Flexible Data Processing Solutions for Space Missions

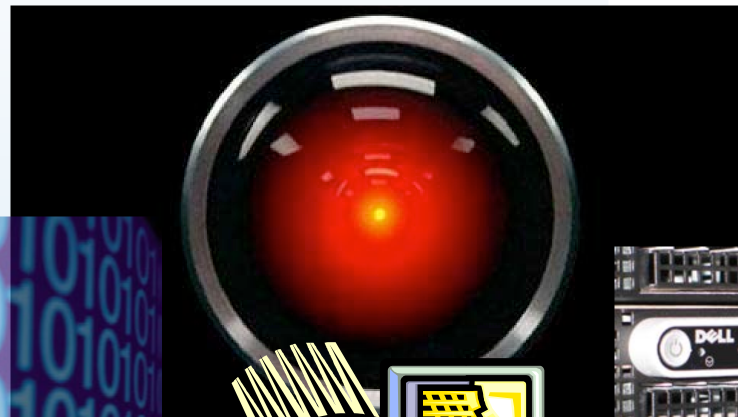

François Arago Centre

Volker Beckmann & Cécile Cavet
François Arago Centre
APC / Université Paris Diderot

CNRS IN2P3 Les deux infinis

UnivEarthS

APC

Challenges

Complex missions:

- data intensive
- CPU intensive
- processing context

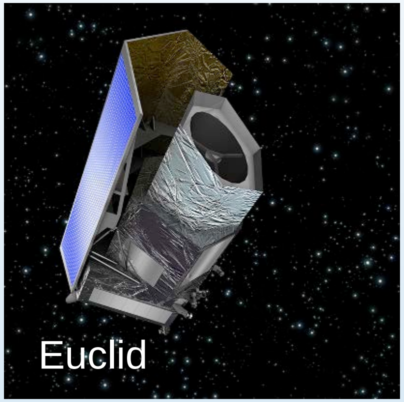

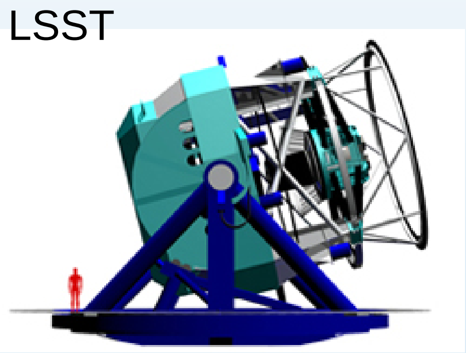

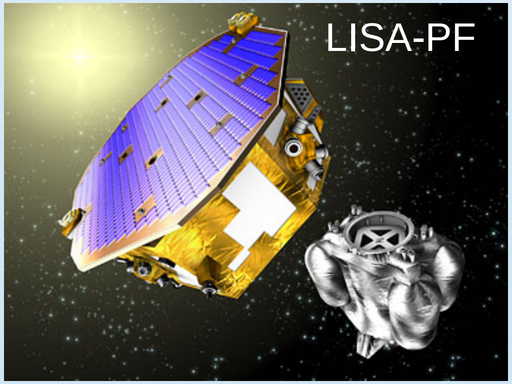

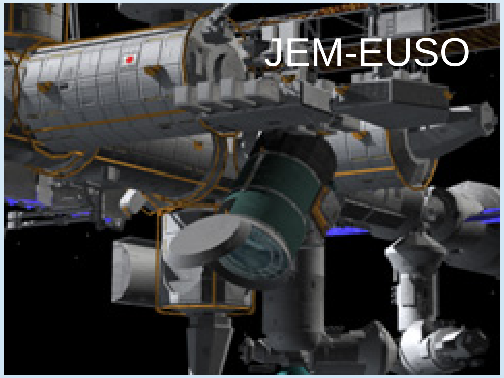

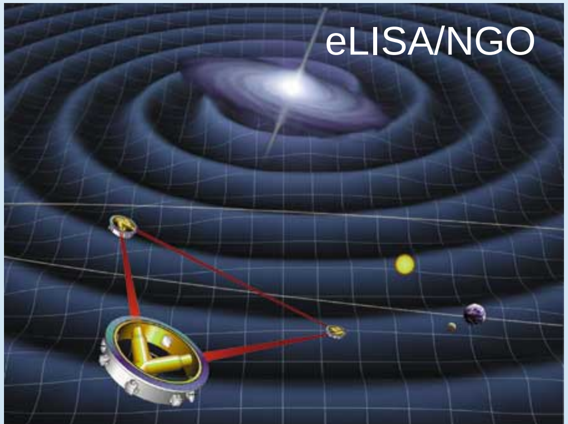

Solutions

Different processing options:

- Local cluster
- Several independent processing centres
- GRID
- Cloud infrastructure
- Cloud infrastructure (commercial)

Access is better than ownership

(Kevin Kelly, *Wired*)

Solutions

Different processing options:

- Local cluster

  e.g. ISDC services for INTEGRAL (see C. Ferrigno’s presentation)
- Several independent processing centres

  e.g. GAIA (S. Els’ presentation)
- GRID

  e.g. LHC (37 TByte/day → 11 Tier 1, 10-15 PByte/year), CTA
- Cloud infrastructure

  e.g. eLISA development (and SDC?)
- Cloud infrastructure (commercial)

  e.g. LHC for Higgs discovery, space?

GRID

- Several participating centres
- Same installation
- Middleware (e.g. gLite, EMI)
- high entry level
- EGI: Heavily supported by FP-6, FP-7, FP-8 (>100 M€ for EGEE and EGEE-II)
- Data intensive processing: needs dedicated infrastructure (e.g. LHC)
- Few, well connected large centres

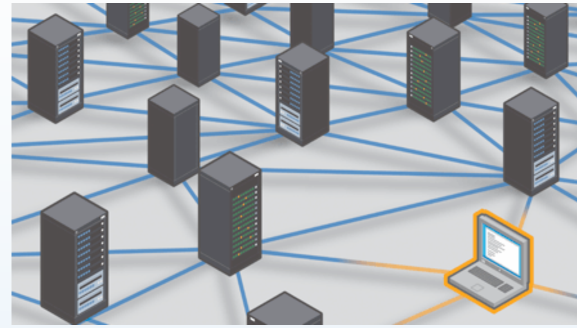

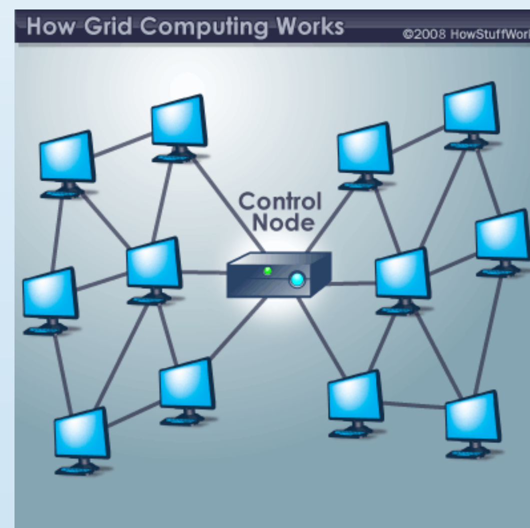

EGI: European Grid Infrastructure
EGEE: Enabling Grids for E-Science in Europe (FP-6 / FP-7)

