Towards High Assurance Networks of Virtual Machines

Fabrizio Baiardi[1] Daniele Sgandurra[2]

[1]Polo G. Marconi - La Spezia, University of Pisa, Italy

[2]Department of Computer Science, University of Pisa, Italy

EC2ND Conference, 2007

Outline

Rootkits

Rootkits have become more sophisticated over the years.

- User-level rootkits: usually, modify system binaries.
- Kernel-level rootkits: for example, a module inserted into the kernel.

Unfortunately, rootkits and IDSes work at the same level. A rootkit can attack or evade the IDS controls.

Virtualization Technology

- Software emulation of the hardware architecture: Virtual Machines (VMs).
- Benefits:
  1. Confinement among the VMs.
  2. Server consolidation.
  3. Centralized management.
- Accessing VM's state from a lower level.
  - Virtual Machine Introspection.

Virtual Machine Monitor

A VMM supports the concurrent execution of multiple OSes.

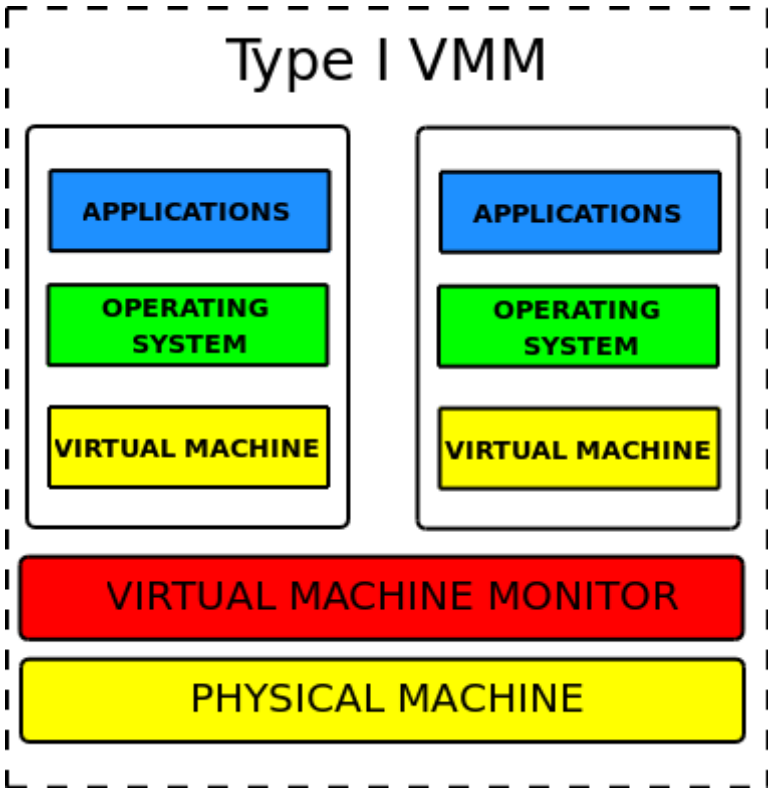

Proposed Approach

Virtual Machine Introspection: Standford University.

- Visibility: access VM's state from a lower level.
- Robustness: detect intrusions from another VM.

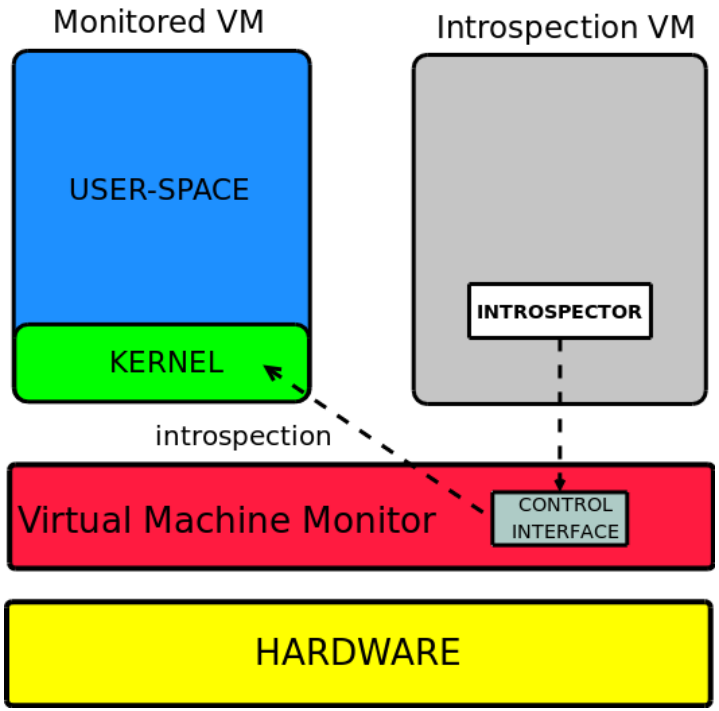

Multi-Level Approach to Intrusion Detection

Semantic gap: introspection can access only raw data.

- It requires an introspection library to rebuild kernel data structures from another VM and check their integrity.
- Introspection can be the first step of a chain of trust:
  1. Introspection protects kernel integrity.
  2. The kernel is extended with further security functions.
     - Customized kernel modules.
     - Security kernel patches: SELinux, LIDS.
  3. Standard IDSes run inside the VM.

