


People Suck. Build Around the Suck.

VOLUME 2
DATA CENTERS



JMAC

SIDE ANGLE WORKS

People Suck. Build Around the Suck.

Volume 2

Data Centers

JMAC

Side Angle Works

First digital edition

Version 1.0

Published October 7, 2026

Publication details

Volume 2: Data Centers

Author: JMAC

Publishing imprint: Side Angle Works

Copyright © 2026 Brian Phillips.

All rights reserved in copyrightable human-authored contributions. No claim is made to public-domain material, third-party works, or purely AI-generated elements.

First digital edition · Version 1.0

Published October 7, 2026.

Complete PDF and EPUB editions: free from sideangleworks.com. A paid print edition is planned separately.

Published digitally in PDF and EPUB by Side Angle Works.

Source credits appear in the chapter references. The following authorship note describes the substantial assistance from ChatGPT and Codex.

Corrections and questions: sideangleworks@gmail.com

A note on AI, authorship, and accountability

This manuscript was built with substantial assistance from OpenAI's ChatGPT and Codex. I used them as thinking partners and working tools: to challenge assumptions, compare drafts, organize research, surface counterarguments, trace sources, and sharpen language.

It would be dishonest to hide that collaboration. It would also be dishonest to hand responsibility for the work over to the tool. The questions, values, judgments, and final choices are mine. AI did not decide what matters, what is true, or what I am willing to stand behind. It can be useful, fast, persuasive, and wrong. Every claim still has to earn its place.

That is consistent with the argument of this book. If we want to understand an outcome, we have to make the choices, tools, and conditions that produced it visible—including those behind this manuscript. Tools and institutions should be built around predictable failure, including our own. I used AI to extend attention and expose blind spots, not to escape responsibility. I remain accountable for this manuscript and for any errors.

Contents

1. FIRST, WHAT ARE WE ASSUMING?	9
2. OKAY, SO WHAT IS THE SYSTEM?	11
3. SO WHAT ARE WE ACTUALLY TRYING TO DO?	12
4. ONE PROJECT IS ACTUALLY A WHOLE CHAIN OF DECISIONS	14
5. THE COMMITMENT LEDGER	16
6. POWER—THE GRID IS NOT A BATTERY WITH INFINITE OUTLETS	17
7. FORECASTS ARE NOT LOAD	19
8. WHO PAYS FOR GRID EXPANSION?	21
9. STRANDED-ASSET RISK	23
10. WATER—DEFINE THE NUMBER BEFORE ARGUING ABOUT IT	25
11. LAND, NOISE, AIR, AND LOCAL INFRASTRUCTURE	27
12. ECONOMIC DEVELOPMENT—ECONOMIC ACTIVITY IS NOT THE SAME AS PUBLIC RETURN	29
13. INCENTIVES—DO NOT COUNT THE RIBBON CUTTING AND HIDE THE CONTRACT	30
14. JOBS—TEMPORARY AND PERMANENT ARE NOT INTERCHANGEABLE	32
15. THE AUTHORITY MAP	33
16. INFORMATION ASYMMETRY — WHO KNOWS MORE HAS POWER	35
17. STRATEGIC FORECASTS	36
18. ZERO REPORTED DOES NOT MEAN ZERO HAPPENED	37
19. THE DESIGN VS. IMPLEMENTATION GAP	38
20. GOOD ACTORS CAN HIDE BAD SYSTEMS	40
21. HOSTILE ACTORS ARE NOT THE ONLY THREAT	41
22. ABUSE SURFACE	42
23. CHECK BEFORE YOU BUILD	43
24. DESIGN REQUIREMENTS	45
25. DESIGN SUBTRACTION	47
26. DOES THE MISSING FACT ACTUALLY MATTER?	48
27. THRESHOLDS ARE INCENTIVES	49
28. TIME DEFINITIONS CAN BE GAMED TOO	50

29. NEW CAPACITY IS A NEW DECISION	51
30. NATIONAL IMPORTANCE DOES NOT ERASE LOCAL COST	52
31. LOCAL CONSENT IS A MECHANISM, NOT A CEREMONY	53
32. CORRECTION IS A DESIGN REQUIREMENT	54
33. GRACEFUL DEGRADATION — CAN ONE FAILURE STAY ONE FAILURE?	55
34. THE HOSTILE-OWNERSHIP TEST	56
35. WHAT THE STRESS TESTS REMOVED	57
36. THE PRECONSTRUCTION DECISION STANDARD	58
37. WHAT DATA CENTERS ADDED TO OUTCOME ANATOMY	59
38. THE TWO MODES OF OUTCOME ANATOMY	60
39. THE COMPLEMENTARY QUESTIONS	61
40. THE DATA-CENTER ANSWER	62
41. WHAT THE FIRST TWO VOLUMES HAVE IN COMMON	64
42. WHERE THIS LEAVES US	65

How to read this book

Start with the question. You do not need to agree with my conclusion to follow how I got there.

Some passages describe what a source reports. Others compare explanations or propose a way a system could work. Those are different kinds of claims. A useful idea still needs testing; a result from one place does not automatically tell us what happens everywhere.

Ask what the evidence shows, what remains uncertain, and what would change the conclusion. The same questions apply to my work.

This volume stands on its own. The other books use different subjects to practice the same method: Outcome Anatomy.

Scope of this volume

This volume examines data-center decisions and the commitments around them. National estimates, modeled scenarios, and a particular project's records answer different questions. Rules, contracts, resource constraints and public costs depend on place and time. Check the actual project before treating a general example as its description.

VOLUME 2 — DATA CENTERS

BUILDING THE SYSTEM BEFORE IT BUILDS ITSELF

A WAY INTO THIS QUESTION

You can start here. The approach used in this book is called Outcome Anatomy: begin with a result, or a result being promised, and ask how it is produced. Follow the choices, resources, incentives, and limits that connect the parts. Then test whether the evidence supports the explanation.

For data centers, that means looking beyond the building. A promise of jobs connects to tax agreements. A demand forecast connects to grid investment. A cooling choice connects to water supply. Decisions made separately can combine into costs, benefits, and obligations that nobody considered together.

Some consequences are already visible. Others are still forecasts. We will keep that distinction clear, compare competing explanations, and ask what would change our minds. The point is not to prove that data centers are good or bad. It is to understand what these decisions could build—and what can still be corrected before the choices harden.

What is the clean question, and what assumptions are hidden inside it?

Here is the question we need to answer:

WHAT RULES SHOULD GOVERN WHETHER, WHERE, AND UNDER WHAT CONDITIONS DATA CENTERS ARE BUILT SO THEIR BENEFITS, COSTS, RISKS, AND DECISION-MAKING POWER STAY VISIBLE, ARE DELIBERATELY ASSIGNED, AND CAN STILL BE CORRECTED?

Data centers are a useful test because the larger system is still forming—and no single actor is designing it.

Electric utilities make power decisions.

Water systems make supply decisions.

Local governments make zoning and land-use decisions.

States make tax and infrastructure decisions.

Companies make siting, cooling, ownership, power, and development decisions.

Regulators address pieces of the environmental and infrastructure footprint.

National policymakers consider computing capacity, security, industrial competition, and technological leadership.

Each decision can look ordinary in isolation.

Stack them together and they shape who receives power and water, who pays, who carries the risk, and which future choices remain open. Those consequences can begin to feel permanent.

The system can form before anyone consciously decides to create it.

That is the challenge Volume 2 gives Outcome Anatomy: follow how today's decisions could produce tomorrow's results, while separating what we can observe from what we can only anticipate.

CAN WE TRACE WHAT TODAY'S COMMITMENTS MAY PRODUCE WHILE THERE IS STILL TIME TO CHANGE THEM?

The foundational rule remains:

PEOPLE SUCK. BUILD AROUND THE SUCK.

Volume 2 adds:

WHILE THERE IS STILL TIME TO BUILD.

