



# CernVM WebAPI

CernVM Users Workshop 2015

Ioannis Charalampidis, PH/SFT




# Agenda

- 1 Origins of CernVM WebAPI
- 2 Components of CernVM WebAPI
- 3 Running **multiple** projects in the same VM
- 4 Hands-on CernVM WebAPI



1

# CernVM WebAPI Origins



**2011** : LHC@Home 2.0 – BOINC Project  
First time where virtualization is used

1

# CernVM WebAPI Origins



Conceived as an additional **burden** to some volunteers

1

# CernVM WebAPI Origins

Want  
to volunteer

Mostly with **good** computer skills

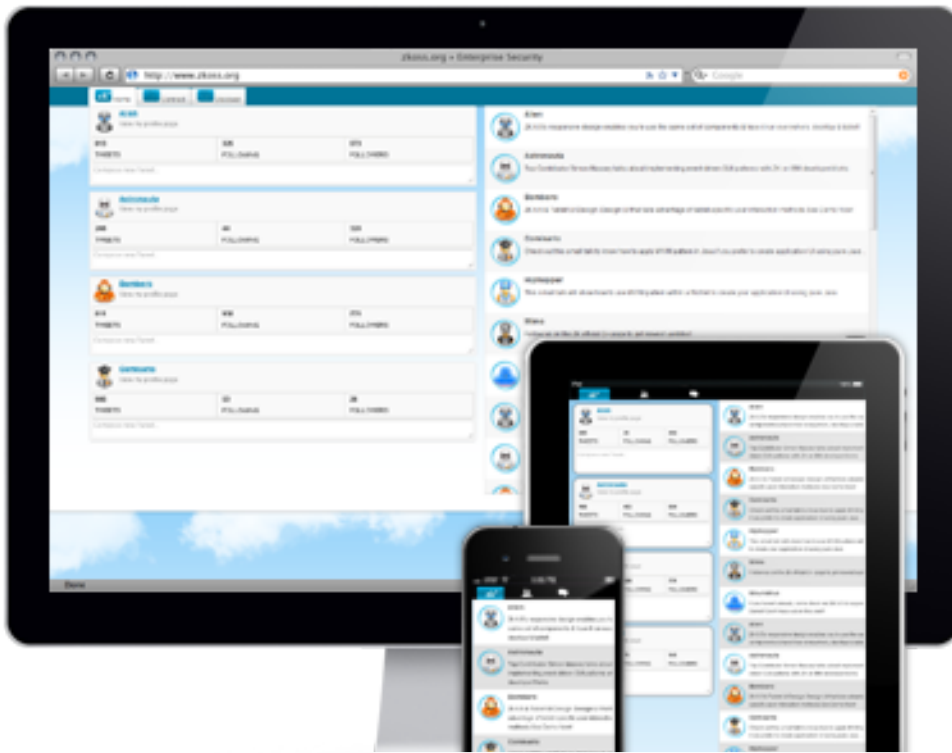
BOINC  
Users

Can we also target the **rest** of the audience?



