



IBM T. J. Watson Research Center

Xen Summit Sept. 7<sup>th</sup> 2006

# **Towards Enterprise-level Security with Xen -- sHype Access Control Module (ACM)**

**Reiner Sailer, Ronald Perez**

**Stefan Berger, Ramón Cáceres, Leendert van Doorn**

**IBM T. J. Watson Research Center, NY**

## Major Goal for a Secure Hypervisor

**sHype = Distributed Workload Protection across  
Workload Balancing and Virtual I/O**

### **Benefits matching Customer / User requirements**

- **Universal protection guarantees across machines**
- **Minimal or no performance overhead**

### **How-to achieve those requirements**

- **Simple, platform-independent protection statements (Policy)**
- **Efficient remote property-attestation**
- **Finer-grained controls layered on top (VM/OS level)**
- **Non-intrusive, easy to maintain in “Internet-speed”**

# Agenda

## **Status Quo**

- **sHype Access Control Module Status**
- **Security Management + vTPM Status**

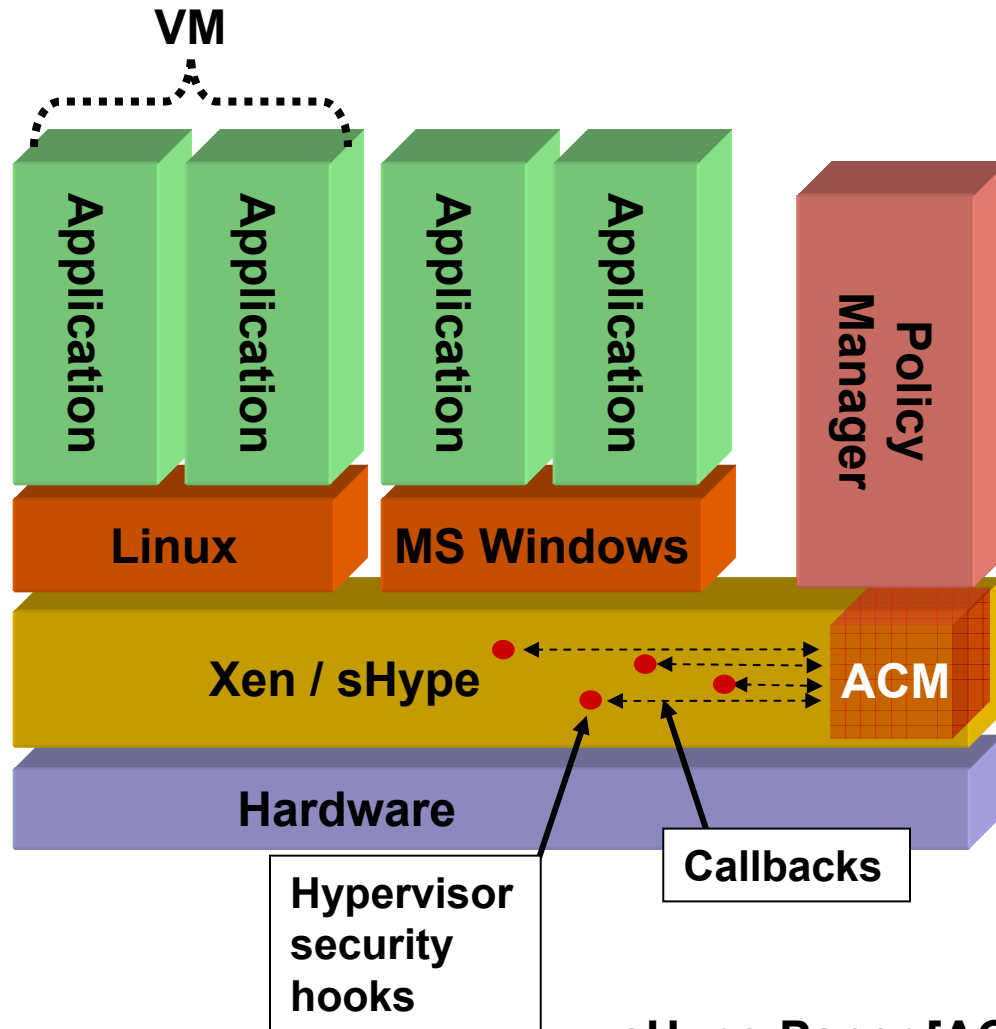
## **Outlook (Discussion)**

- **Long-term Security Architecture for Xen**

## **sHype/Xen Demonstration (Optional)**

- **Create and deploy sHype/Xen workload protection in 5 min**

# Secure Hypervisor Architecture (sHype)



- Flexible framework:  
Supports Multiple Policies
- Access Control Module  
Implements Policy Model
- Hypervisor Security Hooks
  - ✓ mediate inter-VM communication
  - ✓ interact with ACM for access decision
- Implemented for Xen, PHYP, rHype in various stages

sHype-Paper [ACSAC2005]

# sHype: From VM to Workload Protection

Value-added Services


- Quarantining VMs
- Layering OS + App. Security

Granularity


Users/  
Data

Human Resources

sHype  
(controls sharing)




