

Accelerating Desktop Virtualization: Visibility and Insight with Liquidware Labs' Stratusphere™

August 2009



Desktop virtualization is shaping up to be the next frontier in the evolution of widespread virtualization. The major server virtualization platform vendors (VMware in particular) are all expected to ramp up investment and development in the segment significantly over the next 18 months, and continue to innovate to expand the possibilities for desktop virtualization beyond the simplest use cases.

In our view, three thorny issues have stymied desktop virtualization: costs related to server and storage consolidation, performance issues at the endpoint, and management challenges. The first of these is the focal point for a range of new virtual desktop infrastructure (VDI) platform enhancements which dramatically reduce the storage footprint for virtual desktop deployments of virtually unlimited size. VMware View Composer, a feature of VMware View, offers the ability to create hundreds of linked “clone” virtual machines, each taking up a fraction of the storage space of a traditional virtual machine, and to manage pools of these clones, in various flavors, with a few mouse clicks. It is also now possible to consolidate both persistent and non-persistent virtual desktops, with compelling data center storage and server savings that were out of reach just a year ago.

The trickier challenge for most enterprises exploring desktop virtualization, however, is to determine how to offer users equivalent or better performance and application experience with a hosted virtual desktop, and how to manage that user experience without growing administration cost or headcount. Liquidware Labs, a young entrant in the desktop virtualization “assessment and readiness” field, has developed a novel approach that combines process and tools with leading-edge technology. The company’s Stratusphere product and related services promises to shorten the path to desktop virtualization, providing insight into what users are actually doing every day on their desktops and how that behavior can be modeled, base-lined, and maintained in a virtual world.

The Desktop Assessment and Readiness Challenge

The traditional ‘fat’ desktop has become, in too many cases, a bloated wild zone: user-installed applications, un-patched operating systems, out-of-date security tools, and open-access browsing. As long as corporate

applications perform well enough on the existing desktop, however, it’s often easier to leave them alone. But when it comes time to evaluate the readiness of a community of desktops for a VDI conversion, this inconsistency can make it nearly impossible to establish key metrics, minimum performance benchmarks, or acceptable

S O L U T I O N P R O F I L E

service-level thresholds. Which end users could function with a minimal, locked-down non-persistent virtual desktop? Which demand higher graphics performance, or tend to run multiple applications at once? Which groups of users are similar enough to benefit from a shared package of default applications, or might need dedicated file storage? How much memory and CPU resource does each user *need*, regardless of how much they *have*?

These questions are similar to those asked during server virtualization projects, with some key differences. First, desktop virtualization consolidation ratios are ideally much higher than server consolidation ratios – they have to be for the storage and server economics to work – therefore, a greater number of virtual desktops (end users) are contending for limited server and disk resources and are doing so across a larger number of network types: LAN, WAN, public network, and more. Also, server workloads are more predictable and easier to measure, in large part because mature performance management tools exist. Finally, while server workloads experience spikes and troughs caused by fluctuating user demand, the number and type of workloads running on a particular server are under the constant control of an administrator. On the desktop there is no equivalent control; at a given point in time, any particular mix of end-user applications (approved or not) generates a unique set of resource constraints, bottlenecks, and performance problems. Without tools, or even a clear set of established metrics for desktop performance and user experience, answering the

assessment and readiness questions for a VDI project requires some new thinking.

Liquidware Labs

Liquidware Labs (LWL) was formed to help customers shorten the VDI implementation cycle – Assess, Design, Deploy, Manage – and has come out of the gate with solutions focused on the first two phases of the cycle: Assess and Design. LWL combines software, services and training to deliver visibility into a customer's environment. The industry has learned – from server virtualization as well as early VDI projects – that CPU and memory requirements are only the baseline, minimum metrics needed to size a virtual machine; for a complete picture of performance and user experience, a much broader and deeper set of *correlated* metrics are required.

The company has a unique pedigree. Its management team combines expertise in VDI planning methodology from the leading VMware integrator (now VMware subsidiary) Foedus, with a technology platform from vmSight (a leader in virtualization performance technology) and the software tools business savvy of the former CEO and Founder of Vizioncore. We expect this team to deliver tools that push the envelope on performance measurement and deliver real user experience insight, are easy to acquire and customize, form an integral part of a comprehensive services methodology, and are available from leading integrators.

