Smart Grid: E4S – Edge for Smart Secondary Substation Systems

Authors:
Christian Benkeser, Intel
Martin Dauner, Intel
Maik Fox
Dean Samara-Rubio, Intel
Javier Sola Villalobos, Intel
Kelvin Chuang 莊家維, Delta

The electricity distribution grid architecture consists of layers defined by the voltage level of the alternating current (AC) power system. High voltage is used to transfer power over distance efficiently (as high as 1MV or more), whereas power is delivered to consumers for factories, businesses, and residences at lower voltage levels (often tens of kV for manufacturing or less than 1kV for residences). The last layer transition to most end-consumers is from medium (e.g., between 6kV and 15kV) to low voltage (less than 1 kV), and this layer of the grid is often referred to as secondary distribution. **Figure 1** (next page) is a simplified diagram of the voltage levels, the positions of the voltage transformers, and the points of interconnection for generators and consumers.

Some secondary distribution transformers are stand-alone devices, such as the poletop can-shaped transformers that can be seen strung along the single-phase medium-voltage lines threaded through older neighborhoods across the United States. But in many regions, especially in higher population density areas, secondary distribution transformers connect to 3-phase medium-voltage rings. These transformers are co-located with busbars, switchgear, and other devices to form a node referred to as a Secondary Substation (common in Europe and Asia). Secondary substation sites vary considerably in location, size, method of construction, and accessibility – such as roadside buildings, compact prefabricated metal enclosures, closets in building basements, and underground vaults with manhole cover access. As the secondary substations are positioned close to the point of use by consumers, a regional or national utility may have tens of thousands of secondary substations dispersed across its territory. Certainly, for a utility, there will be wide variations in site capacity, configuration, construction type, and local environment for the secondary substations in their network.

Historically, secondary distribution networks have been manually configured and lightly monitored. But in the past two decades, the level of investment in secondary distribution networks has risen. First, the proliferation of distributed generation, energy storage, and active energy management systems in customer premises stress the secondary distribution systems. Second, utilities have installed smart metering infrastructure using powerline communications or short-range wireless communications that connect through meter data concentrators in their secondary substations. As a result, utilities worldwide have begun extending their monitoring, control, and automation systems to reach into the secondary distribution network. They often have an incremental roadmap for applications and capabilities that can be deployed over time. These applications include meter data concentration, feeder monitoring, transformer health assessment, fault indication, busbar monitoring, and circuit protection. The industry practice thus far is to treat these functions independently and install a discrete device for each application. As mentioned earlier, each utility might have tens of thousands of secondary substations, which vary in their configuration. Further, each substation may have five or more independent "controllers," often from different vendors and with different internal designs. The number of vendors, device models, versions of firmware, and configuration management consoles leads to high operational complexity and inflexibility.

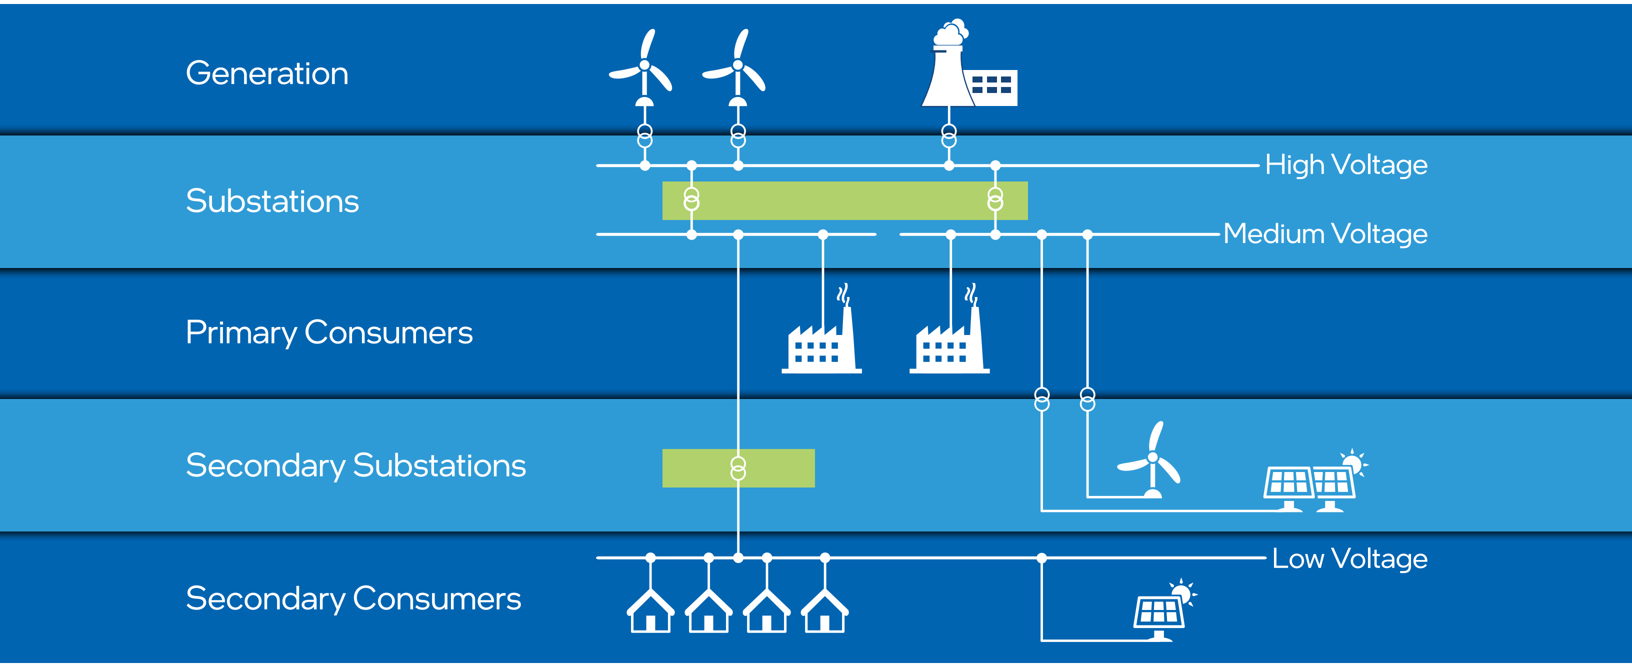

Figure 1: Exemplary grid topology going from generation (power plants, traditional or renewable) through several levels of substations and consumers. In recent years, significant amounts of generation capacity are connected at the medium and low voltage levels. Not all equipment is shown in this simplified diagram.

