



CloudKitty onboarding

Mariusz Karpiarz | mariusz.karpiarz@vscaler.com

Pierre Riteau | pierre@stackhpc.com

Rafael Weingärtner | rafael@apache.org

Outline

- ◆ Introduction
- ◆ Architecture overview
- ◆ Limits of CloudKitty
- ◆ Releases
 - Wallaby
 - Xena
 - Yoga
 - Next releases
- ◆ Conclusion



OpenInfra Summit Berlin

June 7-9, 2022

Introduction

Etherpad links for feedback

- ❖ For today's onboarding session:
<https://etherpad.opendev.org/p/cloudkitty-berlin-users-feedback>
- ❖ Yesterday's ops feedback: <https://etherpad.opendev.org/p/cloudkitty-berlin-ops-feedback>

Introduction



The cloud environment

- > Heterogeneous
- > Vast
- > Complex
- > Dynamic



Object Storage

Databases

Identity management

Git managers

...



XenServer

VMware

OVM

KVM

Bare Metal



Local Disc

iSCSI

Fibre Channel

NFS

Swift

Primary Storage

Secondary Storage



Network Type

Isolation

Firewall

Load Balancer

VPN

Introduction



Constant collect usage data:

- ❖ Storage
- ❖ Network
- ❖ Processing
- ❖ VPN
- ❖ Users
- ❖ And others

CloudKitty does not collect data from monitored systems, but from the metrics storage backend!

Data processing and rating:

- ❖ Has to easily scale
- ❖ On the fly metrics
- ❖ Different pricing for different consumptions
- ❖ Support data transformations
- ❖ Able to change between data scale

CloudKitty metrics and mapping modules/rules

Introduction



Hashmap:

- ❖ Execute the rating process based on resource attributes
- ❖ Can be organized in groups and services
 - Services map a rule to the type of data collected
 - Groups are used to organize rating calculations
- ❖ Use a fields of resources to activate/deactivate a rule
 - Flat and Rate ratings rules
- ❖ Can handle threshold rules

PyScripts:

- ❖ Enable complex rating rules
 - Rules that consider more than one attribute to be activated/deactivated
- ❖ The script needs to be written in Python
- ❖ Takes the datapoint being processed, and outputs the price based on the resource and metric data



CLOUKITTY ONBOARDING

Introduction

Metrics definition:

- ❖ “metrics.yml” files
- ❖ YAML based notation of metrics to be processed/rated
- ❖ Define the attributes retrieved from the backend
- ❖ Allow data manipulation
 - Data mutation
 - Unit conversion
 - Extra operations based on type of backend

```
volume.size: ← Metric name in the backend
  unit: GiB
  groupby: ← Group by section
    - id
    - user_id
    - project_id
  metadata: ← Metadata by section
    - volume_type
  extra_args: ← Extra configurations of the backend
    aggregation_method: mean
    resource_type: volume
    force_granularity: 300 ← Override of default backend config

network.outgoing.bytes.rate:
  unit: MB
  groupby:
    - id
    - project_id
    - user_id
  # Converting B/s to MB/h
  factor: 3600/1000000 ← Data conversion/transformation
  metadata:
    - instance_id
  extra_args:
    aggregation_method: mean
    resource_type: instance_network_interface
```


CloudKitty repositories

Introduction



CloudKitty repositories:

- ❖ **cloudkitty/specs** (<https://opendev.org/openstack/cloudkitty-specs>) -- repository to publish major specification before they are implemented;
- ❖ **openstack/cloudkitty** (<https://opendev.org/openstack/cloudkitty>) -- CloudKitty core components. One will find here the API and processor source code;
- ❖ **openstack/python-cloudkittyclient** (<https://opendev.org/openstack/python-cloudkittyclient>) -- Command-line client;
- ❖ **openstack/cloudkitty-dashboard** (<https://opendev.org/openstack/cloudkitty-dashboard>) -- CloudKitty's Horizon UI plugin;
- ❖ **openstack/cloudkitty-tempest-plugin** (<https://opendev.org/openstack/cloudkitty-tempest-plugin>) -- Tempest plugin. Used for gating tests in Python.

Join us!

