Celebrating kernel diversity with Genode

-

FOSDEM 2025

Alexander Böttcher
<alexander.boettcher@genode-labs.com>

Outline

1. Genode OS framework and Sculpt OS
2. History of supported kernels
3. Yet another kernel
4. What next

Genode OS framework

- Vision of a truly trustworthy operating system

  least privilege ★ *capability-based* ★ *microkernel* ★ *virtualization*

- Open-source tool kit to craft own operating systems
  - Prototype 2006, company founded 2008, first release in 2009
  - Language: C++, License: AGPLv3, optional commercial

- Over 100 ready to use components
  - Sandboxed drivers, multiplexer and applications
  - Several microkernels and one monolithic kernel
  - Kernel-agnostic application binary interface
  - Package management

- Quarterly releases & documentation

Books

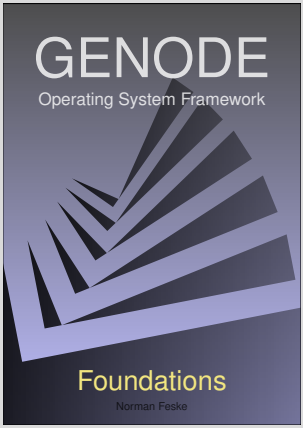
GENODE
Operating System Framework
Foundations
Norman Feske

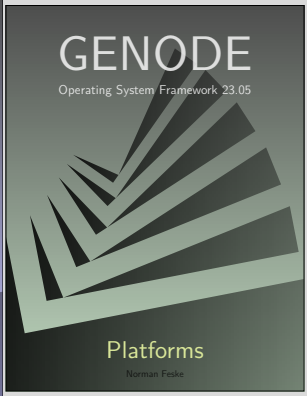
GENODE
Operating System Framework 23.05
Platforms
Norman Feske

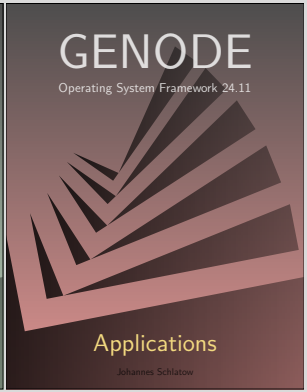
GENODE
Operating System Framework 24.11
Applications
Johannes Schlatow

Genode OS framework components

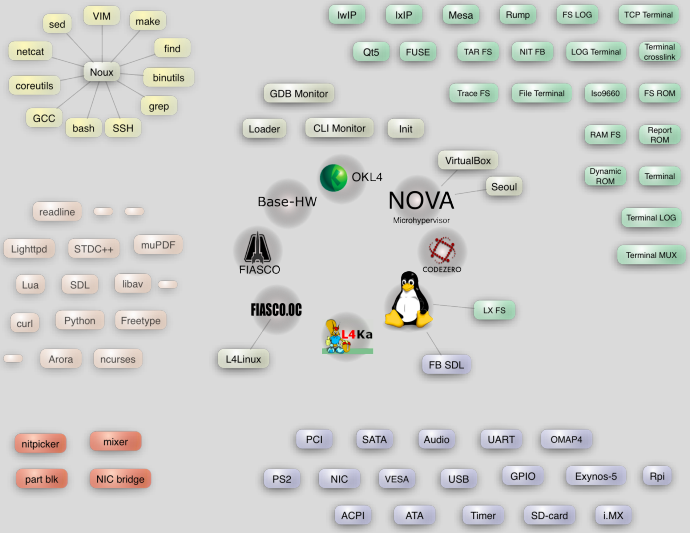

Sculpt OS - dynamic usage of Genode

- **Day to day OS** used by community and developers on consumer devices
- Combining Genode's microkernel architecture, capability-based security, sandboxed device drivers, and virtual machines in a novel OS
  - **Tiny** base OS with automatic device detection and configuration
  - **Tight** GUI as control interface
  - GUI for package manager & **federated** package provisioning
  - Release cycle: 2x time per year since 2017
- Dynamic configurable at runtime - **sculpt your own OS**
  - All federated packages to be downloaded
  - Window manager
  - Native browsers
  - Various VMMs to run VMs
  - Native applications, PDF viewer, GPU applications, ...
  - **Optional**: POSIX runtime, e. g. to run GNU tools

