

Part 1: Information Technology Equipment

Changing Landscape Of Data Centers

BY DONALD L. BEATY, P.E., FELLOW ASHRAE; DAVID QUIRK, P.E., MEMBER ASHRAE

Information technology equipment (ITE) continues to evolve to keep pace with the rate of change of software and the feature-rich services they provide. The HVAC systems and their interactions with the ITE will also need to evolve. While the “occupant” of a data center is the software, at its core, the ITE design dictates how that occupant load manifests itself within the data center. By understanding the ITE, its changes, and how it interacts with the data center HVAC, data center designers will be more equipped to handle this changing landscape.

A tremendous variety of ITE populates today’s data centers, owing to the exponential growth of data generation, data storage, and data consumption that permeates our lives more each day. Traditional servers now cohabitate in the data center space with huge storage arrays and powerful supercomputing clusters, all with widely different workloads that stream our videos, maintain our homes, route our communications, predict our weather, and generally accomplish things previously thought not possible!

And yet, all of these things are accomplished by equipment in data centers, using similar resources and encountering similar challenges. Data may be just a series of 1’s and 0’s, but manipulating trillions of them consumes electricity, and lots of it. That electricity is dissipated in the form of heat, which itself must be dissipated, or the bulk of equipment would almost certainly incinerate.

This is Part 1 in a series that will explore the current state of ITE in the data center, covering the different equipment types, their roles, and their general arrangement within the data center. This will not be comprehensive by any means, but rather serve to create a solid foundation from which to explore in more detail the changing landscape of data centers.¹ Future parts in this

series will cover ITE thermal design and the interactions between IT systems and data center cooling systems.

Data Center Design

Data center infrastructure designers are tasked with wrapping a shell, power, and cooling around ITE. To do this effectively, flexibly, efficiently, and in a scalable manner, we need to better understand how the ITE is designed, how it’s changing, and why it’s changing.

Historical data center design philosophy often considered the ITE as a homogenous load separated into aisles of equipment and arranged in hot/cold aisles for proper thermal management.² The goal of the HVAC was to provide the proper entering air conditions to the ITE.³

Today’s ITE consists of several very distinct types of hardware (servers, networking, and storage), each with their own unique designs and functions for supporting different types of data and services. Each of these types of ITE will correspondingly interact with the data center infrastructure in different ways.

In generic terms the ITE is divided into these basic functions:

- Servers: data computing and processing;

Donald L. Beaty, P.E., is president and David Quirk, P.E., is vice president of DLB Associates Consulting Engineers, in Eatontown, N.J.

- Networking: data transmission; and
- Storage: data storing.

In all cases, the use of data consumes energy and releases heat within the data center. The most efficient and effective management of that heat is dependent upon a more detailed understanding of the ITE design and its interactions with the data center infrastructure.

Servers

In IT terms, “server” broadly refers to equipment that provides utility to a “client.” Servers not only provide the webpages and files that we are most familiar with, but more abstractly provide computational resources across networks to be consumed by services.

The role of the server has grown immensely through the years to meet our demands, and, through specialization, servers may achieve higher degrees of throughput, speed, and energy efficiency. Therefore, the server design of yesterday is not the server design of today, and will not be the server design of tomorrow.

The central processing unit (CPU) has traditionally been the heart of the server, supported by memory, storage, and other peripherals. A CPU consists of millions or billions of individual switches that each produce heat as they operate billions of times per second. As such, they are one of the highest heat producing components within the server and a primary concern for cooling, accomplished by either air or liquid.

The CPU is usually attached via a socket to a motherboard, which provides power to and interconnections between the components of the server located directly on the motherboard as well as other peripherals.

For CPUs to perform their function, they require memory to hold in the inputs and outputs, and it must be incredibly fast to keep up with the demands of the CPU. As a technical result, the memory closest to the CPU is volatile memory, or memory that does not retain its data without power, as this yields the highest performance. Static random-access memory (SRAM) is the fastest and is located on the same die as the CPU, whereas dynamic random-access memory (DRAM) is somewhat slower and typically located on memory chips near the CPU. Later, the important resulting data is written to non-volatile storage media, but at slower speeds.

