DEFINITIONS AND ACRONYMS THAT MAKE UP CLOUD NOMENCLATURE

Cloud can bring flexibility, capability, and predictability to your customers' organization. You need Cloud technology to secure your data, networks, and reduce the burden on your IT team. Here are a few you should know, and a lot more for good measure:

CLOUD TERMS

Active Directory (AD). Centralized database that all devices on a domain go to when needing authentication. Pretty much every single action on the network requires the device to request permission from the AD before it is allowed to read, write, send, etc.

Access Control List (ACL). Rules are used to govern access to digital environments. Organizations use two types of ACLs; filesystem ACLs and networking ACLs; to control traffic flow, grant or deny permissions, and monitor activity in and out of certain systems.

Amazon Web Services (AWS). Third-party Provider of public Cloud computing services, the platform offers over 175 Cloud-native services, including big data tools, database solutions, and Internet of Things (IoT) applications.

Application Modernization. The process of updating legacy software with new capabilities and features to create incremental business value. Organizations typically modernize outdated applications through re-platforming, refactoring, or rehosting efforts, which may involve significant changes to the core architecture.

Application Programming Interface (API). Enable disparate applications to communicate directly with one another according to predefined rules. Types of APIs include Web APIs, Composite APIs, Internal APIs, Open APIs, and Partner APIs. The two most referenced APIs are REST and SOAP APIs, both of which are Web APIs. Organizations use APIs to extend functionality to other systems and gain access to capabilities that fulfill unmet business requirements.

Application Refactoring. Make significant changes to the configuration and source code of an existing application to align with business needs. Refactoring allows organizations to add new features like enhanced performance capabilities and reduce costs. Refactoring doesn't change an application's external behavior; it is a more complex process than re-platforming/rehosting.

Artificial Intelligence (AI). A discipline within computer science that focuses on creating smart machines that can execute tasks that humans typically perform. Advances in Cloud computing have made AI capabilities more accessible. Organizations of all sizes can now build and deploy powerful AI programs that automate manual activities, reduce costs, and create new value.

Auto-scaling. A Cloud computing function where resources are allocated automatically to applications based on real-time demand. Cloud computing enables more organizations to take advantage of auto-scaling and optimize resource consumption across multiple Cloud services.

Availability Zones (AZs). Isolated, logical data centers are available to Amazon Web Services (AWS) customers. AZs come with independent cooling, networking, and power; enabling users to achieve redundancy for critical applications. By hosting applications across multiple AZs, organizations buffer their customers from performance issues and eliminate any single points of failure that would otherwise persist.

Backup as a Service (BaaS). Backing up data to a remote, Cloud-based server.

Bare Metal as a Service (BMaaS). Service model where a vendor deploys dedicated physical IT infrastructure or "bare metal" to the customer's data center with the same on-demand scalability, convenience, and agility of a Cloud service. Unlike an outsourced colocation data center, BMaaS provides a full stack, hosted, on-demand solution, where the procurement of hardware, software, deployment, support, optimization, and lifecycle management services, are all handled by one vendor under one contract.

Big Data. Massive amounts of information are created today with ever-increasing velocity. Organizations collect, store, and process Big Data through advanced data management techniques, mainly through the Cloud. With Big Data analytics, organizations can extract valuable insights from structured, semi-structured, and unstructured datasets.

Blast Extreme. A display protocol built by VMware to deliver an immersive, feature-rich experience for end users across devices, locations, media, and network connections.

Blockchain. An open, immutable, and distributed digital record of information that promotes accountability and transparency amongst all parties. Although originally designed to support digital currencies, organizations use blockchain technology today for numerous applications.

Block storage. A block, sometimes called a physical record, is a sequence of bytes or bits, usually containing some whole number of records, with a maximum length/block size. Data, thus structured, is said to be blocked.

Bring Your Own Device (BYOD). Companies are embracing the idea of employees using their own personal devices in the workplace. With the Cloud, employees can virtually access their workplace desktop from the comfort of their device, using an individual app.