Cloud

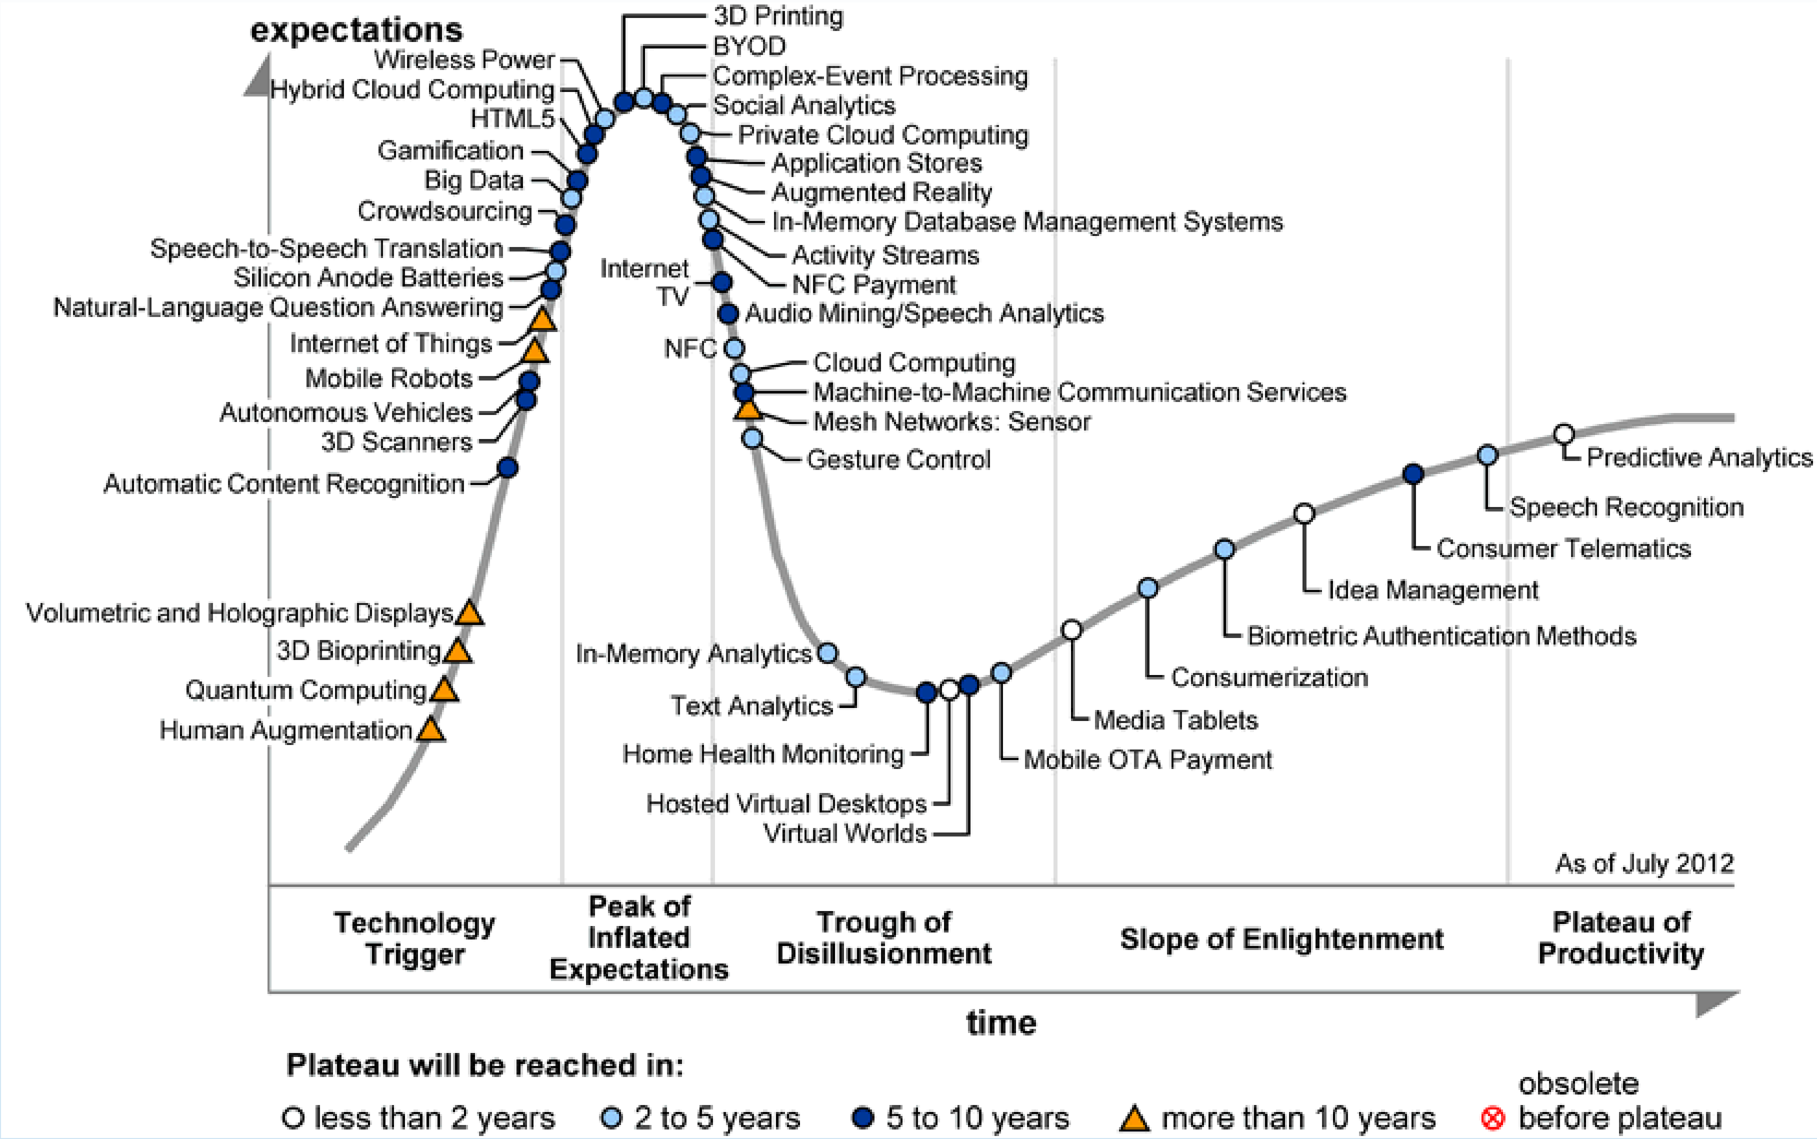

Gartner Inc., Hype Cycle, 2012

Cloud

Cloud computing faces skepticism (Shane Canon, Lawrence Berkeley National Lab):

