Crossload Consulting GmbH

SAP tech. News
- vSAP Stammtisch

30.04.2020

Jens Gleichmann
Geschäftsführer
jens.gleichmann@x-l-c.de
www.linkedin.com/in/jens-gleichmann
www.xing.com/profile/Jens_Gleichmann

Agenda

1. Software Releases
    1. HANA Cockpit
    2. Supported OS for HANA
    3. HANA Maintenance
    4. HANA Rev.47
2. HANA Express Edition
3. HANA Cloud

Releases

- Overview -

Releases

- HANA Cockpit -

- SAP HANA Cockpit SP11 P15
  - fixes Ghostcat AJP tomcat vulnerability (note 2897141 / 2908046)
  - Fixed in XSA version 1.0.124. / tomcat 8.5.51.0
  - affected are also all other XSA installations!
  - Bad news: no delta of the functionalities available
- SAP HANA Cockpit SP11 P16
  - available since 2020/04/22

Releases

- HANA Cockpit -

```
              :/hana/shared/HSC/xs/uaaserver/tomcat> export JRE_HOME=/hana/shared/HSC/xs/sapjvm_8/jre
              :/hana/shared/HSC/xs/uaaserver/tomcat> cd bin/ ; ./version.sh
Using CATALINA_BASE:   /hana/shared/HSC/xs/uaaserver/tomcat
Using CATALINA_HOME:   /hana/shared/HSC/xs/uaaserver/tomcat
Using CATALINA_TMPDIR: /hana/shared/HSC/xs/uaaserver/tomcat/temp
Using JRE_HOME:        /hana/shared/HSC/xs/sapjvm_8/jre
Using CLASSPATH:       /hana/shared/HSC/xs/uaaserver/tomcat/bin/bootstrap.jar:/hana/shared/HSC/xs/uaaserver/tomcat/bin/tomcat-juli.jar
Server version: Apache Tomcat/8.5.42
Server built:   Jun 4 2019 20:29:04 UTC
Server number:  8.5.42.0
OS Name:        Linux
OS Version:     3.10.0-693.61.1.el7.x86_64
Architecture:   amd64
JVM Version:    8.1.060
JVM Vendor:     SAP AG
              :/hana/shared/HSC/xs/uaaserver/tomcat/bin> xs lc

Getting software components in org "HUK" / space "COCKPIT" as XSA_ADMIN...
Found software components:

software component              version
-------------------------------------
XSAC_COCKPIT (sap.com)          2.11.14
XSAC_HRTT (sap.com)             2.11.20062
XSAC_PORTAL_SERV (sap.com)      1.3.2
XSAC_UI5_FESV5 (sap.com)        1.60.18
XSAC_XSA_COCKPIT (sap.com)      1.1.15
```

```
<!-- Define an AJP 1.3 Connector on port 8009 -->
<Connector port="8009" protocol="AJP/1.3" redirectPort="8443" />
```

Releases

- HANA Cockpit -

```
              :/hana/shared/HSC/xs/uaaserver/tomcat/bin> ./version.sh
Using CATALINA_BASE:   /hana/shared/HSC/xs/uaaserver/tomcat
Using CATALINA_HOME:   /hana/shared/HSC/xs/uaaserver/tomcat
Using CATALINA_TMPDIR: /hana/shared/HSC/xs/uaaserver/tomcat/temp
Using JRE_HOME:        /hana/shared/HSC/xs/sapjvm_8/jre/
Using CLASSPATH:       /hana/shared/HSC/xs/uaaserver/tomcat/bin/bootstrap.jar:/hana/shared/HSC/xs/uaaserver/tomcat/bin/tomcat-juli.jar
Server version: Apache Tomcat/8.5.51
Server built:   Feb 5 2020 22:26:25 UTC
Server number:  8.5.51.0
OS Name:        Linux
OS Version:     3.10.0-693.61.1.el7.x86_64
Architecture:   amd64
JVM Version:    8.1.060
JVM Vendor:     SAP AG
```

```
<!-- Define an AJP 1.3 Connector on port 8009 -->
<!--
<Connector protocol="AJP/1.3"
           address="::1"
           port="8009"
           redirectPort="8443" />
```

