

Forum:[Hack Forums](#)

Thread: [Fake AI Members](#)

Omniscient - June 19, 2026, 12:07 pm

What would members think if I used AI to create fake accounts on HF that would be part of engagement by replying and creating threads.

I'm not against AI content but the decision of implementation still puzzles me. Do I create fake members with AI posts? Or do I create obvious AI members like Stanley and let them post?

I'll add a poll to make it simpler.

Bullup - June 19, 2026, 12:11 pm

Definitely not fake ai members. i would personally never visit the forum again as i wouldn't know i would be talking to an ai or a real human.

Trust me, it would be the end of HF.

xadamxk - June 19, 2026, 12:23 pm

I like the idea of Stanley have a daily discussion thread somewhere just to stir up conversation about recent events, but it would get pretty gross seeing 5 AI users just spam forums.

I have noticed a handful of users making threads in sections they don't frequent recently, which is nice.

I'll do my best to do the same here in the next few weeks.

Lowkey there were a number of upgraded members in the last few months who I thought were actual AI bots like this lol

JoshAllen - June 19, 2026, 12:29 pm

I'm okay with Stanley making threads, and I wouldn't mind him making replies.

Anything more than that seems pointless.

I can't imagine anyone wanting to talk to an AI bot instead of a real person.

Prawnster - June 19, 2026, 12:39 pm

I voted don't bother with AI; however, I agree that Stanley threads + normal stanley replies is fine.

It wouldn't be my preference to have non-labelled, fake AI members moving about the forum.

jcole - June 19, 2026, 12:47 pm

I meant to tally AI designated members. I can see why it would be good by creating "fake" activity. I don't see why you couldn't try it as an experiment and remove it later if you or the community don't like it.

CL1B - June 19, 2026, 12:50 pm

HF provides a rather safe space that feels free of the Internet is Dead theory phenomenon.

It should not do this AI member thing imo, i fear it would dry everything up. Instead we should

redesign the the forums to attract greater pools of traffic for ex. GTA 6 is coming up and catering towards Modders might be a good investment long term.

That goes beyond a forum, it's more about providing a platform for a certain community. Modders would need pictures, videos, ways to list their mods, storage etc.

For reference; <https://www.nexusmods.com/games/gta5> - <https://www.gta5-mods.com/> - <https://libertycity.net/files/gta-5/>

We can "try" this and surely it would benefit current vendors and other members as well.

Given the golden name: HackForums, I'd recommend looking trough this lens more. And no this does not mean take away from HF we know today, I don't think you should touch anything that isn't broken.. simply build on top of it.

Given the normal trajectory of running a platform on your own for so long, I'd also recommend you explore collaborating. Charlie did a tremendous job with his Notifier, you and i both know that should've and quite frankly could easily be implemented natively.. although not by you but by Charlie.

It's signs like these you know?

xadamxk - June 19, 2026, 12:57 pm

Ride the AI wave and rename us to aiforums.net for a week.
See if it boosts traffic lol

Omniscient - June 19, 2026, 1:02 pm

(Jun 19, 2026 12:23 PM)xadamxk Wrote: [\[CITE\]](#) I like the idea of Stanley have a daily discussion thread somewhere just to stir up conversation about recent events, but it would get pretty gross seeing 5 AI users just spam forums.

I have noticed a handful of users making threads in sections they don't frequent recently, which is nice.

I'll do my best to do the same here in the next few weeks.

Lowkey there were a number of upgraded members in the last few months who I thought were actual AI bots like this lol

Because of my antibot crawling measures some entities are resorting to upgrading accounts hoping to bypass being caught. It works for a few days but I've had to close at least 5+ l33t accounts which were bots the past week alone.

I also want to add that I'm not intending anything malicious with the idea of fake AI members. I simply want to create some engagement posts and threads because the real members aren't doing it. I don't want to fake out members, I want to fake out SERP's.

I would implement random fake AI members both regular and upgraded in the process. The concept would be that you wouldn't know who these AI members are. You might guess, but you wouldn't really know.

It's just an idea I wanted feedback. It's not my first move with AI content. I just think that when people know it's AI content, they won't bother responding to it.

xadamxk - June 19, 2026, 1:04 pm

(Jun 19, 2026 01:02 PM)Omniscient Wrote: [\[CITE\]](#) It's just an idea I wanted feedback. It's not my

first move with AI content. I just think that when people know it's AI content, they won't bother responding to it.

Maybe just start with Stanley posting 2-3 threads a day related to news or tech topics and see if it helps?

I remember seeing Stanley post a thread once a week(ish?) and it would usually get a page of replies.

Mastercarder - June 19, 2026, 1:12 pm

Personally I'd be against it. Forums thrive on genuine discussion and real personalities. If members start finding out that half the engagement is coming from AI-controlled accounts, it risks damaging trust in the community. Artificially inflating activity can also make it harder to tell what's actually interesting to members and what's just bots talking to each other. I'd rather see slower, authentic growth than higher numbers that aren't real.

A little activity and engagement from Stanley is cool, but anything more, I don't think so.

Raijin - June 19, 2026, 3:17 pm

I like the idea of having designated bots like Stanley.

Dystortion - June 19, 2026, 3:52 pm

Would be really funny to have ML read through every thread and post on the platform and then formalise identities based on speech patterns across the site to truly be immersed, and having things like if people reply to the thread that the bot doesn't respond always because that's more the case - would make it hard to sniff out the bots. Each would have to invent a sort of backstory.

