BCG | ABB

The Strategic Case for Hybrid AC/DC Power

September 2026

By Emmanuel Sissimatos, Monika Saunders, Marc Schmidt, and Tom Piendl

Introduction

The world's energy system is quietly rewiring itself. The dominant new energy sources and loads are natively direct current (DC): solar PV generates it, batteries store it, and data centers and EVs consume it. Yet the distribution infrastructure connecting them still runs on alternating current (AC), the architecture that has prevailed since the 1890s.

That mismatch was once an acceptable inefficiency. Every AC-to-DC conversion loses a small percentage of energy, but the toll is so modest that it barely mattered in the past. The rapid scaling of AI computing power (known familiarly as *AI compute*) is changing that calculus. To keep up with demand, hyperscale campuses will need to operate at gigawatt (GW) scale, magnifying the cost of unnecessary conversions. These facilities force every watt through multiple conversion stages before it reaches a server—stages that exist not because they are necessary but because the infrastructure predates the loads it now serves.

In four sectors—data centers, industrial plants, commercial buildings, and residential buildings—the shift toward greater DC integration is already underway. The destination is not an all-DC grid, but a hybrid architecture: AC remains the backbone of most transmission and regional distribution, while DC increasingly powers high-density applications, mostly behind the meter. Across sectors, the urgency, the economics, and the window for action differ.

This paper, published by BCG in collaboration with ABB, quantifies the economics in areas where decisions about adopting new electrical architecture are already live, and where they soon will be:

- For AI data centers, DC is rapidly becoming the required architecture. It meets the power density demands of next-generation AI hardware, fits more compute behind a scarce grid connection, and turns efficiency gains directly into revenue.
- For automation-intensive industrial plants, the business case for hybrid architecture is positive from the first day of operation, driven by the DC-native loads that most factories already contain.
- For large commercial buildings, DC pays off by creating a backbone for major loads such as elevators, climate-control systems, and energy storage while leaving tenant-facing systems on AC.
- For residential buildings, the business case for hybrid architecture is not yet compelling, but policy shifts on the calendar may improve the economics.

The barriers to faster adoption are systemic rather than technical, which means that this is more of a coordination problem than an engineering issue. Equipment and automation suppliers, hyperscalers, industrial operators, and large developers that commit capital and set specifications over the next two to three years will shape the transition rather than inherit it. Their decisions will influence the cost curves, voltage defaults, and standards that define the next decade. The question that each organization must answer is whether it will help set the standards or merely accept them as developed by others.

How AC Won and Where DC Is Returning

Electrification started with DC. The first commercial power grids, in 1882 in New York and London, ran on 110V DC. These grids worked within strict limits: no practical technology could step DC voltage up for long-distance transmission, so high currents produced losses that kept generation local.

The solution to this problem was the transformer, AC's decisive enabler. By stepping voltage up for transmission and back down for consumption, AC made large centralized generators and national grids viable. The 175-kilometer Lauffen-to-Frankfurt line of 1891 proved the model at scale. Development of the AC induction motor then reinforced the advantage: brushless and robust, this motor needed far less maintenance than the commutator-and-brush DC machines of the era, making AC the natural choice on factory floors. Industrialization then locked in the AC architecture. AC motors, AC generators, and AC protection systems, along with decades of capital investment, created a path dependency that regulation codified.

Although AC prevailed in the war of currents, DC never disappeared. Instead, it retreated to applications where its advantages were decisive: computers, rail traction, marine vessels, and medical devices. Now, however, the structural logic of energy itself has changed. The energy transition is accelerating DC generation and DC consumption simultaneously, making a DC distribution layer self-reinforcing rather than aspirational. Rail and marine provide useful blueprints for what large-scale DC looks like in practice: metropolitan and suburban rail has run on DC traction (typically 750V to 3 kV DC) for over a century, and ships have shifted to low-voltage DC (LVDC) across vessels with up to 20 megawatts (MW) of system power since 2013.

The scale of DC's quiet return is larger than most executives realize. Motor-driven systems consume roughly 70% of industrial electricity, and every variable-frequency drive (VFD) converts that power through an internal DC link. About one in four installed motors globally already carry a VFD, and adoption continues to grow. Therefore, at the component level, modern factories are already partly DC. But the conversion happens invisibly inside each device, and its losses are accepted as overhead.

At higher power densities and loads, DC's physical advantages become increasingly significant, and the practical limits of low-voltage AC (LVAC) become harder to overcome. Leading AI server racks are expected to reach roughly 600 kW in 2027 and to approach 1 MW later in the decade, according to Nvidia's published rack roadmap. At those densities, 800V DC distribution becomes the only practical option. The voltage level is not experimental: 800V DC already operates in electric vehicle powertrains, produced at automotive industry volumes.

At such high voltages, DC can carry the same power as an equivalent AC system through a conductor cross-section that is roughly half as large, an advantage that compounds as power density rises. AC won the war of currents because power electronics did not yet exist. The transformer gave AC an inexpensive, passive way to change voltage at a time when DC could not offer a comparable solution. Semiconductor converters now perform that function for DC, removing the principal constraint that formerly made AC the only viable architecture. That parity now extends across the entire voltage range: low, medium, and high voltage each have AC and DC equivalents, from device-level supply to long-distance transmission. (See Exhibit 1.)

The Hybrid Architecture

The return of DC does not mean the demise of AC. The future grid will emerge not from a binary choice between AC and DC, but from an architecture that uses each type of current where it performs best. AC will remain the synchronous network that ties the system together, from regional distribution to the plug standards locked into consumer end-products. DC will increasingly complement that network at both ends: in high-capacity, long-distance transmission and in application-specific systems behind the meter.

The architecture of the hybrid grid follows the physics and economics of three layers. (See Exhibit 2.) Long-distance transmission increasingly relies on high-voltage DC (HVDC) for controllable, high-capacity transfer. Regional distribution remains predominantly AC, where mature transformers and passive voltage transformation still provide the most economical solution. The fastest change is behind the meter: data centers, industrial plants, commercial buildings, and, in the long term, households will become DC systems connected to the AC grid through a single interlinking converter.

