Metered Memory: Unpaid Labour and the Administration of Personal Archives

Flyxion
Independent Researcher

November 2026

Abstract

Digital platforms invite people to entrust them with correspondence, photographs, and notes, charge for the continued existence of those records, and later instruct the same people to remove material in order to remain within an allocation. This essay argues that the maintenance of personal memory under such arrangements is a recurring form of unpaid labour, concealed by interface language that presents it as a trivial chore. The central difficulty is epistemic: pruning requires a prediction of future relevance, and relevance is partly produced by events that have not yet occurred. The essay distinguishes three claims that are often run together, namely that deletion demands judgment, that an observed judgment can serve as a learning signal, and that a particular provider uses it for a specified purpose. It sets the labour argument against published evidence on the pricing of storage and compute: the gap between rent and hardware cost, the cost of leaving, free tiers that gather records and later count new deposits against an allowance, quotas that bill copies by visible volume, and hardware prices that small buyers cannot negotiate. It examines retention without usable access, the limits of account export, and the gap between a visible deletion and its propagated consequences. It concedes that storage has real costs, and proposes a baseline of storage and compute provided free at the point of use, portable, not revocable, and legible, funded as infrastructure. It concludes that an archive should be judged by how well it preserves the freedom to defer judgment about what mattered.

1 The Promise of External Memory

People keep letters, photographs, and notes for a reason that is older than any computing service: continued availability allows attention to move elsewhere without giving up the possibility of return. A drawer of correspondence is a decision not to carry its contents in mind. Writing and recording extend memory across time, and archives exist largely because the significance of what they hold is unsettled at the moment of deposit. A land survey, a parish register, or a collection of family letters acquires its value from being kept before anyone knows what question will eventually be put to it. Vannevar Bush framed the aspiration in mid-century terms, imagining a device in which an individual stores books, records, and communications so that they may be consulted with speed and flexibility [3]. The ambition was to relieve the person of the burden of holding everything while retaining the power to recover it.

Digital services inherited this expectation and, for a period, appeared to fulfil it. Abundant storage and capable search made continual manual sorting seem unnecessary, and the practical advice that accompanied early webmail was to stop organizing and simply search. The proposal had a clear logic. If retrieval is cheap and retention is cheap, then the labour of deciding what to discard can be avoided without loss, and the person is released from a form of administration that physical archives had always imposed through limits of space.

The expectation is now met with a different message. A storage warning appears and presents the deletion of old correspondence as a modest daily chore, something that a conscientious user would simply do. The instruction is brief, and it carries a set of presuppositions that are rarely stated. It assumes that the user has spare time for review, that the user knows which records are dispensable, and that the user should accept responsibility for keeping the service within limits that the service itself defines. This scene will serve as the concrete problem for what follows, and the broader political economy of the arrangement can be developed from it.

One distinction governs the whole essay. Some claims concern conceptual mechanisms: what an interface demands of a person, what deletion presupposes, what an observed choice could in principle supply to a system that records it. These can be established by analysis. Others concern the documented practices of particular companies: what a named service stores, how it prices, what it does with behavioural data. These require evidence about specific systems, and they are asserted only where a published source supports them and is cited, with the limits of that source stated. Claims about what any

provider does with behavioural data are not asserted at all. Wherever the text speaks of platforms in general, it makes the first kind of claim, and the argument should be read as a basis for scrutiny and not as a finding of misconduct by any provider.

The storage warning is one instance of a wider arrangement that this essay calls the meter. A meter turns a relationship into a quantity. A person who kept a shoebox of letters kept them, whereas a person who stores the same letters with a service rents shelf space from a landlord who may raise the rent, change the terms, or close the account. The sums are small on any single day, which is how the arrangement works: they are too small to contest and too regular to escape, and they accumulate for as long as the person wishes to remember anything.

The essay pursues two arguments of different weight. The first, developed in the early sections, is that the maintenance demanded under the meter is labour, and that its epistemic structure, a forced prediction of future relevance, makes it an unsuitable thing to impose. The second, developed in the later sections, is a proposal: a baseline of storage and compute should be provided free at the point of use, funded as infrastructure in the way roads and libraries are funded. The claim is not that the physical world is free. Disks cost money, operation consumes electricity, and buildings cost land. The claim is that a cost borne in common and recovered through taxation or cooperative levy is a different thing from a price charged to each person at each use by a seller who also sets the terms. The pricing evidence does not establish the cost of an alternative system. It shows only that hardware expenditure alone cannot explain or justify the existing terms of access, and the normative argument, stated in the final sections, carries the proposal without depending on it.

2 The Work Concealed Inside an Interface

Making room is not a single act. To free a meaningful amount of space, someone must inspect material, separate duplicates from variations that differ in ways that matter, determine whether copies exist elsewhere, evaluate attachments whose contents are not visible from a subject line, anticipate what may be needed later, and bear the consequences of any mistake. A bulk selection tool shortens the motion of the hand but not the judgment behind it, because the judgment lies in deciding what the selection should contain. The delete control compresses this entire sequence into one apparently trivial gesture, and the compression is what allows the task to be described as small.

The description of this activity as labour requires a conception of labour that includes

the expenditure of attention and judgment, even where there is no wage, no employer, and no formal assignment. The tradition of research on invisible work has long made the point that effort which sustains a system can go unrecognized precisely because it is distributed, informal, and absorbed into ordinary routines [23, 24]. Terranova's account of free labour in the digital economy extended the observation to online participation, arguing that voluntary and enjoyed activity can nonetheless be structurally necessary to the value of a platform [25]. The present argument is narrower. It concerns the specific operations of uploading, describing, arranging, reviewing, migrating, and deleting, each of which is performed by users and each of which either sustains the service or produces information of potential commercial use.

A distinction is needed between freely undertaken curation and work induced by a service's constraints. Many people enjoy sorting photographs, and no argument here depends on denying that pleasure. The relevant case is the one in which the sorting is prompted by a quota, undertaken under a warning, and performed at a time chosen by the service's limits rather than by the person's interest. That the objects belong to the person does not settle who created the necessity of handling them, who benefits from the handling, or who bears its cost. A tenant who is told to clear a storage unit by a deadline is handling his own possessions, and it would still be a mistake to say that the clearing is purely his own affair.

