BofA SECURITIES

Industrials/Multi-Industry

Who Makes the Data Center 2026

Primer

A global view of a $1.0tr data center market

We estimate data center spending has grown from $275bn in 2022 to nearly $1.0tr in 2026E. BofA analysts forecast this to reach $2.1tr by 2028. This 2028 forecast has been revised upwards from our prior $0.9bn estimate in last year's edition of *Who Makes the Data Center*. Given 18-24 month build times, much of this spending is already financed, permitted, and contracted.

The key products, markets, and vendors for data centers

This report evaluates addressable markets across IT & networking equipment, electrical equipment, thermal equipment, and engineering & construction. We show average content per megawatt (MW) and market shares for 13 separate categories. We estimate the all-in cost of building at data center at $49mn/MW, up $10mn/MW from last year's edition of this report. While inflation for IT equipment drove the majority of this, increasing complexity in data center infrastructure is resulting in higher building costs. On a like-for-like basis, data center infrastructure costs are 10-15% higher y/y.

What's changing in data center infrastructure?

■ New artificial intelligence (AI) focused data center developers, frequently called neoclouds, have become an increasingly large part of this market. We forecast neoclouds to own 6% of the total energized data center market in 2026E, up from less than 0.3% in 2023. We provide a historical context on the size and ownership of the global data center market.

■ AI development in semiconductor and networking is requiring changes in data center infrastructure. This includes a transition from lower voltage alternating current (AC) to higher voltage direct current (DC) in electrical distribution. Initial direct current deployments will likely be by adding power sidecars, which we expect to start shipping in 2027.

■ Given the early stage of solid-state transformers within the data center market, we provide an overview of existing and emerging vendors. We also note there are alternatives to solid-state transformers, which have more established supply chains, but have trade-offs in terms of complexity, integration, and footprint.

■ In response to labor-related bottlenecks, the industry has sought to shift more assembly into factories and away from construction sites. This has driven high-profile orders totaling several billion dollars so far this year. We provide a categorization of the modular market and provide a 2026 market size.

■ The heat rejection architecture used in direct-to-chip liquid cooling continues to evolve. We find more data center developers are building dedicated IT cooling loops. This creates a separate (higher) temperature loop for rejecting heat from semiconductors and a different (lower) temperature loop for traditional air cooling.

23 October 2026

Equity
Americas
Industrials/Multi-Industry

Andrew Obin
Research Analyst
BofAS
+1 646 855 1817
andrew.obin@bofa.com

Vivek Arya
Research Analyst
BofAS
+1 646 855 1755
vivek.arya@bofa.com

Michael Feniger
Research Analyst
BofAS
+1 646 855 1923
michael.feniger@bofa.com

Michael J. Funk
Research Analyst
BofAS
+1 646 855 5664
michael.j.funk@bofa.com

Benjamin Heelan >>
Research Analyst
Merrill Lynch (DIFC)
+44 20 7996 5723
benjamin.heelan@bofa.com

Tal Liani
Research Analyst
BofAS
+1 646 855 5107
tal.liani@bofa.com

Wamsi Mohan
Research Analyst
BofAS
+1 646 855 3854
wamsi.mohan@bofa.com

David Ridley-Lane, CFA
Research Analyst
BofAS
+1 646 855 2907
david.ridleylane@bofa.com

>> Employed by a non-US affiliate of BofAS and is not registered/qualified as a research analyst under the FINRA rules.
Refer to "Other Important Disclosures" for information on certain BofA Securities entities that take responsibility for the information herein in particular jurisdictions.
BofA Securities does and seeks to do business with issuers covered in its research reports. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of this report. Investors should consider this report as only a single factor in making their investment decision.
Refer to important disclosures on page 53 to 55. **13031711**

Timestamp: 23 October 2026 05:00AM EDT

Contents

What has changed	3
Data center market to hit $2.1tn in 2028E	5
Content per megawatt cost analysis	5
Types of data centers	7
Extreme codesign driving M&A	14
Making sense of the modular landscape	16
ELECTRICAL EQUIPMENT	19
An evolutionary path to direct current	19
Uninterruptible power supply (UPS)	26
Switchgear	27
Power distribution equipment	28
THERMAL EQUIPMENT	29
How thermal architecture is evolving	29
Chillers	35
Cooling towers	38
Air cooling equipment	39
Liquid cooling equipment	40
Comparing two reference designs	41
On-site prime power	43
Engineering	47
Construction	48
Backup generators	48
Servers	49
Networking equipment	50

What has changed

1. Higher industry forecasts (pages 5-6)

Upward revisions to cloud service provider capex estimates, capital raising activities by existing and newer developers, and a surge in construction starts all bode well for global data center spending in 2027-28. Our current 2028 forecast is 122% higher than the 2025 edition of *Who Makes the Data Center*.

Exhibit 1: How our forecasts have changed – comparison of 2028E
2028 IT equipment spend forecast has more than doubled, infrastructure spend revised up by 70%

($bn)	2028E - Current	2028E - Prior	% change
AI-related IT Equipment	1,496	629	138%
Non-AI IT Equipment	319	163	95%
Subtotal: IT Equipment	**1,815**	**792**	**129%**
Construction & installation	114	68	67%
Electrical	75	34	123%
Thermal	45	23	97%
Generators	24	14	70%
Engineering	12	8	53%
Sub-total: Infrastructure	**270**	**147**	**84%**
Total	**2,085**	**939**	**122%**

Source: BofA Global Research
Note: Prior forecast as published in Who Makes the Data Center 2025 (September 15, 2025)

BofA GLOBAL RESEARCH

2. Higher voltage direct current nears (pages 20-23)

Relative to a year ago, the path to higher voltage direct current adoption is clearer. Initial deployments will be with power sidecars. These will convert to 400- or 800-volt direct current close to the rack. We continue to see initial shipments of power sidecars in 2027.

3. Look at solid-state transformer suppliers (pages 23-24)

Given the early stage of solid-state transformers within the data center market, we provide an overview of existing and emerging vendors.

4. Extreme codesign driving M&A activity (pages 15-16)

Extreme codesign drives deeper integration across all data center infrastructure. Vendors need to be somewhat agnostic about specific point solutions and willing to invest in R&D, product development, and technology-focused acquisitions. In the past, there was a clear divide among vendors selling into the white space of a data center (i.e., the data hall housing the racks and servers) and gray space (i.e., power & cooling equipment). Extreme codesign merges these areas. We list acquisitions in the data center infrastructure space totaling over $40bn since 2024.

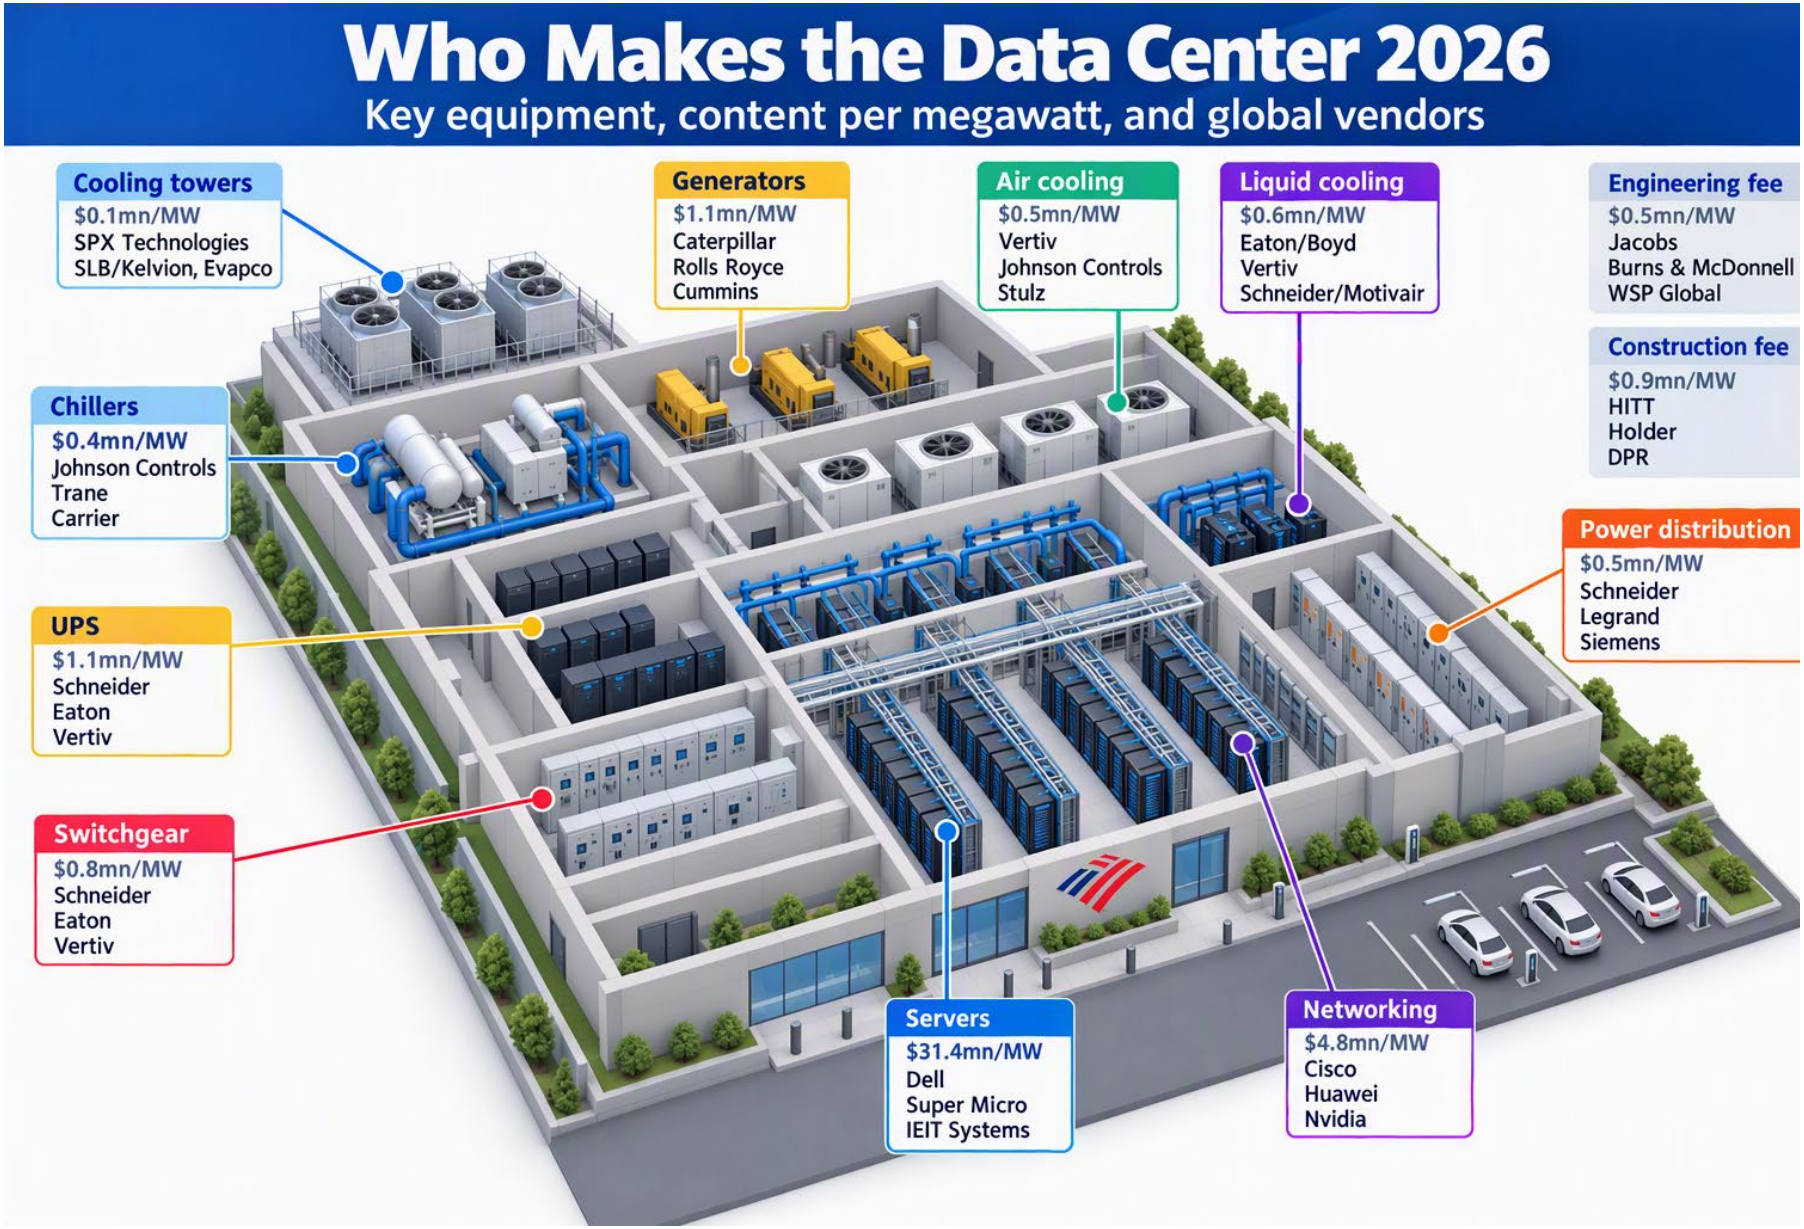
5. Modular products take off (pages 17-20)

In response to labor-related bottlenecks, the industry has sought to shift more assembly into factories and away from construction sites. This has driven high-profile orders totaling several billion dollars so far this year. We provide a categorization of the modular market and provide a 2026 market size.

6. Dedicated IT cooling loops (pages 31-33)

The heat rejection architecture used in direct-to-chip liquid cooling continues to evolve. We find more data center developers are building dedicated IT cooling loops. This creates a separate (higher) temperature loop for rejecting heat from semiconductors and a different (lower) temperature loop for traditional air cooling.

Exhibit 2: Who Makes the Data Center 2026
Estimated capital costs of $49mn per megawatt



Source: BofA Global Research

Data center market to hit $2.1tn in 2028E

We forecast global data center capex to rise from $0.6tn in 2025 to nearly $1.0tn in 2026E. We forecast AI adoption to drive a 37% market-wide CAGR over 2025-30E.

Exhibit 3: We forecast the global data center market to grow at a 37% CAGR over 2025-30E to reach over $2.9tn in 2030
AI adoption drives strong growth in infrastructure

($bn)	2022	2023	2024	2025	2026E	2027E	2028E	2029E	2030E	'25-'30E CAGR
AI Servers	16	49	132	228	445	806	1,172	1,470	1,707	50%
AI Networking	6	10	21	35	96	172	252	329	394	63%
AI Storage	1	3	8	16	28	49	71	90	105	47%
AI-related IT Equipment	23	62	161	278	569	1,026	1,496	1,890	2,206	51%
Non-AI IT Equipment	204	176	173	227	270	301	319	342	356	9%
Subtotal: IT Equipment	**227**	**238**	**334**	**506**	**839**	**1,327**	**1,815**	**2,231**	**2,562**	**38%**
Construction & installation	23	26	35	49	66	89	114	137	159	27%
Electrical	15	18	21	30	40	56	75	90	102	28%
Air-cooling	*6*	*8*	*10*	*14*	*18*	*24*	*30*	*36*	*40*	24%
Liquid-cooling	*1*	*1*	*2*	*5*	*8*	*11*	*15*	*19*	*22*	38%
Thermal	7	9	12	18	26	36	45	54	63	28%
Generators	4	5	7	10	14	19	24	29	34	28%
Engineering	2	3	4	6	8	10	12	14	16	23%
Sub-total: Infrastructure	**51**	**60**	**80**	**112**	**153**	**210**	**270**	**325**	**374**	**27%**
Total	**278**	**297**	**413**	**618**	**993**	**1,537**	**2,085**	**2,556**	**2,936**	**37%**

Source: BofA Global Research

BofA GLOBAL RESEARCH

Content per megawatt cost analysis

Exhibit 4: Average US new build data center costs $49mn per megawatt
Servers are the largest area of data center capex

Category	Average New Build Cost/MW	% of total	AI-focused Cost/MW	% of total
Servers	31,400,000	64%	37,500,000	66%
Networking	4,800,000	10%	5,800,000	10%
Storage	2,100,000	4%	2,100,000	4%
Subtotal IT equipment	**38,300,000**	**78%**	**45,400,000**	**80%**
Uninterruptible Power Supplies (UPS)	1,060,000	2%	743,000	1%
Transformers/Switchgear	750,000	2%	910,000	2%
Busway & other power distribution	495,000	1%	495,000	1%
Rack-level equipment	235,000	0%	512,500	1%
Subtotal electrical equipment	**2,540,000**	**5%**	**2,660,500**	**5%**
Air cooling equipment	510,000	1%	510,000	1%
Coolant distribution units (CDUs)	425,000	1%	637,500	1%
Cold plates, manifolds, etc.	175,000	0%	262,500	0%
Chiller (285-ton capacity)	400,000	1%	400,000	1%
Cooling tower	90,000	0%	135,000	0%
Subtotal thermal equipment	**1,600,000**	**3%**	**1,945,000**	**3%**
Backup diesel generator	**1,050,000**	**2%**	**1,050,000**	**2%**
Engineering fee	475,000	1%	570,000	1%
General contractor overhead & profit	920,000	2%	920,000	2%
Building costs	1,540,000	3%	1,001,000	2%
Installation costs	2,710,000	6%	3,252,000	6%
Subtotal E&C	**5,645,000**	**11%**	**5,743,000**	**10%**
Grand total	**49,135,000**	**100%**	**56,798,500**	**100%**

Source: BofA Global Research. Note: Estimates for delivered projects of 50 MW or more. Excludes land costs, utility connections, and general site works.

BofA GLOBAL RESEARCH

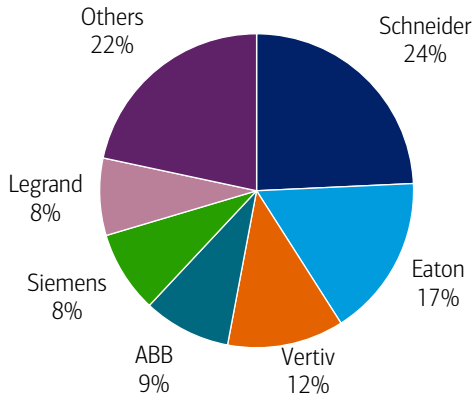
We estimate the all-in cost of building an average new build data center in the US to be $49mn per megawatt. IT-related costs drive 78% of the total cost. Electrical and thermal equipment are 5% and 3%, respectively. These estimates include a mix of traditional data centers, using lower rack densities and traditional networking.

Looking more specifically at AI-only data centers, these have a higher overall cost ($57mn per megawatt). However, higher IT equipment costs drive $7mn of the $8mn increase. We assume building square footage is ~35% lower given greater rack density (electricity usage per IT rack), which helps to offset some of the higher infrastructure-related costs.

Vendor shares infrastructure product categories

Across data center electrical products, we estimate Schneider is currently the market share leader in this $30bn market. Across all thermal products, we estimate Vertiv is the market share leader in this $18bn market.

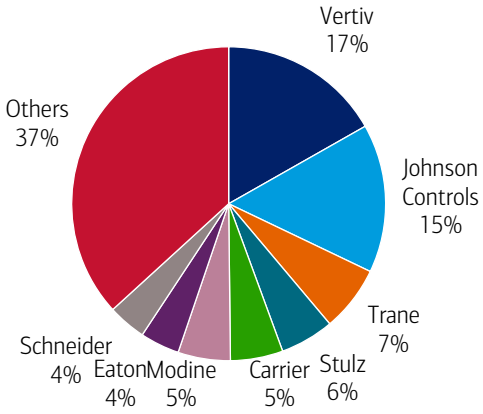
Exhibit 5: Data center electrical equipment is a $30bn market (2025)
Six largest vendors have nearly 80% market share



Source: BofA Global Research

BofA GLOBAL RESEARCH

Exhibit 6: Data center thermal equipment is a $18bn market (2025)
Eight largest vendors have over 60% market share



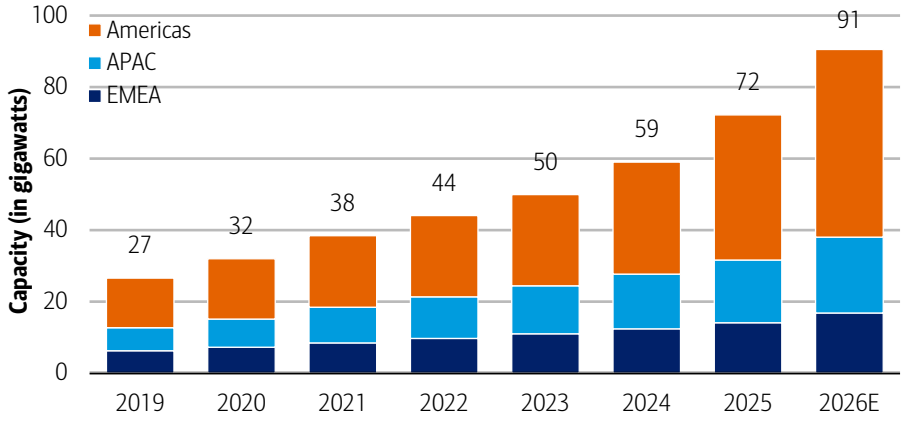
Source: BofA Global Research

BofA GLOBAL RESEARCH

Regional breakout: Americas home to over half of capacity

We estimate data centers (measured by electrical capacity) grew at a 24% CAGR over 2024-26E. EMEA has been a relative laggard (17% CAGR), Asia Pac a touch better (18%), while the Americas region has outperformed (29%).