That single interlinking converter makes the hybrid architecture practical. It isolates the internal DC system from the AC grid, concentrating grid-code compliance at one defined connection point rather than extending it to every component behind the meter. Utilities can continue to operate the upstream network largely as they do today, while facility operators retain control over the architecture within their sites. Because adoption does not require a wholesale redesign of the surrounding grid, DC can spread incrementally in the form of localized systems around dense demand, rather than through a grid-wide conversion.

The relevant question today is, where does the hybrid architecture create enough value to justify its adoption?

EXHIBIT 1

AC and DC Now Span the Full Voltage Range

	AC	DC
Low voltage	**Less than 1 kV** • Household and office supply • Industrial motor and climate-control systems	**0 to 1.5 kV** • AI computing power and industrial DC systems • EV fast charging, solar PV, and battery energy storage
Medium voltage	**1 to 60 kV** • Urban and industrial distribution • Wind and solar grid interconnection	**1.5 to ~100 kV** • Solar PV collection and offshore wind collection (emerging) • Distribution-grid DC links (angle-DC pilot in the UK)
High voltage	**More than 60 kV** • National grid transmission • Large-scale power plant output	**More than ~100 kV** • Offshore wind export and subsea interconnection • Long-distance renewable transmission

Source: BCG analysis.
Note: Standards for high-voltage DC systems and components are well established, but the boundary between medium- and high-voltage DC (proposed at approximately 100 kV) has not yet been formally confirmed. AC = alternating current; DC = direct current; EV = electric vehicle; kV = kilovolt; PV = photovoltaic.

EXHIBIT 2

The Hybrid Grid Assigns Distinct Roles to AC and DC Across Three Layers

Visualization of AC and DC

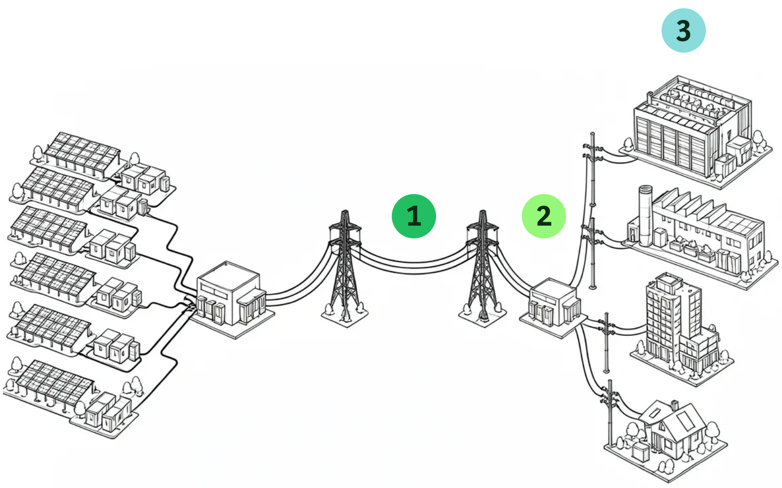

Three layers of hybrid grid

1 Long-distance transmission — DC
- Enables controllable long-distance, high-capacity power transfer
- Minimizes reactive power constraints and reduces loop flows

2 Regional distribution — AC
- Provides cost-efficient infrastructure and reliable protection systems
- Allows passive voltage transformation via transformers

Interlinking converter

3 Behind-the-meter DC systems — DC
- Reduces conversion stages for DC-native loads
- Simplifies storage integration and peak buffering

Data centers	Industrial
Commercial buildings	Households

Sources: European Commission; IEA; BCG analysis.
Note: AC = alternating current; DC = direct current.

Three Myths and Two Real Barriers

Some skeptics challenge this vision. The case against DC rests on five objections heard across boardrooms, engineering teams, and procurement committees. Three of these objections are myths rooted in the DC technology of 10 or 20 years ago; advances in power electronics, protection, and deployment have resolved them. Two of them, however, represent real barriers to adoption. These are the three myths:

- **Myth 1: DC is only for niche or small-scale applications.** The reality is that DC already operates across every voltage tier. HVDC has been commercially mature for decades in offshore wind and grid interconnection. At the LVDC tier, IT equipment, EV batteries, and drives run on DC internally, and AI data centers are transitioning to 800V DC distribution, with AC converted to DC centrally.

- **Myth 2: DC cannot be protected safely.** This concern has a sound physical basis. Unlike AC, DC has no natural current zero-crossing—and as a result, faults are more difficult to interrupt. Until recently, that shortcoming posed a serious engineering challenge, but advances in solid-state and hybrid breakers have largely resolved it. Commercial systems now interrupt DC faults in microseconds, and international standards increasingly define safe operation for LVDC installations.[1] The remaining challenge is whether standardized, solid-state protection can reach the cost and production scale required for broad adoption.

- **Myth 3: DC is inherently more expensive.** Some DC components carry a premium because they are produced at far lower volumes than comparable AC equipment. That premium reflects manufacturing scale, not some fundamental property of DC. The cost differential should continue to narrow as suppliers deploy the same power-electronics platforms that are already scaling in solar inverters, EV drives, and industrial converters. In any case, component prices do not determine total system cost. In deployed configurations, shared DC architectures can reduce conversion equipment, copper use, and space requirements while increasing usable power density.[2]

And these are the real barriers, neither of which is primarily technological:

- **Real Barrier 1: Standards remain fragmented.** Standards exist across the DC voltage range, but they vary in maturity and are inconsistent across regions. As a result, equipment approved in one market may require redesign, retesting, or project-specific approval in another, increasing cost and slowing deployment.[3]

- **Real Barrier 2: The workforce lacks DC-specific expertise.** Designing, protecting, integrating, and commissioning DC systems require capabilities that most electrical engineers and electricians do not possess. The challenge resembles the early years of utility-scale solar, when proven technology and compelling economics outpaced the supply of certified installers. But these capabilities build naturally on existing expertise in power electronics and VFDs, and targeted retraining can occur far faster than building an entirely new workforce. Closing the skills gap must happen in parallel with standards development, however, not after it.[4]

