KENNA Security

Cyentia INSTITUTE

PRIORITIZATION TO PREDICTION

Volume 5: In Search of Assets at Risk



KENNA Security

PRIORITIZATION TO PREDICTION

VOLUME 5: IN SEARCH OF ASSETS AT RISK

Overview & Key Findings 2
Recap of Prior Research 3
Facets of Assets at Risk. 7
– Asset Prevalence 7
– Vulnerability Density 9
– Remediation Coverage 12
– Remediation Velocity. 13
– Remediation Capacity 17
Reviewing What We've Learned 19

This research was commissioned by Kenna Security. Kenna collected and provided the dataset to the Cyentia Institute for independent analysis for this report.

Kenna Security is a leader in predictive cyber risk. The Kenna Security solution enables organizations to work cross-functionally to determine and remediate cyber risks. Kenna leverages machine learning and data science to track and predict real-world exploitations, empowering security teams to focus on what matters most. Headquartered in San Francisco, Kenna counts among its customers many Fortune 100 companies, and serves nearly every major vertical. Find out more at www.kennasecurity.com.

Cyentia INSTITUTE

Analysis for this report was provided by the Cyentia Institute. Cyentia is a research and data science firm working to advance cybersecurity knowledge and practice. We do this by partnering with vendors and other organizations to publish a range of high-quality, data-driven content like this study.

Find out more at www.cyentia.com.

Overview & Key Findings

> "The missile knows where it is at all times. It knows this because it knows where it isn't ... In the event the position it's in is not the position it wasn't, the system has acquired a variation ... If variation is considered significant, it may be corrected ... However, the missile must also know where it is."
>
> *–Adapted from Air Force training on missile guidance*

Apparently, missile guidance and vulnerability management (VM) have something in common: They're both confusing processes that require accurate situational awareness. Indeed, targeting critical risk exposures across a complex and evolving enterprise IT landscape isn't easy. Sometimes you know exactly where risk is. Sometimes you just know where it isn't or at least where it wasn't. Sometimes variation between where risk should be and isn't requires corrective action, but guiding that action depends on knowing where risk actually is. It seems like nonsense, but it's a very real reason why many VM programs veer off target when it comes to remediating assets most at risk.

So far in the Prediction to Prioritization (P2P) series, we've taken a vulnerability-centric view of remediation. That was necessary to properly cover all the bases. But experience and customer feedback suggest that vulnerability management is often asset-centric (or at least asset-influenced). The fact that we manage vulnerabilities in assets rather than in a vacuum requires us to know where risk isn't, where it is now, and where it will eventually be—and which specific risks warrant corrective action.

Because of this, we look at the vulnerability risk landscape through the lens of common asset platforms in this fifth volume of the P2P series. We revisit classic P2P metrics like coverage and capacity but also derive some new measures unique to assets to help VM programs stay on target. Below you'll find a sampling of what we learned.

The typical organization manages vulnerabilities across about 800 active assets, but 10% manage over 35,000. The full range of assets per firm extends from less than 10 to over 1 million.

Windows 10 PCs are the most common asset managed by VM programs, with Linux systems and Cisco appliances rounding out the top three. Overall, half of firms have an asset mix consisting of at least 85% Windows-based systems.

Windows platforms typically have 119 native and third-party vulnerabilities detected in any given month. That's almost 4 times the median number of bugs in the next closest asset category of Macs and 30 times that of network appliances!

70% of Windows systems, 40% of Linux/Unix systems, and 30% of network appliances have at least one open vulnerability with known exploits. If attackers need just one foothold to infiltrate the network, it seems opportunities abound.

Organizations fix the majority of security issues affecting assets, but do so most rapidly for Windows and Mac OS X. The half-life of vulnerabilities in a Windows system is 36 days. For network appliances, that jumps more than tenfold to 369 days!

Good support is critical. A full 50% of the vulnerabilities affecting older Windows platforms no longer supported by Microsoft (e.g., 2000, XP, and 2003 Server) remain open well beyond two years. Free tech refresh justification inside.

It's easier to fix native software. 68% of Microsoft bugs on Windows-based assets are squashed in the first month, while only 30% of non-Microsoft vulnerabilities on those same assets are fixed in that timeframe. Just say no to bloatware.

Good news: VM programs stay ahead of new vulnerabilities for the majority of assets across all platforms. But it's not easy. 32% of Microsoft systems fall behind in the remediation race, while less than 10% of network devices fail to keep up.

Recap of Prior Research

Over the last two years, we've published four volumes of research and learned a lot about vulnerability remediation. While not everything in that body of work is prerequisite to the new material we cover in this report, a refresher on what we've learned so far will help set the stage for what comes next. So, here are FIVE things readers old and new alike should know about vulnerability remediation.

1: There are A LOT of published vulnerabilities

First, it must be acknowledged that a ton of vulnerabilities exist that potentially represent risk to organizations and consumers. Over 130,000 have been published in the National Vulnerability Database (NVD), and many more exist that haven't been officially recognized. Based on sheer volume alone, it's clear why so many vulnerability management programs drown under the deluge.

Figure 1: Number of CVEs added to NVD from 1999 through 2019

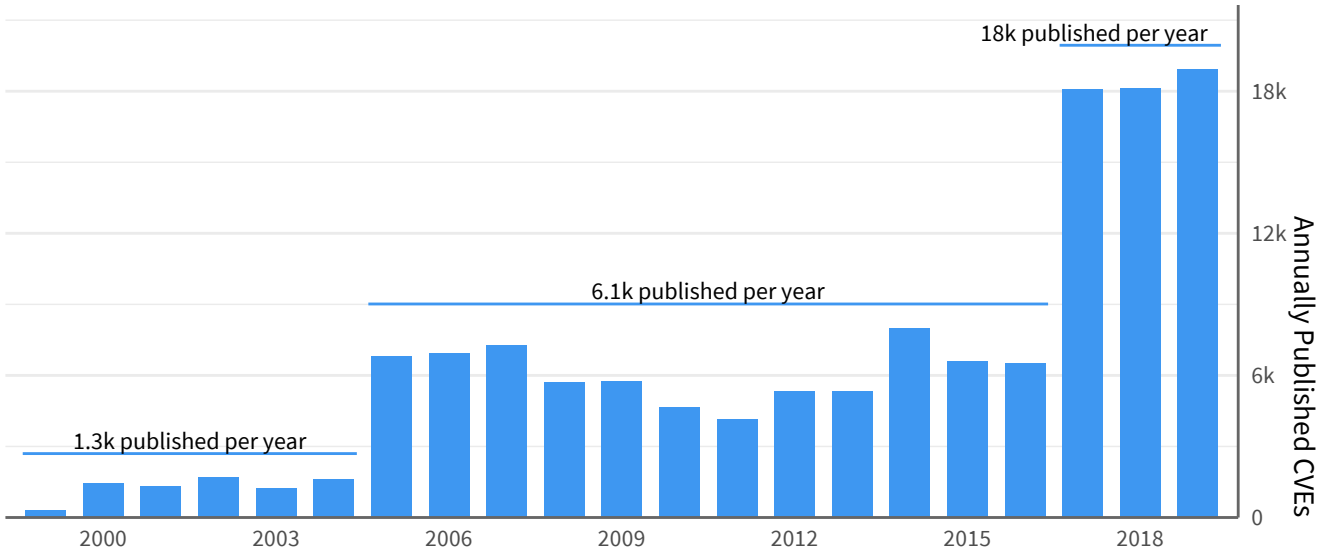


Figure 1 charts the number of vulnerabilities published to the Common Vulnerabilities and Exposures (CVE) list. With the expansion of the CNA process in 2017, the rate of new CVE entries tripled. While that increase may seem scary, remember that this is more a measurement of the CVE process and less about the inherent security of software and hardware over time.

A quick word on non-CVE vulns: It's true that focusing on CVEs misses vulnerabilities that, for one reason or another, fall outside the CVE process. We've clearly stated that limitation from the very first volume in this series. But how much of a limitation is that, actually? We collected data this time to answer that very question.

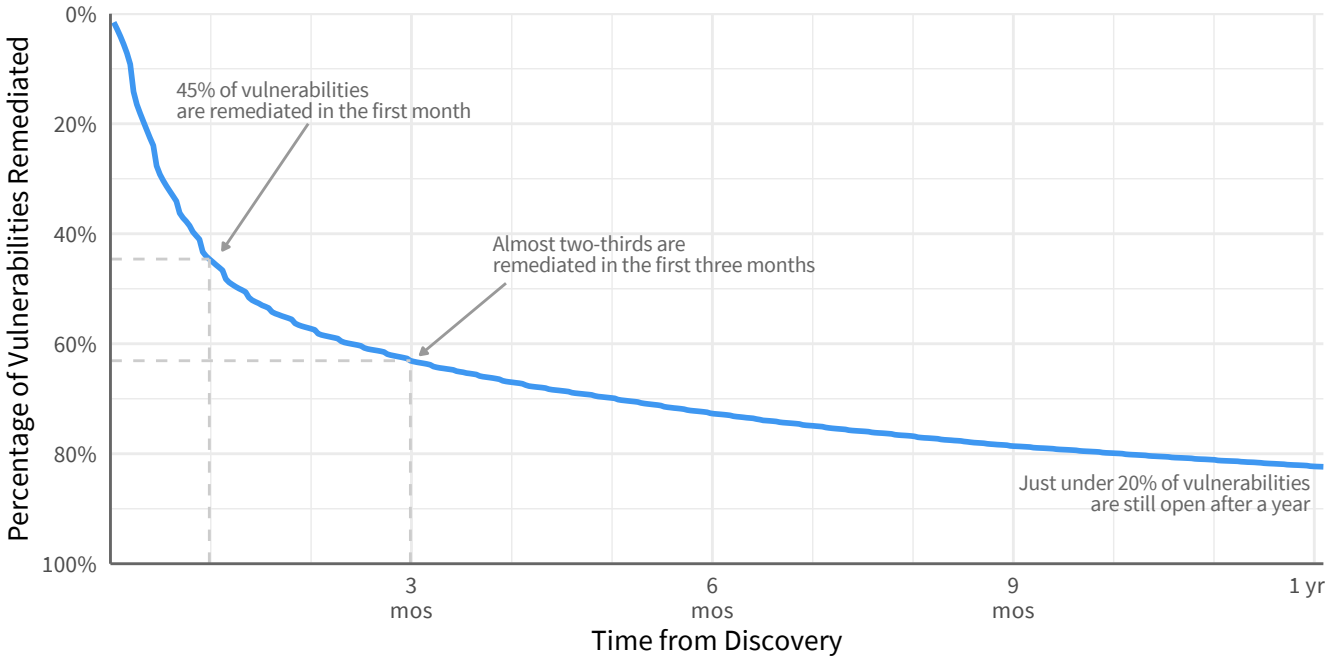
Turns out, about 95% of all observed vulnerabilities in our dataset have a CVE. Interestingly, it's much lower for network appliances and devices (73%). But the overall lesson is that CVEs offer good coverage of the vulnerabilities identified by the various scanners used by organizations around the world as a foundation of their VM programs.

