



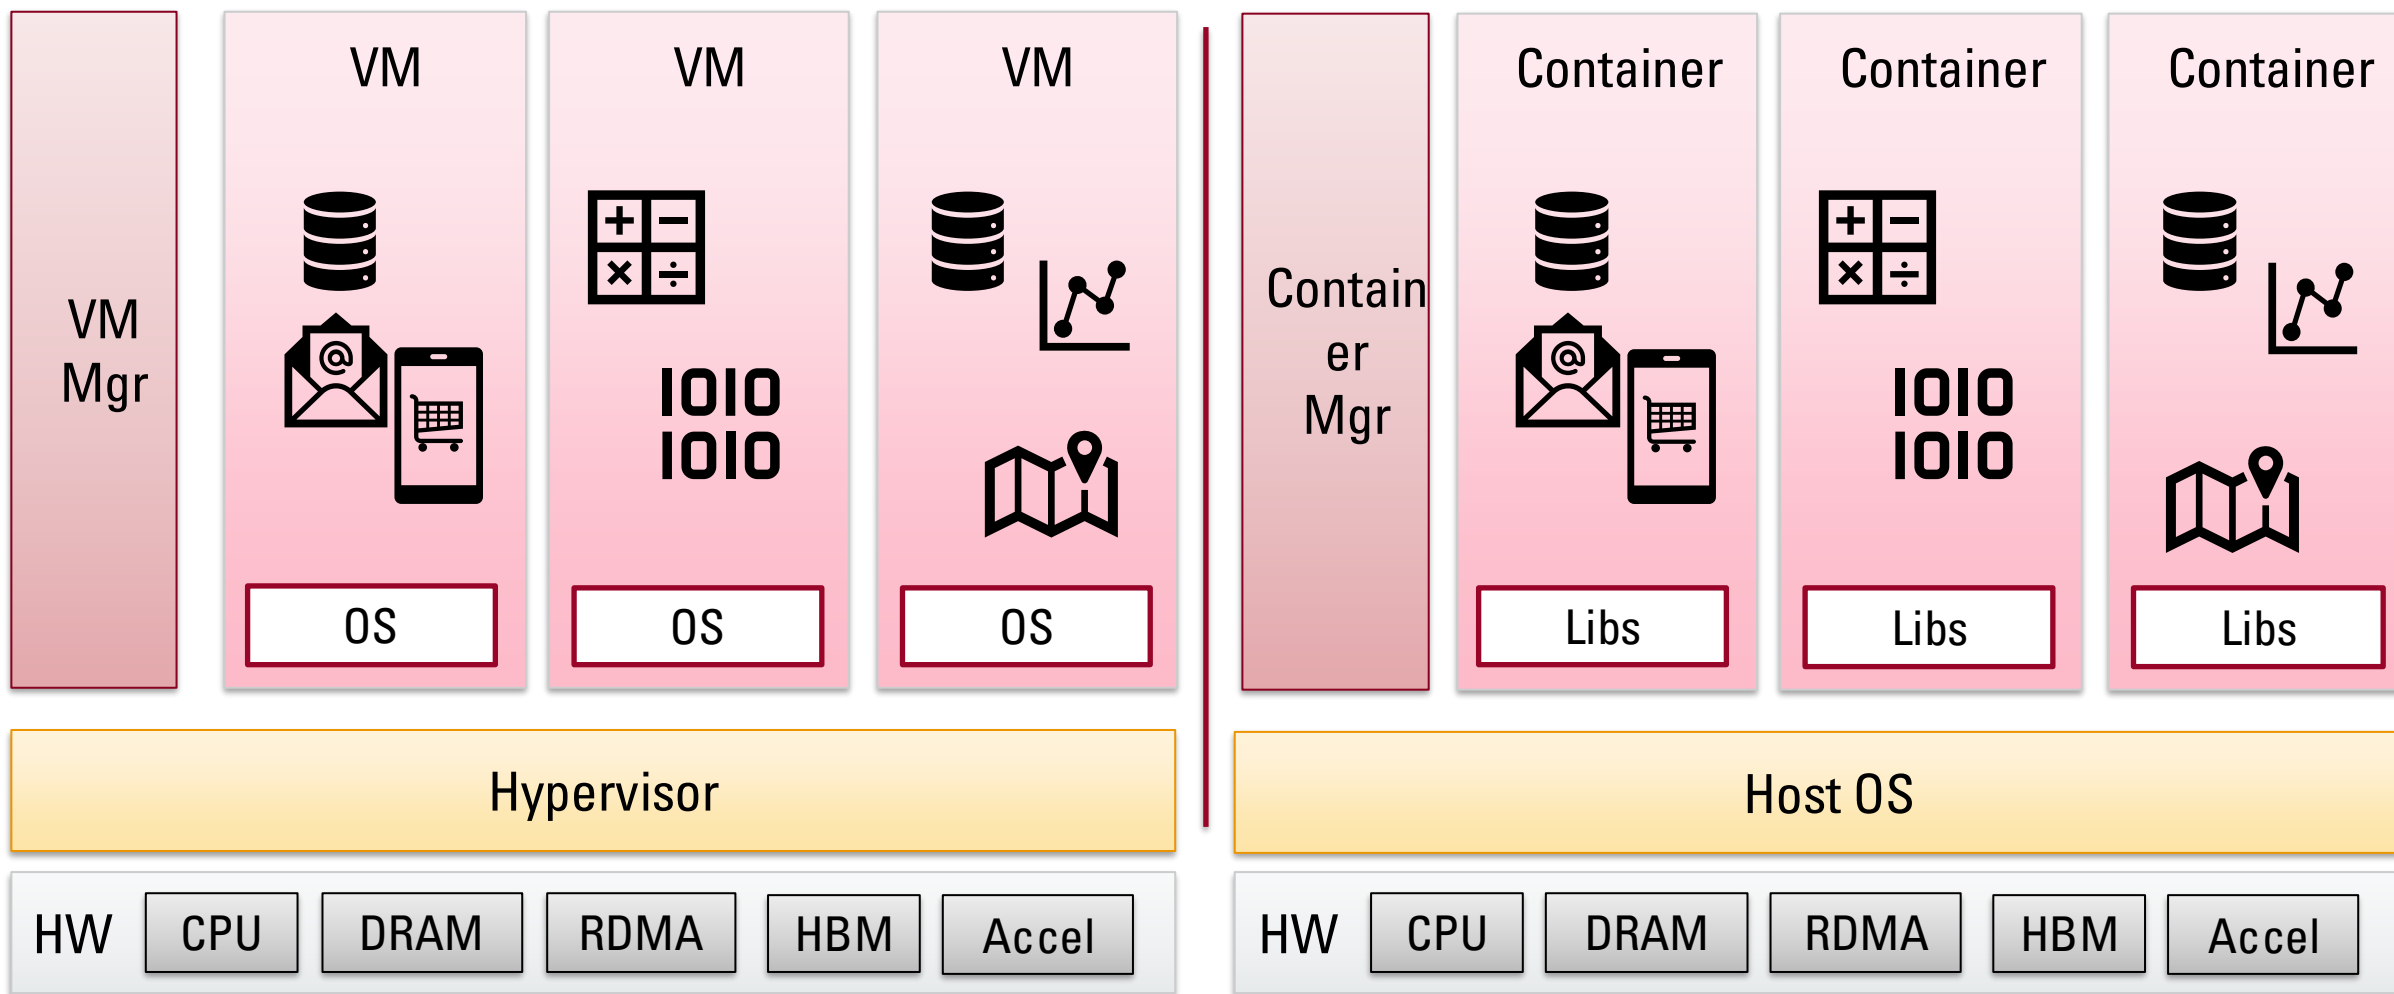
Facing the Complexity: The Challenges of Adopting Microkernels for Cloud Infrastructure

Brüssel, Feb 1, 2026

Michael Müller (michael.mueller@uos.de)

