

Publication status: This preprint has not been published elsewhere.

The Data Economy in the Age of Agents: A Perspective on the Future of Agents as Customers

Hugo Loza-Guzmán

<https://doi.org/10.1590/SciELOPreprints.18349>

Submitted on: 2026-10-08

Posted on: 2026-10-19 (version 1)
(YYYY-MM-DD)

The moderation of this preprint received the endorsement of:

- José Loaiza Torres (ORCID: <https://orcid.org/0000-0002-8164-9420>)

The Data Economy in the Age of Agents: A Perspective on the Future of Agents as Customers

Hugo Loza-Guzmán

Universidad Católica Boliviana "San Pablo" - DTT Lab - Tarija, Bolivia

ORCID: <https://orcid.org/0009-0002-1439-2298>

6 October 2026

Abstract

The expansion of artificial intelligence is changing how knowledge circulates and is monetized on the web. Systems retrieve information from external sources and deliver answers that can replace a user's visit, while those sources continue to bear the costs of producing and maintaining the content. From 2028 or 2029 onward, queries from autonomous agents with spending budgets would become the main monetization channel for a significant part of the knowledge web. The agent would act as a customer capable of selecting and acquiring resources to complete a task on behalf of its user. The literature on information goods and micropayments provides a basis for examining this transition. Access data from sihita.org, a public repository of water information in Tarija, Bolivia, place it in a regional context. HTTP 402, x402, Cloudflare's Monetization Gateway and Pay per Use, together with licensing marketplaces, offer mechanisms for charging for access to and use of content. For institutions that produce specialized knowledge, the period 2026–2028 represents an opportunity to prepare their offerings and explore commercial relationships with these customers, combining paid services with public access to knowledge.

Keywords: x402 protocol, Autonomous agents with spending budgets, AI web crawling, Web micropayments, API monetization, Data economy.

Resumen

La expansión de la inteligencia artificial está modificando la circulación y la monetización del conocimiento en la web. Los sistemas recuperan información de fuentes externas y entregan respuestas que pueden sustituir la visita del usuario, mientras esas fuentes continúan asumiendo los costos de producir y mantener el contenido. A partir de 2028 o 2029, las consultas de agentes autónomos con presupuesto se convertirían en el principal canal de monetización de una parte relevante de la web de conocimiento. El agente actuaría como un cliente capaz de seleccionar y adquirir recursos para completar una tarea por encargo de su usuario. La literatura sobre bienes de información y micropagos aporta una base para examinar esta transición. Los datos de acceso a sihita.org, un repositorio público de información hídrica de Tarija, Bolivia, permiten situarla en un contexto regional. HTTP 402, x402, el Monetization Gateway y Pay per Use de Cloudflare, junto con los mercados de licencias, ofrecen mecanismos para organizar el cobro por acceso y utilización del contenido. Para las instituciones que producen conocimiento especializado, el periodo 2026–2028 representa una oportunidad para preparar su oferta y ensayar relaciones comerciales con estos clientes, combinando servicios remunerados y acceso público al conocimiento.

Palabras clave: Protocolo x402, Agentes autónomos con presupuesto, Rastreo web de IA, Micropagos web, Monetización de APIs, Economía de los datos.

1. The Breakdown of the Internet's Implicit Contract

For years, publishing information on the web meant attracting people. A search engine found a page, a user arrived at it, and that visit could lead to advertising revenue, a subscription, a professional consultation, or a sale. Even sites without commercial activity gained visibility, recognition, and reach. The economic value of publishing rested on that flow of readers.

Artificial intelligence is reorganizing that journey. A person can request a report, a comparison, or a technical answer and receive the result without opening the sources used. Retrieval-augmented generation combines a language model with external documents to produce answers (Lewis et al., 2020). In regular web search, Chapekis and Lieb (2025) observed that Google users who encountered AI summaries clicked less frequently on the results shown below them. Content remains necessary, but the visit shifts to a system that searches, selects, and combines information. Wiederhold (2026) places task delegation at the center of this transformation of the web experience. For those who maintain the information source, the economic question changes: how can they share in the value their information creates when the user receives the answer elsewhere?

From my perspective, this change has a concrete expression: servers hosting local knowledge receive queries from systems operating internationally. What appears to be a traffic problem also signals a new relationship between those who produce information and those who need it to complete a task.

I argue that, from 2028 or 2029 onward, the main monetization channel for a significant part of the knowledge web will shift from human visits to queries from autonomous agents with spending budgets. Transactions will be possible per request, per use, or through licenses, using a combination of HTTP mechanisms, programmable payments, and content marketplaces. The decisive change will be the agent's ability to acquire information on behalf of a person or an organization and use it to achieve an objective.

