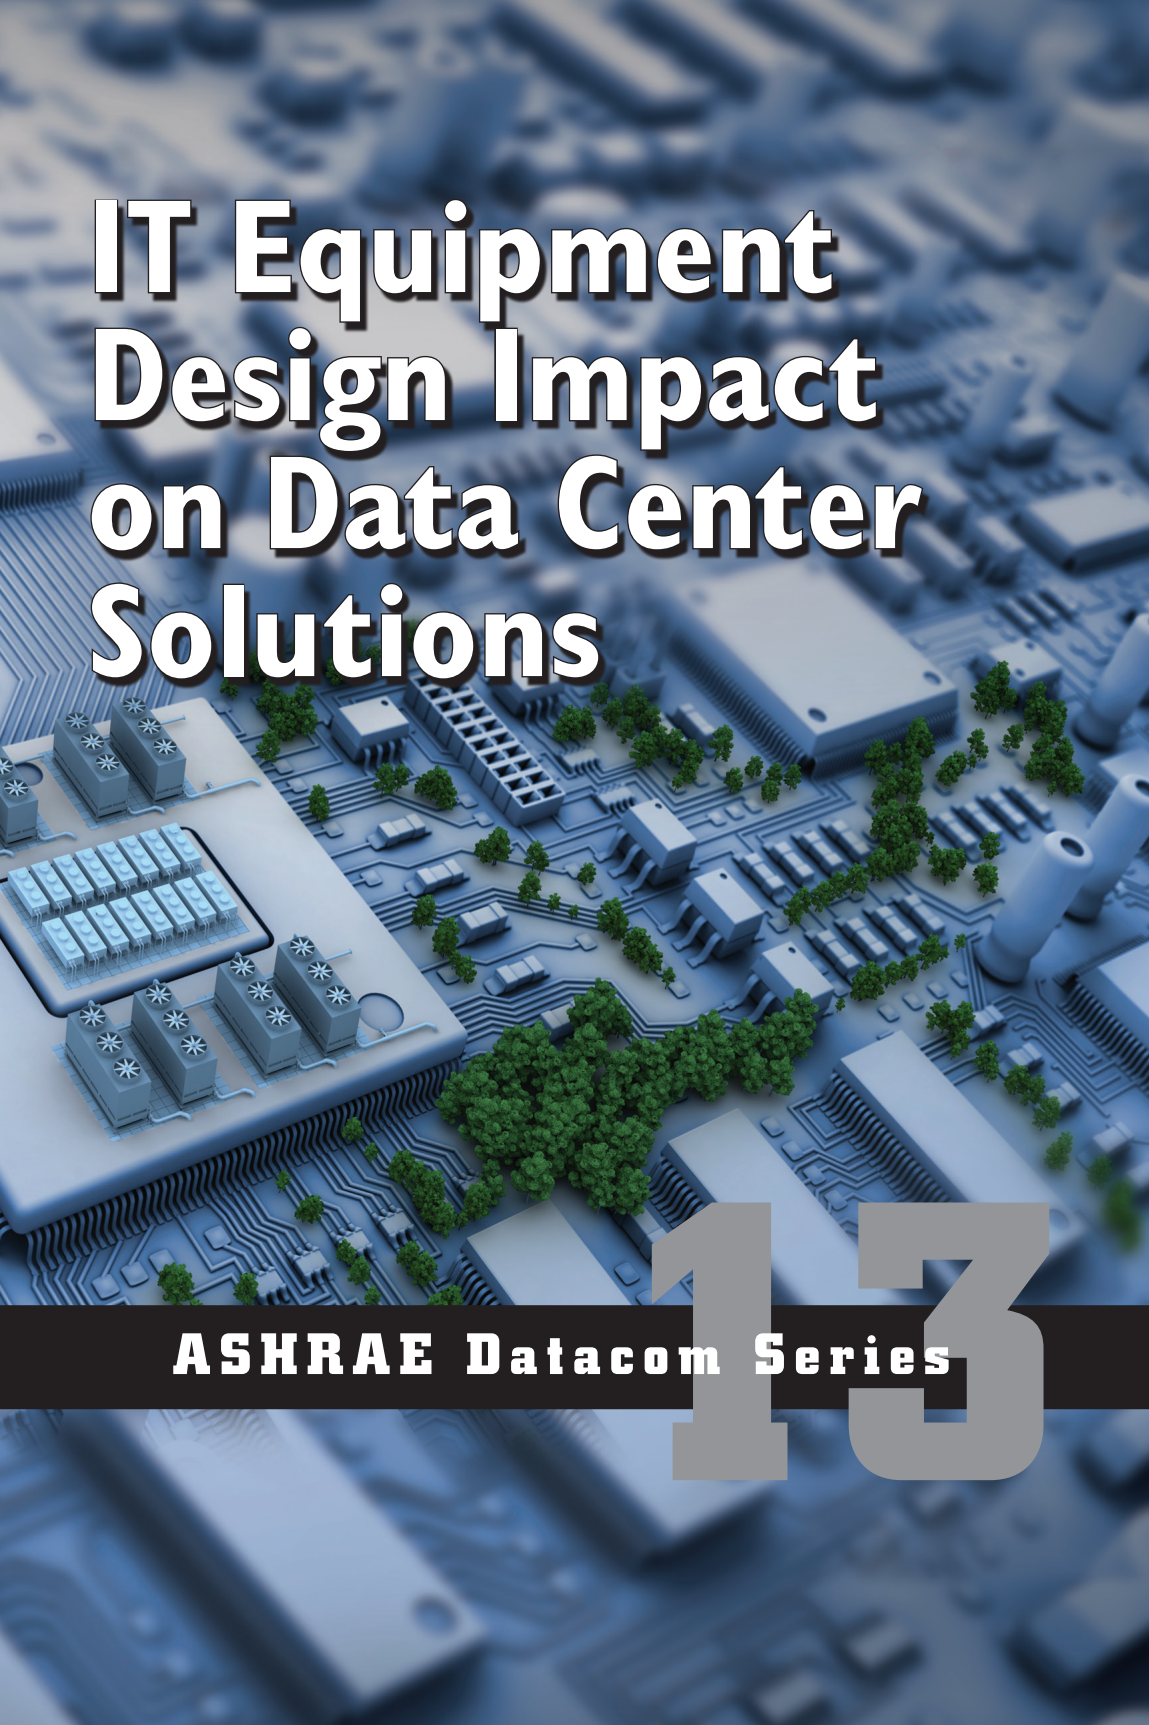

IT Equipment Design Impact on Data Center Solutions
ASHRAE Datacom Series
13

IT Equipment Design Impact on Data Center Solutions

IT Equipment Design Impact on Data Center Solutions *is authored by ASHRAE Technical Committee (TC) 9.9, Mission Critical Facilities, Data Centers, Technology Spaces and Electronic Equipment. ASHRAE TC 9.9 is composed of a wide range of industry representatives, including but not limited to equipment manufacturers, consulting engineers, data center operators, academia, testing laboratories, and government officials who are all committed to increasing and sharing the body of knowledge related to data centers.*

For more information on the ASHRAE Datacom Series, visit
www.ashrae.org/datacenterefficiency

For more information on ASHRAE TC 9.9, visit
https://tc0909.ashraetcs.org

Updates and errata for this publication will be posted on the ASHRAE website at www.ashrae.org/publicationupdates.

IT Equipment Design Impact on Data Center Solutions

ASHRAE Datacom Series
Book 13



Atlanta

ISBN 978-1-939200-19-8 (Paperback)
ISBN 978-1-939200-20-4 (PDF)

© 2016 ASHRAE
1791 Tullie Circle, NE
Atlanta, GA 30329
www.ashrae.org

All rights reserved.
Printed in the United States of America

Cover image by Joe Lombardo of DLB Associates.

ASHRAE is a registered trademark in the U.S. Patent and Trademark Office, owned by the American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.

ASHRAE has compiled this publication with care, but ASHRAE has not investigated, and ASHRAE expressly disclaims any duty to investigate, any product, service, process, procedure, design, or the like that may be described herein. The appearance of any technical data or editorial material in this publication does not constitute endorsement, warranty, or guaranty by ASHRAE of any product, service, process, procedure, design, or the like. ASHRAE does not warrant that the information in the publication is free of errors, and ASHRAE does not necessarily agree with any statement or opinion in this publication. The entire risk of the use of any information in this publication is assumed by the user.

No part of this publication may be reproduced without permission in writing from ASHRAE, except by a reviewer who may quote brief passages or reproduce illustrations in a review with appropriate credit, nor may any part of this publication be reproduced, stored in a retrieval system, or transmitted in any way or by any means—electronic, photocopying, recording, or other—without permission in writing from ASHRAE. Requests for permission should be submitted at www.ashrae.org/permissions.

