The VFS paradigm from the perspective of a component OS



Christian Helmuth

<christian.helmuth@genode-labs.com>







Outline

- 1. Motivation
- 2. History of the VFS in Genode
- 3. Where are we now?
- 4. Main course finished, dessert anyone?





4



- Existing open-source applications can be ported
- Gradual decomposition (kernelization) of sensitive parts



- Existing open-source applications can be ported
- Gradual decomposition (kernelization) of sensitive parts
- Reconstruction of complex (single-purpose) software not appealing



- Existing open-source applications can be ported
- Gradual decomposition (kernelization) of sensitive parts
- Reconstruction of complex (single-purpose) software not appealing
- Consequently traditional applications will always be around



Many attractive applications require a POSIX environment.



Many attractive applications require a POSIX environment.

- Fairly global view on resources as a tree of directories and files
- File abstraction grants access to data on storage, configuration, hardware peripherals, (graphical) user interfaces, to some extent even networking



Many attractive applications require a POSIX environment.

- Fairly global view on resources as a tree of directories and files
- File abstraction grants access to data on storage, configuration, hardware peripherals, (graphical) user interfaces, to some extent even networking
- Mostly traditional access control (user/group permission bits)
- C/C++ runtime
- Extensive system API (only a small share used)



Our playground is Genode.



Our playground is Genode.

- Microkernel-based
- Capability-based
 - No global namespace
 - Fine-grained access control to resources
 - Recursive system structure
- Component-based
 - ► Versatile combination of components according to use case
 - ► Spectrum from strong dependency to loose coupling
 - Only integrate what's needed to solve the task
 - Low software complexity as primary goal



Application sandboxing seems a viable approach.



Application sandboxing seems a viable approach.

- Customized C library
- Runtime environment including VFS



Application sandboxing seems a viable approach.

- Customized C library
- Runtime environment including VFS
- How about the world beyond the frame of the sand pit?
- Connect to Genode services



Outline

1. Motivation

2. History of the VFS in Genode

3. Where are we now?

4. Main course finished, dessert anyone?



Customized C library



- Back end for used subset of C library functions
- Some emulated mechanisms to satisfy applications (e.g., UIDs, permissions, sysctl)



Customized C library

- Back end for used subset of C library functions
- Some emulated mechanisms to satisfy applications (e.g., UIDs, permissions, sysctl)
- Gradually develop C library plugins to access Genode services
 - ► File-system libraries (FFAT, FUSE) use Block service
 - Network-stack libraries (lwip, lxip) use Nic service
 - Terminal service
 - File-system service
 - Special-purpose (e.g., libdrm/gallium)



THE .



Customized C library





The VFS paradigm from the perspective of a component OS



Noux runtime for Unix software







Pipes, process management



- UNIX command-line utilities running in shell
- Fork/exec child processes
- Pipes, process management
- Configure and build with original GNU build system
- From process-local to Noux-wide resource representation



The VFS paradigm from the perspective of a component OS



Plug VFS into C library

The VFS paradigm from the perspective of a component OS



Plug VFS into C library

- VFS and plugins in Noux and stand-alone C applications alike
- Gradually replace C library plugins by VFS plugins



Plug VFS into C library

- VFS and plugins in Noux and stand-alone C applications alike
- Gradually replace C library plugins by VFS plugins
- Process-local VFS instance
- Individual configuration



1			
	The VES paradigm from the perspec	tive of a component OS	15



VFS instance as XML node in component configuration

<config> <vfs>...</vfs> </config>



VFS instance as XML node in component configuration

<config> <vfs>...</vfs> </config>

Directories

<vfs> <dir name="dev">...</dir> </vfs>



VFS instance as XML node in component configuration

```
<config> <vfs>...</vfs> </config>
```

Directories

<vfs> <dir name="dev">...</dir> </vfs>

Plugins as nodes in the tree

<vfs> <dir name="dev"> <log/> </dir> </vfs>



VFS instance as XML node in component configuration

```
<config> <vfs>...</vfs> </config>
```

Directories

<vfs> <dir name="dev">...</dir> </vfs>

Plugins as nodes in the tree

<vfs> <dir name="dev"> <log/> </dir> </vfs>

Configure C library mechanisms to use VFS nodes

```
<libc stdout="/dev/log" stderr="/dev/log"/>
```



```
<config>
<vfs>
<dir name="dev"> <log/> </dir>
</vfs>
<libc stdout="/dev/log" stderr="/dev/log"/>
</config>
```

- Precise declaration of resource representation in VFS
- Plugins can be single files or whole directory sub-trees
- Nodes in a directory organized as stack (or union mount)
- Tweaking of C library behavior



Files backed by ROM service

```
<rom name=".vimrc" label="vimrc.txt"/> <rom name="avatar.png"/>
```



Files backed by ROM service

```
<rom name=".vimrc" label="vimrc.txt"/> <rom name="avatar.png"/>
```

