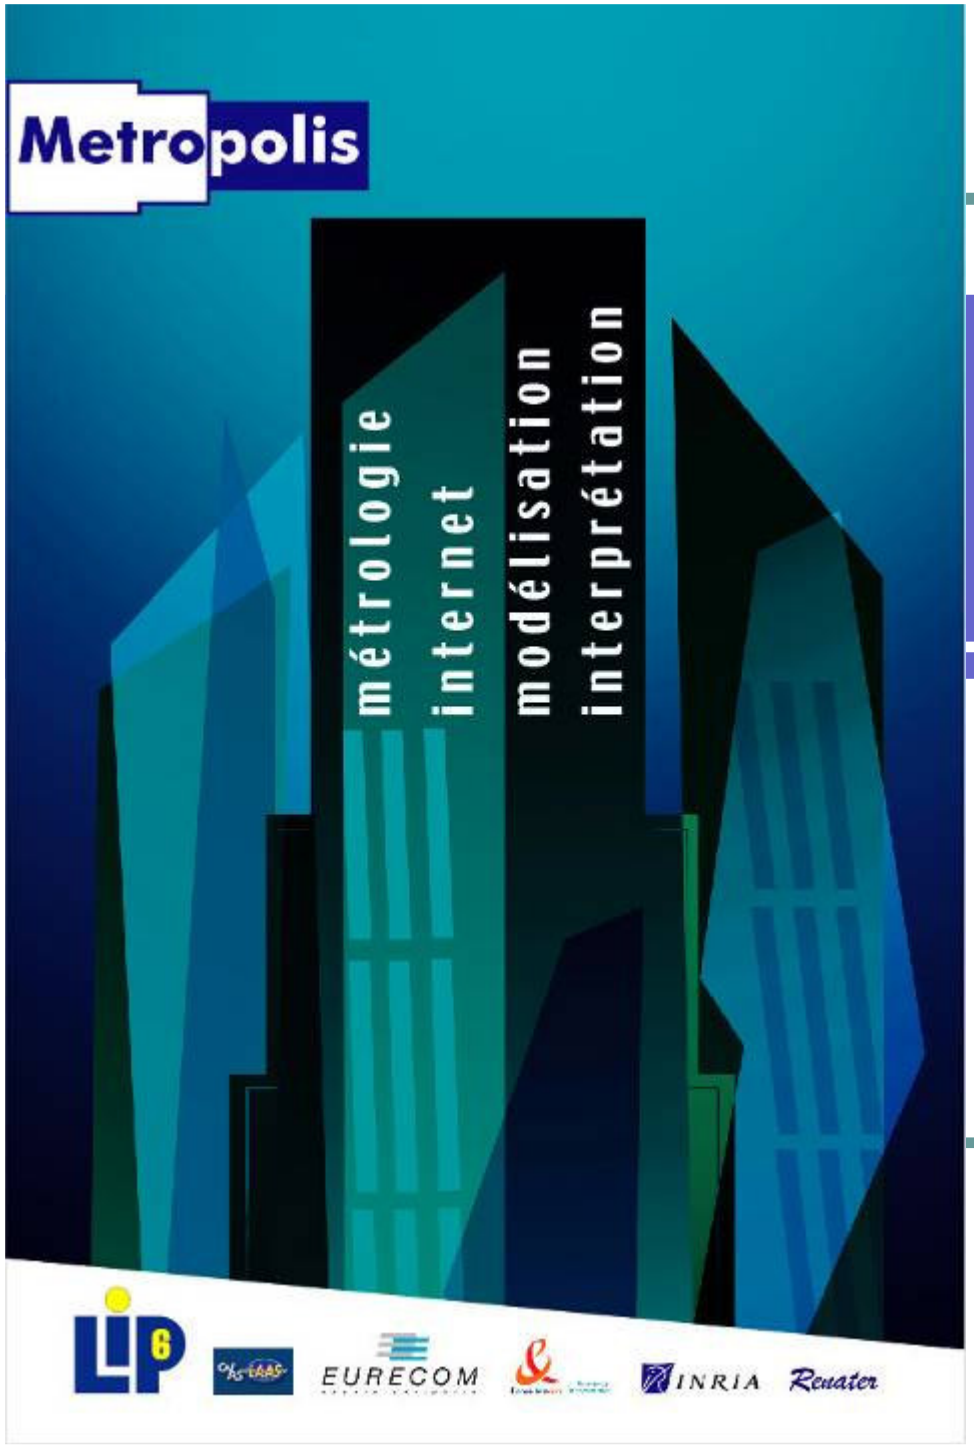

3 June 2004
ISMA Data Catalog Workshop

Metropolis

Experiences from the French measurement infrastructure

Timur Friedman
P&M Curie Univ., Paris
LiP6 lab

Overview

- The Metropolis Project
  - who we are
  - what measurements we conduct
- ISMA data catalog
  - response based upon our experience

Metropolis

Metropolis project

- Funded by the French government
  - The RNRT funding agency
- 36 month project
  - 2 M€
  - 400 man months
  - ending 2004
  - follow-on: MetroSec (3 additional years)

Metropolis

Aims

- To develop a common framework for the metrology of IP networks
  - QoS measurements and SLAs
  - Development of realistic models
  - Protocol analysis and the *in vivo* study of network behaviour
  - Network dimensioning

Metropolis

Metropolis partners

- The LiP6 lab at Pierre & Marie Curie Univ.
  - Project leader: Kavé Salamatian
- France Telecom R&D
- The GET consortium of engineering schools
  - Telecom Paris, ENST Bretagne, Telecom INT
- The INRIA French national research institute
- The Institut Eurecom engineering school
- The LAAS CNRS lab, Toulouse
- Renater
  - The French national high speed research network

Metropolis

External partners

- Sprint ATL lab
  - Exchange of traces, analysis scripts
- ATT Labs-Research
  - Collaboration with MINC (multicast measurements)
- Boston University
- GEANT European research network
  - definition of measurement instrumentation

Metropolis

Ongoing collaborations

- eNext European network of excellence
  - ~50 institutions
  - LiP6: leader of the measurement taskforce
- NGI European network of excellence
  - ~60 institutions
- NIMI
  - Source code access and development
- CAIDA
  - Use of skitter data
- RIPE TTM
  - Use of testbox data