1. IEC technical standard 60479 indicates that AC and DC present different—not uniformly higher or lower—shock risks. At equal voltage, DC has a higher let-go threshold. Hybrid HVDC breakers are already in commercial service, while LVDC protection has been validated through multiyear demonstrations such as DC-INDUSTRIE and Current/OS. Commercial products, including ABB's IEC-certified SACE Infinitus solid-state DC breaker, can interrupt faults in microseconds at up to 1,000V DC.
2. Components that may still carry a premium include DC protection switches and central converters. Even so, system-level savings can outweigh higher component costs in many instances. In one industrial smart-grid assessment, a shared DC bus reduced drive-system costs by about 15%, although additional protection and grid-management equipment offset part of the benefit (Fraunhofer IPA, *Resource and Energy Efficiency Assessment of an Industrial DC Smart Grid*, 2020). Cabling studies report copper savings of approximately 25% when one conductor is eliminated (WSED, *Benefits of Industrial DC Grids*, 2023) and up to 40% when conductor count falls from five to three (ZVEI/DC-INDUSTRIE, *Direct Current Power for Sustainable Factories*, 2020). In data centers, centralized DC distribution can increase usable rack power density per square meter.
3. Organizations including UL Solutions, Current/OS, and the Open Direct Current Alliance are helping close LVDC gaps in the range of 800 to 1,500 V. The consequences of fragmentation are visible in industry surveys: In a 2025 Fraunhofer and University of Stuttgart survey of German manufacturers, 69% identify the lack of standards as the leading barrier to DC adoption, and 81% rate the regulatory framework as poor or average.
4. The expertise shortage is well documented. The International Energy Agency reports that more than 60% of energy companies face labor shortages, and the Institute of Electrical and Electronics Engineers and Kearney estimate that the global power sector will require an additional 450,000 to 1.5 million power engineers by 2030. In the 2025 Fraunhofer and University of Stuttgart survey, the lack of technical experts ranks as the third most cited barrier to DC adoption, at 38%. The missing capabilities include DC system architecture, DC protection and fault-coordination design, multivendor DC integration, and commissioning of live DC installations.

Where the Economics Work Today and Where They Are Headed

Standards and workforce gaps will shape the pace of adoption, but they do not erase the underlying business case. The key is to identify settings where DC already creates enough value to justify investment versus settings where the economics still depend on scale, application, and timing.

Adoption begins where power density is highest and loads are predominantly DC-native already. That puts data centers and automation-intensive industrial plants at the forefront, followed by large commercial buildings, where DC can pay off if used in the right configurations. The hybrid DC home remains a longer-term destination rather than a near-term market. (See Exhibit 3.)

We quantified four reference cases by applying a "greenfield delta" method: comparing otherwise equivalent new-build AC and DC facilities and isolating the difference in capital and operating costs. The analysis offsets the additional investment required for DC against the cost of the AC components it replaces. A capital expenditure advantage therefore means that the DC facility is cheaper to build overall, not merely that its subsequent savings offset a higher upfront outlay. The methodology holds today's equipment prices, electricity tariffs, and regulations constant. Expected changes in technology, policy, and market conditions inform the strategic outlook but are excluded from the modeled returns.

The geographic choices for the reference cases are not random. The data center case is based on Northern Virginia, one of the world's largest and most grid-constrained data center markets. The industrial, commercial-building, and residential cases use European parameters, which are well documented and provide a conservative test of the effects of energy prices and grid reliability.

The underlying sources of value—lower conversion losses, superior energy recuperation, reduced copper requirements, and more efficient integration of batteries and solar—apply globally. However, their financial impact varies with local electricity prices, grid charges, reliability, and other market conditions.

Data Center

The AI buildout has created a power-density problem that conventional AC distribution cannot solve. AI training depends on highly parallel processing, which requires large numbers of processors to sit close together. That design concentrates more power in less space and leaves little room inside the rack for power conversion or energy storage. As rack density increases, ride-through time collapses—from roughly 20 milliseconds in a conventional rack to only a few milliseconds at 1 MW—and conventional AC distribution and protection become increasingly impractical. The first response was the sidecar—a separate power module mounted alongside the rack. The next is to distribute 800V DC directly to the rack; this is the architecture that Nvidia's Kyber racks will use starting in 2027.

Compounding the urgency of DC adoption is a second constraint: limited grid capacity. Globally, the International Energy Agency projects that data centers' electricity consumption will more than double—from 415 terawatt-hours in 2024 to 945 terawatt-hours by 2030, even as grid connection waiting times extend to as much as seven years in Northern Virginia and as much as ten years across Europe's major hubs. For a data center operator committing capital today, the approved grid connection is often the binding constraint, and DC extracts the most value from it. (See Exhibit 4.)

EXHIBIT 3

DC Adoption Begins Where Power Density Is Highest

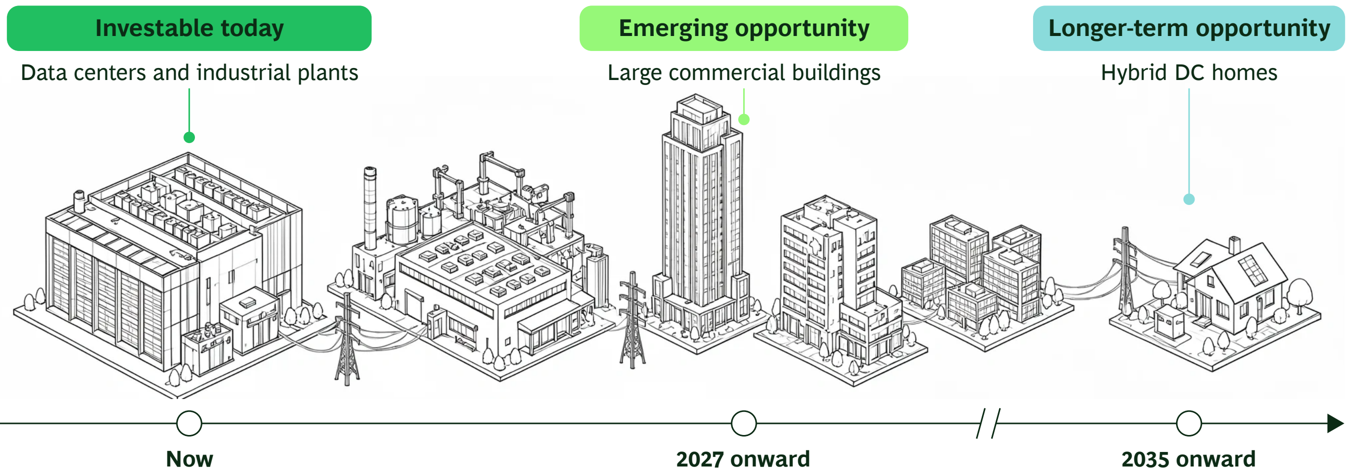

Source: BCG analysis.
Note: DC = direct current.