This is not an assumption that people are uniformly malicious.

People will work with incomplete information, protect their interests, respond to pressure, and resist changes that cost them something. Authority will be split. Forecasts may be optimistic. Costs may land elsewhere. The design has to account for those possibilities before they become excuses.

Do not build a system that works only when every actor is unusually wise, generous, accurate, transparent, and long-term oriented.

1. FIRST, WHAT ARE WE ASSUMING?

Before choosing a side, we took the claims everybody kept repeating and asked one rude question:

What exactly do you mean?

“Data centers are necessary.”

Necessary for which services, users, locations, time horizons, and national objectives?

“Data centers are economic development.”

Which benefits, over what period, compared with the tax incentives, infrastructure costs, what else the public gives up, and the risks it takes on?

“Data centers create jobs.”

Construction or permanent jobs? Jobs at the facility, jobs at suppliers, or jobs estimated elsewhere in the economy? Local hires or regional estimates? Temporary activity or durable employment?

“Data centers raise household utility rates.”

Because new power plants are needed? New transmission or local grid equipment? More backup capacity? Financing costs? Or because some of the added cost is shifted to other customers?

“Data centers consume too much water.”

Which cooling design? Which climate? Which water source? Withdrawal or consumption? Potable water or total footprint? Under average conditions or drought?

“AI is the demand driver.”

How much demand comes from training AI, running AI for users, ordinary cloud services, storage, networking, conventional computing, or capacity that companies have reserved but may never fully use?

“Renewable procurement solves the energy problem.”

Does it add usable power at the time and place it is needed? What supplies electricity when the contracted source is not producing enough? What new transmission is required? What happens when the contract and the facility's actual demand do not line up?

"Local approval equals local consent."

Who received notice? What information existed? Which costs and risks extend beyond the approving jurisdiction? What had already become difficult to reverse?

None of these questions makes the analysis anti-data-center.

It does not make it pro-development either.

Importance, necessity, benefit, harm, and appropriate response are separate propositions.

They do not get to borrow proof from one another.

2. OKAY, SO WHAT IS THE SYSTEM?

Look beyond one building and several systems collide:

computing demand and facility development

electric generation, transmission, local distribution, and reliability

large new power connections and who pays for them

water supply, cooling, treatment, and reuse

land use, zoning, noise, air emissions, and roads

tax breaks, grants, contracts, and performance promises

forecasts about future demand

local, state, and federal government

utilities, markets, companies, and communities

Different jobs require different expertise. A utility should understand the grid; a locality should understand land use.

But suppose each does its assigned job well and nobody checks what their decisions produce together. Who benefits? What gets harder? Who carries the cost? That is where the cracks tend to live: between the boxes.

A utility may evaluate interconnection while a locality evaluates zoning and a state evaluates incentives.

Each decision can be individually defensible while the combined system shifts costs, risks, or authority in ways no institution explicitly approved.

SYSTEMS OFTEN FAIL AT THE BOUNDARIES BETWEEN COMPETENT INSTITUTIONS.

To explain those outcomes, we cannot examine the facility alone.

We need to follow the network of commitments required to build, supply, subsidize, regulate, and eventually correct the project.

3. SO WHAT ARE WE ACTUALLY TRYING TO DO?

We are not trying to answer:

More data centers or fewer data centers?

That is too crude.

Useful projects should be able to proceed. Before we commit, we should know what they need, who pays, who takes the risk, and who decides. We also need a way to respond when a forecast turns out to be wrong.

That requires:

credible demand measurement

power and transmission planning

water planning

land-use compatibility

cost allocation

risk allocation

environmental compliance

transparent incentive accounting

enforceable commitments

ownership visibility

public notice

monitoring

correction

exit and failure planning

A lesson from the campaign-finance inquiry applies immediately:

DO NOT REMOVE OR ADD A MECHANISM UNTIL YOU UNDERSTAND THE FUNCTION IT MUST PERFORM.

Volume 2 adds:

DO NOT APPROVE A COMMITMENT UNTIL YOU UNDERSTAND WHAT FUTURE OPTIONS IT NARROWS.

The goal is not friction for its own sake.

It is to place useful friction before irreversible commitments and remove pointless friction from decisions that remain easy to correct.

4. ONE PROJECT IS ACTUALLY A WHOLE CHAIN OF DECISIONS

A data center does not appear because somebody stamped APPROVED on one piece of paper.

Its lifecycle can include:

site search

land acquisition

zoning and permitting

incentive negotiation

utility service requests

generation and transmission planning

water planning

construction

partial energization

expansion

operational changes

ownership changes

retrofits

closure or repurposing

Different commitments become difficult to reverse at different stages.

Land may be rezoned before the final load is known.

Incentives may be authorized before operating assumptions are proven.

Utilities may begin infrastructure investment before full buildout occurs.

Communities may count expected revenue before mature operation.

Developers may incur costs that create political pressure for later approvals.

Timing is therefore part of the mechanism.

An action that is easy to revise before construction can become financially, politically, or physically difficult to reverse later.

We need a check at each commitment that matters. One approval cannot tell us that every later expansion, contract, or operating change is ready too.

THE RELEVANT QUESTION IS NOT ONLY WHETHER A PROJECT HAS BEEN APPROVED.

IT IS WHICH COMMITMENTS HAVE ALREADY BECOME COSTLY TO REVERSE.

5. THE COMMITMENT LEDGER

Every material commitment should answer:

Who is committing?

What exactly is being committed?

What evidence supports the assumption?

Who receives the benefit?

Who bears the cost if the assumption fails?

What happens if the project changes size, timing, ownership, technology, or lifespan?

Can the commitment be paused, revised, or exited before the next irreversible step?

This became the Commitment Ledger.

Before one institution commits, it needs to know which assumptions another institution has actually checked. Copying the same assumption into another document does not check it again.

The ledger should follow the commitment across agencies and project stages. A forecast should not become a fact merely because it was copied into a permit, tariff, incentive agreement, or public presentation.

A useful rule emerged:

PAUSE ONLY THE NEXT MATERIAL COMMITMENT AFFECTED BY THE MISSING OR FAILED PROOF.

Do not freeze the entire system because one fact is unresolved.

Do not continue an irreversible commitment as if uncertainty does not matter.

Scale the response to the missing evidence and the reversibility of the next decision.

6. POWER—THE GRID IS NOT A BATTERY WITH INFINITE OUTLETS

A data center does not merely “use electricity.”

It affects where power comes from, the lines and equipment that carry it, backup capacity, connection planning, reliability, when the facility uses power, and who pays.

The relevant questions include:

How much load?

At what hour?

At what location?

How firm is the demand?

How quickly can it change?

Can it curtail during stress?

Which infrastructure must be built?

Who pays?

What happens if the projected load arrives late, shrinks, relocates, or disappears?

Annual energy use does not answer these questions.

A statewide percentage can conceal a severe local constraint.

An installed megawatt figure does not reveal hourly use, whether that use happens at the same time as the system peak, or what the facility does during grid stress.

NERC's July 2025 large-load white paper identifies planning and reliability risks from emerging loads, including data centers. Peak demand is only one relevant characteristic; operating behavior and the quality of planning data matter too. That does not establish that every data center threatens the grid.

Two cross-domain rules followed:

ABSOLUTE SCALE AND RELATIVE SCALE MUST BOTH BE REPORTED.

AVERAGE CONDITIONS DO NOT SUBSTITUTE FOR STRESS CONDITIONS.

A system can look comfortable on annual averages and become fragile during peak demand, outage conditions, extreme heat, fuel constraints, or transmission congestion.

Chapter references →

7. FORECASTS ARE NOT LOAD

A forecast can help us plan for demand. It cannot make the demand arrive. Once we start treating the projected load as if it is already plugged in, expensive decisions can get ahead of the facts.

A development announcement is not an energized facility.

An interconnection or service request is not final operating load.

Reserved capacity is not realized demand.