- *Overhead to convert to Cloud environments*
- *Virtual instances underperform bare-metal systems*
- *Less cost effective than most large centers*

Distinguish between

- Commercial cloud
- Cloud as a virtualised infrastructure

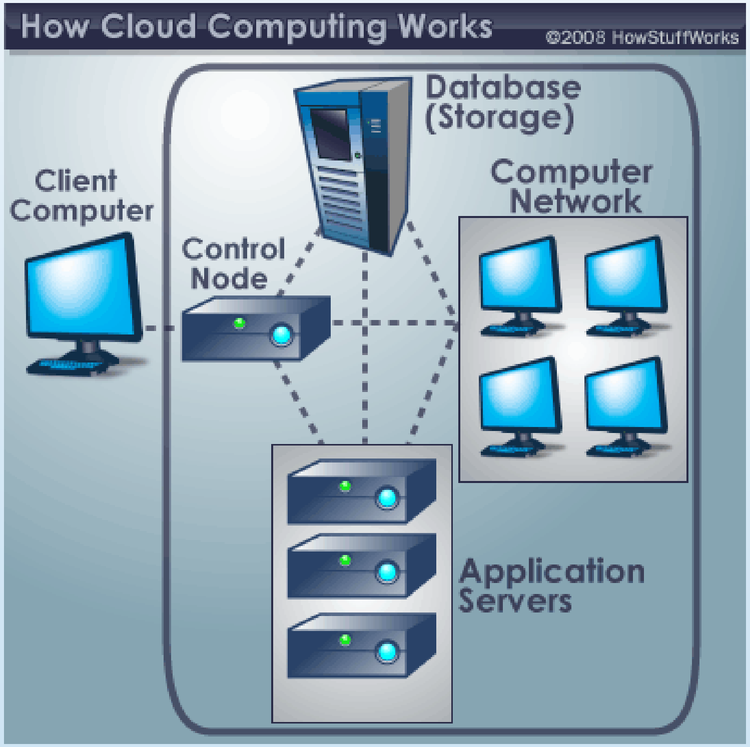

Cloud vs. Cluster

Overhead to convert to Cloud environments

Steps to be done

- Create a disk image of your operation system
- Upload it to cloud environment and
  set parameters of processing (#cores etc.)
- Install whatever s/w you like

Running your first task on the stratuslab cloud is not more challenging than to learn how to use the local cluster.

Disk images can be provided to consortium (e.g. Marketplace in stratuslab)

Cloud vs. Cluster

Virtual instances underperform bare-metal systems

Table 1. Physical machine comparison.

Description	StratusLab Cloud	Arago cluster
Nodes	10	11
Cores/node	24	16
Memory/node	36 GB	48 GB
Interconnexion network	1 GbE/s	10 GbE/s

How does the performance compare between a local cluster and a cloud environment?

Performance test

- Scaling: speedup, classical metric efficiency, Karp-Flatt metric efficiency
- Memory bandwidth
- I/O access
- Benchmarking: NASA paralell benchmark (NPB)
- High Performance Linpack (HPL)

