TOP MISTAKES 3 FIXES & TIPS

BECOME A CITRIX HERO

from
DJ Eshelman
CTXPro.com

#CITRIXHERO

Introduction- Why I'm Giving This Guide Away

Between 2017 and 2018, I decided to catalog the recommendations I was making as a Citrix expert and consultant. These recommendations were being made for companies paying well over $15,000 for me to be there. But the pattern I'm seeing over and over again is that I tend to be making the same recommendations as primary findings. This tells me that this is information that the community as a whole really needs… but it's not proprietary information! In fact, this is all freely available information! So, I have decided to give you the top 3 recommendations that repeated themselves more often among the more than 350 recommendations made in the last year.

But more than that- I'm seeing a lot of pain out there. A lot of people I meet are not confident about Citrix. They spend late nights trying to deal with problems that could have been avoided. They wonder about their job future. I think that's a shame. I decided to do something about it.

These are the highest-impact findings, the ones that cost the least to deploy and have the highest results. What is unique about this guide is that I'm going into much more detail than usual and including a lot of related recommendations that could merit their own chapter at times. So, this was too much for a blog series. And at just under 10,000 words a book seemed the best format.

Regardless, I encourage you to not just check your environment for these things but more importantly to understand WHY they are considered leading practices.

I'm providing these for free, but I'm also making a subscription available where every month I'll send you a new recommendation via email. If you are not getting these emails already- You can sign up at https://ctxpro.com/Top3

If you have come across this eBook from another method than from my email service, I'm flattered that someone would share it- but you'll get better value from joining my email list!

I hope you not only benefit from these tips, but that you become known in your company as…
The Citrix Hero.

About the Author

DJ Eshelman lives in the Nashville, TN area (USA) and works as an independent consultant and Ad-Hoc/Resident Consultant with Citrix Consulting Services where he's served companies of many sizes, including 20 of the Forbes Top 100 companies.

- Citrix Certified Expert
- CTA (Citrix Technology Advocate) 2017-2018
- CUGC Leader - Nashville
- Creator of CTXPro.com
- CitrixCoach.com – Creator & Lead Coach

Contents

How to Use This Guide

This reference is meant to be updated on a yearly basis to my subscribers at ctxpro.com
This first publication is meant to introduce you to me and how I lay out information to help you learn and grow in your career.
Every month, I'll be publishing a new tip either from myself or trusted colleagues, CTAs and CTPs.

Scoring

I use a scoring system for each tip so you can see at a glance:

- The overall importance to your company leadership.
- The level of impact to your users
- How much of a security impact the tip
- How difficult the solution is to implement (the skill level involved)

Layers

I use a system of defining the focus areas for each recommendation called Layers. This system is similar (nearly the same) as Citrix uses and contains the following layers:

- Business (how the solution meets business objectives of the company)
- User (the needs, applications, devices and ways users work)
- Access (the means by which users get information and applications)
- Resource (the applications, desktops and documents and how they are structured)
- Control (Services, Policies and back-end configurations that control resources and access)
- Cloud (The physical and virtual host, hypervisor, network, storage and other considerations)
- Security (the methods by which information will stay under corporate control, etc)
- Operations (The team, tools and other methods used to maintain the environment)

Success Lanes

I'll also include a skills inventory system that aligns to my membership program. Here's how this works. My members learn within 4 primary areas that I call Success Lanes.

- Understand (knowing what Citrix solutions are and why they are important)
- Maintain (managing the day-to-day operations for Citrix technology)
- Build (Engineers and others tasked with making changes or building new environments)
- Design (persons tasked with the guidance and overall planning for Citrix technology)

Where appropriate, I'll also include any Prerequisite knowledge and Links for each tip.
If you need help, the CTXPro community is here for you. You can become a member at https://ctxpro.com/membership or email me at CoachDJ@CTXPro.com for a consultation. Though it is in beta as of this writing, a plan is in the works to do live Q&A sessions with me for each tip for a minimal monthly fee. If you'd like to participate, please see the membership page. If there is adequate interest, I may also create an online course series that does a deep dive into each monthly tip- email me if you are interested in that and I will get you onto the wait list!

Tip #1: Windows OS Tuning

It is one of the most popular web searches regarding Citrix.
It is one of the most well-documented and blogged about topics…
Yet I still manage to find this one ignored recommendation more than any other thing. This really amazes me, because I think this one thing when properly implemented can save companies in some cases tens of thousands of dollars every year- and the information on how to do so is free and has been for more than a decade!
So why isn't this being done? Because in the 'click next, next, next' nature of install in the modern Citrix world it is becoming so easy to install that those without experience are now able to install and run with Citrix effectively. As you'll see later, the answer of simply making another thing to click next in would be risky for Citrix to do, so I don't see it ever being done.
Why?

Because this one thing is that it involves the default settings of the Microsoft Operating system.

Layers	• Resource
Success Lanes	• Design • Build
Prerequisites	• Windows Desktop and Server OS Knowledge • Image Deployment Design

SCORES

Importance

★★★★★

User Impact

★★★★★

Security Impact

★★☆☆☆

Difficulty

★★☆☆☆

The #1 Culprit of Bad Performance: Operating System Defaults

Did you know that the 'out of the box' configuration for every Microsoft OS is NOT optimized for virtual delivery? Microsoft builds the operating system for compatibility, not performance.
Why? This mostly has to do with the number of background services and tasks that ship with Windows. The intention is good- it serves a wider range of needs and is better to have 'on' by default rather than try to make the literally thousands of mechanisms to detect when a service is appropriate. This is why you'll see Bluetooth services on server OS on a Virtual Machine (VM).
Microsoft basically assumes YOU will fine-tune the services and settings for your use case (if you want to)! Of course, try to find this on the box and you won't. On their website you will find a guide (which fellow CTA Dennis Span recently pointed out that the article says that you don't need to optimize and then goes on to point out several pages of optimizations…).
Yet, with each release of the Operating system (or in the case of Windows 10, each iteration or build) it seems the amount of resources (CPU, Memory, IO) increases… all because there are so many things running in the background. YOUR JOB is to figure out which of these is beneficial and which ones are not in a Virtual environment.

Leaving Services that do not need to be running active uses CPU, Memory and IO without giving any benefit to your company. It is wasteful and should be avoided.

