Stratusphere SpotCheck for Citrix DaaS and Virtual Apps and Desktops

Last Updated: 23/02/2026

SpotCheck Methodology Guide

Contents

Document Purpose: To define metrics and thresholds for a SpotCheck as it relates to User Experience in a Citrix Virtual Apps and Desktops (formerly XenApp/XenDesktop) environment utilizing Liquidware Stratusphere UX.

This document is designed to bring together recommendations from many experts in the industry about the metrics that need to be monitored and the thresholds that are deemed acceptable as it relates to User Experience. This document does not make recommendations on changes needed due to the many industry, usage, costing, and application variables that are in play.

What is a "SpotCheck"

A SpotCheck is a point in time health check that focuses on key User Experience metrics with known acceptable performance levels. The review of data from multiple dates and times is critical before making recommendations or changes to the environment. These thresholds represented below are taken at a one hour level of granularity unless otherwise specified in the description and are key areas that affect User Experience. Normal/High Usage dates and times should be examined based on the industry and user requirements.

Critical Notes:

A. Know your company!

- Know your industry/company/department work habits, loads and applications are critical for data interpretation and threshold evaluation.
    1. Example: Moderate/High storage latency may be acceptable during shift changes with large number of users logging in and out, but this is not acceptable during normal work hours as this impedes productivity.
    2. Example: Law firms and healthcare organizations generally need/want sub ten second login times whereas most other industries are satisfied with under 30 second.

B. Know your data!

- There are many monitoring and diagnostic solutions out on the market. Each of the solutions collect data differently and have different levels of granularity. All of these solutions render/report the data in differently and with unique granular roll ups that can drastically change the data and perspective for the user. For this reason, the metric values represented in this document are only for Liquidware Stratusphere UX and may not apply well to other products.
    1. Example: Depending on the view, you could be looking at averages, peaks or peak averages. Did the data come from the Broker, Hypervisor, "In-Guest", "In Band" or "Out of Band"? How much impact did the "In-Guest" Agent put on the OS? How much impact and time lag is on the "Out of Band" Agent, Broker and Hypervisor?

C. Good Blogs

- SpotCheck Methodology
- Grey Matter is Required – Automated Solutions don't work
- Monitoring vs. Diagnostics

D. Liquidware Community Site and Important Links

- **Liquidware Community** - Slack.
- **Important Links** - SE Field Articles, Product Documentation, etc…

Information Needed for Analysis, Conclusions and Recommendations:

A. Multiple Spot Check Dates

- MM/DD/YYYY (Monday), MM/DD/YYYY (Wednesday), MM/DD/YYYY (Friday)

B. Multiple Time Frames Examined on each date

- (Time frames for review are based on business requirements)
  9-10AM, 10-11AM, 2-3PM, 4-5PM

The system(s) should be examined on multiple dates and times for the following information based on the max values shown below. **Please do not make a change based on a single data point.**

Critical Sections for Review:

A. ESXi Host Criteria: (VMware vSphere Best Practices)

- CPU % - (Max 80% Average over 60 minutes)
- CPU % Ready (Max 3% Average over 60 minutes)
- Memory Utilization (Max 85% at any point in time)
- Memory Swapping and Ballooning (Should always be ZERO)
- Datastore Latency (Should be under 10 Milliseconds – Max 15 Milliseconds)

B. Machines/OS Criteria:

- Machine Last Boot – Critical Question – How long has the machine been running?
  1. See Appendix R for more details.
- Login Delay (Industry Average is under 30 Seconds – This is a company preference
  1. See Appendix L for more details.
- Application Load Time (Industry Average is under 3 Seconds – Company Preference)
- CPU Utilization (Max 80%) – Higher than 50% generally is bad over 60 Minutes
  1. This generally denotes stuck or run-away process(es) on the machine.
- CPU Queue (Should not be more than 1 per vCPU assigned to VM)
  https://technet.microsoft.com/en-us/library/Cc940375.aspx
- Memory Usage (Should be less than 80%)
- Best Practice is to reduce Windows Paging
- Page File Usage (Should be as close to zero as possible)
  1. Windows paging cannot be stopped.
  2. Do not turn off the paging file in windows. Set to minimum and maximum size of the page file.
  3. Do not use “System Managed” – Set the page file start size to ¼ the memory.
  4. Windows Paging causes CPU and Disk Overhead and should be reduced whenever possible. To reduce paging, allocate more memory to the virtual machine.
  5. Soft Page Faults occur in memory and Hard Page Faults occur to the disk.
- Disk Queue (Should be ZERO for 99% of Users)
  1. Disk Queue shows that the OS is waiting on disk reads/writes.
  2. This can be caused by antivirus holding up the IO or latency of the disk sub system.
  3. See Appendix D for more details.
- Graphics Intensity will be noted has high when over 100 for more than 1/3 of users.
  1. This must be examined to see if graphics off load processor would help
  2. See Appendix G for more details.
- Applications Non-Responsive – (1 per Day/Per Machine/App is OK)
  1. Any more than this requires investigating the apps and services used by the application.