EXHIBIT 4

Illustrative Hybrid Data Center Architecture

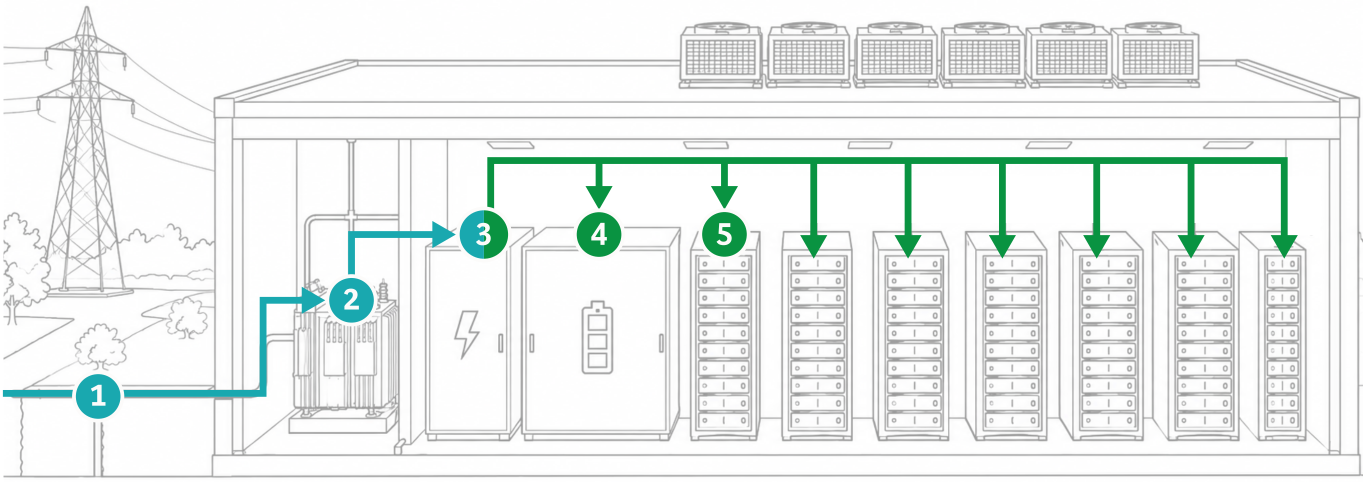

Source: BCG analysis.
Note: AC = alternating current; DC = direct current; MV = medium voltage; LV = low voltage.

Hyperscalers have been the earliest adopters of 800V DC because it removes the power constraint on scaling AI compute through 2030. The economics point in the same direction. Centralizing power conversion improves grid-to-rack efficiency by 5%. It also frees roughly 3,000 square meters of floor space on our 100MW reference campus, which is equivalent to about one-fifth of the area that a sidecar design would require. These gains matter because both power and space set ceilings on compute capacity. The efficiency improvement releases 5 to 6 MW for revenue-generating IT load, while the freed floor space provides room to deploy it. Together, the additional capacity is worth about $92 million in annual revenue and $35 million in operating profit. That is why the argument for 800V DC is fundamentally about capacity rather than efficiency.

The capital case is also positive. The drop in copper content of up to 45% is the largest source of savings, although the higher cost of rack-level DC protection offsets part of the benefit. The overall result is a net reduction of $3.5 million to $6 million in one-time capital expenditure. In settings where the data center does not fully monetize the additional capacity, lower conversion losses reduce annual operating costs by approximately $3.5 million a year. Batteries connected to the 800V bus through a simpler conversion stage provide several minutes of ride-through, which improves site resilience while also shaving peak demand on the grid connection. DC therefore lowers capital costs from the outset and continues to generate operating savings over the life of the facility.

Density determines when facilities need to move to DC, but the architecture they replace determines the size of the economic gain. Above roughly 200 kW per rack, conventional AC distribution becomes increasingly impractical, which is why AI hyperscalers are moving first. Many colocation and enterprise facilities running lower-density cloud and storage workloads can remain on AC for years. But colocation providers and neoclouds that are adding AI capacity face a different combination of problems: rack densities that require a new architecture and legacy power systems that hyperscalers have already moved beyond. Replacing a low-voltage uninterruptible power supply with centralized conversion and 800V storage can reduce capital expenditure by up to $35 million for a 100MW facility and cut the annual energy bill by as much as 9%. The incremental efficiency gain for such facilities can therefore be even greater than that for hyperscalers.

Industrial Facility

The modern factory is already partly DC, even if it doesn't look that way from the outside. Industrial VFDs, which control the motors that power conveyors, robots, and automated handling equipment, are DC-native internally. DC distribution simply consolidates these devices onto a shared plant-wide DC network, eliminating the conversion stages that currently happen inside each device. (See Exhibit 5.) The result is a business case that is positive from the first day of operation.