3 The Demand to Predict Future Relevance

The epistemic centre of the problem is that pruning is a decision about future usefulness taken before the circumstances that establish usefulness exist. The person is asked to classify records as dispensable or retained on the basis of the present, while the significance of a record is frequently a product of what happens afterward. An email may become evidence of an agreement when a dispute arises years later. A photograph of a street corner may acquire historical weight when the corner is demolished. An incidental note may turn out to answer a question that no one formulated at the time of writing. These cases are not alike. The agreement may have carried evidential significance from the day it was written, so that the later dispute only made an existing property urgent. The photograph's importance depends on circumstances that did not yet exist, so that the later demolition helped constitute the significance the image now has. The note may have been recognizable as useful to a sufficiently informed reader, or it may have become useful only through a question that no one could then have framed.

Relevance can accordingly be presently unrecognized, presently unpredictable, or subsequently constituted. In the first case a more careful reader at the time might have noticed it, and better inspection helps. In the second, the relevant considerations exist but cannot be foreseen from the position of the person deciding. In the third, later events help create the significance. Present inspection cannot reliably settle every future use across these cases, and that is the only conclusion the argument requires. The difficulty is a structural feature of records and not a failure of diligence. It is also the reason archival appraisal is skilled and consequential work. Archives necessarily select, and professional appraisal does not uniformly favour keeping everything uncertain: appraisers weigh volume, cost, redundancy, privacy, and evidential value, guided by training and institutional policy. That is precisely the kind of work that a quick-tidy instruction conceals, and it hands the work to a person who has neither the training nor the policy to rely on.

Preservation leaves such questions open, whereas destruction can foreclose their eventual investigation. The asymmetry is the heart of the matter. A retained record that proves worthless has cost the holder some space and some index entries. A destroyed record that proves essential cannot be recovered by any later effort, however well funded. When a platform asks the user to delete in order to remain within a limit, it assigns the user both the labour of selection and the risk of an inadequately informed decision, and the personal and evidential consequences of a mistaken deletion fall primarily on the user, while the provider receives the immediate benefit of reduced storage obligations. Providers are not wholly insulated, since reputational damage and lost customers remain possible, but those costs are diffuse and delayed by comparison. Mayer-Schönberger has argued, in a different spirit, that default forgetting might be a virtue in a digital age [14]. The argument has force where forgetting is chosen and bounded by the person's own purposes. It is far weaker where forgetting is demanded by an external limit and executed under time pressure.

4 Uploading and Deleting as Successive Economic Demands

The sequence of encouragement to accumulate followed by pressure to reduce can look like a reversal of policy, and for that reason it deserves examination as a single arrangement. Uploading deepens dependence on a service by concentrating records, relationships, and

habits inside it. Each additional year of correspondence and each added album raises the cost of leaving, and the same concentration that makes the service useful makes departure expensive. Later, quotas and migration difficulties can make continued preservation contingent on payment or on further work. The two stages are not contradictory. They are successive points at which the provider can convert the user's accumulated position into revenue, reduced cost, or continued attachment.

This essay uses the term economic joint for such a point of value capture: an interaction at which the provider can turn the user's need into revenue, lower costs, increased dependence, or useful behavioural information. A subscription upgrade offered at the moment a quota is reached is a joint. So is the absence of a simple, complete export at the moment a user considers leaving. The concept is borrowed in spirit from the economics of switching costs, in which an installed base of customers becomes a source of rents once leaving is costly [21], and from Hirschman's treatment of exit and voice, in which the practical availability of exit conditions what a customer can demand [13].

The claim is structural and does not require a single coordinated intention across companies. Business arrangements can reward different demands at different stages of a relationship, and the user absorbs the cost of adjusting to each. A provider that encourages upload because accumulation increases engagement, and later encourages deletion because storage has become a cost, need not be acting duplicitously. It may simply be responding to incentives that change over the lifetime of an account. The consequence for the user is the same, and so is the proper object of analysis.

5 The Free Tier

The free tier is the first economic joint in this sequence, and it is how the public's records are gathered before they are billed. The pattern is visible from the beginning of the consumer cloud. Gmail was announced in 2004 with one gigabyte of storage, several hundred times what Hotmail and Yahoo offered at two or four megabytes [20]. The offer worked as an argument as much as a product: keep everything, delete nothing, search instead of sorting. In May 2013 Google pooled Gmail, Drive and Google+ Photos under a single free allowance of 15 gigabytes, which raised the free total for many users, and extended paid storage to Gmail at the same time [17]. That allowance was still 15 gigabytes when the next change arrived. On 11 November 2020 Google announced that the unlimited free storage of compressed photographs in Google Photos would end

on 1 June 2021, after which new uploads would count against a free allowance of 15 gigabytes [11]. Photographs already backed up under the old terms were exempted, and the announcement said the new allowance would hold, in its words, “roughly three more years worth of memories” for more than eighty percent of users. Beyond the allowance, storage was sold through a subscription plan.

Two features of this episode deserve attention, and neither is a charge of bad faith. The first is the exemption. The company honoured the old uploads, which is to its credit, and the honouring shows that the arrangement was entirely a matter of discretion. A promise that holds for years and is then amended by announcement is a policy and not a right. The second is the sequence. An offer of unlimited free storage collects an enormous number of photographs. Once the collection exists and the habit of depositing has been formed, new deposits count against a finite allowance, with payment required beyond it. The people whose family records sit in the service are not the ones who chose the change, and they are invited to pay to go on adding to the archive past the allowance.

A record held at a provider’s discretion is held conditionally, and conditions can change. The remedy is to move the record out of the category of things a firm gives away to attract customers and into the category of things a society provides because a life without them is diminished. A baseline allowance that cannot be withdrawn by announcement is the difference between a tenancy and a deed.