Payroll



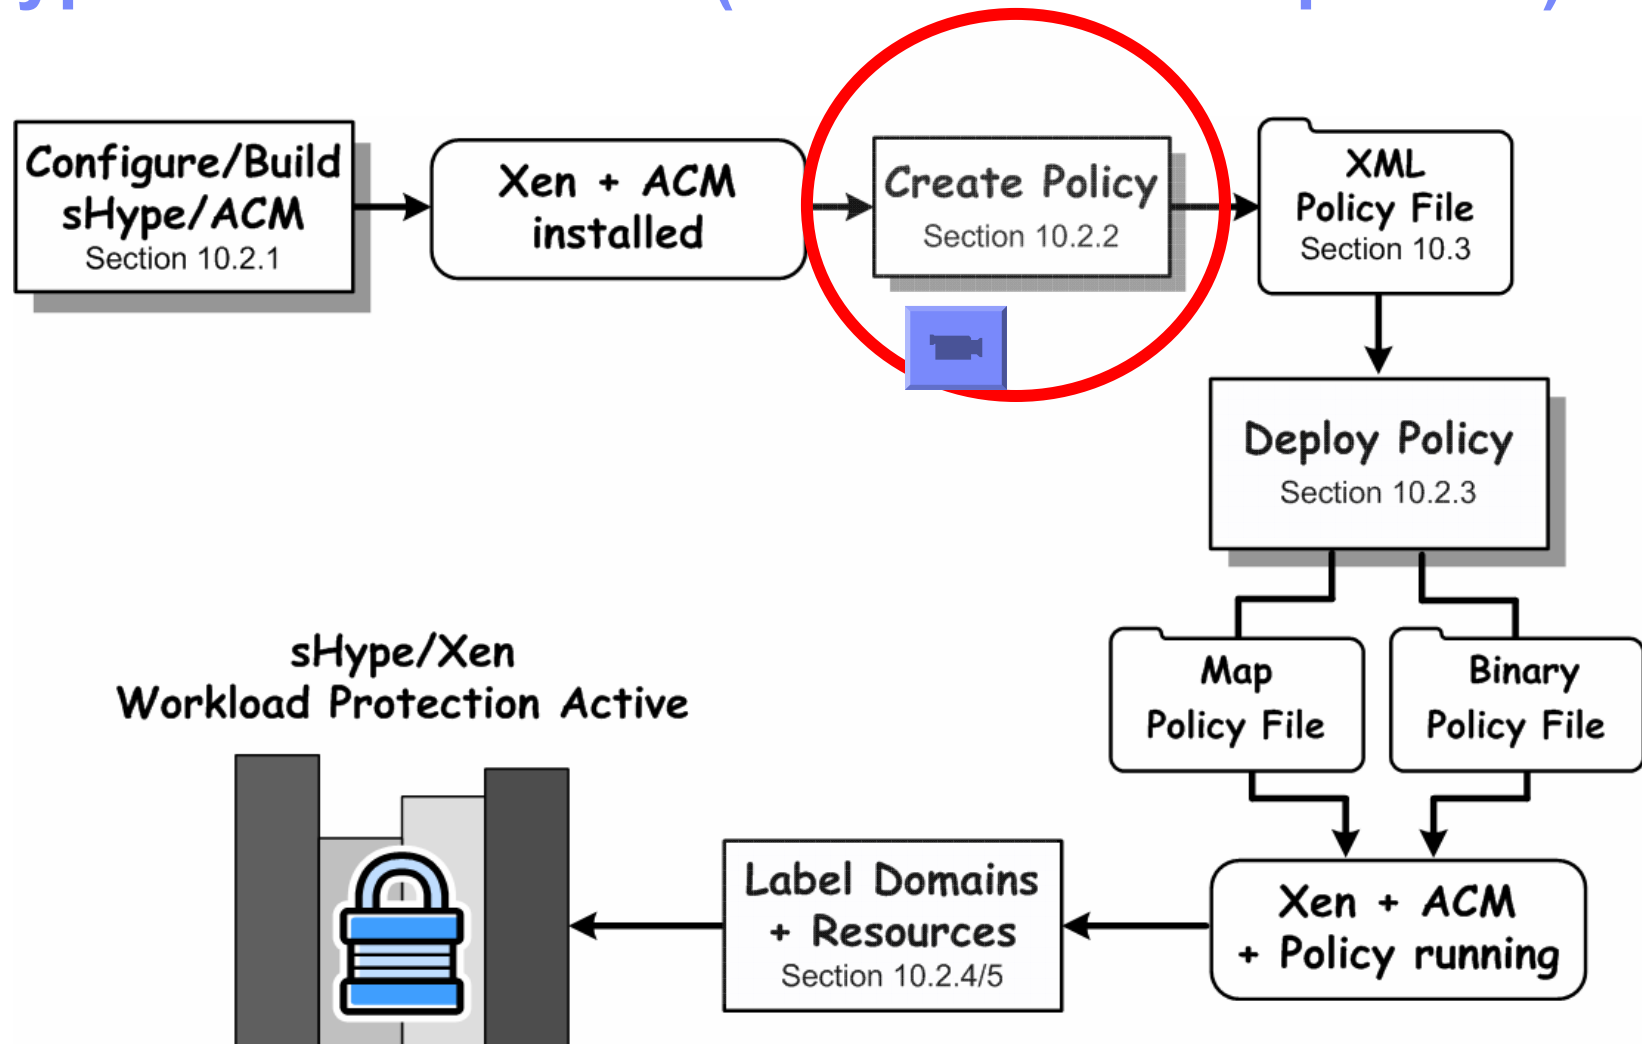
WL

Xen VMM  
(virtualizes + isolates)