Metropolis

Renater network

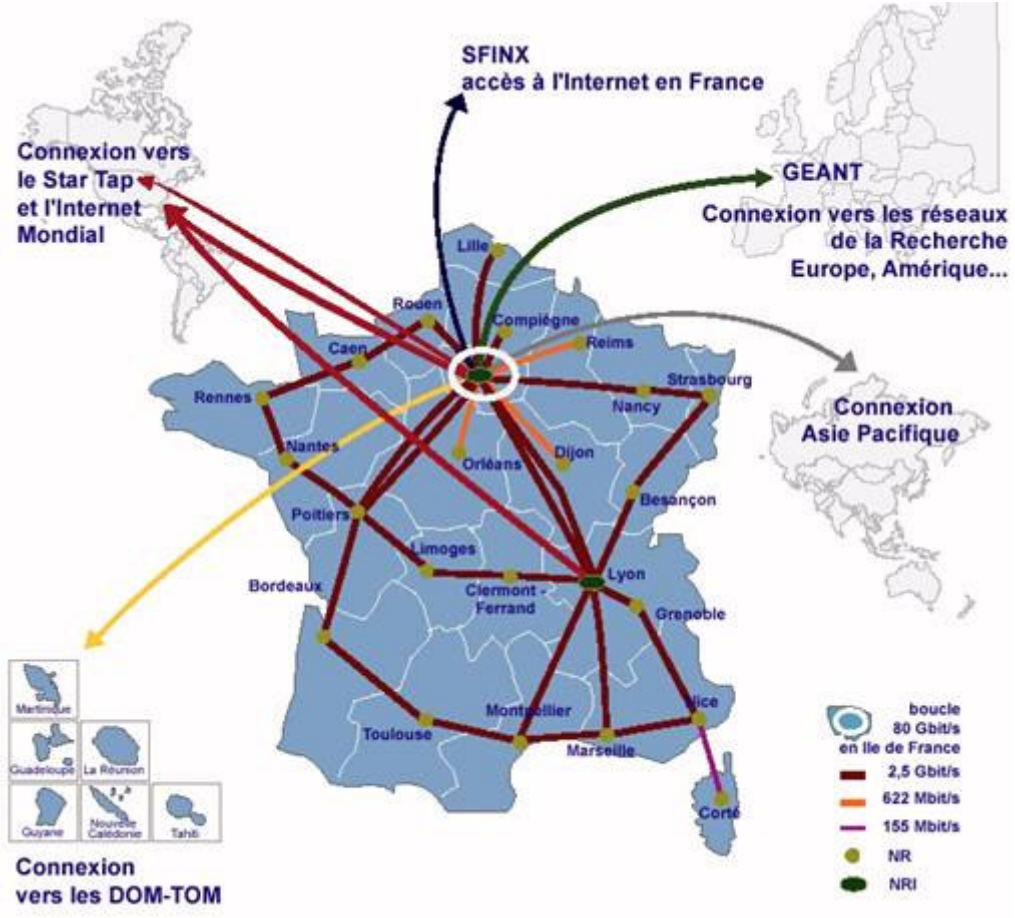

80 Gbps
- Paris region

2.5 Gbps
- Around France
- To EU/US

622 Mbps
- French overseas

40 Mbps
- To Korea

Metropolis

Measurement platforms

- Active measurements
  - RIPE TTM boxes
  - SATURNE boxes (ENST)
  - generic BSD and Linux boxes
- Passive measurements
  - DAG cards (Waikito)
  - QoSMOS boxes
  - Ipanema boxes

Metropolis

Active measurement platforms

- Installed at each of the French partners
  - Extending to eNext European partners
- Hybrid architecture
  - RIPE/Saturne/generic boxes as platforms
  - NIMI for measurement consolidation
- Pandora: A new measurement platform
  - Based on specific μ-kernel for measurement and components architecture
    - Highly scalable, robust, and flexible

Metropolis

Passive measurement platforms

- Microscopic Passive Measurements
  - DAG cards
  - 3 GigaEthernet measurement points
    - Generate around 80 Gbytes of data per day
- Macroscopic passive measurements
  - Ipanema, QosMosMicroscopic passive measurements
  - Flow level measurement
    - QosMos probe
      - Very precise flow classification and application analysis
    - Ipanema probe
      - One way delay measurement

Metropolis

Project publications

- 1 IETF RFC
  - RTCP usage for measurement
- 5 journals
- 22 international conference papers
  - 2004 : 2 Sigmetrics, 3 PAM, 2 ICC, etc…
- 8 French papers
- 14 submissions under review
- 6 common papers between partners

Metropolis

Ongoing research

- Network monitoring
  - Anomaly detection
- Active measurement methodologies
  - Finding good estimator for network parameters
  - SLA validation
- Dimensioning
  - Traffic matrix estimation
  - Access provisioning in presence of P2P
- Wireless measurement characterization
  - GPRS and WIFI

Metropolis

Ongoing research continued

- Traffic engineering
  - Weight assignment
  - Flow classification
- Sampling
  - Scaling the measurement toward OC192 and beyond
- Massively distributed measurement architecture
  - Distributed IDS
  - traceroute@home
    - How to get a realistic view of the network
    - How to fusion distributed topology information
- Localisation
  - How to map IPs to geographical location

Metropolis

Overview

- The Metropolis Project
  - who we are
  - what measurements we make
- ISMA data catalog
  - responses based upon our experience

Metropolis

Our meta-data environment

- Sharing data between partners
  - shared infrastructure, but
    - each user conducts its own measurements
    - each user stores its own measurements
- Making data available to other researchers
- Using data supplied by others

Metropolis

Anonymization

- Heterogeneous requirements
  - site-dependent
    - each partner has its rules about what can be shared
  - dependent upon type of trace
    - passive traces much more sensitive, in general, than active

Metropolis

What was not measured?

