



Components of a modern data center

A glossary of 37 modern data center terms you need to know

Version 1.0

May 16, 2017

Components of a modern data center



Why the modern data center matters

The modern data center as a term is getting more attention among today's IT leaders. And for good reason. Cloud computing, flash storage, software-defined networks, containers, and a blossoming number of orchestration and automation tools are coalescing to form the foundation of a modern data center — a requirement for enterprises in our digital age.

Perhaps most important is the fact that a data center isn't necessarily a physical place. Most enterprises consider a data center as their own facility, a hosted facility, or even a public cloud like AWS, Azure, or GCP. A modern data center stitches this all together as a single infrastructure to run applications and digital services.

To help you navigate all the confusion, we've compiled list of terms and definitions of 37 of the most important modern data center terms.

A glossary of 37 modern data center terms

1. **3D NAND** — A next-generation type of nonvolatile memory technology (Flash) that is becoming more mainstream and prevalent among enterprises. It has the advantage of being able to pack more bits into the same dimensions as older NAND technology. Most Flash memory storage today is still planar, meaning two dimensional. However, as photolithography reaches its practical limits, squeezing more bits into each Flash NAND cell becomes more difficult. Chip makers are going vertical with 3D NAND. Think of 3D NAND as a multistory building and 2D NAND as a single-story building. Both occupy the same amount of real estate (X-Y dimension), but the multistory building is more efficient within the same space because it expands upward.
2. **Application layer** — It is the layer (there are seven layers in the traditional OSI stack) that is closest to the end user in the conceptual model that characterizes the communications functions of a computing system. This means both the application layer and the user interact directly with the given software application being used. The application layer provides full end-user access to a variety of shared network services for efficient data flow. The application layer becomes increasingly important in virtualized and containerized environments where it is abstracted from the physical infrastructure where it runs. It also gives way to applications that program their own infrastructure needs. See "application-specific policies" for more.
3. **Application-specific policies** — Policies (typically as they relate to infrastructure such as servers, storage, networks, and security in the modern data center) that are specifically tied to individual applications and the retrieval of data in a bare metal, cloud, container, or virtual machine environment. Application-specific policies

Components of a modern data center



enable a multi-tenant environment in which each application can have its own unique infrastructure and SLAs.

4. **Automation** — Automation is a key concept in cloud computing. Automation is what sets a cloud infrastructure apart from a virtualized infrastructure. It includes the ability to provision resources on demand without the need for manual, human intervention. Automation is often combined with orchestration to create the ability of a service to be integrated with and fully support the many tools now available to IT to gain control and mastery of their operations. For example, a software-defined storage or software-defined networking solution is one that can easily plug into automation and orchestration tools used in the rest of the data center without requiring customization or modification for a particular environment.
5. **Cloud Foundry** — An open source cloud platform as a service (PaaS) originally created by VMware and is now part of Pivotal Software. It is governed by the Cloud Foundry Foundation and is a PaaS on which developers can build, deploy, run, and scale applications in both public and private cloud environments. The platform leverages containers to deploy applications and enables businesses to take advantage of innovation from projects like Docker and Kubernetes to increase the ease and velocity of managing production-grade applications.
6. **Cluster** — A networked collection of server computers that can, in many respects, be viewed as a single system. The term's meaning can vary depending on the context. However, in the context of the modern data center, a cluster is a group of servers and other resources that function as a single system, sometimes with elements of parallel processing. Many clusters are also distributed systems. See below for related definition.
7. **Container(s)** — Software technology giving a lightweight and portable method for packaging an application that provides isolation from an operating system (OS) and physical infrastructure. Unlike a virtual machine, containers do not include a full OS but instead share the OS of a host. Containers allow an application to be packaged and abstracted to simplify deployment among different platforms. Examples include Docker and Linux Containers (LXC). Containers are often associated with microservices, defined below. Container may also refer to a granular unit of data storage. For example, Amazon S3 (Simple Storage Service) uses the term "bucket" to describe a data container. In certain SDS solutions, the data that makes up virtual disks is stored in logical containers housed on various nodes in a cluster.
8. **Control plane** — Originally a network term, the control plane is generally anything related to the "signaling" of the network. Control plane packets are destined to or locally originated by the router itself. It makes decisions about where traffic is sent

Components of a modern data center



and its functions include system configuration, management, and exchange of routing table information. However, with the rise of software-defined infrastructure, control plane is now a term that extends to server, storage, and security infrastructure. It refers to the programmable set of APIs that govern the configuration, management, and monitoring of the infrastructure.