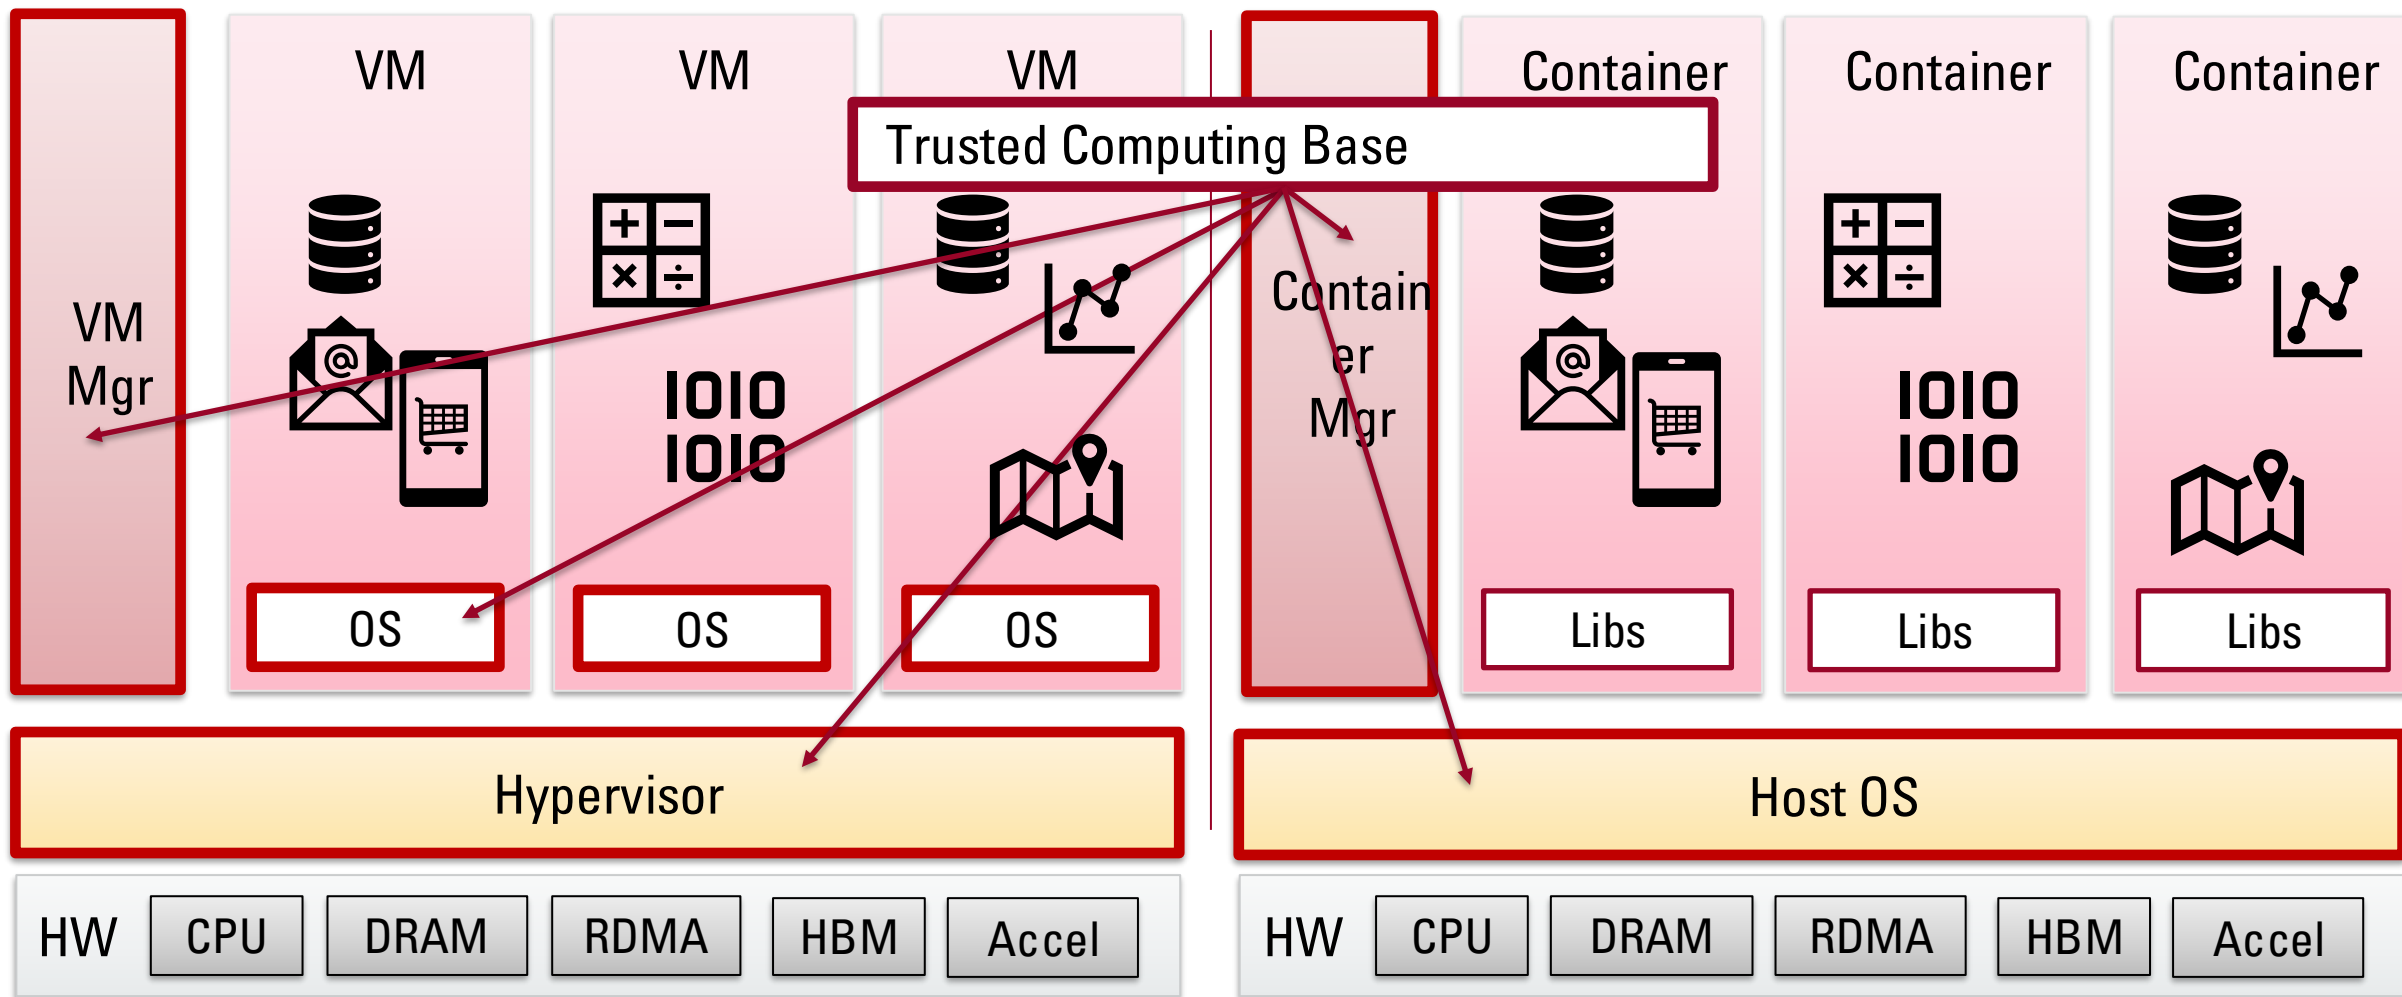


Today's Cloud Software Stack



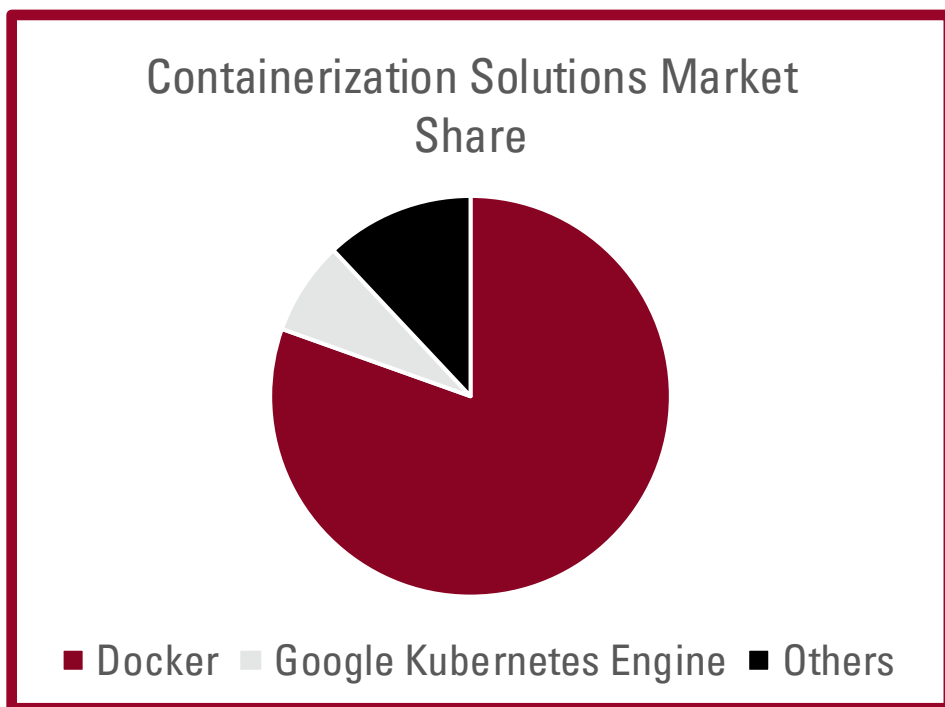


Today's Cloud Software Stack

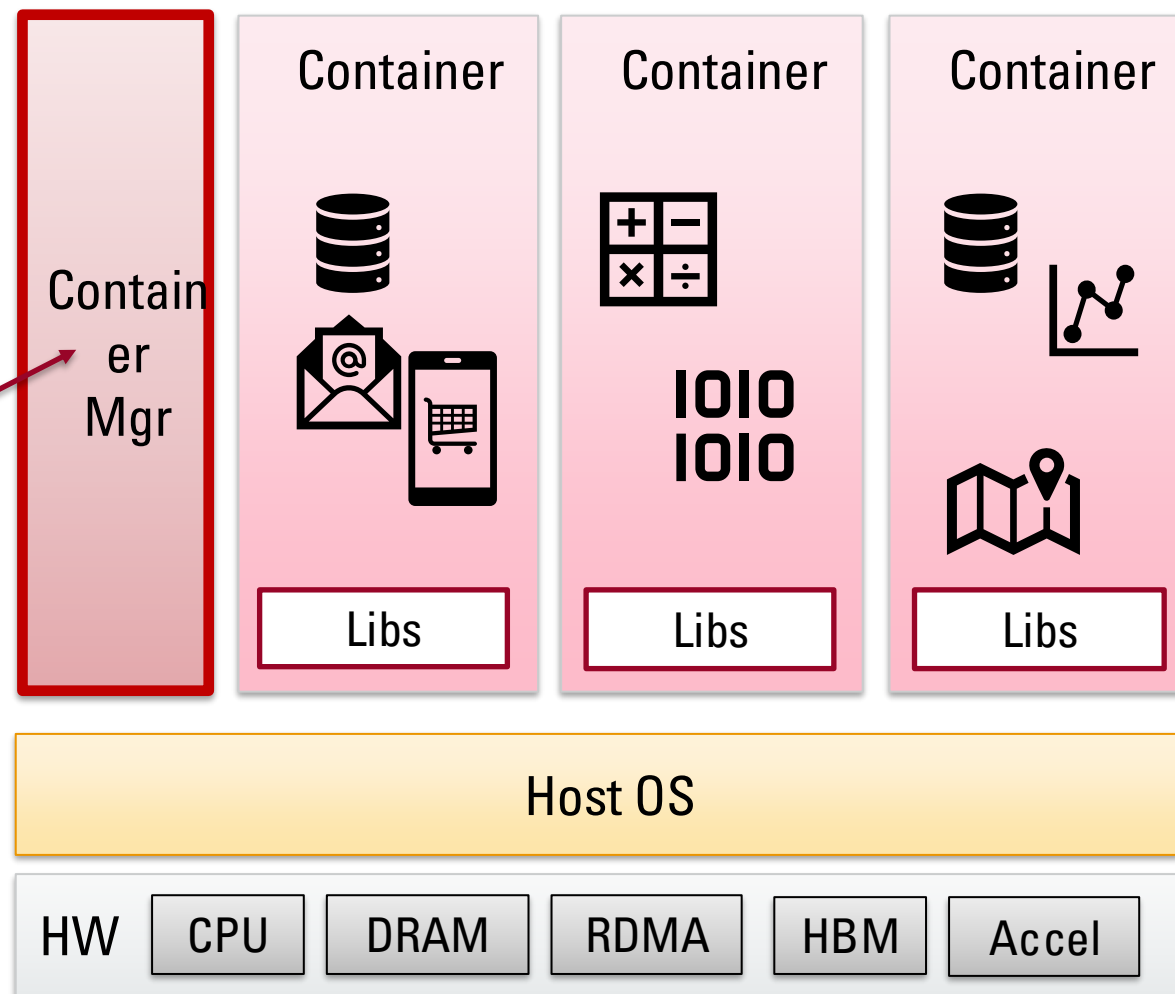




Today's Cloud Software Stack

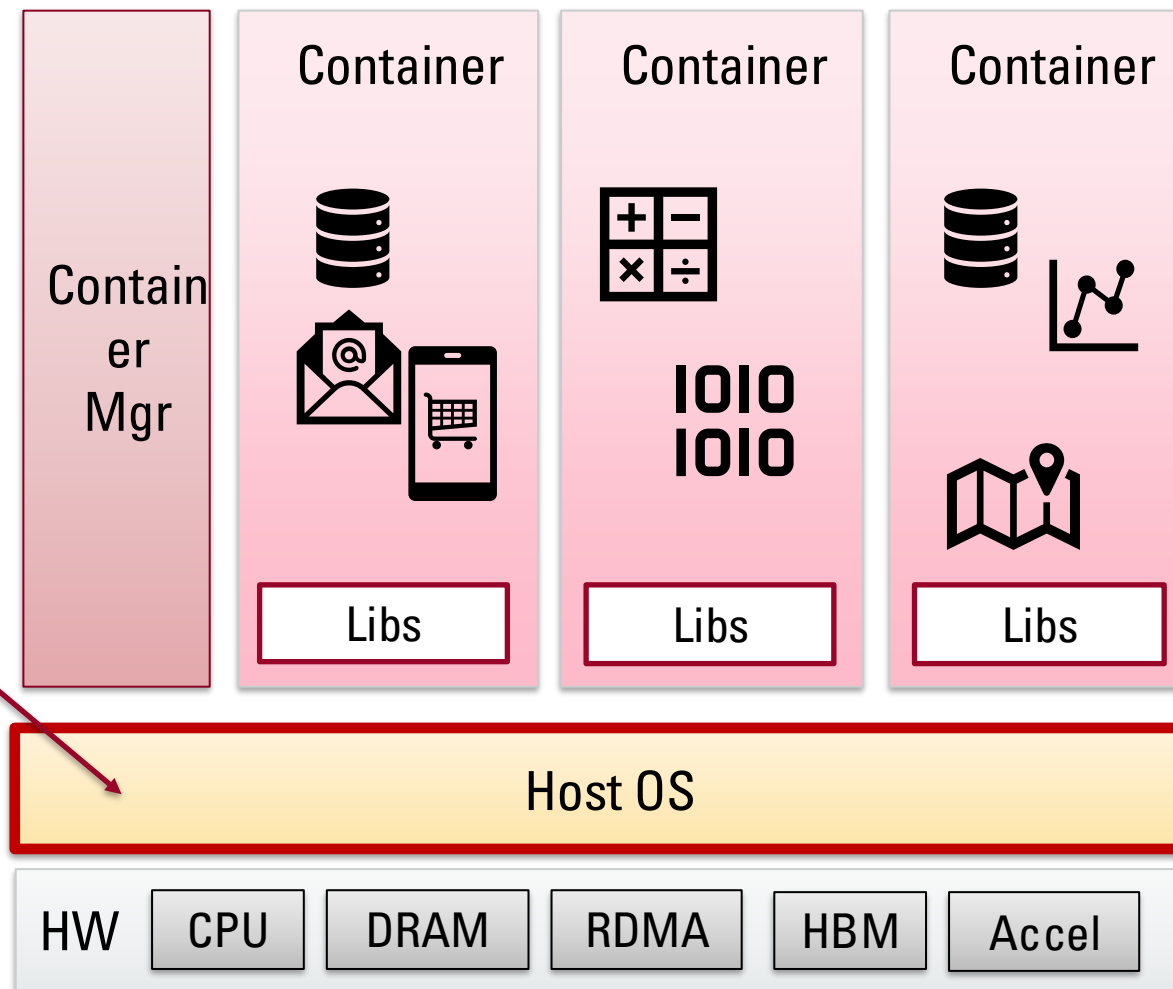



Source: <https://6sense.com/tech/containerization/docker-market-share>





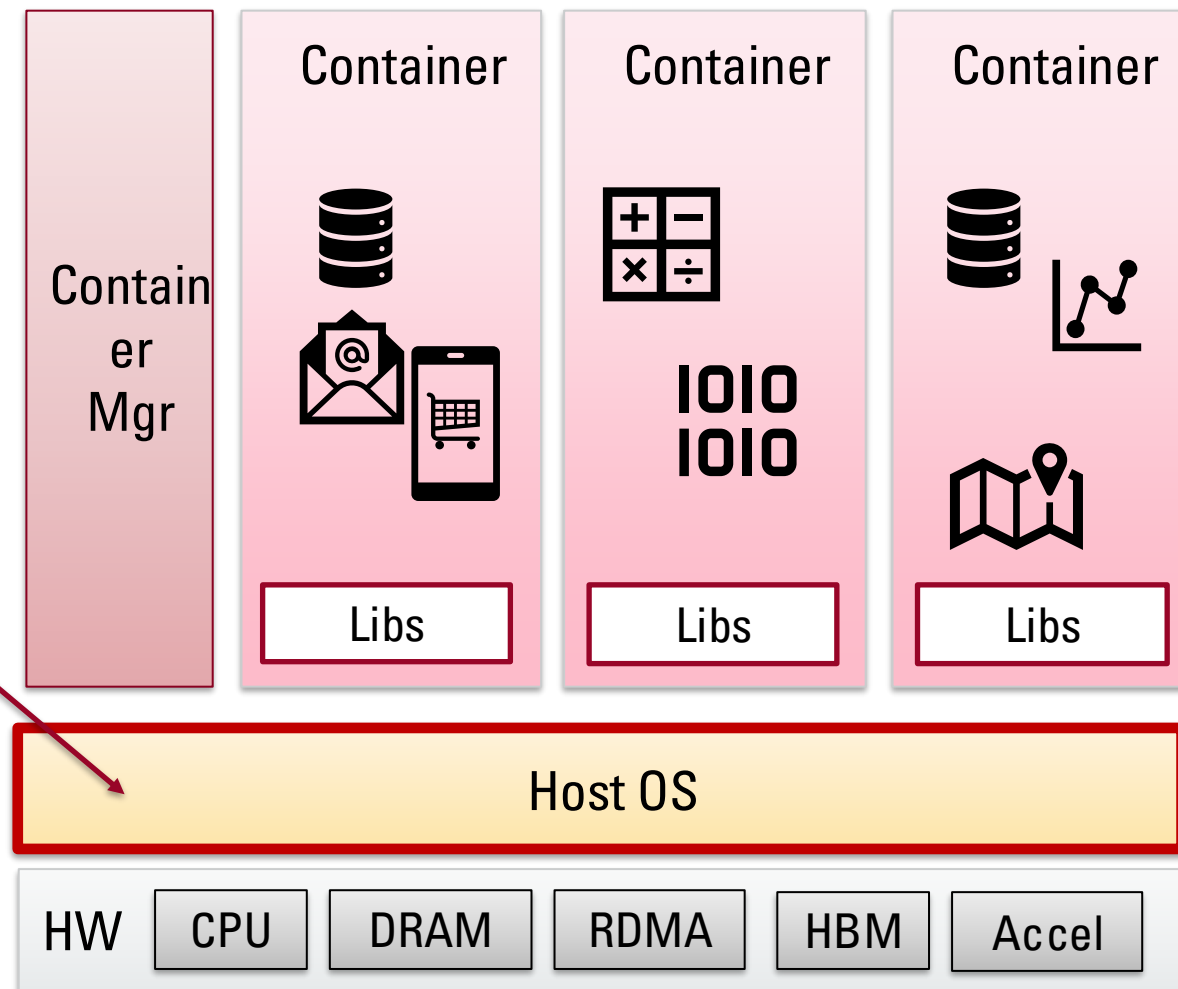

Today's Cloud Software Stack



Source: <https://commandlinux.com/statistics/linux-data-center-market-statistics/>



Today's Cloud Software Stack

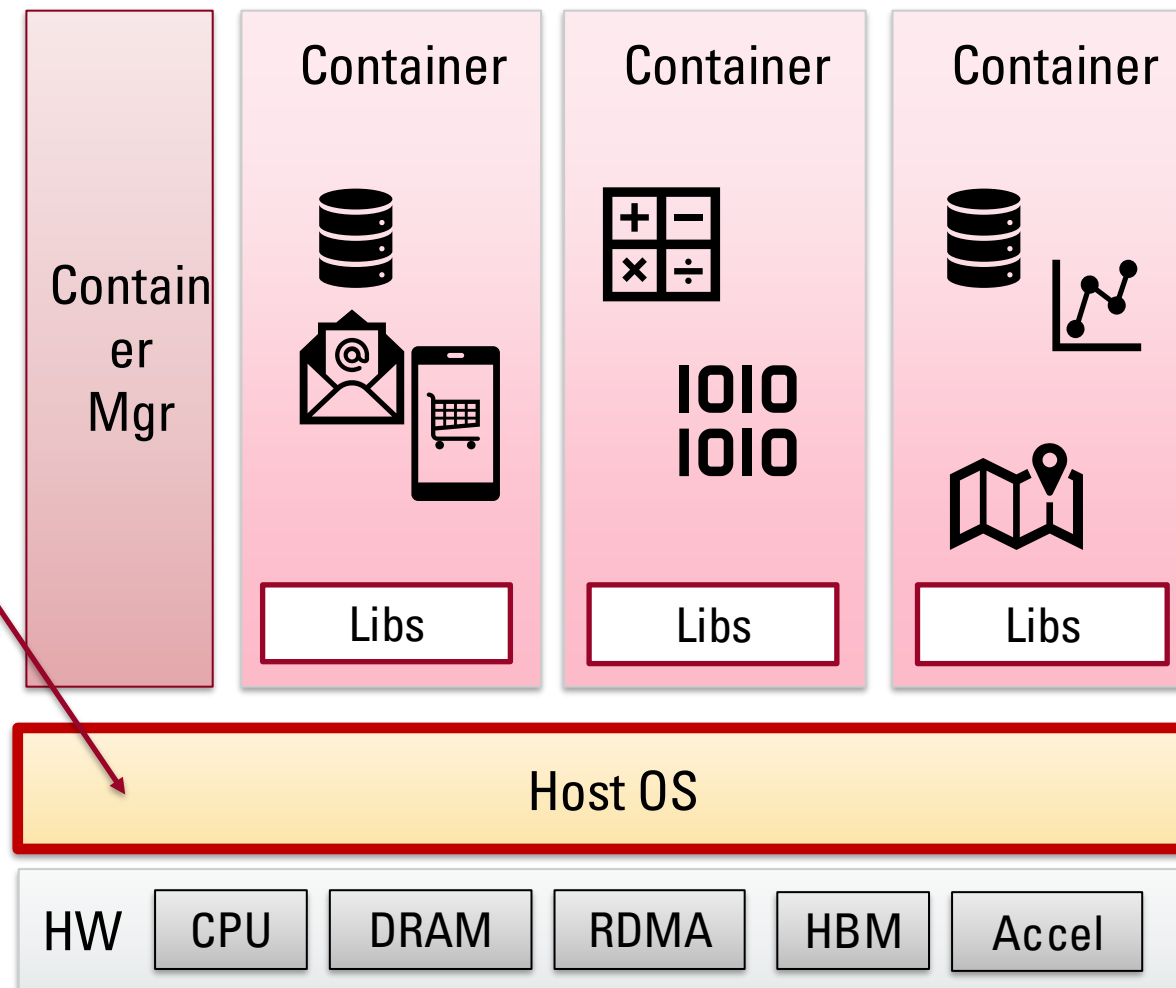




Today's Cloud Software Stack

Isolation via

- Cgroups for resource restrictions
- Namespaces for visibility restrictions (processes, FS)
- Seccomp-ebpf for restricting syscalls