Modernizing the Grid with E4S

Leading utilities recognize that the conventional approach is no longer viable. They have begun a transition to software-defined open platforms constructed using broadly adopted standards of the telecommunications and industrial sectors. The Secondary Substation Platform (SSP) defined by the E4S Alliance (https://e4salliance.com) is leading the way with a comprehensive architecture including substation-grade hardware, software, cybersecurity, device management, and application lifecycle management. The E4S architecture can enable higher reliability and lower costs for grid operations by consolidating applications onto interoperable hardware with open specifications. This, in turn, will open the door for deeper levels of control and automation in the secondary distribution network, including adaptive systems and machine-learning-based optimization.

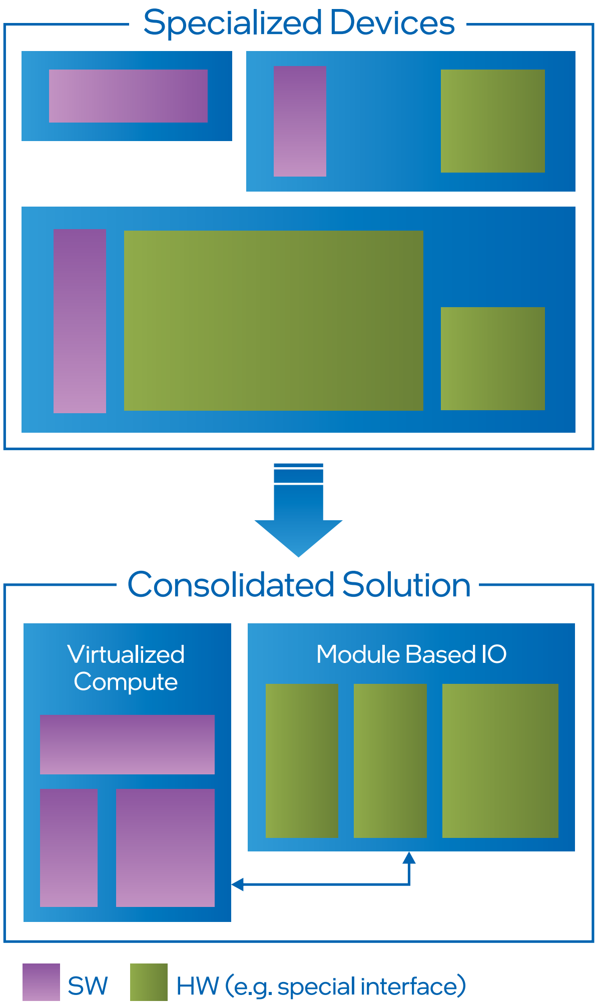

Figure 2: Moving from a group of specialized devices to a consolidated Edge server system within a secondary substation.

The concept is shown in **Figure 2**. Today's specialized devices are single-purpose, fully independent systems, each having all necessary basic elements (power supply, operating system, etc.). The E4S merges these functions into a flexible and scalable consolidated solution using virtualization technologies, shared resources, and open standard interfaces.

Benefits of a Smarter Grid

Resilient, secure and scalable electricity distribution grid architecture is critical to meeting current and future social advancements. Advanced software-defined grid automation and control unlock broader benefits to reach shared sustainability objectives. To learn more:

- » Can Tech Solve the Grid's Climate Change Challenge?
- » How Smart Tech Can Prepare the Grid for Climate Change?
- » Learn about Virtualized Protection and Control Technologies.

The E4S architecture reduces the need for separate data silos for each application. It enables optimized processing of data and smarter reactions to events directly at the substation. The processing of data at the edge benefits the system resiliency, security, and privacy. Building upon modern industry standards in hardware and software will significantly reduce vendor-lock-in and component obsolescence, fostering innovation and ultimately driving costs.

The Hardware Challenges

The E4S hardware system, when deployed to the field, will need to fit into the existing sites: the footprint, the construction, location, and regional and local conditions are not changeable.

The E4S system must be low-maintenance and a local technician and a remote operator should easily diagnose issues that emerge. The systems must be adaptable to the wide range of site configurations and be easily deployable by any installers. For example, one individual should be able to carry and mount the device into position without special tools or in-depth knowledge of the technology. Climatic and environmental variations are also important, and each locale may present its challenge: salt spray, humidity, blown dust, heat, or cold. Insects and gastropods can also cause problems over the long lifetime of the equipment.

Computing and communications components deployed to secondary substations are most often mounted inside sealed plastic or metal enclosures. Many utilities will not deploy equipment with cooling fans or vents since historically these have been associated with reduced reliability. Thus, a fanless E4S system must be operational even when mounted in a sealed cabinet in a hot sunny climate. To achieve this requirement, significant design effort was devoted to the overall thermo-mechanical design.

Additional constraints impacting the thermo-mechanical design include electrical isolation, operation in a challenging electromagnetic environment, and operation around medium-voltage electrical conductors (e.g., 6kV to 13kV or more). Metal heatsinks with strong thermal pathways to internal components are the most effective and common method of heat removal, but care must be taken in the size, placement, and grounding of exposed metallic surfaces in the substation environment. To summarize, the most challenging aspects for the design of a flexible Edge solution that performs within a secondary substation are:

- **Total solution dimensions** – needs to be as compact as the fixed function devices it consolidates.
- **Environmental conditions** – temperature ranges are equivalent to outdoor applications, with up to several kilovolts and different voltage levels.
- **Serviceability and maintenance** – passive cooling is required to avoid moving parts and their servicing; modules need to be easily replaceable.
- **Standards compliance** – utilities expect numerous international standards to be fulfilled, covering topics such as isolation requirements or electromagnetic compatibility (EMC).

The E4S Solution

The diversity of secondary substations in the grid requires a flexible concept with scalability in computing power and the number and type of physical input/output connections. The software and hardware architecture should adhere to open standards and use widely available components wherever possible. While working within the E4S Alliance, Intel created a prototype system aligned to the E4S requirements (E4S Alliance Specification).