Library of Congress Cataloging-in-Publication Data

Names: ASHRAE (Firm), author.
Title: IT equipment design impact on data center solutions.
Description: Atlanta, GA : ASHRAE, [2016] | Series: ASHRAE datacom series ; 13 | Includes bibliographical references.
Identifiers: LCCN 2016001806 | ISBN 9781939200198 (softcover)
Subjects: LCSH: Information technology--Equipment and supplies. | Computer networks--Equipment and supplies. | Data processing service centers--Design and construction. | Data processing service centers--Cooling.
Classification: LCC TK5103 .I85 2016 | DDC 004--dc23 LC record available at http://lccn.loc.gov/2016001806

ASHRAE Staff

Special Publications	Mark S. Owen, Editor/Group Manager of Handbook and Special Publications Cindy Sheffield Michaels, Managing Editor James Madison Walker, Managing Editor of Standards Sarah Boyle, Assistant Editor Lauren Ramsdell, Assistant Editor Michshell Phillips, Editorial Coordinator
Publishing Services	David Soltis, Group Manager of Publishing Services and Electronic Communications Jayne Jackson, Publication Traffic Administrator
Publisher	W. Stephen Comstock

Acknowledgments

PRIMARY CONTRIBUTORS

Dave Moss, book lead/author, Member ASHRAE
Office of the CTO—Data Center Strategist, Dell, Inc., Austin, TX
Robin Steinbrecher, author, Member ASHRAE
Principal Engineer, Intel Corporation, Dupont, WA
Jon Fitch, PhD, author, Associate Member ASHRAE
Data Scientist, DSS Engineering, Dell | Extreme Scale Infrastructure, Austin, TX
Dustin Demetriou, PhD, author, Member ASHRAE
Advisory Engineer, Advanced Thermal Energy Efficiency Lab, IBM Systems, Poughkeepsie, NY
Michael J. Ellsworth Jr, PE, author, Member ASHRAE
Senior Technical Staff Member and Technical Manager, Advanced Thermal Energy Efficiency Lab, IBM Corporation, Poughkeepsie, NY
Jason Matteson, author, Member ASHRAE
Senior Research Staff Member, Lenovo Server Group, Morrisville, NC
Roger R. Schmidt, PhD, primary reviewer and guidance to vision, Member ASHRAE
IBM Fellow Emeritus and Traugott Distinguished Professor, Syracuse University
Retired IBM Chief Engineer for Data Center Energy Efficiency
Member of National Academy of Engineering
IBM Corporation, Poughkeepsie, NY

