

Document Title	Requirements on Security Management for Adaptive Platform
Document Owner	AUTOSAR
Document Responsibility	AUTOSAR
Document Identification No	881

Document Status	published
Part of AUTOSAR Standard	Adaptive Platform
Part of Standard Release	R19-11

Document Change History			
Date	Release	Changed by	Description
2019-11-28	R19-11	AUTOSAR Release Management	<ul style="list-style-type: none"> • Reworded secure channel requirements [RS SEC 04001], [RS SEC 04003] and [RS SEC 04003] • Changed Document Status from Final to published
2019-03-29	19-03	AUTOSAR Release Management	<ul style="list-style-type: none"> • Unnecessary requirement [RS SEC 05006] removed
2018-10-31	18-10	AUTOSAR Release Management	<ul style="list-style-type: none"> • Chapter 2.3 'Protected Runtime Environment' revised
2018-03-29	18-03	AUTOSAR Release Management	<ul style="list-style-type: none"> • Moved the Identity and Access chapter into RS Identity and Access Management (899)
2017-10-27	17-10	AUTOSAR Release Management	<ul style="list-style-type: none"> • Initial Release

Disclaimer

This work (specification and/or software implementation) and the material contained in it, as released by AUTOSAR, is for the purpose of information only. AUTOSAR and the companies that have contributed to it shall not be liable for any use of the work.

The material contained in this work is protected by copyright and other types of intellectual property rights. The commercial exploitation of the material contained in this work requires a license to such intellectual property rights.

This work may be utilized or reproduced without any modification, in any form or by any means, for informational purposes only. For any other purpose, no part of the work may be utilized or reproduced, in any form or by any means, without permission in writing from the publisher.

The work has been developed for automotive applications only. It has neither been developed, nor tested for non-automotive applications.

The word AUTOSAR and the AUTOSAR logo are registered trademarks.

Table of Contents

1	Scope of Document	4
2	Requirements Specification	5
2.1	Functional Overview	5
2.2	Secure Communication	5
2.3	Protected Runtime Environment	6
2.3.1	Inter-Process Separation	6
2.3.2	Process-System Separation	8
2.3.3	Protection against Memory Corruption Attacks	9
2.3.4	Design Rules	9
3	Requirements Tracing	10
4	References	11
A	Protected Runtime Environment	13
A.1	Introduction	13
A.2	Protection against Memory Corruption Attacks	13
A.2.1	Overview	14
A.2.2	Secure Coding	15
A.2.3	Attacks and Countermeasures	16
A.2.3.1	Code Corruption Attack	16
A.2.3.2	Control-flow Hijack Attack	16
A.2.3.3	Data-only Attack	18
A.2.3.4	Information Leak	18
A.2.4	Existing Solutions	19
A.2.4.1	W ⊕ X, Write xor Execute, Data Execution Prevention (DEP)	19
A.2.4.2	Stack Smashing Protection (SSP)	20
A.2.4.3	Address Space Layout Randomization (ASLR)	21
A.2.4.4	Control-flow Integrity (CFI)	23
A.2.4.5	Code Pointer Integrity (CPI), Code Pointer Separation (CPS)	25
A.2.4.6	Pointer Authentication	26
A.3	Isolation	26
A.3.1	Horizontal Isolation	27
A.3.1.1	Virtual Memory	27
A.3.2	OS-Level Virtualization	28
A.4	Vertical Isolation	28

1 Scope of Document

This document specifies the requirements of Adaptive Applications to the functional cluster Security of the AUTOSAR Adaptive Platform. The motivation is to provide standardized and portable security in Adaptive Applications and Platforms.

2 Requirements Specification

2.1 Functional Overview

The AUTOSAR Adaptive Platform Security provides services for Adaptive Applications and other clusters of the AUTOSAR Adaptive Platform. The AUTOSAR Adaptive Platform Security is responsible for all aspects which regards to:

- Secure Communication
- Protected Runtime Environment

2.2 Secure Communication

[RS_SEC_04001]{DRAFT} Secure communication shall be transmitted using secure channels [

Type:	draft
Description:	Communication that is subject to specific security requirements shall be transmitted using a secure channel.
Rationale:	Provide secure communication means
Dependencies:	–
Use Case:	Nodes may need to communicate over a secure channel.
Supporting Material:	–

] ([RS_Main_00503](#), [RS_Main_00140](#), [RS_Main_00514](#), [RS_Main_00510](#), [RS_Main_00280](#), [RS_Main_00200](#))

[RS_SEC_04002]{DRAFT} Secure channels shall be configurable [

Type:	draft
Description:	Secure channels can be realized using various technologies. Therefore the kind and the parameters of a specific secure channel shall be configurable.
Rationale:	Enable modeling system security
Dependencies:	–
Use Case:	Nodes may host several services having different protection requirements thus requiring appropriate secure communication channels.
Supporting Material:	–

] ([RS_Main_00503](#), [RS_Main_01005](#), [RS_Main_00160](#), [RS_Main_00150](#), [RS_Main_00514](#), [RS_Main_00280](#))

[RS_SEC_04003]{DRAFT} The assignment of communication to specific secure channels shall be configurable [

Type:	draft
Description:	Secure channels can be realized using various technologies. Therefore, an application may be deployed in projects supporting different kinds of secure channels. The communication of the application shall be assigned to the respective secure channel.
Rationale:	Increase flexibility of an Adaptive Application
Dependencies:	–
Use Case:	Specific communication shall be assigned to the appropriate secure channel.
Supporting Material:	–

]([RS_Main_00503](#), [RS_Main_01005](#), [RS_Main_00106](#), [RS_Main_00150](#))

[RS_SEC_04004]{DRAFT} Using secure channels shall be transparent on the communication API [

Type:	draft
Description:	Communication through configured secure channels shall be facilitated transparently by the communication stack. Communicating through a secured channel shall make no difference for the application.
Rationale:	Reduce maintenance effort of Adaptive Application code
Dependencies:	–
Use Case:	Application code may be reused to transfer data with different security requirements.
Supporting Material:	–

]([RS_Main_00503](#), [RS_Main_00150](#), [RS_Main_00060](#))

2.3 Protected Runtime Environment

2.3.1 Inter-Process Separation

Each two applications shall be running isolated. There shall be a spatial, time, and resource separation between the two applications.