- Problem acute with distributed monitors

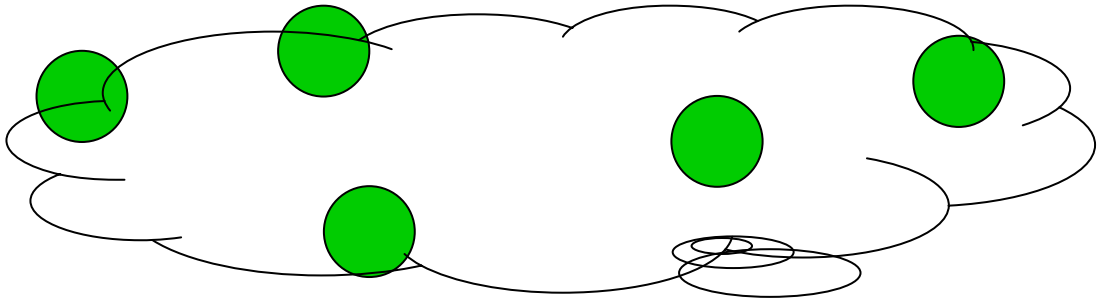

very different from

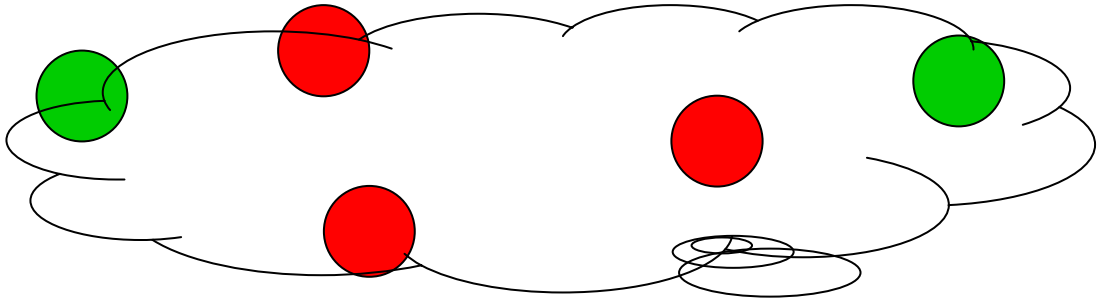

Metropolis

What was not measured? ctd.

- there is the planned experiment
  - which tools?
  - what arguments?
  - on which machines?
  - at what times?
- there is the data that actually comes in
  - command line arguments, return code
  - STDERR as well as STDOUT
  - NIMI good at this
    - build meta-data generation into distributed measurement systems

Metropolis

Identifying data

- support a data publication requirement
- support idea of a collection identifier
  - a citable URI for the data
    - isma://data-owner/id, or
    - isma://data-owner/external/data cataloger/id
  - for citing, not for clicking and downloading

Metropolis

XML

- We find it useful to convert data to XML
  - existing parsing, data manipulation tools
  - XML schema to verify proper format
  - using XML for traceroute@home system
- Most tools do not output XML
- Data compression for large datasets?
- Encourage translation native-to-XML
  - front end, publish XML schema
  - problem: how to encourage people to do this?

Metropolis

Fine-grained measurements

- Want to encourage publication of fine-grained measurements
- Example: basic unit in traceroutes
  - a traceroute?
  - each individual probe packet
- Support standardizing the data collected for basic measurement such as these

Metropolis