Releases

- Supported OS for HANA -

- News for supported OS for SAP HANA
  - RHEL

as of 2020-04-15	OS level	RHEL 7.3	RHEL 7.4	RHEL 7.5	RHEL 7.6	RHEL 7.7	RHEL 8.0	RHEL 8.1
End of Maintenance	HANA SPS	30.11.2020	31.8.2021	~30.04.2022	31.10.2022	31.8.2021	31.12.2020	30.11.2023
Kernel		3.10.0-514.36.5.el7+	3.10.0-693.11.6+	all	3.10.0-957.1.3+	all	4.18.0-80.15.1+	4.18.0-147.5.1+
HANA 1.0	SPS12	120+	122.14+	122.19+	122.23+	-	-	-
HANA 2.0 x86/PPC	SPS02	21+	23+	-	-	-	-	-
	SPS03	30+	30+	32+	36+	-	-	-
	SPS04	-	40+	-	40+	-	40+/45+	45+

Releases

- Supported OS for HANA -

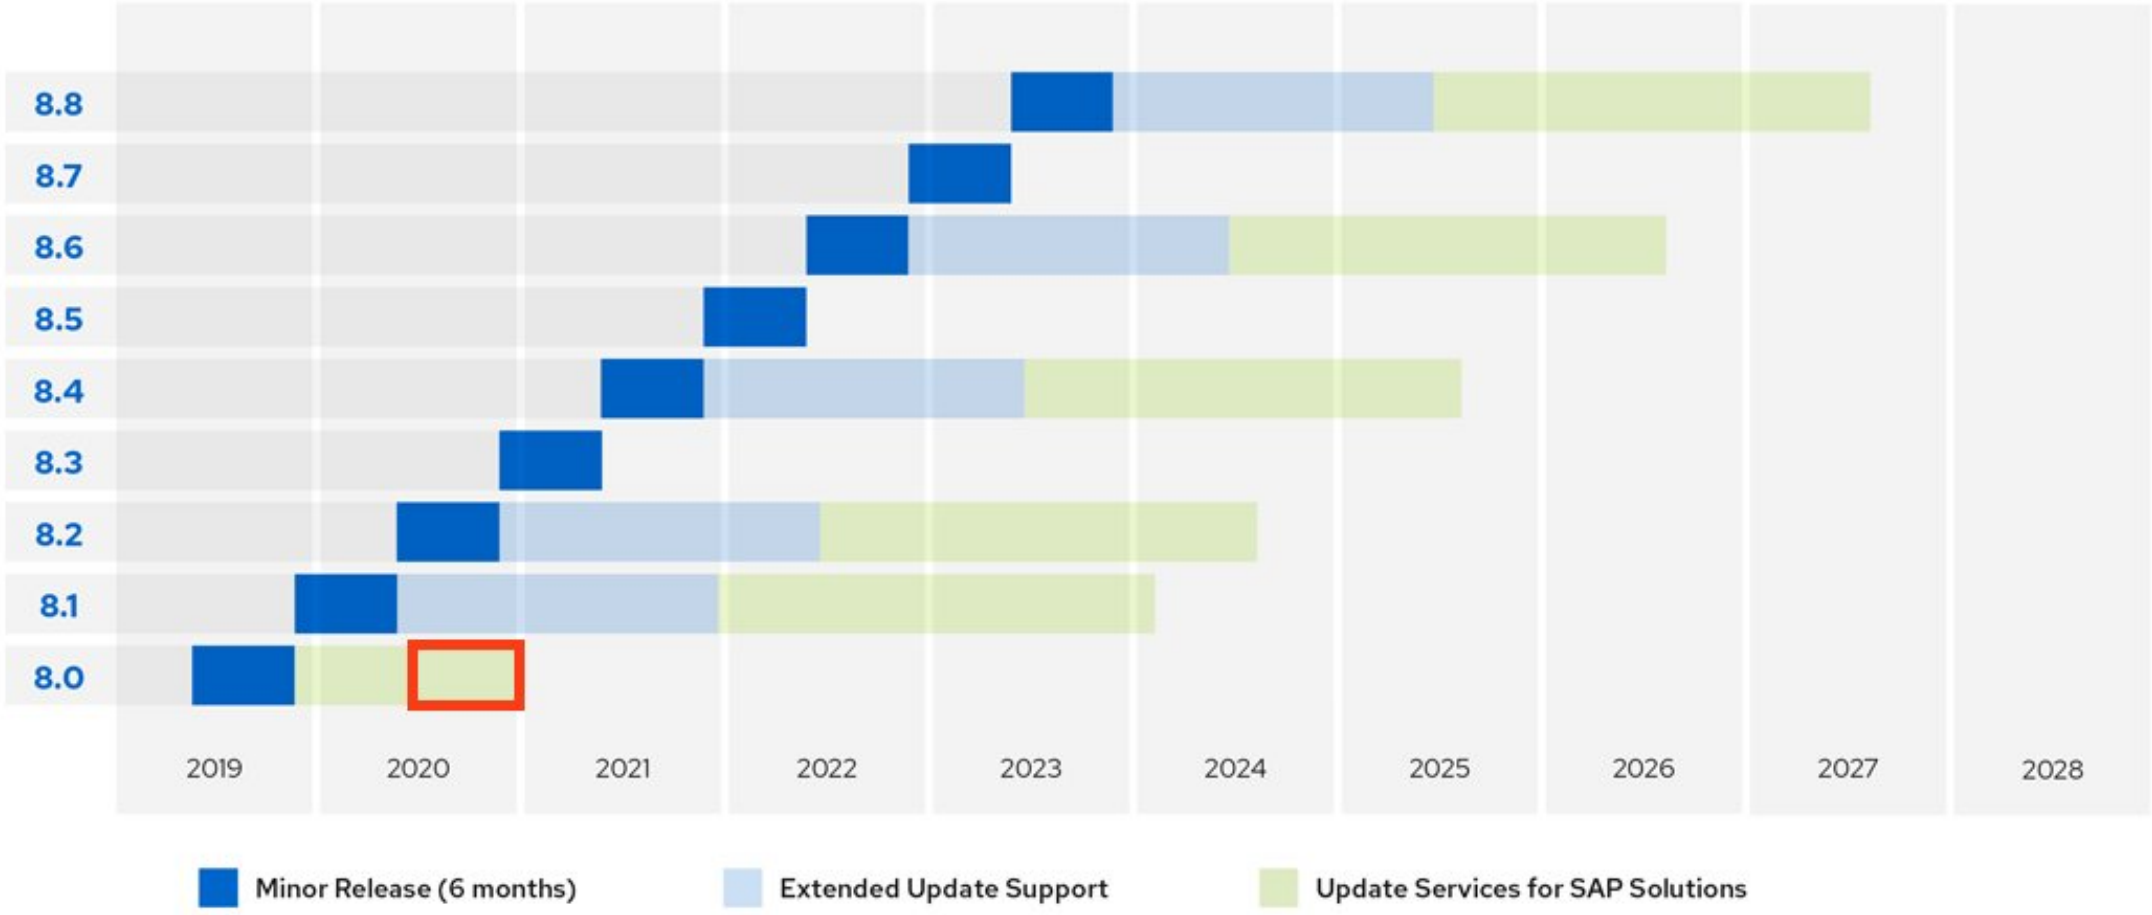

Releases

- Supported OS for HANA -

- News for supported OS for SAP HANA
  - SLES

as of 2020-04-15	OS level	SLES12 SP1	SLES12 SP2	SLES12 SP3	SLES12 SP4	SLES12 SP5	SLES15	SLES15 SP1
End of Maintenance	HANA SPS	31.5.2020	31.5.2021	30.6.2022	30.6.2023	31.10.2024	31.12.2022	?
Kernel		3.12.74-60.64.40.1+	4.4.120-92.70.1+	4.4.120-94.17.1+	4.12.14-95.13.1+	4.12.14-122.17+	4.12.14-150.14.1+	all
HANA 1.0	SPS12	120+	120+	122.15+	122.22+	122.29+	122.21+	122.27+
HANA 2.0 x86/PPC	SPS02	20+	20+	23+	-	-	-	-
	SPS03	30+	30+	30+	35+	-	34+	-
	SPS04	-	-	40+	40+	45+	40+	44+

