

Taming Subsystems

Capabilities as Universal Resource Access Control in L4

Adam Lackorzynski, Alexander Warg
Technische Universität Dresden
Department of Computer Science
Operating Systems Group
{adam, warg}@os.inf.tu-dresden.de

ABSTRACT

The embedded and mobile computing market with its wide range of innovations is expected to remain growing in the foreseeable future. Recent developments in the embedded computing technology offer more performance thereby facilitating applications of unprecedented utility. Open systems, such as Linux, provide access to a huge software base. Nevertheless, these systems have to coexist with critical device infrastructure that insists on stringent timing and security properties. In this paper, we will present a capability-based software architecture, featuring enforceable security policies. The architecture aims to support current and future requirements of embedded computing systems, such as running versatile third-party applications on general purpose and open operating systems side by side with security sensitive programs.

Categories and Subject Descriptors

D.4.6 [Operating Systems]: Security and Protection

General Terms

Design, Security

Keywords

System Architecture, Capability System, Secure System, Small Trusted Computing Base, Mobile Devices

1. INTRODUCTION

In the past, typical mobile devices used closed software architectures as a way to make sure that no (potentially) malicious applications had to be dealt with. However, the flexibility demanded by users and evolving business arrangements prohibits such an approach. The vibrant marketplace for mobile devices is characterized by a short product cycle and highly dynamic security relations. Growing complexity and time-to-market pressure prohibit the in-house development of an entire software stack.

In recent years, a new generation of devices use general purpose operating systems as foundation. Even leading companies draw on commodity systems, such as Linux. They enhance the system by adding proprietary run-time environments, which is necessary because the standard security mechanisms are not deemed trustworthy enough. Those security concerns are alleviated by running applications in additional secured environments. Such additional layers of software compromise the performance. Better performing native applications have lower demands on hardware, which in turn could result in lower hardware costs or extended battery life.

An analysis of the security insufficiencies of contemporary commodity operating systems [6] reveals two principal reasons: First, not all components that are likely to malfunction are properly isolated. Specifically, device drivers reside in the kernel, which acts as single protection domain wherein more fine-grained isolation cannot be enforced. Second, isolation alone is not sufficient for dynamic systems. The notorious security flaws of established commodity operating systems are not caused by holes in the address spaces. Isolation must be complemented by an access-control mechanism, which supports the construction of systems that employ the principle of least authority [16].

We present an operating-system architecture that utilizes object capabilities, which have the reputation of being the preferred model to implement the principle of least authority.

Object capabilities are a method of access control that stands out by having several desirable properties. First, in a capability system security managers can enforce security policies by moderating communication relationships. If a component is entrusted with confidential data, leakage of this data to the outside world can be prevented by not granting that component any capabilities that allow communication to the outside world. Apart from containment, capabilities facilitate the selective propagation of individual rights. Second, a local naming scheme allows multiple instantiation of whole subsystems and provides give fine-grained access control, because access to objects must be explicitly granted instead of being available per default. This policy also prevents a common problem in systems with global names, servers may be the target of denial-of-service attacks as the system does not prevent communication of illegitimate senders. Besides the denial-of-service problem the server is forced to implement communication policies by itself.

Brakensiek et al. [5] describe a security architecture for

This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in

IIES '09 Nuremberg, Germany

Copyright 2009 ACM 978-1-60558-464-5 ...\$5.00.

mobile devices, which is based running different OS personalities on top of a small virtualization layer. The security properties of these systems benefit from a capability-based virtualization layer where interaction of the different OS personalities could be enforced without trusting the OS personalities itself. Additionally, systems with multiple virtual environments benefit from the local naming scheme provided by capability systems, as global names reveal the existence of other applications on the same system.

This work is an evolution of our previous work on L4-based systems. We will present how we extended an existing system with a capability-based access control mechanism.