ADDED THANKS TO OTHER COMMUNITY CONTRIBUTORS

Francisco Romero and Robert Adkins of IBM for writing the forward; David Richardson (Intel), Nishi Ahuja (Intel), Lori Matthews (Dell), Sandeep Ahuja (Intel)

ADDED THANKS TO REVIEWERS

Thomas Davidson (DLB), Greg Jeffers (McKenneys), Andrew Baxter (Page), John Groenewald (Chase), Jean-Christophe Bonnin (ATOS)

SPECIAL THANKS TO DON BEATY AND DLB ASSOCIATES

Neil Chauhan (DLB), Mike Mangan (DLB)

DISCLAIMER

Features and implementations described in this book do not necessarily represent those for any specific information technology equipment (ITE) or ITE manufacturer but rather general concepts and capabilities available in general, volume servers.

Contents

Foreword: The Changing Landscape of Data Centers ix

Chapter 1—Introduction 1

Chapter 2—IT Equipment 3

2.1 Servers 3
2.2 Switches and Networking 9
2.3 Storage Arrays 15
2.4 Tape Storage 19
2.5 ITE Airflow 20
2.6 Industry Environmental Specifications 24

Chapter 3—ITE Thermal Design. 25

3.1 Comparing Air and Liquid Cooling 25
3.2 Thermal Design—ITE Components 28
3.3 Component Power and Thermal Management Features 33
3.4 Thermal Design—Platforms 38
3.5 Thermal Management 55

Chapter 4—Interactions Between IT Systems and the Data Center . . . 75

4.1 Differential Air Temperature Δt:
A Metric of Increasing Importance 75
4.2 Future Flow Rate Increases 78
4.3 Pressure Effect of Coupled Systems 80
4.4 Data Center Transient Response. 94
4.5 Humidity, Contamination, and Other Considerations 96
4.6 Acoustics 101
4.7 Liquid Cooling. 102
4.8 Weight Implications 102

Appendix A—Component Thermal Design 103

A.1 Altitude Correction for Components (Air-Cooling Only) 105
A.2 Heat Sink Design 106
A.3 Heat Spreaders—Air-Cooled Memory 109

Appendix B—System Thermal Design 111

Appendix C—IT Telemetry Use Cases 117

C.1 Use Case: IT Equipment Power Monitoring 118
C.2 Use Case: Monitoring and Troubleshooting the Data Center 119
C.3 Use Case: Intelligent Thermal-Based Scheduling 120
C.4 Use Case: Thermal Modeling and Predictive Analysis 121

References 123

Glossary of Terms 127

Foreword: The Changing Landscape of Data Centers

Moore's Law continues to be a key driver in the economics of IT. Fitting more and more compute power into a shrinking footprint results in an increase in IT density (i.e., W/ft^2 or kW/m^2); however, it also comes with a commensurate increase in performance per watt and a corresponding decrease in the cost of that compute. What does not make the cover story of magazines and publications is that the amount of compute and hardware to perform a specific task may also be growing exponentially because of tasks continuously evolving to expand the user experience with increased functionality. In fact, the compute requirements may even be outpacing the growth of hardware capabilities. If the latest generation operating system was installed on a computer from 2010, how would it perform? If the latest application was put in the newest hardware, would any tasks be faster because the hardware is faster? The performance would probably be the same, except the application would provide a richer user experience and be handling more data. The explosion of data being generated today through cell phones, tablets, watches, thermostats, and online shopping, to name a few, is fueling massive hardware requirements on the back-end IT systems. The cost of running workloads is certainly not following the hardware cost curve. For example, only recently, a standard cloud deployment needed a single 1 Gbps link whereas today most workloads strongly favor a dual 10 Gbps uplink to handle the throughput required by the application. In short, every day brings added demand for more compute, storage, and bandwidth, and hence more data center space.

Disruption is around the corner in all aspects of our lives: social media, online retail, access to information, and entertainment. The applications and code behind today's technology are moving even faster than the hardware they run on and are becoming ever more complex. While the running of legacy applications gets simpler over time, the level of knowledge and understanding required to deploy and manage the disruptive ones—which are the ones the industry seeks—is continuously increasing. The challenge with legacy applications is the constant need for evolution and a certain degree of re-architecting to enable them to run on the newest generation of hardware. The current migration of workloads to the cloud is a perfect example of this. Companies are trying to take advantage of the speed, flexibility, and cost advantage provided by cloud infrastructure as a service, but they are running into serious challenges migrating their applications into a set of standardized infrastructures, which can even be a multitenant environment. This is not an easy problem to solve, and those who are unable to respond and adapt could be at a significant competitive disadvantage in terms of speed and flexibility.

Perhaps the single most important challenge in the IT space currently is, and will continue to be, the security of the massive amounts of data that companies and organizations have access to. There has been no shortage of hacks and security breaches with significant impact and exposure to end customers. The more complex systems get, inherently the harder it is to secure them. In an effort to mitigate this, governments have started rolling out legislation on where data can reside, who can have access to it, and how it can be used. This poses an interesting challenge for organizations, especially those that operate across multiple countries. Cloud environments pose their own unique set of challenges given the extensive use of shared pools of resources to deliver value. Compliance standards have evolved as a way to manage the risk associated with these environments; however, there are more and more risks coming up each day, which makes the landscape increasingly difficult for companies to navigate. Private institutions define some standards, governments define others, and organizations try to deploy all of them, but there is no clear end in sight.