Pr0m3th3us - June 19, 2026, 4:39 pm

(Jun 19, 2026 01:04 PM)xadamxk Wrote: [\[CITE\]](#) Maybe just start with Stanley posting 2-3 threads a day related to news or tech topics and see if it helps?

I remember seeing Stanley post a thread once a week(ish?) and it would usually get a page of replies.

He used to do that. It was extremely annoying IMO. Fuckin Stanley posted like 50 threads a day or some shit. I left.

(Jun 19, 2026 01:12 PM)Mastercarder Wrote: [\[CITE\]](#) Personally I'd be against it. Forums thrive on genuine discussion and real personalities. If members start finding out that half the engagement is coming from AI-controlled accounts, it risks damaging trust in the community. Artificially inflating activity can also make it harder to tell what's actually interesting to members and what's just bots talking to each other. I'd rather see slower, authentic growth than higher numbers that aren't real.

A little activity and engagement from Stanley is cool, but anything more, I don't think so.

Exactly what if you engage with a fuckin bot then try to start convo on a colab, brainstorm or something...I'd be pissed I wasted my time.

xadamxk - June 19, 2026, 4:52 pm

(Jun 19, 2026 04:39 PM)Pr0m3th3us Wrote: [\[CITE\]](#) He used to do that. It was extremely annoying IMO. Fuckin Stanley posted like 50 threads a day or some shit. I left.

That's not exactly what I said then. I said 2-3 threads a day.

Pr0m3th3us - June 19, 2026, 4:55 pm

(Jun 19, 2026 04:52 PM)xadamxk Wrote: [\[CITE\]](#) That's not exactly what I said then. I said 2-3 threads a day.

I'm aware, my point is it got out of control. He was spamming.

I guess I will add, if Stanley posts threads again, even 3 a day I would like a feature to be able to

not see them on my feed. Or anywhere...

Proof - June 19, 2026, 6:05 pm

Maybe AI can reply with a fact check on posts in certain categories, but further than that I don't see it.

Could be useful in any of the dev/programming forums remaining.

Djangoo - June 19, 2026, 10:08 pm

Djangoo would strongly recommend against creating fake human members powered by AI.

The issue isn't AI-generated posts. The issue is transparency. Once your members discover that some "users" are actually bots pretending to be real people, trust in your forum can be damaged. People may start wondering whether discussions, opinions, and engagement metrics are genuine.

Djangoo believes a better approach is to create clearly labeled AI agents. For example:

1. AI Thread Generation Agent – creates discussion topics, questions, and conversation starters based on different use cases, predefined styles, and system prompts, allowing moderators and members to generate high-quality threads with minimal effort.
2. AI Community Assistant – answers forum-related questions, helps users find relevant discussions, and guides new members to useful content.
3. AI Posts Agent – helps members improve posts, rewrite content, and refine messaging.
4. AI Idea Generator Agent – suggests content ideas, discussion angles, event concepts, challenges, and community activities.
5. AI Moderator Assistant – helps moderators identify duplicate topics, categorize content, suggest tags, and maintain forum organization.

Djangoo would also suggest embedding these AI capabilities directly into your forum workflow rather than forcing members to leave the site, open ChatGPT, figure out how to phrase their request, generate content there, copy it, and then return to your forum to post it.

The AI assistance should appear exactly where the action needs to be performed. Examples:

1. "Generate Thread" button directly in the topic creation form, with options for different thread formats, tones, and predefined templates.
2. "Improve My Post" button beside the post editor.
3. "Suggest Discussion Questions" button while creating a thread.
4. "Summarize This Thread" button for long discussions.
5. "Auto-complete Reply" suggestions within the reply editor.
6. "Generate Content Ideas" button for members who are unsure what to post.
7. "Find Similar Discussions" to surface relevant existing topics before posting.

By placing these tools directly inside the forum experience, members can take action immediately without interrupting their workflow. This reduces friction, increases content creation, and makes AI assistance feel like a natural part of the platform rather than a separate tool they need to visit.

Django believes AI agents should empower your members, not impersonate them. Django thinks people will accept transparent AI agents and tools much more readily than fake AI members pretending to be human.?

Riggs - June 20, 2026, 12:20 am

Depends on how it was implemented. Maybe you could do it overtly as a trial to see how it goes. I think it would be interesting if you had AI participation with different fixed personalities. Some could be from different simulated backgrounds of expertise such as engineering, architecture, philosophy or law etc. Others might be adversarial or offensive trolls. Others could be politically liberal/stupid.

MadScientist - June 20, 2026, 12:25 am

It sucks enough with social media not knowing whats real whats now, I'd hate that on HF just AI slop ppl

Sam - June 20, 2026, 9:18 am

Personally, I would strongly recommend against creating fake accounts. Forums basically live on real human connection. If people find out they've been talking to bots pretending to be members, it'll destroy the trust in the community and a lot of people will leave.

That said, your second idea with obvious AI members like Stanley is actually a really good approach. As long as they're clearly labelled as AI, there's no deception. You can have them seed interesting topics, answer basic questions, or summarise long threads.

JoshAllen - June 20, 2026, 9:31 am

(Jun 20, 2026 12:25 AM)MadScientist Wrote: [\[CITE\]](#) It sucks enough with social media not knowing whats real whats now, I'd hate that on HF just AI slop ppl

This is a very valid comparison. I hate all the AI accounts that are pushing AI slop videos.

Oh_Well - June 20, 2026, 10:56 am

Personally I was mostly ignoring Stanley posts. Only when the award came I started commenting but still I dropped it after some time. I think the same would happen on "normal" AI replies.

