

# Nix as a declarative solution for embedded security challenges and system administration problems

UNIVERSITY OF TURKU  
Department of Computing  
Master of Science (Tech) Thesis  
Cyber Security  
January 2025  
Eino Korte

The originality of this thesis has been checked in accordance with the University of Turku quality assurance system using the Turnitin OriginalityCheck service.

UNIVERSITY OF TURKU  
Department of Computing

EINO KORTE: Nix as a declarative solution for embedded security challenges and system administration problems

Master of Science (Tech) Thesis, 57 p.  
Cyber Security  
January 2025

---

Embedded devices are an integral part of our daily lives; household machines, automobiles, and thermal sensors make use of embedded devices. They are subject to the global, developing worlds' security problems. This thesis focuses on those found on public information screens. Embedded devices are particularly vulnerable to security problems as they face issues in receiving constant, reliable updates. This thesis' focal point is maintaining, updating, and upgrading embedded devices. A proposed architecture of a public media screen system is provided with example program snippets to cover most common security issues found in similar setups. The architecture and its content are evaluated through the QuERIES methodology. The central theme of this thesis is NixOS, which is a Linux distribution that forms itself from a set of configuration files, supporting features like atomic rollbacks and reliable dependency handling. The most definitive academic sources in this particular subject are used extensively, as well as papers regarding both embedded security and measuring security.

A quantitative research methodology, QuERIES is used to measure the security of a novel architecture using NixOS. QuERIES contains a number of steps that evaluate the security of a system. The steps are iterated two times, each iteration providing a partially observable Markov decision process (POMDP) output, which is used as a benchmark of the overall security of the reference architecture.

The result of this thesis is that with the use of QuERIES, the overall security of the architecture can be improved methodically with the use of POMDP as a defined attack graph. An economic model of cost estimation to the attacker is gained via the red-blue team setup, which is then used as a tool for revealing the weak spots of the architecture from a chronological standpoint. The output of QuERIES can be generalized with tight constraints; as QuERIES provided tangible improvements in small scale, it could serve well a more complex setup. This is due to the nature of QuERIES in tandem with POMDP being able to handle a number of parameters, which is essential in a larger setting.

Propositions for further work are presented for studying the possibilities of purely functional and declarative solutions in the embedded field. The issues in applying QuERIES are also highlighted, and the development of more accessible tools for measuring cybersecurity is discussed.

Keywords: declarative, nix, nixos, pomdp, security, queries

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Research methodologies and questions . . . . .	3
1.1.1	Literature review . . . . .	4
1.1.2	Research questions . . . . .	5
1.1.3	Data collection and analysis . . . . .	6
<b>2</b>	<b>Declarative vs. imperative systems</b>	<b>7</b>
2.1	Imperative systems . . . . .	8
2.1.1	Debian/Apt . . . . .	9
2.2	Declarative systems . . . . .	10
2.2.1	Non-declarative components . . . . .	12
2.2.2	Home manager and flakes . . . . .	14
2.2.3	Ease of updates . . . . .	14
<b>3</b>	<b>Embedded system security</b>	<b>16</b>
3.1	Common embedded systems pitfalls . . . . .	18
3.2	Nix as an embedded solution . . . . .	19
3.3	Imperative and declarative systems from CIA-triad approach . . . . .	20
<b>4</b>	<b>Proposed architecture solution</b>	<b>23</b>
4.1	Server . . . . .	23

4.2	Client . . . . .	25
4.2.1	MQTT-client . . . . .	26
4.2.2	Weston/XWayland . . . . .	26
4.3	Test device . . . . .	29
4.4	Installing new devices . . . . .	31
<b>5</b>	<b>Security standpoints</b>	<b>33</b>
5.1	A brief history of security metrics . . . . .	33
5.2	Choosing security metrics . . . . .	34
5.3	Measuring security . . . . .	35
5.4	Methodologies in comparison . . . . .	37
5.5	Quantitative metrics . . . . .	38
5.5.1	QuERIES . . . . .	38
5.5.2	QuERIES as a central methodology . . . . .	40
5.5.3	Partially observable Markov decision process . . . . .	41
<b>6</b>	<b>Applying QuERIES</b>	<b>43</b>
6.1	Modeling the problem and quantifying the models . . . . .	44
6.2	Modeling the possible attacks . . . . .	44
6.3	Using the results . . . . .	47
6.3.1	Improving the system using Nix . . . . .	50
6.3.2	Issues with applied methodology . . . . .	51
6.4	Generalising the results . . . . .	52
<b>7</b>	<b>Further research</b>	<b>54</b>
7.1	Further research questions . . . . .	54
<b>8</b>	<b>Conclusion</b>	<b>56</b>
	<b>References</b>	<b>58</b>

# List of Figures

2.1	Terminal output from Emacs installation. . . . .	9
2.2	Relation graph in Nix environment. . . . .	13
3.1	The CIA-triad. . . . .	21
4.1	Working example of NixOS client device. . . . .	29
4.2	An architectural graph of the test setup . . . . .	30
5.1	QuERIES as a flowchart . . . . .	39
6.1	Probabilities to breach into the system. . . . .	48
6.2	Cost benefits of applied economic model. . . . .	49

# List of Tables

6.1	POMDP attack graph. . . . .	46
6.2	Protections applied in each POMDP iteration . . . . .	47

# 1 Introduction

A Linux distribution is a bundle of the Linux kernel and a set of software products called packages [1]. A package manager is an instrument that handles building packages either from source or pre-built binaries, resolving build-time and run-time dependencies of packages and installing, removing and upgrading packages in user environments [1]. Every Linux device must handle its installed programs with their dependencies and configurations either imperatively or declaratively. Overwhelmingly large portion of Linux distributions fall in the first category [2]. Both imperative and declarative distributions have their strengths and weaknesses from administrative and security standpoints.

An imperative package manager updates and modifies packages destructively. Apt, apk, dnf and zypper are some of the popular package managers [2]. Imperative package managers can remove and overwrite existing files which leaves the system in an inconsistent state. Different installs have different states by nature which causes many problems discussed in this thesis. As files are cross-modified through packages with package managers such as apt, upgrading can be disastrous as such distributions do not support atomic rollback capabilities<sup>1</sup>. Due to the unpredictability, the result can be a partially or completely broken system [2].

The reference imperative Linux distribution in this thesis is Debian with its default package manager apt, due to its popularity and relative simplicity. The

---

<sup>1</sup>Some Linux distributions using btrfs filesystem can perform a snapshot and rollback, meaning that a previous state of a system can be recovered without issues [3]

reference declarative distribution is NixOS with its partial namesake package manager Nix which provides declarative configuration of the whole system including the Linux kernel<sup>2</sup>. Nix is configured by the Nix programming language which is inspired by purely functional languages such as Haskell [5]. Declarative systems are those whose operation is based on a set of configuration files that define the system.

The reference architecture depicted in chapter 4 is based on Nix as it is the most popular purely functional and declarative Linux distribution with over 100 000 packages [6]. Architecture in this context refers to a set containing clients, server and the used hardware. The term system means the implementation of the architecture, especially referred in chapter 6. Depending on the context, system can also mean a set of procedures, e.g. a Linux system.

Another good distribution for the architecture would have been Guix, with over 28 000 packages [1]. Guix has some improvements over Nix, including richer and more extensible programming environment with a Lisp-dialect configuration language, Scheme [7]. NixOS remains as the distribution of choice, as the number of packages is greater and general support is found to be better.

This thesis discusses how declarative distributions can be used as an improvement over imperative distributions. In chapter 2 both approaches to package management are compared followed by an introduction to embedded security in chapter 3. Chapter 4 presents an example architecture created with NixOS, followed by chapter 5, which discusses the security metrics field in general and the decision process behind the chosen methodology for this thesis. In chapter 6 a quantitative research is carried out, revealing strengths and weaknesses of the reference NixOS environment setup. The selected methodology, QuERIES provides quantitative results that can be used to improve the security of declarative architectures. Propositions for further research are covered in chapter 7 and finally, chapter 8 concludes the thesis.

---

<sup>2</sup>There exists an experimental project that has succeeded with BSD interoperability [4].

## 1.1 Research methodologies and questions

Research methodologies used in this thesis are:

1. a literature review of central papers on subject themes found with prepared search statements
2. an action research using a laboratory setup
3. a quantitative research process using QuERIES methodology

Literature review will be addressed mostly in the next section and in chapter 3. As this thesis' main research methodology is quantitative, the gathered data points will be addressed as variables that are compared with mathematical methods. Quantitative methods are broken down in chapter 5 and 6. The central methodology is derived from QuERIES, and the information security aspect is investigated with CIA-triad. In section 5.4 QuERIES is compared with other metrics which are explained in section 1.1.2.

Quantitative methodologies are oftentimes used in conjunction with qualitative methodologies, with both approaches having their strengths and weaknesses. One drawback of using qualitative methods in security framework is their inherent subjectivity. For example the Delphi technique, where a set of opinions is gathered and compared from a working group, provides subjective substance for a study instead of objective perspectives [8].

The study design in this thesis is *state based*, which refers to the fact that the research methods focus on different state transitions, e.g. how probable it is for an intruder to gain from partial leverage to a full control of the system. Qualitative research would not alone satisfy the requirements, as investigating different state transitions without quantitative methodologies would be absurd [9].

### 1.1.1 Literature review

This thesis has bibliography from 49 sources, most of it gathered with a carefully prepared search statement. Other sources include manuals, material for research methods and other relevant material. The search statements results presented in next section, provide a quality base for action research and analysis.

The literature focuses on four main concepts: embedded systems security with and without declarative components, imperative systems, measuring security and Nix. The main goal is to find literature that combines these concepts to gain platform for comparing different approaches to support the action research.

Central literature revolves around Nix and multiple texts by Eelco Dolstra [2], [10]–[12] are cited for illustrating the nuances of a Nix ecosystem. Other declarative approaches are discussed by Endres et al. [13] and Van der Burg [14]. These approaches also contain comparison to imperative systems, which is the central viewpoint in chapter 2. Combining cyber security with declarative approaches are discussed by Specht et al. [15] and Kandoi and Artke [16].

Discussion from Ravi et al. [17] and Fysarakis et al. [18] handle embedded security extensively. The approaches itself are too broad for this thesis' scope, so only the most fitting were selected for use.

The most important literature to measure the security of a system is by Carin et al. [19] and Hughes and Cybenko [20]. The topics from these papers revolve around QuERIES, an original approach for measurably improve security.

#### Search statement

The main search statement for this thesis is: "embedded linux" OR "declarative" AND (linux OR \*nix) OR deployment OR "system update" OR (compare\* AND declarative AND imperative AND system\*) OR security.

The search statement was prepared to provide as relevant results as possible for

this thesis. The main goal was to include the hypernyms "embedded linux", "linux" with other terms separated using the "OR" operator. The subterm (compare\* AND declarative AND imperative AND system\*) was chosen to broaden the search to include articles which compare declarative and interactive systems.

As security is a central theme in this thesis, the term "security" was included. Systems security and cyber security materials are also included in the bibliography using search statement "systems security OR cyber security". Separate search "cia-triad" and "partial observable Markov chain" AND "cybersecurity" were used to provide tangible meters for measuring cyber security. To further back up the research for comparing different metrics, term "cyber security metric methodology" was searched.

