TEAL2.0
TECHNOLOGY ENABLED ACTIVE LEARNING

TEAL2.O INFRASTRUCTURE

CONTENTS

This document clarifies the infrastructure solutions for TEAL2.O. The intended audience is the key stakeholders – funders and the founding institutions, as well as the developer teams.

It is recommended to keep the TEAL2.O infrastructure that is used across development, testing and integrated deployment consistent across all participants, across all regions and across all institutions. The TEAL2.O project funding proposal envisages purchase of on-premises support infrastructure including the following hardware components:

Description	Number of Units	Where	Specifications
Development Servers	3	Vellore Institute of Technology, India	Processor: 3.0GHz, 8-Core Intel Xeon E5 with 25MB L3 cache and Turbo Boost up to 3.9GHz
			Software: Mac Server software
			Memory: 16GB (four 4GB) of 1866MHz DDR3 ECC memory Storage:256GB PCI e-based SSD
			Graphics: Dual AMD Fire Pro D700 graphics processors with 6GB of GDDR5 VRAM each Display
			Support: Three dual-cable 5K display, Six thunderbolt display
			Connections: USB3-4, Thunderbolt 2- 6, Dual Gigabit Ethernet, HDMI .4Ultra HD
			Input: Wired Keyboard, Wired Mouse End Point Protector (device control, file tracing, device logging, USB lockdown, DLP)
			Operating System: Mac OS Sierra
Storage Servers	3	University of Peradeniya and University of Moratuwa, Sri Lanka Indian Institute of Information Technology Allahabad	Eight-core Intel Xeon processor. (Xeon Silver 4110), 64GB system memory, RAID 0/1/5/6/10 with controller support up to 8GB cache,
			Hybrid storage supporting SLC, MLC & TLC flash, Storage Pool should be configured with not less than 5TB usable capacity in RAID 5 using with 1.2TB 10K SAS HDDs. (Hot spare drives are Mandatory).
			R/W ratio – 70/30 Block Size – 8K, support Async replication and related licenses should be included can replicate 32 volumes per system, HTML5 Web GUI, CLI & REST, support integration with commonly used virtualization and data archival solutions (e.g. VMware vSphere (ESXi, vCenter); SRM, Microsoft Hyper-V)

Backup Servers	I	Asian Institute of Technology, Thailand	To be determined

The indicated infrastructure should be taken into account during the development process in order to avoid hardware conflicts. This on-premises support infrastructure should be able to meet the requirements of the TEAL2.O platform for up to 10 years starting from the moment when the platform is deployed.

Longer-term solutions may also include using the services of a public cloud provider. It is not expected that a switch to a public service provider would impose any constraints on the technological solution and so it will not affect the development process. The choice of long-term support infrastructure can be made in the course of exploitation.

>> Back to contents

Why choose a Public Cloud Provider?

In the long run, priorities may shift and it may become more reasonable to not spend valuable time in maintaining support infrastructure as such efforts do not provide any value towards expanding the TEAL2.O network. It may also be recommended not to spend valuable time in setting up application servers, databases and keeping the software and operating system updated as again none of these activities contribute towards expanding the TEAL2.O network. If priorities are so realigned, it may be recommended to just focus on functional code and leave all non-functional aspects to an external provider of support infrastructure.

Based on the comparison provided below, the following long-term infrastructure options exist: Public cloud IaaS (Infrastructure as a Service), PaaS (Platform as a Service), SaaS (Software as a Service) and FaaS (Function as a Service).

Which Cloud Provider to Use?

There are several factors to be considered before deciding on a preferred Cloud Provider:

1. Completeness of the vision of Cloud Patterns
2. Ability to execute and keep up with the latest trends
3. Provision of complete coverage of services during development and beyond
4. Provision of quick starts, patterns and solutions out of the box
5. Provision of consistent management of Cloud resources regardless of whether the resources are deployed via portal, command line or SDK
6. Provision of a lightweight contained environment construct that can be used to manage separate budgets across teams, separate components and role-based access for each component
7. Provision of a consolidated hybrid environment where Cloud and on-premise components can be easily integrated, should we chose to go in that direction.

As of the current status-quo, out of all Cloud Providers, obviously AWS, Microsoft Azure and Google Cloud are the market leaders. In terms of (1) and (2) above AWS and Microsoft Azure leads well ahead of Google Cloud according to Garners Report as IaaS Provider. Both AWS and Azure provide (3) and (4) in a comparably matured way, so these requirements are not significant when choosing between AWS and Azure. For (5), (6) and (7), Microsoft Azure leads the way and AWS is still catching up. In (5) for example, ARM (Azure Resource Manager) is architected well to suit all type of deployment, whereas AWS's Cloudformation only works if the initial deployment is done via Cloudformation; thus, for new users of AWS this creates steep learning curve. In (6), Azure Subscriptions and Azure Resource Group provide a lightweight construct to manage separate budgets and separate contained environments, while the AWS account provides the same but comes with the overhead of managing multiple accounts. In (7), Microsoft born on-premise and now focusing on moving to Cloud, the ecosystem in Azure caters for most of the on-premise scenarios and even provides the capabilities and services to host applications

and databases on-premise and expose them securely into Azure environments. AWS on the other hand is born in Cloud and is hence focusing only on migration from on-premise to Cloud rather than on running Hybrid solutions.

On the basis of these considerations, the more reasonable choice would be Microsoft Azure – with an account for the entire project/platform, a subscription for each partner so that it can manage its budgets and a resource group for each component in the component test environment. However, since we are discussing a long-term solution and technological options are quick to change, a similar analysis needs to be done around the 6-7 year mark after the deployment of the TEAL2.O platform.