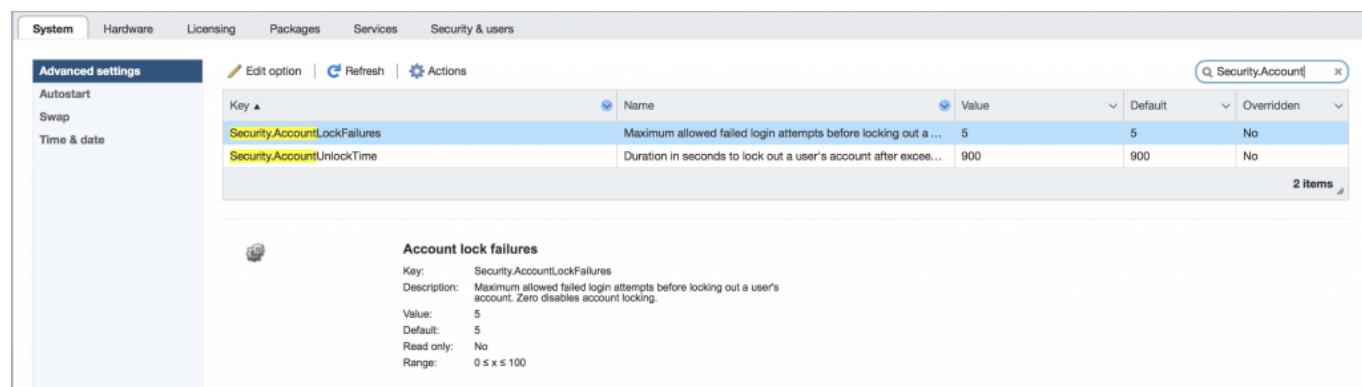


[Resetting vSphere 6.x ESXi Account Lockouts via SSH](#)

VMware vSphere has had a good security feature added since vSphere ESXi 6.0 to add a root account lockout for safety. After a number of failed login attempts, the server will trigger a lockout. This is a good safety measure for when you have public facing servers and is even important for internally exposed servers on your corporate network. We can't always assume that it's external bad actors who are the only ones attempting to breach your devices.

Using the vSphere web client shows us the settings which are used to define the lockout count and duration. The parameters under the Advanced settings are as follows:

Security.AccountLockFailures
Security.AccountUnlockTime



The screenshot shows the vSphere web client interface for the 'Security & users' section. The 'Advanced settings' tab is active, displaying a table of configuration items. Two items are listed:

Key	Name	Value	Default	Overridden
Security.AccountLockFailures	Maximum allowed failed login attempts before locking out a...	5	5	No
Security.AccountUnlockTime	Duration in seconds to lock out a user's account after excee...	900	900	No

Below the table, there is a detailed view for 'Account lock failures' with the following details:

- Key: Security.AccountLockFailures
- Description: Maximum allowed failed login attempts before locking out a user's account. Zero disables account locking.
- Value: 5
- Default: 5
- Read only: No
- Range: 0 ≤ x ≤ 100

Resetting your Failed Login Attempts with pam_tally2

There is a rather simple but effective tool to help you do this. It's called `pam_tally2` and is baked in with your ESXi installation. The command line to clear the lockout status and reset the count to zero for an account is shown here with the root account as an example:

```
pam_tally2 --user root --reset
```

```
vSphere Security documentation for more information.
[[root@turbostack01:~] pam_tally2 --user root --reset
Login          Failures Latest failure      From
root           2754    09/14/17 13:29:54  61.167.12.10
[[root@turbostack01:~]
```

In order to gain access to do this, you will need to have SSH access or console access to your server. Console access could be at a physical or virtual console. For SSH access, you need to use SSH keys to make sure that you won't fall victim to the lockouts for administrative users. In fact, this should be a standard practice. Setting up the SSH keys is relatively simple and is nicely documented in the Knowledge Base article **Allowing SSH access to ESXi/ESX hosts with public/private key authentication (1002866)**

https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externa

[Iid=1002866](#)

Uploading a key can be done with the `vifs` command as shown here:

<https://docs.vmware.com/en/VMware-vSphere/6.0/com.vmware.vsphere.security.doc/GUID-392ADDE9-FD3B-49A2-BF64-4ACBB60EB149.html>

The real question will come as to why you have the interface exposed publicly. This is a deeper question that we have to make sure to ask ourselves at all times. It's generally not recommended as you can imagine. Ensuring you always use complex passwords and 2-factor authentication is another layer which we will explore. Hopefully this quick tip to safely reset your accounts for login is a good first step.

[The Need for IT Operations Agility: Lessons of WannaCry](#)

There is little doubt that the news of ransomware like the recent outbreak of the [WannaCry \(aka Wcry, WannaCrypt\)](#) taking hold in critical infrastructure hits home with every IT professional. The list of affected clients of any ransomware or critical vulnerability is made even more frightening when it means the shutting down of services which could [literally affect people's health like the NHS is experiencing](#).

Would it be any different if it were a small hardware chain? What if it was a bank? What if it was your bank, and your money was now inaccessible because of it? The problem just became very real when you thought about that, didn't it?

