



BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Going For It

Conclusions

BOINC Virtual Machine Controller Infrastructure

David García Quintas

CERN, Switzerland

October 23, 2009



Index

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

What?

Why?

Development
(ie, How?)

Going For It

Conclusions

1 Introduction

- What?
- Why?

2 Development (ie, How?)

- Our Proposed Solution
- Getting Technical

3 Going For It

4 Conclusions



Index

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

What?

Why?

Development
(ie, How?)

Going For It

Conclusions

1 Introduction

- What?
- Why?

2 Development (ie, How?)

- Our Proposed Solution
- Getting Technical

3 Going For It

4 Conclusions



What?

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

What?
Why?

Development
(ie, How?)

Going For It

Conclusions

... are we looking for

A means to *interact* with the system running inside a VM instance.

With *code*, not a human, as the “user” behind this interaction (namely, the BOINC wrapper).



What?

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

What?
Why?

Development
(ie, How?)

Going For It

Conclusions

... are we looking for

A means to *interact* with the system running inside a VM instance.

With *code*, not a human, as the “user” behind this interaction (namely, the BOINC wrapper).



Index

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

What?

Why?

Development
(ie, How?)

Going For It

Conclusions

1 Introduction

- What?
- **Why?**

2 Development (ie, How?)

- Our Proposed Solution
- Getting Technical

3 Going For It

4 Conclusions



Why?

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

What?

Why?

Development
(ie, How?)

Going For It

Conclusions

For one...

The porting or adapting an application isn't always affordable. Or even *possible*.



Why?

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

What?

Why?

Development
(ie, How?)

Going For It

Conclusions

But not only:

Pros

- **Ultimate control over the computation environment.**
- Isolation.
- Flexible resource allocation.



Why?

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

What?

Why?

Development
(ie, How?)

Going For It

Conclusions

But not only:

Pros

- Ultimate control over the computation environment.
- **Isolation.**
- Flexible resource allocation.



Why?

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

What?

Why?

Development
(ie, How?)

Going For It

Conclusions

But not only:

Pros

- Ultimate control over the computation environment.
- Isolation.
- **Flexible resource allocation.**



Why?

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

What?

Why?

Development
(ie, How?)

Going For It

Conclusions

On the other hand:

Cons

- Need for a "bigger" machine (esp. RAM + HD).
- Performance penalty ($\approx 3\% - 7\%$).



Why?

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

What?

Why?

Development
(ie, How?)

Going For It

Conclusions

On the other hand:

Cons

- Need for a "bigger" machine (esp. RAM + HD).
- Performance penalty ($\approx 3\% - 7\%$).



Index

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Our Proposed Solution
Getting Technical

Going For It

Conclusions

- 1 Introduction
 - What?
 - Why?
- 2 Development (ie, How?)
 - Our Proposed Solution
 - Getting Technical
- 3 Going For It
- 4 Conclusions



Requirements

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

**Development
(ie, How?)**

Our Proposed Solution
Getting Technical

Going For It

Conclusions

Cool stuff, I'm all in! What do we need?

To be able to...

- Start / stop the VM.
- Pause / unpause the VM.
- Snapshots (checkpoints).
- VM creation.
- File transfer to/from the VM.
- Arbitrary command execution on the VM.



Requirements

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Our Proposed Solution
Getting Technical

Going For It

Conclusions

Cool stuff, I'm all in! What do we need?
To be able to...

- **Start / stop the VM.**
- Pause / unpause the VM.
- Snapshots (checkpoints).
- VM creation.
- File transfer to/from the VM.
- Arbitrary command execution on the VM.



Requirements

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Our Proposed Solution
Getting Technical

Going For It

Conclusions

Cool stuff, I'm all in! What do we need?
To be able to...

- Start / stop the VM.
- **Pause / unpause the VM.**
- Snapshots (checkpoints).
- VM creation.
- File transfer to/from the VM.
- Arbitrary command execution on the VM.