Exhibit 7: Americas is largest region for data center capacity with over 50% over total
All regions continue to add capacity



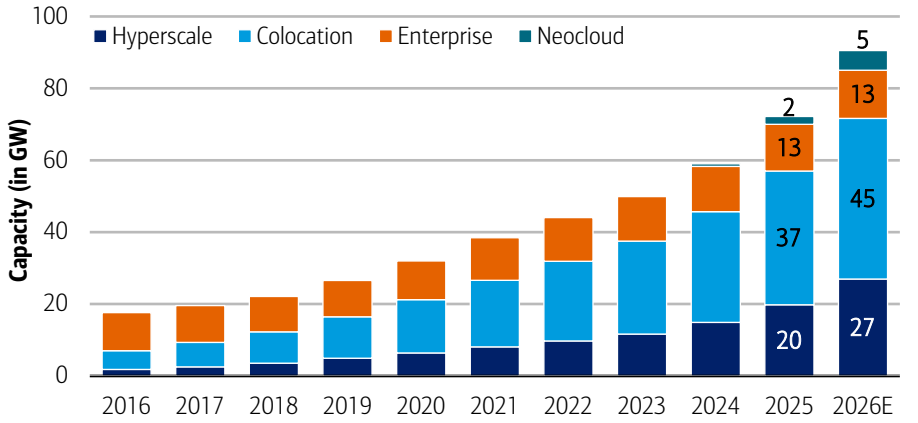
Source: BofA Global Research, Schneider Electric, JLL, IEA

BofA GLOBAL RESEARCH

Types of data centers

- **Enterprise:** a facility owned by a single organization housing its IT infrastructure. Typically, they are owned by large corporations, financial institutions, or government agencies. Over the past ten years, the square footage growth of these data centers has been largely flat, although modernization and IT refreshes continue.
- **Single-tenant colocation:** a facility owned and managed by a third-party and leased to a single tenant. Over the past ten years, cloud service providers have used this method to expand into new geographies. Initial lease terms are typically long with options to extend (e.g., 15-year initial lease with three additional 5-year options).
- **Multi-tenant colocation:** a facility owned and managed by a third-party and leased to multiple tenants. Rents are generally based on a combination of power usage and number of racks. Tenants benefit from shared services (e.g., network connectivity, physical security). This segment can be subdivided into retail (smaller space commitments, less design flexibility) and wholesale (larger commitments, more flexibility in design).
- **Hyperscale:** a large data center (generally above 25 megawatts) engineered to provide maximum uptime (i.e., Tier 4 ranking in Uptime Institute's classification), support distributed computing (e.g., sharing workloads across servers and sites), and scalability. Distributed computing and hyperscale data centers are closely associated with cloud service providers, such as Amazon Web Services and Microsoft Azure. However, not every hyperscale data center is used for cloud services.
- **Neocloud:** a cloud service provider purpose-built for AI workloads, offering GPU-as-a-Service and AI-optimized infrastructure that enables customers to train and run large AI models without owning the underlying hardware. Unlike hyperscalers, neoclouds typically focus on delivering large-scale GPU capacity, faster deployment, specialized AI software stacks, and lower-cost access to compute. Examples include CoreWeave and Nebius.

Exhibit 8: Globally, colocation is the largest data center type
Neoclouds on track to more than double capacity in 2026



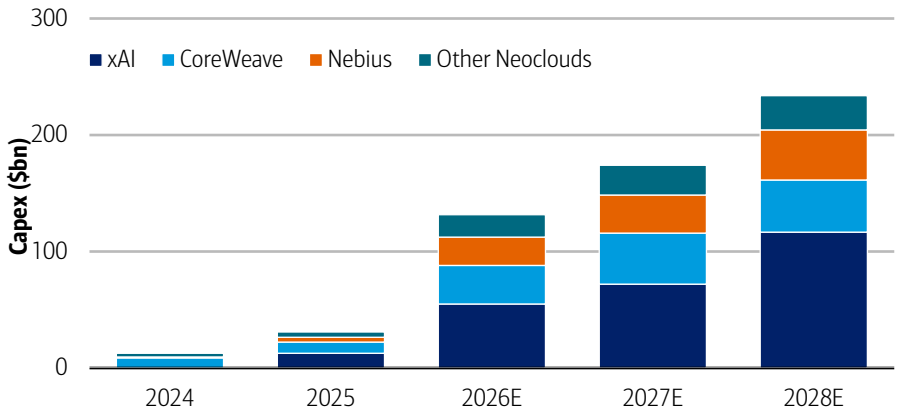
Source: BofA Global Research

BofA GLOBAL RESEARCH

As measured by electrical capacity, we estimate global data centers have grown at an 18% CAGR over 2016-26E. This measures energized IT load capacity and excludes data centers that are under construction.

Neoclouds: new kids on the block

Exhibit 9: US neocloud capex forecasted to top $230+bn by 2028
Capex figures include third-party leases



Source: Company reports, BofA Global Research estimates, Bloomberg

BofA GLOBAL RESEARCH

Hyperscalers shifting position in AI compute

Hyperscalers have the broadest global footprints and integrate compute, storage, databases, analytics, and applications into unified platforms. Their advantage lies in scale, enterprise breadth, and the ability to support multi-tenant workloads from traditional cloud compute through AI within the same environment. However, their limitations tend to be seen in speed, flexibility, and the complexity of balancing internal resource allocation across many competing priorities. Additionally, their architecture reflects CPU-driven compute, with GPU-driven compute retrofitted into existing data centers, impacting their ability to deliver high density, efficient GPU compute clusters.

The rise of Neoclouds

Neoclouds sit at the opposite end of the spectrum from hyperscalers, with architectures purpose-built for GPU-driven AI workloads. Instead of general-purpose data centers, they focus on rapidly deploying GPU-dense clusters optimized for high-throughput, bare-metal performance. They win on speed, specialization, and hardware integration, though their footprints are smaller and their platforms less comprehensive than hyperscalers.

Neoclouds: bare metal compute today, full-stack AI cloud tomorrow

The Neocloud market is often described as a race to build full-stack AI cloud platforms, but today the majority of industry revenue remains tied to a much simpler offering: bare-metal compute. Most neoclouds today derive a majority of their revenue from large hyperscale customers that primarily utilize dedicated GPU infrastructure rather than a broad suite of cloud services. This dynamic highlights a key characteristic of the current AI infrastructure market: access to compute remains a highly valuable commodity. For many hyperscalers and AI labs, the primary objective is securing additional GPU capacity as quickly as possible, reducing the immediate need for higher-layer services, orchestration tools, or managed AI platforms. In this environment, the ability to deliver reliable access to GPUs has been sufficient to generate substantial demand and support rapid revenue growth.

Over the longer term, however, both CoreWeave and Nebius are working to evolve beyond pure infrastructure providers and establish themselves as full-stack AI clouds. This strategy includes offering managed services, orchestration software, storage, networking, inference tools, and other value-added capabilities that improve customer productivity and increase revenue per deployed GPU. While hyperscalers and frontier AI labs currently represent the largest source of demand, we believe enterprises ultimately represent the larger long-term opportunity. As enterprise AI adoption expands, customers are likely to place greater value on integrated platforms that simplify

deployment and management rather than procuring raw compute resources alone. As a result, the evolution from bare metal provider to full-stack AI cloud has become an important strategic objective across the Neocloud ecosystem, even as access to GPU capacity remains the primary driver of demand today.

Table 1: Hyperscalers and Neoclouds have differing data center and compute profiles
AI Compute Vendor Profile Comparison

	Hyperscaler	Neocloud
Data center model	Multi-tenant regions built for general cloud/CPU compute, retrofitted for GPU led compute	Single region GPU sites; some with distributed fabric
Hardware strategy	Standard OEM rack integrated into infrastructure	Custom bare-metal GPU racks
Latency	Low latency within site, region bound	Very low within region. Low latency across regions
Multi tenancy impact	High (workloads share fabric with general cloud)	Low (GPU-native workloads only)
Networking	Low latency RDMA for large clusters, no cross region clusters	High-speed single site fabrics
Elasticity	High elasticity within region	High elasticity within region; some high across entire physical footprint
Utilization efficiency	Lower due to multi-cloud tenancy	High with bare-metal scheduling
Strength	**Ultra large monolithic model training, highly customizable**	**Performance sensitive jobs, workloads that require control and tuning**
Ideal customer	**AI labs and large scale customers who need hundred of thousands of GPUs in one site**	**AI native companies, enterprises who need platform offering**
Biggest Risk	**Training becomes distributed across regions**	**Orchestration and GPU scheduling is open-sourced or brought in-house, Full stack AI platforms becomes more common**

Source: BofA Global Research

BofA GLOBAL RESEARCH

Pricing is largely contract-based, with some exposure to GPU spot prices

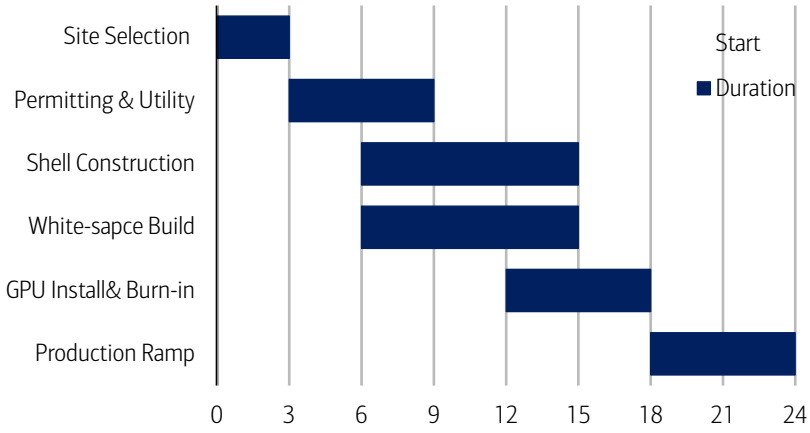
The majority of Neocloud revenue is generated through long-term contracts with large customers, including hyperscalers and AI labs. These agreements are typically structured as multi-year commitments, often spanning five years, and are generally priced on a fixed-rate basis tied to GPU-hour consumption. While contract structures vary and can include provisions for pricing adjustments, hardware refreshes, or capacity expansions, the primary benefit of these agreements is revenue visibility for providers and guaranteed access to scarce compute resources for customers.

A smaller portion of revenue remains exposed to spot-market GPU pricing, where customers purchase compute capacity on a shorter-term or more flexible basis. In the current environment, this exposure allows providers to capture some of the upside from tightening market conditions, while the predominance of long-term contracts continues to provide a stable foundation for revenue growth.

A focus on faster build times

Neoclouds have proved that AI optimized facilities can be brought online in 18-24 months. Neoclouds are optimizing for building on pre-energized land, allowing development to begin on top of existing electric substations. Second, they use repeatable standard high-density designs and modular building blocks that reduce technical complexity. Third, they aggressively overlap construction, utility upgrades, powered shell outfitting and commissioning, enabling the first data halls to go live while later data halls are still being completed.

Exhibit 10: AI/Neocloud Data Centers can take 18-24 months to build
Timeline of Neocloud Data Center Buildout



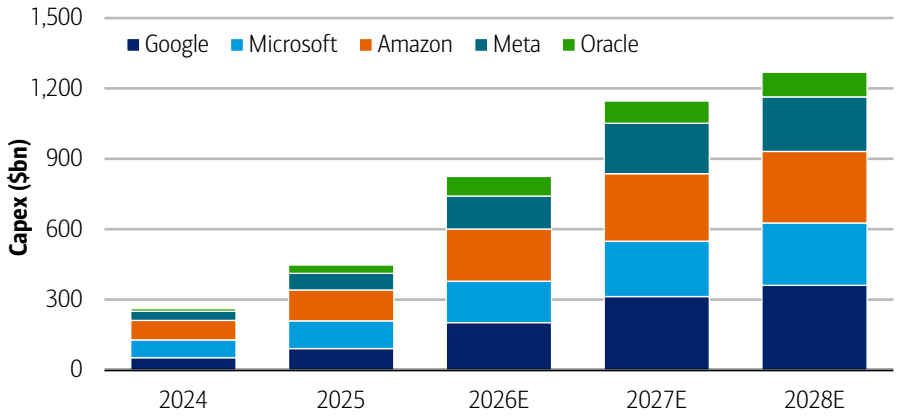
Source: BofA Global Research

BofA GLOBAL RESEARCH

- **Months 0-3: Site selection, land control and pre-work.** This timeline assumes pre-energized industrial land, where electrical substations or high voltage infrastructure already exist. Land control is finalized, civil work begins, and early utility scoping is initiated in parallel. Early civil activities such as soil testing, grading strategy, and road access layout can commence even before full permitting is complete.
- **Months 3-9: Permitting, utility engineering, civil work.** Permitting and utility engineering run concurrently with foundation preparation and early mechanical planning. Zoning and environmental approvals advance while utilities re-rate existing substation.
- **Months 6-15: Powered shell construction and white space build.** By month six, the core building structure is being built and mechanical and electrical infrastructure take shape. Neoclouds take a modular buildout approach, allowing blocks of the data center to enter white space build out (racks, containments, cooling distribution, cable and fiber runs) as soon as it is structurally ready, even if other blocks are still mid-construction. Utilities simultaneously complete upgrades and energize the data center in phases.
- **Months 12-18: First wave of IT equipment installation.** The first blocks of white space begin rack installations. Commissioning occurs in successive waves as different blocks reach white space readiness.
- **Months 18-24: Production ramp and full cluster stabilization.** The first data hall is energized and begins running workloads. Remaining white space blocks continue to be infilled.

Cloud service providers

Exhibit 11: Forecast US cloud service providers' capex to be above $1.2tr in 2028E
Capex figures include third-party leases

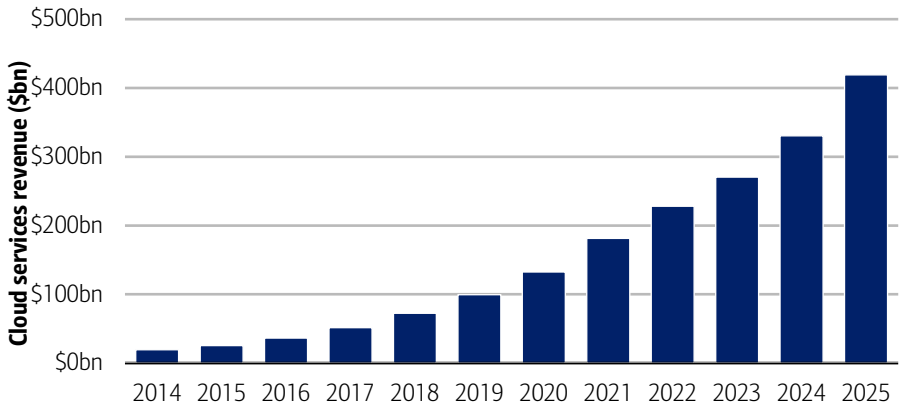


Source: Company reports, BofA Global Research estimates, Bloomberg

BofA GLOBAL RESEARCH

Cloud service providers are scaled, highly profitable firms. In 2025, Amazon Web Services generated $46bn of GAAP operating profit on $129bn of revenue, or a 35% operating margin. For most IT workloads, colocation provides a lower total cost of ownership. However, it requires upfront capex (for services & related IT equipment), multi-year commitments (to colocation firms), and higher level of IT management and support. In contrast, cloud services are flexible and offer higher uptime levels.

Exhibit 12: Cloud services revenue has grown from under $20bn in 2014 to $419bn in 2025
The three largest firms are Amazon Web Services, Microsoft Azure, and Google Cloud



Source: Synergy Research Group, company filings

BofA GLOBAL RESEARCH

Colocation economics

We walk through typical per megawatt project economics for a new build wholesale colocation project (i.e., one leased to a small number of clients on long-term leases).

- **$2.5mn per MW.** Land costs, utility connections, and site works. This would vary by location/site.
- **$11.5mn per MW for the powered shell.** Turner & Townsend's Data Center Cost Index uses real construction data from 300 projects. We also looked at JLL's 2026 construction cost estimates for Atlanta, Dallas, Northern Virginia, and Phoenix. These costs include the building shell, mechanical, electrical, thermal, equipment and installation labor costs, and general contractor margin and contingency.

Typical annual rent is $2-3mn per megawatt; we assume $2.5mn. Current occupancy rates in the US are 96-97%; we conservatively assume 90% over the 20 years. The largest operating expense is electricity ($0.09 per kilowatt hour is average US industrial cost implying $0.8mn for 100% utilization, or $0.7mn given 90% occupancy level). Staffing levels are around two full-time employees per megawatt; we assumed $0.22mn in wages. We assume property taxes of 1% on property value. Typical EBITDA margins are 40-50%; we assume 45%.

For free cash flow, we assume maintenance capex of 1.5% of original $11.5mn powered shell cost. We assume project financing of $7mn in equity (at 10% cost) and $7mn in 20-year amortizing mortgage debt (at 7% rate), yielding at weighted average cost of capital of 8.5%. This would be a 50% loan-to-value. We use 21% corporate tax rate for cash income taxes.

After the 20-year holding period, we assume the data center is sold. Recent data center transactions have been around 5-6% capitalization rates; we assume 6.5% for the exit at the end of year 20 (i.e., 15.4x net operating income). These assumptions yield a 10.7% internal rate of return (IRR) and a $2.0mn net present value (relative to the original $7mn equity investment).

