Volunteer Computing and Virtualization

Ben Segal

Holger Schulz

and:

Predrag Buncic, David Garcia, Jakob Blomer, Pere Mato, Carlos Aguado / CERN
Artem Harutyunyan / Yerevan Physics Institute
Jarno Rantala / Tampere University of Technology
David Weir / Imperial College, London
Yushu Yao / Lawrence Berkeley Laboratory
Rohit Yadav / Bnares Hindu University

2nd Workshop on Multi-Core and Virtualization

CERN

June 21-22, 2010

What is BOINC?

- **"B**erkeley **O**pen **I**nfrastructure for **N**etwork **C**omputing"
- **Software platform for distributed computing using volunteered computer resources**
- **http://boinc.berkeley.edu**
- **Uses a volunteer PC's unused CPU cycles to analyse scientific data**
- **Client-server architecture**
- **Free and Open-source**
- **Also handles DESKTOP GRIDS**

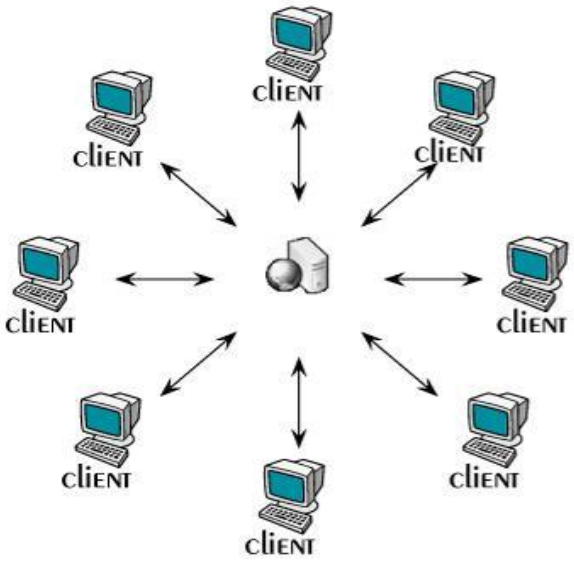

Some volunteer computing projects

SCIENCE
SETI@home (BOINC)
evolution@home
eOn
climateprediction.net (BOINC)
Muon1
LHC@home (BOINC)
Einstein@Home(BOINC)
BBC Climate Change Experiment (BOINC)
Leiden Classical (BOINC)
QMC@home (BOINC)
NanoHive@Home (BOINC)
µFluids@Home (BOINC)
Spinhenge@home (BOINC)
Cosmology@Home (BOINC)
PS3GRID (BOINC)
Mars Clickworkers

LIFE SCIENCES
Parabon Computation
Folding@home
FightAIDS@home
Übero
Drug Design Optimization Lab (D2OL)
The Virtual Laboratory Project
Community TSC
Predictor@home (BOINC)
XGrid@Stanford
Human Proteome Folding (WCG)
CHRONOS (BOINC)
Rosetta@home (BOINC)
RALPH@home (BOINC)
SIMAP (BOINC)
malariacontrol.net (BOINC)
Help Defeat Cancer (WCG)
TANPAKU (BOINC)
Genome Comparison (WCG)
Docking@Home (BOINC)
proteins@home (BOINC)
Help Cure Muscular Dystrophy (WCG)

MATHEMATICS & CRYPTOGRAPHY
Great Internet Mersenne Prime Search
Proth Prime Search
ECMNET
Minimal Equal Sums of Like Powers
MM61 Project
3x + 1 Problem
Distributed Search for Fermat Number Divisors
PCP@Home
Generalized Fermat Prime Search
PSearch
Seventeen or Bust
Factorizations of Cyclotomic Numbers
Goldbach Conjecture Verification
The Riesel Problem
The 3*2^n-1 Search
NFSNET
Search for Multifactorial Primes
15k Prime Search
ElevenSmooth
Riesel Sieve
The Prime Sierpinski Project
P.I.E.S. - Prime Internet Eisenstein Search
Factors of k*2^n±1
XYYXF
12121 Search
2721 Search
Operation Billion Digits
SIGPS
Primesearch

INTERNET PERFORMANCE
Gómez Performance ($)
Network Peer
NETI@home
dCrawl
DIMES
Red Library DLV
Majestic-12
Boitho
PeerFactor
DepSpid
Pingdom GIGRIB
Project Neuron(BOINC)

ECONOMICS
MoneyBee
Gstock

GAMES
ChessBrain
Chess960@home (BOINC)

ART
Electric sheep
Internet Movie Project
RenderFarm@home (BOINC)

Source: http://distributedcomputing.info/

The BOINC community

- Competition between individuals and teams for “credit”.
- Websites and regular updates on status of project by scientists.
- Forums for users to discuss the science behind the project.
- E.g. for LHC@home, the volunteers show great interest in CERN and the LHC.
- Supply each other with scientific information and even help debug the project.

