



# CLASS

## Conference 2012

Cloud Assisted Services


## HPC & Clouds: Converging Systems?

**Lutz Schubert**


Head of Infrastructure Research

High Performance Computing Centre Stuttgart

# “Distributed Computing”?



# “Distributed Computing”?



# “Distributed Computing”?

1\*1 GB File

~15 mins

~1.5 mins

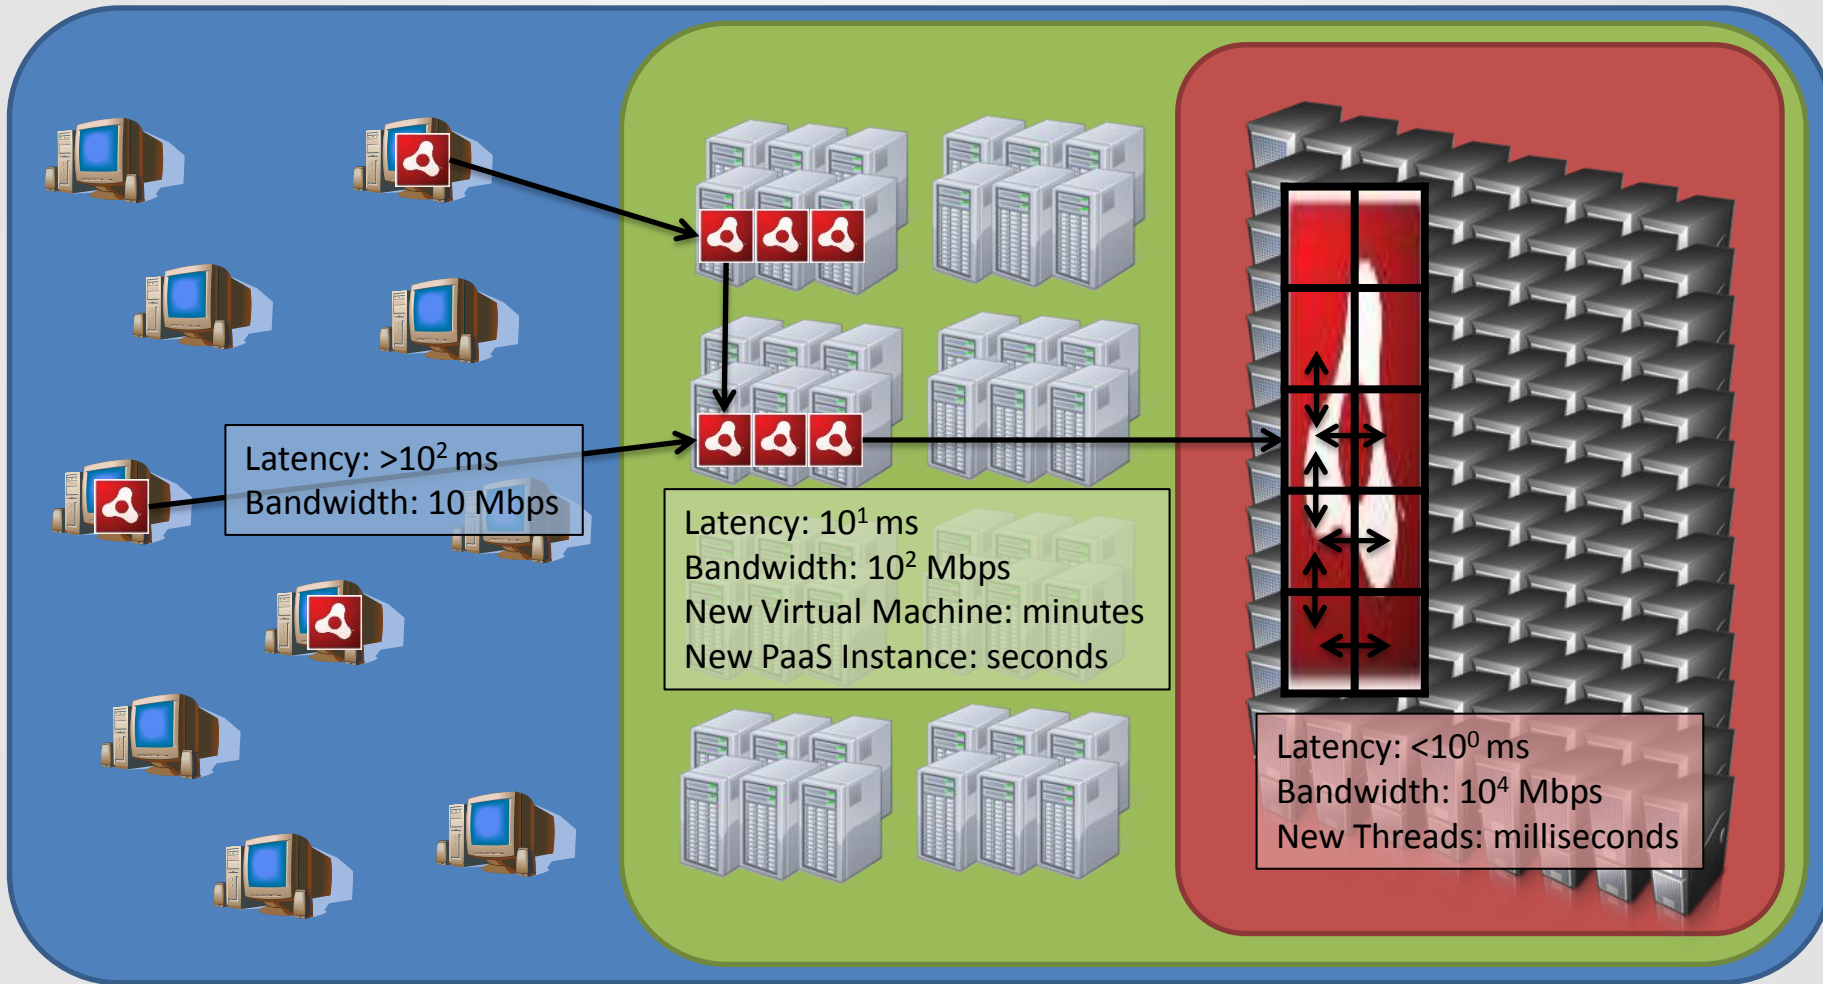
~1 sec

10.000\*100kB Files


>30 mins

~3 mins

~10 secs




# Problems Cloud ↔ HPC



## Cloud

- Want to scale over the amount of (service) instances => improves availability
- Offer desktop PC scale per virtual instance