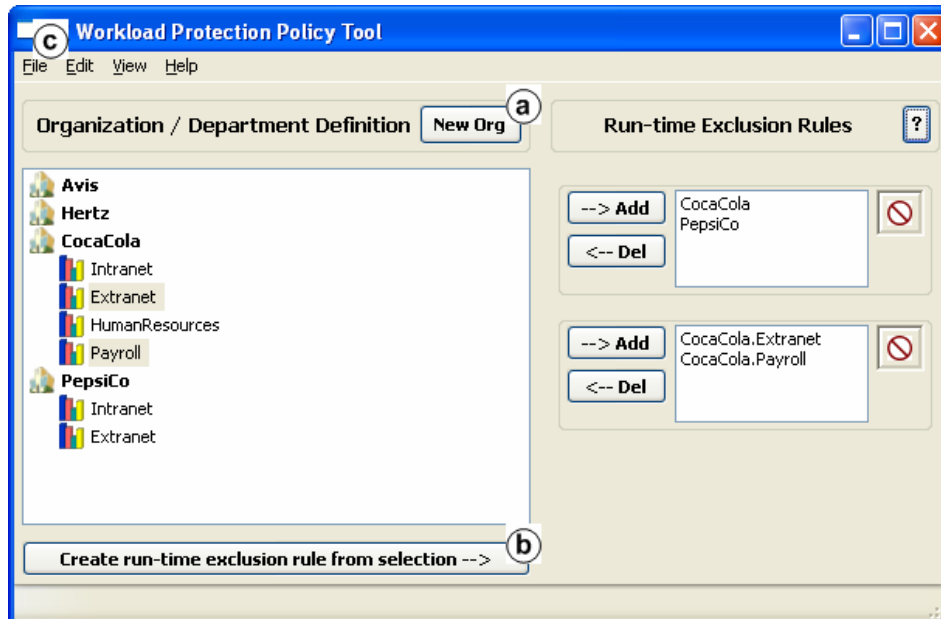

VM

# sHype/ACM Overview (User Guide Chapter 10)

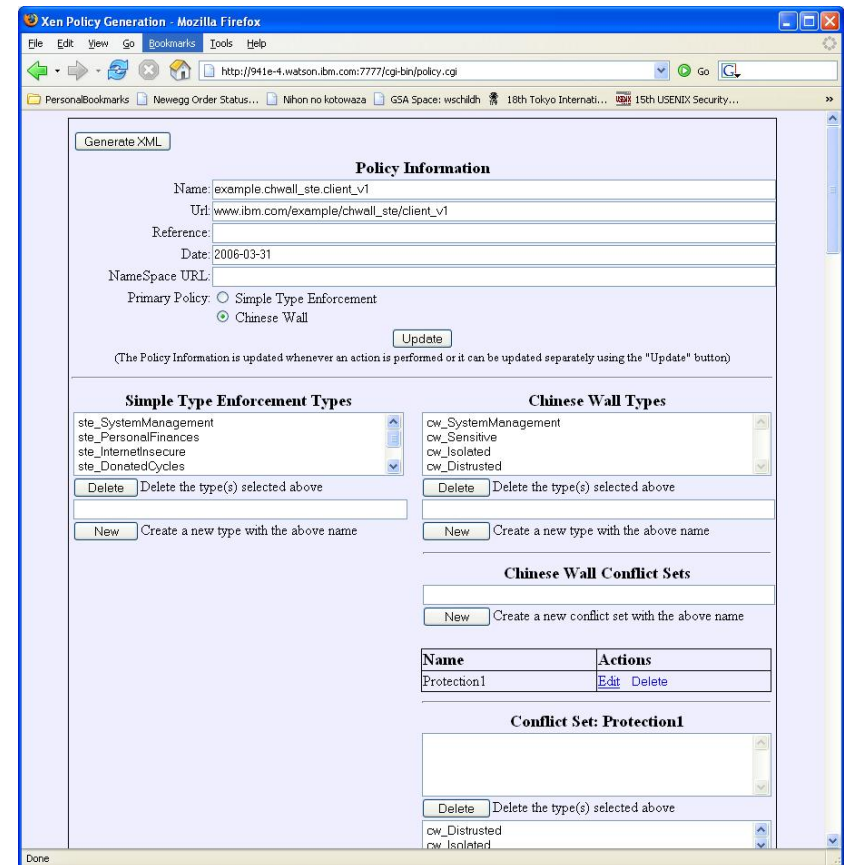


# EZ Policy Management

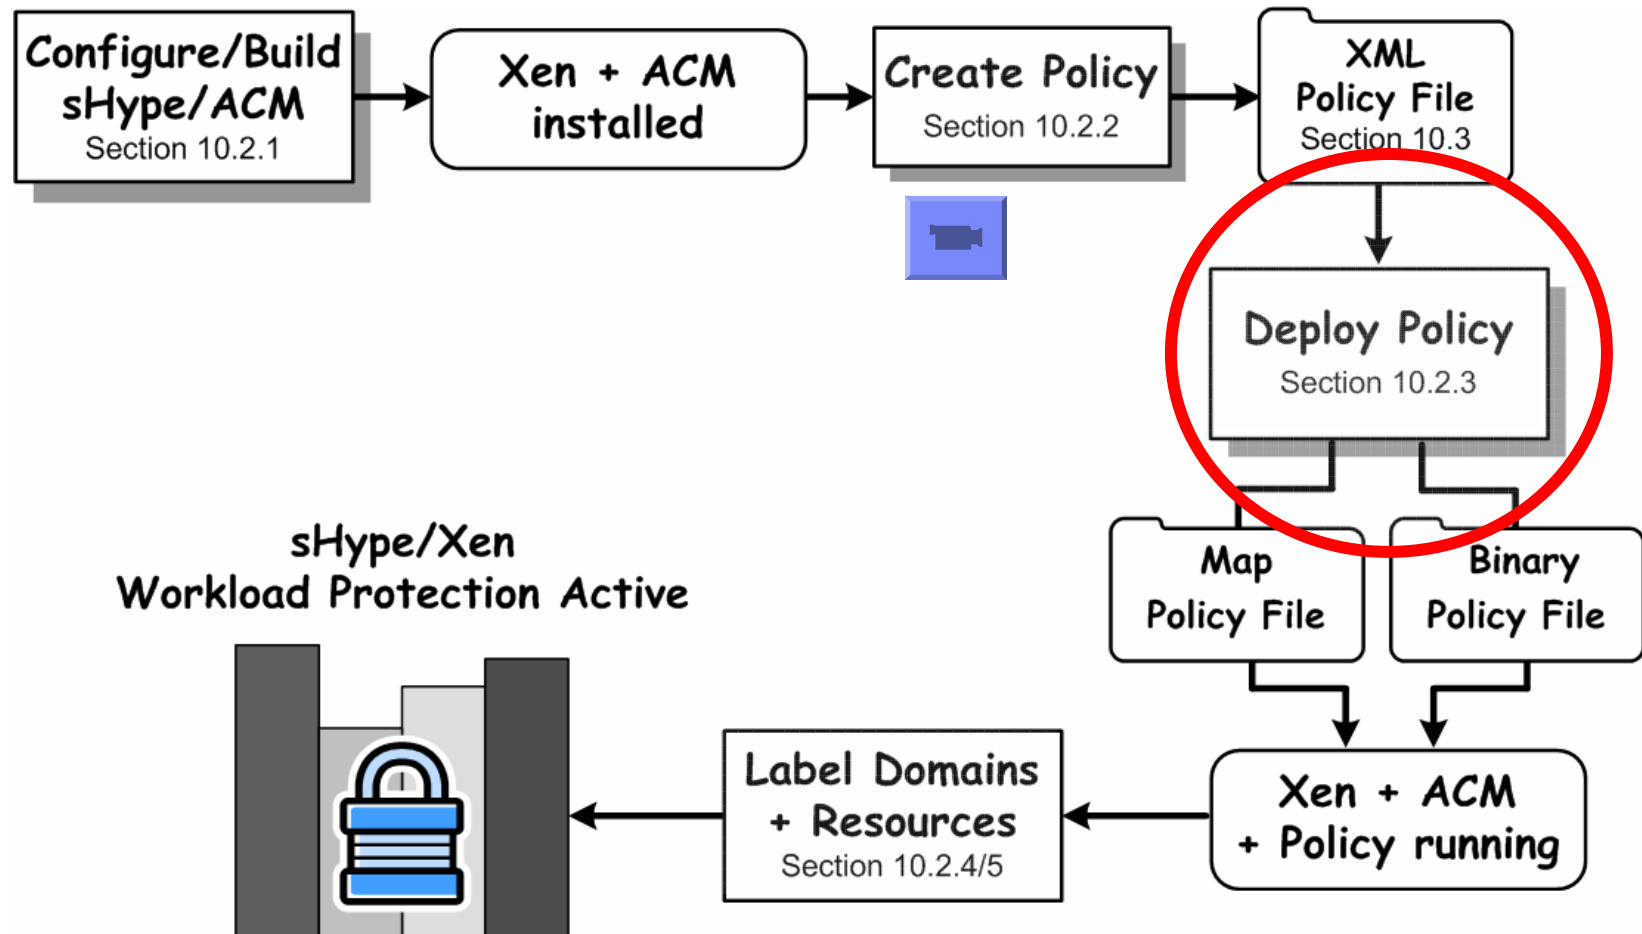
Create policies: xensec\_ezPolicy



Refine policies: xensec\_gen



# sHype/ACM Overview (User Guide Chapter 10)



## Translating & Loading Policies

```
# xm makepolicy example.chwall_ste.test-wld
```