A contract is not necessarily utilization.

And observed operation at one stage is not proof of eventual campus buildout.

Large-load forecasting is difficult because projects can be duplicated across utility territories, proposed at several sites, delayed, resized, or abandoned while still influencing planning.

That does not make forecasts useless.

It makes status and confidence essential.

For every important quantity, ask:

announced?

requested?

studied?

approved?

contracted?

under construction?

partially energized?

fully energized?

average observed?

peak observed?

A number without its comparison base, project status, time period, and confidence level can create an illusion of precision.

THE CLASSIFICATION RULE IS PART OF THE RESULT.

CHANGE THE CLASSIFICATION OR DENOMINATOR AND THE APPARENT SYSTEM CAN CHANGE WITHOUT ANY PHYSICAL CHANGE AT ALL.

Planning should incorporate uncertainty explicitly instead of hiding it inside one authoritative-looking forecast.

Chapter references →

8. WHO PAYS FOR GRID EXPANSION?

Someone has to pay for the added infrastructure. That decision belongs in the main argument, where the people receiving the bill can see it.

A large new load may require generation, transmission, substations, distribution upgrades, reserves, studies, or administrative work.

Possible payers include:

the project

utility shareholders

other commercial customers

households

taxpayers

future customers

The question is not simply whether the project pays a connection fee.

It is whether the full added cost and risk created by the project are assigned appropriately across its lifecycle.

A project can pay the obvious direct costs while leaving future costs, shared-system costs, financing costs, or the risk of unused infrastructure somewhere else.

A requirement that a new customer fund upgrades may reduce cross-subsidy. It does not automatically resolve shared-use assets, later reimbursements, forecast error, creditworthiness, or the treatment of benefits that extend beyond the project.

To understand how the cost lands where it does, ask:

WHO GETS THE UPSIDE, WHO BEARS THE DOWNSIDE, AND DOES THE ACTOR MAKING THE FORECAST BEAR MEANINGFUL CONSEQUENCES IF IT IS WRONG?

Risk transfer became one of Volume 2's strongest contributions.

A subsidy should be named as a subsidy.

An insurance arrangement should be named as insurance.

A risk shifted into future rates should not disappear from the analysis merely because it is absent from the initial invoice.

Evidence note: In July 2025, Ohio's utility regulator approved a data-center tariff for AEP Ohio and identified cost-shifting risks from investments that might be underused. That is a named regulatory example of the risk and a response to it. The regulator's stated protective purpose is not, by itself, evidence that all costs stayed with the data-center customers.

Chapter references →

9. STRANDED-ASSET RISK

Suppose infrastructure is built for expected load.

Then the customer delays, scales down, changes technology, moves, restructures, becomes insolvent, or leaves before the infrastructure is paid for.

The customer may be gone. The steel, debt, and system costs can still be there.

Who absorbs the difference?

That is stranded-asset risk.

A contract can reduce the risk, but the contract itself must be examined.

How long does the obligation last?

Who guarantees it?

Which entity holds the liability?

Can the customer restructure around it?

Does it survive ownership changes?

Does the security remain adequate as exposure grows?

Does its duration match the useful life of the infrastructure?

That produced the Time-Horizon Matching Principle:

THE DURATION AND STRENGTH OF THE PROTECTION SHOULD MATCH THE DURATION AND SCALE OF THE RISK IT IS SUPPOSED TO COVER.

A short promise from a thin project entity is not necessarily meaningful protection against a decades-long infrastructure obligation.

We cannot promise zero stranded-asset risk. We can decide who carries it, make the price visible, require credible financial protection, and specify what changes when the assumptions do.

Tariff note: Ohio Power's October 2026 tariff includes minimum billing-demand, contract-term and financial-security provisions for covered data-center loads. Billing demand measures capacity in kilowatts; it is not a requirement to buy unused kilowatt-hours. For new loads, the initial term includes a ramp period plus eight years. Collateral requirements depend on specified credit and other conditions. Read the coverage, aggregation rules and exceptions before applying this example elsewhere. These provisions reduce particular exposures by design; their existence does not establish that every stranded-cost risk has been removed.

Chapter references →

10. WATER—DEFINE THE NUMBER BEFORE ARGUING ABOUT IT

“Data centers use too much water” is not an analytical statement until the water number is defined.

Relevant distinctions include:

withdrawal versus consumption—water taken from a source versus the portion evaporated, incorporated into products, or otherwise removed from the immediate water environment

potable versus nonpotable—drinking-quality water versus water that is not treated for drinking

surface water versus groundwater

municipal supply versus onsite source

direct cooling use versus upstream power-generation water

average use versus hot-day use

current operation versus permitted maximum

facility use versus full campus buildout

Two people can quote accurate numbers and appear to disagree because they are measuring different things.

Definition is part of the analysis.

The correct question is:

HOW MUCH OF WHICH WATER, AT WHAT TIME, UNDER WHAT OPERATING CONDITION, FROM WHICH SOURCE, WITH WHAT COMPETING USE AND REPLACEMENT CAPACITY?

Cooling technology changes the answer.

Climate changes the answer.

Electricity generation changes the upstream footprint.

Using reclaimed water can reduce demand for potable supplies, but suitable supply, infrastructure, cost, and treatment requirements can limit where it works. JLARC’s 2024 Virginia

report identifies those limits; it does not establish that reclaimed water is feasible at every site.

Again, average conditions are not enough.

A system may be comfortable annually and stressed during drought, peak summer demand, equipment failure, or simultaneous regional growth.

Water scarcity is local and temporal.

A national average cannot approve a local withdrawal.

A permitted maximum cannot be presented as observed consumption.

A low annual average cannot prove resilience under stress.

Accounting note: USGS distinguishes withdrawal, consumptive use and return flow. Report them separately, using the same period and geographic boundary. Potable supply is a water-quality category, not a substitute for total withdrawal or consumption. Returned water also needs a destination, timing and quality description before it can be counted as available to the same community. These definitions organize the comparison; they do not establish a facility's local impact.

Chapter references →

11. LAND, NOISE, AIR, AND LOCAL INFRASTRUCTURE

Some impacts are highly local even when benefits are regional or national.

Land conversion can affect neighboring uses and future development patterns.

Mechanical equipment can create persistent noise.

Construction can affect roads, traffic, dust, and public services.

Backup engines and onsite generation can create local emissions, especially when testing or emergency operation occurs near populated areas.

The existence of an impact does not establish its magnitude.

A dramatic local case does not establish a national average.

A small national share does not establish local insignificance.

JLARC's December 2024 Virginia review found that data-center backup generators contributed a small share of Northern Virginia's regional emissions of the pollutants studied. The report also noted that local air-quality effects were still being investigated. A regional share does not settle the impact at a particular site.

The correct analysis preserves both scales.

A GOOD REGIONAL AVERAGE CAN HIDE A BAD LOCAL OUTCOME.

For the actual site, ask:

Who receives the benefit?

Who experiences the burden?

How concentrated is it?

Who has authority to approve it?

What mitigation is enforceable?

Do the consequences stop at the approving boundary?

When consequences cross jurisdictional lines, local approval alone may not be enough to justify the result for everyone affected.

A project can be acceptable in aggregate and poorly sited.

It can also be controversial locally without imposing an unusual system-wide burden.

Both possibilities belong in the analysis.

[Chapter references →](#)

12. ECONOMIC DEVELOPMENT—ECONOMIC ACTIVITY IS NOT THE SAME AS PUBLIC RETURN

Data centers are frequently described through economic-development language.

That language can hide unlike measures.

Construction spending is not permanent employment.

Regional economic activity is not fiscal return to the jurisdiction granting an incentive.

Assessed property value is not net public revenue.

Gross tax receipts are not net benefit when public infrastructure, service obligations, financing costs, or foregone revenue are omitted.

Temporary construction employment should not be presented as permanent operating employment.

Modeled indirect and induced employment should not be presented as directly observed jobs.

