



VDI FIT and VDI UX: Composite Metrics Track Good, Fair, Poor Desktop Performance

*Key indicators and classification capabilities
in Stratusphere FIT and Stratusphere UX*

Whitepaper

INTRODUCTION

This whitepaper has been authored by experts at Liquidware Labs in order to provide guidance to adopters of desktop virtualization technologies. As more organizations prepare for and deploy hosted virtual desktops, it has become clear that there is a need to measure compatibility and performance for virtual desktop infrastructure both in the planning process and when measuring user experience after moving to virtual desktops. This whitepaper covers best practices and provides an introduction to composite metrics VDI FIT in Stratusphere FIT and VDI UX in Stratusphere UX to provide a framework to measure Good/Fair/Poor desktops.

Information in this document is subject to change without notice. No part of this publication may be reproduced in whole or in part, stored in a retrieval system, or transmitted in any form or any means electronic or mechanical, including photocopying and recording for any external use by any person or entity without the express prior written consent of Liquidware Labs.

Liquidware Labs, Inc.
3600 Mansell Road
Suite 200
Alpharetta, Georgia 30022
U.S.A.
Phone: 678-397-0450
www.liquidwarelabs.com

©2011 Liquidware Labs Inc. All rights reserved. Liquidware Labs, Stratusphere, FlexApp, Stratusphere Designer and ProfileUnity are trademarks of Liquidware Labs. All other products are trademarks of their respective owners.

CONTENTS

Introduction.....	2
Overview.....	4
Emerging Best Practices for Virtual Desktop Measurements.....	4
Three Keys To Useful Desktop Virtualization Measurements	5
Gather A Complete Picture of Machines, Users, Applications	5
Synthesize Data into a Desktop Virtualization Classification System.....	6
Adjust Desktop Classifications Based On Observed Data.....	6
Liquidware Labs' Solutions	7
Introduction to VDI FIT™ and VDI UX™.....	7
The VDI FIT Classification Algorithm.....	8
Notes on VDI FIT Key Indicators	9
The VDI UX Classification Algorithm	12
Notes on VDI UX Key Indicators	14
Conclusion	17
About Liquidware Labs	17

OVERVIEW

As more organizations prepare for and deploy hosted virtual desktops, it has become clear that there is a need to support two related but critical phases. The first is to inventory and assess the physical desktop environment in order to create a baseline for the performance and quality of the user experience for the virtual desktop counterparts. When planning and preparing, organizations would like to know which desktops, users and applications are a good fit for desktop virtualization and which ones are not. The second phase is to track the virtual desktops in production in order to proactively identify when performance and user experience does not meet expectations, as well as to continue to refine and optimize the desktop image and infrastructure as changes are introduced into the environment.

Because virtual desktops live on shared systems, the layers of technology make it more complex to measure and classify fitness and user experience. But with increased industry knowledge and the emergence of best practices, plus new purpose-built products such as Liquidware Labs' Stratusphere FIT and Stratusphere UX, it is now possible to more accurately measure and classify both fitness and user experience. This white paper covers these best practices and provides an introduction to the VDI FIT and VDI UX classification capabilities in Stratusphere FIT and Stratusphere UX.

EMERGING BEST PRACTICES FOR VIRTUAL DESKTOP MEASUREMENTS

When working with standalone desktop computers, planning and support can be challenging, but for the most part it is isolated to understanding the individual user's needs, activities and issues. Planning for an individual user involves deciding on the specifications for their dedicated computer, and fixing their problems means investigating the situation on their dedicated computer, which is, to a certain extent, isolated from other dedicated computers (network and communications issues notwithstanding).

The increased use of virtualization, and in particular desktop virtualization, has created new complexities both in planning and support. When many users are sharing computer resources, either for hosted applications or hosted desktops, and systems are dynamically streamed from remote servers to local computing resources, the impact of users on each other and workloads on shared systems can be complicated to predict or diagnose. The increasing sophistication of virtualization technologies, with the ability to schedule CPU activity, over-commit memory and abstract shared storage locations, along with the increased demands on the network, make the planning and support even more of a challenge.

