



AI Agents:

The Shift From Answering Questions to Getting Work Done

AI Agents are redefining productivity by taking action, orchestrating tasks, and delivering measurable outcomes



Most people are still using AI like a very smart colleague **sitting across the table.**

AI Agents are redefining productivity by taking action, orchestrating tasks, and delivering measurable outcomes

THE OLD WAY

Chatbot Q&A Loop

01

Ask a question

You prompt the AI with a task or query.

02

Read the answer

You review the response and decide what to do next.

03

Give another instruction

You manage every step of the workflow yourself.

Useful, but you are still doing one important job: managing the workflow.

THE NEW WAY

Agent Acts Autonomously

01

Set a goal

Give the agent an outcome to work towards.

02

Provide tools & access

Define what it can use and what it's allowed to do.


03

Agent takes action

It works autonomously, no prompt-by-prompt hand-holding.

That distinction is more important than all the terminology surrounding agents.

A chatbot answers. An agent acts.



Consider a Simple Research Request.



The chatbot approach

You could ask ChatGPT or Claude:

"Tell me about the top competitors in this market."

You receive an answer.

You might then ask: "Check their pricing."

Then: "Look for recent product announcements."

Then: "What are customers complaining about?"

Then: "Compare all of this and give me a one-page briefing."

The AI is doing a lot of useful work.

But you are the project manager.



The agent approach

Now imagine giving an agent one goal:

"Prepare a competitor briefing for my management meeting tomorrow. Identify the top three competitors, check current pricing and products, find recent developments, analyse customer complaints, verify important claims against reliable sources and prepare a one-page comparison."

The agent can break that goal into tasks. It can research each company. It can use search. It can open sources. It can compare findings. It can reject weak evidence. It can assemble the briefing.

You asked for an outcome, not five prompts.

That is the beginning of agentic AI.



What actually makes something an agent?

Honestly, we sometimes make this more complicated than it needs to be. A useful way of thinking about an agent is:



Goal

A clear objective the agent is working toward not a prompt, but an outcome.

Tools

The capabilities it can call on search, code, APIs, files, and more.

Memory

The ability to retain context across steps, so it doesn't start from scratch.

Action

It doesn't just respond; it does things in the world, autonomously.

Check

It evaluates its own output, rejects weak results, and iterates toward quality.



Goal and Tools

An agent only as effective as the clarity of its goal and the access to the right tools



Goal

An agent needs to know what it is trying to achieve.



"Analyze this" isn't much of a goal.



"Identify invoices overdue by more than 60 days, prioritize the accounts by value and collection risk, check the latest customer interactions, and prepare recommended next actions for the collection team" is much better.



The clearer the outcome, the easier it becomes to give the agent sensible boundaries.



Tools

Intelligence without tools can only advise.

An agent becomes useful when it can work with the systems where the job happens.

Depending on the role, those tools might include:



Web search



Databases



Company documents



APIs



CRM



Code



Databases



Calendars



Email



Accounting software




WhatsApp

The agent shouldn't get access to everything. It should get the minimum access required to perform its job.



Memory and Action

An Agent doesn't just think, it remembers and acts. That's what makes it truly effective



Memory

Imagine hiring an employee who forgets yesterday every morning.

They wouldn't be very effective.

Agents have a similar problem.

They need to retain the information that matters:

- ☐ What has already happened?
- ☐ What decision was made?
- ☐ What has already been tried?
- ☐ What worked?
- ☐ What failed?
- ☐ What is still outstanding?

This doesn't mean storing every conversation forever.

Good memory is selective.

It preserves the state required to make the next decision.



Action

This is where we move beyond copilots.

An agent can do something.

- ☐ Create a CRM lead.
- ☐ Check an appointment.
- ☐ Prepare a payment recommendation.
- ☐ Read an invoice.
- ☐ Update a case.
- ☐ Generate a report.
- ☐ Send something for approval.
- ☐ Escalate an exception.

The level of authority matters enormously.


An agent being technically capable of an action doesn't mean it should automatically be allowed to perform it.

Check

This is probably the part we care about most.
An agent shouldn't say:

"I've completed the task."

just because it generated something.
Completed according to what?
You need a definition of success.



Did the figures
reconcile?



Were all required
fields captured?



Did the source actually
support the claim?



Did the transaction remain
within policy?



Was an appointment really created in
the system?


Agents become far more useful when **"done"** means something **measurable**.



A finance example

Imagine a company receiving 3,000 supplier invoices every month. A traditional system extracts fields and pushes them to the ERP.

An agent can go further, end to end.



The same agent then decides what to do next, based on the situation.

Normal Invoice

Workflow continues automatically.
No human intervention needed.
The agent processes and posts without delay.

₹10 Lakh Discrepancy

Agent stops the workflow and sends the evidence package to finance for human approval before proceeding.


Same agent. Different level of authority based on the situation.
That distinction becomes very important in real businesses.

Now consider healthcare

A patient sends a WhatsApp message:

*"Can I see Dr. Arun tomorrow evening?
He operated on me last month and asked
me to come back after four weeks."*

A normal chatbot can understand the question.
An agent can potentially complete the entire
journey, end to end.



