IMPROVE DATA CENTER COOLING FACILITY EFFICIENCY THROUGH PLATFORM POWER TELEMETRY

EDITOR(S):

[Edmund Song, Intel]

[Nishi Ahuja, Intel]

[Mohan J. Kuma, Intel]

CONTRIBUTOR(S):

[Yanbing Sun, Intel]

[Lifei Zhang, Baidu]

[Feng Jiang, Intel]

[Xiaozhong Li, Baidu]

[John Leung, Intel]

[Rahul Khanna, Intel]

Executive Summary

Data center operators usually overprovision facility capacity to ensure enough buffer to fulfill peak demand. The overprovisioning brings great pressure to data center total cost of ownership (TCO). Today, the data center management stack has been widely deployed to monitor data center runtime heath status and it gathered tons of data across power, temperature and resource utilization. These data create opportunity to optimize data center efficiency through data intelligences. In this paper, we introduced our practices in cloud environment for using power trend prediction to improve cooling efficiency. Meanwhile, this paper discussed some key challenges and design considerations while enabling IT platform data driven facility control at hyper scale data center, e.g. telemetry collection, messaging mechanism, and management API. Effective interoperability among IT devices, facility and management system is very critical for solution deployment, and the adoption of Open Compute Project design and Redfish API easier system level integration and reduce deployment costs over different systems and different manufacturers.

Table of Contents

1 Introduction

Modern data centers run a lot of different applications with wide variety of workloads. Because the demand of workload is unpredictable for high dynamics of cloud services, usually, data center operators must reserve excess resources to make sure enough margin to fulfill peak demand. For data center capacity planning, it is not one secret that data centers are overprovisioning from servers to facility power and cooling capacity to avoid system faults from running short of critical resources. The cause of this dilemma is lacking accurate tracking of demands as well as instant control method to adjust supply of capacity, so the only reliable means is to overprovision critical resources giving impact to system uptime. But the waste of capital and operating cost from overprovisioning is very significant and brings big pressure to data center total cost of ownership (TCO). Efficient power and cooling provision are critical for data center efficiency and it is very necessary to accurately analyze demands and manage supply accordingly.

The data center stack consists of two parts, as shown in Figure 1, data center infrastructure management (DCIM) system to manage backend hardware capability, and orchestration system to manage virtual resources. Most of cloud service providers have deployed DCIM system to gather information from a variety of devices and meters, e.g. power consumption, temperature, and humidity, and the data center management software can consolidate such information and present them in dashboard to data center operators. For each day, millions of data points are collected from various IT platforms, network equipment and facility system. These data create new opportunity for operators to optimize data center's performance, efficiency and reliability through data intelligences.

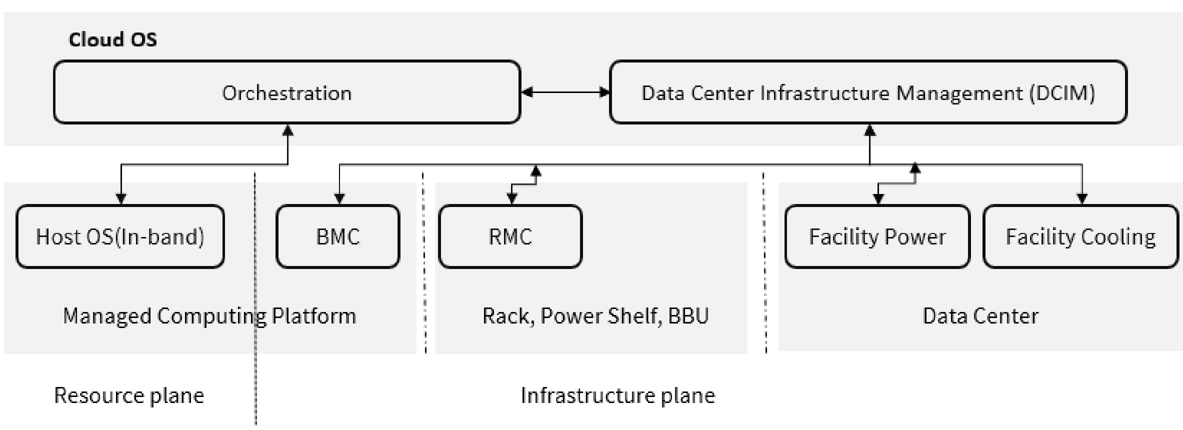

Figure 1 Data Center Management Stack

The overprovisioning of facility capacity causes huge waste of capital and operational cost. To improve data center efficiency and reduce overall cost, the DCIM needs to instantly collect various telemetries, analyze trend of demands, and then supply just right capacity. The telemetry and analytic intelligences can build compressive view of resource utilization (including servers, facility power and cooling) as well as capacity demands. This data view provides important reference to dynamically manage response decisions for capacity request. Such kind of data intelligences help data centers reduce TCO and PUE (power usage effectiveness), moreover it can increase system availability by reducing downtime due to constraints or failures happened in facility system.

