

Info [?]	Name	Homepage	Source Code	Storage Engine	Language	Bytes/point (active Bytes/point (active Dimensions/Tags (Clustering	Comments	NOTE: YOU CAN COMMENT (ADD SUGGESTIONS) THIS SHEET. For more information about the data - see the bottom of the sheet
	Aria	https://github.com/Netflix/aria/wiki/Overview	https://github.com/Netflix/aria	Custom	Scala			Inventory, multi-dimensional
	BigGraphite	https://github.com/etsy/biggraphite	https://github.com/etsy/biggraphite	Cassandra	Python	16 24 [4] No [5]	Yes	its design goals were set with whisper in mind, but based on cassandra
	ClickDB	https://github.com/clickhouse/clickhouse	https://github.com/clickhouse/clickhouse	PostgreSQL	C		Yes	Yes [8]
	graphite-clickhouse	https://github.com/clickhouse/graphite-clickhouse	https://github.com/clickhouse/graphite-clickhouse	Clickhouse	GoLang	2 [6] 6.5 [10] Partial [11]	Yes	Byte/point - depends on nature of the data. Syntactically - 5-ish, real-world - 6-ish. Graphite compatibility layer with partial support for tags. https://github.com/clickhouse/clickhouse and https://github.com/clickhouse/graphite-clickhouse
	graphouse	https://github.com/clickhouse/graphouse	https://github.com/clickhouse/graphouse	Clickhouse	Java	5.4 [12] 6.5 [13]	No	Yes
	Graphite	https://github.com/etsy/graphite-core	https://github.com/etsy/graphite-core	Custom	GoLang		?	Yes
	Influx	https://github.com/influxdb/influxdb	https://github.com/influxdb/influxdb	InfluxDB	Java		?	Yes
	elasticsearch	https://www.elastic.co/	https://github.com/elastic/elasticsearch	Apache Lucene	Java	22 [14] 66 [15]	Yes	Yes
	FluentD	https://github.com/fluentd/fluentd	https://github.com/fluentd/fluentd	Apache Fluentd	Java		?	Yes
	Graphite	https://github.com/etsy/graphite-core	https://github.com/etsy/graphite-core	Spark + Cassandra Scala	Python		?	Yes
	Harvic	https://github.com/etsy/harvic	https://github.com/etsy/harvic	PostgreSQL	Java		Yes	Developed by OpenBank guys. Sounds good on paper. Uses Ceph/S3/Swift API to store time-series data. Requires MySQL or PostgreSQL for tags
	InfluxDB	https://github.com/influxdb/influxdb	https://github.com/influxdb/influxdb	Custom	GoLang	2 [16] 3 [17]	Yes	Commercial
	KaravDB	https://github.com/karavdb/karavdb	https://github.com/karavdb/karavdb	Cassandra	Java	10 [20] 12 [21]	Yes	Yes
	m3db	https://github.com/m3db/m3db	https://github.com/m3db/m3db	Custom	GoLang	1.4 [22] 1.4 [23]	Yes	Yes
	metricbank	https://github.com/metricbank/metricbank	https://github.com/metricbank/metricbank	Cassandra	GoLang	1.3 [24] 3 [25]	Yes (with caution)	Yes [27]
	NetData	https://github.com/netdata/netdata	https://github.com/netdata/netdata	Custom	C		No?	No
	OpenTSDB	https://opentsdb.net/	https://github.com/etsy/opentsdb	Cassandra	Java		Yes?	Yes
	OpenTSDB	https://opentsdb.net/	https://github.com/etsy/opentsdb	Custom	GoLang		Yes	Cassandra-based. Has Graphite-compatible input
	Prometheus	https://github.com/prometheus/prometheus	https://github.com/prometheus/prometheus	Custom + LevelDB	GoLang	12 [28] 39 [29]	Yes	Yes
	QuotedDB	https://github.com/quotedb/quotedb	https://github.com/quotedb/quotedb	Custom	Java	1.3 [30] 3.3 [31]	No?	Yes
	Riak TS	https://github.com/basho/riak-ts	https://github.com/basho/riak-ts	Riak	Erlang		Yes	Yes
	Reali	https://github.com/realtime/reali	https://github.com/realtime/reali	Custom	GoLang		?	Yes
	RRDTool	https://github.com/netmeister/rrdtool	https://github.com/netmeister/rrdtool	RRD	C		No	No
	Scylla	https://github.com/scylladb/scylla	https://github.com/scylladb/scylla	Custom (Inherited C++)	C++		?	Yes
	ScyllaDB	https://github.com/scylladb/scylladb	https://github.com/scylladb/scylladb	Custom	C		?	Yes
	TimescaleDB	https://github.com/timescale/timescaledb	https://github.com/timescale/timescaledb	PostgreSQL	GoLang		Yes?	Yes
	TimescaleDB	https://github.com/timescale/timescaledb	https://github.com/timescale/timescaledb	PostgreSQL	C		?	Yes
	Warp10	https://github.com/longhorn/warp10	https://github.com/longhorn/warp10	LevelDB, HBase	Java	9 [32] 9 [33]	Yes?	Yes
	Whisper (Carbon)	https://github.com/etsy/whisper	https://github.com/etsy/whisper	Whisper	Python	12 [34] 12 [35]	Yes [36]	Yes [37]
	Whisper (in-carbon)	https://github.com/clickhouse/graphite-clickhouse	https://github.com/clickhouse/graphite-clickhouse	Whisper	GoLang	12 [38] 12 [39]	Partial [40]	Yes [41]
	CrashDB	https://github.com/crashdb/crashdb	https://github.com/crashdb/crashdb	Lucene [42]	Java		Yes	Yes
	VictorMetrics	https://github.com/victoriametrics/victoriametrics	https://github.com/victoriametrics/victoriametrics	Custom	GoLang	11 [43] 16 [44]	Yes	Yes [44]
	QuestDB	https://github.com/questdb/questdb	https://github.com/questdb/questdb	Custom	C++		?	?
	LinDB	https://github.com/lindb/lindb	https://github.com/lindb/lindb	Custom?	GoLang		Yes	Yes
	antiDB	https://github.com/anti-db/anti-db	https://github.com/anti-db/anti-db	Apache Parquet	GoLang		Yes	Yes
	TiDB	https://github.com/pingcap/tidb	https://github.com/pingcap/tidb	MySQL	GoLang		Yes	Yes
	tyro	https://github.com/tyro/tyro	https://github.com/tyro/tyro	Custom	GoLang		Yes	Yes
	eyon	https://github.com/eyon/eyon	https://github.com/eyon/eyon	Custom	GoLang		Yes	Yes
	erdb	https://github.com/erdb/erdb	https://github.com/erdb/erdb	Custom	GoLang		Yes	Yes
	weptd	https://github.com/weptd/weptd	https://github.com/weptd/weptd	Custom	GoLang		Yes	Yes
								Abandoned since 2019