Introduction

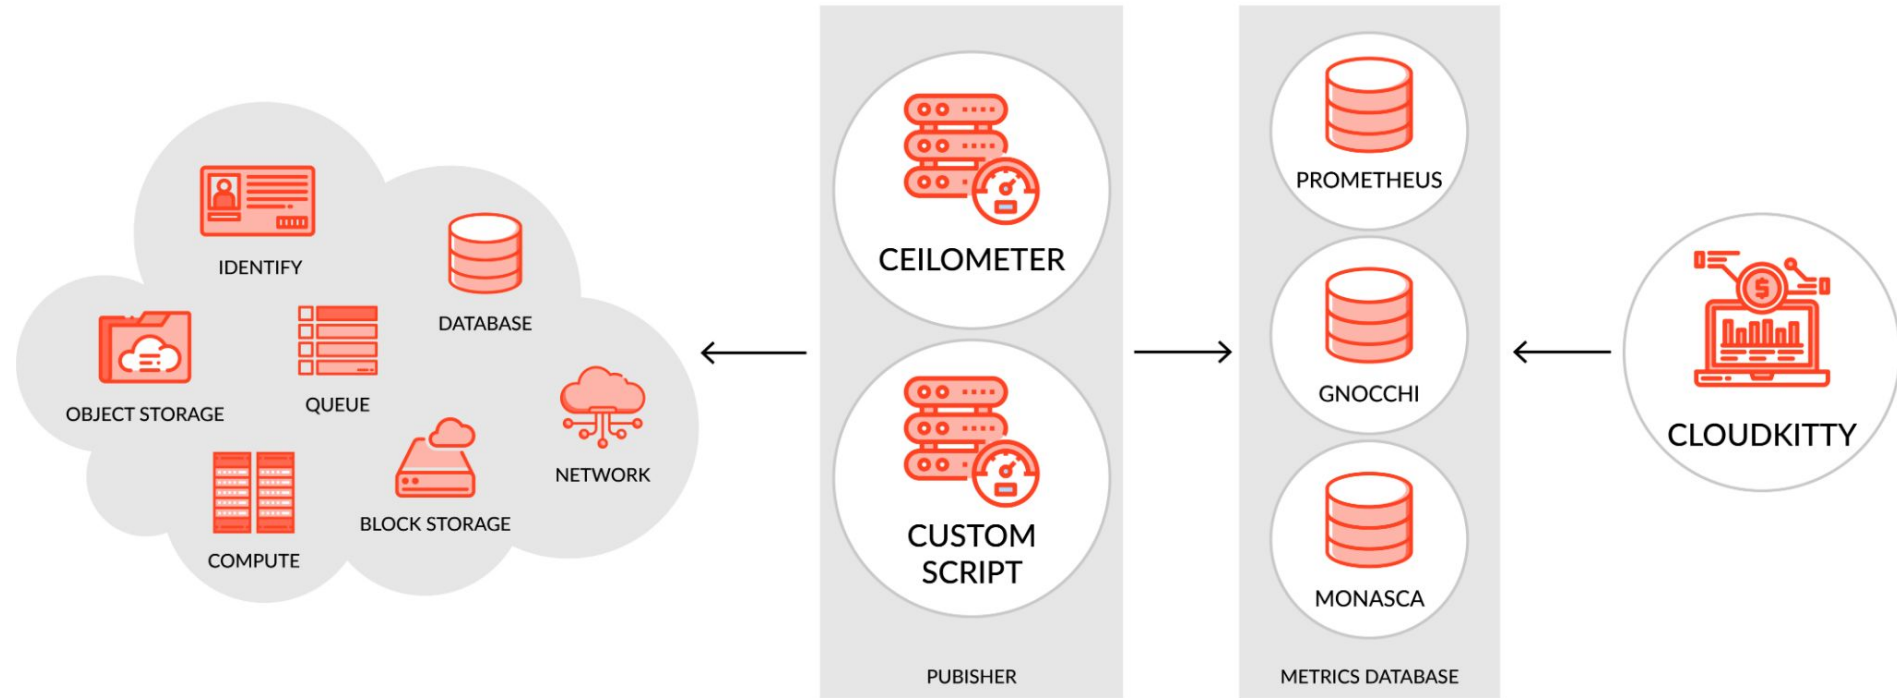


Where to start contributing?

