U.S. NAVAL RESEARCH LABORATORY

Enterprise Scale Separation VMM Systems

Myong Kang

Computer Security Section
Naval Research Laboratory
Washington, DC 20375

myong.kang@nrl.navy.mil

U.S. NAVAL RESEARCH LABORATORY

Issues of Some Secure Systems

- Poorly constructed secure systems
  - Error prone
  - Difficult to understand and maintain
- Security mechanisms/features are too complex or not matching with underlying systems to protect
  - Too difficult to specify security policies
  - Too difficult to understand security policy and enforcement mechanism
- Not enough qualified sys admins / security officers to manage the system in the organization
  - Full power of security mechanisms of the system is not being utilized
    - Security or performance may be sacrificed
    - Complex security policy and mechanism compound the problem

U.S. NAVAL RESEARCH LABORATORY

Xenon Construction

- Replaced Xen's FLASK security architecture with a simpler, more intuitive Xenon security architecture
  - Separation of security policy and enforcement mechanism
    - Intuitive and expressive policy language and interface that are tailored to the hypervisor
    - Easy to validate security mechanisms
    - Easy to understand security policies
- Reduced in size and refactored the source code to significantly reduce its cyclomatic complexity
- Reduced attack surface of control domain (Domain-0)
- Provided a visual management interface for easy and intuitive security policy authoring, VM management and VM monitoring
- Added a network access policy that is easily configurable through the management interface
  - Enforced by Open vSwitch to restrict VM network access

U.S. NAVAL RESEARCH LABORATORY

Xenon Security Features

- Security Tags – e.g., Red, Yellow, Admin, Operational, Entertainment
- Enclave – Container for resources (e.g., VMs, hardware) that inherit the same Tag
- Relationship among Tags
  - A “conflict” between two or more tags can be defined, which will be translated into virtualization connectivity rules enforced by Xenon and network connectivity rules enforced by Open vSwitch
    - Hypervisor does not allow any communication between VMs with different tags
    - Open vSwitch does not allow any network communication between VMs with different tags
  - Option for VMs with two different tags cannot run at the same time on an Xenon host

U.S. NAVAL RESEARCH LABORATORY

Xenon Security Features (cont'd)

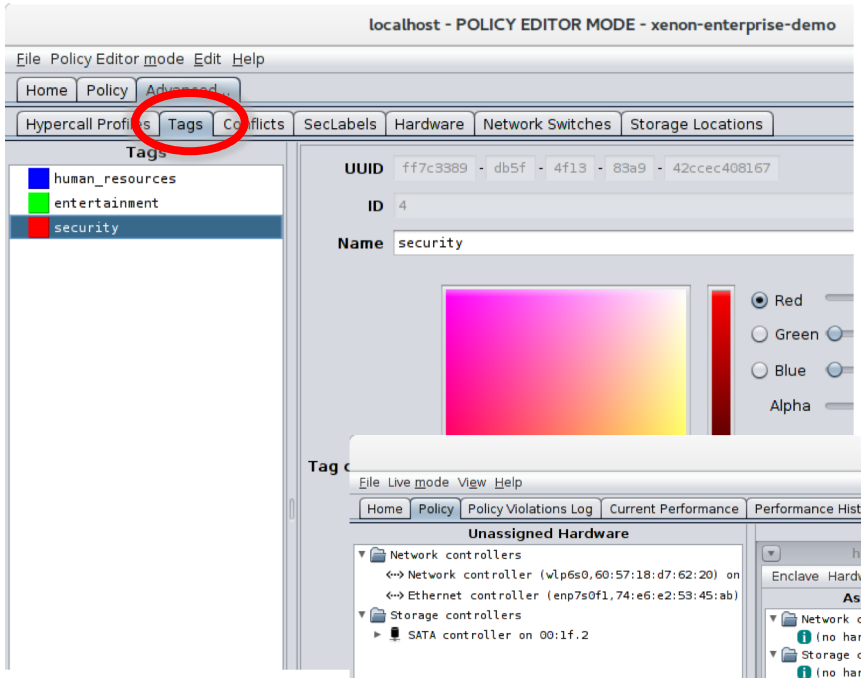

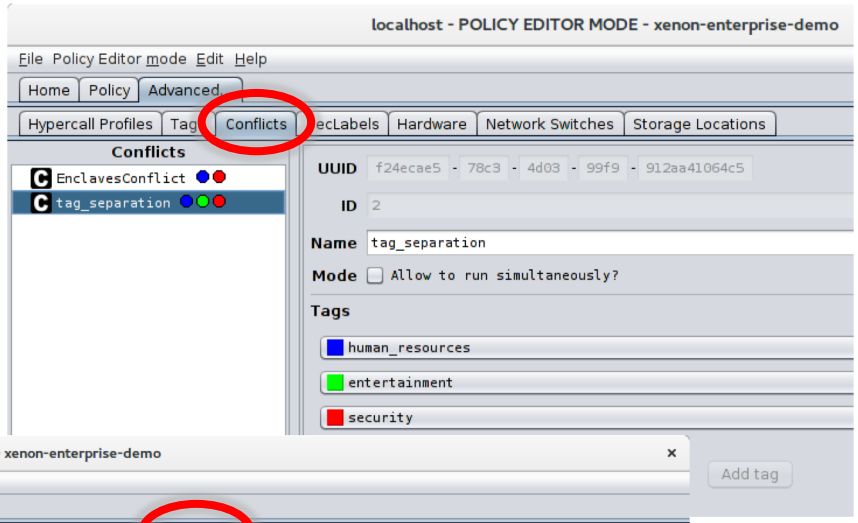

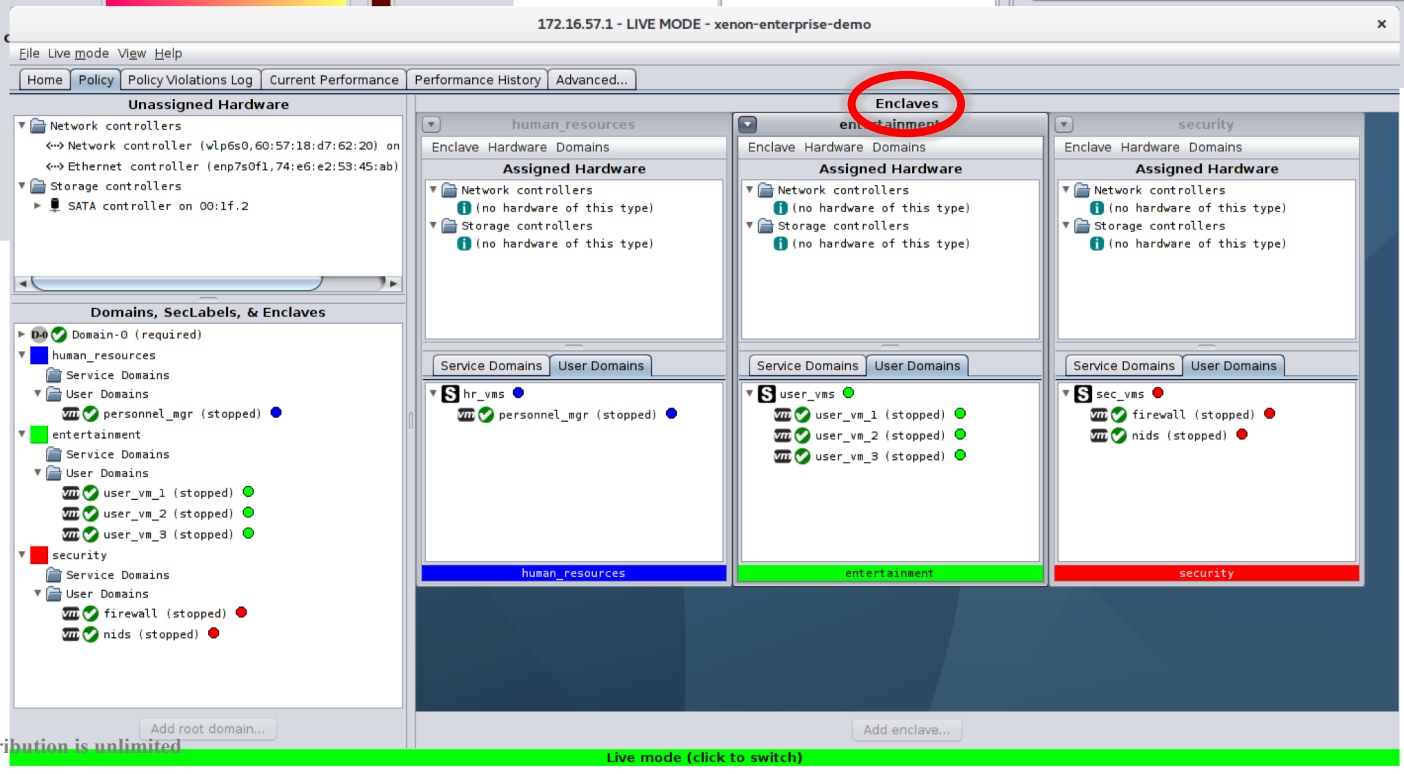

U.S. NAVAL RESEARCH LABORATORY

Xenon Security Features (cont'd)

- Xenon security policy restricts each VM's interface to the hypervisor via *hypercall privileges and profiles*
  - A VM interacts with the hypervisor through hypercalls
  - Hypercall privilege levels
    - Management privileged hypercalls – mainly Dom0
    - Security privileged hypercalls – VM introspection, etc.
    - Regular hypercalls
  - Every VM is required to have an associated hypercall profile
    - VM introspection domain profile, Ubuntu 14.04 profile, etc.

