

# XSEDE Service Provider Advanced Integration Options

Last Revised October 19, 2017

<b>Introduction</b>	<b>1</b>
<b>XSEDE Resource Definition and Identifiers</b>	<b>2</b>
<b>Advanced Integration Options</b>	<b>3</b>
Infrastructure shared by multiple XSEDE resources	3
Resource Descriptions and Identifiers	3
RDR	3
xdresourceid	3
Allocations and Accounting	4
AMIE	4
gateway_submit_attributes	4
xdusage	4
Information Services	5
IPF	5
xdinfo	5
<b>About this Document</b>	<b>6</b>
<b>References</b>	<b>6</b>

## Introduction

Service Providers (SPs) integrating resources into the XSEDE federation must follow a series of [Integration Steps](#)<sup>1</sup> and [Install Software](#)<sup>2</sup>. This document provides SPs with special instructions for handling less common integration configurations.

# XSEDE Resource Definition and Identifiers

XSEDE federates a set of named infrastructure Base Resources that are composed of one or more Specific (Sub-)Resources. Specific (Sub-)Resources may be of type Compute, Storage, Grid, and Other. These resources are described in the Resource Description Repository (RDR).

Allocations and accounting usage information is tracked in XSEDE software using an *XDCDB ResourceName* identifier of the form “bridges.psc.xsede”. Resource descriptive and status information is tracked in XSEDE software using a *ResourceID* identifier in the form “bridges.psc.xsede.org”.

**Requirement:** All XSEDE Level 1, 2, or 3 federated resources, whether allocated or not, must have a unique *ResourceID*. In addition, all XSEDE allocated resource must have a unique *XDCDB ResourceName*. There must be a unique 1-to-1 correspondence between a *ResourceID* and an *XDCDB ResourceName*.

The *ResourceID* and *XDCDB ResourceName* identifiers are assigned in coordination by the XCI Service Provider Coordinator, XDCDB administrators, and information services administrators when a new resource starts the integration process. Both identifiers are derived from SP selected public resource names plus XSEDE or institutional DNS domains.

Currently active XSEDE resources:

<https://info.xsede.org/wh1/warehouse-views/v1/resources-xdcdb-active/?format=html>

The *ResourceID* is a user meaningful globally unique identifier often displayed along with resource descriptive information. It is used by both internal and external developers, such as gateways developers, to discover and access resource information in information services.

# Advanced Integration Options

## A. Infrastructure shared by multiple XSEDE resources

Typically, each XSEDE resource has its own compute and storage infrastructure, and its own login, data transfer, and other service nodes. However, when multiple XSEDE resources share infrastructure, SPs should describe their resources in RDR and install/configure software as summarized in this table and detailed below.

Instructions	Components
Describe each allocated resource as a separate Base Resource	RDR
Install once per allocated resource	xdresourceid gateway_submit_attributes IPF
Install once to support multiple allocated resources	xdusage xdinfo

### 1. Resource Descriptions and Identifiers

#### RDR

Per the [XSEDE Resource Definition](#) above, each allocated resource must have its own unique Base Resource in RDR. Information about shared hardware or services should be entered into all the relevant Base Resources and associated Specific (Sub)Resources. For example:

- A storage (sub)resource shared by multiple allocated resources would be described multiple times as Specific (Sub)Resources under all the relevant Base Resources.

SPs may NOT enter multiple allocable resources as separate Specific (Sub-)Resources under the same Base Resource. For example:

- If a single physical resource has both conventional and GPU hardware that are separately allocated, each allocated resource should have its own Base Resource

#### xdresourceid

The `xdresourceid` tool provides shell users and software a command line interface to access the *XDCDB ResourceName*, *ResourceID*, and other resource descriptors on that resource. For example:

```
# Executed on stampede2
```

```
$ xdresourceid -c
ResourceID="stampede2.tacc.xsede.org"
Name="stampede2"
SiteID="tacc.xsede.org"
Title=""TACC Dell PowerEdge C8220 Cluster with Intel Xeon Phi
coprocessors (Stampede)""
XDCDB_resource_name="stampede2.tacc.xsede"
```

SPs should install one xdresourceid package on each unique allocated resource either on a file-system shared by multiple nodes or local to each node. When nodes are shared by multiple allocated resources they must have multiple xdresourceid packages available. SPs should place one xdresourceid package in the user's default shell environment. When installing, configuring, or running software that relies on xdresourceid (such as IPF and xinfo below) the administrator should first place the appropriate xdresourceid package in the path.

