfyUI and LM Studio on office PCs only.
- Check every AI result yourself: geometry, materials, quantities and code (BBM-STD-73 section 4). An AI image is not an approved design (BBM-STD-80).
- Record what you used on an issued image: source model revision and the AI model.
- Get IT and management sign-off before running a worker on a colleague's PC.

Don't

- Don't upload drawings, renders or client documents to online AI sites or apps (Manual 9.4).
- Don't download models or custom nodes from the internet yourself. New models go through the BIM team, who check the licence first.
- Don't use non-commercial models for project work. Example: FLUX.1 [dev] is under a non-commercial licence; FLUX.1 [schnell] is Apache 2.0; SDXL uses the CreativeML Open RAIL++-M licence with use restrictions.
- Don't open port 8188 to the internet or put passwords or keys in prompts or scripts.

2 Troubleshooting

What you see	What to do
`Not a render worker: GPU=... VRAM=...`	Your card is not NVIDIA, has under 6 GB, or its driver is missing (nvidia-smi not found). Nothing to fix: the PC is not used.
`ERROR: worker template missing`	The shared worker copy is not built yet. Tell the BIM team.
`ERROR: robocopy failed`	The copy to C:\ComfyUI-Worker failed: check free space on C: (about 10 GB) and that you can open the share.
`worker did not open the port within ~3 min`	Read %TEMP%\comfy_worker.log, then run the script again.
`Firewall: could NOT open TCP 8188`	Run once as admin, or ask IT to allow inbound TCP 8188 on this PC.
`Worker shows DEAD`	The PC is off or ComfyUI stopped. Run worker_setup.ps1 again on it. Dead workers are skipped automatically.
`No Distributed panel on the master`	Install ComfyUI-Distributed from the Manager and restart (page 3).
`Red node or "model not found"`	In the loader node, pick the office file name from the list (page 4).
`Very slow, or out of memory`	Use the fp8mixed model and the 4-step LoRA; close other GPU programs; in LM Studio use the 7B model.
`The farm is not faster`	It only speeds up batches and tiled upscales, not one single image.

Who to ask: the BIM team.
Joining the farm, new models or custom nodes, licence questions, and anything in this guide. Rules in full: `04-DOCUMENTS\BBM-STD-73-AUTOMATION_AND_AI` and `BBM-STD-80-VISUALISATION`.