9. **DRaaS** — Disaster recovery as a service (DRaaS) is the replication and hosting of physical or virtual infrastructure by a dedicated provider to enable failover in the event of a man-made or natural catastrophe. DRaaS is one of the primary drivers of cloud computing and often the primary motivation behind adopting a hybrid or multi-cloud architecture.
10. **Data layer** — A term with a number of definitions (including application as a marketing buzzword). However, in the context of a modern data center, the data layer is a data structure that holds all data that needs to be processed and passed within a digital context (as in a website, for example) to other linked applications.
11. **Data plane** — Also known as the Forwarding Plane, it forwards traffic to the next hop along the path to a selected destination network according to the control plane logic (the learned path that the data on the data plane takes). Also originally a networking term, the data plane consists of data (packets) that is sent through the router itself on its way to its next destination. Increasingly, data plane refers to the infrastructure that stores, manages, protects and transmits data for all applications.
12. **Distributed system** — A cluster of autonomous computers networked together to create a single unified system. In a distributed system, networked computers coordinate activities and share resources to support a common workload. The goals of distributed systems are to maximize performance and scalability, ensure fault tolerance, and enable resource availability. Examples of distributed systems include Amazon Dynamo, Google MapReduce, Apache Hadoop, and the Hedvig Distributed Storage Platform.
13. **Docker** — An open-source project that automates the deployment of applications within software containers. Docker containers, like other containers, wrap a piece of software in a complete file system containing everything needed to run: code, runtime, system tools, system libraries, etc. Docker is often synonymous with containers and many use the term interchangeably. It's important to note that Docker is both an open source set of tools as well as a company, which supports the open source technology as well as selling its own proprietary software.
14. **Flash** — A storage device that stores persistent data on nonvolatile solid-state memory chips. Unlike spinning electromechanical disks (i.e., hard disk drives), Flash

Components of a modern data center



drives have no moving parts. Flash also typically produces no noise, stores and accesses data more quickly, has less latency and is more reliable and durable than spinning media. Since the technology is more advanced, the cost of flash is usually higher although the cost of Flash is decreasing as production methods are refined, improved, and scaled.

15. **Hybrid cloud** — A cloud computing environment in which private cloud resources (e.g., on-premise data center) are managed and utilized together with resources provisioned in a public cloud (e.g., Amazon Web Services). Typically applications and data are exchanged across this private/public cloud boundary, creating a single logical infrastructure or set of services.
16. **Hyperconverged** — An architecture that combines software-defined compute and software-defined storage together on a commodity server to form a simplified, scale-out data center building block. The “hyper” in hyperconvergence comes from hypervisor – the server virtualization component of the solution.
17. **Hyperscale** — An architecture where software-defined compute and software-defined storage scale independently of one another. A hyperscale architecture is well suited for achieving elasticity because it decouples storage capacity from compute capacity. Hyperscale architectures underpin web giants including Google and Amazon and is increasingly being adopted by other enterprises as a means to efficiently scale or contract an environment over time.
18. **IaaS** — Infrastructure as a Service (IaaS) is a form of cloud computing in which virtualized computing resources are provided over the Internet. It is considered one of the three main categories of cloud computing, along with Software as a Service (SaaS) and Platform as a Service (PaaS). These computing resources are typically billed on a utility computing basis (pay as you go, pay as much as you use). It is a service model that delivers virtualized infrastructure on an outsourced basis to support organizations. Among its benefits are automated administrative costs, self-serviceability, dynamic scaling, flexibility, and platform virtualization.
19. **Kubernetes** — Another popular open-source system for automating the deployment, scaling and management of containerized applications. Originally designed by Google, it was donated to the Cloud Native Computing Foundation. Kubernetes defines a set of building blocks that collectively provide the mechanisms for deploying, maintaining, and scaling applications. Kubernetes is also designed to be loosely coupled and extensible so it can accommodate a wide range of workloads.

