Letter to New Castle residents about Data Center Ordinance 2026-568
Cynthia Batty, August 27, 2026

Subject: **Hearing on Data Center Ordinance for New Castle City next Monday**

Dear fellow New Castle residents,

Next Monday, August 31, the Planning Commission will hold a public hearing on an ordinance that the City Council has put forward to regulate data centers within our 3.2 square mile city. (https://newcastlecity.delaware.gov/files/2026/08/Ordinance-2026-568-City-of-New-Castle-Data-Center-July-6-2026-edits-2007120.2-002.pdf, and attached) This ordinance is notable for what is included, and what is not included. It is also concerning because New Castle County already has a Data Center ordinance which governs our City area (https://library.municode.com/de/new_castle_county/ordinances/code_of_ordinances?nodeId=1416770, and attached) and it appears that the proposal for the City may have more loopholes than the County ordinance.

You can show up in person for this hearing, or you can send a letter to comment or ask a question to the New Castle City office no later than 3:00 pm on Monday August 31 (no more than 600 words, must be read in 3 minutes). Below is some analysis that you might want to comment on or question.

+++++++++++++

Planning Commission Public Hearing - 08/31/2026 - 06:30 PM (https://newcastlecity.delaware.gov/events/planning-commission-public-hearing-12/)

Location: Community Room New Castle Police Station, Address: 1 Municipal Blvd, New Castle

This hearing will be held in the Community Room at the New Castle Police Department at 6:30PM. You can also attend virtually via Zoom (details on the link above).

If you are unable to attend the Public Hearing, **questions and comments will be taken via email up to 3:00 p.m. on Monday, August 31, 2026 at info@newcastlecity.delaware.gov.**

Ordinance 568 –An Ordinance to Amend the City of New Castle Zoning Code to Establish Standards for the Review, Siting, and Operation of Data Centers (Council Person Joseph Day 07/07/2026)

+++++++++++++++++

Analysis

The City of New Castle is very small at 3.2 square miles. The city zoning contains 14 zones of various types, including industrial, residential, recreational and other types of approved uses. You can see the map of the zones in the New Castle City Comprehensive Plan, Map 4 near the end of this document. (https://ecode360.com/documents/NE1562/public/607645270.pdf, maps only at https://ecode360.com/documents/NE1562/public/576577610.pdf).

Those who follow county politics will know that the County approved a Data Center ordinance on March 10, 2026. This ordinance includes a 1000 foot buffer zone from residential, day care and schools, a broader interpretation of industrial zones, noise not to exceed predevelopment noise, water and energy consumption restrictions, back up power risk management, and strict rules for cleanup of the site in the event it is abandoned which includes a guarantee of funds to cover cleanup before the project breaks ground.

The proposed New Castle ordinance largely restates the County ordinance. It specifically limits the possible development of data centers to the zone I – Industrial, which alone disqualifies a hyperscale (AI) data center (see Map 4 annotated below, showing light yellow I zones and dark yellow actual size of proposed Starwood data center in Delaware City which is currently blocked by DNREC).

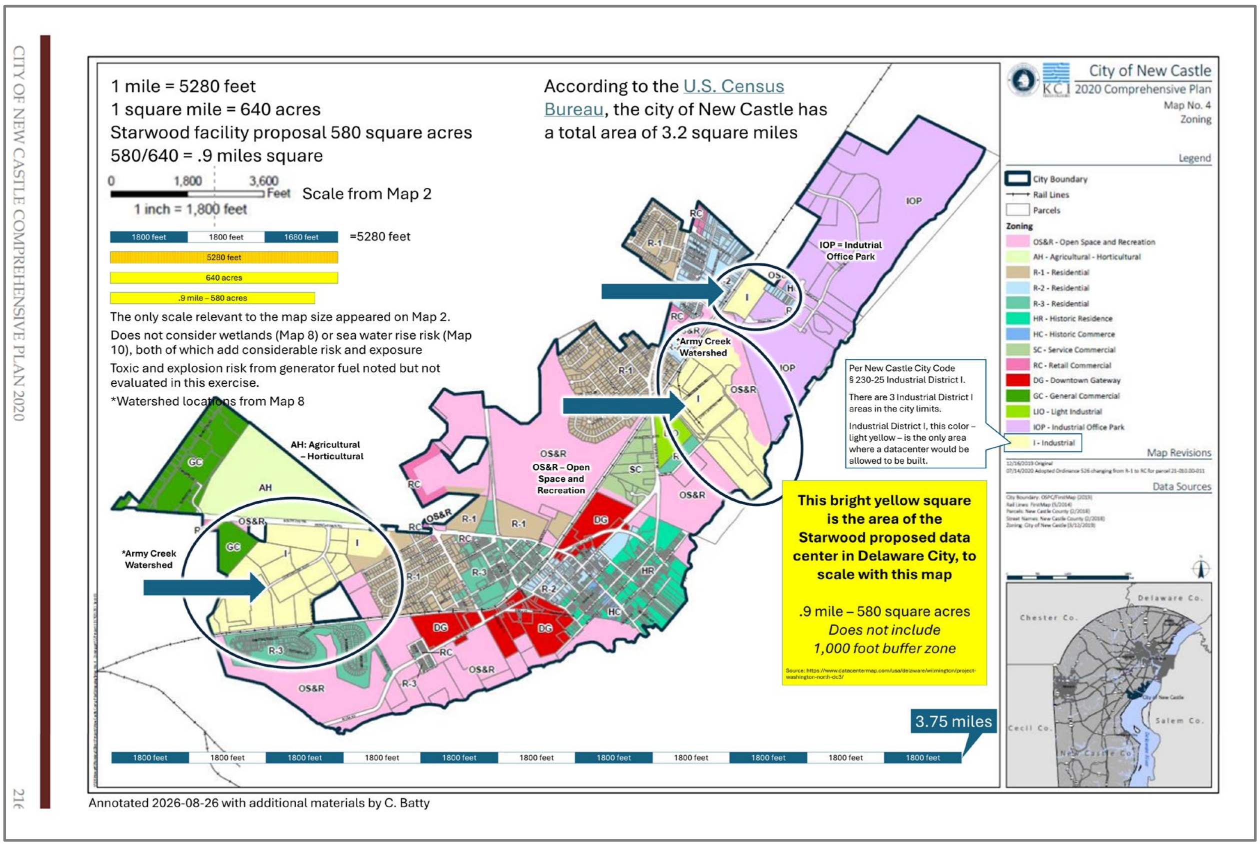

(This map is attached at a larger size NC Map 4 with DC 24 x 36 2026-08-26 v1.pdf so you can zoom in and look at it in detail.)

In principle, the rule that only the I zones can have a data center should be enough to disqualify a hyperscale (AI) data center. If the maps concerning sea water rise (Map 10) and protected wetlands (Maps 8 and 9) are superimposed on this map, it is even more clear that these risks and prohibitions would be enough for DNREC to stop a data center from being built in our tiny, incorporated city area because of fuel storage requirements and risks.

It is not clear why this ordinance does not simply say no to data centers, as all of the zones that are I – Industrial are far too small to accommodate a data center plus the 1,000 foot buffer zone as required in the County ordinance.

However, if for some reason it is required that the City have an ordinance, there needs to be some specific additions to the draft that has been presented.

1. The variance approval by the Board of Adjustment must be severely limited to explicitly state that the BoA may not extend a data center into any other zone other than the **existing I zones**. Because the approval of the data center is expressed as a variance which must be approved by the Board of Adjustment, and the BoA has the ability to allow zone variances, this must be explicitly forbidden in this ordinance. Failing to do this could allow IOP – Industrial Office Park or OS&R – Open Space and Recreation zones to become extensions of the I – Industrial zone, which would open up enough space to build a hyperscale data center.

2. Increase the sound barrier escapement from 1000 feet to 3000 feet. It should be more, as it has been shown that even at a mile and a half (7920 feet) residents have hearing loss near data centers, see attached analysis (Data_Center_Noise_Analysis.pdf).

3. In anticipation of near future uses, exclusion of local nuclear power options in any form must be added to the New Castle ordinance. The industry is talking about small nuclear plants at the data center level to address power concerns and eliminate the huge fields of oil drums filled with volatile fuel. This brings a level of management, maintenance and monitoring that is far beyond the skills and capabilities of data center management or the developer companies seeking to build such sites. This must be explicitly outlawed.

And to help dispel notions that somehow data centers bring money into cities and counties, please see the attached document Economic analysis of datacenters.pdf as well as https://www.brookings.edu/articles/new-evidence-on-data-center-employment-effects/

Letter to New Castle residents about Data Center Ordinance 2026-568
Cynthia Batty, August 27, 2026

I hope you will take some action to either appear at and listen to the public hearing on Ordinance 2026-568 or write a letter to be read into the record expressing your concerns. If you write a letter, keep the word count to about 600 words to allow all of your concerns to be read into the record in 3 minutes.

Thank you,
Cynthia Batty

Attachments:

Ordinance-2026-568-City-of-New-Castle-Data-Center-July-6-2026-edits-2007120.2-002.pdf

County Data Center ordinance 25-101.pdf

NC Map 4 with DC 24 x 36 2026-08-26 v1.pdf

Data_Center_Noise_Analysis.pdf

Economic analysis of datacenters.pdf

Ordinance No. 2026-568

AN ORDINANCE TO AMEND THE CITY OF NEW CASTLE ZONING CODE TO ESTABLISH STANDARDS FOR THE REVIEW, SITING, AND OPERATION OF DATA CENTERS

WHEREAS, the Council of The City of New Castle ("City Council") possesses the authority to adopt, amend, modify, and repeal the City of New Castle Municipal Code ("City Code"); and

WHEREAS, the City Council has determined that it is in the best interests of the City and its residents to proactively regulate data center development to ensure compatibility with surrounding land uses and to protect public health, safety, and welfare; and

WHEREAS, data centers present unique siting and operational considerations, including energy demand, backup power systems, cooling systems, noise, and potential environmental impacts; and

WHEREAS, it is in the public interest to adopt clear standards governing buffers from sensitive receptors, noise performance, cooling technologies, air quality, electrical capacity coordination, abandonment, and decommissioning security; and

WHEREAS, this Ordinance proposes standards addressing data center siting, noise, cooling systems, generator technology, utility coordination, abandonment, and financial assurances, which the City finds instructive for adaptation of data centers and application within the City of New Castle; and

WHEREAS, the City Council further finds that the provisions of this Ordinance are rationally and reasonably related to, substantially advance, and are narrowly tailored to achieve its goal of protecting and preserving legitimate governmental interests, including, but not limited to, the protection and preservation of the public health, safety, prosperity, general welfare and quality of life.

NOW, THEREFORE, making the express finding that the provisions of this Ordinance enhance and promote the health, safety, and welfare of the City of New Castle, be it ordained by the Council of the City of New Castle Council as follows:

SECTION 1. City Code Chapter 230, Zoning, Section 230-1(B), is hereby amended by adding the following definitions to correspond with the alphabetical order of the current list of definitions (with deleted text identified by bold, strikethrough and new language identified by bold, underlined text):

B. Definitions

<u>Data Center. A facility, campus, building, or portion thereof designed, constructed, or operated primarily for the processing, storage, management, transmission, distribution, exchange, or recovery of digital data and information through networked computers, servers, storage systems, telecommunications equipment, networking equipment, and related information technology systems. A Data Center includes all supporting and integral infrastructure necessary for its operation, whether located within or outside the principal building, including but not limited to electrical infrastructure; utility service equipment; substations; transformers; switchgear; power distribution equipment; uninterruptible power supply (UPS) systems; battery energy storage systems; backup power systems; emergency generators and associated fuel storage and delivery systems; cooling, ventilation, air conditioning, heat rejection, and other mechanical systems; water and thermal management systems; fire detection and suppression systems; monitoring, security, and building management systems; telecommunications and fiber infrastructure; and all other ancillary equipment, structures, and improvements that support the continuous and reliable operation of the facility.</u>

SECTION 2. City Code, Chapter 230, Zoning, Section 230-25(B), Industrial District I, is hereby amended as follows (with deleted text identified by bold, strikethrough and new language identified by bold, underlined text):

B. The following uses may be permitted as special exceptions upon approval of the Board of Adjustment in accordance with §230-57B of this chapter:

<u>(3) Data Centers subject to the following requirements:</u>

<u>(a) No portion of any Data Center, including but not limited to, generators, cooling equipment, fuel storage, and equipment yards, may be located within one thousand (1,000) feet of:</u>

(i) **any residentially zoned property;**
(ii) **the property line of any lot with an existing school; or**
(iii) **the property line of any lot with an existing daycare center.**

(b) **Visual screening in the form of walls or fencing in combination with suitable landscaping must be provided to reasonably shield the Data Center use from residential property. The visual screening must consist of a combination of an opaque wall, decorative fence, berm, or similar screening feature together with a landscape buffer designed to provide effective year-round screening of buildings, equipment, parking areas, loading areas, and other operational features of the Data Center from adjacent properties. The landscape buffer must include a mixture of evergreen trees, evergreen shrubs, and other plant materials sufficient to create a continuous visual screen. Evergreen trees must be installed at a minimum height of five (5) feet at the time of planting and must be spaced no more than fifteen (15) feet on center, or an equivalent planting arrangement that achieves a continuous year-round screen within three (3) years of installation.**

(c) **Rooftop mechanical equipment and other appurtenances must be concealed or integrated within the roof such that they are not visible at ground level from public rights of way.**

(d) **The site must be designed to minimize noise impacts on neighboring residential properties, schools, and daycares through siting, shielding, enclosures, and comparable measures.**

(e) **A noise study prepared by a qualified acoustical professional must be submitted with the Board of Adjustment application. The study must establish pre-development ambient noise levels at the property lines of the nearest residential properties, schools, daycare centers, and other sensitive receptors; model operational noise from all Data Center equipment and infrastructure, including but not limited to**

backup generators, cooling systems, transformers, substations, electrical infrastructure, HVAC equipment, battery energy storage systems, and other mechanical or electrical equipment; and demonstrate that operation of the Data Center will not increase pre-development ambient noise levels by more than five (5) dBA at the nearest residential property lines, schools, daycare centers, or other sensitive receptors. The study must also include an analysis of low-frequency noise and infrasound generated by all major operational equipment, evaluating octave-band sound pressure levels between one (1) Hz and two hundred (200) Hz and identifying any tonal, impulsive, or periodic characteristics. The difference between unweighted sound levels (dBZ) and A-weighted sound levels (dBA) at the nearest residential property line, school, daycare center, or other sensitive receptor must not exceed twenty (20) decibels during normal operations unless the Applicant demonstrates through a site-specific acoustical analysis that no adverse impacts will occur. Any approved Data Center will be subject to a condition that the City may require post construction acoustical testing upon receiving a Complaint. Additionally, any approved Data Center will be subject to a condition that the City may require mitigation measures, including acoustical enclosures, barriers, equipment relocation, or operational limitations, as necessary to achieve compliance. Upon receipt of a documented complaint, the City may require post-construction acoustical testing by a qualified acoustical professional, and if noncompliance is demonstrated, the owner must implement corrective measures within a timeframe established by the City. Compliance with the standards of this section must be maintained throughout the operational life of the Data Center.

(f) The Board of Adjustment may require that a Data Center incorporate site-specific noise attenuation and screening techniques, including, but not limited to, the following measures:

(i) Vibration Isolation Systems: including spring and

neoprene isolators to prevent structure-borne noise transmission from generators, chillers, or HVAC units.

(ii) **Acoustic Fencing and Walls: solid or composite barriers with a minimum Sound Transmission Class (STC) rating of 35, extending fully above the line of sight between the noise source and the property line.**

(iii) **Building Orientation and Shielding: siting of Data Center buildings or equipment yards so that structures themselves block direct noise paths toward residential or institutional uses.**

(g) Water Supply. The Applicant must submit a written statement from the applicable water utility confirming the availability of adequate water service and sufficient system capacity to serve the proposed Data Center. The Applicant must also submit a Water Management Plan that identifies estimated daily and annual water consumption, projected peak demand, anticipated sources of supply, water conservation measures, cooling system efficiency, and contingency measures for drought conditions, utility interruptions, emergencies, or other events that could affect the availability of water resources. The Board of Adjustment may require modifications to the plan or additional mitigation measures to minimize impacts on public utilities and ensure the reliable and efficient operation of the facility.

(h) Energy Supply. The Applicant must submit a written statement from the applicable energy utility confirming the availability of adequate energy service and sufficient system capacity to serve the proposed Data Center. The Applicant must also submit an Energy Management Plan that identifies estimated daily and annual energy consumption, projected peak demand, anticipated sources of supply, energy conservation measures, and contingency measures for utility interruptions, emergencies, or other events that could affect the availability of energy resources. The Board of Adjustment may require modifications to the plan or additional mitigation measures to minimize impacts on public utilities and ensure the reliable and efficient operation of the facility.

(i) The Data Center must use closed-loop cooling systems for process and equipment cooling, meaning a system in which the cooling medium is contained within a piping circuit for reuse, unless the Board of Adjustment approves an open-loop or once-through cooling system after coordination with applicable state agencies.

(j) The Applicant must submit to the Board of Adjustment a written statement from the applicable local electric utility confirming either the availability of sufficient grid capacity to meet projected connected and peak loads for all proposed phases, or documentation of onsite generation meeting or exceeding projected load.

(k) Any emergency/backup generator must be certified Tier IV or utilize natural gas with Selective Catalytic Reduction (SCR) and achieve equivalent or superior emissions performance to Tier IV diesel. Emergency generators will be used solely to provide backup electrical power during utility outages, emergency conditions, routine maintenance, and periodic testing necessary to ensure operational readiness. Routine generator testing will be permitted only between the hours of 8:00 a.m. and 6:00 p.m. Emergency generators will not be used for peak shaving, demand response programs, commercial power generation, or the sale or export of electricity to the electrical grid, except as may be necessary during an actual utility outage or emergency event.

(l) The Applicant must submit a Hazardous Materials Management Plan with the Board of Adjustment application identifying all hazardous materials proposed to be stored, used, or maintained on the site, including but not limited to fuels, battery energy storage systems, batteries, refrigerants, lubricants, chemicals, and other regulated substances. The plan must identify the type, quantity, and location of each material and describe all storage, secondary containment, spill prevention, leak detection, fire protection, inspection, and emergency response procedures. The plan must also identify the facility's emergency response personnel, provide

current emergency contact information, and describe coordination procedures with local fire, emergency medical, and hazardous materials response agencies. The plan must be updated as necessary to reflect material changes in operations or hazardous material storage throughout the life of the facility.

(m) Site lighting must be designed and shielded to achieve zero (0) footcandles at adjoining residential property lines.

(n) Prior to the issuance of a Certificate of Occupancy, the owner or operator must submit a Decommissioning Plan for review and approval by the City. The plan must describe the procedures for the removal of all buildings, structures, equipment, generators, battery energy storage systems, fuel storage facilities, utility infrastructure, foundations, and other improvements associated with the Data Center upon permanent cessation of operations. The plan must also describe the restoration of the site to a safe and stable condition in accordance with applicable federal, state, and local laws and regulations. The approved Decommissioning Plan must be maintained throughout the operational life of the facility and updated as necessary to reflect significant modifications to the site or its operations.

(o) Approved Data Centers must include a condition requiring that the operator or property owner must provide written notice to the Building Official whenever the use approved as a Special Exception under this Section is out of active operation for more than six (6) consecutive months. Should the Data Center cease operation for a period of one (1) year or longer, it will be deemed abandoned, unless the Board of Adjustment grants an extension for good cause shown. After a Data Center is deemed abandoned, the operator or property owner must, within six (6) months, either: (a) schedule a site inspection with the Building Official to verify compliance with all applicable approval conditions and use requirements prior to recommencing operation; or (b) remove all equipment, structures, systems, and associated improvements related to the approved Data Center use in

accordance with a decommissioning plan approved by the Building Official.

(p) Periodic Compliance Report. As a condition of Board of Adjustment approval, the owner or operator must submit a compliance report to the City every three (3) years following the issuance of the Certificate of Occupancy. The report must be prepared and certified by the owner or operator and must demonstrate continued compliance with all applicable conditions of approval and the requirements of this Ordinance, including but not limited to noise, lighting, landscaping and screening, hazardous materials management, operational conditions, generator operation, electric consumption, water consumption, and any other applicable Special Use Permit or Board of Adjustment conditions. The City may require supporting documentation, studies, or inspections as necessary to verify compliance.

~~(3)~~ **(4)** Any other industrial use not permitted as a matter of right which would not violate the intent of this section.

SECTION 3. If any section, subsection, sentence, clause or phrase of this Ordinance is for any reason held to be unconstitutional or invalid, such decision shall not affect the validity of the remaining portions of this Ordinance.

SECTION 4. In the event any existing Ordinances or parts of Ordinances are in conflict herewith, the provisions of this Ordinance shall control.

SECTION 5: This Ordinance shall become effective immediately upon its adoption by City Council.

