Vision Pierre Genevier

SYNTHESIS NOTE — PROPOSAL No. 2 Domain name pricing application — architecture, feasibility and data-collection methodology

Document prepared with Claude (Anthropic) — version revised September 12, 2026, incorporating the country-breakdown analysis (Microsoft Copilot) and the proposed cap by company category for the United States (Google Gemini), transmitted and commented on by Pierre Genevier on September 12, 2026, as a supplement to the August 11, 2026 version. Living working document — to be revisited, corrected and expanded at the next session.

1. Purpose and place within the platform

This note addresses Proposal No. 2 of Pierre Genevier's platform: the establishment of a new domain-name pricing system, intended to finance the international organization (IO) dedicated to the governance of the Internet and artificial intelligence, and more broadly the UN missions described in the Vision, as well as the improvement of the Internet's information system.

In line with the common template validated on July 15, 2026, this note focuses on what is specific to Proposal No. 2 — the pricing calculation engine, its application architecture and its data-collection methodology — and refers to the Synthesis Note — Proposal No. 1 for everything relating to the IO's common institutional architecture, its infrastructure scenarios and its overall financing as an institution. This note therefore develops, on the financing side, only the pricing mechanism itself (section 4), while the macro-financial reading through the IO's three infrastructure scenarios remains covered in Note-Proposal No. 1 (section 4.4).

Methodological note: this version fully carries over the content of the August 11, 2026 version and adds to it, without removing anything: (i) in section 2.4, an indicative illustration of the breakdown of the $30 Bn/year target by country (method proportional to GDP, based on a Microsoft Copilot analysis of September 12, 2026) and of the breakdown by company category for the US share, the country that would bear the highest burden (progressive-cap method proposed by Google Gemini on September 12, 2026); (ii) in section 4.1, a paragraph on the benefit provided to large digital groups by the creation of the IO; (iii) a corresponding point of caution in section 6; and (iv) two new annexes (7.5 and 7.6) presenting the detailed figures for these two breakdowns. It remains a living working document, to be expanded in future sessions.

2. Description of the proposal

2.1 General principle: a parametric, continuous calculation engine, and a more detailed and effective Internet information system

The proposal establishes a levy mechanism on domain names, based on a plurality of parameters specific to each liable entity [nature of the entity (company, government body, association, university, ..., individual) and purpose of the site; revenue, profits, headcount, etc. for companies; budgets, headcount, etc. for government bodies, associations, etc.; names, addresses, etc. for individuals; and, for all domain-name holders, connected IoT objects, payment methods, bank details, etc. (for those paying a levy) — section 2.3], rather than on a simple binary revenue threshold. The calculation engine is designed as a configurable rules engine, applying continuous weighting coefficients to each parameter to produce an individualized levy per entity, with no hard-

coded values: the coefficients themselves therefore remain adjustable by political decision of member states, without requiring application redevelopment at each revision of the schedule.

Beyond their role in calculating the levy, this information should also allow the IO to have a more detailed and effective Internet information system, useful in particular for improving Internet security and combating cybercrime (and, where applicable, certain human-rights violations). This objective is in no way intended to establish surveillance of Internet users; the technical arrangements guaranteeing that this data cannot be used for such surveillance remain to be documented, in connection with the phasing described in section 2.5.

2.2 Seventeen functional modules

The application is now broken down into three blocks: six initial modules (Phase 1, modules 1 to 6), covering domain-name holders; six modules centered on individual (human) Internet users via their connected devices (Phase 2, modules 7 to 12); and, since the August 11, 2026 version, five additional modules covering individual income generated without an owned site via third-party platforms (Phase 3, modules 13 to 17 — see section 2.6). Phases 2 and 3 share the same base register of individual users (Module 7), but rely on distinct logics of attachment, collection and verification, detailed below.

Phase 1 — modules 1 to 6

- **Module 1 — Entity register**: central database listing each liable entity (company, government body, association, university, individual site owner) with its identification attributes and associated domain names — see the detail of the types of information collected in section 2.1.
- **Module 2 — Collection and declaration of parameters**: interfaces allowing each entity to declare or have verified its parameters (total and Internet-specific revenue, total and Internet-specific profits, total and Internet-dedicated headcount, attached IoT devices).
- **Module 3 — Parametric calculation engine**: application of the weighting coefficients to produce the individualized levy (cf. section 2.1).
- **Module 4 — Verification and audit**: consistency checks on declared data, cross-referencing with third-party sources where available, anomaly detection. The most sensitive module, since it touches on trust in the system and on national data sovereignty.
- **Module 5 — Billing and tracking**: generation of levy notices, payment tracking, management of reminders and any disputes.
- **Module 6 — Dashboard and reporting**: overview for IO managers, eventually including an interactive simulator for the impact of coefficient changes.

Phase 2 — modules 7 to 12 (individual users, IoT component)

- **Module 7 — Individual user register**: database listing individual (natural person) Internet users, distinct from the site-owner register (Module 1), based on the federated model described in section 2.5. This register is shared by Phases 2 and 3.
- **Module 8 — IoT attachment / interoperable identifier**: mechanism for attaching IoT devices to an existing national identity via an interoperable identifier, without a centralized

global nominal register (cf. section 2.5).

- **Module 9 — Collection and declaration (IoT usage)**: interfaces for declaring or verifying Internet-resource usage linked to IoT devices attached to each individual user.
- **Module 10 — Calculation engine (possible individual levy)**: application, where applicable, of an individual levy linked to Internet-resource usage via attached IoT devices, following the same configurable-engine principle as in section 2.1.
- **Module 11 — Verification and audit**: consistency checks on declared data, with a heightened focus on privacy protection, given the nature of the data involved (natural persons).
- **Module 12 — Dashboard and reporting**: overview for IO managers on the individual-user (IoT) component, distinct from the Module 6 dashboard.

Phase 3 — modules 13 to 17 (individual users, platform-income component — new)

These modules are the counterpart, for income generated via third-party platforms without an owned site, of modules 8 to 12 for the IoT component — see the detail of the principle and justification in section 2.6.