The good news is that desktop virtualization is quickly maturing, and best practices are emerging for how to determine which users and applications are FIT for virtual desktops, as well as best practices for how to measure and diagnose deployed virtual desktops. These best practices are beginning to circulate among desktop virtualization practitioners, and are beginning to breed a new generation of products to assist in this regard, such as Liquidware Labs' Stratusphere FIT and Stratusphere UX.

THREE KEYS TO USEFUL DESKTOP VIRTUALIZATION MEASUREMENTS

Among best practices on how to measure user and application activity, both for desktop virtualization planning and support purposes, the following are three keys that can help guide you to better results.

GATHER A COMPLETE PICTURE OF MACHINES, USERS, APPLICATIONS

The success of desktop virtualization relies on a number of factors, particularly resource requirements including CPU, memory, network and storage, all shared among a group of users. Whether planning for or supporting a desktop virtualization project, to be effective you will need to gather a complete picture of your current environment, user activity and application activity – in the planning phase you will need to do this for your existing (pre-virtualization) physical desktops, and in the support phase you will need to do this for the end-to-end desktop virtualization deployment.

Data you should consider gathering includes:

- Inventory of machines including configuration details (CPU, memory, disk, network, monitors, printers, peripherals, age, location / host, time in use)
- Inventory of applications for virtualized and non-virtualized applications, including versions and patch levels, including time in use
- Inventory of users and groups (user groups may be defined in Active Directory or LDAP)
- CPU consumption including system and user, by user, machine, application
- Memory consumption including swapping and page faults, by user, machine, application
- Network consumption and performance by user, machine, application
- Disk consumption and performance by user, machine, application
- User logon durations (the time it takes a user logon to complete)
- Application load times (the time it takes an application to load and initialize)
- Graphics intensity to identify the level of graphics and screen refresh demands by user, machine, application
- Non-responding applications
- Network latency
- Network application response times
- Failed or dropped network connections

The frequency and scope of data collection can vary according to the needs and purpose. In virtual desktop planning exercises, for example, it may be sufficient to gather data from a sample of users and desktops over a one or two week period. In support situations, you may want to gather data consistently but the frequency might be once an hour, even once a day for remote sites.

There are a number of tools and products that can be used to gather the data, including a number of open source products. With Liquidware Labs' Stratusphere FIT or Stratusphere UX, there are also now two integrated solutions that can gather all this data pre- or post-implementation. In either case, the key is to get the data into a central repository where the data can be analyzed (the Stratusphere Hub is a CMDB designed for this purpose).

SYNTHESIZE DATA INTO A DESKTOP VIRTUALIZATION CLASSIFICATION SYSTEM

Once you have gathered data, the next step is to develop a classification system. In the assessment phase, the classification system should help identify the desktops, users and applications that are a good FIT for desktop virtualization. In support, the classification system should help identify the virtual desktops, users and applications that are experiencing degradation or problems. In both cases, the classification system should ideally be explanatory – in other words, it should not only classify but should also help explain “why.”

If you have successfully gathered the kind of data described above, you have a variety of choices about how to design and implement a classification system. Some important guidelines to keep in mind are:

- Use multiple data elements, since no single element is sufficient for good classification
- Classify over a period of time, individual spikes or lulls can be misleading unless they occur frequently
- Associate classification to groups, so that you can not only see how individual users or desktops are classified, but how groups are classified, including different types of users, or desktops in a specific location or on a shared set of resources

To implement the classification system, you can design a fully- or partially- automated approach. An example of a fully-automated approach would be a system that queries the data elements, performs the classifications and stores the results. A partially-automated approach might involve a user working in an analytical tool such as Excel or SAS, then storing the results. The advantage of a fully-automated approach, such as the one implemented in Stratusphere FIT or Stratusphere UX, is that it can be continuous and up-to-date, and overall once implemented it will save time and effort.

You also have a variety of choices in the classification algorithm. A linear regression for example could be used to define a model for the entire population and then individuals (users, machines or applications) could be modeled and compared to the population, and classified based on their fitness. While this approach has merits, such as strength in anomaly detection, it also has some issues, mainly that it can require statistical expertise to tune for accuracy and it often lacks explanatory power – in other words it can be hard for lay-people to understand exactly why the model is classifying an individual a particular way.