Inline-defined file contents

<inline name="app.config">avatar = /avatar.png</inline>



Files backed by ROM service

```
<rom name=".vimrc" label="vimrc.txt"/> <rom name="avatar.png"/>
```

Inline-defined file contents

```
<inline name="app.config">avatar = /avatar.png</inline>
```

Handy tools

<null/> <zero/> <symlink name="editor" target="/bin/vim"/>



RAM-backed storage (like tmpfs)

<ram/>



RAM-backed storage (like tmpfs)

<ram/>

Integration of package-archive trees

```
<tar name="vim.tar"/>
<tar name="vim-syntax.tar"/>
```



RAM-backed storage (like tmpfs)

<ram/>

Integration of package-archive trees

```
<tar name="vim.tar"/>
<tar name="vim-syntax.tar"/>
```

File-to-service wrappers

```
<log/> <rtc/> <terminal/> <block/>
```



Expandable by custom shared objects

<jitterentropy name="random"/>
<gtotp name="gtop.service.net" secret="IMGLPG6VANGX3UCP"/>



Expandable by custom shared objects

```
<jitterentropy name="random"/>
<gtotp name="gtop.service.net" secret="IMGLPG6VANGX3UCP"/>
```

File-system session to use persistent storage

```
<fs name="home"/>
<fs name="cfg" label="config" writeable="yes"/>
<fs name="bin" label="bin" root="/usr"/>
```



Expandable by custom shared objects

```
<jitterentropy name="random"/>
<gtotp name="gtop.service.net" secret="IMGLPG6VANGX3UCP"/>
```

File-system session to use persistent storage

```
<fs name="home"/>
<fs name="cfg" label="config" writeable="yes"/>
<fs name="bin" label="bin" root="/usr"/>
```

VFS configuration is component-local \rightarrow access control by policies in parent and service components (e. g., file-system service)

```
<policy label="app -> " root="/home" writeable="yes"/>
<policy label="app -> config" root="/app/config" writeable="no"/>
<policy label="app -> bin" root="/" writeable="no"/>
```



Outline

1. Motivation

2. History of the VFS in Genode

3. Where are we now?

4. Main course finished, dessert anyone?





- Robust implementation of file-system service
- File-system itself implemented as VFS plugin



- Robust implementation of file-system service
- File-system itself implemented as VFS plugin
- One VFS configuration for multiple components
- Multiplex access to all aggregated resources
- Differentiate client permissions by policies



- Dynamic reconfiguration of server applies to all clients



- Dynamic reconfiguration of server applies to all clients
- All VFS plugins can be used in the server



- Dynamic reconfiguration of server applies to all clients
- All VFS plugins can be used in the server
- Now, at the latest, HURD translators come into mind...





Shared resources could be provided by large plugins.

Dedicated rump (ext2) server superseded by plugin



Is there more about that server?

Shared resources could be provided by large plugins.

- Dedicated rump (ext2) server superseded by plugin
- Network stack based on lxip/lwip

```
<!-- server -->
<vfs> <lxip/> </vfs>
```

```
<!-- client -->
<vfs> <dir name="/socket"> <fs/> </dir> </vfs>
<libc socket=/socket>
```



Outline

1. Motivation

2. History of the VFS in Genode

3. Where are we now?

4. Main course finished, dessert anyone?



- Audited system resource discovery and access
- C library integration enables existing POSIX applications
- Versatile combination inside component and via VFS server
- Abstraction from Genode services
- Future extension is easy (e.g., USB service adapter for libusb)



Extend VFS interface to support plugins using other plugins



Outlook

- Extend VFS interface to support plugins using other plugins
- Plugin-based filter chains become possible
 - Mangling/routing of mouse/keyboard input events
 - File systems using block-device nodes



Outlook

- Extend VFS interface to support plugins using other plugins
- Plugin-based filter chains become possible
 - Mangling/routing of mouse/keyboard input events
 - ► File systems using block-device nodes
- Split applications scenarios driven by security considerations
 - ► PDF reader with file-to-HTTP plugin in separate component
 - Separate domains for edit-compile-test-push development workflow



Outlook

- Extend VFS interface to support plugins using other plugins
- Plugin-based filter chains become possible
 - Mangling/routing of mouse/keyboard input events
 - File systems using block-device nodes
- Split applications scenarios driven by security considerations
 - ► PDF reader with file-to-HTTP plugin in separate component
 - Separate domains for edit-compile-test-push development workflow

Application-stack architectures range from multiple components connected by file-system sessions to unikernel-like monoliths.



Thank you



Genode OS Framework https://genode.org/

Genode Labs GmbH https://www.genode-labs.com/

Source code on GitHub

https: //github.com/genodelabs/genode

Genode Foundations book

https://genode.org/documentation/genode-foundations-16-05.pdf