- **Module 13 — Attachment to platform accounts / interoperable identifier**: mechanism linking an entry in the individual-user register (Module 7) to the accounts held by that person on monetized platforms (YouTube, Meta, TikTok, Amazon, Etsy, Fiverr, etc.), on the same interoperable-identifier principle as Module 8, without a centralized global nominal register.
- **Module 14 — Collection and declaration (platform income)**: interfaces for receiving declarations transmitted by the platforms themselves (rather than by the individual user), modeled on the European DAC7 directive and the OECD's Model Rules for digital platforms (cf. sections 2.6 and 3.3).
- **Module 15 — Calculation engine (levy on platform income)**: application of weighting coefficients to the levy owed on this income, reusing the Module 3 calculation-engine architecture.
- **Module 16 — Verification and audit**: consistency checks between platform declarations and, where applicable, the national tax declarations of the users concerned.
- **Module 17 — Dashboard and reporting**: overview for IO managers on the platform-income component, distinct from the Module 6 and Module 12 dashboards.

2.3 Parameters taken into account

For Phase 1, the calculation engine relies on four families of parameters declared or verified per entity: identification [legal identity of the site owner (company, government body, individual, etc.) and legal representative for non-individuals, purpose or type of site (commercial, personal, administrative, etc.), domain names held], financial parameters (revenue and profits, total and Internet-specific, budget where applicable), headcount (total and Internet-dedicated), and Internet-usage volume and attached connected devices (IoT). The collection methodology specific to each of these parameters, and the expected degree of reliability, are detailed in section 3.3.

For Phase 2, the calculation engine will necessarily be entirely different, since it will rely (at least initially) solely on the use of Internet resources by the IoT devices associated with the person concerned; the parameters will therefore be identification (legal identity or name of the person, address), attached connected devices (IoT), and the volume of Internet usage linked to those devices.

For Phase 3, the parameters are of a different nature: this is no longer a matter of Internet-resource consumption but of income generated through the Internet without owning a site. The parameters used are therefore the person's identification (name, tax identifier if available), the account identifier on each platform concerned (Module 13), and the amount of income paid by the platform over the reference period (Module 14), as declared by the platform itself rather than by the individual user — see the justification for this choice in section 2.6.

2.4 Revenue target and pricing mechanics to be formalized

The figure of $30 billion per year, mentioned in Pierre Genevier's curriculum vitae, constitutes a pricing-policy parameter set by member states — not a market figure to be verified — intended to fairly distribute the financing of the IO among domain-name holders and Internet users, based, among other things, on their use of resources and the income generated through the Internet. It remains consistent with the current order of magnitude of the global domain-name market ($5 to 6 billion per year), assuming a substantial but not unreasonable overhaul of the existing schedule. What remains to be documented is therefore not this figure itself, but the pricing mechanics for reaching it: progressivity criteria, revenue threshold for liable companies, accounting for actual resource use by individual Internet users, consistency with the parametric calculation engine described in section 2.1 — work to be carried out once the project is approved. The Phase 3 component (section 2.6) constitutes an additional, separate potential revenue source, not included in this $30 Bn/year target at this stage, which will require separate costing once the principle is politically validated.

2.4 ter — Indicative illustration of the breakdown by country and by company category in the United States (new, September 12, 2026)

This subsection addresses a question raised by Pierre Genevier on September 12, 2026: even before the definitive pricing mechanics (progressivity, thresholds) are formalized (section 4.3), it is useful to have an order of magnitude of what the $30 Bn/year target would represent once broken down by country, and then, for the country that would bear the highest burden — the United States — by company category. The two exercises that follow are illustrations, not proposed schedules: they aim to verify that the $30 Bn/year target remains plausible once broken down, and to anticipate the practical obstacles of a breakdown by country and by category of actor.

Country-breakdown method (proportional to GDP)

A first, purely illustrative approach consists of distributing the $30 Bn/year among the 191 states in proportion to each one's nominal Gross Domestic Product (GDP), on the following principle: Amount allocated to a country = $30 Bn × (country's GDP ÷ sum of GDPs used). This calculation was requested by Pierre Genevier from Microsoft Copilot, based on a table of GDP by country (IMF, 2026) consulted on Wikipedia. The detailed amounts for the Top 20 and Top 50 countries by GDP appear in Annex E (section 7.5).

Reliability of the GDP figures used: the 2026 nominal GDP values used by Copilot ($32,384 Bn for the United States, $20,852 Bn for China, $5,453 Bn for Germany, etc.) are consistent with the latest IMF estimates (World Economic Outlook), which place 2026 US GDP at around $31.8 Bn [sic] and China's at around $20.7 Bn [sic] according to the updates available — a discrepancy of roughly 1 to 2% per country, likely due to the version of the Wikipedia table consulted at the time of the request rather than to a methodological error. The order of magnitude and the relative ranking of countries are not affected by this discrepancy. For the US share, the proportional calculation thus comes out at approximately $7.69 Bn/year — a figure to be treated as an indicative order of magnitude to be refined (marginally) with official IMF data at the time the final schedule is drafted, rather than as a definitive value.

Limits of an intra-country breakdown strictly proportional to revenue

A second question, also raised by Pierre Genevier, is how to distribute a given country's share — the US share of $7.69 Bn/year, as an example, being the country with the highest burden — among that country's liable entities. A first simulation, also carried out at Pierre Genevier's request (Microsoft Copilot then Google Gemini), consisted of distributing this amount strictly in proportion to companies' revenue: on the Fortune 500 alone, this would lead Walmart (top revenue in the ranking, approximately $681 Bn) to bear approximately $277 M/year; broadening the base to all US companies (approximately $50,000 Bn in combined revenue, employer companies and independents combined), this amount drops to approximately $105 M/year for Walmart, with the remainder of the US share being diluted across SMEs and independents.

This simulation, while useful for gauging orders of magnitude, mainly illustrates the limits of a raw linear coefficient applied solely to revenue: it disproportionately penalizes large, high-volume but low-margin companies (retail, energy) relative to high-margin, less labor-intensive digital players, and does not reflect the logic of the parametric calculation engine described in section 2.1 (Module 3), which precisely relies on a plurality of weighted parameters — including the digital intensity of the activity — rather than on revenue alone.

