

Welcome to Chukwa!

Table of contents

1 About Chukwa.....	2
2 News.....	2
3 Getting Started	2
4 Getting Involved	2
5 Disclaimer	2

1 About Chukwa

Chukwa is an open source data collection system for monitoring large distributed systems. Chukwa is built on top of the Hadoop Distributed File System (HDFS) and Map/Reduce framework and inherits Hadoop's scalability and robustness. Chukwa also includes a flexible and powerful toolkit for displaying, monitoring and analyzing results to make the best use of the collected data.

2 News

26 April 2010: [Release 0.4 available](#). This is our second release, incorporating many bug fixes, improved documentation, and a variety of new features, including support for processing UDP messages.

See the [Chukwa 0.4.0 Release Notes](#) for details. Alternatively, you can look at the [Jira issue log for all releases](#).

3 Getting Started

Releases are available on the [release page](#). For documentation, check [here](#) and on the [Wiki](#).

4 Getting Involved

Chukwa is an open source volunteer project under the Apache Software Foundation. We encourage you to learn about the project and contribute your expertise. Here are some starter links:

1. Give us [feedback](#): What can we do better?
2. Join the [mailing list](#): Meet the community.

5 Disclaimer

Apache Chukwa is an effort undergoing incubation at The Apache Software Foundation (ASF), sponsored by the Incubator project. Incubation is required of all newly accepted projects until a further review indicates that the infrastructure, communications, and decision making process have stabilized in a manner consistent with other successful ASF projects. While incubation status is not necessarily a reflection of the completeness or stability of the code, it does indicate that the project has yet to be fully endorsed by the ASF.