

TANEJA GROUP

Storage Infrastructure Performance Validation

Load DynamiX Workload Analytics is Key to
Smarter IT

Mike Matchett

11/1/2013

STORAGE INFRASTRUCTURE PERFORMANCE VALIDATION

Load DynamiX Workload Analytics is Key to Smarter IT

NOVEMBER 2013



Unacceptably poor performance can be a career killer and so IT generally “over-provisions” infrastructure as the rule. But how much is this approach really costing us? Today, the biggest line item in IT infrastructure spending is storage. Even with data growth and new performance demands increasing, based on “safe” estimates we still overprovision by 50% or more which results in billions of dollars of wasted storage spending. A more important problem is that we may not be even provisioning the right infrastructure for our application workload requirements, taking serious risks with every new investment.

Equally vital is knowing when to upgrade or refresh. Looking forward, how can anyone know when their current infrastructure will hit its inevitable “wall”? In day-to-day operations, every time a change is made to storage infrastructure, the application or the network, that change could be introducing a deeply rooted problem that might only show up under production pressure. Why do enterprises seem to proceed blindly, willingly rolling the dice when it comes to performance? Here at Taneja Group, we see an obvious correlation between risk of failure and lack of knowledge about how infrastructure responds to each application workload.

Unfortunately, enterprises too often rely on vendor benchmarks produced under ideal conditions with carefully crafted workloads that don’t reflect the real target environment. Or they might choose readily scalable systems so that in times of trouble they can always buy and deploy more resources, although this can be highly disruptive and expensive when buying on a short notice. They might architect for large virtual and cloud environments in an attempt to average out utilization and pool excess capacity for peak demand, but still without knowing how performance will degrade at the upper reaches of VM density. In contrast, we believe that IT managers must evolve from a perspective of *assuming* performance, to one of *assuring* performance.

Typical testing approaches usually involve generating workloads with heavily scripted servers used as load generators. This is an expensive, unreliable, brute force approach, only trotted out when sufficient staff, time and money is available to execute a large-scale performance evaluation. But Load DynamiX has changed that equation for storage, evolving workload modeling and performance load testing into a cost-efficient and practical continuous process. We think that Load DynamiX’s solution supports the adoption of a new best practice of proactively managing infrastructure from a position of knowledge called *Infrastructure Performance Validation (IPV)*.

In this report we will look at Load DynamiX’s workload modeling software and storage performance validation appliances, and walk through how IT can use them to establish effective IPV practices across the entire IT infrastructure lifecycle. We’ll examine why existing approaches to storage performance evaluation fall short and why we believe that successful storage deployments require a detailed understanding of application workload behavior. We’ll briefly review Load DynamiX’s

solution to see how it addresses these challenges and uniquely enables broad adoption of IPV to the benefit of both the business and IT. We'll look at how Load DynamiX generates accurate workload models for storage testing, a key IPV capability, and how limit testing and "what if" scenarios can be run, analyzed, and communicated for high impact. Finally, we'll look at a range of validation scenarios, and how Load DynamiX can be leveraged to reduce risk, assure performance, and lower IT costs.

WHY ENTERPRISES ARE OFTEN FLYING BLIND

The best approach to making performance-impacting decisions in IT comes from a thorough knowledge of an infrastructure's *performance curve* under various levels of load, all the way up until it inevitably degrades beyond usefulness – its breaking point or "where it hits the wall". This is true whether the infrastructure in question is currently deployed or is being evaluated for future implementation.

