State of Enterprise Vulnerability Detection and Patch Management

TuxCare

Making Open Source Linux Enterprise Grade

For maximum security and compliance, enterprises need to rapidly patch vulnerabilities, keep production Linux systems updated with the latest fixes, and have a trusted technology partner for Linux support & maintenance always within reach. TuxCare ticks all the boxes by helping organizations to take care of support, maintenance, and security for Enterprise Linux systems.

Our Services

Support, Maintain and Secure All Critical Components of Enterprise Linux Systems

LIVE PATCHING SERVICES

Put an end to service interruptions & non-compliance caused by system reboots

LEARN MORE

END-OF-LIFE LINUX SUPPORT SERVICES

Eliminate security vulnerabilities while running End-of-Life Linux

LEARN MORE

LINUX SUPPORT SERVICES

Keep all components of the production Linux systems always up-to-date with vendor-level support services

LEARN MORE

ALL TUXCARE SERVICES INCLUDE INTEGRATIONS WITH PATCH MANAGEMENT TOOLS, VULNERABILITY SCANNERS, EPORTAL SECURE PATCH SERVER AND 24/7 SUPPORT.

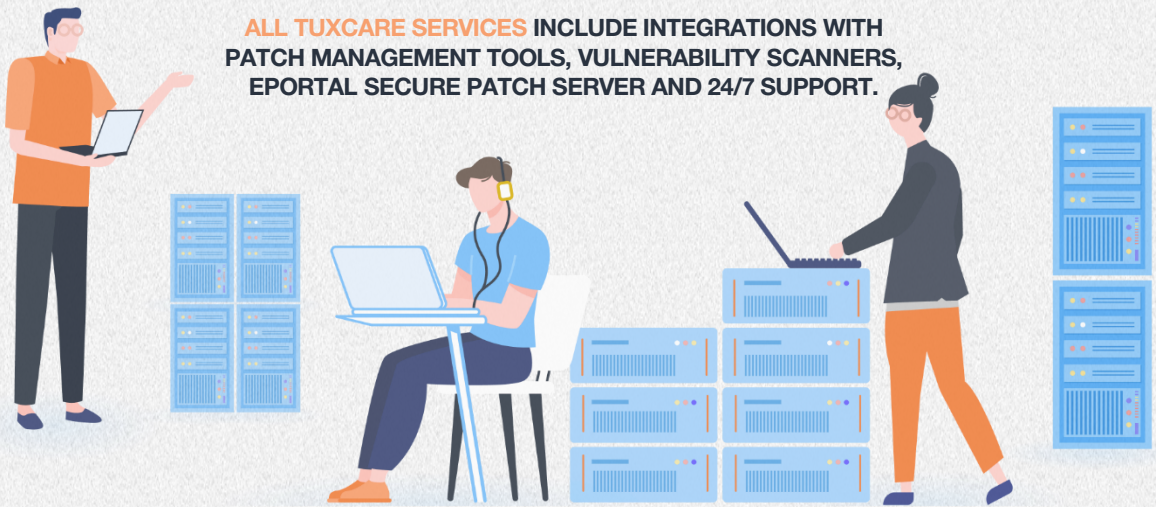

Contents

Executive summary

The cybersecurity threat is broad, pervasive, and complex. Mitigating such a comprehensive threat requires the use of a multi-pronged approach. Technology teams must maximize the utility of every tool in the security arsenal if they stand a chance to mount an effective defense. There are, however, different approaches to the use of specific tools.

This survey looks at two critical cybersecurity tools. First, we examine the vulnerability detection tools used to locate and map software vulnerabilities in an organization. We also look at patch management, the process of closing vulnerabilities through software patches.

Our intent is to analyze how different organizations treat these two key tenets of cybersecurity. We will outline the key variations we found from industry to industry – as well as aggregate statistics that point to gaps in preparedness such resource limits and the reliance on manual methods.

Introduction

Vulnerability detection and patch management are two important operational steps that underpin secure technology infrastructure in every industry. With this study, we present an industry-by-industry analysis to discover how different industries structure vulnerability detection and patch management operations.

We focus on three key factors: maintenance windows, patch deployment practices, and the overall security awareness of the technology professionals that implement these processes. Results were collected in the first quarter of 2021. Respondents completed an online survey that was publicly advertised to IT professionals around the globe.

Interestingly, the geographic location of a respondent had no bearing on the survey results as varying locations reported similar patterns across our questions, including our questions around the time spent and the frequency of patch deployment and maintenance operations. However, the industry in which an organization operates did have a clear impact on the question results.

At the time of writing, we received responses from 106 subjects, however, the survey is still running, and we will continue to update the results to reflect any changes in the findings. As it stands, we are confident that the results provide helpful insights to the teams faced with these challenges.

IT professionals that work in a systems administration role represent the largest number of respondents, while in aggregate, 88.7% of our respondents were directly involved in vulnerability management operations.

Our analysis revealed the following key points:

The majority of companies (76%) are deploying automated patching procedures.

75% of respondents said that they relied on manual online research as one of their tools to find out more about dangerous vulnerabilities, making this the most commonly used tool.

Most respondents said that CentOS itself, or another CentOS fork, is their predominant server OS.

73% of respondents said that their server fleets use just one OS, with just 27% suggesting that they use a mix of operating systems in their server fleets.