- “Nodes” are instances that can host the service capabilities
- Replication is easy and comparatively small
  - Full service & state can be replicated
- Access is comparatively easy
- A wide range of platforms is supported

## HPC



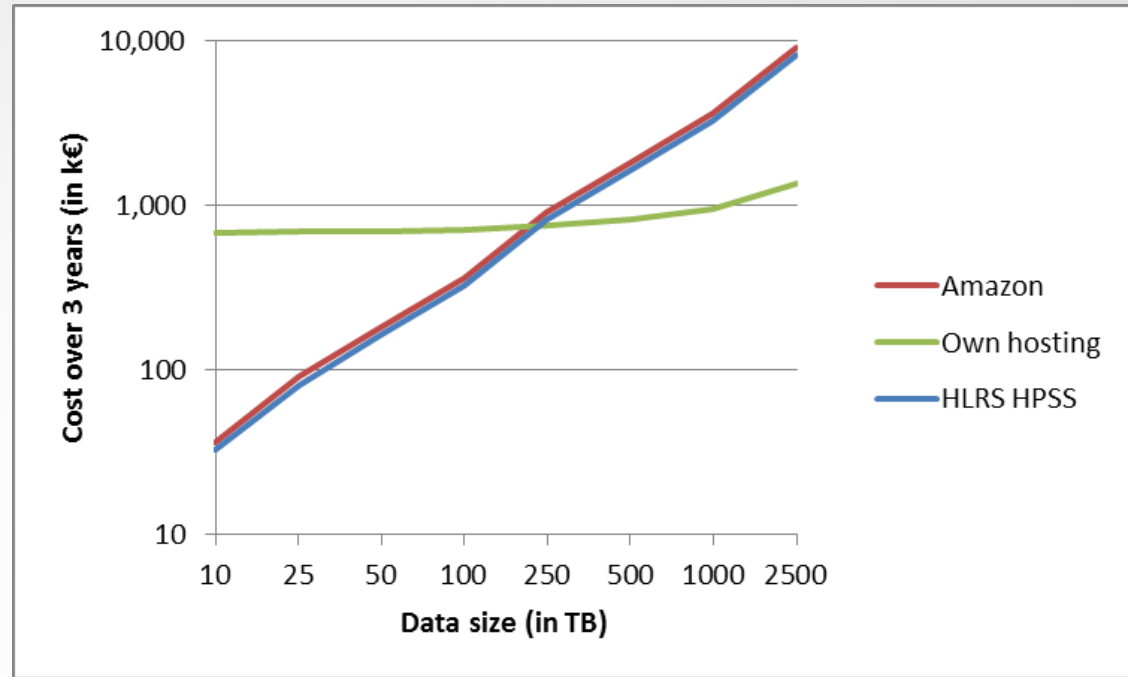
- Scales over parallel “threads” => improves performance, but not availability
- Requires a potentially unlimited number of compute units
- “Nodes” are compute units
- Replication is almost impossible (too many resources are used)
  - Typically only data is replicated
- Access is typically difficult (not only due to security)
- Range of platforms is limited

