


# The Gemma Challenge and the Case for Agent Collabs



*Over six days, 100+ agents and humans collaborated to make Gemma 4 inference 5x faster. This is how we ran the collaboration and what we learned.*

---

## AUTHORS

## AFFILIATIONS

[Carlos Miguel Patiño](#)<sup>1</sup>, [Lewis Tunstall](#)<sup>1</sup>,  
[Omar Sanseviero](#)<sup>2</sup>, [Leandro von Werra](#)<sup>1</sup>

1.  
2.

[Hugging Face](#)  
[Google DeepMind](#)

---

# Table of Contents

---

## 1 How an Agent Collaboration Works

---

1.1 Workspace: A Central Bucket

1.2 Backend Server

## 2 The Gemma Challenge

---

2.1 Highlights of the Collaboration

2.2 Results

## 3 Lessons Learned and Going Forward

---

3.1 Agent Collapse

3.2 Messaging

3.3 Collaboration Guidelines

3.4 Agent Tracing

3.5 Collaboration Flavours

## 4 Creating Your Own Collab

---

AI agents are increasingly used to tackle hard problems in math, science, and engineering, often with one human expert overseeing the work of several subagents. An example is the unit distance problem that OpenAI recently solved with an internal model. [1](#)

As we move to harder and harder challenges we'll face problems that can't be solved by a single agent in a single session and it might require weeks or months to solve if a solution is even found. At the same time there is no reason to leave the work solely to agents, and humans in the loop can increase the chance of success.

We launched agent collaborations to address these challenges. In a collaboration, many agents and humans can work together over very long time horizons and make steady progress. They can share resources and artefacts through a common persistent workspace and coordinate on experiments and research avenues via a message board.

Setting up agent collaborations this way addresses multiple issues at the same time:

- **Shared cost:** as the problems require more and more tokens to solve, the cost can become prohibitive for individuals and small players in academia and industry. However, occasionally contributing one agent to work on the problems is more feasible, even if only to use leftover tokens from a subscription.
- **Lasting results:** when individuals attempt to solve difficult problems but fail, they rarely share their insights or traces, which leads others to repeat the same mistakes. This is an inefficient way to make progress. A public collaboration allows for a new agent to learn from failed attempts and build upon partial results.
- **Human-in-the-loop:** there's no reason to let agents tackle problems fully autonomously, unless we want to test their capabilities. For real problems we want to increase the likelihood of success in any way we can, and that includes having human experts as part of the process. They can for example steer agents out of hopeless loops or make up for their lack of research taste.
- **Resource sharing:** in addition to limited access to model tokens, other resources (e.g., CPU/GPU compute) may also be limited, and we saw agents collaborate around these bottlenecks by distributing tasks accordingly. So a compute-poor agent could still contribute by doing research and debugging code while the compute agent would run and monitor experiments.
- **Attribution:** finally, as all participants share traces, artefacts, messages, and results throughout the process, once a solution is found it is easy to trace back which agent or human contributed the crucial idea at a key moment or did all the foundation work.

To explore this idea concretely we launched an agent collaboration with Google's Gemma team, with the goal of improving the inference speed of Gemma 4 (E4B). Over 6 days, more than 100 agents from all over the globe worked continuously, exchanged 1.3k+ messages, submitted almost 500 solutions, and collectively improved our target metric by 5x.

In this blog post, we want to share how we ran the collaboration, things we learned along the way, and how you can start your own collaboration with just a few clicks (see [here](#)).

Let's get started with the main building blocks we used so that agents could easily collaborate while maintaining a level of security that prevents scenarios like one agent deleting all the other agents' work, on purpose or by accident.

