*CoreGRID Workshop at EUROPAR*
*Dresden, 28-29 August 2006*

# Architecture of a Network Monitoring Element

Augusto Ciuffoletti
INFN-CNAF Bologna – Italy

Michalis Polychronakis
FORTH Heraklion - Greece

# Problem and solution

- Problem: design a scalable (thousands of domains), easily expandable network monitoring architecture
- Solution:
  - design a basic building block
  - design an interconnection framework to compose their functionalities
  - define the provided service
- Requirements:
  - use sustainable (scalable) monitoring techniques
  - provide security mechanisms

# Topology of a Network Monitoring System

- A network monitoring system grows with the square of the number of nodes (does not scale)
- The system is partitioned into domains, that reflect network connectivity
- Grid services that are reachable with comparable performance are included in the same domain:
  - not related with DNS
  - not related to administrative domains
  - possibly dependent on link level structure
  - dynamically configurable (with latency of minutes)
- Alleviates (does not solve) scalability problems

# The Network Monitoring Element

- The Network Monitoring Element(s) concentrates the Network Monitoring capabilities of a Domain.
  - It has access to the directory of Network Monitoring Sensors in the domain.
  - It controls the activity of Network Monitoring Sensors.
  - It manages the directory that associates services to domains.
  - It offers a Network Monitoring Service to user applications
- A Domain may host more than one Network Monitoring Element.

# The Network Monitoring Directory

- The Network Monitoring Directory supports
  - location of Network Monitoring Elements in other Domains;
  - location of local monitoring capabilities (mapped to sensors);
  - location of other kinds of services in other Domains;
  - key distribution to support secure management of the Network Monitoring system.

# The Network Monitoring Service

- Offers an interface to user application in order:
  - to query the public content of the Network Monitoring Directory (e.g. to map services to domains);
  - to modify sensors configuration;
  - to control the production and streaming of observations;
- User authentication is based on an external key distribution service;
- Observations streamed outside sensors bear confidentiality tags.

# All-in-one Network Monitoring Element

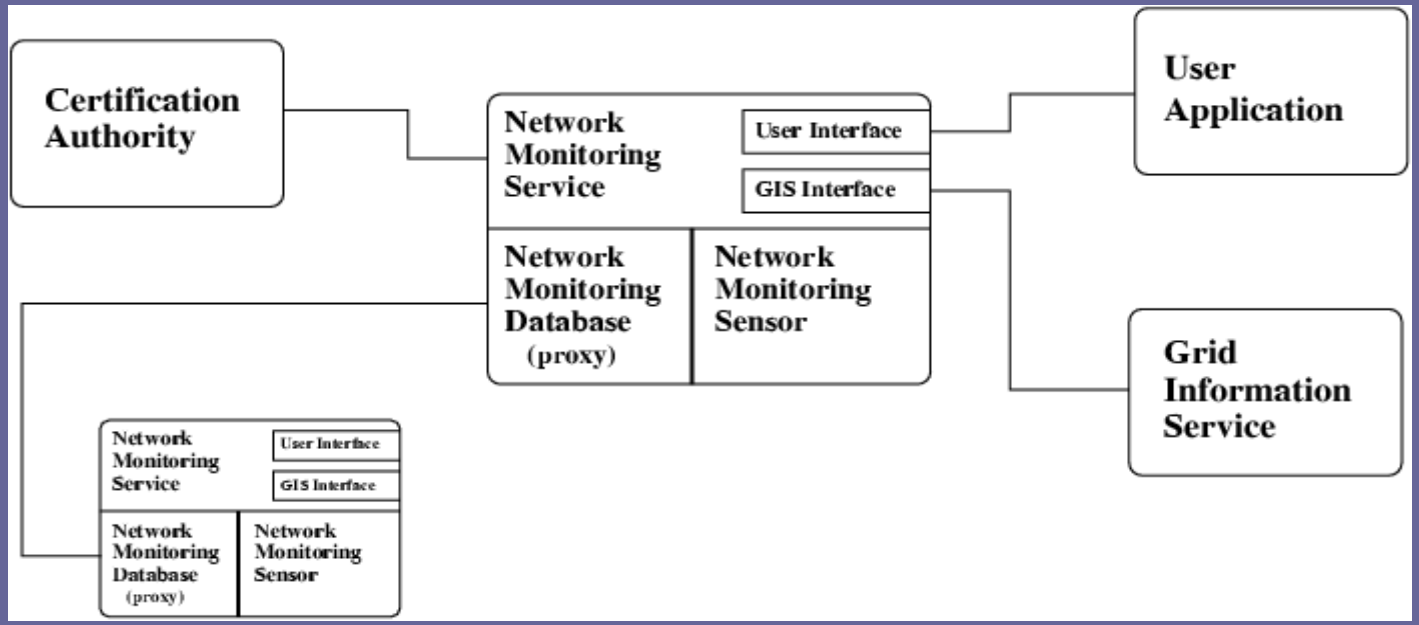


- The three modules (User Application interface, Sensor, Directory) are embedded in a single element.
- A prototype is being designed.

