Cost-Benefit Analysis of Cloud Computing versus Desktop Grids

Derrick Kondo, Bahman Javadi, Paul Malécot, Franck Cappello
INRIA, France

David P. Anderson
UC Berkeley, USA

Cloud Background

- Vision
  - Hide complexity of hardware and software management from a user by offering computing as a service
- Benefits
  - Pay as you go
  - Scale up or down dynamically
  - No hardware management, less software management

Platform Performance vs. Costs

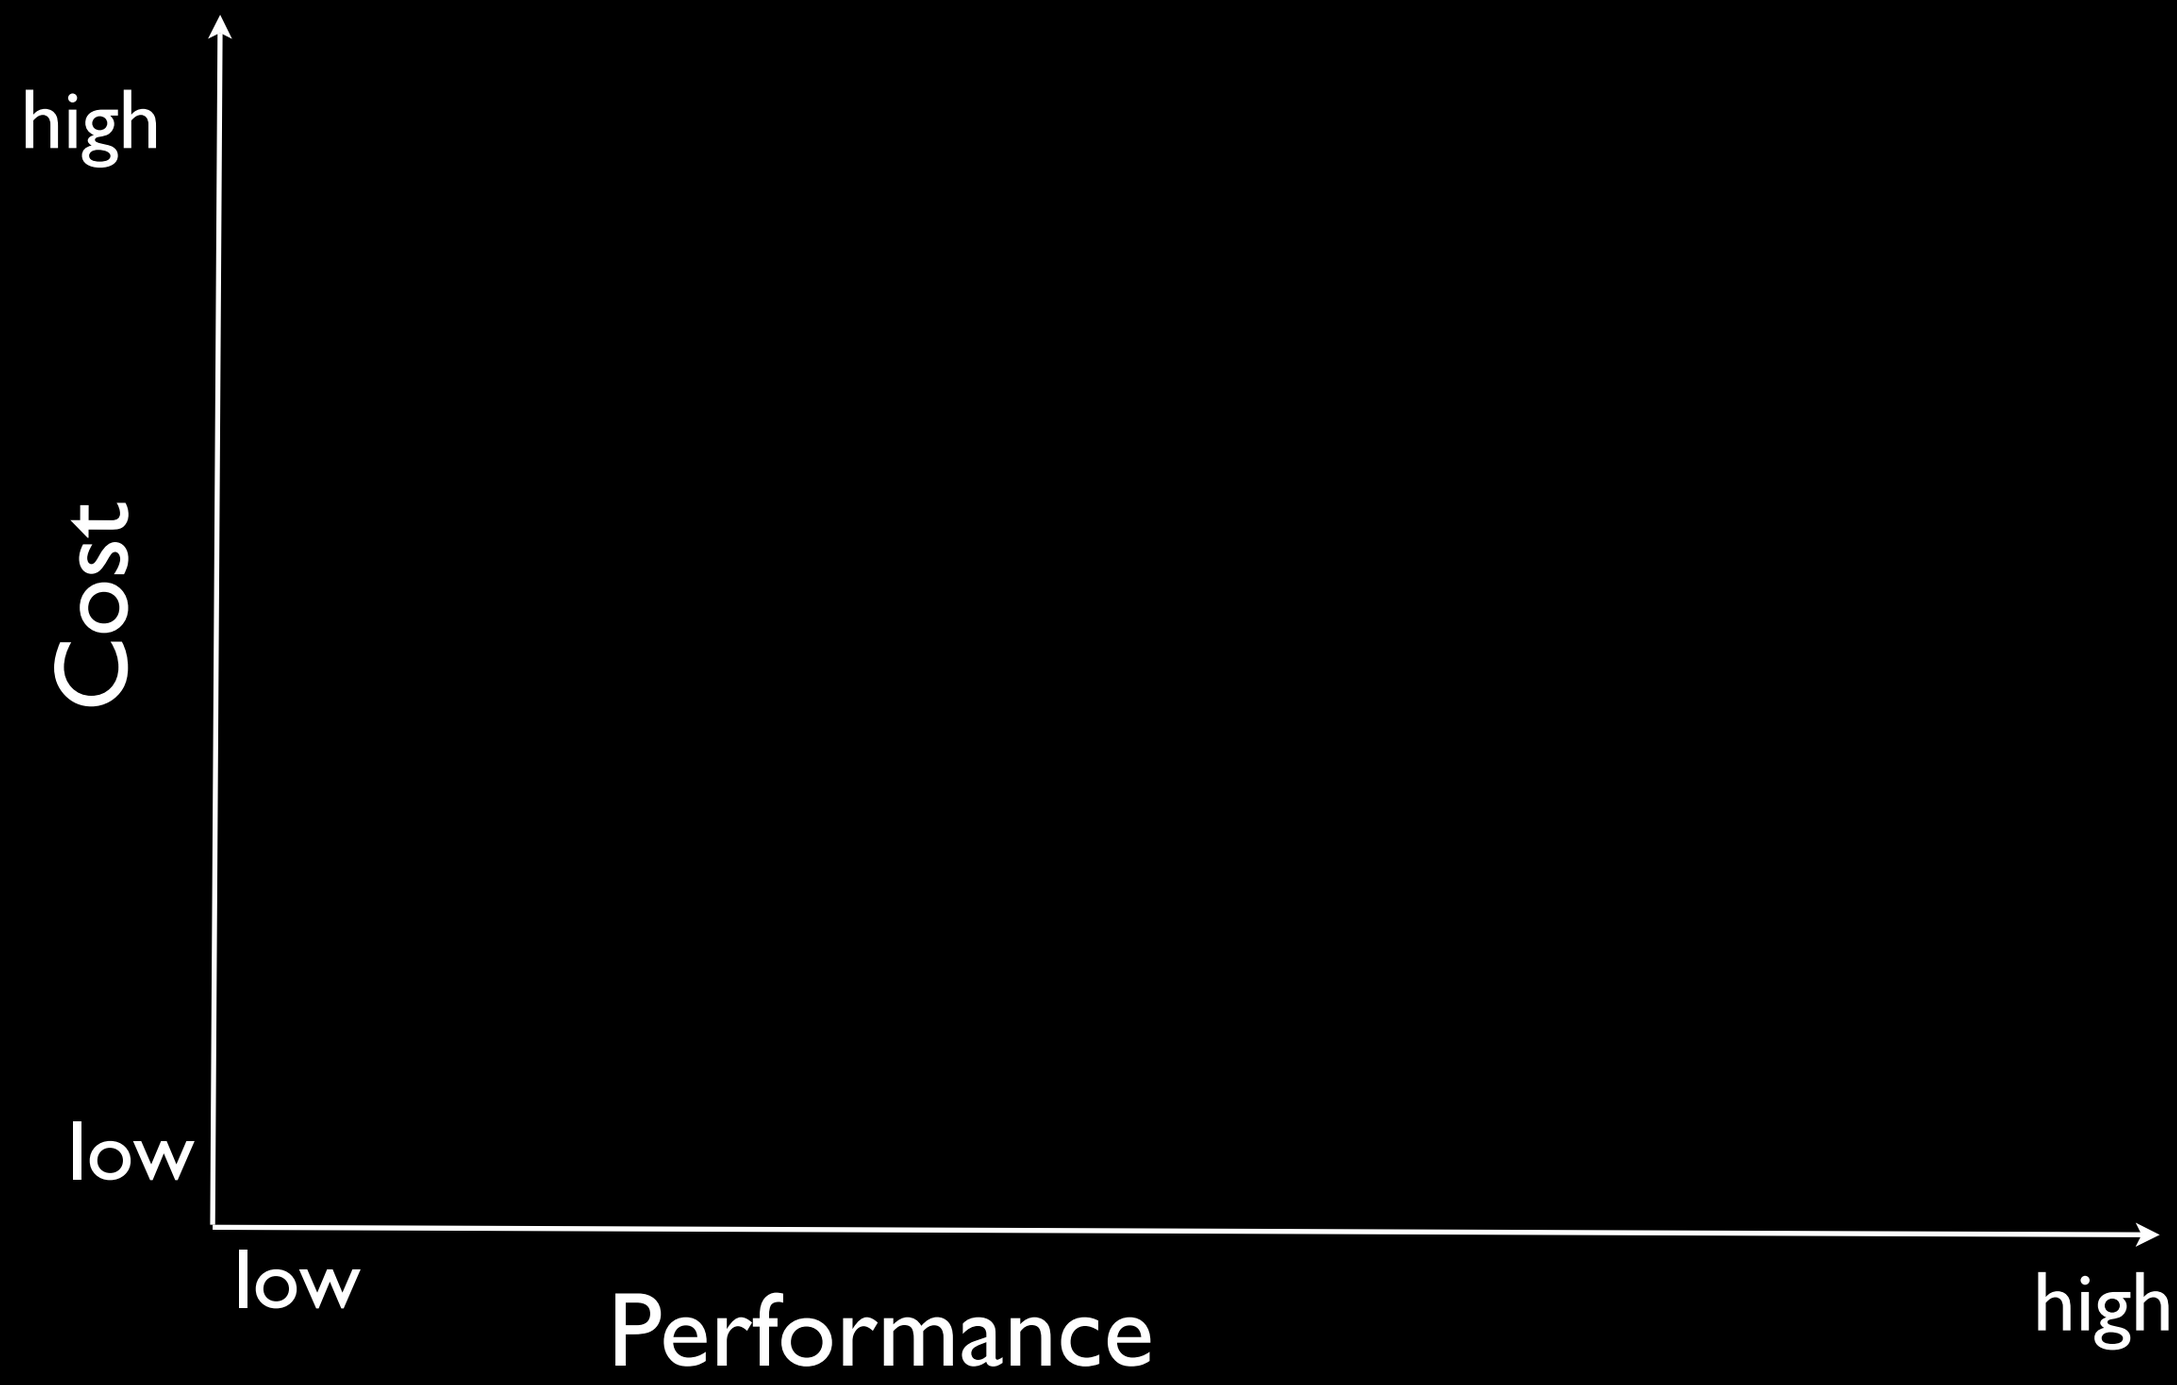

Platform Performance vs. Costs

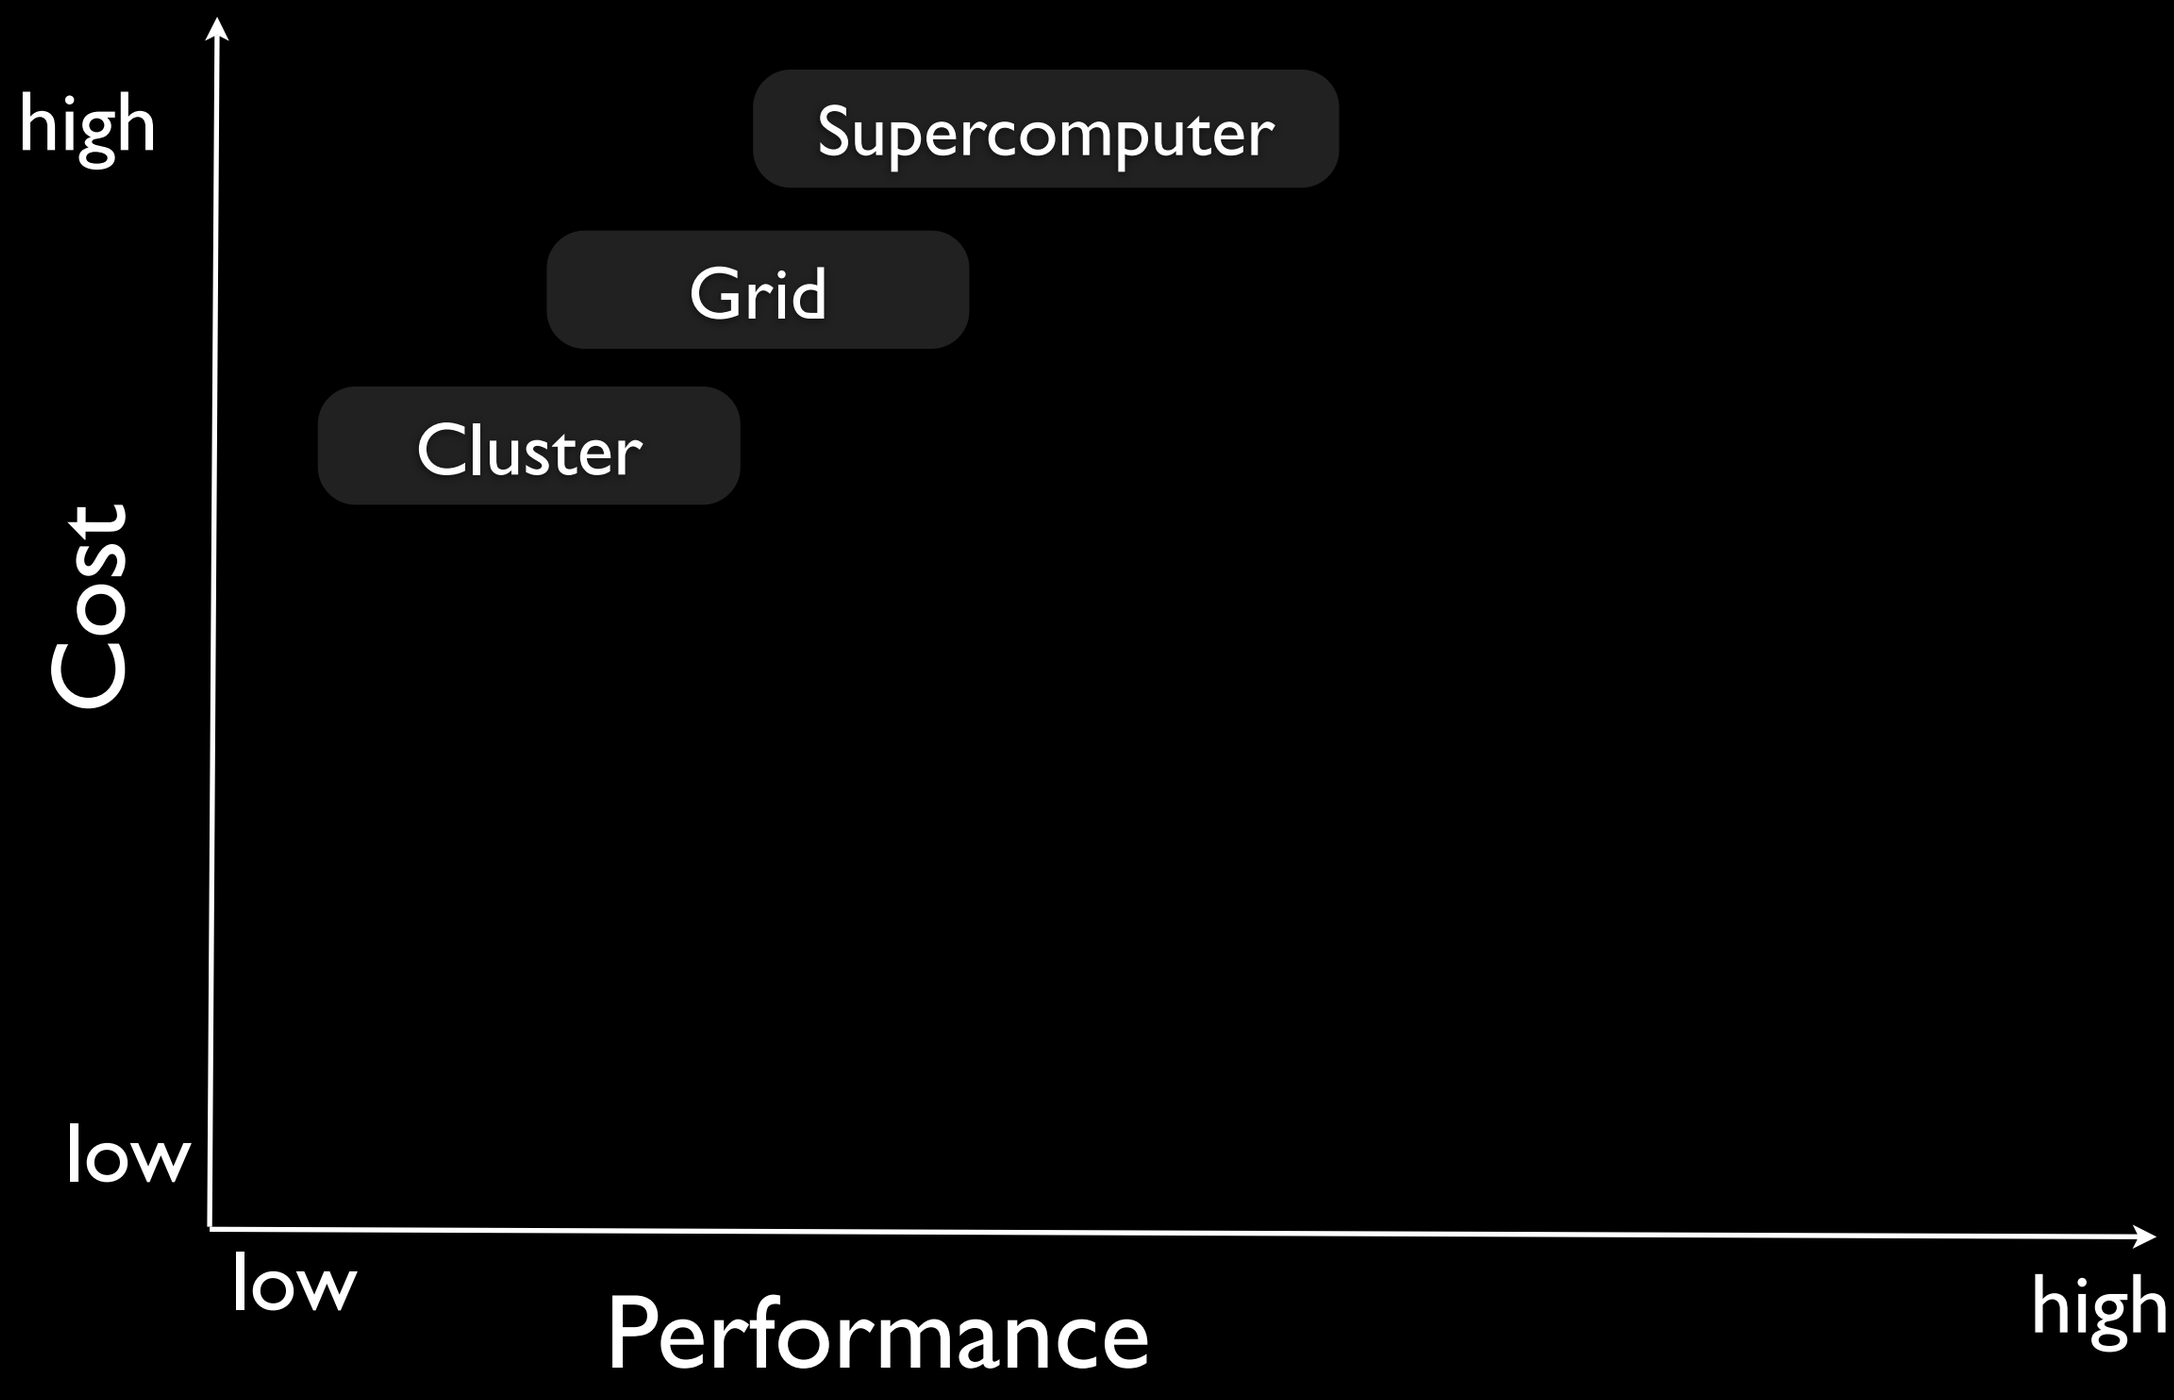

Platform Performance vs. Costs

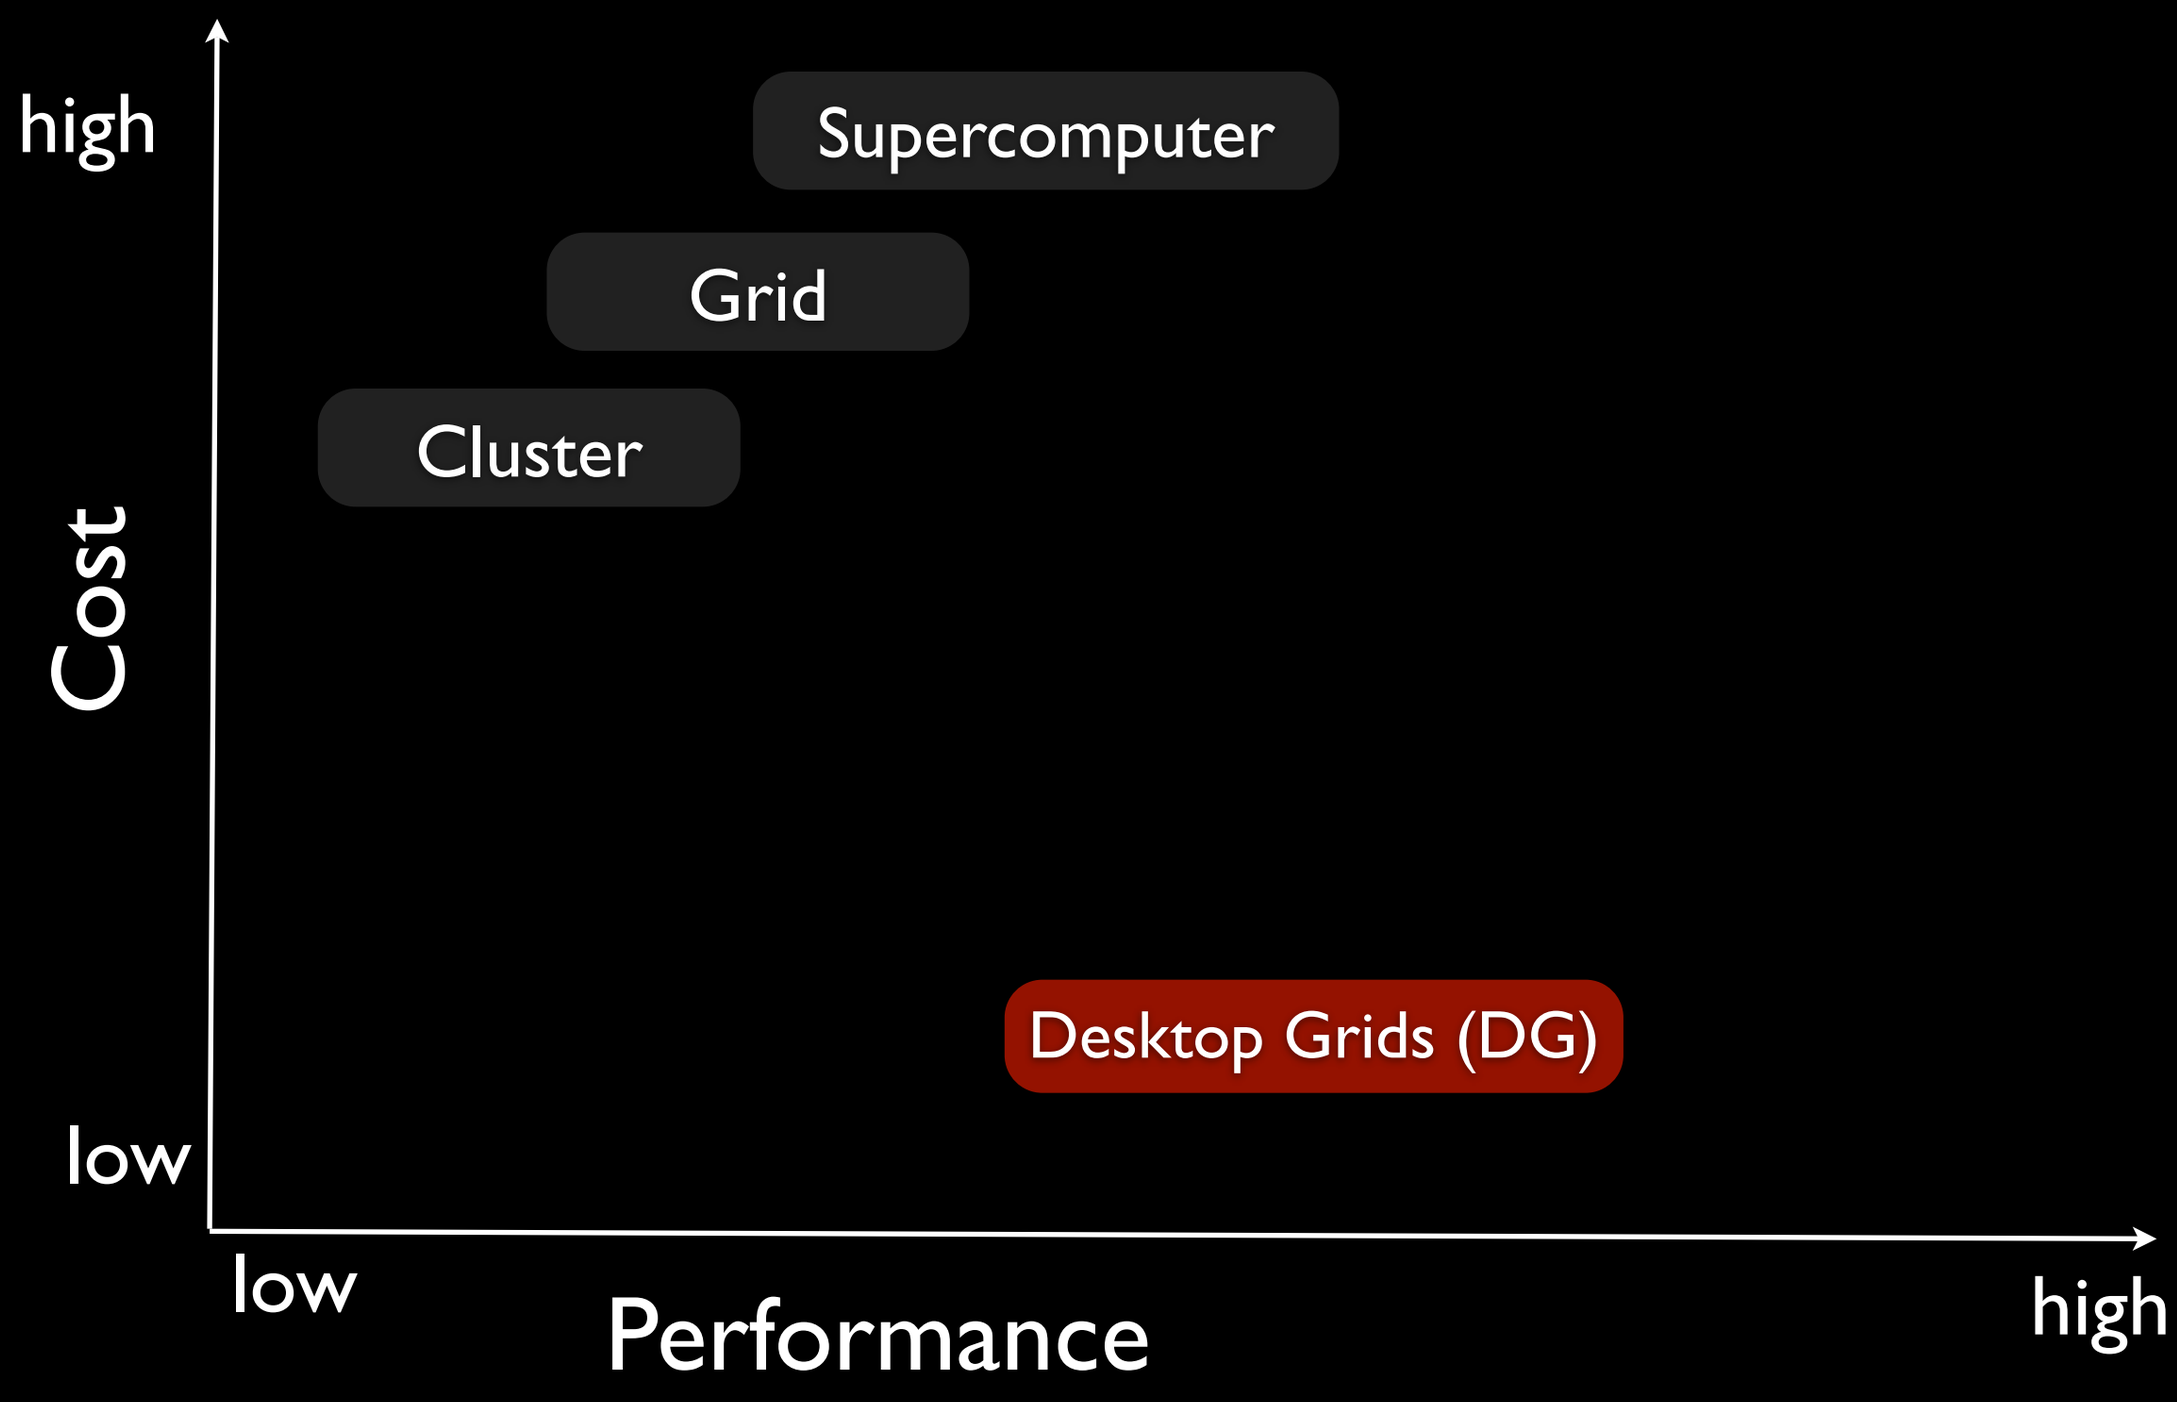

Platform Performance vs. Costs

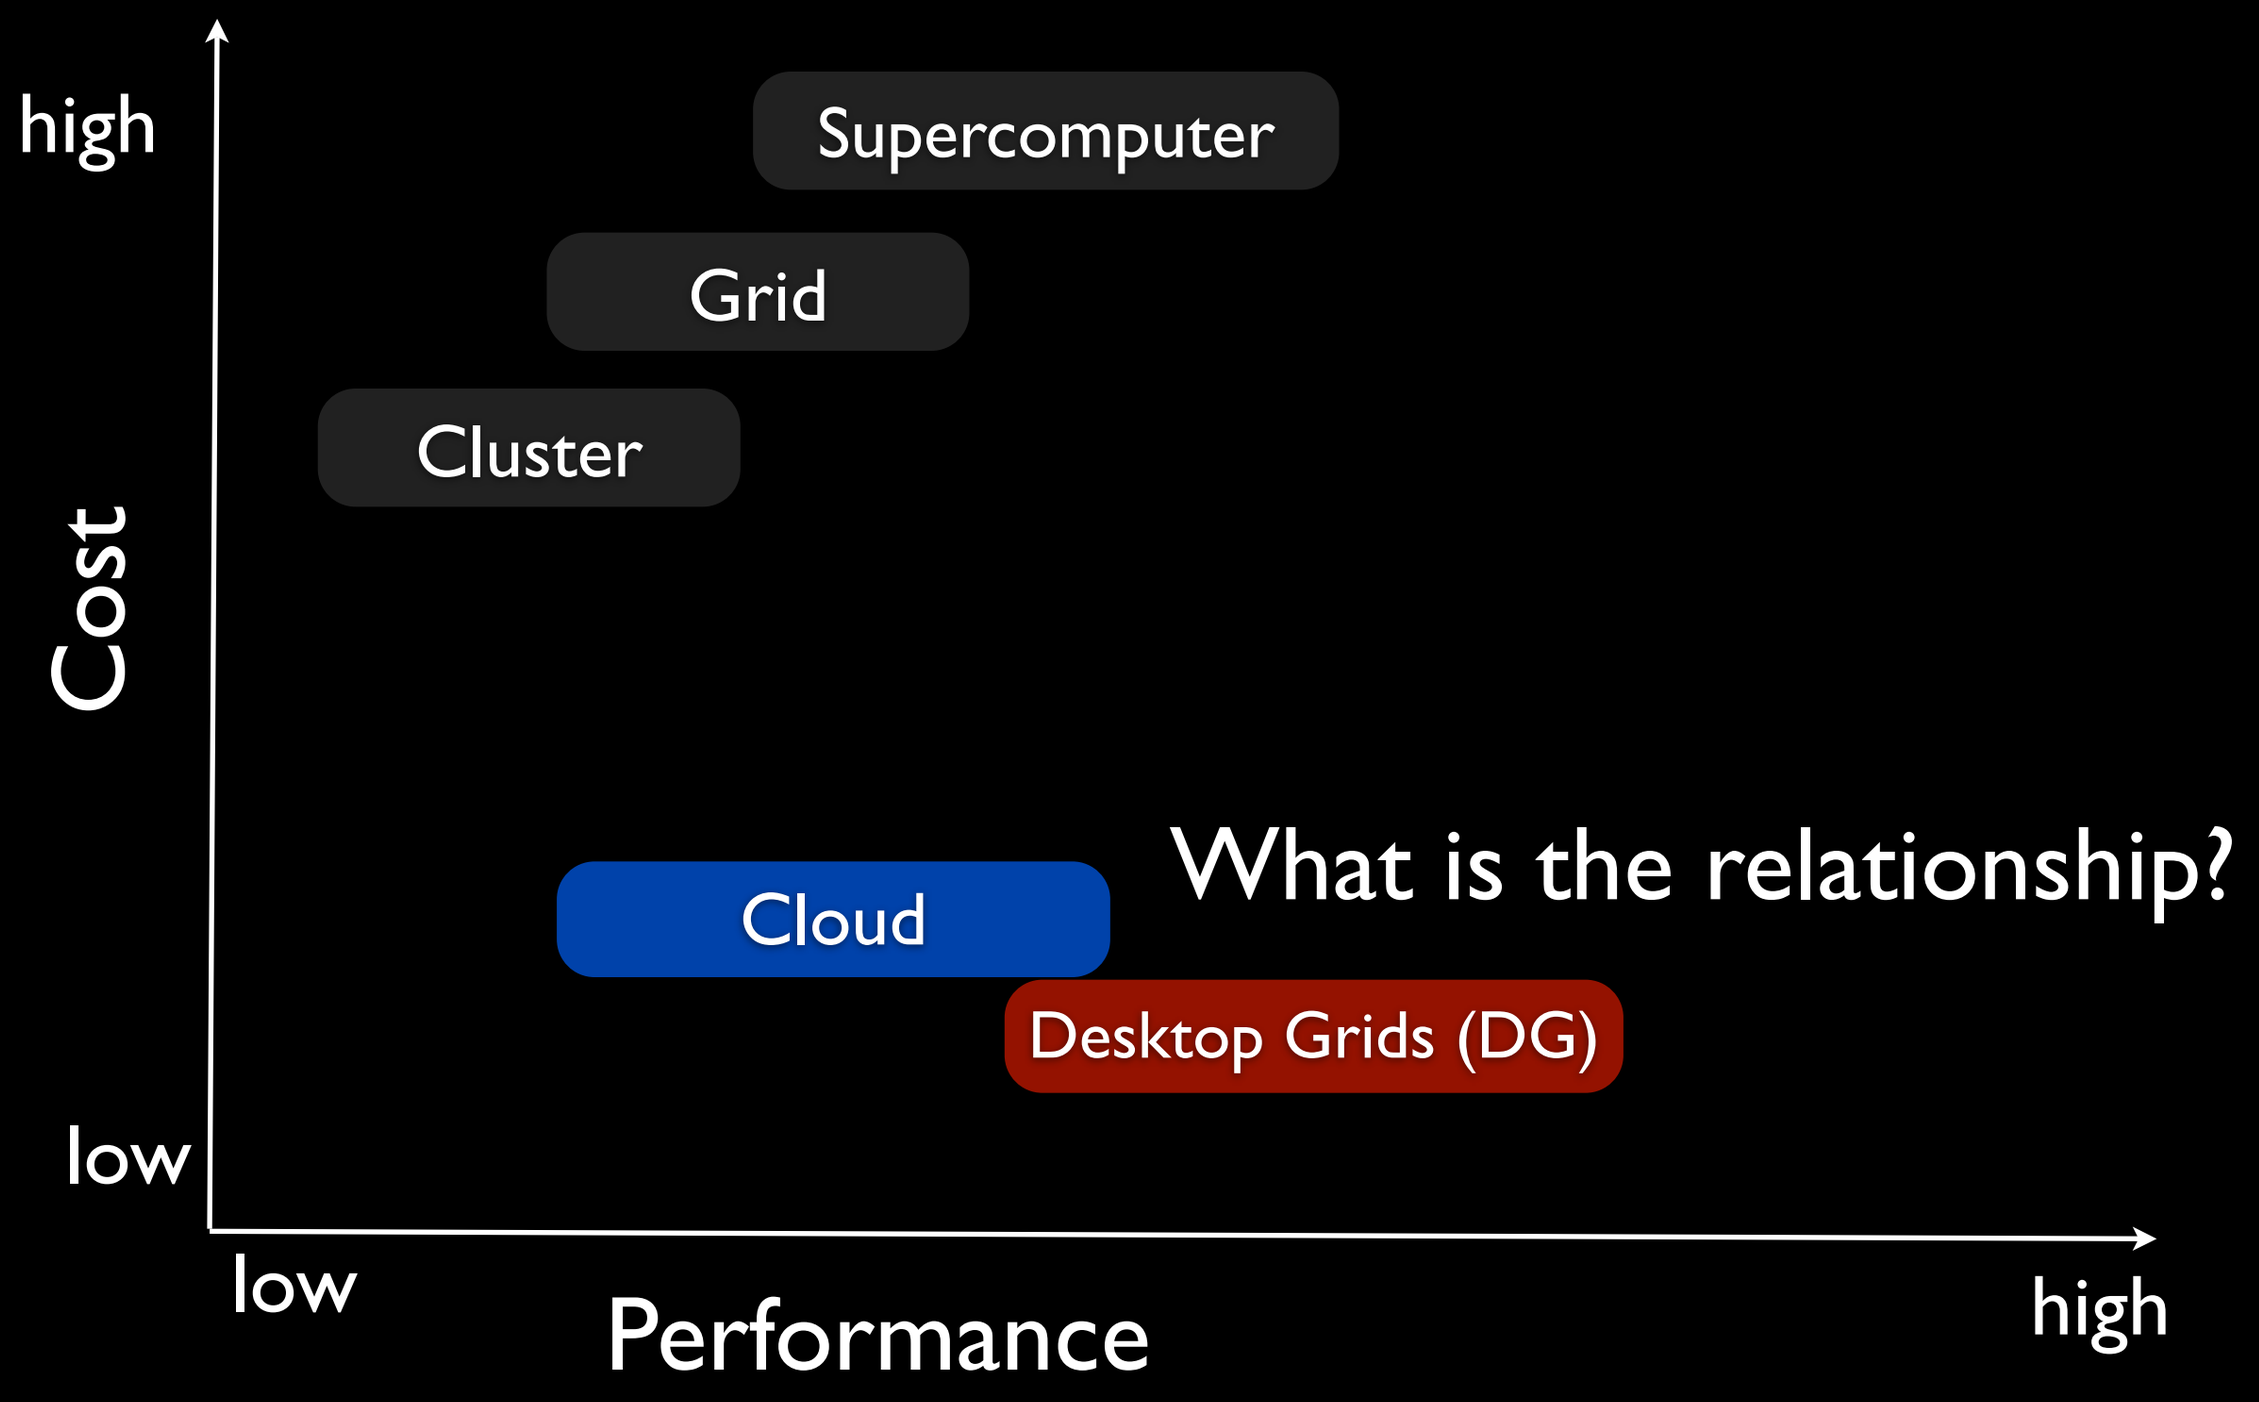

Platform Performance vs. Costs

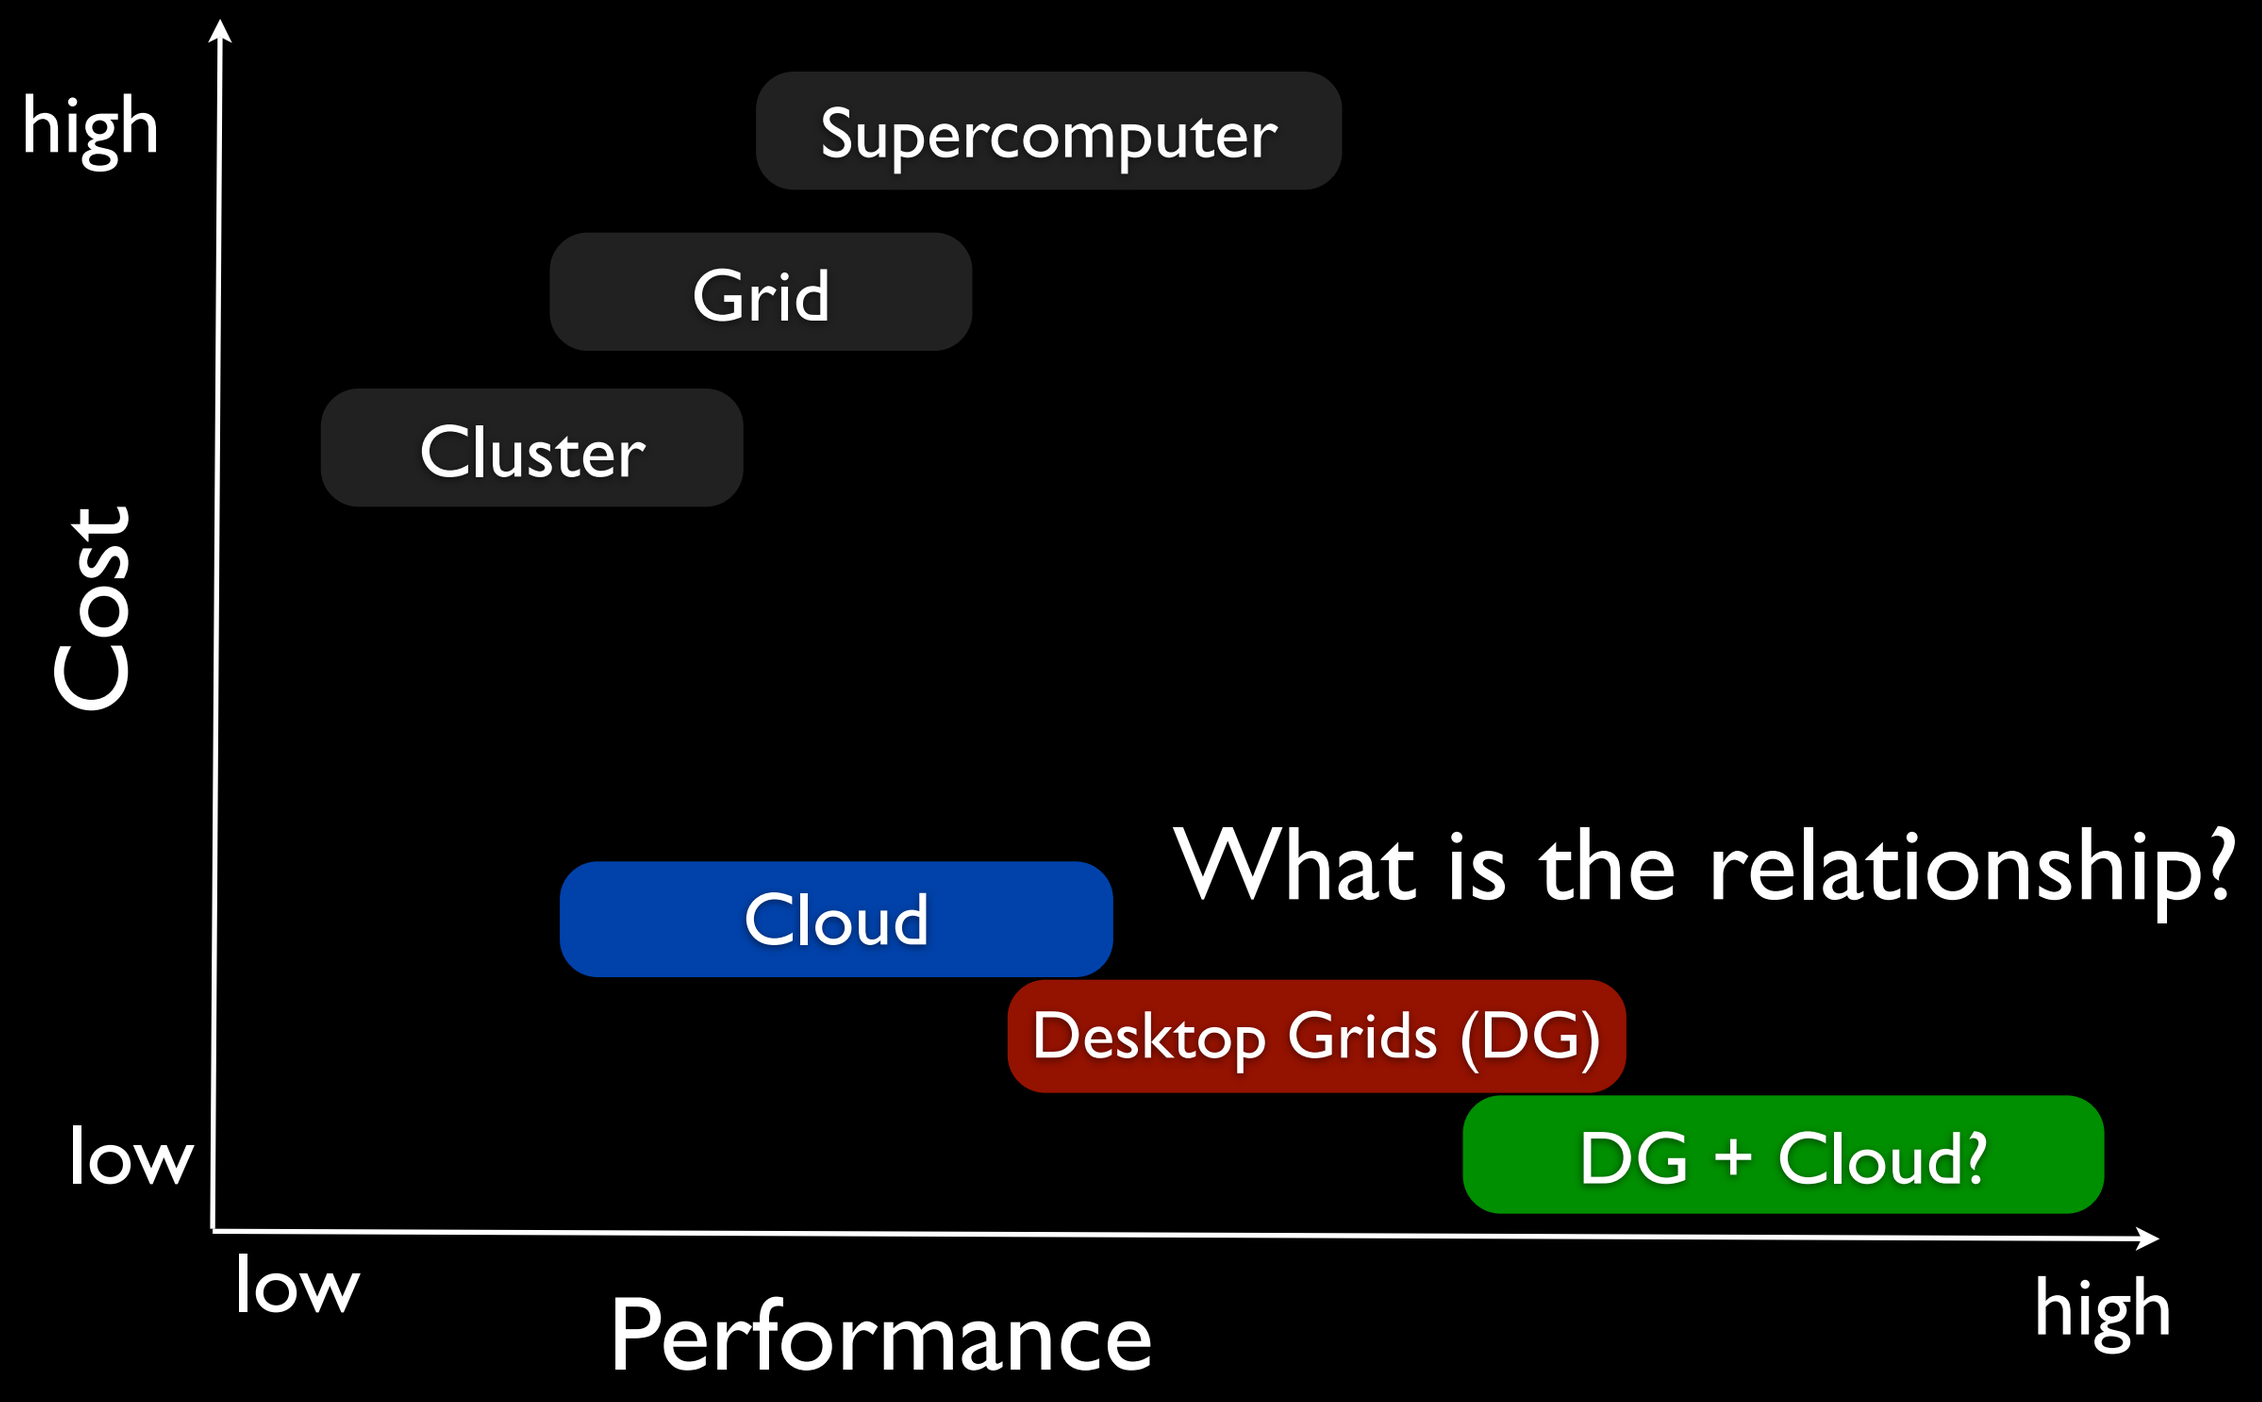

Outline

- Performance tradeoffs
- Monetary tradeoffs
- Client hosting
- Server hosting

Apples to Apples

Loosely-coupled,
high-throughput,
compute-intensive applications

Tightly-coupled,
data-intensive
real-time applications

low complexity

high complexity

Apples to Apples

Loosely-coupled,
high-throughput,
compute-intensive applications

Tightly-coupled,
data-intensive
