

Muen - Toolchain

Reto Buerki Adrian-Ken Rueegsegger

January 25, 2018

無縁

University of Applied Sciences Rapperswil (HSR), Switzerland

Contents

1	Introduction	4
2	Policy	4
2.1	Content	5
2.1.1	Configuration Values	5
2.1.2	Hardware Resources	5
2.1.3	Platform description	5
2.1.4	Kernel diagnostics device	5
2.1.5	Physical Memory Regions	5
2.1.6	Device Domains	6
2.1.7	Events	6
2.1.8	Communication Channels	6
2.1.9	Components	6
2.1.10	Subjects	6
2.1.11	Scheduling Plans	7
2.2	Format	7
2.2.1	Source Format	7
2.2.2	Format A	7
2.2.3	Format B	8
3	Build Process	8
3.1	Policy Merging	8
3.2	Components Build	10
3.3	Components Specification Merging	10
3.4	Policy Compilation	10
3.5	Policy Validation	11
3.6	Structure Generation	11
3.7	Image Packaging	11
4	Core Tools	12
4.1	Merger	12
4.2	Component Joiner	12
4.3	Expander	13
4.4	Allocator	13
4.5	Validator	14
4.6	Structure Generators	14
4.6.1	Page Tables	14
4.6.2	VT-d Tables	15
4.6.3	I/O Bitmaps	15
4.6.4	MSR Bitmaps	16
4.6.5	MSR Stores	16
4.6.6	ACPI Tables	16
4.6.7	Linux Zero Pages	17
4.6.8	Solo5/Ukvm boot info	17
4.6.9	Source Specifications	18
4.6.10	Component Source Specifications	18
4.6.11	Subject Info	18
4.7	Hasher	19

4.8	Packer	19
5	Additional Tools	19
5.1	Kernel ELF Checker	20
5.2	Stack Usage Checker	20
5.3	Hardware Config Generator	20
5.4	Scheduling Plan Generator	20
5.5	Component Binary Splitter	21

List of Figures

1	Build process	9
---	-------------------------	---

1 Introduction

This document describes the process of configuring and building a component-based system running on the Muen Separation Kernel (SK).

2 Policy

A policy is a description of a component-based system running on top of the Muen Separation Kernel. It defines what hardware resources are present, how many active components (called subjects) the system is composed of, how they interact and which system resources they are allowed to access. The following properties are specified by the policy:

- Configuration values
- Hardware resources
- Platform description
- Kernel diagnostics device
- Physical memory regions
- Device domains
- Events
- Communication channels
- Components
- Subjects
- Scheduling plans

The policy serves as a static description of a Muen system. Since all aspects of the system are fixed at integration time the policy can be validated prior to execution, see also section 3.5.

2.1 Content

This section presents the different parts of a system policy and gives an overview what each section contains.

2.1.1 Configuration Values

The purpose of a config section is to specify configuration values which parameterize a system or a component. It allows to declare boolean, string and integer values, e.g. `<boolean name="iommu_enabled" value="true"/>`.

The following sections provide support for configuration values:

- System
- Platform
- Component

During the build process, configuration values provided by the platform are merged into the global system configuration. Component configuration values allow the parameterization of component-local functionality.

2.1.2 Hardware Resources

Systems running the Muen SK perform static resource allocation at integration time. This means that all available hardware resources of a target machine must be defined in the system policy in order for these resources to be allocated to subjects.

Data required by a hardware description includes the amount of available physical memory blocks including reserved memory regions (RMRR), the number of logical CPUs and hardware device resources.

The Muen toolchain provides a handy tool to automate the cumbersome process of gathering hardware resource data, see section 5.3.

2.1.3 Platform description

To enable an uniform view of the hardware resources across different physical machines from the system integrators perspective, the platform description layer is interposed between the hardware resource description and the rest of the system policy. This allows to build a Muen system for different physical target machines using the same system policy.

The config section enables the declaration of platform-specific properties.

2.1.4 Kernel diagnostics device

The Muen SK can be instructed to output debugging information during runtime. The kernel diagnostics device specifies which I/O device the kernel is to use for this purpose.