2: Remediating vulnerabilities can take A LOT of time

As if dealing with 130,000 vulnerabilities wasn't enough, it gets much worse when remediating them in a live environment. It's not simply a matter of squashing a bug and you're done; it's a process of finding and fixing every asset affected by each vulnerability. Because so many systems are affected, remediating security exposures can be a complex, lengthy process. Figure 2 adds some numbers around just how long it can take.

We analyzed time-to-remediation across hundreds of firms and learned that 45% of vulnerabilities are closed in the first month, two-thirds are closed within three months, and just under 20% stick around longer than a year. All the dismissive "just go patch it" types out there should keep this in mind.

Figure 2: Overall remediation velocity across active assets for all firms and vulnerabilities



Note that the remediation timeline shown in Figure 2 is an aggregate view. P2P Volume 3 showed very different curves for different types of organizations and software. We'll break this out across asset platforms later in this report.

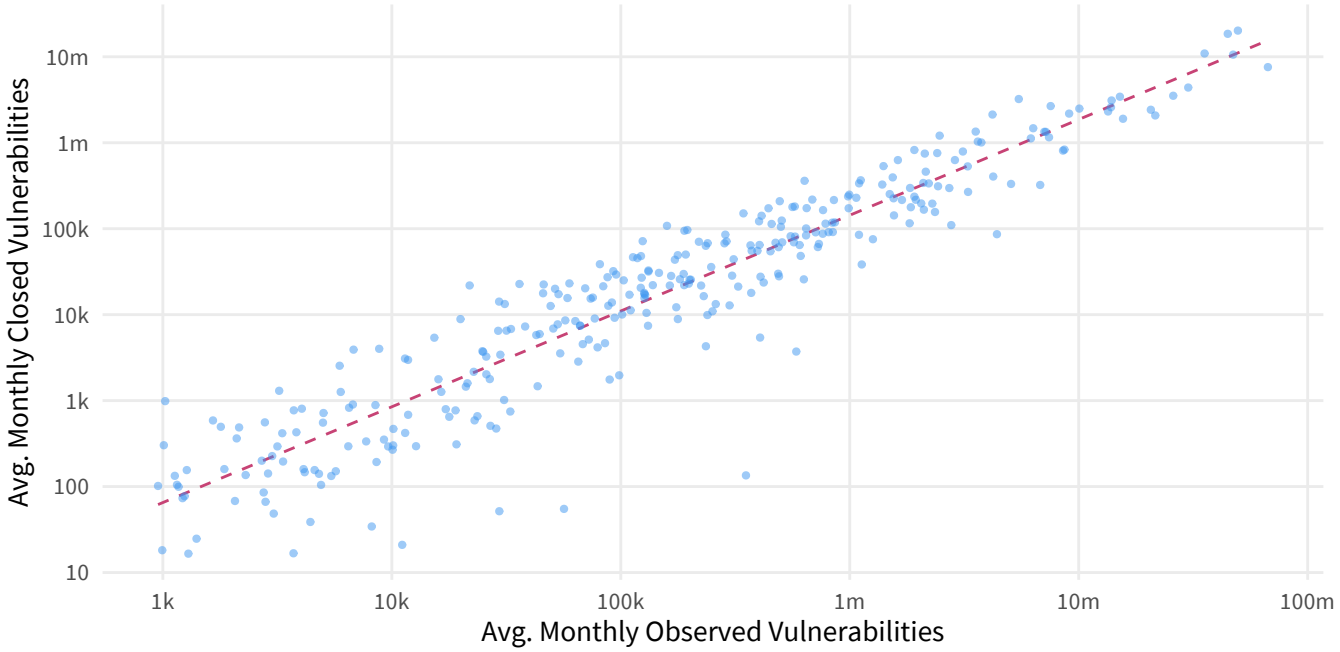
3: Organizations cannot fix ALL vulnerabilities

Given the volume of vulnerabilities affecting infrastructure and the time it takes to fix them, it's not surprising that firms can't consistently remediate them all. We measured the average number of observed (i.e., detected by a scanner) and closed vulnerabilities across hundreds of organizational environments over a two-year period. The critically important outcome is presented in Figure 3.

P2P Volume 3 is chock full of vulnerability survival curves like Figure 2 across organizational dimensions of size and industry.

Every dot in Figure 3 marks the number of observed and closed vulnerabilities for a particular organization. If you follow the red trendline, it's apparent that, on average, firms fix about 1 in 10 issues in their environment. Some firms manage to beat those odds (above the line) and others fare worse, but that overall ratio is remarkably consistent across SMBs and enterprises alike.

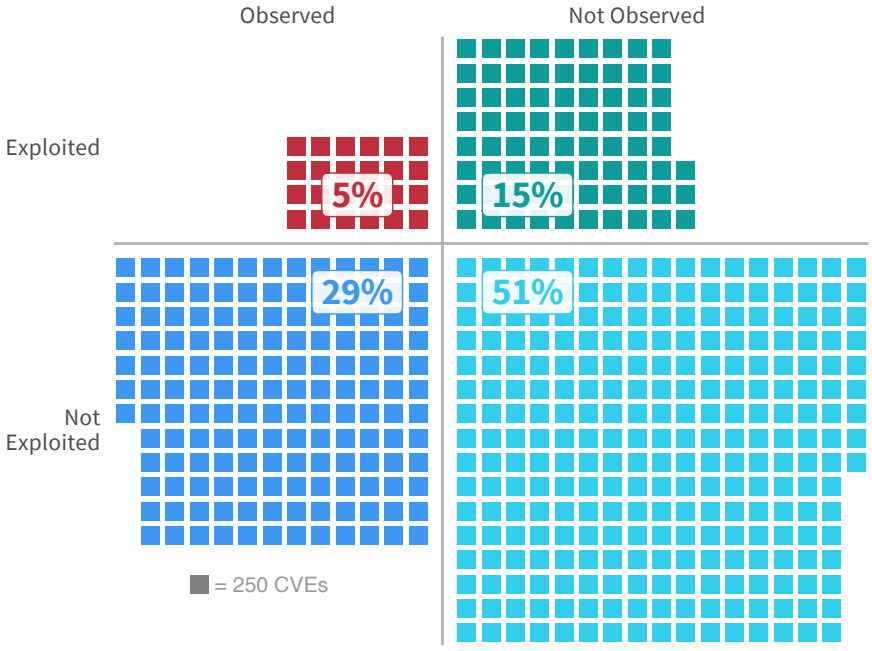
Figure 3: Average number of observed vs. closed vulnerabilities per month



4: Not ALL vulnerabilities need to be fixed right now

It may seem like vulnerability management is a hopeless venture from the charts above, but we found some genuine signs of hope in the data. Firms can't fix everything, but the reality is that they don't need to. The bar representing real risk is quite a bit lower. Figure 4 illustrates this crucial point by separating all published CVEs into a simplified "risk matrix."

Figure 4: Ratio of observed/not observed and exploited/not exploited vulnerabilities



We first separate vulnerabilities into those actually observed in a live environment and those that were never seen by a vulnerability scanner or similar tool. On the right, two-thirds of published CVEs were never detected by the hundreds of firms in our sample. That suggests many vulnerabilities affect technologies not currently used in most enterprise networks.

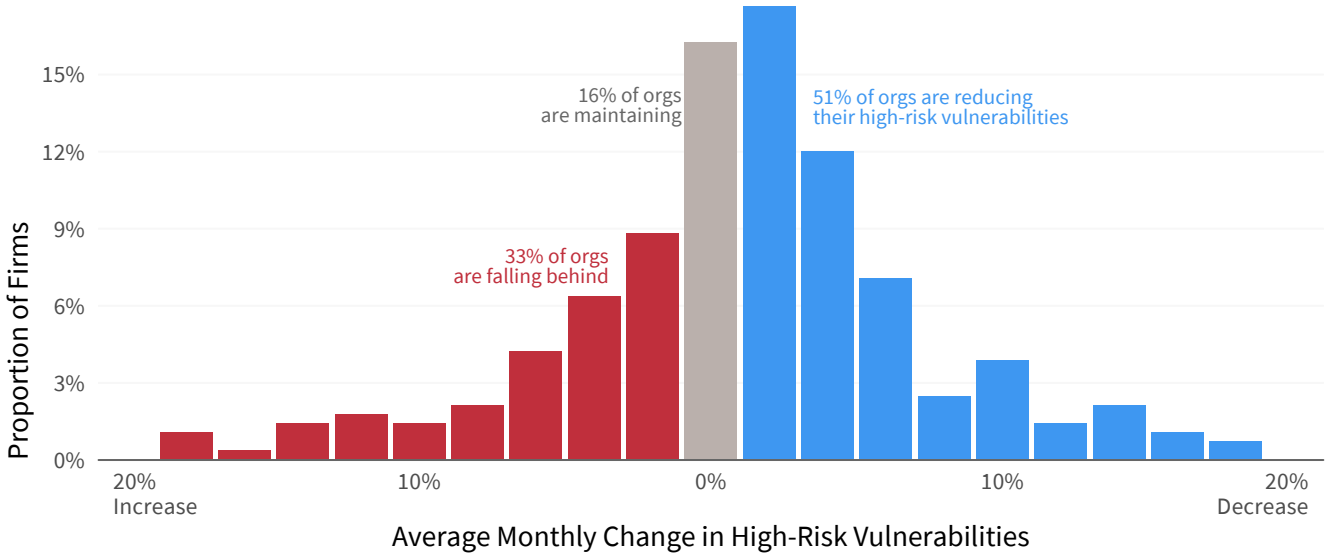
We then separate CVEs into those that have been exploited and those with no known exploits. Even though 20% of published CVEs have a clear threat (either actively exploited in the wild or a published exploit exists), only about 5% of them represent real risk right now for most firms. While it's certainly true that new exploits emerge regularly, it's also true that most firms can safely deprioritize vulnerabilities without known exploits. When you can't fix everything, fixing what matters most is critical.