1


# CernVM WebAPI Origins



**Web apps are simple, run everywhere and require no manual installations**

1

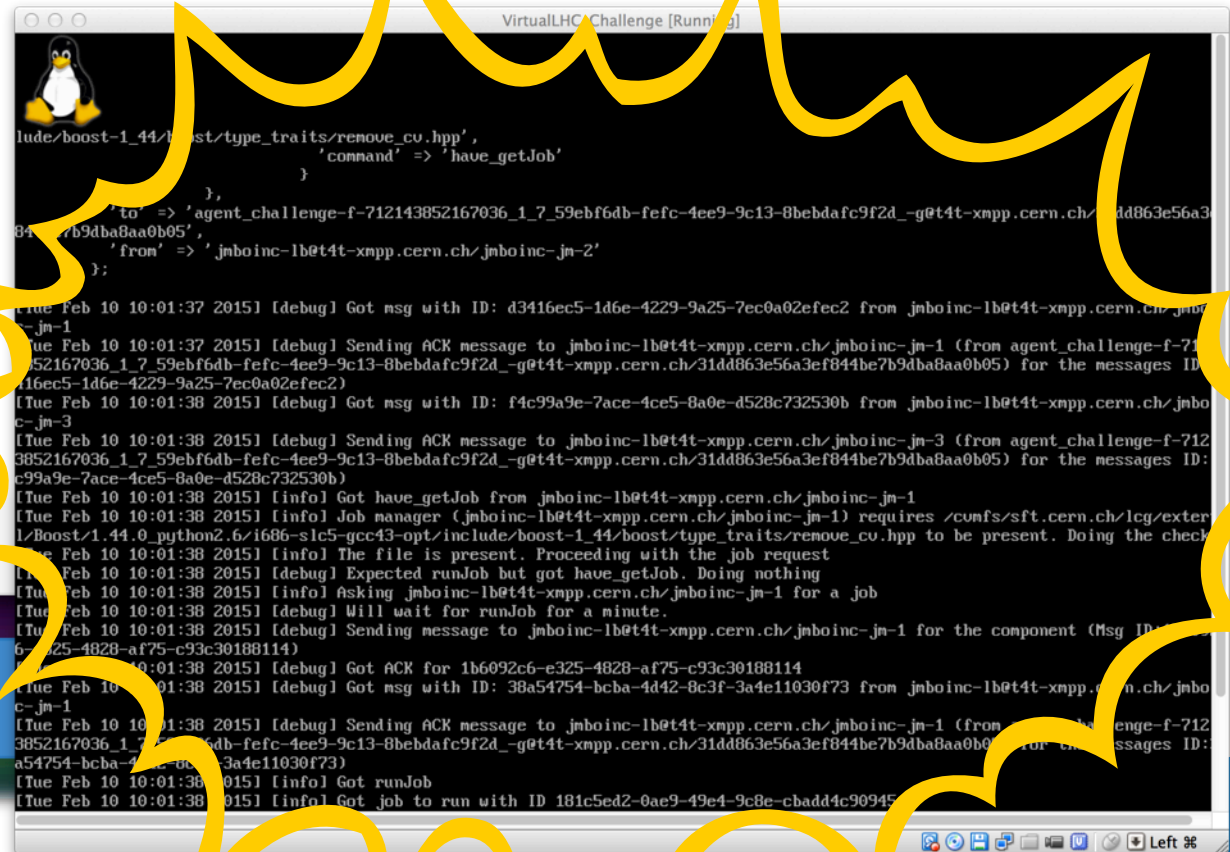
# CernVM WebAPI Origins



**Virtualization** is widespread: Code it once,  
deploy it everywhere

1

# CernVM WebAPI Origins



```
VirtualLHC/Challenge [Running]
include/boost-1.44.0/include/boost/type_traits/remove_cv.hpp',
    'command' => 'have_getJob'
    },
    'to' => 'agent_challenge-f-712143852167036_1_7_59ebf6db-fefc-4ee9-9c13-8bebda9cf2d_g@t4t-xmpp.cern.ch/31dd863e56a3ef844be7b9d8aa0b05',
    'from' => 'jmboinc-lb@t4t-xmpp.cern.ch/jmboinc-jm-2'
  });
[Tue Feb 10 10:01:37 2015] [debug] Got msg with ID: d3416ec5-1d6e-4229-9a25-7ec0a02efec2 from jmboinc-lb@t4t-xmpp.cern.ch/jmboinc-jm-1
[Tue Feb 10 10:01:37 2015] [debug] Sending ACK message to jmboinc-lb@t4t-xmpp.cern.ch/jmboinc-jm-1 (from agent_challenge-f-712143852167036_1_7_59ebf6db-fefc-4ee9-9c13-8bebda9cf2d_g@t4t-xmpp.cern.ch/31dd863e56a3ef844be7b9d8aa0b05) for the messages ID: d3416ec5-1d6e-4229-9a25-7ec0a02efec2)
[Tue Feb 10 10:01:38 2015] [debug] Got msg with ID: f4c99a9e-7ace-4ce5-8a0e-d528c732530b from jmboinc-lb@t4t-xmpp.cern.ch/jmboinc-jm-3
[Tue Feb 10 10:01:38 2015] [debug] Sending ACK message to jmboinc-lb@t4t-xmpp.cern.ch/jmboinc-jm-3 (from agent_challenge-f-712143852167036_1_7_59ebf6db-fefc-4ee9-9c13-8bebda9cf2d_g@t4t-xmpp.cern.ch/31dd863e56a3ef844be7b9d8aa0b05) for the messages ID: f4c99a9e-7ace-4ce5-8a0e-d528c732530b)
[Tue Feb 10 10:01:38 2015] [info] Got have_getJob from jmboinc-lb@t4t-xmpp.cern.ch/jmboinc-jm-1
[Tue Feb 10 10:01:38 2015] [info] Job manager (jmboinc-lb@t4t-xmpp.cern.ch/jmboinc-jm-1) requires /cunfs/sft.cern.ch/lcg/external/Boost/1.44.0/python2.6/i686-slc5-gcc43-opt/include/boost-1.44/boost/type_traits/remove_cv.hpp to be present. Doing the check
[Tue Feb 10 10:01:38 2015] [info] The file is present. Proceeding with the job request
[Tue Feb 10 10:01:38 2015] [debug] Expected runJob but got have_getJob. Doing nothing
[Tue Feb 10 10:01:38 2015] [info] Asking jmboinc-lb@t4t-xmpp.cern.ch/jmboinc-jm-1 for a job
[Tue Feb 10 10:01:38 2015] [debug] Will wait for runJob for a minute.
[Tue Feb 10 10:01:38 2015] [debug] Sending message to jmboinc-lb@t4t-xmpp.cern.ch/jmboinc-jm-1 for the component (Msg ID: 6-7a25-4828-af75-c93c30188114)
[Tue Feb 10 10:01:38 2015] [debug] Got ACK for 1b6092c6-e325-4828-af75-c93c30188114
[Tue Feb 10 10:01:38 2015] [debug] Got msg with ID: 38a54754-bcba-4d42-8c3f-3a4e11030f73 from jmboinc-lb@t4t-xmpp.cern.ch/jmboinc-jm-1
[Tue Feb 10 10:01:38 2015] [debug] Sending ACK message to jmboinc-lb@t4t-xmpp.cern.ch/jmboinc-jm-1 (from agent_challenge-f-712143852167036_1_7_59ebf6db-fefc-4ee9-9c13-8bebda9cf2d_g@t4t-xmpp.cern.ch/31dd863e56a3ef844be7b9d8aa0b05) for the messages ID: 38a54754-bcba-4d42-8c3f-3a4e11030f73)
[Tue Feb 10 10:01:38 2015] [info] Got runJob
[Tue Feb 10 10:01:38 2015] [info] Got job to run with ID 181c5ed2-0ae9-49e4-9c8e-cbadd4c90945
```