So what does this all mean for data centers? It means that data centers have to evolve just as fast as the world around them. They need to have a footprint that is flexible, scalable, and adaptable. They need to be able to move as fast as new applications are developed and keep up with new ideas, new architectures, and new ways of thinking—all in real time.

To start, consider flexibility. Historically, data centers have been constructed for stable workloads and with somewhere around a 10 to 20 year life expectancy. They are architected before they're operational and the infrastructure is made to fit in them. Today's applications or workloads don't even resemble those from 10 years ago, and the development cycle and evolution of these workloads is only accelerating. Can a brand-new application be deployed in a 5- to 10-year-old data center? It certainly can. Will the performance and cost point of the application be optimal? It probably will not. So how can a data center be designed to be more flexible? The software-defined data center is one in which the administrative and operational tasks that today are manual and delivered using software that can dynamically optimize the compute, storage, networking, and facilities infrastructure. Power and cooling can be added, taken out, and rerouted with the click of a button. And, not only that, the infrastructure that runs in the data center will probably look very different in 5 years.

While the concept of scalability is fairly similar to flexibility, it does carry a few other connotations. The fundamental principle of scalability is centered on the ability to add and take out capacity as needed. Currently, the implementation of a scalable data center is perhaps best represented by the containerized data center. If more capacity is needed, another container is added. If less capacity is needed in a location, move the capacity to where it is needed. There are still challenges to solve around these deployments, but they represent what the future will look like. An example of an industry sector that most clearly captures the need for scalable deployments is the video gaming industry. A video game goes live with an expected amount of capacity. If it is successful, capacity needs to ramp up quickly or risk losing revenue. If demand is not as great, it needs to ramp costs down or risk going under. This scalability needs to be underpinned by best-in-class reliability and availability. In the

digital economy, being down and having to bring complex environments back up carries a very steep price.

While flexibility refers to the ability to morph the data center and dynamically reconfigure it to support whatever workloads are running in it, adaptability means it also has to be able to cope with the ever-changing landscape in which it operates. Cost and availability of power sources and fuels, environmental compliance requirements, and data compliance issues are a few of the ones that will directly impact the environment that surrounds the data center. Today, it is estimated that 2% to 3% of the total power consumed in the United States is consumed by data centers (Koomey 2011; DatacenterDynamics 2012). What will it be in the future? What will be the implications of that? While the industry has started getting a taste of this, it still has a long way to go.

This foreword has outlined some of the key guiding principles for the future of data centers. All of them are aligned with the goal in mind of optimizing the cost of the infrastructure layer of technology, fueled by the pressure to reduce the costs of running workloads and thus making them more accessible and used by a broader customer set. The basic concept that underpins this logic is that data centers need to be designed and built for the unknown. Applications need to be the focus. Disruptive industries such as social media, big data, analytics, and online entertainment will drive the requirements. What do they need to thrive tomorrow? They will certainly be better positioned if they reside in a facility that is built with flexibility, scalability, and adaptability in mind. This book focuses on the factors of the IT hardware thermal design that will aid data center operators in achieving desired data center characteristics such as flexibility, scalability, and adaptability.

Francisco Romero and Robert Adkins, IBM

1

Introduction

The ASHRAE Datacom Series along with the ASHRAE Handbook series are important references for those seeking to understand how data center cooling systems are designed and the methods used to optimize those cooling systems (ASHRAE 2015a; 2015b). However, few publications provide useful information to the data center design and owner community describing the information technology equipment (ITE) cooling designs and their interactions with the data center. This book serves to fill this gap through the collective knowledge and experience of ITE thermal engineers, architects, and experts in ITE and data center thermal management.

This book creates a reference for the data center and IT professional covering the following:

- An overview of IT equipment cooling design and thermal management and the differences between servers, storage, and networking equipment (Chapters 2 and 3)
- Power and thermal characteristics trends impacting. data center operation including ΔT (temperature differential across IT equipment) and pressure interactions (Chapter 4)

AUDIENCE

This book can be used in different ways depending upon the following readers' viewpoints:

- Data center infrastructure designers—to better understand how the ITE they need to support within the data center responds to the environment in which it is placed
- Data center owner/operators—to enable selection of features and ITE implementations that will best position the data center to operate optimally
- IT professionals—for a general understanding of the servers and other systems for which they may be responsible for purchasing

2

IT Equipment

Information technology equipment (ITE) is computing and telecommunications equipment designed to store, retrieve, transmit and manipulate data. Different types of ITE serve different functions, use different components, and vary in the level of sophistication of thermal design and optimization. Similar equipment types are further tuned and optimized for the markets and users they serve. This chapter provides an overview of servers, networking, and storage systems and their components.

2.1 SERVERS

Server is a broad term describing a specific piece of ITE that provides computing capability and runs software applications in an environment networked with other ITE, including other servers. Most servers contain the following major hardware building blocks: processors, memory, chipset, input/output (I/O) devices, storage, peripherals, voltage regulators, power supplies, and cooling systems. Additional application-specific integrated circuits (ASICs) may be necessary, such as an onboard redundant array of independent disks (RAID) controller and a server management controller, along with discrete components such as capacitors and resistors, which dissipate relatively little power. The key components that drive the thermal design are discussed in this section.