Company capital investment is not the same thing as community return.

The public question is:

WHAT DOES THE COMMUNITY GIVE, WHAT DOES IT RECEIVE, WHEN DOES EACH OCCUR, HOW CERTAIN IS EACH SIDE, AND WHO BEARS THE RISK IF THE PROJECT UNDERPERFORMS?

That requires a comparison.

What would have happened without the incentive or project?

It also requires a time horizon.

A large one-time construction effect cannot automatically justify a long-lived tax exemption, utility obligation, or land-use commitment.

Before calling it growth worth paying for, follow what the public commits and what it can reasonably expect back.

Chapter references →

13. INCENTIVES—DO NOT COUNT THE RIBBON CUTTING AND HIDE THE CONTRACT

Tax exemptions, abatements, grants, infrastructure support, and other incentives can be rational tools.

They can also become bidding contests in which governments give up value for projects that might have chosen the same location anyway.

Keep the amounts separate:

statutory maximum

actual award

benefit paid only after required performance

foregone revenue

public infrastructure spending

company capital investment

projected activity

verified activity

Timing changes the risk allocation.

A benefit paid after verified performance is different from one granted before construction.

An exemption available automatically to every qualifying project is different from a negotiated package.

Infrastructure that remains useful to the community is different from infrastructure useful only to one customer.

The core question is:

WHAT MUST THE PROJECT ACTUALLY DO TO KEEP THE PUBLIC BENEFIT IT RECEIVED?

If the answer is vague, the commitment is weak.

Every incentive should identify:

the public objective

the required performance

the measurement method

the responsible verifier

the consequence of failure

the duration

the exit plan or way to recover benefits if promises are not met

Do not count the announcement as the benefit.

Do not count the statutory ceiling as the actual cost.

Do not count promised activity as realized return.

Evidence note: JLARC's 2024 incentives report warns that much of the activity associated with incentive programmes might have occurred without them. Its separate data-center report records companies' view that the tax exemption matters to their location decisions. Neither observation establishes how every project would have behaved without an incentive. That comparison must be tested, not assumed.

[Chapter references →](#)

14. JOBS—TEMPORARY AND PERMANENT ARE NOT INTERCHANGEABLE

A large construction project can support substantial employment during buildout.

A completed data center may require far fewer permanent onsite workers than the construction phase suggests.

Both numbers can be true.

Combining them without time labels can mislead.

Employment claims should identify:

construction jobs

construction duration

job-years versus individual workers

permanent operating jobs

direct jobs

indirect jobs

induced jobs

local hires

regional estimates

salary assumptions

Whether a job is “created” also depends on what would have happened otherwise. A worker moving from one regional project to another is not the same economic event as a previously unemployed resident entering durable employment.

The principle is:

TEMPORARY BENEFITS SHOULD NOT JUSTIFY PERMANENT COMMITMENTS WITHOUT MAKING THE TIME MISMATCH VISIBLE.

That does not make construction employment unimportant.

It makes duration, location, and measurement part of the claim.

[Chapter references →](#)

15. THE AUTHORITY MAP

Data-center decisions can be distributed among several authorities. The relevant powers need to be mapped for the jurisdiction and project.

That can provide useful checks.

It can also create accountability gaps.

A locality may control land use but not utility rates.

A utility may plan grid investment but not tax incentives.

A state may pursue investment without bearing every local cost or consequence created by the project.

A company may negotiate separately with institutions that see only part of the project.

Federal and regional authorities may address reliability, transmission, tax, air, or environmental questions without controlling the complete local decision. For example, EPA explains that state and local air agencies issue most data-center air permits under approved Clean Air Act programs.

Air-permitting authority is not the same as land-use, utility-rate, or tax authority.

Fragmented authority can become fragmented accountability.

For every material decision, the Authority Map asks:

Who decides?

What legal power do they possess?

What evidence can they see?

Which consequences are inside their jurisdiction?

Which fall outside it?

Who can challenge the decision?

Who monitors after approval?

Who can correct it if assumptions fail?

Who makes sure the handoff to the next decision-maker actually works?

A SYSTEM WITH DISTRIBUTED AUTHORITY STILL NEEDS OWNERSHIP OF THE INTERFACES.

That does not require one super-regulator.

We need to see who hands what to whom, use definitions the next institution can understand, and assign responsibility for consequences that cross the boundary.

Chapter references →

16. INFORMATION ASYMMETRY — WHO KNOWS MORE HAS POWER

Developers, utilities, governments, and communities do not enter negotiations with equal information.

Companies may know internal demand forecasts, technology plans, location alternatives, financing structures, and contractual flexibility that public institutions do not.

Utilities may possess system models that local officials and the public cannot easily interpret.

Local governments may understand land-use and fiscal constraints that regional institutions overlook.

Communities may understand lived impacts that technical models omit.

The question is not:

Who is lying?

It is:

WHO KNOWS WHAT, WHEN DO THEY KNOW IT, WHO CAN VERIFY IT, AND WHO BEARS THE CONSEQUENCE IF THE INFORMATION IS WRONG?

Information asymmetry can create power without deception.

Confidential information may be legitimately protected.

But confidentiality does not eliminate the need for independent verification of claims used to justify public commitments.

A workable system can separate:

public disclosure

confidential regulatory access

independent verification

aggregated reporting

Not everything has to be public.

The actor accepting public risk needs enough reliable access to check the assumption creating that risk. Protected information still needs a way to be tested.

17. STRATEGIC FORECASTS

Forecasts can be sincere and still be strategic.

A developer benefits from demonstrating a credible project.

A utility benefits from planning enough capacity.

A locality may benefit politically from announcing investment.

A state may benefit from displaying economic momentum.

A community group may benefit from emphasizing a high-impact scenario.

These incentives do not prove the numbers false.

They mean forecasts should remain forecasts.

Keep these kinds of information separate:

reported fact

company projection

utility forecast

scenario

model assumption

observed result

It also asks whether multiple forecasts rely on the same underlying announcement. Apparent agreement among models can be duplicated dependence on one uncertain input.

Forecast confidence should affect commitment strength.

THE STRENGTH AND PERMANENCE OF THE ACTION SHOULD MATCH THE CONFIDENCE OF THE EVIDENCE.

High uncertainty does not always mean do nothing.

It can mean:

stage the commitment

add review gates

require stronger security

preserve exit options

use scenarios rather than one point estimate

monitor the assumption

That is action matched to the level of uncertainty.

18. ZERO REPORTED DOES NOT MEAN ZERO HAPPENED

Compliance data often produce claims such as:

zero violations

zero complaints

zero outages

zero exceedances

zero missed commitments

Before celebrating a zero, ask how we would have found an event. Could monitoring detect it? Could affected people report it? Did the definition include it? Has the system faced the conditions that would expose the failure? A zero means more when those answers are solid.

The Zero-Reported Diagnostic asks:

Was the behavior observable?

Was monitoring capable of detecting it?

Was reporting required?

Did the definition capture the relevant event?

Were affected people able to report it?

Would the responsible institution know if the event occurred?

Could an event be recorded in another category?

A zero can be a finding.

It can also be a result of how the system measured or recorded the event.

The same caution applies to missing national data. Berkeley Lab's 2024 report and its newer update describe important gaps in national data-center energy estimates. Unknown is not zero, and modeled is not observed.

DO NOT CONFUSE ABSENCE OF RECORDED FAILURE WITH EVIDENCE OF PERFECT PERFORMANCE.

Chapter references →

19. THE DESIGN VS. IMPLEMENTATION GAP

A sound rule can fail in practice.

A bad outcome can also reflect bad design rather than poor execution.

Those are different diagnoses.

DESIGN GAP

The rule itself cannot reliably produce the required function.

IMPLEMENTATION GAP

The rule could work, but staffing, data, training, coordination, technology, or enforcement prevents it from operating as intended.

CAPACITY GAP

The assigned institution lacks the authority or resources required to perform the function.