[RS_SEC_05002]{DRAFT} Memory of other processes [

Type:	draft
Description:	The Autosar Platform shall ensure that a process does not access the memory of any other process. This includes access to information stored in memory freed by another process.
Rationale:	Access to memory of other processes can be exploited, even if it is freed by the other process.
Dependencies:	–
Use Case:	–
Supporting Material:	–

]([RS_Main_00514](#))

[RS_SEC_05003]{DRAFT} Access via communication management [

Type:	draft
Description:	The AUTOSAR Platform shall restrict communication between two Adaptive Applications to mechanisms offered by Communication Management.
Rationale:	Inter-application communication is more secure when using the Communication Management.
Dependencies:	–
Use Case:	–
Supporting Material:	–

]([RS_Main_00514](#))

[RS_SEC_05004]{DRAFT} Individual persistency area [

Type:	draft
Description:	The AUTOSAR Platform shall ensure that an application uses only its individual persistency area. The application shall not be allowed to access the persistency area of any other application.
Rationale:	Access to other applications' persistency area could lead to information leaks.
Dependencies:	–
Use Case:	–
Supporting Material:	–

]([RS_Main_00514](#))

[RS_SEC_05005]{DRAFT} Individual process tree [

Type:	draft
Description:	The AUTOSAR Platform shall ensure that an application uses its individual process tree. It shall not be aware of any process belonging to another application.
Rationale:	Otherwise processes gain information from each other.
Dependencies:	–
Use Case:	–
Supporting Material:	–

](RS_Main_00514)

2.3.2 Process-System Separation

[RS_SEC_05019]{DRAFT} Access to Adaptive AUTOSAR Foundation and Services [

Type:	draft
Description:	The AUTOSAR Platform shall ensure that access from an application to the Adaptive AUTOSAR Foundation and Services is restricted to a minimum.
Rationale:	Direct access to the foundation can be abused or misused.
Dependencies:	–
Use Case:	–
Supporting Material:	–

](RS_Main_00514)

[RS_SEC_05008]{DRAFT} No physical memory address [

Type:	draft
Description:	The AUTOSAR Platform shall ensure that an application does not access any physical memory address directly.
Rationale:	Direct access to a physical memory address can be abused or misused.
Dependencies:	–
Use Case:	–
Supporting Material:	–

](RS_Main_00514)

2.3.3 Protection against Memory Corruption Attacks

[RS_SEC_05012]{DRAFT} The Adaptive Platform shall ensure that state-of-the-art protection mechanisms against memory corruption attacks are implemented. [

Type:	draft
Description:	The AUTOSAR Platform shall support the implementation of state-of-the-art protection mechanisms against memory corruption attacks.
Rationale:	Memory corruption is the root cause of nearly all vulnerabilities in today's software. Applications developed using unmanaged languages (such as C or C++) are likelier than managed languages to include errors that allow memory corruption attacks. If such vulnerabilities are exploited, an attacker may trigger for example arbitrary code execution, privilege escalation or extraction of sensitive information. A guideline for implementation is available as an Annex in this document. Note: The guideline is for reference only.
Dependencies:	Hardware, Operating System, Compiler
Use Case:	–
Supporting Material:	EXP_FCDesignSecurityManagement

]([RS_Main_00514](#))

2.3.4 Design Rules

[RS_SEC_05018]{DRAFT} Privileges and permissions [

Type:	draft
Description:	The AUTOSAR Platform shall ensure that an application is only granted the privileges and permissions mandatory for fulfilling its intended purpose.
Rationale:	More privileges and permissions than required can lead to a higher impact in case of a compromised application.
Dependencies:	–
Use Case:	–
Supporting Material:	–

]([RS_Main_00514](#))

3 Requirements Tracing

The following table references the features specified in [1] and links to the fulfillments of these.

Feature	Description	Satisfied by
[RS_Main_00060]	AUTOSAR shall provide a standardized software interface for communication between Applications	[RS_SEC_04004]
[RS_Main_00106]	AUTOSAR shall provide the possibility to extend the software with new SWCs without recompiling the platform foundation	[RS_SEC_04003]
[RS_Main_00140]	AUTOSAR shall provide network independent communication mechanisms for applications	[RS_SEC_04001]
[RS_Main_00150]	AUTOSAR shall support the deployment and reallocation of AUTOSAR Application Software	[RS_SEC_04002] [RS_SEC_04003] [RS_SEC_04004]
[RS_Main_00160]	AUTOSAR shall provide means to describe interfaces of the entire system	[RS_SEC_04002]
[RS_Main_00200]	AUTOSAR specifications shall allow resource efficient implementations	[RS_SEC_04001]
[RS_Main_00280]	AUTOSAR shall support standardized automotive communication protocols	[RS_SEC_04001] [RS_SEC_04002]
[RS_Main_00503]	AUTOSAR shall support change of communication and application software at runtime.	[RS_SEC_04001] [RS_SEC_04002] [RS_SEC_04003] [RS_SEC_04004]
[RS_Main_00510]	AUTOSAR shall support secure onboard communication	[RS_SEC_04001]
[RS_Main_00514]	AUTOSAR shall support the development of secure systems	[RS_SEC_04001] [RS_SEC_04002] [RS_SEC_05002] [RS_SEC_05003] [RS_SEC_05004] [RS_SEC_05005] [RS_SEC_05008] [RS_SEC_05012] [RS_SEC_05018] [RS_SEC_05019]
[RS_Main_01005]	AUTOSAR shall establish communication paths dynamically	[RS_SEC_04002] [RS_SEC_04003]

4 References

- [1] Main Requirements
AUTOSAR_RS_Main
- [2] Explanation of Adaptive Platform Design
AUTOSAR_EXP_PlatformDesign
- [3] SoK: Eternal War in Memory
- [4] Guidelines for the use of the C++14 language in critical and safety-related systems
AUTOSAR_RS_CPP14Guidelines
- [5] The SPARC Architectural Manual, Version 8
<http://sparc.org/wp-content/uploads/2014/01/v8.pdf.gz>
- [6] OpenBSD-3.3 announcement, public release of WX
<http://www.openbsd.org/33.html>
- [7] Smashing The Stack For Fun And Profit
<http://phrack.org/issues/49/14.html>
- [8] The Geometry of Innocent Flesh on the Bone: Return-into-libc without Function Calls (on the x86)
- [9] Jump-oriented Programming: A New Class of Code-reuse Attack
- [10] On the Expressiveness of Return-into-libc Attacks
- [11] Code-Pointer Integrity
- [12] ARM Pointer Authentication
<https://lwn.net/Articles/718888/>
- [13] PaX ASLR (Address Space Layout Randomization)
- [14] Control-flow Integrity
- [15] AMD64 Architecture Programmer's Manual Volume 2: System Programming
<http://support.amd.com/TechDocs/24593.pdf>
- [16] PowerPC Architecture Book, Version 2.02
<https://www.ibm.com/developerworks/systems/library/es-archguide-v2.html>
- [17] PowerPC Operating Environment Architecture Book III
<http://public.dhe.ibm.com/software/dw/library/es-ppcbook3.zip>
- [18] Linux Kernel, Summary of changes from v2.6.7 to v2.6.8
<https://www.kernel.org/pub/linux/kernel/v2.6/ChangeLog-2.6.8>
- [19] PAX
<https://pax.grsecurity.net/docs/pax.txt>

