

A Practical AI Agent Workflow For Companies In 2027 (Guide)

19 MIN · YOUTUBE · [HTTPS://WWW.YOUTUBE.COM/WATCH?V=8RVQUZLRAQ0](https://www.youtube.com/watch?v=8RVQUZLRAQ0)

<https://www.youtube.com/watch?v=8rVQuZLRaQo>

SUMMARY

The speaker discusses how AI agents have transformed productivity in knowledge work, emphasizing a practical workflow that integrates AI into daily tasks. By using a shared workspace and project management tools, individuals can delegate tasks to AI agents, allowing them to focus on higher-level decision-making and quality assurance.

- AI agents are revolutionizing productivity by automating tasks within a shared workspace.*
- The speaker's business is projected to earn over \$400,000 this month by leveraging AI for various tasks.*
- A project management tool, Linear, is used to track tasks with statuses like inbox, next, doing, waiting, and done.*
- Tasks can be tagged as "AI ready," enabling AI agents to autonomously handle them based on predefined parameters.*
- The integration of low-friction capture methods allows for quick task input, enhancing workflow efficiency.*
- Evals are used to assess AI outputs, ensuring they meet quality standards and align with the user's preferences.*
- The speaker emphasizes the importance of scoping tasks and evaluating AI outputs rather than performing all tasks manually.*
- The approach allows for faster execution of ideas and tasks, leading to increased productivity and the ability to focus on strategic initiatives.*

So, the way the most productive people on Earth work today is very different from just maybe 6 months ago. And the entire reason is because of AI agents. But, despite everybody and their mom talking about AI agents, using them in workflows, I don't think anybody's actually really clarified what it looks like in a practical knowledge work style situation. I think Claude made some good headways with, you know, Claude integrated in Slack. But, today I want to show you guys how I am currently getting things done. My business is going to do over \$400,000 this month, and we work with a variety of different things. We work with media, we work with AI implementation, we do consulting, we have a variety of different like verticals and stacks. And I want to show you how you can use AI agents to meaningfully improve both the quantity and then the quality of work that you do in an organization in a very simple and straightforward way. So, the very first thing you need is you need a shared AI and human workspace. And to make a long story short, this is a place where you can have AI working on some task alongside people.

Genuinely, you can use whatever platform or tool you want for this workspace. the thing that is more important than the specific like selection is just the shape of it. And I'm going to run you through what ours looks like right now. I use a tool called Linear, which is just one of the many project management workspaces that you could use for this. And basically, the way Linear is broken down is you have a variety of different statuses that denote

where a task is in the pipeline.

And so, the statuses that we've set up, if I zoom in a little bit over here, is we have one called inbox. We then have another one called next. We have one called doing. And then we finally have one called waiting. And then there's also the done status, which is where we put tasks after they're done, just so that if, you know, we need to take a look at how a task was completed or get some more context or move it back to next or doing, you know, we can do that pretty easily. So, I'm not going to talk all day about this setup because hopefully that's pretty standard. The key thing in here is this isn't just humans working on this to-do list. We now have the ability to weave in AI agents. And rather than have to like prompt an AI agent constantly, all I'll do is I will make a task and I will tag it as AI ready and then I will send it off and actually have it do said task. So one of the tasks on my to-do list today is I need to ideate, you know, another five videos. So I'm going to press C, that's going to open up a new issue and then I'll say ideate five new YouTube videos. And then in the description, because keeping in mind that I'm not just going to give this to a person, I'm going to give this to an AI agent, I'm just going to, you know, ad-lib what I want. Hey, I want you to use context based off videos that I've published before to come up with five new YouTube video ideas including three title options for each, some angles for each, maybe even some brief outlines based off of current trending content on the internet. After I'm done, I'm going to head over here to create issue and click that button.

Now, all I need to do is I just need to tag this and this is a very slow, naive way of doing all this, but I will also show you a much faster way. and once it's an agent ready, we can actually begin working on the task. So the way that, you know, I'm going to do it is I'm just going to drag it over to next. And essentially what occurs when that happens is this sends a request to a server where I have an AI agent living.

This AI agent will pick up that task. It will then consult my knowledge base to learn kind of contextually how previous tasks have been completed, my preferences surrounding tasks and so on and so forth. If we just click on this and scroll down, you can see the Nick OS agent run has actually started. So it's picking it up now, assessing the card, posting the plan and so on and so forth. And then it's even actually writing stuff down.

It's doing so just with my profile cuz I'm cheap and I don't want to spend, you know, a bunch more linear costs, but obviously you could have this be like my fable agent. And then it's going through and actually doing the task for me right now. You know, if I exit out of this, you see we've now changed the status to doing. And so now this is this is occurring almost entirely autonomously. All I do is I'm just like the the project manager, if that makes sense. I add things to a queue, agents complete them according to my specifications, and or, you know, if there's a Q&A step involved, which a task this would probably have a Q&A step involved, you know, the agent will wait for my approval before actually doing, let's say, the publishing or interacting with like the wider internet, creating me a post on some profile. So, while it's doing the task, let me just run you through essentially what's going on under the hood.

