

Data Center Refresh - Using Analytics to Make Informed Decisions

WHITE PAPER

Stu Fisher, Business Development Director, OakGate Products

Rob Dobson, Director Product Marketing, OakGate Products

Date: Jan 2021

Summary

Planning, designing & characterizing Data Center production workloads is critical to the overall business success and efficiency for any Data Center operations team.

Adding more system memory and faster hardware should be made only if the DC will see performance gains.

Characterization of production workloads allows accurate performance modeling and is critical to DC configuration change and hardware refresh.

WorkloadIntelligence™

Analytics provides deeper insight into what is happening at the system and storage layer to what DCs could get previously. This allows the DC to improve overall performance and tune the block layer to achieve better workload performance.

Introduction

By the end of 2020 there will more than 1 trillion connected objects and devices on the planet generating data. There are 2.5 billion gigabytes of data generated every day of which 80% is unstructured. Data has become the most valuable and most strategic resources for any company.

As the amount of data increases exponentially, increasingly enterprises are using Data Centers to make effective decisions and gain a competitive edge. Additionally, Data Centers and enterprises are being challenged to implement new technologies like solid-state storage (SSD) and produce results faster without increasing their budgets. Diverse blended workloads like big data, deep learning, and cloud computing add to this challenge. Moreover, a Data Center must be able to meet challenging 24 by 7 up time requirements and meet the demand of a constantly-changing application workload.

Challenges for the Data Center

To keep up, Data Center server and storage platform refresh cycles are a constant for operations teams and can cause significant anxiety. Knowing when to commit real dollars and resources to a hardware and/or software refresh can be daunting. A refresh can also be extremely costly.



Finding the right set of software tools to assist in analyzing what is best for the company and its customers can be daunting. Planning, designing & implementing a sound hardware

and/or software refresh is critical for overall business success and efficiency of any Data Center regardless of its purpose. Moreover, refresh

performance modeling, workload analysis, and capacity planning are much more accurate with real-world production workload data to help optimize this planning.

Planning a Data Center Refresh?

There are a number of arguments for justifying a hardware and/or software refresh in the Data Center. We have captured a few below:

- Outdated hardware or software can mean increased failures and down time
- Equipment coming off of lease (typically 3 to 5 years)
- Improvements in next generation CPU, memory, storage and software technologies
- New efficiencies in hardware power dissipation and consumption
- Perception that better, faster and more efficient hardware will improve system and storage latencies

This table below outlines how *WorkloadIntelligence™* Analytics can help your team execute a successful DC refresh.

Steps	Description	How Analytics Can Help
1. Determine current hardware specs and performance statistics	<ul style="list-style-type: none"> • Ensure you can characterize the current hardware specs and performance statistics • Identify systems performing nearly at capacity • Identify systems that are running with plenty of performance headroom 	Use Analytics to characterize your current workload performance and efficiency on the systems nearing capacity. Use this information to redirect to other systems with suitable system resources.
2. Document and update business and user needs	<ul style="list-style-type: none"> • Analyze your current systems and the associated business and user needs • Determine if their needs have changed. This will help plan future hardware and software planning 	Use Analytics to determine what processes are driving over/under utilization on your systems.
3. Look at performance and functionality benefits	<ul style="list-style-type: none"> • Characterize performance and functionality increase • Focus on ensuring the data you gathered in Step 2 applies as you select replacement hardware • Don't waste money on unnecessary upgrades to turn low-level servers into high-end systems 	Analytics can help characterize workload performance on the new hardware platforms to confirm increased performance.
4. Consider costs including labor	<ul style="list-style-type: none"> • The time your staff will spend on a refresh project should be factored in when determining the projected costs • Factor in the time spent on the proposed upgrades but also how staff will benefit by saving time not having to troubleshoot and support the old systems 	Give your staff the analytical tools to pro-actively characterize performance issues and fix the problem and/or to pick the right SSD.
5. Implement an orderly upgrade process	<ul style="list-style-type: none"> • Upgrade from the direction of least critical to most critical servers. This allows you to test everything first, back out any changes gone awry, and involve the least potentially negative impact upon business operations. • Establish performance baselines as you implement new systems 	Analytics can help determine where/when problems occur and provide the team detailed data that they will need during the refresh cycle.