- [20] CPI LLVM on github
<https://github.com/cpi-llvm>
- [21] Flipping bits in memory without accessing them: An experimental study of DRAM disturbance errors
- [22] Drammer: Deterministic Rowhammer Attacks on Mobile Platforms
- [23] ANVIL: Software-Based Protection Against Next-Generation Rowhammer Attacks
- [24] A seccomp overview
<https://lwn.net/Articles/656307/>
- [25] Frequently Asked Questions for FreeBSD 10.X and 11.X
<https://www.freebsd.org/doc/en/books/faq/security.html>
- [26] pledge(2)
<https://man.openbsd.org/cgi-bin/man.cgi/OpenBSD-current/man2/pledge.2>

A Protected Runtime Environment

A.1 Introduction

Vulnerabilities in software programs lead to unauthorized system manipulation and access when they are exploited by runtime attacks. Unauthorized system manipulations are, for instance, arbitrary code execution, privilege escalation, or persistent manipulation of storage. The cause of the vulnerabilities are mostly programming mistakes and design flaws. Although design rules are known during the development process or quality assurance processes like static code analysis or fuzzing are performed, vulnerabilities exist statistically in nearly all projects. These measures can be specified only qualitatively by the AUTOSAR Adaptive Platform specification. However, there are technical countermeasures on the operating system level to harden a system against such attacks. Note that the term harden includes that there is still no guaranteed system security but the effort for a successful attack can be raised to a higher level.

The hardening measures are combined as the term *Protected Runtime Environment* (PRE) in the context of AUTOSAR Adaptive Platform. A PRE includes the most important, basic protection mechanisms for a complex software environment. Without it, any other security mechanism will be circumventable, as any compromised process or service will be able to compromise any other running process on a system. The goal of a protected runtime environment is to protect the integrity of a process' control-flow during runtime and to limit the impact of a successful attack. So, two strategies of hardening are considered for AUTOSAR Adaptive Platform: proactive protection that should mitigate the exploitation of vulnerabilities and "post-mortem" measures that isolates an untrusted process.

The present document is a guidance for integrators and implementers who are going to develop an automotive system that is compliant to the AUTOSAR Adaptive Platform specifications. It contains recommendations and hints to fulfill the desired security requirements listed in this document. The actual implementation of a specific measure is yet to be defined and may depend on the concrete system at hand, or may be a combination of multiple measures.

Section [A.2](#) introduces exploit mitigation approaches and presents related integration options. Afterwards, Section [A.3](#) discusses the isolation aspects that limits the action scope of a compromised or untrusted process. In each chapter the general attack and mitigating techniques are detailed and existing countermeasure implementations are presented. Further, technical requirements for the integration are highlighted.

A.2 Protection against Memory Corruption Attacks

Unmanaged languages, such as C or C++, enable programmers to implement their code with a high degree of freedom to manage resources. As such, code can be optimized for runtime performance or memory consumption and access to low level

functions of the operating system is possible. However, programmers are fully responsible for bounds checking and memory management since C/C++ is memory-unsafe. In practice, this often leads to the violation of the *Memory Safety* policy or memory errors, as the manual management of memory is very error prone.

Memory errors in turn can be utilized to cause memory corruption, the root cause of nearly all vulnerabilities in software components. If the vulnerabilities are exploited by an attacker the possible impact is, for example, arbitrary code execution, privilege escalation, or the leakage of sensitive information.

The language of choice in AUTOSAR AP is C++14 as proposed in platform design explanatory document [2, Section 2.3.1]. As such, the designated programming language for Adaptive Applications (AA), AUTOSAR Runtime for Adaptive Applications (ARA), as well as Functional Clusters on the Adaptive Platform Foundation or the Adaptive Platform Services is unmanaged, resulting in a large attack surface. One goal of a PRE is the minimization of this attack surface.

This chapter is structured as follows. Section [A.2.1](#) gives an overview of the potential types of memory corruption vulnerabilities. The typical pitfalls during coding are detailed in section [A.2.2](#). In section [A.2.3](#) an explanation of the basic technical reasons of memory corruption attacks and the corresponding state-of-the-art protection techniques on various attack stages is given. Further, possible practical solutions to implement and integrate the presented protection techniques in terms of the AUTOSAR AP are presented in Section [A.2.4](#).

A.2.1 Overview

The exploitation of vulnerabilities and its mitigation is a complex topic. Computer security researchers continuously develop new attacks and corresponding defenses. A general model for memory corruption attacks and the corresponding protecting techniques is described in [3]. The model (cf. Figure [A.1](#)) summarizes the general causes of vulnerabilities, the way to exploit them according to the targeted impact, as well as mitigation policies on the individual attack stages for four types of attacks: *Code corruption attack*, *Control-flow hijack attack*, *Data-only attack*, and *Information leak*. On each attack stage they define several policies that must hold to prevent a successful attack.

The first two stages are common for all attack types and describe the root cause of vulnerabilities. In the first stage a memory corruption manipulates a pointer. When this *invalid* pointer is then dereferenced, a corruption is triggered. A pointer is invalid if it is an *out-of-bounds* pointer, i.e. pointing out of the bounds of a previously allocated memory area, or if it becomes a *dangling pointer*, i.e. pointing to a deleted object. Commonly known out-of-bounds vulnerabilities are for example: *buffer overflow/underflow*, *format string bug*, and *indexing bugs* like *integer overflow*, *truncation* or *signedness bug*, or *incorrect pointer casting*. Typical dangling pointer vulnerabilities are: *use-after-free* or *double-free*. A collection of C and C++ related issues can be found, for instance, in the *List of Software Weakness Types* of the *Common Weakness*

Enumeration (CWE) from MITRE¹. The exploration of memory errors is the first step of an attack. Subsequently a pointer is dereferenced to read, write or free memory.

