BERNSTEIN
SOCIETE GENERALE GROUP

1 October 2026

US Communications Infrastructure

U.S. Data Center Capacity: Your seat is skewing your numbers. Five ways to forecast DC capacity (and our definitive view)


Madison Rezaei
+1 917 344 8622
madison.rezaei@bernsteinsg.com


David Dai, CFA
+852 2918 5704
david.dai@bernsteinsg.com


Mark L. Moerdler, Ph.D.
+1 917 344 8506
mark.moerdler@bernsteinsg.com


Mark Shmulik
+1 917 344 8508
mark.shmulik@bernsteinsg.com


Chad Dillard
+1 917 344 8469
chad.dillard@bernsteinsg.com


Sunaina Ocalan
.+1 917 344 8503
sunaina.ocalan@bernsteinsg.com


Alasdair Leslie
+44 20 7762 4952
alasdair.leslie@bernsteinsg.com


Varun Govindaraj
+1 917 344 8543
varun.govindaraj@bernsteinsg.com


Nancy Wu
+1 917 344 8545
nancy.wu@bernsteinsg.com

By far the most frequent question that hits our collective inboxes is "how much data center capacity are we really building? How fast?" Well, depending on who you ask, you're going to get a wide range of answers. Across subsectors (and even within a single one), there's no consensus. A survey of 63 available forecasts from 40 sources yielded a spread of 59 to 186 GW by 2030. Today's note looks at five different ways to forecast U.S. data center capacity, points out who likes to use what, unpacks the pros and cons of each, benchmarks credible comps, and articulates the definitive Bernstein range.

The **infrastructure / REIT people** have the most history of the data center forecasters - they tend to model on historical trends and site censuses, with a healthy dose of skepticism mixed in and consequently generally skew conservative. Our builder credibility analysis on the U.S. pipeline is an example of this. Using that as a jumping off point, we forecast a range of 135-167 IT GW in the U.S. by 2030, with a base case of **142 IT-GW**.

Our industrial and equipment friends are all about the incremental additions, usually on a facility or nameplate GW basis (a higher number than IT-GW). They also use a lot of the site-based methods. They are also closely watching the second derivative. Equipment orders can front run some of the IT installed capacity, so this camp's numbers often skew higher. While there are several creative approaches within this group, we've done a growth rate on incremental adds. This approach yields 124-242 IT-GW by 2030 (base **161 IT-GW**).

The tech camp has a few sub-groups. You have your **hyperscale-anchored forecasts**, based on capex estimates, which is a market favorite across many sectors. Hyperscalers represent the majority of capacity across self-builds and colo. However, getting to GWs from capex requires some broad assumptions that are hard to validate - we end up at a base of **152 IT-GW** in 2030 with this approach, but with a very wide range of 115-195.

Then you have the **semis and hardware-based forecasts**, tricky to align for both timing and unknown refresh cycles. For this one, we leaned heavily on David Dai's in-depth AI server power model. Adjusting for U.S. and with some other assumptions, we approach **149 IT-GW by 2030** (range 137-160).

Finally, we have some creative collabs and gating factors, spearheaded by Chad Dillard's craft labor ceiling analysis. Put into apples-to-apples terms with the rest of our forecasts, it is indeed a potential ceiling - **142 GW by 2030** (range 120-204).

We take all of these forecasts into account for our definitive Bernstein estimates: 149 IT-GW by 2030, adding 19 IT-GW (~25 facility GW) in 2027 and 25 IT-GW (~32 facility) in 2028. Model here to tweak the assumptions for yourself.

See the Disclosure Appendix of this report for required disclosures, analyst certifications and other important information. Alternatively, visit our Global Research Disclosure Website.

First Published: 02 Oct 2026 02:47 UTC Completion Date: 02 Oct 2026 02:47 UTC

BERNSTEIN TICKER TABLE

Ticker	Rating	Cur	1 Oct 2026 Closing Price	Price Target	TTM Rel. Perf.	Cur	Adjusted EPS 2025A	Adjusted EPS 2026E	Adjusted EPS 2027E	Adjusted P/E (x) 2025A	Adjusted P/E (x) 2026E	Adjusted P/E (x) 2027E
BXDC (Blackstone Digital Infra)	O	USD	17.28	26.00	NA	USD	0.00	0.32	(1.00)		54.2	(17.2)
CRWV (CoreWeave)	U	USD	88.57	74.00	(49.6)%	USD	(1.20)	(4.56)	(4.77)	N/M	304	38.6
CSQR (Csquare)	O	USD	16.52	27.00	NA	USD	0.00	(1.34)	(1.06)		(12.3)	(15.5)
DLR (Digital Realty)	O	USD	176.49	226.00	(11.3)%	USD	3.87	3.09	3.21	45.6	57.0	55.0
EQIX (Equinix)	O	USD	1,012.33	1,270.00	16.6%	USD	14.96	17.79	20.87	67.7	56.9	48.5
VRT (Vertiv)	O	USD	246.12	368.00	38.1%	USD	4.20	6.60	9.06	58.6	37.3	27.2
NVT (nVent)	O	USD	165.77	229.00	51.3%	USD	3.34	5.15	6.55	49.6	32.2	25.3
CARR (Carrier)	M	USD	55.22	78.00	(21.4)%	USD	2.57	2.85	3.29	21.5	19.4	16.8
JCI (Johnson Controls)	O	USD	151.77	190.00	25.1%	USD	3.78	5.09	6.35	40.2	29.8	23.9
TEL (TE Connectivity)	O	USD	217.76	243.00	(16.0)%	USD	9.26	11.46	12.98	23.5	19.0	16.8
APH (Amphenol)	O	USD	85.67	99.00	23.2%	USD	1.67	2.65	3.31	51.3	32.3	25.9
TT (Trane)	O	USD	458.94	575.00	(6.4)%	USD	13.06	15.28	17.79	35.1	30.0	25.8
MSFT (Microsoft)	O	USD	512.80	660.00	(15.6)%	USD	17.28	19.96	24.01	29.7	25.7	21.4
ORCL (Oracle)	O	USD	138.07	325.00	(66.5)%	USD	7.63	8.37	12.19	18.1	16.5	11.3
ABBN.SW (ABB)	M	CHF	79.62	80.00	28.2%	USD	2.56	2.99	3.60	37.5	32.1	26.7
LR.FP (Legrand)	O	EUR	139.75	170.00	(10.2)%	EUR	5.22	6.13	6.86	29.4	24.7	21.9
SU.FP (Schneider)	O	EUR	292.60	350.00	9.2%	EUR	8.43	10.42	13.24	34.7	28.1	22.1
SIE.GR (Siemens)	O	EUR	273.20	330.00	6.5%	EUR	12.23	11.29	11.64	27.4	23.6	22.2
ENR.GR (Siemens Energy)	O	EUR	143.78	210.00	27.8%	EUR	1.63	4.24	5.94	34.6	19.2	13.8
AMZN (Amazon)	O	USD	248.23	320.00	(1.7)%	USD	7.17	12.65	12.13	34.6	19.6	20.5
GOOGL (Alphabet)	M	USD	338.24	385.00	23.9%	USD	10.81	20.84	15.26	31.3	16.2	22.2
META (Meta)	O	USD	725.93	800.00	(13.0)%	USD	23.49	32.27	36.68	30.9	22.5	19.8
ASML (ASML)	O	USD	1,808.49	2,859.00	66.0%	USD	27.95	44.01	60.58	57.5	36.5	26.5
ASML.NA (ASML)	O	EUR	1,597.00	2,500.00	79.5%	EUR	24.72	38.91	53.56	64.6	41.0	29.8
BESI.NA (Besi)	O	EUR	191.40	320.00	39.5%	EUR	1.66	4.28	6.60	115.3	44.7	29.0
IFX.GR (Infineon)	O	EUR	59.83	102.00	68.0%	EUR	1.38	1.77	2.98	43.2	33.8	20.1
6146.JP (DISCO)	O	JPY	60,440	99,000	(18.3)%	JPY	1,245.79	1,830.28	2,336.94	48.5	33.0	25.9
8035.JP (Tokyo Electron)	O	JPY	12,580	79,300	89.8%	JPY	1,250.88	1,687.17	2,635.81	10.1	7.5	4.8
6857.JP (Advantest)	O	JPY	37,610	45,800	113.9%	JPY	513.30	948.96	1,088.87	73.3	39.6	34.5
6723.JP (Renesas)	O	JPY	3,616.00	6,300.00	78.0%	JPY	181.61	261.36	302.64	19.9	13.8	11.9
VST (Vistra Corporation)	O	USD	139.75	181.00	(44.9)%	USD	2.18	10.39	11.45	64.2	13.4	12.2
CEG (Constellation Energy)	O	USD	258.92	296.00	(40.4)%	USD	7.40	11.82	12.67	35.0	21.9	20.4
BE (Bloom Energy)	M	USD	277.58	282.00	193.5%	USD	0.82	2.85	4.81	336.9	97.4	57.8
NEE (NextEra Energy)	O	USD	76.35	108.00	(17.2)%	USD	3.71	4.31	4.49	20.6	17.7	17.0
SPX			7,666.45									
EDME			1,554.34									
JPL			2,698.24									

O - Outperform, M - Market-Perform, U - Underperform, NR - Not Rated, CS - Coverage Suspended

SIE.GR, ENR.GR, AMZN, GOOGL, META, ASML, ASML.NA, BESI.NA, 6146.JP, 8035.JP, 6857.JP, VST, CEG estimate is Reported EPS; CRWV valuation is EV/EBIT (x); LR.FP, AMZN, GOOGL, META, ASML, ASML.NA, BESI.NA, 6146.JP, 8035.JP, 6857.JP, VST, CEG valuation is Reported P/E (x); ENR.GR valuation is EV/EBITDA (x); MSFT, ORCL base year is 2026;

Source: Bloomberg, Bernstein estimates and analysis.

INVESTMENT IMPLICATIONS

We rate VST, NEE, CEG Outperform, TP $181, $296 and $108 respectively

We rate BE Market-Perform, TP $282

We rate VRT Outperform, TP $ 368

We rate NVT Outperform, TP $ 229

We rate CARR Market-Perform, TP $ 78

We rate JCI Outperform, TP $ 190

We rate TT Outperform, TP $ 575

We rate APH Outperform, TP $ 99

We rate TEL Outperform, TP $ 243

We rate MSFT Outperform TP $660

We rate ORCL Outperform TP $325

We rate SU Outperform, TP € 350

We rate ABB Market-Perform, TP CHF 80

We rate LR Outperform, TP € 170

We rate SIE Outperform, TP € 330

We rate ENR Outperform, TP € 210

We rate META Outperform, TP $800

We rate AMZN Outperform, TP $320

We rate GOOGL Market-Perform, TP $385

Table Of Contents

PM Summary.....4
housekeeping items: why the forecast range is so wide.....5
Site-Based Estimates (the Infrastructure Seat).....6
The INCREMENTAL GROWTH RATE Approach (Popular in Industrials).....9
The Hyperscaler-Capex Method (Technology and Internet Teams).....12
The Server-BASED Method (Semis View).....15
The Craft labor view.....18

DETAILS

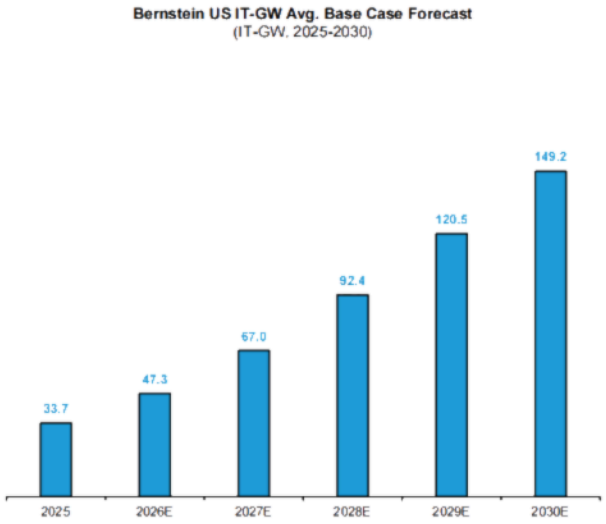
PM SUMMARY

The biggest debate in the data center infrastructure space is simple, but surprisingly complex to answer: **how much capacity is the U.S. actually building?** Depending on the methodology, forecasts for U.S. installed data center capacity in 2030 range from roughly 60 GW to more than 180 GW, with little consensus even among credible industry observers. After reviewing five forecasting frameworks across the infrastructure, technology, industrial, and labor ecosystems, we believe the most reasonable outcome is ~160 IT-GW of installed U.S. capacity by 2030, implying capacity growth of more than 4.5x from 2025 levels.

Our note today highlights that the disagreement is not necessarily about end demand, but stems from differences in starting assumptions, measurement methodologies, timing conventions, and where participants sit within the value chain. Infrastructure investors tend to focus on commissioned facilities and historical deployment trends. Technology analysts often extrapolate from hyperscaler spending plans or hardware shipments. Industrials analysts track orders for electrical and thermal equipment that appear earlier in the construction cycle. Labor and permitting analyses attempt to identify practical ceilings on what can physically be built. Each framework captures a valid piece of the puzzle, but none offers a complete picture on its own.

