

CernVM Co-Pilot: an Extensible Framework for Building Scalable Computing Infrastructures on the Cloud

Artem Harutyunyan¹
Jakob Blomer¹
Predrag Buncic¹
Ioannis Charalampidis¹
Francois Grey²
Anton Karneyeu³
Dag Larsen¹
Josip Lisec¹
Daniel Lombrana González²
Ben Segal⁴
Peter Scands³

¹CERN PH/SFT

²CCC

³CERN PH/TH

⁴CERN IT



INTRODUCTION AND ARCHITECTURE

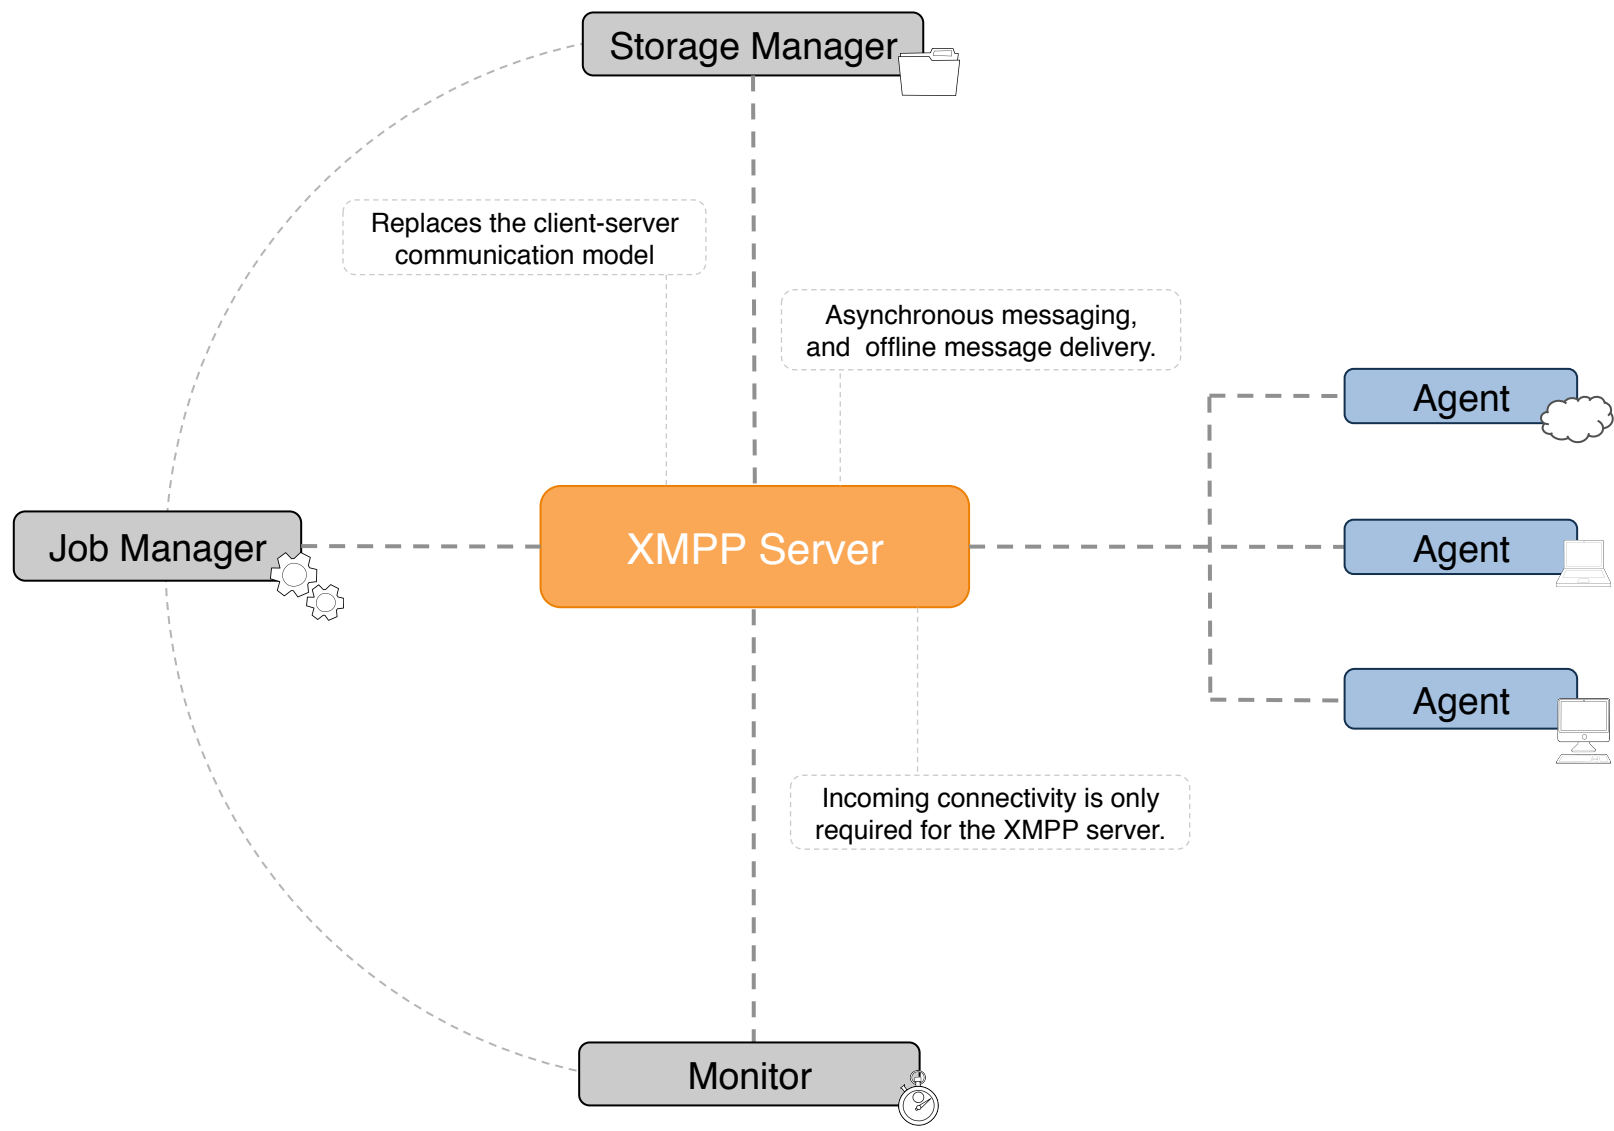


What is CernVM Co-Pilot?