The processor, also referred to as the *central processing unit* (CPU), is the primary source of heat generation within a server. Aside from the basic processing of data and instructions to provide an output result, processors may also have many more features for managing data and power throughout a system. The processor die is generally housed in a package that includes a substrate (i.e., a small printed circuit board [PCB] for breaking out signals) and a lid, as shown in Figures 2.1 and 2.2. The lid or case more evenly distributes heat to an attached cooling component such as a heat sink (air-cooled) or cold plate (liquid-cooled) as shown in Figures 2.3 and 2.4. In most cases, a socket is used to enable removal and replacement of the processor on the motherboard. Some lower-powered processors are lidless, using direct attachment of a heat sink directly on top of the die.

The terms *volatility* and *latency* are important in understanding both memory and storage. *Volatility* refers to the capability for data storage when power is removed from a device. The data is considered to be permanently stored (until it is overwritten) if it is nonvolatile. Alternately, the loss of the data when power is removed implies that the device is volatile. *Latency* refers to the time to access the data for use in the compute stream. Dynamic random access memory devices (DRAMs) have

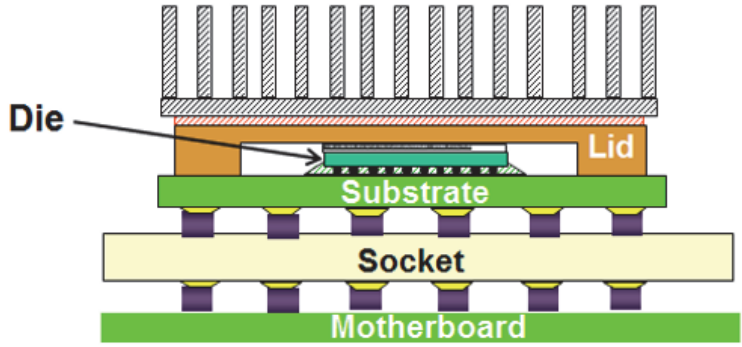


Figure 2.1 Processor package ingredients.

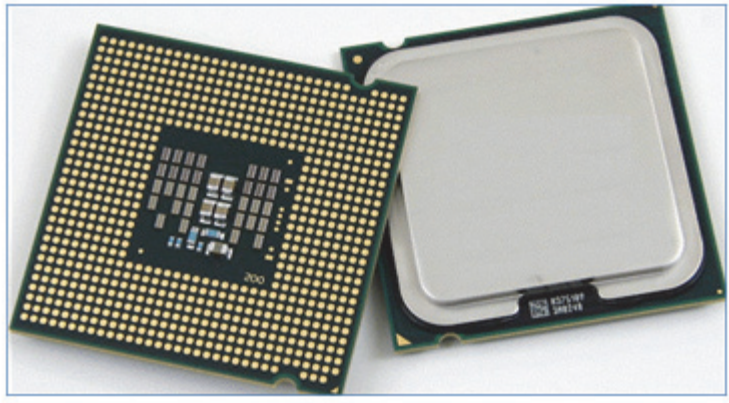


Figure 2.2 Processor package with substrate and lid.

been the traditional method for volatile storage of data directly in the execution path for the workload, enabling low latency (quick access). Hard disk drives (HDD), tapes and optical media are all nonvolatile data storage devices with relatively high latency (slow access).

With the introduction of flash devices, the line between memory and storage has blurred. A generic term for this type of device is *nonvolatile memory* (NVM). NVM can be used in both memory and storage applications due to its nonvolatility and relatively low latency. Solid-state drives (SSD) were the first example and are used in place of HDDs where lower latency is valued over the high-latency HDDs.

Memory temporarily stores data that has been processed or is to be processed. DRAMs are packaged with many devices on one PCB. An example form factor is

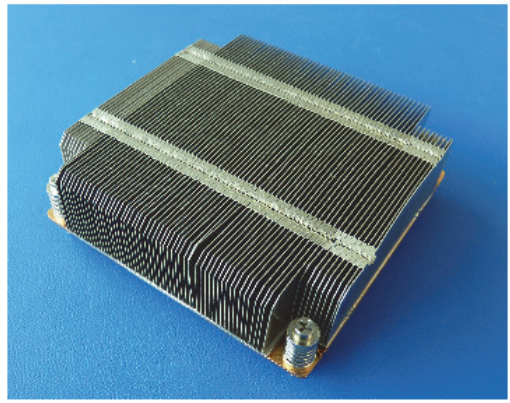


Figure 2.3 Processor heat sink for air-cooling.

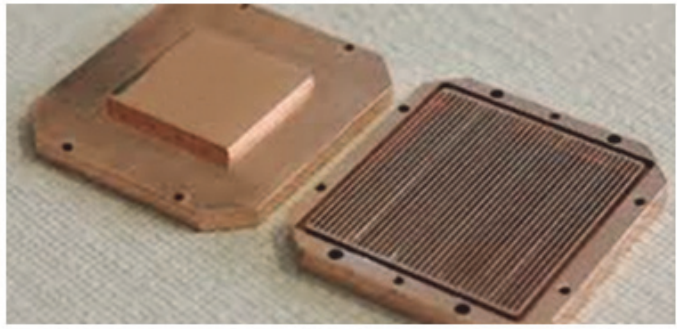


Figure 2.4 Processor heat sink for liquid-cooling (disassembled).

the dual in-line memory module (DIMM) shown in Figure 2.5. These PCBs have edge connectors that allow them to be installed in sockets mounted on the board. Memory hierarchy is driven by the latency to access the data and the capacity to store the data. Onboard memory is one step removed from the on-processor memory that stores data more directly in the execution path (with lower latency) and may be integrated on the processor die or packaged on the processor substrate. NVM-based DIMMs are becoming more commonplace, demonstrating the crossover between memory and storage and the potential for more than one usage for the same device.

