The Public Value AI Footprint

Assessing institutional AI use against public purpose

A framework for Washington's community and technical colleges

Summer Kenesson, August 2026

> *This paper represents the work of its author and is published to support discussion across the system. It does not state the policy or position of the State Board for Community and Technical Colleges, and it does not create any obligation for any college.*

Contents

Summary 2
Part 1: Core Principle 3
Part 2: Why Now 9
Part 3: Context for the Framework 11
Part 4: Background of the Framework 12
Part 5: The Framework 16
Part 6: Worked Examples 20
Part 7: The Washington Context 21
Part 8: An Inherent Conflict Considered 24
Part 9: Implementation 25
Part 10: Limitations and Open Questions 26
A Note on Method: This Paper Assessed by Its Own Framework 27
Part 11: Using and Adapting this Framework 28
Appendix: References and Further Reading 29

CC4.0

Summary

Colleges have adopted artificial intelligence faster than they have built any way to judge whether that adoption serves students. The tools now in common use such as acceptable use policies, privacy reviews, and procurement checklists answer one question well – *may we do this?* They are close to silent on a second question that public institutions cannot avoid – *was it worth what it cost, and who paid?*

This paper proposes a way to answer the second question. The Public Value AI Footprint is an assessment framework for institutions that buy and use AI rather than build it. It has three features that distinguish it from the risk frameworks already in wide circulation.

It cannot be offset. Six conditions function as gates. A use that fails one is stopped in its proposed form, regardless of how strong its benefit case is. This is a deliberate rejection of the logic that governs carbon offsetting, in which harm continues provided it is priced.

It checks where the time went. Nearly every case for adopting AI rests on staff time saved. Almost no institution verifies what happened to that time. Time redirected into work that serves students creates public value; time absorbed into backlog, or used quietly to justify a position left unfilled, does not. The framework requires someone other than the approver to check.

It treats the institution as a buyer with obligations. The environmental and equity literature on AI addresses developers, utilities, and legislatures. It rarely addresses the institution whose purchasing constitutes the demand. In Washington that gap is not abstract; the infrastructure serving this demand sits inside the service districts of colleges in this system, and the communities absorbing its water, energy, and land use include Tribal nations and rural communities with no voice in a college's procurement decision.

What this paper is not

- It is not a policy, although it could inform policy and practice at an institution, agency, or system level.
- It is not a position on data center development, state tax policy, or regional economic development. Section 6 addresses why that distinction matters.
- It is not tested. The framework has not yet been applied at scale

Key findings

1. Public value theory (Moore's strategic triangle) and a growing 2023–2026 literature on public-sector AI support PVAF's premise that legitimacy and authorization, not efficiency alone, govern whether a public institution should adopt AI.
2. Procurement is now widely recognized as the principal lever public agencies have over AI they do not build; this validates PVAF's buyer-as-unit-of-analysis framing and the "redirection audit."
3. Non-compensatory (veto/gate) structures are established in sustainability and multi-criteria assessment; the ethics literature on biodiversity and carbon offsetting supplies a principled rejection of offsetting logic that PVAF embodies.

CC4.0

4. AI's water, energy, and land costs fall disproportionately on Columbia Basin tribes and rural communities inside the SBCTC service area; a live, documented environmental-justice harm.
5. Generative AI systematically renders Native peoples as a single generic culture and reproduces appropriated designs; existing law (IACA 1990) does not reach AI outputs.
6. Indigenous-led AI scholarship and data-sovereignty frameworks (CARE, OCAP, FPIC, Local Contexts/TK Labels) provide ready-made, Indigenous-authored authorities PVAF imports rather than reinvents.

Part 1: Core Principle

1.1 Public Value Theory

Mark H. Moore's *Creating Public Value: Strategic Management in Government* (Harvard University Press, 1995) established the "strategic triangle": a public undertaking must be simultaneously (a) substantively valuable, (b) legitimate and politically sustainable within its authorizing environment, and (c) operationally feasible; if any part is missing, the endeavor fails. Moore's dictum that "successful public managers are explorers commissioned by society to search for public value" reframes an agency's AI decisions as acts of stewardship of public assets (public money and the authority of the state) not merely IT procurement.[1]

A 2023–2026 wave of public-administration scholarship applies this lens directly to institutional AI adoption. Oostvogel et al., "Getting the Priorities Straight: Public Values in AI Adoption" (*Public Administration*, 2025), find that value-setting and visionary leadership, not just top-down efficiency, drive successful AI adoption in hierarchical public organizations.[2] A *Sustainability* article, "Artificial Intelligence and Public Values: Value Impacts and Governance in the Public Sector" (MDPI, 2023), catalogs the public values — efficiency, effectiveness, accountability, transparency, equity — against which public-sector AI must be judged.[3] A cross-national study of 28 national AI strategies (*Government Information Quarterly*, 2024) warns that "viewing AI only as automation may constrain other innovative solutions for public value," and a book chapter in the *Research Handbook on Public Management and Artificial Intelligence* documents "public value destruction" when agencies pursue efficiency values

[1] Moore, Mark H. *Creating Public Value: Strategic Management in Government.* Cambridge, MA: Harvard University Press, 1995. See also Moore, *Recognizing Public Value* (Harvard University Press, 2013); Harvard Kennedy School / Bloomberg Center for Cities, *Public Value Tool Kit*.

[2] Oostvogel, et al. "Getting the Priorities Straight: Public Values in AI Adoption." *Public Administration* (2025). Wiley Online Library, doi:10.1111/padm.70078.

[3] "Artificial Intelligence and Public Values: Value Impacts and Governance in the Public Sector." *Sustainability* 15(6): 4796 (2023). MDPI, mdpi.com/2071-1050/15/6/4796.

CC4.0

alone.[4,5] These sources converge on PVAF's core claim: efficiency public values, if pursued in isolation, can destroy other public values — which is precisely why offsetting logic must be rejected.

1.2 Demand-side / procurement-based governance

A distinct 2024–2026 literature argues procurement is the primary governance lever for institutions that *buy* rather than *build* AI. Tech Policy Press ("Want Accountable AI in Government? Start with Procurement," 2025) states that "in the absence of federal regulation of AI vendors, procurement remains one of the few levers governments have to push for public values, such as safety, non-discrimination, privacy, and accountability."[6] The OECD's *Governing with Artificial Intelligence* (2024) treats public contracts as instruments to align AI acquisition with trustworthy-AI principles, and the World Economic Forum's *AI Procurement in a Box* (2020) frames procurement as a route to demand vendor transparency.[7,8] An arXiv working paper, "AI Governance through Markets" (2025), describes procurement standards as "bottom-up market governance mechanisms that effectively regulate industries," and a *Laws* article (2025) consolidates the view of "public AI contracts not only as objects of procurement but also as instruments of regulation."[9,10] This body of work directly validates PVAF's buyer-as-unit-of-analysis choice: the institutional buyer, exercising purchasing discretion, is the accountable actor.

1.3 Non-compensatory / gate-based assessment

PVAF's design does not let strong scores offset disqualifying findings, and has firm methodological grounding. In multi-criteria decision analysis, non-compensatory methods (limiting-factor, veto/outranking rules) "treat planetary boundaries as hard constraints" rather than tradeoffs, unlike compensatory averages that "allow high income to offset near-zero environmental performance" (arXiv, "A Comparative Review of Methods to Create a Composite Index for Sustainable and Inclusive Wellbeing," 2026).[11] The ethics-of-offsetting literature supplies the moral argument PVAF makes when it rejects carbon-offsetting logic: Karlsson and Edvardsson Björnberg, "Ethics and biodiversity offsetting"

---

[4] "The strategic use of AI in the public sector: A public values analysis of national AI strategies." *Government Information Quarterly* (2024). ScienceDirect.

[5] "Analysis of driving public values of AI initiatives in government in Europe," Chapter 13 in *Research Handbook on Public Management and Artificial Intelligence.* Edward Elgar.

[6] "Want Accountable AI in Government? Start with Procurement." *Tech Policy Press*, 2025. See also "AI Procurement and the Capture of Public Purpose," *Tech Policy Press* (techpolicy.press).

[7] OECD. *Governing with Artificial Intelligence* — "AI in public procurement." OECD Publishing, 2024.

[8] World Economic Forum. *AI Procurement in a Box: AI Government Procurement Guidelines.* WEF, 2020.

[9] "AI Governance through Markets." arXiv:2501.17755 (2025), arxiv.org/abs/2501.17755.

[10] "Transformative Public Procurement of Artificial Intelligence." *Laws* 14(6): 97 (2025), doi:10.3390/laws14060097.