6 Rent on a Disk

Once the allowance is exhausted, the price charged bears an uncertain relation to the cost of the hardware, and the gap between what storage costs to buy and what it costs to rent is the next thing to examine. A published account of Amazon’s S3 Standard storage gives the price as $0.023 per gigabyte per month for the first fifty terabytes in the US East region [8]. That is $23 per terabyte per month. A price survey of new internal hard drives sold on Amazon in October 2026 puts the cheapest per-terabyte price across common capacities at roughly $24 to $41 [12]. A separate tracker using Newegg US listings provides a cross-check from another retail dataset. Its snapshot retrieved on 8 November 2026 reports minimum prices for new desktop hard drives of $29.72 per terabyte at 14 terabytes, $31.67 at 12 terabytes, and $30.50 at 18 terabytes, excluding tax and shipping [5]. These are selected capacity-tier minima and not market averages, but they corroborate the broad scale of the hardware figures used here. A month of renting a terabyte from the largest storage provider therefore is of the same order as the

purchase price per terabyte of selected new high-capacity drives.

This is not a cost accounting, but it does bound the hardware component. The service replicates data across several facilities, so the customer is buying more than one disk, and the provider pays for power, buildings, staff and the engineering that makes the replication work. Standard storage is also the highest-priced tier of a range, and colder tiers are cheaper. Concede all of it. Triple the disk cost to allow for replication, which gives an illustrative hardware figure between $72 and $123 per terabyte. A customer who keeps that terabyte for one year pays $276, between two and nearly four times that figure, and then pays again the next year and the year after, for as long as the file exists. Disks wear out and must be replaced over the life of a service, and tripling is only an illustration of replication, not an accounting of the other costs just listed, so none of this shows what the service costs to run or what margin it earns. What it shows is that the physical cost of the capacity does not determine the billing structure. The rent is perpetual, and it falls on the one thing that cannot be allowed to lapse, which is the record itself.

The point is made sharper by retrieval. The same survey of providers lists internet egress at $0.090 per gigabyte for the first ten terabytes on AWS, $0.087 on Azure and $0.120 on Google Cloud at its premium tier, and $0.00 on Cloudflare R2, which advertises unlimited free egress [6]. At the AWS rate, reading one terabyte back out costs $90, about four months of the $23 monthly charge for storing it. Keeping a thing costs one amount and touching it costs another, larger one, and the second figure is the one that matters to a person who wants to move. The existence of a competitor that charges nothing for egress does not show that delivering bytes is free. It shows that billing for it separately is a commercial choice and not a law of physics, which is all the institutional argument needs.

The limits of this evidence should be stated once and plainly. The sources used in this essay are price comparisons and trade analyses of varying rigor, one of them a vendor's press release, and none of them is a cost study of any provider. They establish orders of magnitude and structure. A full accounting of what it costs to run storage at scale could narrow the gap between rent and hardware, and nothing in the later argument depends on its not doing so.

7 Charging for Copies

A further feature of the arrangement concerns what a quota counts. Where a quota counts retained revisions separately, or counts identical attachments against each recipient's allowance, logical volume can exceed the underlying payload storage required, whatever the provider does underneath. In such cases whether the provider stores the bytes once or many times is the provider's business, and the user pays the same either way. Storage systems that identify content by a hash and keep a single copy of identical content, and that store a revision as the difference from its predecessor, are decades old and standard in version control. Where a provider uses such techniques internally, which the essay does not claim to know for any particular firm, the marginal cost of the payload of one more identical copy is close to nothing. That is not the marginal cost of the copy as a whole, since metadata, access control, availability, indexing, and independent retention commitments remain: a shared attachment must stay available when one recipient deletes it and another keeps it. Billing by logical volume can therefore serve purposes beyond concealing physical savings. The stronger criticism is that the customer cannot assess whether, or how, such savings enter the price.

The consequence is a tax on the habits that make a record valuable. Keeping drafts, keeping every revision and keeping what everyone else also keeps are the things such a quota punishes. The rational response of the customer is to delete, to summarize or to export, which shrinks the record, and the record is what the customer came to the service to keep. A price that scaled with physical consumption would charge for what is new relative to what is already stored. How much that saves depends on the content. Identical files share storage entirely, and successive drafts of a text document differ by small changes that can be represented economically, but compressed photographs, video, encrypted data and heavily rewritten documents share little, and a long history of such material can cost far more than its final version. The essay promises no general ratio. The argument is that logical volume can differ from physical consumption, that the customer ordinarily cannot inspect the relationship, and that a price set on the logical figure alone gives the customer no means of telling how any difference is shared. Billing by visible volume also teaches a habit, which is to treat history as a liability to be trimmed. This is the point at which pricing and the labour of pruning described in section 2 meet, since the quota is what turns the review of one's own history into a recurring chore.

8 Subtraction as Signal

A distinctive consideration is that subtraction can produce information. To select photographs for deletion is to express distinctions among them. Repeated selections can reveal patterns of preference, redundancy, significance, or tolerable loss, and the comparison between what is kept and what is removed has the structure of a labelled contrast. Observed deletion choices can be described as behavioural feedback or as potentially comparative preference data. Where a system observes the user's removals and retains suitable representations of them, it could in principle use them as comparative supervision about which images, messages, or documents a person regards as expendable.

Three claims must be kept apart here, because the argument is easily damaged by running them together. The first is that deletion requires judgment, which is plainly true and was established in the discussion of work above. The second is that an observed judgment can provide a learning signal, which is a general fact about how preference data can be used and does not depend on any particular company. The third is that a particular provider actually uses that signal for a specified purpose. The first two claims establish a mechanism. The third requires evidence about specific systems and data flows, and is not asserted here.

A deletion supplies an observed choice, not necessarily a reliable ranking, and the interpretation of such signals is also compromised by the conditions under which they are produced. Deleting something to free space does not necessarily express dislike or lack of value. A person who removes a photograph in order to receive an incoming message has not thereby ranked it below the photographs that remain. The pressure under which the choice is made is part of its meaning, and a signal harvested without that context risks encoding the quota rather than the person's values. This is a limit on the informational value of deletion under duress, and it is also a further reason to doubt that the labour of pruning is a reliable expression of what the user cares about.