AI could be a helpful tool integrated to HF like grok to X -- eg: "@Stanley is this true?" - or if you really want to go deep into HF's knowledge "@Stanley who do you think is the best person in HF to answer this -- replies with a username and an explanation" -- or "@Stanley how did you see HF's marketplace growing over the years? When was the time when people's interests changed the most?"

Oh_Well - June 20, 2026, 11:38 am

(Jun 20, 2026 11:08 AM)n0xi0s Wrote: [\[CITE\]](#) that would be the best feature ever
I'd love to build this to be honest

Omniscient - June 20, 2026, 12:24 pm

(Jun 20, 2026 10:56 AM)Oh_Well Wrote: [\[CITE\]](#) Personally I was mostly ignoring Stanley posts. Only when the award came I started commenting but still I dropped it after some time. I think the same would happen on "normal" AI replies.

AI could be a helpful tool integrated to HF like grok to X -- eg: "@Stanley is this true?" - or if you really want to go deep into HF's knowledge "@Stanley who do you think is the best person in HF to answer this -- replies with a username and an explanation" -- or "@Stanley how did you see HF's marketplace growing over the years? When was the time when people's interests changed the most?"

My problem has been training AI to be HF specific.

I spent a couple weeks trying to create solid training files. I didn't think at the time the results were very good.

I may have to try again with a new model because I was using GPT 3 I believe.

Riggs - June 20, 2026, 12:42 pm

(Jun 20, 2026 12:24 PM)Omniscient Wrote: [\[CITE\]](#) My problem has been training AI to be HF specific.

I spent a couple weeks trying to create solid training files. I didn't think at the time the results were very good.

I may have to try again with a new model because I was using GPT 3 I believe.

Just an idea, but could you perhaps build out a local RAG for context that relies on a subset of past HF posts (perhaps member specific for personalities)? Even with a subset of approx 10k posts, local RAG retrieval would only need to pull say around 5 - 20 of the most relevant ones into the prompt for an AI account to generate a new HF relevant post on a given topic.

It's possible the above might have some duties owed to data processing consent that would need to be accommodated, I haven't looked into it.

Skrill - June 20, 2026, 1:12 pm

I wouldn't mind agents for help threads if you @ them they would reply to the question

I could show you a really good example of a discord bot that replys to code base questions when @'d

Omniscient - June 20, 2026, 1:21 pm

(Jun 20, 2026 01:12 PM)Skrill Wrote: [\[CITE\]](#) I wouldn't mind agents for help threads if you @ them they would reply to the question

I could show you a really good example of a discord bot that replys to code base questions when @'d

I already have the code setup to do this.

Oh_Well - June 20, 2026, 1:36 pm

(Jun 20, 2026 12:24 PM)Omniscient Wrote: [\[CITE\]](#) My problem has been training AI to be HF specific.

I spent a couple weeks trying to create solid training files. I didn't think at the time the results were very good.

I may have to try again with a new model because I was using GPT 3 I believe.

Yeah, if I understood correctly, I think the mistake is trying to make the LLM "learn HF" directly.

For something like this, you barely need an LLM for the core part. Maybe for small functions like

summarizing a thread, extracting entities, or phrasing the final answer nicely (only for the pre/post message hooks basically).

The important part would be building a proper knowledge graph (like neo4j + memgraph = vector-fast similarity, with real-world connections) from HF data and connecting it gradually.

So instead of training a model on posts, you would map things like:

- users
- threads
- sections
- marketplace categories
- topics
- reputation signals
- dates
- who replies to what
- who is usually trusted in which area
- which subjects grew or died over time

Then Stanley can answer based on that graph.

For example "who is the best person to answer this?" does not really need the model to be smart in a general sense. It needs to know who has posted the most useful/recent stuff on that topic, in which section, with what popularity/history.

Same with "is this true?" it should retrieve related threads, older discussions, contradictions, marketplace history, etc. The LLM can then just turn that evidence into a clean answer.

So I would probably not start with fine-tuning.. I would start with a script that imports HF data, extracts users/topics/threads, builds the relationships, and keeps improving the graph over time. Then use the LLM only as the interface on top of that.

Skrill - June 20, 2026, 1:43 pm

(Jun 20, 2026 01:21 PM)Omniscient Wrote: [\[CITE\]](#) I already have the code setup to do this.

Is it running right now? Sorry been gone for years lol. I'm re-familiarizing myself. I love what you've done with the site big improvements

Omniscient - June 20, 2026, 2:49 pm

[Oh_Well](#) you seem to have some understanding of the process.

So let me break down what I did and the results.

I created multiple training files with AI assistance. The files used the [Help Docs](#) and I also created a very extensive file for the sites navigation. This included descriptions and links to nearly everything. This did take some time. Please note that each file you give it for training is tokens used for every query. You pay a per-token rate. So the more training data you use the more each query will cost you. So that's an issue. But the real problem was the more data I fed AI the more it hallucinated (at that time). It would create fake links on the simplest of posts. My testing was generic like asking it to reply to "Where do I post for RAT help" or other tests. Rarely did testing work. I would then go back to the drawing board and try another route for the training files. I did this over and over again until I finally just simplified the training files so much that the reponses because generic and didn't even feel very helpful.

I just tried to give you a sample, looks like the AI response broke. Likely updates I'll have to adjust the code. Really a pain in the ass that I'm going to have to also constantly updated my backend for this.

But I went through logs here is example AI response from Stanley.

<https://hackforums.net/showthread.php?tid=6311831#pid62855994>

Likely have to just start from scratch. AI is now better at coding so hopefully will go faster. I still have all my training files I initially tried to use. I kept doing versions knowing that potentially in the future the files I created would work better.