To make things worse- many of the default Registry settings and buried background scheduled tasks can also tax or even completely hinder proper operations in a virtual space.
Here's some considerations:

- Are you running a Virtual Machine? You probably are, but it is worth exploring the implications here. For example- if virtual, you won't have attached peripherals. So, a way to optimize right away is to make sure that services that support these are not in use (Wireless, Bluetooth, etc)
- Is your machine running on a SAN? Quite often, a traditional VM running on a SAN will be using SSDs. This is great… but in traditional cases (exceptions will be discussed later in this chapter) performance may actually be *degraded* by several VMs running a defragmentation (Microsoft now refers to this generally as 'optimization') and life of the SSDs can be reduced… with minimal impact. The same goes with those on hyperconverged architecture – rearranging blocks isn't always needed (exceptions exist with PVS, but we'll talk about that later).
- Are your VMs persistent? Some background tasks are a good idea for machines that persist between reboots… but are a *horrible* idea on VMs that reset upon reboot! It is important to know and think about these and determine the benefit individually.

KEY CONCEPT: First Do No Harm

Just because myself or some other expert says you 'should' perform an optimization task… doesn't make it automatically appropriate for you. You MUST do the work of determining which is appropriate in your case. If you're lost or unsure, I would much

rather you reach out to me at citrixcoach.com for a quick (and free) consultation on if you need more help than to take a needless risk.

I mentioned the need is increasing. This may not always show up on an individual VM, but at scale (several or several hundred) VMs can really demonstrate an impact on resources basically being thrown away! Recent tests by LoginVSI proved that Server 2016 is especially bad in this regard- unoptimized VMs can quite literally cost you thousands of dollars because less people can be logged into each VM; and less people can be hosted on each physical server. Add in the effects of the Meltdown and Spectre remediations and virtual hosted servers are able to do far less with the resources available. This means less users per blade which at scale; which is a huge problem.

HERE'S THE BAD NEWS: THIS PAST YEAR MY OBSERVATIONS WERE THAT 90% OF THE COMPANIES I VISITED FOR ASSESSMENTS DID NOT FOLLOW MANY OF THE CITRIX RECOMMENDED GUIDELINES FOR OPTIMIZATION. WHEN THESE COMPANIES IMPLEMENTED THE OPTIMIZATION STEPS PROPERLY, SOME SAW INCREASED USERS PER BLADE OF OVER 30%. ADD IN THE TIPS AND TRICKS I SHARE WITH YOU BELOW AND... YOU'LL BE DRASTICALLY INCREASING THE PERFORMANCE OVERALL!

Needless to say it matters; however fixing it is easy and most tools I talk about are free!
Bottom line- saving your company several thousand dollars = #CitrixHero

Operating System Optimization for Citrix Use Cases

We know we need to optimize the Operating system for remote use, an in some cases for multi-user uses. Truth is some of these optimizations can be used on the Control components as well, but for now we will focus on Resource Components (VDAs).
Server 2016, though it hasn't had a huge amount of adoption yet… is a challenge. Testing from LoginVSI has indicated that unfortunately 2016 just doesn't scale anywhere near as well as 2012R2, even when optimized- though when optimized it is a lot closer. Sorry folks, just the way it is.

Before We Begin

I assume in all things that you are testing, making backups and never EVER putting untested optimizations in production!
All of the optimizations I discuss are essentially free but may be subject to specific usage licensing and caveats. It is your responsibility to read the agreements before you run the tools!
I'm going to assume that you're willing to do some research along with reading this eBook. At the end of this chapter I have several articles I definitely want you to review!

Hit the Easy Button: Citrix Optimizer

There are several options, primary of which is the Citrix Optimizer – a free tool that automatically detects variations from the recommended tunings and lets you select which ones to apply and which ones not to (for example, on some servers you may want Windows Search to run for Outlook – on others you may not, so optimize and test appropriately).

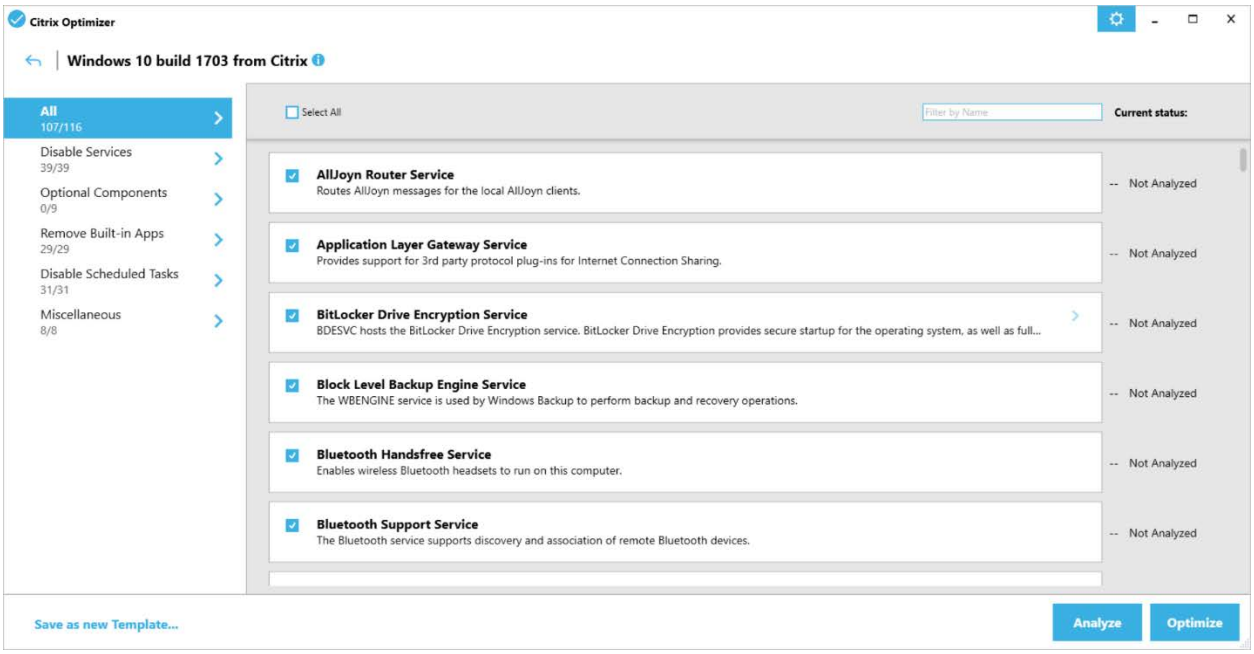

The reality is that simply running this tool will disable a ton of unneeded services, remove built-in apps, and set certain key registry settings with a few clicks. This used to be either a manual process (which was rarely done) or scrips (which often didn't work). This effort, however- is updated frequently so you are always assured it is going to work. It supports rollback and creating your own templates. In all- what this is doing for free is what I used to end up charging clients more than $1500 of efforts to complete. You'd be a fool not to take advantage of this!

