

# New Chapter of Fluentd

Rebranding and New Release Cycle (LTS)

Daijiro Fukuda, Kentaro Hayashi

*ClearCode Inc.*

*Open Source Summit JAPAN 2023*

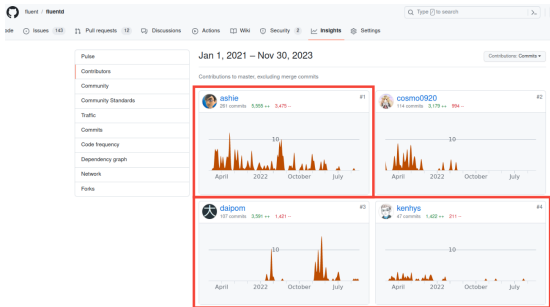
*2023-12-05*

# About us

- ✓ **Maintainers of Fluentd Project**
- ✓ Software Engineers at ClearCode Inc
  - ✓ Driving forward Free Software and Make the business



# Fluentd and us



<https://github.com/fluent/fluentd/graphs/contributors?from=2021-01-01&to=2023-11-30&type=c>

# Agenda

- ✓ What is Fluentd
- ✓ History of Fluentd
- ✓ What's new Fluent Package LTS
- ✓ Remarks of updating
- ✓ Major changes



# Agenda

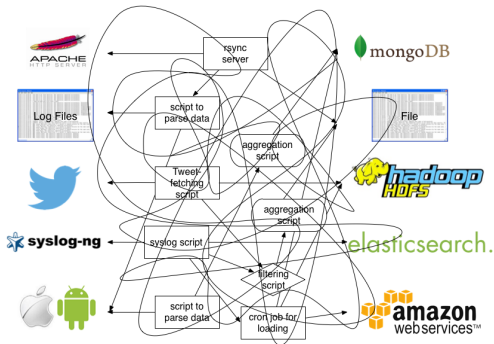
- ✓ **What is Fluentd**
- ✓ History of Fluentd
- ✓ What's new Fluent Package LTS
- ✓ Remarks of updating
- ✓ Major changes

# What's Fluentd

“*Fluentd is an open source data collector that unifies data collection and consumption.*”




# Image of data collection without Fluentd

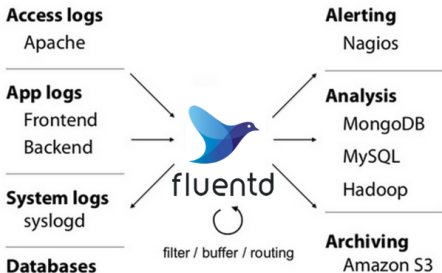


# What Fluentd does

## Without Fluentd



## With Fluentd



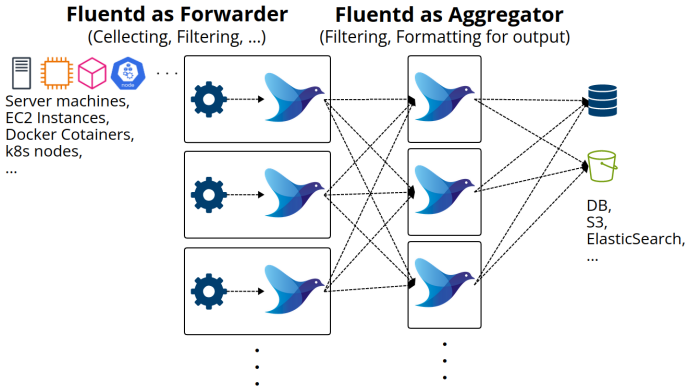
# Why use Fluentd?