Chain of Trust

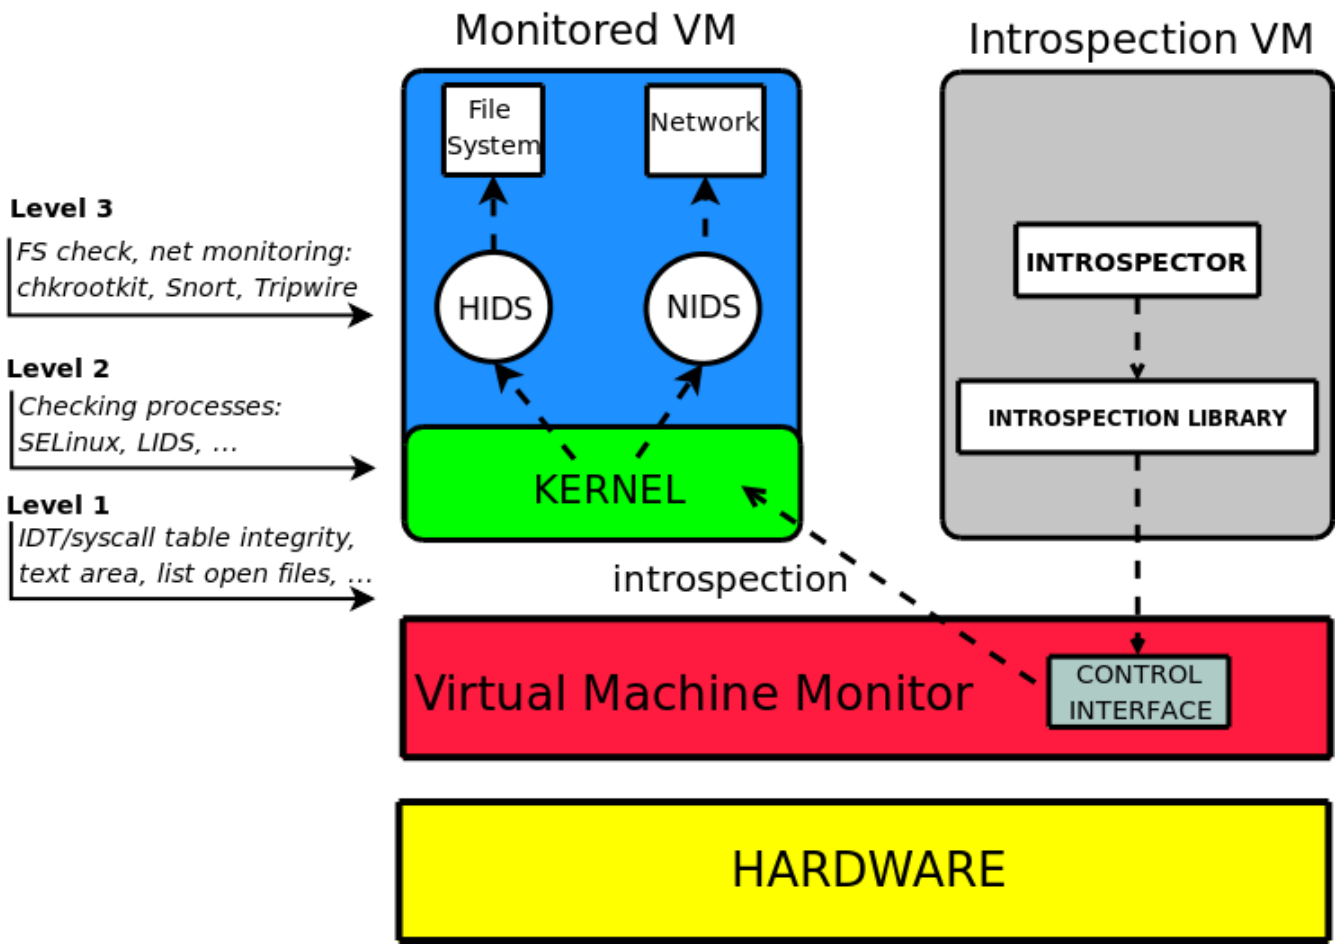

Generalization of the Approach

The assurance can be extended to a network of VMs.

- Each introspection VM monitors the local VMs.
- Introspection VMs under the same administrative domain establish a trust relationship:
  - Certificates.
  - Ex.: when accessing a service of a VM on another node.
- Introspection VMs exchange alerts and information on the local VMs.
  - Distributed attacks, worms.
  - Software attestation.

Extending the Assurance

Join process.

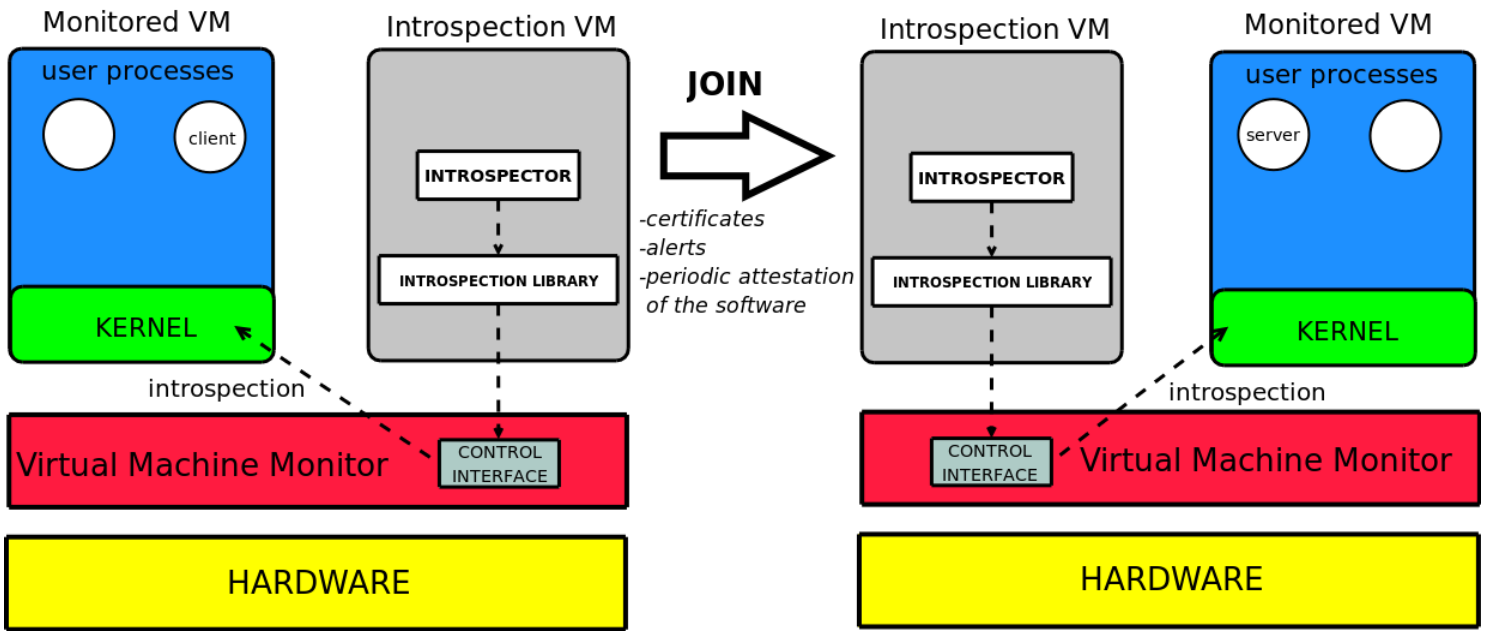

Psyco-Virt Architecture

- The first prototype is based on Xen hypervisor.
- Introspection VM (IVM): monitors all the VMs.
- Monitored VMs (Mon-VMs): execute the systems to be monitored.
- Control Network: supports the exchange of alerts and information among the IVMs.
- Data Network: supports the exchange of application traffic.

