




PENTERA

WHITE PAPER

Aligning Automated Penetration Testing And Risk Management



© 2020 PenteraLtd. All rights reserved. This publication may not be reproduced or distributed in any form without Pentera's prior written permission. While the information contained in this publication has been obtained from sources believed to be reliable, Pentera disclaims all warranties as to the accuracy, completeness or adequacy of such information or analysis. This research is produced independently by Pentera without input or influence from any third party.

Table of contents

EXECUTIVE SUMMARY	4
RISK – A PRIMER	5
Definition	5
Parameters	5
Threat	6
Impact	6
Likelihood	6
Vulnerability	6
Asset	6
Attaching risks	7
Opportunity risk	7
Control risk	7
Hazard risk	7
MANAGING RISK	8
ERM, ORM, ITRM	8
Responsibility	8
1st line of defence	8
2nd line of defence	8
The risk management process	9
Risk response	9
Tolerate	9
Terminate	9
Transfer	9
Treat	9
CONTROLS & RISKS	10
Inherent & residual risk	10
Validating controls	11
Vulnerability scanning	12
Penetration testing	12
Other	12
Continuous control validation	13
The best of both worlds	14
THE EVOLUTION OF SECURITY TESTING	14
Maturity models	14
Comparing maturity of validation techniques	15

Evolution, not revolution	15
Automated	16
Frequent	16
Wider coverage	16
Consistent	16
Ethical hacking techniques	16
Linked to exploitability	16
BENEFITS OF AUTOMATING PENETRATION TESTING	17
Increasing cadence	17
Increasing productivity	18
Efficiency gains	18
Cost savings	18
Reducing risk	19
The test-fix-retest-confirm cycle	19
Penetration test reporting	19
THE POSITIVE EFFECT ON RISK MANAGEMENT	20
Improved information flow – strategic advantage	20
Consistency across business units	20
Confirm your risk assessments	20
Motivation to fix	21
Attacker behaviour	21
Moving up the maturity model	21
ABOUT PENTERA	22
APPENDIX A: TEST, TRAIN & TUNE USING THE MITRE ATT&CK FRAMEWORK	23
Test	23
Train	24
Tune	24

Executive Summary

Risk Management as an organisational function has become much more mature over the last couple of decades, driven by corporate failures such as Enron, Northern Rock and others. Whether it was failure to recognise risks from changing external factors or a failure of internal governance and communication structures, ultimately change was required in order to appreciate the wide range of risks that could affect business objectives and manage them accordingly within corporate constraints.

Consequently, organisations have implemented a range of risk management frameworks to help them deal with uncertainty. Whether it is NIST SP800-37, OCEG GRC Capability Model, COSO ERM Integrated Framework or even a roll your own framework they all share a common theme, which is that they are process driven and iterative. There are very good reasons for this in that organisations have to continue to identify and manage risk in an ever-changing operational environment, both external and internal.

As well as setting out organisational roles, structures and processes, they also guide organisations through common tasks; identifying threats that could impact business objectives, calculating associated risks, defining the control environment, then implementing and monitoring controls on an ongoing basis such that risks are managed and continue to be managed. It is in the monitoring that current challenges lie and that is the focus of this paper. It proposes methods to move away from point in time testing and validation, to an iterative and continuous way so that control environment validation aligns with the risk management processes carried out in accordance with the framework organisations follow and the cadence at which those processes are executed.

In order to deliver on the concept of continuous control validation, new ways of delivering automated testing are required and this paper introduces **automated penetration testing as the evolution of control validation techniques.**

By taking advantage of automated penetration testing, a number of business benefits can be realised. Organisations can accelerate their control validation, report on the control environment in a risk-based and timely manner and embed control validation into their risk management processes.

Ultimately this gets the right data to the right people at the right time so that business decisions can be made, driving down risk and increasing business agility.

Risk – A Primer

Definition

Let's start with the obvious, risk is a complex topic and consequently does not have a single definition. Even across organisations which exist to promote risk management there are differences in the way that risk is defined. For example;

A measure of the extent to which an entity is threatened by a potential circumstance or event, and typically a function of: (i) the adverse impacts that would arise if the circumstance or event occurs; and (ii) the likelihood of occurrence.

[NIST Special Publication 800-30 Revision 1]

The possibility that events will occur and affect the achievement of objectives.

[COSO ERM Framework 2017]

Potential for an unwanted outcome resulting from an incident, event, or occurrence, as determined by its likelihood and the associated consequences.

[DHS Risk Lexicon 2010]

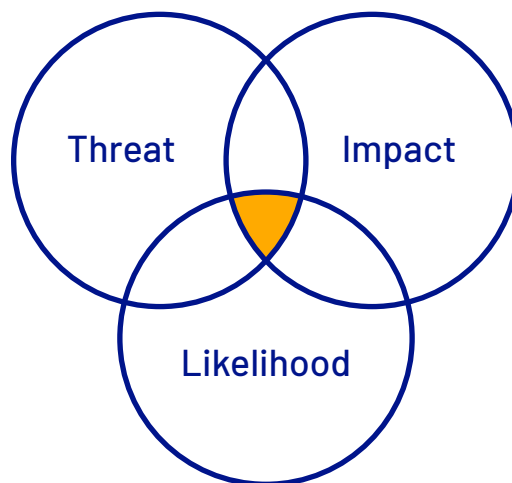


Fig 1: Risk triad

While the definitions are different, we can see commonality running through them. We can also synthesise them down to their simplest component parts; threat, impact and likelihood. While it is recognised there are many different definitions of risk, references to risk in this paper will refer to the above simplified model.

Parameters

Even with a simplified model, there are a number of parameters involved in the assessment of risk, which are worth exploring.

Threat

A threat is something that can adversely affect an objective. It can affect it in different ways depending on what the risk you are calculating is associated with (more on that later). Threats are also calculated using component parts, for example threat actor (further broken down into capability and intent) plus opportunity.

Impact

Impact (or consequence as per the DHS Risk Lexicon) is the effect of the risk event if it occurs. For example, the impact of loss of service to a web-based retailer is loss of revenue, which could ultimately lead to closure of business.

Likelihood

This parameter answers the question 'what are the chances of the threat event happening?'. In risk assessment terms it can be measured quantitatively using calculations such as 'Annualised Rate of Occurrence' or qualitatively using groupings such as 'low, medium, high' for example.