S O L U T I O N P R O F I L E

The Keys to Successful VDI

Successful VDI deployments require tools that support a comprehensive assessment and design project. To start, baseline measurements must be taken over multiple time intervals to track resource utilization by user and application runtime behavior, determine environment needs and limitations, identify user types, and categorize users and applications into manageable subsets. Measurement tools should be easy to install, provide auto-discovery and require minimal configuration, be non-invasive, and be easy to remove. Results should span multiple application tiers (server, storage, network), but present correlated findings simply, with clear identification of target users and use cases and the performance profiles of each.

With these capabilities in place, an iterative data-analysis methodology should be employed. First, gather a complete picture of the desktops in place over various collection time windows. Next, synthesize the data into a classification scheme that estimates how 'fit' each user and application type is for virtualization. Then, repeat the data collection process and compare observed results with the scheme, adjust classification parameters as needed, measure again, and repeat until you've reached an acceptable confidence level. While this methodology delivers proven results, it can be very difficult to implement from scratch. Choosing the types of data to gather, prioritizing and collating, and building a classification scheme can be daunting; these challenges form the basis of the LWL solution.

The Stratusphere™ Solution

Stratusphere, the software element of LWL's solution, targets the keys to successful VDI assessment and design:


Data Collection: Stratusphere collects configuration and usage information from the existing physical environment via a small, dissolvable agent deployed to each desktop to be analyzed. The same agent technology can also be used to monitor on-going performance after a VDI deployment. The agents collect and report configuration data to a central Stratusphere Hub data store (or CMDB, which can be deployed as a virtual appliance) and employ patented "Connector ID Key" technology to identify and track IP packet traffic to monitor network performance metrics down the individual application, user, and desktop.

Stratusphere collects a broad range of data by default, including configuration (CPU, memory, disk, network cards, monitors, peripherals, time in use, and more); application and user/group inventory; performance (CPU, memory, disk); logon times, application load times, graphics utilization; network latency; and more. Stratusphere supports collection of metrics, which are combined into an overall composite metric, as well as an extensive and complete collection of individual key performance indicators relating directly to VDI and end user experience. Collection periods are fully configurable.

S O L U T I O N P R O F I L E

VDI Classification: Data collection is just a starting point. The challenge is correlating and presenting the collected data in a way that quickly helps develop a useful classification system. Stratusphere is designed to simplify the classification process, by providing canned data views and reports, along with full access to the

underlying data for more advanced users. For configuration and resource utilization classification, Stratusphere provides a configuration cluster view as well as a resource utilization view; these enable quick assessment of the range of device types and usage patterns across the measured inventory.

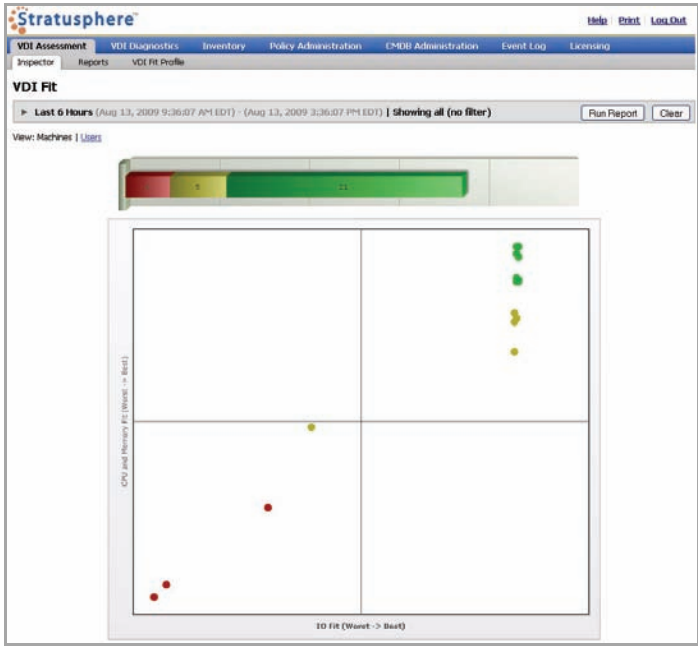


Stratusphere's Configuration Clusters and Resource Utilization Views provide rapid, graphical assessment of large datasets via scatter plots and bar charts.

Detailed resource utilization, by machine, user, and application, is then fed into two predefined classification systems – the 'secret sauce' and key differentiator of Stratusphere. Both systems are based on algorithms that analyze multiple resource utilization dimensions and classify each in terms of its impact on the overall readiness, or 'fitness,' of a user for virtualization. The automated process simplifies analysis by 'categorizing,' or sub-classifying each data element into classes, such as "Good," "Fair," and "Poor."