The modern computing platforms have provided rich telemetry capabilities, including computing resource utilization, power consumption and data center thermal information. Intel has been worked with leading cloud service providers to leverage computing platform telemetry intelligences for facility efficiency optimization. In this whitepaper, we are going to discuss those key challenges and learnings form these practices with following structure –

- Section 2.1: Discussed data center power telemetry hierarchy and design considerations to build one scalable telemetry framework to support following optimization effort;
- Section 2.2: Presented our work for using power trend prediction to reduce overprovision of cooling capacity;
- Section 2.3: Discussed interoperability challenges and Redfish adoption while integrating IT platform telemetry;

2 Data Intelligences Driven Facility Efficiency Evolution

2.1 Data Center Power Telemetry

The telemetry and analytic capability in DCIM are the fundamental to build intelligent data center. But when we are facing a variety of managed devices and thousands of sensor points, how to systematically collect data, organize data for analysis and integrate data intelligences into control process?

- Watch – to expose infrastructure attributes to data center management system, in this phase, the focus is on what kind data to collect and how to get the data report to system;
- Decide – now that we have data, what kind of things behind the data and what kind of actions to take to support the defined objectives?
- Act – How define control loops around the data we have analyzed to make data center to be more efficient and more automated?
- Learn – Moving forward, we can keep learning from the data over time and build evolution model to adapt to changing environments and workloads.

2.1.1 Power Delivery and Telemetry Architecture

Figure 2 illustrates architecture of data center power delivery and telemetry collection. The PDC (power distribution cabinet) is a device to supply electrical energy and redistribute power for each rack through rPDU (rack power distribution unit), the log agent residing in PDC keeps measurement and record of rack power consumption.

To get energy use is more efficiently distributed amongst server nodes, the modern rack designs, e.g. OCP racks, centralize power supplies into single power shelf to allow sharing of power capabilities at rack. The rack with centralized power shelf uses RMC (rack management controller) to monitor input and output of rack power. The traditional rack contains servers that use individual power supply unit, there is no physical agent to collect rack power consumption, and some software agent runs as the virtual aggregator of rack power by aggregating power data from those servers within same rack.

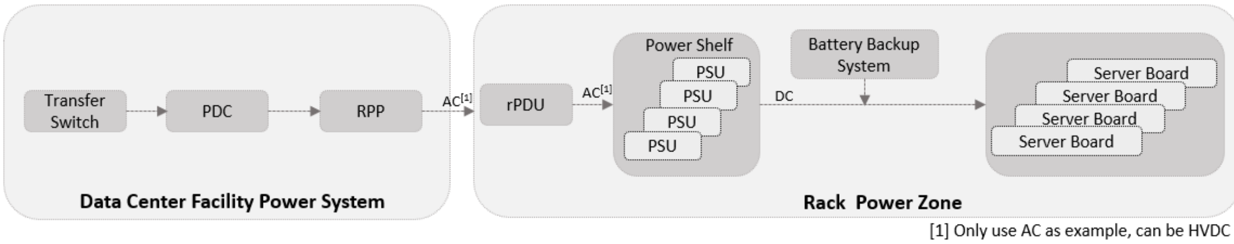

Figure 2 Data Center Power Delivery Architecture (with shared power shelf)

To the server chassis level, the hot swap controller measures system power consumption, and the BMC (board management controller) provides interface to access system power. From statistical data, processor and memory contribute 80% of overall system power consumption [1], and in modern computing architecture, there are digital sensors in processor and chipset to provide finer granularity power telemetry for processor and each memory domain, the processor and memory domain power metrics can be accessed either through Intel Node Manager or in band model-specific register (MSR). The rich power telemetry from data center to server board components can help data center operators to build one comprehensive picture of power demand and allocation.