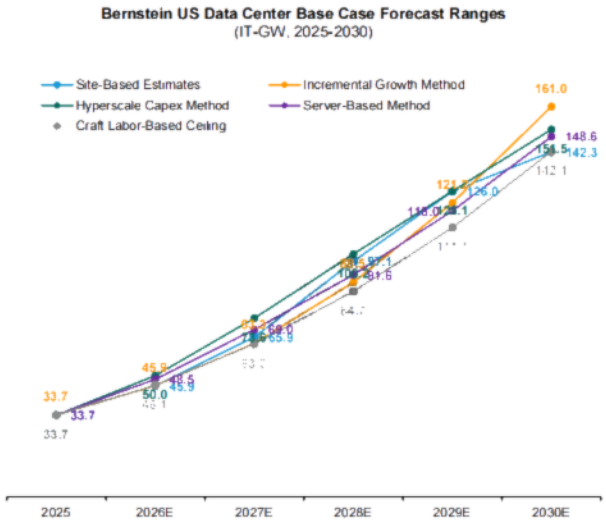
From a data center investment perspective, we continue to favor companies with direct exposure to energized capacity growth rather than whose outlook depends on the current pace of AI adoption. Today's market remains supply constrained, with demand for AI infrastructure exceeding available capacity in many regions. As a result, ownership of, or exposure to, energized capacity remains particularly valuable, as new supply is often absorbed shortly after delivery. If our base case proves correct, data center developers, colocation providers, electrical equipment suppliers, and other infrastructure beneficiaries should continue to benefit from strong demand and multi-year revenue visibility. Conversely, the greatest risk appears concentrated in situations where valuations already assume the most aggressive capacity outcomes materialize and without execution challenges or delays. Ultimately, the range of possible outcomes remains wide, but should support a durable infrastructure build cycle and reinforce our constructive view on the broader AI infrastructure ecosystem.

EXHIBIT 1: **Bernstein US IT-GW Avg. Base Case Forecast (IT-GW, 2025-2030)**



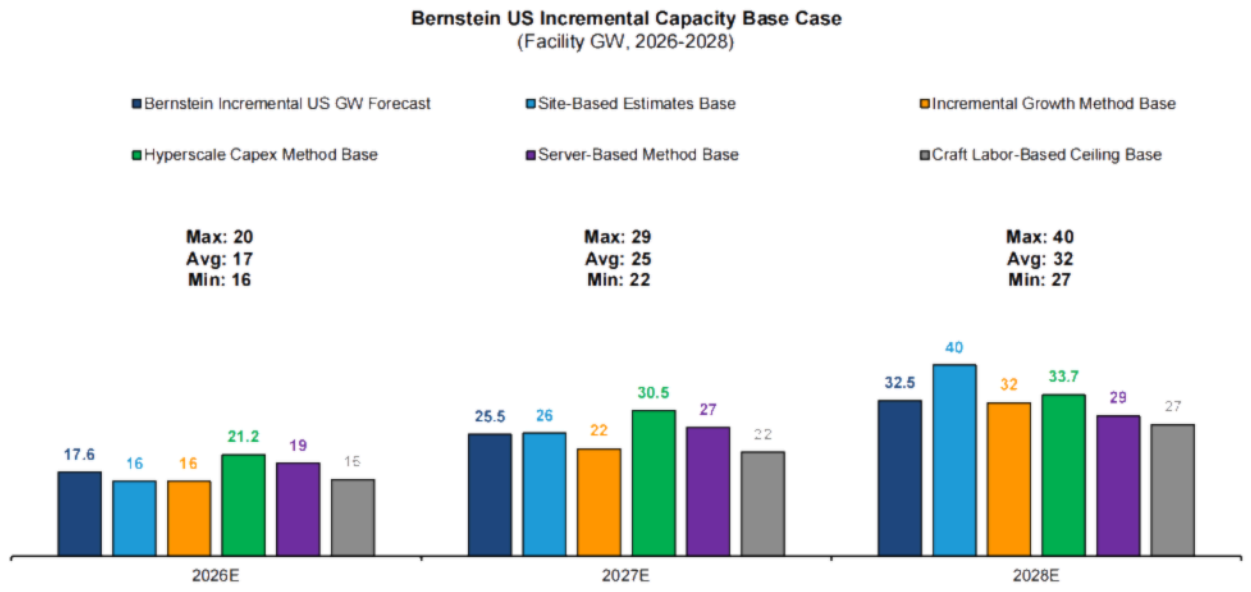
Source: Bloomberg, Company filings and disclosures, Aterio, Bernstein Analysis and Estimates

EXHIBIT 2: **Bernstein US Data Center Base Case Forecast Ranges (IT-GW, 2025-2030)**



Source: Bloomberg, Company filings and disclosures, Aterio, Bernstein Analysis and Estimates

EXHIBIT 3: **Bernstein US Incremental Capacity by Methodology Base Case(Facility GW, 2026-2028)**



Source: Bloomberg, Company filings and disclosures, Aterio, Bernstein Analysis and Estimates

HOUSEKEEPING ITEMS: WHY THE FORECAST RANGE IS SO WIDE

Some important context before we dive into the details here. Forecasting GW constructed by 2030 is hairy: there are supply questions, demand questions, definitional questions, and a ton of noise.

As a market, we absolutely do not agree on a starting point. Credible estimates put capacity between 30 and 80 GW by the end of 2025...a number that was theoretically known ten months ago. But the market is a long tail of private players, and large players (including the hyperscalers who give no meaningful data on existing) so data is sketchy and there's not much to be done. Depending on where a forecast has anchored and its starting point for last year's build, the numbers can skew meaningfully. For our purposes, we have anchored our forecasts to a single 2025 number (35 IT-GW), based on our project pipeline data,

counting only active data centers as of end of 2025 in our data set. **We think this starting point should perhaps be closer to 45GW, but as we deem the growth to be the more meaningful part of this analysis, we have elected to maintain the source historicals over grossing up in a way that would color the incremental numbers.**

There's a ton of definitional noise. A GW is not just a GW - there's installed critical IT capacity (IT-GW, inclusive of a PUE haircut, predominantly used by tech and data center people and how we mostly model in this note), facility GWs (often called nameplate, what is used for electricals, industrials, and power). To convert between the two, we are using a simplified annual Power Use Efficiency assumption of 1.3 in 2026 (trending gently down).

Timing is imperfect. Not only are there major differences in the units themselves, there are also different timing implications based on who in the value chain you're talking to. Facilities have to have power permitted first, then they build (orders for things like cooling systems (Varun Govindaraj), and various power equipment (Chad Dillard, Alasdair Leslie), etc. all frontrun the actual data center activation date. Once the data center is fully built and energized (this would be our "activation date"), hardware goes in, which depending on the vendor can be weeks to months .

Hardware purchases officially hit differently depending again on where you are in the value chain. It could start when the chips are ordered, or when the servers/storage/networking is ordered; or even when it 1st ships to the datacenter[1]. The difference in timing could be weeks to months and also depending on the end datacenter provider (e.g. hyperscaler or neocloud) could be purchased (e.g. chips or servers) timed for delivery when the facility is energized. Note that large vendors including Microsoft and Oracle have hardware always "under construction" given how frequently they are taking ownership of datacenters. Thus the timing can also be off. And then finally, colocation facilities aren't fully "stabilized" for years after that! **In this model, we're counting IT-GW as when a facility is fully energized and ready for hardware, even if it's not fully outfitted yet.**

We are at an inflection point without a clear historical precedent - every year is the biggest year yet. On any metric, 2026 has exceeded any prior year trends, with 2027 and 2028 projects under construction to each continue being the biggest year in history. Data center growth used to follow relatively reliable cycles and trends - those have been thrown out the window given insane demand growth (AI), the proliferation of new providers, geographies, breaking of power bottlenecks, and accelerated build strategies.

Almost unilaterally, forecasts have been revised upward with every update. We looked at publishers year over year and found massive revisions. McKinsey bumped their 2030 forecast by 258% between 2025 and 2026. BCG raised theirs 155%, and EPRI took theirs up 99%. We published a detailed, thoroughly-researched global supply and demand model at the beginning of this year, and candidly, the numbers have ballooned since. For full disclosure: we also changed our sourcing and our research methodology and believe our current outlook to be more reflective of the market.

Why we believe in our range: we have looked at this from the important parts of the value chain and come up with reasonable, consistent assumptions and numbers that tie together. We also have some benefit of time and detailed data in the near term: it takes *at least* 18 months to build a data center facility. More typically this is a 3+ year planning process, only extended with increasing legislation (TX, PA, NY). We track the pipeline closely with monthly updates and have strong views on the credibility of those builds.

Some datacenters are going to delay for regulatory, construction, or business reasons and others will simply never convert from an initial announcement to an operational building for an even wider range of reasons. The net effect is not everything at the top of the funnel will materialize and some believe that at least the hyperscalers are over ordering and building in a natural failure rate in data construction projects (Global Software: What is in a Datacenter contract? Key takeaways from discussion with a legal expert.)

Over the longer term, our range is necessarily wider - we certainly have opinions about what happens by 2030, but don't have a crystal ball - a lot could change between now and then to shift the totals.

SITE-BASED ESTIMATES (THE INFRASTRUCTURE SEAT)

The site-based methodology starts with a simple premise: count the data centers themselves. This approach builds a bottom-up inventory of operating facilities and development projects, then estimates future capacity based on what is already

[1] For example Microsoft and Oracle purchase components and build their own hardware while some neoclouds order hardware from the server, networking and storage vendors.

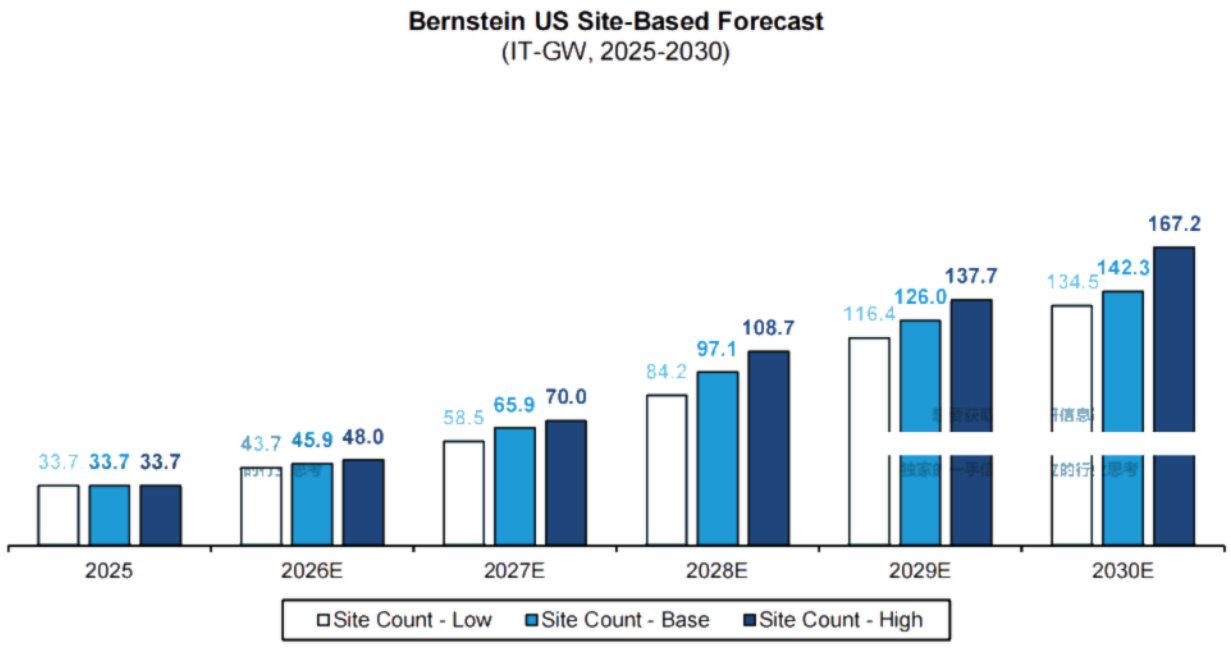
active, under construction, planned, or announced. In our framework, we use the Aterio project database to establish the current installed base and development pipeline, then apply completion assumptions and construction delays to estimate future annual capacity additions. We convert facility power into IT load using our PUE assumptions and generate low, base, and high cases based on project realization timing.

The biggest advantage of the site-based approach is that it is the closest thing to ground truth available. Every data center must eventually exist somewhere geographically / physically, obtain permits, secure land and power, and move through the construction process. As a result, this methodology provides a direct view into the actual development pipeline rather than relying on indirect proxies. It also offers visibility into timing, since projects have construction schedules and expected CODs. Because it starts with individual projects, it is particularly useful for understanding near-term capacity additions and identifying where forecasts may be overly optimistic about construction timelines. In our view, this is the most actionable framework for forecasting the next several years of energized capacity.

The primary limitation is that the method is only as good as the project database underlying it. No project tracker captures every development, definitions of "planned" and "under construction" vary, and pipelines can quickly become stale. While site-level data works very well for the next couple of years, visibility deteriorates meaningfully at longer durations because projects can be delayed, resized, canceled, or replaced by developments that are not yet in the database. Put differently, beyond several years the methodology increasingly reflects what developers currently intend to build rather than what will actually be operational. Forward forecasts therefore require assumptions around conversion rates, delays, and project credibility. Additionally, estimating the pipeline credibility is more art than science - our analysis puts credible GWs at ~35% of total pipeline across all stages (announced to nearly-final), but it's hard to be more exact until construction has physically started.