This perspective is especially relevant to technical, scientific, territorial, socioeconomic, demographic, historical, and cultural sources. An up-to-date piece of information about a watershed, an environmental data series, or a specialized document can contribute value to a task that is difficult to obtain from generic information. That is where the opportunity to treat the agent as a customer begins.

I see a concrete opportunity for those who produce and maintain specialized information: offering services that an agent can query and pay for during a task. Work on the agentic economy and agent-mediated markets describes changes in purchasing, intermediation, and the organization of commerce (Rothschild et al., 2026; Bichler, 2026). My contribution to that discussion is to consider how institutions with technical and territorial knowledge can participate, taking Tarija as a starting point.

2. A Starting Point in Tarija

sihita.org is a public repository of water information from Tarija, hosted in the data center of the Universidad Católica Boliviana. It holds about 70 GB of technical PDF documents and aims to make all information generated about the department's water resources freely available to the community without restrictions. Its content originates in local needs, but the machines that query it bring that content into much broader information networks.

For this work, I calculated two reference periods: February–July 2022 and February–July 2026.

Estimated page views increase from approximately 5,680 to 313,100 (Figure 1).

The first period (February–July 2022) precedes the public launch of ChatGPT on 30 November 2022 (OpenAI, 2022).

Between the two periods, the estimated volume of data served increases from approximately 76 GB to 1,758 GB. Figure 1 shows the change in scale and reach that prompted this reflection.




Figure 1. (a, b) sihita.org. Two periods calculated by the author from log data: February–July 2022 and February–July 2026. Estimated page views and GB served.

Knowledge also circulates across a wider geographical area. In July 2022, Jetpack (WordPress) recorded 448 visitors, and around 90% of readers were in Bolivia. In the Cloudflare telemetry collected on 30 September 2026, 88% of requests came from outside the country, while Bolivia accounted for 11.6%. A repository created to meet regional needs now serves international demand.

Automated access accounts for a substantial part of the repository's activity. Tables 1 and 2 show the composition of traffic and the volume transferred.

Table 1. Traffic composition at sihita.org. Source: origin logs, 6 September–5 October 2026.

Access group	Requests	% of requests	GB transferred	% of volume
Probable and possible human users	64,116	17.5%	15.65	12.3%
Automated traffic	199,919	54.5%	111.17	87.4%
Vulnerability probing	102,552	28.0%	0.42	0.3%
Origin total	366,587	100.0%	127.24	100.0%

Table 2. Automated traffic at sihita.org. Source: origin logs, 6 September–5 October 2026.

Type of automation	Requests	GB transferred
Undeclared automation	104,291	60.66
Declared traditional bots	77,479	38.22
Declared AI	18,149	12.28
Total automated traffic	199,919	111.17

I distinguish two uses of AI in these requests: crawling for training and indexing, and information searches linked to user queries (Figure 2). The operators' documentation helps interpret that distinction.


OpenAI describes GPTBot as a bot that collects content that may be used for training; ChatGPT-User, in turn, accesses pages in response to user actions (OpenAI, n.d.).

Training incorporates knowledge into a model; a query returns to the source when a task requires it. **This second relationship opens up recurring demand.**

PDFs account for 96.8% of the volume transferred to AI bot signatures at the origin. I identified 27 signatures from 19 operators; Figure 2 presents their requests, distinguishing crawling for training and indexing from searches and user queries. The same collection serves both to gather content and to answer queries.

AI requests by bot

6 September–6 October 2026 · estimated total: 23,644



C: crawling for training and indexing · 14,819 requests

S: search and user enquiries · 8,825 requests

Figure 2. Requests recorded for the 27 AI bot signatures. Source: origin logs from 6 to 17 September and Cloudflare from 18 September to 6 October 2026, with a partial final day and adaptive sampling. The count includes blocked requests.

In the origin logs, GPTBot makes 98.5% of its requests between 00:00 and 06:00; ChatGPT-User concentrates its data transfers on PDF documents. These uses suggest two commercial relationships: agreements to collect document collections, and access to documents when a task needs them. For those who maintain the source, recognizing these uses makes it possible to offer terms that match the value each customer seeks.

Requests to missing resources (HTTP 404)

Origin logs · 6 September–5 October 2026




Figure 3. HTTP 404 responses as a share of each selected bot's total requests. Source: origin logs, 6 September–5 October 2026. Labels show requests to resources not found / total requests.

The trend also appears at a larger scale. Cloudflare reports that crawling associated with user actions grew more than 21-fold between January and early December 2025. In June 2026, training accounted for 52% of

requests from crawlers identified by purpose, compared with 22% in spring 2025 (Cloudflare, 2025b, 2026a). The web supports both the development of models and their everyday use.

Across sites in TollBit's network, there was one AI bot visit for every 200 human visits at the beginning of 2025. By the end of the year, there was one for every 31. During the final quarter, bots searching for information to answer users' questions made, on average, more than 4.5 times as many queries as bots collecting content to train models (TollBit, 2026).