Because of the complex interplay between servers, networks, and storage solutions, and the vagaries of real-world production workloads (i.e. all the applications, utilities, and miscellaneous "stuff" that runs in production), the most accurate, most reliable way to understand storage system performance is to stress test the system as a whole. But there are a number of issues and challenges with this type of thorough testing in practice today, including:

1. **Not Enough Time** – With current approaches to testing, it can be a major effort to acquire appropriate equipment, setup a desired test (usually involving custom scripting), run the test, collect the results, analyze them, and then distribute actionable results. We often see multi-week if not multi-month test cycles in enterprises on key storage decisions, which directly impacts the competitive agility of those organizations.
2. **Limited Tools** – Common testing tools are severely deficient in meeting the performance planning needs of storage engineers and architects who should have tools that are designed specifically for storage architects. Fully understanding which storage products and technologies would be best for a specific application workload requires a detailed understanding of the I/O profile associated with each workload. Tools need to be designed to help the storage architecture team create realistic storage workload models, create meaningful tests, execute them, and most importantly, interpret results.
3. **Inaccurate Tests** – Unfortunately well-known freeware storage load testing utilities (e.g. Iometer) require a lot of hand-holding and scripting in order to get beyond the simple read/write benchmark kind of testing. Glaringly, these utilities don't simulate metadata commands, which can easily represent more than half of a workload's command mix. Workloads being tested need to reflect the actual workload in production, or as closely as expected to a future state. They also need to reflect the directory structure and other load parameters. Running a large simulation that requires lots of load generating servers and collating results requires custom test harness development and support. Getting it less than perfect can unintentionally skew results, making them not just wrong, but dangerously misleading.
4. **Limited Scale** – Load testing must stress storage systems to worse case production-scale levels. Current solutions cannot produce sufficient load to achieve this level of scale. Often a response to limited load generation scale is to considerably "downsize" the system under test, then extrapolate back up introducing significant errors of approximation, and usually missing key behavior at realistic scales. Findings from such approaches are not very useful.
5. **Not Easily Repeatable** – Most load tests are conducted once or can be repeated as setup, but not easily reused down the road for subsequent testing cycles – either for regression or

shared with other test clients. Consistency of results can be a significant challenge, especially when making 'apples to apples' comparisons over time.

It's no wonder that we often see critical performance evaluations either ignored or relegated to special event status in IT, despite the continuous decision-making and activity that would be well informed by accurate and timely test results. Certainly new storage evaluations and technology refresh projects are obvious candidates to benefit from reliable performance validation, but so are more common processes like change management, architecture planning, ongoing optimization, and even daily troubleshooting. Basically, what the market needs is a way to practically conduct comprehensive storage performance validation to produce insightful answers to questions about infrastructure behavior in any IT decision-making project or change management process.

Infrastructure Performance Validation with Load DynamiX

Given the large and steadily increasing investments in storage infrastructure, IT needs to have ready access to information that will help them make the right decisions at the right time. Having hard data that a \$1M hybrid storage system will perform just as well as a \$2M all-flash storage system on your specific workloads helps directly support cutting storage CAPEX by 50%. Having the knowledge and confidence that the firmware changes and software updates you are about to make won't degrade performance when deployed into production lowers risk, accelerates productivity, and assures service levels. We could attempt to present a detailed TCO or ROI calculation, but we think that the business benefits of knowing that your deployed infrastructure can handle your constantly changing workloads is obvious given the typical annual multi-million dollar storage spend.

Fortunately, Load DynamiX saw the opportunity to help enterprises implement a comprehensive and continuously practical type of infrastructure performance validation with a purpose-built IPV solution.

To address those challenges, the Load DynamiX solution combines an easy-to-execute, end-to-end workload modeling and performance validation application with a high-powered load generation appliance capable of generating massive loads that can drive storage systems past their maximum capabilities. This solution starts with workload analysis and supports automated generation of test workloads statistically derived from observing live production environments via passive network trace files or standard I/O statistics that are output as part of every storage system.