Vulnerability

Much as there is no single definition of risk, there is no single correct answer to where to add vulnerability to the assessment. For example;

- Add to threat. It could be valid to include the vulnerability in the threat calculation, as a threat actor exploiting one vulnerability could result in a different likelihood or impact than another and needs to be calculated as a separate risk.
- Add to likelihood. It is equally valid to add a vulnerability to likelihood as it is logical that the likelihood of a successful threat event is higher if a vulnerability exists than if it does not.

Asset

Risks are associated with assets and asset identification is one of the first steps in your risk management process (more later). For each asset, you can identify the range of threats the asset is exposed to and then for each threat, calculate the likelihood and impact. This results in a large set of risks mapped to assets, prioritised in whichever way is right for your business.

But what is an asset? An asset can be many things from data to a server, to a person and many things in between. **The type of assets you will consider depend really on where risks are attached**, which is the subject of the next section.

Attaching risks

Consider the following process flow, which starts with a decision at board level to set a strategic objective to pursue an opportunity. The purpose of the opportunity is to deliver value of some sort to the business, whether that is to drive growth, or to eliminate overhead and drive efficiencies.

To deliver the objective, it is determined a new business process is required. In this case it is a business process that can be delivered using a system, so that triggers a project to build and deliver the system. This system can be a combination of people, process and technology.

Once the system is delivered, the process becomes part of business as usual (BAU) operations and enables the objective to be met.

Risks can be attached to this process in the following manner;

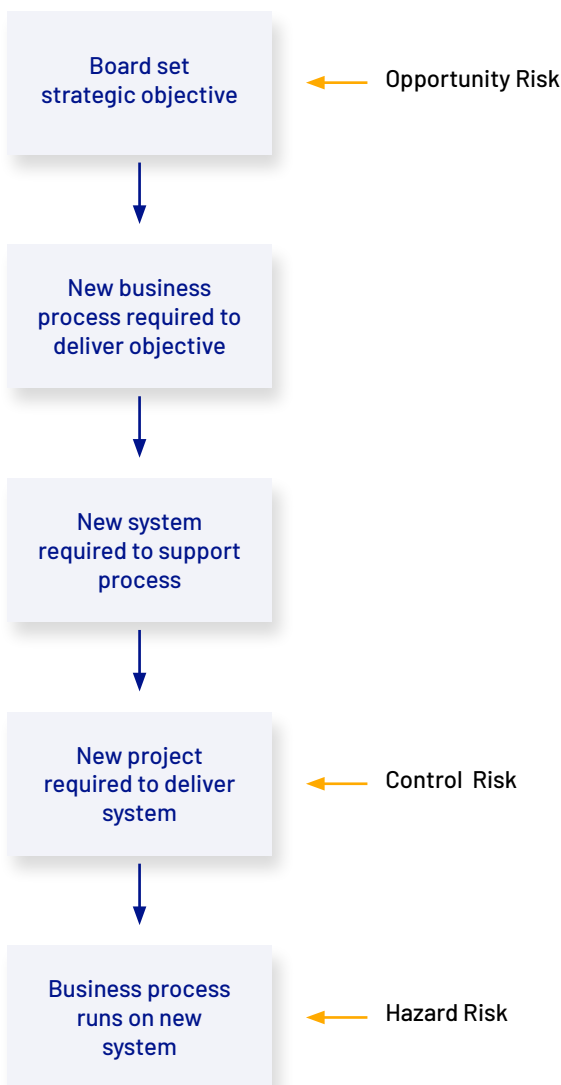


Fig 2: Attaching risks

Opportunity Risk

The board has identified an opportunity that will deliver value to the business in some way and has set an objective to pursue the opportunity. For example, the business is moving into a new geographical region. The opportunity risk is that the value is not realised when the objective has been achieved. Opportunity risk also considers the cost of not pursuing the opportunity, such as a competitor moves into the region instead.

Control Risk

Control risks are based on uncertainty of outcomes, are common in project delivery and are mitigated by sound project and programme management. In this case, delivery of the system requires a new project with a project team, project manager and project/programme board.

Hazard Risk

Hazard risks are the risks involves in the day to day activities of a business. They are always downside risks and this is where operational risk sits. IT risk also sits here as the information flow is up into Operational Risk.

Managing Risk

ERM, ORM, ITRM

Enterprise Risk Management, Operational Risk Management and IT Risk Management are connected within an overall risk management system to expedite information flow and deliver effective risk management throughout the business. In order to make informed business decisions, it is essential that accurate and timely risk-based information flows through the layers of the risk management system efficiently and speedily.

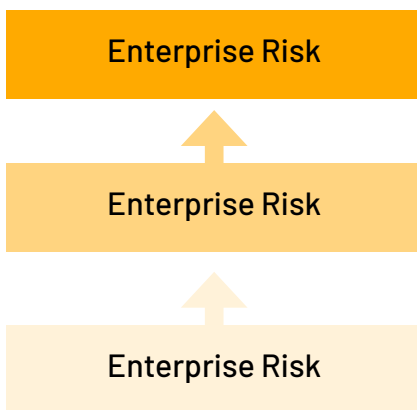


Fig 2: Attaching risks

Business units and teams within each component part will be responsible for the risk management process within their particular area, then reporting upward such that, at the top level, there is a holistic view of risk across the organisation.

Responsibility

Responsibility for risk management activities can reside in different places depending on the size of an organisation, it's vertical and the accompanying regulations that apply. Let's consider a 2LOD (where LOD = line of defence) model:

1st line of defence

The 1st line of defence is generally the line of business that owns, deploys and manages the system, data and/or business process. They are responsible for carrying out the risk management process covering the scope of their activities and often are able to self-assess and accept a level of risk subject to a given limit. In some cases, the responsibility for the underlying systems can be held by another party such as a central IT team or even a cloud service provider.

2nd line of defence

The 2nd line of defence is a relatively independent entity, such as a centralised risk team, or a compliance team. They do not own the system, data or business process but they will undertake an impartial assessment of some sort to validate the assertions the 1st line of defence are making. This avoids the scenario of marking your own homework and is mandated in some circumstances.

There is also a 3rd line of defence which is more independent and typically an audit function, sometimes external.

Now onto the process itself.