The requirements for modules focus on external interfaces, physical size/shape, and power limitations so vendors with differing functional and electrical requirements can create compatible module variants. For example, the compute modules could use a range of different processors and accelerators to scale from advanced automation tasks up to AI/ML-enabled data processing at the Edge. Power modules can be developed to support alternative supply voltage types. One variation is to add energy storage into the module for immunity to external power dropouts. Switch modules can range from simple packet layer switches to managed switches implementing VLANs. The E4S hardware specification enables multiple vendors to deliver modules with the functionality and cost to suit each substation's needs.

The mechanical concept allows for easy configuration and servicing from the system's front face. All modules have their mechanical enclosure and can be handled, installed, and replaced without special equipment. Many utilities require no fans or other moving parts to be used, but fans could be incorporated depending on the needs and policies of each utility.

Figure 3 shows an overview of the E4S hardware architecture. The backplane is indicated by the dotted line box in the diagram. The backplane serves as the mechanical core and provides power rails, high-bandwidth communications links, and a dedicated system management bus. It is fully passive to maximize reliability. The backplane defines the position of the available module types, limiting the placement variability but ensuring predictable thermal and electrical behaviors. In addition to ethernet and serial sideband connections, the backplane includes USB3.0 and PCIe links between the compute modules and from the primary compute module to two of the input/output module positions.

Compute Modules: The primary compute module will host the application software and a system monitoring and management utility that communicates to the other modules for status, state, and health. The primary compute module is a general computing device based on a CPU such as an Intel® Core-i7™ Processor. The secondary compute module can have the same internal design as the primary or it could have, for example, specialty chips for efficiently analyzing data using machine learning techniques. To enhance system resilience and reliability, the secondary compute module can be configured as a backup to the primary compute module.

Input/Output Modules: The data input/output modules are specialized data acquisition and communication modules, e.g., for smart meter data concentration, transformer supervision or switch control. These modules can use Ethernet communication on the backplane or the USB and PCIe communication to the compute elements available in extended slots. One use for the extended slots with high-speed links is for a wireless (5G or WiFi6) modem.

Ethernet Switch Module: The ethernet switch module is the hub of the ethernet backbone.

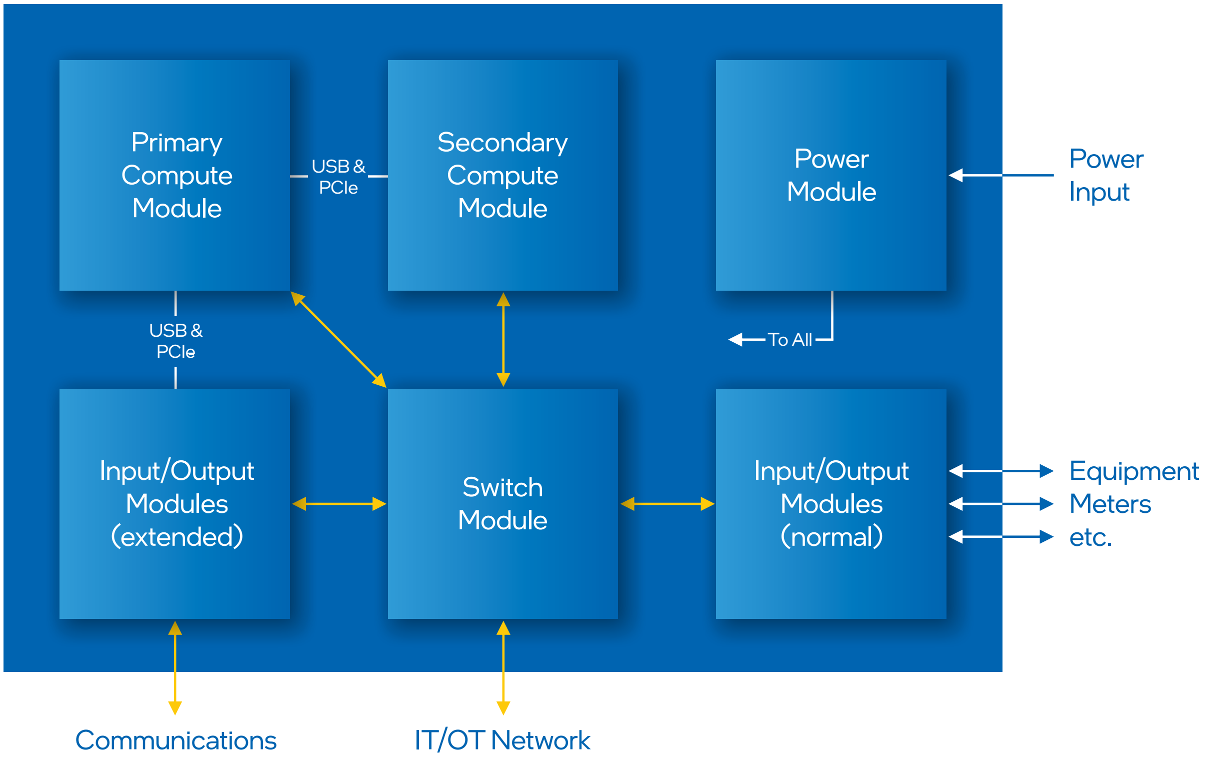

Figure 3: E4S hardware architecture overview.

All modules (except the power module) are connected to the Ethernet switch module for all application data exchange. Each module has a 1Gbps link to the switch and the switch has a pair of externally accessible ethernet connections. Due to the point-to-point connections of Ethernet links, separation of dataflow can be implemented. For example, IT/OT and internal links can be separated with switch modules containing a managed switch device (ASIC) that supports VLAN capability. The switch module contains not only the switch ASIC but also a programmable microcontroller that could be utilized in the future for value-added system capabilities.