DON'T BE A FOOL – USE THE FREE TOOL!

Now- a few caveats:

1) **Take it One Step at a Time**. As you can see in the screenshot above there are literally hundreds of optimizations. If you are interested in a deep dive of these settings, please let me know at CoachDJ@ctxpro.com – if enough are interested in the 'why' of all of these I'll do a deep dive. Generally speaking, however I can tell you that Martin Zugec (creator of the Citrix tool) and the CTA/CTP groups are constantly working to validate these settings. What they've done is created the "Optional Components" section. My suggestion is to validate the base recommendations first, make a backup, then test each of the Optional Components INDIVIDUALLY. Unless directed by a consultant specifically (or you know what you're doing) the components in this section can sometimes even cause VDI to stop working (for example), but are good for back-end services. As you'll often hear with Citrix Consulting: "It Depends…"
2) **TEST**. Always test before going into production with changes like this. Just because you hear me say to do something great doesn't mean that every element of it is what is required for your users. Every use case has little differences which require your attention.
3) **Be careful with certain Optimizer options such as disabling Windows Search**. Search may be something your users need for Outlook or other requirements and their user experience will be diminished without it. In a great many cases, I find that one image may need services that another does not because of the users logging in and applications they run. Does that mean you should skip both? No- It means you should do the optimization when it is appropriate! Search is

a great example of an 'expensive' to run service that actually is valuable in a lot of cases but not all. So, do your research and test!

4) **You *can* run Optimizer on Control servers** (Broker/Controllers, StoreFront, PVS, etc) but be very aware when you do. Monitor event logs closely and as always… TEST BEFORE YOU DEPLOY.

Special note here- if you sign up for Citrix Smart Tools, Optimizer Checks can be scheduled for you! Learn more about that and the other system checks at https://docs.citrix.com/en-us/smart-tools/checks/about-health-checks

PVS Target Device Optimization Tool

I won't linger on this topic a lot but the very first automatic tool for Citrix VDAs was actually the PVS Target Device Optimization Tool. This is typically run when you first capture a vDisk, however- many people don't realize that the tool can be run AGAIN at any time. Here's some caveats, however:

- As mentioned before- while disabling Windows Search makes a lot of sense theoretically for PVS workloads, it is a good idea to make sure your apps will not need Search integrated (ie, Outlook).
- Also do not confuse Search with the Indexing service
- By default, the Optimization Tool will disable Windows Defender. This is not always appropriate, so please be aware!

Additionally, if you are using PVS there are several other tweaks you can test at CTP Carl Stalhood's site.

Bonus Optimizations!

Though I don't always make these recommendations during my assessments, there are a few other things that I just can't ignore to really kick the optimizations into high gear! More settings that if left at the 'default' setting hurts you in the long term, especially at scale. Entire eBooks could be written about these topics, and they move so fast with updates that I hesitated even putting them in here. If you are interested in learning more about these in detail I would love to hear from you! At the time of this writing I am considering creating an affordable deep-dive course system for these topics. Email me at CoachDJ@CTXPro.com and let me know which topics you'd like to see me expand on if you'd like to participate in a beta or initial launch of them!

Sealing the Deal: Image Creation the Right Way

Using the Windows Cleanup Tool

Every time you make changes and updates to an image, you'll want to reboot at least twice, then run the Windows Cleanup utility (cleanmgr) as an administrator to scan thru not only temp files but older updates that no longer need to be there. The process takes quite a while but it is not uncommon for me to see savings of 2-3 GB by doing this process.

Defragmentation – the Hidden Performance Thief

Another overlooked item, especially with Citrix deployments using a central image (MCS or PVS) is that while the new Cache in RAM with Overflow to Disk functionality, is defragmentation. 'But I use SSDs' you say. Good for you. But the Cache in RAM does NOT cache files. It caches blocks. So if you have fragmented blocks, you can use between 2 and 8 times as much memory to cache. This is because the memory cannot fragment, so if only part of the block is used, too bad. It still uses the whole block. So when you have made changes to the base image and are getting ready to deploy the new snapshot or vDisk version… first make sure you defragment the image as part of your sealing process.

*Quick tip on PVS- the best way to Defragment is to first clean up the image, then mount the vDisk on the PVS server, where you can defragment it as an attached disk. It is not only faster, but you are able to defragment files that normally would be in use.
Now- as I mentioned previously, this does not always apply. Here's my general guideline:

DEFRAGMENT YOUR BASE IMAGE BEFORE DEPLOYMENT!

LEAVE EVERYTHING ELSE ALONE UNLESS YOU ARE RUNNING PHYSICAL SPINNING DISK ON YOUR SERVER OR HAVE A SPECIFIC REASON TO KEEP DEFRAG RUNNING

Using Third Party Tools

Confession time. I'm often sad when I do a Citrix Consulting gig because I can't recommend Third Party… anything. The reason is solid; Citrix can't support it! That being said, I'd encourage you to test another tool from some people I know and trust. I've used their tools myself and it works very well. This next tool is one of the best!

BIS-F

BIS-F stands for Base Image Script Framework. The goal was to have a single 'master optimization' script that does everything you should be doing when you 'seal' an image for distribution from either PVS or Machine Creation Services. The list of tasks the script performs is long and distinguished… (don't finish that quote, please) but a few highlights are making sure Windows activation (KMS) is working properly, that the drive is optimized (defragmented).
Another cool feature is running a purge of WEM cache and scanning AntiVirus to mark the drive as safe to a lot of newer programs that can take advantage of this feature (less files to scan = better performance). Also, something that is missed all too frequently: .Net Optimization. What this does is initiate a process that often runs at startup. The problem with this in a non-persistent desktop, it would happen every time the image starts up. I have seen this cause massive issues for host CPU in VDI environments. Again- important details that are frequently missed. Thank goodness for scripts.

