

Scenario-based Distributed Virtualization Management Architecture for Multi-Host Environments

F. Galán (Telefónica I+D), D. Fernández (UPM), M. Ferrer (UPM), Fco. J. Martín (UPM)

DMTF 2nd Workshop on Systems and Virtualization Management (SVM 2008)
Munich, 21-22 October 2008



Index

01 Introduction

02 Design

- General architecture
- Deployment controller
- Operation

03 Practical Results

04 Alignment with DMTF's Standards

05 Summary

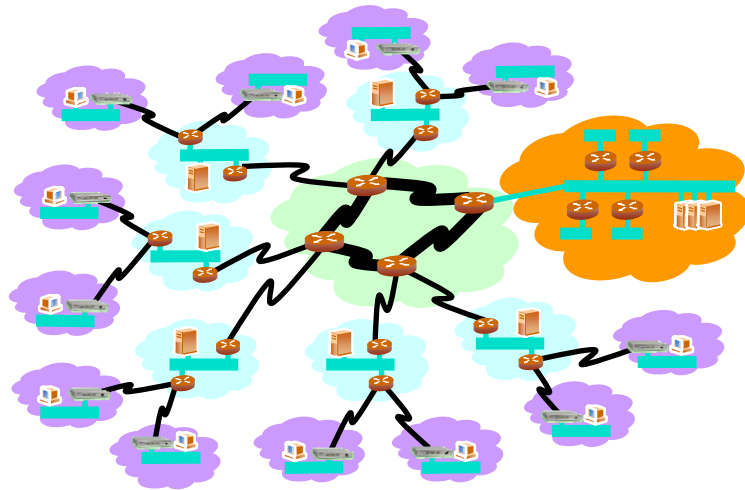
01 Introduction (1/3)

Scenario-based virtualization management principles

- Scenario-based virtualization management tools
 - Aimed at processing **scenario specifications** in order to be deployed in physical infrastructure (hosts) and provide further management (e.g. monitoring, execute command sequences, undeploy, etc.).
 - Scenario specification can be defined as a format representation (e.g., XML) of a **set of virtual machines** along with their **interconnections** in a given topology.
- Some tools nowadays are scenario-based, focused on **experimentation environments**, such as research and educational testbeds
 - VNUML, NetKit, MLN
- Conventional virtualization management tools (e.g. VMware VirtualCenter) for production environments use not to be scenario-based (out of our scope)


01 Introduction (2/3)

Scenario-based management example: VNUML


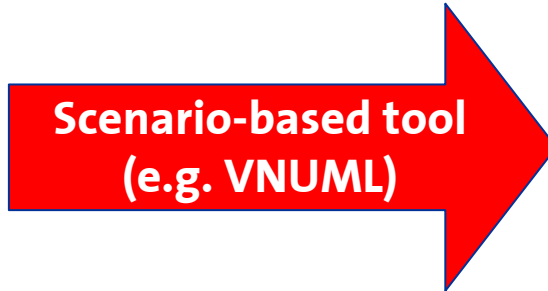


**Desired scenario
(designed by the user)**

**VMs and interconnection
virtual networks**



**Scenario
specification
(in XML)**



Physical host

01 Introduction (3/3)

Problem Statement

- Scenario-based virtualization tools use to consider just **one physical host** deployment
 - Scalability problems (e.g. a 1000 VM scenario)
- We have designed and implemented a distributed virtualization management tool named EDIV which purpose is to deploy scenarios in a **cluster of N physical hosts**.
 - Based on VNUML (in other words, VNUML scenarios can be used in EDIV without modifications)
 - Modular and extensible partition algorithm
 - A partnership project between Telefónica I+D and Universidad Politécnica de Madrid (Spain) which prototype results are publicly available (GPL).