The Risk Management Process

When it comes to managing risk, one of the most important things is that it is a repeatable process and not simply a point in time activity. Fig 4 shows a typical example of such a process and is broken down into a number of steps, from initial identification through to prioritisation, response and reporting/monitoring:

This paper is not meant to explain the entire risk management process, but risk response is worth focusing on.

Risk Response

When it comes to responding to risk, we have a number of options, referred to as the 4 T's.

1. Tolerate
2. Terminate
3. Transfer
4. Treat

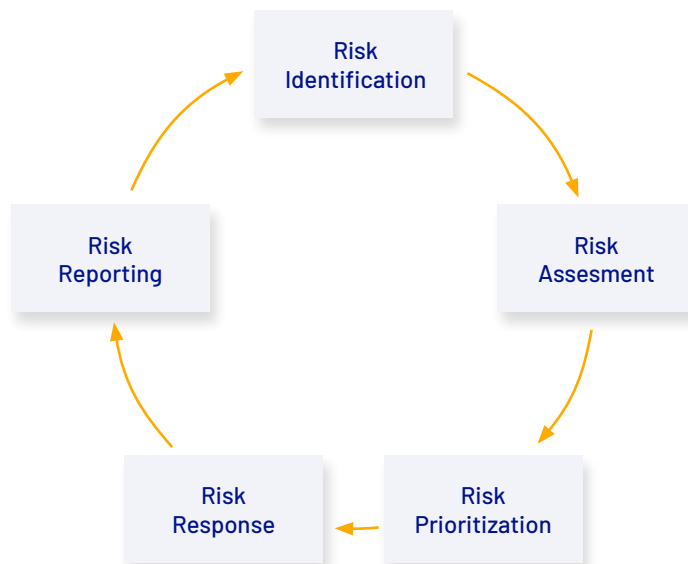


Fig 4: The risk management process

Tolerate

Tolerating a risk means that it falls within the bounds of an accepted risk with no additional mitigating controls. This could be on a permanent or temporary basis.

Terminate

When a risk is terminated, it means the process, system or event that introduced the risk is ceased. This would be the appropriate response for a non-essential service where the impact of ceasing the service is less than impact calculated if the risk materialises.

Transfer

Typically, when we talk about transferring the risk, we are generally talking about insurance. This transfers liability for the costs involved if the risk materialises to the insurer and consequently, the impact to the Organisation is reduced.

Treat

When we treat a risk, we apply a control to mitigate the risk from an unacceptable level (the inherent risk) to an acceptable level (the residual risk). This is the response we will focus on in this paper and the next section looks in more detail at controls, risks and the relationship between them.

Controls & Risks

Inherent & Residual Risk

Controls are the primary mitigating technique used in treating risk. They are used to reduce the risk from an unacceptable level to an acceptable level and can be technical or procedural.

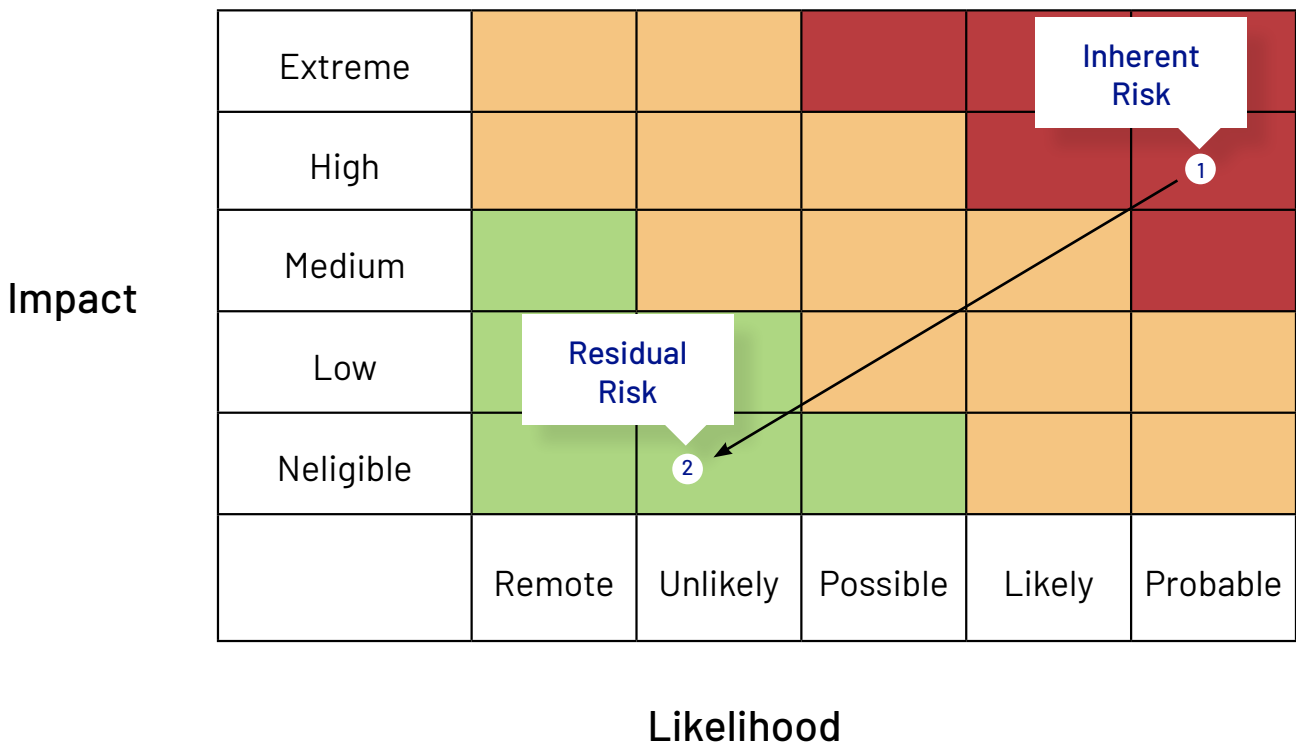


Fig 5: Inherent vs residual risk

The above heat map is specific to a given risk. In this case, the inherent risk (the risk with no control applied) is outside acceptable bounds. A control has been identified and after applying the mitigating control, the residual risk has been calculated and that brings us to within the acceptable limits.

You can already see the importance of the control in that it reduces the risk level but also the confidence you need to have in the control working correctly to mitigate risk as it is expected and calculated in the risk assessment.