- CernVM Co-Pilot is a framework for running jobs on both **unmanaged resources** (such as machines of volunteers), as well as on **managed cloud resources** (such as Amazon EC2).
- Can be used to instantiate a stand alone computing infrastructure.
- Can be used to integrate new resources into existing infrastructures (such as Grid or Batch).

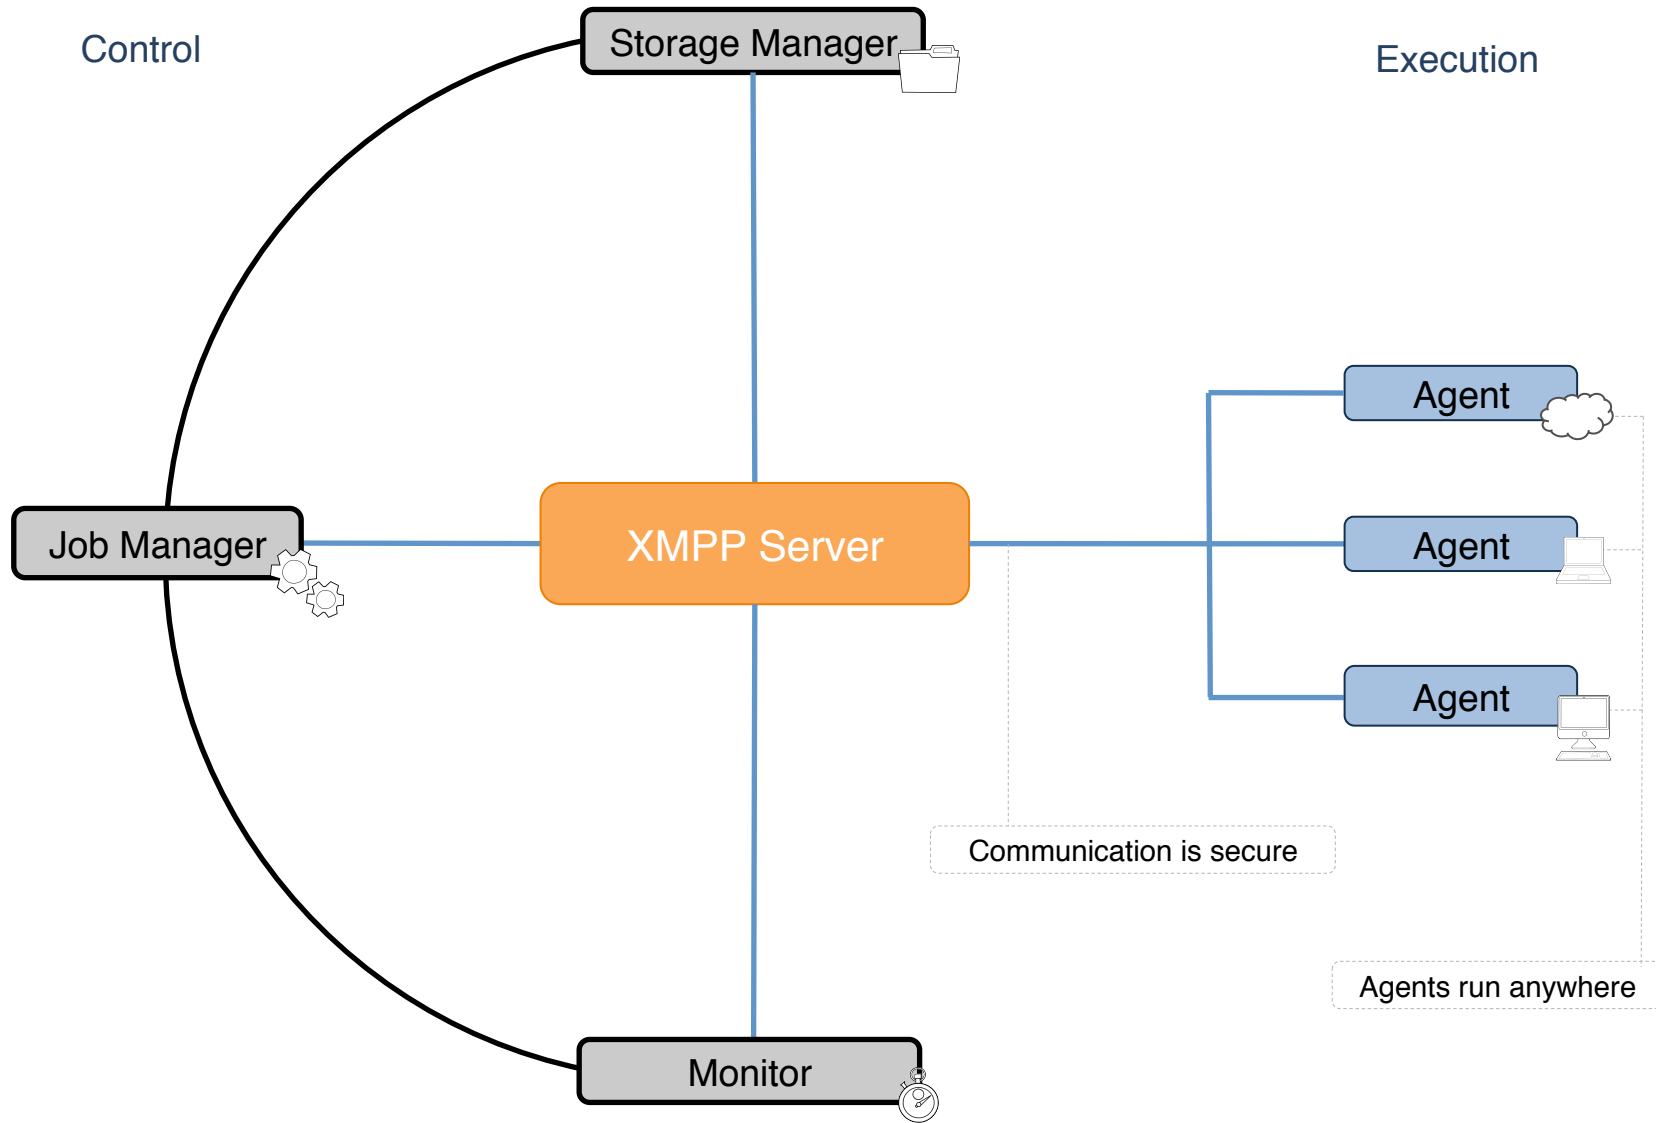


Braking traditional Client-Server model

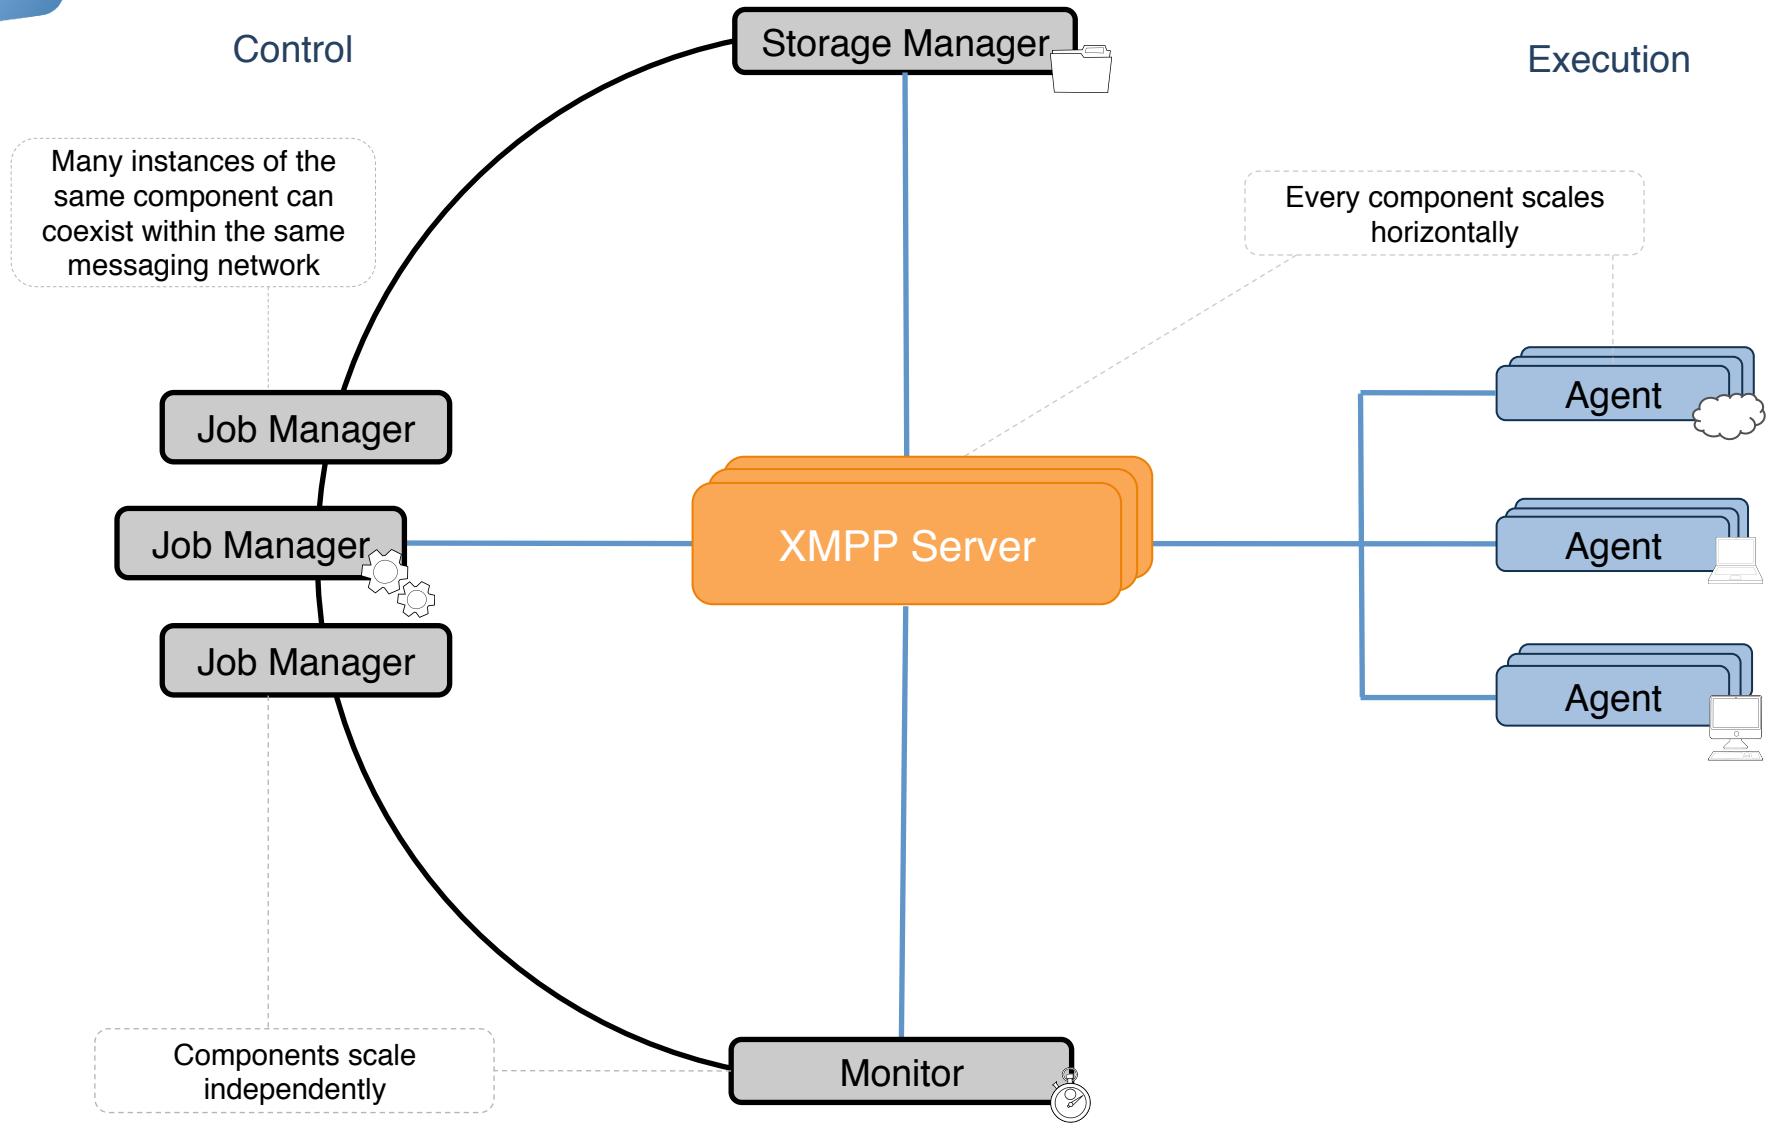




No 'site services' required

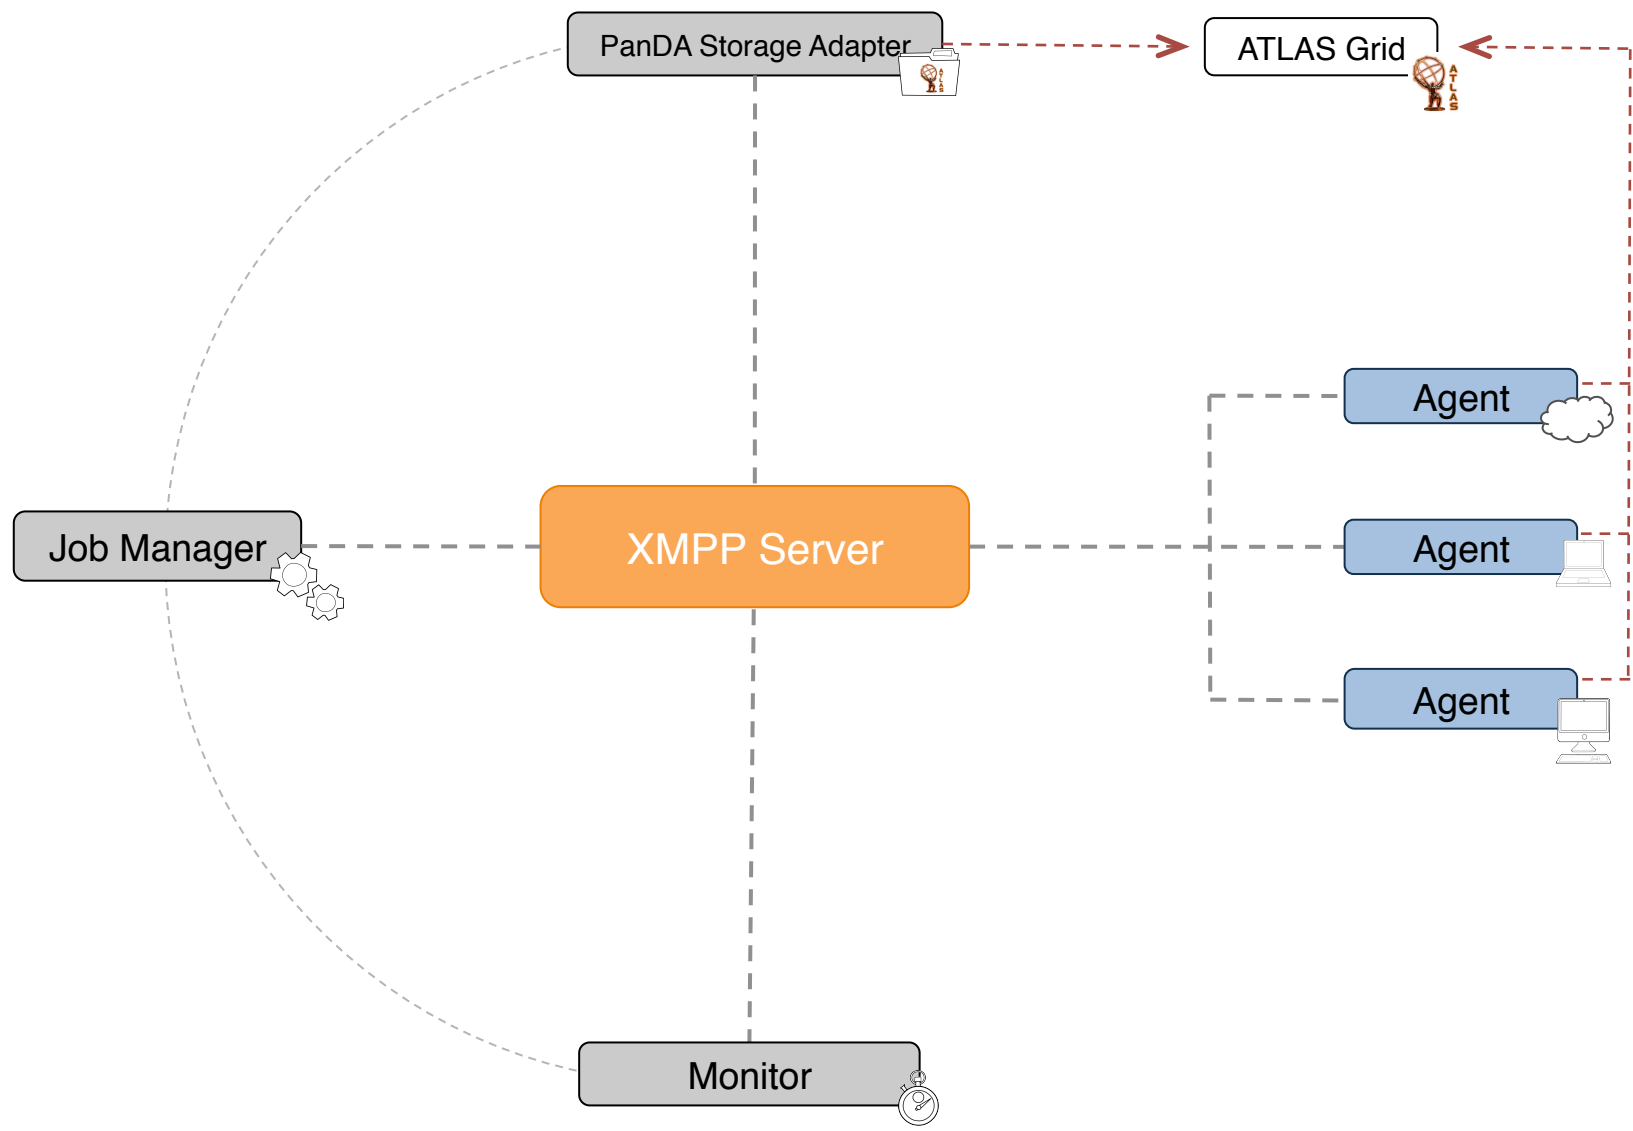