[11] "A Comparative Review of Methods to Create a Composite Index for Sustainable and Inclusive Wellbeing." arXiv:2607.08153 (2026), arxiv.org/abs/2607.08153.

CC4.0

(*Conservation Biology*, 2020), catalog objections including that "losses of nature cannot be compensated for by human interventions" and that offsetting has "negative justice implications"; Ives and Bekessy, "The ethics of offsetting nature" (*Frontiers in Ecology and the Environment*, 2015), demonstrate that offsetting "aligns most easily with a utilitarian ethic, where outcomes rather than actions are the focus."[12,13] PVAF's non-compensatory gates are thus a defensible deontological stance — certain harms are not tradeable — not an idiosyncrasy.

1.4 Environmental and community costs borne by the public

AI infrastructure imposes measurable water, energy, and land costs. Per Li et al. ("Making AI Less Thirsty," *Communications of the ACM* 2025; arXiv:2304.03271), "training GPT-3 in Microsoft's U.S. datacenters can consume a total of 5.4 million liters of water, including 700,000 liters of scope-1 onsite water consumption," and GPT-3 "drinks" roughly a 500 ml bottle of water per 10–50 medium-length responses; the same authors project global AI water withdrawal of 4.2–6.6 billion cubic meters annually by 2027.[14] On energy, the Lawrence Berkeley National Laboratory *United States Data Center Energy Usage Report* found U.S. data centers used 176 TWh (about 4.4% of national electricity) in 2023, rising to roughly 192 TWh (about 4.7%) in 2024, and projected 6.7%–12% by 2028.[15]

In the Pacific Northwest, these costs land on communities inside the SBCTC service area. The Northwest Power and Conservation Council modeled that Columbia Basin data-center growth could strain regional power supply within five years (*Seattle Times*, 2025).[16] The Yakama Nation's *Data Center Findings and Recommendations* (November 2025) documents that basin dam construction already caused the loss of over 80% of salmon and steelhead runs (annual returns falling from an estimated 10–16 million at treaty time to roughly 2 million today) and warns that new data-center demand will intensify pressure to run fossil generation, imposing costs on tribal and other consumers.[17] Columbia Riverkeeper's February 2026 report aggregates data-center impacts across 12 river counties in Washington and Oregon and calls for "avoiding general state permitting that sidesteps much-needed Tribal government and

[12] Karlsson, Mikael, and Karin Edvardsson Björnberg. "Ethics and biodiversity offsetting." *Conservation Biology* (2020), doi:10.1111/cobi.13603.

[13] Ives, Christopher D., and Sarah A. Bekessy. "The ethics of offsetting nature." *Frontiers in Ecology and the Environment* 13 (2015): 568–573, doi:10.1890/150021.

[14] Li, Pengfei, et al. "Making AI Less 'Thirsty': Uncovering and Addressing the Secret Water Footprint of AI Models." *Communications of the ACM* (2025); arXiv:2304.03271 (2023), arxiv.org/abs/2304.03271.

[15] Shehabi, Arman, et al. *United States Data Center Energy Usage Report* (2024 edition, LBNL-2001637; 2025 Update). Lawrence Berkeley National Laboratory.

[16] "PNW data center boom could imperil power supply within 5 years." *The Seattle Times*, 2025.

[17] Yakama Nation. *Data Center Findings and Recommendations* (to the Washington Data Center Workgroup). November 10, 2025.

CC4.0

community engagement."[18] This is precisely the redirection of public goods (water, power, ratepayer costs, sacred land) toward private AI infrastructure that PVAF's redirection audit is built to surface.

1.5 Washington's environmental-justice machinery

The HEAL Act (Healthy Environment for All Act, 2021 c 314; codified at RCW ch. 70A.02) is, per the Washington Attorney General's fact sheet, "the first law in Washington to create a coordinated approach to EJ across state government."[19] It defines "environmental justice" as "the fair treatment and meaningful involvement of all people regardless of race, color, national origin, or income with respect to the development, implementation, and enforcement of environmental laws, rules, and policies," including "addressing disproportionate environmental and health impacts… by prioritizing vulnerable populations and overburdened communities" (RCW 70A.02.010(8)), and defines "overburdened community" as "a geographic area where vulnerable populations face combined, multiple environmental harms and health impacts" (RCW 70A.02.010(11)).[20] It requires covered agencies to "develop a consultation framework in coordination with tribal governments" and to "offer consultation with federally recognized Indian tribes" on strategic plans, community-engagement plans, and significant agency actions affecting tribal lands — "independent of any public participation process" and "regardless of whether the agency receives a request for consultation from an Indian tribe" (RCW 70A.02.100).[21] It also establishes an Environmental Justice Council with dedicated tribal seats (RCW 70A.02.110).[22] SBCTC is not one of the seven core "covered agencies" (Ecology, Health, Natural Resources, Commerce, Agriculture, Transportation, and the Puget Sound Partnership; see note 20), but the HEAL Act supplies the state's authoritative definitions and consultation model that PVAF can adopt by reference.

1.6 Cultural appropriation in AI

The pan-Indigenous homogenization problem

Generative image models collapse hundreds of distinct nations into a single generic "Native" aesthetic. There were 574 federally recognized tribes as of the Bureau of Indian Affairs' August 11, 2023 *Federal Register* notice (88 FR 54654) — 347 in the contiguous 48 states plus 227 in Alaska; the count rose to

[18] Columbia Riverkeeper. "New Report Highlights Data Center Impacts Along the Columbia." February 26, 2026 (columbiariverkeeper.org). See also "Wash. Data Center Workgroup won't recommend clean energy requirement," *Columbia Insight*.

[19] Washington State Office of the Attorney General. "SB 5141: The Healthy Environment for All ('HEAL') Act" (fact sheet), agportal-s3bucket.s3.amazonaws.com — HEAL Act Fact Sheet (PDF); "About the HEAL Act."

[20] Revised Code of Washington, ch. 70A.02 RCW (Environmental Justice), §§ 70A.02.005, 70A.02.010, 70A.02.020 (2021 c 314). Washington State Legislature, app.leg.wa.gov/RCW/default.aspx?cite=70A.02.

[21] RCW 70A.02.100 (Tribal consultation), app.leg.wa.gov/RCW/default.aspx?cite=70A.02.100.

[22] RCW 70A.02.110 (Environmental Justice Council), app.leg.wa.gov/RCW/default.aspx?cite=70A.02.110.

CC4.0

575 on the BIA's December 18, 2025 list — and each is culturally distinct.[23] As one Native photographer told *Native Max*, "AI is just melding everything it learns about Native Americans all into one image… You don't get the kind of concrete differences that you do with the [574] federally recognized tribes."[24]

Victoria Turk's *Rest of World* investigation (October 10, 2023), analyzing 3,000 Midjourney v5.2 images (100 per prompt), found that for "an Indian person," "at least two appeared to wear Native American-style feathered headdresses, indicating some ambiguity around the term 'Indian'… A couple of the images seemed to merge elements of Indian and Native American culture," documenting how models default to stereotyped, conflated renderings.[25] Arizona State University's Future of Being Human project documented Midjourney rendering the real Navajo Nation town of Tuba City as a "run down shanty town… nothing like the real Tuba City."[26] The recurring visual tropes — Plains war bonnets, buckskin, the "stoic warrior" and "vanishing Indian" — reproduce exactly the "pan-Indian" stereotypes that Duke University's library exhibit *American Indians Go Graphic* traces to a century of comic-book and film caricature.[27]

Appropriation of specific designs and sacred imagery

Northwest Coast **formline** — the ovoid/U-form/S-form system codified in Bill Holm's *Northwest Coast Indian Art: An Analysis of Form* (1965) and, in Indigenous terms, described by Rico Lanáat' Worl (Tlingit/Athabascan) as "part of a living culture… It surrounds us, and it holds us up" — is a "traditional cultural expression" governed by complex, esteemed rules (see note 28). In an open letter published by *First American Art Magazine* (November 2021), Lucinda Turner (who carved for 26 years in partnership with Nisga'a master artist Norman Tait) wrote that non-Indigenous production of "'Native-inspired' designs using Formline… hurts immensely to see this exquisite art form degraded into comical and disrespectful caricatures of itself," and that "without stronger Federal legislation in place, it is practically impossible to stop the marketing and sale of these types of uninformed and inauthentic copies that misrepresent this art form worldwide."[28] Although this letter predates the generative-AI wave and addresses appropriation broadly, it is the strongest documented account of formline misappropriation,

[23] U.S. Bureau of Indian Affairs. "Indian Entities Recognized by and Eligible To Receive Services From the United States Bureau of Indian Affairs," 88 Fed. Reg. 54654 (Aug. 11, 2023), federalregister.gov/citation/88-FR-54654; updated list Dec. 18, 2025.

