OGSA/GT3 evaluation

Activity Report
Massimo Lamanna
LCG/GTA

Table of Content

- Introduction
- Activity highlights
  - GT3 ToolKit Experience
  - GT3 Performance studies
  - Integration of Existing Codes/Services
  - Relationships building

Motivation

- The promise of the web services framework
  - New projects are looking to OGSA as the solution to all problems…
  - Input to the strategic planning and architecture activity: 6+ months!
- Here and now
  - Globus α release of the new toolkit in May 2003.
    - OGSI framework and some grid services
    - GT3 out July the 1st (new major release -3.2- 1H04)
  - To provide input to the EGEE middleware activity
- Initial objectives
  - Project approach
    - Clear project scope, goals
    - Limited time span (~ 6 month period)
    - Pragmatic approach (learn by doing)
  - Primary objectives of the OGSA/GT3 evaluation:
    - **Understand** the GT3 offering and its "quality"
    - **Learn** how to create new services in this framework.
    - **Study** how to leverage existing developments in an OGSA context
    - **Create** local know-how on promising technologies

OGSA Engineering Group

- Proposed to the LCG referees (May 2003) by D. Foster
- Started end June 2003
  - M. Lamanna – Overall Coordination (CERN)
  - R. Brito Da Rocha – Service Development (EDG)
  - A. Demichev – Setup (MSU)
  - V. Kalyaev – Service Development (MSU-CERN Summer Student)
  - A. Kryukov – Service Development (MSU)
  - V. Pose – Performance and Testing (JINR Dubna)
  - Tao-Sheng Chen – AliEn (Academia Sinica Taipei)
  - C. Wang – AliEn (Academica Sinica Taipei)
- Most people at CERN only for short periods
  - Variable geometry approach
  - 75% of the people are not always at CERN
  - Open to new collaborators

What does GT3 offer? (NOW)

- The first OGSI implementation (July 2003: 3.0.x)
  - The toolkit itself
    - Build new services and extend existing ones
  - Security Infrastructure
    - GSI (Globus Security Infrastructure)
  - Services
    - GRAM (GT2 implementation wrapped up as a Grid service)
    - IS (Index Service; new GT3 implementation)
    - RFT (Reliable File Transfer; it uses Globus FTP)
    - *RLS (GT2 implementation as a Grid service) , …*
- Explore these **three** lines

TestBeds

- First hand experience on Globus Toolkit 3
  - This can be achieved only by using it!
- The main tool are prototypes, with the following common features:
  - Small
  - Working (with limited functionality)
  - No architectural ambition
  - Engineering approach
    - Mapping of functionality – prototype functions
- GT3 TestBed
  - 4 CERN machines + 1 in Moscow
  - Focus on GT3 basic functionality and performances
    - Performance tests use also some high performance machines and lxplus
- AliEn TestBed
  - 3 CERN TestBed machines
- (ARDA) TestBed
  - Focus on the complexity of future possible architectures
  - Deployment use cases

Example: GT3 Test Bed

ToolKit

- Resource broker and L&B (*Custom service*)
  - Surprisingly fast to set-up
- *A few computing elements (GT3-GRAM, with modifications)*
  - 2 PC boxes in the CERN Computing Centre
  - In a second phase, one PC located in Moscow was added
  - Some problems (solved) in data stage-in/stage-out
  - See GRAM comments in the performance part
- Information service (*GT3-IS*)
  - Native GT3 service
  - In this TestBed talks only with other services

Security

GT3 Services

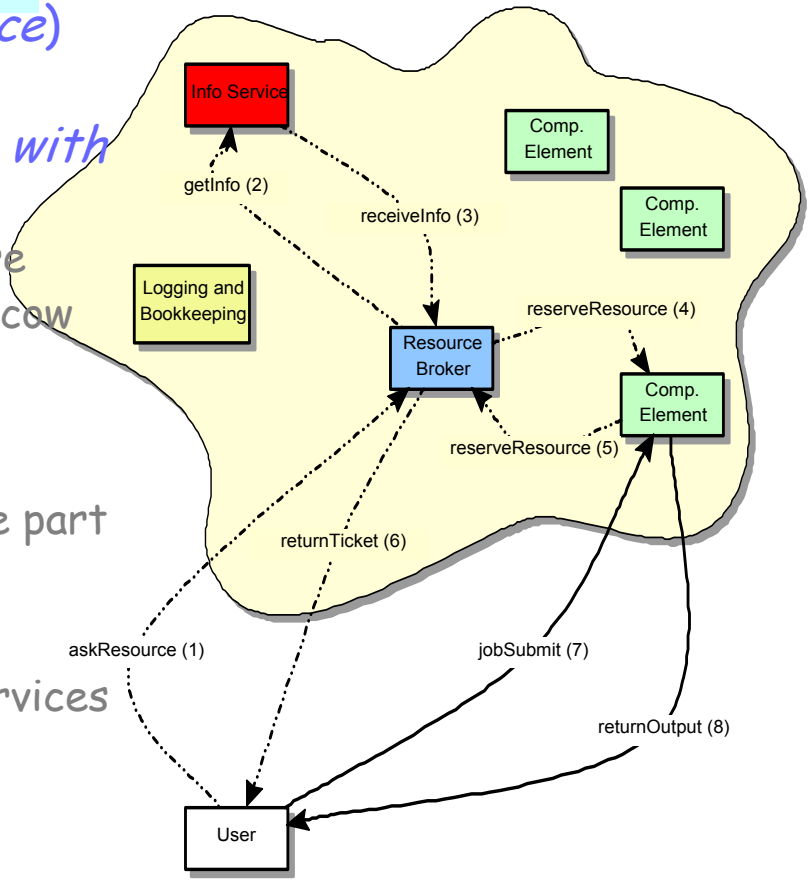

GT3 TestBed coverage…

PortType	Operation	*OSGA/GT3 evaluation*
Gridservice *Every service must implement this PortType*	FindServiceData	•GT3TestBed-RB uses it to retrieve data from IS •IS performance tests (C-client)
	SetTerminationTime	•Not Used Yet (directly)
	Destroy	•Everywhere, e.g. GRAM
NotificationSource	SubscribeToNotificationTopic	•IS perf. Tests (data sources)
NotificationSink	DeliveryNotification	•IS perf. Tests (listener)
Registry	RegisterService	•Code examples
	UnRegisterService	•Code examples
Factory	CreateService	•Via GRAM (first tests) •Specific tests using DummyService
HandleMap	FindByHandle	•Not Used Yet

Modelling activity of this type of service starting

Prototypes developed within the project

- Globus 3 "components" tests
  - GRAM tests
  - Index Service tests
  - Reliable File Transfer tests
  - GSI (Security) tests
- Implementation of deployment use cases
  - Remote installation (via dedicated custom services)
  - Remote management of different version of a service
- Performance Prototypes
  - Dummy Service
  - Dummy Secure Service
  - Dummy Service with Service Data
  - Dummy Service with Notifications
  - Dummy Service + Index Service
  - Index Listener
- Higher Level Prototyping
  - File Catalog Service
  - Metadata Catalog Service
  - Storage Element Service
  - Workload Management Service
  - Computing Element
  - Authentication and Authorization

Globus Toolkit 3 Overview

- The GT 3 is the first complete implementation of the OGSI specification
  - The development process is much easier when compared with GT2.
    - Steep learning curve should be taken into account!
    - New approach to service design and implementation
  - Deployment Tools (not complete)
- Backward compatibility:
  - All GT2 components are shipped with the GT3 full bundle
  - Others are completely independent implementations (eg. MDS2 and MDS3)
  - Could be a problem for GT2 based project (LCG)
- A large user community is being built
- Incomplete documentation
  - Getting better now (tutorials, etc...)
- Several bugs found in these exercises
    - Core implementation related - due to framework short lifetime
    - From tools deployed with the framework – hard to solve (e.g. Axis)
    - From the outside – easy to solve (e.g. Tomcat)
- GT2 GRAM – with an OGSI-compliant but complex architecture behind
  - Worry to lose past experience (gained within the EDG and LCG projects)
  - Confirmed by performance tests (see next slides)

GT3 performance measurements (highlights)

- Goal:
  - explore GT3 under heavy load/concurrency:
    - maximal throughput/rate of GT3 services
    - see the limiting factors
- Highlights from:
  - GRAM
  - Security
  - IndexService
  - RFT

GT3 GRAM performance

- Results: service node
  - Saturation throughput for job submission on the service node: 3.8 jobs/minute with an average CPU user+system usage of 62%

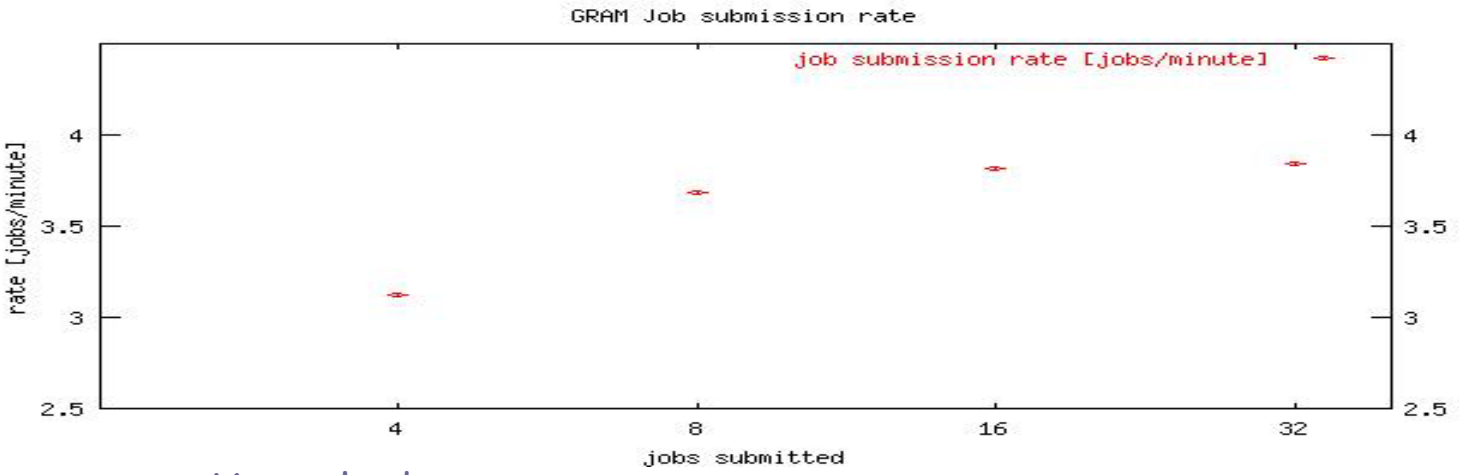

  - Very slow!
  - Scalability issues for (heavily used) servers

DummyService performance

setup	authen-tication	service container	saturation throughput	average CPU u+s usage, %
Push mode	no	GT3 standalone	41 services/s	89
	yes	GT3 standalone	1.3 services/s	88
	no	Tomcat	60 services/s	89
	yes	Tomcat	1.2 services/s	88
Pull mode	no	GT3 standalone	300 method calls/s	96
	yes	GT3 standalone	10 method calls/s	72
	no	Tomcat	290 method calls/s	96
	yes	Tomcat	13 method calls/s	79

- The security overhead needs further investigation (OGSA/GT3 group and Globus)

Index Service (Information service building block) in push mode

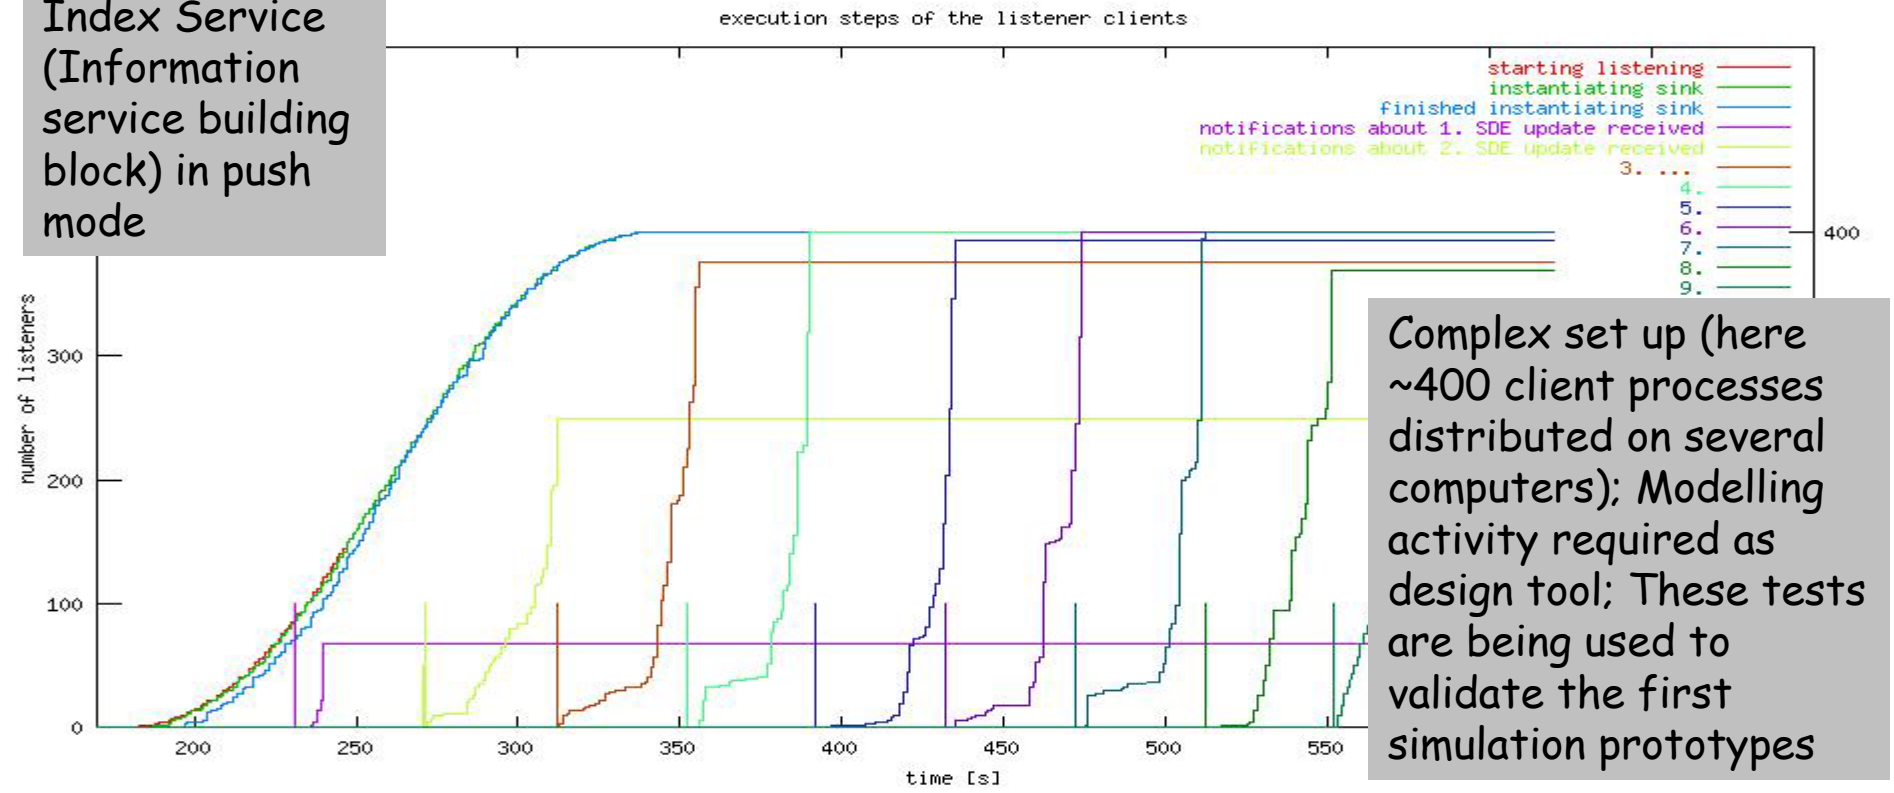

Complex set up (here ~400 client processes distributed on several computers); Modelling activity required as design tool; These tests are being used to validate the first simulation prototypes

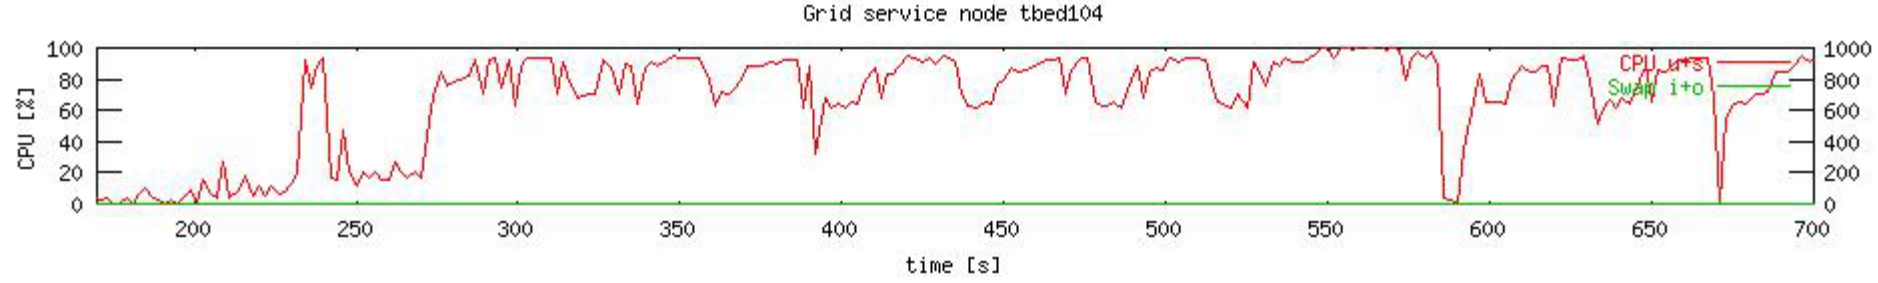

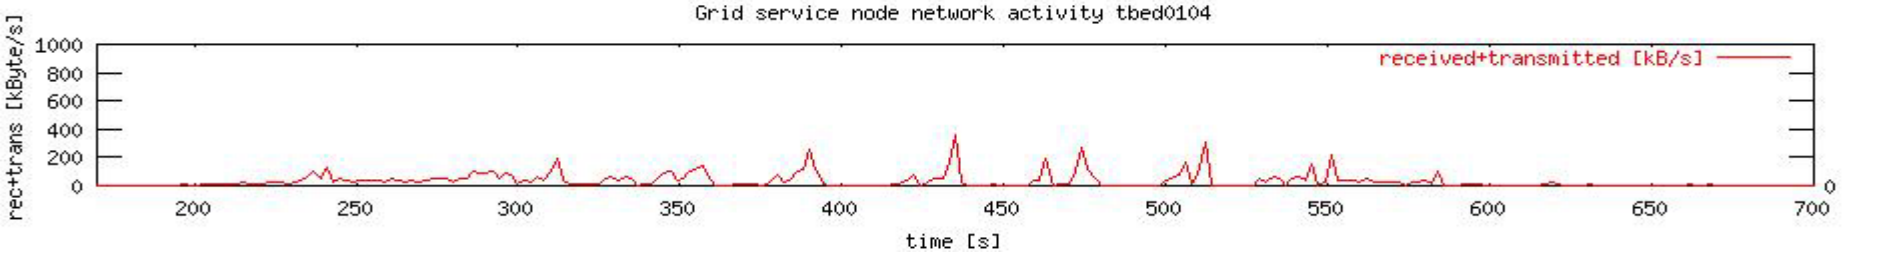

Reliable File Transfer Service

- Emphasis on reliability. Solve problems like e.g.
  - dropped connections,
  - machine reboots,
  - temporary network outages, etc
- Functionality: OK
- Main problem: resource hog
- Comprehensive report submit to Globus
- Fix found by the GT3 team
  - We agreed to test it in detail
- Open chapters
  - gridFTP performances (the RFT "engine")
    - WU-FTP and the new globusFTP

Integration

- GRID mainly concerns about the interoperability among heterogeneous grid components
- Heterogeneous Grid environments
  - AliEn (Alice Environment)
    - Should provide first-hand experience within LCG
- Heterogeneous Grid technologies (non GT3)
  - OGSI .NET, MS .NETGrid (.NET environment)
  - *Unicore, others…*
  - Discussion with some teams at GGF9; to be restarted end of November (after SC2003)
- Necessary to validate GT3 itself!

Bridging between past and future activities

- Assessment of new technology (OGSA and GT3 as a concrete example)
  - EGEE preparation
  - Do not lose the "past" experience (EDG developers, LCG deployment)
  - Clear suggestion from the e-Science Gap Analysis paper
    - "Will GT3 be robust enough to deploy?"
    - Summary: ES-2003-04
  - Coherence between development plans and LCG needs
  - Ensure a dependable fast feedback line
- How to build these relationships?
  - Globus Team (as an example)
  - Other foundation components providers (IBM example)?

Do not lose the "past" experience

- Assessment of OGSA/GT3
  - Strategy defined in coordination with embryonic EGEE teams (last August)
  - EDG
    - Major issue so far: GRAM
  - LCG deployment
    - Major issues so far: GRAM, Information Services and configuration issues
  - eScience gap analysis (Geoffry Fox report)
    - Used to inform the original evaluation plan
  - VDT
    - CondorG/GT3 will be demonstrated at SC2003 (this week). Agreement to use it in our test

Relationship building: Globus Toolkit 3

- Little formalities, working relationship
- Contacts before first results being shown and discussed at CERN
  - Initiated by D. Foster
  - Notably with Ian Foster
- Multiple meetings during GGF
  - Very encouraging
- Regular meetings between Massimo and Lisa Childers since then
  - Lisa is the Technical Product Manager for the Globus Toolkit
- The "preview" page (see GTA Internal Review page)
  - to discuss results with GT3 during the finalization phase
- Discussions list set up
  - Set up by Globus explicitly to allow this kind of quick and open communications between us and them

Relationship building: Globus Toolkit 3

- Common agreed Action List maintained by Massimo and Lisa
  - On the GTA preview page
- Status of the interactions:
  - Access to unreleased software; agreed mechanisms to discuss and give feedback
  - Job Gatekeeper (GRAM)
    - Feedback
    - More priority on performances inside the GT3 team since
  - Reliable File Transfer (RFT)
    - Issues (high CPU consumption) confirmed. Fix available
    - Access to the experimental trunk for verification
  - Index Server (IS)
    - Issues being discussed
  - Security (GSI)
    - Issues being discussed

GTA and GT3-IBM

- The hosting environment plays a role
  - Standalone vs Tomcat
  - Other hosting enviroments (e.g. .NET)
- Some tests already performed (see GTA pages)
- CERN/IBM project to evaluate GT3-WebSphere
  - Another hosting environment
  - Interesting complementary information
  - After a formal preparation stage, this small sub project has started last week
  - Better understanding of the industry commitment in Grid
    - IBM ship a modified GT3 version from IBM