***Policy root: /etc/xen/acm-security/policies***

***Policy files:***

***example/chwall\_ste/test-wld-security\_policy.xml***

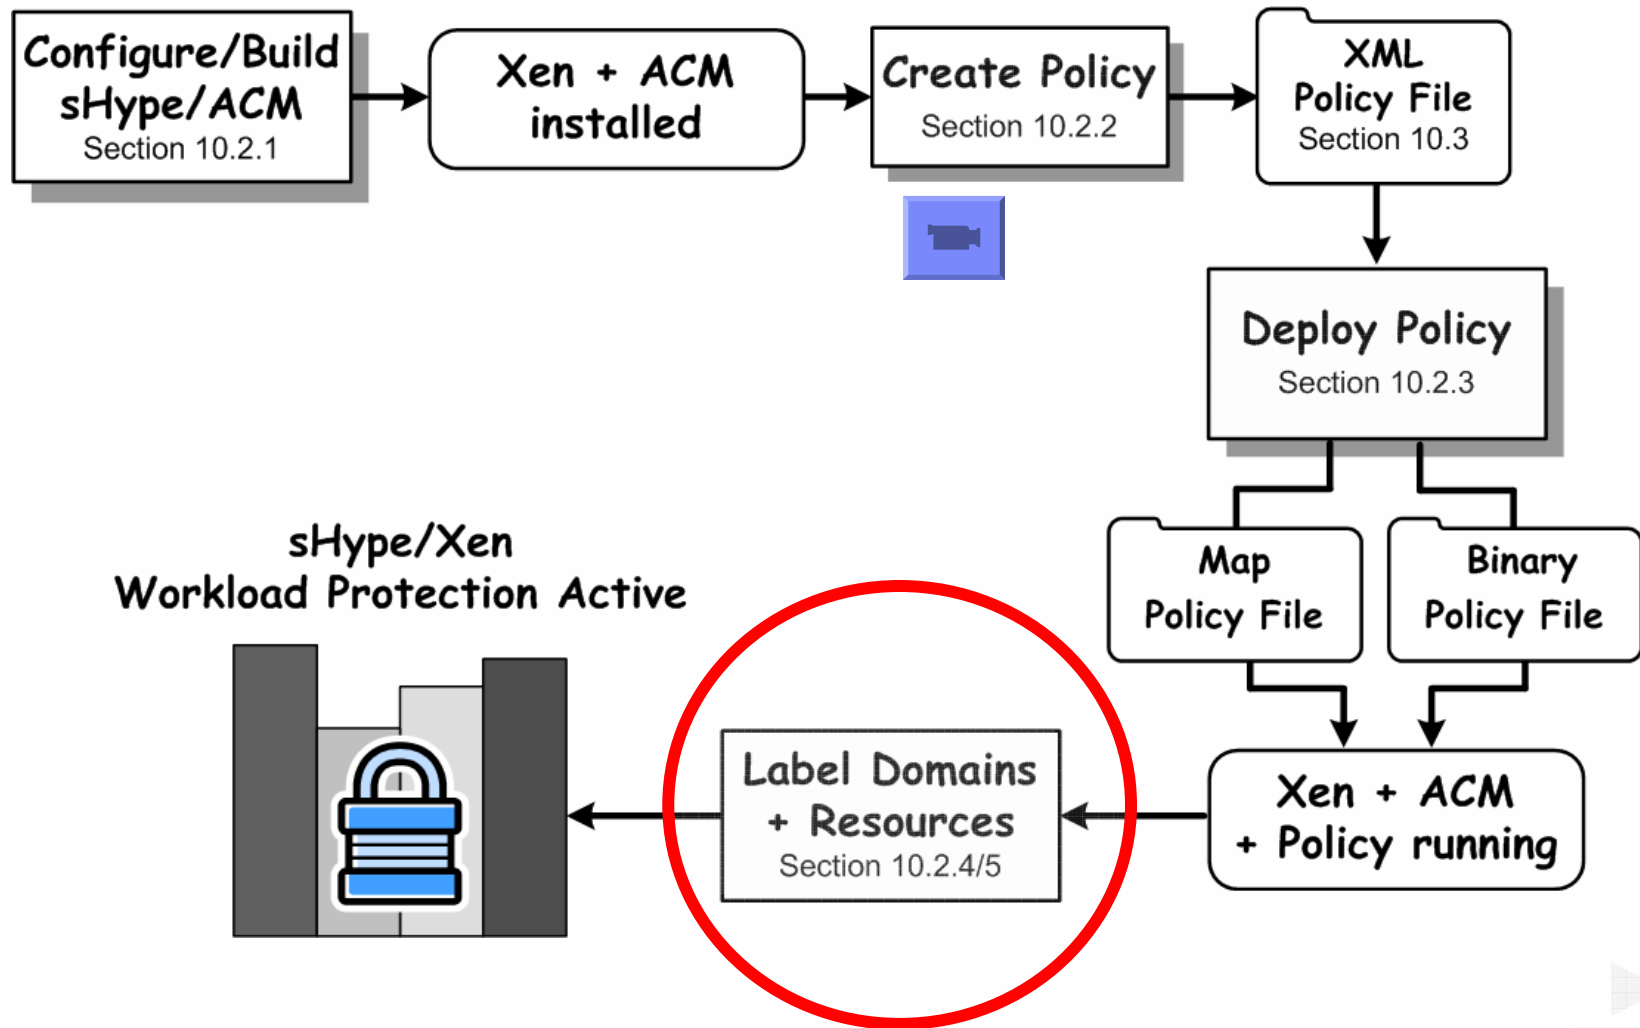
***example/chwall\_ste/test-wld.map***

***example/chwall\_ste/test-wld.bin***

```
# xm loadpolicy example.chwall_ste.test-wld
```

```
# xm cfgbootpolicy example.chwall_ste.test-wld
```

