

# Chronix: Long Term Storage and Retrieval Technology for Anomaly Detection in Operational Data

FAST 2017, Santa Clara

---

Florian Lautenschlager, Michael Philippsen, Andreas Kumlehn, and Josef Adersberger



Florian.Lautenschlager@qaware.de



flolaut

# Detecting Anomalies in Running Software matters

Various kinds of anomalies:


- **Resource consumption:** anomalous memory consumption, high CPU usage, ...
- **Sporadic failure:** blocking state, deadlock, dirty read, ...
- **Security:** port scanning activity, short frequent login attempts, ...

➔ **Economic or reputation loss.**

Detection is a complex task:

- **Multiple components:** Database, Service Discovery, Configuration Service, ...
- **Different technologies:** Go, Java, Java-Script, Python, ...
- **Various transport protocols:** HTTP, Protocol Buffers, Thrift, JSON, ...

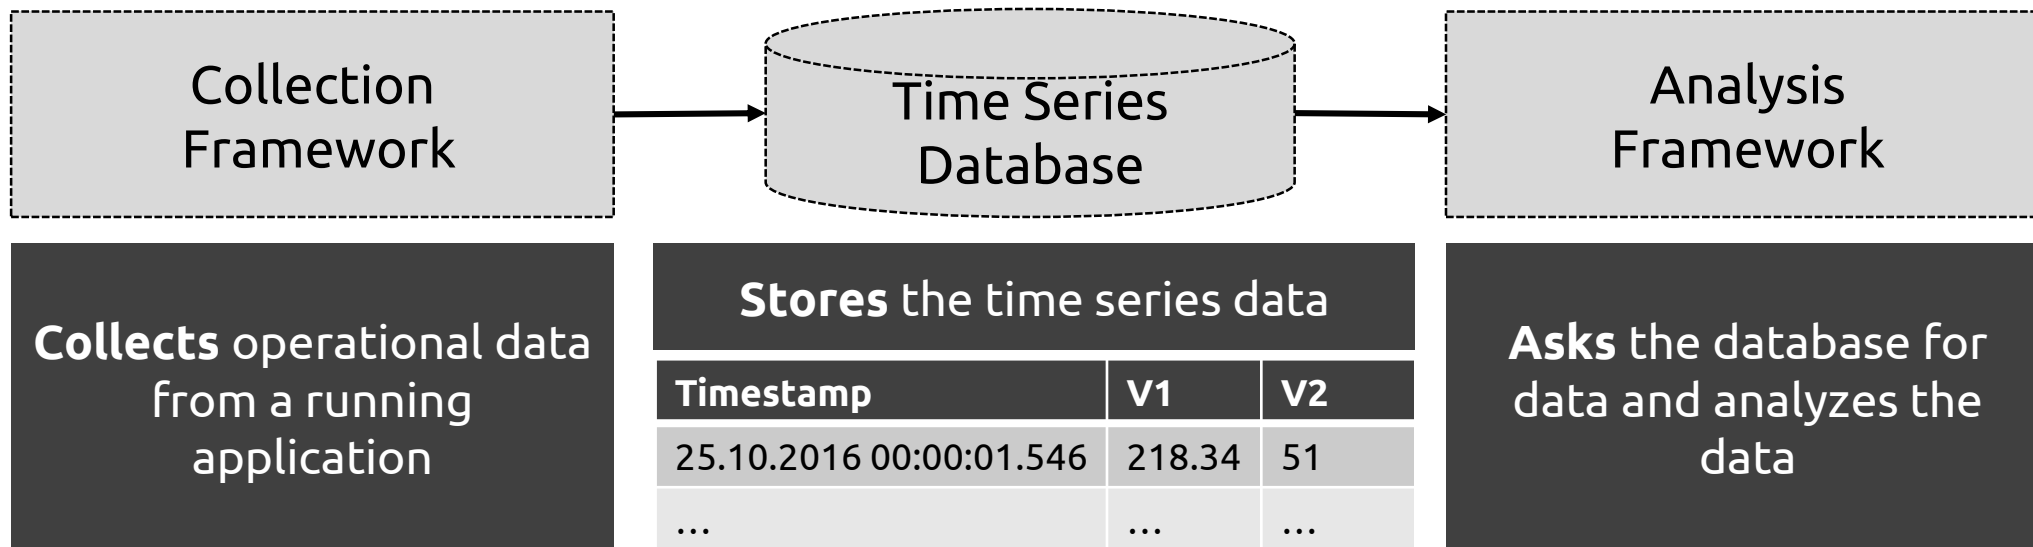
# Anomaly Detection Tool Chain for Operational Data



