VirtualWisdom® Virtual Server Probe

Model ProbeVM

Virtual Instruments products provide an unprecedented scope of diagnosis and prevention capabilities for complex, virtualized infrastructures running mission-critical applications

- → Accelerates the deployment and reduces the risk of implementing virtualized I/O-intensive business-critical applications such as Oracle, SAP, DB2, SQL, and Exchange
- → Reduces time spent on problem resolution with early detection of I/O performance bottlenecks and transmission faults
- → Lowers overall downtime by generating predictive data on potential virtual infrastructure problems
- → Improves virtualized application performance by enabling administrators to optimally balance virtualized workloads based on I/O response time and utilization data, not just CPU and memory utilization
- → Lowers overall operating expenses and capital expenses through ESX server modeling and higher utilization of existing server/SAN assets

The Virtual Server Probe enables comprehensive, cross-domain real-time measurement capabilities that allow IT managers to optimize the performance, utilization, and availability of their virtualized IT infrastructure. The Virtual Server Probe adds deep SAN I/O intelligence and operational visibility into VMware deployments, enabling administrators to model ESX / ESXi server performance and better balance the deployment of virtual machines based on real-time measurements and feedback of I/O performance. The Virtual Server Probe is unique in its ability to eliminate the risk of using virtualization with I/O-intensive Tier 1 applications such as VDI (Virtual Desktop Infrastructure) and those based on Oracle®, IBM DB2®, Microsoft SQL Server® or SAP®.

Use of VirtualWisdom and the Virtual Server Probe results in significantly higher virtual infrastructure utilization and helps server and VMware administrators reduce capital and operational expenses associated with an IT organization's servers and storage. Unlike most server and storage vendor-supplied tools that only look at their own devices, VirtualWisdom looks across all devices: servers, host bus adapters (HBAs), switches, cables, and storage to optimize application performance and the overall utilization of the IT infrastructure.

Virtualization has been heavily deployed in nearly every large enterprise across nearly every industry to lower both capital and operating costs associated with server infrastructure. Most use of virtualization today has been associated with test, development, file serving, web serving, or other non-business-critical production workloads. Unfortunately, virtualization of mission-critical Tier 1 applications has been minimal to non-existent due to fears of performance bottlenecks and the inability to quickly identify and resolve I/O problems associated with virtualized I/O-intensive applications or VDI. VMware states that over 80% of performance problems

experienced by their users are I/O related, yet VMware software has no VM to LUN view of SAN I/O traffic patterns. VirtualWisdom does. A recent Forrester study found that "Maintaining High Performance" was the number one challenge related to storage for virtual server environments.

Given that close to 75% of all VMware deployments are connected to a Fibre Channel SAN, understanding SAN I/O traffic is essential to optimizing the performance, utilization and availability of virtualized mission-critical applications. Only with Virtual Instruments' VirtualWisdom and the Virtual Server Probe can IT gain this understanding. VirtualWisdom enables you to measure and correlate metrics on the VM to the LUN.

The Virtual Server Probe collects information specifically about the utilization and performance of virtual machines within the VMware environment via the VMware vCenter server API. VirtualWisdom is the perfect complement to vCenter as it correlates the vCenter information from the Virtual Server Probe in real-time with actual I/O performance data measured by the Virtual Instruments SAN Performance Probe directly from the SAN infrastructure. Actual measured data from VirtualWisdom enables IT managers to proactively balance the provisioning of applications on virtual machines for maximum application performance with the confidence that systems will not slow down or fail. This high level of virtual I/O visibility makes risk-mitigated optimization and ESX / ESXi server performance modeling possible, enabling mission-critical deployment of virtualized applications possible by verifying Service Level Agreement (SLA) metrics in real-time.

