


Grid Computing:

Status and Perspectives



Alexander Reinefeld
Florian Schintke



„Schwerpunkte der Informatik“
Ringvorlesung am 05.06.2003

1

Outline

- **MOTIVATION**
 - What's a Grid? Why using Grids?
- **TWO TYPICAL APPLICATION DOMAINS**
 - Computation Fluid Dynamics: FlowGrid
 - Particle Physics: DataGrid
- **WHAT'S MISSING IN CURRENT GRID SYSTEMS**
 - Fault tolerance
 - Scalability
 - Distributed data management
 - Advance reservation, co-scheduling
- **CONCLUSION**


2

What's a Grid?

"A computational grid is a hardware and software infrastructure that provides dependable, consistent, pervasive, and inexpensive access to high-end computational capabilities."

Ian Foster, Carl Kesselman (1999)

*strongly influenced by
early metacomputing
projects!*



3

What's a Grid? (2nd Attempt)

Grid computing is concerned with *"coordinated resource sharing and problem solving in dynamic, multi-institutional virtual organizations."*

I. Foster, C. Kesselman, S. Tuecke, „Anatomy of the Grid“ 2000

-> Stresses importance of **protocols** as a means of enabling **interoperability**

*now also addresses
social and
policy issues*

4

What's a Grid? (3rd Attempt)

Ian Foster's Three-Point Checklist in hpcwire (22.07.2002):

„A Grid is a system that

- coordinates resources that are not subject to centralized control
- uses standard, open, general purpose protocols and interfaces
- delivers nontrivial quality of service“

„Ian's checklist may be more **valid for some very large research Grids**, in the future. But this checklist would exclude many contributions from industry, thus likely pushing Grid computing into a niche.“

W. Gentsch, hpcwire 05.08.2002

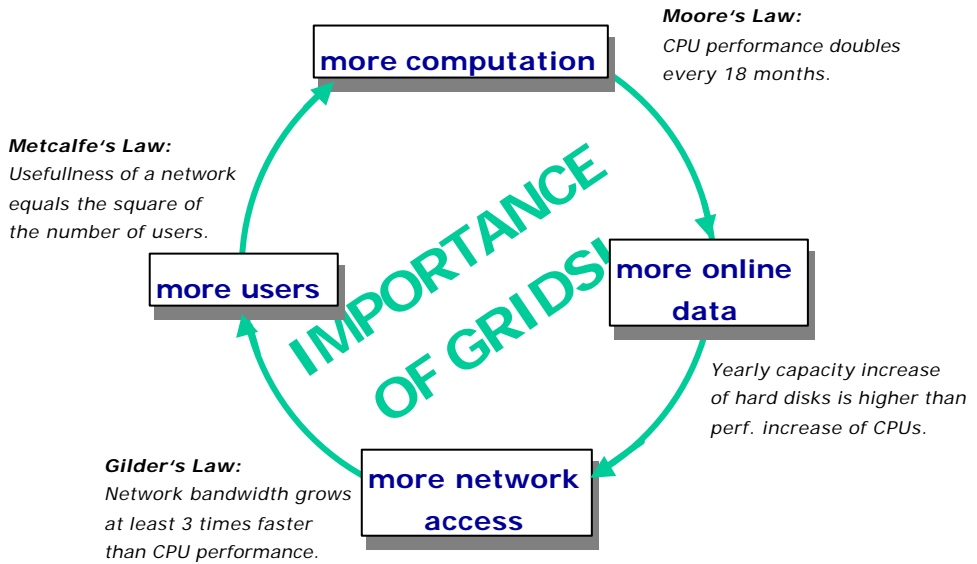
5

What's a Grid? - Our Definition

- Grids are **large**
 - in terms of potentially available resources and distance between them
- Grids are **distributed**
 - substantial latencies in moving data, may dominate the application
- Grids are **dynamic**
 - resources may change during the lifespan of an application
- Grids are **heterogeneous**
 - form and properties of sites (nodes) may differ significantly
- Grids are **across boundaries of organizations**
 - access policies differ at different sites