VMware Optimizer

The VMware OSOT can also be run for your workloads (if you are running on VMware only, please). Even though it will say "Horizon View" the reality is that, once again- these are OS optimizations. So as you can imagine, there is a lot of overlap. But if you are running on a VMware host (as most people are) this is a good idea! I have NOT tested OSOT on non-VMware hosts, nor do I intend to. It is up to you if you want to try it!
A caveat with OSOT is that it is community driven, so the odds of getting bad advice is always real. As always: TEST, TEST and then test again!
But more importantly, what I'll mention with the OSOT is that I was hesitant to mention it because I have on occasion found some template optimizations that prevent you from being able to deploy snapshots for both View and Citrix MCS (irony of which is not lost on me). So because the updates for this will always be more up to date than this eBook- I'm asking, *pleading* with you to do two things:

1) Read up on the template settings and consider the implications
2) As above- take each setting one at a time, making backups along the way

Antivirus Defaults – the Secret Performance Killer

Just like Microsoft's OS settings are made to fit a wide array of solutions 'out of the box' so also are AntiVirus and Anti-Malware programs. I'm a bit off topic here, but I make the recommendation enough that I thought I'd go ahead and give you a reminder here.
While they are getting better (check out Bitdefender's Hypervisor Introspection with XenServer if you have doubts about that) there is still management that needs to be done for every anti-malware software.
I encourage you to look at this Citrix article, but at a high level consider these guidelines:

- Make sure to exclude Citrix executables
- Exclude system files like the page file, print spooler and certain cache directories
- Set scanning to Write Only, especially on non-persistent MCS and PVS workloads
- Do not perform scheduled scans on MCS and PVS workloads (use the BIS-F tool or manually scan before sealing the disk)

There's more here but, but I have one more bonus tip for you.

The Hidden Resource Hog: Ads and Tracking

My friend Dan Allen has been harping on this for a long time and he's absolutely right: with everyone using web browsers all day long, the persistence of advertisements is inevitable even in the workplace. Once again, left to the defaults- the browser will simply consume all it is told to consume, even ads and tracking. But the good news is there are lot of ways this can be reduced or eliminated. Reduced to the tune of over 35% less resource consumption according to this whitepaper.
Here's a list of options, in order of difficulty to deploy:

1) Use a Group Policy Object to enable IE Tracking Protection
2) Use the Ublock plugin (trust me when I tell you to not use Adblock plus – uBlock is more efficient)
3) Attack the problem via DNS or hosts file (again- test this but I've found it VERY effective)

So in summary, remember that the default settings are not there for your benefit, but they are there for your safety. They are the compromise draft- the deal no one was happy with but work around it. If you master these, you are on your way to being a #CitrixHero!
Remember, if you want to dive deeper into this topic, I'd like to know! Let me know if you're interested in live Q&A, Online Courses or even webinars on these topics by emailing me at CoachDJ@CTXPro.com

Resources and References

As Promised- here's some additional reading and resources on Optimizations and general considerations!

- The Citrix VDI Handbook: https://docs.citrix.com/en-us/xenapp-and-xendesktop/7-15-ltsr/citrix-vdi-best-practices.html
- Microsoft RDS-VDI Optimization 'recommendations': https://docs.microsoft.com/en-us/windows-server/remote/remote-desktop-services/rds-vdi-recommendations
- Windows 10 Optimization Guide (PDF): https://support.citrix.com/article/CTX216252
- Server 2016 Optimizations from Citrix Architect Daniel Feller: https://virtualfeller.com/2017/04/11/windows-server-2016-optimizations-for-citrix-xenapp/

- Windows 8/8.1 and Server 2012/2012 R2 Optimization Guide and Script: http://pablolegorreta.com/windows-8-server-2012-optimization-guide/ (based on a CTX document that is no longer available- Pablo wrote the article for Citrix originally)

Tip #2: NetScaler Defaults

In the first chapter we discussed common mistakes that are made with OS Optimization, or more specifically a lack thereof because people leave the default settings.
In this chapter- we will talk more about the second big mistake. Similar in the persistence of default settings - but the cost this time is not performance but security!
Those of you following closely will know that this is not the first time I've talked about problems I'm seeing in the Access Layer in general, specifically the NetScaler (sorry, Citrix ADC. It will take me another year or two for that name transition). One of the biggest problems with this is how rapidly the practices change. So much so I'm hesitant to even put anything out there for fear of it being obsolete… (it already is; there are new solutions with firmware 12.1, but I haven't had time enough to validate them for distribution yet)
Of this year's 350+ individual findings, there were about 13 that were related to the Access Layer. Rather than just giving you one, I'll give them *all* to you as I believe they all work together:

- NetScaler Firmware Vulnerable to attack
- Plaintext StoreFront website vulnerable to man-in-the-middle snooping even if secured from NetScaler Gateway front end
- SSL Labs scores not passing; should be an A+
- Other Leading Practices not yet configured
- Drop Invalid HTTP requests
- Enable Selective Acknowledgement
- Configure Window Scaling
- Use TCP tuning for XenApp & XenDesktop
- Management interface on port 80 and enabled on all interfaces
- ACLs not configured
- And my personal favorite finding (twice this year): NSROOT password still set to default. Talk about an easy hack!

We have a lot to cover so let's dive in!

Layers	• Access
Success Lanes	• Design • Build
Prerequisites	• NetScaler (Citrix ADC) • Microsoft IIS • Citrix StoreFront

SCORES

Change NetScaler Password

Let's start right with my favorite… If you are running the default NSROOT password, for the love of all that is holy change it! If I need to explain why, I'll simply say that this is only a slightly bigger problem than you putting your username and password on a stickynote and sticking it to your monitor. Just change the nsroot password and secure it. Please.

For advanced points, many teams tie the administrative consoles with Active Directory accounts. But I'll be honest- if you are still using the default nsroot password, I don't think describing that particular procedure is beneficial for you.

If you aren't- two bits of advice:

1) If you do change the NSROOT password (which is a good idea on occasion) be sure it isn't used by any monitoring tools or other outside uses.
2) Enabling LDAP (ie Active Directory) authentication is a great way to not have to maintain passwords on the device. Here's how to make it happen: https://support.citrix.com/article/CTX123782 All I ask is that you take the time to design this well before implementing it! I have seen a lot of mistakes and frustration with this process. Reach out to the ctxpro.com community if you need help!

NetScaler Firmware Vulnerable to Attack

Back in September 25th, 2017 Citrix suddenly pulled the firmware releases of nearly every NetScaler from the website. The problem was a very serious flaw (CVE-2017-14602) in older firmware version going back clear to 10.1 that would allow an attacker to bypass the normal security and gain direct access to the administrative interface. Needless to say it got a lot of people's attention… but not enough. I'm STILL finding vulnerable firmware even in enterprise-level deployments that I examined in the last year.

According to CTX227928, if you have firmware at or below the following you are at risk:

- 10.1 – 135.18
- 10.5e – 60.7010.e
- 10.5 – 66.9
- 11.0 – 70.16
- 11.1 – 55.13
- 12.0 – 53.13 (except, oddly enough, 41.24)

The first part of the fix is easy enough- upgrade your firmware. The second part involves changing your ACLs, which we'll talk about in the advanced section. I find this to be optimal… but optional.