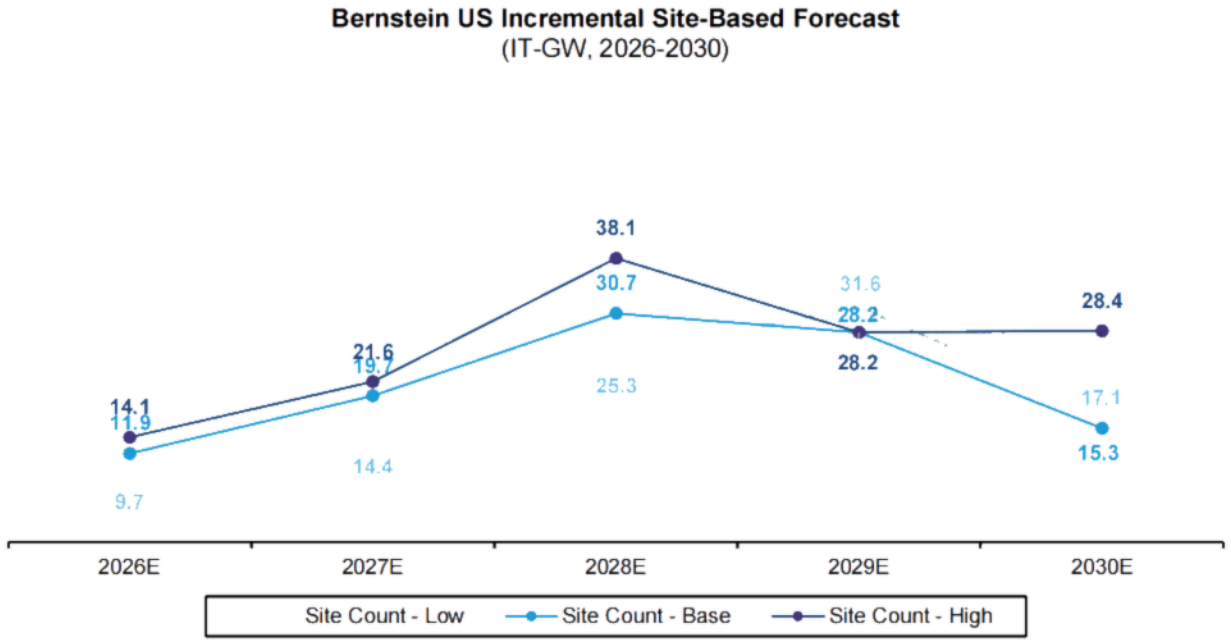
Who uses it? This tends to be the preferred approach of infrastructure investors, data center REIT analysts, and colocation operators, who focus on physical assets and development pipelines. These investors (along with the Communications Infra team here at Bernstein!) generally place greater weight on what is already under construction or has secured a development path, and therefore often arrive at more conservative forecasts than technology- or semiconductor-oriented frameworks. Within our analysis, the site-based methodology produces a range of ~134-167 IT-GW by 2030, with a base case of approximately 142 IT-GW, making it one of the more asset-focused and credibility-based approaches in the forecast universe.

EXHIBIT 4: **Bernstein US Site-Based Forecast (IT-GW, 2025-2030)**



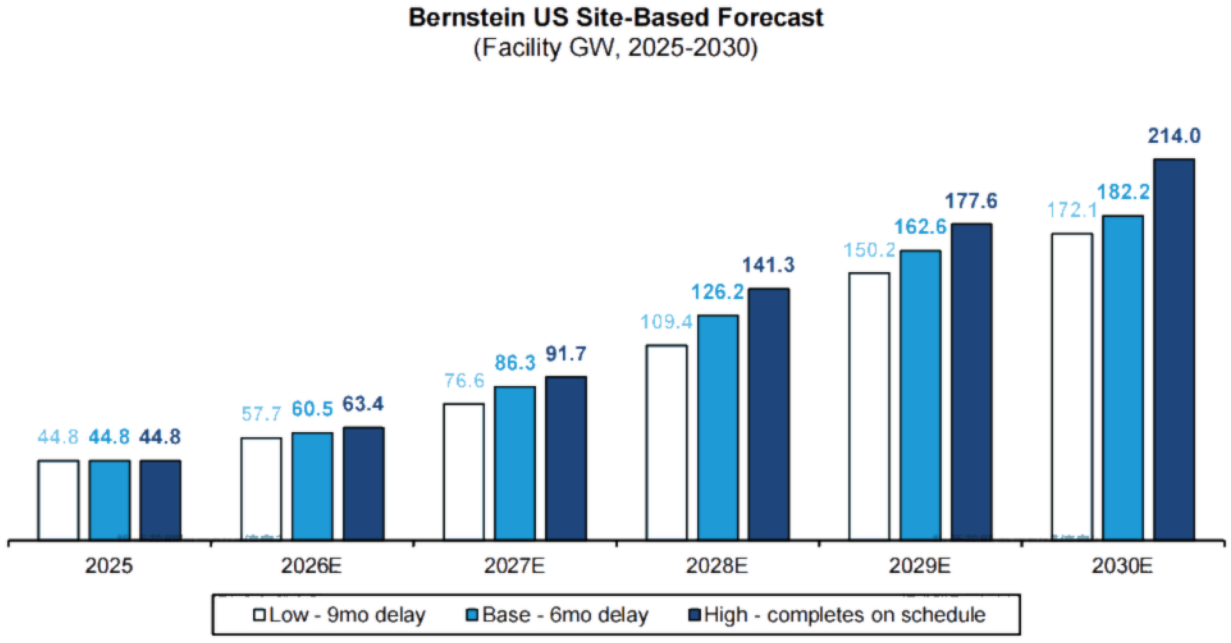
Source: Aterio, Bernstein Analysis and Estimates

EXHIBIT 5: **Bernstein US Incremental Site-Based Forecast (IT-GW, 2026-2030)**



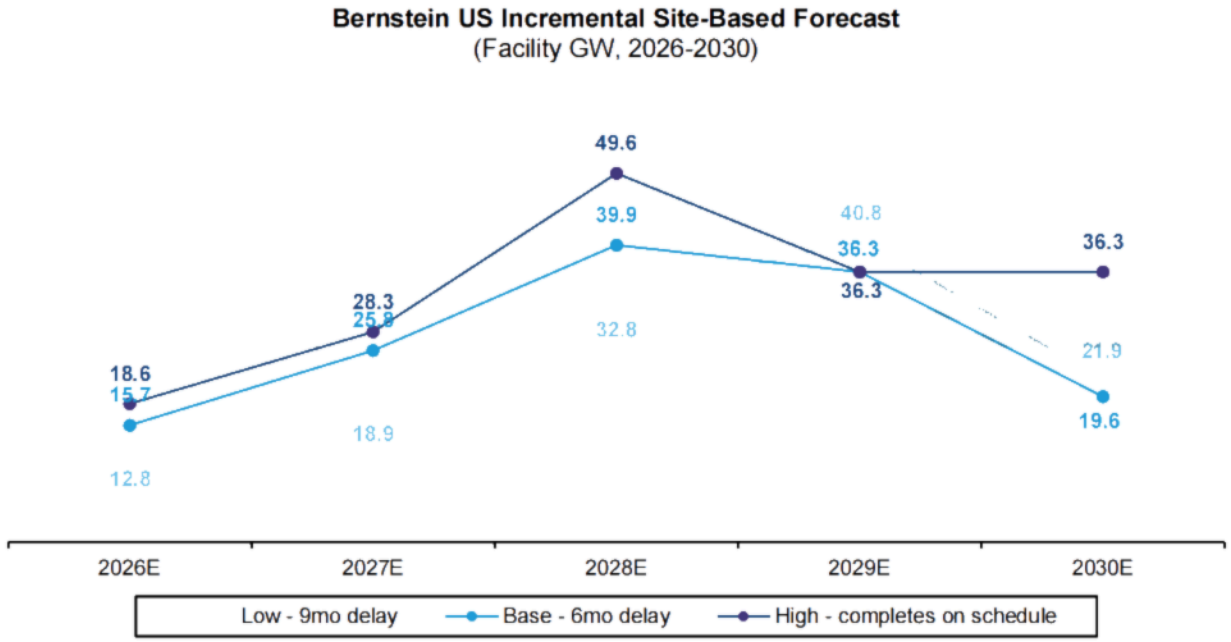
Source: Aterio, Bernstein Analysis and Estimates

EXHIBIT 6: **Bernstein US Site-Based Forecast (Facility GW, 2025-2030)**



Source: Aterio, Bernstein Analysis and Estimates

EXHIBIT 7: **Bernstein US Incremental Site-Based Forecast (Facility GW, 2026-2030)**



Source: Aterio, Bernstein Analysis and Estimates

THE INCREMENTAL GROWTH RATE APPROACH (POPULAR IN INDUSTRIALS)

The incremental growth rate methodology approaches the industry balance from a different angle: instead of counting projects, it examines how quickly annual capacity additions are growing and projects that trend forward —the 2nd derivative of capacity. In our framework, we start with historical estimates of annual U.S. data center capacity additions, identify the current growth trajectory, and apply assumptions around future acceleration or deceleration to estimate installed capacity through 2030. Conceptually, this asks: if the industry continues building at an increasingly rapid pace, where does total capacity ultimately end up?

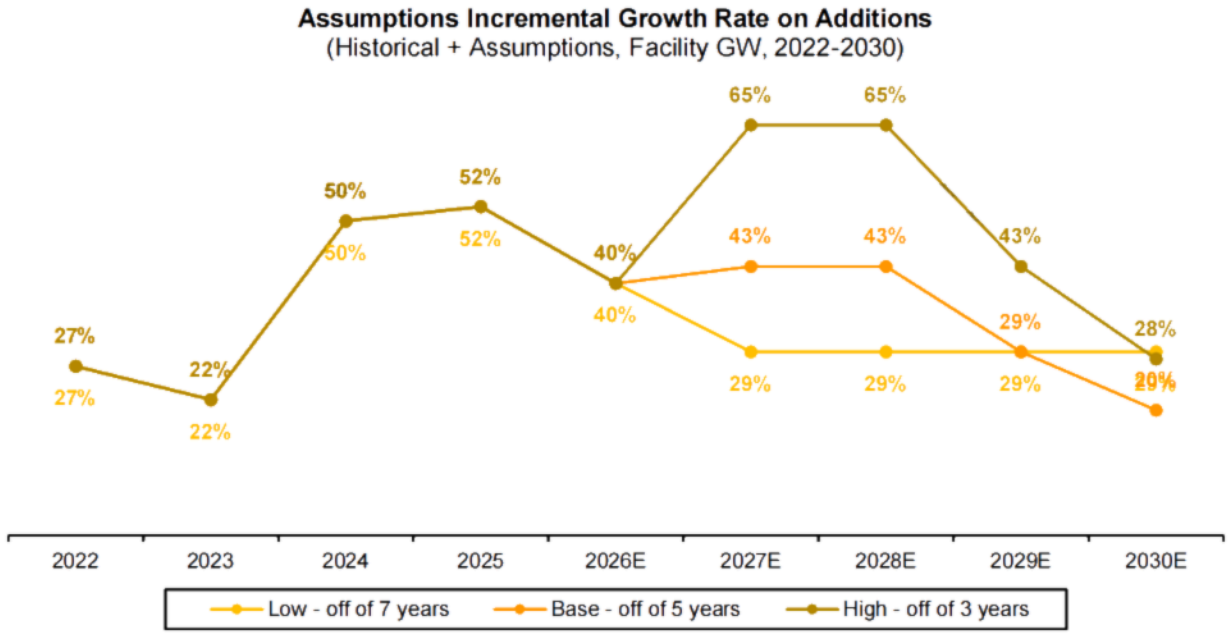
The primary advantage of this approach is that it is highly sensitive to changes in industry momentum. Whereas a site-based model can be constrained by today's visible development pipeline, a growth-rate model can capture the possibility that future demand drives a wave of projects that have not yet been announced. It is therefore particularly useful during periods of rapid market expansion, when historical development databases may struggle to keep pace with new investment activity. The methodology is also relatively simple and transparent. If investors believe the industry has entered a structurally different growth phase driven by AI infrastructure investment, then extrapolating recent capacity-addition trends provides a clear way to quantify what that acceleration could mean for future supply. In many cases, equipment manufacturers and construction-focused investors see evidence of growth inflections before those additions show up in operational capacity statistics, making this framework a useful way to incorporate such signals.

The biggest limitation is that the methodology is highly sensitive to the choice of starting point and growth assumptions. Small changes in annual growth rates can lead to dramatically different outcomes several years into the future. As a result, forecasts can become heavily influenced by whether the analyst views the current environment as a temporary spike, a cyclical upswing, or the beginning of a long-term structural expansion. Unlike a site-based framework, this approach is not directly anchored to identifiable projects or physical constraints. It can therefore generate outcomes that are theoretically achievable mathematically but may prove difficult to realize given power, labor, equipment, permitting, or construction bottlenecks. The methodology is often better at framing a range of possible outcomes than identifying a precise endpoint.