How should a university repository respond when its content begins to feed commercial AI services? SIHITA was created to share water information from Tarija, and that purpose remains. Keeping it available requires storage, connectivity, and technical work. The agents identified include services from companies offering commercial AI products. The university maintains the source, while other organizations can earn revenue by using its information. How should the cost of keeping it available be shared?

As the administrator of this data center, I ask how we can adapt to that demand while preserving the repository's public mission. What terms should accompany intensive use by commercial services? Could they contribute to maintaining the infrastructure or pay for specialized queries and services? Does this moment mark the breakdown of the Internet's implicit contract?

These questions open up a possibility: those who currently consume resources through automated access could also become customers who contribute to sustaining them.

3. Agents as the New Customers

The idea of a program making purchases for us was already being studied in electronic commerce. Maes et al. (1999) described systems able to compare offers, negotiate, and make purchases on behalf of their users. In 2026, the combination of language models, information retrieval, and programmable payments brings that relationship to the knowledge an agent needs during a task.

An agent with a spending budget receives an objective, has resources to achieve it, and decides how to use them within the terms set by the person or organization responsible for it. It can compare sources, select an offer, and pay for information or services. It acts on behalf of a company, an institution, or a person who defines what it can buy, how much it can spend, and what result is expected.

Consider an agent tasked with preparing a water assessment. It could locate public documents, consult an updated data series, purchase a structured data extraction, and commission a specialized analysis. Each purchase would make sense in relation to the result it must deliver. For the source, the unit of value could be a document, a query, an update, a citation to the source, or a dataset with clear provenance and terms of use.

The decision to pay comes from comparing the price of a resource with its contribution to the task. In that assessment, a recent, documented data series can save the effort of searching several sites, extracting tables from a PDF, or checking conflicting versions. The agent could reserve part of its budget for that information and use the rest for processing and analysis. Willingness to pay is thus related to the time saved, the quality obtained, and the importance of delivering an up-to-date result. Someone who buys a query acquires a contribution to the task they commissioned.

The same logic applies when someone at a winery asks an agent to assess the risk of frost. To support a decision, the system could combine a general forecast with recent local observations and historical records for the plots. An institution maintaining those records could offer an updated query or an alert under commercial terms. The value would lie in bringing local knowledge into a concrete decision. The agent would buy that contribution within the budget authorized by the winery.

This relationship already appears in concrete implementations. Ceramic published a demonstration of a research agent that pays 0.001 USDC for each web search through x402 and Cloudflare's Monetization Gateway. The program includes a spending limit per run and checks the budget before authorizing each payment (CeramicTeam, n.d.). Information thus becomes part of the cost of completing a task: the user sets how much can be spent, and the agent uses that budget to obtain the resources it needs.

This gives the client a broader role. In computing, a client requests a resource from a server. In economics, a customer acquires something because they expect to gain value. The agent brings both roles together when its technical request includes a purchasing decision. The server becomes part of a commercial relationship in which information has a price, usage rights, and an accountable counterparty.

Machines can be customers for knowledge, real-time data, analyses, and signals. Each category has a different logic. A document collection can be sold under a license; a series that changes daily can generate repeated queries and be offered through a subscription; an analysis service can charge per result. What they share is that access is linked to a task and a budget.

For institutions that produce knowledge, this relationship changes how they present their offerings. Editorial quality still matters, along with the ability to show what a resource contains, when it was updated, who produced it, and what it can be used for. Wang and Strong (1996) located data quality in its suitability for consumers' needs, including relevance, timeliness, representation, and accessibility. Trust in the source also influences purchasing: it helps an agent choose that source and the person or organization responsible for the agent agree to pay for it.

4. The Infrastructure Is Already Taking Shape

These commercial relationships are emerging through concrete developments taking place as I write in 2026. Cloudflare introduced Pay per Crawl on 1 July 2025, making it possible to charge crawlers for access to content. A year later, it announced its Monetization Gateway, and on 30 September 2026 it opened a closed beta for charging for resource consumption through HTTP 402 (Cloudflare, 2025a, 2026b, 2026d). Access can include a payment offer within the exchange between client and server.

HTTP 402, named Payment Required and reserved for future use in the HTTP standard, provides a starting point for signaling a payment condition (Fielding et al., 2022). x402 develops that possibility through a protocol for programmable payments over HTTP. The x402 Foundation began operating under the Linux Foundation on 14 July 2026 (Linux Foundation, 2026). These mechanisms bring purchasing into the resource request: the client encounters commercial terms, makes the payment, and gains access.

Stablecoins and blockchain architecture allow software to manage and transfer funds directly and automatically. The x402 initiative proposes using them so that software agents can pay for the web services they consume (x402 Foundation, n.d.). For example, when acquiring information or knowledge, the allocated budget can fund a digital wallet that automatically pays for each query as it occurs. The provider then charges directly for its resource, and the agent's supervisor keeps a clear record of spending. With those funds available, the agent can pay at the moment it needs to query a piece of information.