EXHIBIT 5

Illustrative Hybrid Industrial Plant Architecture

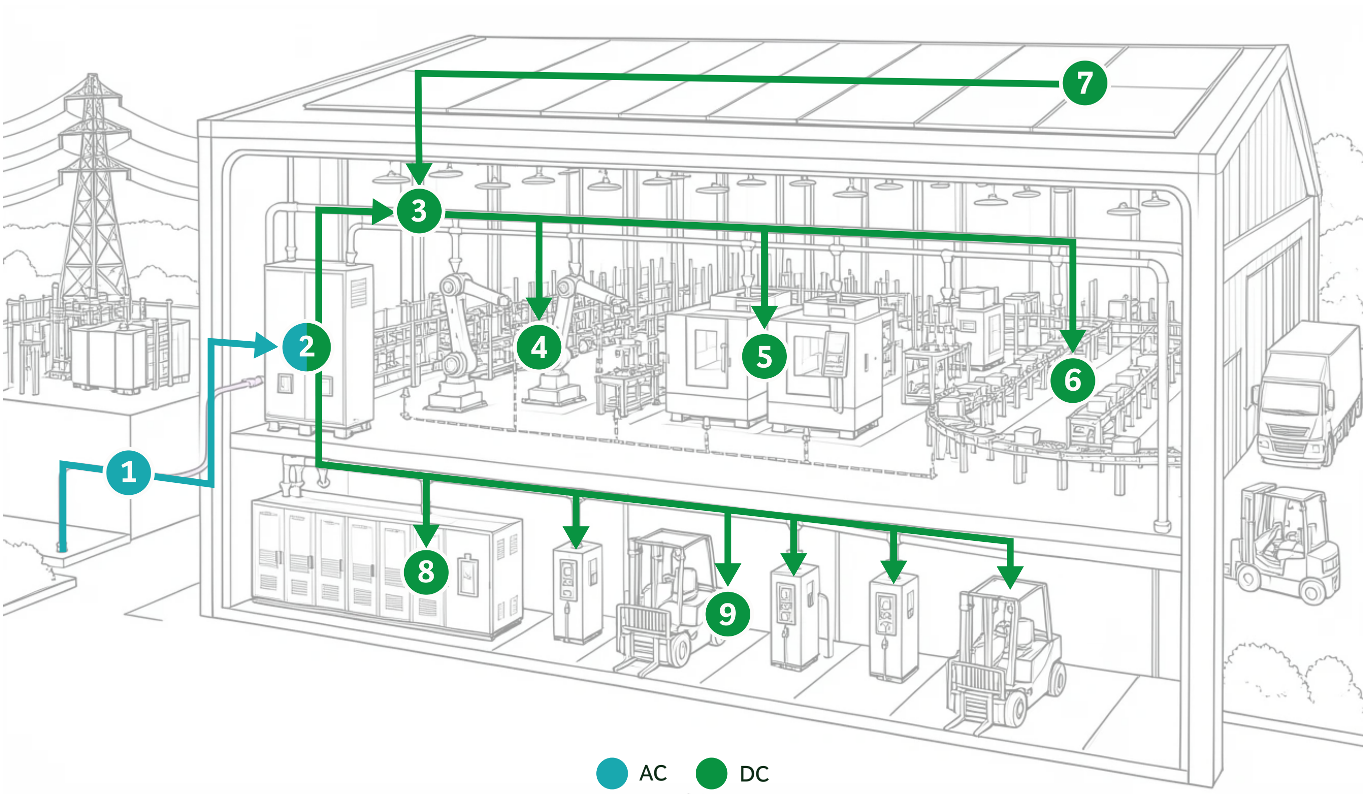

Source: BCG analysis.
Note: AC = alternating current; CNC = computer numerical control; DC = direct current; MV = medium voltage; PV = photovoltaic.

In our 10MW automation-intensive reference plant, DC delivers about $1.2 million in one-time capital expenditure savings, gained through freed floor space, eliminated power distribution units, and reduced use of copper. Annual operating expense savings are roughly $0.5 million, driven primarily by lower conversion losses and lower peak-demand charges through a DC-coupled battery; recovered braking energy is a further contributor.

The economics are strongest in settings where drive intensity, recuperation potential, power-quality needs, and process criticality converge. Automotive body-in-white lines, paper and metal manufacturing, and automated high-bay logistics are examples.

The value architecture holds globally, but the leading driver varies by region. The case is anchored in Germany, where high industrial power prices (modeled at €0.185—or $0.21—per kWh) make the efficiency case most visible. In the US, cheaper power (modeled at $0.09 per kWh) shifts the leading role to demand-charge savings and resilience: the US grid averages roughly 110 minutes of annual interruption versus 11.7 minutes in Germany. In Asia-Pacific, automation density leads, with 74% of global industrial robot installations occurring in the region in 2024.

Adoption will begin with greenfield facilities, in which builders can design the electrical architecture from scratch. Retrofitting existing plants becomes increasingly attractive as electrical infrastructure replacement cycles reach their end-of-life stage and as equipment costs fall.

Commercial Building

Modern commercial buildings are already integrated energy systems in everything but wiring. Rooftop solar, battery storage, EV charging, LED lighting, motor-driven climate-control systems, and elevators are all DC-native, yet conventional AC architecture forces each of these to pass through its own conversion stage. EV charging is the fastest-growing load and requires the most active management. In many cases, it is more economical to add on-site storage than to enlarge the grid connection. A hybrid DC design consolidates the large loads onto a shared 350V DC backbone while tenant plug loads remain on AC—an architectural decision that eliminates most of the conversion stages at once. (See Exhibit 6.)

