



Hypervisor Scheduling Types

examlabpractice.com



<https://t.me/learningnets>



What is Hypervisor Scheduling Types?


- Scheduler types (also known as modes) determine how the Hyper-V hypervisor allocates and manages work across guest virtual processors.
- A Hyper-V host administrator can select hypervisor scheduler types that are best suited for the guest virtual machines (VMs) and configure the VMs to take advantage of the scheduling logic.

Understanding how Hyper-V virtualizes processors

Before considering hypervisor scheduler types, it's also helpful to understand the Hyper-V architecture.


- Hyper-V creates and manages virtual machine partitions, across which compute resources are allocated and shared, under control of the hypervisor.
- Partitions provide strong isolation boundaries between all guest virtual machines, and between guest VMs and the root partition.
- The root partition is itself a virtual machine partition, although it has unique properties and much greater privileges than guest virtual machines.
- The root partition provides the management services that control all guest virtual machines, provides virtual device support for guests, and manages all device I/O for guest virtual machines.
- Microsoft strongly recommends not running any application workloads in the root partition

<https://t.me/learningnets>



Continued...

- Each virtual processor (VP) of the root partition is mapped 1:1 to an underlying logical processor (LP). A host VP always runs on the same underlying LP – there is no migration of the root partition's VPs.
- By default, the LPs on which host VPs run can also run guest VPs.
- A guest VP may be scheduled by the hypervisor to run on any available logical processor. While the hypervisor scheduler takes care to consider temporal cache locality, NUMA topology, and many other factors when scheduling a guest VP, ultimately the VP could be scheduled on any host LP.





Hypervisor scheduler types

Hyper-V hypervisor supports several modes of scheduler logic, which determine how the hypervisor schedules virtual processors on the underlying logical processors.

These scheduler types are:

- The classic scheduler
- The core scheduler
- The root scheduler

The classic scheduler

- The classic scheduler has been the default for all versions of the Windows Hyper-V hypervisor since its inception
- The classic scheduler provides a fair share, preemptive round-robin scheduling model for guest virtual processors.
- The classic scheduler type is the most appropriate for the vast majority of traditional Hyper-V uses – for private clouds, hosting providers, and so on.
- The performance characteristics are well understood and are best optimized to support a wide range of virtualization scenarios, such as over-subscription of VPs to LPs, running many heterogeneous VMs and workloads simultaneously, running larger scale high performance VMs, supporting the full feature set of Hyper-V without restrictions, and more.

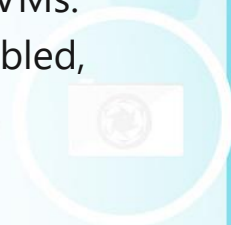
<https://t.me/learningnets>



The core scheduler

- The hypervisor core scheduler is a new alternative to the classic scheduler logic, introduced in Windows Server 2016 and Windows 10 version 1607.
- The core scheduler offers a strong security boundary for guest workload isolation, and reduced performance variability for workloads inside of VMs that are running on an SMT-enabled virtualization host.
- The core scheduler allows running both SMT and non-SMT virtual machines simultaneously on the same SMT-enabled virtualization host.
- The core scheduler utilizes the virtualization host's SMT topology, and optionally exposes SMT pairs to guest virtual machines, and schedules groups of guest virtual processors from the same virtual machine onto groups of SMT logical processors.
- This is done symmetrically so that if LPs are in groups of two, VPs are scheduled in groups of two, and a core is never shared between VMs.
- When the VP is scheduled for a virtual machine without SMT enabled, that VP consumes the entire core when it runs.

<https://t.me/learningnets>



The root scheduler

- The root scheduler was introduced with Windows 10 version 1803. When the root scheduler type is enabled, the hypervisor cedes control of work scheduling to the root partition. The scheduler in the root partition's OS instance manages all aspects of scheduling work to system LPs.
- The root scheduler addresses the unique requirements inherent with supporting a utility partition to provide strong workload isolation, as used with Windows Defender Application Guard (WDAG). In this scenario, leaving scheduling responsibilities to the root OS offers several advantages.
 - For example, CPU resource controls applicable to container scenarios may be used with the utility partition, simplifying management and deployment. In addition, the root OS scheduler can readily gather metrics about workload CPU utilization inside the container and use this data as input to the same scheduling policy applicable to all other workloads in the system.
 - These same metrics also help clearly attribute work done in an application container to the host system. Tracking these metrics is more difficult with traditional virtual machine workloads, where some work on all running VM's behalf takes place in the root partition.

<https://t.me/learningnets>





Selecting the hypervisor scheduler type on Windows Server

The hypervisor scheduler configuration is controlled via the `hypervisorschedulertype` BCD entry.

To select a scheduler type, open a command prompt with administrator privileges:

```
bcdedit /set hypervisorschedulertype type
```

Where `type` is one of:

- Classic
- Core
- Root

<https://t.me/learningnets>