

# Storage Performance Testing: WorkloadWisdom Enables a New, Critical IT Best Practice

## Situation: The Storage Architecture Challenge


Storage is rapidly becoming the biggest IT infrastructure spending component and increasingly influences overall application performance. The need to maximize performance is key to customer satisfaction, sales efficiency, and worker productivity. With its larger influence, storage now represents the biggest source of application performance problems.

IT architects need to take advantage of the newest storage technologies while both optimizing performance and minimizing storage costs, which are often contradictory. They must have an accurate assessment of how a storage system or technology, like flash storage, will perform in their overall environment and determine the impact it will have on application performance. Vendors' benchmark data claims are not representative of real world application environments.

Identifying how storage will perform in a given environment is complicated due to the interaction between servers, network, and storage along with the wide diversity in application workloads. In the past, IT organizations would test "system level" performance with tools like HP's LoadRunner, but they rarely tested storage performance as a standalone process. Application-oriented performance tools are useful, but they cannot identify storage bottlenecks. The lack of meaningful storage performance information means that many organizations are flying blind when it comes to predicting how a given storage solution will influence an application's performance.

For those who do some level of storage load testing, a minimalist approach is to create performance workloads using freeware tools such as Iometer or Vdbench. These tools were created in the past when storage and applications were simpler and had much less at stake. Such tools require significant effort in planning and scripting to approximate an application's workload and aggregate results. Once the script and test plan is created, any changes made to the infrastructure make the scripts obsolete and necessitate a rewrite. In addition to the expense in personnel resources, the investment in the amount of hardware needed to be representative of the production environment is not typically financially or operationally feasible. The net effect is that relatively few useful storage tests are run.

The lack of accurate information means that storage architects end up guessing when it comes to how a given storage solution will affect application performance. Some application providers have tools to generate traffic and measure the performance with their products. These tools can provide some meaningful data on application performance, but these tools cannot characterize or measure other workloads, nor create test cases that identify the performance thresholds of storage products across multiple workloads. This limits the ability to predict storage problems as virtualized workloads increase output



to meet changing business demands. To compensate for the performance risk, many organizations choose to overprovision storage. This adds unnecessary expense as well as the need for additional space, cooling, power, and maintenance. This practice is not only expensive, but it is unknown as to how well it addresses performance issues.

## An Innovative New Solution


WorkloadWisdom, formerly known as Load DynamiX, provides a combined software and hardware performance testing and validation solution that includes storage workload modeling and load generation. With the ability to accurately emulate real-world application workload behavior, WorkloadWisdom enables storage engineers and architects to make intelligent deployment decisions regarding networked storage infrastructure. WorkloadWisdom has traditionally been used by the leading storage vendors in their performance engineering labs and increasingly in their regional competency centers for proof of concepts and benchmarking. IT organizations started to notice these devices, experienced the value of storage performance validation, and started buying the devices for their internal use. To better meet the needs of IT customers, Virtual Instruments developed and released the WorkloadWisdom solution. The primary objective with WorkloadWisdom is to deliver a performance validation toolset for IT organizations that makes it easier to model workloads, create test cases, and evaluate results.

The WorkloadWisdom solution consists of WorkloadWisdom software and the Performance Validation Appliance. The software is a user-friendly platform with a web-based graphical user interface (GUI) that is used to characterize and create workload models, configure and administer tests, and analyze results. WorkloadWisdom works in tandem with 2U load generation appliances.

The appliances are used to generate traffic based on workload models and access patterns that have been configured by WorkloadWisdom software. Supported protocols include NFS, SMB, CIFS, FC, iSCSI, HTTP, HTTPS, OpenStack, Swift, CDMI, and S3 protocols. This diversity of storage connectivity and protocols enables any file, block, or object storage system to be tested and evaluated, eliminating the guesswork that was storage planning.

WorkloadWisdom can create multidimensional workloads that quantify a storage array's capabilities and limitations. Testing can be executed to identify under what I/O profiles and IOPs a storage array's performance will degrade. A performance model can be created by extracting I/O profile data and statistics captured from storage arrays using open source and commonly available vendor tools, or by capturing actual production traffic using Wireshark that generates PCAPs. I/O profile data from either source can then be entered into the GUI of WorkloadWisdom.