INTERFACE GAP

Two competent institutions perform their own roles, but the necessary handoff between them is missing.

The remedies differ.

Do not create a new institution to solve a training problem.

Do not blame implementation for a design that requires impossible behavior.

Do not assign responsibility without authority.

Do not add another reviewer when the failure is an undefined handoff.

These two rules work in order:

DISCOVER BEFORE DESIGNING.

SUBTRACT BEFORE ADDING.

First identify the failed function and why it failed.

Then add only a control that performs a necessary function no existing control performs well enough.

Existing controls count when they genuinely perform the required function.

Understanding the real system should often make the design simpler.

20. GOOD ACTORS CAN HIDE BAD SYSTEMS

Sometimes a project gets through a weak process because somebody refuses to let the missing pieces stay missing.

A careful official catches an assumption nobody checked. A utility employee makes an extra call. A company representative supplies documentation the process never required. A community leader spots the conflict early enough to slow a bad commitment.

Those people deserve credit. Their work also leaves us with a question:

WHAT WOULD HAVE HAPPENED IF THIS PARTICULAR PERSON HAD NOT INTERVENED?

If the answer is failure, we need to look at what they repaired before calling the process a success.

That is the Good-Actor Masking Effect. Something in the design fails. A person works around it. The visible failure never happens, so the design problem stays hidden.

The distortion can run both ways. A bad actor can make a decent system look broken. A good actor can make a broken system look functional. We need to understand what each contributed before deciding what the system can actually do.

So test the design without assuming an unusually capable, connected, or courageous person will always be there to rescue it.

Do not punish the person who found a way through. If the workaround performs a necessary function, ask how to make that function part of the ordinary process.

Excellent people should make the system better. Keeping it competent should not require a rescue every time.

21. HOSTILE ACTORS ARE NOT THE ONLY THREAT

We need to test deliberate abuse. We also need to test Tuesday afternoon, when people are tired, deadlines are close, and nobody is trying to break anything.

Ordinary pressures are enough:

fatigue

ambiguity

career incentives

optimism

political pressure

weak training

convenience

fragmented responsibility

short time horizons

Some failures need someone deliberately exploiting a weakness. Others need only people making sensible choices for their own part of the job. We have to test both.

Outcome Anatomy asks how both kinds of behavior could shape the outcome. Neither a bad result nor a good one tells us, by itself, whether the people or the rules deserve the credit.

The hostile-actor question is:

WHAT WOULD ACTUALLY STOP A FUTURE BAD ACTOR FROM USING THIS POWER IMPROPERLY?

The ordinary-actor question is:

WHAT HAPPENS WHEN DECENT PEOPLE, UNDER NORMAL PRESSURE, MAKE REASONABLE LOCAL CHOICES THAT PRODUCE SYSTEM FAILURE?

A durable system must survive malice, error, pressure, and drift.

22. ABUSE SURFACE

Extraordinary power should not depend on the goodwill of whoever currently holds it.

Before giving someone consequential power, examine:

the breadth of discretion

the clarity of the criteria

who can see the evidence

who can challenge the decision

the standard of proof

the timing of review

the availability of appeal

the consequences of wrongful use

the ease with which an exception can become precedent

A recurring risk pattern is:

Give someone broad discretion, hide the evidence, make independent challenge difficult, and leave few real limits on review. We have created plenty of room for abuse. That is the abuse surface we need to shrink.

The risk can run in more than one direction.

A permitting process can be used to obstruct legitimate development.

An incentive or preemption regime can be used to bypass meaningful local accountability.

A utility rule can conceal favoritism or shift costs without a visible decision.

A public-notice process can satisfy the formal rule while giving affected people little practical notice.

The Actor-Reversal Test is the shortest form of Volume 1's Hostile Takeover Test:

WOULD THIS POWER STILL FEEL ACCEPTABLE IF SOMEONE YOU DISTRUST CONTROLLED IT?

The narrow test examines one power.

The broader test asks what happens when several powers can be combined.

If comfort depends on the actor, the control is not yet safe.

23. CHECK BEFORE YOU BUILD

Analysis eventually has to produce something more useful than “it’s complicated.”

The Preconstruction Audit asks whether the next material commitment is ready. A reviewer liking the project does not answer that. Neither does disliking it.

We need evidence that the necessary functions can work, with stronger evidence where the risk is greater.

The audit begins with four questions:

What is the project’s actual status?

What is the next material commitment?

What becomes difficult or expensive to reverse after that commitment?

What evidence must be true for the commitment to remain defensible?

Then test the system:

Are quantities defined consistently?

Are power and water assumptions stress-tested?

Are infrastructure costs and stranded-asset risks allocated deliberately?

Are public incentives tied to measurable, enforceable performance?

Are ownership and operational control sufficiently known?

Can the responsible institutions monitor what matters?

What correction remains available if assumptions fail?

The output should not default to approve or deny.

It may be:

proceed

proceed with conditions

pause the affected commitment

request specific evidence

redesign a control

reassess when a defined threshold is crossed

That is a response matched to the size of the unresolved risk.

The decision is only as broad as the unresolved risk.

24. DESIGN REQUIREMENTS

The proposed design grew through case analysis and stress tests. Those tests examined its logic and failure modes; they did not demonstrate that it outperforms existing governance in practice.

We stopped adding complexity when it stopped solving problems that mattered. Here are the requirements the remaining design has to meet.

A viable data-center governance system should:

1. Define project status and material quantities clearly.
2. Require evidence before commitments whose risks depend on that evidence.
3. Match infrastructure protection to the duration of the infrastructure risk.
4. Assign added costs and the risk of unused infrastructure deliberately rather than by default.
5. Distinguish average conditions from stress conditions.
6. Distinguish temporary benefits from durable obligations.
7. Make incentive promises measurable, verifiable, and enforceable.
8. Preserve correction points before commitments become difficult to reverse.
9. Map cross-jurisdiction consequences and institutional interfaces.
10. Require monitoring capable of detecting the failures the system claims to prevent.

11. Limit discretionary authority to what the function actually requires.

12. Survive ordinary human error, institutional pressure, and hostile ownership.

13. Reuse competent existing controls before creating new bureaucracy.

14. Remain understandable enough that responsible people can reconstruct why a project passed or failed a gate.

That final requirement is not cosmetic.

A governance system nobody can explain is difficult to challenge, audit, repair, or trust.

25. DESIGN SUBTRACTION

The early design had more controls than the final version needed. Testing them helped us find out which ones did a job of their own and which ones repeated a job another safeguard already did.

That gave us a useful discipline:

SUBTRACT BEFORE ADDING.

Before adding another control, ask what specific failure it prevents. Does an existing institution already handle that job? If we take the proposed control away, what serious failure becomes possible that the remaining safeguards cannot handle?

Then look at what the extra control creates. Another approval to wait for? Another place to pass information? Another rule that conflicts with the first one? Another power someone could use against the people the system is supposed to serve?

A diagram can look reassuring when it has plenty of boxes. Each box still needs to earn its place.

An unnecessary control can add delay, inconsistency, a new way to game the rules, or another place to lose track of who is responsible. More bureaucracy does not automatically buy more safety.

The goal is the fewest controls that reliably do every necessary job. Remove a safeguard only after we understand what it protects and what will still protect that function without it.

26. DOES THE MISSING FACT ACTUALLY MATTER?

Not every missing fact deserves the same response.

One of the most important design corrections was learning to distinguish a missing fact from a missing safeguard.

Some uncertainty will remain. If we wait for every imaginable question to be settled, nothing will move.

The right question is:

DOES THIS MISSING OR FAILED PROOF MATERIALLY AFFECT THE NEXT COMMITMENT?

If no, document the uncertainty, assign it an owner where appropriate, and continue.

If yes, pause only the commitment that depends on it until the evidence is sufficient or the design changes.