Traditional types of non-volatile storage found in the data center include hard disk drives (HDD) and tapes,

although solid state drives (SSD) are increasingly prevalent due to their comparatively low latency and decreasing cost. SSDs use flash memory, a non-volatile form of memory that involves no moving components, unlike HDDs and tapes.⁴

To increase server performance for some types of workloads, a current trend is using SSDs closer to the CPU via peripheral component interconnect express (PCIe) as their costs decrease and performance and capacity increase to where they offer attractive benefits compared to adding more DRAM.

However, HDDs and tapes still play an important role in the data center. HDDs use magnetic storage on rotating disks and offer low cost and fast access, whereas tapes are significantly slower but can store data with unparalleled economy and are the only format to offer an archival-appropriate lifespan (e.g., 30 years).

The most common form factor for servers is the 19 in. (483 mm) wide rack-mount server. The design is a time-less standard defined by the EIA-310 specification⁵ that has filled the majority of the world’s data centers for decades. It is based on 1.75 in. (44.5 mm) tall rack units (U), and equipment is specific as 1U, 2U, 4U, etc., referring to its height. While most equipment makes full use of the width, half-width equipment is an increasingly popular alternative that offers potentially higher density per U.

To further increase density beyond the standard layout, blade servers offer increased flexibility for manufacturers. Blade servers typically consist of a standard-width chassis that occupies between 3U and 10U of a traditional rack. However, the servers are individual “blades” of various dimensions, which plug into the chassis. Additionally, shared resources such as cooling fans and power supplies may be located on the chassis, further broadening the limits of what is possible with blade design (*Figure 1*).

Networking

Networking is the way servers are connected to other servers, other resources, and to the world outside the data center. Network topologies, the ways that networking equipment is connected, is a rapidly evolving topic that varies greatly based on the size and function of a workload. Traditional networking can easily be a source of bottlenecks if not optimized for a workload, but the current trend toward software-defined networking




FIGURE 1 Typical compute server rack and packaging.




FIGURE 2 Possible switching hierarchy based on network size.

(SDN) meets the dynamic demands on the network through abstraction and separation of the control from the hardware.

At the basic level, switches provide communication between devices by receiving data and forwarding it to its destination via the correct port, in contrast to a hub that receives data and forwards it to all other ports. Since it would be impractical for every server in a data center of any considerable size to be connected to a single switch, often multiple switches are connected to form a hierarchy with data flowing up and down the hierarchy as required. Such a hierarchy can be functionally characterized into core, distribution, and edge switches.⁶

Based on the size and design of the data center, network implementations may take many forms. At the simplest level, equipment could be connected to the same (core) switch. As the number of servers increases, it may be optimal to have the servers connected to edge switches, which are in turn each connected to the core switch. If the needs of the network exceed even that configuration, edge switches may be located within the server racks that each communicate with row level distribution switches that each join multiple edge switches to the core switch (Figure 2).

Current switch design, like nearly everything in the data center, is largely driven by performance. At its core, data is received through a port and reaches the physical layer chips. These communicate with an application-specific integrated circuit (ASIC) chip, which is a highly specialized chip designed to route

data as quickly and intelligently as possible. The data flow through the physical layer to the proper outgoing port. Switches may also include a CPU that handles additional data processing such as encryption, but processing data adds significant latency to the otherwise low-latency switching.

In addition to just forwarding data to the right destination, the ASIC chip may also enforce priority flow control (PFC) rules that prioritize some data and deprioritize other data to provide the desired result, such as prioritizing streaming video to be delivered without interruption.

Networking equipment may also be classified based on an abstraction of its function within the data center, using the open system interconnection (OS) seven-layer model. The majority of networking equipment resides at Layers 1 to 3, which are the physical layer, the data link layer, and the network layer, respectively. A description of each layer follows:

- Layer 1: Physical layer equipment such as hubs and repeaters only deal with the raw bit stream.
- Layer 2: Data link layer equipment such as bridges and switches deal with routing entire frames from Point A to Point B.
- Layer 3: Network layer equipment such as a router deals with higher order functions such as addressing, path determination, and subnets.
- Layers 4 to 7: These serve even higher-order functions, ranging from between networks to ultimately communication with the application layer and the human-machine interface.

