THE FOUR GOLDEN METRICS IN VULNERABILITY MANAGEMENT

A NUCLEUS SECURITY WHITEPAPER

Nucleus

FOUR THINGS: KEEPING METRICS SIMPLE.

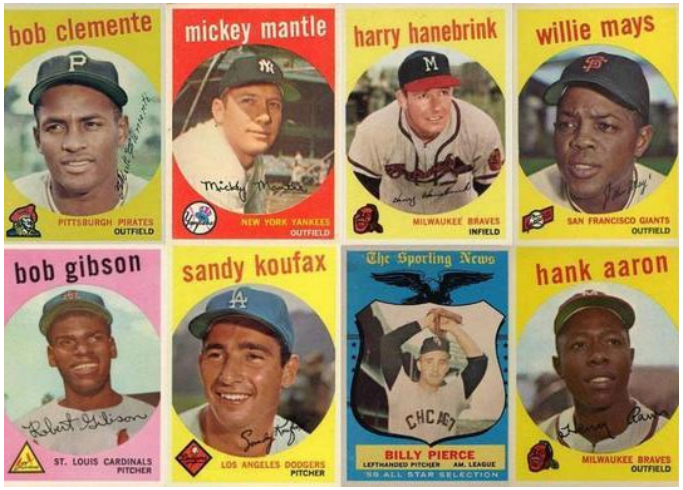

Vulnerability management is a lot like baseball. And if you don't like baseball, or, like much of our audience, hail from a country that doesn't even *play* baseball, I'm talking to you. Because both of them can be extremely difficult to understand or explain – and very easy to argue about. The key to understanding either one is to simplify them. You can simplify either one down to about four statistics or KPIs. Learning how to simplify VM down to four KPI's transformed my career, so I want to share those with you.

TAKING BASEBALL'S EXAMPLE

This rite of passage ended with my generation, but when I was a kid, every child I knew collected baseball cards. If you look on the back of a baseball card, it's an absolute mess of statistics, and with no explanation of what any of them mean. I think we were supposed to ask Dad. There's one question that the back of a baseball card should answer that it rarely does: *is this a good player?*

When you watch baseball on television, when a player steps up to bat, they flash some statistics on the screen. But it doesn't look like the back of a baseball card. They give you about four numbers – and those are generally the four accepted measures that tell you whether the person at the plate is a good player or not.

You can also look at the wrong numbers, or take them out of context, and arrive at a wrong conclusion. Reggie Jackson struck out more times than anyone else who ever played the game. That means he swung and missed more times than anyone, which isn't a good thing in a game whose objective is to hit the ball. So, I could conclude that he was the worst baseball player who ever played. But Reggie Jackson is generally considered one of the best players of his generation, because when he *did* hit the ball, he hit it a very long way. **Context** is important.

If you show me the right numbers, I can tell you in a split second if someone is a good baseball player. The realization that I could do the same thing in the computer security field of vulnerability management changed my career.

KEEPING VM **SIMPLE** WITH STATISTICS.

I can answer the question of whether you have a functioning vulnerability management program with four stats:

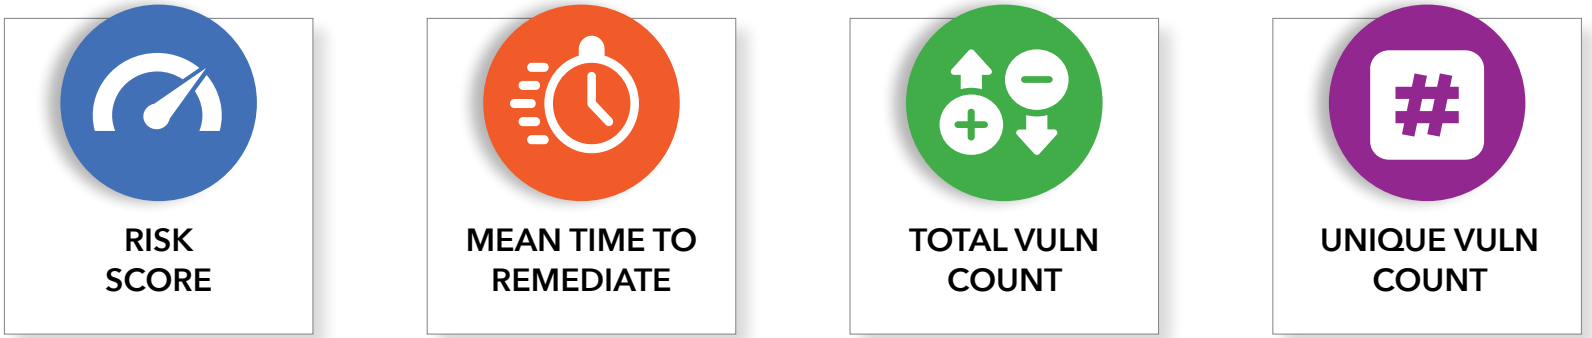

For that matter, as a Security Analyst, those four statistics are enough to tell me at a high level what we need to do to have a better month next month.

