

Scalability in Grids

Thilo Kielmann

Vrije Universiteit, Amsterdam


kielmann@cs.vu.nl



Scalability

“...is a desirable property of a system, a network or a process, which indicates its ability to either handle growing amounts of work in a graceful manner, or to be readily enlarged.”


Grids



Integrating globally dispersed resources that are not subject to centralized control, to deliver non-trivial qualities of service.

Scalability in Grids

- Network delays (due to physical distance)
bandwidth goes up, but speed of light remains the barrier
- Number of resources involved
 - 1000s of CPUs working together
 - But only $O(10)$ machines/clusters at a time
- Number of data integrated / processed
 - Frillions of bytes in remote files/DBs



Scalability in Grids (2)

The “real” issues:

The application that has worked yesterday won't work today.

- Network disconnections (a.k.a. firewalls)
- Service non-interoperability (a.k.a. versioning)
- And even some hardware failures

Scalability problem: $O(10)$ to $O(100)$ independent administrative authorities...

Scalability: challenges (1)

From the GridLab testbed:

no.	machine	last update	uptime	memory	monitoring client	pathrate receiver	pathrate sender	pathrate client	pathchirp receiver	pathchirp sender
0	sierra0.unile.it	21:24:16	7:3:51.2	7.1 (10.0)	ERROR	IDLE	IDLE	IDLE	IDLE	RUNNING
1	rage1.man.poznan.pl	21:23:47	7:3:50:50	3.0 (4.9)	IDLE	IDLE	IDLE	IDLE	RUNNING	CONFUSED
2	mike4.lsu.edu	21:23:26	7:3:48:31	3.3 (5.9)	ERROR	IDLE	IDLE	IDLE	RUNNING	IDLE
3	bouscat.cs.cf.ac.uk	21:23:25	7:3:48:9	2.2 (4.1)	ERROR	IDLE	IDLE	IDLE	ERROR	RUNNING
4	n0.hpcc.sztaki.hu	21:23:50	7:3:50:55	3.1 (5.9)	IDLE	IDLE	IDLE	IDLE	RUNNING	IDLE
5	hn01.ncsa.uiuc.edu	21:23:42	2:13:28:47	2.4 (4.7)	IDLE	IDLE	IDLE	IDLE	RUNNING	IDLE
6	fs0.das2.cs.vu.nl	21:24:15	7:3:49:42	3.5 (5.9)	IDLE	IDLE	IDLE	IDLE	RUNNING	IDLE
7	skirit.ics.muni.cz	21:23:45	7:3:48:10	12.1 (20.1)	IDLE	IDLE	IDLE	IDLE	RUNNING	IDLE
8	grape.man.poznan.pl	21:23:35	7:3:49:23	3.6 (5.1)	ERROR	IDLE	STARTING	IDLE	RUNNING	STARTING
9	onyx3.zib.de	21:23:31	7:3:49:8	3.3 (7.6)	ERROR	IDLE	IDLE	IDLE	RUNNING	IDLE
10	gridentry.uni-paderborn.de	21:23:39	7:3:50:48	3.4 (5.9)	ERROR	IDLE	IDLE	IDLE	RUNNING	IDLE
11	eltoro.pcz.pl	21:23:35	7:3:47:59	3.4 (5.9)	IDLE	ERROR	IDLE	IDLE	RUNNING	IDLE
12	hitcross.lrz-muenchen.de	21:23:23	7:3:50:1	3.1 (5.9)	IDLE	IDLE	IDLE	IDLE	IDLE	RUNNING
13	litchi.zib.de	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
14	origin.aei-potsdam.mpg.de	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
15	inca.cf.ac.uk	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
16	pclab120.telecom.ece.ntua.gr	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A

