



Secure Systems Department

Security for the Xen Hypervisor – Status Quo & Perspective 2006

Reiner Sailer

Xen Summit 2006

1. Access Control Module

2. Virtual Trusted Platform Module

Hypervisor Security Architecture / ACM

Major Goal:

Create distributed confined operating environments

- Allow controlled sharing between domains, building “coalitions” of domains and virtual peripherals

Why: To better protect distributed services / workloads

How: Xen/ACM confinement serving as a

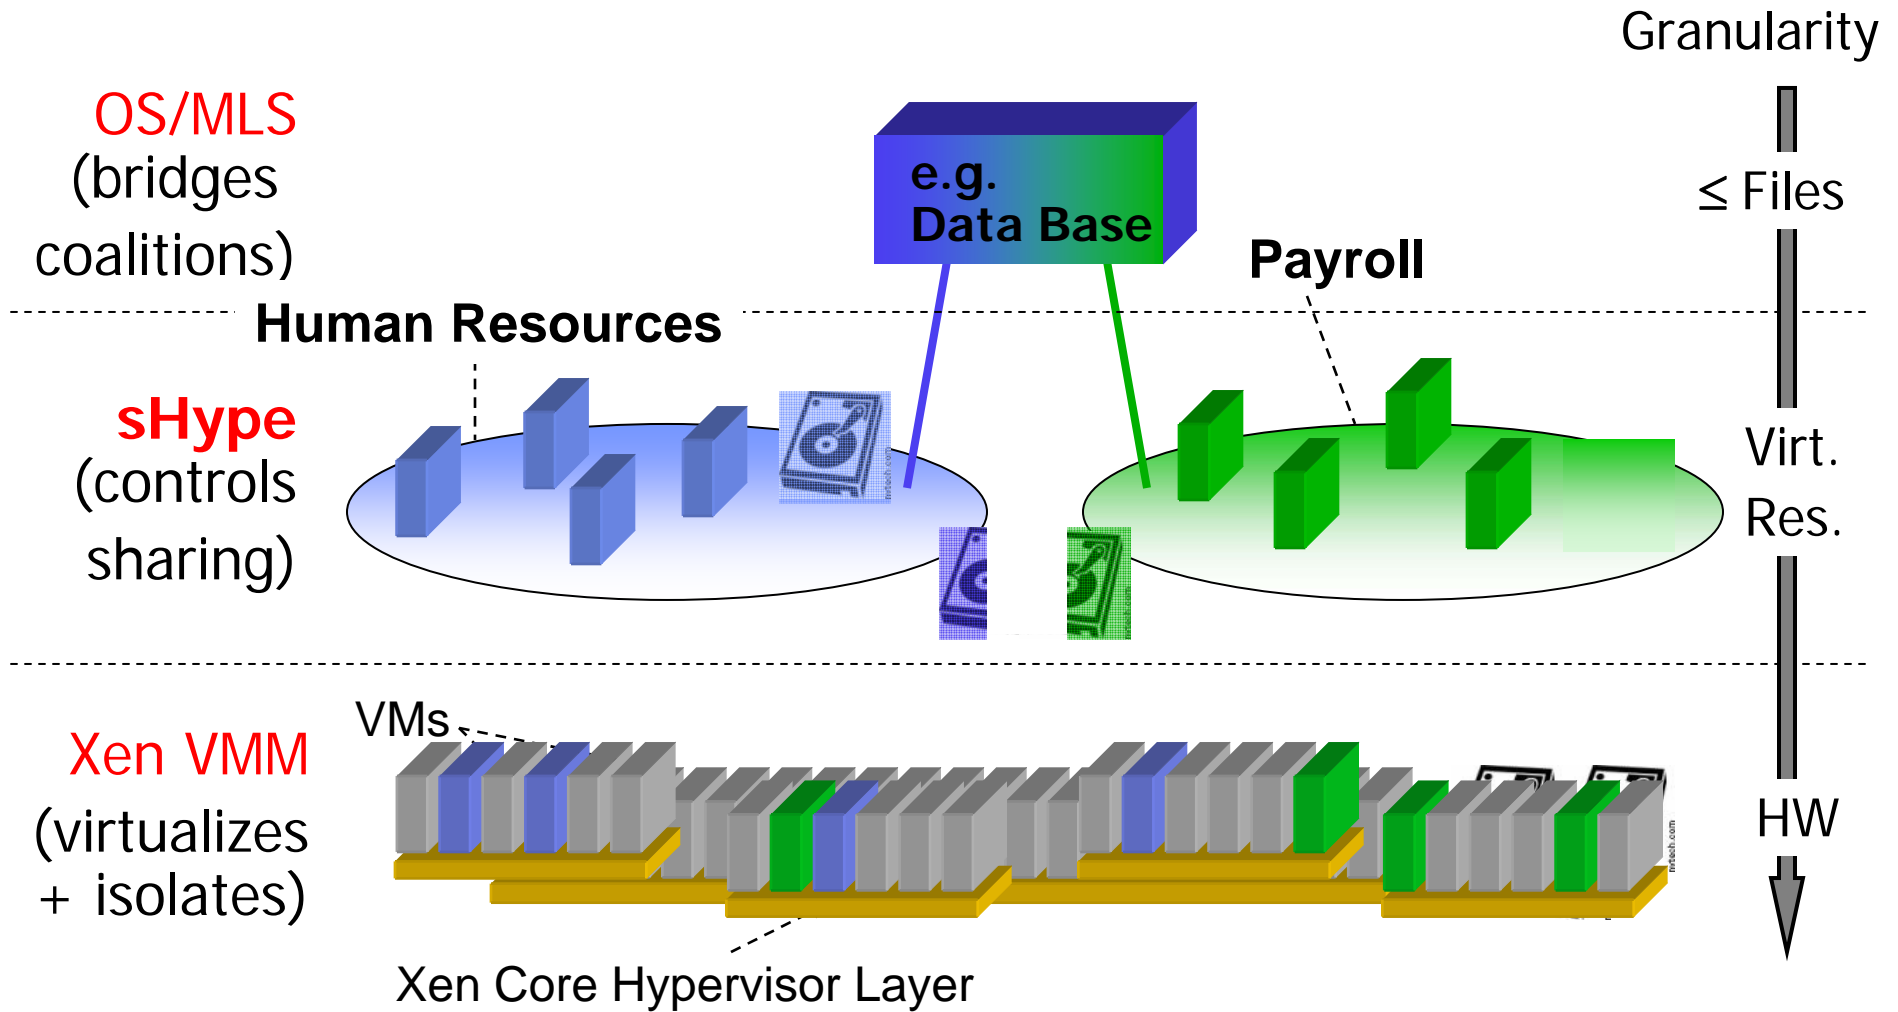
- **Universal foundation** for OS security
- **Safety net** if OS security fails or is missing

