Hamm Institute for American Energy
OKLAHOMA STATE UNIVERSITY

Data Center Employment Forecast Analysis

November 2025

Angela Ryu & Shon R. Hiatt
Zage Business of Energy Initiative, University of Southern California

USC Marshall
Zage Business of Energy Initiative

Data Center Employment Forecast Analysis

Angela Ryu & Shon R. Hiatt

Zage Business of Energy Initiative, University of Southern California

December 3, 2025

This analysis provides a rigorous, field-based forecast of employment associated with the major U.S. data center buildout over the next five years (2025-2030). The model draws on the current inventory of operating and planned hyperscale and colocation facilities (including both wholesale and retail providers) and applies 2024-2025 field benchmarks from recent large-scale developments by major hyperscalers and colocation service providers. These data provide the most up-to-date evidence on workforce requirements by facility size and automation level (see Appendix A1 for sources and examples).

Workforce intensity varies sharply by facility scale. Construction requirements range from 0.7-2.0 workers per MW, while operations employ 0.15-0.35 full-time equivalents (FTEs) per MW once online. The most automated hyperscale campuses (>100 MW) can operate with as few as 20-30 permanent staff per 100 MW, whereas smaller enterprise or regional sites require substantially higher staffing due to lower automation and more complex operations.

1. Benchmark Summary by Facility Size

The data-center employment benchmarks presented below include both direct, on-site workers and indirect and induced employment. However, for consistency across sectors and comparability with the combined cycle gas turbine (CCGT) plant workforce analysis (Ryu & Hiatt, 2025), our ***baseline model relies on direct, on-site employment only***.

Direct employment captures workers physically engaged in construction or operations at each facility, measured as peak construction headcount and operational full-time equivalents (FTEs). Indirect and induced jobs represent the additional employment generated in supply chains and local service industries (see Appendix A1 for detailed sources). These broader effects are incorporated through a multiplier of 2 applied to the direct figures.

Facility Scale	**Capacity Range**	**Total Construction Workers per MW**		**Total Operation Workers per MW**		**Construction Duration**
		Peak Construction (Direct)	***Indirect & Induced***	***Operational FTE (Direct)***	***Indirect & Induced***	
Hyperscale / Efficient (Low)	100+ MW	0.7	0.7	0.15	0.15	18-24 months
Regional / Standard (Mid)	20-50 MW	1.3	1.3	0.24	0.24	24-30 months
Enterprise / Labor-	5-20 MW	2	2	0.34	0.34	30-36 months

Intensive (High)						

2. **Detailed Benchmark Ranges: Workforce (Direct Employment) per MW**
   a. **Construction Staffing Intensity**

Range	Peak Construction Workers per MW	Typical Scenario
Low (Hyperscale / Efficient)	0.6-0.8	Large campuses (≥100 MW) with modular design, prefabricated components, and streamlined contracting.
Medium (Standard / Regional)	~1.0-1.3	20-50 MW builds using standard construction methods.
High (Enterprise / Labor-Intensive)	1.5-2.0	Smaller facilities (<20 MW) or highly customized builds; complex permitting or brownfield conversions.

   b. **Operational Staffing Intensity**

Range	Operational Employees per MW	Typical Scenario
Low (Hyperscale / Automated)	0.15-0.20	Highly automated hyperscale campuses with centralized monitoring and low on-site staff.
Medium (Standard)	0.24-0.28	Regional or colocation data centers with mixed automation and diverse client operations.
High (Enterprise / Manual)	0.30-0.35	Smaller enterprise facilities with manual operations, limited automation, and 24/7 on-site presence.

3. **Distribution of Data Centers by Size and Project Stage**

The data center inventory integrates facility-level information for over 3,000 hyperscale and colocation facilities on operational status (Planned, Under Construction, Operational; as of Q4 2024) together with the imputed year of facility completion and, where available, a decommissioning year. This report relies on reported data from 451 Research (part of S&P Global) wherever possible while imputing the year of completion for projects lacking explicit schedule information. Facilities without reported total utility capacity were excluded from this analysis to ensure consistency in workforce estimation.