- ✓ Flexibility
- ✓ Scalability
- ✓ Open source and community

# Flexibility

- ✓ Over 1000 plugins for many data sources and outputs
- ✓ By combining plugins, you can achieve what you want  
Main functions such as data collection, filtering, formatting, and sending are performed by each plugin
- ✓ You can easily make your plugin with Ruby

# Scalability



# Open source and community

- ✓ You can join to development and make what you want
- ✓ Plugins are maintained in the community
  - ✓ <https://github.com/fluent-plugins-nursery>



# Agenda

- ✓ What is Fluentd
- ✓ **History of Fluentd**
- ✓ What's new Fluent Package LTS
- ✓ Remarks of updating
- ✓ Major changes

# History of Fluentd

2011	v0.9 Released by Treasure Data
2016	Accepted to CNCF
2017	v1.0 Released
2019	Graduated CNCF project

# Needs of distribution package

- ✓ Fluentd is general purpose
- ✓ Plugins allow specific features
- ✓ Ruby is required to use

Users want **convenience**. We need distribution package.

# History of distribution package

Year	Package	Fluentd
2011	td-agent v1	v0.9
2014	td-agent v2	v0.10
2017	td-agent v3	v0.14
2020	td-agent v4	v1.11
2023	fluent- package v5	v1.16

# Reason of rebranding

- ✓ The name td-agent does not associate with Fluentd

The name of distribution package is better to associate or describe the main software

- ✓ td-agent is so called all-in-one package, not limited for Treasure Data CDP

# Fluent Package


- ✓ The new package name reflects community driven development!



# Agenda

- ✓ What is Fluentd
- ✓ History of Fluentd
- ✓ **What's new Fluent Package LTS**
- ✓ Remarks of updating
- ✓ Major changes

# Fluent Package LTS

- ✓ v5.0.2 had been released at Nov 2023  

- ✓ Recommended for stable production use
  - ✓ **Scheduled** support lifecycle
  - ✓ Security fixes, bug fixes for long term

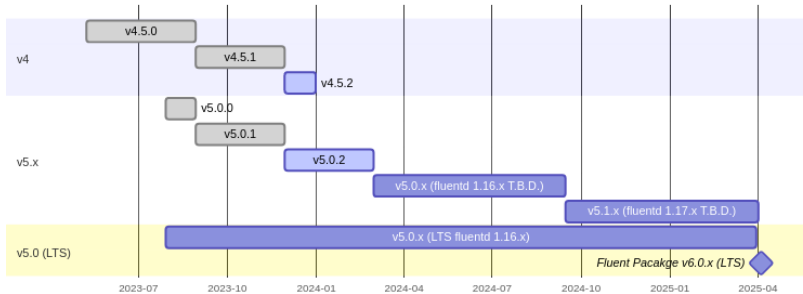


# Standard and LTS

Channel	Support lifecycle
Standard version	Every 3 months
LTS (Long term support)	Every 2 year or so

# Planned release cycle

Scheduled Support lifecycle for Fluent Package



<https://www.fluentd.org/blog/fluent-package-scheduled-lifecycle>

# Profit for updating to LTS

- ✓ **Secure and stable operation!**  
Long Term Support for updates, ready for stable production use
- ✓ **Less hustle for major update!**  
Advanced announcement for scheduled major update  
Next major update is scheduled in April, 2025 (T.B.D.)

# Agenda

- ✓ What is Fluentd
- ✓ History of Fluentd
- ✓ What's new Fluent Package LTS
- ✓ **Remarks of updating**
- ✓ Major changes

# How to update from td-agent v4

- ✓ In most cases, you just install Fluent Package! 🎉
- ✓ **Compatible with td-agent v4!**
- ✓ In some cases, manual migration is necessary

# Compatibility

You can still use ...

✓ your config files as is

✓ `/etc/td-agent/td-agent.conf`

✓ commands with their old names

✓ `td-agent`, `td-agent-gem`

✓ `td-agent` user and group

# Manual migration cases

- ✓ General case
- ✓ Specific to Linux
- ✓ Specific to Windows

See <https://www.fluentd.org/blog/upgrade-td-agent-v4-to-v5> for details!

