

# First steps of a monitoring framework to empower risk assessment on Grids

---

M. Hovestadt, N. Lerch, [H. Nitsche](#), K. Voß  
hn@uni-paderborn.de

Cracow Grid Workshop '06, 17th October 2006



PADERBORN  
CENTER FOR  
PARALLEL  
COMPUTING

# Overview

- ▶ Motivation
- ▶ Short History
- ▶ AssessGrid
- ▶ UMF Design
- ▶ UMF failover
- ▶ Conclusions
- ▶ Contact



Why another monitoring tool?


What's wrong with R-GMA, PCP, Ganglia, Nagios ... ?

**The simple answer is: Nothing!**

We wanted/needed special features:

- ▶ Runtime configuration of probes
- ▶ Flexibility to schedule probes
- ▶ Data coherence / integrity / security
- ▶ (Almost) No loose of probe data

# Short History

Q1 2005	Start development of Nagios Add-On <b>NSCE</b> (Enhanced) <ul style="list-style-type: none"><li>• Inspired by NSCA Nagios Service Check Acceptor</li><li>• Ease of configuration</li><li>• Reduce possible data losses</li></ul>
Q3 2005	Test NSCE on Clusters
Q1 2006	Extend NSCE to use X.509 certificates for use in D-GRID monitoring tests
Q3 2006	Start development of Unified Monitoring Framework ( <b>UMF</b> ) <ul style="list-style-type: none"><li>• From scratch, based on ideas from NSCE</li><li>• Flexible probe schedule on client</li></ul>
Q4 2006	Start integration of UMF into 

# AssessGrid Objectives

## Advanced Risk Assessment and Management for Trustable Grids

- Objective indicators about the quality of the own infrastructure.
- Risk estimation for different situations (low/high loaded resources, vacation time, overloaded network, etc.) which helps to decide on incoming SLAs and to set a penalty fee corresponding to the risk of failure.
- Decision-support for system development, management, and planning: Grid computing needs a complex infrastructure. Bottlenecks are difficult to detect and not removed by simple investment in oversized hardware. Aggregated risk indicators will show, which parts of the infrastructure increase the risk and should be improved.
- Self-organising fault tolerance mechanisms use certain risk indicators as thresholds to increase the reliability. In case of failures and risk above the threshold, the business policy will be adapted. For example, longer slack-times will be negotiated, the penalty fee will be reduced or even SLAs will be rejected. On the other hand, spare resources will be activated or a redundant processing will be activated.



European Commission  
6<sup>th</sup> Framework Programme  
Contract IST- 031772



# Using Risk Assessment

- ▶ compute risk for SLA violation **before** offering an SLA
  - ▶ make a reservation
- ▶ publish the risk in SLA
- ▶ determine charge and penalty fee
- ▶ monitor risk of SLA violation during job execution
- ▶ initiate precautionary FT-actions if risk increases too much
  - ▶ checkpointing and migration: select a suitable resource for migration
- ▶ **these actions are risk management!**
  - ▶ included in RMS modules

# Unified Monitoring Framework features


The key for risk assessment on grid is monitoring!

To fulfil this key role, UMF has built some interesting features:

- ▶ Central configuration management for all clients
- ▶ Deployment of client is done by simply copying one directory
- ▶ The clients probes can be configured at runtime
- ▶ Flexible probe intervals for each probe
- ▶ All existing Nagios plugins can be used (some hundred)
- ▶ Possibility of data loose is minimised by client local backup data
- ▶ Build-in failover for UMC /UMS communication
- ▶ Load balancing for communication
- ▶ Scalable storage solution (DB based)
- ▶ Different priorities for OK, Warn, Critical, Error
- ▶ Depending on priority results can be forwarded immediately

