CENTER OF EXPERTISE FOR ENERGY EFFICIENCY IN DATA CENTERS

BERKELEY LAB

Accessing Onboard Server Sensors for Energy Efficiency in Data Centers

September 20, 2021

This work sponsored by FEMP

Magnus Herrlin, Ph.D.
Lawrence Berkeley National Laboratory
One Cyclotron Road
Berkeley, CA 94720

Federal Energy Management Program (FEMP)

This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Federal Energy Management Program, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Acknowledgements

Chris Payne, LBNL
Steve Greenberg, LBNL
Hannah Stratton, LBNL
Ryan Fogle, EPA
John Clinger, ICF
Mike Patterson, Intel (retired)

Table of Contents **Page**

Executive Summary 4

1. Introduction 5

2. ENERGY STAR for Computer Servers 7

3. Network Protocols and Interfaces 9

4. Data Center Management Solutions 19

5. DCIM and Data Center Networking Tools 21

Conclusions 29

Glossary 31

References 36

Executive Summary

A data center is a facility that houses a compute infrastructure, typically computer servers, storage, and networking equipment as well as critical support infrastructure such as cooling and electrical systems. Monitoring this gear provide data center operators with critical insight that can help them reduce energy use. This report reviews opportunities to leverage existing internal (onboard) server sensors for monitoring rather than discrete external sensors for physical parameters such as intake air temperature and input power. Such an approach has many benefits for data center management, including higher data resolution and lower overall costs. Accessing onboard sensors is potentially a disruptive technology since hundreds of servers can be monitored without any external measurement equipment.

Unfortunately, accessing onboard server sensors is often not well understood. The objective of this report is to make this technology better known to help increase the use of the onboard sensors and thereby be in a better position to manage both IT and facility energy.

Today's servers have nearly universal connectivity through multiple network protocols and interfaces. And, most Data Center Management solutions allow multi-vendor, multi-protocol, and scanning and detection of devices. In addition, most solutions use agentless monitoring. Data Center Management solutions not only start to integrate but also provide connectors to other systems to meet the need for a more holistic view of the data center. All in all, these versatile capabilities make it easier and safer than ever to monitor and manage these critical facilities. By selecting a commercial product from one of the three Data Center Management solution categories discussed, one does not need to become an IT expert since these solutions do most of the work. A number of more specialized but less versatile tools are also listed in the report.

Current and projected future needs should be carefully evaluated before investing in a particular solution. Cost is not insignificant for comprehensive Data Center Management tools. Need for customizations, after-sales support, and professional services can also be costly for sophisticated software.

Although accessing onboard sensor data is generally a superior technology compared with using external sensors, there are some data centers for which this approach may be less suitable. One such example is certain mixed environments, including colocation data centers where access to servers may not be allowed by the server owner.

1. Introduction

A data center is a facility that houses a compute infrastructure, typically computer servers, storage gear, and networking equipment. The infrastructure also includes support systems such as cooling and power systems. Data centers are critical to the mission of many types of organizations. Data center operators are challenged to address key mandates: deliver reliable services, find operational efficiencies, and ensure resiliency and flexibility. This report focuses on improving operational efficiencies of servers in data centers. Data center servers use more energy than all other IT equipment combined, and energy savings at the servers will cascade through the support infrastructure.

Monitoring is a prerequisite to cost-effective server operation. External sensors (A in Figure 1) have traditionally been used for physical parameters such as intake air temperature and input power although internal (onboard) server sensors exist for most of the same parameters. The internal sensors can be accessed either through the server's operating system (B in Figure 1) or through a Baseboard Management Controller or BMC (C in Figure 1).

Tapping data from the onboard sensors has a number of benefits for data center management, including higher resolution/accuracy and lower overall costs. In fact, using onboard server data is potentially a disruptive technology in that hundreds of servers can individually be monitored without any external sensors. But, accessing these sensors is not always well understood. The purpose of this report is to improve this understanding so that data center operators are better positioned to reap the benefits.

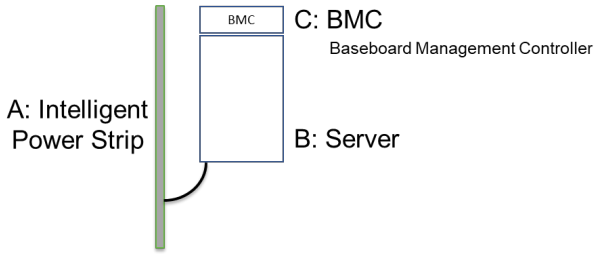

Figure 1. Server input power monitoring.

The advantage of higher resolution measurements offered by onboard server data can be best highlighted with an example. Every server has an onboard sensor for the intake air temperature. The sensor is needed for thermal management of the electronics. The same sensor can also be used for facility thermal management. Traditionally, the temperatures have been measured with a limited number of external sensors mounted on the perforated doors to the IT equipment racks. With onboard sensors, however, all intake air temperatures are measured exactly where they should be measured – at the actual air intakes. The resolution and accuracy increase, and data center thermal management can be improved.

The server intake air temperatures are at the core of data center thermal management (Herrlin, 2008). Thermal management not only helps improve server reliability and longevity but also infrastructure energy efficiency. Due to this importance, several industry documents specify recommended intake air temperatures for IT gear (Telcordia, 2001, 2012 and ASHRAE, 2021).

Utilizing onboard sensors rather than discrete external sensors can also be very cost effective. External wired or wireless sensors for intake air temperatures and external rack Power Distribution Units (PDUs) for measuring input power come with costs for hardware and installation. In short, using external sensors at the servers could become prohibitively expensive.

The flexibility of external sensors is often poor since the data center environment frequently changes, which requires sensor additions or re-arrangements of existing sensors. With onboard sensors, however, another scan of the IT network can find the servers and automatically add them to the monitored devices. Both small and large data centers can easily be managed cost effectively this way.

Finally, for processor (Central Processing Unit - CPU) utilization there is no real alternative to onboard sensor data. This is also true for other IT related information. A host of data is available from onboard sensors over a number of standard management protocols and interfaces. And, the range of platform management information that is accessible is increasing with every new generation of servers.

Although network data exchange with onboard sensors is a superior technology compared with using external sensors, there are currently a number of implementation barriers. The expertise needed can be lacking and/or poorly understood, especially on the facility side in small data centers. Security concerns may also surface when accessing information from servers. Furthermore, in certain mixed environments such as colocation data centers, access to servers may not be allowed. Therefore, some data centers may choose to continue using external sensors. For external intake air temperature sensor placement recommendations, see “Thermal Guidelines and Temperature Measurements in Data Centers” (LBNL, 2020).

In 2009, Lawrence Berkeley National Laboratory (LBNL) led and executed an integration project (IPACK, 2009 and LBNL, 2009) for controlling cooling equipment based on onboard server sensor data. The server intake air temperatures were monitored and made available through each server’s network connection that supported the Simple Network Management Protocol (SNMP) or the Intelligent Platform Management Interface (IPMI).

The data center chosen was Intel’s Santa Clara, CA site where the operational work load was primarily engineering computing. IT, infrastructure, and research team members from the host site were added to complete the coverage from the platform level instrumentation through the data center cooling and control systems. In addition, IBM and Hewlett Packard engineers collaborated in developing the control strategies.

The project demonstrated that in 2009, the IT integration was not straight forward, though the needle has moved since then. This report presents the main alternatives available today to extract server data from the IT network and also provides a discussion around data center management protocols/interfaces to make it easier to digest the core content.

Each alternative is exemplified by free and/or commercial solutions. However, neither LBNL nor DOE endorses any particular products. The products in this document should be considered as examples of general product categories.

This report provides a gateway to accessing onboard server sensor data rather than being a detailed implementation document. The latter information is widely available from vendors. The ultimate purpose of this report is to help increase the use of onboard sensor data and thereby be in a better position to improve the energy efficiency of both the IT and facility infrastructure.

First, the ENERGY STAR specification for computer servers is summarized. It states that access to CPU utilization, intake air temperatures, and input power must be provided. These parameters are at the core of energy efficient data center operation. This summary is followed by a discussion around basic IT concepts, including the main network protocols and interfaces to make it easier to digest the remaining content. Next, an overview of Data Center Management solutions is provided, including Data Center Networking solutions, Data Center Infrastructure Management solutions, and Hybrid solutions. Finally, three practical paths to access the onboard data are presented. Each path or product category is exemplified by a commercial product.

The following bullets are intended as a concise road-map for this report.

- Chapter 2 summarizes the ENERGY STAR specification for computer servers as well as important server parameters for energy optimization in data centers.
- Chapter 3 reviews IT concepts including the main network protocols/interfaces that play a key role in pulling out the onboard sensor data. It also lists specialized data access tools.
- Chapter 4 provides an overview of versatile Data Center Management solutions that can be used for accessing the server and its onboard sensor data.
- Chapter 5 outlines three examples of versatile commercially available Data Center Management solutions that do not require much IT network expertise.

Besides being a free-standing document, it is also an integral part of the DOE Data Center Energy Practitioner (DCEP) certificate training program (DCEP, 2021). The goal of the DCEP program is to accelerate the energy savings in data centers.

2. ENERGY STAR for Computer Servers

ENERGY STAR is a program administrated by the U.S. Environmental Protection Agency (EPA) and U.S. Department of Energy (DOE) that promotes energy efficiency. The program provides information on the energy consumption of products and devices using standardized methods. ENERGY STAR products are third-party certified to be energy efficient. The ENERGY STAR label is found on more than 75 different product categories, including computer servers and other data center equipment.

The ENERGY STAR label makes it easy to find an energy efficient server to meet specific needs. Using the ENERGY STAR (2021) Product Finder, you can select from hundreds of certified energy efficient servers from different companies. There are also ENERGY STAR rated UPSs, storage products, and networking equipment.