Proposal for a progressive, capped schedule by company category

To correct this limitation, an alternative method was proposed by Google Gemini, at Pierre Genevier's request, based on a critical review of this proposal: rather than a single rate on revenue, the pricing mechanics would combine four levers, consistent with the architecture already adopted by the proposal (section 2.1 and Module 3):

- progressivity by bracket and a maximum cap per entity (no actor, even one the size of a very large global group, would owe a confiscatory amount);
- a digital-intensity coefficient for the activity (share of revenue generated online, volume of DNS queries or bandwidth, number of domain names held), already implicit in the principle of the parametric calculation engine (section 2.1);
- mobilization of the Phase 3 component (levy on platform income, section 2.6), which allows a significant share of the target amount to be raised without increasing the burden solely on domain-name holders;
- a genuine worldwide distribution across all of a country's domain-name holders (companies of all sizes, government bodies, universities, associations) rather than solely on the largest companies, which mechanically dilutes the burden per actor.

Applied, for illustration, to the US share ($7.69 Bn/year), this method distributes the burden among company categories according to a progressive, capped schedule, summarized below (full detail in Annex F, section 7.6):

Category	**Capped levy / indicative flat fee**	**Revenue generated (illustrative)**
Very large digital groups (Tech / Cloud / e-commerce)	Cap ~ $25 M/year	~ $0.88 Bn
Large companies (Fortune 500 excluding Tech)	Cap ~ $3 M/year	~ $1.38 Bn
Mid-size enterprises	~ $10,000 to $80,000/year	~ $1.20 Bn
SMEs	~ $1,200 to $5,000/year	~ $0.80 Bn
Small businesses	~ $150 to $600/year	~ $0.45 Bn
Micro-enterprises and independents	~ $15 to $45/year	~ $0.36 Bn
Phase 3 — platform income (US component)	~ 0.5% to 1% withheld at source, exemption below $2,000/year	~ $2.20 Bn
Government bodies, universities, organizations	Flat fee ~ $500 to $10,000/year	~ $0.42 Bn

Under this assumption, a digital giant would see its contribution capped at a marginal amount relative to its results, while a large traditional low-margin company (Walmart, in the example) would go from approximately $277 M (raw calculation proportional to revenue alone on the Fortune 500) to a cap of around $3 M/year. A typical SME would fall in the range of an annual professional software subscription, and an independent in the range of the cost of a domain name.

Scope of this illustration: these figures are orders of magnitude produced by an external simulation exercise (Google Gemini), intended to demonstrate the arithmetic feasibility of the $30 Bn/year target under a progressive, capped mechanism, and not a finalized schedule. The definitive pricing mechanics (exact rates, thresholds, caps, definition of "digital intensity") remain, as indicated above, to be documented once the project is approved (section 4.3), consistent with the parametric calculation engine (Module 3).

2.4 bis Entry into force of the levy and voluntary advance payment

Two distinct mechanisms, not to be confused, allow the IO to anticipate the collection of resources before the mechanism described in section 3 is fully operational.

Immediate entry into force — not a retroactive levy

The principle adopted is not that of a levy retroactive in the legal sense — which would raise a difficulty of principle in most legal systems (non-retroactivity of taxation, nullum tributum sine lege) — but that of a levy whose obligation arises on the date of the General Assembly's decision to create the IO, and whose actual payment is simply deferred until the information system described in section 3 makes it possible to precisely calculate the amount owed by each liable entity. The claim therefore exists from the decision to create the IO; only its settlement is delayed, for practical

rather than legal reasons. This point deserves to be formulated with precision in the final text of the proposal, to avoid any legal ambiguity.

Voluntary advance payment

Independently of the above, a voluntary advance-payment mechanism would be opened, from the project's launch, to domain-name holders wishing to support the IO's financing before the general application of the levy — particularly major Internet players ("Internet giants") with the financial capacity to do so. In return, a reduction would be applied to their future levy, of a probably modest amount and for a limited period (not for life), the exact parameters (rate, duration, cap) remaining to be defined. This mechanism has a twofold objective: to give major Internet players the opportunity to support the project from its launch, and to provide the IO with start-up resources before the general levy mechanism is fully operational. This mechanism must remain clearly distinct, both conceptually and administratively, from the mandatory levy itself, so as not to suggest that the latter is negotiable.

2.5 Phasing: domain-name holders / individual (human) Internet users

The envisaged extension of the mechanism — attaching IoT devices and robots to their individual users, beyond mere pricing for site owners — raises issues of a different nature (privacy, state sovereignty over civil status, risk of mass surveillance), capable of jeopardizing the entire proposal if the two components are not clearly separated. The mechanism is therefore now phased in three stages:

- **Phase 1 — Entities owning Internet sites (or domain names) (immediate)**: covers exclusively government bodies, companies, associations and natural persons owning a site, based on registers already existing in most jurisdictions. This is the phase that enables the implementation of pricing and the generation of the IO's first resources; it accounts for most of the costing in this note (sections 3 and 4).
- **Phase 2 — Individual (human) Internet users, IoT component (conditional, medium-to-long term)**: motivated primarily by the attachment of IoT devices/robots to their users, not by a goal of a generalized individual levy. Conditional on the prior establishment of a data-protection framework negotiated with member states, guaranteeing respect for the privacy of all registered persons — not costed at this stage.
- **Phase 3 — Individual (human) Internet users, platform-income component (see section 2.6)**: unlike Phase 2, Phase 3 can start immediately, including in parallel with Phase 1, without depending on the prior resolution of the same technical and privacy-protection obstacles. Indeed, it does not rely on building a global nominal register of connected devices attached to natural persons — the source of the mass-surveillance risk identified for Phase 2 — but on the IO's reuse of data that platforms already collect for their own compliance obligations (know-your-customer, anti-money-laundering, tax reporting). The principle of verification by the platforms themselves rather than by unverifiable individual declaration (cf. sections 2.6 and 3.3) offers a guarantee of reliability and a legal basis already proven internationally (the European Union's DAC7 directive, the OECD's Model Rules for digital platforms), which makes this phase costable as of this version of the note (section 3.1).

For Phase 2, a federated rather than centralized model is recommended by default: each state would retain control of its own identification system, with the IO holding only a system of interoperable

identifiers allowing an IoT device to be attached to an existing national identity, without a centralized global nominal register — a solution closer to interoperability frameworks already under discussion (ITU, World Bank ID4D) than to the most controversial precedents (centralized registers such as Aadhaar or ID2020). A technical alternative — attachment via certificate or allocation token rather than nominal register — also deserves further exploration before any decision.

Further centralization is not adopted as a Phase 2 objective: it could only be considered, where applicable, if structural — not merely technical — guarantees of privacy protection and non-diversion toward surveillance purposes were established in the future; the federated model remains, as it stands, the recommended default option. This question connects to broader issues on the platform — notably Proposal No. 4 (pension management), which also requires a minimum level of registration of natural persons — addressed, where applicable, in the corresponding notes rather than in this note, whose IoT/levy angle is deliberately kept narrower.

