

Meta-data management issues underpinning Grid and P2P development

*Experiences from GRASP, SWAD-Europe, PELLUCID and CORAS projects
at CCLRC/BITD*

Emphasis: trust & security policy management

Theo Dimitrakos t.dimitrakos@rl.ac.uk

CCLRC Rutherford Appleton Lab.

- RAL has a staff of around 1200 - most are scientists and engineers
- Supports the work of over 10000 scientists and engineers, from industry and University
- Main Facilities Include:
 - World Data Centre for Solar-Terrestrial Physics
 - Molecular Spectroscopy - Infrared, visible and ultraviolet spectroscopy
 - The world's leading pulsed neutron and muon source
- Main facilities enable research into
 - new materials and structures,
 - X-ray laser research,
 - space-based astronomy,
 - particle physics.

CCLRC Business & Information Technology

- To facilitate Technology Translation and Industry Take-Up
- To contribute to the emergence of new ICT paradigms for the European / UK Science and Business
- Host the UK W3C Office
- UK ERCIM member
- To empower new CLRC IT & e-Science projects
 - complement the expertise of the CCLRC eScience Centre
 - focusing on integration of the (Semantic) Web and the Grid Services technologies
 - emphasis on e-government / e-business driven problems
- To support the operation of CLRC facilities as a cross-sector IS department



Overview

Meta-data management issues underlie a number of activities at CCLR across eScience and (core business) Information Technology

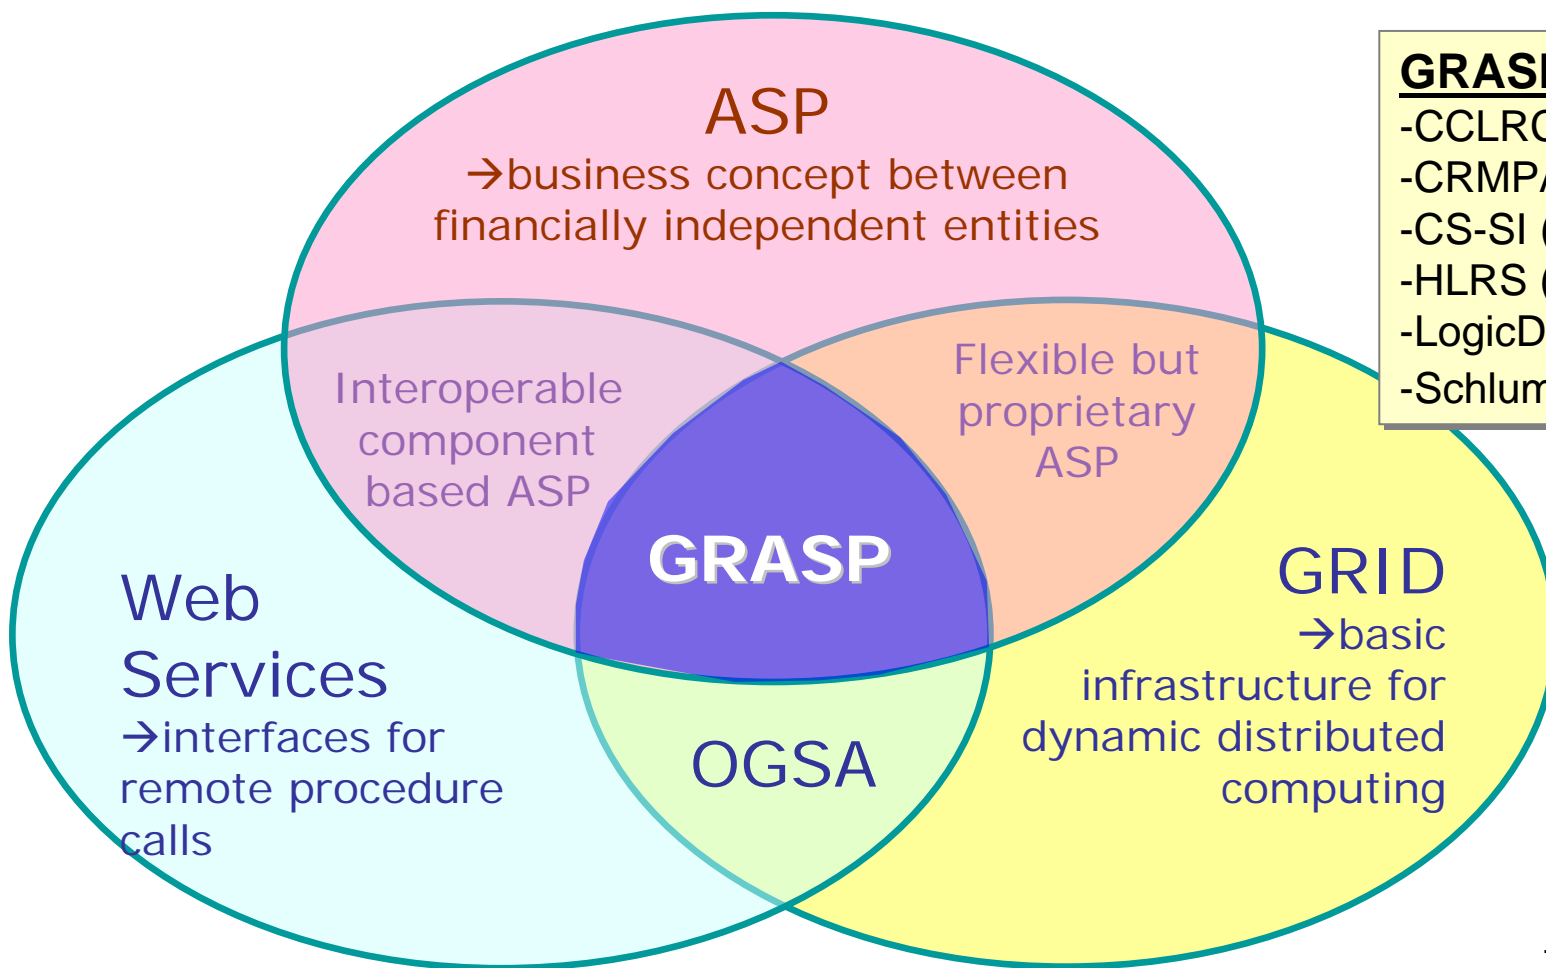
- This talk: an (core business) IT R&D perspective
- Examples:
 - GRASP: Grid-based Application Service Provision
 - CORAS: CASE tool & method support for Security Risk Analysis
 - SWAD-Europe: Semantic Web Technology Development
 - PELLUCID: Agent based platform supporting organisational mobility

GRASP: motivation

- Being mostly used in academic environments, “best-effort” was (and is) a sufficient policy for committing resources to users performing their computational workload.
- Moving into the commercial space, businesses will be bound by commitments. Security, privacy, monitoring and accountability are becoming increasingly important in networked environment. “Best effort” is no longer sufficient.

From “Specifying and Monitoring Guarantees in Commercial Grids through SLA”, Sahai et. al., available at <http://www.hpl.hp.com/techreports/2002/HPL-2002-324.pdf>