real-time applications

low complexity

high complexity

DG's

Apples to Apples

Loosely-coupled,
high-throughput,
compute-intensive applications

Tightly-coupled,
data-intensive
real-time applications

low complexity

high complexity

DG's

Clouds

Apples to Apples

Loosely-coupled,
high-throughput,
compute-intensive applications

Tightly-coupled,
data-intensive
real-time applications

low complexity

high complexity

DG's

Clouds

Comparison assuming
embarrassingly parallel, compute-intensive applications

Method

- Use real performance measurements
  - Exported BOINC (middleware for desktop grids) project data
- Use real costs
  - Large/small BOINC projects (SETI@home / XtremLab)
  - Amazon Elastic Computing Cloud (EC2)

Stages of Project & Application

Platform Construction

Stages of Project & Application

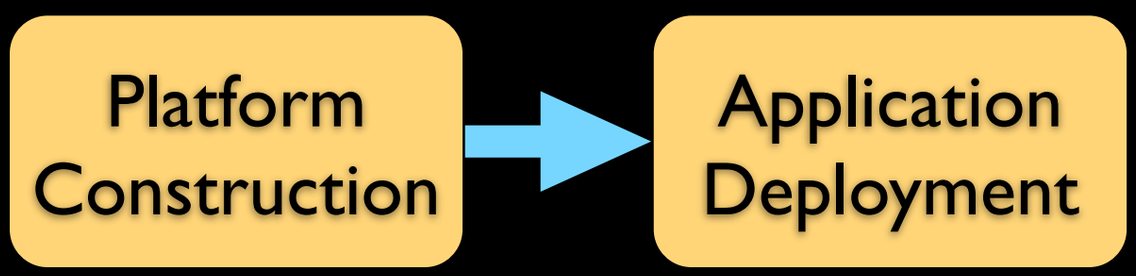

Stages of Project & Application

Stages of Project & Application

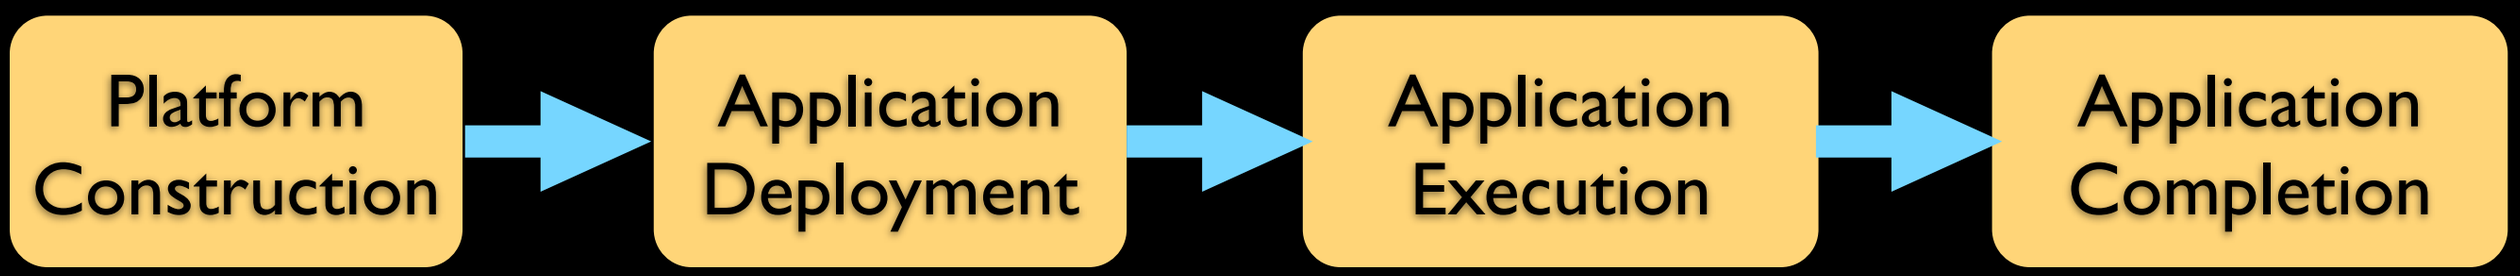

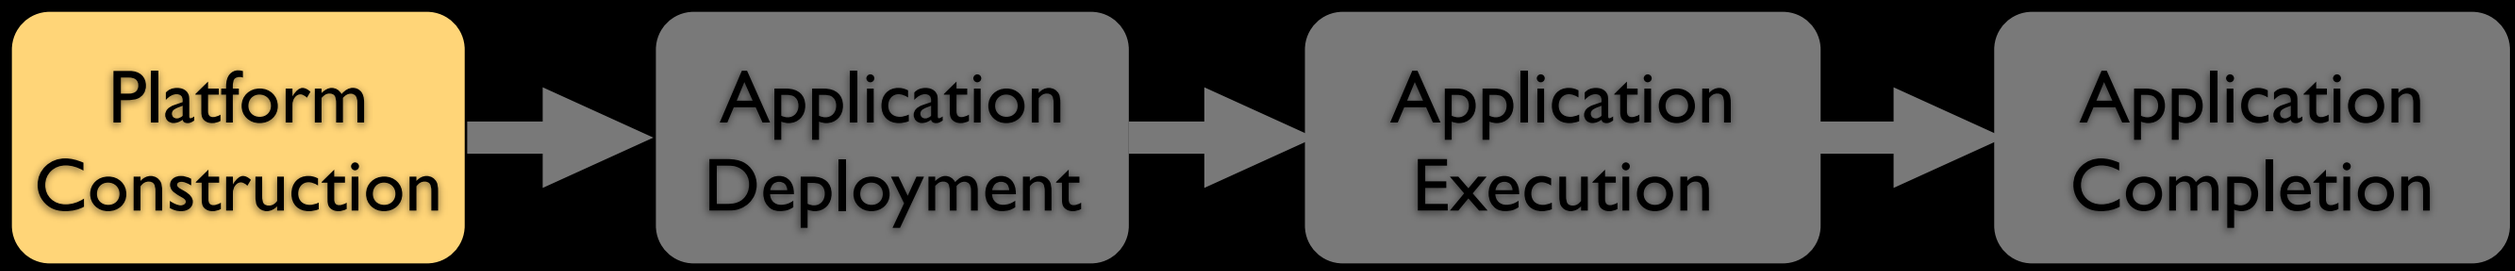

How long before I get X TeraFLOPS?

Platform Construction → Application Deployment → Application Execution → Application Completion

How long before I get X TeraFLOPS?