# Manual migration: general

- ✓ Manually installed plugins
  - ✓ Reinstall those plugins manually
- ✓ Setting of monitoring Fluentd's log file
  - ✓ Update the filepath  
`/var/log/fluent/fluentd.log`



# Manual migration: Linux

- ✓ Environment file
  - ✓ RHEL: `/etc/sysconfig/fluentd`
  - ✓ Debian: `/etc/default/fluentd`
- ✓ Log rotate configuration
  - ✓ `/etc/logrotate.d/fluentd`
- ✓ Systemd unit file

# Systemd unit file

## Update custom unit files (If you use)

e.g. `/etc/systemd/system/td-agent.service`

### ✓ Update Fluentd bin path

✓ `/opt/td-agent/bin/fluentd -> /opt/fluent/bin/fluentd`

### ✓ Update environmental variables

✓ `LD_PRELOAD: /opt/fluent/lib/libjemalloc.so`

✓ `GEM_HOME, GEM_PATH: /opt/fluentd/lib/ruby/gems/3.2.0/`

# Manual migration: Windows

- ✓ Command line arguments (If you use)
  - ✓ Set them again
    - \$ fluentd -reg-winsvc-fluentdopt ...
- ✓ Automatic start as Windows Service
  - ✓ Enable it manually if you need
    - \$ sc config fluentdwinsvc start=delayed-auto

# Fluentd's log rotation setting on Windows

- ✓ No longer need to use command line argument, **--log-rotate-age**, **--log-rotate-size**

Since td-agent v4.4.2, you can use system configuration

- ✓ Use **system** configuration

<https://docs.fluentd.org/deployment/logging#log-rotation-setting>

# Agenda

- ✓ What is Fluentd
- ✓ History of Fluentd
- ✓ What's new Fluent Package LTS
- ✓ Remarks of updating
- ✓ **Major changes**

# Changes for Fluent Package

## ✓ Command name

If updated from td-agent, you can still use the old name

✓ td-agent => fluentd, td-agent-gem => fluent-gem

## ✓ Linux: systemd unit name

You can still use the old name

✓ td-agent => fluentd

## ✓ Windows: no automatic start

✓ Fluentd: Fix **in\_tail** bug

✓ Many other improvements!

See <https://github.com/fluent/fluentd/blob/master/CHANGELOG.md>!

✓ Fluent Package: Update embedded **Ruby** to **3.2**  
from 2.7

✓ Many other improvements!

See <https://github.com/fluent/fluent-package-builder/blob/master/->

[CHANGELOG.md](#)!

# Fix in\_tail bug

## Bug Fixes

`in_tail`: Fix `in_tail` wrongly stopping tailing the current target file and causing handle leaks

This issue has been reported for years. It was looking very serious, but the cause remained unclear until recently.

At last, we found different causes when the `follow_inodes` option is `true` and when it is `false` (default). We fixed the problem of `follow_inodes true` case in v1.16.2, and fixed the problem of `follow_inodes false` case in v1.16.3.

<https://www.fluentd.org/blog/fluentd-v1.16.2-v1.16.3-have-been-released>



# Fix in\_tail bug

- ✓ Finally, the bug is fixed!
- ✓ With **collaboration of many users**, the cause was found, and it is fixed! 🎉

# Fix in\_tail bug

[in\_tail][follow\_inode enabled] found a racing bug that new created watcher be unwatched unexpectedly and thus log stuck #4190

 Closed garyzjq opened this issue on May 31 · 4 comments · Fixed by #4191 or #4208



garyzjq commented on May 31 · edited

Contributor

## Describe the bug

⋮

Here's the racing description:

1. File rotated: previous log file 'container.log' is renamed to "container.backup.log" (inode\_old), new 'container.log' created (inode\_new)
2. refresh\_watchers regularly triggered, mark **old watcher for inode\_old to be detach** and create a new watcher for inode\_new (**@tail[path] now point to new watcher**)
3. old watcher on\_rotate triggers update\_watcher as watcher\_needs\_update = true.
4. in update\_watcher, "rotated\_tw = @Tails[path]" will **return the new watcher to be detached** because tw with same **path** has been updated to step 2.
5. After rotate\_wait, both old and new watchers are detached in step 2 and 4, thus new watcher stuck

The main problem is in #4, when try to detach a watcher, need to determine whether the rotated\_tw is still the old inode, which can be detach, or already refresh to new inode, which need skip detach (detached by other event).