HDDs, SSDs, tape, and optical media are used to provide nonvolatile storage of data for longer periods of time than DRAM- and NVM-based DIMMs. Tape is the only medium rated for archival data storage (30 years). HDDs are not an archival storage medium but are commonly used to store data for several years. Figure 2.6

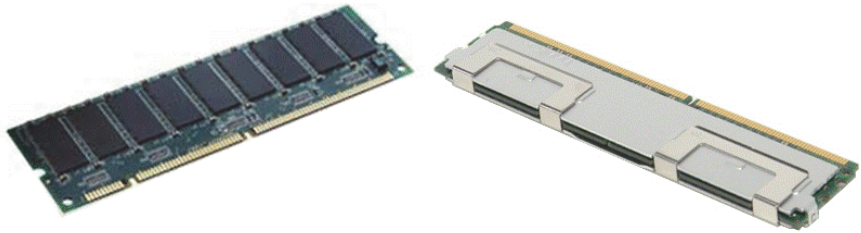
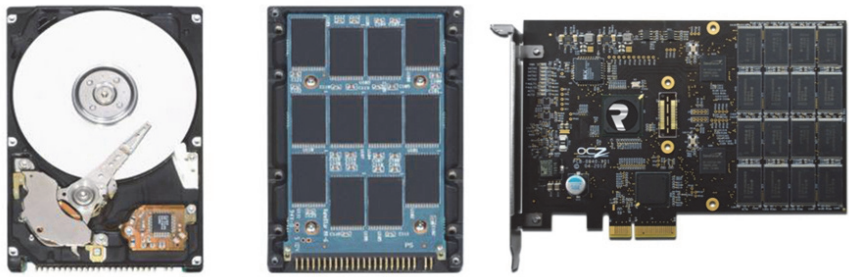


Figure 2.5 Typical dual in-line memory module (DIMM) with and without a heat spreader.



(a) (b) (c)

Figure 2.6 (a) HDD with cover removed, (b) NVM small form factor SSD, and (c) a PCIe NVM SSD.

shows some of the form factors and different types of data storage components. NVM is now being used in SSDs and is transitioning into form factors more closely coupled to the computing core (processor and local memory) in the form of peripheral component interconnect express (PCIe) cards and DIMMs. This transition will drive higher server heat loads than previously projected. SSDs use NVM for information storage, whereas traditional hard drives use disks of rotating magnetic media.

NVM maintains data even when no power is applied. Magnetic storage, such as HDDs, continues on a trajectory of lower cost as measured by cost per byte stored. However, NVM is also trending to lower cost per byte and has been displacing rotating media and, in some cases, volatile memory. SSDs are becoming available in a new form factor: PCIe. Many PCIe SSDs plug into a PCIe slot instead of a drive bay. PCIe allows for much faster data transmission rates than the connection to a typical drive in a drive bay.

A comparison of active power across different types of HDD and SSDs is provided in the ASHRAE storage white paper (ASHRAE 2015c).

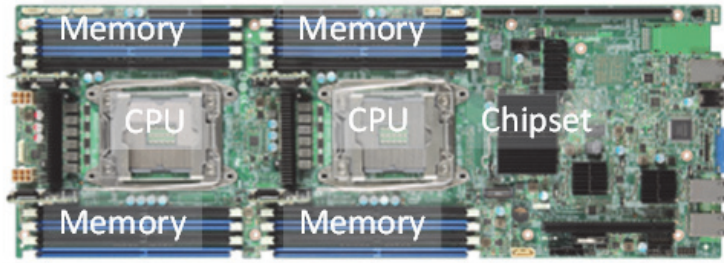


Figure 2.7 Dual processor socket motherboard.

The board (or motherboard) provides interconnections between the various components. Typically, the boards themselves are multilayered, with interconnects and circuitry residing on power and signal layers separated using dielectric layers. Voltage regulators that convert from power supply voltages to required silicon voltages are placed on the board. A typical board layout is shown in Figure 2.7.

The majority of mainstream ITE is designed for connection to an alternating current (AC) source. The ITE's power supply converts the input AC to lower voltage direct current (DC) levels for use by the subsystems in the computer. Power supplies are typically designed to handle a wide range of input voltage levels, enabling use around the world with a single supply that can handle input frequencies in the 50 to 60 Hz range.

Many servers are designed with redundant power supplies whereby the AC-DC conversion subsystem can tolerate the loss of a supply. In the past this may have resulted in lower overall efficiency, but governmental and market influences have driven the creation of new technologies that maximize efficiency throughout the load range, even when operating in redundant mode.

DC-to-DC regulators create the tightly controlled, relatively low voltages used by the silicon devices on the board. A board can easily have more than 25 regulators to create the various voltages required to power the devices in the system. For a detailed description of power delivery, both leading to the ITE and internal to it, see the ASHRAE book on best practices (ASHRAE 2009).

Other supporting components include I/O devices such as add-in cards that enable connection to a network or remote storage devices.

2.1.1 Rack-Mount Servers

Volume rack-mount servers are designed to fit within commonly available rack sizes, such as the 19 in. (0.48 m) rack form factor defined by the EIA/ECA Standard 310-E specification (ECIA 2005). The vertical dimension is expressed in terms of rack units or just "units" (U). One U or 1U represents 1.75 in. or 44.45 mm of vertical height within a rack.