Scalability



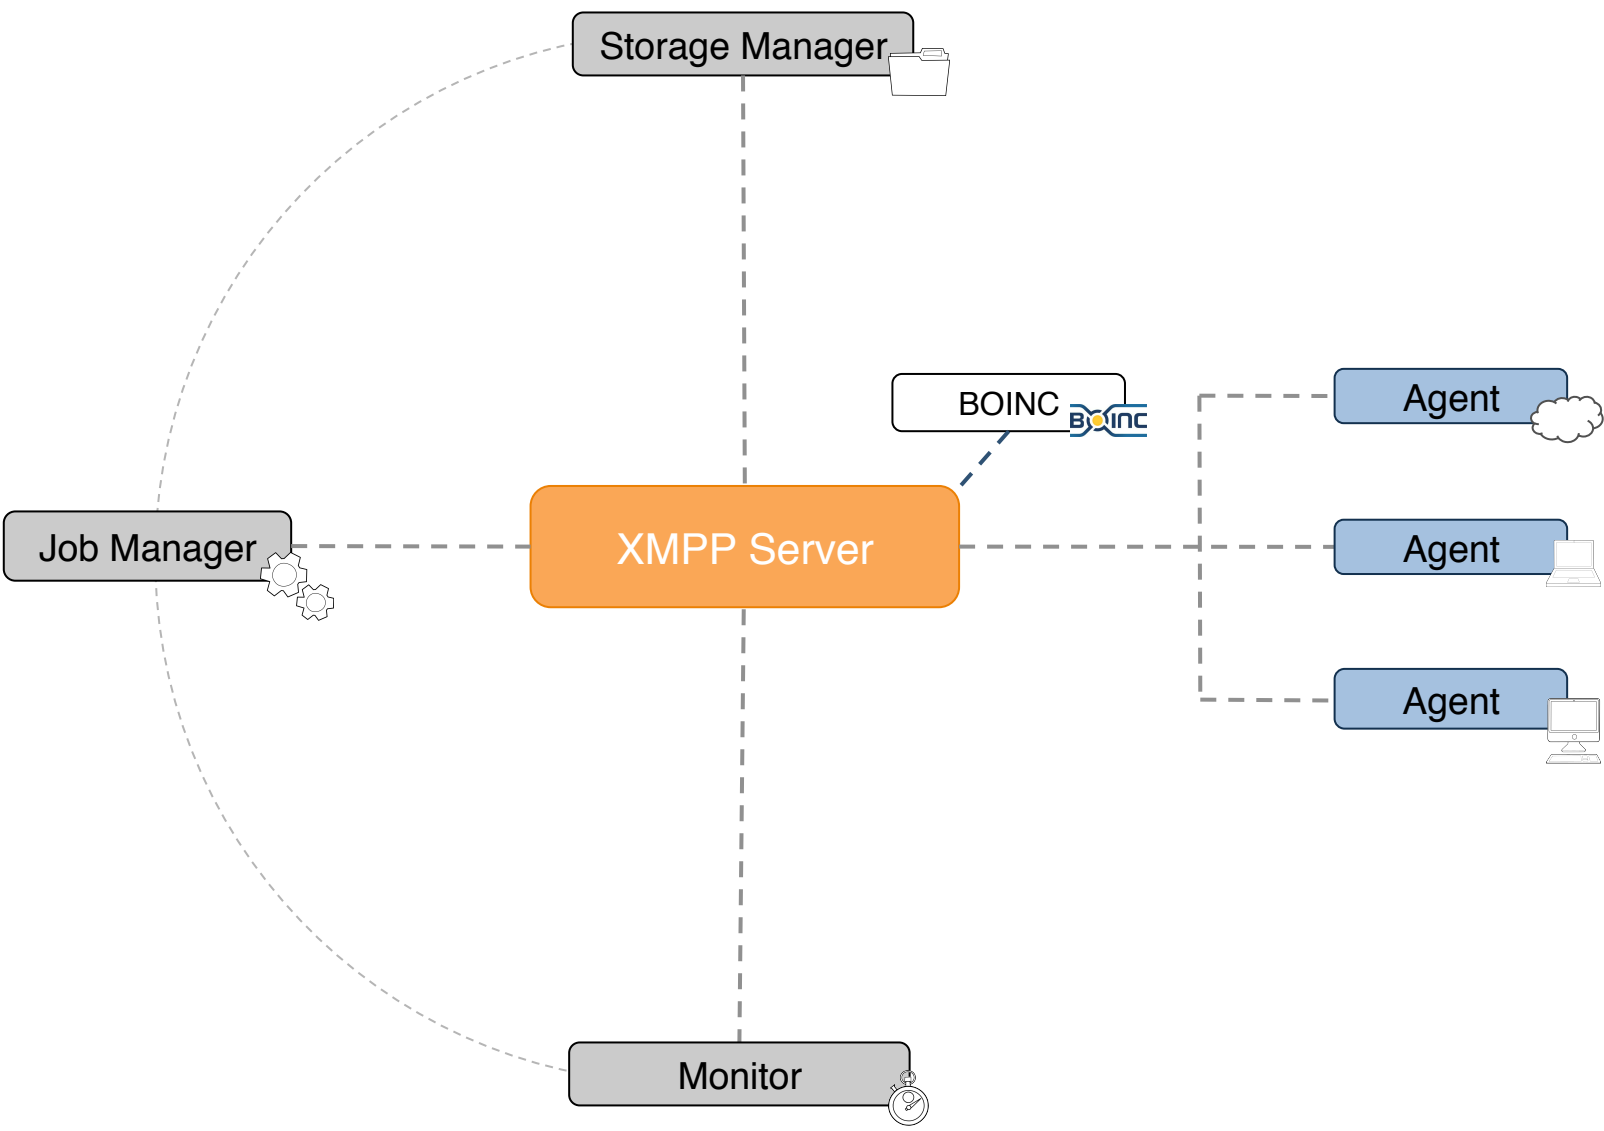


Flexibility and extensibility: ATLAS





Flexibility and extensibility: LHC@home 2.0





Simple and easily extensible protocol

- Components exchange XML formatted messages over Jabber/XMPP
- The protocol currently consists of 10 commands for job management, and 7 commands for monitoring and status queries*

```
<message to='mon@jabber.org'
  from='jm@jabber.org/alpha'
  noack='1'>
  <body>
    <info command='_BASE64:cmVwb3J0RXZlbnQ='
      component='_BASE64:am9ib...scGhh'
      event='_BASE64:c2Vzc2lvbi5zdGFydA==' />
  </body>
</message>
```

- New components can be implemented in virtually any programming language

* Full specs available at <http://cernvm.cern.ch/portal/copilot/documentation>



CURRENT STATE



Off the shelf components

- Documentation
 - <http://cernvm.cern.ch/portal/copilot/documentation>
- Packages
 - Generic
 - copilot
 - copilot-jobmanager-generic
 - copilot-dashboard
 - copilot-monitor
 - copilot-heartbeat
 - copilot-agent
 - copilot-key-manager
 - copilot-util
 - PanDA specific
 - copilot-storagemanager-panda
- Source code
 - <http://cernvm.cern.ch/portal/copilot/downloads>



USE CASE: LHC@home 2.0



The original LHC@home

- The original LHC@home started as an outreach project for CERN 50th Anniversary in 2004.
- Ran for several years and managed to accumulate about 75000 registered users (200 000 PCs).
- Used for calculating stability of proton orbits in the LHC.
- Unfortunately the project was somewhat stagnant during last several years.