Exhibit 13: Illustrative economics for wholesale colocation project
We assume 90% occupancy, $2.25mn of annual initial revenue, 3% inflation, and 45% EBITDA margin

$mn	Yr 1	Yr 2	Yr 3	Yr 4	Yr 5	Yr 6	Yr 7	Yr 8	Yr 9	Yr 10	Yr 11	Yr 12	Yr 13	Yr 14	Yr 15	Yr 16	Yr 17	Yr 18	Yr 19	Yr 20
Rent/MW	2.50	2.58	2.65	2.73	2.81	2.90	2.99	3.07	3.17	3.26	3.36	3.46	3.56	3.67	3.78	3.89	4.01	4.13	4.26	4.38
Occupancy	90%	90%	90%	90%	90%	90%	90%	90%	90%	90%	90%	90%	90%	90%	90%	90%	90%	90%	90%	90%
Revenue	2.25	2.32	2.39	2.46	2.53	2.61	2.69	2.77	2.85	2.94	3.02	3.11	3.21	3.30	3.40	3.51	3.61	3.72	3.83	3.95
Operating expenses	-1.24	-1.27	-1.31	-1.35	-1.39	-1.43	-1.48	-1.52	-1.57	-1.61	-1.66	-1.71	-1.76	-1.82	-1.87	-1.93	-1.99	-2.05	-2.11	-2.17
EBITDA	1.01	1.04	1.07	1.11	1.14	1.17	1.21	1.25	1.28	1.32	1.36	1.40	1.44	1.49	1.53	1.58	1.62	1.67	1.72	1.78
EBITDA margin	45%	45%	45%	45%	45%	45%	45%	45%	45%	45%	45%	45%	45%	45%	45%	45%	45%	45%	45%	45%
Maintenance capex	-0.17	-0.17	-0.18	-0.18	-0.19	-0.20	-0.20	-0.21	-0.21	-0.22	-0.23	-0.23	-0.24	-0.25	-0.26	-0.26	-0.27	-0.28	-0.29	-0.30
Interest expense	-0.42	-0.41	-0.40	-0.39	-0.37	-0.36	-0.35	-0.33	-0.31	-0.30	-0.28	-0.26	-0.24	-0.21	-0.19	-0.16	-0.13	-0.10	-0.07	-0.04
Cash income taxes	-0.05	-0.06	-0.06	-0.07	-0.08	-0.09	-0.10	-0.11	-0.13	-0.14	-0.15	-0.16	-0.18	-0.19	-0.20	-0.22	-0.24	-0.25	-0.27	-0.29
Free cash flow	**0.38**	**0.40**	**0.43**	**0.46**	**0.49**	**0.52**	**0.56**	**0.59**	**0.63**	**0.67**	**0.71**	**0.75**	**0.79**	**0.84**	**0.88**	**0.93**	**0.98**	**1.04**	**1.09**	**1.15**
Debt repayment	-0.15	-0.16	-0.17	-0.18	-0.19	-0.21	-0.22	-0.24	-0.25	-0.27	-0.29	-0.31	-0.33	-0.35	-0.38	-0.40	-0.43	-0.46	-0.49	-0.53
Depreciation	0.37	0.37	0.37	0.37	0.37	0.37	0.37	0.37	0.37	0.37	0.37	0.37	0.37	0.37	0.37	0.37	0.37	0.37	0.37	0.37
Revenue	2.25	2.32	2.39	2.46	2.53	2.61	2.69	2.77	2.85	2.94	3.02	3.11	3.21	3.30	3.40	3.51	3.61	3.72	3.83	3.95
Cash operating expenses	-1.41	-1.45	-1.49	-1.54	-1.58	-1.63	-1.68	-1.73	-1.78	-1.83	-1.89	-1.95	-2.00	-2.07	-2.13	-2.19	-2.26	-2.32	-2.39	-2.47
Net operating income (NOI)	**0.84**	**0.87**	**0.90**	**0.92**	**0.95**	**0.98**	**1.01**	**1.04**	**1.07**	**1.10**	**1.13**	**1.17**	**1.20**	**1.24**	**1.28**	**1.31**	**1.35**	**1.39**	**1.44**	**1.48**
EBITDA	1.01	1.04	1.07	1.11	1.14	1.17	1.21	1.25	1.28	1.32	1.36	1.40	1.44	1.49	1.53	1.58	1.62	1.67	1.72	1.78
Maintenance capex	-0.17	-0.17	-0.18	-0.18	-0.19	-0.20	-0.20	-0.21	-0.21	-0.22	-0.23	-0.23	-0.24	-0.25	-0.26	-0.26	-0.27	-0.28	-0.29	-0.30
Interest expense	-0.42	-0.41	-0.40	-0.39	-0.37	-0.36	-0.35	-0.33	-0.31	-0.30	-0.28	-0.26	-0.24	-0.21	-0.19	-0.16	-0.13	-0.10	-0.07	-0.04
Cash income taxes	-0.05	-0.06	-0.06	-0.07	-0.08	-0.09	-0.10	-0.11	-0.13	-0.14	-0.15	-0.16	-0.18	-0.19	-0.20	-0.22	-0.24	-0.25	-0.27	-0.29
Free cash flow	**0.38**	**0.40**	**0.43**	**0.46**	**0.49**	**0.52**	**0.56**	**0.59**	**0.63**	**0.67**	**0.71**	**0.75**	**0.79**	**0.84**	**0.88**	**0.93**	**0.98**	**1.04**	**1.09**	**1.15**
Electricity	-0.63	-0.65	-0.67	-0.69	-0.71	-0.73	-0.75	-0.78	-0.80	-0.82	-0.85	-0.87	-0.90	-0.93	-0.95	-0.98	-1.01	-1.04	-1.07	-1.11
Staffing	-0.25	-0.25	-0.26	-0.27	-0.28	-0.29	-0.30	-0.30	-0.31	-0.32	-0.33	-0.34	-0.35	-0.36	-0.37	-0.39	-0.40	-0.41	-0.42	-0.43
Real estate taxes	-0.13	-0.13	-0.14	-0.14	-0.15	-0.15	-0.16	-0.16	-0.16	-0.17	-0.17	-0.18	-0.19	-0.19	-0.20	-0.20	-0.21	-0.21	-0.22	-0.23
SG&A costs	-0.23	-0.24	-0.24	-0.25	-0.26	-0.27	-0.27	-0.28	-0.29	-0.30	-0.31	-0.32	-0.33	-0.34	-0.35	-0.36	-0.37	-0.38	-0.39	-0.40
Operating expenses	**-1.24**	**-1.27**	**-1.31**	**-1.35**	**-1.39**	**-1.43**	**-1.48**	**-1.52**	**-1.57**	**-1.61**	**-1.66**	**-1.71**	**-1.76**	**-1.82**	**-1.87**	**-1.93**	**-1.99**	**-2.05**	**-2.11**	**-2.17**

Source: BofA Global Research

BofA GLOBAL RESEARCH

AI: training versus inference

- **AI training:** an iterative process of teaching a code base to identify relationships from a data set so it can make accurate predictions. Training involves iteration over huge datasets with many GPU nodes in parallel with a single large model training run consuming hundreds of MWh of energy. Furthermore, training often pushes hardware to peak utilization with GPUs at 90-100% utilization for days/weeks. However, once the model is trained, that particular task is completed (though models are frequently retrained or fine-tuned), with frequency depending on R&D needs.

- **AI inference:** the use of a trained AI model to make predictions based on new data. Inference power requirements tend to be more distributed and continuous with each inference uses relatively small energy. However, the aggregate energy on millions of inference queries can exceed training over time.

Training power draw continues to accelerate

Training power draw continues to rise despite advances in training techniques (e.g., reinforcement learning). Power draw is defined as the electricity used by a data center for servers simultaneously working on an individual model.

Inference electricity demand to grow with usage

AI reasoning models use a "chain of thought" methodology that breaks the task into parts. This allows the model to do better on tasks related to logic, math, and pattern finding. However, it requires greater computational intensity during inference. According to an MIT Technology Review article published in 2025, the DeepSeek reasoning model used 87% more energy than Meta's Llama 3, a model with the same number of parameters, for a similar inquiry as it generates much longer responses. This is particularly relevant as compute is expected to shift more towards inference versus training over time. (O'Donnell, James. "DeepSeek might not be such good news for energy after all.")

Images & videos are even more energy intensive

Different inferencing tasks require different levels of computing power. An MIT Technology Review study suggests that video generation is by far the most intensive by a factor of ~1000x (versus a large text model).

MIT Technology Review measured energy used for inference for text generation using Meta's Llama models which are open source. The smallest model in the cohort, Llama 3.1 8B, had 8 billion parameters (adjustable inputs in the model). The largest model, Llama 3.1 405B, had 405 billion parameters. More parameters drive more energy per response. Prompts themselves also can drive higher energy needs, with more complicated requests driving more energy usage.

Image generation models require a different architecture called diffusion, which transformers noise/words into images. The energy requirement is not dependent on the prompt, but rather by the size of the model, image resolution, and the number of steps the process takes which can generate higher quality.

MIT used a Chinese AI startup video model called CogVideoX to estimate video generation energy. This went from 109,000 joules (in the August 2024 model) for eight frames per second to 3.4mn joules in November 2024 for 16 frames per second for a five-second video.

Exhibit 14: Inference tasks ranked by energy usage, 2025
Video Generation requires 1000x more energy than large scale text models

Inference	Joules per GPU	Total joules	W per hour
Text Model - Small	57	114	0.016
Image Generation - Small	1,141	2,282	0.317
Image Generation - Large	2,282	4,564	0.634
Text Model - Large	3,353	6,706	0.931
Video Generation	3,400,000	6,800,000	944.445

Source: MIT Technology Review, BofA Global Research

BofA GLOBAL RESEARCH

Extreme codesign driving M&A

Extreme codesign extends to data center infrastructure

Nvidia CEO Jensen Huang's concept of "extreme codesign" reflects a shift from optimizing individual components to designing the AI data center as a single, tightly integrated system. As AI clusters scale, performance increasingly depends not only on the GPU, but on how compute, networking, software, electrical, and thermal equipment work together. Extreme codesign therefore pushes optimization beyond the semiconductor itself and across the broader physical infrastructure supporting it. In our view, this raises the strategic importance of infrastructure suppliers that can engage early in the design process, develop around Nvidia's evolving architectures, and deliver these integrated solutions.

Some suppliers are being invited into early R&D

We believe Nvidia wants suppliers to embed themselves in the semiconductor design process to deliver comprehensive solutions based on the entire rack architecture. They must also scale quickly to support large AI deployments. This is not a theoretical approach. We expect the Vera Rubin platform to tightly link Nvidia's rack architecture with selected electrical & thermal equipment vendors.

Accelerated chip design cycle drives need for systems

Nvidia accelerated its chip development cycle from once every two years to every year, or a CEO Huang said, "We're on a one-year rhythm." This change makes physical infrastructure obsolete more quickly. Data center power and cooling systems now must evolve as fast as the chips they support. For example, Nvidia's roadmap includes a new rack architecture after every two chip releases. For infrastructure equipment vendors, this means that even if you have a great product for the current rack architecture, you could face challenges in just two years. We view this as another push towards integrated solutions that are closely aligned with semiconductor vendors' architecture.

Blurring of white space, gray space, and periphery

Extreme codesign drives deeper integration across all of the data center infrastructure. Vendors need to be somewhat agnostic to specific point solutions and willing to invest in R&D, product development, and technology-focused acquisitions. In the past, there was a clear divide between vendors selling into the white space of a data center (i.e., the data hall housing the racks and servers) and gray space (i.e., power & cooling equipment). Extreme codesign merges these areas. The roadmap also adds mission critical equipment to the periphery of the data center. The periphery could now include battery Energy Storage Systems (BESS), packaged air-cooled chillers, and on-site primary power generation.

Direct current transition accelerates these trends

We think the transition to 800-volt direct current power distribution will show the consequences of this extreme codesign approach. Electrical vendors that are part of the ecosystem and aligned with Nvidia's reference architecture are likely to gain share through new products. These include solid state transformers, power sidecars, BESS, and liquid-cooled busbar. Product-focused companies lacking complete systems offerings are likely to lose share. We see a similar approach by hyperscalers to their data center developers: successful developers get more projects. Hyperscalers want more sites from proven developers versus spreading their spending to more and more developers

Integration of power and cooling

Extreme codesign favors suppliers that offer both power and cooling capabilities. We think this is a meaningful part of the rationale behind Schneider's Motivair acquisition (Feb. 2025) and Eaton's acquisition of Boyd Corp. (Mar. 2026).

Data center infrastructure M&A review

Exhibit 15: Selected acquisitions of data center infrastructure equipment firms since 2024
Ten acquisitions were ~$1bn or more

Company	Target	Announced Date	Close Date	Transaction value (mns)	EV/Sales	EV/EBITDA	Forward revenue (mns)	Forward EBITDA (mns)	EBITDA % margin	Key Products
Eaton	COL Group	9/25/2026	1Q CY27	€ 810	3.2x		€ 250			Medium voltage distribution, switchgear, and modular power solutions
Flex	EPC Power	9/3/2026	10/4/2026	$4,400	3.9x	13.1x	$1,120	$336	30%	Rectifiers, DC-DC conversion, solid-state transformers in development
Vertiv	UtilityInnovation Group	9/2/2026	4Q CY26	$1,450		13.0x		$112		Microgrid controls, onsite generation orchestration, microgrid switchgear
nVent	Maverick Power	8/24/2026	4Q CY26	$1,750	2.5x	11.5x	$700	$152	22%	Low- and medium-voltage switchgear and switchboards
Madison Air	ebm-papst	8/17/2026	4Q CY26	$5,021	1.8x	14.6x	$2,772	$343	12%	Electronically commutated ("EC") fan and motor systems
ABB	Advantics	7/15/2026	TBD							Silicon Carbide power conversion (DC)
Hubbell	NSI Industries	5/4/2026	6/9/2026	$3,000	5.3x	15.5x	$570	$194	34%	Electrical fittings, connectors, and components
Vertiv	BMarko Structures	4/13/2026	4/13/2026							Custom-engineered structural fabrication
Vertiv	ThermoKey	3/23/2026	6/12/2026							Dry coolers and microchannel-based heat-exchange solutions
Ecolab	CoolIT Systems	3/20/2026	7/2/2026	$4,750	8.6x	29.0x	$550	$164	30%	Coolant distribution units (CDUs), cold plates, and direct-to-chip cooling
Trane Technologies	LiquidStack	2/10/2026	3/3/2026							Direct-to-chip and immersion cooling, coolant distribution units (CDUs)
Trane Technologies	Stellar Energy Digital Business	12/2/2025	2/18/2026							Modular cooling plants, central utility plants and coolant distribution units
Legrand	Avtron Power Solutions	11/4/2025	11/4/2025				$350			Load banks and power quality solutions
Daikin	Chilldyne	11/4/2025	11/4/2025							Cooling distribution unit (CDU)
Eaton	Boyd Thermal	11/3/2025	3/12/2026	$9,500	5.6x	22.5x	$1,700	$422	25%	Coolant distribution units (CDUs), cold plates, direct-to-chip cooling, chillers
Vertiv	PurgeRite	11/3/2025	12/4/2025	$1,000		10.0x		$100		Mechanical flushing, purging and filtration services for data centers
Hubbell	DMC Power	8/12/2025	10/1/2025	$825	6.3x	13.8x	$130	$60	46%	High-voltage transmission and substation swage connectors
Daikin	DDC Solutions	8/5/2025	8/18/2025							Cooling cabinets and management software for data centers
Vertiv	Great Lakes Data Racks & Cabinets	7/17/2025	8/20/2025	$200						Data racks, cabinets, advanced cooling systems, enclosures
Eaton	Resilient Power Systems	7/16/2025	8/6/2025	$86						Solid state transformer technology
Eaton	Fibrebond	3/11/2025	4/1/2025	$1,400	3.7x	12.7x	$378	$110	29%	Thermal monitoring solutions for electrical equipment
nVent	Electrical Products Group Business of Avail Infrastructure Solutions	3/10/2025	5/1/2025	$975	2.6x	12.5x	$375	$78	21%	Enclosures, switchgear and bus systems
ABB	Power Electronics business of Gamesa Electric	12/18/2024	12/2/2025							Generators, industrial battery energy storage systems (BESS), and inverters
Schneider Electric	Motivair	10/17/2024	2/28/2025	$850	~4-6x forward					Coolant Distribution Units (CDUs), Cold Plates and Chillers
Munters	Geoclima	7/17/2024	9/30/2024				SEK 455			Air- and water-cooled chillers
Legrand	VASS	6/11/2024	6/11/2024				€ 10			Busbars
nVent	Trachte	6/6/2024	7/16/2024	$695	2.8x	12.0x	$250	$58	23%	Prefabricated structures
Eaton	NordicEPOD	5/31/2024	5/31/2024							Design and assembly of standardized power modules
Eaton	Exertherm	5/21/2024	5/21/2024							Thermal monitoring solutions for electrical equipment
Legrand	Netrack	4/23/2024	4/23/2024				€ 10			Data center racks
Modine	Scott Springfield Manufacturing	2/26/2024	3/13/2024	$190	1.9x		$100			Air handling units (AHU)
Modine	TMG Core	1/2/2024	1/2/2024							Liquid immersion cooling

Source: Company filings, Bloomberg, BofA Global Research

BofA GLOBAL RESEARCH

Making sense of the modular landscape

What modular is (and isn't) in data center construction

We expect modular and prefabricated products to capture a growing share of data center infrastructure spending. Modular and prefabricated products can lessen skilled labor constraints, shorten on-site construction schedules, and perform more testing off-site. Products range from basic electrical skids to eHouses, prefabricated systems installed within the data hall, and fully integrated modular data centers. For investors, we believe the relevant opportunity is therefore not simply the gross revenue associated with modular products, but the incremental premium vendors can capture for engineering, integration, assembly, and testing.

The terminology around modular or prefabricated equipment is not well defined. We see four separate categories as described below.

Exhibit 16: Comparing types of modular equipment
Modular can mean equipment located inside & outside the data center

Feature	All-in-one modular data center	eHouse/ePod	Prefabricated Systems	Power Skid
Integration level	Full data center infrastructure	Full electrical system	Multi-discipline: power, cooling, structure	Component-level integration
Enclosure	Large multi-module building	Fully enclosed, weatherproof	Factory-built modules or containers	Open frame or skid base
Scope	Complete power and thermal infrastructure	Electrical room	Power, cooling and/or white-space modules	Single function (e.g., UPS, switchgear)
Testing	Full system-level factory acceptance testing	Partial system testing	Module factory acceptance testing plus integrated site commissioning	Component-level testing

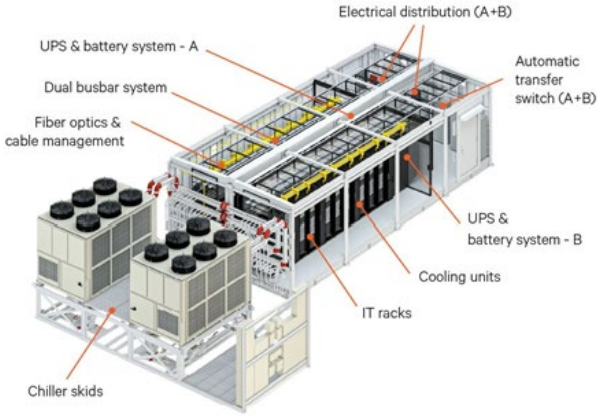
Source: BofA Global Research

BofA GLOBAL RESEARCH

1. All-in-one modular data centers

These historically have been niche product offerings usually used for defense applications, rugged environments (e.g., oil and gas), or by telecom firms in remote areas. Historically, these modules were small deployments of 0.3 MW or less per shipping container-sized enclosures. More recently, these have evolved to contain higher rack level densities.

Exhibit 17: Vertiv's prefabricated containerized data center solution
Vertiv SmartMod Max CW



Source: Vertiv

BofA GLOBAL RESEARCH

Prefabricated solutions are typically built into standard 20-foot or 40-foot shipping containers or custom built enclosures. Built with all the necessary racks, power systems, cooling equipment, and fire and safety systems. One example would be Armada's Leviathan, which features 1.8 MW IT load with separate modules for power distribution and cooling.

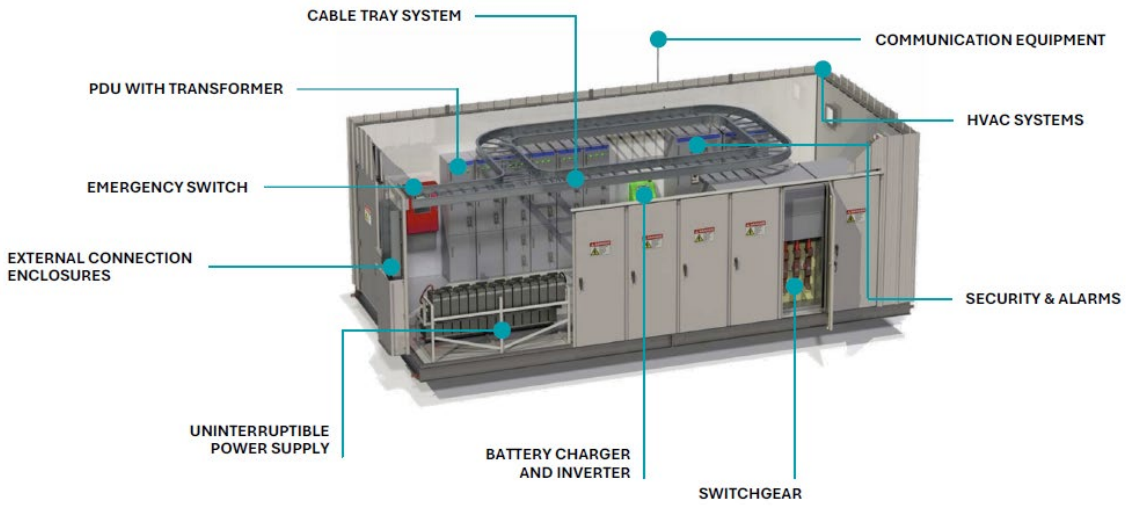
2. eHouse/ePod

An electrical house (or eHouse or ePod) is a prefabricated, modular, climate-controlled enclosure that houses electrical distribution equipment. This can include medium- and low-voltage switchgear, and UPS. Importantly, the eHouse sits outside of the data center and must include its own separate thermal equipment (given the heat generated by the electrical equipment inside).

This market has historically been led by ABB, Schneider Electric, Eaton, Siemens (which has delivered over 1,000 eHouses worldwide), and Powell Industries (in the mining and oil & gas sectors).

However, demand has increased significantly as a result of labor bottlenecks. In October 2025, Regal Rexnord received $0.7bn orders for E-pod offering which features power distribution, switchgear, cooling, and UPS. ABB is partnering with VoltaGrid – a fast-growing provider of on-site power for data centers – to provide eHouses for power control.

Exhibit 18: Illustrative eHouse with electrical and cooling equipment
The enclosure is for deployment outside of the data hall



Source: Forgent Power Solutions

BofA GLOBAL RESEARCH

3. Prefabricated systems

Prefabricated systems are located inside the data hall (i.e., "white space"). While product designs vary by company, these generally include hot aisle/cold aisle containment, overhead network cable trays, and electrical busway/whips. For structures including liquid cooling, they would also include manifolds, piping, and integrated sensors. They are not enclosed and require an existing data hall. However, they combine several field-installed trades into a factory-built, factory-tested system.

