# Accelerating AI Data Center Deployment with Software-Defined Medium-Voltage Switchgear

White Paper 303

Version 1

## Data Center Research & Strategy

By John Slater
Wendy Torell
Venanzio Paolo Ferraro
Maxime Durand
Amita Sonar

## Executive summary

As AI data centers grow in scale and urgency, power infrastructure has become a critical constraint on time-to-power. Yet medium-voltage (MV) switchgear is often still delivered through engineer-to-order (ETO) approaches characterized by project-specific engineering and integration, extensive control wiring, and labor-intensive commissioning. This paper examines how software-defined MV switchgear simplifies deployment by shifting from customized, hardware-bound architectures to more standardized platforms enabled by software-configured functionality. The analysis in this paper demonstrates how this shift can reduce the end-to-end deployment schedule by 64%, accelerating the path to energized AI capacity. While the primary value today lies in accelerating deployment, the architecture can also improve operational flexibility and scalability, reduce personnel exposure during local operations, and reduce embodied carbon.

RATE THIS PAPER ★★★★★

# Key takeaways

1. **Traditional MV switchgear deployment can constrain time to power for AI data centers.** Compute capacity cannot begin producing tokens or generating value until the supporting electrical infrastructure is energized. Project-specific engineering, procurement, manufacturing, installation, and commissioning can make MV switchgear a critical path item.

2. **Software-defined MV switchgear enables configure-to-order delivery.** A standardized hardware platform and reusable configuration templates allow site-specific requirements to be addressed through software rather than changes to hardware, physical integration, and wiring.

3. **In the comparison analyzed in this paper, software-defined MV switchgear reduced deployment time by 64%.** The schedule decreased from 31.5 weeks for traditional ETO switchgear to 11.5 weeks for the software-defined approach. Although the specific reduction will vary by project size and complexity, the comparison demonstrates the potential for substantial schedule improvement.

4. **Time savings from deploying software-defined switchgear extend across both factory and field activities.** The combined engineering, procurement, and manufacturing interval decreased from approximately 26 weeks to 8.5 weeks, while installation and commissioning decreased from 5 weeks to 2.5 weeks.

5. **Faster deployment via software-defined switchgear does not mean compromising protection system resilience.** Essential backup protection and breaker actuation remain local at each panel, while redundant communication paths preserve communications if one network path fails.

6. **Software-defined architectures can provide benefits beyond initial deployment.** Reusable configurations support phased expansion and future changes, spare parts management is simplified, centralized interaction can reduce the need for personnel to approach energized equipment during routine local operations, and fewer control wires and terminals reduce material use and its associated carbon footprint.

# Introduction

As AI data centers grow in scale and urgency, the ability to deliver power infrastructure faster has become a critical determinant of time-to-power. Large campuses are being developed simultaneously across multiple regions, placing significant pressure on grid interconnections, equipment supply chains, and the engineering resources required to design and commission complex electrical systems. Even projects with secured land, permits, and financing can face years-long delays before power infrastructure is ready to support operations[1]. These constraints have elevated deployment speed from a project objective to a strategic imperative.

Much of this pressure originates outside the operator's control. Utilities face growing interconnection queues, manufacturers continue to manage long lead times for critical equipment, and experienced engineering and commissioning personnel are stretched across concurrent projects worldwide. Beyond these market constraints, deployment speed is also affected by how power infrastructure is designed, integrated, and commissioned.

AI data center campuses are reaching gigawatt scale. At these capacities, more extensive medium voltage (MV) infrastructure is needed to distribute power. Depending on the campus architecture, distribution may occur at 13.8 kV, 34.5 kV, or both. Data centers may also integrate on-site generation or energy storage to address grid constraints, strengthen reliability, and manage load changes. This evolution does not eliminate low-voltage distribution near the IT load; rather, it increases the scale and complexity of the upstream MV system supplying it.[2] Yet MV switchgear delivery often remains rooted in engineer-to-order (ETO) practices, relying on project-specific protection schemes, extensive hardwired interconnections, and labor-intensive commissioning that can lengthen the path to energized capacity.

For AI data centers, delivering MV power infrastructure through a more efficient, repeatable, and lower-risk process creates substantial business value by bringing compute capacity online sooner. Software-defined approaches to MV switchgear can accelerate time-to-power by enabling a configure-to-order (CTO) model in which standardized hardware and proven design and configuration templates are adapted to site requirements rather than re-engineering their application and integration for each project. This shift simplifies design and engineering, parallelizes work, eliminates project-specific wiring, reduces material requirements, and streamlines validation, installation, and commissioning. As **Figure 1 illustrates**, the example deployment analyzed in this paper reduced the end-to-end schedule from 31.5 weeks to 11.5 weeks, or approximately 64%. This paper examines the architectural changes that enable these savings, and quantifies their impact across engineering, manufacturing, installation, and commissioning, providing a faster path to energized AI capacity and token production.

**Figure 1**

*MV switchgear deployment timeline comparison for the analyzed 20-panel traditional vs. software-defined configurations*

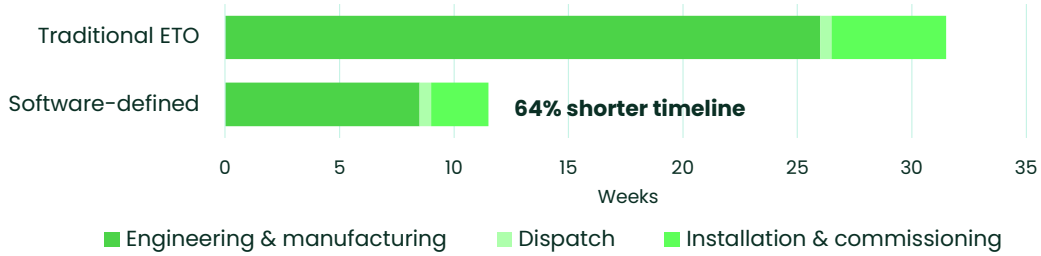


[1] DOE, Recommendations on Powering Artificial Intelligence and Data Center Infrastructure, 2024, connection requests for 300–1,000+ MW hyperscale facilities have lead times of one to three years and are stretching local grids' ability to supply power.