9 Retention Without Usable Access

Storage capacity is only one dimension of an archive. A collection becomes useful through search, organization, context, stable references, and the ability to share meaningful groups of records with others. An album preserves relationships among photographs that no single image carries, and a search interface allows a later question to recover an earlier record without the person having to recall where it was placed. When these functions

degrade, the collection may remain intact in a narrow sense and yet cease to serve the purpose for which it was assembled.

Consider the situation of a person who has stored years of photographs with a service and later tries to share an album of historical material with relatives, or to find a specific message from long ago using remembered fragments. Whether such tasks succeed depends on the capabilities of the particular product at the particular time, and those capabilities change. At the level of structure, a service can continue to hold records while making their retrieval so cumbersome that preservation loses much of its value to the person who supplied them.

The distinction at stake is between a provider's holding of stored material and the person's practical command over it. Holding is a fact about where bytes reside. Command is a fact about whether the person can retrieve, interpret, organize, and migrate records without depending on the original custodian, how quickly, and with what reliability. Sparrow, Liu, and Wegner showed that people adapt their own memory strategies to the availability of external lookup, remembering where information can be found rather than the information itself [22]. Where that lookup is unreliable, the adaptation becomes a liability, because the person has relinquished internal retention on the strength of an external capacity that does not perform as assumed.

10 Export, Exit, and the Transfer of Technical Responsibility

The limits of an account archive illustrate the same pattern in the context of leaving. Regulatory regimes have begun to address the matter. Article 20 of the European General Data Protection Regulation gives a data subject the right to receive personal data that the person has provided to a controller, in a structured, commonly used, and machine-readable format, and to transmit it to another controller, where processing rests on consent or contract and is carried out by automated means [9]. The scope is limited to data the person provided, the provision does not oblige providers to adopt compatible systems (Recital 68), and it guarantees a format and not a functioning archive. It supports the modest claim that a copy must be obtainable, and not the broader account of export developed here, which concerns usability. Receiving a copy is not the same as holding an operationally independent archive. An export may deliver the former while leaving the user to reconstruct a functioning archive from a directory tree,

sidecar metadata files, media in varying encodings, and formats that were designed for the provider's internal use. Search, navigation, contextual relationships between records, and secure access then become the user's technical assignment.

Portability should accordingly be assessed through the effort required to regain ordinary use, rather than through the mere availability of a download. A meaningful export preserves enough structure to support continuity: albums remain albums, conversations remain threaded, dates and places remain attached to the items they describe, and the whole can be browsed without the originating service. A large quantity of files that must be rebuilt into an archive by someone with specialist knowledge is a transfer of responsibility more than a transfer of the archive. The labour argument therefore extends beyond deletion to migration, indexing, maintenance, and recovery, each of which is a task that the service performed implicitly while the account was active and that the user must now perform explicitly.

This is also a point at which unequal capacity matters. A technically skilled person can write scripts to parse an export, build a local index, and host a private interface. Most people cannot, and for them a complete but unstructured export is a form of practical loss. The cost of leaving is thus distributed according to technical ability, which means that the burden of the arrangement falls most heavily on those least able to discharge it.

Fees add a further obstacle. A fee for taking data out is not a fee for storing it, and its main effect is to make leaving expensive. The largest providers have acknowledged as much by retreat. According to the same survey, AWS, Azure and Google Cloud each waived egress fees in 2024 for customers fully leaving their platforms, and the waiver applies only to complete departures, not to routine transfers or the partial movement of a workload [6]. A customer is thus permitted to leave if the customer leaves entirely and says so, and is charged if the customer merely wishes to keep a foot in another door.

Regulators have gone further. Article 29 of the European Data Act ends switching charges from 12 January 2027, with reduced, cost-based charges permitted from 11 January 2024 until then, and Article 2 defines switching charges to include data egress charges incurred in switching [10]. A trade analysis reports that the ban covers infrastructure, platform and software services and binds non-European providers selling into the Union as well as European ones [4]. The Regulation does not abolish every outbound transfer charge: Article 34 allows providers to keep charging for egress used for in-parallel use of two services, though only up to the costs incurred [10]. Nor does it ignore technical lock-in, since the recitals call for transfer in a commonly used, machine-readable format

by means of open interfaces [10]. The trade analysis nonetheless cautions that proprietary interfaces, vendor-specific formats and identity binding can still produce lock-in without any fee. The essay makes no claim about other jurisdictions, whose law is outside the sources used here.

The inference is narrower than it may appear. A charge that recovers a genuine cost is still a charge, and a genuine cost can still be a barrier to leaving. That the legislature assigned the cost of switching to the provider shows that cost recovery does not settle who should bear the cost when charging the customer impedes departure. The same reasoning supports the proposal below. A second lesson comes from the trade analysis: abolishing the toll does not by itself end lock-in, since format, interface and identity can remain the provider's to define. A baseline of free storage that is worth having must therefore come with open formats and portable identity as conditions of provision, and not as favours.

11 Deletion, Propagation, and the Limits of Recall

The word deletion names several different operations, and a single visible command cannot establish that all of them have occurred. Disappearance from an interface is one operation. Removal from active storage is another. Expiry from backups, which may follow on a schedule independent of the user's action, is a third. Removal of copies held independently, whether by correspondents, by synchronized devices, or by downstream systems, is a fourth. These operations differ in scope and in timing, and the person who presses the control sees only the first.

The comparison with newspaper corrections and scientific retractions is instructive. When a newspaper corrects an article or a journal retracts a paper, the source is changed or flagged, and yet descendants of the original persist: quotations, citations, derivative analyses, and decisions made in reliance on the earlier claim. Changing the source does not automatically revise its descendants or undo what they informed. Studies of retracted literature, beginning with Budd, Sievert, and Schultz's examination of retracted biomedical articles, have found that retracted work continues to be cited, mostly without acknowledgment of the retraction [2], which shows how weakly a correction at the origin propagates. Likewise, deleting an original record need not remove every derived representation or every effect it has already had.

The argument does not depend on the claim that companies retain everything indefinitely, and that claim is neither assumed nor needed. The point concerns a mismatch

between the simplicity of the visible operation and the complexity of its consequences. The interface offers a single verb, while the underlying situation involves layered stores, delayed processes, and external copies. A user who relies on the verb to reason about outcomes will systematically misjudge them, in one direction by believing that something is gone when traces remain, and in the other by believing that something remains recoverable when it has been destroyed. Both errors arise from the same compression of a complex process into a gesture.