And that's where we see hope. Our research shows that organizations able to maintain this risk-based focus remediate vulnerabilities better, faster, and more efficiently than those with less-focused programs.

5: Firms can fix all HIGH-RISK vulnerabilities

Thus far, we've learned two key facts: 1) firms can't fix all their vulnerabilities and 2) only a small proportion of vulnerabilities have known exploits. Thus, it's theoretically possible that organizations might be able to remediate all high-risk vulnerabilities across their environment, assuming laser focus. They might even have room to fix vulnerabilities that aren't yet exploited but probably will be in the future! If you're curious about predicting future exploits, check out our Exploit Prediction Scoring System.

Figure 5: Net remediation capacity for high-risk vulnerabilities among firms



Theoretical capacity is one thing. The important question is whether we see evidence that organizations actually accomplish this feat. We do, and Figure 5 proves it. Among the hundreds of firms examined, just over half of them reduced the number of high-risk vulnerabilities in their environment. Another 16% held their ground, meaning two-thirds of organizations (Kenna customers, at least) are successfully managing vulnerability risk in the real world!

How do we know what's exploited or likely to be? Great question! Check out the Exploit Prediction Scoring System.

Facets of Assets at Risk

"See the ball; hit the ball." – Pete Rose

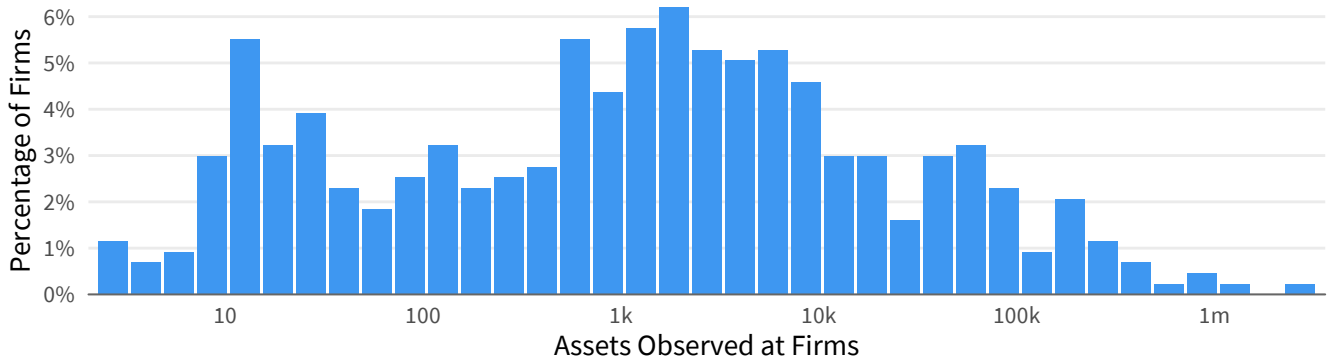
Pete's advice for hitting seems simple enough. And as the all-time hit leader in Major League Baseball, it seems to have worked pretty well for him. But 100+ years of history shows that even the best players fail to execute those simple instructions 70% of the time. Vulnerability management also seems simple on the outside: "Scan the vulnerabilities, fix the vulnerabilities." But anyone actually working in a VM program knows that description is a laughable oversimplification. The process of finding and fixing vulnerabilities varies widely based on the type of asset; desktops have a much different support structure than network appliances, Linux servers are managed differently than Windows servers, and so on.

Because of these stark differences, we'll begin our journey by exploring the many different types of assets observed across the environments of hundreds of organizations. To do this, we've analyzed vulnerability scan results imported into the Kenna platform for management. All told, we have data on over nine million active assets[1] across nearly 450 organizations.

Asset Prevalence

Figure 6 shows the number of assets in use at each firm in our sample. The typical (median) organization has about 800 active assets, but the full range extends from less than 10 to over 1 million. P2P Volume 1 demonstrated that a single CVE can result in open vulnerabilities across hundreds of thousands of assets. This goes back to that point we made earlier about the compounding scope creep of managing numerous vulnerabilities across numerous assets.

Figure 6: Distribution for the number of actively-managed assets per firm

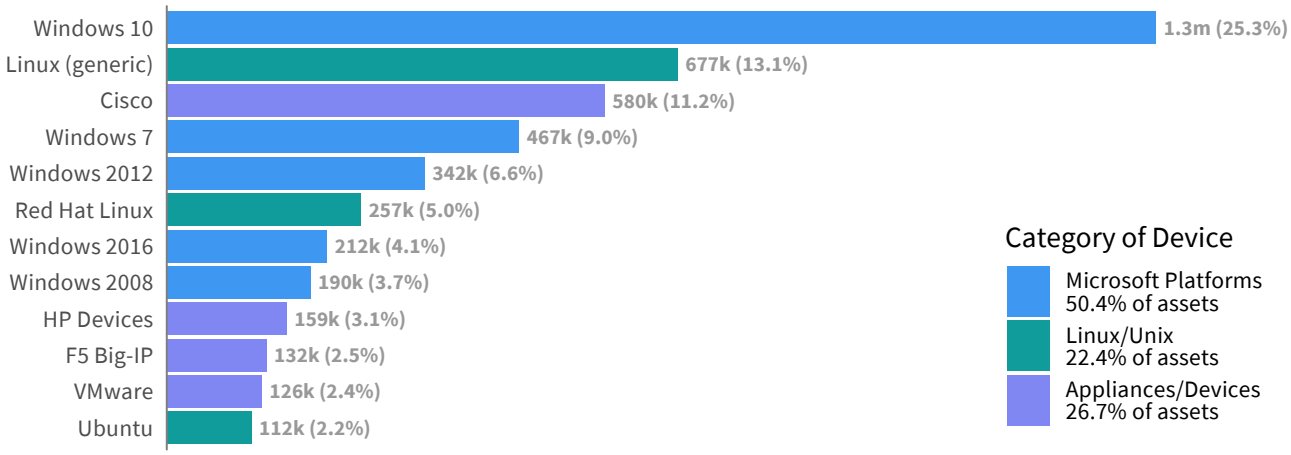


Now we're ready to circle back to the type of assets tallied in the previous figure. After some work parsing the asset detail strings provided by the latest output from vulnerability scanners, we tallied the most prevalent platforms in Figure 7. These 12 platforms account for 87% of all identified assets in our sample. You'll see these top 12 again in many of the charts to come.

[1] *Another 11.5M assets were in inactive status as of Jan. 2020 and were excluded from analysis. We also excluded 3.5M assets that scanners could not identify.*

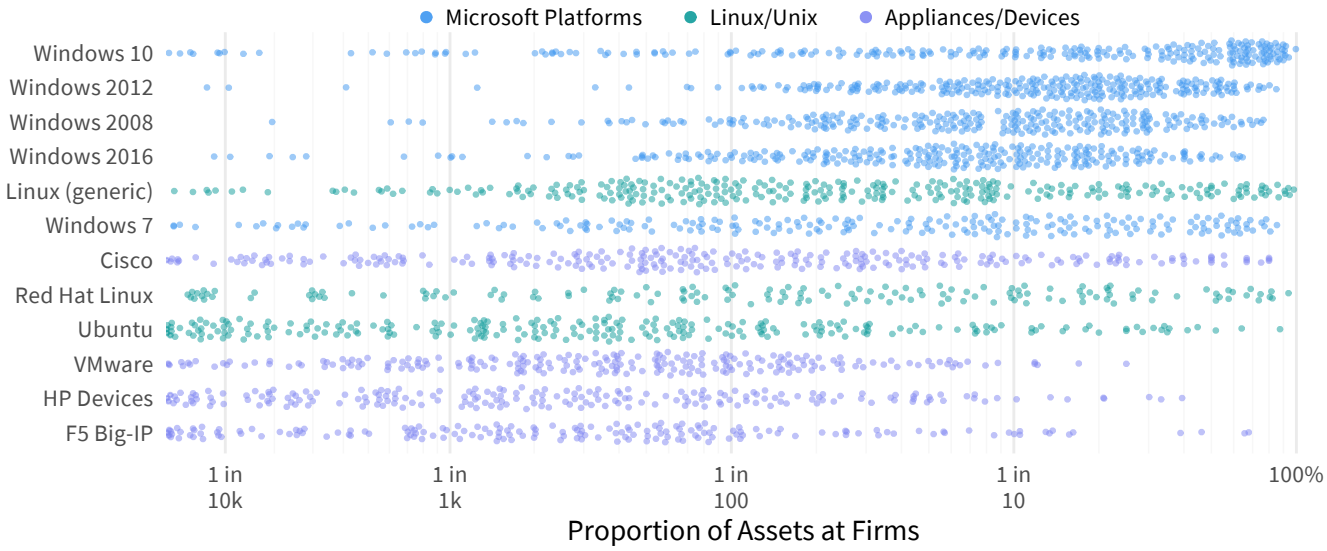
We grouped assets further to derive the major asset families shown in the bottom right of Figure 7, which encompass 99% of all active assets. Overall, about half of enterprise assets are servers and desktops running some version of Microsoft Windows. Just over a quarter fall into a broad spectrum of network devices and a little shy of another quarter consists of various flavors of Unix and Linux. And finally, there's a dash of Mac OS X systems (0.5%) and a smattering of mobile devices (<0.01%) that aren't shown.

Figure 7: Most common platforms among actively managed assets



The overall numbers shown in Figure 7 are interesting, but what we're really after is the proportionality of assets within individual organizations. Figure 8 grants that perspective. The 12 most prevalent platforms are again listed on the left. Each dot is an organization, and the position of that dot along the horizontal axis marks the proportion of the firm's assets represented by each platform. The clustering of the dots indicates where the bulk of organizations lie on the spectrum from "hardly any of our assets are X" to "nearly all of our assets are X."