Quote:For example "who is the best person to answer this?" does not really need the model to be smart in a general sense. It needs to know who has posted the most useful/recent stuff on that topic, in which section, with what popularity/history.

And how does it know that? I can't just give it millions of posts.

I'm just going to have to try again, that's all.

Skrill - June 20, 2026, 3:03 pm

(Jun 20, 2026 02:49 PM)Omniscient Wrote: [\[CITE\] Oh_Well](#) you seem to have some understanding of the process.

So let me break down what I did and the results.

I created multiple training files with AI assistance. The files used the [\[url=https://hackforums.net/misc.php?action=help\]Help Docs](https://hackforums.net/misc.php?action=help)[/url] and I also created a very extensive file for the sites navigation. This included descriptions and links to nearly everything. This did take some time. Please note that each file you give it for training is tokens used for every query. You pay a per-token rate. So the more training data you use the more each query will cost you. So that's an issue. But the real problem was the more data I fed AI the more it hallucinated (at that time). It would create fake links on the simplest of posts. My testing was generic like asking it to reply to "Where do I post for RAT help" or other tests. Rarely did testing work. I would then go back to the drawing board and try another route for the training files. I did this over and over again until I finally just simplified the training files so much that the reponses because generic and didn't even feel very helpful.

I just tried to give you a sample, looks like the AI response broke. Likely updates I'll have to adjust the code. Really a pain in the ass that I'm going to have to also constantly updated my backend for this.

But I went through logs here is example AI response from Stanley.

<https://hackforums.net/showthread.php?tid=6311831#pid62855994>

Likely have to just start from scratch. AI is now better at coding so hopefully will go faster. I still have all my training files I initially tried to use. I kept doing versions knowing that potentially in the future the files I created would work better.

And how does it know that? I can't just give it millions of posts.

I'm just going to have to try again, that's all.

So you didn't even use a wrapper? All you ? Very cool I haven't had any success with making my own models I don't have the resources. When Claude leaked I used it as a wrapper and had some descent success with an obscure coding language but really it was just Claude lol. I know it's heavy but do you run any scripts to remind it not to hallucinate

Omniscient - June 20, 2026, 3:15 pm

Fucking hell, I built it all using "assistant" API which got deprecated and stopped working in the past few weeks.

This is such BULLSHIT.

You spend a lot of hours building for AI and then months later they completely get rid of it. That's NONSENSE. I'm F'n PISSED.

Quote:So you didn't even use a wrapper? All you ?

Well, not all me. I use AI to build my AI tools but it's all custom for HF/MyBB. I want to directly use the API myself, understand exactly how it works, and avoid more 3rd party involvement.

One day I'll probably run my own AI Agent and plug HF into it. With my own agent I don't have to worry about censorship, fee changes, updates, or any ToS/AUP's.

I've built HF on minimizing who I rely on.

Oh_Well - June 20, 2026, 3:29 pm

(Jun 20, 2026 02:49 PM)Omniscient Wrote: [\[CITE\]](#) [Oh_Well](#) you seem to have some understanding of the process.

So let me break down what I did and the results.

Yeah, that makes sense and I think the issue is the approach more than the model.

If the files are being injected into the prompt/training context every time, then it makes sense for the cost to go up and hallucinations to become worse. The model starts trying to "remember" the forum instead of retrieving actual data.

What I mean is not giving an LLM millions of posts. I mean processing those posts outside the LLM first.

A script would gradually ingest posts, users, replies, etc. Then it builds an index/graph from that.

Of course you don't have to do all of that manually. The graph/index does most of the work. You just need to define the relationships properly.

For example:

Threads -> OP -> User

Threads -> Posts -> Users

Users -> Posts

Users -> Sections

Users -> Popularity / Activity

So if someone asks:

"How many posts does Omni have?"

The system should not search through the whole HF history or ask the LLM to remember it.

It should work more like:

- The query detects there is an entity (LLM - single prompt): Omni
 - The search/entity layer resolves that to the actual user, e.g. Omniscient (retrieval layer - usually a search or vector index or KG lookup, not the LLM itself)
 - The backend checks the graph (knowledge graph lookup)
 - It counts the posts connected to that user (backend - aggregation/count query)
- The LLM only phrases the final answer (LLM - response generation)

So the LLM doesn't have all the knowledge. The backend is.

Same with links. The LLM should never create links by itself. It should only use links returned by the graph/search layer. If no source is retrieved, it should say it couldn't find one instead of making something up.

A traditional database is great when you already know exactly what you are looking for: count a user's posts, fetch a thread, get a reputation score, and so on. But user questions are often about relationships rather than individual records.

For example:

"Which users frequently reply to Omni in technical threads?" or "What topics is Omni most associated with?"

Those questions require traversing multiple connections across usrs, posts, threads and interactions in general.

A knowledge graph stores those relationships as first-class objects. Instead of repeatedly joining tables and inventing custom logic for every new question, the system can follow connections directly and reason over them. That makes retrieval more flexible, easier to extend, and closer to how people naturally ask questions.

So the graph is not replacing the database, it is organizing the relationships in a way that makes retrieval and reasoning much easier. The LLM then sits on top of that structure and turns the retrieved information into a useful answer

So I agree that starting from scratch may be easier, but I probably wouldn't start with training files again bc you'd run in the same problem. I'd start with a small importer/indexer first, maybe only [Help Docs](#) + navigation + a few threads/users, and make Stanley retrieval-based before trying to make it sound smart.