Pay per Use works with a different economic unit. In its beta of 30 September 2026, the buyer proposes a price for a specific use, the publisher accepts the offer, and the buyer reports usage so that Cloudflare can invoice and settle the payments (Cloudflare, 2026c). Charging for access and charging for use serve different demands. A crawler can explore many sources before choosing which ones to incorporate into an answer.

Microsoft introduced the Publisher Content Marketplace on 3 February 2026 as a marketplace for licensing content for AI products. Publishers define licensing and usage terms, while AI developers find content to ground answers (Frank, 2026). A transaction can be organized around a collection, a provider, or a usage agreement, as well as an individual request.

Spending authorization is also seeing its own developments. The Agent Payments Protocol (AP2), announced by Google in September 2025, uses signed mandates to record instructions, price limits, and the terms of a delegated purchase (Parikh & Surapaneni, 2025). That authorization layer complements payment mechanisms and links a purchase to the instructions of the person or organization responsible for it.

AI systems need sources, and information producers seek a share of the value they generate. Charging mechanisms make it possible to establish that relationship.

The information market is beginning to have infrastructure for customers able to operate programmatically. Over the coming years, its reach, prices, and ease of integration will influence which sources can participate and how much value they retain.

A knowledge transaction can begin when an agent discovers a source and finds an offer specifying content, version, price, and terms of use. It evaluates that offer within its budget, makes the payment, and receives the resource. The transaction record can preserve provenance and authorization to use it in the result. This process extends what today often ends with a download: it turns access into a relationship that the source can serve, document, and renew.

5. Changes in the Economics of Information

Shapiro and Varian (1999) explained that information goods combine high production costs with relatively low reproduction costs. Their price depends on the value they offer different buyers and on how access is organized. This logic takes a new form when an agent can search for a specialized source and buy the right resource for a task.

In a data center, this economy has a physical dimension. Every available resource requires storage, connectivity, and operational continuity; a query adds data transfer and processing. The repository also depends on those who select documents, correct records, and take responsibility for their quality. Within an institution, these costs are covered by its budget. When other organizations use that information to provide services, there is an opportunity to agree on how they will contribute to maintaining the source from which they gain value.

The value lies in the information a source produces and keeps up to date. Duplicating a digital file costs almost nothing; the costly work is generating reliable data, certifying its origin, and keeping it current. A territorial record, a technical data series, or a specialized analysis retains value because it offers a level of detail that generic sources cannot provide. The economic opportunity lies in that specialization and in the work needed to maintain it.

This model is especially powerful for dynamic information:

- During training: an AI model can absorb a static database as it is trained under an initial license.
- During execution: when solving tasks in real time, that same model will need the rule currently in force, the most recent data, or the latest revision.

Each updated piece of information renews the value of the original source. Thus, while training encourages agreements for fixed bundles of information, real-time queries open the door to recurring subscription services, allowing creators to finance and sustain their ongoing work.

Availability conditions are also changing. Longpre et al. (2024) document the withdrawal of permissions and increasing restrictions on sources used by AI. The open access that supported the initial expansion of models is being renegotiated. For those who maintain information, declaring uses and establishing agreements becomes part of managing the resource.

The history of micropayments raises a decisive question: who takes on the effort involved in buying? Szabo (1999) identified a barrier in the mental cost of evaluating small payments and already considered the role of intelligent agents. Odlyzko (2003) explained the appeal of flat rates and the economic complexity of charging for each small unit of use. Delegating those purchases to an agent can spare us the decision of paying for each query separately. A person defines a budget and preferences; the system manages the queries and reports total spending.

The resulting market can combine micropayments, subscriptions, and bundles. Bakos and Brynjolfsson (1999) show why bundling information goods can be profitable when their marginal cost is low. An agent or company making frequent queries could acquire a license or a bundle, while another pays for an occasional need. For producers, experimenting with these options will help them find a relationship between price, repeat purchases, and the value delivered.

The available budget changes how information is evaluated: each resource must justify its cost as it competes for the funds allocated to completing a task.

For those who produce that information, the way they assess the value of their work also changes. A source with few visits can be decisive in completing a specialized task. It then makes sense to examine which queries it helps answer, how many buyers return, and whether revenue can sustain it. The knowledge an institution holds can generate revenue when it is useful for other people's work. Serving a digital file is the simplest part; the challenge is financing the ongoing work required to produce information and maintain its quality.

6. The Window of Opportunity between 2026 and 2028

For these reasons, I think 2028 or 2029 will mark the beginning of the expansion of agents as customers of the knowledge web. My view is that the combination of automated queries, delegated spending, and payment infrastructure will by then have a sufficient commercial presence to change the model of many data sources. Around 2028, I expect growing interest in selling data and services to machines: the rush that often accompanies an opportunity once it becomes widely known.