Requirements

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Our Proposed Solution
Getting Technical

Going For It

Conclusions

Cool stuff, I'm all in! What do we need?
To be able to...

- Start / stop the VM.
- Pause / unpause the VM.
- **Snapshots (checkpoints).**
- VM creation.
- File transfer to/from the VM.
- Arbitrary command execution on the VM.



Requirements

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Our Proposed Solution
Getting Technical

Going For It

Conclusions

Cool stuff, I'm all in! What do we need?
To be able to...

- Start / stop the VM.
- Pause / unpause the VM.
- Snapshots (checkpoints).
- **VM creation.**
- File transfer to/from the VM.
- Arbitrary command execution on the VM.



Requirements

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Our Proposed Solution
Getting Technical

Going For It

Conclusions

Cool stuff, I'm all in! What do we need?
To be able to...

- Start / stop the VM.
- Pause / unpause the VM.
- Snapshots (checkpoints).
- VM creation.
- **File transfer to/from the VM.**
- Arbitrary command execution on the VM.



Requirements

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Our Proposed Solution
Getting Technical

Going For It

Conclusions

Cool stuff, I'm all in! What do we need?
To be able to...

- Start / stop the VM.
- Pause / unpause the VM.
- Snapshots (checkpoints).
- VM creation.
- File transfer to/from the VM.
- **Arbitrary command execution on the VM.**



Index

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Our Proposed Solution
Getting Technical

Going For It

Conclusions

- 1 Introduction
 - What?
 - Why?
- 2 Development (ie, How?)
 - Our Proposed Solution
 - Getting Technical
- 3 Going For It
- 4 Conclusions



Three technologies to help us reach our goals:

- **Python.**
- Twisted Framework.
- STOMP Protocol
(<http://stomp.codehaus.org/Protocol>).
- Chirp Protocol
(<http://www.cse.nd.edu/~ccl/software/chirp/>).



Three technologies to help us reach our goals:

- Python.
- **Twisted Framework.**
- STOMP Protocol
(<http://stomp.codehaus.org/Protocol>).
- Chirp Protocol
(<http://www.cse.nd.edu/~ccl/software/chirp/>).



Three technologies to help us reach our goals:

- Python.
- Twisted Framework.
- **STOMP Protocol**
(<http://stomp.codehaus.org/Protocol>).
- Chirp Protocol
(<http://www.cse.nd.edu/~ccl/software/chirp/>).



Three technologies to help us reach our goals:

- Python.
- Twisted Framework.
- STOMP Protocol
(<http://stomp.codehaus.org/Protocol>).
- **Chirp Protocol**
(<http://www.cse.nd.edu/~ccl/software/chirp/>).



What we get from these choices:

- **Multiplatform (tested on Windows, Linux and OS X).**
- Powerful event-driven network framework.
- Simple yet flexible message passing protocol.



What we get from these choices:

- Multiplatform (tested on Windows, Linux and OS X).
- **Powerful event-driven network framework.**
- Simple yet flexible message passing protocol.



What we get from these choices:

- Multiplatform (tested on Windows, Linux and OS X).
- Powerful event-driven network framework.
- **Simple yet flexible message passing protocol.**



Wait a minute...

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Our Proposed Solution
Getting Technical

Going For It

Conclusions

How does all this fit into BOINC?



The Overall Picture

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

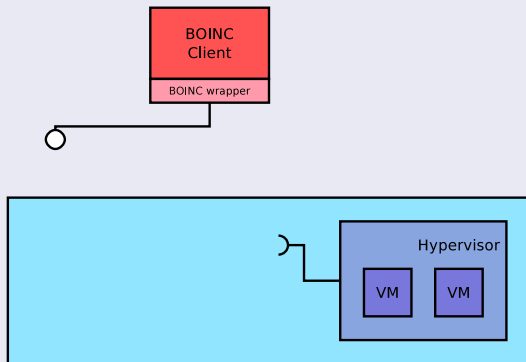
Development
(ie, How?)

