Automated Vulnerability Research

Roadmap

version 3.0

3 February 2021

1 Management summary

On the initiative of the Ministries of Defence, Economic Affairs and Climate Change, and Justice and Security, work has been underway for some time on the topic of Automated Vulnerability Research (AVR)[1]. AVR is a relatively new and unknown topic within cybersecurity. The topic is anchored in the Knowledge and Innovation Agenda (KIA) for Security under the Cybersecurity Mission[2].

The ministries concerned would now like to proceed with the next step and develop a Roadmap for working on AVR. The involvement of the entire cybersecurity chain is crucial: from research and education to application (valorisation). The Roadmap consequently represents an example of the 'programming' of the Cybersecurity Cooperation Platform.

Financial resources will be made available to begin developing the Roadmap. In order to fully implement Roadmap additional funds will need to be found in the years ahead. The Roadmap will contribute to this. This Roadmap also provides an initial indication of how the available funds could be spent, and how the decision-making process could take place.

AVR Roadmap goals
The overall goals of the AVR Roadmap are as follows:

- Clarifying how to work on AVR in the Netherlands in the next six years.
- Coordinating the commencement of AVR activities based on three tracks to boost the development and application of AVR knowledge.
- Clarifying where and how partners in the chain can participate in taking AVR to the next level in the Netherlands.
- Inspiring partners in the chain to give new impetus to AVR.
- Building a vibrant community in the field of AVR.
- Providing insight into what can be achieved in the short and medium term, clarifying both the opportunities and limitations.
- Contributing to the digital autonomy of the Netherlands.

AVR Roadmap ambition
AVR strengthens the preventive component of the cybersecurity chain by automating the detection of known and unknown vulnerabilities in software, developing exploits and patching vulnerabilities. A high ambition has been set for the AVR Roadmap. A 'dot on the horizon'.

The ambition is to achieve the following in six years' time:

- The Netherlands has a *leading position* in the field of Automated Vulnerability Research (AVR) in Europe. This is demonstrated by the combination of publications (thought leadership), PhDs, the

[1] **See Chapter 3 for an explanation of AVR.**
[2] **See Chapter 2 for an explanation of the Security KIA.**

number of students, patents (or open sourced technologies), the volume of commercial services, among other factors.

- *Technology has been developed* for the successful automated detection of vulnerabilities on a commercial basis.
    - In at least one substantial IT development or management activity (TRL 8+).
    - A good starting position has been achieved for the valorisation and industrialisation (TRL6+ demonstrators) of automatic patching.
- Sector-specific applications have been developed, such as the capacity for the automated detection and/or patching of vulnerabilities in software used by banks.

The underlying policy objectives are to promote digital security, including for small and medium-sized enterprises (SMEs), and to strengthen economic earning capacity and strategic autonomy in the digital domain.

Status, needs and application areas

The current status of AVR in the Netherlands and the needs of the parties concerned are summarised in the Roadmap. This includes all partners in the innovation chain: the higher education sector (research universities and universities of applied sciences), applied research institutes, security technology suppliers, security service providers, software developers and IT end users (particularly from the critical sectors).

All parties concerned have a need to take steps in the field of AVR.

The application of AVR in the following four sub-areas is of particular interest[3]:

1. *Secure software development.* By applying AVR to the software development process, vulnerabilities can be resolved at an early stage, before the software is deployed. This is in line with the movement towards focusing attention on security in the early stages of the life cycle of systems ('shift left security' approach).
2. *Security assessments/evaluations.* These evaluations examine security in-depth by performing source code, hardware, side channel analyses, etc. AVR can contribute to further automating these evaluations, performing additional types of analysis, thereby making the process more cost effective.
3. *Patching systems without support*. No patches are released for many IT systems, either because the support period has expired or (in the case of many types of open source software, for example) because no party is responsible for releasing patches. This applies to both 'old systems' and new applications, such as open source applications or applications developed in-house. Automated patch generation, one of the AVR research areas, can offer a solution.
4. *Risk-based patching.* In some environments, implementing patching poses a risk. A patch that has unintended side effects can compromise the stability of such environments. AVR can be used to analyse the effects or side effects of a patch so that, based on a substantiated risk assessment, a decision can be taken on whether or not to implement the patch.

[3] **See Chapter 5 for further details of needs and application areas.**

AVR Roadmap

The AVR Roadmap features four tracks.

1. *Education and training*. The education and training track focuses mainly on expanding curricula and training, and adapting a basic cyber reasoning system (CRS) so that it can be used for educational purposes, and possibly for a cyber challenge.
2. *Research.* The aim of the research track is to strengthen fundamental AVR. Three doctoral researchers are planned to be appointed in the research track in 2021. The intention is to decide on the scope of doctoral research projects in such a way that academic relevance ties in as closely possible with the issues to be addressed in bringing the practical application of AVR closer. The research areas that have been chosen are Automated patch generation / Risk-based patching and Automated vulnerability discovery. It is important that the research institutes concerned jointly define a further research agenda.
3. *Application of AVR*. The application of fundamental AVR knowledge through to validation/demonstration proceeds through three phases: firstly, exploring the available knowledge, technology and application, secondly, applying it to one or more use cases and thirdly, consolidating it in generic techniques or tools. The focus lies on specific application areas. Specific projects will be defined and launched in 2021 with the partners concerned. TNO will play a leading role.
4. *Field lab.* The field lab will support the other tracks. The field lab is primarily responsible for information transfer and cooperation between and within the tracks. It is a virtual space where experiments and data can be exchanged. Furthermore, the field lab will organise cooperation in all tracks and facilitate active communication and the creation of an AVR community.

The figure below summarises the Roadmap (status December 2020). The Roadmap will be updated annually.

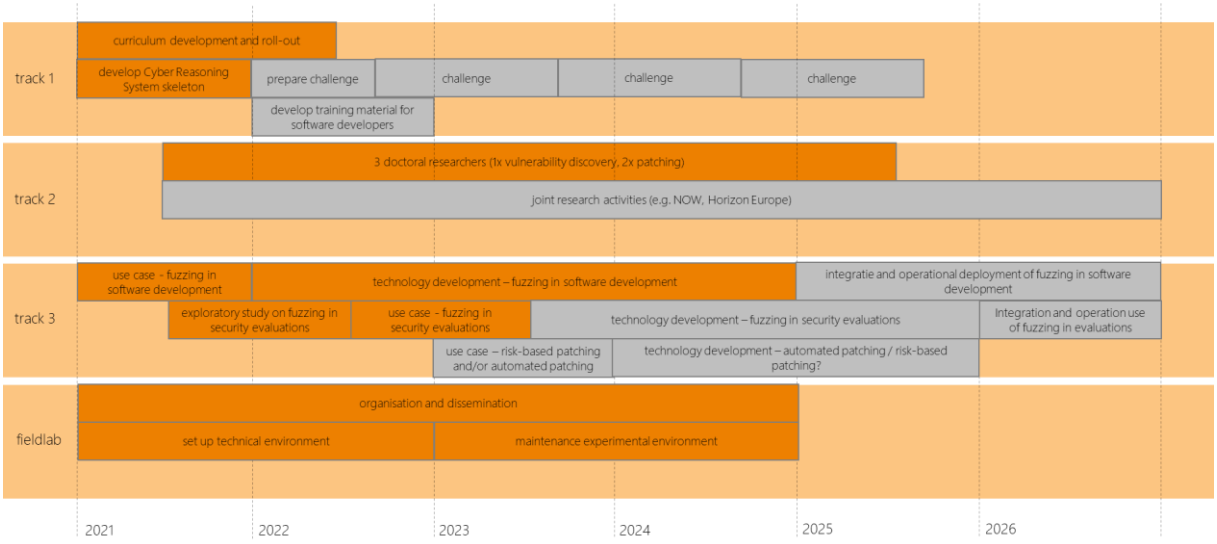

Organisation and funding

The execution and further development of the Roadmap can be regarded as a programme. The programme should be supervised and managed. We will therefore set up a steering committee and appoint programme managers and project leaders.

The steering committee will have a broad composition (similar to the board of the Cybersecurity Cooperation Platform), with representatives, for example, from the Ministries of Economic Affairs and Climate Policy, and Defence, the higher education sector, TNO as the most important applied knowledge institution, cybersecurity companies and end users.

The programme manager is primarily responsible for monitoring progress and achieving the results in the defined (and funded) streams and projects in the AVR Roadmap.

The programme manager will work together with the AVR Cooperation Platform. Essentially, 'innovation brokers' in that platform will be setting up programmes. The innovation brokers are responsible for ensuring the expansion of the Roadmap. They will drive and support the development of new research programmes and new collaborative projects, for instance.

It is crucial to build an AVR community. The Virtual Field Lab will play an important role in this regard.

Important decisions on matters of substance in the Roadmap will have to be taken at various times. It is important to take these decisions pragmatically, objectively and transparently. The publication of this Roadmap is the first step towards transparency.

To achieve the ambitions, €10 to 15 million is estimated to be needed for a six-year period. At present, an initial €2.0 million is available to work on achieving the goals of the Cybersecurity Roadmap. The application of the available funds is outlined in this Roadmap[4].

Far more money is needed to achieve the ambition set out in the AVR Roadmap. The correct use of the funds currently available, the active engagement of the entire chain, the development of an AVR community and a greater focus on AVR should result in more resources being made available: the flywheel will be set in motion.

See figure 13.

2 Introduction and context

On the initiative of the Ministries of Defence, Economic Affairs and Climate Change, and Justice and Security, work has been underway for some time on the topic of Automated Vulnerability Research (AVR).

The work is progressing gradually, primarily because relatively few people and organisations are currently engaged in AVR and, on top of that, they hardly have any contact with each other. Nor is it a focus area for potential AVR user organisations. This was reconfirmed after a series of interviews with various people and organisations before and shortly after the summer of 2020[5].

AVR, and in a broad sense, security automation, are regarded as a logical and necessary means of ensuring cybersecurity. Work in the field of cybersecurity is increasing, and there will never be sufficient qualified resources available to meet demand. The automation of components of the cybersecurity chain is therefore a logical development. AVR focuses primarily on the automated detection and identification of unknown vulnerabilities in software.