Who uses it? This framework is most commonly associated with industrial, electrical equipment, construction, and infrastructure supply-chain analysts, who tend to focus on the rate of new capacity deployment rather than the existing installed base. Because these participants often see order trends and construction activity earlier than REIT or operator data, they are naturally attuned to inflections in growth rates. As a result, growth-rate methodologies frequently produce some of

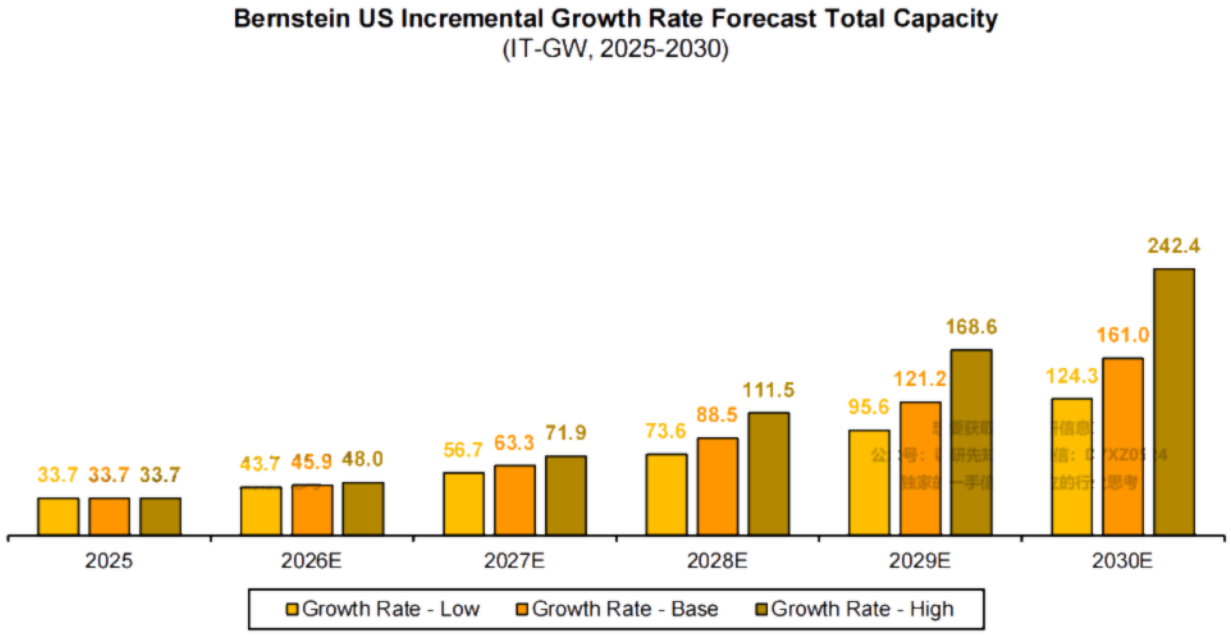
the more aggressive forecasts within the industry. In our analysis, the approach yields a broad range of ~124-242 IT-GW by 2030, with a base case of roughly 161 IT-GW, reflecting the possibility that today's elevated pace of capacity additions persists through the remainder of the decade.

EXHIBIT 8: **Assumptions Incremental Growth Rate on Additions (Historical + Assumptions, Facility GW, 2022-2030)**



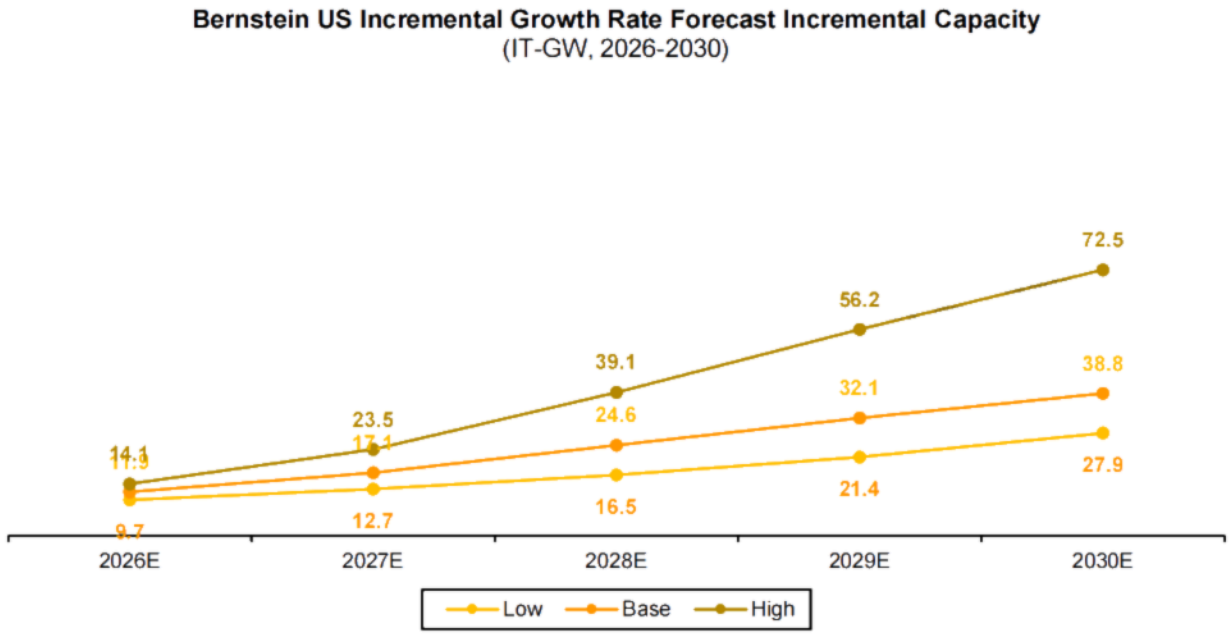
Source: Aterio, Bernstein Analysis and Estimates

EXHIBIT 9: **Bernstein US Incremental Growth Rate Forecast Total Capacity (IT-GW, 2025-2030)**



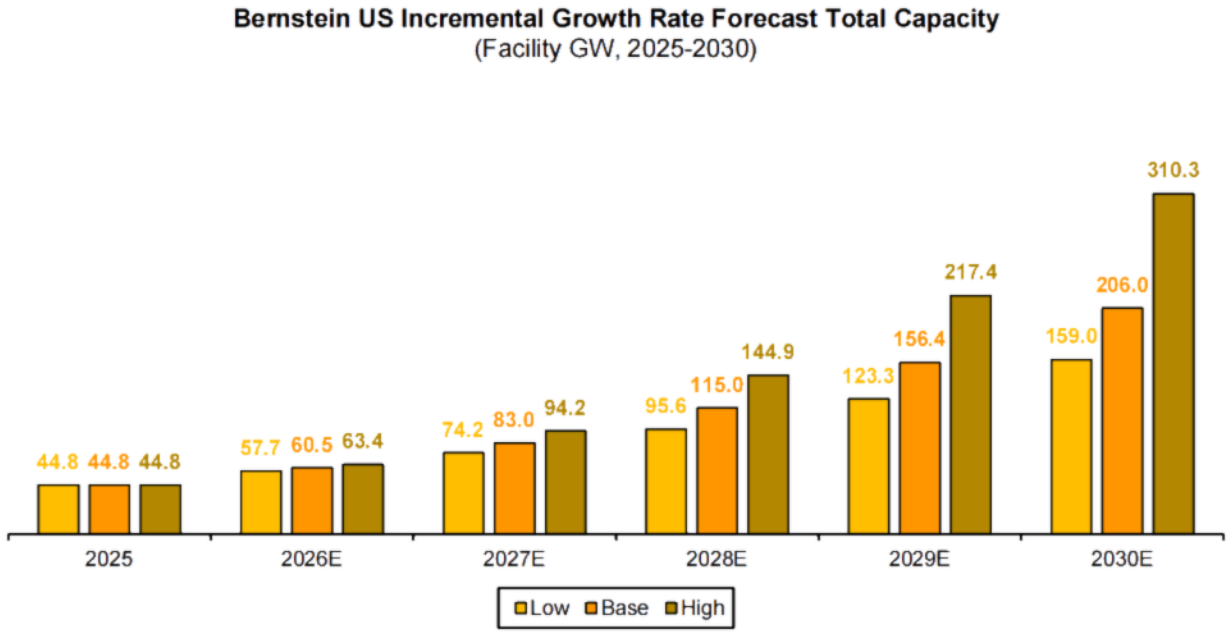
Source: Aterio, Bernstein Analysis and Estimates

EXHIBIT 10: **Bernstein US Incremental Growth Rate Forecast Incremental Capacity (IT-GW, 2026-2030)**



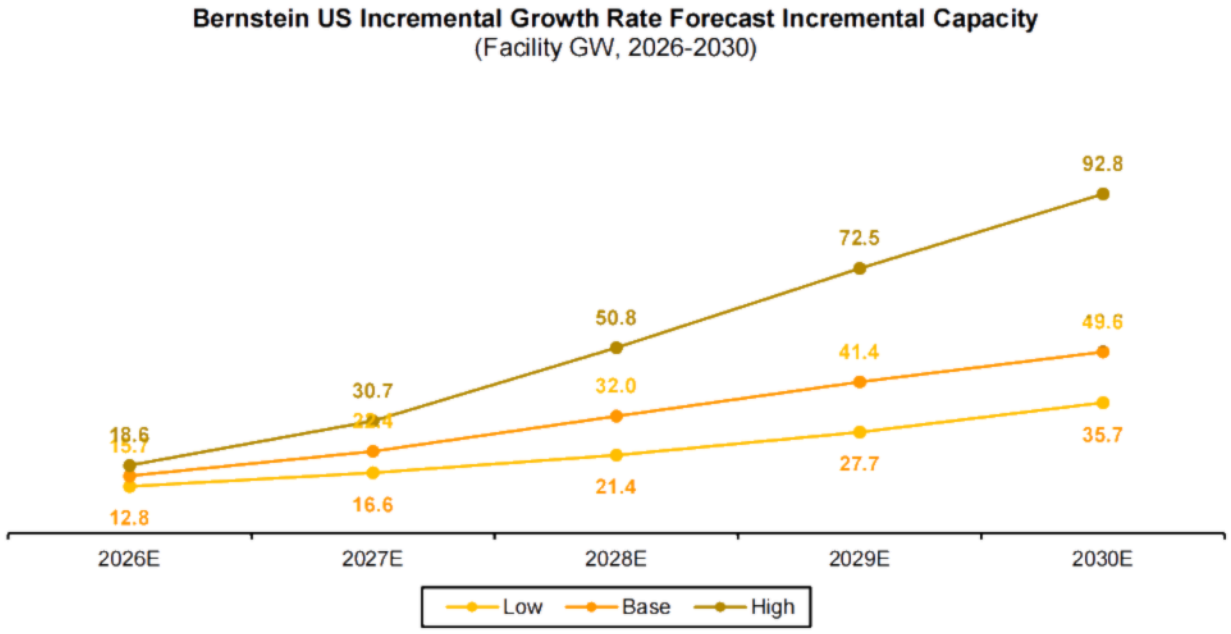
Source: Aterio, Bernstein Analysis and Estimates

EXHIBIT 11: **Bernstein US Incremental Growth Rate Forecast Total Capacity (Facility GW, 2025-2030)**



Source: Aterio, Bernstein Analysis and Estimates

EXHIBIT 12: **Bernstein US Incremental Growth Rate Forecast Incremental Capacity (IT-GW, 2026-2030)**



Source: Aterio, Bernstein Analysis and Estimates

THE HYPERSCALER-CAPEX METHOD (TECHNOLOGY AND INTERNET TEAMS)

This method begins with the largest drivers of data center demand: cloud and AI infrastructure spending by large hyperscalers such as Microsoft, Amazon, Google, Meta, and Oracle. Rather than forecasting capacity directly, this approach starts with expected hyperscaler capital expenditures and translates those investments into data center capacity using assumptions around the share of spending directed toward infrastructure, construction costs per MW, utilization rates, and power intensity. In effect, the methodology tries to answer: if hyperscalers spend what the market expects them to spend, how much data center capacity should that create?

The biggest advantage of this framework is that it is tied directly to the entities driving the majority of incremental demand. Hyperscalers account for much of the industry's recent growth, particularly as AI workloads have accelerated investment in data center infrastructure. As a result, following their capital spending plans provides a natural way to connect end-user demand to future capacity requirements. The methodology also benefits from a relatively robust set of public inputs for highly-covered companies. Investors have access to company disclosures, management commentary, consensus forecasts, and capex estimates, making it possible to update forecasts quickly as spending expectations change. Additionally, unlike site-based methods, it can also capture future demand that has not yet appeared in development pipelines.

On the other hand, the challenge is that capex is not necessarily capacity. Hyperscaler spending encompasses a wide range of investments, including servers, networking equipment, GPUs, land acquisition, power infrastructure, and data center construction. Translating dollars spent into megawatts requires multiple assumptions, each of which can materially influence the final result. The methodology is also sensitive to changes in technology and spending mix. For example, a dollar of capex directed toward AI accelerators may create less incremental physical capacity than a dollar spent on facility construction. Similarly, changes in hardware pricing, server density, deployment timing, or supply chain constraints can affect the relationship between capex and capacity. Disclosures across hyperscalers vary wildly and none give extremely clear breakdowns of their data center capacity. As a result, the framework can be directionally useful while still exhibiting significant uncertainty around the ultimate MW outcome.

