KUMA PUBLISHING
AI Infrastructure Valuation Research | Working Paper 2026-01

The Depreciation Gap

Economic Obsolescence of 2023–2024 Vintage AI Compute Hardware for Ad Valorem Property Tax Purposes: A Demonstration Analysis

Michael Halper · Kuma Publishing · August 2026 · info@kumapublishing.com

Draft for discussion. This is a research study applying a defined methodology to a hypothetical illustrative fleet. It is not a USPAP appraisal of any specific property, and it is not legal or tax advice. See Limitations and Disclosures.

Executive Summary

County assessors across the United States value business personal property using cost-based depreciation schedules that assume equipment loses value gradually as it ages and wears. AI compute hardware violates that assumption. A graphics processing unit (GPU) purchased in 2023 for frontier-class work is displaced not by wear but by architecture: each successive NVIDIA generation (Hopper to Blackwell to Rubin) resets the performance-per-dollar frontier, and the economic value of the prior generation adjusts to whatever the market will pay for its remaining best use. That adjustment is visible, measurable, and large — and standard percent-good tables do not capture it.

This paper tests that proposition against a representative fleet: 512 NVIDIA H100 SXM GPUs (64 HGX server nodes) acquired in the second half of 2023 at an installed cost basis of $16.0 million, valued at the January 1, 2026 lien date. Three independent lines of market evidence — secondary-market transaction pricing, rental-rate economics, and replacement-cost analysis against current-generation hardware — reconcile to a fair market value of roughly $6.0–7.5 million, approximately 38–47% of original cost after 2.3 years.

The finding is more interesting than a one-directional over-assessment story, and we report it as we found it. Where the fleet is classified under California's computer-specific valuation table (AH 581, Table 7), the published factor for 2023-vintage equipment is just 24–30% of cost — a schedule value of about $4.8 million, below our market reconciliation. California's computer schedule, built from mandated valuation studies, actually outpaces observed GPU depreciation. The excess-assessment exposure lives elsewhere, in two large and common situations: (1) classification disputes, where GPU clusters, their supporting infrastructure, or facility-side equipment are assessed under general commercial or industrial tables — or as fixtures — carrying factors of roughly 70% or more at the same age, implying $4–6 million of excess assessment on this single fleet; and (2) jurisdictions without a computer-specific schedule, where generic multi-year equipment lives produce comparable gaps. At typical effective rates, those scenarios represent roughly $50,000–$120,000 of recurring annual tax on a fleet this size, scaling linearly with asset base.

The conclusion this paper defends is therefore not that AI hardware is always over-assessed. It is that schedule-based valuation cannot track an asset class whose value is set by architecture cycles and market clearing prices rather than by age: California's aggressive computer table undershoots the market for this fleet, while generic-life tables overshoot it badly — and the classification decision that selects between those tables can swing taxable value by a factor of more than two on identical equipment. Taxpayers, assessors, and appeal boards need market-evidence-based valuation, in either direction, to reach full cash value. This paper demonstrates the methodology, which we call the Compute Depreciation Curve framework.

1. The Problem: Age-Based Schedules Meet Architecture-Based Value

Most states tax business personal property through a cost approach: the taxpayer reports original acquisition cost by year on an annual property statement, and the assessor applies published index and percent-good factors to estimate current value. In California, the State Board of Equalization publishes these factors annually in Assessors' Handbook Section 581 (AH 581) pursuant to Revenue and Taxation Code section 401.5, including specific valuation factor tables for non-production computers derived from valuation studies under section 401.20. Taxpayers report computer equipment on Schedule A of the Business Property Statement (Form BOE-571-L). Equivalent mechanisms exist in most personal-property states.

These schedules embed an assumption inherited from industrial equipment: value declines as a function of age and physical condition. For a lathe or a chiller, that assumption is roughly right. For AI accelerators it is structurally wrong, for three reasons:

- Value is set by displacement, not wear. An H100 in 2026 is physically identical to an H100 in 2023 — data-center GPUs commonly outlast several architecture generations — but its economic value repriced when Blackwell-class hardware changed the performance-per-dollar and performance-per-watt frontier. The asset's remaining physical life and its remaining economic value have decoupled.
- Repricing is discontinuous. Market observers describe GPU value behavior as regime shifts around generational launches and supply events rather than a smooth annual decline. A schedule that steps down a fixed percentage per year will be wrong in specific, demonstrable amounts in specific years — sometimes high, sometimes low.
- A deep, observable market now exists. Unlike most industrial equipment, AI compute has liquid price discovery: published secondary-market dealer pricing, daily rental-rate indices across dozens of providers, and, as of 2026, exchange-listed instruments referencing GPU compute prices. The evidence needed to test a schedule against market reality is publicly available.