Releases

- HANA Maintenance -

SAP HANA Maintenance Strategy

Revision Strategy for SAP HANA 2.0

- New capabilities are introduced once a year, every time a new SAP HANA **Support Package Stack** (SPS) is released.
- SAP is providing bug fixes and security patches for every SPS for 2 years after RTC*
- SAP will provide Maintenance Revisions for SAP HANA 2.0 SPS05 for a period of 5 years after RTC
- Maintenance Revisions for SAP HANA SPS05 are not scheduled, they are delivered on demand
- We recommend that maintenance timelines and project go live dates are adjusted to this release schedule

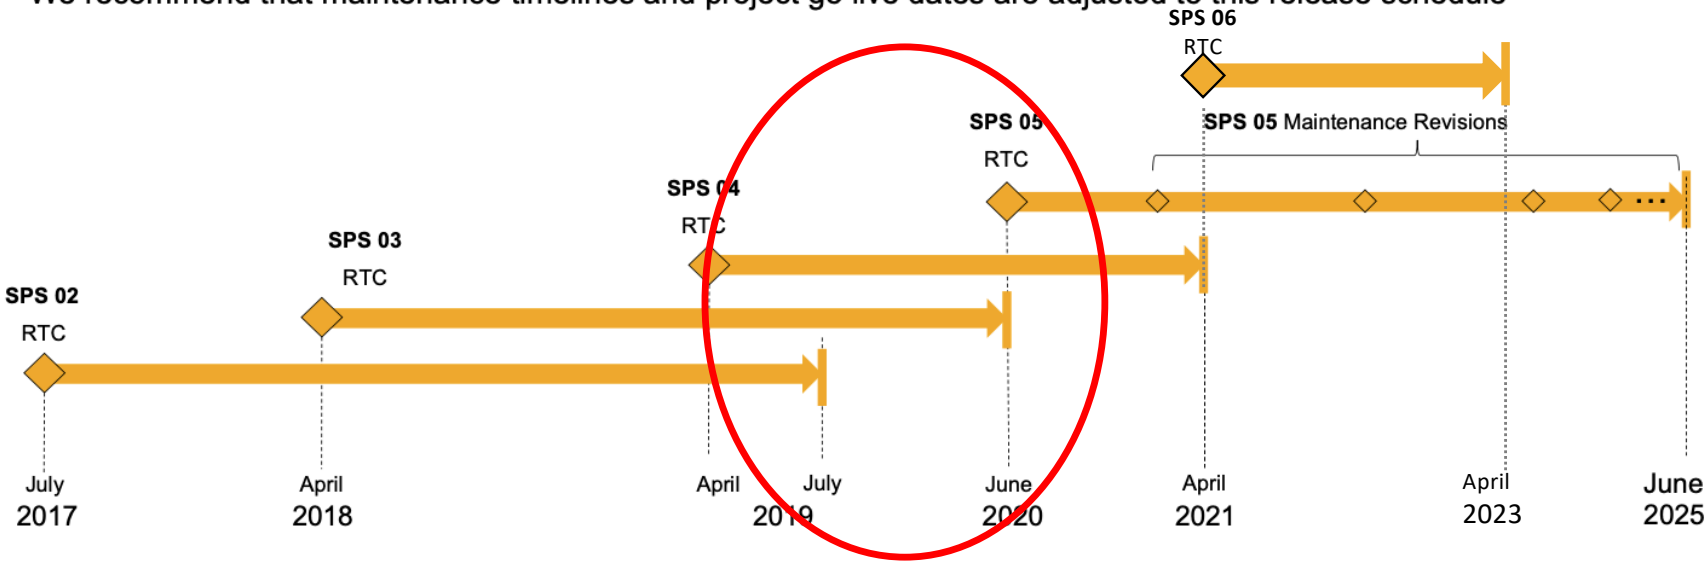

Releases

- HANA Rev. 47 -

- SAP HANA Rev. 47
  - 2893339 - SAP HANA 2 SPS04 Revision 047.00
  - Fixes a bunch of issues regarding memory leaks, HEX, ESX

- My personal recommendation:
  - Wait 4-6 weeks before testing rev. 47 it
  - Wait up to rev. 53/54 before using SPS05