An alternative approach is to use binning, where each data element is sub-classified into “bins”, and then the overall classification is also a set of bins derived from the sub-classifications. Bins can be defined manually, or automatically based on population percentages, value clusters, or other techniques. The assignment of data values into bins, and the assignment of the overall classification for each individual (machine, user or application) can be done by simple boundary value comparisons, or by some degree of fuzzy logic. Binning has three strengths, namely its ease of setup, its speed of calculation and its resulting explanatory power. The classification system in Liquidware Labs’ solutions uses a binning approach with fuzzy logic, as will be described further below.

ADJUST DESKTOP CLASSIFICATIONS BASED ON OBSERVED DATA

Of course no classification system is perfect or completely accurate. It is unrealistic to think that you can gather all the data elements necessary nor establish a powerful enough model to accurately classify all the time, every time. So it is very important that any classification system provide capabilities for the analysts to change, adjust and override the model based on known issues or observed data.

The following are some important capabilities that allow you to adjust, and therefore, improve the classification system over time:

- Feedback, that will show an analyst how the classifications compare to observed data
- Controls, that will allow an analyst to adjust the classification system to better fit observed data, along with the ability to test these changes
- Overrides that will allow an analyst to override the classification system for known cases or exceptions, for example specific users or applications that are known to be a bad fit for desktop virtualization

Again, you can implement these capabilities with a variety of approaches and tools, ranging from fully automated to fully manual. Stratusphere FIT and Stratusphere UX provide integrated, automated systems for feedback, controls and overrides to use in desktop virtualization planning and support.

LIQUIDWARE LABS' SOLUTIONS

Liquidware Labs' Stratusphere FIT and Stratusphere UX are complementary solutions that are designed to classify desktop virtualization fitness and the quality of user experience. Stratusphere FIT is traditionally used for assessments of the physical environment prior to launching a desktop virtualization project. Stratusphere UX is used once the physical desktops have been migrated to virtual desktops in production. Both solutions leverage Liquidware Labs' patented Connector ID (CID) Keys software. For assessments, CID Keys can be deployed to the desktops for a period of time, usually between one week and one month, to gather detailed usage statistics. For virtual desktops, CID Keys and Network Stations can be deployed on the virtual hosts to gather end-to-end diagnostics. The Stratusphere Hub virtual appliance automatically gathers the information, and provides centralized reporting purpose-built for assessment and diagnostics.

INTRODUCTION TO VDI FIT™ AND VDI UX™

VDI FIT (in Stratusphere FIT) and VDI UX (in Stratusphere UX) are built-in classification systems, each relying on set of key indicators and a rating system that can be modified and controlled by the virtual desktop project analyst or administrator. The results of VDI FIT and VDI UX are calculated for each individual machine, each individual user, and individual applications. The classification results can be viewed in graphical dashboard-like display screens and in a variety of roll-up and detailed reports.

THE VDI FIT CLASSIFICATION ALGORITHM

VDI FIT is the Stratusphere FIT classification system that helps identify the users, machines and applications that are a Good, Fair or Poor fit for desktop virtualization. This algorithm is designed to assist in the assessment and planning efforts, helping analysts more accurately and quickly determine where desktop virtualization can best be deployed, and what resources will be required. Data elements are gathered from users and applications on physical desktops, and then are analyzed as to how well they would fit into desktop virtualization. The algorithm is a binning classification system using fuzzy logic.

There are two dimensions of data elements (key indicators) used for VDI FIT – CPU and Memory FIT Indicators and IO FIT Indicators, each weighted equally. The overall VDI FIT classification is into one of three bins – Good, Fair or Poor. So, for a particular time period, you can say that a specific desktop has a “good fit” or a “poor fit” for desktop virtualization. The overall binning is based on fuzzy logic using the values of the two dimensions, each of which themselves is classified into one of three bins: Good, Fair or Poor. The classification of each dimension is also based on fuzzy logic using the binning of a set of key indicators within each dimension.