[2] ASHRAE, "Integrated Design Principles," AI Data Center Energy Performance Framework, sections "Power and Cooling Integration" and "Medium-voltage distribution".

## Traditional MV switchgear and the limits of the ETO model

MV switchgear performs critical functions within the data center electrical system, including fault protection and isolation, protection coordination, metering, source transfer, and control and automation. These functions influence component selection and overall system design. Because they directly affect availability and safety, MV systems have historically been designed around site-specific requirements and validated through extensive testing prior to energization.

In traditional MV switchgear, component selection depends heavily on project-specific characteristics and the associated protection coordination study.[3] Engineers first develop the required system data and select protection, metering, monitoring, and control devices. This selection is influenced by inputs/outputs (I/O) and communication for interlocking schemes, system performance, and preferred function types. They then specify the current transformers (CTs) and voltage transformers (VTs) from hundreds of available references or develop custom specifications as required. Managing the requirements of these functions and their interactions, together with the sizing and selection of the CT and VT references, is a complex and time-consuming process. Because these engineering decisions are prerequisites for manufacturing, the time required to define and approve them can create project execution bottlenecks. Protection settings, control logic, wiring, drawings, configuration files, supporting documentation, and validation procedures are subsequently developed or adapted for the specific site.

As illustrated in Figure 2, traditional ETO deployment involves coordinated activities across engineering, manufacturing, installation, and commissioning, with feedback between stages as designs, configurations, and field implementations are verified and refined. Existing standardization, prefabrication, and repeatable engineering and commissioning practices are already making the process faster. However, the unprecedented size and pace of AI deployments require these practices to support a much greater volume of work.

**Figure 2**

*In a traditional ETO model, project-specific requirements drive time-consuming activities through engineering, manufacturing, installation, and commissioning, increasing the effort required to deliver each new deployment.*

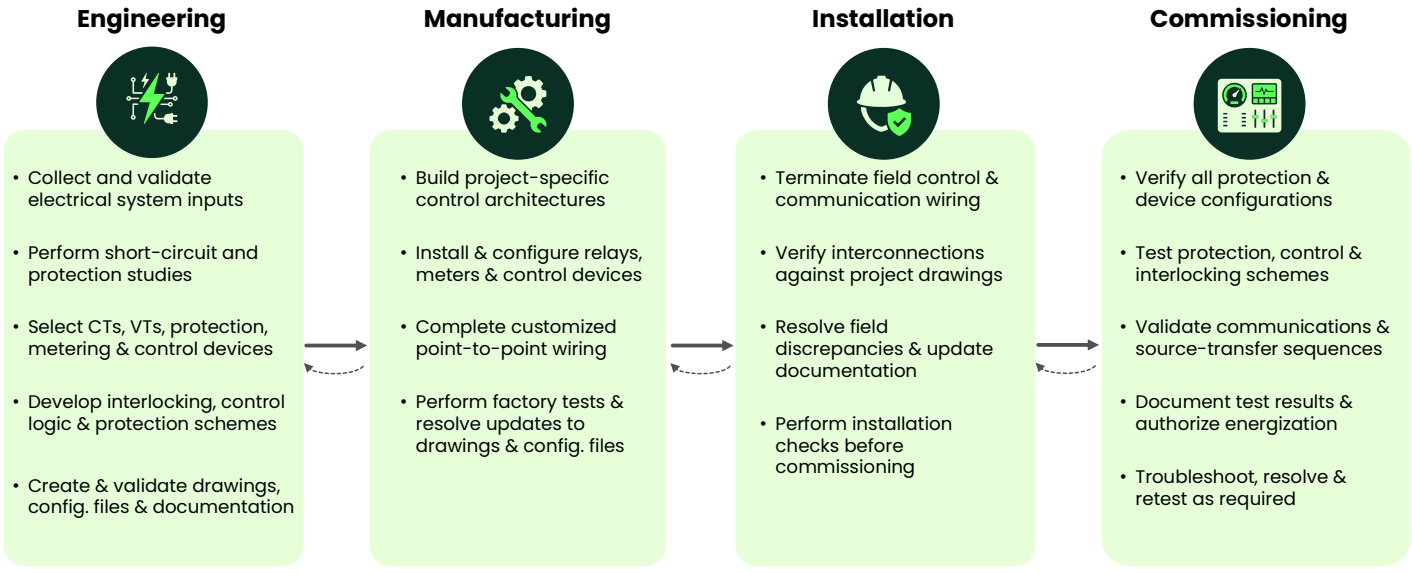


[3] Schneider Electric, "Why are MV component design and integration important?", May 2026

The question is not whether traditional MV architectures can perform their intended functions; they can. Rather, it is whether the current methods used to design, assemble, install, and commission MV switchgear can scale with the pace at which new AI capacity must be delivered. The following section examines how software-defined approaches simplify the underlying architecture and reduce deployment complexity.

## How software-defined MV switchgear simplifies deployment

The software-defined approach does not change the fundamental role of MV switchgear within the electrical system. The equipment and overall system must still comply with all applicable equipment, communication, protection, safety, and local code requirements. Instead, what changes is *how* the switchgear's functions are implemented and delivered, creating opportunities to simplify deployment.

### What is software-defined MV switchgear?

Software-defined MV switchgear uses standardized hardware to provide application-specific protection, automation, metering, monitoring, and control functions through software configuration rather than project-specific hardware selection and integration. Figure 3 compares how functions are implemented in traditional and software-defined switchgear.