Index

01 Introduction

02 **Design**

- General architecture
- Deployment controller
- Operation

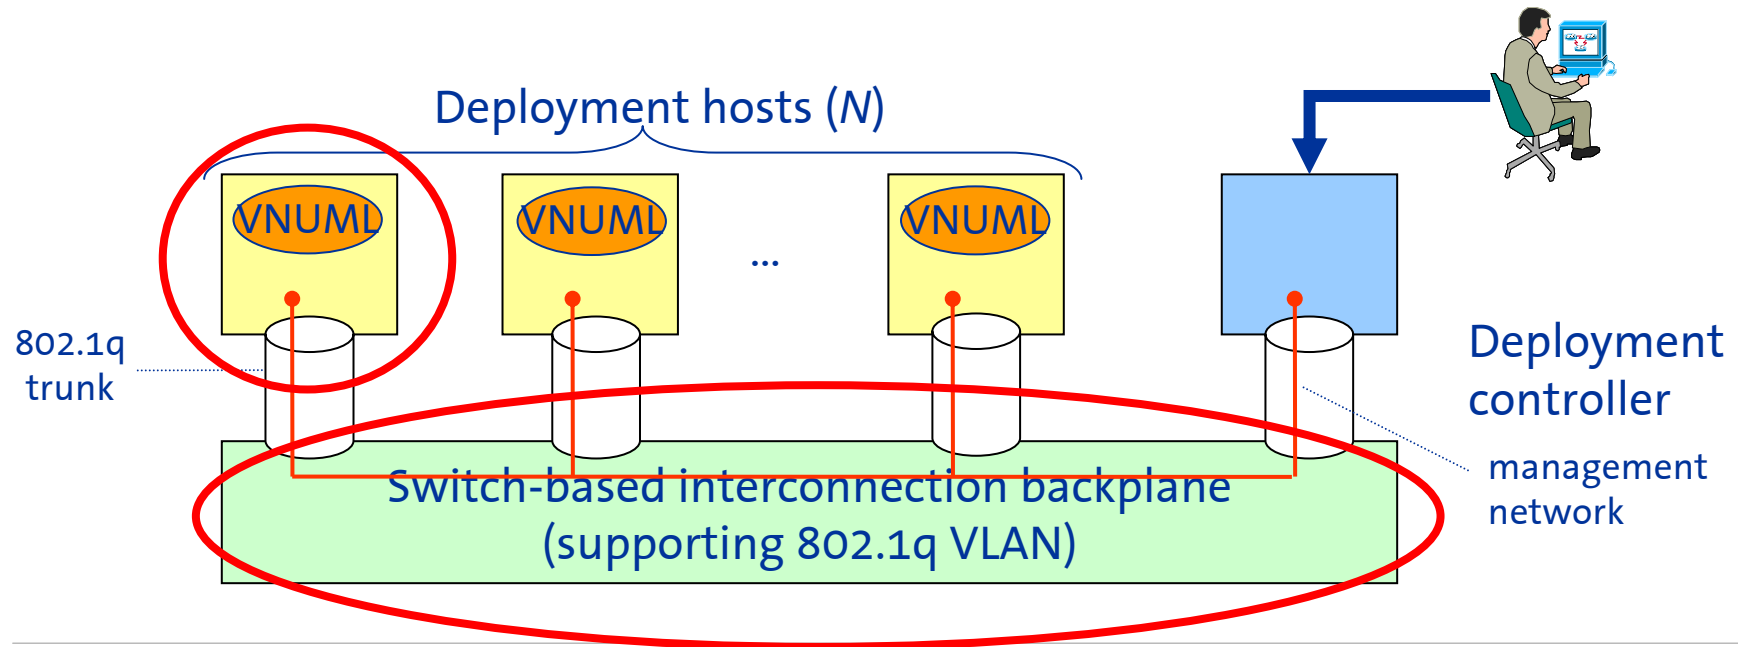
03 Practical Results

04 Alignment with DMTF's Standards

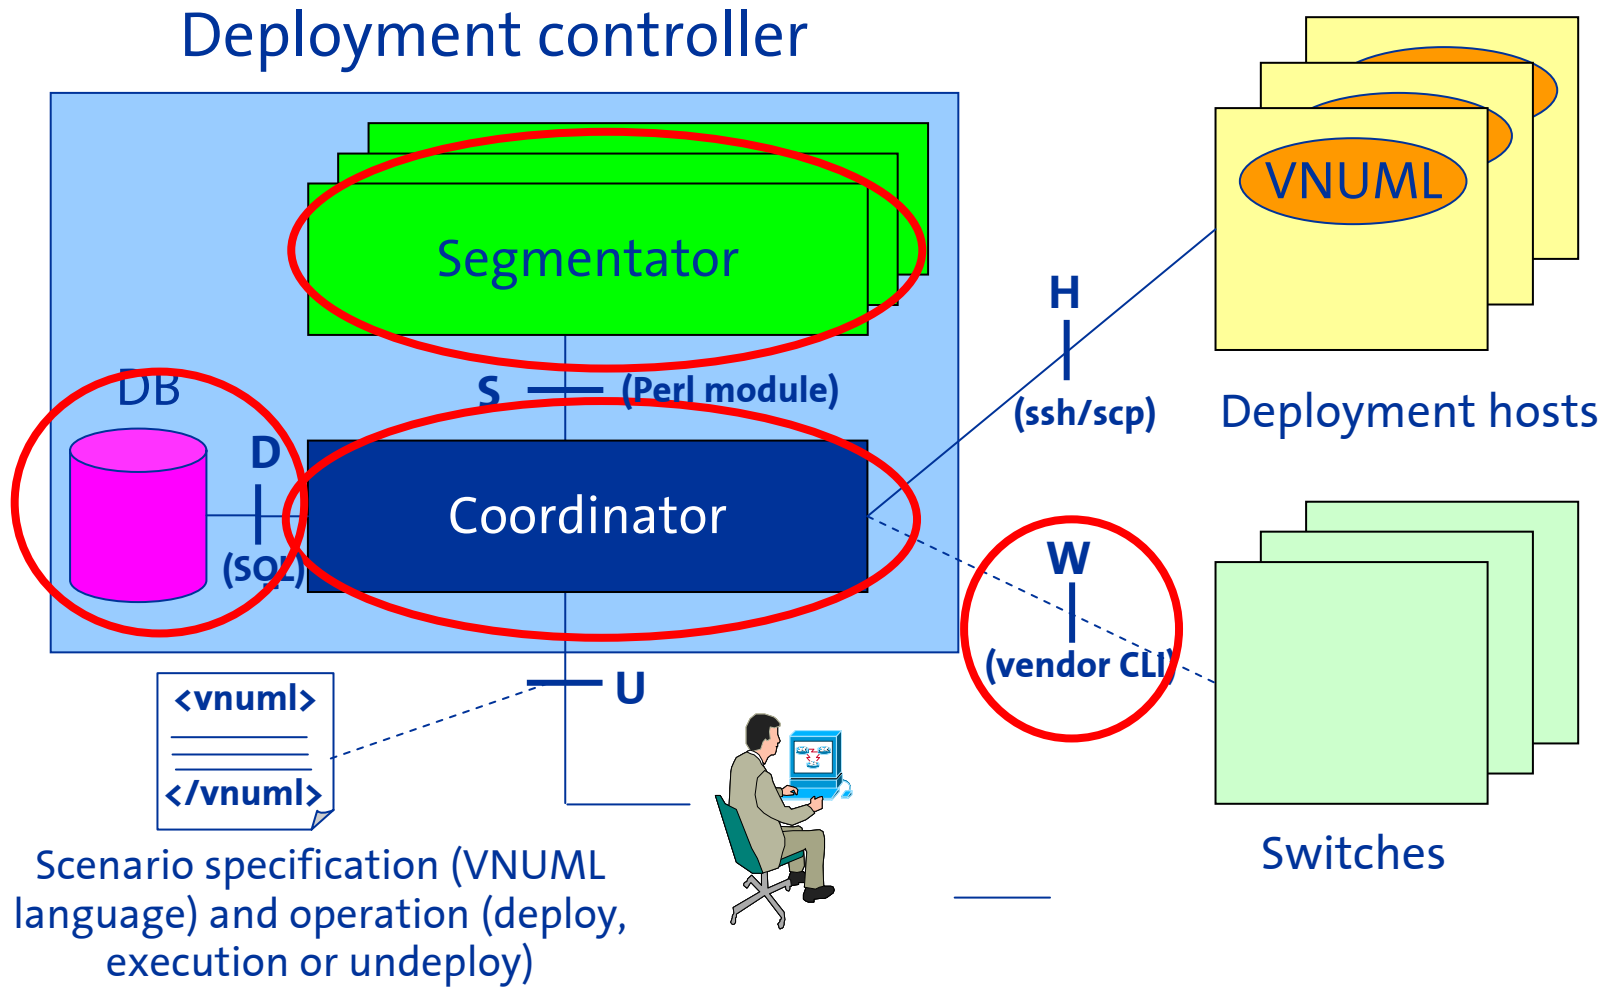
05 Summary

02 General Architecture

- Deployment hosts
 - GNU/Linux systems with VNUML installed along with other utilities (vconfig, sshd, etc.)
 - VNUML provide three operations: deploy, execution and undeploy
- Switch-based interconnection backplane
 - One or several Ethernet switches providing end-to-end 802.1q trunks
 - A management network (physical or VLAN) must be implemented