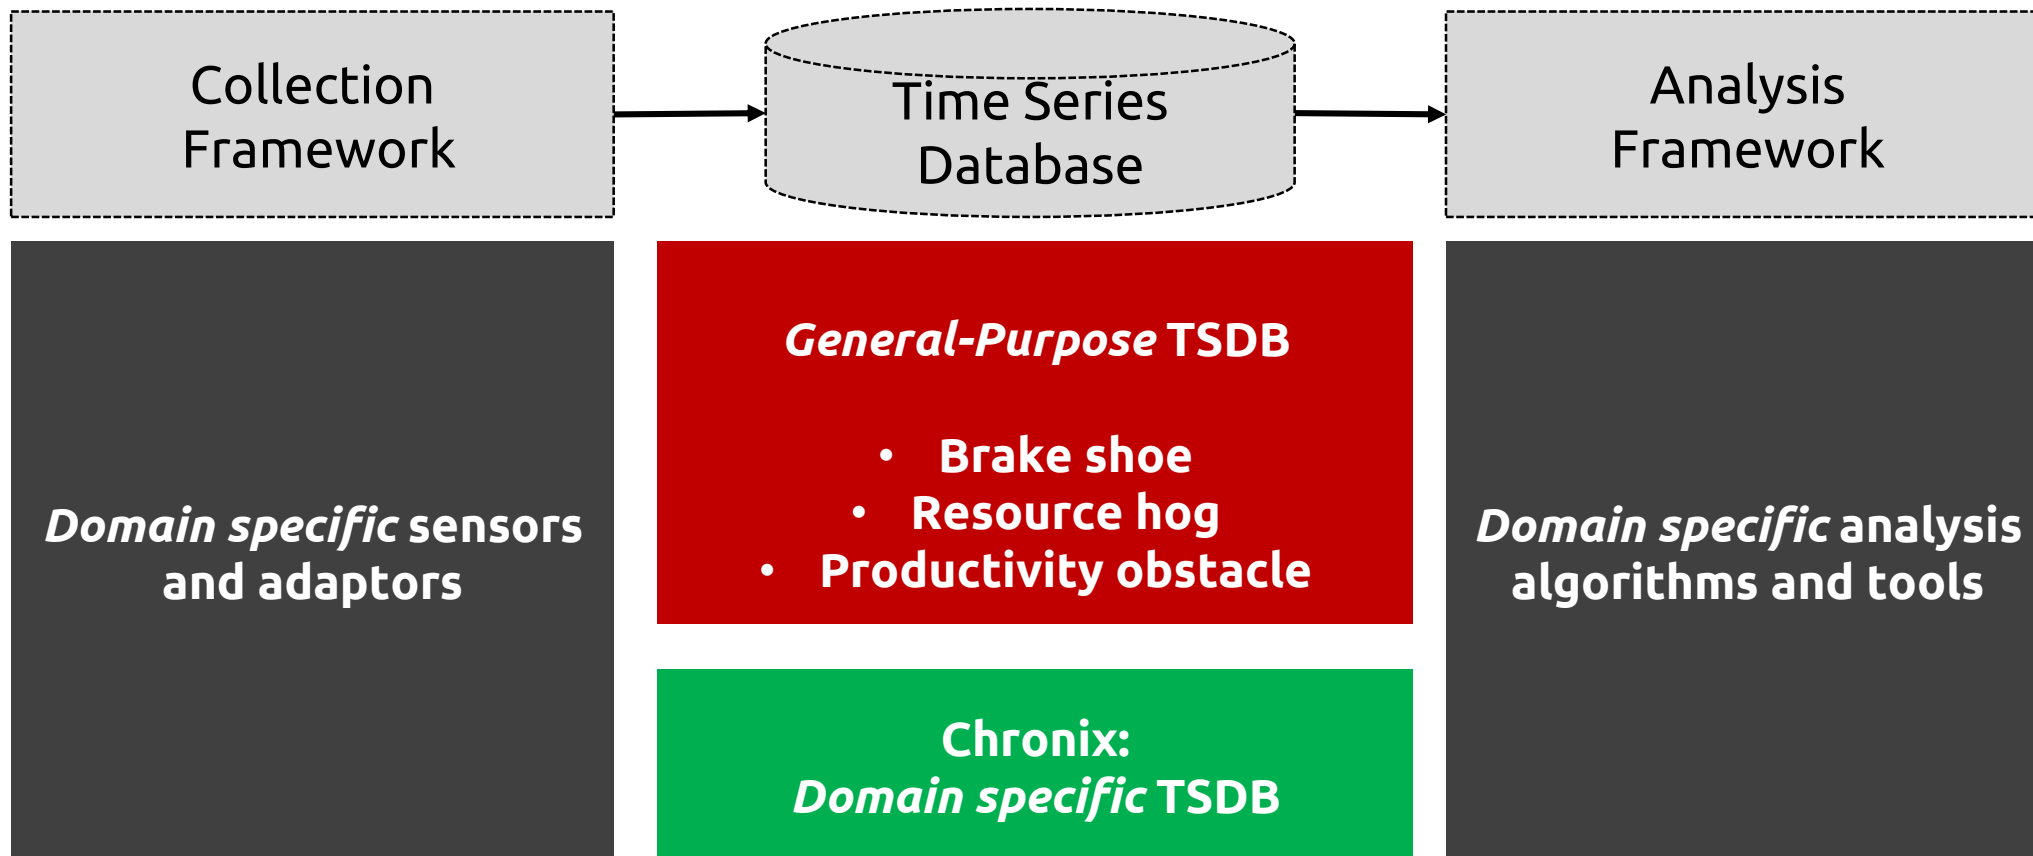
Types of operational data:

- **Metrics:** scalar values, e.g., rates, runtimes, total hits, counters, ...
- **Events:** single occurrences, e.g., a user's login, product order, ...
- **Traces:** sequences within a software system, e.g., the called methods, ...

# Anomaly Detection Tool Chain for Operational Data



# Anomaly Detection Tool Chain for Operational Data



# State of the art: General-purpose TSDBs in Anomaly Detection


	Graphite	InfluxDB	OpenTSDB	KairosDB	Prometheus		Chronix
Generic data model	○	◐	○	◐	○	<p><b>No support for data types</b> = Productivity obstacle</p> <p><b>No support for analyses</b> = Productivity obstacle = Brake shoe</p> <p><b>High memory footprint</b> = Performance hog <b>High storage demands</b> = Performance hog <b>Loss of historical data</b> = Brake shoe</p>	●
Analysis support	◐	◐	○	◐	◐		●
Lossless long term storage	○	●	●	●	◐		●

# 7 Bullets for the domain of Anomaly Detection



- 1 Option to pre-compute an extra representation of the data
- 2 Optional timestamp compression for almost-periodic time series
- 3 Records that meet the needs of the domain
- 4 Compression technique that suits the domain's data
- 5 Underlying multi-dimensional storage
- 6 Domain specific query language with server-side evaluation
- 7 Domain specific commissioning of configuration parameters

# Running Example: Almost-periodic time series with operational data



Timestamp	Value	Metric	Process	Host
25.10.2016 00:00:01.546	218.34	ingester\time	SmartHub	QAMUC
25.10.2016 00:00:06.718	218.37	ingester\time	SmartHub	QAMUC
25.10.2016 00:00:11.891	218.49	ingester\time	SmartHub	QAMUC
25.10.2016 00:00:16.964	218.52	ingester\time	SmartHub	QAMUC
...	...	...	...	...
...	...	...	...	...





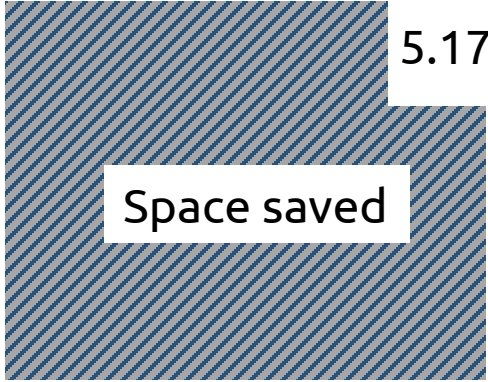
# Option to pre-compute data to speed up analyses


Timestamp	Value	Metric	Process	Host	SAX
25.10.2016 00:00:01.546	218.34	ingester\time	SmartHub	QAMUC	A
25.10.2016 00:00:06.718	218.37	ingester\time	SmartHub	QAMUC	B
25.10.2016 00:00:11.891	218.49	ingester\time	SmartHub	QAMUC	C
25.10.2016 00:00:16.964	218.52	ingester\time	SmartHub	QAMUC	D
...	...	...	...	...	...
...	...	...	...	...	...

- Chronix is **lossless**: it keeps all details because the analyses are ad-hoc and may need them.
- Chronix offers a **programming interface** for adding extra domain specific “columns”. Examples: Fourier transformation, Symbolic Aggregate approxImation (SAX), etc.
- Added “columns” speed up anomaly detection queries.