First Reading

Second Reading

Signed this ________day of ________, 2026

____________________, President of City Council

Attest: ______________________________
, City Clerk

Approved: ________________________
, Mayor

Prime Sponsor(s): Mr. Carter
Co-sponsor(s): Mr. Caneco, Mr. Tackett, Ms. Durham, Mr. Cartier
Date of Introduction: August 5, 2025

SUBSTITUTE NO. 3 TO ORDINANCE NO. 25-101
AS AMENDED BY FLOOR AMENDMENT NO. 3

TO AMEND *NEW CASTLE COUNTY CODE* CHAPTER 40 ("UNIFIED DEVELOPMENT CODE"), ARTICLE 3 ("USE REGULATIONS") AND ARTICLE 33 ("DEFINITIONS") TO ESTABLISH STANDARDS FOR THE REVIEW, SITING, AND OPERATIONS OF DATA CENTERS

WHEREAS, New Castle County currently lacks comprehensive, specific zoning and development regulations tailored to the unique infrastructure demands, environmental impacts, and operational characteristics of data centers; and

WHEREAS, data centers are distinct from other industrial and commercial uses due to their high energy consumption, large-scale backup power systems, significant water usage, and potential impacts on surrounding communities through noise, heat, and traffic; and

WHEREAS, technological advancements have precipitated the growth of data centers nationwide, expediting the need for provisions to guide responsible development and mitigate impacts; and

WHEREAS, it is in the public interest to proactively regulate data center development to ensure compatibility with surrounding land uses and to protect the health, safety, and welfare of County residents; and

WHEREAS, data centers share many characteristics with heavy industry, particularly in terms of energy consumption, building scale, and noise leading to zoning and classification complexities best handled as a separate category of use; and

WHEREAS, the County seeks to encourage economic development while also preserving community character, promoting sustainability, safeguarding public infrastructure capacity, and minimizing land use conflicts; and

WHEREAS, data centers are unique land uses that involve large-scale electrical infrastructure, water-intensive cooling systems, and extensive impervious coverage; and

WHEREAS, data centers typically utilize backup diesel generators, aboveground and underground fuel storage tanks, cooling towers, and large volumes of electronic and battery equipment that may pose environmental, public health, and safety risks if left abandoned or unmaintained; and

WHEREAS, advances in information technology, server design, and power delivery may cause individual data center facilities to become obsolete within a relatively short operating lifespan, resulting in the risk of abandonment of specialized facilities that are not easily adapted to other uses; and

WHEREAS, the County has a responsibility to protect the public health, safety, and welfare of its citizens by ensuring that abandoned or obsolete facilities are properly removed and that sites are restored to a safe and stable condition; and

WHEREAS, financial assurance requirements for decommissioning are a well-established land use tool in other industries with significant environmental or infrastructure impacts, including telecommunications towers, wind and solar energy projects, and mining operations; and

WHEREAS, requiring a decommissioning plan and associated financial assurance for data centers will ensure that the County and its taxpayers are not burdened with the cost of restoring sites and remediating environmental impacts in the event of abandonment; and

WHEREAS, County Council has determined that the provisions of this Ordinance substantially advance, and are reasonably and rationally related to legitimate government interests.

NOW, THEREFORE, THE COUNTY OF NEW CASTLE HEREBY ORDAINS:

Section 1. *New Castle County Code* Chapter 40 ("Unified Development Code"), Article 3 ("Use Regulations"), Division 40.03.110 ("Use Regulations") is hereby amended by adding the material that is underscored and deleting the material that is stricken, as set forth below.

Sec. 40.03.110. – Use table.

Table 40.03.110 lists the type of uses permitted in each district. Refer to Division 40.33.200 for definitions of the various use categories and their components. Should a use not be identified in Table 40.03.110, an interpretation pursuant to Section 40.31.520 can be requested. Table 40.03.522 lists the specific parking requirements for several general uses listed in Table 40.03.110. Limited and special use standards are contained in Table 40.03.210 and Division 40.03.300. *See* Article 4 for district, lot, bulk and other standards.

Table 40.03.110A. General Use Table									Table 40.03.110B. General Use Table								Table 40.03.110C. General Use Table	
Zoning District (Urban and Suburban-Transition Character) Y=permitted, N=prohibited, L=limited review, S=special use review, A=accessory									Zoning District (Suburban and Special Character)								Additional Standards (all districts)	
Land Use	TN *	ST	MM	ON	OR	CR	BP	I	CN	S ***	SE	NC **	HI	EX	SR	P	Parking	Limited & Special Use Standards
Agricultural									**Agricultural**								**Agricultural** …	
Industrial Uses									**Industrial Uses**								**Industrial Uses**	
Compost operations	N	N	N	N	N	N	N	L	N	N	N	N	L	L	S	N	Table 40.03.522	Section 40.03.336
Data Centers	N	N	N	N	N	N	N	S	N	N	N	N	S	S	N	N	Table 40.03.522	Table 40.03.210 Section 40.03.343
Extraction	N	N	N	N	N	N	N	N	N	N	N	N	N	L	N	N	Table 40.03.522	Section 40.03.324 and Table 40.03.210 …

Section 2. *New Castle County Code* Chapter 40 ("Unified Development Code"), Article 3 ("Use Regulations"), Section 40.03.210 ("Buffering, location, bulk and scale standards") is hereby amended by adding the material that is underscored and deleting the material that is stricken, as set forth below.

Sec. 40.03.210. - Buffering, location, bulk and scale standards.

Table 40.03.210 provides standards that increase the degree of buffering, govern the location, and reduce the bulk or scale of limited or special uses. The meaning of the different categories contained in the Table is established as follows:

A. *Location.* …

K. *Other.* …

Table 40.03.210A. Limited and Special Use Standards						Table 40.03.210B. Limited and Special Use Standards						Table 40.03.210C. Limited and Special Use Standards
	Buffer Standards					Locational or Dimensional Standards						Additional Standards
	Bufferyards			Separation		Minimum				Maximum		
Land Use	Location	Increase Opacity or Width	Fence or Berm	Use Protected	Distance from Use to protected parcel	Spacing	Site Area	Road Access	Open Space	Lot Area	Height	Other
Agricultural …						**Agricultural**						**Agricultural** …
Industrial Uses						**Industrial Uses**						**Industrial Uses**
Compost operations												Section 40.03.336
Data Centers	Within 1,000 feet of a residentially zoned parcel or parcel with an established daycare center or school	0.4	Minimum 8 ft. berm	Residential and established daycare center and schools	1,000 ft.*							See Section 40.03.343 * See Section 40.03.343.B
Extraction	All	0.8 (150 ft.)	8 ft. berm	Residential	100 ft. *		25 ac.					Section 40.03.324 * 100 ft. for Recreation, High Intensity Uses …

Section 3. *New Castle County Code* Chapter 40 ("Unified Development Code"), Article 3 ("Use Regulations"), Division 40.03.300 ("Additional limited and special use standards") is hereby amended by adding the material that is underscored and deleting the material that is stricken, as set forth below.

Sec. 40.03.333. - Light Industry uses in OR and CR zoning districts.

A. Light Industry uses as defined in Subsection 40.33.270.E ~~D~~ are considered Limited Uses in the OR zoning district, except that the following uses require a special use permit:

Sec. 40.03.340. - Multifamily conversion.

Multifamily conversions are intended to facilitate and encourage the redevelopment of existing office and commercial sites and increase the diversity in housing opportunities in the county where supporting infrastructure exists. The following requirements apply to all multifamily conversions:

A. A multifamily conversion is not permitted on a parcel that is within one thousand (1,000) feet of a heavy industry zoned parcel or a parcel with a heavy industry use as defined in Subsection 40.33.270.~~C~~ D.

Sec. 40.03.343. – Data Centers.

In addition to the requirements of Section 40.31.430, the following requirements apply to every data center development or the cumulative expansion of an existing data center in excess of 20,000 square feet:

A. The site shall be designed to minimize noise impact on neighboring residential properties and shall not generate noise levels that exceed the predevelopment noise level. A noise study shall be submitted to evaluate this impact and to verify consistency with Chapter 22 of the *New Castle County Code*.

B. No portion of the proposed data center building shall be located within 1,000 feet of a residential district. The Department may approve data center structures located less than 1,000 feet, but no closer than 500 feet from a residential district, provided that noise levels generated will not exceed the allowable noise levels permitted in this section. The Department will require submission of a site-specific noise study demonstrating that the combined effectiveness of the selected measures will ensure compliance with this section.

C. Except for lot lines along an existing office or industrial park or adjacent to an existing industrial use, the bufferyards, street trees, berming and landscaping shall be upgraded, where appropriate to mask the development. The buildings shall meet the requirements for large industrial structures in Section 40.26.260.

D. In addition to meeting all the requirements of Article 11, anticipated trip generation for a data center shall be based on the latest published ITE manual calculations. Data centers shall not be converted to any use that exceeds the approved Transportation Improvement Study (TIS) or equivalent study, for the site unless another TIS or equivalent study is completed and approved in accordance with all other applicable provisions of this Chapter.

E. Open-loop or once-through cooling is prohibited unless reclaimed water is used.

F. All generators shall be Tier IV or other technologies that the applicant can demonstrate as being more efficient and have less impact on the environment.

G. Any energy storage system proposed as part of the data center should be based upon best practices for energy storage and efficiency and have less impact on the environment. Energy storage systems must employ fire-safe technologies as determined and approved by the State Fire Marshall's Office.

H. All new or expanded facilities shall submit documentation outlining the operation of the building as a LEED certified building.

I. All data centers must reserve an area for future integration with onsite or communal microgrid components that include various energy generation sources, energy storage systems, a distribution network, a control system, and metering/billing systems.

J. Abandonment.

1. The operator or property owner shall provide written notice to the Department whenever the data center is out of active operation for more than 6 months. Data centers that cease to operate for 1 year or longer will be considered abandoned.

2. The operator or property owner shall either recommence operations and schedule a site inspection with the Department to verify that all use requirements are still satisfied or shall remove all equipment and systems within 6 months of being considered abandoned.

3. Within 2 years of being considered abandoned, the data center site shall be either restored to its predevelopment condition and inspected by the Department, or the property owner shall submit a land use application consistent with this Chapter to establish a new use for the site. Failure to comply with the requirements of this section shall authorize, but not require, the County to remove the data center and, when abandoned, restore the site to its predevelopment condition and charge the property owner as set forth in Subsection K.

K. Financial assurance. Prior to issuance of a building permit, the operator or property owner shall provide a bond, surety, letter of credit, or other financial assurance in a form and amount acceptable to the Department to secure payment of 100% of the anticipated cost of removal of all associated site improvements and restoration of the site to its predevelopment condition when the data center has been abandoned. Financial assurances for data centers are subject to the following:

1. The operator or property owner shall provide a decommissioning report that includes the following:

   a. An inventory of components installed on the site.

   b. Expected lifetime of the facility.

   c. Decommissioning sequence.

   d. Decommissioning cost estimate based on RSMeans or other software or methodology approved by the Department.

2. The financial assurance shall remain in full force and effect while the data c enter remains in place and operating, and shall be renewed every 5 years and revised as necessary to adjust for inflation and any change in the cost for decommissioning.

L. Annual Report, Enforcement and Compliance. The operator or property owner of the data center shall supply an annual report to the Department indicating compliance with all applicable local, state and federal regulations.

Section 4. New Castle County Code Chapter 40 (“Unified Development Code”), Article 33 (“Definitions”), Section 40.33.270 (“Industrial Uses”) is hereby amended by adding the material that is underscored and deleting the material that is stricken, as set forth below.

A. Compost operations (NAICS 325314). ...

B. *Data center*. A facility composed of networked computers, servers, storage systems, and telecommunications equipment used for data processing, storage, and distribution (NAICS 518210).

C ~~*B*~~. *Extraction.* ...

D ~~*C*~~. *Heavy industry.* …

E ~~D~~ . *Light industry.* …

F ~~E~~ . *Recycling or .* …

G ~~F~~ . *Solar energy system, large scale.* …

H ~~G~~ . *Utility, major.* …

I ~~H~~ . *Utility, minor.* …

J ~~I~~ . *Utilities, maintenance facilities.* …

Section 5. *New Castle County Code* Chapter 40 ("Unified Development Code"), Article 33 ("Definitions"), Section 40.33.300 ("General Definitions") is hereby amended by adding the material that is underscored and deleting the material that is stricken, as set forth below.

Abandonment. …

Closed-loop cooling system. A closed-loop cooling system is one where the cooling medium is contained within a piping circuit for reuse.

Open-loop cooling system or *once-through cooling system*. An open-loop cooling system is one with a cooling method in which water is withdrawn from a surface, groundwater, or public supply source, passed through heat-exchange equipment to absorb heat, and the effluent is then discharged to a surface water body, re-injected into the ground, or discharged into a wastewater treatment system. Open-loop systems typically rely on continuous withdrawals and discharges, frequently resulting in greater thermal or environmental impacts including higher consumptive use of water resources.

Reclaimed water. Water that has been treated to meet specific quality standards, making it safe for a beneficial reuse. This can include treated effluent from wastewater treatment facilities and remediation activities.

Section 6. Consistent with Comprehensive Development Plan. New Castle County Council finds that the provisions of this Ordinance are consistent with the spirit and intent of the New Castle County Comprehensive Development Plan.

Section 7. Inconsistent Ordinances and Resolutions Repealed. All ordinances or parts of ordinances and all resolutions or parts of resolutions that may be in conflict herewith are hereby repealed except to the extent they remain applicable to land use matters reviewed under previous Code provisions as provided in Chapter 40 of the *New Castle County Code*.

Section 8. Severability. The provisions of this Ordinance shall be severable. If any provision of this Ordinance is found by any court of competent jurisdiction to be unconstitutional or void, the remaining provisions of this Ordinance shall remain valid, unless the court finds that that the valid provisions of this Ordinance are so essentially and inseparably connected with, and so dependent upon, the unconstitutional or void provision that it cannot be presumed that County Council would have enacted the remaining valid provisions, without the unconstitutional or void one, or unless the court finds that the remaining valid provisions, standing alone, are incomplete and incapable of being executed in accordance with the County Council's intent. If any provision of this Ordinance or any zoning map or portion thereof is found to be unconstitutional or void, all applicable former ordinances, resolutions, zoning maps or portions thereof shall become applicable and shall be considered as continuations thereof and not as new enactments regardless if severability is possible.

Section 9. Effective Date. This Ordinance shall become effective immediately upon its adoption by County Council and approval by the County Executive or as otherwise provided in 9 *Del. C*. § 1156 and shall only apply to Land Use applications submitted after such date(s) unless the applicant, by written request, agrees to submit to the provisions of the Ordinance.

Adopted by County Council of
New Castle County on: 3/10/26

____
Monique Williams-Johns
President of County Council
of New Castle County

Approved on:

____
Marcus A. Henry
County Executive
of New Castle County

SYNOPSIS: The following is a summary of the revisions that were proposed in the O rdinance, broken down by *New Castle County Code* Division or Section number.

Sec. 40.03.110. The amendment updates the use table to establish which zoning districts permit the data center as a Special Use.

Sec. 40.03.210. The amendment updates the Limited and Special Use Standards Table to include data center uses.

Sec. 40.03.333. The amendment updates references to the definition of Light Industrial uses based on changes proposed to Section 40.33.270.

Sec. 40.03.340. The amendment updates references to the definition of Heavy Industry uses based on changes proposed to Section 40.33.270.

Sec. 40.03.343. The amendment includes Special Use standards for data center uses.

Section 40.33.270. The amendment defined data center uses.

Substitute No. 1. The following changes were included in Substitute No. 1:

Sec. 40.03.110. The amendment further updated the use table to establish which zoning districts permit the data center as a Special Use.

Sec. 40.03.343. The amendment clarified that data centers must comply with State and Federal regulations and allows the Department to permit applicants to substitute other technologies that the applicant has demonstrated as being more efficient and have less impact on the environment. The amendment also included decommissioning, abandonment, and financial assurance provisions for data centers.

Section 40.33.300. The amendment defines closed-loop cooling systems, open-loop c ooling systems, and reclaimed water.

Substitute No. 2. The following changes were included in Substitute No. 2:

Sec. 40.03.210. The amendment referenced Section 40.03.343.B regarding the distancing requirements between data center structures and protected uses.

Sec. 40.03.343. The amendment clarified that noise level restrictions are for those generated by the data center. The amendment also allowed the Department to approve data center structures located less than 1,000 ft. from a residential zoning district, provided that noise levels generated do not exceed the allowable noise levels permitted in Chapter 22 of the *New Castle County Code*. However, no portion of the proposed data center building or accessory uses are permitted within 500 ft. of a residential district.

The amendment changed the effective date of the ordinance to be the date that the ordinance is adopted by County Council and approved by the County Executive.

Substitute No. 3. The following changes are included in Substitute No.3:

Section 40.03.343 has been amended as follows:

1. Clarifies the noise level standards
2. Removes the suggested noise attenuation techniques
3. Removes the Lighting standards since they are found in other Sections
4. Removes the natural gas system requirement for generators
5. Removes the Battery Energy Storage System requirement
6. Requires new or expanded buildings to be LEEDS certified
7. Clarifies areas for onsite or communal microgrid components must be reserved, not designed
8. Removes Sections for State coordination, compliance with State and Federal laws
9. Makes formatting changes to correct subsection citations
10. Modifies the Annual Report requirement