EXHIBIT 6

Illustrative Hybrid Commercial Building Architecture

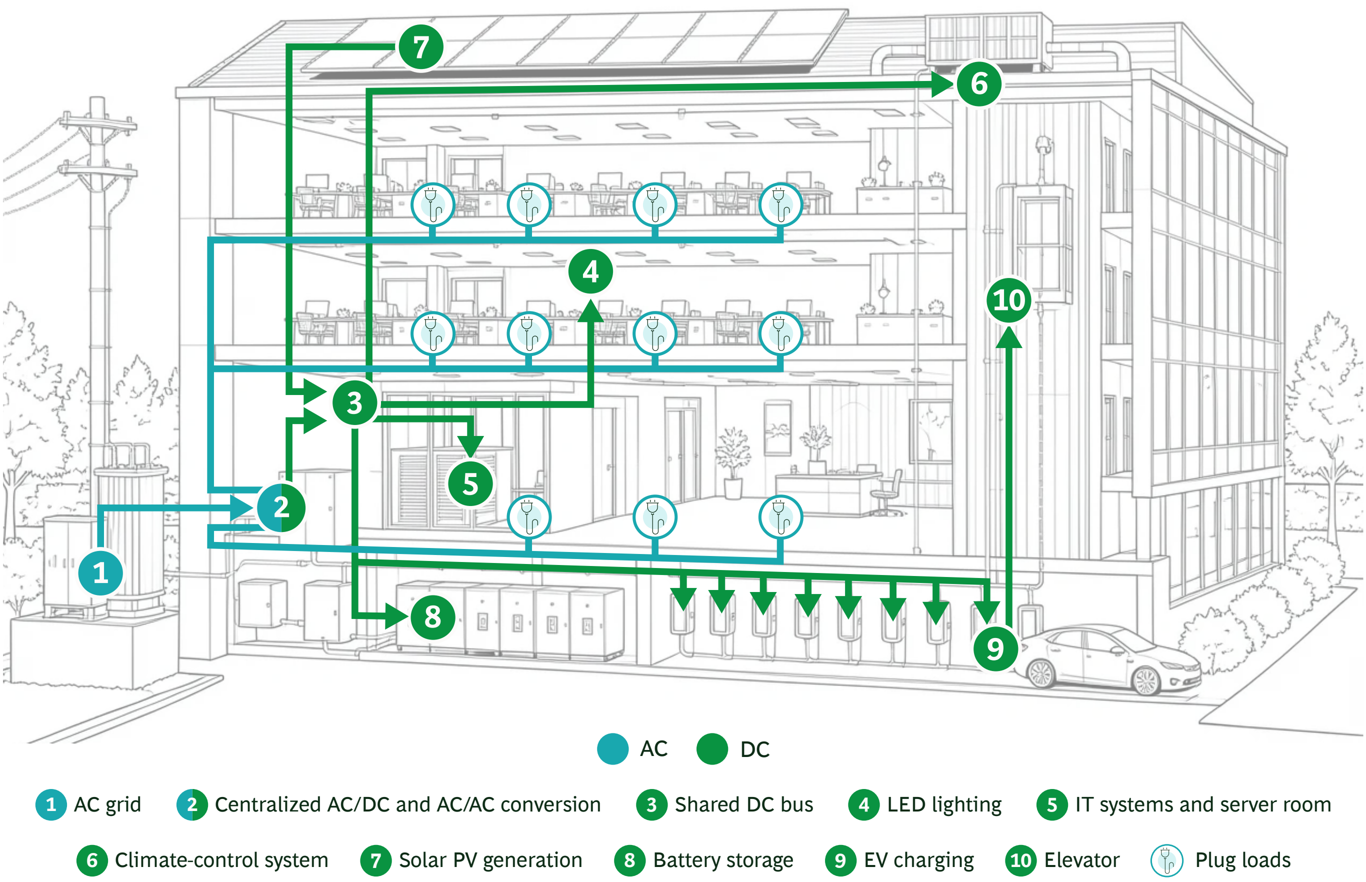

Source: BCG analysis.
Note: AC = alternating current; DC = direct current; EV = electric vehicle; PV = photovoltaic.

The economics already work at a modest scale. Our reference 10,000-square-meter office building consumes roughly 1 GWh per year. It captures approximately $46,000 in one-time capital expenditure savings and roughly $20,000 in annual operating expense savings—about 7% of the electricity bill. Most of the annual benefit comes from reducing conversion losses across LED lighting, IT, and climate control and from using a DC-coupled battery to reduce peak demand charges. Higher solar self-consumption and elevator energy recovery provide additional savings. The office is therefore less expensive to build and to operate.

Unlike a single home, a commercial building has the scale to enable these savings to support a measurable case today. Value concentrates in facilities where large DC-native loads, onsite generation and storage, and critical uptime converge: hospitals and laboratories, where supply resilience is paramount; hotels, airports, and office campuses, where energy intensity is high; and high-power EV charging hubs. Live pilots using the Current/OS open standard in Lille, France, and Delft, the Netherlands, demonstrate the architecture. Commercial facilities are realistically the most likely first building adopters, bridging today's industrial DC applications and tomorrow's hybrid home.

Residential Building

The DC home is the final destination of this trajectory. The single-family reference home runs its largest electrical sources and loads—solar panels, battery storage, the heat pump, and the EV charger—on a shared 350V DC circuit, while ordinary plug-in devices remain on AC. (See Exhibit 7.) For the occupant, nothing changes.

Today the economics do not justify the investment case. The hybrid architecture saves approximately $300 per year despite a capital expenditure premium relative to an equivalent AC home, and parity is likely as components scale. The gains come from lower conversion losses across the major DC loads and from increasing solar self-consumption above its already high 70% baseline. The home is an illustration; the household case alone does not justify adoption.

The systemic case is more important. For an individual home, the annual reduction in emissions is small—about 0.4 metric ton of CO_2e at today's grid mix—and will decline as electricity generation decarbonizes. But across Germany's roughly 13.5 million single-family homes, that amounts to approximately 5.4 million metric tons in the first year, equivalent to roughly 5% of the country's annual building-sector emissions. Taking into account the declining carbon intensity of the grid, the cumulative benefit would be on the order of 54 million metric tons over 20 years.

No household will build a DC home for the carbon benefit alone. The aggregate impact makes residential DC a policy-relevant lever rather than, for now, a compelling consumer proposition. But two shifts that are already registering on policy timelines will narrow the gap in household economics. First, starting in January 2027, newly installed EU charging points must support ISO 15118-20, the standard that enables bidirectional power flow and removes the principal protocol barrier to using an EV as a home battery. An EV typically offers four to six times the storage capacity of a dedicated home battery at a lower cost per usable kilowatt-hour because the buyer purchases the battery pack as part of the vehicle. Second, time-of-use tariffs are becoming more prevalent, rewarding households for buying electricity when it is inexpensive and using it later. A DC home retains more of that price advantage because charging and discharging require fewer conversion stages.