This avoids two opposite failures:

recklessness — proceeding through consequential uncertainty as though it does not exist

paralysis — treating every unresolved question as a reason to stop everything

This rule is not permission to ignore uncertainty.

It is a way to match the response to the consequence.

A well-designed system needs neither false certainty nor permanent indecision.

27. THRESHOLDS ARE INCENTIVES

A threshold does more than sort cases.

It changes behavior around the line.

A project-size threshold can encourage segmentation.

A threshold may create an incentive to stay below the cutoff. Economic research on bunching examines responses to taxes and regulatory thresholds; whether a particular data-center reporting rule produces that response must be tested.

An incentive threshold can alter timing, ownership, or classification.

A review threshold can encourage applicants to change the measured variable without changing the underlying impact.

None of this requires bad intent.

People notice what a rule rewards and adjust what they do.

Therefore:

EVERY THRESHOLD SHOULD BE TESTED FOR THE BEHAVIOR IT REWARDS IMMEDIATELY ABOVE AND BELOW THE LINE.

Ask what can be split, delayed, relabeled, averaged, or shifted outside the measurement window.

Then decide whether the system needs rules that combine related activity, look back over time, block obvious workarounds, or increase requirements gradually.

A threshold should simplify administration without turning the measured number into the real goal.

[Chapter references →](#)

28. TIME DEFINITIONS CAN BE GAMED TOO

Dates look objective.

But which date controls?

announcement date

application date

contract date

permit date

construction date

energization date

commercial-operation date

expansion date

A rule tied to one formal date can produce behavior designed around that date even when the underlying mechanism has already begun—or has not begun at all.

That gives us a simple time rule:

WHEN A RULE DEPENDS ON TIME, VERIFY THAT THE CHOSEN DATE CORRESPONDS TO THE MECHANISM THE RULE IS TRYING TO CONTROL.

A filing date may be easy to measure.

That does not make it the right date.

Define the triggering event, the evidence that proves it occurred, and how phased or restarted activity will be treated.

29. NEW CAPACITY IS A NEW DECISION

Classification becomes dangerous when an existing project expands.

Calling an addition part of an “existing project” can hide the fact that the system is making a new material commitment.

The analytical rule is functional:

IF AN ADDITION REQUIRES NEW MATERIAL RESOURCES, INFRASTRUCTURE, RISK, OR PUBLIC COMMITMENT, TREAT THE ADDED CAPACITY AS A NEW DECISION FOR EACH PURPOSE AFFECTED BY THOSE NEW CONSEQUENCES.

This does not require pretending the entire facility is new.

It requires examining the added decision honestly.

Test added load, added infrastructure, added incentives, added emissions, added water demand, and added exposure on their own terms.

Calling it an existing project does not make the new demand disappear.

30. NATIONAL IMPORTANCE DOES NOT ERASE LOCAL COST

The U.S. Department of Energy describes computing capacity as carrying national economic, technological, and security value.

That does not make every project necessary.

It does not prove every location is suitable.

And it does not justify hiding local costs.

The inverse is also true.

A serious local impact does not prove that a nationally useful activity should be prohibited everywhere.

The system must hold both scales at once.

National value may justify broader coordination, shared standards, or investment.

Local burden may justify mitigation, compensation, siting constraints, or refusal where the burden cannot be reasonably bounded.

Neither scale gets a veto over reality.

Importance is not immunity from analysis.

Harm is not automatic prohibition.

[Chapter references →](#)

31. LOCAL CONSENT IS A MECHANISM, NOT A CEREMONY

Public notice and public hearings can satisfy formal procedure while giving affected people little practical ability to understand or influence a decision.

But unlimited local veto can shift costs or consequences onto the wider region or country too.

The design question is not whether “the community supports the project” as though the community were one person with one preference.

The useful questions are:

Who is affected?

Who was notified—and how?

What information was available when participation occurred?

Which decisions were still genuinely open?

Whose costs and benefits were represented?

Which consequences extend beyond the locality?

What recourse remains after approval?

Participation matters only if it can change something.

If the decisive commitments are already made, what can the meeting still change? Public participation needs a real decision left to influence.

Practical consent requires timely information, choices people can understand, meaningful influence, and a way to challenge or revisit the decision later.

TIMING IS PART OF PARTICIPATION.

32. CORRECTION IS A DESIGN REQUIREMENT

Emerging systems are built under uncertainty.

Demand forecasts will sometimes be wrong.

Technology will change.

Cooling systems will evolve.

Ownership will change.

Projects will expand, contract, fail, or outlive their founding assumptions.

A good governance system therefore cannot depend on perfect prediction.

It needs correction.

For every important commitment, ask:

What happens if the assumption is wrong?

Who can detect the failure?

How quickly will they know?

Who has authority to respond?

Which response is available at that stage?

Can the commitment be adjusted without shifting the loss to someone who did not choose the risk?

Who bears the remaining cost?

Correcting the bad assumption and repairing harm it already caused are not always the same job. A complete system assigns ownership of both.

A system that can only approve is not a complete governance system.

A system that can only punish after irreversible failure is not enough either.

Correctability must exist while meaningful choices remain.

The later the correction point, the more expensive and political the correction becomes.

33. GRACEFUL DEGRADATION — CAN ONE FAILURE STAY ONE FAILURE?

Projects and institutions will fail in ordinary ways.

A company may become insolvent.

A forecast may miss.

A monitoring system may break.

A regulator may be understaffed.

A contract may be litigated.

A political administration may change.

A durable system should not require every safeguard to work perfectly at the same time.

The Graceful Degradation Test asks:

IF ONE CONTROL FAILS, DOES THE SYSTEM RETAIN ENOUGH INDEPENDENT PROTECTION TO PREVENT CATASTROPHIC OR IRREVERSIBLE HARM?

Redundancy matters only when the safeguards are meaningfully independent.

Three controls that depend on the same data, institution, assumption, or enforcement path can fail together.

Failure may still hurt. The test asks whether we can see it, contain it, and correct it before it takes the rest of the arrangement down.

This principle traveled cleanly from campaign finance into physical infrastructure.

The lesson travels because the same question remains useful: how did one failure stay contained—or spread through the rest of the arrangement?

34. THE HOSTILE-OWNERSHIP TEST

The Actor-Reversal Test changes the person holding a power.

The Hostile-Ownership Test changes ownership or operational control and asks whether the surrounding safeguards survive the transfer.

Data-center projects can change owners, financiers, customers, or operators over time.

Public institutions can change leadership too.

Safeguards must survive those reversals.

Would the contract still protect the public if the customer were sold?

Would the disclosure rule remain defensible if it were applied to a politically disfavored company?

Would the emergency authority remain acceptable under an administration you distrust?

Would local preemption remain acceptable when used for a project you strongly oppose?

Would a local veto remain acceptable when used to block infrastructure you consider essential?

Every consequential power should survive hostile ownership.

The test is not whether today's holder can be trusted.

It is whether the power remains bounded when tomorrow's holder cannot.

If the safeguard depends on the identity or goodwill of the actor, narrow the power or strengthen the control.

35. WHAT THE STRESS TESTS REMOVED

Taking controls out improved this design too. Here is what the tests gave us reason to remove or narrow.

Controls that duplicated stronger existing institutions were removed.

Requirements that demanded unavailable precision were narrowed to what could actually be known and verified.

Broad discretionary gates were converted, where possible, into evidence-linked conditions.

System-wide pauses were replaced with pauses tied to the specific material commitment at risk.

Static assumptions were replaced with monitored conditions when uncertainty could not be eliminated.

The design stopped trying to guarantee perfect projects.

It focused instead on:

bounded risk

visible commitments

credible monitoring

enforceable responsibility

independent safeguards

preserved correction

That is a more realistic design target.

A strong system is not the one with the most controls.

It is the one in which every remaining control performs a necessary function.

36. THE PRECONSTRUCTION DECISION STANDARD

So when should the project move forward?