Business Continuity. Activities performed to ensure that critical business functions will be available to customers, suppliers, regulators, and other entities that must have access to those functions. Business continuity is not implemented at the time of a disaster, it refers to activities performed daily to maintain service, consistency, and recoverability.

Business Intelligence. The data and analytics are used to discover new insights, improve decision-making, and create enterprise value. Modern BI practices rely on big data analytics, modern data infrastructure technologies, advanced visualizations, and nuanced reporting to gather and process information quickly at scale, with the hopes of identifying new growth opportunities.

Central Processing Unit (CPU). The electronic circuitry that executes instructions that comprise a computer program, also called a central processor.

Citrix XenDesktop. A Citrix System's tool that offers virtual desktop delivery, allowing users to make various generations of Windows applications available to any device, anywhere. Virtualized desktop resources bring the power of multiple applications and utilities to remote devices for remote work.

Citrix XenApp. An application virtualization service that provides on-demand Windows-based applications to devices and works in the context of the Citrix XenDesktop desktop virtualization service.

Cloud. A vast network of remote servers operating as a single ecosystem. Servers on the Cloud store/manage data, run applications and deliver content/services via Internet-capable devices.

Cloud Application. A web-based program that relies on the power of Cloud computing and related capabilities for data storage, logic processing, etc. Processing for Cloud applications is typically executed by local devices and Cloud computing solutions. Users interact with Cloud applications through Internet browsers.

Cloud Automation. The practice of automating Cloud infrastructure management processes in line with IT resource demand. Cloud automation is commonly used by DevOps, security, and application development teams to free up engineering capacity for more complex aspects of Cloud-native operations.

Cloud Computing. Computing services (data storage, networking, analytics, server hosting, etc.) delivered over the Internet, offer many advantages over on-premises computing, including lower operating costs, flexible resource allocation, and improved scalability.

Cloud Migration. The process of moving all or part of an organization's on-premises IT infrastructure (databases, applications, and other components) to the Cloud. Migrations enable organizations to fulfill ever-evolving business requirements and take advantage of Cloud computing capabilities. These can be highly complex endeavors requiring significant planning and expertise to execute successfully.

Cloud-Native. Software services, business applications, and IT systems are explicitly designed to run in dynamic Cloud environments. Where on-premises applications may need to be modernized for the Cloud, Cloud-native applications work immediately. They are also generally more agile and scalable than legacy technologies.

Cloud Orchestrator. Software that manages the interconnections and interactions among Cloud-based and on-premises business units. These products use workflows to connect various automated processes and associated resources.

Cloud Provider. A company that provides the Cloud-based platform, infrastructure, application, or storage services to other organizations and/or individuals, usually for a fee.

Cloud Provisioning. What Cloud Providers (e.g. AWS and Microsoft Azure) do to deliver Cloud resources and services to customers on an as-needed basis. Cloud provisioning is central to the on-demand nature of the Cloud computing model and represents a key advantage over traditional, more limited approaches to compute resource management.

Cloud Service Provider. A Provider who offers Cloud computing services, networking, and infrastructure over the web. Organizations use third-party Cloud Service Providers to outsource much of the effort associated with maintaining on-premises IT. Today's leading Cloud service Providers offer cost-efficient/scalable data storage and analytical tools through the Internet.

Cloud Storage. The data storage model in which an organization's data is maintained by a Cloud Provider, who's responsible for storing, maintaining, and serving information from a remote repository. Cloud storage frees IT teams from having to set up or manage data infrastructure on-premises, paving the way for enterprises to use modern data architectures (e.g., data lakes) and advanced analytics.

Cluster. A group of computers or hosts that collectively work together to support a specific application or middleware software. In a cluster, individual computing devices are called "nodes," and all nodes work on the same tasks. Clusters are commonly seen in High-Performance Computing (HPC) applications that require significant computing power.

Colocation (Colo). A data center facility in which a business can rent space for servers and other computing hardware. Typically, a Colo provides the building, cooling, power, bandwidth, and physical security while the customer provides servers and storage. Space in the facility is often leased by the rack, cabinet, cage, or room.

