



The pulse of Munich



SWM – Municipal Utility Company

Nils Hendrik Rottgardt
15.12.2021



muenchen.de

Nils Hendrik Rottgardt – BI Architect



Nils Hendrik Rottgardt
BI Architect

Stadtwerke München
Data Analytics & BI



- ▶ I have been working in the SAP BI environment for 15 years. In the beginning, my focus was on BW and integration with third-party SQL databases. The introduction of SAP HANA was an important basis for personal development and ultimately lead into product responsibility of the SAP HANA native subarea of the enterprise data warehouse (UDWH = Unternehmens Data Warehouse) at Stadtwerke München.

Our business segments

Strong legacy: Major business areas



* Electricity, natural gas, district heating, district cooling









Green electricity plants operated by SWM






Including shares; as of 01/2020

Munich and region

-  14 hydropower plants
-  2 wind power plants
-  6 geothermal power plants (one under construction)
-  1 biogas plant, 1 biomass cogeneration plant
-  33 photovoltaic plants
-  12 district cooling plants from ground water/underground streams

Germany

-  3 offshore wind parks (North Sea)
-  Onshore wind parks (Brandenburg, North Rhine-Westphalia, Rhineland-Palatinate and Saxony-Anhalt)
-  2 solar power plants (Bavaria and Saxony)

Europe

-  1 offshore wind park (United Kingdom)
-  Onshore wind parks (Belgium, Finland, France, Croatia, Norway, Poland, Sweden)
-  Parabolic trough power plant (Spain)

Utility company owned by the city of Munich



Customer contracts	1.3 million (energy and drinking water)
---------------------------	---



Shareholder	City of Munich (100 %)
--------------------	------------------------



Offices	Munich, Oslo, London, Copenhagen, Berlin, Brussels
----------------	--



Electricity grid	~13,000 km
-------------------------	------------



District heating grid	~900 km
------------------------------	---------



District cooling grid	~30 km
------------------------------	--------



Natural gas grid	~6,000 km
-------------------------	-----------



Water grid	~3,400 km
-------------------	-----------



Transport network	~700 km
--------------------------	---------

Financials and employees


		2019	2020	
€	Revenues*	in EUR million	10,711.2	10,711.2
	EBITDA	in EUR million	928,0	905,2
	EBIT	in EUR million	458.6	414,4
	Net income	in EUR million	116,1	-152,0
	Investments**	in EUR million	759.8	1,107.2
	Equity ratio	in percent	52	53
👤	Employees		9,444	10,004

* excl. electricity and energy tax | ** in property, plant and equipment and intangible assets

History of SAP BI @SWM

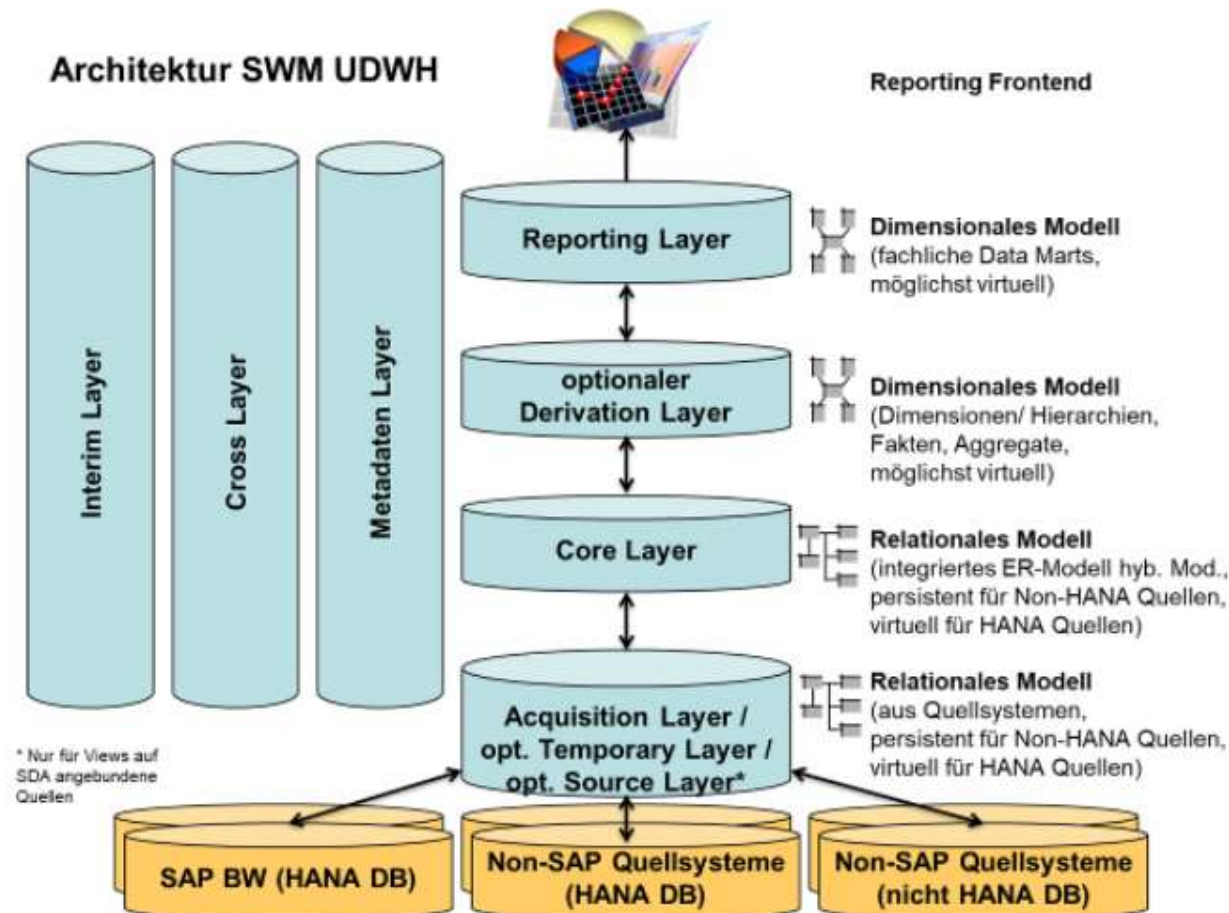
or how we get to agile datawarehousing based on XSA/HANA

- ▶ Heterogenous system architecture
 - ▶ Oracle, SQL Server, PostGre, MariaDB, TM1... → high integration efforts
- ▶ BW since 2003
- ▶ HANA native since 2017
- ▶ XSA since 2019