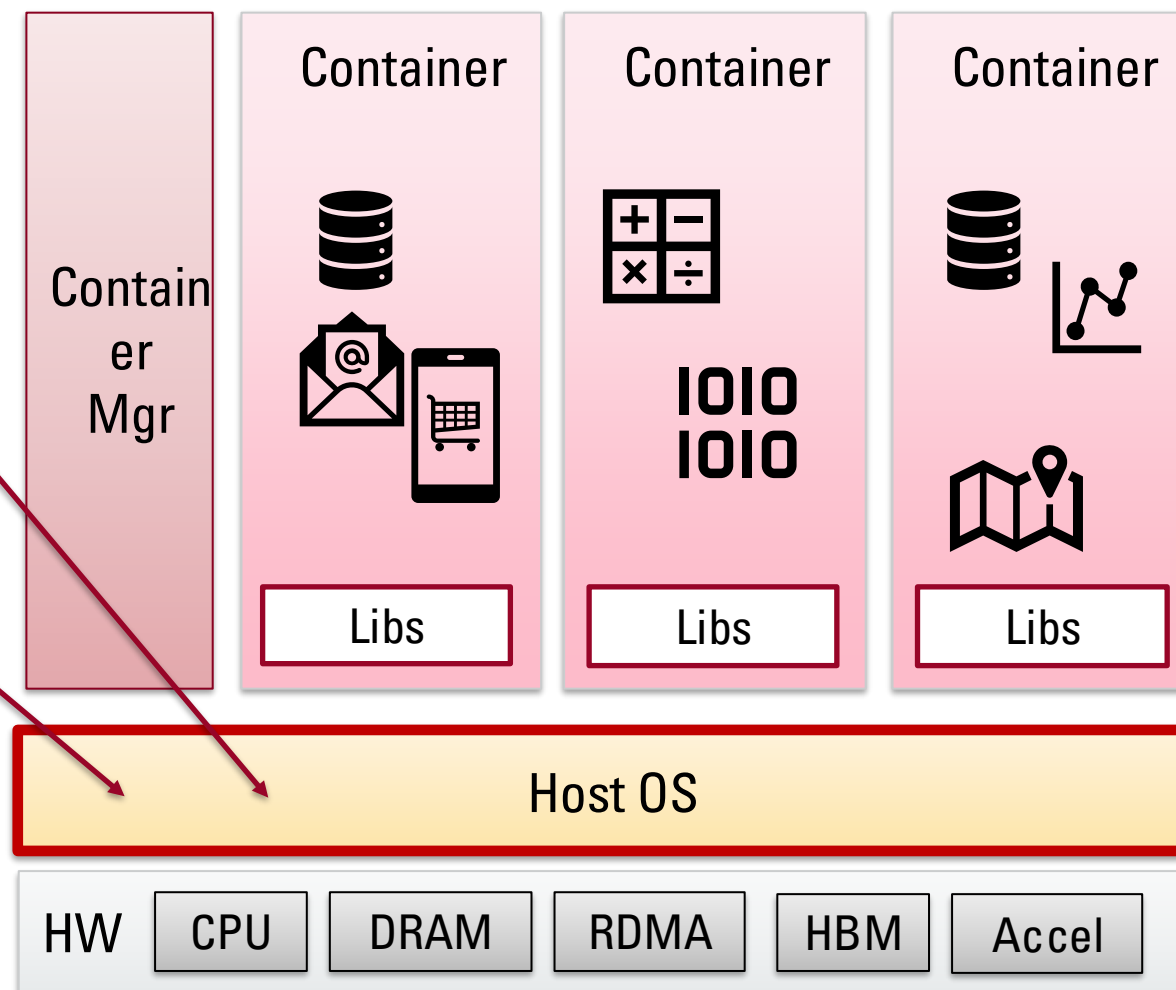


Today's Cloud Software Stack

Isolation via


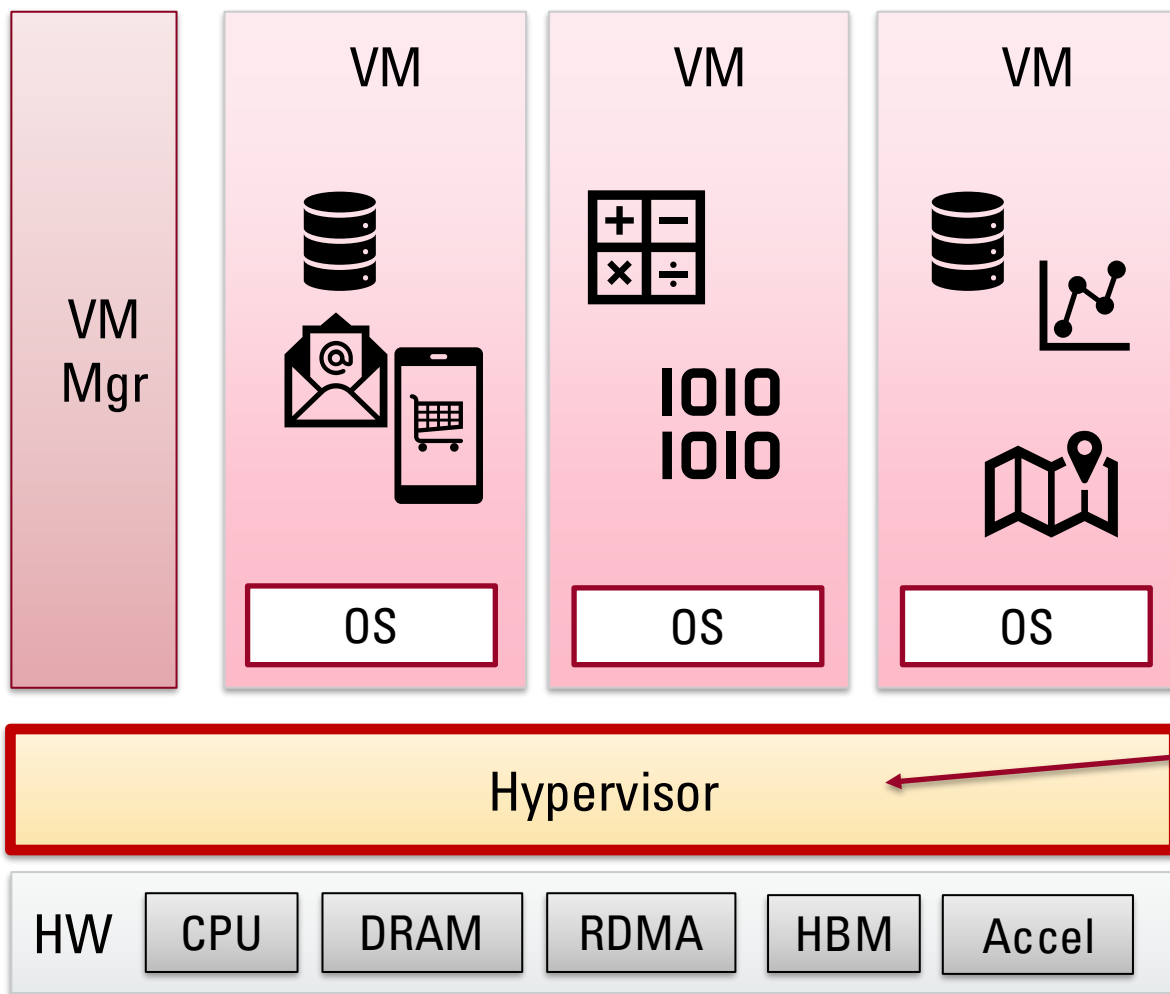
- Cgroups for resource restrictions
- Namespaces for visibility restrictions (processes, FS)
- Seccomp-ebpf for restricting syscalls

- **2,699,812** LoC for Linux 6.7.4 container infrastructure
- **8** critical vulnerabilities for **seccomp**
- **3** critical vulnerabilities for **cgroup**
- **22** critical for **namespaces**
- **33** critical CVEs in **total** [5]







Today's Cloud Software Stack



Source: <https://commandlinux.com/statistics/linux-server-market-share>

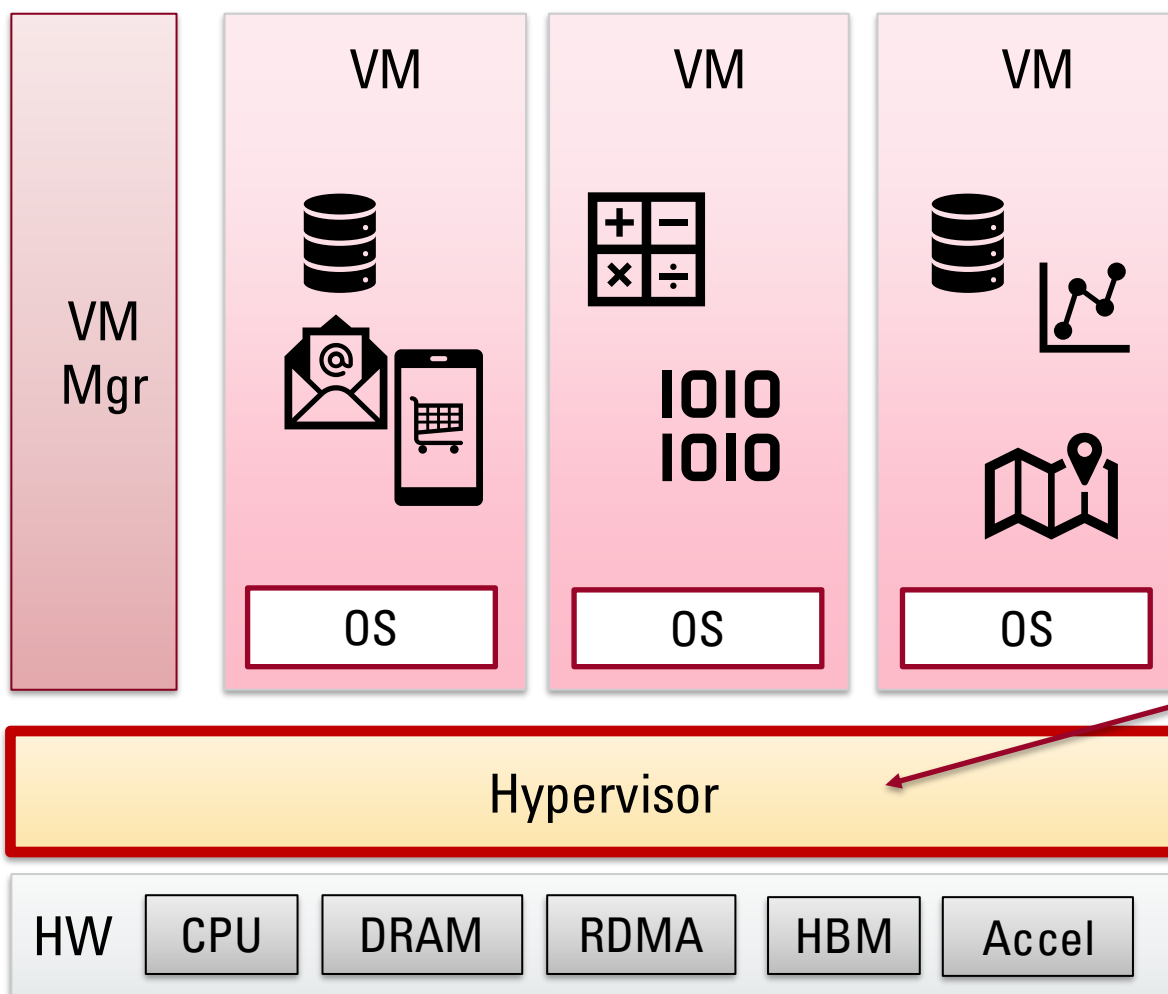


Today's Cloud Software Stack





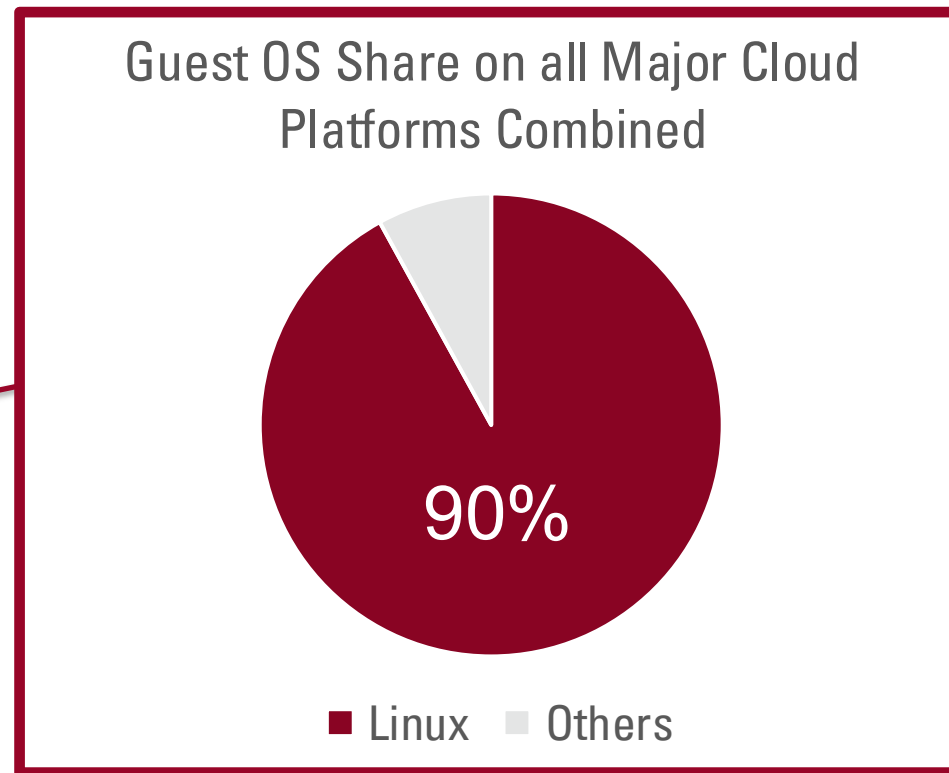
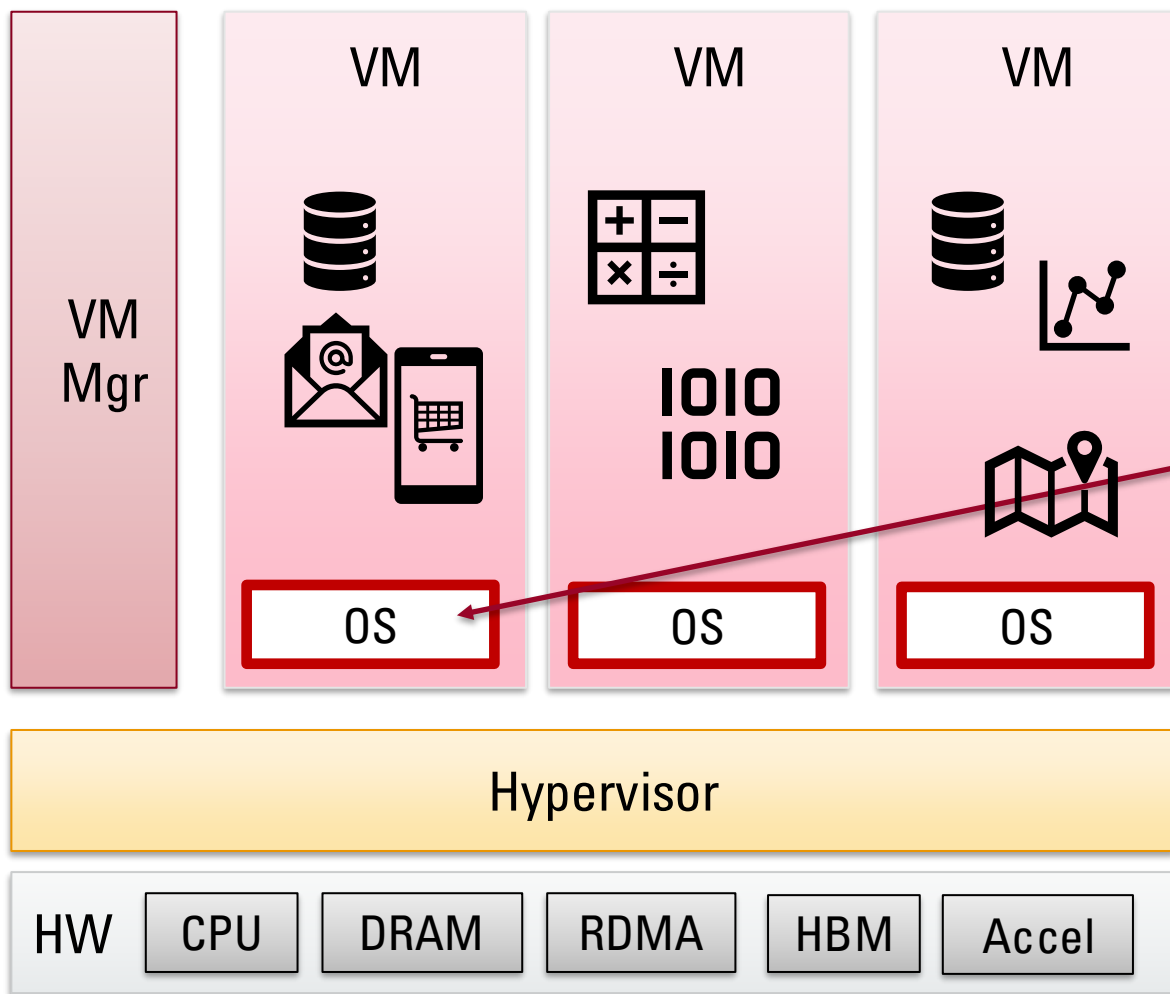
Today's Cloud Software Stack



- > 150,000 lines of code (source LWN.net) for KVM alone
- **24 critical** vulnerabilities for **KVM** in NIST vulnerability database (CVSS > 7.0) since 2008
- **77 critical** vulnerabilities for **Qemu** since 2015 (NIST vulnerability DB)



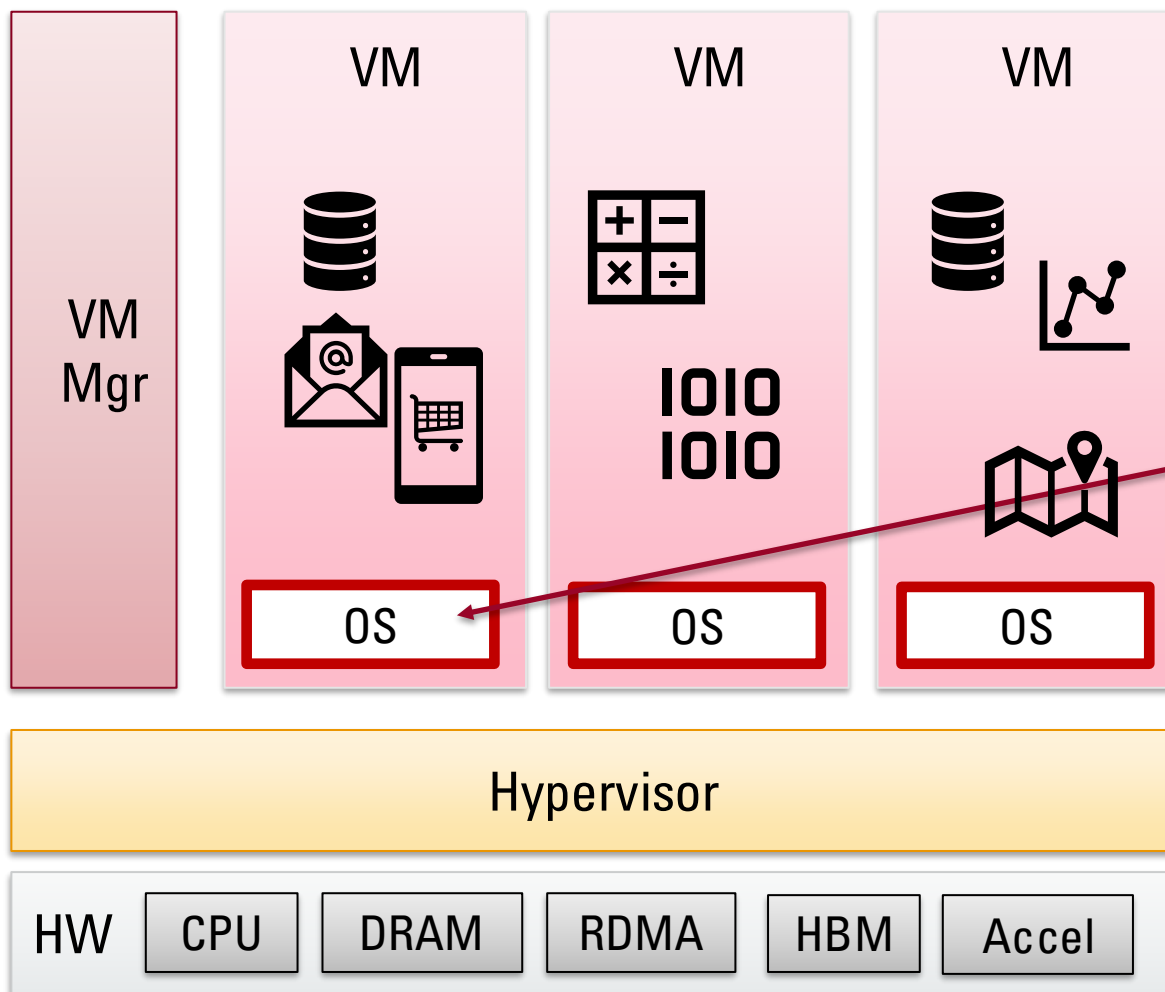
Today's Cloud Software Stack



Source: <https://commandlinux.com/statistics/linux-cloud-infrastructure-market-share>