It allows each facility to transition through a sequence: *Planned - Under Construction - Operating,* over the 2025 - 2030 analysis window, aligned with the typical two-year construction duration observed in large U.S. data-center projects (Crosby, 2014; Rahkonen et al., 2020).[1] When a valid completion year exists, it defines the schedule directly and the COD equals the reported or imputed 'built year.' The construction start year is assumed to be COD - 2, reflecting the typical two-year build duration observed in large U.S. data-center projects. If the completion year is missing, timing is inferred from the facility's reported status.

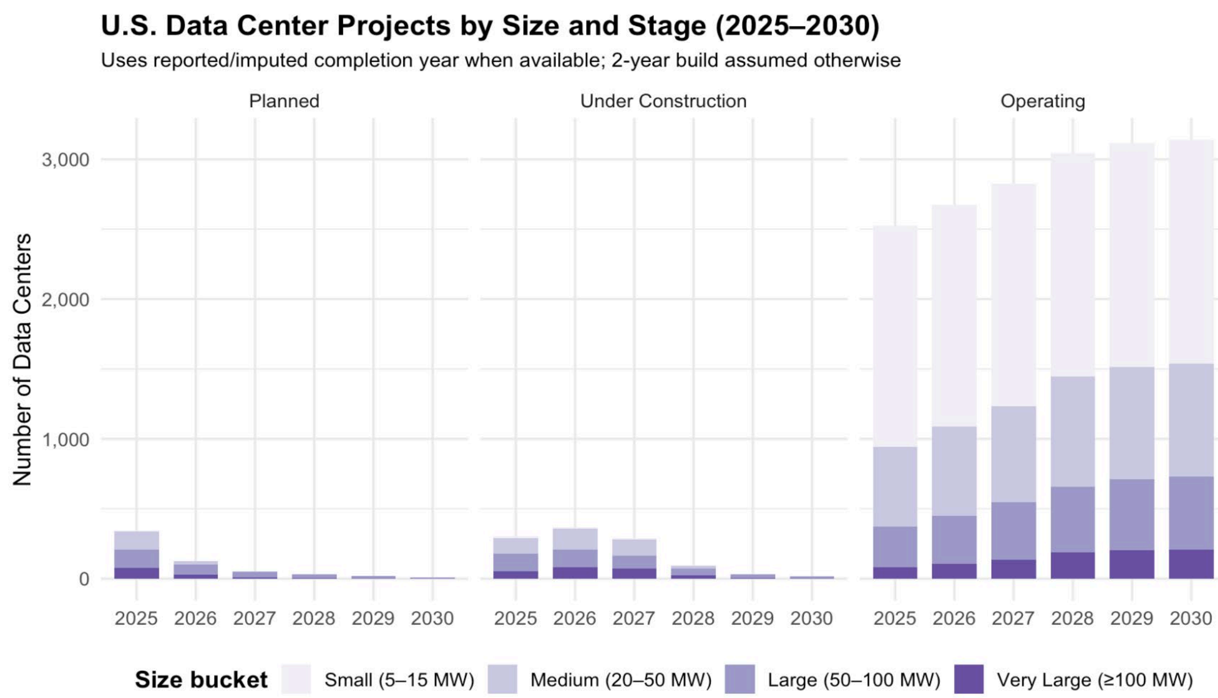

4. Data Center Workforce Projection Model

Similar to the CCGT (Combined Cycle Gas Turbine) workforce projection model, this model estimates future employment associated with data center projects by linking projected total utility additions (2025-2030) with empirically observed staffing intensities that vary by plant size. Both construction and operations workforces are expressed as:

- A_t= MW of new data center facility **starting construction** in year t
- C_t= MW of new data center facility **entering operation** in year t
- W_c= construction workers per MW (size-specific; Appendix A2)
- W_o= operations employees per MW (size-specific; Appendix A2)
- D= average construction duration (years) = 2 years

[1] Industry reports note that full data-center delivery has historically required 18-36 months (Crosby, 2014), while recent modular approaches can shorten physical build times to 8-12 months under ideal conditions (Rahkonen et al., 2020). In the United States, however, permitting, utility interconnection, and power-delivery lead times typically extend overall schedules. Accordingly, the model applies a conservative two-year construction duration as a representative average for large U.S. data-center projects.