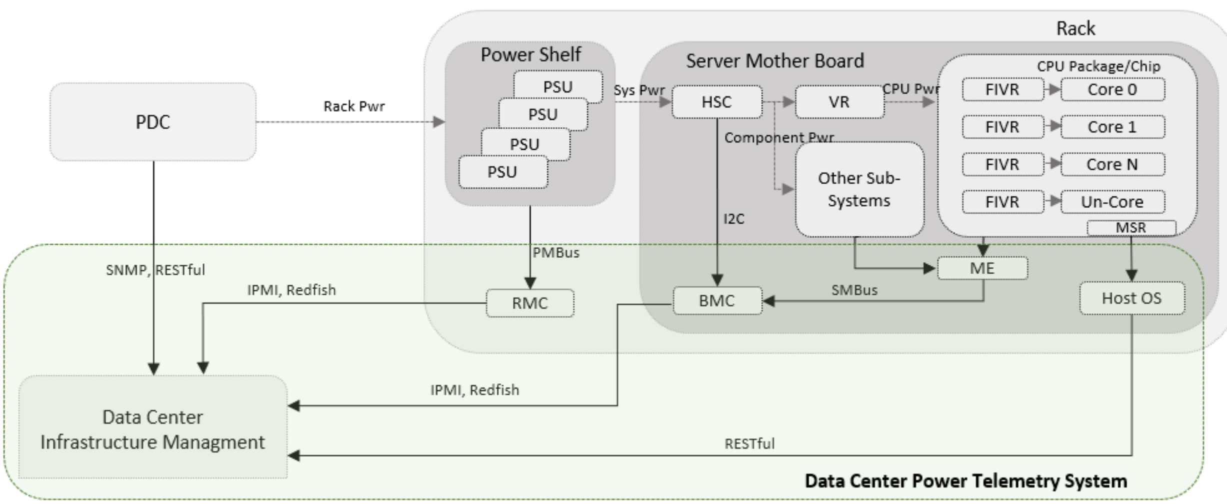

Figure 3 Data Center Power Telemetry Infrastructure

2.1.2 Design Challenges

Nowadays, the DCIM system is mainly to serve monitoring purpose, and the collected power telemetry is only used for health monitor and offline analysis for capacity planning. There are no strict requirements for polling frequency and data accuracy, for example the sample period of power telemetry is configured at minute level to mitigating impact to network traffic. With moving to use power telemetry to build holistic view to manage some critical facility control, those data points gathered from different servers and devices shows outstanding correlations in temporal dimension, and the deviation of time dimension beyond limit may lead to large errors or even wrong control strategy. Time series data collection turns to be one strict requirement. Following sections discuss some major challenges while applying time series data into DC power telemetry system.

- Clock Synchronization: The DCIM gathers millions of data point across thousands of sensors and the data message reported by devices is required to carry timestamp information. The management software extracts timestamp information from data message as the reference to process the temporal relevance. Each device's clock needs to be synced, otherwise, the timestamp extracted from data message will lose meaning. The network time protocol (NTP) is used to synchronize time across multiple servers and facility equipment and to ensure each end-point for data collection, the aggregator, data consumer (e.g. data analysis agent) and the actuator in control path to be synced with same time server.
- Sampling Period: The fixed internal sampling for data collection is one effective way to reduce network traffic and data storage while keeping statistically correct analysis of power insights. The actual power consumption of server board changes promptly but considering the front-end BMC software need a certain of time to poll the power consumption, at 500ms or 1s level, the second level granularity is enough to record measurable changes for server board. The cumulative effect of power consumption at rack wide smooths the intensity of fluctuations, and the required sampling time increases as cumulative level arises. The pie chart in figure 4 shows the distribution of fluctuation time. The PDC can record changes as they occur. By studying the distribution time of fluctuation from PDC power logs, we can understand required sampling time granularity. The pie chart shows more than 99.7% of change time fall within the interval of 20 seconds or more.

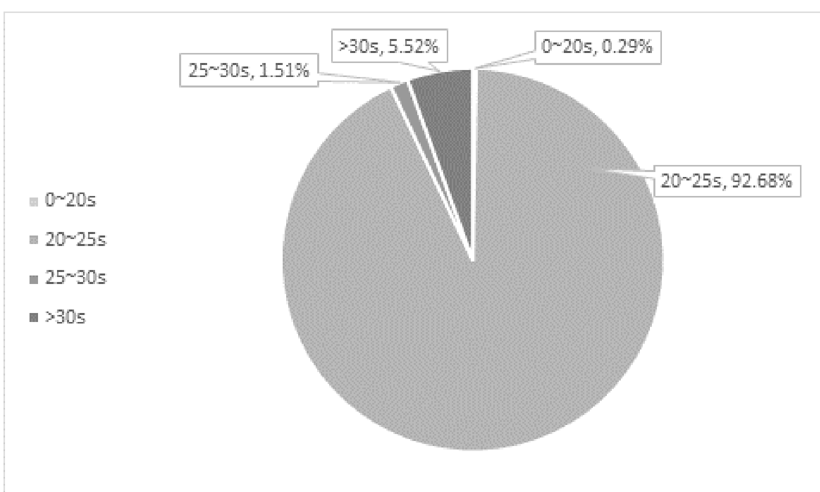

Figure 4 Power Fluctuation Interval Distribution

- Publish-subscribe (Pub-sub): In hyper scale data center with thousands or ten thousand of servers, the publish-subscribe model provides better scalability than traditional request-response model. The client agent that publishes data send it only when data change happens, and the client agent that subscribes to the data receives it from the broker or server only when data changes. The Pub-sub messaging mechanism can dramatically reduce traffics running over management network because it eliminates duplicated network traffic to transmit identical data.

2.2 Power Prediction in Cooling Optimization

Almost all power dissipated in data center is converted to heat, and to ensure data center components running efficiently and safely, the data center need to evacuate heat timely and maintain the environmental temperature within safe range. The heat transfer in modern data center relies on the CRAC (Computing Room Air Conditioner) to provide with cold air through heat exchanger between server exhaust and heat sinks, e.g. chilled water. Failure to exchange heat out of data center causes hotspots generation and may overheat computing components in data center. The overheating higher component failure rates and affect IT equipment lifespan, or directly lead to equipment failures and system down time. To prevent various issues from insufficient cooling, data center operators often overprovision cooling capacity and cause great pressure to efficiency and cost.

