



The Agent Internet A New Evolutionary Phase of the Web and Its Global Impact

Dr.A.Shaji George¹, B.Roshini², Dr.T.Baskar³

¹Independent Researcher, Chennai, Tamil Nadu, India.

²Department of Computer Science and Business System, KPR Institute of Engineering and Technology, Coimbatore, India.

³Professor, Department of Physics, Shree Sathyam College of Engineering and Technology, Sankari Taluk, Tamil Nadu, India.

Abstract – Large language model (LLM) based software agents are transitioning from one shot assistants to long term actors that plan, negotiate and accomplish tasks with multiple services. This transformation is resulting in what researchers call the Agent Internet, the Agentic Web or the Internet of Agents, in which programs are full-fledged players alongside people. This article summarizes academic surveys, industry studies and public market data from 2025 to 2026 to outline the multi-layered nature of this setup and to consider its pros and cons. According to market estimates, the revenue for agent software is estimated to be around eight billion dollars in 2025 and up to 35-53 billion dollars in 2030. Adoption statistics tell a tale of much experimentation and few pilots taking it to scale. We discuss communication protocols that enable agents to find and trust each other, the as yet undeveloped identity and payment layers and the burden of agent traffic on advertising supported content. But the challenges that we see are not just capabilities, but security, accountability, interoperability, and who pays for content and compute. Finally, we highlight some of the current research and policy gaps and suggest avenues for a responsible construction.

Keywords: Agent Internet, Internet of Agents, Agentic Web, Multi-agent systems, Agent communication protocols, Digital economy, Trust and identity, Content monetization.

1. INTRODUCTION

Until recently, the web was an entity with a man or woman behind a keyboard. The pages were created for people to read, the search engines sorted links for people to click on and the advertising paid for it in case a reader saw the ad. That's now changing. Large language models and vision language models can power programs to read a page, compare the options of multiple sites, fill out forms, maintain a multi-step plan in working memory and execute actions to achieve a goal defined by its owner. As these programs start to communicate with each other and with online services at internet scale, they are referred to as the Agent Internet, the Agentic Web or the Internet of Agents.

Timing provides weight to the subject. From 2024 to 2026, the components required to enable agents to work inter-organisational were developed from prototypes to published specifications. Within a short time, open proposals emerged on how tools can be ingested, typed messages can be shared, agents can learn about each other's capabilities, and agents can verify themselves. A 2025 survey categorized four of these proposals, based on the problem they address, ranging from access to local tools to discovery of open networks via decentralized identifiers. Meanwhile, market research companies started to provide revenue

projections and giant consultancies began to give out warnings regarding the proportion of projects that would fail.

On the ground, it's not even. According to one popular annual survey, 40% of companies with over \$1 billion in revenue indicated that they were scaling in 2026, compared to 27% in 2025, while smaller companies remained unchanged at 22%. In the study of enterprise deployments, about ninety-five percent of organizations reported that they did not see a measurable impact on their P&L from generative systems, and the value is in a small minority of organizations. These are not mutually exclusive numbers. They talk about a technology that is good in restricted applications, but does not work well when pushed to wide, unsupervised applications.

The Agent Internet (Internet of Agents)

From websites for people → to a web where AI agents can read, decide, act and work together.




Fig -1: The Agent Internet

This article views the Agent Internet as a socio technical change, and not as a product launch. By the end of the article, readers will have a solid understanding of the different layers that enable agent-to-agent activity, the monetary logic that will determine whether or not it continues, and the potential pitfalls that could slow it down. It is based on sources that are all verifiable, and it calls attention to the differences among the sources, not glosses over them. The names of companies are reserved and the emphasis remains on the mechanisms and evidence, rather than on any one company. There are three questions that are followed throughout the pages that follow. What needs to be constructed to enable agents to safely operate across trust boundaries. Who benefits and who does not as agents shift from one person to another and the services they utilize. But what research and policy are yet to resolve. The responses suggest a new level of infrastructure in the internet that might make many online activities more affordable and subtly change the flow of content, attention, and money on the web.



2. OBJECTIVE OF THE ARTICLE

This article aims to provide a layman, but knowledgeable, reader with a solid description of the Agent Internet. It sets out to describe the concept clearly, explain the mechanics, provide an accurate picture of the size of the market and adoption rates without hyperbole, and identify the security, legal and economic issues that will need to be overcome before it becomes widely used. It also attempts to distinguish between what is reported and what is marketed, enabling a policy maker, student or professional to be able to discern between a measured claim and a marketing claim.

2.1 How the Information Was Gathered and Checked

The content of this material is from three types of sources. The first are academic papers published in 2025 and 2026, which are peer-reviewed and preprint works, including survey papers on agent interoperability protocols, and on the security and privacy of agent networks. The second is industry research, which consists of annual surveys of enterprise use, market-size estimates from research companies and network-level measurements of web traffic. The third is public comments from infrastructure providers and standards bodies about changes in the means of access and payment for content. All quantitative claims were followed up with a cited source and a cited date. If there is more than one forecast, a range of forecasts is given instead of a single forecast. Market forecasts, in particular, are dependent on the method and vendor and hence this article provides the range of forecasts rather than selecting one of them. Claims which could not be corroborated with a dated source were excluded. There are two limits which the reader should note. The differences in the headline figures in the adoption surveys is due to the different definitions of what constitutes an agent in production, and market forecasts are based on assumptions of growth.

3. HOW AGENTS DIFFER FROM THE BOTS THAT CAME BEFORE

Older automation worked on the basis of fixed rules. A script was scheduled to run at a certain time, a chatbot was used to match keywords to canned responses, and a web crawler was used to get pages in along a certain path. None of these were able to alter plans during the execution of a task or reason about a goal not explicitly mentioned in a plan. Not only in degree, but in kind is the current generation different. A model-driven agent is able to read an instruction in natural language, decompose it into steps, select the tool service to invoke in each step, see the outcome, and adjust its strategy.

