

# Inspur Venus works for OpenStack

## ➤ **Background**

- It's time-consuming for log querying while the server increasing to thousands.
- It's difficult to retrieve logs, since there are many modules in the platform, e.g. systems, compute, storage, network and other platform services.
- The large amount and dispersion of log make faults are difficult to be discovered.
- Because of distributed and interaction between components of the cloud platform, and scattered logs between components, it will take more time to locate problems.

## ➤ **About Venus**

In light of the problems and needs of retrieval, storage and analysis etc. of logs on the OpenStack platform, we developed the OpenStack log management module Venus. This module can provide a one-stop solution to log collection, cleaning, indexing, analysis, alarm, visualization, report generation and other needs, which involves helping operator or maintainer to quickly solve retrieve problems, grasp the operational health of the platform, and improve the level of platform management. Additionally, this module plans to use machine learning algorithms to quickly locate IT failures and root causes, and improve operation and maintenance efficiency.

# Application scenario

Provide a simple and easy-to-use way to retrieve all log and the context .

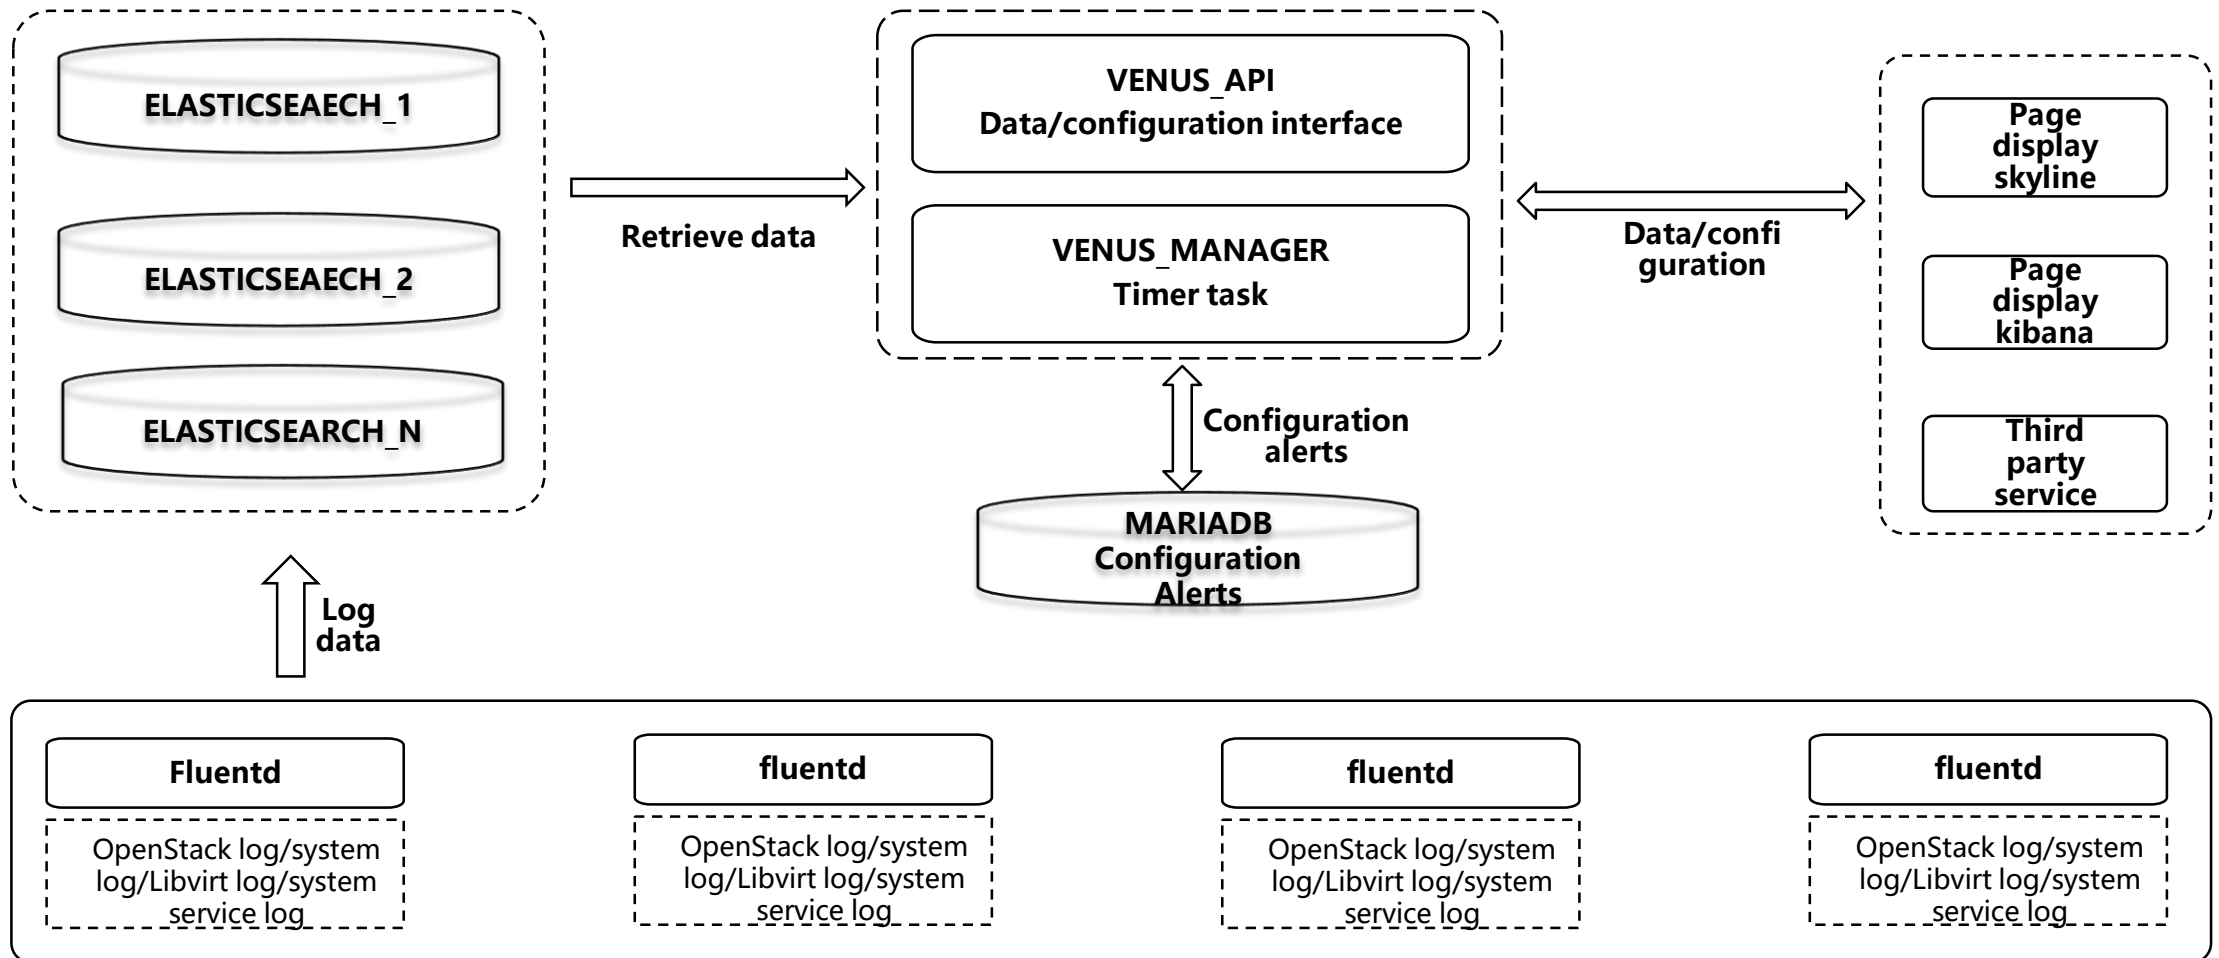


Establish a chain relationship and knowledge graphs to quickly locate problems.