That's a big deal, as one key job of a Vulnerability Analyst is to coach remediation teams through this data. By looking at a team's top fixes and these four stats, a good Vulnerability Analyst should be able to give a simple course of action to allow any team to have a better month next month. And the awesome thing is, these metrics are simple enough to prove to their not-as-technical bosses that they are indeed getting better over time.

And remember, you're on the same team here. Security Analysts and System Administrators often hate each other more than the Boston Red Sox hate the New York Yankees - typically due to misunderstandings. Effective communication, and managing our own expectations, goes a long way toward helping us win together.

To that end, let's talk about each statistic and what they tells us.

RISK SCORE.

A Risk Score is a measure of your typical vulnerability. It's a number between one and 1000, where average is 500 and lower is better. I've received the question everywhere I've worked whether it's possible to reverse this to where higher is better – and the answer is no. The security industry has always measured vulnerabilities on a low-to-high scale, where higher is worse. We measure severity or risk as low, medium, high, or critical, so the numeric scale needs to also follow that nomenclature.

Risk score is useful because it gives you a measure of what your typical vulnerability looks like in your organization. As you get familiar with the vulnerabilities in your environment, those numbers will mean something to you. If I say the words Beverly Hills, certain things come to mind. A 5-digit number is probably one of them.

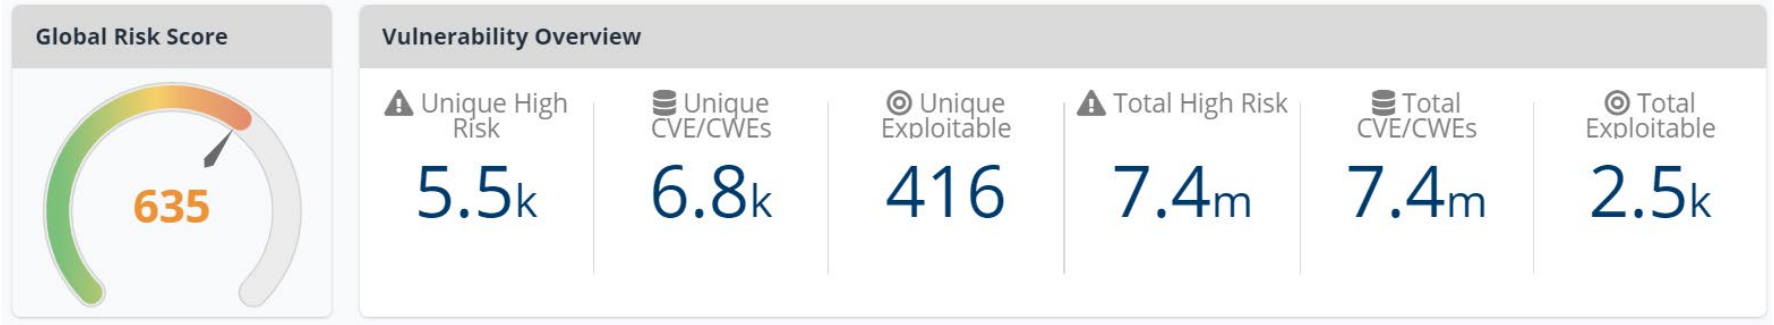

In Nucleus, Risk Score is an average. If you're fixing your most critical vulnerabilities, that number will naturally trend down over time. If you're fixing stuff haphazardly, your score will reflect that approach over time. If nothing else, Risk Score tells you where to start. The biggest problem in vulnerability management is knowing what to fix first. If you fix items that have a score higher than your overall risk score, you won't go wrong.

BEHIND THE NUMBER.

Some people are content to just use the score, while others want to know what goes into it. The Risk Score is generally a composite of vulnerability attributes (what CVSS calls "Metrics"), the business or asset context that you provide, and threat intelligence.