2.1.5 Physical Memory Regions

This part of the policy specifies the physical memory layout of the system. Memory regions are defined by their size, caching, type and are placed by specifying a physical address. Additionally the content of the region can be declared as backed by a file or filled with a pattern.

2.1.6 Device Domains

The physical memory accessible by PCI devices is specified by so called device domains. Such domains define virtual mappings of physical memory regions for one or multiple devices. Device references select a subset of hardware devices provided by the platform.

Device domains are isolated from each other by the use of Intel VT-d. Thus they can only be specified and enforced on systems that provide at least one IOMMU¹.

2.1.7 Events

Events are an activity caused by a subject (source) that impacts a second subject (target) or is directed at the kernel. Events are declared globally and have a unique name to be unambiguous.

Subjects can use events to either deliver an interrupt, hand over execution to or reset the state of a target subject. The first kind of event provides a basic notification mechanism and enables the implementation of event-driven services. The second type facilitates suspension of execution of the source subject and switching to the target. Such a construct is used to pass the thread of execution on to a different subject, e.g. invocation of a debugger subject if an error occurs in the source subject. The third kind is used to facilitate the restart of subjects.

Kernel events are special in that they are targeted at the kernel. The currently supported events are system reboot and shutdown.

2.1.8 Communication Channels

Inter-subject communication is represented by so called channels. These channels represent directed information flows since they have a single writer and possibly multiple readers. Optionally a channel can have an associated notification event (doorbell interrupt).

Channels are declared globally and have a unique name to be unambiguous.

2.1.9 Components

A component is a piece of software executed by the SK. Similar terms are partition or container. They represent the building blocks of a component-based system.

The description of a component specifies the binary program file including the virtual memory location as well as the view of the expected execution environment. This environment is defined in terms of logical resources such as for example communication channels.

2.1.10 Subjects

Subjects are instances of components. A subject specification references a component and maps the declared logical resources to physical resources provided by the system.

Besides the component resource mappings, subjects can specify extra resources such as device and/or memory mappings. This is useful for subjects which are able to enumerate the available resources at runtime via configuration mechanisms like ACPI or the Subject Information Page.

Furthermore, subject specifications enable the declaration of events a subject is allowed to trigger and receive.

Subjects also have an associated profile (e.g. native or Linux) which determines properties of the execution environment provided by the kernel.

¹Input/Output Memory Management Unit

2.1.11 Scheduling Plans

The Muen SK performs scheduling of subjects in a fixed, cyclic and preemptive way according to a user-specified regime. Scheduling information is declared in so called scheduling plans. They specify in what order subjects are executed on which logical CPU and for how long. Multiple scheduling plans can be specified to enable the definition of different system execution profiles which can be switched during runtime.

A scheduling plan is specified in terms of frames. A *major frame* consists of a sequence of minor frames. When the end of a major frame is reached, the scheduler starts over from the beginning and uses the first minor frame in a cyclic fashion. This means that major frames are repetitive. A *minor frame* specifies a subject and a precise amount of time.

The `mugenschedcfg` tool can be used to automatically generate scheduling plans from a given scheduling configuration, see section 5.4.

2.2 Format

The system policy is specified in XML. There are currently three different policy formats:

- Source Format
- Format A
- Format B

The motivation to have several policy formats is to provide abstractions and a compact way for users to specify a system while simultaneously facilitate reduced complexity of tools operating on the policy.

The implementation of such tools is simplified by the absence of higher-level abstractions which would make the extraction of input data more involved. As an example, the page table generation tool can directly access all virtual memory mappings of a subject and must not concern itself with channels. The channel abstraction has already been broken down into the corresponding memory elements during the policy compilation step of the build process (see section 3.4).

2.2.1 Source Format

The user-specified policy is written in the source format. Constructs such as channels or events provide abstractions to simplify the specification of component-based systems. Many XML elements and attributes are optional and will be filled in with default values during later steps of the policy compilation process.

Kernel and $\tau 0$ resources are not part of the source format since they are also automatically added by the policy expansion step.

Additionally the use of configuration values enables easy parametrization of the system policy.

2.2.2 Format A

Format A is a processed version of the source format where all includes are resolved and abstractions such as channels have been broken down into their underlying elements. For example, a channel is expanded to a physical memory region and the corresponding writer and reader subject mappings with the appropriate access rights.