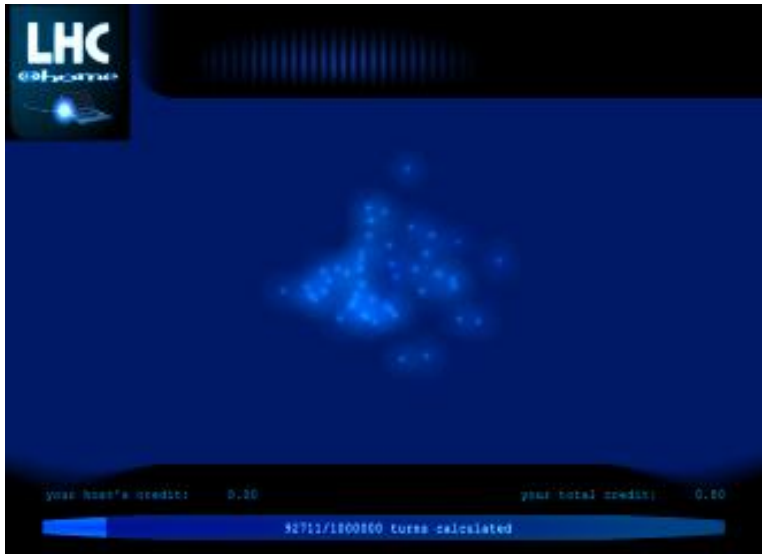

LHC@home screensaver

LHC@home

- Calculates stability of proton orbits in CERN's new LHC accelerator
- System is nonlinear and unstable so numerically very sensitive. Hard to get identical results on all platforms
- About 40 000 users, 70 000 PC's… over 1500 CPU years of processing
- Objectives: extra CPU power and raising public awareness of CERN and the LHC - both successfully achieved.
- Started as an outreach project for CERN 50th Anniversary 2004; used for Year of Physics (Einstein Year) 2005

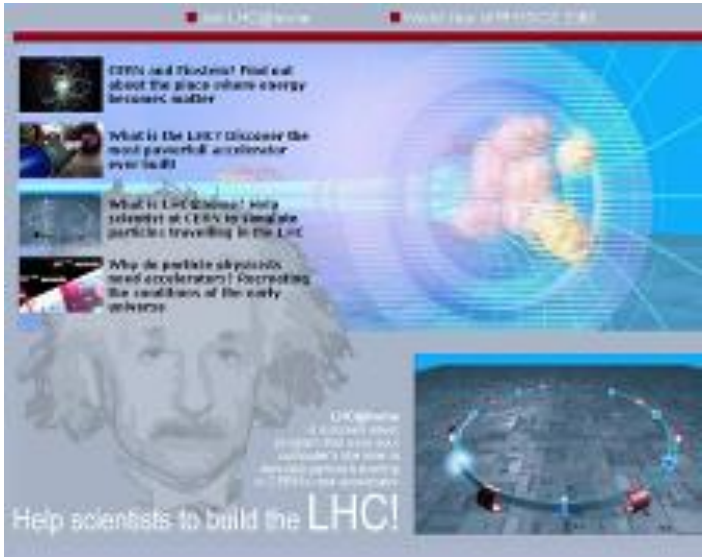

BOINC & LHC physics code

Problems with "normal" BOINC used for LHC physics:

1) A project's application(s) must be ported to every volunteer platform of interest: most clients run Windows, but CERN runs Scientific Linux and porting to Windows is impractical.

2) The project's work must be fed into the BOINC server for distribution, and results must be recovered. "Job submission scripts" must be developed for this, but CERN physics experiments won't change their current setups.

3) Job management is very primitive in BOINC, whereas physicists want to know where their jobs are and be able to manage them.

Solution: a Volunteer Cloud

- **In this case, the Cloud worker nodes are:**

  Volunteer PC's running CernVM…

- **Solves porting problem to all client platforms (Windows, Mac, Linux):**
- **Solves image size problem**
- **Solves job production interface problem**

- **All done without changing existing BOINC infrastructure (client or server side)**
- **All done without changing physicists' code or procedures**

- **How is it done ? … with CoPilot and BOINC …**

BOINC & Virtualization

CernVM and Co-Pilot allow us to solve the three problems of *porting, job submission, and job management,* but to run CernVM guest VM's within a BOINC host we need a cross-platform solution for:

- control of (multiple) VM's on a host, including: Start|Stop|pause|resume|reset|poweroff|savestate
- command execution on guest VM's
- file transfers from guests to host (and reverse)

BOINC & Virtualization

Details of the “VM controller” package:

(developed by David Garcia Quintas / CERN)

- Cross-platform support - based on Python (Windows, MacOSX, Linux…).
  - Uses Python packages:

    Netifaces, Stomper, Twisted, Zope, simplejson, Chirp…
- Does asynchronous message passing between host and guest entities via a broker (e.g. ActiveMQ). Messages are XML/RPC based.
- Supports:
  - control of multiple VM’s on a host, including:

    Start|Stop|pause|resume|reset|poweroff|savestate
  - command execution on guest VM’s
  - file transfers from guests to host (and reverse) using Chirp