Features

- Agent-less architecture enables quick deployment with real-time monitoring of virtual machines, physical servers, and clusters
- Collects data on and calculates 120+ different metrics including CPU utilization and status, memory utilization, disk I/O requests and capacity, network requests and utilization
- Enables monitoring, modeling, and optimization of VMware vMotion and Distributed Resource Scheduler for performance and availability based on historical data and trends
- User definable data collection frequency from VMware vCenter – as low as every 20 seconds
- Policy-based event notifications
- Simple dashboard-view of virtual to physical server I/O from server to storage arrays with drill-down reports for quick problem resolution
- Simple installation into any VMware virtualized environment

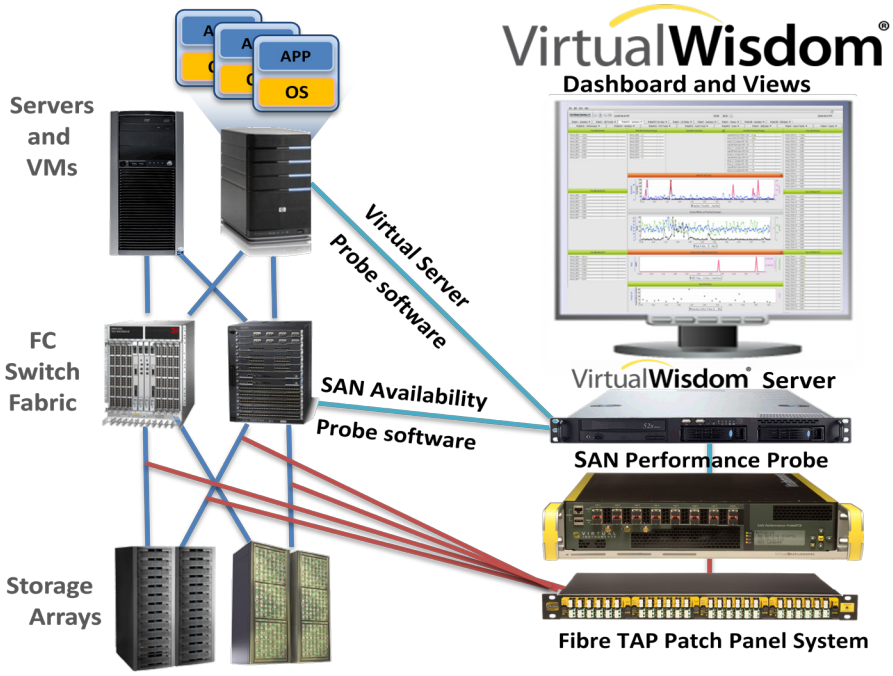

The VirtualWisdom Virtual Server Probe is a virtual software probe that automatically discovers and monitors 120+ CPU, memory, network and disk I/O metrics for VMware ESX / ESXi virtual machines, physical servers, and clusters. The Virtual Server Probe is agent-less and can be configured using User Definable Correlations (UDCs) and filters to monitor and track I/O from any combination of virtual machines or physical servers at any time. I/O data from the virtual machines and servers are automatically stored for historical purposes by the VirtualWisdom Server to enable trend analysis, performance modeling, and policy setting.

The VirtualWisdom Server collects the real-time data and correlates server, network, and I/O metrics from the Virtual Server Probe, The SAN Availability Probe, and the real-time SAN Performance Probe into a complete view of the data path between virtual machines and the SAN storage. Performance deviations or faults are quickly diagnosed utilizing real-time and historical trends and reports. The number of problems and time spent on resolving performance issues or network faults are systematically and dramatically reduced.

In large VMware environments where the VMware Distributed Resource Scheduler initiates multiple virtual machine migrations or VMware vMotion transfers, VirtualWisdom and the Virtual Server Probe play an important part in ensuring the effects on business operations are minimal. During a migration of virtual machines from one server to another, all data synchronization between the two virtual machines takes place within the SAN. The Virtual Server Probe and the SAN Performance Probe monitor these sudden I/O heavy migrations and if thresholds are exceeded, policy-based alerts quickly enable administrators to take action if needed to avoid congestion and faults for other VMs or applications.

VirtualWisdom's configurable dashboard provides a simple, customizable view of I/O from the application on the virtual machine and server to the LUNs on the storage arrays. Metrics, fault warnings and policy deviations are continually updated with quick drill-down reports. SNMP traps can be configured on nearly all of the leading management platforms to receive alarms from VirtualWisdom to ensure maximum performance on mission-critical applications. Finally, alarms can be configured within VirtualWisdom to e-mail reports triggered by an alarm.