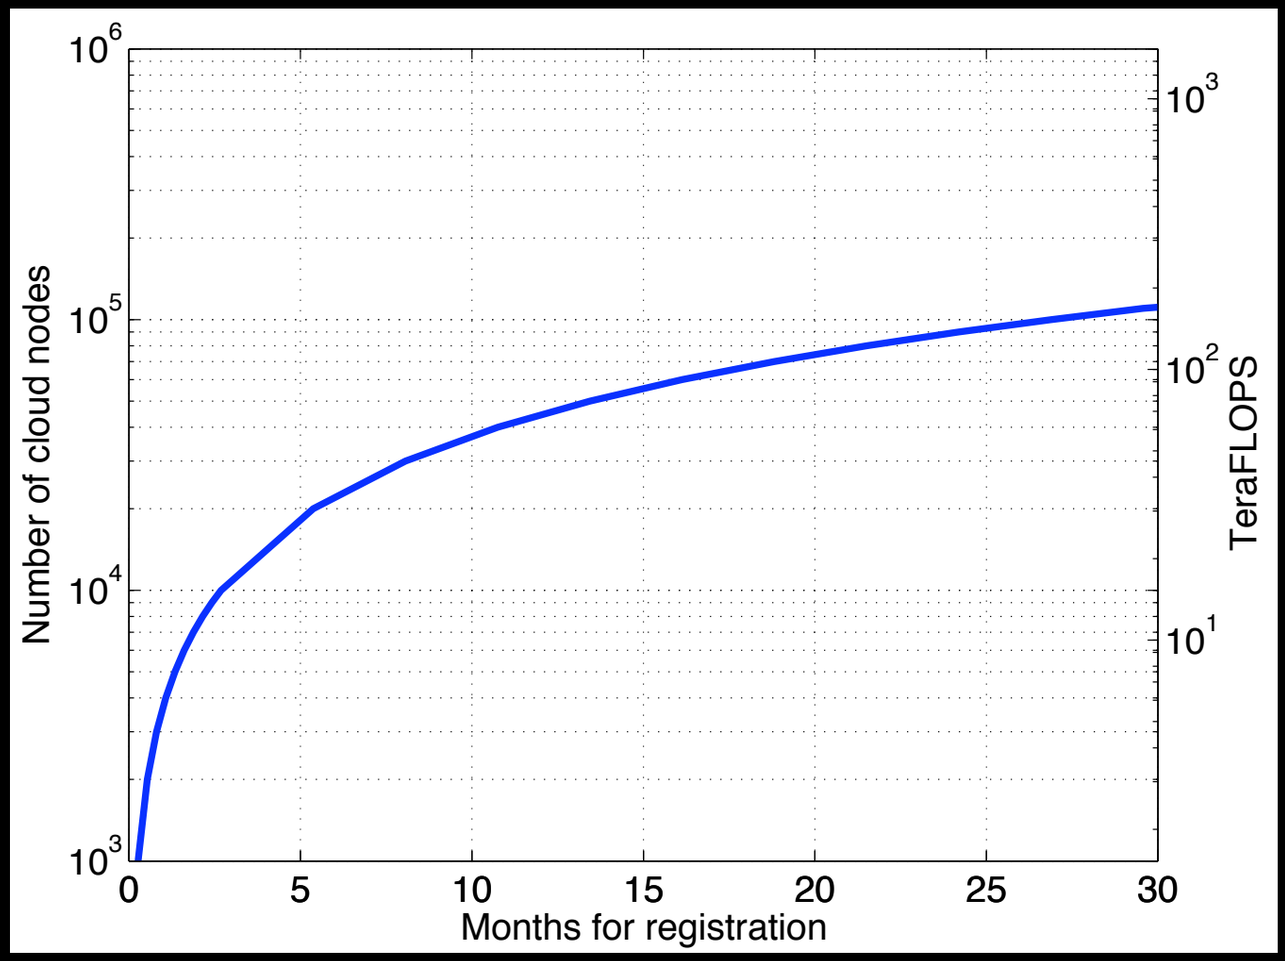

Platform Construction → Application Deployment → Application Execution → Application Completion

How long before I get X TeraFLOPS?

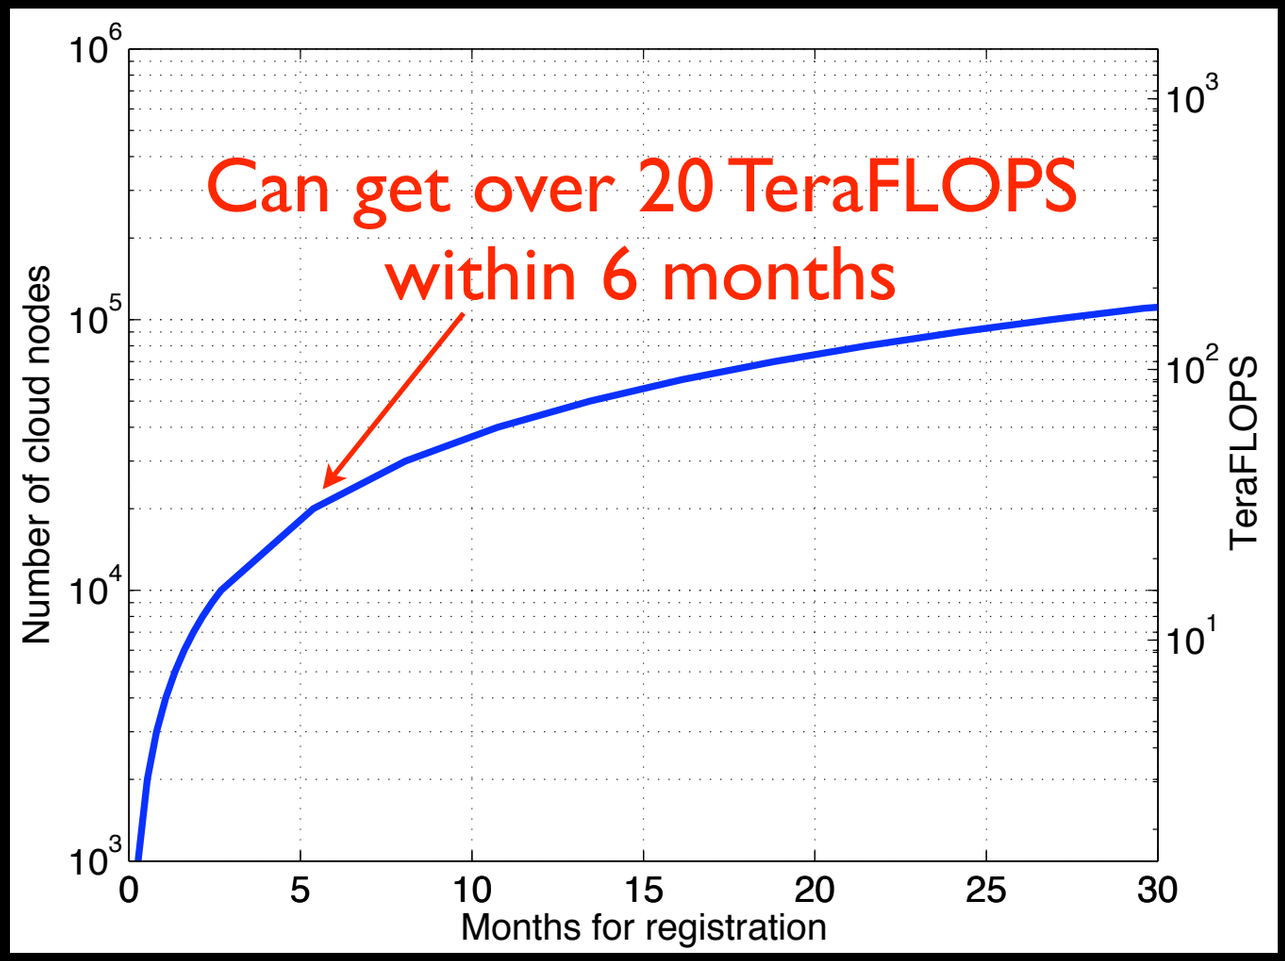

Platform Construction → Application Deployment → Application Execution → Application Completion

How long before I get X TeraFLOPS?

Strategy:
Add to BOINC project list
Press releases
Forum Announcements
Google Ad Sense
Respond to users (leverage volunteers)

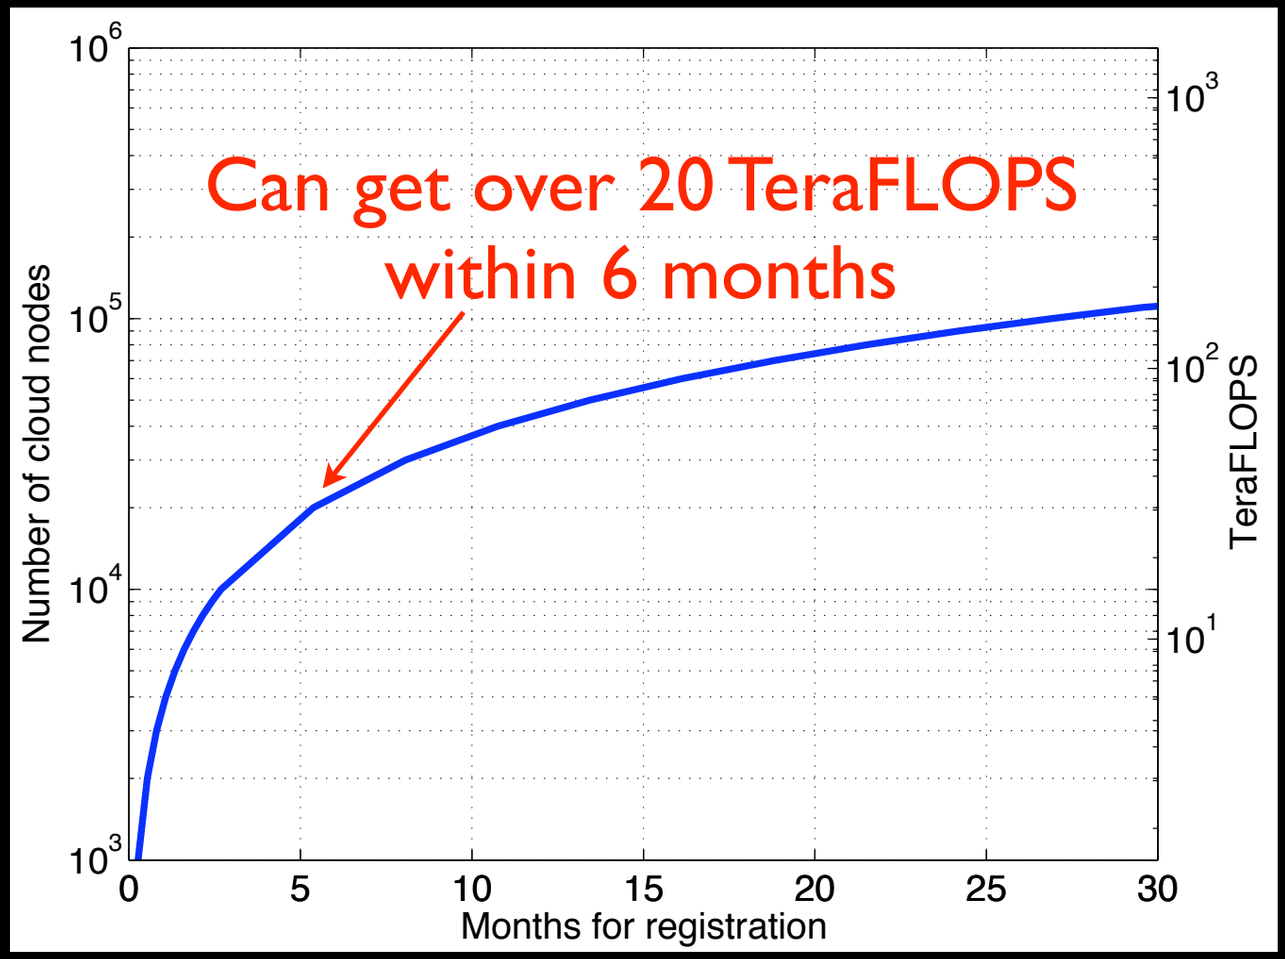

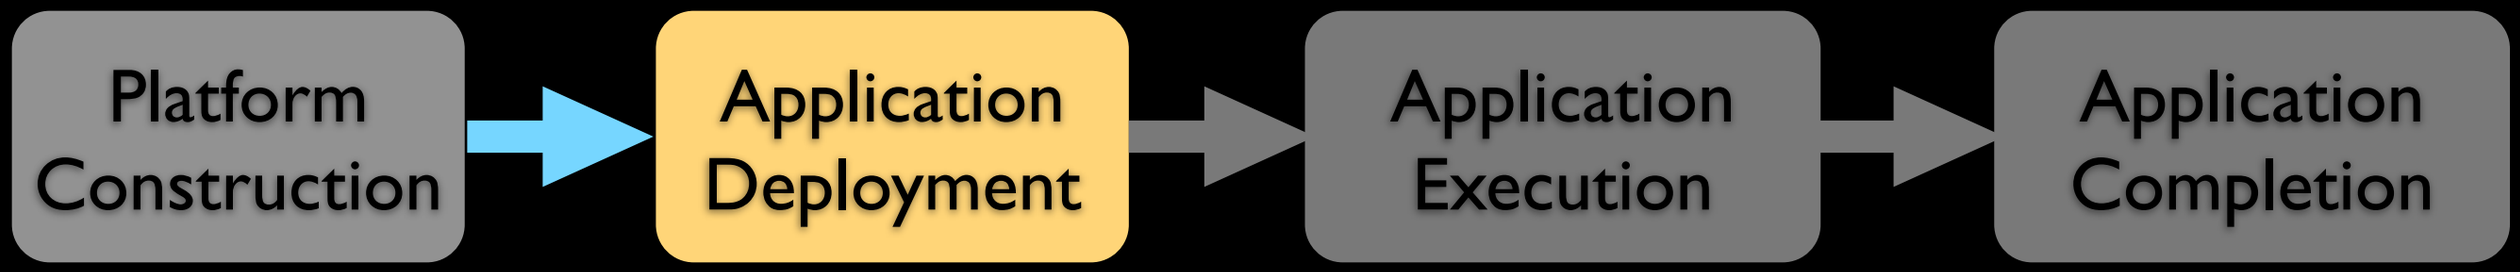

How long to deploy my batch of tasks needing faster response time?

Platform Construction → **Application Deployment** → Application Execution → Application Completion

How long to deploy my batch of tasks needing faster response time?

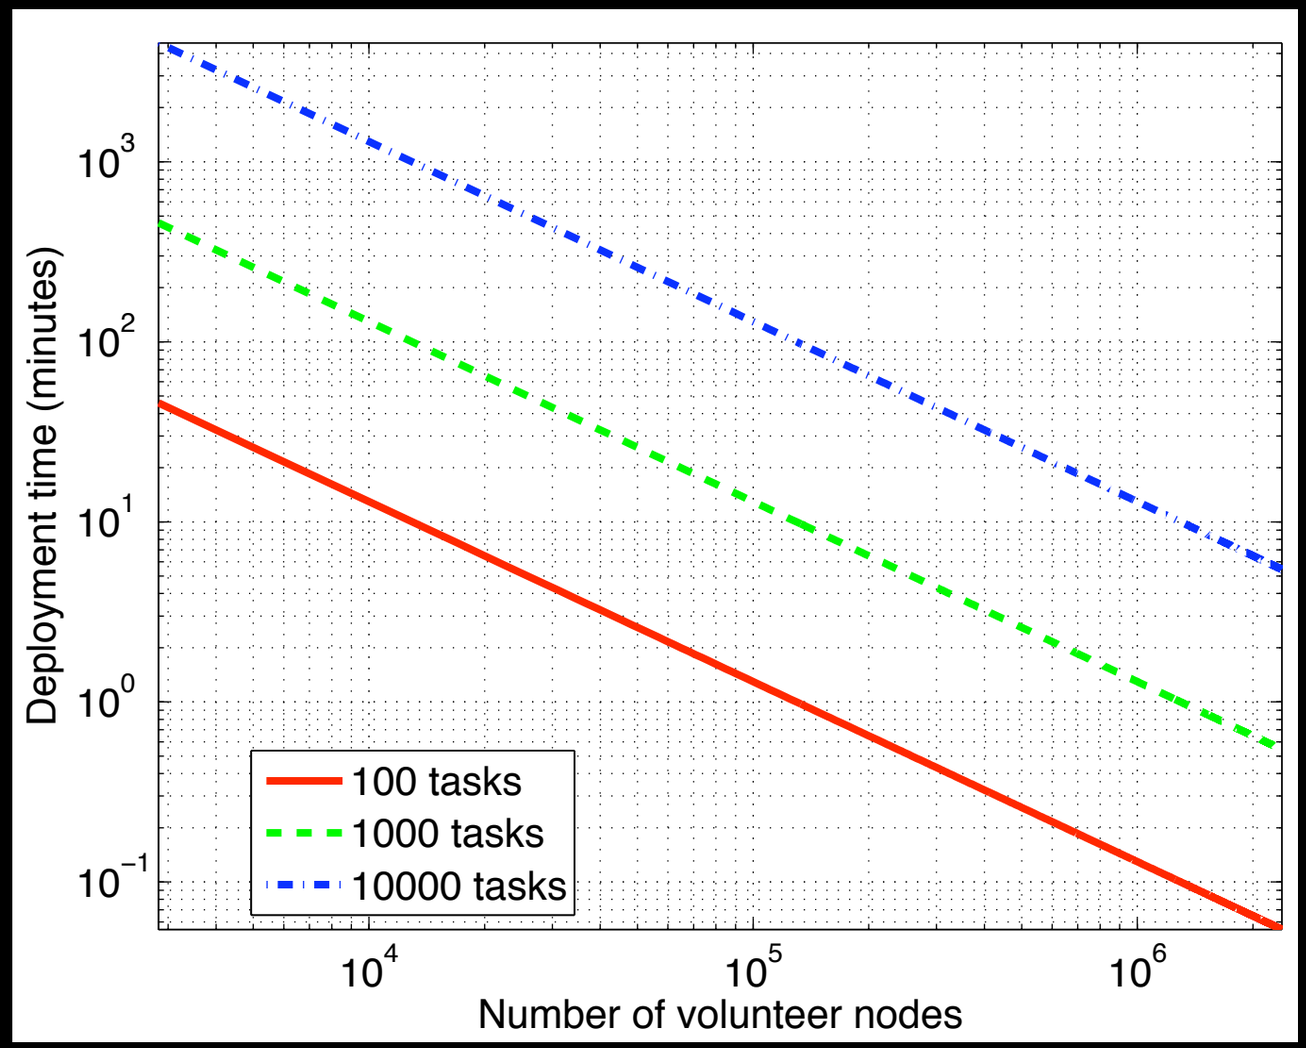

Platform Construction → **Application Deployment** → Application Execution → Application Completion

How long to deploy my batch of tasks needing faster response time?

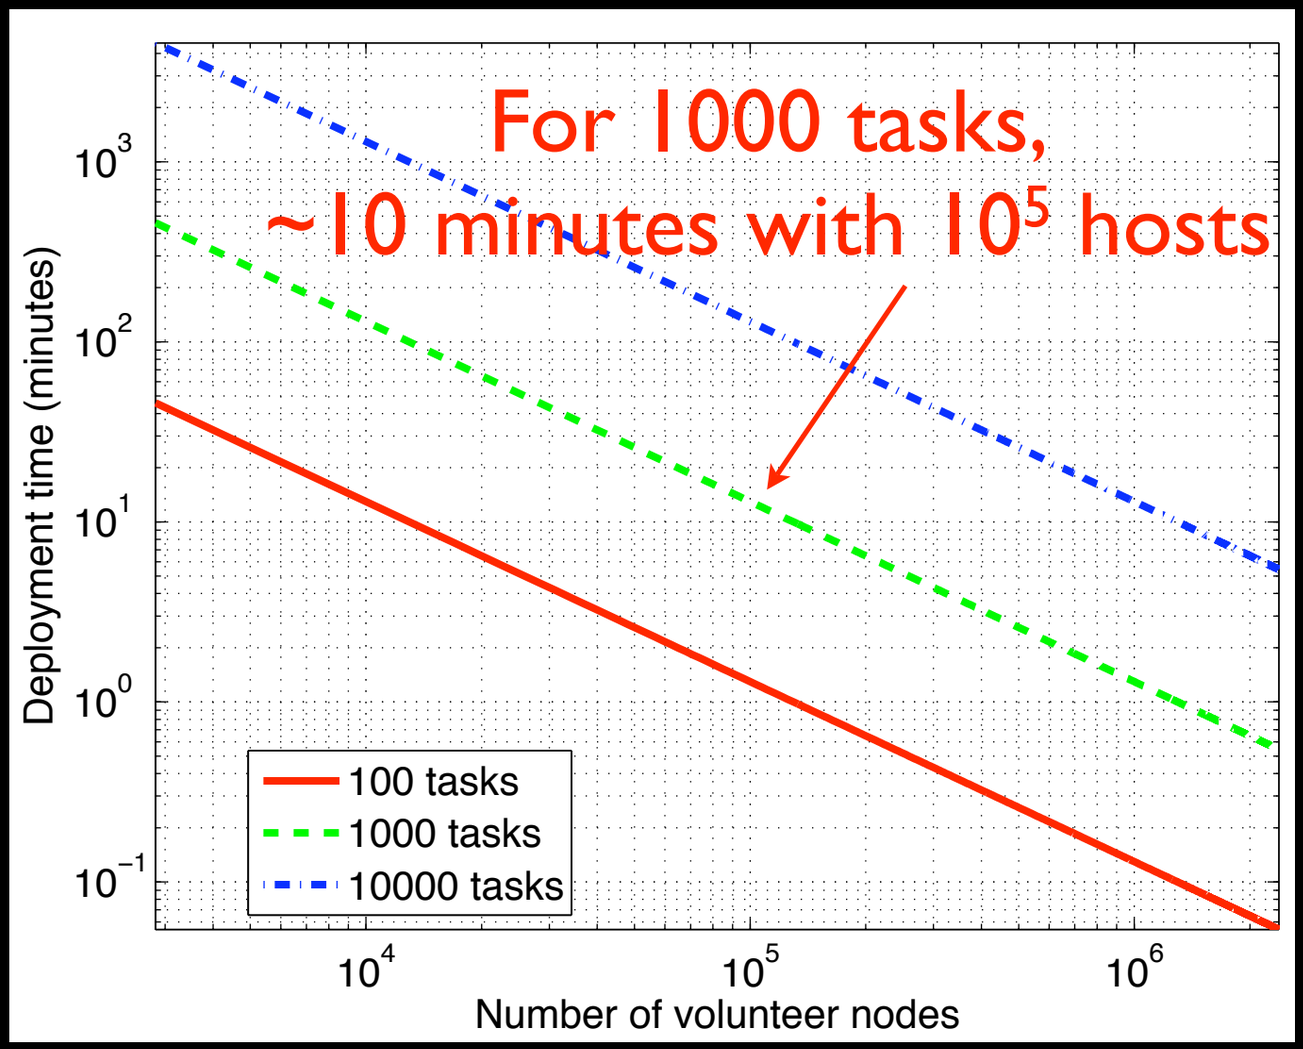

Platform Construction → **Application Deployment** → Application Execution → Application Completion

How long to deploy my batch of tasks needing faster response time?

Strategy: Specify lower latency bounds [Heien et al.]

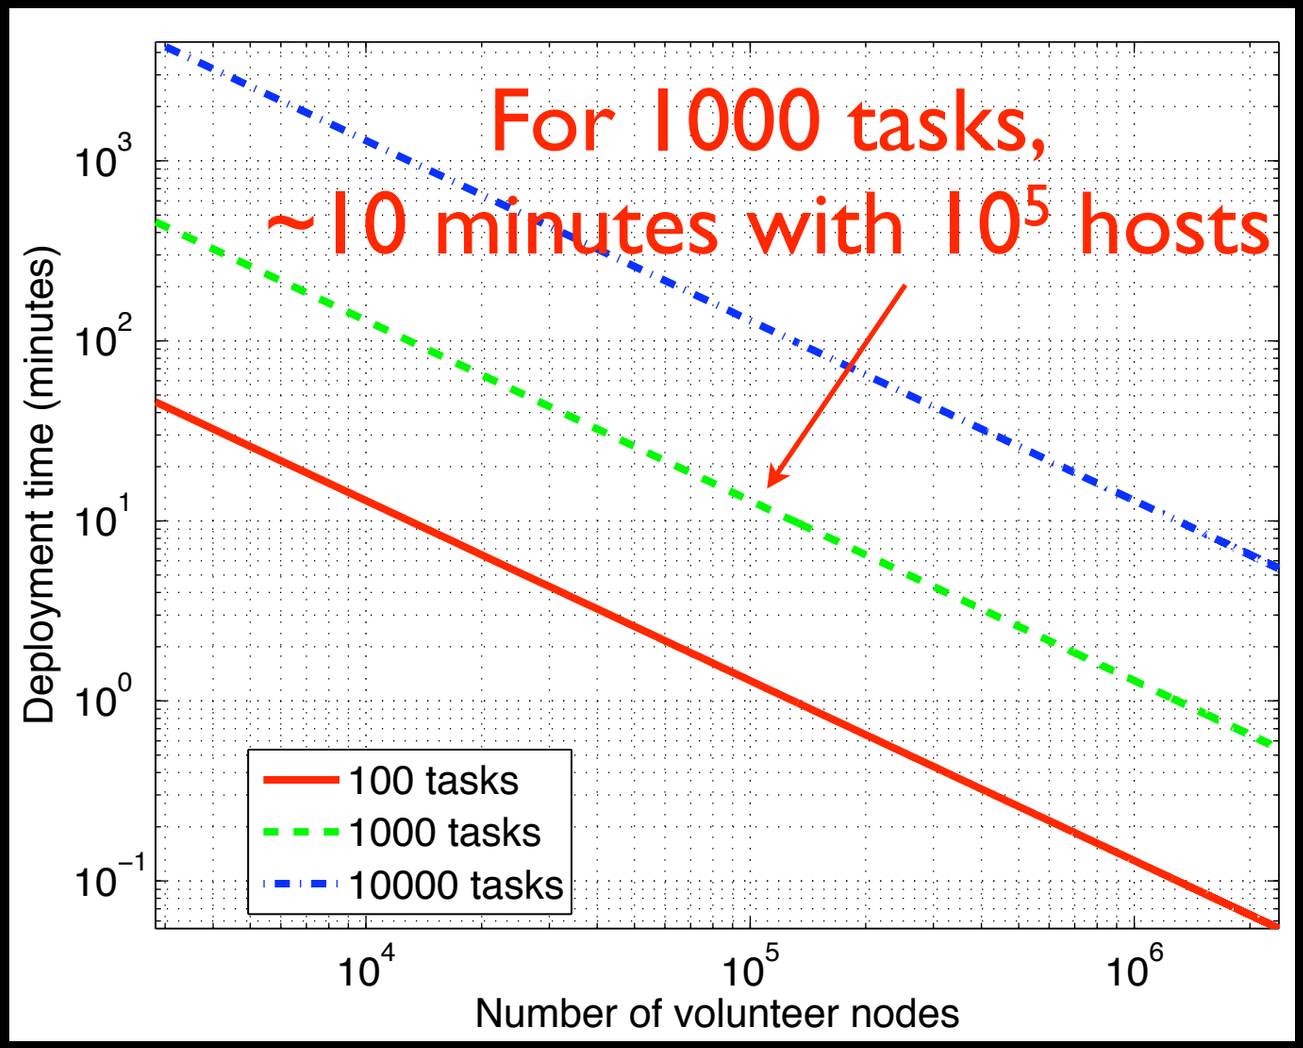

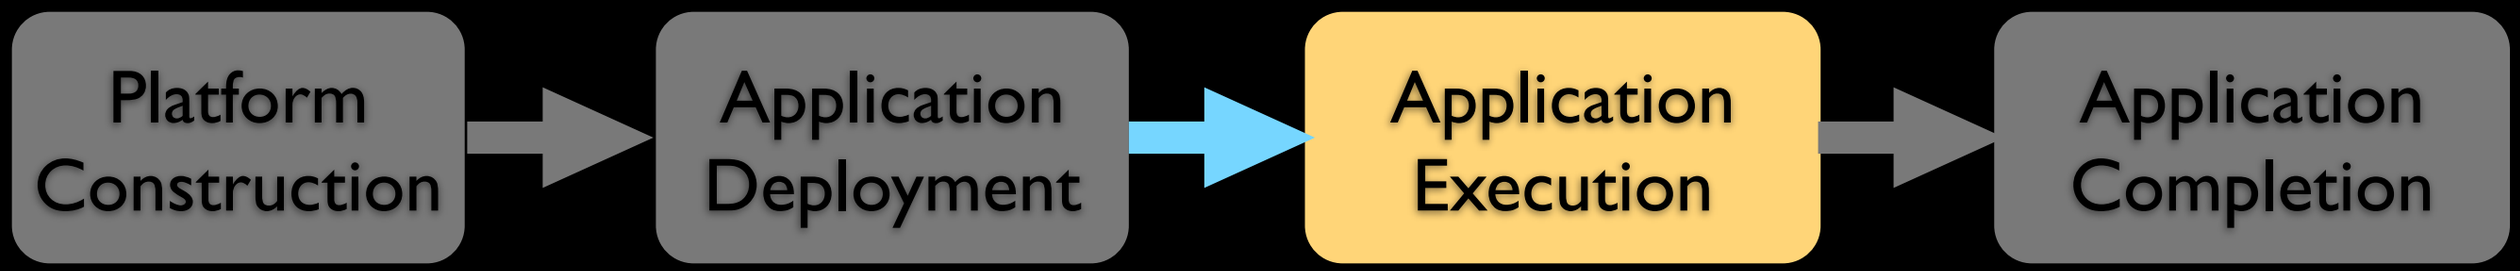
Platform Construction
Application Deployment
Application Execution
Application Completion

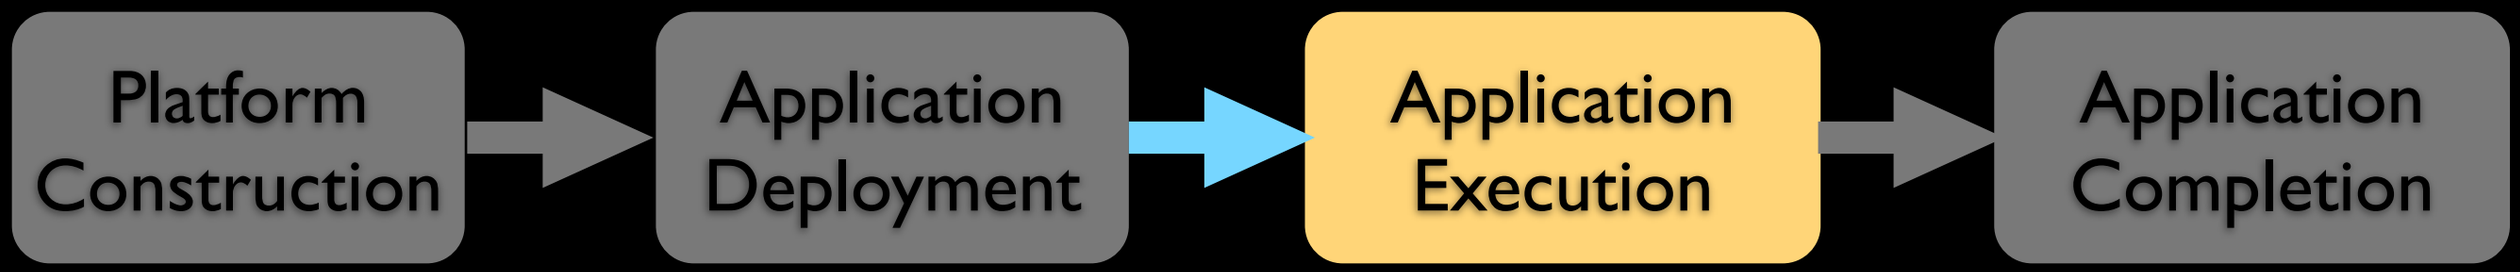

How many volunteer nodes are equivalent to 1 cloud node?