Components of a modern data center



20. **Mesos** — Formally known as Apache Mesos, it is an open-source software to manage computer clusters that was originally developed at the University of California, Berkeley. Apache Mesos abstracts CPU, memory, storage, and other compute resources away from machines (be they physical or virtual) and allows for fault-tolerant and elastic distributed systems to be built and run easily and effectively. It sits between the application layer and the operating system, and eases deploying and managing applications in large-scaled clustered environments. It was originally designed to manage large-scale Hadoop environments but has since been extended to manage other types of clusters.
21. **Microservices** — A method of developing software applications as a suite of independently deployable, small, modular services in which each service runs a unique process and communicates through a well-defined, lightweight mechanism. The idea behind microservices is that some applications are easier to build and maintain when they're broken down into smaller, composable elements. When the different components of an application are separated, they can be developed concurrently while another advantage of microservices is resilience. Components can be spread across multiple servers or data centers; if a component dies one only needs to spin up another component elsewhere and the overall application continues to function. Microservices are similar but differ from a services-oriented architecture (SOA) in that each service can be independently operated and deployed. The rise in microservices' popularity is tied to the emergence of containers as a way of packaging and running the code.
22. **Multi-cloud** — The use of two or more public cloud computing services providers by a single organization. Hybrid clouds can be multi-clouds if two or more public clouds are used in conjunction with a private cloud. Multi-cloud environments minimize the risk of data loss or downtime due to a failure occurring in hardware, infrastructure, or software at the public cloud provider. A multi-cloud approach can also be used as part of a pricing strategy to keep costs under control and prevent vendor lock-in to one cloud provider. This method can increase flexibility by mixing and matching best-in-class technologies, solutions, and services from different public cloud providers.
23. **Multi-tier** — A type of application that is developed and distributed among more than one layer and logically separates the different application-specific, operational layers. The number of layers varies by business and application requirements, but three-tier is the most commonly used. The three tiers are: presentation (user interface); application (core business or application logic); and data (where the data is managed). Also known as N-tier application architecture, it provides a model in which developers can create flexible and reusable applications. Multi-tier can also

Components of a modern data center



refer to the data storage. In this case, multi-tier represents a single storage platform that spans multiple, traditional tiers of storage. In this case, each tier is defined by the specific performance and availability needs of applications. Tier 0 or 1 is often the highest performance, highest-availability applications (often serviced by all-flash arrays), whereas tier 3 or 4 is often the lowest performance, lowest availability applications (often serviced by archive or cold archive storage).

24. **Multi-workload** — A distributed computing environment in which different workloads (all of which may have differing characteristics) are equally supported, managed, and executed. Just as there are different types of bicycles for different uses, different computing workloads place different demands on the underlying infrastructure, whether it be a desktop workload or an SAP system workload. Different workloads have different characteristics in terms of computing capacity, network needs, data storage, backup services, security needs, network bandwidth needs, QoS metrics, among other factors. Multi-workload is gaining prominence as companies look to build cloud environments where a single, shared infrastructure supports all the workload or application needs. This is in sharp contrast to traditional, siloed environments where workloads often have bespoke infrastructures. In a multi-workload cloud, software-defined technologies and application-specific policies enable a single infrastructure to meet the needs of a diverse set of applications.
25. **Multi-site replication** — The ability to natively replicate data among different sites to ensure locality and availability. A site can represent a private cloud data center, public cloud data center, remote office, or branch office. Multi-site replication prevents any one site from being a single point of failure.
26. **Node** — A widely used term in information technology that may refer to devices or data points on a larger network. Devices such as a personal computer, cell phone, or printer are considered nodes. Within the context of the Internet, a node is anything that has an IP address. When used in the context of modern data centers, it can refer to a server computer. Often the different computers that make up a cluster or distributed system are referred to as nodes.
27. **NVMe** — Non-volatile memory express (NVMe or NVM Express) is a specification that allows a solid-state drive (SSD) to make effective use of a high-speed Peripheral Component Interconnect Express (PCIe) bus in a computer. The idea behind NVMe is increased and more efficient performance and interoperability in a broad range of enterprise and client systems. Principal benefits include reduced latency, increased Input/Output operations per second (IOPS), and lower power consumption.