IT architects and storage engineers will find Load DynamiX performance validation solutions indispensable. Broadly, the Load DynamiX approach to IPV enables:

1. **Intelligent Vendor and Technology Selection** – Load DynamiX helps create an accurate model of existing and future workloads that can be run at extreme scales. Potential solutions from competing vendors can be analyzed to determine which will perform best and deliver the most cost-effective service under workloads that represent your production applications. It helps determine the performance impacts from implementing new solid state, deduplication and compression features – before they get deployed into production. Load DynamiX provides independent validation for all those vendor performance claims ensuring investments will deliver as promised, and also ensures IT organizations don't over-provision and over-buy.
2. **Eliminating Problems** – Load DynamiX can be used to simulate the exact workloads in a lab environment that might be causing performance degradation in your production environment. Performance validation can help troubleshoot thorny issues, allowing multiple replays while different remediation solutions are attempted.
3. **Change Management** – Load DynamiX makes it easy to conduct regression testing from a baseline. With the rapid pace of technology changes, storage system updates, and protocol

evolution (e.g. NFS v3 to v4.x upgrade) landing on storage departments, it's more important than ever to be able to validate changes before deployment into production. We believe Load DynamiX will significantly improve an organization's agility and time-to-implement improved services.

4. **Optimizing Investments** – Like change management, having an accurate and easy way to run repeatable test scenarios in Load DynamiX enables trialing out potential optimizations and architectural improvements in a proactive manner. This can be very useful in environments where the storage or networking solutions have lots of knobs and switches for configuration and optimization. In particular, the more the world moves to “software defined” resources, the more it will be important to know how to optimally “define” the resource.
5. **Determining Performance Limits** – Because emulated workloads can easily be scaled up in Load DynamiX, infrastructure can be stressed to its breaking point. Workloads can also be further varied, combined, and aggregated to create tests that represent estimated future workload conditions. It's now possible to realistically predict the impact of future architectural and product alternatives or how server and storage virtualization might impact response times.

LOAD DYNAMIX

When Load DynamiX looked at the existing state of the art in performance validation, they saw an unmet need to help enterprises test at scale, with accurate workloads reflecting actual production. They designed a solution that wouldn't require specialist expertise nor an expensively large scaled-out lab infrastructure. Load DynamiX was intentionally architected to be practical and cost-effective at scale, accessible by many teams for multiple purposes and types of validation including continuous IPV processes.

As storage vendors offer increasingly unified and multi-protocol arrays, performance must also be validated across the same broad mix of storage solutions. And as unified storage in production may be hit with traffic on multiple protocols and interfaces at the same time, validation must also support those same workload mixes. Load DynamiX validates all kinds of storage infrastructure – file, block, and object – over Fibre Channel, iSCSI, NFS, SMB, CDMI, OpenStack Swift, and/or HTTP as required in the same validation test.

Load DynamiX Architecture

Load DynamiX consists of components that correspond with the three main steps of infrastructure performance validation. The main user interface is the Load DynamiX Workload Insight Manager through which most of the following activities are conducted:

1. **Analyze/Characterize Workloads** – Workload Insight Manager and its Workload Analyzer utility help characterize and create accurate workloads from statistics and network trace files collected in the actual production environment.
2. **Model Workloads** – Load DynamiX then helps the user build validation models, maintain libraries of repeatable scenarios and results, schedule validation scenarios for execution, and analyze the results visually.
3. **Emulate Workloads At Scale** – Load DynamiX's Load Generation Appliances are purpose built devices and software that have been optimized to generate and drive high scale storage workloads across a wide variety of protocols and interfaces.

1. CHARACTERIZE WORKLOADS

To fully enable infrastructure performance validation, the first step is to understand current application workloads. Load DynamiX offers the ability to model production workloads based on statistics captured from storage system logs or Wireshark captures of the live operational environment. This extracted data can then be run through the Workload Analyzer utility, which helps dig into the command mix, load patterns, and application data flows with its graphical IO profiling. Workload Analyzer provides a new perspective on how applications interact with the storage infrastructure. There is also an option to be able to directly replay captured workload traces in order to recreate live scenarios for the purpose of troubleshooting, problem remediation, or storage optimization.