Platform Construction → Application Deployment → **Application Execution** → Application Completion

How many volunteer nodes are equivalent to 1 cloud node?

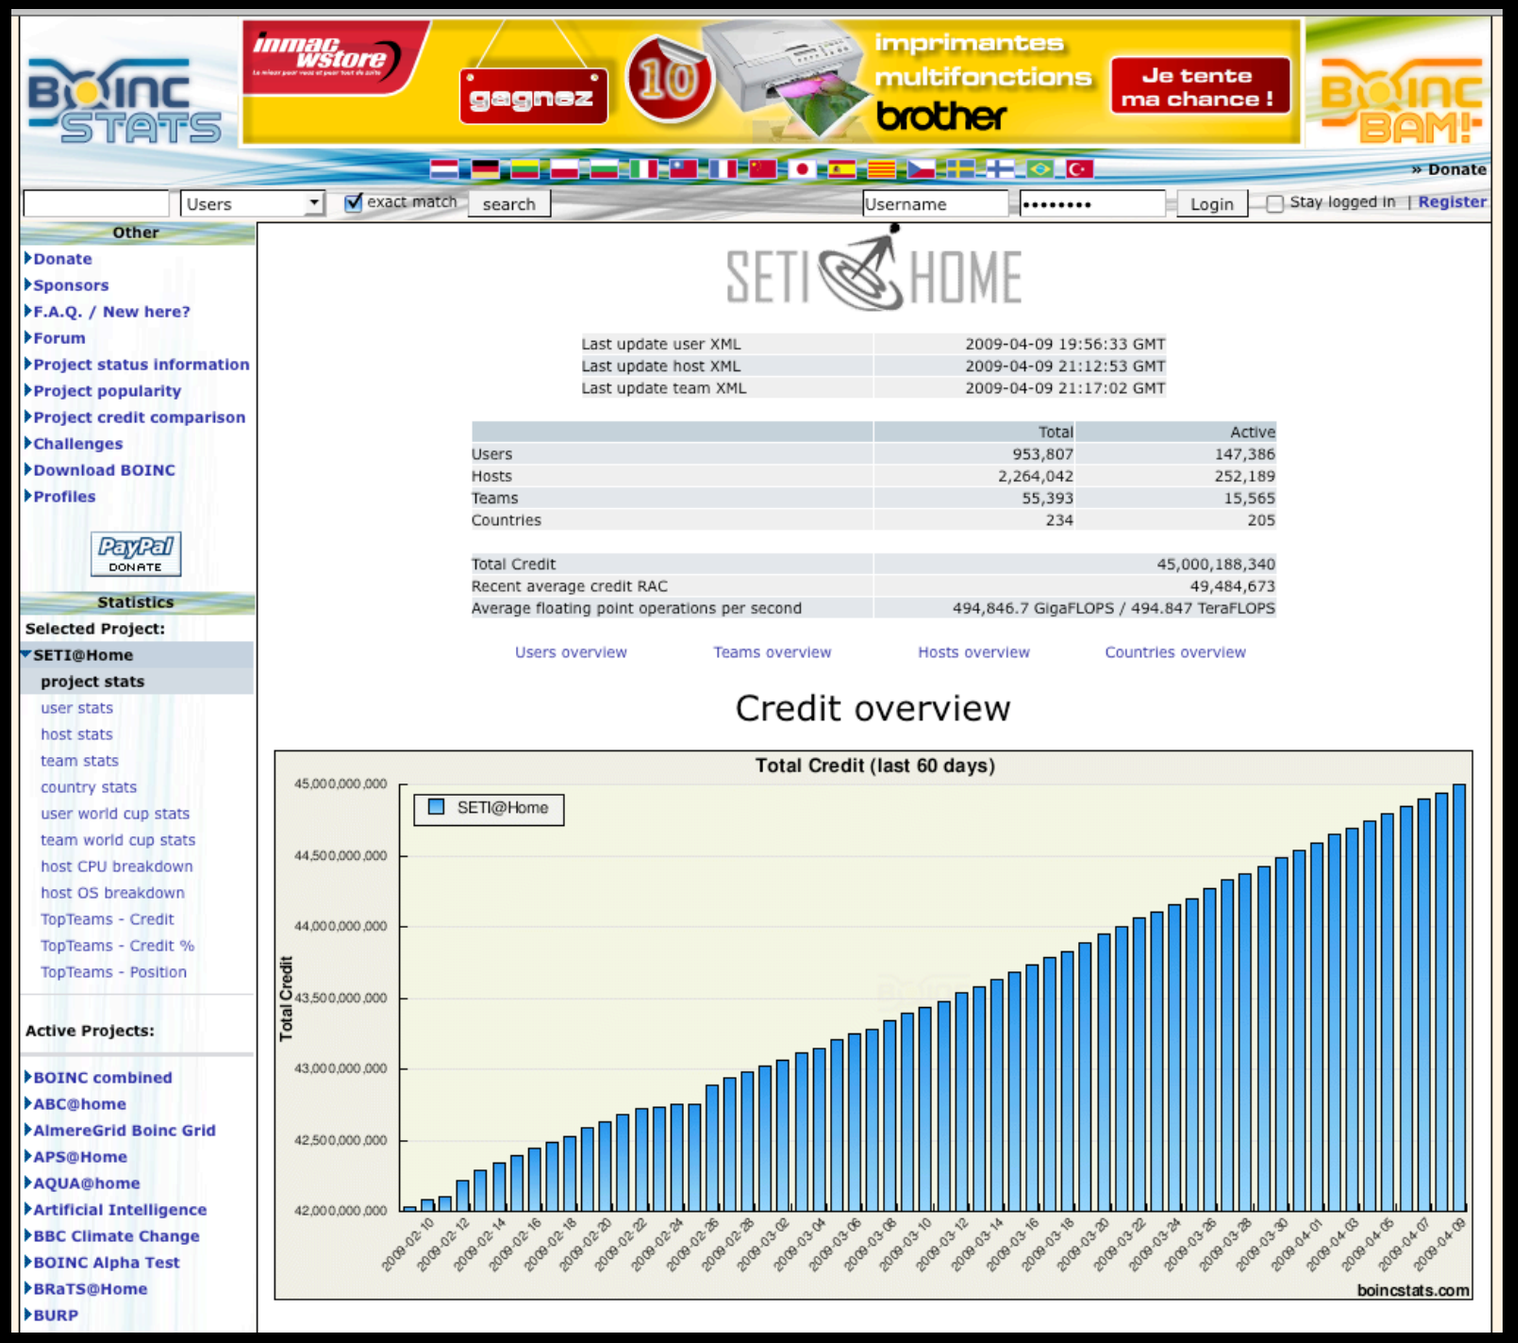

Platform Construction → Application Deployment → **Application Execution** → Application Completion

How many volunteer nodes are equivalent to 1 cloud node?

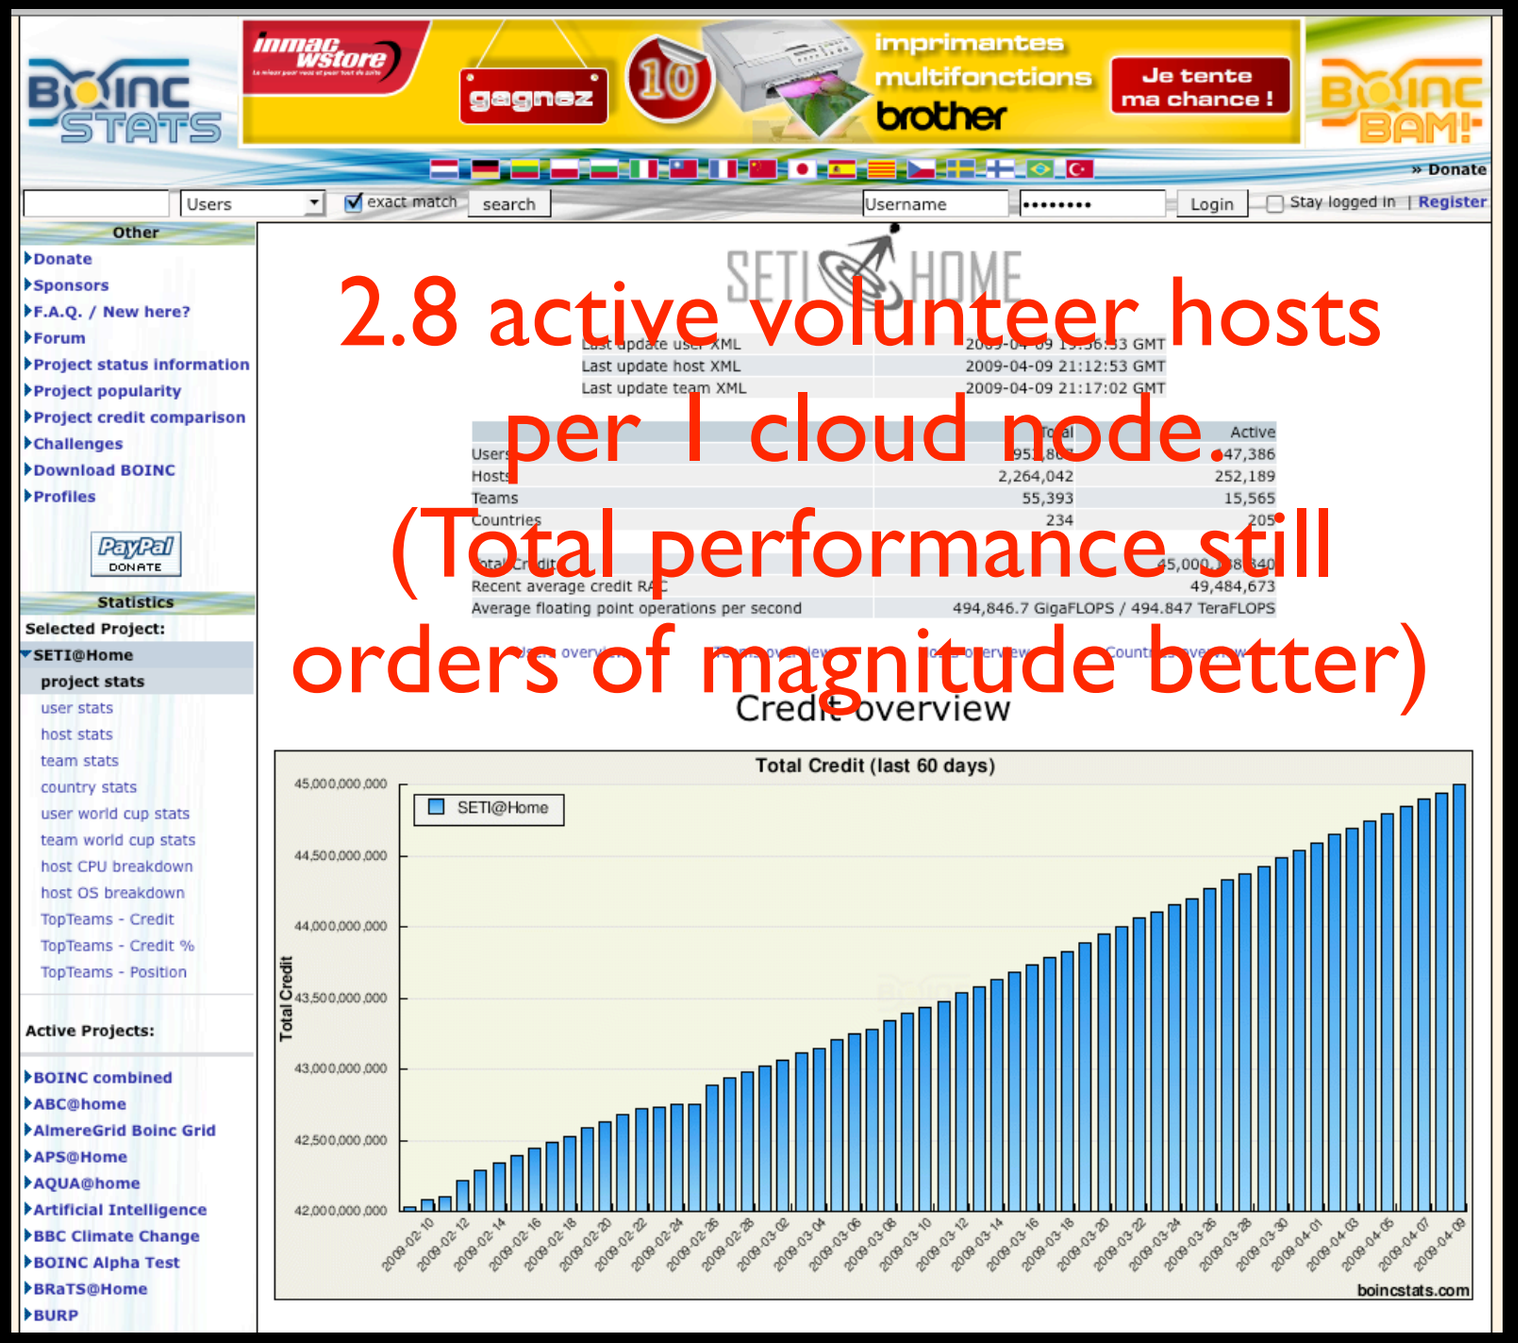

2.8 active volunteer hosts per 1 cloud node. (Total performance still orders of magnitude better)

Platform Construction → Application Deployment → **Application Execution** → Application Completion

How many volunteer nodes are equivalent to 1 cloud node?

Strategy:
Use statistical prediction of availability

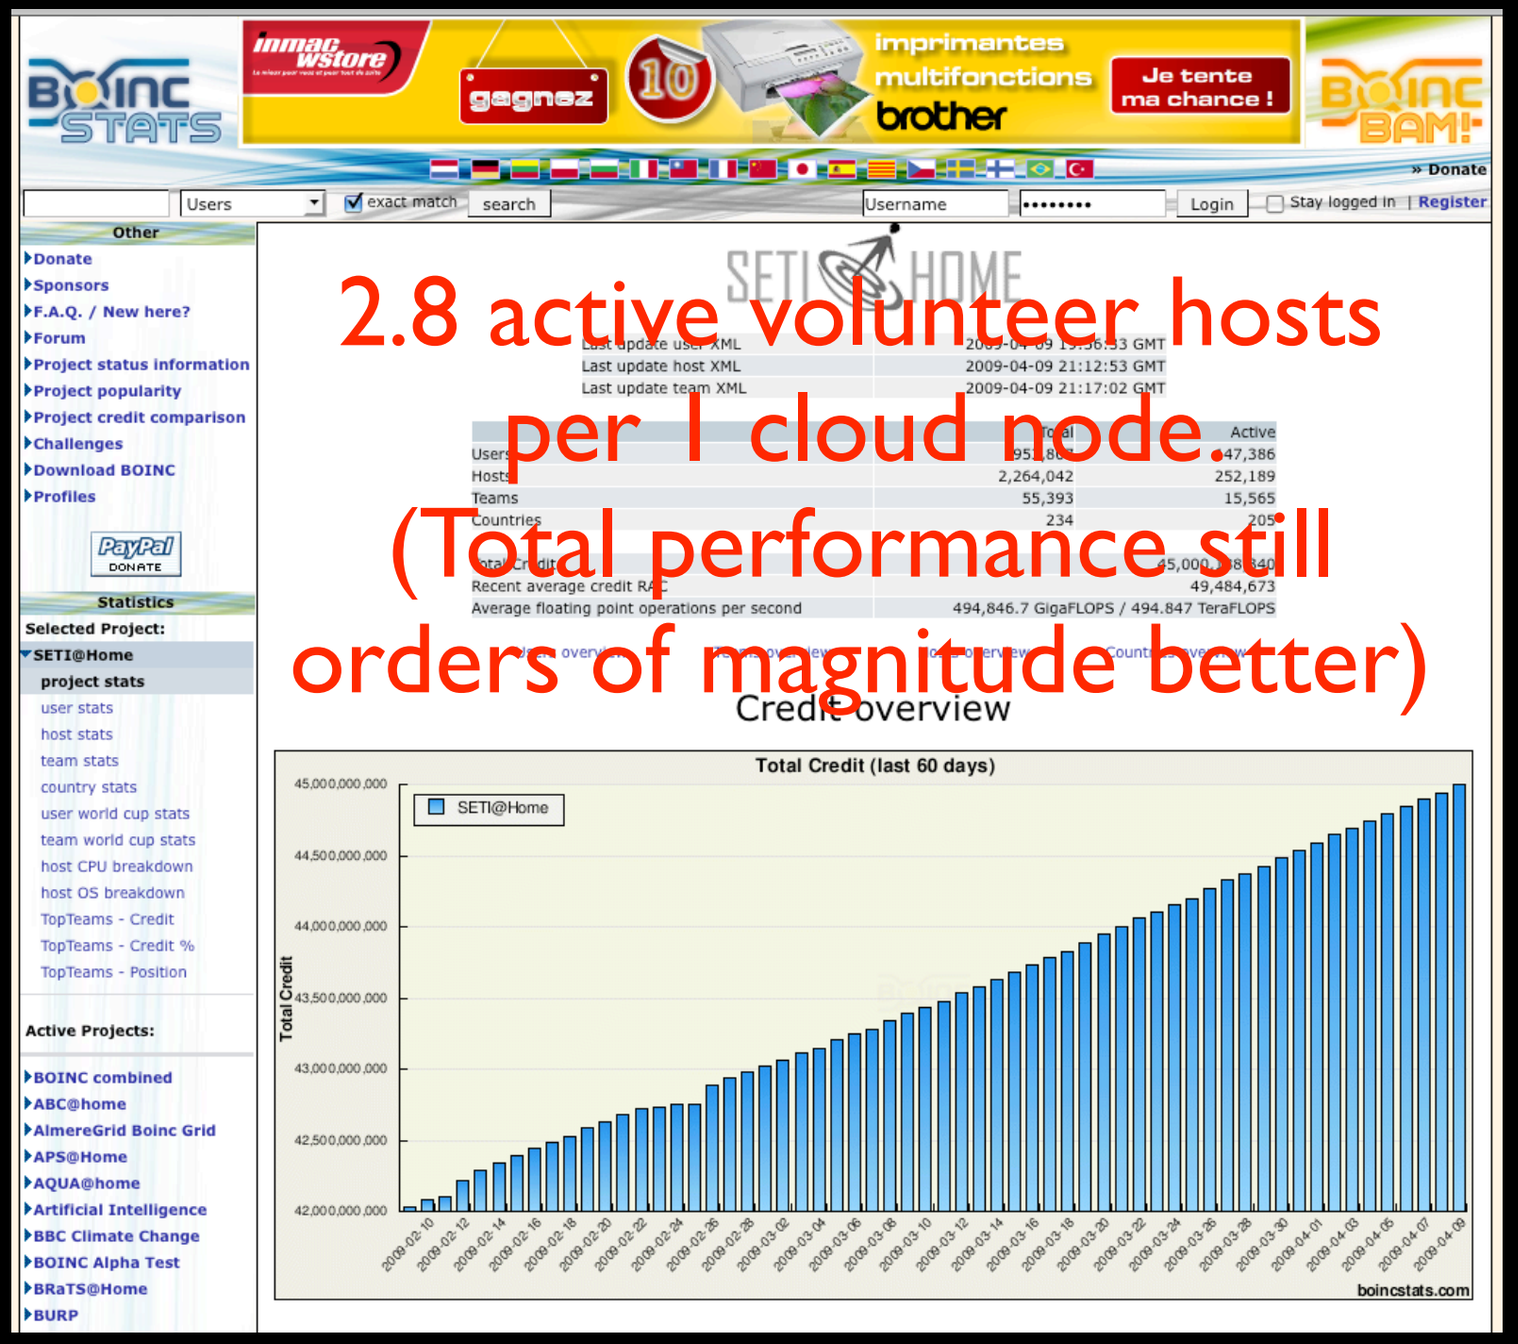

2.8 active volunteer hosts per 1 cloud node.
(Total performance still orders of magnitude better)

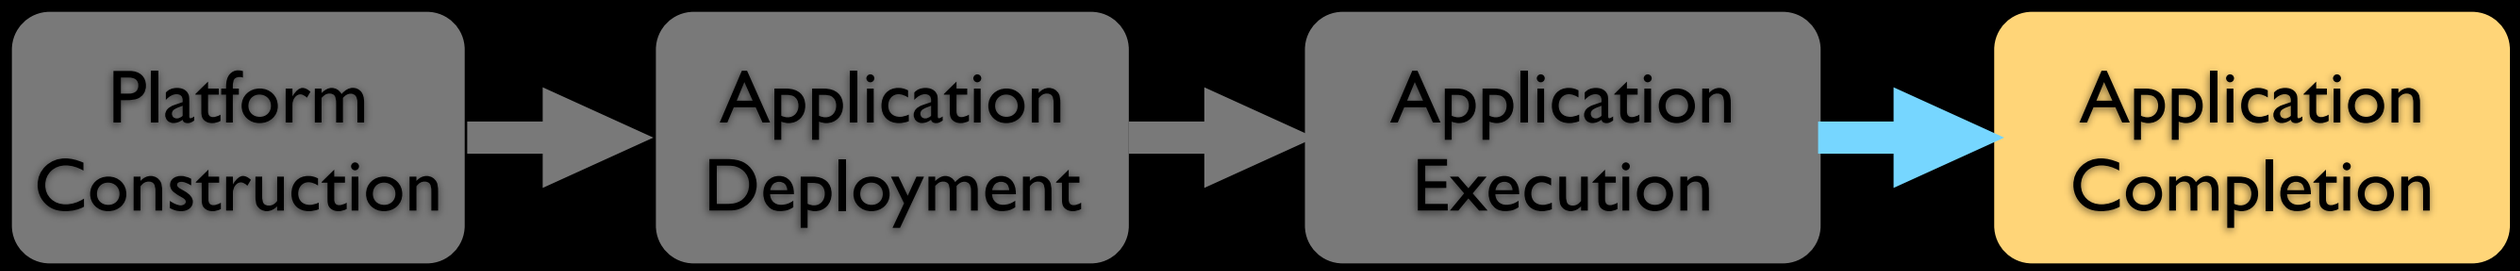
Platform Construction
Application Deployment
Application Execution
Application Completion

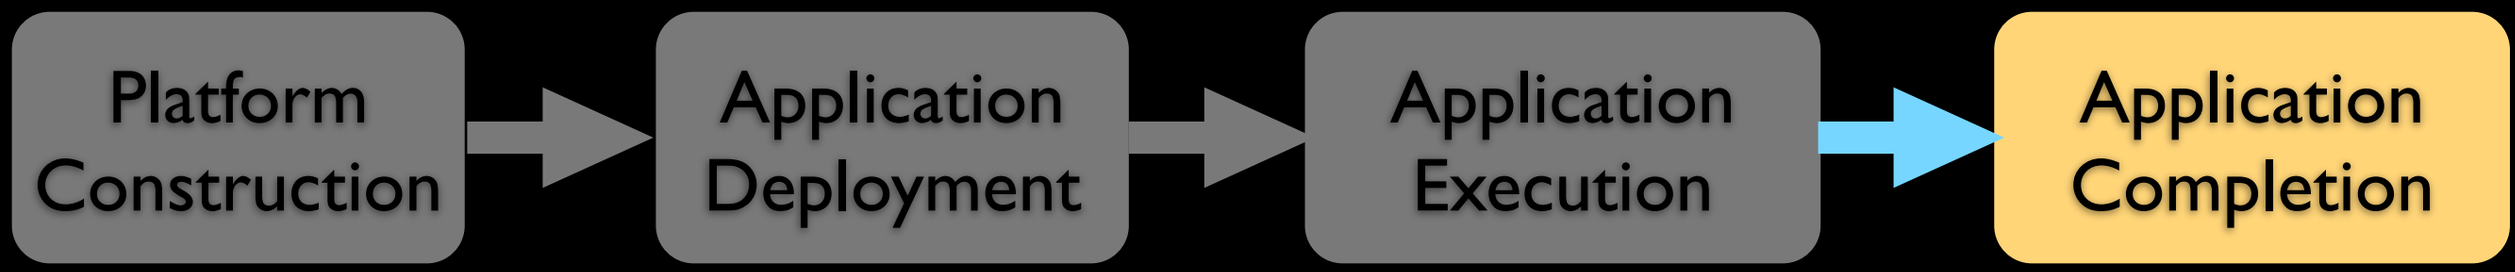

How long should I wait for task completion?

Platform Construction → Application Deployment → Application Execution → Application Completion

How long should I wait for task completion?

Catalog of BOINC Powered Projects – Unofficial BOINC Wiki

http://www.boinc-wiki.info/Catalog_of_BOINC_Powered_Projects

boinc catalog

Babel Fish GTranslate iGoogle Verb Conjugator GGroups BWorkShop08 Calendar GReader GDocs To Do Gmail Grid INRIA – international ANR GTranslate easychair Amazon AConsole CloudWiki CloudWiki ISGC

Projects, Science Applications, and Platforms [edit]

Below is a table of the known BOINC Powered Projects, their Science Applications vs. support by Platform.

Project	Science Application	Operating System	Installed Memory-requirement (MB)	Disk Space Required (MB)	Process time (CPU dependant)	Download Size (MB)	Upload Size (MB)	Deadline (days)
ABC@Home	abc-finder	Windows, Linux, MacOS Intel	104.91	48	0.5 h	0.01	0.01	7
BBC Climate Change Experiment	hadcm3l	Windows, Linux	96.00	600	5 months	0.02	5-6 (16x)	347
ClimatePrediction.net	hadam3	Windows, Linux	1464.85	1024	4 weeks	13	30	180
	hadcm3l	Windows, Linux, MacOS Intel	96.00	600	5 months	0.02	5-6 (x16)	347
	hadsm3	Windows, Linux, MacOS Intel	61.1	573	3 weeks	0.02	5-10 (x3)	345
Seasonal Attribution Project	hadam3	Windows, Linux	256 (uses 430, recommended 1024)	1024	4 weeks	13	30	180
Einstein@Home	einstein-S5R3	Windows, Linux, Mac OS X, Other(*3)	57.22	96	24 h	4-30	0.16	14-21
Leiden Classical	Classical	Windows, Linux, Mac OS X, FreeBSD	?	?	?	?	?	7
	trajtou-cu111	Windows, Linux	?	?	?	?	?	?
	trajtou-pd110paw	Windows, Linux	?	?	?	?	?	?
	trajtou-pt111	Windows, Linux	?	?	?	?	?	?
LHC@Home	garfield	Windows, Linux	?	?	?	?	?	?
	sixtrack	Windows, Linux	57.22	29	1-10 h	0.024	0.034	4-8
MalariaControl.net	malariacontrol	Windows, Linux, MacOS Intel	85.84	191	2h	0.07	?	3.5
	Prediction of Malaria Prevalence	Windows	?	?	0,5 h	?	?	?
	malariacontrol test version	Windows, Linux, MacOS Intel	?	?	?	?	?	?
	Estimation of parameters of infection dynamics	Windows	?	?	?	?	?	?
Rosetta@Home	rosetta	Windows, Linux, Mac OS X	95.37 (recommended 256)	96	3 h (by default, but configurable)	2.8-4	0.02	10

Platform Construction → Application Deployment → Application Execution → Application Completion

How long should I wait for task completion?