We do not need every uncertainty gone. We do need enough evidence for the next commitment, with protections that keep the remaining risks within acceptable limits for as long as they last. Bigger risks need stronger protections.

A developer promising benefits does not settle that question. Neither does an opponent naming a possible harm, an official announcing support, or an approval from an earlier stage. Each may matter. None answers what this next commitment requires.

Ask what we are about to commit to, then follow what happens if an ordinary assumption turns out to be wrong.

Who pays? Which risks are protected against? Who will notice the problem? What can they actually change?

The standard brings those pieces together:

DOES THE NEXT COMMITMENT HAVE THE EVIDENCE, COST ALLOCATION, RISK PROTECTION, MONITORING, AND CORRECTION NECESSARY TO REMAIN ACCEPTABLE WHEN ORDINARY ASSUMPTIONS ARE WRONG?

The more serious the consequences, and the harder the decision is to undo, the stronger the evidence needs to be before we commit.

If a required piece of evidence fails, pause the commitment that depends on it. That failure does not automatically justify stopping every other useful part of the project.

Enthusiasm does not clear the bar. Complete certainty is not the bar either. We need enough evidence and protection for the decision we are actually making.

37. WHAT DATA CENTERS ADDED TO OUTCOME ANATOMY

Chapters 21–36 show how we tested these ideas. The earlier evidence matters too. The tests helped us improve the proposed design; they did not show that it produces better outcomes in practice.

Campaign finance taught us to work backward from political outcomes through substitution, selection, dependency, attention, enforcement, and accountability.

Volume 2 added a different class of lessons.

Physical constraints do not negotiate.

A system cannot bargain away transmission capacity, generation timing, water availability, land limits, or infrastructure life.

Definitions and denominators can determine the apparent result.

Average conditions can hide failure under stress.

Temporary benefits can be used to justify durable commitments.

Risk can be shifted away from the actor receiving the upside.

The time horizon of the safeguard must match the time horizon of the risk.

Interfaces between institutions can be more dangerous than failure inside any one institution.

Material commitments narrow future options.

Correction becomes more expensive as physical, financial, and political dependencies harden.

Information asymmetry creates power.

Timing is a mechanism.

Data centers widened that inquiry. An outcome can be shaped by strategy and incentives, but also by a physical limit that no institution can talk its way around.

38. THE TWO MODES OF OUTCOME ANATOMY

By the end of Volume 2, the same inquiry could run in two directions: back through an outcome already observed, or forward from a decision whose consequences are still taking shape.

DIAGNOSTIC MODE

Begin with an outcome we can observe.

Trace how incentives, rules, selection, dependency, and adaptation helped produce it. Check each connection against evidence; a plausible story is not enough.

Identify what could still change the result inside a system whose expectations and dependencies have already hardened.

Campaign finance lives largely here.

FORMATION MODE

Begin with an emerging decision or commitment, then trace the outcomes it could make more likely before infrastructure, interests, and dependencies harden. These are possibilities to test, not results we can claim to have observed.

Identify the commitments that create path dependence—making future choices harder to change.

Build monitoring, limits, and correction into the system while those choices are still meaningfully reversible.

Data-center governance lives largely here.

The modes can overlap.

A mature system may still contain new formation points.

An emerging system may already carry inherited dependencies.

The core inquiry is the same: which parts interact to produce the result, and what evidence supports that explanation?

The timing—and therefore the available leverage—is different.

39. THE COMPLEMENTARY QUESTIONS

Campaign finance taught:

WHEN YOU REMOVE A MECHANISM, ASK WHAT TAKES ITS PLACE.

Data centers added:

BEFORE APPROVING A COMMITMENT, ASK WHAT SYSTEM THAT COMMITMENT HELPS CREATE.

One question asks where activity moves when we change a rule.

The other asks what becomes harder to undo when we make a commitment.

Together they connect the outcome we see to the choices that shaped it—and to the choices we can still make.

What moves when we intervene?

What hardens when we proceed?

What becomes easier to reverse?

What becomes harder?

Those are not separate problems.

They are the before-and-after views of system change.

40. THE DATA-CENTER ANSWER

Chapters 21–36 explain why we proposed these requirements. We can examine that reasoning now. To know whether the requirements improve outcomes, we still need implementation and a comparison with realistic alternatives.

So what rules should govern whether, where, and under what conditions data centers are built?

The answer is not one national yes-or-no rule.

It is a set of rules for making high-stakes commitments when some facts are still uncertain.

Projects should move forward one major commitment at a time, with evidence strong enough for that next step.

Power, water, land, infrastructure, incentives, ownership, and failure risk should be measured using definitions that correspond to the mechanisms being governed.

Added costs and the risk of unused infrastructure should be assigned deliberately.

Forecasts should be treated as forecasts, not facts borrowed from the future.

Temporary benefits should not be used to justify permanent commitments without an explicit bridge between them.

Public benefits should be tied to measurable, verifiable performance.

Monitoring should be capable of detecting the failures the system claims to prevent.

Consequences and authority gaps that cross between institutions should be mapped.

Correction points should remain available until the relevant commitment becomes difficult to reverse.

Every public and private power created by the system should remain acceptable when controlled by someone we distrust.

And the whole system should remain understandable enough to challenge, audit, and repair.

Useful development should be able to proceed. It should not leave us with a system built by forecasts nobody checked, costs nobody accepted, and separate decisions nobody examined together.

41. WHAT THE FIRST TWO VOLUMES HAVE IN COMMON

Two volumes in, the pattern is clearer.

Campaign finance began with money.

Data centers began with the facility.

Neither visible object explained the outcome on its own.

Campaign finance became a study of how political capacity is created and allocated.

Data centers became a study of how physical limits, infrastructure, costs, risks, and fragmented authority combine into commitments that are hard to undo.

The subjects differ. In both, the visible object sent us looking for the decisions around it.

Ask the clean question.

Identify the outcome we are trying to explain—or the possible outcome we need to examine before committing.

Trace the sequence that connects choices, resources, and constraints to that outcome. Map the institutions where they matter to the explanation.

Separate what we know from what we infer.

Find the bottleneck.

Follow the incentives, dependencies, costs, and risks.

Ask where activity moves when a rule changes and what becomes harder to reverse when a commitment is made.

Test the design under ordinary failure and hostile control.

State what would prove the conclusion wrong.

Preserve the ability to correct course.

Then ask a better question.

Outcome Anatomy is a way to examine how results are produced, not a machine for producing the same answer in every domain.

Future domains must be allowed to refine, limit, or reject its principles—not merely illustrate them.

42. WHERE THIS LEAVES US

Campaign finance taught us to follow what happens when one route to political capacity is removed.

Data centers taught us to follow what today's commitments could make possible—or difficult to change—tomorrow.

One traces how an existing outcome came about.

The other examines how a future outcome could take shape, without mistaking a forecast for a fact.

Together they establish the posture of the series:

UNDERSTAND THE MECHANISM.

FOLLOW THE INCENTIVES AND DEPENDENCIES.

MAKE COST, POWER, AND RISK VISIBLE.

PRESERVE CORRECTION.

DO NOT WAIT FOR FAILURE TO ASK WHAT OUR DECISIONS ARE PRODUCING.

That is systems literacy.

It does not require everyone to agree on values, priorities, or policy.

It requires us to see what our choices actually build, whom they empower, what they displace, and how failure can be corrected.

And underneath all of it remains the less polite rule that gave the series its name:

PEOPLE SUCK. BUILD AROUND THE SUCK.

We get tired, respond to incentives, protect what we control, and miss what our institutions do not make visible. A choice can make sense from our desk and still help produce a disaster beyond it. The design has to account for the people who will actually use it.

Good systems do not demand better humans before they can work.

They make competent behavior easier, abuse harder, failure visible, and correction possible.

The subject will change.

The discipline should travel.

