

CASE STUDY

Cloud Service Provider
Data Center Solutions



SafeDX* Deploys Intelligent Data Center Management Solution to Save Data Center Power

Advanced telemetry features in the latest Foxconn* servers, based on the Intel® Xeon® Scalable processor platform and Intel® Rack Scale Design, improve data center power usage efficiency (PUE)



SafeDX provides IT operations outsourcing services and offers cloud services from its own data centers. SafeDX builds and operates many data centers in Central and Eastern Europe, and is capable of supporting client activities anywhere in Europe and Asia. SafeDX leverages Foxconn's IT technology research and development activities in collaboration with many global enterprises.

Executive Overview

Cloud service providers (CSPs) and enterprises are transforming their data centers to achieve better total cost of ownership (TCO) by taking advantage of the latest advances in compute, storage, and networking technology. Operators are also looking to reduce cost by enabling more automation and improving power usage efficiency (PUE). Since cooling equipment consumes 40% of the total data center power budget¹, a growing trend is to save power by operating data centers at higher ambient temperatures. Employing this technique successfully requires precise monitoring of equipment to keep local temperatures within safe maximums. Intel has collaborated with Foxconn and SafeDX (a Foxconn CSP subsidiary in Europe) to develop and test an end-to-end solution to reduce data center power consumption using real-time server telemetry enabled by the Intel® Rack Scale Design (Intel® RSD) architecture.

The SafeDX solution described here uses Intel's Intelligent Data Center Management, the Intel RSD architecture with its open management APIs, and telemetry capabilities built into the Intel® Xeon® Scalable processor platform (also known as Purley). Together, these technologies provide precise, local thermal and power monitoring of racks and individual servers, enabling an effective feedback and control loop to improve data center cooling efficiency. The project was able to demonstrate a 17% cooling energy reduction using this closed loop approach, and is on track to deliver a PUE of 1.5 in deployment.

Business and Technical Challenges

Increasing the compute density of data center racks to meet rapidly growing processing demands while delivering better service are top concerns for CSPs. These are necessary conditions for continued growth. However, CSPs must also optimize their use of capital resources while reducing operating costs to maintain an attractive value proposition and remain profitable. Because power consumption is a major cost for data centers, there is a clear business mandate to optimize energy efficiency through intelligent monitoring and control.

Data center cooling infrastructure helps ensure the reliability of IT equipment by maintaining equipment at or below a prescribed maximum temperature. At the same time, heat transfer becomes more efficient as the temperature difference

between data center exhaust air and outside temperature increases. Consequently, running a data center at a higher temperature will result in higher efficiency, lowering cooling costs. However, the situation is complicated by temperature variations at different locations, which depend on the design of the data center, air leakage between hot and cold aisles, and how equipment (with different energy dissipation characteristics) is distributed across racks. Energy consumption also varies over time and location as workloads change, and as new equipment is brought into the facility. These fluctuations force operators to set temperature well below the optimum for energy efficiency in order to ensure that local “hot spots” do not exceed maximums.

Reducing this safety margin—allowing the temperature to be set at a higher and more efficient value—depends on the accuracy of temperature monitoring. In most real-world situations, however, the connection between server thermal status and the cooling system is indirect and inaccurate. The size of the required safety margin is typically an educated guess based on a few temperature sensors positioned around the facility. The best solution would be to place temperature sensors at the inlet of every server in the data center, and then set the cooling system in response to the warmest hot spot. However, this would be complex and expensive due to the cost of sensors, wired or wireless infrastructure, and data collection and aggregation components.

Fortunately, such an approach is entirely unnecessary because the required telemetry infrastructure is already built into the Intel Xeon Scalable processor platform and the Intel RSD architecture. The next sections describe an energy saving solution developed by Intel and SafeDX in collaboration with its parent company, Foxconn.

Intel® RSD and Advanced Telemetry API

SafeDX provides IT operations outsourcing services, as well as cloud hosting services from its own data centers. The company has embarked on a multi-phase project to meet its immediate and long-term power efficiency goals. Working with Intel, SafeDX created a pilot solution based on the Foxconn R-5211 server, which uses the Intel Xeon Scalable processor platform (Purley) and the Intel RSD architecture. Purley provides advanced telemetry, including temperature sensing, in silicon, while Intel RSD provides a management software architecture and a rich set of open management APIs that deliver detailed temperature and other data to environmental control software.