Delphoi (Web) service provides network monitoring data

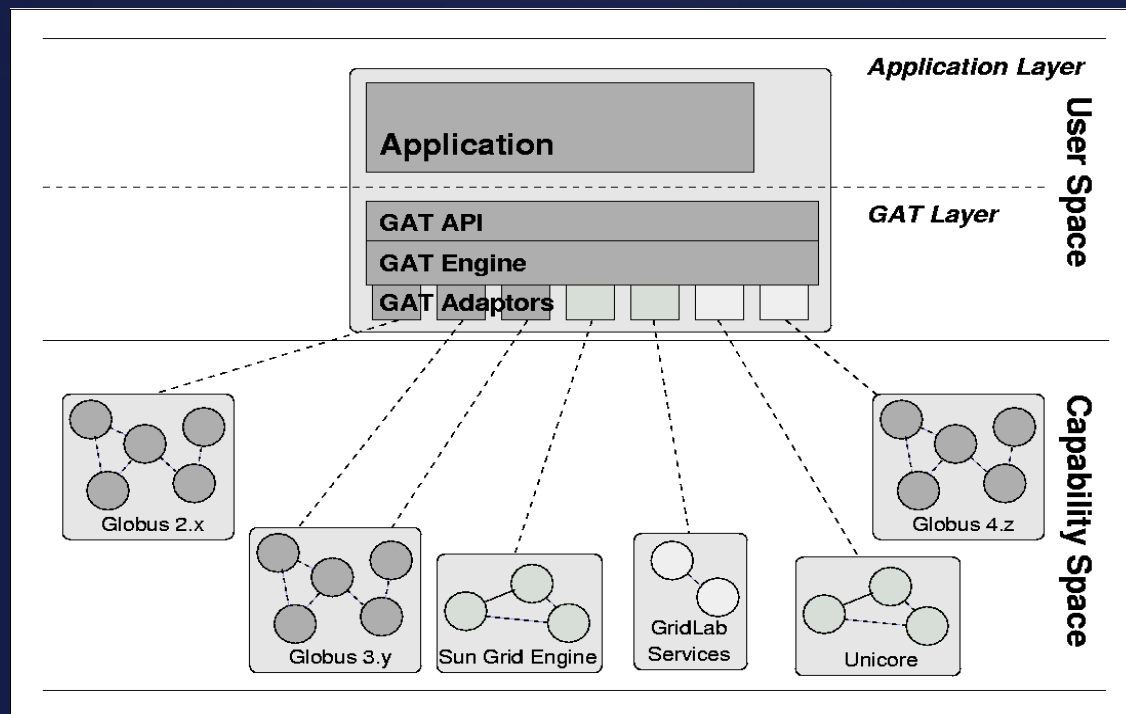
Scalability: challenges (1)

- **Monitoring networks requires $O(N^2)$ measurements**
 - **This becomes practically infeasible, even with $O(10)$ sites.**
- **Challenge:**
build systems that can work without ubiquitous monitoring information

Scalability: challenges (2)

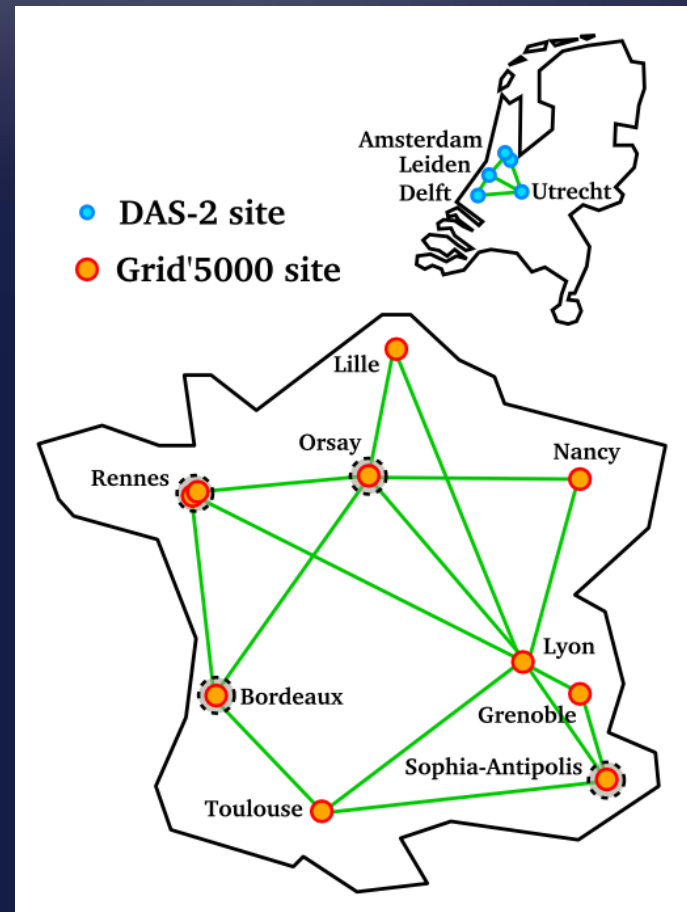
Service non-interoperability

scaling to O(10) middleware/service platforms



Scalability: challenges (3)