Info [45]	Name	Homepage	Source Code	Storage Engine	Language	Bytes/point	(adv Bytes/point)	(wo Dimensions/Tag/Clustering)	Comments	NOTE: YOU CAN COMMENT (ADD SUGGESTIONS) THIS SHEET. For more information about the data - see the bottom of the sheet
	Akumuli	http://akumuli.org	https://github.com	Custom	C++	0.1 [48]	8.5 [49]	Yes	No	No commits for a half year. Might be dead
	Argus	https://github.com	https://github.com	Custom	Java			Yes	No?	DEAD (Archived on github)
	beeringei	https://github.com	https://github.com	Custom	C++			?	No	DEAD. Reference implementation of Corilla whitepaper. In-memory
	BlueFood	https://bluefood.io	https://github.com	Cassandra	Java	8 [50]	8 [51]	Yes	Yes?	DEAD.
	BT4DB	https://blog.aolny.com	https://github.com	Custom	Golang			No?	Yes?	Likely dead. Several years without commits. Comercial version's website doesn't work properly. Academical DB. Papers sounds very nice. Developed for IoT applications
	Catena	https://github.com	https://github.com	Custom	Golang					DEAD. SHOULD BE USED ONLY FOR HISTORICAL REASONS.
	Ceres	http://rj.schlesag.com	https://github.com	Custom	Ceres					Abandonware
	Chronix	http://chronix.io	https://github.com	Custom	Python	8 [52]	8 [53]	Yes	Yes* [54]	DEAD
	Chronix	http://chronix.io	https://github.com	Custom	Java	3 [55]	3 [56]	Yes	Yes	DEAD
	Cube	https://github.com	https://github.com	Custom	Lucene					DEAD FOR OVER 4 YEARS. SHOULD BE USED ONLY FOR HISTORICAL REASONS
	Cyante	https://github.com	https://github.com	Custom	NodeJS					DEAD FOR OVER 4 YEARS. SHOULD BE USED ONLY FOR HISTORICAL REASONS
	Cyante	https://github.com	https://github.com	Custom	NodeJS					DEAD FOR OVER 4 YEARS. SHOULD BE USED ONLY FOR HISTORICAL REASONS
	Cyante	https://github.com	https://github.com	Custom	NodeJS					DEAD FOR OVER 4 YEARS. SHOULD BE USED ONLY FOR HISTORICAL REASONS
	Dalmatiner	https://dalmatiner.com	https://github.com	Custom	Riak				Yes [57]	DEAD. Docs are very bad. No recommendations on clustering, just mentions that it supports it. One man's project. Sounds very good on paper
	EventQL	https://eventql.io	https://github.com	Custom	C++				?	DEAD. BETA QUALITY. General Purpose Column-based analytics db, event-based, needs investigation if its suitable for time-series data
	hawkular	http://www.hawkular.com	https://github.com	Custom	Java				12	DEAD. BETA QUALITY. General Purpose Column-based analytics db, event-based, needs investigation if its suitable for time-series data
	hawkular	http://www.hawkular.com	https://github.com	Custom	Java				12	DEAD. BETA QUALITY. General Purpose Column-based analytics db, event-based, needs investigation if its suitable for time-series data
	Kenshin	https://github.com	https://github.com	Custom	Python				No	DEAD. Design goals - make whisper less I/O hungry
	Kenshin	https://github.com	https://github.com	Custom	Python				No	DEAD. Design goals - make whisper less I/O hungry
	Sidewinder	https://github.com	https://github.com	Custom	Java	3 [58]	4.5 [59]	Yes	Yes	Custom DB with own storage, protocol and API. Accepts InfluxDB proto also, there're Collectd and Grafana integrations.
	Types	https://github.com	https://github.com	Custom	PostgreSQL				8,11 [60]	DEAD
	Types	https://github.com	https://github.com	Custom	PostgreSQL				8,11 [61]	DEAD
	Timely	https://github.com	https://github.com	Custom	Apache Accumulo				?	DEAD
	Timely	https://github.com	https://github.com	Custom	Apache Accumulo				?	No Updates for half a year. Seems to be dead. Created by NSA
	TrailDB	http://traildb.io	https://github.com	Custom	C				???	General Purpose Event-based database, needs investigation if can be suitable for time-series data
	TrailDB	http://traildb.io	https://github.com	Custom	C				???	General Purpose Event-based database, needs investigation if can be suitable for time-series data
	TritanDB	http://www.tritan.com	https://github.com	Custom	Kotlin					Claims to be optimized for IoT Workloads
	Vaultaire	https://github.com	https://github.com	Custom	Haskell				Ceph-based	DEAD FOR OVER 3 YEARS. SHOULD BE USED ONLY FOR HISTORICAL REASONS. Relies on CEPH for data storage
	Vaultaire	https://github.com	https://github.com	Custom	Haskell				?	DEAD. Per-host monitoring/collection framework. Suitable for a single-host out of the box, needs to be combined with something else to be useful on scale
	Vulcan	https://github.com	https://github.com	Custom	NodeJS				?	DEAD. Per-host monitoring/collection framework. Suitable for a single-host out of the box, needs to be combined with something else to be useful on scale
	Vulcan	https://github.com	https://github.com	Custom	NodeJS				Yes	DEAD. Prometheus-compatible long term storage
	Vulcan	https://github.com	https://github.com	Custom	NodeJS				Yes?	DEAD. Prometheus-compatible long term storage