12 Forgetting, Archives, and Funes

The strongest philosophical objection to a defence of preservation is that an inability to let go is itself a harm. Borges's story of Ireneo Funes, who forgets nothing and is crushed by the weight of recollection whose particulars resist abstraction, is the classic image of that harm [1]. Nietzsche's essay on the uses and disadvantages of history for life makes a related claim, that a being incapable of forgetting would be unable to act, and that a capacity for the unhistorical is a condition of health [16]. These authors identify a real danger, and a defence of archives must answer them.

Nietzsche offers a more specific account of forgetting in the opening section of the second essay of *On the Genealogy of Morality*. Forgetting is there an active faculty that regulates access to consciousness, allowing experience to be incorporated without remaining perpetually present [15, II, §1]. His physiological analogy is digestion: an encounter can contribute to the organism without its constituent details remaining objects of attention. Active in this sense does not mean consciously effortful or deliberately chosen, and the activity should be distinguished from the deliberate work of inspecting records and deciding which to destroy. A platform that demands such inspection makes the regulation of an external archive into an explicit cognitive assignment. Nietzsche's account thus helps separate the capacity to cease attending from the labour of deciding what must cease to be available.

The answer is that Funes suffers from compulsory recollection, not from retention. His difficulty, as Borges tells it, is exhaustive particularity: he cannot abstract, because each perception remains distinct from every other, and thought itself is impeded. An archive that remains available without demanding attention is a different kind of thing. It holds material without presenting it, and it yields material only when asked. The nightmare of equal and total presence does not describe a shoebox of letters in a cupboard, and it does not describe a well-organized photograph library that is opened on request.

External records in fact support psychological forgetting. A person can cease rehearsing an event because a recoverable record exists, and the knowledge that the record can be consulted relieves the pressure to keep it active in mind. This is the benefit that archives have always offered, and it is the function that Bush's proposal aimed at. Compulsory pruning can reverse it. A service that requires the user to review stored material repeatedly, and to judge its future relevance in the process, reintroduces the attentional burden that external memory was supposed to relieve. Periodic review does not reproduce exhaustive memory, and the analogy to Funes should not be pressed that far, but it restores a demand for rehearsal and anticipatory judgment that the archive had lifted. The philosophical objection to hoarding therefore does not transfer to the case of a quiet and accessible archive, and may count against the demand to prune.

13 Costs, Constraints, and Responsibility

The strongest objection to the argument so far is practical. Storage, indexing, security, and reliable preservation require resources, and indefinite retention cannot simply be assumed to be costless. Servers consume power, data must be replicated to survive failures, and security obligations grow with the volume of sensitive material held. A provider that offered unlimited free retention without any means of recovering these costs would be unsustainable, and a critique that ignored this would be unserious. The material constraint is conceded. The pricing evidence in section 6 bears on how those costs are allocated, but it is not a cost study of any provider, and it cannot settle what a service actually costs to run.

What the concession leaves open is how costs are allocated and how they are represented to the person who bears them. Several questions follow. Are prices and limits intelligible, so that a user can anticipate when a threshold will be reached and what it will cost to exceed it? Does retrieval remain dependable over the life of the account? Is migration practically achievable, rather than formally permitted? Can users make preservation decisions without recurring pressure, or is each year of retention accompanied by a renewed demand for attention? A service's resource costs do not make the user's time valueless, and they do not render every imposed workflow justified. A provider may need to charge for storage and still be criticized for charging in a form that conceals the labour it additionally requires.

The comparison must account for both sides of the arrangement. Users receive substantial value from these services, including search, synchronization, and protection

against local failure, and a fair assessment includes those benefits. The aim is not to establish that platforms are net extractors in every case. It is to insist that the evaluation include the user's time and risk alongside the provider's expenses, so that neither is treated as free by default.

14 The Price of Being Small

A last pricing question is purchasing power. In September 2026 a server vendor, eRacks, published tracking figures for hard drive prices between March and August of that year [7]. Its table gives cost per terabyte rising from $16.25 to $31.25 for one Seagate Barracuda line, from $24.84 to $41.22 for a WD Red Pro line, from $21.78 to $33.75 for a Seagate IronWolf Pro line, from $29.79 to $44.17 for a WD Gold line, from $22.17 to $30.64 for a Toshiba enterprise line and from $24.33 to $27.92 for a Seagate Exos line. The increases range from fifteen to ninety-two percent. The company attributes them to demand, of the same kind that raised memory and graphics-processor prices earlier in the year, and observes that consumer lines rose more than enterprise lines.

The source needs handling with care. It is a press release from a vendor describing its own price tracking, not an independent market study, and its headline range of forty to ninety percent does not match its own table, which runs from fifteen to ninety-two. The release names demand and does not name the buyers. That the buyers who set the price are the largest ones is therefore the essay's inference and not the source's finding, though it is not an eccentric inference: when a few purchasers can contract capacity in bulk and in advance, the ones who buy a drive at a time buy what is left, at the price the bulk buyers have established. The release also reports that the enterprise lines moved least and the consumer lines most, which complicates the inference. The pattern fits enterprise buyers having contracted ahead while spot buyers absorbed the increase, and it fits equally well consumer demand rising independently, and the source cannot distinguish them.

This bears on the thesis in two ways. It contradicts the easy slogan that storage is cheap, and the essay does not rest on that slogan. Whether a terabyte of disk costs sixteen dollars or forty-four, the price paid by a customer of a storage service is set by a separate decision. And it shows what an unprotected market does to everyone who is not a hyperscaler. The household that wants a disk, the library that wants a server and the researcher who wants a workstation do not negotiate. They take the price. A common baseline, bought at the scale of the purchasers who now set the terms, would put the

public on the same side of the table as the firms that currently dominate the order book.

