

XSEDE Canonical Use Case 7: Subscribe for Resource Information

March 27, 2014

Version 1.2



Table of Contents

- A. [Document History](#)
- B. [Document Scope](#)
- C. [Related Use Cases](#)
- D. [Canonical Use Case 7](#)

A. Document History

Overall Document Authors:

Ian Foster
The University of Chicago and Argonne National Laboratory
Argonne, IL 60439
foster@anl.gov

Morris Riedel Jülich Supercomputing Centre
Forschungszentrum Jülich GmbH
D-52425 Jülich
Germany

Felix Bachmann
Carnegie Mellon University
4500 5th Avenue
Pittsburgh, PA 15213
fb@sei.cmu.edu

Andrew Grimshaw
University of Virginia
PO Box 400740
Charlottesville VA 22904
grimshaw@virginia.edu

David Lifka
Cornell University
512 Frank H. T. Rhodes Hall
Ithaca, NY 14853
lifka@cac.cornell.edu

Warren Smith
Texas Advanced Computer Center
10100 Burnet Road
Austin, TX 78758-4497
wsmith@tacc.utexas.edu

John-Paul Navarro

The University of Chicago and Argonne National Laboratory
 Argonne, IL 60439
 navarro@mcs.anl.gov

	Version	Date	Changes	Author
First use case draft	0.1	3/21/2013	Document created	Foster, Grimshaw, Hossain, Lifka, Riedel, Tuecke
Revised draft	1.0	9/28/13	Clean up formatting; Separate from UCCAN 8; prepare for archiving	Brown
Near final draft	1.1	11/13/13	Split 7,8 into 7,8,11,12; incorporate ADR reviewer feedback	Navarro, Smith
Final Draft	1.2	3/27/2014	Cleaned up some terms.	Smith, Grimshaw

B. Document Scope

This document is both a user-facing document (publically accessible) and an internal working document intended to define user needs and use cases that fall within the overall activities of XSEDE. The definition of use cases is based on a template from Malan and Bredemeyer¹. In general it is in keeping with the approaches and philosophy outlined in “Software architecture in practice.”²

This document is one component of a process that generates at least the following documents, some of which are user-facing, some are as of now intended to be internal working documents:

- ***This document*** - A description of use cases [User facing]
- A set of level 3 decomposition documents, which include:
 - Quality Attributes descriptions
 - Connections diagram in UML

The use cases are presented here using the following format, derived from the Malan and Bredemeyer white paper¹ as follows:

Use Case	Use case identifier and reference number and modification history
<i>Description</i>	Goal to be achieved by use case and sources for requirement
<i>References</i>	References and citations relevant to use case
<i>Actors</i>	List of actors involved in use case
<i>Prerequisites (Dependencies) & Assumptions</i>	Conditions that must be true for use case to be possible Conditions that must be true for use case to terminate successfully
<i>Steps</i>	Interactions between actors and system that are necessary to achieve goal
<i>Variations (optional)</i>	Any variations in the steps of a use case
<i>Quality Attributes</i>	

¹ Malan, R., and D. Bredemeyer. 2001. Functional requirements and use cases. www.bredemeyer.com/pdf_files/functreq.pdf

² Bass, L., P Paul Clements, and Rick Kazman

<i>Non-functional (optional)</i>	List of non-functional requirements that the use case must meet
<i>Issues</i>	List of issues that remain to be resolved

C. Related Use Cases

Canonical Use Cases (UCCAN) 7.0, 8.0, 11.0, and 12.0 reflect two different forms of information system: first an asynchronous publish subscribe mechanism in which consumers can subscribe to topics and publishers can publish on topics, and second a form akin to a centralized registry or database in which publishers synchronously update information that can be subsequently queried by consumers. As a consequence asynchronous pub/sub systems are typically used when high volume message transport is required and data validation is less important. In these systems accepted messages are not immediately available to subscribers. By contrast, synchronous add/update/query systems are typically used when information validation is more important and publishers need confirmation that information was successfully updated and is available immediately to users.