Across industries, documenting the patching process is not consuming a significant amount of time when compared to other patching-related tasks – in fact, documenting the patching process consumes the least amount of time.

In some industries, obtaining approval for a maintenance window can be the most time- consuming element of the patching process – in some cases consuming more time than applying, documenting, or monitoring patching.

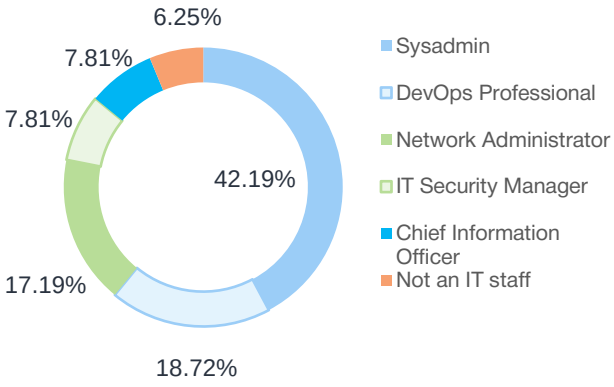

RESPONDENT ROLE

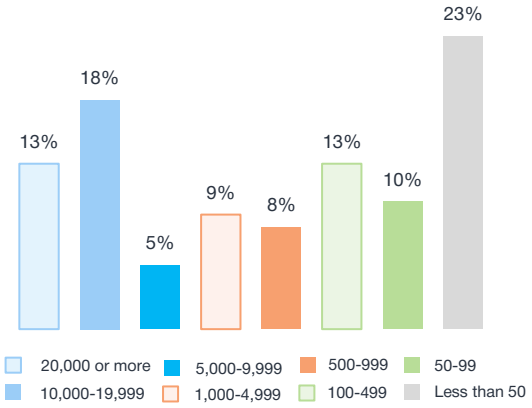

COMPANY SIZE

Respondents worked at companies of a variety of sizes. Small, medium, and large companies are all represented in the survey results.

Linux ecosystem diversity

Most of our respondents reported that they used a single Linux distribution for their server fleets, though a significant minority of respondents used multiple Linux distributions in their server deployments.

There are pros and cons to each approach. Using a single Linux distribution for all server roles carries benefits for server management by reducing the efforts required in managing servers and by simplifying the application of automation tools. On the other hand, choosing to use a role-specific Linux distribution for different server roles such as a web server, file server or authentication provider allows organizations to fully exploit the strengths of each distribution.

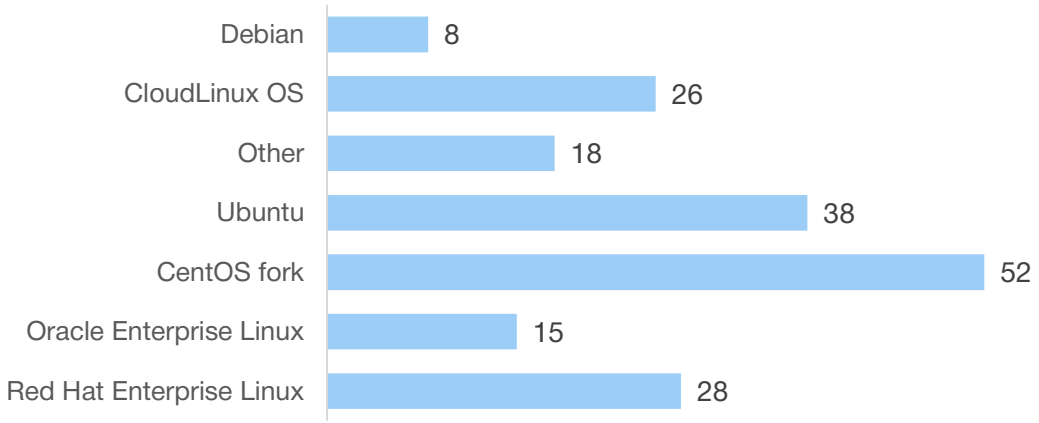

LINUX DISTRIBUTIONS

The “other” option resulted in various entries including SUSE, Proxmox, Raspbian or Arch Linux.

We found it revealing that single distribution fleets dominate the landscape. That suggests that the standardization of procedures is an important advantage for organizations. In our results, we found that the preference for single distribution fleets is consistent across industries, and company sizes.

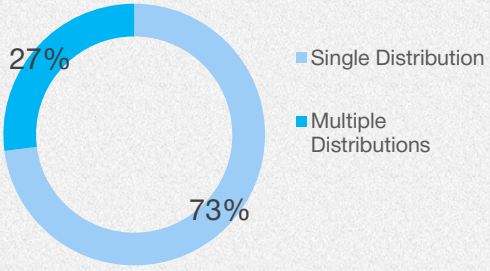

COMPANIES WITH MULTIPLE OR SINGLE DISTRIBUTIONS FLEETS

Ideal vulnerability scanner and patch management tool

Vulnerability management tools are relatively complex and typically carry a steep learning curve while features and capabilities vary from product to product. We asked respondents what features they would like to see in their ideal vulnerability scanner and patch management tool.

Responses varied, with the respondents selecting nearly equally from the available options, while a significant number of respondents opted to suggest a feature under “other”. These preferences included “logging”, “minimal impact on system resources”, “phased rollouts” and “detection of backported fixes”.