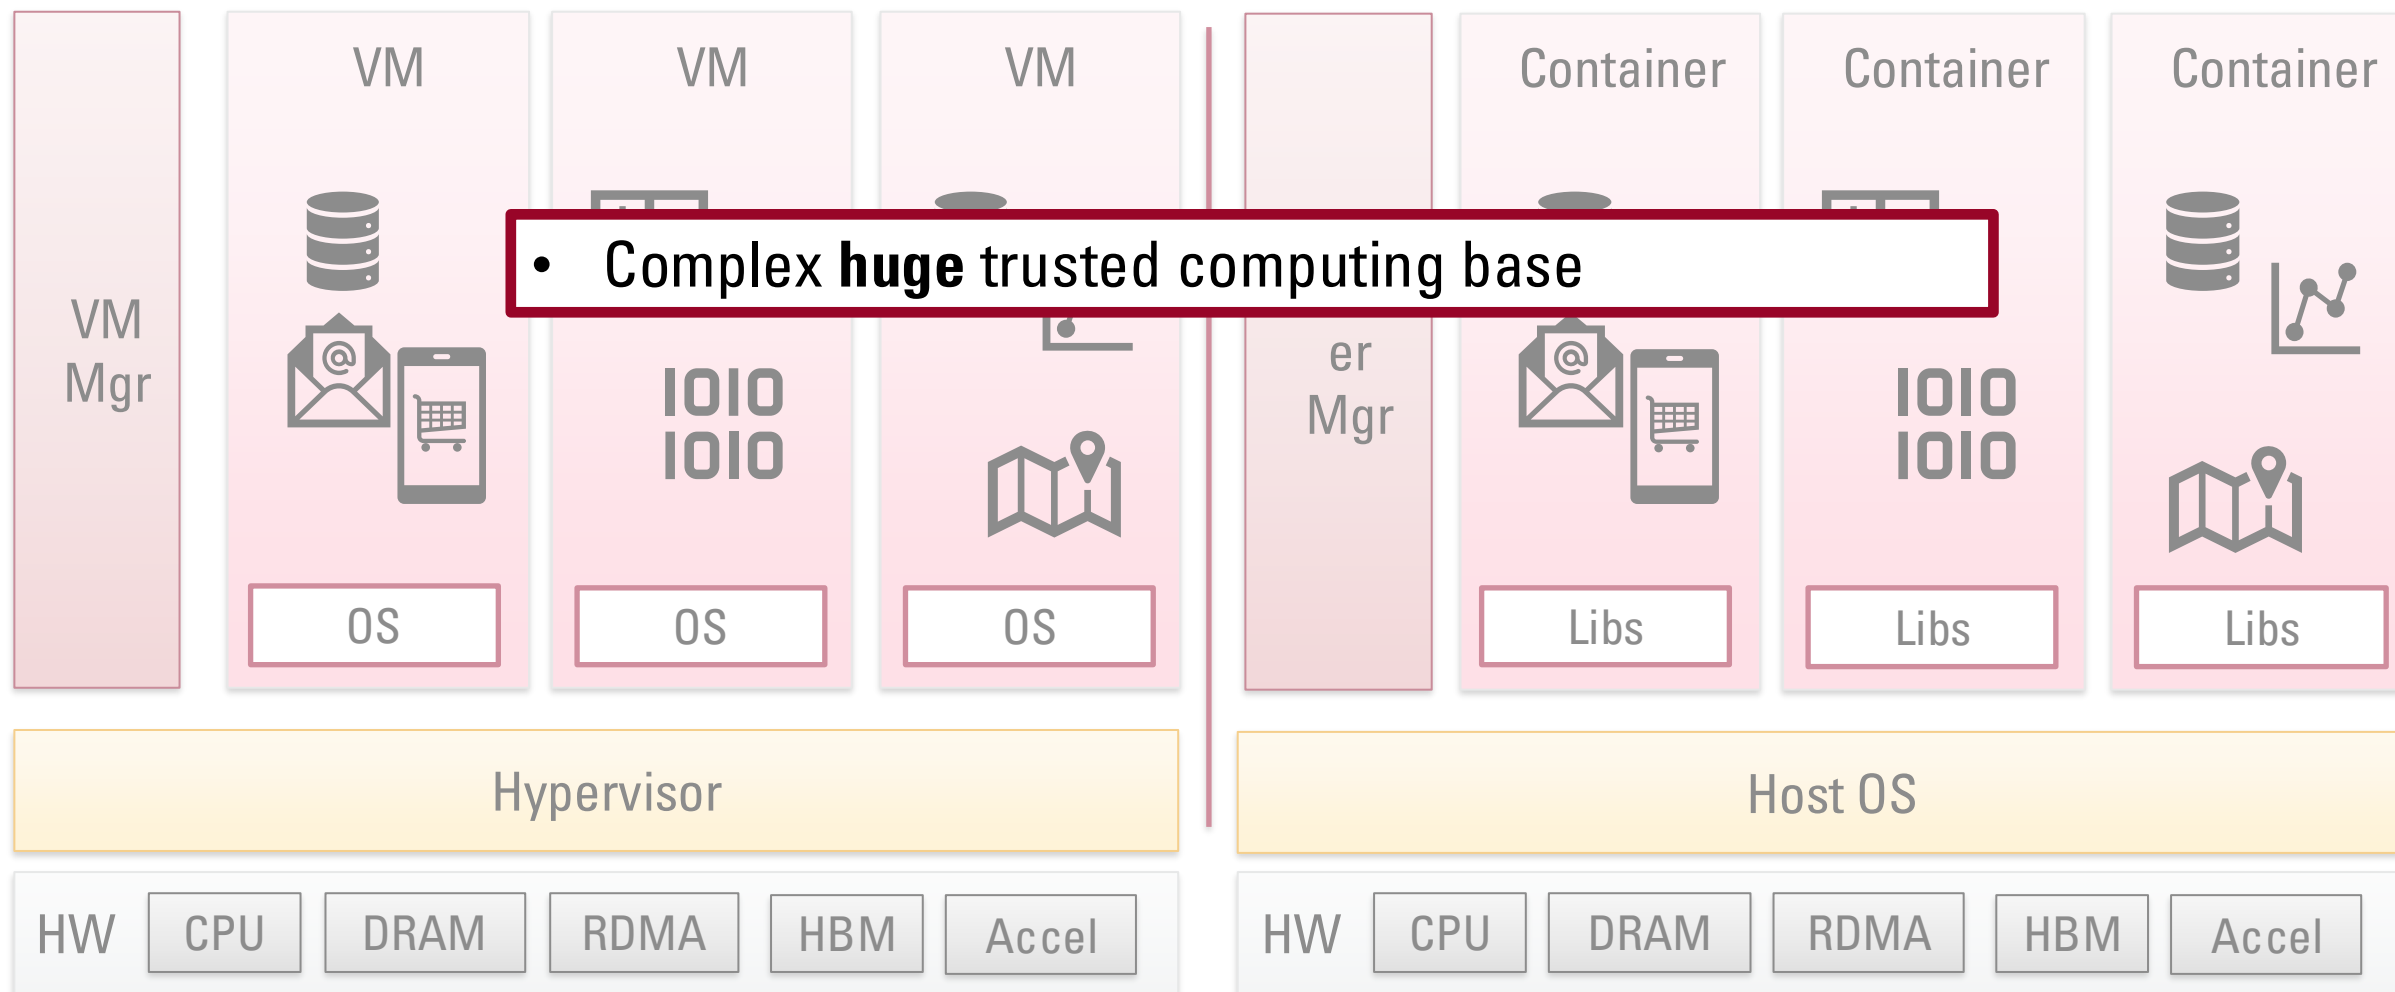
Today's Cloud Software Stack



- > 40 million lines of code for Linux 6.19
- > 1000 critical vulnerabilities (CVSS > 7) since 1999 (according to NIST NVD)