The Metrics Collected or Generated by the Virtual Server Probe include:

Calculated Network Virtual Machine Metrics

Disk IO Utilization % - The average Disk IO Utilization percentage of total Server HBA capacity. Calculation is Avg Disk Usage / Server HBA Capacity *where:*

- Avg Disk Usage = Disk Read Rate + Disk Write Rate
- Total HBA Link Capacity = Sum of HBA Link Rates from HBA ports

Network Utilization % - The average Network Utilization percentage of total Server NIC capacity. Calculation is Avg Net Usage / Total NIC Capacity *where:*

- Avg Network Data Rate = Network Data Transmit Rate + Network Data Receive Rate
- Total NIC Capacity = Sum of NIC Speeds from NIC card Ports

Composite Utilization % - A combination index that shows the average percentage of total Server capacity for four dimensions of resource usage including CPU, Memory, Disk IO, and Network. Calculation is (Disk IO Utilization % + Network Utilization % + CPU Usage % + Memory Usage %) / 4

Max Composite Utilization % - The maximum total Server capacity resource observed for the four dimensions of resource usage including CPU, Memory, Disk IO and Network. Calculation is Max (Disk IO Utilization, Network Utilization %, CPU Usage %, Memory Usage %)

Context Virtual Machine Metrics

Virtual Machine: The name of the Virtual Machine being measured.

Server: The name of the Server being measured.

Cluster: The name of the Cluster being measured

Total HBA Link Capacity (Gb/s) The total Host Bus Adapter link capacity (Gb/s). This is used to calculate Disk IO Utilization %.

Total NIC Capacity (Mb/s): The total Network Interface Card Capacity (Mb/s). This is used to calculate Network Utilization %.

Disk IO Utilization %: The average Disk IO Utilization percentage of total Server HBA capacity

Network Utilization: % The average Network Utilization percentage of total Server NIC capacity.

Composite Utilization: % A combination index that shows the average percentage of total Server capacity for four dimensions of resource usage including CPU, Memory, Disk IO, and Network.

Max Composite Utilization %: The maximum total Server capacity resource observed for the four dimensions of resource usage

CPU Virtual Machine Metrics

CPU Extra (ms): CPU time (ms) that is extra.

CPU Guaranteed (ms): CPU time (ms) that is guaranteed.

CPU Ready (ms):CPU time (ms) spent in ready state.

CPU System (ms): CPU time (ms) spent on system processes

CPU Usage %: CPU usage as a percentage over the collected interval.

CPU Usage (MHz) : CPU usage in MHz over the collected interval. For hosts this can be represented on a per Virtual Machine basis as a stacked graph.

CPU Used (ms): CPU time (ms) that is used.

CPU Wait (ms): CPU time (ms) spent in wait state.

Avg CPU Active % (5 min.): CPU active average % over 5 minute

Max CPU Active % (5 min.): CPU active peak over 5 minute
Max CPU Throttled % (5 min.): Amount of CPU resources over the limit that were refused, average over 5 minute

Avg CPU Running % (5 min.): CPU running average over 5 minute

Max CPU Running % (5 min.): CPU running peak over 5 minute

Memory Virtual Machine Metrics

Memory Active (KB): Amount of memory that is actively used (KB).

Memory Consumed (KB): Amount of host memory (KB) consumed by the virtual machine for guest memory.

Memory Granted (KB): Amount of memory (KB) granted. For hosts this can be represented on a per Virtual Machine basis as a stacked graph.

Memory Overhead (KB): Amount of additional host memory (KB) allocated to the virtual machine.

Memory Shared (KB) : Amount of memory (KB) that is shared.

Memory Swap In (KB): Amount of memory (KB) that is swapped in.

Memory Swap Out (KB): Amount of memory (KB) that is swapped out.

Memory Swapped (KB): Amount of memory (KB) that is swapped.

Memory Swap Target (KB): Amount of memory (KB) that can be swapped.

Memory Usage %: Memory usage as percentage of total configured or available memory.

Memory Balloon (KB): Amount of memory (KB) used by memory control.

Memory Balloon Target (KB): Amount of memory (KB) that can be used by memory control.

Memory Zero (KB): Amount of memory (KB) that is zeroed out.

Disk Virtual Machine Metrics

Disk Bus Resets: Number of bus resets in the period.

Disk Commands Issued: Number of disk commands issued in the period.

Disk Command Aborts: Number of disk commands aborted in the period.

Disk Read Requests: Number of disk reads in the period.

Disk Write Requests: Number of disk writes in the period.

Disk Read Rate (KB/s): Rate of reading data from the disk (KB/s).