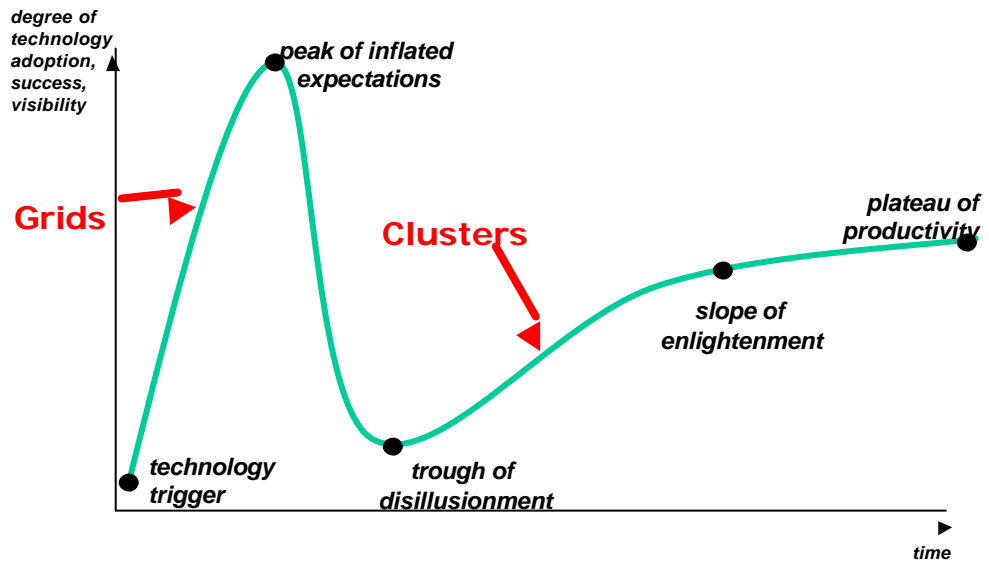
6

Why Using Grids?



7

Emerging Technology Hype Curve



8

3 Approaches to Building Grid Environments

- a collection of tools bundled in a „toolkit“ (like Globus 2.0)
 - o provides access and protocols to use distributed computing resources efficient
 - o applications are loosely coupled
 - o **emphasis on openness**



- a distributed operating system (like Legion)
 - o supports OO distributed programming
 - o applications are tightly coupled
 - o **emphasis on coherence**



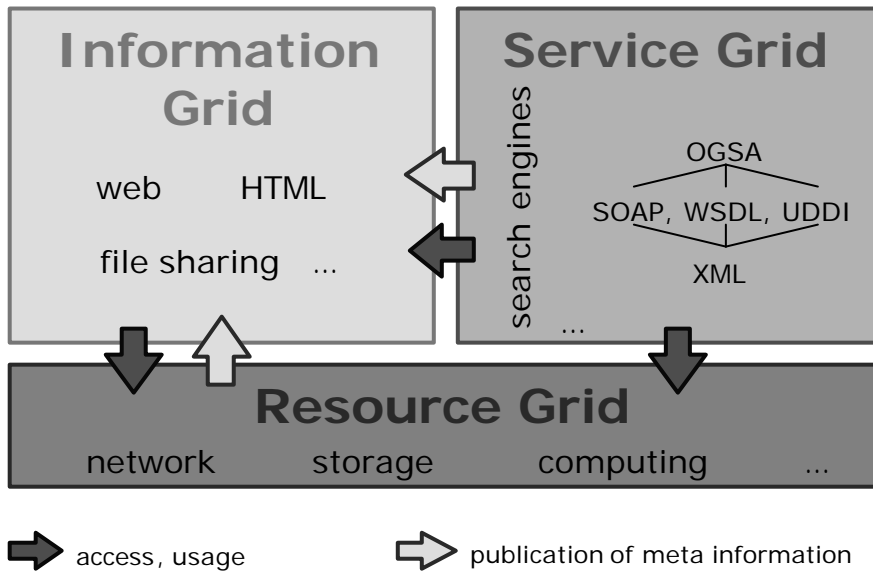
- a distributed resource management framework (like Condor)
 - o harvesting resources for high-throughput computing
 - o **emphasis on work-load balancing**



Many Projects - What did we learn?