As queries linked to users increase, tools are being developed to charge for them. In September 2026, Cloudflare introduced the beta of its Monetization Gateway with use cases for web search, access to financial signals, and document generation through programmable payments. Pay per Use adds another possibility: agreeing on compensation for the use of content and settling it through shared infrastructure (Cloudflare, 2026c, 2026d).

During 2026–2028, information providers can learn to present their resources, set prices, and integrate them into agents' workflows. I expect that experience to facilitate commercial expansion afterward: more sources will be able to offer queries, updates, and services that agents purchase when they need them. The opportunity lies in preparing the available knowledge now to participate in that market.

I expect the transition to be most visible initially in technical information, data requiring constant updates, and specialized services, where it is easier to link a purchase to a concrete need and assess what is received in return. Sources already offering structured queries or licenses can bring these customers into existing commercial relationships; repositories can begin with services added to their collections. The combination of these paths explains my expectation of expansion: demand finds prepared offerings, and experience from the first transactions makes further purchases easier.

Carlota Perez (2002) provides a useful framework for thinking about that sequence. Technological revolutions involve the installation of infrastructure and the transformation of economic and institutional practices. Participating early makes it possible to learn while the market's rules are being defined. The 2026–2028 window is therefore a period of preparation: identifying which knowledge has value and building the capacity to offer it.

Universities, observatories, and specialized repositories hold a particular asset. They produce knowledge rooted in specific places, preserve archives, and maintain relationships with the territories they describe. The environmental history of a watershed, data on a productive activity, or the technical documentation of a city can be useful for tasks requiring local precision. Turning that body of knowledge into an offering accessible to agents requires organization, continuity, and criteria for its use.

This opportunity also broadens the scope of technological development: alongside building applications, preparing the data they need to operate becomes increasingly important. A team can turn a regional archive into a system of structured queries, keep technical records up to date, or issue alerts when relevant changes occur. While programming makes the service possible, the knowledge accumulated by the institution gives its offering a value of its own. For regional producers, organizing that body of knowledge can generate revenue from the knowledge and capabilities they already have.

In Tarija, these offerings could take the form of observations for agriculture, viticulture, and tourism, updated information for organizing tourism experiences, or technical services that process local documents and data series. An agent planning a trip could check the availability and terms of an experience; another could purchase a data extraction for its report.

These examples show that regional knowledge can become part of tasks initiated far from where it was produced. The next customer of a local service could arrive through an agent working for someone in another country.

Preparation can begin with small decisions. Recognizing which resources receive queries helps us understand demand. Identifying operators and purposes makes it possible to define different terms for training, search, and the purchase of services. Metadata, versions, and structured access points make content easier to use. The FAIR principles guide that preparation toward resources that are findable, accessible, interoperable, and reusable by people and machines (Wilkinson et al., 2016). More recently, Lù et al. (2025) proposed web interfaces designed for agents' navigation and tasks. From there, providers can test per-query prices, bundles, and usage agreements, keeping records that show what customers buy and why they return.

To commercialize an information resource, it is essential to define clearly what it delivers and under what terms. For a data-series query, for example, it is useful to specify geographical coverage, available variables, the update date, and the response format. Commercial terms add to this: price, usage limits, and conditions for reuse.

For an institution with a public mission, much remains to be debated. Its offerings can keep documents openly accessible while charging for services that add value: updates, structured queries, processing, or specialized assistance. That combination can help sustain infrastructure and editorial work. The UNESCO Recommendation on Open Science links the circulation of knowledge with sustainable and equitable services and infrastructure (UNESCO, 2021). Defining the rights over each resource and its authorized uses is part of that preparation.

An illustrative case of this coexistence is Bolivia's commercial registry through SEPREC (Servicio Plurinacional de Registro de Comercio). The directory allows users to search for economic units, while the formal registration procedure includes checking for duplicate names and reserving a name, together with payment of the registration fee (SEPREC, n.d.). The information search and the procedure with administrative effects provide different services. That difference helps us consider how an institution can keep information accessible while charging for specific services associated with it.

Open-Meteo provides a practical example of that combination. It maintains data under CC BY 4.0 and open-source code, alongside a free API for noncommercial use and subscriptions with commercial access, dedicated capacity, and support (Open-Meteo, n.d.). The service makes it possible to charge for operational terms and usage capacity while keeping information open. That design is useful for thinking about institutional offerings that finance their maintenance and broaden the circulation of knowledge.

Bolivia has begun updating its framework for virtual assets and financial technology services. The Central Bank's Board Resolution 082/2024 enabled electronic payment channels and instruments for buying and selling virtual assets; Supreme Decree 5384 of 2025 regulated financial technology companies (BCB, 2024; ASFI, 2025). This context creates an agenda for integrating digital contracting, payment settlement, and institutional operations. For local providers, charging international customers will be an important part of service design.