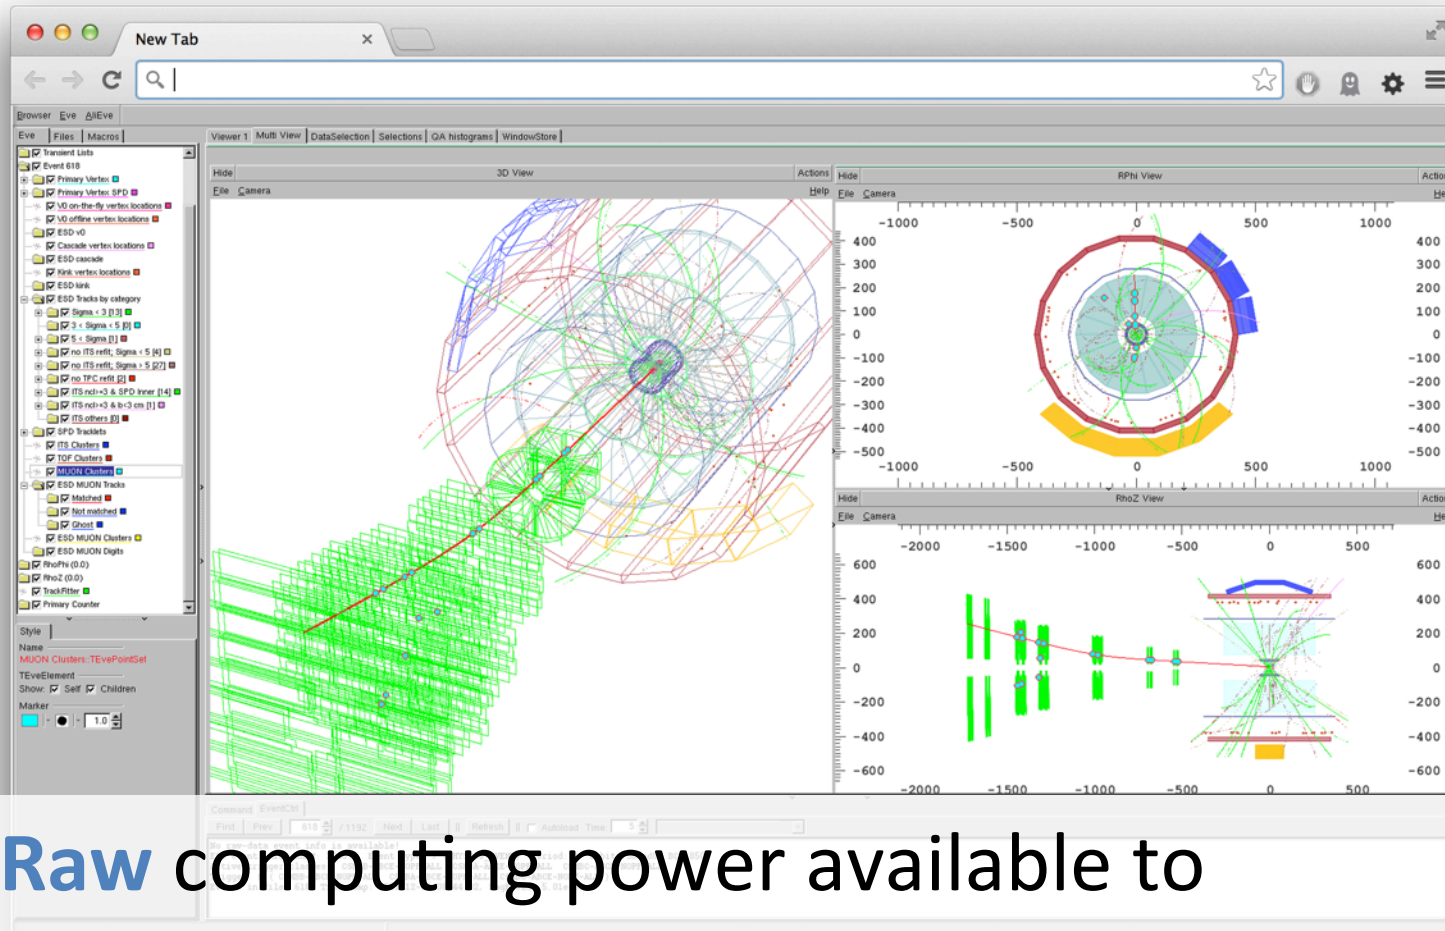
Start





1

# CernVM WebAPI Origins



Raw computing power available to web apps



1

# CernVM WebAPI



Start and Control Virtual Computing  
Resources through your web browser

1

# CernVM WebAPI

Does all the heavy-lifting for the user :

- ✓ Missing **hypervisor**? **Install it!**
- ✓ Improperly **configured** hypervisor? **Fix it!**
- ✓ Missing **resources**? **Download them!**



One-click solution to desktop **virtualization**



1

# CernVM WebAPI

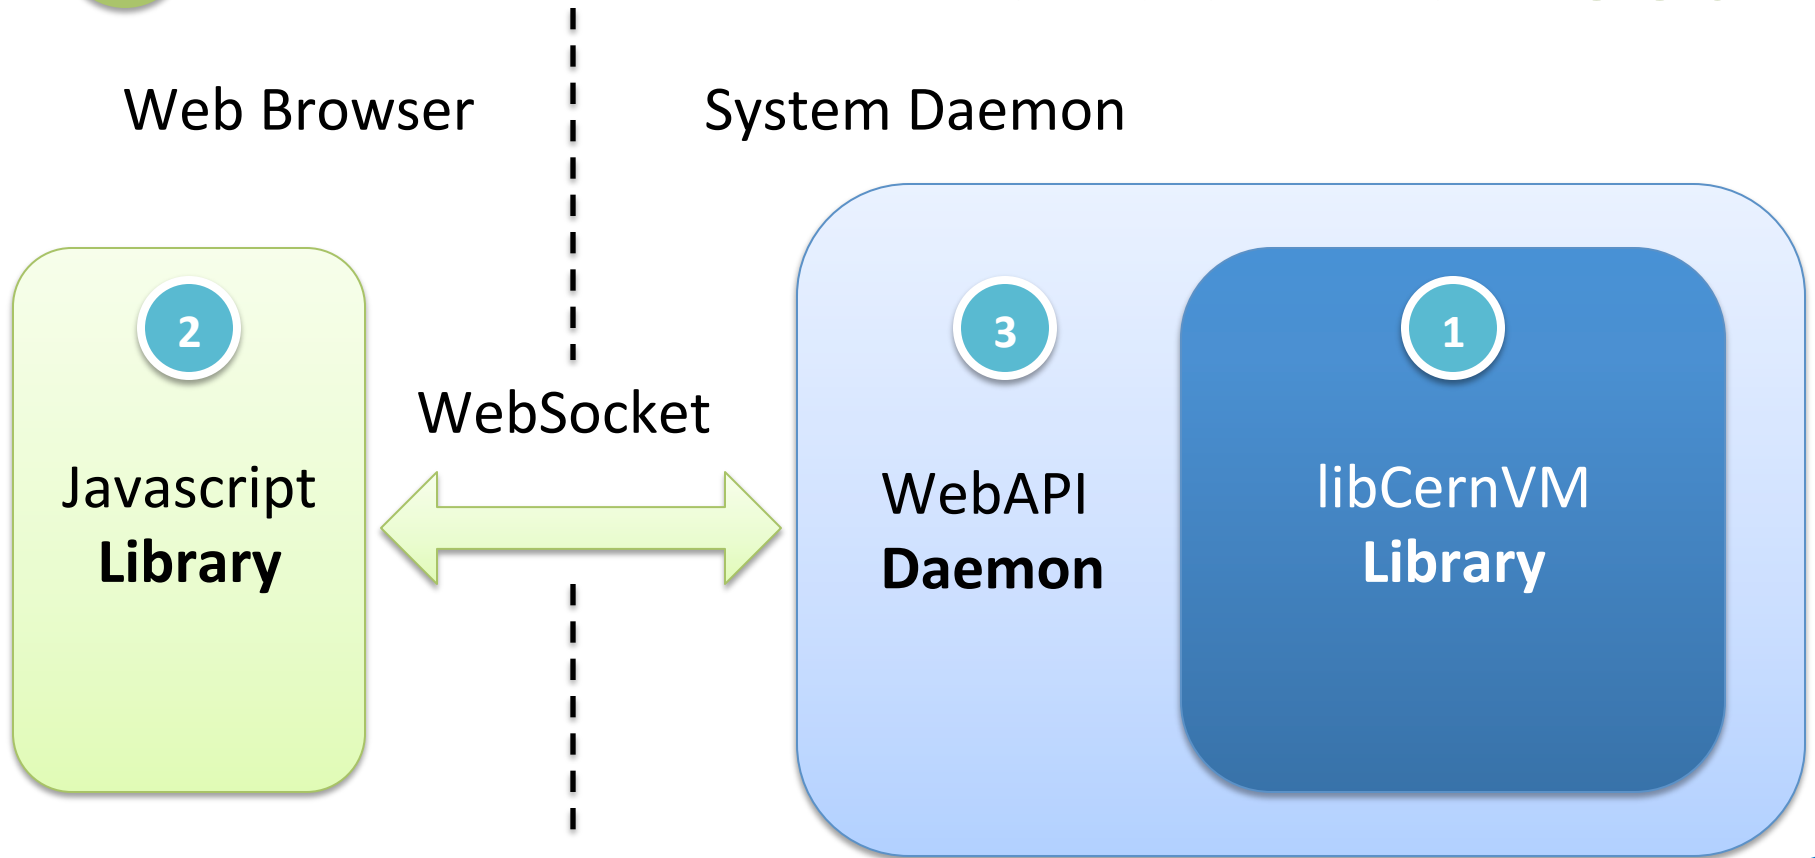
```
// Request API Access
CVM.startCVMWebAPI(function(plugin) {
  // Open Session
  plugin.requestSession("http://domain.com/vmcp?id=1", function(session) {
    // Start VM
    session.start();
  });
});
```