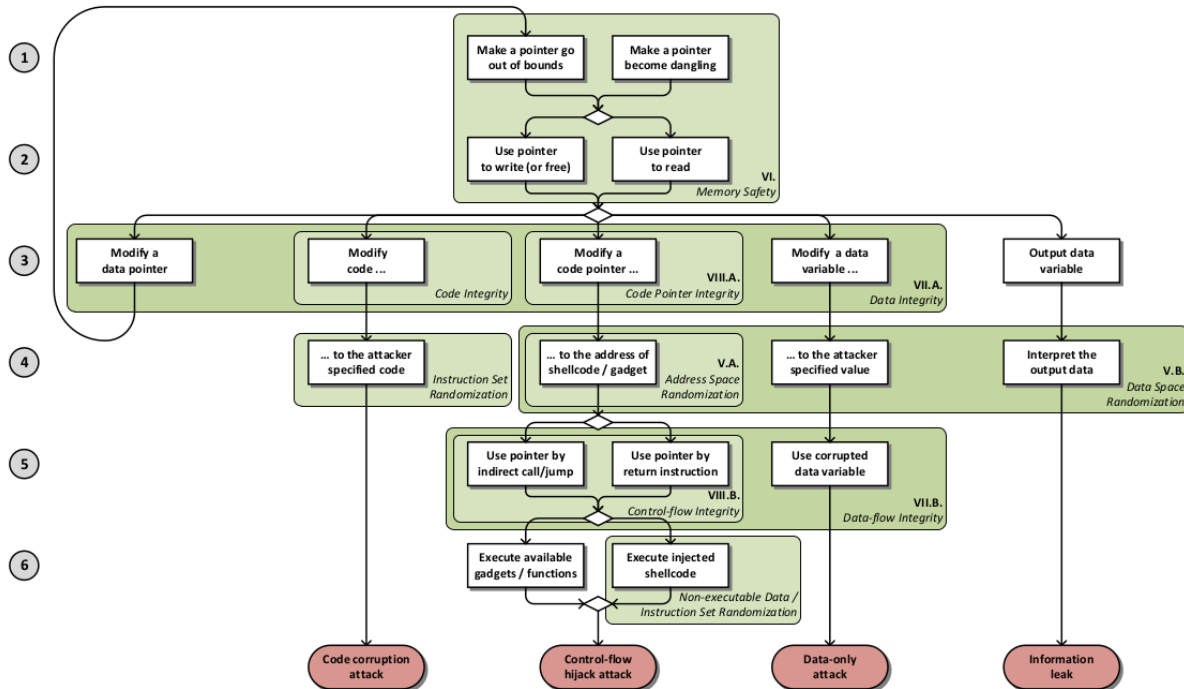


Figure A.1: Attack model from [3] demonstrating four attack types, policies mitigating the attacks in different attack stages

A.2.2 Secure Coding

A first measure to counter vulnerabilities at their root is to avoid mistakes and errors in the first place. To reach this goal programmers have to take care of many pitfalls during the development process. A simple example is the usage of unsafe functions from the standard C library like `strcpy()`. It copies a null character terminated character string to a buffer until the null character is reached. If the allocated destination buffer is not large enough, the function still copies characters behind the end of the buffer and thus overwrites other data. This is one of many pitfalls commonly known as a *buffer overflow* and can be used by an attacker, for example, to overwrite the stored return pointer if the buffer is allocated on the stack. For the given example a programmer should use safer variants instead. To that end, many standard C library functions have been supplemented with versions including a bounds check, for `strcpy()` this is `strncpy()`. Therewith the length of the input is limited and the buffer, if it is allocated properly, does not get overflowed. While there are supposedly more safe functions, such as `strncpy()`, they come with their own quirks and flaws. But also more complex, context related issues must be considered.

¹<http://cwe.mitre.org/data/definitions/659.html>

In practice programmers should have coding guidelines at hand like the *MISRA* for safety-related systems. Unfortunately the AUTOSAR Coding Guideline does not cover a rule set for security related issues [4]. But there are third party guidelines which deals with these issues, like the *SEI CERT C++ Coding Standard*².

Since these guides are very comprehensive only few programmers will follow them in practice due to time reasons. However, there are some tools that can support the check against some rules in a static code analysis, e.g. *Flawfinder*³, *RATS*⁴, or *CodeSonar*⁵. In any case the application of tools does not guarantee vulnerability-free code as runtime conditions (or execution contexts) are out-of-scope of such tools. Similarly, identified vulnerabilities may not be fixed sufficiently with respect to runtime behavior.

A.2.3 Attacks and Countermeasures

For a second line of defense it is assumed that vulnerabilities are present and that some will always be inserted by developers. Therefore a requirement for enhanced countermeasures is independence from written code. The goal is to mitigate exploitation of vulnerabilities.

A.2.3.1 Code Corruption Attack

A *Code Corruption Attack* intends to manipulate the executable instructions in the text segment in the virtual memory space and to breach the *Code Integrity* policy (cf. Figure A.1). The countermeasure is to set the memory pages containing code to read-only or $W \oplus X$ (write xor execute), respectively. It has to be implemented on both system levels, i.e. the processor as well as the operating system. The MMU of the CPU has to provide fine-grained memory permission layout (e.g. NX-bit [5, p.248]) or the operating system should emulate this. Further, the operating system level has to support the underlying permission layout, e.g. like $W \oplus X$ [6] or Data Execution Prevention (DEP). But care has to be taken if self-modifying code and Just-In-Time (JIT) compilation is used, as the generated code must first be written to writeable pages, which are then set to be executable.

A.2.3.2 Control-flow Hijack Attack

A *Control-flow Hijack Attack* starts with the exploitation of a memory corruption to modify a code pointer so that it points to an attacker defined address and the *Code Pointer*

²<https://www.securecoding.cert.org/confluence/pages/viewpage.action?pageId=637>

³<https://www.dwheeler.com/flawfinder/>

⁴<https://security.web.cern.ch/security/recommendations/en/codetools/rats.shtml>

⁵<https://www.grammatech.com/products/codesonar>

Integrity policy is harmed (cf. Figure A.1). This pointer is then used by a indirect control flow transfer in the original code. Therewith the control-flow is diverted from the original and so its *Control-flow Integrity* is violated. The last step is the execution of the exploit payload. The literature distinguishes between two approaches: *code-injection attacks* and *code-reuse attacks*. While code-injection attacks [7] are based on injecting arbitrary and custom instructions (a.k.a. *shellcode*) into the memory as exploit payload, code-reuse attacks, such as *return-oriented program* (ROP) [8], *jump-oriented programming* (JOP) [9], and *return-to-libc* [10], utilize existing code in the process memory to construct so called *gadgets*, which enable the targeted malicious functionality.

All in all, a control-flow hijack attack will be successful if the integrity of a code pointer and of the control-flow are broken. Further, the value of the target address of the malicious functionality must be known and in the case of code-injection, the memory pages holding the injected code must be executable.