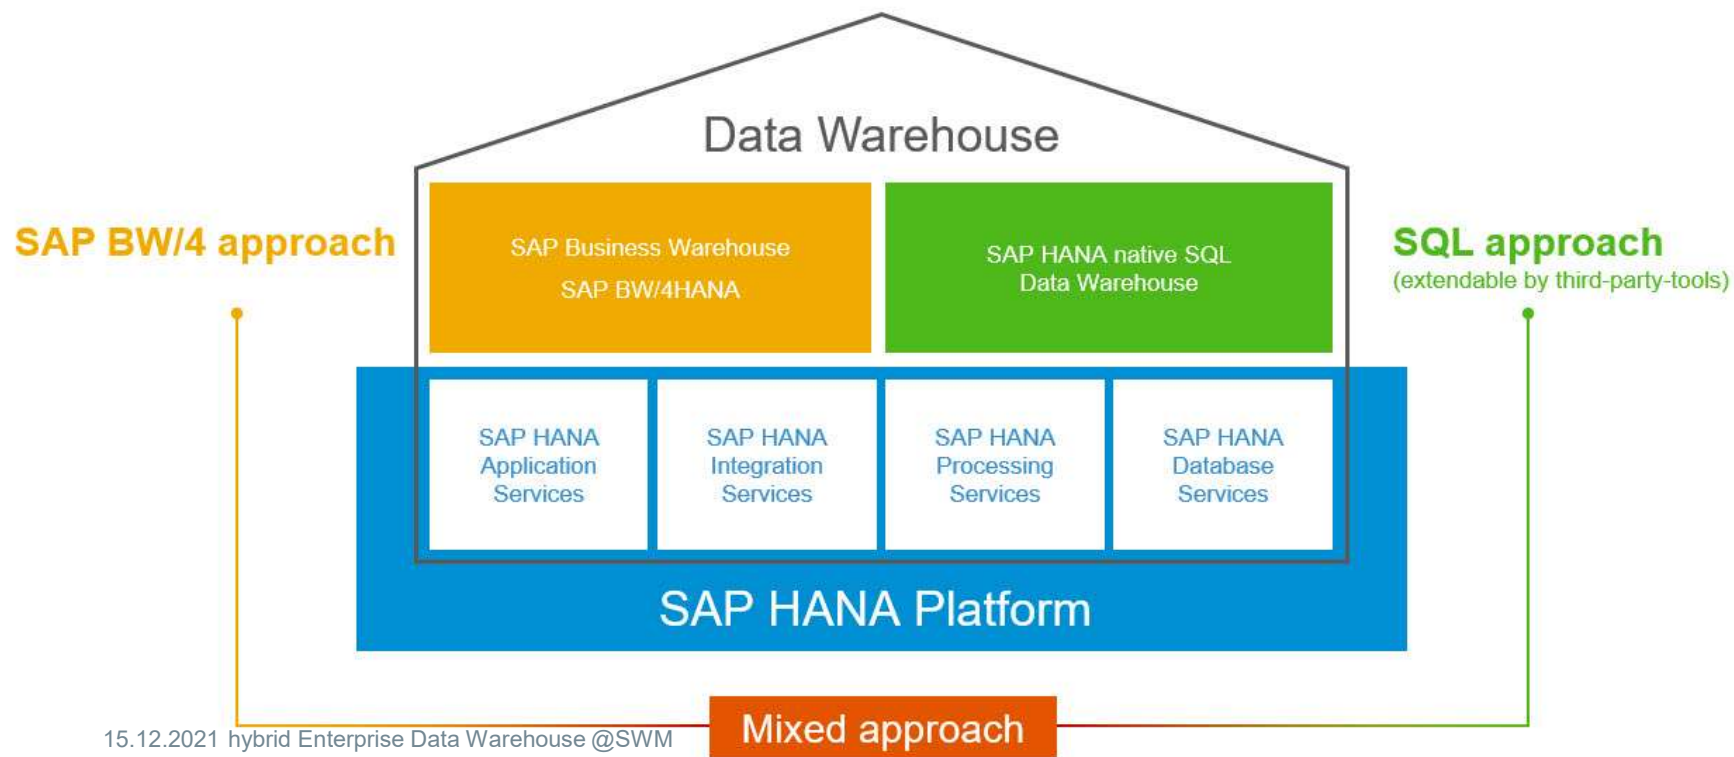
- ▶ Pre-XSA era
 - ▶ Introduction of HANA as central warehouse
 - ▶ manual deployment process, no source versioning, no parallel development in one area
 - ▶ BW und SQL DWH Silos
 - UDWH introduced as single SQL DWH
- ▶ XSA
 - ▶ Introduction of XSA to define a consistent deployment workflow and stabilize development between multiple projects
 - ▶ Containerization of SQL development
 - ▶ Merge of BW und SQL DWH as a single mixed scenario Data Warehouse