NetScaler Gateway Not Scoring an A+ at SSLLabs.com

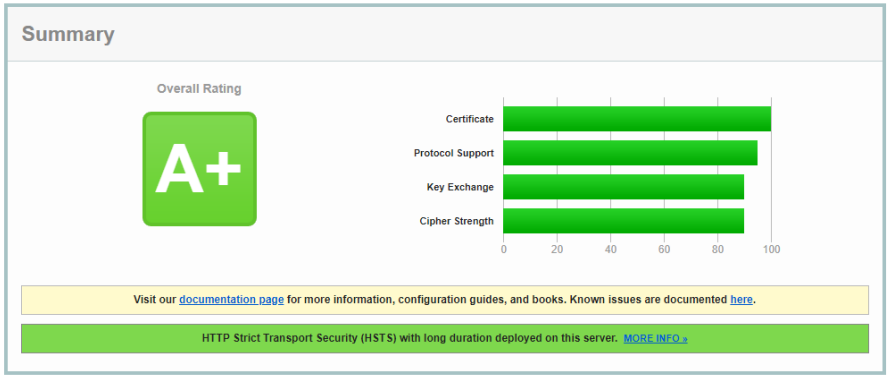

I have actually written about this a few times in 2016 and again in 2017 but basically if you have an external-facing SSL virtual server (NetScaler Gateway or otherwise) you really want to strive for an A+ score to make sure your attack surface is lowered externally. Fortunately, even following the 2016 instructions will still net you a good score and to date my testing has still shown customers using the 2017 advice are still getting A+ scores. Citrix published an article in 2018 as well that is worth the read. However, with the adoption of TLS 1.3 and a few other updates, expect some scoring changes coming up soon, but you know what I'm already seeing out there?

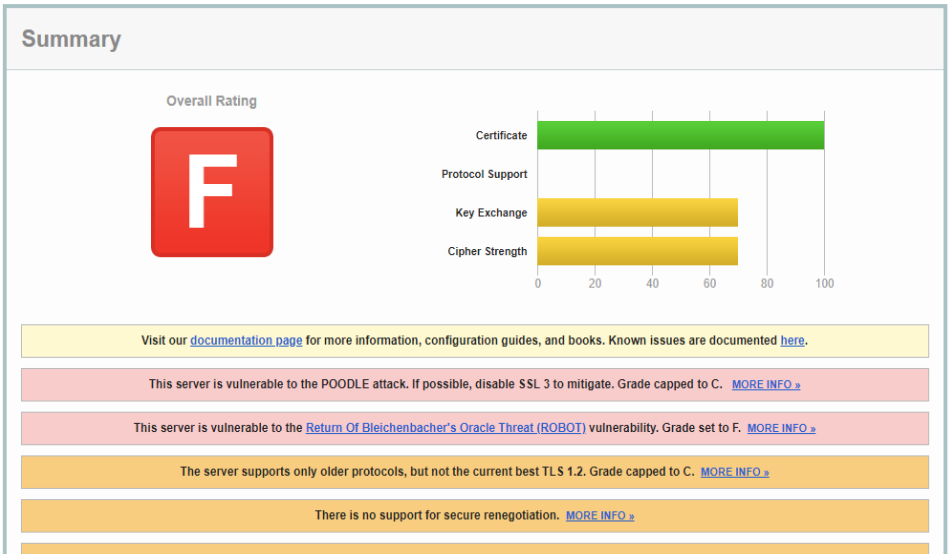

Figure 1: If you see this- time to either fix it or hire someone who can.

That's right. 7 of 10 of the sites I observed in 2018 were still vulnerable in a way that was visible to the world (the SSLLabs test is done externally on what the world sees). In several of these cases the F was from the firmware problem I mentioned above.
This is a massive topic- because I've already covered it please do visit the links above to get full details! Here's what you'll need:

- Perform an analysis at SSLLabs.com and look at the score you are getting and why.
- Keep an eye out for very bad things like still using SSL2 or 3, really old Cipher Suites and vulnerabilities like BEAST, POODLE, and others. The ideal Cipher suite combinations are an interesting discussion and it honestly depends on both your version of NetScaler and what your endpoint devices will support. The good news is that the latest NetScaler firmware builds support a secure profile that takes care of this for you!
- Make sure your SSL server is not vulnerable to ROBOT. If you upgraded to firmware above the ones I listed previously you should be fine.
- If your endpoint devices will support it, remove support for SSL3, TLS1 and TLS 1.1 (this is crucial to test and have a support statement around)
- Check the detection of Secure Renegotiation and Forward Secrecy (PFS). Those setting may indicate a need to adjust your TCP settings on the NetScaler as I recommend.
- Always run an analysis of the NetScaler's support file at the Citrix CIS website (https://cis.citrix.com) to look for other problem spots.
- Disable Client and server side SSL Renegotiation https://support.citrix.com/article/CTX123680 and enable HSTS https://support.citrix.com/article/CTX224172 if possible.

Which is all great. But literally the day I am writing this I found a new rating…