The advantage of preparing now lies in gaining experience before the opportunity attracts greater competition. The Spanish saying “A toro pasado, todos somos toreros”—in hindsight, everyone is a bullfighter—captures how easy it is to recognize a transformation once it has already happened. It makes sense to act while that market is taking shape: organizing the available knowledge, testing ways to sell it, and learning what value the new customers will seek.

7. The Tensions That Will Shape the Market

The development of this market will depend on how several tensions are resolved. The first concerns substitutable and distinctive content. Bakos (1997) showed that reducing the costs of searching for prices and features changes competition in electronic markets. If many sources offer the same information, agents will be able to compare them and put pressure on prices. Sources with continuous updates, specialization, or reliable provenance will be better placed to sustain a differentiated offering. That distinction guides where effort should be invested. The second concerns intermediation. Bailey and Bakos (1997) identified aggregation, matching suppliers with buyers, and providing trust as functions of electronic markets.

Bichler (2026) analyzes agent-mediated markets and shows the economic importance of search rules, the ranking of offers, and platform design. Those decisions influence which providers are visible and which alternatives a buyer considers. For a specialized source, participation means making its quality recognizable and understanding how it reaches the customer. The way purchasing is organized also influences which sources have an opportunity to sell.

Intermediation will also influence the basis for charging. A platform can license a collection, retain it in its systems, and serve many queries under that agreement. A source that updates data continuously can establish a subscription, while another offers occasional purchases of distinctive resources. Licensing marketplaces and usage-based charging accommodate these relationships (Frank, 2026; Cloudflare, 2026c). For the producer, it matters that the agreement recognizes the use of the content and the ongoing work that keeps it valuable.

Advertising can coexist with these payments. An advertising-supported platform can acquire content licenses and offer answers to the user without charging for each query. I also think it is very likely that advertisements will be included in answers or in the content returned, especially when the service is free.

The third brings together trust, security, and spending control. A buyer needs to recognize the source and know the terms of the resource. A seller needs to identify the party responsible for the transaction and verify payment. Delegated budgets require limits, records, and accountability.

The fourth concerns the relationship between openness and sustainability. UNESCO (2021) places collective benefit and equitable access among the principles of open science. Agents can pay for specialized access, while society retains an interest in public knowledge and its circulation. Institutions will need to decide which services to commercialize, which content to keep open, and how to reinvest revenue. That decision will be especially important for universities and repositories funded to fulfill a service mission.

Finally, the pace of adoption will depend on profitability. Search, payment, and processing costs must be proportionate to the resource's usefulness, a tension found in the literature on electronic markets and micropayments (Bakos, 1997; Szabo, 1999; Odlyzko, 2003). In some sectors, broad agreements and subscriptions will predominate; in others, individual purchases. By 2029, the change can be assessed through repeat purchasing, the inclusion of spending budgets in agents, and the share these transactions represent in sources' revenue. Commercial growth will indicate the depth of the transition.

8. Concluding Remarks

Agents as customers offer a way to think about the web's next economic cycle. Knowledge can be found, evaluated, and purchased during the execution of a task. For information producers, that relationship opens the possibility of sharing directly in the value their work contributes to systems able to act and spend on behalf of their users.

I think that, from 2028 or 2029 onward, this relationship will take a central role in a significant part of the knowledge web. Per-request payments, usage agreements, and licensing marketplaces will be complementary ways of organizing it.

The opportunity lies in preparing and valuing the knowledge being produced. Universities, observatories, and repositories can participate as information providers about their own territories and fields of expertise.

Between 2026 and 2028, organizing that offering and learning to transact with the new customers can determine the place they will occupy in the market as it grows.

Author's Note

This article builds on the keynote “La economía de los datos en la era de los agentes: cuando los visitantes se convierten en clientes,” presented at Tarija Tech Week on 1 October 2026 in Tarija, Bolivia.

The figures were prepared by the author.

The interpretation and proposal express the author's position.

The use of AI tools (ChatGPT/Codex, OpenAI) supported literature and source review and translation into English.

License: Creative Commons Attribution 4.0 International (CC BY 4.0).

Author Contribution

Hugo Loza-Guzmán developed the proposal, collected and analyzed the data, prepared the figures, and wrote and revised the article.

Conflict of Interest

The author declares no conflicts of interest.

Data Availability Statement

Aggregate repository-access figures and estimates are presented in the article's tables and figures, with their sources identified. External sources are cited in the text and reference list.