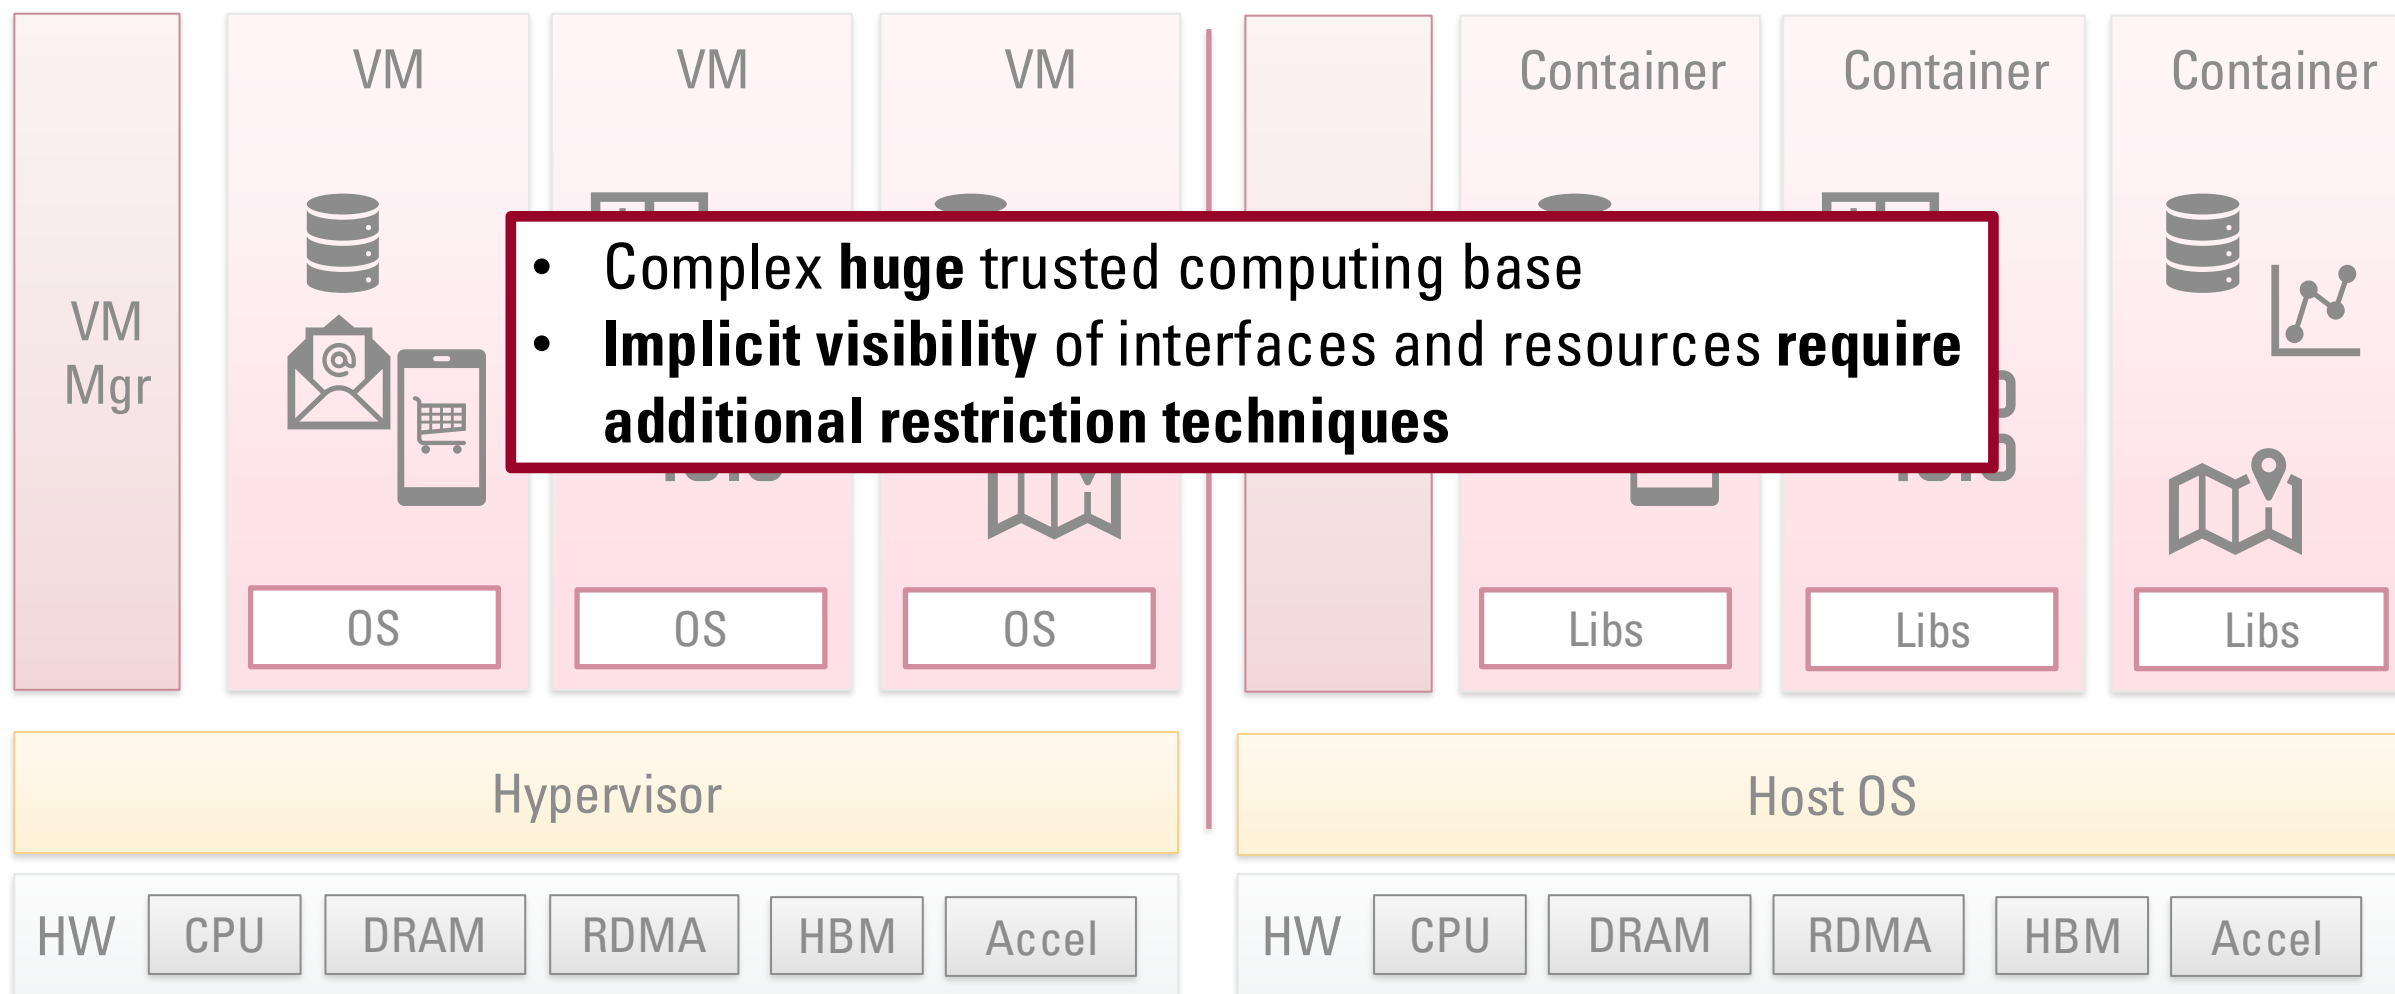


Today's Cloud Software Stack



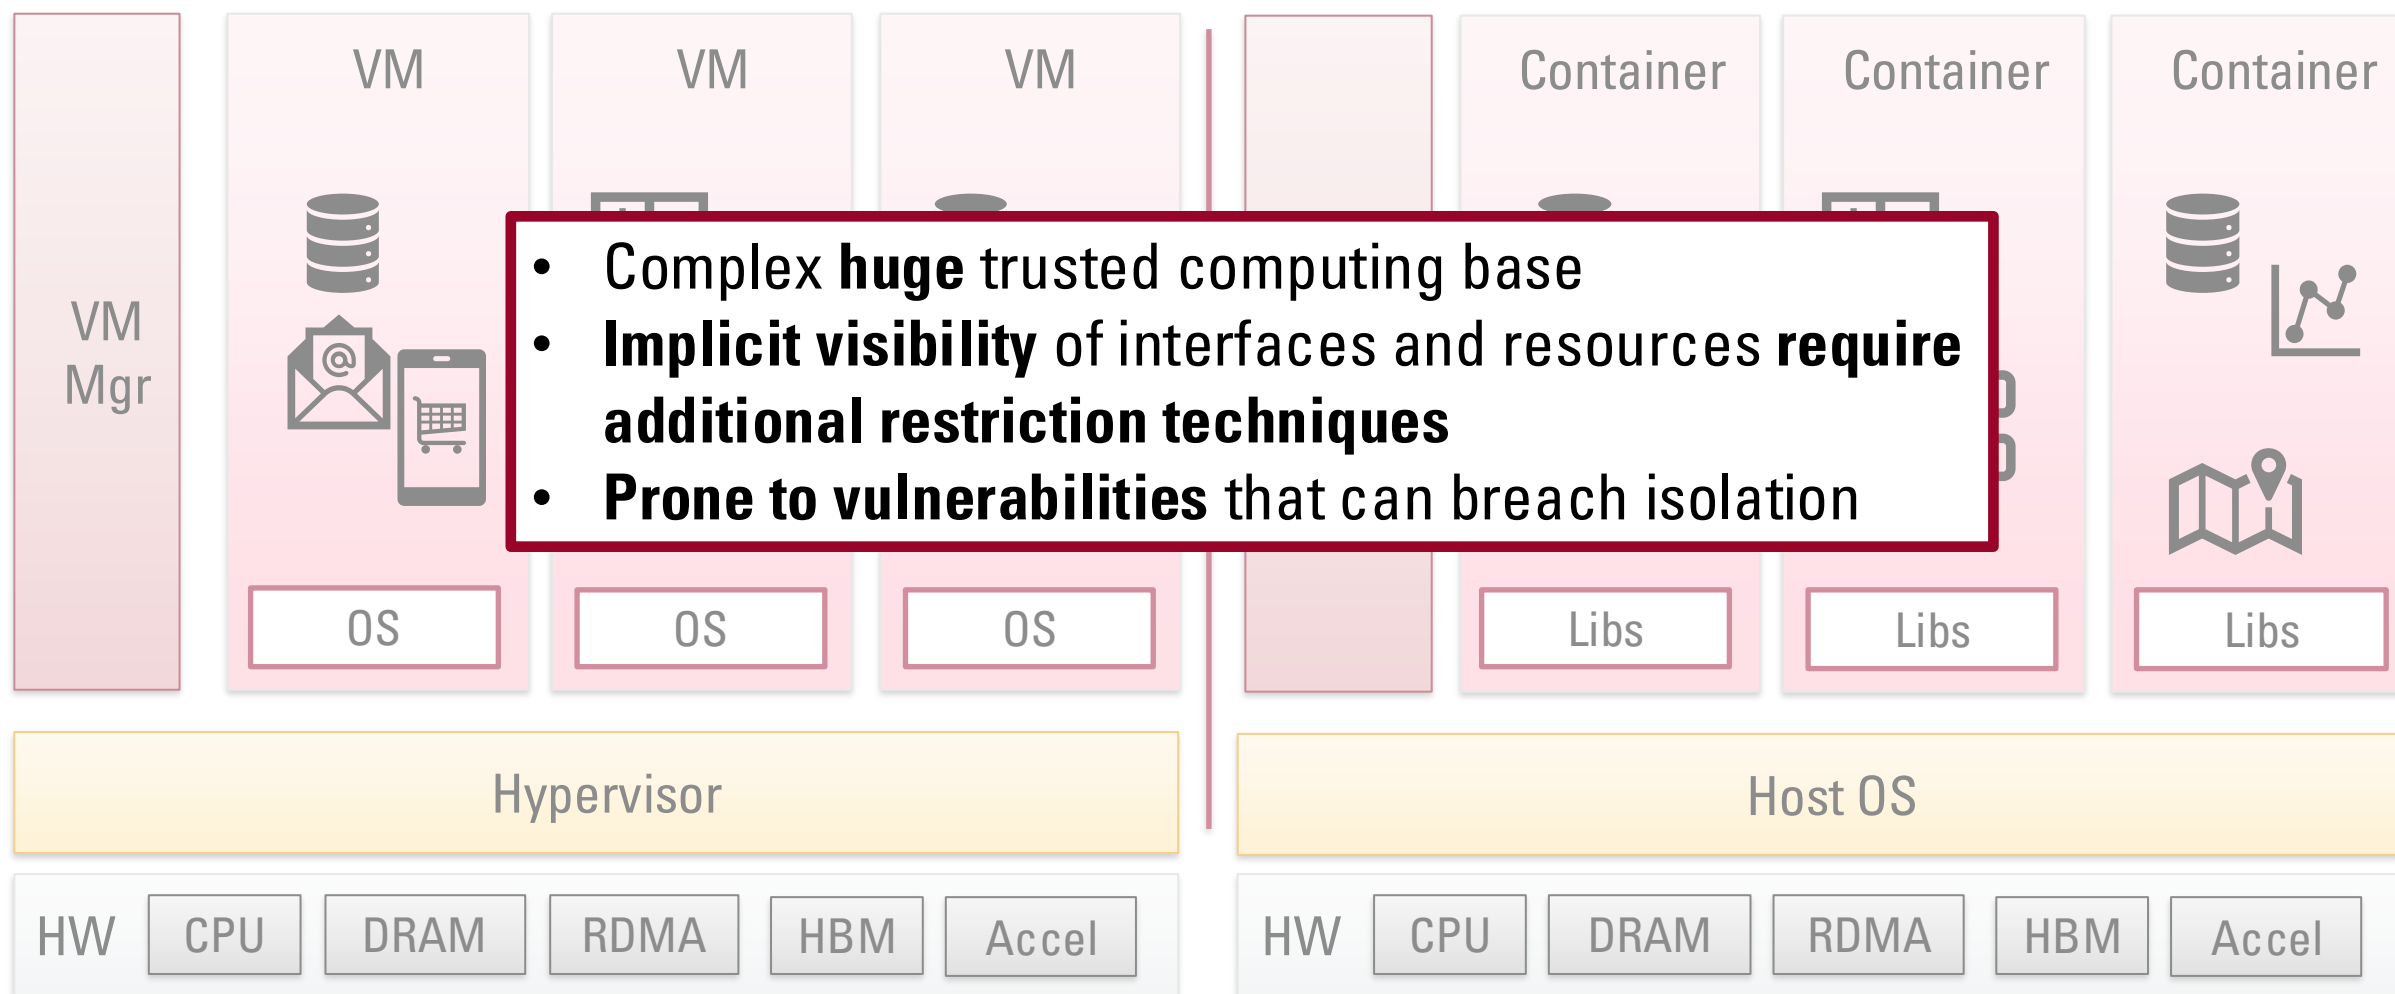


Today's Cloud Software Stack



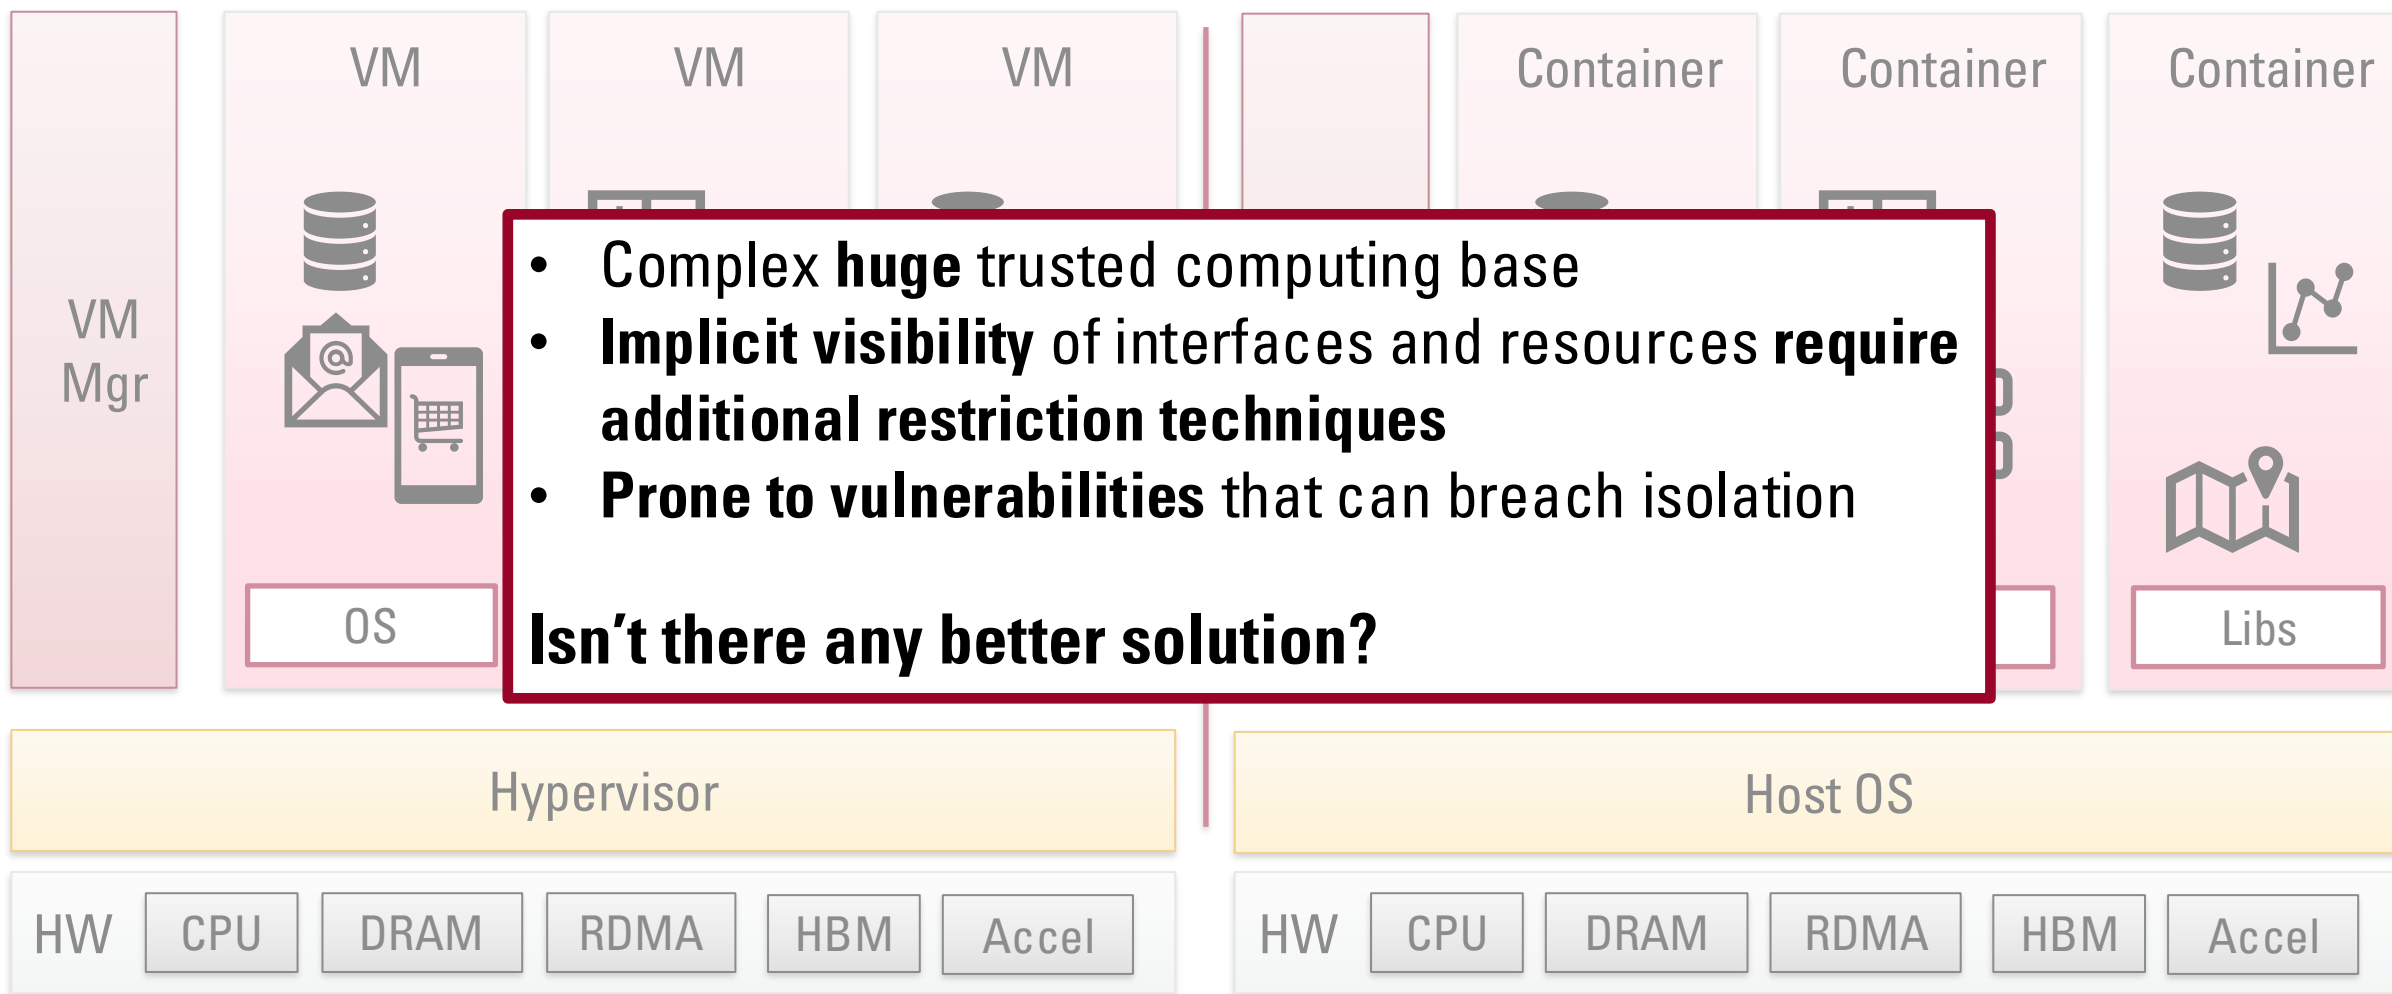


Today's Cloud Software Stack



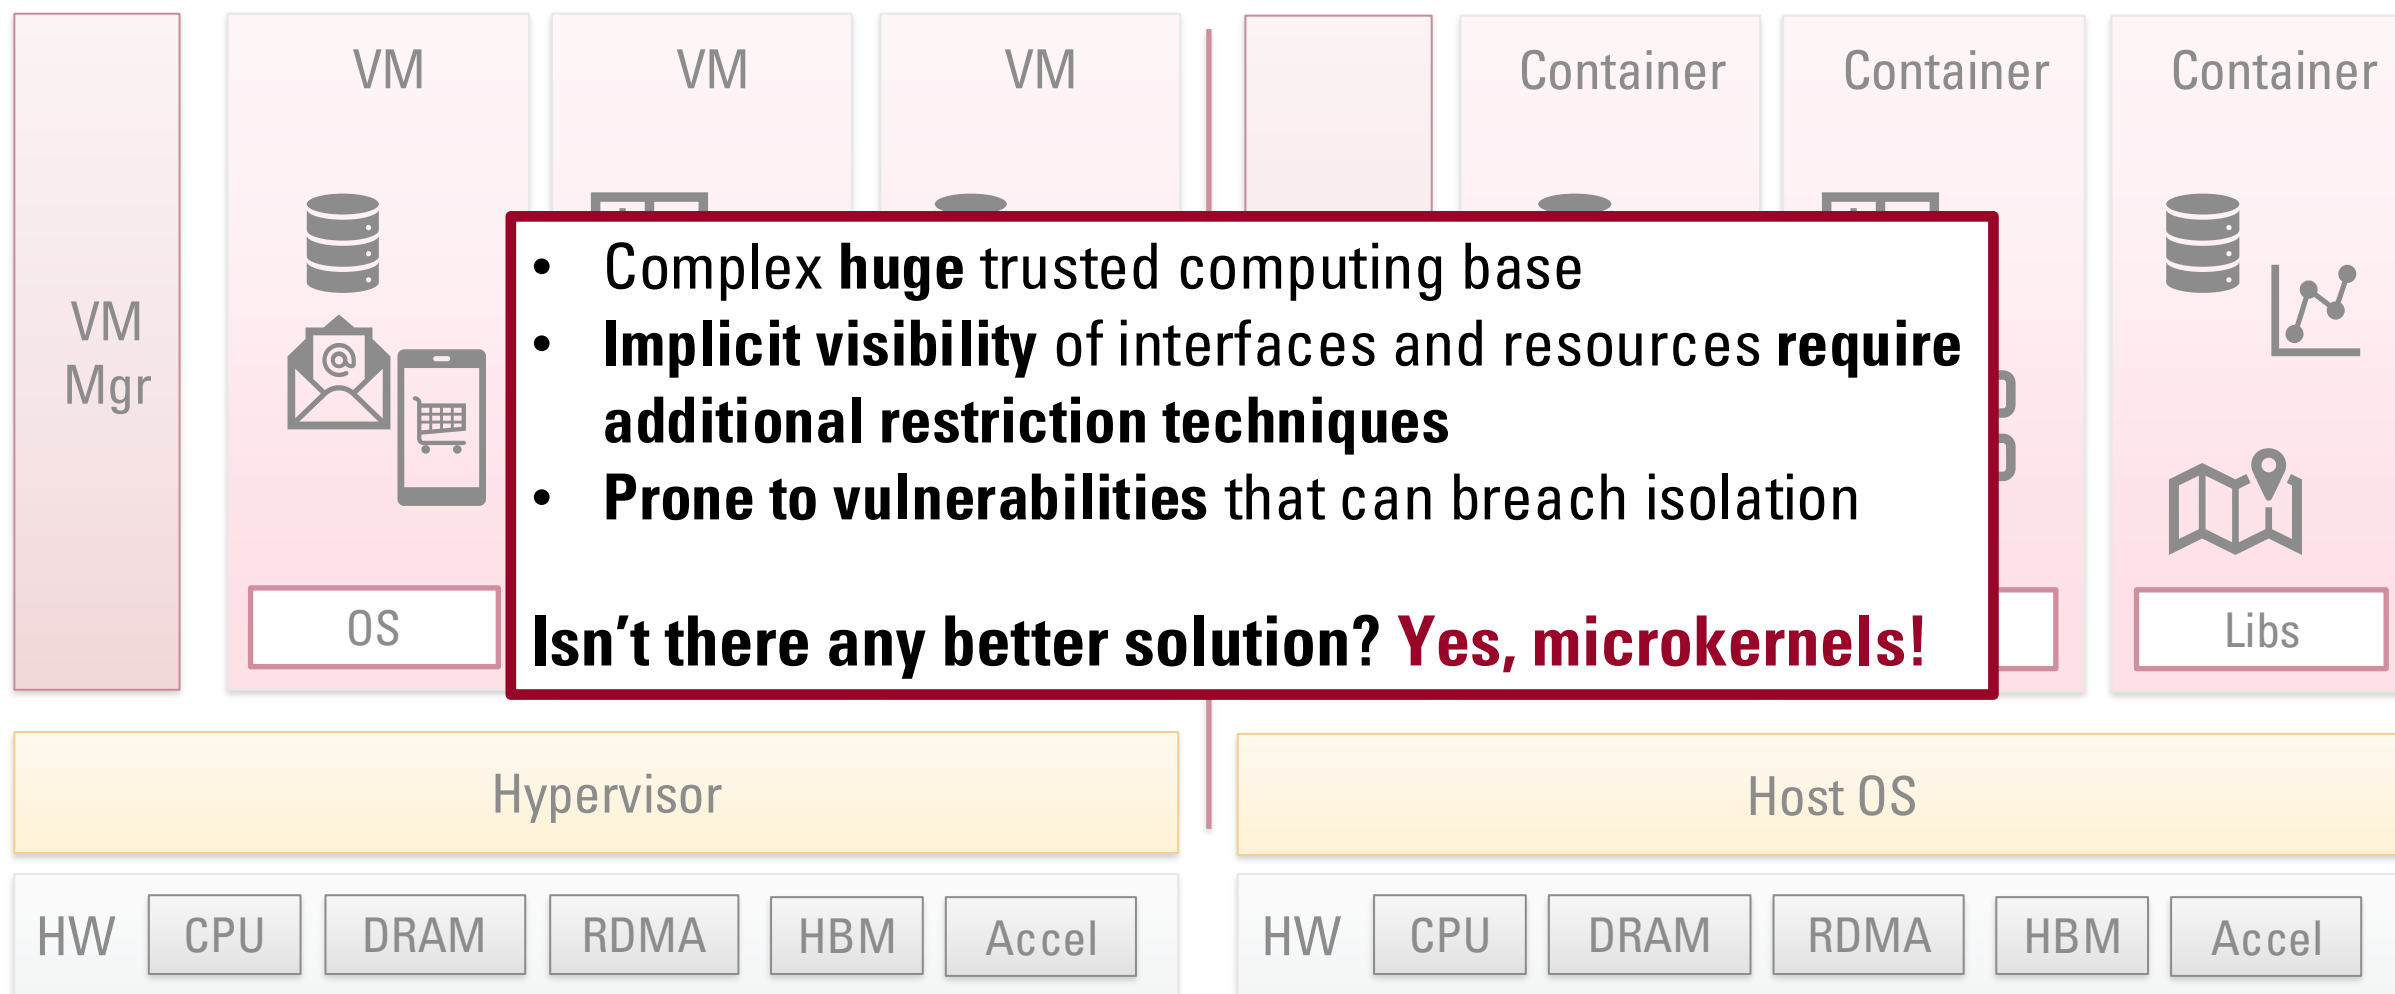


Today's Cloud Software Stack





Today's Cloud Software Stack





Meet




EalánOS

**... an experimental cloud platform based on the Genode OS
Framework**

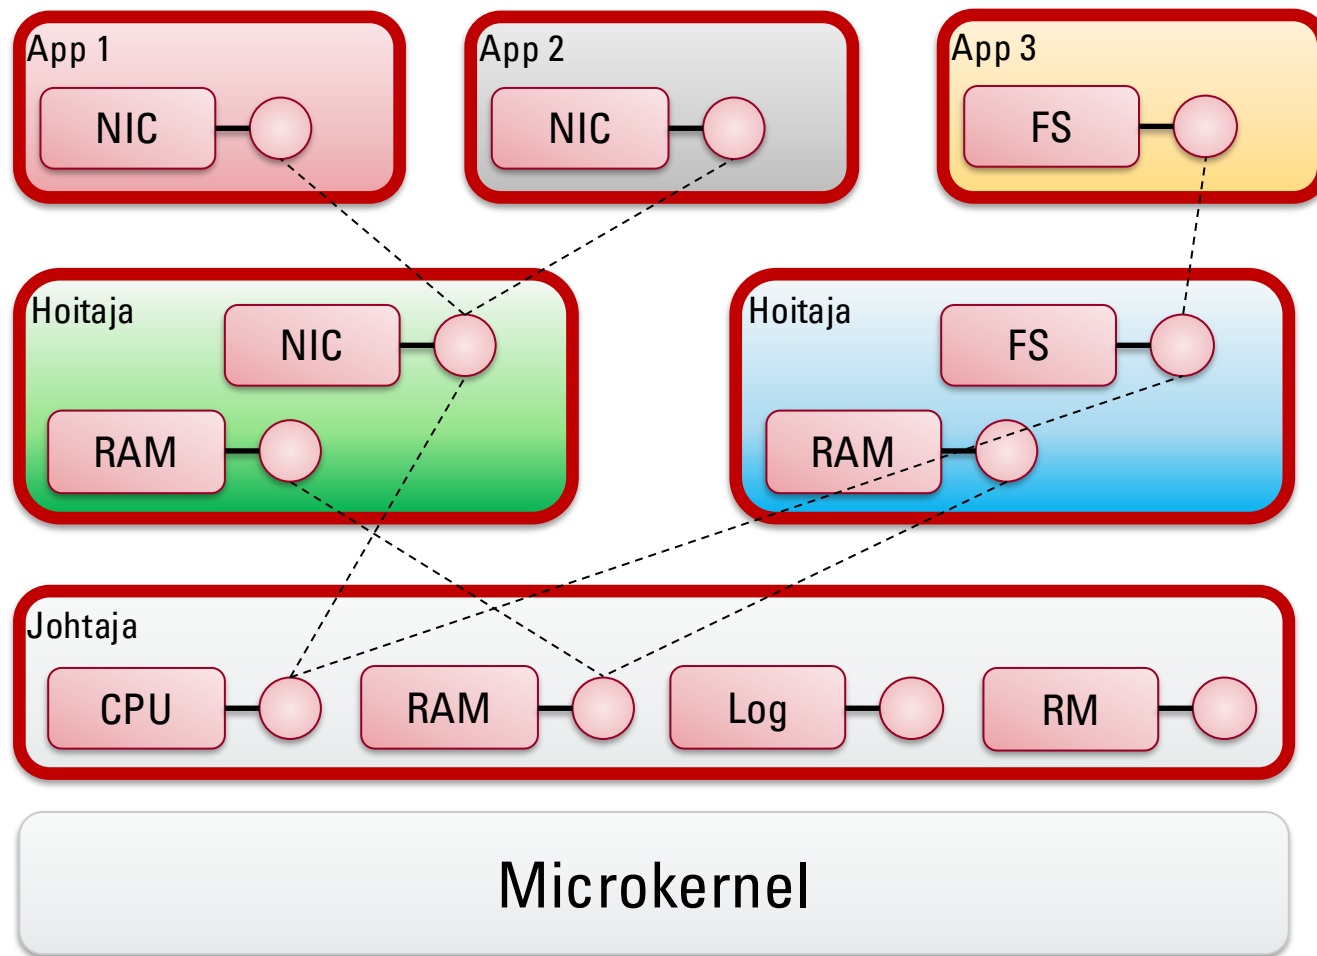


EalánOS Architecture





EalánOS Architecture

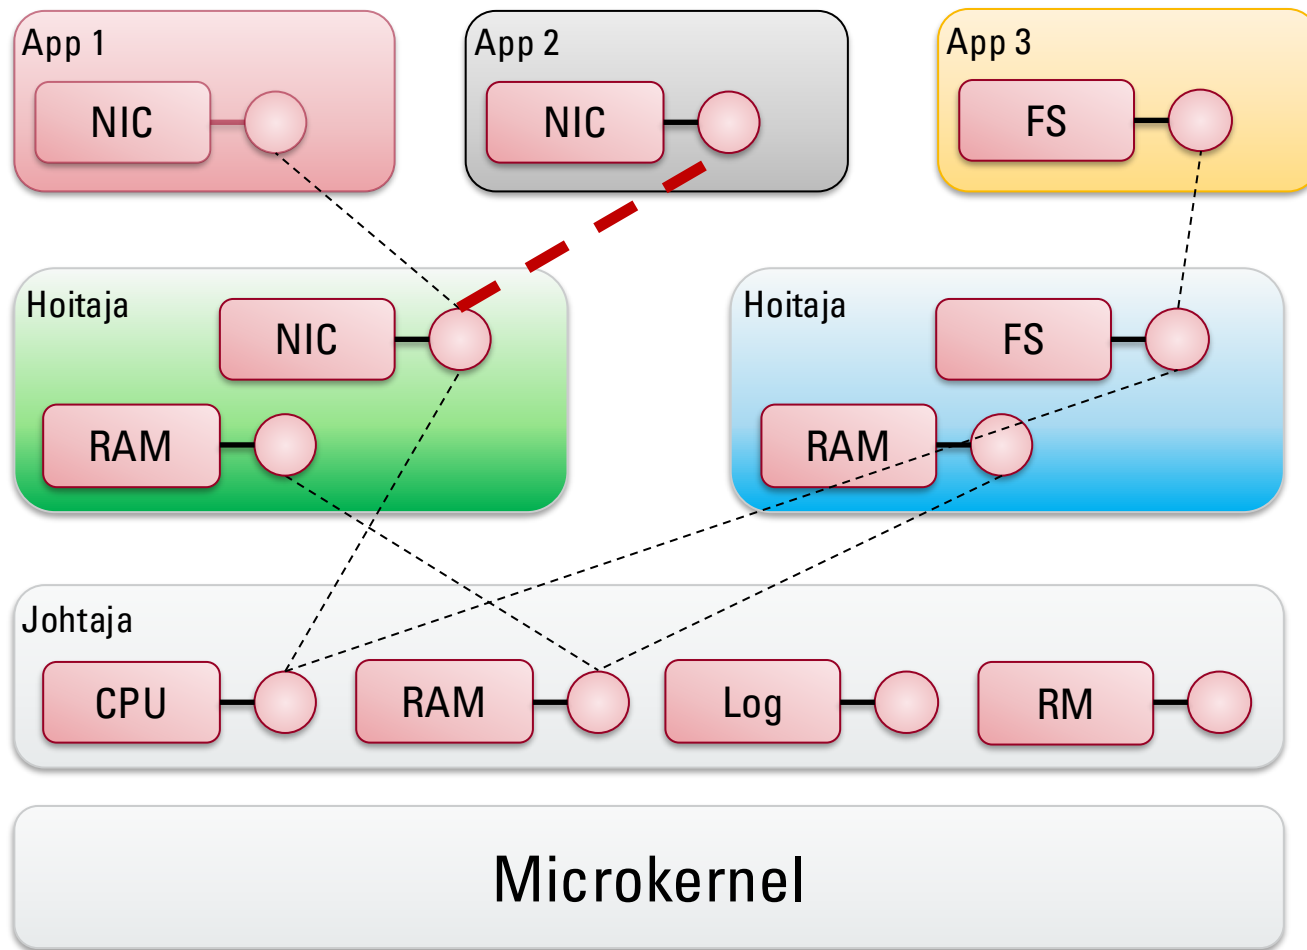