Refreshing a Data Center with latest generation hardware has been proven to provide, on average, anywhere between a two to five percent increase in performance. The question all Data Center operations teams must ask themselves “is this increase enough to move forward with the refresh?”

Given that Data Centers are very cost sensitive, reducing the cost of a single element by a small fraction immediately translates into huge cost savings due to the large scale. When designing a Data Center, it is important to understand the production workloads running in your environment and optimize the hardware and software refresh for these workloads so that the best possible performance per dollar is achieved.

System administrators and Data Center architects need processes and easy-to-use tools to optimize and manage the array of available system and storage resources, whether in the Data Center or in concert with cloud computing environments. Consequently, it has become vitally important to understand workload behaviors in order to further improve the performance of data center computer systems. Detailed knowledge of application workload characteristics can optimize performance of current and future systems.

Using *WorkloadIntelligence*™ Analytics to Optimize Data Center Refresh

Designing optimized systems requires a deep understanding of target workloads and mechanisms to effectively model different design choices. Traditional benchmarking is invalid in hyperscale data stores. Storage or system profiles are hard to obtain and replaying the entire application in all configurations is impractical. Additionally, current workload generators are not comprehensive enough to accurately reproduce key aspects of real application patterns.



The *WorkloadIntelligence*™ Analytics application enables users to optimize their infrastructure and applications in high-performance, hyperscale environments through the analysis of imported workload traces (such as Linux block layer I/O traces). When operating the Analytics application, users can review and analyze I/O data and process identifiers (PIDs); evaluate advanced performance charts, which have detailed analytics and synchronization capabilities; and plot charts with an extensive selection of parameters for reads, writes, and trim IOs.

WorkloadIntelligence™ Analytics enables the user to create canvases of built-in or custom high-performance charts of various parameters. These charts can represent a range of billions of data points to a single data point depending on the selected time range. This charting flexibility and depth provides an extraordinary insight into workload data never available before.

Key Benefits

There are a number of benefits for using WorkloadIntelligence™ Analytics to help optimize any Data Center refresh. We have listed the major benefits below.

- Shows key insights on how the Data Center or application platforms actually perform
- Allows customer to make more informed hardware and application decisions that directly affect their users and customers without totally relying on data sheet specifications
- Helps to deliver cost savings from better hardware utilization and process improvements
- Helps to improve the customer experience through system performance improvements
- Allows the Data Center staff to work directly with the system and storage vendors to optimize their hardware to help improve performance with the Data Center's workload
- Helps to understand if delaying the refresh can actually save on costs

WorkloadIntelligence™ Analytics Workflow

The WorkloadIntelligence™ workflow is highlighted below - **see figure 1**. Once Data Center workloads are imported, the traces can be analyzed across a number of system attributes. The traces can be saved and shared with others.