Know Your (Agile) Enemy

Organizations are struggling with the concept of more rapid delivery of services. We often hear that the greatest enemy of many products is status quo. It becomes even more challenging when we have bad actors who are successfully adopting practices to deliver faster and to iterate continuously. We aren't talking Lorenzo Lamas and Jean Claude Van Damme kind of bad actors, but the kind who will lock down hospital IT infrastructure putting lives at risk in search of ransom.

While I'm writing this, the WannaCry ransomware has already evolved and morphed into something more resilient to the protections that we had thought could prevent it from spreading or taking hold in the first place. We don't know who originally wrote the ransomware but we do know that in the time that we have been watching it that it has been getting stronger. As quickly as we thought we were fighting it off by reducing the attack surface,

The Risks of Moving Slowly

Larger organizations are often fighting the idea of risks of moving quickly with things like patching and version updates across their infrastructure. There are plenty of stories about an operating system patch or some server firmware that was implemented on the heels of its release to find out that it took down systems or impacted them negatively in one way or another. We don't count or

remember the hundred or thousands of patches that went well, but we sure do remember the ones that went wrong. Especially when they make the news.

This is where we face a conundrum. Many believe that having a conservative approach to deploying patches and updates is the safer way to go. Those folks view the risk of deploying an errant patch as the greater worry versus the risk of having a vulnerability exposed to a bad actor. We sometimes hear that because it's in the confines of a private data center with a firewall at the ingress, that the attack surface is reduced. That's like saying there are armor piercing bullets, but we just hope that nobody who comes after us has them.

Hope is not a strategy. That's more than just a witty statement. That's a fact.

Becoming and Agile IT Operations Team

Being agile on the IT operations side of things isn't about daily standups. It's about real agile practices including test-drive infrastructure and embracing platforms and practices that let us confidently adopt patches and software at a faster rate. A few key factors to think about include:

- Version Control for your infrastructure environment
- Snapshots, backups, and overall Business Continuity protections
- Automation and orchestration for continuous configuration management
- Automation and orchestration at all layers of the stack

There will be an onslaught of vendors using the WannaCry as part of their pitch to help drive the value of their protection products up. They are not wrong in leveraging this opportunity. The reality is that we have been riding the wave of using hope as a strategy. When it works, we feel comfortable. When it fails, there is nobody to blame except for those of us who have accepted moving slowly as an acceptable risk.

Having a snapshot, restore point, or some quickly accessible clone of a system will be a saving grace in the event of infection or data loss. There are practices needed to be wrapped around it. The tool is not the solution, but it enables us to create the methods to use the tool as a full solution.

Automation and orchestration are needed at every layer. Not just for putting infrastructure and applications out to begin with, but for continuous configuration management. There is no way that we can fight off vulnerabilities using practices that require human intervention throughout the remediation process. The more we automate, the more we can build recovery procedures and practices to enable clean rollbacks in the event of a bad patch as well as a bad actor.

Adapting IT Infrastructure to be Disposable

It's my firm belief that we should have disposable infrastructure wherever possible. That also means we have to enable operations practices which mean we can lose portions of the infrastructure either by accident, incident, or on purpose, with minimal effect on the continuation of production services. These disposable IT assets (software and hardware) enable us to create a full stack, automated infrastructure, and to protect and provide resilience with a high level of safety.

We all hope that we won't be on the wrong side of a vulnerability. Having experienced it myself, I changed the way that I approach every aspect of IT infrastructure. From the hardware to the application layers, we have the ability to protect against such vulnerabilities. Small changes can have big effects. Now is always the time to adapt to prepare for it. Don't be caught out when we know what the risks are.

Thinking Like the Bad Actors and Prioritizing Security

Assume you've been breached. Period.

The reason that I start there is because I've learned from practice, that we have to work on the assumption that we have had our systems violated in one way or another. The reason that this is important is that we have to start with a mindset to both discover the violation, and prevent it in future.

Who is it that has breached our systems? Well, we have a fun name for them...

Bad Actors



Hey, I like Kirk too, but you have to admit...he's not really a good actor

No, not the kind that you see in SyFy remakes of popular movies, but the ones that have been infiltrating your infrastructure for nefarious purposes. Bad actors are those who have the single-minded purpose of breaching your security, and doing something either inside the environment, or taking something back out.

All too often we hear about breaches long after they have happened. I'm a big fan of [Troy Hunt's](#) web site [Have I Been Pwned?](#) It's a helpful resource, and a reminder of just how important it is that we understand that bad actors exist and are pervasive in the world of internet connected resources.

Bad actors love the internet of things. Just imagine how much simpler it is to access resources when they are interconnected and internet accessible. Physical security is the first place to look, and all the way up the stack to the application layers. Using your mobile to access your bank site when you're in Starbucks? Not a good idea. Seem paranoid to say that? That's what every bad actor hopes you say.

Assume security is failed. Assume you've been breached. The next step comes with how you plan and prepare to discover and recover.

White Hat (aka Ethical) Hacking

Just under a year ago, I attended the [BSides Delaware event](#). This was a very interesting opportunity to go outside of the normal conference circuit that I am used to attending. I would liken this to the VMUG equivalent where DefCon is the VMworld of security. These are great events, and touch on every aspect of security from application, to network, to physical, and even security of yourself including self-defense tactics.