This approach is most commonly employed by technology and internet analysts, particularly those focused on hyperscalers, cloud infrastructure, and AI investment trends. These investors tend to start with demand and spending forecasts rather than development pipelines, making capex a natural forecasting input. Because the methodology reflects expected customer demand rather than currently visible projects, it provides an important complement to site-based forecasting. In our analysis, the hyperscaler-capex framework produces a base-case estimate of approximately 152 IT-GW by 2030, roughly in line

with our overall forecast. However, it also has one of the widest swing of our forecasts, as the results are highly sensitive to the underlying assumptions.

EXHIBIT 13: **Bernstein Annual CAPEX per Hyperscalers Assumptions**

	US Share Estimate	2025	2026E	2027E	2028E	2029E	2030E
Microsoft	60%	$114	$170	$200	$216	$234	$252
Amazon	72%	$131	$222	$304	$306	$309	$312
Alphabet	75%	$97	$207	$334	$337	$340	$343
Meta	88%	$77	$150	$236	$245	$255	$265
Oracle	87%	$45	$107	$92	$97	$97	$97
Total Hyperscaler Capex		**$466**	**$856**	**$1,166**	**$1,202**	**$1,235**	**$1,269**
Estimated U.S. Capex		**$344**	**$642**	**$877**	**$903**	**$926**	**$951**
% of capex allocated to data centers		85%	85%	85%	85%	85%	85%
% capex allocated to hardware refresh		10%	11%	11%	12%	12%	13%
New data center growth capex		**$263**	**$489**	**$664**	**$680**	**$693**	**$707**

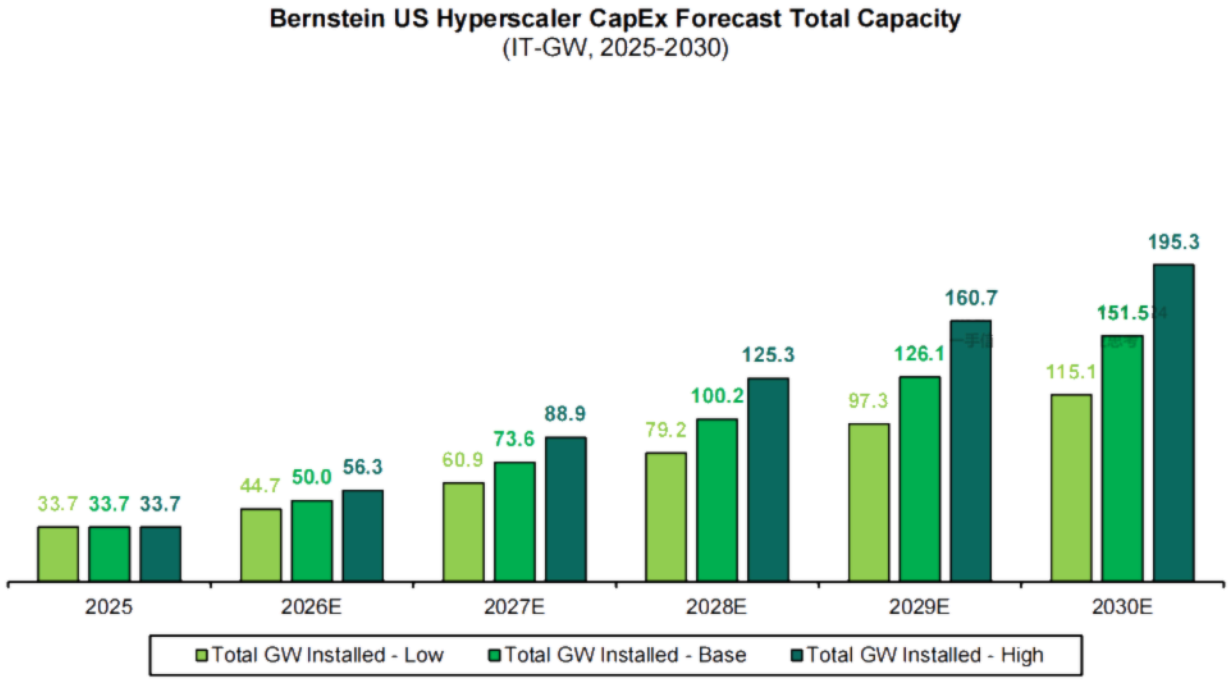
Source: Bloomberg, Company filings and disclosures, Aterio, Bernstein Analysis and Estimates

EXHIBIT 14: **Inputs for Hyperscalers IT-GW Additions Overview**

Inputs - IT-GW Additions	2025	2026E	2027E	2028E	2029E	2030E
Portion of a year's capex energized in-year	50%	53%	55%	58%	60%	63%
Capex energized this year, after lag	**$245**	**$363**	**$555**	**$654**	**$668**	**$681**
Hyperscale GW additions - Low	5.5	7.6	11.0	12.3	11.9	11.6
Hyperscale GW additions - Base	7.0	9.6	13.7	15.1	14.5	14.0
Hyperscale GW additions - High	8.2	11.1	15.6	17.1	16.3	15.6
Hyperscale as % of total market - Low	70%	69%	68%	67%	66%	65%
Hyperscale as % of total market - Base	60%	59%	58%	57%	56%	55%
Hyperscale as % of total market - High	50%	49%	48%	47%	46%	45%

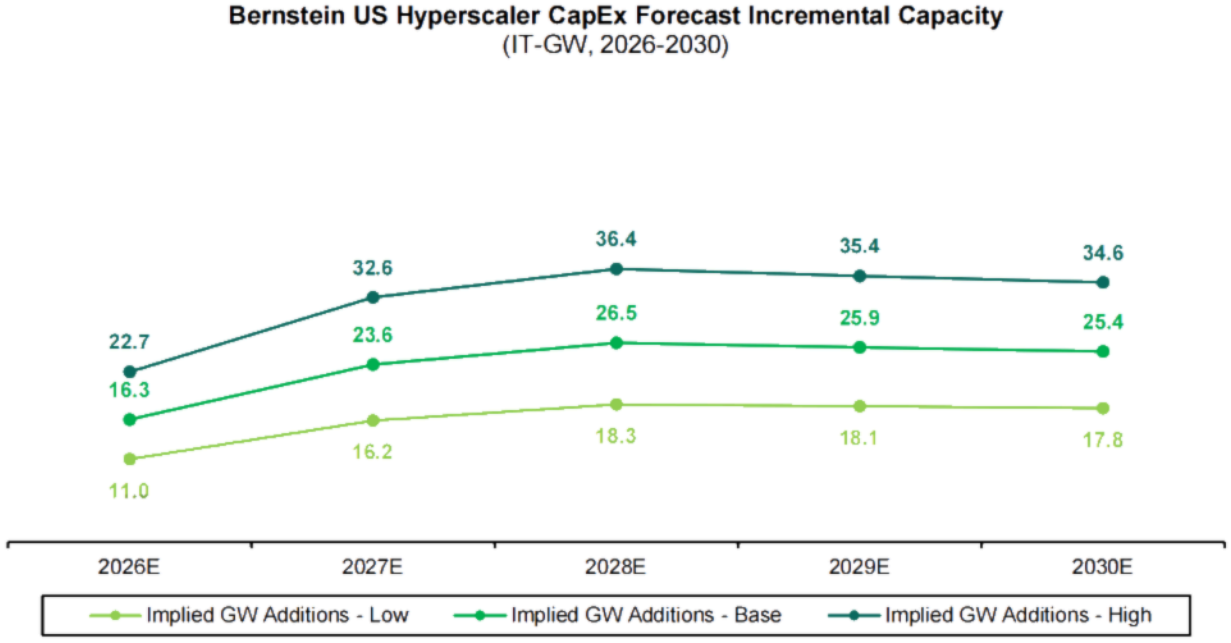
Source: Bloomberg, Company filings and disclosures, Aterio, Bernstein Analysis and Estimates

EXHIBIT 15: **Bernstein US Hyperscaler CapEx Forecast Total Capacity (IT-GW, 2025-2030)**



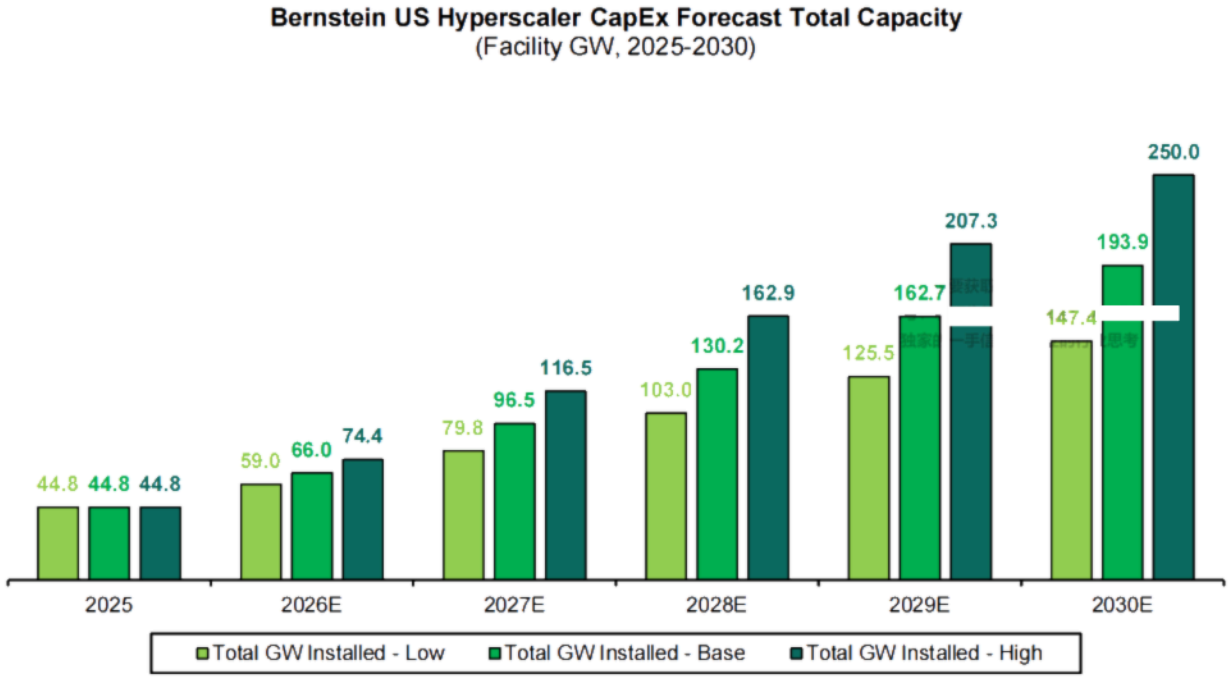
Source: Bloomberg, Company filings and disclosures, Aterio, Bernstein Analysis and Estimates

EXHIBIT 16: **Bernstein US Hyperscaler CapEx Forecast Incremental Capacity (IT-GW, 2026-2030)**



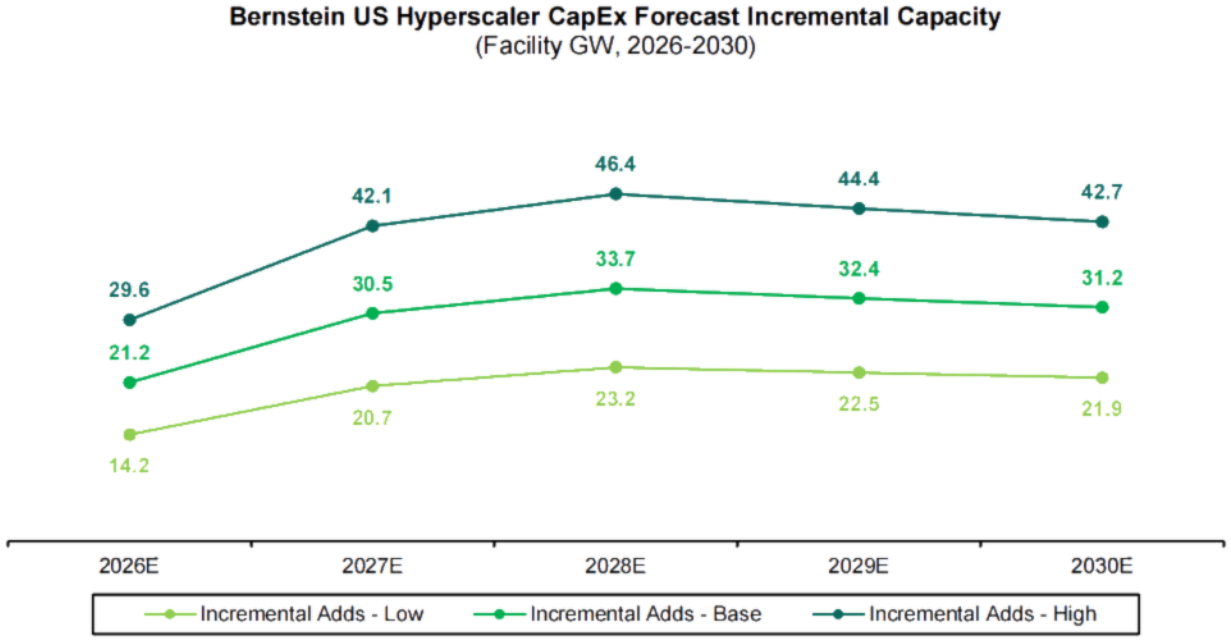
Source: Bloomberg, Company filings and disclosures, Aterio, Bernstein Analysis and Estimates