Key E4S Advantages

- **Suitability** – designed for the thermal, mechanical, and electrical environment of secondary substations. The design considers the installation, configuration, management, and lifecycle maintenance.
- **Scalability** – The compute modules support a range of CPU's such as Intel® Atom Processors, Intel® Core™ Processors, and entry-level Intel Xeon® Processors. Computing accelerator devices are also possible for the secondary compute module slot for analytics or other intensive workloads.
- **Resiliency** – Incorporating two compute modules allows for active standby and active-active configurations to achieve high availability.
- **Interoperability** – Fully open and standards-based interfaces between modules, including ethernet and PCI-express across the backplane, enable modern and efficient data exchange methods for modules and services. This approach enables broad ecosystems of compute elements, accelerators, specialized data acquisition interfaces, communication modules, and more.
- **Configurable** – allows tailoring to individual utility and substation requirements on a large scale and case-by-case basis regarding performance and solution dimensions.

Power Module: The power module takes external DC power, converts it, and distributes it to the other modules on the backplane, each with its own dedicated power rail. The power module monitors the current and voltage levels for each of these power rails, has overcurrent limiting capability, and can switch the rails on and off. The power module is connected to the system management interface for the exchange of data and commands. The hardware design supports the hot plug of input/output modules for easy servicing and upgrades. The power module can be a standard-width or a double-width module (which would reduce the number of input/output modules allowed from 4 to 3). The double-width power module may house ultra-capacitors to maintain internal power rails during transient disturbances or drop out of the external power supply. Other options for a double-width power module include the incorporation of internal redundancy, the addition of a fan, and the ability to accept AC power to the system.

E4S Hardware Design Insights

Intel has created a hardware prototype to assess the feasibility and feed design insights into the specification. In a second design iteration, Delta Electronics, Inc. (https://www.deltaww.com/zh-TW/index/) is adapting the reference design into a commercial system capable of utility field trials and pilot deployments.

Prototype Design and Learnings

The prototype's key goal was to show feasibility in challenging environmental conditions. Based on mechanical design studies and thermal simulations of several operating conditions, the hardware as shown in **Figure 4**, was built and tested. The number of modules was maximized and advanced CPUs with significant power dissipation were chosen to evaluate the system design in the highest power configuration. Configurations with a smaller module count and other CPU's and chipsets can allow for appropriate scaling to meet actual needs.

The design phase included simulations with boundary conditions to predict system behavior and thermal-mechanical performance. For example, a common practice for deploying devices to a secondary substation has been to mount them into a sealed enclosure made of plastic or metal. In some cases, these enclosures can have vents and even circulating fans, but often, they are fully sealed boxes with feedthrough connectors for power cables and communications wiring. These cabinets are used for safety, protection of the equipment from accidental contact by personnel, and limitation of dust, water, insects, and gastropods that can lead to a reduced lifetime of the equipment. For a plastic enclosure, three scenarios were simulated:

1. The prototype system is fully enclosed within a plastic cabinet such that the only path for heat exchange to the environment is through the cabinet walls.

2. The enclosing cabinet has small vents, as is found in some secondary substation designs. This setup allows for convective airflow and passive air exchange with the outside environment.

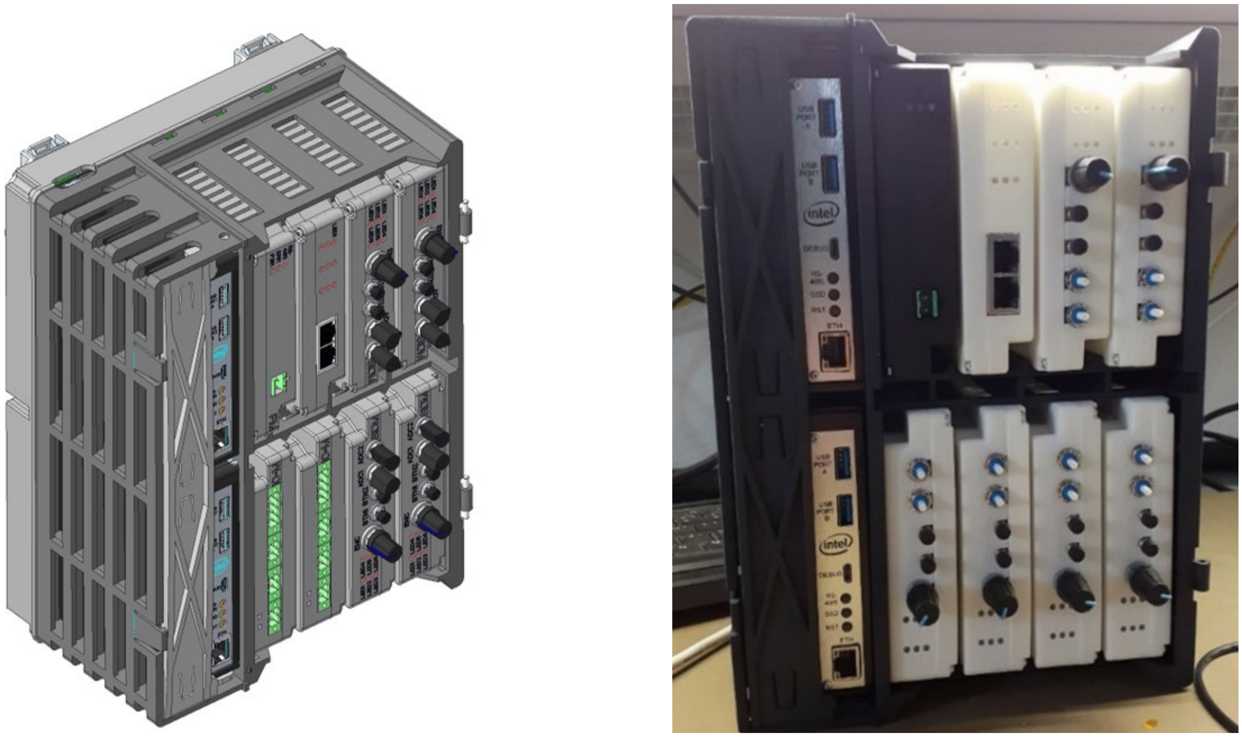

Figure 4: Intel E4S prototype hardware rendering and photo of a fully populated test system. A mechanical restraint on the left partially covers the two compute modules. The power module is in black color, and the remaining modules in white. The switch module is in the center position at the top. All six of the input/output modules shown are for testing and have knobs, buttons, and LED indicators used for thermal and interface testing.