The ministries concerned would now like to proceed with the next step and develop a Roadmap for working on AVR. The involvement of the entire cybersecurity chain is crucial: from research and education to application (valorisation). Financial resources will be made available to begin developing the Roadmap. In order to fully implement the Roadmap additional funds will need to be found in the years ahead. The Roadmap will contribute to this. This Roadmap also provides an initial indication of how the available funds could be spent, and how the decision-making process could take place.

AVR Roadmap goals
The overall goals of the AVR Roadmap are as follows:

- Clarifying how to work on AVR in the Netherlands in the next six years.
- Coordinating the commencement of AVR activities based on three tracks to boost the development and application of AVR knowledge.
- Clarifying where and how partners in the chain can participate in taking AVR to the next level in the Netherlands;
- Inspiring partners in the chain to give new impetus to AVR.
- Building a vibrant community in the field of AVR.
- Providing insight into what can be achieved in the short and medium term, clarifying both the opportunities and limitations.
- Contributing to the digital autonomy of the Netherlands.

[5] The interviews focused primarily on setting up a cyber challenge in the field of AVR.

AVR anchored in the Security KIA and Cybersecurity Mission

The importance of AVR and security automation is widely acknowledged. The importance of the topic of AVR is anchored in the Security KIA and in the Cybersecurity Mission. See Figure 1 for a more detailed explanation

Knowledge and Innovation Agenda (KIA) for Security

The Security KIA is a further elaboration of the societal theme of security as described in the mission-driven top sectors and innovation policy. The missions set out in the KIA were drawn up by line ministries in consultation with various parties and top sectors, and then approved by the government. The elaboration of these missions has resulted in a number of Multi-Year Mission-Driven Innovation Programmes (MMIPs), and should lead to numerous forms of public-private partnerships and economic opportunities for larger and smaller companies. 'First and foremost', the KIAs provide an overview of the knowledge and innovation needs of businesses, ministries, the wider science community, knowledge institutions, the Netherlands Organisation for Scientific Research (NWO) and the regional authorities.[6]

Cybersecurity Mission

The Security KIA was drawn up by the top sectors concerned: High Tech Systems and Materials, Creative Industry, Logistics, Water & Maritime and the Dutch Digital Delta Team. The objective set in the Cybersecurity Mission is that the Netherlands should be able to "safely capitalise on the economic and social opportunities of digitalisation". Knowledge development and innovation in the field of cybersecurity are essential for preventing threats in the digital domain. The Cybersecurity Mission is aimed at strengthening cyber knowledge and skills in the Netherlands, facilitating research and innovation and building an ecosystem of experts and organisations. Active, effective and long-term cooperation across the entire chain is a prerequisite for strengthening the Dutch cybersecurity knowledge base.

Automated Vulnerability Research (AVR)

Despite increasing investments in cybersecurity, most organisations can barely keep pace with the speed and development of digital threats. The Security KIA states that the automation of cybersecurity activities is an important solution for remaining resilient with limited capacity. AVR occupies a central role under this banner, with knowledge and innovation questions centring on "developing methods for the automated detection of vulnerabilities in source code", as well as "automated patching at application level". The capacity to detect and patch vulnerabilities before malicious actors can exploit them is equally essential for digital sovereignty when other state actors are also using automated attack techniques. Furthermore, AVR plays a key role within the 'Offensive Technology' sub-programme. To achieve innovation in the field of AVR, intensive cooperation is required between stakeholders across the entire cybersecurity chain.

Figure 1: AVR anchored in the KIA and Cybersecurity Mission as a priority topic

Drawing up the agenda and programming

Developing a Roadmap ties in well with the thematic approach and the programming process as described in the Advisory Report from the coordinators of the Cybersecurity Cooperation Platform.[7] Figure 2 shows more details of the programming process as described in the advisory report.

[6] Reference KIA

[7] See: Cooperation platform cybersecurity. Advies kwartiermakers, 12 augustus 2020.

The thematic approach has two cycles: drawing up the agenda and programming

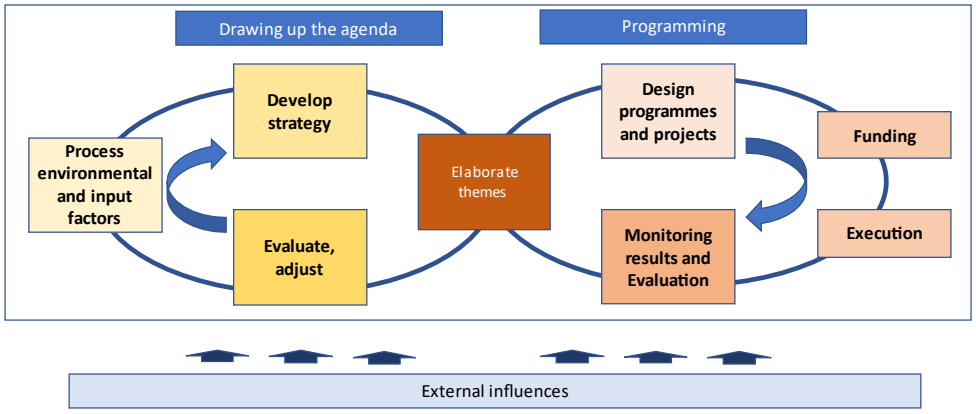

Programming involves elaborating the themes (from the agenda) in concrete programmes and projects. The designs should be sufficiently concrete so that they can be used as a basis for seeking funding. The format of the 'Roadmap' will usually be multi-annual. Different partners play a role in different phases. Supply and demand should come together.

The design of the programmes and projects requires specific knowledge and expertise. Knowledge of the theme will be provided by the various partners in the chain. Knowledge of partners: who are working on the topic, including internationally? Knowledge of the funding possibilities is also required, including national and international subsidy opportunities for knowledge and innovation.

Although a multi-annual roadmap can be developed, experience has shown that the real movement (realisation) always begins with smaller projects. It is crucial to take the first step to ensure success later. The role of the various government authorities will be essential here. But government authorities are certainly not the only initiators. Businesses can also start small. Experimentation should be encouraged.

Programmes will only get off the ground if there are genuine needs among one or more partners. The more specific the programmes are and the more partners involved, the more likely it is that programmes and projects will be funded, according to the theory, and there are plenty of practical examples to back it up. …

If funding is available, execution of the programmes and projects can begin, in phases. An evaluation will take place based on the results and follow-up phases will either go ahead or otherwise. The programming cycle therefore ends with monitoring and evaluation.

Figure 2: Drawing up the agenda and programming [8]

[8] See: Cooperation platform cybersecurity. Advies kwartiermakers, 12 augustus 2020, p. 18-19.

Draft Roadmap

This draft Roadmap was developed during the months of September–November 2020, based on the insights already obtained and with the support of various parties. And in the spirit of the coordinators' Advisory Report.[9]

This draft Roadmap is by no means set in stone:

- Firstly, because it has been developed within a short space of time and various iterations are anticipated in order to make final choices for initial projects.
- Secondly, and more importantly, the topic itself is developing rapidly. It would be utopian to look ahead six years: we will be overtaken by developments. We therefore propose evaluating and revising the Roadmap at least annually.
- And thirdly, because in addition to the tracks and initiatives described, more resources and initiatives will hopefully be provided to add to and successfully implement the Roadmap.

The Roadmap has been written in such a way that it should also be understandable for people with a non-technical background. The AVR Roadmap consists of the following chapters.

In Chapter 3, we briefly explain what AVR actually is. Why it is important and why it is such a difficult topic.

In Chapter 4, we zoom in on the results we aim to achieve. This is a six-year outlook (the 'dot on the horizon') and more concrete objectives for the shorter term.

Chapter 5 forms the core of the document and contains the actual AVR Roadmap.

Chapter 6 describes the organisation and funding.

Chapter 7 outlines the next steps.

In conclusion, Chapter 8 summarises the key risks and possible mitigation measures.

9 On 30 September, with representatives from the financial sector (banks) and on 6 October, with representatives from research universities, universities of applied sciences and the cybersecurity sector. Organisations and people were approached who were known to be working on the topic of AVR. This draft Roadmap will be rediscussed with these parties.

3 What is Automated Vulnerability Research?

3.1 Automation in cybersecurity

With the ongoing advancement of digitalisation and the consequent dependence on digital devices and the data processed, the importance of cybersecurity is also widely acknowledged. Cybersecurity will become ever more important with the roll-out of the Internet of Things, including in our critical infrastructures.

However, cybersecurity faces a number of challenges, such as:

- The wide range and large number of cyberattacks we need to defend ourselves against.
- The growing volume of information that must be processed swiftly in order to deal with attacks.
- The substantial scarcity of cybersecurity experts in relation to the growing need for them. The shortage of staff is only expected to increase in the years ahead.

An important potential solution to these challenges is to automate cybersecurity operations and functions currently performed by scarce experts, where possible.

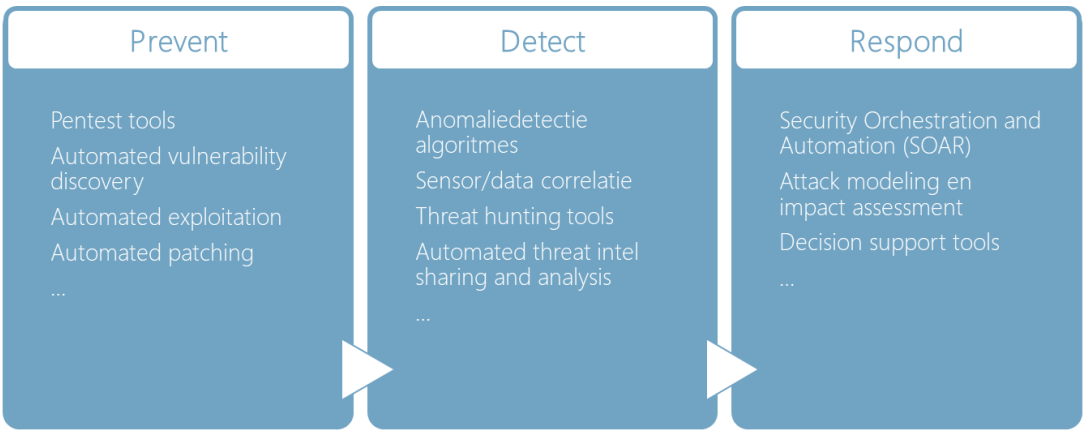

Figure 3: Automated vulnerability research positioning in the cybersecurity chain

Automation can be applied in all phases of the cybersecurity chain. The National Institute of Standards and Technology (NIST) cybersecurity framework distinguishes between the 'protect', 'detect' and 'respond' functions[10], see figure 3. Automation is already being applied to each of these functions, and the possibilities for further automation are being researched. An example is the Automated Security Operations (ASOP) consortium launched in September 2020[11] with support from the Ministry of Economic Affairs and Climate Policy. The consortium aims to develop a platform in the years ahead facilitating organisations with the automated detection of, and thus faster, response to cyberattacks. It focuses primarily on the detect and respond functions in the NIST model. ASOP is

[10] The 'identify' and 'recover' functions are also included in the framework, but serve more to support the other functions. The NIST framework is available on: https://www.nist.gov/system/files/documents/cyberframework/cybersecurity-framework-021214.pdf.