U.S. NAVAL RESEARCH LABORATORY

Xenon Security Features (cont'd)

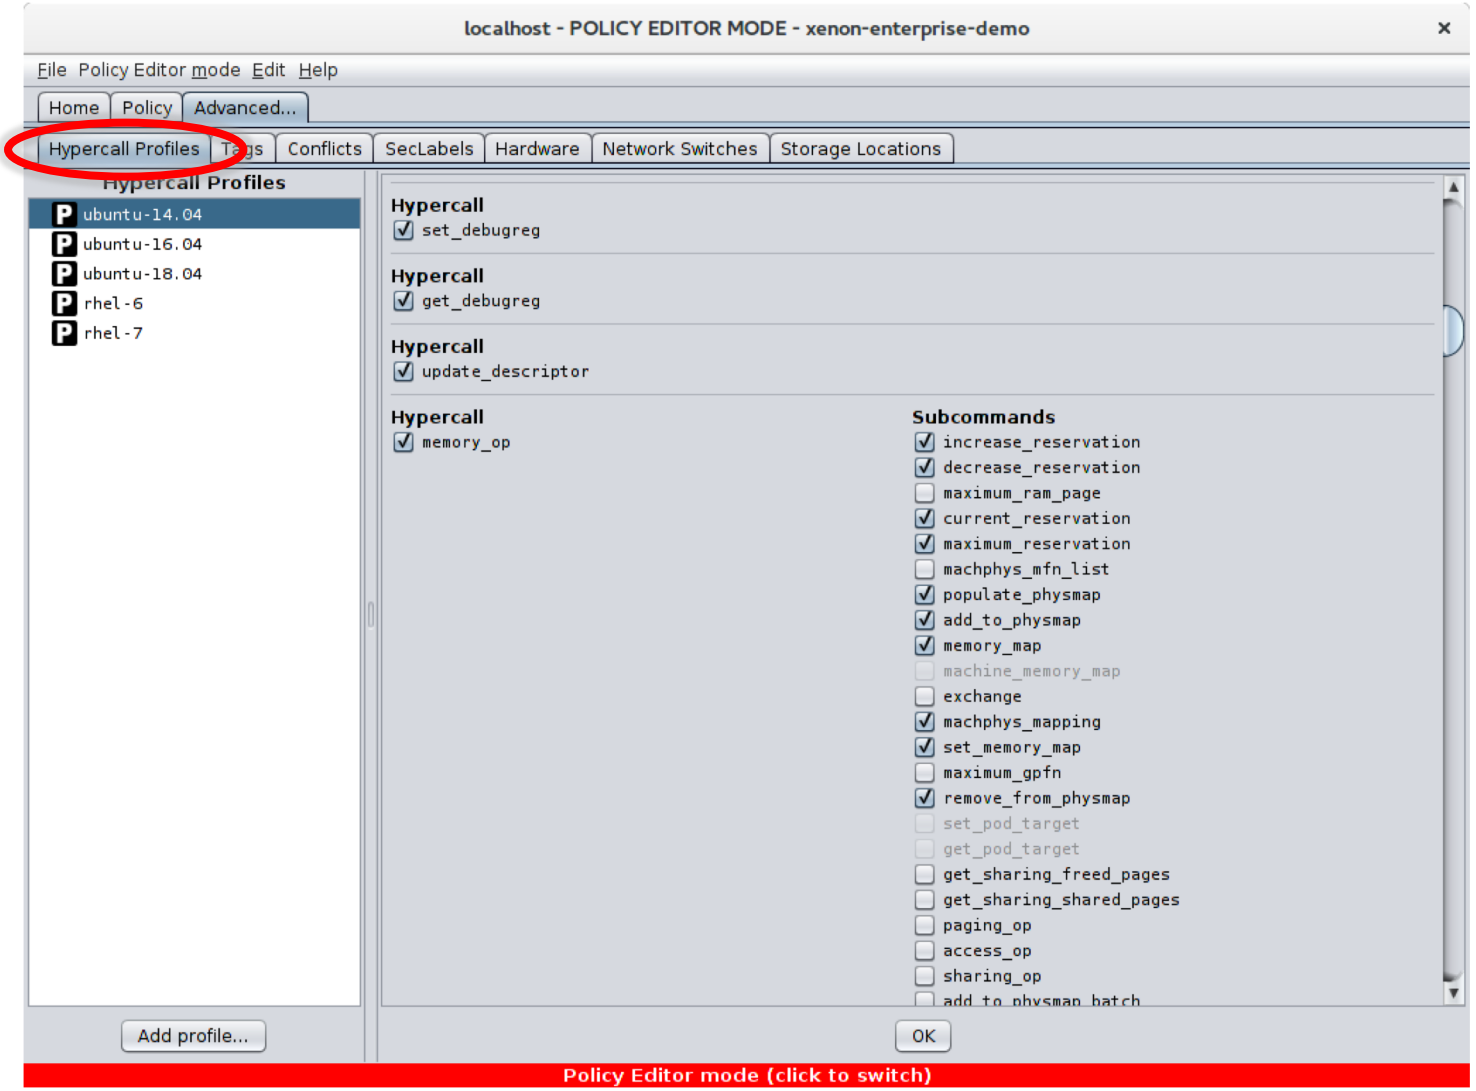

U.S. NAVAL RESEARCH LABORATORY

Xenon Security Features (cont'd)

- Persistent audit logging of all policy violations
  - Logs are stored outside the VM, physically separated from and inaccessible by the violating VM
  - Security policy can be configured to permit a maximum number of violations per VM
- Xenon security policies can be set dynamically (i.e., modified at runtime without rebooting)
  - Supports scalability and significantly reduces the need for downtime
  - Tags and VMs can be added but not removed

U.S. NAVAL RESEARCH LABORATORY

Xenon Security Features (cont'd)