IT vendors offer the ability to view server data in real time using onboard sensors and automation software. Indeed, the ENERGY STAR (2018) for Computer Servers Specification Version 3.0 addresses the importance of server monitoring and specifies that a server sold as an

ENERGY STAR qualified server must expose data on CPU utilization, intake air temperature, and input power. This data must be made available in a published or user-accessible format that is readable by third-party, non-proprietary management software over a standard network. The fact that this information is part of the requirements is a testament of the importance of the data to data center energy efficiency. These parameters are summarized below.

- **CPU Utilization**: Ghost servers (also known as zombie or comatose servers) are powered but are doing little or no productive work. They can waste vast amounts of energy. Unfortunately, they are not uncommon. Data center operators often lack guidance on how to systematically identify and decommission these servers. Accessing the onboard CPU utilization data can help in this endeavor. The DOE Data Center Energy Practitioner (DCEP) training program (DCEP, 2021) has a module looking at ways of improving the energy efficiency of IT equipment operation, including maximizing the CPU utilization.

- **Intake Air Temperatures**: As outlined above, the thermal environment is defined by the temperature of the air drawn into the IT equipment, the temperature the electronics depends on for reliable cooling. The purpose of relevant industry guidance documents is not only to help maintain high IT equipment reliability but also operate the data center energy efficiently. Accurate intake air temperatures are imperative for the success of both objectives. The DCEP program (DCEP, 2021) has modules looking at ways of improving the IT equipment reliability and data center energy efficiency with proper thermal management.

- **Input Power**: Accurate and complete input power readings are important for data center power capacity planning or tracking energy-efficiency improvements with increased CPU utilization or increased intake air temperatures. The latter benefit the energy efficiency of the cooling infrastructure. Traditionally, costly external intelligent rack Power Distribution Units (PDUs) were used to monitor input power to servers. With onboard server data, real-time data can be collected without deploying such infrastructure.

ENERGY STAR qualified servers that include a pre-installed Operating System (OS) must include all necessary drivers and software for end users to access standardized data as specified in the Server Specification. Products that do not include a pre-installed OS must be packaged with printed documentation of how to access registers that contain relevant sensor information. This requirement may be met via either printed materials, electronic documentation provided with the computer server, or information publicly available on the Partner's website where information about the computer server is found.

The ENERGY STAR document was not designed to provide implementation guidance to meet the requirements. Consequently, smaller data centers in particular may need a bit of hand holding to access the data. The document you are reading tries to help in this regard.

3. Network Protocols and Interfaces

Network protocols and interfaces play an important role in pulling out the onboard server sensor data. Readers familiar with basic IT network concepts may want to skip this section, otherwise, it provides information that makes it easier to digest the remaining chapters.

A “network protocol” is the method by which two or more devices understand one another. It defines the rules for exchanging information on the connecting point, like a spoken language. The protocol defines the rules, syntax, semantics, and synchronization of communications and possible error recovery methods.

In modern networking, there are many coexisting network protocols at many levels, each with its own specific purpose. There are data transfer protocols (e.g., the Transmission Control Protocol - TCP), management protocols (e.g., the Standard Network Management Protocol - SNMP), and more. This can be an arduous task. Network Management Systems convey and manage the operations and communications performed on a computer network.

A “network interface”, on the other hand, is the connecting point between two adjacent devices on the network. A physical interface can be a connector and a virtual interface can be an Application Programming Interface (API).

Application Programming Interface (API) (introduced in 2000)

Modern servers are equipped with Baseboard Management Controllers (BMCs) that enable remote (or “out-of-band”) server management via multiple standard built-in Application Programming Interfaces (APIs). The BMC in the managed server can be accessed with a remote console by requesting information and the API will provide a response, for example server intake air temperature. More about the BMCs later in this chapter.

These interfaces are purpose-built software for an express use to allow communication between devices or applications. The primary value of an API is that it allows a company to access information or software capabilities from another source, providing greater value without an additional investment of time, money, and resources. It is used for solutions integration management purposes. If APIs didn’t exist, users would have to do their own nitty-gritty coding. APIs simplify the integration by hiding complexities and extend functionality. Needless to say, APIs play a key role in network management.

An API is a set of defined rules that explains how devices or applications communicate with one another, acting as an intermediary (middleman) layer that processes data transfer between systems. This allows services and products to communicate with each other and leverage each other’s data and functionality through a documented interface. Developers don't need to know how an API is implemented; they simply use the interface to communicate with other devices and services.

Here is how an API specifically works:

1. A client application initiates an API call to get information—also known as a *request*
2. After receiving the request, the API makes a call to the managed device
3. The device sends a *response* to the API with the requested information
4. The API transfers the data to the requesting client.

An API defines all the valid messages that it can accept. It says nothing about the proper ordering of these messages or their interaction with other devices. So, the network protocols sit on top of the APIs. As discussed above, a protocol defines the valid sequence of messages that flow between devices to accomplish some higher-level task.

In addition to the fundamental Internet Protocol (IP) and the Hypertext Transfer Protocol (HTTP), three common network management protocols/interfaces will be reviewed: SNMP, IPMI, and DMTF Redfish. Please note again that multiple protocols run at the same time on the network.

Transmission Control Protocol/Internet Protocol (TCP/IP) (introduced in 1970)

TCP/IP is a suite of communication protocols - a set of rules and procedures - used to interconnect devices on the Internet. TCP/IP is also used as a communications protocol in a private computer network (an intranet or extranet) in the form of a LAN or WAN. In fact, the TCP/IP is a global standard for Internet, LAN, and WAN. IP addresses are assigned to every device (node) in the network and all devices need to be configured with the TCP/IP suite. Table 1 shows the different TCP/IP "layers" and important examples, each providing specific functionality.

TCP/IP uses the client-server model of communication in which a user or machine (a client) is provided a service, like sending a webpage, by another computer (a server) in the network.

TCP/IP Layers	Examples
Application Layer	SNMP, FTP, HTTP
Transport Layer	TCP, UDP
Network (Internet) Layer	IP
Network Access Layer	Ethernet

Table 1. The TCP/IP Layers.

TCP/IP specifies how data is exchanged over the Internet by providing end-to-end communications that identify how it should be broken into packets, addressed, transmitted, routed and received at the destination. TCP/IP requires little central management. The two main protocols in the TCP/IP suite serve specific functions.

TCP (Transport Layer) and IP (Network Layer) are the two main protocols (giving the name to the suite), though other protocols are included in the TCP/IP suite. The TCP is the most popular transport protocol. The UDP is a faster transport protocol but less reliable, and it is used mainly in voice over IP.

The TCP protocol defines how applications can create channels of communication across a network. It also manages how a message is assembled into smaller packets before they are then transmitted over the Internet and reassembled in the right order at the destination address. The IP protocol, on the other hand, defines how to address and route each packet to make sure it reaches the right destination. Each gateway computer on the network checks the IP address to determine where to forward the message.

The TCP/IP model differs slightly from the Open Systems Interconnection (OSI) networking model. The OSI model has seven layers and the TCP/IP model has only four. The two models have four layers in common: Application, Transport, Network, and Link. The conceptual OSI model gives guidelines on how communication needs to be done, while the TCP/IP protocols lay out standards on which the Internet was developed. Thus, for our purposes, TCP/IP is the more practical model.

Hypertext Transfer Protocol (HTTP) (introduced in 1989)

The Internet is a system of interconnected computer networks that use the TCP/IP protocol to link devices worldwide using copper wires, fiber optic cables, or wireless networks. The World Wide Web (WWW) on the other hand refers to the online content that is formatted in the Hyper Text Markup Language (HTML) and accessed via the Hypertext Transfer Protocol (HTTP). In other words, the Internet refers to hardware while the WWW refers to software.

The WWW is an application that runs on the Internet. It refers to a large collection of web pages, which are linked together using hyperlinks. A web browser and the HTTP protocol help to access webpages or websites in the WWW. The HTTP protocol is a set of rules for transferring files such as text, images, audio, video and other multimedia files on the WWW.

The purpose of the HTTP protocol is to provide a standard way for web browsers and servers to talk to each other. For example, when a computer is fetching data, it usually sends an HTTP message called a GET request (see the DMTF Redfish interface later in this chapter).

HTTP is text-based, and it's designed to be readable by humans as well as machines. People creating applications choose HTTP on purpose because it is well understood by many developers. Today, it is the foundation of the WWW.

HTTPS (with an “S”) refers to Hypertext Transfer Protocol *Secure* and, as the name suggests, is a more secure variant of HTTP. Specifically, it opens an encrypted connection so that data can be sent encoded in a way that will not be readable by eavesdroppers.

Simple Network Management Protocol (**SNMP**) (introduced in 1988)

SNMP is an Internet standard protocol for collecting and organizing information about nearly any managed device on IP networks and for modifying that information to change device behavior. The protocol is a vital tool for network management, that is, monitoring and controlling devices on the network. Devices that support SNMP include typical data center equipment such as

servers, storage, switches, and routers. SNMP Agents come preloaded (bundled) on most devices, and they simply need to be enabled and configured. Each SNMP Agent has its own IP address as a unique identifier.

SNMP is part of the TCP/IP model and belongs to the Application Layer, which is the layer closest to the user (see Table 1). It is a software specification which is used across a wide variety of hardware, usually in the Operating System (OS), and it is primarily used “in-band”. In-band means that it only works after the OS has been booted. Compare this with out-of-band management discussed on page 14.

SNMP exposes management data in the form of variables on the managed device organized in a Management Information Base (MIB) which describe the device status and configuration. The Agent assembles the MIB from which the SNMP Manager - often also called SNMP Console or SNMP Station - queries data and, in some circumstances, manipulates the data.

SNMP is a so-called manager-agent model (see Figure 2). In short, an SNMP-managed network consists of four key components:

- Managed Devices (e.g., servers)
- SNMP Agent – software that runs on the Managed Devices
- Managing Device
- SNMP Manager – software that runs on the Managing Device.

