Contribution ID: **203** Type: **Presentation**

KubePie: Leveraging Kubernetes for Scalable End-User Data Access

Wednesday 19 March 2025 12:30 (15 minutes)

Does the idea of running and managing "own"web server for every user with the corresponding user/group IDs and securing access to it sound crazy? Definitely 10 years ago it was, but not in the current Cloud Native world. KubePie facilitates streamlined access to end-user data by harnessing Kubernetes' scalability and deployment capabilities to actually running, managing and securing web-servers on a large scale.

Inspired by the success of JupyterHub on Kubernetes, KubePie implements *Permissions Impersonated Environments ("Pies")*, constructed from open-source components to provide a personalized data consumption experience for end-users through any web browser or HTTP client.

The microservices framework within KubePie is structured around various "Pies":

- **PieData:** The personal web-server Pie, stuffed with data services. Wrapped in a nice Helm chart.
- **PieTrack:** Transfer activities keeping for PieData. Tracks your calories consumption over time.
- **PiePass:** Ensure secure passage to the KubePie bakery services for the customers.
- **PieCut:** The "Control Unit" of KubePie bakery, taking orders and delivering PieData.
- **PieDeck:** The "Data Endpoints Controller for Kubernetes" operates the baking process of the PieDatas for end-users.
- **PieSec:** Admission webhook for the security seasoning of the PieData. Adds the UID/GID and SGID spices. Make sure no rotten eggs are added.
- **PieEat:** "Erases at Timeout" the stale PieDatas.

PiePass serves as the authentication and authorization entry point for end-users, relying strictly on OIDC/OAuth2 protocols and leveraging Apache web servers, equipped with mod_oauth2 and mod_auth_openidc modules from OpenIDC.

PieCut is a main action point for end-users. Technically it is Kubernetes Controller that creates PieData Custom Resource Definitions (CRDs), based on authenticated user information from HTTP headers set by PiePass.

PieDeck is a Helm-based Operator using OperatorSDK to manage PieData deployments via PieData CRDs. It is a central configuration point of common PieData settings (e.g. Images, Resources).

PieSec is the Admission Controller for Kubernetes providing a securityContext based on the JWT claims information and a pluggable mapping backend config (static mapfile and LDAP are implemented). Keeping PieSec separate (compared to defining securityContext by PieCut) improves security, UID/GID comes as an enforcement not as a request and is separated from the main Pod spawning logic.

PieData is taking care of servicing the main data traffic. End-user is redirected to a PieData endpoint via the proper Ingress configuration and uses Apache with OpenIDC/OAuth2 modules as well to authenticate the end-user. Optionally, ephemeral stateless credentials for WebDAV can be generated as well.

PieData is capable of carrying multiple containers inside the same Pod to enhance the functionality. In particular, data analysis services, such as HDF5 viewer or even JupyterLab, can be provided. Thanks to PieSec, all PieData services are running with restricted permissions and storage access.

PieTrack is essentially a custom Prometheus exporter for Apache. This way KubePie offers both, generic observability and an interface for PieEat. With Prometheus metrics it is easy to define a dashboard for data transfers monitoring, including the per-user traffic statistics.

PieEat is another rather simplistic Kubernetes Controller that acts on data-transfer inactivity and removes unused PieDatas. This contributes to both efficient resource utilization and security. KubePie framework is designed to run services on demand, avoiding unnecessary infrastructure usage by idle services.

The framework deployment at MAX IV, illustrates the utilization of KubePie for scientific data access with integrated HDF5 viewer. It runs on a bare-metal Kubernetes cluster with native SpectrumScale data access, Keycloak SSO for authentication and LDAP for user mapping, delivering high-speed data transfer rates via standard HTTPS protocol.

Another use-case KubePie helps to solve at MAX IV is on-demand OpenVSCode Server instances, mounting user own “home”directory, providing easy IDE and web-terminal experience.

Authors: SALNIKOV, Andrii (Lund University (MAX IV)); ERMAKOV, Dmitrii (Lund University)

Presenter: SALNIKOV, Andrii (Lund University (MAX IV))

Session Classification: CS3 Jupyter SIG & Data Science and Visualisation Platforms

Track Classification: Main sessions: User Voice: Innovative Applications, Data Science Environments & Open Data