Finally, Load DynamiX offers a growing library of pre-defined application workload models that can be applied out of the box. Given that workloads can be replayed as a trace, intelligently abstracted from statistics, or selected from a pre-defined library, Load DynamiX can be used in a myriad of ways to build insight into how storage infrastructure performs under real application scenarios.

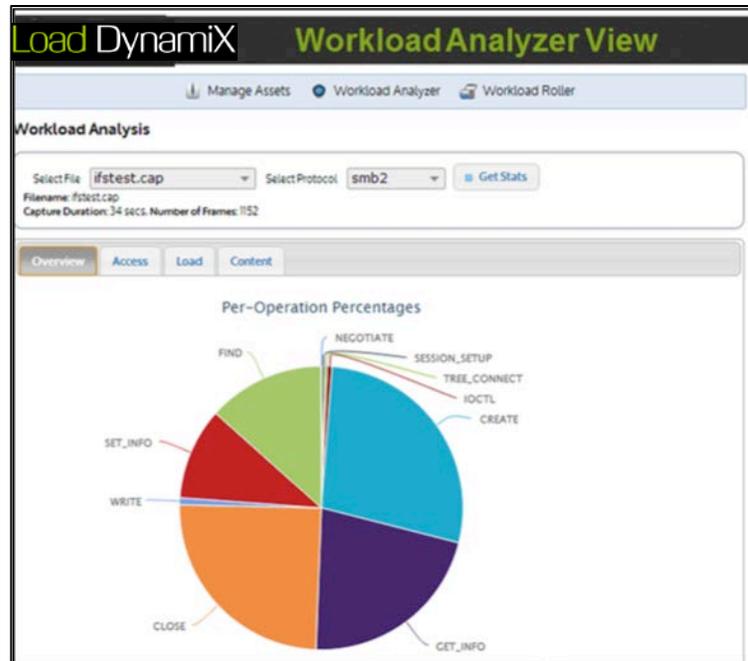


Figure 1 - Load DynamiX Workload Analyzer illustrates how applications interact with the storage infrastructure

2. MODEL WORKLOADS

Load DynamiX supports a simple way to model any workload in further detail within its Workload Insight Manager. “What-if” kinds of scenarios around consolidation, workload growth or infrastructure change can easily be constructed and analyzed.

Load DynamiX provides an ability to edit these workload models at a very detailed level if desired (metadata commands, background processes, overhead calls, access patterns, et.al.), even creating and building workloads up from scratch. While editing at this level might require domain expertise, there are advanced, valuable use cases for validating performance for workloads that don’t exist or might

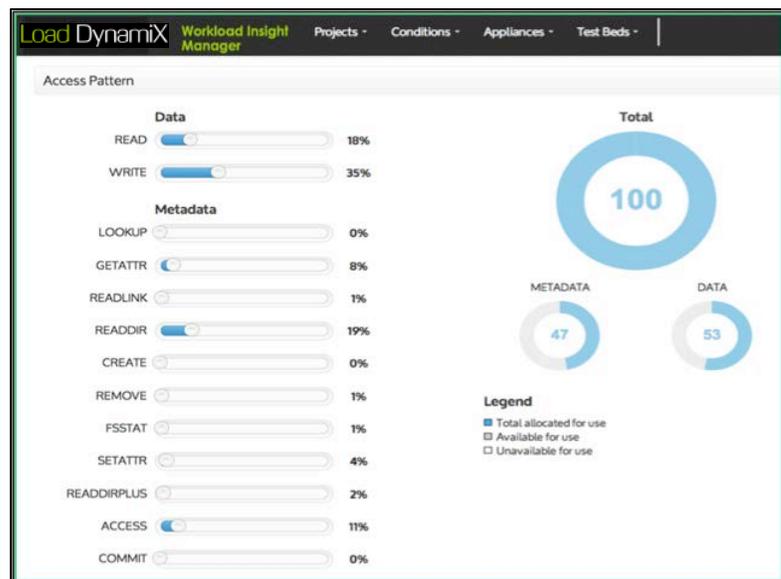


Figure 2 - Load DynamiX Workload Insight Manager enables detailed modeling of storage workloads, including metadata, that freeware options can’t provide