One or more administrative servers (managing devices) have the task of monitoring or managing a group of hosts or devices on a computer network (managed devices). Each managed device executes a software component called an SNMP Agent which reports device-specific information via SNMP to the SNMP Manager. The Agent has local knowledge of management information and translates that information to or from an SNMP-specific form. Agents can also send notification to the manager without being polled, for example if an error is detected.

TRAP is one of five primary types of SNMP messages to communication between the SNMP Agents and the SNMP Manager. A TRAP is sent to the Manager when an issue needs to be reported. SNMP TRAP is a widely used mechanism to alert and monitor devices’ activities across a network.

The SNMP Manager is software that runs on the managing device. The Manager executes applications that monitor and control managed devices. It sends requests and receives Agent responses in return. Most Managers poll the network for information regularly. SNMP Managers range from the very simple to highly complex.

The SNMP Agent is software that is packaged within the network device so no extra hardware is required. This means that the Agent will not be available if the managed device is turned off. Enabling the Agent allows it to collect the Management Information Database (MIB) from the device and makes it available to the SNMP Manager. SNMP allows devices to communicate even if the devices have different hardware and run different software.

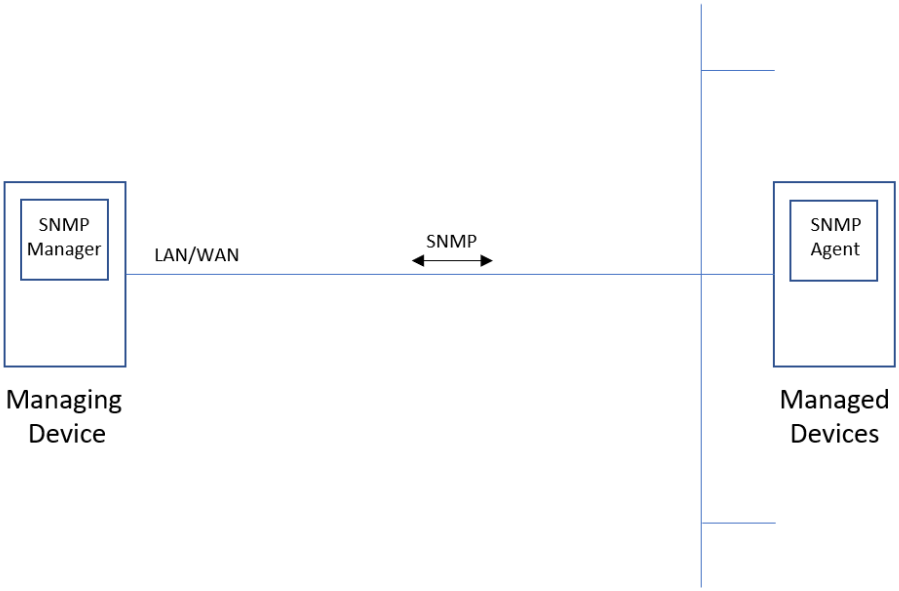

Figure 2. The SNMP model.

SNMP messages (or commands) between the Manager and the Agents are transported via the User Datagram Protocol (UDP) which is part of the TCP/IP protocol and its Transport Layer. There are seven SNMP UDP types: GET request, SET request, GET NEXT request, GET BULK request, RESPONSE, TRAP, and INFORM. Messages go between the Manager and the Agents in a Pull (requested by the Manager) and Push (pushed by the Agent) fashion.

Note that all network management requires an agent, whether the agent is embedded in the management platform, the managed device, or separately installed software. The industry has accepted the de-facto definition of “agentless” as a management agent that is embedded in the software of the device or as a capability of the manager, requiring no separate installation or licensing. Agentless monitoring really means the use of existing, embedded capabilities. SNMP is an example of agentless monitoring.

More specifically, agentless monitoring refers to operations where no service, daemon, or process needs to run in the background on the device. In agentless monitoring, there is no need to install or run new software related to the task itself, and there is no need for intrusive and expensive agents to be installed on the managed devices. This means that managing of the monitoring environment becomes easy. And, some devices don’t allow an agent to be installed, for example, network and storage devices don’t allow anything to be installed.

Agent-based monitoring, on the other hand, typically involves installing an agent (small executable) on, or alongside the system which will be monitored. This agent is often supplied by a vendor with a monitoring solution. An agent-based monitoring system automatically collects data on the performance and availability of hardware resources, operating systems, middleware, and applications in a physical, virtual or cloud environment.

SNMP is simple, yet powerful. It has the ability to help manage the network through the following capabilities:

- Provide Read/Write abilities – for example re-configure IP addresses
- Collect information on how much bandwidth is being used
- Collect error reports into a log, useful for troubleshooting and identifying trends
- Email an alert when your server is low on disk space
- Monitor your servers' CPU and Memory use, alert when thresholds are exceeded
- Text or send an Email message when a device fails
- Perform active polling, i.e., the Manager asks devices for status every few minutes
- Managed devices can send alerts to the Manager on error conditions.

Many user interfaces are available for the SNMP Manager. There are OS Command Line Interfaces (SNMP CLI) for humans sending and receiving information over the network using a specified syntax. Think Windows DOS-type of text based interface. There are both free and commercial tools. For example, SnmpWalk is a simple free tool for querying managed devices. There are also GUI-type of interfaces, both free and commercial.

Intelligent Platform Management Interface (IPMI) (introduced in 1998)

IPMI is one of the most used acronyms in server management. IPMI is a vendor-neutral standard specification defining a set of common interfaces for monitoring server physical health and more. IPMI became popular due to its acceptance as a standard monitoring interface by hardware vendors and developers. The specification is supported by more than 200 computer system vendors, such as Dell, HPE, Cisco, and Intel. Yet, it is not as universal as the SNMP protocol.

IPMI is an agentless hardware-based solution for securing, controlling, and managing mainly servers. As such, IPMI operates independently of the operating system (OS) of the managed device and allows administrators to manage a device remotely and offers enhanced features when used with system management software. Since an "interface" is the connecting point between two adjacent devices, the IPMI is called an interface since it interfaces with the managed server hardware with standardized specifications.

IPMI allows out-of-band management of computer systems. Again, out-of-band management is a way to manage a computer that may be powered off or otherwise unresponsive by using a network connection to the hardware rather than to an OS or login shell. The option of monitoring and managing systems independently of a server's OS is one significant benefit other monitoring tools often lack. With IPMI it is possible to make adjustments to settings such as BIOS without having to log in or seek permission from the OS. Another use case may be installing an OS remotely. Furthermore, IPMI does not rely on any proprietary hardware, which eliminates compatibility issues.

The managed device may be powered off, but must be connected to a power source and to the monitoring medium, typically a Local Area Network (LAN) connection. IPMI can also function

after the OS has started to expose management data and structures to the system management software.

The IPMI model is shown in Figure 3. An IPMI-managed network consists of four key components:

- Managed Devices (e.g., servers)
- BMC – specialized hardware on the Managed Devices running IPMI software
- Managing Device
- IPMI Utility – software that runs on the Managing Device to manage IPMI devices.

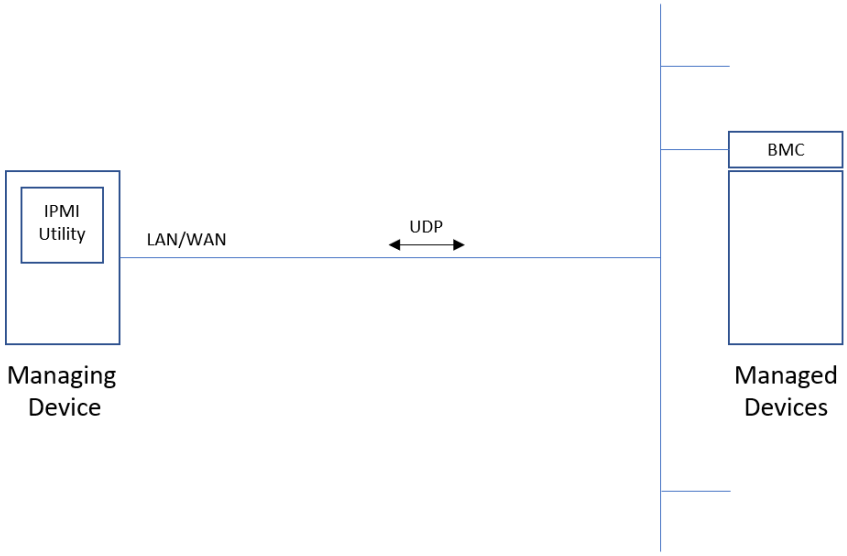

Figure 3. The IPMI model.

IPMI runs on separate, specialized hardware attached to the motherboard of a managed device. This hardware is the Baseboard Management Controller (BMC) – the heart of the system with the main processor where IPMI is implemented through an API. The BMC is key to IPMI's out-of-band capability. The BMC acts like an intelligent middleman, which operates on stand-by power. A BMC is included on just about every server motherboard. It can be accessed remotely either via a dedicated or a shared network connection (see Figure 4).

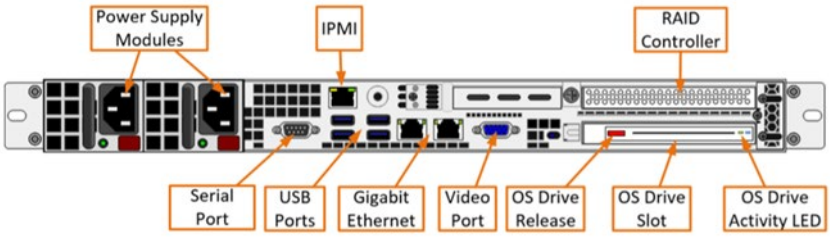

Figure 4. BMC (here labeled IPMI) dedicated management port.

IPMI monitors, via the BMC, vital data that define the working status of a server's hardware. For example, it monitors power, fan speed, server health (e.g., temperature), CPU utilization, security details, and the state of the operating system.