Javascript code for starting a VM

2

# Under the Hood



CernVM WebAPI components

<https://github.com/wavesoft/cernvm-webapi>



2

# Under the Hood - libCernVM

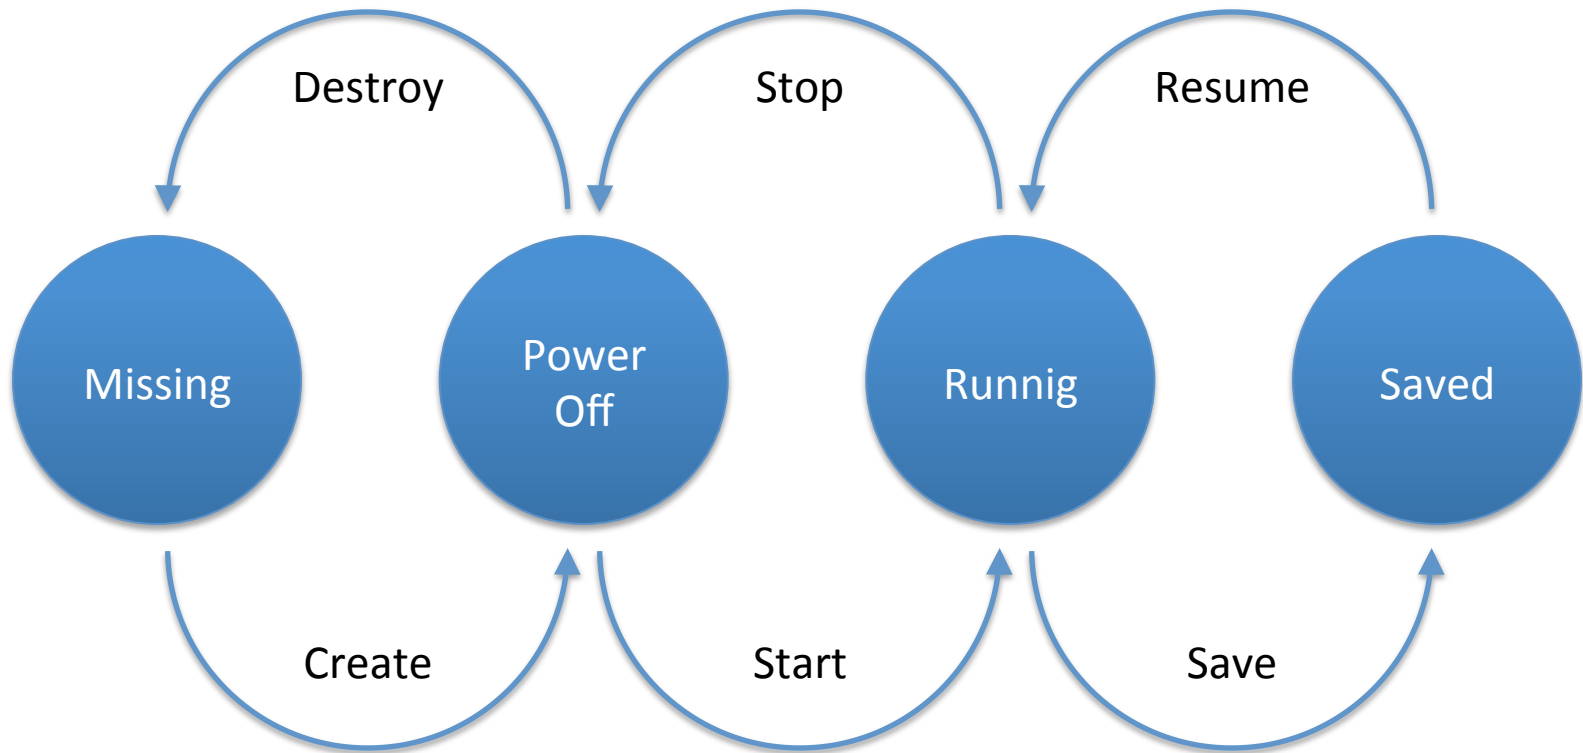


**libCernVM** interfaces with the Hypervisor

<https://github.com/wavesoft/libcernvm>

2

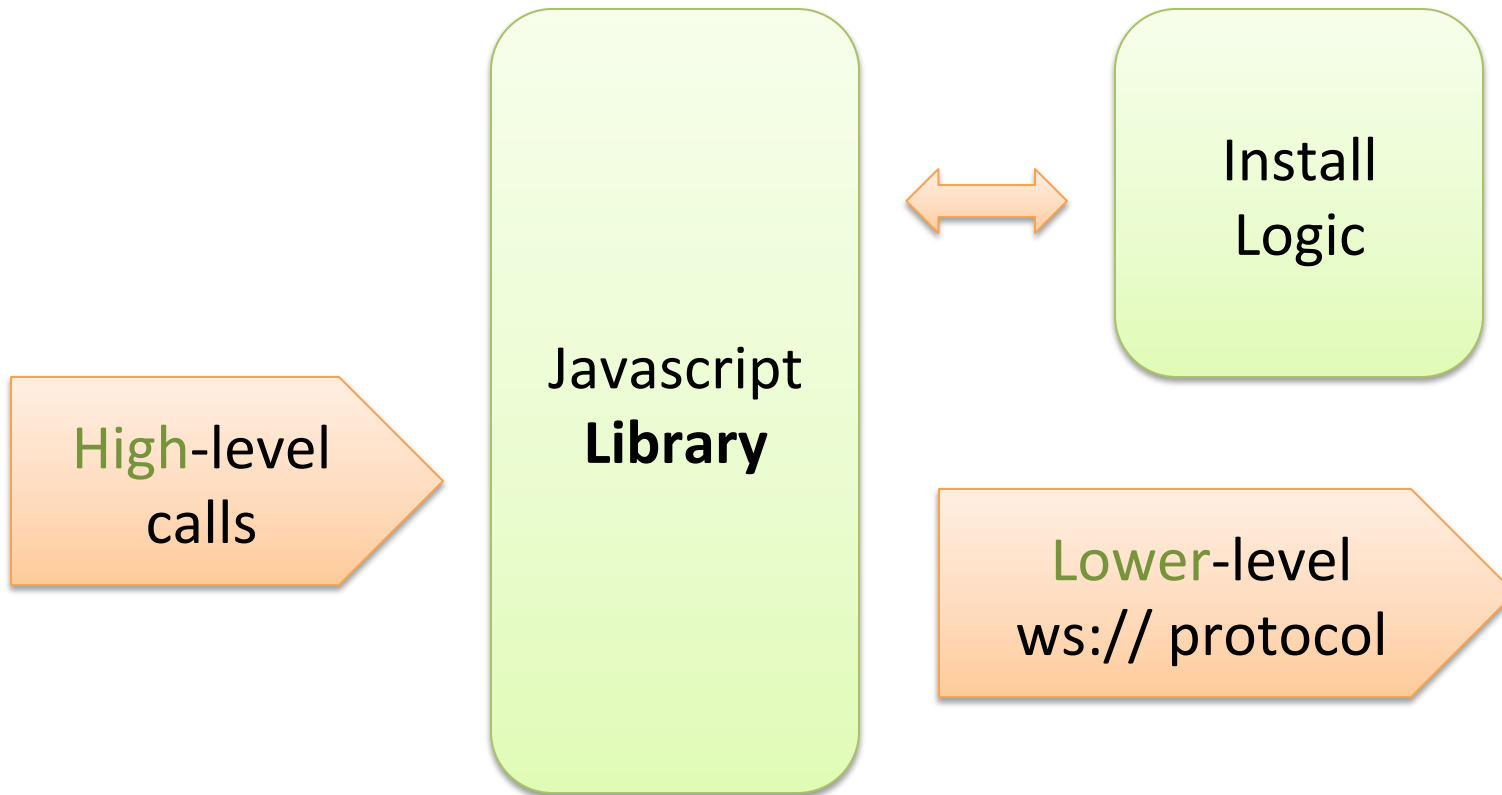
# Under the Hood - libCernVM



Each session is **implemented** as an FSM

2

# Under the Hood – Javascript



**cvmwebapi.js** provides the higher-level abstraction

# Under the Hood – Javascript

The screenshot shows the CernVM WebAPI installation page. At the top, there is a