There are three technical advances which made this possible. Function calling and structured tool use enable a model to translate its reasoning into real-world function calls the model can request a report or query a database, instead of just describing what to do. Reasoning and acting loops enable the model to reason and then act, and then reason again and act again, etc., verifying each action. Persistent memory allows an agent to retain information from one session to the next, such as a user's preferences, or the status of a long task, without having to start from scratch again.

These skills are not free, as they were in the case of plain chatbots. An agent which can act can also do it wrong and can be directed by text it reads. A national request for comment on current agent development tools found that while there are some basic safeguards, including tool allowlists, sandboxed execution, and filtered inputs, most of the tools lack mature capabilities to separate privileges among agents or to authorize one agent to act through a chain of other agents. The capabilities came before the controls, so much of the recent work is about not so much as what agents can do, but how to keep what they do within safe limits.

How Agents Differ from the Bots That Came Before

From fixed rules to flexible intelligence — AI agents can read, reason, use tools, remember, and adapt to achieve your goal.



Fig -2: How Agents Differ from the Bots That Came Before

The bottom line for a reader is that an agent is a program that has judgment and reach, but for which there is no accountability. It is as if it took its authority and its guilt from the one or the group that operates it. All of the layers below are in existence to make that borrowed authority legible, limited and reversible.

4. THE LAYERS THAT HOLD AN AGENT WEB TOGETHER

The Agent Internet can be thought of as a stack, each layer being able to solve a problem that the layer below it cannot. Compute, storage and networking at the base with some reasoning occurring locally and some occurring in a data center. It's not just a matter of personal taste where a computation is performed. It determines latency, cost and the amount of private data that is outside of the user's control. Over the base is the agent layer itself, which consists of personal assistants, workflow agents within the company, agents that are integrated into physical devices, or experts who are specialized in a particular industry like finance, health or logistics. All have a combination of perception, planning and tool calling abilities. These agents are seldom used singly. A multi-agent system is a cooperative system of multiple agents that can be used to accomplish a single task within a single trust boundary. If agents of different owners and different platforms have to cooperate beyond these borders, an interoperable multi-agent system emerges, that a 2025 architectural survey considers as the technical implementation of the Internet of Agents.

Communication layer provides agents with a common means of discovering one another, characterizing their capabilities, sending messages and negotiating. Otherwise, you'd have to have a custom integration for each pair of agents, which isn't scalable as the numbers increase. The trust and identity layer provides an answer to a more difficult question. How is the agent on the other side to check the agent's claim of representing a bank or a hospital and determine what to allow. The proposals here are based on

cryptographic identifiers, verifiable credentials and audit logs of what was done by whom on behalf of whom.



Fig –3: The Layers That Hold An Agent Web Together

The economic layer is located close to the top. When agents are able to move, they are also able to buy, sell and bargain, requiring payment, reputation and rules to make sure that an agent's actions are aligned with its owner's interests. None of the layers are completed. The areas with the fastest progress are communication, with the slowest being trust and economics the opposite of what a safe rollout would want.

5. PROTOCOLS THAT LET STRANGERS AGENTS TALK

The communication layer has been the most tangible and a survey in 2025 has clustered four proposals according to the problem they try to solve. It's interesting to understand the pattern as it reflects the way the field is thinking. The least problematic is the access to tools. Structured remote calls using a client-server design allows an agent to pull in context and call tools within a common trust boundary. It supports the scenario of an agent of one application accessing the data and functions around it, and provides a clean session life cycle of setup, operation and shutdown.

The next challenge is enhanced messaging. A design that is based on typical web request patterns has multi-part, mixed-media messages as well as immediate and long running, streaming replies. Great for local groups of agents who need to share images, files and structured data back and forth while a human or other system monitors the communications. The third issue is cooperation between agents which are not in a trust relationship. Peer to peer means one agent can pass a task to another agent, in a machine-readable card, specifying what each agent can do, so that delegations can take place across company lines without revealing the internals of either agent. Open discovery is the most prevalent issue. A

decentralized approach is based on identifiers that are managed by an open web standard and linked data descriptions, enabling an agent to discover and verify another agent that it has not met before, via an open network, not a walled environment. It's the layer that might enable open marketplaces of agents and is the least mature.

Protocols That Let Strangers' Agents Talk

Standard ways for AI agents to find, message, share tools, and delegate tasks — even across different companies.
Like the internet for people, these protocols help AI agents work together safely, step by step.



Fig -4: Protocols That Let Strangers Agent Talk

There is an emergent useful order. If you're new to the tools, then structured messaging, then task delegation, then open discovery. That's the order of trust needed for each step, as it's easier and safer to talk to your own tools than to trust a stranger's agent on an open network. This advancement creates a vulnerability that is security depth. These protocols address primarily low-level issues like authentication and transport encryption, a 2026 analysis said. They do not yet address higher level issues like secure delegation, trust relationships between agents and what to do with privilege when one agent is delegating through a chain of agents. Safety valves are not in place yet plumbing is in.

6. PROVING WHO AN AGENT IS AND WHO IT WORKS FOR

The Agent Internet either gains or loses trust in the area of identity. A second factor and a password can be used to check a person visiting a site. An agent that acts for that person must be able to prove its own identity, and the extent of authority it has on behalf of the person as well as the boundaries of that authority. A few of the proposals employ identifiers under an open web standard, and credentials that can be verified by a relying party without a call back to a central registry. There is a real-life risk at stake. Agentic systems are the biggest growth in the enterprise attack surface since the shift to cloud computing, according to a 2026 report based on a survey of 1,625 tech and security leaders. It discovered that 80% of leaders said they had good visibility into agent activity, but 86% said that they would see agent growth outpace their security

controls in a year. It also explained that there is an explosion of non-human identities, many of which are not multi-factor protected, and a lack of ability to roll back an agent's unintended actions. Not having a **undo** button is a known issue.