15 Compute by the Token

The argument extends from storage to computation, since stored records are now read, searched, and transformed by machine. Compute is metered more finely than storage. A computation is billed by the second of processor time, by the hour of an accelerator, or by the token in a language-model interface, and each unit is small enough that no single charge invites a complaint. The fineness is the instrument. A price that is paid thousands of times a day is never examined, and a meter that runs at the speed of thought makes thinking itself an itemized expense. This section is structural only. It argues from the form of public price lists, offers no price comparison or cost estimate, and should be weighted accordingly.

The token is a billing unit whose relationship to computational work depends on the model, implementation, and serving conditions, so that price and work can come apart in the way the quota shows. Where providers discount repeated input, the discount concedes that some repetition is cheaper to serve, though a small change to a request can invalidate reusable computation, and a repeated prefix does not make a whole interaction cheap. The argument does not rest on any such meter. It rests on access, on predictable provision and on the value of iteration: most intellectual work proceeds by revising, re-asking and thinking again, and a meter that runs at each pass makes the person doing the work bear the price of every pass, at a rate the seller sets.

The case for treating compute like storage is, if anything, stronger, because compute is the more rapidly spreading medium of ordinary intellectual work. Searching a body of text, translating a letter, checking an argument, summarizing a report and drafting a reply are now done by machine, in the same way that arithmetic was once done on paper, and a person who cannot afford the machine is in the position of a person who cannot afford paper. Charging by the unit means that the cost falls hardest on the most active users, who are students, small researchers and anyone working on a long project, and that the rate is whatever the seller chooses. A customer cannot compare the price of a token against the cost of producing it, because that cost is not published.

None of this denies that large-scale computation is expensive. A training run at the frontier costs what it costs. The claim concerns the baseline: the computation an individual needs to read, write, search, calculate and keep a record is small by the standard of modern hardware, and the machinery that supplies it, whether in a public

facility, a cooperative or the person's own device, can be provided at a flat level without metering. Above the baseline, supply can be priced at what it costs to provide it. The proposal is a floor and not a ceiling.

16 An Account of the Service's Actual Cost

A proper account of what a memory service costs combines several items that are usually reported separately or not at all. There are the subscription payments, which are visible. There is the attention spent on sorting, the effort of annotation and arrangement, the work of migration when circumstances change, the uncertainty of not knowing what will remain recoverable, and the potential loss of records whose significance had not yet emerged. Each of these is a real cost to the person, and only the first appears on a statement.

This account is offered as a framework for evaluation. It is not a claim that every interaction benefits the provider more than the user. Establishing unequal benefit requires examining particular arrangements, including the alternatives available to the user and the costs of leaving, and conclusions will differ between services and between periods of a single service's history. A platform that offers generous export, stable search, and transparent pricing stands in a different position from one that does not, and the framework is meant to distinguish them.

Within that qualification the accusation can be stated precisely. Routine interface language can conceal consequential transfers of work and responsibility. A message that describes a deletion as a quick tidy compresses an uncertain, consequential, and unpaid evaluation into a phrase that sounds like housekeeping. The apparent convenience of a service can depend in part on excluding such transfers from the advertised account of what the service costs. Making them visible does not by itself settle whether they are justified, but it is a precondition for any serious discussion of whether they are.

17 Capacity Without Competence

Free capacity delivered as an opaque service would repeat the enclosure under a different name. A person given an allotment by an operator whose storage works by mechanisms the person cannot inspect, in formats the person cannot read, behind an interface that explains nothing, holds a result and not the means of producing it. The distinction matters because the lock-in that survives the abolition of exit fees is of exactly this kind:

proprietary interfaces, vendor-specific formats and identity binding remain in force after a ban on charges [4]. What a society keeps, when it keeps anything, is the ability to rebuild the thing, and that ability lives in people and not in archives. Storage whose workings are published and taught leaves a population able to run its own, and storage whose workings are withheld leaves a population dependent on the operator's continued goodwill, however cheap the service.

Deduplication shows the difference. A provider that stores identical content once saves real money, and the saving can be kept as margin, passed on as a lower price, or explained to the public so that anyone can build the same system. The stored bytes are the same in all three cases. What differs is who can reproduce the efficiency afterward. Computation follows the same pattern. An assistant that supplies answers and an assistant that shows how its answers were reached may be equally accurate, and the population served by the second ends up more capable than the population served by the first. A baseline of free capacity should therefore be provided on terms that favour the second kind of system, with documented mechanisms, inspectable behaviour and interfaces that teach.

This is the sense in which the legibility condition is more than a matter of portability. Portability lets a person leave an operator, whereas legibility lets a public reproduce what the operator does, including the efficiencies it is paid for. In practice it would require published specifications of storage and replication design, documented deduplication and retention behaviour, and reporting of the savings that such techniques yield, so that the question of how savings enter the price can be answered by anyone and not only by the seller. Without such reporting, the customer's inability to assess the relation between logical volume and cost, noted in the discussion of copies, becomes permanent.

18 What Free Means

Free here means free at the point of use for a defined baseline, provided as infrastructure and recovered through taxation or levy, in the way public libraries lend books without a till or public roads carry traffic without a toll. It does not mean that no resources are consumed, or that anyone may consume without limit. The analogy to roads and libraries concerns why a good is provided in common and not how economists classify it. Storage and compute are consumed in aggregate, as roads are congested, and the case for common provision rests on the same ground: exclusion cuts a person off from ordinary life. The provision has four conditions, and a scheme lacking any of them does not qualify.

The first condition is a baseline allotment, large enough to hold a person's lifetime of documents, photographs, correspondence and work, and, separately, to run ordinary computation such as search, translation, and drafting, set as a quantity and not as a promotion. Since a lifetime of video can exceed any fixed figure, the baseline is defined by principle and then fixed as a number: sufficient for documents, correspondence and still images across a lifetime, with bulk video above the baseline unless the number is raised, and reviewed periodically as hardware and habits change. An illustrative storage figure is on the order of a terabyte. The compute baseline requires its own definition, since a terabyte says nothing about processing. It would specify a guaranteed share of processor and accelerator time, an availability standard, and a ceiling on waiting time for ordinary tasks. Scarce capacity can be allocated by fair scheduling, so that iteration is unmetered per interaction without implying unlimited simultaneous computation. The second is portability: open formats, published interfaces, no charge for export and an account identity the holder controls, so that the allowance can be moved to a different operator without penalty. The third is non-revocability within the baseline. A person's stored record cannot be deleted for non-payment, since there is no payment, and cannot be suspended at an operator's discretion; the sanction for abuse is a narrow one, applied through a procedure the holder can contest, and sanctions against an account preserve access to, or transfer of, its lawfully retained records. Unlawful material is handled by the existing apparatus of law, and the guarantee is not absolute in every circumstance. The fourth condition is legibility: the mechanisms by which the capacity is provided, and the efficiencies they yield, are published and documented, so that the allowance can be understood and rebuilt as well as used.