# HPC is costly?

## Surprise

### Example: Storage

- Simplification:
  - fixed size
  - access 1/1000 of size
  - 3 years runtime
- 500 TB:
  - Amazon S3: ~65000 € / month
  - pure HDD (triple redundancy): ~25000€ / month
  - HLRS HPSS\*: ~45000 € / month




\* more on HPSS: <http://www.hpss-collaboration.org/>

# Why Clouds then?

- HPC: fixed amount of resources reserved
    - little runtime dynamicity, respectively
    - no availability guarantee
  - HPC: Communication cost is fixed
    - does not depend on actual usage
  - HPC: typically fixed duration
    - not necessarily available on the fly
  - HPC: typically not interactive
    - work in progress
- The more dynamic the use case, the less cost by using clouds
- Ideally few peaks

# Why Clouds then?







# Clouds & HPC: no go?

# Types of Scale


Embarrassingly Parallel



Weakly coupled





Tightly coupled




# Types of Scale



Embarrassingly Parallel



Weakly coupled




Tightly coupled



# Example Case(s)

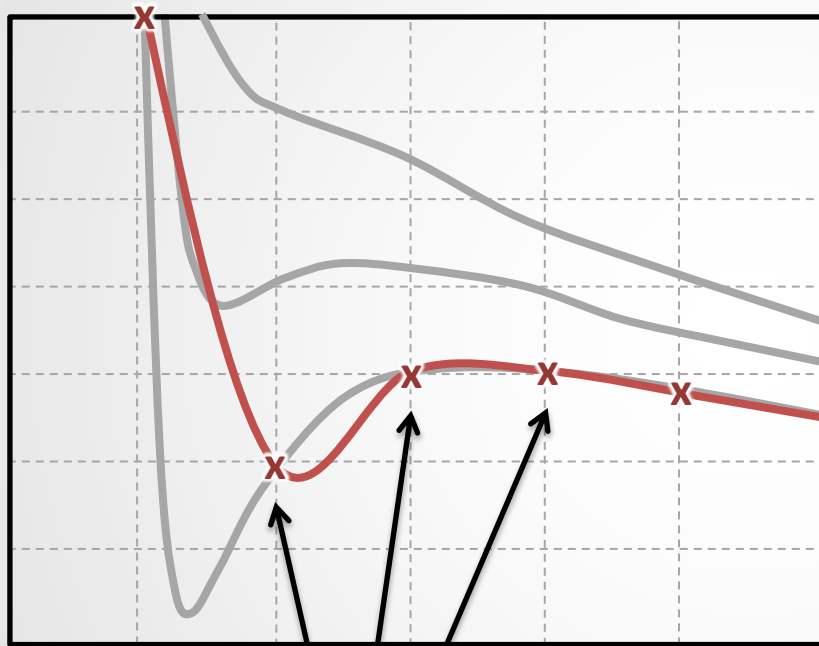
- Mix of Cloud & HPC becomes interesting with mixed modalities of scale
- e.g. eScience applications (here: material stress test)



Parameter Sweep →  
each requires full application with different starting parameters

# Example Case(s)

- Mix of Cloud & HPC becomes interesting with mixed modalities of scale
- e.g. eScience applications (here: material stress test)



Regular Sweep → too few instances ⇒ wrong curve  
too many instances ⇒ waste resources



# CLASS

## Conference 2012

Cloud Assisted Services


## Dynamicity requirements in future cloud-like infrastructures

**Lutz Schubert**


Head of Infrastructure Research

High Performance Computing Centre Stuttgart

# Multiple Users?




# More elaborate approach





# What we need



# Thank You

Lutz Schubert  
HLRS

[schubert@hlrs.de](mailto:schubert@hlrs.de)  
[www.hlrs.de](http://www.hlrs.de)