2.2.1 Power Trend Prediction

Therefore, understanding heat changes as it occurs and allocating proper capacity adaptively prevent risks from both aspects of overheating and overprovision. The precise control of cooling supply is essential to maximize IT equipment availability and data center efficiency with two objectives: (i) to minimize the energy consumption from facility cooling; (ii) to reduce of overheating risk from improper supply of cooling capacity, e.g. time lag in cooling control response, or unbalanced distribution. Diagram 5 shows control logic of using power trend forecast to improve the precise of cooling control.

- The DCIM software collects the power telemetry information, from server node, RMC and PDC;
- The power trend prediction module analyzes power trend and predict power consumption and corresponding heat at time $(\mathrm{T}+\ \boldsymbol{\tau})$;
- The heat variation trend predicted in step2 is used as reference to forecast cooling demand at time $(\mathrm{T}+\boldsymbol{\tau})$, and then drive cooling system act earlier.

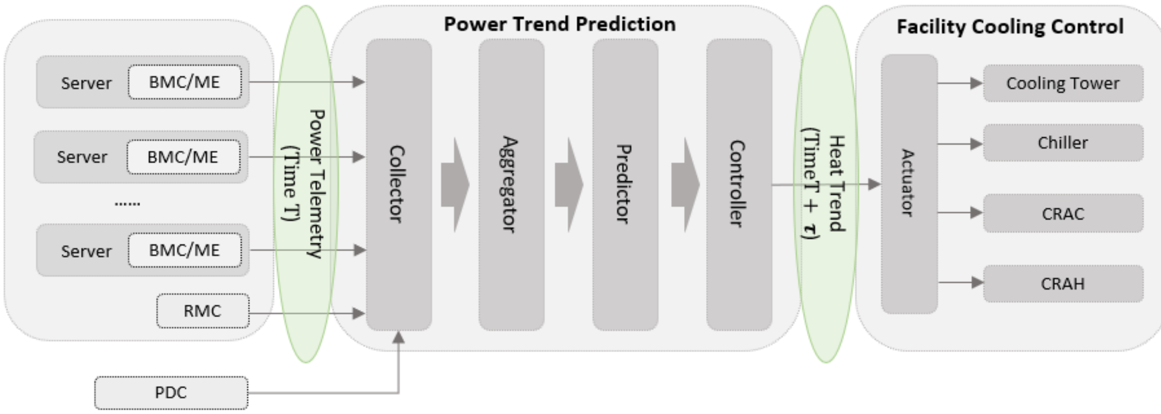

Figure 5 Power Trend Prediction

Equation 2.1 indicates the total amount of produced heat (H_T) within datacenter, which is one summation function of energy ($\boldsymbol{P_i}$) dissipated into data center by all IT devices, for example computing server, storage, network switches etc.

$$H_T = f \times \sum_i^n \boldsymbol{P_i} \qquad (2.1)$$

Where $\boldsymbol{H_T}$ represents the total amount of heat at time $\mathbf{T}$, $\boldsymbol{f}$ is the conversion factor from energy consumption to heat, and $\boldsymbol{P_i}$ is the amount of power consumption for specified IT equipment at time $\mathbf{T}$.

The power is usually reported in watts (W) and heat is reported in British thermal unit (BTU) per hour (BTU/hr) and they are interchangeable. The conversion between power and heat is 1w = 3.412 BTU/hr [2] .

This (Figure 6) is framework of power variation trend prediction. There are three phases: training phase, prediction phase, and calibration phase. The prediction relies on the analysis of historical power relevant telemetries ($\boldsymbol{I_{[1..T]}}$). The power relevant telemetry set ($\boldsymbol{I_T}$) is one feature vector with thermal information, power, and resource utilization (e.g. CPU and memory) as below:

$$I_T = [Inlet\ Temp_T, System\ Power_T, CPU\ Power_T, CPU\ \ Util_T, Memory\ \ Util_T\ \ ...]$$