In this format all implicit elements, such as for example automatically generated page table memory regions, are specified. The kernel and $\tau 0$ configuration is also declared as part of format A.

The only optional attributes are addresses of physical memory regions.

2.2.3 Format B

Format B is equivalent to Format A except that all physical memory regions have a fixed location (i.e. their physical address is set).

3 Build Process

The build of a system is divided into the following steps:

- Policy merging
- Components build
- Components spec merging
- Policy compilation
- Policy validation
- Structure generation
- Image packaging

The toolchain is composed of several tools that operate on a user-specified system policy. Following the Unix philosophy "A program should do only one thing and do it well" each of the tools performs a specific task. They work in conjunction to process a user-defined policy and build a bootable system image. An in-depth description of the involved tools is given in section 4 while figure 1 gives an overview of the whole build process.

3.1 Policy Merging

The Merger tool outlined in section 4.1 is responsible to merge XML files stored at different locations on the file system into one system policy in source format.

The tool reads a system configuration in XML format to locate the following files:

- System policy
- Hardware specification
- Additional hardware specification
- Platform specification

The tool also provides an implementation of the XML XInclude mechanism². Using includes, the policy writer is able to separate and organize the system policy as desired. Instead of specifying the whole policy in one file, the subject specifications can be split into separate files,

²<http://www.w3.org/TR/xinclude-11/>

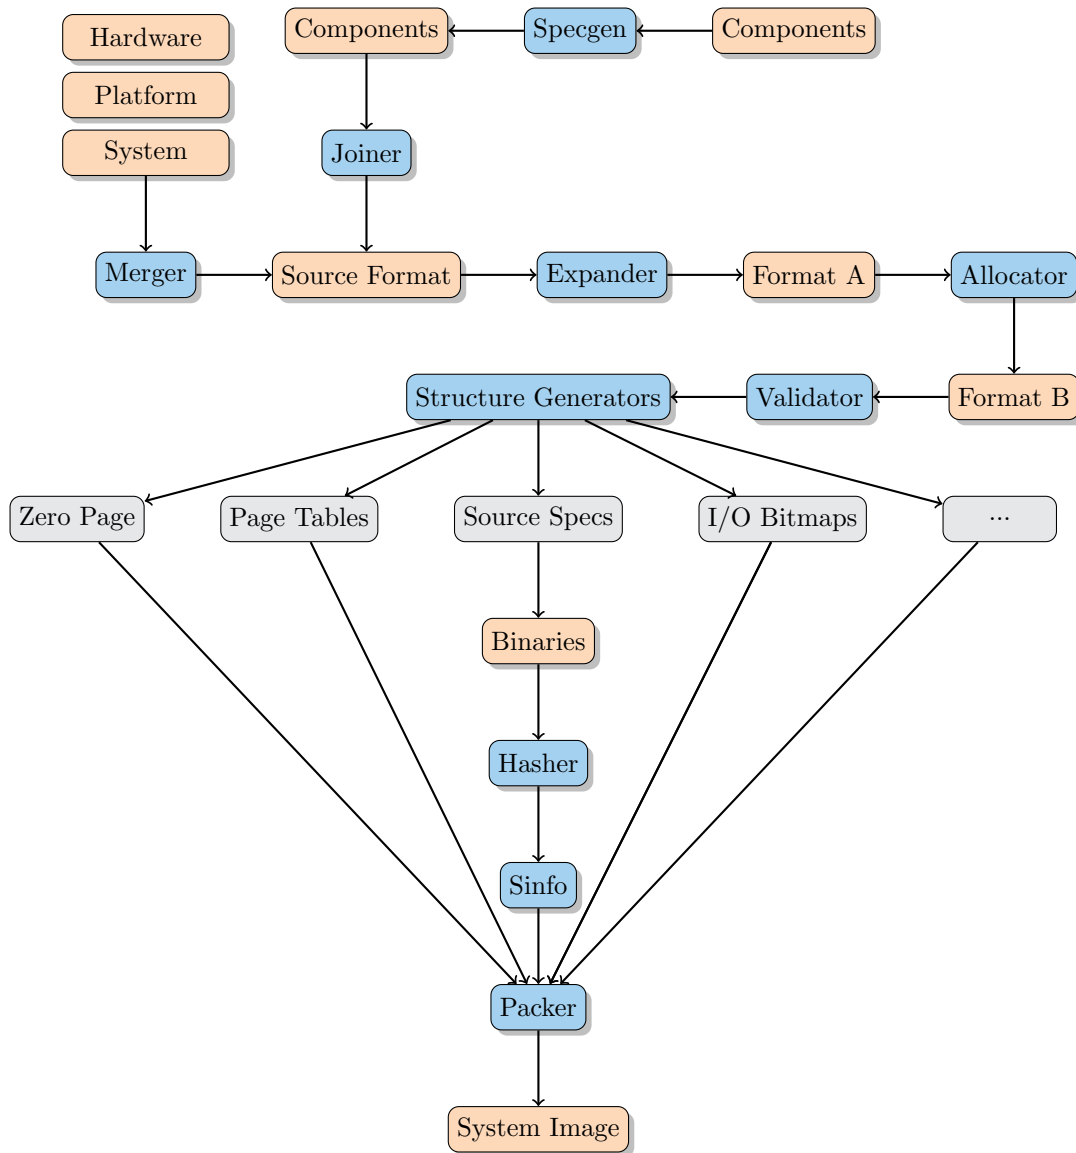


Figure 1: Build process

or common parts shared by different system descriptions can be extracted. See section 4.1 for more information about the Merger tool.

After the merge step, the resulting policy is well formatted to minimize the difference in the generated policies resulting from the subsequent tasks. This allows the user to easily review (`diff`) and therefore verify the results of each policy compilation task.

Expressions can be used to formulate (nested) boolean terms using the numeric equality/inequality and logical operators. They are evaluated to boolean config values prior to processing conditionals.

The use of conditionals enables selective activation of parts of the source policy depending on

the value of a given config variable. This allows flexible customization of a system during policy compilation time by setting the value of a config variable or formulating an appropriate boolean expression.