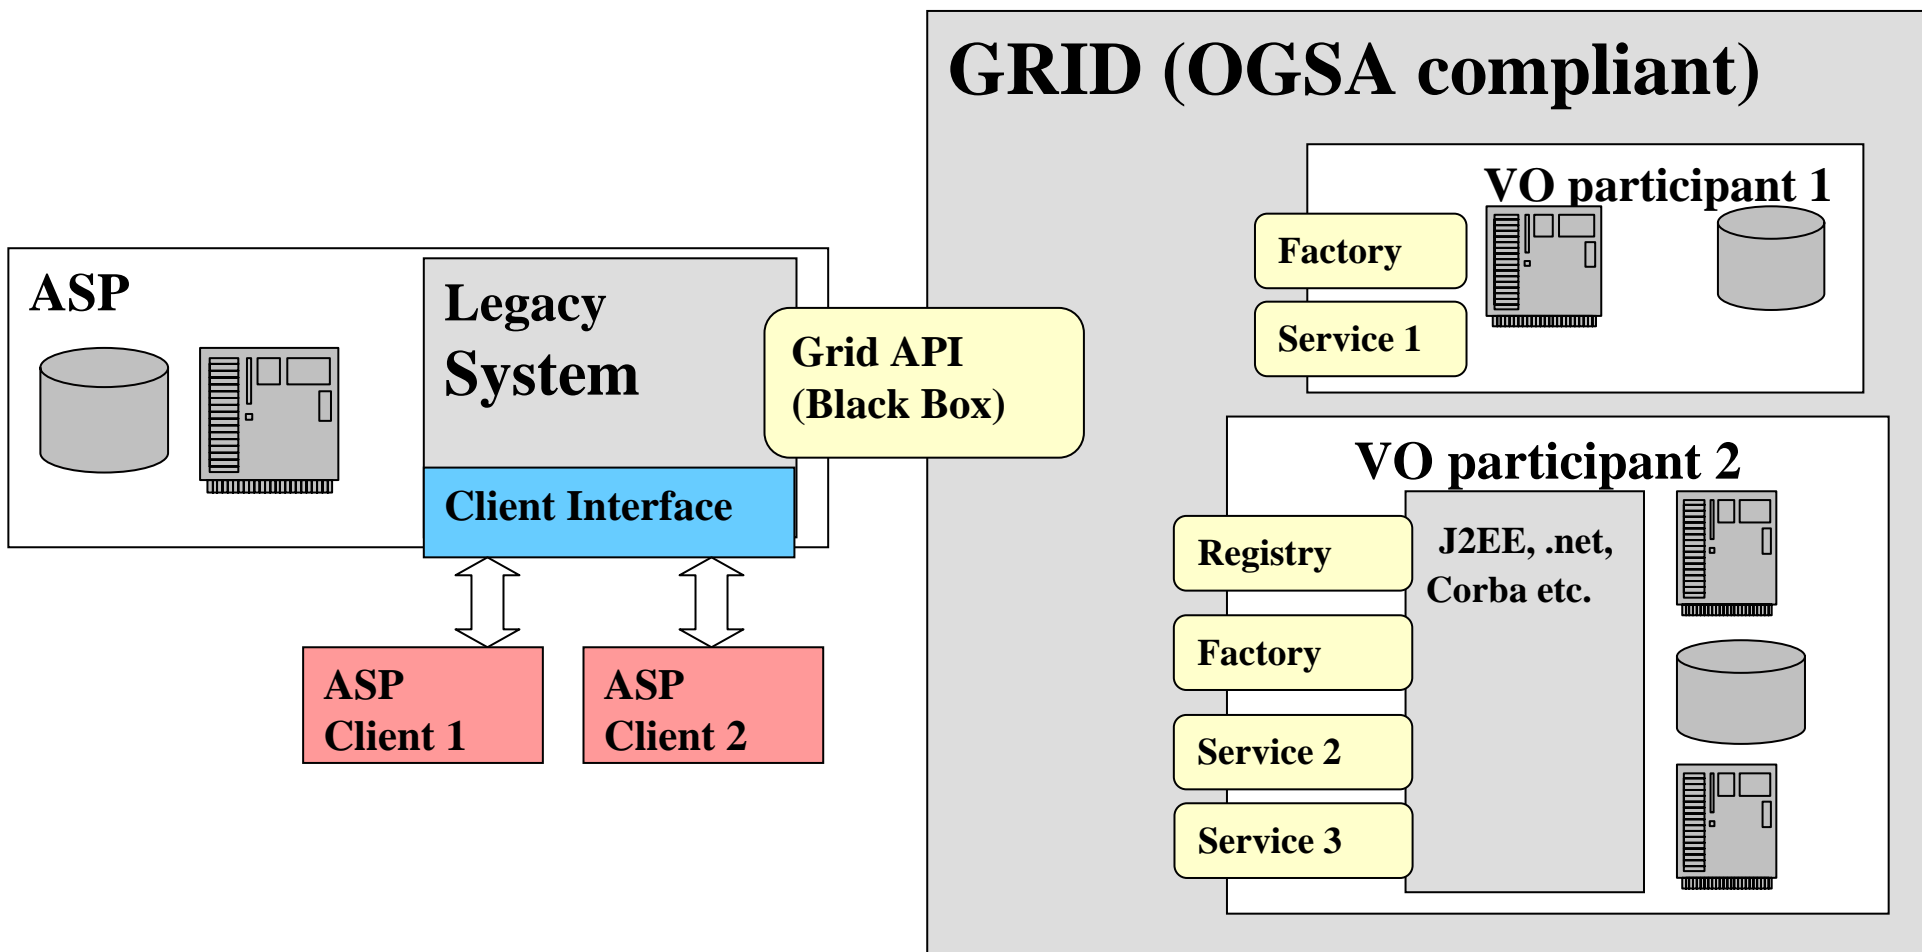
GRASP: approach



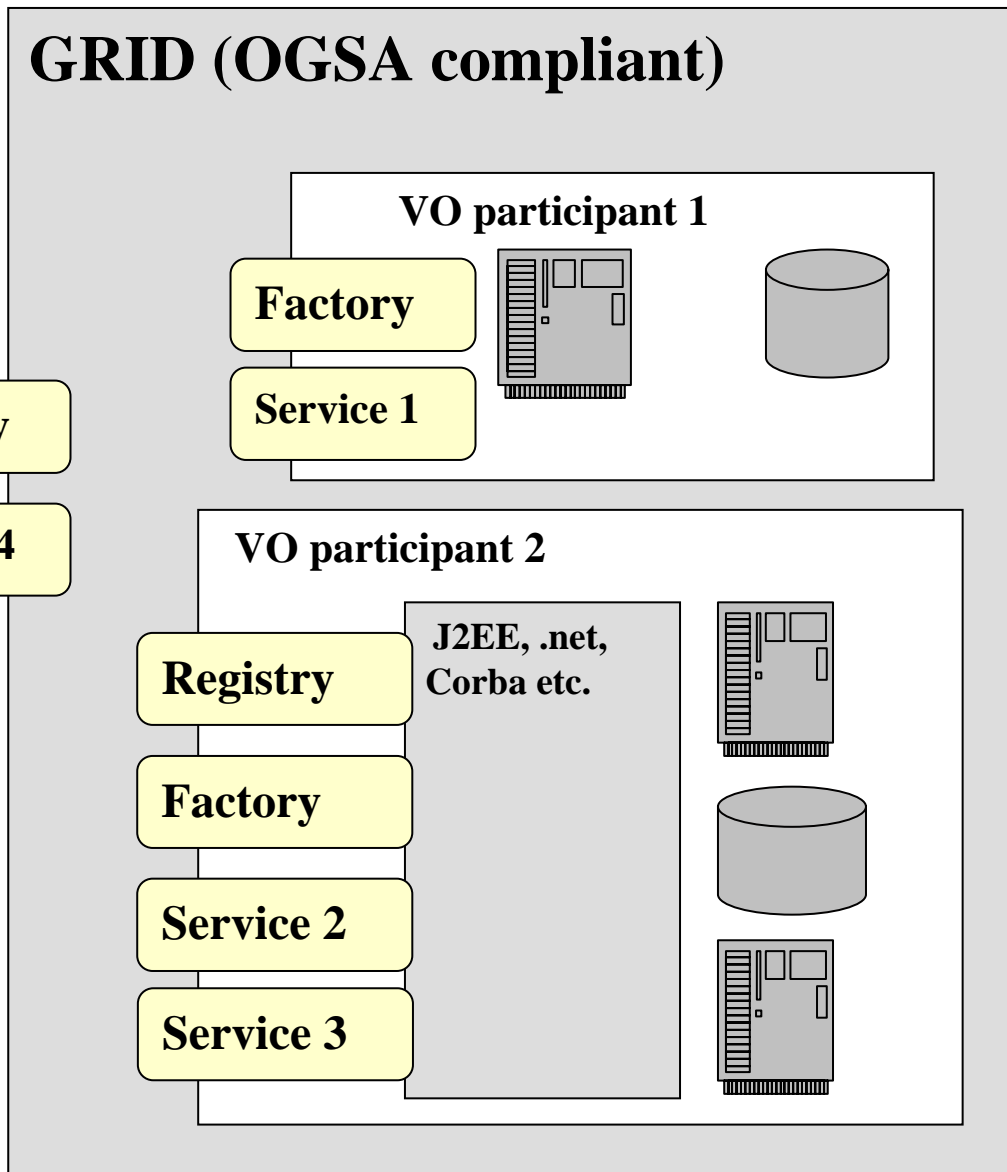
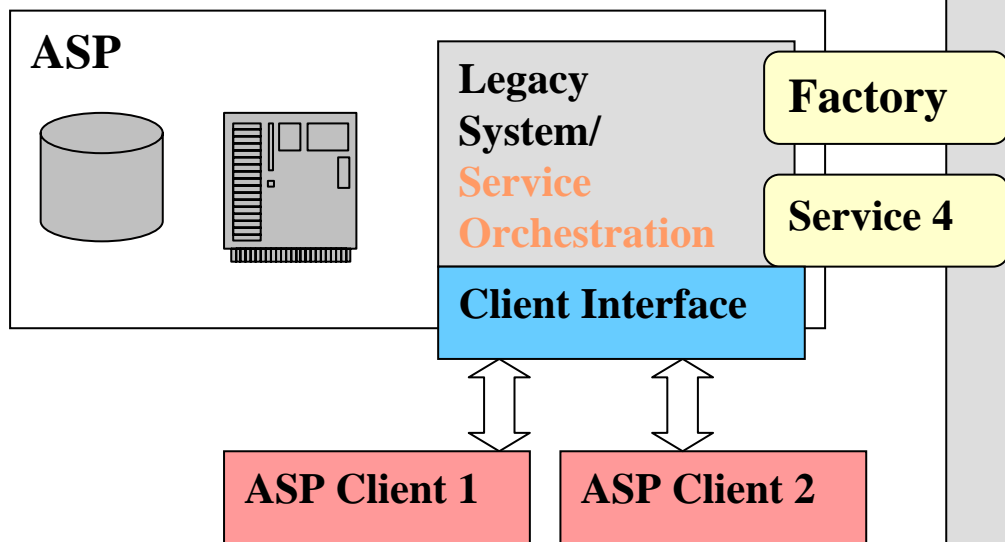
GRASP consortium

- CCLRC (UK)
- CRMPA (Italy)
- CS-SI (France)
- HLRS (Germany)
- LogicDIS (Greece)
- SchlumbergerSema (Spain)

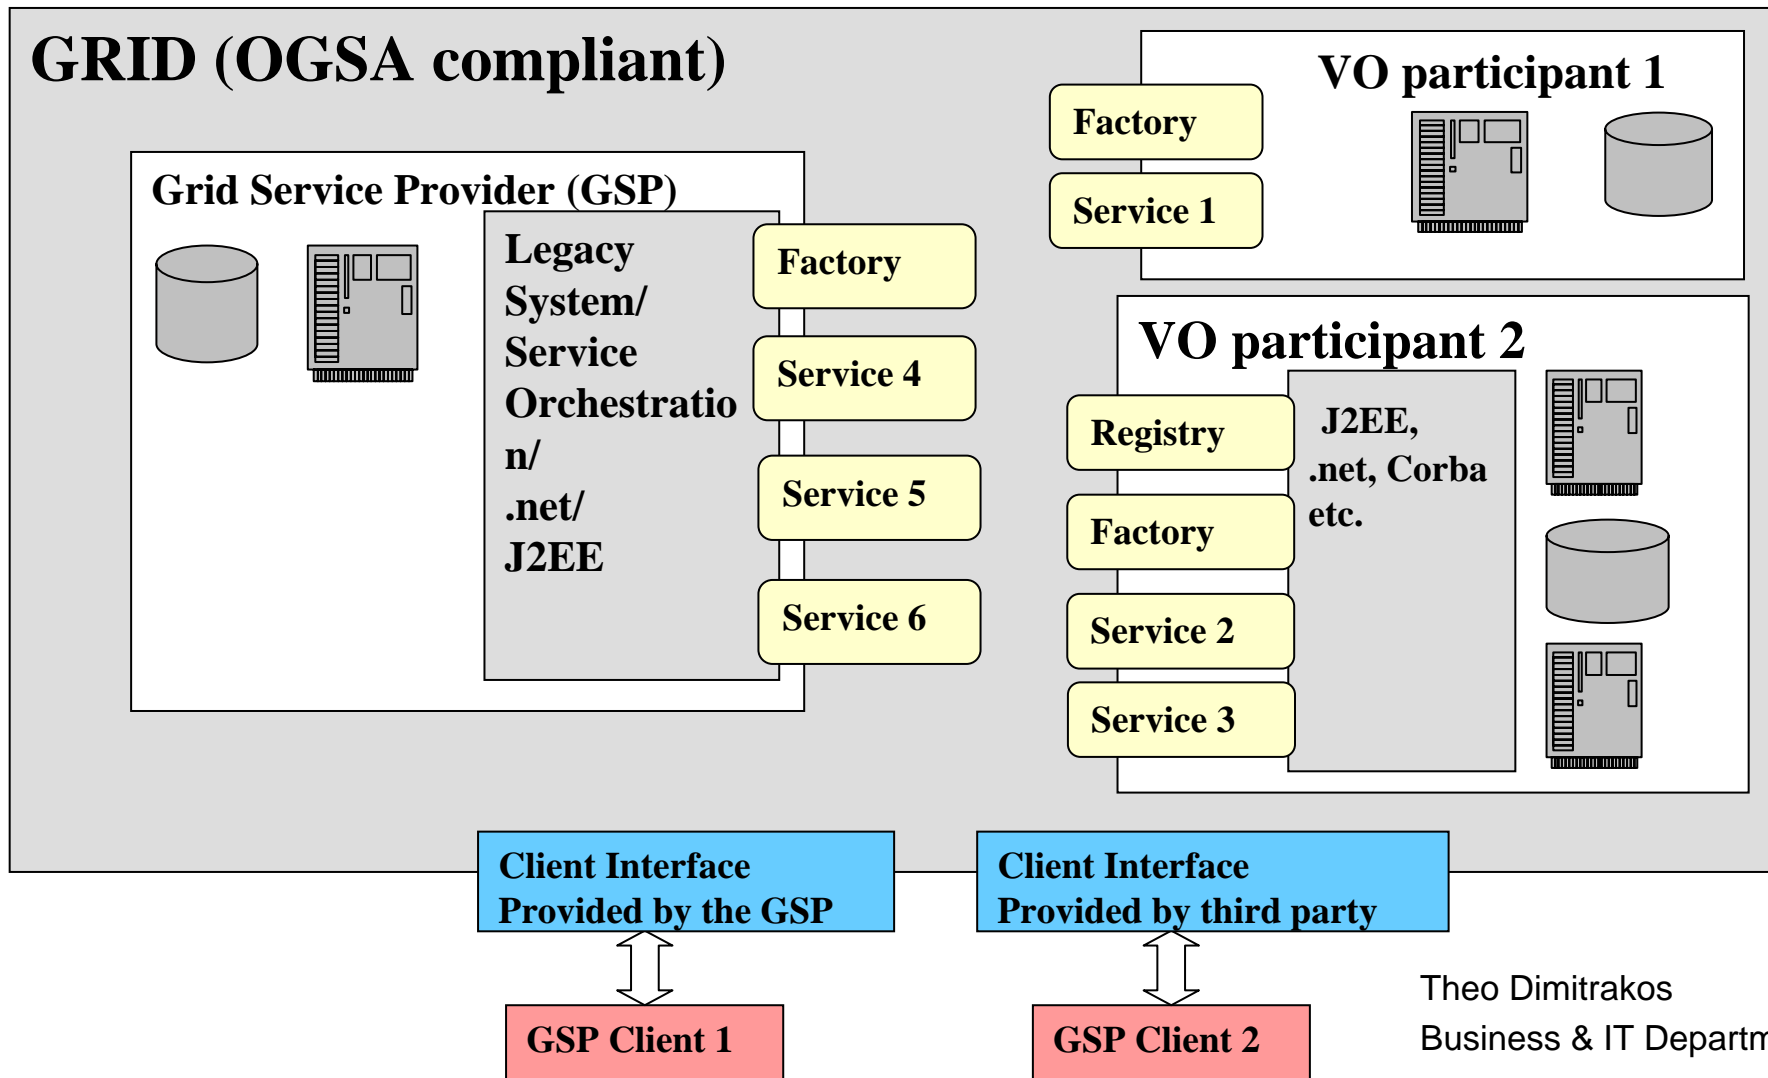
“Grid User”



