



Utilization vs. Time In-Use

Formerly TSO Logic


Overview

TSO Logic has two concepts around reporting historical usage of virtual and physical server instances:

- Utilization
- Time In-Use

Utilization

How much of the provisioned resources were used? Each server can have multiple resources each with different utilization values (i.e., CPU, memory, storage).



Time In-Use

What percentage of time did the server provide business value? Each server, at any given time, can only be in one classification.

OS INSTANCE CLASSIFICATION ...



For time in-use, the TSO Logic Collector classifies each server, at 15-minute intervals, by examining the CPU utilization. By default, the system is configured to use thresholds of 300 mHz CPU Usage for VMware virtual machines and 5% CPU for Hyper-V and bare-metal machines.

For each 15-minute interval, a server is determined to be in one of only four states:

- In-Use
 - The average metric value for this 15-minute interval exceeds the configured threshold.
- Idle (not in-use)

- The average metric value for this 15-minute interval does not exceed the configured threshold.
- Unclassified
 - No metric values are available for determination and the server instance is not determined to be powered off.
- Expected Idle
 - If the asset was tagged in the 'business data' as expected to be idle (normally part of DR) or is a VMware template. These assets will not be counted as zombies.

While looking over a report period, servers may be classified as:

- 100% expected idle
 - asset tag in the 'business data' or is a VMware template.
- 100% a zombie
 - was not expected idle
 - was never determined to be In-Use for the entire report period but has at least 1 idle classification
 - or was never powered on and power state was never always unknown
- Some % of unclassified, idle and in-use
 - total will be 100%
 - based on number of 15-minute buckets in each category over the report period