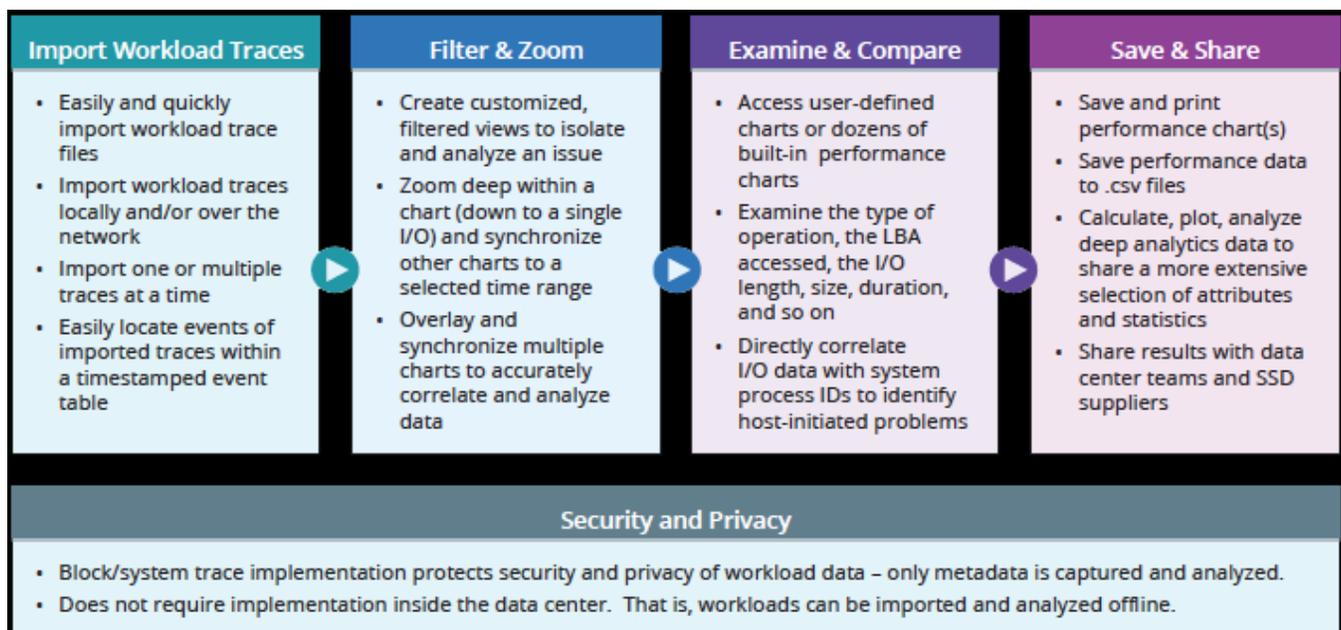


Figure 1 – WorkloadIntelligence™ Analytics Workflow

WorkloadIntelligence™ Analytics at Work – Real-world Example

With the WorkloadIntelligence™ Analytics application our team systematically analyzed a customer’s hyperscale Data Center workload for anomalies in an effort to characterize their workloads on the high-performance storage systems they were developing internally. We identified a number of problem areas in the workload and passed on our recommendations to our client which helped them improve the overall efficiency of their application and internally-developed high performance storage system.

Below are several of the charts generated by WorkloadIntelligence™ Analytics during our analysis. Figure 2 displays the process ID’s vs time. Our team was able to identify a “rogue” process that repeatedly started and stopped. Figure 3 highlights CPU vs latency. From this graph we were able to ascertain an inefficient utilization of the 56-core CPU on one of their systems. Eight of the CPUs were maxed out processing IOs while the others remained underutilized. Figure 4 highlights IOPs vs time and shows IOPs dropping significantly while trying to deal with the latency spike caused by the multi-second discards (trims) occurring on the solid-state drive. Finally, Figure 5 displays the large latency spikes caused by the discards.

