Preserving Recomputability of Results from Big Data Transformation Workflows

Matthias Kricke (Leipzig University)
Martin Grimmer (Leipzig University)
Michael Schmeißer (mgm)

München/HQ Bamberg Berlin Đà Nẵng Dresden Grenoble Hamburg Cologne Leipzig Nuremberg Prague Washington Zug

Information is constantly acquired

Information from external sources is used to create value

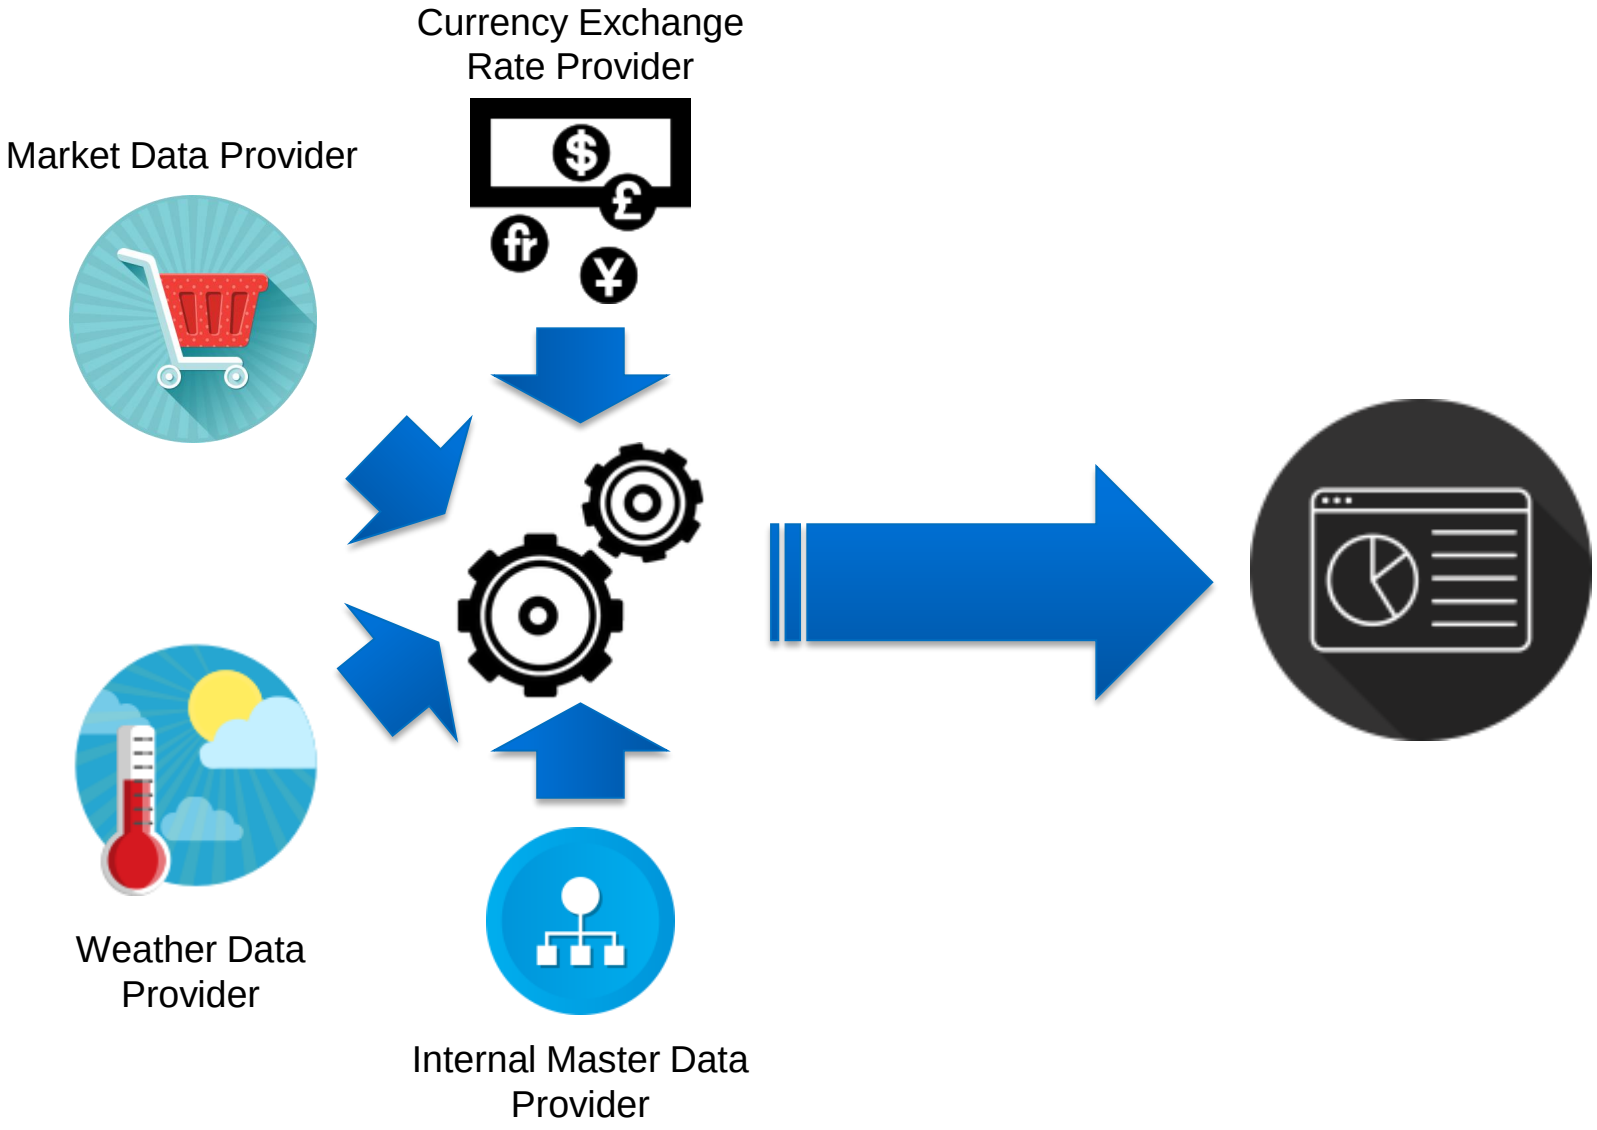

Market Research Use Case

- Storage and processing of highly diverse event data from external sources
- Fully automated production line despite heterogeneous data quality
- Asynchronous integration of manual process steps

mgm

Requirements for Recomputability

- Possibility to recompute delivered products at any time from the raw data, for instance to deliver them again or adapt them selectively based on customer demands
- The originally computed result needs to be annotated with all information required to reproduce it
- The recomputation should be able to take place fully automatic