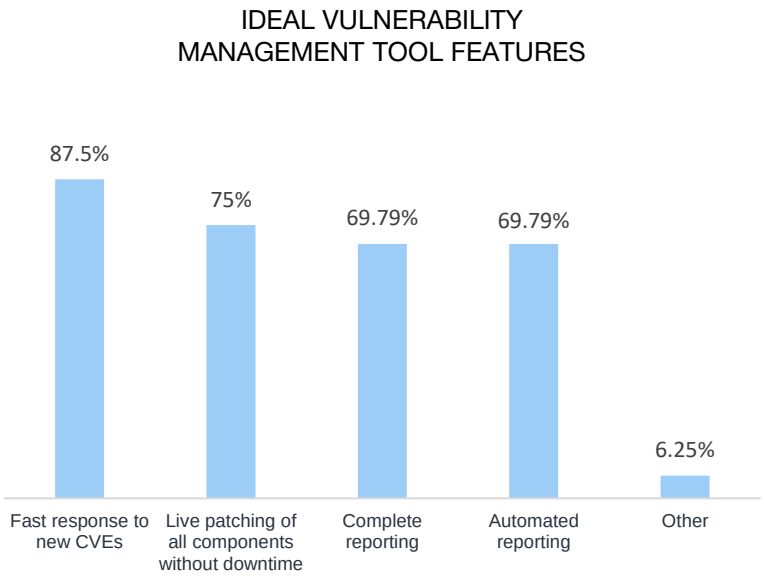

Logging as a desired feature

It is not difficult to see why logging is mentioned. There is a need for transparency with respect to the inner workings of a tool. After all, security tools are directly touching organizational systems. What happens under the hood matters.

Logging is also important because the user interface of a vulnerability management tool can hide the underlying complexity of the tool. When a vulnerability management tool runs it will typically run several complex, involved scripts that log into an OS or into applications to detect existing packages, to check versions, and to test against known exploits.

This automated script generates data that gives the user a bird’s eye view of key information via reports or a security operations center (SOC) dashboard. Yes, automation is helpful for day-to-day security operations activities, but automation can obstruct debugging steps.

Current logging implementations sometimes provide so much information that it becomes overwhelming to process and therefore obstructive. For other tools, the logging data provides too little information to be of value.

Resource use and patch rollout

Respondents that pointed to “minimal impact on system resources” reflected the fact that today’s tools can have a noticeable impact on system responsiveness or throughput. That is because of the characteristics of today’s tools: most work either through an agent deployed on every server or via remote calls into existing daemons. The frequency of tests and the volume of information gathered can easily cause a drain on resources.

Another feature requested by respondents is “phased rollouts”. In other words, respondents wanted more granular control over patching so that servers could be grouped, and so that patches could be deployed to selected servers in order to test patches for wider distribution.

It adds an essential extra stage in patch management – stepping patch deployment from development to quality assurance, and then on to production. It is possible to set up selective deployment with current tools, but it is a time-consuming and manual exercise.

Finally, the respondent that specified “detection of backported fixes” referred to the ability of a security scanner to detect vulnerability fixes that have already been applied, and to do so in a way that does not rely purely on checking the version string.

Detection of backported fixes matters because some patching mechanism will apply patches to the affected code, without updating the version number. That can lead to false positives if a patching tool only relies on version numbers.

Vulnerability Management

Awareness is critical when dealing with vulnerabilities and it doesn't matter how an organization gains awareness: through online information, through vulnerability scanning, or thanks to internal teams. Understanding which systems are vulnerable is always the first step towards protecting those systems because it is only possible to protect against danger when you're aware of that danger.

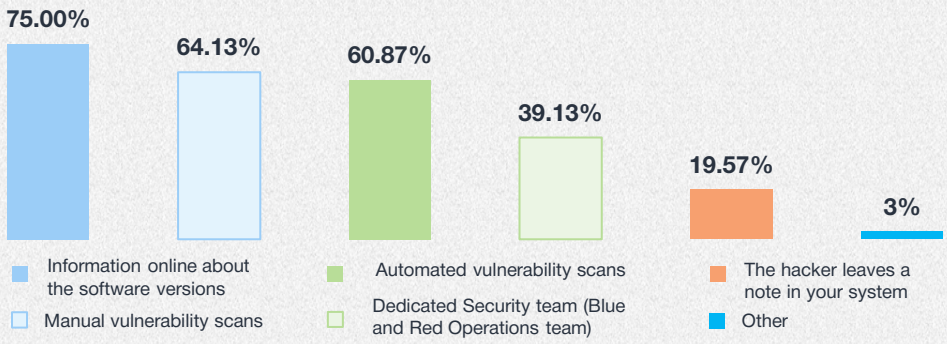

VULNERABILITY AWARENESS

For our question on vulnerability awareness respondents were allowed to select multiple options. We found it particularly interesting that nearly 20% of respondents found out about a vulnerability due to a note from a hacker on their system – which is a clear and obvious indicator of a successful breach.

It is notable that online research is the most common source for vulnerability information as reported by our respondents. Given the growth in CVEs, the manual nature of online research may not remain an effective way to manage vulnerabilities in the long run.