If a control is not effective in mitigating the risk, it could be for a number of reasons, such as failure of the technology, or administrator misconfiguration (in fact some risk assessment methodologies call this out as a specific risk). This is where control validation comes into play.

Validating Controls

Controls typically are recorded in the risk treatment plan, which contains risks in prioritised order along with a number of fields including the control, the owner and how the control will be validated and monitored.

Controls are validated to make sure they work. If a control fails, then the risk assessment is no longer accurate, so it is essential they are validated by testing on deployment and continue to be regularly tested on an iterative basis.

There are a number of ways in which controls can be tested on a regular basis, including:

- vulnerability scanning
- penetration testing
- other (configuration checks, documentation review etc)

The following graphic shows the relationship between the risks, controls and validation techniques.

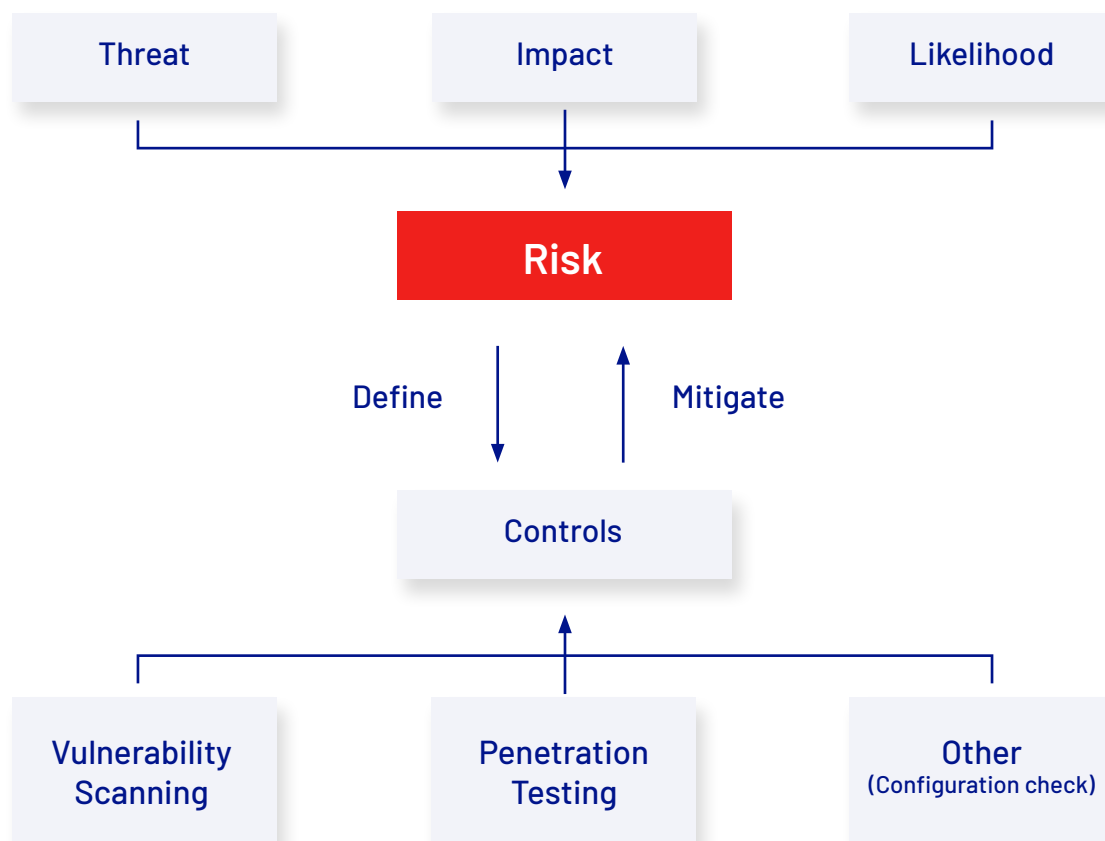


Fig 6: Connecting risk, controls & validation

Let us now dig deeper into each of these techniques.

Vulnerability Scanning

Delivery	Automated
Frequency	Regular cadence to suit operational cycles
Coverage	Wide, due to speed and scale of automation
Consistency	High, using automated tooling
Methods	Static Vulnerabilities
Linked to exploitability	No
Primary Issue	Volume of remediation effort required

Table 1: Vulnerability Scanning characteristics

Penetration Testing

Delivery	Manual
Frequency	Typically annually to suit compliance mandates
Coverage	Narrow, based on defined scope due to cost, resources or both
Consistency	Low to medium, dependent on provider, individual testers skillset and methodology
Methods	Static vulnerabilities + ethical hacking techniques
Linked to exploitability	Yes
Primary Issue	Time & resource intensive limits scope and coverage

Table 2: Penetration Testing characteristics

Other

(Configuration checks in this example)

Delivery	Manual
Frequency	Infrequent. Typically on creation & during penetration test
Coverage	Narrow, based on penetration test scope
Consistency	Low to medium if manual, higher if tooling is used
Methods	Manual system check
Linked to exploitability	No
Primary Issue	Time & resource intensive, irregular cadence

Table 3: Configuration checking characteristics

The conclusion is that different validation techniques offer different advantages and disadvantages over each other. For example;

Vulnerability scanning uses automation, therefore offers a speed, scale and consistency advantage over manual penetration testing. However, **that is offset by the main issue of organisations struggling to remediate the high volume of vulnerabilities** because they are not linked to exploitability.

Penetration testing on the other hand is not commonly conducted in the same iterative and continuous way as vulnerability scanning, mainly due to the manual nature of delivery. Nor can consistency be guaranteed, due to firstly, the mandate many organisations have for multiple service providers, but also down to the individual skill sets of the testers (also an issue for inhouse teams). But remediation becomes much more efficient because of the link to exploitability. The outcome is a reduction in the volume of work and effort can be focused on where it has most impact.

As we noted earlier, as the risk management process is a repeatable process, it would not make sense for validation of the mitigating controls to be a one-off activity at the point of deployment, or simply to be an annual exercise. Control validation should be repeatable and continuous and ideally aligned with the risk management process, as it is this validation that is proving the assertions made in the risk assessments.