- α= ratio of average annual to peak construction workforce = 0.6

a. Construction Workforce

Under-construction MW (stock): This term represents the total megawatts of projects still under construction in year t. Because each facility takes about D = 2 years to complete, we sum all projects that began construction in the current year and the last year.

$$U_t = \sum_{k=0}^{D-1} A_{t-k}$$

Peak construction workforce (stock): Calculated by multiplying under-construction MW in each facility size class by its construction intensity (workers per MW):

$$W_t^{\text{peak}} = \sum_{\text{size}} W_c(\text{size}) \cdot U_t(\text{size})$$

New construction hires in year t (flow): Estimates the annual inflow of new construction jobs created in year t by projects starting that year. Because not all workers are on-site for the full year, the new construction hires in year t is scaled by α = 0.6:

$$W_t^{\text{flow,constr}} = \sum_{\text{size}} \alpha \cdot W_c(\text{size}) \cdot A_t(\text{size})$$

b. Operation Workforce

Once facilities are completed and enter service, they require a permanent operational staff, typically expressed as full-time equivalents (FTEs) per MW.

Operational workforce (stock): Calculated by multiplying total MW of operating data centers in year t by its operations intensity (employees per MW), aggregated across all facilities online

$$W_t^{\text{stock, ops}} = \sum_{\text{size}} W_o(\text{size}) \cdot MW_{\text{oper}}(\text{size}, t)$$

New O&M hires in year t (flow): Incremental O&M positions added in year t as newly built facilities begin operation.

$$W_t^{\text{flow,ops}} = \sum_{\text{size}} W_o(\text{size}) \cdot C_t(\text{size})$$

5. Data Center Workforce Estimation Results

We then apply the equations from previous section to the data center project pipelines in 2025-2030. We provide the estimation under three scenarios (each with size-specific workforce intensities; see Appendix A2 for the full list of parameters and weights used): Efficient Builds / Highly Automated (***Low***), Standard Projects /Typical Automation (***Mid***), and Complex or Labor-Intensive Sites (***High***).

The data center workforce outlook for 2025-2030 shows a clear two-phase pattern: a construction-driven surge followed by steady growth in long-term operations employment. Construction activity dominates the workforce profile, peaking around 2026-2027 when overlapping projects are at their height.

Under the standard build scenario, total active construction employment ("stock") reaches roughly 30,000 workers, with annual new hires ("flow") rising over 15,000 at peak. By 2027, as the pace of new projects starts to decline, construction employment gradually decreases, though ongoing multi-year builds sustain a portion of the workforce into 2030. As new capacity enters service, operations employment becomes the dominant, long-term source of jobs. The total active operations staff ("stock") grows steadily to about 50,000 full-time equivalents (FTEs) by 2030, with annual new hires ("flow") peaking at about 6,000 FTEs in 2028 (under the standard scenario).

When indirect and induced employment is included using the conservative 2.0, total job impacts approximately double. Under this broader measure, total active construction employment in a given year ("stock") peaks at roughly 60,000 total jobs in 2026, while total operations-related employment approaches 110,000 by 2030 (see Appendix A3). These results highlight both the near-term labor intensity of data-center construction and the substantial, sustained operations workforce required to support the large-scale digital infrastructure coming online over the second half of the decade.