EXHIBIT 17: **Bernstein US Hyperscaler CapEx Forecast Total Capacity (Facility GW, 2025-2030)**



Source: Bloomberg, Company filings and disclosures, Aterio, Bernstein Analysis and Estimates

EXHIBIT 18: **Bernstein US Hyperscaler CapEx Forecast Incremental Capacity (Facility GW, 2026-2030)**



Source: Bloomberg, Company filings and disclosures, Aterio, Bernstein Analysis and Estimates

THE SERVER-BASED METHOD (SEMIS VIEW)

The server-based methodology starts with the compute deployed inside data centers rather than the fa
themselves. This approach forecasts shipments of servers, AI accelerators, GPUs, racks, or related hardware a

人在草木间

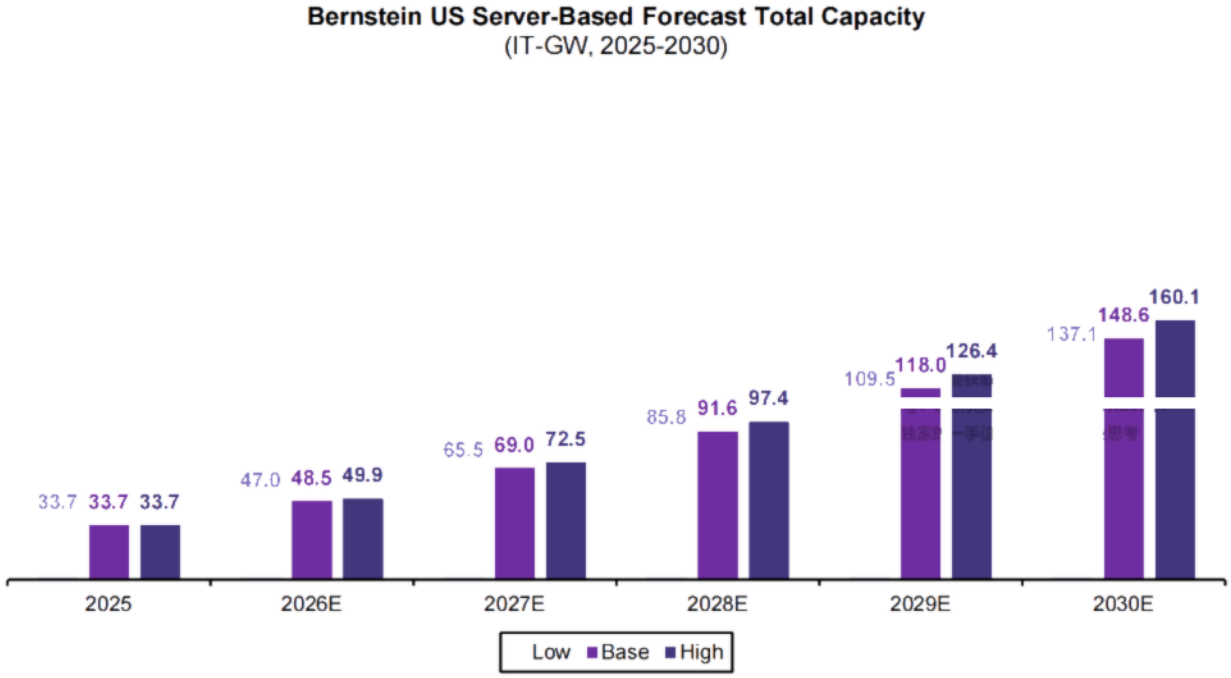
converts that installed compute base into power consumption and, ultimately, data center capacity requirements. In our framework, we rely heavily on server and AI hardware deployment forecasts, including accelerator shipment expectations, and translate those deployments into IT load using assumptions around power draw, utilization, refresh cycles, and geographic allocation. Put simply, the methodology asks: if the industry deploys the amount of compute hardware currently expected, how much data center capacity will be required to support it?

The primary advantage of this approach is that it is closely tied to the hardware ultimately filling the data centers. Every AI cluster, server rack, or accelerator shipment must eventually be deployed somewhere, making hardware forecasts a useful way to estimate future infrastructure requirements. This methodology can also capture demand signals before they become visible in development pipelines or construction activity, particularly during periods of rapid technology adoption. The approach is especially relevant in the current AI cycle, where growth is being driven by increasingly power-dense compute deployments. Semiconductor shipment forecasts often provide one of the earliest indicators of future infrastructure needs, allowing analysts to assess potential demand before capacity additions are formally announced.

The principal challenge is converting hardware deployments into data center capacity. Several assumptions stand between a server shipment forecast and a megawatt estimate, including utilization rates, geographic mix, deployment timing, replacement cycles, rack density, and future improvements in efficiency. Small changes in these assumptions can meaningfully affect the resulting capacity forecast. The methodology also faces timing challenges. Hardware shipments do not necessarily coincide with facility commissioning, and some equipment may replace existing infrastructure rather than drive entirely new capacity demand. As a result, the resulting capacity forecast is particularly sensitive to assumptions around deployment timing, refresh cycles, utilization, and the conversion of compute demand into physical infrastructure.

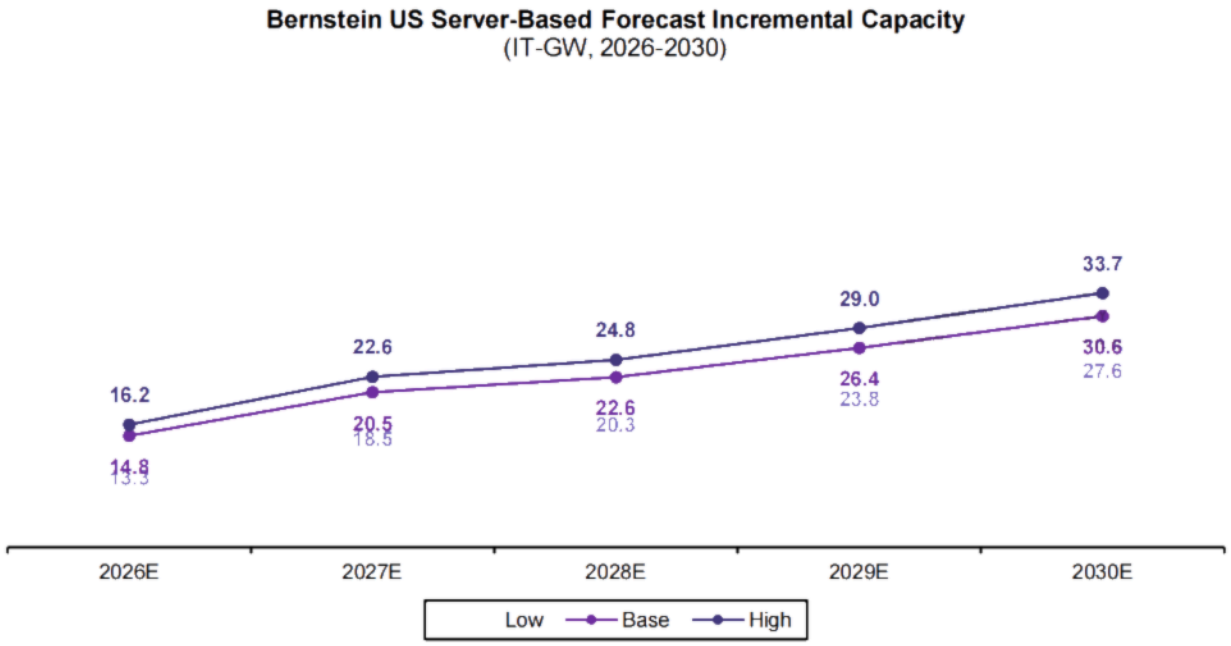
This framework is most commonly associated with semiconductor, hardware, and AI infrastructure analysts, who track accelerator demand, server deployments, supply chains, and advanced packaging capacity. These investors tend to focus on the technologies driving data center growth rather than the facilities themselves, making a server-based framework a natural extension of their coverage universe. Because the methodology begins with end-demand for compute, it often serves as a useful complement to both site-based and hyperscaler-capex approaches. In our analysis, the server-based framework produces a base-case estimate of approximately 149 IT-GW by 2030, making it one of the higher forecasts in our survey and broadly consistent with our 149 IT-GW Bernstein estimate. This result reflects the substantial infrastructure requirements implied by current expectations for AI compute deployment.

EXHIBIT 19: **Bernstein US Server-Based Forecast Total Capacity (IT-GW, 2025-2030)**



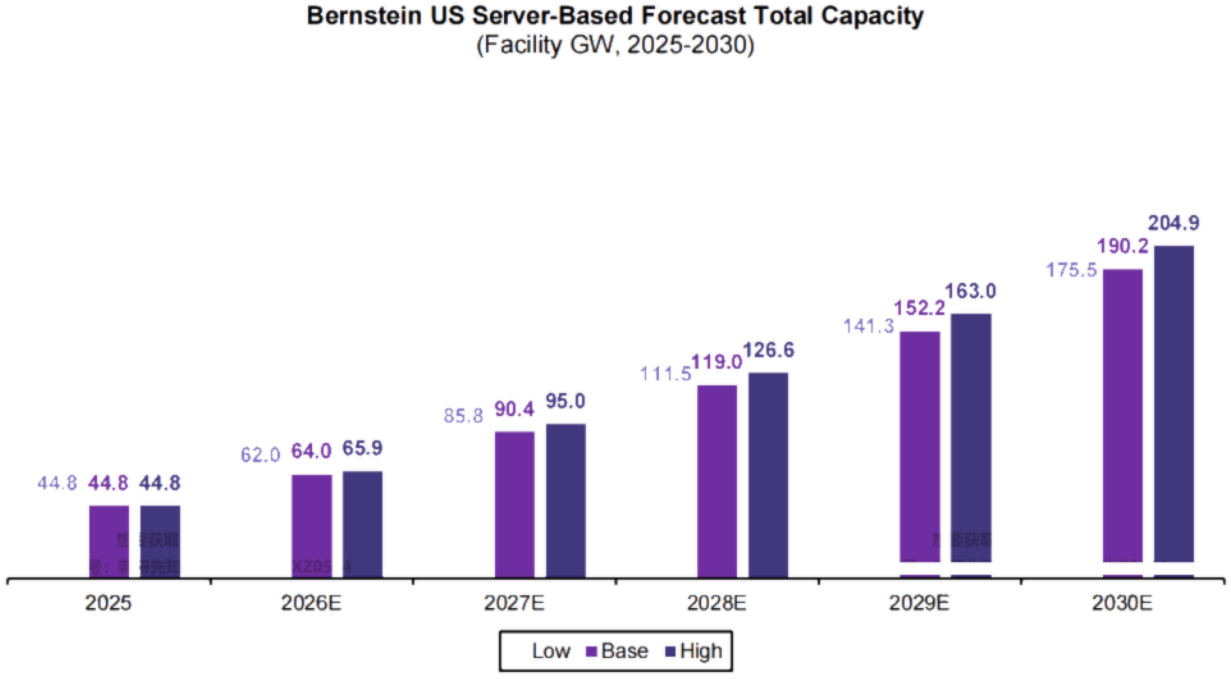
Source: Bernstein Analysis and Estimates

EXHIBIT 20: **Bernstein US Server-Based Forecast Incremental Capacity (IT-GW, 2026-2030)**



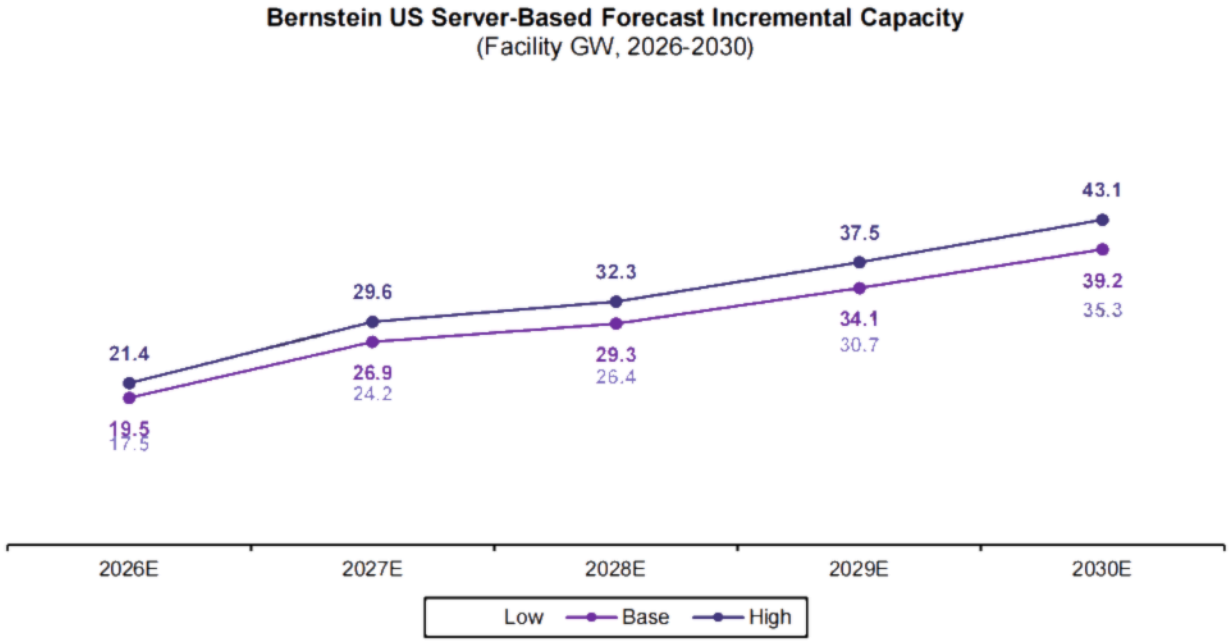
Source: Bernstein Analysis and Estimates

EXHIBIT 21: **Bernstein US Server-Based Forecast Total Capacity (Facility GW, 2025-2030)**



Source: Bernstein Analysis and Estimates

EXHIBIT 22: **Bernstein US Server-Based Forecast Incremental Capacity (Facility GW, 2026-2030)**



Source: Bernstein Analysis and Estimates

THE CRAFT LABOR VIEW

As a team, we believe craft labor may be one of the biggest ceilings limiting data center construction. Chad Dillard's craft labor approach looks at the market from the supply side and assesses how much capacity the industry can physically build with the available workforce. In our framework, we begin with estimates of the skilled labor required to construct data center capacity and compare those requirements against the available pool of electricians, pipefitters, HVAC technicians, and other specialized trades. The resulting analysis establishes a practical ceiling on deployment based on workforce availability and productivity assumptions.