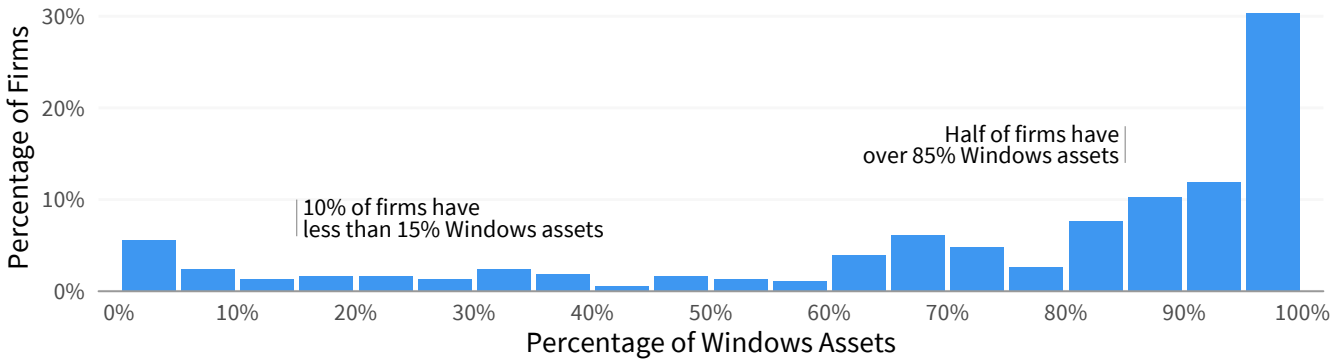
Figure 8: Proportion of all active assets in each firm represented by top platforms



From this, it's obvious that a strong majority of the actively managed assets in most firms run some version of Windows. But it's also clear that's not a universal rule; a good number of organizations also exhibit sizable proportions of non-Windows assets. And you'll notice a few curious standouts, like firms that appear to have a predominance of network appliances in their environment (or at least within the scope of their VM program).

The fact that Microsoft platforms dominate corporate environments probably won't come as a shock to anyone. Given the extent of that domination, let's do one more view that separates enterprise assets into just two categories: Windows and non-Windows. Figure 9 makes it a little easier to, for instance, determine that about half of all firms have an environment consisting of at least 85% Windows assets. We'll see what effect this has on the performance of the vulnerability management program as we progress through this report.

Figure 9: Proportion of assets within firms categorized as Microsoft Windows



We should qualify these results with a reminder that asset identification in this sample comes from vulnerability scanners. So, it's a direct reflection of the scope of VM programs. Firms may not be scanning production data centers, cloud environments, mobile devices, etc., which would obviously change the representation shown above. We suspect this amplifies the proportion of Windows systems in our results. But this fact in no way diminishes our findings, because assets in scope for the VM program are exactly those we want to study.

Vulnerability Density

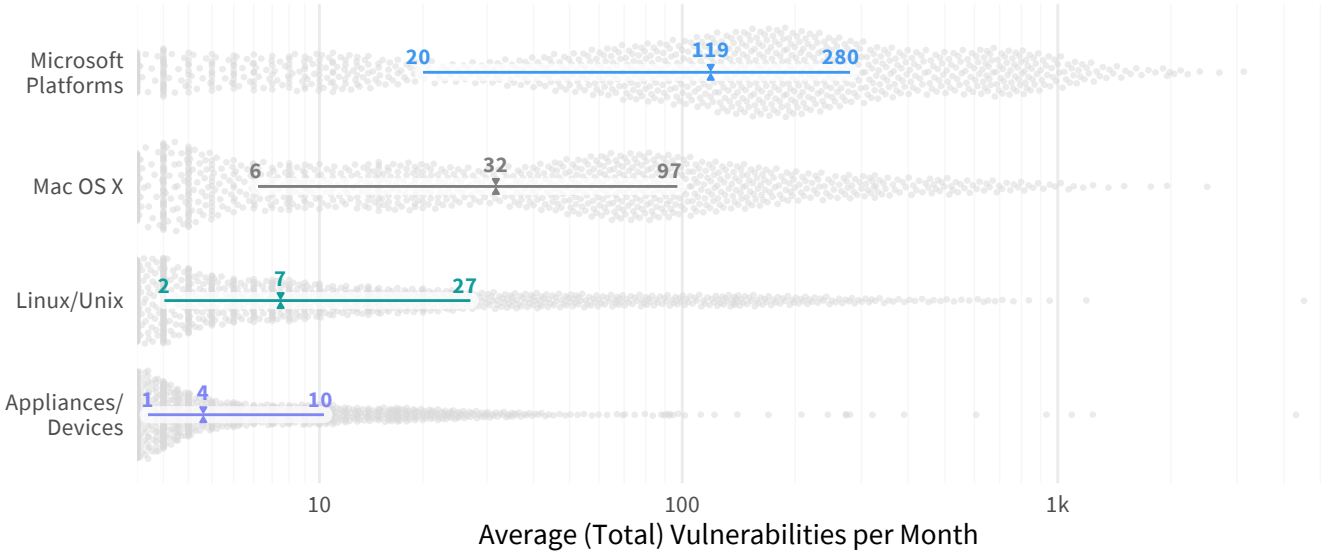
We now turn our attention to security exposures affecting enterprise assets. To do that, we'll need to introduce a new concept to our growing body of VM metrics analyzed in the P2P series – vulnerability density. In short, vulnerability density measures the number of vulnerabilities that exist in an asset. Lower density is generally better, but higher doesn't necessarily indicate poor security. Some assets will have more exposures simply based on the platform, installed software, function, configuration, and a host of other factors. But it is a good indicator for the volume of issues needing triage to securely manage an asset.

> Vulnerability Density: Measures the number of vulnerabilities that exist in an asset. This is inclusive of the base operating system, other software installed on the asset, and any misconfigurations identified by a vulnerability scanner.

Definition in hand, let's get to the more interesting part of actually measuring vulnerability density. We could calculate a generic measure across all assets in an organization, but the result would be highly dependent upon their relative mix of platforms. Instead, we calculated the median number of observed vulnerabilities each month for each asset over a two-year period. Figure 10 compares vulnerability density across the four major categories of assets introduced in the last section.

The dots in Figure 10 represent assets this time rather than organizations. The spread and shape of those dots form a distribution around the number of vulnerabilities observed per month per asset. Because vulnerability density varies so widely across assets (many with <10 but some >1000), we've added lines and labels to mark the 25th, 50th (median), and 75th percentiles for each category. As you review the numbers, keep in mind they include unpatched vulnerabilities carried over from the prior month, plus any that were newly discovered.

Figure 10: Distribution of vulnerability density across major asset categories



Using those guidelines, it's easy to see that a Windows-based asset typically has 119 vulnerabilities to manage in any given month. That's 30 times the median number of vulnerabilities in network appliances and almost 4 times the next closest asset category of Macs. That makes for compelling soundbites, but treating all vulnerabilities as equal isn't in line with the risk-centric philosophy we espouse around here. Let's rectify that.

We've repeatedly said in the P2P series that organizations should prioritize vulnerabilities that have known exploits or those likely to be exploited soon rather than pursuing the unrealistic goal of fixing everything indiscriminately. Applying this logic to vulnerability density, one might conclude that the presence of high-risk vulnerabilities in an asset is more critical than the overall density shown in Figure 10. And we agree with that astute observation. Bravo – you've earned your official P2P badge!

Before we run our vulnerability density metric through the risk filter, let's take an intermediate step. It's often said that just one opening is needed to successfully compromise a system and gain access into a target organization's environment. Following that logic, defenders might well wonder where attackers are most likely to find that foothold within enterprise infrastructure. Figure 11 supports that line of inquiry by determining the proportion of assets in each major category that have at least one exposed high-risk vulnerability.

We said this above, but it bears repeating: Statistics for asset categories and platforms shown in this report include not only vulnerabilities affecting the base operating system, but also other software installed on the asset. So, that median of 119 vulnerabilities on Microsoft platforms would include Windows plus Microsoft Office, Adobe Acrobat, Mozilla Firefox, etc.

The proportions in Figure 11 will obviously change for each individual organization, but this gives a reasonable baseline of what's typical across our sample. 70% of Windows systems and 40% of Linux/Unix systems offer an attacker that single exploitable vulnerability they need to succeed.[2] Network appliances and Macs exhibit somewhat lower exposure rates, but the basic impression from the data is that ample opportunity exists across the board. This is why we encourage firms to eliminate these easy footholds first and then move on from there.

Figure 11: Percentage of assets in each category with at least one open high-risk vulnerability