For searching specific material about embedded systems, the search statement "embedded AND security" was used. All searches were done on Google Scholar platform.

### 1.1.2 Research questions

The research questions for this thesis are:

1. How can a declarative system be used to measurably improve the security of an embedded system used for displaying public media?
2. What are the advantages and/or disadvantages of such system from a system administrator standpoint?
3. How can a system running a different Linux distribution be securely and seamlessly migrated to NixOS?

Research question 1 is the most important and it traverses through themes of the whole thesis. The hypothesis is that traditional imperative embedded device fleets

have problems that can be solved with the use of modern declarative systems. First we aim to gain information from a specific scenario presented in chapter 4. Then in chapter 6 the gained information is analysed and generalised as suitably as possible.

Research question 2 brings up the human element; how can a system administrator use a new palette of features adequately to provide a more secure system. Research question 3 handles a situation where an existing system should be replaced with a NixOS system. How this could be done securely without risks and preferably easily with existing tooling, is answered in section 4.4. The next chapter compares imperative approaches to declarative approaches and provides insight for understanding the central differences, the emphasis being on how declarative systems can be used to solve problems better than with imperative systems.

### 1.1.3 Data collection and analysis

Data collection is done with simulated red–blue team setup, where both teams are provided by a time frame where they must conduct a series of tasks. These tasks are formalised as partially observable Markov chain parameters, and analysed with QuERIES methodology. This methodology is used to gain knowledge and make the system more reliant and better with multiple iterations. Chapter 6 answers research questions 1 and 2 and provides analysis for the reference system. Research question 3 is answered in chapter 4 section 4.4.

## 2 Declarative vs. imperative systems

Different approaches to declarative modeling of systems design are presented. Endres et al. [13] compare declarative and imperative standpoints in a cloud computing context and collect systematic information on what are the strengths and weaknesses of TOSCA, IBM Bluemix, Chef, Jujy, and OpenTOSCA. Van der Burg and Eelco Dolstra [14] use NixOS as a solution for declaratively distributing into cloud, executing integration and system tests. Most approaches researched through literature review focus on distributing to cloud. Distributing to embedded clearly remains a niche.

Breitenbücher et al. [21] focus on deploying into embedded and discuss the challenges an IoT user face when deploying a system. It is proven that setting up devices with mandatory scripts and other actions is a challenging task, when a number of devices should be set up [21]. Cloud is something that is practical to be used in tandem with IoT, but this thesis focuses on an *in-premises* reference solution.

In this chapter, we focus on comparing different declarative approaches to the more traditional imperative models, highlighting the strengths and weaknesses of both. Specifically, examples are provided to illustrate the limitations often observed in imperative systems, particularly in terms of reproducibility, scalability and administration standpoints. Cloud-oriented approaches serve as a key reference point for the effective distribution and automation of declarative systems. It can be ar-

gued that similiar approaches as those taken in cloud, should be taken in embedded and IoT to increase security through updatability and upgradability. This argument is backed up in section 6.3.1, where a concrete use case is demonstrated.

## 2.1 Imperative systems

Imperative deployment models base their functionalities through a process in which the order of events have a critical significance to the output [21]. In the context of virtualisation, imperative tooling can be used to form all activities to be executed: the control flow, execution order and the data flow between them [13]. This kind of process is best to be used in conjunction with a formalised workflow or standard such as BPEL [13]. In contrast, declarative models don't have such specific requirements, as these models formalise the processes in the configuration files [13].

Imperative systems have inherent problems regarding administrative traits contributing to a framework where the underlying system has **no traceability**: the implication that reproducibility is impossible, as changes to a system are not traced. Nix provides a solution for this problem with its Nix generation system [11]. With imperative systems, upgrading is more error-prone than installing from scratch. This is due to the fact that imperative systems have **unpredictable** state, from where the system should migrate to a predictable state. This causes major issues regarding upgradability.

**The inability to run multiple configurations side-by-side** is an inherent side effect of a *stateful* system [11]. Declarative systems don't have this problem: an arbitrary number of configurations can exist side by side, as the system is defined only by the configuration, not with the state as a component.

### 2.1.1 Debian/Apt

An example of imperative systems' problematic nature is provided with the following demonstration. Executing shell command

```
apt install emacs
```

installs a text editor wrapped as a .deb package. The package emacs has a dependency, emacs-gtk, which can be removed with command

```
apt remove emacs-gtk
```

Another dependency, emacs-lucid can be removed with command

```
apt remove emacs-lucid
```

we can see that after removing, apt automatically installs emacs-gtk to avoid breaking the application. The package manager warns: "emacs-lucid has dependency problems, but removing anyway as you requested" as shown in figure 2.1.1. It is also noteworthy, that the manual page for apt does not tell anything about a possible installation side-effect of a package removal command [22]. We could forcibly remove the package by invoking

```
dpkg --remove --force-depends emacs-lucid
```

leaving the system in an unreliable state. Dpkg is a low-level tool associated with apt, which does not automatically handle dependency resolutions or further package relations [23]. What happens if we had a large number of devices, in-premises or cloud, where all system commands are done imperatively? We would

```
dpkg: emacs-lucid: dependency problems, but removing anyway as you requested:
emacs depends on emacs-gtk (>= 1:27.1) | emacs-lucid (>= 1:27.1) | emacs-nox (>= 1:27.1); however:
 Package emacs-gtk is not installed.
 Package emacs-lucid is to be removed.
 Package emacs-nox is not installed.
```

Figure 2.1: Terminal output from a Debian system when installing an Emacs package.

have a large number of devices that differ from each other, because as shown, the order of commands affect the state of the system. Time is also a factor that causes systems to diverge as packages are not up-to-date by default. Invoking

```
apt update
```

updates the local repositories to match the download mirrors. If by technical reasons or possible user error the commands are conducted in the wrong order, this will lead to divergent systems.

Implementing a deployment model with only Debian would be a gruesome task as the order of events, which occur during the setup phase, is critical. As presented by Endres et al. [13] a formalised workflow graph would be needed to set up a reliable system. However, Debian could be used as a host to user-space application deployment such as Bluemix or Chef, where common DevOps practices can be used.

## 2.2 Declarative systems

Presented problems in section 2.1.1 can be solved using a reproducible, reliable and atomic package manager. In Nix, package installations are isolated from one another to prevent conflicts, ensuring they function consistently even if the underlying installations differ. As packages are declared in a single set of configuration files, it is trivial to reproduce the system in a different environment. The demonstrated effect in snippet 2.1.1 was a problem due to lack of isolation. When dependencies are scattered in the system instead of declared explicitly in an installed package, a faulty state could be achieved. Nix assures, that these kinds of problems are out of the question. A result of this is that in a Nix system installs of same program can reside side-by-side with varying versions [2].

As presented by Endres et al. [13] systems can be declared, even if the underlying

infrastructure is imperative by nature. This thesis focuses on purely functional methodologies which fix the most prevalent issues found in imperative models. Tools such as Chef focus on deploying on an imperative system, which causes an inherent problem with cohesion in a system that should work regardless of the underlying machine or network. Alternative deployment tools are discussed in section 2.2.1.

One benefit from Nix is its lightweight option to enable system tests. Integrating system tests with a Debian system would require a considerable amount of work, as setting up such system needs a lot of configuration and executing commands in a correct order [24]. Debian definitely fits in an imperative deployment strategy but the requirement for explicit detail of every step would be prone to errors even for a seasoned administrator [21].

It is also noteworthy that many imperative package managers do not support rollback mechanisms. If the Nix configuration file is changed and the system is rebuilt with command

```
nixos-rebuild switch
```

the previous state could be recovered by

```
nix profile rollback
```

This is an important feature since the Nix configuration files control the whole system, they can also leave the system in an undesired state. Nix switches between *profiles*, which is a way to provide different configurations for different user environments as shown in figure 2.2, providing atomic upgrades and rollbacks [25].

A fundamental component of the ecosystem is Nix, a domain-specific language designed for configurations, distinguished by its functional nature and lazy evaluation. The concept of purity is central to Nix, where values remain unchanged throughout computation and every function consistently yields the same output regardless of input [10]. The security implications of using Nix vs. an imperative system are discussed in the next chapter.

### 2.2.1 Non-declarative components

Declarative distributions such as Nix cannot do everything in the system in stateless manner. Some components of the system such as databases must have a distinct state, which cannot be practically declared with package manager apart from initial configurations [5]. Home directories can vary as much as the system administrator desires. For example, a configuration file for text editor vim is usually declared in the file

```
/home/<user>/.vimrc.
```

Nix provides multiple ways to perform the whole configuration process from the Nix configuration files. One way is declaring the desired `.vimrc` in the Nix configuration, as in the following program snippet:

```
{
  environment.systemPackages = [
    (pkgs.vimConfigurable.customize {
      vimrcConfig.customRC = ''
        " arbitrary vim config
      '';
    })
  ];
}
```

Nix also provides ways to fetch content to the system from remote URLs. If the administrator does not want the system to remain "pure", they can also build the system by

```
nixos-rebuild switch --impure
```

This results in the system having mutable components, which can be desirable from an accessibility point of view, but can cause unpredictable behaviour if the impure components are modified. Purity means that the components are read-only and immutable [26].

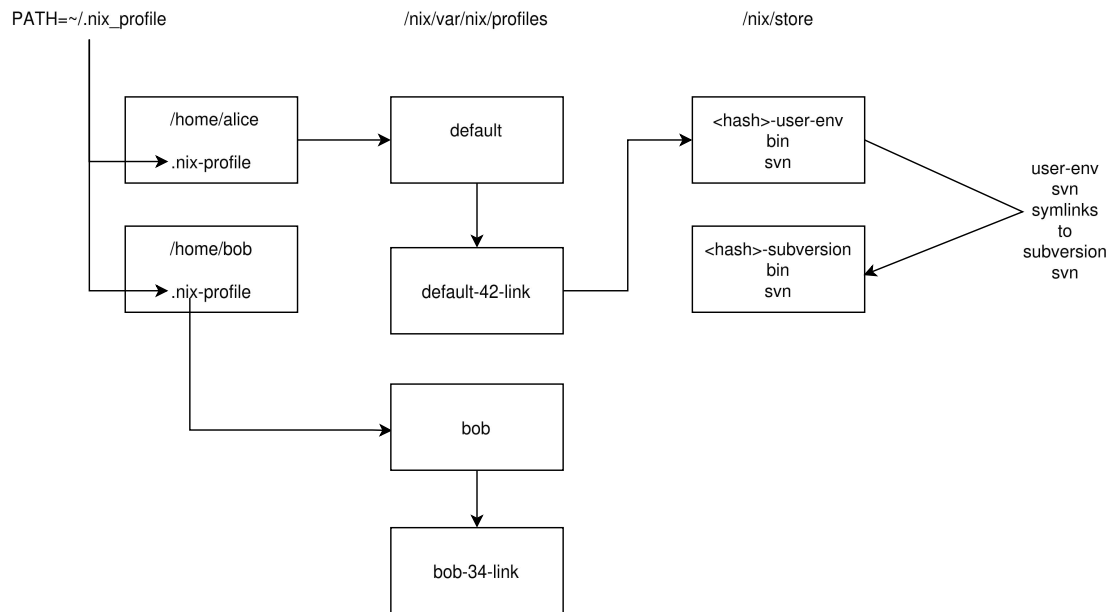


Figure 2.2: Relations between different user environments and installed programs [27].

