

Cloud Hosting Glossary

Cloud Hosting Glossary

Access Control Entry (ACE): A record in the [ACL](#).

Access Control List (ACL): A list of permissions attached to an object that specifies who or what is allowed to access the object and what operations are allowed to be performed on the object. AppNexus uses ACLs to set firewall rules on each customer VLAN. For information on changing AppNexus ACLs, see [How to Set Firewall Rules](#).

Ad-optimized: A type of load balancing pool. See [HTTP Fast](#).

Anycast: Two or more devices simultaneously announce the same destination IP address. Requests are then routed to the "nearest" device.

Apache Tomcat: a Servlet container that implements the Java Servlet and the JavaServer Pages specifications from Sun Microsystems, and provides a "pure Java" HTTP web server environment for Java code to run in. Tomcat should not be confused with the [Apache web server](#), which is a C implementation of a HTTP web server.

Apache HTTP Server: A web server developed and maintained by an open community of developers under the auspices of the Apache Software Foundation.

API: Application Programming Interface. A set of routines, protocols, and tools for interacting with a software environment. AppNexus APIs are the primary means of interacting with your virtual infrastructure. AppNexus creates its own APIs to make this interaction as efficient as possible.

Availability: Making sure equipment, applications, and data are available and running when you need them. To ensure availability, preparation must be made for routine hardware failures, usage spikes, disaster, etc.

BGP: Border Gateway Protocol. The core routing protocol of the Internet, BGP attempts to send traffic by the shortest path and is used by most Internet service providers to establish routing between one another. BGP works by maintaining a table of IP addresses and makes routing decisions based on path, network policies, and/or rulesets.

CLI: Command-Line Interface. The command-line tools that allow you to control the AppNexus [API](#). See [API Documentation for Cloud Services](#) for details.

CNAME: Canonical Name record. A record in a DNS database that indicates the "true," unique host name of a computer or server. A CNAME is used to direct multiple URLs, such as [www.fakesite.com](#) and [support.fakesite.com](#) to a single IP address, and to redirect users to sites after their URLs change.

Colocation: A colocation company runs the buildings that house their clients' IT equipment. The colocation company is responsible for environmental issues such as cooling, power, and security.

Configuration Type: A type of server with uniform specifications. Configuration types can be grouped by function and revision: eg "weba," "webb," "dba". "Web" servers are generic workhorses with reasonable CPU, memory, and disk. "DB" servers have more memory, disk, and CPU, and are designed for heavier IO-intensive jobs such as databases and file-servers.

Datacenter: A physical location that houses servers and other hardware. At AppNexus, we sometimes refer to our datacenters as a [Nexus](#).

Floating IP: An IP address which has two or more networking devices attached to it. One device acts as a "master" and answers requests sent for this IP address; the other takes over the IP address in case the master fails.

Fully Layer 7: Refers to standard HTTP load balancing pools, which deal with packets in a "smart" way, as opposed to HTTP-Ad Optimized pools, which do not open any packets.

GigE: Gigabit Ethernet. Transmission of Ethernet frames at a rate of a gigabit per second. This is about 667 times as fast as a T-1 connection.

Global Traffic Management: Distribution of traffic between multiple datacenters to reduce latency and ensure availability. Also known as GSLB, or Global Server Load Balancing.

Global Traffic Manager (GTM): The F5 device that AppNexus uses within its datacenters to distribute traffic between multiple datacenters. Also known as a GTM.

GSLB: Global Server Load Balancing. See [Global Traffic Management](#).

HTTP Method: An action to perform on a resource. (Essentially a command.) Some of the methods available in HTTP are GET, POST, and PUT.

HTTP Fast: A streamlined version of the HTTP type load balancing pool optimized for speed under ideal traffic conditions by accelerating certain types of HTTP connections and reducing the number of connections opened to the back-end HTTP servers. Also called ad-optimized.

Hypervisor: A virtualization platform, or the software that allows multiple operating systems to run on a host computer at the same time. AppNexus currently uses the Xen hypervisor.

i18n: Abbreviation for "internationalization" where 18 is the number of letters between the i and the n. Internationalization is the process of designing a

software application so it can be adapted to various languages and regions without engineering changes.

Instance: A single, virtualized, fully functional operating system. Also known as a Guest Operating System, Virtual Machine, Virtual Server, or DomU. Each instance is tied to a specific server.

LAMP: A solution stack of software used to run websites or servers. Refers to Linux, Apache, MySQL, and PHP (or Perl, Python).

Latency: A time delay. Latency is often due to the time it takes data to travel over a network.

Load Balancing: Refers specifically to **local** load balancing, which means pooling identical instances into a pool to create redundancy, increase capacity, or improve response time and distributing traffic among the instances, or nodes, in the pool.

Load Balancer: A physical device that distributes traffic across a series of nodes for a specified external address. See also [Load Balancing](#).

Local Traffic Manager (LTM): The type of local load balancer AppNexus uses within its datacenters. Manufactured by F5, the LTM manages routing of traffic from one external Virtual IP address to multiple local servers.

Method: A specialized command. Also called a "member function."

NAT: Network Address Translation. Allows a single IP address to represent a network.

Network Attached Storage: File-level computer data storage attached directly to a network for fast access.

Nexus: A physical or virtual AppNexus location, where exchange of information/services happens.

NIC: Network interface card. A piece of hardware that allows computers to communicate over a network.

Nodes: The members of a [pool](#). At AppNexus, nodes are most commonly [instances](#).

Pool: A collection nodes (usually instances) that share an IP address for load balancing purposes.

Power Distribution Unit (PDU): A device that distributes electric power. Each rack contains a pair of redundant PDUs.

PTR: Pointer Record. A type of DNS record that points to a [canonical name](#). The most common use is for reverse DNS lookups.

Rack: A physical cabinet that contains a number of servers. Each rack contain a redundant pair of network switches and redundant power distribution units (PDUs) and up to 20 servers.

Redundancy: The duplication of network components (such as hardware, data, or an instance) in order to have a backup or failsafe.

Server: A physical machine. Also known as a [Xen Host](#).

SSL: A method for encrypting data sent over the Internet.

Stateful Inspection: A firewall architecture that works at the network layer. Also referred to as dynamic packet filtering.

TTL: Time to live.

VIP: Virtual IP address.

Xen: A free software [Hypervisor](#).

Virtual Private Network: A virtual network formed over a shared or public network, such as the Internet, that uses encrypted data for security.

Virtual Server: See [Instance](#).

Virtual Machine (VM): See [Instance](#).

VLAN: Virtual Local Area Network. A group of hosts that communicate as if they were attached to the same wire, regardless of their physical location. A VLAN has the same attributes as a physical LAN, but is configured through software rather than hardware.