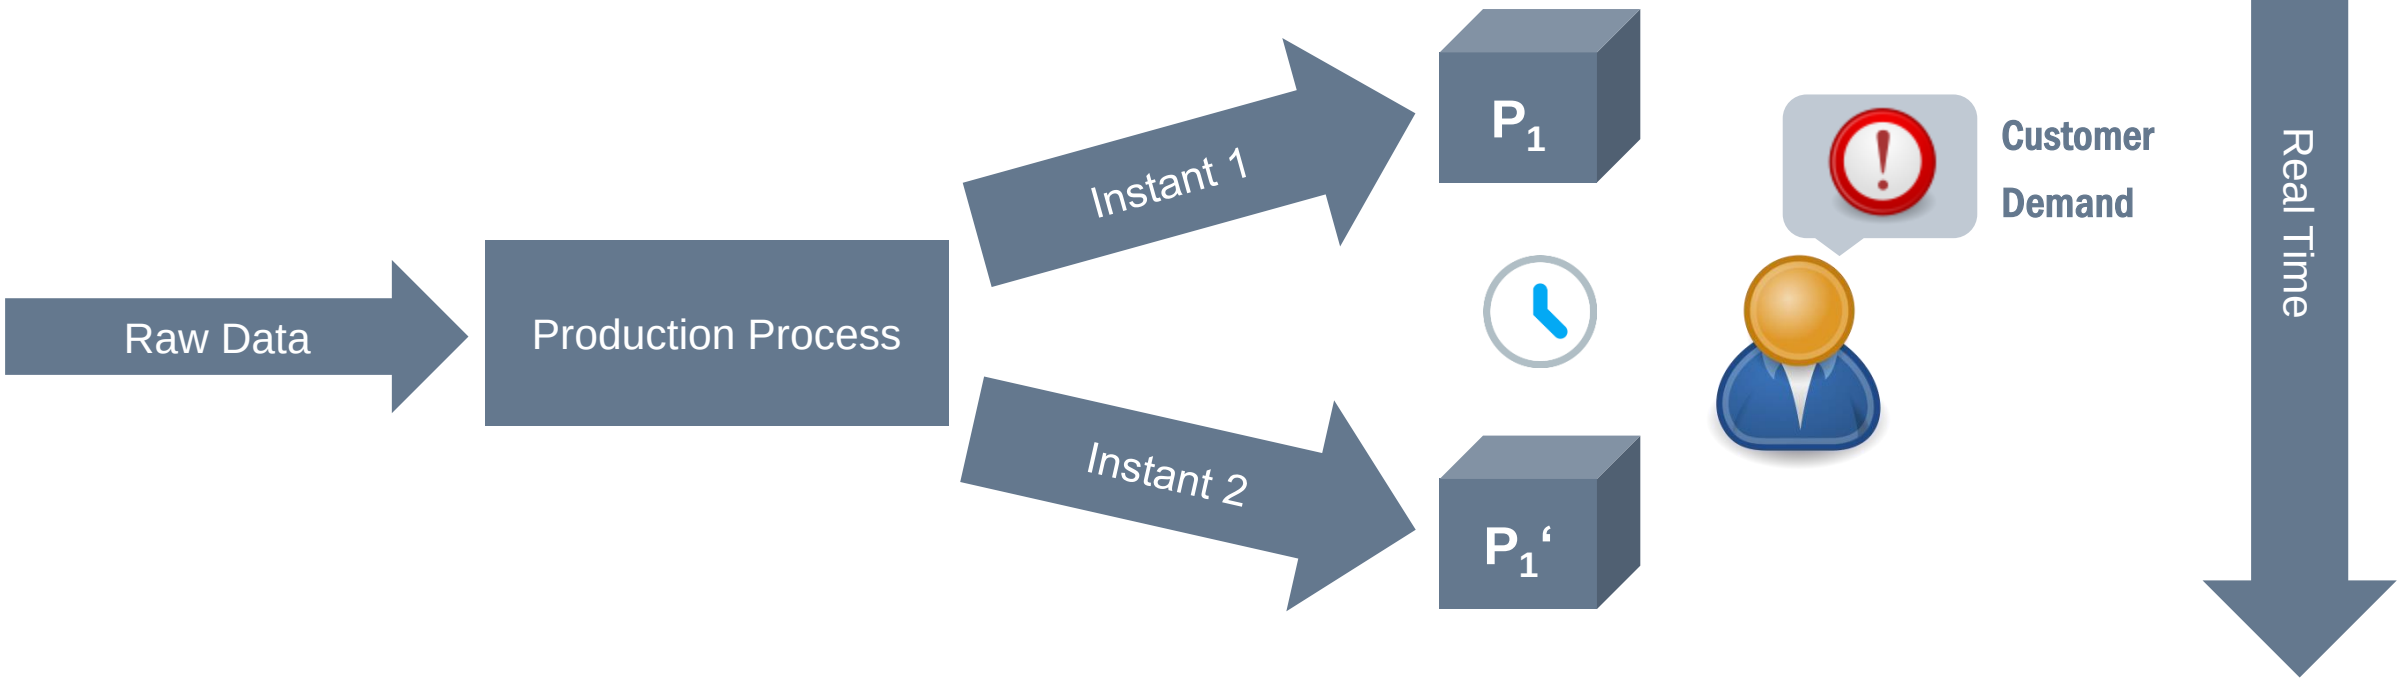

Customers expect stability of delivered data products

Turnover in € 02/17	Germany	United Kingdom	France	Total
TVs	523,239	499,021	607,201	**1,629,461**
Smartphones	1,239,402	1,340,023	1,234,481	**3,813,906**
Tablets	829,012	1,022,339	1,032,211	**2,883,562**
Total	**2,591,653**	**2,861,383**	**2,873,893**	**8,326,929**