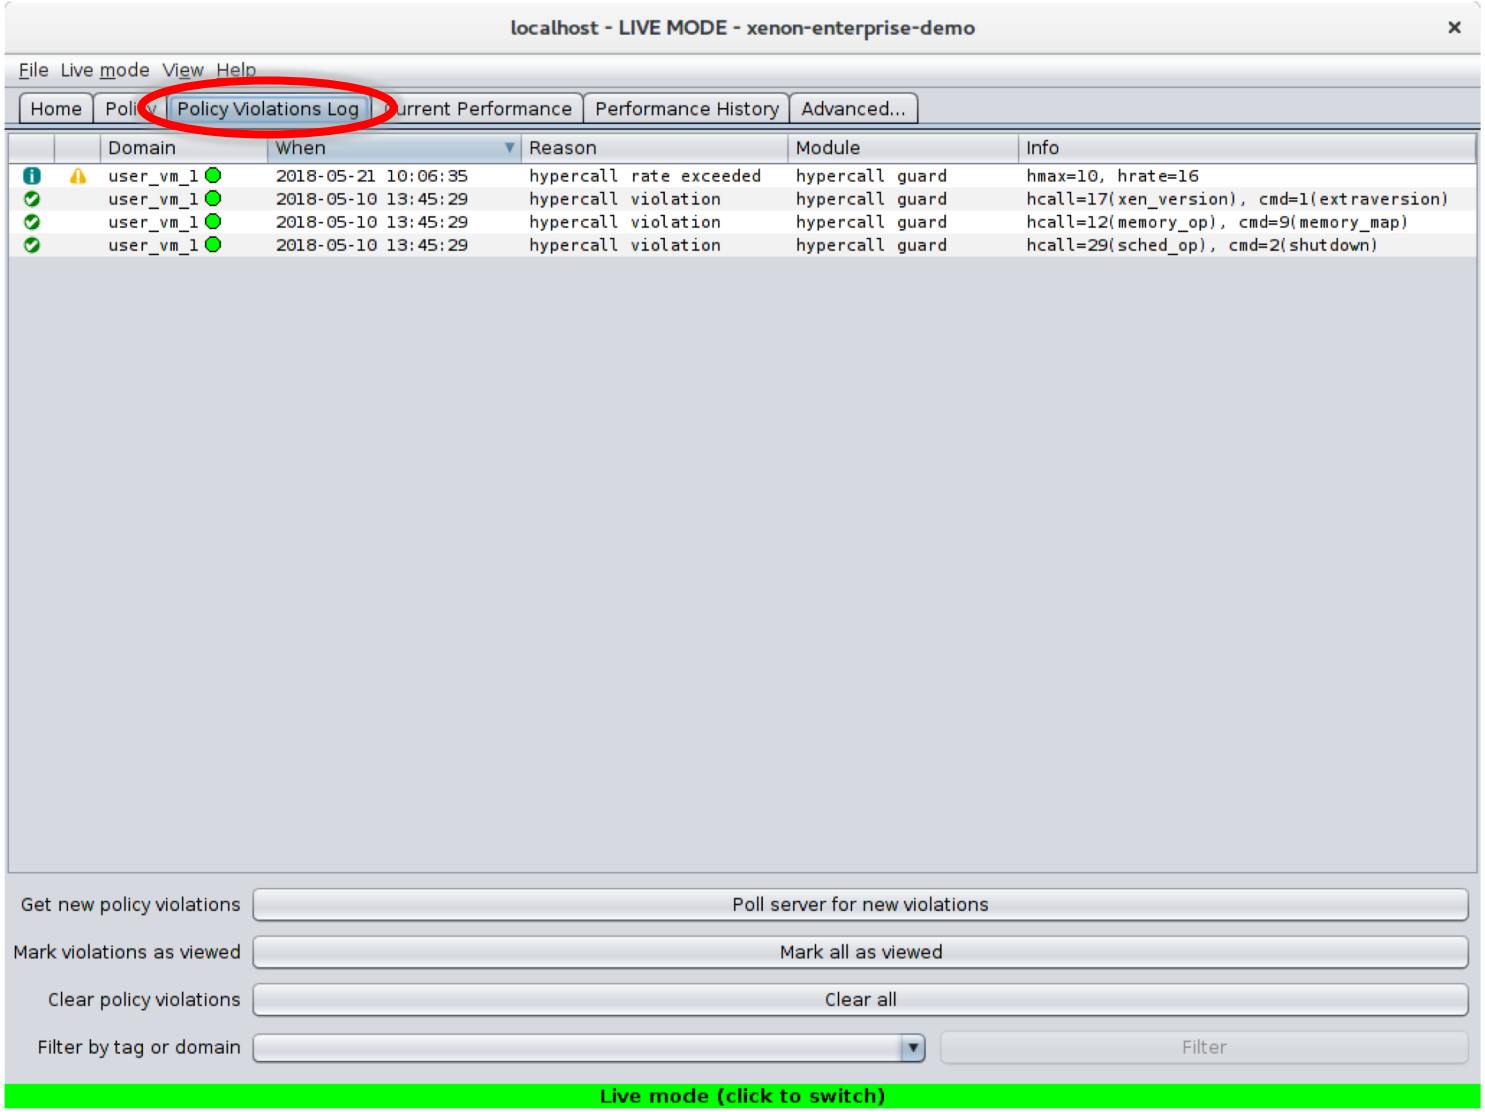

U.S. NAVAL RESEARCH LABORATORY

Xenon is More Than A Secure Hypervisor

- Out of the box, customers gain security and performance best practices
  - Preconfigured service VMs (e.g., network service VM)
  - Distribution of security policy through a signed XML policy document
  - Role-based access control for different administrative tasks
    - System admin role (e.g., starting, migrating VMs)
    - Security policy admin role (e.g., setting up security policies)
    - System monitoring role
- Guide users to configure virtual computing environment securely through intuitive visual management tool
  - Provide hardware-independent secure configuration mechanism
    - Export & import policy “templates”
  - Visual management interface provides
    - An easy and intuitive way to author and understand security policy
    - An easy way to manage and monitor user domains

U.S. NAVAL RESEARCH LABORATORY

Xenon is More Than A Secure Hypervisor (cont'd)