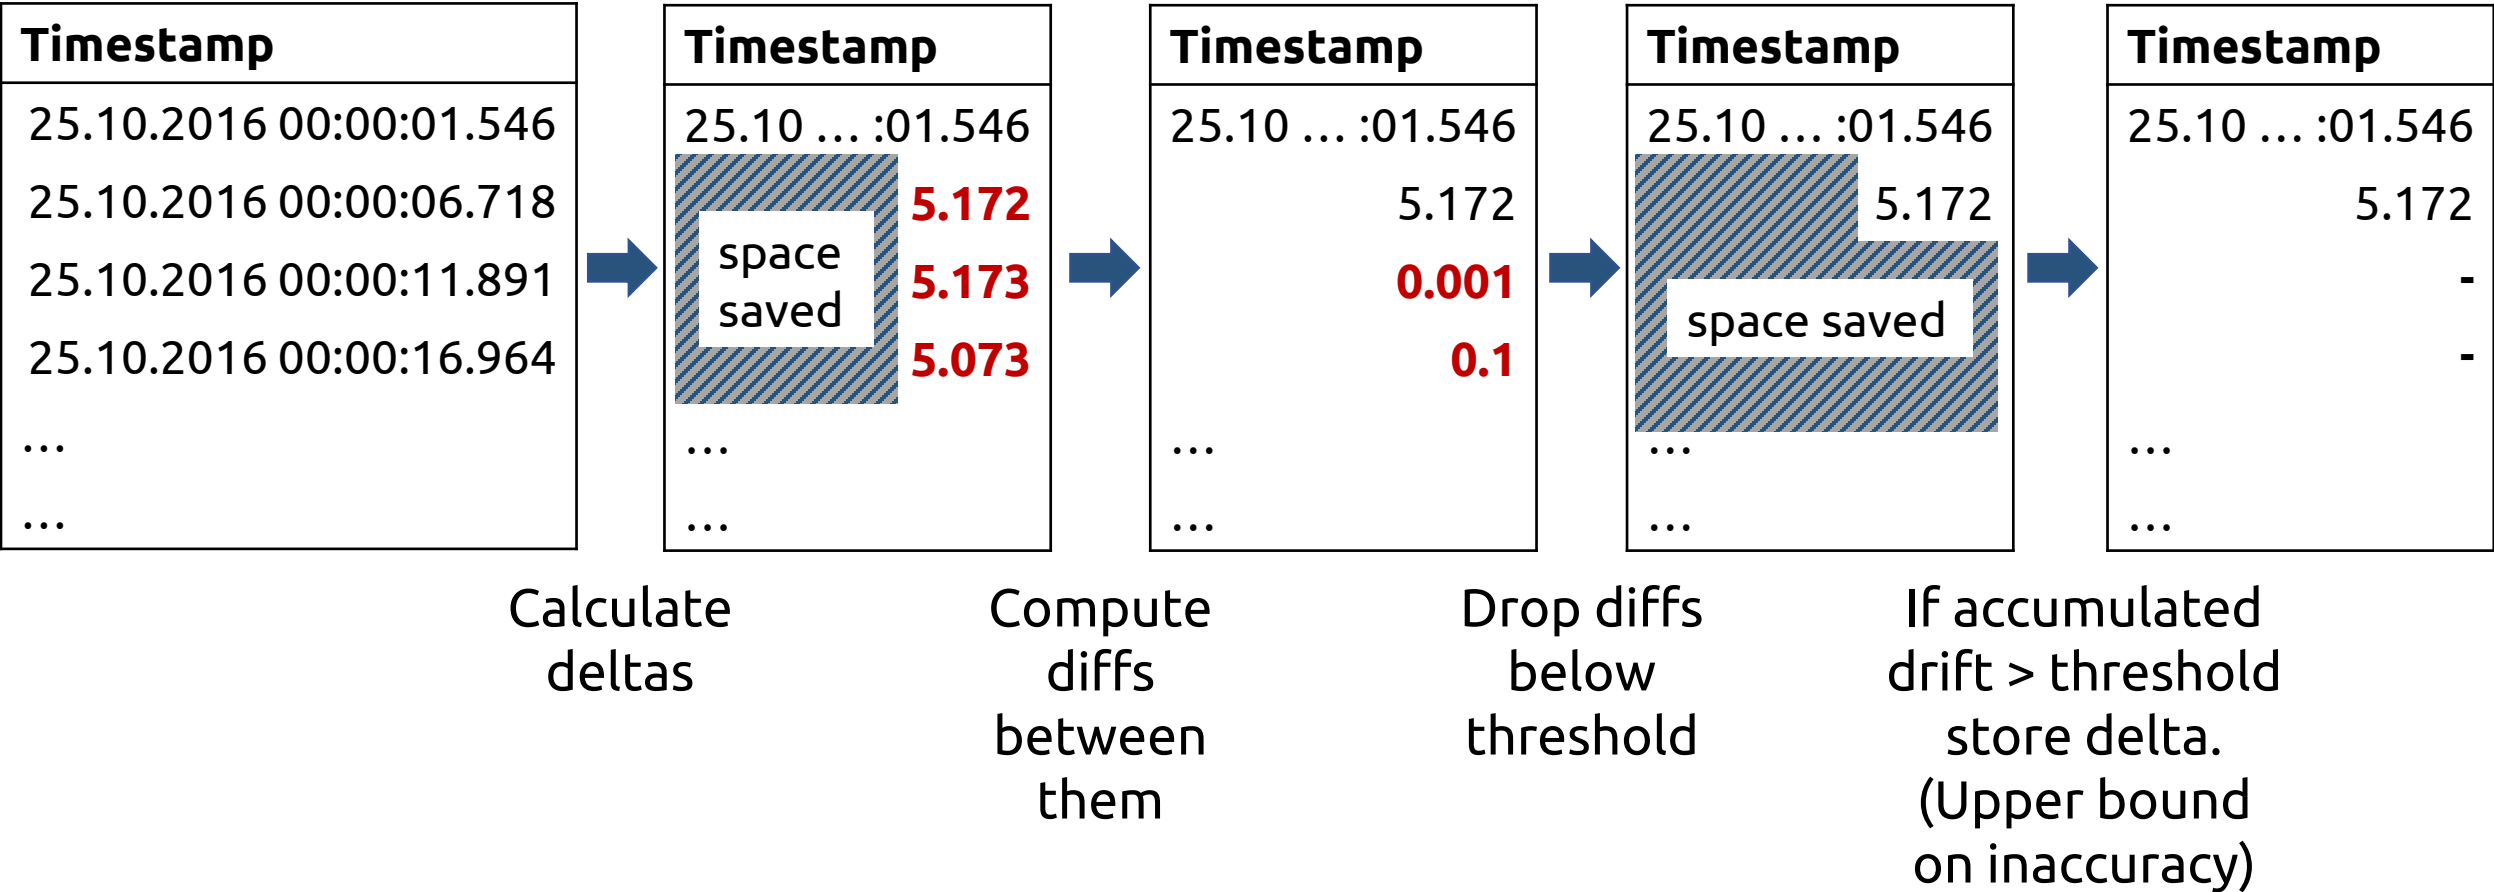


# Optional timestamp compaction

Timestamp	Value	Metric	Process	Host	SAX
25.10.2016 00:00:01.546	218.34	ingester\time	SmartHub	QAMUC	A
 5.172	218.37	ingester\time	SmartHub	QAMUC	B
-	218.49	ingester\time	SmartHub	QAMUC	C
-	218.52	ingester\time	SmartHub	QAMUC	D
...	...	...	...	...	...
...	...	...	...	...	...