Appendix:

Appendix A: HDX (formerly ICA) - Remote Display Protocol

- Image Quality:
  1. This is a good "Base" metric to gage/monitor the user connection quality. HDX will lower the image quality if it has packet loss, high latency or a low bandwidth connection to the end user.
  2. Modern supported versions of Citrix VAD 7.15+ has a default image quality of medium.
  3. Adjusting image quality higher than the default is only needed for clients that may be viewing broken bones and need to see very small fractures.
  4. Image quality directly affects the number of frames per second sent from the VM to the end client. This can dramatically affect the Host CPU and network bandwidth required per user.
- Session Latency:
  1. General Max Observations:
     - New York to California – 30-50 Milliseconds
     - USA to India – 150-200 Milliseconds
     - Interoffice same city – 10 Milliseconds
     - Inter0ffice same building – 5 Milliseconds
     - Huge Latency Numbers in Stratusphere shows users dropping on and off the network. (Huge Denotes 800+ ms)
- Protocol: (Good and Bad… This is Just info for you)
  1. HDX is a combination of TCP and UDP when allowed and TCP only as a failback.
     - UDP Packets have a lower priority than TCP Packets on most networks.
     - UDP is dynamic and bursty by nature of the protocol.
     - UDP is faster than TCP because there is no error-checking for packets
     - UDP is lightweight. There is no ordering of messages, no tracking connections, etc.
     - UPD can perform error checking if enabled but there is no recovery. Packets must be resent and with no ordering it is up to HDX to request large blocks for retransition.
- Packet Loss:
  1. Packet Loss with HDX can cause users poor experiences: mouse lag, screen artifacts, slow screen redrawing, typing latency, etc…
- General Recommendations:
  1. QOS (Quality of Services) should be implemented on all routers.
     - HDX should be right under Voice Over IP and Video.
  2. Lower the Maximum Image Quality using the HDX Policy settings to best fit identified business use cases and applications.
     - USB and Audio Channels:
       Disable and lower the priority of these channels as it suites your business requirements. Disabling USB or lowering Audio quality can dramatically lower VM/Host CPU and Network requirements.
  3. There are many options for HDX Tuning. Test all scenarios and consult the best practice guides from Citrix and tune/monitor users for best user experience based on the environment.
- HDX Technical Overview - Overview and Thinwire Graphics Modes

Appendix D: Disk Queue/Disk Response/Disk Latency

1. Terms:
   - I/O: This is Throughput - Number of Megabytes per Second (MBps)
   - IOPs: This is Input/Output Operations Per Second (typically related to storage/disk performance)
   - Disk Queue is the number of processes in the queue to read/write in the disk file system
     Note* This is affected by disk response time (Latency) and filter drivers that are in between the physical disk and the file system.
     Example: Click on a Microsoft Excel Spreadsheet in file explorer. Antivirus will scan the document first then let Excel read the data. When Excel saves the document, it must go back through the Antivirus scanner before being written to the disk.
     Example: When viewing a machine, the Disk Queue may be high as an overall average for the day. This is normal as the login process lays down the user profile data and registry settings. Drilling down on the machine may show that disk queue was only high during the login process. Generally, this can be ignored for machine performance but should be investigated in login breakdown for more optimization of the login process. This is normal as the login process lays down the user persona and registry settings. Generally, this can be ignored for machine performance but should be investigated in login breakdown for more optimization of the login process.
2. Disk Response (Latency) is the Read/Write time from the operating system file system to the underlying physical/virtual disk.
3. Machine Thresholds:
   - Disk Response/Latency should generally be less than 1-2 milliseconds
   - Disk Queue is preferred to be 0.02 or less over a one hour timeframe
     Note: Cloud Machine I/O and IOP Limits vary by provider and instance types with bursting based on current load of the provider.