Disk Usage (KB/s): Aggregated storage performance statistics (KB/s). For hosts this can be represented on a per Virtual Machine basis as a stacked graph.

Disk Write Rate (KB/s): Rate of writing data to the disk (KB/s).

Network Virtual Machine Metrics

Network Packets Received: Number of packets received in the period.

Network Packets Transmitted: Number of packets transmitted in the period.

Network Data Receive Rate (KB/s): Rate at which data is received (KB/s).

Network Data Transmit Rate (KB/s): Rate at which data is transmitted (KB/s).

Network Usage (KB/s): Aggregated network performance statistics (KB/s). For hosts this can be represented on a per Virtual Machine basis as a stacked graph.

Other Virtual Machine Metrics

Avg CPU Active (1 min.): CPU active average over 1 minute.

Max CPU Active % (1 min.): CPU active peak over 1 minute.

Max CPU Throttled % (1 min.): Amount of CPU resources over the limit that were refused, average over 1 minute.

Avg CPU Running % (1 min.): CPU running average over 1 minute.

Max CPU Running % (1 min.): CPU running peak over 1 minute.

Group CPU Sample Count: Group CPU sample count.

Group CPU Sample Period (ms): Group CPU sample period (ms).

Heartbeat: Number of heartbeats in this period.

Resource CPU Usage (MHz): Resource CPU usage (MHz).

Uptime (sec): Total time (sec) elapsed since last startup.

Context Server Metrics

Server: The name of the Server/host being measured.

Cluster: The name of the Cluster being measured.

Total HBA Link Capacity (Gb/s): The total Host Bus Adapter link capacity (Gb/s). This is used to calculate Disk IO Utilization %.

Total NIC Capacity (Mb/s): Network Interface Card Speed (Mb/s).

Disk IO Utilization %: The average Disk IO Utilization percentage of total Server HBA capacity.

Network Utilization %: The average Network Utilization percentage of total Server NIC capacity.

Composite Utilization %: A combination index that shows the average percentage of total Server capacity for four dimensions of resource usage including CPU, Memory, Disk IO, and Network.

Max Composite Utilization %: The maximum total Server capacity resource observed for the four dimensions of resource usage including CPU, Memory, Disk IO, and Network.

CPU Server Metrics

CPU Idle (ms): CPU time (ms) spent in idle state.

CPU Reserved Capacity (MHz): Total CPU capacity (MHz) reserved by the virtual machines.

CPU Usage %: CPU usage as a percentage over the collected interval.

CPU Usage (MHz): CPU usage (MHz) over the collected interval. For hosts this can be represented on a per Virtual Machine basis as a stacked graph.

CPU Used (ms): CPU time (ms) that is used

Avg CPU Active % (5 min.): CPU active average % over 5 minute

Max CPU Active % (5 min.): CPU active peak over 5 minute Max CPU Throttled % (5 min.) 3 Amount of CPU resources over the limit that were refused, average over 5 minute

Avg CPU Running % (5 min.): CPU running average over 5 minute

Max CPU Running % (5 min.): CPU running peak over 5 minute

Memory Server Metrics

Avg Memory Active (KB): Amount of memory (KB) that is actively used.

Avg Memory Granted (KB): Amount of memory granted (KB). For hosts this can be represented on a per Virtual Machine basis as a stacked graph.

Avg Memory Heap (KB): Amount of memory (KB) allocated for heap.

Avg Memory Heap Free (KB): Free space in memory heap (KB).

Avg Memory Overhead (KB): Amount of additional host memory (KB) allocated to the virtual machine.

Memory Reserved Capacity (KB): Amount of memory (KB) reserved by the virtual machines.

Avg Memory Shared (KB): Amount of memory (KB) that is shared.

Avg Memory Shared Common (KB): Amount of memory (KB) that is shared by common.

Memory State: Memory state.

Avg Memory Swap In (KB): Amount of memory (KB) that is swapped in.

Avg Memory Swap Out (KB): Amount of memory (KB) that is swapped out.

Avg Memory Swap Used (KB): Amount of memory (KB) that is used by swap.

Avg Memory Unreserved (KB): Amount of memory (KB) that is unreserved.

Memory Usage %: Memory usage as percentage of total configured or available memory.

Avg Memory Zero (KB): Amount of memory (KB) that is zeroed out.