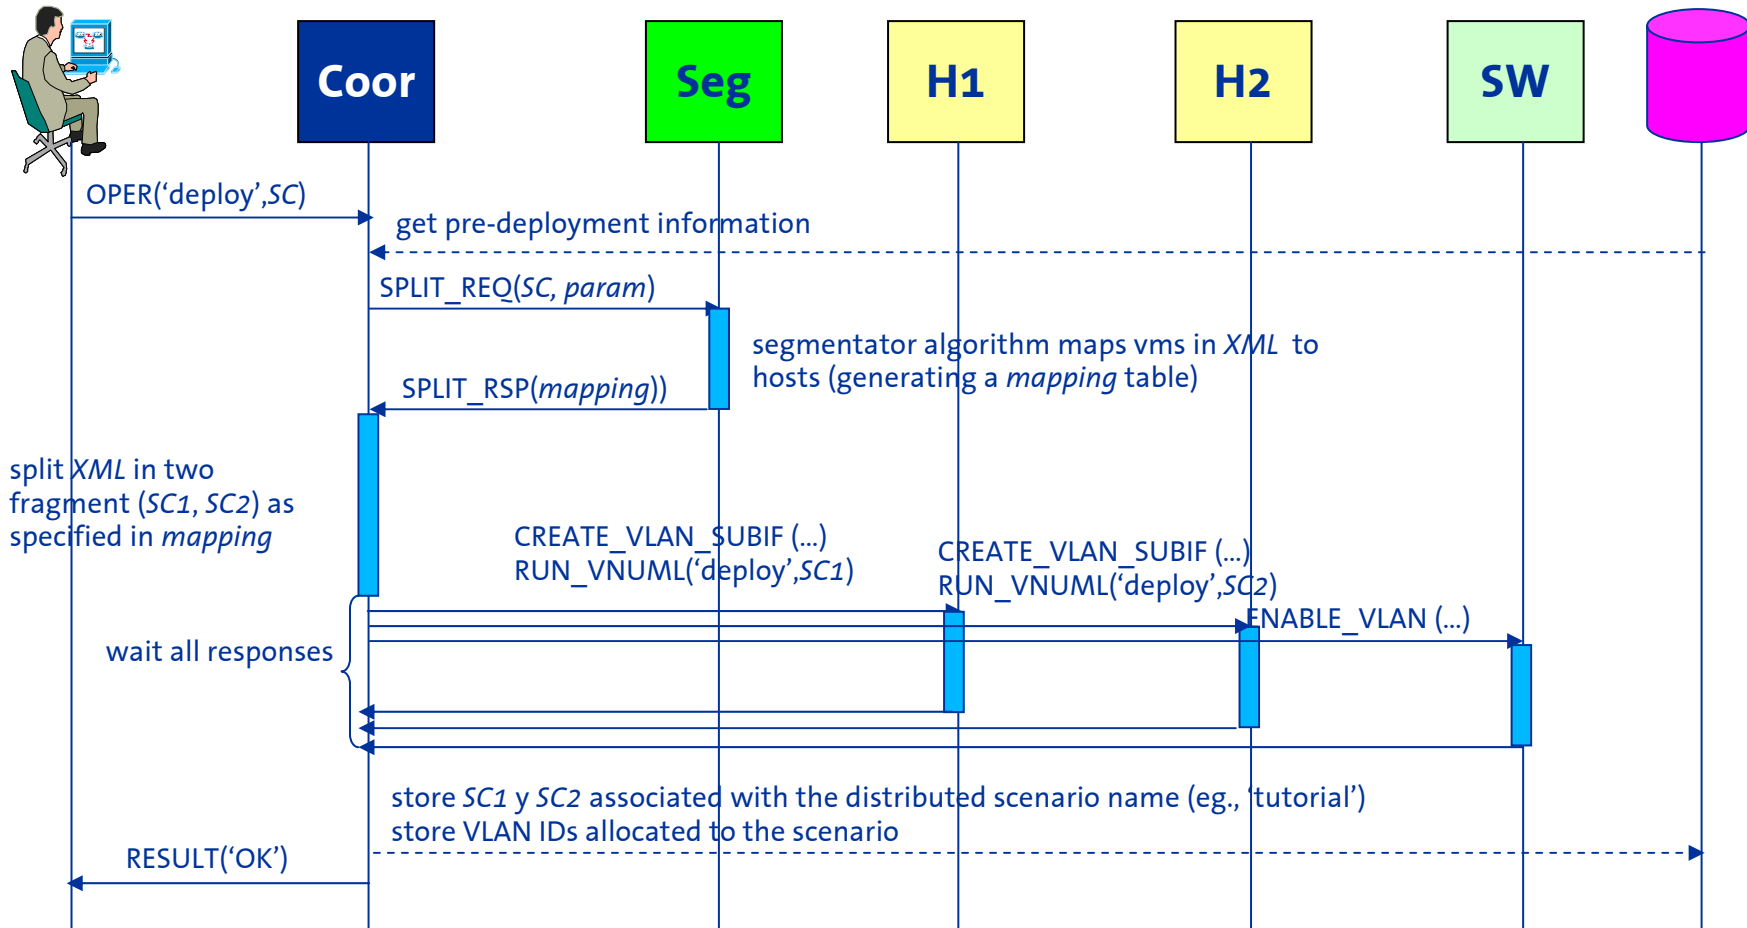


02 Deployment Controller



02 Operation

Deployment



Index

01 Introduction

02 Design

- General architecture
- Deployment controller
- Operation

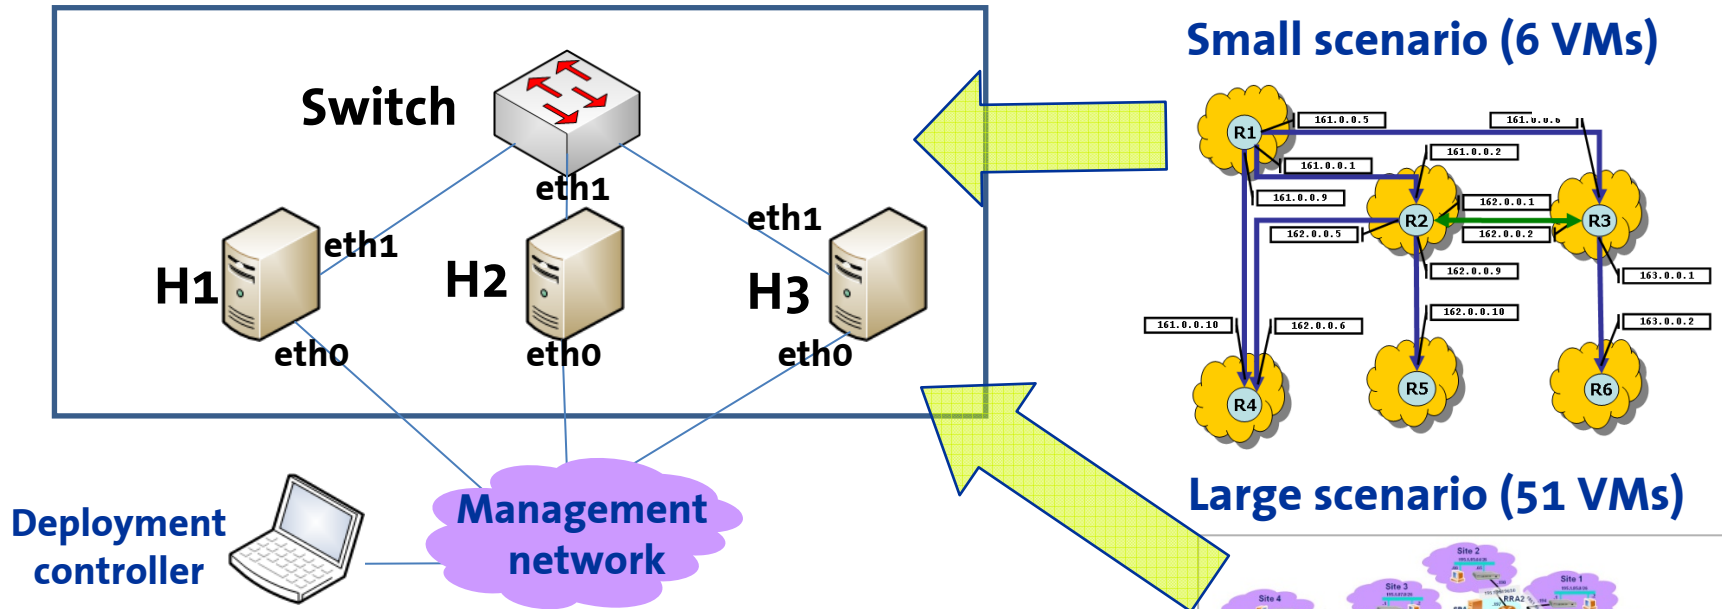
03 **Practical Results**

04 Alignment with DMTF's Standards

05 Summary

03 Practical Results

Experiment Setup



- EDIV tool implemented in Perl
- Physical management network
- Partition algorithms
 - Round robin
 - Weighted round robin (based on CPU load)
 - Deployment constraints (“VM A and B in the same physical host”)

03 Practical Results

Virtual machine distribution

	Deployment host loads			Round-Robin			Weighted Round Robin		
	L1	L2	L3	H1	H2	H3	H1	H2	H3
Small scenario (6 VMs)	L	L	L	2	2	2	2	2	2
	M	M	M	2	2	2	2	2	2
	H	L	L	2	2	2	1	3	2
	H	M	M	2	2	2	2	2	2
Large scenario (51 VMs)	L	L	L	17	17	17	17	17	17
	M	M	M	17	17	17	17	17	17
	H	L	L	17	17	17	7	22	22
	H	M	M	17	17	17	13	19	19

L (low) ⇒ 0% load

M (medium) ⇒ 50% load

H (high) ⇒ 100% load

03 Practical Results

Time comparison

- How much does it take to deploy each scenario?

- With VNUML

- Small scenario: **42 s**

- Large scenario: **282 s**

- With EDIV (low mean load)

- Small scenario: **17 s**

- Large scenario: **130 s**

- With EDIV (high mean load)