EXHIBIT 7

Illustrative Hybrid Home Architecture

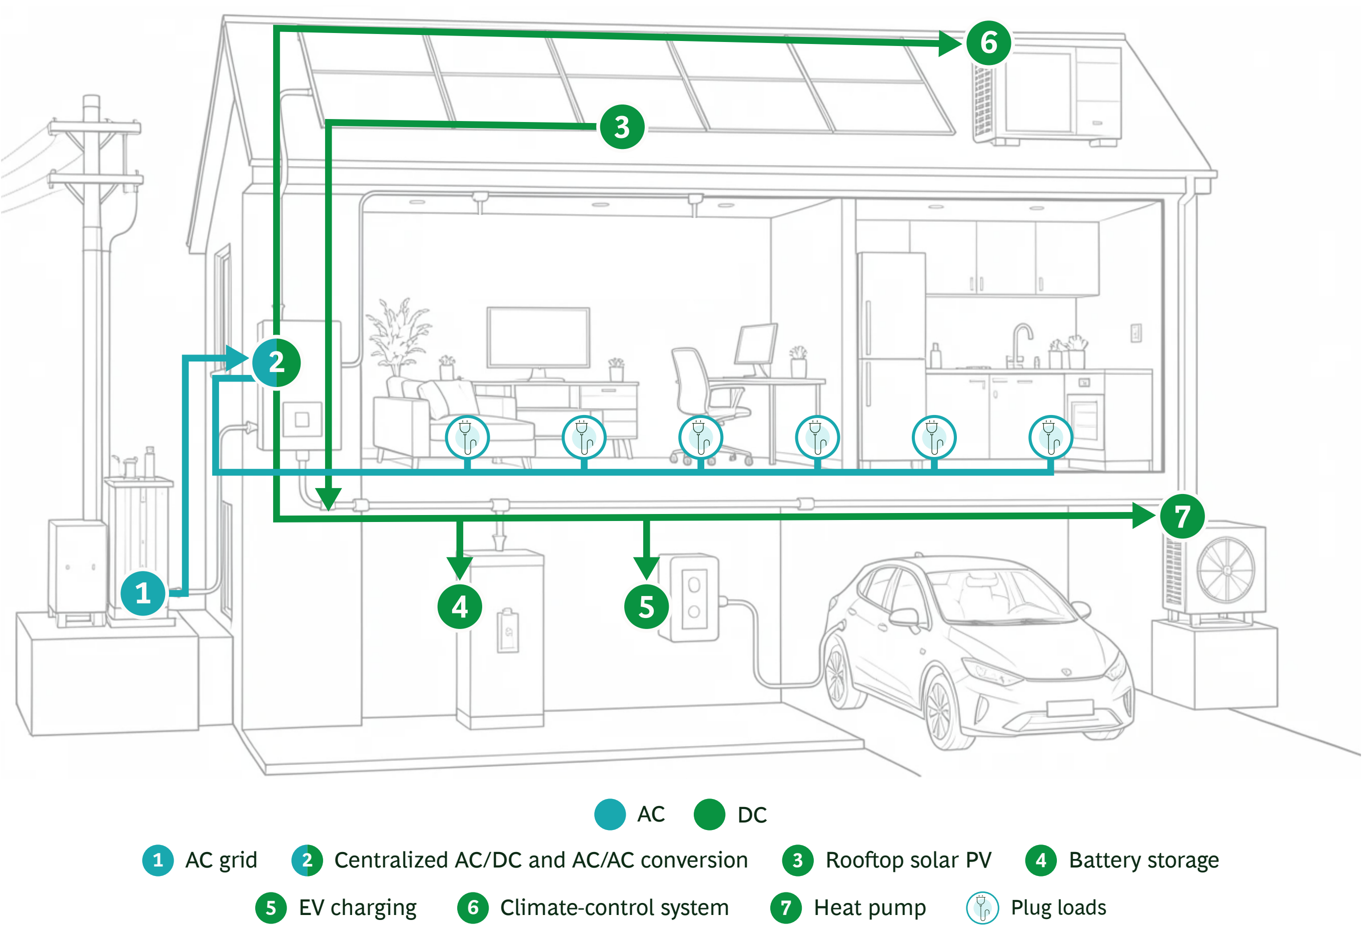

Source: BCG analysis.
Note: AC = alternating current; DC = direct current; EV = electric vehicle; PV = photovoltaic.

What Will Determine the Transition's Pace

The technology questions are largely answered, and the economics are increasingly compelling. The pace of adoption will therefore depend on how quickly the surrounding ecosystem evolves: from demand creation and manufacturing to regulation, standards, and workforce development. Four factors will play key roles in determining the speed of the transition. (See Exhibit 8.) They are as follows:

- **Demand and supply need to scale together.** The learning curve for power semiconductors depends on committed volume. Framework agreements, buyer consortia, and standardized specifications from hyperscalers, industrial companies, and large developers provide manufacturers with the confidence to industrialize production ahead of demand. The transition accelerates fastest when demand creation and supply expansion occur in parallel rather than sequentially.

- **Engineering friction must fall.** The core building blocks—including 800V DC rack power, 350V DC facility distribution, and DC-coupled storage—are already commercially available. The remaining challenge is to offer them as integrated systems and reduce the engineering effort necessary to deploy them. Fully specified reference architectures, certified design requirements, and multivendor interoperability standards shorten implementation cycles and lower project risk. Industry initiatives such as Current/OS and the Open Direct Current Alliance are moving in this direction.

- **Regulation and standards can converge.** Although convergence is crucial, neither regulation nor standards needs to wait for the other. Industry baselines already exist, from Current/OS to the Open Compute Project's 800V DC specifications. Since December 2025, Ireland has required every new data center connection to provide dispatchable generation or storage equal to its import capacity. This offers a blueprint that the European Union could replicate within a single legislative cycle, and comparable approaches are available to US state regulators (which exercise primary authority over electricity regulation in that country) and Asian authorities. Regulation can advance ahead of international harmonization, but it cannot substitute for it: scaling requires codified low-voltage DC defaults and aligned voltage bands. The speed at which these two tracks converge will have materially different competitive implications for early- and late-moving jurisdictions.

- **The workforce should develop specialized DC capabilities.** The principal workforce constraint is not the number of engineers but the availability of expertise in DC system architecture, protection, multivendor integration, and commissioning. As was the case with the early expansion of solar PV, deployment is likely to be constrained by an insufficient number of qualified installers before it is constrained by technology-related issues.

These four factors are mutually reinforcing rather than sequential. Treating any one as a strict precondition—a common flaw in early DC roadmaps—delays the entire chain rather than mitigating risk.

EXHIBIT 8

Four Reinforcing Conditions Will Determine the Pace of the Transition

Demand and supply scale together	Engineering friction falls	Regulation and standards converge	Workforce develops DC capabilities
Committed volumes give manufacturers the confidence to industrialize production	Integrated systems, reference architectures, and common standards reduce deployment effort and project risk	Regulation can move first, but scaling still requires internationally harmonized standards	Specialized expertise in DC design, protection, integration, and commissioning expands deployment capacity

Source: BCG analysis.
Note: DC = direct current.

