

# Repository Management with Pulp

*Presented by:*

Michael Hrivnak

*Principal Software Engineer*

mhrivnak@redhat.com

@michael\_hrivnak



<http://www.pulpproject.org>



Distributing software is a messy business.



# What is Pulp?

- Manage repositories of content
- Supports many content types
- Pull-through cache
- OSS under GPL2
- <https://github.com/pulp>
- Python web app



# Create New Repositories

- Create a new repository in Pulp
- Sync content from a remote repo, manually or on a schedule
- Upload your own content
- Copies are cheap



# Publish Your Repository

- Mix and match content, then publish
- Publish can mean many things



# DEMO



# Content Types

- RPM family
- Docker Images
- Puppet Modules
- Python Packages
- OSTree
- Regular Files
- Debian Packages (community)
- NPM?



# Who uses Pulp?

- Red Hat Release Engineering
- Public Clouds
- Katello / Red Hat Satellite 6
- Community






# Use Case: Mirror Python Packages

- Sync packages from PyPI to a local repository
- Add or remove packages as desired
- Retain old versions
- Control which versions appear together



# Use Case: Dev / Test / Production

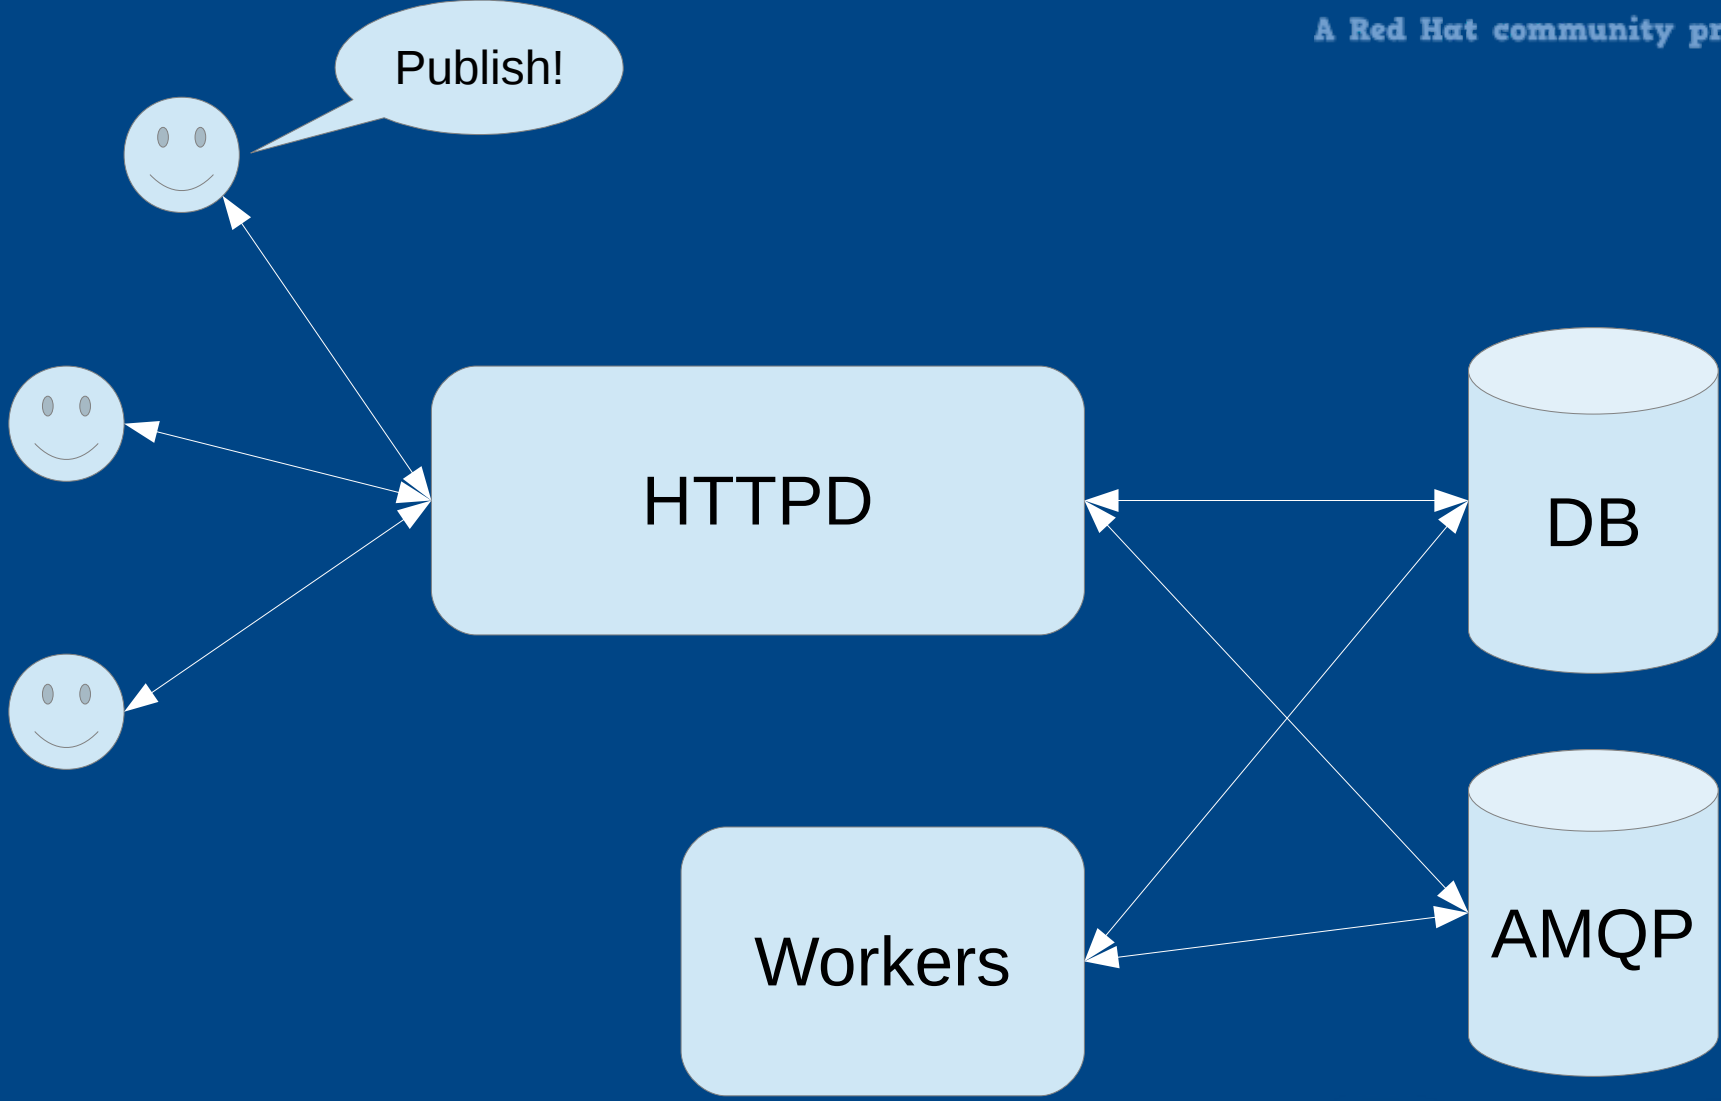
- Sync content into a development repository
- “Promote” by copying to a Test repository, and then to a Production repository
- Useful for testing upstream content, such as new RHEL point releases