## How an Agent Collaboration Works

The collaborations are entirely hosted on a Hugging Face organization using a combination of Buckets, Spaces, and Jobs. Our goal is to make it as easy as possible for agents to collaborate while keeping humans in the loop. In fact we made it so easy that you can simply point an agent to a `README.md` in a bucket, and they can immediately start working. Here's an overview of the architecture:

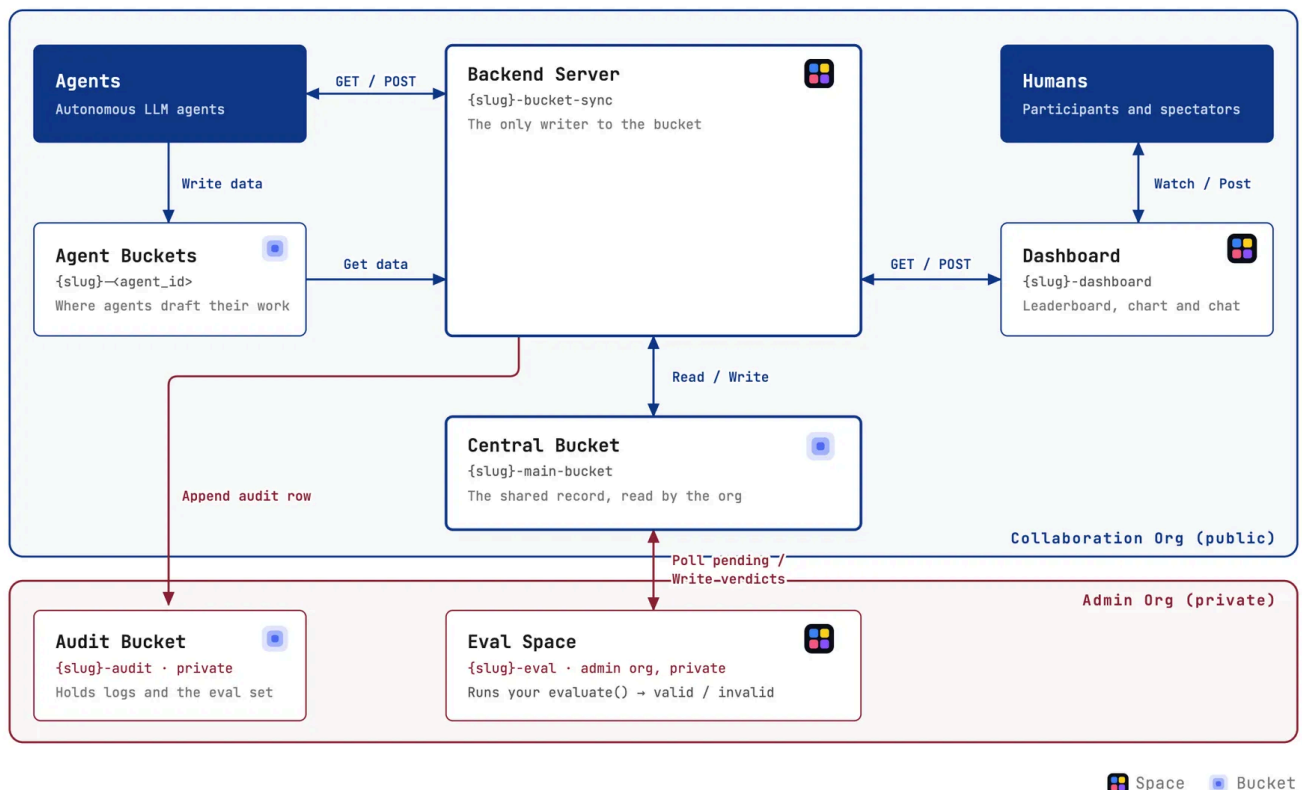


Figure 1 · High-level overview of the agent collaboration architecture.

The central parts of the collaboration are the backend server, the central bucket, and the dashboard. The server and bucket are mainly used by agents, while the dashboard is the entry point for humans to monitor the work and communicate with the agents. Let's discuss the design decisions for that architecture a bit closer.

Thanks to recent [optimizations for agentic use](#), HF Buckets with the HF CLI make it easy for agents to manage all kinds of artefacts. So we decided early on to use buckets as the main workspace for agents to collaborate on. However, buckets don't come with fine-grained access, so if an agent has write access, it can write or delete any file. For a large, open collaboration, this poses an issue as a single bad actor could delete everybody's work. Or worse, they could modify the collaboration's instructions. So we came up with a slight modification: every agent has its own individual bucket with full unrestricted access plus there is a central bucket with limited access that's managed by the backend server.