User environments (Nix profiles) can be used in multiple environments for different users, in which the user can operate as shown in figure 2.2. User environments are a successor to the concept, where installed programs either reside in

`/usr/bin`

or

`/usr/sbin`

or have a symbolic link to the said directories. They can be figured as trees of symbolic links that reside also in the Nix store hence referred packages are called "activated packages" [2]. The installed programs reside usually in

`/nix/store`

Continuous build and integration services such as Hydra, which include Nix-compatible support for handling runtime configuration and tools such as Disnix and Charon<sup>1</sup>, focus on setting up complementary infrastructure. Van Der Burg

<sup>1</sup>Charon is now called NixOps [28]

[5] presents these new tools to replace Cfengine, Puppet and Chef, which execute operations in convergent manner, meaning that they capture what changes should be done to the machines in a specified network.

These approaches have two central problems: imperative nature of handling environment difference and inability to guarantee configuration compatibility with a machine. Disnix is a Nix derivative that can overcome these challenges by separating logical properties from physical and capturing the essential aspects which form a system [5].

### 2.2.2 Home manager and flakes

Nix environments can be built from a single `configuration.nix` file, but there are two significant configuration tools for managing Nix systems: home manager and flakes. Home manager is an extension for managing user profiles with a declarative Nix syntax [29]. However home manager has issues with atomic rollbacks and for this reason they are not used in this thesis' examples.

Flakes is an experimental feature of Nix, which provides environments, where dependencies are pinned in a lock file, further improving reproducibility of Nix systems [30]. A flake is a file system tree which contains a root directory with the Nix lock file specification called `flake.nix`. The usage of flakes is a favorable method for organising different environments within a Nix system where it can consist of multiple flakes. Flakes is an experimental feature, thus out of this thesis' scope.

### 2.2.3 Ease of updates

Due to atomic rollbacks, updating with Nix is easy and riskless. Nix handles software providing through channels. A channel is a set of latest Git commits in a Nixpkgs repository, where they are divided to stable/unstable and large/small channels [31]. Unstable channels (large and small) have the latest commits on a rolling basis, but

include less conservatively checked functionalities. Stable channels are submitted through a version number (e.g. 23.11), where a new release is published every six months. Large channels contain a full set of Nixpkgs binaries, when small ones include a subset. If a system administrator decides to submit to a small channel, they have more recent updates at their disposal, but have to resort to compiling some needed packages from source.

Updating a Nix system is just a manner of invoking command

```
sudo nix-channel --update
```

and if stable release is chosen, changing the

```
system.stateVersion
```

from the `configuration.nix` file [25]. Nixpkgs is a repository of working Nix packages using a continuous integration service called Hydra. Hydra evaluates the needed Nix expression of a package and ensures its functionality [25].

It is apparent that purely functional and declarative approaches are not as popular as imperative systems due to the need of steep learning curves and obscure syntax. This can be seen as historical payload: imperative models have been in use a longer time than purely functional approaches such as Nix.

In this chapter we revealed, that declarative systems have inherent strengths in both deployment strategies and system upgrading. The next section brings up the security viewpoint of declarative systems specifically in an embedded Linux setting.

## 3 Embedded system security

In the previous chapter, the core differences between imperative and declarative approaches were examined. This chapter first discusses the embedded Linux security in general, following up on how Nix can be used to improve embedded systems' security, reflecting the CIA-triad in section 3.2.

According to Serpanos et al. [32] the use of embedded devices can be divided into four fields: industrial systems, nomadic environments, private spaces and public infrastructure. This thesis focuses on public infrastructure, specifically in information screens in a public environment. Implementing security mechanisms and policies is essential for information screens to function securely from both organisational and technical viewpoints. Implementing those policies and assuring compliance is trivial with declarative approaches with increased benefits from reliability perspective.

Embedded systems are distinct from other types of systems due to their varying nature ranging from programmable logic controllers (PLC) to larger systems such as servers or routers [18]. A usual embedded device conducts a specific task and possibly demands networking capabilities. Working with embedded systems is typically working with a limited set of resources, demanding careful design when a multitude of features are needed. Even services such as SSH have had history of vulnerabilities, which prove that upgradability is a fundamental base of a secure system [33]. It can be argued that the security aspect of embedded devices could be improved significantly with the use of declarative systems as seen in section 2.2

and in the section 3.2.

Embedded devices demand precision and security as their function may be very critical for variety of safety reasons, e.g. in automotive industry or healthcare applications [18], [34]. Reliability is a defining requirement for number of embedded applications; a pacemaker that does not function reliably all the time is completely useless. While a declarative solution itself cannot fulfill all security needs, it definitely could improve the *reliability* of such systems.

As stated by Fysarakis et al. [18] implementing access control is essential for any system to prevent unauthorised access. Defining complex access control is trivial with Nix, as the configuration files denote completely which user has access to which resources. Access control in a modern day embedded environment could be hard to implement in traditional imperative systems, as scaling such system that spans multiple devices and changing environments would require a lot of manual intervention. This is definitely one of strengths of declarative approaches: scalability is never an issue when a centralised configuration defines the systems. A declarative approach is often adopted in the cloud as discussed in section 2.2, but its implementation still requires significant effort in the embedded Linux field.

Improving information security aspects using a policy modeling standard, CIA-triad is discussed in section 3.3 and Nix is reflected with the use of the triads axes. Dolstra et al. [12] states that Nix is policy-free meaning that it contains a set of mechanisms which allow policies to be constructed with configuration files.

Embedded systems being a broad field, devices are limited to those which can run Linux kernel and provide the most basic networking capabilities. These cover architectures i686, x86\_64 and arm64 supported by NixOS. PLCs and microcontrollers are outside of scope as NixOS needs a functional Linux kernel and a specific architecture to work.

### 3.1 Common embedded systems pitfalls

Common issues regarding embedded devices are their lack of updates, weak data integrity and the multitude of features [18], [35]. For example, a toy teddy bear may have an audio recorder, data transfer capabilities and ability to geolocate itself. These kinds of devices may lack firmware or software updates and the data-transfer may be insecure.

A solution for secure data transfer would be TLS-encrypted messaging between clients. This could be achieved with MQTT-protocol, but configuring certificates poses extra effort. Multitude of features is a definite security problem as the user may not be aware of them at all times [18]. In an increasing global world, importing embedded devices from unreliable sources can prove to be a security issue. The household items may or may not adhere to latest security compliance.

Attack surface of embedded systems in general range from physical access to network and geolocation problems. One way of manipulating a device, apart from directly gaining access to the operating system, is side channel attacks. Analysing the power or electromagnetic properties of device input/output can be used to determine critical aspects of a device, e.g. key lengths or algorithms of security measures [18], [32]. Attack surface may be used to gain access or performing denial of service attacks. Geolocating is both a privacy and security issue as location data may be used to trace identities of device users, which can lead to e.g. blackmailing, physical intrusion or other methods [18]. This means that the principals of this thesis' reference architecture could theoretically be targeted with such malicious intents.

Embedded systems have problems regarding monitoring and system administration. It is very different to have home automation system with less than 20 nodes than to have public transport embedded fleet in a big city with 2000 nodes. As the number of devices grows, so does the challenge of monitoring and administrative tasks. Home automation has usually one person dedicated to the task: the

home owner. Monitoring should be trivial to automatise (e.g. by using tools like Prometheus), but administrative tasks are harder to automatise due to tasks being potentially very challenging, even for dedicated system administrators. This is where Nix comes to play as updating any number of devices becomes trivial.

## 3.2 Nix as an embedded solution

Declarative systems have advantages over imperative systems in reliability and safety aspects due to two facts: rollout and rollback that are equally trivial tasks as well as desired configuration that can be tested in a sandbox environment. It is very accessible to manage a rollout strategy as the rollout/rollback can be done multiple times or executed completely in a replicated sandbox environment. Simpler and more straight forward practical steps give space for easier strategical planning [16].

Kandoi et al. [16] argue that with declarative systems, it should be nearly impossible to misconfigure the system in the first place and if a faulty state was achieved, a simple rollback could undo the changes. As stated in section 2.2, it is definitely possible to achieve faulty systems with Nix. It can be argued that these problems can be mitigated with a well thought rollout/rollback strategy with a dedicated test environment.

Updatability is possible with many different platforms, but it becomes problem when updating is a sole duty of a consumer who may or may not have the adequate knowledge on how or why they should update their systems. Lightweight updatability comes out-of-the-box with Nix, and that is something that inherently should make it more secure. However, consumer products are out of this thesis' scope.

Nix is a double-edged sword for system administration tasks. On one hand, it has a steep learning curve, but on the other hand it can make tasks that could be very challenging with traditional systems trivial. In a well built Nix ecosystem, security actions such as updating or modifying user or kernel space can be used

to enhance security and in such system, any changes could easily be replicated to multiple devices without the need for manual intervention.

Some other clear disadvantages for Nix in embedded Linux is the fact that a purely functional and declarative system inherently must use disk space more than its imperative counterparts. In the worst case scenario, if one derivation of a system takes up 1Gb of space, the changes could result in the system needing 2Gbs of space [11]. The worst case scenario rarely occurs, but due to Nix's indestructive nature, this formula of disk space demand has to be considered in an embedded Linux setting.

### 3.3 Imperative and declarative systems from CIA-triad approach

CIA-triad can be used as a tool to show conflicts between different points of information security interests. It consists of three meters: confidentiality, integrity and availability as seen in the figure 3.1. Confidentiality can be seen as a superset of privacy. Confidential data is classified with technologies such as data encryption and user privileges [36]. Integrity means that the data has not been tampered with and remains untouched by unauthorised parties while it is in transit or stored e.g. in a server [36]. A way of providing integrity is checking hashes of downloaded files. Availability is a user viewpoint to the accessibility of the system [36]. When confidentiality and integrity are pushed to the extreme, availability aspect suffers, e.g. when a service enforces a strict multi-factor authentication.

Systems with an imperative package manager are more accessible than declarative systems as learning a new programming language with esoteric paradigm can pose extra effort. Configuring a whole Nix system demands a thorough knowledge of Nix language and that hinders the ease of access to a Nix system from a sys-




Figure 3.1: The CIA-triad, a way to demonstrate conflicting security measures [20].

tem administrator standpoint. With NixOS, an easy extent of accessibility can be achieved via planting sufficient configuration files during device setup.

Atomic systems such as Nix have great benefits towards integrity. As the Nix store, where every installation is read-only, it is impossible for attackers to modify the store. That is not the case in traditional approaches, where a user with root privileges can arbitrarily modify installed programs and files.

According to Nix manual, it has three main strengths in relation with security. **Security by obscurity:** combined with the unusual file system and the usage of user-environments, some malware that rely on the usual locations of installed programs may fail [37]. This is a very thin layer of security, as targeted malware has usually no problems navigating in an unfamiliar environment. **Multi-user installations:** the requirement for root access nearly always widen the potential attack surface. Nix provides a way for multiple users installing their programs through the use of user environments hence mitigating the need for root access. This both lessens the availability aspect as well as mitigates the program's root access. When file changes are made in user-specific scope, a thin layer of isolation is achieved [25]. **Data integrity** is achieved both by installed programs residing in the read-only Nix store, but also them having been checked against SHA256

checksums. Moreover, the core installation resources for NixOS are GPG-signed by an administrative Nix team [37].