**Figure 3**

*Attributes of a 20 panel software-defined MV vs. traditional MV switchgear. The same attributes apply to smaller lineups.*

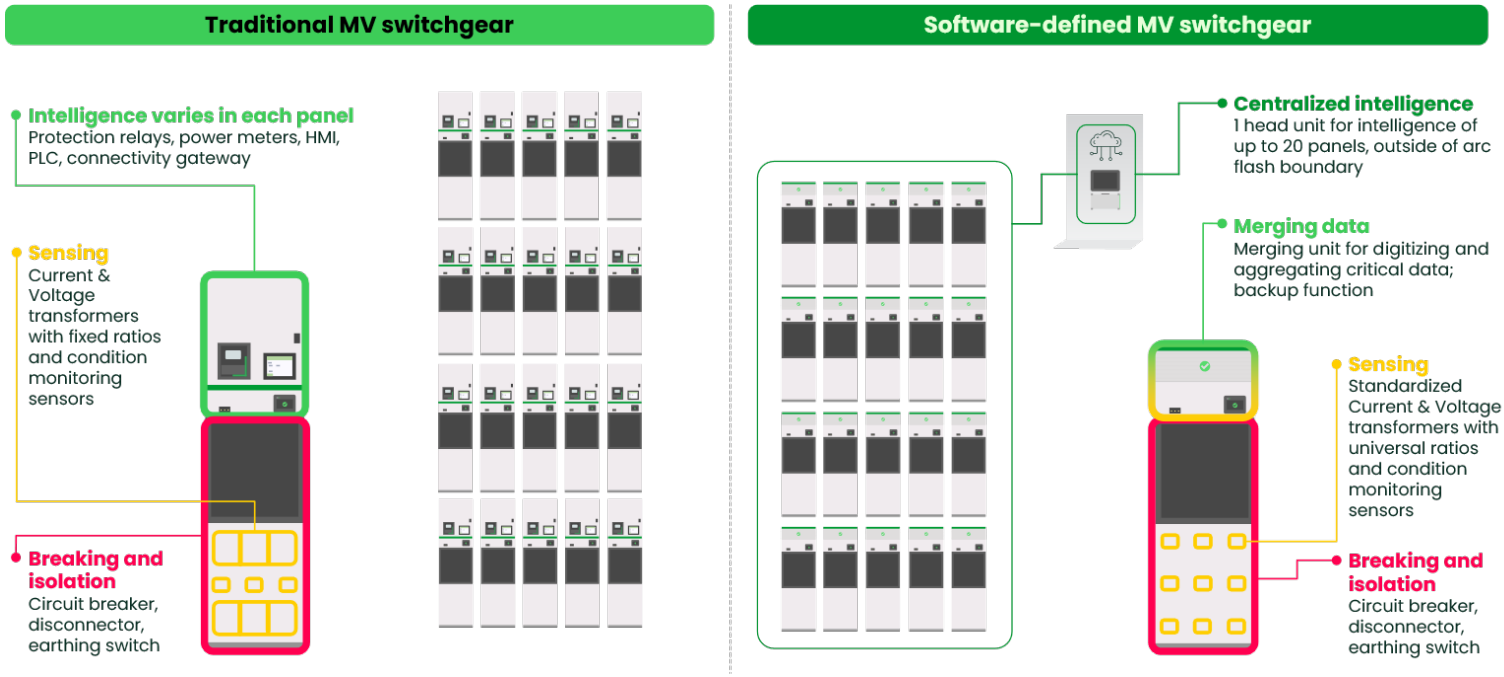


Each software-defined switchgear panel (also referred to as a cubicle) uses low-power current transformers (LPCTs), low-power voltage transformers (LPVTs), and a merging unit. Unlike conventional instrument transformers (CTs and VTs), a single variant of LPCT and LPVT covers the full range of switchgear protection and measurement requirements, without project-specific hardware variants. They are also substantially lighter than conventional CTs and VTs, simplifying handling and installation and reducing the overall material required for the sensing architecture. Figure 4 compares the size and weight of the low-power sensors with conventional instrument transformers.

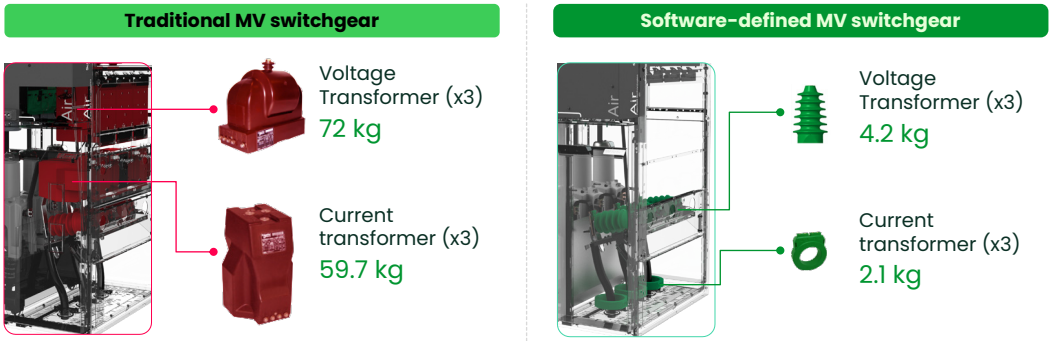


**Figure 4**

*Comparison of conventional instrument transformers and low-power sensors used in MV switchgear.*

The merging unit converts the analog measurement signals received from the LPCTs and LPVTs, together with other panel status and auxiliary signals, into digital data. This data is transmitted within the software-defined scheme using a closed-loop redundant Ethernet network, with IEC 61850 process bus communications, assuring the needed real-time latency.[4] This digital architecture reduces the point-to-point wiring and terminations traditionally required in the low-voltage (LV) compartment of each panel. The merging unit also hosts critical back-up functions, improving the resilience of the protection scheme. The examples in Figure 5 illustrate how a traditional LV compartment contains multiple protection, metering, and control devices with their associated wiring and terminals, whereas a software-defined design uses a merging unit and a simpler wiring architecture.