Policy Support Status Quo

- **Simple Type Enforcement Policy**
 - Controls which domains can share
[Coalitions for Domains]

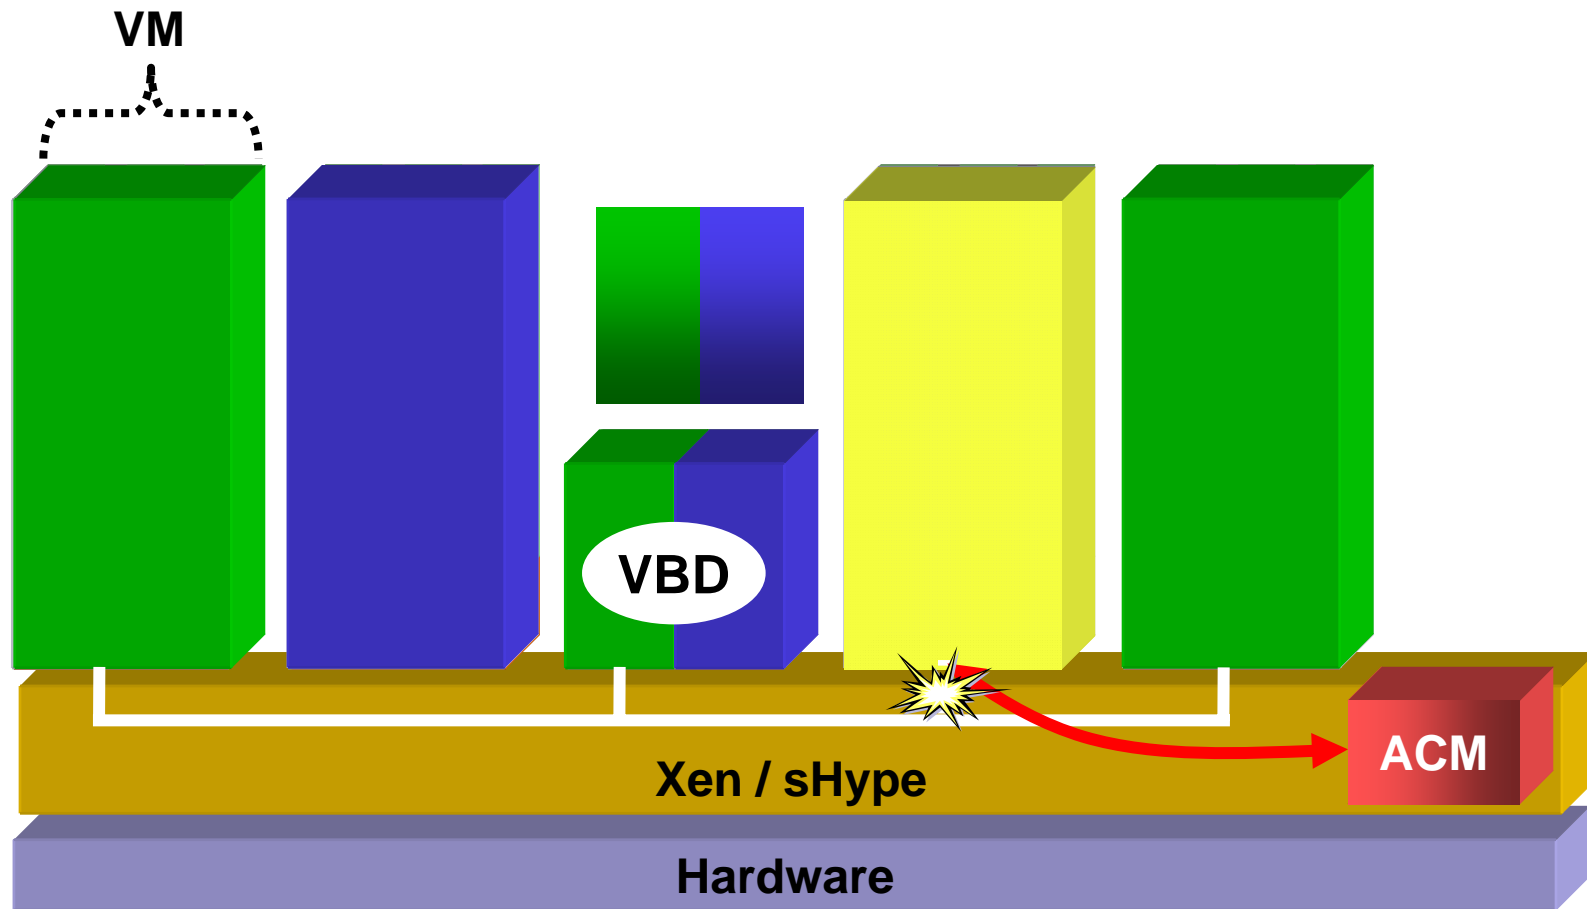
- **Chinese Wall Policy**
 - Controls which domains can run simultaneously on the same system
[Approximated “Air-Gap” between Domains]