Introspection VM

Introspection VM: monitors all the VMs.

- The introspector protects kernel integrity.
- Join process: authentication/attestation.

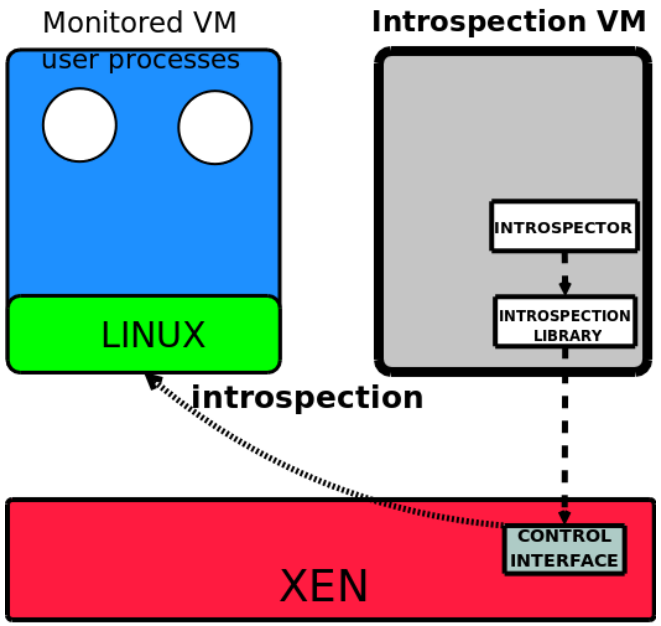

General Approach for Introspection

1. Pause the execution of the VM.
2. Map the memory pages containing the kernel data structure to be monitored.
3. Cast the raw memory into the correct data structure.
   - Linux kernel headers.
4. If there are pointers, map the page containing the referenced address.
   - Ex.: linked list.
5. Apply the consistency checks.
6. Resume the execution of the VM.

General Approach for Introspection

Map the memory pages containing the kernel data structure.

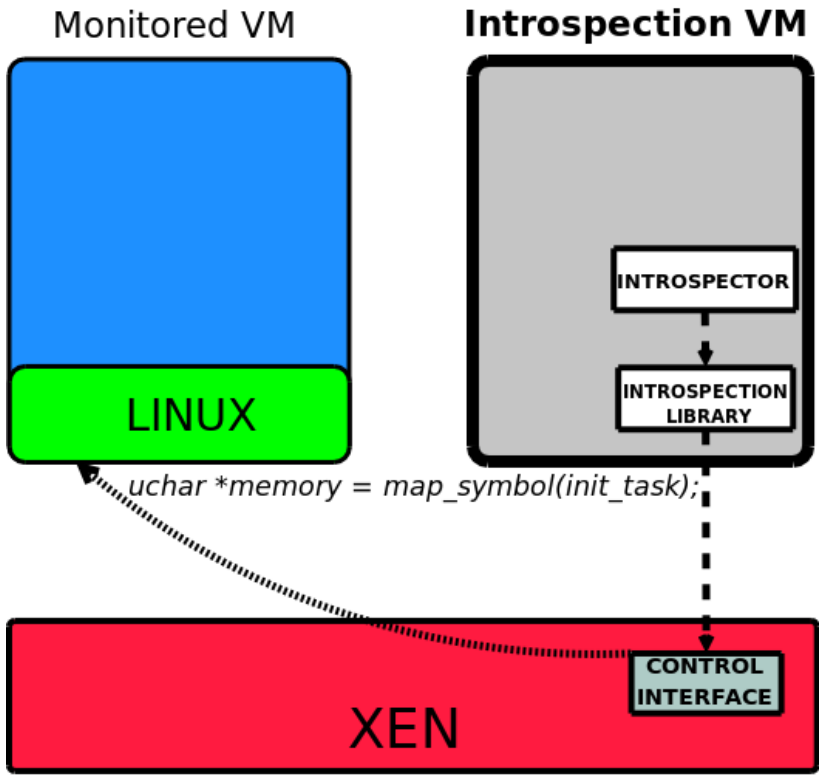

General Approach for Introspection

Map the memory pages containing the kernel data structure.