Appendix G: Graphics Intensity

4. Graphics rendering is a large part of the user experience. Depending on the application it can use MS GDI, DirectX, OpenGL, CUDA, etc… or many other video interface drivers/protocols.
5. There is always a misconception that since there are no extremely graphic intensive applications that GPUs (Graphics Processing Units) are not needed. This is not true; Windows and normal Microsoft Office applications have a lot of graphics requirements. All desktops/laptops built in the last 10 years have a GPU. These processors are used by the OS and applications to offload drawing of boxes, circles and other complex shapes from the main CPU and rendering them on the monitor.
6. GPUs are not all the same! Manufactures pick from many vendors to meet a cost point for the desktop or laptop they are selling.
   - Laptops: Tend to have energy/heat constrained GPUs.
   - Desktops: Have many tiers and options for expansion with more power and cooling available.
   - Driving multiple monitors at high resolution can often overload the built in GPU and then offload that back onto the main CPU.
   - Improperly installed video drivers and older versions can also cause off load back to the main CPU.
   - Graphic rendering does not show up in Task Manager, Resource Monitor or Stratusphere because this is a Kernel process and very hard to break out.
   - When looking at a physical machine with no obvious constrains on memory or disk we must look at CPU Utilization and CPU Queue. Low to moderate CPU utilization with HIGH CPU Queue is sign of overloaded graphics process. Also examine the GDI (Graphics Device Interface) objects in Task Manager or Stratusphere. GDI Objects average for the machine over 1 hour greater than 100 is consider high graphics intensity.

     - Example Application GDI Usage: Microsoft Outlook:
       First Monitor (1024x768) – 800-900 GDI Objects
       Second Monitor (1320x1024) – 1,200-1,400 GDI Objects
7. This is a complex topic and often difficult to identify. Stratusphere does show GPU utilization for many of the manufactures on the market. If you see no GPU load in Stratusphere for a physical machine the GPU is not reporting information, supported, drivers are bad, or resolution is not supported by the GPU/Driver.
8. If you find that you are overloading the GPU in the machine you have 2 options. First, turn off hardware acceleration for the applications or second, buy machines with faster GPUs.
9. Microsoft Office, Google Chrome and Mozilla Firefox all have Group Policy settings to disable hardware acceleration.

Appendix I: Important Links

- Liquidware SE Field Articles
- Liquidware Documentation Site - Official Product documentation
- Liquidware Community - Slack, Blogs, LinkedIn, X, Facebook, YouTube, etc…
- Liquidware Stratusphere Training Site – Training Site

Appendix L: Login Delay

- Time consumed with users logging into a machine is a large part of the user experience. Stratusphere can breakdown the machine boot and login processes. Due to the complexity of active directory and the environments we can only offer a few guiding hints in this document. For a complete login breakdown session please engage Liquidware SE/support or partner.
- Domain Controller (DC) Discovery Time
    1. DC Discovery happens at boot and login time.
    2. Healthy response times are 300-500 milliseconds.
- Changing of the DC during boot and login shows a potential issue.
    1. DC Discovery Times over 500ms:
        - DC Overloaded – Cannot process request fast enough.
        - Network latency from the machines to the DC.
        - Sites and Services – Machine/User is talking to a DC in another location.
- Long running processes
    1. AD GPOs, Item Level Targeting and Scripts.
        - Need to review these in Stratusphere Login Breakdown.
        - AD Lookups and Local machine WMI Queries are very slow.
        - Mapping a drive/printer to a machine that does not exist, or the user does not have access can make the login excessively long.
    2. Antivirus Scanning
        - Don't forget that batch files, PowerShell, VB Scripts are all interpreted languages. Meaning that each line in the batch file or script is executed one line at a time. AV systems scan each line then all the previous lines of the script to ensure it is not a virus.
- Domain Overview
    1. Understand which Domain Controllers are processing logins.
    2. How long was the average authentication process on each Domain Controller?
    3. Understand which Domain Controllers have a large number of abnormal events.
- Physical Desktops and Persistent virtual machines need to be treated differently than non-persistent virtual desktops.
    1. Broken and/or Corrupt GPOs.
        - A yearly (at a minimum) review of the GPOs should be performed. Example: IE7 GPOs should not be applied to Windows 10.
        - Conducting GPO reviews can help with login times and with security.
    2. Sites and Services
        - This is one of the top issues found with Stratusphere login break down.
        - A machine in New York it should not be authenticating from a domain controller in Canada.
        - With the speed on needing to provide Work from Home/Work from Anywhere new virtual desktop pools or new VLANs were deployed support these initiatives and zoning properly in the correct sites and services for authentication can be missed.
- Animated GIF on how to get to Login Breakdown - Login Breakdown
- Animated GIF on how to get to Domain Overview – Domain Overview
- Link to Training Video on Login Breakdown – Login Breakdown Training Video