It is apparent that Nix can be used to improve embedded Linux security with some hardware constraints. A purely functional declarative solution can improve the rollout/rollback capabilities, testing pipelines and provides features such as read-only Nix store. The next chapter proposes an architecture solution which works as an example of a Nix ecosystem. The main purpose of the architecture solutions is to demonstrate the security implications of declarative systems which are then discussed in chapter 6.

## 4 Proposed architecture solution

For gaining data points for the research, a test setup must be configured. This chapter presents a Nix ecosystem which goes through a red team testing in chapter 6. The proposed architecture solution uses a completely declarative approach introduced in section 2.2, where strengths and weaknesses of both approaches are discussed.

The proposed architecture is a referential framework for a Nix ecosystem and provides only a basis for a subset of features needed in such system. For example, remote access and tunneling would require extra work in a production setting. Note that presented solutions' scalability could be improved with the use of tools referred in section 2.2.1.

The main purpose of the architecture is to construct a system that focuses on reliability and fluid deployment tasks. Kandoi and Hartke [16] discuss the operating large scale IoT solutions through declarative configuration APIs and conclude that it can solve multitude of scalability and extensibility challenges of IoT systems thus ensuring reliable and safe operations. As embedded security is the focal point of this thesis, a further analysis is presented in chapter 6.

### 4.1 Server

Van der Burg et al. [5] provide a reference architecture with OpenSSH, Quake 3 server and Transmission services. Dolstra et al. [10] build a declarative system

with simple web server. This thesis' architecture is moderately complex in comparison with these approaches, as it contains multiple components and a functioning server–client implementation.

A proposed architecture can be viewed in figure 4.2. In simplicity, the NixOS server runs a MQTT-broker for publishing images encoded in base64 format for the client devices. The clients are called device fleet and the server is called central server. The client devices submit to configured topics, displaying them using a Wayland compositor, Weston.

The central server has two main functions: receiving SSH-connections for admin usage and forwarding bitmaps formatted in base64 to the clients. In production environment, the server would be the element located in a public network and the devices would be accessible locally (through a tunnel). SSH-connections to the machine would then be forwarded through the central server to the clients. The central server currently does not generate the imagery, but this could be achieved via headless browser or other graphical tools.

MQTT is a an extremely lightweight machine-to-machine publish/subscribe protocol which can be used on virtually every platform including microcontrollers [38]. The chosen MQTT broker and client for this project is Mosquitto. The following program snippet shows how the Mosquitto server is configured in the NixOS server.

```
services.mosquitto = {
  enable = true;
  listeners = [
    {
      users.<user> = { acl = [ "readwrite #" ]; }; settings = {
        cafile = "<ssl-path>/ca.crt";
        certfile = "<ssl-path>/myhostname.crt";
        keyfile = "<ssl-path>/myhostname.key";
      };
    }
  ];
};
```

The "services" statement tells us that a SystemD service is being defined. The settings section specify which certificate and key files are to be loaded to the service.

The MQTT-broker (Mosquitto in this case) publishes a message that is forwarded to the subscribing clients via a following script.

```
nix shell nixpkgs#mosquitto --command mosquitto_pub -h localhost -t
images/test -m "$IMG_BASE64"
```

Sending could be automatised with a service using SystemD timer:

```
systemd.timers.publish-image = {
  timerConfig = {
    OnCalendar = "*-*-* *:*:00";
  };
  wantedBy = [ "timers.target" ];
};
```

which invokes a specified service.

## 4.2 Client

Some IoT solutions prefer relationship with client and server, where the client automatically searches for suitable server dynamically [16]. This thesis' client structure is static, meaning it follows static addresses, which forms initial connections that require manual intervention.

Both server and client are running NixOS. The client has three main functions: subscribing to media, receiving and displaying the gained media which can be arbitrary. Currently, the image refreshes every second and through configuration, even displaying animations with this setup could be possible. Media display happens with feh, an image showing tool, that works with X server. However, this example is using Wayland, so compability layer XWayland must be used [39]. Image data messaging functions through MQTT-protocol, which is explained more in the next

subsection.

### 4.2.1 MQTT-client

The MQTT client subscribes to a topic from the following script. The image is received as a base64 string, and is converted back to PNG format with the following script:

```
IP="<server ip>"
TOPIC="images/test"
nix shell nixpkgs#mosquitto --command mosquitto_sub -h $IP -t
  $TOPIC >" <image directory path>/image.base64" base64 -d "<image
  directory path>/image.base64" > images/latest.png
```

### 4.2.2 Weston/XWayland

Wayland is a display protocol aiming to replace partially or fully the old X window system [39]. Wayland functions through a "compositor" (server) and that provides a surface for the device to draw graphics. Wayland was selected for this project due to increased security, as the X window system has support for network transparency which broadens the attack surface. Wayland has combined server and client rendering with the Wayland compositor, so that safety-critical throughput between display server and window manager is not a concern.

The example project displays an image from a directory via script:

```
sleep 5 && /nix/store/qc9j6pm6ykyx531s4kb06084mczy2l6g-feh
  -3.10.1 /bin/feh -F -Z -R 1 <image-path>/latest.png
```

As the programs must be found from an absolute path, the system must generate the scripts accordingly. This is done via a SystemD service, specified as:

```
{ config, pkgs, ... }: let
  # Create the feh.sh script to launch feh
  fehLaunch = pkgs.writeText "feh.sh" ''
    echo "sleep 5 && ${pkgs.feh}/bin/feh -F -Z -R 1 <image
      directory>/latest.png" > /home/user/abzug-receiver/weston/
      img.sh
  '';
  # Create the initImg.sh script to compile the C code
  initImg = pkgs.writeText "initImg.sh" ''
    echo "${pkgs.gcc}/bin/gcc <source path> img.c -o <binary path>/
      img" > /home/user/abzug-receiver/weston/init.sh
  '';
in {
  systemd.services."weston1" = {
    enable = true;
    unitConfig = {
      Type = "oneshot";
    };
    serviceConfig = {
      # Set environment variables
      Environment = "XDG_RUNTIME_DIR=/var/run/user/1000";
      ExecStartPre = [
        "${pkgs.bash}/bin/bash ${initImg}"
        "${pkgs.bash}/bin/bash <init.sh path>/init.sh"
        "${pkgs.bash}/bin/bash ${fehLaunch}"
        "${pkgs.bash}/bin/bash <init.sh path>/init.sh"
      ];
      ExecStart = "${pkgs.weston}/bin/weston --config=<weston
        configuration directory>/weston.ini";
      RestartOn = "failure";
    };
    wantedBy = [ "graphical-session.target" ];
  };
}
```

This would not be a problem in a traditional Linux distribution. However, program locations vary from machine to machine in NixOS [26]. A program resides in Nix store with a cryptographic hash of all build inputs in its directory path [26]. For that reason, one way of proceeding is to write a SystemD service to generate the configuration files. Note that this SystemD configuration differs in how SystemD scripts are declared usually. Most SystemD distributions have SystemD files in

```
/lib/systemd/
```

system directly or via a symbolic link. In the beginning of the configuration, variables are defined and program locations expanded from Nix package paths. Then, in the "serviceconfig" part of the configuration, the strings are forwarded to shell which in part compiles sources and executes scripts. After the "ExecStartPre" section, in the "ExecStart", Weston is launched with a very basic kiosk configuration:

```
[core]
idle-time=0
xwayland=true
[shell]
panel-location=""
panel-position=none**
[autolaunch]
path=<feh launcher path>/img
```

the "[autolaunch]" only functions with compiled binaries thus the shell script is not directly executed. Instead, a program written in C is invoked, which in part invokes the shell script with parameters. The autolaunch path can handle switches, but unfortunately not parameters. With these workarounds the kiosk successfully can display media as shown in figure 4.1. Feh is launched with -R 1 parameter, which causes it to refresh the image every second. That way when a new image is uploaded, the display is also refreshed.




Figure 4.1: Working example of a NixOS client displaying image with Wayland and Feh.

The sources and scripts are downloaded from Github, again via a SystemD service. Following snippet shows the "ExecStart" part of the service:

```
ExecStart="${pkgs.git}/bin/git clone <git url> <installation path  
>";
```

This part of the configuration is not purely functional, as the downloaded scripts and configuration can be arbitrarily changed with correct permissions. The SystemD services' impurity do not trigger the Nix language evaluation itself, thus "--impure" switch is not mandatory.

## 4.3 Test device

A huge benefit from declarative systems is their broad possibilities of automated system tests [24]. The phrase: "if it works on one machine, it will work on another" builds a stable foundation for such tests [25]. Different deployment strategies benefit from slightly different approaches, but as deployment strategies generally are out of this thesis' scope, only a minimal test setup is configured.