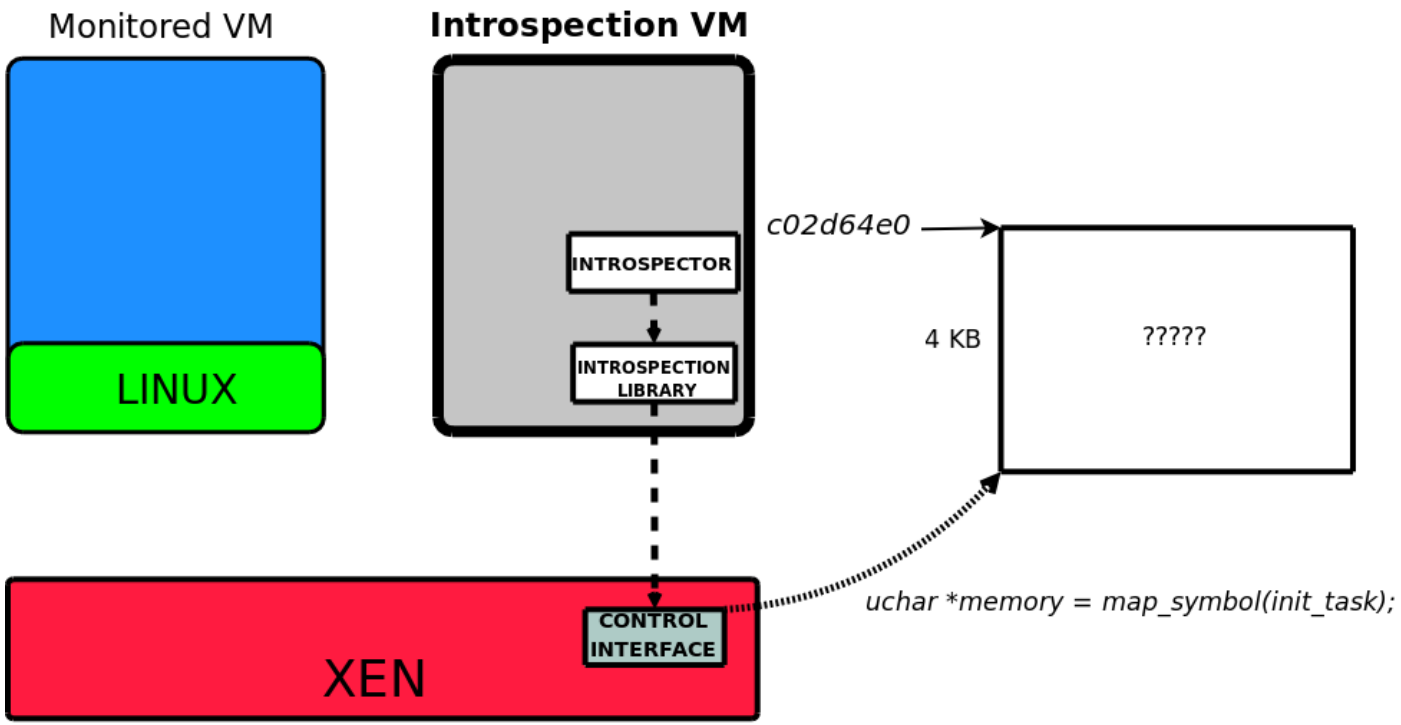

General Approach for Introspection

Cast the raw memory into the correct data structure.

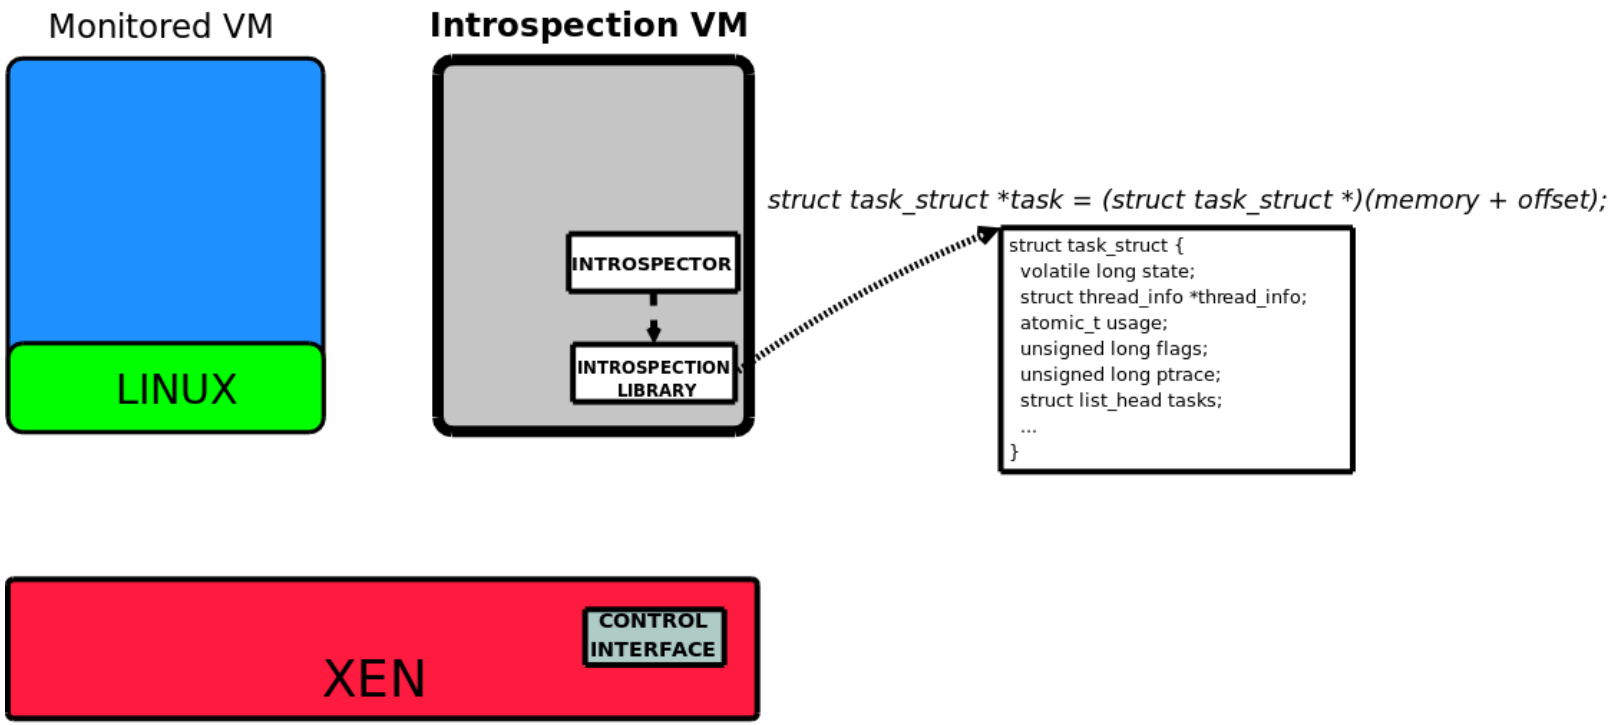

General Approach for Introspection

Map the pages containing referenced addresses.