Appendix R: Machine Last Reboot Time

- Knowing how long a machine has been running is a critical question. Applications can have memory, graphics and CPU process "Leaks" over time which can/will degrade performance. Machines running longer than one month are also missing critical security/feature patches that put them out of security compliance and at risk.
- Know the difference between Restart vs Shutdown. In modern versions of Windows 10, by default there is a feature called "Fast Startup" enabled. With "Fast Startup" enabled when the machine is issued a Shutdown command the state of the Windows kernel is saved to disk after logging the current user off the system to speed up the machine boot. This can lead to the behavior of Windows Updates not being installed. Restart fully flushes everything running out of a machine. If the desired behavior to have Shutdown completely flush everything running out of the system, this can be configured via GPO or registry key.
- Below is a recommendation only of reboot policies based on experience of Liquidware engineers. This is not a Liquidware recommendation as there are no official recommendations from Microsoft.

  Note: The below recommendations also must conform to company business practices and change control policies.

  1. ***Domain Controllers:***
     - Monthly Reboot – Primarily for OS Security Patches
  2. ***Critical infrastructure machines running Windows Server OS:***
     - Monthly Reboot – Primarily for OS Security Patches
  3. ***Single User Virtual Machines (persistent and non-persistent):***
     - Minimum of a Weekly Reboot – Your mileage will vary based on the applications being used by the users. A Daily reboot is ideal to ensure users have the best experience.
     - Minimum of a Monthly Reboot for OS Security Patching.
- Minimum of a Monthly Reboot for OS Security Patching.

Appendix V: VoIP – Voice over IP

- Voice over IP Solutions are critical to business meetings and user to user calls. There are many solutions on the market for VoIP and team chat solutions, but they all rely on the network to provide good call quality.
- Most voice over IP solutions and chat systems can sustain a good voice quality up to 200 milliseconds of latency.
  Poor voice quality is introduced when "Jitter" is over 5 Milliseconds.
  Jitter: Is the difference in latency millisecond to millisecond.
- CPU being overloaded can cause latency and this is commonly overlooked. See Machines/OS Criteria section for more information on CPU utilization.
- Why does Jitter Happen:
  1. User network overloaded with other apps downloading/uploading information.
     Note: Many VoIP solutions can offload voice connections from a virtual machine back to the end user device thereby reducing latency and jitter.

Appendix G: Graphics Intensity

10. Graphics rendering is a large part of the user experience. Depending on the application it can use MS GDI, DirectX, OpenGL, CUDA, etc… or many other video interface drivers/protocols.
11. There is always a misconception that since there are no extremely graphic intensive applications that GPUs (Graphics Processing Units) are not needed. This is not true, Windows and normal Microsoft Office applications have a lot of graphics requirements. All desktops/laptops built in the last 10 years have a GPU. These processors are used by the OS and applications to offload drawing of boxes, circles and other complex shapes from the main CPU and rendering them on the monitor.
12. Non-vGPU enabled Virtual Machines:
    - Turn off hardware acceleration for all applications. Even though you don't have a vGPU in the host Hypervisor (VMware, VirtIO, Citrix Hypervisor and Hyper-V) Guest tools still has a driver that looks like a GPU to the OS and applications.
13. Applications that have the option to disable "Hardware Graphics Acceleration" should be done unless you have a GPU installed in the host. Most modern application have a GPO that can turn this off. Note that this is generally a per user-based GPO. Microsoft Office, Google Chrome and Firefox all have GPO settings to turn off hardware acceleration.
    - Note: These simple application changes can result in a 10% CPU reduction on your host operating system. Your results will vary based on the OS, application and host. You can monitor this with Stratusphere.
14. vGPU enabled Virtual Machines - (Machines that have access to a vGPU in the hosts)
    - vGPUs are expensive and sometimes difficult to determine if you are getting the most out of them. The resources are allocated per machine and most settings are around the framebuffer (aka GPU RAM) allocation. Stratusphere can determine if the machine/app is using the GPU memory that is assigned to it.
    - Example: Allocated 2,048MB of GPU framebuffer but it is observed that only 768MB is being used with a non-consistent burst to 1,024MB will allow for lowering the framebuffer allocation to 1,024MB will allow for more vGPU enabled machines.