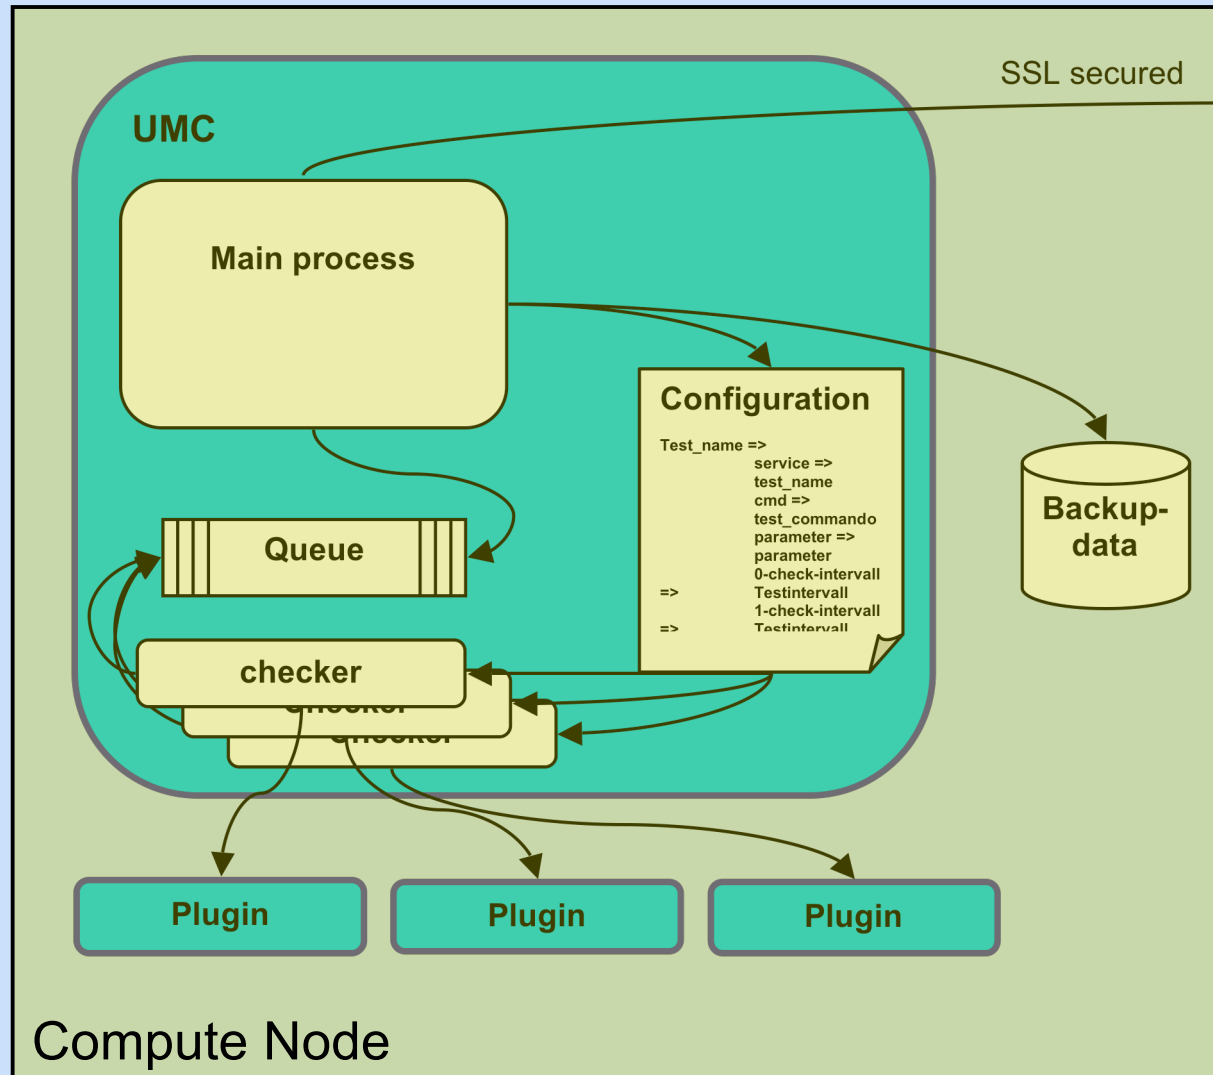
# UMF System architecture

Compute Nodes



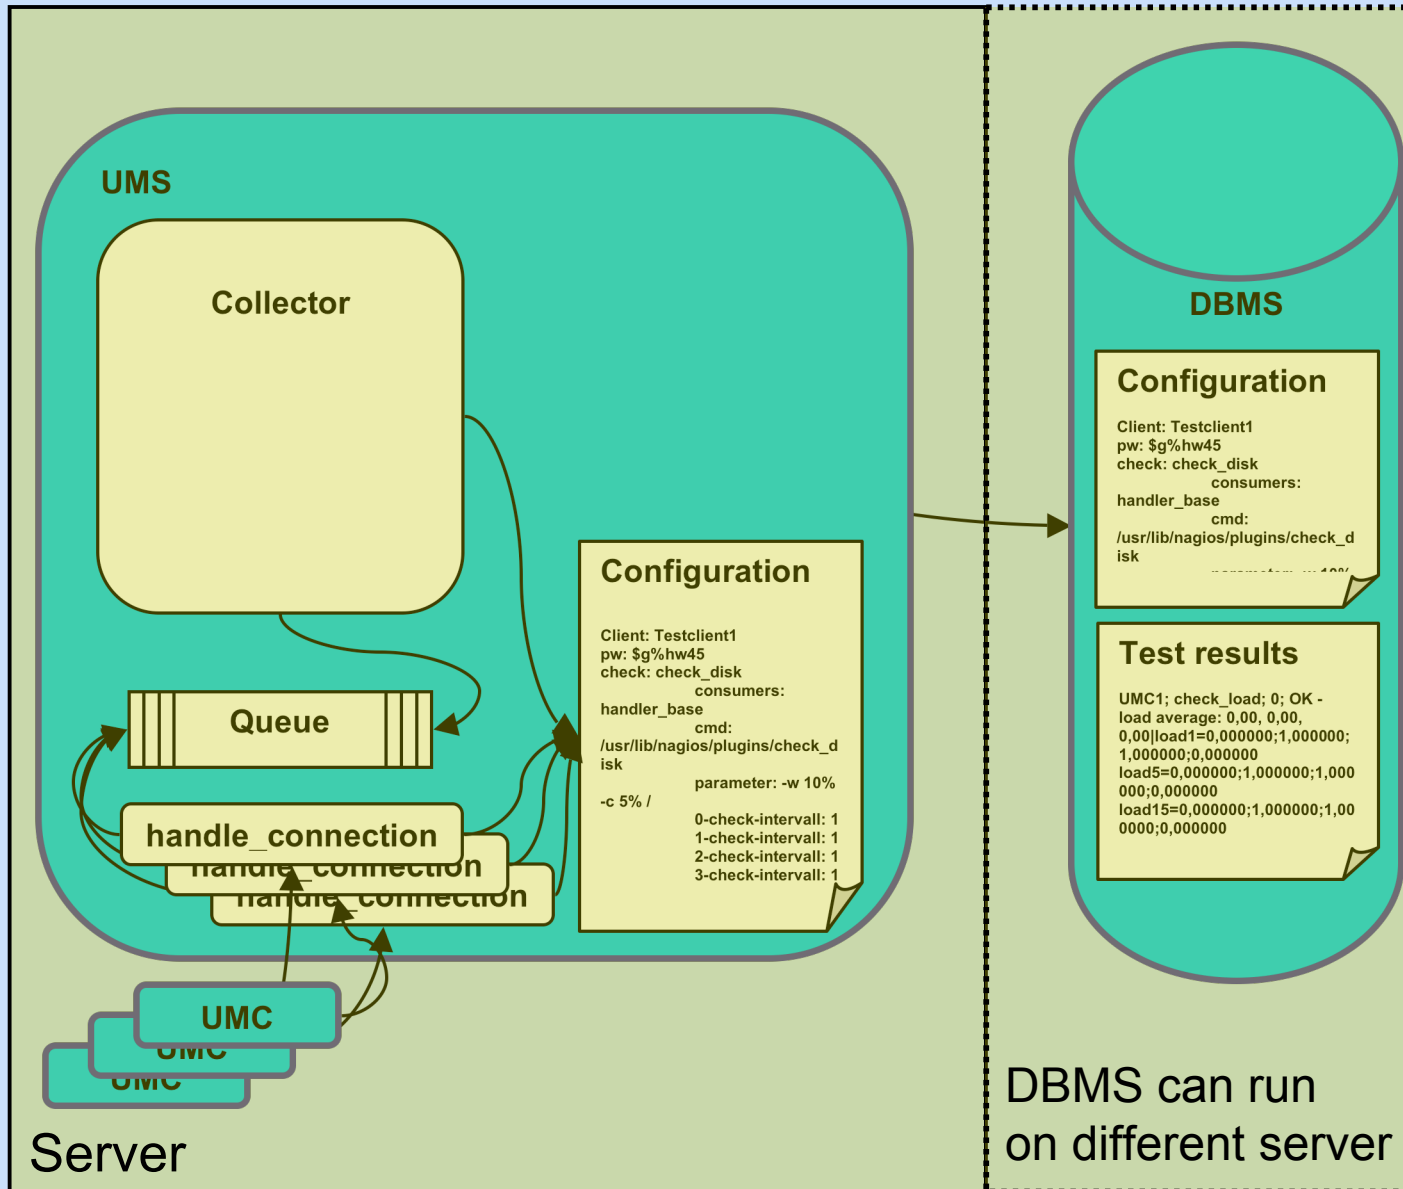


# Unified monitoring client



- UMC fetches configuration
- Configuration is stored on disk
- Configuration is parsed
- Checker threads are started
- Plugin results put in a queue
- Main fetch results from queue
- Save data to backup data file
- Transfer data to UMS
- Ask UMS for new configuration
- Cleanup backup data file

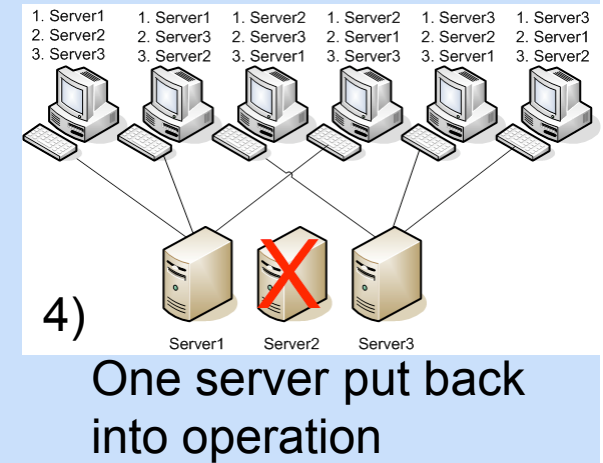

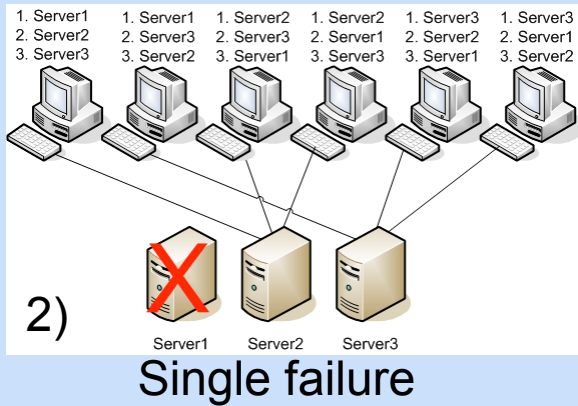

# Unified Monitoring Server




- UMS fetches configuration
- Configuration is parsed