Realize log association, field value statistics, and provide multi-scene and multi-dimensional visual analysis reports.

Convert retrieval into active alerts to realize the error finding in massive logs.

# Overall structure



**vuenus\_api:** API module, provide api、restapi service

**venus\_manager:** Internal timing task module to realize the core functions of the log system

# Code structure

api	RestAPI service
cmd	RestAPI service and command of background periodic task
common	Public base library
db	Database ORM
modules	RestAPI realization
custom_config	Personalized configuration sub-module
backend	Database read and write operations
action.py	Interface layer
controller.py	Interface controller layer
search	Data retrieval sub-module
action.py	Interface layer
controller.py	Interface controller layer
task	Background periodic task
backends	Database read and write operations
core	Task realization
delete_es_index_task.py	Delete Elasticsearch expired index
unittest	Unit test

# Current progress

## collection

Develop fluent collection tasks based on collected to read, filter, format and send plug-ins for OpenStack, operating systems, and platform services, etc.

## Index

Develop the API interface and restapi interface to deal with multi-dimensional index data in elastic-search, and provide more concise and comprehensive authentication interface to return query results.

## analysis

Develop task scheduling code to analyze and display the dispmodule errors, mariadb connection errors, rabbitmq connection errors and so on.

## alerts

Develop alarm task code to set threshold for the number of error logs of different modules at different times, and provides alarm services and notification services.

## location

Develop the call chain analysis function based on global\_requested series, which can show the execution sequence, time and error information, etc., and provide the export operation.

## management

Develop configuration management functions in the log system, such as alarm threshold setting, timing task management, and log saving time setting, etc..

# Current function



Full text/multi-dimensional retrieval

Keyword retrieve for full text or indexed dimensions (fuzzy matching is possible)



Multi-dimensional statistics/display

Count and display the number of logs of one or more dimensions that have been indexed



Typical error analysis

Integrate typical error analysis, such as mariadb connection error, etc.



Error log alerts

Give an alarm and notification when the number of error logs in a period exceeds the threshold.



Call chain analysis

Analyze the process of calls in the system according to global\_requestid



API execution status/time-consuming

Analyze the status and time consumption of API execution for modules that record the status of API calls.

# Current log source

## ➤ **OpenStack logs**

- Integrated most modules of OpenStack system.
- Including nova、 neutron、 keystone、 cinder、 glance, etc.

## ➤ **Operating system logs**

- The log of operating system on which OpenStack platform runs.
- Take CentOS7 as an example, including messages, secure, boot, btmp, dmesg, etc.

## ➤ **Platform service logs**

- Platform service logs which are depending on OpenStack system.
- Including mariadb、 rabbitmq、 zookeeper、 kafka、 haproxy, etc.