Sculpt OS - supported consumer hardware

- MNT Reform i.MX8MQ, PinePhone (ARMv8) and x86 notebooks

Sculpt OS - supported consumer hardware

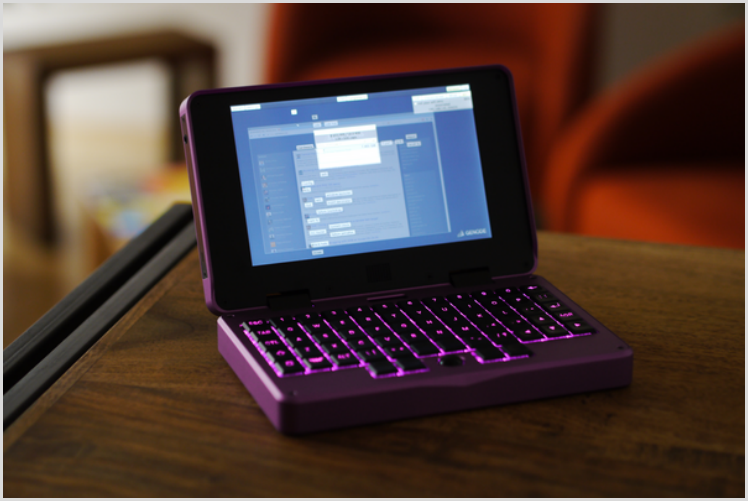

- MNT Pocket Reform, i.MX8MP (ARMv8)

Sculpt OS 24.10 - user visible highlights

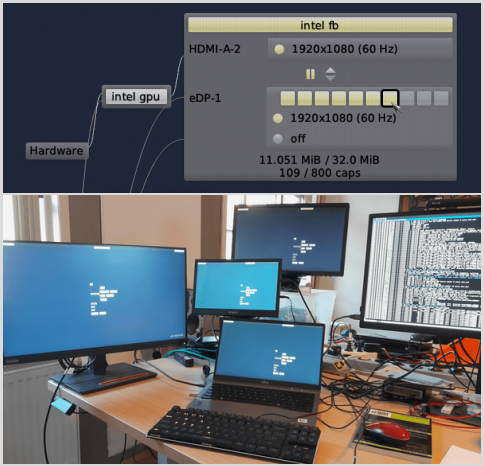

Holistic multi-head & panorama support

Sculpt OS 24.10 - user visible highlights

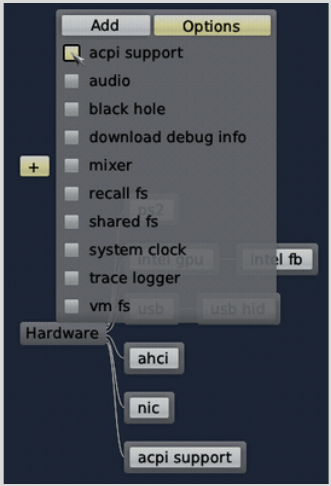

Suspend & resume support for x86

Joint FOSDEM stand - Qubes OS and Genode

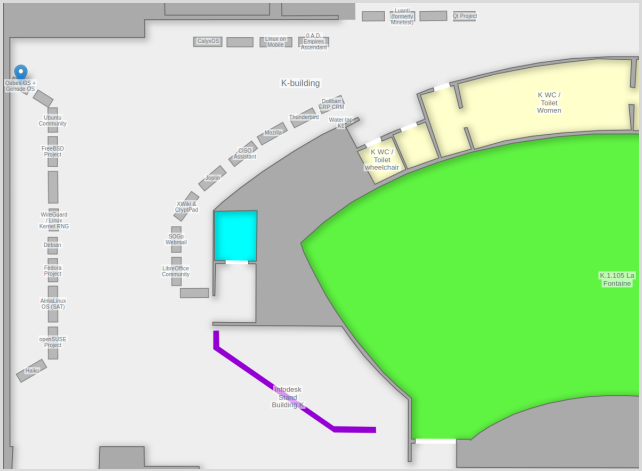

Worth a travel - K building, level 1, stand 1

Outline

1. Genode OS framework and Sculpt OS
2. History of supported kernels
3. Yet another kernel
4. What next

Expectations towards a kernel

- Microkernel
- Kernel protected capabilities
  - Delegation during Genode RPC calls, e. g. cap as argument of an RPC
  - Recognition of identical argument capabilities after RPC calls
- Strict accounting of physical resources
  - No static/fixed kernel memory pools
- Non-blocking kernel interfaces
  - Supports event style programming in Genode
- Multicore: cross core IPC support
- Optional: virtualization