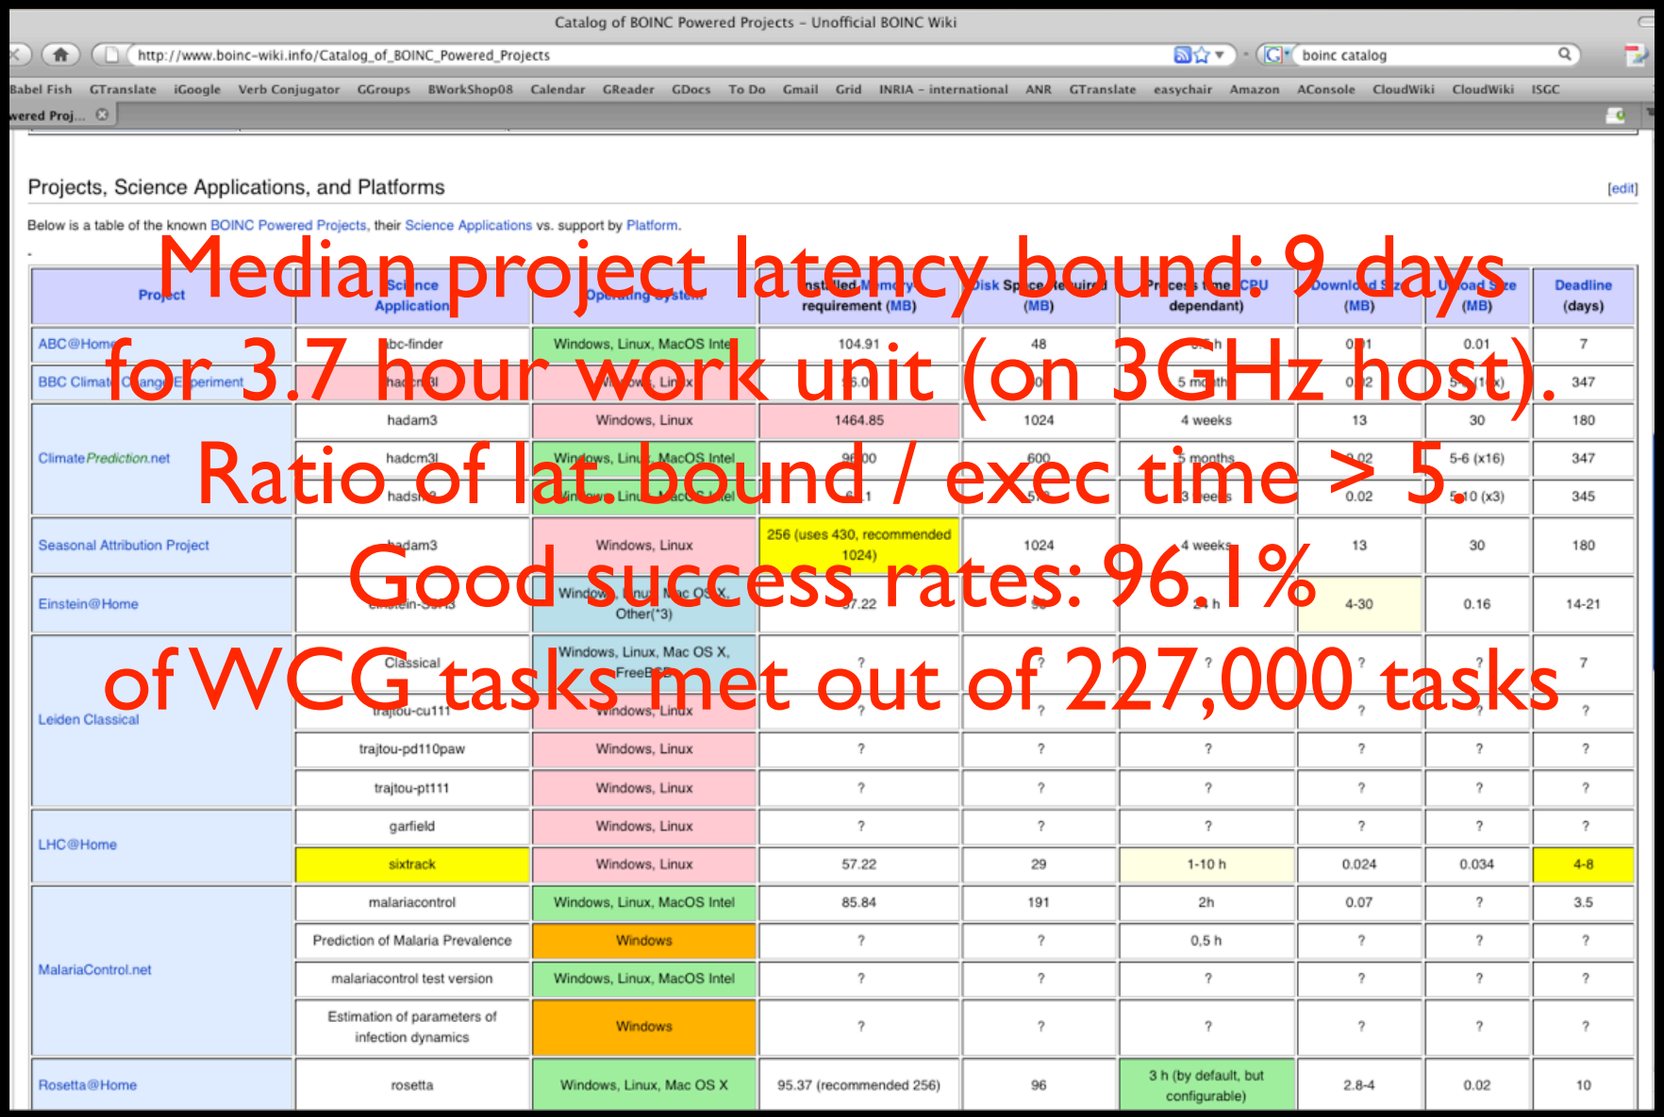

Project	Science Application	Operating System	Installed Memory requirement (MB)	Disk Space required (MB)	Processor time (CPU dependant)	Download Size (MB)	Upload Size (MB)	Deadline (days)
ABC@Home	abc-finder	Windows, Linux, MacOS Intel	104.91	48	3-4 h	0.01	0.01	7
BBC Climate Change Experiment	hadcm3l	Windows, Linux	1536	600	5 months	0.02	5-6 (x16)	347
ClimatePrediction.net	hadam3	Windows, Linux	1464.85	1024	4 weeks	13	30	180
	hadcm3l	Windows, Linux, MacOS Intel	1536	600	5 months	0.02	5-6 (x16)	347
	hadsm3	Windows, Linux, MacOS Intel	921	512	3 weeks	0.02	5-10 (x3)	345
Seasonal Attribution Project	hadam3	Windows, Linux	256 (uses 430, recommended 1024)	1024	4 weeks	13	30	180
Einstein@Home	Einstein-S5R3	Windows, Linux, Mac OS X, Other(*3)	57.22	50	2-4 h	4-30	0.16	14-21
Leiden Classical	Classical	Windows, Linux, Mac OS X, FreeBSD	?	?	?	?	?	7
	trajtou-cu111	Windows, Linux	?	?	?	?	?	?
	trajtou-pd110paw	Windows, Linux	?	?	?	?	?	?
	trajtou-pt111	Windows, Linux	?	?	?	?	?	?
LHC@Home	garfield	Windows, Linux	?	?	?	?	?	?
	sixtrack	Windows, Linux	57.22	29	1-10 h	0.024	0.034	4-8
MalariaControl.net	malariacontrol	Windows, Linux, MacOS Intel	85.84	191	2h	0.07	?	3.5
	Prediction of Malaria Prevalence	Windows	?	?	0,5 h	?	?	?
	malariacontrol test version	Windows, Linux, MacOS Intel	?	?	?	?	?	?
	Estimation of parameters of infection dynamics	Windows	?	?	?	?	?	?
Rosetta@Home	rosetta	Windows, Linux, Mac OS X	95.37 (recommended 256)	96	3 h (by default, but configurable)	2.8-4	0.02	10

Platform Construction → Application Deployment → Application Execution → Application Completion

How long should I wait for task completion?

Strategy:
See BOINC Catalog for typical deadlines and compute/comm/mem ratios.

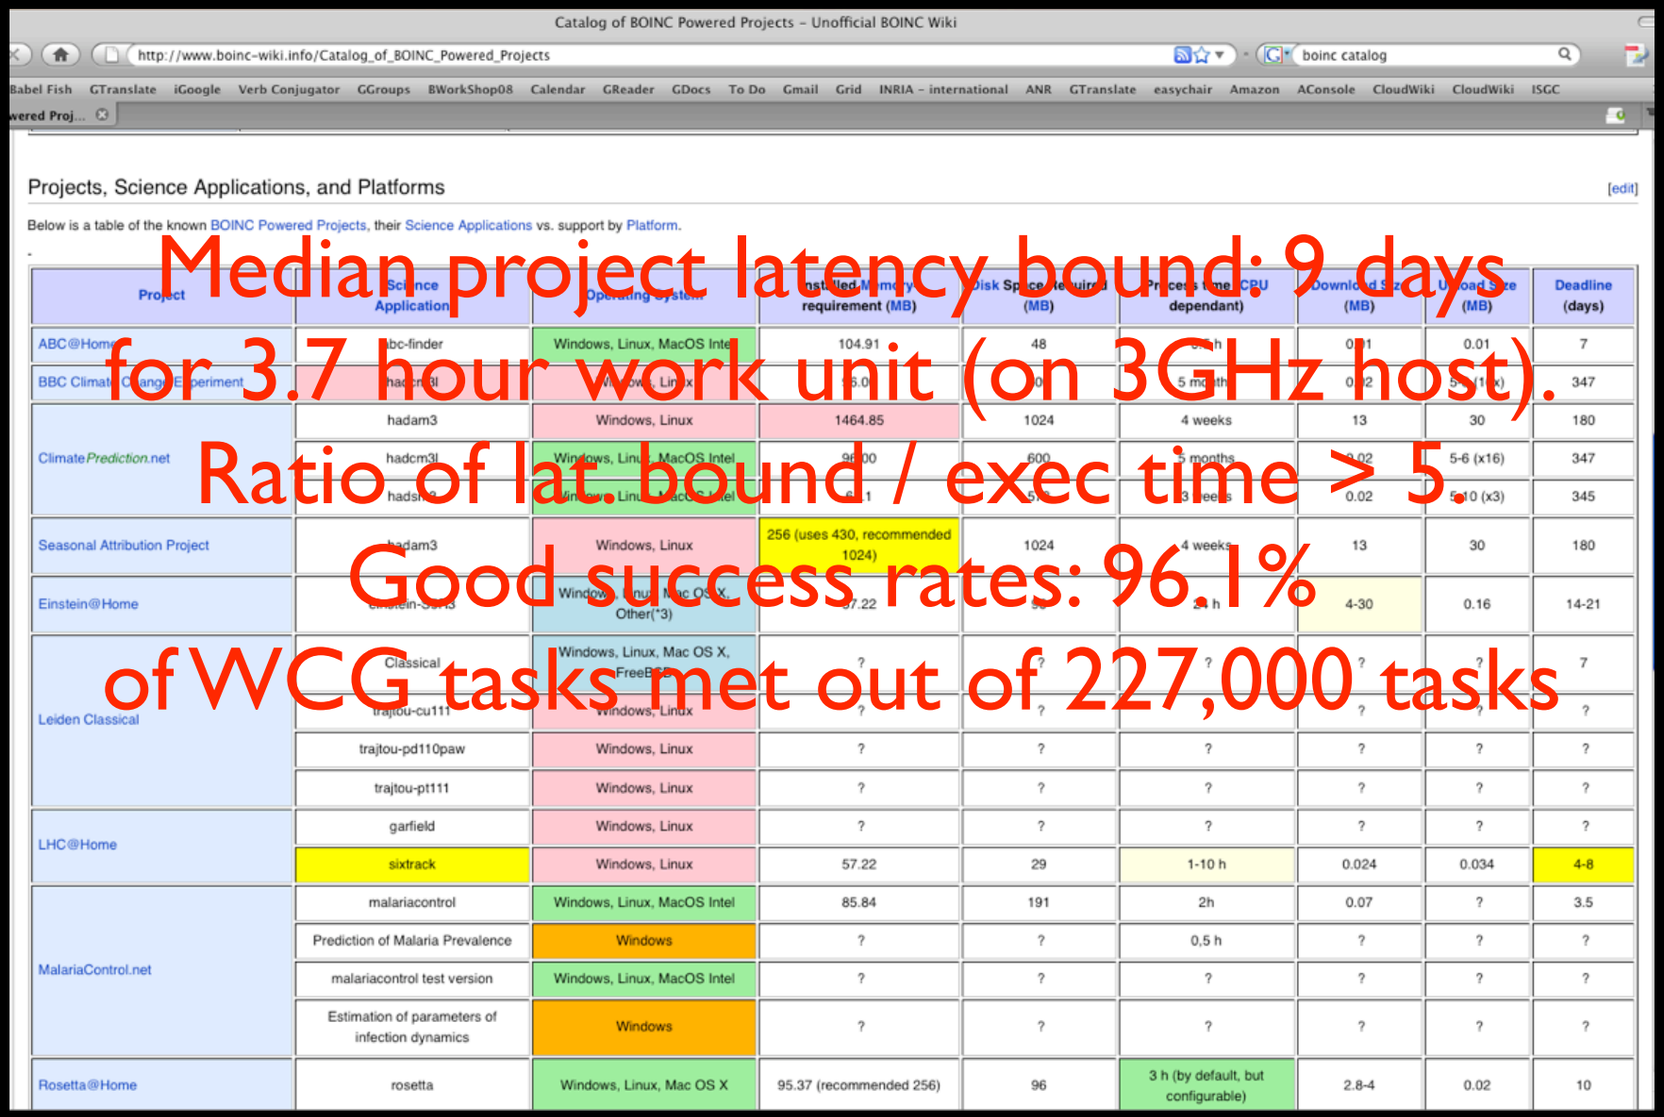

Median project latency bound: 9 days
for 3.7 hour work unit (on 3GHz host).
Ratio of lat. bound / exec time > 5.
Good success rates: 96.1%
of WCG tasks met out of 227,000 tasks

Monetary Tradeoffs

- Client hosting on cloud
  - Not worth it and never will
- Server hosting on the cloud
  - Possible solution

Monthly Project Costs

Component	Project	
	SETI@home	*XtremLab*
Salaries	10K for sys admins	5K
Electricity	90 for 6 servers	15
Network	2K for 100 Mbit	covered by university
Hardware	18K for servers, 25K for air conditioner	4K
Total startup	43K	4K
Total monthly	12K	5k

EC2 Pricing

Instance Type	Cost/hour (USD)
Standard Small	0.10
Standard Large	0.40
High-CPU	0.20

Table 1. Pricing for EC2 Instances

Transfer Type	Cost/GB-Month (USD)
Inbound transfer	0.10
first 10 TB	0.17
next 10-50TB	0.13
next 50-150TB	0.11
over 150 TB	0.10

Table 2. Pricing for EC2 Data Transfer

Resource	Rate (USD)
Storage	0.10 / GB-Month
IO request	0.10 / million

Table 3. Pricing for EBS

What about Client Hosting on the Cloud?

Cost of Clouds versus DG

How long until DG's more cost effective than Clouds?

Cost of Clouds versus DG

How long until DG's more cost effective than Clouds?

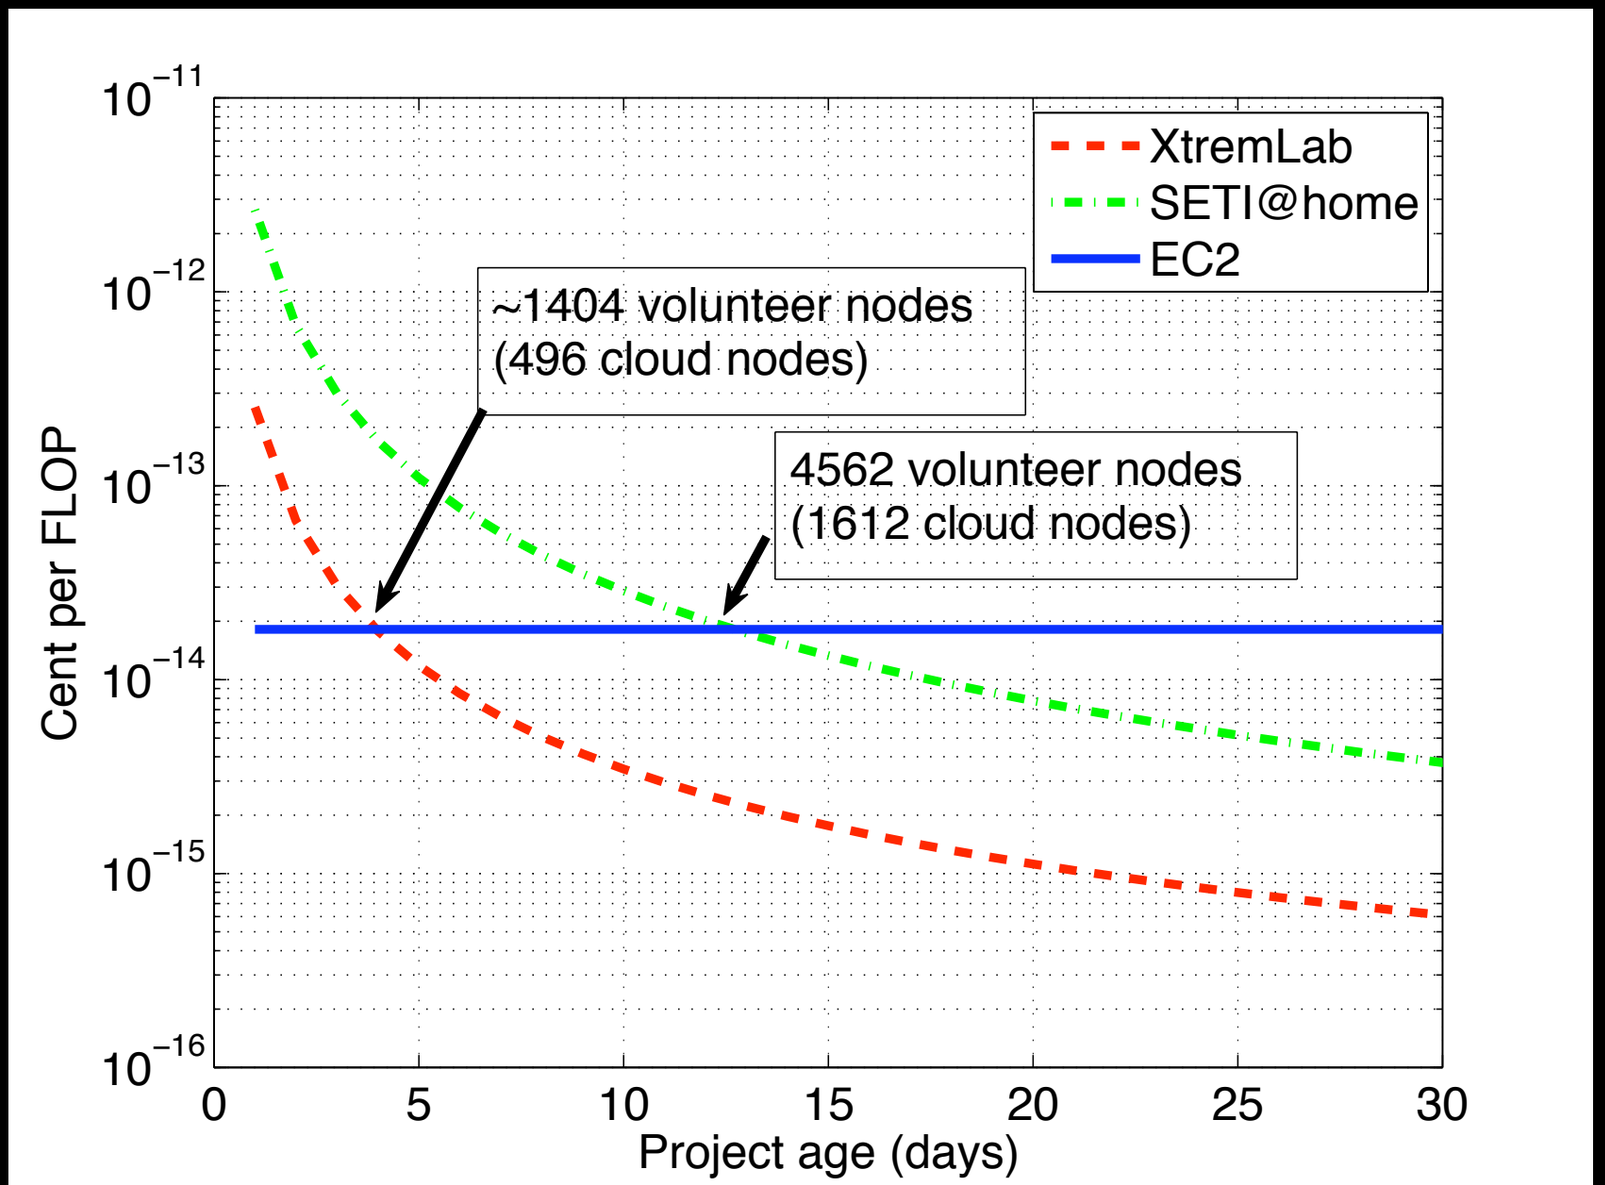

Cost of Clouds versus DG

How long until DG's more cost effective than Clouds?

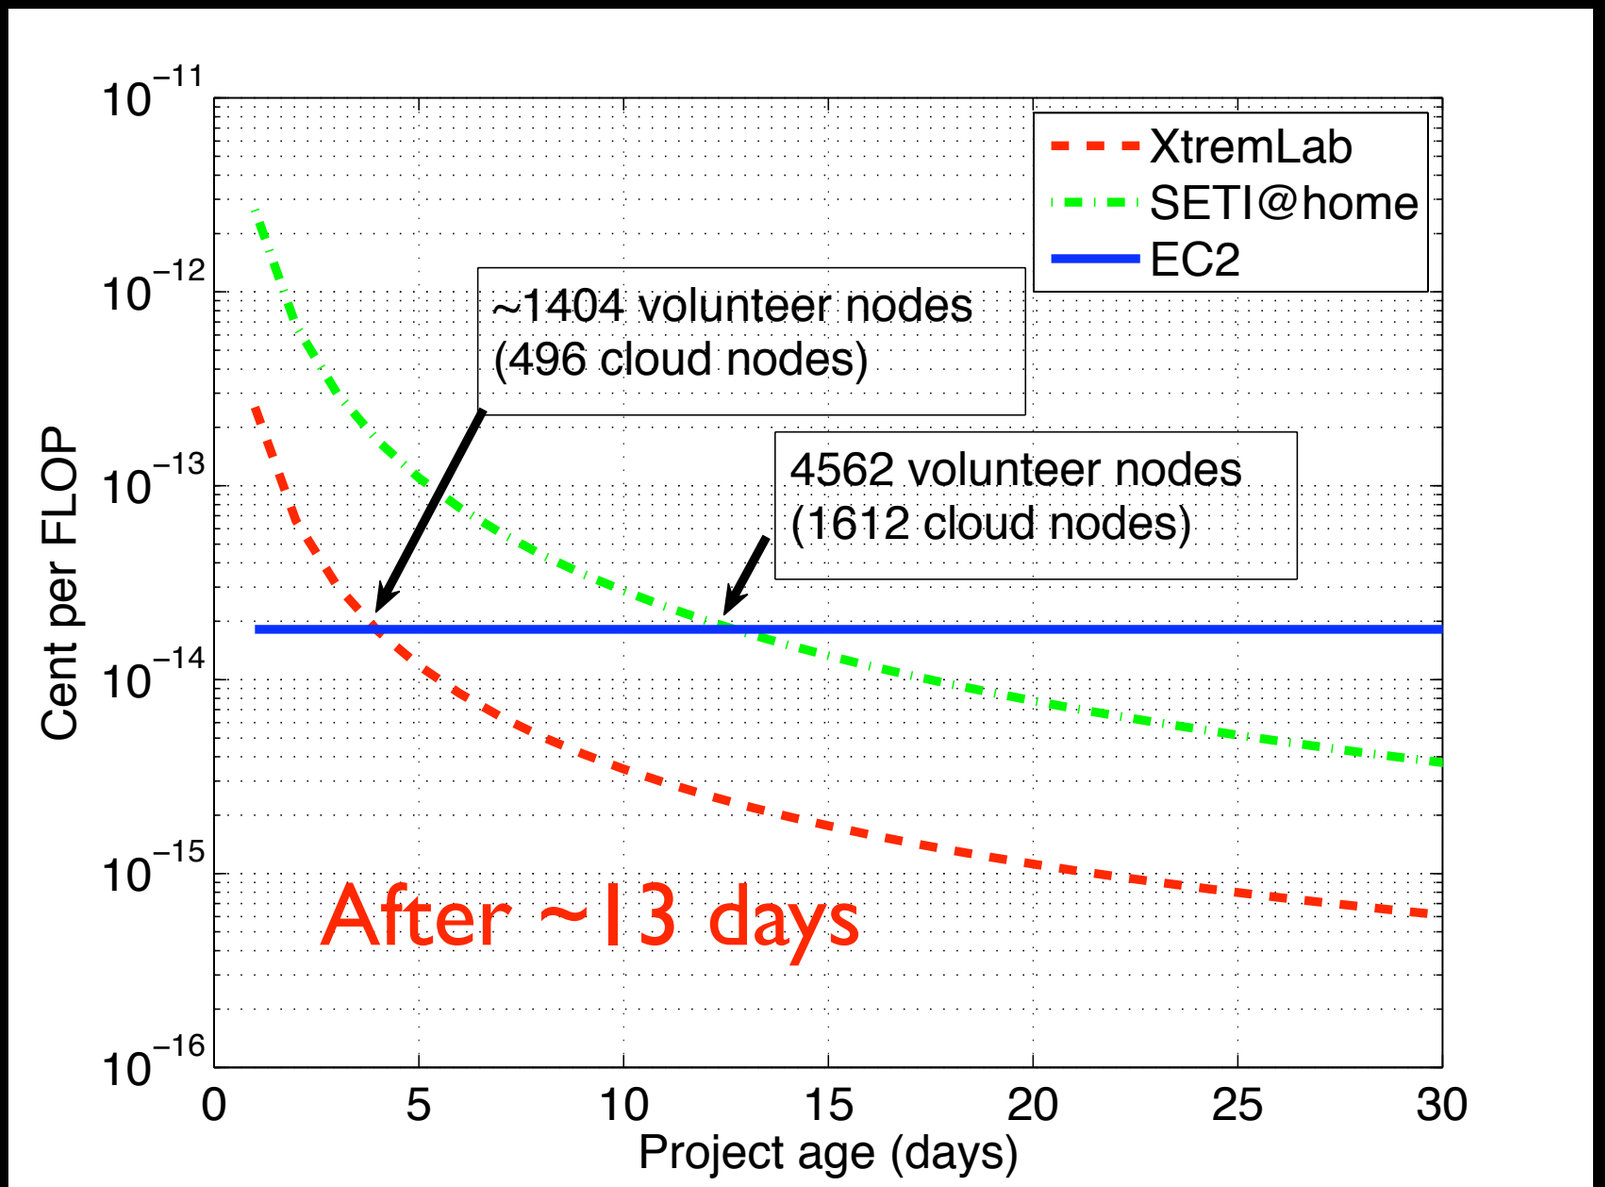

Cost of Clouds versus DG (2)

What are total costs over time?

Cost of Clouds versus DG (2)

What are total costs over time?