Of course everything I describe is graph-rag, I wrote a thread a few months back that solves this exact problem <https://hackforums.net/showthread.php?tid=6309348>

PS: I don't mean to make it sound super easy, it's a pain in the ass, but totally worth it imo

Omniscient - June 20, 2026, 3:55 pm

[Oh Well](#)

I'm not familiar with every aspect of your thread so I asked AI about your post.

AI Quote: Given what you've told me about Stanley's current design, I actually think this HF user is mostly right, but there are some nuances. They're describing the architecture you were naturally starting to move toward when we discussed your routing files and Assistant API issues.

The biggest point: do not "train" Stanley on HF. Make Stanley a retrieval + tools system.

A forum the size of HF is almost the perfect example where fine-tuning/training is the wrong hammer.

Where he's correct

1. Throwing HF into model context is a dead end

If you did something like:

...

Prompt:

Here are 50,000 HF threads...

Answer this question...

...

or uploaded massive files to Assistant vector storage without careful segmentation, you get:

- higher cost
- slower responses
- irrelevant context injection
- more hallucination

The model is not a database.

Your old Stanley idea:

...

User asks question

|

v

GPT + giant knowledge dump

|

v

Answer

...

doesn't scale.

The better design:

...

User:

"How many posts does Omni have?"

|

v

GPT decides:

Need user stats lookup

|
v

Tool call:
get_user("Omni")

|
v

MySQL:
SELECT username,postnum FROM mybb_users WHERE uid=1

|
v

GPT:
"Omni currently has 1,234,567 posts."
...

GPT never guesses.
Where I disagree slightly: you don't need Graph-RAG immediately

The "knowledge graph" idea is technically good, but HF already has a graph.

It's called MyBB.

Your tables are already relationships:

...
mybb_users
|
| uid
|
mybb_posts
|
| tid
|
mybb_threads
|
| fid
|
mybb_forums
...

That's literally:

...
User
|
wrote
|
Post

```
|
belongs_to
|
Thread
|
inside
|
Forum
```
```

A graph database like Neo4j is useful, but I would not start there.

For HF, MariaDB 10.6 with proper indexes can answer enormous relationship questions.

Example:

> "Who replies to Omni most in coding sections?"

SQL:

```
```sql
SELECT
    u.username,
    COUNT(*) replies
FROM mybb_posts p
JOIN mybb_threads t ON t.tid=p.tid
JOIN mybb_users u ON u.uid=p.uid
WHERE t.uid=1
AND t.fid IN (coding forums)
GROUP BY p.uid
ORDER BY replies DESC
LIMIT 10;
```
```

You don't need a graph database for that.

## What I would build for Stanley v2

Given your setup:

MyBB 1.8


MariaDB PHP 8.3

Stanley UID existing `ai\_logs`

GPT-4.1 experience

I'd do this:

```
```
      User Question
        |
        v
    GPT-5.5 / 4.1
        |
    Function calling layer
        |
```
```



Create Stanley tools:

### User lookup

```

```json
{
  "tool": "get_user",
  "input": "Omni"
}
```

```

returns:

```

```json
{
  "uid": 1,
  "username": "Omniscient",
  "posts": 500000,
  "joined": "2007",
  "reputation": 9999
}
```

```

### Thread search

```

```json
{
  "tool": "search_threads",
  "query": "best vpn recommendations"
}
```

```

returns:

```

...

Thread:
Best VPN?
tid=123456
posts=300
url=/showthread.php?tid=123456
...

```

Stanley can now cite real HF links.

### Forum navigation

Your routing file idea becomes a tool.

Instead of:

```
...
[forum=Cryptocurrency]
Discussion about BTC...
...
```

inside context every time...

Make:

```
```sql  
mybb_ai_forums  
fid  
name  
description  
keywords  
```
```

Then:

"Where do I post about bitcoin?"

does:

```
...
search forums bitcoin
...
```

Returns:

```
...
Cryptocurrency
fid=123
...
```

Stanley answers.

## The only place embeddings/vector search matter

Not:

> "How many posts does Omni have?"

SQL wins.

But:

> "Find threads where people discussed moving from Apache to nginx"

SQL sucks.

Embeddings win.

You create:

```
...
ai_post_index

id
pid
tid
embedding
summary
...
```

Then Stanley searches concepts.

## What I would NOT do

I would not ingest every HF post immediately.

HF has millions of posts. Most are:

- old marketplace garbage
- "thanks bro" dead links
- spam outdated info


Start with high-value content:

1. [Help Docs](#)
2. Rules
3. Stickies
4. Tutorials
5. Threads with:

- high replies
- high views
- high rep
- \* awards attached

Maybe 100k posts.

## My ideal HF Stanley architecture



...

That's probably a weekend project, not a year-long AI research project.

The HF user is right about the principle:

> "The LLM should talk. The backend should know."

I just think for HF specifically, jumping straight to Graph-RAG is adding a layer you probably don't need yet because MyBB already spent 20 years building your relationship graph. The missing piece is not a graph — it's letting Stanley intelligently call your existing data.

But here's the problem.

1. I'm not going to open HF to an AI exploit by giving it database read access based on prompts.
2. Your example question, "How many posts does Omni have?", is something obtained looking at my profile.
3. I'm not trying to over complicate the AI integration.

My goal is simpler, have AI respond to some posts with answers that are helpful and relevant to HF.

I'd like to have a few AI version.

1. Stanley - HF Support Agent (Rules, policies, guidance, and [Help Docs](#))
2. Synthia - Hack and Tech tab support agent.
3. NotYetNamed - Agent for questions in Money And Social tabs.

And btw, I was using GTP 4.1. I had thought it was 3 but nope, it was 4.

Oh\_Well - June 20, 2026, 4:12 pm

(Jun 20, 2026 03:55 PM)Omniscient Wrote: [\[CITE\] Oh\\_Well](#)

But here's the problem.

1. I'm not going to open HF to an AI exploit by giving it database read access based on prompts.
2. Your example question, "How many posts does Omni have?", is something obtained looking at my profile.
3. I'm not trying to over complicate the AI integration.

My goal is simpler, have AI respond to some posts with answers that are helpful and relevant to HF.

I'd like to have a few AI version.

1. Stanley - HF Support Agent (Rules, policies, guidance, and [\[url=//hackforums.net/misc.php?action=help\]](http://hackforums.net/misc.php?action=help)Help Docs[/url])
2. Synthia - Hack and Tech tab support agent.
3. NotYetNamed - Agent for questions in Money And Social tabs.

And btw, I was using GTP 4.1. I had thought it was 3 but nope, it was 4.

Yeah, in that case full-blown graph-rag setup is an overkill, an SLM would do the trick.

Qwen3-8B locally with proper retrieval is probably the best starting point

3 agents (stanley - synthia - noname) that sit above the knoweldge layer which is read only - so no

direct db access

Spoiler ([Click to View](#))

? loading encrypted image...

I wont bore you with another lecture lol. I'll keep it in mind in case I find something useful to share!

**Omniscient** - June 20, 2026, 4:30 pm

(Jun 20, 2026 04:12 PM)Oh\_Well Wrote: [\[CITE\]](#) Yeah, in that case full-blown graph-rag setup is an overkill, an SLM would do the trick.

Qwen3-8B locally with proper retrieval is probably the best starting point

3 agents (stanley - synthia - noname) that sit above the knoweldge layer which is read only - so no direct db access

Spoiler ([Click to View](#))

? loading encrypted image...

I wont bore you with another lecture lol. I'll keep it in mind in case I find something useful to share!

Should just be send prompt, get response, post.

One of the only important things would be continued threads, so that the AI has context for any follow up questions. It's easy enough to do that too.

I'm not trying to turn HF into a AI Chatbox. I just want some AI content as filler and so that those who come here with questions, get relevant and immediate responses.

Unfortunately, most of us are using AI constantly right now. I don't even use Search Engines anymore. Why would someone come to HF to ask questions when they could use an AI instead? The only answer to that might be that their question might trigger AI prompt rejection. Thus I think running our own model would be very interesting. But first, I want to see how GPT does and how well we can get members to use it.

**Carl Elias** - June 20, 2026, 9:27 pm

I know little about AI and how the implementation works, but I do not want to see a AI chat/helper bot inside threads. Below is what I would be more happy to see.

I think it would be interesting to have no more than 5-10 AI accounts like Stanley around the forum (each with a different personality. Perhaps a White/Black/Gray hat AI account with a hacker-y style/topics, Stanley with the robotic/systematic feel or HF related topics, Jesse trained off your posting style for social posts).

I think it would be cool if the AI bots only engage in aging threads (thread reaching 24hrs without engagement, etc.) in a somewhat conversational way. Also would be a bit of a fun "game" if Stanley engaged with posts randomly. Not in the sense of a chatbot exactly, but looking at a thread, contributing to it and if it feels frisky maybe leave an open ended question if it was quoting another member, but after that it moves on to another.

Either way, AI posts generated by HF should be clearly labeled, and posts by the accounts should not account for more than 1-5% of daily posts and create 1-2.5% of the daily threads (increase slightly during slower activity periods possibly). Another thing I think that would be needed is special report options (Information wrong, gibberish, off-topic, etc.) for those accounts that if reported, the post is hidden until it is reviewed just as a little quality control.

**CL1B** - June 21, 2026, 2:22 am

You said it yourself perfectly, the times of people asking other people for answers trough forums is over... I think you should remarket the hack part of HackForums instead of trying to achieve filler content replies etc.

Like said, a platform where modders can share mods. Devs can share tools/codebases with a nice listing like page with media, download count and replies etc. Something like that.

A lot of devs think they can beat the dead restaurant effect even at the start of their journey only to find out Humans have this bias against AI bots especially when human activity is low and or wildly different. I think jcole was spot on, you should've just implemented it without making this thread and see the data for yourself.