What's a Grid? Our View




11

Example of a Resource Grid: Globus 2.0 Toolkit™

GLOBUS' PROTOCOLS

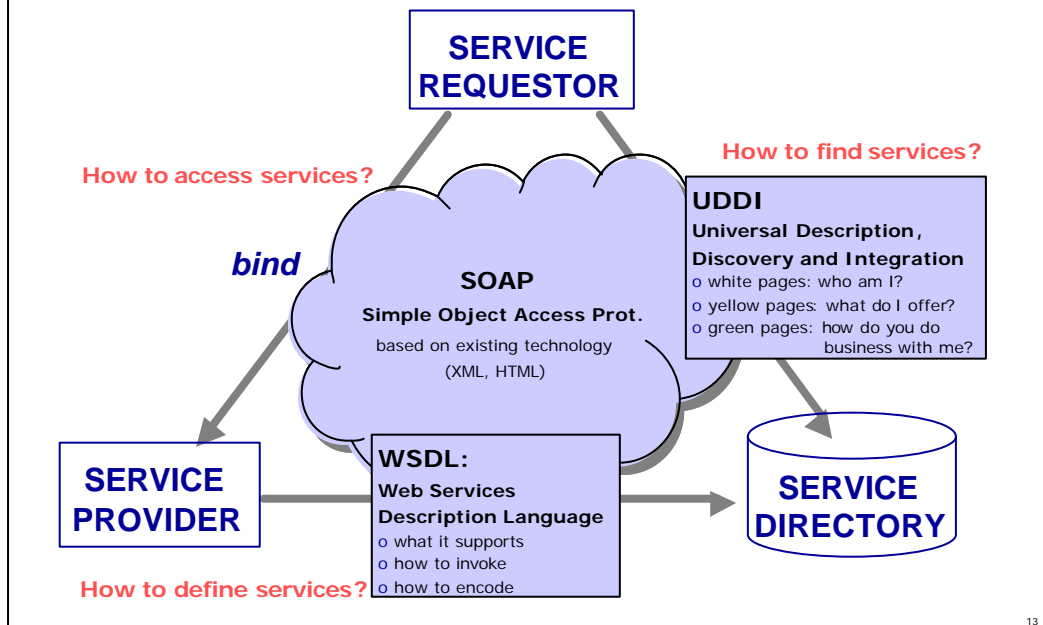
- **Connectivity layer**
 - **GSI:** Grid Security Infrastructure
- **Resource layer**
 - **GRIP:** Grid Resource Information Protocol
 - **GRAM:** Grid Resource Allocation Management
 - **GridFTP:** Grid File Transfer Protocol
- **Collective layer protocols**
 - Info Services, Replica Management, etc.

Hourglass Model



12

Towards a Service Grid: Web Services



WEB Services versus GRID Services

- **WEB Services** address discovery & invocation of **persistent services**
 - Interface to persistent state of entire enterprise
- **GRIDs** must also support **transient service instances** that are dynamically created/destroyed
 - need interfaces to the states of distributed activities

Significant implications for how services are managed, named, discovered, and used

OGSA
Open Grid Service Architecture

OGSA - Framework for Grid Services

- Each OGSA-compliant service should support standard interfaces for:
 - **creation of service instances** (Factory)
 - **lifetime management** (Get-/SetTerminationTime, explicit destruction)
 - **registration & discovery** (Query, extensible query language)
 - **authorization** (remote access control policy management)
 - **notification** (observe service existence, lifetime, information, ...)

SOAP, WSDL, and UDDI **can** build the base for an OGSA implementation.

Globus3.0 (GT3) contains OGSA compliant services for all **(most?)** GT2 services