Cloud vs. Cluster

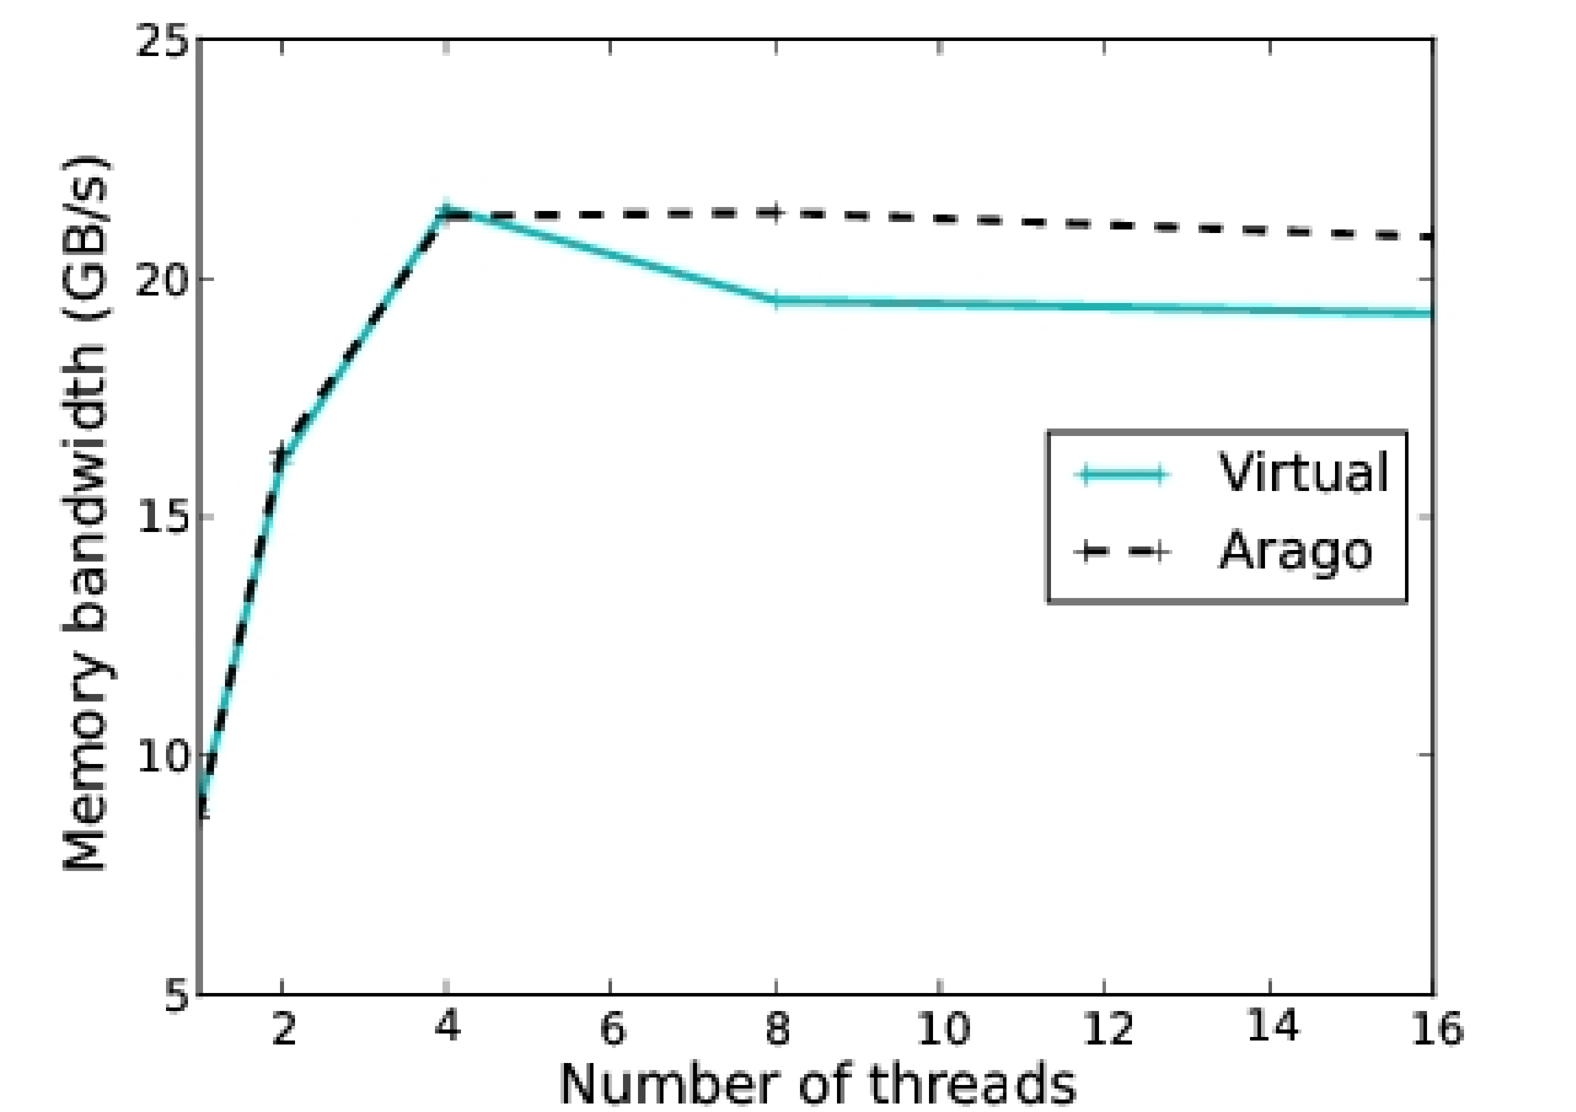

Cloud & cluster both approach band-width saturation in a similar fashion

Cloud vs. Cluster

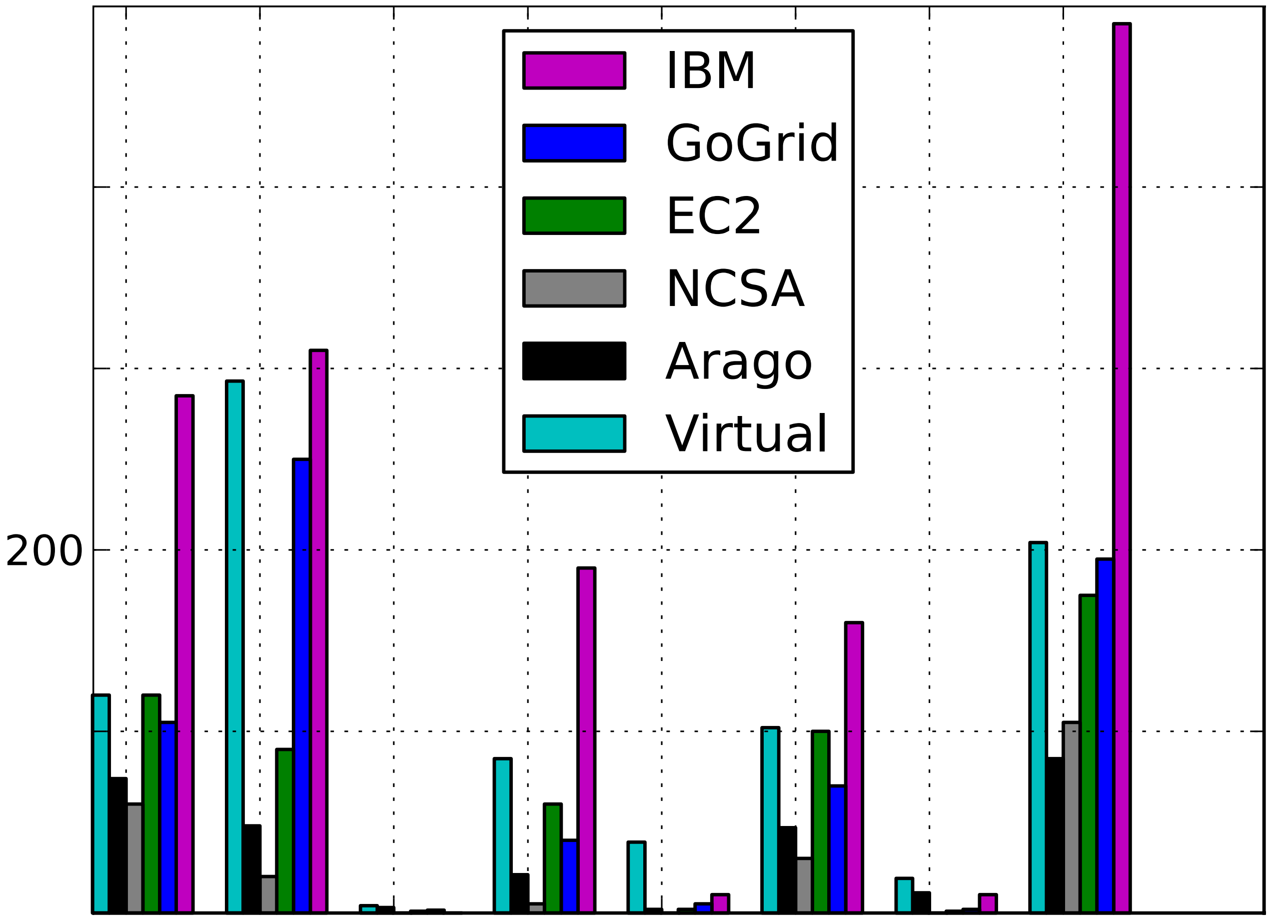

Cloud environments under-perform for processes with large inter-node message transfer

Cloud vs. Cluster

Table 3. Summary of processing resources on Amazon EC2.

type	arch.	CPU	cores	memory (GB)	network	storage	price
m1.small	32 bit	2.0–2.6 GHz Opteron	1/2	1.7	1 Gbps Ethernet	local	US$0.10 h^{-1}
m1.large	64 bit	2.0–2.6 GHz Opteron	2	7.5	1 Gbps Ethernet	local	US$0.40 h^{-1}
m1.xlarge	64 bit	2.0–2.6 GHz Opteron	4	15.0	1 Gbps Ethernet	local	US$0.80 h^{-1}
c1.medium	32 bit	2.33–2.66 GHz Xeon	2	1.7	1 Gbps Ethernet	local	US$0.20 h^{-1}
c1.xlarge	64 bit	2.0–2.66 GHz Xeon	8	7.5	1 Gbps Ethernet	local	US$0.80 h^{-1}

Table 4. Summary of processing resources on the Abe high-performance cluster.