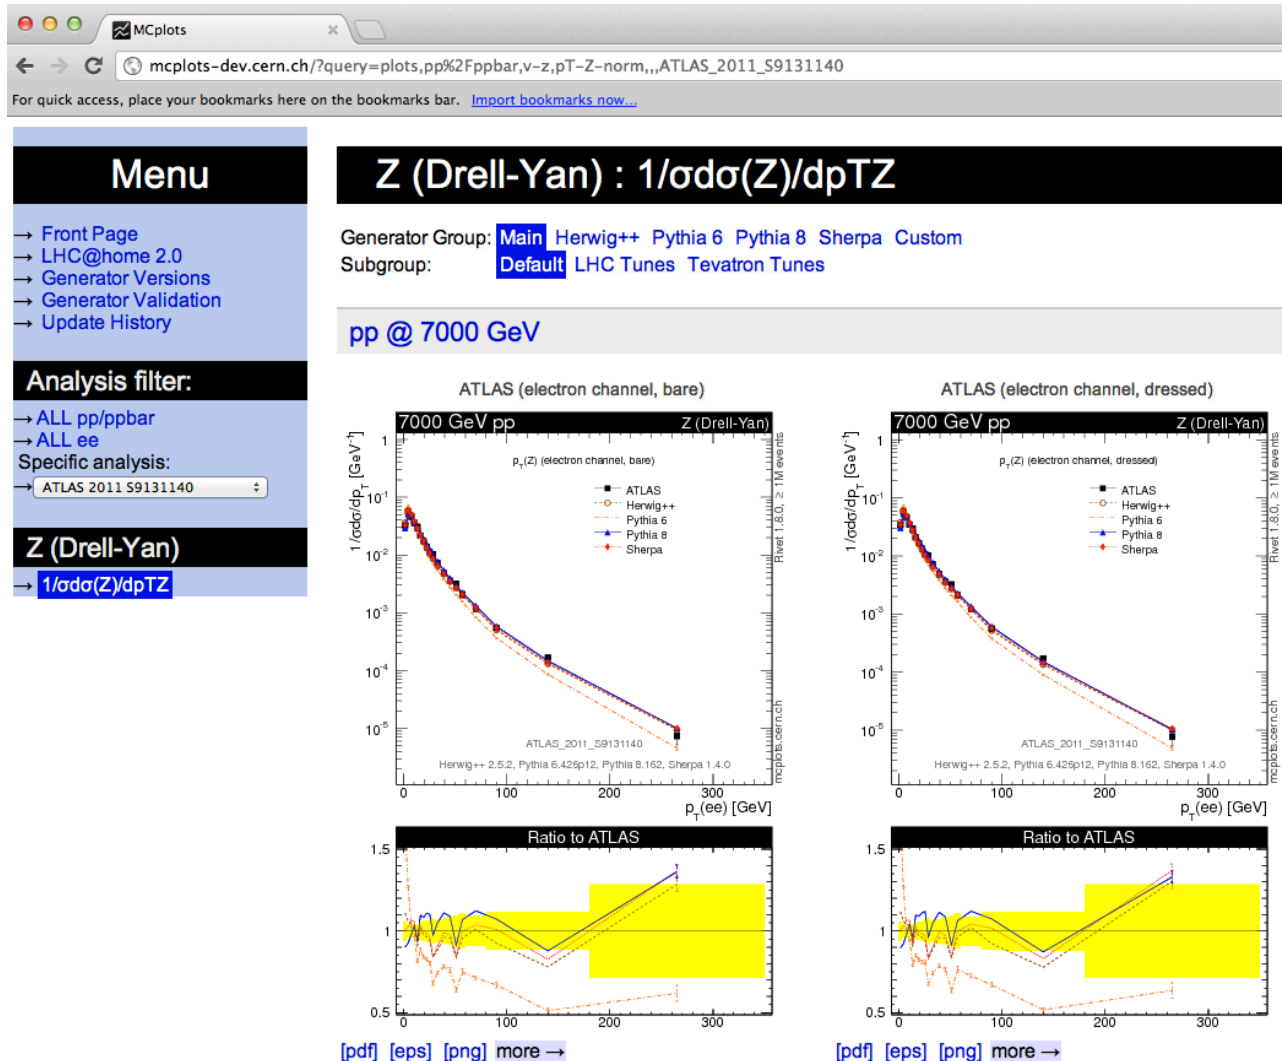
The new LHC@home 2.0

- CernVM Co-Pilot (along with other components developed within CernVM) made it possible to give new life to the LHC@home project.
- In November 2010 the resurrected project (called LHC@home 2.0) started in Alpha testing mode (private, invitation only).
- LHC@home 2.0 is now running a Monte-Carlo simulations for CERN Theory group.



LHC@home 2.0: results of calculations

- Results of calculations are used to populate the MCPlots web site - a browsable repository of Monte Carlo plots comparing High Energy Physics event generators to a wide variety of experimental data (<http://mcplots.cern.ch>)