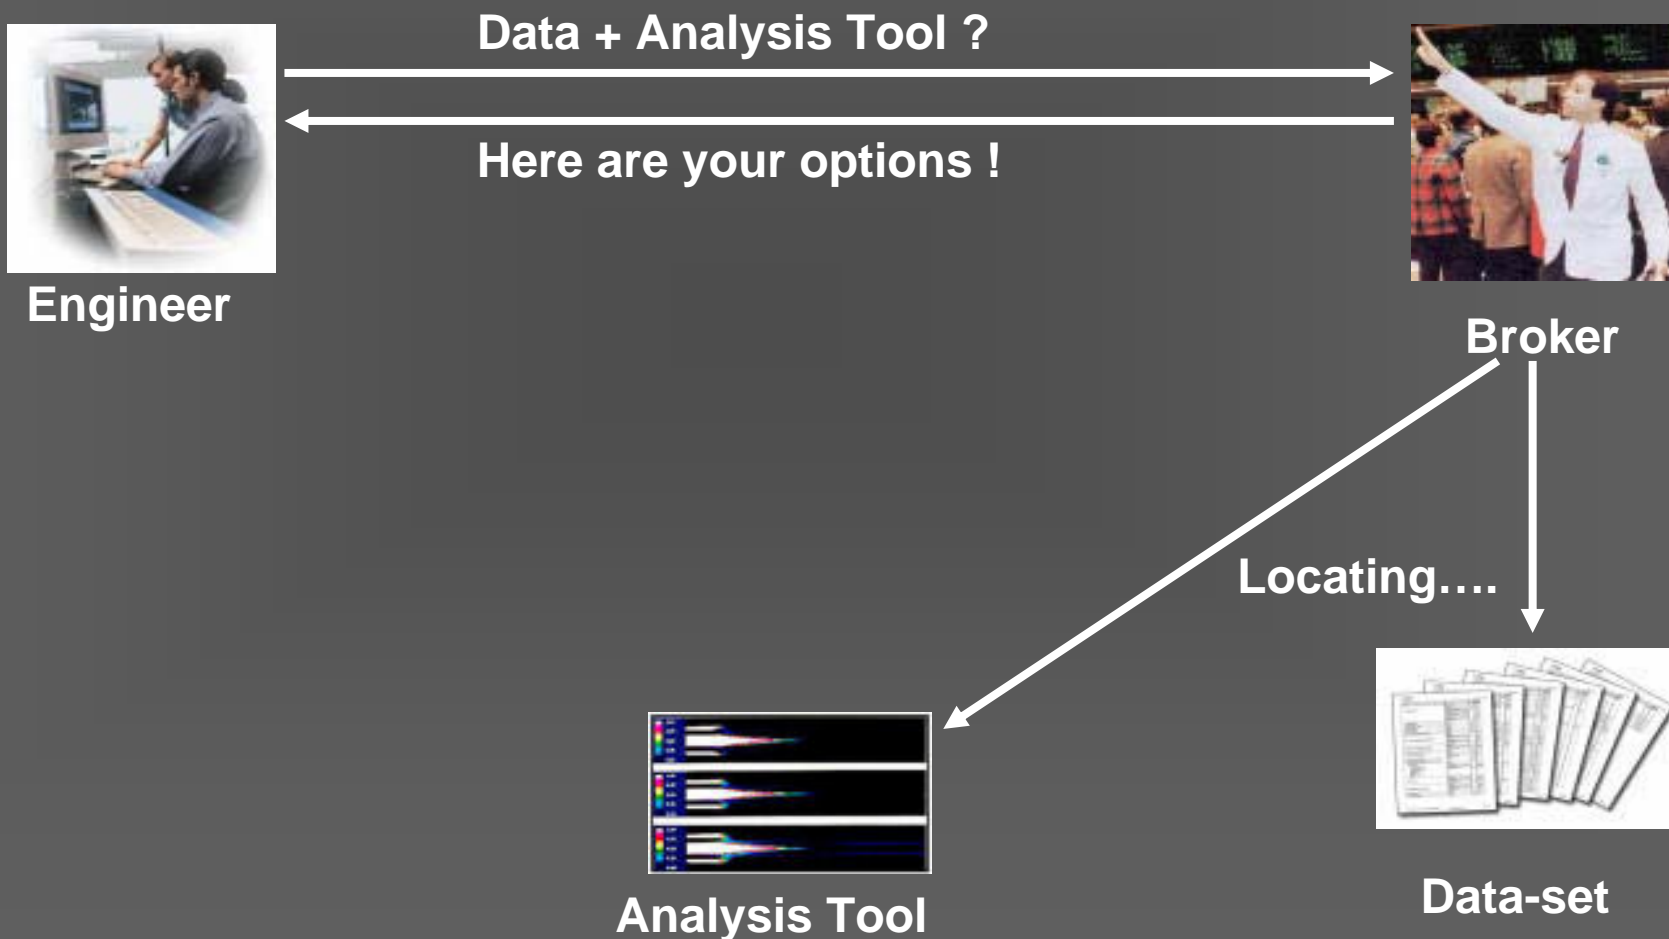
“Grid Enabler”



“Grid Builder”