[24] "Native Fashion in AI." *Native Max Magazine*, nativemaxmagazine.com/native-fashion-history-identity-in-ai.

[25] Turk, Victoria. "How AI reduces the world to stereotypes." *Rest of World*, October 10, 2023, restofworld.org/2023/ai-image-stereotypes.

[26] "Advanced AI is creating challenges for Native American communities – but also opportunities." Future of Being Human, Arizona State University.

[27] *American Indians Go Graphic* (exhibit). Duke University Libraries, Rubenstein Library.

[28] Turner, Lucinda. "Northwest Coast Indigenous Art: Theft and Appropriation" (open letter, Nov. 2021). *First American Art Magazine* (firstamericanartmagazine.com). Formline definitions per Bill Holm, *Northwest Coast Indian Art: An Analysis of Form* (University of Washington Press, 1965), and Rico Lanáat' Worl, "A Basic Guide to Northwest Coast Formline Art."

CC4.0

and its warning applies directly to AI systems that can now mass-produce "Native-style" designs on demand.

Diné/Hopi scholar Dr. Angelo Baca (Rhode Island School of Design) addresses Navajo/Diné intellectual property and AI directly in his TEDxRISD talk "Fighting cultural appropriation in the age of AI."[29] A February 2026 *Native America Calling* panel (via Indianz.Com), featuring Baca, law professor Trevor Reed (Hopi, UC Irvine School of Law and associate justice of the Hopi Tribe Court of Appeals), and Dr. Tamika Worrell (Gamilaroi, Macquarie University), warned that AI-generated images "present vaguely Native visual and audio characteristics with little to no authentic cultural connections," eroding decades of work to counter stereotypes.[30]

Training data, economic harm, and the IACA gap

These outputs are possible because models scrape Indigenous cultural materials without attribution or permission. Economic harm to Native artists follows. The Indian Arts and Crafts Act of 1990 (P.L. 101-644) is a truth-in-advertising statute prohibiting the sale or display of goods that falsely "suggest" they are Indian-produced, with penalties up to $250,000 / 5 years for an individual and $1,000,000 for a business.[31] But the IACA reaches the *marketing of goods represented as made by a tribal member*; it does not squarely reach AI-generated "Native-style" imagery produced by no identifiable person and not marketed as tribally made — an analytic gap (not a settled judicial holding) that leaves Native artists economically exposed and consumers deceived.[32]

Indigenous-led responses

The Indigenous Protocol and Artificial Intelligence Working Group *Position Paper* (edited by Jason Edward Lewis; Initiative for Indigenous Futures and CIFAR, 2020) is "a starting place for those who want to design and create AI from an ethical position that centres Indigenous perspectives," authored by a multinational group spanning Hawaiian, Māori, Anishinaabe, Cree, Lakota, Mohawk and other communities.[33] The foundational essay "Making Kin with the Machines" (Lewis, Arista, Pechawis & Kite, *Journal of Design and Science*, 2018) anchors the field now called Indigenous AI.[34] Abundant Intelligences (co-directed by Lewis) is "an Indigenous-led research program that reimagines how to

[29] Baca, Angelo (Diné/Hopi). "Fighting cultural appropriation in the age of AI." TEDxRISD (Rhode Island School of Design), YouTube.

[30] "Native America Calling: The growing artificial intelligence threat." Indianz.Com, February 19, 2026 (panelists Angelo Baca, Trevor Reed, Tamika Worrell), indianz.com.

[31] Indian Arts and Crafts Act of 1990, Pub. L. No. 101-644, 104 Stat. 4662. U.S. Department of the Interior, Indian Arts and Crafts Board, "The Indian Arts and Crafts Act of 1990," doi.gov/iacb.

[32] "The Indian Arts and Crafts Act of 1990: Fit for Purpose?" Center for Art Law (itsartlaw.org).

[33] Lewis, Jason Edward, ed. *Indigenous Protocol and Artificial Intelligence Position Paper.* Indigenous Protocol and AI Working Group / Initiative for Indigenous Futures and CIFAR, Honolulu, 2020, indigenous-ai.net/position-paper.

[34] Lewis, Jason Edward, Noelani Arista, Archer Pechawis, and Suzanne Kite. "Making Kin with the Machines." *Journal of Design and Science*, MIT, 2018, jods.mitpress.mit.edu/pub/lewis-arista-pechawis-kite.

CC4.0

conceptualize and design AI based on Indigenous Knowledge systems," and the Wíhaŋble S'a Center for Indigenous AI at Bard College — led by Dr. Suzanne Kite (Oglala Lakota) and funded by a $500,000 NEH grant in 2024 — models Indigenous-governed AI research.[35,36] The Lakota AI Code Camp (co-founded by Mason Grimshaw, Sicangu-Rosebud Lakota, and Michael Running Wolf) builds Indigenous AI capacity and Indigenous-language work (the FLAIR / First Languages AI Reality initiative).[37] These bodies are the appropriate authorities and partners for a public institution assessing appropriation risk.

Part 2: Why Now

Three developments over the past eighteen months have moved this from a question institutions could defer to one they will be asked to answer.

2.1 Adoption has outrun policy

EDUCAUSE research published in January 2026, produced with AIR, NACUBO, and CUPA-HR and drawing on a survey of nearly two thousand staff, administrators, and faculty across more than eighteen hundred institutions, found that 94% of higher education workers had used AI tools for work in the preceding six months while only 54% knew their institution's AI policies.[38] Separately, EDUCAUSE found institutional AI acceptable use policies rose from 23% in 2024 to 39% in 2025, and that a majority of staff have used AI tools their institution did not provide and has not evaluated for privacy, security, or compliance.[39]

The gap is not primarily one of policy absence. It is that the policies that exist address permission and say almost nothing about warrant.

2.2 Governance is becoming a compliance and funding question

The regulatory picture firmed considerably during 2026. Transparency obligations under the EU AI Act requiring disclosure when users interact with AI systems took effect in August 2026, and the Act treats AI used in admissions screening, student assessment, and performance monitoring as high-risk, requiring

[35] Abundant Intelligences (Indigenous-led research program; co-directed by Jason Edward Lewis), abundant-intelligences.net.

[36] Wíhaŋble S'a Center for Indigenous AI, Bard College (Dir. Dr. Suzanne Kite, Oglala Lakota); NEH-designated AI research center, $500,000 grant, 2024. See also "Lakota-Led AI Research Explores Ethics Through Indigenous Knowledge."

[37] "Indigenizing artificial intelligence and programming." ICT News; "Lakota Youth Use Coding and AI to Save Language," *Lakota Times* (Lakota AI Code Camp; co-founders Mason Grimshaw and Michael Running Wolf; FLAIR).

[38] EDUCAUSE. The Impact of AI on Work in Higher Education (with AIR, NACUBO, and CUPA-HR), January 12, 2026, educause.edu/research/2026. The survey drew nearly 2,000 staff, administrator, and faculty respondents across more than 1,800 institutions. See also "Data Shows AI 'Disconnect' in Higher Ed Workforce," Inside Higher Ed, January 13, 2026.

[39] EDUCAUSE. 2025 AI Landscape Study; EDUCAUSE, The Impact of AI on Work in Higher Education (January 2026).

CC4.0

conformity assessment, bias testing, human oversight, and auditability. Institutions enrolling students from the EU may fall within scope regardless of location.[40]

Domestically, federal guidance has indicated that grant funds may support AI in instruction, student support, and workforce programs, but conditioned on civil rights and privacy compliance and on humans remaining accountable for high-stakes decisions.[41] Accreditors have begun signaling expectations around transparency, accountability, and bias. The practical effect is that undocumented AI adoption is becoming an accreditation and funding exposure rather than only an ethical one.

2.3 The infrastructure question arrived in Washington's own communities

This is the development most specific to this system, and Section 6 treats it in detail. In brief: during 2026 the Legislature considered comprehensive data center regulation and did not pass it; a state Data Center Workgroup convened by the Governor is expected to issue a final report that will inform the 2027 session; and local governments across Washington — including Moses Lake, where a college in this system is located — adopted moratoriums on new data center development while they study the question.[42]

2.4 Consent and Indigenous data sovereignty

Knowledge extraction from Native peoples without consultation or consent has a long arc — anthropological and museum collection, colonial consumption of cultural objects, boarding-school-era collections, and research abuses. The paradigmatic case is *Havasupai Tribe v. Arizona Board of Regents*: between 1990 and 1994, ASU researchers Therese Markow and John Martin collected DNA from approximately 400 Havasupai members for diabetes research, then used the samples for studies of schizophrenia, migration, and inbreeding that the tribe never authorized. The April 20, 2010 settlement returned the samples and paid $700,000 to 41 tribal members; as Mello and Wolf reported in the *New England Journal of Medicine* (June 9, 2010), ASU "agreed to pay $700,000 to 41 members of the Havasupai Indian tribe to settle legal claims that university researchers improperly used tribe members' blood samples in genetic research."[43] The Native American Graves Protection and Repatriation Act (NAGPRA, 1990) is a partial corrective — the National Park Service reports almost 4,000 published

[40] Regulation (EU) 2024/1689 (Artificial Intelligence Act), Article 50 (transparency obligations, applicable August 2, 2026) and Annex III(3) (education and vocational training uses classified as high-risk, including admission, evaluation of learning outcomes, and monitoring), eur-lex.europa.eu/eli/reg/2024/1689/oj.