[11] https://cyberveilignederland.nl/consortium-automated-security-operations-asop-van-start/

an example of multi-year programming and a public-private partnership between knowledge organisations, the government and the cybersecurity industry.

However, automation can also play a pivotal role in the "prevent" step. The following sections explain what AVR is in greater detail and how it contributes to the prevention of cybersecurity incidents.

3.2 Automated vulnerability research

AVR is aimed at the semi or fully automated detection of vulnerabilities. This means that AVR forms part of the 'prevent' component of the cybersecurity chain. (see figure 3).

The 'prevent' function involves the prevention of cyberattacks. An important point of attention is to keep networks, computers and applications secure, among other things, by promptly identifying and revolving vulnerabilities. In current practice, patching (the systematic implementation of software updates offered by suppliers) and vulnerability scanning are used (searching for old, unsafe software versions on the organisation's own systems) and pentesting/red teaming (asking technical security experts to examine or even crack the security of the organisation's own systems). All of these activities centre almost entirely on the detection and remediation of *known* vulnerabilities: vulnerabilities known to the software supplier and/or the security community.

Of all software vulnerabilities, we are not familiar with the majority. This applies to virtually all types of software. A small group of cybersecurity researchers are engaged in researching or searching for *new, unknown* vulnerabilities in software. This discipline, *vulnerability research*, demands highly specialised skills that are scarce even within the wider cybersecurity field. A patch can be developed for a vulnerability once it has been detected and will then be included in the regular security patching processes.

Clearly, new zero-day vulnerabilities are of great value to cyber attackers. A zero-day vulnerability is practically impossible for users to protect against, as there are no patches available yet. For this reason, when a vulnerability is detected that it is actively being misused, it is essential to patch and disclose it as quickly as possible. This recently occurred with a vulnerability detected by Google's Project Zero[12].

The Defence Research Institute DARPA in the USA has been investing in AVR research for many years. AVR knowledge and capacities are therefore of a strategic nature and will not be shared between countries automatically.

---

[12] **https://arstechnica.com/information-technology/2020/10/googles-project-zero-discloses-windows-0day-thats-been-under-active-exploit/#:~:text=Google's%20Project%20Zero%20says%20that,almost%20two%20weeks%20from%20now.&text=Attackers%20were%20combining%20an%20exploit,recently%20fixed%20flaw%20in%20Chrome**

By building up AVR capacity in a timely manner, the Netherlands can continue to protect its own critical IT systems, including against attackers who are capable of exposing our country's vulnerabilities in an automated and large-scale manner.

3.3 AVR process steps

AVR encompasses various research areas, which although stand-alone have areas of overlap and mutual dependencies. These research areas are jointly working towards a process flow for the semi-automated analysis of software targets. An overview of the process flow is shown in figure 4.

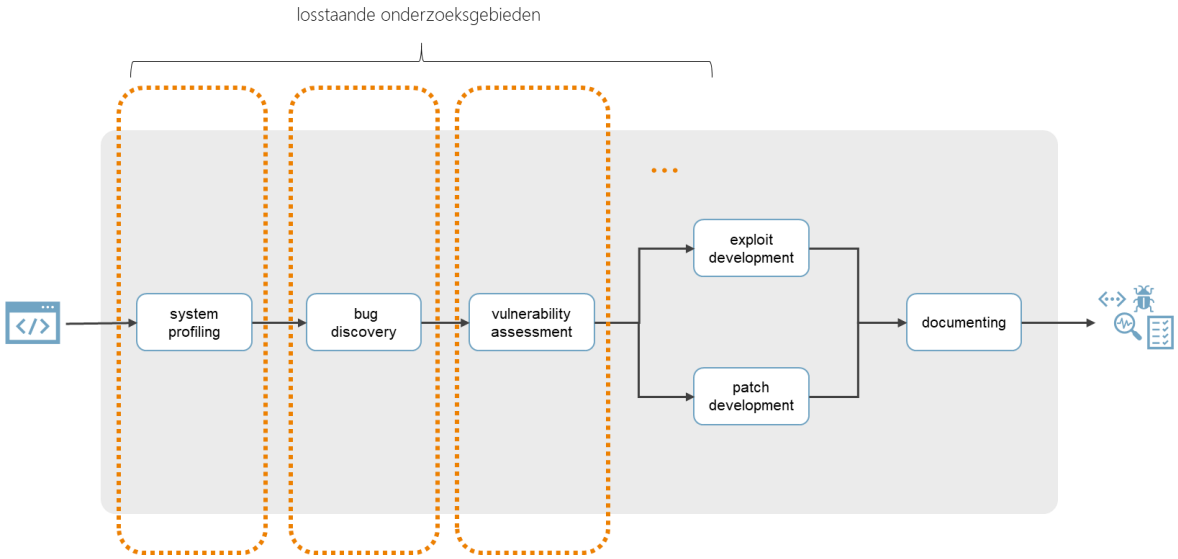

Figure 4: The various process steps within AVR.

The process steps to be completed are:

- analysing the software target or the available information on it (profiling);
- identifying design or programming errors that may have introduced a vulnerability (bug discovery);
- assessing whether a vulnerability actually exists, the type of vulnerability, and how it can be exploited (vulnerability assessment);
- developing an exploit for the vulnerability detected (exploit development);
- developing a patch that eliminates the vulnerability (patch development).

Research into automated bug discovery has undergone rapid development in recent years. A large number of open source tools are available for source code analysis[13] and fuzzing[14]. The effective use

[13] https://owasp.org/www-community/Source_Code_Analysis_Tools
[14] https://arxiv.org/pdf/1812.00140.pdf

of these tools and techniques still poses various research challenges. The automation of later process steps is still in its infancy.

4 AVR Roadmap ambition

AVR strengthens the preventive component of the cybersecurity chain by automating the detection of known and unknown vulnerabilities in software, developing exploits and patching vulnerabilities. A high ambition has been set for the AVR Roadmap. A 'dot on the horizon'.

The ambition is to achieve the following in six years' time:

- The Netherlands has a *leading position* in Automated Vulnerability Research (AVR) in Europe. This is demonstrated by the combination of publications (thought leadership), PhDs, the number of students, patents (or open sourced technologies), the volume of commercial services, among other factors.
- *Technology has been developed* for the successful automated detection of vulnerabilities on a commercial basis.
    - In at least one substantial IT development or management activity (TRL 8+).
    - A good starting position has been achieved for the valorisation and industrialisation (TRL6+ demonstrators) of automatic patching.
- Sector-specific applications have been developed, such as the capacity for the automated detection and/or patching of vulnerabilities in software used by banks.

The underlying policy objectives are to promote digital security, including for small and medium-sized enterprises (SMEs), and to strengthen economic earning capacity and strategic autonomy in the digital domain.

To achieve the ambition, more people, more expertise and more services are required.
The process objectives are follows:

- The entire valorisation chain is involved.
- Research and development are carried out in association with several sectors of society.
- Knowledge and technology will be adopted by at least two Dutch companies (preferably more) for further commercial development. In the unlikely event that no company wishes to act as the technology provider, the formation of one or more start-ups will be encouraged.

According to an initial estimate, €10-15 million will be needed to achieve the ambition. In implementing the Roadmap, it remains to be seen whether and how we can continue to put *the Netherlands* and Dutch interests first when it comes to a highly international topic (examples include export control). Funding for the Roadmap is discussed in greater detail in Chapter 7.

The goals of the AVR Roadmap are in line with the goals of the Cybersecurity Cooperation Platform and have also been derived from them. The goals of the Cybersecurity Cooperation Platform are shown in figure 5.

Cybersecurity Cooperation Platform
The platform has opted for a target and result-oriented approach based on a mission, vision and objectives. The approach chosen always demonstrably contributes to the achievement thereof.

- **Mission**: To contribute to a safer, smarter, digitally autonomous and economically stronger Netherlands.
- **Vision**: The Cybersecurity Cooperation Platform effectively brings together supply, demand and funding for cybersecurity education, research, innovation and application.
- **Objectives**:
    - There are a sufficient number of well-trained cybersecurity officers in the Netherlands; and
    - we generate leading international cybersecurity expertise in the Netherlands; and
    - cybersecurity expertise leads to effective application in Dutch products and services: there is valorisation.

Figure 5: Overview of the mission, vision and objectives of the Cybersecurity Cooperation Platform

The ambition can be implemented into more specific objectives, for example:

- *Cybersecurity officers*
    - AVR teaching materials are available for the higher education sector and for cybersecurity specialists working in the business sector. The teaching materials will at least be used in all relevant Master's programmes.
    - A technical environment is available where students can experiment with AVR, based on a Cyber Reasoning System.
    - An annual AVR challenge will be held in which teams from various research universities and universities of applied sciences will participate.
    - There is a constant flow of supply and demand of relevant work placements for cybersecurity students
    - In six years' time, the number of AVR specialists among Master's students, PhD candidates, permanent staff in higher education, and specialists working for cybersecurity service providers and end users has increased significantly.
- *Cybersecurity expertise*
    - The Netherlands has an active community of AVR experts who know each other and maintain regular contact. These experts are from the entire cybersecurity chain.
    - Every year, PhD candidates will embark on research projects in the field of AVR.
    - There are clearly defined joint research programmes in the field of AVR.
    - The research groups involved are renowned for their AVR expertise in Europe and in the world.
    - AVR expertise is reflected in both scientific publications and nationally developed tools and products in the field of AVR.
- *Cybersecurity application*
    - Academic vulnerability research leads to (a usable basis for) tools and products, which are known to knowledge institutions, cybersecurity service providers and end users.
    - Sector-specific applications. For banks, for example, and/or for cybersecurity service providers.