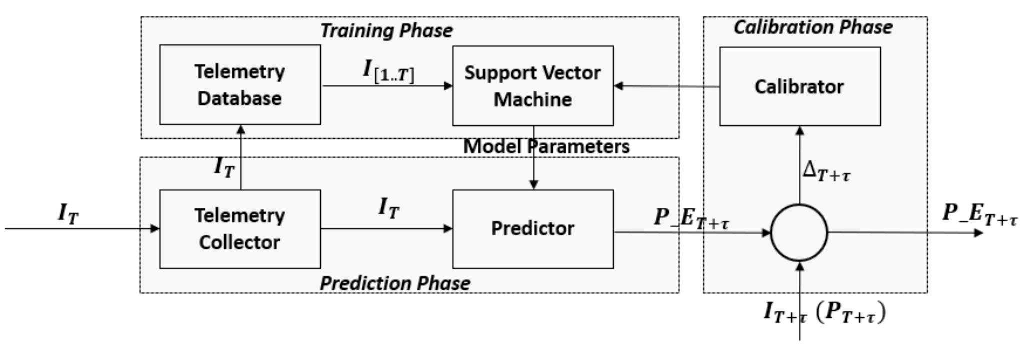

Figure 6 Power Trend Prediction Algorithm

The collector collects $\boldsymbol{I_T}$ from DCIM management system and store data into time series database. The support vector machine stream historical data ($\boldsymbol{I}_{[\boldsymbol{1}..\boldsymbol{T}]}$) into memory to build classification function and regression model, which is feed into feed into prediction phase. The predictor estimates the power trend variation ($P\_E_{\boldsymbol{T+\tau}}$) for next prediction cycle (τ). The calibrator phase is designed to capture the delta between forecast power consumption ($\boldsymbol{P\_E_{T+\tau}}$) and observed value ($\boldsymbol{P_{T+\tau}}$) at time $\boldsymbol{T}$ and feedback the delta ($\Delta_{\boldsymbol{T+\tau}}$) to support vector machine to continually calibrate the parameter of prediction model.

In our experiment, considering the power consumption change model at rack or data center wide is a long and slow procedure, the time window for historical telemetry analysis is set as month based (e.g. 30 days historical data). On the other hand, the prediction cycle (τ) is set as minutes granularity in the light of the speed at which IT equipment dissipates heat. The experiment in following sections evaluated 5 minutes and 10 minutes prediction separately.

2.2.2 Proactive Cooling Method

Figure 7 depicts the transformation from reactive cooling model to proactive cooling model through adding power trend prediction. The reactive cooling method is widely adopted by conventional air-cooling system at data center. The data center management monitors heat fluctuation through data ambient temperature sensors deployed at data center. The cooling controller uses the observed thermal information as the input to drive valve opening ratio or/and air flow rate to adjust cooling supply.

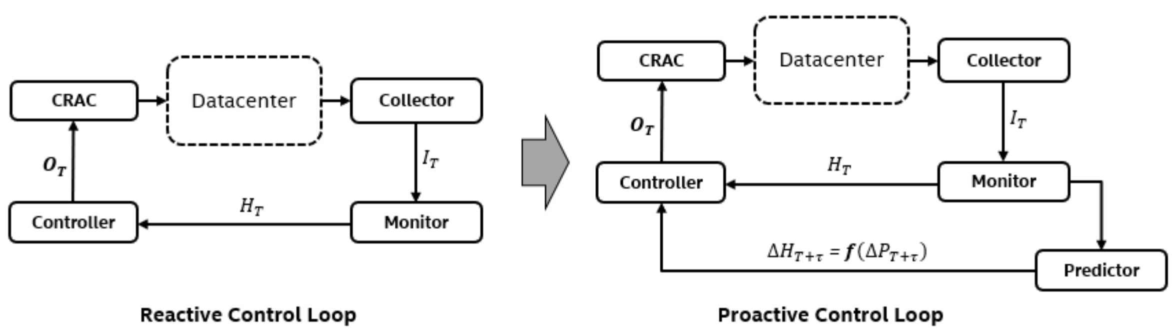

Figure 7 Reactive Control vs. Proactive Control

For proactive cooling control method, the power predictor is added to into the cooling control loop. The power predictor estimates the heat fluctuation trend at prediction window ($\boldsymbol{T+\tau}$) through analyzing the power fluctuation at time $\boldsymbol{T}$. Thus, the cooling controller can use the observed thermal information at time $\boldsymbol{T}$ and estimated thermal fluctuation trend to drive cooling system to supply required cooling capacity ahead of one prediction window ($\boldsymbol{\tau}$), as shown in Figure 8:

- Upper Diagram: The dash dot line is the power consumption curve (P_{total}), and the dot line is the demand of cooling capacity (C_{demand}). The change of the cooling demand can be estimated through the power variation analysis.
- Bottom Diagram: The dash dot is the supplied cooling capacity ($C_{reactive\_supply}$) from reactive cooling control method and the dot line indicates the supplied cooling capacity ($C_{proactive\_supply}$) by proactive cooling control method. The result is an earlier response ($\tau = \mid T_{re} - T_{pro} \mid = \mid T_2 - T_1 \mid$) in supplied cooling capacity to meet the cooling demand.