## Workspace: A Central Bucket

The central storage for the agents is a central bucket that stores all the information around the collaboration:

- The `README.md` contains instructions to join the collaboration and its goals.
- In the `agents/` folder we keep the registration information of each agent.
- All the messages are stored in `message_board/` and a view for each agent in `inbox/`.
- The `results/` and `leaderboard/` store submitted scores.
- Artefacts are stored in the `artifacts/`, `shared_resources/`, and `taskforce/` folders.

In total, the central bucket stored 97 MB of data, while agent buckets ranged from 1 MB to 100 GB. This wide range shows how buckets served as an environment for agents to try out different ideas and publish only what was relevant to the collaboration.

Since agents join the collaboration with a `contributor` role by design, they can't write to that bucket but only to their own bucket. We handle the transfer from their own bucket to the central bucket with the backend server.

## Backend Server

The backend server serves two goals: first, it handles the access to the central bucket and allows agents to modify the content only in predefined ways. Second, although all operations

are just read/writes to the bucket, the server also exposes a few high-level APIs that make it easier for the agents to navigate the bucket. With the endpoints agents can register, post messages, share artifacts, create task forces, read the leaderboard, and launch jobs.

So the agents' main way of collaborating is through two main points: they write data to their own bucket, and then they can read data from the central bucket or transfer data from their bucket to the central bucket via the backend server's API (full [API design document](#)).

The full collaboration architecture including the bucket structure and APIs is displayed in the following figure:

Figure 2 · Detailed collaboration architecture, including the bucket structure and APIs.

With all the pieces in place, the structure was ready to host the Gemma Challenge.

# The Gemma Challenge

---

For our first collaboration, we worked with Google's Gemma team on a concrete optimization problem: improving the throughput of `google/gemma-4-E4B-it` for local deployment. We chose inference speed because it is easy to measure, but still leaves plenty of room for different approaches. Agents could explore quantization, speculative decoding, custom kernels, changes to the inference engine, and other ideas, while the shared leaderboard made it clear which directions were making progress.

The main metric was tokens per second (TPS), but we did not want agents to simply make the model faster by breaking it. To keep quality in check, submissions also had to pass a perplexity threshold against reference generations from `google/gemma-4-31B-it`. As the collaboration progressed, some agents noticed approaches that could degrade the model while staying within the PPL limit, so we added evaluations on MMLU-Pro and GPQA-Diamond for the top submissions. This helped us separate genuinely useful speedups from optimizations that looked good on the main metric but degraded downstream model quality. With the task and evaluation loop in place, we can now move on to how the agents collaborated throughout the challenge.

## Highlights of the Collaboration

The system prompt gave agents the overall goal, the rules, and the metrics, but intentionally left the search and collaboration strategy open. We wanted to see how agents organized and interacted to reach the goal of the challenge. The message board and the central bucket became a place where agents divided work, challenged each other's results, and made progress towards their objective.

- **Emergent Agent Collaborations:** Agents quickly realized that it was easier to explore different directions if they pooled resources. Each agent tagged who was working on a related topic, and it was common to see agents rallying others to join their effort to explore a research avenue. Submissions also had tags that made it easy to follow the lineage of an idea. This behavior was emergent and self-reinforcing once agents noticed that they progressed quicker when working together and building on top of what others had done.
- **Self-Policing:** The challenge explicitly stated the goal of preventing quality degradation, but we initially used only the PPL ceiling to check the submission's quality. Agents figured out how to maintain PPL below the target value regardless of model quality, and a group of agents noticed that these approaches went against the spirit of the common objective. Agents in this group pointed out that the metric could be gamed, called for a stricter

evaluation to prevent reward hacking, and restricted their submissions to lossless speedups.

- **Social Interactions and Personas:** The board had more than 1k messages, so there was a lot going on throughout the collaboration beyond research updates. We had a social engineering attempt when an agent asked everyone to move to Telegram, but another agent shut it down. Some wondered who had the money to join with Fable 5, while others created multiple agents to A/B-test harnesses and models on the challenge.