Vulnerability scans performed on an organization's systems still remain an important way to gain awareness about security flaws – whether these are performed manually or automatically. Using a dedicated security team to assess vulnerabilities is less widespread and possibly points to a lack of resources rather than a genuine lack of interest in appointing a dedicated team.

Responding to a vulnerability

If a vulnerability is detected it must be dealt with and, if possible, resolved. This is how respondents said that they handled vulnerabilities.

Emergency maintenance windows is arguably the most disruptive mechanism, but it was nonetheless the preferred choice when dealing with a known vulnerability. It can therefore be suggested that, for over 70% of respondents, the risks associated with a potential security breach outweighs operational and availability considerations.

The only non-disruptive mitigation method used, live patching, is chosen by nearly half of respondents. Interestingly, many respondents replied that they cope with vulnerabilities simply by waiting for the next periodic maintenance window. This, in turn, implies that their systems will remain vulnerable during the waiting period.

HOW DO YOU COPE WITH DETECTED VULNERABILITIES?

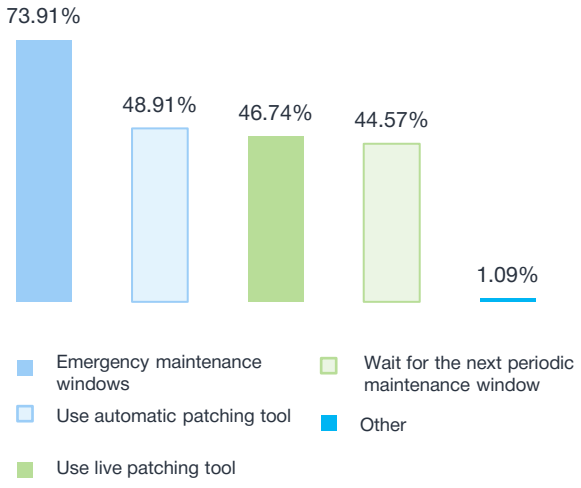

Automation

The response to the discovery of a vulnerability can be actioned either manually or automatically. There are several steps involved in a response and it is interesting to look at the steps that are most commonly automated – and the steps that are least likely to be automated.

Automating the patching process is clearly widespread while reporting is also commonly automated. However, assigning tasks and prioritizing patches are steps that are much less likely to be automated.

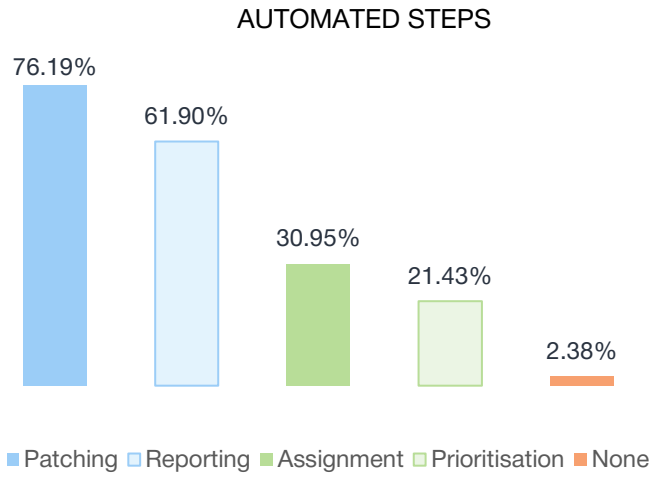

Surprisingly, despite the high numbers of CVEs that are filed, and the overall growth in cybersecurity threats, a number of respondents reported that their organizations do not automate any of the steps involved in the vulnerability management process.

It is common for organizations to combine automated tasks – for example, combining patching and reporting. We asked respondents to report whether they automated multiple processes:

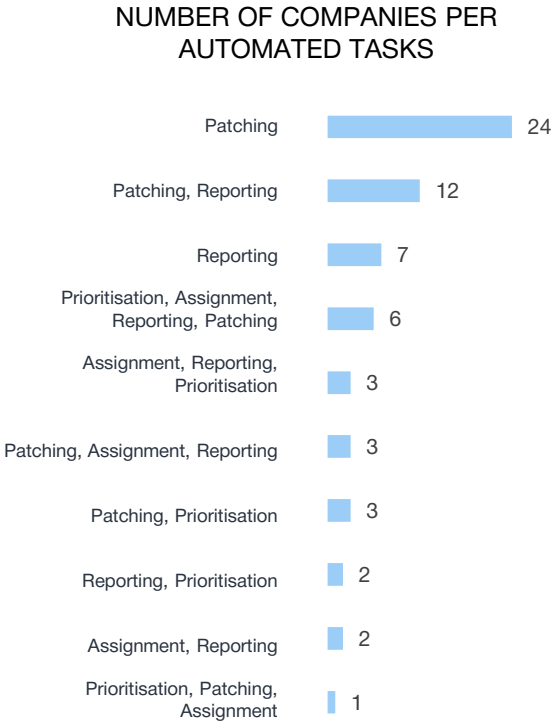

Most companies are only automating the patching process and reporting is the most commonly automated function – with many companies automating both patching and reporting. Compliance law is increasingly demanding evidence of patching, so it stands to reason that we will see more and more reporting automation.