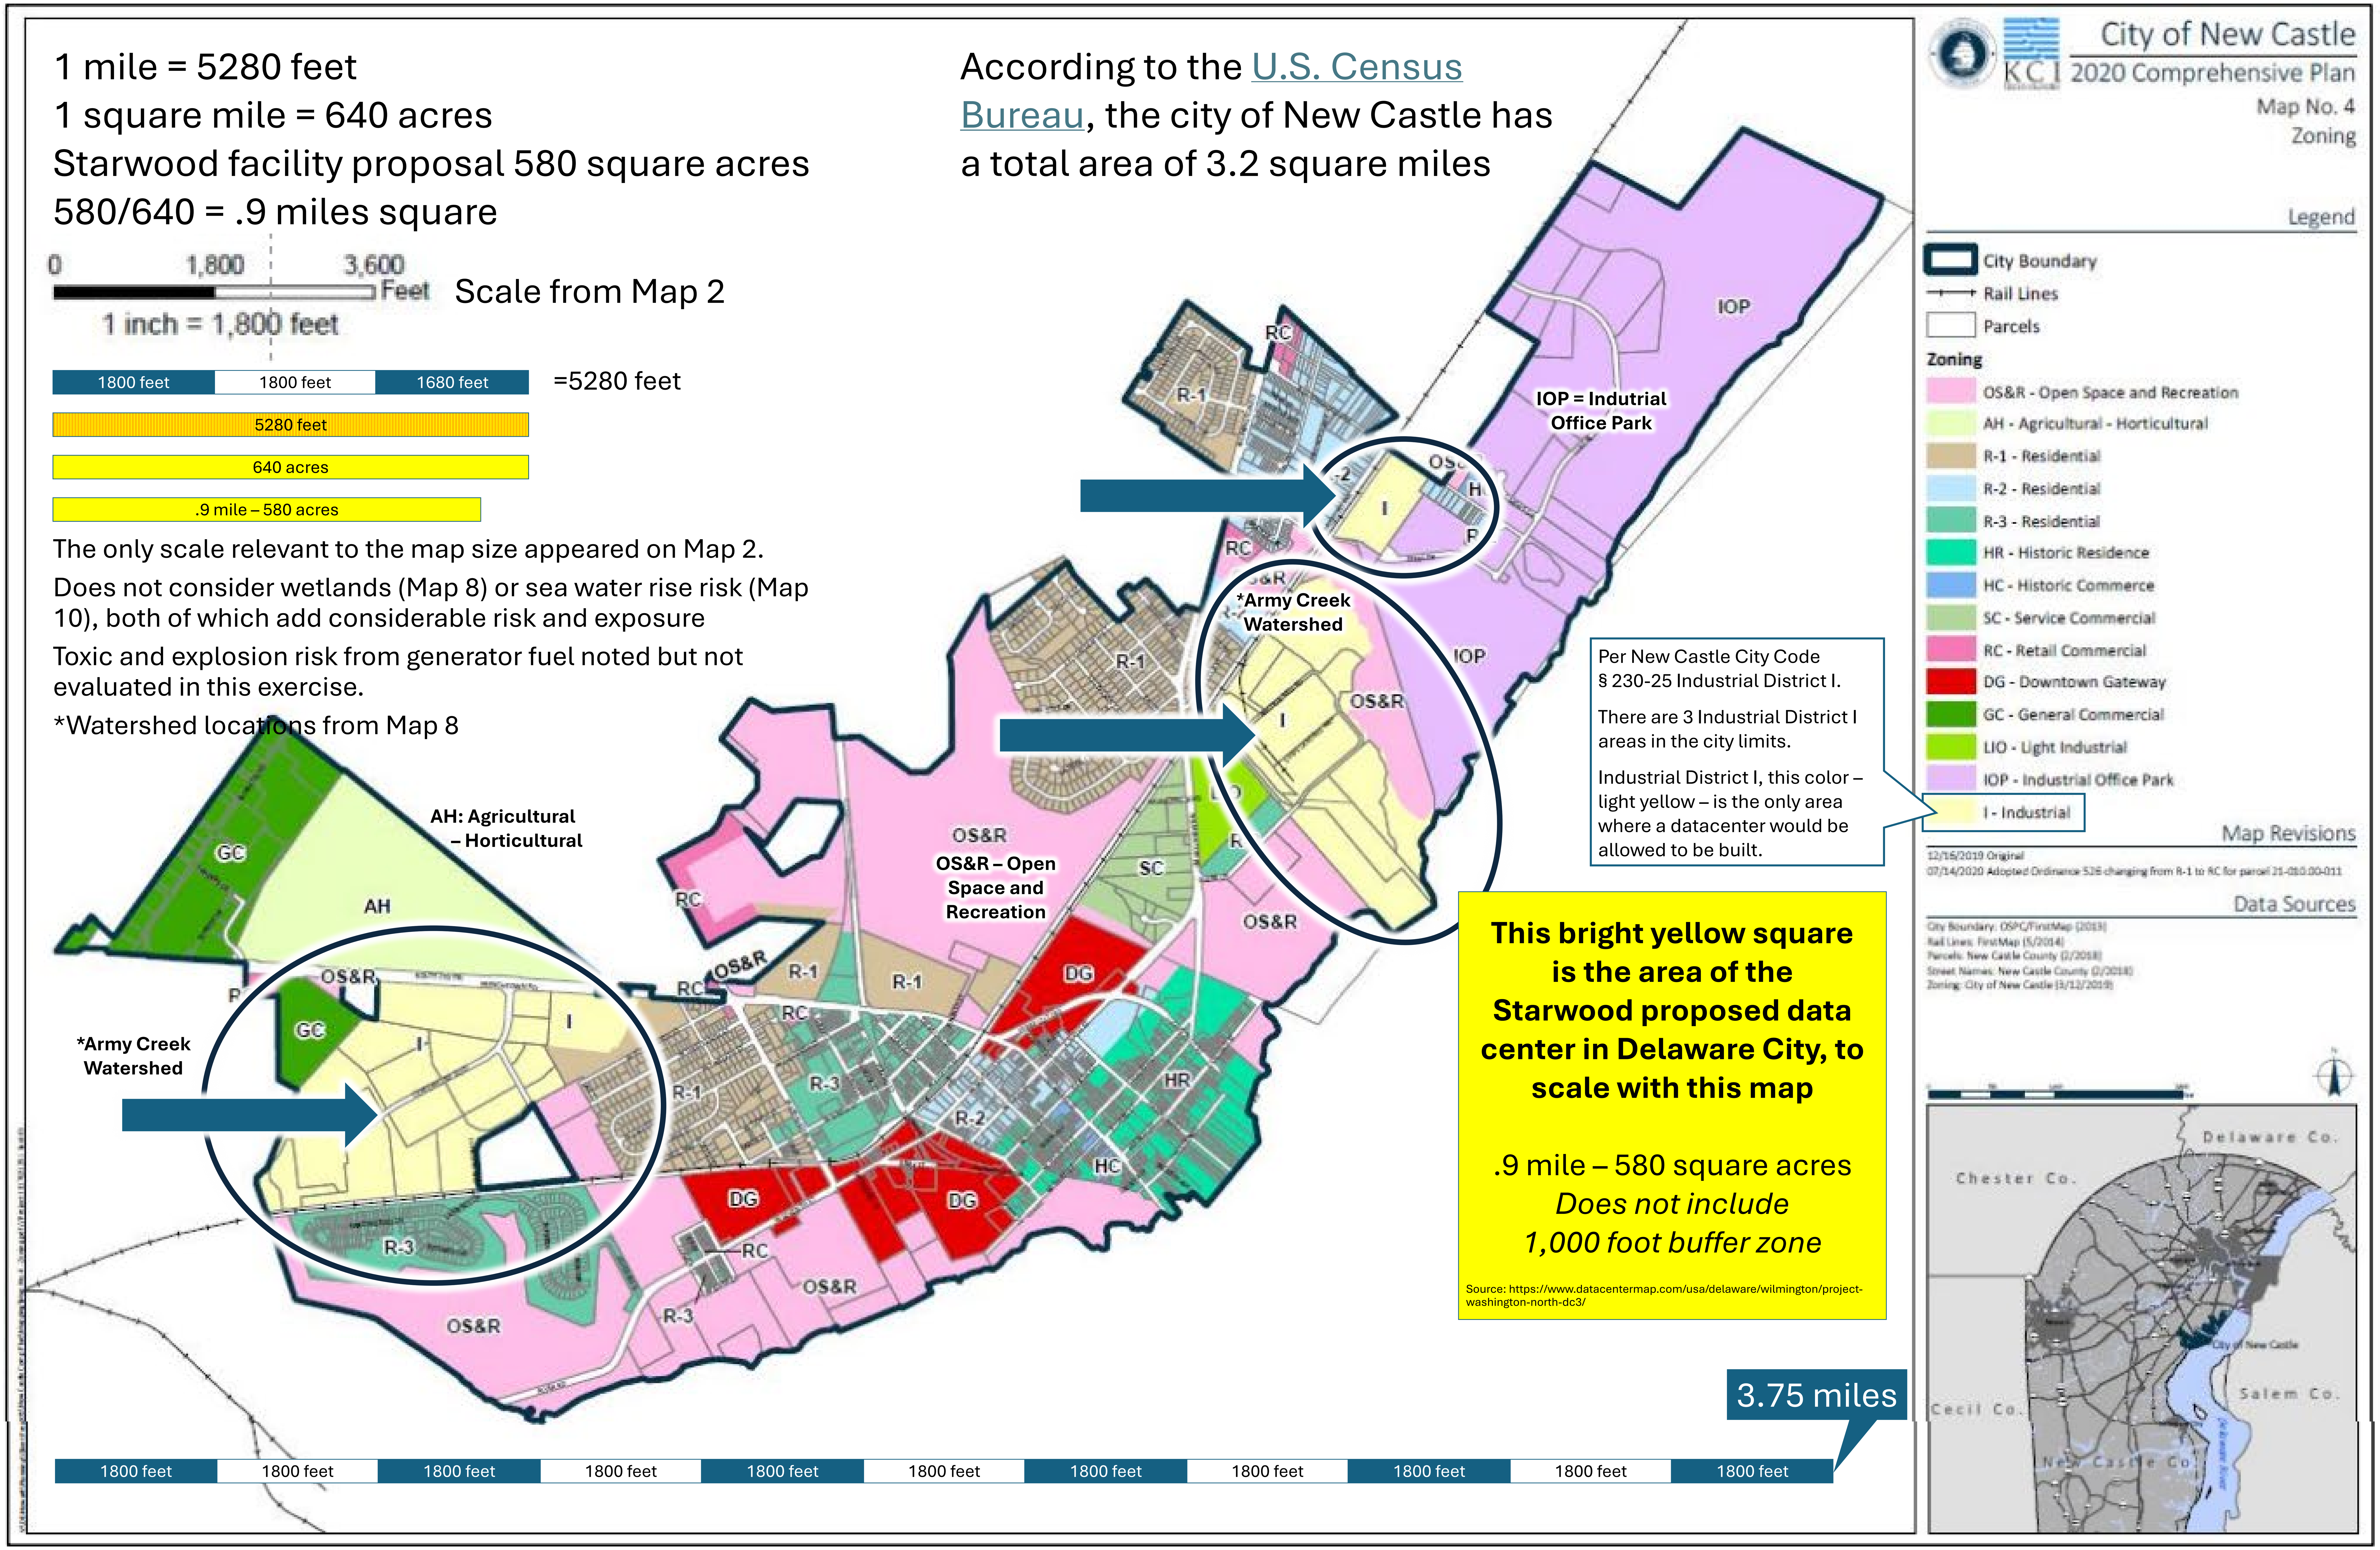

Annotated 2026-08-26 with additional materials by C. Batty

If the Board of Adjustment were allowed to rezone for a data center

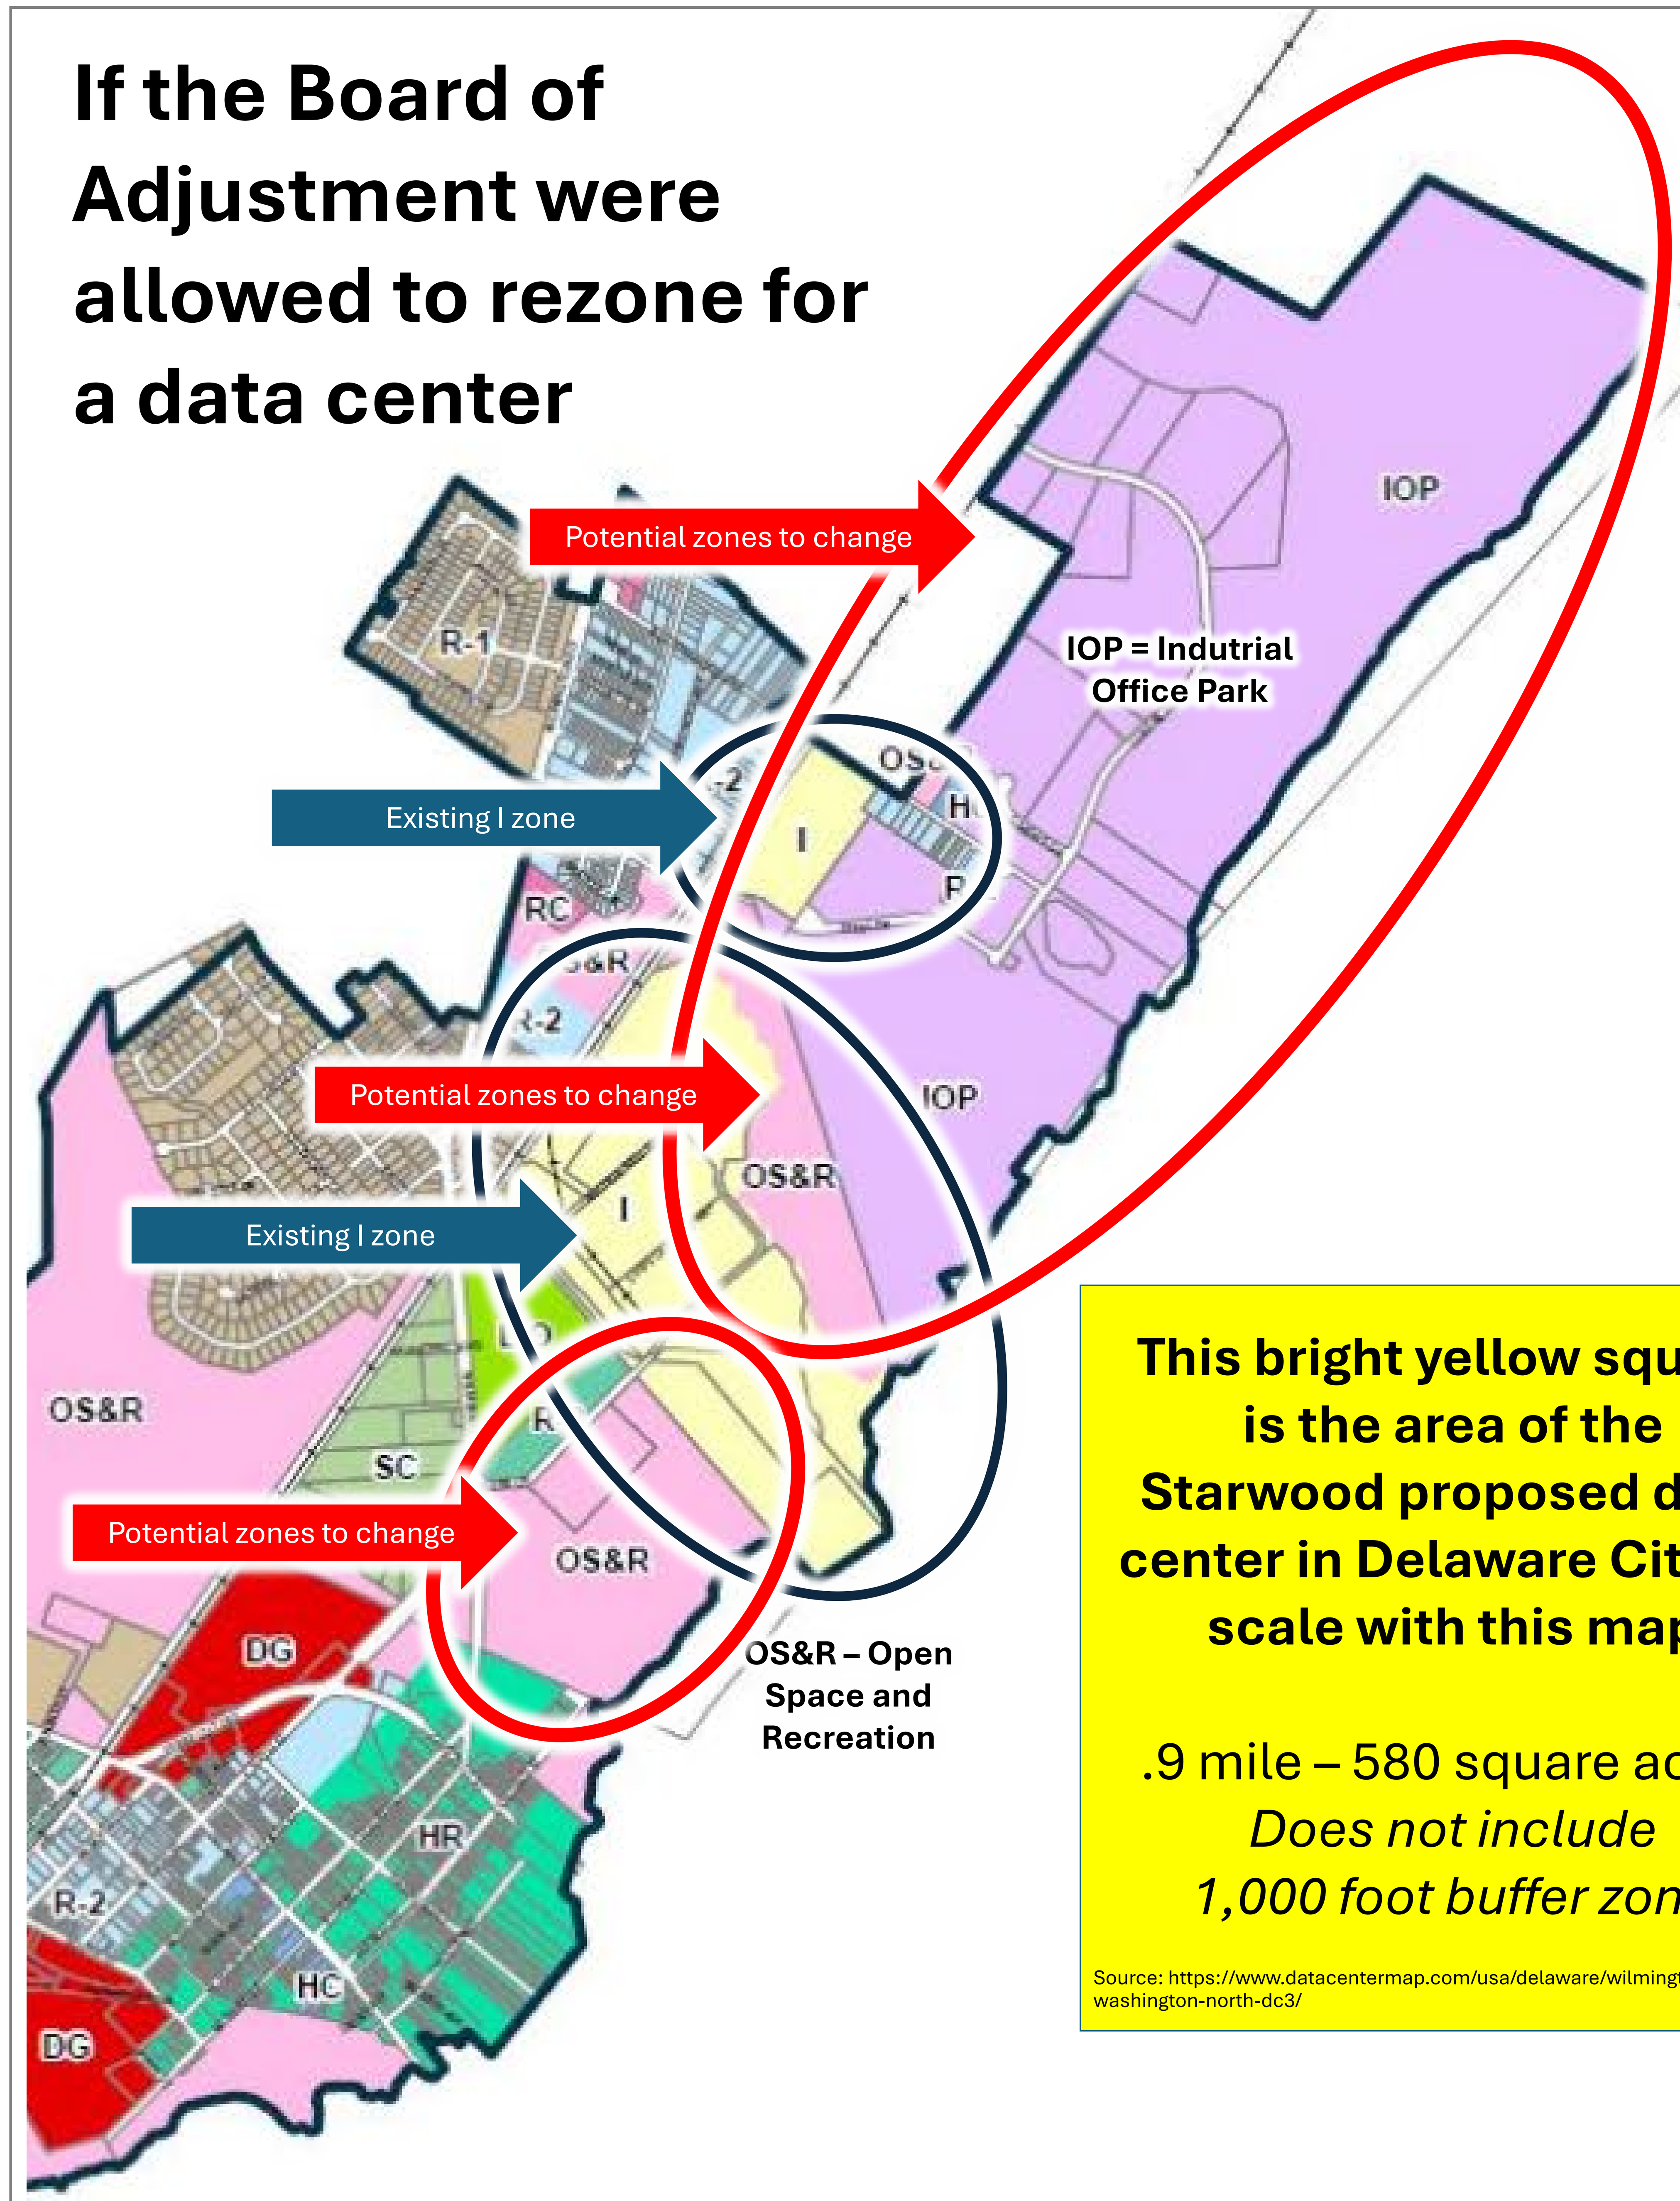

This bright yellow square is the area of the Starwood proposed data center in Delaware City, to scale with this map

.9 mile – 580 square acres
Does not include
1,000 foot buffer zone

Source: https://www.datacentermap.com/usa/delaware/wilmington/project-washington-north-dc3/

Data Center Noise Analysis
By Dr. David Gilbert, D.Sc.
As of 06/14/2026

Introduction:

Data Centers and the construction thereof, are causing increased concern on local residences living by or near such facilities. These concerns include, but not limited to, power and water usage, heat produced, air quality and noise levels produced from such facilities. This synthesis paper looks at the effects of noise, specifically low frequency noise, produced by such facilities.

Problem: The problem is people living near or by a data center perceive low frequency, high decibel noise will negatively affect both health and property values.

Argument: The closer residences are to data center noise sources the greater the negative impact on both residency health and property values.

Research Questions:

1. At what distances and rate do low frequency, high decibel sound waves, like those produced by a data center, travel before they become non-impacting to human physiology?
2. At what distances and rate do low frequency, high decibel sound waves, like those produced by a data center, travel before they become non-impacting to property values?

Limitations include the use of secondary data and information from sources that may not be consider as authoritative. In such cases, multiple sources have been noted, whenever possible, in order to make the facts presented both true and strong. This does not replace the need for primary or more formal research to occur.

Disclaimer – The information provided in this paper is for general informational and educational purposes only. While every effort has been made to ensure the accuracy of the content, the author makes no representations or warranties of any kind, express or implied, regarding the completeness, accuracy, or reliability of the information provided. In addition, the views, thoughts, and opinions expressed in this paper are those of the author alone. The information provided is for informational purposes only and should not be construed as professional, legal, or financial advice. Readers should consult with a qualified professional before making any decisions based on the content.

Methodology:

The methodology of research includes the use of relevant secondary data and information gathered from both authoritative and non-authoritative sources through the use of the internet and with the assistance of AI for gathering and conformation purposes only and when applicably. The data and information gathered was reviewed, sorted, and analyzed using established and accepted academic research methodologies as best can be applied given the type data and information. Then organized and presented in a manner also consistent with accepted academic research methodologies. AI has not been quoted or cited.

Background:

Sound as a phenomenon is defined as a mechanical disturbance that propagates through a medium. A sound wave is the measurement of that disturbance in terms of frequency, wavelength, and amplitude. [1,2,6] Where frequency is the number of times a sound vibrates per second and wavelength is the distance between each wave height. See Figure 1. As such, there is an inverse relationship between frequency and wavelength. Meaning the higher the frequency the short the wavelength. Therefore, high frequency sounds have shorter wavelengths and low frequency sounds have longer wavelengths. See Figure 2.

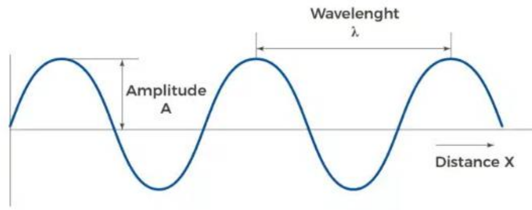

Figure 1: Amplitude, Wavelength. [1]

Frequency, for humans is between 20 and 20,000 Hz [5,7], See Figure 3, and is organized into three basic categories low, medium and high. Where:

- Low Frequency 20 – 500 Hz: Have long, wide wavelengths carrying immense energy. These wavelengths can easily bend around obstacles, pass through walls, and travel vast distances depending on medium. [3,7]
- Medium Frequency 500 – 2,000 Hz: Are at the core of Human/Animal/ Insects detection. Medium frequency sounds can travel up to 2 miles depending on frequency, amplitude, and medium. [3,7]
- High Frequency 2,000 – 20,000 Hz: Have short, closely packed wavelengths that don't bend easily around solid objects, are easily blocked and absorb quickly into the air. High frequency sounds don't travel far less than 1000 feet depending on medium and volume. [3,7]

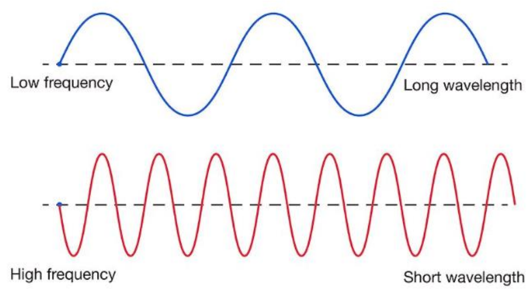

Figure 2: Long and Short Wavelengths. [2]

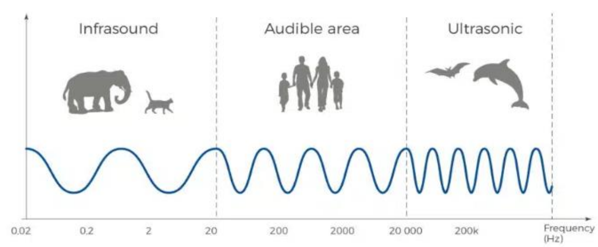

Figure 3: Human Low to High Frequency Range.

