World's first performance analysis tool that identifies root causes of performance issues in virtual environments
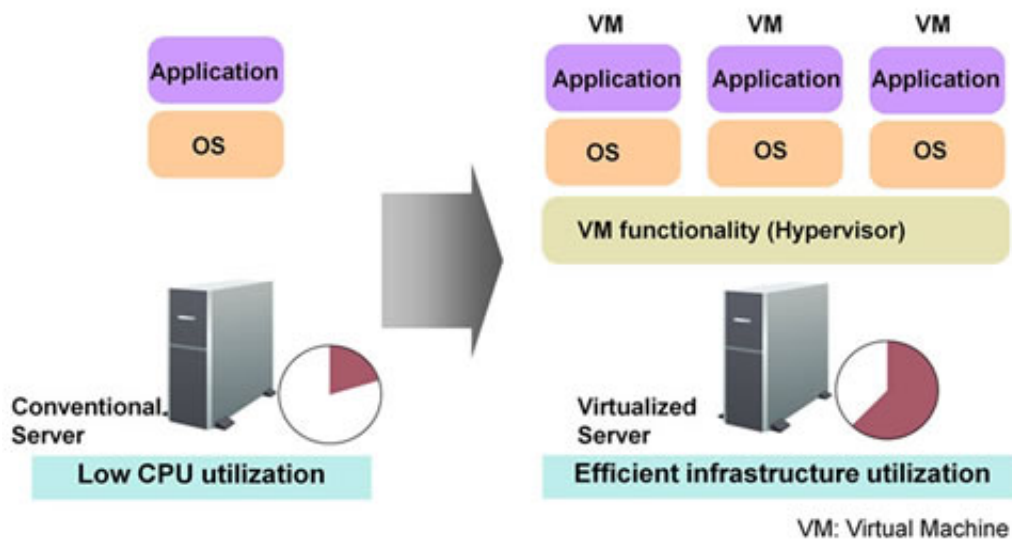
March 12 2013

Figure 1: Infrastructure optimization through server virtualization.

Fujitsu Laboratories announced the development of the world's first technology that can identify the root cause of performance problems in the widely used virtual environments that are part of the infrastructure underlying cloud computing and enterprise information systems.

Up until now, when problems arose within a virtual environment, the absence of a tool to analyze the behavior of the virtual environment as a

whole made it difficult to identify the root cause of slowdowns at the application or program level. Fujitsu Laboratories has developed a technology that performs overall program behavior analysis by collecting execution information from user programs in each virtual machine (VM) on the virtual platform's hypervisor.

This technology can quickly and accurately analyze performance characteristics from the hypervisor level to the OS and user programs (business applications) on the VM, thereby grasping overall behavior. This should help reduce costs by improving the efficiency of server resource utilization.

In server virtualization, a single physical server can run multiple VMs simultaneously, with each VM running its own operating system and user programs. This makes use of the continued improvements to server CPUs and makes it possible to reduce equipment costs by consolidating on one server the operations that had previously required multiple servers.

With the popularization of virtual environments, instances of end users experiencing a sudden drop in performance without knowing why, or experiencing a less-than-expected performance output, have become common problems.

The reasons for these issues occurring could be among the following:

- Having multiple VMs and a hypervisor environment that provides VM functions all on one physical server makes the overall operating environment more complex, thus making it difficult to identify the cause of the problem.
- Due to the VMs being constantly switched in and out of operation by the hypervisor, it is not possible to accurately

monitor and analyze the performance by a single VM at a time.

In mission-critical applications, interruptions or stoppages due to system failures are unacceptable. As such applications become increasingly virtualized, performance expectations for virtual environments have risen. The same is true of the requirement for program-level optimization of overall system performance, which includes applications.