Many of these components tend to be stable over time, as the vulnerability attributes rarely change. The business or asset context won't change unless you change it. But vulnerability intelligence is constantly changing. The appearance of new exploits, and exploits gaining or falling out of favor with attackers, can cause scores to change dramatically between scan or reporting cycles. It is very common for a vulnerability to have a low risk score when it first appears, only to progress up the ladder to medium, high, and ultimately, critical. The difference between a low and a critical rating can just be how reliable the exploits are and how popular they are with attackers.

It's also possible for a vulnerability's risk score to fall. If attackers stop using exploits against a particular vulnerability, that vulnerability's risk score will fall. This is a bit unusual, but if a new vulnerability appears and it gives similar results to an existing one, but works more reliably, or makes less noise on the network, attackers will naturally migrate to that new exploit, the same way we migrate from older software to newer software that works better for us.

ONE **CAVEAT** TO RISK SCORES.

Under some odd circumstances, it's possible for you to do exactly the right thing, and have your overall risk score go up. There are usually two explanations for this, and both involve factors that are beyond your control. This is why risk score isn't the end all be all statistic. It's useful, but I need more than risk score to explain what's going on in any given vulnerability management program.

One way this happens is when you deploy an update that supersedes other updates. You deploy one update, and a dozen other vulnerabilities go away. But imagine a scenario where you deploy a critical update, and that update supersedes three updates of low criticality. Since your risk score is an average, getting rid of those three lows and only one critical may cause your overall average to go up rather than down. This problem is more likely to occur with new customers than with established ones. Unfortunately, this kind of discouragement is usually the last thing a new customer needs.

It's important to not get discouraged when this happens, because it's also common for a critical update to supersede other criticals and highs. If you keep deploying the right updates, your risk score will inevitably trend downward over time. I'm much more concerned about what your risk score has done for the last three months than what it did this month.

The other issue that can cause your risk score to go up in spite of patching happens when a new high criticality vulnerability appears during your patching cycle, and is severe and/or numerous enough to overshadow the remediation work that you did that month. This is more likely to happen in a mature program.

Provided you can prove this is a temporary situation, I'm fine with a risk score being bad for a month, but I need other numbers to confirm it's a circumstantial occurrence.

MEAN TIME TO REMEDIATE.

This statistic answers two questions. Are you patching? If this is a tangible number, and it's not increasing by 30 days every month, then the answer is yes. The second question is, are you patching fast enough? You probably have some idea how long you're willing to have a vulnerability exist in your environment. If your policy is to fix vulnerabilities within 30 days, and your average time to remediate is 60 days, then you're not following policy. If the average time to remediate is 30 days, then that means you are following that policy at least some of the time. If the average time to remediate is below 30 days, that means you're following that policy most of the time.

You may have different policies for each level of severity. Especially early on, it's important that you do, so I like to measure this statistic for each level of criticality. This statistic can be useful in shaping your policy. If your teams cannot hit their numbers, they may have constraints that prevent them from complying with your policy. In that case, you have two choices. Loosen your policy, or remove those constraints. You can probably guess which of those things is usually easier.

Let's take a (purely hypothetical) situation of a company with a policy that critical vulnerabilities need to be fixed within seven days. They only have one maintenance window per month, and they have a large backlog of critical vulnerabilities that built up over years. It's not realistic to fix a problem that accumulated over the course of years in a single month, let alone a week. In this case, relaxing the policy temporarily is more realistic than getting a maintenance window every week and hiring a bunch of people for the remediation team. You can come back and revise the policy once the backlog is under control.

Having this kind of flexibility - to accommodate for circumstances that will not fit the policy and need exception - is a goodwill gesture that makes it easier to get buy in across teams. Taking a hardline stance rarely leads to good things.

TOTAL VULNERABILITY COUNT.

Raw vulnerability counts are perhaps the most popular statistic. They're also problematic, yet they do have uses – particularly psychological ones. I estimate I fixed 800,000 vulnerabilities during my System Admin career. That's a number almost anyone can understand. My MTTR or risk score requires more explanation.

I've used vulnerability counts to get people promoted. At a previous job, one action by a System Admin reduced our vulnerability count by 25% in a single month. The CISO noticed, and he told the CIO. That awesome Admin got a raise, and eventually a promotion. He knew that I made sure the CISO knew he deserved the credit. I'm pretty sure if I were to call him and tell him I was in jail and needed him to bail me out, he would do that for me. He would tell me I was an idiot, but he would help me, because I helped him once.

This count is more useful if you break it down by severity, and count both open and closed vulnerabilities. If I tell you that you have a million open vulnerabilities, your mind goes someplace – probably not a good place. But if I also tell you that you've fixed 10 million vulnerabilities, your confidence is probably solid. If you've fixed 10 million, I know you're capable of fixing another million.