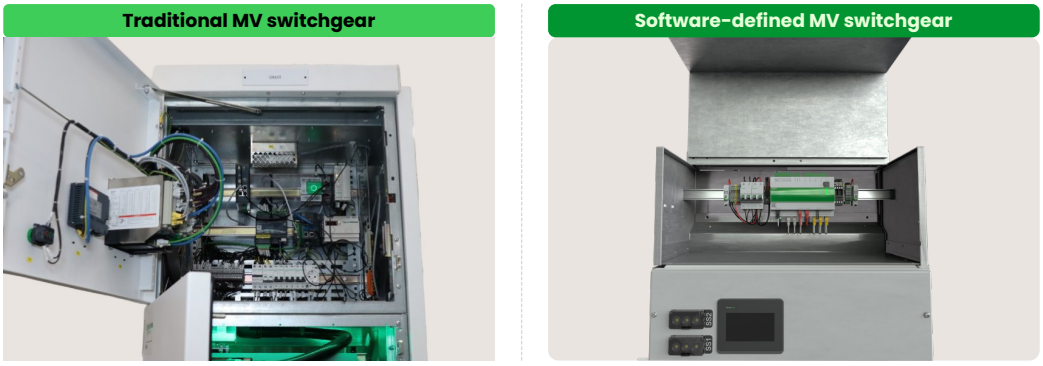


**Figure 5**

*Example traditional and software-defined LV compartments, illustrating the reduced device count, wiring, and terminations enabled by the software-defined architecture.*

Intelligence functions traditionally distributed across multiple protection relays, meters, monitoring and control devices in each panel can instead be hosted as software applications on a central computing platform, referred to in this paper as the head unit. This rugged industrial computer serves as the digital core of the switchboard, hosting, orchestrating, and communicating all protection, automation, metering, and monitoring functions for up to 20 panels. Communication upstream of this head unit can use a number of different protocols, such as IEC 61850, OPC-UA, or Modbus TCP.

## Enabling configure-to-order deployment

By standardizing the hardware architecture, this approach enables a CTO model that shifts project-specific work from hardware selection, physical integration, and hardwiring to software-based configuration. Reusable logic, configuration, and validation templates allow previously validated designs to be adapted across deployments without recreating the underlying architecture each time.

[4] International Electrotechnical Commission (IEC). IEC 61850: Communication Networks and Systems for Power Utility Automation. IEC Technical Committee 57. The IEC 61850 series defines communication protocols, data models, and engineering processes that enable interoperable communication among intelligent electronic devices (IEDs) in electric power systems.

Protection studies, settings, interlocking schemes, control logic, drawings, documentation, and validation remain necessary. However, protection and control functions can be selected and configured in software rather than tied to project-specific relays, meters, and control devices. This work can proceed in parallel with manufacturing of the standardized equipment. Because the physical architecture is common and no project-specific connections are necessary, wiring diagrams and other engineering documents also become repeatable and faster to produce.

## Simplifying the implementation

As illustrated by the component and LV-compartment examples in Figure 4 and Figure 5, the software-defined approach can reduce equipment weight and simplify device and wiring arrangements. The software-defined approach does not fundamentally change the switchgear. Instead, it standardizes more of the hardware and interfaces used for sensing, protection, metering, and control. Functions previously distributed across multiple protection relays, meters, and control devices can now be hosted as software applications in the head unit. Measurements and status signals are digitized by the merging units, while many signals that previously required point-to-point wiring are exchanged over Ethernet.

This approach does not eliminate all devices or hardwiring. Connections required for primary power, local operation, manual breaker actuation, the merging unit, and selected backup or safety functions remain hardwired as appropriate. However, reducing project-specific devices and control connections simplifies cabinet layouts and decrease the quantity of control wiring, terminals, and related components.

We performed a comparative analysis of a 20-panel traditional ETO switchgear configuration and the equivalent software-defined switchgear configuration to illustrate the magnitude of these reductions, as shown in Table 1.[5]

**Table 1**

*Material comparison of control wiring and instrument transformers of a 20-panel switchgear lineup.*

** Assumes only 4 panels require VTs*

*** Assumes all panels are standardized with LPVTs*

| Metric | Traditional ETO | Software-defined | Reduction |
|---|---|---|---|
| Control-wire length[6] | 4,800 m | 635 m | 87 % |
| Number of terminals[6] | 12,000 | 1,120 | 91 % |
| Copper conductor mass[7] | 32.3 kg | 5.0 kg | 85 % |
| VT / LPVT mass | 288 kg* | 84 kg** | 71% |
| CT / LPCT mass | 1,194 kg | 42 kg | 96% |

The shift from point-to-point wiring in the traditional ETO approach to networked communications in software-defined switchgear substantially reduces control wiring, terminal requirements, and associated copper use. These reductions improve material efficiency at a time when copper demand is increasing across electricity networks, data centers, and other electrification applications.[8]

[5] Values based on comparison of a typical 34.5 kV, 2,500A MV switchgear for IEC applications (traditional and software-defined) and aligned with experience of data center architectures.

[6] Assumes traditional ETO panels each with 600 control wiring terminals [average 0.8 m wire length], and software-defined panels each with 56 terminals [average 1 m wire length] plus 3.75 m of Ethernet.

[7] Copper mass is based only on control wiring using 0.75 mm² stranded copper wire at 6.9 g/m, and 75m Ethernet wiring at 14.6 g/m (1.1 kg) added for software-defined panels. CT and VT wiring and other copper sources were not analyzed.

[8] International Energy Agency, Copper prices have hit record highs, but smelters face mounting strategic pressures, March 2026. The IEA identifies grids, data centers, transportation, construction, and industry as sources of increasing copper demand and highlights material efficiency as one measure for addressing potential supply constraints.

These changes simplify several of the time-consuming activities identified in Figure 2. Manufacturing teams procure, mount, and wire fewer devices. Installation teams terminate and verify fewer physical connections, reducing both the opportunity for wiring errors and the number of discrepancies that must be resolved.

Site-specific verification remains necessary, but commissioning teams can spend less time performing point-to-point continuity checks and troubleshooting wiring. Where software-based simulation of the configured MV system is performed in parallel, selected protection and control functions and their interactions can be validated before physical assembly, reducing on-site testing effort and de-risking the project at an earlier stage.

Taken together, these changes simplify work across engineering, manufacturing, installation, and commissioning. Figure 6 illustrates how standardized hardware production and application-specific software configuration can proceed in parallel, while reduced wiring and earlier validation decrease the work required later in the deployment process.