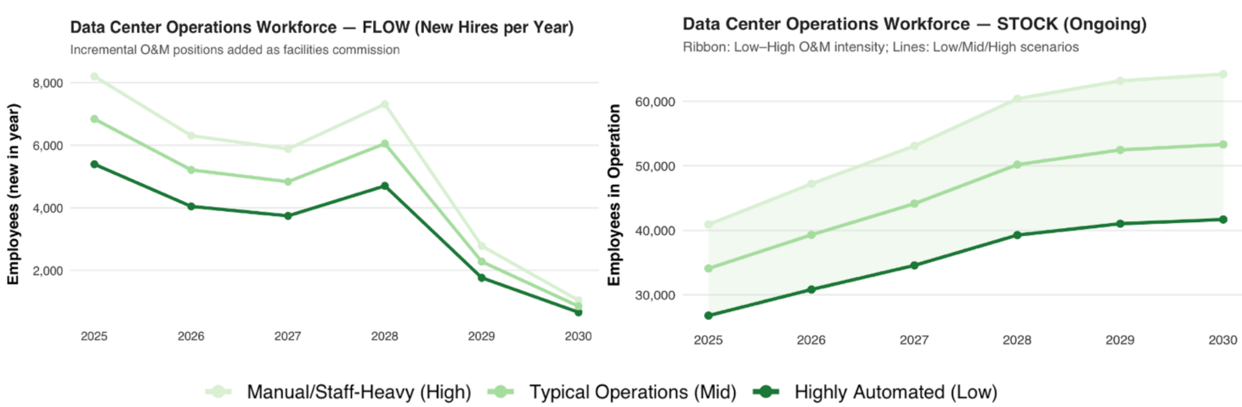

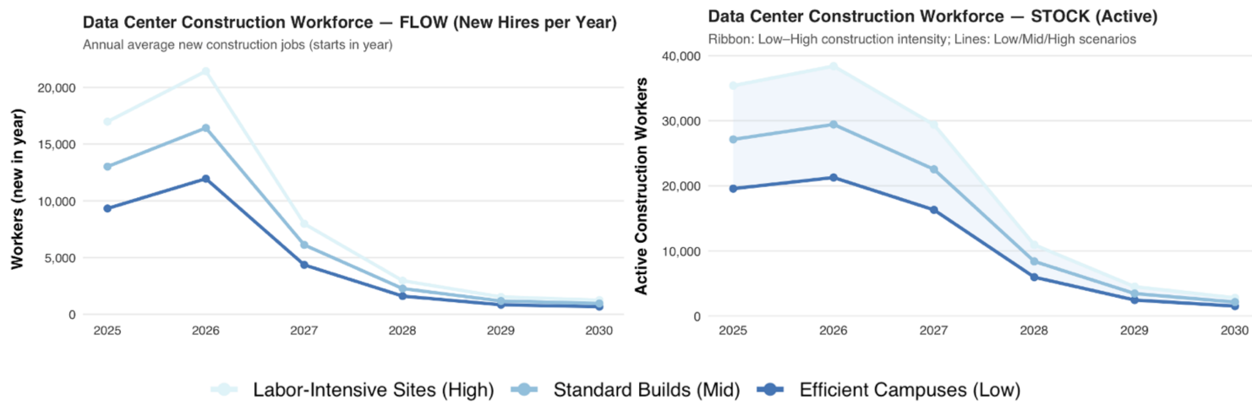

6. Conclusion

The U.S. data center buildout for hyperscale and colocation facilities through 2030 represents a sustained source of infrastructure employment in the modern economy. Construction activity will continue to drive near-term labor demand through 2027, while long-term operations employment expands as new facilities come online. Even as automation improves efficiency, the unprecedented scale of hyperscale and colocation development ensures substantial ongoing workforce needs. These findings complement broader industry warnings, such as OpenAI's projection that upcoming AI infrastructure could require one-fifth of the nation's skilled trade workforce, and highlight the urgency of addressing emerging construction and technical labor shortages in the years ahead.

References

Amazon. (2025, October 29). *AWS's Project Rainier: The world's most powerful computer for training AI*. About Amazon. https://www.aboutamazon.com/news/aws/aws-project-rainier-ai-trainium-chips-compute-cluster

Crosby, C. (2014, July 30). *Data Center Cost Myths: SCALE*. Uptime Institute Blog. https://journal.uptimeinstitute.com/the-myth-of-data-center-cost/

DPR Construction. (n.d.). *Richland Parish Data Center*. DPR Construction. Retrieved December 3, 2025, from https://www.dpr.com/richland-parish-data-center

Econsult Solutions. (2019). *The Economic and Revenue Impact of Data Centers in Pennsylvania*. Econsult Solutions. https://econsultsolutions.com/wp-content/uploads/2019/06/Data-Center-Final-Report_Econsult-Solutions.pdf

Microsoft. (2025, August 19). Datacenter investments drive economic growth in Central Washington. *Microsoft*. https://local.microsoft.com/blog/datacenter-investments-drive-economic-growth-in-central-washington/

O'brien, M. (2025, September 23). *OpenAI shows off Stargate AI data center in Texas and plans 5 more elsewhere with Oracle, Softbank*. AP News. https://apnews.com/article/openai-stargate-oracle-data-center-0b3f4fa6e8d8141b4c143e3e7f41aba1