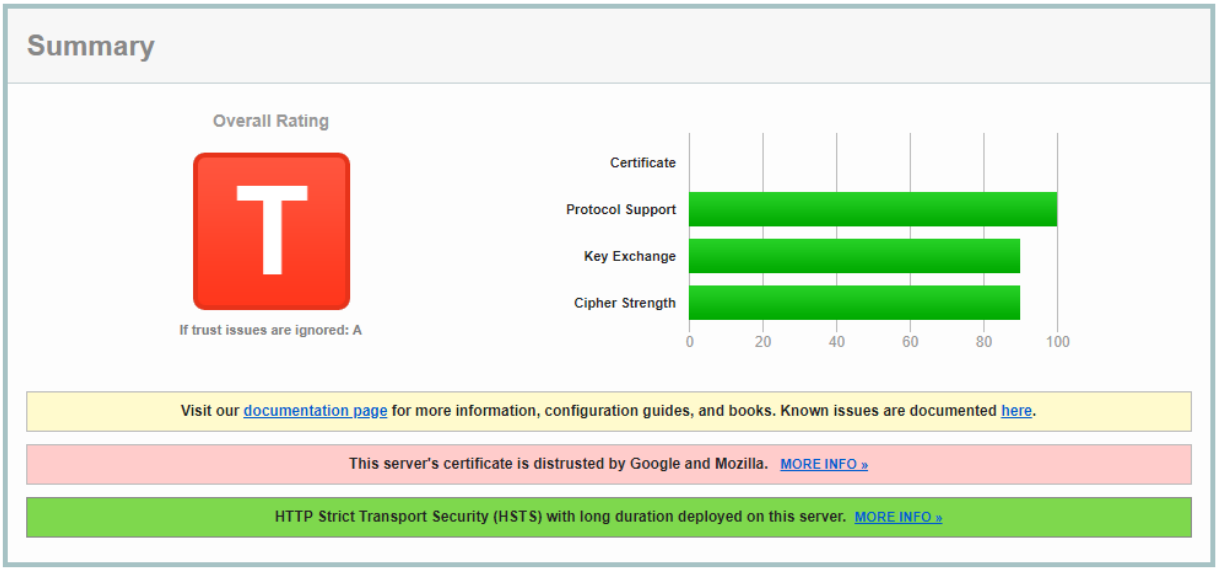

Why? If you are using a certificate that can't be trusted by Google (typically a Symantec certificate of a certain generation) you may, forgive the pun, have Trust Issues.

SSL Performance and Security

An additional word about SSL Ciphers. The security element of cipher sets is largely talked about the most… however, a key you should know is that performance is also very important. Specifically, performance in the case of what I'm talking about – the ICA protocol.

In other words, ciphers for general web traffic may be different from a performance perspective than those for NetScaler Gateway. Be sure when researching that you take those into consideration. And honestly- I don't want to muddy the water with detail about the way these things work.

That would be yet another book!

That, and it changes really rapidly.

That and there are also implications for MPX vs VPX and differences with VPX on SDX… throw in FIPS and you've got yourself basically a cocktail party with an accountant, movie stuntman, PETA activist and Donald Trump. Everyone's going to have an opinion about how to best approach things.

Okay maybe not that bad.

But I tell you what- what I really want you to do is get in the habit of checking the Citrix documentation for the particular firmware version you have. For example, check out the Ciphers available topic for 12.1.

The possibilities are nearly endless. I will say if you're interested in learning more, looking more into ECDSA ciphers may be an exciting topic. Not to mention learning more about DTLS and the proper cipher support for that functionality (if you support teams overseas, for example- you really should learn more about DTLS because that is what allows UDP based protocols like EDT to work over SSL.

For now, my suggestion is to keep it simple until you know for a fact all endpoints will support the advanced ciphers. Fortunately, the SSLLabs testing tool will actually show you how the ciphers configured for the tested URL are performing. Even a slightly better cipher group is better than the usual defaults. Again, defaults are meant to be there to support the widest range of devices and configurations, not to protect you.

Leading Practice Configuration

Just as in Chapter 1- the defaults for the NetScaler are meant for initial deployment. Once you have deployed, the idea is that you'd lock the device down. This is easy to miss, and would explain why over 90% of customers I've assessed over the last 3 years have missed it. But not you, #CitrixHero! You can find more information about these settings at https://support.citrix.com/article/CTX121149 - but this is crucial: TEST and be aware of your settings. While enabling Nagle's Algorithm and Enabling Selective Acknowledgement are beneficial and rarely have effects, settings such as Window Scaling and dropping invalid HTTP packets may need further examination based on your networking setup. 9 times out of 10, however- I find that all the recommended settings make sense.

Lock Down Management Interfaces

Again, the default setting for NetScaler is to allow for setup to occur cleanly with the assumption it will be locked down. So, a lot of people miss that they have management functionality on interfaces where this is not intended. Disable Enable Management Access control, Telnet, SSH, and GUI on all Non-Management IPs. More information here about restricting NSIPs to Only allow management applications: https://support.citrix.com/article/CTX126736 Also see the "Enable Secure Access to NetScaler GUI" at https://support.citrix.com/article/CTX111531.

Optional – Define Interface ACLs

One of the best ways to defend against future attacks is to lock down the administrative (management) interfaces so that only certain IPs are allowed. We call this defining Access Control Lists (ACLs). Obviously, be careful with this setting. See my 2017 article for more details about how to design for this, but basically you can set IPs or a range of IPs as trusted utility servers or PCs to allow access to the admin functions. Everything else is blocked, including internal requests, making it very secure. Note- you'll need to do this on all NetScalers as this setting does not get copied via HA.

Other Access Layer Issues

StoreFront Using Plaintext

This issue actually affects two different areas I will point out today. The first is more obvious; when you have a NetScaler Gateway you are using port 443. Great! But what I'm seeing out there is that the StoreFront connection to the NetScaler is still using port 80 (Plaintext HTML). This makes this communication susceptible to Man in the Middle attacks – where information can be intercepted mid-stream but still passed along. The problem here? Personally identifiable information including passwords can pass along this path.

So sure, it is a little more work to create SSL to the internal StoreFront servers but if you are using them, secure them at all points!

Bonus: XML Service Using Plaintext

The second way that StoreFront uses plaintext is, believe it or not, worse. The XML services that live on controllers are an aging but still viable method by which the StoreFront (formerly, WebInterface) sends information to the Controller regarding the user and the XML service communicates back to the StoreFront server as to which icons to load… *and a token to login to a server.* So, if intercepted, a slew of things could happen far too unpleasant to go into. Fortunately, this is rare and would only really occur in an inside attack… but my notion is to always treat endpoint devices as threats and secure the network as much as possible!

In older days you typically had IIS installed to your Controllers- if you do the instructions to secure this traffic is here: https://support.citrix.com/article/CTX200415

But if you do not have IIS installed (in an ideal world, you wouldn't) you will probably want to examine my article on the topic here: https://ctxpro.com/securing-citrix-broker-xml-service-without-iis/

To give a general idea, your goal is to bind a certificate using a command line for each controller, but make sure to use a certificate that can be either loaded or trusted by the StoreFront servers.

As you can imagine, there are dozens more things I typically recommend on these topics- and I don't even consider myself an authority on NetScaler. However, keep an eye out on ctxpro.com. If you for some reason aren't subscribed to the newsletter, you really should be. I'm organizing educational events for my subscribers and members several times over the next year.

Tip #3: Workload Placement and Sizing

Okay great, you've tuned the OS of the VMs and your access layer. But for around 60% of the companies I assessed last year, another big problem was robbing them of value: not paying attention to the special requirements and tunings for Citrix Virtual Apps within their physical server (hypervisor) setups. While this can often become an almost religious debate- I stand behind these recommendations, having seen improvements in dozens of environments since 2010 when I started drawing a hard line on doing things this way. You can debate with whitepapers all day- but the bottom line is that I've seen this work even when it seems to make no sense. Sometimes you just must swallow your pride and trust, and that's what I'm asking a lot of you with doubts to do now. Trust that myself and over 100 other experts have it right here. Hypervisor tuning matters. You might have a fight ahead if you don't run the physical servers.