Supported Kernels - timeline

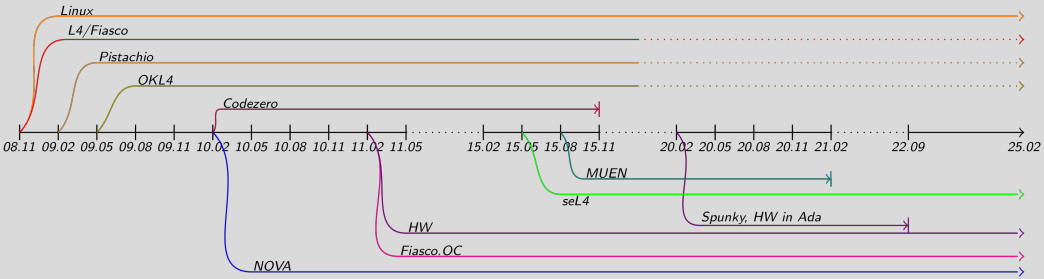

First kernels - characteristics

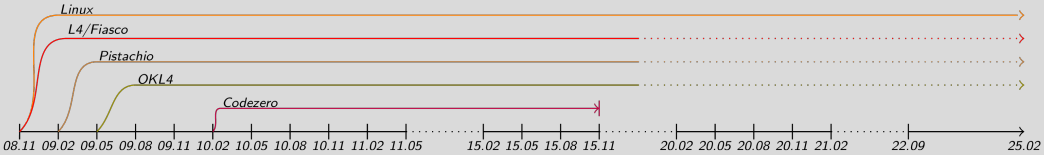

	Linux	L4/Fiasco	OKL4	Pistachio	Codezero
x86/ARM	x / x	x / -	x / x	x / -	- / x
kernel type	**monolithic**	micro	micro	micro	micro
→ language	C	C++	C	C	C
→ SMP	yes	no	yes	yes	yes
capabilities		**no**	**no**	**no**	**no**
virtualization		L4/Linux	OK/Linux	Afterburner	
→ support			09.11 - 13.11		
usage state	active	CI	CI	CI	-

Modern kernels - characteristics

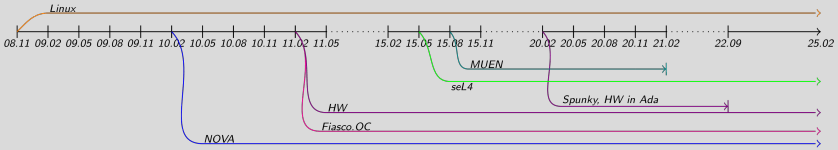

kernel - version	seL4 12.1	Fiasco.OC ~2019	hw	hw Spunky	NOVA Genode	MUEN
kernel	micro	micro	micro	micro	micro	**separation**
→ language	C	C++	C++	Ada C++	C++	Ada Spark
→ memory	**dynamic**	static	**dynamic**	**dynamic**	static	static
→ locking	**big**	fine	**big**	**big**	fine	
roottask/kernel	separate	separate	**merged**	**merged**	separate	
usage state	CI, **EXP**	CI	active	-	active	-
scenarios	static	dynamic	dynamic	-	dynamic	-

Outline

1. Genode OS framework and Sculpt OS
2. History of supported kernels
3. Yet another kernel
4. What next

Motivation - the endeavor

- Familiarize with upstream NOVA kernel development
  - My personal spare time project, very low priority, since ~8 months

  → Immediate goal: **easy** NOVA kernel upstream **testing**

- Just a minor update, right ?
  - **No**, same roots, but diverged much - **unfortunately**
  - Genode's version **ready**, only extensions on a need by need basis

- Further reasons:
  - Upstream NOVA kernel under active development
  - Several modern CPU features, trusted computing, see FOSDEM 2020, 2022, 2023
  - Ongoing formal **verification** of this **C++** kernel

- Working title: base-novae
  - NOVAe → abbreviation for `experimental`, but also plural

Genode@NOVAe - working items

- Copy repository, base-nova → base-novae
  - Wipe all extensions done since 2012 - :'(
- Roottask - core
  - Review documentation NOVAg vs NOVAe
  - Adjusting resp. writing syscall bindings
  - Get hold of major capabilities
  - Resource parsing of Multiboot 2, setup of core allocators
  - Timer service adjustment
  - One thread per IRQ in core ;-(, due to blocking syscall
- Packaging of base-novae
  - Pre-requisite for advanced scenarios

Genode@NOVAe - capability delegation during RPC ?!