**Figure 6**

*Sequential ETO vs. parallel software-defined CTO approach. Traditional ETO includes a longer sequence of project-specific device selection, engineering, procurement, and wiring. The software-defined approach decouples standardized hardware manufacturing from application-specific configuration, letting these proceed in parallel before integration & validation.*

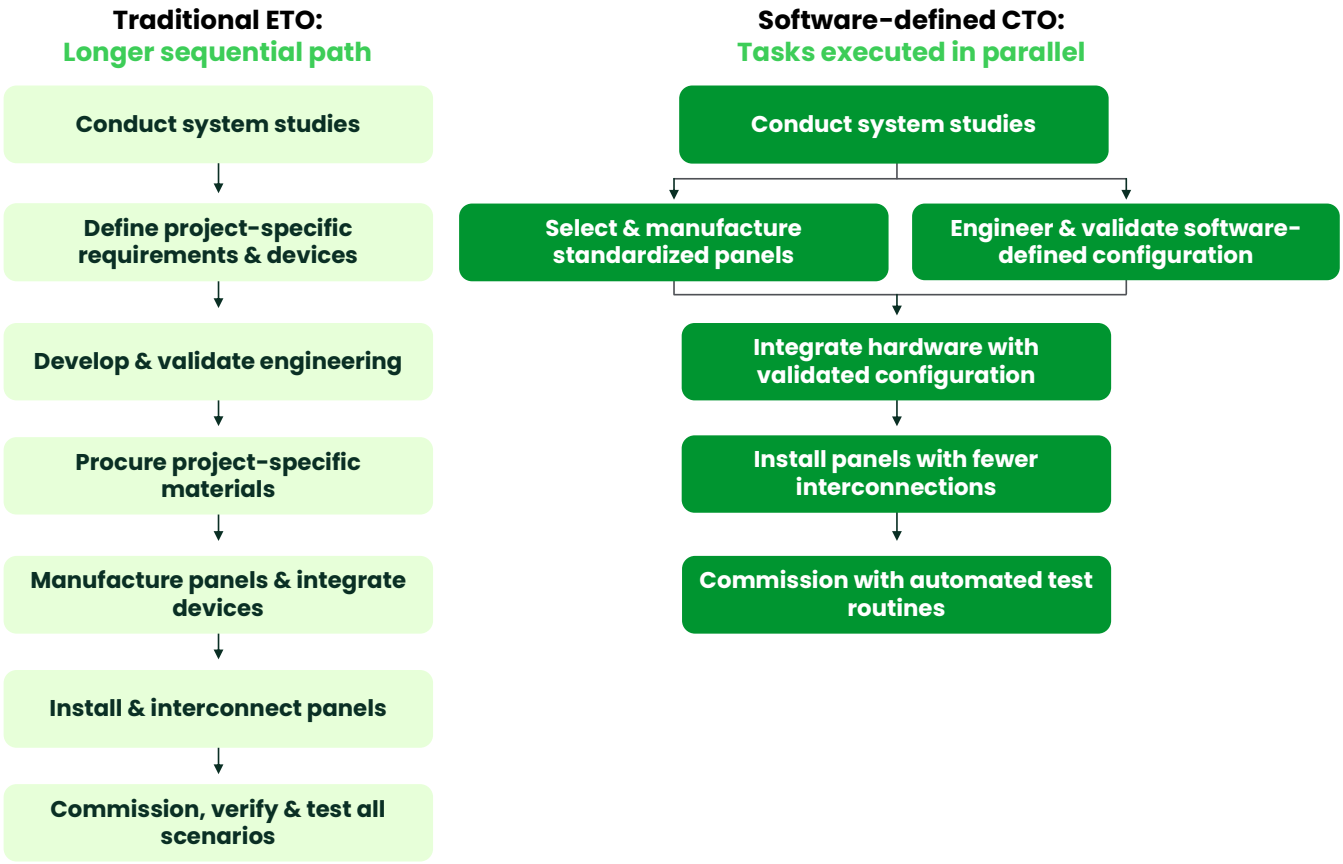

## Designing for protection system resilience

Centralizing does not mean that all protection functions are in a single location or along a single communication path. Selected backup protection must remain at the panel level within the merging unit, while breaker actuation, local operation, and safety functions retain direct connections. The allocation of functions should be determined through the protection and availability studies for the application.

Network-dependent functions also require communication redundancy. Parallel redundancy protocol (PRP)[9] and high-availability seamless redundancy (HSR)[10] are standardized approaches for maintaining communication following a single network-path or network-element failure. Their use does not conflict with retaining hardwired connections for selected local and backup functions.

In the architecture analyzed in this paper, backup protection, tripping, and manual breaker actuation are retained locally. Temporary loss of the head unit or communication network therefore does not eliminate the local ability to isolate a fault.

## Simplified condition monitoring

Condition monitoring has evolved to include sensors embedded in critical components that monitor equipment health and provide actionable data to maintenance teams. These sensors are typically combined with communication devices that collect and transmit condition data to centralized monitoring systems.

The software-defined architecture applies the same standardization principles, with primary intelligence centralized in the head unit. In this approach, sensors communicate through a Zigbee wireless network, and a single gateway aggregates data from up to 20 panels before forwarding the condition monitoring data stream to the head unit.

This architecture maintains the role-based separation of functions and other requirements for cybersecurity compliance, with protection and monitoring systems remaining logically segregated. It also reduces system complexity by reducing devices and unused access points, such as programmable logic controllers (PLCs) or other processing devices.

## Cybersecurity and configuration governance

Because software-defined MV switchgear relies on networked communications and software-configured functions, cybersecurity and configuration governance must be incorporated into the system design. As with conventional protection systems, engineering access should be controlled, protection-system networks appropriately segmented, and changes affecting protection, control, and interlocking subject to the same engineering validation applied to changes in conventional wiring.