[41] U.S. Department of Education, guidance on the use of federal grant funds for artificial intelligence in instruction and student supports, and the Department's proposed supplemental priority on advancing AI in education (July 2025); see also Exec. Order No. 14277, Advancing Artificial Intelligence Education for American Youth (Apr. 23, 2025).

[42] Municipal Research and Services Center (MRSC). "What the Data Center Boom Means for Washington Local Governments," August 2026, mrsc.org/stay-informed/mrsc-insight/august-2026/data-centers; Washington Governor's Data Center Workgroup (Exec. Order 25-05), preliminary report, December 2025 (final report pending).

[43] Mello, Michelle M., and Leslie E. Wolf. "The Havasupai Indian Tribe Case — Lessons for Research Involving Stored Biologic Samples." *New England Journal of Medicine* (June 9, 2010), doi:10.1056/NEJMp1005203. See also *Havasupai Tribe v. Arizona Board of Regents* (settled April 20, 2010).

CC4.0

notices enabling repatriation of over 84,000 ancestral remains and over 1.5 million funerary objects — yet testimony indicates over 117,000 ancestral remains still remain in collections, the majority not yet culturally affiliated.[44,45] AI training-data extraction constitutes a new layer of this same extractive logic: intellectual and informational appropriation that omits Tribal consent and ethical collaboration.

> *A system that consumes AI while its colleges sit in communities pausing to study what that consumption requires is in an unusual position. It can either develop a considered account of its own demand or wait to be asked for one.*

Part 3: Context for the Framework

Institutional AI governance has converged on a recognizable form. In place are components such as an acceptable use policy defining prohibited categories, a data classification scheme, a procurement security review, a privacy review, and where an institution is mature, a committee that reviews exceptions.

This apparatus answers whether a use is *allowed*. It does not address whether it was *warranted*. The distinction matters because the two have different failure modes. A permission framework fails when it authorizes something harmful. A justification framework fails when it authorizes something merely permissible, that is, it is lawful, private, secure, but pointless. The second failure is more common and almost entirely unmeasured.

3.1 Three reasons the drift is hard to see from inside

The benefit is near; the cost is far

An institution experiences AI's benefits directly such as a report drafted faster or a document translated. It experiences the costs as an invoice. The energy, water, land, materials, and labor consumed are borne by parties with no relationship to the institution and no voice in its decisions. Ordinary institutional feedback cannot register that cost.

The efficiency claim checks itself

When a unit reports that a tool saved eleven hours, no one audits the claim, and no one audits its consequence. Eleven hours saved is an input. Whether it became public value depends on what happened to those hours, which is a question outside the boundary of almost every efficiency instrument in ordinary use. The central justification offered for AI adoption is the one element subject to virtually no verification at all.

---

[44] National Park Service. "Native American Graves Protection and Repatriation Act of 1990" (NAGPRA overview), nps.gov/subjects/nagpra. See also U.S. Department of the Interior, testimony on NAGPRA implementation.

[45] U.S. Department of the Interior, Office of Congressional and Legislative Affairs, "Native American Graves Protection" (statement noting over 117,000 ancestral remains still in collections).

CC4.0

Capability erosion takes years and reporting takes quarters

Whether AI builds or depletes professional expertise has a multi-year answer and a quarterly reporting cycle. In the short run, delegation and development look identical because the work gets done. The difference appears when the institution needs the underlying judgment and finds it has not been maintained. Colleges and state agencies are exposed here, because accumulated professional judgment is much of what they employ people for.

Part 4: Background of the Framework

Frameworks are easier to publish than to justify. This section states plainly which parts of this one are built upon of existing work, creating a foundation for what is new.

4.1 Existing designs and concepts and how they inform the framework

Element	Status in the literature	Treatment here
Public value applied to public-sector AI	Established. Mark Moore's public value theory; a 2023 systematic review of public-sector AI and public values with a national survey of U.S. government employees; ongoing work in public administration journals.	Cited and built upon.
Net benefit-versus-cost assessment for AI	Established. Net Impact Assessment methods compare ecological costs against attributable benefits, corrected for rebound effects and additionality, using counterfactual reasoning. Sustainable AI impact assessment makes a parallel argument.	Counterfactual and additionality tests adopted directly.
Risk, privacy, fairness, transparency	Comprehensively covered. NIST AI Risk Management Framework, OECD AI Principles, UNESCO recommendation, EU AI Act.	Imported by reference.
Higher education AI governance	Active and crowded. EDUCAUSE landscape and workforce studies, institutional policy templates, system-level GenAI policy work already underway in Washington.	Cited. Framework positioned as a complement.
Environmental justice of AI infrastructure	Established and growing, much of it Indigenous-led.	Cited. Framework positioned as the demand-side complement.

CC4.0

Element	Status in the literature	Treatment here
Indigenous data governance	Established and Indigenous-led. CARE Principles (Global Indigenous Data Alliance, 2019; Carroll et al., 2020), building on OCAP® (First Nations Information Governance Centre, 1998) and FPIC under UNDRIP (2007); stewarded by the Indigenous data sovereignty networks in 4.2 (USIDSN, Te Mana Raraunga, Maiam nayri Wingara, GIDA-Sápmi).	Adopted directly, not paraphrased. Operationalized in gate G6.
Indigenous cultural expression and AI appropriation	Established and Indigenous-led. Indigenous Protocol and AI Working Group position paper (Lewis, ed., 2020); documented pan-Indigenous stereotyping in generative image models; the Indian Arts and Crafts Act's analytic gap. See 1.6.	Cited. Operationalized in gate G5.
Consent-based cultural licensing tooling	Established. Traditional Knowledge (TK) Labels / Local Contexts (Anderson and Christen, 2010); Mukurtu CMS (Washington State University). See 4.2.	Adopted by reference; routed through gates G5 and G6.
Verifying where released capacity went	No operational instrument identified in this review.	Developed here.
The institutional buyer as accountable actor	Underdeveloped.	Developed here.
Non-compensatory gate structure	Uncommon in this domain.	Adopted as a design commitment.

This review was conducted in mid-2026 and is not exhaustive. Readers who know of prior work operationalizing either of the two elements claimed as new are asked to send it to the contact address; the paper will be corrected.

4.2 Indigenous authorities applied

Consistent with PVAF's design, the framework cites these Indigenous-authored authorities directly rather than restating them:

- **CARE Principles for Indigenous Data Governance (Collective benefit, Authority to control, Responsibility, Ethics)** Carroll et al., *Data Science Journal* (2020); released by the Global Indigenous Data Alliance (GIDA) in September 2019 as a "people and purpose-oriented" complement to the FAIR principles.[46] GIDA's 2025 communiqué, *CARE Directs Us Home: Prioritizing Indigenous Peoples'*

[46] Carroll, Stephanie Russo, et al. "The CARE Principles for Indigenous Data Governance." *Data Science Journal* 19:43 (2020), doi:10.5334/dsj-2020-043. Global Indigenous Data Alliance (GIDA), September 2019, gida-global.org/care.

CC4.0

Community Standards, insists that "there is no justice in Indigenous data without Indigenous standards."[47]

- **OCAP® (Ownership, Control, Access, Possession)** — First Nations Information Governance Centre (FNIGC); established in 1998, a registered trademark since 2015, predating and informing CARE. OCAP asserts that "First Nations… own and control how information can be stored, interpreted, used, or shared."[48]
- **FPIC (Free, Prior, and Informed Consent)** under UNDRIP (Articles 10, 11, 19, 28, 29) and ILO Convention 169 — the right of Indigenous peoples to give or withhold consent to actions affecting them and their resources, "defined by Indigenous Peoples," with consent that can be withdrawn at any stage.[49,50]
- **Networks:** the United States Indigenous Data Sovereignty Network (USIDSN), formed in 2016 and based at the Native Nations Institute, University of Arizona; Te Mana Raraunga (Māori Data Sovereignty Network, chartered 2016); Maiam nayri Wingara (Aboriginal and Torres Strait Islander Data Sovereignty Collective, 2017); and GIDA-Sápmi.[51,52]
- **Traditional Knowledge (TK) Labels and Local Contexts** (Jane Anderson, NYU, and Kim Christen, Washington State University; founded 2010) and the Mukurtu content management system (Christen, WSU): practical tooling that lets communities set consent-based, culturally determined

[47] GIDA / FNIGC. *CARE Directs Us Home: Prioritizing Indigenous Peoples' Community Standards* (communiqué), 2025, gida-global.org.