The overall algorithm can therefore be depicted as:

VDI FIT = Good / Fair / Poor based on 50% CPU and Memory FIT = Good / Fair / Poor based on binning of Key Indicators according to profile

and

50% IO FIT = Good / Fair / Poor based on binning of Key Indicators according to profile –

and

Good / Fair / Poor adjustment using Fuzzy Logic

The binning of the key indicators into Good, Fair and Poor is based on boundary values defined by the analyst in the VDI FIT Profile (default values are provided in the system) and a weighting for each indicator (the weight can be set to zero to exclude specific indicators). For example, for the IO FIT key indicator “Network Latency,” the analyst may specify that anything less than 150ms is Good, anything between 150ms and 500ms is Fair, and anything greater than 500ms is Poor. More information on the individual key indicators is provided below.

The following is a step-by-step guide to how the algorithm is used and data is classified:

1. data elements are gathered from a desktop for a period of time (example 10 minutes)
2. gathered data is sent back to the Stratusphere Hub
3. for each key indicator, Stratusphere FIT examines the value reported and bins the value as Good, Fair or Poor for that period based on the bin boundaries setup by the analyst
4. using the bin classifications of each of the key indicators and the weightings and fuzzy logic, Stratusphere FIT determines a Good, Fair or Poor rating for each of the two dimensions (CPU and Memory FIT, IO FIT) for that period; each of the key indicators is weighted equally within its dimension in this process
5. using the bin classifications for the two dimensions and the weightings and fuzzy logic, Stratusphere FIT determines a Good, Fair or Poor rating for the machine and user over the specific period of time (example 10 minutes)

So for each individual period of time, Stratusphere FIT records a specific VDI FIT rating for each machine and each user. The analyst can examine this data individually, or can view in a variety of roll-up reports (some of which average the ratings over a specified period of time).

Here are some examples of how the classification system works (example a specific machine over 10 minutes):

- If all key indicators are Good, then VDI FIT is Good
- If between 25% - 33% of the key indicators in a dimension are Poor and the rest are Good, the dimension will be Fair
- If 50% or more of the key indicators in a dimension are Poor, the dimension is Poor
- If both dimensions are Fair, the overall VDI FIT is Fair
- If both dimensions are Poor, the overall VDI FIT is Poor

The VDI FIT algorithm also includes a key override element, "Known Bad Applications." The analyst can use this to list out any applications that are known to be a poor fit for desktop virtualization (perhaps because of their heavy graphics requirements and interactivity for example). This override causes Stratusphere FIT to automatically classify a machine or user as a Poor fit during any time period where the application is observed in use.

NOTES ON VDI FIT KEY INDICATORS

The following is information on the key indicators used for VDI FIT, along with the default Good / Fair / Poor boundary values in the system. Each of these key indicators was chosen for its usefulness in helping to identify machines, users, and applications as fit for desktop virtualization. Note that by default each data element is weighted equally within its dimension for any observed period of time. Adjusting the boundary values of a single element, causing values to move from one bin to another, will affect the overall rating of that dimension relative the total number of key indicators in the dimension. For example, the IO FIT dimension currently uses three key indicators, so each indicator weights the result by 33%.

Key Indicator:	System CPU Load
Dimension:	CPU and Memory FIT Indicators (default weight 20%)
Default Boundaries:	Good <= 15%; Fair <= 30%; Poor > 30%

Notes:

This indicator is the measure of the average amount of system CPU used by a machine, user or application during a given time period. High use of system CPU is of concern when virtualizing desktops, indicating system level services that will be competing for CPU resources on a shared host. If system CPU use is high (such as averaging above 30% during any observed period) the situation should be carefully examined to determine the cause and the potential implications for virtual desktops. Consider setting the boundaries even lower if CPU resource contention is a strong concern, or if you want to better identify desktops with low CPU usage to increase virtual desktop density on specific virtual hosts.