Software-defined switchgear and its components should also be developed and maintained to ensure a secure operational environment using a Secure Development Lifecycle (SDL) aligned with IEC 62443 requirements. This incorporates core cybersecurity principles, including role-based access control, defense in depth, zero-trust architecture, perimeter protection and network hardening. For further

[9] Parallel Redundancy Protocol (PRP) connects devices to two independent networks and transmits duplicate frames over both, enabling uninterrupted communication if one network path fails. Defined in IEC 62439-3:2021

[10] High-availability Seamless Redundancy (HSR) transmits duplicate frames in opposite directions through a ring network, enabling uninterrupted communication if a network path or element fails. Defined in IEC 62439-3:2021

guidance, see the IEC 62351 series[11], the IEC 62443 series[12], and Schneider Electric White Paper 216, *Cybersecurity Guidance for Data Center Power and Cooling Infrastructure Systems*.

### Technology maturity enables the shift

The technologies underlying software-defined approaches are well established. IEC 61850[13] and other communication protocols, digital substation concepts, centralized protection and control, redundant Ethernet networks, and edge computing have matured over decades, particularly in utility and high-voltage applications. Industry experience has also grown in coordinating data across devices and edge platforms while keeping time-critical protection functions local.[14]

**What is changing is the integration of these technologies into a standardized architecture at the MV switchgear level.** This integration makes it practical to separate the common physical platform from the software configurations that define application-specific behavior. At the same time, compressed project schedules and the need to replicate large power blocks across multiple sites have increased the value of architectures that can be configured and reused.

Together, these factors make software-defined CTO approaches a practical alternative to traditional ETO implementations. In the following section, we quantify how this architectural shift affects engineering, manufacturing, installation, and commissioning, and examine its cumulative impact on deployment speed.

## Quantifying the impact on deployment speed

The architectural simplifications described in the previous section affect activities across the MV switchgear deployment process. Time savings in engineering, procurement, manufacturing, installation, and commissioning can compound to significantly shorten the switchgear deployment interval. To estimate the potential impact, we compared traditional ETO and software-defined deployments using the scope and assumptions described below.

### Comparison scope and assumptions

The comparison represents a typical 20-panel primary, three-phase 34.5 kV, 2,500A switchgear application of over 100 MVA. Both cases were assumed to provide equivalent primary ratings and protection, control, interlocking, metering, monitoring, safety, and availability performance suitable for a large-scale data center application.

The traditional ETO baseline used conventional instrument transformers with panel-level protection, metering, monitoring and control devices connected through project-specific point-to-point wiring. In contrast, the software-defined case used low-power current and voltage transformers (LPCTs/LPVTs) and merging units connected over an IEC 61850 process bus with redundant communication. All data connects to a single centralized head unit serving all 20 panels, and hosting protection, automation, metering and monitoring functions.

---

[11] IEC 62351-6:2020, Security for IEC 61850; IEC 62351-8:2026, Role-based access control

[12] IEC 62443-4-1:2018, Secure product development lifecycle requirements; IEC 62443-4-2:2019, Technical security requirements for IACS components; and IEC 62443-3-3:2013, System security requirements and security levels.

[13] Schneider Electric White Paper 285, Comparing IEC 61850 vs. Traditional Monitoring, Protection and Control Solutions in Large Data Centers, evaluates the application of IEC 61850-based monitoring, protection, and control to a large, complex MV data center electrical system.

[14] CIGRE Working Group B5.60, Protection, Automation and Control Architectures with Functionality Independent of Hardware, Technical Brochure 891, CIGRE, 2023

The software-defined case assumes that the reusable configuration files and test routines have already been created. Under these conditions, software-defined architectures enable functional testing to begin significantly earlier in the schedule because protection, control, metering, and automation interactions can be fully simulated and verified in a software environment well before the physical hardware is available.

Durations were derived from a modeled schedule developed with input from engineering, manufacturing, installation, and commissioning subject-matter experts. The comparison begins when project requirements are available and ends when installation and commissioning are complete and the lineup is ready for energization. It includes project-specific engineering, manufacturing, dispatch, installation, testing, and commissioning. It excludes initial platform development, utility interconnection, permitting, site civil work, upstream equipment, and energization. Activities performed concurrently are represented as elapsed schedule time rather than summed task durations.

## Engineering, procurement, and manufacturing

Data center owners may already reuse established electrical schemes across sites, but the physical implementations are rarely identical. Incoming utility characteristics, geographic and local requirements, transformer selection, connected equipment, and the choice of backup source, such as batteries or generators, can vary from one deployment to another. In traditional ETO implementations, these differences often affect hardware selection, physical layouts, wiring diagrams, protection settings, control logic, and documentation.

The software-defined approach allows most of this project-specific variation to be addressed through configuration rather than changes to the underlying hardware. The same hardware platform can therefore support applications that use similar electrical schemes but differ in their detailed protection and control requirements.

The schedule reduction results from three primary changes:

1. **Activities can proceed in parallel.** Traditional manufacturing generally cannot begin until project-specific device selections and wiring designs are sufficiently complete. Standardized panels in software-defined architectures can be manufactured while application-specific protection, interlocking, and control functions are configured in software. Testing of protection and control functions can also be performed in software early in the project, without waiting for hardware availability.
2. **Procurement and physical integration time is reduced.** Fewer project-specific relays, meters, and control devices reduce the effort and elapsed time associated with device selection, ordering, delivery, mounting, integration, and inspection. Standardized hardware can also be procured in advance of the final software configuration.
3. **Engineering and validation become more repeatable.** Validated configuration, logic, drawing, documentation, and test templates can be adapted to project requirements rather than recreated for each deployment. Fewer terminals and point-to-point connections also reduce wiring and inspection work during manufacturing.

**The combined engineering, procurement, and manufacturing interval, from the initial review of customer requirements through factory acceptance testing, decreased from approximately 26 weeks for traditional ETO switchgear to 8.5 weeks for the software-defined approach, a reduction of approximately 67%.**

## Installation and commissioning

The physical placement of the switchgear, main power connections, civil work, and required safety checks remain largely unchanged. The field savings are concentrated in control wiring and commissioning.