# Distributed Application

- REST API
- Content served via HTTP
- Worker Processes for Async Jobs





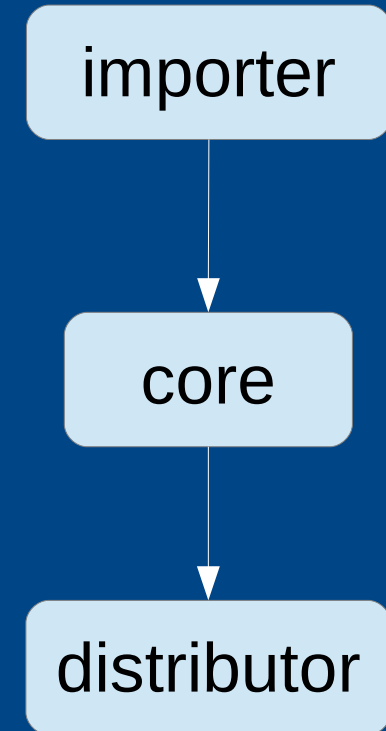
# Extensible

- Standard ways to support new content types
  - How does content flow in?
  - How does content flow out?



# Plugins

- Type Definition
  - Model Class (using mongoengine)
- Importer
  - Used to pull in upstream content
  - One per repository
- Distributor
  - Used to publish content
  - Many per repository



# CLI Extensions

- Hierarchical layout
- Many generic commands available for re-use and customization




# Integration

- REST API
- Events published to AMQP topic exchange
- HTTP callbacks





# Pull-through Cache



- Available in 2.8.0 Beta
- Adds metadata to database without downloading files
- Retrieves files on-demand or in the background



# Consumer Tracking

- Consumer machines can register
- Report list of installed packages
- Pulp calculates which machines need updates



# Documentation

<http://www.pulpproject.org/docs/>



# Stickers!



# Questions?



Michael Hrivnak  
mhrivnak@redhat.com  
@michael\_hrivnak

<http://www.pulpproject.org>