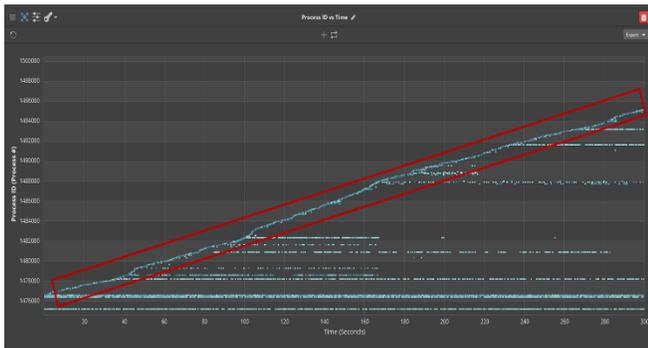


Figure 2 - Process ID vs Time

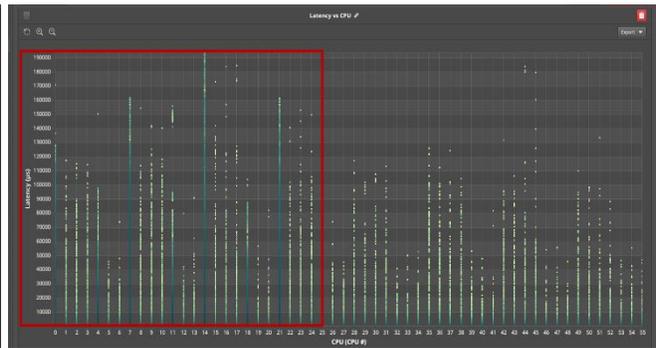


Figure 3 - CPU vs Latency Graph

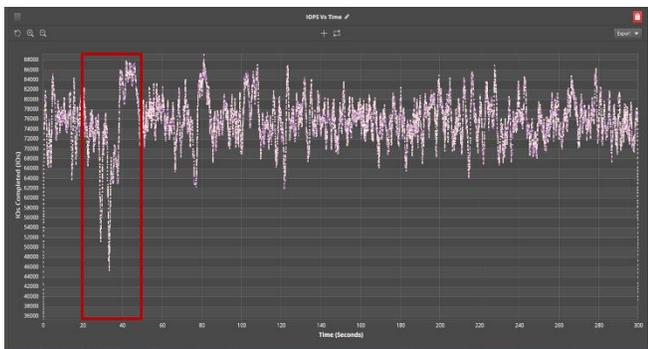


Figure 4 - IOPs vs Time Graph

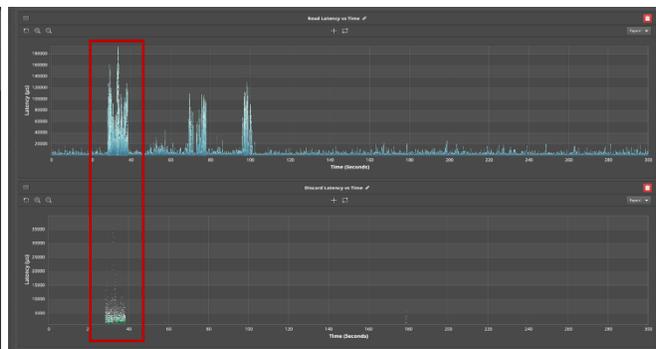


Figure 5 - Discards vs Time Correlated with Latency Graphs

Our recommendations included:

1. Reviewing the Linux OS Kworker process settings that initiated the multi-second Trims
2. Examining the additional processes that were kicked off by host at the time of the latency events

3. Investigating the rogue or unoptimized process to determine whether it's actually a rogue process or a process that could be doing more before being recycled
4. Reviewing the CPU core allocation with the application team to understand whether the cores could have been allocated more efficiently
5. Tuning the Application and Linux block layer to achieve better workload performance
6. Assessing a number of other enterprise solid-state drives to determine if they work better with the business's production workloads

Summary

For Data Center refresh planning, the characterization and analysis of production workloads is a critical factor for the business success of any Data Center operations team. As we stated earlier, the data explosion has created a new series of data-centric workloads to handle huge amounts of data being created and utilized each day. Additionally, the data is an increasingly blended rich media, which includes images, audio and video.

Given that the Data Centers hosting the online data and running these data-centric workloads are very much cost driven, it is important to understand how this emerging class of applications affects some of the design decisions in the Data Center. However, understanding real-world production workload profiles has been one of the most challenging barriers to system and storage performance, provisioning and validation.

WorkloadIntelligence™ Analytics helps engineers and operators analyze workload performance and make highly accurate and analytics-driven evaluations of their storage architecture to assure the performance of business-critical applications and data. Some of the world's largest Data Centers are just now deploying *WorkloadIntelligence*™ and discovering previously unknown, underperforming and poorly designed application workloads in their facilities.

Key Benefits:

- Shows key insights on how the Data Center or application platform performs
- Allows customer to make more informed hardware and application decisions that directly affect their users and customers
- Helps to deliver cost savings from better hardware utilization and process improvements
- Helps to improve the customer experience through system performance improvements
- Allows the Data Center staff to work directly with the system and storage vendors to optimize their hardware to help improve performance with the Data Center's workload

For more information and/or demonstration of *WorkloadIntelligence*™ and other Teledyne LeCroy products, visit <https://www.teledynelecroy.com/oakgate> or contact us by email at oakgate_marketing@teledyne.com