OAR Documentation - REST API - Administrator's guide

Dedication: For OAR administrators whishing to set up the REST API

Abstract: OAR is a resource manager (or batch scheduler) for large clusters. By it's functionnalities, it's near of PBS, LSF, CCS and Condor. It's suitable for productive plateforms and research experiments.

BE CAREFULL : THIS DOCUMENTATION IS FOR OAR >= 2.5.0
PDF version : OAR-DOCUMENTATION-API-ADMIN.pdf

Table of Contents

Introduction

The OAR REST API is currently a cgi script being served by an http server (we recommend Apache) that allows the programming of interfaces to OAR using a REST library. Most of the operations usually done with the oar Unix commands may be done using this API from the favourite language of the users. But this API may also be used as an administrator portal as it provides you a convenient way to create resources, edit configuration variables or admission rules.

Installation

...To be written...

Authentication setup

The API authentication relies on the authentication mechanism of the http server used to serve the CGI script. The API may be configured to use the IDENT protocol for authentication from trusted hosts, like a cluster frontend. In this case, a unix login is automatically used by the API. This only works for hosts that have been correctly configured (for which the security rules are trusted by the admistrator). If IDENT is not used or not trusted, the API can use the basic HTTP authentication. You may also want to set-up https certificates. In summary, the API authentication is based on the http server's configuration. The API uses the **X_REMOTE_IDENT** http header variable, so the administrator has to set up this variable inside the http server configuration. Look at the provided apache sample configuration files (api/apache2.conf of the OAR sources or the installed /etc/oar/apache-api.conf of packages) for more details.