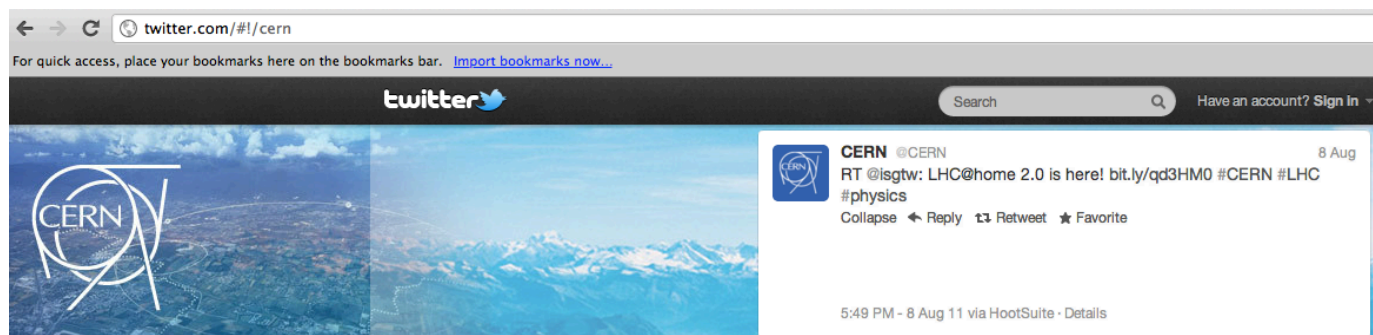
- CERN press release August 2011

<http://press.web.cern.ch/press/pressreleases/Releases2011/PR13.11E.html>



- CERN twitter

<http://twitter.com/#!/CERN/status/100594401251823617>

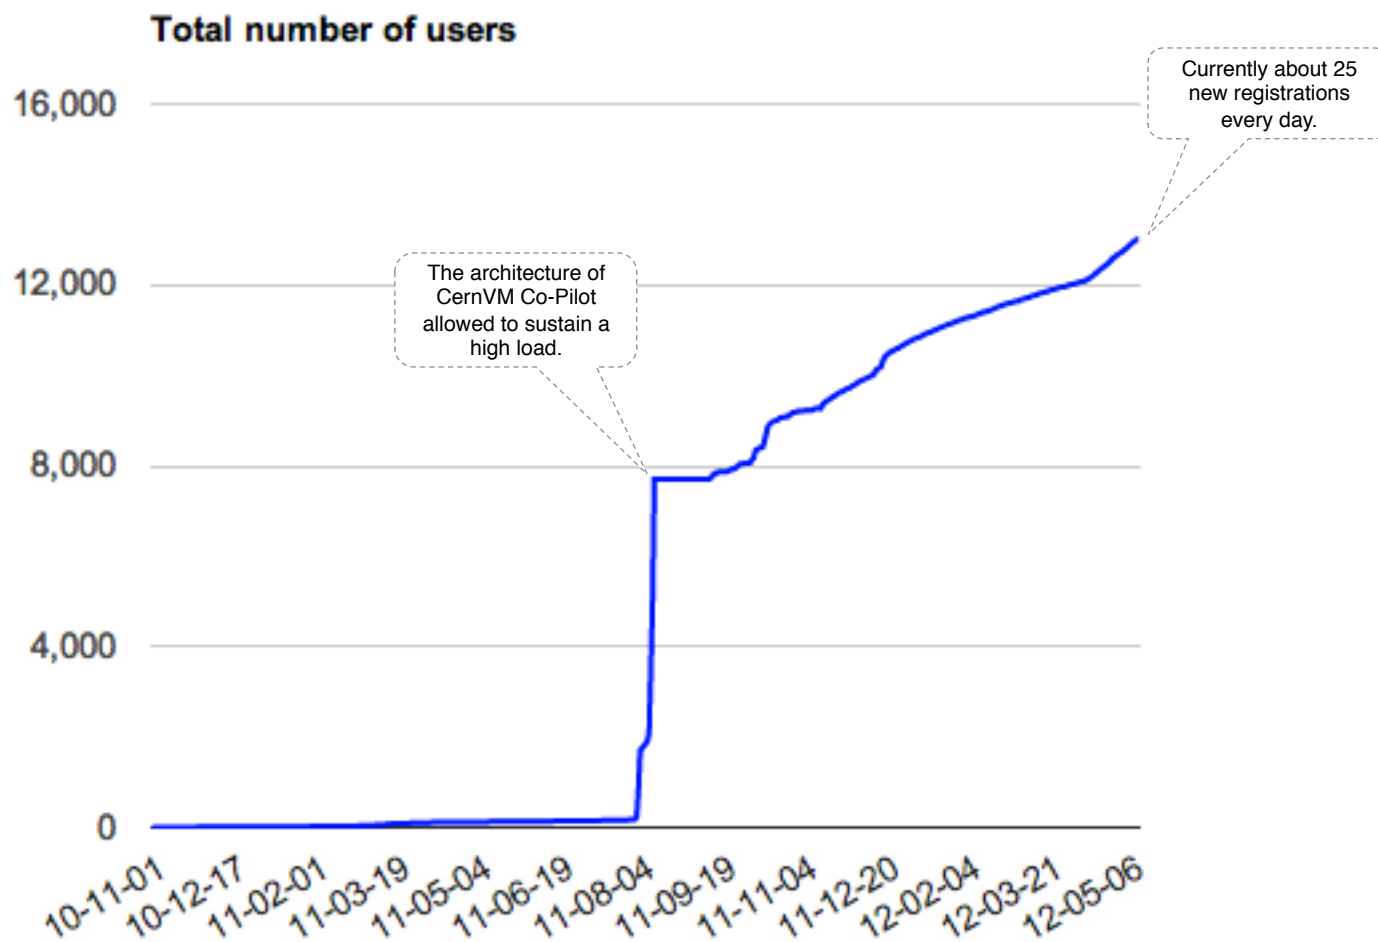


- The next day


<http://lhathome2.cern.ch/media>



Observed an (unexpected) growth of number of users (from ~100 to ~8000 in 3-4 days)



Number of machines



Where CernVM was booted (before the press release)



Where CernVM was booted (after the press release)






Current stats (as of May 2012)

Processing ~ 100000 ($\pm 10\%$) jobs per week (a main computing resource for MCPlots)

Jobs rates



Example week (16.04.2012 - 22.04.2012)

Executed:


- 96240 'good' jobs (exit status 0)
- 4688 'bad' jobs (non zero exit status)

Spent:

- ~ 105000 CPU hours per week for good jobs
- ~ 33 CPU hours per week for bad jobs

Contributed CPU time

CPU time contributed by volunteers for jobs processing






Current stats (as of May 2012)

Processing ~ 100000 ($\pm 10\%$) jobs per week (a main computing resource for MCPlots)

Jobs rates



Example week (16.04.2012 - 22.04.2012)

Executed:


- 96240 'good' jobs (exit status 0)
- 4688 'bad' jobs (non zero exit status)

Equivalent to:

~ 678 machines crunching 24/7 for a week

Connected machines

Number of volunteers machines connected to copilot host



- Co-Pilot central services are hosted on 5 (small) virtual machines in CERN PH/SFT Group (about to be migrated to CERN IT).
- How much it would cost to rent 600 machines on Amazon EC2 (about 2.750.000 **good CPU/hours**) and execute the same workload?

Lower limits**:

Instance type	Bootstrap price	Total price
On-demand Micro	-	55 kUSD
On-demand Small (default)	-	220 kUSD