Key Indicator: User CPU Load

Dimension: CPU and Memory FIT Indicators (default weight 20%)

Default Boundaries: Good <= 30%; Fair <= 70%; Poor > 70%

Notes:

This indicator is the measure of the average amount of user CPU used by a machine, user or application during a given time period. High use of user CPU is of concern when virtualizing desktops, indicating applications that will be competing for CPU resources on a shared host. If user CPU use is high (such as averaging above 70% during any observed period) the situation should be carefully examined to determine the cause and the potential implications for virtual desktops. Consider setting the boundaries even lower if CPU resource contention is a strong concern, or if you want to better identify desktops with low CPU usage to increase virtual desktop density on specific virtual hosts.

Key Indicator: Memory Used

Dimension: CPU and Memory FIT Indicators (default weight 20%)

Default Boundaries: Good <= 700K; Fair <= 1,200K; Poor > 1,200K

Notes:

This indicator is the measure of the average amount of memory used by a machine, user or application during a given time period. High memory use is of concern when virtualizing desktops, indicating applications or users that have large memory requirements that may be more difficult to meet if relying on shared memory resources. If memory use is high (such as averaging above 1,200KB during any observed period) the situation should be carefully examined to determine the cause and the potential implications for virtual desktops.

Key Indicator: Graphics Intensity

Dimension: CPU and Memory FIT Indicators (default weight 20%)

Default Boundaries: Good <= 200; Fair <= 500; Poor > 500

Notes:

This indicator is the measure of the average graphics intensity of each application used on a machine or by a user, obtained by examining the number of loaded GDI objects. High Graphics Intensity indicates users and applications that may have extreme graphics needs, and therefore may not be a good fit for desktop virtualization where frequent screen refreshes and highly interactive graphics may be compromised by streaming, or could negatively impact network performance. Consider adjusting the boundary settings based on the target virtual desktop environment, for example a WAN environment will be less likely to provide good support for graphics-intensive applications than a LAN, and the use of desktop virtualization technologies designed to assist graphics-intensity may also affect the fitness rating. Consider setting the boundaries lower if you want to identify low-graphics machines and users if, for example, you are targeting a desktop virtualization deployment with remote connection over the WAN and no special acceleration.

Key Indicator: Monitors in Use

Dimension: CPU and Memory FIT Indicators (default weight 20%)

Default Boundaries: Good <= 1; Fair <= 2; Poor > 2

Notes:

This indicator measures the active monitors in use on each machine. Machines and users that require more than two monitors may not be a good fit for desktop virtualization. Base the thresholds on the known limitations of your planned desktop virtualization environment, and whether you plan to change configurations for migrated users.

Key Indicator: Disk Load

Dimension: IO FIT Indicators (default weight 33%)

Default Boundaries: Good <= 15 IOPS; Fair <= 30 IOPS; Poor > 30 IOPS

Notes:

This indicator is the measure of the average disk IO per second (IOPS) by machine or by user for a period of time. This indicator covers both read and write activities. Machines, users and applications that have a high rate of IO may not be a good fit for desktop virtualization, especially if there are many such workloads competing for shared storage. Base the boundaries on the known limitations and acceptable levels for the planned storage system in the target desktop virtualization environment.

Key Indicator: Network Load

Dimension: IO FIT Indicators (default weight 33%)

Default Boundaries: Good <= 10 KB/s; Fair <= 30 KB/s; Poor > 30 KB/s

Notes:

This indicator is the measure of the average network KB per second by machine or by user for a period of time. This indicator covers both send and receive functions. Machines, users and applications that have high network bandwidth requirements may not be a good fit for desktop virtualization, especially if there are many such workloads competing for shared network on a single virtual host. Base the boundaries on the known limitations and acceptable levels for the planned network connections on your target virtual hosts.

Key Indicator: **Network Latency**

Dimension: IO FIT Indicators (default weight 33%)

Default Boundaries: Good <= 150ms; Fair <= 300ms; Poor > 300ms

Notes:

This indicator is the measure of the network latency (full round-trip, not one-way) between a machine and the Stratusphere Hub at the time of measurement. If the user is in the location where their planned virtual desktop terminal will be, and the Stratusphere Hub is in the data center where desktops and/or applications will be hosted, then this latency is useful to measure the projected latency based on the current network for the virtual desktop environment. In most LAN environments the numbers will be small and not of concern, but this measure is much more important for distributed and remote locations. For many desktop virtualization implementations it is generally thought that roundtrip latency above 300ms (above 150ms one-way) can result in poor user experience for end users, however base your boundary settings on the known concerns and limitations related to your target virtual desktop environment.