- Listener thread is started
- Handle\_connection is spawned threads are started
- UMC data put in a queue
- Collector fetch data from queue
- Collector stores data in DBMS
- Collector locks for new configuration for UMC
- Transfer configuration to UMC

# UMF failover

Initial scenario



Restore operation to normality



# First results

- ▶ Functionality was verified in a small setup (4 nodes)
- ▶ Developed Nagios plugin “oprofile” to test granularity of scheduled probes

- With a “normal” set of probes we can scale up to more than 2000 clients on Gigabit Ethernet server
- We can do probes at ~1ms
- We can cope with large data sets per client ~290MB/min
- Failover does work

Überwachte Komponente	Tests / Minute	MB / Minute (800 B / Test)
<b>Standarddienste</b>		
AFS	1	0,0008
DHCP	0,1	0,00008
NTP	0,1	0,00008
LDAP	1	0,0008
SSH	0,1	0,00008
Datenbank	1	0,0008
DNS	1	0,0008
<b>Hardware</b>		
Temperaturen	6	0,0048
Lüfterdrehzahlen	6	0,0048
Festplattenzustand	6	0,0048
Spannungen	6	0,0048
<b>Sonstiges</b>		
CPU-Auslastung	600	0,48
Speicherauslastung	600	0,48
Festplattenbelegung	600	0,48
Netzwerkauslastung	600	0,48
Logfiles	600	0,48
<b>Infrastruktur</b>		
Switche	60	0,048
Router	60	0,048
Umgebung	60	0,048
MB / Minute pro UMC:		2,56664