Figure 1 shows the Intel RSD architecture, which provides an infrastructure to collect temperature data from every server in the data center. The Intel RSD software components include:

- **Pooled System Management Engine (PSME)** collects and relays server telemetry data via Intel RSD APIs over an out-of-band management network (Ethernet).
- **Rack Management Module (RMM)** performs a similar function by providing rack-level resource management and relaying data in response to queries.
- **POD Manager (PODM)**, running on an independent server, acts as the central point of control for communication with the other Intel RSD components within one or more racks. PODM discovers all assets and catalogs their capabilities in its own database, making this information available through its northbound open APIs. It also gathers telemetry information from the PSMEs and RMMs, and relays the data to the Intelligent Data Center Management software on request.

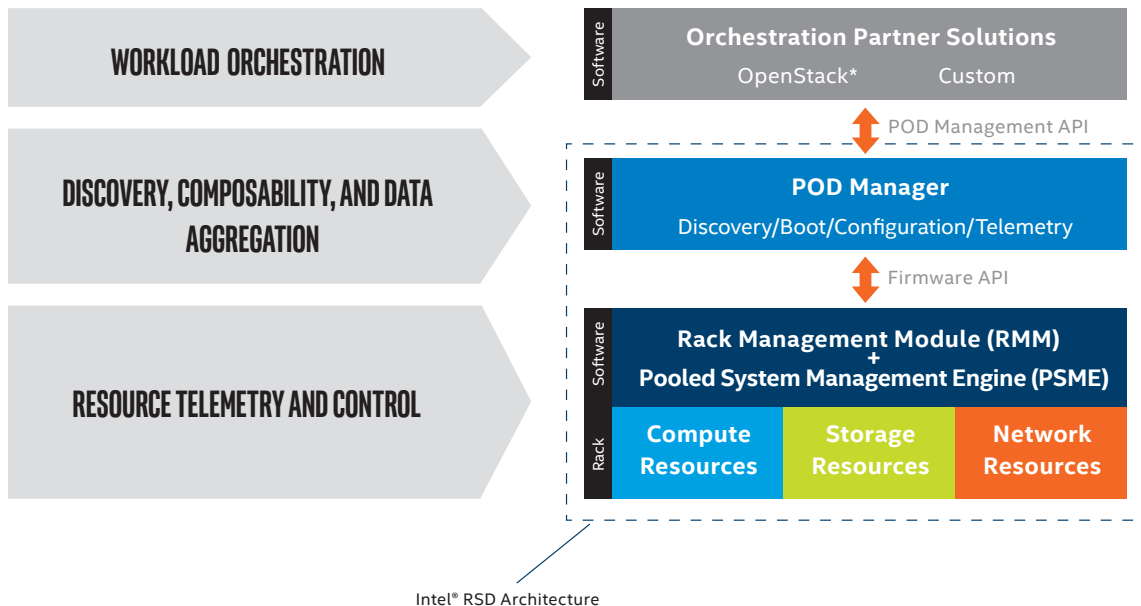


Figure 1. The Intel® Rack Scale Design architecture.

Case Study | Cloud Service Provider Data Center Solutions

Figure 2 shows the overall energy efficiency solution architecture, which employs advanced telemetry (e.g., temperature sensing) at the device level, Intel RSD software components and APIs at the control level to access and relay the sensor data, and the Intelligent Data Center Management software at the management layer.

