

Feature Status

This page will try to highlight different features and their maturity based on stories gathered from the community.

States are

- ✓ Recommended - feature is known to be in daily use in multiple places, and recommended for use. The API is stable.
- ⚠ Beta - feature is usable, known to be in use at one or to places, but may have some issues or limitations. The API is close to stable.
- ❓ Experimental/Development - feature is in development and will likely have issues. API may change significantly. Be prepared to go into code.
- ℹ Testing - for debugging and testing purposes only.

v0.9.4 (pre-apache release)














Flume Master features

Feature	Status	Notes
Master Mode		
Single Master ZK-backed mode	✓	Default. Recommended mode of use
Single Master Memory-backed mode	ℹ	Useful for development and debugging
Multi-Master ZK-backed mode	❓	This is being used by a handful of users but has known limitations. (e.g. no support for auto-agent collector mapping)
Automatic configurations		
Automatic agent-collector mapping	❓	Must use Single Master mode. All logical nodes must be up before this feature will work. May have problems with reconfiguration or if configuration written before logical nodes report
Automatic flow isolation	❓	Must use Single Master mode. Dependent upon Automatic agent-collector mapping
Metrics		
Master JSON Metrics	⚠	Looking for feedback and nagios/munin/etc integration stories.

Flume Node features

Client	Status	Notes
Linux clients	✓	Core development and most users are here
Windows client	❓	This is fairly early

Feature	Status	Notes
Agents and Collectors		
agentBESink	✓	
agentDFOSink	✓	
agentE2ESink / agentSink	✓	
collectorSource	✓	
collectorSink ("hdfs://...",...)	✓	HDFS is the recommended target file system.
collectorSink ("s3n://...",...)	✓*	In use but has known issue due to how s3 files are "closed"
collectorSink ("file:///...",...)	ℹ	Intended for testing but could be used in production
agentBEChain	⚠	Need known production use stories
agentDFOChain	⚠	Need known production use stories
agentE2EChain	⚠	Need known production use stories

auto*Chain		Limitation on at the master, to be address in future versions
Sources		
thriftSource		Default RPC. Recommended mode of use
avroSource		Need known production use stories
syslogTCP source		
tail/multiTail/tailDir	 *	This is known to have some duplication issues but has been used in production settings. If this is encountered, known workarounds include using exec source, or having applications write to flume via RPC
exec source		
log4j appender sources		
text		
Sinks		
thriftSink		Default RPC. Recommended mode of use
avroSink		Need known production use stories
attr2hbase, hbase sink plugin		Known to be in use in a few places
seqfile sink		Used internally (E2E mode)
dfs sink		Used internally
Metrics		
Node JSON metrics		Looking for feedback and nagios/munin/etc integration stories. This api is likely to change as we receive feedback