Memory Sharing Revisited

Jorrit N. Herder, Herbert Bos, Arun Thomas, Ben Gras, Andrew S. Tanenbaum
{jnherder,herbertb,arun,beng,ast}@cs.vu.nl
Vrije Universiteit, Amsterdam, The Netherlands

Shortcomings of Existing Protection Schemes

System V IPC and POSIX Shared Memory lack flexibility and offer limited memory protection. Some shortcomings include:

- Coarse-grained, page-based protection
- Protection based on UID not process
- Access rights cannot be delegated
- No seamless integration for safe DMA

Fine-grained,Delegatable Memory Grants

Process that wants to share memory creates a grant table, builds a memory grant, and sends index to other party: (proc ID, grant ID) identifies the grant.

Recipient of a memory grant must call the kernel in order to perform privileged grant operations. Kernel validates the access rights and performs the request.

Definition of API and implementation of operations are a work in progress:

- Grant creation and revocation
- Memory copying
- Memory mapping
- Direct memory access (DMA)

Take-home Messages

- Memory grants are a novel alternative to existing memory protection models.
- Precise per-proces access control for byte-granularity memory regions.
- Used in MINIX 3 to protect against memory corruption by buggy drivers.

Memory Grant Data Structure

Direct Memory Grant

Indirect Memory Grant

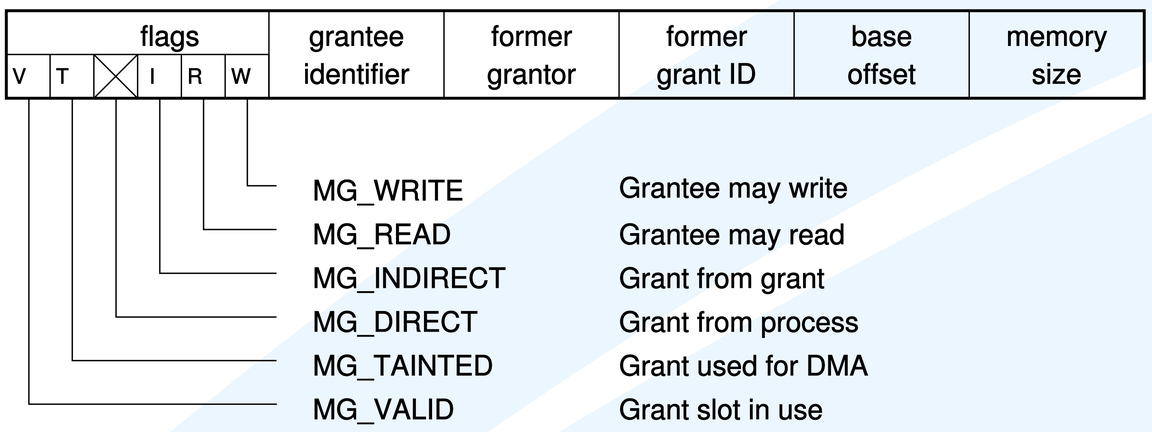

Grant Tables and Delegation

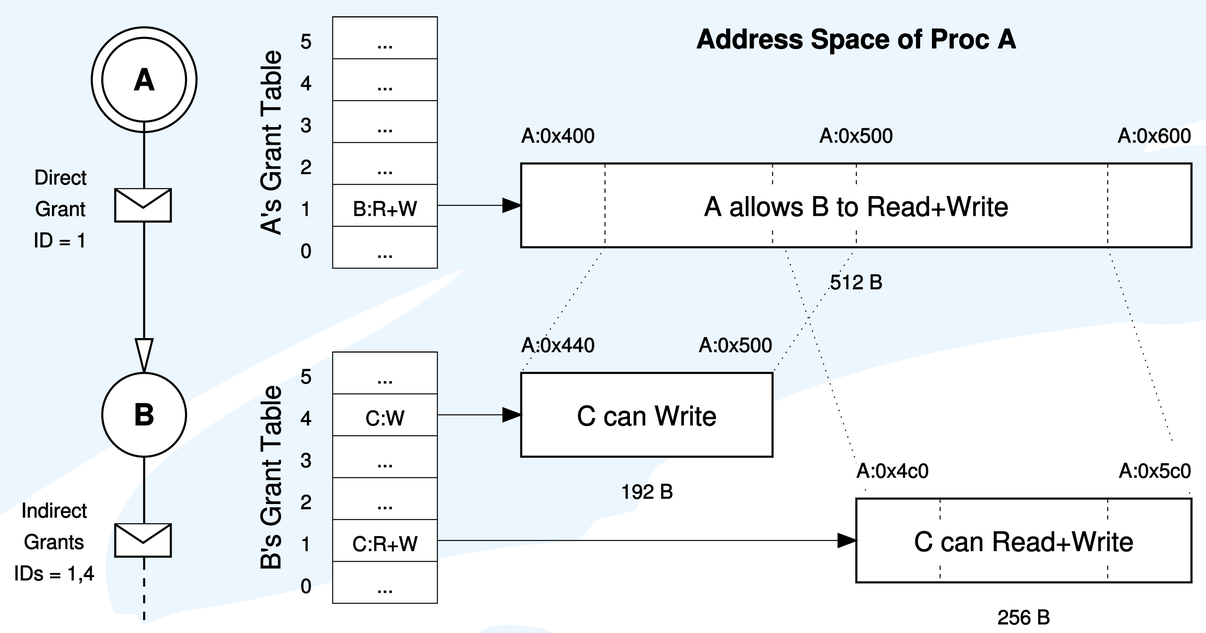

Grants for Above Scenario

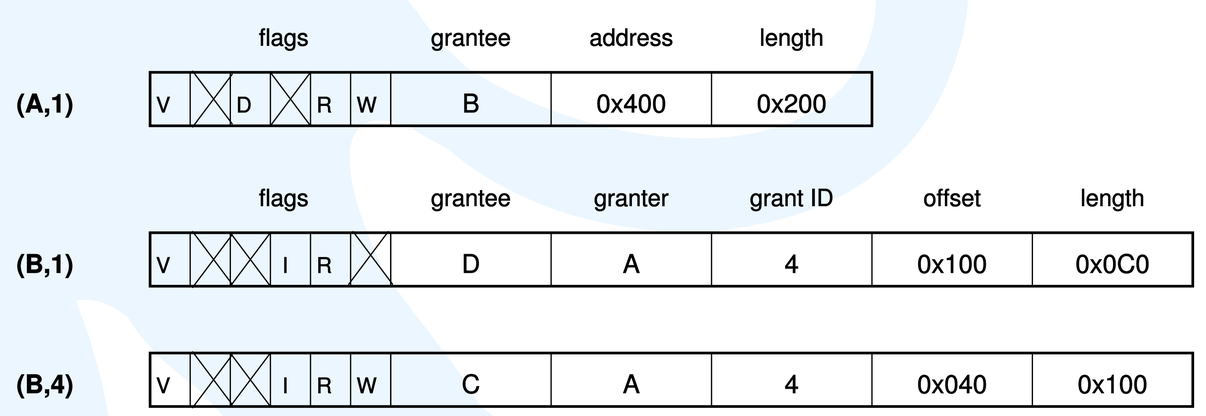

Supported by the Netherlands Organization for Scientific Research (NWO) under grant 612-060-420.