Continuous Control Validation

There are a number of reasons why control validation should be a continuous process;

- **Change in IT landscape** – whether it is at the macro level, such as post M&A integration, standing up new systems to support organic growth etc, or at the micro level, where users and individual devices change rapidly, all companies undergo some sort of change in their IT landscape, which changes the attack-surface
- **Change in threat landscape** – new threat actor groups, new campaigns, new strains of malware and of course, new vulnerabilities appear on a daily basis.
- **Change in regulatory landscape** – many regulations, such as GDPR, are now mandating regular, repeatable security testing. In fact, non-compliance with mandated regulations should sit as a risk in itself as there is a potential impact to non-compliance, such as financial penalties.

You can see it is constantly shifting sands, and the only way to keep up is to conduct regular repeatable testing. We saw above that different types of testing had advantages and disadvantages so how best to be able to test controls on a continuous basis?

The Best of Both Worlds

What if we take the best of Vulnerability Scanning and add it to the best of Penetration Testing? If we construct a table similar to the ones above, the result looks like this:

Delivery	Automated
Frequency	Regular cadence to suit operational cycles
Coverage	Wide, due to speed and scale of automation
Consistency	High, using automated tooling
Methods	Static Vulnerabilities + ethical hacking techniques
Linked to exploitability	Yes

Table 4: Optimal control validation characteristics

Penetration Testing

Now we have a testing capability that benefits from the speed and scale that automation brings and additionally, uses ethical hacking techniques and is linked to exploitability. These represent the characteristics we want in an automated penetration testing platform, characteristics which will allow us to improve on the way we validate security controls and consequently, validate our risk assessments.

The Evolution of Security Testing

Maturity Models

When it comes to improvement, maturity models give us an indication of where we are starting from and where we want to get to. There are a number of different maturity models but one of particular relevance to this paper is the Cybersecurity Maturity Model mandated by the United States Department of Defence. Like many maturity models it consists of a number of levels with underlying requirements that organisations can follow to improve their working practices. In this case, the levels are thus:

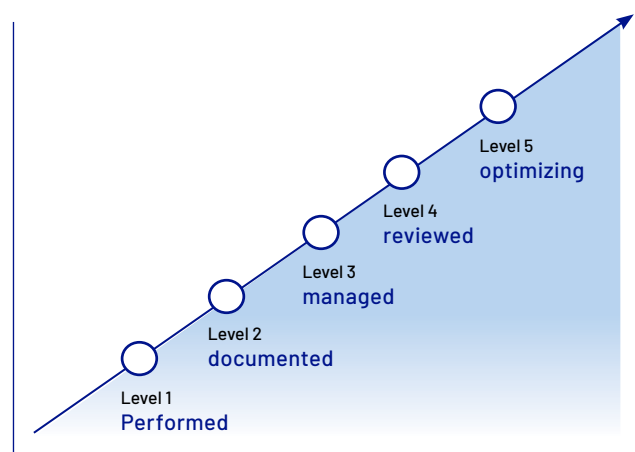


Fig 7: Maturity Models

The maturity model can be applied to any underlying workflow or process, so if an organisation follows the NIST Risk Management Framework, they can apply the model to the practices and processes they are implementing and following as part of that framework.

For each given process, the model will indicate what the current level is, from Level 1 which is more of an ad-hoc approach, all the way through to level 5 where it is documented in detail, understood, embedded in standard working practices and uses feedback for constant process improvement.

Comparing maturity of validation techniques

Generally, an organisation would have a more mature vulnerability management programme compared to penetration testing. Two of the reasons for this are;

- use of automated tooling
- owned and operated by in-house teams

The use of automated tooling alone means that vulnerability scans can be configured to run regularly with limited manual input. The reports generated can then be input into various workflows owned and executed by a number of internal teams such as those responsible for remediation as well as risk reporting. Additionally, the tooling and process allows for a high level of consistency and repeatability. This places the vulnerability management program at a medium to high level of maturity depending on specifics.

Penetration testing on the other hand often suffers from being viewed as a compliance exercise and is therefore a once a year activity for many organisations. Those organisations with larger teams and budgets may operate at a more frequent level of testing but either way, given the manual nature of delivering penetration tests, it is a project that needs managing every time. This tends to increase cost, increase duration, reduce scope and of course loses consistency particularly where different service providers and even individual testers skillsets are taken into account. For a lot of organisations this represents too onerous a task to make it anything more than an annual or even an ad-hoc activity, which places it in the lower levels of maturity.

Evolution, not revolution

The characteristics in table 7 represent the best elements of vulnerability scanning, namely:

- Automated
- Frequent
- Wider coverage
- Consistent

Added to the best elements of penetration testing:

- Ethical hacking techniques
- Linked to exploitability

This is where we want to get to with our control validation, so let us look at these individually in greater detail.

Automated

Automation can be delivered using different methods from rules-based engines to ML and AI, depending on the nature of the data, the process and what decisions need to be made on the data. Sometimes it makes sense to use AI, sometimes not and a rules-based engine is more appropriate. Whatever method is used, it is the underlying principle that drives many of these other characteristics.

Frequent

When you automate a process, it allows the process to be iterated through at speed with minimal to no manual input (in some circumstances, a degree of manual input is mandated). In this case, while it is possible to automatically trigger a vulnerability scan, that's not the case with the manual penetration test, so let's take that capability and apply it to penetration testing as well. Once we do that, we want to be able to schedule a penetration test to start at a time and date to suit, as well as let the penetration test run in the background while the operator gets on with their day job.

Wider Coverage

It stands to reason that if I can increase the frequency at which I run my penetration tests I can execute more penetration tests. So, if I target those tests across different parts of my network then I can cover much more ground than I could using manual techniques.

Consistent

When you run a vulnerability scanner, it will execute its tasks in the same way to the same level every time. That is not necessarily true for manual penetration testing for a couple of reasons. Firstly, you may have a mandate to cycle through different service providers who will undertake their testing using their own methodologies, plus the experience & skillsets of the individual testers varies. If penetration testing is delivered using an automated platform, that removes the inconsistencies leading to repeatable testing.

Ethical hacking techniques

Vulnerability scanners scan for static vulnerabilities, with an entry on the CVE database and a CVSS score. Penetration tests on the other hand will involve an element of static vulnerability scanning but additionally will use dynamic vulnerabilities such as misconfigurations, abuse of protocols and services, credential related attacks and more. Not only does a penetration test offer more in terms of the techniques it uses, but those techniques are the same as malicious third-parties use, as we can map them to the MITRE ATT&CK Framework.

