

CA MICS® Resource Management Managing z/VM and Linux Platforms

Challenge

Mainframe resource management requires being able to report on all CPU measurements with a high degree of accuracy. Yet, when running Linux for zSeries under z/VM, obtaining accurate CPU measurements can be challenging. z/VM reports on total CPU time for Linux guests, but may not correctly capture the workload level data for individual processes executing on Linux for zSeries. The CPU times reported may be proportionately accurate, but wrong in absolute terms. Linux guests under zSeries are not aware that they are guests in a virtualized environment sharing CPU (and other resources). What sources can be used to get accurate workload level (processes) CPU time for capacity planning, application performance analysis, and chargeback?

Opportunity

CA MICS Resource Management offers a comprehensive solution for resource management, supporting such activities as performance planning, capacity management, financial accountability, availability management, and service level management. The CA MICS solution has historically provided best of class resource management for both mainframe and distributed applications.

In recent years, more CA MICS customers have been taking advantage of the cost and performance benefits provided by Linux application deployment on the z/VM platform, using IFL processors on zSeries mainframe servers.

In many cases, well-tuned business applications have exclusive use of a Linux guest under z/VM. For these Linux instances, the traditional z/VM Monitor and Account records may provide all you need. But for Linux instances running multiple applications, applications shared by different cost centers, or applications with misbehaving processes, you need to see the resource utilization at the workload (process) level. CA MICS is designed to provide accurate data at both the Linux guest and the individual Linux process level.

Solution

To meet the needs of customers requiring accurate process level measurements for Linux applications hosted by z/VM, CA engaged in an informal partnership with Velocity Software to develop an innovative customer solution, which benefitted all three parties. Velocity Software's ESALPS™ product includes pioneering software that uses the process level information available from native Linux facilities; and, with a low overhead solution, is designed to correct invalid (but proportionally correct) CPU times reported by Linux for each process.

The Velocity Software ESALPS product makes use of the z/VM Monitor records that report CPU times for each Linux guest. The total z/VM CPU time for the guest, including CP overhead, is distributed across the internal Linux processes proportionately, based on the amount of CPU for each process reported by Linux for z/Series. ESALPS creates interval-based records for each Linux process which identify the time boundaries for the interval (e.g., 10:30 - 11:00), and the time boundaries within the interval that the process was active (e.g., 10:32 - 10:34). The data record contains numerous process identification metrics as well as additional resource metrics related to memory use, page size, and page faults.

The CA MICS Analyzer Option for VM/CMS reads the process level data generated by the Velocity Software ESALPS product, along with other z/VM Monitor records. From the ESALPS data, CA MICS builds the LNXLXP Linux Process Activity file. This file can be used for accounting and chargeback at the Linux guest process level, but also for process level performance analysis and capacity planning.

One technical caveat—that is well explained in the CA MICS data element documentation—relates to how the CPU times are reported by the VM Monitor for z/VM guests versus the way CPU times are reported by Linux for zSeries processes. Linux for zSeries monitors and reports on two types of CPU time: the Linux Process User CPU Time (LXPVTIME) and the Linux Process System CPU Time (LXPSTIME). These are analogous to the TCB (task) and SRB (system services) CPU time that the IBM z/OS mainframe system reports for z/OS workloads.

z/VM Monitor and z/VM Account data reports two types of CPU time for z/VM guests: virtual (attributable to the guest) and total (virtual plus overhead associated with virtualization). In general, the total CPU time reported by Linux for zSeries (LXPSTIME+LXPVTIME) would normally be equal to the virtual time reported for the guest by z/VM Monitor and Account. But the algorithms of ESALPS distribute the total time used by the guest (virtual plus overhead) over all the process records. To sum up, using the ESALPS process records for a given guest over an interval, you will find that the total CPU time for a guest (virtual plus overhead) is equivalent to the total CPU time across all active processes in the interval (user plus system) with minor variations due to process activity occurring near the interval boundaries.

With this capability, CA MICS provides data to assist customers taking advantage of the many benefits provided by Linux applications deployed on z/VM—meeting workload service level, cost accounting, and performance analysis requirements.