Target Architecture as of 2017

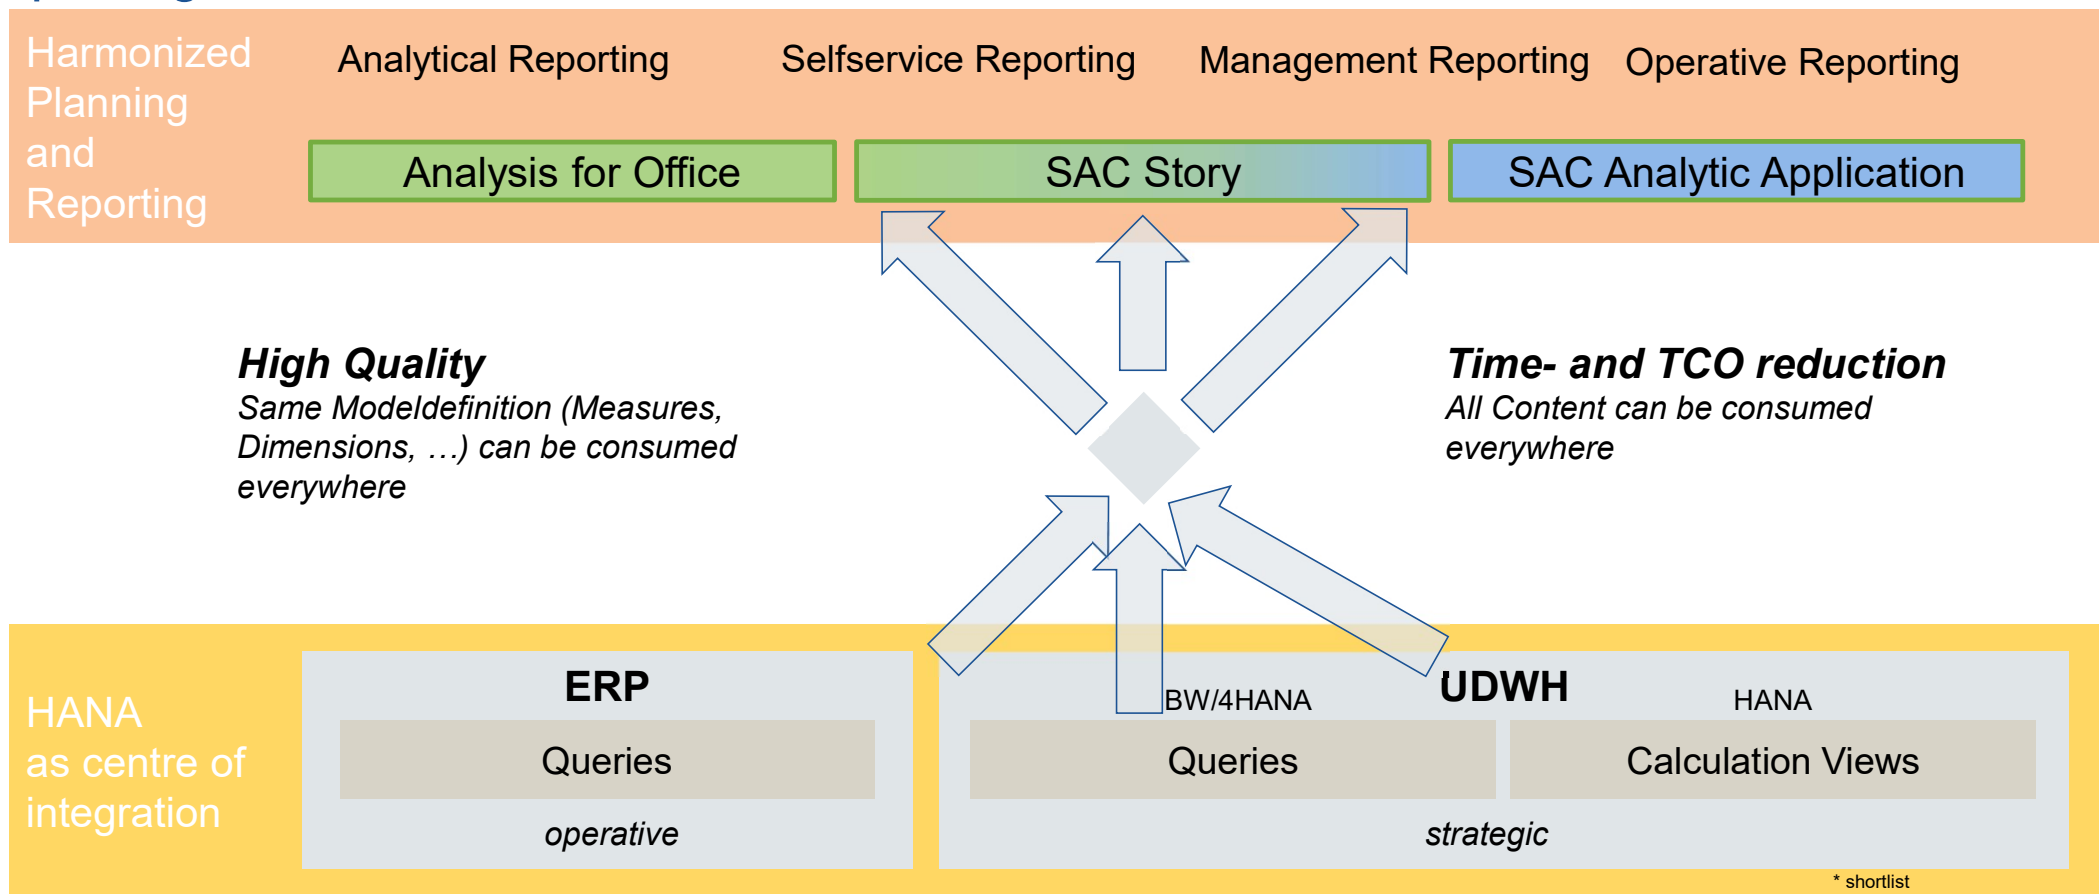


Target Architecture as of 2019

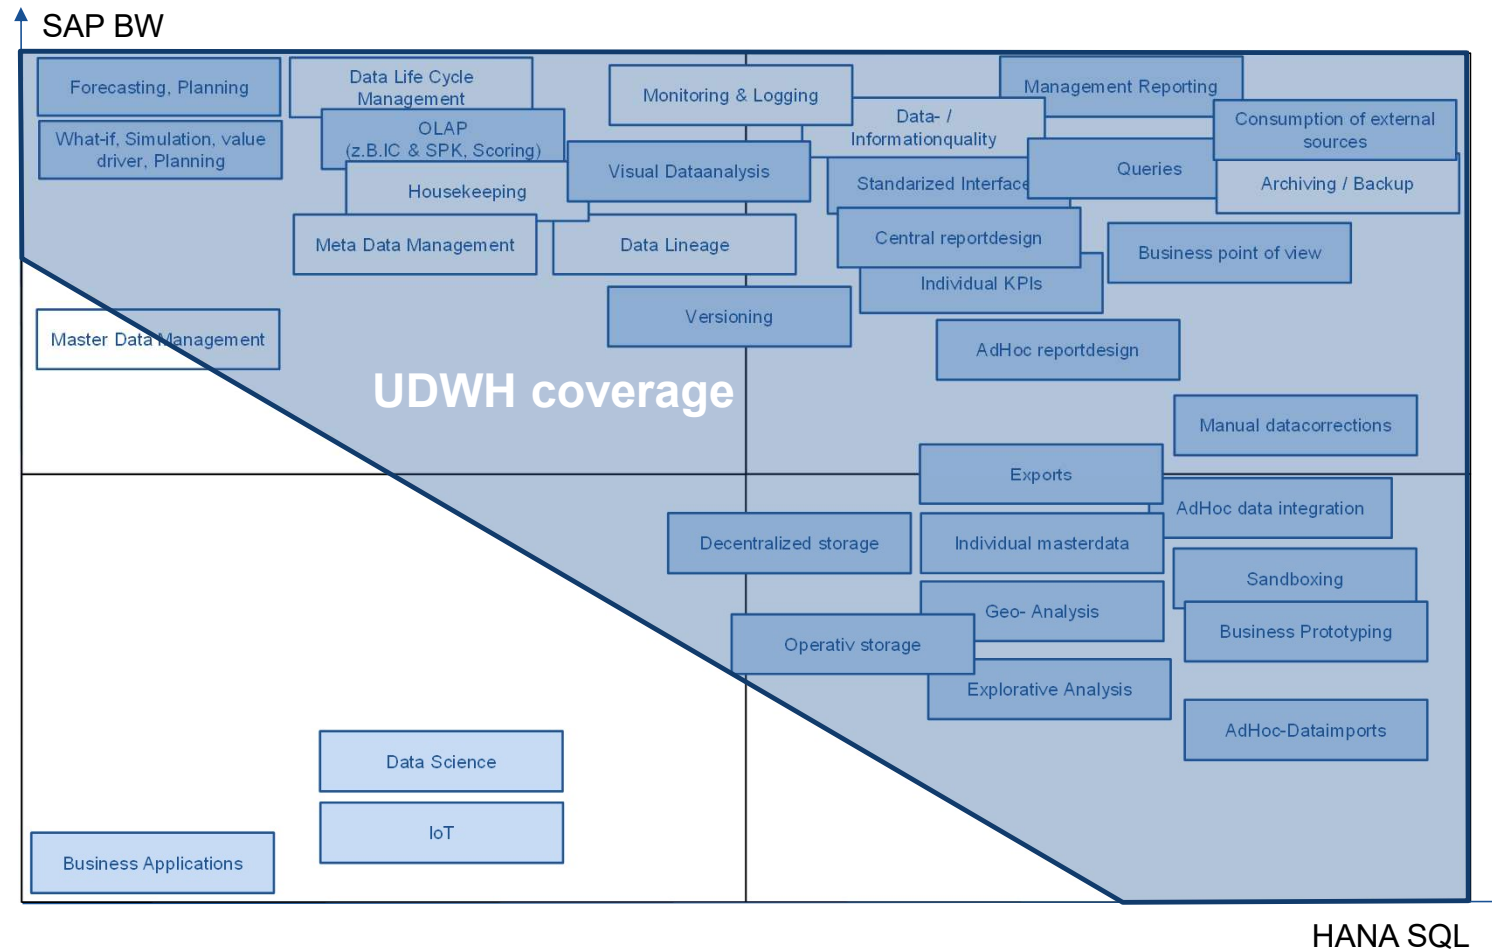
The three approach-strategy for SAP HANA Data Warehousing