- Composed of components*

*) Based on the architecture of the Genode OS Framework (www.genode.org)



EalánOS Architecture

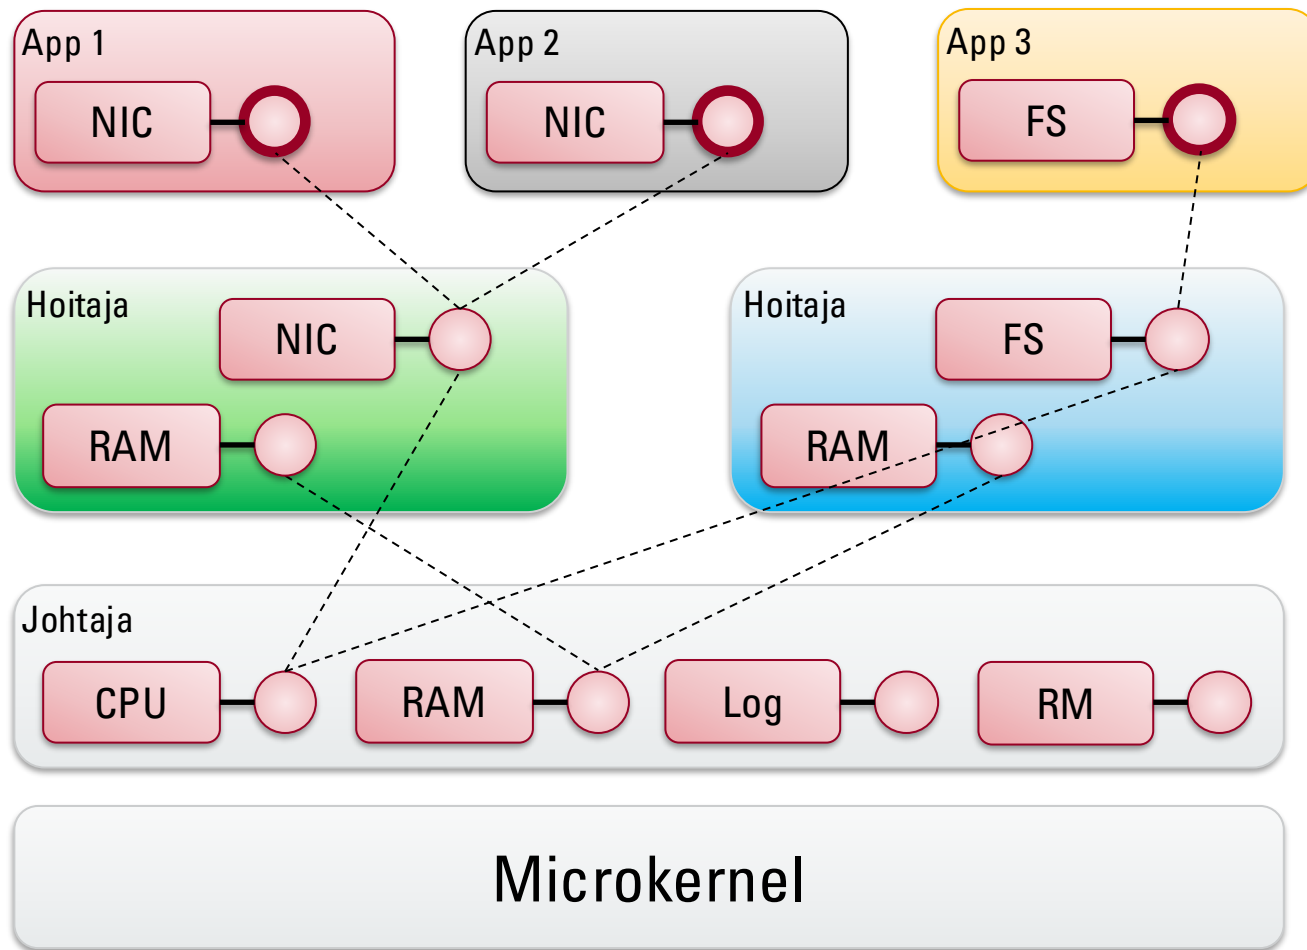


- Composed of components*
- Components interact via RPC

*) Based on the architecture of the Genode OS Framework (www.genode.org)



EalánOS Architecture

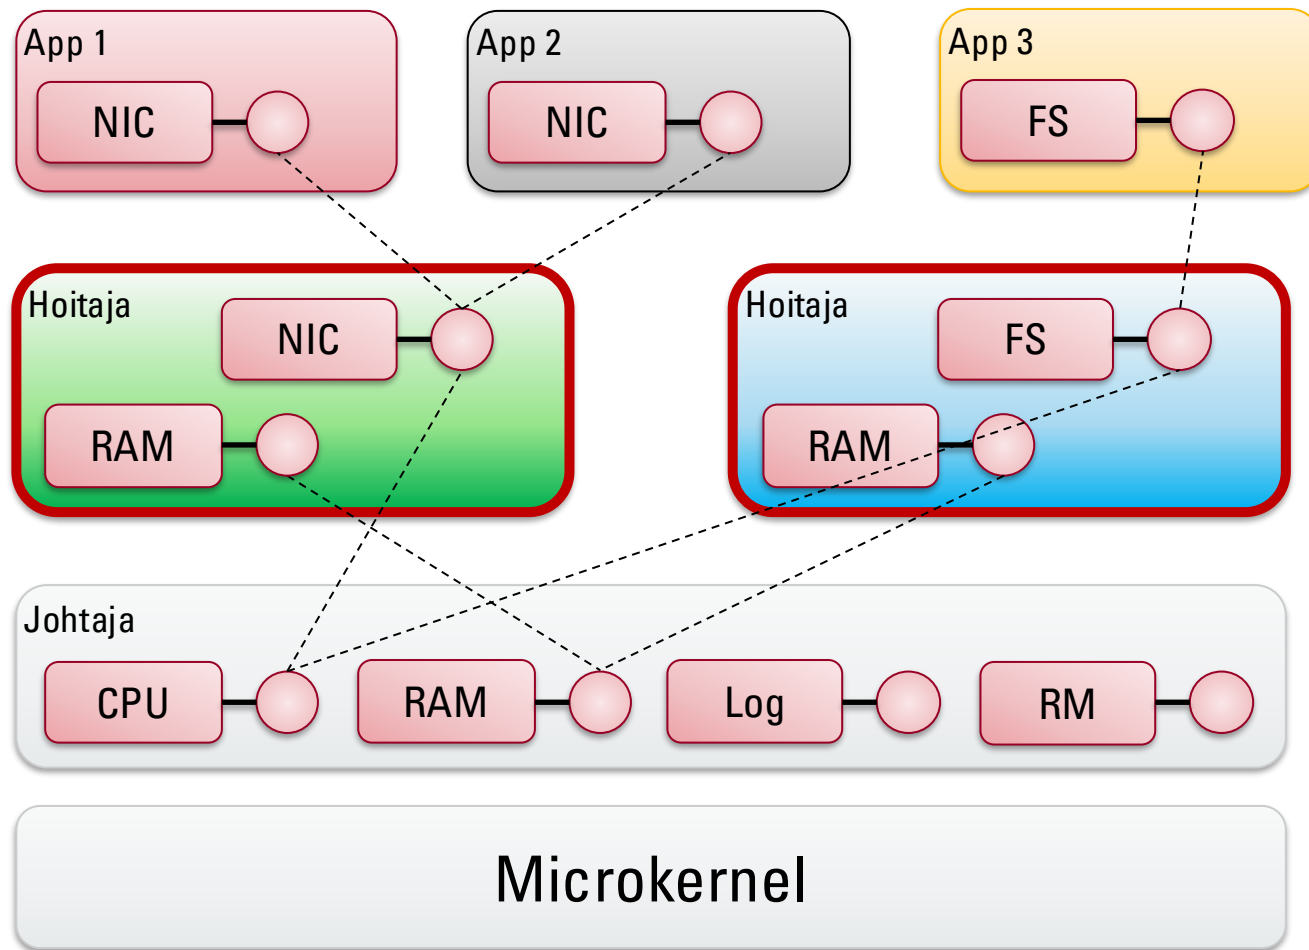


- Composed of components*
- Components interact via RPC
- Access only to resources for which capability was granted

*) Based on the architecture of the Genode OS Framework (www.genode.org)