Examples of information that may be available thru publish subscribe:

- Job status and accounting information
- Operational test pass/fail information
- System load and dynamic monitoring information

Event information

D. Canonical Use Case 7

Use Case UCCAN 7.0	Subscribe for Resource Information
<i>Description</i>	Subscribe for resource information from the information system, with asynchronous, best-effort delivery.
<i>References</i>	Patrick Th. Eugster, Pascal A. Felber, Rachid Guerraoui, and Anne-Marie Kermarrec. 2003. The many faces of publish/subscribe. <i>ACM Comput. Surv.</i> 35, 2 (June 2003), 114-131. DOI=10.1145/857076.857078 http://doi.acm.org/10.1145/857076.857078
<i>Actors</i>	Subscriber: registers interest in resource information with the information System in order to receive information Information System: allows Subscribers to register interest in resource information, receives published resource information, delivers information to Subscribers
<i>Prerequisites (Dependencies) & Assumptions</i>	<ul style="list-style-type: none"> ● The Subscriber knows how to contact and communicate with the Information System. ● The Information System is an abstraction for different possible information delivery mechanisms including direct or via intermediate services. ● When required the Subscriber authenticates with the Information System and is authorized. ● Authentication and Authorization are consistent with UCCAN 9.0 and the relevant security architecture. ● Subscribers must not rely on receiving information published to the Information System prior to their subscriptions. ● Subscribers must be able to communicate securely with the information system when they wish to. ● Subscribers will comply with XSEDE wide information policies.
<i>Steps</i>	<ol style="list-style-type: none"> 1. The Subscriber registers interest in resource information with the Information System, typically including a description of what information is of interest, how the information should be delivered, and if (and for how long) information should be buffered if the Subscriber is not ready for delivery 2. At some later time, the Information System has resource information of interest to the Subscriber: <ol style="list-style-type: none"> a. If the Subscriber is ready for delivery, the Information System delivers the information to the Subscriber. b. If the Subscriber is not ready for delivery, has requested that information be buffered, and the Information System supports buffering, the Information System buffers the resource information for later delivery. 3. At some later time, the Subscriber becomes ready for delivery of information.

	<ol style="list-style-type: none"> a. The Subscriber informs the Information System that it is ready. b. If the Information System supports buffering, it delivers the information that it has buffered for the Subscriber. c. The Subscriber optionally informs the Information System that it is no longer ready to receive information. <p>4. The Subscriber unregisters interest or the subscription times out.</p>
<i>Variations (optional)</i>	<i>Related use case UCCAN 8.0 describes an alternate approach for obtaining resource information. Information obtained using this use case may not be available thru variants.</i>
<i>Quality Attributes</i>	<ul style="list-style-type: none"> • The rate at which the Information System can deliver information to a single Subscriber must be: at least 10 messages per second [source: A&D] • The minimum aggregate rate at which the Information System can deliver information to all Subscribers: 150 messages per second and 3 MB/sec [GLUE2 and Inca running on XSEDE and OSG] (duplicates one in 7.a) • Subscribers can contact the Information System 99.9% of the time [source: A&D] • Messages are delivered to subscribers without modification (integrity) 99.9% of the time [source: A&D] • The amount of time from receipt of resource information by the Information System to receipt of that resource information by active Subscribers is less than 1 second. [source: A&D] • Information buffered in the Information System will be available for delivery to Subscribers within the configured buffering rules (e.g. maximum buffer length or maximum buffer duration) 99% of the time unless it is explicitly deleted. [source: A&D]
<i>Non-functional (optional)</i>	<i>Publish-subscribe software is easy to install and support. [source SPs]</i> <i>Publish-subscribe interfaces are simple and well documented. [source SPs]</i>
<i>Issues</i>	