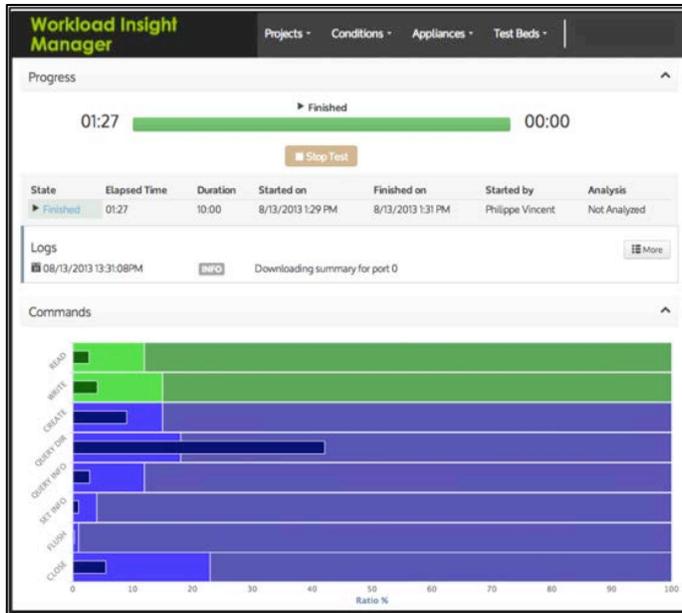


Figure 3 - Load DynamiX Test Run Dashboard

Load DynamiX also generates VAAI commands to emulate VMware optimizations, navigates VLANs, supports IPv6, and speaks a range of authentication protocols. In addition, Load DynamiX provides a User Parameter function to help scale tests with realistic end-user patterns of access and authentication.

As shown in Figure 2, Load DynamiX makes it easy to edit the IO command mix, including control over the percent and mix of IO metadata overhead. This is a major weakness in freeware and seat-of-the-pants scripting approaches as metadata traffic can easily be 50% or more of the overall loading on storage systems or create a controller bottleneck (just consider the work a handful of VAAI primitives can kickoff in an array).

The Load DynamiX Workload Insight Manager provides an online validation environment in which scenarios can be constructed, saved in a library, edited, executed, and analyzed.

Results can be compared to baseline, stored, and shared within the Load DynamiX environment. The interface design supports inter-disciplinary validation, and broad dissemination and leverage of both workload/test creation efforts and performance validation execution results.

3. EMULATE WORKLOADS AT SCALE

Load DynamiX offers a range of performance validation load generating appliances with varying network connectivity. These appliances have been highly optimized, including a completely re-written TCP/IP implementation, for

only be theoretical.

Load DynamiX validates a wide range of storage workload types that include:

- **File** – CIFS/SMB 2/SMB 3, NFS 2-4.1. Test scenarios can respond to NFS error code responses. Advanced file load capabilities include pNFS and NFS 4.1 threading.
- **Block** – Fibre Channel with P2P or AL, and NPIV support, iSCSI with DCB, and fine grained SCSI controls including streaming to emulate tape clients (VTL).
- **Object** – HTTP REST commands including dynamic form submissions and encrypted HTTPS. CDMI, OpenStack Swift, and AWS S3 validations also supported.

Load DynamiX also generates VAAI



Figure 4 – Workload Insight Manager view of validation results and analysis

intense load generation. Due to the purpose-built approach, a single appliance can replace dozens of high-end load generating servers. This directly leads not only to significant lab CAPEX and OPEX savings, but also enables massively scaling up the maximum lab test capability without custom scripting and synchronization concerns.