- There is a stream of projects to make the tools and products suitable for use. Various partners in the chain work on this together.
- Various tools will be deployed commercially by cybersecurity service providers and end-users.
- National parties will be responsible for maintenance, further development and support for tools and products.

5 The AVR Roadmap

5.1 The innovation chain

Figure 6 below provides an overview of the innovation chain, in which the foundation is laid through academic research and various phases are completed in order to apply this knowledge, resulting ultimately in deployment by end users in operational environments.

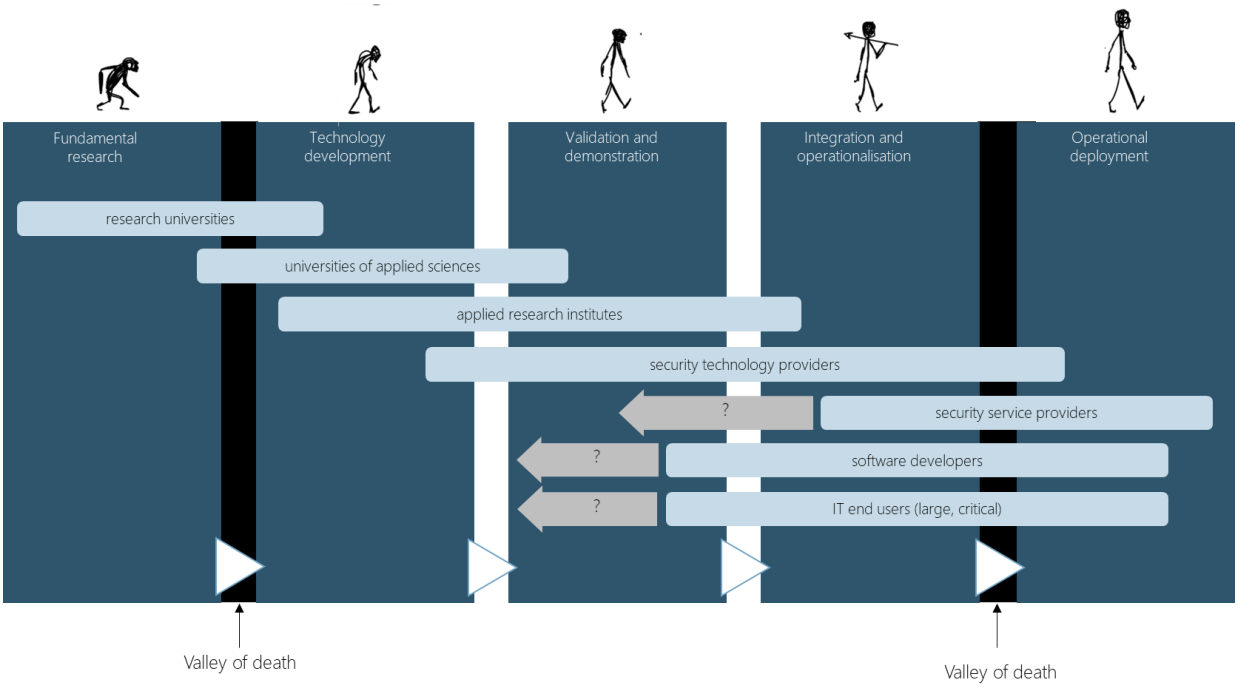

Figure 6: The AVR Roadmap in relation to the innovation chain

The figure also indicates the types of organisations that can play a role in AVR in the innovation chain. Known challenges are the "valleys of death":

- from basic research to technology development;
- from demonstration to operationalisation.

The first valley of death can be bridged by close cooperation between universities and applied research institutes. In addition, it is important that the end users of AVR technology are involved at an early stage to bypass the second valley of death.

In AVR we can distinguish between two types of end users:

- Organisations that want to protect their own systems and software. This includes not only IT end users who rely on IT and data for their business operations, but also software developers who want to monitor the quality and security of their products.
- Cybersecurity companies serving the actual end users, in which a distinction can be made companies between engaged in the development of security technology and companies that provide security services.

Clearly, these are roles and there are organisations that combine several roles.

The AVR Roadmap initially focuses on Dutch organisations that can play a role in the development of AVR. Parties from abroad may also be involved, for example, if they supply security products that are widely used by Dutch organisations. Examples are antivirus software suppliers and the developers of operating system kernels.

5.2 What is the current status?

First and foremost, experienced vulnerability researchers are required to be able to develop technology for automated vulnerability research. Experts who know the techniques, understand how automation can play a role and can contribute to their development. In the Netherlands, very few experts are actually engaged in vulnerability research. Only a handful of organisations are engaged in vulnerability research on a more or less structural basis.

At Dutch universities, AVR is recognised as an important field of research. Vulnerabilities exist in countless forms, but when the focus lies on vulnerabilities in software/firmware binaries (as in the DARPA Cyber Grand Challenge), in-depth knowledge of software and systems is a prerequisite. The cybersecurity group at VU Amsterdam is specialised in this area and occupies the strongest knowledge position among the Dutch universities.
Other research universities and universities of applied sciences have also included software vulnerabilities and their exploitation in their curricula. Research that is more or less closely related to AVR is also being conducted at various research universities. Nevertheless, the total volume of academic vulnerability research in the Netherlands is small, especially when compared with other important countries, and a substantial additional stimulus is required to ensure that the knowledge gap does not grow.

Besides the educational institutions, a few years ago, TNO's cybersecurity research group conducted automated vulnerability research. TNO's research focuses on applying the current AVR technology and creating the conditions to further develop it for practical application.

In the series of interviews referred to earlier, which were conducted among Dutch cybersecurity companies this year, the majority stated that they are not currently engaged in vulnerability research. They focus mainly on known vulnerabilities, the remediation of which has already proven to be a major challenge. And as long as regular patching is not in order, it has the highest priority. However, these companies regard AVR as an important research topic. A number of the organisations also stated that they may add AVR technology, upon maturity, to their service portfolio.

In conclusion, there are numerous hackers in the Netherlands, some of whom are extremely capable, who are engaged in cybersecurity either professionally or in their spare time. They constitute a valuable pool of knowledge. Talented hackers who would be capable of conducting vulnerability

research are currently earning their money in practice by performing other cybersecurity activities because that is where the demand is in the current market. Some driven vulnerability researchers earn 'bug bounties' by reporting vulnerabilities to the software supplier. The most well-known platform is HackerOne[15], an online platform where software suppliers offer their products to the hacker community to detect vulnerabilities in return for a fee. HackerOne's headquarters are based in Silicon Valley and the company is worth almost USD1.0 billion. HackerOne was founded by two hackers from Groningen in 2012 out of frustration that they were not allowed to report the software vulnerabilities they detected themselves.

5.3 Stakeholders

In 2020, a survey of parties who have an interest in and/or are active in the field of AVR was carried out on behalf of the Ministries of Economic Affairs and Climate Policy, Defence, and Justice and Security. These parties and their potential roles are shown in figure 7 below. The parties who are already active in the field of AVR research/application, and those are interested in becoming active in the field of AVR or have concrete plans in this regard are also shown.

		Fundamental research	Technology development	Demonstration and validation	Integration and	Operational deployment
Research						
	VU Amsterdam	√	√			
	University of Twente	√				
	TU Delft	√				
	Radboud University	√				
	University of Groningen	√				
	Open University of the Netherlands	√				
	Amsterdam University of Applied Sciences		√			
	Leiden University of Applied Sciences		√			
	TNO		√	√	√	
Cybersecurity companies						

[15] https://www.hackerone.com/company

	KPN			√	√	√
	Deloitte			√	√	√
	CapGemini			√	√	
	Riscure		√	√	√	√
	Secura				√	√
	QBit				√	√
	Cybersprint				√	√
	Zerocopter				√	√
	Tesorion				√	√
	SIG and CIP government				√	√
	NBV (Ministry of Foreign Affairs)				√	√
IT end users and software developers						
	ING				√	√
	Rabobank				√	√
	ABN AMRO Bank				√	√
	Ministry of Defence			√	√	√

Green: Already active in AVR

Orange: Has shown interest in AVR

Figure 7: AVR interest and activities among organisations (status 2020).

There are more organisations that could be involved in the years ahead. This table only shows the organisations that have been contacted.

5.4 What are the needs?

During the months of September and October, sessions were held with a number of the above stakeholders to discuss their interest in, and the potential applications of AVR. During the various discussions, specific applications were discussed along with the specific needs and contributions the parties can make. Both are discussed in the following sections.

5.4.1 Applications

During the discussions, the following applications were identified as the most valuable:

Application 1: Secure software development

By applying AVR to the software development process, vulnerabilities can be resolved at an early stage, before the software is deployed. This is in line with the movement towards focusing attention on security in the early stages of the life cycle of systems ('shift left security' approach).

Large software developers in particular already use static and dynamic security testing, and it is also applied to popular open source libraries. GitLab, one of the leading providers of CI/CD software development tool chains, has recently bought two security companies that specialise in fuzzing, with the aim of integrating it into their portfolio.

Application 2: Security assessments/evaluations
A few companies in the Netherlands specialise in carrying out security evaluations of IT components, the security of which is crucial. This is because they are used in classified networks, for instance, or because they form the foundation on which other security measures are built. These evaluations examine security in-depth by performing source code, hardware, side channel analyses, etc.
AVR can contribute to further automating these evaluations, performing additional types of analysis, thereby making the process more cost effective.

Application 3: Patching systems without support
Patching is a regular process in which suppliers make software updates available to their customers. Vulnerabilities can consequently be resolved shortly after they have detected or reported. However, no patches are released for many IT systems, either because the support period has expired or (in the case of many types of open source software, for example) because no party is responsible for releasing patches. This applies to both 'old systems' and new applications, such as open source applications or applications developed in-house.
Automated patch generation, one of the AVR research areas, could offer a solution. The financial sector has shown considerable interest in this topic.