Turnover in € 02/17	Germany	United Kingdom	France	Total
TVs	523,239	499,021	607,201	**1,629,461**
Smartphones	1,239,402	1,340,023	1,234,481	**3,813,906**
Tablets	829,012	1,022,339	1,032,211	**2,883,562**
Convertibles	11,428	9,210	17,329	**37,967**
Total	**2,603,081**	**2,870,593**	**2,891,222**	**8,364,896**

Customers expect stability of delivered data products

Turnover in € 02/17	Germany	United Kingdom	France	Total
TVs	523,239	499,021	607,201	**1,629,461**
Smartphones	1,239,402	1,340,023	1,234,481	**3,813,906**
Tablets	829,012	1,022,339	1,032,211	**2,883,562**
Total	**2,591,653**	**2,861,383**	**2,873,893**	**8,326,929**

Turnover in € 02/17	Germany	United Kingdom	France	Total
TVs	523,239	499,023	607,201	**1,629,463**
Smartphones	1,239,402	1,340,026	1,234,481	**3,813,909**
Tablets	959,012	1,012,341	1,022,211	**2,993,564**
Convertibles	21,428	19,211	27,329	**67,968**
Total	**2,743,081**	**2,870,601**	**2,891,222**	**8,504,904**

External systems may not offer everything that is needed by our data transformation process