Exhibit 19: Illustrative prefabricated system
Vertiv's SmartRun combines power distribution, hot aisle containment, liquid cooling, and cable management into a single, integrated product



Source: Vertiv

BofA GLOBAL RESEARCH

4. Skidded kit

Combines multiple electrical components into a single deliverable skid. Skid is then delivered to inside of exiting data center building for rapid deployment and commissioning. Benefits include testing the skidded equipment at the factory level, reducing wiring time, and simplifying commissioning.

Exhibit 20: Skidded electrical equipment
Steel skid and side bracing provide support



Source: Eaton

BofA GLOBAL RESEARCH

How big is the modular market? It depends...

Looking at the total revenue or orders tied to prefabricated products overstates the total market size. We focus on the premium that the modular offering generates over the constituent parts. This premium captures the engineering and customization work, enclosure/skid, internal cabling/piping, assembly labor, and factory testing. This varies from skidded equipment on the low end to modular all-in-one data centers at the high end.

Below we show a top-down estimate for modular data center product revenue in 2026E. We start with the total product $66bn market size for electrical and thermal equipment. We assume 45-55% of this is addressable content for modular products. Key components, like chillers and medium voltage transformers, are not part of prefabricated offerings. We assume an adoption rate of between 5-9% and modular premiums of 35-55%. This yields a gross modular revenue of $2.0-5.1bn, or 3-8% of total electrical & thermal equipment spending. However, this overstates the true value-add of modular products. This is between $0.5-1.8bn, or 1-3% of total electrical & thermal equipment spend.

Exhibit 21: Estimated market size for modular products, 2026E
Modular products are likely the fastest-growing part of data center equipment market

($bn)	Low		Medium		High
Electrical & thermal data center products	66.1	66.1	66.1	66.1	66.1
Addressable content	45.0%	47.5%	50.0%	52.5%	55.0%
Modular addressable market	29.8	31.4	33.1	34.7	36.4
Modular adoption rate	5.0%	6.0%	7.0%	8.0%	9.0%
Component-level revenue	1.5	1.9	2.3	2.8	3.3
Modular product premium	35%	40%	45%	50%	55%
Gross modular product revenue	**2.0**	**2.6**	**3.4**	**4.2**	**5.1**
As % of total electrical data center products	3.0%	4.0%	5.1%	6.3%	7.7%
Modular value-add	**0.5**	**0.8**	**1.0**	**1.4**	**1.8**
As % of total electrical data center products	0.8%	1.1%	1.6%	2.1%	2.7%

Source: BofA Global Research

BofA GLOBAL RESEARCH

ELECTRICAL EQUIPMENT

An evolutionary path to direct current

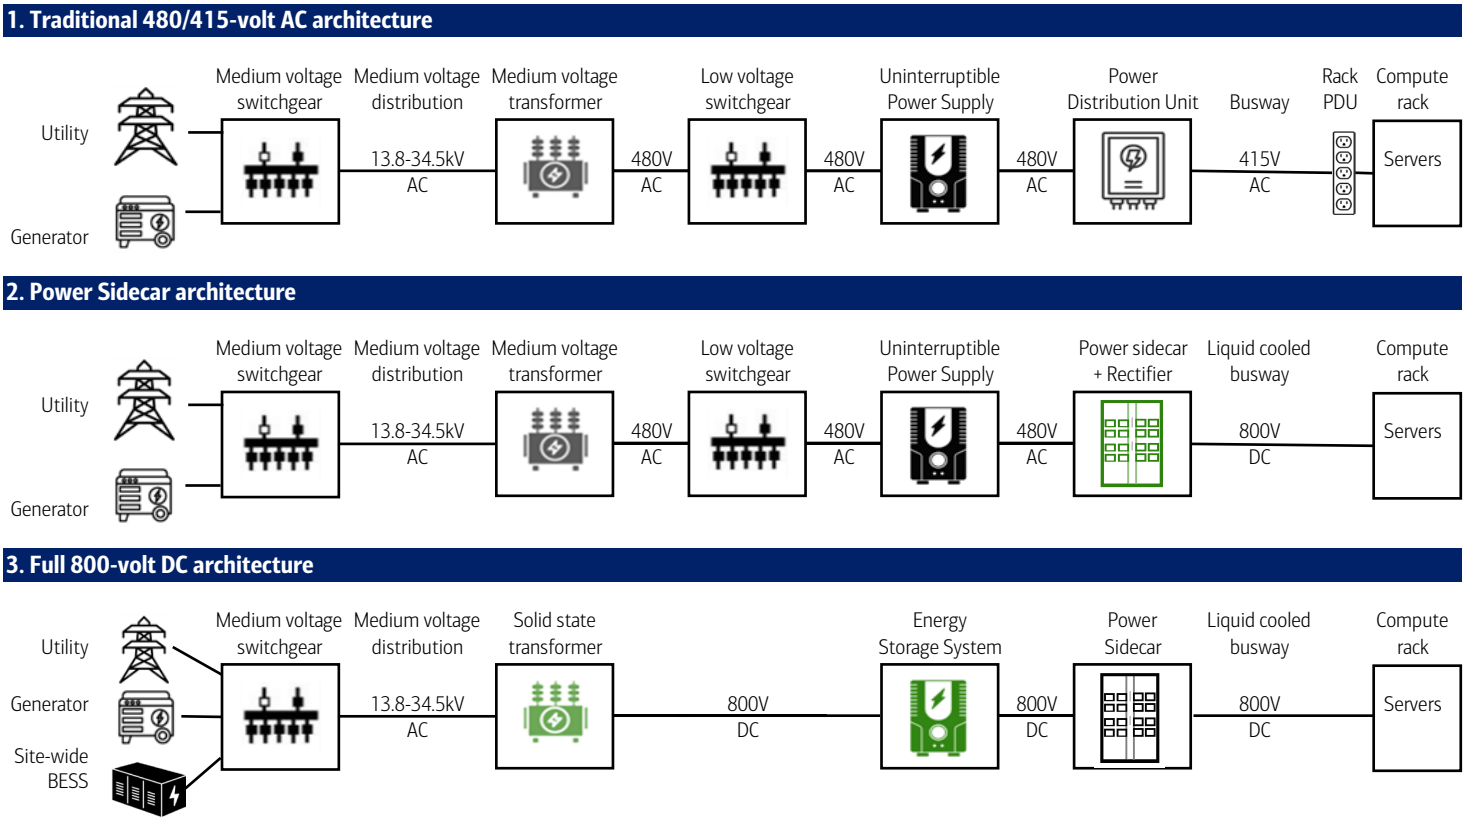
Why shift to higher voltage direct current?

Data centers are moving from alternating current (AC) architectures to direct current (DC) because rack power density (watts/server rack) is increasing. Nvidia's Rubin Ultra Kyber rack is expected to draw over 1,000 kilowatts (kW). The Open Compute Project's (OCP; consortium of data center owners, operators, and equipment providers) Mount Diablo rack supports up to 1,100kW per rack. Both Nvidia's 800-volt and OCP's 400-volt direct current architectures will require sizeable changes to electrical equipment.

The shift to higher voltage direct current reduces conductor requirements, electrical losses due to resistance, and simplifies overall power delivery. However, developers are likely to proceed cautiously given a new equipment ecosystem, need for technician training, and lack of broad industry standards.

> "The transition to a fully realized 800 VDC architecture will occur in phases, giving the industry time to adapt and the component ecosystem to mature." — *Nvidia Technical Blog "Building the 800 VDC Ecosystem for Efficient, Scalable AI Factories"*

Exhibit 22: How data center electrical architecture changes in evolution to direct current
Power sidecar and solid state transformer are key additions; low voltage and rack power distribution units (rPDUs) are likely to be eliminated



Source: BofA Global Research, Nvidia Technical Blog "Building the 800 VDC Ecosystem for Efficient, Scalable AI Factories". Note: AC: alternating current, BESS: Battery Energy Storage System, DC: direct current; kV: kilovolts, PDU: power distribution unit, V: volts

BofA GLOBAL RESEARCH

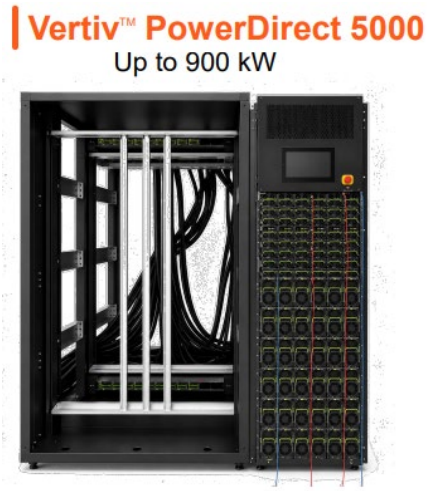
1. Traditional 480/415-volt AC

Medium-voltage utility power (e.g., 13,800-34,500-volt AC) is fed into medium-voltage switchgear. Transformers step this down to 480-volt AC for further distribution. Low-voltage switchgear (including both main switchboard and additional switchgear) routes the electricity to various circuits. This passes through double-line conversion uninterruptible power supplies (UPS), which typically provide ~15-30 minutes of ride-through full-load power. The conditioned power then goes through a power distribution unit, a remote power panel, or directly through busway. Servers inside compute racks connect to the power through rack power distribution units (rPDUs) typically mounted on the outside the rack. Power supply units (PSUs) convert AC electricity to DC inside the compute rack.

2. Adding a power sidecar

There is no change to medium-voltage equipment or UPS. After the UPS, a rectifier converts the conditioned 480-volt AC to 800-volt DC. Some configurations include the rectifier inside the power sidecar, while others put it outside the sidecar. A liquid-cooled busway delivers electricity to the compute rack. The power sidecar also includes new rack-level energy storage equipment. Inside the compute rack, there is still an additional DC-DC conversion (from 800-volt to 12-volt).

Exhibit 23: Sidecars to support higher wattage racks in existing AC infrastructure
Vertiv's upcoming PowerDirect 5000 sidecar solution to come in early 2027



Source: Vertiv

BofA GLOBAL RESEARCH

3. Adding a solid state transformer and energy storage

There is no change to medium-voltage switchgear and distribution. A solid state transformer replaces the traditional transformer, which converts the AC medium-voltage feed to 800-volt DC. With fewer components, the UPS is greatly simplified, and becomes more of an energy storage system. However, this would still need bypass switches, power controls, and communication. The power sidecar now receives 800-volt DC power, eliminating the need for a rectifier.

Tracking the content per megawatt

Exhibit 24: Total content per megawatt is similar, but with more MV equipment and less LV equipment
Key additions include solid-state transformers, power sidecar components, and site-wide BESS

1. Traditional 480/415-volt AC		2. Power Sidecar		3. Full 800-volt DC	
Equipment	**Cost/MW**	**Equipment**	**Cost/MW**	**Equipment**	**Cost/MW**
Medium voltage switchgear	70,000	Medium voltage switchgear	70,000	Medium voltage switchgear	70,000
Medium voltage distribution	75,000	Medium voltage distribution	75,000	Medium voltage distribution	75,000
Medium voltage transformer	140,000	Medium voltage transformer	140,000	~~Medium voltage transformer~~	0
				Solid-state transformer	840,000
Low-voltage switchboard	390,000	Low-voltage switchboard	390,000	~~Low-voltage switchboard~~	0
UPS bypass switchboard	150,000	UPS bypass switchboard	150,000	~~UPS bypass switchboard~~	0
Busway & tap boxes	420,000	Busway & tap boxes	420,000	Busway & tap boxes	420,000
Power distribution units (PDUs)	235,000	~~Power distribution units (PDUs)~~	0	~~Power distribution units (PDUs)~~	0
Uninterruptible Power Supplies (UPS)	1,060,000	Uninterruptible Power Supplies (UPS)	1,060,000	~~Uninterruptible Power Supplies (UPS)~~	0
				Row-level energy storage	318,000
				Site-wide BESS (1 MWh)	425,000
No power sidecar		Power sidecar		Power sidecar	
		Liquid-cooled busway (within rack)	352,500	Liquid-cooled busway (within rack)	352,500
		Battery backup unit(s)	40,000	Battery backup unit(s)	40,000
		Supercapacitor unit	120,000	Supercapacitor unit	120,000
Total	2,540,000	Total	2,817,500	Total	2,660,500

Source: BofA Global Research

BofA GLOBAL RESEARCH

Why our content/MW is likely conservative

- **Custom versus standard:** Our pricing estimates are based on similar products, but not those certified by semiconductor manufacturers. Pricing from Nvidia partners or OCP Accepted (TM) products, for example, may receive a meaningful pricing premium.

- **Devil may be in the (component) details:** Subcomponent pricing may result in substantially higher overall system costs. To give one example, traditional AC hot swap controllers rated for 90 volts cost $2-3 each. However, an 800-volt hot swap controller will likely require a silicon carbide (SiC) solution. Existing (i.e., non-data center specific) SiC controllers cost over $20 each, of 5-10x the comparable lower voltage component. Another example would be changing from traditional mechanical circuit breakers (acts in milliseconds) to solid-state circuit breakers (acts in microseconds (e.g., 100x faster). Directionally, these would be 5-10x higher costs.

- **History would suggest a costly initial rollout:** This is not the first time that companies have tried DC architectures. In 2011, ABB acquired a majority stake in DC power distribution firm Validus DC. They built several DC-based data centers globally. The absence of standards and equipment meant that the initial cost to deploy was much higher.

New equipment in 400- or 800-volt direct current

- **Solid state transformers (SSTs):** In applications where size/weight matter, engineers have long used power electronics, such as solid-state transformers (e.g., many electric vehicle rely on silicon carbide (SiC) power inverters). Traditional transformers are passive devices. In contrast, SSTs can actively regulate current and convert from AC to DC. The trade-off is a significantly higher cost on a per-megawatt basis.

Exhibit 25: SSTs enable direct current infrastructure throughout the hall
Eaton's immersion cooled MV SST



Source: Eaton

BofA GLOBAL RESEARCH

Given the early stage of the solid-state transformer market for data centers, we provide an overview of various vendors and those who have announced product development. We find four companies have already shipped solid-state transformers for the data center market, two that have product demonstrations, and three others that are undergoing product validation by customers.

Exhibit 26: Overview of solid state manufacturers
Focus on megawatt-scale and up for data center applications

Company	Status	Notes
Delta Electronics	Shipping product	Deployed at a Chindata data center in July 2026
ABB / Mitsubishi Heavy / DG Matrix	Shipping product	$60mn Series A in Feb. 2026
Eaton / Resilient Power Systems	Shipping product	acquired RPS for $86mn in August 2025
SIFANG Electric	Shipping product	2.4 megawatt launched in April 2026
Hyosung Heavy	Product demonstrated	22.9 kilovolt version demo in April 2026
Amperesand	Product demonstrated	Commercial units to ship by YE26; $80mn Series A funding led in Nov. 2025
GE Vernova	Product validation	Expected to start taking orders in 2027

Exhibit 26: Overview of solid state manufacturers
Focus on megawatt-scale and up for data center applications

Company	Status	Notes
Jinpan Technology	Product validation	Prototypes delivered to clients
Heron Power	Product validation	$140mn Series B funding led by a16z in February 2026
Enphase Energy	Product development	Customer pilots expected in 2027; production started in Sept. 2026
SolarEdge	Product development	Partnership with Infineon announced in 2025
Flex / EPC Power	Product development	Acquired by Flex in October for $4.4bn
Siemens Energy	Product development	Prototype ready in 2027
Kehua	Product development	Under development

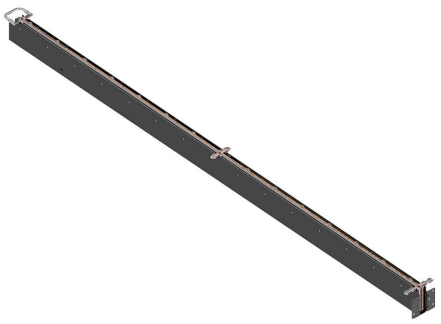
Source: BofA Global Research, company filings, company presentations
Note: Status as of 3Q26

BofA GLOBAL RESEARCH

As an alternative to solid-state transformers, data centers can combine traditional transformers (to step up to higher voltage) with a rectifier (to convert to direct current). These transformers rectifier units (TRUs) have more established supply chains, but require additional integration, protection, and footprint versus SSTs.

- **Liquid-cooled busway:** Allows for 5x more electrical current versus air-cooled busway. The busbar "free rides" on the existing direct-to-chip liquid cooling architecture, with fittings at the bottom (chilled entry) and top (warm exit) of each full-rack height busbar. Servers connect to the busbar through blind-mate connectors (versus traditional plugs that go into rack PDUs).

Exhibit 27: Liquid cooled busbars enable greater electrical currents needed for DC architecture
Molex's liquid cooled busbar



Source: Molex

BofA GLOBAL RESEARCH

- **Row-level energy storage system:** The energy storage role of the UPS in the traditional AC architecture shifts to two different areas. With no need for AC-DC and DC-AC conversions (i.e., double conversion), the UPS can evolve to a focused row-level energy storage.

Exhibit 28: MV UPS work as power buffering and power conditioning systems
ABB HiPerGuard MV UPS



Source: ABB

BofA GLOBAL RESEARCH

- **Site-wide Battery Energy Storage System (BESS):** BESS can serve to smooth out load fluctuations and reduce the strain on the broader grid or on-site power generation. BESS can also fulfill peak shaving requirements. We view this as fully incremental to existing data centers, few of which have site-wide energy storage.

Exhibit 29: Increasing energy needs and complexity of data centers includes BESS as standard design
Vertiv EnergyCore Grid - Battery Energy Storage System



Source: Vertiv

BofA GLOBAL RESEARCH

- **Supercapacitor unit:** These provide high-power, short-duration energy storage. The “supercap shelf” provides the first line of defense against micro-seconds to seconds-long energy “spikes” by IT equipment.

Exhibit 30: Sharp spikes in inference loads require rapid power delivery provided by supercapacitors
XTM-18 Supercapacitor module



Source: Eaton

BofA GLOBAL RESEARCH

- **Battery Backup Units (BBUs):** Likely with lithium-ion for compact deployment, these would provide 60-120 seconds of backup power. This is consistent with Nvidia's "multi-timescale" outline. Depending on application (e.g., training versus inference versus mixed load), it may be more efficient to emphasize BBUs over row-level storage or vice versa.

Exhibit 31: Battery backup units provide seconds to minutes power smoothing
Vertiv Trinergy 3-Phase Large UPS



Source: Vertiv

BofA GLOBAL RESEARCH

Uninterruptible power supply (UPS)

Uninterruptible power supply (UPS) provides automated backup electrical power for a data center. UPS systems can also perform power conditioning, including voltage fluctuations, under- or over-voltage conditions, and frequency variations. We focus on the large-scale UPS market, which is most applicable to data centers.

Exhibit 32: Example of uninterruptible power supply (UPS)
The UPS contains both batteries and power conditioning equipment



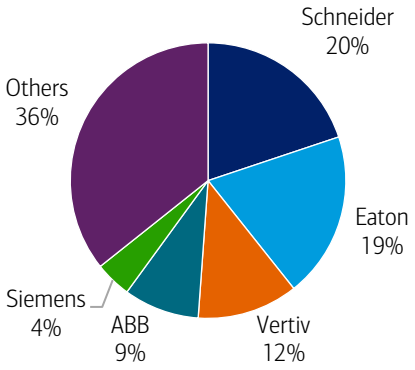
Source: Company website

BofA GLOBAL RESEARCH

UPS provides short-term backup power (typically 15-30 minutes) in the event of a power failure until the backup generators are up and running. The size of the UPS battery array is therefore proportional to the supported data center electrical load.

We estimate a $12-16bn market size, which includes a significant amount of replacement revenue. Data center operators typically replace UPS every 10 years, given increased risk of failure after this period. Based on the 13 GW of data center capacity added in 2025, we estimate the new install market size at closer to $12bn.

Exhibit 33: We estimate data center UPS to be a $12-16bn market (2025)
We estimate Schneider is the market leader



Source: BofA Global Research, company filings

BofA GLOBAL RESEARCH

Switchgear

Switchgear is equipment that controls, protects, and isolates electrical equipment. Common components include switches, fuses, isolators, relays, and circuit breakers. Data centers typically have two kinds of switchgear. A set of medium-voltage switchgear for the incoming electrical supply from the utility before it goes to the step-down transformer. Then a set of low-voltage switchgear before the electrical supply reaches the UPS backup batteries.