Releases

- SUM DMO -

- SUM SP07
- Kernel 7.77
- ENQ2

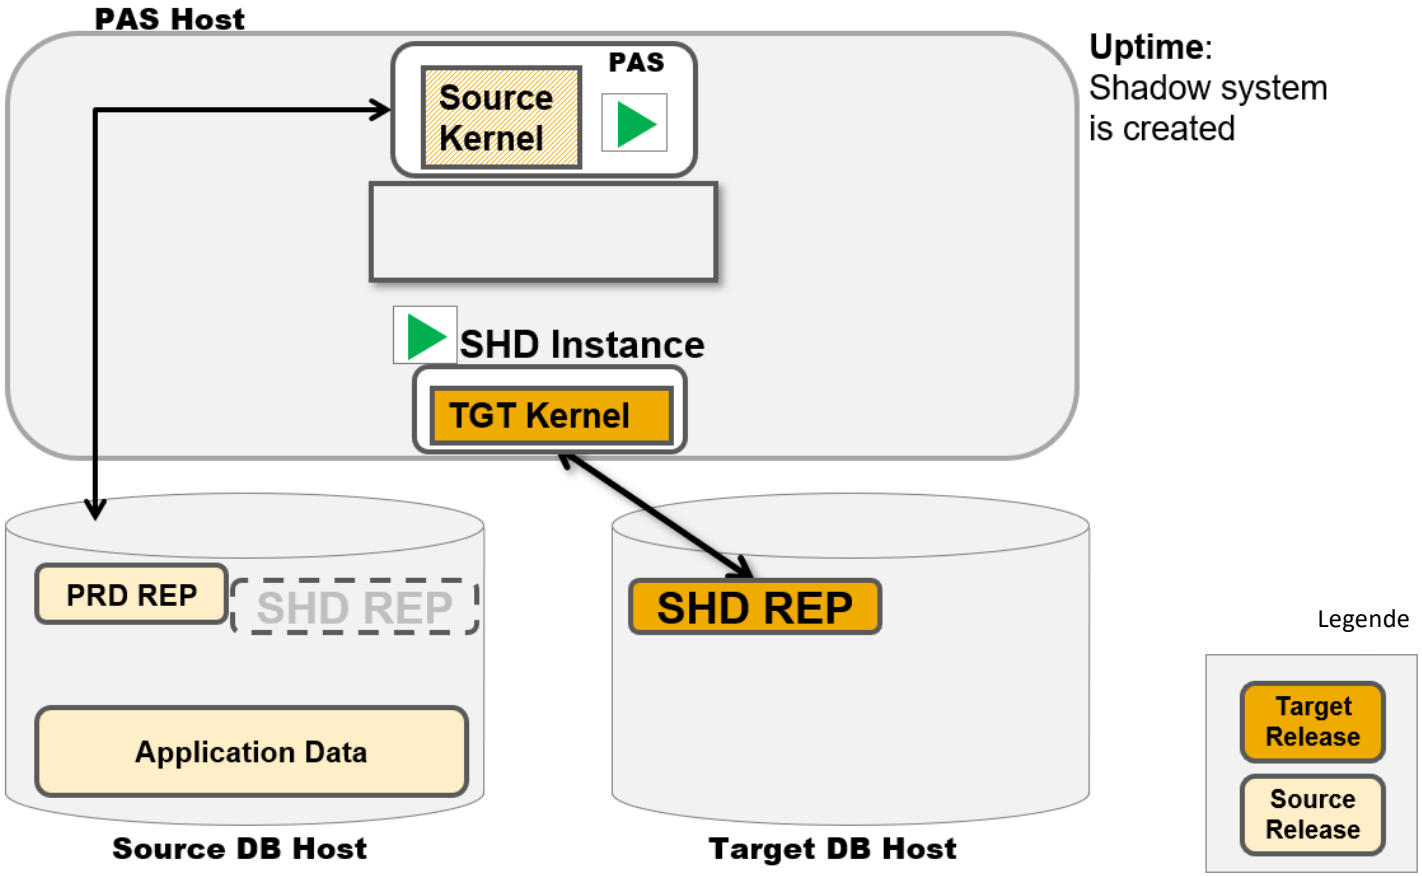

Database Migration Option (DMO) of SUM 2.0 SP07
https://launchpad.support.sap.com/#/notes/2840346
Add. info. on converting to SAP S/4HANA using SUM 2.0 SP07
https://launchpad.support.sap.com/#/notes/2840218