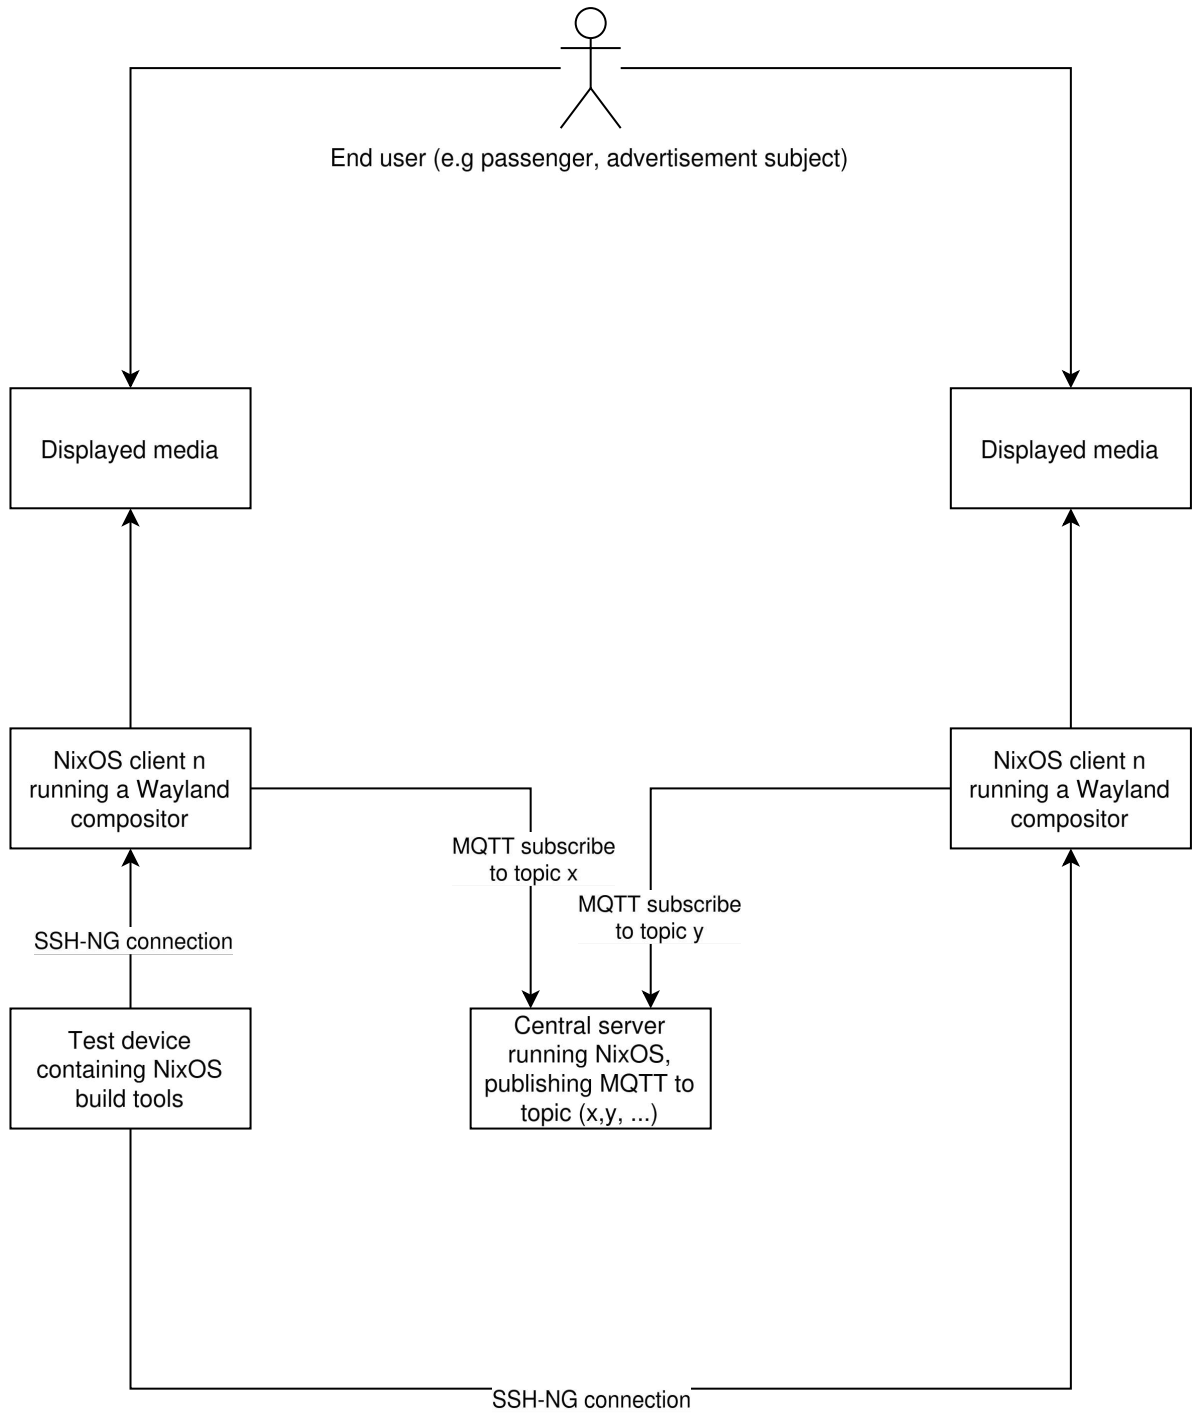


Figure 4.2: An architectural graph of the test setup.

NixOS devices can be upgraded and updated either through a server or locally with physical access [25]. In this thesis' example, a separate test machine is configured, and its programs are replicated to other devices as seen on figure 4.2. Updates are done manually with the following command:

```
nix build --eval-store auto --store ssh-ng://remote-host
```

The test device is thus replicated to all remote hosts manually, but an automation script would be useful as the setup scales.

## 4.4 Installing new devices

Installing devices could be made trivial as it is trivial to generate new NixOS images from pre-existing configuration files [25]. The problem is, however, that in many cases a legacy architecture exists which needs to be overwritten. This section focuses on the research question 3, providing a solution for migrating from existing architectures.

If the whole setup needs to be replicated to an existing device fleet, an open source project "NixOS anywhere" would be of great value. NixOS anywhere documentation defines minimal specifications: "Unless you're using the option to boot from a NixOS installer image, or providing your own kexec image, it must be running x86\_64 Linux with kexec support. Fortunately, most modern x86\_64 Linux systems have kexec support. By providing your own image you can also perform kexec for other architectures e.g. aarch64" [40]. NixOS anywhere also has a requirement that the devices should be available in a public network (not only wireless LAN), which likely would not be a problem in production environments. However NixOS anywhere only works with NixOS flakes. In that case, this thesis' configuration would need to be contained in a flake.

The installation would contain the following steps:

1. Run

```
curl -L https://github.com/nix-community/nixos-images/releases/download/nixos-unstable/nixos-kexec-installer-noninteractive-x86\_64-linux.tar.gz | tar -xzf- -C /root /root/kexec/run for booting separate kernel
```

for installation.

2. Generate a minimal configuration file:

```
nixos-generate-config --no-fileSystems --root /mnt''
```

3. Add ssh-keys to the configuration

4. Upload the flake from the server to the device:

```
nix run github:nix-community/nixos-anywhere -- --flake <path to configuration>#<configuration name> root@<ip address>''
```

Another way would be setting up NixOS from the installation medium and manually setting up or creating an installation media with NixOS generators project (<https://github.com/nix-community/nixos-generators>). NixOS generators is available from Nixpkgs and can be installed to a NixOS development machine or server. Invoking

```
nixos-generate -f iso -c <configuration.nix path>
```

would result in an ISO format image ready for installation [40].

The next chapter provides insight on how to measure security as a whole, transferring the necessary metric systems for this thesis' purposes. The architecture presented in this chapter then is gone through further analysis in chapter 6.

# 5 Security standpoints

Security is inherently challenging to measure adequately due to its complex and chaotic nature. Qualitative analysis may also result in subjective output. As such, no unambiguous standard of measuring security can be provided [8]. To overcome these challenges, several metric systems are compared and the one that provides the most precise answers specifically for this thesis' cause is chosen.

In this chapter, a selection process for meters which in part are used for gaining quantified data, is presented. When finally quantitative metrics are gained with the help of two paradigms, GQM (goal, question, metric) and specifically its superset SMART, a red-blue team layout is erected using QuERIES methodology. SMART is used in conjunction with the literature review to reveal the most suitable methodology for this thesis. SMART is opened more in the section 5.2.

Using an exact metric system is important, as this this thesis' emphasis is on measuring the security of a declarative approach. The presented architecture in the previous chapter works as an example on how a declarative system can be improved by reconfiguration, utilising a red-blue team setup.

## 5.1 A brief history of security metrics

The history of security metrics begin from Trusted Computer System Evaluation Criteria (TCSEC) also known as the Orange Book from 1983, which popularised many terms still in use today, such as identification, authentication and authorisation

[41].

In the 1990s, when the US National Bureau of Standards (NBS), later known as the National Institute of Standards and Technology (NIST), tried to standardise security, it became clear that systems needed to adhere closely to the definitions outlined in the Orange Book and the subsequent Common Criteria project [41]. Over time, various other standards gained popularity, such as the System Security Engineering Capability Maturity Model (SSE-CMM), which serves as a checklist for system design from the ground up.

Later it was observed, that systems design is only a part of a successful security strategy and operational practices played a bigger role than expected. This was something to be addressed in 1995 NIST Computer Security Handbook which evolved to provide ground to combat modern issues [41].

As metrics can be used as a tool for decision making, the strategical approach of the mentioned publishes is important. It is noteworthy, that the strategies (the Orange Book, SSE-CMM, etc.) begin to measure security by compliance to defined ratings [41]. Later in 2000s more mathematical approaches were taken, one which is delved deeper in section 5.4.

## 5.2 Choosing security metrics

Security is something that is challenging to measure due to its complex nature. A GQM (Goal, question, metric) paradigm helps to choose appropriate metrics: first there must be a set goal to an organisation, then a formulated question for each goal [42]. These answers are then reflected to gain the desired metric. This strategical approach is perhaps too broad for this thesis' scope, but it aligns well with a usual organisational strategy.

A more appropriate tool for this task would be SMART – a set of inputs to evaluate meter systems' suitability. These inputs describe how specific, measurable,

attainable, relevant and timely the methodology is [43].

In cyber security, being specific is very important. A common issue with security meters is, that they either cover too many topics and are without precise definitions or they are too specific to be generalised to a broader scope of situations [8]. The results of this thesis are crucial to measure, as the research focuses on system states and aims to assess the outputs produced by state transitions.

The concept of attainability is important in this context because this thesis has a narrower scope compared to a large organisation. Therefore, the proposed setup must be evaluated to ensure that the metrics' goals are achievable. Relevance has to do with risk assessment and how important it is to measure something related to its value. Risk assessment is explored in detail in section 6.1. Time-bounding signifies the importance of time as a meter; a system that can be penetrated in a minute can certainly be seen as weaker than a system that takes years to be compromised.

### 5.3 Measuring security

In this section, different methodologies and perspectives gained through literature review for cybersecurity are discussed, and potential methodologies are compared to gain the most adequate metric system for usage through SMART process [43].

Security metrics can be categorised into four themes: **system vulnerabilities**; measuring vulnerabilities can be applied to user-related, interface-induced, password, and software vulnerabilities [44]. Users are constantly at risk of threats such as phishing attacks or malware infections, where the user of any system becomes the primary attack vector. Interface-induced vulnerabilities refer to attack vectors associated with open ports and endpoints.

Password vulnerabilities refer to instances where a password can be cracked through computational methods [44]. This is fairly easy to measure, as it is possible to estimate the time required to crack a password or assess its vulnerability using

statistical password guessability. Software vulnerabilities are a common cause of security breaches. These vulnerabilities can be measured and estimated based on past exploitations. A key metric in this context is the time taken to patch a software vulnerability.

**Defense** measures can be applied to strengthen reactive, preventive, proactive and overall defenses. Reactive measures include blacklisting which is a lightweight mechanism to prevent e.g. for preventing botnet to harm the protected system by blacklisting IP-addresses related to the botnet. For measuring defence, the reaction time is essential and most importantly, it is a gained meter to measure preventive defense. Blacklisting can also be used as a preventive and a proactive measure, as a pre-filled blacklist can be used with desired parameters [9], [44].

Overall defenses can be measured with the combination of all defensive measures and with the use of penetration testing in a red–blue team setup. Penetration testing aims to gain a result, also known as penetration resistance, which is a meter indicating cost or time that the red team must spend in case of a successful system compromisal [9], [44].

**Threats:** zero-day vulnerabilities can be measured from two perspectives: life-time of zero-day vulnerability and the number of nodes that are compromised as a result. Malware spreading can be traced with the parameter infection rate, which is defined as infected node per a time unit. Attack evasion is assessed using either the obfuscation prevalence metric or the structural complexity metric. These metrics offer insights into the obfuscation of acquired samples, such as through encryption, or the complexity of the target system, which is measured by its runtime [9], [44].

**Situations** can refer to the security state, security incidents, and security investments. The security state encompasses various parameters, such as the incident rate and the blocking rate. Security investments measure the percentage of the budget allocated to security and the return on those investments [44].

## 5.4 Methodologies in comparison

Quantified mathematical models for cybersecurity metrics, which form the foundation of this thesis, are currently well-established. Three different methodologies are discussed, and one is picked for measuring the security of this thesis' architecture implementation. All the following metric systems are **measurable**. However, some fit better in relation with **time-related** and **relevance** axes.

The literature review provided three central methodologies from different perspectives. Complex mathematical models presented by Alshammari et al. [45] are too broad. This thesis' scope is limited and this methodology would fit better in a wider cyber security setting.