15

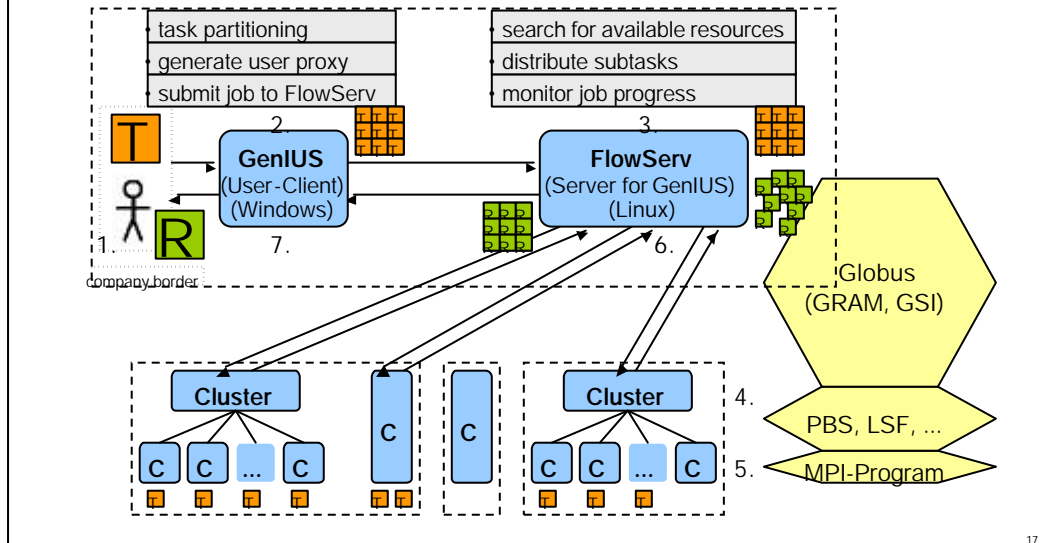
Outline

- **MOTIVATION**
 - What's a Grid? Why using Grids?
- **TWO TYPICAL APPLICATION DOMAINS**
 - Computation Fluid Dynamics: FlowGrid
 - Particle Physics: DataGrid
- **WHAT'S MISSING IN CURRENT GRID SYSTEMS**
 - Fault tolerance
 - Scalability
 - Distributed data management
 - Advance reservation, co-scheduling
- **CONCLUSION**

16

Example 1: FlowGrid

A typical run, starting with the specification of a CFD task *T* and ending with the presentation of the results *R* to the user



17

Resource Description in FlowGrid


```

<System SystemSite="Symban">
  <ResourceManager>
    <ManagerName>zeus.flowgrid.com</ManagerName>
    <CPUInfo>
      <CPUSpeed>500</CPUSpeed>
      <CPUCount>2</CPUCount>
      <CPUAvailable>1</CPUAvailable>
    </CPUInfo>
    <MachineMemory >2000</MachineMemory >
    <MachineDiskSpace>40000</MachineDiskSpace>
    <MachineNetSpeed>100</MachineNetSpeed>
    < InstalledSoftware >genius.exe</InstalledSoftware >
    < InstalledLicense>genius.lic</InstalledLicense>
  </ResourceManager>
  <ResourceManager>
    <ManagerName>hercules.flowgrid.com</ManagerName>
    <CPUInfo>
      <CPUSpeed>2400</CPUSpeed>
      <CPUCount>1</CPUCount>
      <CPUAvailable>1</CPUAvailable>
    </CPUInfo>
    <MachineMemory >1000</MachineMemory >
    <MachineDiskSpace>60000</MachineDiskSpace>
    <MachineNetSpeed>100</MachineNetSpeed>
    < InstalledSoftware >genius.exe</InstalledSoftware >
    < InstalledLicense>genius.lic</InstalledLicense>
  </ResourceManager>
</System>
  
```

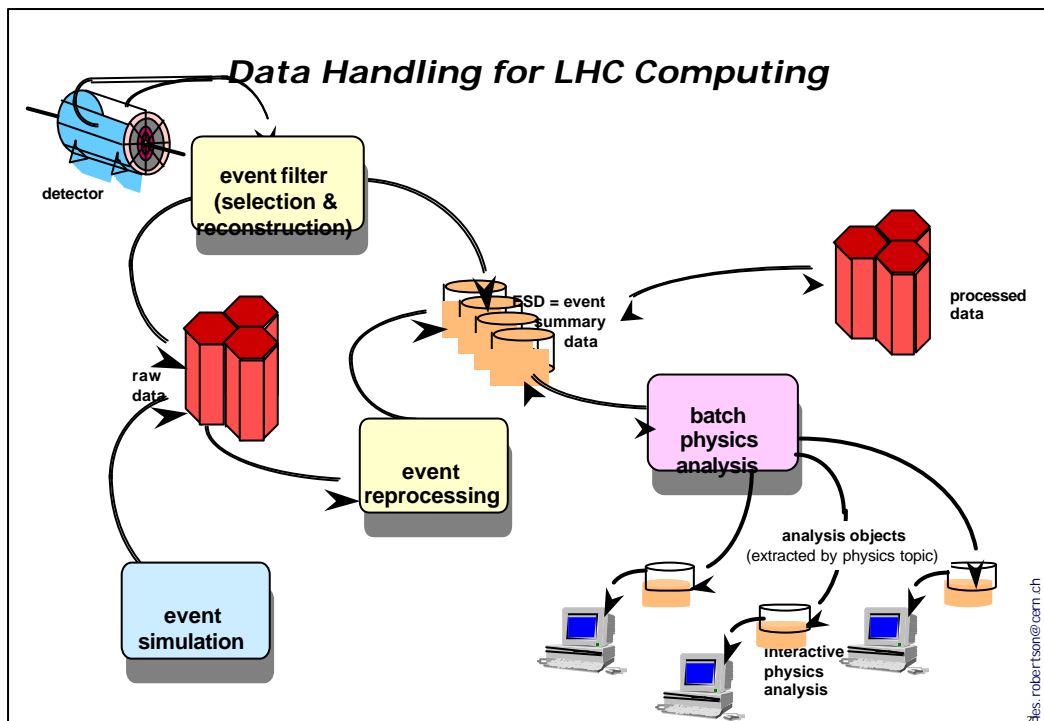
18

Example 2:

The 4 detectors of the *LHC* at Cern will produce $\sim 10^9$ events per year = 4 PB



19



Data Access Pattern

Cern worldwide has

- o 267 institutes in Europe, 4603 users
- o 208 institutes elsewhere, 1632 users
- o all will be accessing the LHC Data Grid!



Layered Architecture

- 1 Tier-0 center (Cern) with $O(10^4)$ PCs and $O(10^4)$ TByte storage
- ~ 10 Tier-1 centers with $O(10^3)$ PCs and $O(10^3)$ TByte
- > 10 Tier-2 centers with $O(10^2)$ PCs
- > 100 Tier-3 centers with some PCs or SMPs

Data stored distributedly

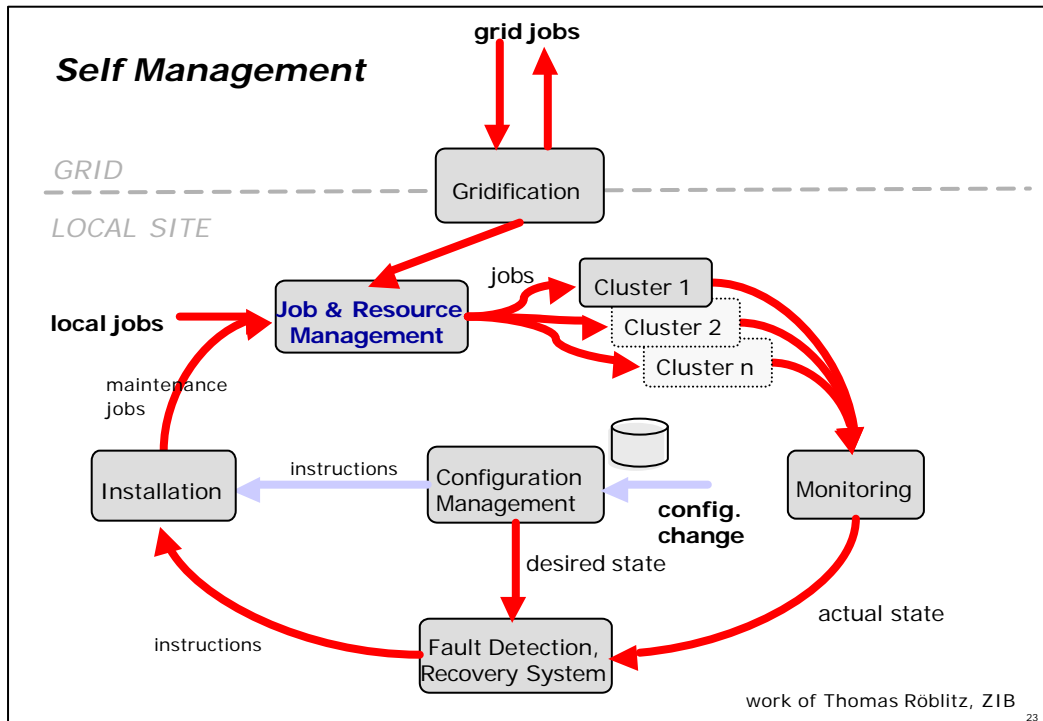
21

Large PC Farms: Cost Effective? Reliable?

assuming PCs with one 500 GB disk, GbE, 250 Watt, 3000 €

nodes	Tier 1: 3,000	Tier 0: 10,000	all LHC: 50,000
daily faults (avg. PC lifetime is 5 years)	1.6 / d	5.5 / d	27 / d
summed disk storage	1.5 PB	5 PB	25 PB
disk read errors (every 8 y = 70.000 h)	1.0 / d	3.4 / d	17 / d
one bit error per transferred TBit (GbE)	0.2 Mio / d	0.8 Mio / d	4.3 Mio / d
power consumption	750 kWatt	2.5 MWatt	12 Mwatt
power cost (10 cent/kWh)	650 k€/y	2.2 M€/y	11 M€/y
monthly hardware reinvest	144 k€/m	500 k€/m	2.4 M€/m