type	arch.	CPU	cores	memory (GB)	network	storage
abe.local	64 bit	2.33 GHz Xeon	8	8	10 Gbps InfiniBand	local
abe.lustre	64 bit	2.33 GHz Xeon	8	8	10 Gbps InfiniBand	lustre

Berriman et al. 2013, « The application of cloud computing to scientific workflows: a study of cost and performance », Phil. Trans. R. Soc. A 2013 371

Cloud vs. Cluster

Table 3. Summary of processing resources on Amazon EC2.

type	arch.	CPU	cores	memory (GB)	network	storage	price
m1.small	32 bit	2.0–2.6 GHz Opteron	1/2	1.7	1 Gbps Ethernet	local	US$0.10 h^{-1}
m1.large	64 bit	2.0–2.6 GHz Opteron	2	7.5	1 Gbps Ethernet	local	US$0.40 h^{-1}
m1.xlarge	64 bit	2.0–2.6 GHz Opteron	4	15.0	1 Gbps Ethernet	local	US$0.80 h^{-1}
c1.medium	32 bit	2.33–2.66 GHz Xeon	2	1.7	1 Gbps Ethernet	local	US$0.20 h^{-1}
c1.xlarge	64 bit	2.0–2.66 GHz Xeon	8	7.5	1 Gbps Ethernet	local	US$0.80 h^{-1}

Table 4. Summary of processing resources on the Abe high-performance cluster.

type	arch.	CPU	cores	memory (GB)	network	storage
abe.local	64 bit	2.33 GHz Xeon	8	8	10 Gbps InfiniBand	local
abe.lustre	64 bit	2.33 GHz Xeon	8	8	10 Gbps InfiniBand	lustre

Berriman et al. 2013, « The application of cloud computing to scientific workflows: a study of cost and performance », Phil. Trans. R. Soc. A 2013 371

Cloud vs. Cluster

Table 3. Summary of processing resources on Amazon EC2.

type	arch.	CPU	cores	memory (GB)	network	storage	price
m1.small	32 bit	2.0–2.6 GHz Opteron	1/2	1.7	1 Gbps Ethernet	local	US$0.10 h^{-1}
m1.large	64 bit	2.0–2.6 GHz Opteron	2	7.5	1 Gbps Ethernet	local	US$0.40 h^{-1}
m1.xlarge	64 bit	2.0–2.6 GHz Opteron	4	15.0	1 Gbps Ethernet	local	US$0.80 h^{-1}
c1.medium	32 bit	2.33–2.66 GHz Xeon	2	1.7	1 Gbps Ethernet	local	US$0.20 h^{-1}
c1.xlarge	64 bit	2.0–2.66 GHz Xeon	8	7.5	1 Gbps Ethernet	local	US$0.80 h^{-1}

Table 4. Summary of processing resources on the Abe hi

type	arch.	CPU	core
abe.local	64 bit	2.33 GHz Xeon	8
abe.lustre	64 bit	2.33 GHz Xeon	8

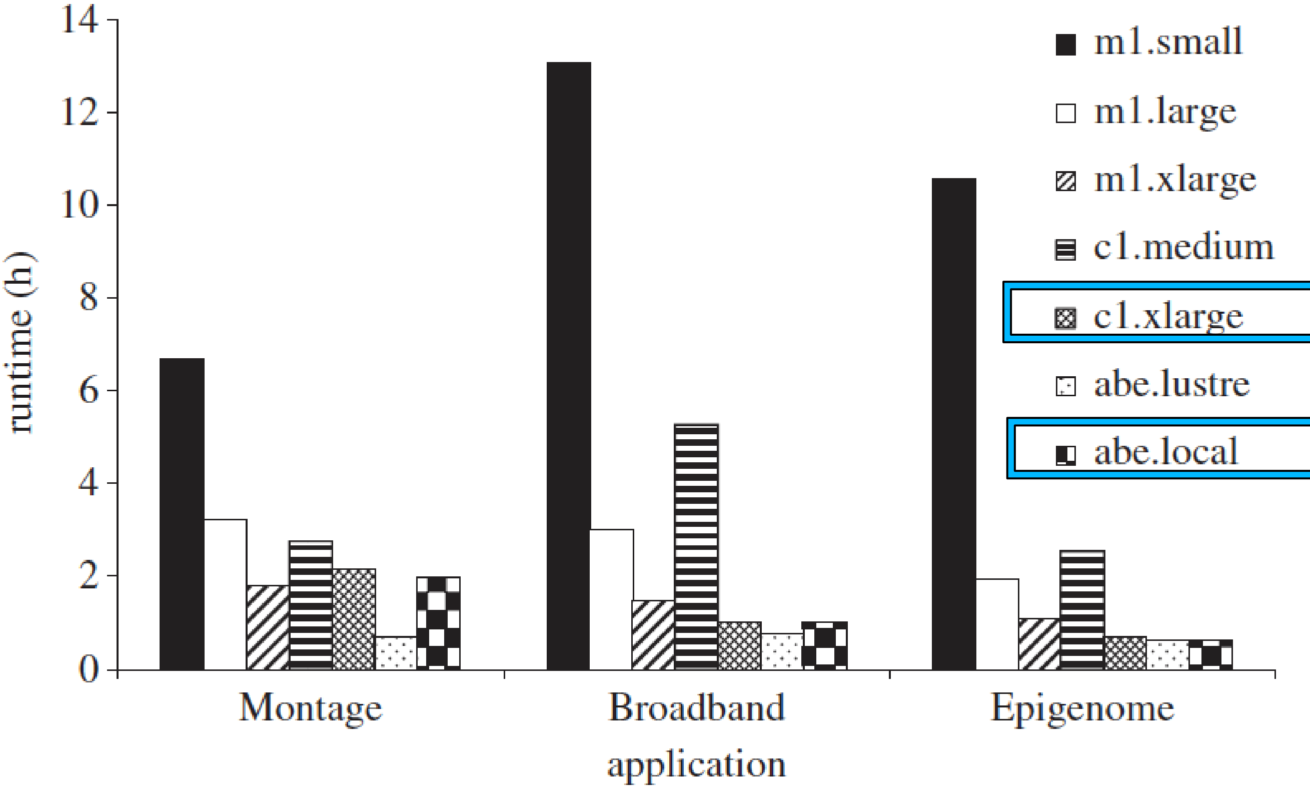

Cloud environments perform similar for CPU- and memory-bound processes.