Linked to exploitability

One of the main advantages of penetration testing over vulnerability scanning is related to the volume of remediation activity. A good penetration test should exploit a range of vulnerabilities, both static and dynamic, to successfully execute attack techniques and generate a report to show those underlying vulnerabilities that need to be fixed so that the attack technique is no longer possible. This link to exploitability, a major advantage over vulnerability scanning, means a large reduction in remediation effort, which historically has been a constant problem with vulnerability scanning and vulnerability management programmes.

Benefits of automating Penetration Testing

At this point, we have looked at the importance of continuous validation of your security controls and how they are key to the accuracy of your risk assessments. We have also looked at the best of vulnerability scanning, the best of penetration testing and how they can be brought together in an automated penetration testing platform.

Now let's take a look at the benefits we can realise when we use automated penetration testing. From our discussions with our customer base, the benefits synthesise down into these areas:

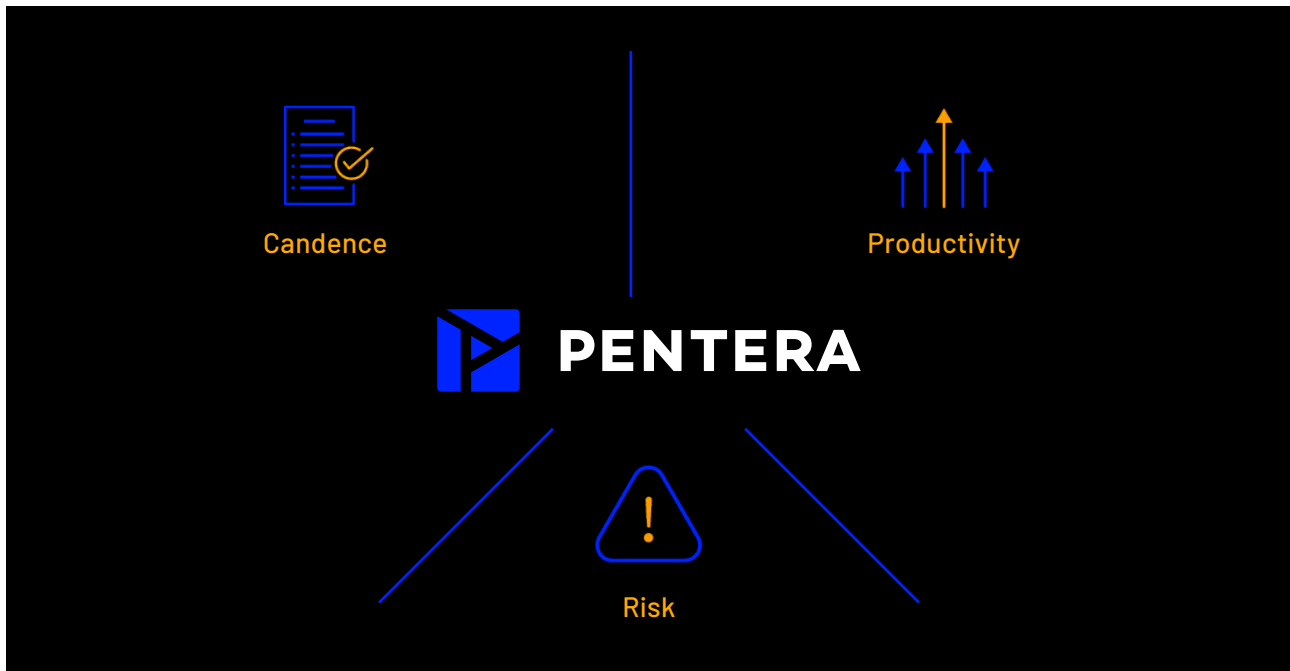


Fig 8: Business benefits

Increasing Cadence

As we discussed earlier, automation allows you to increase the cadence at which you can run your penetration tests. This can allow for:

- The same part of your network to be tested on a much more frequent basis
- More of your network to be tested
- Retesting after remediation activity

The third point is really important as it enables you to execute on the concept of continuous security validation as discussed earlier.

Increasing Productivity

Efficiency gains

When companies embrace automation, one of the main drivers is to increase efficiencies and reduce costs. When you apply automation to a traditionally manual activity such as penetration testing, a lot of the efficiency gains are realised in two places.

Firstly, there is the reduced amount of remediation activity, which is reduced because of the link to exploitability. Ops teams and other teams who the remediation activity falls to will no longer be expending time and effort applying changes which have limited material impact. Instead they can focus their effort on the lower number of underlying vulnerabilities which are exploitable and have led to a successfully executed attack during the test.

Secondly, there are productivity gains because of changes to the penetration testing teams workload. In-house penetration testing teams can automate as much of their manual activity as they can which frees them up to concentrate on other tasks such as social engineering or testing esoteric devices and applications. Not only this but the automated tests can be configured, scheduled, executed and reported on in the background and in parallel to whatever other work needs to be carried out. This makes it a very efficient process.

Cost savings

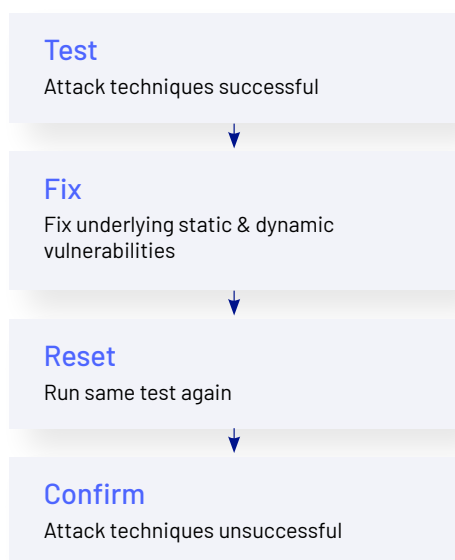
If penetration testing is outsourced, we see customers automating the majority of their penetration testing activity leaving a much smaller scope for the service provider, which reduces costs. What we see very often is an organisation will run a range of automated penetration tests throughout the year themselves and then use a third-party provider to deliver the annual compliance test, in accordance with mandated regulations for example.

Reducing Risk

We can look at risk reduction in different ways.