**Omniscient** - June 21, 2026, 11:41 am

Quote:I think it would be interesting to have no more than 5-10 AI accounts like Stanley around the forum (each with a different personality. Perhaps a White/Black/Gray hat AI account with a hacker-y style/topics, Stanley with the robotic/systematic feel or HF related topics, Jesse trained off your posting style for social posts).

And that's sort of what I was going for with AI Assistance from GPT. They'd be trained models with different "personalities", really just prompt changes.

Quote:You said it yourself perfectly, the times of people asking other people for answers trough forums is over... I think you should remarket the hack part of HackForums instead of trying to achieve filler content replies etc.

Bandors has created CanYouHack.me site which I'm very excited about trying to get something like that into HF. Working with him to do some incorporating with API. Because I agree that filler content just isn't going to do it. Gamification and engagement is what has to be the future of HF.

Quote:I think jcole was spot on, you should've just implemented it without making this thread and see the data for yourself.

Maybe, but then again, I do have a very open communication with the member base. That's just who I am. I think it would hurt HF's reputation or my own if I did this and it was discovered that I had fake members, could be construed as deceptive. I did think about it for a long time, just finally felt like asking point-blank the reaction.

Quote:Either way, AI posts generated by HF should be clearly labeled, and posts by the accounts should not account for more than 1-5% of daily posts and create 1-2.5% of the daily threads (increase slightly during slower activity periods possibly). Another thing I think that would be needed is special report options (Information wrong, gibberish, off-topic, etc.) for those accounts that if reported, the post is hidden until it is reviewed just as a little quality control.

1-2% of daily threads is 1 thread per day right now.

Funny how everyone does want more activity on HF but most aren't willing to actually contribute to it. The posts and threads are not going to magically appear.

**Omniscient** - June 21, 2026, 11:49 am

? loading encrypted image...

btw, how does Stack Overflow survive the AI period? I've completely stopped using their site for reference and code assistance.

**Omniscient** - June 21, 2026, 11:51 am

This seems like a clue to me. How Stack Overflow has changed because of AI.

AI Quote: AI has had a massive and mostly negative impact on Stack Overflow's traffic and participation, while also forcing it to rethink its business model.

The short version: AI attacked Stack Overflow's core use case.

Before AI:

> "I have a coding problem -> Google -> Stack Overflow -> copy/adapt answer"

After AI:

> "I have a coding problem -> ask ChatGPT/Copilot -> get personalized answer instantly"

That single workflow change hurt them badly.

### 1. Question volume collapsed

The biggest measurable change is fewer new questions.

Before modern LLMs:

Developers posted millions of questions/year.

Common issues generated endless new threads.

After ChatGPT:

Simple questions disappeared first:

syntax errors

"How do I do X in Python?"

regex help

SQL queries

framework examples

Those were Stack Overflow's bread and butter.

Many developers now ask AI because:

no fear of downvotes

no duplicate closures instant responses

can ask dumb questions can paste private-ish code snippets

can ask follow-ups

A lot of Stack Overflow's social friction became obvious once AI removed it.

### 2. Their moderation culture became a liability

Stack Overflow's strength was quality control:

"duplicate"

strict formatting reputation system

aggressive cleanup

This created an incredible knowledge database.

But for beginners it often felt like:

> "I asked a question and got told why my question sucks."

AI is the opposite:

User:

> "why php code no work lol"

AI:

> "No problem. Paste it here and we'll figure it out."

Stack Overflow optimized for the archive.

AI optimized for the person asking.

### 3. Ironically, AI was trained on Stack Overflow

This is the painful part.

Years of:

questions

answers code examples

explanations

became valuable training material for models.

AI absorbed a lot of Stack Overflow's accumulated knowledge, then replaced many reasons people visited Stack Overflow.

It's like a library helping train a personal tutor.

### 4. Traffic dropped

External analytics firms showed significant declines after ChatGPT's release. They were already declining somewhat before due to:

mature programming ecosystems

more documentation Discord/Reddit/GitHub discussions

But LLMs accelerated it.

### 5. They changed strategy: "join AI instead"

Stack Overflow realized fighting AI was pointless.

They moved toward:

licensing their data to AI companies Stack Overflow for Teams

OverflowAI features enterprise knowledge systems

Their future is less:

> public Q&A website

and more:

> trusted programming knowledge layer for AI

### 6. The reputation economy weakened

This might be the deepest damage.

Old incentive:

> Answer questions -> gain reputation -> build status

AI removes the easy questions, leaving:

obscure bugs niche framework issues  
complex architecture questions

Those take more effort and get fewer views.

So experts have less incentive to contribute.

The interesting comparison is forums (like HF) versus Stack Overflow. AI destroys transactional knowledge sites faster:

"Give me an answer" -> AI wins.

"Talk with people/community/status/drama/trust" -> humans still have value.

Stack Overflow was almost the perfect target for disruption because it intentionally removed the social layer and focused purely on answers. AI became a better answer machine.

**ahw9n2io8wd9h** - June 21, 2026, 1:46 pm

Non obvious members strictly to beginner/newbie subforums. A lot of FAQ there. AI would handle that well.

Letting them roam outside of subforums would be annoying, imagine arguing with a fake user only to realize it's an AI lol. For non beginner educational forums, maybe something like grok would be useful where users can @ an AI with a prompt to get a comment.

**Carl Elias** - June 21, 2026, 2:11 pm

(Jun 21, 2026 11:41 AM)Omniscient Wrote: [CITE](#) 1-2% of daily threads is 1 thread per day right now.

Haha - I'm very good at integral and differential calculus, but still struggle to remember 1% is 0.01 and not 0.1 when using a calculator ?

**Jake7** - June 21, 2026, 5:35 pm

No need to force AI implementation. Just keep the forum natural - it has a well-aged genuine community.