References

- ASFI. (2025). Decreto Supremo N.º 5384, de 7 de mayo de 2025. Decretos supremos referidos a servicios financieros. <https://www.asfi.gob.bo/pb/decretos-supremos-referidos-servicios-financieros>
- Bailey, J. P., & Bakos, Y. (1997). An exploratory study of the emerging role of electronic intermediaries. *International Journal of Electronic Commerce*, 1(3), 7–20. <https://doi.org/10.1080/10864415.1997.11518287>
- Bakos, J. Y. (1997). Reducing buyer search costs: Implications for electronic marketplaces. *Management Science*, 43(12), 1676–1692. <https://doi.org/10.1287/mnsc.43.12.1676>
- Bakos, Y., & Brynjolfsson, E. (1999). Bundling information goods: Pricing, profits, and efficiency. *Management Science*, 45(12), 1613–1630. <https://doi.org/10.1287/mnsc.45.12.1613>
- Banco Central de Bolivia [BCB]. (2024). El BCB, en coordinación con ASFI y UIF, actualiza normativa en torno a los activos virtuales. Resolución de Directorio 082/2024, de 25 de junio. <https://www.bcb.gob.bo/?q=content/el-bcb-en-coordinaci%C3%B3n-con-asfi-y-uif-actualiza-normativa-en-torno-los-activos-virtuales>
- Bichler, M. (2026). Agentic markets. *Electronic Markets*, 36, Article 55. <https://doi.org/10.1007/s12525-026-00906-y>
- CeramicTeam. (n.d.). ceramic-x402-search-agent [Demonstration code and documentation]. GitHub. Retrieved October 6, 2026, from <https://github.com/CeramicTeam/ceramic-x402-search-agent>
- Chapekis, A., & Lieb, A. (2025, July 22). Google users are less likely to click on links when an AI summary appears in the results. Pew Research Center. <https://www.pewresearch.org/short-reads/2025/07/22/google-users-are-less-likely-to-click-on-links-when-an-ai-summary-appears-in-the-results/>

- Cloudflare. (2025a, July 1). Introducing pay per crawl: Enabling content owners to charge AI crawlers for access. <https://blog.cloudflare.com/introducing-pay-per-crawl/>
- Cloudflare. (2025b). The 2025 Cloudflare Radar Year in Review: The rise of AI, post-quantum, and record-breaking DDoS attacks. <https://blog.cloudflare.com/radar-2025-year-in-review/>
- Cloudflare. (2026a, July 1). Content Independence Day, one year on: Building the business model for the agentic Internet. <https://blog.cloudflare.com/agentic-internet-bot-report/>
- Cloudflare. (2026b, July 1). Announcing the Monetization Gateway: Charge for any resource behind Cloudflare via x402. <https://blog.cloudflare.com/monetization-gateway/>
- Cloudflare. (2026c, September 30). Pay Per Use: When AI uses your work, you should get paid. <https://blog.cloudflare.com/pay-per-use/>
- Cloudflare. (2026d, September 30). Monetization Gateway beta: Charge AI agents for consumption with HTTP 402. <https://blog.cloudflare.com/monetization-gateway-beta/>
- Fielding, R., Nottingham, M., & Reschke, J. (2022). HTTP Semantics. RFC 9110. <https://www.rfc-editor.org/rfc/rfc9110>
- Frank, T. (2026, February 3). Building toward a sustainable content economy for the agentic web. Microsoft Advertising. <https://about.ads.microsoft.com/en/blog/post/february-2026/building-toward-a-sustainable-content-economy-for-the-agentic-web>
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-T., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems*, 33. https://papers.nips.cc/paper_files/paper/2020/hash/6b493230205f780e1bc26945df7481e5-Abstract.html
- Linux Foundation. (2026, July 14). Linux Foundation announces operational launch of x402 Foundation to standardize internet-native payments for AI agents and applications. <https://www.linuxfoundation.org/press/linux-foundation-announces-operational-launch-of-x402-foundation-to-standardize-internet-native-payments-for-ai-agents-and-applications>
- Longpre, S., Mahari, R., Lee, A., Lund, C., Oderinwale, H., Brannon, W., Saxena, N., Obeng-Marnu, N., South, T., Hunter, C., Klyman, K., Klamm, C., Schoelkopf, H., Singh, N., Cherep, M., Anis, A., Dinh, A., Chitongo, C., Yin, D., . . . Pentland, S. (2024). Consent in crisis: The rapid decline of the AI data commons. *arXiv:2407.14933*. <https://doi.org/10.48550/arXiv.2407.14933>
- Lù, X. H., Kamath, G., Mosbach, M., & Reddy, S. (2025). Build the web for agents, not agents for the web. *arXiv:2506.10953*. <https://doi.org/10.48550/arXiv.2506.10953>
- Maes, P., Guttman, R. H., & Moukas, A. G. (1999). Agents that buy and sell. *Communications of the ACM*, 42(3), 81–91. <https://doi.org/10.1145/295685.295716>
- Odlyzko, A. (2003). The case against micropayments. In R. N. Wright (Ed.), *Financial Cryptography, FC 2003, LNCS 2742* (pp. 77–83). Springer. <https://www-users.cse.umn.edu/~odlyzko/doc/case.against.micropayments.pdf>
- Open-Meteo. (n.d.). Pricing. Retrieved October 6, 2026, from <https://open-meteo.com/en/pricing>
- OpenAI. (2022, November 30). Introducing ChatGPT. <https://openai.com/index/chatgpt/>
- OpenAI. (n.d.). Overview of OpenAI crawlers. <https://developers.openai.com/api/docs/bots>
- Parikh, S., & Surapaneni, R. (2025, September 16). Powering AI commerce with the new Agent Payments Protocol (AP2). Google Cloud Blog. <https://cloud.google.com/blog/products/ai-machine-learning/announcing-agents-to-payments-ap2-protocol>
- Perez, C. (2002). *Technological Revolutions and Financial Capital: The Dynamics of Bubbles and Golden Ages*. Edward Elgar. <https://carlotaperez.org/books/>
- Rothschild, D. M., Mobius, M. M., Hofman, J. M., Dillon, E., Goldstein, D. G., Immorlica, N., Jaffe, S., Lucier, B., Slivkins, A., & Vogel, M. (2026). The agentic economy. *Communications of the ACM*, 69(2), 39–42. <https://doi.org/10.1145/3747175>