Application 4: Risk-based patching
Implementing patching poses a risk in some environments. One example is the operational technology widely used in industry and in critical infrastructures. A patch that has unintended side effects can compromise the stability of such environments.
On the one hand, AVR can be used to analyse the effect of a patch. On the other hand, AVR can be used assess the vulnerability the patch is intended for.
A substantiated risk assessment can therefore be made to determine whether or not to implement the patch.

The first two applications use the steps at the front-end of the AVR process: the automated analysis of a software component to detect whether it has any errors/bugs. The last two applications focus on the back-end of the AVR process: automating the patching process after the vulnerability has already been detected. This is visualised in figure 8:

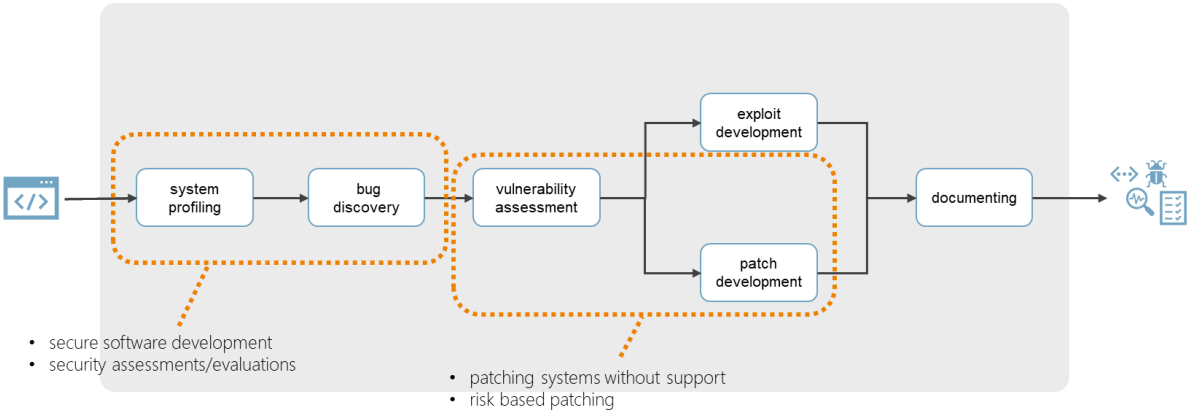

Figure 8: Potential applications in relation to the AVR process steps

Another difference is that many research results and tools are already available that focus on vulnerability detection. Although a great deal of research is still required to improve the effectiveness of these techniques and tools, it is already possible to achieve results based on the current state of the art. The automation of patching is far less advanced in the innovation chain. Fundamental research is primarily required to study methods and feasibility.

5.4.2 Needs matrix

During the discussions with the various parties, all parties indicated their specific needs and the specific contributions they can make. The input received is summarised in the needs matrix in figure 9.

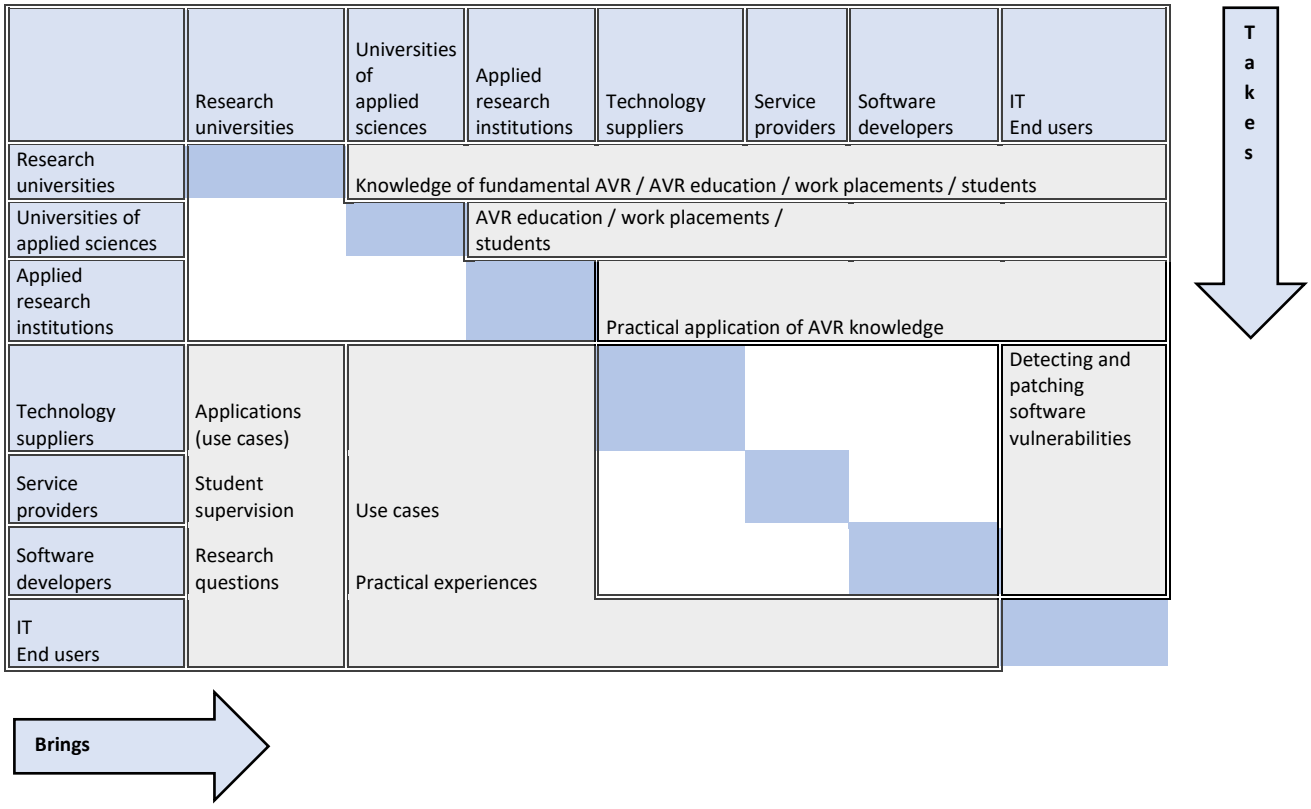

Figure 9: AVR needs matrix

5.5 The Roadmap

The aim of the Roadmap is to strengthen and expand national capacities in the field of AVR and to reduce or even eliminate the gap between the Netherlands and other countries (see Chapter 4 for further details). This requires a multi-year investment in knowledge and technology development.

The Roadmap focuses on the period 2021 to 2026 because a substantial impetus can be given and concrete results can be achieved within a period of six years.

Three tracks have now been defined within the Roadmap:

1. Education and training
2. Strengthening fundamental research
3. Application

A 'field lab' is planned to support these tracks and the cross-fertilisation between them. Figure 10 shows the positioning of the three tracks and the field lab in the innovation chain.

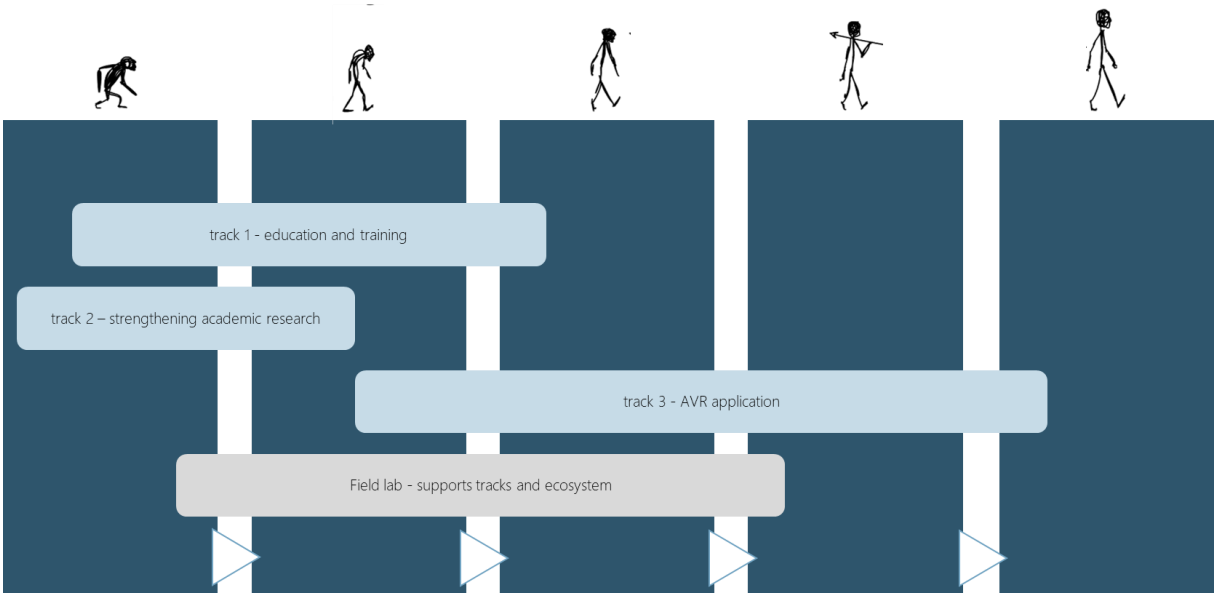

Figure 10: the tracks of the AVR Roadmap across the valorisation chain

5.5.1 Track 1: Education and training

The education and training track focuses mainly on:

- Expanding curricula and training.
- Adapting a basic cyber reasoning system (CRS) so that it can be used for educational purposes, and possibly for a cyber challenge.

Expansion of curricula

To strengthen and broaden education on the subject of AVR, a teaching package will be compiled that can be used within the various cybersecurity programmes. Material is already available at various research universities and universities of applied sciences. There is no need to start from scratch. The aim is to use the teaching materials already being used at the various research universities and to integrate, consolidate and make these materials more widely available. It is important that the resulting module can be applied to different levels so that it will be used by both

research universities and universities of applied sciences. A further aim is to make AVR a more attractive subject, so that it will appeal to more cybersecurity students. The Cyber Reasoning System will also contribute to this (see the next section). Organising a challenge in the future is also expected to have a positive impact on students.
Finally, training material will be developed focusing on the application of AVR in businesses. This will contribute to the integration and acceptance of AVR technology among end users. Businesses have stated that they need well-trained people in this field.
In terms of working method, it would be logical to start with the teaching material for research-oriented higher education and to derive the material for higher professional education and businesses from this. It is also important to ensure that the programme is not only theoretical but also provides opportunities for acquiring hands-on experience (see also CRS development).