Appendix I: Important Links

- Liquidware SE Field Articles
- Liquidware Documentation Site - Official Product documentation
- Liquidware Community - Slack, Blogs, LinkedIn, X, Facebook, YouTube, etc…

Appendix L: Login Delay

- Time consumed with users logging into a machine is a large part of the user experience. Stratusphere can breakdown the machine boot and login processes. Due to the complexity of active directory and the environments we can only offer a few guiding hints in this document. For a complete login breakdown session please engage Liquidware SE/support or partner.
- Domain Controller (DC) Discovery Time
    3. DC Discovery happens at boot and login time.
    4. Healthy response times are 300-500 milliseconds.
- Changing of the DC during boot and login shows a potential issue.
    2. DC Discovery Times over 500ms:
        - DC Overloaded – Cannot process request fast enough.
        - Network latency from the machines to the DC.
        - Sites and Services – Machine/User is talking to a DC in another location.
- Long running processes
    3. AD GPOs, Item Level Targeting and Scripts.
        - Need to review these in Stratusphere Login Breakdown.
        - AD Lookups and Local machine WMI Queries are very slow.
        - Mapping a drive/printer to a machine that does not exist, or the user does not have access can make the login excessively long.
    4. Antivirus Scanning
        - Don't forget that batch files, PowerShell, VB Scripts are all interpreted languages. Meaning that each line in the batch file or script is executed one line at a time. AV systems scan each line then all the previous lines of the script to ensure it is not a virus.
- Domain Overview
    4. Understand which Domain Controllers are processing logins.
    5. How long was the average authentication process on each Domain Controller?
    6. Understand which Domain Controllers have a large number of abnormal events.
- Physical Desktops and Persistent virtual machines need to be treated differently than non-persistent virtual desktops.
    2. Broken and/or Corrupt GPOs.
        - A yearly (at a minimum) review of the GPOs should be performed. Example: IE7 GPOs should not be applied to Windows 10.
        - Conducting GPO reviews can help with login times and with security.
    3. Sites and Services
        - This is one of the top issues found with Stratusphere login break down.
        - A machine in New York it should not be authenticating from a domain controller in Canada.
        - With the speed on needing to provide Work from Home/Work from Anywhere new virtual desktop pools or new VLANs were deployed support these initiatives and zoning properly in the correct sites and services for authentication can be something that is missed.
- Animated GIF on how to get to Login Breakdown - Login Breakdown
- Animated GIF on how to get to Domain Overview – Domain Overview

Appendix R: Machine Last Reboot Time

- Knowing how long a machine has been running is a critical question. Applications can have memory, graphics and CPU process "Leaks" over time which can/will degrade performance. Machines running longer than one month are also missing critical security/feature patches that put them out of security compliance and at risk.
- Know the difference between Restart vs Shutdown. In modern versions of Windows 10, by default there is a feature called "Fast Startup" enabled. With "Fast Startup" enabled when the machine is issued a Shutdown command the state of the Windows kernel is saved to disk after logging the current user off of the system to speed up the machine boot. This can lead to the behavior of Windows Updates not being installed. Restart fully flushes everything running out of a machine. If the desired behavior to have Shutdown completely flush everything running out of the system, this can be configured via GPO or registry key.
- Below is a recommendation only of reboot policies based on experience of Liquidware engineers. This is not a Liquidware recommendation as there are no official recommendations from Microsoft.

  Note: The below recommendations also must conform to company business practices and change control policies.

  4. ***Domain Controllers:***
     - Monthly Reboot – Primarily for OS Security Patches
  5. ***Critical infrastructure machines running Windows Server OS:***
     - Monthly Reboot – Primarily for OS Security Patches
  6. ***Single User Virtual Machines (persistent and non-persistent):***
     - Minimum of a Weekly Reboot – Your mileage will vary based on the applications being used by the users. A Daily reboot is ideal to ensure users have the best experience.
     - Minimum of a Monthly Reboot for OS Security Patching.
  7. ***Multi-User Virtual Machines (persistent and non-persistent):***
     - Weekly Reboot – At a minimum a scheduled reboot weekly is recommended.
     - Minimum of a Monthly Reboot for OS Security Patching.

Appendix V: VoIP – Voice over IP

- Voice over IP Solutions are critical to business meetings and user to user calls. There are many solutions on the market for VoIP and team chat solutions, but they all rely on the network to provide good call quality.
- Most voice over IP solutions and chat systems can sustain a good voice quality up to 200 milliseconds of latency.

  Poor voice quality is introduced when "Jitter" is over 5 Milliseconds.
  Jitter: Is the difference in latency millisecond to millisecond.
- CPU being overloaded can cause latency and this is commonly overlooked. See Machines/OS Criteria section for more information on CPU utilization.
- Why does Jitter Happen:
  2. User network overloaded with other apps downloading/uploading information
     Note: Many VoIP solutions have the ability to offload voice connections from a virtual machine back to the end user device thereby reducing latency and jitter.