2. L4 CLASSICS

The starting point of our work is based on the L4 operating system developed in our group: TUD:OS. It comprises the L4/Fiasco [1] microkernel, the user-level infrastructure called L4Env [2], and *L⁴Linux* [3]. *L⁴Linux* is a paravirtualized version of the popular Linux operating system kernel together with a broad range of software components on top.

We have been using this software stack for our own research as well as for teaching and it has been deployed in several projects together with external partners. The stability of this system reached a level that enabled sophisticated real-life scenarios, as described in [18, 19, 5].

Yet it lacked essential features whose relevance increased over time, namely kernel-memory management and communication control. Kernel-memory management controls the amount of kernel memory that can be occupied for a particular user activity. Communication control is a mechanism to define whether tasks are allowed to communicate with each other, or not. We will shortly address the first point whereas the remainder of this work will focus on communication control.

2.1 Brief Introduction to L4

Before we describe our new approach let us summarize the traditional L4 in more detail. L4 is a microkernel family that has been originally started by Jochen Liedtke [12]. It follows the rules of minimality and least authority. The system consists of microkernel and a user-level software layer that implements services in different subsystems, isolated by address spaces. Functionality is only integrated in the kernel if it cannot be implemented securely at user level or doing so would significantly decrease performance. Protection-domain construction is a kernel service on the grounds of security, whereas scheduling is included because no user-level scheduler has yet met reasonable performance goals.

The basic building blocks of L4 are tasks and threads. A task is a protection domain and a container for resources. Threads are the unit of execution, which are scheduled by the kernel. Interprocess communication (IPC) is a crucial part of the system as it allows to exchange messages between threads, however, it is also used for fault reflection, interrupt processing, and rights delegation.

Regarding kernel-memory management *L4/Fiasco* implements a quota mechanism to prevent denial of service attacks on the kernel. Each user-initiated kernel operation that requires kernel-memory is accounted to the quota of this user activity. Memory quotas can be shared among tasks as well as split, creating new quotas. This mechanism just gives control over the available kernel memory and is rather

pragmatic and simple, compared to much more sophisticated mechanisms, described in [8], [11].

Work on IPC control started with an approach called *ipcmon* which implements an in-kernel cache of communication relationships. For each new communication relationship, the kernel invokes a security monitor in user-level that decides whether this relationship is policy complying or not. The relationships are cached and succeeding IPCs are allowed without interaction of the monitor. *Ipcmon* provides a solution for fine-grained communication control but still uses a global naming structure that made parts of the system structure visible to all components.

The solution for the problem of a global name space is to give each component a local name space so that objects become only visible if needed. In the remainder of this paper we concentrate on the design of our capability-based kernel and user-level software, featuring communication control.

3. KERNEL DESIGN

The principle of least authority implies the possibility of fine-grained access control to objects. First, we will look at possible objects in the system. Afterwards, we describe how user activities can access these objects.

3.1 Objects

An object encapsulates functionality that is accessible through a well-defined interface only. In a client-server scenario, a server implements an object whereas the client uses the associated interface to access the object. In our system objects are implemented either by the kernel or by user-level components. To the user of an object there is no difference where the object is implemented or of which type an object is. In principle there are three different places where an object might be implemented: in the same task, in the kernel, or in another user-level component. Being implemented in the same task allows to the use of normal function call semantics and will not be pursued further here. The other two methods will be examined in the next.

3.1.1 Kernel Objects

The objects implemented in the kernel are offering the kernel functionality to the user. The following main objects are implemented by the kernel:

Task A task is a protection domain and provides the means for isolating different components. It consist of an virtual-memory address space, represented by the task's page table, as well as the capability table—holding local references to objects.

Thread A thread is the unit of execution. Multiple threads can execute within a single task—sharing resources, such as virtual memory and capabilities.

IRQ IRQs are senders of asynchronous notification events. Either a user or a hardware interrupt can flag an asynchronous event and continue running, while the receiver can use blocking or non-blocking mechanisms to receive the event. Receiving IRQ messages is based on IPC.

