

Could Samvera Use an Infrastructure Working Group?

Samvera Connect 2019

Facilitator: Michael Klein

Introduction

- Forming a WG is very likely to happen
- What types of topics do we want to see discussed?
- What would we like to see maintained?
- There is much duplication when it comes to Docker containers for certain infrastructure services

Defining Infrastructure

Primary Elements

- Deployment
- Runtime
- Dev. Practices
- Monitoring
- Observability
- Metrics
- We should include...
 - Anything that is required for the application to run in a particular environment
 - Everything from environment variables, documentation, containers, CI scripts...
 - Anything in the deployment and runtime environment
 - Much larger in scope than application
- There is also overlap with software architecture
- Should we address files as if they are accessed over an API or on a shared file system?
- Cannot assume that there is a /tmp subdirectory (it might not be there)

12 Factor is a good starting point

- <https://12factor.net/>

- List of key practices which promote clean separate between an application and its environment
- Number of attendees confident in the current state of their infrastructure: 0

Monitoring and Observability

- Defining the term
- Monitoring is watching things which are already easy to watch
- Observability is a series of practices around ensuring that your application is easy to observe, to debug on a scientific basis (rather than a reactive, alert-based basis)
- Building in the exposure of information

- Monitoring: Watching what is running
- Observability: Are your pieces working together (and what are they doing? and if they aren't working, why not?)
- Examples:
 - Monitoring: Pinging the home page of the app.
 - Observability: Background jobs are created but failing, but the servers aren't crashing

- Example with AWS:
 - AWS Neptune: Cannot talk to the world
 - AWS Lambdas need to be able to communicate with a firewalled/isolated Neptune
 - The Neptune may be up and this may be confirmed using monitoring
 - ...but observability would ensure that the Lambdas can communicate with Neptune

- Monitoring Tools:
 - Nagios
 - NewRelic

- Observability:
 - Build in application hooks to an application for a monitoring database
 - Use these types of approaches will permit the measurement of throughput
 - Without these approaches, observability tends to be passive (greeting through log files)
- Observability Utilities and Services:
 - Honeycomb

Working Group Organization

- There will be different levels of deliverables
- Documentation specific to languages and systems which the majority of people are using
- There will also be concrete deliverables, such as Docker containers for core components (and how to use these with other Samvera pieces)

Assessment

- How many attendees are using on-prem deployment?: 7
- How many are dependent upon Library IT?: 6
- How many are dependent upon Central IT?: 1
- How many are using cloud-based deployment: 6
- How many on AWS: 6
- How many are on Google: 2

Community Assessment

- Sharing deliverables with organizations with the developer(s) lacking DevOps engineers as resources would be beneficial
- Environment
 - This becomes infrastructure when parity between development/staging/testing and production is maintained
 - Some create releases daily and deploy them to production
 - However, for their workflow, they don't need the development environment to align exactly with the production environment
- Proposed Best Practices
 - Have available a Docker Compose which includes the services which are available within a production environment (e. g. Solr)

DockerHub Organization for Samvera

- Tom Johnson, Michael Klein, and Rob Kaufman are administrators
- Are there other third-party repositories in common use where we should claim the Samvera common namespace
- Ansible

- Very popular, but some are trying to move away from using this for server provisioning
- Registering with Ansible Galaxy might be worth exploring
- Proposed Charter Objective: At least one member volunteers to maintain the registry of members and projects for Docker Hub, Ansible Galaxy...

Lifespan of the Working Group

- Continuing this Working Group may or may not follow the approach undertaken by the Core Component Maintenance Working Group (with different phases, each being renewed with new membership and an updated charter)

Samvera Infrastructure Working Group

Call for Participation

Michael and Brian McBride are interested in leading this

Who is using or desires to use:

- Docker
- Kubernetes
- Terraform
- Puppet
- Ansible
- Chef

There should be a survey directed towards the community assessing this

- Will there be costs associated with participating?
 - Docker Hub costs when using private repositories
 - Cloud providers also charge an hourly rate
 - Monitoring and Metrics: There are some services which do not charge
 - Investigating this may be something which the WG does

We should also aim to align with the Documentation Working Group

- Are there other Interest or Working Groups which intersect with this domain? *The consensus was no, although there was a DevOps Interest Group which became inactive in 2017*
- #devops Channel is very inactive as well

Will this WG focus on Hyrax?

- It is good to use Hyrax as an example for deriving use cases, but we should aim to generalize beyond these
- What of Nurax?
 - In collaboration with Notch8 and Data Curation Experts, Tom Johnson has been exploring using Docker deploys with helm (Kubernetes) to assist with the deployment of Nurax
- Proposed Approach for addressing Hyrax:
 - Try and have this WG focus upon generic solutions and avoid anything specific to solution bundles like Hyrax
 - Hyrax developers and maintainers should be treated as stakeholders
 - (This is the approach undertaken by the Core Component Maintenance Working Group, where Hyrax issues are not scoped for maintenance sprints, but Tom Johnson attends and contributes as a member regardless)

Cloud-Native Deployments

- Well-documented Dockerfiles provides a human-readable guide for deploying services even if the Kubernetes or Docker infrastructure is not available (e. g. guidance for Ansible Roles/Playbooks, Puppet Manifests...)

Action Items

- Draft a Charter
- Schedule an initial meeting
- Use #devops on Slack (?)
- Documentation
- Assess the state of infrastructure for organizational members within the community
- Converging on a set of common Docker-based solutions (and deliver these as concrete solutions)

Session concluded at 10:37CDT