The goal of *Code Pointer Integrity* is satisfied if all dereferences that either dereference or access sensitive pointers, such as direct and references to code pointers, are not modified (cf. [11]). There are a few recent feasible approaches which detect the alteration of code pointers at this early stage. The *Code Pointer Integrity* (CPI) [11] mechanism provides full memory safety but just for direct and references to code pointers. According to the authors, this approach protects against all control-flow hijack attacks. CPI is a combination of static code analysis to identify all sensitive pointers, rewrite the program to protect identified pointers in a safe memory region, called *safe-stack*, and instruction level isolation that controls the access to the safe region. These mechanisms require both compiler and runtime support and comes with an overhead of ca. 8% on runtime. An additional requirement is Code Integrity. *Code Pointer Separation* (CPS) [11] is a relaxation of CPI. Among others, CPS limits the set of protected pointers to code pointer (no indirections) to lower the overhead to ca. 2%. It still has strong detection guarantees. A further approach to detect modification on code pointers is *Pointer Authentication* [12]. This approach uses cryptographic functions to authenticate and verify pointers before dereferencing. Pointers are replaced by a generated signature and cannot be dereferenced directly. This requires compiler and instruction set support. To reduce overhead, hardware acceleration for the cryptographic primitives is required.

With *Stack Smashing Protection* (SSP) stack-based buffer overflows, which overwrite the saved return address, can be detected. Therefore, a pseudo-random value, also known as *Stack Canary* or *Stack Cookie*, is inserted on the stack before the saved base pointer and the saved return address as part of the function prologue. When the function returns and the function epilogue is performed, the value is compared to a protected copy. If the value is overwritten by a buffer overflow targeting the return address, the program aborts, because the values do not match anymore.

As mentioned before, code-injection attacks require executable memory pages for the injected instructions. With principle of $W \oplus X$, also called *Data Execution Prevention* (DEP), a memory page is either flagged as writable or executable, but not both. This prevents that instructions overwrite data memory such as stack or heap and ex-

ecute it afterwards. The approach requires a fine-grained page permissions support either by the MMU of the CPU and the so called *NX-bit* (No-execute bit) or emulated in Software as described in Section [A.2.3.1](#). Moreover, the employed operating system must support it.

Code-reuse attacks are not affected by the $W \oplus X$ mechanism since existing, executable marked, memory regions are utilized and no additional code must be injected to the memory. To mitigate such kind of attacks currently deployed countermeasures are implemented on previous attack stages. On the fourth attack stage it is stated that the target address of the malicious functionality must be known. In general, the attacker knows or can just estimate an address in the virtual address space since it is static for a binary after compilation. A countermeasure in practice is the obfuscation of the address space layout by *Address Space Layout Randomization* (ASLR) [13]. Therewith the locations of various memory segments get randomized which makes it harder to predict the correct target addresses. ASLR requires high entropy to prevent brute-force de-randomization attacks and depends on the prevention of unintended Information Leak (Section [A.2.3.4](#)) that are used by dynamically constructed exploit payloads. To guarantee high entropy ASLR should be implemented on 64-bit architectures (or above). Additionally, every memory area must be randomized, including stack, heap, main code segment, and libraries.

In addition to ASLR the policy Control-flow Integrity intends to detect a diversion of the original control-flow. Established techniques are: *Shadow Stack* and *Control-flow integrity (Abadi)* (CFI) [14]. The idea of shadow stack is to push the saved return address to a separate shadow stack so that it is compared upon a function return. In addition, CFI also protects indirect calls and jumps as well. The original CFI creates a static control-flow graph by determining valid targets statically and give them a unique identity (ID). Afterwards, calls and returns are instrumented to compare the target address to the ID before jumping there. It is required to protect valid targets from overwrite by $W \oplus X$.

A.2.3.3 Data-only Attack

A memory corruption can also be exploited to modify security critical data that is not related to control-flow data. For instance, exploiting a buffer overflow to alter a conditional construct can lead to unintended program behavior. Therewith the policy *Data Integrity* is violated. Techniques such as *Data Space Randomization* and *Write Integrity Testing* (WIT) makes it harder to perform such kinds of attacks but they are not established in practice yet.

A.2.3.4 Information Leak

Memory corruption attacks are also used to leak memory contents. Therewith probabilistic countermeasure like ASLR can be circumvented by the knowledge of randomly

generated data. As for the data-only attack *Data Space Randomization* might help to mitigate information leakage.

A.2.4 Existing Solutions

In this section current state-of-the-art solutions of implemented countermeasures mentioned in the sections before are presented. For each approach the system level (compiler, operating system, or hardware) at which an approach is enforced is called and which technical and security requirements are expected.

A.2.4.1 $W \oplus X$, Write xor Execute, Data Execution Prevention (DEP)

System Level: Hardware, Operating System

The idea of this approach is to flag memory pages either writable or executable, but not both at the same time. Therewith code-injected attacks are mitigated. At the lowest system level a mechanism for fine-grained memory page permissions is required. Further the operating system must support the hardware mechanism or even emulate a memory page permission mechanism.

Architecture	Instruction Set	Enforcement
x86	AMD64 [15, p. 56] Intel64	No-eXecution bit (NX-bit), Page Table eXecute Disable (XD-bit), Page Table execute never bit (XN-bit), Page Table PEN privileged execute never, PAN privileged access never
ARM	ARMv6 ARMv8-A	Translation Storage Buffer (optional)
SPARC	Oracle SPARC Architecture 2011	Segment Lookaside Buffer (SLB)
PowerPC	IBM PowerISA [16] [17, p. 33]	

Examples of Hardware Support for Execution Prevention

Family	Name	Implementation
Linux	Linux kernel PaX	Linux kernel mainline \geq 2.6.8 [18] Patch for the Linux kernel, uses hardware support or emulates memory page permission [19]
	Exec Shield	Patch for the Linux kernel, part of Fedora Core 1 through 6 ⁶ and Red Hat Enterprise Linux \geq 3 ^{7,8} , uses hardware support or emulates memory page

⁶https://archives.fedoraproject.org/pub/archive/fedora/linux/core/1/x86_64/os/RELEASE-NOTES.html

⁷<http://people.redhat.com/mingo/exec-shield/>

⁸https://www.redhat.com/f/pdf/rhel/WHP0006US_Execshield.pdf

	grsecurity	Patch for the Linux kernel, uses hardware support or emulates memory page permission ⁹ ¹⁰
Unix	Android	Android \geq 2.3 ¹¹
	OpenBSD	OpenBSD \geq 3.3 ¹²
	NetBSD	NetBSD \geq 2.0 ¹³
	FreeBSD	FreeBSD \geq 5.3

Examples of Operating System Support for Execution Prevention

A.2.4.2 Stack Smashing Protection (SSP)

System Level: Compiler

Stack Smashing Protection (SSP) mechanisms place pseudo-random values (*Stack Canaries* or *Stack Cookie*) on the stack before the saved base pointer and the saved return address as part of the function prologue and compare the value again before a function returns. SSP is enforced at compile-time. Due to performance reasons, the compiler has to decide which function has to be protected. If the compiler implementation makes the wrong decision and the concerned function is vulnerable, SSP fails. Further, information leaks enable an attacker to read the pseudo-random value and integrate it to her exploit so that the buffer is overwritten with the correct value.