Full History

Low Latency

High Throughput

Availability

Time-to-consistency bound

External systems are used via an External Sytem Adaptor

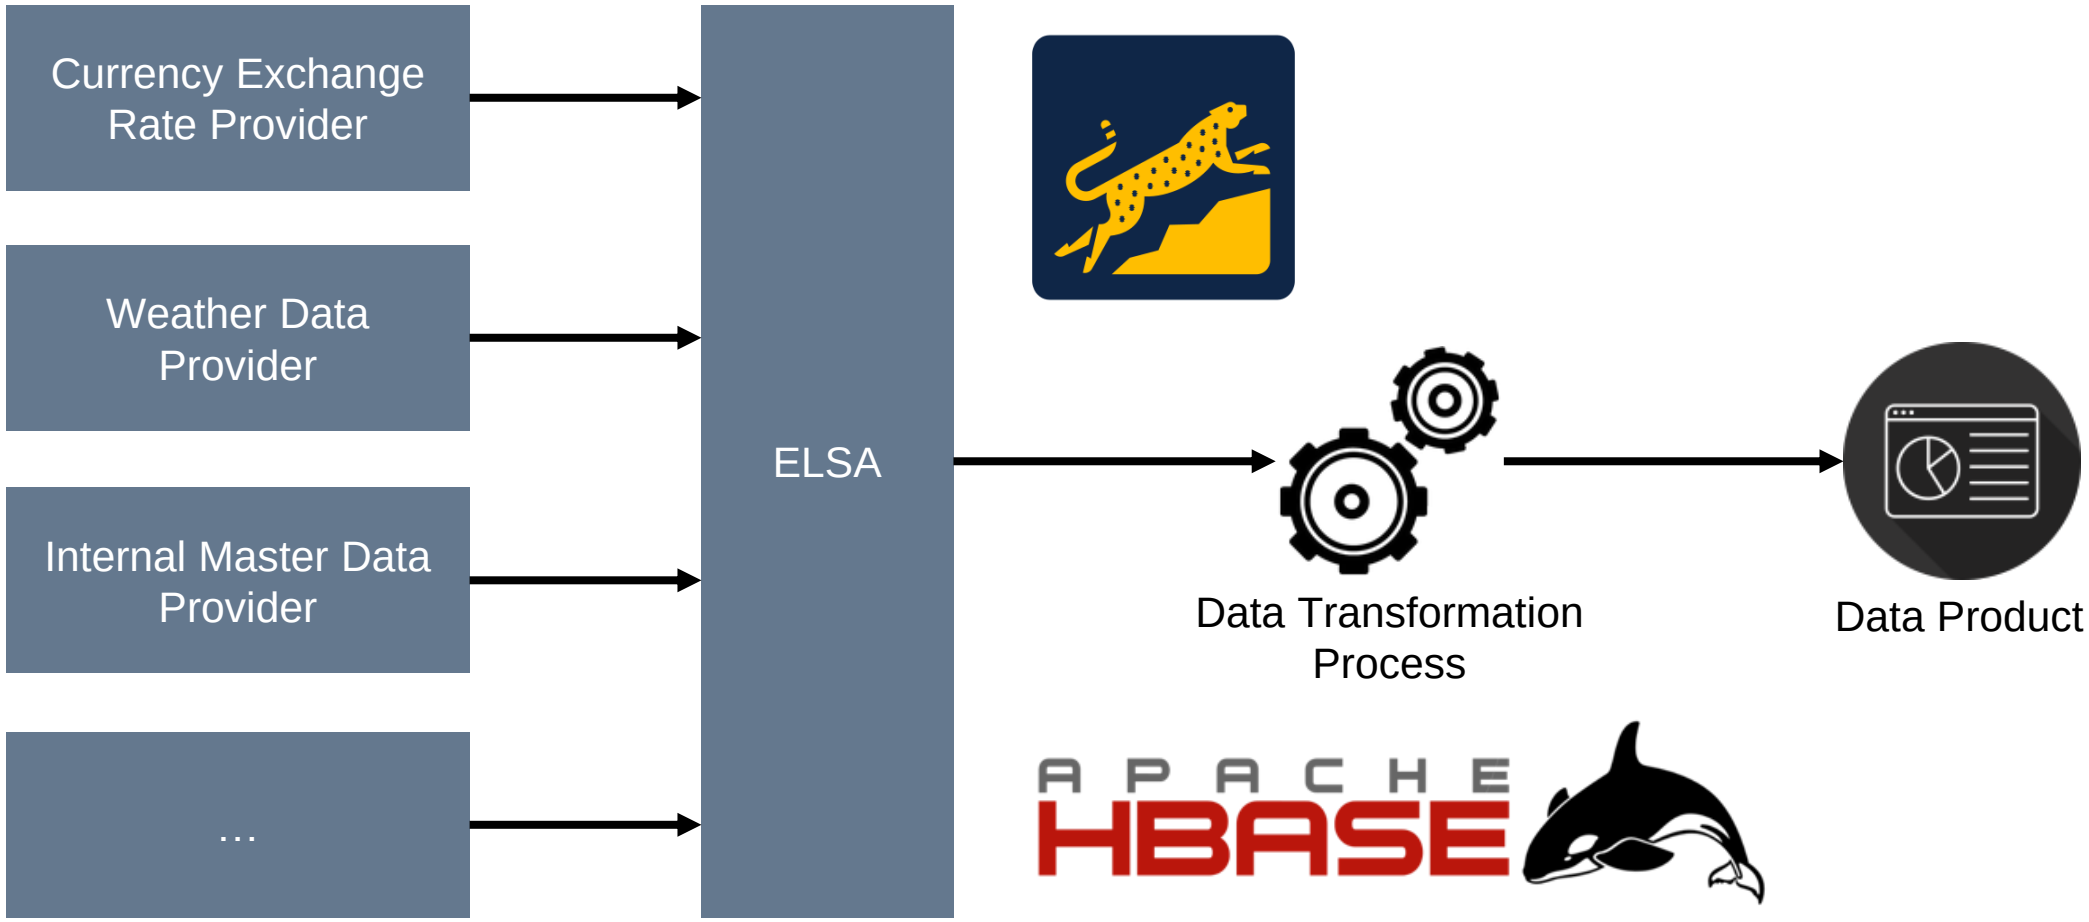

A time-to-consistency bound is required for recomputability

- Time-to-consistency t_{con} is the maximum duration that it may take for a write operation to become and stay visible for all reading processes, starting with the ingest timestamp of the write operation
- Write operations use the current time for the ingest timestamp
- Read operations use at most the current time minus the time-to-consistency as the requested ingested timestamp
- $t_{con} > 0$
- Normally, the time-to-consistency needs to be lower than the transaction timeout for relational databases
- For CP-type distributed databases (HBase, Accumulo), the write timeout can be used, because successful writes are immediately visible to all readers
- If a write operation fails, the retry should use a new timestamp if possible, because then time-to-consistency restarts

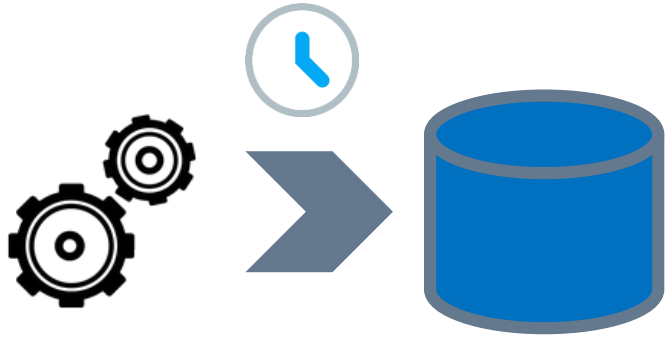

Using the modification timestamps of the external systems can endanger recomputability