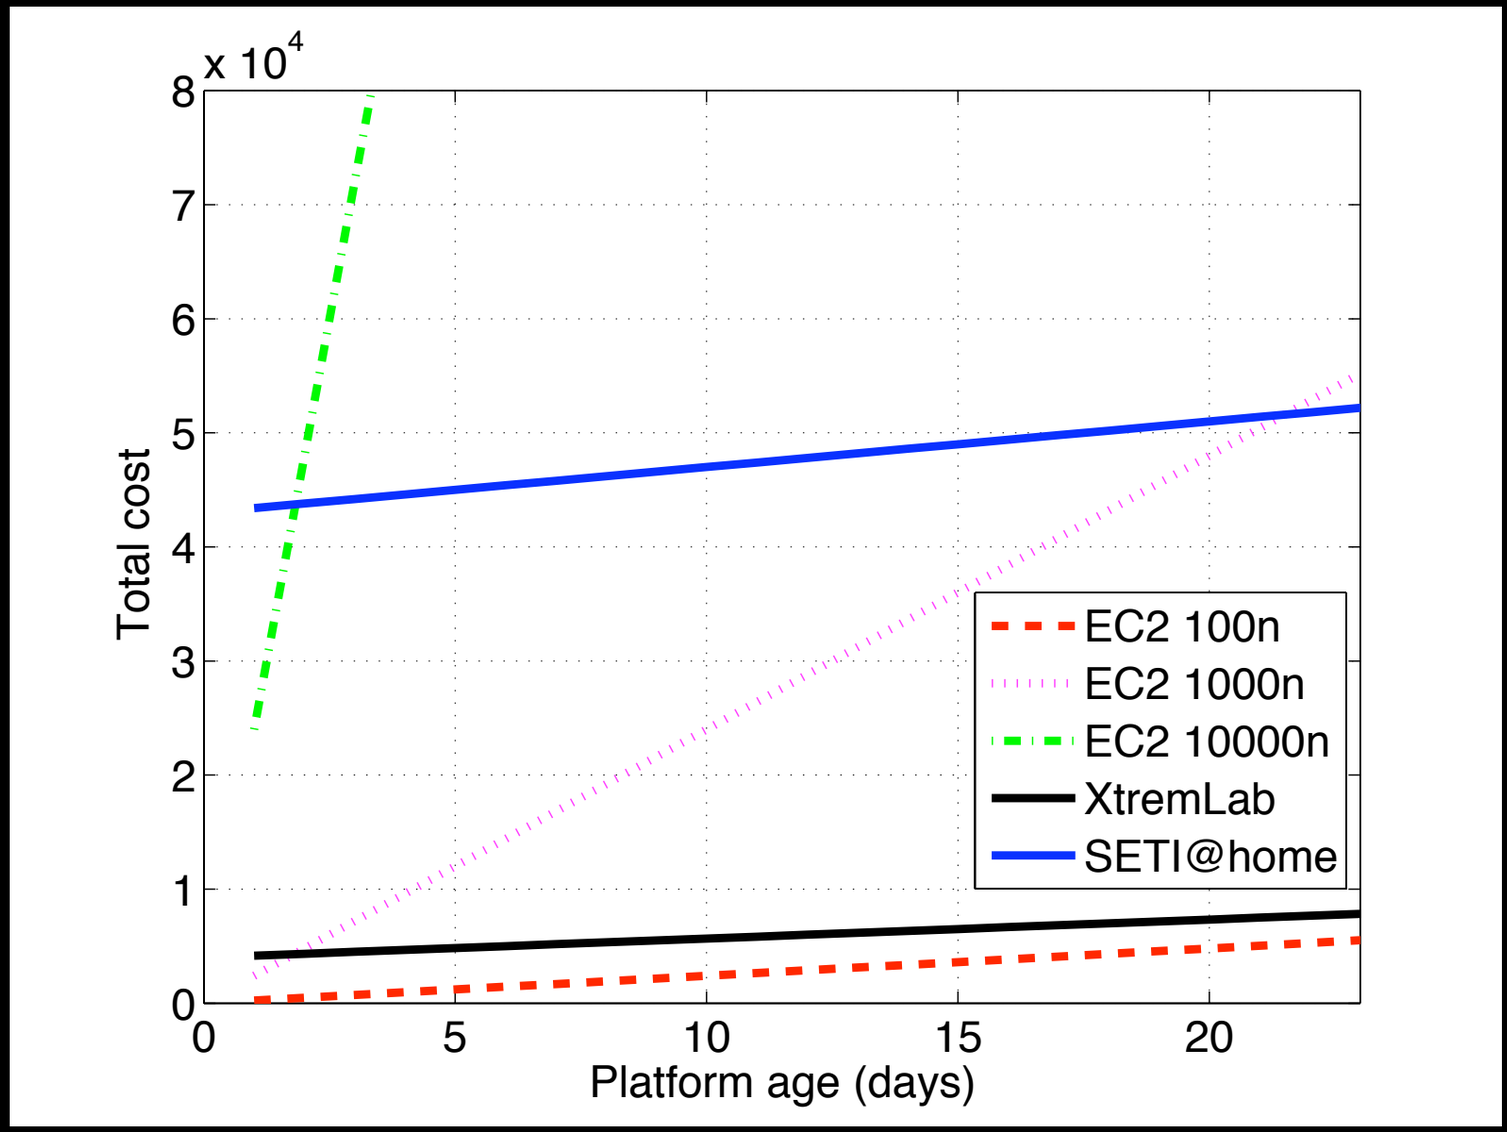

Cost of Clouds versus DG (2)

What are total costs over time?

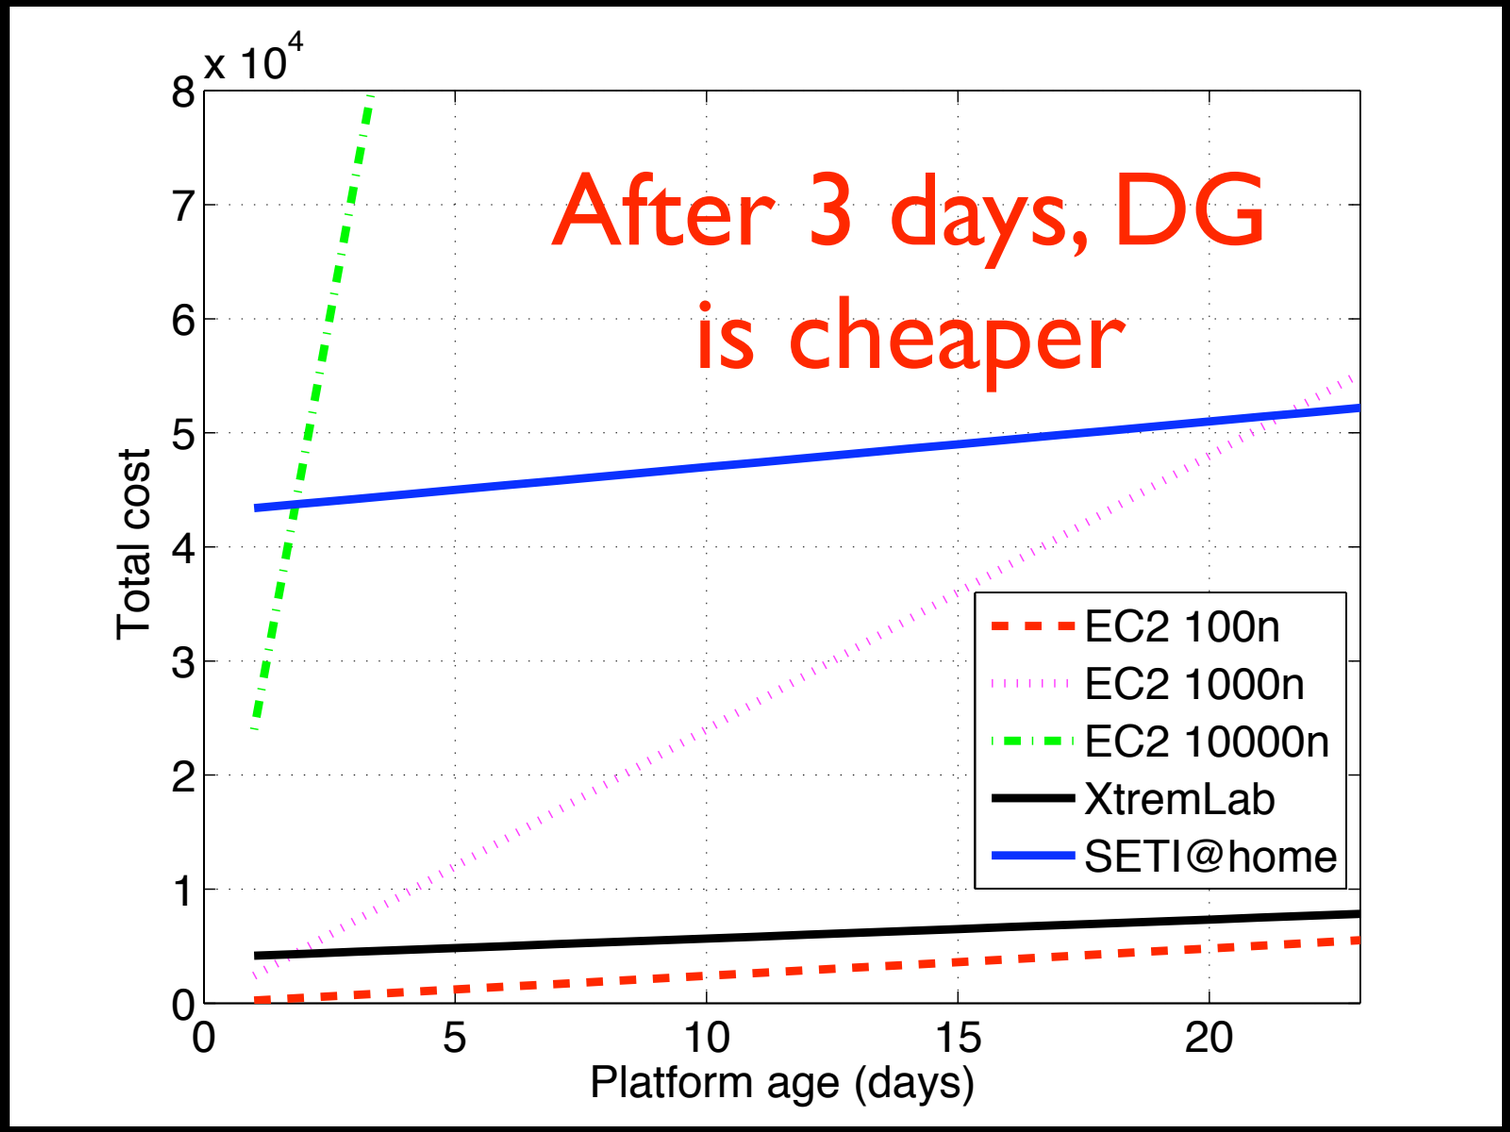

Cost of Clouds versus DG (3)

How many months
of DG can
X months of Cloud
buy me?

Cost of Clouds versus DG (3)

How many months of DG can X months of Cloud buy me?

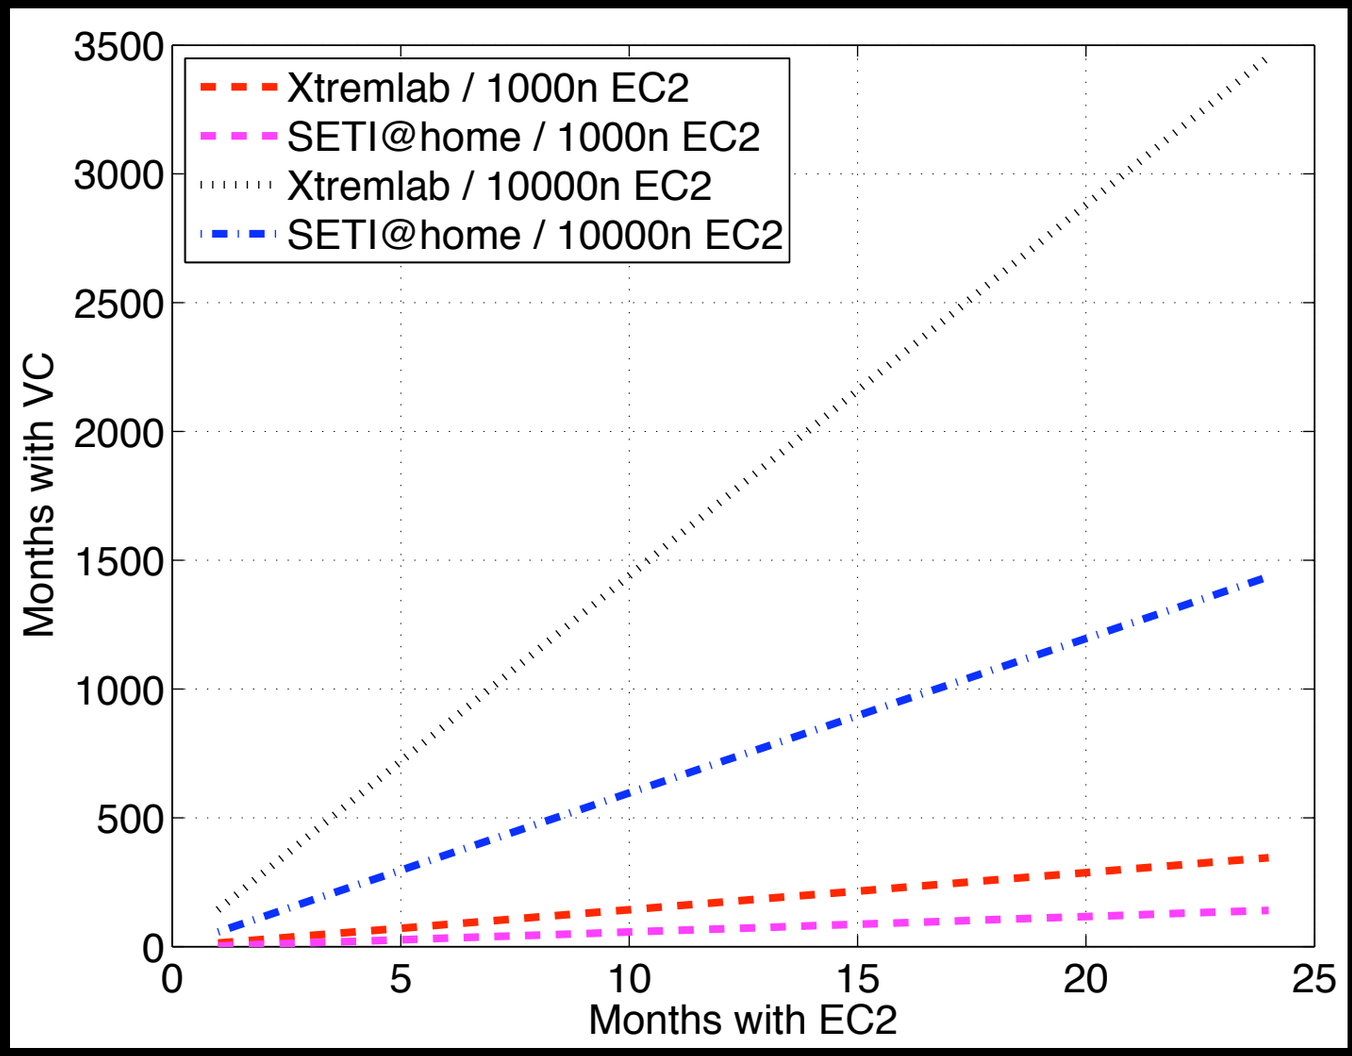

Cost of Clouds versus DG (3)

How many months of DG can X months of Cloud buy me?

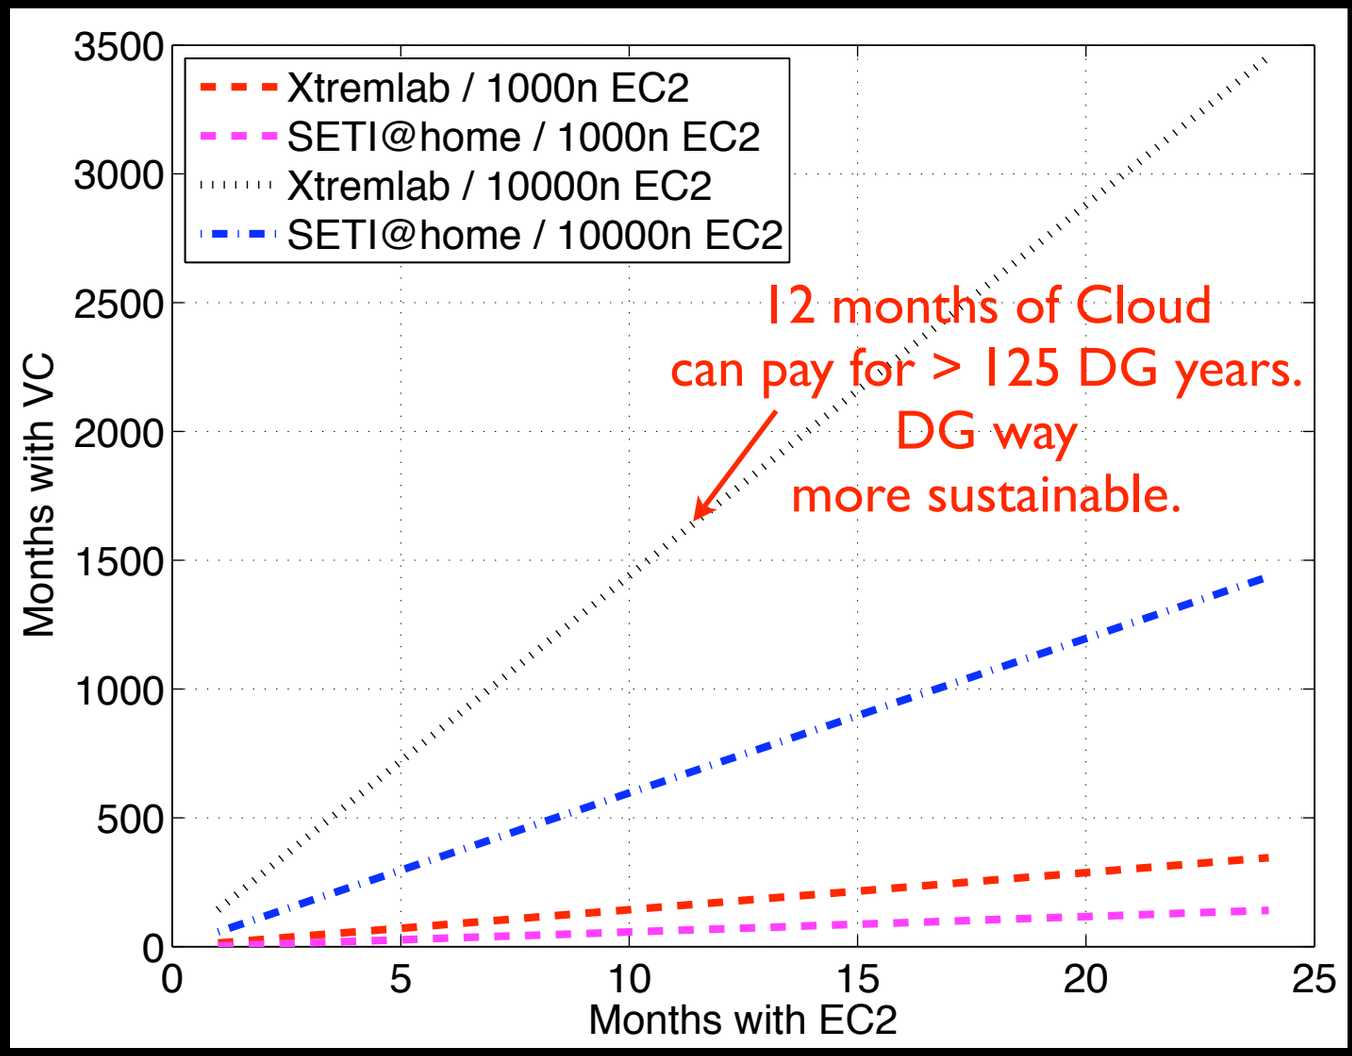

Equivalent Clouds given BOINC Costs

Given BOINC
costs, what size
Cloud can I buy?

Equivalent Clouds given BOINC Costs

Given BOINC costs, what size Cloud can I buy?

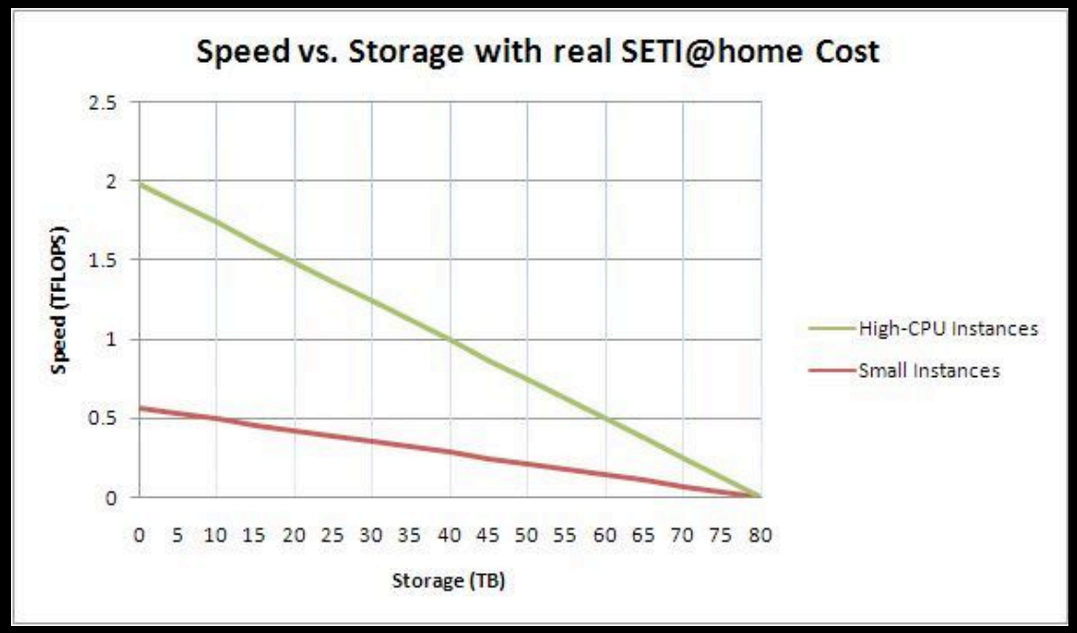

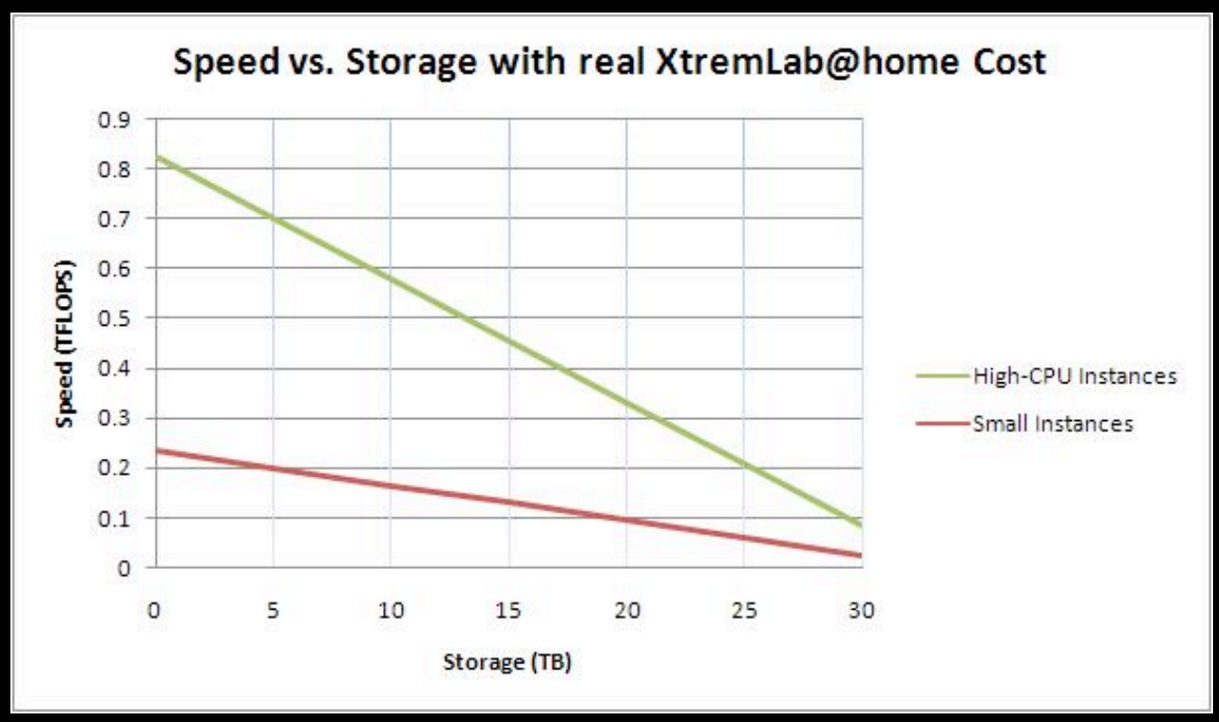

Equivalent Clouds given BOINC Costs

Given BOINC costs, what size Cloud can I buy?

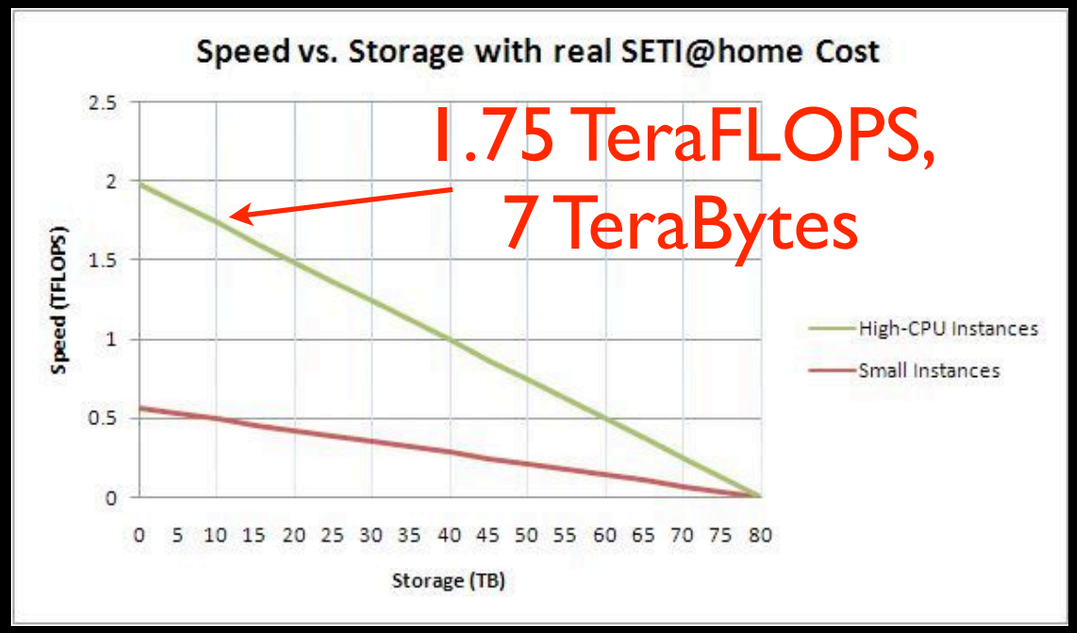

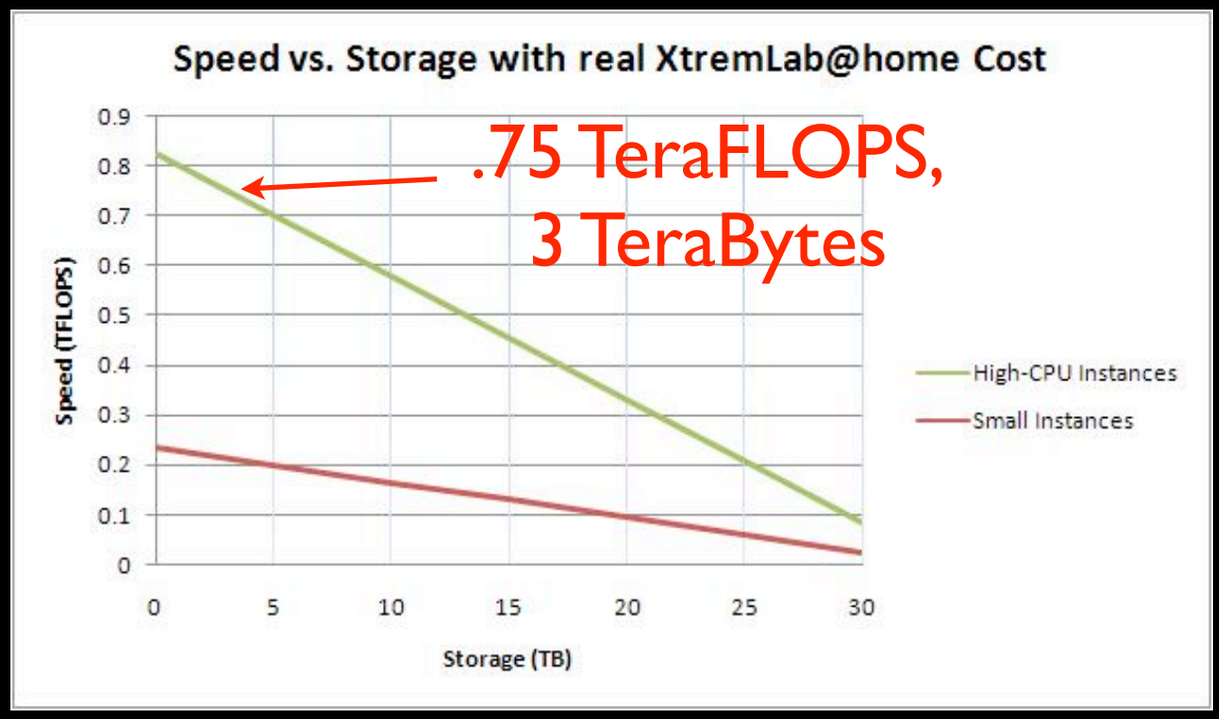

Equivalent Clouds given BOINC Costs

Given BOINC costs, what size Cloud can I buy?