Common Internet File System (CIFS). Remote file system access protocol that allows groups of users to work together and share documents via the internet/corporate intranets. CIFS is an open, cross-platform technology based on the native file-sharing protocols built into Microsoft Windows and other operating systems.

Compute. Activities that require processing resources beyond what is available through internal memory. Organizations must be aware of their existing computing capacity and the power they need to support critical business activities.

Containers. Software units that enable organizations to run their applications quickly and reliably in different computing environments. Containers group all runtime elements together (including code, system libraries, and settings) into lightweight and secure packages. Organizations use containers to decouple applications from their native environments, so they can be deployed easily and consistently anywhere.

Content Delivery Network (CDN). Group of geographically distributed servers collaborating to deliver content over the web. CDNs enable the rapid transfer of assets (HTML pages, images, stylesheets, videos) to end users. CDNs deliver most of the world's web content.

Continuous Integration/Continuous Development (CI/CD). A set of practices used by DevOps teams to automate activities related to application building, testing, and deployment. Through CI/CD, DevOps teams can constantly innovate, deliver new features to the market, and iteratively deploy updates. CI/CD is considered a best practice in modern Cloud computing.

CPU Cores. A multi-core processor is a microprocessor on a single integrated circuit with two or more processing units, called cores, each of which reads/executes program instructions.

CPU Socket/Slot. Contains one or more mechanical components providing mechanical and electrical connections between a microprocessor and a printed circuit board (PCB). This allows for placing and replacing the Central Processing Unit (CPU) without soldering.

Customer Relationship Management (CRM). A system for managing relationships with customers, keeping track of interactions, data, and notes about customers or potential customers. Data is stored in a central database and accessible to multiple people within an organization. CRM helps streamline sales, marketing efforts, customer service, accounting, and management for growing companies. Multiple people can access and edit the information.

Database Administrator (DBA). DBAs use specialized software to store and organize data. Their role may include capacity planning, installation, configuration, database design, migration, performance monitoring, security, troubleshooting, as well as backup and data recovery.

Data Center. A centralized repository, either physical or virtual, for the storage, management, and dissemination of data and information organized around a particular body of knowledge or about a particular business.

Datacenter N+1 architecture. If N equals the amount of capacity needed to run the facility, N+1 indicates an additional component added to support a single failure or required maintenance on a component.

Desktop as a Service (DaaS). A form of virtual desktop infrastructure (VDI) in which the VDI is outsourced and managed by a third party. Also called hosted desktop services, DaaS is frequently delivered as a Cloud service, along with the apps required for the virtual desktop.

Disaster Recovery. The process, policies, procedures, and preparations implemented to protect IT technology, infrastructure, and data to ensure recovery of data and continuation of mission-critical network services vital to an organization's business continuity in the event of a natural or human-caused disaster.

Distributed Resource Scheduling (DRS). A type of VMware, vSphere cluster provides load balancing by migrating VMs from a heavily loaded ESXi host to another host that has enough computing resources, all while the VMs are still running.

Domain Name System (DNS). A hierarchical and distributed naming system for computers, services, and other resources on the internet or other Internet Protocol networks. It associates various information with domain names assigned to each of the associated entities.

Domain Controller. A server that takes care of managing and hosting the Active Directory (AD) database. If a server is a Domain Controller, it has Active Directory running on it.

Dynamic Host Configuration Protocol (DHCP). A client/server protocol that automatically provides an Internet Protocol (IP) host with its IP address and other related configuration information such as the subnet mask and default gateway. RFCs 2131 and 2132 define DHCP as an Internet Engineering Task Force (IETF) standard based on Bootstrap Protocol (BOOTP), a protocol with which DHCP shares many implementation details. DHCP allows hosts to obtain required TCP/IP configuration information from a DHCP server.