## 2. Allocations and Accounting

### AMIE

With AMIE SPs can exchange transactions with XDCDB for managing XSEDE related accounts, allocations, and accounting records. SPs should deploy a **single** instance of AMIE for use by all their resources. Each AMIE instance can process transaction for multiple allocated resources and *XDCDB ResourceName* values.

### gateway\_submit\_attributes

The gateway\_submit\_attributes package provides science gateways a way to tell XSEDE which gateway user is associated with each batch job run under a community account. Each gateway\_submit\_attributes package is pre-configured to submit information for a specific allocated resource and *XDCDB ResourceName* value.

SPs should install one gateway\_submit\_attributes package on each unique allocated resource where science gateways can run jobs under community accounts. When login or batch submit nodes are shared by multiple allocated resources those nodes must have multiple gateway\_submit\_attributes packages, one per allocated resource. The SP must inform Science Gateway developers of the gateway\_submit\_attributes packages they can use for the allocated resource they are submitting jobs to.

### xdusage

The xdusage tool provides shell users a command line interface to view allocations information. Each xdusage package is configured with a default allocated resource identified by *XDCDB ResourceName*.

SPs may install a single `xdusage` tool on login nodes that are shared by multiple allocated resources and should configure it with a default *XDCDB ResourceName* of their choosing. Since `xdusage` has command line options to access information for alternate *XDCDB ResourceNames*, SPs should provide resource specific user documentation for looking up allocation usage information for resources other than the default resource configured in `xdusage`. For example, this command can be used on any XSEDE resource:

```
$ xdusage -r maverick.tacc.xsede
Project: TG-IRI160006/maverick.tacc.xsede
PI: Hazlewood, Victor
Allocation: 2017-06-28/2018-06-27
Total=1,000 Remaining=994.13 Usage=5.87 Jobs=105
  Navarro, John-Paul portal=navarro usage=0 jobs=0
```

### 3. Information Services

#### IPF

The Information Publishing Framework (IPF) package enables SPs to publish dynamic local software, services, and batch scheduler information to XSEDE's central information services. Each IPF package is configured to publish information about a specific *ResourceID*.

SPs must install one IPF package for each XSEDE integrated resource, allocated or not, and configure it to publish information about one *ResourceID*. SPs may install multiple IPF packages on a single node and configure each one to publish information about a different *ResourceID*. To do this the SP should place the appropriate `xdresourceid` tool in their path when installing, configuring, and running each IPF deployment.

Each IPF deployment is configured to run a set of workflows that gather and publish specific types of information. For example, the *computing\_activity* workflow publishes information about the batch scheduler, such as queues definition and node information, and the *computing\_activities* workflow publishes information about batch jobs. When an SP has a single batch system that manages queues, nodes, and runs jobs for multiple allocated resources, the *computing\_activity* and *computing\_activities* workflows will not be able to publish filtered subsets of information relevant to each allocated resource. Service providers can customize these workflows to filter and report separate information for each allocated resource.

#### xdinfo

The `xdinfo` tool provides shell users and software a command line interface to view software, services, resources, and outages information. An `xdinfo` deployment can show output about the local resource based on the `xdresourceid` found in the user's path. `xdinfo` can also be run with a command line argument to show information about any resource. For example:

```
# Executed on stampede2
% xdinfo services on bridges
```

Type	Name	ResourceID	URL
gridftp	org.globus.gridftp	bridges.psc.xsede.org	gsiftp://...
openssh	org.globus.openssh	bridges.psc.xsede.org	bridges.ps..

SPs may install a single xinfo package on nodes regardless of how many allocated resources those nodes support.

## About this Document

This document was created by the XCI Team.

Please send requests for additions, corrections, or clarifications to [help@xsede.org](mailto:help@xsede.org) with the Subject: Feedback on XCI Service Provider Integration documentation.

## References

<sup>1</sup> <https://software.xsede.org/sp-resource-integration-status>

<sup>2</sup> <https://software.xsede.org/sp-resource-integration-status-0>