2.6 Levy on individual income generated via third-party platforms, without an owned site

This section addresses a question raised by Pierre Genevier on August 11, 2026: beyond domain-name holders (Phase 1) and individual users via their IoT devices (Phase 2), many natural persons today generate significant income through the Internet without owning a site in their own name — content creators, influencers, marketplace sellers, freelancers on platforms — relying exclusively on third-party platforms (YouTube, Meta/Instagram/Facebook, TikTok, Amazon, Etsy, eBay, Fiverr, Upwork, etc.).

Principle adopted

It is consistent with the overall logic of Proposal No. 2 (a levy based both on the holding of Internet resources and on the income derived from them) to extend the contribution principle to this category of users. This component is, however, treated as a mechanism distinct from Phase 2 rather than as an extension of it: Phase 2 measures consumption of technical resources (connected devices), while this component measures income, which brings its calculation logic closer to that of Phase 1 (section 2.1) — the difficulty lying not in the calculation principle but in the collection point, since a platform user by definition holds neither a domain name nor a registrar to which collection could be attached.

Collection mechanism: relying on the platforms rather than on the individual user

Rather than requiring individual declarations from several hundred million creators and sellers scattered around the world — an unrealistic and hard-to-verify hypothesis — the proposed mechanism relies on an already-established institutional precedent: the European DAC7 directive (Directive 2021/514, transposed since January 1, 2023) requires digital-platform operators (marketplaces, service platforms, including those located outside the European Union where they have European sellers) to collect, verify and transmit to tax authorities the identity and income of their monetized users, above a de minimis threshold (€2,000 or 30 transactions per year under DAC7). The OECD has also formalized Model Rules adopting this principle internationally, outside the European Union. The mechanism therefore does not need to be invented: it involves transposing, for the benefit of the IO, a principle of platform self-declaration that is already proven and accepted by the sector's major players.

Two implementation options are conceivable and remain to be decided: (a) on the DAC7 model, the

platform declares the income paid to the IO (or to the national tax authority, which would transmit it to the IO), which then issues the levy notice to the individual user; (b) on a withholding-at-source model, the platform directly deducts the levy amount before paying the user, which is administratively simpler and more reliable for the IO, but politically harder to negotiate with the platforms — resistance comparable to that encountered by national digital-services taxes (DSTs) is to be expected (cf. section 6).

Orders of magnitude (sourced data, replacing the unsourced preliminary analysis of July 31, 2026)

The previous version of this analysis relied on an analysis produced by a general-purpose AI assistant (Microsoft Copilot), several of whose figures were not verifiably sourced, in particular the estimate of 80 to 120 million people living from these activities. A specific, sourced search was conducted for this note; the following orders of magnitude are better established, but remain themselves marked by wide dispersion across research firms, which must be preserved in the wording:

- Size of the global "creator economy" market (total income earned by individual creators from digital platforms): specialized research firms converge on an order of magnitude of $250 to $320 Bn for 2025-2026 (Grand View Research: $252 Bn in 2025 / $310 Bn in 2026; Research and Markets: $256 Bn in 2025 / $323 Bn in 2026), with projections of $800 Bn-$1,350 Bn by 2030-2033 depending on the source. The gap between firms (roughly double depending on the scope used) illustrates the absence of a standardized definition of this sector, to be treated as a range rather than a single figure in any external communication.
- Number of active creators worldwide: on the order of 200 to 300 million people, across all platforms (convergent order of magnitude across several sources).
- Very unequal income distribution: according to available data, only about 4% of active creators earn more than $100,000 per year in income, and about half earn less than $15,000 per year. This finding matters for designing a de minimis threshold (see below): the mass of very-low-income creators does not justify the cost of collection or of disputes, and a threshold comparable to DAC7's appears transposable at the international level.
- Influencer-marketing market, a subset of the creator economy corresponding to brand spending: approximately $32.5 Bn in 2025, with a projection of approximately $40.5 Bn in 2026 (Influencer Marketing Hub / Mordor Intelligence).
- Direct payouts to creators by a single platform, as an illustration of the order of magnitude available for a levy based on income paid: YouTube has announced having paid out more than $100 Bn to creators, artists and media over the 2021-2025 period, i.e., an annual flow on the order of $20 to $30 Bn for this platform alone (source: public announcements by Alphabet/YouTube, 2025).

These figures confirm that the order of magnitude is significant and warrants further work on this component, but they should not, at this stage, be used to set a quantified revenue target for Phase 3: the pricing mechanics (rate, de minimis threshold, possible cap) remain to be documented, on the same principle as for Phase 1 (section 4.3), once the principle is politically validated.

Point of caution

This component exposes the IO to a risk of overlap with digital-services taxes (DSTs) already in force or under discussion in several jurisdictions, as well as to a politically sensitive negotiation with platforms possessing considerable legal and lobbying power — a point to be handled with the same caution applied to the sensitive wording of Proposal No. 5 (cf. session memo).

3. Technical feasibility

3.1 Development estimate by module

Indicative estimate of complexity and development workload by module, on the assumption of a team experienced in developing large-scale transactional applications — to be refined once the technical architecture is validated.

Module	**Complexity**	**Estimate (person-months)**	**Risk points**
1. Entity register	Moderate	8–14	Reconciliation with heterogeneous national identifiers
2. Collection / declaration	Moderate to high	12–20	Multiplicity of languages, jurisdictions, source data formats
3. Calculation engine	Moderate	6–10	Design of the configurable weighting model
4. Verification / audit	High	14–24	Absence of reliable third-party sources for certain parameters (cf. section 3.3)
5. Billing / tracking	Low to moderate	6–10	Interfacing with multi-currency payment systems
6. Dashboard	Low	4–8	—

Total estimate for Phase 1: on the order of 50 to 85 person-months for a full version, i.e., with a team of 6 to 10 people (developers, architect, data engineer, security expert, project manager), a development schedule, running modules in parallel, of 10 to 16 months, excluding large-scale testing and phased rollout.