Config variable substitution enables the policy writer to set the value of attributes to those of referenced config values. Attributes that start with a dollar sign followed by a config value name are substituted by the value of the config variable.

3.2 Components Build

After hardware, platform and high-level system policy are merged, components can extract relevant information from it using whatever means suit best. For example an XSL transformation (XSLT) script could extract the I/O port of a specific device and create a corresponding configuration value based on it, which is then included in the component specification.

The `mucgenspec` tool described in 4.6.10 implements the blue *Specgen* task shown in figure 1. It is used to process component specifications and, similar to the policy merger, supports conditionals, expressions and configuration value substitutions. It also generates Ada/SPARK packages containing constants derived from the declared component resources and config values. These constants can be used to reliably address specific or configurable resources in the source code.

After the component specification has been processed, the component source code is compiled into a binary.

The `mucbinsplit` tool described in 5.5 can be used to extract ELF sections of the component binary into separate files. It automatically extends the component specification by adding a corresponding memory region with the appropriate access rights (e.g. executable, writable).

3.3 Components Specification Merging

The processed component specifications are merged into the system source policy by the Muen component specification joiner tool described in section 4.2.

This step is optional as static component specifications which need no processing can also be manually specified in the system policy directly.

3.4 Policy Compilation

Policy compilation encompasses the tasks involved to transform the policy from source format to format A and finally to format B, which is the fully expanded format with no implicit properties.

The Expander tool takes care of completing the user-specified policy with additional information and abstractions only available in format source are resolved to low-level constructs.

For example, the concept of *channels* only exists in format source. Therefore a channel specified in format source must be expanded to shared memory regions with optional associated events in format A. Also, the Expander tool inserts specifications for the Muen kernel itself so the user is lifted from that burden. Generally, the aim of the expansion task is to make the life of a policy writer as easy as possible by expanding all information which can be derived automatically. Section 4.3 explains the Expander tool in detail.