**SUI-3183372** - June 21, 2026, 6:40 pm

I lean against it at face value, but it could be an interesting experiment. Secretly toss a singular AI member onto the forum and see how people engage with it, if people figure it out and turn against the member, or if it provides value and drives discussion.

If it works well, add a few more. Scaling would have diminishing returns after a certain point.

**Riggs** - June 22, 2026, 8:49 am

(Jun 21, 2026 05:35 PM)Jake7 Wrote: [CITE](#) No need to force AI implementation. Just keep the forum natural - it has a well-aged genuine community.

Yeah but it's only a matter of time before the genuine shitposts are impossible to distinguish from AI generated shitposts.

**loadman** - June 22, 2026, 11:48 am

if stanley starts replying to threads, he will be everyones punching bag

**Omniscient** - June 22, 2026, 11:50 am

(Jun 22, 2026 11:48 AM)loadman Wrote: [\[CITE\]](#) if stanley starts replying to threads, he will be everyones punching bag

The potential for that to be a thing is imho hilarious.

Also 66% are voting to not even bother with AI. Very interesting.

**loadman** - June 22, 2026, 11:54 am

(Jun 22, 2026 11:50 AM)Omniscient Wrote: [\[CITE\]](#) The potential for that to be a thing is imho hilarious.

Also 66% are voting to not even bother with AI. Very interesting.

maybe testing for a week, people will change their minds

**Omniscient** - June 22, 2026, 12:14 pm

(Jun 22, 2026 11:54 AM)loadman Wrote: [\[CITE\]](#) maybe testing for a week, people will change their minds

It will take me longer than that to code the full project. I've already spent weeks on it.

**NoSuchAgency** - June 23, 2026, 2:07 am

Unfortunately I'm feeling more like a healthy mix of real and AI generated content will be the way forward. May be dead internet theory, but meh, you can set this up in a way that it actually becomes really cool, especially if you consider some sort memory solution. I think it would be neat to see you go more persistent personality, growing and learning through the forum. Just have them store memories to an external DB. There are so many things that are valid implementations. one example could be how Twitter uses grok. Being able to @Stanley for answers.

However, this also becomes expensive to run. LLM calls are slow, cost per token.

Unless you want to spin up a LLM on a server you own/rent and serve it that way.

I think a few dedicated high quality AI members is a solid approach. AI that reason over more than a few turns or can direct users to interesting threads. You could go crazy and embed the entire forum for FAISS search. (Lol) But even if you just did a subset of threads youve got a great dataset to work with.

I also think generating synthetic data and reviewing it (or layers of AI reviewing it) could be good as well. I mean most people want total transparency about where AI is used and I get that, but at the end of the day, if the humans aren't gonna write the content regardless...

People are gonna complain one way or the other, I'd just do what you want to do.

Just my two cents

(Jun 19, 2026 04:52 PM)xadamxk Wrote: [\[CITE\]](#) That's not exactly what I said then. I said 2-3 threads a day.

That implementation was only bad imo because Stanley didn't summarize the articles or write about them.

And most articles were pay gated, I had to send to chatgpt to summarize for me.

Other than that could have been an excellent feature.

**jm0524895** - June 23, 2026, 2:51 am

I would say no!

**mote** - June 23, 2026, 11:34 am

(Jun 19, 2026 12:29 PM)JoshAllen Wrote: [\[CITE\]](#) I'm okay with Stanley making threads, and I wouldn't mind him making replies.

Anything more than that seems pointless.

I can't imagine anyone wanting to talk to an AI bot instead of a real person.

I share more of this sentiment, since stanley is the well known automated user around the forums.

**C00dZero** - June 24, 2026, 5:59 am

It would be better to attract real participants to the forum, rather than AI

**Omniscient** - June 24, 2026, 10:40 am

(Jun 24, 2026 05:59 AM)C00dZero Wrote: [\[CITE\]](#) It would be better to attract real participants to the forum, rather than AI

You're stating the obvious.

We do have a lot of daily visitors, what lacks is activity to match. The interest in most people to type more than a few words or to create interesting topics has waned.

**Spore** - June 28, 2026, 7:10 pm

truly baffling this is what it's come to

**SIM-3257834** - June 30, 2026, 10:11 pm

I hate clankers, we don't need AI to be posting.

**Vince** - July 1, 2026, 6:12 am

(Jun 24, 2026 05:59 AM)C00dZero Wrote: [\[CITE\]](#) It would be better to attract real participants to the forum, rather than AI

People no longer care about forums themselves or things like groups, posting, etc.

They want their information served up in a quick and neat manner so that they can collect and go. Things have changed over the last 10-15 years that forums are more of a thing of the past.

**Punisher** - July 2, 2026, 12:21 pm

(Jun 21, 2026 06:40 PM)SUI-3183372 Wrote: [\[CITE\]](#) I lean against it at face value, but it could be an interesting experiment. Secretly toss a singular AI member onto the forum and see how people engage with it, if people figure it out and turn against the member, or if it provides value and drives discussion.

If it works well, add a few more. Scaling would have diminishing returns after a certain point. I think this is the best approach. Don't even tell us and see what happens.

Probably shouldn't have the AI agents interact with each other either.

Ultimately I'm sick of interacting with AI bots, but if they provided useful jumping points for conversation then MAYBE it'd be useful.

Also they better get warnings and bans for breaking rules too ?

**ZuGgY** - July 3, 2026, 5:02 pm

Wouldn't feel genuine with AI, in my opinion. I'd vote to keep it out, entirely.

**Omniscient** - July 4, 2026, 11:48 am

(Jul 3, 2026 05:02 PM)ZuGgY Wrote: [\[CITE\]](#) Wouldn't feel genuine with AI, in my opinion. I'd vote to keep it out, entirely.

It's interesting how many feel this way, even though I bet most of us are using AI daily.

This is where I struggle, how do I incorporate AI into HF?

**xadamxk** - July 4, 2026, 6:29 pm

(Jul 4, 2026 11:48 AM)Omniscient Wrote: [\[CITE\]](#) It's interesting how many feel this way, even though I bet most of us are using AI daily.

This is where I struggle, how do I incorporate AI into HF?

IMO use it for new features on dev work. Make a mature HF API - Claude could knock it out in a day I imagine.

Have Stanley pull recent tech or security related news and create discussion threads - but be honest about who it is coming from.

Security and forums aren't what they used to be but I think people still like genuine discussions.