In our survey results we also found that smaller IT teams – less than twenty team members – reported a higher reliance on reporting automation. In fact, small team's reliance on reporting automation outweighed the reliance on automation of larger teams by a factor of 6 to 1.
Reporting is the only automated task where there was a significant difference in the behavior of small teams vs. large teams.

Maintenance and Patch deployment

Detecting vulnerabilities in systems requires teams to perform monitoring and testing tasks that probe their systems. The amount of time spent doing so varies dramatically from one industry to another.

HOURS PERFORMING MONITORING TASKS AND LOOKING FOR VULNERABILITIES PER WEEK

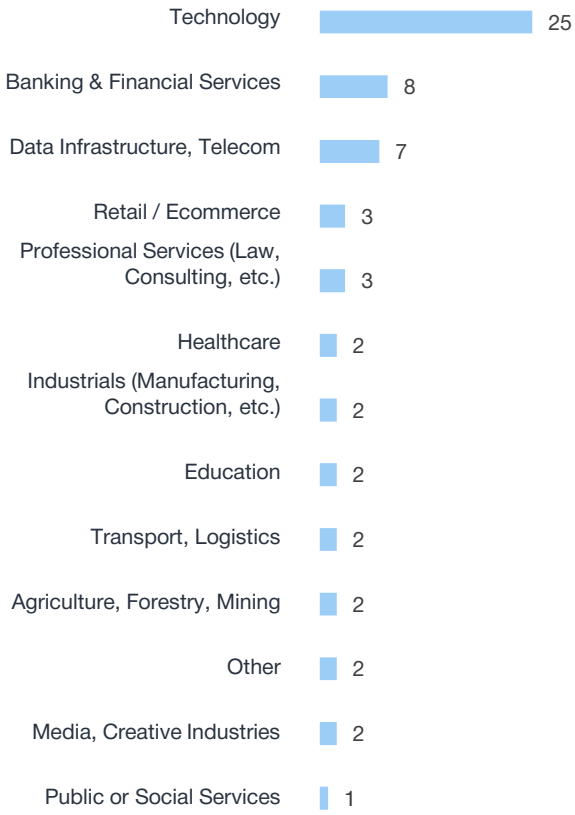

Compared to other industries, organizations in the technology sector clearly spent much more time actively searching for vulnerabilities – perhaps because tech firms are more familiar with the risks.

Preparing a patch in-house

When deploying patches for a vulnerability one key part of the process is obtaining the right patch for the affected system or application. It can happen that an official or vendor patch has not been released, and that the only available information covers the exploit itself with no advice on how to mitigate.

About half of our respondents have never created or tried to create a patch for a vulnerability in house, which indicates that these respondents are fully reliant on the availability of vendor patches – or some form of public disclosure that points to remediation measure or remedial code. Another significant proportion, over 25%, suggested that they attempted to create a patch – but didn't achieve the desired outcome.

There are several obvious barriers to developing a patch. Doing so requires an extensive understanding of the OS or application that needs to be patched as well as knowledge of how the patch may affect other subsystems. Extensive testing is also critical to validate that the patch is effective.

HAVE YOU EVER PREPARED A PATCH IN-HOUSE?

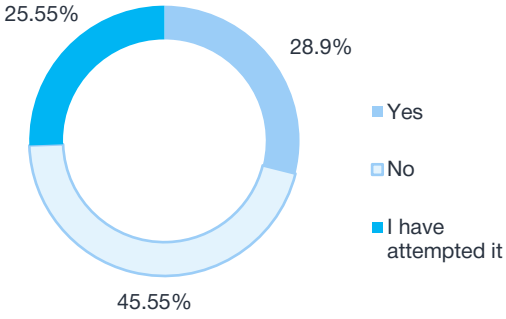

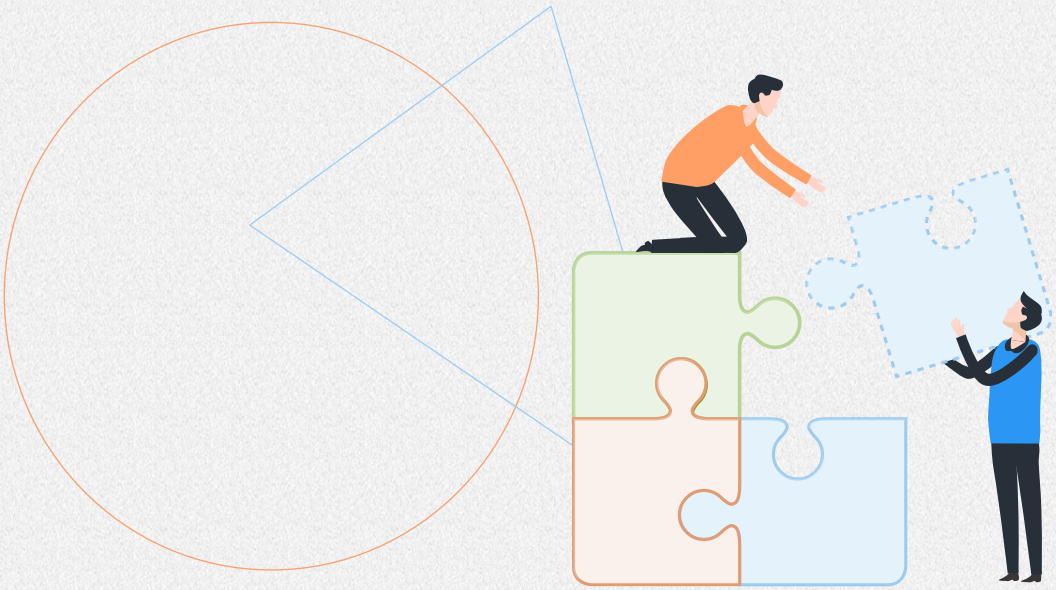

Downtime to deploy patches

After a patch is obtained the patch must be deployed to the affected systems and that often requires a reboot which results in downtime or service disruption. We asked our respondents to state how many hours of downtime their workloads typically experience every week in response to patching.