Amplitude is the measurement of energy contained within the soundwave, its intensity, i.e. loudness, measured in decibels (dB). [8,9,10] Soundwaves with a higher amplitude carry more energy thus create a louder sound.

- 10 dB: Sound of a pin drop.
- 20 dB: Barely perceptible sound, registering just slightly above the threshold of human hearing
- 30 dB: A quiet library or whispering.
- 60 dB: Normal conversation.
- 85 dB: The threshold for potential hearing damage with prolonged exposure.
- 120 dB: A live rock concert or a siren, which can cause immediate physical discomfort or pain.

Normal back ground noise range between 30 dB in the wilderness to 85 dB in cities and other urban areas. With a normal conversation ranging between 60 – 70 decibels. [10,11] Continued exposure of over 80 dB will cause hearing loss. [11,12] The volume (dB) and the length of exposure gives indication as to how harmful the noise is. “In general, the louder the noise, the less time required before hearing loss will occur.” [13] See Table 1. OSHA requires hearing protection for 8-hour exposure above 85 dBA [17]

Noise Level dBA	Maximum Exposure Time	Noise Level dBA	Maximum Exposure Time
85	8 hours	106	3.7 minutes
88	4 hours	109	1.85 minutes
91	2 hours	112-118	Less than 1 minute
94	1 hour	121-124	Less than 10 seconds
97	30 minutes	127	1 second
100	15 minutes	130-140	Less than 1 second
103	7.5 minutes	140+	NO EXPOSURE TIME

MDHearing

Table 1: Noise Exposure [13]

Data Center: A data center is a centralized physical facility, generally operating 24/7, that securely houses vast amounts of computing equipment and infrastructure support systems` used to facilitate Cloud Computing, Web Hosting, and Artificial Intelligence. [14, 15] See Figure 4.

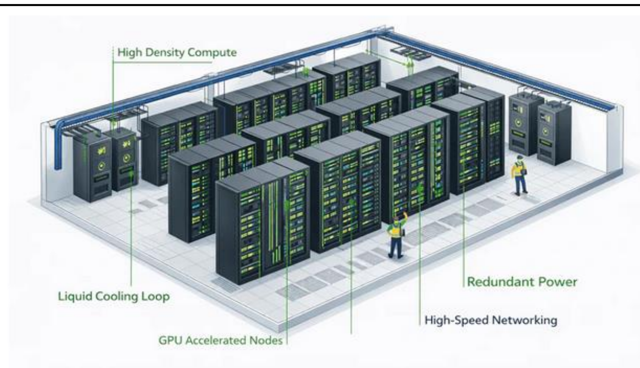

Figure 4: Data Center [15]

There are 4 basic types of data centers: Enterprise Data Centers – owned and operated by a single organization strictly for their own use and generally located in or near a company headquarters [43,44,45]; Colocation (Colo) Data Centers – shared buildings where multiple organizations rent space, power, and cooling systems housing their own servers and networking systems [48,49,50]; Hyperscale Data Centers – large scale facilities typically located in or near a major city occupying thousands of square feet containing thousands of severs and miles of connecting equipment. Some these hyperscale data centers can be as large as 60,000 square feet [41,43,44,45] housing as many as 10,000 server racks with a capacity in excess of 80 megawatts (MW) [51]; and Edge Data Centers – smaller data centers ranging from a single server rack to as many as a 100 generally operating in the 500 kilowatts (kW) to 2 megawatts (MW) range. [51] Within each of these data centers are a large number of servers, storage systems, and networking equipment all configured within server racks. [40,41,42]

There are two basic types of server racks. The standard 42U server rack containing about 35 individual servers and the 48U server rack typically housing about 43 servers. Each individual server within a server rack can produce 1000 to 3,000 BTUs of heat per hour, with an average of 1600 BTUs per hour. Each server rack can produce up to 70,000 BTUs per hour. [47,48]

1600 BTUs per Server per Hour x 43 Servers per Server Rack = 68,800 BTUs/Hour (70,000 rounded)

Each server rack on average occupies a space of approximately 6 sq. ft. on a foot print of 3 by 2ft. Meaning, a 60,000 sq. ft. Hyperscale Data Center may contain as many as 10,000 server racks each producing up to 70,000 BTUs/Hour resulting in temperatures over 200°F without an effective cooling system.

In order to maintain proper cooling temperatures most data centers use backward-curved centrifugal cooling fans. The average output for one of these fans, in terms of airflow, is approximately 3000 CFMs. Using a rule of thumb for data center cooling of 1 to 1.5 CFMs per square foot. This equates to needing one cooling fan for every 2000 sq. ft. of data center space. When in operation each these cooling fans can produce an overall noise level of approximately 85 decibels at an average frequency between 80 and 100 Hz. [18,21,23] Thus, a 60,000 sq. ft. data center will require at least 30 backward-curved centrifugal cooling fans to maintain an ASHRAE-recommended temperature between 64 – 81°F.

A Data Center, when in operation, produces sound frequencies between 20 – 20,000 Hz but generally in the range of 60 to 250 Hz with an approximate average of 80 to 100 Hz. These low frequencies are the "dominant characteristic of data center" [19] putting out significant amounts of low frequency sounds [22] ranging from 60 decibels to as high 100 decibels or more [21,24] depending the number of cooling fans and/or HVAC units. Most of which are located on the data center's roof, allowing for low-frequency sound waves to travel greater distances unrestricted; impacting nearby communities [20] including those at much as 3 miles away particularly at night when overall ambient noise decreases and cooler temperatures [24] provide a more favorable medium in which to transmit. [32,33,34]

Discussion:
Data Center: As noted above a Data Center's most dominant characteristic is the generation of low frequency sound waves, 80 to 100 Hz, at high decibel levels; 85 dBs or greater. This frequency places data center noise directly in line within a human beings' low frequency hear range as shown in table 2.

	Low Frequency						Medium Frequency	
Frequency Hz	0	50	100	150	200	250	300	350
Data Center								
Human								

Table 2: Data Center, Human Low Frequency Range

With additional cooling fans decibel levels can increase to over 100 dBs on a logarithmic scale. Meaning that for each doubling of a data center cooling fan the number of decibels increases by 3. [26,27,28] Therefore, there is little difference between having one fan or 100. See Figure 5.

No. of Fans	Decibels
1	85
2	88
4	91
8	94
16	97
32	100
64	103
128	106

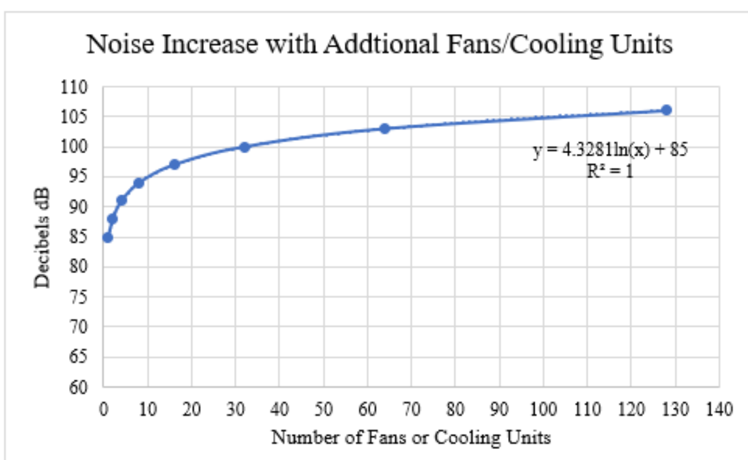

See Figure 5: Additional Cooling Fan Effects

Low frequent sound dissipation, in normal atmospheric conditions, is also logarithmic dropping approximately 6 decibels for each doubling of feet traveled (See Inverse Square Law). [4,29,31]

100 Hz Cooling Units (Fans) in Decibels				
Feet	1 Unit	30 or More	Subjective	Comments
0	85	101	Very Load	Hearing Protection Required
3	79	95		
7	73	89		
13	67	83		
26	61	77	Loud	Disruptive/ Maybe Painful
52	55	71		
105	49	65	Moderate	Normal Conversational Range
210	43	59		
420	37	53		
840	31	47		
1680	25	41	Faint	Quiet/ Whispering
3360	19	35		
6719	13	29	Very Faint	Barely Perceptible
13438	7	23		

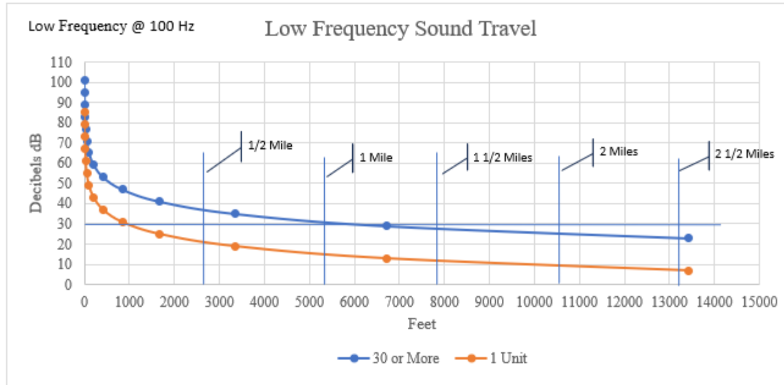

Figure 6: Low Frequency Sound Travel

As shown figure 6, a low frequency 100 Hz sound beginning at 85 decibels, equivalent to a data center with 1 cooling fan, would most likely dissipate to the very faint level (barely perceptible 30 dBs or lower) after traveling approximately 1,680 ft or about ⅓ of a mile. A low frequency 100 Hz sound beginning at 101 decibels, equivalent to a data center with 30 or more cooling fans or units, would most likely dissipate to the very faint level (30 dBs or lower) after traveling approximately 6,719 ft or about 1¼ miles.

Human Effects: Low frequency, high decibel sound waves can cause serious health issues, especially to those closer to the sound source. Health issue reported include, but not limited to: increased level of anxiety and depression, increased stress levels, hypertension, headaches & migraines, reduced memory, concentration, sleep deprivation, and reduced ability to work or exercise. [16,34,35,36,37]

Land Value Effects: Low frequency, high decibel sound waves, like those produced by data centers, may have a negative impact on established neighboring residential land and property values especially those living in close proximity, a mile or less, although there is little published data to support it.

In fact, one such study that actual refutes that proposition conducted by George Mason University's Schar School of Policy and Government found homes, whether single-family residences, townhomes, or condos, closer to a data centers actually sold for more. Reason; data centers are typically built in areas with strong infrastructure, including good roads, reliable utilities, and easy access to jobs and airports, all of which are features that make nearby neighborhoods more desirable. [38]

Although, according to documentation published by Own Luxury Homes ® [39] any resident within a mile of a data center may experience as much as an 18% discount. Those between 1 and 2 miles of a data center may actually see a 1 to 4% increase in property value. Those between 2 and 5 miles may see a 2% increase in value and residencies greater than five miles are unaffected with home values set at current market value. See Table 3.

Distance from Facility	OLH Designation	Noise Level (est.)	Visual Impact	Employment Premium	Value Impact	OLH Guidance
0–500 ft	Immediate Adjacency	55–65 dB continuous	High — direct view	Minimal	Discount 8–18%	Avoid for residential purchase. Industrial neighbor discount is permanent.
500 ft–0.25 mi	Near Adjacency	40–55 dB	Moderate	Low	Discount 3–8%	Due diligence required. Verify cooling system type (air vs water cooled) and noise barrier.
0.25–0.5 mi	Transition Zone	Under 40 dB	Low–None	Low–Moderate	Neutral ±1%	Standard due diligence. Verify no expansion planned within 0.25 mi.
0.5–1 mi	OLH Proximity Sweet Spot	None	None	Moderate	Neutral to +2%	Optimal range. Employment benefit without adjacency discount.
1–2 mi	Employment Catchment	None	None	High	+1–4%	Strong position. Captures DC employer demand; residential character unaffected.
2–5 mi	Corridor Catchment	None	None	Moderate	Neutral to +2%	Broad employment benefit. Useful for investor analysis of employment-driven appreciation.
5+ mi	General Metro	None	None	Low	Market rate	No specific DC proximity effect. Standard market analysis applies.

OLH Data Center Proximity Valuation Framework, May 2026. Distance ranges and discount/premium estimates are approximate and market-specific. Northern Virginia markets have higher discount/premium sensitivity due to density of DC facilities. Noise levels are estimates based on EPA industrial facility guidance and published hyperscale facility specifications. Value impact estimates based on OLH analysis of comparable sales data in Northern Virginia, Phoenix, and Columbus DC corridors.

Table 3: Residential Home Value Impact. [39]

When above information is plotted in relationship to low frequency sound travel it appears that homes that within a mile of the data center may experience negative home values and the closer to the data center the more negative the value becomes. See Figure 7.

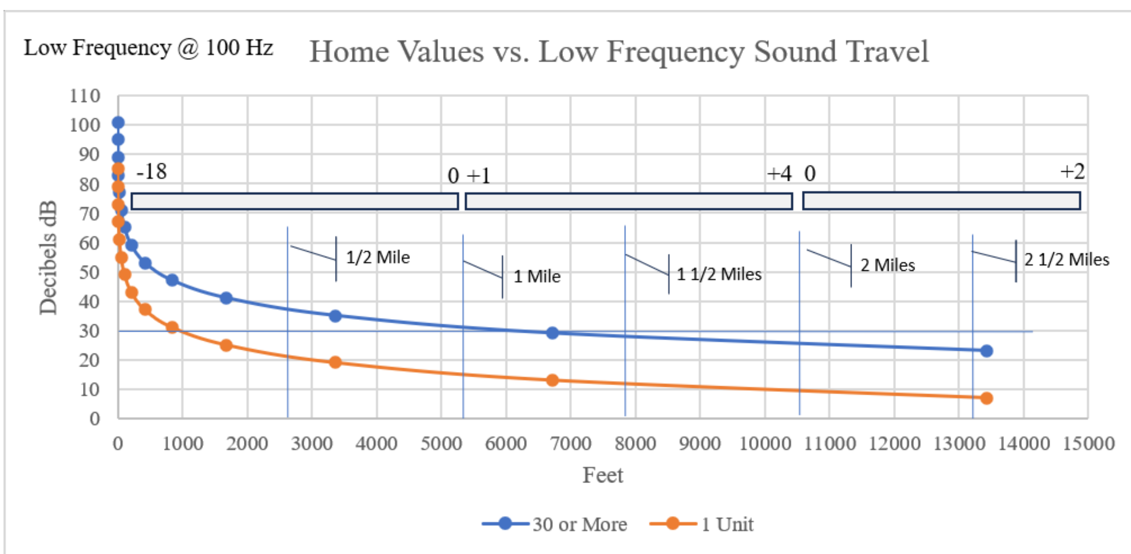

Figure 7: Home Values vs. Low Frequency Sound Travel

Homes between one and two miles away from a data center may see a positive increase in home values, as much a 4% the closer the home is to the two-mile mark benefitting from a what is called the employment catchment area defined as a geographic area from which the data center draws its workforce. After the two-mile mark and before the five-mile mark, home values may increase in value only about 2% due to the home being in what is known as the corridor catchment area. Past the five-mile mark local market rates apply.

Conclusion:
From the data and information provided a data center when in operation can generate a significant amount of low frequency sound waves ranging between 60 and 250 Hz [18,21,23] with an average of approximately 80 – 100 Hz [21,25] at decibel levels of 85 dB or greater depending on number of cooling fans or units applied.

Unlike high frequency sound waves that generally travel short distances and are easily blocked by barriers such as walls or other structures. [3,7] Low frequency sound waves can travel great distances and cannot be easily by the blocked by walls or other structures. [3,7] Low frequency, high decibel sound waves, like those produced by data centers, can travel up to 1½ miles before sound dissipation falls below the 30 dB level depending on the environmental medium.

Those living within 1¼ miles and moving towards a data center, consisting of 30 or more cooling fans producing low-frequency 100 Hz sound waves at 100 dB will most likely experience ever increasing hearing loss, increased health issues, along with decreasing property values as much as an 18% or more.

Those living greater than 1¼ miles and moving outwards from a data center, consisting of 30 or more cooling fans producing low-frequency 100 Hz sound waves at 100 dB will most likely not experience the same health effects as those living within 1¼ miles and may find property values increase.

Therefore, for those residing within 1¼ miles of a data center the argument "the closer residences are to data center noise sources the greater the negative impact on both residency health and property values" is true and the problem statement "people living near or by a data center perceive low frequency, high decibel noise will negatively affect both health and property values" is verified.

References:

[1] https://svantek.com/academy/sound-wave/
[2] https://sciencephotogallery.com/featured/wavelength-and-frequency-science-photo-library.html
[3] https://acousticalsolutions.com/sound-frequency-chart?srsltid=AfmBOorYLMZMtae9QICCQmZLk9SNhFAJXdUfK2H5baVHxRxY6iib6WDQ
[4] https://avcsstechworld.com/the-inverse-square-law/
[5] https://www.hearingyourbest.com/blog/how-human-hearing-stacks-up-in-the-animal-kingdom/
[6] https://mecanum.com/frequency-wave-length/
[7] https://www.soundproofcow.com/difference-high-middle-low-frequency-noise/?srsltid=AfmBOoqoLl3moD5vueQDWZbO9_WblH0GdD_yZEWlGTLNsah4h4iQiZy-
[8] https://int.livhospital.com/what-decibel-level-is-harmful-to-your-hearing/
[9] https://www.johndcook.com/blog/2008/03/11/how-loud-is-the-evidence/
[10] https://ia.cpuc.ca.gov/environment/info/aspen/cresseygallo/fmnd/5-12_noise.pdf
[11] https://ehs.yale.edu/sites/default/files/files/decibel-level-chart.pdf
[12] https://noiseawareness.org/info-center/common-noise-levels/
[13] https://www.mdhearingaid.com/blog/decibel-chart/?srsltid=AfmBOooWSrVC2IoVMchJa2WM9O7dbOnZumFk2EKlPs8PTTfv71BAj5yy
[14] https://aws.amazon.com/what-is/data-center/
[15] https://www.supermicro.com/en/glossary/data-center

[16] https://www.eesi.org/articles/view/communities-are-raising-noise-pollution-concernsabout-data-centers
[17] https://acousticalsolutions.com/data-center-noise-pollution?srsltid=AfmBOoquM3pBoQAtDN3x-6Y0jDAYxCO0QcPkgpqzKsPoiZIti1XU9eAG
[18] https://ketchumandwalton.com/noise-control-in-data-centers-what-it-is-why-it-matters-and-how-to-get-it-right/
[19] https://noisemonitoringservices.com/data-center-noise-control/
[20] https://www.datacenterknowledge.com/data-storage/what-are-the-5-main-causes-of-noise-in-data-centers-
[21] https://www.lsars.com/data-center-noise
[22] https://lab.acentech.com/2025/11/06/good-noise-neighbor/
[23] https://energylinkinternational.com/areas-of-focus/noise-management-solutions/data-centers/
[24] https://pragmaticenvironmentalistofnewyork.blog/wp-content/uploads/2025/10/physical-attributes-of-a-typical-loudoun-county-v.pdf
[25] https://www.environmentalhealthproject.org/post/the-dangers-of-data-centers
[26] https://www.sonitron.be/combined-sound-sources-calculator/
[27] https://www.ussi.com/acoustical-rules-and-demystifying-decibel.php
[28] https://www.firgelliauto.com/blogs/calculators/noise-pollution-db-addition-calculator
[29] https://isostore.com/research/understanding-the-inverse-square-law/?srsltid=AfmBOor1zi3V4Rsat84vJL7jFsqsLOkhM2mxcz9-7QamIvkNEtLnUBcn
[30] https://www.mdpi.com/2076-3417/10/15/5205
[31] https://www.firgelliauto.com/blogs/engineering-calculators/decibel-distance-calculator-sound-level?srsltid=AfmBOoqd5ohpVLXIOvW9fIcGFd1p6EhLjDEsi_N4SYDbqmvKyAsP_RVJ
[32] https://www.peacefulpeculiar.org/uploads/1/5/0/3/150368424/health_issues_46196868514666.pdf
[33] https://www.sierraclub.org/sites/default/files/2025-07/data-center-factsheet-4.pdf
[34] https://www.pwcva.gov/assets/2024-06/240605%20Presentation%20to%20DCOAG-JMT%20%28Lyver%29.pdf
[35] https://www.epa.gov/archive/epa/aboutepa/epa-identifies-noise-levels-affecting-health-and-welfare.html
[36] https://lailluminator.com/2026/06/09/data-centers-3/
[37] https://gerrymcgovern.com/data-centers-are-noisy-and-smelly/
[38] https://www.housebeautiful.com/design-inspiration/real-estate/a71253889/data-center-property-values/
[39] https://www.ownluxuryhomes.com/markets/ai/guide/data-center-adjacent-residential-values
[40] https://encoradvisors.com/data-center-hardware/
[41] https://www.fortinet.com/resources/cyberglossary/data-center
[42] https://datacenteruniversity.be/whats-inside-a-data-center/
[43] https://www.ibm.com/think/topics/data-centers
[44] https://www.acceldata.io/blog/data-center-types-enterprise-to-edge-explained
[45] https://www.nlyte.com/faqs/what-is-a-data-center/
[46] https://www.jll.com/en-us/insights/why-smaller-data-centers-are-taking-off
[47] https://www.serverroomenvironments.co.uk/blog/how-to-calculate-server-room-air-conditioner-sizes?srsltid=AfmBOoo4mekIlTDw05rgymxlyvbnpG6d14uqQgziUCliQN3NxHApKhwN
[48] https://www.techtarget.com/searchdatacenter/tip/Data-center-cooling-optimization-in-the-virtualized-server-world.
[49] https://scispace.com/pdf/low-frequency-noise-and-its-main-effects-on-human-health-a-1mi8gdwyfc.pdf
[50] https://pmc.ncbi.nlm.nih.gov/articles/PMC9008144/
[51] https://www.tri-cityaudiology.com/hearing-loss-articles/sounds-can-have-a-bigger-impact-on-your-health-than-you-know/
[52] https://pmc.ncbi.nlm.nih.gov/articles/PMC10001830/

The Local Economic Impact of Data Centers

Dany Bahar

Brown University, Providence, RI 02912, USA.
dany_bahar@brown.edu

Greg C. Wright*

University of California, Merced, 5200 North Lake Road, Merced, CA 95343, USA.
gwright4@ucmerced.edu
*Corresponding author.