Some jurisdictions have begun to respond. Utah's own fiscal analysis of data-center growth assumes $1.0 billion of GPU equipment falls to $620 million of taxable value after one year and roughly $70 million after five — an implicit admission that AI hardware depreciates on an economic curve far steeper than legacy schedules. Where a taxpayer's jurisdiction has not made a comparable adjustment, the difference between schedule value and market value is a live valuation dispute. Valuation practitioners already report material recoveries from related corrections: Kroll's published work on data-center taxation finds that unbundling non-taxable

soft costs and correcting classifications alone can reduce tax owed by 10–30%. Obsolescence is the larger, less-litigated half of the same problem.

2. The Illustrative Fleet

To make the analysis concrete, we define a hypothetical fleet representative of a mid-tier AI operator, a managed-cloud provider, or an enterprise on-premises deployment. All subsequent calculations refer to this fleet. The fleet is hypothetical; every valuation input applied to it is drawn from cited public market evidence.

Parameter	Assumption
Equipment	64 HGX H100 server nodes; 8× NVIDIA H100 SXM 80GB per node (512 GPUs total), incl. chassis, CPUs, memory, storage, InfiniBand networking
Acquisition period	Q3–Q4 2023 (reported cost year 2023)
Installed cost basis	$250,000 per node · $16.0M total (approx. $31,250 per GPU fully burdened)
Valuation date	January 1, 2026 lien date (equipment age ≈ 2.3 years)
Jurisdiction frame	California (AH 581 non-production computer factors; ~1.1% effective tax rate). Methodology generalizes to other percent-good states.

Table 1. Illustrative fleet definition. The per-GPU burdened cost is consistent with published 2023 H100 card pricing of roughly $25,000–$31,000 plus node-level infrastructure.

3. The Schedule Value: Three Very Different Answers

Under a standard cost-approach workflow, the assessor takes the reported 2023 cost of $16.0M and applies the published factor for the equipment's classification and acquisition year. Which factor applies turns out to matter far more than the equipment's age.

3.1 California's computer-specific table

For the 2026 lien date, AH 581 Table 7 (Non-Production Computer Valuation Factors) publishes, for 2023-acquisition-year equipment (age 3): 24% for personal computers and 30% for local area network equipment plus mainframe computers — the column under which enterprise GPU server nodes most plausibly report on Schedule A of Form BOE-571-L. Applied to the fleet:

Table 7 schedule value: **$16.0M × 30% = $4.8M (24% column: $3.8M)**

Note the finding, because it cuts against the naive version of this paper's thesis: $4.8M is below the market-evidence reconciliation developed in Sections 4–6 ($6.0–7.5M). California's computer factors, derived from mandated valuation studies under section 401.20, depreciate faster than the observed GPU market. A correctly classified H100 fleet in California is not over-assessed by this table; if anything, the schedule is favorable to the taxpayer. We report this plainly — a methodology that only ever finds over-assessment is advocacy, not valuation.

3.2 The classification doorway

Table 7's favorable factors apply only if the assessor accepts the equipment as non-production computers. That doorway is genuinely contested for AI infrastructure: a liquid-cooled, high-density GPU cluster with dedicated electrical and cooling plant invites arguments that some or all of the installation is (a) general commercial or industrial equipment rather than computers, or (b) fixtures constructively annexed to the realty. General equipment percent-good factors at age 3 commonly remain on the order of 70% or more of indexed cost, and fixture treatment can be worse. Under such a classification:

Illustrative misclassification value: **$16.0M × ~70–80% = $11.2M – $12.8M — roughly 2.5× the Table 7 figure on identical equipment**

The same doorway swings the treatment of node infrastructure, networking, and especially facility-side cooling and electrical equipment, where bundling into real-property improvements both inflates value and forfeits the computer schedule. This is the terrain where practitioners already document 10–30% recoveries from unbundling and reclassification alone.