Proving Who an Agent Is and Who It Works For

An agent must prove its identity, show who it works for, and only do what it is allowed to do.
Every step is recorded, verifiable, and can be revoked if needed.



Fig –5: Proving Who an Agent Is and Who It Works For

Delegation clarifies the issue. If a user authorizes an agent, and that agent passes on part of the task to a second agent, and that second agent passes on part of the task to a third agent, the original authorization of the user continues to the third agent, with whom the user has never dealt. Every hop must have a record of the permissions and who gave them, and each must decrease, not increase, the permissions delegated. A number of sources say that current development tools don't cope with this well, and provide allowlists and sandboxes, but not fine grained control of privilege as it passes down a chain. The simple moral for a reader is that an agent's identity is a kind of "borrowed permission. What is important about the whole layer is that it makes that record honest, checkable and revocable. A verifiable identity that can't be revoked quickly is almost useless, since the whole idea of an identity in an adversarial environment is to be able to terminate use of an identity if it's compromised or used improperly. Part of the reason it lags, is that it's slower and less glamorous to build that revocation and audit capability, than to build the agents themselves.

7. THE MONEY LAYER OF PAYMENTS, REPUTATION, AND NEGOTIATION

The convenience is an agent that can read only. If the agent can afford to pay, then it is an economic agent and that alters the design problem. After agents transact, the web must provide methods for them to pay

small amounts, to establish and view reputation, and to negotiate pricing, access and service conditions without having to have a human verify every action.

The Money Layer of Payments, Reputation, and Negotiation

How AI agents can pay, build trust, and negotiate on the web — automatically, safely, and at internet scale.




Fig -6: The Money Layer of Payments Reputation and Negotiation

Scalability questions arise when it comes to payments for agents, questions that human checkout never had to consider. If the agent purchases on behalf of the user, then the user should not be allowed to exceed the spending limits one or both category rules, this is because the agent can be manipulated by the content it reads. Reputation is as important. In an open network, where an agent may be dealing with counterparties with which it has not previously had any interactions, it is the hard to fake record of past behavior that enables one agent to trust the other's claim or turn down the deal. The negotiation is the most different aspect, compared to today's fixed pricing. In theory, agents will be able to negotiate price and service level in real time, indicating a potential for data, compute and small services micro-markets. This is both good-looking and dangerous. It might be able to balance supply and demand more closely than menu pricing, and it might be able to help well-resourced agents secure terms that a human would not accept or for two agents to collude in a manner neither owner would have intended.

Content is under a live warning. Some agent crawlers were found to be returning tens of thousands of pages for each visitor they sent back, with one provider returning nearly seventy thousand pages for each visitor and another returning nearly one thousand pages for each visitor. Where this happens, and the return flow is insignificant compared to the consumption, the old barter of content for traffic breaks down and some form of payment must be substituted or the quantity of content will decrease. At least one big infrastructure provider had started to default to blocking training and agent crawlers from pages with ads by late 2026, on the assumption that an ad is a human eye signal that the page was priced for human eyes. In short, the money layer is a must have. It is the lifeline of the web that agents depend on if it is not there, the parts of the web that agents depend on lose the means to pay for their own upkeep.

8. WHAT THE MARKET NUMBERS ACTUALLY SAY

The revenues are big numbers and should be interpreted with caution as they come from various methods. According to three research companies, the agent software market size is estimated to be around eight billion dollars in 2025. Their 2030 numbers vary from approximately \$35 billion to \$53 billion, with compound annual growth rates of approximately 43 percent to 46 percent. One consultancy proposed a minimum bottom line and a maximum top line, implying that the autonomous agent market can be as large as thirty-five billion dollars by 2030 and maybe forty-five billion dollars if organizations coordinate agents well and manage the risks.

What the Market Numbers Actually Say About AI Agents

Big opportunities. Real progress. But also real risks. Here's what the latest research shows — in simple terms.




Fig -7: What the Market Numbers Actually Say About AI Agents

That upside came with a stark warning, too, More than 40 percent of the projects agentic work is currently doing could be dropped by 2027 due to unanticipated cost, scaling challenges, or unforeseen risk, the same consultancy estimated. That is NOT a footnote. A big portion of the spending that the forecasts rely on could not withstand the production test, it says. Numbers of adoptions complete the picture and tellingly, disagree. One prediction predicted that by the end of 2026, 40% of enterprise applications will have task-specific agents in them, compared to less than 5% at the time of the forecast. Fifty-two percent of the 3,466 senior leaders surveyed across 24 countries indicated that their organization was using agents in production and 39 percent had deployed more than 10. However, a research of three hundred deployments discovered about ninety-five per cent of organizations experienced no quantifiable profit or loss impact, while value was focused in approximately one in twenty integrated efforts.

These are not so contradictory numbers as they seem. The presence of an agent in an app does not equate with an agent carrying real load and reporting productivity gains doesn't equate to the agent booking profit. The division was captured in a single annual global survey 80% of respondents said the technology

enhanced their individual output and half said it helped them make better decisions, with the proportion of firms saying it had a real impact on earnings remaining around the same level as the previous year. The pattern is due to the business logic behind it. The revenue streams in this market are from the sale of agent software, the sale of the compute that agents utilize and the sale of tools to monitor and manage agents, and the sale of content and services that agents desire. The first is one that's gaining momentum. The second is the amount that firms that ultimately reach production invest in evaluation and monitoring, which is higher. The third is unsettled, and is treated in the next section but one.