Releases

- SUM DMO -

- SUM SP07
- Kernel 7.77
- ENQ2

SAP S/4HANA up to 1709 *SAP BW/4HANA up to 1.0*	*SAP S/4HANA 1809 and higher* *SAP BW/4HANA 2.0 and higher*
Windows Server 64-Bit	Windows Server 64-Bit
Linux on x86_64	Linux on x86_64
AIX	AIX
Linux on Power Little Endian	Linux on Power Little Endian
Linux on Power Big Endian	- Support discontinued -
HP-UX on IA64	- Support discontinued -
Solaris on SPARC and Solaris on x86	- Support discontinued -
IBM i	- Support discontinued -
IBM z/OS	- Support discontinued -
Linux on IBM Z	- Support discontinued -

Database Migration Option (DMO) of SUM 2.0 SP(
https://launchpad.support.sap.com/#/notes/28403
Add. info. on converting to SAP S/4HANA using SU
https://launchpad.support.sap.com/#/notes/28402

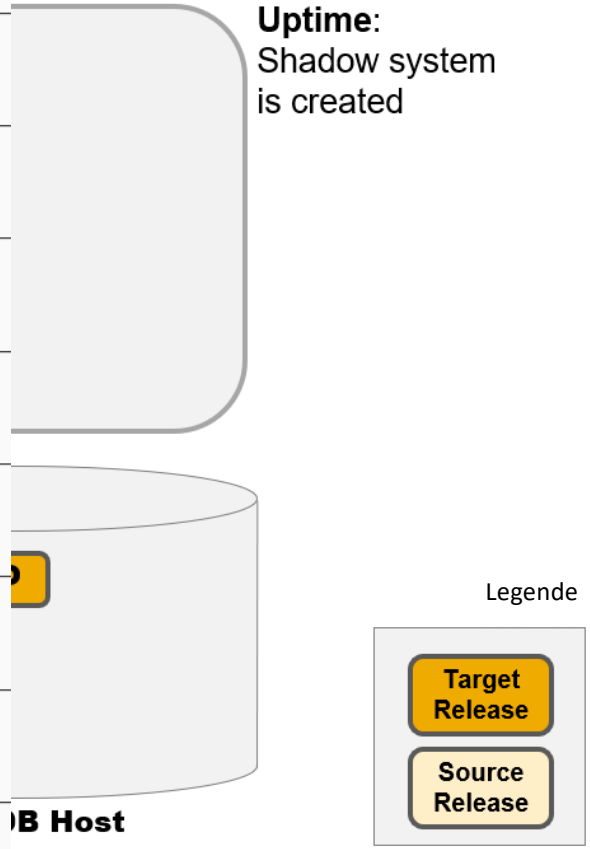

SAP HANA Express Edition

- Possibilities -

SAP HANA Express Edition

- Possibilities -

- available for free, for development and productive use, up to 32GB of RAM on premise (min. 8GB, use of XSA: min. 16GB)
- On Azure: 32GB , 64GB , 128GB , 256GB

The table shows current software and infrastructure pricing for services hosted in **EU (Frankfurt)**. Additional taxes or fees may apply.

SAP HANA, express edition (128GB)

	EC2 Instance type	Software/hr	EC2/hr	Total/hr
◉	r3.4xlarge ★*Vendor Recommended*	$1.99	$1.70	$3.69
○	r5.4xlarge	$1.99	$1.366	$3.356
○	r5a.4xlarge	$1.99	$1.246	$3.236
○	r5d.4xlarge	$1.99	$1.534	$3.524

Source: https://help.sap.com/viewer/32e387b229ec4ec0bbfdcd200af0ac5b/2.0.040/en-US/68b3c27e3d0649a09e2160e0d4f3d3c5.html

SAP HANA Express Edition

- Possibilities -

- available as docker image
- SAP HANA, express edition for Docker has been tested on the following Linux operating system versions:

Linux OS	OS Version	Docker Editions
Ubuntu	17.04 (Zesty Zapus)	Community, Enterprise
openSUSE	openSUSE Leap	Enterprise
CentOS	7 (Core)	Community, Enterprise
Debian	9 (Stretch)	Community
Fedora	28 (Server Edition)	Community

This installation does not support Docker for Windows or Docker for Mac!

