

XSEDE Operations Hybrid Cloud Plan

1 November 2018

Version 1.2




Table of Contents

A.	Document History	iv
B.	Document Scope	v

List of Figures

Figure 1: XES Hybrid Cloud..... **Error! Bookmark not defined.**

A. Document History

Relevant Sections	Version	Date	Changes	Author
Entire Document	1.00	08/20/2011	Initial draft	V.Hazlewood
	1.1	10/9/2018	Updated based on Greg feedback and Figure 1 update after discussion with Ester	V. Hazlewood G. Peterson G. Rogers
Section C	1.2	11/1/2018	Updated based on J. Towns comments	V. Hazlewood

B. Document Scope

This document covers the XSEDE2 hybrid cloud plan for the XSEDE Enterprise Services for XSEDE2 program years 8 to 10 (PY8 to PY10) operated by the XSEDE Operations Systems Operations (SysOps) group.

C. XSEDE Enterprise Services

XSEDE is a virtual organization that supports the allocation, use, and evolution of a national cyberinfrastructure (CI) made up of a variety of resources and service providers funded by the National Science Foundation. Like many organizations an enterprise information technology (IT) infrastructure exists to support the mission of the XSEDE organization. XSEDE enterprise services (XES) are comprised of approximately 50 information technology services which includes an XSEDE enterprise database called the XDCDB, an RT ticket system, a website, a user portal, and a Confluence-based wiki site, to name a few. The XES are operated on a variety of information technology resources located at University of Illinois (at NCSA), Indiana University (on the Jetstream resource), University of Tennessee (at NICS), University of Texas (at TACC), Carnegie Mellon University (at PSC), and in the public cloud provided by Amazon Web Services (AWS). Each of the data centers and IT environments provided by each of these partners and providers are world-class. By world-class we mean that the power, space, cooling and physical security environments are state-of-the-art, networks provided are at 100 gigabit per second or greater and have contingency backup network capabilities, that the IT components are 5 years old or less with vendor maintenance provided, and that professional staff are the custodians and administrators of these components.

In XSEDE1 most XES were operated on compute and storage hardware purchased by the partner site with XSEDE funds to instantiate a VMWare virtual machine (VM) environment that can initiate, operate, and decommission VM servers that would then run any individual XES. These environments operated essentially separate from each other with primary and backup VM services that provided the role of any individual XES. The current XES VMWare infrastructure for XSEDE does not provide for failover capabilities between partner site data centers. VMWare has features that provides high availability and fault tolerance where if a server fails, the affected VM is restarted on other available servers in the VMWare cluster. There are no VMWare clusters that cross partner data centers across the WAN.

In the case of XDCDB, there was a primary, secondary, and tertiary XDCDB XES operated at three different partner sites. In XSEDE2 PY7, a pilot project was initiated to evaluate moving and operating the XDCDB in the public cloud using AWS. This pilot project was successful and the XDCDB was moved to AWS and the primary XDCDB XES is now provided by AWS. The XDCDB XES failover is provided by AWS via automated failover AWS cloud technologies. The XDCDB was chosen because of its high-availability requirements due to many service dependencies on the XDCDB. Not all XES meet the high-availability need threshold, and therefore, not all XES are candidates to move to the cloud. Evaluation of XES suitability to move to a VMWare cluster or to the cloud will continue. Cost and security implications of running XES in the cloud at AWS compared to running VMs for XES on partner-housed equipment using VMWare clusters in their data centers shall continue to be evaluated.

D. XES Hybrid Cloud

XSEDE Operations will transition XES to a hybrid cloud model for PY8 to PY10. The hybrid cloud model will be made up of two major components: the XSEDE private cloud which will be geographically dispersed across the XES partner sites and the AWS public cloud used for critical XES that need high-availability and has significant dependencies with other XES. The XSEDE provide cloud for XES will be made up of geographically diverse XES locations running VMWare with those sites participating in an XSEDE VMWare cluster. The XES deployed and used in the public cloud will be determined by the SysOps Manager and approved by the Operations Director and shall be XES that have high-availability requirements along with significant dependencies on other XES. The XSEDE public cloud will be made up of VMWare components at University of Illinois (at NCSA), Indiana University (on the Jetstream resource), University of Tennessee (at TACC), University of Texas (at TACC), and Carnegie Mellon University (at PSC) shown in Figure 1. Some XES, such as security-related XES, will remain on dedicated server hardware and will not participate in the XSEDE private cloud.