Layers of Isolation and Sharing



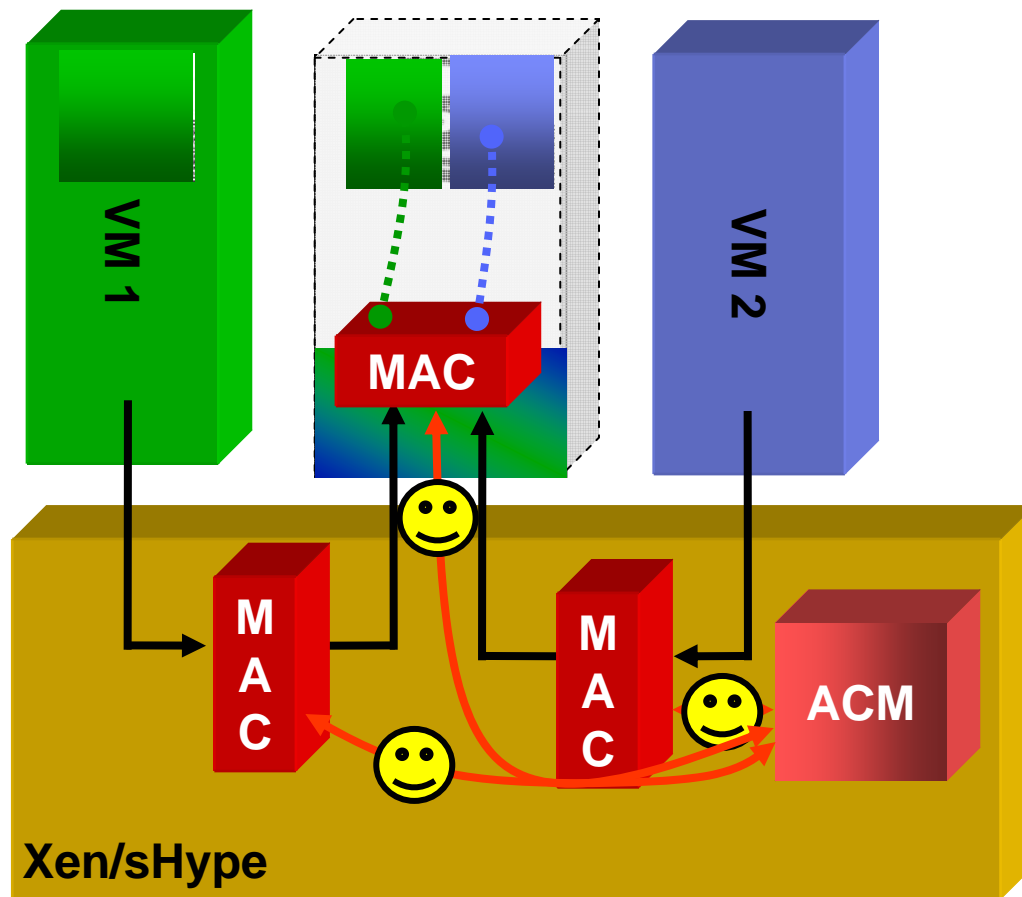
Simple Type Enforcement Policy (STE)

Example: 3 Coalitions – Yellow, Blue, Green



Sharing of HW Devices Through Isolated Virtual Devices

sHype coalition isolation no stronger than MAC OS Isolation



- Create isolated virtual devices to prevent over-provisioning
- Xen/sHype controls which VM can connect to MAC VM
- **Xen/sHype defers MAC enforcement to MAC-OS**
- Xen/sHype can provide access control decisions

Perspective 2006 – Access Control Module

- **Encourage Security Community to Build on top of Xen/ACM Policy (e.g., SELinux/MLS)**
- **Extending Access Control Enforcement onto virtual peripheral devices (VBD, V-NIF)**
- **Extending Access Control across multiple platforms**

Challenges – Security Adds to Trend of Refactoring Dom0

- **Refactoring of Domain0 Essential**

- It's too large (LOC) → no (strong) guarantees possible
- It's too powerful → no confinement in case of error
- It's ever changing → very hard to stay non-intrusive