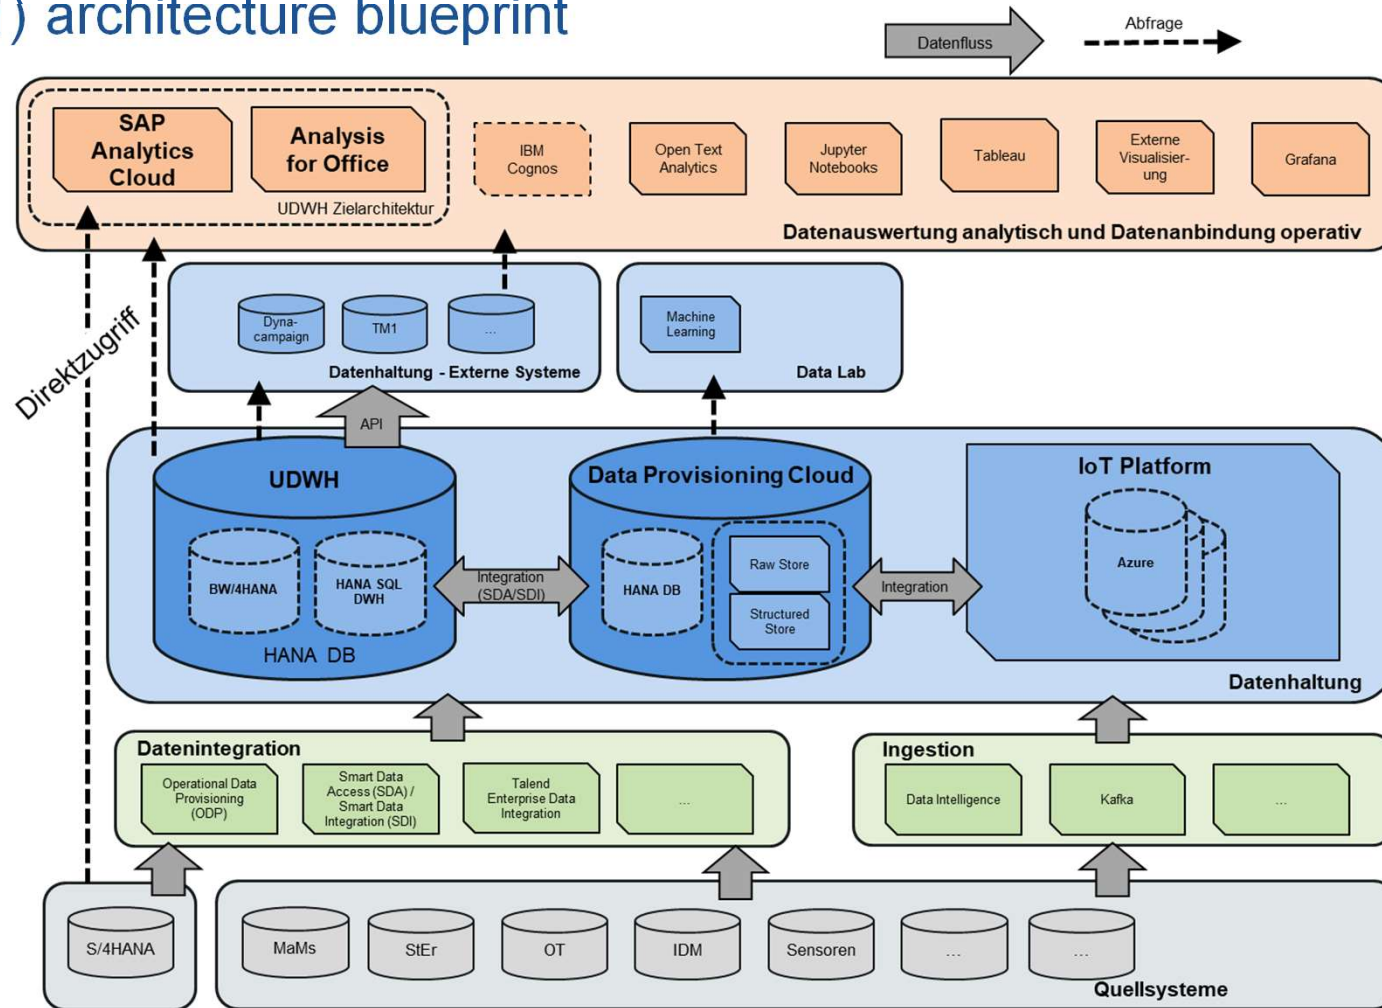
Satisfy Poweruser to Endusers with one Frontend during the whole lifecycle of the reporting



What is our focus of SAC together with our UDWH (BW/4HANA, HANA XSA)

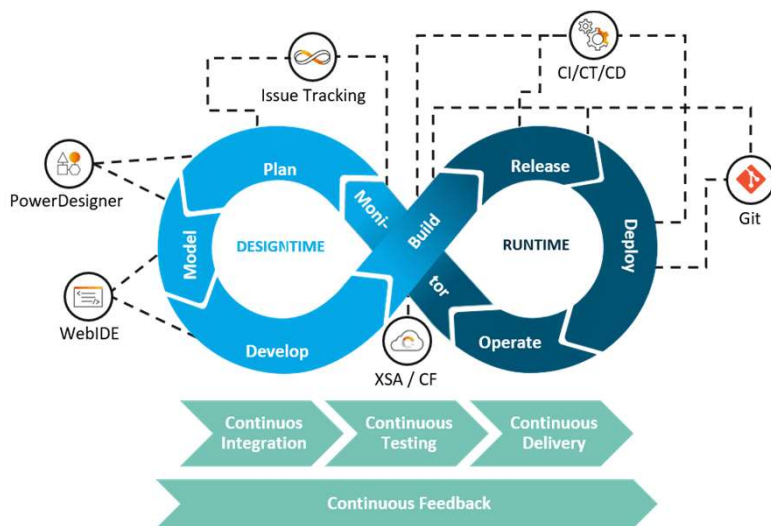


Actual (2021) architecture blueprint



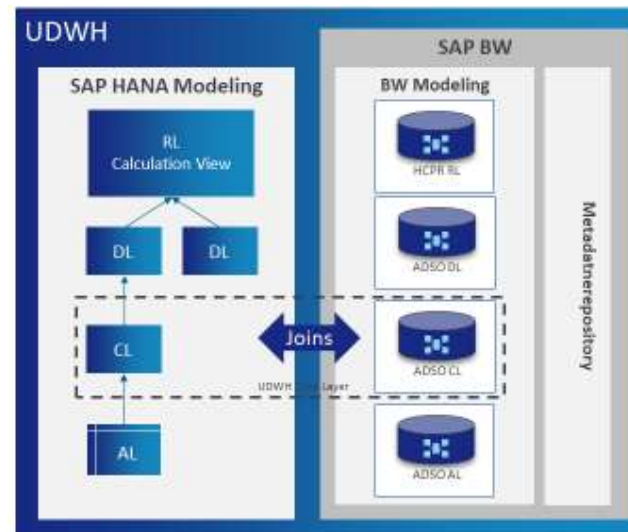
* nach Strategie Data Management SWM: dort heißt dieser Datentopf „Data Lake“ und enthält nur anonymisierte Daten und ist grundsätzlich SWM-weit zugreifbar

Why two pillars of Data Warehousing? The best of two worlds!