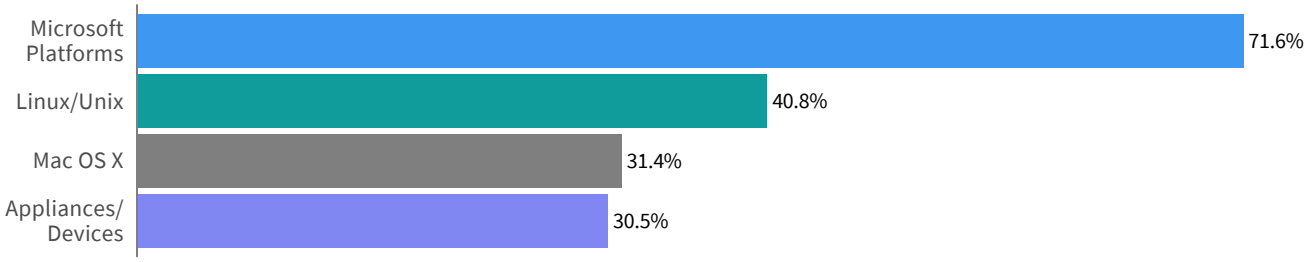
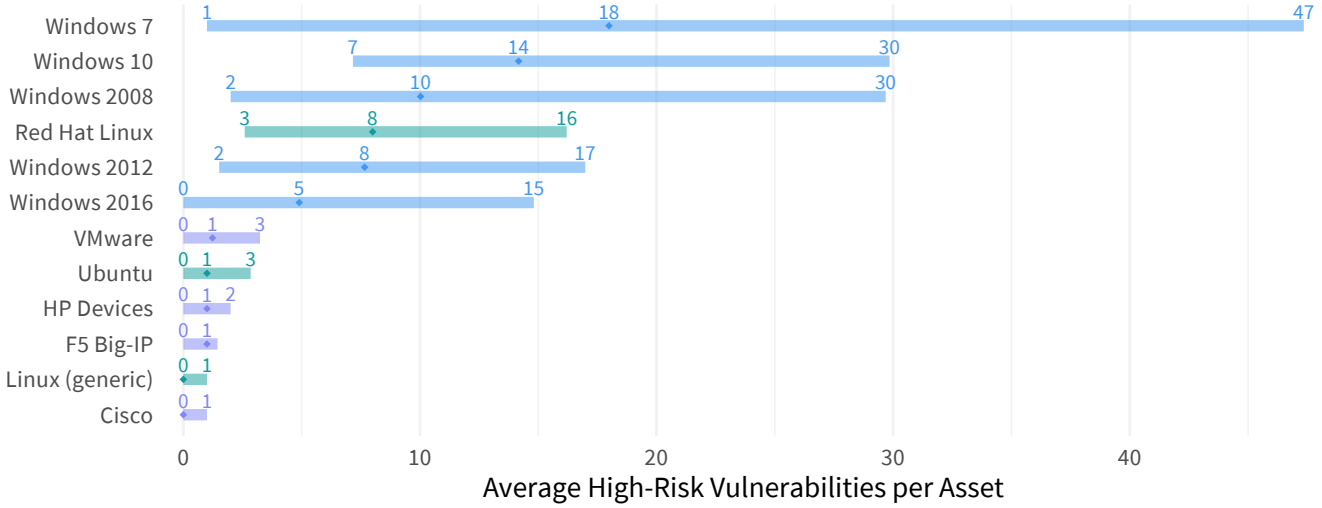


Figure 12 offers another way of visualizing high-risk vulnerability density. It's similar to Figure 10, except we've removed the dots (individual assets) and show only the bars representing the 25th to 75th percentile range to make things a bit cleaner. Information overload? It's okay to simply key in on the middle value corresponding to the median number of high-risk vulnerabilities per asset among the top platforms. Returning to the concept of attack surface, Figure 12 can be thought of as revealing the number of working keys attackers have to open doors into various platforms. Thankfully, the numbers are a lot smaller than those in Figure 10.

Figure 12: High-Risk vulnerability density across the top asset platforms



It appears from Figure 12 that desktops tend to have a higher density of risky vulnerabilities than servers –at least for Windows. The average Windows 10 PC has 14 weaponized bugs; the 75th percentile jumps to 47 high-risk vulnerabilities for Windows 7 (which ended mainstream support in 2015). The typical Windows 2016 server has 5, which is half the density of exploitable vulnerabilities compared to 10 on a Windows 2008 server (also no longer supported). If you're getting a "latest is greatest" vibe here, so are we. More on that later.

[2] *We're only considering assets that have open vulnerabilities with known exploits here. We do not expect all of these systems to be Internet-facing such that an attacker would be able to probe them. We're also not considering whether or not the exploit is just proof-of-concept or integrated into a tool like Metasploit. This is meant to be a simple view of potential exposures from an asset-centric perspective.*

Now let's talk about Red Hat systems, which seems almost out of place in Figure 12's otherwise Windows-dominated list. The median density of high-risk exposures for Red Hat systems is significantly higher than other *nix assets. To understand a key reason why, one must remember that this includes native platform vulnerabilities and all other vendors' software on the system. The median number of unique vendors detected on a Red Hat system is 41, which ranks highest among all platforms. Windows assets, by comparison, range in the teens of non-Windows vendors, and that drops to low single digits for most network appliances.

Overall, we found a moderate positive correlation between the number of non-native applications installed on assets and vulnerability density. More vendors, more vulns isn't terribly surprising, but it's good to know. And it's a good, evidence-based justification to minimize non-essential software on critical systems.

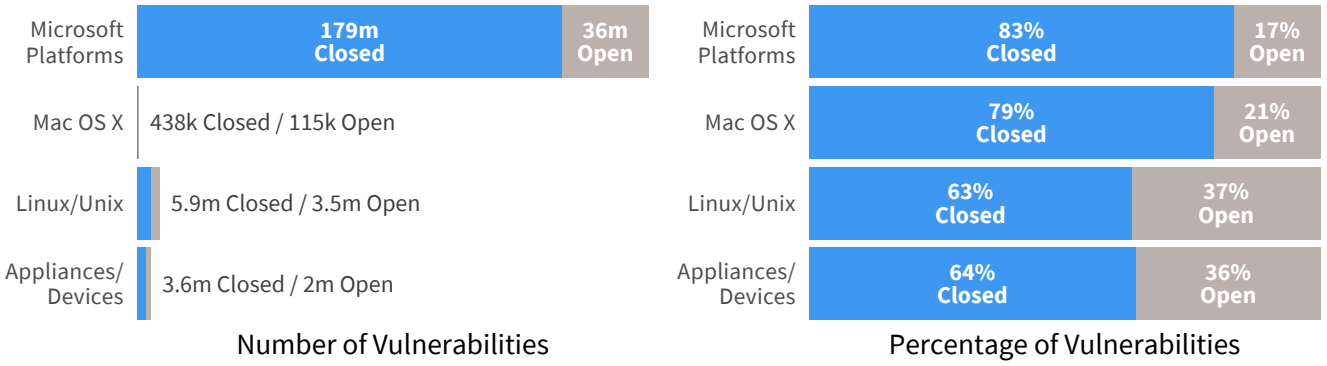
Remediation Coverage

We now turn to how organizations address exposures affecting their digital infrastructure. Remediation coverage is a good place to start because it picks up where we left off in the last section with high-risk vulnerabilities. It's also one of the first metrics we discussed in the P2P series way back in Volume 1. Coverage relates to the comprehensiveness of remediation efforts by measuring the percentage of vulnerabilities with known exploits that are closed or otherwise mitigated.

> **Remediation Coverage:** Measures the comprehensiveness of remediation. What percentage of exploited or "high-risk" vulnerabilities have been remediated?

Figure 13 presents two perspectives of coverage among the four asset categories. The chart on the left incorporates the raw count of high-risk vulnerabilities affecting each category. This emphasizes the fact that Microsoft platforms have many more exploited vulnerabilities observed across enterprise environments than any other category by a long shot. Taken alone, that makes it seem like Microsoft assets are in dire straits with respect to risk exposure. But isolated facts aren't what we do here; there's another side to the story (and figure).

Figure 13: Comparison of remediation coverage across major asset categories



The chart on the right offers a proportional viewpoint on coverage rates. This casts Microsoft platforms in a more favorable light with the highest percentage (83%) of exploited vulnerabilities that have been successfully remediated. That's even more impressive when you look at the raw number of issues that proportion signifies. OS X workstations aren't far off that mark with a coverage of 79%. Linux/Unix systems and network appliances lag behind, but still show coverage rates approaching two-thirds of high-risk vulnerabilities closed.

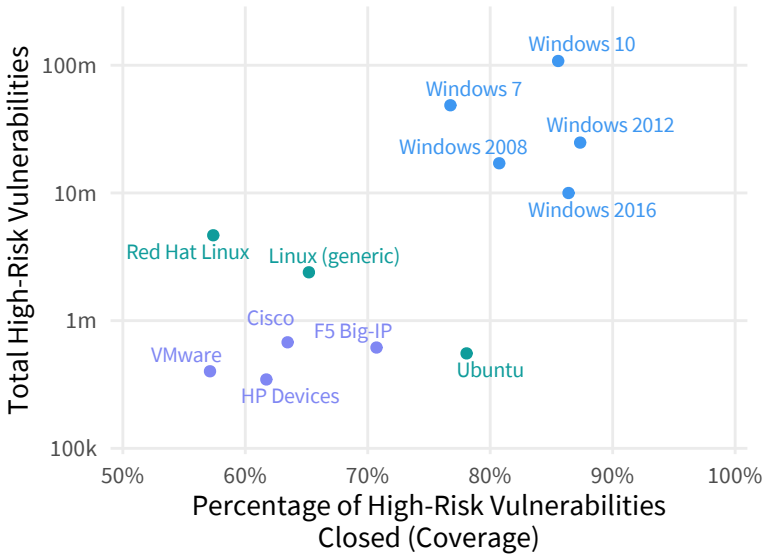


Figure 14: Number of high-risk vulnerabilities vs. remediation coverage across platforms

Figure 14 transforms the two different charts from Figure 13 into a single scatterplot and features the 12 platforms rather than the four asset categories. This helps contextualize the relative number of vulnerabilities that must be closed (y-axis) to achieve comparable rates of coverage (x-axis). The chart is counter-intuitive to a certain degree because it shows a positive correlation between the number of exploited vulnerabilities and proportion of closures. Perhaps necessity is the mother of remediation as well as invention?

Remediation Velocity

Platform prevalence and vulnerability density establish the foundation for the next set of questions we want to answer: Does the speed of remediation vary across assets? Are certain platforms better suited to enabling VM programs to stamp out vulnerabilities in a timely manner? To answer these questions, we'll reach for a tool we've used before in the P2P series – survival analysis.

> Remediation Velocity: Measures the speed and progress of remediation. How quickly are issues addressed and how long do they persist within and/or across assets?