One thing that you learn about hacking, is that it takes a hacker to find and prevent a hacker. White Hat hacking has been a practice for many years, and it is an important part of the security and networking ecosystem. If you aren't already engaging an organization to help with penetration testing or some form of security analysis, you absolutely should.

The same skills that drive the bad actors have been embraced by white hat hackers to provide a positive result from that experience. We use real users to provide UX guidance, so it only makes sense that we should use the same methodology for our security strategy.

Make Security Part of Infrastructure Lifecycle

Whether it's your application lifecycle, or your infrastructure deployment, security and automated testing should very definitely be a part of the workflow. I was lucky to have a great conversation on my Green Circle Live! podcast recently with Edward Haletky.



We chatted about how there is a fundamental flaw in both the home and the data center. The whole podcast is a must listen if you ask me, and I encourage folks to rethink security as something that should be top of mind, not an after thought.

There are lots of bad actors out there. I prefer to keep them in the movies and out of my data, how about you?

SDN challenges - “You can keep your networking gear. Period.”

You may recall a statement regarding some big U.S. legislation that led us to the forever quoted phrase: “You can keep your insurance. Period.” that has caused quite a ruckus in the insurance industry both for providers and customers because it was found to be untrue.

So just imagine that a similar situation that is about to come up in the enterprise networking environment. With Software Defined Networking (SDN) being the hottest buzzword and most aggressively marketed paradigm shift in recent months, we are about to hit a crossroads where adoption may leave many customers taking on unexpected costs despite being pitched a similar line that SDN will simply run as an overlay, but you can keep your existing networking hardware.

Let’s take a look three particular challenges which are present as companies take a look at SDN and figuring out the cost/benefit and how it relates to existing infrastructure.

Challenge 1 - No reduction of ports

This is one of the most common misconceptions around SDN. The idea that ports will be reduced is unfounded because the number of uplinks that will exist into host systems, virtualized or not, will continue to be the same. If anything, we will have more uplinks as scale-out commodity nodes are utilized in the data center to spread the workloads around more.

The reduction in ports will happen as a result of the migration to higher speed ports like 40GbE and up, but the consolidation level will be limited for physical endpoints. SDN is a great enabler for creating and leveraging overlay networks and making physical configuration less of a factor in the logical design of the application workloads.

In order to get the savings on per-port utilization, the move to 40GbE and higher ports will trigger the rollover of existing hardware and expansion to new physical networking platforms. In other words, you need to change your existing hardware. Hmmm...that wasn’t in the original plan.

Another interesting shift in networking is the new physical topology which includes ToR (Top of Rack) switches which are connected to a centralized core infrastructure. The leaf-spine design is being more widely used and continues to prove itself as an ideal way for separation of workloads and effective physical isolation which has other benefits also.

Challenge 2 - Policy-based delivery requires policies

This is the business process part that can add a real challenge for some organizations. Putting a policy-based framework into place is only truly going to add value when you have business policies that can leverage it. Many CRM and Service Desk implementations fail because of the lack of adoption which stems from a lack of understanding of existing processes.

Many organizations are having difficulty adapting to cloud implementations because it is a very process-oriented technology. As more and more companies make the move to embrace cloud practices, the move towards SDN will be more natural. There is much more awareness now about where the efforts are needed to make SDN deployments successful.

Challenge 3 - Your physical gear doesn't support your SDN platform

Other than the previous limitations where we mentioned the port speed issues for higher consolidation levels, there is also the issue of firmware and software capability on existing ASIC hardware. As an example, you can use Cisco ACI as your SDN product of choice, but if you are running all Cisco Catalyst equipment I have some bad news for you. (*UPDATE 11/21*: Thanks to [@jonisick](#) for the tip that there are smaller physical investments to allow the use of ACI. It is not a full rip and replace, but more some additional hardware to augment the current deployment in most cases).

There will be a barrier to entry for many SDN products because there are requirements for baseline levels of hardware and firmware to support the enhancements that SDN brings. This will be less of an issue in a few years I am sure, but for right now the move to embrace an SDN architecture may be held back by the need to upgrade physical hardware to prepare.

Have No Fear! SDN will work...No seriously, it will

While these scenarios may be current, realistic barriers to the adoption of a SDN platform, we are also dealing with hardware and software lifecycles that are becoming shorter and more adaptive.

The hardware platforms you are running today are inevitably going to be upgraded, extended, or replaced within a reasonable time frame. During that time we will also see the shift the way that we manage and deploy the networking inside organizations. This fundamental shift in process will align with the wider acceptance of SDN platforms which are being regarded as only accessible to agile organizations sometimes.

What SDN brings to us is really the commoditization of the underlying physical hardware platforms. Not necessarily the reduction of quality or cost of the hardware, but the commoditization of its role in the networking architecture.

What is important for us all as technologists is that we are prepared for the arrival of these new products and methodologies. We have a responsibility to stay ahead of the curve as much as possible to get to the real benefit of SDN which is to enable agility for your business.