- Small scenario: **43 s**

- Large scenario: **135 s**

EDIV can achieve a **50-60%** saving regarding the mono-host case with VNUML

Index

01 Introduction

02 Design

- General architecture
- Deployment controller
- Operation

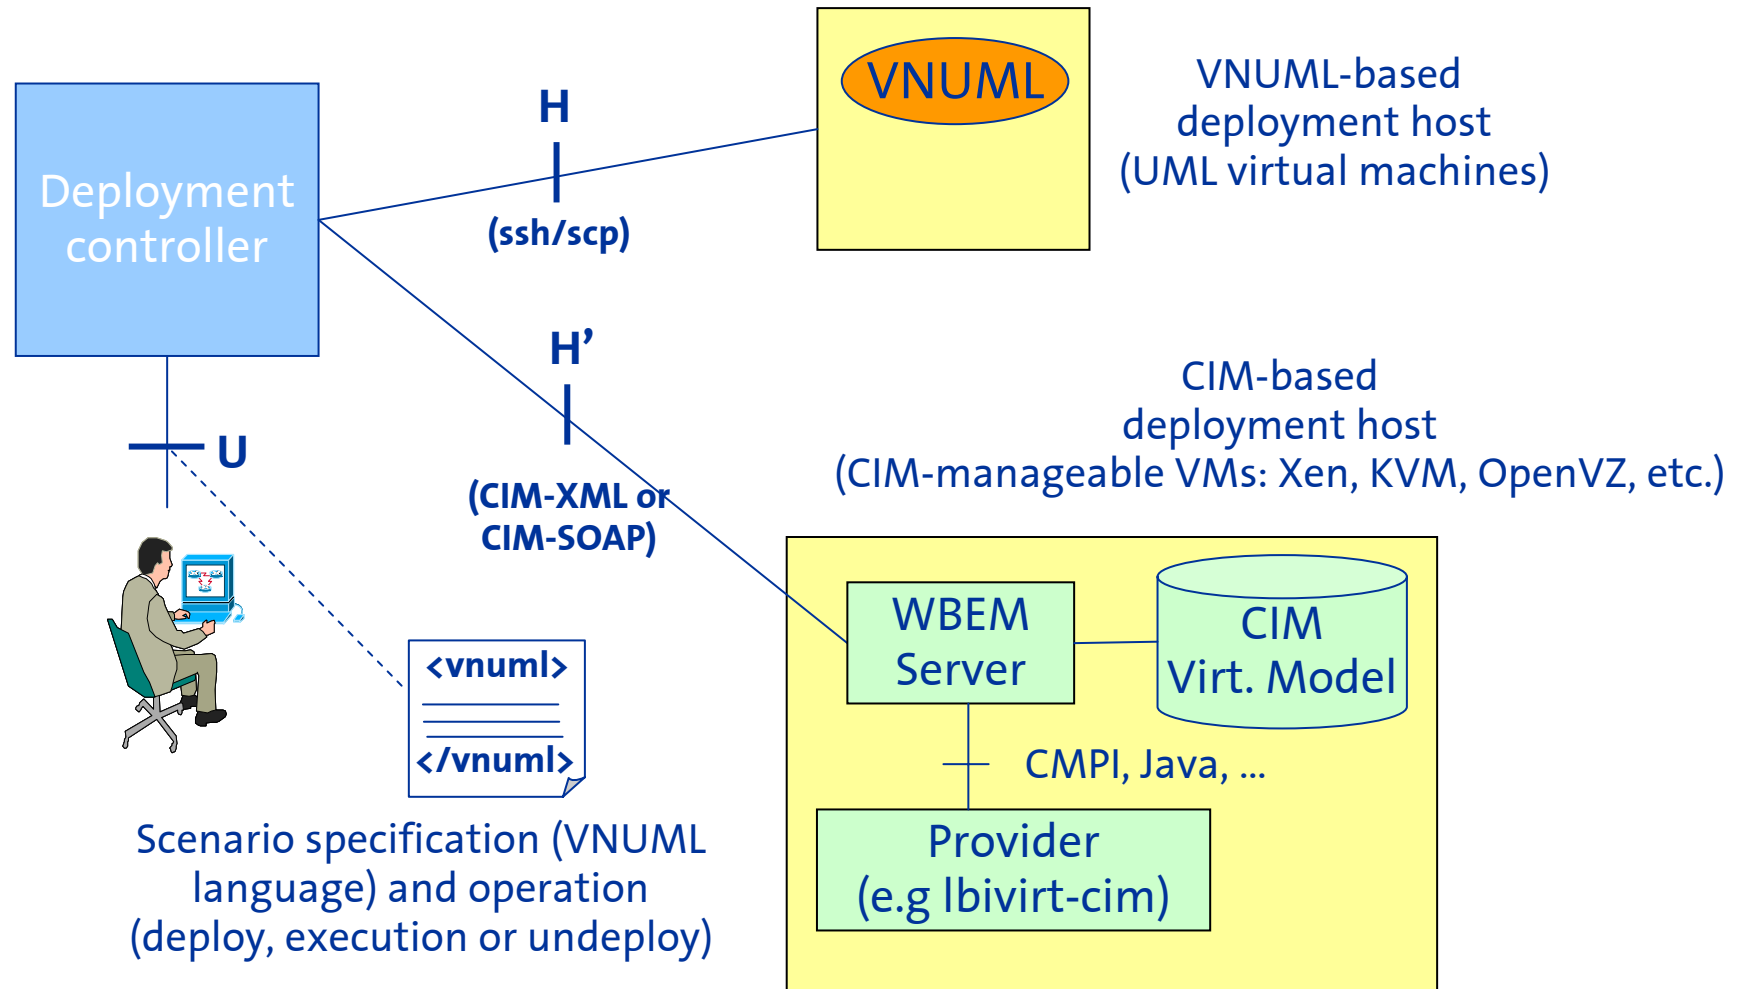
03 Practical Results

04 **Alignment with DMTF's Standards**

05 **Summary**

04 DMTF's Standards Alignment

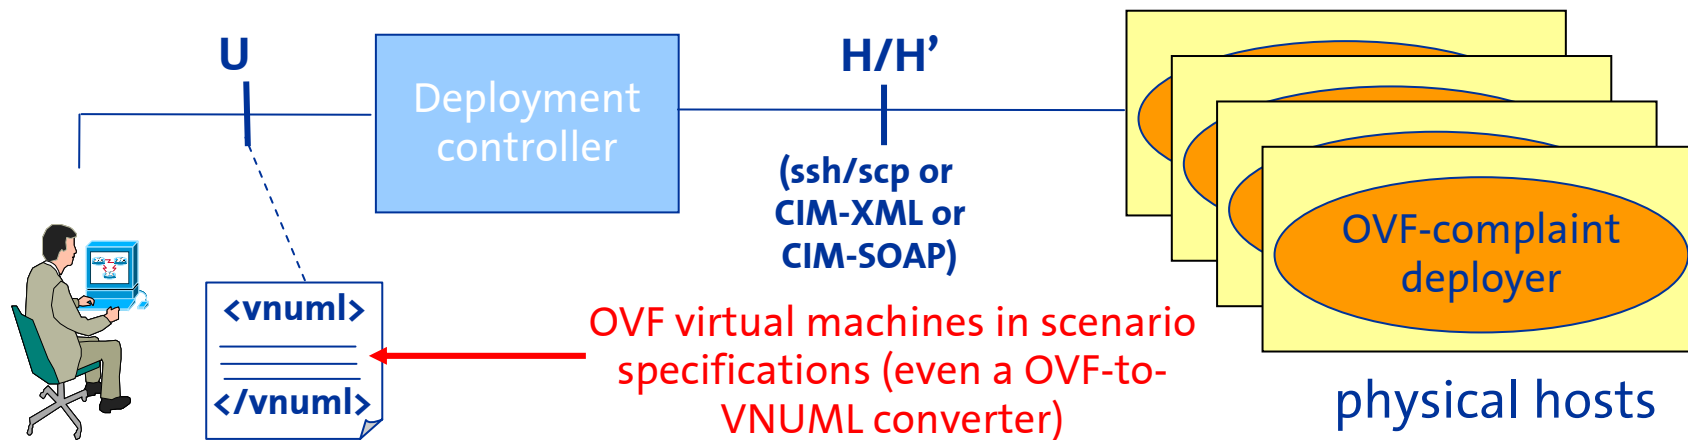
CIM-based virtual machine management



04 DMTF's Standards Alignment

OVF based virtual machines in EDIV scenarios

- Open Virtualization Format (OVF)
 - Recently published DSP0243 1.0.0d (September 2008)
- Why introduce OVF in EDIV?
 - To use virtual appliances highly optimized for specific purposes in EDIV scenarios (e.g., firewalls, dynamic routing stacks, etc.)
 - To smooth the migration from development/pre-production environments to production
- How to do it?



05 Summary

■ Main conclusions

- EDIV architecture solves the problem of scenario-based virtualization management in distributed multi-host environment (as checked by ours experiments)
 - Scalability
 - Dramatic reduction (50-60% saving) in deployment time
- Simple and open API to develop new partition algorithms
- Transparency to VNUML users (the specification language is the same)

■ Work in progress

- Wide-area deployment environment (instead of a local cluster)
 - This is being working out in PASITO (a distributed experimentation platform within RedIris, the Spanish National Research and Education Network)

■ Future working lines

- Alignment with DMTF's standards
- Virtual machine filesystem management

Telefonica

**EDIV tool and documentation is
freely available at
<http://www.dit.upm.es/vnuml>**