Progress® MarkLogic®

Progress MarkLogic Server Core Licensing

March 2024

Contents

License Model

Progress® MarkLogic® products are licensed Per Core. This allows for a precise measure of computing power and a consistent licensing metric regardless of deployment environment, whether physical or virtual, on-premises or in the cloud. Customers must license all the compute power accessible by the instance of MarkLogic Server.

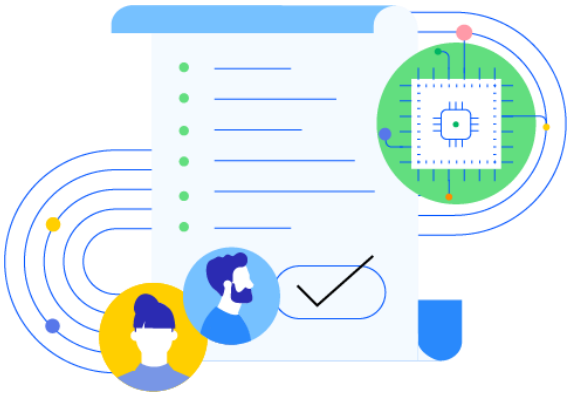

Licensing Physical Servers

When licensing for a Physical Server running a single Operating System Environment, license all the physical cores on the server.

For example: HP DL380 with two Intel 4310 processors requires 24 cores of license, 12 per processor.

If a processor is installed but is not accessible to the operating system (i.e., turned off in BIOS), the cores on that processor do not need to be licensed.

Licensing Virtual Machines

When deploying MarkLogic Server in a private cloud, you must license the maximum compute power available to that virtual machine given the hypervisor configuration. The examples below are just two possible examples, as many hypervisors offer other options.

Licensing for Performance

License each virtual core with a Core License. This will allow each virtual core of the virtual machine to run on a dedicated physical core, if available. This is the default configuration for most hypervisors and facilitates maximum performance and license compliance.

Licensing for License Efficiency

License each set of two virtual cores with a single core license (8 virtual cores = 4 licenses, etc.). This will require configuring the virtual machine to place pairs of virtual cores onto a single physical core.

In VMWare ESXi, use the setting `numa.vcpu.preferHT=TRUE`.

With OpenStack, use `hw_cpu_thread_policy=require`.

RHEV does not have a way to set hard limits on CPU utilization at this time.

Hyper-V NUMA Affinity should be turned on when Hyperthreading is on to promote license efficiency and to help prevent NUMA spanning which will slow performance.

When oversubscribing resources on a host (multiple guests running multiple MarkLogic Server instances per host), it is best practice to associate a single license key with the host so all the cores on the host are licensed.

The cores allocation for the guests can then be assigned, dynamically, according to the workload. This allocation does not need to be implemented through license key reassignment. Since all the cores on the host are allocated, the total workload of the guests running on that host cannot exceed the license boundaries.

Licensing in the Public Cloud

Below are guidelines for applying the core license definitions for Universal Licenses based on the publicly available documentation.

When assigning licenses to Cloud instances, it is best practice to assign adequate cores for the maximum number of cores that instance can consume, since control of the underlying hardware is typically beyond the control or monitoring capabilities of the customer.

Amazon Web Services

https://aws.amazon.com/ec2/instance-types/

- One Core License per two AWS vCPU.
- PVM instances with Hyperthreading off require one license core for each vCPU.
- Instance types with only a single vCPU require one core license.

Azure

https://azure.microsoft.com/en-us/pricing/details/virtual-machines/linux/

- One Core License per each Instance Core (non-Hyperthreaded Instances).
- One Core License per two vCPUs (Hyperthreaded Instances).
- Instance types with only a single instance core require one core license.

Google Compute Engine

https://cloud.google.com/compute/docs/machine-types

- One Core License per two GCE vCPU.
- Instance types with only a single vCPU require One Core License.

Licensing Containers/Docker

By default, a container has no resource constraints and can use as much of a given resource as the host's kernel scheduler will allow. As a result, unless configured with limits, each container running MarkLogic software must be licensed for all the physical cores on the host.

When applying resource limits for containers, use fractional CPUs, CPU quotas or CPU "sets" to limit how much CPU is available to MarkLogic Server. CPU "shares" will still allow a container to utilize 100% of the resources of the host as long as there are no competing workloads.

Reference 1: https://docs.docker.com/engine/admin/resource_constraints/

Reference 2: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/#resource-requests-and-limits-of-pod-and-container

It is best practice to associate a single license key with the host so all the cores on the host are licensed. In Docker containers, the allocation of cores for each container can adjust dynamically based on the workload. This helps prevent the total cores used across the host from exceeding the total licensed cores available for all the guests.