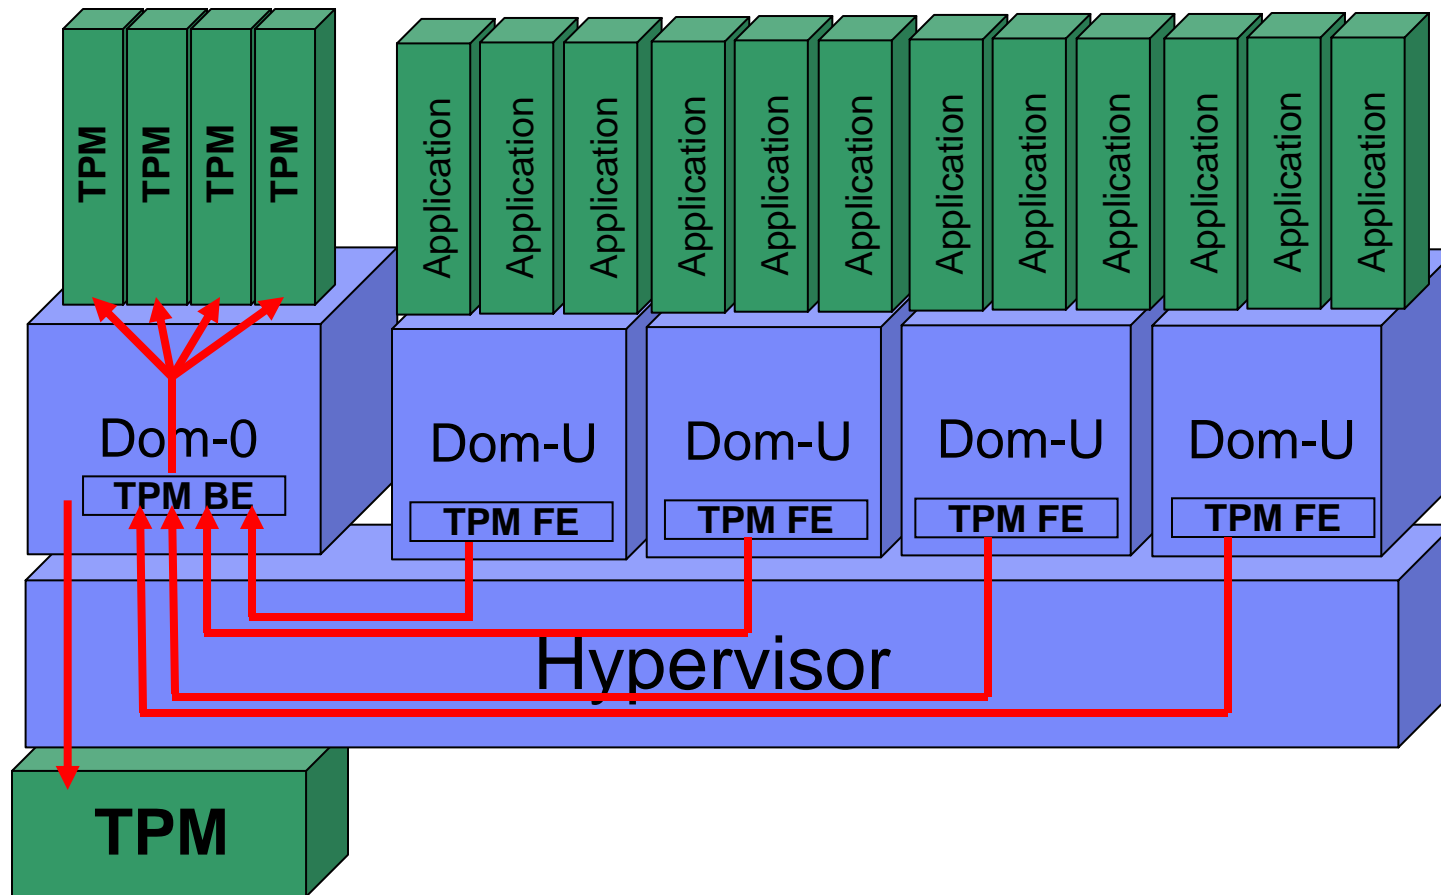
- **Driver aka Device aka MAC Domains**

- Small domains own hardware and create isolated virtual peripherals that can be assigned to different coalitions (assumes IO-MMU)

1. Access Control Module

2. Virtual Trusted Platform Module

Virtual TPM Support in Xen



Perspective – Virtual TPM

- **vTPM drivers in close co-operation with Intel**
 - Migration support for domains with vTPM
- **Intel driving authenticated boot**
 - Minimal startup partition to initialize the access control and vTPM environment
 - Measuring Xen components
- **Co-operation with AMD on dynamic root of trust**
- **Attestation support for domains (IMA)**

Summary

- **Policy and enforcement in Xen is stable**
- **Labeling and policy enforcement for virtual VM resources continuing (network + VBD)**
- **Extending sHype across multiple platforms continuing**
- **vTPM support expanding (dynamic root of trust)**
- **Security Requirements add to the Trend of Refactoring Domain0**

Available Xen/ACM Security Tools

- **Policy Management Interface (Web-based)**
- **Command line policy translation tools (xml policies → Xen/ACM binary policies)**
- **Labeling tools (setlabel, getlabel)**

- **IBM LTC: Tom Lendacky**
- **IBM Research: Stefan Berger, Reiner Sailer**