Output [63]	Diamond [64]	Telegraf [65]	Collectd [66]	Snap [67]	Netdata [68]
amon		V			
ampq		V			
cassandra				V	
cloudwatch	V	V			
datadog	V	V			
elasticsearch		V		V	
etcd				V	
graphite	V	V	V	V	V
grafana				V [69]	
graylog		V			
hana				V	
hawkular				V	
heapster				V	
heka				V	
hostedgraphite	V				
http	V		V		
influxdb	V	V		V	V
kafka		V	V	V	V
kinesis		V			
kairosdb				V	V
librato		V			
mongodb			V		
mqtt	V	V			
mysql	V			V	
nats		V			
nsq		V			
opentsdb		V		V	V
prometheus		V	V		V
rabbitmq	V			V	
riemann	V	V		V	
redis			V		
rrdtool	V				
sentry	V				
sensu			V		
statsd	V				
Name	Source Code	Language	License	Supported Outputs	Extra
Diamond	https://github.com/python-diamond	Python	MIT	Graphite, InfluxDB, mqtt, m	Outputs and collectors are python scripts
Telegraf	https://github.com/influxdata/telegraf	Go	MIT	amon, ampq, cloudwatch, d	Plugins are compile-time
Collectd	https://github.com/collectd/collectd	C	MIT	graphite, http, kafka, mongc	Plugins are .so files
Snap	https://github.com/intelsdi-x/snap	Go	Apache 2.0	influxdb, graphite, opentsdb	All plugins (collectors, outputs) are separate pieces of software that talks over gRPC
Netdata	https://github.com/firehol/netdata	C	GPLv3	graphite, kairosdb, influxdb,	Can work as a standalone monitoring+alerting+dashboard system. Have Alerting capabilities.