3.3 Jurisdictions without a computer-specific schedule

California's Table 7 is the exception, not the national rule. Many personal-property jurisdictions value computer equipment under generic economic-life tables — frequently multi-year industrial lives with percent-good factors near or above 70% at age 3, and some with statutory floors. There, the schedule-versus-market gap is structural: a table built for equipment that wears cannot see equipment that gets displaced. The remainder of this paper develops the market side of that comparison; Section 6 quantifies all three scenarios side by side.

4. Market Evidence

4.1 Secondary-market transaction pricing (sales comparison indication)

A liquid resale market for H100-class hardware has developed through IT asset disposition firms, specialist brokers, and refurbishers. Published dealer and marketplace pricing through 2026 places used H100 boards broadly in the $15,000–$28,000 range depending on form factor (PCIe vs. SXM), condition, warranty, and configuration — roughly 60–70% of original card price at the favorable end, materially less at the volume end. For fleet-scale dispositions, three adjustments push realizable value toward the lower half of quoted retail ranges: (a) bulk-lot discounts versus single-unit dealer pricing; (b) the SXM form factor trades as complete nodes, and node-level infrastructure (chassis, CPUs, networking) recovers a much smaller fraction of its cost than the GPUs themselves; and (c) dealer ask prices embed refurbisher margin and warranty that a fleet seller does not receive.

Applying installed, fleet-scale resale economics of roughly $110,000–$140,000 per node (GPUs at wholesale-realizable levels plus modest recovery on node infrastructure):

Sales-comparison indication: **64 nodes × $110k–$140k = $7.0M – $9.0M (midpoint $8.0M)**

Standing alone, this indication overlaps the low end of the schedule range — which is precisely why a defensible study does not stop at listed resale prices. Retail ask prices are the ceiling of value, not its center, and the income evidence below explains why fleet buyers will not pay the ceiling.

4.2 Rental-rate economics (income indication)

The dominant commercial use of H100 fleets is renting compute. Rental rates are therefore direct evidence of the income the asset can command, and their trajectory since acquisition is the cleanest public record of economic obsolescence:

Period	H100 rental market	Source basis
2023 – early 2024	~$8/hr typical; major-cloud on-demand at times above $12/hr	Introl; Silicon Data retrospective
Mid-2025	AWS cuts H100 on-demand ~44%; market follows as Blackwell ships	CloudZero market analysis
Dec 2025	$2.85–$3.50/hr typical on-demand; ~64% below peak	Introl
Oct 2025 – Mar 2026	1-yr reserved contracts rebound ~40%: $1.70 → $2.35/hr	SemiAnalysis rental index
Q2–Q3 2026	On-demand spread ~$1.40–$11/hr across providers; cohort median ~$3.15/hr vs. >$7 in early 2024	Ornn AI OCPI; multi-provider index

Table 2. H100 rental-rate trajectory, acquisition to valuation date. Cited sources appear in Section 8.

Two facts matter for valuation. First, the level: at a mid-2026 achievable rate of roughly $2.00–$2.50/hr for committed fleet capacity, a GPU running at a strong 70% utilization generates roughly $12,000–$15,000 of gross annual revenue — before power, facilities, financing, and operations costs that industry analyses place near a $1.65/hr breakeven floor for providers. Capitalizing the thin residual economics of a two-generation-old asset with a limited remaining economic life supports a per-GPU value at or below the wholesale resale levels in 4.1, not above them. Second, the trajectory: the asset's income-generating power on the valuation date stands roughly 60–70% below where it stood at acquisition. No age-based schedule anticipated that path.

4.3 Replacement economics (cost-approach obsolescence adjustment)

The cost approach itself, correctly applied, requires deducting all forms of depreciation — physical, functional, and economic — from replacement cost new. The relevant replacement asset in 2026 is not a new H100; it is Blackwell-class hardware delivering substantially better performance per dollar and per watt. When a buyer can rent B200 capacity at $4.00–$5.50/hr or acquire current-generation hardware with a mature supply chain, the maximum they will rationally pay for prior-generation equipment is capped by the equivalent capacity cost of the replacement, discounted for the older asset's higher operating cost per unit of work and shorter remaining economic life. This functional-obsolescence deduction is exactly what generic percent-good tables omit, and it is the analytical core of the taxpayer's argument: the schedule

measures age; the statute demands full cash value; and full cash value is what the market — not the table — says it is.