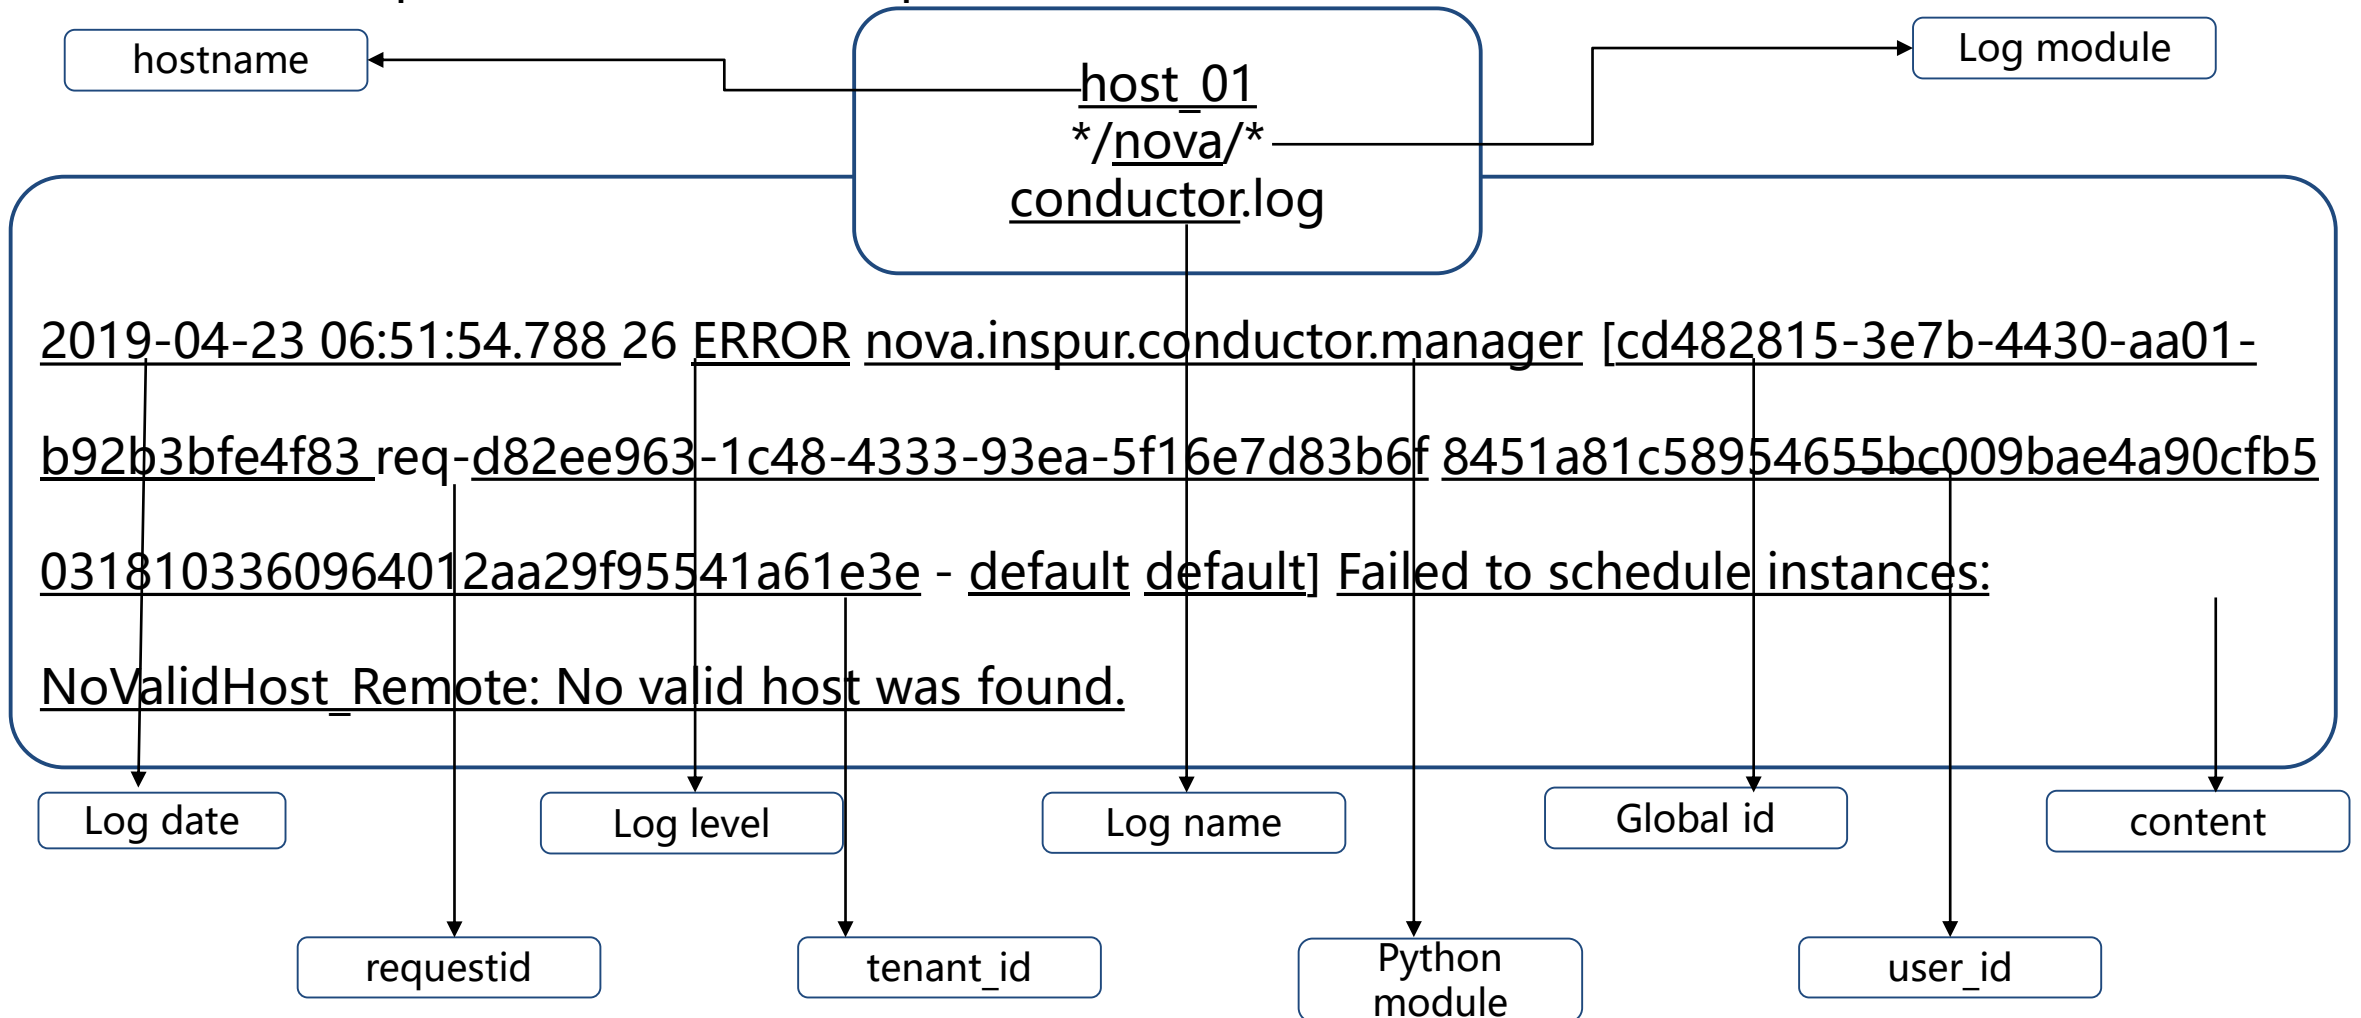
## ➤ **Virtualization logs**

- Virtualization logs at the bottom of OpenStack.
- Mainly libvirtd.log.



# OpenStack log index dimensions

In addition to full-text indexing, OpenStack logs will be indexed from the following dimensions to facilitate subsequent retrieval in multiple dimensions..



# Next step

extending  
plugins

Explore more operation and maintenance scenarios, and conduct statistical analysis and alarm on key data.

developing  
show

Form clustering log and construct knowledge map, and integrate algorithm class library to locate the root cause of the fault.

collection

analysis

display

location

In addition to fluent, other collection plugins such as logstash will be integrated.

exploring  
Scenes

The configuration, analysis and alarm of Venus will be integrated into horizon in the form of plugin.

Integrating  
algorithm

**Thank you**

***inspur* 浪潮**