

TIER Grouper - Docker Reference Implementation

TIER Reference Implementations are designed to enable rapid evaluation of a component or set of interconnected components and provide a starting point for a full campus deployment. Reference Implementations are developed for Docker Swarm but dependencies are minimized within the containers in order to facilitate the use of other container orchestration mechanisms.

Introduction

Like the other TIER distributions for larger production environments, the TIER Grouper solution will be built using Docker Swarm mode. TIER will provide generic instructions and pointers on setting up the Docker environment for this solution but will not provide a pre-built virtual machine. Schools needing a vm-based solution should focus on the TIER Grouper Appliance instead. This deployment specification may be implemented in stages and all requested functions may not be available in the initial release.

1. Background & Requirements

- a. The migration of a Grouper installation to new platform typically requires:
 - i. Full database backup and restore.
 - ii. Filesystem configuration files backup and restore
 1. Hopefully sites have some form of configuration management in place. How to leverage this in a TIER context is a question.
- b. High vs. Standard Availability
 - i. While we believe that the majority of schools operate Grouper without High Availability, some schools do operate the web services components of Grouper in HA mode.
 - ii. The production TIER-version will account for high availability.
- c. Expandability
 - i. Users will need the ability to install additional modules
 - ii. [Provisioning and Integration](#)
- d. Consider the use of [supervisord](#) for multi-process containers. We are using this in the COmanage TIER implementation and the COmanage team has used it in other projects.
- e. Scripting will exist, when practical, to guide users through the startup process. Documentation can cover scripting gaps, especially in early versions. The scripting is focused on Grouper and not building the Docker Swarm environment.

2. TIER Production Distribution Components

The Grouper components listed below will be included in the TIER distribution at the indicated availability levels.

- a. Grouper Database (HA)
 - i. The database may be supplied by the campus.
 - ii. A TIER MARIADB implementation will be available.
 - iii. Configuration options to either use the TIER-provided database or (perhaps manually) configure a campus database will exist.
- b. Grouper Web Services (HA)
- c. Grouper user interface (HA)
- d. Grouper message bus to AMQP (soon) HA
 - i. TIER has selected RabbitMQ for the reference build.
 - ii. We need to nail down the reference AMQP implementation once implementation starts.
- e. Grouper PSPNG (HA) for LDAP provisioning
- f. Shibboleth SP for authentication to Grouper UI (HA)

3. Configuration Defaults and Options

- a. Web Services Authentication
 - i. Apache basic
 - ii. LDAP
 - iii. future: certificate
- b.

4. Implementation

This section contains the preferred implementation structure for the application.

- a. Containers
 - i. Grouper UI Container
 1. Grouper UI
 2. Apache
 3. Shibboleth SP
 - a. Preconfigured for InCommon; simple scripting
 - ii. Grouper Web Services Container
 - See also 3.a above
 1. Grouper Web Services
 2. Apache
 3. Shibboleth SP
 - iii. Database
 1. TIER-provided MARIADB
 2. Ability via documentation or scripting to use campus-provided database.
 - iv. Grouper Loader
 1. Grouper Loader
 2. Just needs Java
 3. Include jars for: PSPNG, AWS, RabbitMQ, ActivMQ
 - v. Load Balancing
 1. Docker swarm native vs. haproxy
 2. Sticky sessions are needed for the Grouper UI
- b. Misc
 - i. All containers need to run ntp
 - ii. Externally exposed ports

1. Goal: 443 for UI only
 2. Can use second port for web services
 3. May be some other ports in the future
- iii. Logging
 1. Goal: send logs to Docker stdout
 - a. Swarm log aggregation
 - b. Logs will be formatted as per the TIER container logging definition.
 2. Short term: potentially logging to files in containers.
- c. Default Configuration
 - i. Goal is to include default settings based on the Grouper Deployment Guide.
 - ii.