XSA

- Flexible
- Short implementation cycles (agile)
- Evolutional approach

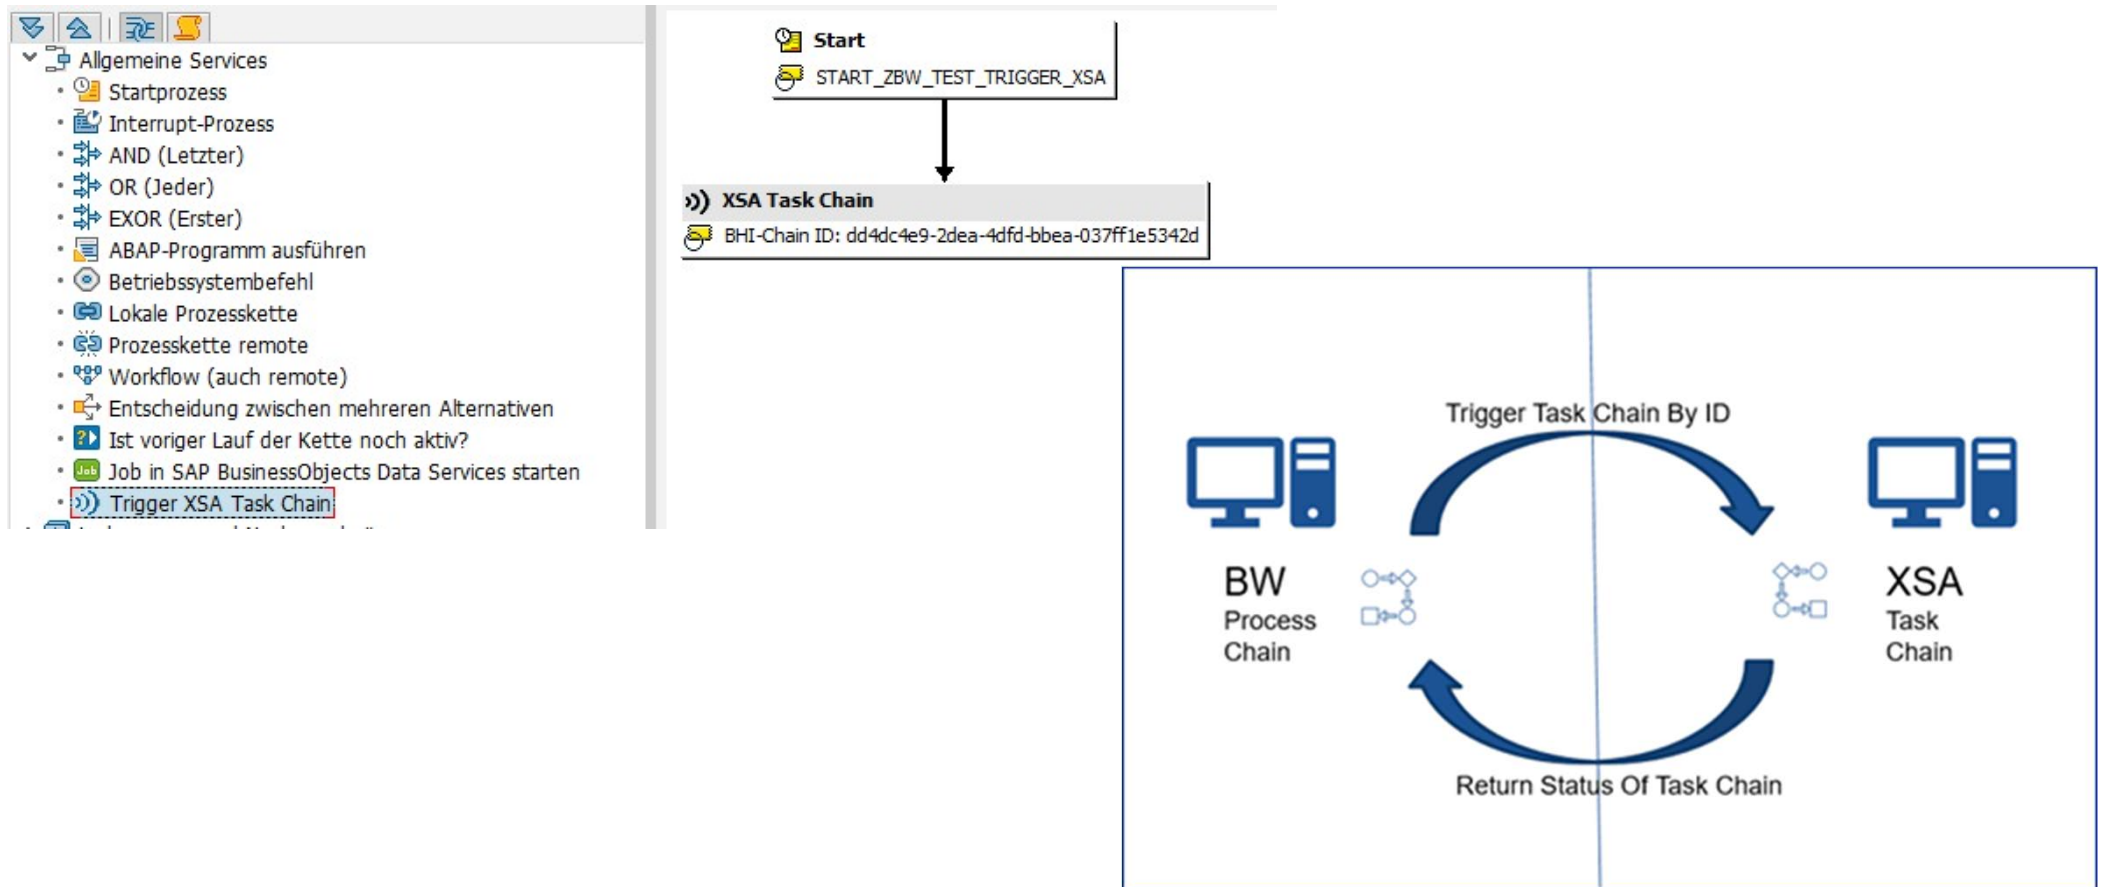


Similar architecture according LSA++ in both worlds and strong integrated data model. Unified Core Layer, so ADSO and also HANA Objects could used in both ETL processes independent where the object is managed (HANA/DWF or BW).

BW

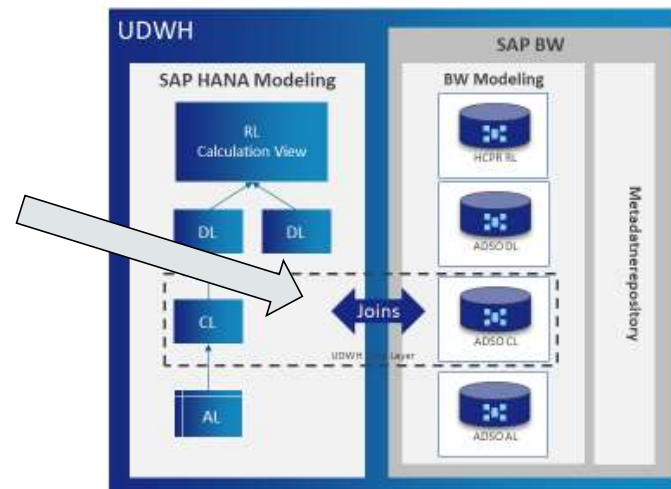
- Data Warehouse Framework
- High automatization
- Stable approach
- Metadata driven

Integration of DWF Task Chains in BW Process Chains for a seamless orchestration



The glue in between... Cross Database Access (CDA)

- ▶ Single Core Layer managed by BW and XSA
- ▶ Bidirectional access between BW and native developments
 - ▶ BW OpenODS View → CDA → SQL Data Warehouse
 - ▶ SQL Data Warehouse → CDA → BW DB Artifacts (ADSO View 7, ADSO View 8, InfoObject ext. HANA Views)
- ▶ No shared access rights → Full Core Layer Access for developers
- ▶ Access Management is done in Reporting Layer
 - ▶ BW Analysis Authorization
 - ▶ HANA Analytical Privileges



Similar architecture according LSA++ in both worlds and strong integrated data model. Unified Core Layer, so ADSO and also HANA Objects could be used in both ETL processes independent where the object is managed (HANA/DWF or BW).