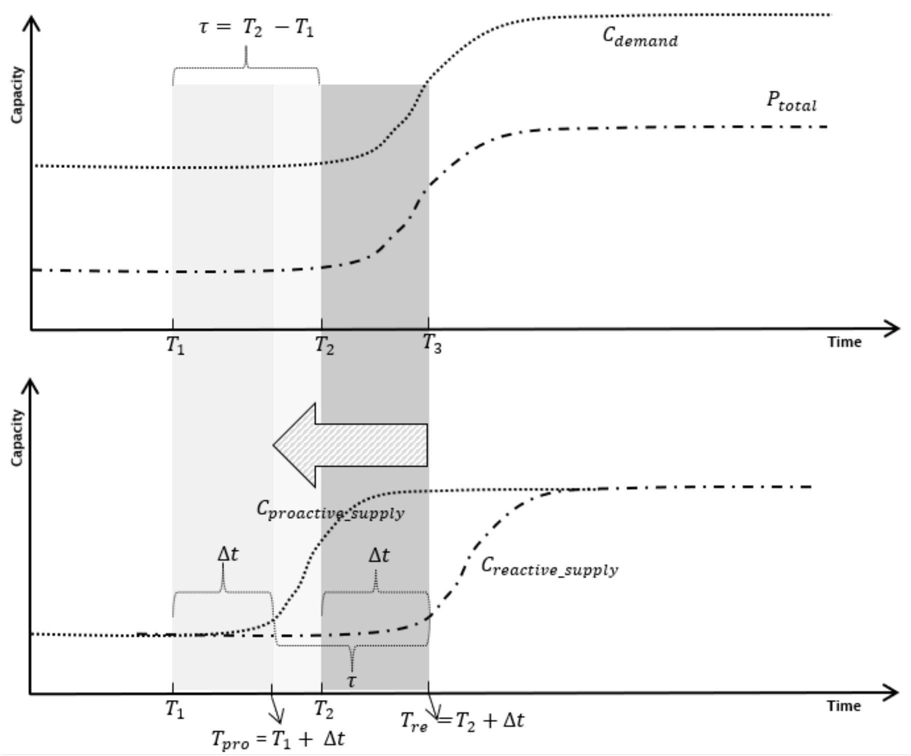

Figure 8 Proactive Cooling Control

2.2.3 Evaluation and Results

For conventional air-cooling method, the return air temperature (T_{return}) directly reflects the cooling effect. If the return air temperature is higher than the setting point, the IT equipment run in a overheated risk status because of insufficient cooling capacity. From time dimension, the longer the duration of return air higher than the setting point, the higher the risk of running into thermal issue. The area surrounding by T_{return} curve and $T_{setpoint}$ can be used to measure the risk level of cooling management. In figure 9, the undershoot area below setting point ($T_{setting\ point}$) is over provision cooling capacity, and the overshoot area above setting point is over heat IT device.

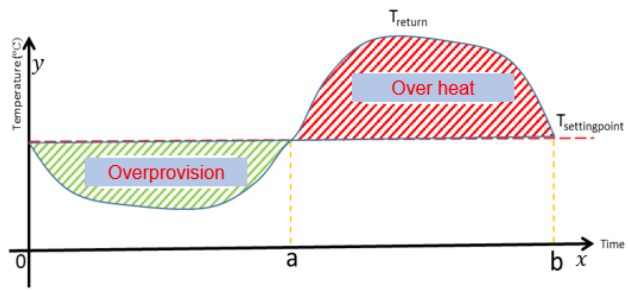

Figure 9 Risk Level Evaluation for Cooling Management

Cooling Risk Indicator (R_c) is the integral calculus of over shooting area:

$$R_c = \int_b^a (T_{return\_x} - T_{setting\ point}) \times D(x) \quad (2.2.3)$$

where: T_{return_x} is the temperature of return air (ºC);

$D(x)$ is its duration (minutes);

$T_{setting\ point}$ is set point of the return temperature(ºC).

Figure 10 shows the temperature of return air ($\mathbf{T_{return}}$) under different scenes, no power prediction, 5 minutes prediction window, 10 minutes prediction window and 15 minutes prediction window. The overshoot area is reduced after adding power prediction. To avoid over heat risk, we usually reserve excess cooling buffer via overprovision cooling capacity. By adding the power prediction, the over shoot area is reduced, which implicates reduce of over provision area and improve cooling efficiency.

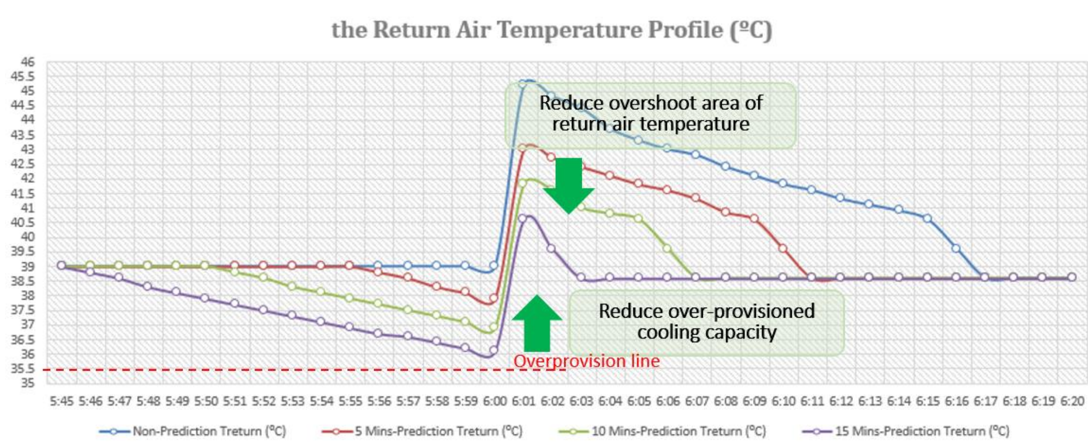

Figure 10 Return Air Temperature Profile

2.3 Data Center Management API

Deploying and managing hyper scale data center with a large number of servers, network devices, and facility equipment are facing a lot of challenges: complicated, time-consuming, and security breaches. In data center management, the IPMI (Intelligent Platform Management Interface) is the most adopted industry standard, but it is complex, and outdate. The cloud service providers prefer RESTful style protocols for cloud service communication, for its adaptability for different applications, simplicity on management, and extensibility to support service evolution. DMTF's Redfish® is a standard API designed to deliver simple and secure management for modern scalable data center infrastructure. Redfish leverages RESTful interface and JSON payload to expose information to the modern tool chains [3].