5. The Honest Complication: Volatility Cuts Both Ways

A credible study must confront the counter-evidence, because opposing counsel and assessors' staff certainly will. Three items deserve direct treatment:

- Reserved-rate rebound. One-year H100 contract rates rose almost 40% from their October 2025 bottom through March 2026 as reserved capacity tightened. This is real, and it demonstrates that H100 economic value had not gone to zero — but a 40% rebound from a bottom roughly 75% below peak still leaves value far below the acquisition-era income level, and spot pricing subsequently softened again by mid-2026.
- Long service lives. Major clouds have historically run accelerators for many years (V100 instances ran roughly 7.5 years), and one operator has reported 2022-era H100 capacity re-booking near original contract pricing. These data points argue for remaining economic life; they do not restore acquisition-era value, and single-operator anecdotes are not market clearing prices.
- The symmetry obligation — demonstrated, not merely promised. Section 3.1 is this paper's own proof: California's Table 7 values the illustrative fleet below our market reconciliation, and we report that result even though it favors the assessor's table. If market evidence governs, it governs in both directions — in a supply-squeezed year, or under an aggressive schedule, market value can exceed the schedule figure. We regard this as a feature of the methodology, not a concession: the taxpayer's entitlement, and the assessor's duty, is full cash value, and a factor table is simply the wrong instrument for an asset repriced by architecture cycles and capacity markets. The appropriate response is annual market-evidence valuation on the lien date, in whichever direction it points — which is exactly what this framework performs.

6. Reconciliation and Indicated Value

The three market indications reconcile as follows:

Market value indication	Range	Weight
Sales comparison (fleet-scale resale)	$7.0M – $9.0M	Primary
Income (rental economics, remaining life)	$5.5M – $7.5M	Primary
Cost less all obsolescence (replacement-capped)	$6.0M – $8.0M	Supporting
Reconciled fair market value	**$6.0M – $7.5M (midpoint $6.75M ≈ 42% of cost)**	

Table 3. Reconciliation of market-evidence value indications for the illustrative fleet, January 1, 2026 lien date. Ranges reflect cited public evidence; an engagement-grade appraisal would tighten them with fleet-specific data.

Setting that market value against the three schedule scenarios from Section 3 quantifies where assessment risk actually lives:

Scenario	Schedule value	vs. market ($6.75M)	Annual tax effect*
CA Table 7, correctly classified (30%)	$4.8M	−$1.95M (below market)	Favorable to taxpayer
Misclassified: general equipment / fixtures (~70–80%)	$11.2M – $12.8M	+$4.5M – $6.0M excess	≈ $50k – $66k / yr
Generic-life jurisdiction, no computer table (~70%+)	$11.2M+	+$4.5M+ excess	≈ $45k – $120k / yr†

*Table 4. Schedule scenarios versus reconciled market value, illustrative fleet. *At a 1.1% effective rate. †Range reflects effective rates of roughly 1–2% across personal-property jurisdictions. Amounts recur each lien date and scale linearly: the same percentages on a $160M asset base imply ten times these figures.*

Two conclusions follow. First, in the misclassification and generic-schedule scenarios — the common cases for AI infrastructure nationally — the excess assessment on this single mid-sized fleet runs $4.5–6.0 million of value, roughly $50,000–$120,000 of recurring annual tax, with classification of facility-side equipment adding further exposure not modeled here. Second, the structural point: every input on the market side of the ledger is public, citable, and current, while the schedule side rests on tables that never observed a GPU transaction — and, as California's own Table 7 shows, even a well-researched table lands materially away from market in a given year. In a dispute over full cash value, that asymmetry of evidence is the taxpayer's case; where the table undershoots, it is the assessor's.