The result of the expansion task is a policy in format A which is the input for the Allocator tool. This tool is responsible to assign a physical memory address to all memory regions which are not already explicitly stated. By querying the hardware section of the policy, the tool is aware of the total amount of available RAM on a specific system and allocates regions of it for memory elements with no explicit physical address. The Allocator tool also implements optimization

strategies to keep the resulting system image as small as possible. For example, file-backed memory regions (e.g. a memory region storing a component executable) are preferably placed in lower physical regions. See section 4.4 for a description of the Allocator tool.

After the allocation task is complete, the policy is stored in format B. This format states all system properties explicitly and is used as input for the Validation step discussed in the following section.

3.5 Policy Validation

Before structures required to pack the final system image are generated, the policy must be thoroughly validated to catch errors in the system specification. Such errors might range from overlapping memory, undefined resource references to incomplete scheduling plans etc. The Validator task performs checks that assure the policy in format B is sound and free from higher-level errors that are not covered by XML schemata restrictions.

It is important to always run the Validator as the system could otherwise exhibit unexpected behavior. This is especially true if a policy writer decides to specify the system directly in format B which is also possible but not advised. Section 4.5 explains the Validator tool and lists some example checks performed by the tool as illustration.

3.6 Structure Generation

The structure generation step encompasses various tools which extract information from a policy in format B and generate files in different formats (see figure 1).

While some generated files are directly linked into the Muen kernel (i.e. Source Specs, see 4.6.9), most of them are packed into the final system image by the packer tool.

For example, the tool responsible to generate page table structures queries memory mappings and the associated physical memory regions from the policy and creates page table structures in accordance to the format specified by the Intel Software Developer's Manual (SDM). The resulting files are packed into the system image and only applied by the kernel. The kernel itself does not care about memory management, all required tables are pre-built during system integration.

For more information about the structure generators, see section 4.6.

3.7 Image Packaging

The Packer tool assembles the final system image by first allocating a memory buffer which is initialized to zero. The size of the buffer is large enough to hold the complete system image, which consists of all file-backed memory regions specified in the policy:

- Kernel binary
- Kernel page tables
- I/O bitmaps
- MSR bitmaps
- MSR store
- Subject binaries
- Subject page tables

- VT-d tables
- ACPI tables for VM subjects
- Initial Ramdisks for Linux subjects
- Zero Page structures for Linux subjects
- Solo5/Ukvm boot info structures for MirageOS subjects
- Muen subject information structures

It then simply iterates over all file-backed memory regions and inserts the contents of the specified files into the allocated buffer. After performing various post-checks on the created image, it is written to a file. The resulting image can be booted by any Multiboot³ compliant bootloader.

For more information about the Packer tool, see section 4.8.

4 Core Tools

This section describes the tools which form the core of the Muen toolchain.

4.1 Merger

The Merger combines user-provided system policy files into a single XML document.

Name

`mucfgmerge`

Input

System configuration as XML, Colon-separated list of include paths

Output

System policy in format source (merged)

This tool reads the system configuration and merges the specified system policy, hardware and platform files into a single file. It evaluates boolean expressions, resolves conditional parts of the policy and substitutes attribute configuration value references. Included files are inserted at the corresponding locations in the merged file. The XML content is re-formatted so changes to the policy by subsequent build steps can be manually reviewed or visualized by diffing the files.

It also merges the platform configuration section (if any) into the global configuration section, removing the platform configuration section in the process.

4.2 Component Joiner

The Muen component specification joiner adds component XML specifications to the component section of a specified system policy and writes the result to a specified output file. Each given component/library specification is loaded and validated against the components XML schema. If it is correct the content is added to the components section of the system policy specified as input file. If the given system policy does not yet contain a components section, it is created. The result is written to the file specified by the `-o` parameter. In-place processing is supported by passing in the same value for input and output file.

³<https://www.gnu.org/software/grub/manual/multiboot/multiboot.html>

Name

mucfgcjoin

Input

System policy in format source, comma-separated list of component specs

Output

System policy in format source (joined)

4.3 Expander

The expander completes the user-provided system policy by creating or deriving additional configuration elements.