2.3.1 Redfish in DCIM

As mentioned in last chapter, the DCIM software collects telemetries from IT platforms and facility equipment, conducts data analysis, and builds control sequences to drive facility cooling to supply proper capacity. Figure 11[4] depicts fundamental components involved into this procedure, like IT devices, power and cooling system, data center management system. It is difficult to get all these components to talk each other because of lacking common management protocol. For example, the cooling system relies on fieldbus interface, such as mod-bus or 485; the power system uses SNMP, IPMI, and even web-based API for modern PDUs. The situation is much better for IT equipment because most of IT equipment support IPMI and now is moving to Redfish. The diagram shows scenes to transform all communications to Redfish schema, which is quoted from Redfish data center infrastructure management (DCIM) specification [4]. Here, the Redfish protocol help create a holistic telemetry and control interface throughout data center components. Those components with legacy management protocols can interact with DCIM core system via protocol proxy. The proxy translates the format of telemetry data and control commands from/to Redfish protocol. No matter what kind of protocol or data format behind, the data is converted to Redfish and run over RESTful API and JSON format, which greatly simplifies the complexity of integration and the cost of deployment.

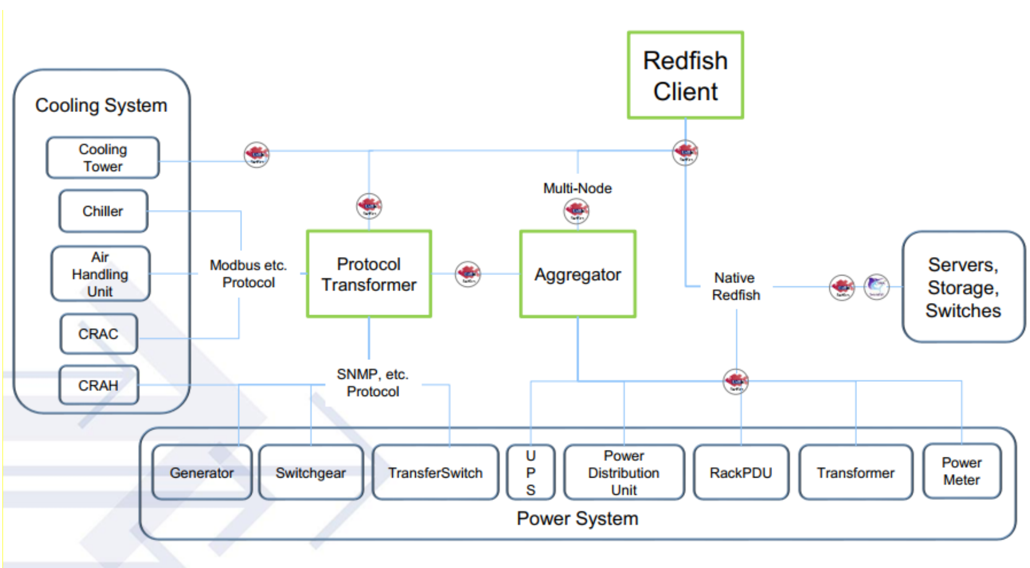

Figure 11 Redfish vs. Data Center Infrastructure Management [4]

2.3.2 Mockup Examples

The Redfish is creating one possibility to establish common management interfaces to connect facility management with IT equipment based on same protocol and resource model to interoperate each other. In this section, we will introduce some Redfish mockup examples to use Redfish schema for rack power metrics and prediction.