7. The Compute Depreciation Curve Framework

The methodology demonstrated above generalizes into a repeatable framework we apply to any AI-compute asset register:

- Step 1 — Asset decomposition. Separate the register into accelerators, node infrastructure, networking, and facility-side equipment; each class carries a different obsolescence profile and, in many jurisdictions, a different classification treatment (with real/personal classification itself a frequent source of over-assessment).
- Step 2 — Generational displacement mapping. For each accelerator vintage, identify the displacing architecture, its ship date, and its performance-per-dollar and performance-per-watt delta. Displacement events, not calendar years, are the framework's depreciation triggers.
- Step 3 — Three-market evidence assembly. On the valuation date, assemble (a) secondary-market transaction and dealer pricing at fleet scale, (b) rental-rate levels and utilization economics for the vintage, and (c) replacement-capacity economics for the displacing generation.
- Step 4 — Reconciliation to full cash value. Weight the indications, document the schedule gap, and quantify the excess assessment and recurring tax effect.
- Step 5 — Annual refresh. Because repricing is regime-driven, the analysis is re-run each lien date — in whichever direction the market moved.

The framework's output is designed to serve as evidence: an exhibit an appeal practitioner can attach, an assessor can test, and a board can weigh, with every figure traceable to a cited public source or a documented fleet-specific input.

8. Data Sources and Evidence Base

Primary public sources relied upon in this demonstration (full citations available on request; hyperlinked in the digital edition):

- California State Board of Equalization, Assessors' Handbook Section 581 (2026), Table 7, Non-Production Computer Valuation Factors, via the BOE Open Data Portal: 2026 lien date, 2023 acquisition year (age 3) — 24% (personal computers), 30% (LAN equipment plus mainframe computers); published under Rev. & Tax. Code §§ 401.5, 401.20; Form BOE-571-L, Schedule A.
- Introl, 'GPU Cloud Prices Collapse' (Dec. 2025): H100 on-demand $8/hr → $2.85–3.50/hr; ~$1.65/hr provider breakeven; ownership breakeven 14–16 months at full utilization.
- SemiAnalysis, H100 1-Year Rental Price Index (Apr. 2026): reserved rates $1.70/hr (Oct. 2025) → $2.35/hr (Mar. 2026).
- Ornn AI OCPI (May 2026, via Bloomberg Terminal listing): on-demand H100 spread $1.38–$11.01/hr; ICE GPU compute futures announcement.
- CloudZero, 'Nvidia H100 Price in 2026' (Aug. 2026): AWS mid-2025 ~44% H100 price cut; B200 rentals $3.99–$5.50/hr; ~$31,000 card pricing; CoreWeave re-booking anecdote; Azure V100 ~7.5-year service life.
- Multi-provider rental index (July 2026): 25-month monthly snapshots, 67 providers; H100 cohort median ~$3.15/hr vs. >$7 in early 2024.
- Published secondary-market dealer pricing for used H100 hardware, 2025–2026 ($15,000–$28,000 range by form factor and condition).
- Utah data-center fiscal analysis (Aug. 2026): $1.0B GPU equipment → $620M taxable value after one year, ~$70M after five.
- Kroll, 'Data Center Tax Strategies' (2025): unbundling and classification corrections saving 10–30% of tax owed.

9. Limitations and Disclosures

This paper is a methodological demonstration applied to a hypothetical fleet. It is not an appraisal performed under the Uniform Standards of Professional Appraisal Practice, not an opinion of value of any actual property, and not legal, tax, or accounting advice. Schedule factors shown are illustrative pending insertion of current published tables; fleet-specific engagements require actual asset registers, acquisition documentation, configuration detail, and jurisdiction-specific schedules and procedures. Market data are drawn from the cited public sources as of their publication dates and are subject to revision and to the volatility this paper itself documents. Property tax appeal deadlines are jurisdiction-specific and short; in California, assessment appeal applications are generally due September 15 or November 30 depending on

the county. Taxpayers should consult qualified property tax counsel or registered consultants regarding any specific matter.

About the author: Michael Halper is the founder of Kuma Publishing's AI Infrastructure Valuation Research practice, applying institutional valuation methodology — discounted cash flow analysis, market-evidence reconciliation, and a published intrinsic-value framework — to AI compute assets. National USPAP coursework in progress.

Inquiries: info@kumapublishing.com.