9. THE GAP BETWEEN PILOTS AND PRODUCTION

The one thing that is consistent in 2026 surveys is the disconnect between project initiation and scaling up a project. One survey of enterprise technology leaders found that 78 percent had at least one agent pilot, while 14 percent had scaled an agent to organization wide operation more than half of the workload that the agent is targeted for. Another put 83 percent of the projects into production and 41 percent of the production deployments into another story, a 42-point difference that the researchers dubbed the story of the year.




Fig -8: The Gap Between Pilots and Production

The reason for the stall is now well understood and is not due to the quality of the model. The leaders interviewed cited governance as an issue ownership is not clearly defined across departments audit infrastructure is lacking, there is no clear escalation pathway to a human reviewer. One survey revealed that the total amount of money spent by organizations at production was not more than money spent by the organizations in pilots. They allocated their resources differently, investing more in infrastructure for



evaluation, monitoring and staff to run the system and less in the selection of models and in timely tuning. That reading of the problem isn't an issue of underinvestment, but rather a build-versus-operate imbalance.

An impressive contemporary counterpoint is found. Many organizations indicated that the agents needed to be monitored more manually than what they were intended to save time on. If so, then that means that a lot of deployments are net negative on labour today, which is precisely what you would expect from a technology that is behind the eight ball on control tooling. It is the firms that make operation not the demo the hard part that are the winners. The sector breakout provides some detail. The production rates for financial services were the highest, almost twenty-one per cent in one survey, with early work on document processing and compliance. The lowest was in healthcare, at almost eight per cent, where regulatory prudence is exercised in clinical decision making. In between were manufacturing and retail. As much as anything else, the pattern reflects risk tolerance and the cost of error as opposed to technical readiness, and in this regard, the Agent Internet has brakes that are as much organizational and legal as they are computational. If you're considering the hype, the straight scoop is this. Agents are rapidly being adopted and applied in a limited fashion. They do more to assist people than they do to balance books. Where money and most of the disappointment now resides is between a working pilot and a reliable production system.

10. WHAT AGENTS SAVE PEOPLE AND FIRMS IN EVERYDAY WORK

The most obvious current advantage is time back. An agent that can search, compare, summarize and then act on multiple services eliminates the little, repetitive hassles that make up a work day, Reconciling invoices, triaging alerts, drafting first drafts of routine documents, getting data from multiple systems into one place. One vendor analysis found that agents who do invoice reconciliation and alert triage spend over sixty percent less time on manual tasks in the cases they analyzed. More general survey evidence is more conservative, showing productivity effects on some tasks to be big and on others to be small, and even negative on some tasks for some workers.

The variance is the true part of the story. The more well-defined, repeated and checkable the task is, the more agents help, and the more open-ended and the higher the cost of a wrong answer, the less they help. A study of 12,000 European companies revealed a 4% increase in labour productivity as a result of the adoption, and an even bigger increase where employees were trained. That's a valuable increase, albeit a small one, and the emphasis on training implies that the benefit is to be derived from the person learning to use the tool, rather than the tool supplanting the person. The advantage for firms is that agents are able to maintain and take action on context that a static workflow can't. With a rigid script, a change in the form or a late supplier will cause a break in the script. An agent can, to a certain extent, modify. That's important for businesses with frequent minor changes, such as supply chains or customer support, where the flexibility of re-planning on the go outweighs sheer speed.

Less work is not the benefit; more about doing different work. The human jobs that expand as agents take on routine research, booking and first draft work are goal setting, output checking and addressing cases agents tag as **uncertain**. This changes the type of skills that are required to focus on judgment, oversight and designing the agents themselves. A bigger proportion of respondents 39 percent to one large survey believe that their agents will reduce their headcount in the next year, compared with 32 percent, and a larger proportion believe that they will not change their headcount. Notably, this year's expectations of a

decline were roughly double the actual decline, which occurred the year before, as a reminder that fears of displacement have so far been greater than reality.

What Agents Save People and Firms in Everyday Work

AI agents use software tools to search, compare, summarize and take action, so people can spend less time on repetitive work and more time on higher-value work.




Fig -9: What Agents Save People and Firms in Everyday Work

The measured view is that agents today are a means for improving individual productivity and changing the character of routine work; and that gains at the firm level remain skewed towards the few firms that have overcome the operating challenge.

11. WHO GAINS WHEN COUNTRIES PUT AGENTS TO WORK

The argument for agents at the national economy level is the same as that for the previous generations of software the cost of coordination is reduced and labor and capital can be freed up for other purposes. It is where it really counts that this is relevant in the services that governments and citizens are in constant contact with. Agents who locate the benefit that the individual is entitled to, fill out a tax return or simply direct the individual to the appropriate office might shorten the time and aggravation involved in public services, and decrease the mistakes generated by misunderstanding forms. There is some macro-level evidence that is worth citing and is early. A national index of measures related to AI recorded nearly double the average growth of the previous decade, reaching 2.7 percent in 2025 in the United States. This isn't entirely agents fault as there are many factors that drive productivity, but it's in line with a time of fast adoption of the underlying technology. The European firm-level productivity increase of 4% from adoption is in the same direction, but with the proviso that training is important.

The distribution of gains is not a given, but a policy choice, for countries. When agents are primarily found in larger organisations with the capacity to implement and maintain them, the advantage becomes focused and the disparity between larger and smaller organisations becomes even greater. This is already

reflected in the 2026 survey results 40% of large firms compared to 22% of smaller firms are scaling agents, and this gap hasn't been closing. If a nation is interested in achieving widespread gains, it must make tools for agents affordable and secure enough for small businesses and public agencies to use, which favors public standards and infrastructure over each organization building their own.

Who Gains When Countries Put Agents to Work

AI agents reduce the cost of coordination, make public services easier, increase productivity, and create more opportunities for citizens, businesses and governments.




