Vsys: A programmable sudo

Sapan Bhatia (Princeton), Giovanni Di Staasi (University of Naples), Thom Haddow (Imperial College), Steve Muir (Juniper), Andy Bavier (Princeton), Larry Peterson (Princeton)

What is Vsys

- Tool for restricting access to privileged operations
- Sometimes user demands not conveniently satisfied by default security model
- Vsys helps satisfy these demands safely
- Operations:
  - Simple: Open a raw socket, Access system logs
  - Complex: Create private overlay network, shape traffic

Vsys as a tool

- Typically runs in a chroot environment
- Outside chroot: privilege extensions implemented as executable files
- Inside chroot: FIFO pipes, UNIX domain sockets
- Users use these to communicate with extensions by reading from and writing to to these

Vsys as a tool

- Typically runs in a chroot environment
- Outside chroot: privilege extensions implemented as executable files
- Inside chroot: FIFO pipes, UNIX domain sockets
- Users use these to communicate with extensions by reading from and writing to to these
- Advanced:
  - Access control policies (ACPs)
  - Passing control structures (sockets, file descriptors)

Vsys architecture

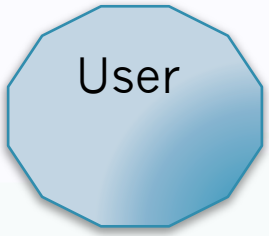

Vsys architecture

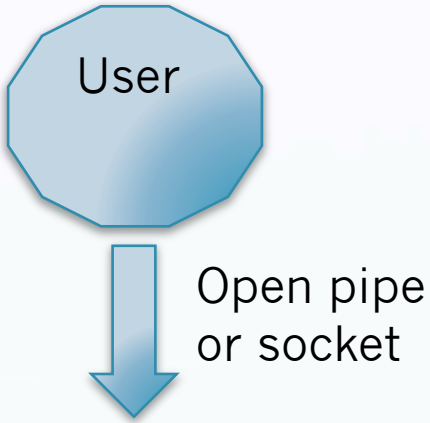

Vsys architecture

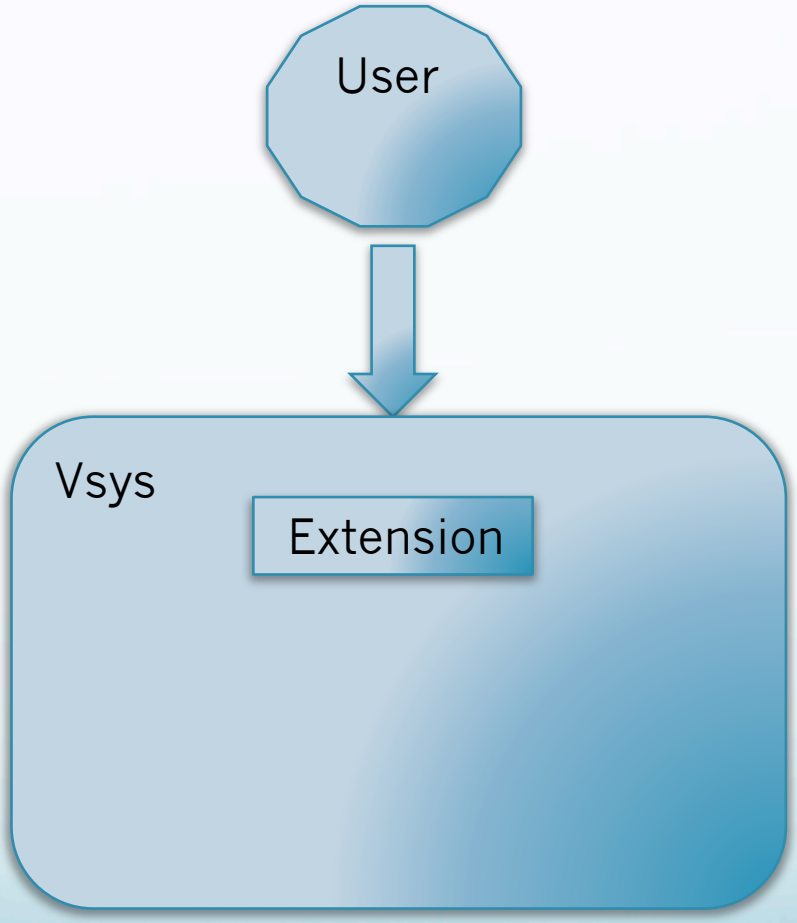

Vsys architecture

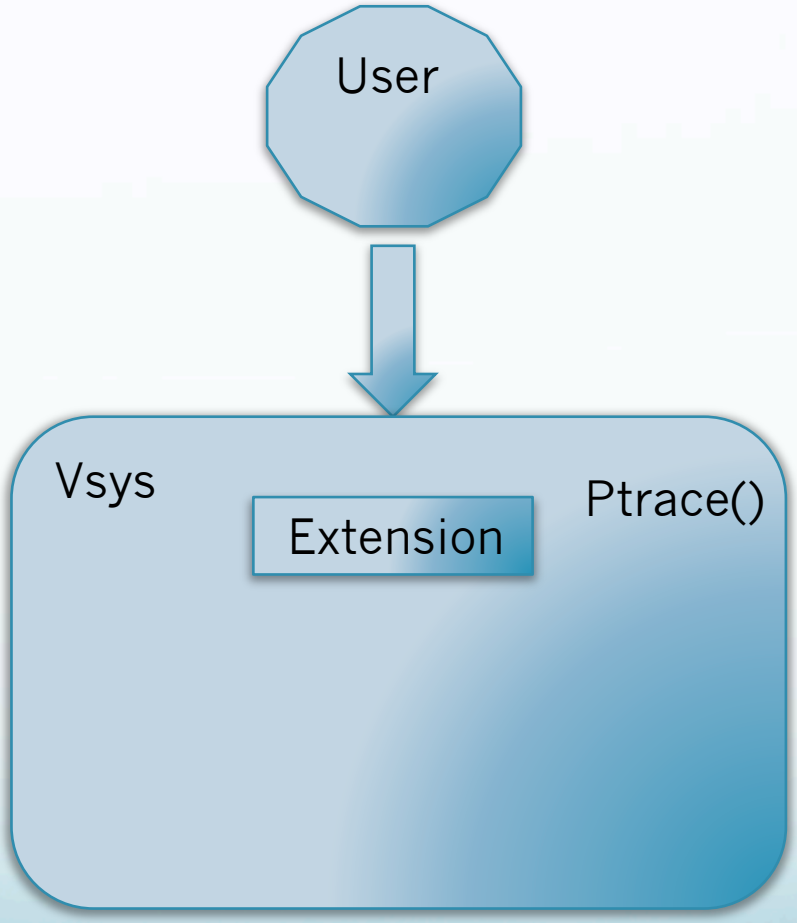

Vsys architecture

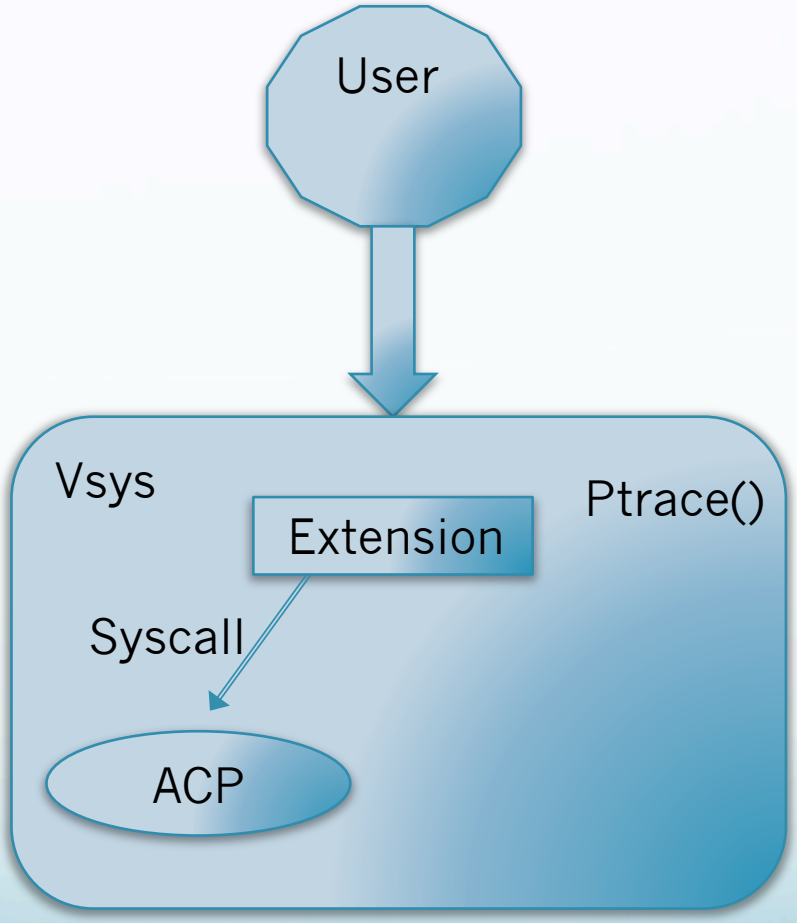

Vsys architecture

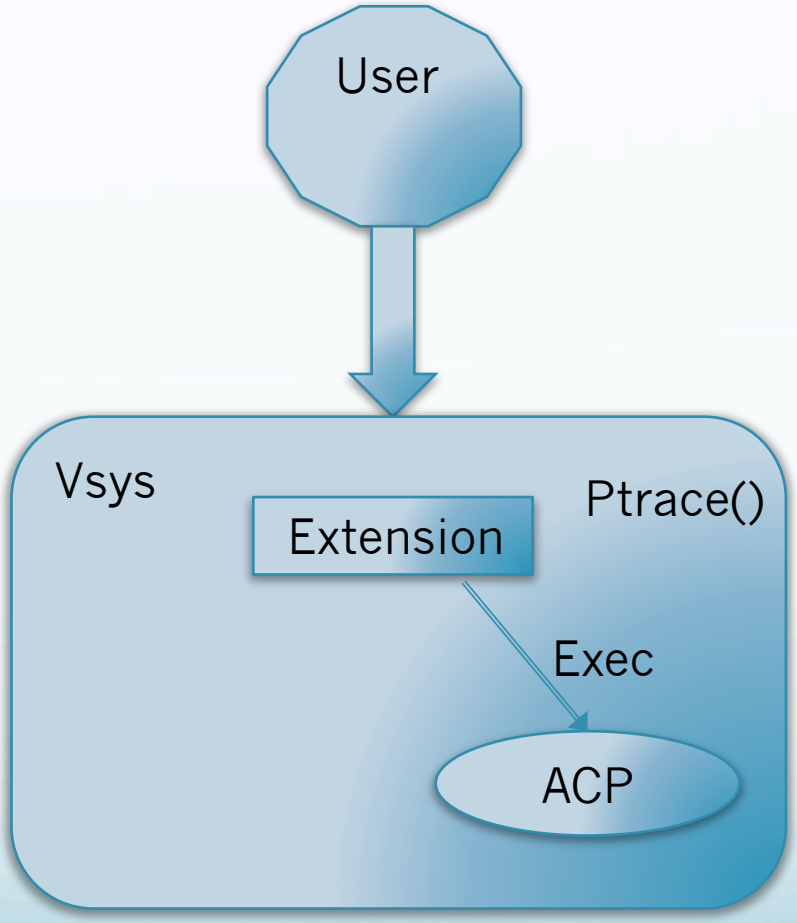

Vsys as an approach

- Don't extend the OS
- Reuse existing building blocks for isolation:
  - Processes
  - File descriptors
  - Network interfaces
  - Packet tags
- Combine them:
  - Network interface + Packet tag = Isolated Interface
  - Raw socket + Packet tag = Isolated raw socket

Example: Sliceip

- Version of “ip route” tool open to non-administrators
- Creates user-specific routes

User side	Vsys side
Request new virtual interface	Allocate private IP, Set interface name, configure firewall rules
Request new route	Modify route to apply to user’s packets and/or interface

Other Vsys extensions on PlanetLab

- Fusemount: Mount and use userspace filesystems created by other users
- Socketops:
  - Large TCP/UDP buffers
  - Raw sockets
  - Control sockets
    - (e.g. read packet headers from kernel)
  - QoS settings
  - Etc.
- Vtuntap: Create and manage virtual devices

Lesson: Creating new OS abstractions is hard

- Easy to prototype, hard to run over long periods
- E.g. Vsys networking Vs Linux namespaces
- Initially went with Linux namespaces project
- Disadvantages:
  - Bad interactions with other components (iptables, linux-vservers)
  - Bugs
  - Missing tool and library support

[BONUS SLIDE]

Conclusion

- Vsys is a flexible sudo
- Rapid deployment of new isolated functionalities
- Encourages grassroot abstractions
- Experience: supported surprisingly powerful extensions
- Highly successful as tool and approach
  - 50+ privilege extensions
  - Supported papers and PHD dissertions: NIPS, NSDI, OSDI, …
  - 10+ external developers