The limiting factor in developing teaching packages is not so much money but rather the time needed from the scarce experts. A pragmatic way of combining knowledge and materials must be found. It is also recommended to compile modular teaching packages, focusing on the different AVR process steps. Finally, attention should also be paid to keeping the teaching material up to date: the AVR world is changing rapidly.

Execution:

- VU Amsterdam, TU Delft, University of Twente (Master's curriculum)
- Universities of applied sciences
- TNO (training material for businesses)

Start: 2021

Cyber reasoning system
TNO built a first version of a cyber reasoning system (CRS) in the above track. The CRS consists of modules that collectively cover all the steps involved in detecting, exploiting and patching vulnerabilities. The CRS contains a lot of "basic technology".

In Track 1, the CRS will be further adapted for use in the following applications:

- Curriculum support. Students can work with and on the CRS during the programme to gain a better understanding of the different steps required in conducting vulnerability research.
- Final projects. Final projects can be defined in which students develop improvements to the CRS or develop new modules for it. This should preferably be carried out in consultation with the cybersecurity companies.
- Cyber challenge. A cyber challenge can be organised annually. In this cyber challenge involving a competition, students can work on new modules for the CRS.
- Data sets. There is an urgent need for suitable data sets that students can work with. The possibility of making these available, for example as a component of the CRS, should be examined.

Execution:

- TNO, VU Amsterdam

5.5.2 Track 2: Strengthening fundamental research

The aim of Track 2 is to strengthen fundamental vulnerability research. The knowledge position of VU Amsterdam as a leading research group will be further strengthened for this purpose. A further aim is to raise the profile of the subject in other cybersecurity research groups and consequently broaden the basis.

In this track the research that is already being conducted should be more clearly identified, as knowledge about this is still fragmented. This concerns not only the research being conducted at research universities and universities of applied sciences. The Ministry of Defence, TNO and possibly other government institutions are also conducting research in the field of AVR.
The aim is to gain more insight and develop a better overview before jointly defining new research programmes.

Three doctoral researchers are planned to be appointed in the research track in 2021. The intention is to decide on the scope of doctoral research projects in such a way that academic relevance ties in as closely possible with the issues to be addressed in bringing the practical application of AVR closer. The doctoral researchers are therefore expected to actively contribute to the field lab, and to transfer the interim results to the other tracks, where possible.

Track 2 focuses on the following research areas:

- Automated patch generation/Risk-based patching
- Automated vulnerability discovery

As stated in 5.4.1 the first areas referred to are still entirely in the fundamental research phase. However, automated vulnerability discovery also poses fundamental research challenges, and it is possible to link this theme closely with the projects in Track 3.
The doctoral researchers will then focus on the themes that are still entirely in the fundamental research phase.

The proposal is to engage two doctoral researchers in 2021 to start working on the research themes of patching and one doctoral researcher on automated vulnerability discovery. To ensure the continuity of the research, it would be advisable to appoint doctoral researchers on a regular basis (on average at least one each year) from 2022 onwards.

More importantly, further choices for new research incentives for specific sub-themes should be made jointly. Research proposals should also be formulated and funding acquired, for example, in the context of a call under the Dutch Research Agenda or European research programmes.

Automated/Risk-based patching

Under the Automated patch generation/Risk-based patching track, fundamental research will be conducted on the automated development of patches for vulnerabilities detected, the automated assessment of the severity of a vulnerability, and analysing the effects of a patch (patch qualification).
This research can be used to prioritise vulnerabilities according to risk, provide decision support on whether or not to patch, and to reduce the costs and the lead time for developing patches.

The interim research results are anticipated to be used in the course of 2023 in Track 3, possibly for a first use case.

Start: in 2021, two doctoral researchers

Automated vulnerability discovery
Automated vulnerability discovery focuses on the detection of vulnerabilities not yet known. This research builds on earlier research at VU Amsterdam. The aim of this track will be to explore how to detect vulnerabilities as early as possible in the software development process by, for example, combining fuzzing with other analysis techniques.

The knowledge acquired will be important in 2021 for Track 3, in which work will start on applying automated vulnerability discovery in the software development process, based on the techniques and tools available at that time.

Start: 2021, one doctoral researcher

5.5.3 Track 3: Application of AVR

The application of fundamental AVR knowledge through to validation/demonstration proceeds through three phases:

- Exploring the available knowledge, technology and application
- Applying this in one or more use cases
- Consolidating in generic techniques or tools

These phases form an hourglass model. In the exploration phase, knowledge and experimental experience will be acquired of a wide range of current AVR techniques and tools, for a specific application (e.g. software development). This will provide insight into possibilities and limitations, configuration options, combination possibilities, etc. In the second phase, a specific use case will be chosen (e.g. the identification of vulnerabilities in a specific software target). Only the knowledge, techniques, configurations, etc. that can contribute to the use case will be combined and, use-case specific adjustments will be made, where necessary. After one or more use cases have been successfully completed, the techniques applied can be consolidated and further developed into tools that can be deployed more generically. This process is shown in Figure 11.

In phases 2 and 3, input from, and testing by, potential end users is necessary. They can submit relevant use cases as well as preconditions for integrating technology into their business process. In a

sense, the use cases also form the basis for Track 2 (research). Essentially, the research should, if possible, focus on applications that are interesting for the professional field in order to increase the chances of valorisation. The relationship between Tracks 2 and 3 will have to evolve over time. Businesses are primarily interested if a clear return on investment can be achieved. Projects must be sufficiently innovative and/or have potential. The focus should primarily be on subjects and technologies that businesses are unable to develop or have difficulty in developing themselves. Or those that are not yet available off the shelf.
Businesses may also want to participate because of knowledge development. And to be recognised as a business at the forefront of innovation and research.

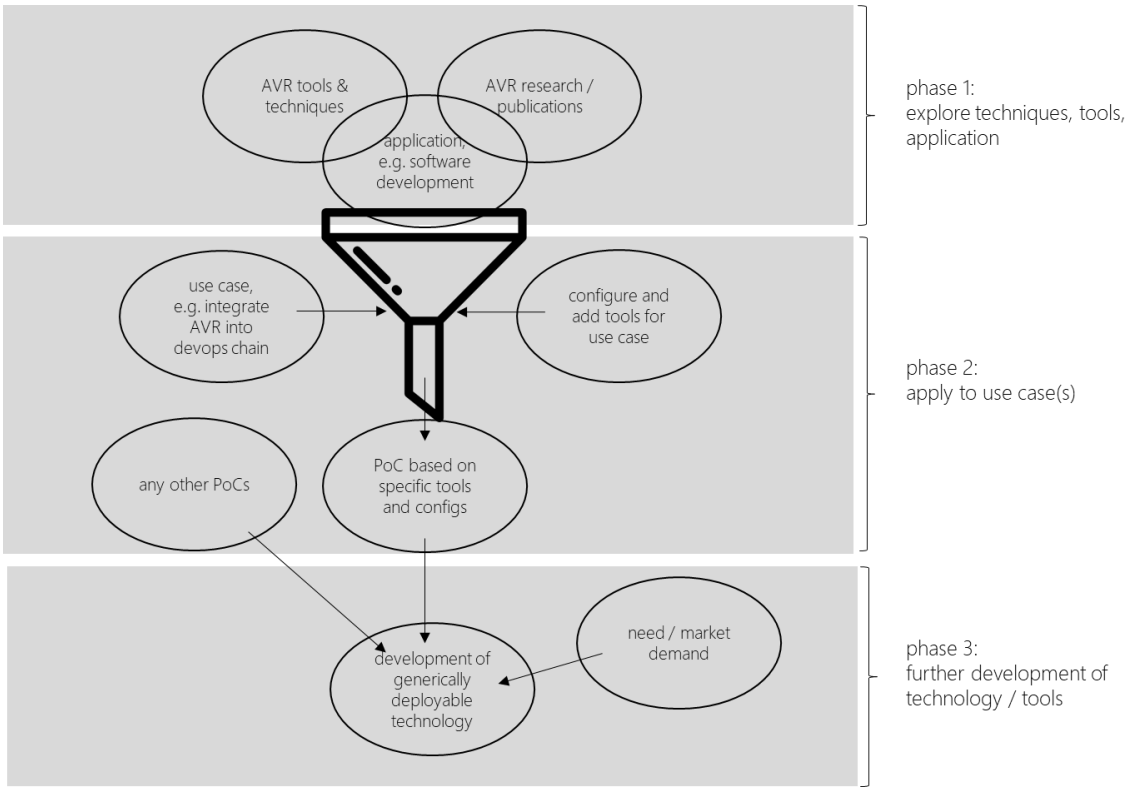

Figure 11: Applied research from exploration to technology development

The AVR applications described earlier do not have the same maturity: a PoC based on the current state of the art is possible for AVR in secure software development. Usable tooling for automated patch generation is not yet available.

The phases referred to can therefore be completed for all AVR applications, but phased over time:

- (2021) Fuzzing in software development, use case
- (in the course of 2021) Fuzzing in software assessments, start of exploration
- (2023/24) Automated patch generation and/or Risk-based patching, depending on the results of Track 2

Fuzzing in software development

A software development line is used for software development. Varying levels of automation are applied. In its most advanced version, every developer's commit is automatically stored, compiled and tested on a number of predefined tests, on different (predefined) platforms. The optimisation of the development process (DevOps, continuous integration) is receiving considerable attention in the market and is supported by software packages such as Gitlab. Security testing is also receiving increasing attention, for example, through the application of source code analysers. Gitlab has recently taken tentative steps to incorporate fuzzing tests into its development line[16].

In Track 3, the available fuzzing tools will be combined with the latest results of academic vulnerability research. A representative software development line will be set up to acquire experimental experience and to carry out use cases. In both designing the development line and in formulating of use cases, cooperation will be sought with Roadmap partners who are engaged in software development. In view of the maturity of this application, a use case is feasible in the short term.

If the results of the use case are positive, a follow-up phase is planned from 2022 onwards in which generically usable tooling will be developed. TNO and a security technology end user can work together on this. If a security technology developer is interested, the developer will also be involved.

Aim: to demonstrably improve the software development process as a result of detecting bugs prior to release (shift left security approach).