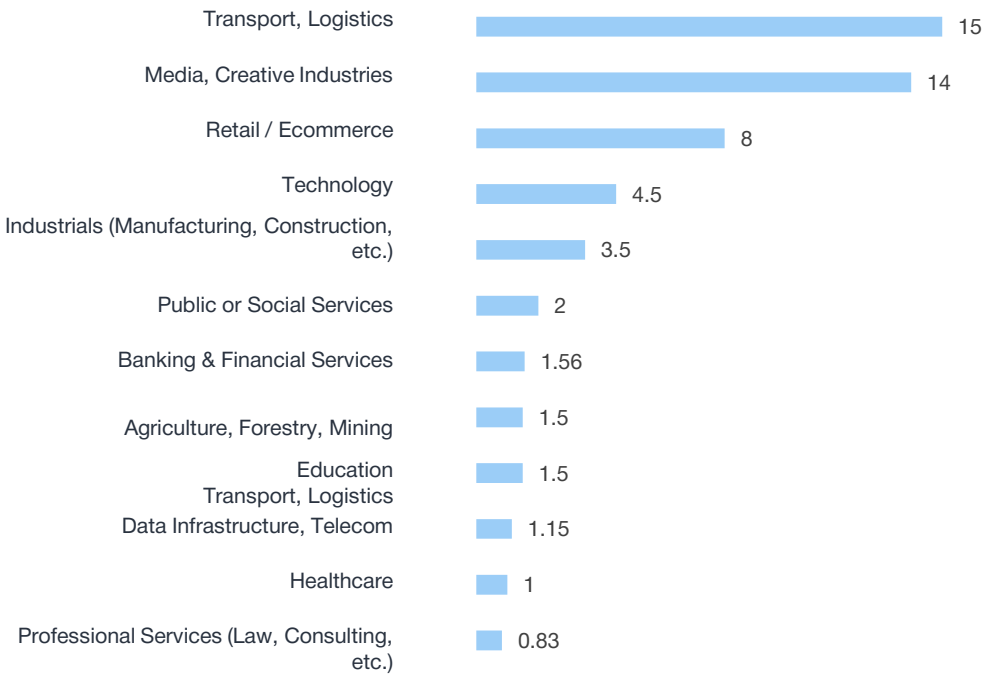

AVERAGE HOURS OF DOWNTIME FOR PATCHING PER WEEK, PER INDUSTRY

Most industries reported less than two hours per week lost to patching procedures. However, two industries reported outsize numbers – transport and logistics and media and creative industries both reported considerably higher hours lost due to patching.

Staff time consumed by patching tasks

Patching can be divided into subtasks. Teams must coordinate downtime with stakeholders, processes need to be documented, and some patches need to be installed outside the maintenance window.

The overall time spent on patching, as broken down by industry, does not tell us that much because different industries will have different software stacks and varying regulatory requirements. The expectations of end users around availability also vary, so some organizations may not be able to arrange for downtime as easily as others.