THE VDI UX CLASSIFICATION ALGORITHM

VDI UX is the Stratusphere UX classification system used to specify the users, machines and applications that appear to be Good, Fair or Poor for user experience (UX). This algorithm is designed to assist in the support and diagnosis of desktop virtualization deployments, helping administrators more easily and automatically identify when user experience may be poor or degrading. Data elements are gathered from users and applications on virtual desktops and the virtual network (on each virtual host), with visibility into the streaming protocols and the data storage systems (SAN) as well. The data elements are then analyzed as part of the VDI UX classification. The algorithm is a binning classification system using fuzzy logic.

There are two dimensions of data elements (key indicators) used for VDI UX – Machine Experience Indicators and IO Experience Indicators, each weighted (equally by default, zero weighting can be used to exclude specific indicators). The overall VDI UX classification is into one of three bins – Good, Fair or Poor. So, for a particular time period, you can say that a specific user is rated as having “good user experience” or a “bad user experience.” The overall binning is based on fuzzy logic using the values of the two dimensions, each of which themselves is classified into one of three bins: Good, Fair or Poor. The classification of each dimension is also based on fuzzy logic using the binning of a set of key indicators within each dimension. The overall algorithm can therefore be depicted as:

VDI UX = Good / Fair / Poor based on

50% Machine Experience = Good / Fair / Poor based on binning of Key Indicators according to profile

and

50% IO Experience = Good / Fair / Poor based on binning of Key Indicators according to profile

and

Good / Fair / Poor adjustment using Fuzzy Logic

The binning of the key indicators into Good, Fair and Poor is based on boundary values defined by the analyst in the VDI UX Profile (default values are provided in the system). For example, for the IO Experience key indicator “Network Latency,” the analyst may specify that anything less than 150ms is Good, anything between 150ms and 300ms is Fair, and anything greater than 300ms is Poor. More information on the individual key indicators is provided below.

The following is a step-by-step guide to how the algorithm is used and data is classified:

1. data elements are gathered from the virtual desktop environment for a period of time (example 10 minutes)
2. gathered data is sent back to the Stratusphere Hub
3. for each key indicator, Stratusphere UX examines the value reported and bins the value as Good, Fair or Poor for that period based on the bin boundaries setup by the analyst
4. using the bin classifications of each of the key indicators and the weightings and fuzzy logic, Stratusphere UX determines a Good, Fair or Poor rating for each of the two dimensions (Machine Experience, IO Experience) for that period; each of the key indicators is weighted within its dimension in this process
5. using the bin classifications for the two dimensions (weighted equally) and fuzzy logic, Stratusphere UX determines a Good, Fair or Poor rating for the machine and user over the specific period of time (example 10 minutes)

So for each individual period of time, Stratusphere UX records a specific VDI UX rating for each machine and each user. The analyst can examine this data individually, or can view in a variety of roll-up reports (some of which average the ratings over a specified period of time).

Here are some examples of how the classification system works (example a specific machine over 10 minutes):

- If all key indicators are Good, then VDI UX is Good
- If between 25% - 33% of the key indicators in a dimension are Poor and the rest are Good, the dimension will be Fair
- If 50% or more of the key indicators in a dimension are Poor, the dimension is Poor
- If both dimensions are Fair, the overall VDI UX is Fair
- If both dimensions are Poor, the overall VDI UX is Poor

NOTES ON VDI UX KEY INDICATORS

The following is information on the key indicators used for VDI UX, along with the default Good / Fair / Poor boundary values in the system. Each of these key indicators was chosen for its usefulness in helping to identify issues that directly affect user experience on virtual desktops. Note that by default each data element is weighted equally within its dimension for any observed period of time, and the weightings can be adjusted (weight of zero causes an indicator to be excluded). Adjusting the boundary values of a single element, and causing values to move from one bin to another, will affect the overall rating of that dimension relative the total number of key indicators and the weightings in the dimension. For example the IO Experience dimension currently uses three key indicators, so by default each indicator weights the result by 33%.