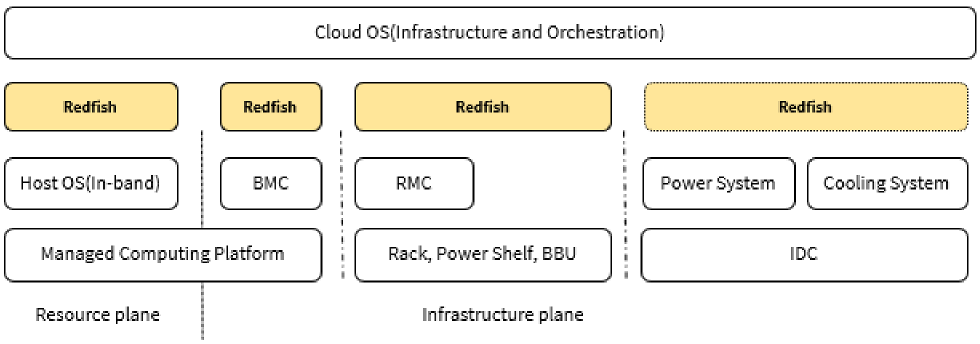

Figure 12 Redfish as Data Center Stack Management API

Mockup Example – Rack Power Metrics and Prediction:

```
{
        ...
        "PowerControl": [
        {
                "@odata.id": "/redfish/v1/.../Power#/PowerControl/0",
                "MemberId": "0",
                "Name": "Rack Power Control",
                        . . .
                "PowerMetrics": {
                        "Interval": “P10S”,
                        "MinConsumedWatts": 7200,
                        "MaxConsumedWatts": 7900,
                        "AverageConsumedWatts": 7600
                },
                "PowerPrediction": {
                        "PredictionWindow": “P600S”,
                        "PredictedInstantWatts": 7300,
                        "PredictedAverageWatts": 7200
                },
         . . .
        }
  ],
 . . .
}
```

PowerMetrics represents power telemetry information gathered from server platform or any other devices with power meters, e.g. power centrialized cabinet or power distribution unit. The property Interval indicate the sampling period in which the PowerMetrics is measured. The Interval is duration format from ISO 8601[5], P10S stands for the duration to collect the power metrics is 10 seconds:

- P is the duration designator (for period) placed at the start of the duration representation.
- S is the second designator that follows the value for the number of seconds.

PowerPrediction represents the time series prediction for power consumption. It is used to describe the estimation of power consumption for next prediction window. The PowerPrediction define 3 properties:

- PredictionWindow (read-write): The time interval over which the power metric prediction applies. Using ISO 8601 format. P600S stands for the prediction window for rack power estimation is 600 seconds (5 minutes).
- PredictedInstantWatts(read): The predicted power consumption value, the unit is watts.
- PredictedAverageWatts(read): The predicted average power consumption value, the unit is watts.

3 Conclusion

The overprovisioning of facility capacity has been taken as one best practices at before to prevent system faults from running short of mission critical resources, e.g. cooling or power. The excess resource provision consumes capital investment and operational cost and lower down ROI (return of investment). The data intelligences driven data center efficiency optimization is gradually accepted and becoming possible with most of cloud service providers deployment data center infrastructure management system. This paper introduced our practices with cloud service providers to optimize facility cooling efficiency with the help of holistic data center power telemetry. The data intelligences from power telemetry analysis can help the facility cooling system to adjust cooling capacity earlier and reduce overprovisioning. The interoperability among IT devices, facility equipment and DCIM system is supposed to be one big challenge to impact integration complexity and deployment, the adoption of Redfish help builds one common resource model and management API to support power telemetry driven facility cooling optimization. Driven by the values of sharing, openness and collaboration, the Open Compute Project has been establishing effective ecosystem to facilitate innovations for data center. As part of the OCP, the Data Center Facility (DCF) project is developed to focus on power, cooling, facility management and control as well as operations.

4 Glossary

TCO: Total Cost of Ownership.

DCIM: Data Center Infrastructure Management.

OCP: Open Compute Project

RMC: Rack Management Controller

BMC: Board Management Controller

ME: Intel Management Engine

OS: Operation System

DC: Data Center

BBU: Battery Backup Unit

PUE: Power Usage Effectiveness

PDC: Power Distribution Cabinet

PDU: Power Distribution Unit

rPUD: Rack Power Distribution Unit

RPP: Rack Power Panel

VR: Voltage Regulator

FIVR: Fully Integrated Voltage Regulator

PSU: Power Supply Unit

HSC: Hot-swap Controller

CRAC: Computing Room Air Conditioner

NTP: Network Time Protocol

Pub-sub: Publish Subscription

BTU: British Thermal Unit

IPMI: Intelligent Platform Management Interface

DMTF: Distributed Management Task Force

CRAH: Computing Room Air Handler

RESTful: Representational State Transfer

JSON: JavaScript Object Notation

SNMP: Simple Network Management Protocol

API: Application Interface

5 References

1. [Increase Your Data Center's IQ] - https://www.intel.ca/content/dam/www/public/us/en/documents/solution-briefs/cpg-modern-infrastructure-telemetry-business-brief-for-distribution-only.pdf
2. [Data Center Power and Cooling White Paper]. https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/unified-computing/white_paper_c11-680202.pdf
3. [Redfish]. https://www.dmtf.org/standards/redfish
4. [DSP-IS0005] https://www.dmtf.org/content/updated-redfish-dcim-schemas-released-public-review
5. [ISO_8601]. https://en.wikipedia.org/wiki/ISO_8601

6 License

7 About Open Compute Foundation

The Open Compute Project Foundation is a 501(c)(6) organization which was founded in 2011 by Facebook, Intel, and Rackspace. Our mission is to apply the benefits of open source to hardware and rapidly increase the pace of innovation in, near and around the data center and beyond. The Open Compute Project (OCP) is a collaborative community focused on redesigning hardware technology to efficiently support the growing demands on compute infrastructure. For more information about OCP, please visit us at http://www.opencompute.org

8 Appendix A. [Title]