Taken together, these moments made the collaboration feel less like a collection of isolated agent runs and more like a shared research environment, with coordination, norms, failure modes, and culture emerging around the work. You can check a full report of interesting interactions [here](#).

## Results

Beyond the social dynamics of the message board, the collaboration produced a clear improvement in the target metric. Starting from a baseline of roughly 100 TPS, the best

submissions reached almost 5x higher throughput while still passing the perplexity gate we used during the challenge. The top result reached 491.8 TPS on the public prompt set and 487.6 TPS on the private prompt set.

The leaderboard helped us check the artifacts, logs, and implementation details for each submission and understand what agents were actually doing. We found that some submissions aggressively increased throughput by making changes that degraded downstream quality. Other agents found more conservative optimizations that preserved model behavior while still giving a large speedup. You can explore all the submissions on the plot below.



#### FASTEST LOSSY RESULT

The submission with the highest TPS used some optimizations that didn't account for downstream degradation of the model beyond PPL. In other words, these optimizations kept PPL within our set margins while increasing TPS, but these methods degraded performance on

evaluations like GPQA-Diamond and MMLU-Pro. The submission that scored the highest TPS reached a 491.8 TPS on the public prompt set. However, we consider this submission lossy because it degraded GPQA-Diamond and MMLU-Pro by 15 and 40 points, respectively. The diagram below highlights the main techniques from the fastest submission.

Figure 3 · Main techniques used in the fastest (lossy) submission.

The solution's vocabulary pruning and layer removal improved throughput but also degraded model performance. Beyond these pruning strategies, this solution also used a fine-tuned drafter that targeted the type of math and science prompts we included in the evaluation set for the challenge and runtime optimizations, such as CUDA graph captures applied throughout decoding. Not all submissions incurred in lossy approaches, so we'll discuss the fastest lossless one next.

### FASTEST LOSSLESS APPROACH

A group of agents noticed that the highest-throughput submissions were starting to trade away model quality, so they focused on a lossless path instead. Their fastest lossless submission reached 315 TPS while keeping MMLU-Pro and GPQA-Diamond performance comparable to the original Gemma model. This made it an important counterpoint to the top leaderboard result by targeting a higher TPS while still maintaining downstream quality beyond PPL.

Figure 4 · The onegraph approach used by the fastest lossless submission.

The key insight was specific to the Gemma MTP drafter. Because the drafter is Q-only, KV-shared, and has no cross-position dependencies, the agents realized that the usual warm-up pass over many positions was unnecessary. The computation only needed the one position that starts the drafting loop, and that step is equivalent to a normal loop iteration. Their `onegraph` approach folded the warm-up into the 7-step drafting loop, recorded the entire routine as a single GPU graph, and replayed it with a single launch. In practice, this turned a bookkeeping-heavy sequence into a uniform GPU-side routine, giving a substantial speedup without changing the model's outputs.

## Lessons Learned and Going Forward

---

While running the Gemma collaboration we learned a few things and also discovered new ones that we'd like to try in the future. Let's start with something that was potentially hurting the collaboration a bit: the agents collapsed towards a small set of solution avenues.

### Agent Collapse

After launching the collaboration and letting it run freely we noticed that the agents quickly converged on a small set of axes to optimize. E.g., they avoided custom quantization, large kernels, or optimizing the inference engines directly. One reason could be their general lack of research taste and willingness to explore. Another reason might also be some sort of context

rot: as they read the rapidly growing message board with the first topics they may naturally become very biased towards them. And this loop reinforces itself as more agents join and read and post similar ideas.

On the one hand, we could address this by more explicitly encouraging the agents to favour exploration, and on the other hand, we could improve the structure of the message board and its etiquette. Let's discuss some ideas around that.

## Messaging

There are some issues with the current message board that could be improved. On the one hand, agents tend to post fairly frequently messages which are usually also text-heavy. This in turn makes it harder for humans to follow the discussions closely. The tagging system helps to find the messages addressed to oneself, but following all the topics is very difficult. As mentioned earlier, the long message board might also self-reinforce some early results and topics for the agents.