Example RPC:

- void sigh(**Signal_context_capability** sigh);

First kernel without direct IPC + cap delegation support anymore

- NOVA ctrl_pd syscall + target PD cap → solely available to core

→ Add capability `re-routing` functionality via core

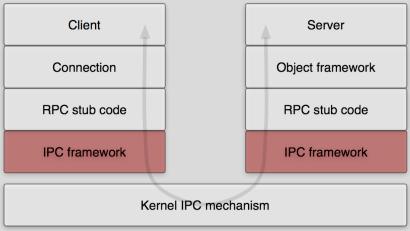

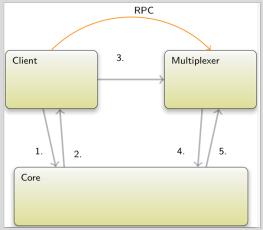

Genode@NOVAe - current state

Highly experimental by now - x86 only

- Drivers up and running
  - Restarting/killing not supported by now
- Minor kernel patches
  - ACPI RSDP lookup via Multiboot 2
  - Location of very first UTCB in roottask
  - Increase of compile time static kernel memory size

→ No SMP

→ No virtualization support

→ No Genode priorities

→ No improved (efficient) Genode signal sending

Beside that - Sculpt@NOVAe is alive :)

Demo - Sculpt 24.10 image with Genode@NOVAe & friends

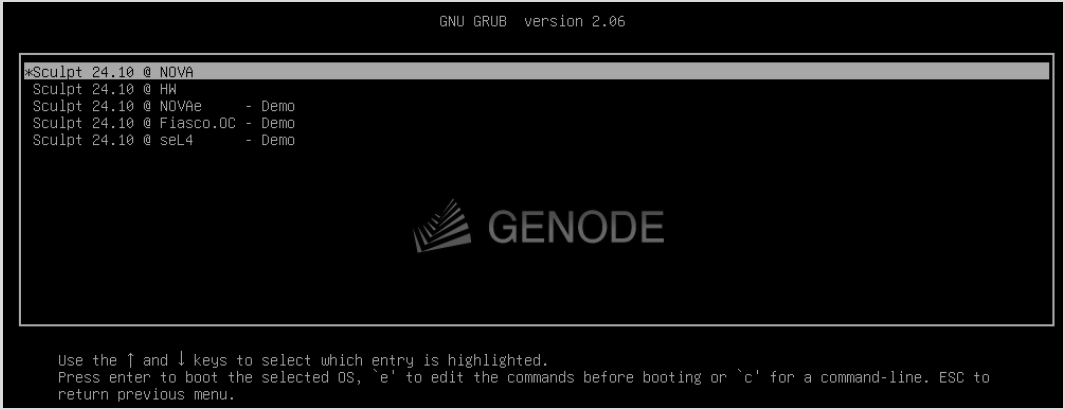

USB image - https://depot.genode.org/alex-ab/images/fosdem_2025.img

Genode branch - https://github.com/alex-ab/genode - branch: fosdem_novae_2025

Kernel branch - https://github.com/alex-ab/nova - branch: r24.35

Thank you

Genode Discourse forum

`https://genode.discourse.group`

Genodians.org community blog

`https://genodians.org`

Genode OS Framework

`https://genode.org`

Sculpt OS download and manual

`https://genode.org/download/sculpt`

Backup slides

Modern kernels - challenges

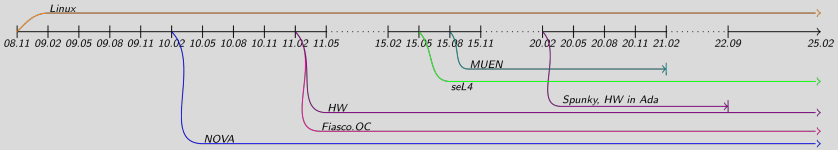

kernel - version	seL4 12.1	Fiasco.OC ~2019	hw	NOVA 2012	NOVA Genode
Multicore: cross core IPC	yes	yes	yes	no	yes
Thread exception handling by e. g. paging, breakpoint, etc - required threads	core 1	core 1	core 1	core N	core C

N: overall number of component threads in system

C: number of CPUs

Modern kernels - challenges

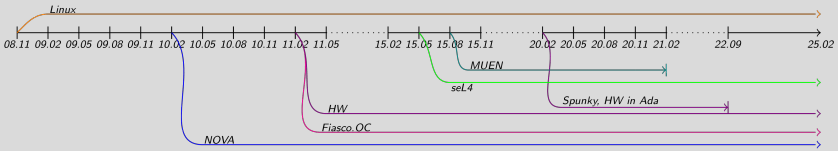

kernel - version	seL4 12.1	Fiasco.OC ~2019	hw	NOVA 2012	NOVA Genode
Cap delegation during RPC	yes	yes	yes	yes	yes
- cap translation during RPC	no	no	yes	(yes)	yes
- **extra** cap compare	yes	(yes)			

bracket: principal support by kernel, minor kernel patches required
no : no adequate support by kernel that fits Genode expectations
extra: challenge as surmounted by base-<kernel> specific solution