Name

mucfgexpand

Input

System policy in format source

Output

System policy in format A (expanded)

The Expander performs the following actions:

- Pre-check the system policy to make sure it is sound
- Expand channels
- Expand device resources
- Expand device isolation domains
- Expand kernel sections
- Expand τ_0 subject
- Expand additional memory regions
- Expand hardware-/platform-related information
- Expand additional subject information
- Expand profile-specific information
- Expand scheduling information
- Post-check resulting policy

4.4 Allocator

The Allocator is responsible to assign a physical address to all global memory regions.

Name

mucfgalloc

Input

System policy in format A

Output

System policy in format B (allocated)

First, the Allocator initializes the physical memory view of the system based on the physical memory blocks specified in the XML hardware section. It then reserves memory that is occupied by pre-allocated memory elements (i.e. memory regions with a physical address or device memory). Finally it places all remaining memory regions in physical memory. In order to reduce the size of the final system image file-backed memory regions are placed at the start of memory.

4.5 Validator

The Validator performs additional checks that go beyond the basic restrictions imposed by the XML schema validation. Currently over 110 checks are performed.

Name

`mucfgvalidate`

Input

System policy in format B

Output

None, raises exception on error

Examples of checks include:

- Assert that references between policy elements are correct (e.g. a physical memory region referenced by a virtual memory region exists)
- Assert that memory regions do not overlap
- Assert that device interrupts are unique
- Assert that no subject has access to system or kernel memory
- Assert that scheduling plan major frames have the same length on each CPU

4.6 Structure Generators

These tools do not change the policy and use it read-only.

4.6.1 Page Tables

Generate page tables for kernel(s) and subjects.

Name

`mugenpt`

Input

System policy in format B

Output

Page tables of kernels and subjects in binary format

Output format

- IA32-e paging structures, Intel SDM Vol. 3A, section 4.5
- EPT paging structures, Intel SDM Vol. 3C, section 28.2

The tool generates paging structures for subjects and kernels running on each CPU. These page tables are used to grant access to physical memory according to the virtual memory layout of the subject. The rest of physical and device memory is isolated from the subject.

An IA32-e page table is generated for each kernel running on a logical, active CPU. Depending on the subject profile either native 64-bit IA32-e or Extended Page Tables (EPT) are generated.

Page tables are used by the memory management unit (MMU) to enforce isolation of physical memory according to the system policy.

4.6.2 VT-d Tables

Generate VT-d tables for each device isolation domain.

Name

mugenvtd

Input

System policy in format B

Output

VT-d tables of device domains in binary format

Output format

VT-d tables according to Intel VT-d specification, section 9

The tool creates root, context and second-level address translation tables for Intel VT-d DMAR (DMA⁴ remapping) hardware (see Intel VT-d specification, section 3). DMAR is used to restrict direct hardware device access to physical memory via DMA. Devices are put in so-called device security or device isolation domains and are only allowed to access physical memory as granted by the policy.

Interrupt remapping tables are also generated for Intel VT-d IR to ensure that physical devices can only generate interrupt requests as specified by the system policy.

4.6.3 I/O Bitmaps

Generate I/O bitmaps for each subject.

Name

mugeniobm

Input

System policy in format B

Output

I/O bitmaps of subjects in binary format

Output format

Intel SDM Vol. 3C, section 24.6.4

The tool generates I/O bitmaps for each subject. Access to device I/O ports is granted according to the device I/O port resources assigned to a subject.

I/O bitmaps are used by the hardware (VT-x) to enforce access to I/O ports according to the system policy.

⁴DMA - Direct Memory Access

4.6.4 MSR Bitmaps

Generate MSR bitmap for each subject.

Name

mugenmsrbm

Input

System policy in format B

Output

MSR bitmaps of subjects in binary format

Output format

Intel SDM Vol. 3C, section 24.6.9

The tool generates MSR bitmaps for each subject. Access to Model-Specific Registers (MSRs) is granted according to the MSRs assigned to a subject.

MSR bitmaps are used by the hardware (VT-x) to enforce access to Model-Specific Registers according to the system policy.

4.6.5 MSR Stores

Generate MSR store for each subject with MSR access.

Name

mugenmsrstore

Input

System policy in format B

Output

MSR store files of subjects in binary format

Output format

Intel SDM Vol. 3C, table 24-11

The tool generates MSR stores for each subject. The MSR store is used to save/load MSR values of registers not implicitly handled by hardware on subject exit/resumption.