Servers used for computing are available in standard rack mount and custom configurations. Typical dimensions and sizes for standard rack mount compute serv-

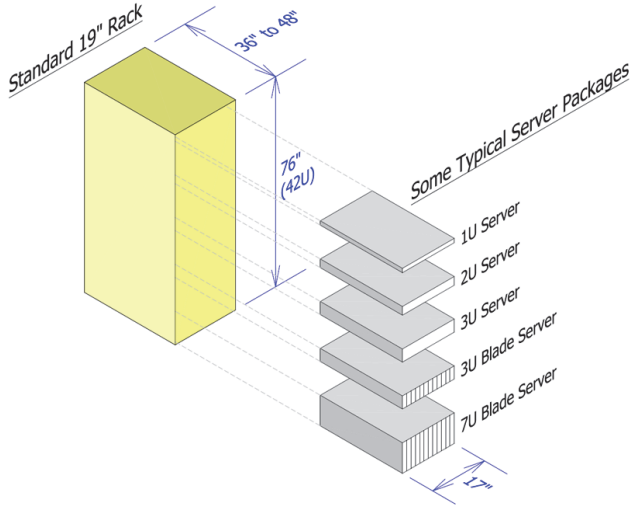


Figure 2.8 Typical compute server rack and packaging.

ers are shown in Figure 2.8. In most cases, a single server is housed in a full-width 1U or 2U chassis. A single server chassis may contain multiple server nodes where the nodes may share some of the chassis infrastructure. Half-width servers are now becoming popular because of the capability to populate two servers per U. An example 1U server is shown in Figure 2.9 with most of its major internal components labeled.

2.1.2 Blade Servers

To achieve even higher compute density than the 1U form factor, blade servers were introduced. Each manufacturer designs their blades based on their own packaging and design goals. Blade servers are the result of technology compaction, which allows for a greater processing density in the same equipment volume. The greater processing density also results in a greater power and heat density, further complicating data center cooling.

Each blade typically has a board with processor, memory, and I/O devices. Storage is either distributed on the blades or centralized with all components contained within a minimal amount of packaging, which results in blade-like dimensions (Figure 2.10).

Server components that had previously been packaged inside the tower/pedestal and rack mount form factors (e.g., fans, power supplies, etc.) are still required, but these components are now located within a chassis that is designed to house multiple blade servers in a side-by-side or stacked configuration (Figure 2.11). These chassis are typically 3U to 10U tall and can house many blades.

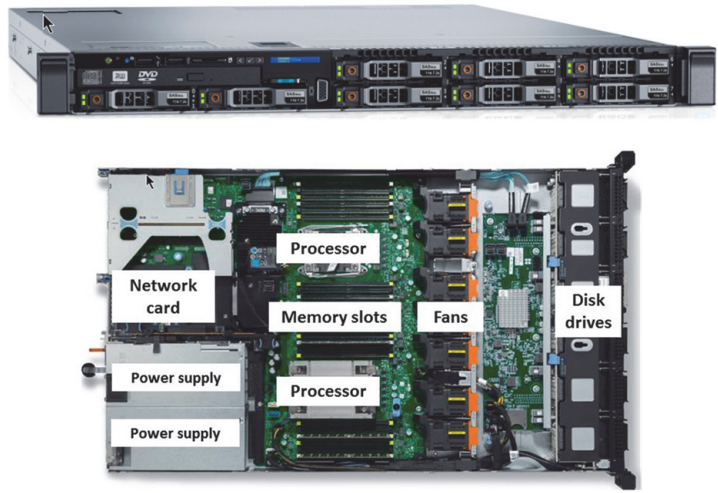


Figure 2.9 1U dual-socket volume server

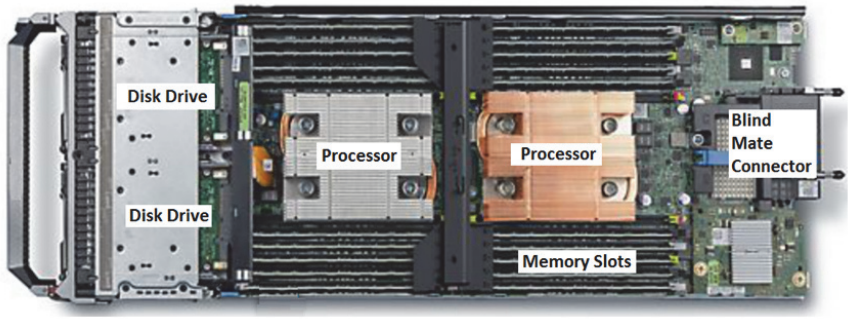


Figure 2.10 Example of a blade.

2.2 SWITCHES AND NETWORKING

Switches enable efficient communication between devices connected on a network. Many networks are designed in a hierarchical manner, with many devices (in the case of a data center, servers and storage devices) connected to a switch, which is connected to another switch at the next level of the hierarchy, and so on. Another common topology is a mesh configuration, in which many peer network switches are connected to one another to form a single level of the hierarchy. For the purposes of this overview, a hierarchical topology is assumed with three different levels: core, distribution, and edge. For general information on networking technol-

ogy and more detail on networking devices, refer to the ASHRAE white paper on network equipment (ASHRAE 2014a).

Figure 2.12 shows labeled functional elements inside of a typical top-of-rack (TOR) switch. Data enters through the ports and travels to the physical layer (PHY) chips. The ASIC chip decides where that data needs to go and sends it back out through the correct port. The CPU controls both the PHY and the ASIC. The CPU



Figure 2.11 Example blade server chassis.