Modern kernels - challenges

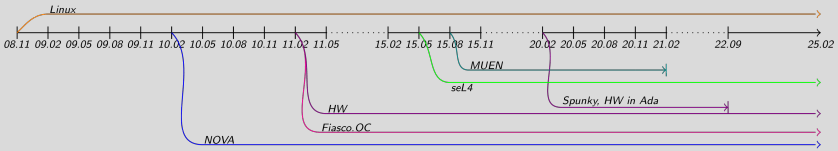

kernel - version	seL4 12.1	Fiasco.OC ~2019	hw	NOVA 2012	NOVA Genode
NB IRQ notification	(yes)	(yes)	yes	no	yes
- **extra** IRQ threads	N	1		N	
NB virtualization interface	no	no	yes	yes	yes
- **extra** thread per vCPU	yes	yes			

N : number of GSIs+MSIs in-use

NB : non blocking

no : no adequate support by kernel that fits Genode expectations

Modern kernels - supported architectures

kernel - version	seL4 12.1	Fiasco.OC ~2019	hw	NOVA Genode	MUEN
RISC-V			16.02	-	-
ARM_v6		11.05	12.11	-	-
ARM_v7	17.08	12.05	12.08	-	-
ARM_v8		19.05	19.08	-	-
x86_32	15.05 - 23.05	11.02	-	10.02	-
x86_64	17.08	11.05	15.05	12.08	15.08-21.02

- Empty field: not enabled by Genode, kernel may support it

Modern kernels - virtualization

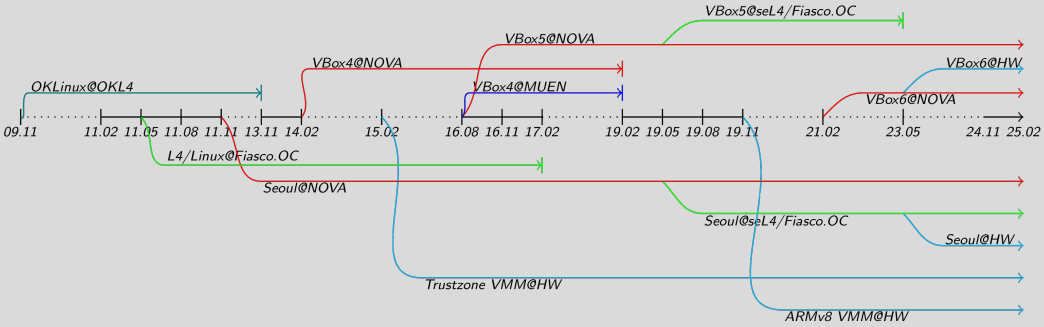

- PC: harmonized VM interface since 19.05 for seL4, Fiasco.OC, NOVA
  - base-hw: AMD/SVM in 23.05, Intel/VMX in 24.05
  - VMMs: VBox5, VBox6, Seoul
- ARM: VMM solely for base-hw

Modern kernels - extended features

kernel - version	seL4 12.1	Fiasco.OC ~2019	hw	NOVA Genode
SMMU ARM IOMMU Intel IOMMU AMD	 -	 -	(x) 23.11 -	- 13.02 20.11
Guarded MSR access for a user component e. g. CPU Power and frequency tuning			-	23.04
SMP - thread migration	x x	x (x)	x -	x (x)

- Dates: since when the feature is available
- Empty field: not enabled by Genode, kernel may support it
- Brackets: support got enabled, but either partially supported or discontinued/disabled

Device driver environments

Genode repository	Source
dde_ipxe	IPXE project
dde_bsd	Free BSD
dde_linux	Linux kernel
dde_zircon	Google Fuchsia, 18.08 - 20.08

Genode's NOVA version vs NOVA version from 2012 I/II

- Improved kernel resource management
  - Freeing up of all kernel objects
  - Kernel quota per PD and quota trading between PDs
  - Dimension kernel memory at boot time according to available system RAM
- UEFI support
- Scheduling
  - Yield support
  - Priority inheritance for Genode Locks
- PD address space management
  - Explicit manipulation of target PD
- PD capability space
  - Extended translate support of capabilities
  - Accessibility of a cap in space_obj but keep it in mapping

Genode's NOVA version vs NOVA version from 2012 II/II

- Semaphore
  - Support to block thread on IPC reply
  - Support to wait for multiple asynchronous events (aka chained semaphores)
    - No need to have multiple threads per (IRQ) SM anymore
    - Used to implement Genode signals efficiently
- Cross core IPC
- 64bit VM support
- NX bit support
- Eager FPU switching support
- Export kernel log messages via shared memory to roottask
- Bugfixes