Advertisement formerly in this space.

Advertisement formerly in this space.

Storage Arrays

A third major component of ITE within data centers is for the purpose of storing data, which is typically stored for one of four reasons. Data may be stored online, meaning it is current and able to be accessed rapidly. Data may be for backup to prevent the loss of information. Data may be for archival purposes, which should be accessible in the long term but is rarely needed. Or it may be for disaster recovery, which is similar to archival except that it is stored in a different physical location to guard against natural disasters.

A typical consumer's computer may include one or two hard drives that are individually managed by that machine. However, when servers need to store data, especially where that data needs to be accessible by tens to thousands of servers, it simply is not feasible for each server to manage individual hard drives. Storage arrays are the solution to this problem.

Storage arrays abstract individual storage devices (HDDs or SSDs) and aggregate them into larger arrays that appear as a single logical unit. The presentation of devices may be as just a bunch of disks (JBOD), but commonly the storage devices are grouped together into a redundant array of independent disks (RAID) to improve performance or enhance redundancy.

RAID works by "striping" data across multiple devices, which can reduce the data written by each disk, increasing how quickly the data can be written (RAID 0). The RAID array may also use additional devices to write data redundantly across multiple disks (RAID 1).

More complex RAID configurations can provide both performance and redundancy. For example, RAID 4 and 5 do this by striping data across multiple devices, and then calculating the parity of each bit in each corresponding stripe of data (essentially whether the sum of bits is even or odd) and writing the parity bit to a redundant device. If any stripe is lost, the parity of the remaining stripes can be used to determine the lost bit. RAID 4 uses a dedicated parity device, whereas RAID 5 distributes the parity bits across all of the devices.

Beyond RAID itself, storage arrays typically have redundant power supplies, RAID controllers, and ports to provide resiliency and reduce points of failure. Each device in the storage array is accessible by both controllers, which may use either power supply, and which may use either set of ports to communicate.

As mentioned before, tape storage still plays a role in the data center, particularly for archive and backup, in spite of newer technologies. Tape formats have been updated throughout the years with higher capacities and faster speeds, now nearing 10 terabytes each. Although individual tape drives are available, tape drives in the data center are most often found within tape libraries that manage a large number of tapes at once.

Summary

Today's ITE consists of several very distinct types of hardware (servers, networking, and storage), each with its own unique designs and functions for supporting different types of data and services. Each of these types of ITE will correspondingly interact with the data center infrastructure in different ways.

It is reasonable to believe that the IT industry is currently undergoing more rapid change than almost anything it has experienced. Not only is data within existing segments being created, stored, and processed more quickly than ever before, but entirely new IT segments are emerging, such as virtual reality (VR) and self-driving vehicles. These changes are placing increased demands on data centers and the ITE in them.

By understanding the fundamental roles of ITE (servers, networking, and storage), its changes, and how it interacts with the data center HVAC, data center designers will be better equipped to handle this changing landscape.

References

1. ASHRAE. 2016. *IT Equipment Design Impact on Data Center Solutions*. Atlanta: ASHRAE.
2. ASHRAE. 2009. *Design Considerations For Datacom Equipment Centers*, 2nd Edition.

Atlanta: ASHRAE.

3. ASHRAE. 2015. *Thermal Guidelines For Data Processing Environments*, 4th Edition. Atlanta: ASHRAE.

4. ASHRAE. 2015. "Data Center Storage Equipment Thermal Guidelines, Issues, and Best Practices."

5. EIA/ECA-310-D-1992, *Cabinets, Racks, Panels, and Associated Equipment*.

6. ASHRAE. 2014. "Data Center Networking Equipment—Issues and Best Practices." ■

Advertisement formerly in this space.