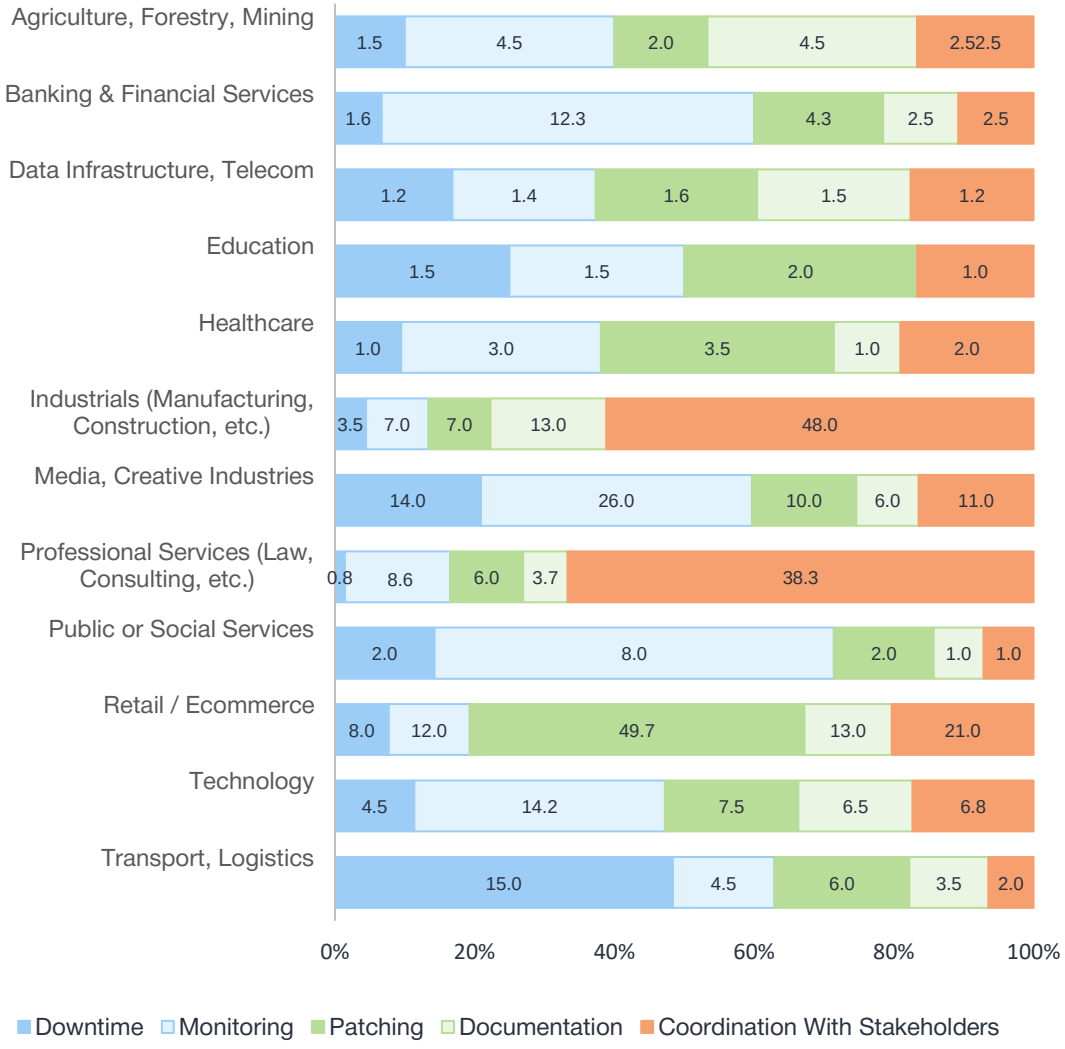

However, it is interesting to note that, for each industry, the proportion of time devoted to each individual subtask varies dramatically.

When patching is delayed

In an ideal world, patching will occur immediately after a vulnerability is disclosed, but there are several factors that affect this process. In the graph below, we allowed respondents to choose multiple options – the numbers on the graph represent the number of times an option was selected.