Fig -10: Who Gains When Countries Put Agents to Work

Life goes on in more subdued ways. When routine, digital chores get reduced, people find time and mental space that the current web takes away. Citizens who are now unable to deal with complex portals are the ones who benefit most if public services are more easily accessible via an agent who can convert a simple request into the appropriate actions. The downside to this is that it can lead to addiction. When a population is dependent on agents to access the services they need, a failure of agents, manipulation of agents, or the inaccessibility of agents due to pricing, can make a population vulnerable, thus, national adoption and national resilience must be planned together, not in sequence.

12. ACCESS FOR PEOPLE THE OLD WEB LEFT BEHIND

One of the more down-to-earth advantages is ease of access. The current Web requires a lot from its users knowing how to read a specific language, being comfortable with complex interfaces, and patience in taking the multiple steps required to complete the process. The three barriers are reduced simultaneously by an agent that converts an intent, expressed in speech or writing, into the appropriate series of actions. Limited sight, limited mobility, limited digital skill and limited command of the language a service is written in can mean the difference between reaching a service and giving up, when an agent is available to serve as a patient intermediary.

This isn't a small community. The speed and complexity of digital services is a barrier to many in ageing societies, especially, and an agent that takes care of the mechanics while the person says the goal is a way of regaining a degree of independence. This is also true of individuals with chronic diseases who need to schedule appointments, share records and prescriptions with systems that aren't communicating. It's an agent that puts those systems together on behalf of the user that does work that the platforms themselves have not made easy.



Fig -11: Access for People the Old Web Left Behind

The environment question is more balanced, and calls for honesty, not optimism. The upside is that agents which coordinate across systems could help cut waste in the physical processes they interact with, from optimizing logistics routes, to better aligning energy supply and demand, and coordination is what agents excel at. This is offset by the cost of the models' energy. Agent workloads are more than just simple requests, an Agent can think in a few steps, call a few tools and can run in the background, not just on demand. Whether or not the net impact on energy consumption is agreed upon and will depend on whether the savings in the physical world that result from the coordination outweigh the cost in compute of the reasoning that achieves those savings. Anyone who thinks that they've got a clear environmental winner at this stage is ahead of the facts. The fair reading is that agents provide a genuine improvement of access for those for whom the current Web is not a great tool, and an unanswered question on energy. This benefit of accessibility should be measured as it increases, not assumed, and the environmental accounting is not as apparent as the accessibility.

13. THE SQUEEZE ON CONTENT AND THE SEARCH FOR NEW PAYMENT

The most obvious area of Agent Internet that is causing harm now is its impact on agents' reading material. The web was based on an unstated deal. Creators opened up their pages to search, search drove people

to the pages, and those people viewed the ads, clicked on the affiliates, and performed the analytics actions that got creators paid. Each of these links assumed that there was a person on the page. An agent can break all the links at once, as it takes out the answer and gives it to a user elsewhere and then there is a log entry.

The Squeeze on Content and the Search for New Payment

How AI agents changed the web, why it hurts creators, and the solutions being built.




Fig -12: The Squeeze on Content and the Search for New Payment

The figures are shocking. According to one estimate, the percentage of zero-click search no site visited increased from fifty-six percent in May 2024 to sixty-nine percent in May 2025. One research firm found click-through rate dropping from fifteen percent to eight percent or a relative drop of nearly forty-seven percent, when AI-generated answers were present in results. Randomized experiments showed that there was an approximately forty percent decrease in outbound organic clicks when an AI answer was displayed and aggregate data showed that AI answer features were correlated with an approximately twenty-five percent decrease in referral traffic. Losses from publishers have been reported to be as high as 90 percent or as low as 20 percent; the smaller sites are hit the hardest.

The unit economics provide the rationale as to why this cannot go on unchanged. The publisher's costs for hosting, writing and editing are identical for human and machine readers. However, all of the revenue events had to involve a human. Agent traffic costs money, but doesn't generate any revenue. No business can supply at a negative margin supply side repricing. There are a number of responses in the making. Some publishers are now tweaking their pages to be machine-readable in addition to human readable. Licensing agreements have demonstrated that there is a demand for paid access, but these are still ad hoc, and, by most accounts, don't yet offset lost referral and ad revenue. The latest concept is to pay the creator at the time of the fetch and only when the content is included in the agent's answer, with reporting back to the creator of which queries the content was a part of. Meanwhile, at least one large infrastructure provider has started to block training and agent crawlers from pages containing ads by default from

September 2026. It's not just about any one publisher. The open knowledge commons that agents rely on would be without a way of sustaining itself if it were no longer supported, and so would be the agents, who would have no source of answers. So, a payment mechanism that provides value to creators is not charity to publishers. It's ground maintenance, the entire setup is dependent on.

14. NEW OPENINGS FOR ATTACK AND MANIPULATION

An attack surface exists where software has the ability to act that does not exist in a reading only system. In a 2025 survey of the security and privacy of agent networks, the threats were broken down into four categories, and this breakdown is a good guide to the threats defenders now face.




Fig -13: New Openings for Attack and Manipulation

The first one is identity authentication. When an agent can pretend to be another agent, or pretend to be from a trusted institution, then all following decisions based on the agent's identity are undermined. That's why it's important that so much design effort is put into verifiable identifiers and credentials and why the absence of being able to quickly revoke a compromised identity is considered a serious problem.

The second is trust amongst agents. In an open network, an agent has to determine how much to trust another agent with which it has no prior interactions. A malicious partner may make false capability claims, or may attempt to poison a common context or may coordinate with other agents to influence a target's actions. Reputation systems are designed to help mitigate this, but there is a way to game reputation.

Third is physical device embedded agents' security. An attack on an agent controlling a robot, a vehicle or industrial equipment, can move from data loss to physical damage, making it much more than the typical breach.