# sHype/ACM Overview (User Guide Chapter 10)



## Labeling Domains

```
#xm addlabel Avis.HR dom avis_hr.xml
```

```
kernel = "/boot/vmlinuz-2.6.16.13-xen"  
ramdisk="/boot/avis_hr_ramdisk.img"  
memory = 164  
name = "avisHR"  
vif = [ ' ' ]  
dhcp = "dhcp"  
disk = [ 'phy:sda3,sda3,w' ]
```

```
#####SHYPE-Labeling#####
```

```
access_control =  
    [ 'label=Avis.HR,  
      policy=example.chwall_ste.test-wld' ]
```

avis\_hr.xml

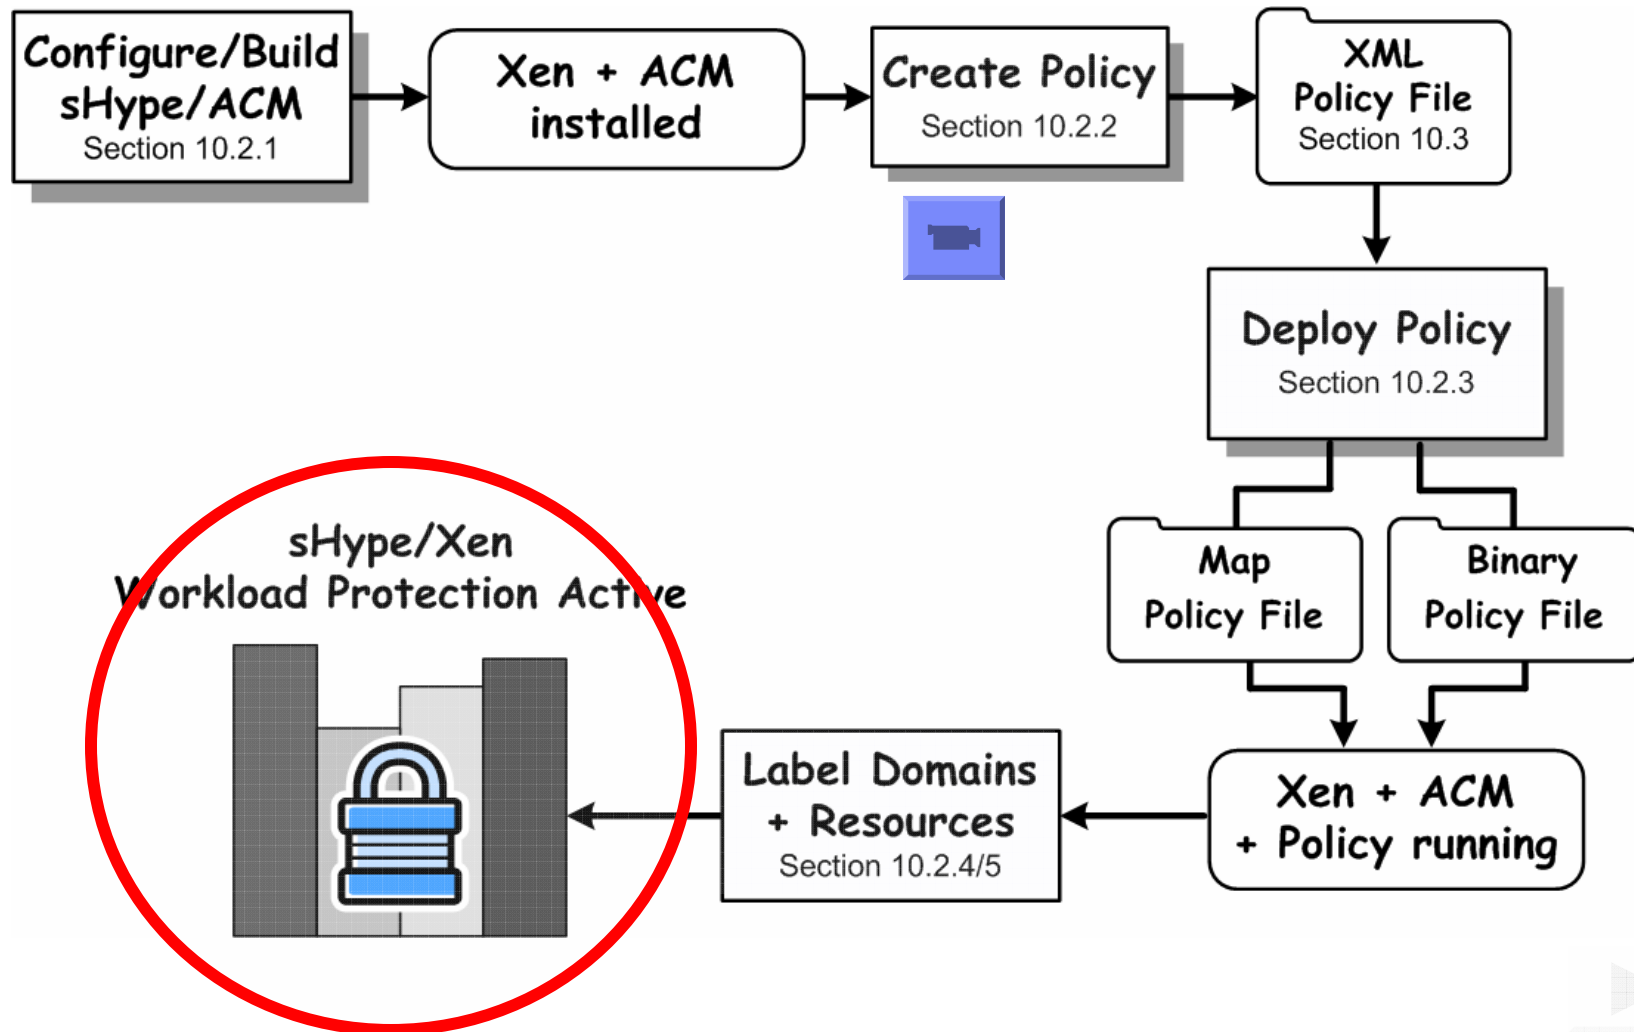
## Labeling Resources

```
#xm addlabel Avis.HR res phy:sda3
```