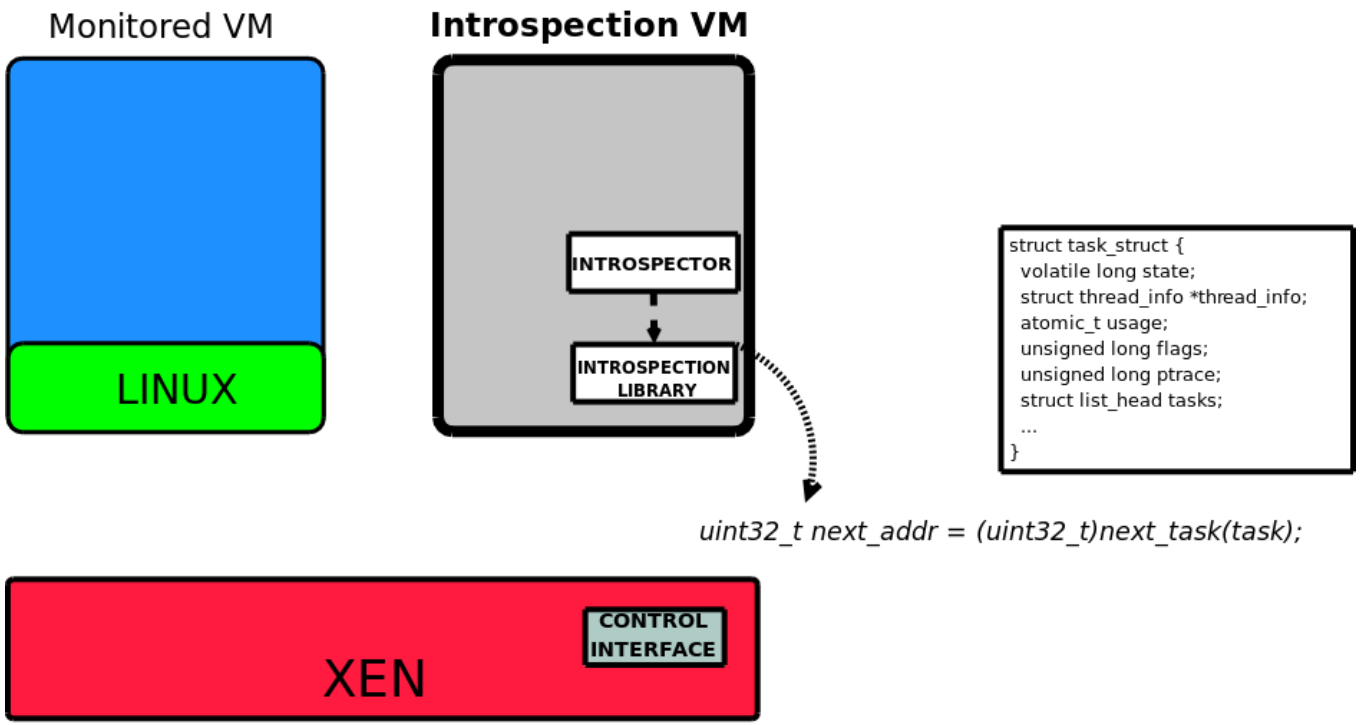

General Approach for Introspection

Map the pages containing referenced addresses.

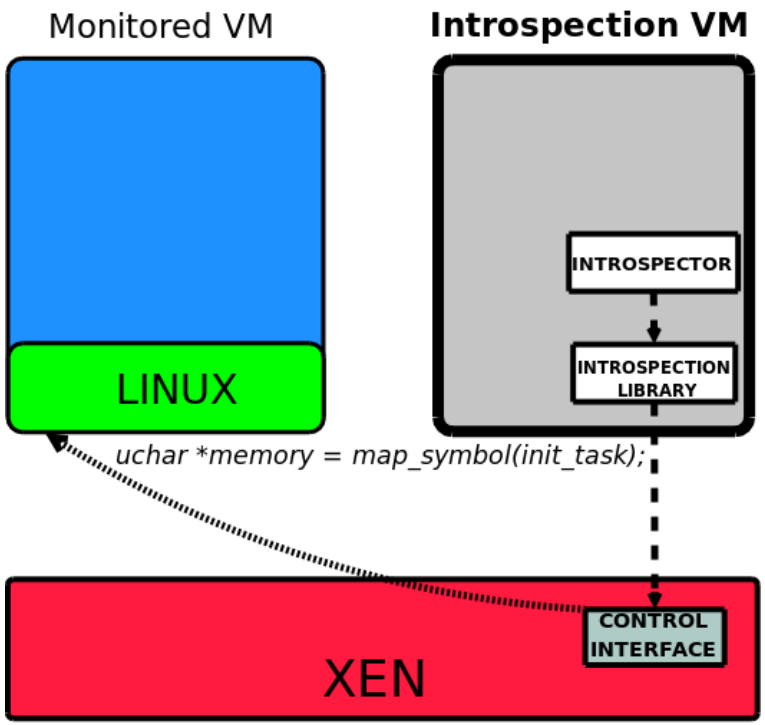


```
struct task_struct {
  volatile long state;
  struct thread_info *thread_info;
  atomic_t usage;
  unsigned long flags;
  unsigned long ptrace;
  struct list_head tasks;
  ...
}
```

General Approach for Introspection

Map the pages containing referenced addresses.