A methodology based on object-oriented thinking and UML diagrams is suitable in many contexts. However, as mentioned in the paper, this measurement methodology is used to compare similar alternative designs [45]. As noted, the focus of this thesis is not comparative, since all critical comparisons have already been made in chapter 2.

The Hidden Markov models presented by Wang et al. [46] are closely aligned with the end goal of this thesis. The time-related aspect is suitable, as the Hidden Markov model incorporates a time parameter. However, the issue lies in the generality of the methodology, and a more specific approach would be better suited for this thesis. Nevertheless, we apply a similar approach to that of Wang et al. for calculating the POMDP parameters.

The last and the most fitting methodology would be the one presented in papers by Carin et al. [19] and Hughes et al. [20]. The QuERIES methodology is delved deeper in section 5.5.1 and its time-related, relevance and specific axes are a near-perfect match for our goals as the model itself is relatively simple and provides shifting probabilities from states, which serves this thesis' study design well.

## 5.5 Quantitative metrics

Carefully selecting a metrics system includes asserting our goals and questions. Our goal is to discover this thesis' architecture proposals' tenacity in a simulated setting. The main goals reside in this thesis' two research questions:

- How can a declarative system be used to measurably improve the basic security needs of an embedded system used for displaying public media?
- What are the advantages and/or disadvantages of such system from system administrator standpoint?

The proposed architecture solution presented in chapter 4 will go through a red and blue team inspection, complying with the QuERIES model in chapter 6.

### 5.5.1 QuERIES

QuERIES model consists of number of steps that

1. model the problem - by conducting a risk assessment of the attack surface and the value of the possible intrusion
2. model the possible attacks - build an attack graph of intruding though vulnerabilities or other means
3. quantify the models - by conducting a controlled red team attack and provide quantified results for the said attack
4. use the results - use blue team methodologies to provide increased protection against the exposed problems

First, the blue teams' task contains the risk assessment of the attack surface. It is particularly important for the blue team to acknowledge the most crucial points

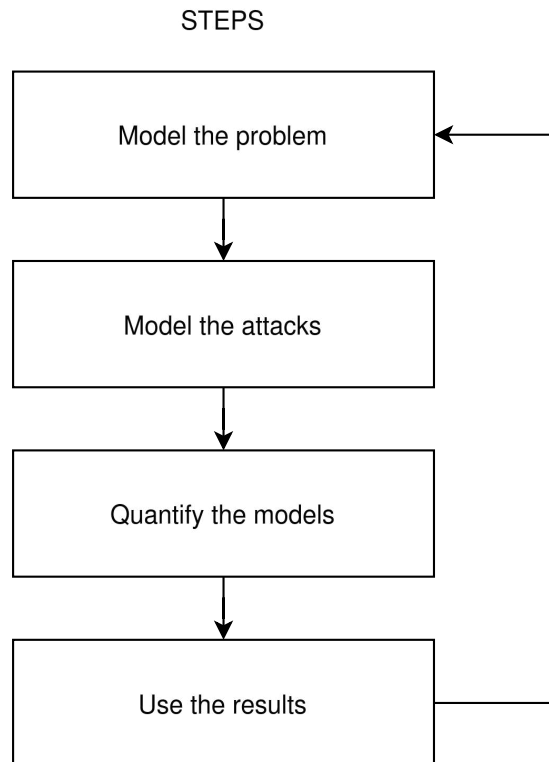


Figure 5.1: The QuERIES methodology is used as a reference flowchart for evaluation of security. [20]

of the attack surface, that is also used as the base for quantitative analysis. As seen in figure 5.1, the methodology is applied *iteratively*, i.e the steps are repeated as many times as needed for the system to be secure.

Modeling the possible attacks is a task for the red team – by constructing an attack graph, the opposing forces have a plan which can be used as a template for analysis. In this thesis, the models are quantified with the use of time framing. Both teams have a limited amount of time to conduct their tasks and the probability for succeeding a certain task is calculated with the formula

$$\frac{t_e}{t_t}$$

where  $t_e$  stands for elapsed time and  $t_t$  for maximum time that can be used which is the same for all tasks.

### 5.5.2 QuERIES as a central methodology

QuERIES draws inspiration from computer science, game theory, control theory and economics, thus is a complex answer to a complex question. It has been stated, that it can be used as an alternative to popular methodologies such as red teaming or black-hat analysis, used commonly in risk-assessment [19].

QuERIES is proposed to have potentially significant usage in DoD (Department of Defense) and in the private sector [19]. Initial testing of QuERIES in small-scale and realistic scenarios presented by Carin et al. [19] suggest that the methodology can in fact be used as to improve risk-assessment of more complex settings. This thesis follows similar steps: first the QuERIES methodology is used to assess risks followed by generalisation with strict constraints in mind.

As stated by Hughes et al. [20] the result of QuERIES is not binary: the attacker must think about the most optimal timeframe to stop the operation. The strategy uses **open** and **closed** loop decision algorithms for deciding when to stop trying. Closed loop decision algorithm constantly evaluates when is the optimal time to stop trying and open loop means that the system has pre-defined goals for evaluation [19]. As a positive side-effect, by gaining the probabilities through red-team evaluation, the system is thoroughly tested and improved. This is a valuable metric, offering insight into how the value of the system changes over time in terms of the cost of breaching it, thereby reflecting the true cost of the attack. As stated by Hughes and Cybenko [20], if the value is high enough related to the value of the gained value, they will perform the attacks less likely.

In this thesis, the value of the whole system architecture is defined as 1, constituting a value of holded intellectual property (IP). Costs are defined as fractions of the value, deflating 0.1 every hour. This is contrary to the original paper where the value of the intellectual property has a certain value of 30 000\$, and the value for the possible intruder is defined as 60\$ per hour [19]. This divergency is due to the

fact that it is impossible to define a certain value for our system. It also provides clarity as it is ergonomic to see how an attack estimates the relation to the value of the intellectual property.

### 5.5.3 Partially observable Markov decision process

Lastly, using the results of the POMDP can be used to increase the protection against the discovered problems [19]. The original QuERIES methodology used economic models for estimating POMDP parameters instead of calculating them manually [19]. We gain the parameters from POMDP by calculating state transitions with given observations. Attacks and defenses are quantified through a partially observable Markov decision process (POMDP) which contains the following six steps:

1. Define possible states the system can be in
2. Define the actions the system can take
3. Define the possible observations the system can take
4. Define the transition probabilities of the system
5. Define the observation probabilities of the system
6. Rewards: guide the system towards the desirable actions and states.

A POMDP is used widely in these kinds of applications, as both blue and red teams have only partial observations in relation to the system [47]. The blue team cannot be completely certain that the system is secure, and the red team cannot perform fully reliably as systems and environments differ from each other. As mentioned in the previous section, the focus of the QuERIES analysis is the time to reverse-engineer the system thus emphasising the importance of only partial observations.

In the next section the architecture presented in chapter 4 is improved using the applied QuERIES methodology. An attack graph is constructed and POMDPs are

utilised. The result will provide information about the architecture's security, and it will undergo several iterations of the QuERIES pipeline with each improving the overall security of the system.

## 6 Applying QuERIES

As mentioned in section 5.1, the QuERIES methodology is used for the analysis of this chapter. Table 6.1 consists of all the possible states (S), defender actions (D), attacker actions (A), and observations (O). Then, each transition from one state to another state is calculated as a probability. Carin et al. [19] use an economic model to estimate POMDP parameters. We deviate from this methodology and calculate the parameters as transition probabilities from each state for a given observation, as it provides a better construction of attack graph and yields more specific data.

QuERIES as a model is described as multidisciplinary as it takes inspiration from traditional red-blue team approaches and mathematical models to economic models [20]. In other literature, particularly that of Bojanc and Jerman-Blažič [48], it is argued that modeling cybersecurity with economic models can provide substance for minimising risks organisation-wide. In this chapter, we will apply the model and discuss its implications in both security and economic standpoints. However, the main focus is on how we can improve the security of the architecture using QuERIES.

## 6.1 Modeling the problem and quantifying the models

The project example of this thesis is an image showing architecture that could be used e.g. for advertisements, public transport timetables or practically anywhere static media should be presented. If an intruder gains unauthorised access, the results could range from anywhere from displaying explicit imagery to succeeding in displaying propaganda or other unwanted content. Unauthorised access would negatively impact service provider's economic standing, as every organisation displaying media wants to remain credible among stakeholders.

The attack surface of the example focuses on physical access and vulnerabilities in remote connections. With MQTT-messaging, SSH and display protocols, internal and external messaging takes place.

As mentioned in section 5.5.2, the value of intellectual property is capped at 1, for which other parameters refer to as fractions of it. If this was applied to a real setting, the value of intellectual property would be calculated appropriately for the scenario.

## 6.2 Modeling the possible attacks

In this section, the possible attacks are modeled using an attack graph, depicted as a POMDP which is a modeling tool presented earlier. In the original paper where QuERIES is presented, the time to reverse-engineer the system without prior information about the protection scheme is an important parameter [19]. Our approach is different; the attacks are considered as successful, if they gain further leverage in the attack graph e.g. transitioning from state "idle" to "partial loss of system". This is to maintain cohesion in the study, for we don't need to define what it means to "reverse engineer" the system. In our case, the system configurations are publicly available

as an open source project, thus the information of the system being available also to the attacker.

In table 6.1, the first column describes states the system can be in. The second and the third columns state defensive actions and observations of the system. The fourth column contains a template of the attacks that could be conducted. After forming the attack graph depicted in the table, the attacker tries to breach the system, leveraging through the A0-A4 column. The probabilities are then calculated with the model presented in section 5.1.

The algorithm produces a reward value for each QuERIES iteration, which is used to improve the setup. For calculating reward values, R script with library "pomdp" was used.

S0-S7	D0-D4	O0-O4	A0-A4
Idle	Monitor system	Normal operation	Intercept MQTT messaging
Receive media through MQTT	Isolate system	Detected suspicious activity	Compromise Github repository
Set up SystemD services	Shutdown system	Detected unauthorised access	Gain physical access to device
Start Weston	Isolate device	Detected unusual media display	Exploit vulnerabilities in display
Display media			Exploit vulnerabilities in SSH connections
Partial loss of system			
Complete loss of system			

Table 6.1: Different states, defensive measures, observations and attack measures for the system.

A weight of 1 was used for positive results and -100 was used if something was to be compromised. This weight distribution is due to the fact that even if blue team succeeds most of time, the effect of failure is more drastic than a succeeding result. The discount constant influences the priority of immediate vs. future rewards [47]. Our case signifies the importance of both, so value of 0.75 was used. Note that the maximum reward value is 4 and the minimum is -400.

### 6.3 Using the results

The reward score is taken in account on how successful or unsuccessful the setup is from a security perspective. The score itself is rather abstract; its main function is to demonstrate the *overall security* of the system. In short, negative numbers imply that there are flaws in security and positive numbers mean, that the system is more secure [47].

As stated by Hughes et al. using the results means evaluating the gained results to decide, if proposed protections are adequate for our means [20]. The QuERIES model was iterated 2 times and the results were placed in the table 6.2.

The reward function in the first iteration is calculated with the mentioned script. The second value is calculated as the maximum accumulated points; the highest value is 4 points without any discount, as the red team failed to provide any results for the second iteration. This is partially due to time limit, as there was only one attacker with very limited time in the test setup. It can be argued, that with more time, it would have been possible for the attacker to breach.