Host to VM Guest communication

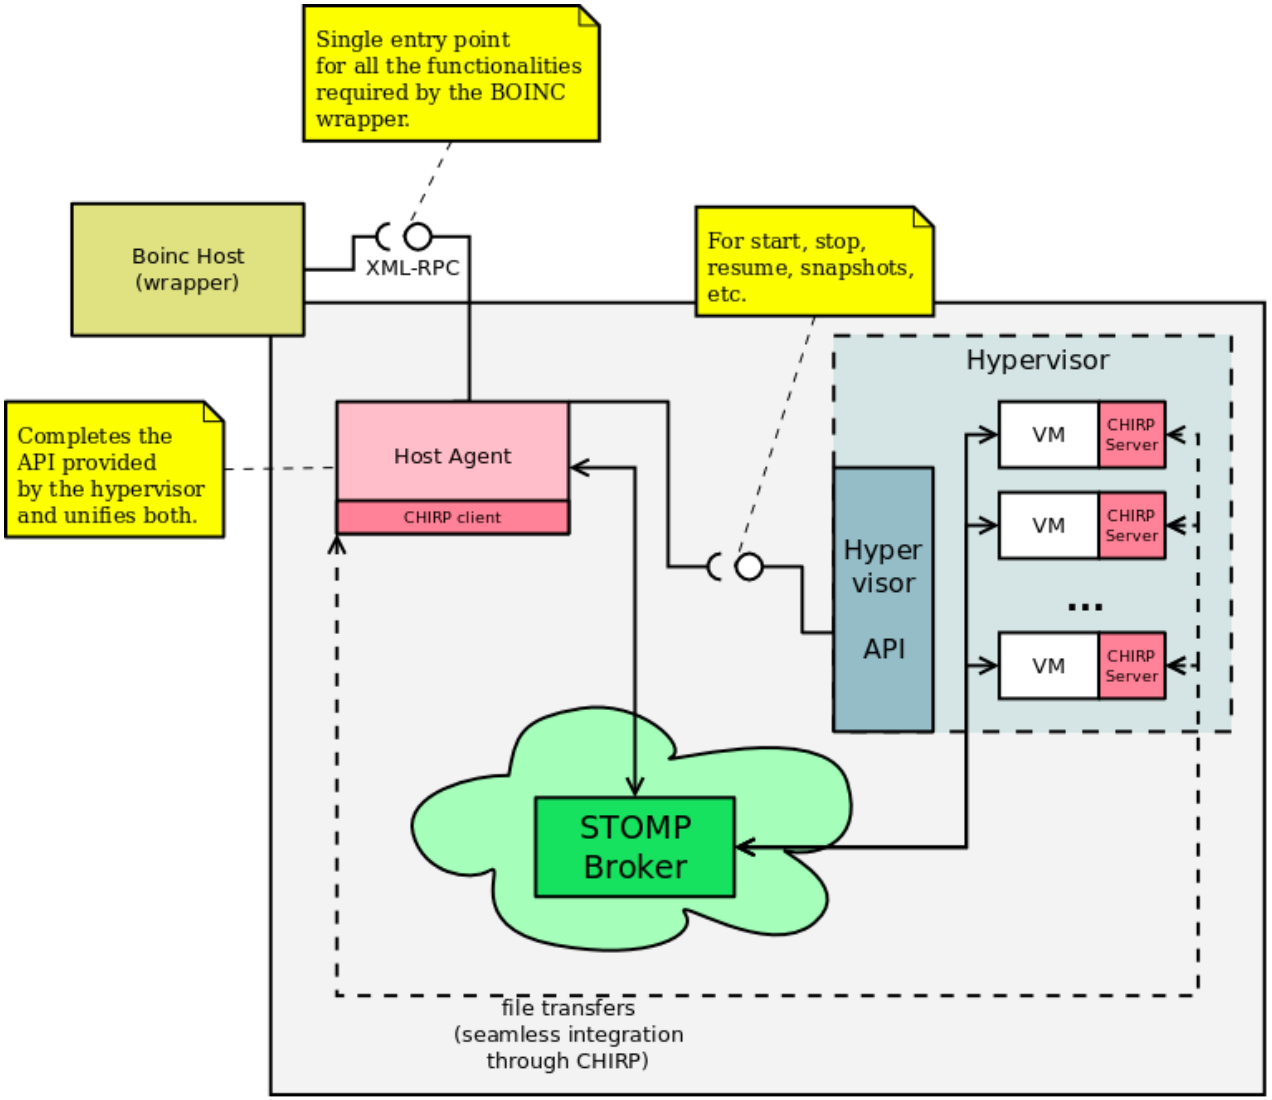

BOINC & Virtualization

Details of the new BOINC "Vmwrapper":
(developed by Jarno Rantala / CERN openlab student, 2009)

- Written in Python, therefore multi-platform
- Uses "VM controller" infrastructure described above
- Back-compatible with original BOINC Wrapper
  - Supports standard BOINC job.xml files
  - For VM case, supports extra tags in the job.xml file
- Able to measure the VM guest resources and issue credit requests
  - ..including "partial credits" to allow very long-running processes/jobs