But like any hero, a #CitrixHero sometimes has to fight for what is right.

Layers	• Cloud
Success Lanes	• Design • Build
Prerequisites	• Hypervisor • BIOS and Physical Hardware

SCORES

Workload Placement

What I'm seeing emerge once again these days is an old practice which I thought for sure we had done away with: All the VMs in a single cluster and let the Hypervisor sort them out. This presents a… host… of problems (the puns will continue until Leading Practices are followed folks):

1) You are no longer able to predict how much adding new users will cost in terms of hardware.
2) You cannot accurately predict how many Server VDA VMs you will need.
3) You cannot predictably assure performance from one user to the next due to other workloads that co-habitat the same host. For example, when a SQL server goes into freakout mode on the same CPU as your Server VDAs- you'll have users complaining even though CPU is not showing any signs of issue.
4) Different workloads can tolerate different overcommit ratios. With a mixed workload style you may have your user workloads on hosts that are actually overcommitted.

I am sure that even as people were reading my statement above tension started to show up. Doubts. But the reality is that user-based workloads behave very differently than the average hypervisor administrator realizes. An RDSH server (XenApp) with 30 users is going to behave very differently than a SQL or Exchange server. So placing them on the same host means you are making the hypervisor essentially pick favorites. But placing the same kind of workloads on a physical host has a kind of magic, reducing the conflicts and more importantly making the scale predictable. We'll talk more about how many VMs and their configuration later- a very important aspect to avoid oversubscription. But to do that properly, we first need to make sure that our Resources will always have the right backing and no conflicts. This can't be done without isolation.

The solution to this first problem is to isolate your user workloads from the backend servers completely. This is typically done by dedicating Resource clusters for Server VDAs and Desktop VDAs with all Control components located in the main infrastructure cluster. This has the benefit of allowing you many times to use a different licensing, version or even type of hypervisor for your resources which typically do not need things like HA and backups. If your team has skillsets of supporting multiple hypervisor types but you primarily use VMware, you may even want to consider a higher-performing but simplified hypervisor solution for the Resource hosts such as XenServer (Citrix Hypervisor) or Nutanix Acropolis. Even Hyper-V. However- I do NOT recommend learning a new hypervisor just for this purpose. If your team only knows one hypervisor well, keep the course and just simplify the Resource cluster.

If you don't have a lot of hosts available or are concerned about keeping N+1 for two clusters rather than one, the second method is to use Host DRS (or it's equivalent). This will place the workloads to preferred hosts when they boot. So in most cases you would have two Host DRS groups- Control and Resource. In the event of a failure you will still be able to load VMs onto the other hosts temporarily, but only when capacity is exceeded. I'm seeing this option a lot in smaller and mid-sized environments but I never recommend this if you have more than 6 hosts. At that point, you're usually better off isolating workloads completely. You should also keep in mind that this method may require some 'babysitting' during maintenance to be sure workloads end up back on their intended hosts.

Workload Sizing

Now to sizing. This is the other area I'm seeing massive amounts of fail lately. Someone read a whitepaper that said to configure their XenApp servers with 4 vCPU and 8 GB RAM… and wonder why they can only get 15 users before it starts slowing down. The reality here is that you need to default to scaling UP with your users per VM until you reach a performance threshold and then scale OUT with more servers. The way to do this is to increase the resources per VM – namely CPU, Memory and

Storage. So how do you know how much you can use? This is an 'it depends' answer if there ever was one. But once again, this is why we isolate workloads. We need to know how many CPU cores and RAM we have comfortably available (N+1 or N+2 is typically the threshold here).
So let's say we have hosts with two 14-core CPUs and 256 GB RAM. To be safe, we say that about 200 GB RAM is our safely available amount, to give the hypervisor some room and some 'just in case'. We also ALWAYS want our user workloads, be it Desktop or Server OS, to fully reserve the RAM in the hypervisor. This prevents the usage of a paging disk (required in VMware if you don't use this option, sucking down storage space) and is… you guessed it, another reason to isolate the workloads. We'll use Server OS for our workload examples.
Some feasible options in terms of memory (just examples- you can use any memory value you want in Windows these days) are:

- 6 VMs at 32 GB
- 4 VMs at 48 GB
- 12 VMs at 16 GB

Next we turn to CPU. Here's where it gets interesting. An often overlooked bit of math that you must do to be successful in terms of Citrix is division. I won't go into the full explanation of NUMA vs UMA but in a nutshell you want your workloads to co-habitate the same physical CPU as much as possible. This prevents the need for crossing memory bus lanes and a host of other implications that can slow down the hypervisor itself.
We do this by setting our VMs to a value that matches the CPU's NUMA values. You should always confirm these, but they are typically divisible numbers of the physical (not virtual) cores. I will tell you that I chose the 14 core processor for a very good reason- it only has 4 valid vCPU NUMA values: 1, 2, 7 and 14. While you *can* use less- research has shown that you are better off having fewer VMs with more CPUs than you are with more VMs with smaller CPUs because of the way the hypervisor essentially is forced to arrange workloads on the CPUs.
So in our case, we know that we are best served by configuring our VMs with 7 vCPUs – so because this will be a much larger VM, we are looking at either 32 or 48 GB RAM. To determine the right sizing we really need to know our tolerable CPU Overcommit ratio. In most cases I have observed, 1.2:1 is about right for Server OS, whereas Desktop OS can often scale to 5:1 or up to 12:1 in some cases I have seen. Again- 'it depends'. For us, we know that 1.2x28 (the total of physical CPU cores on the machine) is 33.6. Dividing by 7 gives us 4.8, which we safely round DOWN to 4.
So- your CPU Blade should have FOUR VMs with 7 vCPU and 48 GB RAM.

BUT… how many users per VM?

This is another common Citrix mistake! We shouldn't really care about users per VM anywhere near as much as we need to determine how many users per PHYSICAL BLADE/HOST.
Think of it this way. If I was to put a pair of virtual machines on your laptop- how many users do you think would safely be able to use it before it became, well, unusable? Even if these are VMs intended for a lot or not a lot of users, the physical hardware can only do so much. The hypervisor doesn't magically fix this. Any host will have limits that must be anticipated and respected. So now that you've isolated your workloads- good news! You can figure out those numbers easily by simple division- Number of users per host divided by the number of VMs. Now, keep in mind that mileage may vary here based on the applications and Operating system… however there is one mistake that you should NOT make:

If Per-VM performance is bad because of the amount of users, you need to adjust the number of HOSTS, not the number of VMs. Adding VMs to a physical host will degrade performance for all users because the physical limitations have not changed.