MSR stores are used by hardware (VT-x) to enforce isolation of MSR (i.e. subjects that have access to the same MSRs cannot transfer data via these registers).

4.6.6 ACPI Tables

Generate ACPI tables for all Linux subjects.

Name

mugenacpi

Input

System policy in format B

Output

ACPI tables of all Linux subjects

Output format

Advanced Configuration and Power Interface (ACPI) Specification⁵

ACPI tables are used to announce available hardware to VM subjects. A set of tables consists of an RSDP, XSDT, FADT and DSDT table. See the ACPI specification for more information about a specific table.

⁵<http://www.acpi.info/DOWNLOADS/ACPIspec50.pdf>

4.6.7 Linux Zero Pages

Generate Zero Pages for all Linux subjects.

Name

mugenzp

Input

System policy in format B

Output

Zero pages of all Linux subjects

Output format

Linux Boot Protocol⁶

Zero Page⁷

The so-called Zero Page (ZP) exports information required by the boot protocol of the Linux kernel on the x86 architecture. The kernel uses the provided information to retrieve settings about its running environment:

- Type of bootloader
- Map of physical memory (e820 map)
- Address and size of initial ramdisk(s)
- Kernel command line parameters

4.6.8 Solo5/Ukvm boot info

Generate Solo5/UKVM boot info structures for Mirage unikernels running on the Solo5 platform with the unikernel monitor backend.

Name

mugenukvm

Input

System policy in format B

Output

Solo5/Ukvm boot info for all MirageOS subjects

Output format

struct ukvm_boot_info⁸

The Ukvm boot info structure exports information required by Solo5. The unikernel uses the provided information to retrieve settings about its runtime environment:

- Memory size in bytes
- Address of end of unikernel
- Command line parameters
- TSC frequency

⁶<https://www.kernel.org/doc/Documentation/x86/boot.txt>

⁷<https://www.kernel.org/doc/Documentation/x86/zero-page.txt>

⁸https://github.com/Solo5/solo5/blob/master/ukvm/ukvm_guest.h

4.6.9 Source Specifications

Generate source specifications used by kernel and subjects.

Name

`mugenspec`

Input

System policy in format B

Output

Source specifications in SPARK, C and GPR format

Gathers data from the system policy to generate various source files in SPARK, C and GNAT project file (GPR) format. Created output includes constant values for memory addresses, device resources, scheduling plans, etc.

4.6.10 Component Source Specifications

Process component description and generate source specifications from it. Write processed description to specified output file.

Name

`mucgenspec`

Input

Component description in XML, colon-separated list of include paths

Output

Component source specifications in SPARK, processed component description in XML

The component spec generation tool processes the given component description by evaluating XIncludes, boolean expressions and resolving conditional parts. Furthermore, it performs substitutions of attributes with configuration values.

It also generates Ada/SPARK packages containing constants of the declared logical component resources. The generated specifications can be used in the component source code to access the declared resources.

The resulting processed component description is written to the given output location.

4.6.11 Subject Info

Generate subject information data for each subject.

Name

`mugensinfo`

Input

System policy in format B

Output

Subject info data in binary format

Output format

As specified in `common/musinfo/musinfo.ads`

The Sinfo page is used to export subject information data extracted from the system policy to VM subjects. Currently, information about available memory regions, communication channels and assigned PCI devices is provided.

4.7 Hasher

The Mucfgmemhashes tool is used to add memory integrity hashes to a given policy.

Name

mucfgmemhashes

Input

System policy in format B

Output

System policy in format B with memory integrity hashes

The Mucfgmemhashes tool appends a hash to all memory regions with fill and file content. It must run after all files have been generated by the structure generator tools.

The actual hash is generated using the SHA-256 algorithm and is intended to be used to verify the integrity of memory regions during runtime.

Note that no hashes are generated for sinfo memory regions. Since the hash information will be exported via sinfo, and the sinfo region is itself part of the memory information of a subject, this hash would be self-referential.

The tool also replaces all occurrences of `hashRef` elements. A hash reference element instructs the tool to copy the hash element of the referenced memory region after message digest generation.