VDI Fit™: Multiple data elements are collected and classified, based on LWL's best-practice experience in the field and can be adjusted by the virtual architect, then given weights and combined into an overall rating. The first rating, **VDI Fit**, generates a quadrant view of Best and Worst Fits based on the Good, Fair and Poor binning of multiple data elements, or Key Indicators. This view provides a clear, high-level classification view, hiding the complexity of the underlying data while enabling drill-down into specifics for each user desktop.

S O L U T I O N P R O F I L E



VDI Fit scatter plot for a population of VDI candidate users


Key Indicators for **VDI Fit** are equally weighted between CPU and Memory Indicators (a set of four, also equally weighted) and Network I/O Indicators (a set of three, equally weighted), as shown below. The Good, Fair, and Poor classifications are defined by default values provided by Stratusphere, which may be adjusted by an administrator. For example, as delivered, the Key Indicator Network Latency is classified as follows by default:

- Good: less than 40ms
- Fair: between 40ms and 140ms
- Poor: greater than 140ms.

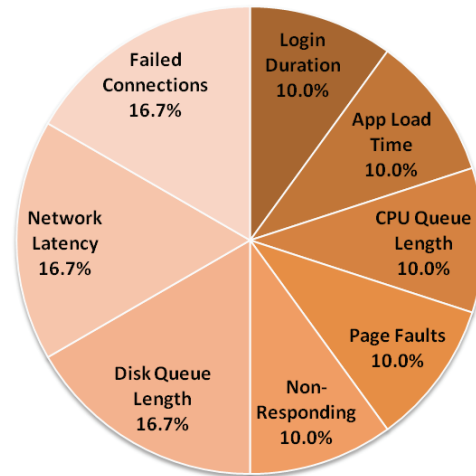
VDI UX™: The VDI UX (User Experience) classification is also created by codifying LWL’s field expertise and best practices into the Stratusphere Hub. The goal in this case is to assist with both pre-virtualization classification and post-virtualization monitoring for user interaction experience. The **VDI UX** classification is equally weighted between Key Indicators for Machine Experience (Login Duration, Application Load Time, CPU Queue Length, Page Faults, Non-Responding Applications) and I/O Experience (Failed Connections, Network Latency, and Disk Queue Length).

S O L U T I O N P R O F I L E

VDI Fit Key Indicators



VDI UX Key Indicators



The default weightings for each Key Indicator in the VDI Fit and VDI UX classifications.

Model Refinement: The VDI Fit and VDI UX classifications provide an out-of-the-box starting point for assessment, but they are easily tweaked and refined by an administrator with little training required. Ongoing model refinement is based on feedback, gathered by collecting additional Key Indicator data and comparing it to Stratusphere’s earlier Best Fit analyses; controls, which enable tweaking of the classification parameters and weightings; and overrides, which allow the user to override the classification system for a particular user or application which is known to exhibit outlier behavior. These model refinement capabilities extend Stratusphere’s usefulness beyond the Assess and Design phases of VDI implementation

and into the Deploy and Manage phases, where on-going monitoring and quick identification of performance problems is critical.

Taneja Group Opinion

Application performance management tools exist in several flavors, and generally suffer from being either too specific (to a particular transaction type, for example) or too general, flooding the user with data collected along several dimensions and across application and infrastructure tiers. And none that we’re aware of have targeted the virtual desktop infrastructure specifically.

S O L U T I O N P R O F I L E

LWL has staked out a unique position with Stratusphere: a solution that's simple to install and configure; is able to collect deep metrics that span CPU, memory, and I/O performance, along with both User and I/O experience; and ships with pre-defined classification algorithms that are easy to comprehend, communicate, and adjust.

Large-scale VDI projects will almost always include an early, iterative internal selling phase – with Stratusphere in the enterprise or consultant's toolkit, we expect the process to be greatly simplified. LWL has combined years of implementation experience with built-for-purpose software to quickly make sense of a large user base, through efficient Best Fit analysis and comprehensive data collection. The vendor has already announced plans to integrate advanced user

profile management capabilities into its product line (another sticky challenge in VDI implementations), and has launched an on-line community at VDI.com to provide intelligence, best practices, and advice from VDI gurus and users alike.

We are impressed by LWL's approach to the challenges of large-scale enterprise VDI and have yet to see a competing offering that compiles such a rich data set or delivers such a powerful classification algorithm. Stratusphere takes the heavy lifting out of user, desktop and application best fit analysis, and presents its results in an innovative display that's easy to understand and socialize. We will be watching LWL closely as it matures and expands its solution reach.

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 provided without warranty of any kind. All product names 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.