The fourth is privacy. Agents build up context, preferences and histories to perform their tasks and that repository is a target itself. An agent that can remember a user to serve their well will also have a detailed record which can be stolen or misused.

The unique weakness is that agents can be directed by the content they work on, across all four. Prompt injection, both direct and indirect, is still an open problem and one 2026 audit discovered that the reasoning core of mainstream platforms remained open to prompt injection despite providers mitigating lower-level risks. An attacker who can put the text where the agent reads it has a foothold which no firewall can close because the agent's power is in what it reads. The level of concern is recorded. 86% of the leaders surveyed said that they predicted the growth of agents would outpace their security controls within one year, and most were unable to reverse an agent's actions. The honest and uncomfortable situation on the field is the defenders know that the threat is coming at them quicker than their defenses.

15. WHO ANSWERS WHEN AN AGENT CAUSES HARM

The one thing technology cannot fix and may have the most to do with the extent of Agent Internet is accountability. If an agent commits an expensive mistake, breaches a rule or does damage, the answer to **who is responsible** is never clear and the challenge increases as the delegation is passed along.




Fig -14: Who Answers When an Agent Causes Harm

Let's take a look at a basic chain. An agent is authorized by a user. To complete that, that agent delegates a portion of that to a second agent from different company which delegates to a third party service. Who is responsible if the consequence affects someone's well-being could reasonably be the user who set the goal, the author of the first agent, the operator of the second, the service provider, or the individual whose manipulated content caused the agent to go astray. Consumer protection, liability and data protection



laws were designed for human users and fairly static software. They do not work well when applied to persistent, adaptive programmes that work across organisational boundaries. This is no ordinary lawyer's puzzle. It's a direct impact on adoption and it is evident in the survey data. Most often cited reason why agent projects never get to production is governance, not model quality lack of clear ownership, lack of audit trails and no clear escalation path to a human decision. In essence, until a question is answered, organizations are refusing to deploy agents on a large scale and that question is Who will be responsible if something goes wrong. Before the answerability is, the technology is ready.

There are two ideas from recent normative work that point towards a path. The first is the delegation which is made visible and readable everything an agent does is accompanied by a record of the delegation and its boundaries that can be checked. The second is two-way transparency, that is, a service can see enough information about an agent to determine what access to give it, and an agent's owner can see enough about the service to trust it. These would all be part of the chain of responsibility and traceable afterwards, the prerequisite for a fair liability. An accountability problem lurks also within the market-structure risk. If only a few big guys can afford the audit and control machinery that safe deployment requires, then those guys become the gatekeepers that agent activity must go through and the power imbalance of the current web resurfaces in a new format. It is important that the design of accountability is also designed with an eye to small operators being able to comply with it, therefore, if they are not, the market will not remain open.

16. WHAT READERS GAIN AND WHY IT MATTERS FOR THE NEXT GENERATION

A reader who follows the evidence above will have several things to take away from the public discussion that is often confused. The first is a working definition which distinguishes agents from the bots that came before them an agent plans, selects tools, remembers context, and takes action to achieve a goal and in the process delegates authority which it is not held to. That's the only thing that makes identity, delegation and accountability the stuff of serious books, and raw capability a no-go.

The second is a sense of scale, which is calibrated. The predictions of the market are big and, even when they agree, they disagree from around thirty-five to around fifty-three billion dollars by 2030, and even the companies that make the prediction warn that a good portion of the projects in existence could be canceled first. Adoption is real, but shallow, and the impact of production and firm-level profits remains in the hands of a few. These numbers will enable a reader to distinguish an evidenced claim from a promotional one.

The third is an awareness of the location of the value and harm. It becomes apparent as time is restored to people and as flexibility in operations is gained that are subjected to frequent small disturbances. The damage is initially in the content economy, with measured referral drops of 20 to 90 percent putting at risk the very same content that agents rely on for their news and information, and then in an expanded attack surface that defenders know they cannot keep up with.

The practical relevance is a reality for the next generation. Skills that don't go out of fashion are the ability to set goals well, to check outputs and to supervise systems that act, not to do the routine steps that agents will learn. The rights to defend are the right to delegate access as one chooses, the right to see and revoke what an agent does on one's behalf, and the right of creators to receive payment if their work is used to provide an answer from an agent. These are not far off, theoretical concepts. They are being determined now, in protocol specifications, in infrastructure defaults, and in governance decisions of what projects ship.



Fig -15: How AI Agent Work and What They Mean for Our Future

The end result is the practice of measuring. On closer examination, nearly all assertions regarding agents, their productivity, energy consumption or job destruction are actually a spectrum, which varies according to the job, the industry and the effectiveness of the supervision. A reader who can ask what does this number mean. What is the range. When was this number reported. What is its definition. is more likely to be able to judge the Agent Internet than the reader who takes the biggest promise or the most dire warning on face value.

17. RESEARCH GAP

There are a number of holes in the current knowledge which are clearly evident naming them is the honest end to a survey such as this. The most significant is the difference between the security of the low-level protocol and the desired security of the higher level protocol that the Agent Internet requires. The published communication protocols include authentication and transport encryption, but currently lack secure delegation, trust between agents and privilege management across chains of agents, according to the 2026 analysts. Until those are addressed, the interoperable, open-network vision is too risky for widespread deployment and it is the open-network case the one where agents would be most useful that is least advanced.

The second gap is an empirical gap and not a technical gap. The reason for the huge discrepancies between the headline figures of adoption and productivity is that the surveys used different definitions of who is an agent in production and different samples. The current state of independent and longitudinal measurement of the effects of agents on labour, on markets and on content ecosystems is very limited and policy is being developed based on snapshots, not trends. The environmental accounting is even more



sparse there is no agreed upon value of the coordination savings that the agents generate on the physical level vs compute cost of the reasoning.