Compiler	Option	Description
GNU Compiler Collection (GCC) ¹⁴	<code>-fstack-protector</code>	Emit extra code to check for buffer overflows, such as stack smashing attacks. This is done by adding a guard variable to functions with vulnerable objects. This includes functions that call <code>alloca</code> , and functions with buffers larger than 8 bytes. The guards are initialized when a function is entered and then checked when the function exits. If a guard check fails, an error message is printed and the program exits.
	<code>-fstack-protector-all</code>	Like <code>-fstack-protector</code> except that all functions are protected.
	<code>-fstack-protector-strong</code>	Like <code>-fstack-protector</code> but includes additional functions to be protected – those that have local array definitions, or have references to local frame addresses.

⁹https://en.wikibooks.org/wiki/Grsecurity/Appendix/Grsecurity_and_PaX_Configuration_Options#Enforce_non-executable_kernel_pages

¹⁰https://en.wikibooks.org/wiki/Grsecurity/Appendix/Grsecurity_and_PaX_Configuration_Options#Paging_based_non-executable_pages

¹¹<https://source.android.com/security/#memory-management-security-enhancements>

¹²<http://www.openbsd.org/33.html>

¹³<http://www.netbsd.org/docs/kernel/non-exec.html>

¹⁴<https://gcc.gnu.org/onlinedocs/gcc-7.2.0/gcc/Instrumentation-Options.html#Instrumentation-Options>

	<code>-fstack-protector-explicit</code>	Like <code>-fstack-protector</code> but only protects those functions which have the <code>stack_protect</code> attribute.
Clang ¹⁵	<code>-fstack-protector-all</code> <code>-fstack-protector-strong</code> <code>-fstack-protector</code>	Force the usage of stack protectors for all functions. Use a strong heuristic to apply stack protectors to functions. Enable stack protectors for functions potentially vulnerable to stack smashing.
Intel C++ Compiler ¹⁶	<code>-fstack-security-check</code>	This option determines whether the compiler generates code that detects some buffer overruns that overwrite the return address. This is a common technique for exploiting code that does not enforce buffer size restrictions.
Keil ARM C/C++ Compiler ¹⁷	<code>-protect_stack</code> <code>-protect_stack_all</code>	Use <code>-protect_stack</code> to enable the stack protection feature. Use <code>-no_protect_stack</code> to explicitly disable this feature. If both options are specified, the last option specified takes effect. The <code>-protect_stack_all</code> option adds this protection to all functions regardless of their vulnerability.

Examples of Stack Smashing Protection Compiler Support

A.2.4.3 Address Space Layout Randomization (ASLR)

System Level: Compiler, Operating System

Address Space Layout Randomization (ASLR) obfuscates the address space layout of a process. Therewith the locations of various memory segments get randomized which makes it harder to predict the correct target address which is needed to perform a code-reuse attack. ASLR requires high entropy to prevent brute-force de-randomization attacks and depends on the prevention of unintended Information Leak (Section A.2.3.4) that are used by dynamically constructed exploit payloads. To guarantee high entropy ASLR should be implemented on 64-bit architectures (or above). Additionally, ASLR requires position-independent executables (PIE) which implements a random base address for the main executable binary.

ASLR is enforced by the operating system primarily but requires position-independent executables for binaries and position independent code for shared libraries generated by the compiler.

Family	Name	Implementation
--------	------	----------------

¹⁵<https://clang.llvm.org/docs/ClangCommandLineReference.html#cmdoption-clang-fstack-protector>

¹⁶<https://software.intel.com/en-us/node/523162>

¹⁷http://www.keil.com/support/man/docs/armcc/armcc_chr1359124940593.htm

Linux	Linux kernel	Linux kernel mainline ≥ 3.14 ¹⁸ . ASLR for user-space programs and Kernel Address Space Layout Randomization (KASLR) ¹⁹ for the kernel itself.
	PaX	Patch for the Linux kernel ²⁰
	Exec Shield	Patch for the Linux kernel ²¹
	grsecurity	Patch for Linux kernel ²²
BSD	Android	Android ≥ 4.1 ²³
	OpenBSD	OpenBSD ≥ 4.4 ²⁴
	NetBSD	NetBSD ≥ 5.0 ²⁵
	FreeBSD	FreeBSD as patch ²⁶

Examples of ASLR Operating System Support

Compiler	Option	Description
GNU Compiler Collection (GCC) ^{27,28}	<code>-fpie, -fPIE</code>	These options are similar to <code>-fpic</code> and <code>-fPIC</code> , but generated position independent code can be only linked into executables. Usually these options are used when <code>-pie</code> GCC option is used during linking. This is especially difficult to plumb into packaging in a safe way, since it requires the executable be built with <code>-fPIE</code> for any <code>.o</code> files that are linked at the end with <code>-pie</code> . There is some amount of performance loss, but only due to the <code>-fPIE</code> , which is already true for all the linked libraries (via their <code>-fPIC</code>).
	<code>-fPIC</code>	If supported for the target machine, emit position-independent code, suitable for dynamic linking and avoiding any limit on the size of the global offset table. This option makes a difference on AArch64, m68k, PowerPC and SPARC. Position-independent code requires special support, and therefore works only on certain machines.
Clang ²⁹	<code>-fpie, -fPIE</code>	See GCC.

¹⁸<https://lwn.net/Articles/569635/>

¹⁹http://selinuxproject.org/~jmorris/lss2013_slides/cook_kaslr.pdf

²⁰<https://pax.grsecurity.net/docs/aslr.txt>

²¹http://www.redhat.com/f/pdf/rhel/WHP0006US_Execshield.pdf

²²https://en.wikibooks.org/wiki/Grsecurity/Appendix/Grsecurity_and_PaX_Configuration_Options#Restrict_mprotect.28.29

²³<https://source.android.com/security/enhancements/enhancements41>

²⁴<https://www.openbsd.org/plus44.html>

²⁵<https://netbsd.org/releases/formal-5/NetBSD-5.0.html>

²⁶<https://hardenedbsd.org/content/freebsd-and-hardenedbsd-feature-comparisons>

²⁷<https://gcc.gnu.org/onlinedocs/gcc-7.2.0/gcc/Code-Gen-Options.html#Code-Gen-Options>

²⁸https://wiki.debian.org/Hardening#gcc_-pie_-fPIE

²⁹<https://clang.llvm.org/docs/ControlFlowIntegrity.html>

Intel C++ Compiler ³⁰³¹	-fPIC	See GCC.
	-pie	Determines whether the compiler generates position-independent code that will be linked into an executable.
Keil ARM C/C++ Compiler ³²	-fpic	Determines whether the compiler generates position-independent code.
	-bare_metal_pie	(Bare-metal PIE support is deprecated. There is support for <code>-fropi</code> and <code>-frwpi</code> in <code>armclang</code> . You can use these options to create bare-metal position independent executables.) A bare-metal Position Independent Executable (PIE) is an executable that does not need to be executed at a specific address but can be executed at any suitably aligned address.
	-fropi, -fno-ropi	Enables or disables the generation of Read-Only Position-Independent (ROPI) code.
	-frwpi, -fno-rwpi	Enables or disables the generation of Read/Write Position-Independent (RWPI) code.