Delivery can take several forms, and the argument does not require one. Operators can be public bodies, cooperatives, or private firms paid by a public authority to supply the baseline on the four conditions, with competition among them. Above the allotment, supply is priced at cost, including energy, and the market continues to exist where the market is useful. What disappears is the layer that charges the public rent for the continued existence of its own memory.

Two bodies of work bear on the design. Studies of long-lived common resources find that they tend to be governed through clearly defined boundaries, monitoring accountable to the users and graduated sanctions [18], which matches the bounded allotment and the narrow, contestable sanction proposed here. And the digital preservation literature has long held that digital records are fragile because formats and media become obsolete, so that keeping a record means keeping the means of reading it [19], which is the ground of

the portability and legibility conditions.

19 Objections

The first objection is that nothing is free, and that someone must pay. This is true and is not an objection. Someone pays for roads and libraries, and the question is who and by what mechanism. The proposal is that the public pays in common, in proportion to means, through the same fiscal arrangements that provide other infrastructure, and not individually and perpetually at the seller's rate. That the payer is a taxpayer who did not choose each item is true of every common good, and the answer is the political one, a public decision on a defined and bounded capability with the levy designed in proportion to means; the essay does not design it. It offers no budget, and a serious one would have to be built from operating data that the large providers do not publish; the figures in section 6 show only that the hardware price of a baseline allotment is small beside the rent charged for it. At the illustrative figure given there, a terabyte of replicated hardware for each of a million people would cost between $72 million and $123 million before power, buildings, staff and replacement, which gives an order of magnitude and no more.

The second objection is abuse. A free allotment invites misuse, such as storage of illicit material, mining, resale of capacity and automated sign-up for multiple allotments. The answer is that a quota is not a price. A cap per holder and metering above it limit consumption, mining and resale, and the existing apparatus of law handles illicit content. The harder problem is the identity needed to enforce a per-person cap, which raises its own privacy costs. A household or cooperative account without real-name registration is one way to reduce them, and the problem is acknowledged here as open.

The third objection is that public provision will be poor, that public information systems have often been slow and badly designed. It is a fair description of some history. The conditions in section 18 are meant to make it survivable: with open formats and portable identity, a bad operator can be abandoned without loss, and competition among paid operators can be sustained. The concentration the proposal targets is held by the sellers who now define the formats.

A harder objection concerns procurement. A public authority buying capacity at scale may find that only the largest firms can supply it, and the baseline could end by routing public money to the same sellers. The four conditions are the protection, since portability and legibility let the authority change operators and let smaller operators and

cooperatives enter, but the risk is real and the essay does not claim to have removed it.

The fourth objection is that free baselines would reduce the incentive to build new services. The layers above the baseline remain open to commerce, and public provision of basic capacity leaves substantial scope for technical development, including storage reliability, compression, accessibility, and efficient computation, as well as for commercial services. Photographs, mail and search were built on top of storage and compute, and they do not become unbuildable when the underlying capacity is a utility. Nothing in the history of those services shows that they depended on the capacity beneath them being sold by the month. A bounded baseline leaves demand above it intact, and it widens the number of people able to build, since an allowance that does not lapse removes a cost that now stops small projects before they start.

The fifth objection is environmental. Free use without limit would waste energy, and the metering that now restrains use does at least send a price signal. The reply is that the baseline is bounded, that use above it is priced at cost including energy, and that a price set by what the market will bear tracks the seller's margin and not the energy drawn, so that it is a poor proxy for the cost to the grid. The figures needed to compare the two are not published, which is the limit of this reply. Above the baseline, energy can be priced directly, which is the better signal. Because the baseline is bounded per person, its energy draw can be planned for, whereas metered consumption has no ceiling other than what customers will pay.

The last concession belongs to the evidence, whose limits were stated in section 6. A fuller accounting could narrow the gap between rent and hardware, and the case would remain, because it rests as much on discretion, exit and dependence as on the arithmetic.

20 Purchasing Power and the Lifelong Record

The load-bearing claim of the proposal does not concern margins, and it can be stated without any of the pricing evidence. Suppose providers incur every cost they say they incur and earn only ordinary margins. A competitively priced and efficiently run service can still be an unsuitable foundation for a lifelong record, because continued access then depends on individual purchasing power, and the interruption of access can destroy the conditions of a person's intellectual and personal continuity. A record held on terms that must be met each month is a record that lapses whenever a person's circumstances change, whether through poverty, illness, bereavement, or a few years of attention elsewhere.

That is a reason for a society to fund a defined capability in common, in the way it

funds other things whose absence would cut a person off from ordinary life. The price evidence in the earlier sections does not establish the cost of an alternative system; it shows that hardware expenditure alone cannot explain or justify the existing terms of access. The normative argument shows that the arrangement would be unsuitable even if those terms were fully justified by cost.

21 Preservation as the Freedom to Defer Judgment

The opening storage warning can now be reconsidered. What appeared to be a small maintenance task asks someone to evaluate an accumulated life under constraints established elsewhere, using knowledge that may be unavailable until much later. The request treats the person's records as a stock to be trimmed, when for the person they are a set of open questions whose answers have not yet arrived. The instruction to make room is, in effect, an instruction to answer those questions prematurely.