Exhibit 34: Example of low-voltage switchgear
Provides power distribution to data center racks

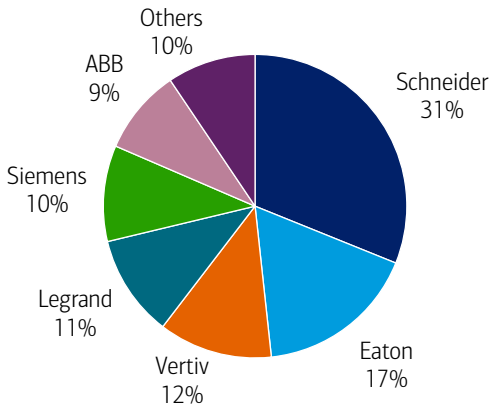


Source: Company website

BofA GLOBAL RESEARCH

Based on the 13 GW of data center capacity added in 2025, we estimate an $8-11bn market size. Switchgear is ubiquitous throughout electricity distribution and the broader market is over $100bn. The broader market is dominated by ABB, Eaton, Legrand, Schneider, and Siemens.

Exhibit 35: We estimate data center switchgear to be an $8-11bn market (2025)
We estimate Schneider is the market leader



Source: BofA Global Research, company filings

BofA GLOBAL RESEARCH

Power distribution equipment

Traditionally, electrical power went from the UPS system to power distribution units (PDUs). PDU components typically include circuit breakers, power monitoring panels, power metering, and cabling to each rack. PDUs have drawbacks, including taking up floor space in the data hall and generating waste heat.

In high-density data centers, busway is an alternative approach to PDUs. Busway is typically mounted overhead, providing power to each rack through plug-in units with breakers. The busway draws power directly from low-voltage switchgear. While busway takes up less floor space, it is typically more expensive to install and less flexible to change in rack location.

Exhibit 36: Example of rack power distribution unit
Provides plugs for servers, storage, and networking equipment inside the rack



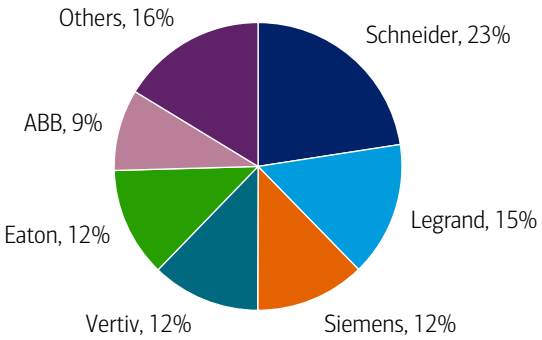
Source: Company website

BofA GLOBAL RESEARCH

Finally, rack power distribution units (rPDUs) are mounted on the rack itself. rPDUs provide outlets to plug in servers, storage, and networking equipment. These are the last step in power distribution to IT equipment. In tier 3 and tier 4 data centers, there are two rack PDUs for each rack, providing redundancy.

Based on the 13 GW of data center capacity added in 2025, we estimate an $8-11bn market size. Similar to UPS, data center operators typically replace rack PDUs every 10 years, given increased risk of failure after this period, so this includes a portion of replacement revenue.

Exhibit 37: We estimate data center power distribution to be an $8-11bn market (2025)
We estimate Schneider is the market leader



Source: BofA Global Research, company filings

BofA GLOBAL RESEARCH

THERMAL EQUIPMENT

How thermal architecture is evolving

Direct-to-Chip liquid cooling adoption path

Investor expectations for liquid cooling adoption continue to outpace real-world deployments. Conversations at industry conferences and with suppliers suggest that scaled deployment of coolant distribution units (CDUs) among colocation operators continues to ramp.

The pace may be slower than expected, but the ultimate penetration could be higher. Rather than deploying liquid cooling only where immediately required, many operators are designing facilities with optionality. Even facilities expecting rack densities below roughly 70 kW per rack, levels that could still be cooled entirely with air, are increasingly installing infrastructure that enables future migration toward liquid cooling.

Industry estimates suggest direct-to-chip cooling represented roughly 5-10% of energized IT load during 2024 and increased to approximately 15-20% during 2025. Current projections imply adoption could exceed 30% during 2026.

The dominant architecture remains single-phase direct-to-chip cooling. While two-phase cooling continues to attract attention because of its superior heat-transfer characteristics, discussions with operators, engineering firms, and cloud providers suggest adoption remains limited. Most participants believe the industry will spend considerable effort extending the usefulness of single-phase cooling before embracing more complex refrigerant-based solutions.

First-Generation Liquid Cooling Architecture

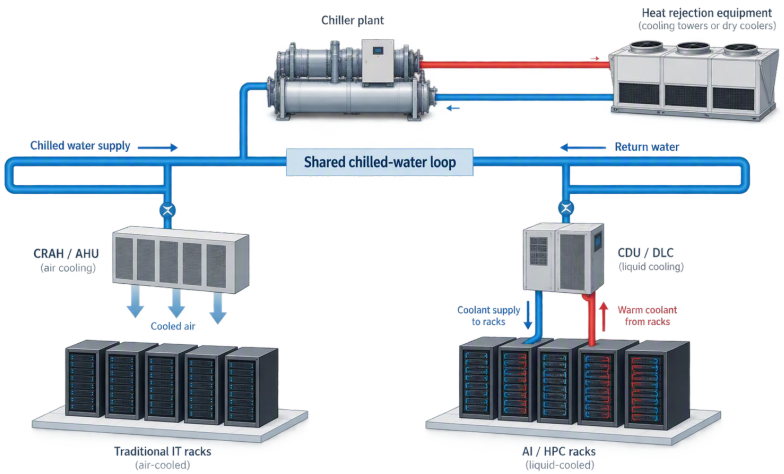
Most early direct-to-chip deployments were layered onto existing facility cooling designs. In a typical single-loop configuration, the chillers, cooling towers, or dry coolers produce chilled water. Traditional air-cooling infrastructure, such as CRAHs and fan walls, relies on facility water to cool large volumes of air.

For direct-to-chip liquid cooling, a CDU circulates coolant, typically a water-glycol mixture, through cold plates attached to processors. Heat is absorbed from the chip and returned to the CDU. Rather than actively cooling the fluid, the CDU transfers heat through a liquid-to-liquid heat exchanger into the facility chilled-water system.

The primary benefit of single-loop architecture is simplicity. By utilizing the existing facility water loop, operators minimize piping, pumps, controls, valves, and heat exchangers. Fewer components generally translate into lower capital costs, a smaller physical footprint, and reduced system complexity. This architecture allowed operators to integrate liquid cooling while minimizing changes to underlying facility infrastructure.

However, since both air-cooled and liquid-cooled systems share the same water circuit, the liquid-cooled load may receive water that is colder than necessary, reducing opportunities for free cooling and increasing reliance on chillers and mechanical refrigeration.

Exhibit 38: In single-loop cooling, both air-side and liquid-cooling systems draw from the same chilled water loop
Operators can minimize component (piping, pumps, controls) costs



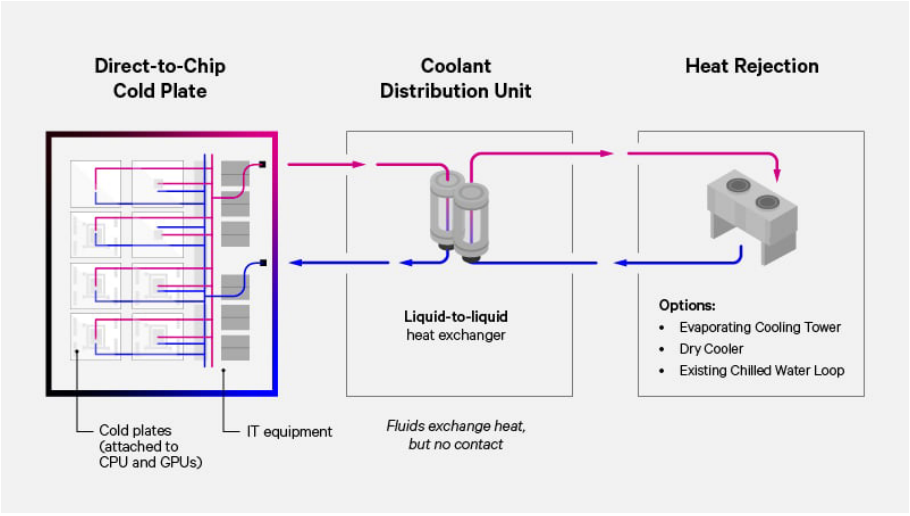
Source: BofA Global Research

BofA GLOBAL RESEARCH

Industry Moving Toward Dedicated IT Cooling Loops

The next stage of thermal architecture development increasingly involves separating technology cooling from traditional building chilled-water systems. As rack densities increase and liquid-cooled racks represent a larger share of total IT load, data center operators are finding it increasingly beneficial to divide the air-cooled and liquid-cooled infrastructure into separate water circuits connected through heat exchangers.

Exhibit 39: Multi-loop cooling systems use CDUs to circulate liquid coolant to cold plates
Direct-to-chip cooling is crucial for maintaining the longevity of hardware



Source: Vertiv

BofA GLOBAL RESEARCH

Under these designs, liquid-to-liquid heat exchangers transfer heat into a dedicated heat-rejection loop rather than into the chilled-water system used for air conditioning and dehumidification. Vertiv referred to this concept as a dedicated secondary fluid network connected to separate heat-rejection infrastructure. The industry often refers to these designs as dedicated IT loops.

Multi-loop architecture separates the cooling system into distinct temperature domains: the low-temperature loop and high-temperature loop. The low-temperature loop serves air-cooling infrastructure such as CRAHs and fan walls. The chilled water is typically

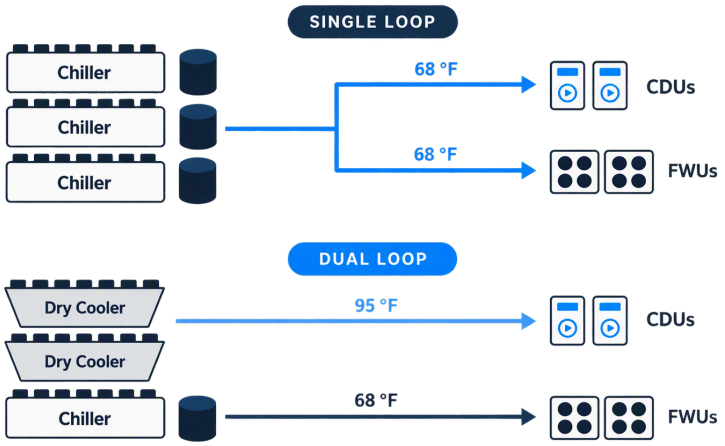
supplied by mechanical chillers. The high-temperature loop serves direct-to-chip systems through CDUs, and cold plates attached directly to GPUs and CPUs. Heat can be rejected through cooling towers or dry coolers for much of the year, which reduces reliance on mechanical cooling.

Separating these loops allows operators to optimize each independently. The technology loop can operate at elevated temperatures to maximize free cooling opportunities and reduce mechanical cooling requirements. Simultaneously, the facility air-cooling system can continue operating at lower temperatures necessary for occupant comfort, networking equipment, storage infrastructure, and other air-cooled devices.

Vertiv highlighted this tradeoff extensively during its investor day. The company argued operators increasingly face a "two-temperature problem" in which liquid-cooled AI hardware benefits from higher water temperatures, while remaining air-cooled equipment continues to require substantially cooler conditions. Their proposed solutions generally involve either dual-loop architectures or hybrid systems capable of serving both requirements efficiently.

Regardless of the specific implementation, the directional trend appears clear: liquid cooling systems are increasingly becoming independent thermal networks rather than extensions of building chilled-water infrastructure.

Exhibit 40: Single-loop vs dual-loop cooling
Dual-loop architecture allows the DTC loop to operate at warmer temperatures, which means lower energy consumption



Source: BofA Global Research

BofA GLOBAL RESEARCH

Higher temps reduce cooling, but do not eliminate chillers

Much of the industry's recent discussion has centered on Nvidia's Vera Rubin architecture and its expected use of 45°C supply-water temperatures.

At CES, Nvidia CEO Jensen Huang stated that no water chillers would be necessary for Vera Rubin deployments because of the platform's elevated supply-water temperature requirements. While technically accurate under certain operating conditions, we believe investors often overinterpret what this means for thermal infrastructure demand.

Liquid-cooling loop temperatures are only one component of facility thermal management.

Chillers are primarily used to produce low-temperature water for air-cooling systems. Even if technology loops can reject heat directly through dry coolers or other heat-rejection equipment, facilities generally continue to require cooling systems for networking gear, power electronics, storage equipment, and residual air-cooled compute infrastructure.

Vertiv's own cooling architecture discussion reinforced this point. The company repeatedly emphasized that air cooling and liquid cooling will coexist for the foreseeable future. Even facilities that are 80-90% liquid cooled still contain meaningful quantities of equipment requiring air cooling.

Moreover, thermal reliability concerns encourage operators to maintain conservative operating margins. The semiconductor industry has long recognized that higher operating temperatures accelerate device degradation. Given the cost of modern AI clusters, operators are unlikely to push thermal limits aggressively, particularly when service-level agreements and uptime requirements remain paramount.

Shift from water-cooled to air-cooled chillers

Historically, large data centers favored water-cooled chillers due to superior efficiency at scale. Water-cooled systems typically deliver better thermodynamic performance, particularly in large campus environments with substantial cooling loads.

However, several industry participants have pointed to increasing adoption of air-cooled chillers. The rationale is largely practical. Water-cooled systems require auxiliary heat rejection equipment, condenser-water infrastructure, extensive piping networks, and significant site preparation. These projects can become challenging in an environment characterized by rapid deployment schedules and supply-chain constraints.

Exhibit 41: Comparing water- versus air-cooled chiller plants
Water-cooled win on efficiency, capacity, and useful life, but air-cooled wins on initial capex, speed-to-market, and lower complexity

Issue	Water-cooled chiller plant	Air-cooled chiller plant
Electrical efficiency	Typically, more energy efficient at large, stable loads. Benefits from additional heat rejection equipment.	Rejects heat against outdoor temperature, so power requirements rise during hot weather.
Capacity	Large water-cooled centrifugal machines allow a multi-megawatt cooling load to be served by relatively few chillers.	Historically requires more units and more outdoor heat-exchanger area, but newer products are becoming much larger.
Useful life	Indoor water-cooled machines avoid sun, precipitation, and outdoor temperature cycling. Large centrifugal equipment can also be overhauled rather than completely replaced.	Outdoor coils, coatings, fan systems and controls are exposed to weather and airborne contaminants.
Upfront capex	Requires auxiliary heat rejection equipment, condenser-water pumps, piping, controls, and structural support.	Packaged equipment avoids the cooling tower and condenser-water system, simplifying installation and commissioning.
Time to market	More complex design, controls, and commissioning. Central plants cannot be added incrementally.	Factory-packaged modules can be installed on pads, connected to power, and added in phases.
Mechanical complexity	More pieces of equipment and more plant-level control interactions.	Packaged architecture is mechanically simpler.

Source: BofA Global Research

BofA GLOBAL RESEARCH

Air-cooled chillers offer several advantages: 1) lower upfront complexity, 2) smaller installation footprints, and 3) faster deployment timelines. The tradeoff is lower efficiency and often shorter useful life. Industry estimates generally suggest air-cooled chillers carry modest capital cost advantages but consume more electricity over time.

For hyperscalers focused on absolute efficiency, water-cooled systems remain attractive. For colocation and AI-focused developments prioritizing rapid deployment and flexibility, air-cooled systems are becoming increasingly competitive.

Water concerns driving migration toward dry cooling

Beyond chiller selection, public concern around water consumption is increasingly influencing data center cooling strategies. Traditional evaporative cooling towers achieve high efficiency by using water evaporation to reject heat. While thermodynamically attractive, these systems consume substantial water, a growing concern in many regions facing drought conditions or water infrastructure constraints.

As a result, operators are increasingly evaluating dry cooling technologies. Dry coolers reject heat directly to ambient air without evaporative water consumption. The primary benefit is obvious: near-zero ongoing water usage. However, dry coolers become less effective as ambient temperatures rise, particularly during peak summer conditions.

Exhibit 42: Comparing heat rejection technologies
There are tradeoffs between water usage, electricity efficiency, and footprint

	Advantages	Disadvantages
Wet cooling	• Can cool process water close to ambient wet-bulb temperature • Reliable operation in hot ambient conditions • High cooling capacity per unit	• High equipment and installation cost • High operating costs (e.g.. frequent water treatment) • Relatively high-water consumption
Dry cooling	• Low operating and initial maintenance costs • Reliable operation in most climate conditions • No water consumption and no water loss	• Cannot cool process water below dry bulb temp without evaporative (adiabatic) pre-cooling • Cooling performance declines in hot weather
Adiabatic cooling	• Improved cooling capacity relative to a dry cooler • Reliable operation in dry environments • Consumes less power than mechanical AC units	• Increases a data center's water usage compared to dry cooling • Lower efficiency in humid climates • Typically requires larger footprint than wet cooling towers

Source: BofA Global Research

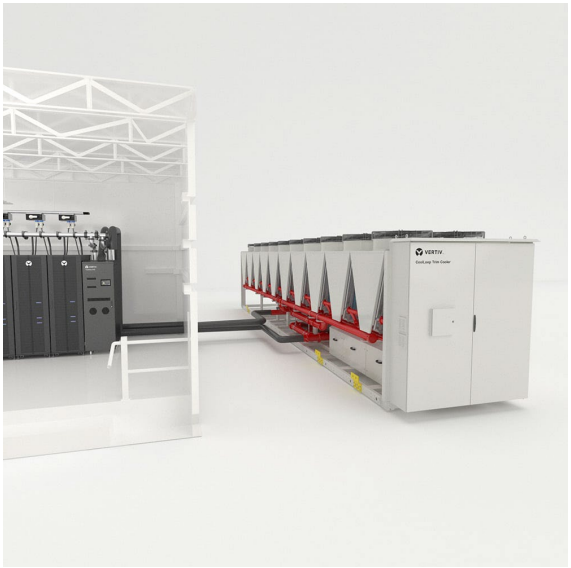
BofA GLOBAL RESEARCH

Trim coolers emerging as a hybrid alternative

A trim cooler combines elements of dry cooling and mechanical chilling within a single system. Under favorable ambient conditions, the equipment behaves similarly to a dry cooler, rejecting heat efficiently without significant compressor operation. During hotter conditions, mechanical cooling stages activate to maintain performance.

Vertiv positioned its CoolLoop Trim Cooler as a way to bridge the gap between traditional chillers and dry coolers. The company argues that trim coolers can substantially expand the number of hours during which facilities operate in free-cooling mode while preserving cooling performance during challenging weather conditions.

Exhibit 43: Vertiv's CoolLoop trim cooler
Trim coolers provide supplemental cooling when ambient conditions exceed the capability of the dry cooler



Source: Vertiv

BofA GLOBAL RESEARCH

Absorption chillers another path for energy efficiency

Unlike conventional vapor-compression chillers, which rely primarily on electricity-driven compressors, absorption chillers use heat sources to drive the cooling cycle. Potential inputs include waste heat, steam, district energy systems, or industrial process heat.

Absorption chillers can reduce electrical consumption and improve facility energy economics when sufficient low-cost heat sources are available. On-site power generation, such as natural gas turbines or reciprocating engines, generate heat which can be used to power absorption chillers.

We argue the potential market for absorption chillers will grow as the number of data centers with on-site power generation grows.

Exhibit 44: JCI's YHAU-C Single Effect Steam Chiller
The absorption chiller uses steam rather than electric compressors to produce chilled water



Source: Johnson Controls

BofA GLOBAL RESEARCH

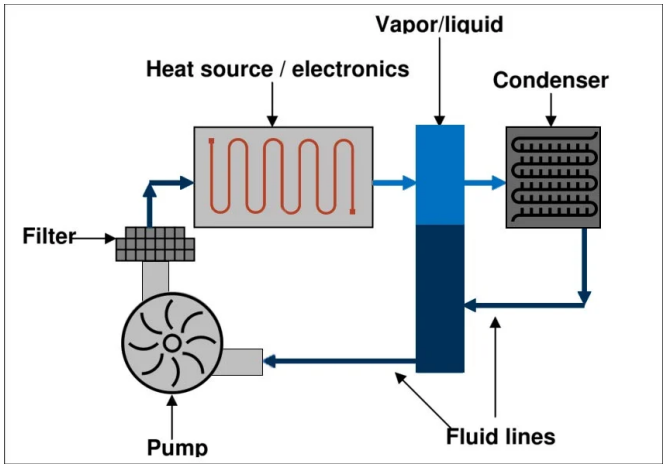
Two-phase cooling remains a long-term opportunity

While single-phase direct-to-chip cooling is rapidly scaling, two-phase cooling remains on the horizon. In two-phase systems, the cooling fluid boils during heat absorption, providing substantially greater heat-transfer efficiency. The theoretical performance benefits are attractive, especially as rack power moves toward several hundred kilowatts and potentially megawatt-scale configurations.

Despite these advantages, industry adoption remains limited. Comments from conference participants, engineering firms, cloud providers, and thermal equipment manufacturers consistently pointed to the same concerns: 1) higher system complexity, 2) uncertain operating costs, 3) limited maintenance experience, and 4) lack of mature operating standards.