On-demand Medium (High CPU)	-	$450 / 2 = 225$ kUSD
Reserved*** Small	117 kUSD	$117 + 44 = 161$ kUSD
Reserved*** Medium	237 kUSD	$237 + 44 = 281$ kUSD
Reserved*** Medium (High CPU)	300 kUSD	$300 + 55 = 355$ kUSD

* <http://aws.amazon.com/ec2/pricing/> (as of May 16th, 2012).

** Only the CPU time is taken into account, the price for traffic and storage is neglected.

*** 1-year reservation of 600 'Heavy Utilization Reserved Instances'.





Summary

- Developed a generic framework for building an ad-hoc distributed computing infrastructure.
- Can be transparently integrated into existing computing infrastructures (e.g. ATLAS Grid)
- Individual components of the system can be scaled horizontally so the overall system can sustain a high load.
- Makes possible exploiting vast volunteer computing resources so far untapped by LHC experiments.
- Eliminated the need to modify the applications for running on unmanaged resources.
- It opens new possibilities for outreach, and promotes the image of CERN as an open institution.
- Allows public to be a part of 'something big', and engage into something that they hear about on the news.



- Do not forget to visit the CERN PH/SFT Group Booth in Kimmel Center (right in front of coffee table on 4th floor)
- To learn more about CernVM File System
 - ‘Status and Future Perspectives of CernVM-FS’ (by Jakob Blomer, 05.22.2012 at 14:20 – 14:45, Room 802, Kimmel Center)
- To learn more about the CERN Virtual Machine
 - Poster 134: Managing Virtual Machine Lifecycle in CernVM Project
 - Poster 135: Long-term preservation of analysis software environment
- To learn more about the volunteer computing service at CERN
 - Poster 281: BOINC service for volunteer cloud computing