Elastic Computing. The system's ability to scale processing, memory, and storage capacity with changes in demand. Organizations that implement elastic computing don't have to worry about capacity planning or peak usage scenarios. Instead, they can trust their IT infrastructure to acquire computing resources dynamically.

Encryption. The encoding of your information and data to keep it secure.

Endpoint. A remote computing device or node that communicates and receives information across a network. Can be data terminals, host computers, modems, bridges, and other commonly used infrastructure. Endpoints are particularly valuable in IoT and smart applications that depend on "Edge" devices to gather information from the surrounding environment which can then be used to support new applications, offerings, or business models.

Enterprise Resource Planning (ERP). Type of software that organizations use to manage day-to-day business activities such as accounting, procurement, project management, risk management/compliance, and supply chain operations.

ESXi (formerly ESX). Enterprise-class type-1 hypervisor developed by VMware for deploying and serving virtual computers. As a type-1 hypervisor, ESXi is not a software application that is installed on an operating system (OS); instead, it includes and integrates vital OS components. External Cloud. A Cloud service provided by a third party; can be public or private.

Extract, Transform, and Load (ETL). The process of ingesting and integrating data from diverse sources into a single, consolidated data store. ETL is particularly important for organizations gathering information from remote endpoints and edge devices that may not share the same data management protocols. For organizations that aim to leverage big data analytics and AI/ML, ETL is a crucial step in the early stages of the data pipeline.

Failover. A backup operational mode in which the functions of a system component (such as a processor, server, network, or database, for example) are assumed by secondary system components when the primary component becomes unavailable through either failure or scheduled downtime.

Fibre Channel (FC). A high-speed data transfer protocol providing in-order, lossless delivery of raw block data. Primarily used to connect computer data storage to servers in storage area networks in commercial data centers.

Fibre Channel over Ethernet (FCOE). Computer network technology that encapsulates Fibre Channel frames over Ethernet networks. This allows Fibre Channel to use 10 Gigabit Ethernet networks while preserving the Fibre Channel protocol.

File Storage (File-Level/File-Based Storage). Stores data in a hierarchical structure (files and folders) presented in the same format to both the system storing it and the system retrieving it

File Transfer Protocol (FTP). A standard communication protocol is used for the transfer of computer files from a server to a client on a computer network. Built on a client–server model architecture using separate control and data connections between the client and the server.

Firewall. Set of related programs, located at a network gateway server, which protects the resources of a private network from other network users.

Global Server Load Balancing (GSLB). The practice of distributing internet traffic amongst a large number of connected servers dispersed around the world. The benefits of GSLB include increased reliability and latency reduction.

Google Cloud Platform (GCP). A third-party public Cloud computing services Provider, GCP offers solutions for data management, infrastructure modernization, smart analytics, etc.

Graphics Processing Unit (GPU). Computer component that excels in rendering graphical content. allowing a system to display visually intense videos, images, and animations from software or video games.

Hadoop (from Apache). An open-source framework that enables organizations to store massive volumes of data in an efficient manner, which also facilitates clustering so engineering teams can quickly analyze large datasets in parallel. Includes four modules: Hadoop Distributed File System (HDFS); Yet Another Resource Negotiator (YARN); MapReduce; and Hadoop Common.

High Availability. Refers to the quality of an application or infrastructure to continue performing despite disruptions. High Availability systems use redundant hardware and software to minimize service interruptions and mitigate single points of failure. When failures do occur, highly available infrastructure relies on failover processes and backups to maintain operations.

Hosted Application. Software that runs on third-party infrastructure rather than on-premises and can be accessed from anywhere in the world through the Internet. More organizations are using hosted applications to minimize the complexities and costs of maintaining on-premises infrastructure.

Hosted Application. Software programs that are accessible over the Internet and run on a server (SaaS).

Hybrid Cloud. A computing environment that uses a combination of private and public Cloud services and/or on-premises infrastructure. Organizations use the hybrid Cloud approach to optimize IT architecture around digital transformation goals. E.g. a company might use a public Cloud Provider for its on-demand Cloud resources, a private Cloud for security purposes, and on-premises infrastructure for compliance reasons.