[48] First Nations Information Governance Centre (FNIGC). "The First Nations Principles of OCAP®" (established 1998; trademark 2015), fnigc.ca.

[49] United Nations Declaration on the Rights of Indigenous Peoples (UNDRIP), 2007, Articles 10, 11, 19, 28, 29, un.org — UNDRIP; ILO Convention No. 169. UN OHCHR, "Consultation and free, prior and informed consent (FPIC)."

[50] FAO. *Free, Prior and Informed Consent: An Indigenous Peoples' Right and a Good Practice for Local Communities* (2017 manual).

[51] United States Indigenous Data Sovereignty Network (USIDSN), formed 2016, Native Nations Institute, University of Arizona, usindigenousdata.org.

[52] Te Mana Raraunga — Māori Data Sovereignty Network (charter 2016), temanararaunga.maori.nz; see also Kukutai, Tahu, and John Taylor, eds., *Indigenous Data Sovereignty: Toward an Agenda* (ANU Press, 2016); Maiam nayri Wingara (2017); GIDA-Sápmi.

CC4.0

terms for cultural materials, recasting "intellectual property as culturally determinant and dependent upon cultural consent to use of materials."[53,54]

4.3 Washington's consultation obligations

Washington recognizes 29 federally recognized tribes.[55] The Centennial Accord (1989) and Millennium Agreement (1999), implemented through RCW 43.376, establish full government-to-government relationships and require each state agency to "make reasonable efforts to collaborate with Indian tribes in the development of policies, agreements, and program implementation that directly affect Indian tribes," to designate a tribal liaison reporting to the agency head, and to ensure consultation training.[56] SBCTC's Tribal Government Affairs office states that the agency "joins other state agencies and public education institutions to build government-to-government relationships with Washington's 29 federally recognized tribes," and that "the disproportionality of American Indian/Alaska Native student recruitment, retention, and completion rates can be addressed through building collaborative partnerships and consultation," including efforts to "incorporate indigenous knowledge systems" and to "train Washington community and technical college leaders about RCW 43.376."[57] SBCTC's system-wide vision, adopted June 27, 2019 and embedded in its 2020–2030 Strategic Plan, commits to "leading with racial equity" across its 34 colleges "in service to our diverse communities."[58] Together these establish that, for a Washington higher-education agency, consulting tribes about AI touching their data, knowledge, or imagery is not a courtesy but a legal and relational duty grounded in existing government-to-government relationships.

[53] Anderson, Jane, and Kim Christen. Local Contexts / Traditional Knowledge (TK) Labels (founded 2010), localcontexts.org. See Anderson and Christen, "'Chuck a Copyright on it': Dilemmas of Digital Return and the Possibilities for Traditional Knowledge Licenses and Labels," *Museum Anthropology Review*.

[54] Christen, Kimberly. "Tribal Archives, Traditional Knowledge, and Local Contexts: Why the 's' Matters." *Journal of Western Archives* 6(1) (2015); Mukurtu content management system, Center for Digital Scholarship and Curation, Washington State University, mukurtu.org.

[55] Centennial Accord Between the Federally Recognized Indian Tribes in Washington State and the State of Washington (Aug. 4, 1989); Millennium Agreement (1999); Washington Governor's Office of Indian Affairs (GOIA), goia.wa.gov.

[56] RCW 43.376 (Government-to-Government Relationship with Indian Tribes). Washington State Legislature, app.leg.wa.gov/RCW/default.aspx?cite=43.376. See also Environmental Law Institute, "Tribal Consultation in Washington State."

[57] Washington State Board for Community and Technical Colleges (SBCTC), Tribal Government Affairs office, sbctc.edu/about/tribal-government-affairs.

[58] SBCTC. Vision Statement (adopted June 27, 2019) and 2020–2030 Strategic Plan; Equity, Diversity, and Inclusion (EDI) office, sbctc.edu/about/edi.

CC4.0

Part 5: The Framework

5.1 Four commitments

No offsetting. Gate findings stop a use; they do not subtract from a total. No overall score is produced, because producing one would let offsetting back in through the scoring.

Proportionality. Assessment scales with consequence. Most AI use in a college or agency is routine and should be governed by general policy, not assessed. A framework that generates paperwork for drafting an email with AI assistance is inefficient and ineffective.

Evidence over intention. Benefit claims require a measure and a review date. Time-saving claims require an account of where the time went. Unevidenced claims are recorded as unevidenced.

Name the gap. Where a cost cannot be quantified, which, given current vendor disclosure, will usually be the case, the framework requires that the gap be recorded rather than treated as zero. An unmeasured cost is not an absent cost.

Step 1: Triage

Tier	What it covers	What is required
Tier 0	Individual, routine, low-stakes use of an approved tool, with no student-affecting output and no non-public data. Drafting correspondence, notetaking, formatting.	Nothing. General AI use policy applies.
Tier 1	Recurring unit-level use; any use touching non-public institutional data; any output informing a decision that affects people.	One-page assessment. Reviewed annually.
Tier 2	System-level deployment; new procurement; any use affecting student status, aid, admission, placement, discipline, or progression; any use involving Tribal data or data about Indigenous communities.	Full assessment with burden ledger. Consultation screening. Post-implementation review.

Step 2: Gates

Six conditions must hold. Failure at any one stops the use in its proposed form. Gates are not scored and cannot be offset.

CC4.0

	Condition	What failure looks like
G1	Lawful and privacy-sound	The use conflicts with FERPA, state law, contract, or the institution's data classification scheme — including student data entering a system without a documented legitimate educational interest determination.
G2	Named human accountability	No identified person is answerable for the output. Accountability cannot rest with a committee, a process, or the tool. If the output is wrong, there needs to be a named sponsor.
G3	No delegation of owed judgment	The tool substitutes for judgment the institution owes a person as a matter of right or professional duty – e.g. admission, aid, accommodation, academic standing, employment. Advisory use with genuine human determination passes. Nominal review of a machine recommendation does not. In other words, providing a service the college or agency is obligated to do anyway is not a benefit.
G4	Burden acknowledged and addressed	The proposer cannot say what is known and unknown about the use's impacts on energy, water, siting, materials, labor, and data provenance, or has taken no available step to reduce or disclose it. This gate does not require zero burden. It requires that the burden not be invisible.
G5	Cultural appropriation	The institution can demonstrate that the output will not (1) depict Native/Indigenous peoples generically or through stereotyped tropes (headdresses, buckskin, "stoic warrior," "vanishing Indian"); (2) reproduce tribally specific or sacred designs — Northwest Coast formline, Diné/Navajo weaving patterns, beadwork, regalia, or culturally protected motifs — without documented permission from the relevant nation or Native arts authority; or (3) rely on training data or produce outputs derived from Indigenous cultural materials without attribution and consent.
G6	Consent/Indigenous data sovereignty	The institution can demonstrate alignment with the CARE Principles and OCAP across the full data lifecycle; FPIC-consistent consent (free, prior, informed, and revocable); and Washington's government-to-government consultation duties under RCW 43.376 and the Centennial Accord, using the HEAL Act consultation framework (RCW 70A.02.100) as the operational model.

Step 3: Six dimensions, reported separately

These are assessed and reported as a profile, not summed. A reader sees six findings and forms a judgment, which is the right place for a judgment of this kind.

D1 — Public benefit

What improves, for whom, by how much, measured how, reviewed when. The dimension fails not when benefit is small but when it is unspecified. Required: the beneficiary named as a population rather than a

CC4.0

function; the outcome stated as a change rather than an activity; a measure that could show the claim false; a review date.

Counterfactual test. What would have happened without this use? Only the difference counts. A report that would otherwise have been written more slowly yields a time benefit, not an information benefit. This test does substantial work deflating inflated claims.

Equity test. Does the benefit reach students the system has historically served least well, or does it accrue where capacity already exists? An efficiency that accelerates well-resourced units and leaves others behind has widened a gap the system exists to close.

D2 — Human capability

Not whether a human is involved (a human almost always is) but whether the human involvement requires and therefore maintains expertise. Reviewing an AI-drafted analysis maintains analytical judgment only if the reviewer could have produced the analysis. Where that stops being true, the review becomes ceremonial and the capability is gone.