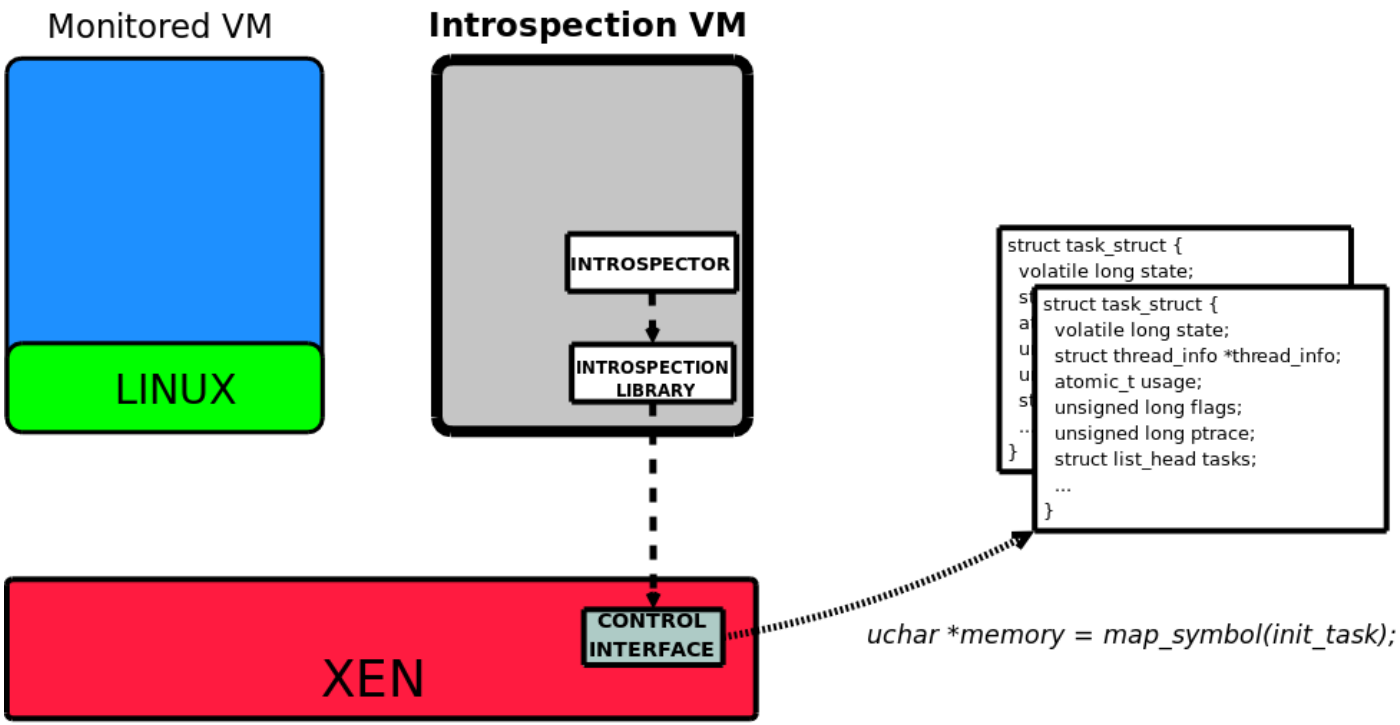

General Approach for Introspection

Apply the consistency checks.

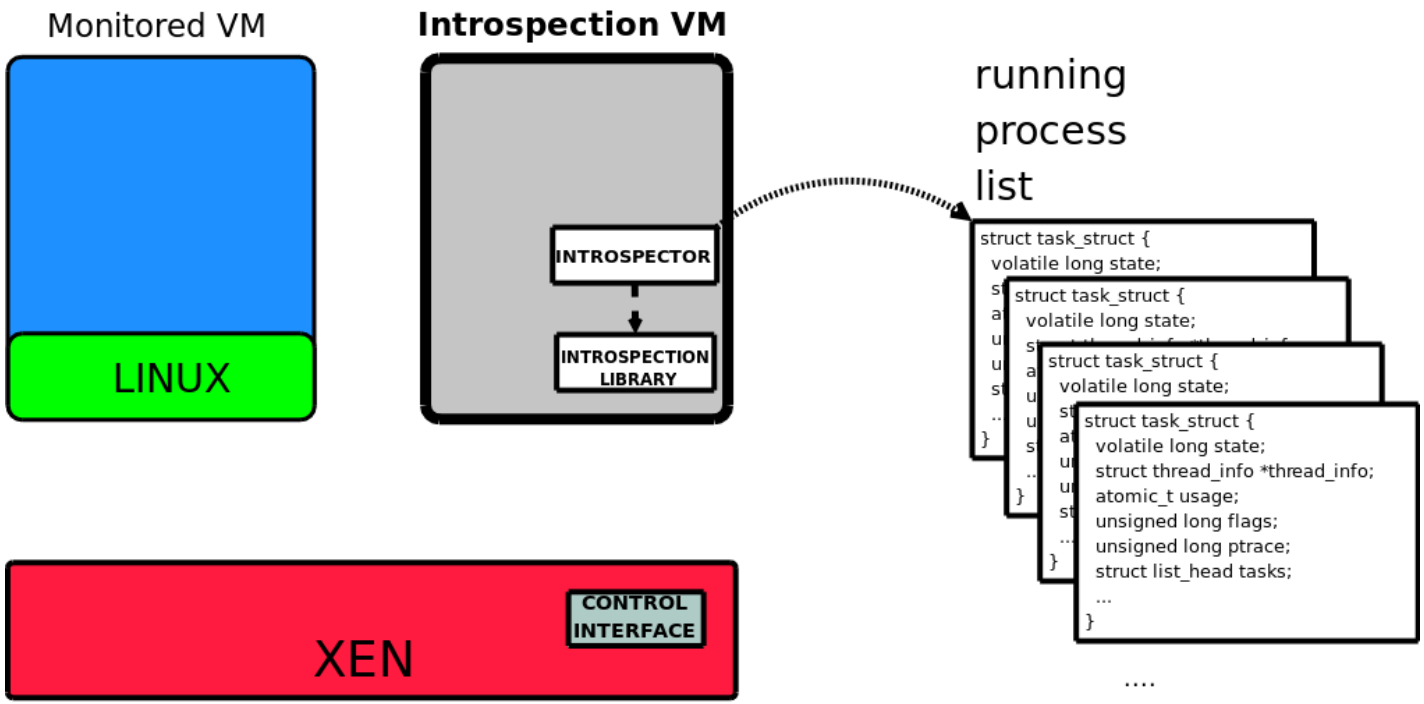

Monitored VM

Monitored VM: executes the system to be monitored.

- It may run IDSes to detect attacks/intrusions.
  - The collector receives all the alerts from the local IDSes.
  - The kernel checks IDSes integrity.

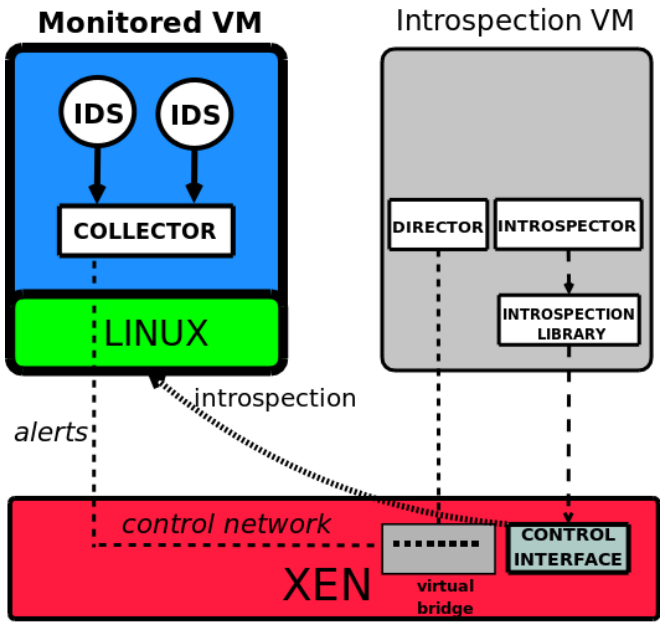

Control Network

Control Network: supports the exchange of alerts and commands among the VMs.

