BodenTypeDC Cooling Innovation

For air cooled data centers, energy can be wasted by providing too much air to the IT space from the cooling equipment. Traditionally, air cooled systems are controlled by providing a set temperature to the cold aisles, which invariably leads to slow response of the cooling systems to the thermal swings in the data center created by large changes in the IT workload.

LEVEL1 – INSTANT RESPONSE OF THE COOLING EQUIPMENT TO THE VARIATION OF SERVER WORKLOAD (ACHIEVED)

main goal: immediate and accurate control of the coolers based on the current status/expected cooling requirements (used energy is proportional to the cube rule of the fan's revolutions per minute)

Experiment 1: POD1 – monitoring of the power draw of the 480 servers provides a good indication of the heat load for the cooling system – the kW draw, average server fan speed, average CPU temperature are all provided via registers that are accessible to the cooling system enabling instant response to changes in the IT space power draw. This in turn maintains a stable instantaneous Power Usage Effectiveness (PUE) as seen in the following graph.

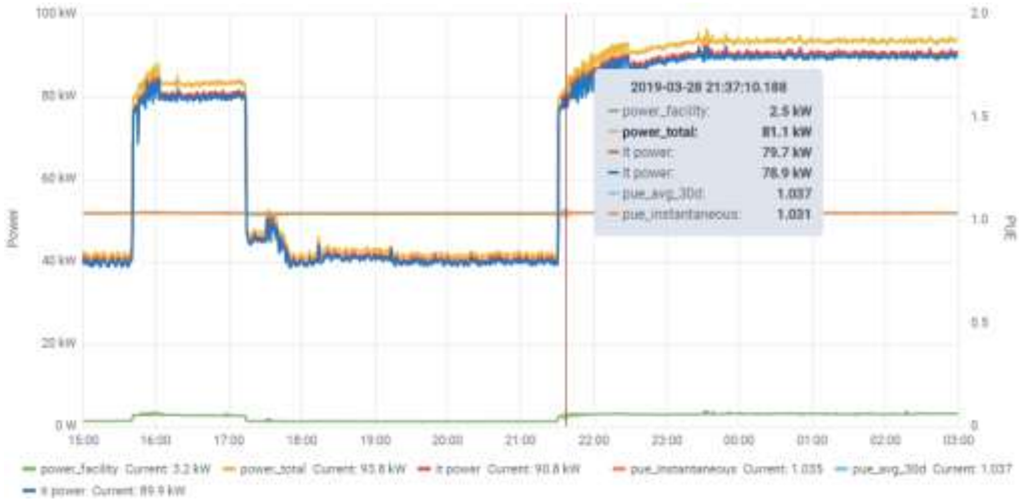

This cooling setup has enabled POD1 of the BodenTypeDC project to run with a thirty-day rolling average PUE of 1.03 meaning that for each kWh of energy consumed by the IT systems, 30Wh is consumed by the facility, namely the cooling fans, power distribution losses and the lights.

Experiment2: POD3 - reading all information from all Application Specific Integrated Circuits (ASICs) of the test partner and providing information to the coolers (this experiment also exemplifies information about the weaknesses of PUE, since higher server fan usage is demonstrating a lower PUE, but possibly more energy usage)

LEVEL2 – CONSTANT TEMPERATURE IRRESPECTIVE OF THE WORKLOADS (TO BE EXECUTED)

The 480 servers in POD1 each contain two central processing units (CPUs) and two fans. The BodenTypeDC project has developed a mechanism to easily provision a variety of IT workloads that would cause the CPUs to increase their thermal flux and at the same time the two fans respond by increasing their speed and power consumption in response to this increase in heat load. The Open Compute Project servers in POD1 have each been fitted with a fan card that sits between the two fans per server and has many functions, the two main functions in operation currently is to pass through the fan signal from the motherboard to the fans so that their speed is controlled by the BIOS of the motherboard, but also to monitor directly the fan energy and RPM. The fan energy can be used to modify the PUE to what is closer to a total usage effectiveness (TUE) introduced by the Energy Efficient High-Performance Computing Working Group by including the fan energy as part of the cooling and hence not in the denominator of the PUE.

main goal: when the servers in POD1 have no workload their CPU temperatures are around 40 to 50 degC, but at peak workload they ramp up to around 70 to 80 degC. Overiding the usual control of the fans the objective would be to maintain CPUs at a constant temperature of say 70 degC.

Experiment 1: creating a control algorithm that runs as a daemon on each of the 480 servers to control the fan speed of the server to maintain both CPUs betwen 70 degC at the front and 80 degC at the back irrespective of the workload would generate interesting and until now never achieved energy saving arrangement. At low and part workloads we would anticipate a similar PUE value, but a lower energy consumption at peak workloads the PUE would continue to be maintained by the energy savings over level 1 would be insignificant. This is good because most data centers run at part load all of the time.

Experiment 2: involving a workload distribution approach that attempts to orchestrate workloads on servers at their sweet spot (this is temperature, pressure and humidity dependent) on top of the previous experiment would be considered the „holly grail".

LEVEL3 – REFINED COOLING CONTROL (TO BE EXECUTED)

main goal: reach the same cooling results as with level 2, but with even lower energy usage by applying simple physics in a sophisticated way.

Experiment1: fan speed distribution in the POD that results in a closer to zero pressure difference between the cold and hot aisles when server airflow consumption varies from close to the coolers to that at the „other end" of the cold aisle

Experiment2: workload distribution, orchestrating deployment on servers to their optimal operation in the POD that will have the same cooling result as the previous experiment - not with the control of the server fans but with the distribution of the workload amongst the servers.