August 7, 2026

Abstract

A hyperscale data center can cost more than a billion dollars while employing only a few dozen people. Using a registry of 341 hyperscale facilities, we test whether this investment generates local spillovers through labor pooling, supplier linkages, or firm clustering and find little evidence that it does. Nighttime lights rise sharply at the parcel when construction begins, and the increase fades within five kilometers. Announced projects of similar scale that were never built show no comparable change. Advertised salaries and new-firm registrations do not rise, and we can rule out increases above 5.5 and 1.6 percent. Hiring and supplier estimates are also near zero but less precise. County data show increases in data-processing employment and establishments, which include the facility itself, but no consistent growth across related industries. Rent estimates are positive but imprecise. Compute-using firms locate near data centers, but most were already there before the facility opened.

Keywords: data centers, agglomeration, economic geography, place-based investment.

JEL Codes: R11, R23, O33, H71.

1 Introduction

The physical infrastructure behind digital services has become a major component of US capital investment. Data centers consumed roughly four percent of US electricity in 2023, and the rise of artificial intelligence has accelerated new construction [Shehabi et al., 2024, International Energy Agency, 2025]. A hyperscale campus can cost more than a billion dollars and draw hundreds of megawatts of continuous power, yet it may employ only a few dozen people on site [Alvarez et al., 2026]. At least thirty-five states now offer tax incentives specifically for data centers, with cumulative subsidies approaching twenty billion dollars.[1] These subsidies rest partly on the same promise of local economic gains that motivates place-based incentives more broadly [Slattery and Zidar, 2020, Kline and Moretti, 2014]. We study whether data-center entry produces those gains.

The natural benchmark comes from the literature on large plant openings. Greenstone et al. [2010] compare counties that won a new plant with runner-up counties that lost it and find that incumbent establishments become more productive in the winning counties. They attribute the gain to a deeper shared labor pool, denser supplier linkages, and knowledge spillovers among nearby firms. From a policy perspective, this productivity gain helps motivate place-based investments. At the same time, later work shows that winner-versus-loser estimates are sensitive to the choice of comparison counties [Patrick, 2016], while the average large plant produces little evidence of productivity spillovers [Patrick and Partridge, 2022]. Plant openings therefore do not provide a settled benchmark, even though manufacturing offers the strongest traditional case for agglomeration.

Data centers separate the location of capital from many of the workers, suppliers, and customers that use it. This contrast lets us ask whether a large investment generates local spillovers. Labor pooling, local supplier linkages, and knowledge spillovers can all make proximity productive [Duranton and Puga, 2004, Ellison et al., 2010], and these channels allow a large employer to generate a local multiplier [Moretti, 2010]. A data center may engage each channel only weakly. It employs dozens of people rather than thousands, purchases many inputs from regional or national suppliers, and delivers computing services to distant users without requiring them to locate nearby. The scale of the investment may therefore reveal little about the size of its local spillovers.

Data centers do not locate at random. Operators screen sites on power, land, and fiber, so places that host a facility differ from places that do not in ways that also shape local outcomes. A credible instrument for this siting decision is difficult to find. We therefore organize the analysis around two complementary facility-level designs. The within-site design compares changes near a completed facility with changes farther from the same site. At completed sites, we date construction from satellite-detected land clearing and restrict the capital analysis to facilities without an earlier data center within two kilometers. This comparison absorbs characteristics shared by the surrounding area, and we use it only when the outcome has a flat pre-trend. The built-versus-near-miss design compares completed facilities with 84 announced projects of hyperscale size that were never built, 52 of which have satellite coverage. These projects crossed an observable announcement and siting

[1] Authors' tally from the Good Jobs First Subsidy Tracker [Good Jobs First, 2024], which records state and local data-center subsidy awards and is one of the sources for our facility registry (Section 2).

threshold, although they tend to be newer, farther from existing data-center clusters, and less likely to have a named hyperscale sponsor. We therefore do not treat them as otherwise identical counterfactuals. Finally, county panels provide descriptive evidence on industry outcomes, although pre-period trends limit what those panels can establish.

Construction visibly transforms the site. At isolated sites, nighttime lights rise by 38 percent within the first kilometer when construction begins, an increase that fades to near zero by five kilometers [Henderson et al., 2012]. Daytime imagery confirms that vegetation clears at the same time, while announced projects that were not built show no increase in lights. The primary question is whether this activity spreads beyond the parcel.

We find little evidence of propagation through these channels. For labor pooling, advertised salaries do not rise, with an upper confidence bound of 5.5 percent. Hiring is less precisely estimated, and the job-posting interval permits an increase of up to 17 percent. Workplace employment begins rising before opening rather than breaking at the event. Supplier postings rise by 5.9 percent, but their confidence interval ranges from a 15 percent decline to a 32 percent increase, leaving the supplier-linkage evidence inconclusive.

For firm clustering and potential knowledge spillovers, we do not observe knowledge transmission directly. New-firm registrations and business applications do not rise, with upper confidence bounds of 1.6 and 5.7 percent. Compute-using firms locate closer to data centers, but 69 percent predate the nearest facility, suggesting shared metropolitan advantages rather than a cluster created by the data center. County data reinforce this distinction. Data-processing employment rises by 26 percent and establishments by 27 percent, but related industries show no consistent response.

Other local outcomes may adjust even without a broad worker or firm response. Nearby rents may rise by a few percent, but the estimates are imprecise. We therefore treat the rent evidence as suggestive. Neither multifamily permitting nor net migration rises. School property-tax revenue and county highway spending have positive point estimates, but neither is statistically significant after correcting for the six fiscal outcomes we test. We find no robust change in foot traffic or commercial spending, and residential electricity prices do not rise. The evidence leaves open modest effects on rents or public revenue, but not broad adjustment in population, construction, consumer activity, or electricity prices.

Taken together, these results distinguish a large local investment from a broad local response. Existing county studies find positive effects on employment, establishments, income, construction, and house prices [Alvarez et al., 2026, Yue and Zeng, 2026]. Our finer data separate activity at the facility from changes among nearby workers, suppliers, and firms and establish whether compute users arrived before or after the data center. At this scale, the sites are transformed, but we find little evidence of a new local cluster. This matters as US place-based investment shifts from labor-intensive plants to labor-light digital capital.

2 Data

Hyperscale data-center registry. We construct a master registry of US data centers from several sources. We begin with the Open Source Data Center Atlas, produced by Pacific Northwest National Laboratory's IM3 project, which identifies 1,242 facilities from OpenStreetMap footprints tagged as data centers and provides coordinates along with square footage and operator when available. We supplement the atlas with operating dates from operator press releases, state economic-development announcements, the Good Jobs First subsidy tracker, and trade publications. Targeted searches of state records, specialty wholesale operators, artificial intelligence operators, and real-estate investment trust filings add facilities that are missing from the atlas. The resulting master registry contains 1,494 facility records. Our analysis begins with the 341 records that we classify as hyperscale based on operator identity. This category includes Amazon, Microsoft, Google, Meta, Apple, Oracle, and a hand-curated group of large specialty and artificial intelligence campuses. We do not impose a megawatt threshold because capacity is not consistently observed. Of the 341 hyperscale records, 334 have an analysis date within the 2012–2024 VIIRS window and 315 have a high- or medium-confidence date.

For the built-versus-near-miss comparison, we separately assemble 84 announced projects that were not built. Of these, 52 have the coordinates and satellite coverage needed for the nighttime-lights design, and 49 are sufficiently far from operating campuses to enter the isolated capital sample. Section 3 describes the sponsors of these projects and shows that the construction-year estimate is 72 percent when we restrict the control group to projects with a named hyperscale tenant, compared with 70 percent in the full isolated sample.

Table 1: Hyperscale sample construction

Sample restriction	N
Panel A: Hyperscale registry	
Hyperscale facility records	341
with an analysis date in 2012–2024	334
with high- or medium-confidence date	315
with satellite-detected construction date	87
Panel B: Capital-footprint design	
Announced hyperscale projects that were not built	84
with coordinates and satellite coverage	52
Isolated, construction-dated built sites	**43**
Isolated announced near-miss sites	**49**

Notes: The 341 hyperscale records are drawn from the broader 1,494-record registry and include the six major cloud operators and hand-classified large specialty and artificial-intelligence campuses. High- and medium-confidence dates exclude projected openings, unverified planned-completion dates, and unresolved source conflicts. The capital design also requires a satellite-detected construction date and no earlier data center within two kilometers.

Nighttime lights. We extract monthly mean radiance from the Visible Infrared Imaging Radiometer Suite (VIIRS) Day/Night Band Monthly Composite from April 2012 through December 2025. Around each facility, we measure radiance in concentric rings of 0–1, 1–2, 2–5, 5–10, 10–25, and 25–50 kilometers and aggregate the monthly observations to annual means. We define a facility as isolated when no data center within two kilometers opened earlier, and the headline capital-footprint estimate uses this isolated, construction-dated sample.

Local labor, housing, and spending panels. Several proprietary data products measure local activity around each site. WageScape provides job postings and advertised salaries by geocoded location [WageScape, 2022]. LinkUp Job Records provide employer-level postings used in the firm analysis [LinkUp, 2025]. RentHub Rental Data provide monthly median asking rents by ring [RentHub, 2022]. SafeGraph Spend Patterns measure consumer spending [SafeGraph, 2022], and Advan Weekly Patterns measure foot traffic [Advan Research, 2025]. We analyze rents, spending, and foot traffic in a within-site ring design dated to the operating year, with state-by-month fixed effects. The posting and salary panels extend to 150 kilometers, providing the comparison area used in the labor analysis (Section 3.2).

Coverage differs across these sources, so each analysis uses the largest subset of the 315 cleanly dated hyperscale records for which its outcome and design are available. As a robustness check, Appendix Table D.3 restricts the national facility-ring outcomes to a common universe of 125 hyperscale sites. The construction, two-state business-registration, and county analyses retain their design-specific samples because they require different dates or use different geographic units.

County employment. The county-level analysis uses employment and establishment counts from the Quarterly Census of Employment and Wages. We retain the Bureau of Labor Statistics disclosure flag because the public files code suppressed employment as zero even though the underlying value is unobserved. We impute these cells using disclosed state-industry and county-parent totals and report alternatives that leave them missing or allocate them using state totals alone. The analysis covers data processing (NAICS 518), telecommunications (517), web and information services (519), IT consulting (5415), and professional, scientific, and technical services (541). We end the 519 series in 2021 because the 2022 NAICS revision moved major components of that sector to 513 and 516. Section 5.5 reports the imputation, disclosure rates, and pre-trend diagnostics that limit causal interpretation of the county results.

Instrument-search inputs. The instrument search draws on predetermined county features, namely 345 kV transmission-substation density, the InterTubes long-haul fiber network [Durairajan et al., 2015], the 1980 county share of urban college population (NHGIS), station-level climate normals (CustomWeather), and the EIA-860 generation fleet (Appendix B).

2.1 Dating construction and operations

We preserve announcement, groundbreaking, construction, planned-completion, and operating dates as distinct events. Operating dates reported by operators or public agencies receive high confidence. Less detailed sources receive medium confidence when the year agrees with an independent registry or satellite evidence. We classify announcement-based projections, planned completion dates, and unresolved conflicts as low confidence and exclude them from analyses that require a reliable date. This process produces 315 high- or medium-confidence hyperscale dates, 19 low-confidence dates, and 7 facilities dated outside the 2012–2024 window.

For a smaller sample, we recover the start of construction from daytime satellite imagery. A hyperscale build begins with land clearing, which appears in Landsat and Sentinel-2 imagery as a sustained decline in the parcel's vegetation index, a rise in its built-up-surface index, or both. Vegetation remains stable before the detected construction year and then falls sharply. We validate this timing with Federal Aviation Administration obstruction-evaluation filings, which are recorded independently of the satellite data. The imagery sample contains 87 hyperscale facilities. The FAA records match 64 facilities within 500 meters and 119 within one kilometer (Section 3.5). Together, these independent dates support the satellite construction timing used in the capital design.

3 Research Design

Data centers do not locate at random. Operators screen locations on electricity prices, available grid capacity, large parcels of low-cost land, fiber proximity, and tax policy. Counties that attract a data center therefore differ from other counties in ways that may also affect the local economy, making a simple cross-county comparison difficult to interpret. Appendix B documents this selection by testing a range of predetermined siting predictors as potential instruments. Power infrastructure, industrial land, and fiber strongly predict siting, but each is also correlated with local industrial conditions. Climate and historical human capital pose fewer obvious exclusion concerns, yet neither predicts hyperscale locations.

We organize the analysis around two complementary facility-level designs. The within-site design compares changes close to a completed facility with changes farther from the same site, thereby absorbing time-invariant siting differences. The built-versus-near-miss design asks whether changes around completed facilities exceed those around announced projects that were not built. We report both where the data allow. Unlike the runner-up counties in Greenstone et al. [2010], the announced projects are not known finalists in the same site-selection process, so we do not assume that they are balanced with built sites. County panels separately describe aggregate industry changes, although suppression and pre-opening trends prevent us from interpreting most of those estimates as causal effects on related employment.

3.1 Announced near-miss facilities

An announced project that reached the siting stage but was never built provides an external comparison for a completed facility. Both groups crossed an observable announcement and siting threshold, although projects that proceed may still differ from those that do not. We identify 84 announced projects that were later cancelled and were large enough to qualify as hyperscale. The sample comes from trade publications, regulatory filings, and operator announcements. Of these, 52 projects have the coverage required for the nighttime-lights design. For each outcome, we estimate

$$y_{i,r,t} = \sum_k \left[\beta_k \, \mathbf{1}[t - t_i^c = k] \times \text{Near}_{i,r} \times \text{Built}_i + \delta_k \, \mathbf{1}[t - t_i^c = k] \times \text{Near}_{i,r}\right] + \alpha_{i,r} + \lambda_t + u_{i,r,t}, \quad (1)$$

where Built_i identifies completed facilities, t_i^c is the satellite-detected construction start for built sites and the announcement date for cancelled sites, $\alpha_{i,r}$ is a facility-by-ring fixed effect, and λ_t is a calendar-year fixed effect. In the capital analysis, $y_{i,r,t}$ is the log of mean radiance, and the path of β_k measures the change around built facilities relative to the change around cancelled sites. Timing is not symmetric across the two groups because cancelled projects lack an observed or consistently reported construction start. Announcements usually precede construction by at least a year, placing some announcement-related activity for built sites in their own pre-period and, if anything, attenuating the estimated break. This limitation is less consequential for the capital outcome because a parcel that is never built does not light up, regardless of why the project was cancelled.

The announced near misses also differ in sponsor commitment. Fourteen of the 84 projects had a named hyperscale operator at cancellation, 19 were sponsored by established data-center operators without a named tenant, and the remaining 51, roughly 60 percent, were speculative developments led by real-estate sponsors. Projects without a signed tenant or confirmed grid allocation are less informative about the counterfactual for a committed facility. Even so, restricting the capital comparison to the nine isolated projects with a named hyperscale operator produces a 72 percent increase in the groundbreaking year, compared with 70 percent in the full isolated sample. The named-sponsor subset has a somewhat less stable pre-period and is too small to support longer-run estimates. More generally, many cancellations followed local opposition, so the near-miss sites may be located in less development-friendly places.[2] These limitations prevent us from treating the near misses as balanced runner-up sites. We therefore use them alongside the within-site comparison below. Built sites are dated from satellite-detected land clearing, whereas near-miss timing comes from less reliable schedules in filings and trade publications. Independent regulatory filings validate the built-site dates (Section 3.5).

Table 2 describes these differences in the exact capital-footprint sample. Built and near-miss

[2] Because most near-miss projects were announced late in the sample, the near-miss group contributes few post-event observations. Coverage at event years zero through four is 48, 21, 5, 3, and 1 sites, so the built-versus-near-miss path at longer horizons is identified primarily from changes among built sites relative to the common pre-event trend rather than from post-period movement among near-miss sites. Section 4 quotes magnitudes only through the year after groundbreaking for this reason.

sites have similar local populations and access to high-voltage infrastructure, but they differ substantially in event year, baseline radiance, proximity to an earlier data center, region, and sponsor commitment. Facility-by-ring fixed effects absorb the level differences, while the pre-event path reveals whether the inner and outer rings were already diverging. The imbalance in timing remains a concern. The near-miss group contains only one to five projects per calendar year before 2020 and 21 to 49 thereafter, whereas 26 of the 43 built sites break ground in 2017 or 2018. Section 4 therefore reports results by construction cohort and shows that the within-site design, which does not use near-miss controls, reproduces the same cohort pattern. The observable imbalances prevent us from interpreting the near misses as conventionally balanced runner-up sites.

Table 2: Built and announced near-miss sites in the capital-footprint sample

Characteristic	Built mean	Near miss mean	Standardized difference
Event year	2018.81	2024.37	-4.21
Log population within 10 km	10.54	10.42	+0.06
Pre-event log radiance, 0–1 km	0.54	2.26	-1.28
Log distance to earlier data center	2.35	3.64	-1.19
Log distance to $\geq$230 kV substation	1.87	1.90	-0.04
Log distance to long-haul fiber	2.15	1.87	+0.29
Share Northeast	0.00	0.06	-0.36
Share Midwest	0.23	0.41	-0.38
Share South	0.35	0.43	-0.16
Share West	0.42	0.10	+0.76
Share with named hyperscale sponsor	1.00	0.18	+2.95

Notes: Exact isolated-site sample used at event year zero in Figure 1, namely 43 built sites and 49 announced near-miss sites. Built sites are dated by satellite-detected construction start. Near-miss sites are dated by announcement. The standardized difference is the difference in group means divided by the square root of the average of the two group variances, so a value of one means the groups differ by one pooled standard deviation. Distances are in kilometers before applying $\log(1 + x)$. Population is based on 2020 Census tract population centroids. High-voltage access uses substations rated at least 230 kV. Long-haul fiber uses the InterTubes conduit map. The table is a diagnostic, not a claim that the groups are balanced.

3.2 Within-site rings

The near-miss comparison is available when the same outcome can be reconstructed around never-built sites. For outcomes observed only around operating facilities, the within-site design compares an inner ring around a facility with the 25 to 50 kilometer ring around the same site,

$$\log(1 + \mathrm{NTL}_{i,r,t}) = \sum_k \beta_k \, \mathbf{1}[t - t_i^c = k] \times \mathrm{Near}_{i,r} + \alpha_{i,r} + \gamma_t + \varepsilon_{i,r,t}, \qquad (2)$$

where i indexes the facility, r the ring, t the year, and t_i^c the facility's construction-start year. The indicator $\text{Near}_{i,r}$ marks the inner ring, $\alpha_{i,r}$ is a facility-by-ring fixed effect, and γ_t is a calendar-year fixed effect. The facility-by-ring fixed effect absorbs any time-invariant difference between a site's inner and outer ring, including siting preferences, local industrial structure, and access to grid and fiber infrastructure. The coefficients are therefore identified by changes in the inner ring relative to changes in the outer ring of the same site.[3]

Two changes address the pre-trend present in a simpler version of the design. First, we date the event to construction rather than operations, so that activity during the building phase does not appear in the pre-period (Section 2.1). With this timing, inner-ring brightness is flat before groundbreaking and rises when construction begins. Second, the headline sample includes only isolated facilities with no earlier data center within two kilometers. A new facility built beside an existing campus starts with a brighter inner ring and can create a spurious pre-trend. In the isolated sample, the event study is flat before construction, rises by 38 percent in the construction year, and continues to build over the following years. A separate linear trend fitted to each site's inner-minus-outer path during the four pre-construction years produces the same 38 percent observed-minus-predicted change, with a campus-bootstrap 95 percent interval of 15 to 71 percent.

We interpret ring estimates only when the outcome is stable before construction. The capital footprint, rents, advertised wages, and firm entry satisfy this condition. Workplace employment does not, and we place no interpretive weight on its inner-versus-outer estimate. Labor and supplier effects may also extend beyond the immediate site. For advertised wages, job postings, and supplier postings, we therefore compare the pooled 0 to 25 kilometer area with locations 50 to 150 kilometers away. This geography allows for effects across a commuting zone while keeping the comparison area outside the immediate data-center cluster.

3.3 Demand-side outcomes

We apply the same within-site comparison in equation (2) to residential rents, consumer spending, and foot traffic, dating these outcomes to the operating year rather than the start of construction. Because the monthly panels span 2020 to 2023, we replace the calendar-year fixed effect with state-by-month fixed effects that absorb differences across states in work-from-home patterns and the urban recovery. We also report built-versus-near-miss estimates where the data permit. Rents rise modestly at built sites relative to near misses, while foot traffic shows no robust change and commercial spending is unchanged in either design. Section 6 uses these estimates to examine who bears the local incidence of entry.

[3]With two rings, replacing γ_t with a ring-by-year term would absorb the average inner-versus-outer gap in each calendar year and leave the event-study coefficients identified only from differences in construction timing across sites. We use the calendar-year specification throughout. Section 4 reports the estimate under the log of radiance as well, so that its magnitude can be compared with the near-miss design.

