Job Provenance Deployment

Jiří Sitera

CESNET

JRA1 All-Hands, 21.2.2008, Amsterdam

Agenda

- Job Provenance architecture
- Deployment goals
- Deployment strategy
- Involved components and people

Job Provenance

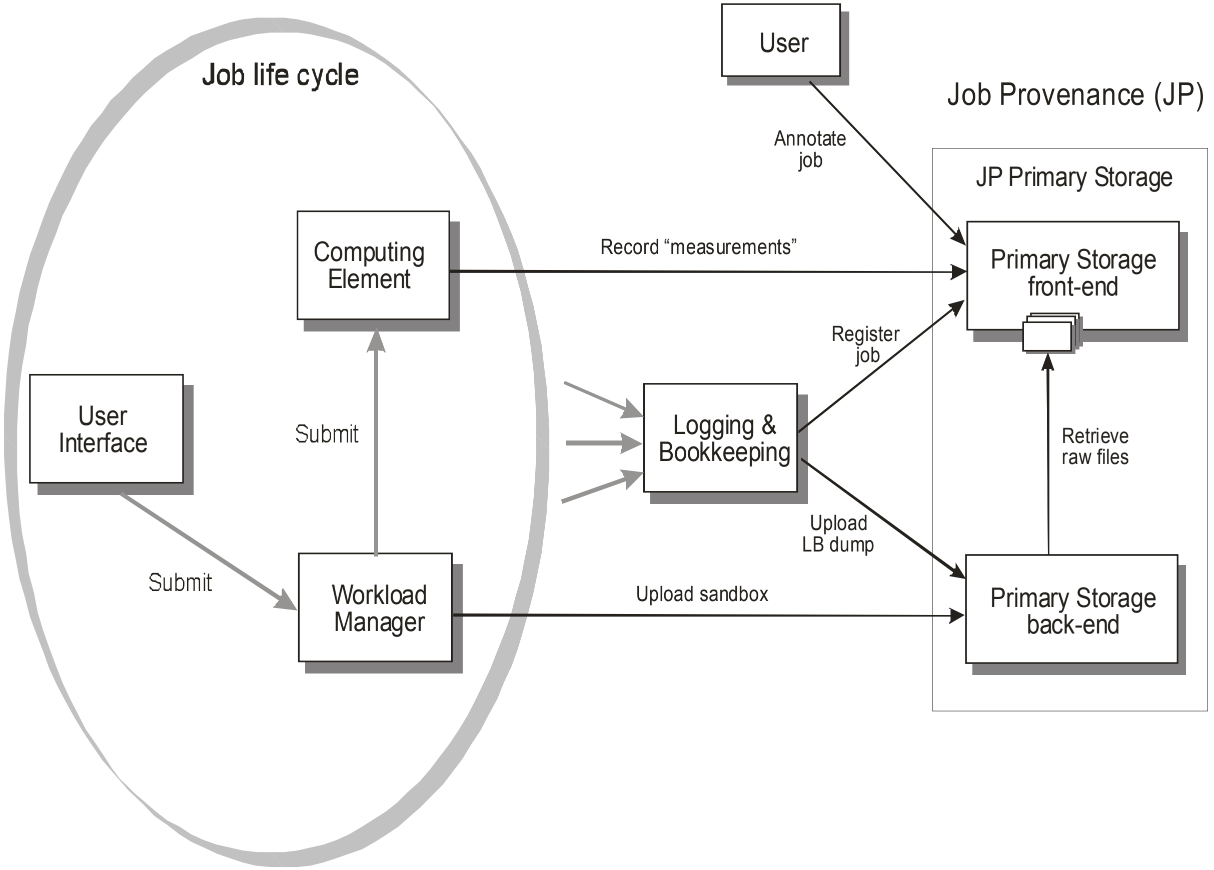

- **Job related data**
  - System job lifecycle track (LB)
  - User (annotation)
- **Long-term storage**
- **Efficient search**
- **Generic engine for customized tools**

- **Current state**
  - Working on JRA1 Preview testbed
  - Finished demonstration and evaluation projects

Job Provenance Architecture

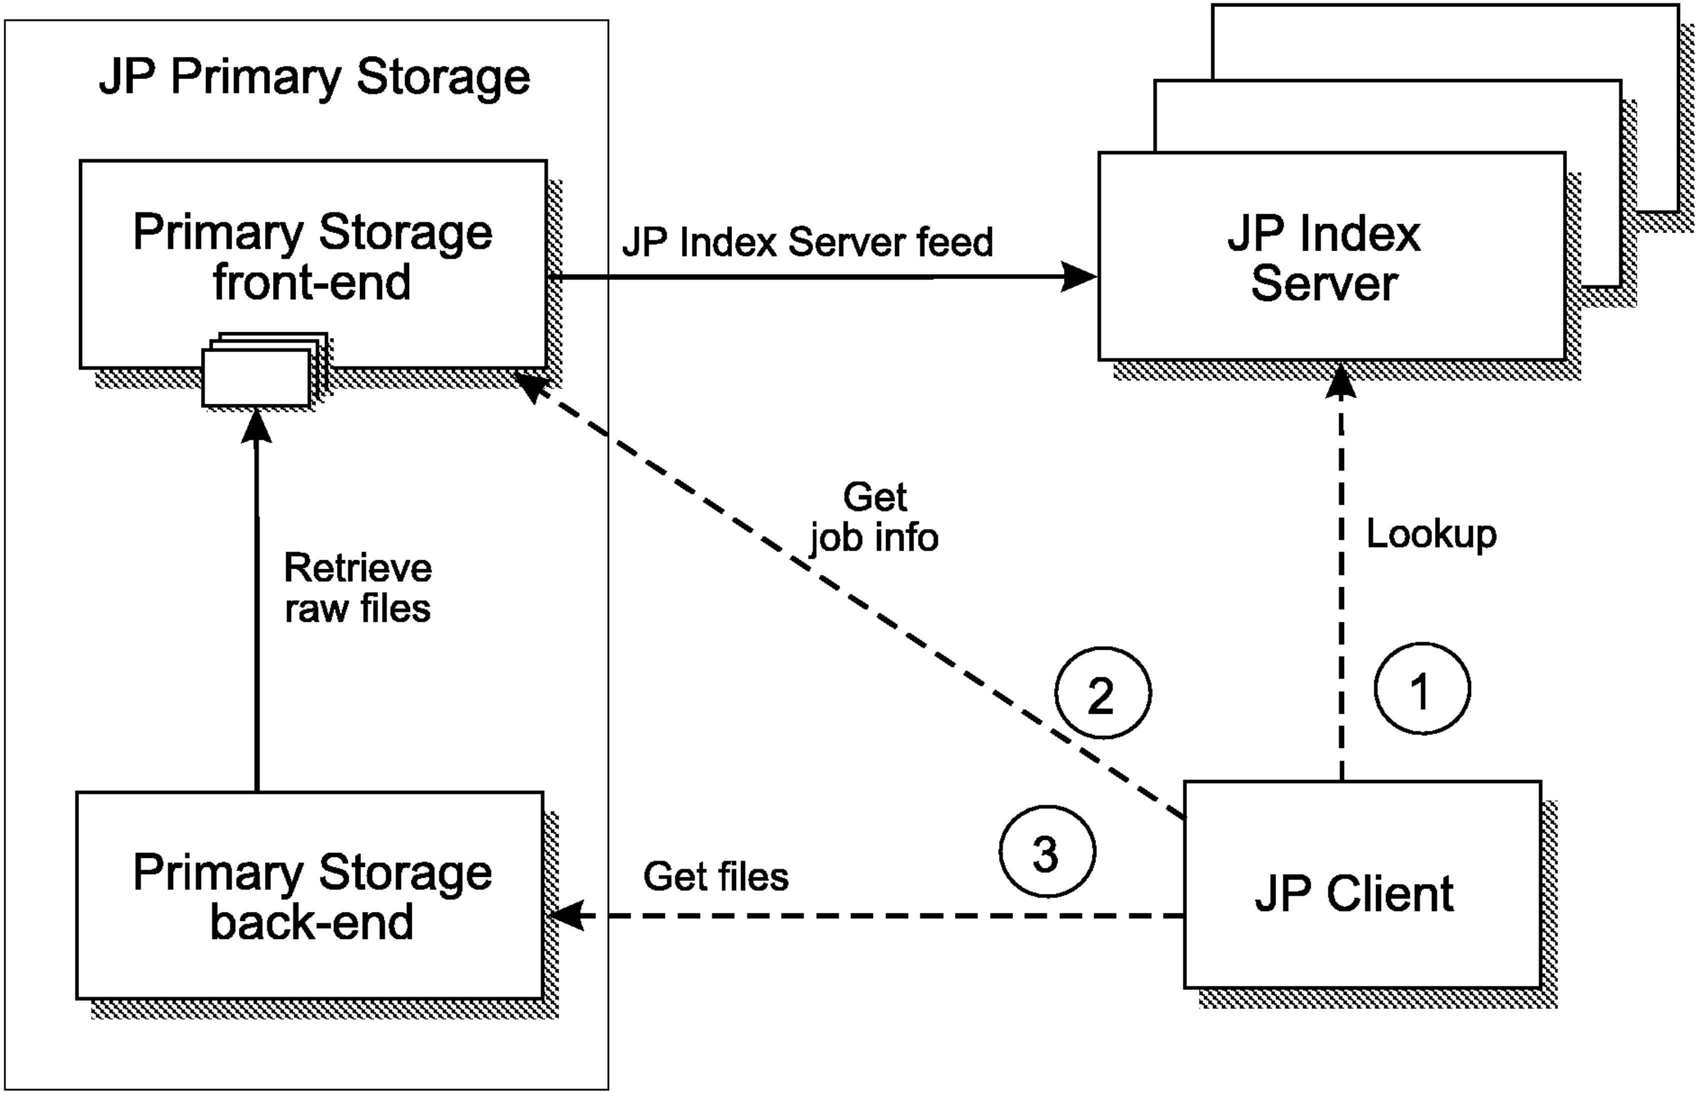

Goals

- **Infrastructure integration, availability for users and experiments**
  - One JP Primary Storage per large VO or region
  - A set of "system" JP Index Servers (prescribed configuration)
    - Support basic usecase "get info about job"
    - System JP Index Servers designed to work in tandem with LBs for this queries
    - Not only for the first phase of deployment, should form a stable minimal set of JP Index Servers to be maintained
  - A package to setup purpose specific JP Index Servers

Enabling Grids for E-sciencE

Job Tracking Data Lifecycle

- ## Without Job Provenance

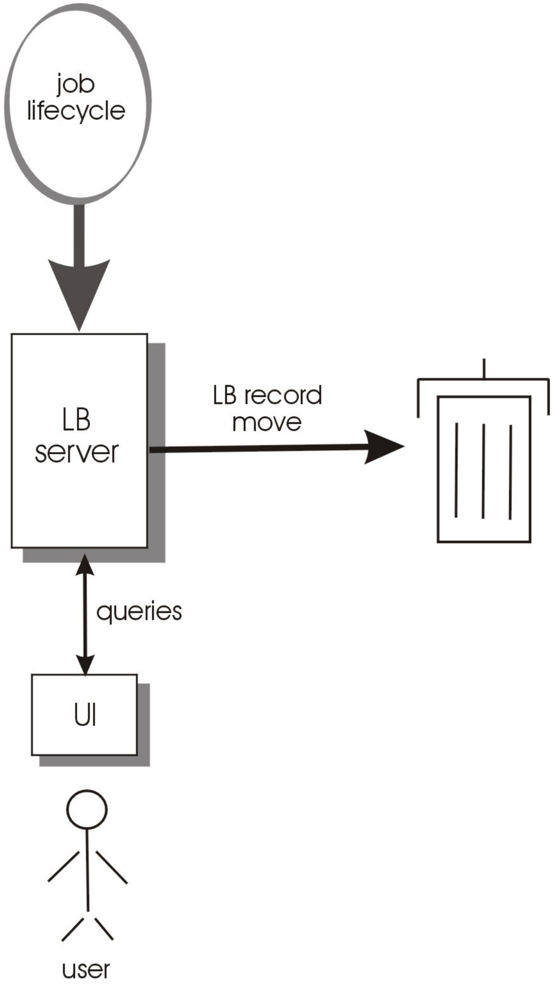

- ## With Job Provenance

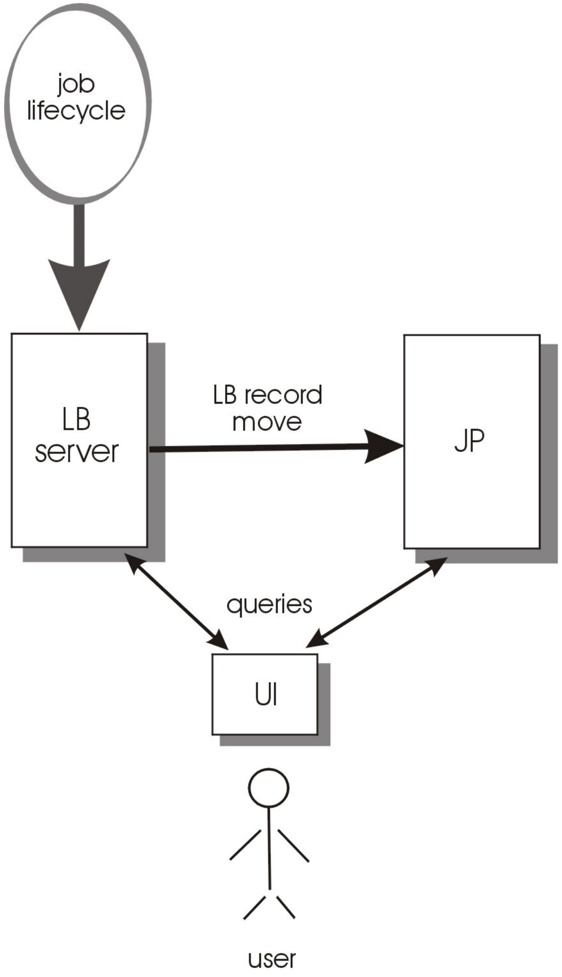

When Migrate Data To JP?

- **Days after job reaches its terminal state**
  - Designed behavior (when JP deployed), but:
    - Many LB users will miss their jobs
    - If you are not aware about JP, you cannot find anything based on JobId
- **Many months after job termination**
  - JP searches not functional – show stopper for end-user customized tools
  - LB database still can remain too big
- **Solution: use the first option but enhance it**

Zombies in LB!

- **Each job stays in LB even after purge (move to JP)**
  - In a form of a zombie
  - Easy work for LB, important hint for users
- **Zombie job**
  - Very compact record
  - *"Such job (JobId) existed; its track moved into JP server."*
  - Actual address of the JP server not stored here
- **How to find appropriate JP server?**
  - Configuration parameter of UI
  - Service discovery can be used here

Get Job Info – System JP-IS

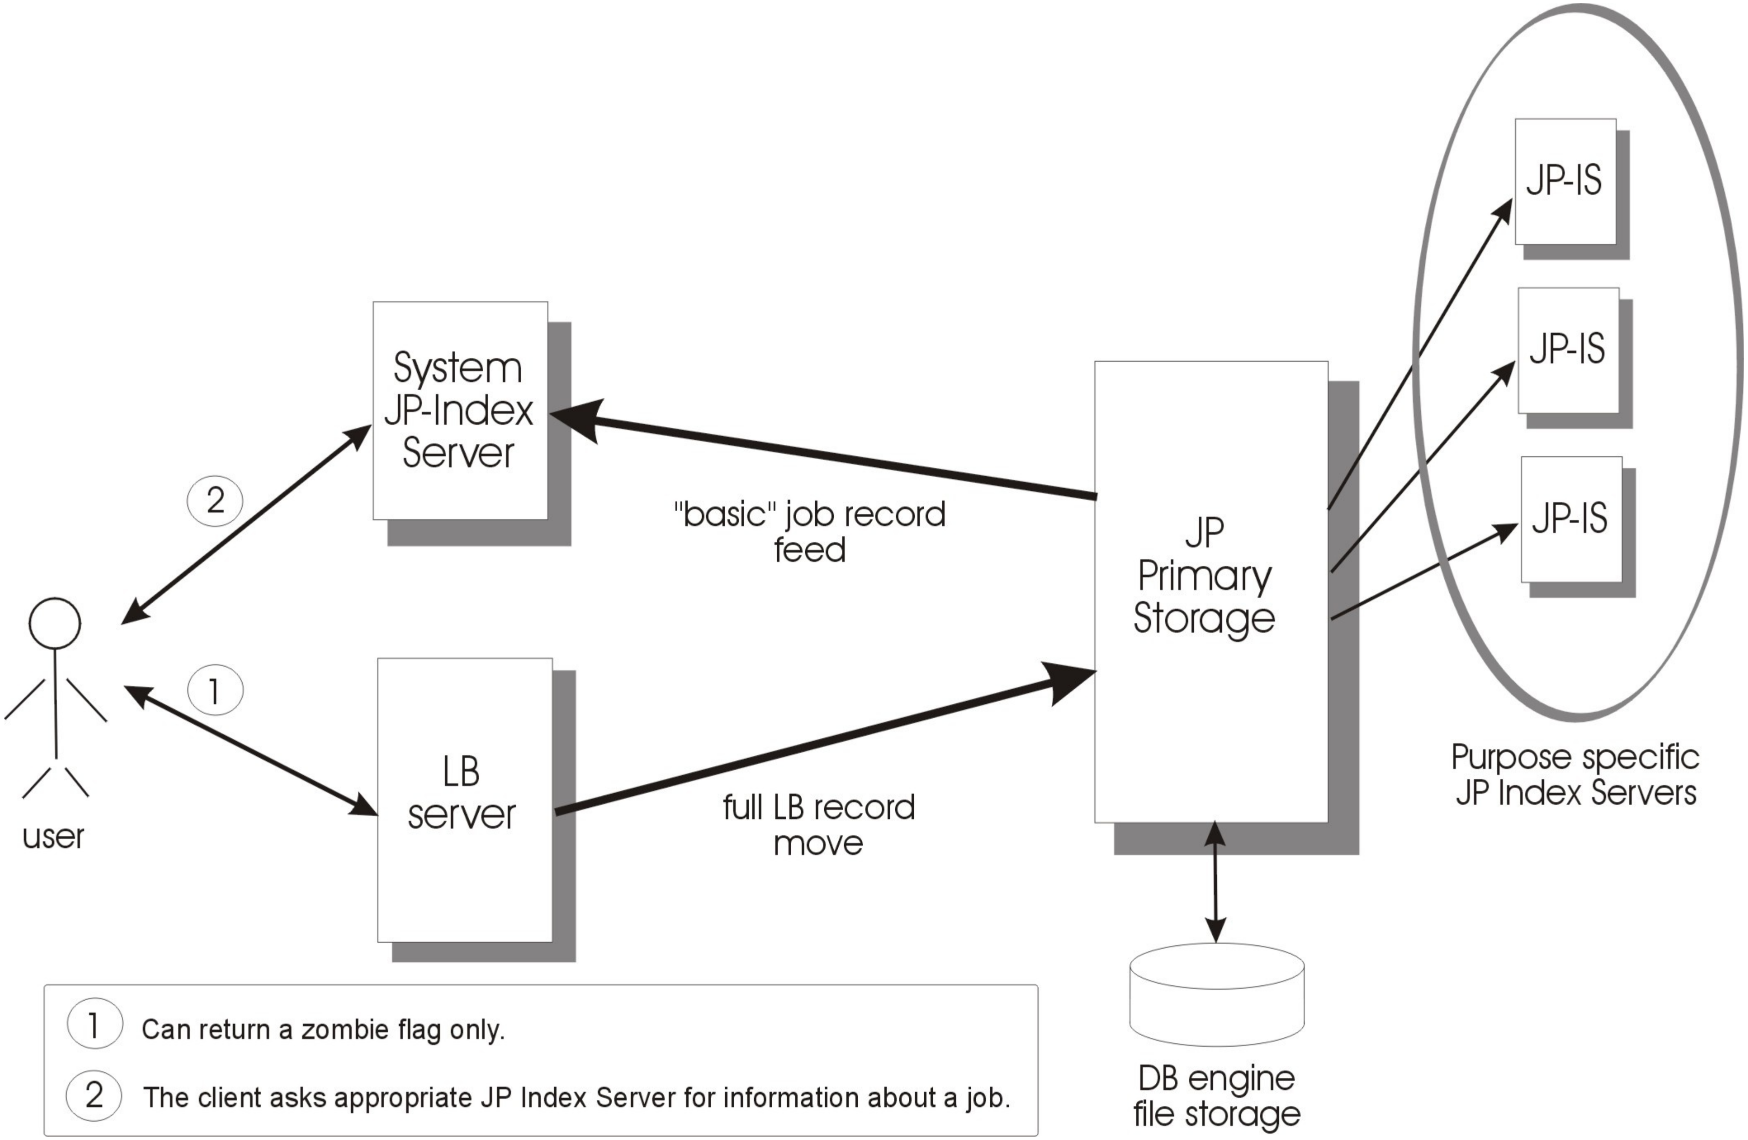

1 Can return a zombie flag only.

2 The client asks appropriate JP Index Server for information about a job.

New Components To Be Deployed

- **Primary Storage**
  - Released in RESPECT, standard gLite deployment & config (YAIM)
  - Standalone server without preconfigured knowledge about other components
  - Needs database & file storage with backups, etc.
    - Majority of data stored as files (LB dump), database contains just one simple record per job
- **Index Server**
  - JP-IS package provided with default configuration (system JPIS)
    - Static configuration (condition ”jobs from LB server”, basic attributes, feeding JP location)
  - The package can also be used to install purpose specific Index Servers (custom configuration)
  - We provide example JP-IS package (example configuration)
    - Index all jobs for a given VO (covering a set of LBs)

Involved Components

- **LB server**
  - JP feeder is already part of LB server
  - The “right” LB database purge procedure
  - Goal: LB with purger & JP feeder certified gLite component
  - Config: JP server address, purge timeouts
- **UI**
  - Job state queries
    - Ask JP if zombie flag returned from LB API
  - JP client needed for advanced JP searches

Involved People

- **Administrators**
  - New data flow – LB to JP
    - Automates LB database clean-up, keeps it reasonably sized
  - New services to maintain
- **Users**
  - Information about old jobs remain available for long-term
  - JP available
    - Advanced job searches possible
    - Application specific notebook about jobs
- **Developers**
  - Handle LB zombies in job state queries where appropriate
  - Developers of statistics & QA tools
    - Declaration *"SHOULD use LB notifications or JP, not direct access to LB database."* changed to *MUST*
    - We will provide support from LB/JP side
  - JP ready to support building of tools for experiments

User Point Of View

- **Job Provenance demonstration and evaluation projects**
  - Auger, chemical application (biomolecule docking), Atlas
  - Application specific attributes of jobs
  - Application specific workflow support tools (GUI)
- **Stress test done**
  - Overall JP: 3-4 megajobs/day
  - One bottleneck identified: GridFTP 150k jobs/day
    - We have a way how to remove this bottleneck (connection sharing)

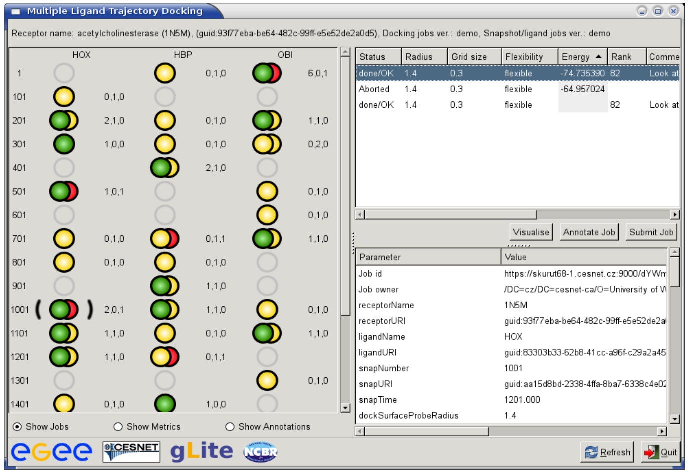

Receptor name: acetylcholinesterase (1N5M), (guid:93f77eba-be64-482c-99ff-e5e52de2a0d5), Docking jobs ver.: demo, Snapshot/ligand jobs ver.: demo

	HOX	HBP	OBI
1		0,1,0	6,0,1
101	0,1,0		
201	2,1,0	0,1,0	1,1,0
301	1,0,0	0,1,0	0,2,0
401		2,1,0	
501	1,0,1		0,1,0
601			0,1,0
701	0,1,0	0,1,1	1,1,0
801	0,1,0	0,1,0	
901		1,1,0	
1001	( ) 2,0,1	1,1,0	0,1,0
1101	1,1,0	0,1,0	1,1,0
1201	1,1,0	0,1,1	
1301			0,1,0
1401	0,1,0	1,0,0	

● Show Jobs ○ Show Metrics ○ Show Annotations

Status	Radius	Grid size	Flexibility	Energy ▲	Rank	Comme
done/OK	1.4	0.3	flexible	-74.735390	82	Look at
Aborted	1.4	0.3	flexible	-64.957024		
done/OK	1.4	0.3	flexible		82	Look at

Parameter	Value
Job id	https://skurut68-1.cesnet.cz:9000/dYWm
Job owner	/DC=cz/DC=cesnet-ca/O=University of W
receptorName	1N5M
receptorURI	guid:93f77eba-be64-482c-99ff-e5e52de2a
ligandName	HOX
ligandURI	guid:83303b33-62b8-41cc-a96f-c29a2a45
snapNumber	1001
snapURI	guid:aa15d8bd-2338-4ffa-8ba7-6338c4e02
snapTime	1201.000
dockSurfaceProbeRadius	1.4

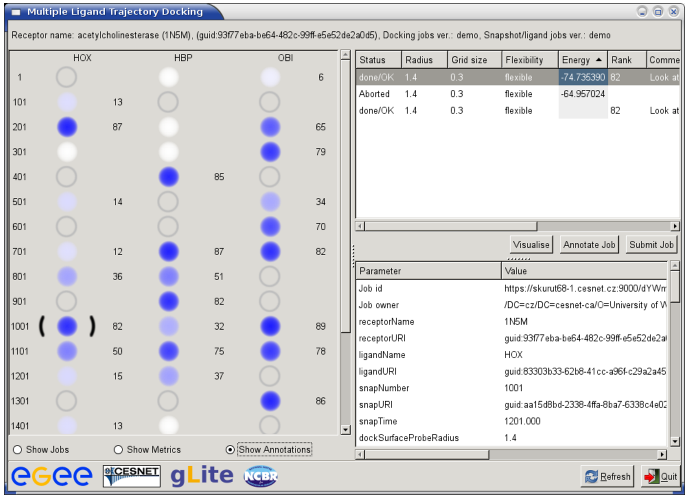

Receptor name: acetylcholinesterase (1N5M), (guid:93f77eba-be64-482c-99ff-e5e52de2a0d5), Docking jobs ver.: demo, Snapshot/ligand jobs ver.: demo

	HOX	HBP	OBI
1			6
101	13		
201	87		65
301			79
401		85	
501	14		34
601			70
701	12	87	82
801	36	51	
901		82	
1001	( ) 82	32	89
1101	50	75	78
1201	15	37	
1301			86
1401	13		

○ Show Jobs ○ Show Metrics ● Show Annotations

Status	Radius	Grid size	Flexibility	Energy ▲	Rank	Comme
done/OK	1.4	0.3	flexible	-74.735390	82	Look at
Aborted	1.4	0.3	flexible	-64.957024		
done/OK	1.4	0.3	flexible		82	Look at

Parameter	Value
Job id	https://skurut68-1.cesnet.cz:9000/dYWm
Job owner	/DC=cz/DC=cesnet-ca/O=University of W
receptorName	1N5M
receptorURI	guid:93f77eba-be64-482c-99ff-e5e52de2a
ligandName	HOX
ligandURI	guid:83303b33-62b8-41cc-a96f-c29a2a45
snapNumber	1001
snapURI	guid:aa15d8bd-2338-4ffa-8ba7-6338c4e02
snapTime	1201.000
dockSurfaceProbeRadius	1.4

First Phase - Steps

- **Certify LB with purger (March)**
- **Release JP in RESPECT (April)**
- **Maintain JP at JRA1 Preview testbed**
  - New JP-PS and JP-IS based on released JP packages working here in April

Summary – JP Deployment

- **Strategy: provide new potential at low deployment cost**
  - Users can see their jobs in JP and start to explore JP features
- **Benefits**
  - LB + JP works in tandem as designed (JP database cleaning)
  - JP ready for users
    - Job info available for long-term
    - Experiments (custom tools to analyze and automate experiments)
- **Costs**
  - JP Primary Storage deployment and maintenance
  - JP Index Servers with default configuration deployment and maintenance
  - Job status queries internals to be changed

Questions?

Thank You

for your attention!