Avg Memory Consumed (KB): Amount of host memory (KB) consumed by the virtual machine for guest memory

Avg Memory Balloon (KB): Amount of memory (KB) used by memory control

Disk Virtual Machine Metrics

Disk Bus Resets: Number of bus resets in the period.

Disk Commands Issued: Number of disk commands issued in the period.

Disk Command Aborts: Number of disk commands aborted in the period.

Avg Physical Device Command Latency (ms): The average time (ms) taken to complete a command from the physical device.

Avg Physical Device Read Latency (ms): The average time (ms) taken to complete a read from the physical device.

Avg Physical Device Write Latency (ms): The average time (ms) taken to complete a write from the physical device.

Avg Kernel Disk Command Latency (ms): The average time (ms) spent in ESX Server VMKernel per command.

Avg Kernel Disk Read Latency (ms): The average time (ms) spent in ESX Server VMKernel per read.

Avg Kernel Disk Write Latency (ms): The average time (ms) spent in ESX Server VMKernel per write.

Disk Read Requests: Number of disk reads in the period.

Disk Write Requests: Number of disk writes in the period.

Avg Queue Command Latency (ms): The average time (ms) spent in the ESX Server VMKernel queue per command.

Avg Queue Read Latency (ms): The average time (ms) spent in the ESX Server VMKernel queue per read.

Avg Queue Write Latency (ms): The average time (ms) spent in the ESX Server VMKernel queue per write.

Disk Read Rate (KB/s): Rate of reading data from the disk (KB/s).

Avg Disk Command Latency (ms): The average amount of time (ms) taken for a command from the perspective of a Guest OS. This is the sum of Kernel Command Latency and Physical Device Command Latency.

Avg Disk Read Latency (ms): The average amount of time (ms) taken for a read from the perspective of a Guest OS. This is the sum of Kernel Read Latency and Physical Device Read Latency.

Avg Disk Write Latency (ms): The average amount of time (ms) taken for a write from the perspective of a Guest OS. This is the sum of Kernel Write Latency and Physical Device Write Latency.

Disk Usage (KB/s): Aggregated storage performance statistics (KB/s). For hosts this can be represented on a per Virtual Machine basis as a stacked graph.

Disk Write Rate (KB/s): Rate of writing data to the disk (KB/s).

Max Disk Command Latency (ms): The highest latency value across all disks used by the host

Network Server Metrics

Network Packets Received: Number of packets received in the period.

Network Packets Transmitted: Number of packets transmitted in the period.

Network Data Receive Rate (KB/s): Rate at which data is received (KB/s).

Network Data Transmit Rate (KB/s): Rate at which data is transmitted (KB/s).

Network Usage (KB/s): Aggregated network performance statistics (KB/s). For hosts this can be represented on a per Virtual Machine basis as a stacked graph.

Network Receive Packets Dropped: Number of receive packets dropped in the period

Network Transmit Packets Dropped: Number of transmit packets dropped in the period

Other Server Metrics

Avg Agent Memory Used %: Memory Used as percentage of total configured or available memory.

Avg Agent Memory Swap In (KB): Amount of memory (KB) that is swapped in.

Avg Agent Memory Swap Out (KB): Amount of memory (KB) that is swapped out.

Avg Agent Memory Swap Used (KB): Amount of memory (KB) that is used by swap.

Avg CPU Active % (1 min.): CPU active average over 1 minute.

Max CPU Active % (1 min.): CPU active peak over 1 minute.

Max CPU Throttled % (1 min.): Amount of CPU resources over the limit that were refused, average over 1 minute.

Avg CPU Running % (1 min.): CPU running average over 1 minute.

Max CPU Running % (1 min.): CPU running peak over 1 minute.

Group CPU Sample Count: Group CPU sample count.

Group CPU Sample Period: Group CPU sample period (ms).**Avg Resource CPU Usage (MHz):** Resource CPU usage (MHz).

Uptime (sec): Total time (sec) elapsed since last startup.

CPU Fairness: Fairness of distributed cpu resource allocation

Memory Fairness: Aggregate available memory resources of all hosts within a cluster

VIRTUAL INSTRUMENTS

Corporate Headquarters
25 Metro Drive Suite 400
San Jose, CA 95110
Phone: 408-579-4000
Fax: 408-579-4001

Sales
sales@virtualinstruments.com
Phone: 408-579-4081

Support
support@virtualinstruments.com