Phase 2 (modules 7 to 12) remains uncosted at this stage, conditional on the prior negotiation of the data-protection framework (section 2.5). Phase 3 (modules 13 to 17), however, which can start immediately per the clarification given in section 2.5, is the subject of a first indicative estimate:

Module	**Complexity**	**Estimate (person-months)**	**Risk points**
13. Platform-identifier attachment	Moderate	8–14	Technical negotiation with each major

Module	Complexity	Estimate (person-months)	Risk points
			platform (proprietary API formats)
14. Collection / declaration (platform income)	Moderate to high	10–16	Multiplicity of platforms and their reporting formats
15. Calculation engine	Moderate	6–10	Partial reuse of the Module 3 architecture
16. Verification / audit	High	12–18	Cross-referencing with national tax declarations, heterogeneous across states
17. Dashboard	Low	4–8	—

Total estimate for Phase 3: on the order of 40 to 66 person-months — an estimate to be refined once the collection principle (platform declaration vs. withholding at source, cf. section 2.6) is decided with Pierre Genevier.

Recommended development phasing

- **Phase 1 — MVP**: modules 1 and 3 on a reduced scope (companies from a limited number of pilot jurisdictions, parameters already publicly available). MVP = "Minimum Viable Product": the smallest possible version of the system that remains functional and demonstrates feasibility, before investing in the full version. Objective: demonstrate the feasibility of the calculation engine and obtain real figures for political negotiation.
- **Phase 2 — Extension**: full addition of module 2 and extension to all jurisdictions and entity categories.
- **Phase 3 — Consolidation**: modules 4, 5 and 6 — verification, billing, steering. Worldwide scale-up phase.

Overall timeline estimate: 3 to 5 years between the MVP and a fully operational worldwide system, consistent with the usual deployment timelines for new international administrative infrastructure.

Cross-cutting risk factor: collection and verification (modules 2 and 4) are the limiting factor of the project, more so than the technology itself — the application can be developed quickly, but its usefulness depends entirely on the reliability of declared data, particularly for parameters with no third-party source. This point steers the negotiation toward the need for an inter-state verification cooperation mechanism, more than toward a purely technical challenge (cf. sections 3.3 and 6).

3.2 Infrastructure needs

Volume estimate

Sizing assumption: a base on the order of 400 to 500 million entities eventually (companies + government bodies + associations), based on the following orders of magnitude — to be refined:

approximately 300 to 350 million active companies worldwide; approximately 190 states and a much higher number of sub-national government bodies; several tens of millions of associations, with very uneven documentary coverage across countries; approximately 350 to 370 million active domain names worldwide, across all extensions.

Storage and computing

For a database of this size, with several dozen fields per entity (identity, domains held, financial parameters, levy history, attached IoT devices), the estimated storage volume is on the order of several tens of terabytes for the system as a whole — modest by current cloud-infrastructure standards. The parametric calculation engine itself represents a low-to-moderate workload (weighted arithmetic calculations, without heavy AI model training or inference), with a peak load expected during annual declaration periods, requiring batch processing rather than permanent real-time computation.

Key point, already established in Note-Proposal No. 1 (section 3.4): this application does not require frontier AI computing capacity. It is a classic transactional database combined with a parametric calculation engine, falling at most under Scenario B as defined in Note-Proposal No. 1, and whose scale remains stable regardless of the overall infrastructure scenario chosen for the IO. The Phase 3 component (section 2.6), relying on receiving aggregated declarations transmitted by a limited number of major platforms rather than on direct collection from hundreds of millions of individuals, adds only a marginal infrastructure burden compared to Phase 1.

Security and resilience

Since the database contains sensitive financial data (profits, revenue) for hundreds of millions of entities, it will require a level of security comparable to that of an international tax administration — encryption, audit trail, granular access control. The question of data sovereignty remains to be settled: single centralized hosting or regional distribution with synchronization, the latter option probably being preferable for political acceptability among member states, echoing the question raised for the natural-persons database in section 2.5.

3.3 Data-collection methodology

The collection mechanism relies on an online form, filled in by each site owner (Phase 1) or, where applicable, by each individual user concerned (Phase 2). The data entered is verified, as far as possible, almost immediately using the sources listed in the tables below. Where the information is already available via these sources, a pre-filled form — fully or partially — is provided to the entity concerned, which is responsible for completing and validating it. For Phase 3, the logic is reversed: the data is transmitted directly by the platform (cf. section 2.6) and does not require a pre-filled form to be validated by the individual user, on the model of withholding at source or automatic declaration already used in tax matters (DAC7).

For each parameter identified in section 2.3, the table below specifies the potential source(s), the current level of availability, the envisaged collection mechanism and the expected level of reliability.

Entity identification

Parameter	Potential source(s)	Current availability	Mechanism / reliability
Legal identity of companies	National business registers (SIRENE, Companies House...)	High in developed jurisdictions, uneven elsewhere	Cooperation agreements/API — high reliability where available
Identity of government bodies	UN databases, official state lists	High at national level, incomplete at sub-national level	Declaration by member states — high reliability
Identity of associations	National association registers	Low to moderate, very uneven	Self-declaration + partial cross-referencing — moderate reliability
Domain names held	Registrars, WHOIS/RDAP databases	High (standardized data)	Agreements with ICANN and registrars — high reliability

Financial parameters

Parameter	Potential source(s)	Current availability	Mechanism / reliability
Total revenue	Tax authorities, business registers, stock-market authorities	High for large companies; moderate for SMEs	Agreements/mandatory declaration — high to moderate reliability
Total profits	Same sources	Same	Same
Internet-specific revenue / profits	No existing third-party source	Very low — not currently required	Self-declaration only, with a normative definition to be established — low reliability, critical point (cf. section 6)

Headcount

Parameter	Potential source(s)	Current availability	Mechanism / reliability
Total headcount	National social-security registers	High in most jurisdictions	Cooperation agreements — high reliability
Internet-dedicated headcount	No existing third-party source	Very low — not separately recorded	Self-declaration, normative definition to be established — low reliability, same limitation as above

Internet usage volume and IoT devices

Parameter	Potential source(s)	Current availability	Mechanism / reliability
Bandwidth / traffic, storage	CDNs, hosting providers, cloud providers, ISPs	Technically high but held by private actors, not centralized	Cooperation agreements — complex negotiation, sensitive data — moderate reliability
Attached IoT devices	Manufacturers, telecom operators (M2M SIMs), cloud IoT platforms	Technically high but fragmented among many actors	Self-declaration then agreements with major providers — low to moderate reliability, expected to improve

Platform income (Phase 3)

Parameter	Potential source(s)	Current availability	Mechanism / reliability
Income paid to the individual user	Direct declaration by the platform (YouTube, Meta, TikTok, Amazon, Etsy, Fiverr, etc.)	High — major platforms already hold this data for their own compliance obligations (KYC, DAC7, 1099)	Mandatory platform declaration, on the DAC7/OECD model — high reliability, subject to the agreement of the platforms concerned

Reminder: DNS root servers are not a relevant source for actual site traffic, since they only see the initial resolution of top-level domain names.