EalánOS Architecture

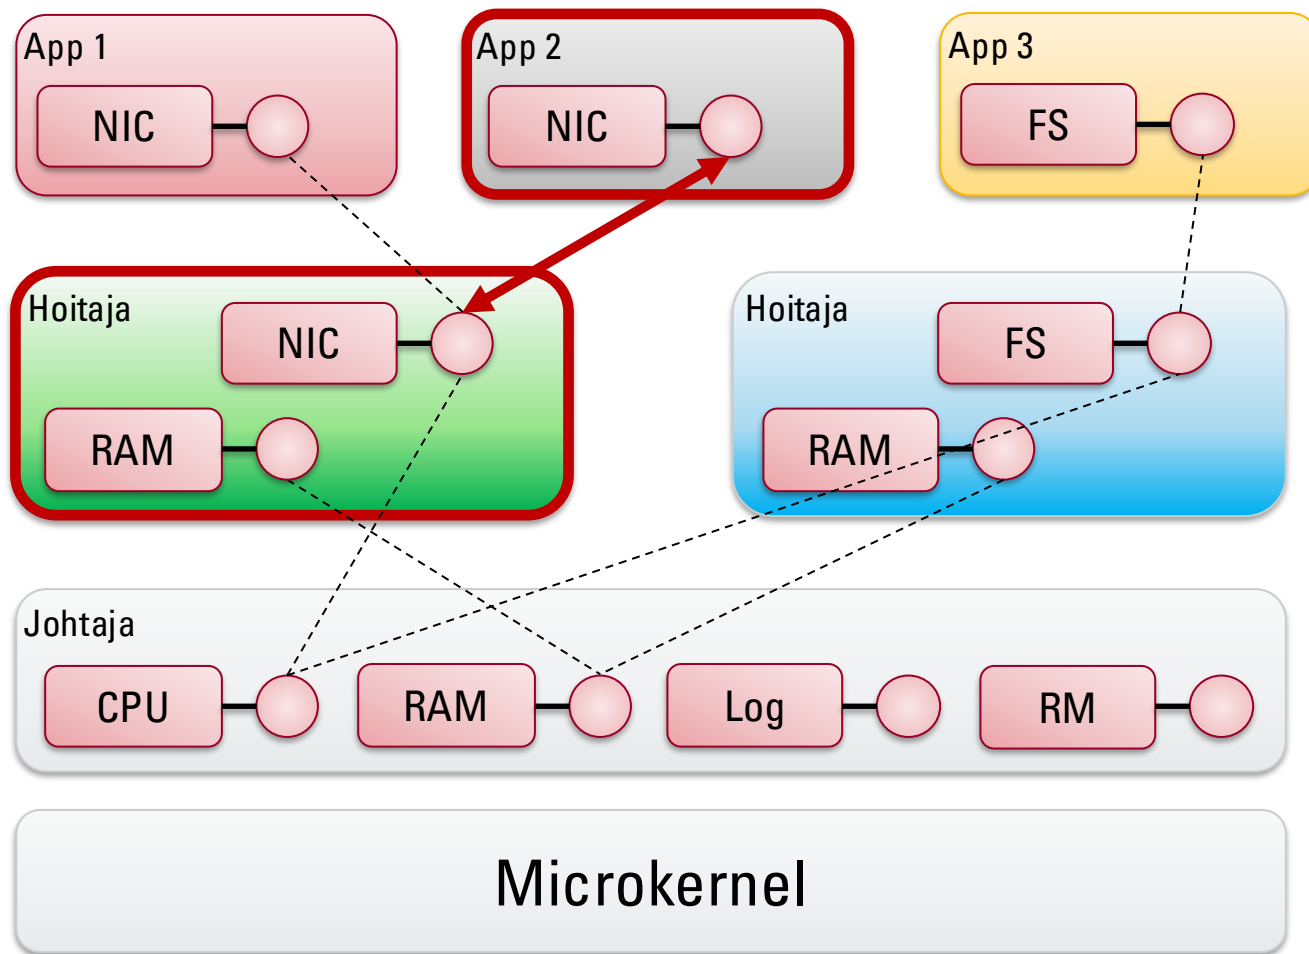


- Composed of components*
- Components interact via RPC
- Access only to resources for which capability was granted
- Parents control access to services and resources

*) Based on the architecture of the Genode OS Framework (www.genode.org)



EalánOS Architecture

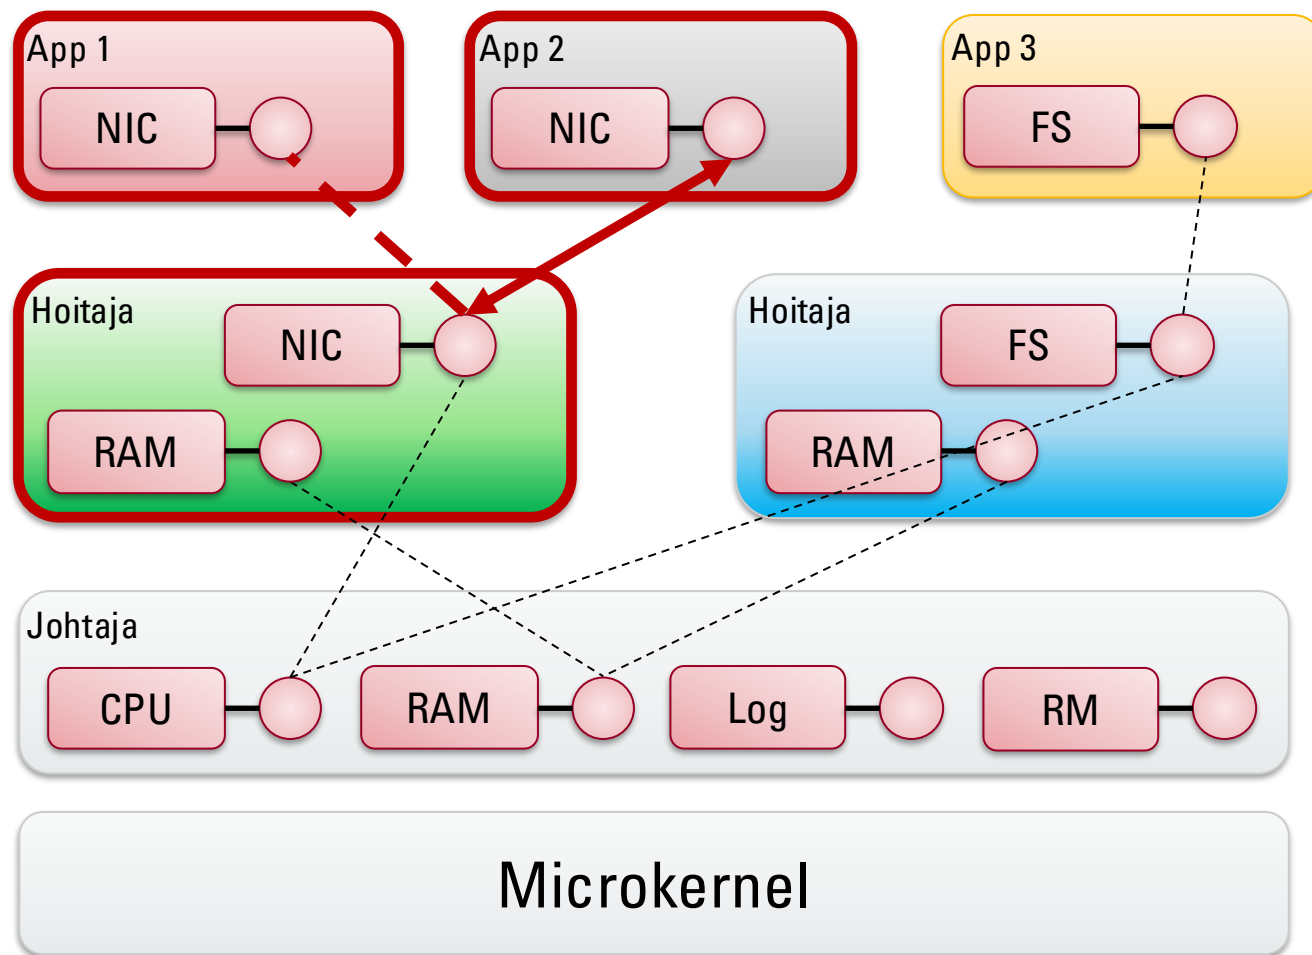


- Composed of components*
- Components interact via RPC
- Access only to resources for which capability was granted
- Parents control access to services and resources
- Children request resource access from parent

*) Based on the architecture of the Genode OS Framework (www.genode.org)



EalánOS Architecture

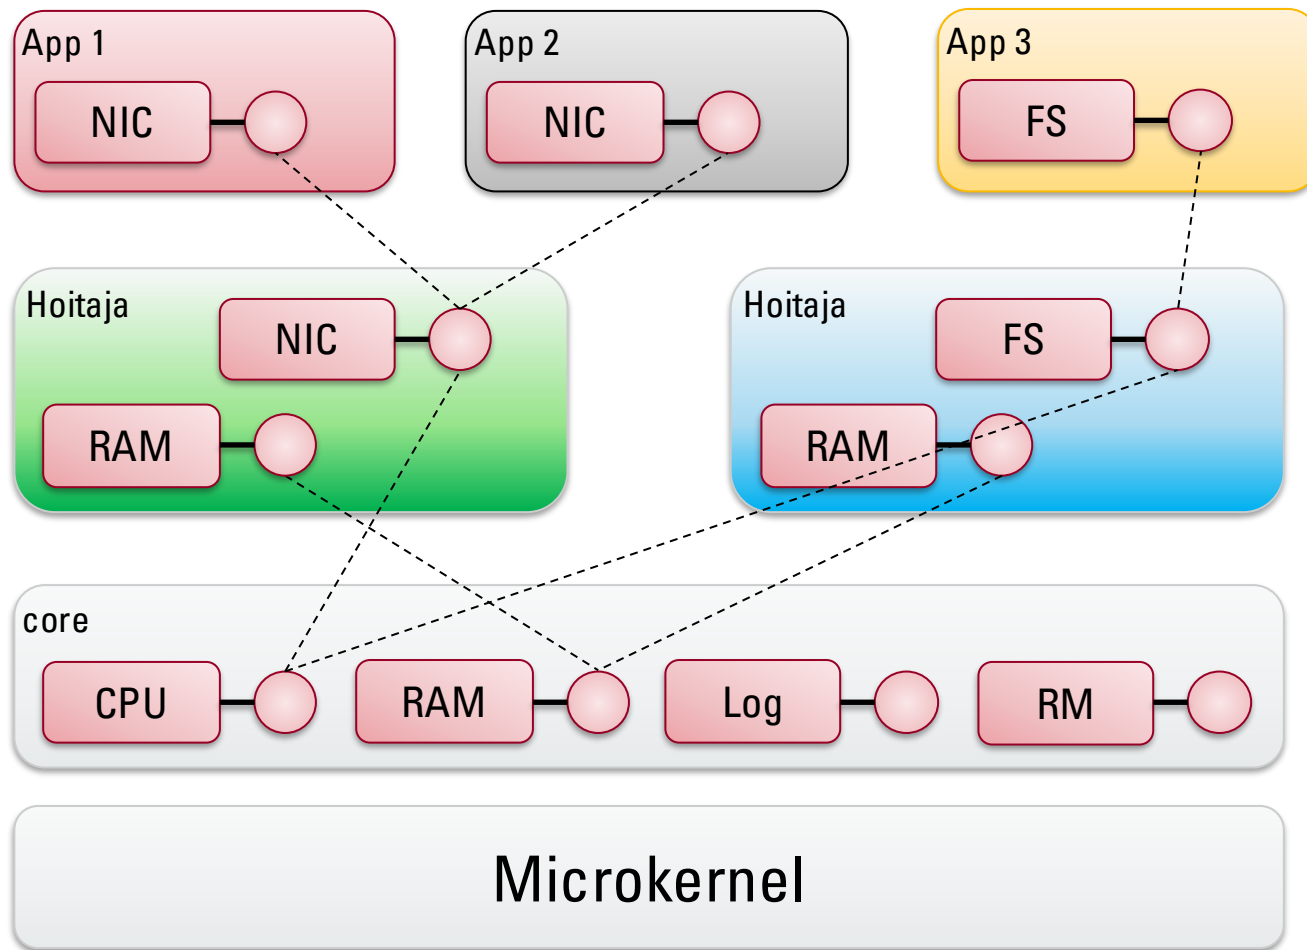


- Composed of components*
- Components interact via RPC
- Access only to resources for which capability was granted
- Parents control access to services and resources
- Children request resource access from parent
- Habitats to isolate tenants

*) Based on the architecture of the Genode OS Framework (www.genode.org)



EalánOS Architecture




- No global namespace

Advantages



EalánOS Architecture




- No global namespace
- Explicit access control via capabilities

Advantages



EalánOS Architecture




- No global namespace
- Explicit access control via capabilities
- Small trusted computing base

Advantages



EalánOS Architecture

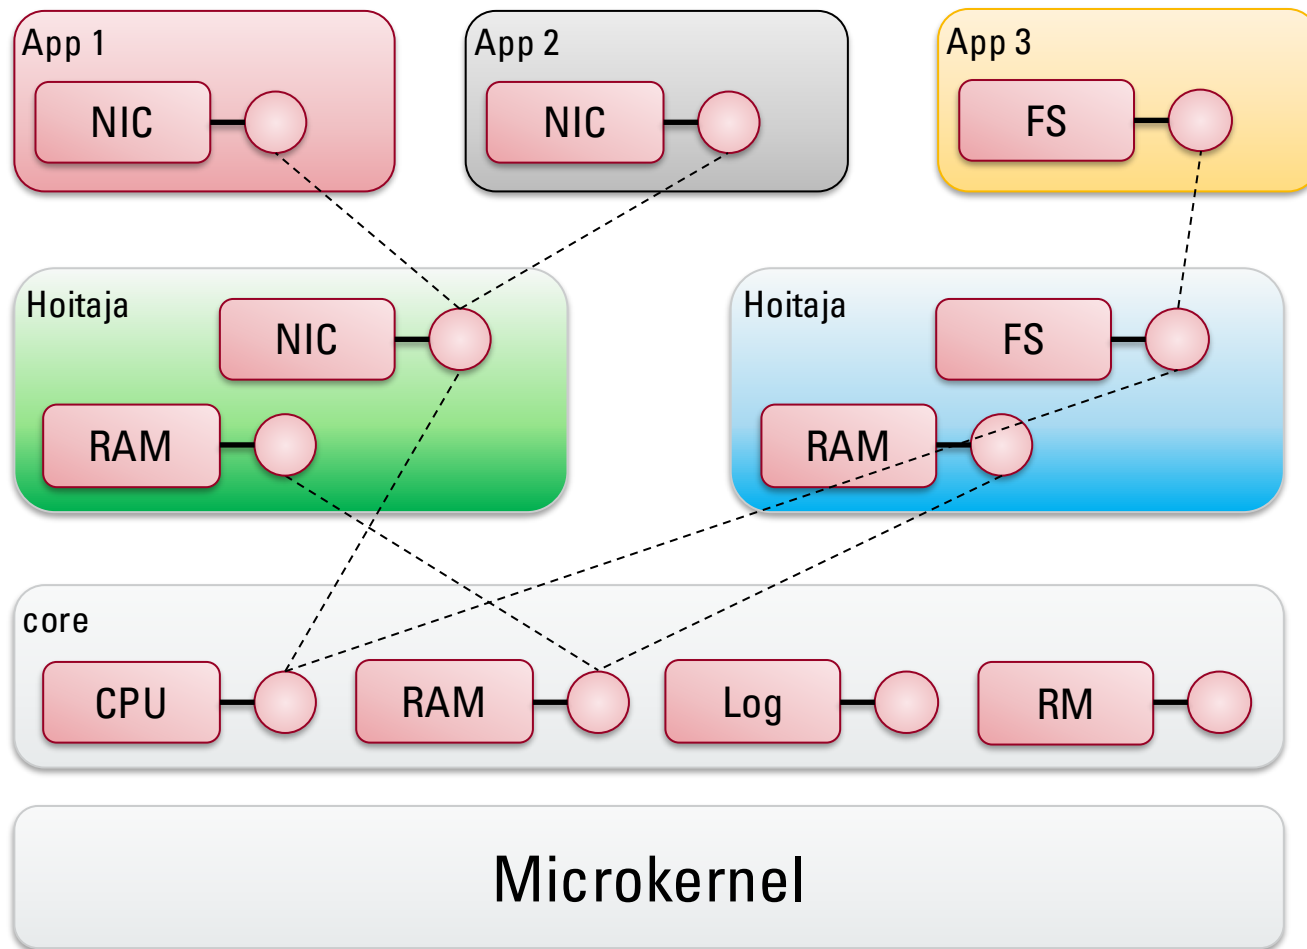


- No global namespace
- Explicit access control via capabilities
- Small trusted computing base
- Fault containment and modularity

Advantages



EalánOS Architecture

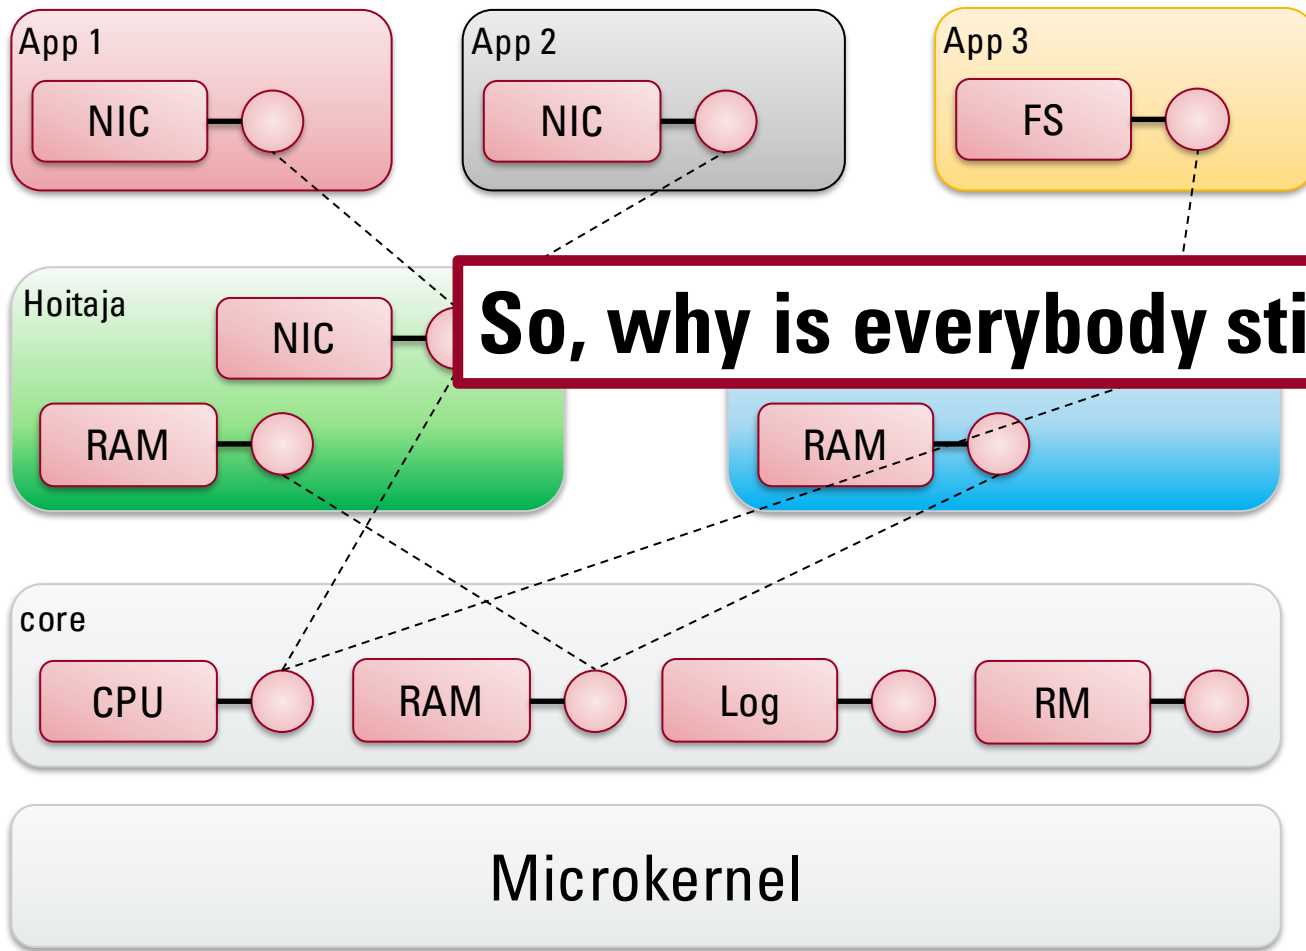


- No global namespace
- Explicit access control via capabilities
- Small trusted computing base
- Fault containment and modularity
- Certification of microkernel possible