This list will contain just names of databases I think would be interesting to test:					
Requirements:					
Tags (Metrics 2.0 or InfluxDB Line Protocol)					TimescaleDB
Reliability					clickhouse
Retentions					m3db
					eventql
					siridb
					Roshi
					DalmatinerDB
					Sidewinder

[1] Self-link: <https://goo.gl/BRqzdG>

This sheet is an attempt to structure basic information about time-series databases (or stuff that can be used as them to some extent).

At this moment I'm against providing performance evaluation of those databases, because it's very hard to create a fair environment, and if it's not fair - test will be useless.

Useful link with some thoughts on TSDBs: <https://misfra.me/2016/04/09/tsdb-list/> (updated once in a while)

I'll only add an OpenSource database (at least basic functionality must be opensource under any OSI approved license)

[2] average for large metrics. If database have compression - better to provide either values (range preferably) and links to the experiments (if any), otherwise - should be a value from whitepaper/docs and if there are any notes - they should be added on a "Comments" field.

[3] For tags support I'll place Yes if and only if there is a history of changes there, e.x. query by dimentions will return the real state of things, not the current one.

[4] <https://docs.google.com/presentation/d/17opE2U2ale1TFYJgr0Z8Dd2fJhy2x0YbReCVeCz6uhA/edit?usp=sharing>

FOSDEM 2018 talk

[5] Tags are not supported at the moment, as per <https://github.com/criteo/biggraphite/blob/master/biggraphite/plugins/tags.py>

[6] It relies on PostgreSQL to store data, without any custom compression, etc

[7] It relies on PostgreSQL to store data, without any custom compression, etc

[8] Custom (not Postgres-Based)
Highly depends on data model (thx. to dzhdanov)

[9] Depends on Clickhouse version and config.