The detailed typology of IoT devices currently in use (eight broad categories, from the connected home to service robotics) appears in the annex (section 7.2); the "service robotics" category is the one that connects most directly to the argument developed in section 2.5 in favor of an identifier attachable to the user.

Two cross-cutting findings emerge from this methodology, to be foregrounded explicitly rather than minimized: "total" data are largely collectable through institutional cooperation with already-existing sources, at a reasonable marginal cost; "Internet-specific" data, on the other hand, currently have no reliable third-party source and will necessarily rely on self-declaration at launch — which is not a dealbreaker (many international tax systems initially operate this way, with ex-post control) but must be presented as an accepted limitation of the system at launch, with a trajectory of progressive improvement rather than as a solved problem. The Phase 3 component partially escapes this limitation since it relies, by construction, on declaration by a verifiable third party (the platform) rather than by the liable party itself.

4. Financial feasibility

4.1 Current market and revenue target

The current global domain-name market generates an order of magnitude of approximately $5 to 6 billion per year across all registries and registrars (a figure to be confirmed and precisely sourced in a later version). ICANN's current annual budget is on the order of $150 to 200 million (figure to be confirmed). The $30 billion/year target set for the new pricing (section 2.4) thus represents a substantial, but not unreasonable, overhaul of the existing schedule relative to these orders of magnitude.

Benefit for large digital groups (added September 12, 2026): this financing target should not be presented as a mere additional tax, but as the counterpart of substantial benefits that the creation of the IO would bring, first to the actors who would contribute the most — the large digital groups — and, more broadly, to all companies using the Internet: legal stability and harmonization of rules across jurisdictions, strengthened security and enhanced fight against cybercrime (section 2.1), international governance of artificial intelligence, and improvement of the Internet's information system itself. For companies whose results run into the tens of billions of dollars, a capped levy (section 2.4 ter) represents a marginal cost relative to securing a global digital environment from which these same companies derive a significant share of their revenue. This reciprocity argument — financing in exchange for the stability and security of the ecosystem — deserves to be developed in the version intended for political negotiation, in addition to the macro-financial sustainability argument already covered in section 4.4 and in Note-Proposal No. 1.

4.2 Application development cost

Based on an international team experienced in developing large-scale transactional applications, mobilized by or for an international organization (and therefore subject to higher multi-jurisdictional security and compliance requirements than a purely private team), an average fully loaded cost on the order of $18,000 to $28,000 per person-month is used as a working assumption — including loaded salary, social contributions, project management and overhead, but excluding infrastructure (section 3.2).

Applied to the estimates in section 3.1, this assumption gives the following orders of magnitude:

- Phase 1 (50 to 85 person-months): development cost estimated between $0.9 M and $2.4 M.
- Phase 3 (40 to 66 person-months, cf. section 3.1): development cost estimated between $0.7 M and $1.85 M.

These amounts are very modest relative to the annual revenue target ($30 Bn/year, section 2.4) and therefore do not constitute a limiting factor for the project; the limiting factor remains, as indicated in section 3.1, the reliability of collection rather than the cost of application development. This average-cost-per-person-month assumption remains to be validated with Pierre Genevier and harmonized with the same question raised for Proposal No. 3 (cf. Note-Proposal No. 3, section 6).

4.3 Pricing mechanics to be formalized

The schedule for moving from the current order of magnitude ($5 to 6 Bn/year) to the $30 Bn/year target remains to be documented: progressivity criteria, revenue threshold for liable companies,

accounting for actual resource use by individual Internet users, consistency with the parametric calculation engine (section 2.1). This work, distinct from application development, directly determines whether the revenue target is met and will be carried out once the project is approved, in accordance with section 2.4. The breakdown illustration by country and by company category presented in section 2.4 ter, and detailed in Annexes E and F, constitutes a first basis for this costing work, without replacing it.

4.4 Cross-reference — macro-financial reading

The reading of the IO's financing through the three infrastructure scenarios (share of the $30 Bn/year target devoted to CAPEX and OPEX for each scenario) falls under the overall financing of the IO as an institution and is covered in Note-Proposal No. 1 (section 4.4), to which this note refers.

5. Implementation timeline

- **Phase 1 — MVP** (modules 1 and 3, pilot scope): demonstration of the calculation engine's feasibility, obtaining first real figures for the political negotiation of the schedule (section 4.3).
- **Phase 2 — Extension** (module 2, generalization to jurisdictions and entity categories): scale-up of collection, in parallel with negotiating the pricing mechanism.
- **Phase 3 — Consolidation** (modules 4, 5, 6): verification, billing and steering at global scale — it is at the end of this phase that the system becomes fully operational.

Total estimated duration between the MVP and a fully operational worldwide system: 3 to 5 years (section 3.1). This timeline directly conditions that of Note-Proposal No. 1: the move from Scenario A to Scenario B of the IO assumes that the pricing mechanism has reached sufficient maturity to finance its CAPEX and OPEX (Note-Proposal No. 1, section 4.4).

The IoT-device component (modules 7 to 12, section 2.2) remains conditional on the prior negotiation of a data-protection framework and is not included in this timeline. The platform-income component (modules 13 to 17), on the other hand, can be launched as early as the MVP phase, in parallel with Phase 1, in accordance with the clarification given in section 2.5 — its own timeline remains to be established once the principle is politically validated and the collection mode (declaration vs. withholding at source, section 2.6) decided.

6. Risks and open questions

- Absence of a reliable third-party source for Internet-specific data (revenue, profits, headcount): will rely on self-declaration at launch — a central risk factor for the project, to be documented as an accepted limitation rather than minimized (section 3.3).
- Moderate-to-low reliability of data on associations, traffic/storage and IoT devices: largely dependent on cooperation from fragmented, non-centralized private actors.
- Collection of payment methods and bank details for liable entities (section 2.1): data of a sensitivity of a different order than identity or revenue, exposing the system to a specific security risk (fraud, theft of bank data at the scale of hundreds of millions of entities);

requires a dedicated protection level (PCI-DSS or equivalent), separate from the rest of the register.