IPC-Gate The IPC-Gate is a special kernel object as it is the only one that does not implement any particular interface. It is rather used to establish communication channels for implementing user-level objects.

Factory A factory is an object that can create new kernel objects such as tasks, threads, IPC-Gates, or new factories. Factories implement kernel memory quotas.

3.1.2 User-Level Objects

User-level objects implement any kind of object a user-level service might provide. The implementor defines an interface the client can use to interact with the object. The interesting case here is how a communication channel between two tasks is established given that both tasks have their own local name space. In this case the IPC-Gate comes into play.

The IPC-Gate designates a receiver. The implementor of an objects creates an IPC-Gate and hands it over to clients so that these can communicate with the object. Each IPC-Gate has an associated thread that receives messages sent through the same.

IPC-Gates carry a secure, kernel protected label that is set by the server during IPC-Gate creation. On the receiver side this label can be used to attribute the message to the IPC-Gate it was sent through. Using this mechanism a receiver is able to distinguish between different senders, if each sender is supplied with a uniquely labeled IPC-Gate. Another use case is the identification of different objects on the receiver side by using one IPC-Gate per object that is carrying the object ID in the label. Receiving a message gives the receiver a single automatic reply capability to the sender so that replies can be sent.

3.1.3 User-Object References

Up to now we did not mention how tasks actually reference objects. At the lowest level all accessible objects map to kernel objects, might they be threads or IPC-Gates—upon which user-level objects are built. The kernel maintains a per-task capability table that holds the object references. Such a reference makes an object available to the particular task. Figure 1 depicts a schematic view on how capabilities and the capability table relate to objects in the kernel. The user refers to objects by their index in the table. An entry in the table can point to a kernel object (7) or can be invalid (8). The same table index in different tables might point to different kernel objects (2).

3.2 Object Invocation

Calling an object through a capability is called capability invocation. Compared to previous L4 implementations, where several system calls exist [14, 7, 13], the only remaining system call is the capability invocation.

The invocation of a capability transfers payload to the receiving object. The content of this payload is not generally defined as it is specific to the object being invoked. For example, to set the program counter of a thread, the interface of the *thread* object defines a specific layout of the payload.

For user-level applications the invocation of a capability is modeled as a normal function call of a (C++) object. Arguments of calls to members are marshalled into the payload that can be supplied when invoking a capability. A returning invocation can also supply data in the payload that will be

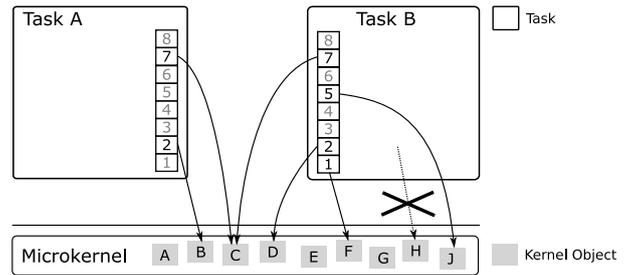


Figure 1: Schematic view of the relationship between capability tables of tasks and kernel objects.

consequently unmarshalled to the return arguments of the function call.

Figure 2 depicts the principle of an IPC-Gate used between a server and a client. Both the server and the client have a capability to the IPC-Gate, invocations on the gate are forwarded to the receiving thread.

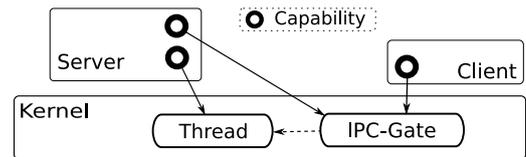


Figure 2: Cross task communication. An IPC-Gate is always connected to a thread. A thread invoking an IPC-Gate will send the message to the thread attached to the IPC-Gate.