Servicio Plurinacional de Registro de Comercio [SEPREC]. (n.d.). Directorio de Unidades Económicas e inscripción de sociedad comercial. Retrieved October 6, 2026, from <https://miempresa.seprec.gob.bo/> ; <https://www.seprec.gob.bo/index.php/tramite2/>

Shapiro, C., & Varian, H. R. (1999). *Information Rules: A Strategic Guide to the Network Economy*. Harvard Business School Press. <https://faculty.haas.berkeley.edu/shapiro/>

Szabo, N. (1999, May). Micropayments and mental transaction costs. Author's text archived at the Satoshi Nakamoto Institute. <https://nakamotoinstitute.org/library/micropayments-and-mental-transaction-costs/>

TollBit. (2026). State of the Bots: Q3 & Q4 2025 — The Leaky Pipes. <https://tollbit.com/state-of-the-bots/q3-q4-2025/>

UNESCO. (2021). Recomendación de la UNESCO sobre la Ciencia Abierta. <https://www.unesco.org/en/legal-affairs/recommendation-open-science>

Wang, R. Y., & Strong, D. M. (1996). Beyond accuracy: What data quality means to data consumers. *Journal of Management Information Systems*, 12(4), 5–33. <https://doi.org/10.1080/07421222.1996.11518099>

Wiederhold, B. K. (2026). From human web to agentic web: When machines become the Internet's primary users. *Cyberpsychology, Behavior, and Social Networking*, 29(5), 278–280. <https://doi.org/10.1177/21522715261432850>

Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J.-W., da Silva Santos, L. B., Bourne, P. E., Bouwman, J., Brookes, A. J., Clark, T., Crosas, M., Dillo, I., Dumon, O., Edmunds, S., Evelo, C. T., Finkers, R., . . . Mons, B. (2016). The FAIR guiding principles for scientific data management and stewardship. *Scientific Data*, 3, 160018. <https://doi.org/10.1038/sdata.2016.18>

x402 Foundation. (n.d.). x402. Retrieved October 6, 2026, from <https://x402.org/>

This preprint was submitted under the following conditions:

- The authors declare that the necessary Terms of Free and Informed Consent of participants or patients in the research were obtained and are described in the manuscript, when applicable.
- The authors declare that the preparation of the manuscript followed the ethical norms of scientific communication.
- The authors declare that they are aware that they are solely responsible for the content of the preprint and that the deposit in SciELO Preprints does not mean any commitment on the part of SciELO, except its preservation and dissemination.
- The authors declare that the data, applications, and other content underlying the manuscript are referenced.
- The deposited manuscript is in PDF format.
- The authors declare that the research that originated the manuscript followed good ethical practices and that the necessary approvals from research ethics committees, when applicable, are described in the manuscript.
- The authors declare that once a manuscript is posted on the SciELO Preprints server, it can only be taken down on request to the SciELO Preprints server Editorial Secretariat, who will post a retraction notice in its place.
- The authors agree that the approved manuscript will be made available under a [Creative Commons CC-BY](#) license.
- The submitting author declares that the contributions of all authors and conflict of interest statement are included explicitly and in specific sections of the manuscript.
- The authors declare that the manuscript was not deposited and/or previously made available on another preprint server or published by a journal.
- If the manuscript is being reviewed or being prepared for publishing but not yet published by a journal, the authors declare that they have received authorization from the journal to make this deposit.
- The submitting author declares that all authors of the manuscript agree with the submission to SciELO Preprints.