You know, we start with the new task over here, and that's the one that I just created talking about, you know, the content that I wanted to make, the the YouTube titles and and angles. Then it's caught by a webhook. For those of you guys that don't know, that's just a place on the internet that a request gets sent to. And the cool part

about Linear, and one of the reasons I like it, is just because it has that functionality built in. You can just like send a request off when it's tagged with something.

And now it goes to Fable 5. The thing is, Fable 5 has access to three things. It has access to workspace context. So, this is everything that, you know, the we are currently using or doing in the workspace, all the files, all the other tasks, you know, our content pipeline, and so on and so forth. It then has access to a knowledge base, which is a bunch of additional textual information that it can draw upon if necessary.

And then finally, it has it has access to credentials. And these credentials are things like, you know, passwords to various services and platforms that maybe it needs to use Chrome DevTools to sign into, or you know, API keys, and so on and so forth. And all this is stored relatively securely, to the point where Fable can use all of it in conjunction with its own gigantic galaxy brain intelligence to actually do the work. And you can kind of see that in the task. I mean, after the update, it's just confirmed that this is how it understands it. Make me some videos.

Well, I've given it context as to my highest performing courses in the past. I've given it I've given it context as to my knowledge base. I've given it context as to like my strategy docs, so how I typically create ideas and angles, and so on. And you can also see that it just actually changed the label to waiting, and changed the status to waiting. So, let's take a look at what that looks like. Now, the thing to know about AI agents is, you know, their outputs aren't incredible. You can't just trust an AI agent to do everything entirely on its own. What you need to do is you need to verify its outputs. You need to essentially have some sort of final line in the sand where you will check on the outputs and then select the ones that you like the most. You'll basically force ideation, let's say, over a large solution space and then apply your human taste to selecting like the best winners. I just did that in my ads video that I published earlier. And you can see here that I don't really like all of these ideas. I don't think they're all brilliant. But this one down here, I let AI agents run a YouTube channel for 90 days. This one here's pretty solid. You can see it's come up with some different title formats, a source trail. It's given me some reasonably good reasoning and rationale behind why something like that would work. And I think that might actually be a video now that I am going to do. In a typical day, my pipeline might look reasonably like this. I'll have a couple of tasks in next, I'll have a few in waiting. you know, let's say I make another one right off the top of my head. to sign up to Anthropic partner network, find out everything I need and pre-write draft application. you know, I'll I'll just be rolling through here selecting agent ready on tasks that make sense, but then also adding tasks that maybe other people in my organization or I need to do. this is really the the idea behind a collaborative and shared workspace. And the main benefit there is I no longer just have to sit down and then look at a terminal, wait for its outputs, and then finally when it, you know, gives me an output, I proceed.

I'm capable of operating like the speed of thought. I'm capable of going very fast here. You know, I can bang out 20 of these ideas simultaneously of 20 different agents all operating on various tasks using my knowledge base, and then I can just check in, you know, once or twice a day when they're done to like assess in batches. This basically solves context switching. Now, once you're done, you can actually just give it some information like, "Hey, I like this one. I'd like to generate 10 thumbs for video Alongside adding this plus alternatives plus everything into content pipeline.

The reason why this is valuable is you just don't need a text box open all day and you can actually work off of a single source of truth. I wanted to buy an ergonomic chair for my home setup. Well, I had to go through and just do a tremendous amount of research cross-referencing Reddit and so on and so forth to get me a bunch. Then I actually ended up buying one of these here, which took like 30 seconds. And you can do this for any task, you know, email rest of people and clear of a CRM cancel Rogers internet for family. You can see that was back when I was trying to master my English accent. That didn't work very well. Looking to how to achieve US resident status and fill out forms. The whole idea is you're basically taking the agent out of the chat box and then you're actually integrating it into the core of your business. But believe it or not, as it is, this isn't very valuable. I mean, what we've done is we basically just created a way to dump ideas in where we haven't really created a way to extend past that. The thing that makes it valuable is when you combine that with capture. Now, the first time I encountered the concept of capture was when I was reading through David Allen's *Getting Things Done* from Forever ago.

And to make a long story short, what capture means is it's just a simple and easy and low-friction way of getting ideas into some to-do list. Now, we actually have our to-do list, right? That to-do list, in essence, is the project management shared human and AI workspace that I showed you earlier. Well, the valuable thing about having a low-friction capture method is, like a good project managers, anything that comes to mind over the course of the day, any task that you need to working on, any concept or idea for anything, you can get into that system extremely easily, whether or not you're on the go, let's say, like, walking around your city or something like that, going from meeting to meeting, or sitting down at the computer. So, obviously, a really low-friction capture method is just using linear and hot keys. So, maybe create WikiData entry for Left Click AI Incorporated. And maybe this is something that, you know, I want to tag as agent ready, create an issue for, and then actually have it proceed with the task like we just did a moment ago.