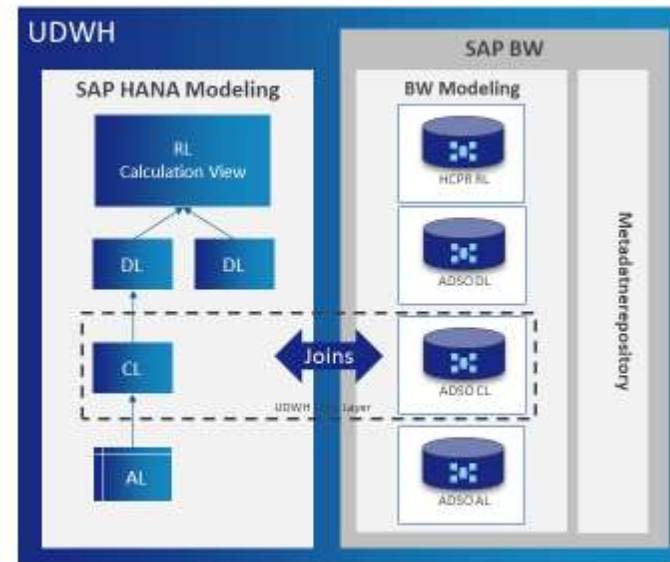
Fast access to BW Objects from XSA/HANA to reduce data redundancy

Advantages

- ▶ Replication of data is not necessary anymore
- ▶ Direct integration of data between BW and HANA
- ▶ No redundancies leads into a small data footprint and a cheaper architecture

Preconditions

- ▶ Introduction of BW/4 (or an ADSO driven BW 7.5 on HANA)
- ▶ Unified Core-Layer
- ▶ Implementation of an SID (Integer) and/or BK based HANA modelling
(Standard-Data Vault with HASH, etc. is not possible)
- ▶ CDA connection between both Tenants



Similar architecture according LSA++ in both worlds and strong integrated data model. Unified Core Layer, so ADSO and also HANA Objects could be used in both ETL processes Independent where the object is managed (HANA/DWF or BW).

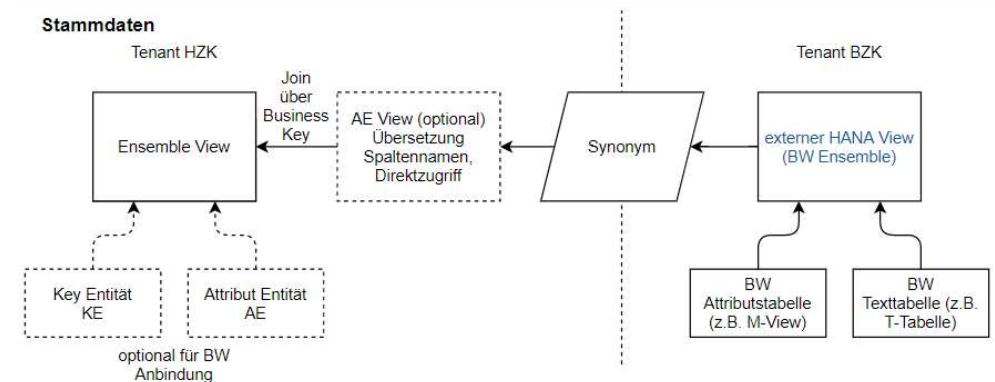
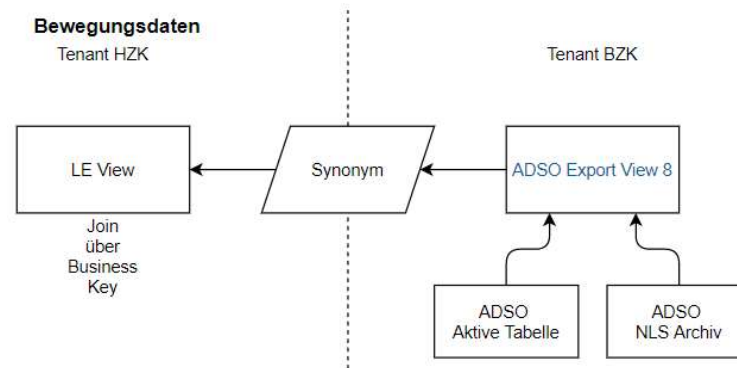
Fast access to BW Objects from XSA/HANA to reduce data redundancy

Advantages

- ▶ Replication of data is not necessary anymore
- ▶ Direct integration of data between BW and HANA
- ▶ No redundancies leads into a small data footprint and a cheaper architecture

Preconditions

- ▶ Introduction of BW/4 (or an ADSO driven BW 7.5 on HANA)
- ▶ Unified Core-Layer
- ▶ Implementation of an SID (Integer) and/or BK based HANA modelling
(Standard-Data Vault with HASH, etc. is not possible)
- ▶ CDA connection between both Tenants



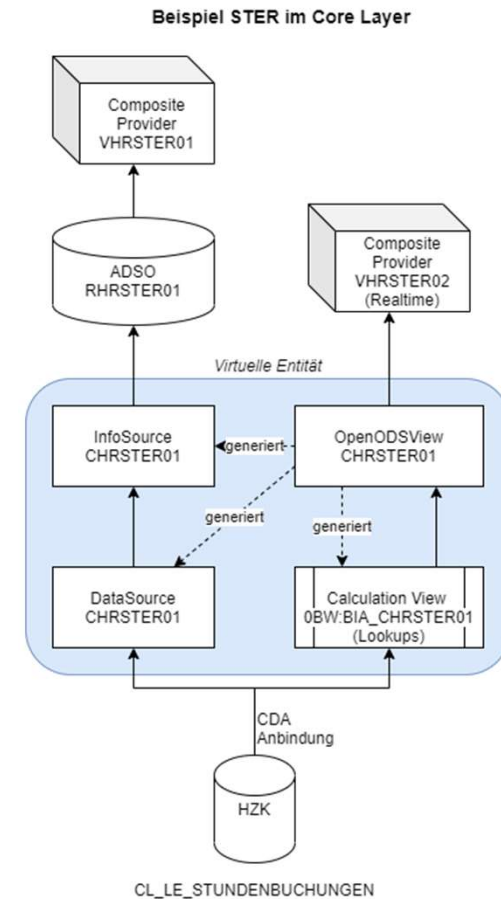
Fast access to XSA/HANA Objects from BW to reduce data redundancy

Advantages

- ▶ Replication of data is not necessary anymore
- ▶ Direct integration of data between BW and HANA
- ▶ No redundancies leads into a small data footprint and a cheaper architecture

Preconditions

- ▶ Introduction of BW/4 (or an ADSO driven BW 7.5 on HANA)
- ▶ Unified Core-Layer
- ▶ Implementation of an SID (Integer) and/or BK based HANA modelling
(Standard-Data Vault with HASH, etc. is not possible)
- ▶ CDA connection between both Tenants



Official Smart Data Access FAQ – Joins are bad

SAP Smart Data Access Modeling Recommendations

- **Avoid JOINS between tables from multiple remote sources**


Joining multiple remote sources would create performance issue. It would require moving data from one remote source to another remote source via HANA to execute the join.

For example, lets say the analytic use requires building HANA model to combine data from HANA, Teradata, and IQ.

In this case joining Teradata tables with IQ tables in a model would create performance overhead because data from either Teradata or IQ will need to be brought over to HANA, then push it to IQ or Teradata where JOIN will execute.