- Amplifying: extends expertise the institution retains and continues to develop.
- Neutral: performs work requiring no institutional expertise.
- Depleting: performs work that is the reason the role exists, with nothing maintaining the underlying skill.

Depletion is not automatically disqualifying, but it must be declared and paired with a maintenance strategy.

D3 — Resource proportionality

Whether the computational intensity of the method matches the difficulty of the task. Substantial institutional AI use applies large general-purpose models to work a search query, a template, a spreadsheet formula, where a small, specialized model would handle at a fraction of the cost. That is an environmental point, but it is first a competence point. Reaching for the largest available tool regardless of task signals that no one examined the task.

D4 — Redirection

Capacity saved is capacity released, and released capacity has a destination. Public value depends entirely on that destination. Redirection is therefore committed at approval and verified at review, by someone other than the approver.

Category	What it means	Effect on public value
Reinvested	Capacity moved to higher-value work not previously being done — more student contact, deeper analysis, quality improvement.	Positive. The efficiency became value.
Absorbed	Capacity consumed by existing backlog or higher volume of the same work.	Neutral to weakly positive. Legitimate, but the gain is throughput and should be claimed as such.

CC4.0

Category	What it means	Effect on public value
Displaced	Capacity used to justify unfilled vacancies, reduced hours, or role elimination.	Contested. May be a defensible budget decision, but it is a service reduction and should be reported as one.
Dissipated	No identifiable destination.	None. The efficiency claim is unsupported.
Negative	Verification and correction effort exceeded time saved.	Negative. More common than reported, because verification effort is rarely logged.

A necessary condition. Redirection findings must not be used as the sole or primary basis for position reductions, and that commitment must be stated in policy (section 8).

D5 — Displaced burden

This dimension records, at the point of decision, costs borne by parties outside the institution. The framing is deliberately modest. A college or agency cannot control where a data center is sited, what watershed it draws from, or whose land it occupies. It can control how much AI service it consumes, which vendors it selects, what it requires those vendors to disclose, and whether it is honest about what it does not know.

Burden	What to record	What the institution controls
Energy	Vendor-disclosed intensity where available; grid mix of serving infrastructure; whether load is additive to existing generation.	Volume; vendor selection; contract disclosure terms.
Water	Consumptive use for cooling; watershed and its stress condition; whether the basin is over-allocated or subject to treaty-reserved rights.	Volume; vendor selection.
Siting and community	Location of serving infrastructure; jurisdiction and character of the host community; documented opposition, moratorium, or consent process; ratepayer cost-shift evidence.	Vendor selection; procurement disclosure; public reporting.
Materials	Hardware lifecycle: extraction, manufacture, replacement cycle, e-waste pathway.	Refresh cycles for owned hardware; procurement standards.
Labor	Conditions of data annotation, evaluation, and content moderation work in the supply chain.	Procurement disclosure; vendor selection.
Data provenance	Whether training data was collected without consent; whether Indigenous data or knowledge is implicated; alignment with the CARE Principles.	Vendor selection; refusal; consultation.

CC4.0

On the blanks. Many of these fields cannot currently be completed. Vendor disclosure is poor, and independent measurement of AI's environmental impacts lacks both transparency and scientific consensus. That is a finding, not an excuse. The ledger records each burden as quantified, characterized, or undisclosed. An accumulation of "undisclosed" establishes that the institution is deciding without information it has a legitimate interest in having, which belongs in the record and in the next procurement conversation.

On consent. The equity question is not quantity but consent. A burden accepted by a community through a process it controlled is categorically different from one imposed on a community that was not asked. Purely quantitative environmental accounting cannot see that difference, which is why the ledger records jurisdiction and process alongside volume. Where serving infrastructure sits on or near Tribal land, the question is not only environmental but sovereign. Native nations are sovereign governments, and several have exercised that authority through moratoriums and formal impact studies. An assessment that treats Tribal land as a location rather than a jurisdiction has already erred.

D6 — Long-term capacity

- Knowledge: does the use produce reusable method and documentation, or one-off outputs?
- Literacy: does it build capacity to evaluate AI critically, including recognizing when it is wrong?
- Dependency: what happens if the vendor changes terms, raises price, degrades the model, or exits?
- Portability: can institutional data, method, and outputs be recovered in usable form?

Part 6: Worked Examples

The following examples are illustrative rather than empirical. Each shows the framework producing a result that differs from what an ordinary approval process would produce.

Legislative bill summarization

A policy unit uses a general-purpose model to summarize bills affecting the college system during session, with staff review before distribution. Tier 1.

All six gates pass (G5 and G6 are not triggered; no Indigenous depiction, data, or knowledge is implicated), with G4 conditional — burden characterized qualitatively, vendor disclosure absent and recorded as absent. The consequential finding is at D1. The counterfactual test reveals that summaries were previously produced selectively under time pressure, so the benefit is coverage rather than existence. The claim is narrowed from "improves policy awareness" to "extends coverage from roughly a third of relevant bills to substantially all." Redirection intent is recorded and scheduled for verification.

Result: proceed, with the benefit claim narrowed. The narrowing is the assessment doing its work.

Predictive student risk flagging

A system-level tool identifies students predicted to be at risk of non-completion and surfaces them to advisors. Tier 2.

CC4.0

G3 is decisive. The gate passes only if flags are advisory inputs to advisor judgment and cannot restrict access to courses, programs, aid, or opportunity. If flags gate anything, the use fails outright — regardless of the model's accuracy and regardless of how strong the benefit case is. That is the non-compensatory structure doing precisely what it was built to do.

At D1, the framework requires evidence that flagging changes outcomes, not merely that it predicts them accurately; predictive accuracy is not benefit, intervention efficacy is. At D2, depletion risk is significant, since advisors may substitute the flag for their own read of a student. At D3, many risk models are matched by a small number of directly observable indicators, so the chosen method needs justification. At D5, data provenance is implicated: whose historical outcomes trained the model, and does it reproduce the disparities it learned from?

Result: proceed only under the G3 constraint, with equity audit of model behavior across student populations and mandatory redirection verification.

Heritage-month campaign imagery

A college communications office proposes using a generative image tool to produce artwork for Native American Heritage Month and for recruitment outreach to Native students — posters, web banners, and social media graphics "in a Native style." Tier 1 by volume and data sensitivity; the depiction of Indigenous peoples and cultural material triggers the G5 screening regardless of tier.

G5 fails as proposed. The office cannot say which nations' visual traditions the model would draw on; cannot document permission from any nation or Native arts authority; and cannot obtain a vendor warranty that training data excludes scraped Indigenous cultural materials. The predictable output is the pan-Indigenous composite documented in 1.6 — generic "Native-style" regalia and motifs attributable to no nation, addressed to students who belong to particular ones. Visual quality is not the question; provenance and permission are, and neither can be shown. Nor is lawfulness the bar: the use would likely survive the Indian Arts and Crafts Act's analytic gap, which is precisely why the gate exists.

The failure is generative rather than terminal for the underlying goal. The assessment routes the office to the alternative the gate protects: commission the work from Native artists — ideally artists of the nations in the college's district, or Native students and alumni — with attribution and payment, and route the campaign through the college's Tribal relations function for review. At D1, the stated benefit (Native students see themselves represented accurately) is served better by the commissioned alternative than by the proposal, at modest cost; the counterfactual test shows the AI use would not have produced the claimed benefit at all, since a generic composite does not represent any actual student's nation.

Result: the use fails G5 and does not proceed in its proposed form. The redirected version — commissioned Native artwork with consultation — proceeds outside the framework's AI scope entirely. That is the gate working as designed, not a workaround: the non-compensatory structure did not weigh cultural harm against outreach benefit; it took the harm off the table and left the benefit reachable by other means.

Part 7: The Washington Context

CC4.0

The mission already contains the standard

Most institutions adopting an AI ethics framework must first argue that ethical criteria apply to an operational decision. The SBCTC system does not have that problem. Its work is grounded in Guided Pathways, ensured learning, and an explicit commitment to equity. The benefit dimension does not require a new definition of benefit; it requires asking whether a given AI use advances outcomes the system already commits to measuring.

An existing environmental justice method the system may adopt

The Healthy Environment for All Act (RCW 70A.02) requires seven state agencies — Agriculture, Commerce, Ecology, Health, Natural Resources, Transportation, and the Puget Sound Partnership — to embed environmental justice into agency operations. Obligations include environmental justice implementation plans, community engagement plans, Tribal consultation frameworks under RCW 70A.02.100, environmental justice assessments for significant agency actions, and integration of environmental justice into budget and funding decisions. The Act sets a goal that 40% of grants and expenditures creating environmental benefits reach Tribes, vulnerable populations, and overburdened communities, and it establishes an Environmental Justice Council with dedicated Tribal seats.