The BMC manages the interface between the managed device and the managing device. The IPMI Utility is installed on the managing device. The BMC receives reports from sensors within a server and acts on these reports. The BMC also have an SNMP Agent to make it backward compatible with SNMP. The Agent has to be enabled and configured to work.

IPMI supports remote monitoring over the Internet, LAN, or WAN. It is part of the application layer of the TCP/IP model but it does not belong to a single Open System Interconnection (OSI) layer, but rather several. The BMC uses IP to communicate with the IPMI Utility (also called IPMI Remote Console). IPMI uses the User Datagram Protocol (UDP), which prepares and forwards data packets over the network. Since IPMI is on top of UDP, requests and responses can be transferred inside a UDP datagram.

Besides the BMC, every vendor has their unique out-of-band systems management adapter. For example, Dell EMC has its Integrated Dell Remote Access Controller (iDRAC), HPE has its Integrated Lights-Out Adapter (iLO), and IBM has its Remote Supervisor Adapter (RSA). These controllers/adapters provide extended functionality.

Disadvantages of using IPMI center on security and usability. Cybersecurity challenges led to the addition of encryption and firmware firewall features in later versions of IPMI. There are also documented configuration and updating challenges.

These disadvantages led the IPMI promoters to issue a statement in 2019 (Intel, 2019). It stated that no further updates to the IPMI specification should be expected. They encouraged the adoption of a more modern interface which can provide better security, scalability, and features for existing data centers and be supported on the requisite platforms and devices. DMTF Redfish is given as an example of one such interface (see next section). The statement applies only to the IPMI specification, and should have no impact on existing IPMI implementations.

Again, the IPMI Utility is software installed on the managing device for managing IPMI servers. A wide variety of remote access services are supported, including Command Line Interfaces (CLIs) and Graphical User Interfaces (GUIs). There are many available IPMI utility tools, both free and commercial. IPMITool, for example, is a free command line utility to access IPMI functionality. An IPMITool command line for showing temperature senor data is *ipmitool sdr type temperature* where sdr stands for Sensor Data Repository. There are also GUI-based tools such as IPMIView from Supermicro.

DMTF Redfish (API) (introduced in 2015)

The Intelligent Platform Management Interface (IPMI) protocol has held a key position in server management in data centers for many years, in large part, due to its ability to use common commands to control servers from multiple manufacturers located in multiple physical locations. Nevertheless, IT administrators are up against the limits of efficient automation using IPMI. IT

automation is the use of software to create repeatable instructions and processes to replace or reduce human interaction with IT systems**.**

DMTF Redfish (or simply Redfish), the de-facto successor to IPMI, is a relatively new (2015) global, open server hardware management standard where interoperability is key. Similar to IPMI, Redfish is an interface serving as a connecting point between two adjacent devices on the network. It is nationally and internationally recognized by ANSI and ISO, respectively. Redfish provides functionality well beyond IPMI, including secure and scalable server management automation. Most servers come with both IPMI and Redfish. Redfish was initially focused on servers but has now extensions to storage (Swordfish) and networking (Yang) devices.

The agentless management interface Redfish is developed by the largest global server manufacturers who have unified under the umbrella of the Distributed Management Task Force (DMTF, 2021). DMTF is a 501(c)(6) standards organization. Its board of directors is led by industry-leading companies, including Broadcom, Cisco, Dell, Hewlett Packard Enterprise (HPE), Hitachi, Intel, Lenovo, and NetApp. DMTF standards enable a more integrated and cost-effective approach to management through interoperable solutions.

The Redfish model shown in Figure 5 looks similar to the IPMI model since Redfish also uses the Baseboard Management Controllers (BMCs) infrastructure. That means that Redfish operates independently of the operating system (OS) of the managed device. Therefore, much of the BMC technology and discussion under IPMI in the previous section applies to Redfish as well. Specifically, Redfish servers are managed by using Redfish APIs, which reside on the BMC in the managed device. The Redfish Utility (also called Redfish Remote Console) is used to monitor the managed device through the BMC. Redfish support on servers include the following companies and their BMCs:

- Dell iDRAC BMC
- HPE iLO BMC
- HPE Moonshot BMC
- Lenovo XClarity Controller (XCC) BMC
- Supermicro X10 BMC
- IBM Power Systems BMC
- Cisco Integrated Management Controller.

The Redfish API is an interface for data center management that leverages the well-known, web hypertext transfer protocol HTTP/HTTPS as a request-response mechanism to enable communications between clients and servers. The API is called by issuing an HTTP request: CONNECT (starts two-way communications), DELETE (deletes resource), GET (requests representation of resource), HEAD (requests header information of resource), OPTIONS (requests permitted communication options), PATCH (applies partial modifications to resource), POST (sends data to server), PUT (creates new resource or replaces resource), and TRACE (performs a message test to the resource). The data is contained in JSON (Java Script Object Notation). The HTTP/HTTPS protocol is part of the Application Layer of the TCP/IP protocol.

The Redfish API ecosystem is an open industry-standard specification and schema designed to help customers integrate solutions within their existing solutions. Thus, Redfish servers come with multiple other standard APIs loaded on the BMC, such as SNMP and IPMI. This results in nearly universal server connectivity.

Again, the Redfish Utility is software installed on the managing device for managing Redfish servers. A wide variety of remote access services are supported, including Command Line Interfaces (CLIs) and Graphical User Interfaces (GUIs). There are many available Redfish utility tools, both free and commercial. The DMTF Redfish Utility, for example, is a command line interface to access Redfish API functionality. The DMTF Redfishtool is another command line interface.

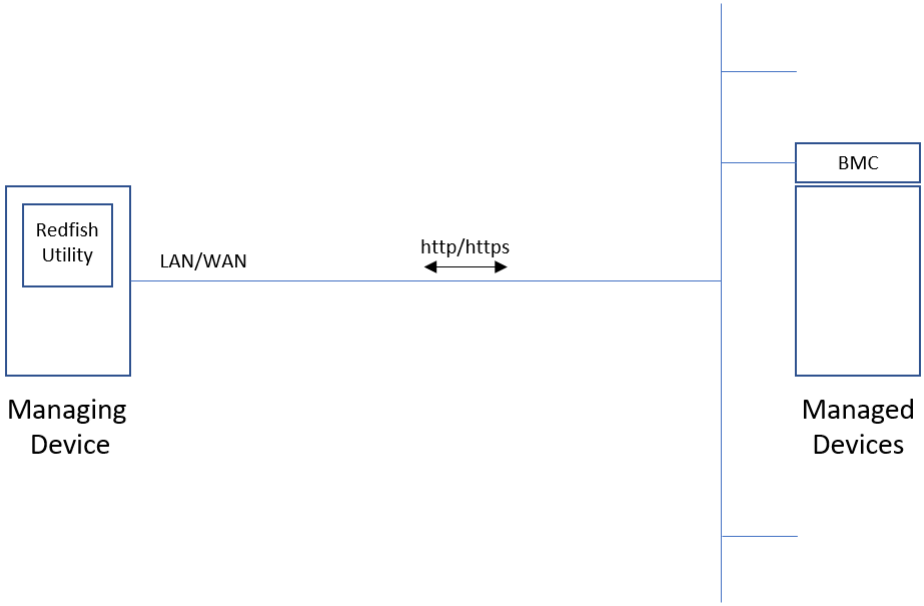

Figure 5. The Redfish model.

4. Data Center Management Solutions

Data Center Management solutions can be divided into Data Center Networking tools and Data Center Infrastructure Management (DCIM) tools. The former focuses on the IT equipment and IT connectivity whereas the latter focuses on the support infrastructure, such as cooling and electrical systems.

These tools start to integrate to meet the need for a more holistic view of the data center and its systems and components. Most servers provide a host of baseboard information to the IT network. For example, the intake air temperature sensor data from the servers can be funneled to the facility (support infrastructure) management system to control the cooling system.

In the previous chapter, network protocols and interfaces that are used to fetch data from servers or other IT devices were reviewed. Today's servers have BMCs with APIs for IPMI, Redfish, and more. In addition, they have SNMP Agents. In other words, they have nearly universal connectivity. Any of the SNMP Managers and IPMI or Redfish Utilities listed in the previous chapter can be used to access the functionality of SNMP Agents and IPMI or Redfish APIs, respectively. However, each one of these tools is *specific* to a particular network protocol or interface.

In contrast, Data Center Management tools provide more *versatile* solutions with two distinctly different emphases. The first type of tools (Data Center Networking tools) focus on data found in IT devices and their virtual components. The second type (DCIM) deals with data found in data centers' thermal, mechanical, and electrical infrastructure support systems.

Data Center Networking tools facilitate communication and information exchange between their connected devices and internal/external networks. Companies use these tools to create stable connections between data centers and their connected devices. They let you talk to the motherboards to gather motherboard-level data. These tools often integrate with DCIM solutions (see below) to better manage hardware components and optimize resource usage.

Devices on a network commonly communicate via traditional networking protocols such as SNMP or through APIs. Servers also provide data on physical parameters such as input power, power supply status, operational status, internal fan speeds, intake air temperatures, and CPU usage from different places within the device.

An example of a Data Center Networking tool is OpenManage Network Manager from Dell (2021). This solution will be discussed in some detail in the next chapter. Other companies with similar products include HPE and Extreme Networks.

Data Center Infrastructure Management (DCIM) tools are the counterpart of Data Center Networking tools at the support infrastructure level. They are primarily used to manage, organize, and monitor data center power, cooling, and physical space. Companies utilize these tools to optimize performance of their data center support infrastructure.

They often integrate with other data center systems from a wide range of vendors to create an increasingly dynamic solution. The DCIM solution becomes the aggregation point and correlation engine that interprets (raw) data.

The Building Management Systems (BMS), for example, hold a wealth of information and when integrated into a DCIM suite will enable easy monitoring of cooling resources as well as power and environmental factors. These systems communicate with any one of a number of different protocols, including SNMP, MODbus, and BACnet. Many DCIM solutions can also tap into onboard server data from Data Center Networking tools.