3.2.1 IPC

IPC as known on L4 is still available, and indeed, invoking a capability is an IPC operation on the capability. The IPC sub-operations offer the same range of functionality as traditional L4 IPC. Examples are *call*, *reply_and_wait*, *send*, and *receive* operations.

3.3 Object Access

So far we have assumed that a task has capabilities to objects. However, we did not answer the question of how they appear in a task in the first place. Capabilities are granted to other tasks in the same way memory pages can be granted to tasks. In L4, a task can delegate access rights to resources, to which it has access, to other tasks by an operation called *mapping*. Mappings are established using IPC by offering rights that are transferred to the receiver if the receiver consents. Mapping a capability to another task gives both tasks access to the referenced object (*map*) or to the receiver only (*grant*). The method of mapping rights leads to a hierarchical mapping tree. Rights to objects can be revoked by any reference holder in the tree, leading to revocation in the whole subtree.

3.4 Objects and Interfaces

An interface describes the operations that can be invoked on an object. As references to objects are represented to the user as capabilities, an interface specifically describes the payload when invoking an object. An object is not strictly bound to a specific or a single interface, it might implement multiple interfaces. Furthermore, a client using an interface does not know, nor care, which object is implementing the interface, nor where the object is located. Objects may be implemented directly in the kernel or may be located in a user-level component. For user-level implementations the IPC-Gate is important as it is the kernel object that proxies invocations between user-level threads.

Taking this a step further allows us to virtualize objects in a client-transparent fashion. An example is a thread monitor that monitors thread creation of an application in order to gather runtime information of the monitored application. To do so, it can be configured to interpose in the factory of the application, which the application utilizes to create its threads. The monitor will then receive invocations of this interface and can thus track created threads, it will create the thread on behalf of the application and pass it on to the application using a map operation in the reply. The application will not notice that the thread creation was interposed. This mechanism allows a user-level component to implement a more sophisticated policy than available in the kernel. The scenario is depicted in Figure 3.

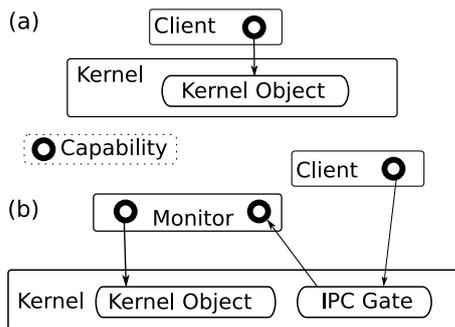


Figure 3: Part (a) shows calling an object without a proxy. The capability in the client points directly to the kernel object. In part (b) the call to the kernel object is interposed by the proxy. There is no difference for the client that is still invoking the capabilities (an IPC-Gate) with the interface of the kernel object.

This mechanism together with local naming also allows to virtualization of whole subsystems with the possibility to apply dynamic resource policies or just aid debugging. Given a proxy application that mimics the kernel's startup protocol, it is possible to launch an operating-system personality not directly on the kernel but on this proxy. Consequently, OS personalities can be nested, allowing to host several different OS personalities, the same personalities with different characteristics, or a mix of both.

3.5 Summary

Let us summarize the achievements of our changes to the kernel interface. Local names improve the access control of objects in the system significantly. Subsystems can be

nested and instantiated multiple times, as the layout of a local name space is not affected by the environment of a subsystem.

Kernel services are implemented as objects, any invocation on an object is itself a message passing primitive allowing interception by intermediate components. Instead of evolving the kernel, a user-level manager can offer its service transparently by intercepting kernel-object invocations and providing customized service.

4. USER-LEVEL SERVICES

Using the mechanisms described in the previous chapter we have built a user-level infrastructure that uses the capability-based kernel mechanisms. This user-level infrastructure is the software layer that provides basic abstractions for programs to execute. The runtime environment implements essential services, such as a memory manager, a program loader, and a naming service, as well as common application libraries.