SBCTC is not a covered agency, and other agencies may opt in. This is the most consequential implementation fact. The system need not invent environmental justice methodology, community engagement standards, or a Tribal consultation framework. Mature templates exist, built by peer agencies, with an established review body.

Posture	What it involves	Assessment
Borrow the method	Use HEAL-style assessment structure for AI decisions with no formal relationship to the Act.	Lowest cost, available now, no external approval. Also weakest - no accountability relationship, and gates erode without one.
Voluntary alignment	Publicly commit to HEAL-consistent methodology for AI procurement and system-level deployment; report on it.	Practical, credible, reversible, and it produces a record. Recommended starting point.
Formal opt-in	Opt in under RCW 70A.02 with the full obligations across agency operations.	Strongest and most consequential. Extends well beyond AI. Should be decided on its own merits, not adopted as a byproduct of an AI framework.

Tribal consultation is a standing obligation, not an AI question

Washington maintains government-to-government relationships with federally recognized Tribes that predate any AI consideration and operate under their own statutory and accord-based framework. For this framework the design requirement is specific - consultation is triggered by the character of the decision, not by whether AI is involved. Where a deployment would use data about Indigenous students or communities, implicate Indigenous knowledge, or rely on infrastructure with documented effects on Tribal land or treaty-

CC4.0

reserved resources, the existing obligation applies and the framework routes to it rather than substituting for it.

The CARE Principles for Indigenous Data Governance (collective benefit, authority to control, responsibility, and ethics) should be adopted directly for any use involving Indigenous data. They were developed by Indigenous data sovereignty networks precisely because open-data principles ignored power differentials and historical context. Paraphrasing them into a system-specific scheme would recreate the authority problem they exist to correct.

The infrastructure is in the system's own districts

Washington cultivated data center development deliberately, and the resulting concentration is in the Columbia Basin. By 2025, seven of Grant County's ten largest taxpayers were data centers — six in Quincy and one in Moses Lake — with the assessed value of the top ten taxpayers reaching roughly $6.1 billion and generating about $54.3 million in property taxes, an increase of approximately 1,300% since 2006. Quincy's property tax levy rate fell by roughly 70% over that period, and state reporting indicates the added revenue funded substantial public building and infrastructure work.[59]

The policy environment shifted sharply in 2026. The Legislature considered House Bill 2515, which would have required utilities serving large energy users to adopt tariffs preventing cost shifts to other ratepayers, imposed clean energy requirements, and mandated reporting of energy and water use and emissions. The bill did not pass; its sponsor has indicated intent to continue the work in 2027. The Governor's Data Center Workgroup, whose preliminary report in December 2025 provided the backdrop for that bill, is expected to issue a final report that will inform the 2027 session.[60] Separately, 2026 legislation ended the sales and use tax exemption for replacement and refurbished data center equipment.[61]

Local governments moved faster than the Legislature. Moratoriums on new data center development have been adopted in Burien, Federal Way, Marysville, Moses Lake, Renton, Seattle, Skagit County, Snohomish

[59] Washington Department of Revenue, Data Center Workgroup, Tax Revenue and Economic Development Subgroup, adopted findings (October 2025), dor.wa.gov — Data Center Workgroup adopted findings (PDF) (in 2006, before data center siting, the assessed value of Grant County's top ten taxpayers was roughly $313 million, paying about $4.2 million in property taxes); "Where agriculture meets technology: Inside the small Washington town that became a data center hub," The Spokesman-Review, July 26, 2026, spokesman.com, July 26, 2026.

[60] H.B. 2515, 2026 Reg. Sess. (Wash.) (did not pass); Washington Governor's Data Center Workgroup, preliminary report, December 2025 (final report pending). During 2026 the Utilities and Transportation Commission separately opened work on rate guidance for large energy users.

[61] Engrossed Substitute S.B. 6231, ch. 266, Laws of 2026 (repealing the retail sales and use tax exemption for replacement and refurbished data center equipment, effective July 1, 2026); Washington Department of Revenue, dor.wa.gov — 2026 Tax Legislation.

CC4.0

County, and Spokane, with Yakima County commissioners directing staff to draft one in August 2026 after residents raised water scarcity concerns without any applications pending there.[62]

> *Big Bend Community College is in Moses Lake, which has adopted a moratorium. The environmental dimension of this framework is not imported from national advocacy. It describes conditions in the system's own service districts, in a community that has chosen to pause and study the question.*

The implementation vehicle exists

The eLearning Council AI Task Force connects eLearning and instructional design professionals across the system and produces practical guidance for responsible AI use in teaching, learning, and operations.[63] System AI policy work is already underway, addressing shared governance, the valuing of human creativity and labor, and documented attribution of AI contributions. This framework could be introduced through that existing structure as an extension of work already begun, not as a parallel initiative.

Part 8: An Inherent Conflict Considered

Colleges in the Columbia Basin have workforce development relationships with regional industry, including trades that serve data center construction and operation. Data center investment has been a major source of rural county revenue, and industry representatives, local chambers, and organized labor have all argued against rolling back the tax preferences that attracted it.

A framework that treats data centers primarily as an environmental harm therefore asks colleges to adopt a stance in tension with their own regional mission and their communities' economic interests. Careful wording will not dissolve that. Three responses:

Separate the object of assessment

This framework assesses the system's own AI consumption. It does not assess whether data centers should be built, takes no position on state tax preference policy, and does not evaluate colleges' workforce programs. The demand-side obligation is real and narrow. It does not extend to a systemwide position on regional economic development.

Acknowledge that the question is genuinely contested

Views in affected communities are not uniform. Some see economic opportunity, energy sales, ownership stakes, and revenue that could fund a generation of local investment. Others see AI development repeating a familiar pattern. Both positions are held sincerely by people with direct stakes, and a state education system

[62] MRSC, "What the Data Center Boom Means for Washington Local Governments," August 2026 (listing adopted moratoriums in Burien, Federal Way, Marysville, Moses Lake, Renton, Seattle, Skagit County, Snohomish County, and Spokane); "Data centers can raise electricity prices. WA wants to prevent that," The Spokesman-Review, August 29, 2026.

[63] Washington State eLearning Council AI Task Force, GenAI institutional policy resources.

CC4.0

is not the body to adjudicate between them. The framework requires the system to know what its consumption costs and who bears it. It does not require the conclusion that the answer is always no.

Treat it as curriculum

The strongest response is educational rather than defensive. A system that trains the workforce building this infrastructure, serves the communities hosting it, and consumes its output can teach the full picture in a way almost no other institution can. The tension is a pedagogical asset and suppressing it would waste this system's most distinctive contribution.

Part 9: Implementation

Three functions, not necessarily one body

Function	What it does	Suggested placement
Standard-setting	Maintains the framework, triage criteria, gate definitions, templates. Reviews annually.	eLearning Council AI Task Force, with data governance and policy staff.
Assessment review	Reviews Tier 2 assessments, confirms gate findings, escalates disclosure gaps.	Existing data stewardship or AI advisory structure; college-level equivalent for college-initiated work.
Redirection verification	Confirms at review what happened to released capacity.	Deliberately separate from the approving body.

The third row matters more than it looks. A unit that both approves a use on the basis of a projected efficiency and later confirms the efficiency occurred has no incentive to find otherwise.

Federated authority

SBCTC coordinates 34 colleges across 30 districts; it does not operate them. The framework distinguishes:

- Systemwide: assessments for system-level procurement, statewide deployment, and use of shared data infrastructure. These are SBCTC decisions.
- Available for adoption: triage criteria, gates, templates, and toolkit offered as a shared resource, with local discretion.
- Reserved to colleges: instructional use, academic freedom questions, college-level operational decisions. A system framework reaching into classroom practice will be resisted, correctly.

The system's comparative advantage is a shared method individual colleges cannot economically develop alone, particularly the vendor disclosure work, which benefits from being done once for 34 institutions rather than 34 times.

Sequencing

CC4.0

Phase	Focus	Content
Phase 1 Months 1–4	Establish and pilot	Adopt triage criteria and gates. Apply to three to five current or proposed uses across tiers, including at least one Tier 2. Do not announce systemwide. The purpose is to find where the instrument needs development or refinement.
Phase 2 Months 5–9	Revise and extend	Revise from pilot findings. Begin vendor disclosure inquiry for existing AI contracts and record the gaps. Draft consultation trigger criteria in coordination with the system's Tribal relations function — not in advance of it.
Phase 3 Months 10–15	Adopt and offer	Adopt for system-level decisions. Publish the toolkit. Run the first redirection verifications. Report findings, including negative ones.
Phase 4 Month 16+	Assess the framework	Evaluate against the tripwire below. Revise or discontinue rather than maintain.