For now, the industry's focus remains firmly on extending the capabilities of single-phase liquid cooling.

Exhibit 45: Operating principle of a pumped two-phase cooling system
Two-phase cooling removes heat through refrigerant evaporation and condensation



Source: Eaton

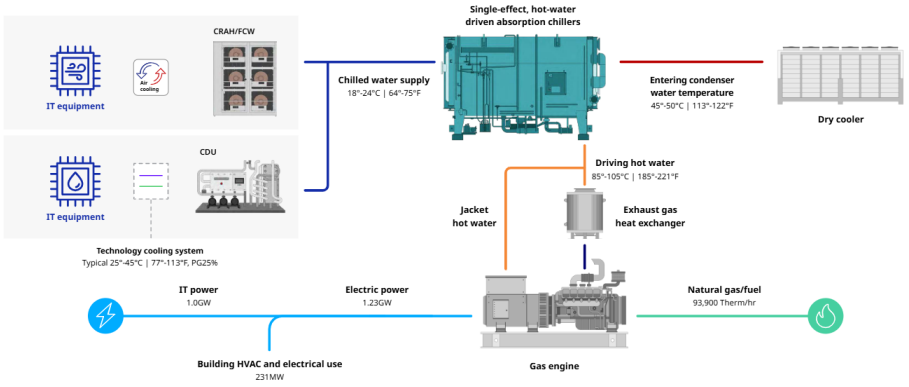
BofA GLOBAL RESEARCH

Chillers

Chillers are refrigeration units that chill facility water for distribution inside the building. Heat is absorbed from the facility water into the chilled coolant through the evaporator heat exchanger. The coolant is then run through a compressor, increasing the pressure and temperature. The condenser heat exchanger transfers the heat to either outside air or water. Finally, the refrigerant runs through an expansion value, lowering the pressure and temperature before going into the evaporator heat exchanger. There are four varieties:

- **Air-cooled chillers** are typically located outside the building. The condenser transfers heat to the outside air.
- **Water-cooled chillers** are typically located inside the building (mechanical space). It uses a liquid-to-liquid heat exchanger to transfer heat to a dedicated cooling loop. That loop is connected to a cooling tower outside the building.
- **Trim coolers** are located outside the building, either on the rooftop or on nearby pads**.** Trim coolers combine large dry coolers (i.e., air-to-air heat exchangers) with an air-cooled chiller backup. Hot facility water circulates to the trim cooler. On days when the ambient temperature is low, the facility water can then return to the building without mechanical cooling. This is also known as “free-cooling.” On hotter days, the air-cooled chiller will activate, “trimming” the remaining few degrees off the facility water. The result is lower energy consumption (compared to traditional air-cooled chillers alone).
- **Absorption chillers** are typically located inside the data center. They use a specialized refrigerant, typically lithium bromide salt mixed with water. They do not use traditional compressors or mechanical cooling. Instead, they use changes in the pressure and concentration of the lithium bromide salt to provide cooling. A facility loop gathers heat and receives chilled water from the evaporator/absorber chamber. A separate loop runs outside of the building to provide heat rejection (e.g., cooling tower or air-to-air heat exchanger). The refrigerant itself circulates between the evaporator/absorber chamber and the condenser/generator chamber. Instead of using electricity to run a compressor, absorption chillers use a heat source to heat the lithium bromide salt and water refrigerant, driving the absorption cooling cycle. One source of heat would be from natural gas engines. As a result, absorption chillers use significantly less electricity compared to traditional systems.

Exhibit 46: Diagram of 1 GW data center with absorption chillers and on-site power generation
Waste heat from the gas engine drives the absorption cooling cycle



Source: Johnson Controls

BofA GLOBAL RESEARCH

- **Compressors:** equipment that increases refrigerant pressure by reducing the volume. Compressors are the most critical part of the chiller and largely determine its capacity, efficiency, and power usage. There are many different types of compressors, but the largest chillers tend to use centrifugal compressors. These compressors pull refrigerant using centrifugal force, and compress it using an impeller. They are more energy efficient, particularly in large capacity applications.

Exhibit 47: Example of water-cooled centrifugal chiller
Chills facility water to circulate to computer room air handlers



Source: Company website

BofA GLOBAL RESEARCH

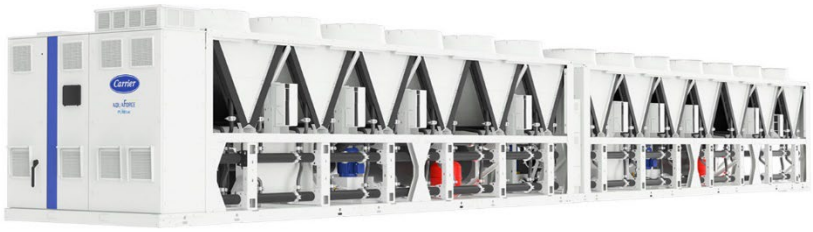
Exhibit 48: YORK high-capacity chiller
The YK family is one of JCI's flagship products because centrifugal chillers are used in many large chilled-water plants



Source: Johnson Controls

BofA GLOBAL RESEARCH

Exhibit 49: Example of air-cooled screw chiller
Uses ambient air to produce chilled water for data center cooling applications

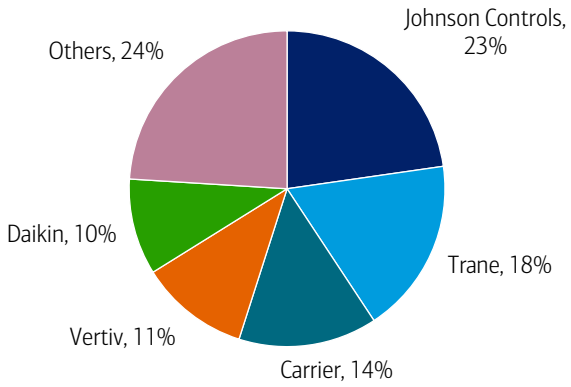


Source: Carrier

BofA GLOBAL RESEARCH

Every one megawatt of power supplied to a data center requires approximately 285 tons of cooling, similar to the requirements for a 115,000 square foot commercial building. A 285-ton chiller is roughly $380-420,000. Based on the 13.2 GW of data center capacity added in 2025, we estimate a $4.4-6.2bn market size.

Exhibit 50: We estimate data center chillers are a $4-6bn market (2025)
We estimate JCI is the market leader



Source: BofA Global Research, company filings

BofA GLOBAL RESEARCH

Cooling towers

Cooling towers are heat rejection equipment that dissipates heat into the outside air. Cooling towers are giant heat exchangers used to reject heat created by IT equipment inside the data hall. Cooling towers take condenser water (hot) water and cool it using outside air. They come in three varieties.

- **Wet cooling towers** use water evaporation to create additional cooling capacity. However, this consumes a large amount of water (e.g., "open circuit"). Wet cooling towers are efficient in hot and dry regions; however, the costs are higher for the equipment & installation and the water consumption.

- **Dry cooling towers** are closed-circuit (e.g., no water loss) towers where there is no direct contact between the ambient air and the fluid being cooled. Facility water transfers to the air through radiators. Dry towers have lower initial maintenance costs and can work in most climate conditions, however there is a lower capacity and can't cool below a certain temperature.

- **Adiabatic cooling** uses evaporation to reduce energy consumption in cooling. Warm outdoor air is drawn through water-moistened pads or a fine mist. As the water evaporates, it absorbs heat. The resulting cooler air passes over closed-loop heat exchanger coils. The result is better heat rejection relative to a dry cooler. As a result, the overall system uses less electrical power However, adiabatic cooling increases a data center's water usage. The trade-off between water usage and energy savings varies by location and climate. Adiabatic cooling is highly effective in dry environments.

We estimate cooling towers costs $330,000 on average, but these costs vary with size. A single cooling tower can provide heat rejection for 3-4 MW of supplied electrical power (e.g., 1,000 tons of cooling). We estimate that the cooling tower market for data centers was $1.0-1.4bn in 2025.

Exhibit 51: Example of cooling tower
This tower is 32 feet high with a cooling capacity of 2,400 tons

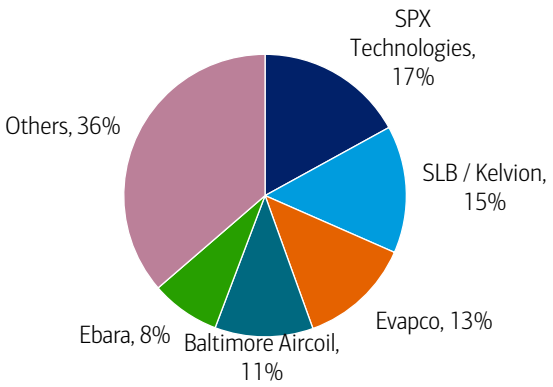


Source: Company website

BofA GLOBAL RESEARCH

Wet vs dry cooling towers is something to be considered for data centers. Wet cooling towers are open-circuit cooling towers where water comes in contact with the ambient air. Dry cooling towers are closed circuit cooling towers where there is no direct contact between the ambient air and the fluid being cooled. Wet cooling towers are reliable in hot ambient temperatures and have high cooling capacity, however the costs are higher for the equipment and installation and the water consumption. Dry towers have lower initial maintenance costs and can work in most climate conditions, however there is a lower capacity and can't cool below a certain temperature.

Exhibit 52: We estimate data center cooling towers to be a $1.0-1.4bn market (2025)
We estimate SPX is the market leader



Source: BofA Global Research, company filings

BofA GLOBAL RESEARCH

Air cooling equipment

Key components

- **Computer room air conditioners (CRACs):** used in small data centers, these are full air conditioning units located inside the data hall. They are tied to a condenser located outside of the building. CRACs have small cooling capacities.
- **Computer room air handlers (CRAHs):** used in larger data centers. These blow air over a coil with chilled facility water to lower the temperature inside the data hall. They are connected to a chiller. Often used in data centers with raised floors, where cold air is blown under the IT equipment and hot air rises to the top of the data hall.
- **Fan walls:** a type of high-capacity CRAH. Large fans blow air over chilled water in coils. Designed to lower air temperatures by 15-25 degrees (with water temperatures rising 15-25 degrees). Often used in data centers without raised floors.

This equipment includes computer room air handlers (CRAHs; use chilled water and fans to deliver conditioned air), fan walls (high-capacity CRAHs), air handling units (AHUs; large air volumes generally connected to economizer systems), air-side economizers (systems using cool outside air when conditions permit), and hot-aisle/cold-aisle containment systems.

Exhibit 53: Computer room air handling unit
CRAHs are typically located around the walls of the data hall



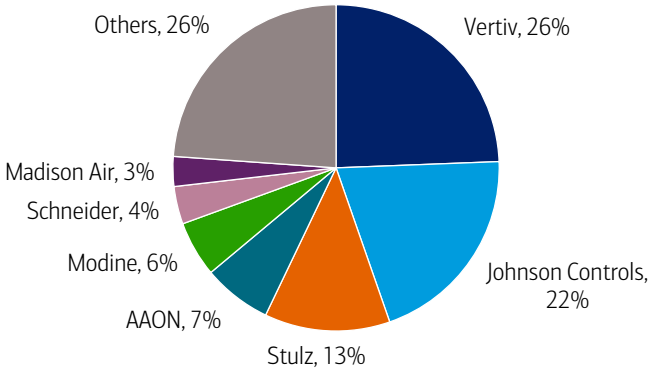
Source: Company filing

BofA GLOBAL RESEARCH

CRAHs, fan walls, and hot-aisle/cold-aisle containment systems are located in the "white space" of a data center (i.e., where IT equipment is located). Air handling units and air-side economizers sit outside of the data center. Other air cooling equipment includes in-rack cooling fans and related sensors & controls.

We estimate the air cooling equipment and related equipment is a $6-8bn market.

Exhibit 54: We estimate data center air cooling equipment to be a $6-8bn market (2025)
We estimate Vertiv is the market leader



Source: BofA Global Research, company filings

BofA GLOBAL RESEARCH

Liquid cooling equipment

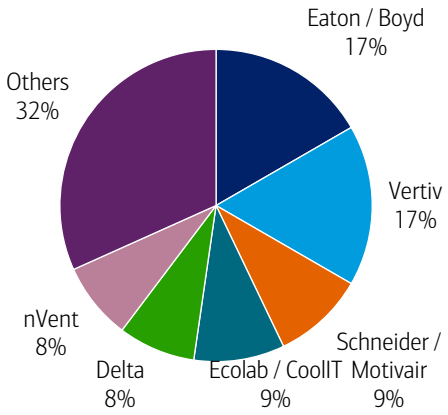
- **Coolant Distribution Units (CDUs):** circulates and pumps coolant in a closed-loop system to row manifolds, rack manifolds, and through either cold plates or rear door heat exchangers. The coolant (typically a water-glycol mix) returns to the CDU where it runs through a heat exchanger. The heat is there transferred either to air ("liquid-to-air CDU"), or facility water/dedicated cooling loop ("liquid-to-liquid CDU") for heat rejection. CDUs will typically come with two pumps, offering redundancy around mechanical failure. CDUs will also have sensor to monitor and control temperature, pressure, and flow rate. CDUs come in two forms:

- **Liquid cold plate:** a metal block specifically designed to sit on top of a chip with microchannels for coolant to flow through the chip. The liquid cold plate facilitates heat exchange from the semiconductor into the cooling fluid.
- **Quick disconnectors:** couplings with zero leaking of coolant. Used to connect the liquid cold plates to the rack manifold and to connect the rack with the row manifold. Quick disconnectors in data center applications typically have a latch to secure the connection.

The largest part of a direct-to-chip liquid cooling system is the coolant distribution unit (CDU), which comprises approximately 65-75% of the total equipment cost. Manifolds deliver coolant to the racks. Cold plates sit on top of the semiconductor or server and facilitate heat absorption by the coolant. A rear door heat exchanger (RDHx) attaches to the back of a server rack to remove heat from exhaust air. We estimate a $4-5bn market size in 2025, which is likely to more than double in 2026.

Our market size estimate excludes thermal buffer tanks (generally located outside), immersion cooling equipment, and quick disconnectors (used in coupling cold plate to rank manifold).

Exhibit 55: We estimate data center liquid cooling equipment to be a $4-5bn market (2025)
We estimate Eaton/Boyd is the market share leader



Source: BofA Global Research, company filings

BofA GLOBAL RESEARCH

Given the high-growth prospects, it is not surprising that the liquid cooling market has seen many new entrants. By our count, there are 30 vendors offering over 100 CDU variants in the market today. New entrants, such as Carrier, JetCool, LG Electronics, Munters, Madison Air, and Trane, may have garnered more revenue in 2026, but our focus is on 2025 given data availability.

Given the conservative nature of data center operators, we argue that reputation and service capability will play a major factor in decision making. We think this bodes well for Vertiv, which has over 5,300 field service engineers globally offering same-day service in most locations.

Comparing two reference designs

Johnson Controls reference design

Exhibit 56: JCI's absorption reference design requires the most equipment
JCI product amounts by reference design type

	Air-cooled	Water-cooled	Absorption
Chillers	352	96	672
CDUs	1,056	528	2,112
CRAHs	784	392	1,568

Source: Company reports

BofA GLOBAL RESEARCH

Air-cooled

Released May 5, 2026, as the second guide in the series, RDG 401 is a blueprint for cooling a 1 GW-scale AI factory using an entirely air-cooled chiller architecture. The architecture integrates high-efficiency air-cooled YORK centrifugal chillers (YDAM and YVAM models), fan coil walls (FCWs), and coolant distribution units (CDUs) to manage both air- and liquid-cooled IT loads, with sizing references built around 220 MW compute clusters. Each compute cluster is 90% liquid cooled and 10% air cooled. The design uses 352 chillers (288 high temperature chillers and 64 medium temperature chillers, 1,056 CDUs and 784 CRAHs. The medium temperature chillers pair with the CRAHs and the high-temperature chiller array pairs with the CDUs. A traditional data center would use just high-temperature chillers, but the JCI design uses both medium and high-temperature chillers to maximize tokens per watt resulting in up to 50 MW returned to the data center.

The design has water-free heat rejection (eliminating cooling towers to save 12+ million gallons daily), a 32% improvement in annual energy consumption via intelligent use of redundant chillers, and 20 MW of peak power savings from mitigating heat-island effects. The design also uses 27% fewer chillers from raising the chilled water temperature to support warm-water Technology Cooling System (TCS) loop. The max ambient temperature for the design is 60C/140F.

Water-cooled

Launched on February 2, 2026, as the inaugural guide in the series, RDG 402 is a blueprint for cooling a 1 GW-scale AI factory using an entirely water-cooled chiller architecture. The reference design is built around 220 MW compute clusters. The compute clusters have a single networking-core capable of 105MW of computer with a flexible liquid-to-air ratio. The design uses 96 chillers (76 high-temperature, 12 swing chillers, and 8 medium temperature chillers), 216 dry coolers, 528 CDUs, and 392 CRAHs.

The design is positioned around three outcomes: zero water usage in the cooling process, future-readiness for high-temperature coolant loops, and optimization for high-density AI workloads with energy-efficient performance, while also reducing heat-rejection equipment requirements. By using a mix of chillers/ bifurcated loops it maximizes tokens per watt resulting in up to 40 MW returned to the data center. The design of four compute quadrants with a central core results in a 35% reduction in required heat rejection equipment, Capex, and space resulting in lower costs and quicker deployment.

Absorption

Introduced July 28, 2026, RDG 407 is blueprint for thermally driven absorption chillers with on-site power generation for a 1GW AI Factory. The design converts the ~57% of energy typically lost as waste heat from on-site power generation (natural-gas turbines, fuel cells, etc.) into productive cooling. By cutting cooling-related electrical demand by up to 44%, the design frees power for compute supporting up to 97 MW of additional AI capacity in a 1 GW factory. The repeatable, modular architecture scales from 100 MW campuses to gigawatt facilities without redesign, targeting a PUE as low as 1.23 and up to 43% lower cooling-system CO_2 emissions. The design uses 576 reciprocating gas engines, 672 chillers (576 hot-water driven absorption chillers and 96 direct gas-fired absorption chillers), 2,112 CDUs, and 1,568 CRAHs.

Trane reference design

Exhibit 57: TT's air-cooled reference design requires the most equipment
TT product amounts by reference design type

	Air-cooled	Water-cooled
Chillers	340-350	100-108
CDUs	410	410

Source: Company reports

BofA GLOBAL RESEARCH

Air-cooled

Trane Reference Design #504, published March 16, 2025, outlines a high-efficiency cooling architecture designed to support a 1 GW AI data center using air-cooled chillers. The design's primary benefits include zero water consumption by eliminating cooling towers, a 14% cooling-system efficiency improvement vs a single-loop architecture, and the ability to free up ~35 MW of power for additional compute capacity through dedicated temperature loops that allow 93% of load to operate at a more efficient 26°C supply temperature. The design has reduced installation complexity and faster deployment through fully packaged air-cooled systems requiring no condenser-water infrastructure. Based on the published sizing assumptions, each cooling block supports approximately 102.4 MW of liquid-cooled load and 13.7 MW of air-cooled load, requiring roughly 35 high-temperature chillers (102.4 MW ÷ 2.93 MW/chiller), 5 medium-temperature chillers (13.74 MW ÷ 2.75 MW/chiller), and about 8 CDUs (102.4 MW ÷ 14 MW per CDU) per cooling block. Scaling this architecture to a full 1 GW AI factory implies approximately 340 to 350 air-cooled chillers and ~410 CDUs.

Water-cooled

Trane Reference Design #501, published November, 6, 2025, outlines a high-efficiency cooling architecture designed to support a 1 GW AI data center using water-cooled chillers. The design is engineered around the NVIDIA Omniverse DSX blueprint for AI factories capable of running NVIDIA's next-generation Vera Rubin chips. The design uses high-efficiency CenTraVac water-cooled chillers (up to 11 MW / 3,000 tons each), dry coolers for heat rejection, fan coil walls for air-cooled loads, and CDUs for liquid-cooled AI racks, with dedicated temperature loops that optimize efficiency across both air- and liquid-cooling applications. Key benefits include lower cooling power consumption that frees additional electrical capacity for compute nodes, improved PUE, zero water consumption through the use of dry coolers instead of cooling towers, and reduced capital costs via the use of swing chillers that can serve either high- or low-temperature loops, reducing the total number of chillers required while maintaining redundancy.

The design supports approximately 13.7 MW of air-cooled load and 102.4 MW of direct-to-chip liquid-cooled load per cooling block, with individual chillers serving roughly 6.9 MW (air-side loop) or 10.2 MW (liquid-cooling loop). The design utilizes ~10 liquid-cooling chillers and 2 air-cooled chillers per cooling block or 116.14 MW per block implying 100 – 108 chillers for a 1GW data center. Based on Trane's published sizing assumptions, the 1 GW AI Factory design appears to require ~410 CDUs, each rated at up to 14 MW, to support the liquid-cooled portion of the data hall.