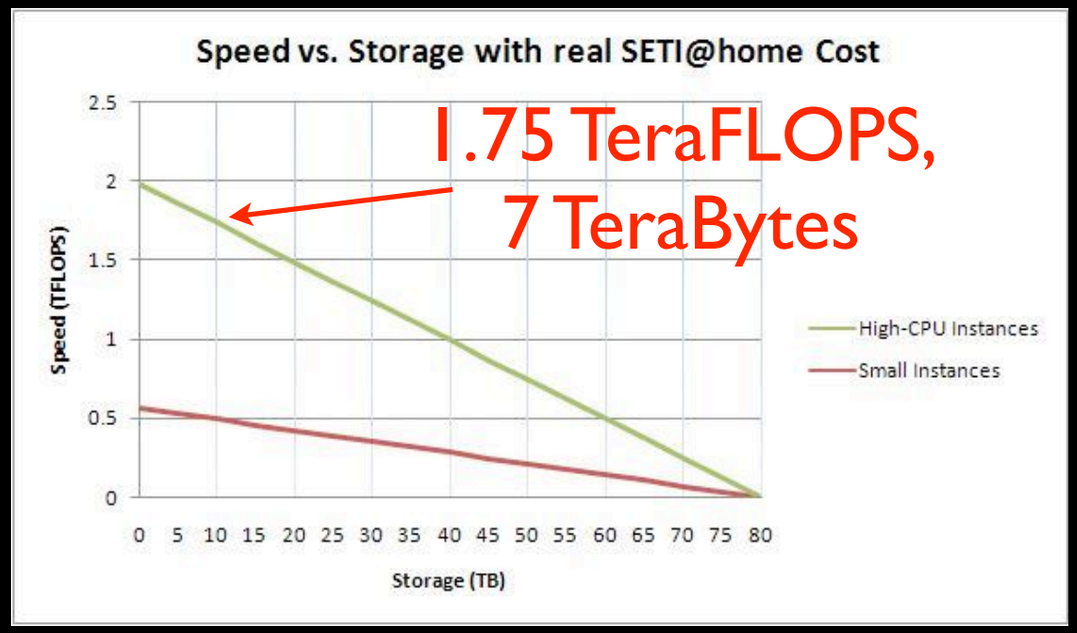

2 orders of magnitude lower than BOINC

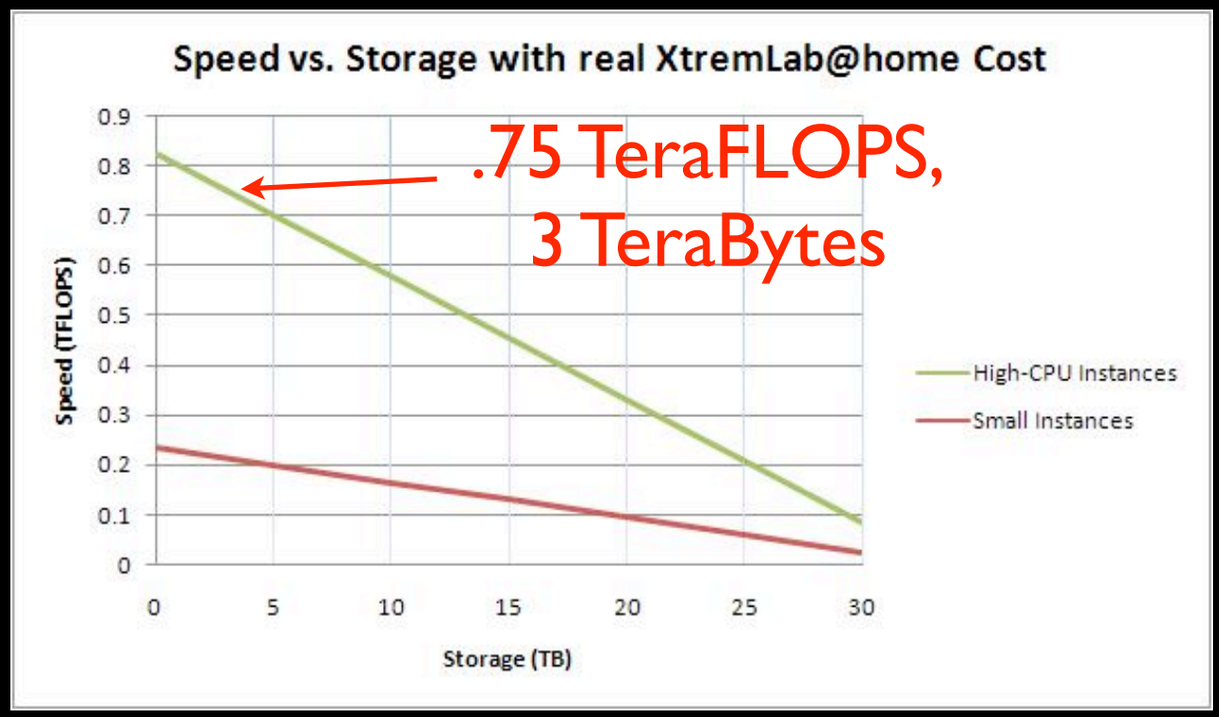

What about Server Hosting on the Cloud?

Number of Hosts over Time

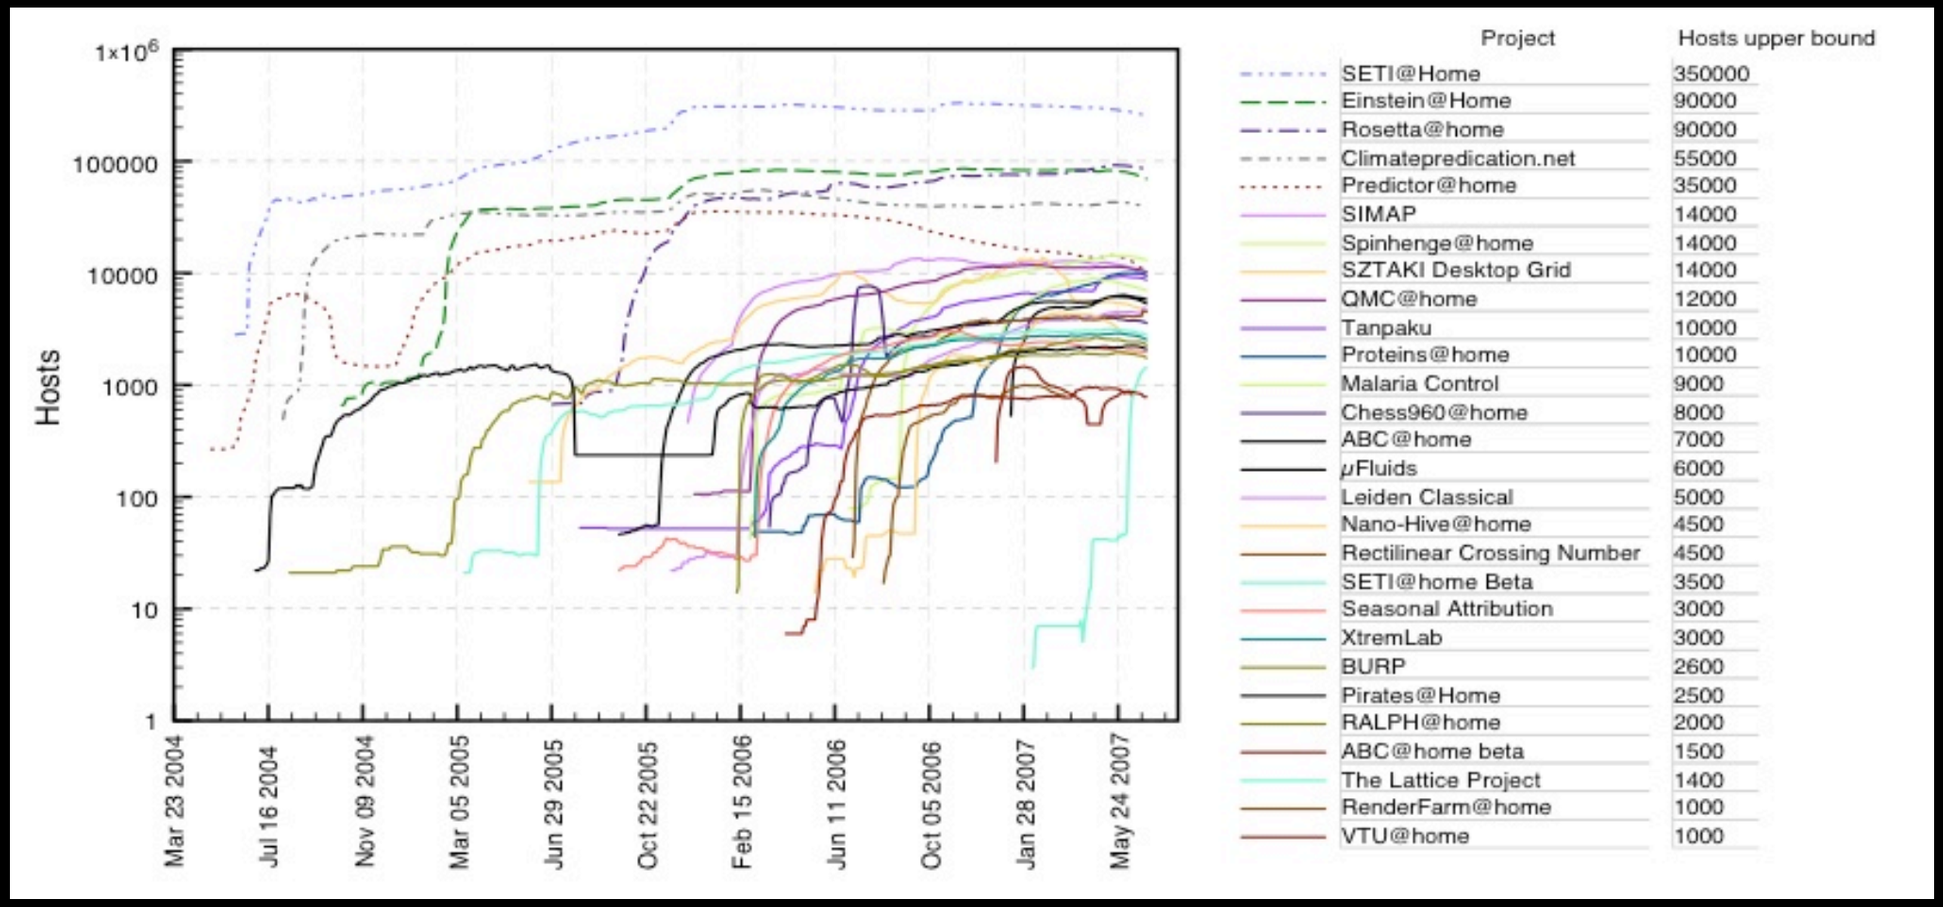

Number of Hosts over Time

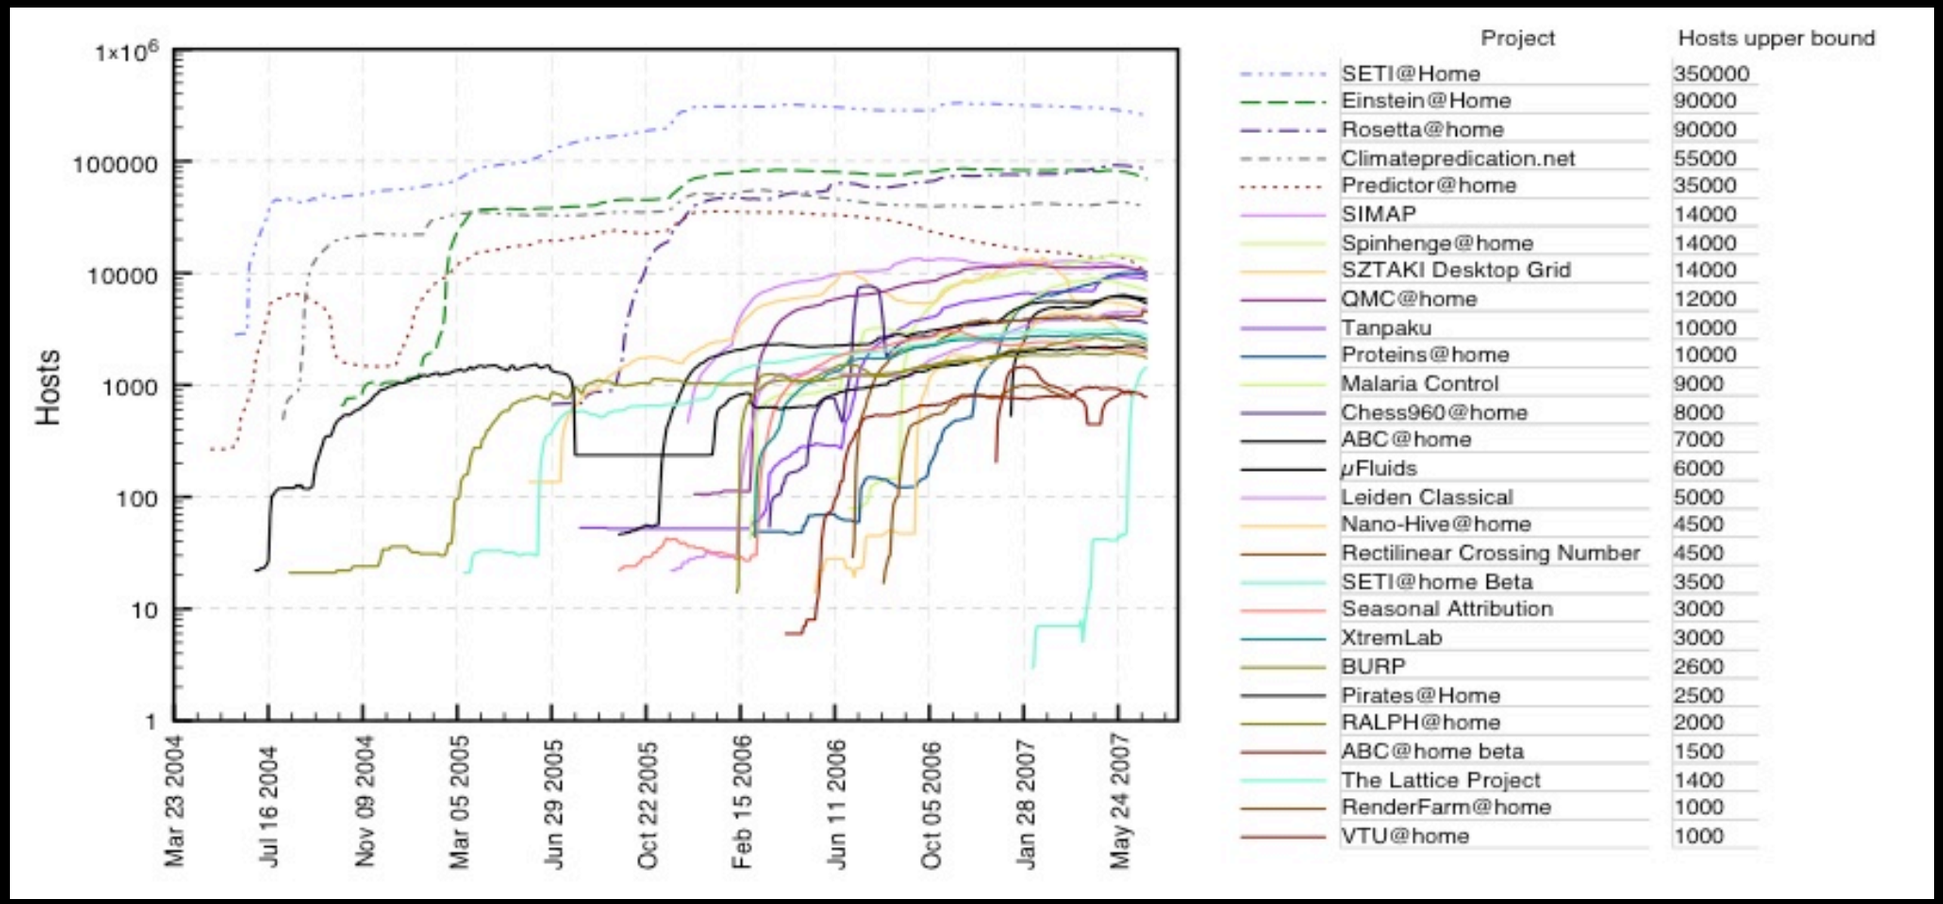

Project	Hosts upper bound
SETI@Home	350000
Einstein@Home	90000
Rosetta@home	90000
Climatepredication.net	55000
Predictor@home	35000
SIMAP	14000
Spinhenge@home	14000
SZTAKI Desktop Grid	14000
QMC@home	12000
Tanpaku	10000
Proteins@home	10000
Malaria Control	9000
Chess960@home	8000
ABC@home	7000
μFluids	6000
Leiden Classical	5000
Nano-Hive@home	4500
Rectilinear Crossing Number	4500
SETI@home Beta	3500
Seasonal Attribution	3000
XtremLab	3000
BURP	2600
Pirates@Home	2500
RALPH@home	2000
ABC@home beta	1500
The Lattice Project	1400
RenderFarm@home	1000
VTU@home	1000

Load variation exits w/ publicity, projects run out of work, etc. Clouds take care of server management. Potential to exploit clouds.

Server Costs on a Cloud

How much to host
BOINC server on cloud?

Server Costs on a Cloud

How much to host
BOINC server on cloud?

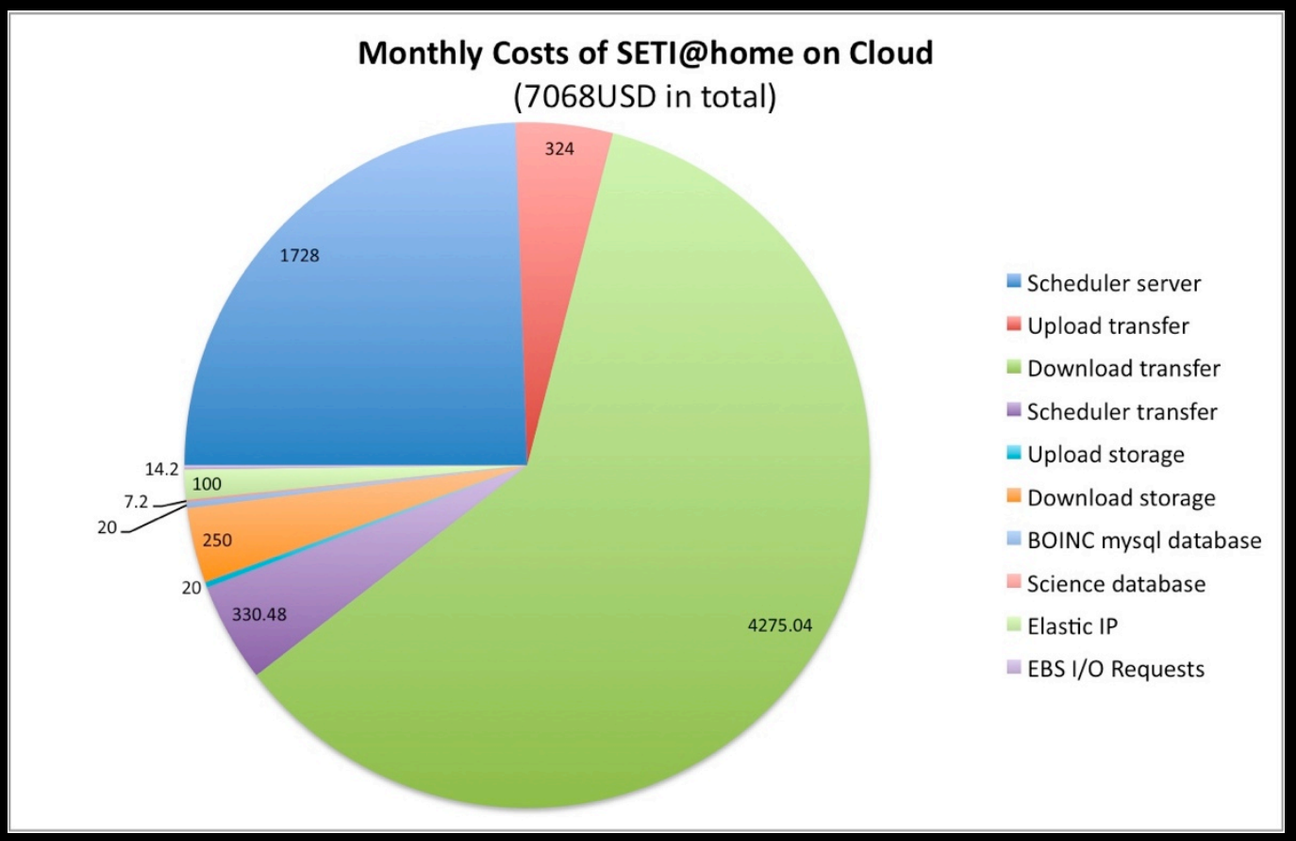

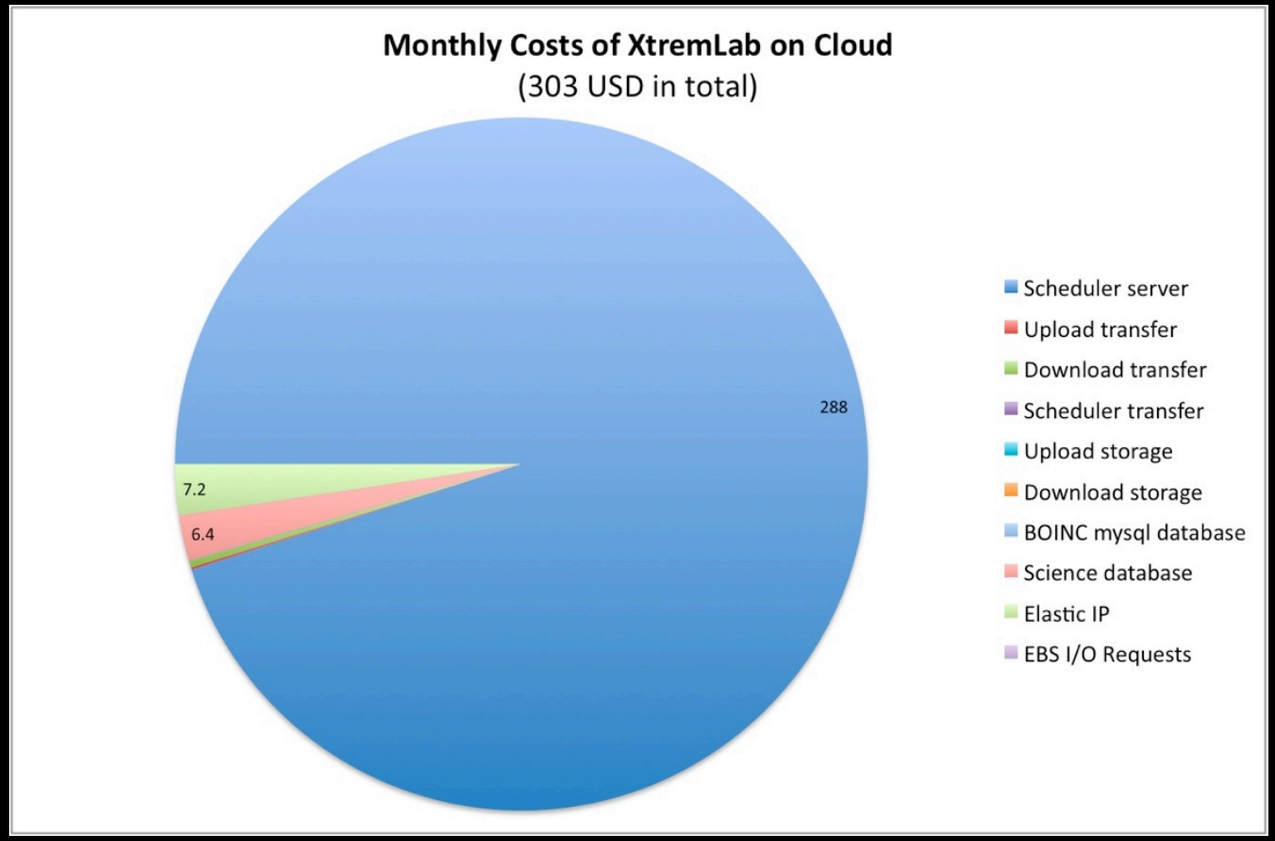

Server Costs on a Cloud

How much to host
BOINC server on cloud?