- Grid plugtests
- Grids @ Work 2005
- building systems for $O(1000)$ machines



Scalability: challenges (3)

Changes in our Ibis system due to scalability problems:

- **Spread all-to-all connection setup over the runtime**
(false positive denial-of-service, TCP socket limits)
- **Optimize central registry**
(multi threading, message combining)

Not many grid software is built for large-scale use.


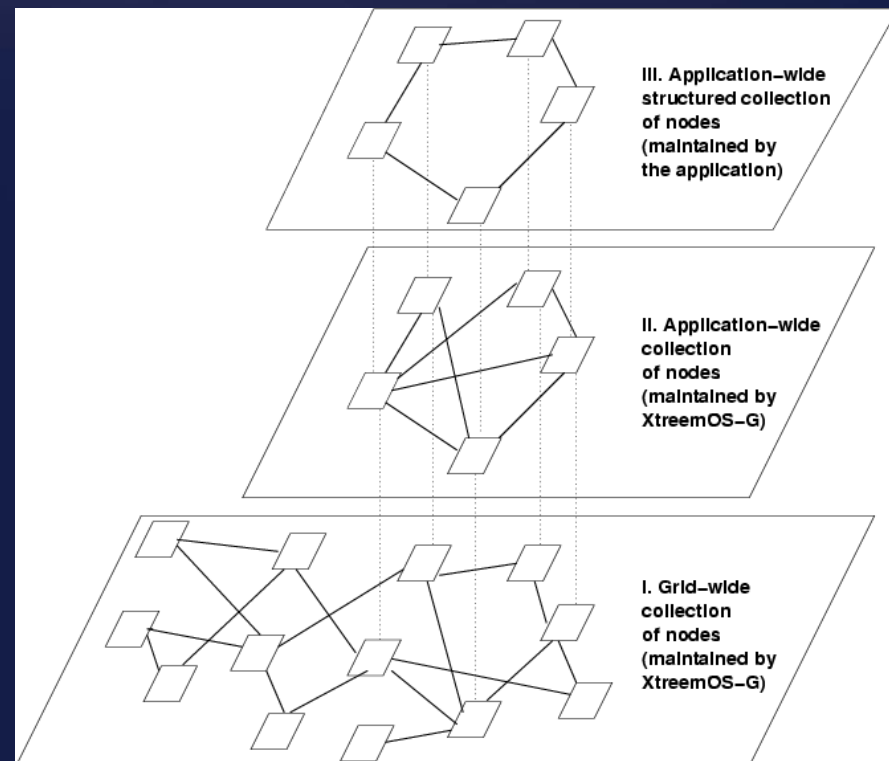
Scalability: challenges (4)

“What can go wrong, will go wrong.” (Murphy)

- Due to the large amount of components (hardware, networks, middleware, applications...) there will **always** be something that is not working (as expected)
- The real challenge is to build systems that are
 - self-*
 - autonomic
 - simply working in a reliable manner...

Scalability: challenges (5)

- single/simple Web services and servers provide critical points of failures
- XtremOS is building virtual servers with hand-over based on peer-to-peer technology



Scalability: challenges (6)

- **Security, AAA**
 - Authentication and authorization for users and services
 - Should be a trivial (straight forward) problem
 - ... if it would not be for scaling to **MANY** users
- **Manual granting of credentials, non-technical but human (or legal) issues**
- **Challenge: automated security mechanisms that are powerful, flexible, and trustworthy**

Conclusions

- Many scalability problems can be circumvented by “doing our homework” (write solid, well-designed software)
- The real challenge is to build systems that are
 - self-*
 - autonomic
 - simply working in a reliable manner...as this is the key to addressing the scalability issues!