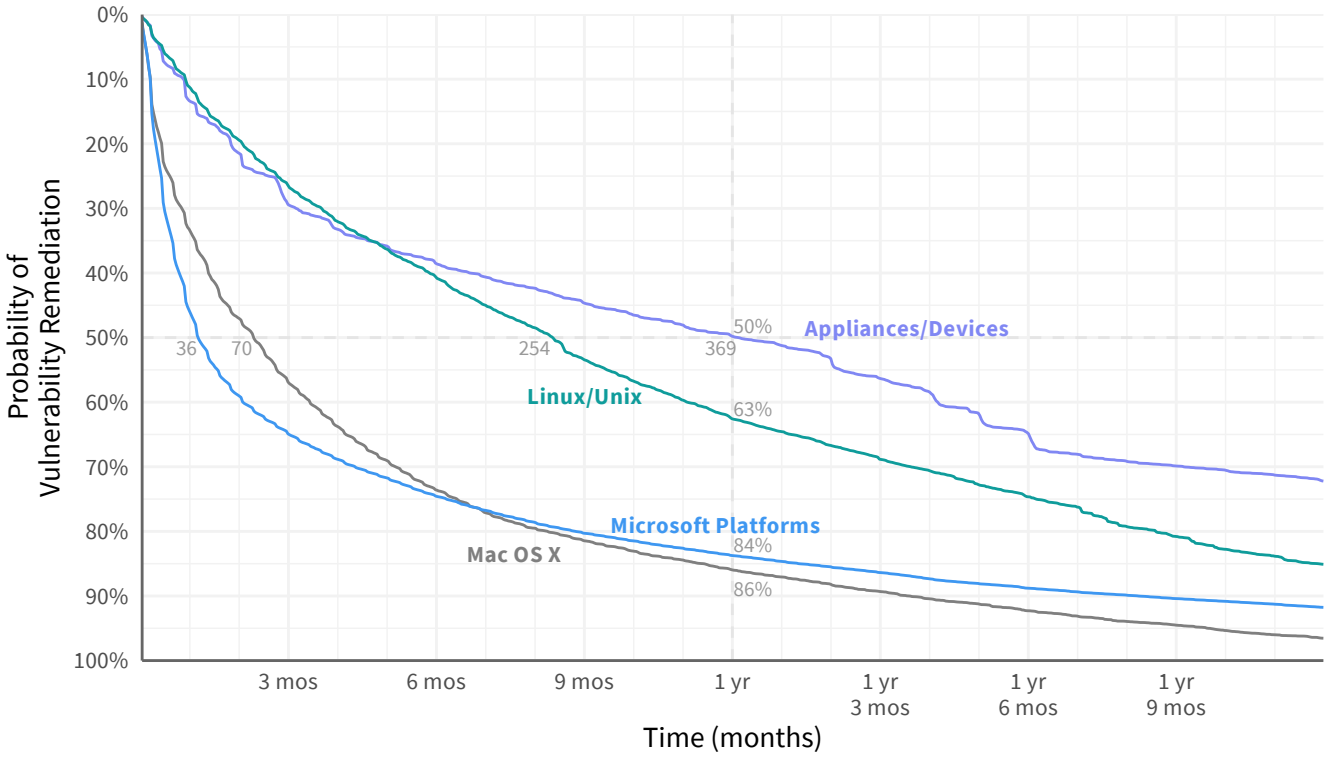
Survival analysis is a set of methods to study the time duration to an event (i.e., remediation of a vulnerability). If an organization finds 100 open vulnerabilities today and fixes 10 of them, the survival rate would be 90%. As time passes and vulnerabilities continue to be fixed, that rate will drop and produce curves like those shown in Figure 15 for the four major asset categories.

It's helpful to think of remediation activity as a race. Some are quick off the starting line, closing large numbers of vulnerabilities in a short timeframe. Windows assets look to run in that way. Some hit their stride a little later and make up for lost ground in the long run. Macs fit that description, passing Windows-based assets just before the seven-month mark and hanging on to the lead for the duration. Others (*nix and appliances) never seem to kick remediation into a higher gear and run the remediation race at their own steady pace.

Curious how to improve remediation velocity? P2P Volume 4 tests a dozen or so factors to determine which help and hurt.

We've added some values to Figure 15 to highlight these differences. The half-life of vulnerabilities in a Windows system is 36 days. For network appliances, that jumps tenfold to 369 days! After one year, only about 15% of bugs remain on Mac and Windows assets, but Linux/Unix systems and network appliances still have a lot of work ahead of them.

Figure 15: Comparison of remediation velocity across major asset categories



Let's see whether all Windows platforms do so well and whether all appliances lag so much. The next three charts split out remediation curves for specific platforms in each category (except Macs) from Figure 15. As you'll see, this is where things get really interesting.

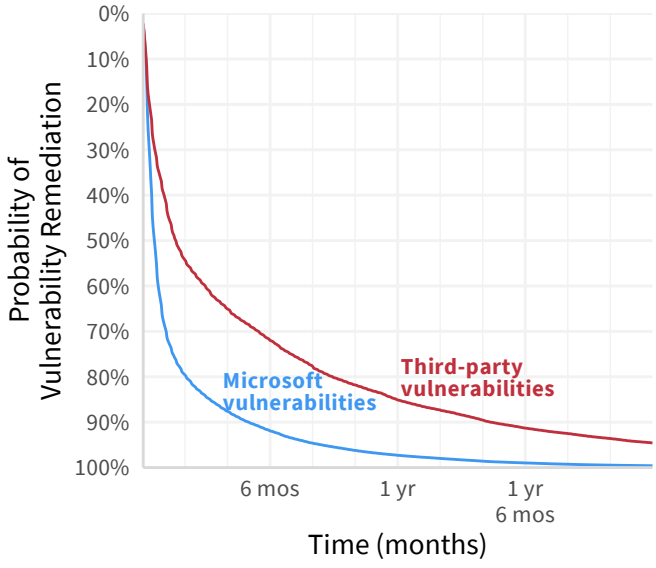


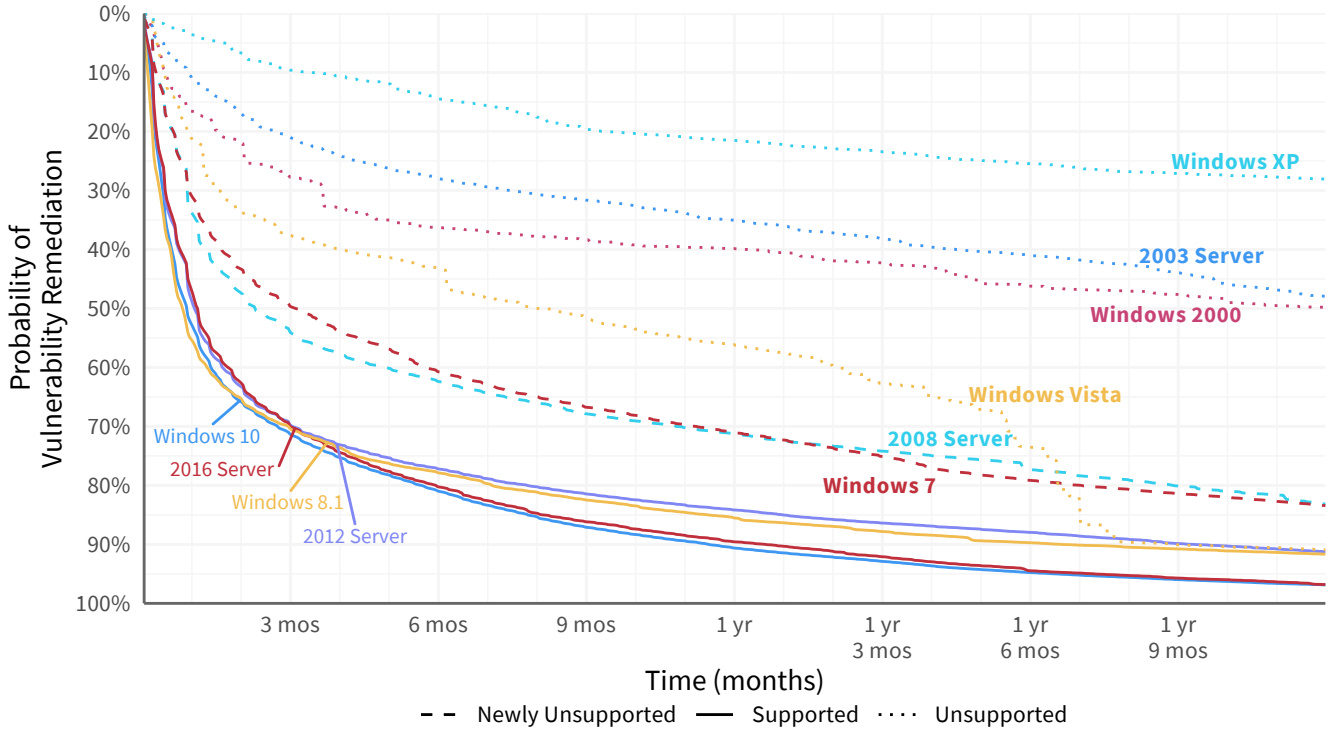
Figure 16: Remediation velocity for native vs. non-native vulnerabilities in Windows 10

Microsoft's leading vulnerability half-life is even more impressive when you remember assets have loads of additional non-Microsoft software installed. We thought it would be interesting to compare remediation velocity for native Microsoft flaws to those affecting non-native software on assets with a Windows OS. That's done in Figure 16 and it leaves us with even more respect for what Microsoft is doing to help VM programs succeed. 68% of Microsoft bugs are squashed in the first month compared to 30% for non-native software on those same assets! Notice to all vendors and platforms – the flag is planted.

Figure 17 compares vulnerability survival curves among Windows platforms. The difference in remediation velocity between the fastest version (Windows 10; 2016 Server) and the slowest (Windows XP) is pretty incredible. Let's start at the top, where (indicated with a dotted lines) you'll find older versions of Windows that are no longer supported. Clearly, that milestone in a platform's life makes it much more difficult to manage vulnerabilities in a timely manner (or even at all). If you're looking for data to show upper management that squeezing a little more life out of those legacy Windows servers and desktops is a bad idea, this should do the trick.

2008 Server and Windows 7 run in the middle of the pack in Figure 17. Mainstream support ended for both of these years ago, but extended support was available until January 2020. Since we pulled data for this report at the beginning of 2020, extended support was available for the duration of the timeline shown in the chart.