22



- ## Outline
- **MOTIVATION**
 - What's a Grid? Why using Grids?
 - **TWO TYPICAL APPLICATION DOMAINS**
 - Computation Fluid Dynamics: FlowGrid
 - Particle Physics: DataGrid
 - **WHAT'S MISSING IN CURRENT GRID SYSTEMS**
 - Fault tolerance
 - Scalability
 - Distributed data management
 - Advance reservation, co-scheduling
 - **CONCLUSION**
- 24

Scalability

- MDS – the Grid metadata catalog uses hierarchical LDAP servers
→ not scalable for large Grids!

Latency of „remote“ DoNothing() call on localhost


System	Language	Latency (ms)
JavaRMI	Java	1.2
CORBA	Java	1.5
MS SOAP Toolkit	Visual Basic	16.8
SoapRMI	Java	19.5
SOAP::Lite	Perl	42.0
Apache SOAP	Java	23.4
Apache Axis	Java	15.6

- OGSA based on SOAP calls
→ too slow for HPC environment and massive accesses.

25

How to Achieve Scalability?

A P2P System = A Component



P2P systems are simple.

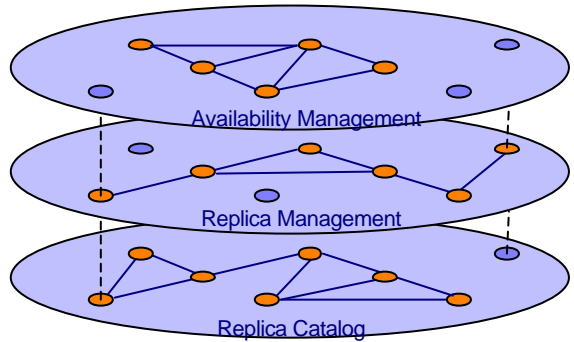
Together, they may perform complex tasks.

26

Combining P2P Systems

Components

- have their **own overlay network**
- are **loosely coupled** or **independent**



FUNCTIONAL COMPONENTS provide user functionality

SELF MANAGEMENT COMPONENTS monitor and control the grid

27

Distributed Data Management

- **replication** to improve reliability by introducing redundant file copies,
- **synchronization** for keeping replicas in a consistent state,
- **placing** to determine optimal replica placement for fast access,
- **caching** and **pre-fetching** to benefit from spatial and temporal locality,
- **staging** to improve job execution time by scheduled data transfers,
- **locating** nearest replica for fast processing.

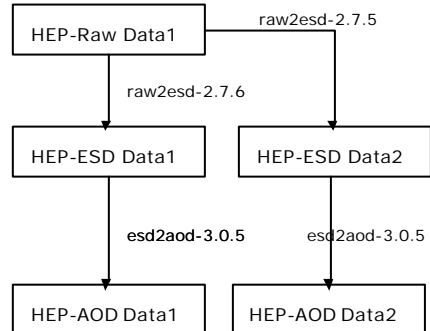
28

Metadata Management and Co-Scheduling

- **Keep track of replicas**
 - maintain file dependency graph
 - co-locate all dependent/derived files
- **Co-schedule data and computation**

e.g. by modifying PBS to ...

 - schedule jobs to data
 - replicate data to jobs
 - predict future data usage
 - prefetch data



29

Outline


- **MOTIVATION**
 - What's a Grid? Why using Grids?
- **TWO TYPICAL APPLICATION DOMAINS**
 - Computation Fluid Dynamics: FlowGrid
 - Particle Physics: DataGrid
- **WHAT'S MISSING IN CURRENT GRID SYSTEMS**
 - Fault tolerance
 - Scalability
 - Distributed data management
 - Advance reservation, co-scheduling
- **CONCLUSION**

30

Conclusion

- **No future without Grid technology!**

- more computation, more distribution, more collaboration




HU-Antrag Graduiertenkolleg

- **No future without autonomic computing!**

- Build self-regulating systems.

- **“Data is the Grid’s Killer App!”**

(Fran Berman, director SDSC and NPACI, 4/2002 at ICCS)



HU-Schwerpunkt

„Große Datenmengen im Web“