



live optics

PERFORMANCE OVERVIEW



PROJECT: KC-APOLLO 2020-09-01 14:50 -05:00

PREPARED FOR: Reaha Dahlke, RDahlke@kcata.org

PREPARED BY: Reaha Dahlke, RDahlke@kcata.org,

CREATED: 9/2/2020



Intro






This document is an assessment of the server and storage workloads as observed at KANSAS CITY AREA TRANSPORTATION AUTHORITY, MISSOURI based on servers, server disk activity, and storage capacity. The data depicted in this report may be used to make recommendations on future configurations and solutions.

This report may be used to identify times of peak IO activity, and marks the critical 95th IOPS thresholds, a key metric used as a baseline for sizing. This report includes details on servers, CPU and memory usage, as well as free and used storage capacities.

If you have any questions regarding the information presented in this report, please contact your Live Optics team for immediate assistance.






Executive Summary

Time Recorded	23 Hour(s), 18 Minute(s), 9/1/2020 - n/a		
Total Combined Capacity	31.54 TB	Peak Aggregate Network ...	35.44 MB/s
Total Used/Free Capacity	21.44 TB / 10.10 TB	Total Memory	537.41 GB
IOPS	10401 at peak, 2409 at 95%	VM Guest Count	0
Read/Write Ratio	90% / 10%	Peak page fault of 11699 for server KC-GFI	
Average Daily Write	643.42 GB	A peak of 25.1 (6.1%) of 413.2 NetCPU Cycles (GHz)	
Total Cores / Total Proces...	214 / 71	50% of IOPS falls on 20% of your capacity (4.29 TB)	



Executive Storage Summary

Time Recorded	23 Hour(s), 18 Minute(s), 9/1/2020 - n/a		
Total Local Capacity	31.05 TB	Total Shared Capacity	500.00 GB
Free/Used Local Capacity	10.08 TB (32%), 20.97 TB (68%)	Free/Used Shared Capacity	18.00 GB (4%), 482.00 GB (96%)
Disk Throughput	319 MB/s	Total Disks	62



Definitions

CPU

The CPU shown in the individual server reports measures the utilization of the CPU as a percentage of the total available CPU clock cycles. When shown in the aggregate summary, the CPU aggregate represents the sum of the clock cycles utilized. The saturated core count shows the number of CPU cores that were observed to have a usage level of 90% or more during the sampling period. This count provides insight into how your CPU workloads are distributed across all the cores. If a core is saturated, the task that it is processing is bottlenecked on the CPU. By spreading tasks out over multiple threads (and thus multiple cores) the bottleneck may be removed. Live Optics supports saturated core counts for Windows and Linux hosts. Saturated core counts for other host operating systems are not supported.

IO Transfer Size

This unit of measure records the average payload size of a given set of read or write operations. The average IO transfer size may be used to describe the nature of the IO workload and is an important factor in understanding the character of an IO workload.

IOPS

Live Optics monitors the read and write operations at the block storage layer of the host operating system. IOPS are the Input / Output Operations per Second performed by a physical server host or a virtual machine.

Latency

The average time required for the host to complete an IO operation is measured as its average latency. Live Optics monitors latency at the host level. Like queue depth, the host's latency will differ from latency reported by external storage arrays, and individual disks within RAID sets. Intermittent spikes in latency should not be a matter of concern. Also, average latency reported in Live Optics is only accurate when the sustained IOPS are over 100 IOs per second. High latencies reported during very low periods of IOPS should be discounted and often can be attributed to counter skew factors associated with the latency calculation. Latency is an important factor when investigating a storage solution that is demonstrating known problems at the application layer. Most applications will not tolerate abnormally high levels of latency on a sustained basis.

MB/s

Megabytes per second represents the average rate of read and write IO transfers in terms of payload size per second. This measures the IO throughput of a given system and is an important factor in sizing storage network bandwidth requirements.

Memory

The memory shown in the individual server reports measures the total available (free) and used memory for the server. When shown in the aggregate summary, the total memory shown represents the aggregate memory usage (and free memory) of all servers in the report.

Participation

This simple chart shows the number of servers that were monitored for the given sampling period. If the connection between Live Optics and a server was lost, the participation chart will show a drop during the period when the server was out of communication with Live Optics.

Queue Depth (QD)

The number of outstanding IO operations that are the host is processing at any given time is the queue depth. Live Optics monitors queue depth at the host level. Host level queue depth will differ from the queue depth reported by external storage arrays, and individual disks within RAID sets. Storage devices tend to perform faster when queue depths are higher. However, high queue depths can also lead to higher than acceptable latencies.

Definitions Continued

Page Faults

Page Faults per Second show the rate of hard page faults that required a virtual memory page to be read from disk. High hard page fault rates may suggest insufficient physical memory allocation to the OS and can be signs of application performance problems.

Workload Concentration Line

Workload Concentration plots the total used capacity of the project by disk along the X-axis and the IOPS activity of each disk along the Y-axis. The disks are sorted by IOPS with the disks with the highest IOPS being closest to the Y-axis on the left. The vertical line drawn splits the graph into two sections, with the section on the left indicating the region that comprising 20% of the total used capacity, and the region on the right showing the remaining 80% of the capacity. The label shows you the portion of IOPS that fall on either region. Note: the IOPS shown are observed only at the disk scope. Live Optics does not analyze IO activity at a block level based on the address of individual IOs on a given disk. Additional IO concentration may or may not be present at the block level within the individual disks.

Workload Concentration Bubble

Workload Concentration plots a bubble chart of your disks and sorts them by their IO activity. The vertical placement of the bubble represents the amount of IO. The size of each bubble is proportional to the capacity of the disk. The horizontal axis demonstrates the total amount of IO that falls within a certain percentage of total capacity.