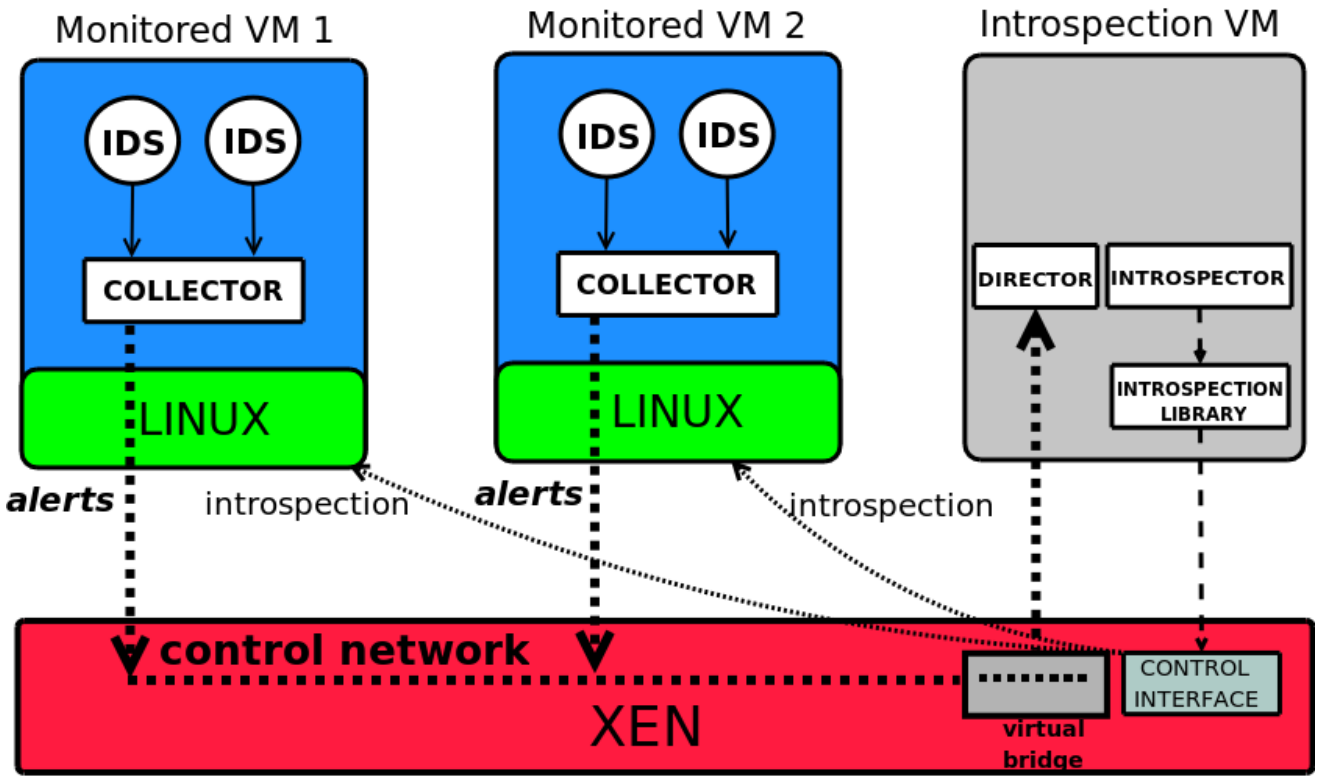

Architecture

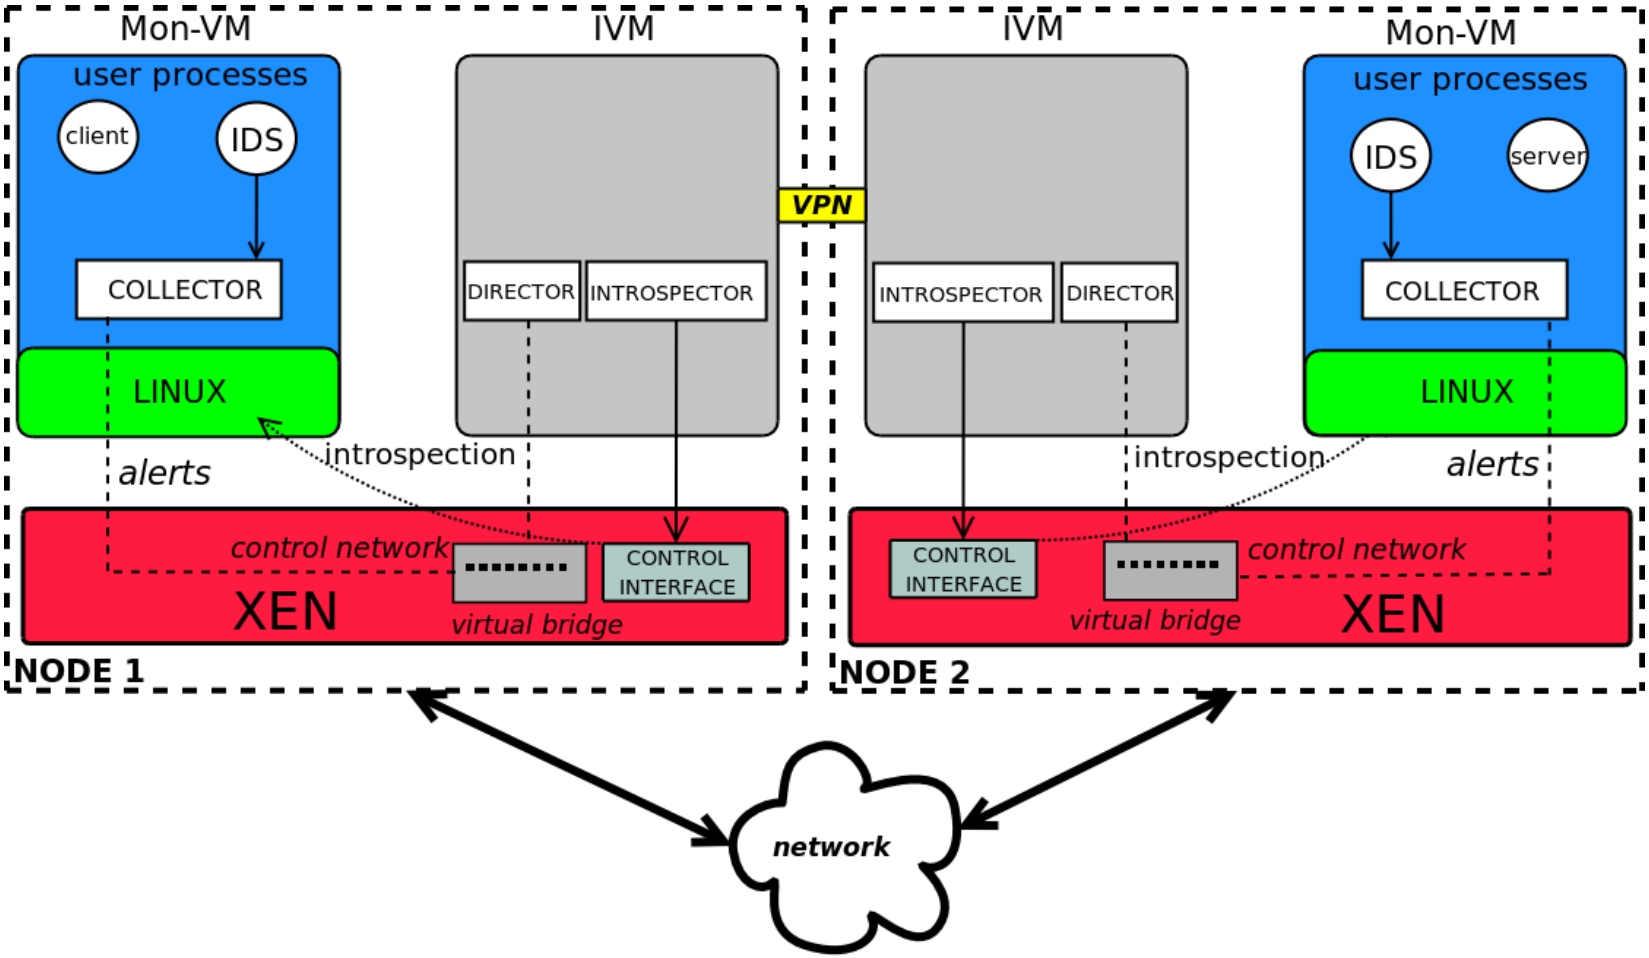

High Assurance Network of VMs

- Each IVM has a list of authorized partner IVMs.
  - PEM certificates.
- We assume VMMs and IVMs are trusted.
  - Small size, no Internet services.
- Each node is authenticated and protected.
- Remote software attestation.

Attacks Detected

Currently, Psyco-Virt detects:

- Attacks to the kernel code also those inserting a malicious module.
- Updates to the IDT and syscall table.
- Updates to the text area of a critical process.
- Replacing ps and lsof.
- Interfaces set into promiscuous mode.

IOzone Read Performance

We used the IOzone Benchmark Tool to run NFS performance tests on a client Mon-VM.

- Overhead is less than 10%.

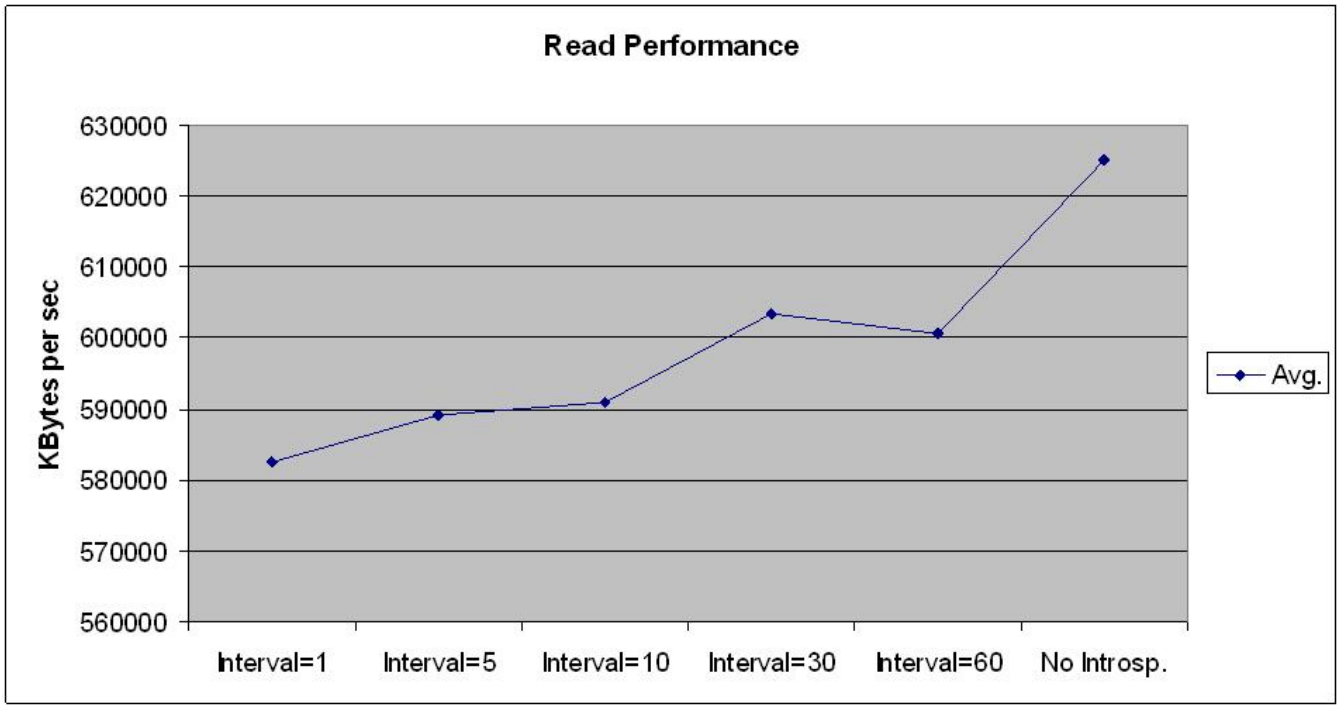

Limitations

Current limitations of the prototype:

- No checks on kernel dynamic data, such as stack.
- Other critical kernel data structures, besides IDT and syscall table, have to be protected.
- Attacking the kernel between each execution of the checks.
- The join process is static.
- Distributed attacks: work in progress.

Results

- Preventing evasion of the controls and attacks to IDSes.
- Multi-level approach to form a chain of trust:
  1. VMM.
  2. Kernel.
  3. IDSes.
- Extending the TCB over multiple physical nodes.
- Acceptable overhead.

Future Works

- A master VM manages and configures the whole network.
- Checking at runtime kernel invariants.
  - Using abstract interpretation of kernel code.
- Tracing a VM, such as with ptrace.
  - Verifying system call parameters.
- Enhanced cooperation among the nodes:
  - Dynamic join process.
  - Worm detection.