Hypervisor. A host operating systems program, also called a Virtual Machine Monitor (VMM), that controls and manages compute tasks on the host hardware in the Cloud, allocating resources to multiple operating systems (OS's) that share the host hardware to ensure uninterrupted service for all OS's residing on the host hardware.

Hyperscaler. Large Cloud service Providers, that offer services like computing and storage at enterprise scale. While there is no universal standard for what should be classified as a hyperscaler, major Cloud Providers such as Amazon Web Services, Google Cloud, Microsoft Azure, IBM Cloud, and Alibaba Cloud fit the description.

Hypertext Transfer Protocol Secure (HTTPS). An extension of the Hypertext Transfer Protocol, it uses encryption for secure communication over a computer network and is widely used on the Internet. In HTTPS, the communication protocol is encrypted using Transport Layer Security (formerly, Secure Sockets Layer).

Independent Computer Architecture (ICA). A proprietary protocol for an application server system, designed by Citrix. The protocol lays down a specification for passing data between servers and clients but is not bound to one platform. Citrix's ICA is an alternative to Microsoft's Remote Desktop Protocol (RDP).

Infrastructure-as-a-Service (IaaS). One of the primary types of Cloud services provides users with instant computing, storage, and other IT infrastructure delivered through the Internet. IaaS solutions typically scale with demand, allowing organizations to pay only for what they use, minimizing the complexity of having to purchase and manage on-premises infrastructure.

Internet Protocol Version 4 (IPV4). The fourth version of the Internet Protocol is one of the core protocols of standards-based internetworking methods on the internet and other packet-switched networks.

Internet Protocol Version 6 (IPV6). The most recent version of the Internet Protocol provides an identification and location system for computers on networks and routes traffic across the Internet.

Internet Small Computer Systems Interface (iSCSI) Storage. An Internet Protocol-based storage networking standard for linking data storage facilities. It provides block-level access to storage devices by carrying SCSI commands over a TCP/IP network.

JavaScript Object Notation (JSON). A data interchange format that makes it easy for companies to store and transport data across the web in a way that both humans and machines can understand. JSON represents data in two ways; through key-value pairs; and as arrays representing ordered collections of values. Because of JSON's popularity, the format is a common output for APIs and data that gets sent from a server to a webpage.

Kubernetes. An open-source platform from Google that organizations use to manage containerized workloads and services. In addition to being portable and extensible, it comes with support, tools, and services to help developers run their production workloads at scale.

Latency. The time it takes for a packet of data to get from one designated point to another.

Linux. Open-source and community-developed Operating System (OS) for computers, servers, mainframes, mobile devices, and embedded devices. It's supported on most major computer platforms, including x86, ARM, and SPARC, making it one of the most widely supported OS.

Load Balancing. Process of spreading network traffic over multiple servers to ensure no one server is entirely responsible for supporting an application – allowing organizations to distribute processing resources (hardware and/or software) as needed to improve the performance and responsiveness of modern applications. Load balancing techniques include Round-robin, Least Connection, Resource Based, and Weighted Response Time.

Machine Learning (ML). A branch of Artificial Intelligence (AI) is concerned with building smart computer algorithms that improve over time. Organizations use ML to identify patterns in massive datasets and use those insights to enhance performance. Machine learning is responsible for many software services today, including recommendation engines, social media feeds, and voice assistants.

Managed Service Provider (MSP). A third-party company that offers ongoing services to help organizations maintain their IT infrastructure. MSPs generally offer network, security, and application support services through an existing data center or third-party IaaS Provider.

Management and Governance. Implementing adequate protections and oversight for IT infrastructure, management, and governance allows organizations to monitor the integrity of their applications, perform audits, analyze resource consumption, manage costs, etc.

Microservices. Software development method that aims to compartmentalize application functions so they can deploy, run, and scale independently. Unlike monolithic applications, microservices are loosely coupled and flexible when it comes to implementing updates or fixing errors.