Rahkonen, T., Lawrence, A., & Ascierto, R. (2020). Best-in-class data center provisioning. *Uptime Institute Intelligence*, *36*. https://intelligence.uptimeinstitute.com/sites/default/files/mammoth/2023/09/12/Keynote%20report%2036_Best-in-class%20data%20center%20provisioning.pdf

Reilly, G. (2025, October 30). *First part of Amazon project west of South Bend now operational, running AI models*. South Bend Tribune. https://www.southbendtribune.com/story/news/local/2025/10/30/first-part-of-amazon-web-services-project-near-new-carlisle-operational/86983620007/

Ryu, A., & Hiatt, S. R. (2025). *Combined Cycle Gas Turbine (CCGT) Plants Employment Forecast analysis*. Zage Business of Energy Initiative, University of Southern California.

Srivalli, K. V. (2025, November 21). *Top 10 Upcoming Data Centers in the USA (2025)*. Market Research Company - Blackridge Research & Consulting™. https://www.blackridgeresearch.com/blog/upcoming-largest-data-center-projects-in-united-states-usa

The Office of Alabama Governor. (2024, May 2). Governor Ivey Announces Meta Plans to Build $800 Million, Next-Generation Data Center in Montgomery - [5/2/2024]. *The Office of Alabama Governor Kay Ivey*. https://governor.alabama.gov/newsroom/2024/05/governor-ivey-announces-meta-plans-to-build-800-million-next-generation-data-center-in-montgomery/

Underwood, J. (2018, April 9). Google kicks off construction on $600M Alabama data center. *Made in Alabama*. https://www.madeinalabama.com/2018/04/google-kicks-off-construction-on-alabama-data-center/

WSBT 22. (2025, October 31). *Amazon's "Project Rainier" starts operations in New Carlisle*. WSBT. https://wsbt.com/news/local/amazons-project-rainier-starts-operations-create-jobs-ai-billion-dollar-project-near-new-carlisle-st-joseph-county-indiana

WXYZ Detroit. (2025, October 30). *Whitmer, OpenAI announce new AI datacenter for Saline Township, construction set for early 2026*. WXYZ 7 News Detroit. https://www.wxyz.com/news/region/washtenaw-county/whitmer-openai-announce-new-ai-datacenter-for-saline-township-construction-set-for-early-2026

Appendix

A1. Source and Methods for Benchmark Ratio Calculation

Public reporting on data-center workforce varies substantially across operators and projects. No single source provides a complete, consistent breakdown of direct vs. indirect employment, construction vs. operations staffing, or MW-specific headcounts. Some announcements disclose only total job creation, others separate construction and operational roles but do not report MW, and some provide MW capacity without corresponding staffing detail. Because of these inconsistencies, this analysis relies on facility-level and regional examples that disclose at least one dimension of workforce information (peak construction headcount, permanent operations staffing, or total jobs), and uses these values to construct conservative benchmark ranges applied uniformly across the model. The full benchmark values and how they enter the model are provided in Appendix A2.

1. **Direct Construction Jobs/MW** are from (reported construction headcount)/(MW delivered) for actual mega-campus projects.
    - Meta Richland Parish (Louisiana): 2,000 MW capacity, 5,000+ peak construction workers = 2.5 jobs/MW (DPR Construction, n.d.)
    - OpenAI Stargate (Abilene, TX): 4,500 MW capacity, 6,000+ daily workers = 1.3 jobs/MW (O'brien, 2025)
    - Vantage Frontier (Texas): 1,400 MW capacity, 5,000+ construction jobs = 3.6 jobs/MW (Srivalli, 2025)
    - Amazon Project Rainier (Indiana): 2,200 MW capacity, $11 billion project with 1,000+ construction and operations jobs (Amazon, 2025; Reilly, 2025; WSBT 22, 2025)
        - Note: Amazon's Project Rainier reports "1,000+ construction and operations jobs" without separating the two, so it is used as a secondary source.