From an abstract point of view, the `hashRef` element is a way to link multiple memory regions by declaring that the hash of the content is the same. The hash may serve as an indicator on how to reconstruct the (initial) content of a memory region. This mechanism is heavily used by the subject loader (SL) during subject init and reset operation. The subject loader expander remaps writable memory regions of the loadee (the subject under loader control) to SL and replaces the original regions with new ones containing a hash reference to the associated physical memory region. This way SL is able to determine the intended content of the target memory region by looking up the region in its sinfo page by using the hash value as key.

4.8 Packer

The Packer is responsible to assemble the final system image.

Name

mupack

Input

System policy in format B, Input directories, System image filename

Output

System image file

The Packer calculates the size of the resulting system image by querying the file-backed memory region with the highest physical memory address. It allocates a buffer of that size which is initially filled with zeros. It then iterates over all file-backed memory regions in the policy and adds the content of the files to the buffer. Before writing the buffer to a file specified on the command line, the packer tool performs post-checks on the buffer to make sure it is sound.

5 Additional Tools

This section lists additional helper tools which simplify the process of generating and validating a Muen system.

5.1 Kernel ELF Checker

The `Mucheckelf` tool enforces that the format of a given Muen kernel ELF binary matches the kernel memory layout specified in a system policy.

Size, VMA (Virtual Memory Address) and permissions of binary ELF sections are validated against kernel memory regions defined in the policy. The following table lists the correspondence of ELF section names to logical kernel memory region names.

ELF Section	Memory Name
.text	kernel_text
.data	kernel_data
.rodata	kernel_ro
.bss	kernel_bss

5.2 Stack Usage Checker

The `Mucheckstack` tool statically calculates the worst-case stack usage of a native Ada/SPARK component or the Muen kernel compiled with the `-fcallgraph-info` switch⁹.

The tool takes a GNAT project file and a stack limit in bytes as input. All control-flow information (.ci) files found in the object directory of the main project and all of its dependencies are parsed. Once the control-flow graph is constructed the maximum stack usage of each subprogram is calculated and checked against the user-specified limit. The tool exits with a failure if a stack usage exceeding the limit is detected.

Note that the tool is not applicable to arbitrary software projects as it does not handle dynamic/unbounded stack usage and recursion. In the context of the Muen project these cases can not occur since they are prohibited by the following restriction pragmas:

- `No_Recursion`
- `No_Secondary_Stack`
- `No_Implicit_Dynamic_Code`

Additionally, the `-Wstack-usage` compiler switch warns about potential unbounded stack usage.

5.3 Hardware Config Generator

The `Mugenhwcfg`¹⁰ tool has been created to automate the process of gathering all necessary hardware information. To collect data for a new target hardware all that is required is to run the tool on a common Linux distribution. See the project README for more information.

5.4 Scheduling Plan Generator

The `Mugenschedcfg`¹¹ tool generates scheduling plans for Muen based on a given scheduling configuration. The configuration allows the user to specify the following scheduling properties:

- Number of CPU cores

⁹https://www.adacore.com/uploads/technical-papers/Stack_Analysis.pdf

¹⁰<https://git.codelabs.ch/?p=muen/mugenhwcfg.git>

¹¹<https://git.codelabs.ch/?p=muen/mugenschedcfg.git>

- The tick rate of the CPUs
- Security constraints to meet
 - Same CPU domains
 - Simultaneous execution domains
- Subject specifications
- Score functions
- Number of plans to generate
- Plans
 - Weighting of plan importance
 - Levels
 - Subjects of a plan
 - Chains with throughput metric

Consult the project's README and example plans on how to use the tool.

5.5 Component Binary Splitter

The `mucbinsplit` tool splits component binaries into multiple files per ELF section.

Name

`mucbinsplit`

Input

Component description in XML, Component ELF binary

Output

Binary files corresponding to ELF sections, processed component description in XML

The component binary splitter tool processes component binaries and creates a separate file for each ELF section. The component XML description is extended by adding a file-backed memory region for each ELF section with the appropriate size and access rights.

The resulting processed component description is written to the given output location while the binary section files are written to the specified output path.