



Storage Workload Analysis

Performance analytics enables smarter storage acquisitions and reduces deployment risk

Introduction

Understanding the real application I/O workload profile has traditionally been one of the most challenging barriers to storage performance provisioning and validation. According to Howard Marks, founder of DeepStorage.net, “Most organizations know surprisingly little about the demands their applications present to their storage systems. Storage managers routinely specify systems with expensive performance headroom to account for their lack of workload knowledge”. In a 2015 survey of large enterprise storage engineers and architects, we found that 64% “don’t know their I/O profile”. Applications, networking, and storage teams rarely know their own workload profiles, and their vendors really aren’t in a position to help. Consequently, performance planning is guesswork, sometimes resulting in under provisioning, but more often, over provisioning. If workload performance planning is done at all, it’s nearly always done by trial and error.

Background

While many storage engineers have a pretty good understanding of their current performance (e.g. latency) from the perspective of their storage arrays, this doesn’t help them much when trying to size new arrays, or build new storage infrastructures, or effectively troubleshoot existing storage infrastructure problems. What’s missing is the analysis from the point-of-view of their applications, that is, workload analysis. Workload Analysis of production infrastructure is the ability to statistically analyze application workload data over a period of time and create workload models that can be used as a basis for storage performance planning and troubleshooting.

The ability to do this offers organizations many benefits, including:

- **Storage (right) sizing** – storage professionals are often asked to recommend and deploy new capacity with little or no idea of the performance requirements of the application workload(s). They often rely on rules of thumb and make best guesses, even when the cost differences between storage tiers can mean millions of dollars. Over-provisioning may provide a safe choice, but it clearly unnecessarily wastes valuable financial resources.
- **Storage migration** – before migrating workloads to new infrastructure, including the cloud, it’s important to be able to describe the profile of the applications (e.g. read/write mix, block/file size distributions, IOPS, throughput, latency, random/sequential mix, temporality, LUN activity, etc) in order to match the workload to the right tier of storage.
- **Storage consolidation** – as with migration, understanding the workload I/O profile is essential to avoiding surprises when consolidating workloads onto fewer target devices, as often happens in virtual server and private cloud environments.
- **Problem avoidance or resolution** – understanding how workload changes affect key performance indicators is often the first step in either avoiding performance bottlenecks or identifying potential hotspots and adhering to Service Level Agreements. Troubleshooting and resolving performance problems can be greatly accelerated with a better understanding of workload behavior.

Benefits

- Match your workload with the right tier of storage
- Reduce performance risk when virtualizing or moving to private cloud
- Avoid performance problems by understanding the effect of workload changes on application performance
- Make better purchase and deployment decisions via simplified storage performance planning

“*Using Load DynamiX has limited our downtime and gives us a huge cost benefit. We’re maximizing the use of the storage arrays and configuring them in the most efficient way.*”

Brian Walker

PRINCIPAL ARCHITECT

GE

Storage Performance Analytics Overview

The Load DynamiX Enterprise 5.0 release introduced Storage Performance Analytics, a capability of Load DynamiX Enterprise that allows storage engineers to analyze temporal and spatial workload behavior via powerful visualization to better understand workload I/O patterns that affect storage performance. The new Workload Data Importer feature of Load DynamiX Enterprise accesses historical production data from storage arrays at pre-determined time intervals, usually at a granularity of minutes. These workload models can then be applied to any file, block or object storage system via a Load DynamiX Workload Generation Appliance to fully evaluate its performance, to conduct what-if scenario analysis with modified workloads, or to efficiently troubleshoot performance problems.

Load DynamiX 5.0 offers storage teams an automated and simple way to acquire storage array and storage network workload data, and analyze that data for a holistic approach to architecting, optimizing, and troubleshooting storage infrastructure. The storage performance analytics capabilities within Load DynamiX Enterprise include new solutions for real-time and historical workload acquisition and analysis. The automated analysis aids in the creation of individual and composite storage workloads, useful for workload generation and realistic simulation of complex workloads in the lab, whether running on physical or virtual hosts. These solutions build upon traditional Load DynamiX strengths in storage workload modeling and workload generation.

Guide to the illustration on the next page, figure 1:

The illustration on the next page shows the Load DynamiX Workload Analyzer module of Load DynamiX Enterprise. Various KPIs like latency, throughput, or (as the illustration shows), IOPS over time can be reported. Such historical visualization can cover any time period. The Workload Analyzer includes easy to use analysis policies to see: IOPS, throughput, latency, R/W mix, random/sequential mix, block size distributions, temporality, and more.

The Access Pattern report presents several useful metrics, starting with a granular view of IOPS, Latency, and Throughput across several commands you see here in the index of the IOPS graph.

Below the real-time graph are a series of summary charts, including read/write ratio. These summary charts show the values you would get if a single constant workload is created. They are averages for the time period you select above.

The Request Size Distribution is all of the block sizes and their frequency of use. This graph updates every 30 seconds. You simply cannot get this level of data from legacy software-based polling tools.

“Load DynamiX is like Iometer on steroids! I can test 5 different storage arrays simultaneously, configure a base workload and test it, and then change the attributes and test again. I’m so much more efficient using Load DynamiX - I’ll never go back to freeware tools again.”