Our Proposed Solution
Getting Technical

Going For It

Conclusions

Where Do We Fit?



- BOINC (wrapper) exports an API.
- The Hypervisor exports an API.



The Overall Picture

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

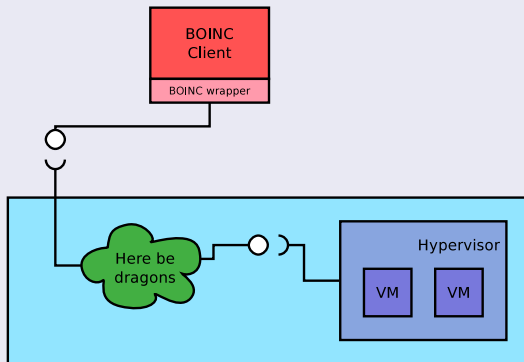
Development
(ie, How?)

Our Proposed Solution
Getting Technical

Going For It

Conclusions

Where Do We Fit?



All that's left to do is to interface both APIs, providing them with what they need from each other.



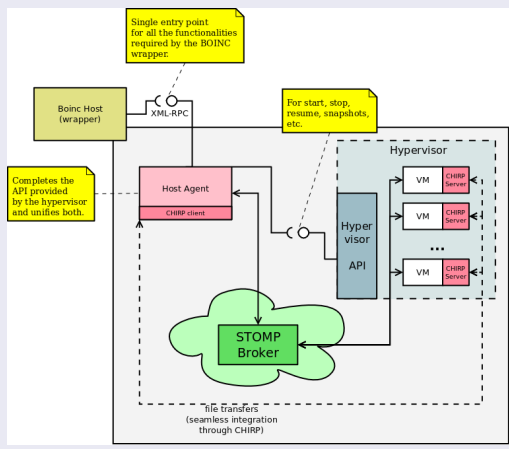
In Detail

BOINC Virtual Machine Controller Infrastructure
David García Quintas

Introduction
Development (ie, How?)
Our Proposed Solution
Getting Technical
Going For It
Conclusions

And this is how this interfacing is realized:

Architecture





What we get out of all this

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Our Proposed Solution
Getting Technical

Going For It

Conclusions

Main advantages of doing things this way:

- A highly decoupled system (could be “clouded”).
- Easily extensible.
- Multiplatform.
- Scalable (basically, as much as the broker).
- Hypervisor agnostic.



What we get out of all this

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Our Proposed Solution
Getting Technical

Going For It

Conclusions

Main advantages of doing things this way:

- A highly decoupled system (could be “clouded”).
- **Easily extensible.**
- Multiplatform.
- Scalable (basically, as much as the broker).
- Hypervisor agnostic.



What we get out of all this

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Our Proposed Solution
Getting Technical

Going For It

Conclusions

Main advantages of doing things this way:

- A highly decoupled system (could be “clouded”).
- Easily extensible.
- **Multiplatform.**
- Scalable (basically, as much as the broker).
- Hypervisor agnostic.



What we get out of all this

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Our Proposed Solution
Getting Technical

Going For It

Conclusions

Main advantages of doing things this way:

- A highly decoupled system (could be “clouded”).
- Easily extensible.
- Multiplatform.
- **Scalable (basically, as much as the broker).**
- Hypervisor agnostic.



What we get out of all this

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Our Proposed Solution
Getting Technical

Going For It

Conclusions

Main advantages of doing things this way:

- A highly decoupled system (could be “clouded”).
- Easily extensible.
- Multiplatform.
- Scalable (basically, as much as the broker).
- **Hypervisor agnostic.**



Let's Get Technical

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Our Proposed Solution
Getting Technical

Going For It

Conclusions

Enough with the marketing!