- It suffices to be able to reconstruct approximate timestamps for almost-periodic time series.
- **Date-Delta-Compaction**
- Chronix is **functionally lossless** as it keeps all *relevant* details.
- The tolerable degree of inaccuracy is a Configuration Parameter of 

# Date-Delta-Compaction



# Domain specific data characteristics

Timestamp	Value	Metric	Process	Host	SAX	
25.10.2016 00:00:01.546	218.34	ingester\time	SmartHub	QAMUC	A	
	5.172	218.37	ingester\time	SmartHub	QAMUC	B
	-	218.49	ingester\time	SmartHub	QAMUC	C
	-	218.52	ingester\time	SmartHub	QAMUC	D
	...	...	...	...	...	
		...	...	...	...	

Many anomaly detection tasks need blocks of data rather than "lines".

Some compression techniques work better than others.

"Columns" with repetitive values.



# Records that meet the needs of the domain

Therefore:

**Record := Attributes + Start + End + Type + Data Chunk**

Timestamp	Value	Metric	Process	Host	SAX
25.10.2016 00:00:01.546	218.34	ingester\time	SmartHub	QAMUC	A
5.172	218.37	ingester\time	SmartHub	QAMUC	B
-	218.49	ingester\time	SmartHub	QAMUC	C
-	218.52	ingester\time	SmartHub	QAMUC	D
1	1	2	2	2	1
...	...	...	...	...	...

chunk & convert



Record

metric: ingester\time  
process: SmartHub  
host: QAMUC  
start: 25.10.2016 00:00:01.546  
end: ...  
type: metric  
data:

Timestamp	Value	SAX
25.10.2016 00:00:01.546	218.34	A
5.172	218.37	B
-	218.49	C
-	218.52	D

**BLOB**

- Chronix offers a **programming interface** to implement time series specific records.
- Chronix exploits repetitiveness and bundles "lines" into data chunks.

- The chunk size is a Configuration Parameter of





# Compression technique that suits the domain's data

Record

metric: ingester\time  
process: SmartHub  
host: QAMUC  
start: 25.10.2016 00:00:01.546  
end: ...  
type: metric  
data:

Timestamp	Value	SAX
25.10.2016 00:00:01.546	218.34	A
5.172	218.37	B
-	218.49	C
-	218.52	D

serialize & compress



Record

metric: ingester\time  
process: SmartHub  
host: QAMUC  
start: 25.10.2016 00:00:01.546  
end: ...  
type: metric  
data: **00105e0 e6b0 343b 9c74 080  
7bc 0804 e7d5 0804 00105f0**

**Compressed BLOB**

- Chronix exploits that domain data often has small increments, recurring patterns, etc.
- Chronix uses a lossless compression technique that minimizes (record sizes + index sizes).
- The choice of compression technique is a Configuration Parameter of



# Underlying multi-dimensional storage

Record

metric: **ingester\time**  
process: SmartHub  
host: **QAMUC**  
start: 25.10.2016 00:00:01.546  
end: ...  
type: **metric**  
data: 00105e0 e6b0 343b 9c74 080  
7bc 0804 e7d5 0804 00105f0

Record

metric: **ingester\methods**  
process: SmartHub  
host: **QAMUC**  
start: 25.10.2016 00:00:01.546  
end: ...  
type: **trace**  
data: d65fa01 7ab2 433c 7c8e f123  
2ca 0713 a8f5 926b 01006e1

```
q=host:QAMUC AND metric:ingester*  
AND type:[metric OR trace]  
AND end:NOW-7MONTH
```

By using a multi-dimensional storage ...

- ... Chronix supports explorative analyses.
  - Attributes are visible to the storage and indexed.
  - Users can use any combination to find a record.
- ... Chronix supports correlating analyses.
  - Every type of data can be stored.
  - Queries can use and combine types.



# Domain specific query language with server-side evaluation

	<i>Graphite</i>	<i>InfluxDB</i>	<i>OpenTSDB</i>	<i>KairosDB</i>	<i>Prometheus</i>	<i>Chronix</i>
Basic						
distinct	×	✓	×	×	×	✓
integral	✓	×	×	×	×	✓
min/max/sum	✓	✓	✓	✓	✓	✓
count	×	✓	✓	✓	✓	✓
avg/median	✓	✓	✓	✓	✓	✓
bottom/top	×	✓	×	×	✓	✓
first	✓	✓	×	✓	×	✓
last	✓	✓	×	✓	×	✓
percentile (p)	✓	✓	✓	✓	✓	✓
stddev	✓	✓	✓	✓	✓	✓
derivative	✓	✓	×	✓	✓	✓
nnderivative	✓	✓	×	×	×	✓
diff	×	✓	×	✓	✓	✓
movavg	✓	✓	×	×	×	✓
divide/scale	✓	✓	×	✓	✓	✓
High-level						
sax [33]	×	×	×	×	×	✓
fastdtw [38]	×	×	×	×	×	✓
outlier	×	×	×	×	×	✓
trend	×	×	×	×	×	✓
frequency	×	×	×	×	×	✓
grpsize	×	×	×	×	×	✓
split	×	×	×	×	×	✓