Hardware overview at-a-glance

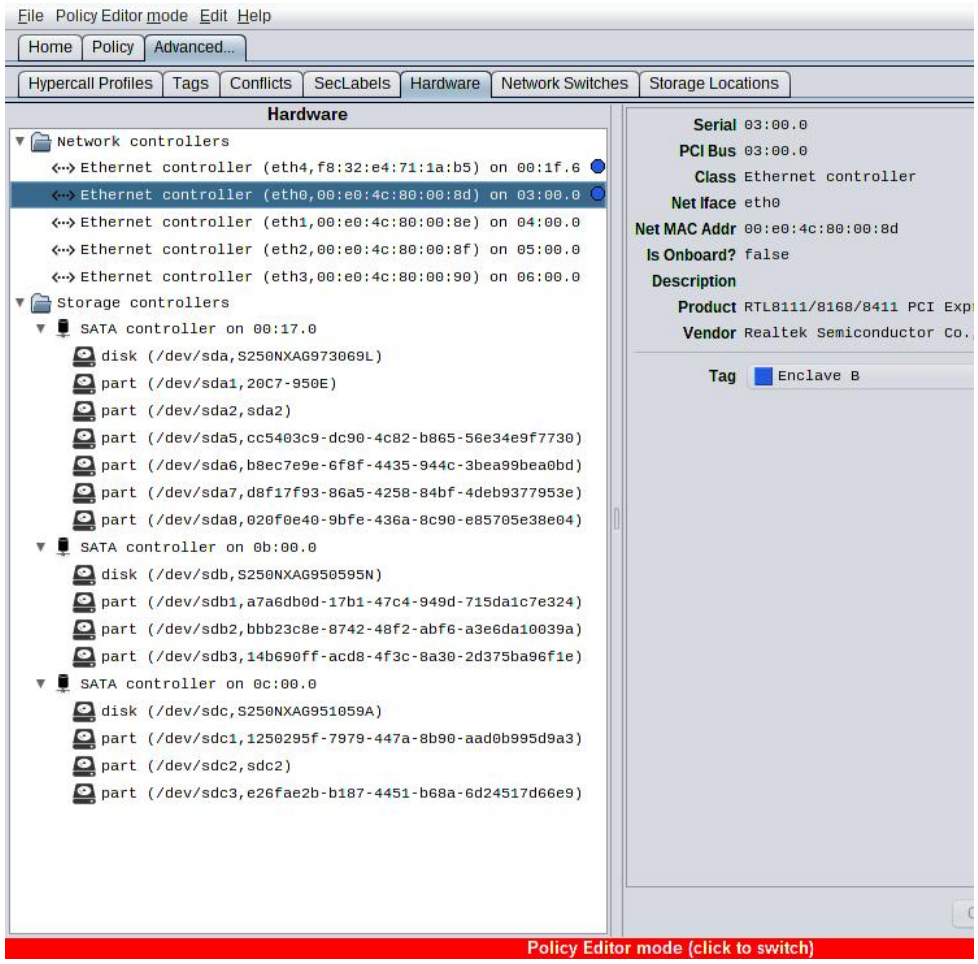

Simple configuration of resources

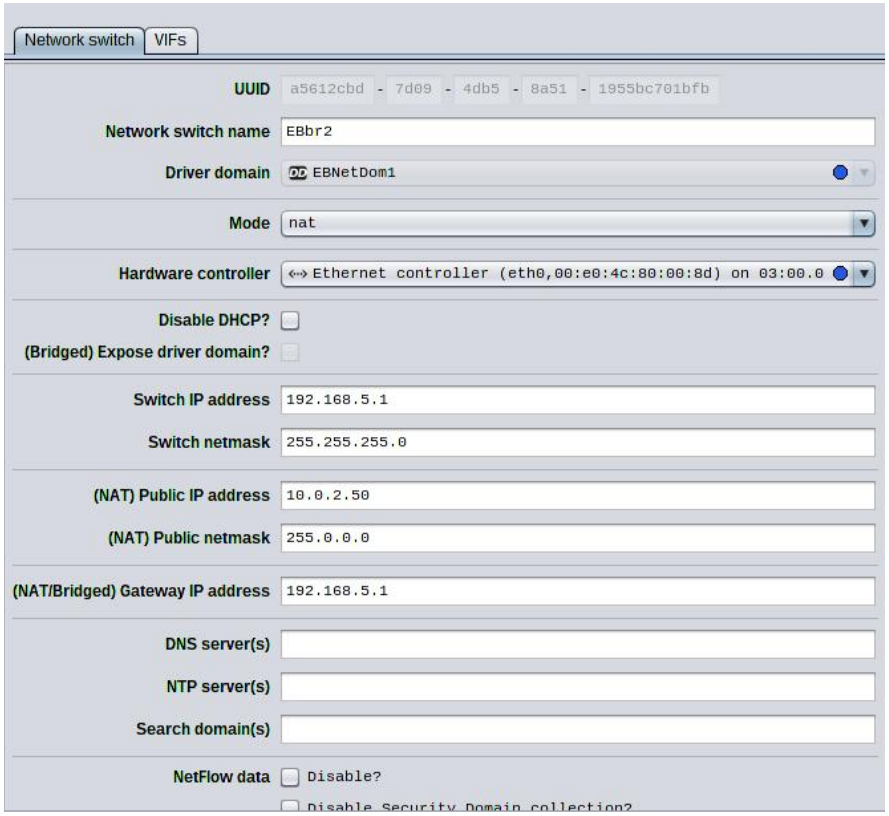

Xenon is More Than A Secure Hypervisor (cont'd)