An example of a DCIM tool is Nlyte from Nlyte (2021). This solution will be discussed in some detail in the next chapter. Other companies with similar products include Schneider Electric and Sunbird.

Hybrid tools are positioned somewhere in between the two main tool categories discussed above and often have a more targeted scope. These software solutions often collect and analyze the real-time health, power, and thermals of a variety of devices in data centers helping you improve the overall efficiency and uptime.

An example of a hybrid tool is the Intel DCM (Data Center Manager) from Intel (2021). This solution will also be discussed in some detail in the next chapter.

5. DCIM and Data Center Networking Tools

We will now explore a few of the commercial tools mentioned in the previous chapter that could be used to collect onboard sensor data: Dell OpenManage Network Manager, Nlyte DCIM, and Intel DCM. Selecting a commercial tool for accessing motherboard data makes the collection process simpler, without requiring much IT network expertise. What the three tools have in common is that they allow multi-vendor, multi-protocol, and scanning and detection of devices. They also use agentless monitoring. These particular features make it easier for end users to monitor and manage their data centers.

Dell OpenManage NM (Data Center Networking solution)

The IT department of most data centers has Data Center Networking tools that can extract motherboard-level data one needs for overall energy efficiency efforts, including ENERGY STAR required data, namely CPU utilization, intake air temperature, and input power.

Dell OpenManage is a set of systems management applications built using industry-standard protocols and specifications. It is not a product within itself, but rather a suite of products. This suite provides systems management solutions designed to simplify, automate, and optimize the IT operations.

One of these products is OpenManage Network Manager or OpenManage NM (OMNM). This agentless offering discovers (finds and registers), configures, monitors, and manages multi-vendor devices from vendors such as Dell, Aruba, Cisco, Brocade, HP, and Juniper. It provides a unified management system and automates common network management operations.

The OMNM provides a central area for network management for the Dell networking portfolio as well as multi-vendor infrastructures. Thousands of specific networked devices can be managed, as long as it has an IP address. The OMNM provides a number of protocol/interface types to discover and manage devices. They include Redfish, HTTP, HTTPS, IPMI, SNMP, and more.

A server can be considered to consist of three layers: Apps/Services/VMs, OS/Hypervisor, and the physical hardware. The OMNM has used OS (e.g., Windows or Linux) level access to discover, manage, and monitor CPUs, discs, memory, and more for a long time. A relatively new feature is physical hardware level access for discovery, management (reboot, turn ON/OFF), and monitoring from the BMC. This feature provides management access to a multitude of server performance data.

Again, a dedicated management port on the servers provides access to the onboard data via a number of protocols/interfaces. The related server BMC is discovered by the OMNM discovery feature selecting the Redfish interface. Most modern servers support this relatively new interface. See Chapter 3 for the benefits of using the BMC rather than the OS-level access.

OMNM provides a detailed picture of the network, including health and performance (including alarms for out-of-bound data) as well as configuration and change management. It displays potential problems in the network, such as under-performing devices. There is a single source to

manage configuration files (backups, updates, deployment, restoration, schedules, etc.) through scripts.

OMNM helps network administrators manage and monitor the network from a single web GUI interface console (see Figure 6) using Internet Explorer (IE) or Chrome. The tool comes with a number of default monitors/reports and custom monitors/reports can be quickly be created by using a number of report templates. Customizable physical and logical topology maps can be viewed with detailed information about the devices and their connections. Network traffic in real time or over time can also be viewed for different protocols, applications, and devices.

OMNM can be installed on a single server on Windows, Linux, or as a virtual appliance. The licenses are based on the number of devices under management. Comes in one-, three-, or five-year licenses. All core features are included in the license fee.

Figure 6: Dell OMNM Web GUI interface.

Nlyte DCIM (full-featured DCIM solution)

If you have a full-featured DCIM tool, you likely already have the capability to access the data you need for your data center energy efficiency efforts. If you do not have a DCIM tool in place that provides this type of data, it might be better to consider one of the other types of tools discussed in this chapter. Investing in a full-featured DCIM tool may be overkill if only limited data is desired. Cost is significant for comprehensive DCIM solutions but should be weighed against benefits such as limiting labor intensive activities and improving resource management. Need for customizations, after-sales support, and professional services can also be costly for any sophisticated software.

Once the data center reaches a certain size, it is increasingly difficult to manage the data center support infrastructure manually. Data centers provide vast amounts of data about the health and operational status of the facility. With hundreds or thousands of devices communicating telemetry data, all this information quickly becomes unmanageable with ad-hoc solutions like

Excel, Visio, or CAD drawings. As a result, data center professionals often turn to DCIM tools to more effectively manage operations.

The IT and facility functions have traditionally had two different teams and two different fields of expertise. This gap needs to be bridged to make the overall data center operation optimized. Therefore, a full-featured DCIM solution may morph into something that resembles an Integrated Data Center Management (IDCM) solution, which brings together different disciplines to deliver better outcomes. An IDCM solution provides integration between critical facility infrastructure, the devices in the data center (servers, storage, network, and many more), and the application workloads running on those resources. IDCM connects the capabilities and features of Building Management Systems (BMS), DCIM, and IT operations. A BMS monitors and controls a building's facility systems, including electrical, lighting, and mechanical systems.

Nlyte is one of many full-featured DCIM solutions, and the following is a closer look at this particular offering. Data centers throughout the Federal government must comply with the optimization mandates set out in the Data Center Optimization Initiative (DCOI), which includes a requirement to use DCIM solutions. Nlyte has an enterprise agreement with the Department of Energy (DOE). Under this agreement, Federal data centers have access to no-cost licenses and annual technical support paid by DOE. However, implementation services are paid by the data centers.

The three main Nlyte solutions are Asset Optimizer (NAO), Energy Optimizer (NEO), and Asset Explorer (AE). These solutions provide a rich set of functionality right out of the box. The functionality includes the physical layer for servers, storage, network, power, space, and cooling and their utilization, including server CPU utilization.

In addition, Nlyte offers pre-built "connectors" to integrate with third-party tools to enhance the DCIM capabilities in the areas a user may need the most (see Figure 7). It also provides deep integration with third-party tools through RESTful APIs. Nlyte also offers their own additional "modules" that enhance the basic capabilities. One module of special interest is the Nlyte System Utilization Monitor, which identifies under-utilized or ghost servers. In short, the connectors and the additional modules allow a broad integration of IT and facility data.

Note the Intel DCM connector in the third row/column in Figure 7. The Intel DCM will be discussed in the next section.

Nlyte provide agentless asset discovery for finding and registering IT devices on the network. No manual tracking is generally necessary no matter server brand/model (hardware agnostic). It has real-time power and environmental monitoring, which supports all major protocols (although SNMP is the primary protocol) and major devices. A unique feature is DCOI relevant dashboards, including DCOI PUE, DCOI Server Utilization, and DCOI Virtualization. This feature helps show compliance with the Federal DCOI requirements.

Nlyte has a web-based Graphical User Interface (GUI) with both pre-defined and custom dashboards/reports to communicate performance data, including heat and power maps at the room, rack, and device level; view physical layout of the space; look inside the equipment racks;

and view the power and network connections. It has the ability to run on different computer operating platforms and web browsers.

Lastly, a number of data centers can cost-effectively be tied together with one instance of the DCIM tool with remote device control. Their SaaS (Software as a Service) or perpetual pricing can be based on the number of racks, assets monitored, or total power of the facility.

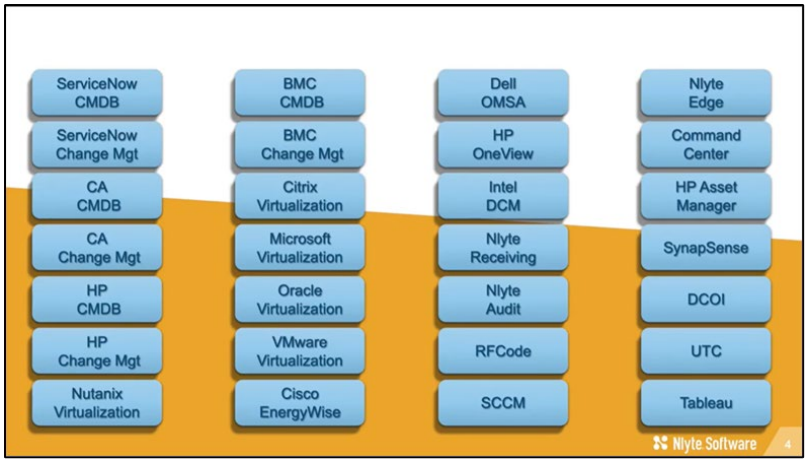

Figure 7: Nlyte Connectors and Modules (partial list).

Intel DCM (Hybrid solution)

The Intel Data Center Manager (DCM) provides a platform for IT and facility professionals alike to assess energy consumption, one of the fastest growing costs in data centers.

If you do not have access to a comprehensive Data Center Networking or DCIM solution, you could consider investing in a nimbler and less expensive solution. DCM is specifically designed to help lower operational costs and optimize infrastructure. It bridges the gap between IT equipment and facility infrastructure. Intel categorizes this as a hybrid of a Server Management Tool and a DCIM tool. It is a targeted solution for dedicated data centers or colocations.

DCM allows real-time device power monitoring, thermal monitoring (intake air temperature and other device-level thermal data) to detect cooling problems and save money on cooling and increase data center uptime, and CPU and power utilization monitoring to identify ghost servers and under-utilized servers for decommissioning or virtualization. It also allows health monitoring of the devices as well as device management such as power control of servers, including policy-based power capping to avoid catastrophic power loss or to save energy at low utilization.

It can monitor not only servers but also Network and Storage equipment as well as power chain equipment such as Power Distribution Units (PDUs) and Uninterruptable Power Supplies (UPSs) giving it its hybrid status.

DCM takes advantage of the onboard sensors in devices ("platform telemetry"). In other words, it talks directly with the devices' motherboards and then use analytics to extract actionable information. DCM replaces manual, hardware-defined application provisioning and management with an automated, software-defined resource model.