Instead you should create encapsulated calc view for only IQ tables, and create encapsulated calc view for only Teradata tables, and use these encapsulated models to build final analytic model.
- **Avoid large data transfer between SDA remote sources**

In case of joins on high-cardinal its columns, there will be significant data transfer from SDA remote source to HANA. Consider moving these joins to the remote source, by replicating the HANA table to the remote source

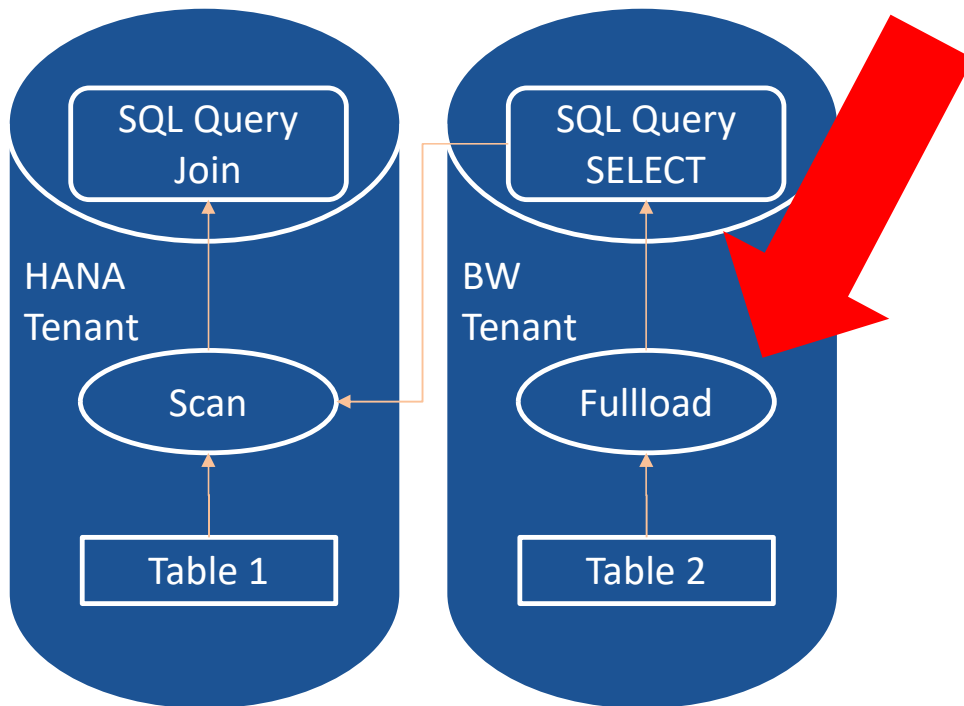


© 2016 SAP SE. All rights reserved.