Load DynamiX appliances are currently available with 8 1GbE ports, up to 8 10GbE ports, or 2 4/8Gb Fibre Channel ports. In addition to supporting 16Gb FC and eventually 40GbE, we expect the number of ports on each appliance to increase significantly in early 2014. We imagine that organizations will find that a small combination of Load DynamiX appliances addressing the range of deployed storage architectures will prove incredibly valuable while saving large swaths of rack space.



Figure 5 - Load DynamiX 10G Series Appliance (left), Load DynamiX Fibre Channel Series Appliance (right)

Load DynamiX appliances can be aggregated together for larger test scenarios, and can also be shared on a port-by-port basis in a multi-tenant fashion between multiple users - even at remote sites. This many-many user-to-appliance design provides a flexible and agile testing capability that enables productive, cost-efficient performance validation for large organizations with many varied validation needs.

LOAD DYNAMIX FOR BEST PRACTICE IPV

Once Load DynamiX is deployed, we think it will quickly become one of those tools that IT storage architects and engineers simply can't live without. Specific Infrastructure Performance Validation tasks for storage might include:

- **Storage array or network upgrades** – Before patches or upgrades are rolled into production, Load DynamiX can validate that the changes won't break the system or degrade performance.
- **Capacity or performance expansion** – Before adding new disks, shelves, scale-out nodes, additional controllers, or other hardware intended to expand capacity or boost performance, validation can prove (or disprove) that the investment will have the desired effect.
- **New technologies** – Given the pace of new storage technologies that vendors are introducing like flash, storage virtualization, "auto"-tiering, deduplication, compression, or replication that promise to improve performance, or at least not degrade it, IT needs more than ever to be able to validate new technologies before deployment (if not before investment) to understand what's really going to happen. Often a "simple" new capability has unexpected consequences.
- **Storage Consolidation** – Data center consolidation is an ongoing project at many enterprises, and one of the main opportunities as well as pain points in consolidating storage. Load DynamiX based validation based on actual production workloads could prove invaluable to assuring performance as consolidation activities progress.
- **Cloud Migration** – Similarly, projects to migrate storage services into private clouds could proceed faster and with less risk to performance if IT could assure performance to users through performance validation.

- **DR Scenarios** – Load DynamiX can be used to test DR failover conditions, such as performance under failover/failback. It can also validate what happens when backups occur or runover, and what happens to applications when and if recovery activities kick in.
- **VDI Storage** – What is the optimal number of virtual desktops that can boot on a given storage infrastructure? Load DynamiX can help establish the performance versus loading curve, all the way up to where boot storms would swamp any given storage solution.
- **Load Balancers** – Load DynamiX can validate that load balancing schemes actually balance, what levels of traffic they can handle before degrading, and then what happens when they become overwhelmed.

Simply adding up the expected opportunities for validation listed above could very well lead to productive validation exercises on a weekly if not daily basis. With a cost-effective performance validation solution like Load DynamiX that supports quick execution while providing detailed actionable information, there is now plenty of justification for adopting an “validate first” approach when dealing with any kind of change in architecture, infrastructure, or observed performance.

Test For Tomorrow

To take full advantage of the opportunity presented by cost-effective and accurate performance validation, organizations should establish an ongoing practice dedicated to addressing important questions about their storage infrastructure. Establishing a central lab running Load DynamiX makes sense to centralize validation best practices and ensure wide sharing of validation results.

For example, most organizations approach storage optimization piecemeal and haphazardly, only taking time to seek “just enough” improvement when performance (or budget) becomes constrained. By employing Load DynamiX proactively, both IT and business stakeholders will benefit from recurring strategic optimization experiments.

A recommended question for IT storage shops might be “How do I leverage lower cost storage or commodity platforms?” Load DynamiX can help address an increasing data deluge by helping evaluate alternative architectures and solutions, validating lower cost designs for practical implementation.