header with the text "Contacting CernVM WebAPI" and a green button that says "Log-in and keep track of your progress". A blue diagonal banner in the top right corner says "Public Beta". The main content area features the CernVM WebAPI logo and the text "You need to install the CernVM WebAPI app to get started." Below this is a large blue button labeled "Install app" with the text "Version 2.0.12 for Macintosh" underneath it. At the bottom of the page, there is a link that says "Learn more information about the CernVM technolo...".

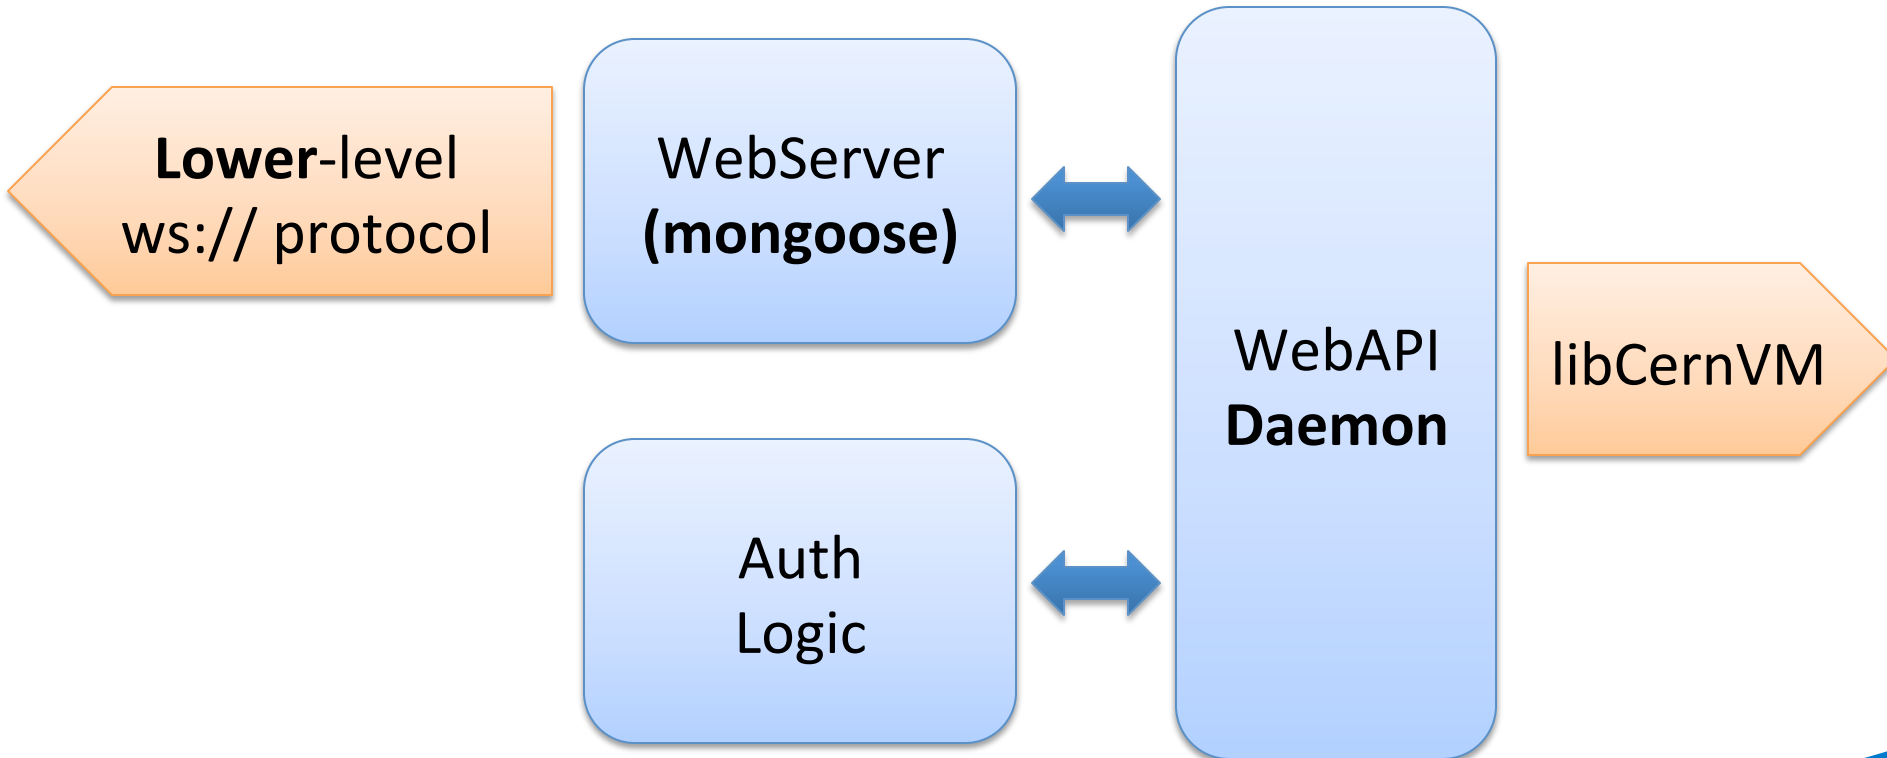
Annotations on the screenshot include:

- An orange arrow pointing to the "Install app" button with the text "First-time **install** instructions".
- An orange arrow pointing to the "Install app" button with the text "Embedded UI".
- A yellow box containing a list of instructions: "1. Click the file at the lower-left of your browser window **cvmwebapi-2.0.12.pkg**. 2. Follow the on-screen instructions to install the CernVM WebAPI app." Below this box is a blue link that says "Click here to retry the download if it failed."

**Micro-UI** injected in the website for **guiding** the user

2

## Under the Hood – Daemon



The **daemon** is just the glue between the **javascript** interface and **libCernVM**



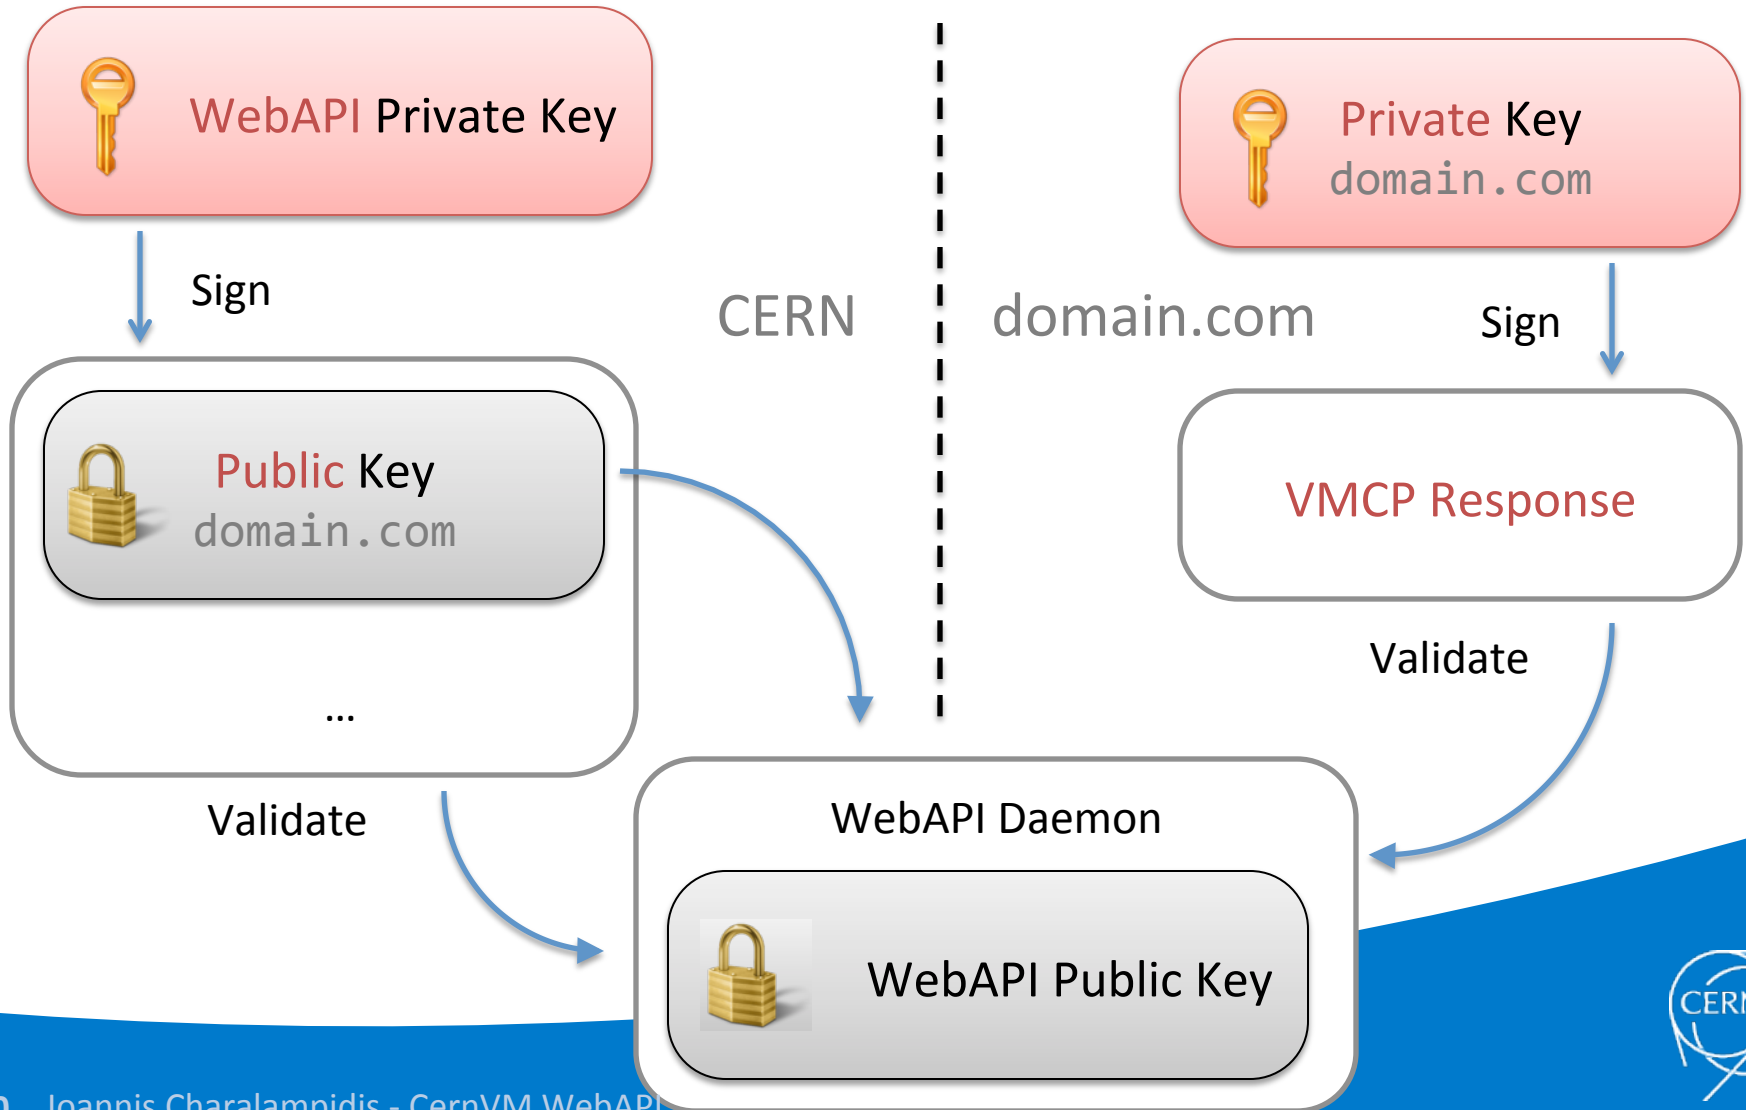
# Security & Trust



Be **secure**, or you are a potential **botnet** node ...