Überwachte Komponente	Tests / Minute	MB / Minute (800 B / Test)
<b>Standarddienste</b>		
AFS	60	0,048
DHCP	0,1	0,00008
NTP	0,1	0,00008
LDAP	60	0,048
SSH	0,1	0,00008
Datenbank	60	0,048
DNS	60	0,048
<b>Hardware</b>		
Temperaturen	6	0,0048
Lüfterdrehzahlen	6	0,0048
Festplattenzustand	6	0,0048
Spannungen	6	0,0048
<b>Anwendung</b>		
Cache Misses	60000	48
Pipeline flushes	60000	48
MemoryAccess	60000	48
Interrupts	60000	48
Bus locked Cycles	60000	48
Pos. im Programm	60000	48
<b>Infrastruktur</b>		
Switche	60	0,048
Router	60	0,048
Umgebung	60	0,048
<b>Sonstiges</b>		
CPU-Auslastung	600	0,48
Speicherauslastung	600	0,48
Festplattenbelegung	600	0,48
Netzwerkauslastung	600	0,48
Logfiles	60	0,048
MB / Minute pro UMC:		290,32344

# Conclusion

- ✓ We have developed a first prototype of a Unified Monitoring Client
- ✓ The UMF has some unique features for data acquisition and transmission
- ✓ We showed promising scalability
- ✓ We could show a probe rates in Milliseconds

# Thank you !

---

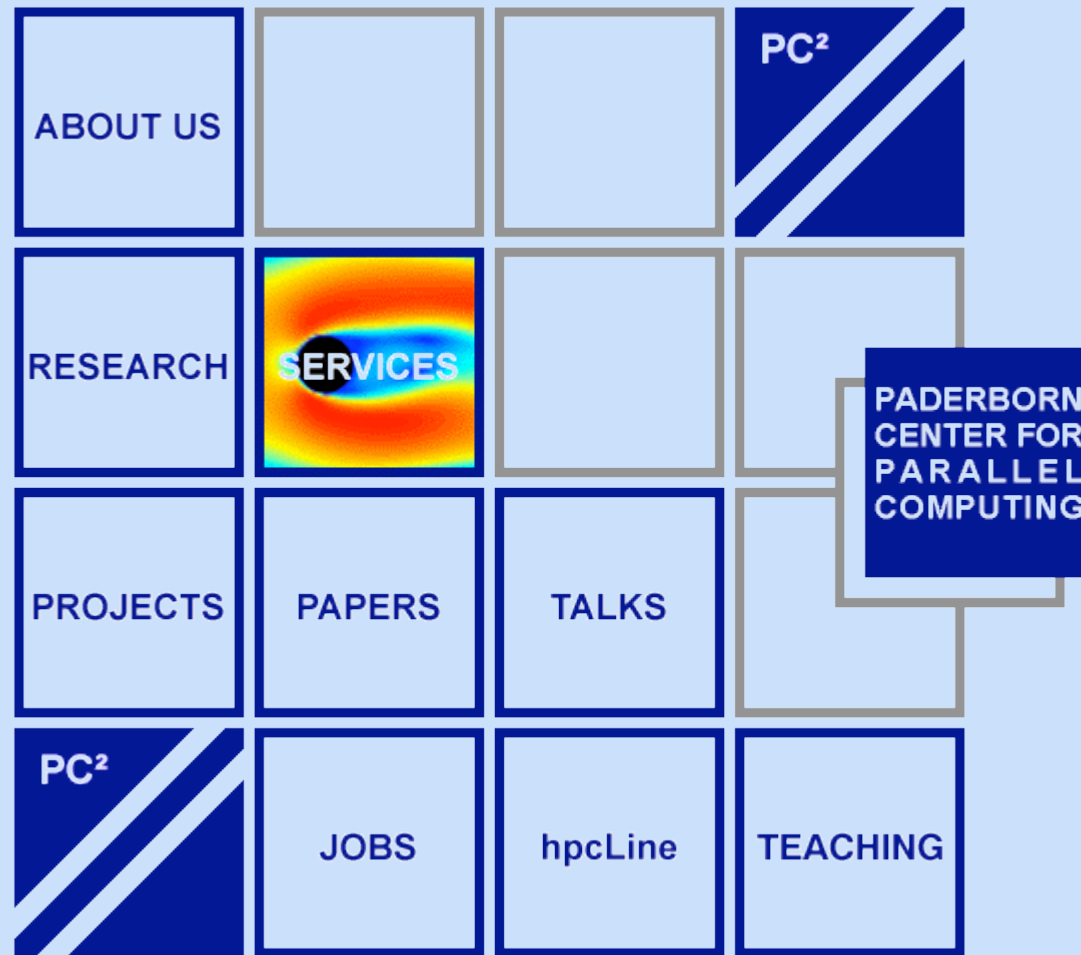
Any remarks or questions?