# Design for scalability
# Passive monitoring

- Passive monitoring is introduced to limit the impact of network measurements.
- No injected traffic.
- The MAPI library is used, which allows the remote configuration of sensors.
- Sensors are allowed to configure peer sensors.
- Available metrics include:
  - round-trip delay
  - packet loss rates
  - file transfer rates
- Various forms of aggregation are available

# Measurement of packet loss rates

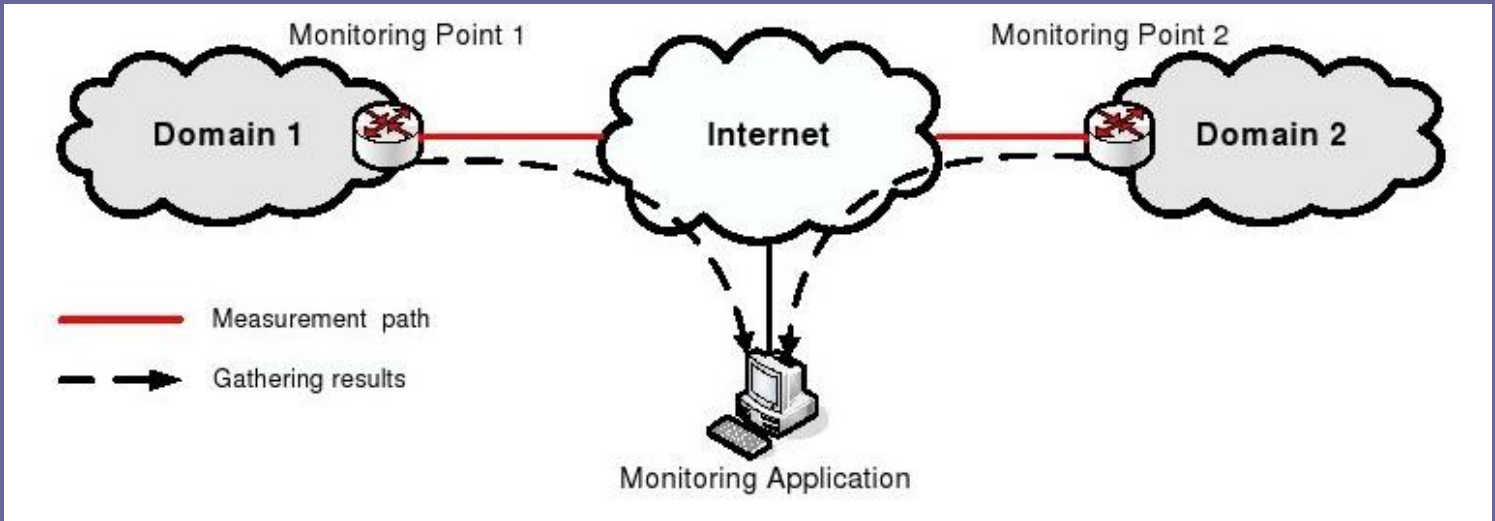


- Packet loss measurement is based on collaboration between two partners
- Sent/Received packets are counted
- Measurement is related to expired flows
- Paper with early experimental results submitted at the Integration Workshop

# Design for scalability
# Gossip based distributed directory

- The key distribution scheme that supports internal security is based on a peer to peer scheme:
  - public keys are certified by a certification authority,
  - certified public keys are downloaded only once,
  - they are circulated using tokens,
  - tokens are encrypted using public keys.
- The number of tokens circulating in the system scales with the logarithm of number of domains.
- Token length scales linearly.

# Distributed directory updates distribution

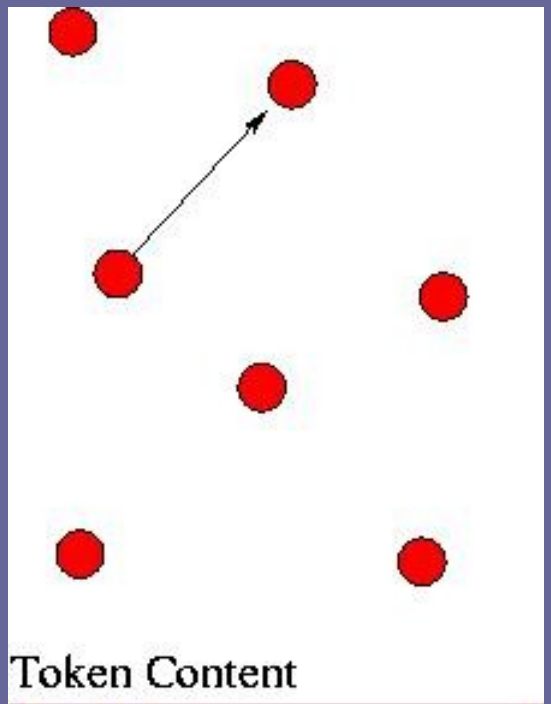


- The token contains authentication data and recent database updates
- Token passing uses a reliable 3-way protocol (over UDP)
- The number of tokens dynamically adapts to system size
- The number of tokens is controlled by a self-stabilizing algorithm
- Simulation results accepted at DAPSYS (Innsbruck)

# Design for scalability
# Application driven monitoring

- This approach ensures the scalability of our design under the following assumptions:
  - the number of applications grows linearly with the size of the system.
  - the amount of monitoring related activity required by an application is constant on system size.
- Policing and shaping of monitoring requests is applicable to control monitoring activity:
  - avoid activity bursts induced by network problems
  - associate priorities to monitoring tasks

# Conclusions

- The design is still on paper, but based on running prototypes and deployed solutions.
- It targets Grid systems hosting thousands of domains, and is designed for scalability.
- Security is kept into account, paying special attention to access to network monitoring configuration.
- A prototype implementation is being developed.

# Help Wanted

- Measurement publication
- Application Driven monitoring
- Deployment of a testbed for distributed database testing (small overhead, plug&play installation)