Chapter references

6. POWER—THE GRID IS NOT A BATTERY WITH INFINITE OUTLETS

North American Electric Reliability Corporation, Large Loads Task Force. (2025, July). *Characteristics and risks of emerging large loads* [White paper].

Source: https://www.nerc.com/globalassets/who-we-are/standing-committees/rstc/3_doc_white-paper-characteristics-and-risks-of-emerging-large-loads.pdf

[Return to chapter →](#)

7. FORECASTS ARE NOT LOAD

North American Electric Reliability Corporation, Large Loads Task Force. (2025, July). *Characteristics and risks of emerging large loads* [White paper].

Source: https://www.nerc.com/globalassets/who-we-are/standing-committees/rstc/3_doc_white-paper-characteristics-and-risks-of-emerging-large-loads.pdf

[Return to chapter →](#)

8. WHO PAYS FOR GRID EXPANSION?

Federal Energy Regulatory Commission. (2026, July 22). *Interconnection of large loads to the interstate transmission system* (Docket No. RM26-4-000) [Webpage, last-updated date].

Source: <https://ferc.gov/rm26-4>

Public Utilities Commission of Ohio. (2025, July 9). *PUCO orders AEP Ohio to create data center specific tariff* [News bulletin; Case No. 24-508-EL-ATA].

Source:
<https://content.govdelivery.com/accounts/OHPUC/bulletins/3e8bb79>

[Return to chapter →](#)

9. STRANDED-ASSET RISK

Ohio Power Company. (2026). *Schedule DCT (Data Center Tariff)*, P.U.C.O. No. 22, first revised sheets 223-1-223-8, issued and effective August 28, 2026. In *Ohio Power tariff book*, October 2026 compilation.

Source: https://www.aepohio.com/lib/docs/ratesandtariffs/Ohio/October_2026_Ohio_Power_Tariff_Book.pdf

[Return to chapter →](#)

10. WATER—DEFINE THE NUMBER BEFORE ARGUING ABOUT IT

U.S. Department of Energy, Federal Energy Management Program. (n.d.). *Cooling water efficiency opportunities for federal data centers*.

Source: <https://www.energy.gov/cmei/femp/cooling-water-efficiency-opportunities-federal-data-centers>

Shehabi, A., Smith, S. J., Hubbard, A., Newkirk, A., Lei, N., Siddik, M. A. B., Holecek, B., Koomey, J., Masanet, E., & Sartor, D. (2024, December). *2024 United States data center energy usage report* (LBNL-2001637). Lawrence Berkeley National Laboratory.

Source: https://eta-publications.lbl.gov/sites/default/files/2024-12/lbnl-2024-united-states-data-center-energy-usage-report_1.pdf

Joint Legislative Audit and Review Commission. (2024, December 9). *Data centers in Virginia* (Report 598). Commonwealth of Virginia.

Source: <https://www.appomattoxcountyva.gov/DocumentCenter/View/768/JLARC-Data-Center-Study?bidId=>

U.S. Geological Survey, Water Science School. (2018, June 15). *Water science glossary*.

Source:
<https://www.usgs.gov/water-science-school/science/water-science-glossary>

[Return to chapter →](#)

11. LAND, NOISE, AIR, AND LOCAL INFRASTRUCTURE

U.S. Environmental Protection Agency. (2026, October 1). *Clean Air Act resources for data centers* [Webpage, last-updated date].

Source: <https://www.epa.gov/stationary-sources-air-pollution/clean-air-act-resources-data-centers>

Joint Legislative Audit and Review Commission. (2024, December 9). *Data centers in Virginia* (Report 598). Commonwealth of Virginia.

Source: <https://www.appomattoxcountyva.gov/DocumentCenter/View/768/JLARC-Data-Center-Study?bidId=>

[Return to chapter →](#)

12. ECONOMIC DEVELOPMENT—ECONOMIC ACTIVITY IS NOT THE SAME AS PUBLIC RETURN

Joint Legislative Audit and Review Commission. (2024, December 9). *Data centers in Virginia* (Report 598). Commonwealth of Virginia.

Source: <https://www.appomattoxcountyva.gov/DocumentCenter/View/768/JLARC-Data-Center-Study?bidId=>

[Return to chapter →](#)

13. INCENTIVES—DO NOT COUNT THE RIBBON CUTTING AND HIDE THE CONTRACT

Joint Legislative Audit and Review Commission. (2024, November 7). *Economic development incentives 2024: Spending and performance* (Report 597). Commonwealth of Virginia.

Source: <https://jlarc.virginia.gov/pdfs/reports/Rpt597.pdf>

[Return to chapter →](#)

14. JOBS—TEMPORARY AND PERMANENT ARE NOT INTERCHANGEABLE

Joint Legislative Audit and Review Commission. (2024, December 9). *Data centers in Virginia* (Report 598). Commonwealth of Virginia.

Source: <https://www.appomattoxcountyva.gov/DocumentCenter/View/768/JLARC-Data-Center-Study?bidId=>

[Return to chapter →](#)

15. THE AUTHORITY MAP

U.S. Environmental Protection Agency. (2026, October 1). *Clean Air Act resources for data centers* [Webpage, last-updated date].

Source: <https://www.epa.gov/stationary-sources-air-pollution/clean-air-act-resources-data-centers>

[Return to chapter →](#)

18. ZERO REPORTED DOES NOT MEAN ZERO HAPPENED

Smith, S. J., Hubbard, A., Newkirk, A., Ganeshalingam, M., Holecek, B., Sartor, D., Mills, M., & Shehabi, A. (2026, June 18). *United States data center energy usage report: 2025 update* (LBNL-2001758). Lawrence Berkeley National Laboratory. <https://doi.org/10.71468/P1RP4F>

Source: <https://www.energy.gov/documents/united-states-data-center-energy-usage-report-2025-update>

[Return to chapter →](#)

27. THRESHOLDS ARE INCENTIVES

Kleven, H. J. (2016). Bunching. *Annual Review of Economics*, 8, 435–464. <https://doi.org/10.1146/annurev-economics-080315-015234>

Source: <https://www.annualreviews.org/doi/abs/10.1146/annurev-economics-080315-015234>

[Return to chapter →](#)

30. NATIONAL IMPORTANCE DOES NOT ERASE LOCAL COST

U.S. Department of Energy. (n.d.). *Powering America's AI future—Data center resource hub*.

Source: <https://www.energy.gov/powering-americas-ai-future-data-center-resource-hub>

[Return to chapter →](#)

About the author

JMAC is the writing name of Brian Phillips, an independent systems researcher and writer building Side Angle Works.

Before this work, he spent years in professional kitchens. That experience still shapes his approach: understand the whole operation, find the bottlenecks, work with the resources available, and account for real human behavior.

Through Outcome Anatomy, he examines how choices, conditions, incentives, rules and resources combine to produce results.

Teaching people to ask better questions.
Team Human.

Corrections and contact

A correction is useful. If you find an error, a missing explanation, or evidence that changes the picture, include the volume, chapter and passage so I can find it.

sideangleworks@gmail.com

Other volumes in the series

People Suck. Build Around the Suck.

- Volume 1: Campaign Finance
- Volume 2: Data Centers
- Volume 3: Foster Care and Child Welfare
- Volume 4: Addiction-Related Homelessness

JMAC · Side Angle Works

A data center gets approved. What else have we committed to?

Behind the building are decisions about electricity, water, land, tax breaks, and public infrastructure. Different institutions make those decisions. Their combined consequences can reach far beyond the site—and last longer than the promises that helped secure approval.

This book follows the commitments: who benefits, who pays, who carries the risk, and what happens when a forecast fails. It asks what evidence belongs before each major decision, which promises can be enforced, and which future choices are becoming harder to change.

Outcome Anatomy traces how choices and conditions combine to produce results. Here, it asks us to examine the system taking shape while there is still time to change it. A useful project should be able to answer difficult questions about the obligations it creates.

Teaching people to ask better questions.

SIDE ANGLE WORKS

Team Human.