As seen, the reward function is growing as the proposed protections are applied

Iteration	Reward function	Proposed protections
1	-25	Multiple
2	4	None

Table 6.2: Protections applied in each POMDP iteration




Figure 6.1: Probabilities to breach into the system.

in both iterations. In the figure 6.1 we can see, that the time to breach peaks in the first hour.

The figure 6.1 has a similar shape as the figure in the paper by Carin et al. [19] meaning that attacks in their initial phase succeed more frequently. The figure 6.2 is used to decide, when is the optimal time to stop the attacks from the attacker's perspective. We use this as well to reflect the cost estimate of the blue team.

This economic model can be used to generalise risk in these kinds of situations. In the graph 6.2 we can see, that the cost for the attacker peaks at about 6 hours

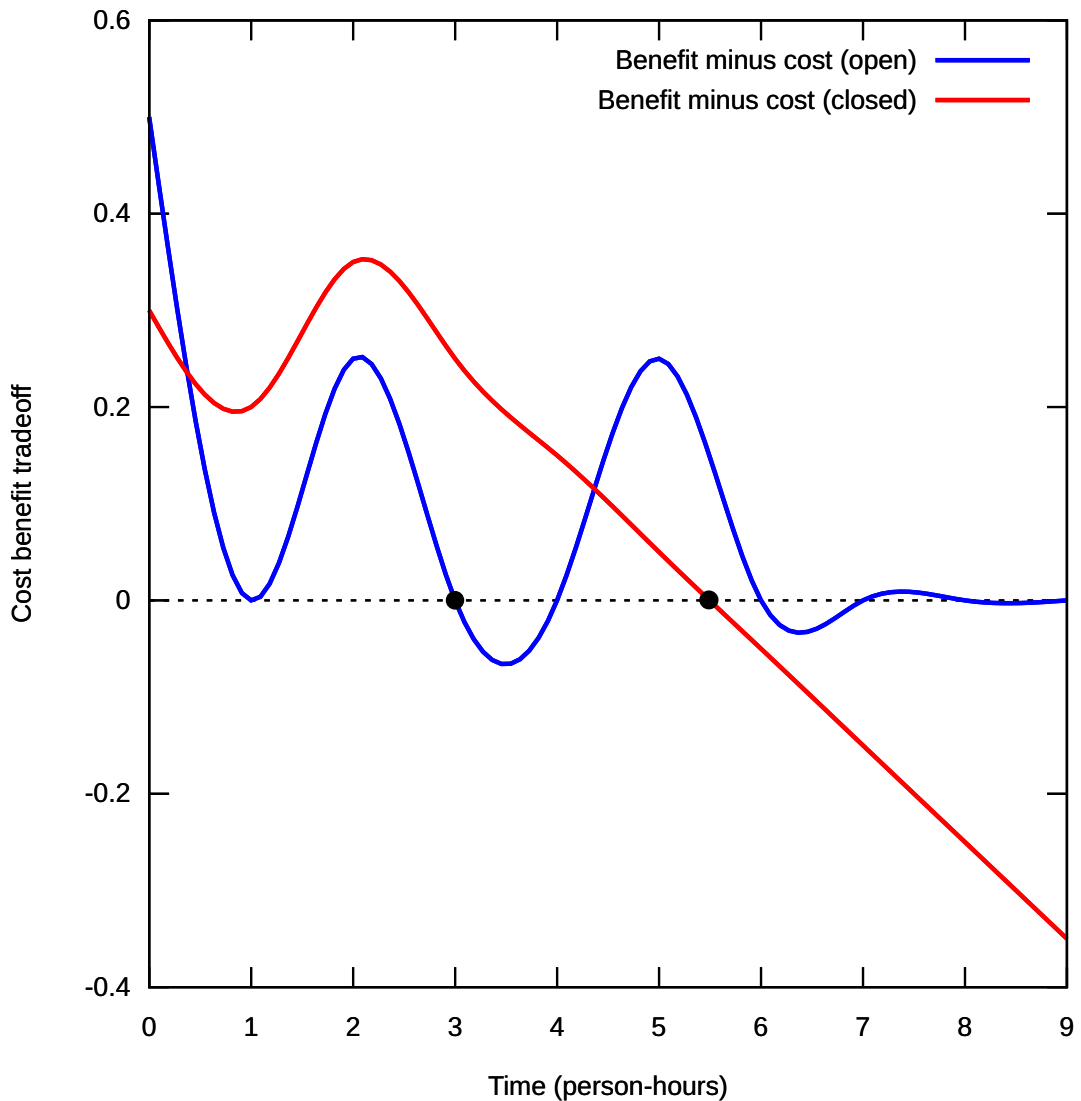


Figure 6.2: Cost benefits and reductions. The black dots indicate the most optimal times to stop the attack.

using closed algorithm and at 3 hours using open algorithm. Open loop algorithm means, that the cost estimate never evaluates feedback from a system [49]. Closed loop behaves in an opposite manner; it constantly modifies its behaviour based on the feedback it gets [49]. In our case, the cost values using the open algorithm remain unchanged, while with the closed algorithm, the value decreases by 0.1 every hour.

The blue team should be aware that in the first QuERIES iteration, most of the

breaches appeared in the first half of the time segment. The blue team can use this information to estimate the future distribution of breaches. In our case, the second iteration provided no breaches, but it can be argued, that in different experimental design, the attacker could have succeeded. The results do not mean that the system is completely secure; instead we can use the first iteration's estimates to decide on how could the blue team be improved in the future. The blue team measures depicted in figure 6.1 did not have the planned effect to protect the system.

We can see that with closed loop algorithm, the attacker can gain more results with worse cost tradeoff, but with open loop, it gets immediate results with a cheap cost. This could be important for the holder of the intellectual property as it must protect itself against attackers using both algorithms. In our case, where potential threats include hacktivist groups, the cost estimate is less significant. What matters more is the understanding that an attacker with no economic interests could have virtually unlimited time to attempt a system breach.

### 6.3.1 Improving the system using Nix

The following measures were taken in the first iteration: disable USB ports, change MQTT password to token based authentication and change SSH to key based authentication. As we can see, the QuERIES output provided concrete results on what are the weak points of the system.

As an example, the exposed USB-ports of the system were found to be a problem due to the possibility of infecting the system. This can be mitigated easily by modifying the client devices' configuration.nix with the following changes:

```
{ config, pkgs, ... }: {
  boot.kernelPackages = pkgs.linuxPackages_latest;
  boot.kernelParams = [
    "nousb" # Disables USB at the kernel level
  ];
  boot.kernelModules = [];
  boot.extraModulePackages = [];
  services.udev.extraRules = ''
    SUBSYSTEM=="usb", ACTION=="add", OPTIONS+="ignore_device"
  '';
}
```

The configuration could be applied to any number of clients, proving that Nix could be used to rapidly address arising security issues. This supports the argument presented in chapter 2, that through updatability the security could be improved. This highlights the scalability of Nix systems; it does not matter if we have one or thousands of devices – updating them is equally simple. A huge contrast between imperative and declarative systems is found, as imperative systems would need linear administration time in relation with the number of devices in the worst case. The configuration is only an example to mitigate hardware access attack vectors, as there are many more ways for the attacker to leverage the access, e.g. by replacing the whole device, depending on the resources of the attacker.

### 6.3.2 Issues with applied methodology

One issue found using QuERIES with POMDP was that applying the reward functions in our case is in rather arbitrary. We assumed, that for positive results the reward function is defined as 1 and for the negative outcomes as -100. This is due to the negative results deemed as catastrophic and the positive results being slightly positive. The choice for both parameters could have been any integer, but the issue

is, that it demands a "gut-feeling" of the author to select the appropriate parameters. This is due to the deviation from the QuERIES, where the POMDP parameters were not applied from the economic model. In our case, a trade-off to obtain specificity forced us to rely on the set reward and discount values. Using another model, for example the one presented by Wang et al. [46] using hidden Markov models, would in retrospective provide us a more specific layout of the attack graph, similar on how we applied the QuERIES methodology. The combination of both hidden Markov models and QuERIES could possibly be a "best of both worlds" solution.

One remark using POMDP is that it produces very generalised output. This is why we use methodologies such as QuERIES to improve the system, POMDP being just one component. Other benefit from using POMDP is that it provides us a concise attack graph and it can be argued, that using the POMDP calculations may even be redundant. In POMDP, negative rewards signify that the system has issues, and positive meaning that the system is more secure [47]. A more ergonomic approach would be calculating rewards without positive outcomes, due to them possibly shadowing the most critical issues.

## 6.4 Generalising the results

The results of this chapter could be generalised to a more complex setting, using a similar methodology. The POMDP would scale well, if more parameters were supplied. This study demonstrates that although measuring cybersecurity is challenging, staying within strict constraints allows for adequate estimates that can be used to apply economic models for the red team, further analysing the risk vs. reward.

Carin et al. [19] argue that QuERIES can be applied in both public and private sectors to help improve the security of both the software and hardware. It can be argued that a hand-tailored application of QuERIES can be used as a powerful model

that is ergonomic to use in the right hands. However, the use of complex applied mathematics models requires proficiency in both cybersecurity and mathematics, which can be difficult to achieve in a real organisational setting. This is why it is proposed in the next chapter, that a simpler, more pessimistic and more concise model would be easier to reach for an organisation.

## 7 Further research

One big limitation of Nix is the fact, that it mainly supports only x86\_64, i686 and arm64 platforms [25]. This is not generally enough as many different architectures are used in embedded, thus limiting the use-cases of NixOS [18].

A declarative approach could be used with other processor architectures. This would require work on a declarative package manager, which would be used in conjunction with base systems created with Buildroot or Yocto.

As far as security is concerned, the main methodology of this thesis' research, QuERIES, proved to be an adequate tool that helped improve the security of the reference architecture. However, some issues were raised in section 6.3.2. These issues could be addressed by developing a new methodology inspired by QuERIES, combining the acquisition of POMDP parameters methodically, without estimating them in conjunction with a more reachable calculation approach.

### 7.1 Further research questions

Questions for further research include:

1. What kind of new methodology would be better, reflecting the found issues in QuERIES for similar use-case as in this thesis?
2. What set of tools would be optimal in creating a secure true cross-platform declarative package manager

Developing a new Linux distribution would be a challenging task. Implementing support for different architectures for Nix or Guix for embedded use would be the focus of further work. These approaches would hold the same arguments as in this thesis to improve the security of embedded devices.

## 8 Conclusion

Nix can be a useful base for issuing a purely functional and declarative systems while providing an advanced rollback mechanism. Configuring it, depending on the system administrator, may be gruesome. However, it can dodge the usual pitfalls of more popular systems. While the usage of Nix language demands proficiency, it is a solution for distributing, updating and administrating more secure systems. NixOS handles features such as Wayland and audio driver setup effortlessly. Ultimately, NixOS is very predictable and the statement "if it works on one machine, it works on another" proves to be true.

Problems presented in this thesis could be solved with many different solutions, but the Nix approach is rather unique as it makes reproducible systems straight forward. As the system is defined in the `configuration.nix` file, there is little that could go wrong even if underlying hardware varies. The same does not apply to the mentioned Debian distribution.

Measuring and improving the architectures solution proved to be challenging but fruitful. With multiple iterations used with QuERIES, the architecture developed from insecure to more secure. This is a testament for applying a certain methodology for a specific task successfully.