Implementation: TNO
Partners: software developers (e.g. financial sector)
Facilities: experimental development line in field lab
Start: 2021

Fuzzing in security assessments/evaluations
Assessing the security of software is an increasingly complex task. In the current situation, this can be performed on the basis of a source code (which is not always available) or on the basis of documentation and certification of the development process. Security assessments are performed by specialised parties, and on systems requiring a high level of assurance. During the round of interviews, Riscure (a commercial provider of security assessments) and the NBV (government agency charged with security evaluations for the Ministry of Defence, among others) showed interest in AVR as a possible addition to their toolbox.

In Track 3, the fundamental knowledge developed on fuzzing and automated vulnerability discovery will be applied within an environment in which software evaluation is a primary task. Close cooperation is desirable between TNO (contribution of AVR knowledge) and an organisation

[16] https://docs.gitlab.com/ee/user/application_security/coverage_fuzzing/

specialised in performing software security assessments. Because limited knowledge is available, research into this application will begin in the first phase (exploration of application possibilities).

Aim: to demonstrably improve the quality of software security assessments by applying new techniques to detect vulnerabilities.

Execution: TNO
Partners: Riscure, other security evaluation service providers
Facilities: IT facilities for remote cooperation between partners (field lab)
Start: in the course of 2021

Automated patch development/Risk-based patching
In 2021, Track 2 will start with fundamental research on patching. The first results of this research are expected to be published from 2023. The knowledge acquired will then be applied in Track 3.

The use case for this project, and the timelines, will largely depend on the results of Track 2. Within this project, a partner will be sought for whom automated patch generation is a valuable application, after which a use case can be carried out. In the needs assessment, parties from the financial sector showed interest in patching systems without support, including open source.

Aim: to demonstrably improve the patching process by making patches available for unsupported software, and decision support on whether or not to implement patches.

Execution: TNO
Partners: VU Amsterdam, financial sector
Start: 2023/2024

After academic AVR knowledge has been converted into technology validated in use cases, the final two phases in the valorisation chain follow: integration/operationalisation and actual application. The second valley of death lies between demonstration and operationalisation. In order to bridge this successfully, parties wishing to apply AVR technology must be closely involved in the earlier research phases. In this regard, timely agreements must be made/provided for on the following:

- Proprietary rights: a degree of exclusivity may form a precondition for businesses to invest in the technology. This may clash with the desire of universities or other parties to share knowledge and to publish results as open source.
- Maintenance: a party must be willing to maintain the technology and benefits must be provided in return. There is a strong preference for maintenance to be carried out by a Dutch party to prevent the acquired knowledge from disappearing abroad.

5.5.4 Field lab

The field lab will support the other tracks. The field lab is primarily responsible for information transfer and cooperation between and within the tracks. The field lab is not a physical space in which

experiments are carried out, but rather a virtual space in which experiments and data can be exchanged.

A connecting role is foreseen for the field lab, both in terms of cooperation between parties and in technical terms, for the integration of results. The field lab consists of the following components:

- Virtual/IT environment for storing codes and carrying out experiments.
- Organisation of cooperation.

Virtual/IT environment for storing codes and carrying out experiments.
To speed up the transfer between the various tracks, each track will make the results available in a shared environment. This is aimed at providing all parties involved access to these data and to enable them to repeat the research conducted themselves. An important component of the virtual environment is the CRS (see 5.5.1). The results of research, use cases and challenges can be integrated, shared and reused in the CRS.

This environment will not only be used for storing codes and data (e.g. using a Git environment), but also for performing experiments (e.g. by providing computing capacity for experiments). The environment can also be used to share all AVR teaching materials.

Parts of the field lab will be open for participants in work placements or challenges, including the skeleton CRS. It is important to make agreements on the confidentiality of data and knowledge and to ensure that it is safeguarded.

Execution: TNO, VU Amsterdam
Partners: Track 3 use case partners, other partners, challenge participants
Start: 2021

Organisation of cooperation
Meetings will be organised to promote cooperation between the tracks. Regular meetings are planned between the persons involved in the Roadmap. Plans and results, including partial results, will be shared during these meetings. The objective is to promote the transfer of knowledge. Particular attention will be paid to applying fundamental knowledge from Track 2 within concrete use cases in Track 3, and to feed practical experiences from Track 3 back to Track 2.
Various meetings will be organised:

- Technical meetings focusing on presenting and discussing the technical results. These meetings can be organised both by track and by topic.
- Meetings with all parties involved to create a community on the topic of AVR.

It is important to organise face-to-face meetings again coronavirus restrictions permitting, where people can actually work together. In addition, virtual collaboration environments can be created, for example, equipped with slow-messaging and other communication facilities.

Execution: programme management, TNO, VU Amsterdam
Partners: partners from all tracks
Start: 2021

5.5.5 Activities and phasing

The activities that will take place in the Roadmap tracks were explained in the preceding sections. The initial time schedule is shown in figure 12. The activities that will be carried out within the available budget from November 2020 are indicated in dark orange. The activities for which budget will have to be found in the period ahead (in good time) are shown in grey.

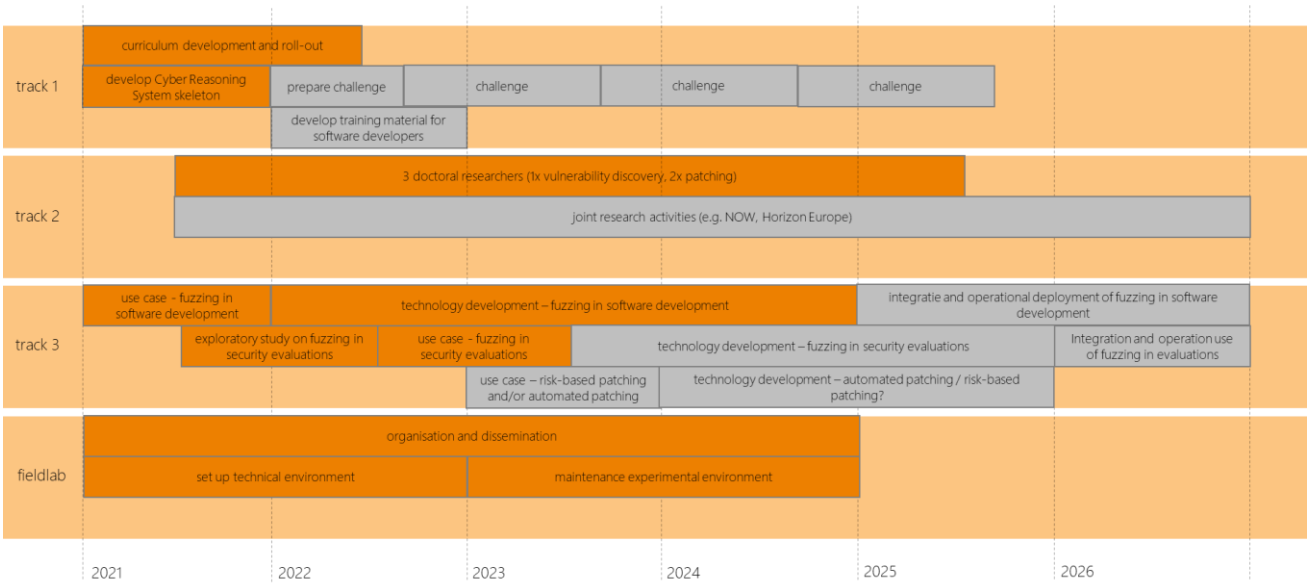

Figure 12: Activities and initial schedule

5.5.6 Internationalisation and links with other initiatives

In view of the complexity of the automated vulnerability research area, and the current national knowledge position, the desired capacities will not be able to achieved entirely independently. The Netherlands at least has an interest in seeking cooperation, within Europe for instance. VU Amsterdam and other universities within and outside Europe are already working together in the field of academic vulnerability research.

The Roadmap initially aims to create an ecosystem of Dutch parties. This will create a good starting position for later international cooperation and will also facilitate a quick start.

In the Netherlands, there are also other programmes in the field of automation in cybersecurity, such as the ASOC programme referred to earlier. It is important to keep a sharp focus on the distinctive position of the AVR Roadmap and to ensure that the same activities are not carried out in different

programmes and that the programmes tie in well with each other. The Cybersecurity Cooperation Platform will play an important role in this regard.

The programme manager responsible for implementing the AVR Roadmap will look into where it would be worthwhile seeking links with European/international partners and initiatives.

6 Organisation and funding

6.1 Organisation

The execution and further development of the Roadmap can be regarded as a programme. The programme should be supervised and managed. We propose setting up a steering committee and appointing programme managers and project leaders for this purpose.

If the development of the Roadmap proceeds successfully, various other collaborations and projects are likely to emerge over which neither the steering committee nor the programme manager can exercise control. We consider this a sign of success. However, it would be beneficial to maintain oversight of these collaborations and developments so that we can assess the extent to which the ambition and goals of the AVR Roadmap are being achieved.

These roles are explained in more detail below.

Steering committee

The AVR Roadmap steering committee:

- Is passionate about the topic of AVR.
- Has a direct interest in the success of the Roadmap. The steering committee members prioritise the overall goals that must achieved in the Roadmap.
- Monitors progress in work flows and the development of the Roadmap in general, based on the programme manager's reports.
- Promotes the topic of AVR within and outside the AVR community and consequently strengthens the prominence of AVR.
- Undertakes efforts to make available additional resources for the Roadmap.
- Provides direction and solicited and unsolicited advice to the programme manager.

The steering committee will convene at least four times a year. The following composition is proposed so that all partners in the chain are represented. Representatives from:

- Ministry of Economic Affairs (as the main financier and initiator)
- Ministry of Defence (as initiator and driver)
- Higher Education sector (proposed: VU Amsterdam Sec, as initiator and recognised leader in the field of AVR)
- Knowledge institutions (proposed TNO, as budget holder, initiator and leader in the field of AVR)
- Cybersecurity companies (two representatives are proposed: one person from a large company and one person from an SME)
- End users (a representative from a financial institution/bank is proposed)

The steering committee will elect a chair and ensure that new members are sought and found as necessary to remain effective, ensuring that all partners continue to be represented.

