Welcome to Chukwa!

Table of contents

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 #exible and powerful toolkit for displaying, monitoring and analyzing results to make the best use of the collected data.

2. News

Chukwa 0.3 released. This is the first public release of Chukwa. Release notes are available here. Note that we've been updating the documentation, in response to early user feedback. You should consult the Apache copy of the documentation, rather than the copy included with the release.

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. It is a subproject of Hadoop. 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.