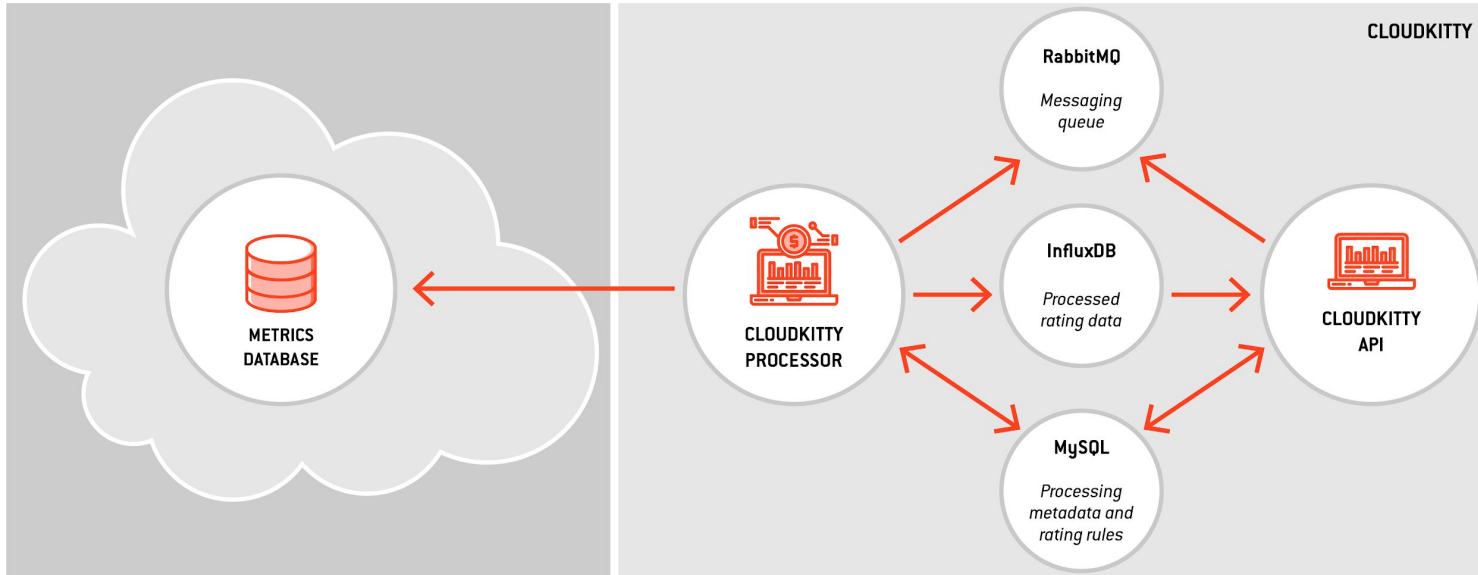
- ❖ Help us reviewing and testing smaller patches
- ❖ Join our bi-weekly IRC meeting -- https://meetings.opendev.org/#CloudKitty_Team_Meeting
- ❖ Ping us on for questions
 - Mailing list is the preferred communication channel
 - #cloudkitty on OFTC IRC
- ❖ **Join the forum discussions!**