Berriman et al. 2013

Science Clouds

Cloud computing skepticism (Shane Canon, Lawrence Berkeley National Lab):

- *Overhead to convert to Cloud environments*

Not more heavy than training colleagues on clusters (depends also on application)

Science Clouds

Cloud computing skepticism (Shane Canon, Lawrence Berkeley National Lab):

- *Overhead to convert to Cloud environments*

Not more heavy than training colleagues on clusters (depends also on application)

- *Virtual instances underperform bare-metal systems*

True (especially for heavy i/o), but portability, collaboration in consortium, service, long-term possibility to process data (R. Rohlfs: “s/w must not include h/w specific routines”)

Science Clouds

Cloud computing skepticism (Shane Canon, Lawrence Berkeley National Lab):

- *Overhead to convert to Cloud environments*

Not more heavy than training colleagues on clusters (depends also on application)

- *Virtual instances underperform bare-metal systems*

True (especially for heavy i/o), but portability, collaboration in consortium, service, long-term possibility to process data (R. Rohlfs: “s/w must not include h/w specific routines”)

- *Less cost effective than most large centers*

Might be true when considering commercial clouds (again depends on application). Science cloud: in comparision with clusters, probably less costs for IT

Cloud

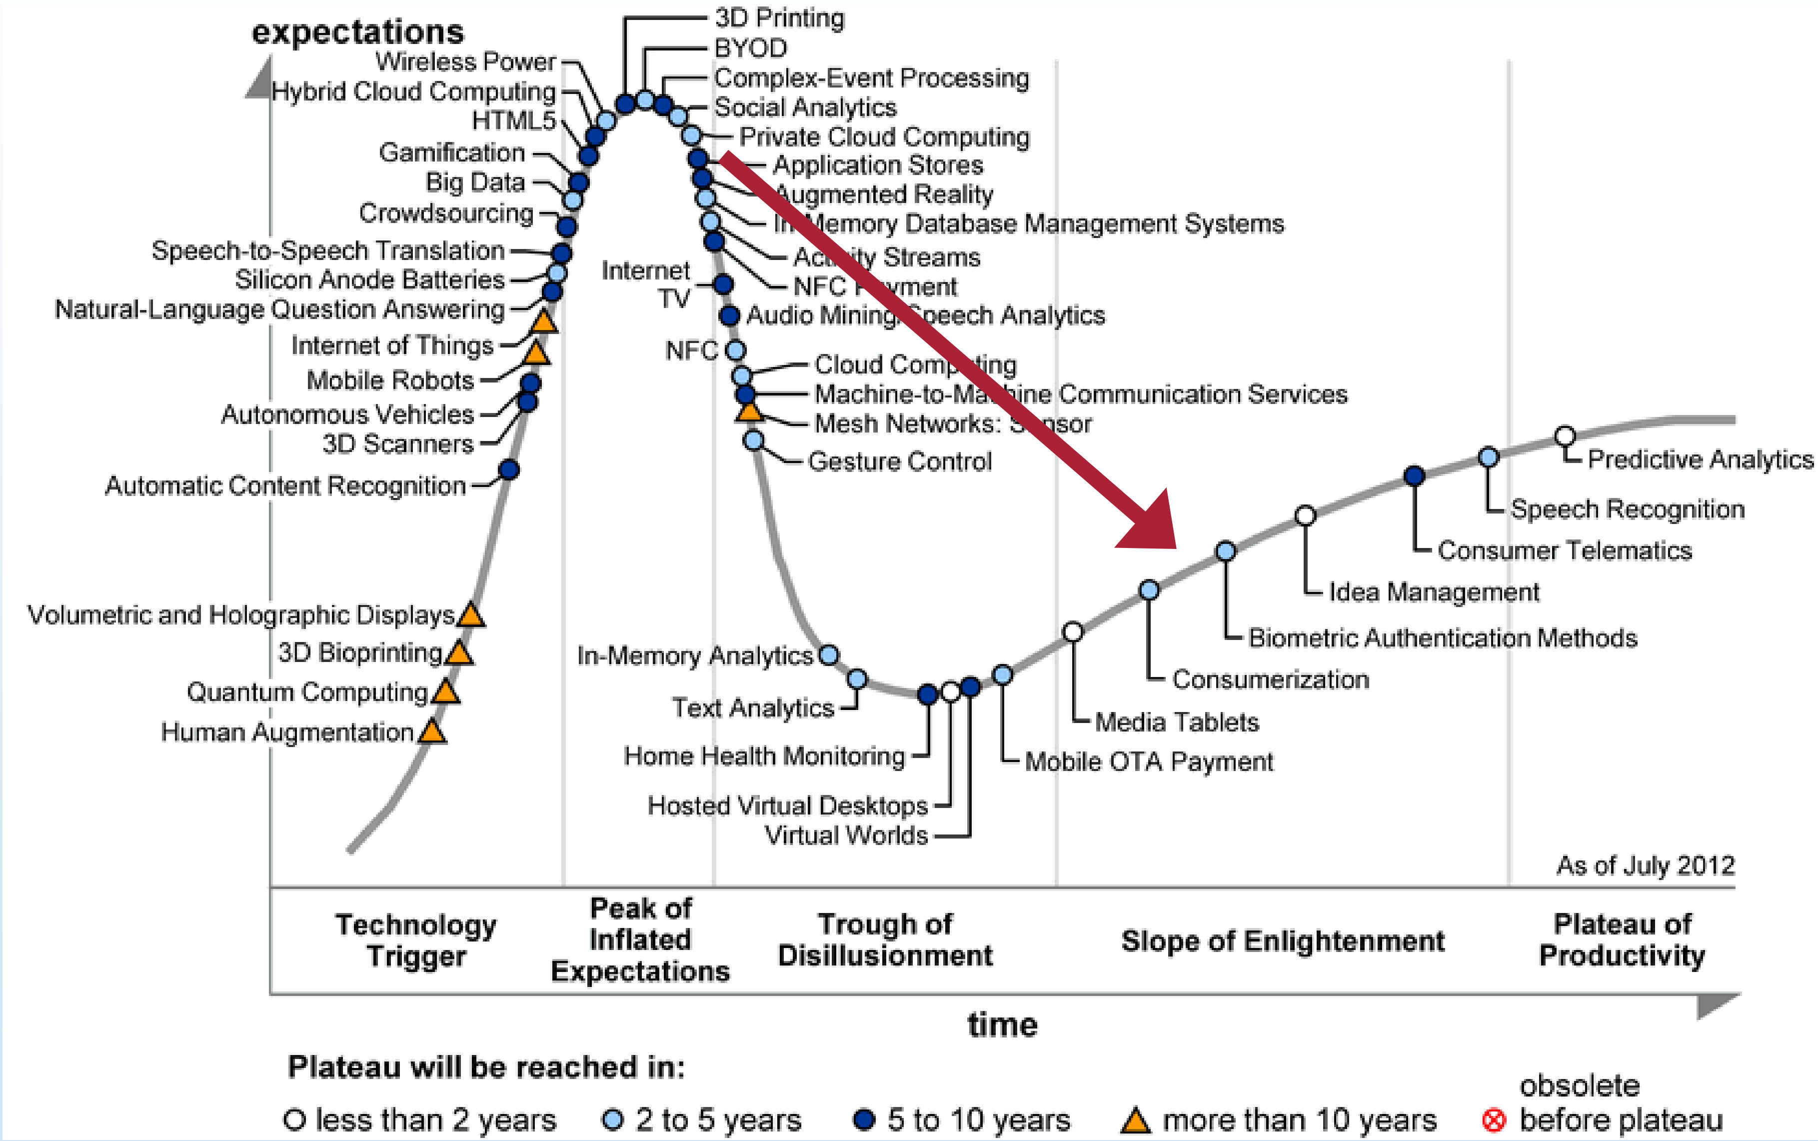

Gartner Inc., Hype Cycle, 2012

Conclusion

Best solution depends on task + politics

GRID approach for heavy + long term + well financed tasks

Cloud environments can be a flexible solution for space projects

But: “The more communication, the worse the performance becomes” (Jackson et al. 2010)