On-site prime power

Why look at on-site power? Time to token

On-site power generation has never been common for US data centers. Diesel generators are only for backup, not prime, power. Supply is lagging further and further behind. For example, power generation projects going live in 2025 took a median of 5.1 years to connect to the grid (e.g., late 2019 initial request for a 2025 go live). As a result, data center developers are turning to "bring your own power" solutions while waiting for firm grid connections, or bridge-to-grid solutions.

158 GW of announced behind-the-meter projects

Using several sources (Bloomberg New Energy Finance, Wood Mackenzie, Platts, company announcements), we piece together 109 discrete announced projects representing 158 GW of planned onsite gas generation projects. This is the closest proxy to a behind-the-meter addressable market, though some are technically selling power into the grid as well as to a datacenter.

Approximately 2 GW of projects are operational already. We segment the remaining pipeline into four categories:

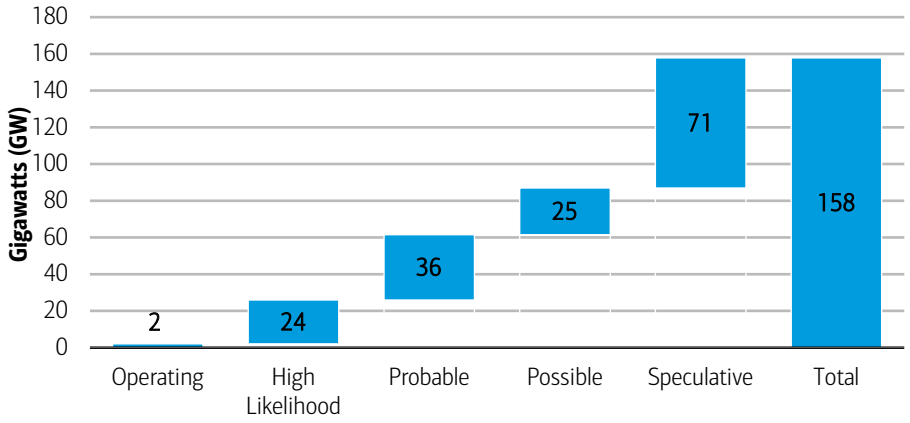
1. **High likelihood.** 24 GW are under construction, with a start date and an announced power developer or original equipment manufacturer

2. **Probable.** 36 GW have a listed start date, an equipment provider or power developer and are in pre-construction status

3. **Possible.** 25 GW have a start date, listed equipment provider or power developer and are in announced status

4. **Speculative.** The remaining 71GW has no start date or no announced power

Exhibit 58: We track 158 GW of planned on site gas generation projects
Behind the meter projects by capacity in gigawatts

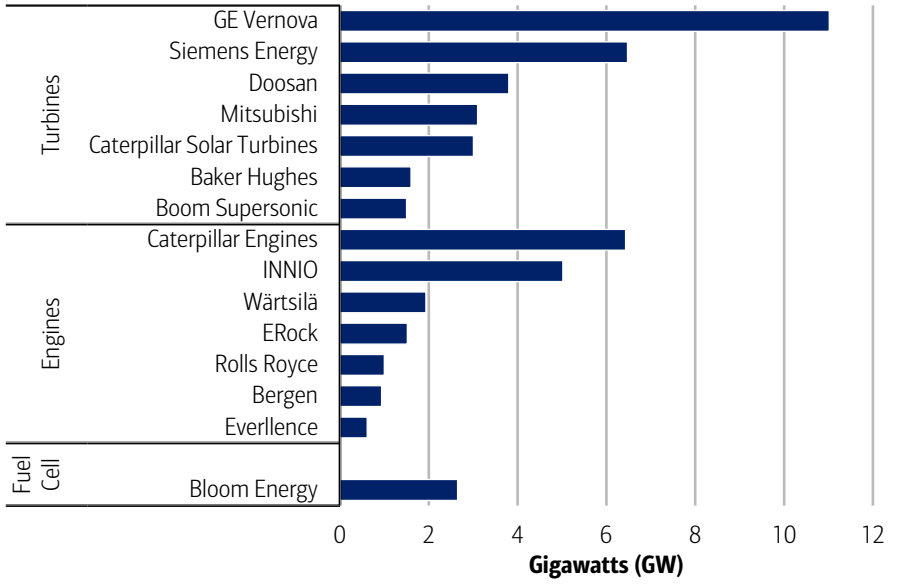


Source: BofA Global Research Estimates, Bloomberg New Energy Finance, Wood Mackenzie, Platts, Company Filings

BofA GLOBAL RESEARCH

On 51 GW of these projects, a specific equipment provider is listed. To be clear this is not a random sample and is a biased view of market share in behind the meter projects. Larger projects are more likely to use large gas turbines and be announced projects.

Exhibit 59: Mix of original equipment manufacturer on 51 GW of announced projects
Projects tied to each manufacturer in gigawatts (GW)

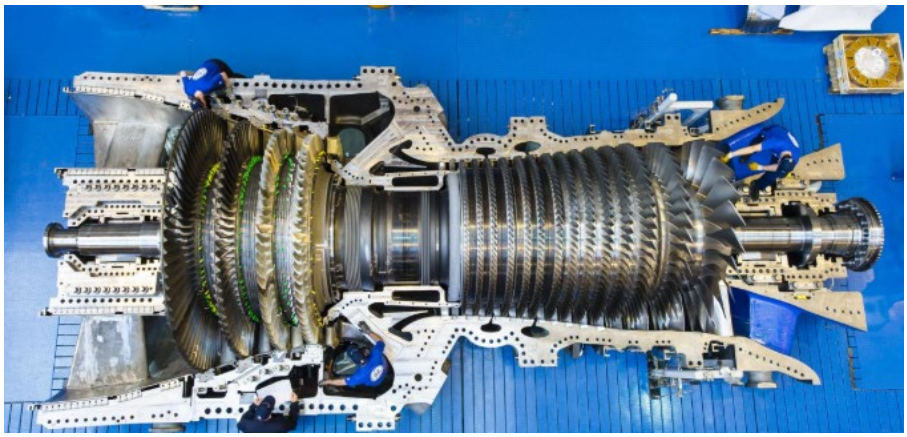


Source: BofA Global Research Estimates, Bloomberg New Energy Finance, Wood Mackenzie, Platts, Company Filings

BofA GLOBAL RESEARCH

Heavy-duty natural gas turbines

Exhibit 60: GE Vernova 9HA turbine being built (note three employees in frame for scale)
9HA turbines are typically used in combined cycle applications



Source: Company website

BofA GLOBAL RESEARCH

Heavy-duty combined cycle turbines turn 64% of the energy from burning natural gas into electricity. The initial ignition powers the combustion turbine. Then the exhaust gas is used to heat water creating steam which in turn drives a steam turbine. This fuel efficiency makes them preferred by utilities. They are large-scale plants. For example, operating two GE Vernova 9HA.02 turbines and one steam turbine (i.e., 2x1 configuration) would have a 1,680 MW output.

For on-site power applications, the key drawbacks are their current long lead times, large format sizes (making them less modular), high initial capex costs, long start-up times (i.e., 30 minutes or more), and the need for a natural gas connection.

Aeroderivative natural gas turbines

Exhibit 61: GE Vernova TM2500 DLE installed in Veresegyház, Hungary
Installation in less than three weeks at manufacturing plant



Source: Company website

BofA GLOBAL RESEARCH

Aeroderivative natural gas turbines are based on the designs of jet engines. They are typically used in simple cycle configurations for peaking applications, given their lower efficiency. For the current generation of aeroderivative turbines, 43% of energy from burning natural gas is converted into electricity. They are smaller plants, generating 30-120 MW of electricity per turbine. Trailer-mounted versions (e.g., GE Vernova's TM2500 DLE) can be quickly installed after natural disasters.

For on-site power applications, the key drawback are medium start-up times (i.e., 5-10 minutes), medium format sizes, ongoing operating costs, and the need for a natural gas connection.

Industrial natural gas turbines

Exhibit 62: Caterpillar's subsidiary Solar Turbines' Titan 130 installed in Tarkwa, Ghana
Providing 16.5 MW of electricity to a mining location



Source: Company website

BofA GLOBAL RESEARCH

Industrial natural gas turbines are typically used in robust settings and have longer operating times between required overhauls. However, they have lower efficiency versus combined cycles. For the current generation of industrial turbines, 39-40% of energy from burning natural gas is converted into electricity. They are larger plants, generating 50-300 MW of electricity per turbine.

For on-site power applications, the key drawback are longer start-up times (e.g., 20 minutes), larger format sizes, and need for a natural gas connection.

Reciprocating natural gas engines

Exhibit 63: Caterpillar CG170B-20 offers 2.3 MW of power generation in 120 sq. ft.
Up to 45% electrical efficiency



Source: Company website

BofA GLOBAL RESEARCH

Reciprocating natural gas engines are used in industrial, commercial, and institutional facilities for both power generation and combined heat and power applications. The US has over 6 GW of installed capacity with an 18% capacity factor, suggesting a mix of prime and backup power applications. Best-in-class reciprocating engines can turn 50% of the energy from burning natural gas into electricity. They are typically smaller plants, generating 1-30 MW of electricity per engine. They ramp quickly and are efficient even with partial loads (e.g., 25% of nameplate capacity).

For on-site power applications, the key drawbacks are the initial capex costs, the need for a natural gas connection, and higher lifetime maintenance costs with scheduled downtime.

Fuel cells

Exhibit 64: Bloom Energy solid oxide fuel cells in Ulsan, South Korea
These fuel cells produce 0.1 MW and are part of a 1 MW project



Source: Company website

BofA GLOBAL RESEARCH

Fuel cells convert chemical energy into electricity through a reaction, without combustion. There are a wide variety of fuel cells, but the most common are solid oxide fuel cells (which typically use natural gas) and proton exchange membrane fuel cells (which typically use hydrogen). CoreWeave deployed 14 MW of solid oxide fuel cells at a site in Volo, Illinois. Microsoft, Ballard Power Systems, and Caterpillar demonstrated a pilot to use hydrogen fuel cells in backup power applications.

For on-site power applications, the key drawbacks are the initial capex costs, the need for a natural gas connection, and limited large-scale operational history of the technology.

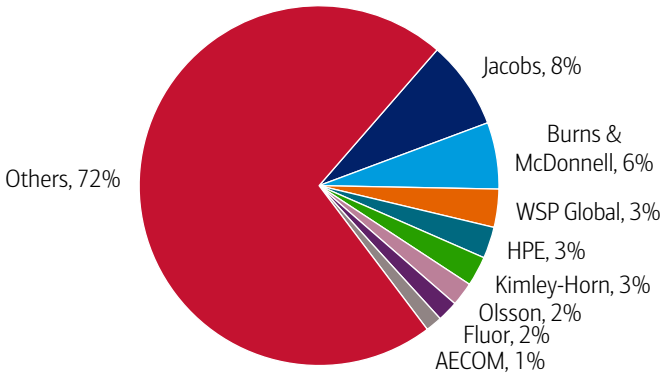
Engineering

Design & engineering services generally cost 4.5-6.5% of the infrastructure costs of a data center (e.g., excluding IT equipment). Publicly traded US firms covered by our colleague Michael Feniger include Jacobs Solutions (J), Fluor (FLR), and AECOM (ACM). These firms are highly diversified among end markets, with data centers representing a small percentage of total revenue.

These engineers plan the electrical, mechanical, cooling, fire protection, and physical security systems of the data center. Importantly, they must understand the IT infrastructure (e.g., network, routing, storage), which has implications for the physical infrastructure requirements.

Based on the 13 GW of data center capacity added in 2025, we estimate a $5.3-7.3bn market size.

Exhibit 65: We estimate data engineering services to be a $5.3-7.3bn market (2025)
We estimate Jacobs is the market leader in a fragmented market



Source: BofA Global Research

BofA GLOBAL RESEARCH

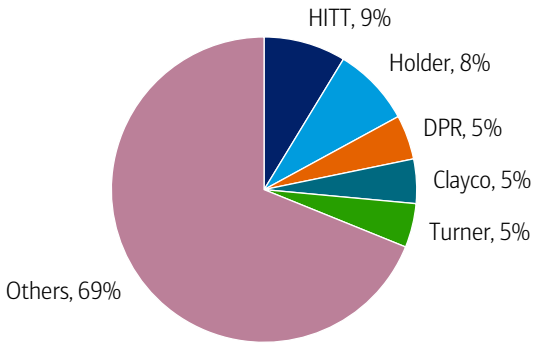
Construction

Construction firms oversee all aspects of the construction project including project management, specialty contractors, material purchasing, and equipment rental. These firms report modest operating margins, given the large amount of pass-through costs. The average operating margin is approximately 4% among publicly traded contractors with data center exposure.

Publicly traded firms include Balfour Beatty and Skanska. There are also smaller, private construction firms that specialize in data centers, such as US-based T5 Construction Services and Ireland-based Mercury Engineering.

Based on the 13 GW of data center capacity added in 2025, we estimate a $60-80bn market size. However, this would include material & equipment pass-through costs. Using an average margin of 4%, this implies $2.4-3.2bn of operating profit.

Exhibit 66: We estimate data center construction to be a $60-80bn market (2025)
We estimate the top five firms have a 31% market share



Source: BofA Global Research

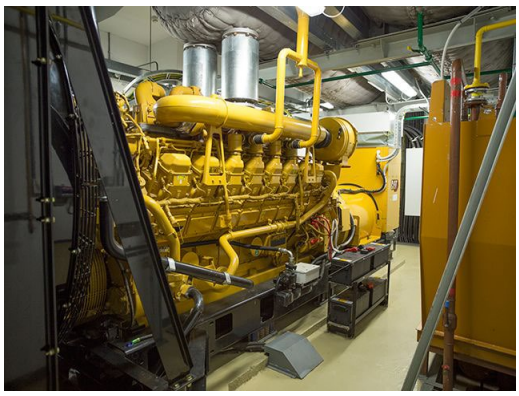
BofA GLOBAL RESEARCH

Backup generators

Typical diesel backup generators cost $1mn/MW. Total system costs would include fuel tank, fuel pump, and installation costs, would increase the all-in costs to closer to $2mn/MW. However, this is only a rough approximation. The actual costs will vary based on requirements (e.g., selective catalytic reduction systems, sound-attenuated enclosures, level of redundancy). Generators are typically sized to fully supply the

electrical consumption of the data center. For example, a 100MW datacenter will typically have 100MW worth of generator power on site to ensure 99.999% uptime for clients.

Exhibit 67: Example of a diesel generator
This generator creates 1.25MW of electricity and consumes 86 gallons of diesel per hour

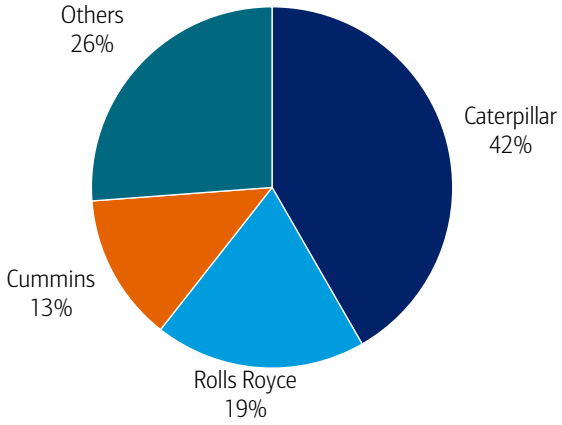


Source: Company website

BofA GLOBAL RESEARCH

Based on the 13 GW of data center capacity added in 2025, we estimate a $12-16bn market size for generator equipment only (excluding ancillary products and installation costs).

Exhibit 68: We estimate data center generators to be a $12-16bn market (2025)
We estimate Caterpillar is the market leader



Source: BofA Global Research, company filings

BofA GLOBAL RESEARCH

Servers

Servers are the largest single product category of data center capex. In 2025, data centers spent over $415bn of servers. On a dollar basis, AI servers comprised over half of this spending, but traditional servers represented the vast majority on a unit basis.

Exhibit 69: Example of a data center server
This server contains eight Nvidia H200 GPUs and two Intel CPUs and draws 3,200 watts



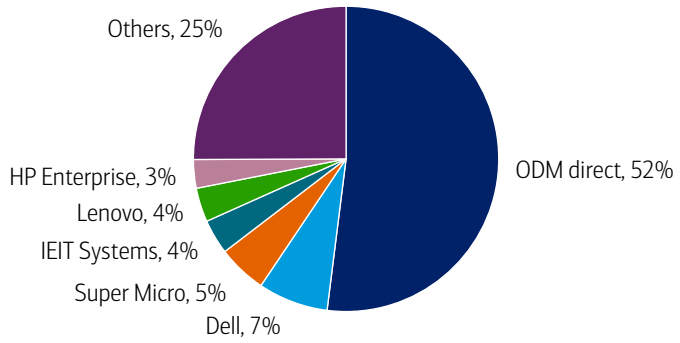
Source: Company website

BofA GLOBAL RESEARCH

Original design manufacturers (ODMs) are companies that produce servers to their own design. For example, Google's Tensor Processing Units (TPUs) are custom semiconductors and these are put into servers at Google Cloud data centers. Similarly, Amazon Web Services has its own custom semiconductors (Graviton) and servers. Original equipment manufacturers (OEMs) are companies building servers to clients' specifications.

The largest OEM is Dell Technologies, with an estimated 7% market share.

Exhibit 70: We estimate data center servers to be a $415bn market (2025)
Dell is the largest OEM



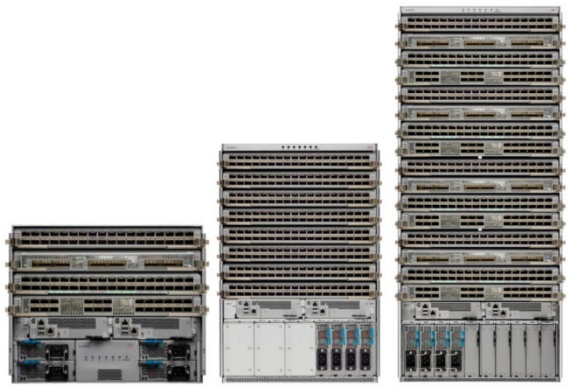
Source: BofA Global Research

BofA GLOBAL RESEARCH

Networking equipment

Networking equipment includes several different pieces of equipment. Switches communicate within the data center or local area network. Some cloud service providers use "white box" networking switches (e.g., manufactured by third parties to their specifications). Typically, each rack would have a networking switch. Routers handle traffic between buildings, typically using internet protocol (IP).

Exhibit 71: Example of routers
Routers can handle terabytes of data throughput per second



Source: Company website

BofA GLOBAL RESEARCH

The rise and acceleration of agentic and generative AI workloads have created unprecedented demand on the network, given the data and compute-intensive nature of these workloads. Fundamentally, AI workloads are bandwidth-intensive, connecting hundreds of processors with gigabits of throughput. As these AI models grow, the number of GPUs required to process them grows, meaning bigger networks are required to interconnect the GPUs..

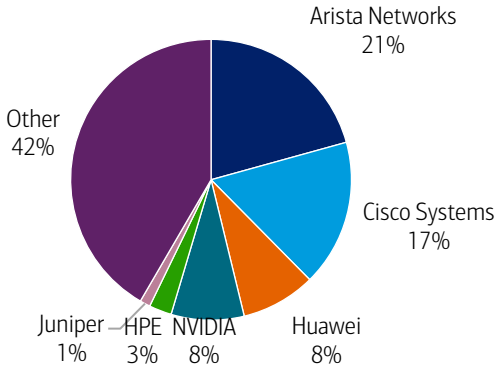
Per Gartner, the total addressable market for AI infrastructure, inclusive of network fabrics, semiconductors, servers and more, is expected to reach near $3.9tn by 2030. At a high level, networking vendors are benefitting from elevated investments in both back-end and front-end Cloud networks following major capex increases by hyperscalers. For example, the top five hyperscalers (Meta, Microsoft, Amazon, Google, and Oracle) are expected to grow their capex by a remarkable 96% y/y in 2026 after an impressive 72% y/y in 2025. In optical equipment, we expect elevated investments to sustain in the mid-to-long term, as Cloud providers catch up on prior under-investments from the last few years in wide area network (WAN) infrastructure.

The AI network fabric market is expected to reach ~$76bn by 2030, up from $18bn in 2025, across scale-up, scale-out, and scale-across networks. On the switching front, back-end AI networks have largely migrated towards enhanced flavors of Ethernet, carrying several advantages over InfiniBand, such as scalability, resilience, and interoperability. Dell'Oro forecasts that AI back-end Ethernet data center switch (including scale-out and scale-up network fabrics) revenue would grow over 10x in size in the next five years, from $11mn in 2025 to $157mn by 2030. Hyperscalers are a significant portion of the market opportunity, accounting for ~70% of AI data center back-end switching share.