- Pricing mechanics (progressivity, thresholds, accounting for actual usage) not yet documented: this work is now deferred to project approval (sections 2.4 and 4.3) rather than undertaken immediately.
- Average-cost-per-person-month assumption ($18,000 to $28,000, section 4.2) based on an indicative order of magnitude, not yet validated by Pierre Genevier.
- Sovereignty and hosting of sensitive financial data: choice between centralized hosting and regional distribution not settled (section 3.2).
- Phase 2 (individual users, IoT component): conditional on a data-protection framework not yet negotiated; choice between a federated model and a certificate/token mechanism not settled; any move toward further centralization would remain subject to structural, not merely technical, guarantees (section 2.5).
- Phase 3 (individual users, platform-income component): risk of overlap with digital-services taxes (DSTs) already in force or under discussion in several jurisdictions; politically sensitive negotiation with platforms possessing considerable lobbying power, comparable to the resistance already encountered by national DSTs; de minimis threshold and collection mode (declaration vs. withholding at source) not yet decided (section 2.6).
- Worldwide replication of DNS governance itself (number of anycast instances, redundancy): distinct from the volume figures of this application, already flagged as an open point in Note-Proposal No. 1.
- Breakdown illustrations by country and by company category (section 2.4 ter, Annexes E and F, new): these figures come from external simulation exercises (Microsoft Copilot for the GDP breakdown, Google Gemini for the company-category cap in the United States), not independently verified; the GDP data used shows a discrepancy of roughly 1 to 2% from the latest direct IMF estimates depending on the source. They must be presented, in any external communication, as illustrative orders of magnitude and not as a finalized schedule.

7. Annexes

7.1 Annex A — Table of modules and workload estimate (Phase 1)

Reproduced from section 3.1 for quick reference.

Module	**Complexity**	**Estimate (person-months)**
1. Entity register	Moderate	8–14
2. Collection / declaration	Moderate to high	12–20
3. Calculation engine	Moderate	6–10
4. Verification / audit	High	14–24
5. Billing / tracking	Low to moderate	6–10
6. Dashboard	Low	4–8

7.1 bis Annex A' — Table of modules and workload estimate (Phase 3)

Module	Complexity	Estimate (person-months)
13. Platform-identifier attachment	Moderate	8–14
14. Collection / declaration (platform income)	Moderate to high	10–16
15. Calculation engine	Moderate	6–10
16. Verification / audit	High	12–18
17. Dashboard	Low	4–8

7.2 Annex B — Typology of IoT devices currently in use

Lists the broad categories of existing connected devices, useful for scoping the declarations expected under Module 2 and for anticipating the evolution of the device fleet concerned.

- **Consumer / connected home**: home automation (thermostats, lighting, locks, cameras), connected appliances, voice assistants and smart speakers, connected leisure devices (TVs, consoles, toys).
- **Wearables**: connected watches and bracelets, wearable medical devices (glucose monitors, heart monitors), connected clothing.
- **Industrial IoT (IIoT)**: production and predictive-maintenance sensors, factory automation, connected industrial robots, asset tracking (RFID, GPS).
- **Connected health**: implantable or worn connected medical devices, telemedicine and remote monitoring, connected hospital equipment.
- **Smart cities / infrastructure**: traffic management, smart public lighting, smart electrical grids, connected meters, environmental sensors.
- **Connected agriculture**: precision soil and irrigation sensors, livestock tracking, agricultural drones.
- **Vehicles and mobility**: connected vehicles and driving assistance, real-time tracked transport fleets, connected charging stations.
- **Service robotics**: domestic robots (vacuum cleaners, autonomous mowers), companion/assistant robots (emerging), commercial service robots (reception, logistics, delivery) — a category directly linked to the argument in section 2.5.

7.3 Annex C — Detailed volume figures

Category	Order of magnitude
Active companies worldwide	~ 300 to 350 million
States + sub-national government bodies	~ 190 states, several hundred thousand entities at the municipal level
Associations / non-profit organizations	Several tens of millions (highly uncertain

Category	Order of magnitude
	estimate)
Active domain names worldwide	~ 350 to 370 million
Total base used for sizing (Phase 1)	~ 400 to 500 million entities

7.4 Annex D — Sources and figures on income generated without an owned site

Figures cited in section 2.6, with their sources, replacing the unsourced estimates from the July 31, 2026 version. These figures remain orders of magnitude, to be treated as ranges rather than single values in any external communication of the note.

Data point	Order of magnitude	Source(s)
Size of the global creator-economy market, 2025	~ $250 to 255 Bn	Grand View Research; Research and Markets
Size of the global creator-economy market, 2026 (projection)	~ $310 to 325 Bn	Grand View Research; Research and Markets
Number of active creators worldwide	~ 200 to 300 million	Multiple convergent research firms
Share of creators earning more than $100,000/year	~ 4%	Aggregated creator-economy analyses, 2026
Share of creators earning less than $15,000/year	~ 50%	Aggregated creator-economy analyses, 2026
Global influencer-marketing market, 2025	~ $32.5 Bn	Influencer Marketing Hub / Mordor Intelligence
Global influencer-marketing market, 2026 (projection)	~ $40.5 Bn	Mordor Intelligence
Cumulative YouTube payouts to creators, 2021-2025	> $100 Bn	Public announcements by Alphabet / YouTube (2025)

7.5 Annex E — Indicative breakdown of the $30 Bn/year target by country (new, September 12, 2026)

Calculation carried out at Pierre Genevier's request by Microsoft Copilot, based on nominal GDP-by-country data for 2026 (IMF, as reproduced on a Wikipedia page consulted on September 12, 2026). Method: Amount allocated to a country = $30 Bn × (country's GDP ÷ sum of GDPs used). See the methodological caution in section 2.4 ter on the 1 to 2% discrepancy between these values and the latest direct IMF estimates for certain countries — the order of magnitude and relative ranking are not affected.