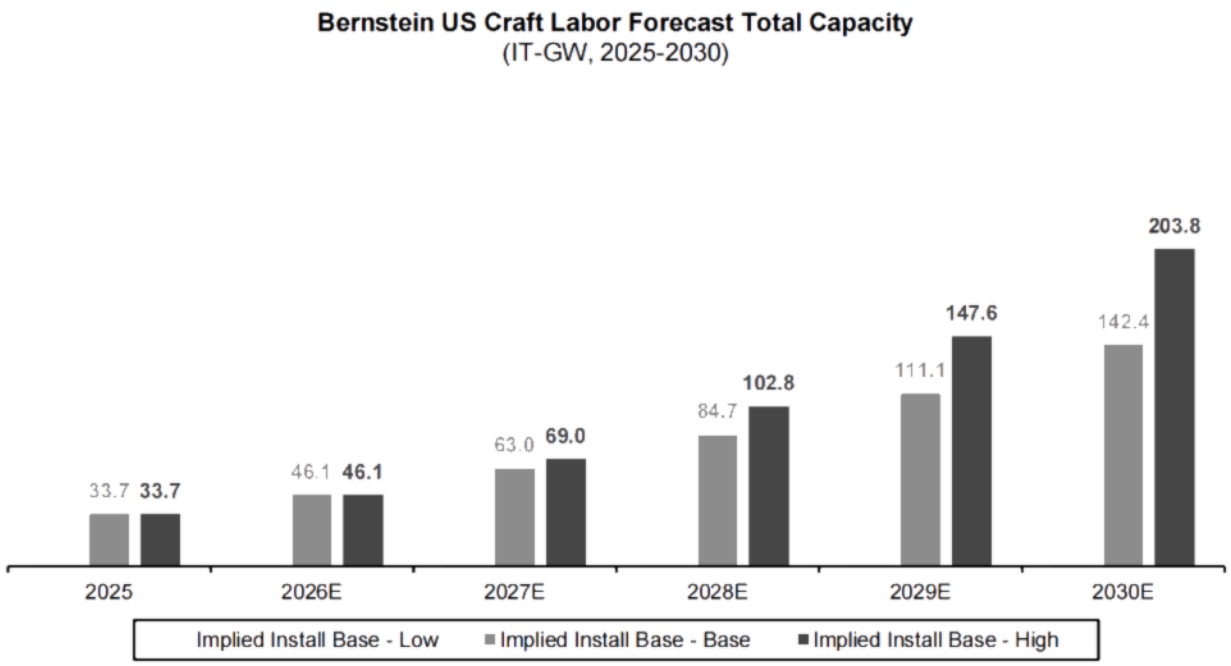
This approach explicitly incorporates one of the industry's most important physical constraints. While many forecasting methodologies focus on demand signals, project pipelines, or capital availability, all data centers ultimately require skilled labor to move from planning to operation. By grounding forecasts in workforce capacity, this methodology provides a useful reality check on projections that may implicitly assume unlimited construction resources. The framework is particularly valuable during periods of rapid industry expansion, when labor shortages can become a meaningful bottleneck. It also captures an important dimension that is often underrepresented in other forecasting approaches: even if power, capital, and demand are available, projects cannot be delivered without the workers required to build them.

The main limitation is that labor availability is itself difficult to forecast and can evolve over time. Workforce participation, regional migration, training programs, productivity improvements, modular construction techniques, and changing project designs can all influence the amount of capacity that can be built with a given labor pool. In addition, labor is only one of several potential constraints on development. A labor-based forecast may underestimate future capacity if the industry successfully expands the workforce or improves construction productivity. Conversely, it may overestimate achievable deployment if other bottlenecks, such as power infrastructure, equipment availability, or permitting delays, prove more restrictive than labor availability. As a result, forecasts can be sensitive to assumptions regarding labor productivity, workforce growth, and the relationship between labor inputs and delivered capacity.

This framework is most commonly associated with construction, engineering, labor, and industrial analysts, as well as investors focused on the practical constraints of infrastructure deployment. It is typically used less as a demand forecast and more as a test of feasibility, helping determine whether projected growth can realistically be delivered by the industry's available workforce. In our analysis, the craft labor framework produces a base-case estimate of approximately 142 IT-GW by 2030, making it one of the more conservative forecasts within our methodology set, along with site-based. Importantly, it is

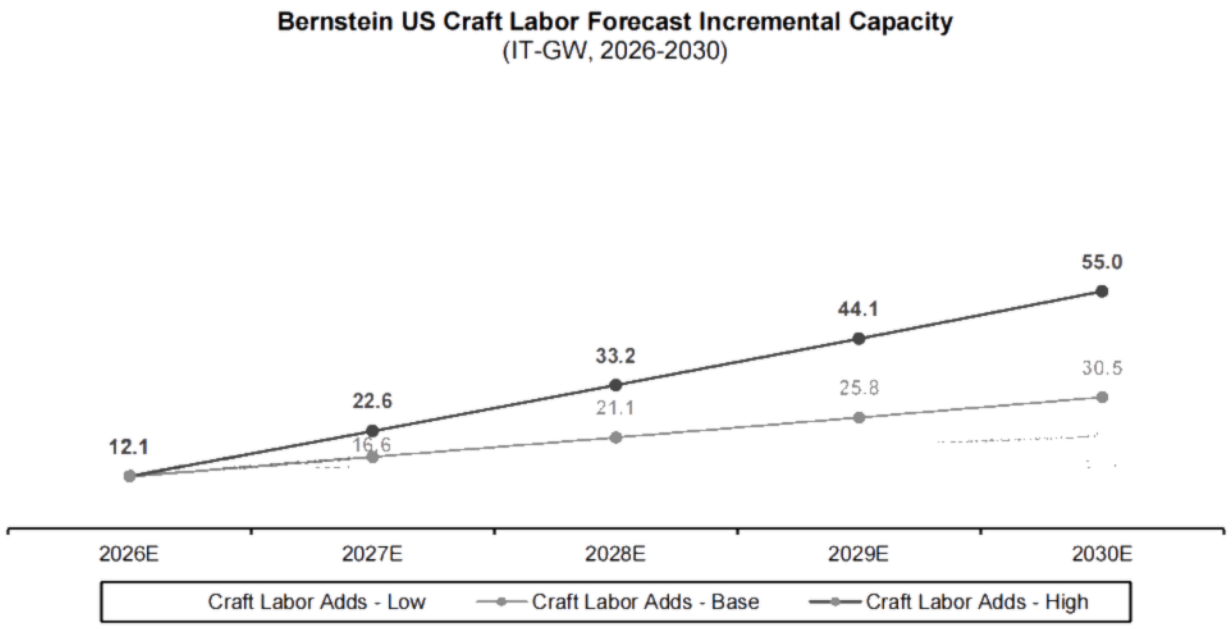
best interpreted as a potential deployment ceiling rather than a demand forecast, reflecting the amount of capacity that could reasonably be supported by available labor resources under our given set of assumptions.

EXHIBIT 23: **Bernstein US Craft Labor Forecast Total Capacity (IT-GW, 2025-2030)**



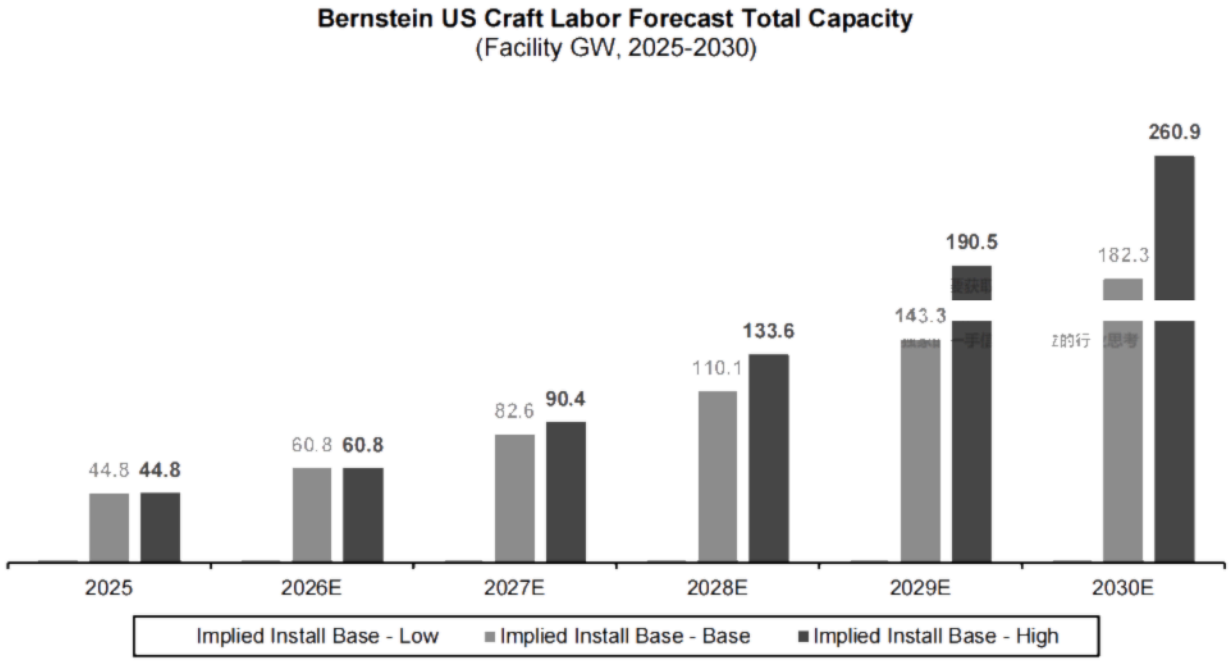
Source: Bernstein Analysis and Estimates

EXHIBIT 24: **Bernstein US Craft Labor Forecast Incremental Capacity (IT-GW, 2026-2030)**



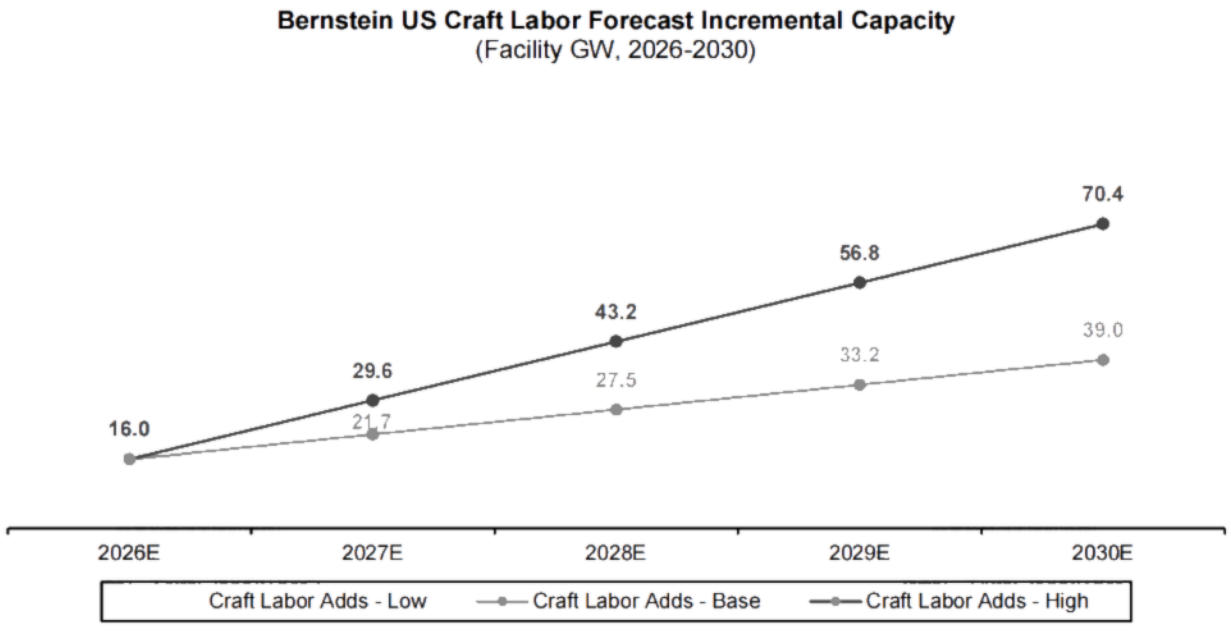
Source: Bernstein Analysis and Estimates