BOINC VMwrapper architecture

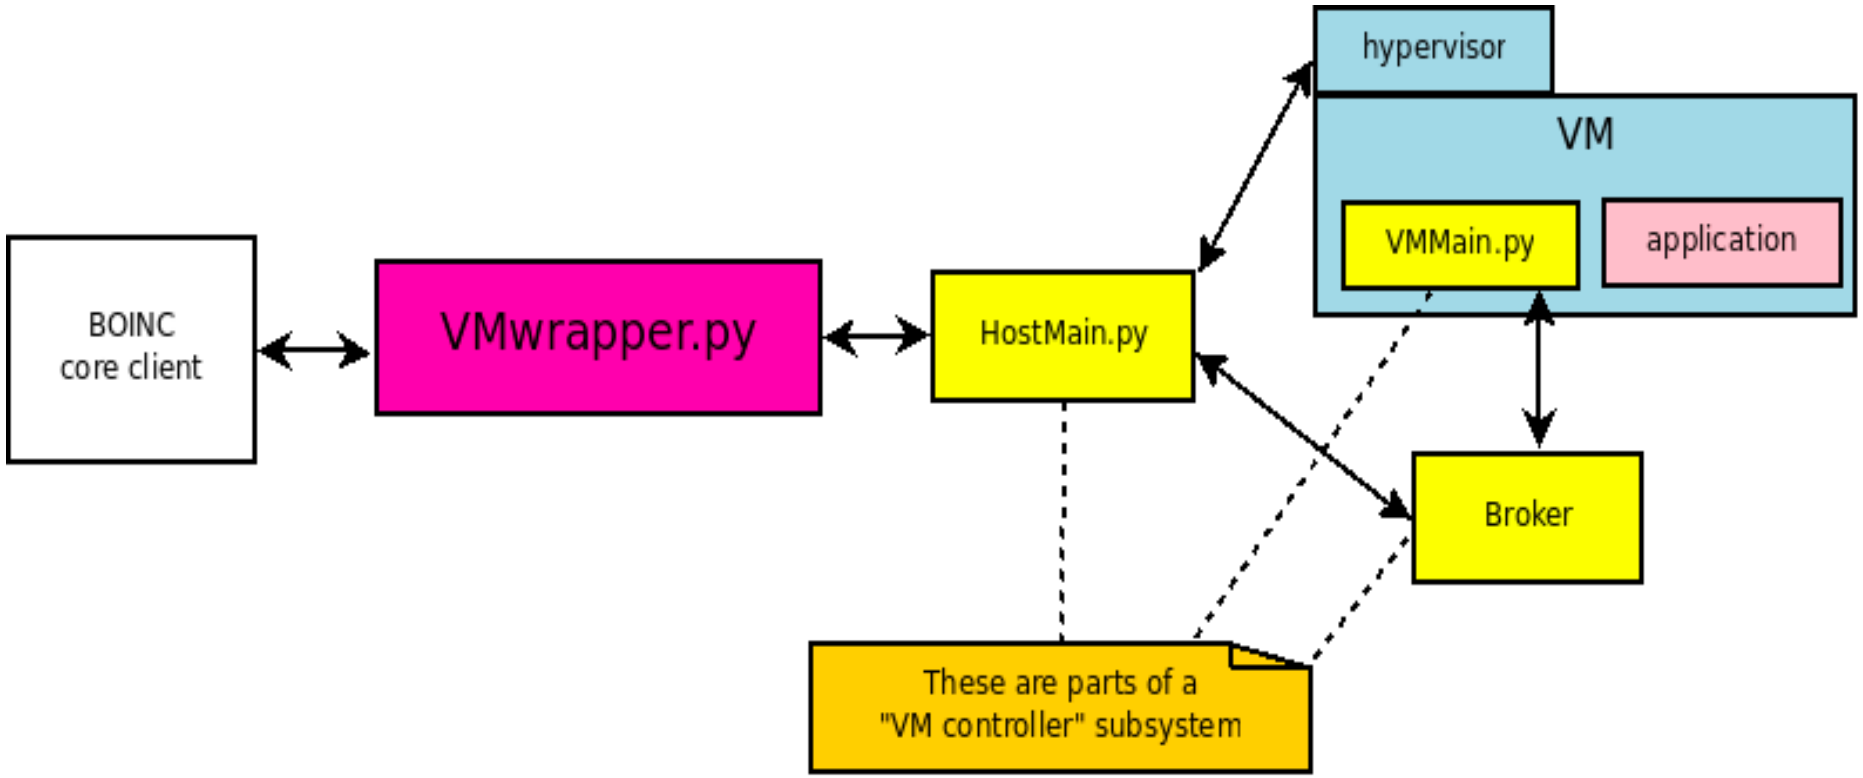

BOINC Virtual Cloud

Summary of the method:

- **New BOINC wrapper (VMWrapper) used to start a guest Virtual machine in BOINC client PC, and execute a CernVM image.**
- **The CernVM image has all LHC software and CoPilot code.**
- **Host-to-VM communication/control provided for any BOINC PC.**
- **The new Vmwrapper gives BOINC client and server all the functions they need - they are unaware of VM's…**
- **The CoPilot allows LHC job production to proceed without changes.**