Traditional ETO commissioning requires teams to verify point-to-point wiring and confirm that protection, control, interlocking, communications, and source-transfer functions operate as specified. Project-specific wiring and logic increase the number of physical connections and interactions that must be checked and can increase troubleshooting and rework when discrepancies are found.

The software-defined configuration reduces this workload by replacing many hard-wired control connections with networked communications. As shown in Table 1, the software-defined configuration uses 91% fewer terminals and 87% less control-wire length. These reductions do not translate directly into equivalent reductions in total installation time because the primary equipment and power connections remain, but they substantially reduce the number of control connections that must be terminated and checked.

The software-defined comparison also assumes that a software engineering tool is used to validate selected protection, interlocking, control, and communication interactions before switchboard assembly. Because all data are digitized, test scenarios can be simulated and performed through software interactions during factory or site acceptance tests. This can be done in advance of installation and reused across similar panels and subsequent deployments. Site-specific verification remains necessary, including inspection of physical connections, safety checks, communications verification, and confirmation that the installed system performs as required.

In the comparison, **the combined installation and commissioning interval decreased from approximately five weeks to 2.5 weeks, a 50% reduction.** Because these activities overlap and were evaluated together, the result is reported for the combined elapsed interval. Most of the reduction is expected to occur during control-wiring verification and commissioning rather than physical installation of the primary switchgear.

## Overall deployment impact

Across the complete comparison, the MV switchgear deployment interval decreased from approximately 31.5 weeks to 11.5 weeks, a reduction of approximately 64%. This translates to 2.7 times faster deployment. Dispatch remained unchanged at 0.5 weeks; the improvement resulted from shorter engineering, procurement, manufacturing, installation, and commissioning intervals. Table 2 summarizes the schedule comparison.

**Table 2**

*Deployment activity and time for traditional ETO vs. software-defined MV switchgear*

| Deployment activity | Traditional ETO | Software-defined | Time reduction |
|---|---|---|---|
| Engineering and manufacturing* | 26 weeks | 8.5 weeks | 67% |
| Dispatch** | 0.5 weeks | 0.5 weeks | 0% |
| Installation and commissioning | 5 weeks | 2.5 weeks | 50% |
| **Total MV switchgear deployment interval** | **31.5 weeks** | **11.5 weeks** | **64%** |

* Includes the initial review of customer requirements through factory acceptance testing. Since activities partly overlap, durations represent elapsed schedule intervals rather than summed task durations.
** Assumes zero dispatch difference and excludes delivery time which varies by project and location.

The magnitude of the schedule reduction achievable depends on the size and complexity of the application, project-specific requirements, maturity of the reusable templates, implementation practices, delivery distance, and availability of trained personnel. When MV switchgear lies on the data center project's critical path, shortening this interval can contribute directly to improved time-to-power.

## Additional lifecycle benefits

The primary focus of this paper is the effect of software-defined MV architectures on deployment speed. However, standardized hardware, reusable configurations, and networked communications can also provide benefits during operation.

### Operational and functional flexibility

Changes in operating practices, utility requirements, maintenance strategies, safety policies, or business priorities may require updates to protection, control, and interlocking behavior. In a software-defined architecture, many of these changes can be implemented through configuration updates on the existing hardware platform rather than through the installation of new physical devices and wiring. Examples include modifying operating philosophies, updating transfer schemes, introducing new protection functions such as arc-flash protection, strengthening safety interlocks, replicating control functions across panels, and responding to changing utility or grid requirements. Many of these updates can be implemented without requiring switchgear outages, reducing disruption while enabling the system to adapt to evolving operational requirements. All changes affecting protection, control, and interlocking must continue to follow appropriate engineering review, validation, documentation, and change control processes.

In addition, as data center campuses grow and the number of deployments increases, operators may need to support more equipment with a limited pool of experienced personnel. Standardized platforms, consistent interfaces, centralized monitoring and diagnostics, and reusable operating procedures can reduce equipment-specific variation and make training and knowledge transfer more scalable. Where cybersecurity policies permit, remote diagnostics can also allow specialized personnel to support multiple sites. Together, these capabilities can help reduce human error and shorten mean time to repair (MTTR), although they do not eliminate the need for qualified personnel.

### Simplified spare parts management

Traditional MV switchgear lineups may incorporate multiple relay, control, communication, and monitoring devices with different hardware versions and spare part requirements. Maintaining inventories for these devices can increase procurement, storage, and lifecycle management complexity.

A software-defined architecture reduces the number of device types required across the lineup by standardizing hardware platforms and consolidating functions. This approach reduces the number of spare part references that operators must stock and manage, while simplifying maintenance planning and replacement strategies. As deployments increase across multiple sites, greater hardware commonality also improves consistency in spare parts inventories and support practices.

### Simplified phased expansion

Software-defined power architectures can also simplify physical expansion by separating much of the site's functional variation from the underlying hardware design. As additional power blocks are deployed, previously validated hardware designs, configurations, and test procedures can be reused and adapted rather than

recreated. This approach enables standardization, reduces engineering effort, improves consistency across installations, and supports phased capacity growth while maintaining established operational practices.

## Reduced exposure during local operations

In traditional MV switchgear, interfaces for protection relays, metering, and controls are commonly located on the equipment front. Routine operating, monitoring, and diagnostic activities may therefore require personnel to approach energized equipment and, depending on the task, enter the arc flash boundary.

In the software-defined architecture, current and voltage measurements[15] are digitized by the merging units and communicated over the process bus. Routine metering, monitoring, configuration, and control interactions across multiple panels can be consolidated through a shared head unit. Depending on functional and customer requirements, the head unit can be located within view of the switchgear or in a nearby location *outside* the arc flash boundary.

This arrangement can reduce how frequently personnel must approach energized MV equipment for routine local operations.[16] It does not eliminate the need for local inspection, switching, or maintenance, nor does it reduce the underlying arc flash hazard or calculated incident energy. Appropriate arc flash assessments, personal protective equipment, and safe work practices remain necessary.