Index

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Our Proposed Solution
Getting Technical

Going For It

Conclusions

- 1 Introduction
 - What?
 - Why?
- 2 Development (ie, How?)
 - Our Proposed Solution
 - Getting Technical
- 3 Going For It
- 4 Conclusions



A Simple Example

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

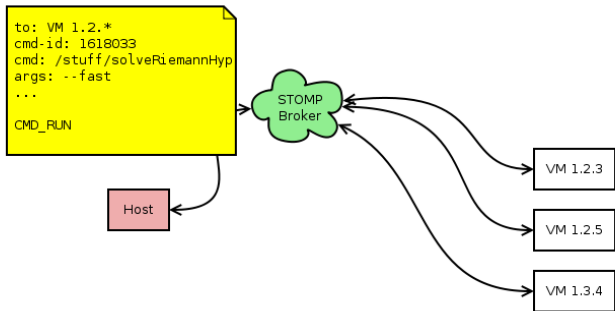
Our Proposed Solution

Getting Technical

Going For It

Conclusions

A Simple Example



Host (VM Controller) requests to *some* of the VMs the execution of a command.



A Simple Example

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)


Our Proposed Solution

Getting Technical

Going For It

Conclusions

A Simple Example



The addressed VMs process the incoming msg.



A Simple Example

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)


Our Proposed Solution

Getting Technical

Going For It

Conclusions

A Simple Example



They stay busy for a while...



A Simple Example

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

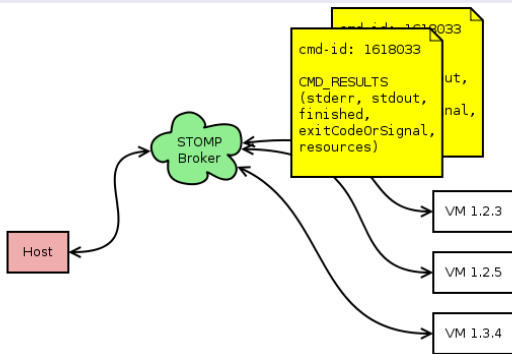
Our Proposed Solution

Getting Technical

Going For It

Conclusions

A Simple Example



... and eventually they come up with something. At any given moment! All operations are asynchronous.



A Simple Example

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

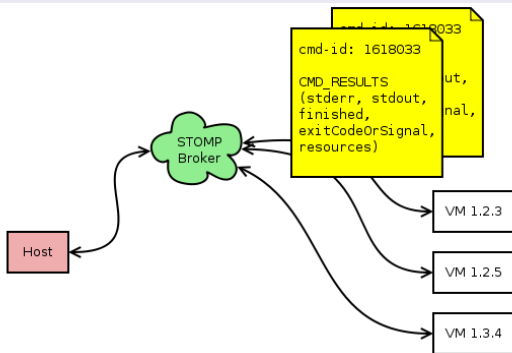
Our Proposed Solution

Getting Technical

Going For It

Conclusions

A Simple Example



... and eventually they come up with something. At any given moment! All operations are asynchronous.



A Simple Example

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

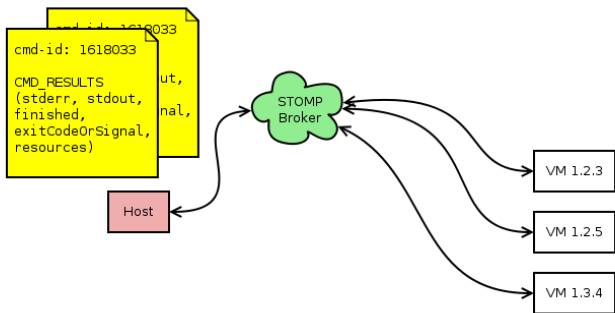
Our Proposed Solution

Getting Technical

Going For It

Conclusions

A Simple Example



The VM controller (thus the BOINC client) eventually receives the result(s) of the computation (together with some extra data).