CernVM Co-Pilot architecture

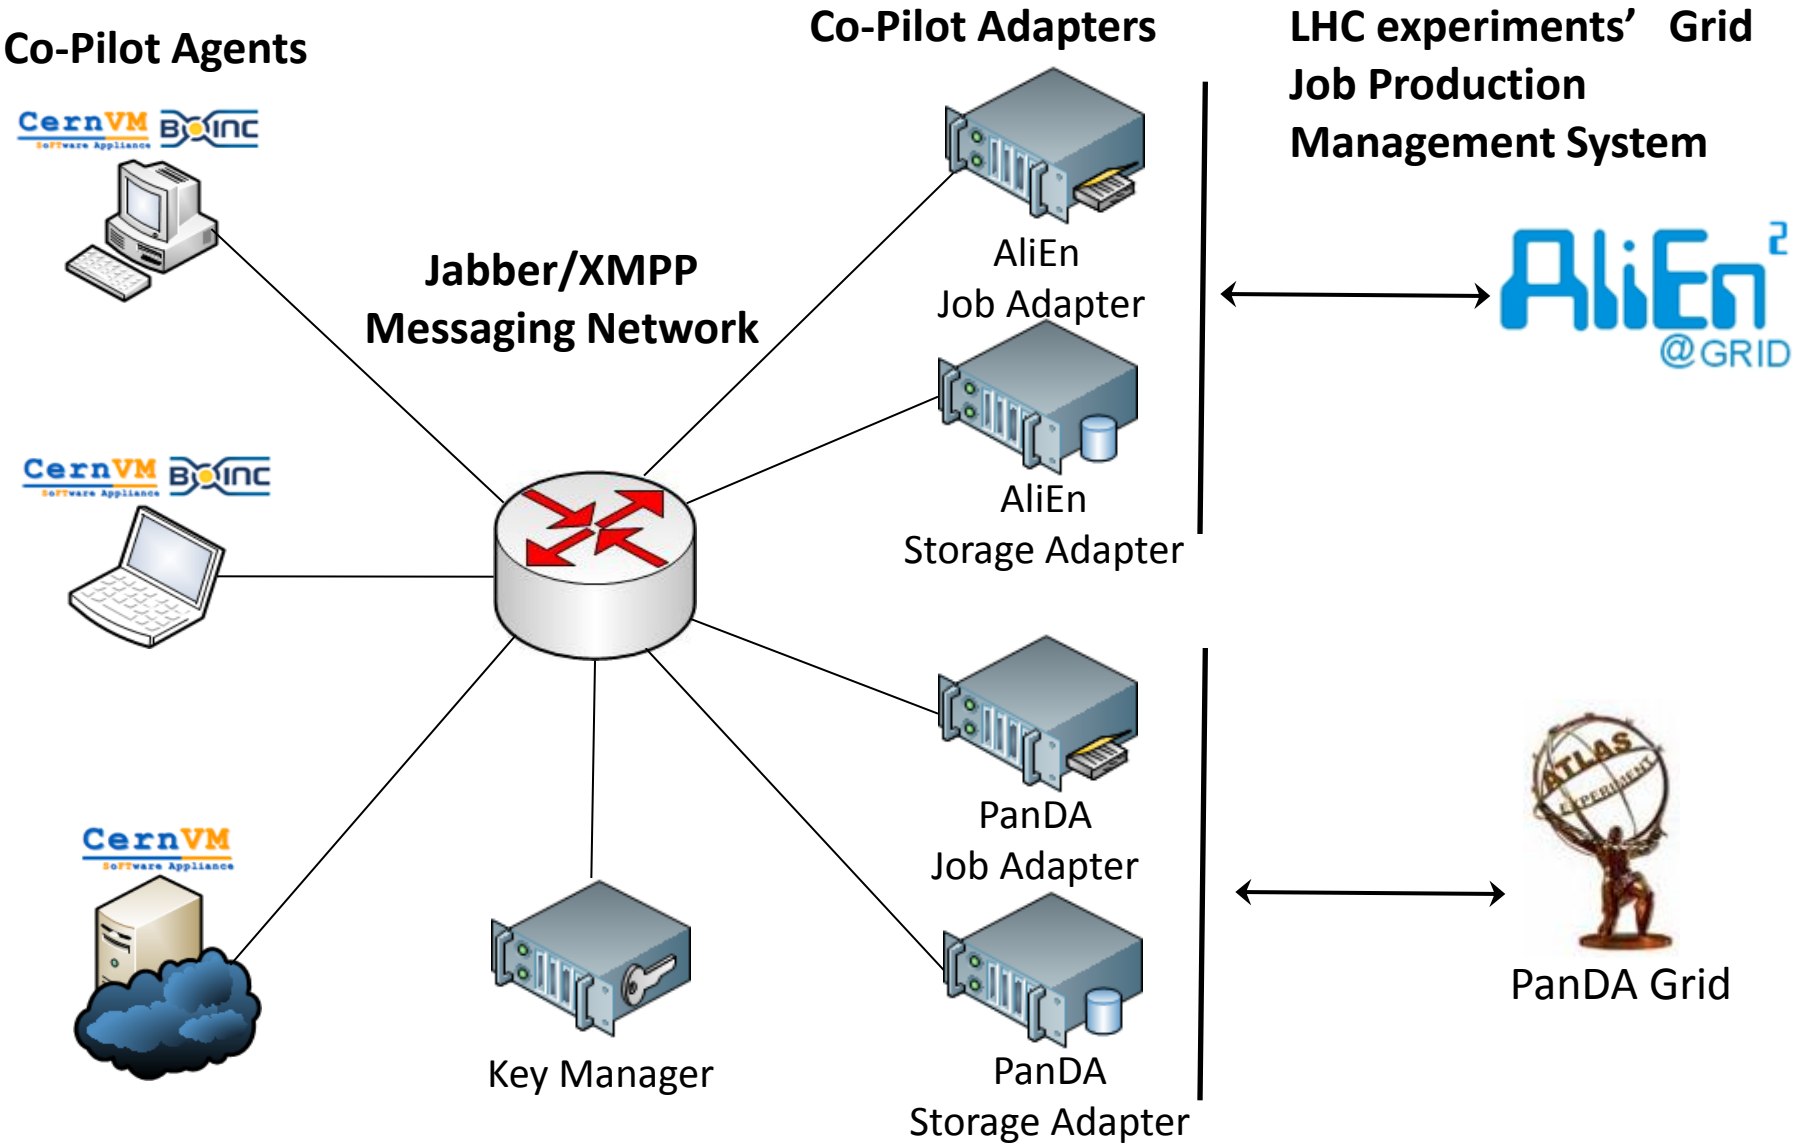

The use of Jabber/XMPP allows to scale the system in case of high load by just adding new Adapter instances

The work of CernVM Co-Pilot Adapters

- Job Adapter
  - Receives job execution requests from an Agent
  - Contacts Grid services (e.g. AliEn/PanDA/Dirac) and gets a job for execution
  - Gets the necessary input files from the Grid file catalogue
  - Makes input files available for the Agent via Chirp
  - Instructs Agent to download the input
  - Instructs Agent to execute job command (e.g. 'root myMacro.C')
- Storage Adapter
  - Receives job completion request
  - Provides Agent with the Chirp directory to upload job output
  - Puts the job output to the Grid file catalogue
  - Contacts Grid services and sets the final status of the job
- **NOTE: Grid credentials are handled by Adapters, not sent to Agents**

CernVM Co-Pilot job execution (ALICE)