Create pre-configured Service VMs

Assign resources to unprivileged VMs

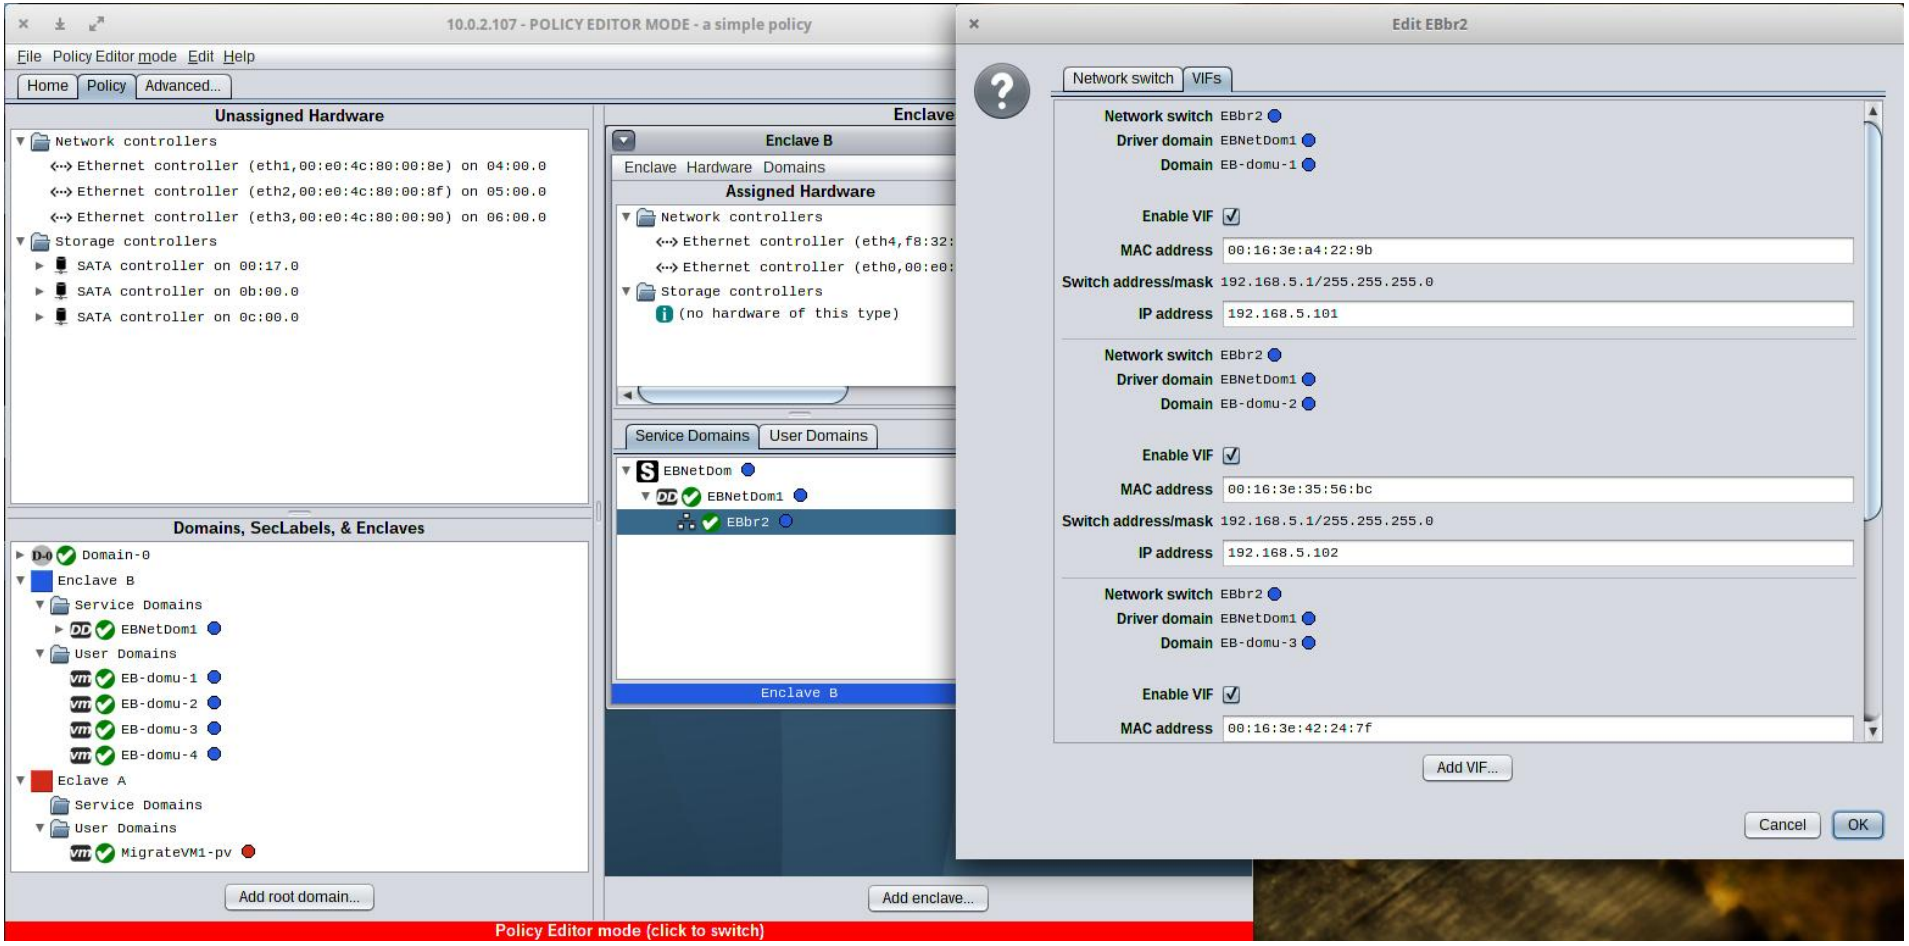

U.S. NAVAL RESEARCH LABORATORY

A Typical Xenon Host Configuration