Components of a modern data center



28. **OpenStack** — A free and open-source software platform for cloud computing, deployed mostly to underpin private or public cloud Infrastructure-as-a Service (IaaS). The software platform consists of interrelated components that control diverse, multi-vendor hardware pools of processing, storage and networking resources throughout a data center. Users manage it via a web-based dashboard, command-line tools, or through a RESTful API.
29. **Orchestration layer** — Consisting of programming that manages the interconnections and interactions among cloud-based and on-premises components. In this layer, tasks are combined into workflows so the provisioning and management of various IT components and associated resources can be automated with several tools or managers such as Puppet, Chef, Ansible, Salt, Jenkins, among others. Traditional data center infrastructure management tools like VMware vSphere, Microsoft Hyper-V, and OpenStack are also considered part of the orchestration layer.
30. **PCIe** — An abbreviation of Peripheral Component Interconnect Express, it is a serial expansion bus standard for connecting a computer to one or more peripheral devices. With PCIe, data center managers can take advantage of high-speed networking across server backplanes, and connect to Gigabit Ethernet, RAID and Infiniband networking technologies outside of the server rack. It offers lower latency and higher data transfer rates than parallel buses such as PCI and PCI-X.
31. **Private cloud** — A type of cloud computing designed to deliver similar advantages of the public cloud (including scalability, flexibility and self-service) but is dedicated to a single organization. A large multinational enterprise, for example, may establish its own private cloud that mimics the characteristics of those offered by public cloud providers, which deliver services to multiple companies concurrently. Private clouds can be deployed in wholly-owned data center facilities or hosted in outsourced facilities. Thus, private cloud does not necessarily mean on-premises, although most are deployed as such.
32. **PaaS** — An application platform, Platform as a Service (PaaS) is a category of cloud computing services providing a platform that allows customers to develop, run, and manage applications without the complexity of building and maintaining the infrastructure typically associated with developing and launching an app. There are different types of PaaS, including public, private and hybrid. Originally intended for applications on public cloud services, PaaS has expanded to include private and hybrid options.
33. **Scale-out** — Used to describe a type of architecture that may apply to storage, networking, or applications. In general, scale-out refers to adding more components

Components of a modern data center



in parallel to spread out a workload. In most cases, scale-out adds more controllers with each node added to the scale-out system. This enables higher levels of scalability, performance, and resiliency. This contrasts with scale-up, which refers to adding more capacity to the system without adding more controllers. Most scale-up systems are a dual-controller model and represent scaling, performance, and resiliency limits based on this constraint.

34. **Software-defined** — An increasingly widely used term in storage, networking and other information technology applications, it generally refers to a new class of products where the software is deployed on commodity hardware to provide a capability. Traditional, or hardware-defined, systems have a tight coupling of software to proprietary hardware components or designs. Software-defined abstracts physical resources, automates actions, and enables the programming of infrastructure to meet specific application and workload needs.
35. **Stretch(ed) clusters** — Deployment model in which two or more virtualization host servers are part of the same logical cluster but which are in separate geographical locations. In stretched clusters, the servers act as a single system to provide high availability and load balancing despite not being in the same facility. They have the advantage of enabling easier migration of virtual machines from one physical location to another while maintaining network connections with the other servers in the cluster.
36. **Tier(s)** — Can refer to multi-tiered architecture (as defined above) but in context of storage, determines the priority and importance of an organization's data. Tier 1 data, for example, is the data that an organization or computing environment must have immediate access to for the most mission-critical applications. Tier 2 data most often includes business-critical application data and the type of storage will depend on performance and availability requirements. Tier 3 data typically refers to backups, whereas archived data is typically Tier 4 or higher.
37. **UDP** — Conceived by Hedvig, the Universal Data Plane (UDP) is a single, programmable data management layer spanning workloads, clouds and tiers that capitalizes on the distributed systems approach being adopted by many organizations. It's a virtualized abstraction layer that enables any workload to store and protect its data across any location. It also dramatically simplifies operations by plugging into modern orchestration and automation frameworks like Docker, Kubernetes, Mesos, Microsoft, OpenStack, and VMware.