Well, the simplest and easiest way of sorting this out is actually just by, you know, binding a hot key or something, and then, you know, adding a task as follows. And so, now I actually have this built into my computer. So, maybe I'm out and about, and I don't know, I'm watching some YouTube videos for my daily updates, and I realize that the thumbnail sort of generation prompt is a little bit off. Well, now what I can do is I can actually kick off an AI agent workflow literally without having to stop a beat just by opening this and saying, "Fix thumbnails for daily updates channel." Well, now this has actually been sent off to my linear. And if I wanted to tag this as AI agent ready, you know, I could actually send this to my agent with a brief. I just hold command and press G. So, that's pretty easy. I can give it some additional information if I want. I can tell it, "Right now, the thumbnail text seems to be a little bit small." At any point in time, I basically have woven in AI into my my day-to-day flow, such that it exists with me, next to me, and I can work on essentially whatever tasks. I mean, if you think about it, I could fire off 50 of these simultaneously. I can now operate at the speed of thought.

The bottleneck is no longer, you know, how much work can I do in a day, it's how quickly and accurately can I scope the work that I want done in a day. I just ran over here to get my phone, so I could show you what this looks like as well. You could build a shortcut into your phone such that you can actually just hold this button, create a demo task in Linear, and it'll actually go through, and I don't know if you guys could tell, but I just had a little chime noise, and then I can go back over to Linear, and now you can see that I just I just added that in the

system. you know, it's like tied to my action button, it's extraordinarily low friction.

These little hacks, these aren't necessary, to be clear, but the lower friction that you make something like this, the more you end up relying on sort of your to-do list, your project manager, your shared AI and human workspace. And you can also tie this in such that, you know, maybe if you click the action button twice, it automatically starts it as an agent task, whereas if you click it once, you reserve that for a human task. You can also set notifications so that when, you know, an AI agent needs you for something, it'll actually pop a notification up on your screen. I did this a while back and it was super valuable. Having push notifications when agents are waiting on you, and then all you do is you click on it, and then voice transcribe, like, that is pretty freaking sexy. And not only is it sexy, it's also, obviously, the future of work. If right now our ability to assess and verify the outputs of agents is the bottleneck, we need to make our ability to do that as easily as humanly possible. What's cool, too, is you get total task visibility, so there's like a trail, and it's persistent. It doesn't disappear at the end of every prompt.

So, for instance, you could see here that I actually moved a task from waiting back to next. I removed a label that said waiting on Nick, because, you know, I wanted to provide some more context to the task. And this is going to get picked up just like any other task. All the context is going to be fed into this async runner, and then, you know, things are just going to kind of work on its own autonomously without my direct oversight. The third thing you need in a system like this is you need what are called evals.

Now, in case you guys didn't know, evals are a set of evaluations that you run an output through, or a model through, in order to determine whether or not it is giving you the sorts of things that you want. We're essentially assessing its performance. And for project management and modern AI-based productivity, you know, I define evals as basically a standardized checklist of steps that I run all outputs through before they give it to me. In that way, I know that, you know, we have my tone of voice on every project. We have all of the basic LLM-isms taken out of the text, like m-dashes and stuff like that. We have my reasoning applied to everything. So, I fed in a big knowledge base basically based off my the way that I make decisions, and I was like, "Was this a sort of decision that I would reasonably have made under these circumstances?"

>> >> And the whole idea is rather than giving work off to AI completely and then just hoping it does a good job, what you do is you give it a set of guardrails if it assesses that it's fallen out of those guardrails, it will just iterate and continuously retry the project until it eventually gets within the constraints. Once it's in the constraints, my actual work to, I don't know, touch it up or change something or, I don't know, pick better YouTube titles, pick better tasks, change the work, whatever, is significantly reduced. So, for instance, here's one on the left-hand side here called visual asset. Is this render usable? This applies to demos, thumbnails, diagrams, and generated imagery before it reaches my folder. So, you know, in order for a render to be considered waiting on Nick, basically in order before it even gets to the point where I'm ready to make a Q&A check, the text has to render clean.