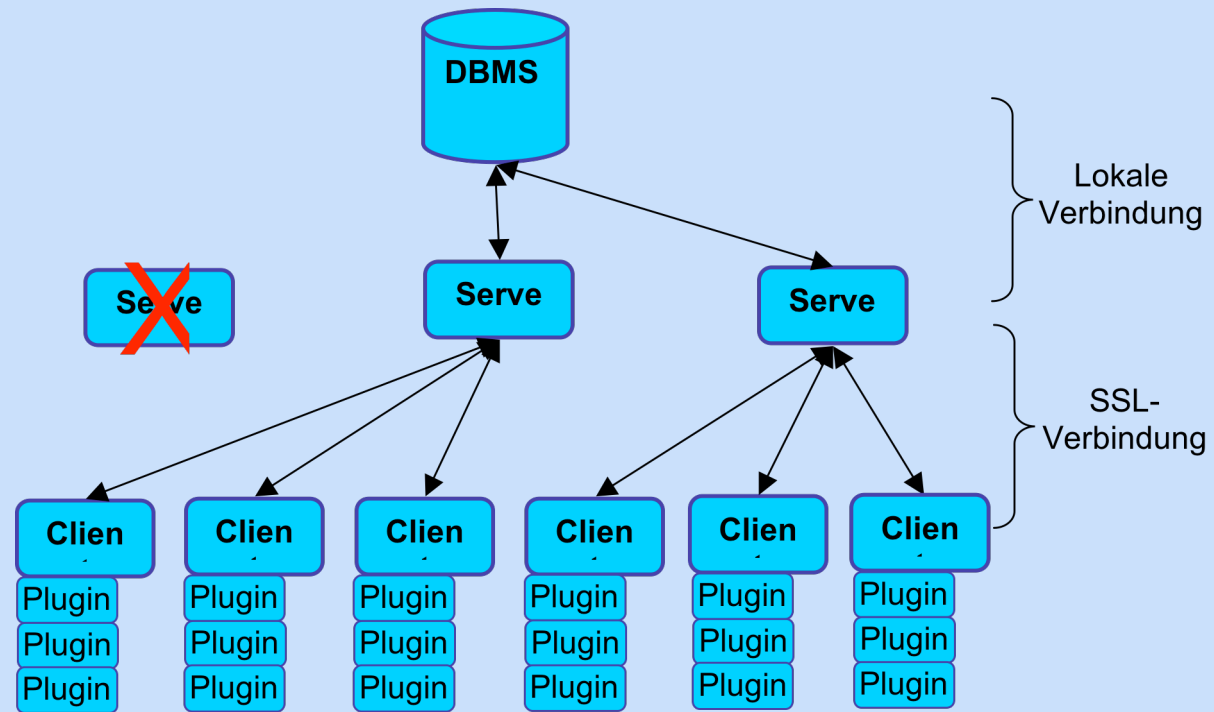


PADERBORN  
CENTER FOR  
PARALLEL  
COMPUTING

# Contact

- ▶ [mailto: hn@uni-paderborn.de](mailto:hn@uni-paderborn.de)
- ▶ PC<sup>2</sup> homepage at <http://www.upb.de/pc2>







# Motivation

- ▶ commercial users have specific demands on service quality level for job execution
  - ▶ SLAs define relationship between end-user and provider (performance, price and penalty fee)
  - ▶ agreeing an SLA is a business risk for providers
- ▶ Therefore: SLA is currently only a concept