In 2026 and beyond, AI workloads should have a multiplicative factor on bandwidth, with back-end AI networks expected to drive 20% growth of traffic in the front-end, requiring continued investments between both ends of the Ethernet network. The high-bandwidth requirements for AI networks are also driving the need for higher-gigabit port speeds for switching and optical equipment, such as 800 gigabit, 1.6 terabit, and 3.2 terabit, where demands call for the highest available packet per second transport.

We estimate a $63bn market size for networking equipment including both AI and traditional data center markets. Arista and Cisco Systems are the two largest vendors.

Exhibit 72: We estimate data center networking to be a $63bn market (2025)
Arista and Cisco Systems are the two largest vendors



Source: BofA Global Research

BofA GLOBAL RESEARCH

Special Disclosures

BofA Securities is currently acting as financial advisor to Flex Ltd in connection with its proposed spinoff of its Power and Cloud portfolio, which was announced on May 5, 2026. The proposed transaction is subject to approval by shareholders of Flex.This research report is not intended to provide voting advice, serve as an endorsement of the proposed transaction, or result in the procurement, withholding or revocation of a proxy.

BofA Securities is currently acting as joint financial advisor to ABB Ltd in connection with the sale of its Robotics Division to SoftBank Group Corp, which was announced on October 8, 2025.

Disclosures

Important Disclosures

FUNDAMENTAL EQUITY OPINION KEY: Opinions include a Volatility Risk Rating, an Investment Rating and an Income Rating. VOLATILITY RISK RATINGS, indicators of potential price fluctuation, are: A - Low, B - Medium and C - High. INVESTMENT RATINGS reflect the analyst's assessment of both a stock's absolute total return potential as well as its attractiveness for investment relative to other stocks within its Coverage Cluster (defined below). Our investment ratings are: 1 - Buy stocks are expected to have a total return of at least 10% and are the most attractive stocks in the coverage cluster; 2 - Neutral stocks are expected to remain flat or increase in value and are less attractive than Buy rated stocks and 3 - Underperform stocks are the least attractive stocks in a coverage cluster. An investment rating of 6 (No Rating) indicates that a stock is no longer trading on the basis of fundamentals. Analysts assign investment ratings considering, among other things, the 0-12 month total return expectation for a stock and the firm's guidelines for ratings dispersions (shown in the table below). The current price objective for a stock should be referenced to better understand the total return expectation at any given time. The price objective reflects the analyst's view of the potential price appreciation (depreciation).

Investment rating	Total return expectation (within 12-month period of date of initial rating)	Ratings dispersion guidelines for coverage cluster[R1]
Buy	≥ 10%	≤ 70%
Neutral	≥ 0%	≤ 30%
Underperform	N/A	≥ 20%

[R1]Ratings dispersions may vary from time to time where BofA Global Research believes it better reflects the investment prospects of stocks in a Coverage Cluster.

INCOME RATINGS, indicators of potential cash dividends, are: 7 - same/higher (dividend considered to be secure), 8 - same/lower (dividend not considered to be secure) and 9 - pays no cash dividend. *Coverage Cluster* is comprised of stocks covered by a single analyst or two or more analysts sharing a common industry, sector, region or other classification(s). A stock's coverage cluster is included in the most recent BofA Global Research report referencing the stock.

BofA Global Research personnel (including the analyst(s) responsible for this report) receive compensation based upon, among other factors, the overall profitability of Bank of America Corporation, including profits derived from investment banking. The analyst(s) responsible for this report may also receive compensation based upon, among other factors, the overall profitability of the Bank's sales and trading businesses relating to the class of securities or financial instruments for which such analyst is responsible.

Other Important Disclosures

From time to time research analysts conduct site visits of covered issuers. BofA Global Research policies prohibit research analysts from accepting payment or reimbursement for travel expenses from the issuer for such visits.
Prices are indicative and for information purposes only. Except as otherwise stated in the report, for any recommendation in relation to an equity security, the price referenced is the publicly traded price of the security as of close of business on the day prior to the date of the report or, if the report is published during intraday trading, the price referenced is indicative of the traded price as of the date and time of the report and in relation to a debt security (including equity preferred and CDS), prices are indicative as of the date and time of the report and are from various sources including BofA Securities trading desks.
The date and time of completion of the production of any recommendation in this report shall be the date and time of dissemination of this report as recorded in the report timestamp.

Recipients who are not institutional investors or market professionals should seek the advice of their independent financial advisor before considering information in this report in connection with any investment decision, or for a necessary explanation of its contents.
Officers of BofAS or one or more of its affiliates (other than research analysts) may have a financial interest in securities of the issuer(s) or in related investments.
Refer to BofA Global Research policies relating to conflicts of interest.
"BofA Securities" includes BofA Securities, Inc. ("BofAS") and its affiliates. Investors should contact their BofA Securities representative or Merrill Global Wealth Management financial advisor if they have questions concerning this report or concerning the appropriateness of any investment idea described herein for such investor. "BofA Securities" is a global brand for BofA Global Research.
Information relating to Non-US affiliates of BofA Securities and Distribution of Affiliate Research Reports:
BofAS and/or Merrill Lynch, Pierce, Fenner & Smith Incorporated ("MLPF&S") may in the future distribute, information of the following non-US affiliates in the US (short name: legal name, regulator): Merrill Lynch (South Africa): Merrill Lynch South Africa (Pty) Ltd., regulated by the Financial Sector Conduct Authority; MLI (UK): Merrill Lynch International, regulated by the Financial Conduct Authority (FCA) and the Prudential Regulation Authority (PRA); BofASE (France): BofA Securities Europe SA is authorized by the Autorité de Contrôle Prudentiel et de Résolution (ACPR) and regulated by the ACPR and the Autorité des Marchés Financiers (AMF). BofA Securities Europe SA ("BofASE") with registered address at 51, rue La Boétie, 75008 Paris is registered under no 842 602 690 RCS Paris. In accordance with the provisions of French Code Monétaire et Financier (Monetary and Financial Code), BofASE is an établissement de crédit et d'investissement (credit and investment institution) that is authorised and supervised by the European Central Bank and the Autorité de Contrôle Prudentiel et de Résolution (ACPR) and regulated by the ACPR and the Autorité des Marchés Financiers. BofASE's share capital can be found at www.bofaml.com/BofASEdisclaimer; BofA Europe (Milan): Bank of America Europe Designated Activity Company, Milan Branch, regulated by the Bank of Italy, the European Central Bank (ECB) and the Central Bank of Ireland (CBI); BofA Europe (Frankfurt): Bank of America Europe Designated Activity Company, Frankfurt Branch regulated by BaFin, the ECB and the CBI; BofA Europe (Zurich): Bank of America Europe Designated Activity Company, Zurich Branch, regulated by the Swiss Financial Market Supervisory Authority FINMA, the ECB and CBI; BofA Europe (Madrid): Bank of America Europe Designated Activity Company, Sucursal en España, regulated by the Bank of Spain, the ECB and the CBI; Merrill Lynch (Australia): Merrill Lynch Equities (Australia) Limited, regulated by the Australian Securities and Investments Commission; Merrill Lynch (Hong Kong): Merrill Lynch (Asia Pacific) Limited, regulated by the Hong Kong Securities and Futures Commission (HKSFC); Merrill Lynch (Singapore): Merrill Lynch (Singapore) Pte Ltd, regulated by the Monetary Authority of Singapore (MAS); Merrill Lynch (Canada): Merrill Lynch Canada Inc, regulated by the Canadian Investment Regulatory Organization; Merrill Lynch (Mexico): Merrill Lynch Mexico, SA de CV, Casa de Bolsa, regulated by the Comisión Nacional Bancaria y de Valores; BofAS Japan: BofA Securities Japan Co., Ltd., regulated by the Financial Services Agency; Merrill Lynch (Seoul): Merrill Lynch International, LLC Seoul Branch, regulated by the Financial Supervisory Service; Merrill Lynch (Taiwan): Merrill Lynch Securities (Taiwan) Ltd., regulated by the Securities and Futures Bureau; BofAS India: BofA Securities India Limited, regulated by the Securities and Exchange Board of India (SEBI); Merrill Lynch (Israel): Merrill Lynch Israel Limited, regulated by Israel Securities Authority; Merrill Lynch (DIFC): Merrill Lynch International (DIFC Branch), regulated by the Dubai Financial Services Authority (DFSA); Merrill Lynch (Brazil): Merrill Lynch S.A. Corretora de Títulos e Valores Mobiliários, regulated by Comissão de Valores Mobiliários; Merrill Lynch KSA Company: Merrill Lynch Kingdom of Saudi Arabia Company, regulated by the Capital Market Authority.
This information: has been approved for publication and is distributed in the United Kingdom (UK) to professional clients and eligible counterparties (as each is defined in the rules of the FCA and the PRA) by MLI (UK), which is authorized by the PRA and regulated by the FCA and the PRA - details about the extent of our regulation by the FCA and PRA are available from us on request; has been approved for publication and is distributed in the European Economic Area (EEA) by BofASE (France), which is authorized by the ACPR and regulated by the ACPR and the AMF; has been considered and distributed in Japan by BofAS Japan, a registered securities dealer under the Financial Instruments and Exchange Act in Japan, or its permitted affiliates; is issued and distributed in Hong Kong by Merrill Lynch (Hong Kong) which is regulated by HKSFC; is issued and distributed in Taiwan by Merrill Lynch (Taiwan); is issued and distributed in India by BofAS India; and is issued and distributed in Singapore to institutional investors and/or accredited investors (each as defined under the Financial Advisers Regulations) by Merrill Lynch (Singapore) (Company Registration No 198602883D). Merrill Lynch (Singapore) is regulated by MAS. Merrill Lynch Equities (Australia) Limited (ABN 65 006 276 795), AFS License 235132 (MLEA) distributes this information in Australia only to 'Wholesale' clients as defined by s.761G of the Corporations Act 2001. With the exception of Bank of America N.A., Australia Branch, neither MLEA nor any of

its affiliates involved in preparing this information is an Authorised Deposit-Taking Institution under the Banking Act 1959 nor regulated by the Australian Prudential Regulation Authority. No approval is required for publication or distribution of this information in Brazil and its local distribution is by Merrill Lynch (Brazil) in accordance with applicable regulations. Merrill Lynch (DIFC) is authorized and regulated by the DFSA. Information prepared and issued by Merrill Lynch (DIFC) is done so in accordance with the requirements of the DFSA conduct of business rules. BofA Europe (Frankfurt) distributes this information in Germany and is regulated by BaFin, the ECB and the CBI. BofA Securities entities, including BofA Europe and BofASE (France), may outsource/delegate the marketing and/or provision of certain research services or aspects of research services to other branches or members of the BofA Securities group. You may be contacted by a different BofA Securities entity acting for and on behalf of your service provider where permitted by applicable law. This does not change your service provider. Please refer to the Electronic Communications Disclaimers for further information.

This information has been prepared and issued by BofAS and/or one or more of its non-US affiliates. The author(s) of this information may not be licensed to carry on regulated activities in your jurisdiction and, if not licensed, do not hold themselves out as being able to do so. BofAS and/or MLPF&S is the distributor of this information in the US and accepts full responsibility for information distributed to BofAS and/or MLPF&S clients in the US by its non-US affiliates. Any US person receiving this information and wishing to effect any transaction in any security discussed herein should do so through BofAS and/or MLPF&S and not such foreign affiliates. Hong Kong recipients of this information should contact Merrill Lynch (Asia Pacific) Limited in respect of any matters relating to dealing in securities or provision of specific advice on securities or any other matters arising from, or in connection with, this information. Singapore recipients of this information should contact Merrill Lynch (Singapore) Pte Ltd in respect of any matters arising from, or in connection with, this information. For clients that are not accredited investors, expert investors or institutional investors Merrill Lynch (Singapore) Pte Ltd accepts full responsibility for the contents of this information distributed to such clients in Singapore.

General Investment Related Disclosures:

Taiwan Readers: Neither the information nor any opinion expressed herein constitutes an offer or a solicitation of an offer to transact in any securities or other financial instrument. No part of this report may be used or reproduced or quoted in any manner whatsoever in Taiwan by the press or any other person without the express written consent of BofA Securities.

This document provides general information only, and has been prepared for, and is intended for general distribution to, BofA Securities clients. Neither the information nor any opinion expressed constitutes an offer or an invitation to make an offer, to buy or sell any securities or other financial instrument or any derivative related to such securities or instruments (e.g., options, futures, warrants, and contracts for differences). This document is not intended to provide personal investment advice and it does not take into account the specific investment objectives, financial situation and the particular needs of, and is not directed to, any specific person(s). This document and its content do not constitute, and should not be considered to constitute, investment advice for purposes of ERISA, the US tax code, the Investment Advisers Act or otherwise. Investors should seek financial advice regarding the appropriateness of investing in financial instruments and implementing investment strategies discussed or recommended in this document and should understand that statements regarding future prospects may not be realized. Any decision to purchase or subscribe for securities in any offering must be based solely on existing public information on such security or the information in the prospectus or other offering document issued in connection with such offering, and not on this document.

Securities and other financial instruments referred to herein, or recommended, offered or sold by BofA Securities, are not insured by the Federal Deposit Insurance Corporation and are not deposits or other obligations of any insured depository institution (including, Bank of America, N.A.). Investments in general and, derivatives, in particular, involve numerous risks, including, among others, market risk, counterparty default risk and liquidity risk. No security, financial instrument or derivative is suitable for all investors. Digital assets are extremely speculative, volatile and are not regulated in a manner similar to U.S. securities markets. In some cases, securities and other financial instruments may be difficult to value or sell and reliable information about the value or risks related to the security or financial instrument may be difficult to obtain. Investors should note that income from such securities and other financial instruments, if any, may fluctuate and that price or value of such securities and instruments may rise or fall and, in some cases, investors may lose their entire principal investment. Past performance is not necessarily a guide to future performance. Levels and basis for taxation may change.

This report may contain a short-term trading idea or recommendation, which highlights a specific near-term catalyst or event impacting the issuer or the market that is anticipated to have a short-term price impact on the equity securities of the issuer. Short-term trading ideas and recommendations are different from and do not affect a stock's fundamental equity rating, which reflects both a longer term total return expectation and attractiveness for investment relative to other stocks within its Coverage Cluster. Short-term trading ideas and recommendations may be more or less positive than a stock's fundamental equity rating.

BofA Securities is aware that the implementation of the ideas expressed in this report may depend upon an investor's ability to "short" securities or other financial instruments and that such action may be limited by regulations prohibiting or restricting "shortselling" in many jurisdictions. Investors are urged to seek advice regarding the applicability of such regulations prior to executing any short idea contained in this report.

Foreign currency rates of exchange may adversely affect the value, price or income of any security or financial instrument mentioned herein. Investors in such securities and instruments, including ADRs, effectively assume currency risk.

BofAS or one of its affiliates is a regular issuer of traded financial instruments linked to securities or digital assets that may have been recommended or discussed in this report. BofAS or one of its affiliates may, at any time, hold a trading position (long or short) in the securities and financial instruments recommended or discussed in this report.

BofA Securities, through business units other than BofA Global Research, may have issued and may in the future issue trading ideas or recommendations that are inconsistent with, and reach different conclusions from, the information presented herein. Such ideas or recommendations may reflect different time frames, assumptions, views and analytical methods of the persons who prepared them, and BofA Securities is under no obligation to ensure that such other trading ideas or recommendations are brought to the attention of any recipient of this information.

In the event that the recipient received this information pursuant to a contract between the recipient and BofAS for the provision of research services for a separate fee, and in connection therewith BofAS may be deemed to be acting as an investment adviser, such status relates, if at all, solely to the person with whom BofAS has contracted directly and does not extend beyond the delivery of this report (unless otherwise agreed specifically in writing by BofAS). If such recipient uses the services of BofAS in connection with the sale or purchase of a security referred to herein, BofAS may act as principal for its own account or as agent for another person. BofAS is and continues to act solely as a broker-dealer in connection with the execution of any transactions, including transactions in any securities referred to herein.

Copyright and General Information:

Copyright 2026 Bank of America Corporation. All rights reserved. iQdatabase® is a registered service mark of Bank of America Corporation. This information is prepared for the use of BofA Securities clients and may not be redistributed, retransmitted or disclosed, in whole or in part, or in any form or manner, without the express written consent of BofA Securities. This document and its content is provided solely for informational purposes and cannot be used for training or developing artificial intelligence (AI) models or as an input in any AI application (collectively, an AI tool). Any attempt to utilize this document or any of its content in connection with an AI tool without explicit written permission from BofA Global Research is strictly prohibited. BofA Global Research utilizes AI, including machine learning and other technologies, to enhance the services we provide to our clients. These technologies assist our analysts in various aspects of their work, including but not limited to data analysis, content extraction, content creation, data aggregation and summarization and identifying relevant information from diverse sources. All AI-driven processes are subject to review by BofA Global Research employees. BofA Global Research information is distributed simultaneously to internal and client websites and other portals by BofA Securities and is not publicly-available material. Any unauthorized use or disclosure is prohibited. Receipt and review of this information constitutes your agreement not to redistribute, retransmit, or disclose to others the contents, opinions, conclusion, or information contained herein (including any investment recommendations, estimates or price targets) without first obtaining express permission from an authorized officer of BofA Securities.

Materials prepared by BofA Global Research personnel are based on public information. Facts and views presented in this material have not been reviewed by, and may not reflect information known to, professionals in other business areas of BofA Securities, including investment banking personnel. BofA Securities has established information barriers between BofA Global Research and certain business groups. As a result, BofA Securities does not disclose certain client relationships with, or compensation received from, such issuers. To the extent this material discusses any legal proceeding or issues, it has not been prepared as nor is it intended to express any legal conclusion, opinion or advice. Investors should consult their own legal advisers as to issues of law relating to the subject matter of this material. BofA Global Research personnel's knowledge of legal proceedings in which any BofA Securities entity and/or its directors, officers and employees may be plaintiffs, defendants, co-defendants or co-plaintiffs with or involving issuers mentioned in this material is based on public information. Facts and views presented in this material that relate to any such proceedings have not been reviewed by, discussed with, and may not reflect information known to, professionals in other business areas of BofA Securities in connection with the legal proceedings or matters relevant to such proceedings.

This information has been prepared independently of any issuer of securities mentioned herein and not in connection with any proposed offering of securities or as agent of any issuer of any securities. None of BofAS any of its affiliates or their research analysts has any authority whatsoever to make any representation or warranty on behalf of the issuer(s). BofA Global Research policy prohibits research personnel from disclosing a recommendation, investment rating, or investment thesis for review by an issuer prior to the publication of a research report containing such rating, recommendation or investment thesis.

Any information relating to sustainability in this material is limited as discussed herein and is not intended to provide a comprehensive view on any sustainability claim with respect to any issuer or security.

Any information relating to the tax status of financial instruments discussed herein is not intended to provide tax advice or to be used by anyone to provide tax advice. Investors are urged to seek tax advice based on their particular circumstances from an independent tax professional.
The information herein (other than disclosure information relating to BofA Securities and its affiliates) was obtained from various sources and we do not guarantee its accuracy. This information may contain links to third-party websites. BofA Securities is not responsible for the content of any third-party website or any linked content contained in a third-party website. Content contained on such third-party websites is not part of this information and is not incorporated by reference. The inclusion of a link does not imply any endorsement by or any affiliation with BofA Securities. Access to any third-party website is at your own risk, and you should always review the terms and privacy policies at third-party websites before submitting any personal information to them. BofA Securities is not responsible for such terms and privacy policies and expressly disclaims any liability for them.
All opinions, projections and estimates constitute the judgment of the author as of the date of publication and are subject to change without notice. Prices also are subject to change without notice. BofA Securities is under no obligation to update this information and BofA Securities ability to publish information on the subject issuer(s) in the future is subject to applicable quiet periods. You should therefore assume that BofA Securities will not update any fact, circumstance or opinion contained herein.
Certain outstanding reports or investment opinions relating to securities, financial instruments and/or issuers may no longer be current. Always refer to the most recent research report relating to an issuer prior to making an investment decision.
In some cases, an issuer may be classified as Restricted or may be Under Review or Extended Review. In each case, investors should consider any investment opinion relating to such issuer (or its security and/or financial instruments) to be suspended or withdrawn and should not rely on the analyses and investment opinion(s) pertaining to such issuer (or its securities and/or financial instruments) nor should the analyses or opinion(s) be considered a solicitation of any kind. Sales persons and financial advisors affiliated with BofAS or any of its affiliates may not solicit purchases of securities or financial instruments that are Restricted or Under Review and may only solicit securities under Extended Review in accordance with firm policies.
Neither BofA Securities nor any officer or employee of BofA Securities accepts any liability whatsoever for any direct, indirect or consequential damages or losses arising from any use of this information.