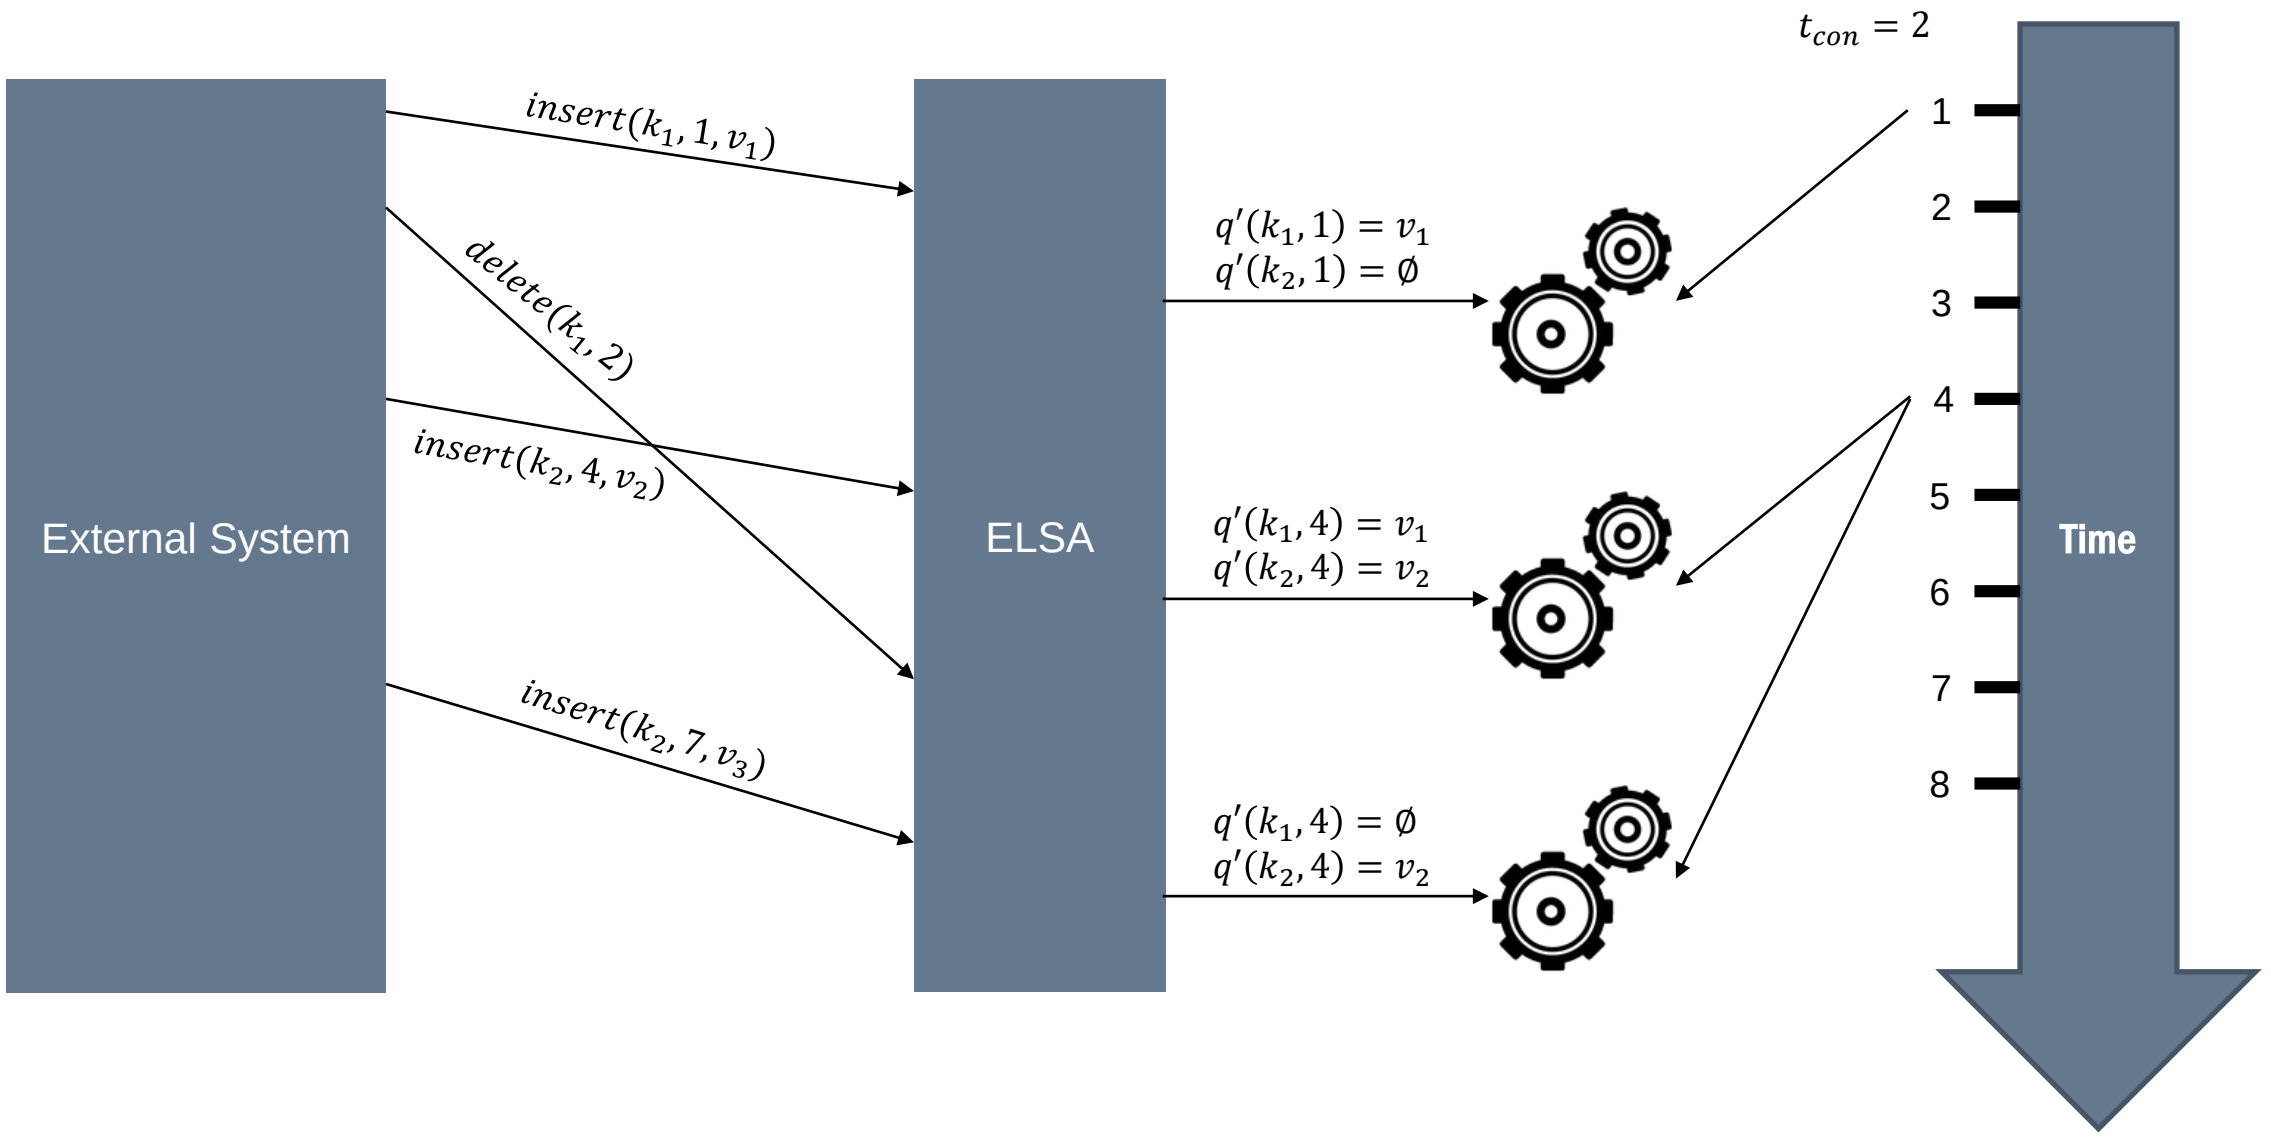

Bitemporal versioning is required for recomputable results

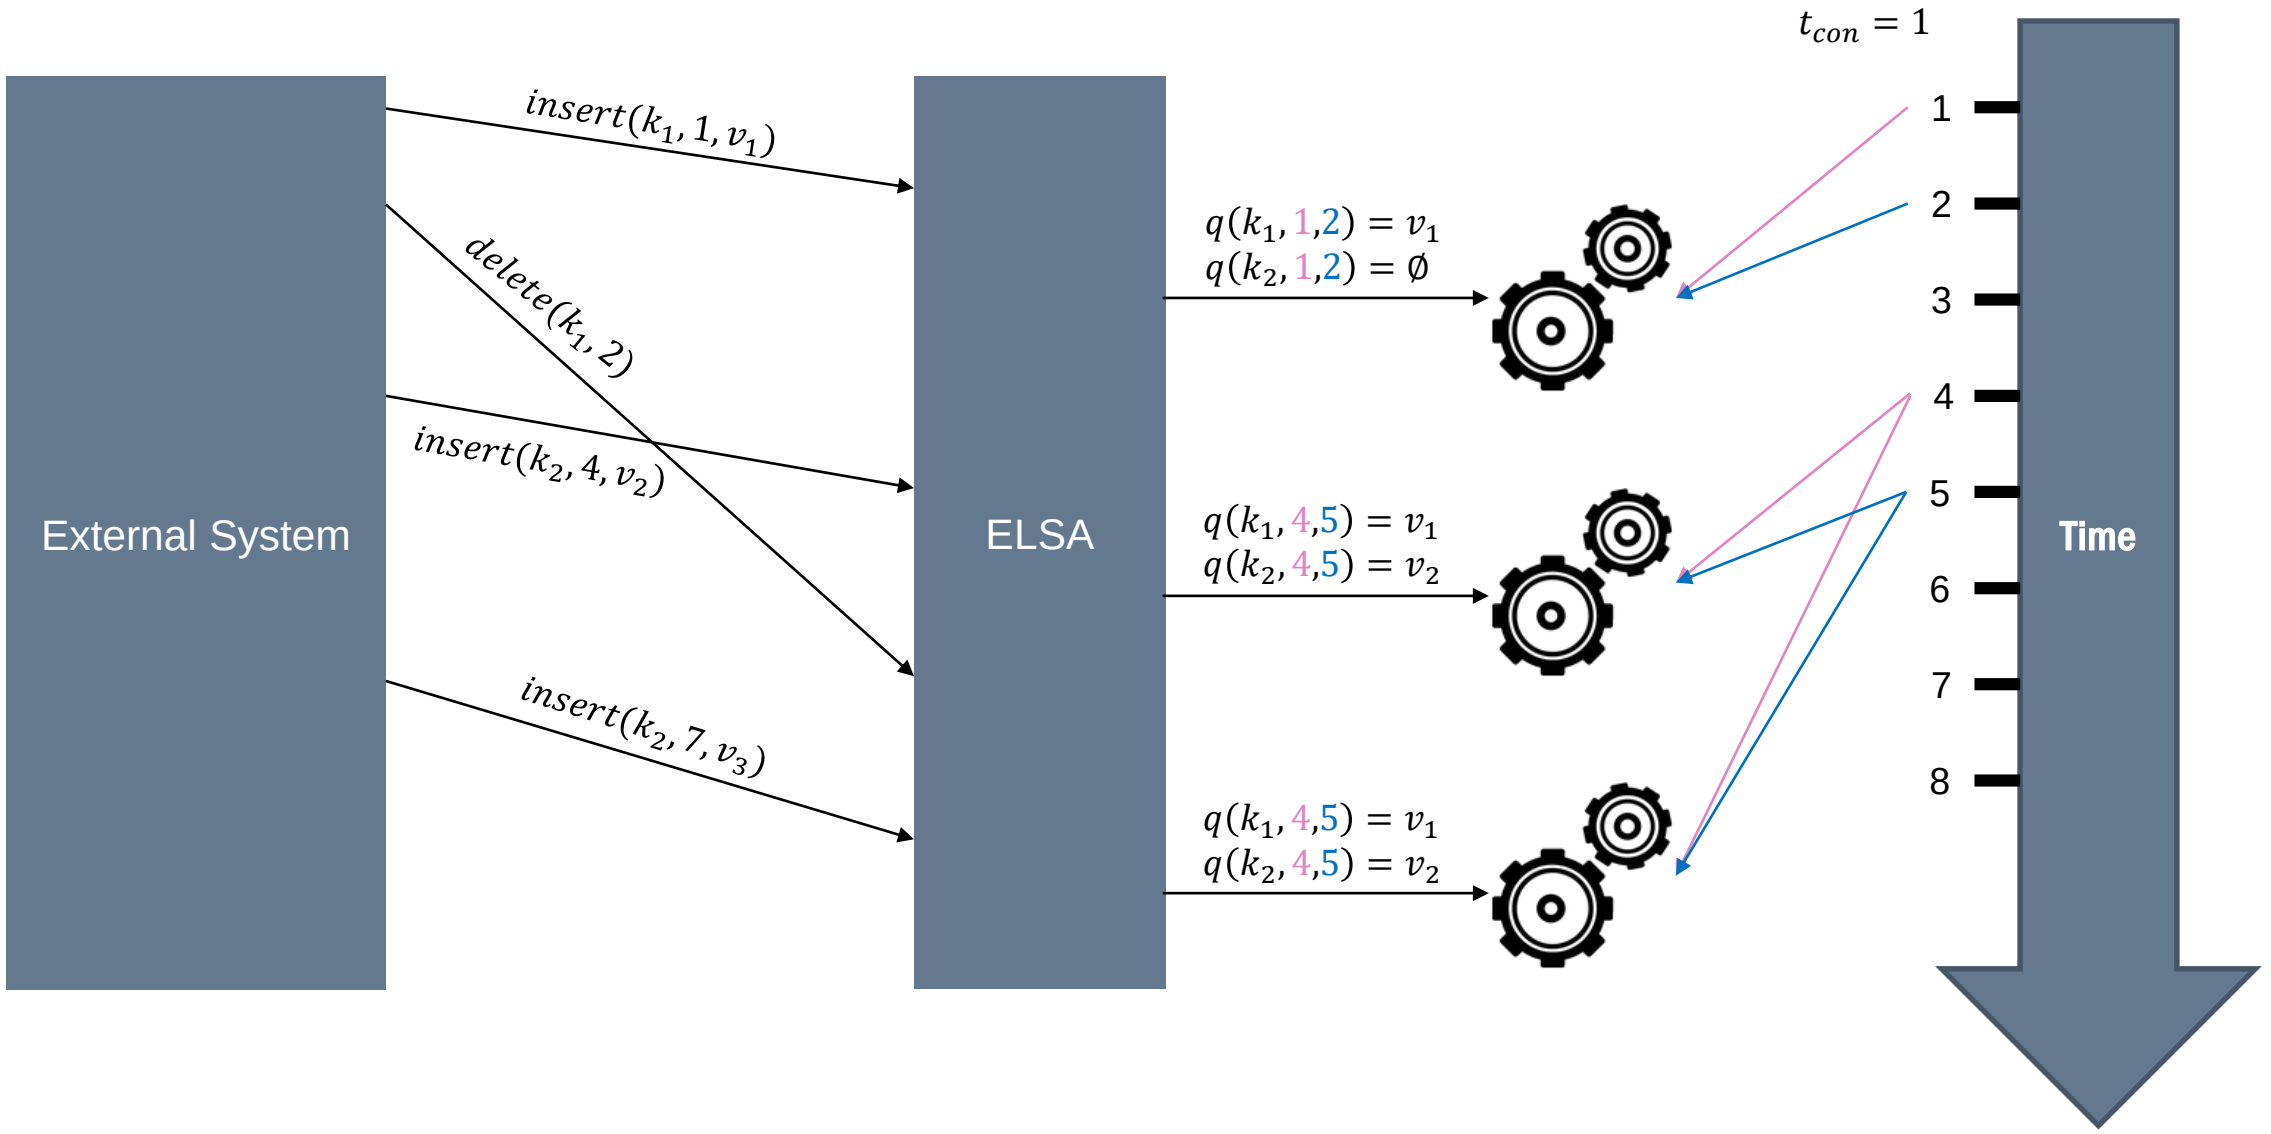

The ELSA Data Synchronization keeps the data up to date

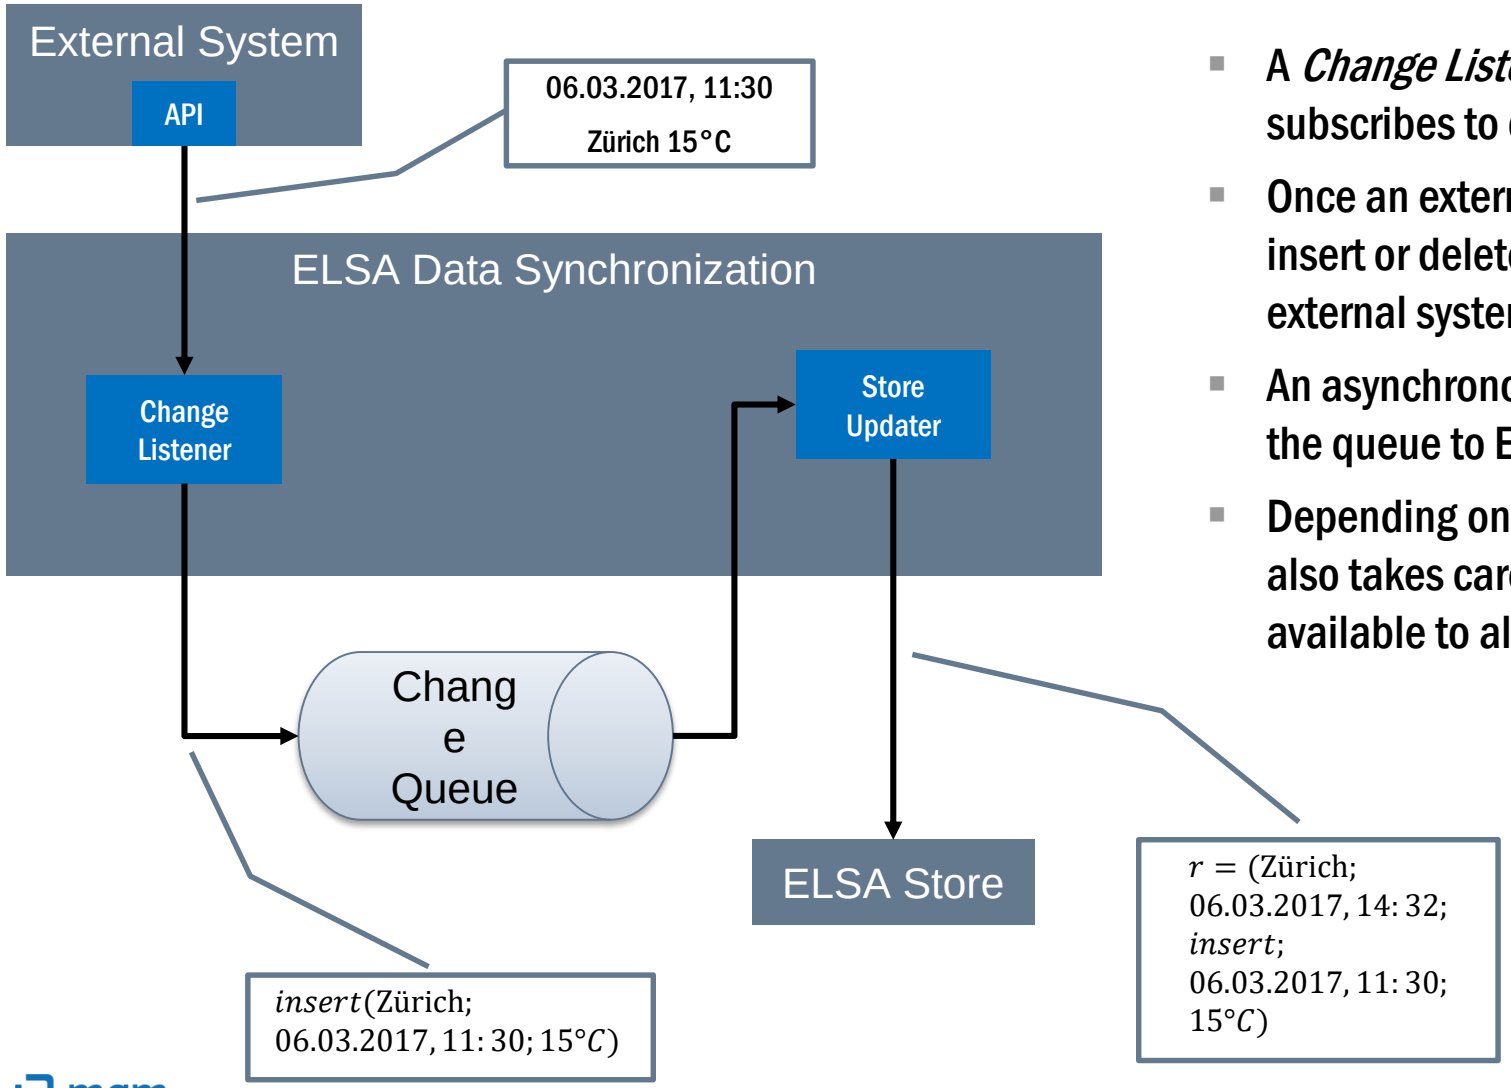

- A *Change Listener* in the ELSA Data Synchronization service subscribes to changes in each external system
- Once an external change arrives, it is transformed to an insert or delete and stored in the change queue for the external system
- An asynchronous *Store Updater* transforms the changes from the queue to ELSA Store records
- Depending on the Store technology used, the Store Updater also takes care that the updated store files become available to all nodes

The ELSA Store provides a queryable history of the external systems‘ state

Record r	Row Key k	Column Family External Store	Column Qualifier t_e	Version t_i	Value Operation & v
r_1	x	ext_1	5	10	insert &v_1
r_2	x	ext_1	10	30	delete
r_3	x	ext_1	12	20	insert &v_2
r_4	x	ext_1	35	40	insert &v_3

$q_1 = (x, 15{,}35)$
$r_1 \rightarrow select$