I'm also a lot less concerned if those million vulnerabilities are low severity than I would be if the majority of them are high or critical. Some people take a stance that "a vuln is a vuln". That stance is not conducive to getting things done – you have to prioritize patching to be productive. Obviously, the usefulness of this statistic changes over time. Initially, it's most useful for convincing yourself you have a problem you need to fix. Once you have a track record of fixing things, it becomes a tool for convincing yourself what's possible.

UNIQUE VULNERABILITY COUNT.

Security people understand vulnerability counts – they think of each vulnerability as a unique entity that an attacker may be able to exploit. As Security Analysts, that's our job to think that way.

System Administrators think in terms of tasks – and while each vulnerability could be treated as a separate task, you'll never get it fixed if you do. To remediate large numbers of vulnerabilities at scale, you have to automate. The way patching tools typically work, you can find an update, and instruct the tool to deploy it to all applicable systems. It's the same amount of work whether the update applies to one system or 100,000 systems. Your unique vulnerability count is a much more realistic measure of how much work needs to be done than the overall vulnerability count.

Your unique vulnerability count will probably approximate the number of updates you need to deploy to run your vulnerability count to 0. It likely won't be exact, but it's close enough for estimation purposes.

I've been in monthly metrics meetings where the remediation teams and the security team got into arguments because the security team didn't have this number. This is the number remediation teams think in terms of, and since the security team had the data that made sense to them, the two teams just talked past each other. While this number isn't terribly useful to a security team, it helps the remediation team estimate how much work remains to be done. So, from that perspective, if the remediation team acts on this number, it could end up being the most useful statistic of all.

KEEP IT SIMPLE.

There are, of course, other statistics that can be useful. For example, knowing your number of systems can certainly come in handy. Vulnerability density, which is the number of vulnerabilities divided by the number of systems, is another one I've used in the past. If I have a million vulnerabilities spread out over half a million systems, that paints a very different picture than if I have a million vulnerabilities spread out over 1,000 systems.

But at the core of this whitepaper, I want to caution you against over-complicating your key guiding metrics. There are hundreds of statistics you could collect, but that doesn't mean it's a good idea. I have seen companies build dashboards in other tools that scroll on for several screen lengths. These dashboards look impressive, but they aren't very useful.

Many other statistics may seem like a good idea, but you need to be able to answer and account for what that statistic tells you. Vulnerability management is really about answering three key questions: are you patching, are you patching effectively, and how much work remains? If a statistic doesn't help you answer one of those three questions, and doesn't identify a blocker standing in the way of patching effectively, there's not really any value in collecting it.

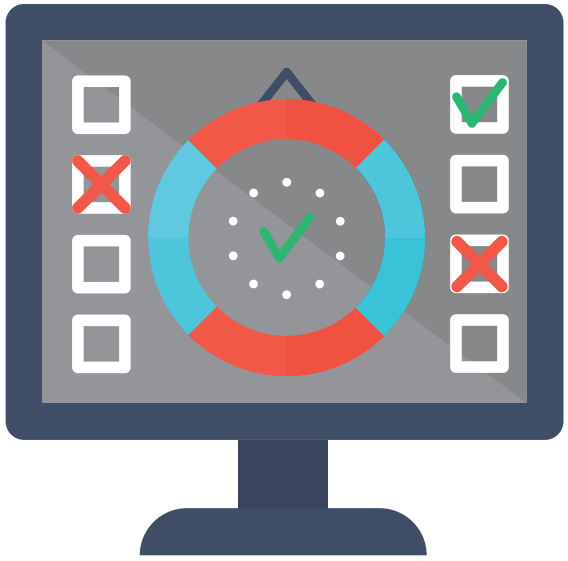

The simpler you can keep it, the more effective you are going to be. Collecting lots of statistics will make you seem smart, and being smart is a good thing. But being understood is an even better thing – and is more likely to lead to results than simply being smart.

ABOUT THE AUTHOR.

Dave Farquhar, CISSP. Dave is a security veteran and Solutions Architect at Nucleus Security. Certified in Qualys, Tenable, and Rapid7 tools, Dave has experience deploying and running all three tools either independently or in conjunction with a risk aggregation tool such as Nucleus. As a System Administrator, Dave patched 800,000 vulnerabilities, and helps today's System Administrators do even better.

TAKE THE **NEXT STEP** WITH NUCLEUS.

Keep track of these four core metrics and much, much more with Nucleus. See the power of our unified VM platform for yourself, on your own time, with a Demo on Demand.

WATCH A DEMO