Top 20 countries by nominal 2026 GDP (IMF) and corresponding share of the $30 Bn

Rank	Country	2026 GDP ($ Bn)	Share of $30 Bn ($ Bn)
1	United States	32.384	9.60
2	China	20.852	6.20
3	Germany	5.453	1.60
4	Japan	4.379	1.30
5	United Kingdom	4.265	1.27
6	India	4.153	1.23
7	France	3.596	1.07
8	Italy	2.738	0.81
9	Russia	2.656	0.79
10	Brazil	2.636	0.78
11	Canada	2.507	0.74
12	Australia	2.124	0.63
13	Mexico	2.121	0.63
14	Spain	2.091	0.62
15	South Korea	1.931	0.57
16	Turkey	1.640	0.49
17	Indonesia	1.540	0.46
18	Netherlands	1.450	0.43
19	Saudi Arabia	1.389	0.41
20	Switzerland	1.147	0.34

Top 50 countries by nominal 2026 GDP (IMF) and corresponding share of the $30 Bn

Sum of the GDPs of the top 50 countries used by Copilot: $126,295,331 M; proportion factor: 30 / 126,295 ~ 0.0002376.

Rank	Country	GDP ($M)	Share ($M)
1	United States	32,383,920	7,694
2	China	20,851,593	4,955
3	Germany	5,452,858	1,296
4	Japan	4,379,253	1,041
5	United Kingdom	4,264,794	1,014
6	India	4,153,191	987
7	France	3,596,094	855
8	Italy	2,738,164	651
9	Russia	2,656,452	631

Rank	Country	GDP ($M)	Share ($M)
10	Brazil	2,635,912	626
11	Canada	2,507,340	596
12	Australia	2,123,963	505
13	Mexico	2,120,855	504
14	Spain	2,091,222	497
15	South Korea	1,931,008	459
16	Turkey	1,640,223	390
17	Indonesia	1,539,872	366
18	Netherlands	1,449,704	344
19	Saudi Arabia	1,388,676	330
20	Switzerland	1,146,911	272
21	Poland	1,134,248	270
22	Taiwan	976,719	232
23	Ireland	779,381	185
24	Belgium	776,730	184
25	Sweden	760,481	181
26	Israel	719,848	171
27	Argentina	688,378	164
28	Singapore	659,572	157
29	Austria	623,719	148
30	United Arab Emirates	621,546	148
31	Norway	599,406	142
32	Thailand	579,996	138
33	Colombia	539,530	128
34	Vietnam	527,266	125
35	Malaysia	516,428	123
36	Philippines	512,222	122
37	Bangladesh	510,705	121
38	Denmark	503,772	120
39	Romania	480,834	114
40	South Africa	479,964	114
41	Pakistan	452,000	107
42	Hong Kong	450,138	107
43	Czech Republic	432,597	103
44	Egypt	429,645	102

Rank	Country	GDP ($M)	Share ($M)
45	Chile	407,850	97
46	Peru	380,900	91
47	Portugal	380,637	91
48	Nigeria	377,365	90
49	Kazakhstan	360,456	86
50	Finland	337,669	80

Additional ranks, for illustration of the lower end of the ranking

Rank	Country	GDP ($M)	Share ($M)
51	Iran	300,293	71
52	New Zealand	278,636	66
71	Croatia	116,574	28
72	Côte d'Ivoire	112,115	27
101	Sudan	44,688	11
102	Iceland	43,800	10
121	Trinidad and Tobago	26,836	6
122	Chad	25,628	6
190	Nauru	196	0.05
191	Tuvalu	65	0.02

7.6 Annex F — Illustration of the cap by company category for the US share ($7.69 Bn/year) (new, September 12, 2026)

Illustration carried out at Pierre Genevier's request by Google Gemini, based on U.S. Census Bureau data (approximately 6 million employer companies, approximately 30 million independents/non-employers). Objective: to demonstrate that a progressive, capped schedule by company category, consistent with the parametric calculation engine (section 2.1, Module 3), makes it possible to reach the US share of the $30 Bn/year target (~ $7.69 Bn, Annex E) without imposing a confiscatory amount on any given company. See the scope and limitations of this illustration in section 2.4 ter.

Company category	Indicative liability criterion	Approximate number of entities in the US	Capped levy / indicative annual flat fee	Revenue generated (illustrative)
Very large groups (Tech / Cloud / e-commerce)	Revenue > $50 Bn + high digital dependence	~ 30 to 40 entities	Cap ~ $25 M/year	~ $0.88 Bn
Large companies (Fortune 500	Revenue > $10 Bn	~ 460 entities	Cap ~ $3 M/year	~ $1.38 Bn

Company category	Indicative liability criterion	Approximate number of entities in the US	Capped levy / indicative annual flat fee	Revenue generated (illustrative)
excluding Tech)				
Mid-size enterprises	Revenue between $50 M and $10 Bn (or > 500 employees)	~ 30,000 entities	~ $10,000 to $80,000/year	~ $1.20 Bn
Medium-sized companies (SMEs)	Revenue between $5 M and $50 M (20 to 499 employees)	~ 200,000 entities	~ $1,200 to $5,000/year	~ $0.80 Bn
Small businesses	Revenue between $500K and $5 M (1 to 19 employees)	~ 1.3 million entities	~ $150 to $600/year	~ $0.45 Bn
Micro-enterprises and independents	Revenue < $500K (with an owned site or domain)	~ 12 million entities	Flat fee ~ $15 to $45/year	~ $0.36 Bn
Phase 3 — platform income (US component)	Creators / sellers without an owned site (YouTube, Amazon, Etsy...)	Millions of people	~ 0.5% to 1% withheld at source (exemption below $2,000/year)	~ $2.20 Bn
Government bodies, universities, organizations	Public institutions, large NGOs, universities	~ 80,000 entities	Flat fee ~ $500 to $10,000/year	~ $0.42 Bn

Total illustrative figure for the United States: ~ $7.69 Bn, consistent with the share attributed to the United States by the GDP-proportional breakdown (Annex E).

Reading of the proposed caps

- Digital giant (cap ~ $25 M/year): a group such as Alphabet or Amazon would pay a fixed amount, marginal relative to its results, but together they would provide close to a billion dollars to the IO's budget.
- Traditional large company (cap ~ $3 M/year): a player such as Walmart or Home Depot would no longer pay $277 M (raw calculation proportional to revenue on the Fortune 500 alone), but a much lower capped amount.
- SME (~ $1,200 to $5,000/year): an amount comparable to a standard professional software subscription.
- Micro-business / independent (~ $15 to $45/year): an amount comparable to the annual cost of a domain name with its management services.