This is a vendor agnostic solution that utilizes multiple protocols/interfaces such as Redfish, HTTPS, IPMI, and SNMP without using agents (agentless) to access the corresponding exposed SNMP Agents or APIs on the managed devices. The tool essentially functions as a SNMP Manager or an API Utility.

It also has the capability to scan networks and add devices. Although an automatic discovery and registration of devices for different protocols/interfaces is not available it can be accomplished programmatically. This feature discovers supported devices given their IP-address range and credentials.

The product comes in three flavors:

- Bundled with DCIM offerings (Siemens, Schneider Electric, etc.)
- Integration into OEMs' (Novel, Dell, Supermicro, etc.) monitoring capabilities
- Stand-alone "Console".

The DCM Console is a stand-alone, web-based GUI dashboard application that allows you to access, monitor, and manage CPUs, thermals, power, and health of devices in the data center (see Figure 8). In other words, it provides features mostly focusing on data center energy and reliability management. An API package (without a GUI) is provided for deep integration through RESTful API for DCIM solutions and Original Equipment Manufacturer (OEM) management products.

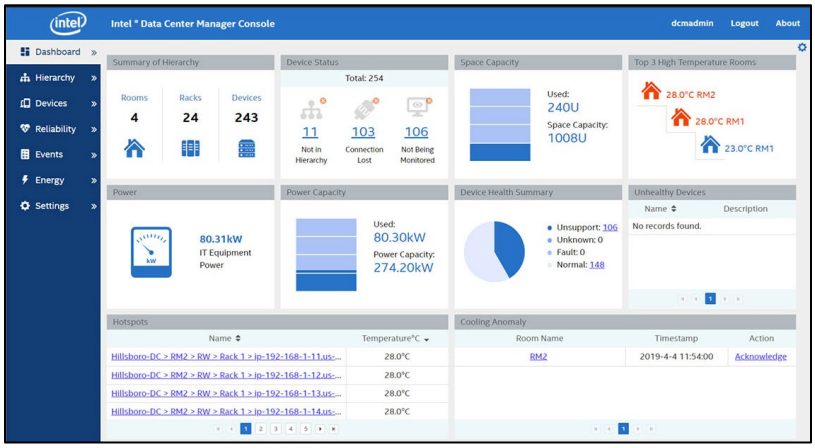

Figure 8: Intel Data Center Manager Console.

The tool allows analytics and trending of thermal and power data (to understand the data and provide optimization suggestions), including gradients to predict potential emergency situation (predictive thermal problem detection). It also allows mapping of room layout and equipment racks with overlays of intake temperatures, power, etc.

The DCM Console has been validated on most major Operating Systems (OS) and web browsers. DCM is licensed by the device (node), which makes it affordable for small data centers as well.

Other Solutions

If you consider using an in-house developed solution to access energy-related data, the ENERGY STAR Power and Performance Data Sheet has some useful information. The overall ENERGY STAR program is covered in Chapter 2, whereas the following section takes a look at the Data Sheet, which provides information on compatible network protocols/interfaces for data collection.

The ENERGY STAR Active Certifications (2021) database provides an up-to-date listing of all certified servers. As was discussed earlier, qualified servers must provide direct access to CPU utilization, intake air temperature, and input power. Furthermore, the data must be in a user accessible format that is readable by third-party, non-proprietary software over a standard network.

The ENERGY STAR Power and Performance Data Sheet contains a wealth of information on system characteristics, system configurations, power data, power and performance for benchmarks, power savings features, power and temperature measurement and reporting, and thermal information. These sheets can be found on the vendors' websites generally under Server Specifications. The next two pages (Figure 9) shows an example of an ENERGY STAR Power and Performance Data Sheet for a Dell server. The Power and Temperature Measurement and Reporting section is most interesting for our current purposes. It covers availability and accuracy of CPU utilization, temperature, and power.

A detail of this particular section is shown in Table 2. In this case, the data collection can be done through the IPMI interface, which is one of the main network interfaces used for accessing onboard server sensors (see Chapter 3 for more on this interface).

Input Power Available & Accuracy?	Yes, +/- 5% for 20%-100% of max PSU load
Input Air Temp Available & Accuracy?	Yes, +/- 2°C
Processor Utilization Available?	Yes
Other Data Measurements Available & Accuracy?	
Compatible Protocols for Data Collection	IPMI
Averaging method and time period	Power: 1 min running average of 2s interval samples. Temperature: no

Table 2: Example of ENERGY STAR Power and Temperature Measurement and Reporting section.

ENERGY STAR® Power and Performance Data Sheet

Dell PowerEdge R815 Featuring the Dell 1100W PSU and AMD Opteron 12C 115W Procs

System Characteristics

Form Factor	2U
Available Processor Sockets	4
Available DIMM Slots / Max Memory Capacity	32 /256
ECC and/or Fully Buffered DIMMs	ECC Registered DIMMS
Available Expansion Slots	6
Minimum and Maximum # of Hard Drives	1, 6
Redundant Power Supply Capable?	Yes
Power Supply Make and Model	Dell 1100W
Power Supply Output Rating* (watts)	1100W
Minimum and Maximum # of Power Supplies	1, 2
Input Power Range (AC or DC)	100 - 240 VAC 50-60Hz
Power Supply Efficiency at Specified Loadings*	81.33%@10%,89.25%@20%,92.30%@50%,90.80%@100%
Power Supply Power Factor at Specified Loadings*	0.79@10%, 0.89@20%, 0.95@50%, 0.98@100%
Operating Systems Supported	Microsoft Windows® Server 2008 R2 Microsoft Windows® Server 2008 SP2 Microsoft Windows Essential Business Server 2008 Microsoft Windows Small Business Server 2008 Red Hat Enterprise Linux 4 and 5 Citrix XenServer 5.x[3] Vmware ESXi 3.5[3] SUSE Linux Enterprise Server 10 and 11
Installed Operating System for Testing	Microsoft Windows® Server 2008 SP2

* Note: Power supply information is for a single power supply only

System Configurations	Minimum	Typical	Maximum
Configuration ID			
Processor Information	4x AMD 6168	4x AMD 6172	4x AMD 6174
Memory Information	16 RDIMMS, 4GB, 1333 MHz	16 RDIMMS, 4GB, 1333 MHz	32 RDIMMS, 4GB, 1333 MHz
Internal Storage	1x 146GB 10k SAS HDD	2x 146GB 10k SAS HDD	6x 146GB 10k SAS HDD
I/O Devices	4x 1Gb LOMs 1x PERC H700 0x quad port 1Gb NIC 0x PERC H800 0x single port FC HBA	4x 1Gb LOMs 1x PERC H700 1x quad port 1Gb NIC 1x PERC H800 1x single port FC HBA	4x 1Gb LOMs 1x PERC H700 2x quad port 1Gb NIC 2x PERC H800 2x single port FC HBA
Power Supply Number and Redundancy Configuration	1 non-redundant	2, 1+1 redundant	2, 1+1 redundant
Management Controller or Service Processor Installed?	Yes	Yes	Yes
Other Hardware Features / Accessories	n/a	n/a	n/a

Power Data	Minimum	Typical	Maximum
Idle Category (1S and 2S only)		N/A (3S or 4S)	
ENERGY STAR Idle Power Allowance (1S and 2S only)	N/A (3S or 4S)	N/A (3S or 4S)	N/A (3S or 4S)
Measured Idle Power (watts)	320.2	375.6	462.0
Power at Full Load* (watts)	550.4	620.3	725.9
Benchmark / Method Used for Full Load Test		Sandra Drystone 2010 SP1	
Test Voltage and Frequency for Idle and Full Load Test		115 V/60 Hz	
Range of Total Estimated Energy Usage ** (kWh/year)	5,610 to 9,642	6,580 to 10,868	8,094 to 12,717
Link to Detailed Power Calculator (if available)		WWW.Dell.com/CALC	

* Note: Full load power represents the sustained, average power at 100% load of the given workload, and does not necessarily represent the absolute peak power or the highest average, sustained power possible for other workloads.

** Note: Estimated kWh/year gives the absolute range of energy use a user could expect from continuous operation (24x7x365) and ranges from 100% idle usage to 100% full load operation. The calculation also includes typical data center overhead at a ratio of 1 watt of overhead to every 1 watt of IT load (corresponding to a PUE of 2.0). Closer approximations may be found by using established power calculators and specific information about the intended operating environment (e.g., average time at idle, data center PUE, etc.).