Normal Follow-Up Request

Agent completes the entire booking journey automatically; no human intervention needed.



Urgent Clinical Signal

The agent stops immediately. A properly designed workflow recognizes the exception and escalates to the appropriate human or clinical process; not a booking form.



Sometimes the smartest thing an agent can do is stop.

Manufacturing gives us another example

A camera detects a defect on a production line. That's computer vision, useful, but the story doesn't have to end there.

An agent can take the detection and investigate:



Which batch is running?



Did the defect begin after a machine setting changed?



Was there recent maintenance?



Has this problem happened before?



What does the SOP require?



Who is responsible for this line?



Does production need to be paused?



Agent

Assembles all the evidence such as batch data, maintenance logs, SOP requirements, prior incidents, in seconds.



Human Supervisor

Reviews the assembled picture and makes the consequential decision, whether to halt production or continue.



That's a pretty good division of labour.



There are different levels of agents

Not every company needs fully autonomous agents tomorrow morning.

Think of this as four levels.



1 

Level 1: Chat

You ask. AI answers.
You control every next step.



2 

Level 2: AI with tools

AI can search, read documents, query databases or run code when the task requires it.



3 

Level 3: Multi-step agent


You provide the goal. The agent plans the work, uses multiple tools, executes several steps, checks its results and returns the finished outcome.



4 

Level 4: Triggered agent

The agent doesn't need someone to start each job. An event starts it. A new invoice arrives. A customer sends a WhatsApp message. A payment becomes overdue. A quality alert appears. A scheduled report becomes due. The agent wakes up, performs its bounded role and either completes the work or escalates it.



Companies don't have to jump from Level 1 to Level 4.
In fact, it is not recommended.



Where should companies start

Start with one piece of work.

Not *"let's build an AI employee."*

Not *"let's automate the finance department."*

Pick something boring.

Something repetitive
Something measurable. Something where people are currently spending time moving information between systems, checking the same things repeatedly or following a reasonably consistent decision process.



Then run it manually with AI first. See where it fails. See what information is missing. See where humans genuinely add judgment. See which exceptions appear. Only then give the agent more tools and more authority.

Apply a simple test before building any business agent:

Is there a clear goal?

Can the agent access the information required to achieve it?

Can we objectively check whether it succeeded?

If one of those answers is no, be cautious.

Don't automate a broken process.

AI doesn't magically repair unclear responsibilities, bad master data, missing policies or poorly designed workflows. It may simply make the confusion happen faster.



Humans still matter

There's a strange assumption in some AI discussions that the destination is "no humans."

The better question is:

Where is human judgment valuable? Let's divide agent decisions into three zones:



Low-Risk, Reversible Actions

Can be autonomous.

An agent can categorize an enquiry without asking anyone; the action is cheap to undo if wrong.

Medium-Risk Actions

Prepared by agent, approved by human.

It may prepare a ₹5 lakh credit adjustment but require human sign-off before posting.

High-Risk, Regulated or Irreversible

Remain human-led.

No agent should independently make a clinical diagnosis or sign a legal agreement, even if the model can produce an answer.

Good agent design isn't about giving AI maximum freedom. It's about giving AI the **right** amount of freedom.

From one agent to a digital workforce

Once one bounded agent works reliably, businesses can start defining additional digital roles. They don't all need to be connected; each can stand alone with its own inputs, tools, and rules.

A healthcare company might eventually have:

- ☐ AI Receptionist
- ☐ Enquiry Agent
- ☐ Care Coordinator Agent
- ☐ Scheduling Agent
- ☐ Billing Agent
- ☐ Patient Follow-up Agent
- ☐ Quality Agent

A finance function might have:

- ☐ Invoice Agent
- ☐ Collection Agent
- ☐ Reconciliation Agent
- ☐ Expense Review Agent
- ☐ Tax Agent
- ☐ Reporting Agent

Each agent should have a job. Inputs. Tools. Permissions. Rules. Outputs. Checks. Escalation paths. And an accountable human owner.

That's beginning to look less like "using AI" and more like **designing a new workforce** where people and digital workers operate together.

The technology isn't the hardest part

Models will keep getting better.

What will separate successful agent deployments from expensive demonstrations will be everything surrounding the model:



Process design



Permissions



Data quality



Memory



Integration



Verification



Exception handling



Monitoring



Cost control



Human accountability

This is why the conversation about agents needs to move beyond *"Which model should we use?"* The model matters. But the **operating model matters more.**

One definition to remember

If all the terminology around agents becomes confusing, come back to this:

Agent = Goal + Tools + Memory + Action + Check

Build from here. Give the agent the minimum authority it needs. Measure whether it genuinely works. Build trust through evidence. Then give it the next job.



A chatbot helps you think.

An agent helps you get work done.

The smartest starting point isn't maximum autonomy, it's one valuable, clearly defined job.

That's how businesses should approach the agent era. One workflow at a time.



Let's build what's next.

Partner with us to accelerate transformation and deliver measurable impact.



www.10xDS.com



www.10xDS.ai



Connect@10xDS.com