basic functions

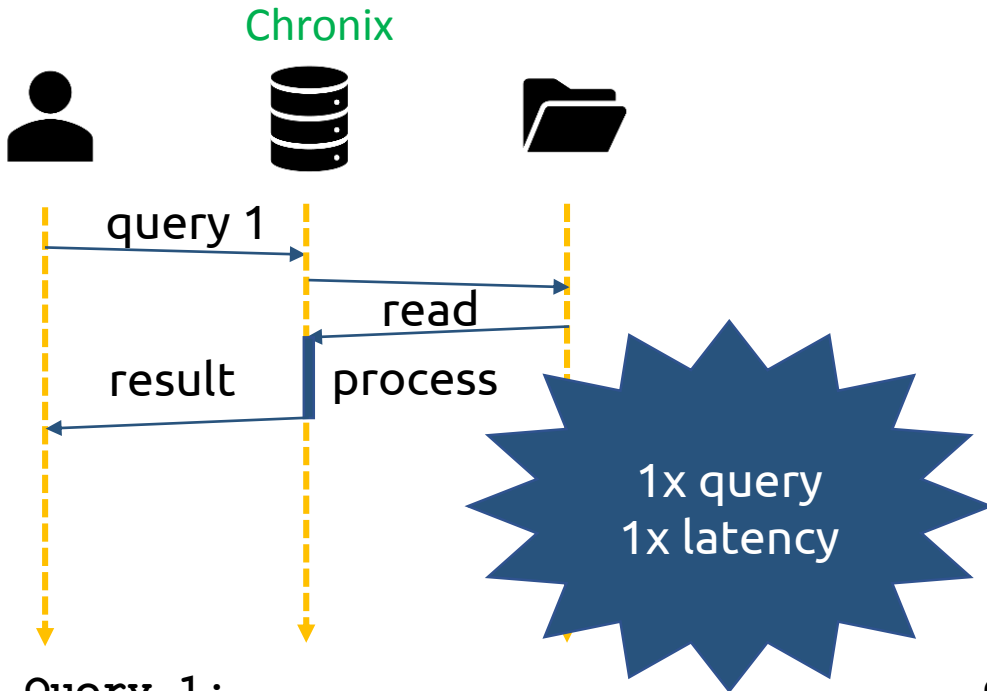
also needed for anomaly detection

- Chronix offers not just basic functions but also high-level built-in domain specific analysis functions.
- Chronix evaluates functions server-side for speed.
- Chronix offers a **plug-in interface** to add functions.



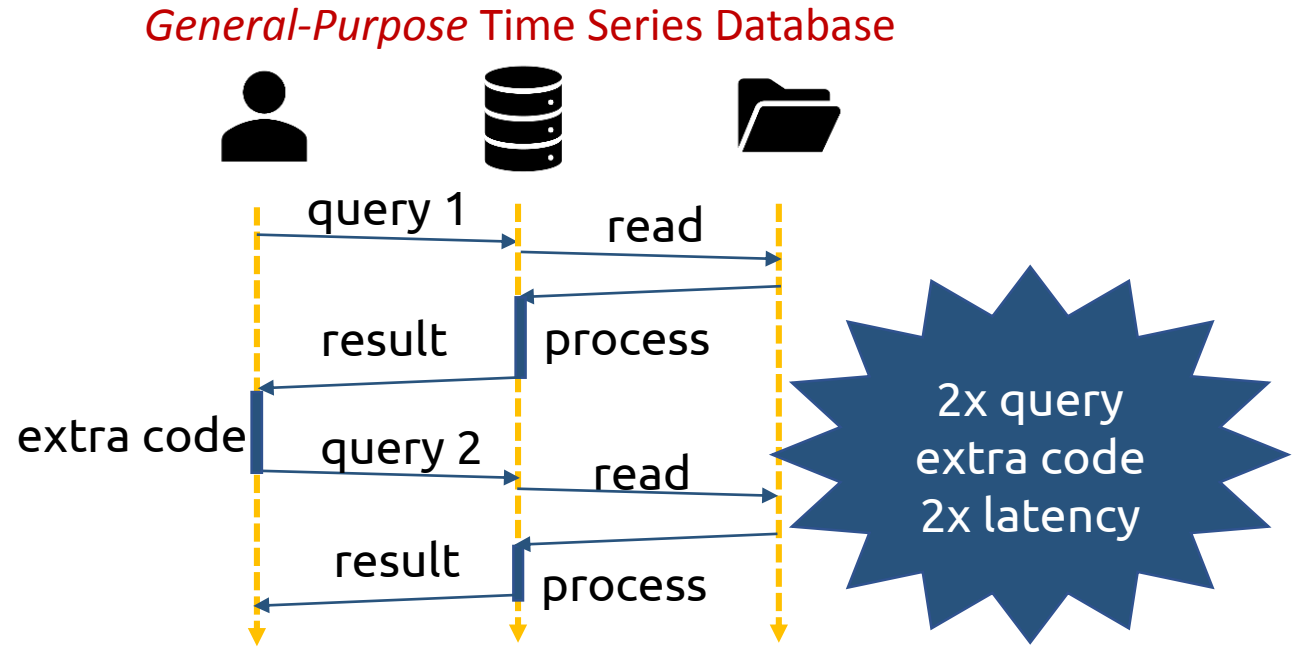


# Domain specific query language with server-side evaluation



Query 1:  
`q=metric:ingester\time`  
`& cf=outlier`

high-level  
function



Query 1:  
`select q(0.25,time),q(0.75,time) from ingester`

Calculate threshold

extra code

Query 2:  
`select time from ingester where time >= threshold`

- Chronix achieves more programming comfort & fast results.

# Operational data of 5 industry projects

	Description	Interval (sec)	Pairs (mio)	Time series
<b>P1</b>	Application for searching car maintenance and repair instructions. (8 app sever, 20 search server)	30	2,4	1,080
<b>P2</b>	Retail application for orders, billing, and customer relations. (1 database, 2 app server)	60	331.4	8,567
<b>P3</b>	Sales application of a car manufacturer. (1 database, 2 app servers)	30	162.6	4,538
<b>P4</b>	Service application for modern cars (music streaming)	1	metric 3.9 lsof 0.4 strace 12.1	500
<b>P5</b>	Manage the compatibility of software components in a car.	60	3,762.3	24,055
<b>Total</b>			<b>4,275.1</b>	<b>38,740</b>