3. The cabinet of scenario two, with additional equipment, cohabits the enclosure with the E4S system. This scenario can be seen in the left panel of **Figure 5** where the E4S modular system is shown towards the bottom, and additional active equipment is in the upper section of that enclosure.

The prototype system design has the two compute modules stacked vertically on the left side of the chassis. The simulations indicated that the upper compute module would have a higher temperature than the lower compute module but that the hottest region of the system would be within the power and switch modules. This can be seen in the left panel of **Figure 5**. However, the measurements indicate the upper compute module is hotter than the lower compute module but is also hotter than any other region of the system. Overall, the system performance of the prototype was worse than expected. Several mitigation possibilities were examined in subsequent thermal simulations based on the measurements from the prototype. These optimizations included:

- Improving the size, distribution, and airflow around the metal cooling fins of heat sinks.
- Computing module size and placement.
- Improving heat transfer from the silicon die of the processor to the compute module heat sink.
- Optimization of the internal positioning of components within the compute module.

The findings of these simulation runs were fed back into the E4S hardware specification. After this step, the specification was mature enough to develop a device that can be used for end-customer lab evaluations and then for live field trials.

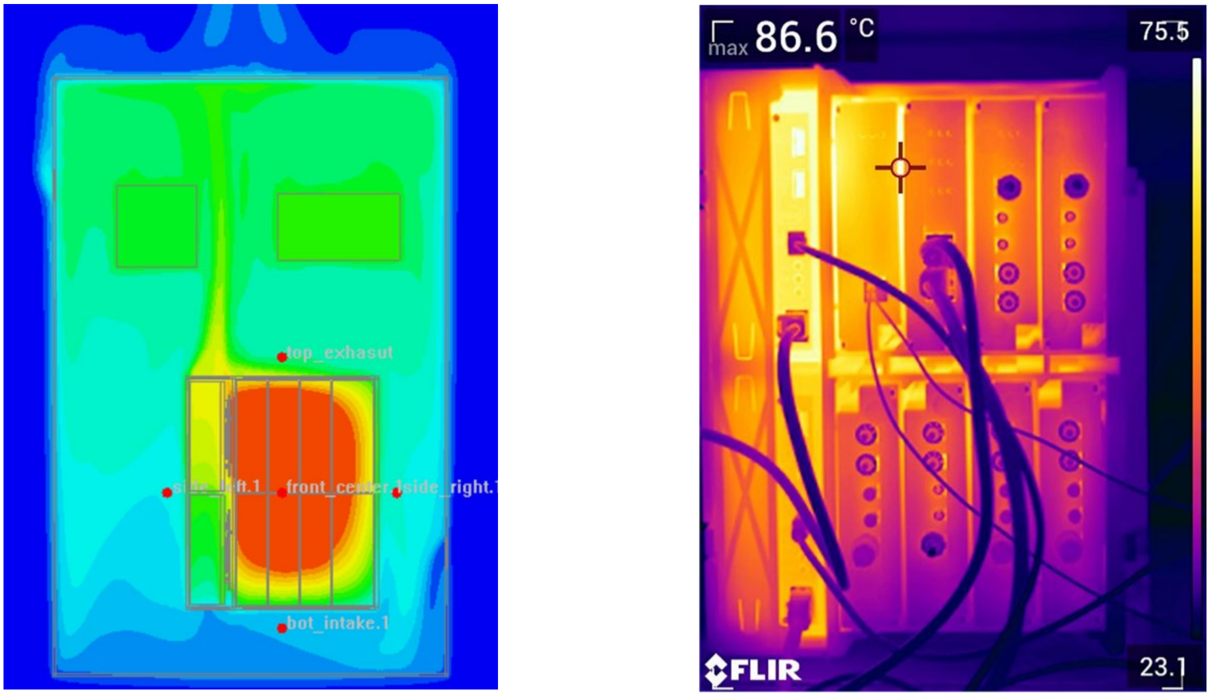

Figure 5: Thermal simulation (left) in an enclosed plastic cabinet box with additional hardware – qualitative results show the relative temperature distribution. On the right is an image of the real prototype provided by a FLIR thermal imaging camera. The system is under full load in an open lab environment.

Industrialized Design and Field-test Hardware

The experience with the prototype has allowed the Delta engineers to design a product suitable for field-testing and pilot deployments. This design, as shown in **Figure 6**, features an airflow-optimized outer frame that is sturdy enough for handling and mounting the system while still having enough openings to allow for efficient passive cooling. Mechanical mounting of individual modules was improved, a stabilizing horizontal center bar was added, and the chassis material was upgraded, as was the din-rail mounting design. The compute modules are made larger and repositioned to the left and right edges. By repositioning the compute modules, they are thermally separated, and their heat sinks are oriented to maximize convection. Overall system dimensions were not changed, so this design has at most four input/output modules (as opposed to six input/output modules of the prototype).

While the system architecture was not changed, thermal design constraints influenced the system's physical appearance. These changes are derived from diverse thermal simulations that integrate computer models of common substation requirements, including mechanical design, environment, heat sources, and airflow parameters.

For the secondary substations with sealed enclosures and in hot climates, the design requirement is for a 70°C ambient temperature within the enclosure, whereas a maximum ambient temperature of 60°C is specified in moderate climates. A corollary design challenge, therefore, was to see if a single high-performance Intel CPU such as the Core-i7 1185GRE could be used in the compute module