## Lower material use and embodied carbon

Replacing project-specific devices and point-to-point control connections with standardized merging units and networked communications can reduce requirements for control wiring, terminals, electronic devices, and associated panel space. As shown in Table 1, the software-defined configuration analyzed used 87% less control-wire length, 91% fewer terminals, 85% less copper conductor mass, 71% less voltage transformer mass, and 96% less current transformer mass than the traditional ETO configuration.

These reductions could lower the material-related embodied carbon associated with the switchgear. Any resulting reduction in Scope 3 CO2e emissions is not quantified here; it depends on the specific configuration and should be evaluated at the project level based on the selected architecture and components.

[15] IEC 61869-9:2016, Instrument Transformers, Part 9: Digital Interface for Instrument Transformers, 1st ed., 2016

[16] National Fire Protection Association, NFPA 70E: Standard for Electrical Safety in the Workplace, 2024 ed., Section 130.5 and Informative Annex O

# Next steps

Software-defined MV switchgear offers a path to faster deployment. It does not remove the need for application engineering, protection coordination, or rigorous validation. Instead, it shifts most of this work from repeated project-specific design toward standardized hardware, reusable configurations, and repeatable validation processes. This shift can shorten time-to-power while preserving the protection and control capabilities required by AI data centers.

Achieving these outcomes requires software-defined switchgear to be considered early in planning and procurement. To maximize value, data center owners, designers, and electrical system providers should take the following actions:

1. **Evaluate applicability based on size, scale, and complexity of the MV architecture and identify deployment bottlenecks.** This approach typically applies to large data centers and AI factories, where repeatable power blocks provide opportunities for scale. Determine whether MV infrastructure is on the critical path and which project-specific activities create lead times incompatible with deployment objectives.

2. **Define a reusable electrical design baseline.** Most large-scale data center operators already implement varying degrees of design standardization across deployments. A software-defined architecture allows these standard practices to be captured as reusable hardware that can be adapted to site-specific requirements with configuration and validation templates. The baseline should enable controlled variation and change control.

3. **Establish reusable protection and control schemes**. Identify redundancy and backup functions for protection, automation and control, and their required sequences of operation. Define local and centralized functions, failure responses, network redundancy, cybersecurity, configuration controls, and validation requirements.

4. **Validate the approach through testing and prepare to scale.** Begin with a defined power block or site, measure schedule and implementation results, and refine the standardized design. Incorporate lessons learned and build skills to accommodate local requirements and implement future changes consistently across regional or global deployments.

5. **Establish a digital lifecycle strategy.** Integrate software-defined MV switchgear into the broader digital strategy for the data center. Define how configurations, test records, change history, and operational data will be maintained throughout the asset lifecycle and reflected in digital representations used for commissioning, operations, and maintenance. This foundation can support future digital twin initiatives while improving visibility, consistency, and governance across deployments.

This document is to be considered as an opinion paper presenting general and non-binding information on a particular subject. The analysis, hypothesis and conclusions presented therein are provided as is with all faults and without any representation or warranty of any kind or nature either express, implied or otherwise.

## About the authors

**John Slater** is a Senior Manager in Thought Leadership Content specializing in medium voltage switchgear and power systems. With more than 25 years of experience in the electrical industry, he has developed and delivered technical and strategic content across high voltage (HV) and medium voltage (MV) markets.

**Wendy Torell** is a Senior Research Analyst in Schneider Electric's Data Center Research & Strategy group with 30 years of data center experience. She analyzes the value of emerging technologies and translates industry trends into practical design and operational guidance. Wendy leads development of Schneider Electric's interactive, web-based TradeOff Tools, helping clients quantify decisions related to availability, sustainability, and cost. She holds a BS in Mechanical Engineering from Union College, an MBA from the University of Rhode Island, and is an ASQ Certified Reliability Engineer.

**Venanzio Paolo Ferraro** is a Senior Principal Architect within Schneider Electric's Energy Management business. He has extensive experience in R&D and marketing, including product, project, and technical management roles focused on medium-voltage products. He is also active in standardization as chairman of National technical committee and member of IEC and CIRED/CIGRE Working Groups.

**Maxime Durand** is a Principal Solution Technical Leader within Schneider Electric's Energy Management business. He leads the technical definition and development of digital solutions for medium-voltage applications, with a focus on protection, control, and asset monitoring.

**Amita Sonar** is a Global Product Owner focused on software-defined power and digital solutions for power distribution systems. Drawing on deep expertise in connected technologies, IoT-enabled equipment, and product innovation, she drives the development of intelligent solutions that improve operational performance and asset reliability.

## Acknowledgements

Special thanks to **Brice Martinot-Lagarde, Kevin Castel, Julien Moreau, Sebastien Chabanet, Steve Heighington, Geoffroy Moens, Sebastien Cruz-Mermy, Herve Bigeard, Francisco Jose Ramirez Fajardo, and Maria A. Torres Arango** for their expertise and technical review of this white paper.

## RATE THIS PAPER

# Resources

How 6 AI Attributes Change Data Center Design
**White Paper 110**

Retrofitting Existing Power Systems for AI Clusters
**White Paper 126**

Cybersecurity Guidance for Data Center Power and Cooling Infrastructure Systems
**White Paper 216**

Comparing IEC 61850 vs. Traditional Monitoring, Protection and Control Solutions in Large Data Centers
**White Paper 285**

Scaling AI-Enabled Digital Twins Across the Data Center Lifecycle with OpenUSD
**White Paper 302**

Browse all white papers and executive reports
**Data Center Research & Strategy Papers**

Browse all TradeOff Tools™
**Data Center Research & Strategy** Tools

Note: Internet links can become obsolete over time. The referenced links were available at the time this paper was written but may no longer be available now.

## Contact us

For feedback and comments about the content of this white paper:

Schneider Electric Data Center Research & Strategy
dcsc@se.com

If you are a customer and have questions specific to your data center project:

Contact your Schneider Electric representative at
https://www.se.com/ww/en/work/support/country-selector/sales-contact/

© 2026 Schneider Electric. All rights reserved.