The design of the runtime environment is based on a scheme where an IPC-Gate is created for each user-level object (see Section 3.1.2). The secure label of the IPC-Gate contains a pointer to the user-level data structures. In the following, when we refer to user-level objects there is always a corresponding IPC-Gate. If a task has access to an IPC-Gate it has access to the represented user-level object.

The basic interfaces provided are:

Dataspaces Dataspaces are an abstraction of memory objects. They can be used to represent normal memory (RAM), files within a file system or memory mapped I/O regions of devices. Shared memory is implemented by sharing dataspaces among multiple tasks.

Memory Allocator A memory allocator makes anonymous memory available to clients using dataspaces.

Region Management The region manager manages the address space of a task. It offers functions like reserving ranges of virtual memory, making dataspaces visible and resolving page faults. The functionality provided by the region management is similar to the POSIX `mmap` and `munmap` functions.

Name space A name space implements a mapping from textual names to objects and thereby makes services available to tasks. A name space is an important component as it implements the security policy of the system.

These interfaces are the building blocks for applications executing on top of the runtime environment.

Security wise, the most important interface is the name space. The name space implements a mapping from logical names to objects, in practical terms a name-space object delegates access rights to clients. A security policy governs which users are eligible for which services. Servers announce their services through their name spaces.

4.1 System/Program Startup

Upon system startup, the kernel starts one task that it equips with a base set of capabilities, such as the root factory. This program is called `roottask` and is the root of the program hierarchy (similar to the `init` process in

UNIX systems). It implements all of the basic runtime-environment objects and starts further programs based on configuration scripts. A configuration script has two main parts: the definition of name-space objects, which we shall describe in more detail afterwards, and a list of programs that shall be started including a description of the initial set of capabilities for each program. As the initial set of capabilities completely describes the environment of a running application, it is the place where all access-control and resource-management policies are anchored.

The roottask creates a set of kernel and user-level objects for each program to execute. Mandatory objects are a task, a thread, as well as a region-manager object. The task acts as protection domain for the application. The thread executes the code of the application binary. The region-manager object is responsible for resolving page faults and for providing memory mappings for the new program. As memory is handled in terms of dataspace we also need a set of dataspace objects for the program binary. Read-only parts of the binaries may be shared among multiple instances of a program whereas writable parts, such as data and BSS sections, require a private copy.

To enforce a policy on the kernel-resource usage, a program may be equipped with a particular factory capability to put constraints on kernel memory usage. Similar to that, each application may get a capability to an individual memory-allocator object that controls the policy for dynamic allocation of user-level memory (e.g., for `malloc`). For controlled access to other services in the system each application needs a capability to a name-space object.

4.1.1 Name-Space Objects

As mentioned, at the beginning of Section 4, name-space objects play a central role in defining the access-control policies for running applications. The configuration language allows to create new name-space objects and to add links to objects in other name spaces to them. Objects implemented by user-level services are announced to other applications by linking them to their name space.

4.2 Interception

An example for providing extended features using interception of kernel objects is a *Load Balancer* that implements the interface of the kernel factory object in order to monitor the creation of threads. The Load Balancer periodically gathers statistics information of the threads and of the CPUs on which they run allowing it to implement dynamic policies for distributing threads to available CPUs. The configuration of program loading allows to specify a user-implemented factory that replaces the default one.

4.3 Applications

We ported a set of applications from L4Env and adapted them to the capability-based access control. These applications are for example our GUI servers: L4Con, a full-screen console server for multiple clients; and the DOpE windowing system [10], a multi-client windowing system. Both servers multiplex access to a frame buffer and offer virtual frame buffers to their clients, allowing to transparently compose them in the program loading configuration.

4.3.1 Device Management

We implemented a device manager that is responsible

for providing selective access to the peripheral devices of the hardware platform. It can be configured to create virtual-bus objects with a subset of the available hardware peripherals. The name-space mechanism is used to grant access to device resources to different applications that provide drivers for peripherals.