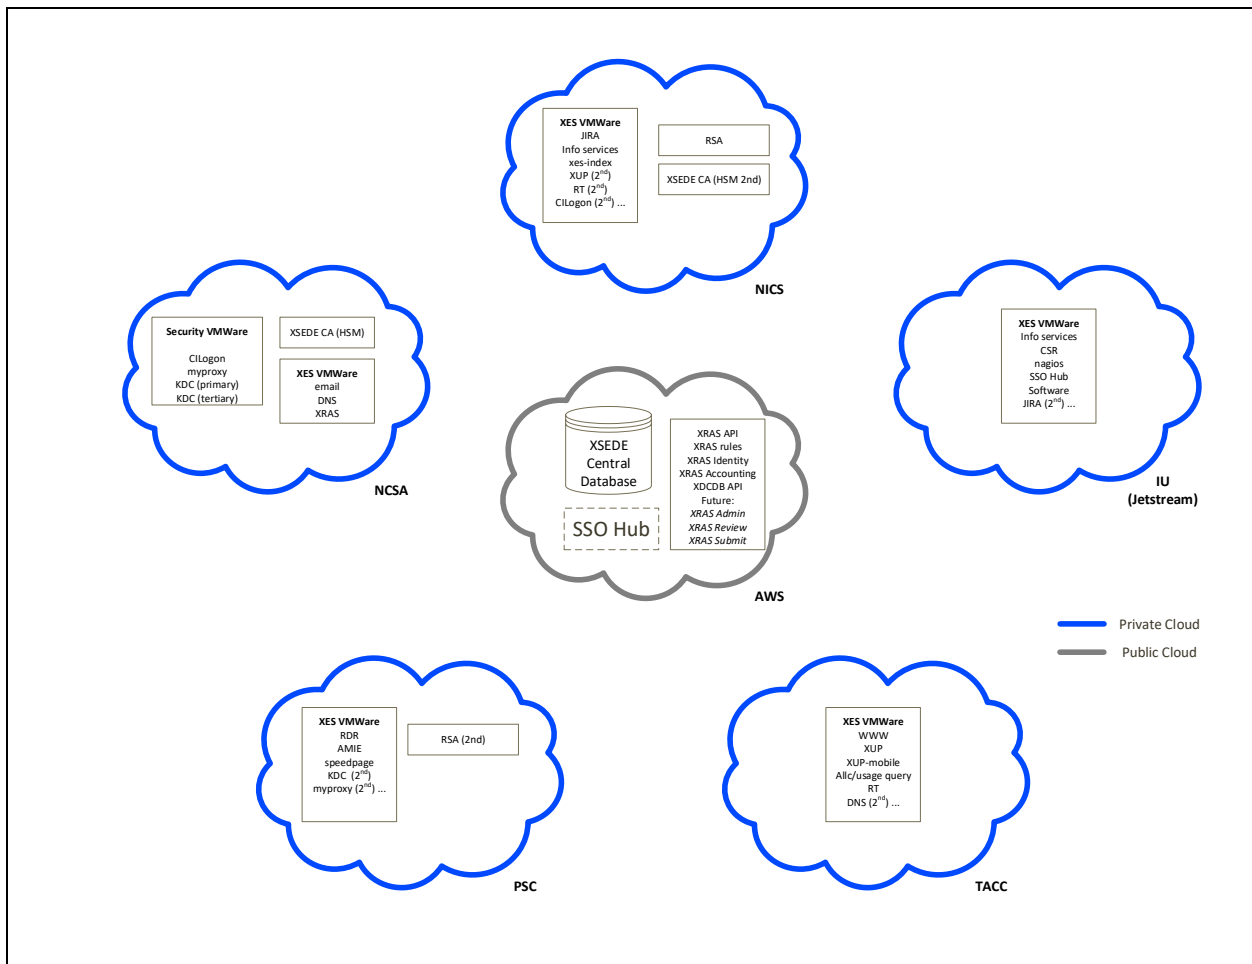


Figure 1: XSEDE Hybrid Cloud