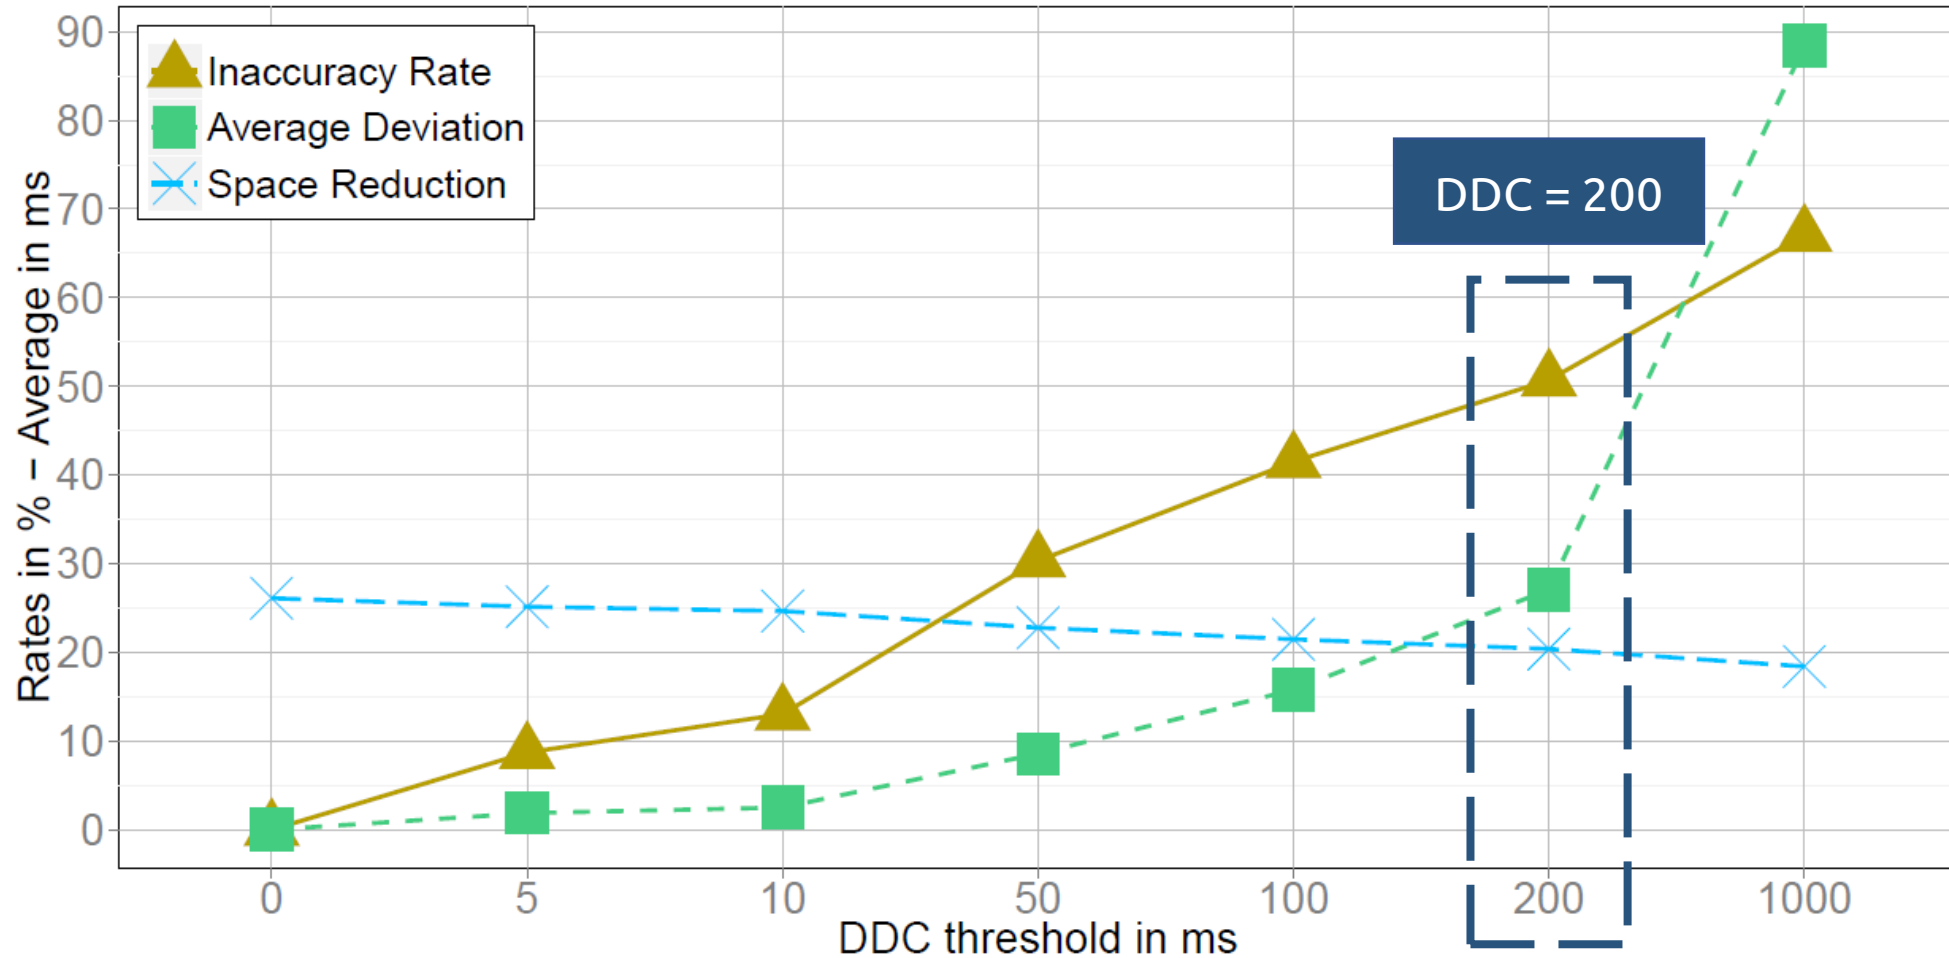
used for



used for the Evaluation



# Best threshold for the Date-Delta-Compaction



# Operational data of 3 (of 5) industry projects →

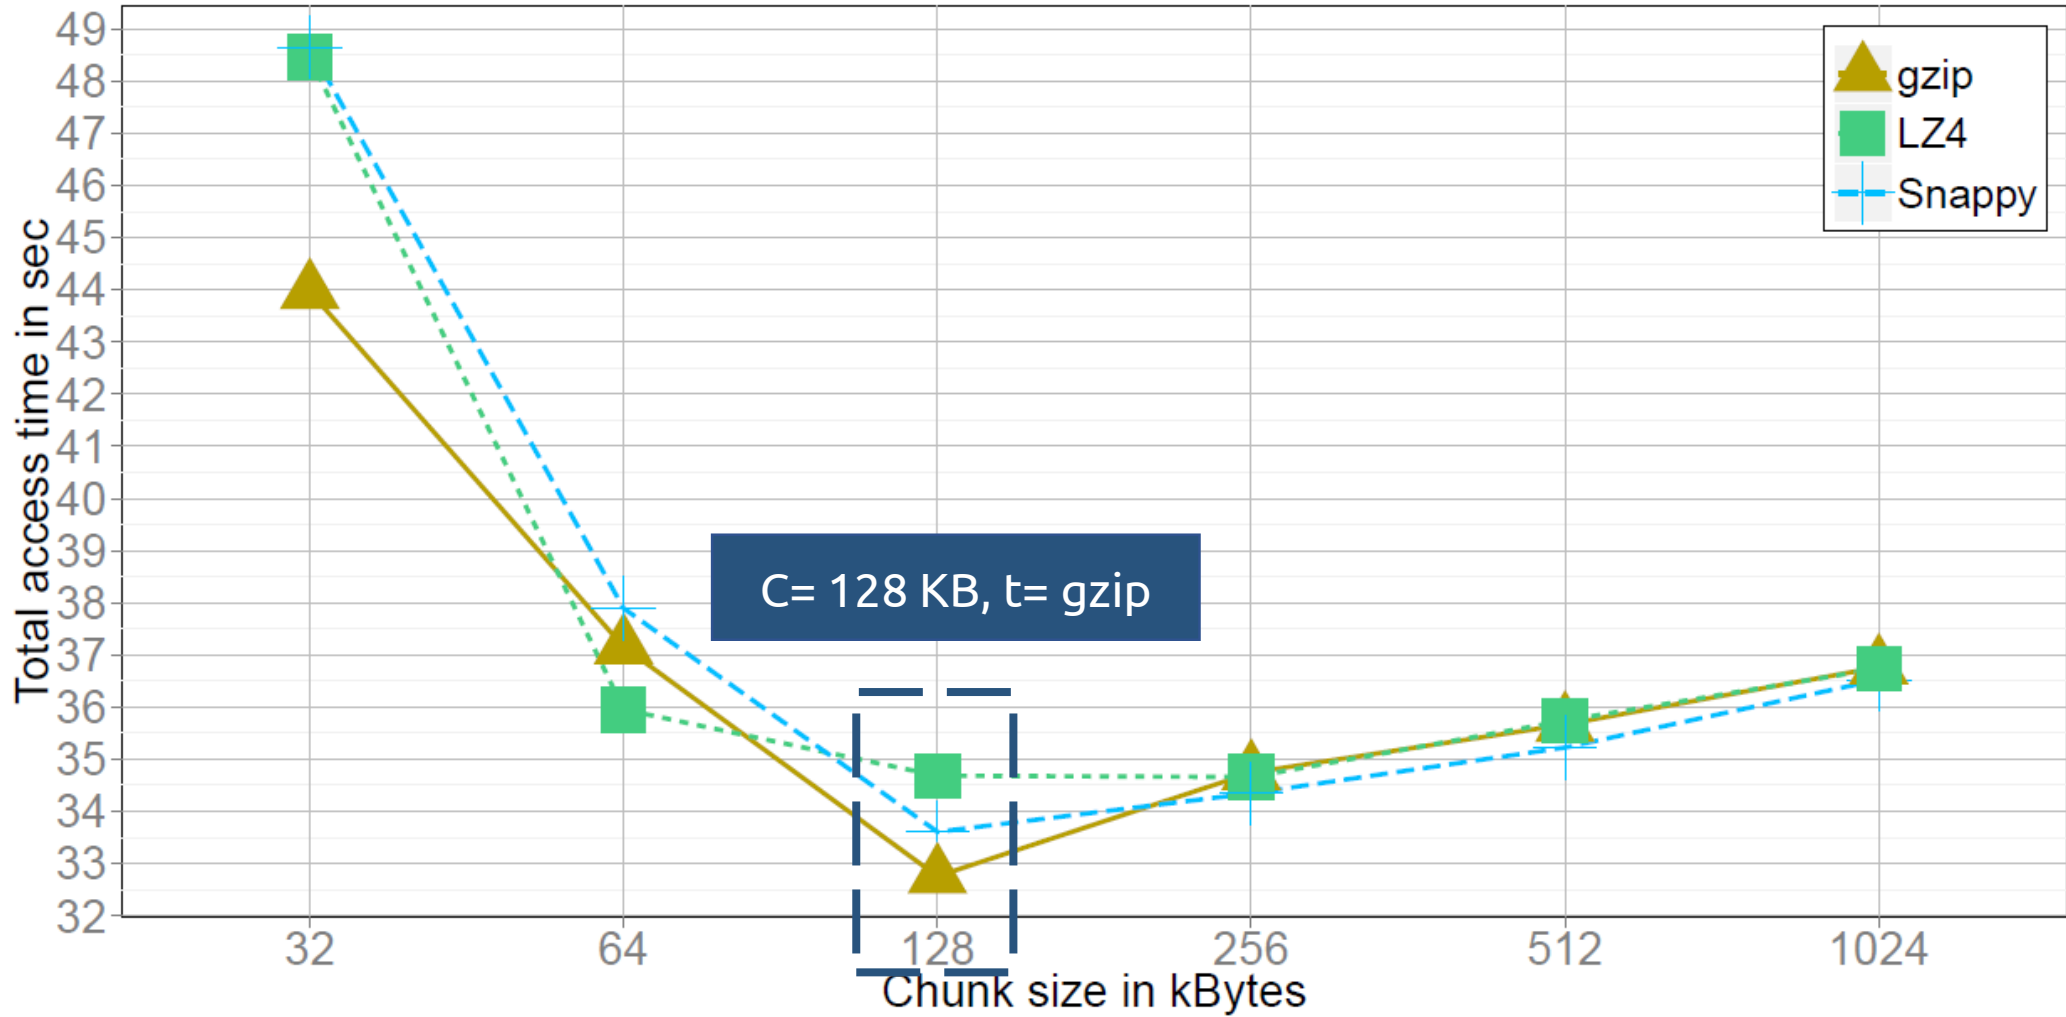


	Description	Interval (sec)	Pairs (mio)	Time series	Query Mix	
					r	q
<b>P1</b>	Application for searching car maintenance and repair instructions. (8 app sever, 20 search server)	30	2,4	1,080	91	2
					56	1
					28	3
<b>P2</b>	Retail application for orders, billing, and customer relations. (1 database, 2 app server)	60	331.4	8,567	21	5
					7	30
					0.5	15
<b>P3</b>	Sales application of a car manufacturer. (1 database, 2 app servers)	30	162.6	4,538	1	30
					0.5	15
					...	...
<b>P4</b>	...	...	...	...	...	...
<b>P5</b>	...	...	...	...	...	...
<b>Total</b>			<b>4,275.1</b>	<b>38,740</b>		