4.3.2 Virtualization

L⁴Linux supports unmodified legacy Linux applications on top of the capability system. The *L⁴Linux* kernel is an adapted version of the Linux kernel and runs as ordinary application, using objects, such as dataspace, name spaces, and the virtual frame-buffer interface implemented by our GUI servers. It is possible to use native device drivers in *L⁴Linux* given that access permissions are granted by the device manager.

5. EVALUATION

To evaluate the performance impact of a capability lookup we conducted a low-level benchmark comparing the non-capability version of *L4/Fiasco* with the new capability-based version. The benchmark was run using an AMD PhenomTM 9550 Quad-Core CPU on an ASUS M3A78-EM motherboard equipped with 2GB DDR2-800 RAM.

The experiments show that a simple `l4_myself()` system call on the non-capability *L4/Fiasco* takes approximately 95ns and an invocation of a kernel object with an invalid protocol ID takes approximately 106ns. The results show that both versions perform equally well. Using a faster kernel entry method¹ even decreases the invocation time to 72ns, which is not possible on the non-capability system as this implements system calls with different kernel entries and the fast method is exclusively used for IPC. The difference can be explained by the fact that the capability invocation takes an additional branch to decide on the operation on the object. Overall this is the expected behaviour as the capability lookup is merely a memory reference that branches to the implementing object.

6. RELATED WORK

Using capabilities as access-control mechanism has been explored in various systems.

A pioneer in its field, Mach[4] could not draw on experiences of microkernel design and did not strike the best balance between functionality and performance. A profound issue was its mediocre IPC performance, the reason of which was a rather complex model that did not lend itself to well performing implementations. For example, reliable asynchronous message passing is convenient at higher levels but introduces buffer management into the kernel. Later microkernels designed from scratch, such as L4, showed that a dedication to minimal functionality could reconcile performance and generality.

The seL4 microkernel [9] from UNSW is similar to our approach. It has also been developed based on the experience with previous L4 systems. seL4 provides capability-based communication control and sophisticated means of kernel-memory management. One main difference is our design goal of a virtualizable kernel interface, which supports transparent interception by user-level services.

¹IA32 SYSENTER/SYSEXIT instructions

EROS[17] is an operating system based on persistence and capabilities. It uses capabilities to arbitrate access to all resources in the system, including communication rights and memory pages. EROS does not consider features, such as real-time that have been addressed by L4.

Flicker[15] is an approach for minimizing of the size of trusted computing base by using hardware-provided mechanisms to execute secure code. It uses newly available security extensions in recent IA32 CPUs to securely switch between the host operating system and Flicker code. This approach allows to add a single secure and hardware protected program to existing commodity systems.

7. CONCLUSION

The demands and complexity of software in the embedded and mobile devices market grows. The OS design for these systems moves towards OS-virtualization technologies besides having small secure and real-time applications on a single platform. This paper describes how our capability-based and object-oriented approach leads to better security properties with respect to isolation of multiple untrusted software components. We show how local naming, provided by capabilities, eases virtualization and component-based design for OS-paravirtualization as well as for small native applications.

The strictly object-oriented nature of the microkernel interface, as well as of the user-level components, provides the flexibility of transparent interception of all kernel primitives and enables the implementation of enriched versions of the kernel primitives at user level.

8. ACKNOWLEDGEMENTS

We would like to thank Michael Peter and the reviewers for their constructive criticism leading to this version of the paper. The authors also would like to thank the European Commission for their support through the European Research Programme FP7².