$r_2 \rightarrow select$
$r_3 \rightarrow select$
$r_4 \rightarrow terminate$
$result = r_3$

$q_2 = (x, 11{,}40)$
$r_1 \rightarrow select$
$r_2 \rightarrow select$
$r_3 \rightarrow terminate$
$result = r_2$

$q_3 = (x, 15{,}15)$
$r_1 \rightarrow select$
$r_2 \rightarrow skip$
$r_3 \rightarrow skip$
$r_4 \rightarrow terminate$
$result = r_1$

Other Factors which influence produced results

Configuration

- Configuration changes may have an impacted in the produced results, e.g. which correction steps are automatically applied
- *Solution:* Annotate the computed results with the configuration values used to produce them
- *Alternative:* Configuration as data – stored in its own versioned store

Version of the software

- *Solution:* Annotate the computed results with the software version used
- *Pitfall:* Old versions may no longer be available to reproduce results! In this case, you could pull up a new cluster with the old version.

Machine learning models

- Might provide different answers to the same questions, e.g. if they have been retrained or reconfigured
- *Solution:* Version them as if they were regular data or configuration

Probabilistic transformations

- Using RNGs
- Hash-based partitioning
- Different amount of partitions
- Rounding errors
- *Solution:* Don‘t do it

Summary

- External systems often don‘t offer what is needed for a distributed data transformation process that shall produce recomputable results
- For system landscapes which need recomputability and scalability, ELSA offers an architecture for integrating external systems
- CP-type columnar databases are good candidates as ELSA store technologies because of their scalability, consistency guarantees and lookup performance
- However, the additional system complexity of the ELSA store and synchronization process may sometimes not be worth the benefits
- Right now, ELSA is limited to key value lookups

Matthias Kricke
kricke@informatik.uni-leipzig.de
Leipzig University

Martin Grimmer
grimmer@informatik.uni-leipzig.de
Leipzig University

Michael Schmeißer
michael.schmeisser@mgm-tp.com
mgm technology partners GmbH

Sources

- https://www.iconfinder.com/icons/134164/cash_currency_exchange_money_icon#size=256
- https://www.iconfinder.com/icons/383986/basket_buy_cart_order_sale_shop_shopping_icon#size=374
- https://www.iconfinder.com/icons/63467/database_storage_icon#size=128
- https://www.iconfinder.com/icons/763237/bubble_comment_communication_conversation_message_other_review_talk_icon#size=128
- https://www.iconfinder.com/icons/18282/browser_earth_global_globe_international_internet_network_planet_world_icon#size=256
- https://www.iconfinder.com/icons/1886958/diagram_hierarchical_hierarchy_order_organization_structure_team_icon#size=256
- https://www.iconfinder.com/icons/667368/celcius_clouds_farenheit_sunshine_temerature_thermometer_weather_icon#size=256

Innovation Implemented.


Munich


Berlin


Dresden


Grenoble


Hamburg


Cologne


Leipzig


Nuremberg


Prague

Michael Schmeißer
mgm technology partners GmbH
Frankfurter Ring 105a
80807 Munich
Tel.: +49 (0) 89 / 35 86 80-0
Fax: +49 (0) 89 / 35 86 80-288
http://www.mgm-tp.com
Michael.Schmeisser@mgm-tp.com