Examples of ASLR Compiler Support

A.2.4.4 Control-flow Integrity (CFI)

System Level: Compiler

Control-flow Integrity (CFI) is a current research topic. However, Clang includes an implementation of a number of control flow integrity (CFI) schemes, which are designed to abort the program upon detecting certain forms of undefined behavior that can potentially allow attackers to subvert the program's control flow. These schemes have been optimized for performance, allowing developers to enable them in release builds. The CFI implementation in Clang has a performance overhead of ca. 1% and a binary size overhead of ca. 15% (only forward-edges)³³.

Compiler	Option	Description
----------	--------	-------------

³⁰<https://software.intel.com/en-us/node/523278>

³¹<https://software.intel.com/en-us/node/523158>

³²http://www.keil.com/support/man/docs/armclang_dev/armclang_dev_chr1405439371691.htm

³³<https://clang.llvm.org/docs/ControlFlowIntegrity.html>

GNU Compiler Collection (GCC) ³⁴	<code>-fvtable-verify=[std preinit none]</code>	<p>This option is only available when compiling C++ code. It turns on (or off, if using <code>-fvtable-verify=none</code>) the security feature that verifies at run time, for every virtual call, that the vtable pointer through which the call is made is valid for the type of the object, and has not been corrupted or overwritten. If an invalid vtable pointer is detected at run time, an error is reported and execution of the program is immediately halted.</p> <p>This option causes run-time data structures to be built at program startup, which are used for verifying the vtable pointers. The options <code>std</code> and <code>preinit</code> control the timing of when these data structures are built. In both cases the data structures are built before execution reaches <code>main</code>. Using <code>-fvtable-verify=std</code> causes the data structures to be built after shared libraries have been loaded and initialized. <code>-fvtable-verify=preinit</code> causes them to be built before shared libraries have been loaded and initialized.</p> <p>If this option appears multiple times in the command line with different values specified, <code>none</code> takes highest priority over both <code>std</code> and <code>preinit</code>; <code>preinit</code> takes priority over <code>std</code>.</p>
Clang ³⁵	<code>-fsanitize=cfi-cast-strict</code> <code>-fsanitize=cfi-derived-cast</code> <code>-fsanitize=cfi-unrelated-cast</code> <code>-fsanitize=cfi-nvcall</code>	<p>Enables strict cast checks.</p> <p>Base-to-derived cast to the wrong dynamic type.</p> <p>Cast from <code>void*</code> or another unrelated type to the wrong dynamic type.</p> <p>Non-virtual call via an object whose <code>vptr</code> is of the wrong dynamic type.</p>

³⁴<https://gcc.gnu.org/onlinedocs/gcc-7.2.0/gcc/Code-Gen-Options.html#Code-Gen-Options>

³⁵<https://clang.llvm.org/docs/ControlFlowIntegrity.html>

<code>-fsanitize=cfi-vcall</code>	Virtual call via an object whose vptr is of the wrong dynamic type.
<code>-fsanitize=cfi-icall</code>	Indirect call of a function with wrong dynamic type.
<code>-fsanitize=cfi</code>	Enable all the schemes.

Examples of CFI Compiler Support

A.2.4.5 Code Pointer Integrity (CPI), Code Pointer Separation (CPS)

System Level: Compiler

Code-Pointer Integrity (CPI) is a property of C/C++ programs that guarantees absence of control-flow hijack attacks by requiring integrity of all direct and indirect pointers to code. Code-Pointer Separation (CPS) is a simplified version of CPI that provides strong protection against such attacks in practice. SafeStack is a component of CPI/CPS, which can be used independently and protects against stack-based control-flow hijacks.

CPI/CPS/SafeStack can be automatically enforced for C/C++ programs through compile-time instrumentation with low performance overheads of 8.5% / 1.9% / 0.05% correspondingly. The SafeStack enforcement mechanism is now part of the Clang compiler, while CPI and CPS are available as research prototypes. For the current status please see [20].

Compiler	Option	Description
Clang ³⁶³⁷	<code>-fsanitize=safe-stack</code>	SafeStack is an instrumentation pass that protects programs against attacks based on stack buffer overflows, without introducing any measurable performance overhead. It works by separating the program stack into two distinct regions: the safe stack and the unsafe stack. The safe stack stores return addresses, register spills, and local variables that are always accessed in a safe way, while the unsafe stack stores everything else. This separation ensures that buffer overflows on the unsafe stack cannot be used to overwrite anything on the safe stack.

Examples of CPI and CPS Compiler Support

³⁶<http://dslab.epfl.ch/proj/cpi/>

³⁷<https://clang.llvm.org/docs/SafeStack.html>

A.2.4.6 Pointer Authentication

System Level: Hardware, Compiler

Pointer Authentication uses cryptographic functions to authenticate and verify pointers before dereferencing [12]. Pointers are replaced by a generated signature and cannot be dereferenced directly. This requires compiler and instruction set support. ARM recently introduces the Pointer Authentication extensions in ARMv8.3-A specification³⁸. The GCC introduces basic support for pointer authentication in version 7³⁹ for the *AArch64* target (ARMv8.3-A architecture). The overhead is negligible because of hardware acceleration for the cryptographic primitives⁴⁰.

Compiler	Option	Description
GNU Compiler Collection (GCC) ⁴¹	<code>-msign-return-address=scope</code>	Select the function scope on which return address signing will be applied. Permissible values are <code>none</code> , which disables return address signing, <code>non-leaf</code> , which enables pointer signing for functions which are not leaf functions, and <code>all</code> , which enables pointer signing for all functions. The default value is <code>none</code> .

Examples of Pointer Authentication Compiler Support

A.3 Isolation

Isolating software components within a system is a common protection measure to protect other components from erroneous ones, either through unintentional programming errors, or intentional harm caused by an attacker taking over a corruptible component. While the protective measures detailed in section A.2 are a preventive measure, intended to impede an attacker from taking over a software component by exploiting programming errors, isolation intends to limit the influence an attacker might have to other software components *after* taking over a software component. As such, isolation is only an effective measure, if the architecture of the system is appropriately designed, dividing and isolating different functional aspects of the system accordingly. Note that this guide does not cover this architectural aspect of isolation, but the technical aspect, i.e. *how* isolation can be implemented.