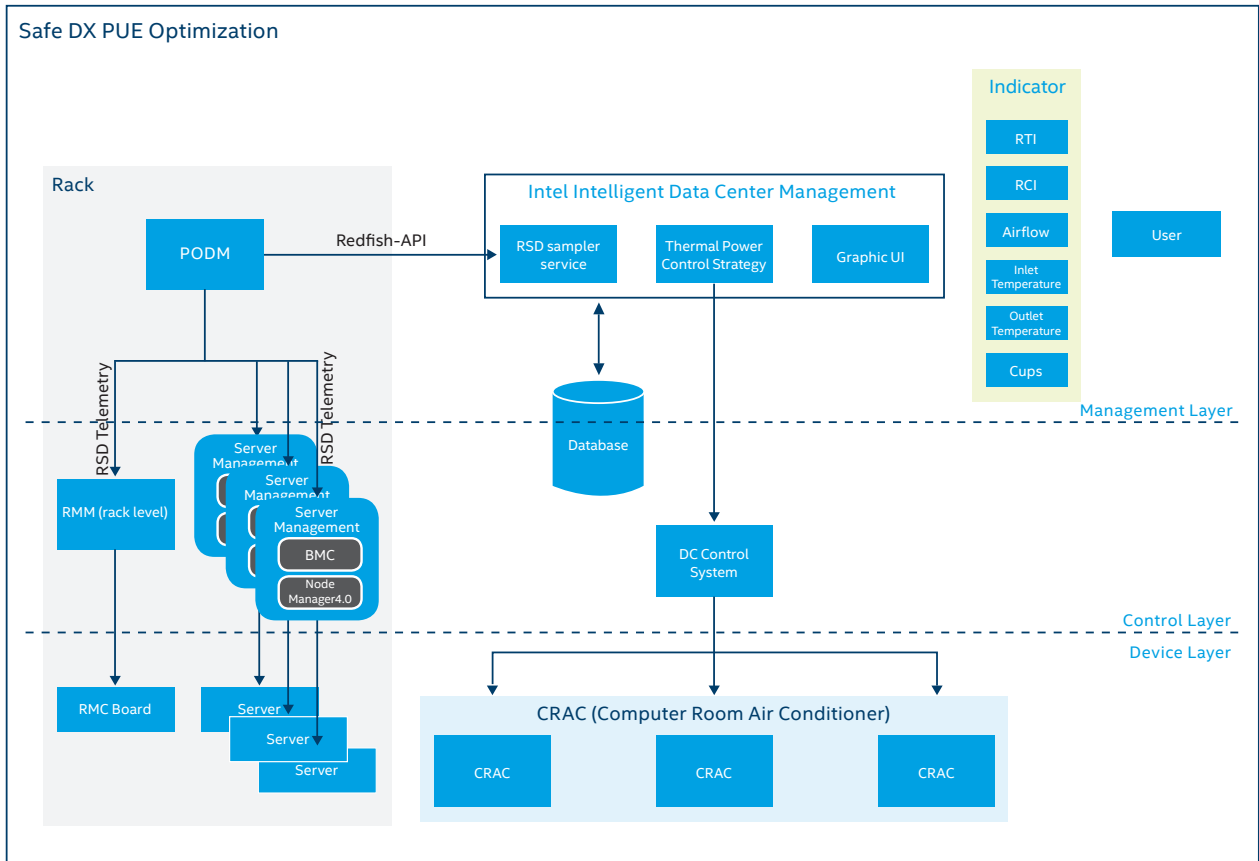


Figure 2. The Energy Efficiency Solution architecture.

Working together, PSME, RMM and PODM collect telemetry data such as fan speed, airflow, CPU thermal margin, DIMM thermal margin, inlet and outlet temperatures, CPU/memory/system utilization, and power consumption info from each server node. PODM relays all this information to Intel's Intelligent Data Center Management, which maintains overall control of the cooling system.

Intelligent Data Center Management provides back-end services including an RSD sampler service, Thermal Control Strategy (TCS) service, and a database. The RSD sampler service leverages PODM APIs to gather thermal and power telemetry data from the racks (servers and rack infrastructure) and stores it in the database. The TCS service analyzes the data and calculates efficiency metrics such as return temperature index (RTI) and rack cooling index (RCI). Based on its analysis, it sends commands to the air conditioning units to adjust thermal settings. Real-time data and metrics are also presented to operators through a GUI dashboard.




Figure 3. Data center configuration with temperature measurements showing the temperature variability across the data center.

Proof of Concept—Current Situation

The SafeDX data center has a typical hot aisle/cold aisle layout. The diagram in Figure 3 shows temperature measurements at various points around the facility, demonstrating the variability of local temperature. The default control system measures return air temperature and attempts to achieve a constant ambient temperature within the room, with the target temperature set to 26C. With just this simple cooling control, the measured data center PUE was 2.0, and the actual temperature at many server inlets was found to be below 24C, i.e., lower than required to meet requirements.

Proof of Concept—Advanced Control Experiment

To test the telemetry control concept, Foxconn R-5211 servers were placed in one cooling zone (red box in Figure 3) for the PUE optimization project. The Intel RSD telemetry APIs collect data at the servers. Once the data is aggregated and forwarded by PODM, the Intelligent Data Center Management software evaluates the data against set rules and algorithms, and generates a range of actions such as adjusting individual rack and server fans speeds, migrating workloads to servers located in areas with lower local temperatures, or, if necessary, lowering the temperature of air from the cooling units. The telemetry capabilities allow operators to continuously monitor servers to make sure they stay within thermal limits as the data center's ambient temperature is raised. The solution can also issue alarms or alerts when temperatures exceed thresholds.

The derived metrics, along with direct server sensor data, also help identify and address data center energy efficiency issues like hot spots, cold spots, recirculation, and bypass. The real-time data from servers can be archived and

analyzed over time to determine if results meet original data center design goals. For example, this information is valuable to data center architects when they need to design a new data center or upgrade an existing data center.

Experimental Results

Based on the data collected for this deployment of Foxconn R-5211 servers with advanced real-time server monitoring based on Intel RSD, SafeDX concluded that the set point for return temperature could be raised from 26°C to 28°C while keeping all local temperatures within ASHARE's recommended temperature range over time. This resulted in a significant improvement in cooling efficiency, with PUE dropping from 2.0 to 1.69. SafeDX believes additional savings can be achieved by exploring an increase in outlet temperature from 28°C to 30°C, and that the project is on track to deliver a PUE of 1.5, representing cooling energy savings of up to 25% (see Figure 4).