A scenario driven walkthrough



A scenario driven walkthrough



Locate




Outburst



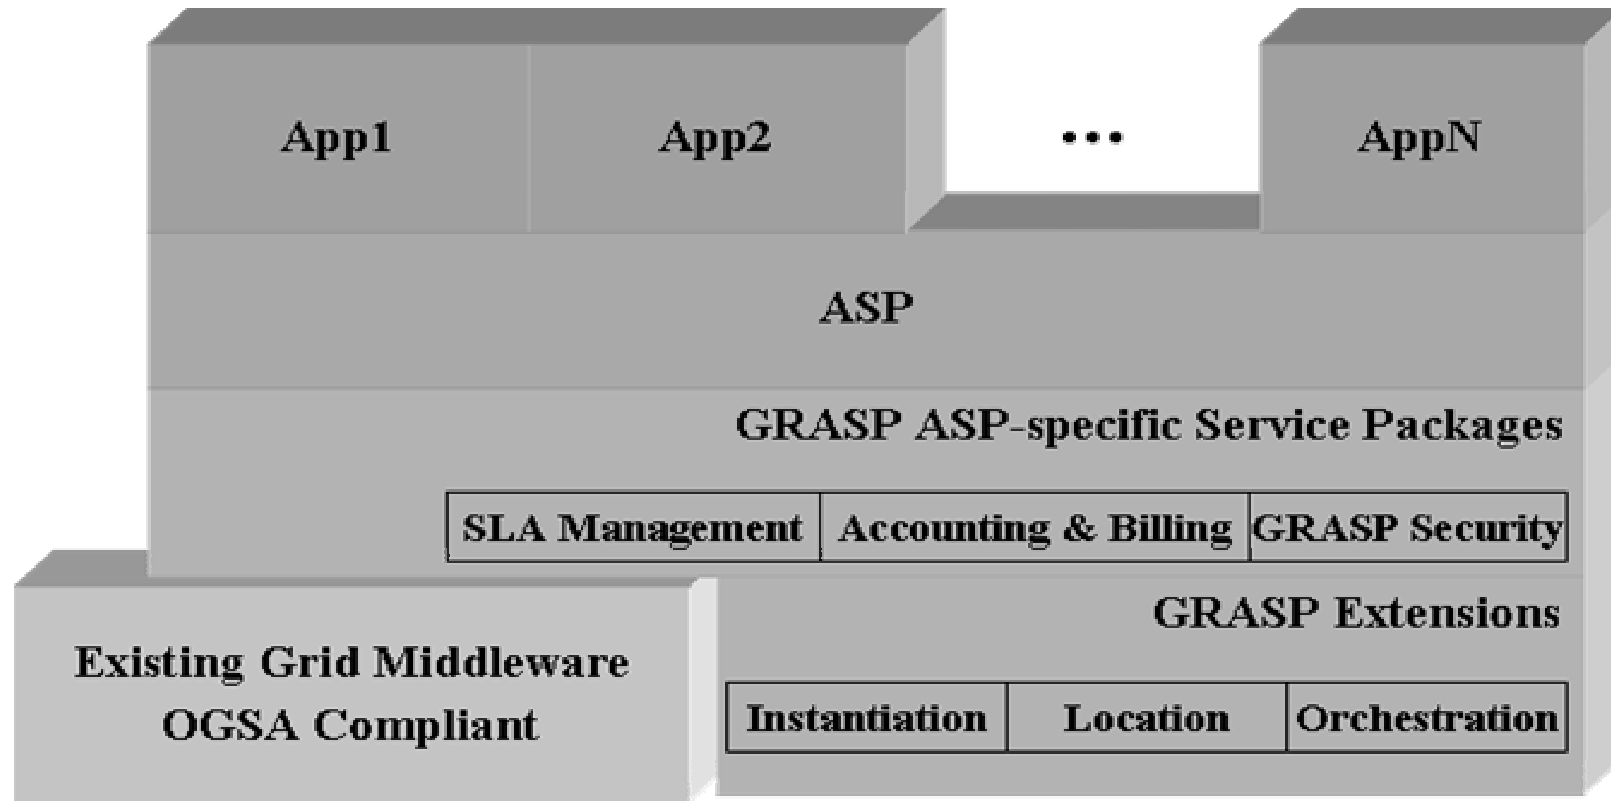
Broccoli

Set-up

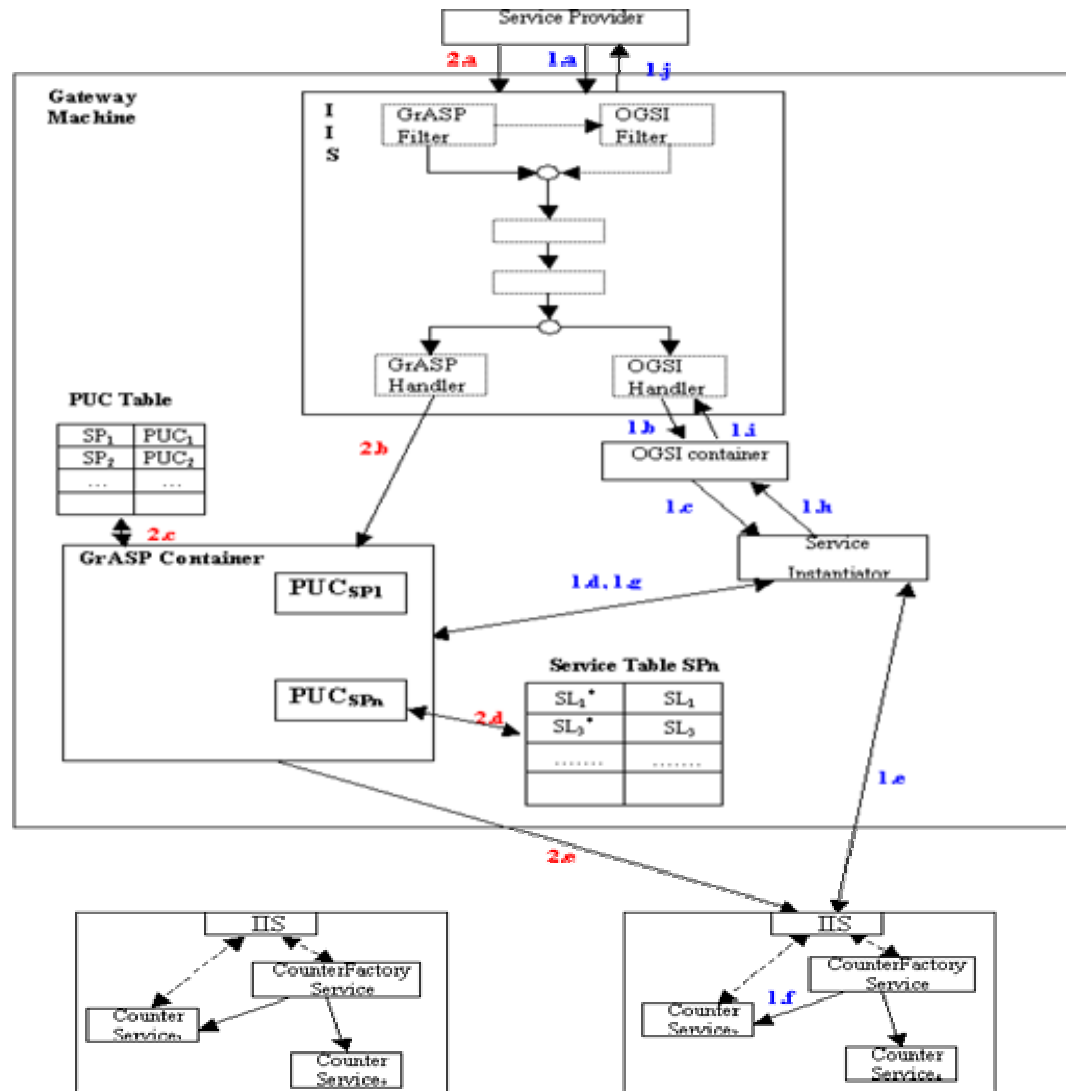


Theo Dimitrakos
Business & IT Department

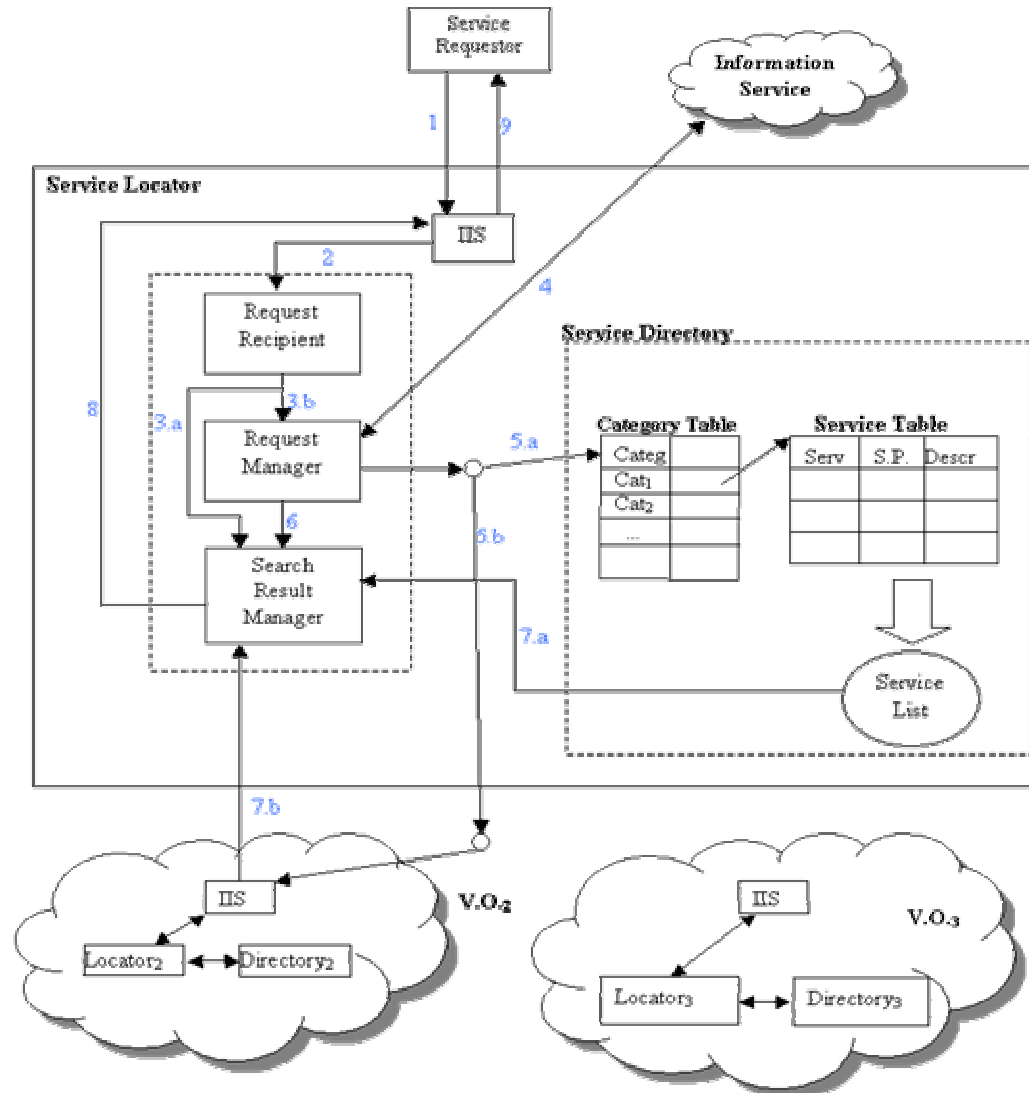
GRASP Architecture



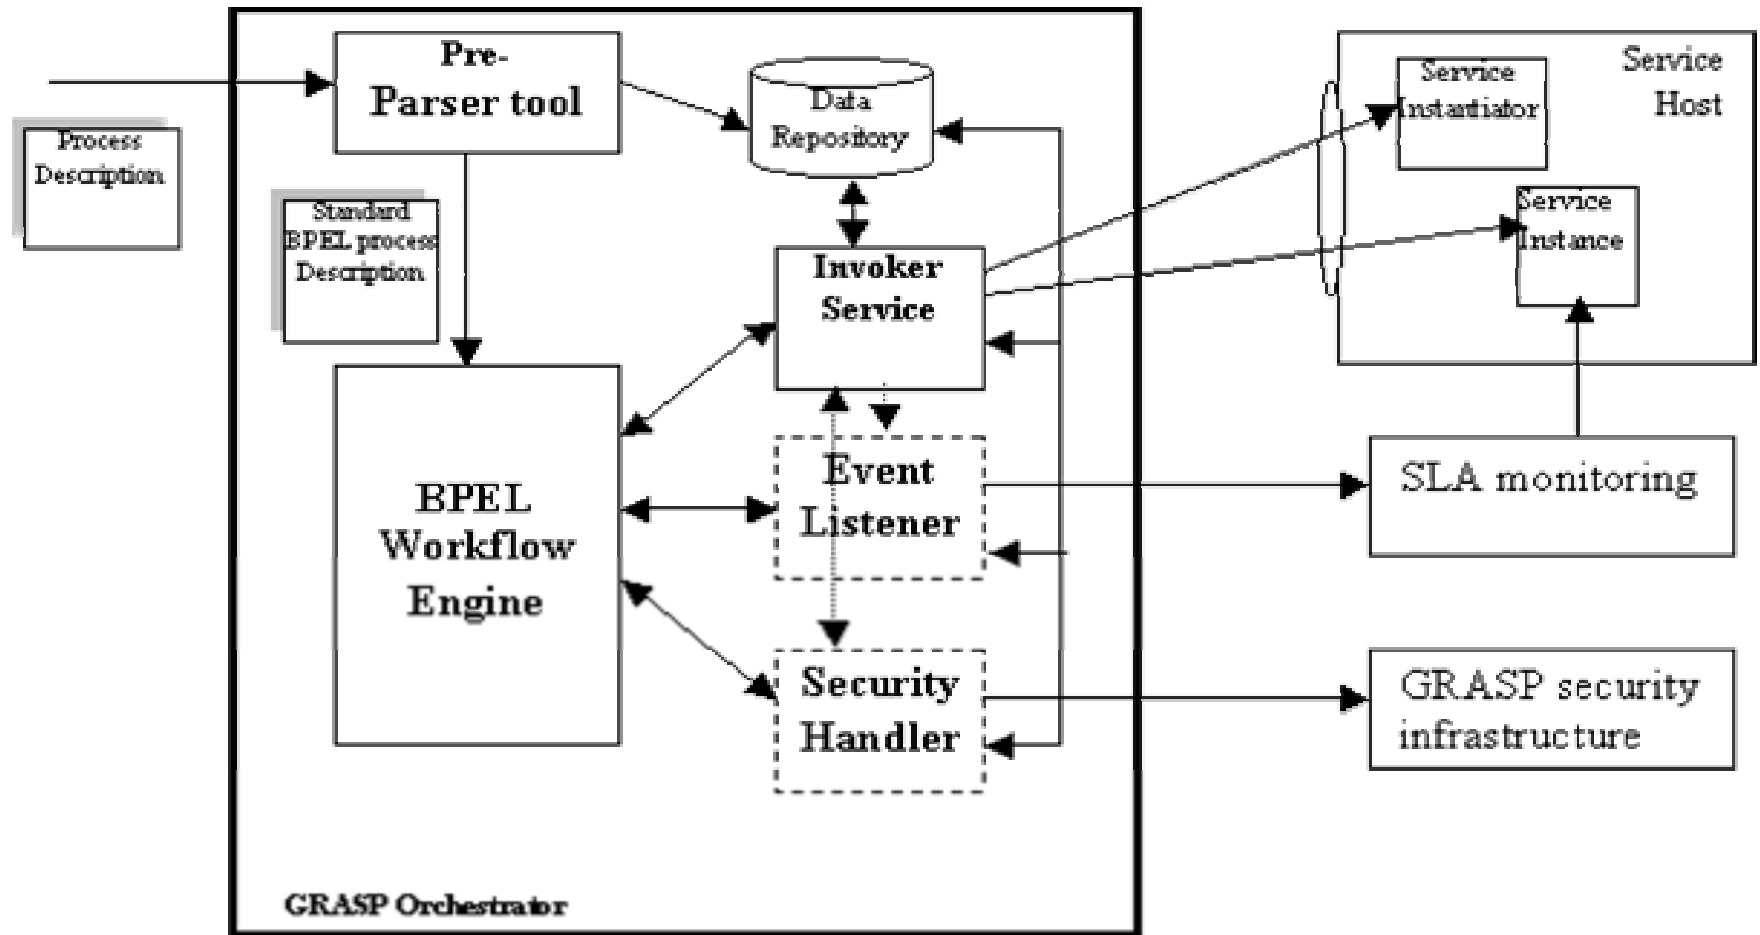
GRASP Architecture: GS-Instantiation



GRASP Architecture: GS-Location



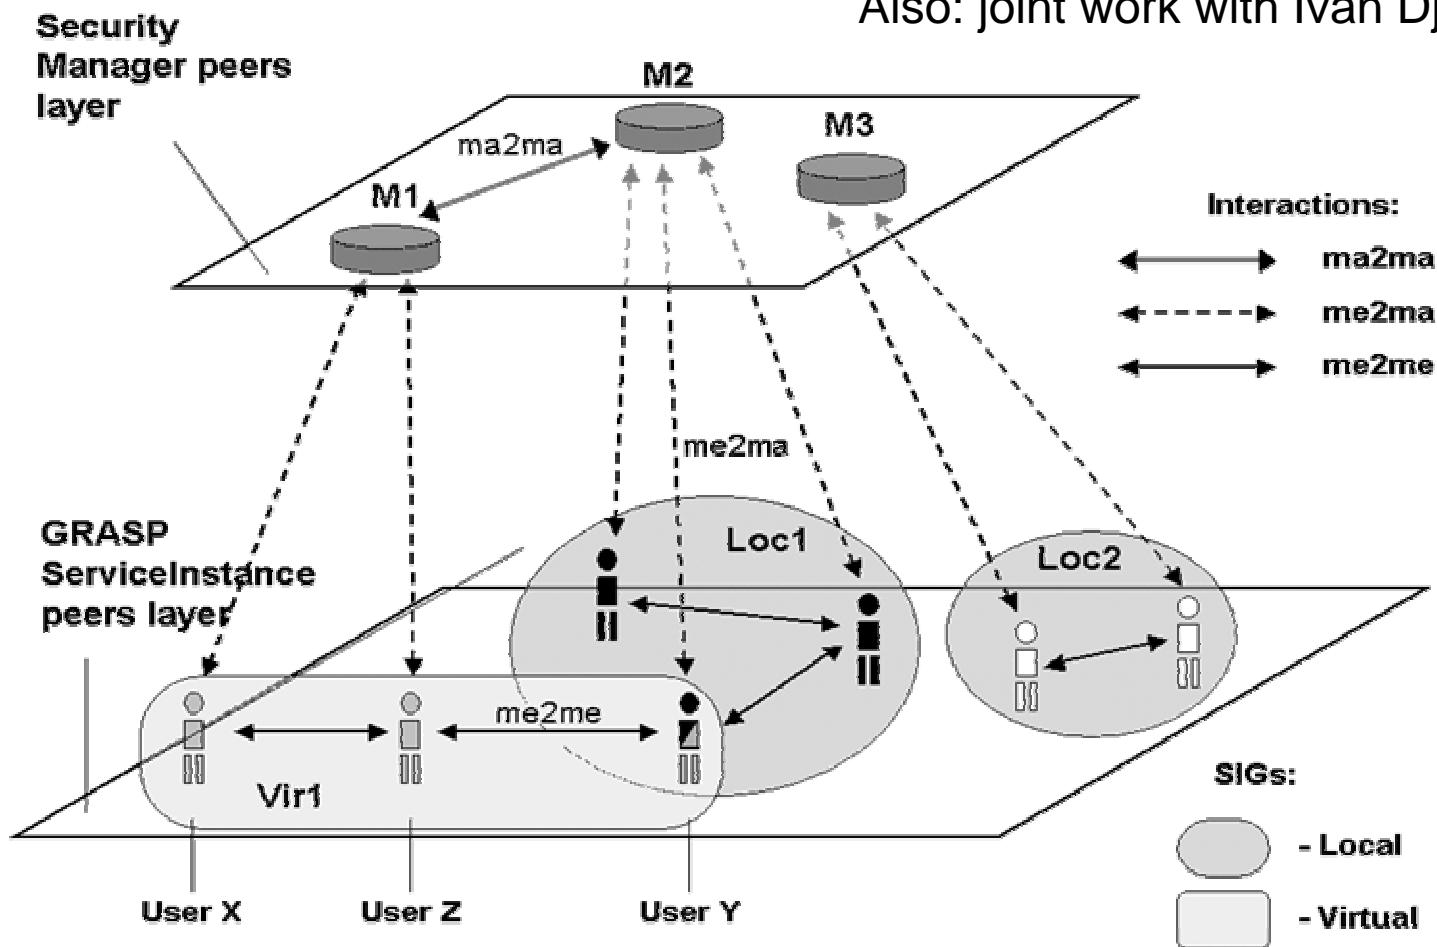
GRASP Architecture: GS-Orchestration



GRASP Architecture: Security

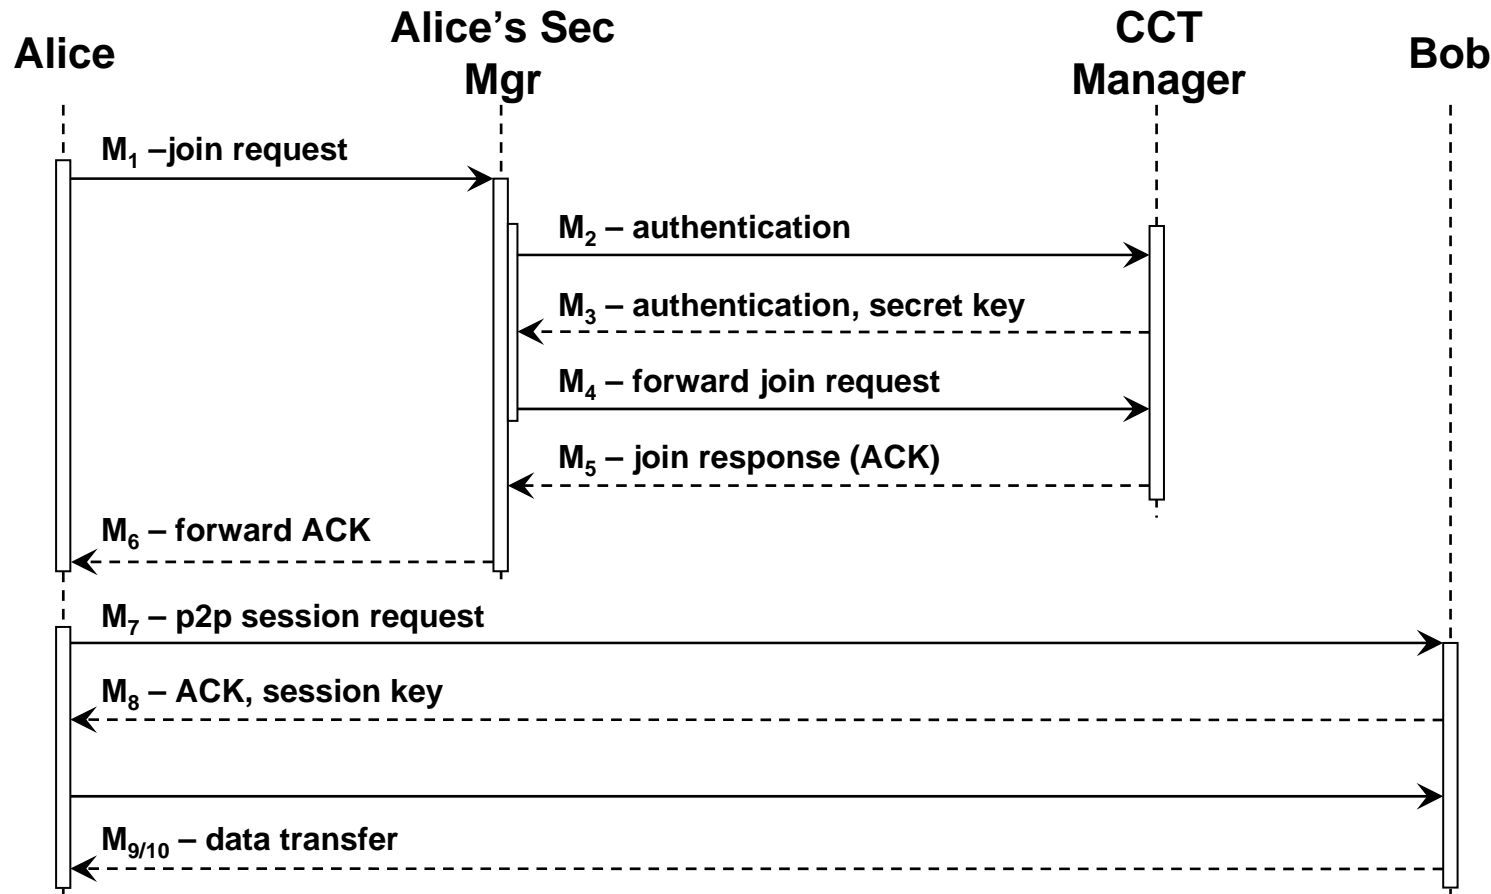
dynamic collaboration networks

Also: joint work with Ivan Djordjevic @ QMUL



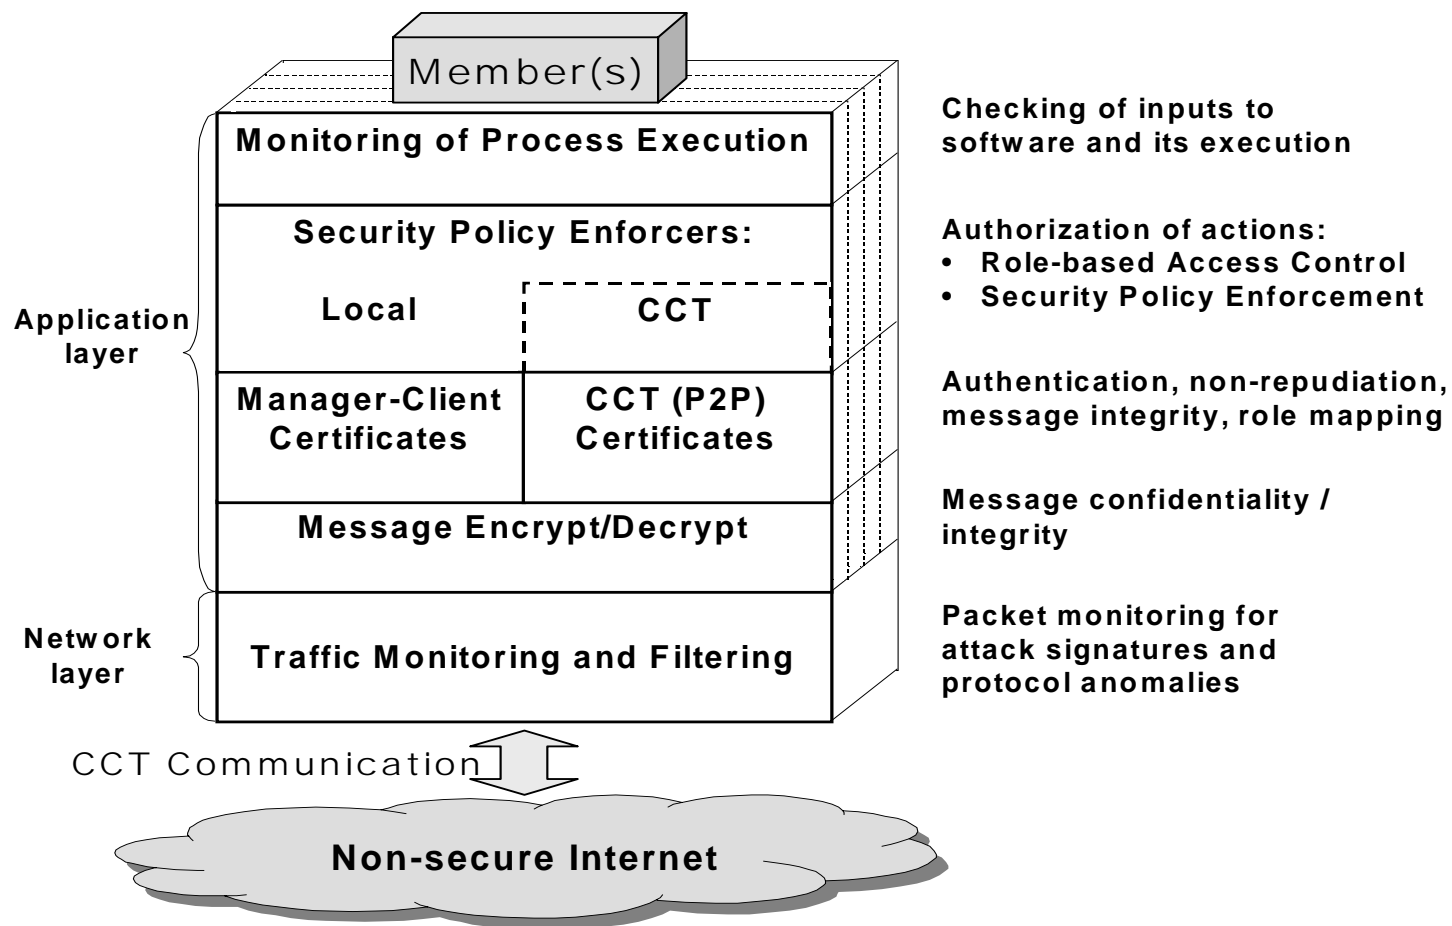