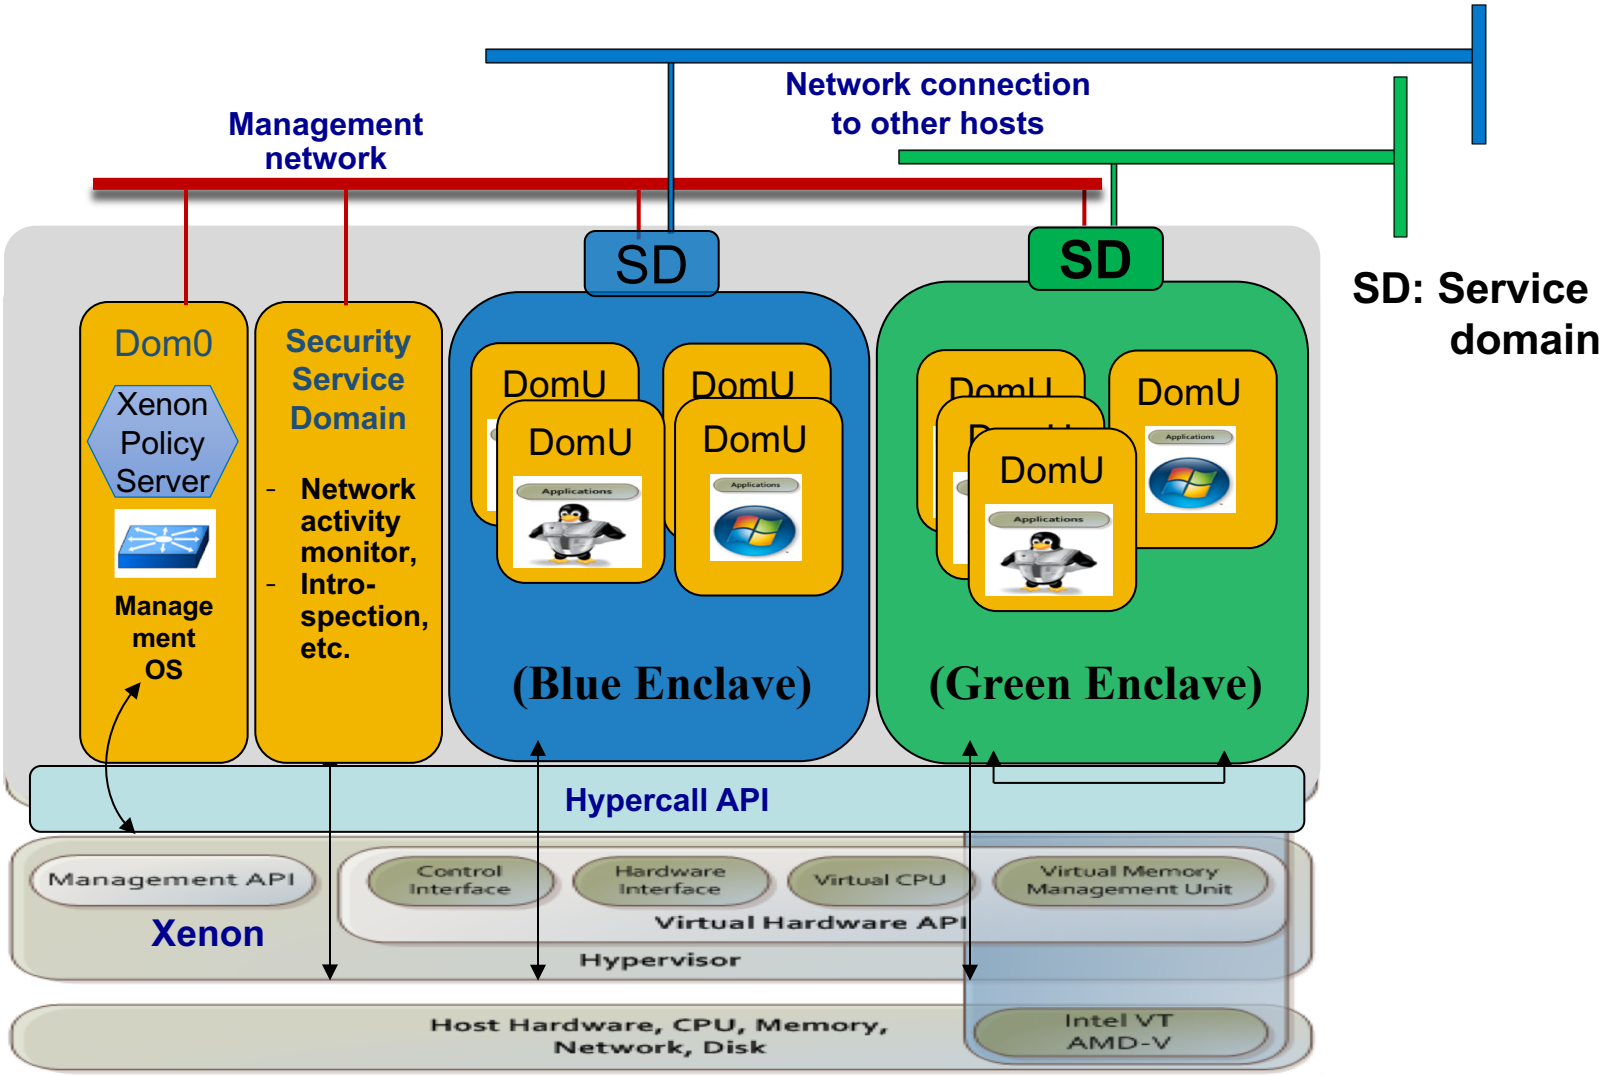

U.S. NAVAL RESEARCH LABORATORY

Xenon Enterprise

Manage multiple Xenon hosts

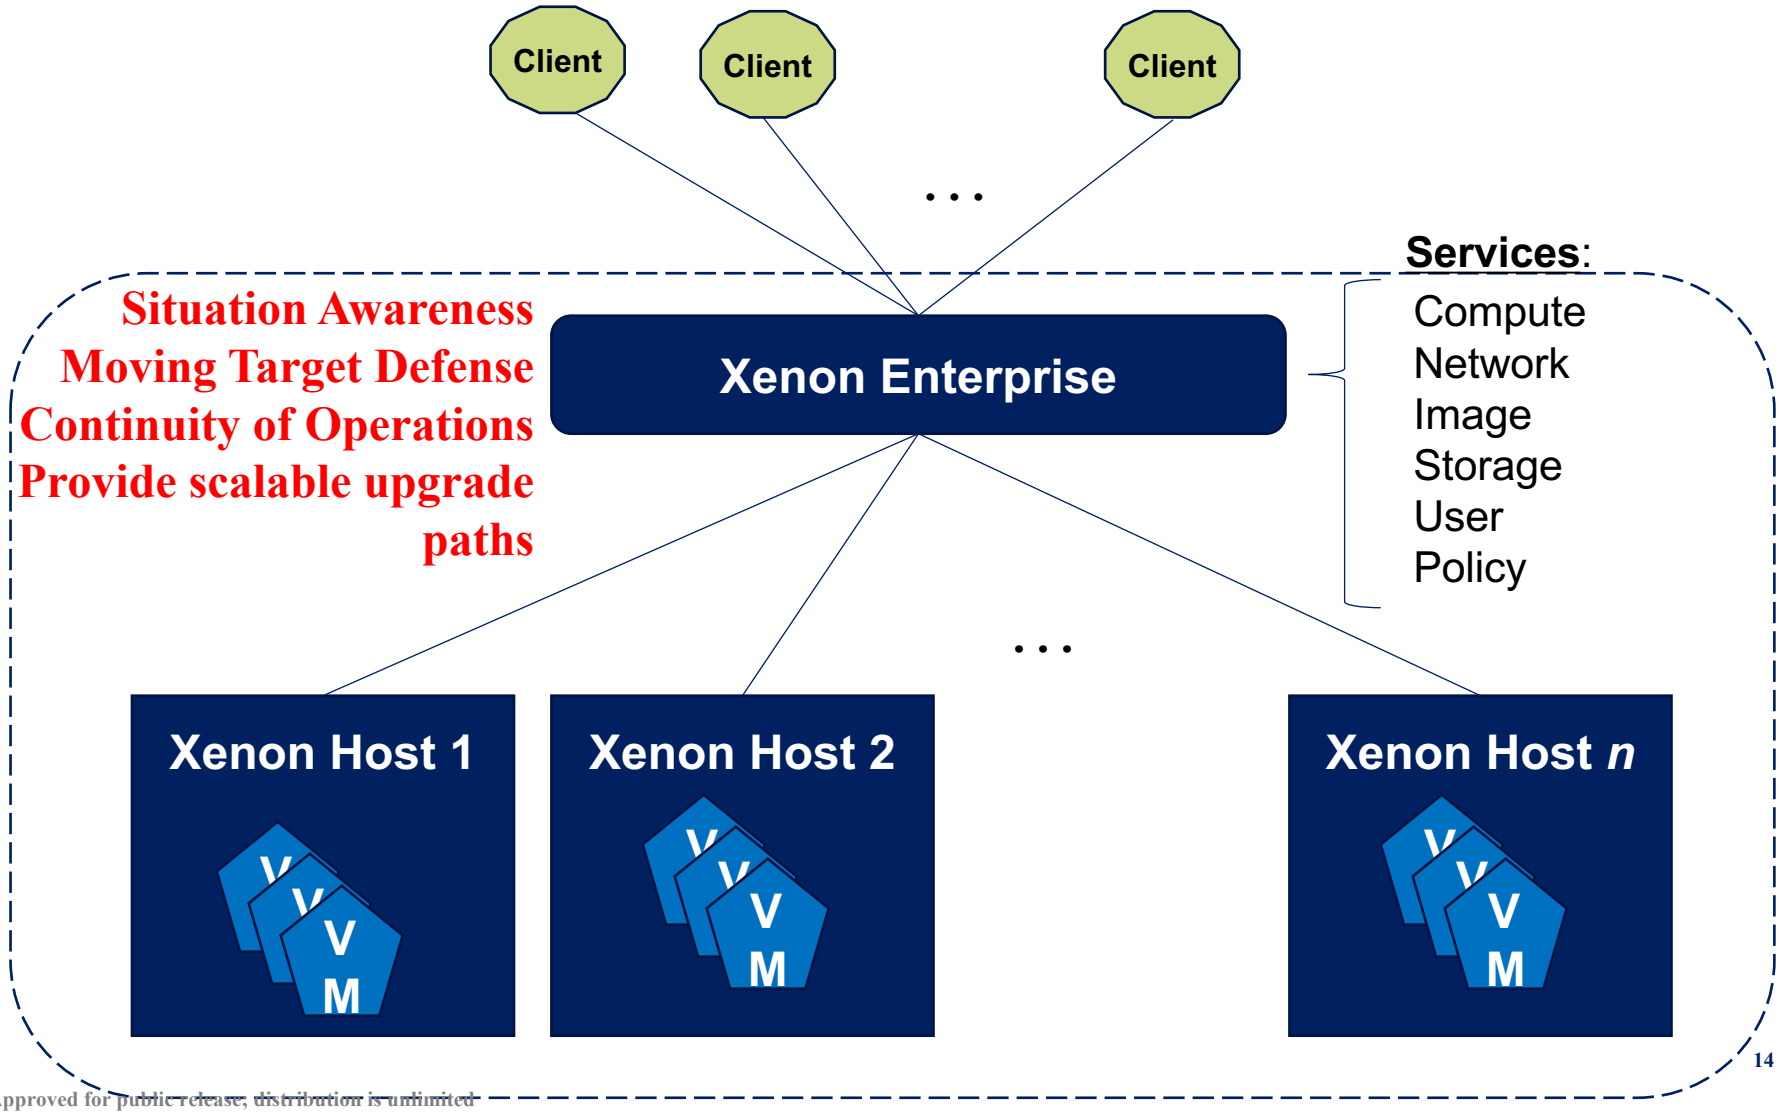

U.S. NAVAL RESEARCH LABORATORY

Xenon Enterprise (cont'd)