As far as security implications are concerned, Nix is a good tool for asserting compliance. In purely functional environment, anything that needs compliance can be checked from the configuration files. As disk encryption is trivial and read-only,

Nix store provides increased integrity. A system administrator who is proficient with Nix, can assess the states of the systems very quickly. For example, viewing the open ports of each device can be answered quickly and in a precise manner by simply viewing the configuration files.

Popular embedded Linux distribution creating tools, Buildroot and especially Yocto have also steep learning curves as well. Furthermore, they lack the ease of updatability provided by Nix, which is a significant problem regarding security.

The research proved, that the purely functional Nix system is an easily updatable and upgradable solution for maintaining a secure system. While NixOS is a niche Linux distribution, and probably remains as such, the concepts however will likely live on in the future. A purely functional declarative system with atomic rollbacks is something to be iterated in the future with further user availability in mind.

# References

- [1] GNU Guix contributors, *Packages & GNU Guix* — *packages.guix.gnu.org*, <https://packages.guix.gnu.org/>, [Accessed 24-09-2024].
- [2] E. Dolstra and A. Löh, “Nixos: A purely functional linux distribution”, in *Proceedings of the 13th ACM SIGPLAN international conference on Functional programming*, Delft, The Netherlands, 2008, pp. 367–378.
- [3] OpenSUSE community, *System Recovery and Snapshot Management with Snapper | Reference | openSUSE Leap 15.0* — *doc.opensuse.org*, <https://doc.opensuse.org/documentation/leap/archive/15.0/reference/html/book.opensuse.reference/cha.snapper.html>, [Accessed 16-09-2024].
- [4] A. Tosini, *GitHub - nixos-bsd/nixbsd: An unofficial NixOS fork with a FreeBSD kernel* — *github.com*, <https://github.com/nixos-bsd/nixbsd>, [Accessed 09-11-2024].
- [5] S. van der Burg, *A Reference Architecture for Distributed Software Deployment*. Delft, The Netherlands: Citeseer, 2013.
- [6] NixOS foundation, *NixOS Search* — *search.nixos.org*, <https://search.nixos.org/packages>, [Accessed 24-09-2024].
- [7] L. Courtès, “Déploiements reproductibles dans le temps avec gnu guix”, *GNU/Linux Magazine*, 2021.

- 
- [8] A. J. A. Wang, “Information security models and metrics”, in *Proceedings of the 43rd annual Southeast regional conference-Volume 2*, Marietta, GA, USA, 2005, pp. 178–184.
- [9] A. Ramos, M. Lazar, R. Holanda Filho, and J. J. Rodrigues, “Model-based quantitative network security metrics: A survey”, *IEEE Communications Surveys & Tutorials*, vol. 19, no. 4, pp. 2704–2734, 2017.
- [10] E. Dolstra, R. Vermaas, and S. Levy, “Charon: Declarative provisioning and deployment”, in *2013 1st International Workshop on Release Engineering (RELENG)*, IEEE, Atlanta, Georgia, USA, 2013, pp. 17–20.
- [11] E. Dolstra and A. Hemel, “Purely functional system configuration management.”, in *HotOS*, Utrecht, The Netherlands, 2007.
- [12] E. Dolstra, M. De Jonge, E. Visser, *et al.*, “Nix: A safe and policy-free system for software deployment.”, in *LISA*, vol. 4, Utrecht, The Netherlands, 2004, pp. 79–92.
- [13] C. Endres, U. Breitenbücher, M. Falkenthal, O. Kopp, F. Leymann, and J. Wettinger, “Declarative vs. imperative: Two modeling patterns for the automated deployment of applications”, in *Proceedings of the 9th International Conference on Pervasive Patterns and Applications*, Xpert Publishing Services (XPS), Stuttgart, Germany, 2017, pp. 22–27.
- [14] S. van der Burg and E. Dolstra, “Declarative testing and deployment of distributed systems”, *Technical Report Series TUD-SERG-2006-020*, 2010.
- [15] E. Specht, R. M. Redin, L. Carro, L. d. C. Lamb, E. F. Cota, and F. R. Wagner, “Analysis of the use of declarative languages for enhanced embedded system software development”, in *Proceedings of the 20th annual conference on Integrated circuits and systems design*, Porto Alegre, Brazil, 2007, pp. 324–329.

- 
- [16] R. Kandoi and K. Hartke, “Operating large-scale iot systems through declarative configuration apis”, in *Proceedings of the 2021 Workshop on Descriptive Approaches to IoT Security, Network, and Application Configuration*, New York, NY, United States, 2021, pp. 22–25.
- [17] S. Ravi, A. Raghunathan, P. Kocher, and S. Hattangady, “Security in embedded systems: Design challenges”, *ACM Transactions on Embedded Computing Systems (TECS)*, vol. 3, no. 3, pp. 461–491, 2004.
- [18] K. Fysarakis, G. Hatzivasilis, K. Rantos, A. Papanikolaou, and C. Manifavas, “Embedded systems security challenges”, in *Measurable security for Embedded Computing and Communication Systems*, SCITEPRESS, vol. 2, Chania, Greece, 2014, pp. 255–266.
- [19] L. Carin, G. Cybenko, and J. Hughes, “Cybersecurity strategies: The queries methodology”, *Computer*, vol. 41, no. 8, pp. 20–26, 2008.
- [20] J. Hughes and G. Cybenko, “Quantitative metrics and risk assessment: The three tenets model of cybersecurity”, *Technology Innovation Management Review*, vol. 3, no. 8, 2013.
- [21] U. Breitenbücher, K. Képes, F. Leymann, and M. Wurster, “Declarative vs. imperative: How to model the automated deployment of iot applications?”, *Proceedings of the 11th advanced summer school on service oriented computing*, pp. 18–27, 2017.
- [22] Canonical, *Ubuntu Manpage: Apt - command-line interface — manpages.ubuntu.com*, <https://manpages.ubuntu.com/manpages/bionic/man8/apt.8.html>, [Accessed 22-05-2024].
- [23] G. K. Thiruvathukal, “Gentoo linux: The next generation of linux”, *Computing in science & engineering*, vol. 6, no. 5, pp. 66–74, 2004.

- 
- [24] S. Van Der Burg and E. Dolstra, “Automating system tests using declarative virtual machines”, in *2010 IEEE 21st International Symposium on Software Reliability Engineering*, IEEE, Delft, The Netherlands, 2010, pp. 181–190.
- [25] NixOS foundation, *NixOS Manual — nixos.org*, <https://nixos.org/manual/nixos/stable/>, [Accessed 22-05-2024].
- [26] E. Dolstra, A. Löh, and N. Pierron, “Nixos: A purely functional linux distribution”, *Journal of Functional Programming*, vol. 20, no. 5-6, pp. 577–615, 2010.
- [27] NixOS foundation, *User Environment - NixOS Wiki — nixos.wiki*, [https://nixos.wiki/wiki/User\\_Environment](https://nixos.wiki/wiki/User_Environment), [Accessed 09-11-2024].
- [28] NixOS foundation, *Nix/Nixpkgs/NixOS — github.com*, <https://github.com/nixos>, [Accessed 08-11-2024].
- [29] NixOS foundation, *Home Manager Manual — nix-community.github.io*, <https://nix-community.github.io/home-manager/>, [Accessed 24-05-2024].
- [30] NixOS foundation, *Flakes - NixOS Wiki — nixos.wiki*, <https://nixos.wiki/wiki/Flakes>, [Accessed 22-05-2024].
- [31] NixOS foundation, *Nix channels - NixOS Wiki — nixos.wiki*, [https://nixos.wiki/wiki/Nix\\_channels](https://nixos.wiki/wiki/Nix_channels), [Accessed 22-05-2024].
- [32] D. N. Serpanos and A. G. Voyiatzis, “Security challenges in embedded systems”, *ACM Transactions on embedded computing systems (TECS)*, vol. 12, no. 1s, pp. 1–10, 2013.
- [33] SecOpsSolution Inc., *History of SSH / SecOps® Solution — secopsolution.com*, <https://www.secopsolution.com/blog/history-of-ssh>, [Accessed 18-03-2024], Middletown, DW, USA.

- 
- [34] N. Turab and Q. Kharma, “Secure medical internet of things framework based on parkerian hexad model”, *International Journal of Advanced Computer Science and Applications*, vol. 10, no. 6, 2019.
- [35] R. A. Kemmerer, “Cybersecurity”, in *25th International Conference on Software Engineering, 2003. Proceedings.*, IEEE, Santa Barbara, California, USA, 2003, pp. 705–715.
- [36] G. Pender-Bey, “The parkerian hexad”, *Information Security Program at Lewis University*, 2019.
- [37] NixOS foundation, *Security - NixOS Wiki — nixos.wiki*, <https://nixos.wiki/wiki/Security>, [Accessed 22-05-2024].
- [38] MQTT, *MQTT Version 3.1.1 — docs.oasis-open.org*, <http://docs.oasis-open.org/mqtt/mqtt/v3.1.1/os/mqtt-v3.1.1-os.html>, [Accessed 22-05-2024].
- [39] Freedesktop, *Wayland — wayland.freedesktop.org*, <https://wayland.freedesktop.org/docs/html/>, [Accessed 18-09-2024].
- [40] NixOS foundation, *GitHub - nix-community/nixos-anywhere: Install nixos everywhere via ssh [maintainer=@numtide]* — *github.com*, <https://github.com/nix-community/nixos-anywhere>, [Accessed 01-10-2024].
- [41] J. Bayuk and A. Mostashari, “Measuring systems security”, *Systems Engineering*, vol. 16, no. 1, pp. 1–14, 2013.
- [42] Y. V. Papazov, “Cybersecurity metrics”, *NATO Science and Technology Organization (STO)*, pp. 1–18, 2019.
- [43] S. C. Payne, “A guide to security metrics”, *SANS Institute Information Security Reading Room*, 2006.

- 
- [44] M. Pendleton, R. Garcia-Lebron, J.-H. Cho, and S. Xu, “A survey on systems security metrics”, *ACM Computing Surveys (CSUR)*, vol. 49, no. 4, pp. 1–35, 2016.
- [45] B. Alshammari, C. Fidge, and D. Corney, “Security metrics for object-oriented class designs”, in *2009 Ninth International Conference on Quality Software*, IEEE, Brisbane, Australia, 2009, pp. 11–20.
- [46] H. Wang, S. Roy, A. Das, and S. Paul, “A framework for security quantification of networked machines”, in *2010 Second International Conference on Communication Systems and NETWORKS (COMSNETS 2010)*, IEEE, Bangalore, India, 2010, pp. 1–8.
- [47] J. C. M. Ashely S. M McAbee Murali Tummala, *The use of partially observable Markov decision processes to optimally implement moving target defense*, <https://scholarspace.manoa.hawaii.edu/server/api/core/bitstreams/86f32196-380c-443a-88b5-07f07649137e/content>, [Accessed 23-05-2024], Manoa, HI, USA.
- [48] B. Jerman-Blažič *et al.*, “An economic modelling approach to information security risk management”, *International Journal of Information Management*, vol. 28, no. 5, pp. 413–422, 2008.
- [49] R. Bars *et al.*, “Theory, algorithms and technology in the design of control systems”, *Annual Reviews in Control*, vol. 30, no. 1, pp. 19–30, 2006.