2. **Operations FTE/MW** reflect permanent on-site workforce levels and shift/work schedule structures at major hyperscale campuses. Because operators vary in how they disclose operations staffing (and often do not report MW capacity and O&M headcount in the same document) the FTE/MW values below pair publicly disclosed staffing numbers with separately reported or analyst-estimated MW capacities. Examples include:
    - Meta Richland Parish (Louisiana): ~2,000 MW campus with 500+ operational jobs, $10B investment = 0.25 job/MW (DPR Construction, n.d.)
    - Meta Montgomery (Alabama): An $800 million AI-optimized data center is expected to support 100 operational jobs once complete. = 0.83 job/MW (The Office of Alabama Governor, 2024); campus capacity estimated at ~120 MW (451 Research)
    - Google Jackson County (Alabama): A $600 million data center project is expected to

create around 100 "full-time jobs" when operational (Underwood, 2018) ; The campus capacity is estimated at ~100 MW = 1 job/MW (451 Research)

- Note: Public announcements of "full-time jobs" often include roles beyond the core O&M workforce (e.g., engineering, administrative, or regional support positions) and therefore tend to overstate MW-specific staffing levels. These ratios are treated as upper-bound examples rather than facility-level O&M staffing requirements.

3. **A multiplier of 2.0** is used as a conservative total-employment multiplier, meaning total jobs are estimated as 2× direct on-site employment. In this model, each direct job is assumed to support one additional off-site job in the regional supply chain and local economy. This conservative value reflects the lower bound of employment effects observed in regional data-center impact studies and provides a unified assumption across construction and operations, since project-level studies rarely disclose separate multipliers for these two phases in a consistent way.
    - Microsoft Central Washington (regional study): For construction, 1.1 additional jobs per direct construction job (implied multiplier of 2.1); for operations, two additional jobs per direct on-site job (implied multiplier of 3) (Microsoft, 2025)[2]
    - Meta Pennsylvania statewide study: 9,549 direct data-center jobs support approximately 30,000 total jobs (implied multiplier of 3.1) (Econsult Solutions, 2019)

Recent project announcements show similar or stronger spillovers:

- OpenAI (Michigan): 450 permanent on-site, 1,500 additional community jobs (implied multiplier of 4.3) (WXYZ Detroit, 2025)
- Meta Richland Parish (Louisiana): 500 operational + over 1,000 indirect/induced (implied multiplier of 3) (DPR Construction, n.d.)

A2. Parameters Used in the Model

Size Buckets Used in the Model

Facility Scale	Total Utility Range (MW)	Description
Small	5-15 MW	Enterprise or edge data centers with high labor intensity and limited automation.
Medium	20-50 MW	Standard regional or colocation facilities built with conventional methods and moderate automation.

[2] See infographics for the report at https://local.microsoft.com/wp-content/uploads/2025/06/Central-Washington-Infographic-FINAL-2.pdf

Large	50-100 MW	Multi-building regional campuses that leverage modular construction and partial automation.
Hyperscale	≥100 MW	Advanced large-scale campuses with highly efficient, prefabricated builds and centralized monitoring.

Construction Labor Intensity (W_c =construction workers per MW)

Size Bucket	**Low (Efficient Builds)**	**Mid (Standard Projects)**	**High (Complex Sites)**
Hyperscale (≥100 MW)	0.70	1.00	1.30
Large (50-100 MW)	0.80	1.15	1.50
Medium (20-50 MW)	1.00	1.30	1.70
Small (5-15 MW)	1.30	1.70	2.00

Operations Staffing Intensity (W_o = operation workers per MW)

Size Bucket	**Low (Highly Automated)**	**Mid (Typical Automation)**	**High (Labor-Intensive)**
Hyperscale (≥100 MW)	0.18	0.24	0.30
Large (50-100 MW)	0.20	0.26	0.32
Medium (20-50 MW)	0.22	0.28	0.33
Small (5-15 MW)	0.25	0.30	0.35

How the Weights Are Applied

1. Each data center facility in the dataset is assigned to a size bucket based on its total utility capacity (MW).
2. For each year and scenario (Low, Mid, High), the model multiplies the megawatts of new or operating capacity in that bucket by the corresponding construction (Wc) and operations (Wo) workforce intensity values.
3. Construction employment is estimated for both *stock* (projects active across multi-year build periods) and *flow* (new starts in a given year). Operations employment is also calculated for both *stock* (cumulative operating staff) and *flow* (incremental hires as facilities begin service).

4. Results are then aggregated across all size buckets to produce national workforce estimates under each automation and efficiency scenario, capturing variation in labor intensity across hyperscale, large, medium, and small data center development.

A3. Data Center Workforce Estimation Results (Total Workforce; includes indirect and induced)