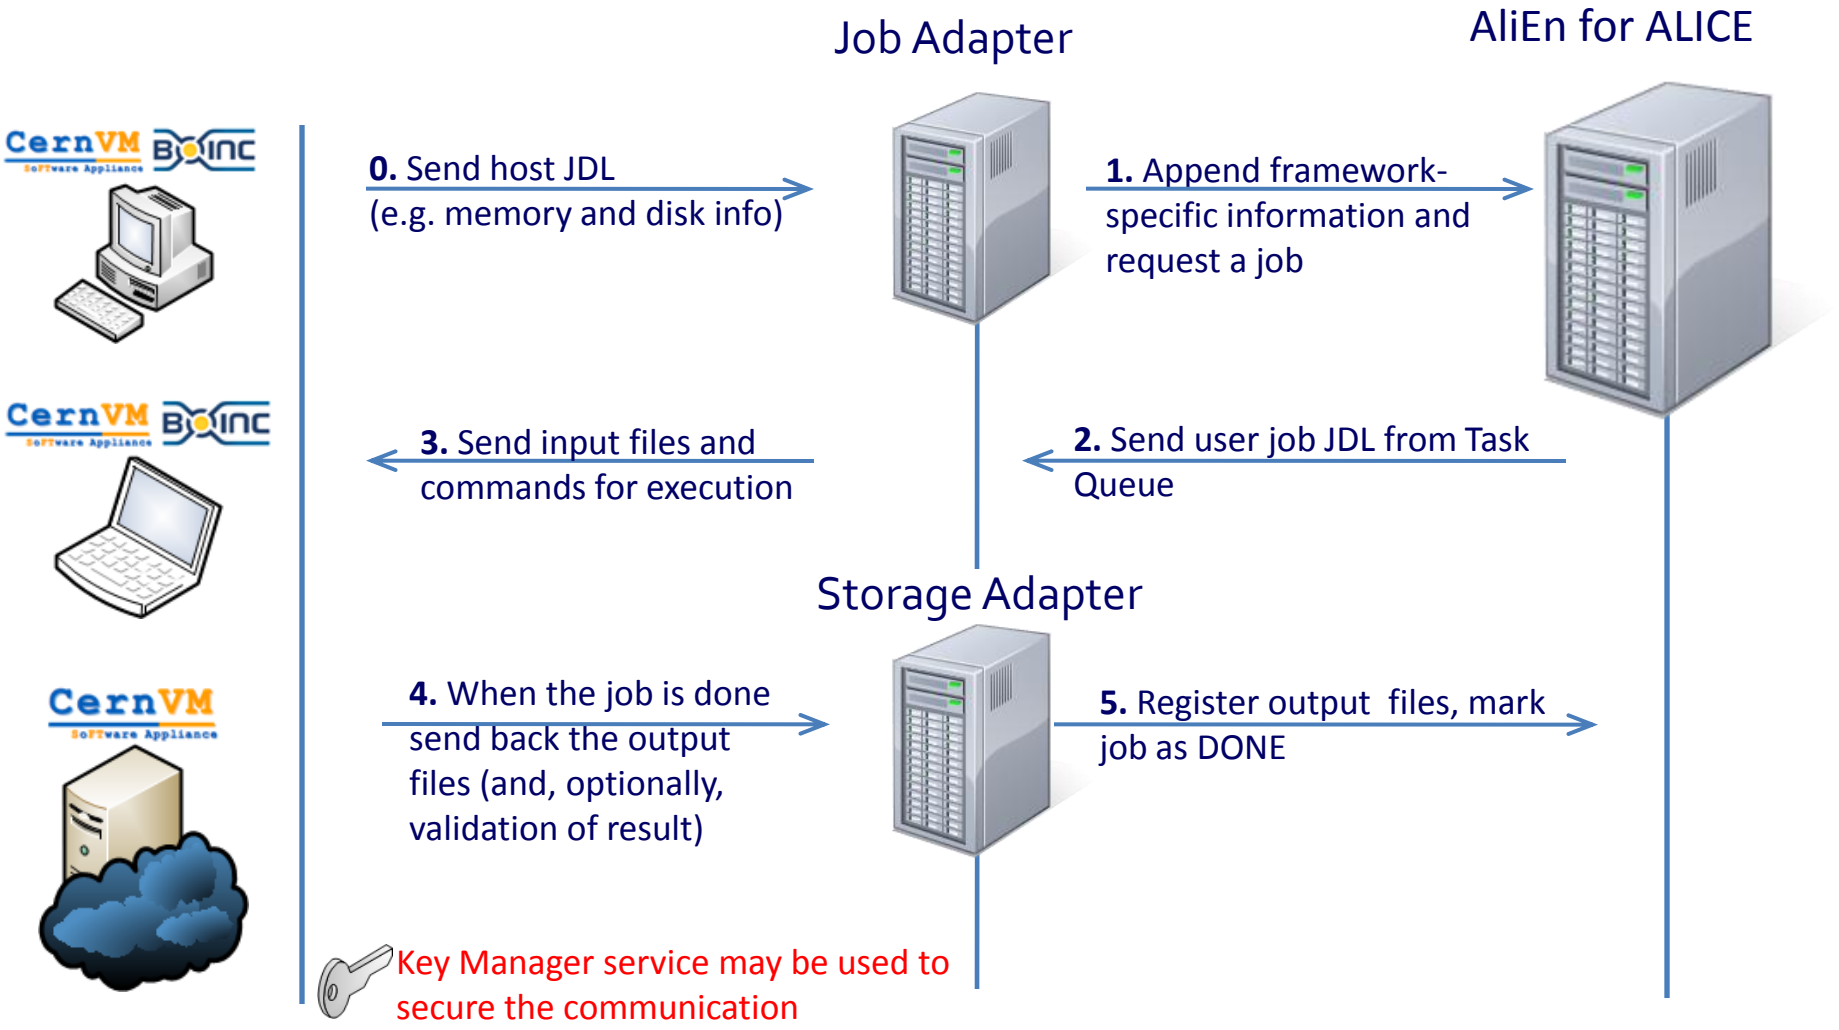

For the detailed description of the communication protocol please see:
https://cernvm.cern.ch/project/trac/cernvm/wiki/CoPilotProtocol

CernVM Co-Pilot job execution (ATLAS)