Key Indicator: **Login Duration**

Dimension: Machine Experience Indicators (default weight 20%)

Default Boundaries: Good <= 15s; Fair <= 60s; Poor > 60s

Notes:

This indicator is the measure of the average duration of each user login, measuring the time from when a user logs in to the time that the login process and login scripts have completed. Slow user logins can be a key indicator of overloaded or mis-configured desktop virtualization systems and are one key contributor to poor user experience. Adjust the boundaries according to expected durations, observed data and target performance in the virtual desktop environment.

Key Indicator: **Application Load Time**

Dimension: Machine Experience Indicators (default weight 20%)

Default Boundaries: Good <= 10s; Fair <= 30s; Poor > 30s

Notes:

This indicator is the measure of the average time each user application takes to load, measuring the time from when a user launches the application to the time that the application is fully initialized and ready for input. Slow application load times can be a key indicator of overloaded or mis-configured desktop virtualization systems and can be a key contributor to poor user experience. Adjust the boundaries according to expected load times, observed data and target performance in the virtual desktop environment.

Key Indicator: CPU Queue Length

Dimension: Machine Experience Indicators (default weight 20%)

Default Boundaries: Good ≤ 3 ; Fair ≤ 6 ; Poor > 6

Notes:

This indicator is the measure of the average CPU queue length in individual virtual desktops over a period of time. CPU queue length is the number of CPU instructions currently in the queue waiting for available CPU. This number should typically be below six for most virtual desktop workloads. If the CPU queue length grows, then the VM is waiting for CPU resources, for most desktops anything at or above six is a concern that CPU resources may be overloaded on the virtual host, and anything above six indicates that the system may be seriously overloaded. It is possible that some workloads do stress the CPU more and may therefore have higher CPU queue lengths that are normal and expected; if so, adjust the boundary parameters accordingly.

Key Indicator: Page Faults

Dimension: Machine Experience Indicators (default weight 20%)

Default Boundaries: Good $\leq 2,000$; Fair $\leq 10,000$; Poor $> 10,000$

Notes:

This indicator is the measure of the number of page faults observed for individual virtual desktops during a period of time. Page faults occur when the virtual desktop cannot find a page it needs in resident memory and needs to swap pages to and from disk. Higher numbers of page faults are indicators that applications are over-taxing memory resources or that the virtual desktops have insufficient memory assigned, potentially contributing to application performance degradation and poor user experience. You will need to consider and adjust the boundary values according to the machine inspection interval you are using in Stratusphere UX. For example, if your interval is to report hourly, you would expect the boundaries for Page Faults to be higher than if your interval is ten minutes.

Key Indicator: Non-Responding Applications

Dimension: Machine Experience Indicators (default weight 20%)

Default Boundaries: Good ≤ 2 ; Fair ≤ 3 ; Poor > 3

Notes:

This indicator is the measure of the number of non-responding applications observed for individual virtual desktops during a period of time. Non-responding applications can be a key indicator of overloaded or mis-configured desktop virtualization systems, and can be a key contributor to user frustration and degraded user experience. You will need to consider and adjust the boundary values according to the machine inspection interval you are using in Stratusphere UX. For example, if your interval is to report hourly you might set the boundaries for Non-Responding Applications to be higher than if your interval is ten minutes.

Key Indicator: Disk Load

Dimension: IO Experience Indicators (default weight 25%)

Default Boundaries: Good ≤ 25 IOPS; Fair ≤ 50 IOPS; Poor > 50 IOPS

Notes:

This indicator is the measure of the average disk IO per second (IOPS) by machine or by user for a period of time. This indicator covers both read and write activity. Virtual desktops, users and applications that have a high rate of IO may be an indicator of intensive workloads that can result in poor user experience. Base the boundaries on the known limitations and acceptable levels for the virtual desktop environment storage system.