r = range (days)  
q = # of queries



# Best compression technique & Best chunk size for query mix

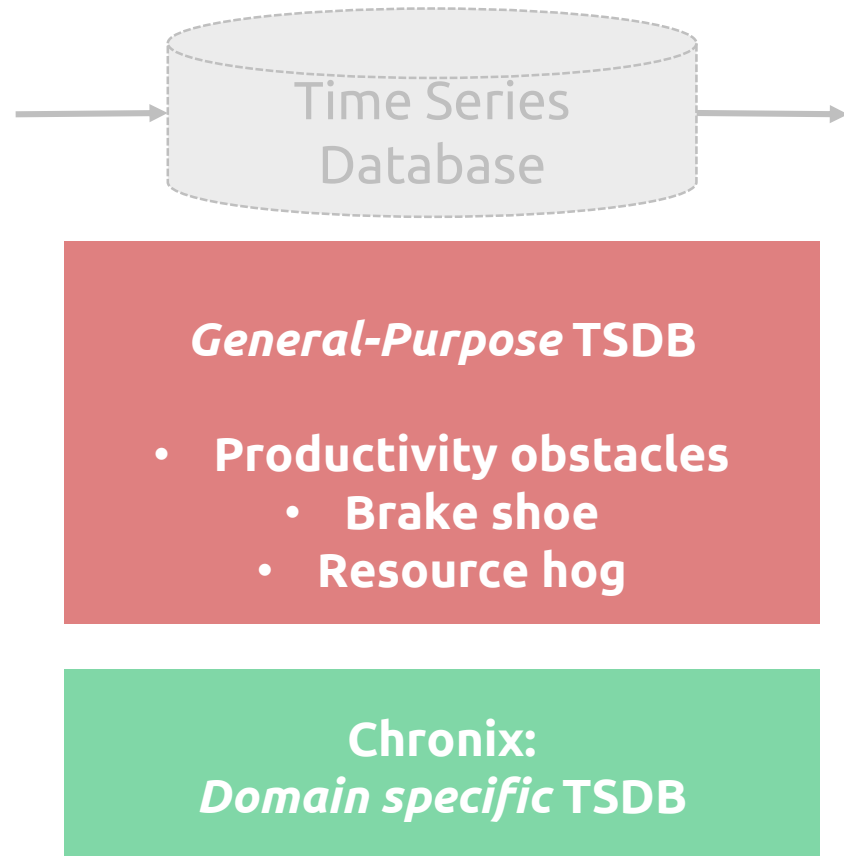


# Operational data of 2 of (5) industry projects → Evaluation

	Description	Interval (sec)	Pairs (mio)	Time series	Query Mix r q b h
<b>P1</b>	...	...	...	...	⋮ ⋮ ⋮ ⋮
<b>P2</b>	...	...	...	...	⋮ ⋮ ⋮ ⋮
<b>P3</b>	...	...	...	...	⋮ ⋮ ⋮ ⋮
<b>P4</b>	Service application for modern cars (music streaming)	1	metric 3.9	500	180
			lsof 0.4		2
			strace 12.1		2
					1
					2
<b>P5</b>	Manage the compatibility of software components in a car.	60	3,762.3	24,055	21
					12
					2
					6
					6
	14				
	8				
	7				
	15				
	5				
	10				
	1				
	11				
	6				
	6				
	0.5				
	1				
	1				
	2				
<b>Total</b>			<b>4,275.1</b>	<b>38,740</b>	

**r** = range (days)  
**q** = # of queries  
**b** = # of basis queries  
**h** = # of high-level queries

# Quantitative comparison



TSDBs under test

Comparisons

**InfluxDB**  
**OpenTSDB**  
**KairosDB**

**Chronix**

- Memory footprint
- Storage demand
- Data retrieval times
- Query mix runtimes

## a) Memory footprint

Memory footprint of the databases (in MB)

	<b>InfluxDB</b>	<b>OpenTSDB</b>	<b>KairosDB</b>	<b>Chronix</b>
<b>Initially after startup (processes up and running)</b>	33	2,726	8,763	446
<b>Maximal memory usage during import</b>	10,336	10,111	18,905	7,002
<b>Maximal memory usage during query</b>	8,269	9,712	11,230	4,792

 **Chronix has a 34% – 69% smaller memory footprint.**



## b) Storage demand

Storage demand (in GB)


	<b>Raw data</b>	<b>InfluxDB</b>	<b>OpenTSDB</b>	<b>KairosDB</b>	<b>Chronix</b>
Project 4	1.2	0.2	0.2	0.3	0.1
Project 5	107.0	10.7	16.9	26.5	8.6
<b>total</b>	108.2	10.9	17.1	26.8	8.7

 **Chronix saves 20% – 68% of the storage space.**

## c) Data retrieval times

Data retrieval times for 20 · 58 queries (in s)

<b>r</b>	<b>q</b>	<b>InfluxDB</b>	<b>OpenTSDB</b>	<b>KairosDB</b>	<b>Chronix</b>
<b>0.5</b>	<b>2</b>	4.3	2.8	4.4	0.9
<b>1</b>	<b>11</b>	5.5	5.6	6.6	5.3
<b>7</b>	<b>15</b>	34.1	17.4	26.8	7.0
<b>14</b>	<b>8</b>	36.2	14.2	25.5	4.0
<b>21</b>	<b>12</b>	76.5	29.8	55.0	6.0
<b>28</b>	<b>5</b>	7.9	3.9	5.6	0.5
<b>56</b>	<b>1</b>	35.4	12.4	24.1	1.2
<b>91</b>	<b>2</b>	47.5	15.5	33.8	1.1
<b>180</b>	<b>2</b>	96.7	36.7	66.6	1.1
<b>total</b>		343.8	138.3	248.4	27.1

 **Chronix saves 80% – 92% on data retrieval times.**

## d) Query mix runtimes

Runtimes of 20 · 75 b- and h-queries (in s)

	q		InfluxDB	OpenTSDB	KairosDB	Chronix
Basic (b)	4	avg	0.9	6.1	9.8	4.4
	5	max	1.3	8.4	9.1	6.0
	3	min	0.7	2.7	5.3	2.8
	3	stddev.	6.7	16.7	21.1	2.3
	5	sum	0.7	6.0	12.0	2.0
	4	count	0.8	5.5	10.5	1.0
	8	perc.	10.2	25.8	34.5	8.6
High-level (h)	12	outlier	30.7	29.1	117.6	18.9
	14	trend	162.7	50.4	100.6	30.2
	11	frequency	47.3	23.9	45.7	16.3
	3	grpsize	218.9	2927.8	206.3	29.6
	3	split	123.1	2893.9	47.9	37.2
	75	total	604.0	5996.3	620.4	159.3

more important {

➔ **Chronix saves 73% – 97% of the runtime of analyzing queries.**

# → Chronix unleashes Anomaly Detection tasks

## 7 domain specific levers to unleash Anomaly Detection

1. Option to pre-compute an extra representation of the data
2. Optional timestamp compression for almost-periodic time series
3. Records that meet the needs of the domain
4. Compression technique that suits the domain's data
5. Underlying multi-dimensional storage
6. Domain specific query language with server-side evaluation
7. Domain specific commissioning of configuration parameters

## 4 beneficial performance effects

- Chronix has a 34% – 69% smaller memory footprint.
- Chronix saves 20% – 68% of the storage space.
- Chronix saves 80% – 92% on data retrieval time.
- Chronix saves 73% – 97% of the runtime of analyzing queries.



[www.chronix.io](http://www.chronix.io)  
open source