Microsoft Azure. A public Cloud computing platform commonly referred to as Azure. Their service enables organizations to build, test, launch, and manage modern applications hosted in Microsoft-managed data centers.

Middleware. Software that sits between various applications and the operating system in use. Its purpose is to allow for seamless communication, functionality, and data management across diverse systems without detracting from the user experience.

MultiCloud. For circumstances in which an organization uses more than one Cloud vendor for the same type of Cloud deployment. E.g. a company might use one public Cloud service for on-demand computing needs but another public Cloud Provider for a unique application that fulfills a specific business need. Many organizations implement multi-Cloud deployments to gain redundancies and avoid vendor lock-in.

Multi-Tenant Architecture (Multitenancy). A software architecture commonly used in Cloud computing to deploy several single instances of software from one physical server. Through multitenancy, organizations can securely and dynamically serve multiple customers on one server via independent instances.

MS Hyper-V. Microsoft's hardware virtualization product, lets you create and run a software version of a computer, called a virtual machine. Each virtual machine acts like a complete computer, running an operating system and programs.

MS Remote Desktop Protocol (MS RDP). Proprietary Microsoft protocol, providing users with a graphical interface to connect to another computer over a network connection. The user employs RDP client software while the other computer must run RDP server software.

MS SQL Always On Availability Groups. A high-availability and disaster-recovery solution that provides an enterprise-level alternative to database mirroring. Always On Availability groups maximize the availability of a set of user databases for an enterprise.

MS SQL Database Mirroring. A solution for increasing the availability of a SQL Server database. Mirroring is implemented on a per-database basis and works only with databases that use the full recovery model.

MS SQL Server. A Microsoft relational database management system. As a database server, it's a software product with the primary function of storing and retrieving data as requested by other software applications—running on the same computer or another across a network.

MS SQL Server Clustering. A collection of two or more physical servers (nodes), connected via a LAN, each of which hosts a SQL server instance and has the same access to shared storage. Clustering SQL servers provide high availability and protection from disasters whenever a server hosting the SQL Server instance fails.

MS SQL Server Log Shipment. Automatically send transaction log backups from a primary database on a primary server instance to one or more secondary databases on separate secondary server instances. The transaction log backups are applied to each of the secondary databases individually. An optional third server instance, known as the monitor server, records the history and status of backup and restore operations and, optionally, raises alerts if these operations fail to occur as scheduled.

Network Attached Storage (NAS). File-level computer data storage server connected to a computer network providing data access to a heterogeneous group of clients. NAS can refer to both the technology and systems involved, or a specialized device built for such functionality.

Network File System (NFS). A mechanism for storing files on a network. It is a distributed file system allowing users to access files and directories located on remote computers and treat those files and directories as if they were local.

On-Premises. IT infrastructure that organizations manage onsite; includes hardware and software; that could otherwise sit on a remote server farm or be provided by a Cloud services vendor. On-premises IT often requires more effort to maintain from an operational standpoint.

OpenStack. Free Cloud computing platform (by Rackspace), allowing organizations to manage public and private Cloud environments via Infrastructure-as-a-Service (IaaS) deployments. OpenStack is one of the top three open-source projects in the world based on current activity.

Open Virtualization Format (OVF). Open standard for packaging and distributing virtual appliances or, more generally, software to be run in virtual machines.

Operational Expense (OpEx). Money an organization spends on an ongoing, day-to-day basis to run a business or system.

Oracle Cloud. A service offered by Oracle that provides servers, storage, network, applications, and services through a global network of managed data centers. Services are provisioned on demand over the Internet.

Oracle Database (Oracle DB). A multi-model database management system produced and marketed by Oracle. It's a database commonly used for running online transaction processing, data warehousing, and mixed database workloads.

Oracle Data Guard. Provides a comprehensive set of services that create, maintain, manage, and monitor one or more standby databases to enable production Oracle databases to survive disasters and data corruption.

Orchestration. The process of scheduling and integrating automated tasks across disparate systems. Organizations can orchestrate workflows between on-premises and Cloud infrastructure, as well as streamline the execution of complex, interconnected workloads.