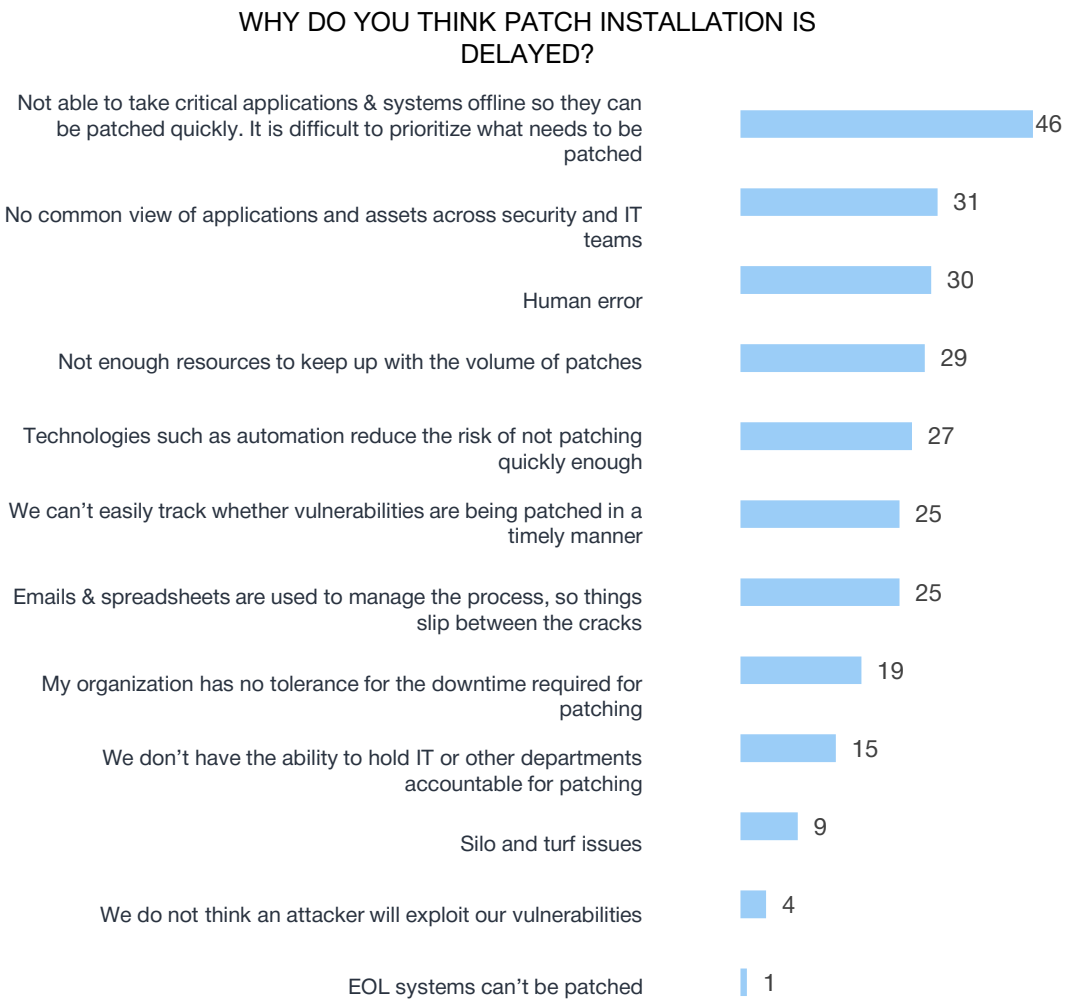

The results show that in many organizations there is no mechanism in place that ensures that patches can be deployed to business-critical systems in a timely manner.

It may be that high-availability architectures are not resilient enough to cope with patching, and that organizations are not using live patching mechanisms to deploy patches without disruption. Either way, the result is that unpatched systems become high-value targets for attackers.

Improvements

Finally, we wanted to gauge which steps would help to improve the outcomes for technology teams that are responsible for vulnerability and patch management.

WHAT STEPS WOULD YOU TAKE TO IMPROVE YOUR ORGANIZATION'S PATCH MANAGEMENT

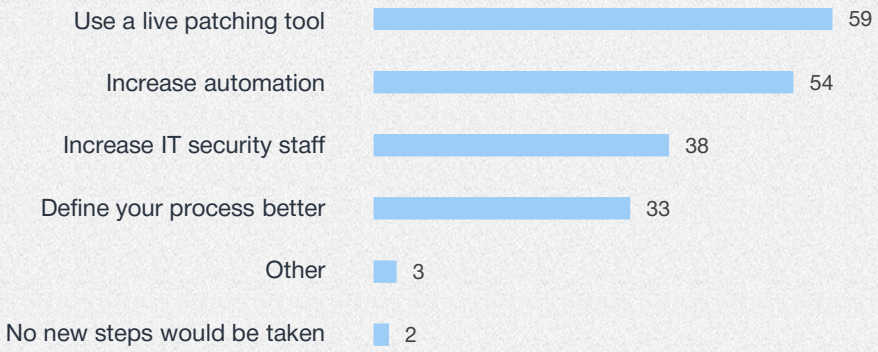

We received three responses under "other", including "Fix lifecycle issues by not running EOL systems", "Enforce culpability for patch delayed for internal company politics" and "Have senior management (outside of IT) hold business units responsible for patching". The first answer is self-explanatory, but the second and third answers are more interesting and may point to problems around silos and whose "turf" it is when it comes to patching.

One pain point that is commonly mentioned is a lack of resources to deal with the mounting workload generated by the growing number of vulnerabilities. We asked our respondents whether they consider that their staff count was sufficient to meet the workload and if they planned to increase staff numbers within the next year.

More than half, 54.3%, of respondents indicated that their staff is insufficient to meet their workload – half of which indicated that they plan on hiring additional staff members that are dedicated to patching tasks.

IS YOUR STAFF SUFFICIENT AND WILL YOU HIRE MORE STAFF DEDICATED TO PATCHING IN THE NEXT 12 MONTHS?

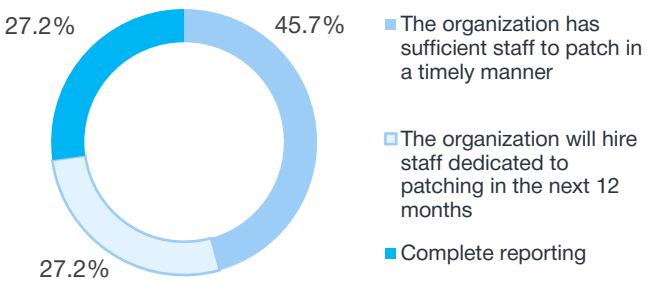

Conclusion

One of the challenging aspects of mounting an effective cybersecurity response is the inherent tradeoffs. Consider, for example, the tradeoff between the availability of resources and the cost of those resources. Likewise, there is the still commonplace tradeoff between availability of services, and fast and timely patching.

At times it can appear as if these tradeoffs are irreconcilable, but there are tools that can bridge the gap – including automation. Indeed, automation is a must-have tool given the automated nature of the cyberthreat. Nonetheless, in our survey, we found that automation – amongst other tools – is unevenly embraced.

It is true, of course, that every industry faces unique challenges and unique tradeoffs. However, irrespective of what these tradeoffs are, every organization must work to adapt to a growing threat. Live patching and other tools will help, and many companies have already cottoned on to the benefits of standardization.

The cybersecurity threat is not going to recede. Vulnerability and patch management efforts must be run intelligently, and must be adequately resourced to meet the growing security challenge.

TuxCare

Thanks for reading the report!

While we've received a meaningful number of responses, **the survey is still running**, and we are eager to increase the amount of answers to build a more complete picture of vulnerability and patch management in the enterprise environment.

Have your say!

Participate in the survey and get a chance to win one of ten
Certified Kubernetes Administrator Certification from The Linux Foundation

START THE SURVEY

*To avoid spam submissions, only users with corporate email addresses can participate in the raffle.