3.4 Employment and county-industry diagnostics

Local labor demand is harder to identify than the capital footprint, so we measure it at two geographic scales. The local analysis counts vacancies advertised near each facility and uses the 50 to 150 kilometer comparison area described above. A second analysis moves to the county level, where the data center's own workers and any related firms need not fall inside a single ring. Using employment and establishment counts by industry from the Quarterly Census of Employment and Wages, we estimate

$$\log(1 + \text{Emp}^s_{c,t}) = \sum_k \theta^s_k \mathbf{1}[t - t^o_c = k] \times \text{DC}_c + \mu_c + \lambda_t + u_{c,t}, \tag{3}$$

where s indexes industry, c county, and t^o_c the county's first data-center year. The specification includes county and year fixed effects. We also estimate cohort-by-event-time effects using not-yet-treated counties as controls and state-cluster bootstrap inference, following Callaway and Sant'Anna [2021]. Two data features constrain the interpretation. Depending on the sector, 27 to 68 percent of detailed employment cells are suppressed. We impute these cells from the disclosed industry and geographic margins, but the estimates remain sensitive to how the missing employment is allocated. Establishment counts in several industries also begin rising before opening. Section 5.5 reports these diagnostics and results in full.

3.5 Credibility of the capital-footprint design

The capital-footprint estimate is stable across the standard checks for a ring design. Dropping each facility in turn leaves the coefficient in a narrow range, and excluding the campus itself and using the 1 to 2 kilometer ring as the inner comparison produces the same spatial decay. Because hyperscale facilities often share a campus, we cluster standard errors at the campus rather than the facility level. This widens the confidence interval but leaves the effect statistically significant.

Federal Aviation Administration obstruction-evaluation filings provide an independent check on construction timing. We match facilities by location to these filings, whose timestamps are recorded before any radiance response and do not use satellite information. Re-dating sites with these filings produces a capital footprint that is positive and grows after filing. Because an FAA filing usually precedes physical construction, the comparison establishes the ordering of events rather than an exact alternative groundbreaking date.

Both the ring and near-miss designs use staggered construction dates. As an additional check, we compare each construction cohort only with never-built sites observed in the same calendar years (Section 4). This cohort-robust estimate is imprecise and fails its lead test, so we treat it as an inconclusive diagnostic rather than corroborating evidence.

The rent analysis requires an outcome-specific adjustment. Between 2020 and 2023, rents rose faster near metropolitan cores than in surrounding areas. Because a data center's inner and outer rings occupy different points on this gradient, the ring comparison may capture the metropolitan gradient rather than the facility. Section 6 therefore adds a synthetic control for each facility from population-matched locations without a data center. This comparison absorbs common growth and isolates the additional change at the data-center location.

Table 3 summarizes which design carries the interpretation for each outcome. No single comparison is credible for every outcome.

Table 3: Outcome estimates across research designs

Outcome (hyperscale)	Built vs. near miss	Spatial comparison	County evidence
Capital (nighttime lights)	**+70% ($p < 0.01$)**	+38% ($p < 0.01$)[a]	–[b]
Residential rent	+4.3% ($p = 0.23$)	**+2 to +5%**[c]	–[b]
Advertised salary	−11.6% ($p = 0.54$)	**+1.3% ($p = 0.53$)**[d]	–[b]
Job postings	+4.0% ($p = 0.75$)	**+5.2% ($p = 0.35$)**[d]	–[b]
Supplier postings	–[e]	**+5.9% ($p = 0.61$)**[d]	–[b]
Foot traffic	−12.0% ($p = 0.12$)[f]	−3.3% ($p = 0.48$)[l]	–[b]
Consumer spending	**+4.6% ($p = 0.60$)**	+12.1% ($p = 0.25$)[h]	–[b]
Firm entry (SoS registrations)	–[e]	−0.5% ($p = 0.62$)[i]	+1.4% ($p = 0.49$)[j]
Workplace employment	–[e]	excluded[g]	inconclusive[k]

Notes: Post-event estimates for hyperscale facilities at the innermost ring each design supports. Bold marks the single design used to interpret each outcome. The built-versus-near-miss capital estimate is the isolated-sample construction-year jump. The two groups are not balanced on several observable characteristics (Table 2), and magnitudes at longer horizons rest on few near-miss controls (Section 4).

[a] Agrees with the near-miss estimate once both use the same transform. The gap is $\log(1 + x)$ against $\log(x)$, not the comparison group (Table 4).

[b] Not estimated at the county level. There is no county analog of the ring outcome.

[c] Synthetic-control headline. The within-site estimates (+6.6 to +9.4%) are an upper bound (Section 6).

[d] Pooled 0–25 km area against a 50–150 km comparison area, allowing labor-market effects to extend beyond the immediate site (Section 3.2).

[e] No panel around near-miss sites exists for this outcome (business registrations and LODES cannot be reconstructed around never-built sites at ring scale).

[f] Attenuates to −3.0% ($p = 0.68$) on the balanced panel, so we do not interpret it as reliable evidence of displacement.

[g] Excluded because the pre-opening coefficients show an inner-versus-outer trend that the ring design cannot separate from an effect (Section 3).

[h] Cleaned-panel trend-break estimate using the v9 opening dates. It is positive but imprecise, and the built-versus-near-miss design carries the interpretation.

[i] Hyperscale-only sample using updated opening dates. The pooled all-archetype estimate is +0.7% ($p =$ 0.24).

[j] Census Business Formation Statistics, national county design.

[k] Data-processing employment and establishments rise with direct entry after suppressed employment is imputed. Estimates for related industries are not consistent across sectors, pre-trend tests, and alternative allocations of the suppressed cells (Section 5.5).

[l] Cleaned-panel trend-break estimate. The plain ring comparison has a pre-opening trend, so the built-versus-near-miss design carries the interpretation.

4 The capital footprint

4.1 Within-site ring design

The within-site estimate compares the inner ring around a facility with the 25 to 50 kilometer ring around the same site, using the construction start recovered from satellite-detected land clearing (Section 2.1). Two aspects of this design determine how we estimate the effect. First, a hyperscale campus takes two to three years to build. Dating the event to the operating year would place the construction phase, when the parcel is already changing, in the pre-period and misclassify that activity as a pre-trend. Observing construction and operations separately avoids this problem.

Second, the headline sample excludes sites built next to an existing data center. Roughly half of hyperscale facilities with a satellite construction date lie within two kilometers of an earlier facility and therefore inherit some of its brightness before construction begins. At these non-isolated sites, inner-ring brightness rises by about 50 percent during the pre-period. By contrast, the inner ring at an isolated site contains only the future parcel. Its pre-construction excess is 7 percent, and the event study is flat before groundbreaking. A joint test does not reject equality of the pre-construction coefficients to zero ($p = 0.37$), after which brightness rises by 38 percent in the construction year. Fitting a separate linear trend to each site's inner-minus-outer path over event years -4 through -1 and extrapolating to event year zero produces the same 38 percent increase, with a campus-bootstrap 95 percent interval of 15 to 71 percent.

4.2 Built versus announced near misses

Figure 1 reports the built-versus-near-miss comparison from equation (1). The comparison follows radiance within one kilometer of 43 isolated built facilities and 49 announced projects located away from operating campuses that were not built. As Table 2 shows, the two groups differ on several observable characteristics.

Before construction, the coefficients range from 8 to 11 percent below the omitted year, but they are jointly indistinguishable from zero when clustered by campus ($p = 0.24$), and their linear slope is also insignificant (+3.5 log points per year, $p = 0.22$). Radiance then rises by 70 percent in the construction year and reaches 231 percent above the near-miss sites one year later. Campus-level clustering widens the confidence intervals by roughly 60 percent, yet every post-groundbreaking coefficient remains statistically significant. The nighttime-light footprint captures activity at the parcel, including security and parking lights, substations, cooling equipment, and eventually continuous operation.

Two features determine how this benchmark should be interpreted. First, the near-miss cohort is concentrated in recent announcement years. Although 48 control sites identify the construction-year contrast and 21 identify the year-one contrast, five or fewer remain at later horizons. We therefore end the figure after year one and place no weight on subsequent coefficients.

Second, the estimate describes a campus rather than a single building. Thirty-two of the 43 isolated built sites add another building within two kilometers during the event window, which

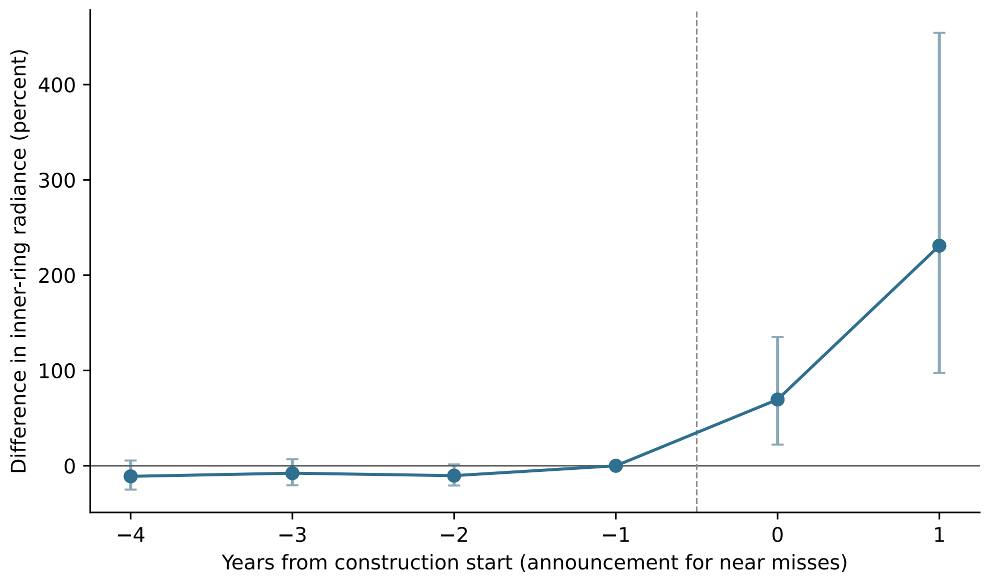

Figure 1: Capital footprint at built versus announced near-miss sites

Alt text: The event-study line is flat for four years before construction, rises sharply in the construction year, and rises further one year later. Confidence intervals widen after construction.

Notes: Event study based on equation (1). The inner-ring (0 to 1 km) log nighttime-lights path at 43 isolated built hyperscale facilities is measured relative to 49 isolated announced near-miss facilities, omitting the year before. Built facilities are dated by satellite-detected land clearing. Near-miss facilities are dated by announcement year. Standard errors are clustered on two-kilometer campus groups. Near-miss coverage is 48 sites in the construction year and 21 one year later. Table 2 reports observable differences between the groups, and Table 4 reports the estimate across designs, transforms, and construction cohorts.

helps explain why the full-sample path continues rising through year four. The 11 sites that remain single-building projects are also much smaller, and their construction-year estimate is 7.7 percent with a standard error of 6.9 percent. Phased construction cannot explain this initial gap because no sibling building has yet been added in the construction year. Instead, the two groups differ in project scale from the outset. We therefore interpret the estimate as the physical response to a large multi-building campus, which is also the project that a jurisdiction recruits.

4.3 Reconciling the two designs

The ring and near-miss designs produce the same timing and shape, and their magnitudes align once the outcomes are placed on a common scale (Table 4). The apparent difference between the 70 and 38 percent estimates comes from the transformation of radiance rather than the control group. The near-miss design uses the log of mean radiance, whereas the ring design uses $\log(1 + \text{NTL})$, which compresses proportional changes at parcels with very low pre-construction radiance. When the ring design also uses log radiance, it produces increases of 63 percent in the construction year and 226 percent one year later, compared with 70 and 231 percent in the near-miss design. A Poisson model of radiance levels yields 59 percent. Thus, three specifications that differ in control group and functional form place the initial footprint in a similar range.

Both designs also reveal the same gradient across construction cohorts. Among sites that broke ground by 2018, radiance rises by 97 percent in the near-miss design and 95 percent in the ring design. The corresponding estimates for sites that began in 2019 or later are 33 and 35 percent. Because the ring design does not use near-miss controls, this gradient cannot be explained by the concentration of near misses in recent years. Instead, it reflects the initial location. Earlier projects were built on much darker parcels, with median pre-event radiance of 1.1 compared with 4.5 for later projects, so the same amount of new lighting creates a larger proportional change. A cohort-robust estimator limited to 17 treated sites with never-built controls in the same calendar years yields 27 percent, with a 95 percent interval of $[-2, +92]$ percent. Because its pre-construction leads reject jointly ($p = 0.008$), we treat that estimate as inconclusive and rely on the two better-supported designs for the cohort pattern.

Restricting the near-miss group to the nine isolated projects with a named hyperscale operator yields a construction-year estimate of 72 percent, compared with 70 percent in the full isolated sample. The smaller subset has a somewhat less stable pre-period ($p = 0.15$) and cannot support longer-run estimates.

4.4 Spatial decay

The footprint is sharply localized. Its magnitude is largest within one kilometer, falls by more than half between one and two kilometers, and becomes small between two and five kilometers, with no detectable response beyond that distance. Omitting the central kilometer and treating the one-to-two kilometer ring as the inner zone produces the same decay. The nighttime-light increase therefore extends beyond the parcel itself but not beyond its immediate surroundings.

4.5 Colocation facilities as a comparison

Colocation facilities are generally smaller, multi-tenant projects that often occupy or expand existing buildings, although some are purpose-built. They produce no comparable change in nighttime lights. In the construction-dated within-site design for the first five kilometers, the construction-phase estimate is -1.0 percent ($p = 0.65$) across 21 colocation facilities, compared with $+8.0$ percent ($p = 0.02$) across 43 hyperscale facilities. The colocation estimate during operations is also near zero (-1.4 percent, $p = 0.74$). We therefore focus the remaining analysis on hyperscale facilities and report colocation results where the data support a meaningful comparison.

5 Test of agglomeration

Section 4 shows that hyperscale construction creates a large, spatially concentrated capital shock. We now ask whether that investment generates the local spillovers historically associated with large plant openings. Greenstone et al. [2010] organize these spillovers around a deeper shared labor pool, denser supplier linkages, and knowledge spillovers among nearby firms. Each can raise productivity when firms cluster [Combes et al., 2012]. Our outcomes capture responses associated

Table 4: The capital footprint across estimators and cohorts

Estimator	Construction year	Year 1	Built sites	Lead p
Panel A: Built versus announced near miss, log(NTL)				
All isolated sites	+70%	+231%	43	0.24
broke ground 2018 or earlier	+97%	+305%	26	0.41
broke ground 2019 or later	+33%	+140%	17	0.69
controls with a named hyperscale sponsor	+72%	+259%	43	0.15
Poisson on radiance levels	+59%	+153%	43	–
cohort-robust, never-treated controls	+27%	+154%	17	0.01
Panel B: Within-site rings, log(NTL)				
All isolated sites	+63%	+226%	43	0.51
broke ground 2018 or earlier	+95%	+294%	26	0.16
broke ground 2019 or later	+35%	+142%	17	0.89
Panel C: Within-site rings, $\log(1 + \text{NTL})$*, as reported in the text*				
All isolated sites	+38%	+132%	43	0.37
site-specific linear pre-trends	+38%	–	43	–
broke ground 2019 or later	+25%	+99%	17	0.90

Notes: Change in inner-ring (0–1 km) radiance at isolated hyperscale sites, in the construction year and one year later, transformed from log points to percent. Panels A and B use the same outcome, so their magnitudes are directly comparable. Panel C applies the $\log(1+x)$ transform used in the text, which compresses changes at the low pre-construction radiance of these parcels. The within-site rows use no near-miss control at all, and they reproduce the cohort gradient in Panel A, so that gradient reflects differences between early and late construction cohorts rather than the thin near-miss coverage before 2020. Early cohorts sit on darker land, with a median pre-event radiance of 1.1 against 4.5 for later cohorts, which enlarges a log change. The cohort-robust row compares each construction cohort with never-built sites observed in the same calendar years, using the 17 sites with such support, and inference is a campus cluster bootstrap with 999 draws. Its construction-year 95 percent interval is [-2, +92] percent, and its lead test rejects, so we treat that row as inconclusive rather than as corroboration. The site-specific-trend row fits a separate linear pre-trend to each site's event years -4 through -1 and reports the observed-minus-predicted event-zero change. Its campus-bootstrap 95 percent interval is [+15, +71] percent. Lead p jointly tests event years -4 through -2. Standard errors are clustered on two-kilometer campus groups throughout.

with each channel, not the underlying exchange of workers, inputs, or knowledge. Because labor-market effects may extend beyond the immediate site, the labor outcomes compare the pooled 0 to 25 kilometer area with locations 50 to 150 kilometers away (Section 3.2). Colocation facilities provide an additional comparison, although they are located in denser markets than hyperscale campuses.

Figure 2 places the preferred estimates and confidence intervals on a common scale. Advertised salaries and firm entry are estimated precisely enough to rule out modest increases. The job-posting interval is wider, while supplier postings are too sparse to provide a meaningful bound. Appendix Table D.1 reports the confidence intervals and minimum detectable effects underlying the figure.

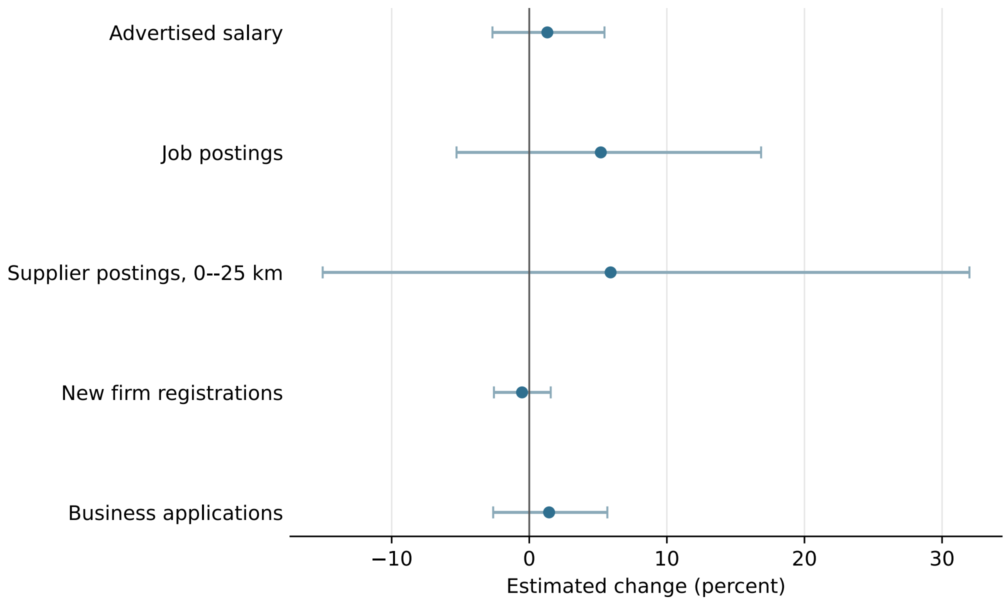

Figure 2: Precision of the agglomeration-channel estimates

Alt text: Horizontal confidence intervals for five estimates. Salary, firm-entry, and business-application intervals are narrow around zero. Job-posting and supplier intervals are wider and include sizable positive effects.

Notes: Preferred estimate for each outcome with 95% confidence intervals, transformed from log points to percent. Salary, job postings, and supplier postings compare the pooled 0 to 25 km zone with a 50 to 150 km control. Firm registrations use the Virginia and Texas within-site trend break. Business applications use the national county trend break.

5.1 Labor pooling

If a data center deepens the local labor pool, hiring and wages should rise nearby. Across 120 hyperscale facilities, job postings increase by 5.2% ($p = 0.35$), with a 95 percent interval of $[-5.3, +16.8]$ percent. The built-versus-near-miss estimate is similar at +4.0% ($p = 0.75$), with a flat pre-period ($p = 0.96$). Poisson counts, a long difference that conditions on pre-period volume, and construction dating also produce estimates near zero.

The interval is not a precise zero. Before opening, the median hyperscale facility has roughly 760 advertised vacancies per year within five kilometers. The upper bound of 17 percent corresponds to about 130 additional postings annually. Because an operating campus employs only a few dozen people, the design cannot isolate the facility's own hiring from a modest local multiplier. The confidence interval rules out a large hiring boom but allows smaller gains.

Advertised salaries provide a tighter test. Across 119 facilities, the mean log salary of nearby postings rises by 1.3% ($p = 0.53$), with a 95 percent interval of $[-2.7, +5.5]$ percent. Because the panel is concentrated in posted salaried positions, this estimate is most relevant for white-collar workers and may not capture compensation in the construction trades. The share of postings that report a salary is unchanged at opening ($p = 0.43$), making differential disclosure an unlikely explanation for the salary estimate. Workplace employment from the Longitudinal Employer-Household Dynamics data faces a different problem. Employment within five kilometers rises by 8 percent, but the increase begins before opening, and detrending reduces the opening-year change to -0.3%.

5.2 Supplier linkages

The second channel operates through local suppliers. We identify postings in the industries that build and equip data centers, including construction and specialty trades, electrical-equipment wholesale, architectural and engineering services, and machinery manufacturing. Relative to the 50 to 150 kilometer control, supplier postings within 25 kilometers rise by 5.9% ($p = 0.61$), with a 95 percent interval of $[-15, +32]$ percent. Because these postings form a thin series, the interval is too wide to distinguish a substantial increase from a decline. Supplier postings therefore do not establish whether local demand for these industries rises.

5.3 Firm clustering and knowledge spillovers

The third channel is knowledge exchange among co-located firms. We do not observe knowledge transmission directly. We instead ask whether compute-intensive firms enter or cluster around new facilities, an observable implication of a new local cluster. Appendix A maps firms that make intensive use of computing capacity. These firms are closer to data centers than the average firm, but much of the relationship reflects a shared metropolitan location, and 69 percent of nearby compute firms predate the nearest facility. Exurban data centers, by contrast, often have few compute firms nearby. Although these descriptive patterns do not identify an opening effect, they do not show a compute cluster forming around the facility.

Firm entry also does not rise. We geocode secretary-of-state business registrations in Virginia and Texas to the facility rings and find no average break at opening. Within five kilometers, registrations rise by 0.7% ($p = 0.24$) in the pooled sample and fall by 0.5% ($p = 0.62$) around hyperscale facilities. A national county design using Census Business Formation Statistics yields a similarly small estimate of $+1.4\%$ ($p = 0.49$). The upper bounds of the corresponding 95 percent intervals are 1.6 and 5.7 percent. These results rule out systematic entry effects above those bounds, while allowing that individual locations may respond differently.