Figure 17: Comparison of remediation velocity across Windows platforms

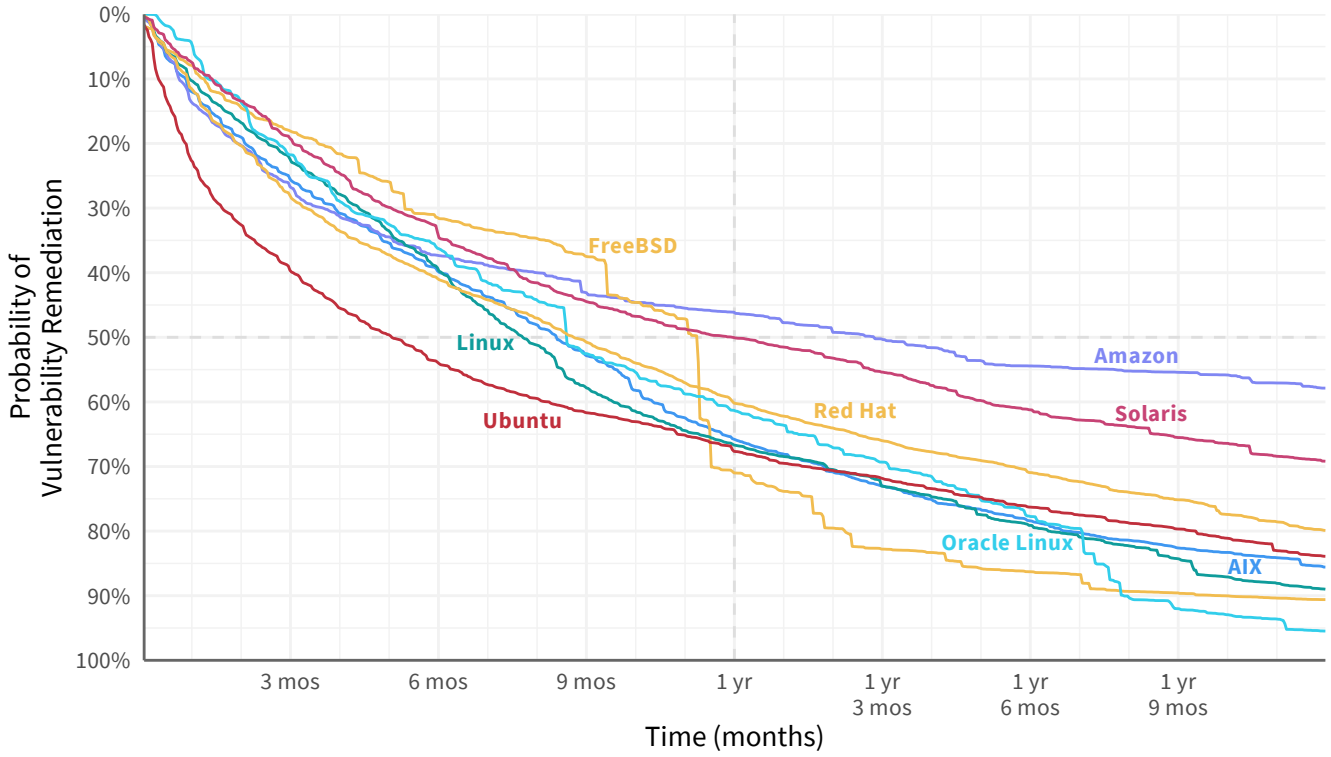


The latest desktop and server versions all cluster tightly together at the bottom in Figure 17. We interpret this as strong evidence that Microsoft (and the ecosystem of third-party tools for Microsoft products) is making it increasingly more efficient for organizations to manage vulnerabilities affecting supported versions of their operating systems. The recipe of a regular cadence for patch releases, automatic/forced updates (love them or loathe them), efficient tools for deploying patches, etc., appears to be yielding good fruit for fast fixes.

We have one more observation before moving on from Figure 17. There's less difference between Windows desktops and servers than anticipated. This may tie back to the reliable and streamlined cadence of patch release and remediation easing classic concerns of patch deployments blue-screening critical servers over time.

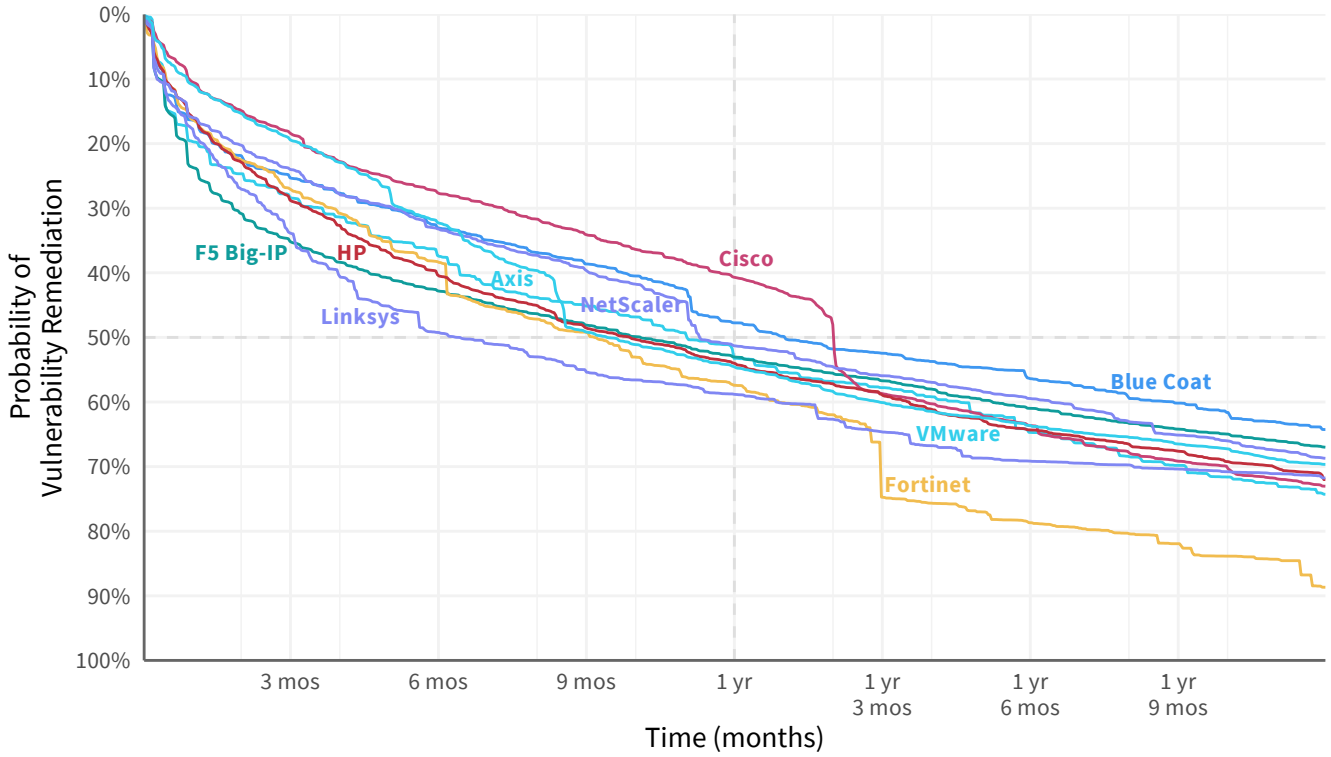
Moving on to Linux and Unix platforms, Figure 18 presents survival curves for that segment of enterprise infrastructure. Two observations stand out. First, remediation velocity is slower across the board, which isn't unexpected given what we saw in Figure 15. That concave shape of early momentum seen for actively supported Windows assets is completely absent here. Second, there's much less variation among platforms. In all fairness, however, we attribute much of that to the fact that we're not splitting out different historical versions of *nix flavors like we did for Windows.

Figure 18: Comparison of remediation velocity across Unix/Linux platforms



A cynical explanation of the contrast between Figures 16 and 18 is that Windows admins are more accustomed to regular reboots after decades of conditioning. And there might be some truth to that hypothesis. A less cynical explanation is that the tooling for managing Linux fleets at scale generally lags behind that of Windows. Whatever the reasons, patching of Linux and Unix systems seems plateaued somewhere between Windows 2000 and 2008 levels. Perhaps they'll figure out how to package whatever FreeBSD admins are doing between months 9 and 12. For now, let's proceed to network appliances in Figure 19.

Figure 19: Comparison of remediation velocity across network appliances



One major thing to keep in mind regarding Figure 19 is that vulnerability density is much lower for this asset class (see Figure 10). That means it takes more time to close fewer issues. The main takeaway is that appliances generally have a remediation half-life of about a year, and even after two years, few achieve a 70% closure rate.

The silver lining here is that exploits targeting network devices are far less common in the grand scheme of potential targets across the enterprise. That's not at all grounds for complacency. But if you're looking to focus remediation efforts on maximum risk reduction, starting with other asset categories likely gets you closer to that goal.

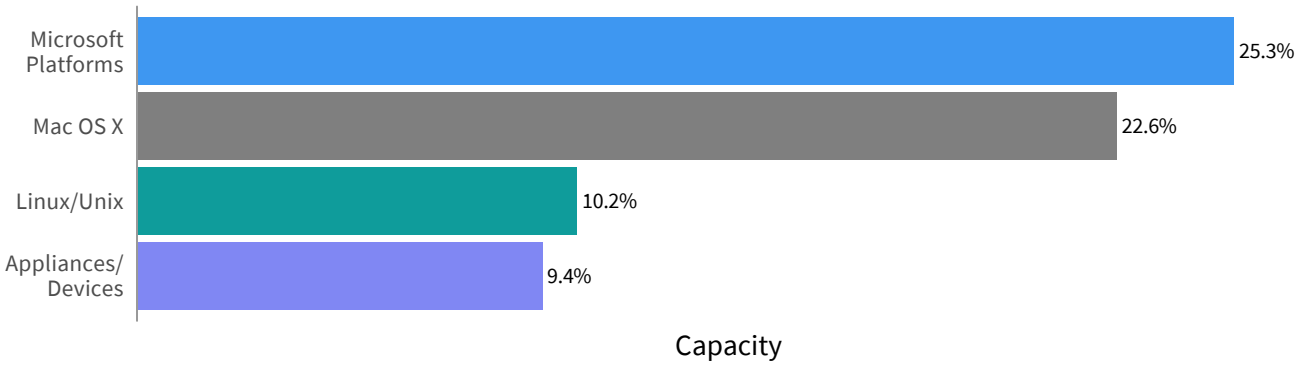
Remediation Capacity

With a study of vulnerability density, coverage, and remediation velocity under our belts, we're now ready for what might be the defining question facing VM programs –are we focused and fast enough? Pursuing the answer to that question takes us into the realm of remediation capacity. It essentially measures whether the rate of closing known vulnerabilities keeps pace with the discovery of new ones.