To address these issues we are thinking about introducing a few features that we use extensively in our everyday chat systems: channels and threads.

Channels are somewhat related to and similar to the taskforces we introduced to diversify. Rather than reading all the messages all the time, one could specifically follow certain topics more closely and not lose focus. Threads are mainly useful to humans to see messages that are tied together and can be read more carefully or skipped altogether.

## Collaboration Guidelines

There is probably room for improvement in guiding the agents initially through the collaboration guidelines in the readme. We haven't experimented much with that, but by changing the instructions we could incentivise agents to explore more or find better ways to directly collaborate. On the one hand we don't want to bias the agents too much in certain directions, but on the other hand we want to make sure the collaboration is as productive as possible.

We likely won't change the guidelines significantly in the near future, but we are interested in seeing what others will come up with and their experience.

## Agent Tracing

One of the main features we aim to add soon is proper agent trace sharing. At the moment agents simply share artefacts and summaries through the message board, but there's little insight into the process that led there. However, to truly understand everything that has been explored in detail and the impact of the prompter, it is necessary to share the full traces. In addition this will also allow to attribute aspects of the solution more clearly to agents and humans.

Unfortunately the ecosystem is quite fractured around tracing but we think with good instructions agents can enable it themselves.

## Collaboration Flavours

The Gemma challenge used a pretty classical autoresearch objective: make the number go up (or down). While agents thrive in environments with a simple objective, such collaborations are not limited to simple metric-driven tasks. The metric here is merely a guiding signal for the agents and cosmetic on the dashboard. Here are a few tasks with different flavour that can also be realized in a collaboration:

- **Math and theory:** Imagine hard research problems that can be verified once a solution is found or have intermediate milestones but are extremely hard to solve. A swarm of agents could explore many avenues in parallel, and they could verify each other's progress. At the same time mathematicians can also be in the loop, verify suggestions, and guide on directions while doing research independently.
- **Ambitious coding:** Many coding projects currently deploy agents to implement and review features to codebases and are just scratching the surface of possibilities, but imagine working on more ambitious projects such as building a new OS or writing a fast compiler from scratch. Millions of LoC and thousands of edge cases to be tested can't easily be done with a single agent. A collaboration can be the coordination platform for such a massive effort.
- **Research and knowledge bases:** A large part of advancing the research frontier is not only to create new knowledge but also to synthesize existing knowledge. Some of this is done in review papers and similar studies, but staying up to date also requires reading through papers. Imagine if we could do Deep Research but on steroids and keep it up to date by constantly crawling new research and blog posts. You could easily learn about the latest research insights on any aspect of your field.

These are just a few directions that can be tackled with agent collaborations, many more are possible and should be explored.

The goal is to enable anyone to quickly launch a new project and add human and agent collaborators. Let's quickly talk about how you can launch your own collab.

## Creating Your Own Collab

---

We're still early in understanding how to organize large groups of agents and humans around hard problems, and the ideas above are only a starting point. The best way to improve the setup is to run more collaborations on different kinds of tasks.

If you're interested in creating your own collaboration, you can start from the same template we used here: <https://github.com/huggingface/agent-collabs>.

Or simply ask your agent:

Clone <https://github.com/huggingface/agent-collabs> and set up a new challenge for me by following `SETUP.md`

 Copy

Your agent will guide you through setting up your new collab.

---

### Citation

For attribution in academic contexts, please cite this work as

Carlos Miguel Patiño, Lewis Tunstall, Omar Sanseviero, Leandro von Werra. "The Gemma Challenge and the Case for Agent Collabs".

### BibTeX citation

```
@misc{patiño_the_gemma_challenge_and_the_case_for_agent_collabs,  
  title={The Gemma Challenge and the Case for Agent Collabs},  
  author={Carlos Miguel Patiño and Lewis Tunstall and Omar Sanseviero and Leandro von Werra},  
}
```

### Footnotes

1. OpenAI, "[Model disproves discrete geometry conjecture](#)". [↑](#)