The following sections describe two approaches to isolation: the isolation of multiple software components between each other, and the isolation of software components and the operating system itself. These two approaches are orthogonal – i.e. horizontal

³⁸<https://community.arm.com/processors/b/blog/posts/armv8-a-architecture-2016-additions>

³⁹<https://gcc.gnu.org/gcc-7/changes.html>

⁴⁰<https://lwn.net/Articles/719270/>

⁴¹<https://gcc.gnu.org/onlinedocs/gcc-7.1.0/gcc/AArch64-Options.html#AArch64-Options>

isolation between applications and vertical isolation between applications and the OS – and can be combined accordingly.

A.3.1 Horizontal Isolation

A.3.1.1 Virtual Memory

The oldest and most prevalent isolation mechanism is the concept of virtual memory, i.e. presenting each running software component of a system with its own virtual address space, which is mapped to the available physical memory by the operating system. The origins date back to the 1950s, with the original intention of hiding the fragmentation of physical memory, but offers the possibility of isolating software components. As each component operates in a virtual address space, it cannot access the memory of other components (unless explicitly allowed by the operating system). This feature requires hardware support, e.g. by a Memory Management Unit (MMU), but this support is nearly ubiquitous in most computer systems except very small micro-controllers.

An extension of this concept is that of *virtual machines (or virtualization)*, whereby multiple virtual machines (VM) are emulated by a *hypervisor*. Each VM may run a complete operating system, depending on the degree of virtualization even in an unmodified state. The hypervisor controls the access of each VM to the physical hardware components, or even emulate certain components such as network interfaces. Similar to virtual memory, the virtualization requires dedicated hardware support to achieve an appropriate level of performance and security.

Note that both approaches have limitations. The isolation provided by virtual memory or virtualization is only as strong as the operating system or hypervisor itself, as a malicious application might take control over the OS or hypervisor through a programming error (the measures described in section [A.4](#) intend to minimize this attack surface). Similarly, a malicious application might use its access to other hardware components to circumvent the isolation, as some hardware components may have unrestricted access to the systems memory. “IOMMUs”, a technique which presents hardware components with a virtual address space, can be used to counter this. Lastly, an attacker might use the volatile properties of physical memory itself to circumvent the isolation. For example, in [21] the “rowhammer” attack is described, which is capable of flipping bits in memory locations usually inaccessible to an application. The attack uses prolonged reads to a memory location performed in quick succession, which in the case of “DRAM” memory will cause neighboring memory cells to change state. This effect has been shown to be capable of raising the privileges of an application in Linux and Android (cf. [22]) systems. System designers must consider using one or more mitigations to such attacks, for example, by using memory with error correction, or software mitigations as shown in [23].

A.3.2 OS-Level Virtualization

A more lightweight form of virtualization, often called operating-system-level virtualization or *containerization*, is available in modern OS. Prominent examples are “LXC”⁴², which is built on top of the Linux Kernel Namespacing functionality, or the “Jail”⁴³ functionality provided by the FreeBSD operating system. These tools only virtualize certain resources of an operating system, for example, file system, the process tree, or the network stack. This way, the operating system creates a *container* with a tightly controlled access to system resources or other containers. In contrast to full virtualization, these containers cannot execute a different operating system, albeit a completely separate user-space instance can be run side-by-side.

OS Family	Solution	Description
UNIX	chroot	This is a system call available in many UNIX alike systems, which can be used to set up a an execution environment with a different root filesystem. As such, it only isolates the filesystem of the host from the container. If the container contains privileged processes, they can easily affect the system
Linux	Namespaces, LXC, Docker, systemd-nspawn	Many solutions building upon the Linux Namespacing functionality exist, many allowing the setup of isolated containers. Resources of other containers or the host are not visible to the processes running inside these containers, unless assigned. This way, even privileged processes inside a container are not capable of affecting the rest of the system.
FreeBSD	Jails	A solution similar to chroot, except for improved security. For example, these containers (or called “jails” in this context) offer an isolated network, as well as restricting the capabilities of privileged processes inside (e.g. privileged processes cannot affect the rest of the system).

Operating System Virtualization / Container Implementations

A.4 Vertical Isolation

Isolating the operating system from applications, often called *sandboxing* is an another important aspect of a protected runtime environment. The basic idea is to limit the

⁴²<https://linuxcontainers.org/>

⁴³<https://www.freebsd.org/doc/handbook/jails.html>

capabilities of a process, i.e. restricting what a process can do⁴⁴. The classic way to do this is by dropping privileges as soon as they are not needed anymore, i.e. *privilege revocation*. For example, the `ping` command requires root privileges on UNIX systems to create a raw network socket, but will drop its privileges to a regular user after creating it. Ideally, a software component should drop (or never be given) any privileges it does not require, or drop them as soon as they are no longer required. As with the example of `ping`, any software component must then be structured with a setup phase, in which all advanced privileges are used and subsequently dropped. The following shows a few examples of operating system functionalities, which allow an application to drop capabilities or privileges.

OS Family	Solution	Description
Linux	Seccomp Mode 1 (Strict) [24]	Processes on the Linux operating system can limit their set of allowed system calls in a very easy manner. The allowed subset is extremely strict, limiting the process to <code>read</code> , <code>write</code> , <code>exit</code> and <code>sigreturn</code> . Once activated, this Seccomp mode cannot be deactivated again.
Linux	Seccomp Mode 2 (Filter) [24]	A more recent version of the Seccomp mode, the seccomp filter mode, allows a much more fine-grained control. The concept is to allow a process to attach a small filter program, which will check each system call. This filter program must be a “Berkley Packet Filter” (BPF), a very restricted form of byte-code, which will be executed by the kernel before each system call. This filter can then examine each system call, for example, check which system call is made or what the parameters are set. The filter then returns a decision as to what the kernel should do. The system call can be allowed or blocked and if the call is blocked, several choices can be made as to how it is blocked. The filter may decide to immediately kill the process, to simply let the system call return an error, send a signal to the offending process or to notify a tracer attached to the program (such as a debugger). A notable property of these filters is, that they are inherited by the spawned child processes, which enables a setup of a filter before a potentially dangerous process is started.
FreeBSD	Securelevel [25, chapter 13]	This functionality limits the possibilities of all processes running on a system in incremental levels, which cannot be lowered once entered (until a reboot of the system).

⁴⁴Not to be confused with “what a process does”, as the behaviour of a process is changed by an attack.

OpenBSD

Pledge [26]

This functionality can be used to limit the system calls a process can execute. Certain system calls can be prohibited completely, others can be restricted in their functionality. For example, the `open` system call can be limited to only open a small set of system files, or the `mprotect` cannot set memory to be executable. Once the limitations are in place, the process cannot lift them again.

Sandboxing Mechanisms in Operating Systems