Todd Gleason, IT

INFRASTRUCTURE MANAGER
FIREHOST

“Load DynamiX made it easy to migrate to a private cloud. We were able to determine real-world behavior and save money. It doesn’t get better than that!”

SYSTEM OPERATIONS
DIRECTOR, ELLIE MAE

“ We can assess the hottest storage technologies like SSDs, caching, tiering, and de-dupe, against our full production requirements, faster and more accurately. Load DynamiX puts us in the driver’s seat when it comes to our storage roadmap and our cost structure.”

Justin Richardson
STORAGE ENGINEER
GO DADDY



Figure 1: Load DynamiX Workload Analyzer module of Load DynamiX Enterprise

The Command Mix graph tells us any command executed. This can even help in root cause analysis if unexpected commands are occurring. An example might be a ‘mode sense’ that bifurcates different latency figures.

The Latency vs. IOPS can establish the relationship between the two and by watching trends, you can often uncover problems. For instance, if in a workload the latency is going up at the point where IOPS are increasing, the system is overloaded.

Latency vs. Request Size gives a good indicator of the underlying block size of the array.

You can use this analysis to improve your ability to match your user/application requirements to storage infrastructure deployments. You no longer have to guess at the impact a workload will have on your storage. And beyond this first phase analysis, you can now build simulated models of your workload, and run synthetic workloads in the lab, against block, file or object storage, to determine optimal technologies, protocols, tiering models, vendors, products, and configurations.

In the example below, the storage team was challenged to select the best performing NAS array from two distinct vendor proposals, which were nearly identical in price. As the graph below shows, when two representative workloads were tested on two vendors NAS arrays, the workloads performed much better on Vendor B, when comparing latency. Not only was average latency lower, but it was much more predictable. You simply cannot get this level of understanding merely by reading the performance specs from a vendor datasheet.

Ask your vendor to provide this level of analysis in their Proof of Concept lab. Before the availability of Load DynamiX Enterprise, understanding and modeling existing workloads was a long and tedious process, often stretching POCs into many weeks or months. But today, there's no reason you can't do this in hours, and be much more intelligent about your primary storage deployment and configuration decisions. It all starts with workload analysis.

Conclusions

Simple to deploy and configure, Load DynamiX Enterprise presents a dynamic analysis & clustering of workload I/O behavior across time and locality, offering a powerful real-time or offline visualization & analytics solution that helps storage professionals:

- Make better purchase and deployment decisions via simplified storage performance planning,
- Avoid many performance related issues, and
- Resolve problems faster.

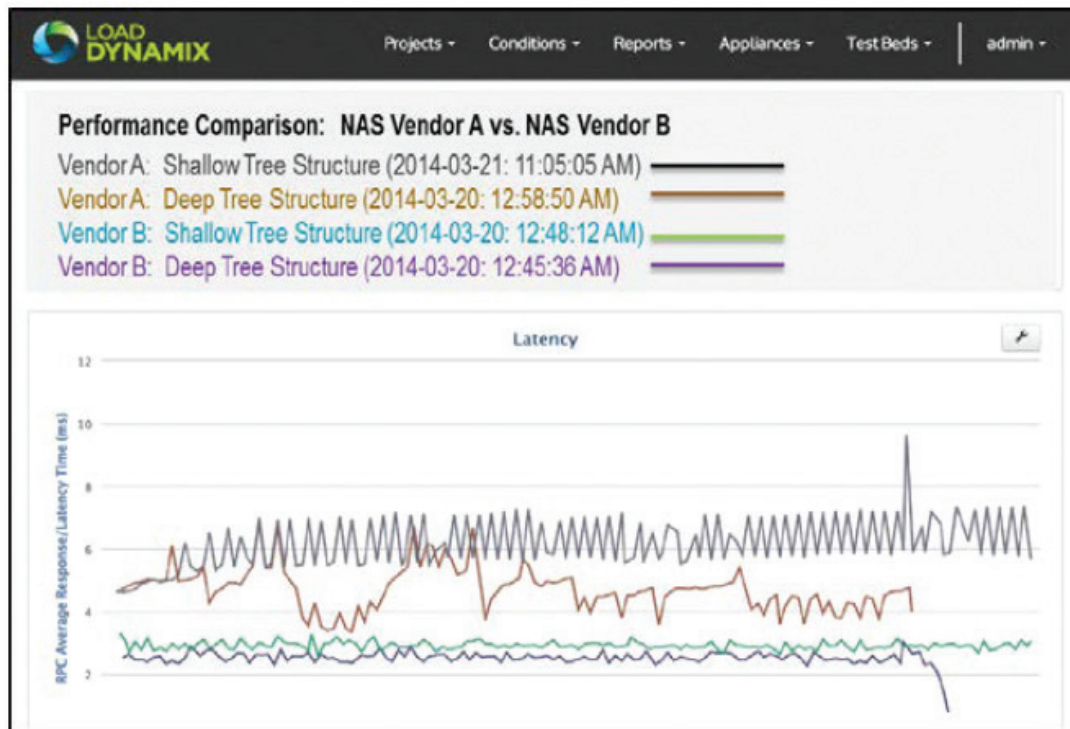


Figure 2: Storage performance comparisons