Analysis output based on their reporting tools is very configurable and can be presented as a summary table, time-based graphs, or as a histogram. In the performance comparison example at right, vendor B shows significantly lower latency in both shallow and deep tree directory structures. The Workload Iterator is another output and provides a graphical tabular format, enabling true performance profiling—a deeper dive into dozens of metrics that identify the strengths and weaknesses of an array.



## Key Benefits and Use Cases

- **Storage Product Evaluation** – Storage products and configurations from different vendors can be evaluated with the same I/O profiles and traffic volume. The ability to quantify latencies and IOPs takes the performance guesswork out of making a product choice.
- **Validate Flash and Other Technologies** – Flash has proven to be effective in removing storage bottlenecks and increasing application performance. Not all applications benefit equally from flash; determining which specific applications will benefit will optimize the flash investment. As an example, a hybrid flash system may provide similar performance results to an all-flash system for a given application.
- **Optimize Storage Investment** – Under and over provisioning can be avoided by aligning workload requirements to deployment decisions. This ensures that storage resources are optimized and confirms at what point additional resources will benefit an application. Hot spots can be identified and addressed through workload simulation before production begins.
- **Deduplication & Compression** – Deduplication and compression can provide significant capacity and network cost savings. By supplying desired compression and deduplication ratios (see right) the impact of data reduction results can be verified and quantified. The processing required for data reduction can negatively affect system performance. WorkloadWisdom measures the performance at different data reduction rates, including with data reduction turned off. This assists in both the product choice as well as which data should be used for data reduction.
- **Change Management** – The effect that patches, firmware updates, OS updates, and hardware or process changes have on performance can be measured. This greatly reduces the risk of problems for the operations teams with making changes and enables periodic regression testing to be done.
- **Maximize IT Resources** – Performance testing and analysis can consume a significant amount of planning and execution time, as well as incremental hardware, software, and operational expense. WorkloadWisdom not only frees up staffing from this labor intensive work, but prevents problems before they happen and accelerates the time it takes to troubleshoot should they occur. WorkloadWisdom makes storage load testing economically viable for the first time; a true test lab in a box.

Data Parameters

Use  data reduction data content

Compression ratio 2.0 :1 percent 50%

Deduplication ratio 4.0 :1 percent 75%

Number of unique duplicates 100

## EMA Perspective

Because the storage staff may be oversubscribed, there is a temptation to depend on the storage vendor's word for how a solution will perform or rely on system level testing. Storage performance testing is too critical to both overall IT cost and performance to rely on such approaches. Quantifying how storage will perform with applications can greatly reduce risk, lower cost, and increase the satisfaction of internal and external customers. Based on experiences gained as the storage performance validation standard for storage vendors, the WorkloadWisdom solution leverages this DNA in an easy to use software platform with an intuitive web-based GUI. The WorkloadWisdom solution is capable of testing the performance of the largest physical, virtual, and cloud environments. The technology is exceptional, but it is the business benefits that provide the most benefit to the storage planners. The WorkloadWisdom solution lowers storage costs, mitigates business risk, and increases storage staff productivity while avoiding incremental lab expenses. These benefits contribute directly to a company's bottom line and justify the investment in the solution.

### About EMA

Founded in 1996, Enterprise Management Associates (EMA) is a leading industry analyst firm that provides deep insight across the full spectrum of IT and data management technologies. EMA analysts leverage a unique combination of practical experience, insight into industry best practices, and in-depth knowledge of current and planned vendor solutions to help EMA's clients achieve their goals. Learn more about EMA research, analysis, and consulting services for enterprise line of business users, IT professionals and IT vendors at [www.enterprisemanagement.com](http://www.enterprisemanagement.com) or [blogs.enterprisemanagement.com](http://blogs.enterprisemanagement.com). You can also follow EMA on [Twitter](#), [Facebook](#) or [LinkedIn](#). 3100.021115