EXHIBIT 25: **Bernstein US Craft Labor Forecast Total Capacity (Facility GW, 2025-2030)**



Source: Bernstein Analysis and Estimates

EXHIBIT 26: **Bernstein US Craft Labor Forecast Incremental Capacity (Facility GW, 2026-2030)**

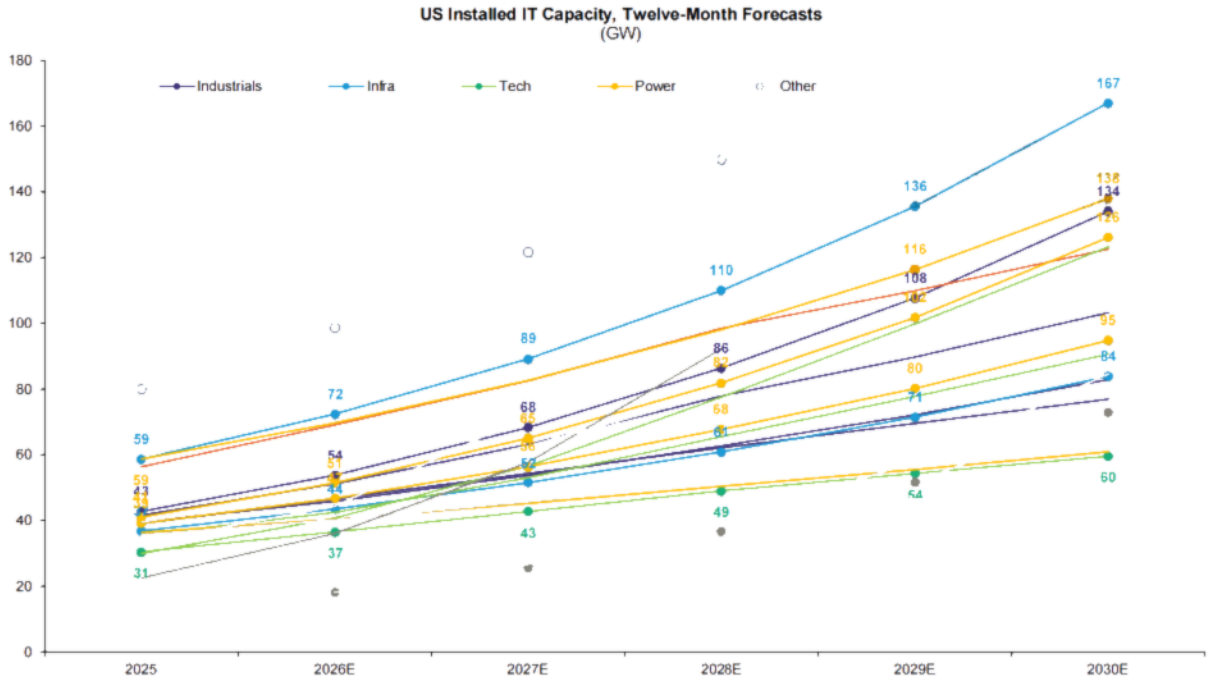


Source: Bernstein Analysis and Estimates

APPENDIX - FINANCIAL FORECASTS

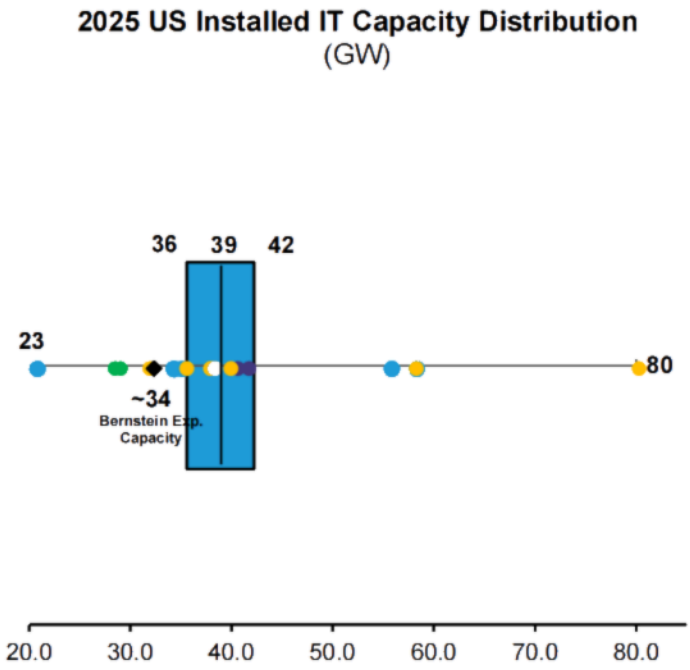
RANGE OF EXISTING FORECASTS

EXHIBIT 27: **US Installed IT Capacity, Twelve-Month Forecasts (GW)**



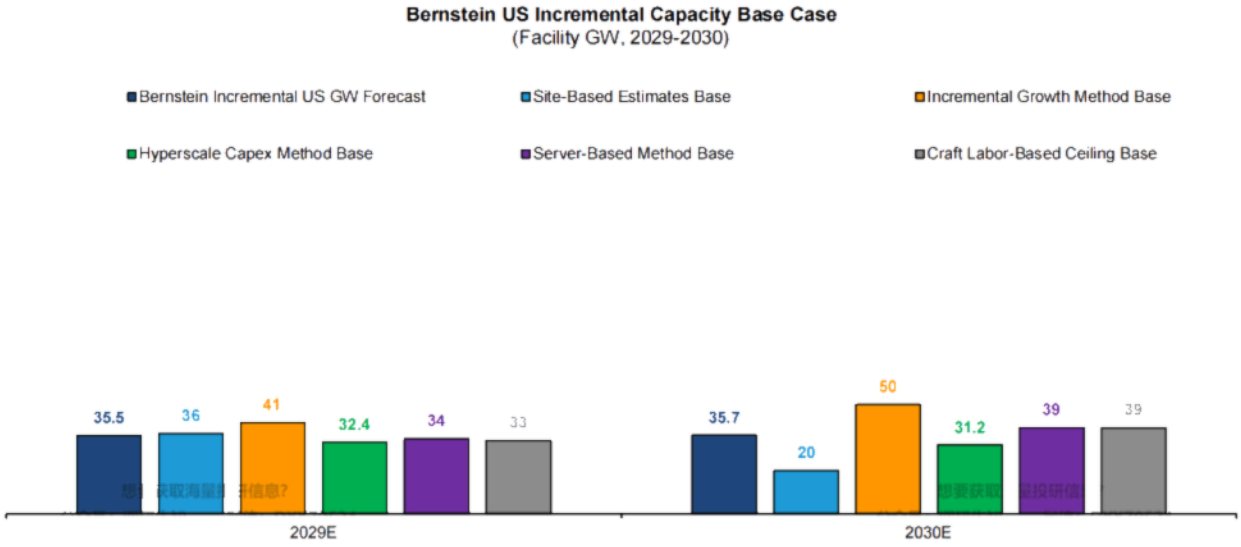
Source: LBNL, Cleanview, Cushman, JLL, S&P Global, Bloomberg, Grid Strategies, EPRI, Bain & Co, Bernstein analysis

EXHIBIT 28: **2025 US Installed IT Capacity Distribution (GW)**



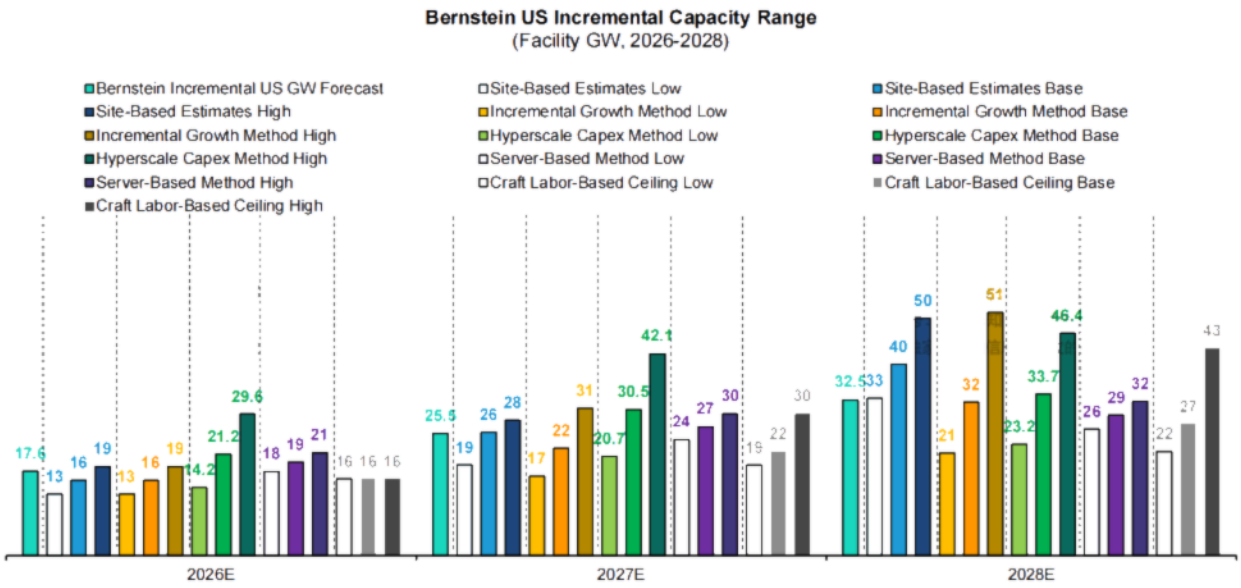
Source: LBNL, IDC, IEA, S&P Global, SemiAnalysis, Cushman, DC Byte, Cleanview, JLL, Bloomberg, BCG, Grid Strategies, EPRI, Bain & Co, McKinsey, Bernstein Analysis and Estimates

EXHIBIT 29: **Bernstein US Incremental Capacity by Methodology Base Case (Facility GW, 2029-2030)**



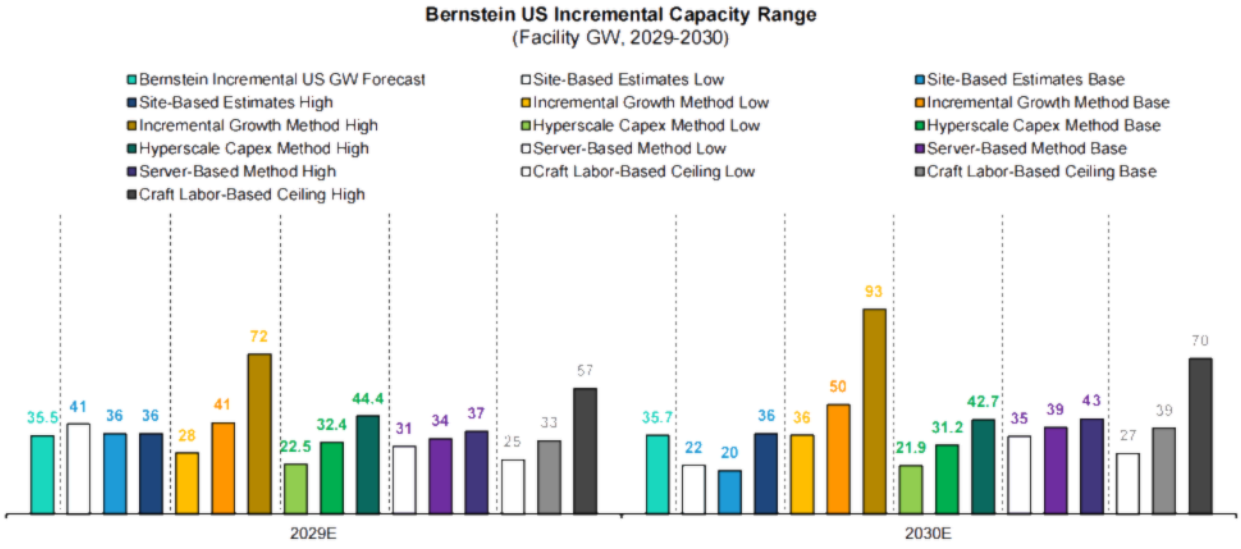
Source: Bloomberg, Company filings and disclosures, Aterio, Bernstein Analysis and Estimates

EXHIBIT 30: **Bernstein US Incremental Capacity Range (Facility GW, 2026-2028)**



Source: Bloomberg, Company filings and disclosures, Aterio, Bernstein Analysis and Estimates

EXHIBIT 31: **Bernstein US Incremental Capacity Range by Methodology (Facility GW, 2029-2030)**



Source: Bloomberg, Company filings and disclosures, Aterio, Bernstein Analysis and Estimates