policies/resource\_labels

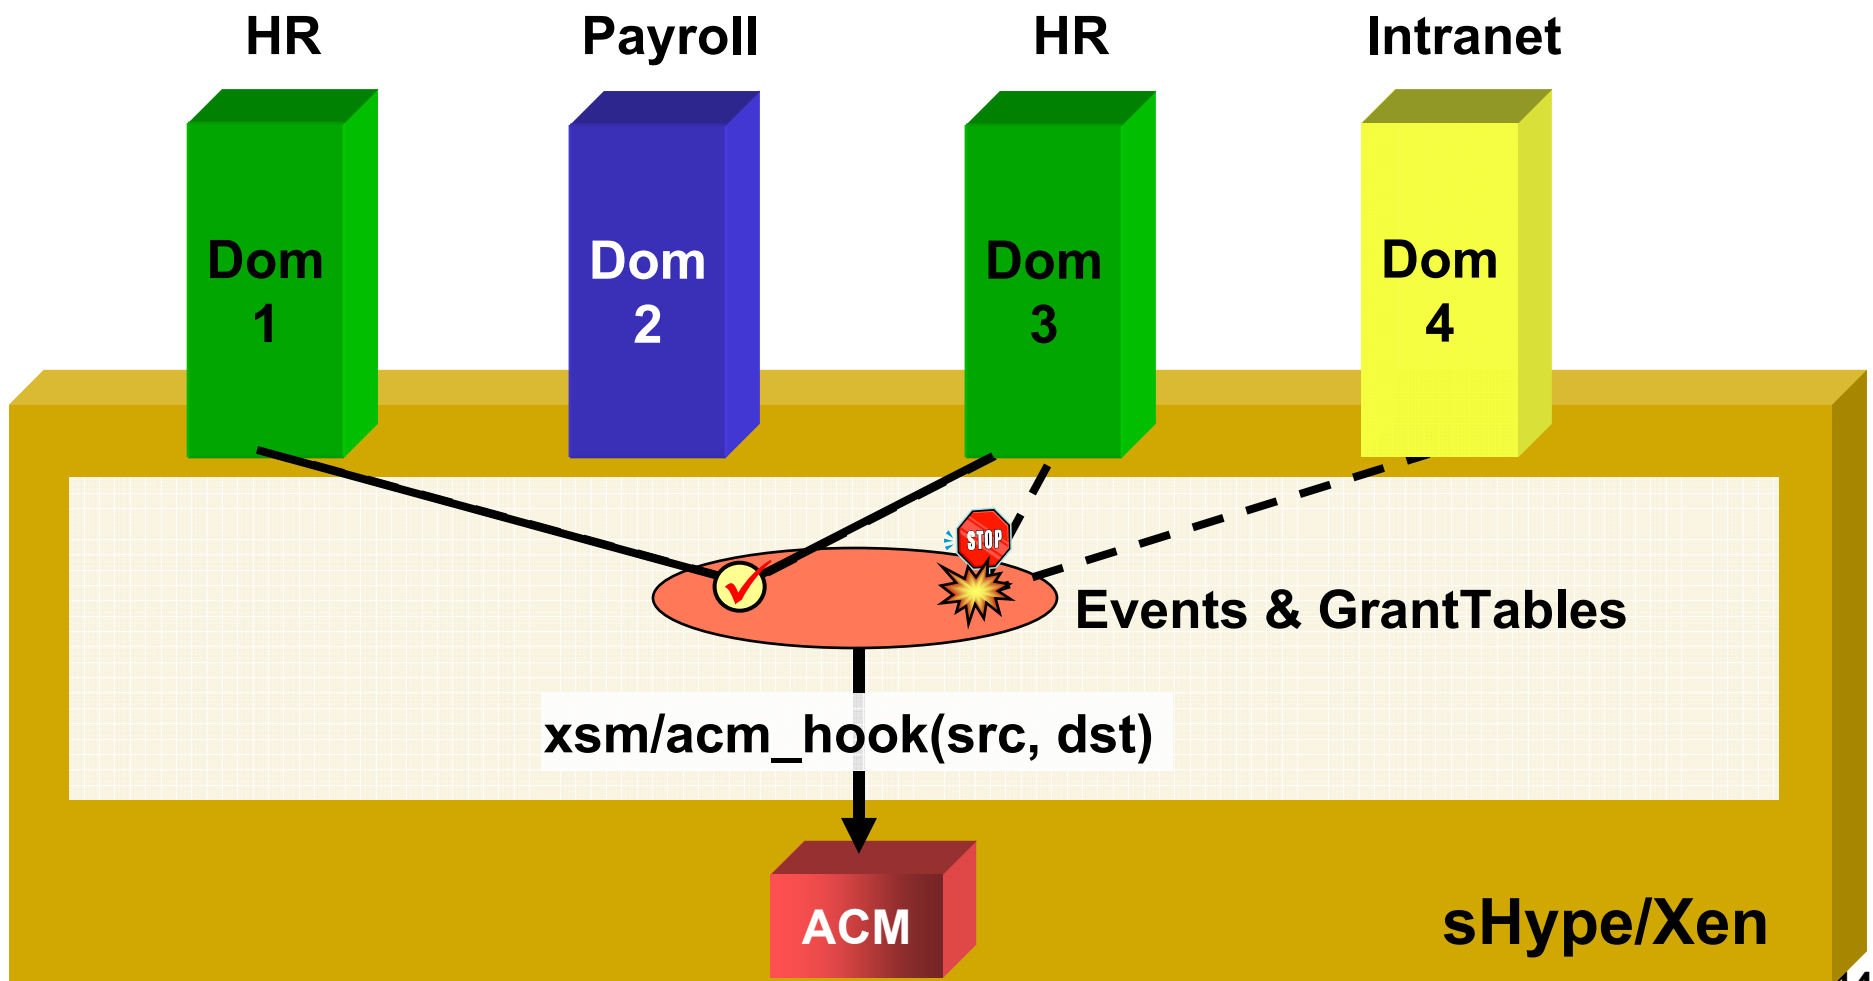
```
resources = {  
  'file:/xen/cocacola.swap': ('example.chwall_ste.test-wld', 'CocaCola'),  
  'phy:/dev/sda4': ('example.chwall_ste.test-wld', 'CocaCola '),  
  'phy:/dev/sda3': ('example.chwall_ste.test-wld', 'Avis.HR'),  
}
```