# Fix in\_tail bug

in\_tail: Use inode for key of TailWatcher when follow\_inodes #4185



Open

kattz-kawa wants to merge 1 commit into `fluent:master` from `kattz-kawa:issue3614`



**kattz-kawa** commented on May 26 • edited ▾

Contributor



**Which issue(s) this PR fixes:**

Partially fixes [#3614](#) (follow\_inode true case)

**What this PR does / why we need it:**

Thank you for reading to the great developer's!

This PR is fixed for unexpected file close after logs rotate in fluentd v1.16.1.

I found the [#3614](#)'s reproducer.

# Update embedded Ruby to 3.2

- ✓ You can try Ruby's performance improvements **YJIT!**

[https://docs.fluentd.org/deployment/system-config#enable\\_jit-experimental](https://docs.fluentd.org/deployment/system-config#enable_jit-experimental)

- ✓ It is still experimental in Fluentd  
Please try it and give us feedback!

# Fluent Package and future

- ✓ What happened?
- ✓ What will happen?

# The fact from statistics (1)

- ✓ Recent Top 5 requests:
  - ✓ United States (44.74%)
  - ✓ Netherlands (9%)
  - ✓ **Japan (6.79%)**
  - ✓ France (5.96%)
  - ✓ Ireland (5.58%)

# The fact from statistics (2)

Recent trends:

- ✓ Top 5 location Before LTS release
  - ✓ United States, Netherlands, Japan, France, Ireland
- ✓ Top 5 location After LTS release
  - ✓ United States, Netherlands, Japan, France, **Taiwan**

# The fact from statistics (3)

- ✓ The undesirable facts:
  - ✓ The majority of td-agent users: 3.x (unmaintained) or 4.x series.
  - ✓ There is still an access for 2.x (unmaintained)



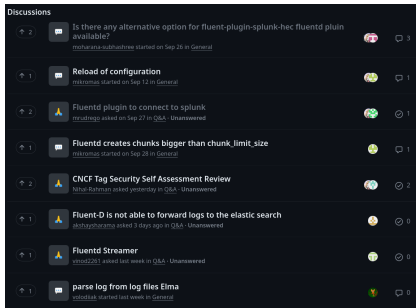
# The fact from statistics (4)



- ✓ The most widely used package:
  - ✓ td-agent 4.x (aarch64) for Amazon Linux 2.

# The fact from statistics (5)

✓ GitHub Discussions low activity



The screenshot displays a list of GitHub Discussions for the 'fluent-plugin-splunk-hec' repository. The list includes the following items:

- Discussion 1:** "Is there any alternative option for fluent-plugin-splunk-hec fluentd plugin available?" by moharana-subhashree, started on Sep 26 in General. It has 2 upvotes and 3 replies.
- Discussion 2:** "Reload of configuration" by mikromas, started on Sep 12 in General. It has 1 upvote and 1 reply.
- Discussion 3:** "Fluentd plugin to connect to splunk" by rruudrigo, asked on Sep 27 in Q&A - Unanswered. It has 2 upvotes and 1 reply.
- Discussion 4:** "Fluentd creates chunks bigger than chunk\_limit\_size" by mikromas, started on Sep 28 in General. It has 1 upvote and 1 reply.
- Discussion 5:** "CNCF Tag Security Self Assessment Review" by Nihal-Rahman, asked yesterday in Q&A - Unanswered. It has 2 upvotes and 2 replies.
- Discussion 6:** "Fluent-D is not able to forward logs to the elastic search" by abhysbarama, asked 3 days ago in Q&A - Unanswered. It has 1 upvote and 0 replies.
- Discussion 7:** "Fluentd Streamer" by vmod2261, asked last week in Q&A - Unanswered. It has 1 upvote and 0 replies.
- Discussion 8:** "parse log from log files Elma" by volodkak, started last week in General. It has 1 upvote and 0 replies.

# The future

- ✓ No more old package users
- ✓ Renewal about distribution channel (Sponsored by CNCF?)
- ✓ Encourage users to be a fluentd supporter  
Join GitHub discussions!
- ✓ Continuous release

# Wrap up

- ✓ Interested in and understand Fluent Package LTS!
- ✓ Encouraged switching to fluent-package LTS further more!
- ✓ Encouraged joining Fluentd community!