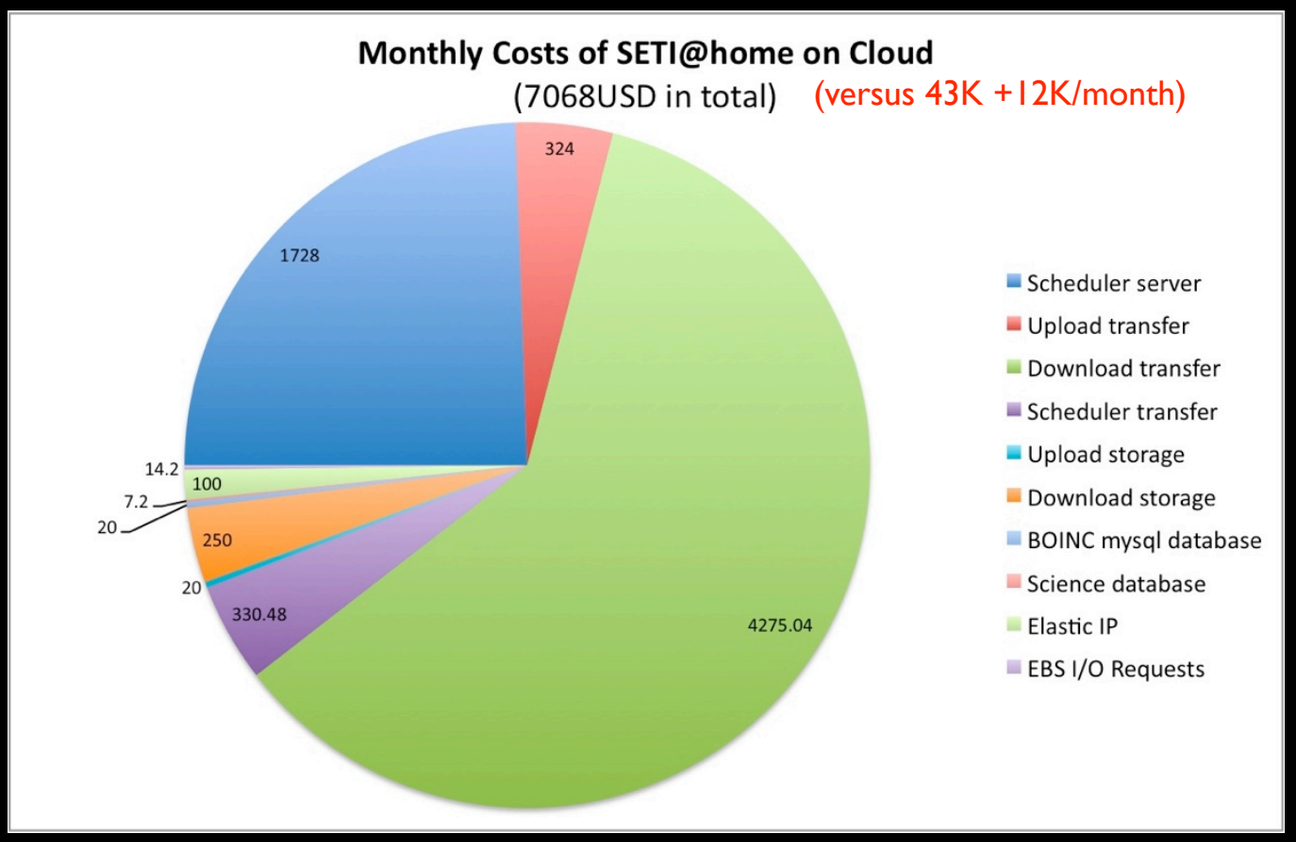

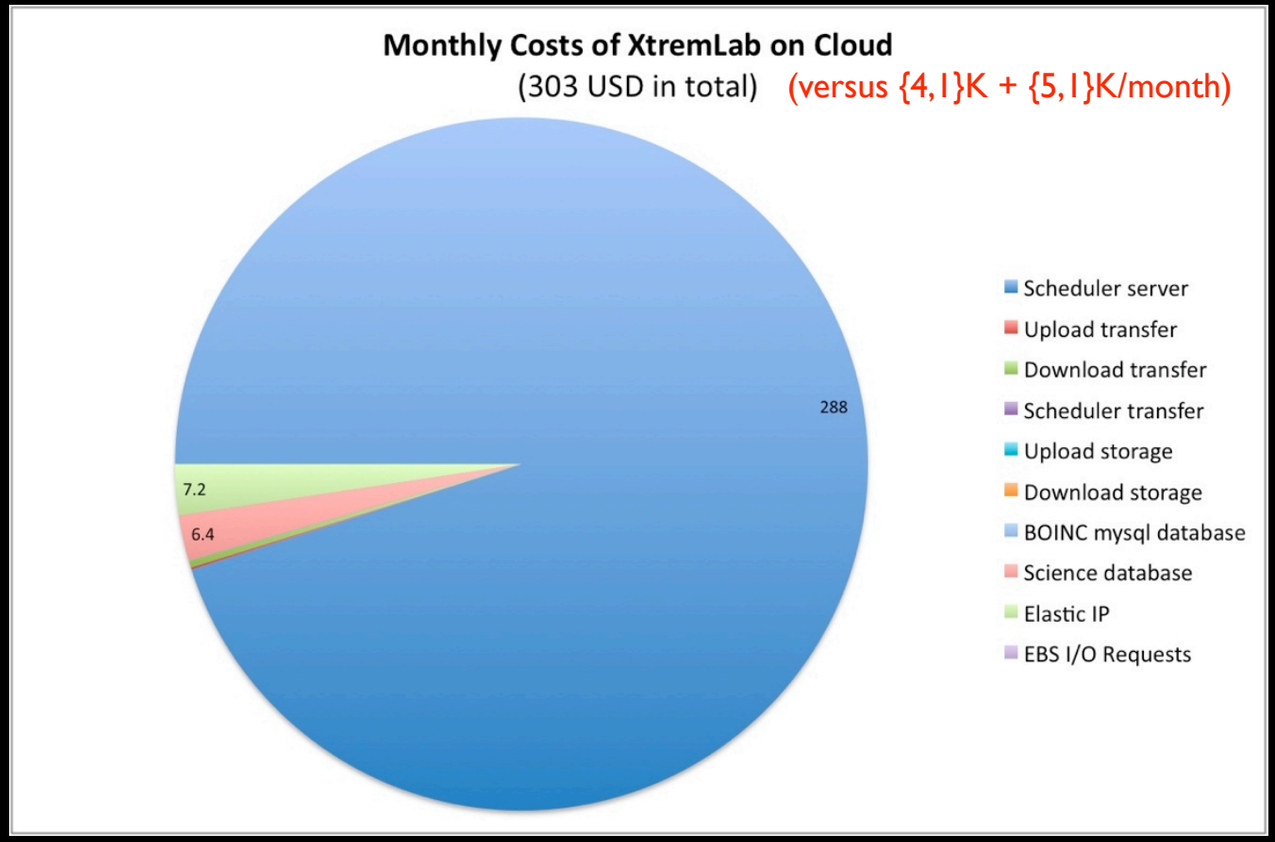

Server Costs on a Cloud

How much to host
BOINC server on cloud?

1.7 - 16.5 times cheaper to
host on Cloud,
but bandwidth is expensive.

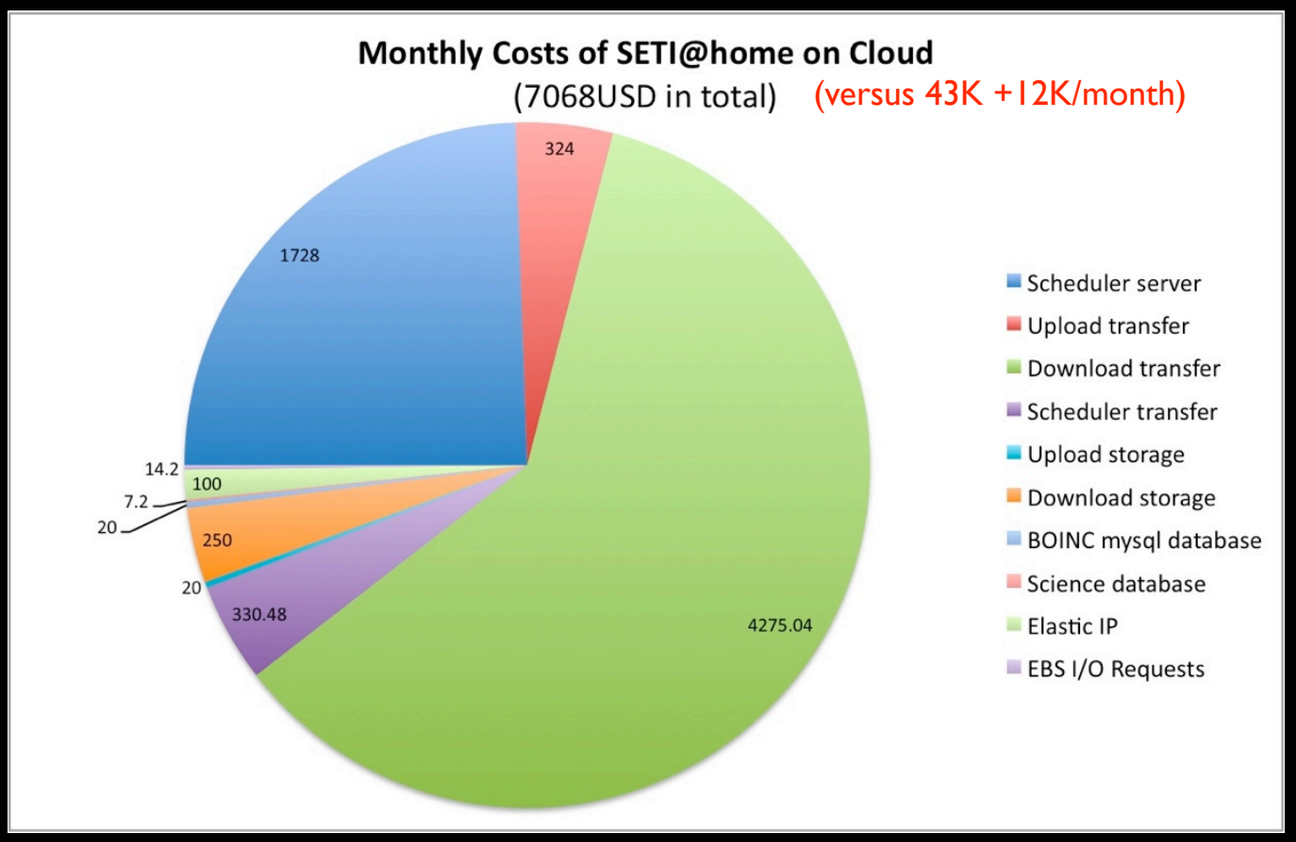

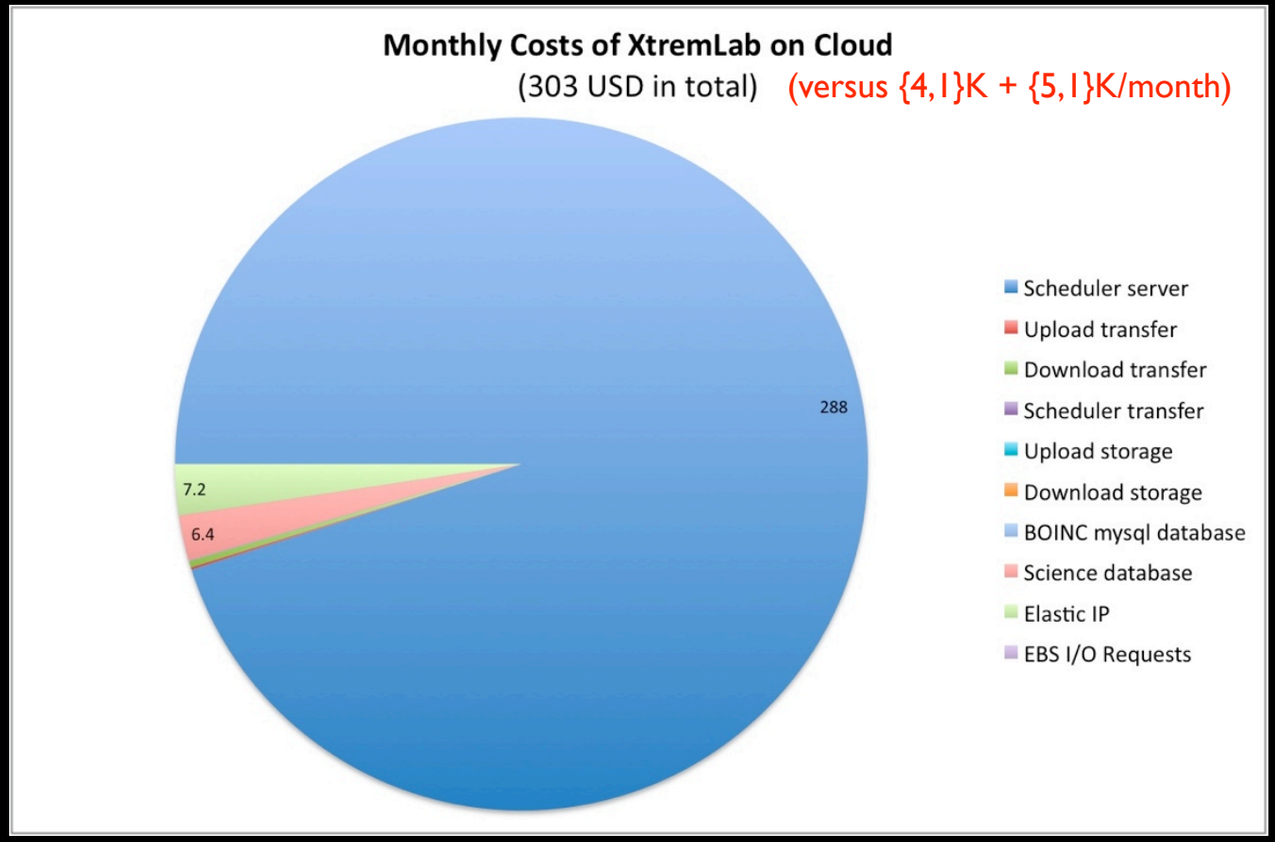

Cloud Resources with Given Budget

How big of a server can I support with given budget?

Cloud Resources with Given Budget

How big of a server can I support with given budget?

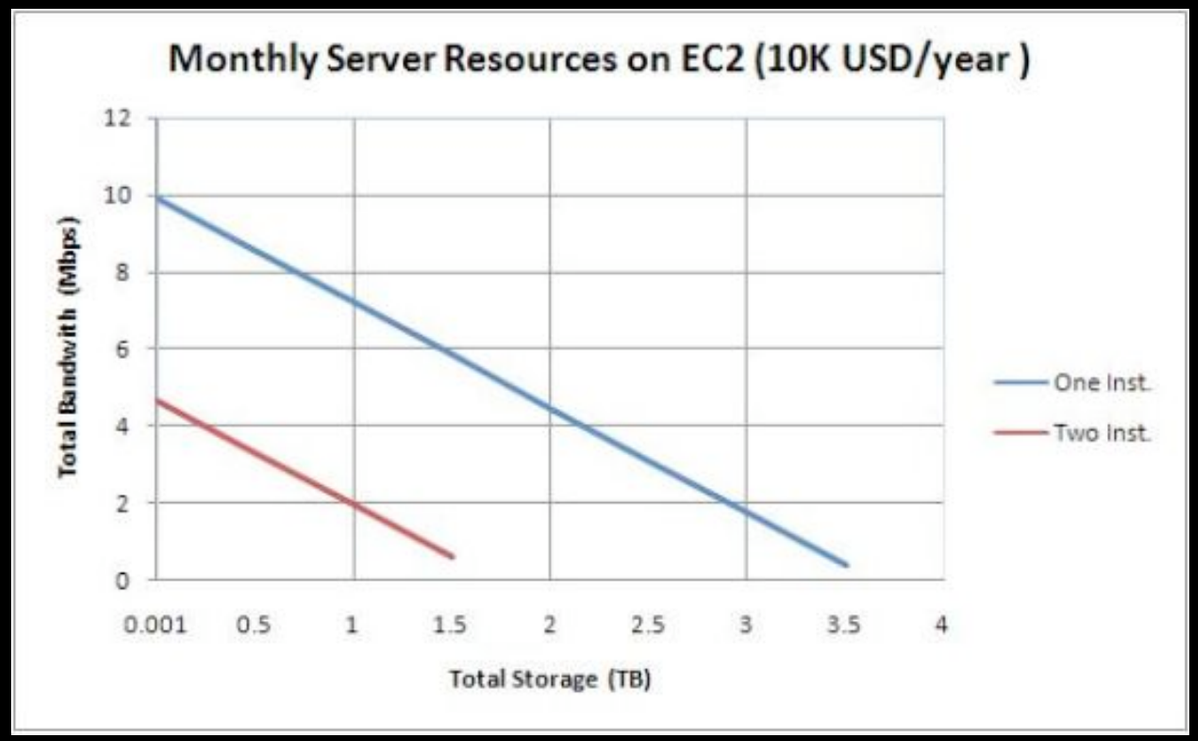

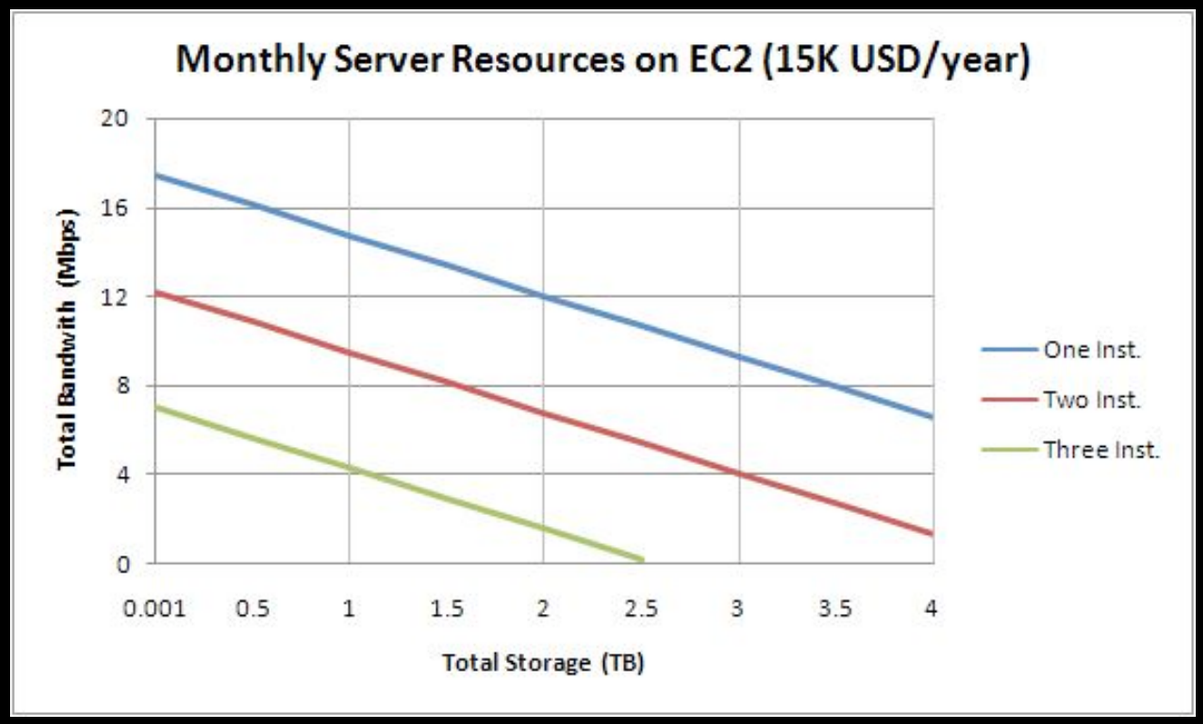

Cloud Resources with Given Budget

How big of a server can I support with given budget?

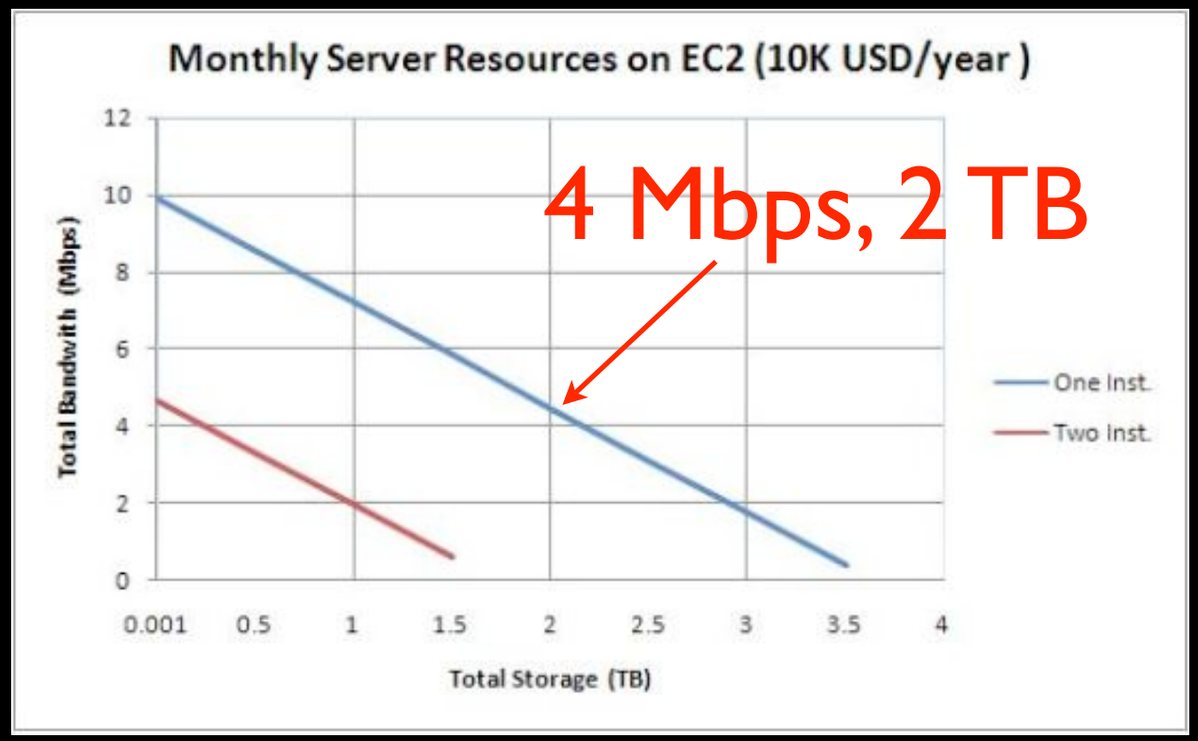

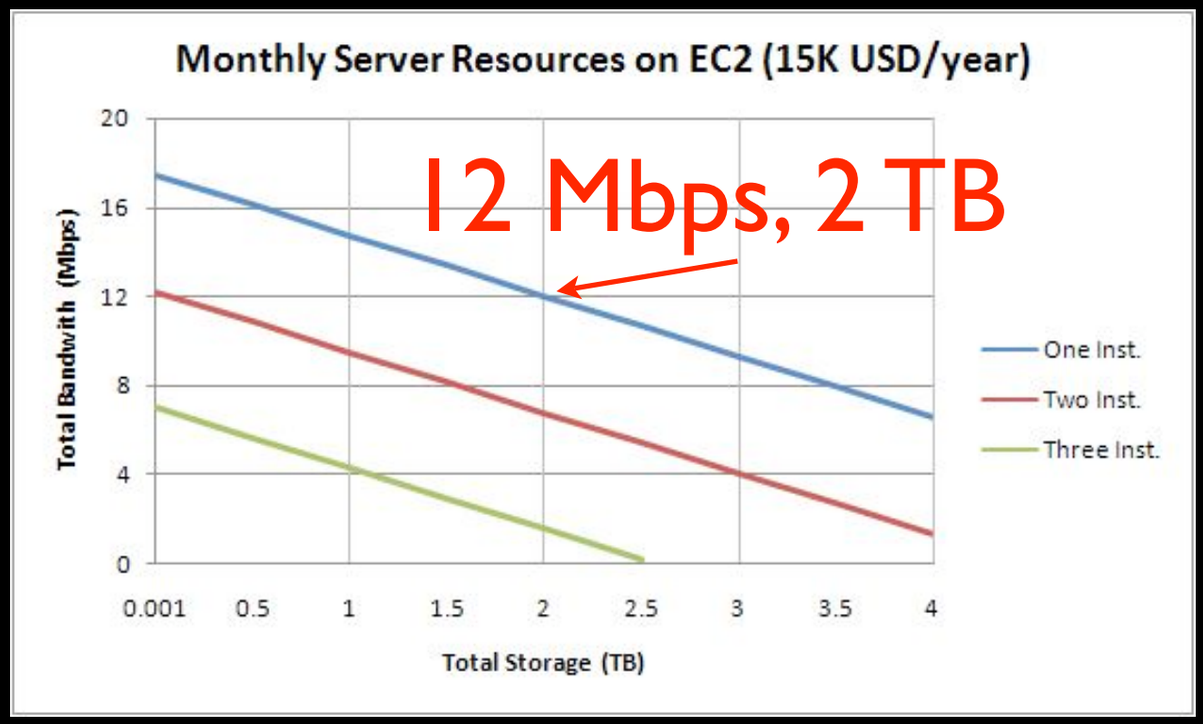

Cloud Resources with Given Budget

How big of a server can I support with given budget?

Many project servers are sustainable on the cloud

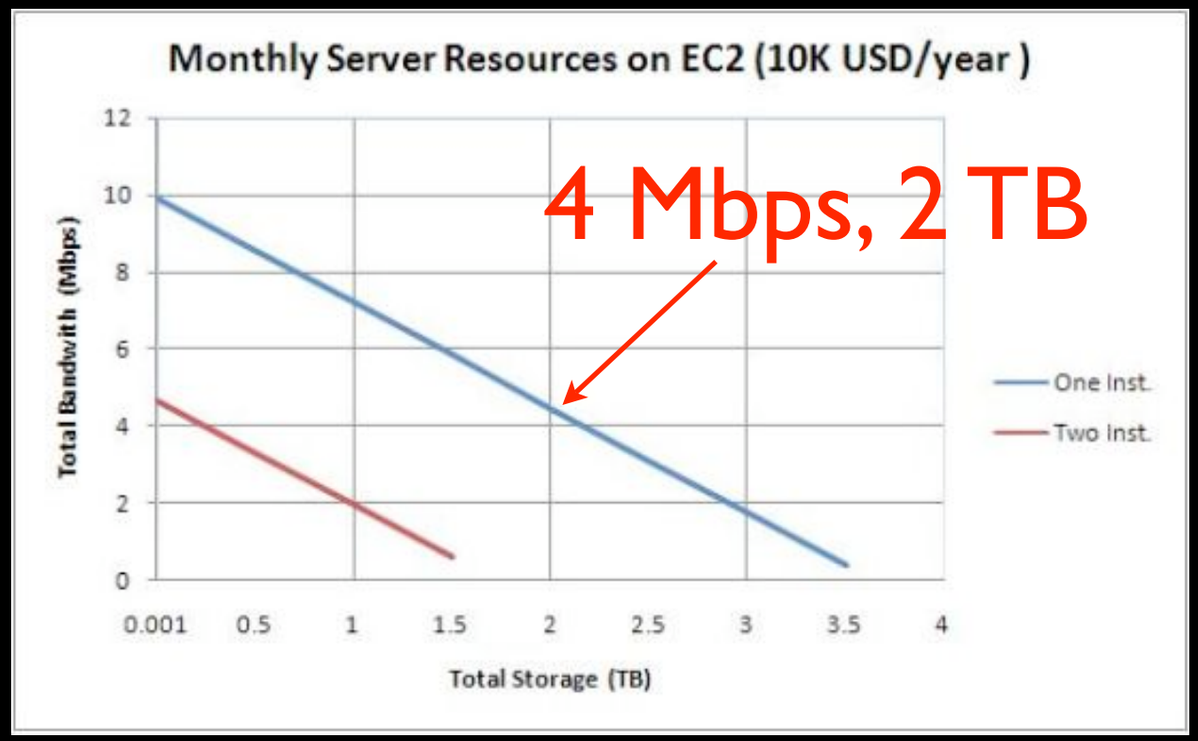

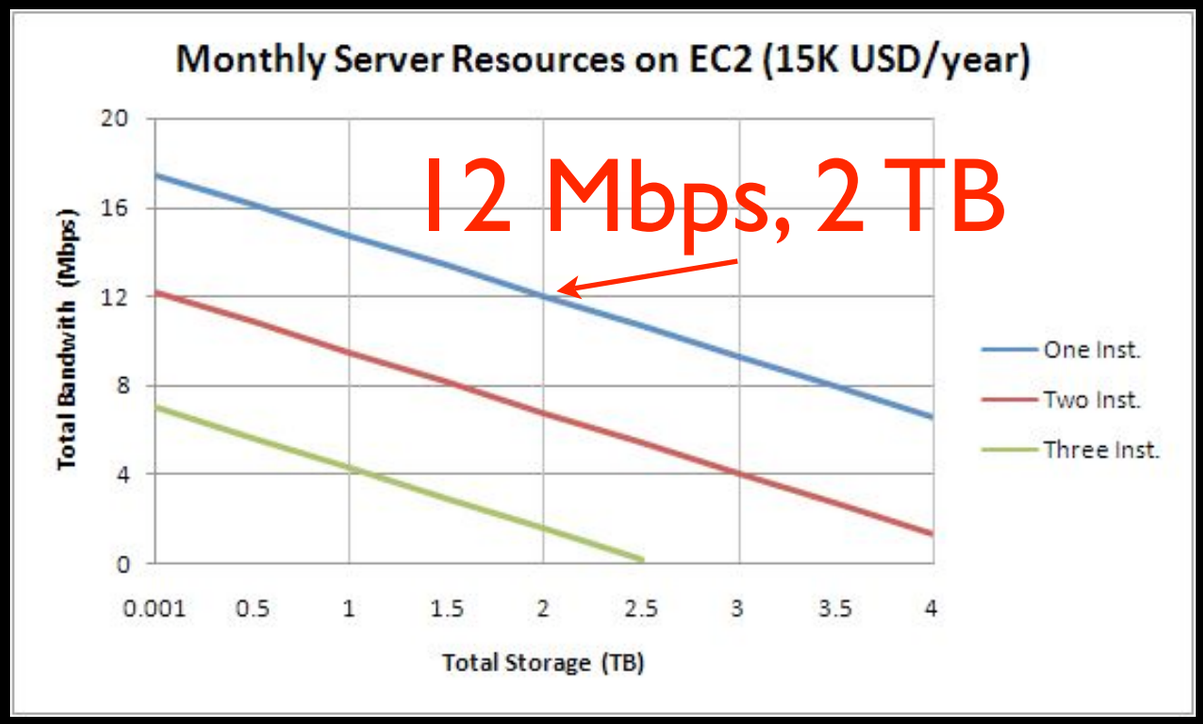

Summary

Summary

- Performance tradeoffs
  - 20 DG TeraFLOPS within 6 months

Summary

- Performance tradeoffs
  - 20 DG TeraFLOPS within 6 months
- Monetary tradeoffs
  - Client hosting
    - After 13 days, DG more cost effective
  - Server hosting
    - DG server on cloud is cost-effective
      - Best for small-medium sized projects
      - Savings of at least 40%
- http://mescal.imag.fr/membres/derrick.kondo/cloud_calc.xlsx

Summary

- Performance tradeoffs
  - 20 DG TeraFLOPS within 6 months
- Monetary tradeoffs
  - Client hosting
    - After 13 days, DG more cost effective
  - Server hosting
    - DG server on cloud is cost-effective
      - Best for small-medium sized projects
      - Savings of at least 40%
  - http://mescal.imag.fr/membres/derrick.kondo/cloud_calc.xlsx
- BOINC server image for EC2 / Xen available
  - http://boinc.berkeley.edu/trac/wiki/CloudServer

Future Work

- Clouds@home
- Statistical and predictive methods for ensuring host availability
- Lightweight mechanisms for virtual machine migration

- wu flop calc
- bw in cloud
- how is cloud equiv calculated