Index

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Going For It

Conclusions

- 1 Introduction
 - What?
 - Why?
- 2 Development (ie, How?)
 - Our Proposed Solution
 - Getting Technical
- 3 **Going For It**
- 4 Conclusions



Going For It

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Going For It

Conclusions

How much is it going to cost me?

On the VM image:

- Python Runtime (≥ 2.4 . 2.6 recommended).
- Twisted Framework.
- Zope Interfaces.
- Stomper.
- Netifaces.
- Chirp.



Going For It

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Going For It

Conclusions

How much is it going to cost me?

On the VM image:

- Python Runtime (≥ 2.4 . 2.6 recommended).
- Twisted Framework.
- Zope Interfaces.
- Stomper.
- Netifaces.
- Chirp.



Going For It

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Going For It

Conclusions

How much is it going to cost me?

On the host side (ie, BOINC client):

- Python Runtime (\geq 2.4. 2.6 recommended).
- Twisted Framework.
- Stomper.
- Chirp.
- Your hypervisor of choice (we've been using VirtualBox).



Going For It

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Going For It

Conclusions

How much is it going to cost me?

On the host side (ie, BOINC client):

- Python Runtime (\geq 2.4. 2.6 recommended).
- Twisted Framework.
- Stomper.
- Chirp.
- Your hypervisor of choice (we've been using VirtualBox).



Trials at CERN

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Going For It

Conclusions

No vaporware!

Put to the test in the context of the ALICE experiment at CERN.
(successfully so)



Trials at CERN

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Going For It

Conclusions

No vaporware!

Put to the test in the context of the ALICE experiment at CERN.
(successfully so)



Index

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Going For It

Conclusions

- 1 Introduction
 - What?
 - Why?
- 2 Development (ie, How?)
 - Our Proposed Solution
 - Getting Technical
- 3 Going For It
- 4 Conclusions



Wrapping up

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Going For It

Conclusions

Something to take away from this talk:

- **Virtualization eases the deployment of certain computations. In some cases, making them possible altogether.**
- It is now possible to BOINC-ify applications strongly tied to a certain environment (by virtualizing it).
- You're not limited by your choice of hypervisor.



Wrapping up

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Going For It

Conclusions

Something to take away from this talk:

- Virtualization eases the deployment of certain computations. In some cases, making them possible altogether.
- **It is now possible to BOINC-ify applications strongly tied to a certain environment (by virtualizing it).**
- You're not limited by your choice of hypervisor.



Wrapping up

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Going For It

Conclusions

Something to take away from this talk:

- Virtualization eases the deployment of certain computations. In some cases, making them possible altogether.
- It is now possible to BOINC-ify applications strongly tied to a certain environment (by virtualizing it).
- **You're not limited by your choice of hypervisor.**



Wrapping up

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Going For It

Conclusions

It's even documented!

<http://boinc.berkeley.edu/trac/wiki/VmApps>

Overall architecture and VirtualBox specific details:

<http://boinc.berkeley.edu/trac/wiki/VirtualBox>



Wrapping up

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Going For It

Conclusions

It's even documented!

<http://boinc.berkeley.edu/trac/wiki/VmApps>

Overall architecture and VirtualBox specific details:

<http://boinc.berkeley.edu/trac/wiki/VirtualBox>



What's next?

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Going For It

Conclusions

What's next?

- **Support for more hypervisors.**
- Automate the installation process, on both sides.
- Testing, testing, testing!



What's next?

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Going For It

Conclusions

What's next?

- Support for more hypervisors.
- Automate the installation process, on both sides.
- Testing, testing, testing!



What's next?

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Going For It

Conclusions

What's next?

- Support for more hypervisors.
- Automate the installation process, on both sides.
- **Testing, testing, testing!**



The End

BOINC Virtual
Machine
Controller
Infrastructure

David García
Quintas

Introduction

Development
(ie, How?)

Going For It

Conclusions

Thank you!