Third is the lack of legal and institutional. There is no clear answer on who is liable when a chain of delegating agents are involved, and the survey results indicate that this is the top obstacle to deployment. Delegation and transparency in both directions is being worked on, but has not been tried on a large scale or been codified in enforceable rules.

The fourth gap is related to the structure of the market. The tools needed to deploy agents safely, the evaluation, monitoring and audit infrastructure are costly and the risk that only a few large operators can afford them and thus become the gatekeepers of agent activity has been mentioned but not addressed. There is little research on how to make safe deployment affordable for small organizations public agencies.

Last, there is no solution that has been tested for the content payment problem. Pay-per-answer schemes and licensing have emerged, but it is too early to see if any of these schemes can provide the lost referral and advertising revenue on a large scale, and who pays for the open knowledge that agents rely on is an open question. There is a relationship between these gaps. The five issues that hinder the way to a working demonstration and a trustworthy public layer are security, accountability, measurement, market openness and content payment, and none of these is close to being solved.

18. CONCLUSION

The Agent Internet is not so much a faster Web, as a Web that's augmented by the presence of software that can be directed toward goals, that can be given authority to act on behalf of its owner, and that can be steered by what it reads. Most of that argument is in that description. It offers a reason why the focus of the serious work is on identity, delegation and accountability and not on what a model can do for a firm's productivity, and it offers a reason why so far, the technology has not boosted the earnings of most firms. The evidence collated herein is in support of a measured position. Agents save time in real time in well-defined, checkable tasks and they increase access for those for whom the current web doesn't work so well. The communication layer which allows agents to locate and converse with each other has been changed from prototype to published specification. The market is expanding rapidly, with estimates ranging from thirty-five to fifty-three billion dollars by 2030, however, many of the projects in existence today will be cancelled prior to shipping, according to the same sources. The adoption is widespread and superficial, with the majority of pilots stalling and only a small number of pilots with profit impact, who found a solution to the operating problem, not the model problem.

The barriers are common to all the sources. Security controls are coming after the capabilities they need to defend against, and security defenders predict they will be outpaced in a year. There is no clear answer as to who is accountable for what within a chain of delegating agents, and that is the single issue that is more of a constraint than any technical constraint on the deployments' ability to reach scale. Agents are sucking the life out of the content economy by traffic that sucks without paying back, and a payment model has yet to be shown to be effective in this regard. What is still to be done is only too obvious. The delegation, revocation and audit system that makes borrowed authority legible and reversible needs to be developed and must be affordable so that small organizations and public agencies can afford to have it just like large organizations. There is a need to move beyond the current mosaic of single one off surveys to independent, long-run measurement of the labour, market, content and energy effects. The legal issue of who is responsible for the actions of an agent that harms has to be addressed, not after agents start acting at



scale in high stakes environments. The Agent Internet will be more a result of the ability of the agents to solve these five problems security, accountability, measurement, market openness, and content payment than of how good they become. The technical is going the quickest, the safety and economic the slowest, the opposite of what a well-thought out rollout would do. That's the work that still remains to be done to get that order right.

REFERENCES

- [1] McKinsey & Company (2025). The state of AI: How organizations are rewiring to capture value. [online] McKinsey & Company. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>.
- [2] Pan, T. (2026). Your agent read the page. Nobody saw the ad. [online] tianpan_co. <https://tianpan.co/blog/2026/07/02/your-agent-read-the-page-nobody-saw-the-ad>.
- [3] Zero Labs, R. (2026). The state of the agent: Understanding adoption, risk, and mitigation. [online] Rubrik.com. <https://zerolabs.rubrik.com/reports/state-agent-understanding-adoption-risk-and-mitigation>.
- [4] A survey of agent interoperability protocols: Model Context Protocol (MCP), Agent Communication Protocol (ACP), Agent-to-Agent Protocol (A2A), and Agent Network Protocol (ANP). (n.d.). <https://arxiv.org/html/2505.02279v1>
- [5] AI Agents market to grow 43.3% annually. (2026, January 5). BCC Research LLC. <https://www.bccresearch.com/pressroom/ait/ai-agents-market-to-grow-433-annually?>
- [6] AI Agents Statistics 2026: Adoption, ROI, and Enterprise Deployment data | Tensorway. (n.d.). <https://www.tensorway.com/post/ai-agents-adoption>
- [7] Beyond context sharing: a Unified Agent Communication Protocol (ACP) for secure, federated, and autonomous Agent-to-Agent (A2A) orchestration. (n.d.). <https://arxiv.org/html/2602.15055v1>
- [8] Bijani, N., & Bijani, N. (2026a, June 8). The state of AI agent adoption in US Enterprise 2026: Findings from 50+ CTOs. Codiste Pvt Ltd. <https://www.codiste.com/ai-agent-adoption-us-enterprise>
- [9] Content Independence Day, one year on- building the business model for the agentic Internet. (2026, July 15). Cloudflare Blog. <https://blog.cloudflare.com/agentic-internet-bot-report/>
- [10] Ehtesham, A., Kent State University, Singh, A., Cleveland State University, Gupta, G. K., Youngstown State University, Kumar, S., & Northeastern University. (2025). A SURVEY OF AGENT INTEROPERABILITY PROTOCOLS: MODEL CONTEXT PROTOCOL (MCP), AGENT COMMUNICATION PROTOCOL (ACP), AGENT-TO-AGENT PROTOCOL (A2A), AND AGENT NETWORK PROTOCOL (ANP) [Journal-article]. <https://www.rivista.ai/wp-content/uploads/2025/05/2505.02279v1.pdf>
- [11] George, D. (2025). How AI Agents Are Reshaping the Internet from Human-Centered to Machine-Mediated Commerce. Zenodo (CERN European Organization for Nuclear Research). <https://doi.org/10.5281/zenodo.17064239>
- [12] Ehtesham, A., Singh, A., Gupta, G. K., & Kumar, S. (2025). A survey of agent interoperability protocols: Model Context Protocol (MCP), Agent Communication Protocol (ACP), Agent-to-Agent Protocol (A2A), and Agent Network Protocol (ANP). arXiv (Cornell University). <https://doi.org/10.48366/r1596648>
- [13] George, D. (2026a). The Agent Economy: How autonomous AI systems are restructuring knowledge work, capital allocation, and the architecture of enterprise value. Zenodo (CERN European Organization for Nuclear Research). <https://doi.org/10.5281/zenodo.18555676>
- [14] MarketsandMarkets, & MarketsandMarkets. (n.d.). AI Agents Market Report 2025-2030, by Application, Geo, Tech. MarketsandMarkets. <https://www.marketsandmarkets.com/Market-Reports/ai-agents-market-15761548.html?hl=en-GB>
- [15] George, D., Dr.T.Baskar, Srikanth, P. B., & Pandey, D. (2026). From AirGaP to Zero Trust-A Systematic Review of Operational Technology Cybersecurity Technologies for Critical Infrastructure in the era of ICS-Specific Malware. Zenodo (CERN European Organization for Nuclear Research). <https://doi.org/10.5281/zenodo.22688581>
- [16] MarketsandMarkets Research Pvt. Ltd. (2026, August 20). AI Agents Market Surges to \$52.62 billion at a CAGR 46.3% by 2030 | Report by MarketsandMarkets&sup>TM. Yahoo Finance. <https://finance.yahoo.com/technology/ai/articles/ai-agents-market-surges-52-150000335.html>