The test-fix-retest-confirm cycle

Firstly, if we remember the simplified model of risk there were 3 component parts; threat, impact and likelihood. Consider the cycle of an automated penetration test;



- I. When you first run the automated penetration test, it is possible to successfully execute attack techniques. Again, these should be mapped to the MITRE ATT&CK framework and executed using the same tactics and techniques as attackers are known to use.
- II. The gaps which allow the attack techniques to be executed are exposed in reports and remediation activity is undertaken.
- III. The automated test is rerun with the same scope and settings for consistency.
- IV. The attack techniques are now unsuccessful, your security controls have been validated as working effectively and your risk assessments are correct.

Fig 8: Business benefits

Now that the automated penetration test cycle has confirmed the attack techniques are no longer possible, if an attacker was able to compromise your network and run those same attack techniques, they would be unsuccessful as well. **You have reduced the impact of a successful attack and consequently reduced the risk.**

Penetration test reporting

Any automated penetration testing platform should offer a range of reports immediately, from the detailed to the executive summary. As you increase the cadence at which you conduct penetration testing, it follows that reporting will be much more up to date. Whether risk reports are taken directly from the automated penetration testing platform or whether there is integration into a holistic GRC tool, you will be able to benefit from a clearer, timelier picture of where your risks lie and you no longer have to rely on a penetration testing report that could possibly be up to a year old.

Where outsourced providers are concerned, firstly you don't have to wait for them to process and send the report and deliver the presentation. Connected with that, you no longer have your sensitive security related data leaving the building on a penetration testers laptop while they upload it to their own system where it could persist far beyond the life of your penetration test.

The Positive Effect on Risk Management

We have seen there are clear benefits to taking the best of vulnerability management and manual penetration testing and building it into an automated penetration testing platform. But how does using that platform help us improve our risk management process?

Improved information flow – strategic advantage

We saw earlier the connection between IT risk, operational risk and enterprise risk and also how the more accurate and up to date the risk information is, the better informed the decision makers are and better business decisions can be made in a timely manner. By moving to an automated, frequent, iterative process, the speed at which the risk related data moves through the levels of the risk organisation increases and the quicker it is made available to the business decision makers. Making quicker business decisions based on up to date, accurate data allows for the organisation to increase agility and flexibility leading to strategic advantage.

Consistency across business units

Consider Fig.10, where IT risk exists and is owned and managed across multiple business units

If we automate the penetration testing process across all business units using common tooling and processes, this increases consistency of testing and reporting. This can result in operational risk now only having to deal with a common reporting format and taxonomy which in itself can lead to further efficiency improvements. Operational risk teams could then take further advantage of this new consistency and commonality by improving the efficiency of their own workflows or even automating their own processes.

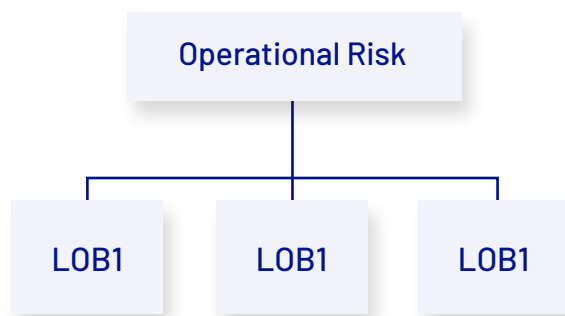


Fig 10: Multiple lines of business

Confirm your risk assessments

Earlier we saw how risks are assessed based on a number of variables, in our case threat, impact and likelihood. Once we have calculated the inherent risk, we use one or more controls to reduce the risk down to an acceptable level, the residual risk. The residual risk level is only valid if the controls work, and continue to work, as designed. Automating the penetration testing process allows organisations to test effectively and validate the risk assessments are indeed correct both at the point they are deployed and also continue to be tested continuously and iteratively, aligned with risk management processes.

Motivation to fix

If we remember the 4 T's, tolerating risk is indeed a valid risk response option and again, is appropriate to be used where risk is already at an acceptable level. Sometimes however, risk is accepted because the barrier to treat the risk is too high for some reason. One of the reasons is linked to the time and effort required to fix the underlying vulnerabilities. The result is that risks which are outside acceptable levels are now being accepted, although often timebound.

If we instead focus on those exploitable vulnerabilities that allow the automated penetration testing platform to execute attack techniques, there are a number of advantages. Firstly, we end up dealing with a much lower number of vulnerabilities, but these vulnerabilities are the really important ones. These are the vulnerabilities that lead to the most impactful return on time and effort spent on remediation activity. Secondly, because testing is repeatable and frequent, the amount of work over a period of time is broken down into packages which are manageable and can realistically fit into an organisations operational cycle.

This in itself increases the motivation to fix by lowering the barrier to treat the underlying vulnerabilities.

Attacker behaviour

Of course, if you are going to test your security controls, you are going to want to use real attacker behaviour to do so. That is how malicious third parties undertake their work, how manual penetration testers undertake their work and how any automated tooling should also undertake its work. An attacker won't ask you to open ports or switch off controls and an attacker won't install agents. If an attacker is going to enumerate and attack all live hosts within a given range, then your automated platform should enumerate and attack all live hosts within a given range. If an attacker moves laterally by instantiating a DCOM object, your automated platform should also move laterally by instantiating a DCOM object.

Additionally, the MITRE ATT&CK Framework provides information on adversary behaviour. Whether it is a given malware strain such as EMOTET, or a specific threat actor such as APT33, the MITRE ATT&CK framework shows the taxonomy of the methods they are known to use in an attack. By automating the same behaviour, carried out in the same way, your controls are tested against techniques that you would be subject to if you were under a real attack and you can have confidence that your risk related data, derived from the output of your automated penetration tests, is fit for purpose. See Appendix A for further information.

Moving up the maturity model

We saw earlier in this paper how maturity models can be used to gauge the maturity of your current processes, identify the maturity level you would like them to be at and plan how to move from current state to future state. We also saw that in general, vulnerability management programmes tend to operate at a higher level of maturity than penetration testing for various reasons, although they lack some of the positive characteristics of penetration testing such as the link to exploitability which leads to the advantages detailed above.

Introducing an automated penetration testing platform enables you as an organisation to firstly, bring your penetration testing and continuous control validation processes to a higher level of maturity covering a number of areas such as incident response, security assessment and more. Secondly, by aligning iterative, automated penetration testing with your risk management processes, you embed continuous control validation into your organisational culture, leading to overall security improvements.

