

UNICORE 7 – Middleware Services for Distributed and Federated Computing

Krzysztof Benedyczak¹, **Bernd Schuller**², Maria Petrova El-Sayed², Richard Grunzke³

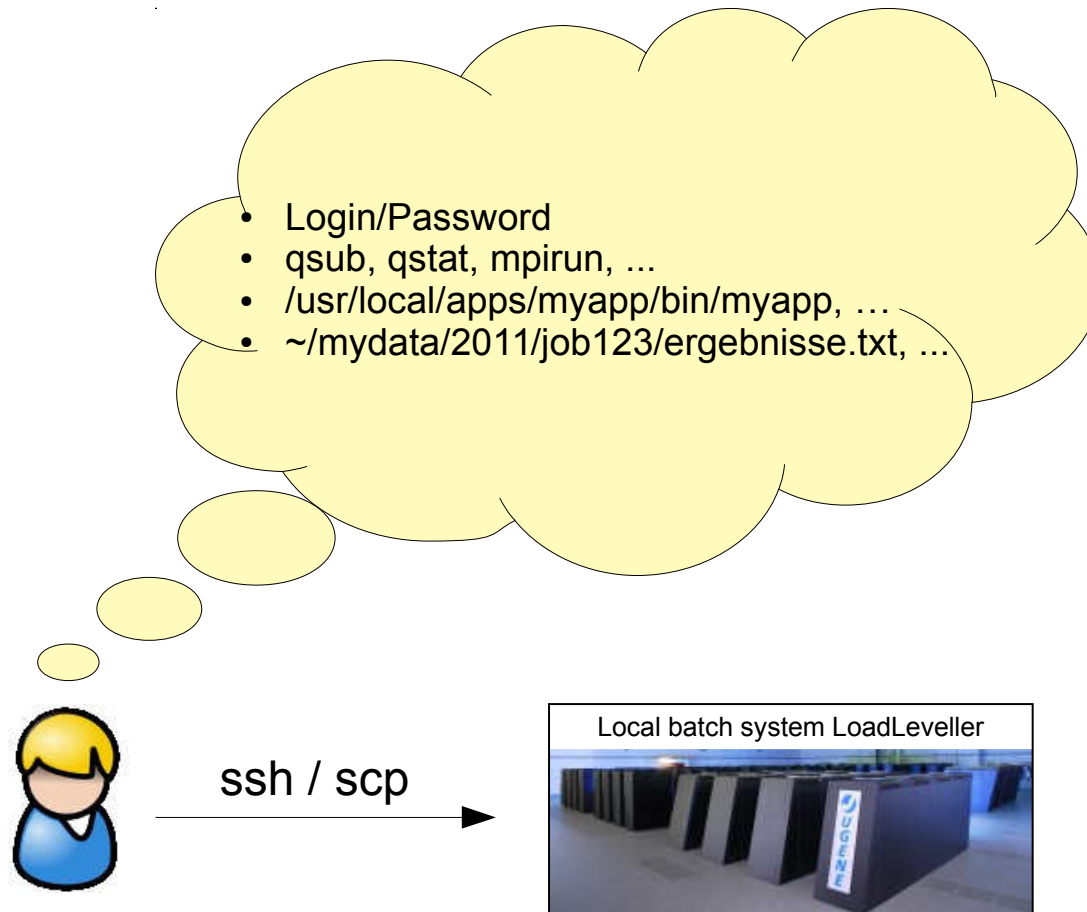
¹ Interdisciplinary Center for Math. and Comp. Modelling, Warsaw, Poland


² JSC, Forschungszentrum Jülich GmbH, Germany

³ ZIH, Technische Universität Dresden, Germany

Outline

- UNICORE – UNiform Interface to COmputing and data REsources
- Services for Distributed and Federated computing
 - UNICORE : overview
 - Unity : user authentication and identity management
 - UNICORE Web Portal
 - RESTful APIs
- Summary and outlook





How can I ...

- ... use multiple, heterogeneous systems seamlessly,
- ... manage my job input data and results?
- ... across multiple systems? Workflows?
- ... integrate HPC and big data into custom applications and portals?

UNICORE

Web Command line GUI API

Clients

Users

Workflows Jobs Data Management Discovery

Services

Federations

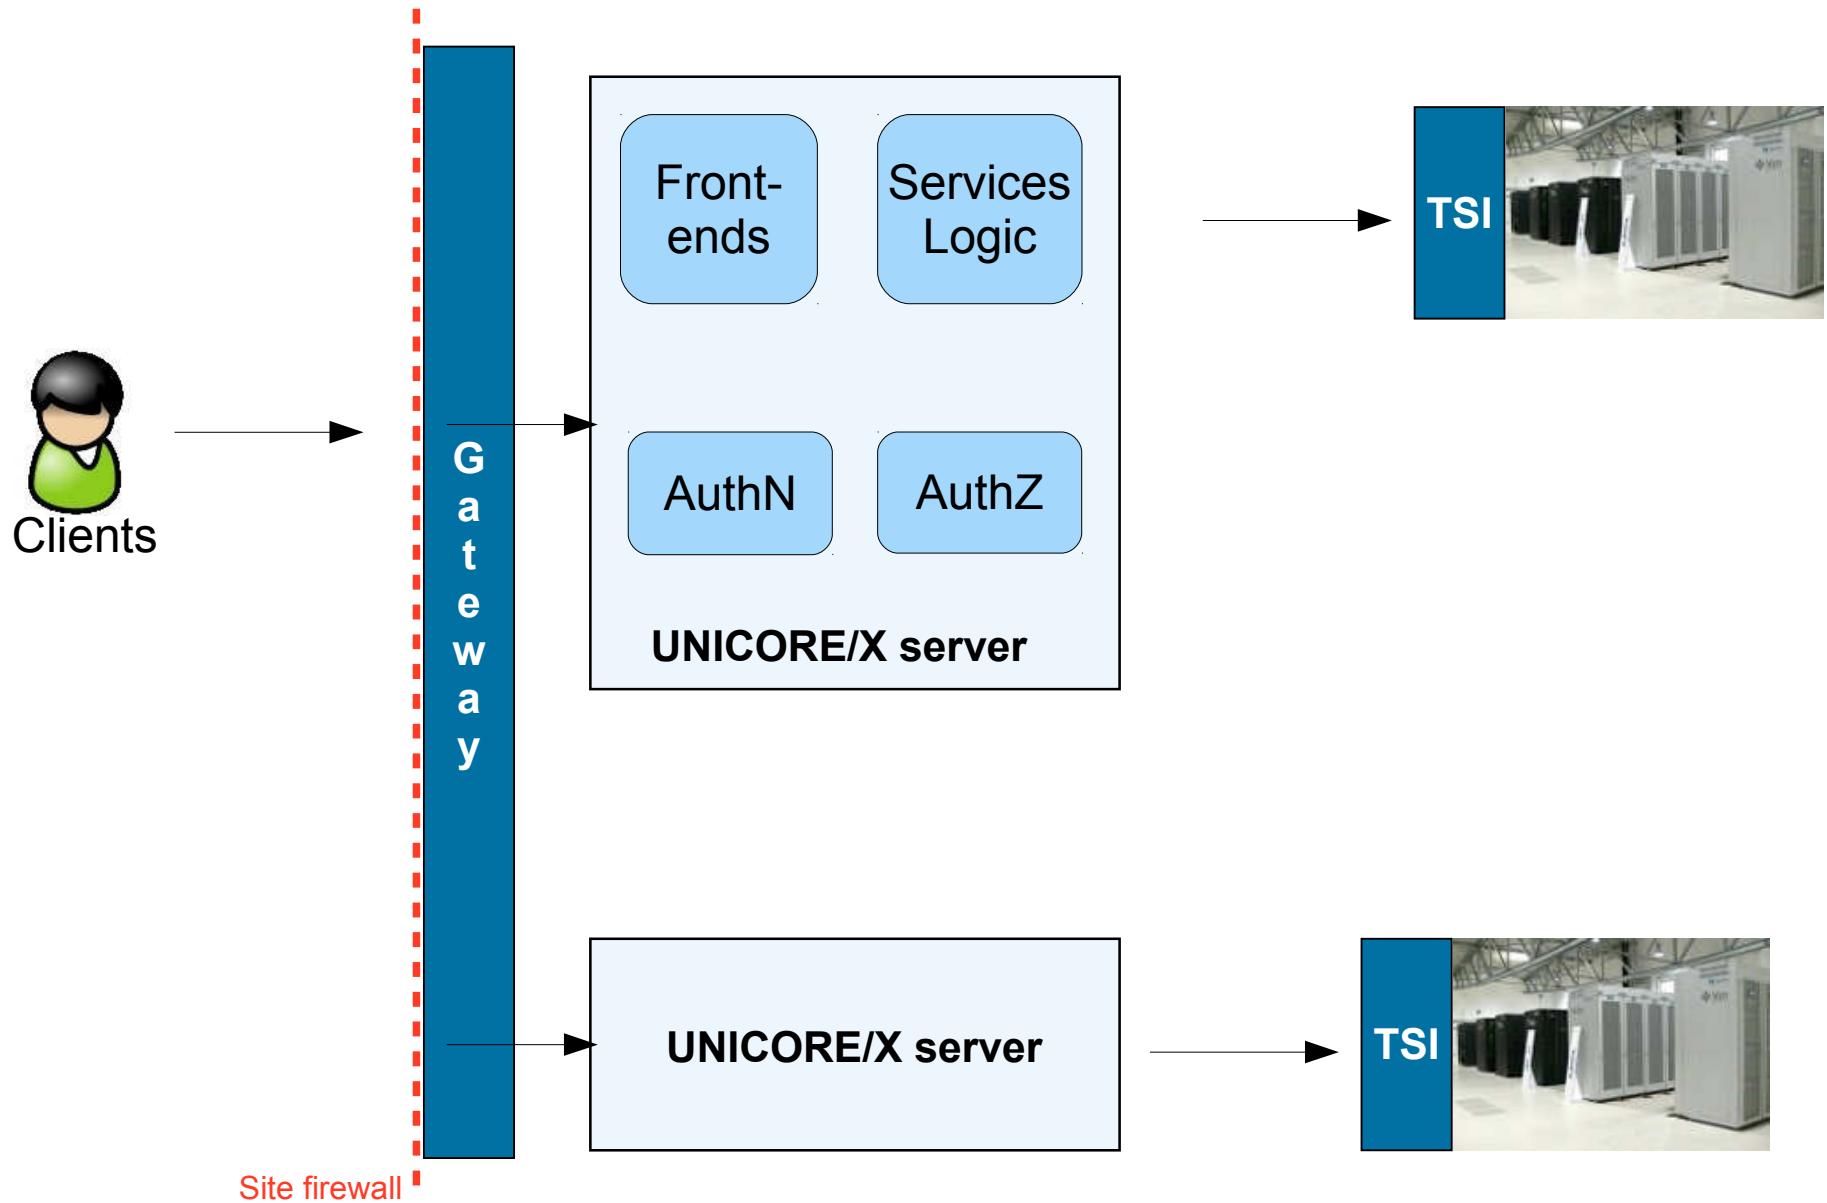
Compute Storage

Resources

Policies

Security

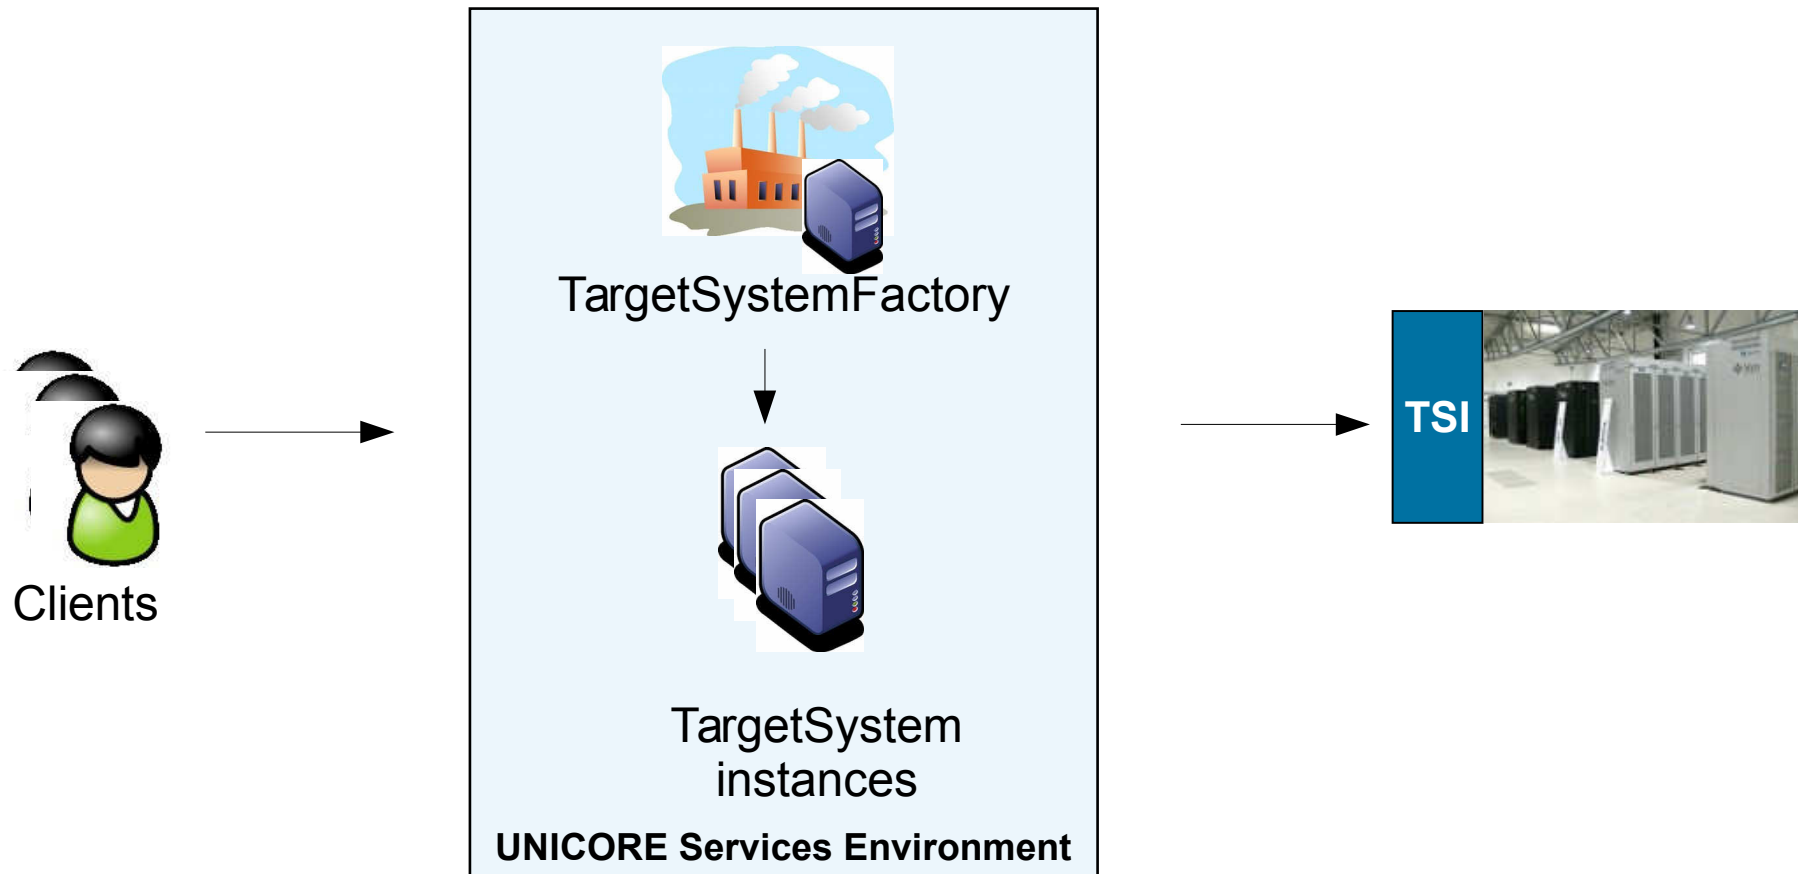
Typical UNICORE installation at a HPC center





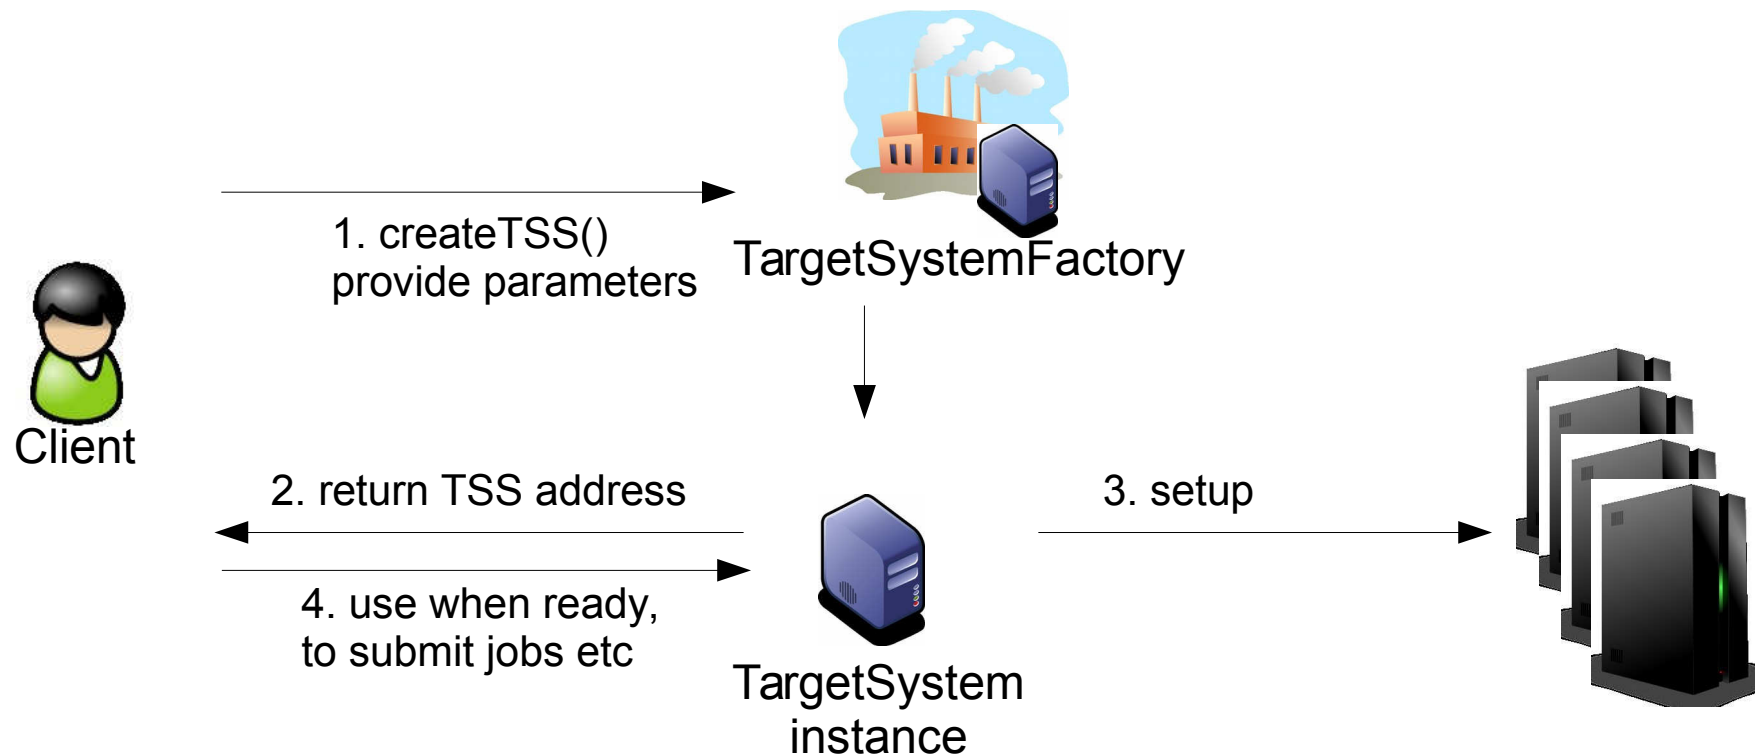
- Workflow enactment
- Task execution
- TargetSystemFactory
- TargetSystem
- JobManagement
- Reservations
- StorageFactory
- StorageManagement
- FileTransfer
- Metadata
- Registry
- Resource Broker

Default setup



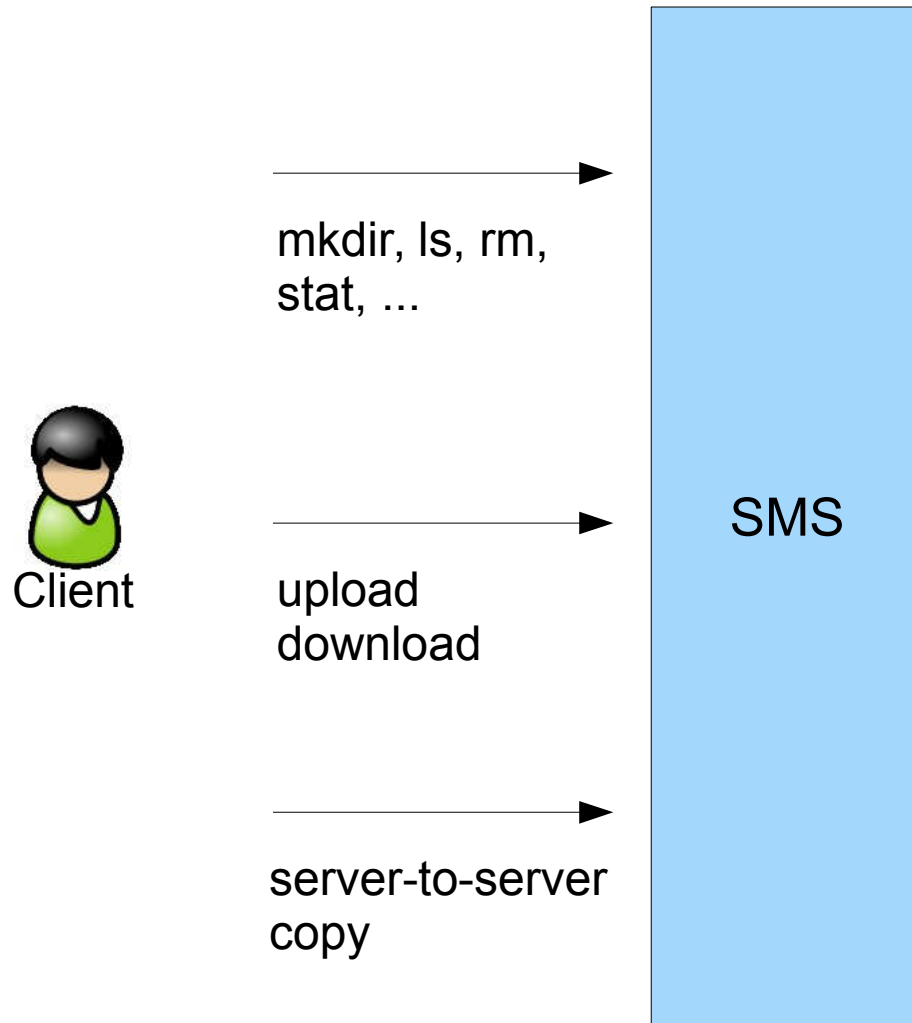
- Access to resource manager and file system via TargetSystemInterface (TSI) daemon installed on the cluster login node(s)

Factory services: virtualisation support



- Set up a virtual image during initialisation phase
- Aim at OpenStack VMs, Amazon EC2, ...

Storage Management Service



- File systems

- Apache HDFS



- S3



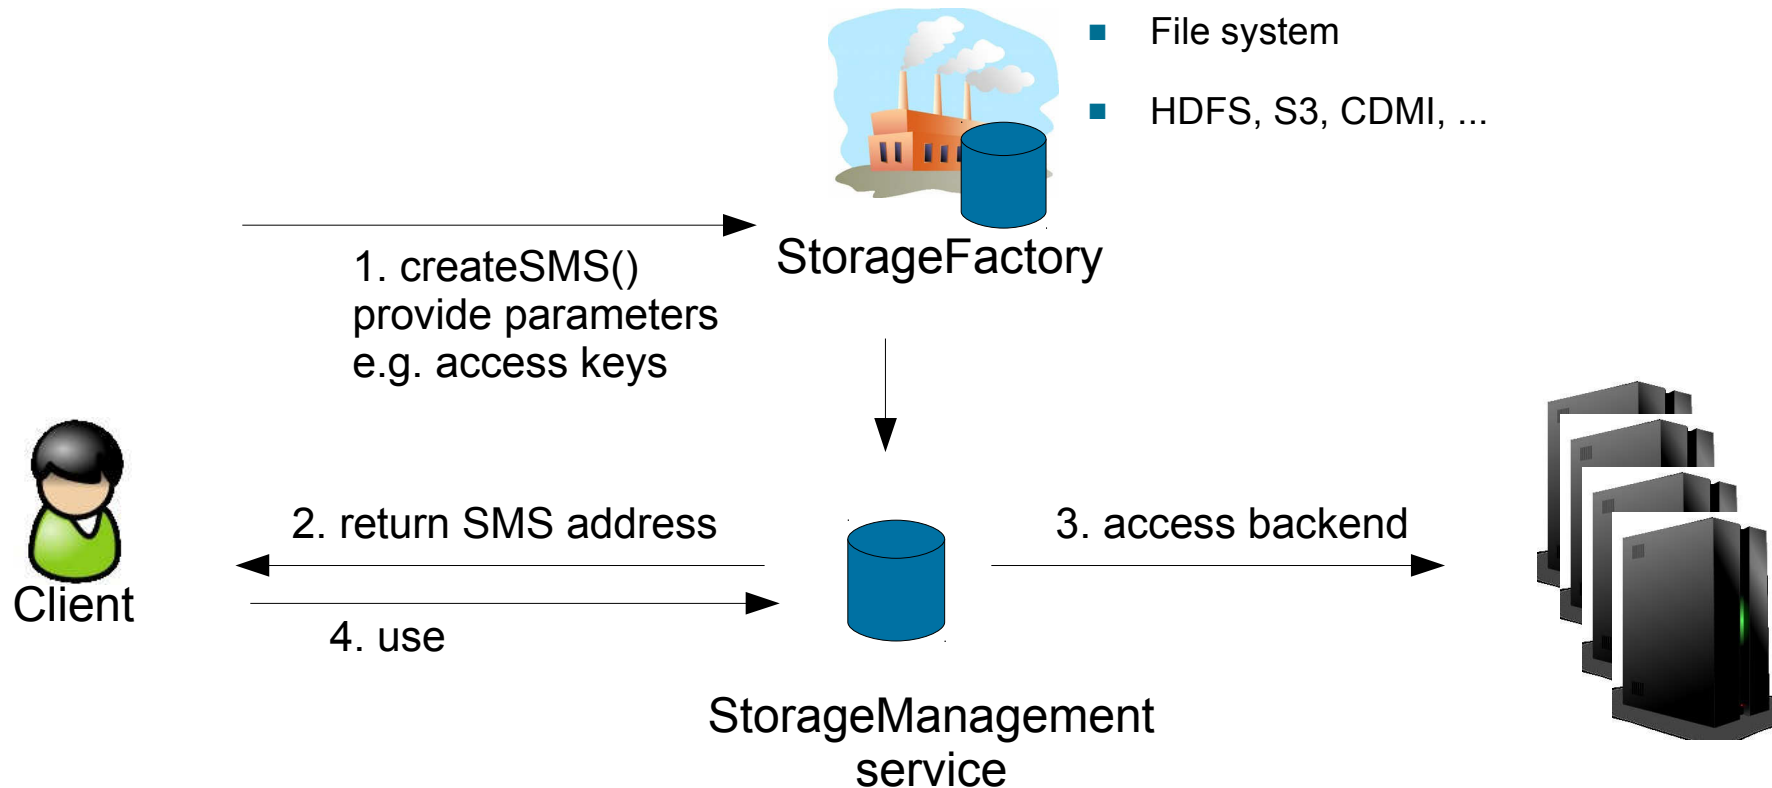
- CDMI (prototype)

- ...

Storage Management Service: more than a file system

- Initiate file transfers
 - Multi-protocol support
 - Scheduled server-to-server copy
- Metadata management
 - Schema-free, key-value
 - Indexed via Lucene, searchable
- Rule-based data processing
 - New files automatically trigger actions
 - e.g. metadata extraction, compression, etc

Factory services: virtualisation support



- Different types of storage backends can be supported
- User can select and provide required parameters



Compute





Storage

Resources

- Batch systems (Torque, Slurm, LoadLeveler, GridEngine, ...)
- Apache Hadoop (YARN)
- Direct execution (e.g. on Windows)
- ... (extensible)
- File systems
- Apache HDFS
- S3
- ... (extensible)


Federated access: security is key



- **Service invocation:** a web service call is made to a UNICORE service
- **Authentication:** who is the user?
 - Results in the user's X.500 DN („CN=..., O=..., OU=..., C=...“)
- **Assign attributes** to the user
 - Standard attributes: role, Unix ID, groups, etc.
 - Custom attributes: (e.g. S3 access and secret keys)
- **Authorisation**
 - Add context: e.g. who owns the resource?
 - Check resource policies (ACLs)
 - Check server policies (XACML)
- → **Allow or deny** the request

Delegation

- Allow Service to work on behalf of the user
- UNICORE solution based on SAML
 - Use chain of signed assertions
 - Trust always delegated to particular server
 - Can be validated and audited



End-user authentication in UNICORE

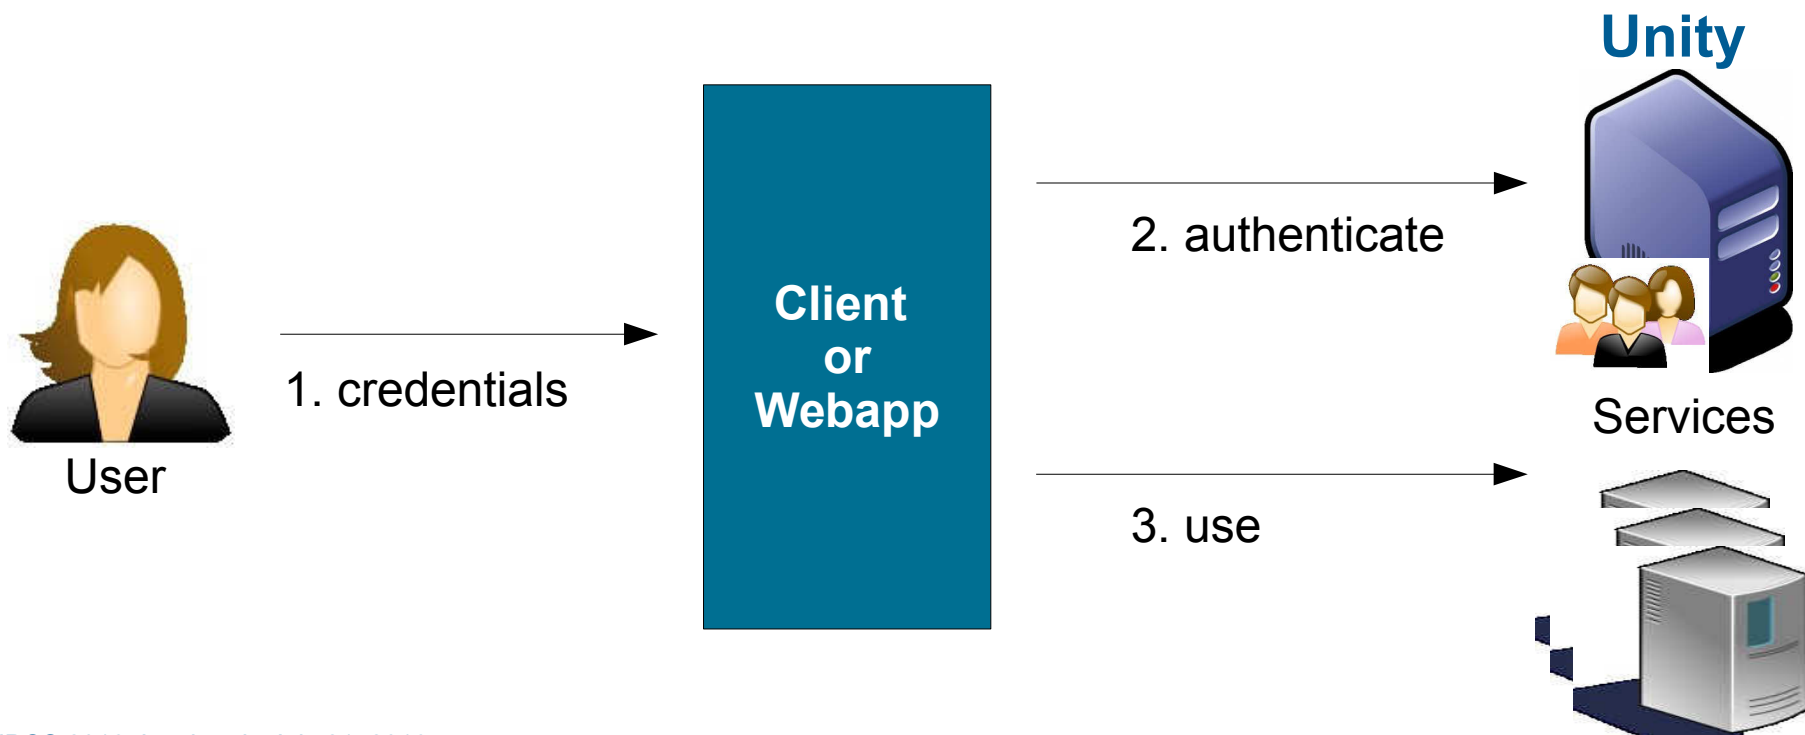
- Pre-UNICORE 7: X.509 client certificates **REQUIRED** for end-users
- Users tend to hate them
 - All sorts of usage issues
- Lack of understanding leads to lack of security (sending keys via email etc)
- Users understand passwords
 - and it is relatively easy to teach basic security measures

Certificate-less end-user authentication

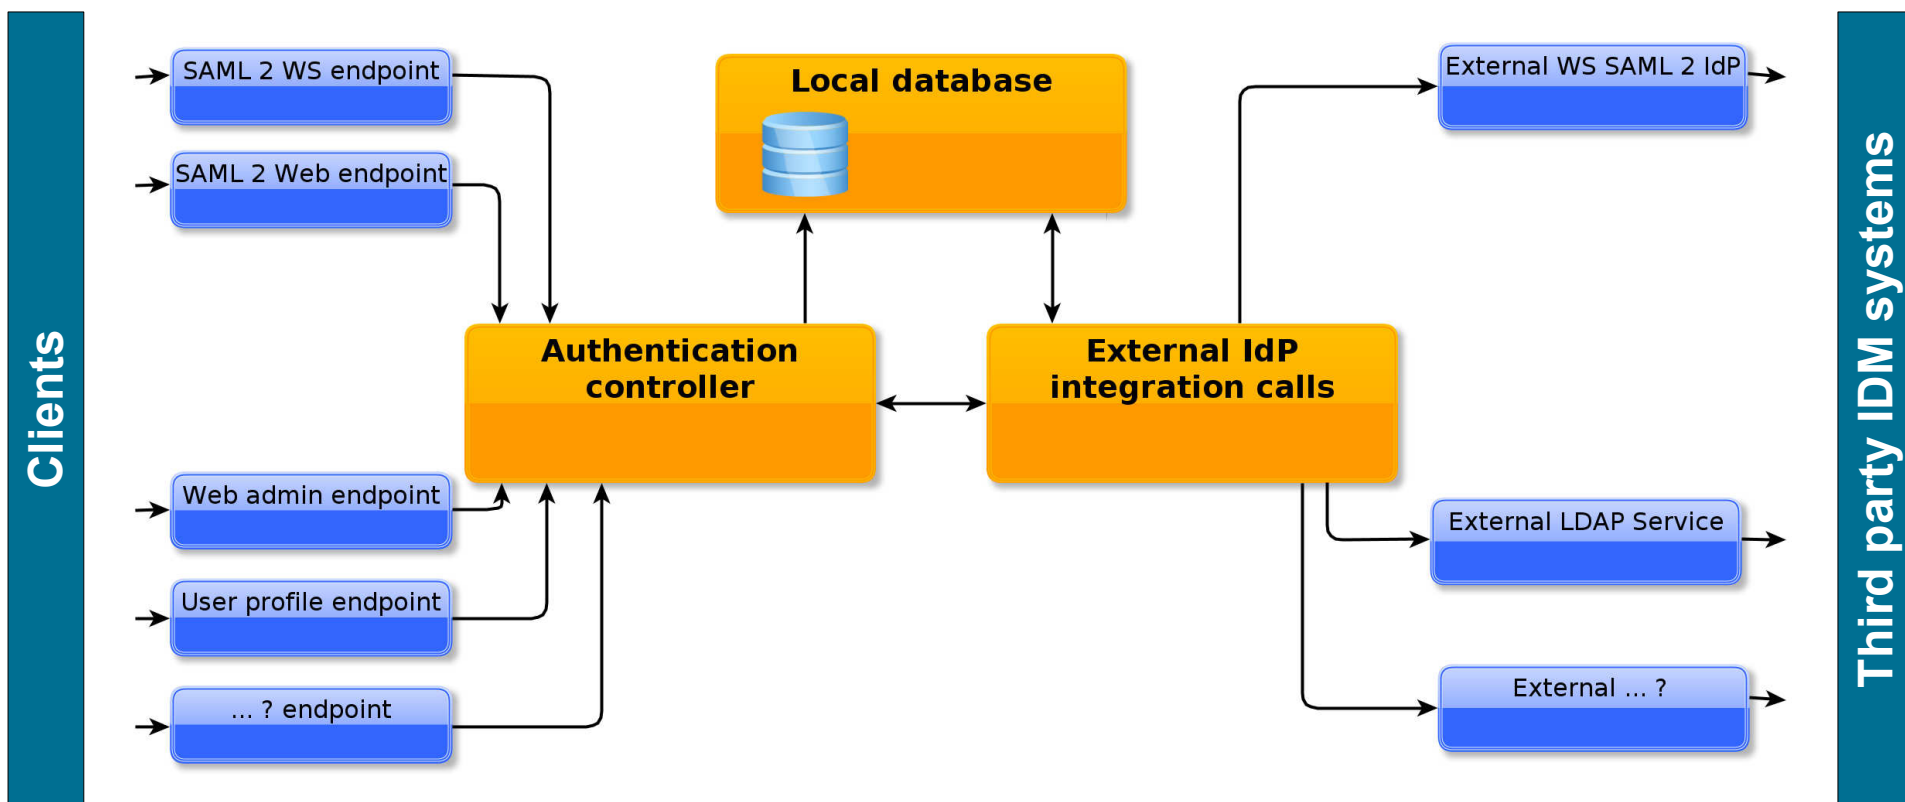
- Goal: **no end-user certificates** (not even short-lived)
- Approach
 - Use **signed SAML assertions** for authentication
 - Issued and signed by a trusted service
 - Flexible solution is required: e.g. want support for existing SAML Identity providers , federations like DFN AAI, OAuth, etc
- Implications
 - Client – server TLS is not client-authenticated
 - End-user cannot sign anything (no „non-repudiation“ guarantee)

Introducing Unity

- Complete **Authentication and Identity Management** solution
- Manage users and user attributes, group membership
- Separate, standalone product: www.unity-idm.eu
- Increasing take-up: e.g. HBP, EuDAT






Unity architecture



Managing Unity via web application

UNITY administration interface

Logged as: Default Administrator [entity id: 1]   

Contents management | Registrations management | Schema management | Server management

Groups

- Root (/)
 - A
 - D
 - portal

Group /portal members

Group by entities Show targeted identities Search:

ENTITY	IDENTITY TYPE	IDENTITY	S
[3]	userName	demo	ENA
[3]	persistent	5c1e8334-e268-4ddd-a7c7-3097bc320813	ENA
[3]	x500Name	CN=Demo User,O=UNICORE,C=EU	ENA

Group /portal details

GROUP'S ATTRIBUTES CLASSES
UNICORE portal attributes

ATTRIBUTE STATEMENTS

Attributes of entity [3] in group /portal

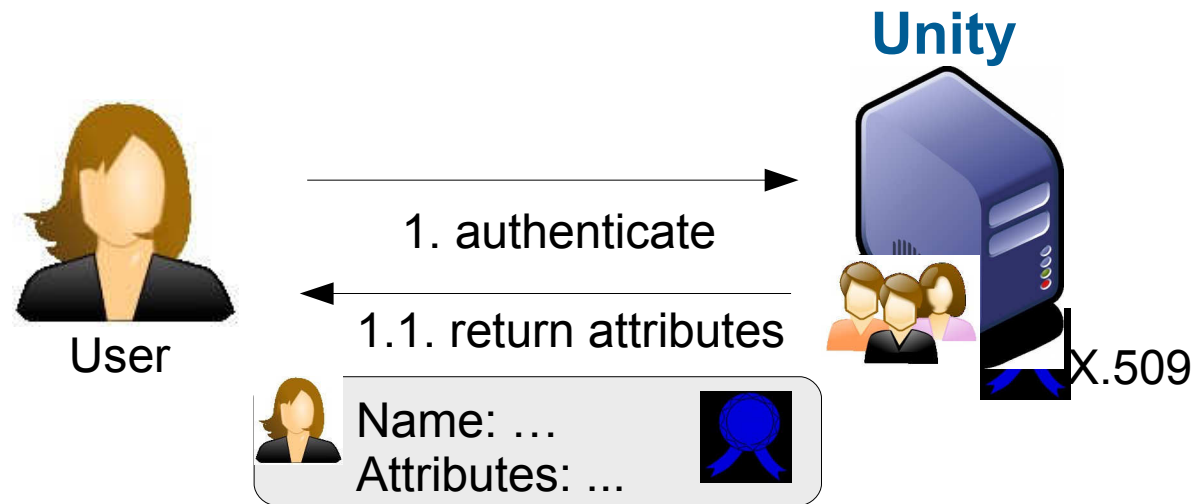
Effective Internal **Required in bold**

ATTRIBUTES
email
cn

Information
Directly defined
Created at 11/16/14 10:19 AM updated at 11/16/14 10:19 AM

Value
test@example.com

Example: Unity authentication assertion



```

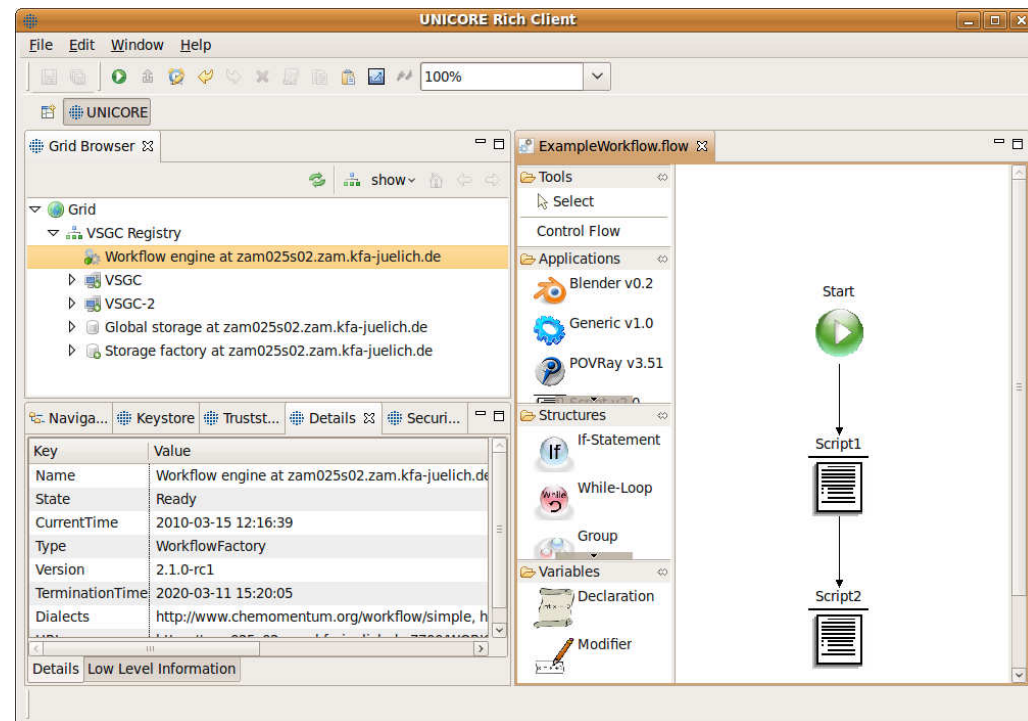
<urn:Assertion>...
  <dsig:Signature... </dsig:Signature>
  <urn:Subject>
    <urn:NameID
      Format="urn:oasis:names:tc:SAML:1.1:nameid-format:X509SubjectName">CN=Demo User,O=UNICORE,C=EU</urn:NameID>
    <urn:SubjectConfirmation Method="urn:oasis:names:tc:SAML:2.0:cm:sender-vouches">
      <urn:SubjectConfirmationData NotOnOrAfter="2014-11-16T10:30:23.334Z"/>
    </urn:SubjectConfirmation>
  </urn:Subject>
  <urn:AttributeStatement>
    <urn:Attribute Name="cn">
      <urn:AttributeValue>Demo User</urn:AttributeValue>
    </urn:Attribute>
    <urn:Attribute Name="email">
      <urn:AttributeValue>test@example.com</urn:AttributeValue>
    </urn:Attribute>
    <urn:Attribute Name="memberOf">
      <urn:AttributeValue>/portal</urn:AttributeValue>
      <urn:AttributeValue>/</urn:AttributeValue>
    </urn:Attribute>
  </urn:AttributeStatement>
</urn:Assertion>
  
```



- Portal
- UCC
- Eclipse-base Rich Client
- RESTful API
- Third-party science gateways
- UFTP
- Custom clients
- Java APIs

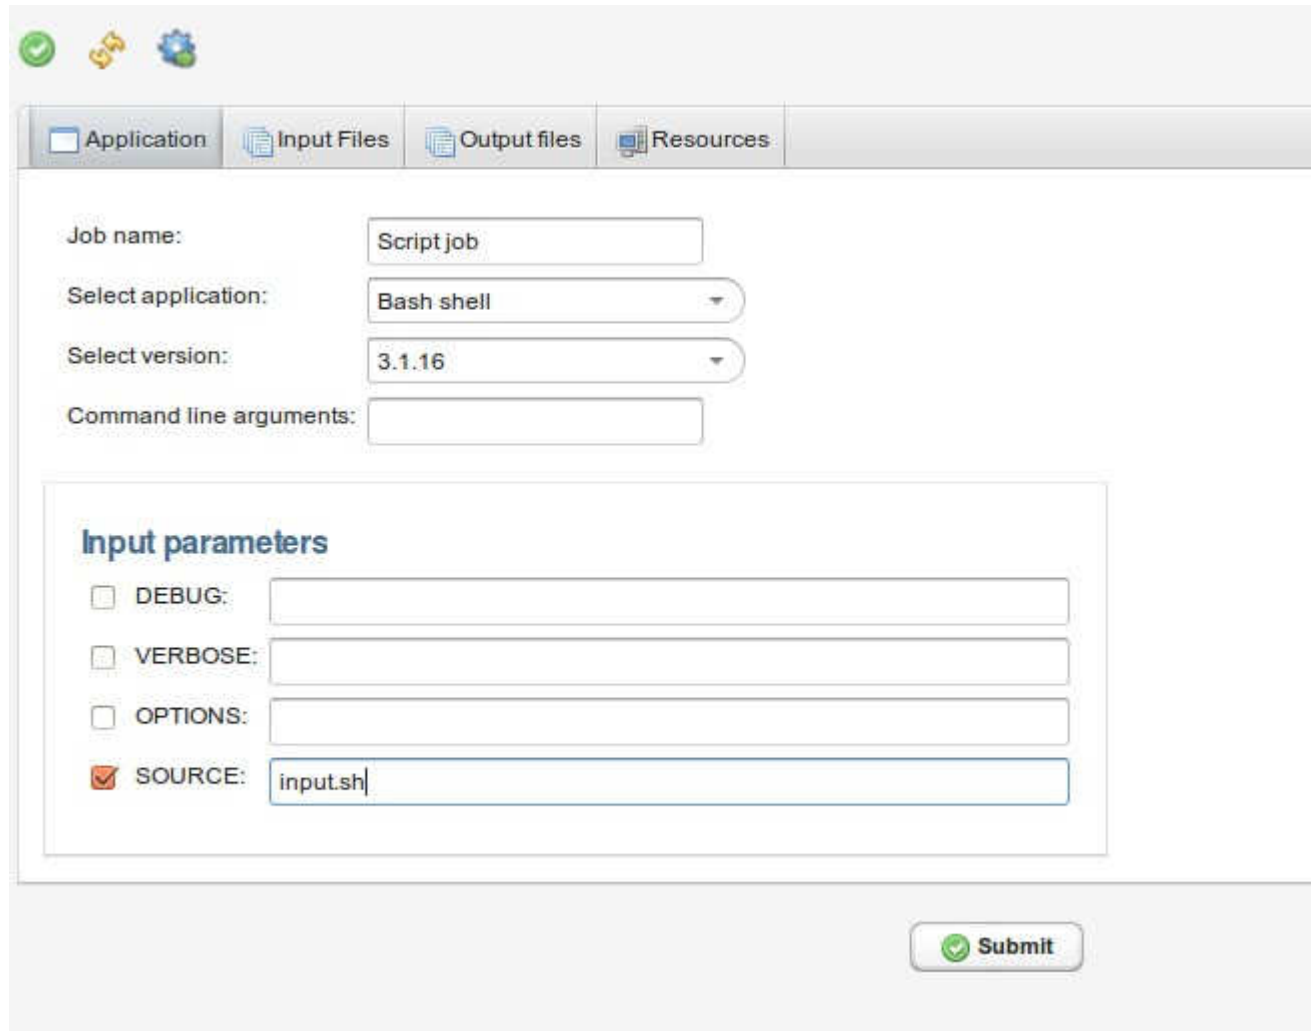
UNICORE Rich client

- Based on the Eclipse framework
- Building, submitting and monitoring jobs and workflows
- Integrated data and storage management
- X.509 and Unity for AuthN
- *Mostly targeted at expert users*



- Aim for a simple, easy-to-use web application
- Simple use cases
 - Job and (limited) workflow management
 - Data management
- Less details exposed to user
- Implementation choices
 - Java-based, VAADIN web framework
 - Use UNICORE SOAP/WS APIs

UNICORE Portal – Job creation view



The screenshot shows a web-based form for creating a job. At the top, there are three icons: a green checkmark, a yellow dollar sign, and a blue gear. Below these is a navigation bar with four tabs: 'Application' (selected), 'Input Files', 'Output files', and 'Resources'. The main form area contains the following fields:

- Job name:** A text input field containing 'Script job'.
- Select application:** A dropdown menu with 'Bash shell' selected.
- Select version:** A dropdown menu with '3.1.16' selected.
- Command line arguments:** An empty text input field.

Below these fields is a section titled 'Input parameters' with a light blue border. It contains four rows, each with a checkbox and a text input field:



- DEBUG:** [empty text input]
- VERBOSE:** [empty text input]
- OPTIONS:** [empty text input]
- SOURCE:** [input.sh]
















At the bottom right of the form is a 'Submit' button with a green checkmark icon.

UNICORE Portal – various

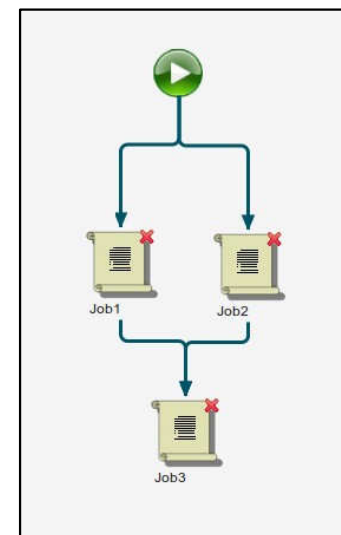
- Several „list“ views, e.g. jobs, sites

Jobs Browser

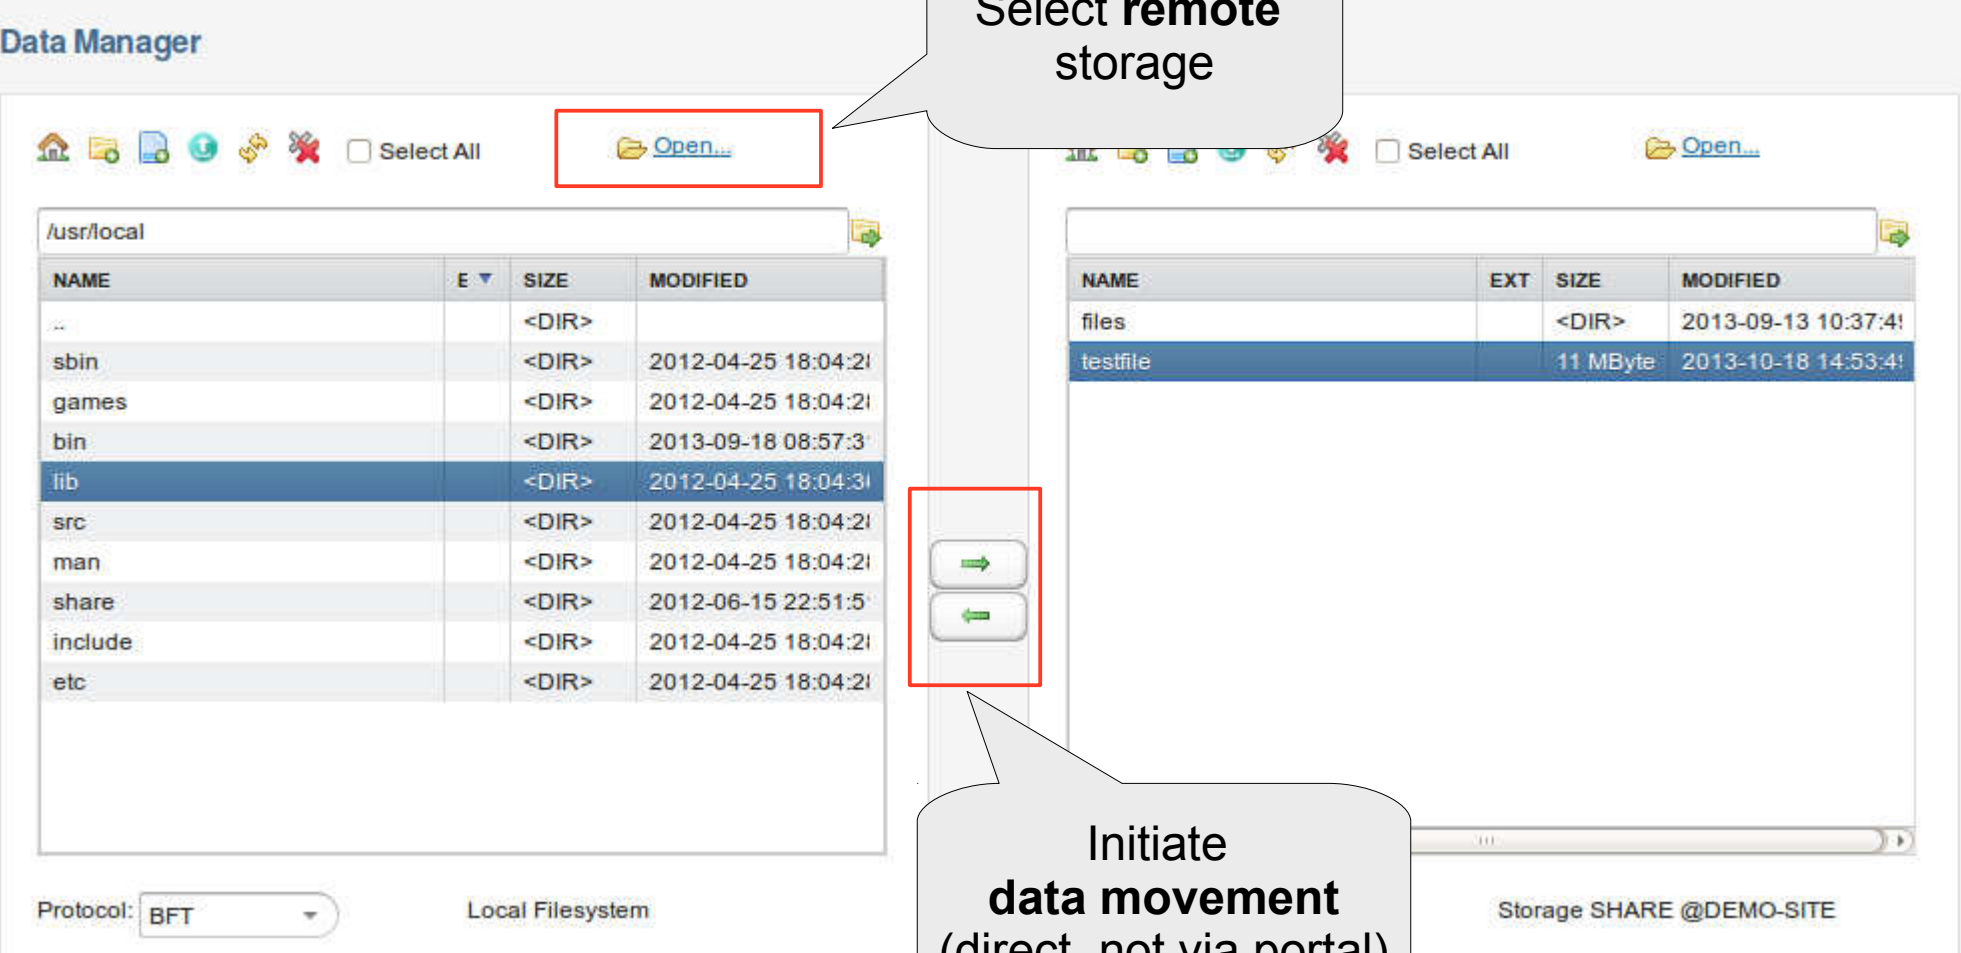


 Select All
 Items per page:

	NAME	JOB STATUS	SITE	QUEUE	ESTIMATED FINISH TIME	ACTIONS
	Job1	SUCCESSFUL	DEMO-SITE	N/A	unknown	   
	Test 1	QUEUED	DEMO-SITE	N/A	unknown	   
	Example job	QUEUED	DEMO-SITE	N/A	unknown	   

- Simple workflow creation
- JavaScript



UNICORE Portal: Data manager



Select remote storage

NAME	E	SIZE	MODIFIED
..		<DIR>	
sbin		<DIR>	2012-04-25 18:04:21
games		<DIR>	2012-04-25 18:04:21
bin		<DIR>	2013-09-18 08:57:3
lib		<DIR>	2012-04-25 18:04:31
src		<DIR>	2012-04-25 18:04:21
man		<DIR>	2012-04-25 18:04:21
share		<DIR>	2012-06-15 22:51:5
include		<DIR>	2012-04-25 18:04:21
etc		<DIR>	2012-04-25 18:04:21

Protocol: BFT Local Filesystem

NAME	EXT	SIZE	MODIFIED
files		<DIR>	2013-09-13 10:37:4!
testfile		11 MByte	2013-10-18 14:53:4!

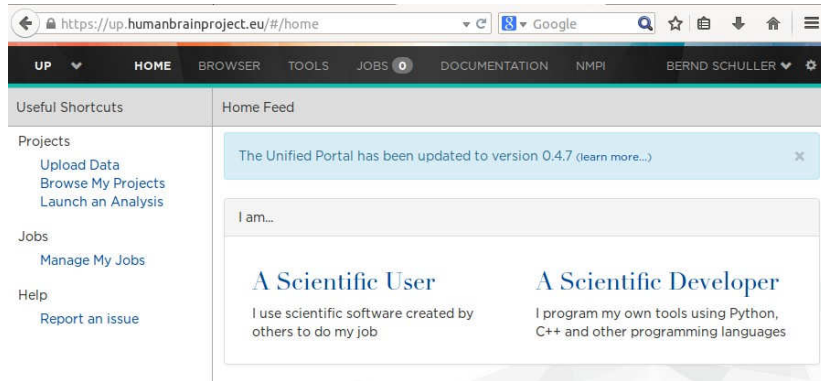
Storage SHARE @DEMO-SITE

Initiate data movement (direct, not via portal)

RESTful APIs to UNICORE services

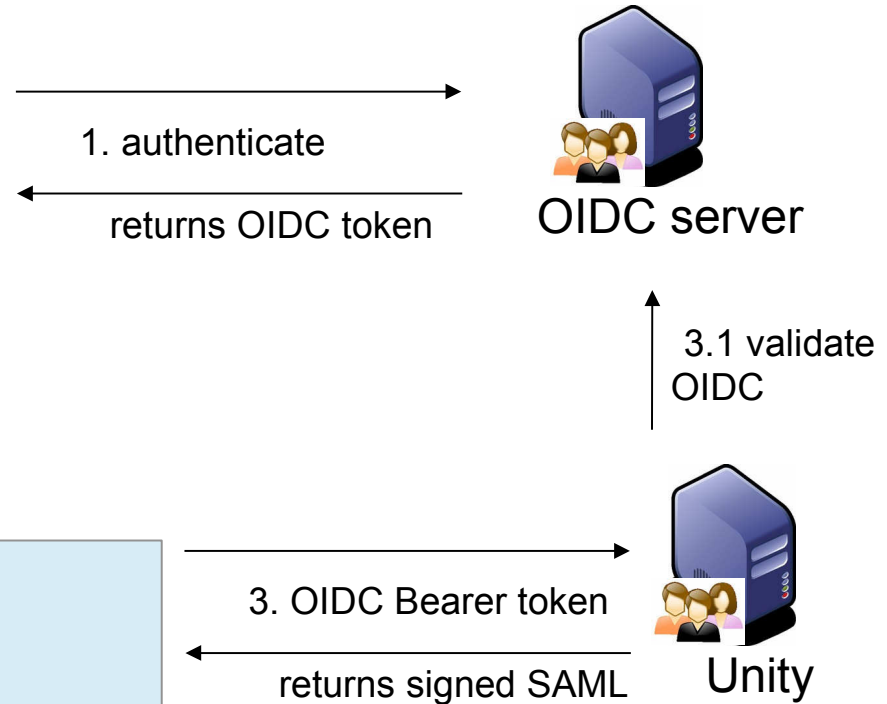


Use case: the Human Brain Project's HPAC platform



<https://collab.humanbrainproject.eu>

2. pass OIDC Bearer token



BSC

HPC site

CINECA

HPC site

CSCS

HPC site

JSC

HPC site

KIT

S3 storage

SOAP and WS(RF) – Defining UNICORE since 2004/2005

- Pros
 - Strongly typed – XML schema based
 - SOAP header/envelope mechanism
 - WS-Security, SAML, etc are well established

- Cons
 - CPU intensive (XML processing, XML signatures)
 - Complex interface (look at a typical WSDL!)
 - Only Java and C# can be realistically used on the client side

RESTful APIs to UNICORE Services

- REST
 - Document / Resource oriented approach
 - HTTP semantics (GET, PUT, POST, DELETE, error codes, caching, ...)
 - Multiple message formats and resource representations can be used
 - *JSON, XML, HTML, ...*
 - Several authentication options (HTTP basic, OAuth, ...)
 - Clients in all languages (even *curl* or *wget*)
- Keep SOAP/WS (for backwards compatibility), fully in sync with RESTful APIs

Example: job submission

```
job.u:  
{  
  "Executable": "/bin/echo" ,  
  "Arguments": ["Hello World"],  
}
```

```
$> curl -X POST -H "ContentType: application/json"  
  --data-binary @job.u  
  https://localhost:8080/DEMOSITE/rest/core/jobs
```

```
HTTP/1.1 201 Created  
ContentType: application/json;charset=ISO88591  
Location: https://localhost:8080/DEMOSITE/rest/core/jobs/ \  
          74198236e970429db55ca7d59c831f14
```

- UNICORE – a complete solution for building federations
- Main progress in **UNICORE 7**
 - Simplify user experience
 - *Make X.509 user certs obsolete*
 - **Web portal** targetted at non-expert users
 - Simplify integration options
 - Complete set of **RESTful APIs** for computing and data
→ *bringing HPC to the Web!*
 - Widen integration options
 - **Unity** as a universal solution for federated identify management solution

Outlook

- Consolidate and simplify
 - Installation and configuration
 - Packaging and automation of deployment
- Add/extend support for
 - Cloud resources (OpenStack, EC2, ...)
 - Virtualised applications (Docker)
- <http://www.unicore.eu>



- Björn Hagemeyer, Valentina Huber, André Giesler, Boris Orth, Mariya Petrova, Jędrzej Rybicki, Rajveer Saini, Bernd Schuller and many others at JSC
- Krzysztof Benedyczak, Marcelina Borcz, Rafał Kluszczynski, Piotr Bała and others at ICM / Warsaw University
- Richard Grunzke and others at Technical University Dresden
- Students: Burak Bengi, Maciej Golik, Konstantine Muradov
- ... many others who reported bugs, suggested features, contributed code and provided patches

The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 604102 (Human Brain Project)

A federation software suite

- Secure and seamless access to compute and data resources
- Focus on scientific applications and workflows
- Complies with typical HPC centre policies
- Complete solutions: APIs, clients, services, ...
- Java/Python based, supports UNIX, MacOS, Windows and many resource management systems (Torque, Slurm, SGE, ...)
- Long development history (since 1997)
- **Open source, BSD licensed, visit <http://www.unicore.eu>**