GRASP Architecture: Security

secure intra-/inter group communication

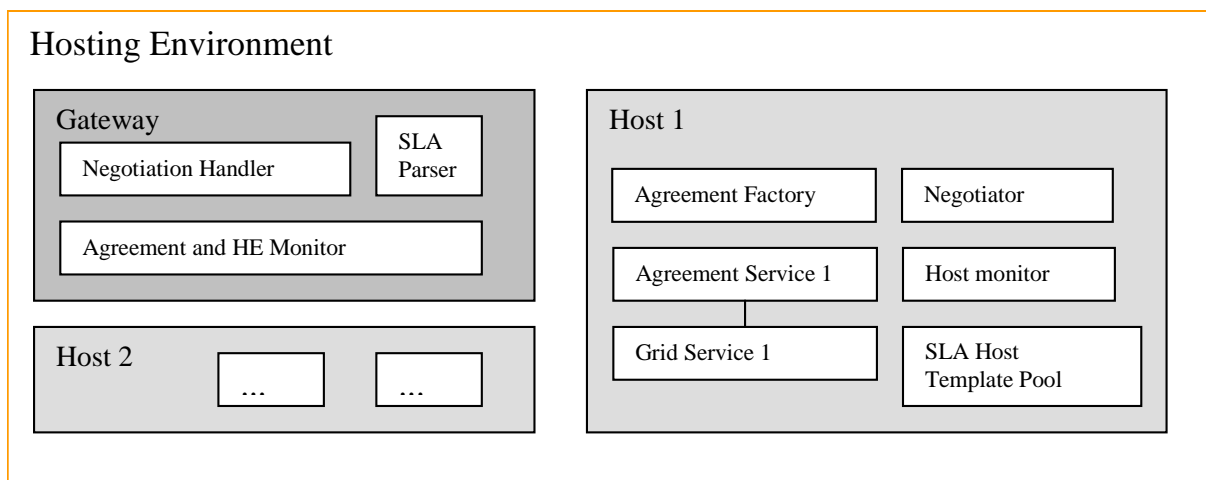


GRASP Architecture: Security

enforcing dynamic service security perimeters



GRASP Architecture: SLA monitoring



GRASP Architecture: SLA monitoring


Monitoring scheme options

Integrated with the Security Perimeter / Community Management Model

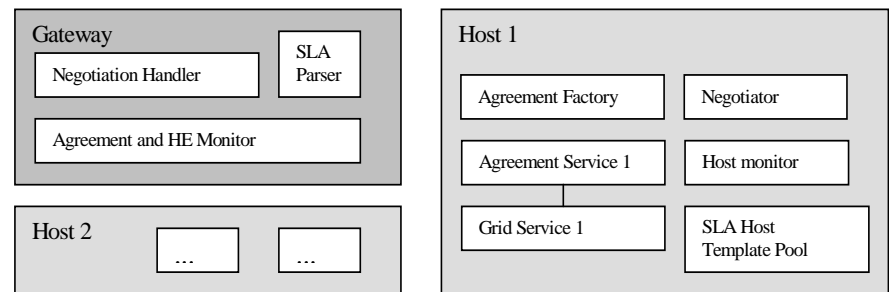
Compatible with OGSi-Agreement

Leverages BCA high-level concepts

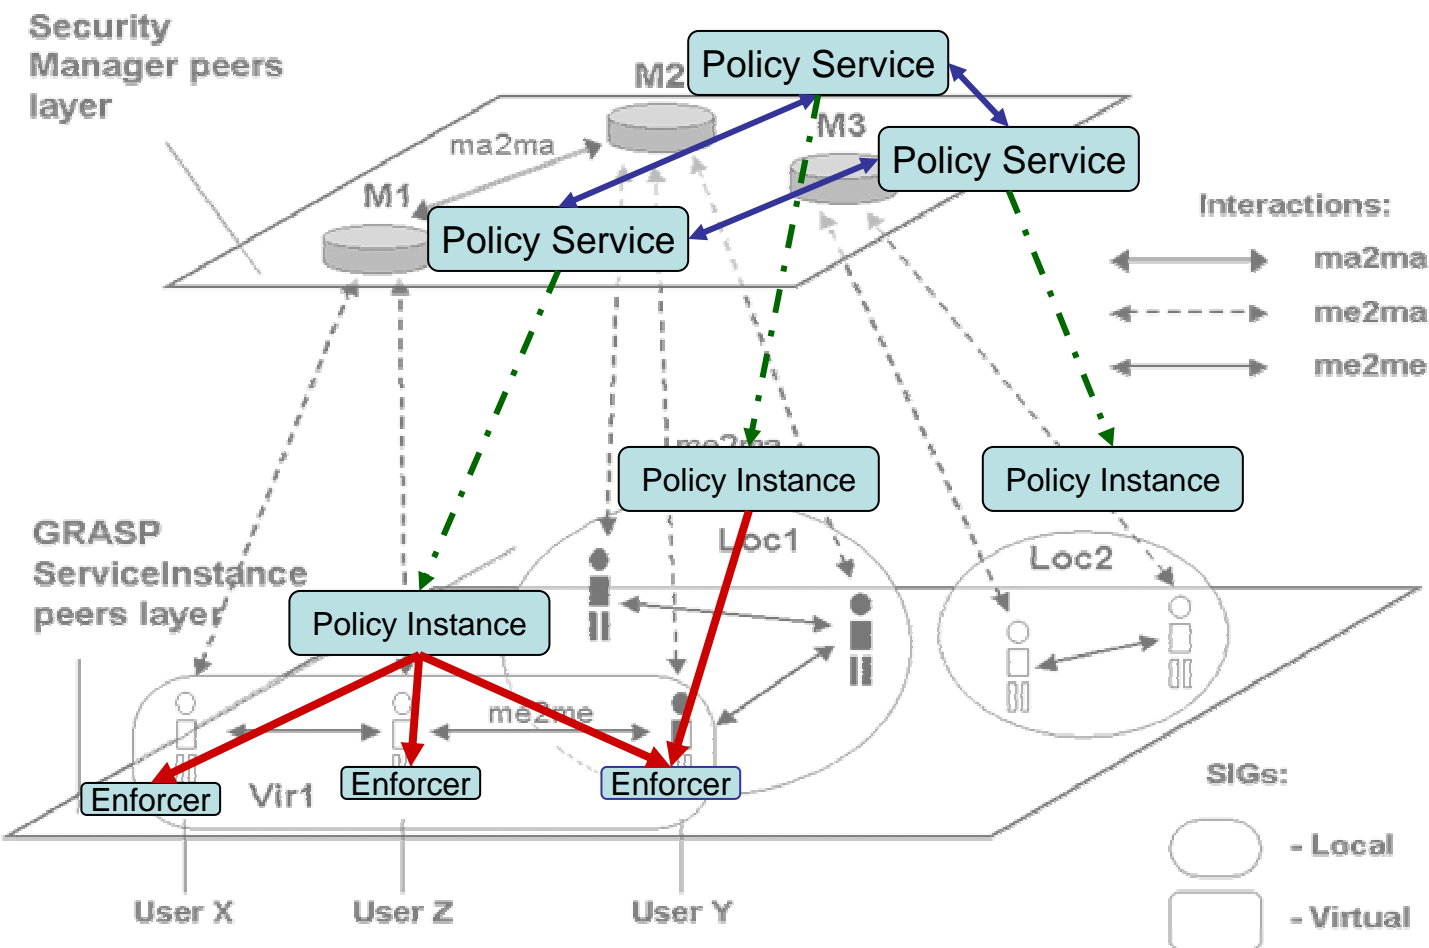
Leverages GeneSyS low level admin



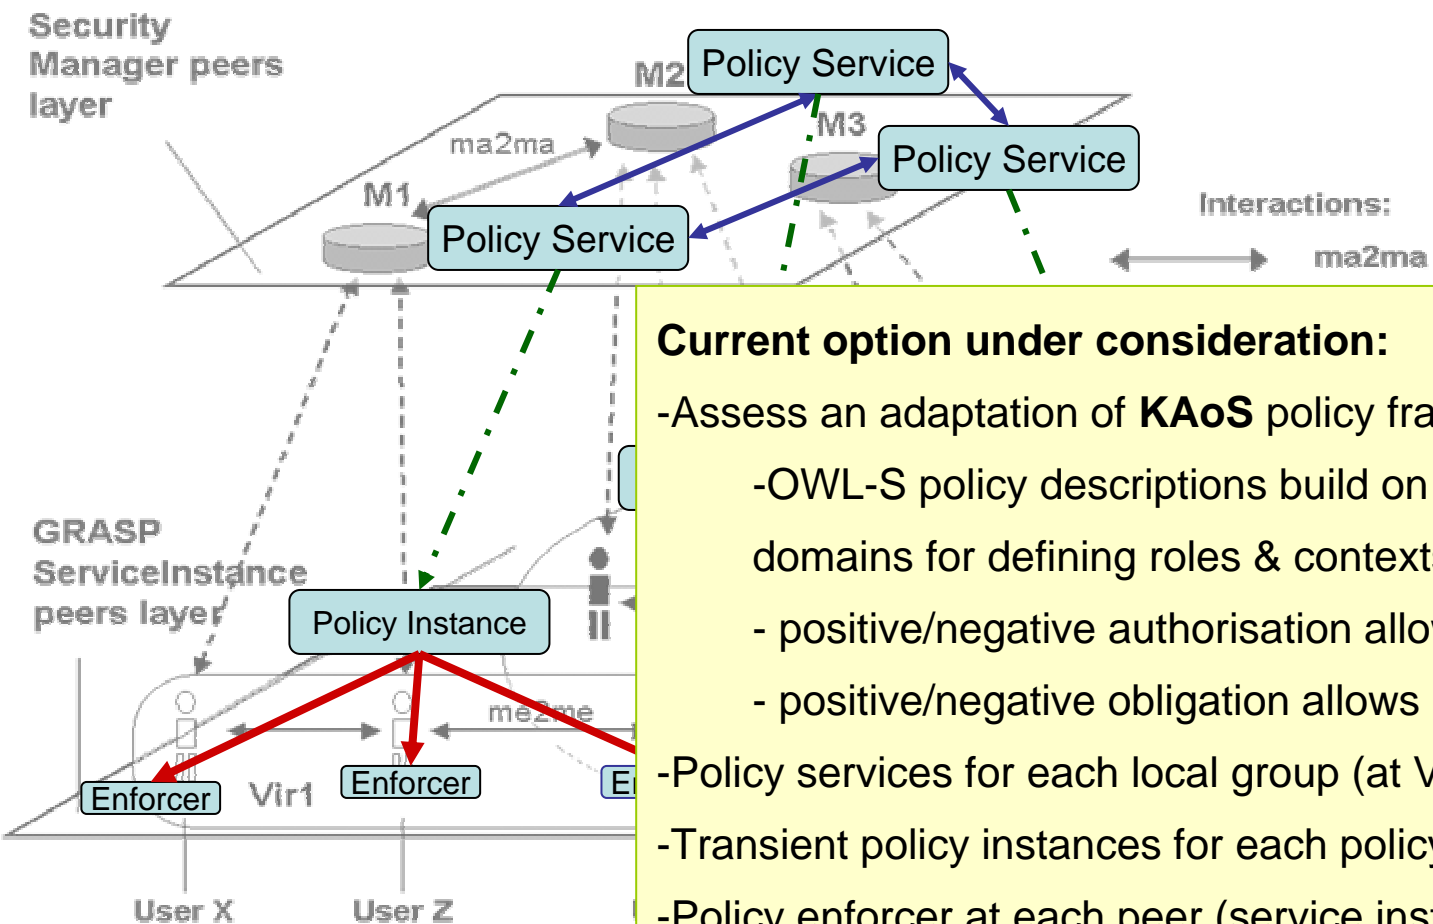
Hosting Environment



GRASP Architecture: Policy Management



GRASP Architecture: Policy Management



Current option under consideration:

- Assess an adaptation of **KAoS** policy framework
- OWL-S policy descriptions build on four basic policy types + domains for defining roles & contexts
- positive/negative authorisation allows controlling access
- positive/negative obligation allows enforcing SLA clauses
- Policy services for each local group (at VHE)
- Transient policy instances for each policy clause for each group
- Policy enforcer at each peer (service instance) in a group



Semantic Web technologies addressing the Trust Management problem

URI variable



- 1) If X is AC rep of Y, X can delegate W3C member access rights in Y.
- 2) *Kari* is AC rep of *Elisa*.



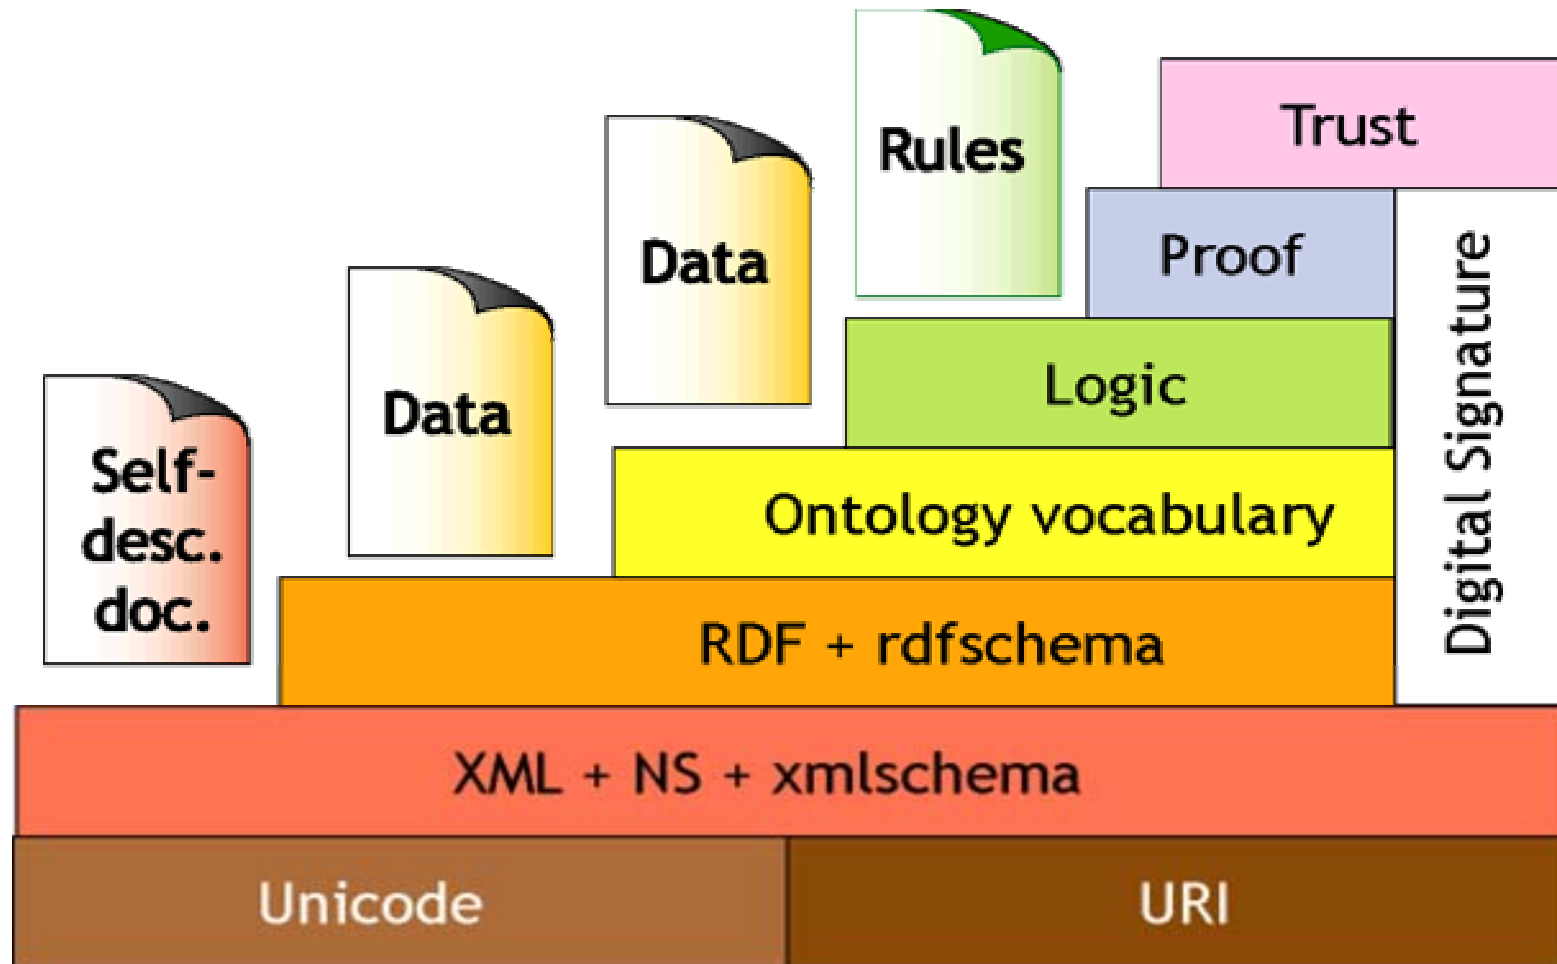
- 1) If X is employee of *Elisa*, X has W3C member access rights.
- 2) *Tiina* is employee of *Elisa*.



Tiina: I have W3C member access rights
Proof: Alan 1, Alan 2, Kari 1, Kari 2



Semantic Web Vision

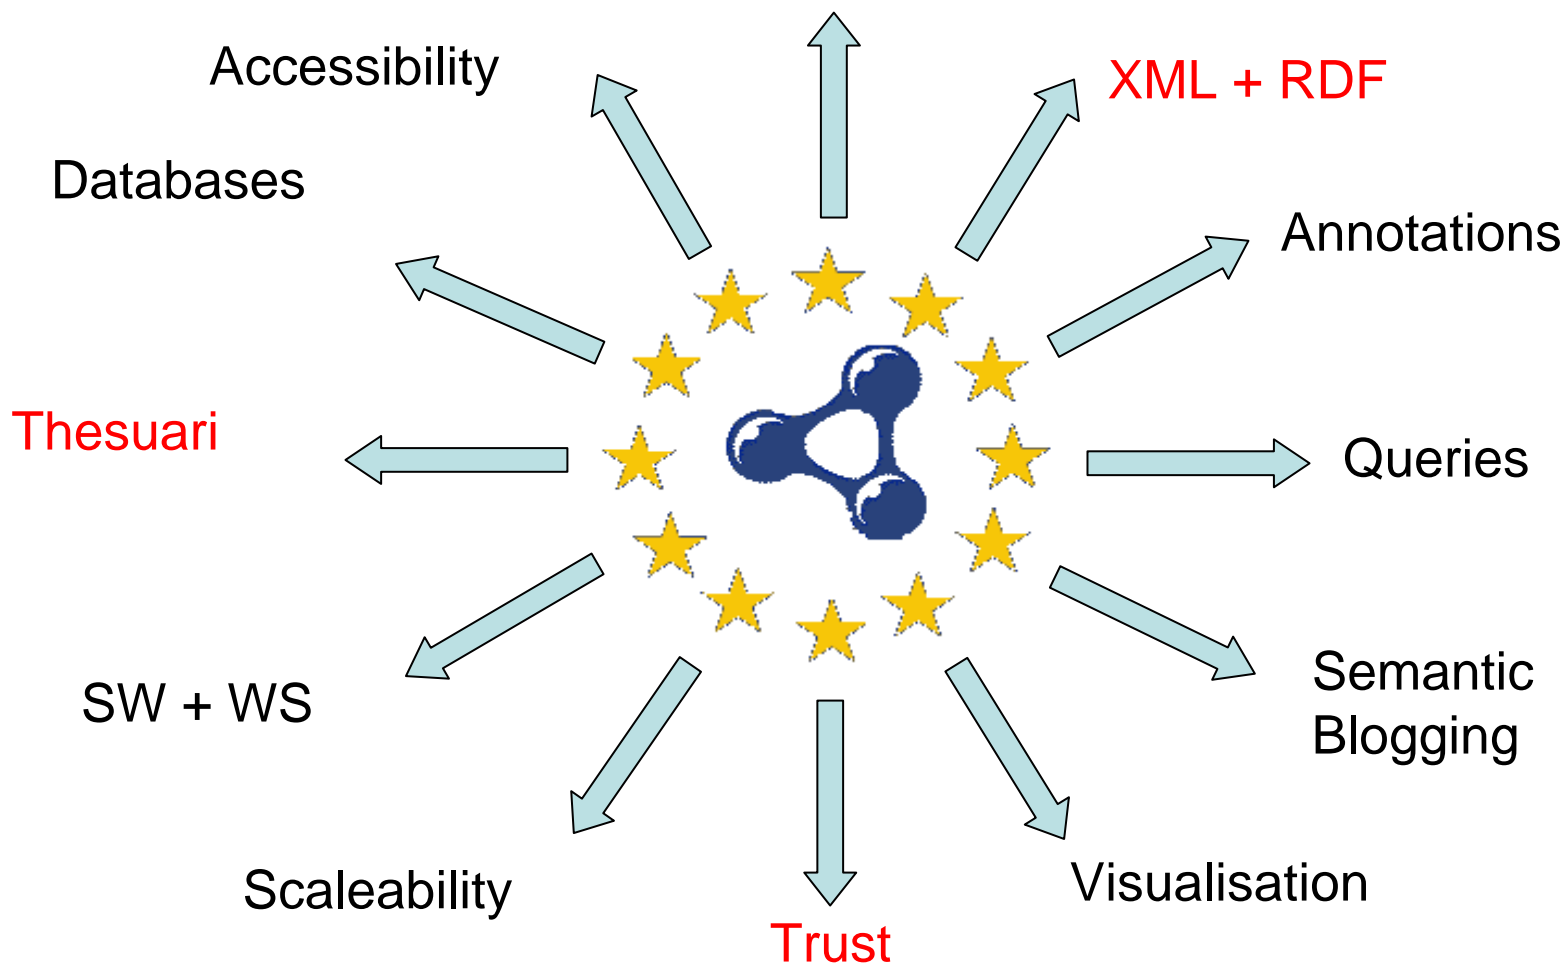


SWAD-Europe

Semantic Web Advanced Development in Europe

- Purpose is to encourage the use of Semantic Web tools and techniques now:
 - By an outreach programme
 - By developing practical demonstrators
 - By providing tools and standards
- Partners:
 - Univ. of Bristol,
 - W3C-INRIA,
 - CCLRC,
 - HP Labs,
 - Stilo

Overview of activities



CLRC in SWAD-Europe

- Three major areas
 - Developing XML Schemas from the Semantic Web
 - Developing tools and techniques for representing thesauri in the Semantic Web
 - Especially Multilingual Thesauri
 - Developing tools and techniques for representing and processing Trust relationships in the Semantic Web.

Pellucid overview

The Pellucid project (IST-2001-34519) is developing a customisable software platform for knowledge management systems to aid organisationally mobile employees.

It integrates several advanced information technologies, including autonomous cooperating agents; ontologies; workflow and process modelling; organisational memory; document indexing and metadata for accessing document repositories.

The Pellucid platform is agent-based and has three layers:

- the interaction layer, concerned with managing the interface with the employee (end-user) and the external world;
- the process layer, concerned with managing tasks and workflows;
- the access layer, concerned with search and retrieval of a wide range of documents.

Pellucid overview

Each of these layers comprises a collection of agents with defined competences and communications, acting together in a dynamic, flexible way.

An organisational memory will allow for monitoring of the overall behaviour of the system and a learning capability for continuous improvement.

The competences of the agent classes are as follows:

- **Personal assistant agents.** Responding to explicit requests for information; presentation of information both spontaneously and on request.
- **Role agents.** Monitoring performance of roles in work process; matching appropriate forms of advice to the user through the Personal assistant agent.
- **Task agents.** Instantiating particular forms of advice selected by the role according to the working context of the particular task. Working context encompasses both position in workflow and domain-specific attributes.
- **Information search and access agents.** Locating and retrieving information on request from diverse repositories.
- **Monitoring agents.** Monitoring users' passage through workflow and communicating between Pellucid system and workflow management or tracking system.

Pellucid overview

The aim of Pellucid is to provide **experience management**, disseminating the knowledge of more experienced employees to those who are less experienced, a situation that is increasingly common owing to **organisational mobility**.

The vessel for experience management in Pellucid is the **active hint**: a particular piece of advice presented spontaneously to the user and tailored to the working context. Active hints are constructed in a variety of ways based on templates appropriate for different situations.

There are three end-user organisations in the Pellucid project, with very different applications but all with experience management needs:

- **The Comune di Genova** (Italy), whose application is the process of evaluating, planning and executing the installation of traffic lights in the city.
- **SADESI** (Spain), a company that operates the call centre for the telephone network of the regional government in Andalusia—the application is the operation of the call centre itself, where high staff turnover means that experience management is a high priority.
- The **Mancomunidad de Municipios del Bajo Guadalquivir**, an association of local governments in the south of Spain, whose application is the process of management of projects and services.

CORAS Overview

- Eleven institutions from four European countries.
- Developed a tool-supported methodology for model-based risk analysis of security-critical systems.

The CORAS tool-supported methodology provides:

- A methodology for model-based risk assessment integrating aspects from partly complementary risk assessment methods and state-of-the-art ICT systems engineering
- A UML based specification language targeting security risk assessment.
- A library of reusable experience packages.
- A web-enabled software tool that supports the methodology and provides two repositories; an assessment repository and a repository for the reusable experience packages.
- An XML mark-up for exchange of risk assessment data.
- A vulnerability assessment report format.

Theo Dimitrakos
Business & IT Department

CORAS impact

2003-12: The CORAS UML profile for security assessment, submitted as part of the proposal OMG Document ad/2002-01-07, has now been adopted as an OMG standard by the OMG.



2003-09: The first release of the CORAS Risk Assessment Platform has been made available to the public as Open Source via SourceForge.net



CORAS relevance

CORAS process integrates a standardised Risk Management process with the OMG Unified Process (c.f. RUP), so as to fully incorporate risk analysis into the design & development of critical ICT systems.

CORAS platform supports the documentation, evolution and maintenance of risk analysis results and their correlation to system models during this process.

CORAS platform architecture is based on the ability to create, correlate and manage meta-data both about Risk Analysis and about System Designs

Meta-data is used ...