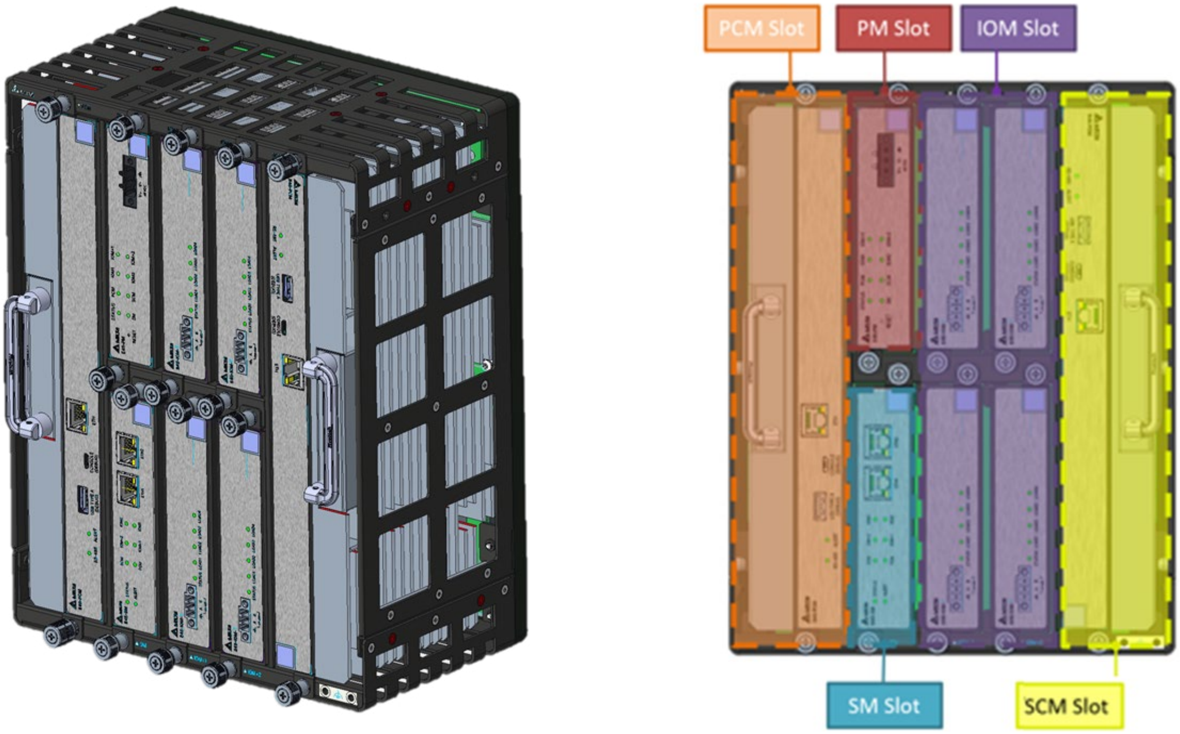

Figure 6: Industrialized mechanical design for the field tests. The compute modules are located at the left and right sides. There are the power module, switch module, and four IO modules. The outer housing is opened in strategic locations to optimize airflow and passive cooling.

for both climate scenarios by taking advantage of the scalability of the CPU power dissipation. The thermal simulation overview is shown below in **Figure 7**.

For the thermal simulation, a minimal and laminar airflow of 0.3 m/s is assumed to flow upwards from below the chassis. (This is a level of airspeed expected by convective currents for this device size and temperature range.). The simulation system is fully equipped with all module slots populated and operating at the maximum specified power level. All modules have a heat sink with fins running vertically. The chassis has numerous ventilation openings to facilitate airflow across all the module heatsinks.

The smaller modules' power density (power dissipation per unit volume) is approximately 6.8 mW/cm3. The power module has 21 mW/cm3. A compute module with an Intel CPU rated at 15 Watts will have 15 mW/cm3. A two-PCB approach was chosen for the compute module: an off-the-shelf standardized Computer-on-Module with the CPU and memory is landed on a carrier board with the remaining system components and interfaces to the backplane connectors. The CPU package surface is thermally connected to a copper heat spreader and then to an aluminum heat sink that is an integral piece of the module housing.

Figure 8 (next page) shows the compute module design that was submitted to the simulations.

Depending on the geometry of the compute module heat sink, a bigger surface area generally facilitates heat exchange. However, when investigating economical solutions and physical constraints, the fins' height and clearance between fins are limited. When the space for convection between the fins becomes too narrow, a turbulent flow pattern emerges, reducing effectiveness. Extending fin height gives a measurable advantage but with diminishing returns. In this design, the aluminum heatsink of the compute module has 17 fins of 2mm width, 5mm spacing, and 256cm

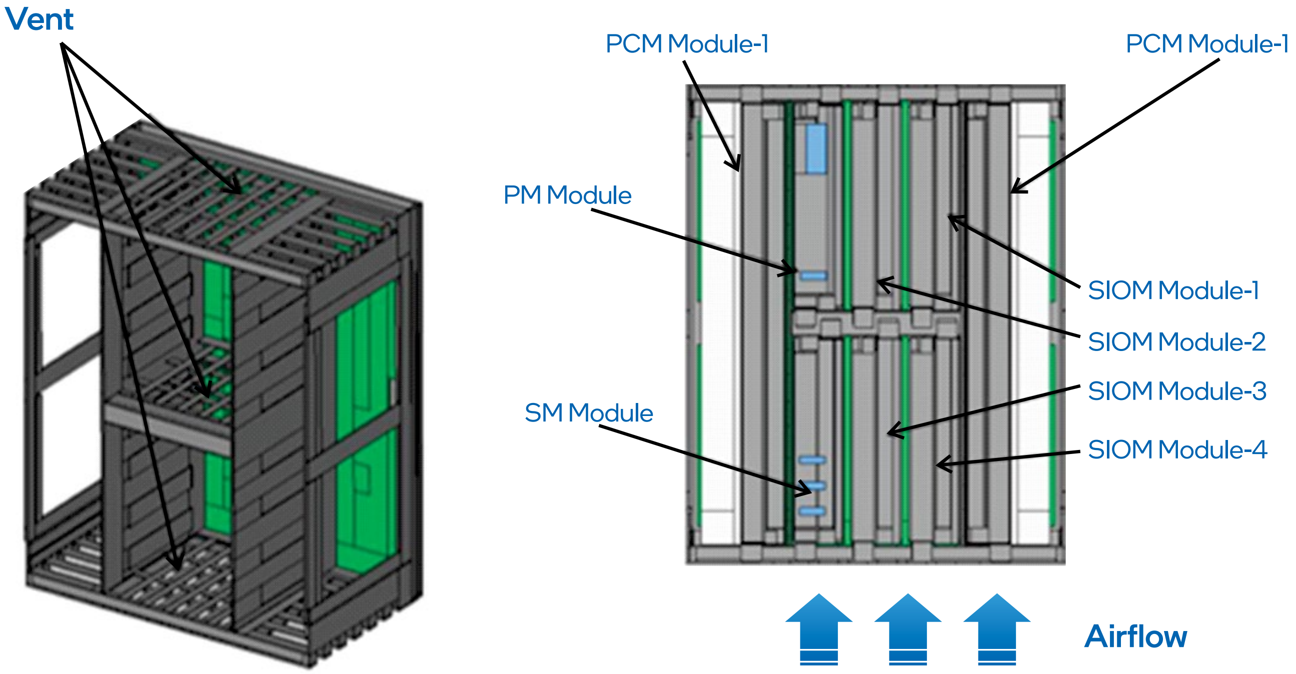

Figure 7: Basic thermal simulation model setup shows airflow direction and a mechanical view of the empty chassis. For the simulation, one of the input conditions is a laminar airflow rate of 0.3 m/s

length. A detailed simulation experiment was designed to find the optimum design parameters for the fin height, as shown in **Figure 9**.