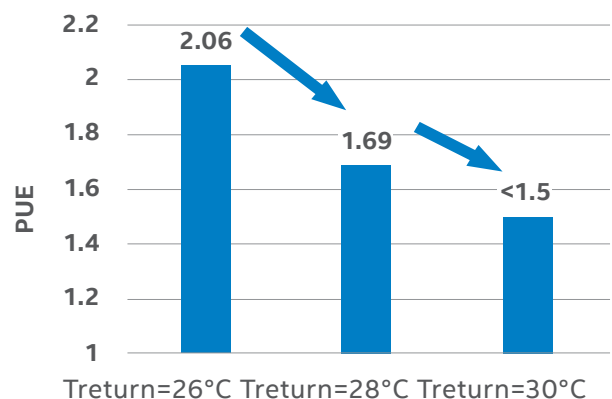


Figure 4. PUE improvements with better telemetry and control.

Conclusion

Intel's Intelligent Data Center Management working with the Intel RSD infrastructure can monitor individual server temperatures in real time across a data center. These metrics help identify hot spots, measure the level of overcooling, detect thermal anomalies, and help quantify the level of hot and cold air mixing. The solution provides more precise and accurate server inlet air temperature data to enable tighter control of cold aisle supply air temperature and airflow (fan speeds)². More precise data over time and space makes more accurate cooling control possible, which enables energy savings without risking thermal safety margins.

Where to Get More Information

For more information on Intel data center solutions, visit intel.com/cloud.

For more information on Intel Rack Scale Design, visit intel.com/intelRSD.

For more information on SafeDX, visit www.safedx.eu.

For more information on Foxconn R-5211 server, visit www.foxconn.com.

Solution Ingredients

Intel's Intelligent Data Center Management solution uses telemetry, metrics, and analytics to solve performance, infrastructure efficiency, and reliability issues in the data center.

Foxconn R-5211 server with Intel RSD solution.

Hon Hai / Foxconn Technology Group is the most dependable partner for joint-design, joint-development, manufacturing, assembly, and after-sales services to global Computer, Communication and Consumer-electronics (3C) leaders. Recognized for its legendary green manufacturing execution, uncompromising customer devotion, and award-winning business model, eCMMS, Hon Hai is the world's most trusted name in contract manufacturing services, including CEM, EMS, ODM, and CMMS.

Find more information at www.foxconn.com.



Intel® Rack Scale Design: A New Architecture for Higher Efficiency, Productivity, and Agility

Intel® Rack Scale Design (Intel® RSD) is an industry-aligned architecture for the data center that provides physically disaggregated, modular building blocks that can be managed as resource pools and dynamically composed to meet specific workload requirements. It includes open, scalable and secure management APIs that support a software-defined infrastructure with interoperability across hardware and software vendors. A composable disaggregated infrastructure reduces overprovisioning, which results in higher utilization and reduced Capex. It makes equipment refresh more economical by avoiding the need to replace entire servers. It also reduces Opex by placing the entire infrastructure under software control. It also makes detailed telemetry data and other critical features easily accessible to a variety of data center management software, including orchestration software for cloud and virtual data center environments.

Intel RSD makes the data center more economical, flexible, simpler to manage, and easier to scale out on demand. It also sets the stage for advanced analytics-based data center optimization solutions. Intel RSD-based products are available today, providing significant cost savings for both new capacity scale out and technology refresh to meet growing demand.

¹ <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=8217258>

² Proper calibration of volumetric air flow shall be performed as described in Intel SPS Services Integration Guide. Find more information at <http://www.intel.com/content/dam/www/public/us/en/documents/technical-specifications/intel-power-node-manager-v3-spec.pdf>

This document and the provided information are only for Intel customers' reference and do not provide any form of warranty (whether it is expressed or implied), including merchantability, fit for a particular purpose, and implied warranties of non-infringement of intellectual property rights. Obtaining and holding this document is not an authorization of the right to use any of the intellectual property rights that are described, shown, or included here. Intel® products are not to be used in medical studies, to save lives, to sustain life, for critical controls, in security systems, or on nuclear energy facilities.

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps. No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request. No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document. Intel technologies' features and benefits depend on system configuration and may require enabled hardware or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software, or service activation. No computer system can be absolutely secure. Learn more at intel.com, or from the OEM or retailer.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Results have been estimated or simulated using internal Intel analysis or architecture simulation or modeling, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance.