Web Services and SGE – Open DRMAA Service Provider (OpenDSP)

Piotr Domagalski
piotr.domagalski@man.poznan.pl
Krzysztof Kurowski
office@fedstage.com

Overview

- Why did we develop OpenDSP?
- DRMAA
- Main features
- AAA mechanisms included
- Other components
- Performance
- FedStage

Motivations

- Consistent access to many DRMs (we started with SGE)
  - remote
  - multi-user
  - policy-based
- Extensible Authentication, Authorization, Accounting (AAA)
- Interoperability – try to use only accepted standards
- High performance
- Portability

OGF DRMAA

- Distributed Resource Management Application API
- Interface to submit jobs, control them, obtain their status and exit information
- Bindings for many languages (C, Java, Perl, Python, Ruby)
- Single interface to many DRMs
  – SGE, Condor, Torque, PBS, LSF
- Preferably uses internal API – performance gain comparing to other solutions
- But only local interface with limitations (administration, sessions)

DRMAA vs OpenDSP

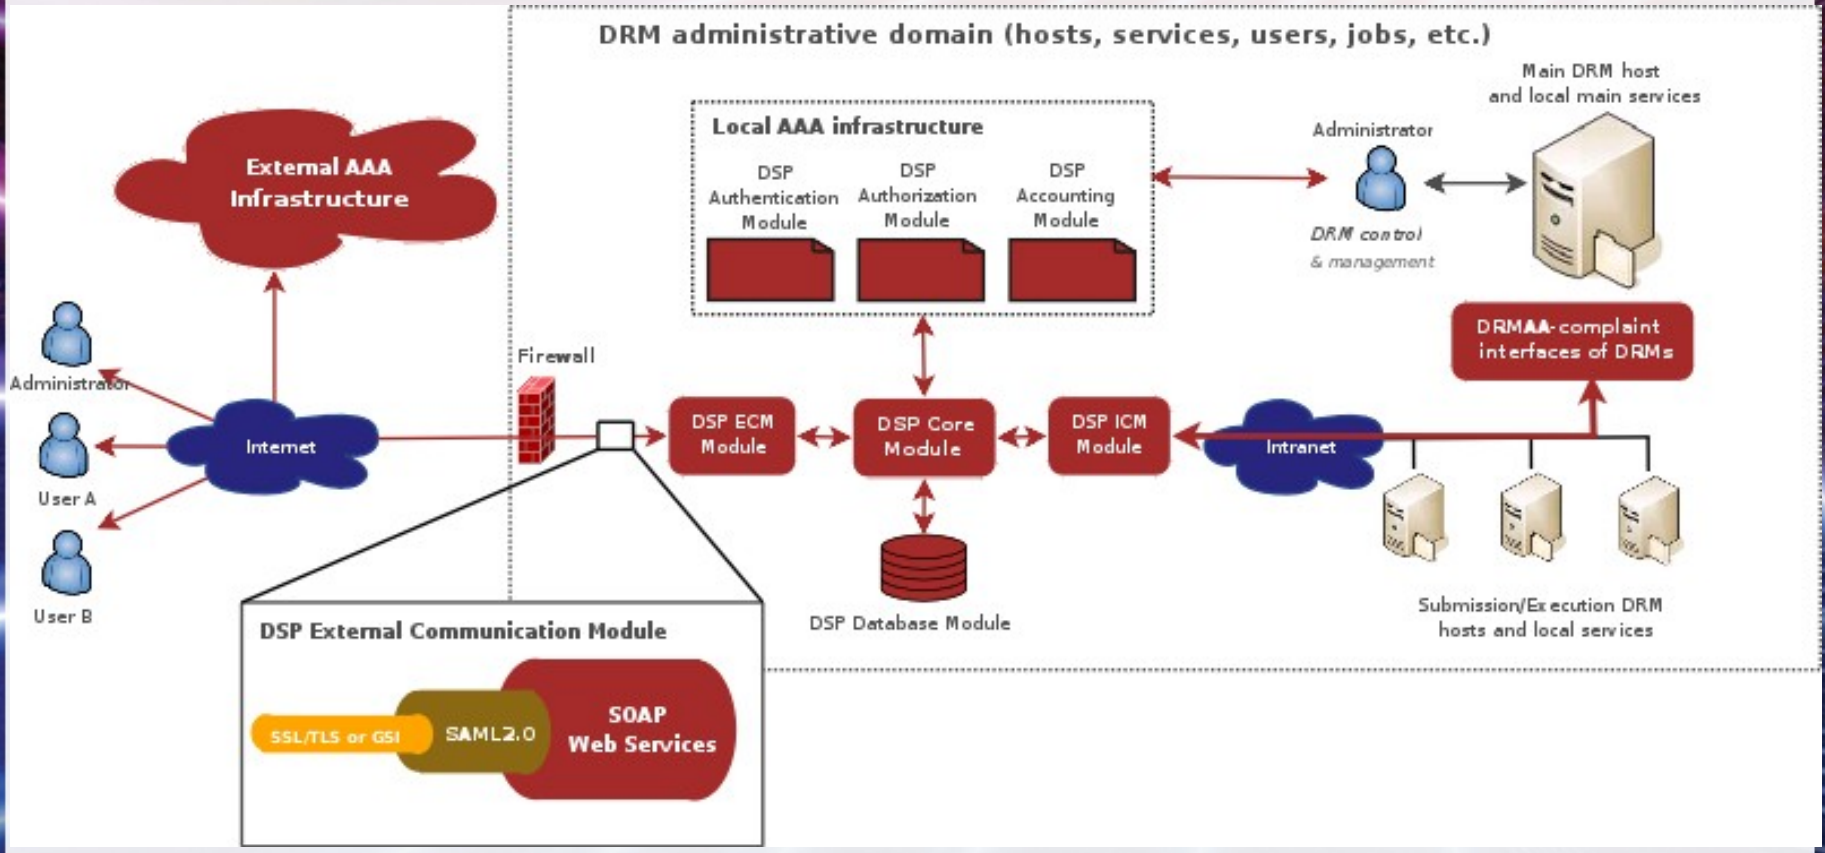

OpenDSP features

- Written in C using lightweight components (gSOAP) with modular approach
- OGSA-BES-like Web Service interface
  - Submit single jobs and job arrays (also parallel jobs)
  - Control the job (hold, suspend, terminate)
  - Monitor the job state (push and pull)
  - Input, output and error file staging
  - Retrieve job description
  - Retrieve DRM system information and control it
- Jobs described with JSDL documents, extensible through DRM-dependent options
- At most once submission with two-phase commit (UUID on create in the upcoming version)

OpenDSP features, cont.

- Strong privilege limiting model - no code running as root
- Audit logging to ODBC-compliant database
  - Job owner and job description
  - Used resources: computing and time
- Works on Linux, Solaris, FreeBSD
- Successfully tested with SGE, Condor, Torque and LSF

Authentication

- Modular architecture allows to mix
  - Transport level authentication and security
    - Plain HTTP
    - SSL with or without client authentication (X.509)
    - GSI (X.509)
  - Message level authentication
    - WS-Security Username Token Profile
    - WS-Security SAML Token Profile
    - WS-Security X509 Token Profile

Authorization

- Authorization decision and local user mapping
- Policy Decision Point (PDP) callout chain
- Available modules for
  - Globus-like mapfile
  - Gridge Authorization Service (GAS)
  - Anonymous authorization
  - Time/connection rate limits, etc.

Accounting

- Track for how long and how many processors the software runs on
- Computing resource providers and software owners can easily adopt new utility-based pricing schemes that charge customers based on how much use they get out of their software
- Broader adoption of utility computing and detailed accounting - businesses can adjust the amount of computing resources devoted to specific applications as needed
- OpenDSP is using its own accounting database and reporting/billing service integrated with PayPal (check out FedStage for more details)

State access: push model

- WS-Notification Broker (OASIS specification) enables dynamic registration of various Web Service publishers, notification topics; supports XPath queries
- OpenDSP as a Notification Producer generates messages for
  - Job submission
  - Job termination
  - Job control event
- Notification Consumer subscribes to the Notification Broker

Monitoring Service

- Used for service discovery purposes and job/computing resources monitoring
- Two modules
  - Ganglia
  - Our own monitoring system
    - Based on WS-Notification Broker
- Publish metrics with current load statistics (e.g. average queue waiting time, reservations)

Identity Provider

- Our own identity and authentication service
- Issues SAML assertions

Performance

- No limits on maximum active jobs (DRM or DRMAA limits)
- Controllable maximum concurrent requests
- Using SSL at transport level
  - can serve over 30 requests per second – with no request failure
  - can serve peak load of over 600 requests at a time
- HTTP keep-alive and SSL session caching

Performance, cont.

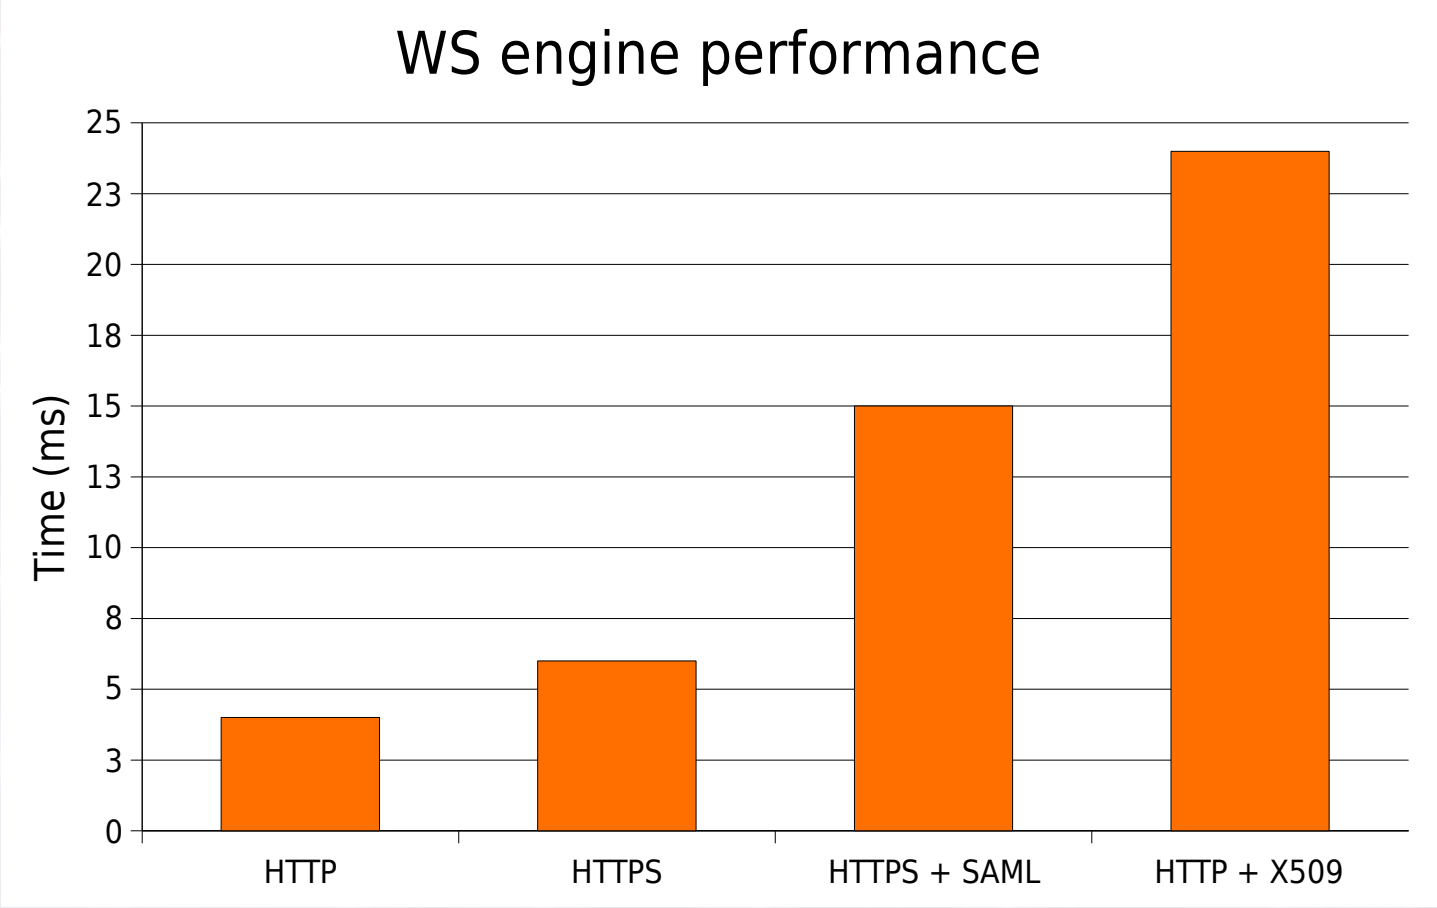

OpenDSP vs GT4 GRAM

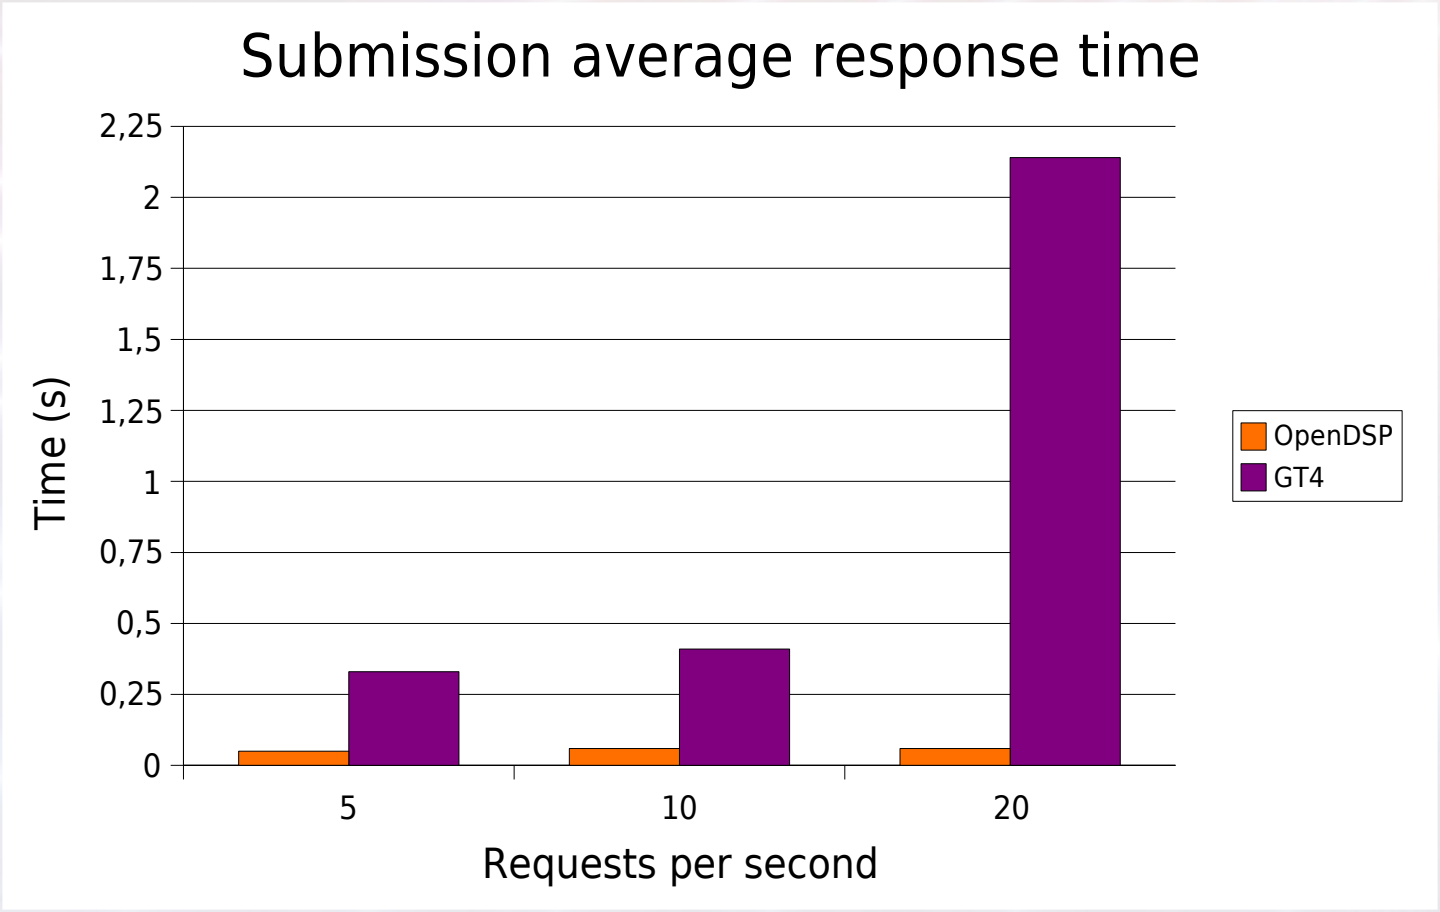

OpenDSP vs GT4 GRAM

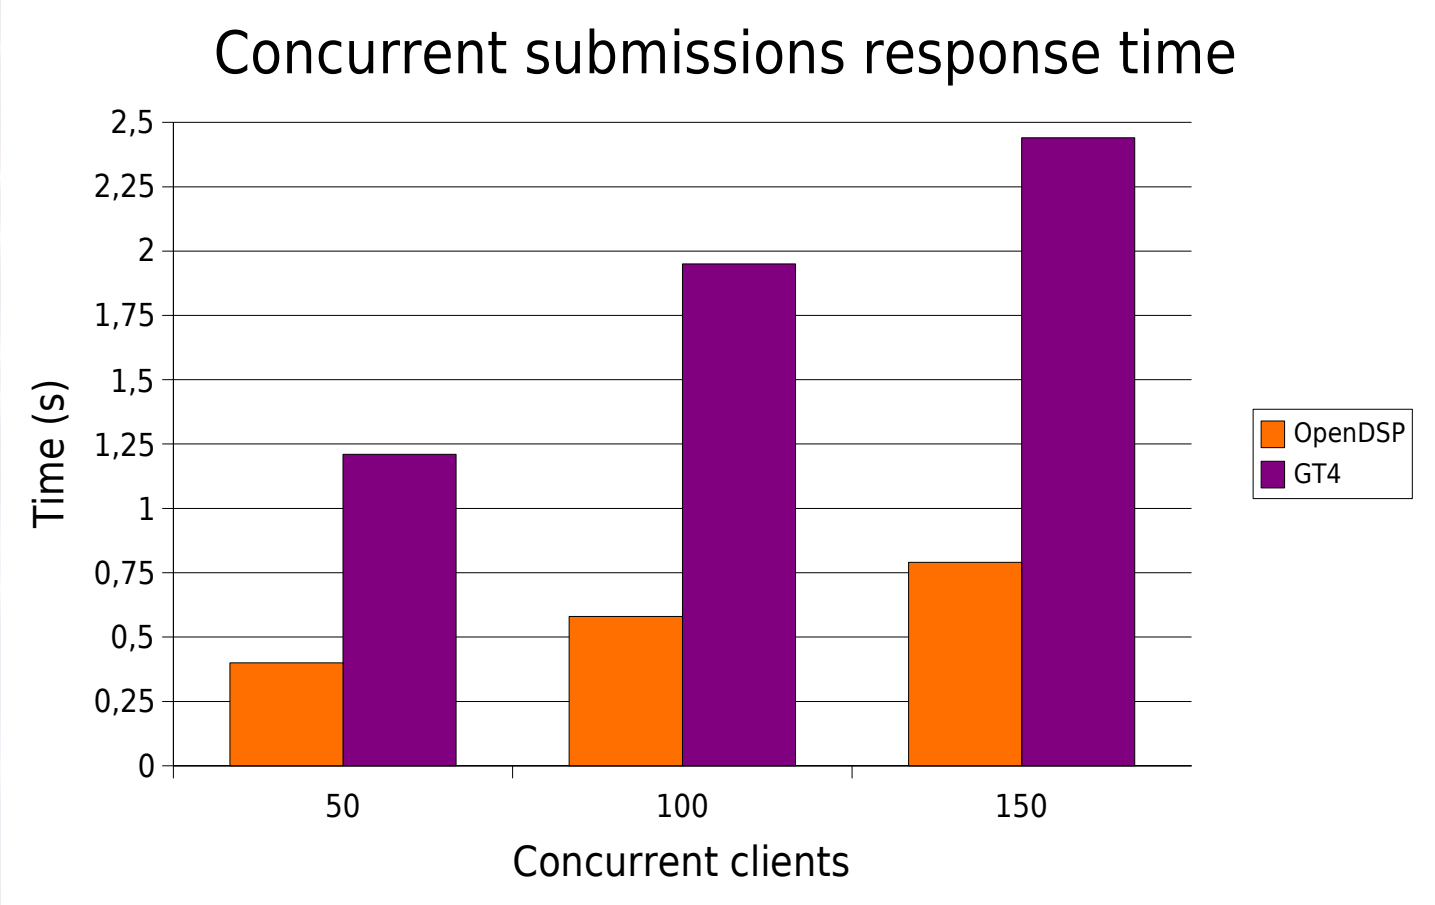

FedStage systems

OpenDSP was announced at **SunSource.net** in 2006 and since then we have developed many useful components and services around it.

We decided to create a spin-off company to support our open source products on the market globally. We now offer **professional customization, consulting**, **training** and **knowledge-transfer services.**

Making secure protected, integrated federated, and fast faster...

Our team has worked closely with Sun professionals for the last several years to optimize our technologies for use with Sun N1 Grid Engine. FedStage allows Sun N1 Grid Engine users to easily and securely expose standard based (Liberty Alliance, WS-I, OASIS, OGF, IETF, W3C) Web Services to offer utility computing capabilities.

OpenGridForum OPEN FORUM | OPEN STANDARDS

solaris

PayPal CERTIFIED Developer

SLA

www.fedstage.com

FedStage systems

A new wave of open source SOA and market-driven solutions for utility computing.

We provide a broad portfolio of:
- **software-as-a-service,**
- **federated computing services,**
- **Web 2.0 applications,** and
- **mobile products,**

based on **Sun Microsystems** products to meet today's demanding grid computing and cross-enterprise IT integration needs.

Our team helped Sun customers worldwide to use OpenDSP and its extended version called FedStage Computing Provider (FCP).
Part of FedStage Fedareted Service Bus, FCP is the only open source utility computing solution on the market which offers a full support for various AAA models, certified PayPal integration, Web 2.0 access and much more...

LIBERTY ALLIANCE PROJECT

Java

Sun microsystems

www.fedstage.com

Current work

- Integration with Sun products, e.g. Access Manager and Java Web Infrastructure Suite
- Resource requirements and reservations
  - New DRMAA version or use DRM dependent reservation interfaces?
  - Adapt JSDL specification
- OpenMPI integration
  - Adapt JSDL SPMD Application extension
  - Offer more capabilities than DUROC/MPICH-G2 for large scale experiments

Conclusions

- OpenDSP 1.0 available at SourceForge
- Next version with a lot of new features will be released in October, 2007
- If you are interested in utility computing open source solutions, contact us at

office@fedstage.com