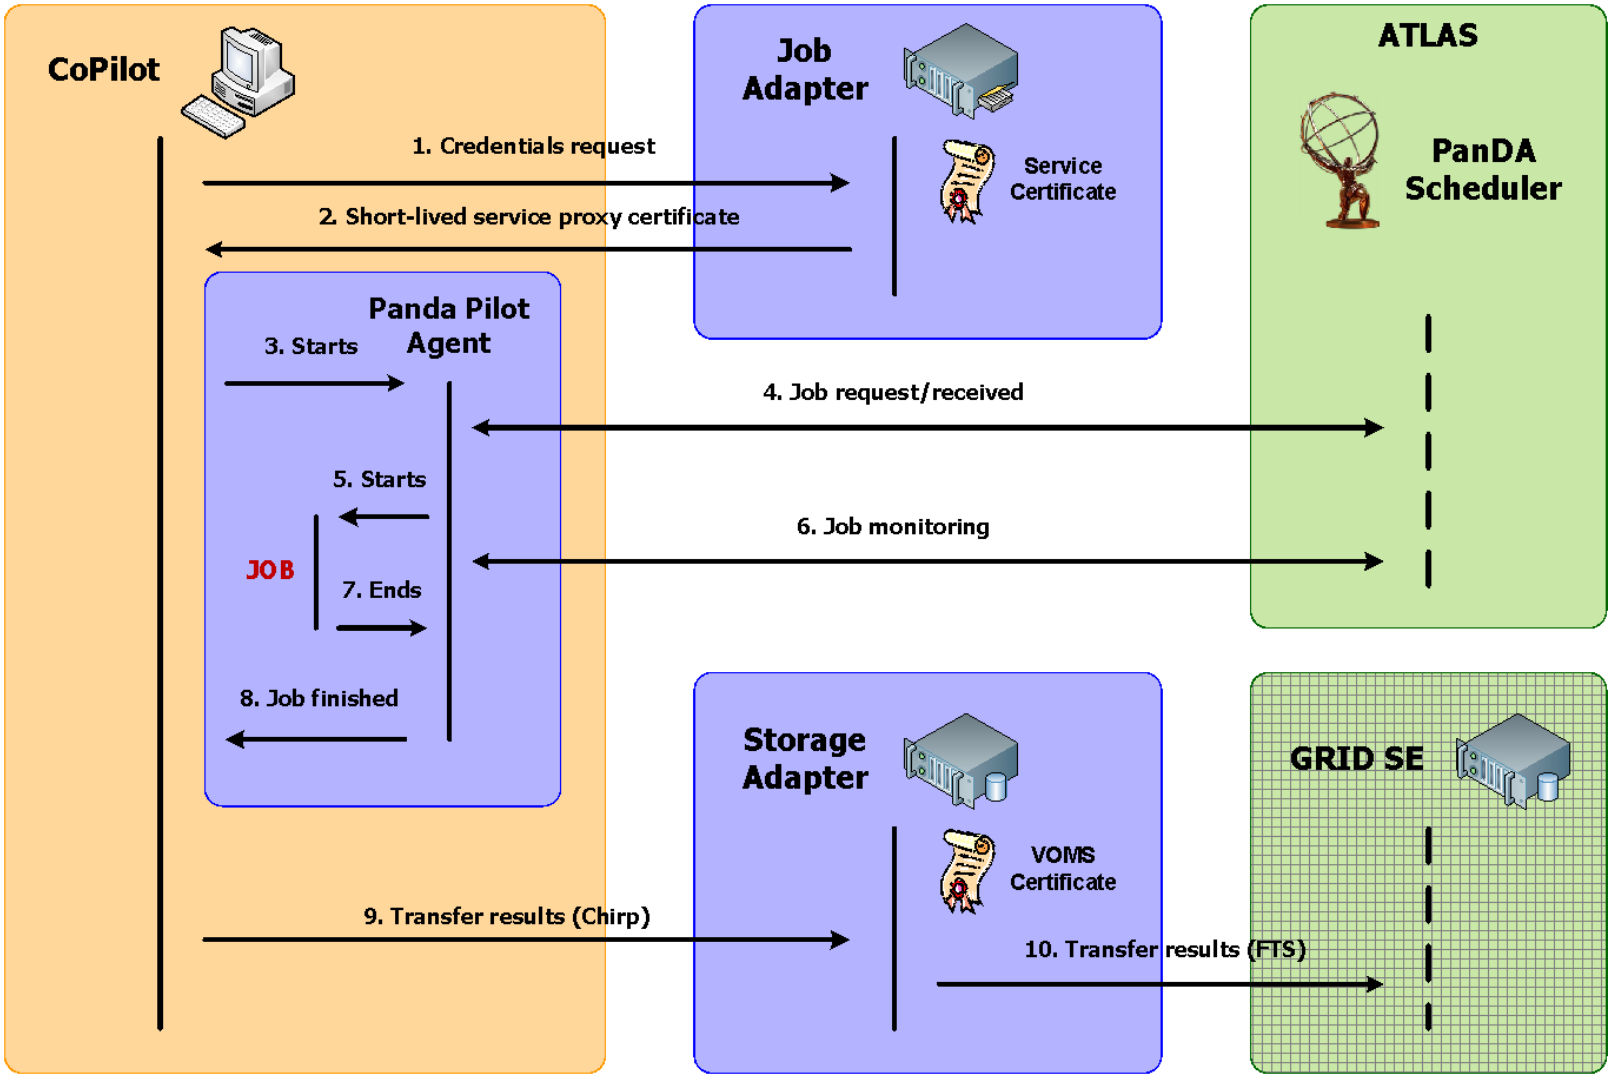

For the detailed description of the communication protocol please see:
https://cernvm.cern.ch/project/trac/cernvm/wiki/CoPilotProtocol

Building a Volunteer Cloud

- **Final Summary:**

- **Solved porting problem to all client platforms:**
- **Solved image size problem**
- **Solved job production interface problem**

- **All done without changing existing BOINC infrastructure (client or server side)**
- **All done without changing physicists' code or procedures**

- **We have built a "Volunteer Cloud" …**