The positive principle that organizes the critique is that an archive preserves our ability to discover what mattered after the circumstances of its recording have passed. This is the function that distinguishes an archive from a pile, and it is the function that justifies keeping material whose use cannot presently be named. It depends on retention, and it equally depends on retrievability, on the integrity of context, and on the practical possibility of moving the collection when its custodian changes. A service that offers to remember on behalf of its users should be evaluated by how well it sustains that possibility, including the labour required to retrieve, preserve, move, or relinquish their records. Judged in this way, the question for any such service is whether it leaves the person free to defer judgment, or whether it quietly requires the judgment to be made now, by the person, at the person's expense.

The proposal developed in the preceding sections follows from that standard. It is a baseline of storage and compute, free at the point of use, portable, not revocable, and legible, funded as the infrastructure it has become, with everything above the baseline sold at what it costs. What should not be sold, by the month and at the provider's discretion, is the permission to remember.

References

[1] Borges, Jorge Luis. "Funes el Memorioso." In *Ficciones*. Buenos Aires: Sur, 1944.

[2] Budd, John M., MaryEllen Sievert, and Tom R. Schultz. "Phenomena of Retraction: Reasons for Retraction and Citations to the Retracted Articles." *JAMA* 280, no. 3 (1998): 296–297.

[3] Bush, Vannevar. "As We May Think." *The Atlantic Monthly* 176, no. 1 (1945): 101–108.

[4] cloudmagazin.com. "EU Data Act: When Cloud Switching Fees Are Abolished, What CIOs Need to Examine." 6 July 2026. `cloudmagazin.com/en/2026/07/06/eu-data-act-when-cloud-switching-fees-are-abolished-what-cios-need-to-examine/`

[5] ComponentPrices. "Hard Drive Prices per TB." Price-table snapshot retrieved 8 November 2026. `https://componentprices.com/hard-drives/price-per-tb`. Tracks new desktop hard-drive listings at Newegg US; prices are in US dollars, excluding tax and shipping.

[6] ecorpit.com. "Cloud Egress Fees in 2026: What AWS, Azure, Google Cloud and Cloudflare R2 Actually Charge." `ecorpit.com/cloud-egress-fees-ai-inference-bill-comparison-2026/`. Egress rates verified June 2026; notes on the 2024 exit waivers.

[7] eRacks Open Source Systems. "Hard Drive Prices Up 40 to 90 Percent Since March: eRacks Publishes Its Own Six-Month Tracking." Press release, 14 September 2026. `pr.com/press-release/978414`. Vendor's own price tracking; its headline range does not match its table, which shows increases of 15 to 92 percent.

[8] Filebase. "AWS S3 Pricing in 2026: What You'll Actually Pay." `filebase.com/blog/aws-s3-pricing-in-2026-what-youll-actually-pay/`. Reports S3 Standard at $0.023 per GB-month for the first 50 TB in US East; the publisher sells competing S3-compatible storage.

[9] European Parliament and Council of the European Union. "Regulation (EU) 2016/679 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data (General Data Protection Regulation)." *Official Journal of the European Union* L 119 (2016), Articles 17 and 20; Recital 68.

[10] European Parliament and Council of the European Union. "Regulation (EU) 2023/2854 of 13 December 2023 on Harmonised Rules on Fair Access to and

Use of Data (Data Act).” Articles 2(36), 29 and 34; Recitals 88, 97 and 99. eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32023R2854

[11] Google. Announcement of Changes to Google Photos Storage, 11 November 2020. blog.google/products/photos/storage-changes/

[12] HDDHunt, a commercial price-comparison site. “HDD Price per TB, August 2026.” hddhunt.com/blog/hdd-price-per-tb-august-2026/. The page is titled for August 2026, but the price table it carries is a snapshot dated October 2026 of the cheapest new internal drives on Amazon US, roughly $24 to $41 per TB by capacity.

[13] Hirschman, Albert O. *Exit, Voice, and Loyalty: Responses to Decline in Firms, Organizations, and States*. Cambridge, MA: Harvard University Press, 1970.

[14] Mayer-Schönberger, Viktor. *Delete: The Virtue of Forgetting in the Digital Age*. Princeton, NJ: Princeton University Press, 2009.

[15] Nietzsche, Friedrich. *On the Genealogy of Morality*. Edited by Keith Ansell-Pearson, translated by Carol Diethe. Cambridge: Cambridge University Press, 2007. Originally published 1887.

[16] Nietzsche, Friedrich. “On the Uses and Disadvantages of History for Life.” In *Untimely Meditations*, translated by R. J. Hollingdale. Cambridge: Cambridge University Press, 1997. Originally published 1874.

[17] Official Gmail Blog. “Bringing It All Together: 15 GB Now Shared Between Drive, Gmail and Google+ Photos.” May 2013. gmail.googleblog.com/2013/05/bringing-it-all-together-15-gb-now.html

[18] Ostrom, Elinor. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge: Cambridge University Press, 1990.

[19] Rothenberg, Jeff. “Ensuring the Longevity of Digital Documents.” *Scientific American* 272, no. 1 (1995): 42–47.

[20] ScienceAlert. “When Gmail Was First Announced, People Thought It Was an April Fools’ Joke.” sciencealert.com/when-gmail-was-first-announced-people-thought-it-was-an-april-fools-joke. Gmail’s one gigabyte against two or four megabytes at Hotmail and Yahoo; the

article calls this about a hundred times, but the figures it gives imply a factor of several hundred.

[21] Shapiro, Carl, and Hal R. Varian. *Information Rules: A Strategic Guide to the Network Economy*. Boston: Harvard Business School Press, 1999.

[22] Sparrow, Betsy, Jenny Liu, and Daniel M. Wegner. "Google Effects on Memory: Cognitive Consequences of Having Information at Our Fingertips." *Science* 333, no. 6043 (2011): 776–778.

[23] Star, Susan Leigh, and Anselm Strauss. "Layers of Silence, Arenas of Voice: The Ecology of Visible and Invisible Work." *Computer Supported Cooperative Work* 8, no. 1–2 (1999): 9–30.

[24] Suchman, Lucy. "Making Work Visible." *Communications of the ACM* 38, no. 9 (1995): 56–64.

[25] Terranova, Tiziana. "Free Labor: Producing Culture for the Digital Economy." *Social Text* 18, no. 2 (2000): 33–58.