Hosts & VMs **Enclaves, security labels, and VMs**

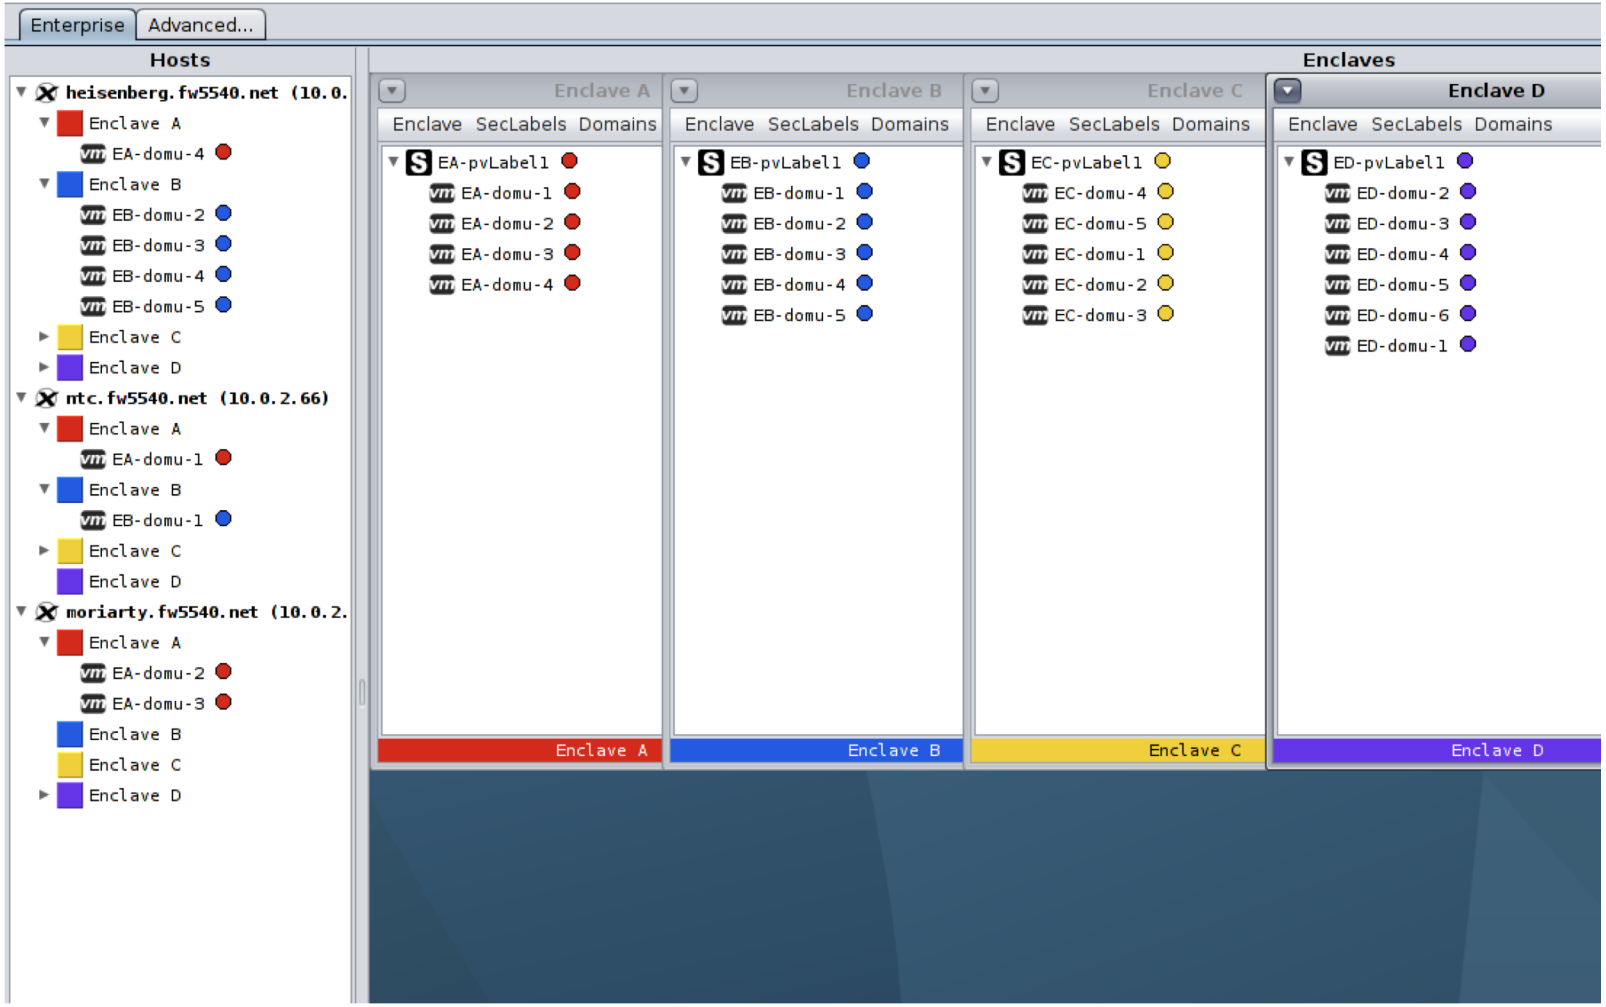

U.S. NAVAL RESEARCH LABORATORY

Xenon Enterprise Provides Templates for Application-centric Cloud

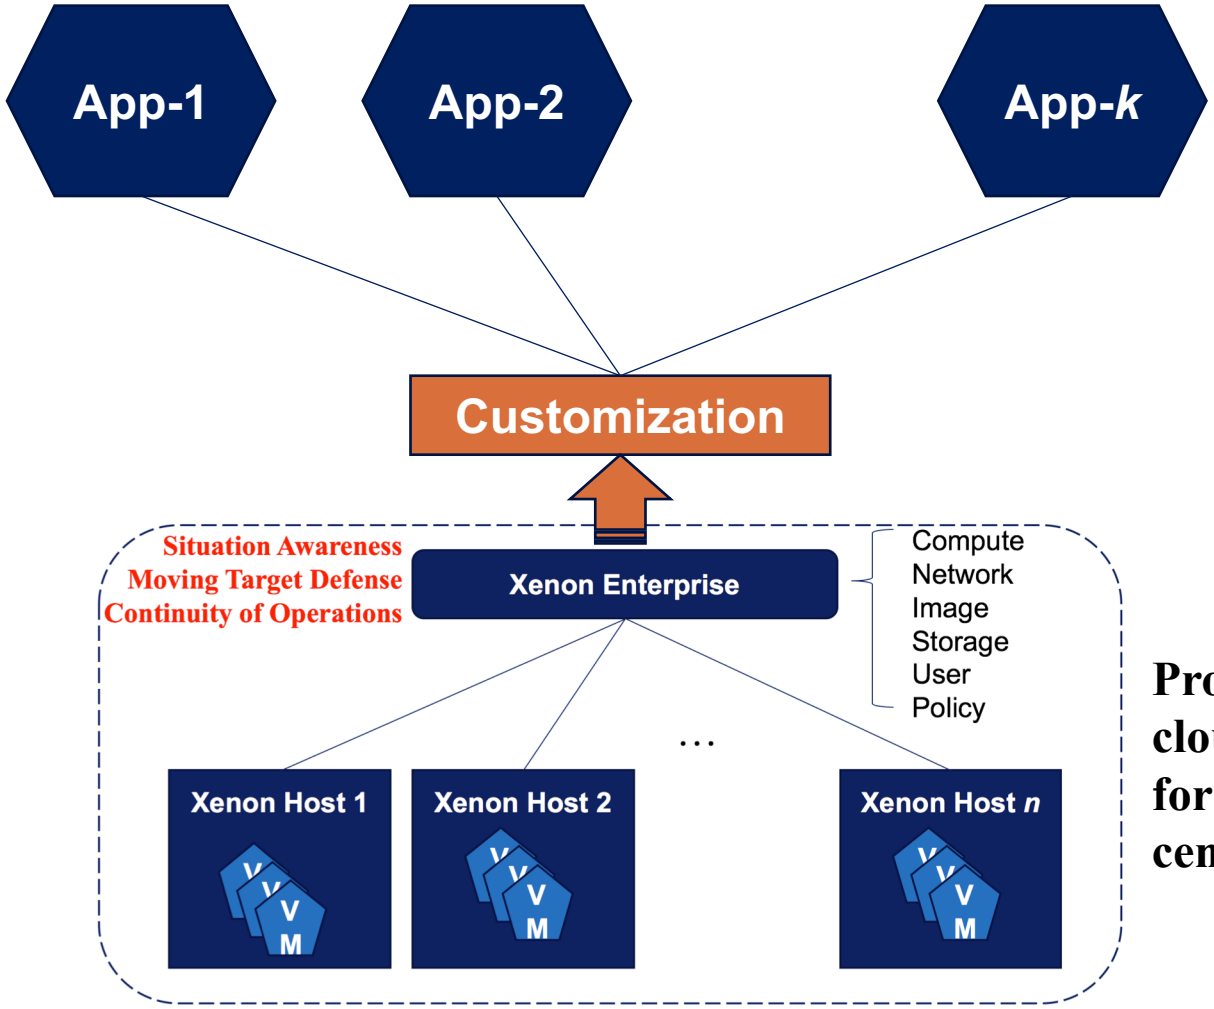

Provides a secure cloud template for application centric cloud

U.S. NAVAL RESEARCH LABORATORY

Xenon Enterprise with Xenon's Enclaves Provides Secure Cloud Computing Platform for Multiple Applications