Pay-As-You-Go. Billing system in which Cloud services are paid for as they are bought, or as they are used, depending on whether it is subscription-based or consumption-based.

PC over IP (PCoIP). A remote display protocol developed for delivering remote desktops and applications to endpoints. Works by rendering client desktops on a network or Cloud server, the desktop pixels are then compressed, encrypted, and transmitted to the client device.

Physical Server. A hardware server with a motherboard, CPU, memory, and IO controllers. It's considered a bare-metal server because the hardware is used directly by an OS instead of a virtualization platform. A physical server is used to run a single instance of an OS.

Platform-as-a-Service. Cloud-based service model in which Providers, like AWS, offer all hardware and software needed to develop and deploy modern applications. PaaS frees companies from having to manage their own servers and infrastructure, creating additional capacity to focus on delivering modern applications. Examples include Windows OS and Cloud-based environments.

Private Cloud. A Cloud environment/resources used exclusively by a single organization. Private Clouds may be deployed from an on-site data center or hosted by a third-party managed services Provider. The advantage of using a private Cloud is that organizations can customize management, governance, and other operating elements to their unique needs.

Public Cloud. A Cloud environment owned and operated by a third-party Provider, with resources delivered over the Internet to "tenants" that all share hardware, storage, and network devices. The advantage of using a public Cloud Provider is that organizations don't have to purchase or maintain critical IT infrastructure.

Redundant Array of Independent Disks (RAID). Way of storing the same data in different places on multiple hard disks or solid-state drives (SSDs) to protect data in the case of drive failure.

Recovery Point Objective (RPO). Maximum allowable time in which data may be lost due to a disaster or disruption.

Recovery Time Objective (RTO). Maximum desired time allowed between a disaster or disruption and the restoration of normal operations.

Redundancy. A system design in which a component is duplicated, so if it fails, there will be another system to take its place.

Replication. The continuous copying of data changes from one database (publisher) to another database (subscriber). The two databases are generally located on different physical servers, resulting in a load-balancing framework by distributing assorted database queries and providing failover capability. The server for the subscriber database may be configured as a backup in the event of failure of the server for the publisher database.

RVTools. A Windows .NET (4.6.2 or higher) application that uses VMware vSphere Management SDK 7.0U3 and CIS REST API to display information about your virtual environments. Interacting with VirtualCenter 4.x, ESX Server 4.x, VirtualCenter 5.x, ESX Server 5.x, VirtualCenter 6.x, ESX Server 6.x, VirtualCenter 7.0, and ESX server 7.0. RVTools can list information about VMs, CPU, memory, disks, partitions, network, CD drives, USB devices, snapshots, VMware tools, vCenter, Resource pools, Clusters, ESX hosts, HBAs, Nics, Switches, Ports, Distributed Switches, Distributed Ports, Service consoles, VM Kernels, Datastores, multipath info, license info, and health checks.

Scalability. The ability of a Provider, application, network, organization, or process to adjust up and down, instantly, and automatically, to provide computing capacity to meet spikes in demand. For example, scalable applications can support rapid increases in utilization, thus providing high-quality experiences for users regardless of network traffic or resource demand.

Schema. Description of how data is structured in databases or XML files. Database schemas refer to the tables and fields that help organize data. These schemas are often represented as visual diagrams. XML schemas highlight data included in XML files and provide structure for that information.

Security, Identity, and Compliance. Terms related to concerns of securing workloads and applications adequately in the Cloud. Typical priorities include protecting data, managing permissions, safeguarding infrastructure, monitoring Cyber-threats, and maintaining data privacy compliance.

Serverless Computing. An approach in which users rely on third-party Providers to dynamically allocate machine resources from their own services. Organizations pay only for the computing resources they use without having to manage, provision, or maintain any servers themselves.

Shared Security Model. Framework to help Cloud Service Providers (CSPs) and customers determine how to divide up security responsibilities. Different types of Shared Security Models depend on how much operational support the CSP provides. Shared Security Models cover everything from hardware and infrastructure to data, network, and endpoints.