About Pentera

Established in 2015, Pentera delivers Pentera, the automated network penetration testing platform that assesses and helps reduce corporate cyber security risks. Hundreds of security professionals and service providers around the world use Pentera to perform continuous, machine-based penetration tests that improve their immunity against cyber-attacks across their organizational networks. In 2020 the company completed its B round funding of \$25M with Insight Partners and existing investors, the Blackstone Group and AWZ Ventures, bringing the total investment to \$40M. With 170 enterprise global customers across all industries, Pentera is the fastest-growing cyber security startup in Israel. In 2020, Pentera was chosen by Gartner as a Cool Vendor in the Security Operations and Threat Intelligence report.

Our solution - Pentera, the Automated Penetration Testing Platform

Requiring no agents or pre-installations, Pentera™ platform uses an algorithm to scan and ethically penetrate the network with the latest hacking techniques, prioritizing remediation efforts with a threat-facing perspective. The platform enables organizations to focus their resources on the remediation of the vulnerabilities that take part in a damaging “kill chain” without the need to chase down thousands of vulnerabilities that cannot be truly exploited towards data theft, encryption or service disruption.

With Pentera™ a company can maintain the highest resilience posture by performing penetration tests as frequently as needed - daily, weekly or monthly. Pentera allows you to validate your cyber security posture as often as you need, keeping your guard up at all times and maintaining consistency across your organization.

Appendix A: Test, Train & Tune using the MITRE ATT&CK Framework

The MITRE ATT&CK Framework contains a taxonomy of real attacker behaviours found in the wild. When you introduce regular automated penetration testing into your security testing strategy, the automated penetration testing platform should use the same tactics and techniques as real attackers so that your security controls are tested effectively and completely. This test, train, tune approach detailed below covers the people, processes and technology involved in your detection, response and alerting system.

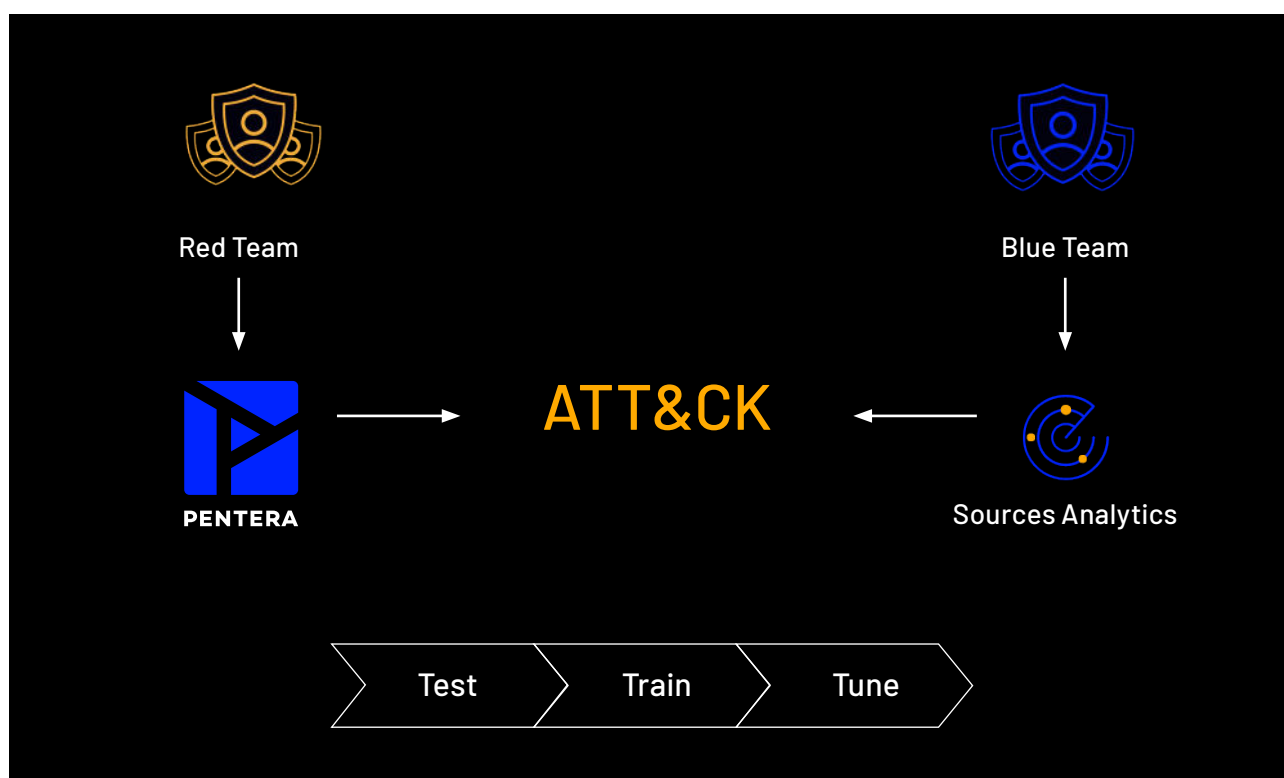


Fig A1: Test, train, tune using the MITRE ATT&CK Framework

The graphic shows how to use Pentera to execute real attack techniques leading to improvements in your detection and response capabilities.

Test

Test repeatedly with Pentera. Pentera will execute the full attack lifecycle mapped to the tactics within the MITRE ATT&CK Framework and the techniques and sub-techniques within each tactic. From reconnaissance through to execution, lateral movement and more, Pentera can test during working hours but you can additionally schedule tests outside of working hours to see how your detection, response and alerting system would cope in the event of an attack at, for example, 1am in the morning.

Train

Train the people. You really don't want the first time your SOC operators see an alert they have to take action on to be during a real attack. By regularly testing, you can train your personnel to identify the correct alerts and take the appropriate action based on what is in front of them. Then when and if the time comes that they are subjected to a real incident, their training kicks in and they know what to do.

Tune

Tune the process and the technology. It may be that regular testing identifies ways to make your SOC team better at their jobs by making the processes they need to follow more efficient, easier to execute, or simply they can iterate through them quicker.

Also, tuning the technology stack can not only tune the event management and corresponding alerts so that events of interest are identified at the right point in an attack, but conversely, there is not too much noise. Not only is the important data at risk of being lost in the noise but, additionally, SOC operator fatigue can be a real problem. By tuning the system down, alerts are generated at a lower frequency but at a frequency good enough to meet the alerting requirements and action to be taken.