5.4 Productivity implications

Each of the three channels can raise productivity among incumbent firms. Greenstone et al. [2010] find total factor productivity gains of roughly 12 percent among incumbent establishments. We do not observe plant-level productivity, so we instead examine advertised salaries and hiring at firms that existed before the data center, two margins through which a large productivity gain might appear. The salary interval has an upper bound of 5.5 percent. It would rule out a 12 percent productivity gain only if at least 45 percent of that gain passed through to posted salaries. With lower pass-through, such a gain would remain inside the interval. The salary result therefore bounds the wage response, not productivity itself. Incumbent hiring also shows no positive break at opening, but the pooled and colocation samples have a downward pre-trend that the design cannot fully remove. We treat this hiring pattern as corroborating evidence rather than a causal estimate and do not claim to measure productivity directly.

5.5 What county-industry data can establish

County-industry data show whether the facility's direct employment is accompanied by growth in local suppliers or related digital services. The detailed Quarterly Census of Employment and Wages cells create two problems in this setting. The first is disclosure. The Bureau of Labor Statistics suppresses employment when publishing a value could reveal an individual employer, and the public files record these unobserved cells as zero. Retaining the disclosure flag shows that suppression affects 68 percent of county-year records in data processing, 51 percent in telecommunications, and 66 percent in web and information services. Treating the suppressed cells as observed zeros would mechanically generate employment growth when a value first becomes publishable.

We therefore impute the suppressed cells using a fixed-point procedure adapted from Autor et al. [2026]. We initialize each cell using its establishment count and prior-year employment per establishment, then rescale the missing values to match state-industry and county-parent employment totals. Published county-industry values are never altered. The fixed point converges to both margins in 92 percent of the state-year groups with suppressed cells. For the remainder, we retain the allocation that matches state-industry totals. The procedure fills the panel and allows all five industries to use the same 46 treated counties. Appendix Table D.2 compares these estimates with the disclosed-cell sample and with a simpler allocation that matches only state-industry totals.

The second problem is that host counties were already changing before the facility opened. We estimate cohort-by-event-time effects relative to the year before opening, using not-yet-treated counties as controls and a state-cluster bootstrap [Callaway and Sant'Anna, 2021]. Table 5 reports the average estimate over event years zero through five together with a joint test of the pre-period leads. Five of the ten hyperscale outcome series change significantly before opening. The establishment pre-trend fails in telecommunications. Both employment and establishments fail in web and information services, while only employment fails in computer systems design and only establishments fail in professional services. Data-processing employment and establishments both have flat pre-trends.

Table 5: County-industry disclosure and pre-trend analysis for hyperscale facilities

Industry (NAICS)	Employment suppressed	Employment effect	Pre-trend p	Establishments effect	Pre-trend p
Data processing (518)	68%	+25.9%*	0.64	+27.2%***	0.11
Telecommunications (517)	51%	−1.4%	0.34	−2.6%	0.04
Web and information services (519)	66%	+11.2%	< 0.01	+18.3%***	0.03
Computer systems design (5415)	40%	+4.8%	0.04	+5.7%	0.11
Professional and technical services (541)	27%	+6.0%**	0.27	+7.6%***	< 0.01

Notes: Employment suppression is the share of reported county-sector-year records carrying the BLS non-disclosure code. Employment effects use the hierarchy imputation described in the text, while establishment counts are published without suppression. Effects are cohort-weighted mean log changes over event years zero through five, transformed to percent, relative to not-yet-treated counties. Pre-trend p is a joint test for event years −4 through −2. Inference uses 1,999 state-cluster bootstrap draws. The 519 series ends in 2021 because the 2022 NAICS revision moved major components to sectors 513 and 516, so its post-period average is taken over a shorter window than the other sectors and is weighted toward earlier cohorts. Stars refer to the post effect. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Benjamini–Hochberg q-values across the ten sector-outcome tests are < 0.001 for data-processing establishments, 0.011 for web and information-services establishments, 0.049 for professional-services employment, and < 0.001 for professional-services establishments.

Data-processing employment rises by 26 percent and establishments by 27 percent. The employment estimate has a bootstrap p-value of 0.10 and does not survive the multiple-testing adjustment, while the establishment estimate does. Because data processing contains the facility itself, the establishment result confirms direct entry and the employment estimate is consistent with direct hiring. Neither result shows that a surrounding cluster formed.

The other industries do not show a consistent response. The employment and establishment estimates for telecommunications are near zero, although the establishment series has a pre-trend. Web and information-services establishments rise by 18 percent, but both outcomes in that sector were already increasing before opening. In computer systems design, the employment pre-trend fails at $p = 0.04$, while the establishment pre-trend test does not reject at $p = 0.11$; neither post-opening estimate is statistically significant. Professional-services employment rises by 6 percent in the hierarchy-imputed panel, but its establishment count was already increasing before opening, and the employment pre-trend fails when we use the state-only allocation. The county records therefore provide some evidence of direct data-center employment but no consistent evidence that local supplier industries or related digital services grew in response.

6 Incidence outside the labor market

Having examined whether the investment propagates through labor pooling, local suppliers, and firm clustering, we turn to housing, fiscal, and consumer outcomes.

Within-site estimates place rent increases at +6.6% within 5 kilometers, +9.4% between 5 and 10 kilometers, and +8.5% between 10 and 25 kilometers. Unlike the capital footprint, this gradient does not decline with distance, raising the possibility that it reflects broader metropolitan rent growth rather than the facility. A density-matched comparison shows that similar locations without a data center also experienced substantial rent growth during the 2020–2023 work-from-home period. Facility-level synthetic controls yield smaller increases of two to five percent across the three rings, although the synthetic controls do not support formal placebo inference. The built-versus-near-miss estimate is +4.3% ($p = 0.23$), with a flat pre-period ($p = 0.44$). We therefore treat a rent increase of a few percent as suggestive and the larger within-site estimates as upper bounds. Colocation facilities show no detectable rent response.

Housing supply and net migration do not show a corresponding increase. In the Census Building Permits Survey, permits for buildings with at least five units rise by 3.4% ($p = 0.87$). For migration, we assign each county its first reliably dated hyperscale opening within 25 kilometers, including openings before the IRS panel, and compare treated counties with those not yet treated. Gross inflows rise by 5.4% ($p = 0.02$) and gross outflows by 7.6% ($p < 0.01$), but outflows and migrant income were already increasing before opening. We therefore do not interpret the gross-flow estimates as effects of the data center. The net measures are close to zero. The log inflow-to-outflow ratio moves by -1.4% over event years zero through five ($p = 0.52$), while the analogous adjusted-gross-income ratio moves by +0.4% ($p = 0.91$). Their pre-trend p-values are 0.14 and 0.09, respectively. Net migration therefore does not rise, and the pre-trends prevent a causal interpretation of the gross flows.

Foot traffic and commercial spending also show no robust response relative to near-miss sites. Foot traffic falls by roughly 12 percent in the full panel ($p = 0.12$), but the estimate attenuates to 3 percent in the balanced panel ($p = 0.68$). Commercial spending rises by 4.6% ($p = 0.60$). Although the cleaned-panel within-site estimate is also positive at 12.1% ($p = 0.25$), it is imprecise and is not confirmed by the built-versus-near-miss design.

The fiscal estimates are concentrated in school finance rather than broad local government spending. Table 6 uses the same not-yet-treated estimator and state-cluster bootstrap as the county analyses. Local school property-tax revenue rises by 16.6% ($p = 0.07$) without a detectable pre-trend, and capital outlay per pupil rises by 28.5%, although the latter estimate is imprecise ($p = 0.18$). Total local revenue and local revenue per pupil remain unchanged. County highway spending also has a positive, marginally significant estimate, but its magnitude changes materially when we update the opening dates, and the broader public-works series fails its pre-trend test. Neither the property-tax nor the highway estimate survives adjustment across the six fiscal outcomes ($q = 0.21$ for each). After this adjustment, neither fiscal estimate is statistically distinguishable from zero.

Table 6: Public-finance effects after hyperscale openings

Outcome	Geography	Effect	Post p	FDR q	Pre-trend p
Local property-tax revenue	school district	+16.6%	0.07	0.21	0.31
Total local revenue	school district	−2.3%	0.83	0.83	0.55
Local revenue per pupil	school district	+2.0%	0.76	0.83	0.14
Capital outlay per pupil	school district	+28.5%	0.18	0.35	0.57
Highway current + construction	county, balanced	+38.9%	0.07	0.21	0.15
Broad public works	county, balanced	+9.3%	0.40	0.61	< 0.01

Notes: Cohort-weighted mean log changes over event years zero through five, transformed to percent, relative to not-yet-treated geographies. School data are NCES/Census F-33 finance files for 2012–2023. County data are Census local government finance files for 2012–2023. Outcomes are adjusted for state-by-year means before the cohort comparisons. Inference uses 1,999 state-cluster bootstrap draws. Pre-trend p jointly tests event years −4 through −2. FDR q applies the Benjamini–Hochberg adjustment across the six outcomes in the table. The highway estimate moves materially across alternative opening dates and is treated as a sensitivity result rather than a headline effect.

The non-labor outcomes do not indicate a broad increase in local demand. Rents may rise modestly, while school revenue and road spending have positive but imprecise estimates. Net migration, new multifamily permits, foot traffic, and commercial spending do not rise.

7 Interpretation

Nighttime lights establish that construction occurred at the location and time recorded in our data and that the visible footprint ends within a few kilometers. The lights capture the facility itself, including security and parking lights, substations, and cooling equipment, rather than local economic activity. We organize the possible propagation around labor pooling, local supplier linkages, and firm clustering with potential knowledge spillovers.

The evidence is strongest against broad labor-pooling and firm-clustering responses. Advertised salaries and new-firm registrations do not rise at opening, and their upper confidence bounds are 5.5 and 1.6 percent. The salary bound excludes the corresponding wage response unless little of any productivity gain reaches posted salaries (Section 5.4). The hiring confidence interval rules out increases above 17 percent but allows smaller gains. Supplier linkages are less precisely estimated because supplier postings are too sparse to provide a useful bound. Finally, the firm-entry bound excludes a response of the size commonly associated with agglomeration. Firm entry is an observable implication of clustering rather than a direct measure of knowledge transmission.

Three limitations could allow a conventional local multiplier to go undetected. The response could occur beyond the zones we measure, although the 50 to 150 kilometer comparison is designed to capture that possibility. It could appear in employment that is neither advertised nor registered, some of which should nevertheless enter the county records. The response could also arrive outside our event windows, a possibility that future data on the post-2022 campuses will test. The geography of compute in Appendix A points instead to weak local linkages. Data centers and compute firms value the same metropolitan infrastructure and skilled labor, but 69 percent of compute firms within 25 kilometers predate the nearest facility. Their co-location therefore does not show that the data center created the surrounding cluster.

Synthetic-control rent estimates range from two to five percent, while net migration and multi-family permitting remain unchanged. School property-tax revenue and highway spending also have positive point estimates, although neither is statistically significant after correcting for the six fiscal outcomes we test. Residential electricity rates do not rise for the cloud-era cohorts in our sample. We do not observe land prices, environmental damages, or the terms of utility and tax agreements, so we cannot determine whether local welfare rises.

Finally, hyperscale and colocation facilities differ in both technology and location. Hyperscale campuses are generally larger and more likely to involve greenfield construction, whereas colocation facilities are multi-tenant and concentrated in denser markets. The hyperscale-colocation contrast describes the projects that jurisdictions encounter, but it combines facility and market differences. It does not identify the effect of randomly assigning one facility type rather than the other.

8 Conclusion

A hyperscale data center rivals the manufacturing plants in the agglomeration literature in capital scale, but the local response is much narrower. At isolated sites, nighttime lights rise when construction begins and fade within five kilometers, while daytime imagery confirms that the physical change starts with construction. Announced projects that were not built show no comparable break, although these near misses are not balanced runner-up sites.

The evidence differs in precision across the three spillover channels. Advertised salaries do not rise, and the hiring estimates rule out a large labor-pooling response while allowing smaller gains. Supplier postings are too sparse to determine whether local supplier linkages deepen. New-firm registrations do not rise, and most compute firms near data centers arrived first. The data therefore provide little evidence that new facilities create a surrounding cluster, although they do not measure knowledge transmission directly.

For the cloud-era cohorts we study, a large physical investment therefore does not produce the nearby worker, supplier, and firm responses needed to support a broad local-multiplier claim. Data-processing employment and establishments rise, but these series include the facility itself. Rent and public-finance estimates are positive but imprecise. Two limits define the scope of this conclusion. The hiring and supplier-posting estimates allow for moderate gains, and the post-2022 artificial intelligence campuses are larger and more power-intensive than any facility in our sample. Whether those new campuses generate a different local response remains an open question, even as states continue to write subsidies for them.

CRediT authorship contribution statement

Dany Bahar: Conceptualization, Methodology, Writing – review and editing. **Greg C. Wright:** Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Software, Visualization, Writing – original draft, Writing – review and editing.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work, the authors used OpenAI ChatGPT and Codex, Anthropic Claude, and Google Gemini to assist with code debugging, robustness checks, and editorial revision. After using these tools, the authors reviewed and edited the content as needed and take full responsibility for the content of the article.

References

Advan Research. Foot traffic / weekly patterns plus [dataset]. https://doi.org/10.82551/C103-N851, 2025.

Fernando E. Alvarez, David Argente, Joyce Chow, and Diana Van Patten. Data centers and local economies in the age of AI: A shift-share approach. NBER Working Paper 35194, National Bureau of Economic Research, May 2026. URL https://www.nber.org/papers/w35194.

David Autor, Anne Beck, David Dorn, and Gordon H. Hanson. Help for the heartland? the economic and political effects of the trump tariffs in U.S. local labor markets. Working paper, June 2026. URL https://www.ddorn.net/papers/ABDH-Heartland.pdf.

Brantly Callaway and Pedro H. C. Sant'Anna. Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230, 2021.

Pierre-Philippe Combes, Gilles Duranton, Laurent Gobillon, Diego Puga, and Sébastien Roux. The productivity advantages of large cities: Distinguishing agglomeration from firm selection. *Econometrica*, 80(6):2543–2594, 2012.

Ramakrishnan Durairajan, Paul Barford, Joel Sommers, and Walter Willinger. InterTubes: A study of the US long-haul fiber-optic infrastructure. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication (SIGCOMM)*, pages 565–578, 2015.

Gilles Duranton and Diego Puga. Micro-foundations of urban agglomeration economies. In *Handbook of Regional and Urban Economics*, volume 4, pages 2063–2117. Elsevier, 2004.

Glenn Ellison, Edward L. Glaeser, and William R. Kerr. What causes industry agglomeration? Evidence from coagglomeration patterns. *American Economic Review*, 100(3):1195–1213, 2010.

Good Jobs First. Subsidy tracker: Data center industry. https://subsidytracker.org/, 2024.

Michael Greenstone, Richard Hornbeck, and Enrico Moretti. Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings. *Journal of Political Economy*, 118(3): 536–598, 2010.

J. Vernon Henderson, Adam Storeygard, and David N. Weil. Measuring economic growth from outer space. *American Economic Review*, 102(2):994–1028, 2012.

International Energy Agency. Energy and AI. Technical report, International Energy Agency, Paris, April 2025. URL https://www.iea.org/reports/energy-and-ai.

Owen Kay, Robert Reaser, and Reid Taylor. Processing power: The effect of data centers on wholesale electricity markets. Working Paper 2606, Federal Reserve Bank of Dallas, March 2026. URL https://doi.org/10.24149/wp2606.

Patrick Kline and Enrico Moretti. Local economic development, agglomeration economies, and the big push: 100 years of evidence from the Tennessee Valley Authority. *Quarterly Journal of Economics*, 129(1):275–331, 2014.

LinkUp. Job records [dataset]. https://doi.org/10.82551/V091-HN53, 2025.

Robyn Meeks, Jacquelyn Pless, Zhiyuan Qi, and Zhenxuan Wang. Who pays for growth? Evidence from data centers and the grid. Working Paper WP-2026-13, MIT Center for Energy and Environmental Policy Research, June 2026. URL https://ceepr.mit.edu/workingpaper/who-pays-for-growth-evidence-from-datacenters-and-the-grid/. SSRN No. 6967558.

Enrico Moretti. Local multipliers. *American Economic Review*, 100(2):373–377, 2010.

Nicholas Z. Muller. Measuring the impact of data centers in the United States economy: Monetary damage from air pollution and greenhouse gas emissions. NBER Working Paper 35100, National Bureau of Economic Research, April 2026. URL https://www.nber.org/papers/w35100.

Carlianne Patrick. Identifying the local economic development effects of million dollar facilities. *Economic Inquiry*, 54(4):1737–1762, 2016.

Carlianne Patrick and Mark Partridge. Agglomeration spillovers and persistence: New evidence from large plant openings. Working Paper CES-22-21, U.S. Census Bureau Center for Economic Studies, 2022.

RentHub. Rental data [dataset]. https://doi.org/10.82551/5MMV-7M10, 2022.

SafeGraph. Spend patterns [dataset]. https://doi.org/10.82551/NSF5-R186, 2022.

Arman Shehabi, Sarah Josephine Smith, Alex Hubbard, Alexander Newkirk, Nuoa Lei, Md AbuBakar Siddik, Billie Holecek, Jonathan G. Koomey, Eric R. Masanet, and Dale A. Sartor. 2024 United States data center energy usage report. Technical Report LBNL-2001637, Lawrence Berkeley National Laboratory, December 2024. URL https://doi.org/10.71468/P1WC7Q.

Cailin Slattery and Owen Zidar. Evaluating state and local business tax incentives. *Journal of Economic Perspectives*, 34(2):90–118, 2020.

Veridion. Firmographics (company core profiles) [dataset]. https://doi.org/10.82551/RTSB-1K62, 2026.

WageScape. Job postings with salary [dataset]. https://doi.org/10.82551/TSAG-NZ08, 2022.

Asa Watten, John Bistline, and Geoffrey Blanford. Have data centers raised your electric bill? Causal evidence from the United States. arXiv:2606.19777, June 2026. URL https://arxiv.org/abs/2606.19777.

Daniel Yue and Yiyang Zeng. The local economic effects of data center entry. Working paper, SSRN No. 6497238, March 2026. URL https://ssrn.com/abstract=6497238.

A The geography of compute

A common rationale for data-center incentives is that a facility will anchor a local digital economy by attracting firms that use cloud capacity. We examine this claim by mapping the location of heavy compute users relative to data centers and asking whether any shared location predates the facility. We use the Veridion firmographic database [Veridion, 2026], which geolocates 6.5 million US establishments and classifies them by industry, employment, and founding year. We define the compute-native core to include software publishing, computing infrastructure and hosting, computer systems design, and web and streaming services. A broader definition adds finance, scientific R&D, and semiconductors. These descriptive exercises use all 1,087 reliably dated data centers in the continental registry, rather than only the hyperscale analysis sample.

A.1 Compute users cluster near data centers

Compute users are closer to data centers than the average firm. The median compute-native firm is 11 kilometers from the nearest data center, compared with 20 kilometers for non-compute firms. Similarly, 69 percent of compute users are within 25 kilometers of a data center, compared with 55 percent of other firms (Table A.1). Proximity also varies across compute-intensive industries. Cloud and hosting firms are nearest, at a median distance of 6 kilometers, followed by software firms at 9 kilometers. Telecommunications and streaming firms, whose operations are more geographically dispersed, are farther away.

The proximity of compute users to data centers is not an artifact of headquarters location. When we measure a firm's operating footprint using the geography of its job postings instead, compute users remain closer to data centers. The distances are somewhat larger because branch offices extend beyond the headquarters metropolitan area. The proximity gap also does not simply reflect the concentration of both groups in a handful of technology hubs. After removing the 45 largest metropolitan areas, compute users remain closer to the remaining data centers than non-compute firms are, with median distances of 77 and 89 kilometers, respectively. Software firms are among the closest, and the broader association remains visible in second-tier cities and rural areas.

Table A.1: Distance from firms to the nearest data center

	Firms	Median km	≤ 25 km	≤ 50 km
Panel A: Headquarters location (Veridion)				
Non-compute firms	5,692,641	20.6	55%	70%
Compute users (core)	309,760	11.0	69%	80%
Cloud and hosting	26,188	6.1	75%	84%
Software	101,757	8.9	72%	84%
Computer systems design	154,934	11.7	68%	81%
Finance	193,540	15.9	62%	75%
Telecommunications	10,125	22.2	53%	66%
Panel B: Establishment footprint (job postings)				
Compute users (core)		15.1	63%	76%
Panel C: Outside the 45 largest metros				
Non-compute firms	2,603,195	89.0	20%	32%
Compute users (core)	103,356	76.8	30%	40%

Notes: Great-circle distance from each firm to the nearest reliably dated data center in the full continental registry. Panel A places each firm at its Veridion headquarters coordinate. Panel B replaces the headquarters with the firm's job-posting locations (WageScape), weighted by the number of distinct hiring firms at each location, so it reflects the operating footprint rather than the headquarters. Panel C excludes all firms and data centers within 50 km of the 45 largest metropolitan areas and technology hubs and measures distance to the nearest remaining data center. Compute-native "core" sectors are software publishing, computing infrastructure and hosting, computer systems design, and web and streaming services.