Hybrid cloud solutions appear to satisfy many of the demands of space missions

Commercial cloud for temporary needs only

Read more:

- Berriman et al. 2013, Phil. Trans. R. Soc. A, 371
- Magellan report on Cloud Computing for Science, DoE, 2011
- Jackson et al. 2010, IEEE 2nd International Conf. on Cloud Comp. (Cloud Com)

See also presentation by **Jorgo Bakker** today!

Conclusion

Best solution depends on task + politics

GRID approach for heavy + long term + well financed tasks

Cloud environments can be a flexible solution for space projects

But: “The more communication, the worse the performance becomes” (Jackson et al. 2010)

Hybrid cloud solutions appear to satisfy many of the demands of space missions

Commercial cloud for temporary needs only

Read more:

- Berriman et al. 2013, Phil. Trans. R. Soc. A, 371
- Magellan report on Cloud Computing for Science, DoE, 2011
- Jackson et al. 2010, IEEE 2nd International Conf. on Cloud Comp. (Cloud Com)

See also presentation by **Jorgo Bakker** today!

matériel supplémentaire
additional slides
noch mehr Folien
materiale aggiuntivo
aanvullend materiaal

Cloud vs. Cluster

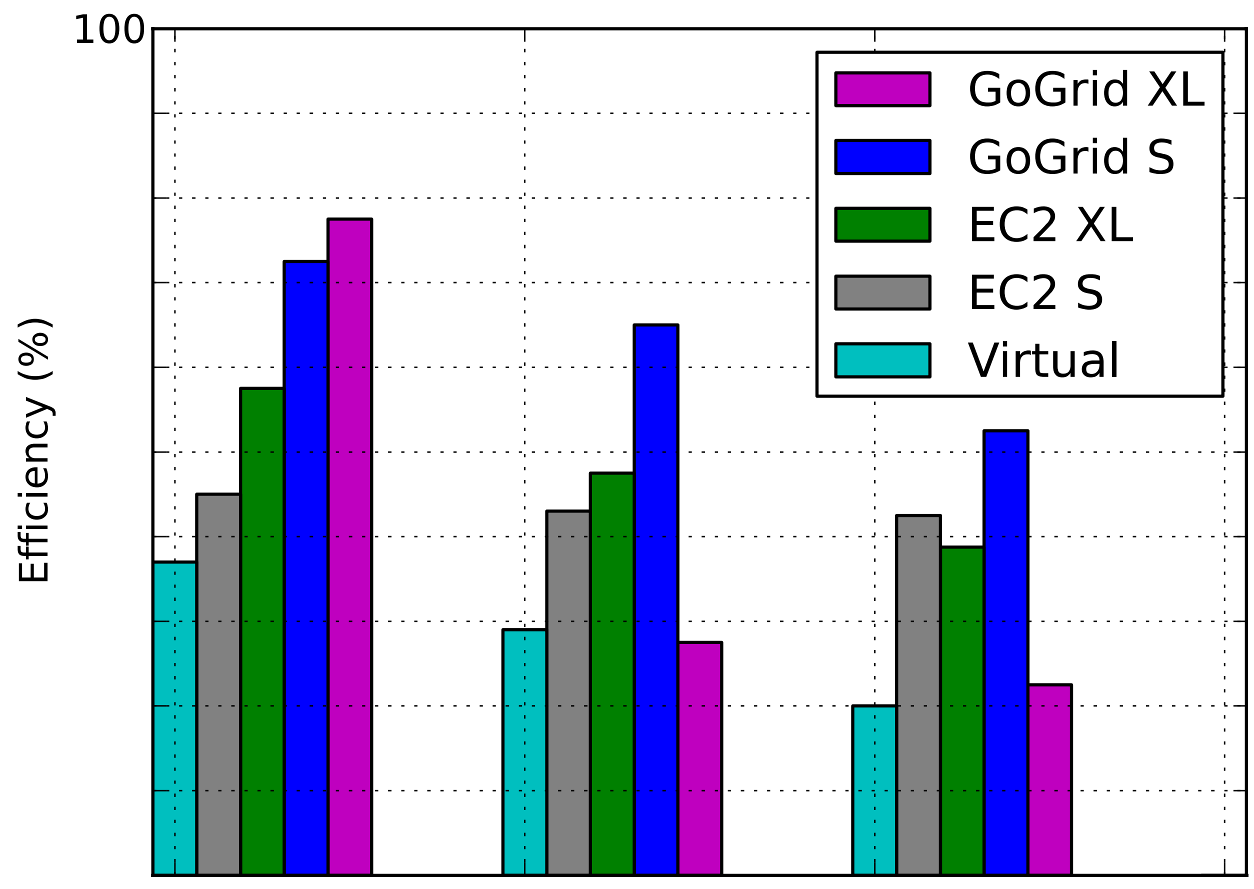

Cloud vs. Cluster

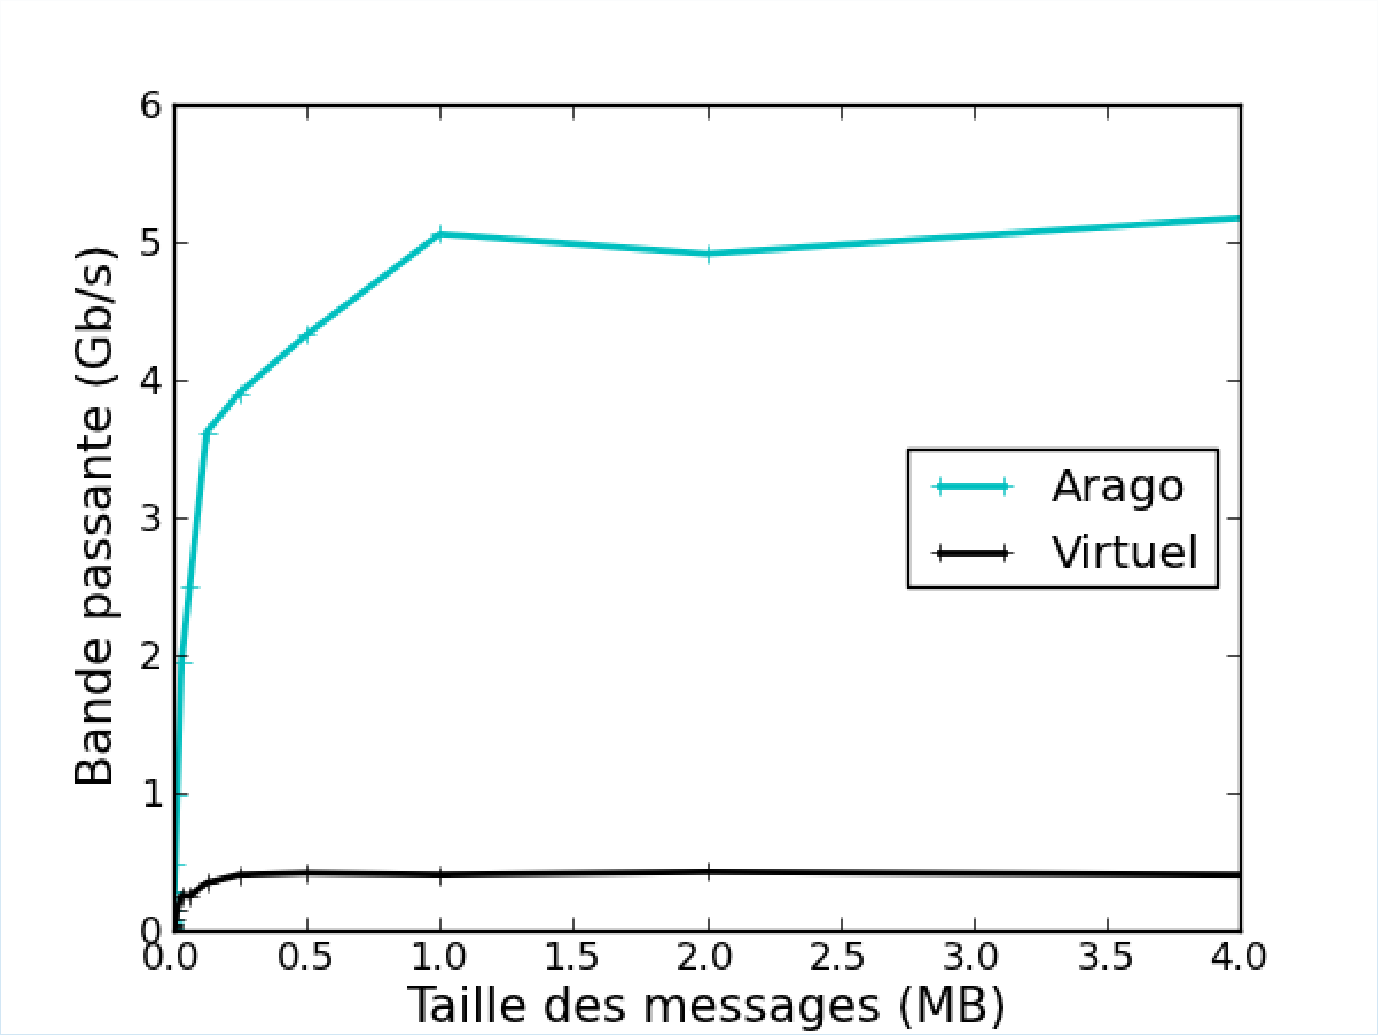

Cloud vs. Cluster

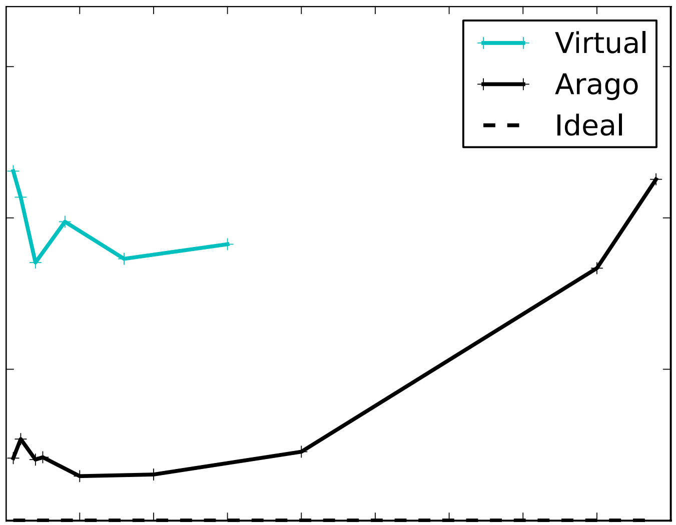

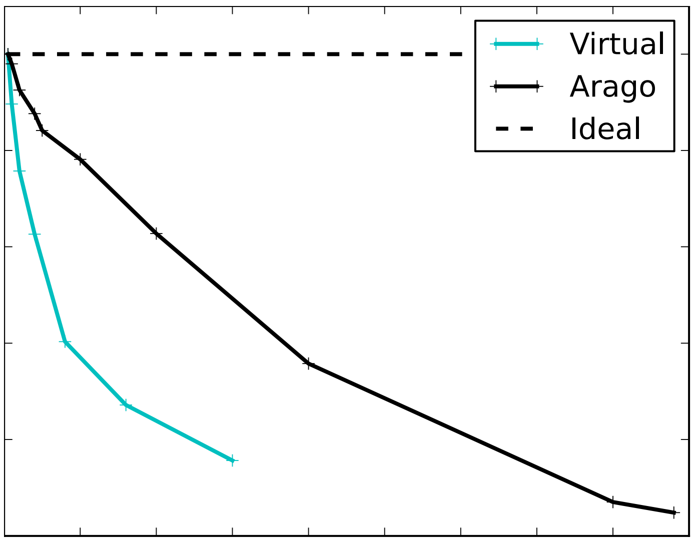

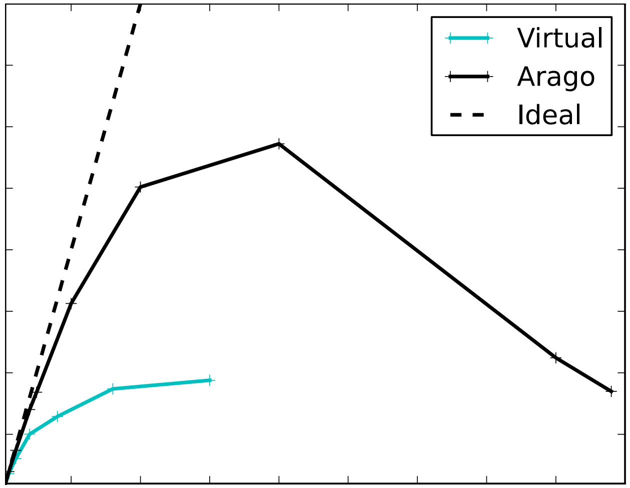

Infrastructure

- Clusters, 620 CPU, 100 kW refroidissement, 100 TByte disque dur
- 10 Gbit/s connection
- 2 salles de conférence vidéo
- 2 salles de réunion
- Bureaux à la demande
- Support logiciel et matérielle
- Concurrent Design Facility