9. REFERENCES

- [1] Fiasco website. URL: <http://os.inf.tu-dresden.de/fiasco/>.
- [2] L4 Environment website. URL: <http://os.inf.tu-dresden.de/l4env/>.
- [3] L4Linux website. URL: <http://os.inf.tu-dresden.de/L4/LinuxOnL4/>.
- [4] M. J. Accetta, R. V. Baron, W. Bolosky, D. B. Golub, R. F. Rashid, A. Tevanian, and M. W. Young. Mach: A new kernel foundation for unix development. In *USENIX Summer Conference*, pages 93–113, Atlanta, GA, June 1986.
- [5] J. Brakensiek, A. Dröge, H. Härtig, A. Lackorzynski, and M. Botteck. Virtualization as an enabler for security in mobile devices. In *Proceedings of the First Workshop on Isolation and Integration in Embedded Systems (IIES 2008), EuroSys 2008 Affiliated Workshop*, pages 17–22, Glasgow, Scotland, UK, April 2008.
- [6] A. Chou, J. Yang, B. Chelf, S. Hallem, and D. Engler. An empirical study of operating systems errors. In *SOSP '01: Proceedings of the eighteenth ACM symposium on Operating systems principles*, pages 73–88, New York, NY, USA, 2001. ACM.
- [7] U. Dannowski, J. LeVasseur, E. Skoglund, and V. Uhlig. L4 experimental kernel reference manual, version x.2. Technical report, University of Karlsruhe, 2004. Latest version available from: <http://l4hq.org/docs/manuals/>.
- [8] Dhammika Elkaduwe, Kevin Elphinstone and Philip Derrin. Kernel design for isolation and assurance of physical memory. In *Proceedings of the First Workshop on Isolation and Integration in Embedded Systems (IIES 2008), EuroSys 2008 Affiliated Workshop*, Glasgow, Scotland, UK, April 2008.
- [9] D. Elkaduwe, P. Derrin, and K. Elphinstone. Kernel data – first class citizens of the system. In *Proceedings of the 2nd International Workshop on Object Systems and Software Architectures*, pages 39–43, Victor Harbor, South Australia, Australia, Jan 2006.
- [10] N. Feske and H. Härtig. Demonstration of DOpE — a Window Server for Real-Time and Embedded Systems. In *24th IEEE Real-Time Systems Symposium (RTSS)*, pages 74–77, Cancun, Mexico, Dec. 2003.
- [11] B. Kauer. L4.sec Implementation - Kernel Memory Managment. Master's thesis, TU Dresden, May 2005.
- [12] J. Liedtke. On μ -kernel construction. In *Proceedings of the 15th ACM Symposium on Operating System Principles (SOSP)*, pages 237–250, Copper Mountain Resort, CO, Dec. 1995.
- [13] J. Liedtke. L4 reference manual (486, Pentium, PPro). Arbeitspapiere der GMD No. 1021, GMD — German National Research Center for Information Technology, Sankt Augustin, Sept. 1996. Also Research Report RC 20549, IBM T. J. Watson Research Center, Yorktown Heights, NY, September 1996.
- [14] J. Liedtke. L4 nucleus version x reference manual (x86). Technical report, University of Karlsruhe, Sept. 1999.
- [15] J. M. McCune, B. J. Parno, A. Perrig, M. K. Reiter, and H. Isozaki. Flicker: an execution infrastructure for tcb minimization. In *Eurosys '08: Proceedings of the 3rd ACM SIGOPS/EuroSys European Conference on Computer Systems 2008*, pages 315–328, New York, NY, USA, 2008. ACM.
- [16] M. S. Miller, K.-P. Yee, and J. Shapiro. Capability Myths Demolished. Technical report, 2003.
- [17] J. S. Shapiro, J. M. Smith, and D. J. Farber. Eros: a fast capability system. In *In Symposium on Operating Systems Principles*, pages 170–185, 1999.
- [18] L. Singaravelu, C. Pu, H. Härtig, and C. Helmuth. Reducing TCB complexity for security-sensitive applications: three case studies. *SIGOPS Oper. Syst. Rev.*, 40(4):161–174, 2006.
- [19] C. Weinhold and H. Härtig. VPFS: building a virtual private file system with a small trusted computing base. *SIGOPS Oper. Syst. Rev.*, 42(4):81–93, 2008.

²<http://cordis.europa.eu/fp7/>