Power and Performance for Benchmark #1		Minimum	Typical	Maximum
Benchmark #1	Benchmark Used and Type of Workload		Sandra Drystone 2010 SP1	
	Avg. Power Measured During Benchmark Run	550.4	620.3	725.9
	Benchmark Performance Score	322	354	373
	Power Performance Ratio (perf score/avg. power)	0.59	0.57	0.51
	Link to Full Benchmark Report (Where Available)			

Power and Performance for Benchmark #2 (optional)		Minimum	Typical	Maximum
Benchmark #2	Benchmark Used and Type of Workload			
	Avg. Power Measured During Benchmark Run			
	Benchmark Performance Score			
	Power Performance Ratio (perf score/avg. power)			
	Link to Full Benchmark Report (Where Available)			

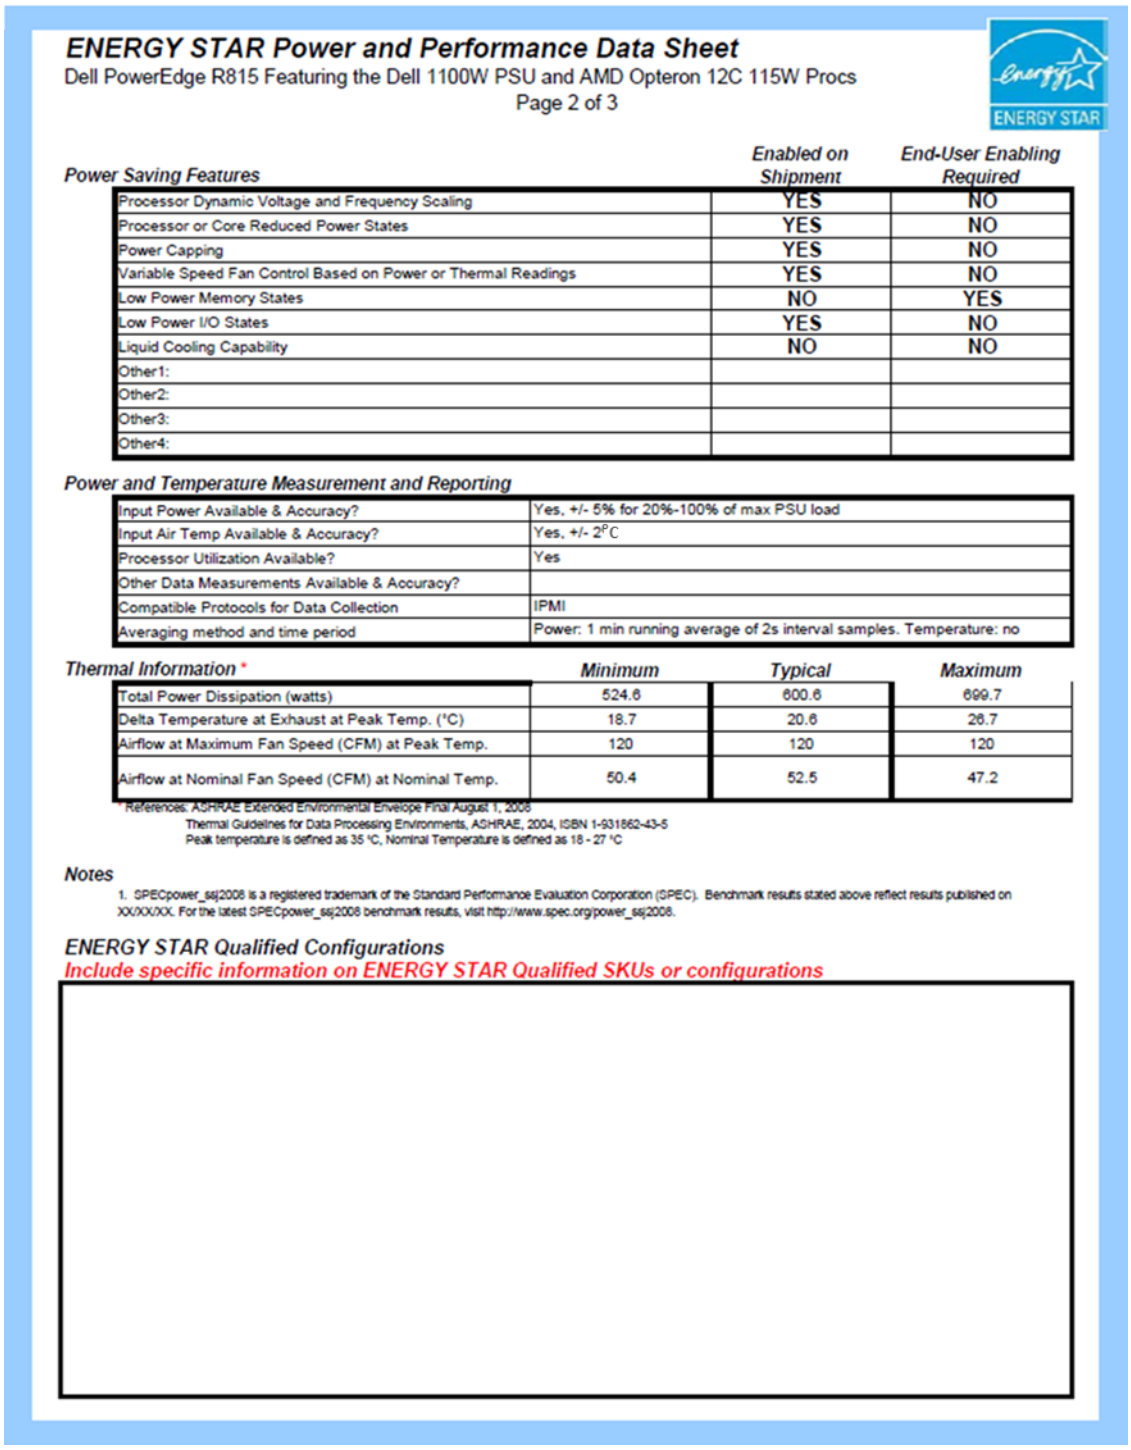

ENERGY STAR Power and Performance Data Sheet

Dell PowerEdge R815 Featuring the Dell 1100W PSU and AMD Opteron 12C 115W Procs

Page 2 of 3

Power Saving Features	Enabled on Shipment	End-User Enabling Required
Processor Dynamic Voltage and Frequency Scaling	YES	NO
Processor or Core Reduced Power States	YES	NO
Power Capping	YES	NO
Variable Speed Fan Control Based on Power or Thermal Readings	YES	NO
Low Power Memory States	NO	YES
Low Power I/O States	YES	NO
Liquid Cooling Capability	NO	NO
Other1:		
Other2:		
Other3:		
Other4:		

Power and Temperature Measurement and Reporting

Input Power Available & Accuracy?	Yes, +/- 5% for 20%-100% of max PSU load
Input Air Temp Available & Accuracy?	Yes, +/- 2°C
Processor Utilization Available?	Yes
Other Data Measurements Available & Accuracy?	
Compatible Protocols for Data Collection	IPMI
Averaging method and time period	Power: 1 min running average of 2s interval samples. Temperature: no

Thermal Information *	Minimum	Typical	Maximum
Total Power Dissipation (watts)	524.6	600.6	699.7
Delta Temperature at Exhaust at Peak Temp. (°C)	18.7	20.6	26.7
Airflow at Maximum Fan Speed (CFM) at Peak Temp.	120	120	120
Airflow at Nominal Fan Speed (CFM) at Nominal Temp.	50.4	52.5	47.2

References: ASHRAE Extended Environmental Envelope Final August 1, 2008
Thermal Guidelnes for Data Processing Environments, ASHRAE, 2004, ISBN 1-931862-43-5
Peak temperature is defined as 35 °C, Nominal Temperature is defined as 18 - 27 °C

Notes

1. SPECpower_ssj2008 is a registered trademark of the Standard Performance Evaluation Corporation (SPEC). Benchmark results stated above reflect results published on XX/XX/XX. For the latest SPECpower_ssj2008 benchmark results, visit http://www.spec.org/power_ssj2008.

ENERGY STAR Qualified Configurations

Include specific information on ENERGY STAR Qualified SKUs or configurations

Figure 9: Example of an ENERGY STAR Power and Performance Data Sheet.

Conclusions

Various data center equipment needs monitoring to operate cost effectively. For computer servers, external sensors have traditionally been used for physical parameters such as intake air temperature and input power although internal (onboard) sensors can also be used to gather information for most of the same parameters. Tapping the data from these onboard sensors has a number of benefits for data center management, including lower overall costs and higher data resolution.

Unfortunately, accessing onboard server sensor data is not always well known nor well understood. This report sought to make this technology more accessible to data center operators and facilitate the use of onboard sensor data, with the ultimate goal of reducing IT and facility energy use.

The ENERGY STAR specification for computer servers requires that access to onboard server CPU utilization, intake air temperatures, and input power data must be provided. These data are imperative to overall energy efficiency in data centers. This specification is highly recommended for data center operators since it contains a wealth of relevant information.

Leveraging onboard server data requires a basic knowledge of IT concepts - including the main network protocols and interfaces. The greater an individual's understanding of IT concepts, the easier it will be to select the most appropriate Data Center Management Solution to access the server and its onboard sensor data. To make the available solutions a bit more approachable, this report reviewed three commercially available software packages.

Modern computer servers have almost universal connectivity through multiple standard network protocols and interfaces, such as SNMP, IPMI, and Redfish. Furthermore, modern Data Center Management solutions have multi-vendor (vendor agnostic) and multi-protocol capabilities as well as semi- or fully-automatic scanning and detection of networked devices that greatly simplifies the tracking of the server inventory. Different management solutions not only start to integrate but also provide connectors to other types of data center systems to meet the need for a more holistic view of the data center. Besides, most such solutions use agentless monitoring, that is, no intrusive agent needs to be installed on the server. All in all, these capabilities make it easier and safer than ever to monitor and manage data centers.

Once the data center reaches a certain size, it is increasingly difficult to manage the infrastructure. Data centers provide vast amounts of data about the health and operational status of the facility. With hundreds - maybe thousands - of devices communicating telemetry data, all this information quickly becomes unmanageable with ad-hoc solutions. Data center professionals often turn to Data Center Management software.

The current and projected future needs should be carefully evaluated before investing in a particular solution, as the cost can be significant. If the interest is purely on energy efficiency, investing in a full-featured management tool may not be the best path forward.

By selecting a commercial Data Center Management product from one of the three product categories discussed in this report (Data Center Networking, Data Center Infrastructure

Management, and Hybrid solutions), there is no need to become an IT expert to access the onboard server data, as these solutions do most of the work for you.

If network protocol/interface versatility is not a priority, then one of the dedicated SNMP Managers, IPMI Utilities, or Redfish Utilities listed in this report can be used for accessing onboard server sensors via a particular protocol/interface. A bare-bones approach could be to use the data access drivers and software that ENERGY STAR servers come with.

If you choose to go it all alone, the ENERGY STAR Power and Performance Data Sheets for servers - also outlined in this report - provide information on compatible standard network protocols and interfaces for accessing the onboard server sensors.