CloudKitty architecture overview and limits

Architecture overview



Architecture overview

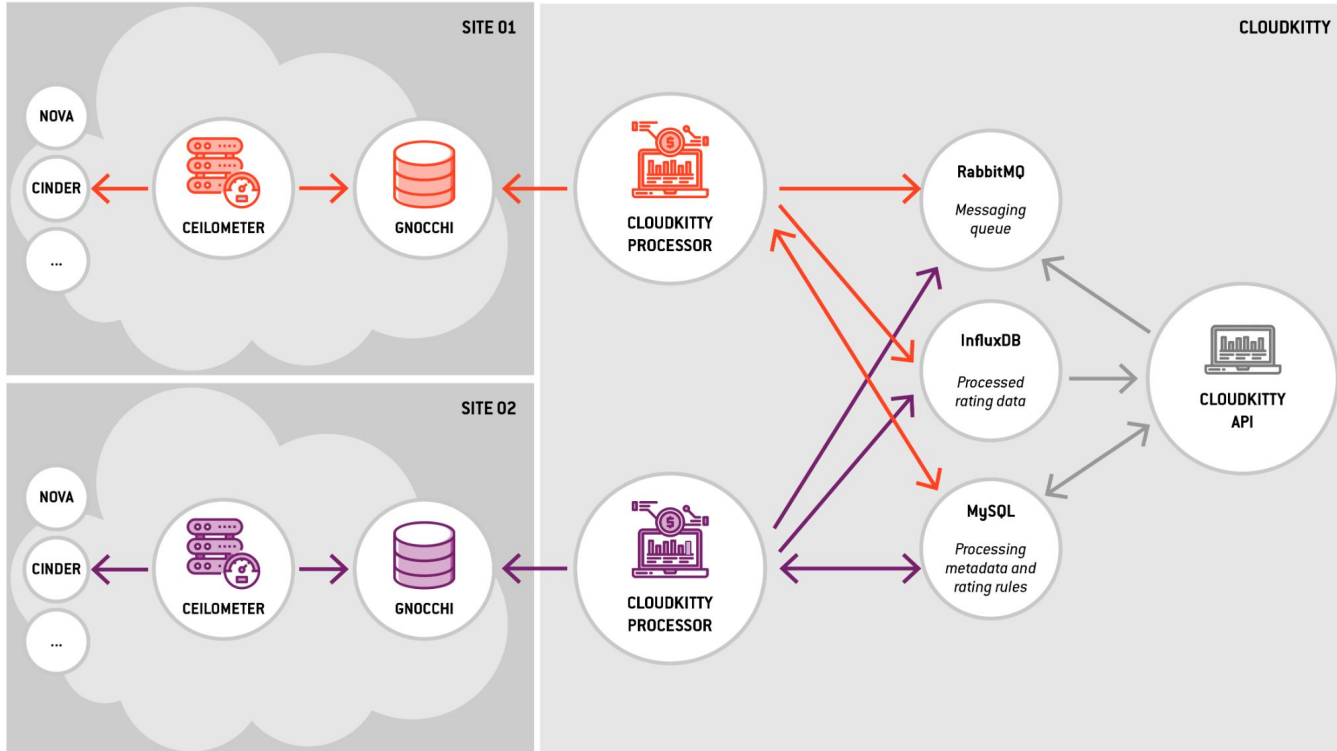


All CloudKitty components in a standard deployment



CLOUKITTY ONBOARDING

Architecture overview

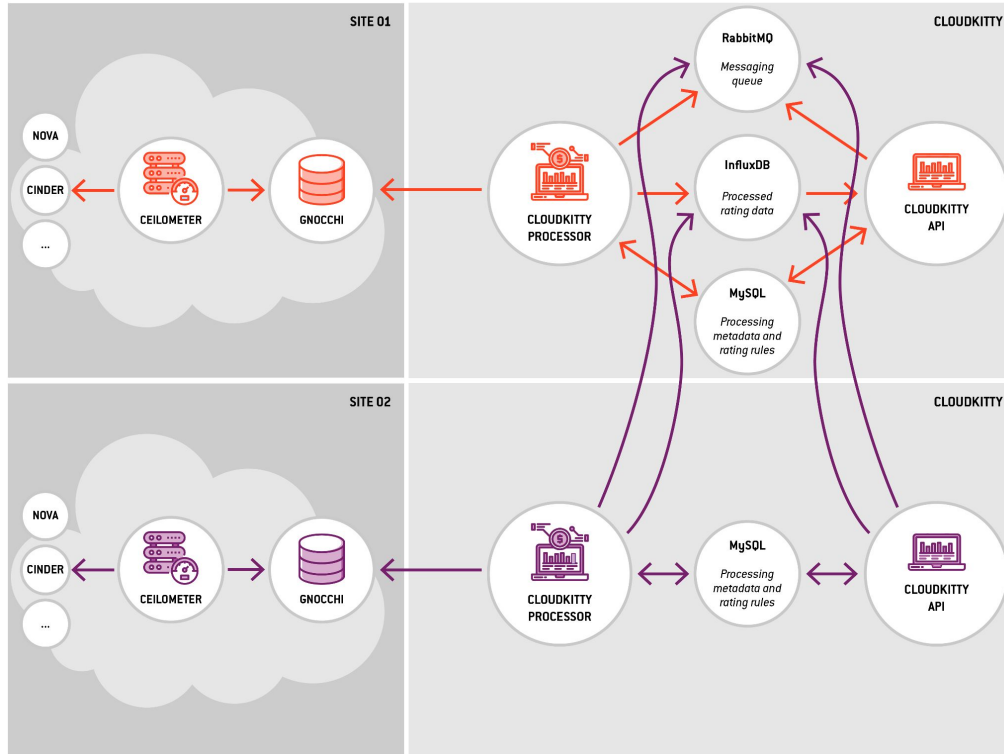


A multi site deployment with a central view/export point



CLOUKITTY ONBOARDING

Architecture overview



A multi site deployment with a distributed site view/export point



CLOUKITTY ONBOARDING

Limits of CloudKitty

- ❖ Depends on the metrics storage backend
- ❖ Single attribute matching with hashmap rules
- ❖ Does not execute invoicing
- ❖ Depends on data collection executed by other systems

Releases



Wallaby

- ❖ Create the option 'use_all_resource_revisions' for Gnocchi collector
- ❖ Fix create_threshold method when using cost as 0
- ❖ The CloudKitty dashboard now inherits the interface type from Horizon

Xena

- ❖ Enable the use of custom queries with the Gnocchi collector
- ❖ The new "NOTNUMBOOL" mutator has been added
- ❖ Fix the definition of the admin_or_owner policy expression

Yoga

- ❖ Options "ignore_disabled_tenants" and "ignore_rating_role" added in the fetcher_keystone
- ❖ Add active status option in the storage state table and API
- ❖ Introduce the reprocessing schedule API
- ❖ Add support for multiple value filters in the summary GET V2 API
- ❖ Adds support for specifying optional prefix and/or suffix to add to Prometheus queries
- ❖ Introduce response_format option for the V2 summary API

Releases



Zed +

- ❖ Add V2 functional tests on Zuul
- ❖ Improve the default configurations to other systems such as Monasca and Prometheus
- ❖ Compound rules for hash mapping rating rules
- ❖ Focus on the community development

Conclusion



Onboarding session

- ❖ Ping us on for questions
 - Mailing list is the preferred communication channel
 - #cloudkitty on OFTC IRC
- ❖ CloudKitty is steadily evolving over the year
- ❖ We still need to develop further CloudKitty community and user base



Thank you!