The CEO Agenda: Shaping the Transition, Not Inheriting It

The ecosystem conditions that will determine the pace of the transition lie largely outside the control of any single company. Even so, they point to strategic imperatives for companies that supply hybrid-power technologies and for those that deploy them. (See Exhibit 9.) The strategic opportunity is to help shape the transition.

For Electrification and Automation Suppliers

Several moves are essential for a supplier that aims to establish early leadership:

- **Capture emerging revenue pools.** Hybrid power will create successive waves of demand: hyperscale AI infrastructure first, automation-intensive industry next, and commercial and residential buildings over time. Suppliers should establish positions in each market as it emerges, leading in hyperscale AI infrastructure, securing industrial opportunities where DC delivers day-one savings, and preparing to scale when commercial and residential buildings become the next growth engines.
- **Reinvent the power-conversion stack.** Suppliers should build integrated portfolios that span capabilities, such as 800V DC active front ends, DC switchgear, battery-integrated DC buses, real-time energy-management software, grid-forming capabilities, digital twins, multiyear service offerings, and a certified DC workforce. The opportunity is to become a systems integrator of hybrid power rather than just a component supplier.
- **Match the go-to-market model to the opportunity.** Each market demands a different commercial model. Hyperscalers require deep co-development relationships. Industrial deployment depends on engineering, procurement, and construction firms and system integrators. Buildings will require wholesalers, installers, and OEM partnerships. Success will depend as much on commercial execution as on technology.
- **Move now to scale later.** The objective is to secure hyperscaler and industrial proof points in 2026 and 2027, scale commercial building deployment starting in 2027, and prepare for broader residential adoption in the next decade. Companies that establish market positions before deployment accelerates will shape customer expectations and industry economics rather than responding to them.
- **Build the value story.** Hybrid power is about more than efficiency. Suppliers should treat the hybrid architecture as an enabler of the energy transition and of AI infrastructure. Building recurring service revenue through multiyear contracts will reinforce the equity story for investors.

EXHIBIT 9

Strategic Priorities for Suppliers and Operators in the Hybrid Power Transition

Electrification and automation suppliers

- **Capture emerging revenue pools**
  Start with data centers and industrial plants, then scale to buildings
- **Reinvent the power-conversion stack**
  Sell integrated systems, not single components
- **Match the go-to-market model to the opportunity**
  Co-develop with hyperscalers; partner for buildings
- **Move now to scale later**
  Take early positions to shape the market
- **Build the value story**
  Seek recurring revenue to strengthen the equity story

Data center, plant, and building operators

- **Treat DC as a capacity and margin decision**
  Fit more computing power per grid connection; improve day-one margin in plants
- **Decide how much DC capability to own**
  Retain strategic capabilities; source execution under performance guarantees
- **Co-develop the market with suppliers**
  Shape the specifications before they become fixed
- **Align capital cycles with the transition**
  Avoid locking long-lived assets into AC architectures
- **Extend the economics beyond individual projects**
  Value DC as strategic infrastructure across the portfolio

For Data Center, Plant, and Building Operators

On the deployment side, the priority is to preserve strategic flexibility while the architecture takes shape:

- **Treat DC as a capacity and margin decision.** The source of value differs by application. Hyperscale data centers can fit more compute behind an approved grid connection. Industrial plants can improve margins from the first day of operation. Large commercial buildings can establish a DC backbone for major loads. Leaders should evaluate options in terms of capacity, margin, and asset performance—not simply energy efficiency.

- **Decide how much DC capability to own.** Suppliers will typically handle integration and commissioning, so adopters don't need to build full in-house DC engineering teams. They do, however, need to determine which capabilities to retain internally—such as system specification, reference designs, and operational readiness—and which to source under clear performance guarantees.

- **Co-develop the market with suppliers.** Reference architectures, framework agreements, and buyer consortia can compress the learning curve, reduce project-specific engineering, and mitigate risks associated with multivendor integration. Early participation also gives adopters greater influence over the technical specifications and procurement practices that will later become industry norms.

- **Align capital cycles with the transition.** Facilities designed for delivery during the period from 2027 to 2030 will remain in operation for decades. Adopters should avoid locking those assets into AC architectures that may require mandatory retrofitting before the end of their useful lives as power densities, standards, and regulatory requirements evolve.

- **Extend the economics beyond individual projects.** Planners should reframe DC as strategic infrastructure for capacity, resilience, and decarbonization. Adopting that broader perspective will enable them to make investment decisions on the basis of portfolio-wide and system-level value, not just with regard to the payback of an individual facility.

The companies that move in the next two to three years will set the cost curves, voltage defaults, and installation standards of the next decade. Those that wait will inherit the design choices that others make. For suppliers, key actions include capturing emerging revenue pools and establishing the technologies, capabilities, and commercial positions that define the market. For adopters, best practice is to treat hybrid power as strategic infrastructure and avoid decisions that lock long-lived assets into yesterday's architecture. In both cases, the advantage is asymmetric: it will accrue to those that help shape the transition rather than simply respond to it.

About the Authors

Emmanuel Sissimatos
Managing Director & Senior Partner
Düsseldorf
sissimatos.emmanuel@bcg.com

Monika Saunders
Managing Director & Partner
Zurich
saunders.monika@bcg.com

Marc Schmidt
Partner
Düsseldorf
schmidt.m@bcg.com

Tom Piendl
Associate
Stuttgart
piendl.tom@bcg.com

Acknowledgments

We gratefully acknowledge the contributions of ABB's experts who shared their insights, technical knowledge, and practical experience throughout the development of this paper.

About Boston Consulting Group

Boston Consulting Group bridges the gap between ambition and outcomes for the world's leading companies and organizations. We are built for this era of unprecedented change—bringing strategic clarity rooted in over 60 years of deep domain knowledge, combined with applied AI shaped by our practitioners. BCG works shoulder-to-shoulder with CEOs across industries and geographies to deliver transformative impact at scale: stronger returns, transferred capabilities, and change that sticks. For more information, visit bcg.com.

About ABB

ABB is a global technology leader in electrification and automation, enabling a more sustainable and resource-efficient future. By connecting its engineering and digitalization expertise, ABB helps industries run at high performance, while becoming more efficient, productive and sustainable so they outperform. At ABB, we call this 'Engineered to Outrun'. The company has over 140 years of history and around 110,000 employees worldwide. For more information, visit www.abb.com.