Although accessing onboard server sensor data is generally the preferred technology compared with using external sensors, there are a few potential implementation barriers. The expertise needed can be lacking, especially on the facility side of the organization (this report is intended to rectify that barrier). And, in mixed environments such as colocation data centers, access to some servers may not be allowed by their owners. Nonetheless, onboard server sensor data offers a valuable and often untapped resource that data center operators can leverage to improve both data center energy efficiency and IT equipment operation.

Glossary

Agentless Monitoring
Agentless monitoring means the use of existing, embedded capabilities. There is no need for intrusive and expensive agents to be installed on the monitored devices.

API
An Application Programming Interface acts as an interface between two different applications so that they can communicate with each other.

ASHRAE
American Society of Heating, Refrigerating and Air-Conditioning Engineers.

BMC
Baseboard Management Controller - Also known as RAC (Remote Access Controller). Enables remote server management via built-in Application Programming Interfaces (APIs).

BMS
Building Management System – A building management system (BMS), also called a building automation system (BAS), is a computer-based control system installed in buildings that controls and monitors the building's mechanical and electrical equipment.

CLI
Command Line Interface - For humans sending and receiving information over the network using a specified syntax. For example, Windows DOS text-based interface is a CLI.

Client-Server Model
Client–server (host) architecture is most frequently identified with the request-response model: a client sends a request to the server, which performs some action and sends a response back to the client, typically with a result or acknowledgment.

CPU
Central Processing Unit - Commonly called the computer “chip”.

Data Center Networking Tools
Facilitate communication and information exchange between the connected devices and internal/external networks. These tools talk to the motherboards to gather data.

DCEP
Data Center Energy Practitioner program – A DOE funded certificate training program for saving energy in data centers.

DCIM
Data Center Infrastructure Management – DCIM is an integration of information technology (IT) and facility management disciplines to centralize monitoring, management, and control of data centers’ critical systems.

DCOI
Data Center Optimization Initiative - DCOI is part of the Federal Information Technology Acquisition Reform Act that requires the Federal Government to consolidate and optimize agencies' data centers.

DNS
Domain Name Service - Each DNS server holds information about the devices that are part of an organization's network. In cases where a sending device needs to find the address of a device that belongs to a different network, the DNS will locate the appropriate DNS server anywhere on the Internet to give you the appropriate target device's IP address information.

DOE
US Department of Energy.

ENERGY STAR
ENERGY STAR is a voluntary certification program initiated by the Environmental Protection Agency (EPA). It is designed to help consumers identify energy-efficient and environmentally friendly products. ENERGY STAR certification was first applied to computer products in 1992.

Ethernet
One of the most common technologies in use for local area networks (LANs). The Ethernet is a wired (with RJ45 connectors) network technology to connect devices.

FEMP
Federal Energy Management Program.

FMS
Facility Management System – Software designed to help businesses save time and money by properly managing their buildings, assets, and occupants more efficiently and effectively.

Ghost Servers
Computer servers that are powered on but are doing little or no productive work. They are also called zombie or comatose servers. They can waste vast amounts of energy.

GUI
Graphical User Interface - A system of interactive visual components for computer software. It displays objects that convey information, and represent actions that can be taken by the user.

HTML
Hyper Text Markup Language - A standardized system for tagging text files to achieve font, color, graphic, and hyperlink effects on World Wide Web (WWW) pages.

HTTP/HTTPS
Hypertext Transfer Protocol - A set of rules for transferring files such as text, images, audio, video and other multimedia files on the World Wide Web (WWW).

ICT
Information and Communications Technology - ICT is often abbreviated IT.

IDCM
Integrated Data Center Management solution – These solutions provide integration between critical facility infrastructure, the devices in the data center (servers, storage, network, etc.), and the application workloads running on those resources. IDCM brings together the capabilities and features of BMS, DCIM, and IT operations.

Internet
A global system of interconnected computer networks that use the TCP/IP protocol to link devices worldwide.

IP
Internet Protocol – Principal part of the Network Layer of the Internet protocol suite (TCP/IP), which is the global standard for Internet, LAN, and WAN. TCP/UDP and IP work together to transmit data over the Internet but at different levels (layers).

IP Address
A device can be located by its IP address or by its device or host name. If a device needs to connect to a device, but only knows its name (e.g., http://datacenters.lbl.gov) it can ask a computer configured with DNS software to find the IP address by its host name.

IPMI

Intelligent Platform Management Interface – The IPMI specification defines a set of common interfaces that system administrators can use to monitor server health and manage the system. IPMI is a vendor-neutral standard and one of the most-used acronyms in server management.

LAN and WAN
A Local Area Network (LAN) - A computer network that interconnects devices within a limited area such as a data center. By contrast, a wide area network (WAN) covers a larger distance. Ethernet and WiFi are the two most common technologies in use for local area networks.

LBNL
Lawrence Berkeley National Laboratory.

MIB
Management Information Base - Management data in the form of variables on the managed device organized in a data base, which describe the device status and configuration.

Network Device and Networked Device
A "network device" is a component that makes up the network infrastructure such as modems, routers, and switches. A "networked device" on the other hand refers to equipment that connects to a network, which includes computers, printers and most A/V gear (receivers, media hubs and servers, Blu-ray players, etc.), which operate in an Ethernet or Wi-Fi network or both.

Network Interface
The connecting point between two adjacent devices on a network.

Network Protocol
The method for two or more devices to understand one another.

OEM
Original Equipment Manufacturer - Generally perceived as a company that produces parts and equipment that may be marketed by another manufacturer. However, the term is also used in other ways, which causes ambiguity.

Onboard Sensors
Sensors on IT equipment motherboards.

OS
Operating System - Software that supports a computer's basic functions, such as scheduling tasks, executing applications, and controlling peripherals.

OSI
Open Systems Interconnection networking model – OSI differs slightly from the TCP/IP model. The conceptual OSI model gives guidelines on how communication needs to be done, while the TCP/IP protocols lay out standards on which the Internet was developed.

Redfish
The successor to IPMI, it is a relatively new (2015) global, open server hardware management standard where interoperability is key.

RESTful API
An API (Application Programming Interfaces) acts as an interface between two different applications so that they can communicate with each other. One of the most popular types of API is REST or, as they are sometimes known, RESTful APIs. A RESTful API provides a relatively high level of flexibility and freedom for developers.

Routers
Just as a switch connects multiple devices to create a network, a router connects multiple switches, and their respective networks, to form an even larger network. The router also allows networked devices and multiple users to access the Internet.

SNMP
Standard Network Management Protocol – SNMP is a standard way (protocol) of monitoring hardware and software from nearly any manufacturer. It comes preloaded on most devices. SNMP is a way for devices on a network to share information with one another. The protocol is a vital tool for network management.

Switches
Switches facilitate the sharing of resources by connecting together all the devices, including computers, printers, and servers, in a small business network. Thanks to the switch, these connected devices can share information and otherwise communicate with one another.

TCP
Transmission Control Protocol – Principal part of the Transport Layer of TCP/IP, which is a global standard for Internet, LAN, and WAN. TCP and IP work together to transmit data over the Internet but at different levels (layers).

UDP
User Datagram Protocol – Principal part of the Transport Layer of TCP/IP, which is a global standard for Internet, LAN, and WAN. UDP and IP work together to transmit data over the Internet but at different levels (layers).

WiFi
A wireless networking technology to connect devices. One of the most common technologies in use for local area networks.

WWW
World Wide Web - Refers to the online content that is formatted in HTML and accessed via the Hypertext Transfer Protocol (HTTP).

References

ASHRAE 2021. Special Publication, Thermal Guidelines for Data Processing Environments, American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc., Atlanta, GA. www.ashrae.org

DCEP, 2021. DOE Data Center Energy Practitioner certificate training program. http://datacenters.lbl.gov/dcep

Dell, 2021. OpenManage Network Manager https://www.delltechnologies.com/en-us/solutions/openmanage/index.htm

DMTF, 2021. www.dmtf.org

ENERGY STAR, 2021. ENERGY STAR Product Finder ENERGY STAR Certified Enterprise Servers | EPA ENERGY STAR

ENERGY STAR Active Certifications (2021) https://data.energystar.gov/Active-Specifications/ENERGY-STAR-Certified-Version-3-0-Enterprise-Serve/qifb-fcj2/data

ENERGY STAR, 2018. ENERGY STAR Computer Server Specification Version 3.0 www.energystar.gov/products/spec/enterprise_servers_specification_version_3_0_pd

Herrlin, M. K. 2008. Airflow and Cooling Performance of Data Centers: Two Performance Metrics. ASHRAE Transactions, Volume 114, Part 2. http://www.ancis.us/publications.html

Intel, 2021. Intel DCM (Data Center Manager). https://www.intel.com/content/www/us/en/software/data-center-overview.html

Intel, 2019. Statement from IPMI promotors http://www.intel.com/design/servers/ipmi/index.htm

IPACK, 2009. Energy-Efficiency through the Integration of Information and Communications Technology Management and Facilities Controls, Proceedings of InterPACK'09, July 19-23, 2009, San Francisco, California, USA

LBNL, 2020. Thermal Guidelines and Temperature Measurements in Data Centers http://datacenters.lbl.gov/sites/default/files/FINAL%20Thermal%20Guidelines%20and%20Temp%20Measurements%209-15-2020.pdf

LBNL, 2009. LBNL-3137E, Control of Computer Room Air Conditioning using IT Equipment Sensors. http://datacenters.lbl.gov/resources/control-computer-room-air-conditioning

Nlyte, 2021. DCIM tool from Nlyte.
https://www.nlyte.com/solutions/data-center-infrastructure-management-dcim/

Telcordia. 2012. (Kluge, R.) Generic Requirements NEBS GR-63-CORE, NEBS Requirements: Physical Protection, Issue 4, April 2012, Telcordia Technologies, Inc., Piscataway, NJ.
www.telcordia.com

Telcordia. 2001. (Herrlin, M.) Generic Requirements GR-3028-CORE, Thermal Management in Telecommunications Central Offices, Issue 1, December 2001, Telcordia Technologies, Inc., Piscataway, NJ.
www.telcordia.com