Key Indicator: Disk Queue Length

Dimension: IO Experience Indicators (default weight 25%)

Default Boundaries: Good ≤ 1 ; Fair ≤ 3 ; Poor > 3

Notes:

This indicator is the measure of the average disk queue length in individual virtual desktops over a period of time. Disk queue length is the number of read and write operations currently in the queue waiting because the disk IO channel is too busy. This number should typically be below one for most desktop workloads. If the disk queue length grows then the VM is waiting for disk resources, for most desktops anything above one is a concern that disk resources may be overloaded on the virtual host, anything above three indicates that the storage system may be seriously overloaded. It is possible that some workloads do stress the storage system more and may therefore have higher disk queue lengths that are normal and expected, if so adjust the boundary parameters accordingly.

Key Indicator: Network Latency

Dimension: IO Experience Indicators (default weight 25%)

Default Boundaries: Good ≤ 150 ms; Fair ≤ 300 ms; Poor > 300 ms

Notes:

This indicator is the measure of the network latency (full round-trip, not one-way) for any connections longer than two minutes observed to the virtual desktop during a period of time. This would typically be the latency seen during desktop streaming connections (RDP, ICA) from endpoint terminals. High network latency can be a key contributor to poor user experience. In most LAN environments, the latency will be small and not of concern, but this measure is much more important for distributed and remote locations. For many desktop virtualization implementations, it is generally thought that roundtrip latency above 300ms (above 300ms one-way) can result in poor user experience for end users. However base your boundary settings on the known concerns and limitations related to your existing virtual desktop environment and network capabilities.

Key Indicator: **Failed Connections**

Dimension: IO Experience Indicators (default weight 25%)

Default Boundaries: Good ≤ 5 ; Fair ≤ 15 ; Poor > 15

Notes:

This indicator is the measure of the number of failed network connections observed for individual virtual desktops over a period of time. Failed connections, including failed RDP or ICA connections, can be a key indicator of overloaded or mis-configured virtual desktop infrastructure systems, and can be a key contributor to user frustration and degraded user experience. You will need to consider and adjust the boundary values according to the machine inspection interval you are using in Stratusphere UX. For example, if your interval is to report hourly you might set the boundaries for Failed Connections to be higher than if your interval is ten minutes. The values may also vary based on the applications and normal behavior in the environment, so adjust the boundary values accordingly.

CONCLUSION

If you are planning a desktop virtualization project, your first step should be assessing the current physical environment to classify physical desktops' fitness for the virtual environment. Assessment is critical to the success of your project. Not only will you save time and reduce cost, you will also be fairly well assured that you will design the correct image(s) and infrastructure to support the type and number of user(s) you have selected. The assessment metrics can also be used to establish a baseline for performance of the virtual infrastructure and set the bar for the user experience. Once the project is moved well into production, metrics can be used to provide insight into how well the system is performing and what the user experience is like on the virtual desktops. When scaling or changing the virtual desktop production environment too, it is important to consistently track and measure key indicators outlined in this paper to keep the system performing optimally. The effort to develop a solid process and system pays off both in the short- and long-term, and adhering to a set of best practices such as those defined in this white paper will help make sure that your efforts have the intended results. There are a variety of tools and approaches you can use in this effort. For a fully integrated and purpose-built solution, we invite you to investigate Liquidware Labs' Stratusphere FIT and Stratusphere UX solutions for your desktop virtualization planning and support needs.

ABOUT LIQUIDWARE LABS

Founded in 2009, Liquidware Labs™ is the leader in desktop transformation solutions for next-generation physical and virtual desktops, including VMware View, Citrix XenDesktop, and Microsoft Windows 7. The company's Stratusphere™ and ProfileUnity™ solutions have been described by analysts as the industry's first 'On-Ramp to DESKTOP VIRTUALIZATION,' providing a complete methodology and software that enables organizations to decouple users and applications from the operating system and to cost-effectively assess, design, migrate, and validate the user experience for next-generation desktop infrastructure. Liquidware Labs products are VMware and Citrix certified, and are available through a global network of certified partners. Visit www.liquidwarelabs.com for further information.