2

# Security & Trust




2

# Security & Trust - VMCP

## 1) Open Session

Referrer: `domain.com`

VMCP: `vmcp.domain.com/vm2`



## 2) Request Config



## 3) Sign response



## 5) Respond

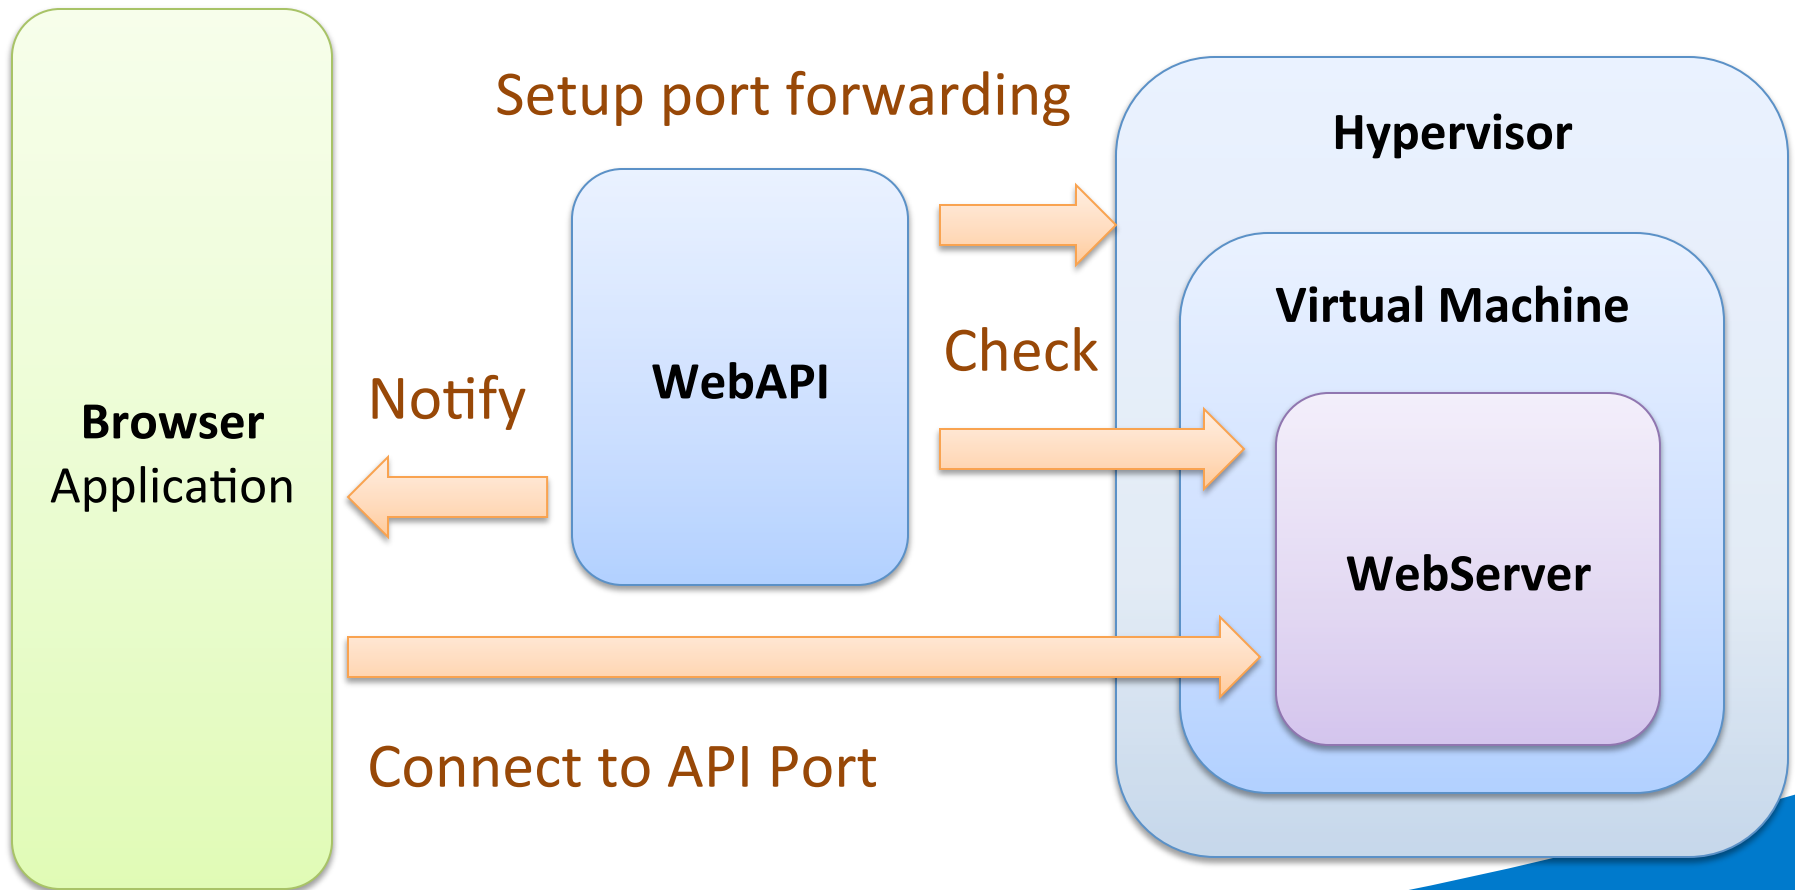


## 4) Forward request



2

# Interaction with the VM



# WebAPI in Production

**Starting virtual event generator** John Charalampidis Progress details Public Beta

Event Rate	Progress	Jobs Completed	Ranking
How many events (particle collisions) your virtual machine is simulating per minute. 500 0	What fraction of the current job is completed. (Typically, 1 job = 100,000 events) 100% 0%	How many simulation jobs your virtual machine has processed. 1 0	Your ranking among the volunteers in this challenge, by number of jobs completed. 1 0

Learn more:

- Introduction to high energy physics simulations
- See the simulations produced by your computer
- Learn about the software that does the simulations

Generator: **pythia8** Collisions analyzed with: **ATLAS\_2011\_S9131140**

Energy: **2000.00 GeV**

Starting virtual event generator

While waiting... wanna help ATLAS look for Higgs?

Stop [Settings] [Trash]

<http://test4theory.cern.ch/vlhc>

**CERN 60** – Public Computing Challenge



# WebAPI in Production

## HALL OF FAME

Here is a name cloud of all the volunteers who signed in and contributed at least one job. Thanks to you, and t

Patrick Werber Hendrik Richter Michael Claes Indy Gilstrøm Jaime Farill Mike Hamilton Kirsten  
 W Hallam Albert Booth Evert adsgafg Ryder Bluhm Harrison Totty Nicolò Gottardello Neil M  
 Katz Wade Gillingham nekdo z jamian Dimock Michelle Greenlee Lex Imperatoris Billy Vier  
 Damme Justin Pekular Jason Lorsung Joshua Lee AXfactOR AXfactOR Giovanni Siragusa Carson  
 Marek kupinsmoke 1973 Maelstrom Scott McDermott Ryan Ford Naomi Cathcart Brian Bond Geoff  
 Blapkinz Ian Peter Braun RJ Hill Bertil Spolander Cody Wang Carl Michael Stojanovic duG Varrette Arturo Saura Ben Segal Eralp Ersoy Alex  
 Copero Tim Weinert Jan Füsting thomas noé Leo Wright Joe Tursi david9000 Borja González Herrero Benjamin Bertrand Pieter  
 Orlando Andrew Smith John Jones Rouslan Korneychuk Jasper Homann Matthew Snow Sylvain  
 Manzi Elemental Brain Adam Paugh Ninette Kelly Cameron Phillips Erick Erickson Emma  
 Bidema Ex Cool Francois Grey Clara Á. Luna Jim 'Artless' Merrill Mikey Babb Chris Ir  
 Engebretsen Adrian Ellingsgaard Jonathan Chan Aika Code Arthur Molnar Clemens Der-ganze  