If zero-timestamp is used and double delta compression on Clickhouse side - could be lower than 2.

2 is however average for optimal settings according to the Author (lomik).

According to my own tests with clickhouse and default compression, with "zero-timestamp" turned off it gave:

5.4 - based on autogenerated data

6.5- based on real-world data set

Depends on a dataset even with generic compression algorithm it can go down to about 2.5 b/point (for system metrics for example)

[10] According to my own tests
5.4 - based on autogenerated data
6.5- based on real-world data set

[11] <https://github.com/lomik/graphite-carbon> only supports current state of dimentions.

Though it's possible to implement history of changes

[12] May be a bit better than graphite-clickhouse, but likely not significant as both stores same data in a mostly same way

[13] May be a bit better than graphite-clickhouse, but likely not significant

[14] According to <https://www.elastic.co/blog/elasticsearch-as-a-time-series-data-store> - post optimize 508M for 23M points, $508/23 \approx 22$

Valid for 2.0, might be better in 5.x

[15] According to <https://www.elastic.co/blog/elasticsearch-as-a-time-series-data-store> - post optimize 2200M for 23M points = $2200/23 \approx 96$

Valid for 2.0, might be better in 5.x

[16] <https://julien.danjou.info/talks/storing-metrics-at-scale-with-gnocchi.pdf>
6.25 - avg

[17] <https://julien.danjou.info/talks/storing-metrics-at-scale-with-gnocchi.pdf>
6.25 - avg

[18] 2 bytes/point according to official response: <https://community.influxdata.com/t/how-much-disk-does-influxdb-consume/31>

Depends on the data set it can go below 1 byte/point

[19] According to official documentation, depends on randomness of data

[20] <https://groups.google.com/forum/#!topic/kairosdb-group/B77fsHCyFtk> - according to some random tests.

On a good data it can be 4 bytes per point, on worse data can be worse

[21] <https://groups.google.com/forum/#!topic/kairosdb-group/B77fsHCyFtk> - according to some random tests.

On a good data it can be 4 bytes per point, on worse data can be worse

[22] according to one of the authors.

[23] according to one of the authors.

[24] According to one of the authors

[25] According to one of the authors

[26] Can store them but current version don't have Documented API to query them

[27] Cassandra as storage is ok

Metriktank's own clustering is very basic and relies on kafka partitions (or graphite relays), HA is also possible, but master promotions are manual

[28] <http://opentsdb.net/faq.html>

12 bytes with compression, without HDFS replication.

39 - minimal amount without compression, + 6b/tag

[29] <http://opentsdb.net/faq.html>

12 bytes with compression, without HDFS replication.

39 - minimal amount without compression, + 6b/tag

[30] Depends on compression algorithm and data
<https://prometheus.io/docs/operating/storage/>

[31] Depends on compression algorithm and data
<https://prometheus.io/docs/operating/storage/>

[32] Uncertain, not from docs

[33] Uncertain, not from docs

[34] Excluding metadata overhead

[35] Excluding metadata overhead

[36] TagDB in current master (what will become 1.1.0).

[37] there are several projects that provides clustering for that kind of metrics - github.com/go-graphite/ for whisper-based, [graphite-web](https://github.com/graphite-project/graphite-web) can also do some sort clustering

[38] Excluding metadata overhead

[39] Excluding metadata overhead

[40] Only current state with <https://github.com/kanatohodets/carbonsearch>

Requires separate daemon, query language is rudimental

[41] there are several projects that provides clustering for that kind of metrics - github.com/go-graphite/ for whisper-based, [graphite-web](https://github.com/graphite-project/graphite-web) can also do some sort clustering

[42] ElasticSearch to be precise

[43] Main article about Compression: <https://www.google.com/url?q=https://medium.com/@valyala/victoriametrics-achieving-better-compression-for-time-series-data-than-gorilla-317bc1f95932&sa=D&ust=1559117797555000&usg=AFQjCNELWLSFF1ETBmLTqEOeDLOsWFiKwg>