The Rule of 5 and 10

Here's the thing. Unless you have the tools, time and patience to figure out EXACTLY how many users to have per blade- you need a simple rule to start with, then further optimize from there.

"More what you'd call guidelines than actual rules…"

-Hector Barbosa (Pirates of the Caribbean)

Thanks here go to Citrix Consulting and Nick Rintalan for figuring the math out on this one. As a matter of fact, him and I worked together for a while trying to figure out a very overly complicated spreadsheet, so I was really pleased when he noticed the way the numbers always seemed to work out. I could go into the underlying arithmetic here, but let's keep it simple. You can determine how many users can be on a host based on the *physical* CPU cores. Not the threads (virtual), but Physical.

- Desktop OS workloads: 5 users per pCPU core
- Server OS workloads: 10 users per pCPU core

So again in our example, blades with two 14-core processors (28 pCores) we can expect 280 Server OS users or 140 VDI users. Note, this is *active* users, not VMs. In our example, it means we should expect a maximum of 70 users per VM. While this is possible- it may not always be practical. Load testing is important to determine the exact number.

Now that we know the hardware can handle it, we end up in the domain of the OS itself. Typically I have seen a properly sized and tuned Server 2012 R2 VM safely handle easily 120 users running fairly light apps, but when using a published desktop the numbers dropped to about 50 users per VM. Sometimes, you may need to add an additional VM- just keep to the NUMA values and try not to exceed the overcommit ratio and you'll usually be fine. Test, Test, Test!

So if we have 3000 users, we know we need 11 host servers (+1 for redundancy = 12). And, as an added #CitrixHero moment, when the VP of IT asks you how much it will cost to add another 500 users… you'll be able to give them a REAL answer! Scalability is fun!

The Citrix article is https://www.citrix.com/blogs/2017/03/20/citrix-scalability-the-rule-of-5-and-10/

So the question you're asking yourself is does this work in the real world? Yes. Yes it does.

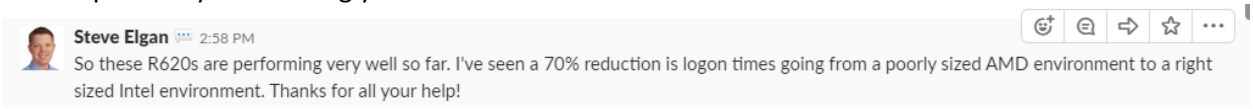

BONUS Recommendation: Hardware Virtualization Settings

I was going to save this for another chapter, just because it has the kind of impact that can't be denied- but it really is quite simple. So, you get it here, right now.

Several times this year I was surprised to find that the underlying BIOS settings for physical hosts were not set correctly. Depending on the setting, this can be a MASSIVE problem, so while it isn't common… I thought I'd include this because of its additional impact.

If you find your VMs are slow and running 100% much of the time, check your BIOS settings. Actually, scratch that. *Just go check your BIOS settings anyway*. You may find that C-States are enabled. This

seems like such a great idea… until you put a Hypervisor on top of it. Once again… they ship for compatibility. You have to tune them once you get them! (sensing a theme here?)

Common Mistakes (not just Citrix, but any Hypervisor):

- C-States enabled (this allows the CPU to throttle and changes the CPU percentage calculation)
- Virtualization not enabled (breaks the hypervisor's ability to function as, well, a hypervisor)
- Hyperthreading not enabled (yes, the rule of 5 and 10 assumed HT is on)
- Power settings not set properly (typically should be the "Power" or "no power management" setting)
- Also very much related to the NUMA settings above, make sure to check for the correct QPI 'snoop' modes and don't always trust the 'auto' setting to give you Cluster On-Die. Failure to do this, for example on a 14 core processor can be problematic (this processor tends to array the silicon in a 6+8 configuration and uses COD to present the proper NUMA values).

Power Management Matters

An article that is definitely worth reviewing is from Jasper Geelen (LoginVSI):
https://www.loginvsi.com/blog/834-influence-of-power-management-on-vdi-performance
In a nutshell Jasper notes that power management settings specifically are different per vendor, as are the names used- but almost all allow options to either let the OS handle power or various other settings. While you can let your hypervisor control these settings, you should know that it isn't typically dynamic. My recommendation is to simply lock the BIOS in it's highest power state for any VDI or RDSH workloads. In some cases, this simply means turning power management off completely; this depends on your vendor.

> *"Faulty power management is the most common but easiest to fix VDI mistake. Configuring this properly can save your users a lot of energy (and) user experience will increase"*
>
> *Mark Plettenberg, LoginVSI and fellow Citrix CTA*

LoginVSI found that this can be a difference of 64% of performance.
Go ahead and read that again, I'll wait.
This means that in a great many cases more than half of the available performance the machines expect simply isn't there. And this holds up in the real world, in my case even *better* than the synthetic testing showed.
My favorite case of this last year was one of those times when an assessment being performed paid for itself three times over, simply because we caught this one thing that an engineer had made the assumption was correct. The client was able to cancel an order for over $45,000 because upon enabling the power management correction, their problems with VMs reporting 100% CPU went away. They were immediately able to more than double the number of users on each blade and still had better performance than previously there. When combined with other tuning suggestions, they estimate that they will be able to go another three years without additional purchases. So it seems simple enough, but one miss or assumption can literally cost that much.
So, save yourself the $20,000 consulting bill and double-check these settings.
Or- you know, don't… and call me. I'll take the money.

Additional Reading

My friend Helge Klein talked about this in 2013 using HP servers as an example.
If you have a Cisco UCS system, I highly suggest having a look at their article, paying special attention to Table 4.

Conclusion

On behalf of my team and the members of CTXPro.com, thanks for taking the time to read this eBook. As I glance at the word count just under 10,000 I realize I may have gotten a bit long-winded about it. This project was an expansion of three much smaller blog posts and personally, I'm glad I did this!

I would appreciate your feedback- it will determine if I continue this project with additional resources and refinements. It will also help me determine if a subscription service or engaging sponsors for this series will make sense.

Finally- make sure you are subscribed to my email list. I have a lot of fun webinars coming up- one of which is my "Dirty Dozen" where I describe nine more of the top problems I'm finding on assessments. It is a lot of fun to open people's eyes to all the problems that exist and just how common a lot of them really are!

Cheers!

-DJ Eshelman
September 2018
Franklin, TN