It needs to, like, literally visually assess every output and determine that there's no garbled or melted characters. It needs to figure out, "Hey, you know, is this in the style Nick likes?" Which right now is this ink editorial thing,

black ink on pure white. I just love that. Is it legible at thumbnail size? Basically, is it like sufficiently zoomed in? Is it faithful to the brief? And I doesn't have the right shape? AKA, you know, I gen my thumbnails at 1280 by 720. You know, if I do diagrams, it's at 920 1920 by 1080. Here's another one here. Was this done the way that Nick thinks? And this is loosely, you know, related to the knowledge base that I fed it in a while ago, but it's, you know, "Hey, there are five questions. You got to score each of these zero to two. If the total is less than seven, AKA, if one of them is like a zero, then this whole project is a fail and we need to continuously iterate until it is good." So, one of my principles is first principles. Did the work reason from the actual mechanics of the problem or did it just pattern match to what people usually do?

EV discipline. Are the conclusions and choices high expected value? if so, then it's good. And if not, then it's not. Is Nick's time minimized? You know, if there's any part of this that could have been done without Nick in the loop, make sure to go back and do it according to this so that Nick doesn't have to do any additional work. Is it verified, not plausible? Is there a leverage check? You know, if there's a big lever ignored because I had a skill in my workspace that I want you to rerun this. And so, I mean, like, this is these are tokens, right? This isn't free, but typically this is worth far less than the time that I would spend getting an average task up to snuff. And the key is, you know, having a knowledge base of some kind that it consults to like get the the right shape, and then forcing that shape into that guardrail container that I talked about earlier. And finally, there's quality and assurance. And this is really where you come in as, you know, the final, as mentioned, gate to the task.

The reality is, the nature of work has changed a fair amount. You know, before, I was actually doing a lot of the work. But now, what I'm typically doing is I'm spending time scoping the work so that, you know, the guardrails are set really efficiently and effectively, as well as a clear definition of what it is that I want. And then I'm spending most of my time evaluating the results of the work doing Q&A. It's kind of a mind shift from back when I was actually like freelance writing in my content marketing agency. And then later on, I ended up managing a bunch of content writers. It's very similar. I'm just going one level up the abstraction chain. So, not actually directly involved with the deliverables. You know, my fingers and my hands aren't the the things that are doing the heavy lifting, moving the products, and so on and so forth. It's my mind assessing the quality and ensuring that my taste sort of applied to everything that I want.

And so, obviously, you can take what I showed you today, and you can apply it in any way, shape, and form. You could, for instance, go down to the description, expand it, show the transcript, go to the top right-hand corner, copy and paste the entire thing, and put it into Fable, and it could actually recreate something like this fairly straightforwardly for you. Probably without burning more than 25% to 50% of your session limit. you can do the same thing with GPT-5.6 or really any model. This is just my own architecture. But I just wanted to begin the the wider conversation of how to actually apply AI in a modern workplace.

I think a lot of people here have sort of pieced together bits and tiny atomic habits from various workflows that they see on the internet, but this is a pretty cohesive way to weave this into your organization. and I know a lot of you guys are probably going to look at what I've done and be like, "Huh, we could make that better." and that was sort of sort of the idea. I'm sure there are many ways you can make this better, and I'm keen to see, you know, what you guys do with us. Since implementing this inside of Left Click, Clervo, and Maker School, you

know, my day-to-day responsibility has also shifted a fair amount, too. Rather than necessarily be involved with every little thing, I've sort of had to go up one level of abstraction, and then be responsible for the the fleet of managed agents, essentially. I've also had to start thinking about things like my token budgets. These are things that I never used to think about before because I was typically constrained to operating within one little like terminal window, having some input go in, and then waiting for the output to come out. I've also become significantly more productive because no longer do I have to just sit and wait for an output.

you know, I can just fire off outputs basically about as fast as the speed of thought. And I think that's a major unlock. You have like a little ideation session, you sit down, ideate the top five tasks that need to be done, assign agents to as many of these as can be done. They do a bunch of pre-lifting while you focus on something else, and then maybe three or four times throughout the day, every couple of hours, you check back in on the outputs.

that sort of gets you out of just staring at the terminals all day. Actually gets you, I don't know, building relationships with key vendors, selling, which is obviously the number one lever that any founder or, you know, top-level executive in a company really needs to be doing, and so on and so forth. Okay, if you like this sort of thing, I actually have a bunch of files down below in the description, in the line that says Maker Zero, where I you can just copy and paste everything that I write there to an agent, and you can have it do it just like that transcript idea earlier. I also want to make it clear that this isn't prescriptive. I'm not telling you have to do it all linear, you have to do it in a particular way. But this is just how I'm managing it in my organization and the results that I'm seeing. And if you want to learn how to implement strategies like this into an AI and automation service business, basically a business where you get paid to build things like this for other people, definitely check out Maker School, as well. It's my 90-day automation road map, where I actually guide you through everything you need to do to start as a total beginner and actually get your first client. And I guarantee you that first client in 90 days, or I give you all your money back. It's very straightforward, it's very tried and tested, and it's been finished by over 10,000 people to date. Thank you very much for watching. I'm looking forward to seeing all y'all in the next video.