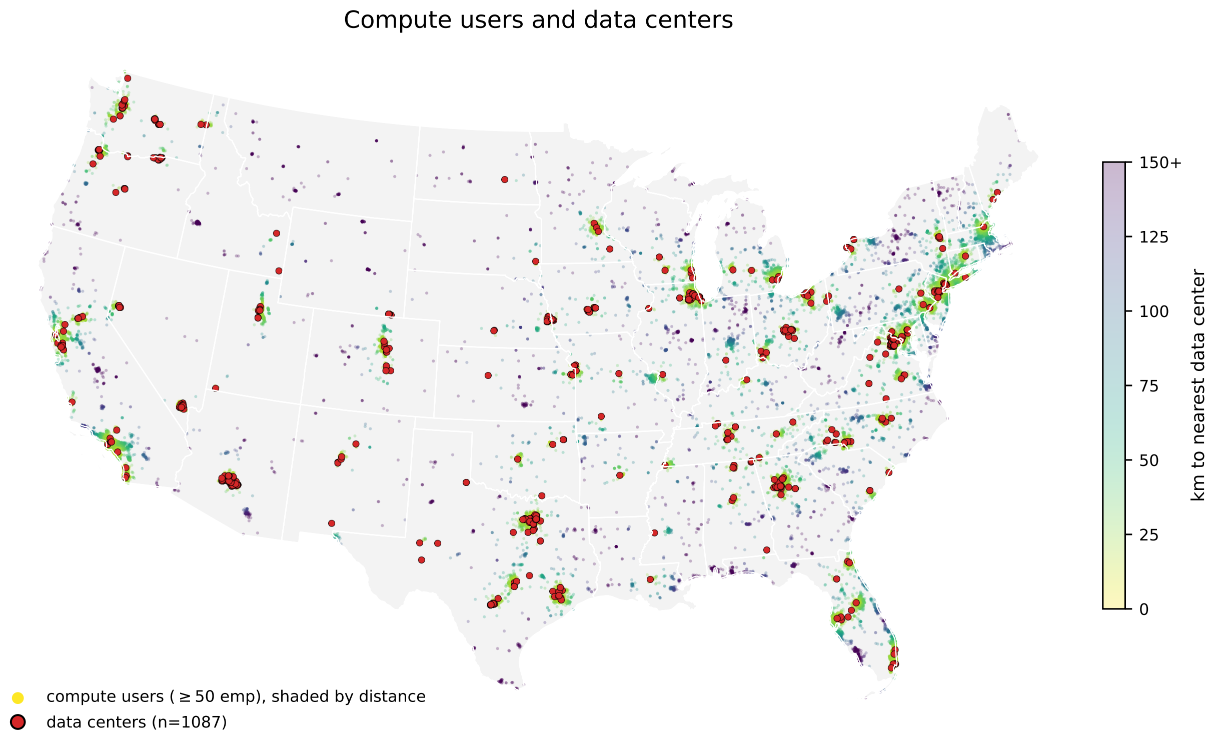

Figure A.1: Compute-user firms and data centers across the continental United States. Points are compute-native firms with at least 50 employees, shaded by distance to the nearest data center, with data centers marked in red. Both concentrate in the same coastal and metropolitan corridors. Points in the interior are darker because they are farther from a facility.

Alt text: A US map shows compute firms and data centers concentrated in coastal and metropolitan corridors, with sparse coverage across much of the interior.

The two groups trace the same map (Figure A.1), with both concentrated in coastal and metropolitan corridors and becoming less common across the interior. This association is strongest among large firms. As firm employment increases, the median distance to a data center falls monotonically, from 13 kilometers for the smallest firms to 4 kilometers for firms with more than ten thousand employees. The distribution narrows as well. Only about 5 percent of compute firms with at least a thousand employees are more than 50 kilometers from any data center. Those exceptions are generally software companies that grew in place in second-tier cities, such as Esri in Redlands, Jack Henry in rural Missouri, and Ansys outside Pittsburgh. Their experience nevertheless shows that even large compute users can access cloud capacity without locating near a data center.

A.2 Two kinds of data-center hub

The geographic overlap is primarily a metropolitan phenomenon. Grouping data centers into hubs reveals a sharp divide between metropolitan and exurban locations. Metropolitan hubs are embedded in dense clusters of sector-matched compute users. The 236 data centers around Ashburn, Virginia sit among thousands of government IT contractors, with DXC, Leidos, Booz Allen, SAIC, and CACI all within 25 kilometers. The Bay Area hubs are ringed by Google, Broadcom, Synopsys, and ServiceNow. By contrast, the large exurban hyperscale clusters are isolated. The 48 data centers around Hillsboro, Oregon and the cluster near Quincy, Washington each have on the order of four compute firms within 25 kilometers. The same type of physical infrastructure coincides with a dense compute neighborhood in a technology center and a sparse one in the exurbs.

A.3 Shared location does not mean that data centers created the cluster

Part of the explanation is shared infrastructure. Data centers and compute users both require access to the high-voltage grid, and both sit closer to it than ordinary firms. The median data center is 3.4 kilometers from a 230-kilovolt substation, the median compute user 5.0 kilometers, and the median other firm 6.4 kilometers. High-voltage substation density independently predicts where compute users locate. But the grid is only part of the story. Compute users remain closer to data centers after controlling for substation access. A shared pull toward the power backbone, fiber, and skilled labor can therefore place both in the same markets without one drawing the other.

Three patterns weigh against the interpretation that a data center creates the local compute cluster. First, the compute users near data centers predate them. Among compute firms within 25 kilometers of a data center, at least 69 percent were founded before the nearest data center opened. Because the firm registry over-represents recently founded firms, the share is likely higher. The local compute base usually existed first. Ashburn's history is consistent with this ordering. Northern Virginia's government technology firms and early internet exchange preceded the region's data-center cluster. Second, the job-posting analysis finds no detectable increase around opening, either in compute-using industries or overall, although the estimates are imprecise. Third, exurban data centers are located in areas with few compute firms. The median metropolitan data center has roughly 2,700 compute firms within 25 kilometers, while the median exurban data center has 47 (Table A.2).

Table A.2: Compute users near metropolitan versus exurban data centers

	Data centers	Median compute firms $\leq$ 25 km	Share with $<$ 5 nearby
Metropolitan data centers	814	2,714	0%
Exurban data centers	273	47	10%

Notes: The sample contains all 1,087 reliably dated data centers in the continental registry. Each data center is classified as metropolitan if it lies within 50 km of one of the 45 largest metropolitan or technology-hub centers, and exurban otherwise. The middle column reports the median count of compute-native firms (Veridion) within 25 km of the data center. The last column reports the share of data centers with fewer than five compute firms within 25 km. Separately, among compute firms within 25 km of any data center, 69 percent were founded before the nearest data center opened.

A.4 Implication for place-based policy

Data centers, especially exurban facilities, do not appear to draw compute users to the surrounding area. Firms that consume cloud capacity are attracted to transmission, fiber, and skilled labor, assets that are concentrated in metropolitan areas and are not created by a single facility. Most compute users located near data centers were already present, and we find no systematic increase in firm entry after opening. The capital footprint is large and local, while rent estimates are smaller and less precise. We do not detect the formation of a digital-firm cluster.

B The instrument search

A natural way to address endogenous site selection would be to find an instrument for data-center location. We investigated several candidates but did not find one with a credible exclusion restriction. This appendix documents that search and explains why the paper instead relies on the within-site and built-versus-near-miss designs. These tests also clarify the limits of common predictors of data-center siting.

Predetermined siting instruments. We tested several predetermined predictors of data-center siting, interacting them with the national capacity wave where appropriate. We also applied a manufacturing falsification test. A valid instrument for data-center activity should not independently predict manufacturing employment. Table B.1 summarizes the results.

Transmission-substation density strongly predicts siting, but it fails the falsification test because it also captures the availability of industrial land. The 1980 urban college share produces a moderate county first stage, but its reduced form for information employment is negative. Historical human capital may also affect local sectoral growth directly, so it does not provide a credible exclusion restriction. Proximity to long-haul fiber predicts siting across metropolitan areas but not within a county, so it cannot identify the sub-county effect studied here. Climate, measured by wet-bulb temperature and free-cooling days, does not predict siting. Operators can engineer around climate conditions when pursuing cheap power, and the largest US hyperscale hub is hot and humid. Finally, the timing of generation capacity has no predictive power within counties.

The same features that predict hyperscale siting, including power, industrial land, and fiber, are also correlated with a location's broader industrial character. By contrast, climate and historical human capital are less likely to violate the exclusion restriction but do not predict where hyperscale facilities locate. We therefore do not find a credible predetermined instrument for siting at the sub-county scale.

Table B.1: Predetermined siting instruments and why each fails

Instrument	First stage	Why it fails
345 kV substation density	strong	predicts manufacturing employment
1980 urban college share	moderate	negative information-sector reduced form
InterTubes fiber proximity	metro-scale	no within-county effect
Climate (wet-bulb, free-cooling)	none	does not predict siting
Generation-capacity timing	none	does not predict siting within counties

Notes: The county panel exercises interact each predetermined feature with the national hyperscale capacity wave where applicable. The fiber row also uses a separate within-county test based on distance to the nearest conduit. Predicting manufacturing employment violates the exclusion restriction.

A shift-share instrument also fails. We also construct a shift-share instrument similar to that used in county-level work [Alvarez et al., 2026], interacting predetermined fiber and college exposure with national growth in data-center demand. The instrument has predictive power. With square footage as the endogenous variable, the cross-sectional first-stage F-statistic is about 32, while the county-panel first stage is about 10. It does not, however, resolve the exclusion concerns described above. Public measures of data-center scale, including capacity, square footage, and facility count, are lumpy and supply-driven, while the underlying exposure measures also predict broader local development.

Because the shift-share design does not resolve the exclusion concerns, we do not use its estimates as evidence. The instrument also predicts population-serving sectors such as health, reinforcing the concern that it captures general area development. We therefore report the exercise as documentation of the search rather than as an identification strategy.

C Incidence on electricity rates

A central concern in data-center siting debates is that the load from a hyperscale facility will raise electricity bills for nearby households. Existing estimates are mixed. Wholesale-market modeling finds higher wholesale prices [Kay et al., 2026], utility-level estimates find higher retail rates, especially at investor-owned utilities [Meeks et al., 2026], and county-level instrumental-variable estimates find higher local electricity prices [Alvarez et al., 2026]. In contrast, Watten et al. [2026] estimate that data-center growth modestly reduced average retail rates from 2015 through 2024. We estimate whether residential customers in a host utility's territory pay more. Data centers select utilities with cheap power and available capacity, so a comparison of treated and untreated utilities must account for the facility's siting choice.

We treat each utility within each state as a separate unit. Using Energy Information Administration Form 861, we compute average residential prices as revenue over megawatt-hours sold from 2012 through 2024. A unit is treated in the year the first hyperscale facility opens in a county it serves. The hyperscale sample contains 52 utility-states first treated between 2015 and 2021, none of which also hosts a colocation facility. We compare them with roughly 2,000 utility-states that host no facility in the registry and, separately, with 54 utilities serving counties with an announced hyperscale near miss. A separate colocation sample of 60 treated utility-states is reported alongside.

Figure C.1 reports a cohort-by-event-time estimator using not-yet-treated controls. Against never-treated utilities, the mean effect over years zero through five is -1.0% ($p = 0.42$), with a joint pre-trend p-value of 0.97. The balanced panel, which retains utility-states observed in at least twelve of the thirteen years, gives -1.5% ($p = 0.17$). Against near-miss utilities, the full sample estimate is -0.9% ($p = 0.37$), and the year-five estimate is -1.9% ($p = 0.27$). Colocation estimates are small and positive, with a mean effect of $+0.7\%$ ($p = 0.54$).

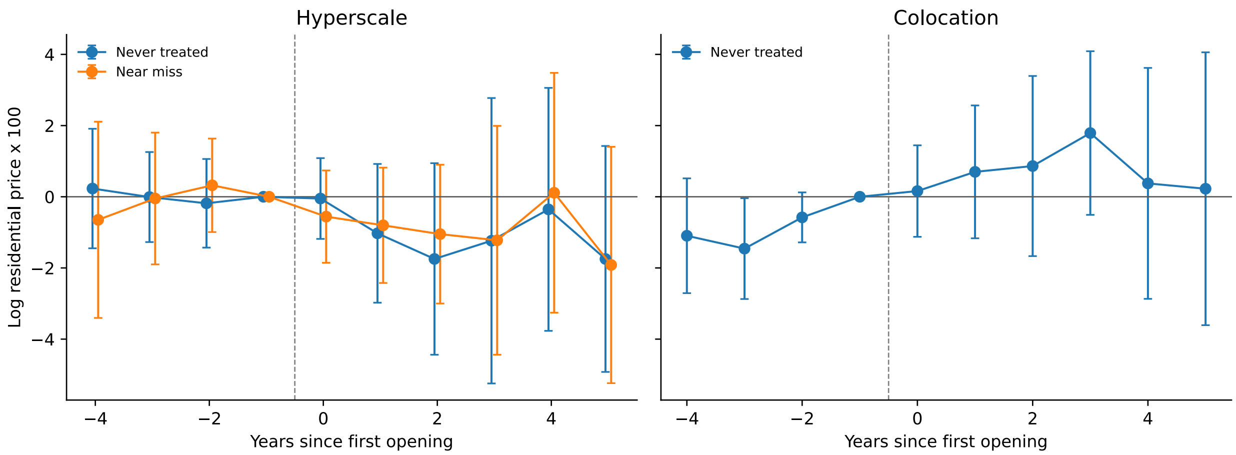

Figure C.1: Residential electricity prices around the first data-center opening in a utility's territory

Alt text: Two event-study panels show residential-price estimates around hyperscale and colocation openings. Hyperscale estimates become modestly negative after opening. Colocation estimates fluctuate around zero with wide intervals.

Notes: Cohort-by-event-time estimates of log average residential price (revenue over MWh, EIA Form 861 bundled-service utilities, 2012–2024). The hyperscale sample contains 52 treated utility-states, 2,040 never-treated controls, and 54 near-miss controls. The colocation sample contains 60 treated utility-states and the never-treated controls. The never-treated comparisons remove state-by-year means. The near-miss comparison removes year means. Inference uses 1,999 state-cluster bootstrap draws. The reference period is the year before opening, and 95% confidence intervals are shown.

The estimates provide no evidence that residential rates rise in the host utility over this horizon. One possible explanation for the negative point estimates is that a large new load spreads fixed network costs across more electricity sales. The full-sample and near-miss averages are imprecise, however, and we do not observe each utility's cost allocation or contract. The design therefore cannot identify that mechanism. Retail rates are also only one incidence margin. Electricity generation creates environmental costs that may be borne outside the host utility [Muller, 2026]. The sample covers cloud-era cohorts through 2021 followed to 2024, not the larger post-2022 artificial intelligence build-out.

D Supplementary results

Table D.1 reports confidence intervals and minimum detectable effects for the agglomeration estimates. Table 4 reports the capital footprint across estimators, transforms, and construction cohorts.

Table D.1: Precision of the agglomeration-channel estimates

Outcome	Comparison	Units	Estimate	95% CI	80% MDE
Advertised salary	0–25 vs 50–150 km	119	+1.3%	[-2.7, +5.5]	5.9%
Job postings	0–25 vs 50–150 km	120	+5.2%	[-5.3, +16.8]	16.2%
Supplier postings, 0–25 km	0–25 vs 50–150 km	120	+5.9%	[-15.0, +32.0]	37.0%
New firm registrations	Within-site trend break	117	-0.5%	[-2.6, +1.6]	3.0%
Business applications	National county	38	+1.4%	[-2.6, +5.7]	6.0%

Notes: Preferred estimate for each outcome, transformed from log points to percent. MDE is the minimum two-sided effect detectable with 80 percent power at the five-percent level, holding the reported standard error fixed. Salary, job postings, and supplier postings compare the pooled 0–25 km area with a 50–150 km comparison area, allowing labor-market effects to extend beyond the immediate site. Firm registrations use the Virginia and Texas within-site trend break. Business applications use the national county trend break. The unit count is facilities for the ring designs and treated counties for the business-application design.

Appendix Table D.2 compares the county-industry employment estimates under alternative treatments of cells suppressed by the Bureau of Labor Statistics.

Table D.2: QCEW employment estimates under alternative treatments of suppressed cells

Industry (NAICS)	Disclosed cells		Hierarchy imputation		State-only imputation	
	Effect	Pre-trend p	Effect	Pre-trend p	Effect	Pre-trend p
Data processing (518)	+5.2%	0.17	+25.9%*	0.64	+31.6%**	0.63
Telecommunications (517)	+5.2%	0.22	−1.4%	0.34	−5.9%	0.06
Web and information services (519)	+9.3%	0.12	+11.2%	< 0.01	+20.4%*	0.31
Computer systems design (5415)	−0.9%	0.17	+4.8%	0.04	−0.5%	0.02
Professional and technical services (541)	+6.2%**	< 0.01	+6.0%**	0.27	+6.6%***	0.02

Notes: Entries are cohort-weighted mean log changes over event years zero through five, transformed to percent, for hyperscale openings relative to not-yet-treated counties. The disclosed-cell specification treats suppressed employment as missing. The hierarchy imputation allocates suppressed cells to match state-industry and county-parent totals, using establishment counts and prior-year employment per establishment as starting weights. The fixed point reaches both margins in 92 percent of state-year problems; the remaining cells use the state-industry allocation. The state-only benchmark matches state-industry totals but does not use county-parent totals. Pre-trend p jointly tests event years −4 through −2. Inference uses 1,999 state-cluster bootstrap draws. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The outcome panels do not cover identical facilities. Table D.3 therefore repeats the national facility-ring estimates on a common hyperscale-site universe. The salary and total-posting estimates barely change. Supplier postings remain positive but imprecise, and foot traffic and spending remain statistically indistinguishable from zero. Rent remains positive and is marginally significant on the common subset, although its confidence interval remains wide. The qualitative conclusions are therefore not driven by differences in facility coverage across these panels.

Table D.3: Common-site robustness for hyperscale facility-ring outcomes

Outcome	Outcome-specific coverage	Common-site coverage
Advertised salary	+1.3% [-2.7, +5.5] ($N = 119$)	+1.3% [-3.1, +5.9] ($N = 110$)
Job postings	+5.2% [-5.3, +16.8] ($N = 120$)	+6.4% [-5.2, +19.4] ($N = 111$)
Supplier postings	+5.9% [-15.0, +32.0] ($N = 120$)	+12.6% [-11.2, +42.7] ($N = 111$)
Foot traffic	-3.3% [-11.8, +6.1] ($N = 187$)	-4.2% [-14.9, +7.9] ($N = 106$)
Commercial spending	+12.1% [-7.8, +36.4] ($N = 109$)	+3.0% [-12.0, +20.6] ($N = 80$)
Residential rent	+7.4% [-0.9, +16.4] ($N = 180$)	+10.0% [+0.4, +20.4] ($N = 118$)

Notes: Each cell reports the estimated opening-year jump, its 95% confidence interval, and the number of facilities. The common-site column restricts every outcome to the same 125 hyperscale facilities that have high- or medium-confidence dates, opened by the first quarter of 2024, and appear in all five national facility-ring panels. Salary, job postings, and supplier postings use the 0–25 versus 50–150 km trend-break design. Foot traffic, spending, and rent use the 0–5 versus 25–50 km trend-break design with state-by-month fixed effects. These outcomes use the trend-break form because foot traffic and spending do not have the flat pre-period that the plain within-site comparison requires (Table 3). Standard errors are clustered by facility. Cell-specific counts can be lower than 125 when an outcome is missing in one of the required rings or periods. The capital, two-state registration, and county analyses retain their design-specific samples because they use different timing requirements or geographic units.

Data Availability Statement

This study combines publicly available data with several proprietary datasets obtained under license. All code that constructs the analysis datasets from the raw sources and reproduces every table and figure will be deposited in a public data and code repository. The deposit includes the publicly available inputs and all author-constructed files, namely the data-center registry, the operating and construction-start dates, and the announced near-miss cohort. The proprietary inputs cannot be redistributed. For each, we report the provider and access terms and include the exact product identifiers, date ranges, geographic filters, and extraction code, so that a licensed user can regenerate the facility-by-ring panels and reproduce all results.

Publicly available data.

- *Satellite nighttime lights.* VIIRS Day/Night Band Monthly Composite (NOAA, product VCMCFG), accessed through Google Earth Engine.
- *Daytime imagery for construction dating.* Sentinel-2 and Landsat surface reflectance (ESA and USGS), accessed through Google Earth Engine.
- *Data-center registry.* The starting point is the IM3 Open Source Data Center Atlas (Pacific Northwest National Laboratory), augmented by the authors from operator press releases, state economic-development announcements, the Good Jobs First subsidy tracker, trade press,

DatacenterMap, Cloudscene, and data-center REIT filings. The compiled registry, dates, and announced near-miss cohort are in the deposit.

- *Construction timing.* Federal Aviation Administration Obstruction Evaluation / Airport Airspace Analysis filings (Form 7460-1).
- *Employment and establishments.* Quarterly Census of Employment and Wages (Bureau of Labor Statistics) and County Business Patterns (Census Bureau).
- *Migration and income.* IRS Statistics of Income county-to-county migration data.
- *Housing supply and population.* Census Building Permits Survey and 2020 Census Centers of Population (tract population-weighted centroids).
- *Workplace employment.* LEHD Origin-Destination Employment Statistics (LODES, Census Bureau).
- *Business formation.* Census Business Formation Statistics and state business-registration bulk files for Virginia, Texas, and Iowa (Secretaries of State and data.iowa.gov).
- *Local public finance.* Census of Governments and the Annual Survey of School System Finances (F-33).
- *Electricity.* U.S. Energy Information Administration Form 861.
- *Instrument-search inputs.* The InterTubes long-haul fiber map [Durairajan et al., 2015], the 1980 county urban-college share (IPUMS NHGIS), and transmission-substation locations (Homeland Infrastructure Foundation-Level Data).

Proprietary data (available under license, not redistributable).

- *Job postings.* LinkUp Job Records [LinkUp, 2025].
- *Advertised salaries.* WageScape salary postings [WageScape, 2022].
- *Residential rents.* RentHub Rental Data [RentHub, 2022].
- *Consumer spending.* SafeGraph Spend Patterns [SafeGraph, 2022].
- *Foot traffic.* Advan Weekly Patterns [Advan Research, 2025].
- *Firmographics.* Veridion firm master data.

The five Dewey-distributed datasets are available to researchers through a subscription to Dewey Data (`https://www.deweydata.io`). Veridion firmographics are available under license from Veridion (`https://veridion.com`). The authors accessed these data under institutional license and are not permitted to redistribute the raw files. The deposit provides the dataset identifiers and the extraction and aggregation code required to rebuild the analysis files from a licensed copy.