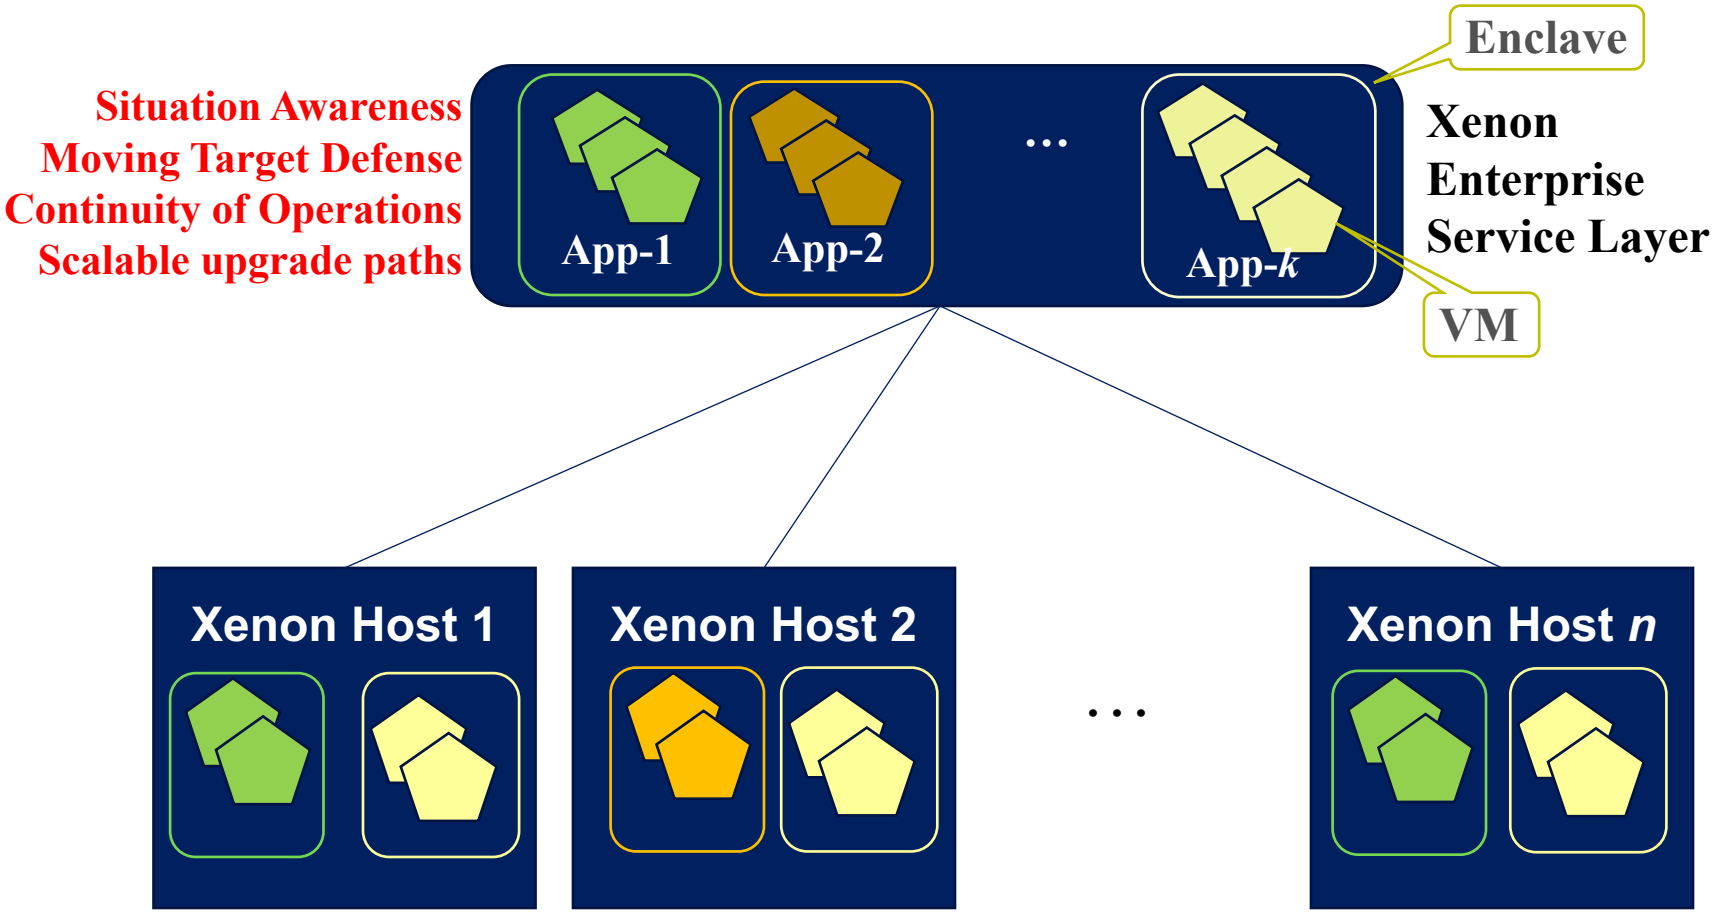

Security Implementation Details

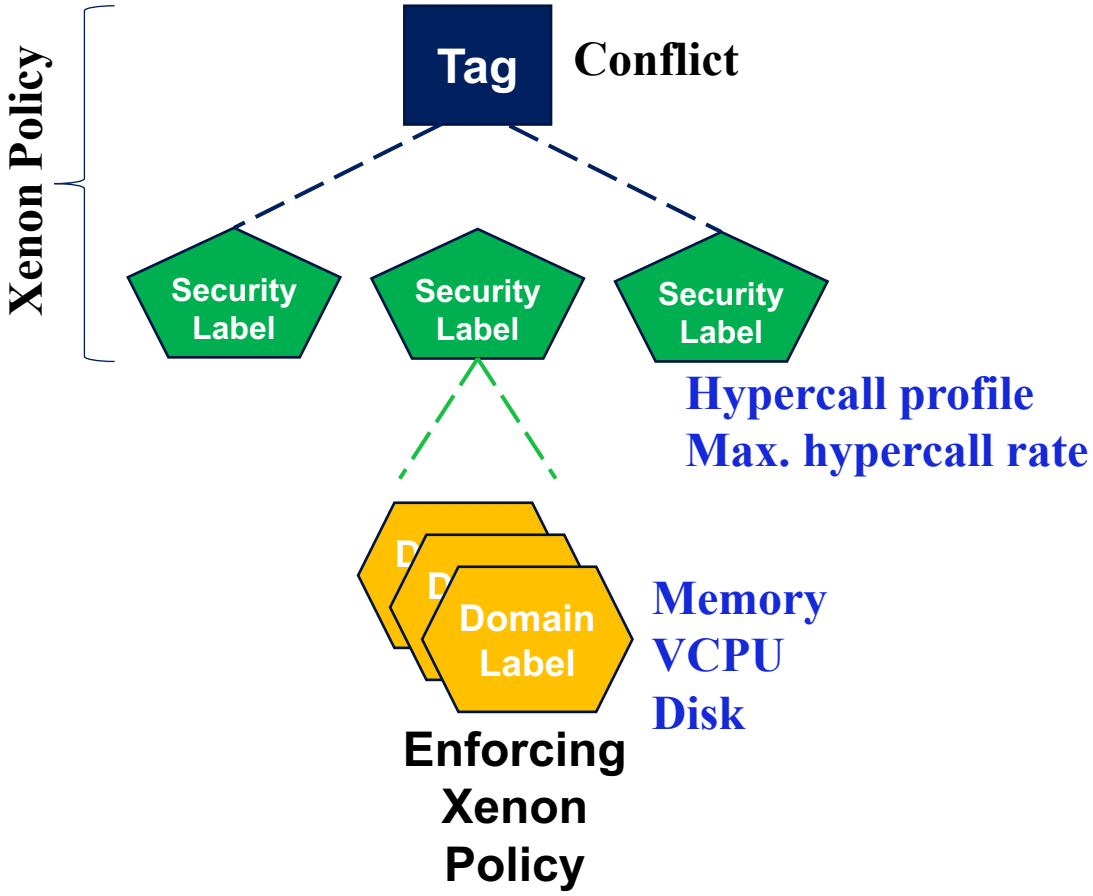

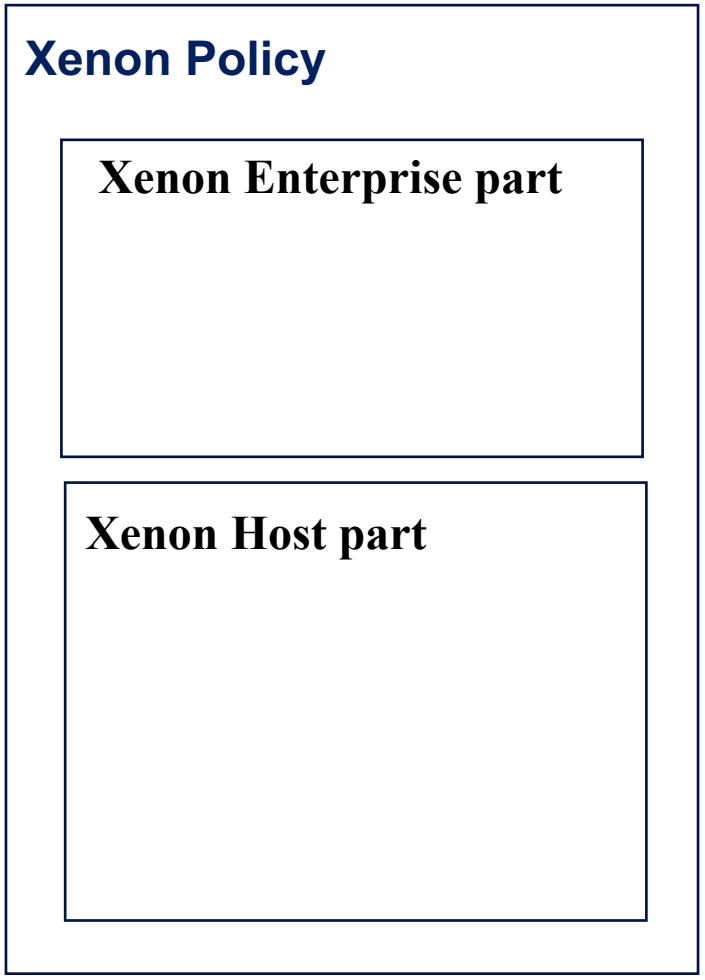

U.S. NAVAL RESEARCH LABORATORY

Security Implementation Details (cont'd)

Enterprise-wide policy

Per-host configuration (including policy)

Policy + configuration distributed as signed XML document

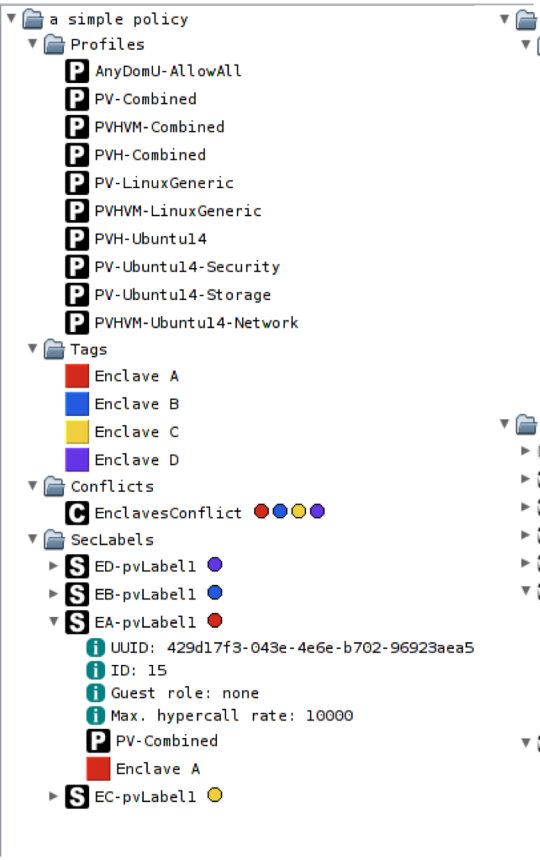

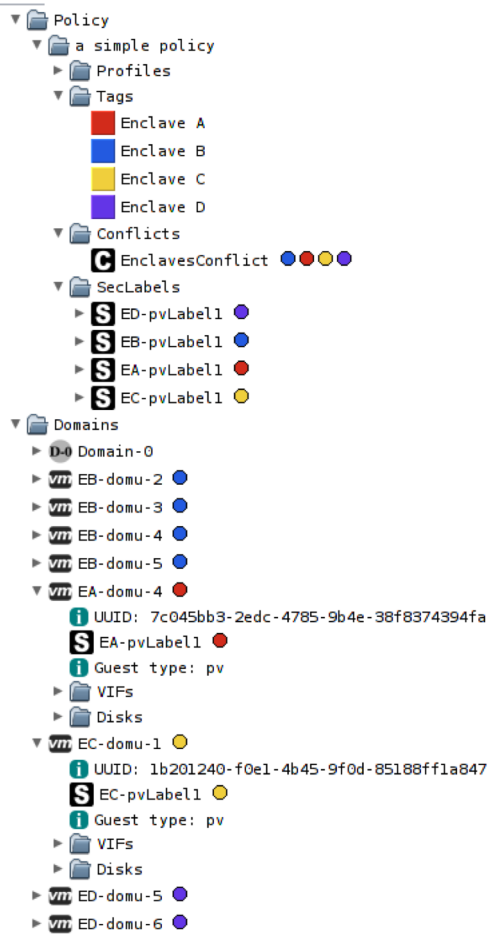

Questions?