Programme manager

The programme manager:

- Is primarily responsible for monitoring progress and achieving results in the defined (and funded) streams and projects in the AVR Roadmap.
- Is directly responsible for the Virtual Field Lab.

- Acts as a ‘linchpin’, as the central contact for all parties involved and therefore regularly coordinates matters with them. Actively seeks and involves parties that can play a role in the Roadmap.
- Seeks opportunities and resources to continue implementing the Roadmap.
- Maintains contact with the Cybersecurity Cooperation Platform.
- Based on his/her role, reports to the Board.

The programme manager will fulfil the role part-time for the time being even though the above tasks can easily lead to full-time deployment. The programme will focus primarily on the activities for which budget is available (see figures 12 and 13). Projects will be defined for this purpose. The successes achieved in the programme will have to generate additional resources to expand and coordinate the role of programme manager.

The programme manager will work together with the AVR Cooperation Platform. Essentially, 'innovation brokers' in that platform will be setting up programmes. The innovation brokers are responsible for ensuring the expansion of the Roadmap. They will drive and support the development of new research programmes and new collaborative projects, for instance.

Project leader

Project leaders will be appointed for the activities in the streams, where necessary. We currently anticipate that the project leader role will be required for the following projects:

- Stream 1: Education and training
    - Development of teaching materials. Project leader from a university who is in charge developing these materials.
    - Adapting CRS for use in education (TNO).
- Stream 2: Strengthening research
    - No project manager required
- Stream 3:
    - Project management for applied research (from TNO).
    - Project management for each of the collaborative and co-financing projects to be defined later
- Field lab: TNO project management during the first phase of the programme. Further development under the responsibility of the programme manager.

The costs of project management are covered by the budget available for the project.

AVR community

The development of an AVR community is extremely important. The Virtual Field Lab will play an important role in this regard. The development of the AVR community falls under the responsibility of the programme manager. Several cybersecurity communities are expected to be formed through the Cybersecurity Cooperation Platform. The programme manager will maintain direct contact with the Cybersecurity Cooperation Platform in order to organise matters as efficiently and effectively as possible.

Decision-making

Important decisions on matters of substance in the Roadmap will have to be taken at various times. Example: Which university will appoint doctoral researchers? What topics will the doctoral researchers work on? Who will develop the training material? Etc. It is important to take these decisions pragmatically, objectively and transparently. A pragmatic process will be developed for this purpose.

6.2 Funding

To achieve the ambitions, €10 to 15 million is estimated to be needed for a six-year period. At present, an initial €2.0 million is available to work on achieving the goals of the Cybersecurity Roadmap. Further agreements will be made between the Ministry of Economic Affairs as the financier and TNO as the budget holder on rendering account of the funds.

In addition, a remaining budget of €200,000 is available from a project to develop a CRS and organise an AVR cyber challenge.

This is 'new funding'. Furthermore, various parties are already investing in the field of AVR (various universities, Ministry of Defence, a number of companies). The amount of these funds is unknown.

The preceding chapters show that numerous matters still need to be worked out in greater detail. Chapter 7 describes this process. However, the Roadmap will always need to be further elaborated and expanded. The Roadmap will be updated at least annually and additional projects can be launched at any time. Funding will always be a condition.

The application of the available funds is outlined in Figure 13. Further activities for which funding is not yet available have also been included. This overview serves as the basis for further elaboration.

Co-financing options will be examined in early 2021 based on more detailed project proposals. Co-financing may also be possible for start-ups that will or may be created in the field of AVR. This possibility will also examined.

Far more money is needed to achieve the ambition of the AVR Roadmap. The correct use of the resources currently available, the active engagement of the entire chain, the development of an AVR community and a greater focus on AVR should result in more resources being made available: the flywheel will be set in motion.

Track	Activity	Budget (€) available	Budget (€) required	2021	2022	2023	2024	2025	2026
Track 1: Education and training	Develop and roll out curriculum	€ 100.000							
	Adaptation skeleton CRS	€ 200.000							
	Teaching materials for software developers		TBD						
	Cyber challenge		TBD						
	Subtotal	€ 300.000							
Track 2: Strengthening research	Further analysis of current and desired research	€ 20.000							
	Appointment of 3 doctoral researchers	€ 730.000							
	AVR research programme, including doctoral researchers		TBD						
	Subtotal	€ 730.000							
Track 3: Application	Applied research	€ 650.000							
	Further projects		TBD						
	Co-financing projects	€ 250.000							
	Further projects		TBD						
	Subtotal	€ 900.000							
Track 4: Field lab and programme	Programme manager (part time)	€ 150.000	TBD						
	Set up and promote field lab (part time & out of pocket	€ 100.000	TBD						
	Subtotal	€ 250.000							
	Total	**€ 2.180.000**	**TBD**						

Funding available

Funding available from CRS/cyber challenge budget

Funding not yet available

Figure 13: Overview of activities and funding over time

7 Next steps for the Roadmap

This document is a second detailed draft of the AVR Roadmap. It is proposed to initiate actions in the months ahead, on the basis of which work on the first projects in the Roadmap can actually begin and the required funds can be made available. This includes appointing a part-time programme manager.

Figure 14 outlines the proposed next steps.

Track	Topics	Actions	When/who
Track 1: Education and training	Development of teaching package	Develop initial proposal for: Pragmatic exchange of available material Scope and depth of material to be developed	S. Verwer (TU Delft), E. Kok, Jan./Feb. 2021
	Adaptation of CRS	Develop project plan	TNO, Jan./Feb. 2021
Track 2: Strengthening research	Appointment of 3 doctoral researchers	1. Set out process, preconditions and criteria (feedback from all universities involved) 2. Submission of proposals 3. Assessment of proposals and selection	Optimistic time schedule: step 1: E. Kok, Jan. 2021 step 2: Feb. 2021 step 3: March 2021
	AVR Research programme	1. Explore possibilities and ideas in broad outline 2. Prepare and organise meeting with representatives from reserach universities/universities of applied sciences	E. Kok, TNO, ? Q1 2021
Track 3: Application	Applied research	Work out concrete projects/proposals. Individual work sessions TNO companies. Deepen areas of interest and objectives.	TNO, start Q1 2021
		Project calendar to be specified by TNO, as leader of this stream	TNO, Q1 2021
	Co-financing projects	Work out rules in further detail / co-financing possibilities. Agreements on IP and collaboration should also be taken into consideration.	TNO, Q1 2021
		Explore funding opportunities during individual work sessions.	TNO, start Q1 2021
Track 4: Field lab and programme	Set up field lab	Set up technical environment Coordinate with parties Create environment	TNO, start Q1 2021
		Organise the 'community' announcement, opening Draw up activity calendar	TNO, E. Kok start Q1 2021
Organisation	Steering committee	Set up/expand steering committee	Current steering committee, E. Kok, Jan. 2021
	Programme manager	Appoint programme manager	Steering committee, Jan. 2021

8 Realisation and risks

The key risks we currently anticipate and the mitigating measures that can be taken are shown in the table below.

Risk	Explanation and mitigating measures
Insufficient support for the Roadmap	Transparency is and will remain crucial: Explaining the Roadmap Involving all parties Opening the field lab Ongoing communication
Insufficient co-financing	Is not an immediate problem given the availability of resources. The Roadmap objectives cannot be achieved. The programme manager and steering group must make this a recurring agenda item and actively seek additional resources from the network.
The teaching package will not be used	The Dutch education model contains numerous rules regarding the use of teaching materials. The teaching package will be created in a modular format as far as possible so that it can be widely integrated. The community will always have to monitor use of the material.
Europe will overtake us	That might not be a bad thing. If more attention is paid to AVR internationally in Europe, this Roadmap can be quickly linked to these developments. The Netherlands already has an organisation in place. The parties involved know each other and have often worked together before.

9 Appendix 1. Overview of parties involved

The following people were involved in developing the Roadmap:

Steering committee

Ministry of Defence	LKol – Ilse Verdiesen
	LTZ1 TD – Jos van den Burg
Ministry of Economic Affairs and Climate Policy	Lars van Willigen
	Timon Domela Nieuwenhuis Nyegaard
TNO	Patrick de Graaf
Research universities	Herbert Bos (from December 2020)

Project team

TNO	Bert-Jan te Paske
	Gerben Broenink
	Dana Tiggelman
Cocus Consult	Edwin Kok (independent project manager)

Education sector

Research university / university of applied sciences	**Person**
VU Amsterdam	Herbert Bos
University of Groningen	Fatih Turkmen
University of Twente	Andreas Peter
University of Twente	Andrea Continella
TU Delft	Sicco Verwer
Radboud University Nijmegen	Erik Poll
Radboud University Nijmegen	Ileana Buhan
Open University of the Netherlands	Harald Franken
Hanze University of Applied Sciences, Groningen	Arne Padmos
Hanze University of Applied Sciences, Groningen	Ellen van Hegelsom

Business community (cybersecurity sector)

Company	**Representative**
Cyberveilig NL	Liesbeth Holterman
Tesorion	Rogier van Agt

Tesorion	Rick Hofstede
Tesorion	Daniel Jansen
Zerocopter	Edwin van Andel
Cybersprint	Vincent van Thiele
Fox-IT	Michelle Postma
Fox-IT	Dennis de Hoog
KPN	Edwin Bron
KPN	Oscar Koeroo
Capgemini	Geert van der Linden
Capgemini	Sanne Kuijpers
Philips	Maarten Bodlaender
Riscure	Marc Witteman
Riscure	Erwin in't Veld
Deloitte	Jelle Niemantsverdriet
Deloitte	Cas van Cooten
Deloitte	Kevin Jonkers
Deloitte	Daan van Moorsel
Secura	Ralph Moonen

Financial sector

Company	**Representative**
ING	Yehenew Gizaw
ING	Stefan Petrushevski
ABN AMRO Bank	Olaf Streutker
ABN AMRO Bank	Coen Klaver
ABN AMRO Bank	Alexander den Engelsman
ABN AMRO Bank	Robert van Lierop
ABN AMRO Bank (Clearing)	Jeroen Schilders
Rabobank	Johan Romkes

Other

Organisation	**Representative**
FME cybersecurity working group	Stijn Bouwhuis
National Vehicle and Driving Licence Registration Authority (RDW)	Marc de Bruin
Enexis	Mauriche Kroos
NLNCSA	Jan Verschuren