# sHype/ACM Overview (User Guide Chapter 10)

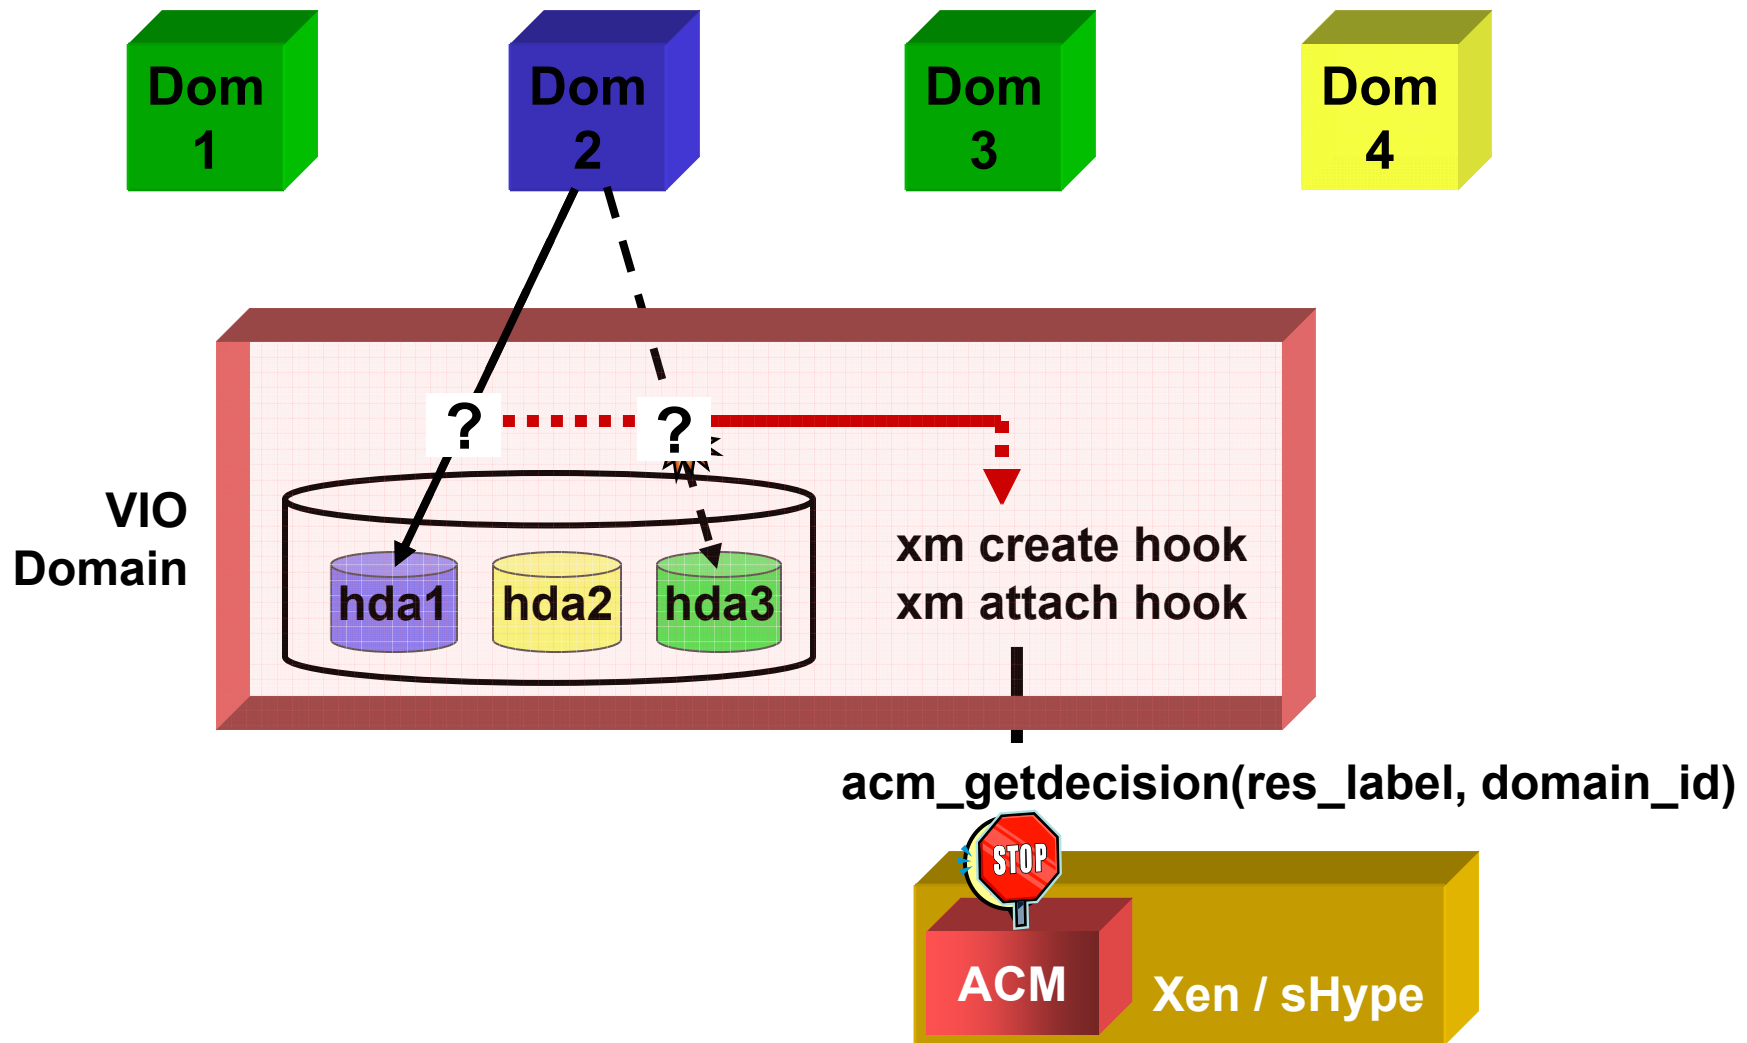


100 %

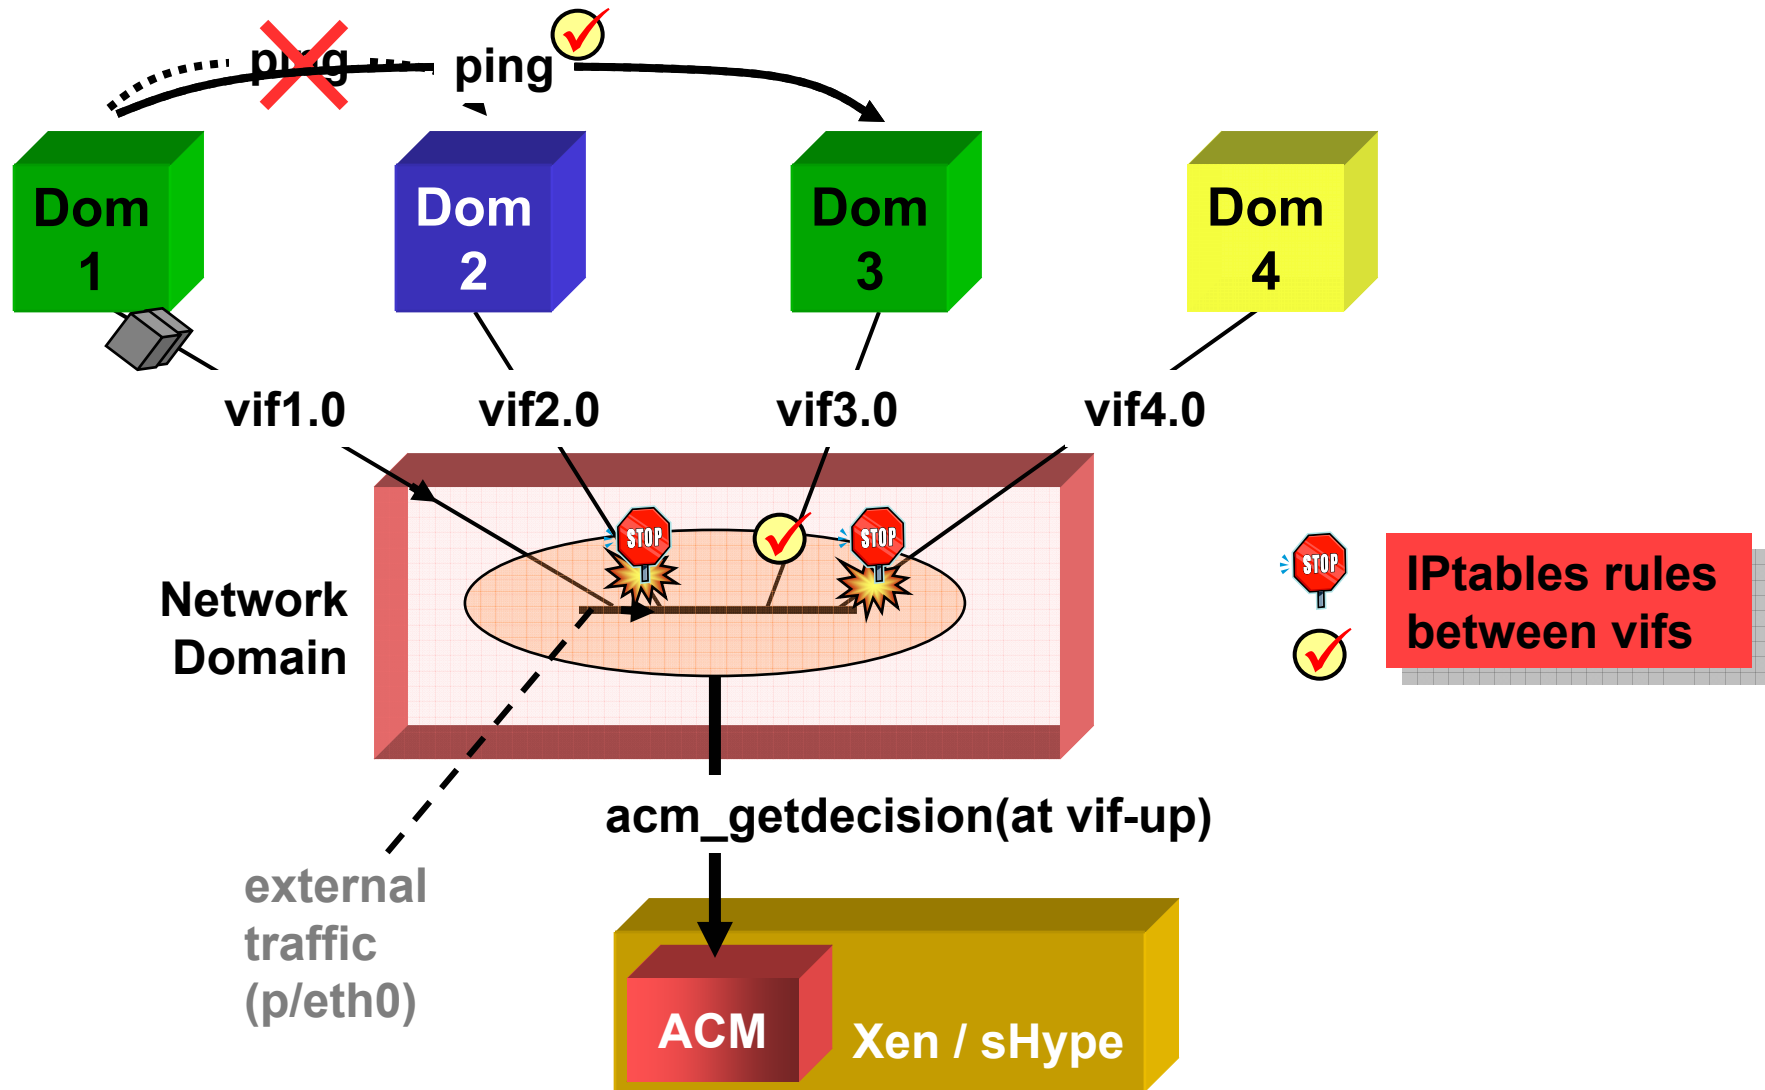
# sHype Distributed Workload Protection I: Simple Type Enforcement Policy (STE)




# sHype Distributed Workload Protection II: Policy Enforcement on Shared Resources



# sHype Distributed Workload Protection III: Policy Enforcement on Local Network Traffic



# sHype Distributed Workload Protection III: Policy Enforcement on Local Network Traffic



Chain FORWARD (policy DROP 24 packets, 6753 bytes)

- ACCEPT PHYSDEV match --physdev-in peth0
  - ACCEPT PHYSDEV match --physdev-out peth0
- } external traffic
- 
- ACCEPT PHYSDEV match --physdev-in vifX.0 --physdev-out vif0.0
  - ACCEPT PHYSDEV match --physdev-in vif0.0 --physdev-out vifX.0
- } 8 rules
- 
- ACCEPT PHYSDEV match --physdev-in vif3.0 --physdev-out vif1.0
  - ACCEPT PHYSDEV match --physdev-in vif1.0 --physdev-out vif3.0

## sHype Status Quo

- **Xen sHype/ACM User Guide Chapter**
- **ezPolicy Workload Protection Policy Creation Tool**
- **Resource Labeling and Enforcement**
- **IPtables Access Control on Local Network Traffic**

## What's Next

- **sHype/Xen Performance Evaluation + Optimization**
- **External Network Traffic + new Xen HW features**
- **Quarantining VMs, e.g., for problem isolation**

## CIM-Based Security Management Status Quo

- **Policy, domain and resource labels, and vTPM need to be managed**