According to the simulations, for the CPU consuming 15W at an ambient temperature of 65°C, the fin height of 40mm can maintain the CPU surface temperature at 100°C (which is the maximum CPU temperature in the product specification). The model was extrapolated to estimate that a 90mm fin height would be needed to allow for a 70°C ambient air temperature.

To reach the worst-case required ambient temperature of 70°C, the CPU power dissipation must be reduced below 15W. Thus, the next step of analysis focused on estimating the allowable CPU power dissipation level with the heatsink determined previously.

Compute Module Overview

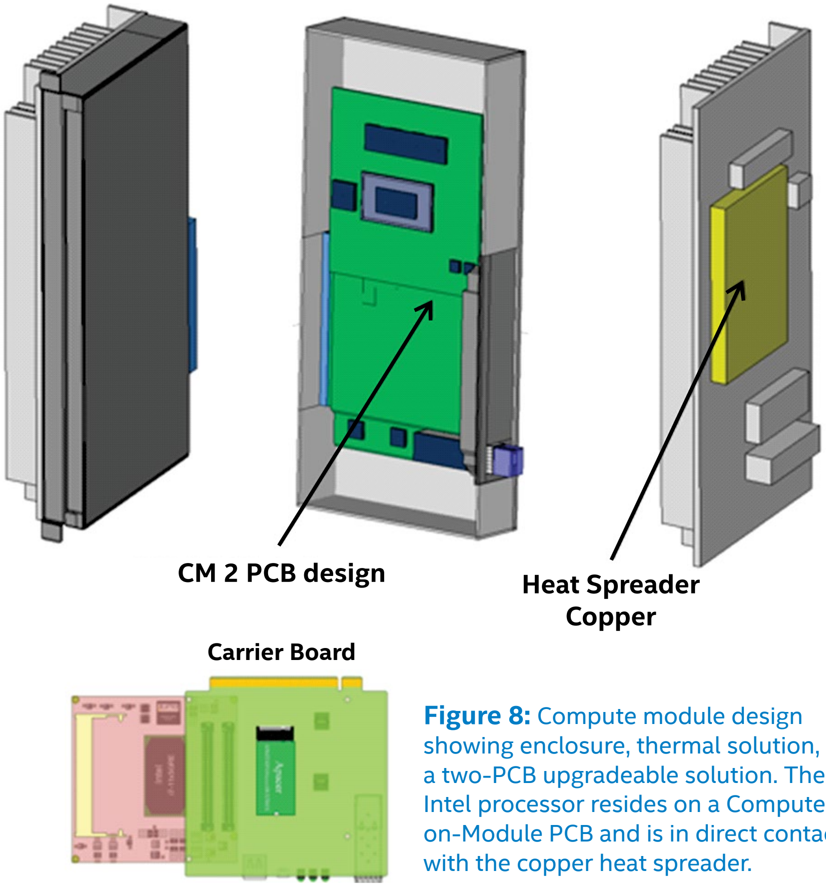

Figure 8: Compute module design showing enclosure, thermal solution, and a two-PCB upgradeable solution. The Intel processor resides on a Computer-on-Module PCB and is in direct contact with the copper heat spreader.

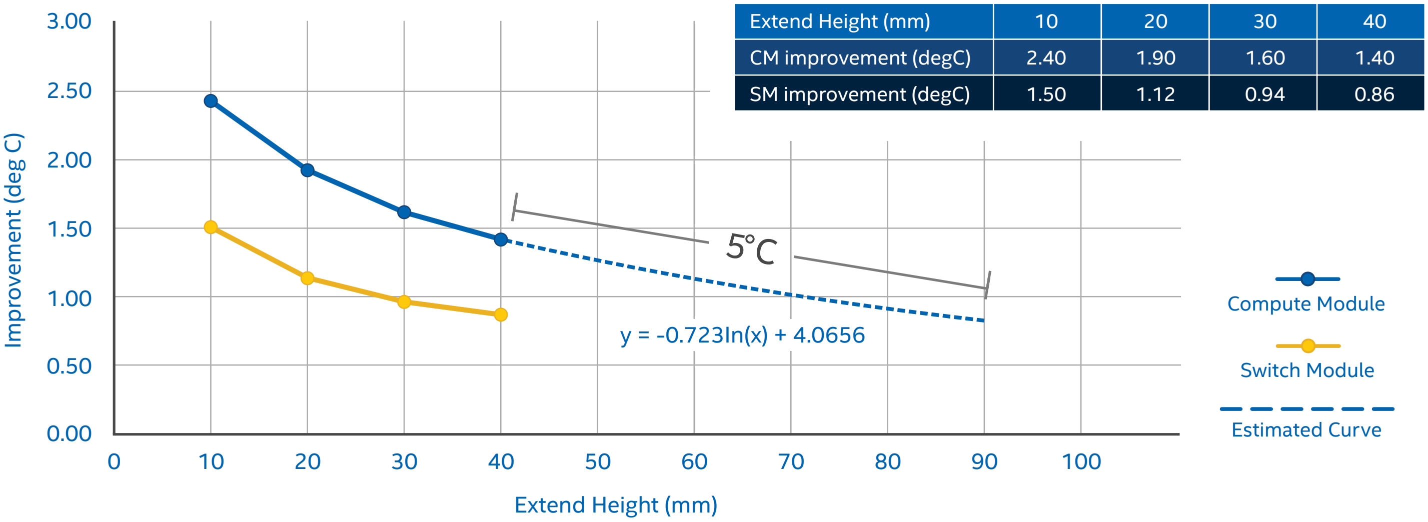

Extend Height (mm)	10	20	30	40
CM improvement (degC)	2.40	1.90	1.60	1.40
SM improvement (degC)	1.50	1.12	0.94	0.86

Figure 9: Simulation results of the fin height variation experiment. The temperature numbers given represent the additional improvement per step (temperature delta per 10mm fin height extension). So, for the compute module, the 10mm fin height gives 2.4 degrees improvement compared to a flat sheet, moving from 10mm to 20mm height gives an additional 1.9 degrees improvement, etc. Curve fitting was used to extrapolate to 90mm fin height.