> **Remediation Capacity:** Measures the ratio of new and closed vulnerabilities in a given timeframe. Below capacity indicates the VM program cannot keep up with newly discovered vulnerabilities over time. Above capacity means the program is generally able to close enough vulnerabilities to offset the new ones.

We mentioned in the recap section that organizations remediate about 10% of the vulnerabilities in their environment during any given month. Is that a universal constant or do some types of assets exhibit higher (or lower) than average remediation capacity? Figure 20 holds the answer.

Figure 20: Median proportion of vulnerabilities closed per month

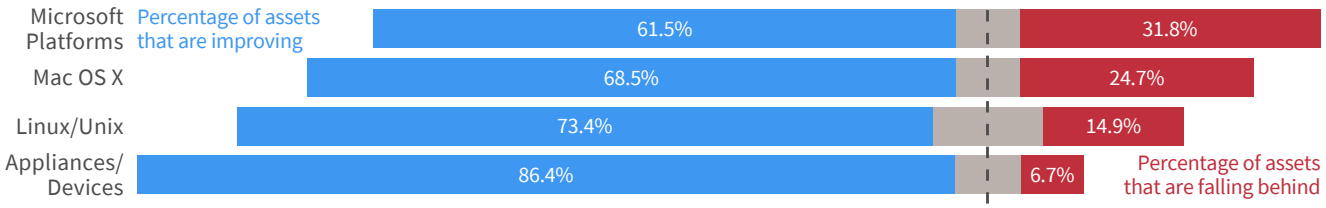


We once again see Microsoft Windows systems achieving impressive levels of remediation performance. Kudos to Apple as well. Assets in those categories more than double the overall average for raw fix capacity. Network appliances and *nix systems, on the other hand, can't break through the 1-in-10 glass ceiling for raw fix capacity.

Results for raw fix capacity shouldn't be too surprising, since it relates strongly to velocity. If vulnerabilities can be remediated quickly, as we saw for Windows and OS X, it stands to reason that a larger proportion of them could be successfully closed. But there's another key aspect that plays into remediation capacity beyond sheer speed – the total volume of vulnerabilities affecting the asset.

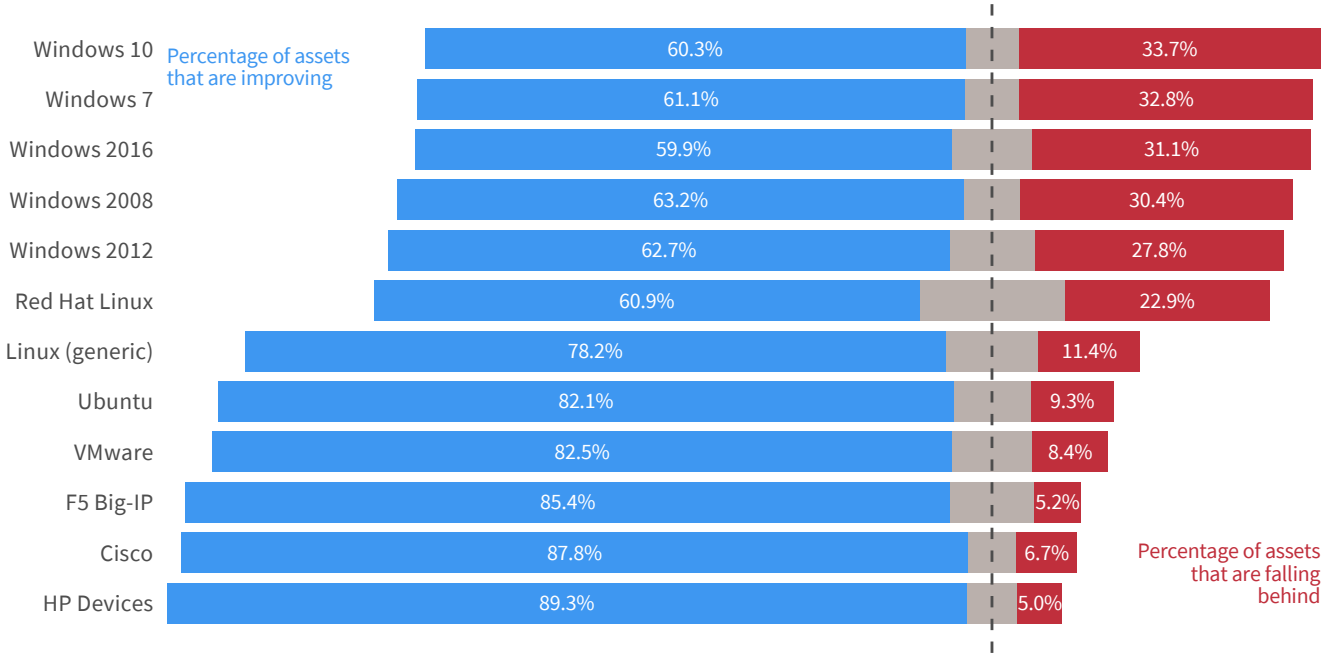
Figure 21 harkens back to Figure 5 (and Bomb Pops), except assets are the unit of measurement rather than organizations. We want to know for which asset categories the VM program is typically falling behind (new high-risk vulnerabilities exceed closed), maintaining (new = closed), or gaining ground (closed exceeds new) over a two-year period.

Figure 21: Comparison of high-risk remediation capacity across major asset categories



Before we proceed to the details, let's first appreciate the fact that far more assets show a positive capacity than negative capacity across the board. We think congratulations are in order on that accomplishment. Beyond that, the results in Figure 21 pretty much invert what we've seen for prior metrics. Organizations struggle most to keep Windows and Mac-based assets ahead of new vulnerabilities and demonstrate the highest net capacity for Linux/Unix systems and network appliances. We see the same basic pattern in Figure 22 when splitting out these categories into particular platforms.

Figure 22: Comparison of high-risk remediation capacity across the top 12 asset platforms



Though inverted, this outcome is intuitive if we recall what we learned about asset prevalence and vulnerability density. Yes, Windows systems close out issues impressively fast, but the predominance of Windows assets in the enterprise combined with the sheer volume of vulnerabilities affecting those platforms makes it infinitely more difficult to keep up over the long haul. We're honestly impressed that so many manage to do it successfully. All in all, it reinforces the absolute necessity of using every drop of remediation capacity on vulnerabilities that matter most. Otherwise, there's truly no hope of staying on target to reduce risk to the organization.

Reviewing What We Learned

We've covered a lot of metrics and statistics during this analysis, so let's bring things to a close by reviewing and reflecting on what we've learned. To that end, we thought one last "Grand Finale" visualization would help us better see the forest for the trees. Figure 23 compares the four major asset categories across key metrics presented in each section of this report.

Figure 23: Comparison of vulnerability and remediation metrics across major asset categories

	Vulnerability Density	Assets at Risk	High-Risk Density	High-Risk Coverage	Vulnerability Half-Life	Fix Rate	Net Capacity
Microsoft Platforms	119.1	71.6%	11.8	83.2%	35	25.3%	68.2%
Mac OS X	31.7	31.4%	3	79.2%	69	22.6%	75.3%
Linux/Unix	7.5	40.8%	1	63.2%	253	10.2%	85.1%
Appliances/ Devices	3.6	30.5%	0	64.0%	368	9.4%	93.3%

- **Vulnerability Density**: Median number of vulnerabilities detected on assets (Figure 10)
- **Assets at Risk**: Percentage of assets with at least one open high-risk vulnerability (Figure 11)
- **High-Risk Density**: Median number of high-risk vulnerabilities detected on assets (Figure 12)
- **High-Risk Coverage**: Percentage of all high-risk vulnerabilities remediated (Figure 13)
- **Vulnerability Half-Life**: Days required to remediate half the vulnerabilities in assets (Figure 15)
- **Fix Rate**: Median percentage of total vulnerabilities closed in a given month (Figure 20)
- **Net Capacity**: Percentage of assets where high-risk vulnerability closures meet/exceed new vulns (Figure 21)

It may seem rather obvious, but our overall takeaway from this volume in the P2P series is that assets matter for vulnerability management. They matter so much that the vendors and platforms comprising the infrastructure under management **could** have more **innate** influence on performance than the VM program itself. We bold "could" and "innate" in that last sentence to underscore that we're not implying that you have no control over your VM destiny. Rather, we're merely making the point that the materials used to build a thing have a lot to do with whether that thing meets the purpose of its design and how long it lasts.

If your goal is to minimize opportunities for attackers to exploit systems on the network, an environment filled with Windows servers and desktops will make that more challenging from the start. But if you're aiming to maximize the ability to remediate vulnerabilities quickly and comprehensively, then a heavy Windows-based architecture gives the best chance of hitting that mark. What's best for your organization comes down to aligning risk tolerance, strategy, and capabilities.

We understand that most organizations don't get to start from scratch and choose assets that support the mission of the VM program. Assets and VM programs support the mission of the business. But understanding the inherent strengths and weaknesses of different platforms in use will greatly help set strategies that leverage those strengths and compensate for weaknesses so the VM program is better able to succeed in its vital mission. We hope Figure 23 and the rest of this report help VM teams in considering and developing those strategies.

One more word before we conclude. We've spoken mostly to organizations, but these findings emphasize the critical role played by vendors in the VM ecosystem. We hope this serves as both a pat on the back where it's deserved and a kick in the rear where it's needed. Vendors would do their customers a great service by reflecting on what can be done to improve their products' capabilities with respect to the various metrics presented in this report. We're all in this together. Let's win this together.

PRIORITIZATION TO PREDICTION

VOLUME 5: IN SEARCH OF ASSETS AT RISK

KENNA Security

119 Cyentia INSTITUTE

> "The vendors and platforms that comprise the infrastructure under management **could** have more **innate** influence on performance than the VM program itself."

© 2020 Cyentia Institute, LLC. All Rights Reserved. The Cyentia Institute and Kenna Security names and logos are the property of their respective owners.