- **Xen CIM/API offers management consolidation**
- **Currently CIM/API does not support sHype/vTPM**

### What's Next

- **Co-operate with Xen CIM/API group**
- **Define security extensions (ongoing)**
- **Enable security extensions in Xen CIM/API**

## vTPM Status Quo (Collaboration with Intel)

- **vTPM Device Integrated With Xen-Linux (hot-plug, drivers)**
  - no additional requirements on para-virtualization
- **Support For Fully Virtualized Domains**
  - Atmel qemu device model (submitted, pending)
- **Intel: Support For vTPM Migration**

## What's Next

- **TCG-BIOS guest support: CRTM, tGrub (ready)**
  - vTPM support when booting HVM
  - use vTPM, extend ACPI measurement logs for TPM
- **CIM support for vTPM management**

## Selected Related Work

- **IBM Converged Power Hypervisor**
- **High Assurance Platform (US Gov and vendors)**
- **Open Trusted Computing (European Framework)**
- **Intel Collaboration on vTPM**
- **Data Diode (NRL, Network Pump)**
- **University Relations w.r.t. sHype/Xen**
  - Pennsylvania State Univ, Georgia Tech, Carnegie-Mellon Univ, Brown Univ, Worcester Polytechnic Institute, ...**

## Xen Security People @ T. J. Watson

- **Stefan Berger – vTPM**
- **Ramón Cáceres – CIM-API**
- **Reiner Sailer\* – sHype Access Control Module**
- **Guest: Bryan D. Payne\* (Georgia-Tech)  
Resource Labeling**
- **Ronald Perez\* + Leendert van Doorn**

(\* .. attending this Xen summit)