Measures for the first year

Deliberately modest and chosen because each can come back negative.

- Tier 2 assessments completed, and how many had a gate finding that altered the proposal.
- Share of burden ledger fields returned undisclosed by vendors, tracked as a trend.
- Redirection verification outcomes by category, reported without penalty for absorbed, dissipated, or negative findings.
- Assessments where a lower-intensity method was substituted after the proportionality question.

Deliberately absent: adoption rate, colleges using the toolkit, staff trained. Those measure activity, not effect, and the framework should not model the behavior it exists to correct.

Part 10: Limitations and Open Questions

Readers should weigh the following before adopting anything here.

1. The framework is untested. It has not been applied at scale, and its behavior under real institutional pressure is unknown. The most likely failure is that gates soften into scores.
2. The burden ledger depends on disclosure that mostly does not exist. Its near-term value is as a procurement instrument and a record of ignorance, not as measurement. The intention is to spur this information from vendors and providers.
3. Redirection auditing raises legitimate surveillance concerns. It will produce useful data only where reporting is at unit rather than individual level and where the no-penalty commitment is real and stated in policy.

CC4.0

4. The demand-side argument is weakest at the level of a single institution. The argument rests on aggregate demand and on the expressive obligations of a system.
5. Non-compensatory gates will occasionally block uses with strong benefit cases. As the pilot progresses, the weight and definition of the gates should be evaluated, although strong benefit will not always outweigh the cost.

Where evidence would help most in further development

- Longitudinal tracking of released capacity across a set of institutional AI deployments, categorized as in D4. This would be the first evidence base on whether public-sector AI efficiency claims survive contact with what actually happened.
- Method development for institutional-level burden estimation under vendor opacity, that is, what can be inferred from regional infrastructure, and what genuinely cannot be known.
- Whether non-compensatory structures survive institutional pressure, and under what governance conditions.
- What obligations, if any, a public institution incurs toward communities hosting the infrastructure it consumes; a question with legal, ethical, and treaty dimensions largely unexamined at this unit of analysis.

A Note on Method: This Paper Assessed by Its Own Framework

This paper was prepared with substantial AI assistance: research synthesis, source retrieval, drafting support, and document design were performed with Anthropic's Claude, working from the author's design, direction, and prior drafts. A paper proposing that institutions account for their AI use cannot exempt its own production, so this section applies the framework to the work of writing it. The assessment is the author's. Highlighted brackets mark claims only the author can verify; they are published only once verified.

Triage. Tier 1: recurring, unit-level knowledge work. No student data, no non-public institutional data, and no output informing a decision about a person entered any AI system.

Gates. G1 passes: only public sources and the author's own drafts were processed. G2 passes: the named author is answerable for every claim; AI-assisted drafting does not dilute that accountability and it is the reason the byline carries obligations. G3 passes: no judgment owed to a person was delegated, and the positions taken are the author's. G4 is conditional, and honestly so. The vendor (Anthropic) does not disclose the energy or water footprint of the sessions that produced this paper, so the burden is recorded as undisclosed (the same finding the burden ledger anticipates for most institutional use). G5 passes with a note: no imagery was generated, and Indigenous scholarship is cited to Indigenous authors rather than paraphrased into the author's voice; where AI assistance touched those sections, the sources remained the authority. G6 is not triggered by authorship (no Indigenous data or knowledge was ingested) but the obligation it protects attaches to adoption. Any system-level use of this framework touching Tribal data,

CC4.0

knowledge, or imagery routes to consultation under RCW 43.376 before it proceeds, and this paper's existence does not substitute for that.

Dimensions, briefly. D1: the counterfactual is that this paper would exist without AI assistance, later and with narrower source coverage; the benefit claimed is therefore time and breadth, not existence. D2: The synthesis was reviewed and verified against sources; retained capability is asserted and the reader should weigh it as an assertion. D4: the released time went to accelerating essential work on the system's campus Climate Assessment surveys and associated materials to meet a prioritized and otherwise unattainable deadline, thereby facilitating beneficial outcomes for student and staff equity; per the framework, that claim is recorded at publication and is checkable. D5: undisclosed, as recorded at G4. D6: the method is documented and repeatable, which is the long-term capacity claim.

Two things this disclosure is not. It is not a claim that the framework blesses its own production; a self-assessment by the assessed is exactly the arrangement Part 9 warns against, which is why the redirection claim above is stated in verifiable form. And it is not a disclosure made because AI assistance is suspect. It is made because the framework's first commitment is that an unexamined use is the failure mode — whoever's use it is.

Part 11: Using and Adapting this Framework

The framework and its templates may be adapted and redistributed with attribution. No permission is required. Two requests:

- If you adapt it, record what you cut. The pattern of what practitioners remove is the best available evidence about which parts are not earning their place, and the author would welcome hearing about it.
- If you apply it and it fails, that is more useful to report than if it succeeds. Negative findings will be incorporated into future versions and attributed if you wish.

Companion materials — a toolkit of ten templates covering triage, assessment, the burden ledger, redirection verification, vendor disclosure, and consultation screening — are available at [URL].

CC4.0

Appendix: References and Further Reading

Frameworks and standards

- National Institute of Standards and Technology. AI Risk Management Framework (AI RMF 1.0).
- OECD. Recommendation of the Council on Artificial Intelligence.
- UNESCO. Recommendation on the Ethics of Artificial Intelligence, 2021.
- European Union. Artificial Intelligence Act. Transparency obligations under Article 50 effective August 2026; high-risk classification applies to admissions screening, student assessment, and performance monitoring.
- Global Indigenous Data Alliance. CARE Principles for Indigenous Data Governance, 2019. gida-global.org/careprinciples
- Carroll, S. R., et al. "Operationalizing the CARE and FAIR Principles for Indigenous data futures." Scientific Data, 2021.

Public value and public-sector AI

- Moore, M. Creating Public Value: Strategic Management in Government. Harvard University Press, 1995.
- "Artificial Intelligence and Public Values: Value Impacts and Governance in the Public Sector." Sustainability 15(6), 2023.
- Tan, J., and Chandra, Y. "AI as a framework for public service innovation." Public Money & Management 45(7), 2025.

Environmental and net-impact assessment

- "Greening AI or AI's Greenwashing? A Net Impact Assessment on Environmental Sustainability," 2026.
- Walton, J. L. "Sustainable AI Impact Assessment," 2025.
- "Navigating the Environmental Paradox of AI: A Decision Framework for Clean Technology Practitioners." Clean Technologies 8(2), 2026.
- Federation of American Scientists. "Reporting AI's Impacts to Build Public Trust and Advance AI," 2025. fas.org

Higher education

- EDUCAUSE. The Impact of AI on Work in Higher Education, January 2026 (with AIR, NACUBO, CUPA-HR).
- EDUCAUSE. 2025 AI Landscape Study.
- EDUCAUSE. 2026 Top 10; 2026 Horizon Report, Teaching and Learning Edition.

CC4.0

- Washington State eLearning Council AI Task Force. GenAI institutional policy resources.

Washington policy and context

- RCW 70A.02, Healthy Environment for All Act, including RCW 70A.02.100 (Tribal consultation).
- Washington Attorney General's Office. HEAL Act guidance and significant agency action determinations. atg.wa.gov/about-heal-act
- Washington Department of Health. HEAL Act Community Engagement Guide; Tribal Engagement Guide.
- Washington State Legislature. House Bill 2515 (2026), emerging large energy use facilities. Did not pass.
- Governor's Data Center Workgroup. Preliminary report, December 2025; final report pending.
- Washington Department of Revenue. 2026 tax legislation summary. ESSB 6231 (ch. 266, Laws of 2026), effective July 1, 2026.
- Municipal Research and Services Center. "What the Data Center Boom Means for Washington Local Governments," August 2026. mrsc.org
- Spokane Journal of Business. Reporting on Grant County data center property tax effects, July 2026.

Environmental justice and Tribal sovereignty

- American Indian Policy Institute. "Tribal Sovereignty in the Age of AI," 2025. aipi.asu.edu
- Cultural Survival. "Progress or Digital Colonization? AI Data Centers Spark Debate on Native Lands," 2026.
- Climate Justice Alliance. Data Centers and Artificial Intelligence Infrastructure. climatejusticealliance.org/datacenter
- Front and Centered. "Centering Indigenous and Frontline Perspectives on Data Centers and AI," 2026.
- Honor the Earth. Tracking of proposed data center projects on and near Tribal and rural lands.

CC4.0