The three simulation trials are tabulated in **Figure 10**. The figure also shows the temperature profiles of the compute module with the three surface points indicated (top, front, and center).

Modification	Trial 1	Trial 2	Trial 3
Ambient (degC)	60	60	70
CPU (Watt)	12	15	9

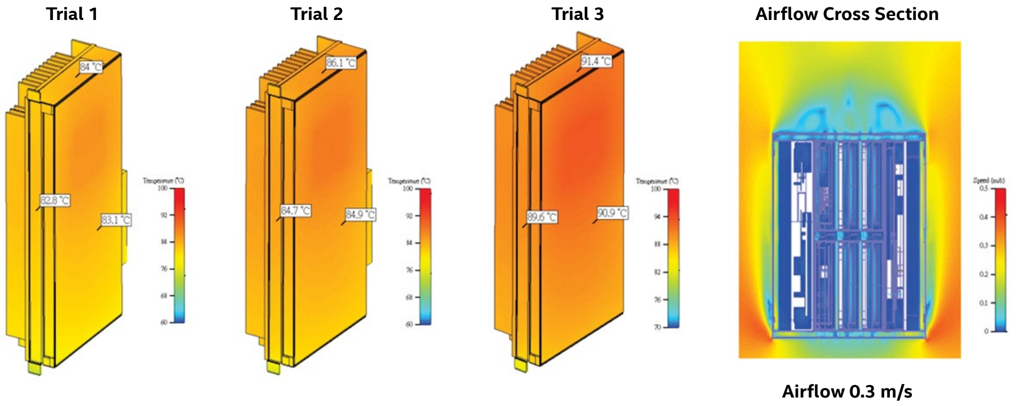

Trial 1

Trial 2

Trial 3

Airflow Cross Section

Airflow 0.3 m/s

Figure 10: Overview of the thermal simulation results for three scenarios to find maximum allowable CPU power dissipation at differing ambient temperatures.

Figure 11 provides the simulation results in graph format in which the data is plotted against a 9W processor power configuration at 70°C ambient temperature. The CPU power is a strong influencer of the processor package temperature. Still, the heat from the CPU is effectively channeled to the heatsink since the other points within the module see significantly less rise in temperature.

Intel processors' thermal design and protection features can throttle the power dissipation dynamically if specified by the end-user. The Intel Core i7 1185GRE processor provides sufficient computing capability for the set of applications and fits the power envelope with a maximum power dissipation of 15W (for design calculations), and a throttled power level of approximately 9W.

Temperature Rise on the Compute Module with increasing CPU Power Dissipation, at 70°C Ambient

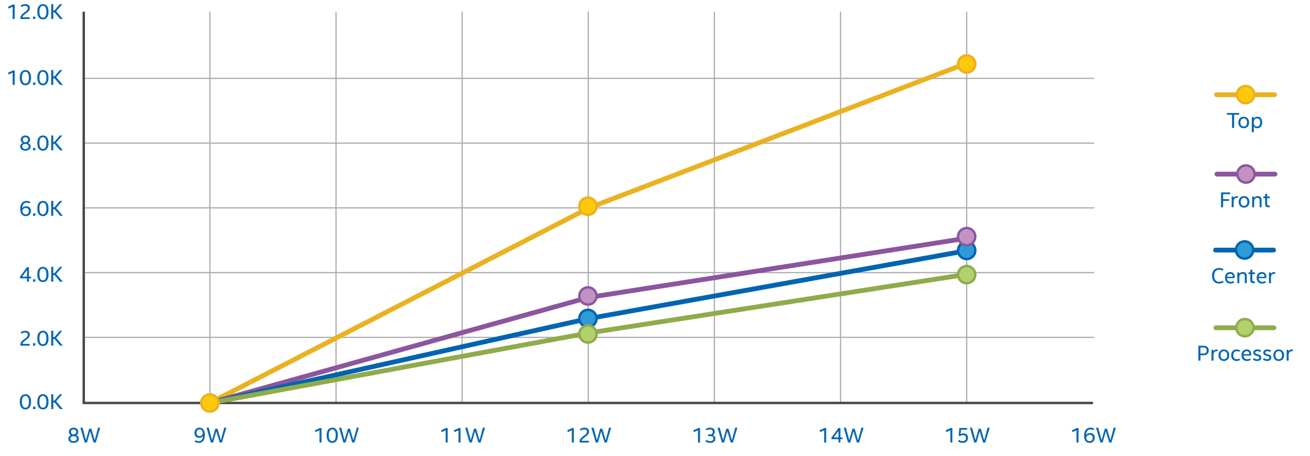

Figure 11: Relative temperature changes for four key points within the compute module as the processor power dissipation rises from 9W to 15W.

Conclusion and Outlook

As the energy sector evolves, the operation of the grid must adapt. The E4S Alliance is a forum where leading utilities and the vendor community have joined to define requirements and create reference designs for the next generation of secondary distribution automation and control. Intel has created an initial hardware reference design based on the E4S Alliance hardware specifications and co-invested with Delta Electronics to commercialize the concept with the intent to use the design in live trials and pilots in the coming years. The platform features two Core-i7-based compute modules and up to four input/output modules with an integrated ethernet switch and integrated power supply. The increasing complexity of electricity distribution to generators and consumers calls for innovative engineering to address immutable physical limitations with sensible, adaptive designs. The E4S Edge Alliance solution offers a step forward by addressing the significant engineering challenge with an open architecture and an effective thermo-mechanical design.