According to:

<https://medium.com/@valyala/high-cardinality-tsdb-benchmarks-victoriametrics-vs-timescaledb-vs-influxdb-13e6ee64dd6b>

and

<https://medium.com/@valyala/measuring-vertical-scalability-for-time-series-databases-in-google-cloud-92550d78d8ae>

[44] Cluster version is separate from single node one: <https://github.com/VictoriaMetrics/VictoriaMetrics/blob/master/cluster/README.md>

[45] Self-link: <https://goo.gl/BRqzdG>

This sheet is an attempt to structure basic information about time-series databases (or stuff that can be used as them to some extent).

At this moment I'm against providing performance evaluation of those databases, because it's very hard to create a fair environment, and if it's not fair - test will be useless.

Useful link with some thoughts on TSDBs: <https://misfra.me/2016/04/09/tsdb-list/> (updated once in a while)

I'll only add an OpenSource database (at least basic functionality must be opensource under any OSI approved license)

[46] average for large metrics. If database have compression - better to provide either values (range preferably) and links to the experiments (if any), otherwise - should be a value from whitepaper/docs and if there are any notes - they should be added on a "Comments" field.

[47] For tags support I'll place Yes if and only if there is a history of changes there, e.x. query by dimentions will return the real state of things, not the current one.

[48] Depends on dataset, can be worse than 4 bytes on dataset with random floats

[49] Depends on dataset, can be worse than 4 bytes on dataset with random floats

[50] <https://github.com/rackerlabs/blueflood/wiki/FAQ>

[51] <https://github.com/rackerlabs/blueflood/wiki/FAQ>

[52] Excluding constant metadata overhead

[53] Excluding constant metadata overhead

[54] there are several projects that provides clustering for that kind of metrics - github.com/go-graphite/ for whisper-based, graphtie-web can also do some sort clustering

[55] It's UNCLEAR what's the characteristics, because they doesn't specify that in bytes/point. On <https://www.usenix.org/system/files/conference/fast17/fast17-lautenschlager.pdf> they claim to be 20% better than InfluxDB at that point, that was 4 b/point

[56] It's UNCLEAR what's the characteristics, because they doesn't specify that in bytes/point. On <https://www.usenix.org/system/files/conference/fast17/fast17-lautenschlager.pdf> they claim to be 20% better than InfluxDB at that point, that was 4 b/point

[57] Requires external PostgreSQL for tags.

[58] <http://sidwinder.srotya.com/docs/#!/designs/compression>

[59] <http://sidwinder.srotya.com/docs/#!/designs/compression>

[60] Random source:
<https://news.ycombinator.com/item?id=13247598>

[61] Random source:
<https://news.ycombinator.com/item?id=13247598>

[62] Custom, not postgres based

[63] This table is based on what software says about its Outputs. There was no sanity checks on that done by author - e.x. no checks were made to verify if there is any use case for having data in Hana for example

[64] Language: Python
License: MIT

Outputs and collectors are python scripts

[65] Language: Go
License: MIT

Plugins are compile-time

[66] Language: C
License: MIT

Plugins are .so libraries

[67] Language: Go
License: Apache 2.0

All plugins (collector and outputs) are separate pieces that talk over gRPC

[68] Language: C
License: GPLv3

Can work as a standalone system
Have Alerting capabilities.

[69] Direct output to Grafana (<https://github.com/raintank/snap-app>) - it doesn't store data for very long, but allows you to use Grafana to get current state of metrics

See <https://grafana.com/blog/2016/03/31/using-grafana-with-intels-snap-for-ad-hoc-metric-exploration/> for more details

[70] Aerospike has community edition money-free (sends telemetry) and two kinds of commercial. All three supports HA for needed nodes number and replication factor. HA is the main feature. Therefore they opened HA + Clustering for all editions.

[71] As it's a framework - it's not strictly defined by the database.

[72] As it's a framework - it's not strictly defined by the database.