To address the problem of deteriorating performance and other performance-related issues in virtual environments, it has become possible to confirm where bottlenecking occurs in the CPU, memory, or I/O resource levels. However, there have been no performance-analysis tools that could accurately and comprehensively analyze the characteristics from processing that takes place within the hypervisor to the processing of the user programs and the operating system running on the VMs. As a result, one of the only ways to resolve problems was to move VMs or reallocate resources, raising concerns about the ineffective use of resources and potentially higher operational management costs.
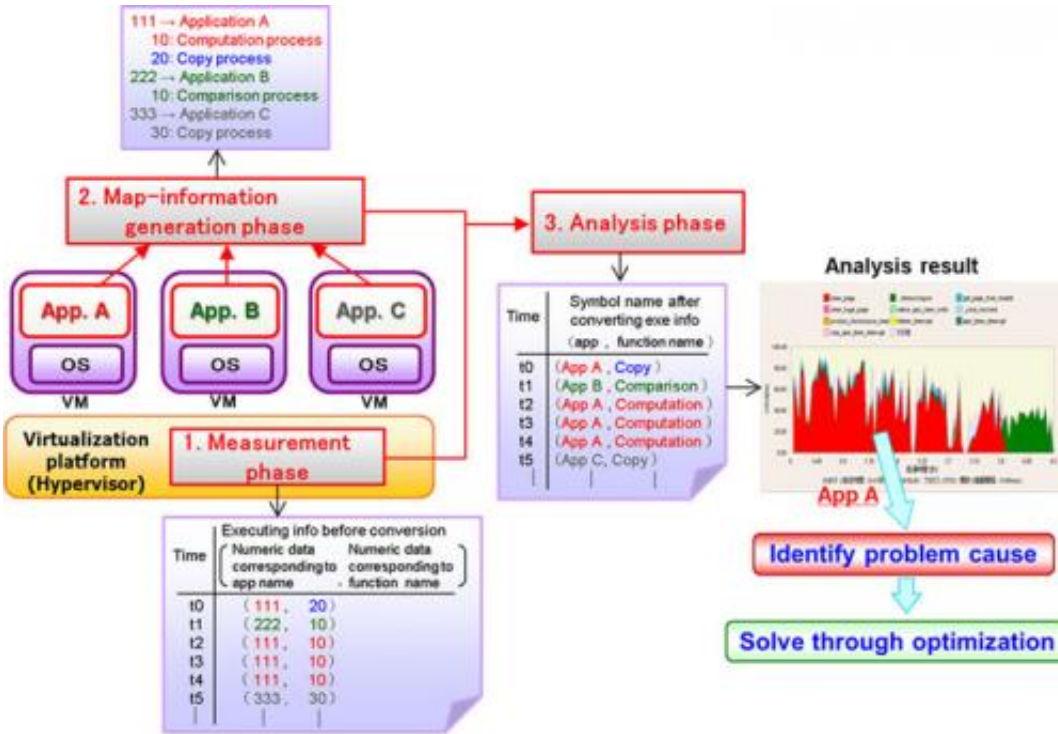
Figure 2: VM performance-analysis process flow.

Fujitsu Laboratories has developed a new tool that collects execution information from the user programs on each VM via the hypervisor and can then quickly and accurately analyze performance characteristics, from the hypervisor processing level to the processing level on the OS and user programs running on the VM. The process from the information gathering stage to the analysis stage is as follows:

1. Measurement phase

Execution information on each user program running on the VMs is centrally sampled at the hypervisor level, rather than the VM in operation.

2. Map information generation phase

Map information is generated in each VM, for converting the data collected in the measurement phase to the symbol names of applications, process functions, etc.

3. Analysis phase

Execution information on the user programs in each VM collected in the measurement phase is cross-referenced with the data from the map information generation phase, producing a central hypervisor time-series analysis that breaks down the active program run states in each VM.

This centralized data sampling and analysis makes it possible to accurately grasp the actual states of complex VMs. Also, simply by combining this functionality with the hypervisor, it becomes possible to analyze the performance of "black box" VMs without having to modify the operating systems or applications in each VM.

This technology makes it possible to identify the cause of problems in virtual environments that had previously been difficult to identify, such as why a particular performance issue arose or which processing program is the cause of the problem that needs to be resolved. It makes it possible to quickly and accurately analyze performance and understand behaviors, opening the way to more efficient resource utilization, which, in turn, should help lower costs.

The new technology will be used by Fujitsu's Linux technology support services beginning in April 2013.

Provided by Fujitsu

Citation: World's first performance analysis tool that identifies root causes of performance issues in virtual environments (2013, March 12) retrieved 23 June 2024 from https://phys.org/news/2013-03-world-analysis-tool-root-issues.html