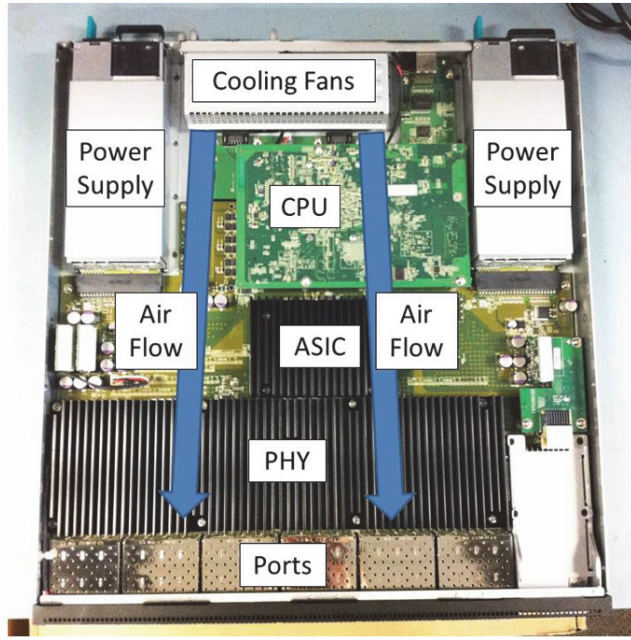


Figure 2.12 TOR switch with top cover removed.

can take data off of the network, process it, and then send it back out onto the network. The CPU is considered out-of-band because it is off the network and does not function at the full I/O speed. Encryption is an example of data processing done by the CPU. One of the key functions of the ASIC is to handle priority flow control (PFC) rules. PFC rules prioritize which data is sent first and which data can wait. One example of where PFC rules come into play is the transmission of video. Video is nearly always given a high PFC priority, because any delays in transmission could cause a glitch that would be noticeable to the end user. The ASIC also performs valuable functions such as blocking network traffic and stopping a flood of data, if this should occur.

The functional elements of large data center or telecom central office networking equipment are very similar to those of the small TOR switch, except that the size and scale of the system is larger. Also, larger switches tend to use optical ports instead of copper Ethernet cable type ports.

Functions of a networking system are characterized in terms of abstraction layers. An abstraction layer (or abstraction level) is used as a way of hiding the implementation details of a particular set of functionality. A majority of data center networking equipment and software uses the ISO/IEC 7498-1 standard Open Systems Interconnection (OSI) 7-layer model (ISO/IEC 1996) for computer network protocols. In this model, similar communication functions are grouped into logical layers. A layer serves the layer numerically above it and is served by the layer numerically below it. A description of the purpose and functionality of the more commonly used OSI layers is given as follows:

- Layer 1: Generally considered the physical layer. A network hub, or repeater, is an example of a simple Layer 1 network device. Any packet entering a port of the hub will be "repeated" on every other port, except for the port of entry.
- Layer 2: Generally used for basic point A to point B routing and is considered the data link layer. A network bridge is a good example of a Layer 2 device. The bridge may interconnect a small number of devices in a home or office. Single bridges can also provide extremely high performance for specialized applications such as storage area networks.
- Layer 3: Intended to handle higher-level functions such as PFC and spanning tree protocol.
- Layer 4: The exact meaning of layer 4 is vendor dependent. However, it almost always starts with a capability for network address translation.
- Layers 5 and 6: These layers are not commonly used.
- Layer 7: Includes the graphical user interface layer of a personal computer or server. Layer 7 switches may distribute loads based on Uniform Resource Locator (URL) or by an installation specific technique to recognized application-level transactions. A Layer 7 command typically spans many commands across many different layers.

Common types of networking equipment, their functions and the OSI layers they use are given in Table 2.1.

Table 2.1 Common Types of Networking Equipment and Their Functions

Equipment Type	OSI Layers	Description
Repeater	OSI layer 1	Amplifies or regenerates the digital signals received while sending them from one part of a network into another.
Hub	OSI layer 1	Connects multiple Ethernet segments, making them act as a single segment.
Bridge	OSI layer 2	A device that connects multiple network segments along the data link layer.
Switch	OSI layer 2	A device that allocates traffic from one network segment to certain lines which connect the segment to another network segment. Unlike a hub, a switch splits the network traffic and sends it to different destinations rather than to all systems on the network.
Router	OSI layer 3	A specialized network device that determines the next network point to which it can forward a data packet towards the final destination of the packet.
Gateway	OSI layers 4 and 7	Device is placed at a network node and interfaces with another network that uses different protocols.

The core network function can be thought of as the gateway through which all data entering and exiting the data center must pass. As such, the core network equipment is connected either directly or indirectly to every device in the data center. The core switch is also connected to a service provider (SP), which is the "pipe" through which all data passes from the data center to the Internet. The distribution level of the network acts as an intermediate level between the core and edge levels of the network, and as such can offload some of the work the core network equipment needs to do. Specifically, the distribution level is useful for passing data between machines inside the data center or aggregating ports to reduce the number of physical ports required at the core. The edge network level consists of switches that connect directly to devices that are generating or consuming the data, and then passing the data up to the distribution or core. A data center network implementation may have all three of these levels, or it may combine or eliminate some of the levels.

Figures 2.13 through 2.15 are simplified diagrams showing how these concepts might be implemented in a data center, and where particular types of networking equipment may be located. Understanding where the equipment is located assists the equipment supplier and those responsible for deploying the network to understand the environment that the network equipment will be exposed to.

For each of the types of networking equipment shown in these figures, a preferred infrastructure configuration is shown. The infrastructure configurations