 Anhgarin Phil Webb Greg Lockett Dubois Alexandre Rohan Kundu Kostas Rakitzis Ayush Jha José Rafael  
 (zeroXten) Amkazan Amkazan Jan-niclas Graumann Not Erison Veshi Matthew Gregg Simone G  
 Poe Felix Wustrack Ivan Koeff Daniel Reynolds Jacob Wolf Nel-gez Andreas Roy van Odenbaeckes ymush101 Robby Wilson Johan  
 Torselius Alexandr Samsonov Steve Kuntz Mike Bentley Daniel Lindmark mininukewarrior Ivan Bisol Frederik Vander Biest Thomas  
 Kristensen Marco Nardes Craig Pekar Mike Wilkinson John Smith Bernie Telalovic Kilazur Dargnaith Laurence Woolford Guillaume  
 Toussaint Adam Wheeler Jade Woods Nicolas Segal Anders Hvidberg Frandsen James Alexander Gardner Efi Psomopoulou Kalle  
 Bexhorn Joseph Flagler Ben Vaughn Anton Stangaard Jakob Rigsby Eric Plummer Josh Robbins Ryan Moodey Kristof Ballet Jack Eden  
 Patrik Frank Dan Shays Shuang Ma Michael Ron Beske Sandro  
 Isaza  
 Boch karevmads Kristensen Lomax francis Marcus Vithuuh bunov Joshua Karim (Samintell) Konrad  
 Mgi  
 Stephen

16.000 sessions

8.000 users

108 countries – 90 languages

Avg. of 400 sessions per day

5,400 people booted a VM

1,100 had problems

(about 80%: slow network)

During the challenge detailed analytics were collected



# WebAPI in **Production**

- Used in **cernvm-online.cern.ch**  
For testing your contextualization image directly from the browser
- Used in **Virtual Atom Smasher** game  
For starting the worker nodes (simulation agents) that players need



3

## Multiple projects with WebAPI

A VM can be easily booted from the browser

- However, starting a VM per project can consume **resources**
- Why not use **lightweight virtualization** (Linux Containers) inside a single VM?

*A 'dumb' scheduler: **DumbQ***

<https://github.com/wavesoft/dumbq>



3

# Multiple projects with WebAPI

50% Chance

```
# Server configuration
project1:50:sft.cern.ch:sft.cern.ch/bootstraps/pj1/init.sh
Project2:30:sft.cern.ch:sft.cern.ch/bootstraps/pj2/init.sh
```

30% Chance

**cernvm-fork**  
pj1/init.sh

Free CPU Slot

Free CPU Slot

3

## Multiple projects with WebAPI

Very **simple integration** with your current setup

- Your script runs in a **standard** CernVM environment
- This is **not a scheduler**, you have to use your own (as your current set-up perhaps)
- DumbQ provides the information required to **identify** the volunteer



# Hands-on (?)

Step-by-step tutorial:

<https://github.com/wavesoft/cernvm-webapi/wiki/Tutorial-Intro>