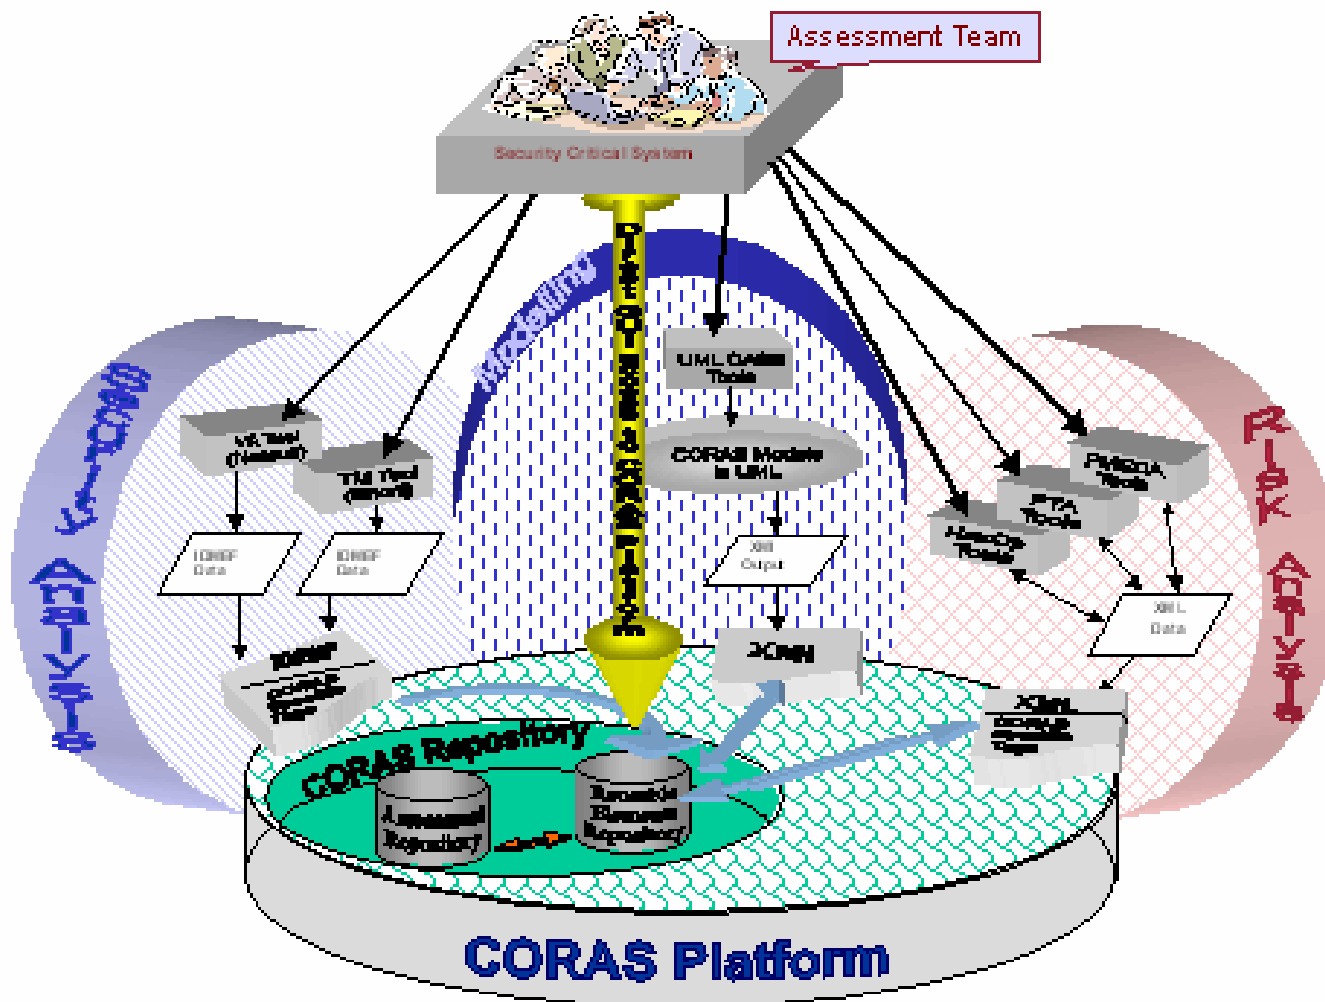
... as stored or exported “output” for the:

- (a) Internal representation and book-keeping of the Risk Analysis results
 - (b) Internal representation and book-keeping of the UML diagrams
- using tailored XML based notations**

... as stored or imported “input” for the

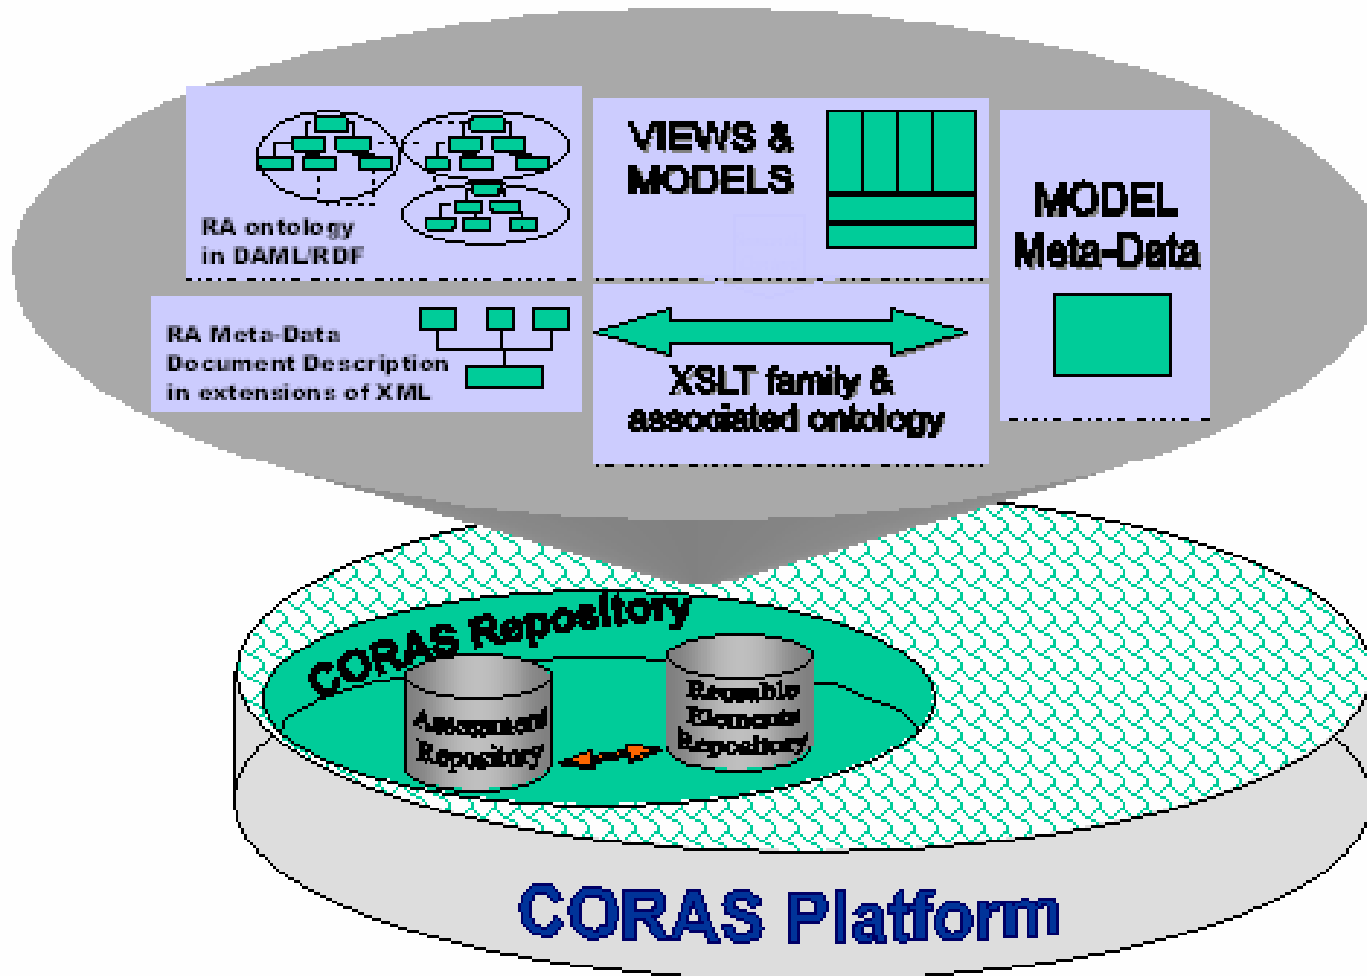
- (a) Visualisation of Risk Analysis results
 - (b) Visualisation of systems engineering diagrams based on UML
- using XSL-based technology**

CORAS architecture



CORAS meta-data management

(as we would have liked it to be...)



CORAS lessons

Experience with developing the current Open Source version of CORAS platform software indicates that:

- (a) **XML DTD** and **XML Schema** definitions are very useful for book-keeping Risk Analysis results, but **lack the semantic content** that would allow effective cross-referencing and manipulation of **RA meta-data** during the CORAS process
- Defining a core ontology for Risk Analysis data and RA-technique specific extensions in **RDF** or **OWL** (DAML+OIL) could provide a solution
- (b) **The above is particularly relevant for supporting solutions about**
- How to **transfer RA knowledge** from one technique to another (e.g. HaZOp to FTA to Markov Analysis) in relation to the same target system
 - How to **manage the correlation of RA results with** parts of the **system model** throughout design and development
 - How to **dynamically generate presentations of RA results** that are relevant to one specific concern and one specific view of the system.

Following the successful completion of the CORAS project we are interested in continuing the development of the CORAS platform & its architecture so as to appropriately address the above issues.

CCLRC contacts for more information

- GRASP: Theo Dimitrakos
t.dimitrakos@rl.ac.uk
- SWAD-Europe: Brian Matthews
b.m.matthews@rl.ac.uk
- Pellucid: Simon Lambert s.c.lambert@rl.ac.uk
- CORAS: Theo Dimitrakos
t.dimitrakos@rl.ac.uk

What does the future hold?

- **TRUSTCOM:** trust/security & contract management framework for dynamic Virtual Organisations
 - Theo Dimitrakos t.dimitrakos@rl.ac.uk
 - Michael Wilson m.d.wilson@rl.ac.uk
- **E-LeGI:** elements of a European Learning Grid Infrastructure, focusing on experiential learning applications:
 - Damian Mac Randal d.f.mac.randal@rl.ac.uk
 - Theo Dimitrakos t.dimitrakos@rl.ac.uk
- **Integration of Grid middleware and Pervasive / Nomadic Computing over heterogeneous networks (emphasis on mobility)**

Get involved: forthcoming events

Spring 2004



CCLRC Rutherford Appleton Laboratory
CCLRC
Second International Conference On
TRUST MANAGEMENT
29 March - 1 April 2004
St Anne's College, Oxford UK

An annual event of



Working group on
Trust Management in
Dynamic Open Systems

www.itrust.uoc.gr

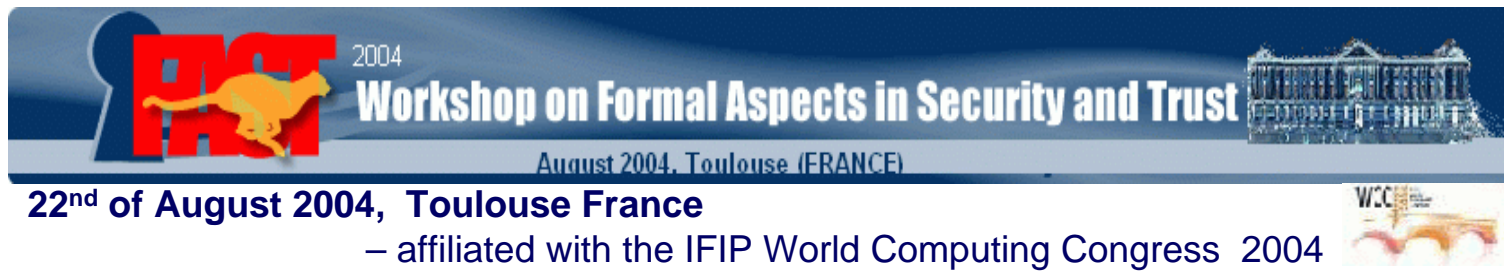
www.trustmanagement.clrc.ac.uk

Supported by



www.w3c.rl.ac.uk

Summer 2004



FAST 2004
Workshop on Formal Aspects in Security and Trust
August 2004, Toulouse (FRANCE)
22nd of August 2004, Toulouse France
– affiliated with the IFIP World Computing Congress 2004

Autumn 2004

Towards a European Learning Grid Infrastructure

1st International Conference
5th LeGE-WG workshop

14-16 September 2004
St Anne's College, Oxford UK

Learning Grid of Excellence
Working Group