Snapshot. The state of a system at a particular point in time, the term was coined as an analogy to a photograph. It can refer to an actual copy of the state of a system or a capability provided by certain systems.

Software-as-a-Service (SaaS). Software delivery model where vendors license access to their data/applications over a network, usually through the Internet. Generally, SaaS vendors host and maintain their own code, databases, and services. Customers then pay for on-demand access, enabling them to fulfill certain business requirements without building applications in-house or assuming long-term contracts. Examples include Gmail and other Cloud-based email services.

Storage. Digital space leased from third-party Cloud vendors. With Cloud storage, organizations don't have to purchase or maintain storage infrastructure themselves, relying instead on vendors managing capacity, security, etc. – paying only for what they use.

Storage Area Network (SAN)/Storage Network. Computer network which provides access to consolidated, block-level data storage. SANs are primarily used to access data storage devices, such as disk arrays and tape libraries from servers so the devices appear to the operating system as direct-attached storage.

Thin Client. Computer with minimal local storage, designed to deliver programs and applications to the user from a centrally located terminal server. They often still host a minimal operating system and store configuration settings in flash memory. Most data and applications utilized by the end user are not stored directly on the device. Thin clients connect to a central server on boot and rely on that server to provide users with the necessary software applications and data, with the computer acting as a terminal to facilitate connection. The user sees and works within the applications run by the server.

Thick Client (Heavy/Fat). Essentially a traditional computer, it has a storage drive with an operating system and applications installed directly onto it. Thick clients can operate without dependencies on external resources, with everything loaded from the local storage drive inside the system, meaning you can use it even if you unplug the ethernet or turn off Wi-Fi.

vCenter. An advanced server management software that provides a centralized platform for controlling vSphere environments for visibility across hybrid Clouds.

vCloud Director (VCD) VMware. A platform with multi-tenant support for managing Software-Defined DataCenters (SDDC) and providing Infrastructure as a Service (IaaS) to customers. Virtual CPU (vCPU)/Virtual Processor. A physical Central Processing Unit (CPU) is assigned to a Virtual Machine (VM).

Vendor Lock-in. Situation in which a customer is dependent on a vendor for products and services, unable to use another vendor without substantial switching costs.

Video Random Access Memory (vRAM/VRAM). A dual-ported variant of dynamic RAM was once used to store the frame buffer in graphics adapters. Most computers/game consoles don't use this form of memory. Dual-ported VRAM shouldn't be confused with other forms of video memory.

Virtualization. Technology organizations use virtualization to deploy virtual instances of something abstracted from physical hardware. With virtualization, organizations can use IT infrastructure more efficiently by distributing capacity that would otherwise go unused across different tenants or environments. Virtualization is at the base of Cloud technology, enabling individuals with an internet connection to access data anywhere and anytime.

Virtual Desktop Infrastructure (VDI). An untethered virtual workspace that can be accessed from anywhere at any time on a variety of devices for a predictable monthly fee.

Virtual Machine. A digital computing environment that behaves like a physical computer, using software, rather than hardware, to run apps and programs, enabling developers to test applications in isolated environments.

Virtual Private Cloud. An isolated environment with access to on-demand computing resources within a broader public Cloud environment. Organizations use virtual private Clouds to gain privacy and control over their data, applications, and code without sacrificing scalability and the other advantages of using public Cloud platforms.

VMware. A virtualization and Cloud computing software, with server virtualization. A hypervisor is installed on the physical server to allow for multiple Virtual Machines (VMs) to run on the same physical server. Each VM can run its operating system, which means multiple OS can run on one physical server. All VMs on the same physical server share resources, such as networking and RAM.

Windows Server. Group of operating systems for servers developed by Microsoft.

Zero Client. A computer with no local storage that relies on a connection to a local server. When a zero client turns on, it uses firmware to boot off of the network and connect to the server. It then downloads what it needs directly into memory rather than onto a hard drive.