9

CDA behavior is like everything is stored on the same Tenant

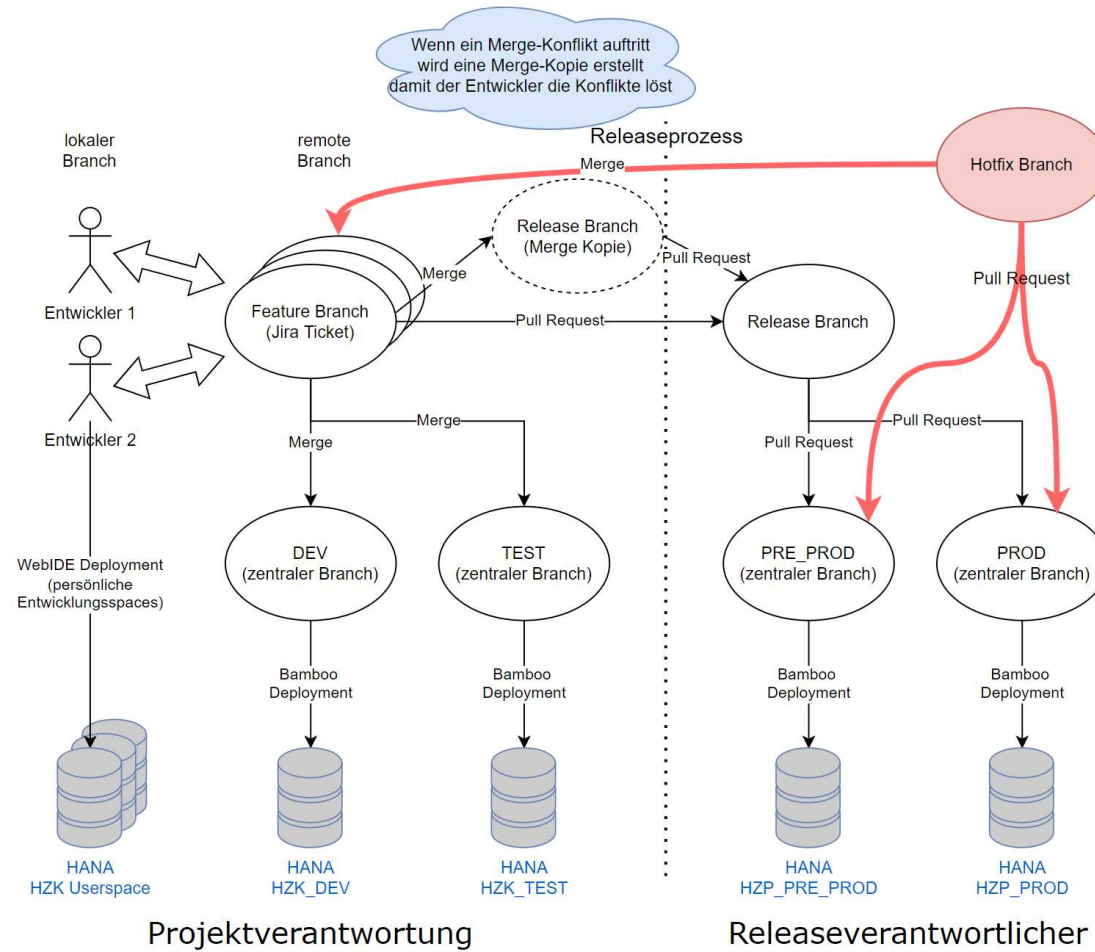
SDA



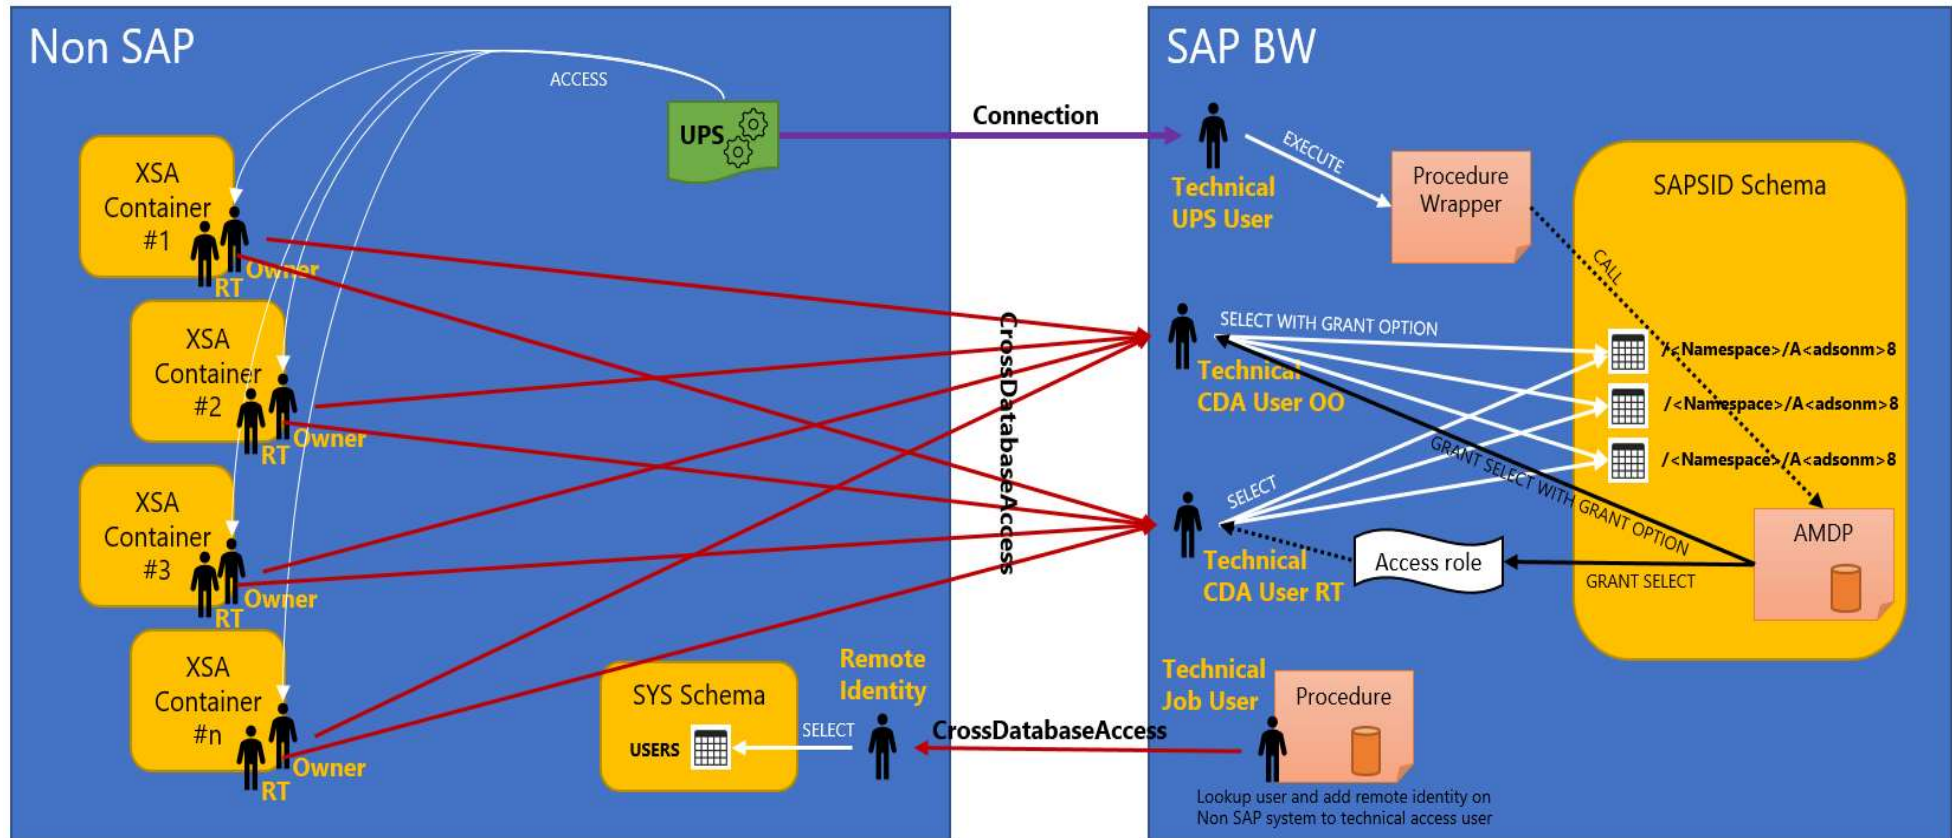
CDA



How we work in XSA (with Git)



CDA automation with HDI Containers – a little complex but running fine



Unified Core Layer – an example

[BZK] OPM_ORDER

Allgemein: Merkmal OPM_ORDER

Allgemein

Technischer Name: OPM_ORDER

Beschreibung: Instandhaltungsauftrag

Kurzbeschreibung: Inst.Auftrag

☒ Externer SAP HANA-View für Stammdaten [system-local.bw.bw2hana.0/PM_ORDER](#)

☒ Externer SAP HANA-View für das Reporting [system-local.bw.bw2hana.0/PM_ORDER_REPORTING](#)

Dictionary

Datentyp: CHAR - Zeichenfolge

Länge: 12

Ausgabelänge: 12

Konvertierungsroutine: ALPHA

☐ Hohe Kardinalität

☐ Groß-/Kleinschreibung beachten

[Klammerungen \(0\)](#)

Eigenschaften

☒ Stammdaten [Stammdatenpflege](#)

☒ Texte

☐ Hierarchien [Hierarchiepflege](#)

☒ Als InfoProvider nutzbar

☐ Berechtigungsrelevant

☐ Nur Attribut

```


1 {
2   "SWM_SAP_BZ_CDA": {
3     "object_owner": {
4       "object_privileges": [{
5         "name": "system-local.bw.bw2hana.0/PM_ORDER",
6         "privileges_with_grant_option": ["SELECT"],
7         "schema": "_SYS_BIC"
8       }]
9     },
10    "application_user": {
11      "object_privileges": [{
12        "name": "system-local.bw.bw2hana.0/PM_ORDER",
13        "privileges": ["SELECT"],
14        "schema": "_SYS_BIC"
15      }]
16    }
17  }
18 }
    
```

```

*extermer_zugriff.hdbgrants x *core_NC.hdbsynonymconfig x *CL_BZ_AUFTRAG_ENSEMBL... x
1 VIEW CL_BZ_AUFTRAG_ENSEMBLE_V AS
2 SELECT
3   A."OPM_ORDER" as BK_CL_KE_AUFTRAG_AUFTRAG,
4   A."0WORKCENTER" as BK_CL_KE_ARBEITSPLATZ_ARBEITSPLATZ,
5   A."0PMACTTYPE" as BK_CL_KE_INSTANDHALTUNGSLEISTUNGSART_INSTANDHALTUNGSLEISTUNGSART,
6   A."0BUS_AREA" as BK_CL_KE_GESCHAFTSBEREICH_GESCHAFTSBEREICH,
7   A."0MAINTLOC" as BK_CL_KE_STANDORT_KONTIERUNG_STANDORT_KONTIERUNG
8   ....
9 FROM "BZ_system-local.bw.bw2hana.0/PM_ORDER_SY" as A
10
11 ;
    
```

```

*extermer_zugriff.hdbgrants x *core_NC.hdbsynonymconfig x
1 {
2   {
3     "BZ_system-local.bw.bw2hana.0/PM_ORDER_SY": {
4       "target": {
5         "object": "system-local.bw.bw2hana.0/PM_ORDER",
6         "schema": "_SYS_BIC",
7         "database.config": "css-service/database"
8       }
9     }
10  }
11 }
    
```



Lessons learned

- ▶ Technical preconditions must be fulfilled like Cross Database Access. SDA as great restrictions for a Mixed Scenario
- ▶ XSA and BW developers must “trust” each other → Unified Core-Layer (at best one developer for both technologies)
- ▶ Architects has to understand BW and XSA/HANA to decide objective where to implement a data flow (HANA, BW or hybrid approach)
- ▶ XSA is oriented on software development processes which is not easy to understand for colleagues from the business
- ▶ The native modelling strategy has to be aligned with the requirements of the overall model. So e.g. a straight data vault approach is not possible as it is not compatible with BW modelling
- ▶ Most important: reserve time for SAP incidents and actual missing functions like automatic revalidation of artefacts

What you get with Mixed Scenario Data Warehouse by SAP

A highly flexible Data Warehouse covers most of the customer needs like standard reporting, SelfService-Reporting (OLAP) up to data wrangling (SQL) with less implementation effort. Supporting classic and agile projects best.