- [17] George, D., Dr.T.Baskar, & Srikanth, P. B. (2026). Biometric Authentication Contemporary Trends and Systemic Failures-A Review. Zenodo (CERN European Organization for Nuclear Research). <https://doi.org/10.5281/zenodo.22055699>
- [18] Playwire. (2025, December 6). How AI crawling Affects your ad Revenue: A Data-Driven Analysis. <https://www.playwire.com/blog/how-ai-crawling-affects-your-ad-revenue-a-data-driven-analysis>
- [19] George, D., & Dr.T.Baskar. (2025). AI Agents Are Reshaping the Internet from Human-Centered to Machine-Mediated Commerce. Zenodo (CERN European Organization for Nuclear Research). <https://doi.org/10.5281/zenodo.17212467>
- [20] PR Newswire. (2025, September 5). AI Agent Market Forecast to Reach \$42.7 Billion by 2030: North America is Leading with 40% Market Share - MarkNtel Advisors. Yahoo Finance. <https://finance.yahoo.com/news/ai-agent-market-forecast-reach-135600682.html>
- [21] George, D. (2026c). Killer acquisitions and the blind spot of merger control: economic logic, contested evidence, and implications for innovation policy. Zenodo (CERN European Organization for Nuclear Research). <https://doi.org/10.5281/zenodo.22250776>
- [22] Prefactor. (2026, September 16). 79% of companies run AI agents: 65 adoption stats (2026). <https://prefactor.tech/learn/ai-agent-adoption-statistics>
- [23] George, D. (2026b). Big data, behavioral biometrics, and the era of continuous authentication - a review. Zenodo (CERN European Organization for Nuclear Research). <https://doi.org/10.5281/zenodo.22070510>
- [24] Security Considerations for Artificial Intelligence Agents (Perplexity response to NIST/CAISI Request for Information 2025-0035). (n.d.). <https://arxiv.org/html/2603.12230>
- [25] George, D. (2026e). Modern data center network architecture Spine-and-Leaf topologies, virtualization technologies, and socioeconomic implications for digital infrastructure. Zenodo (CERN European Organization for Nuclear Research). <https://doi.org/10.5281/zenodo.23077164>
- [26] Security of Internet of agents: attacks and countermeasures. (n.d.). <https://arxiv.org/html/2505.08807>
- [27] George, D. (2026d). Cost to Company and Revenue to Company-Reading Labour Value from Both Sides in an Age of Rented Machines and Metered Software Agents. Zenodo (CERN European Organization for Nuclear Research). <https://doi.org/10.5281/zenodo.22924283>
- [28] Team, D. A. (2026, March 25). AI Agent scaling Gap March 2026: Pilot to production. DigitalApplied. <https://www.digitalapplied.com/blog/ai-agent-scaling-gap-march-2026-pilot-to-production>
- [29] Toward a safe internet of agents. (n.d.). <https://arxiv.org/html/2512.00520v1>
- [30] Unlocking exponential value with AI agent orchestration. (2026, September 9). Deloitte Insights. <https://www.deloitte.com/us/en/insights/industry/technology/technology-media-and-telecom-predictions/2026/ai-agent-orchestration.html>
- [31] Wibowo, J. A., & Polyzos, G. C. (2025). Toward a safe internet of agents. arXiv (Cornell University). <https://doi.org/10.48550/arxiv.2512.00520>
- [32] Yaqoob, M. (2026, February 2). How GPTBot, ClaudeBot, and PerplexityBot actually crawl your site, and how to optimize for AI search. DEV Community. <https://dev.to/msmyaqoob25/how-gptbot-claudebot-and-perplexitybot-actually-crawl-your-site-and-how-to-optimize-for-ai-search-4862>
- [33] You, T. E. I. C. C. a. M. F. T. R. O. I. Y. W. N. L. U. T. R. O. T. Y. P. F. F. T. C. Y. P. S. W. C. B. A., & Playwire. (2026, June 4). Google AI mode is eating publisher traffic. Here's what to do. <https://www.playwire.com/blog/google-ai-mode-is-eating-publisher-traffic-heres-what-to-do>