During any evaluation/validation, instead of testing just one level of simple workload as with other tools, Load DynamiX can readily be used to establish full price/performance curves for the given infrastructure based on the organization’s own application usage profiles. These curves can be maintained as valuable design references for the evolution or expansion of either the infrastructure or the application workload.

Infrastructure Performance Validation Lifecycle

As a system management process, IPV is organized around the lifecycle of an enterprise’s infrastructure. At each step in the infrastructure’s lifecycle, performance validation should be used to produce reliable, accurate decision-making information. For example, proposed changes to infrastructure should be first validated before rolled into production.

Performance Validation for the Entire Infrastructure Lifecycle



As infrastructure management is a continuous process, IPV should also be one for maximum benefit.

Load DynamiX In Summary

As we've seen, there are many things about Load DynamiX solutions that differentiate it from other performance validation approaches, but here are a few key points we'll summarize:

- **High Performance** – One Load DynamiX appliance can generate storage workloads at scales equivalent to dozens of servers running scripts. They can find the performance limits of any storage system. This enables high load testing where it wasn't feasible before and reduces the costs of building and operating a large lab of test equipment.
- **Practical, Accurate Workloads** – Load DynamiX workload modeling incorporates both the data and metadata commands in real I/O storage communications, in the mixes as observed in actual usage, providing accurate emulated workloads that enable reliable validation results.
- **Decision-Making Speed** – As an IPV solution, Load DynamiX workload analytics provides actionable information in a quick timeframe that supports intelligent and timely design, investment, troubleshooting, and change management decisions.

TANEJA GROUP OPINION

The speed of change within the storage industry is accelerating. There are more changes, new technologies, increasing data and performance-based competition showing up every single day. Storage organizations have long been slow to adapt and evolve, but they now desperately need to become more agile and intelligent about their infrastructure. And to do that they need more information and knowledge about how to architect, manage, and proactively control an evolving and ever-changing storage infrastructure.

Load DynamiX represents an opportunity to roll-out more services, with more agility and less risk. Imagine characterizing and emulating your live production storage workloads for use in test and development at any time of your choosing. Imagine being able to replay exact data and metadata sequences on block, file and object storage protocols for troubleshooting and optimization. Imagine being able to quickly validate any proposed change in your storage infrastructure from application workload, server, protocol, network, to storage array before making the changes in production.

The power in knowing that proposed changes, application upgrades, and new architectures will perform as desired before you commit, and producing that information quickly and cost-effectively, is why we think Load DynamiX will see quick adoption by IT enterprises. Load DynamiX is enabling continuous storage IPV to go from a "wish I could" to a "must have".

Since 2010, nearly all of the leading storage system and networking vendors have used Load DynamiX in development, test, and QA departments. That is where the company started. It's definitely time that IT enterprises are similarly "armed", particularly as production storage workloads are unique to each enterprise. Understanding how application workloads interact with the infrastructure is now essential knowledge for an IT architect.

Using Load DynamiX enables 'Smart IT'. It allows the storage architect to make intelligent technology and purchase decisions based on real data that reflects actual workloads. In a world where IT is increasingly being outsourced, Load DynamiX brings job security because it enables internal IT to proactively optimize infrastructure investments – existing and new – while assuring service to each business-critical workload.

Bottom line, Load DynamiX removes any reason for not adopting the best practices of storage Infrastructure Performance Validation. It is eminently affordable in relation to the storage budget it helps optimize, produces highly accurate realistic workloads, and scales to any level. Most of all, Load DynamiX enables IT make decisions from a position of knowledge rather than risk. You'll have peace

of mind knowing how storage technologies actually work in your environment and where the limits of your infrastructure are hiding.

NOTICE: The information and product recommendations made by the TANEJA GROUP are based upon public information and sources and may also include personal opinions both of the TANEJA GROUP and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. The TANEJA GROUP, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors that may appear in this document.

TANEJA GROUP