Advantages



EalánOS Architecture



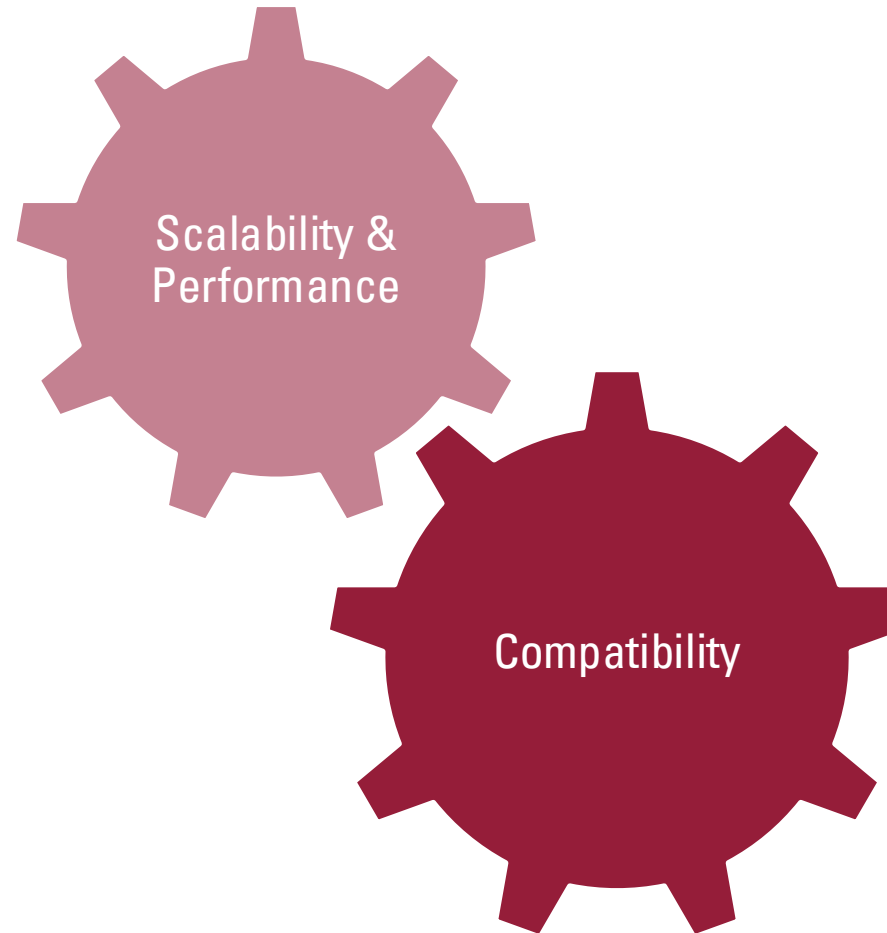
So, why is everybody still using Linux?

- No global namespace
- Explicit access control via capabilities
- Modularity
- Certification of microkernel possible

Advantages



Challenges for Microkernel adoption





Challenges for Microkernel adoption



- Support for **server** hardware required

Hardware



Challenges for Microkernel adoption



- Support for **server** hardware required
 - Non-Uniform Memory (**NUMA**)

Hardware



Challenges for Microkernel adoption



- Support for **server** hardware required
 - Non-Uniform Memory (**NUMA**)
 - **NICs**, RAID controller, ...

Hardware



Challenges for Microkernel adoption



- Support for **server** hardware required
 - Non-Uniform Memory (**NUMA**)
 - **NICs**, RAID controller, ...
 - **Accelerators** (GPGPUs, NPUs, ...)

Hardware



Challenges for Microkernel adoption



- Support for **server** hardware required
 - Non-Uniform Memory (**NUMA**)
 - **NICs**, RAID controller, ...
 - **Accelerators** (GPGPUs, NPUs, ...)
 - TPM, SGX

Hardware



Challenges for Microkernel adoption



- Support for **server** hardware required
 - Non-Uniform Memory (**NUMA**)
 - **NICs**, RAID controller, ...
 - **Accelerators** (GPGPUs, NPUs, ...)
 - TPM, SGX
- Solution: Use Linux drivers in **user-space**

Hardware



Challenges for Microkernel adoption



- Support for **server** hardware required
 - Non-Uniform Memory (**NUMA**)
 - **NICs**, RAID controller, ...
 - **Accelerators** (GPGPUs, NPUs, ...)
 - TPM, SGX
- Solution: Use Linux drivers in **user-space**
 - Genode's DDE Linux (Linux 6.12)
 - Provides compatibility layer

Hardware



Challenges for Microkernel adoption



- Support for **POSIX-style** applications

Software



Challenges for Microkernel adoption



- Support for **POSIX-style** applications
- **Tool** support required (for **porting** and **administration**)

Software



Challenges for Microkernel adoption



- Support for **POSIX-style** applications
- **Tool** support required (for **porting** and **administration**)
- Genode/EalánOS offers **Goa** toolkit
 - For **porting** 3rd party software
 - **Debugging, Testing, Publishing**

Software



Challenges for Microkernel adoption



- Support for **POSIX-style** applications
- **Tool** support required (for **porting** and **administration**)
- Genode/EalánOS offers **Goa** toolkit
 - For **porting** 3rd party software
 - **Debugging, Testing, Publishing**
- EalánOS provides **Kuori** remote shell for
 - **managing** components at **runtime**
 - **Monitoring** resource **utilization**


Software



Challenges for Microkernel adoption



- Larson memory allocation benchmark [4]
- 5,000,000 malloc() and free() calls each
- 8B-1kB chunks randomly selected






Challenges for Microkernel adoption



- TCP echo benchmark
- Comparison against state-of-the-art DPDK-based network stack Caladan [3]



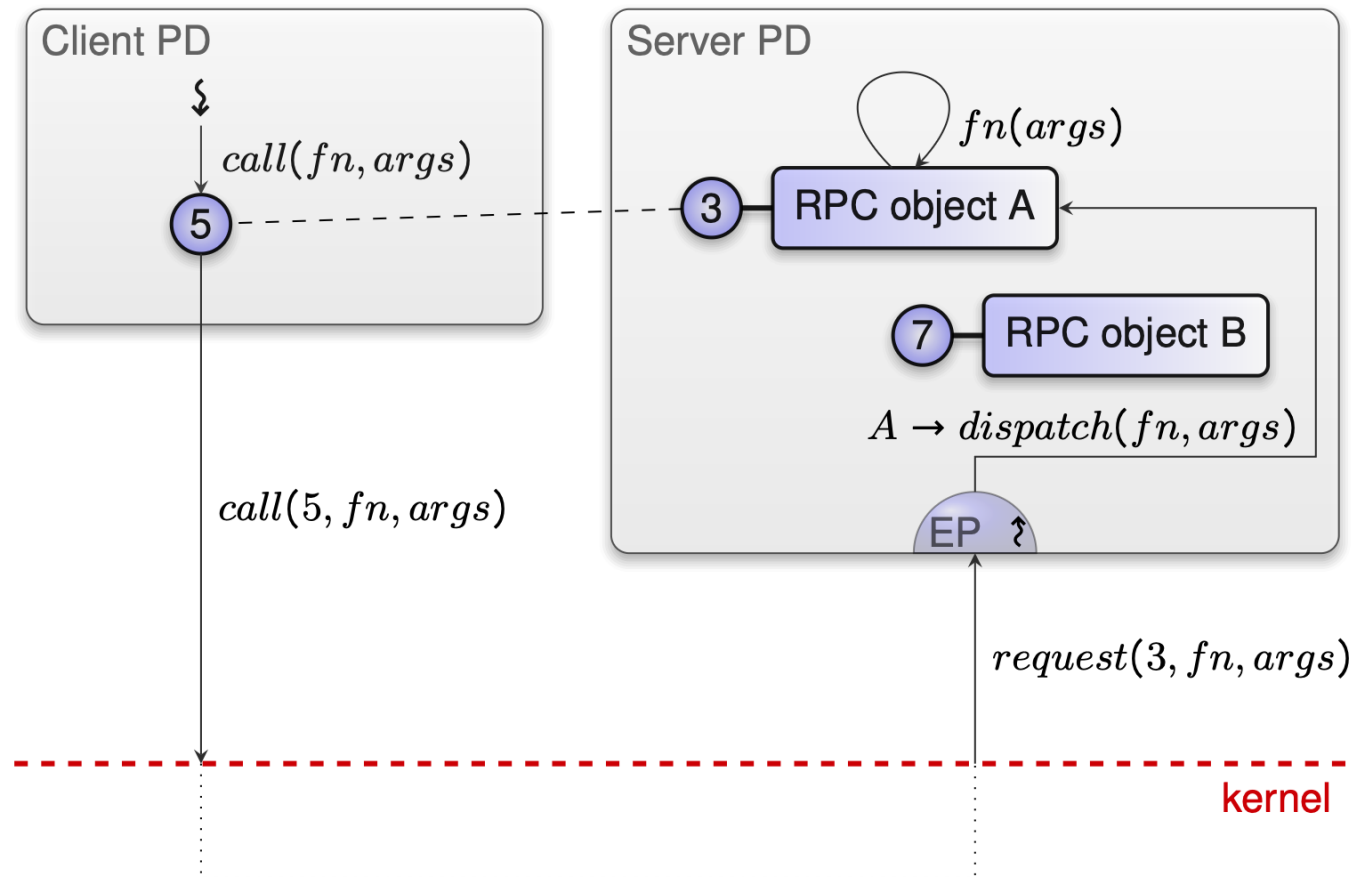
TCP performance



Challenges for Microkernel adoption



- Replace non-scalable locks
- Use wait-free synchronization



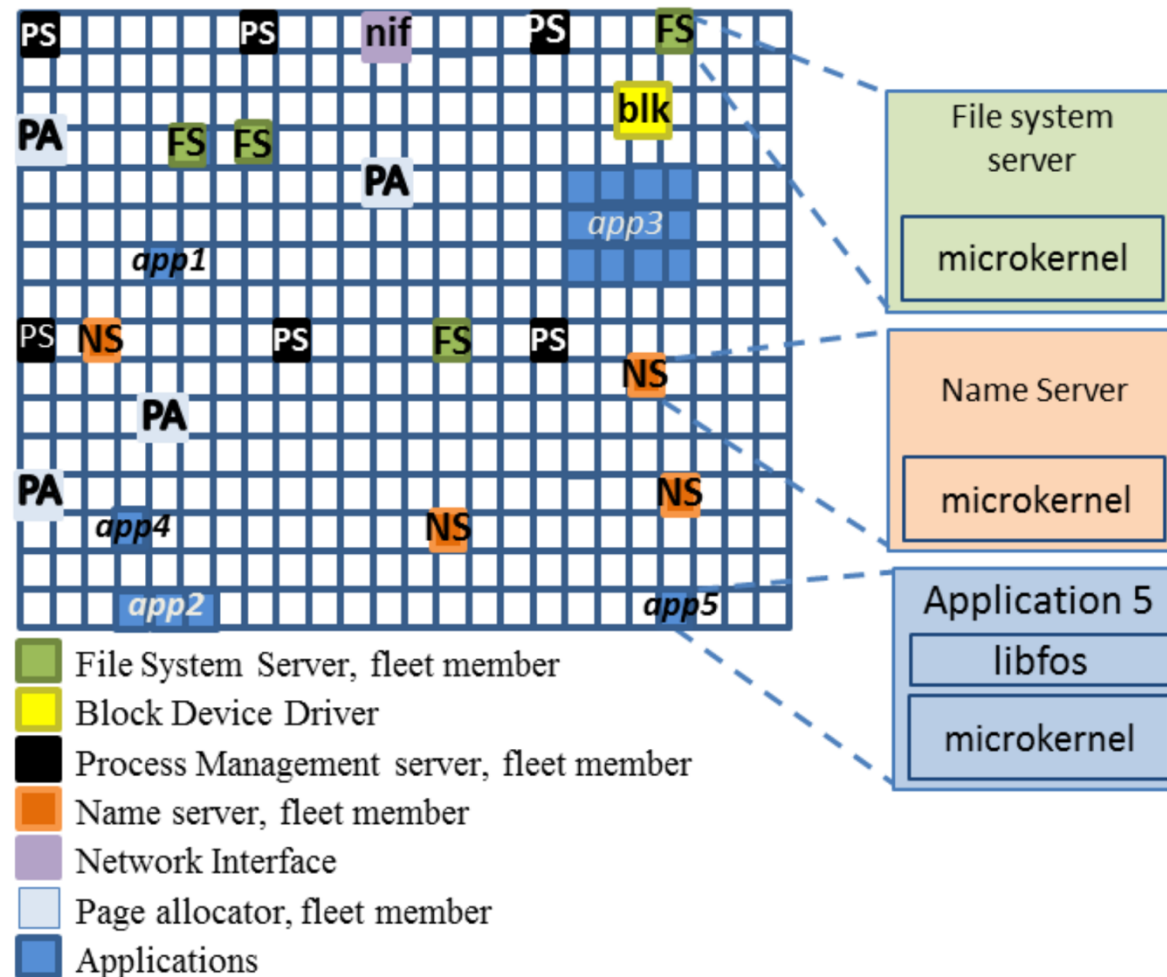
Source: Genode Foundations [2]



Challenges for Microkernel adoption



- Multi-threaded service fleets
- Like in Factored Operating System (fos) [1]





Future Roadmap for EalánOS

- Porting of common server applications (nginx, postgres, ...)
- Advancing of tooling for remote management
- Redesign of Genode's network infrastructure
- Evaluation and redesign of Genode's page (dataspace) allocator



Summary

- Current cloud infrastructure is based mostly on Linux
 - Huge and complex trusted computing base
 - Riddled with critical vulnerabilities (33 for containers and 24+77 for KVM and qemu)



Summary

- Current cloud infrastructure is based mostly on Linux
 - Huge and complex trusted computing base
 - Riddled with critical vulnerabilities (33 for containers and 24+77 for KVM and qemu)
- Microkernels offer greater security, safety and modularity



Summary

- Current cloud infrastructure is based mostly on Linux
 - Huge and complex trusted computing base
 - Riddled with critical vulnerabilities (33 for containers and 24+77 for KVM and qemu)
- Microkernels offer greater security, safety and modularity
- However, no wide adoption yet due to



Summary

- Current cloud infrastructure is based mostly on Linux
 - Huge and complex trusted computing base
 - Riddled with critical vulnerabilities (33 for containers and 24+77 for KVM and qemu)
- Microkernels offer greater security, safety and modularity
- However, no wide adoption yet due to
 - Lack of hardware drivers for server-grade hardware



Summary

- Current cloud infrastructure is based mostly on Linux
 - Huge and complex trusted computing base
 - Riddled with critical vulnerabilities (33 for containers and 24+77 for KVM and qemu)
- Microkernels offer greater security, safety and modularity
- However, no wide adoption yet due to
 - Lack of hardware drivers for server-grade hardware
 - Lack of tool and software support



Summary

- Current cloud infrastructure is based mostly on Linux
 - Huge and complex trusted computing base
 - Riddled with critical vulnerabilities (33 for containers and 24+77 for KVM and qemu)
- Microkernels offer greater security, safety and modularity
- However, no wide adoption yet due to
 - Lack of hardware drivers for server-grade hardware
 - Lack of tool and software support
 - Issues with scalability of system services



Summary

- Current cloud infrastructure is based mostly on Linux
 - Huge and complex trusted computing base
 - Riddled with critical vulnerabilities (33 for containers and 24+77 for KVM and qemu)
- Microkernels offer greater security, safety and modularity
- However, no wide adoption yet due to
 - Lack of hardware drivers for server-grade hardware
 - Lack of tool and software support
 - Issues with scalability of system services
- But there's ongoing work to overcome those challenges by Genode Labs, Gapfruit (Genode) and Kernkonzept, Barkhausen Institute (for L4Re) and me



Thank you for your attention

- github.com/mmueller41/genode
- github.com/mmueller41/NOVA

The logo for EalánOS, featuring a green fern frond on the left and the text 'EalánOS' in a green, stylized font on the right.

Prototypical implementation of a
microkernel cloud architecture

Get in touch

Embedded
Software Systems



- ess.cs.uos.de
- www.uni-osnabrueck.de



- mxkernel.org
- dfg-spp2037.org



References

- 1) Agarwal, A., Miller, J., Wentzlaff, D., Kasture, H., Beckmann, N., Gruenwald Iii, C. and Johnson, C. 2012. FOS: A Factored Operating Systems for High Assurance and Scalability on Multicores. (Aug. 2012).
- 2) Feske, N. *Genode Operating System Framework Foundations*.
- 3) Fried, J., Ruan, Z., Ousterhout, A. and Belay, A. 2020. Caladan: Mitigating Interference at Microsecond Timescales. *14th USENIX symposium on operating systems design and implementation, OSDI 2020, virtual event, november 4-6, 2020*(2020), 281–297.
- 4) Larson, P.-Å. and Krishnan, M. 1998. Memory allocation for long-running server applications. *Proceedings of the 1st international symposium on memory management*(Vancouver, British Columbia, Canada, 1998), 176–185.
- 5) Miemietz, T., Reusch, V., Hille, M., Wrenger, L., Eisoldt, J., Klötzke, J., Kurze, M., Lackorzynski, A., Roitzsch, M. and Härtig, H. 2025. {MettEagle}: Costs and benefits of implementing containers on microkernels. *19th USENIX symposium on operating systems design and implementation (OSDI 25)*(2025), 979–996.