Source: https://help.sap.com/viewer/32e387b229ec4ec0bbfdcd200af0ac5b/2.0.040/en-US/68b3c27e3d0649a09e2160e0d4f3d3c5.html

SAP HANA Express Edition

- Possibilities -

- Compatibility:

Processing Services

Feature	SAP HANA, express edition
Spatial/GIS	Yes
Predictive Analytics, R	Yes (PAL, R-Connector)
Graph Engine	Yes
Fuzzy Text and Fuzzy Text Search	Yes
Text Analytics and Data Mining	Yes
Smart Data Quality	No
Series Data	Yes
AFL	Yes
Python Client API for machine learning algorithms	Yes (available with SAP HANA, express edition only)
R Client API for machine learning algorithms	Yes

Integration Services

Feature	SAP HANA, express edition
Smart Data Access	Yes
SAP HANA streaming analytics	Yes
Smart Data Integration	Yes
Hadoop	Yes

Database Services

Feature	SAP HANA, express edition
Columnar, ACID compliant store for OLTP and OLAP	Yes (32 GB maximum)
Highly Optimized Compression	Yes
Multitenant Database Containers	Yes
Data Warehouse Foundation	Yes
SAP HANA dynamic tiering	No
Multiple Hosts/Scale Out	No
Backup Recovery	Yes
System Replication	Yes

Source: https://help.sap.com/viewer/32e387b229ec4ec0bbfdcd200af0ac5b/2.0.040/en-US/68b3c27e3d0649a09e2160e0d4f3d3c5.html

SAP HANA Express Edition

- Possibilities -

- Compatibility:

Application Services

Feature	SAP HANA, express edition
Extended Application Services	Yes
CDS	Yes
Server-side Javascript	Yes
Modeling, Information Composer	Yes

Development Tools

Feature	SAP HANA, express edition
HANA Studio	Yes
Web Based Development Workbench	Yes
SAP Web-IDE	Yes*

Administration/Security

Feature	SAP HANA, express edition
Monitoring and Troubleshooting	Yes
SAP HANA Cockpit	Yes*
Security	Yes
Solution Manager	No

i Note

*Requires the applications installation package. Not available in the smaller server-only installation package.

Source: https://help.sap.com/viewer/32e387b229ec4ec0bbfdcd200af0ac5b/2.0.040/en-US/68b3c27e3d0649a09e2160e0d4f3d3c5.html

SAP

- Possibilit

Component versions included in the SAP HANA, express edition 2.0 SPS 04 Revision 45 release are:

Component Name	Version	Comment
SAP HANA	HANA 2.0 SPS 04 Revision 45	Updated
XSA	1.0.121	Updated
XSA Messaging Services	1.4.5	Updated
WebIDE	4.4.18	Updated
HRTT	2.11.64	Updated
Cockpit	2.11.11	Updated
SHINE	1.7.6	Updated
SAP HANA streaming analytics	2.0.044	Updated
SAP HANA Smart Data Integration	2.4.2	Updated
SAP HANA smart data integration – Data Provisioning Agent (Linux)	2.4.2.1	Updated
SAP HANA smart data integration – Data Provisioning Agent (Windows)	2.4.2.1	Updated
EPMMDS	2.00.045.00	Updated
APL	1911	Updated
Python ML	1.0.7	Updated
R ML	1.0.4	Same
DWF	2.5.1	Same

SAP HANA Cloud

- Overview -

Cloud Buzzword Bingo

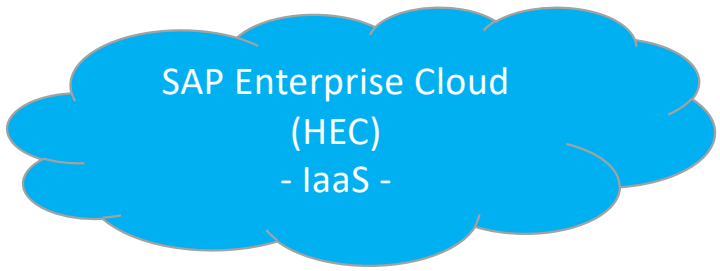

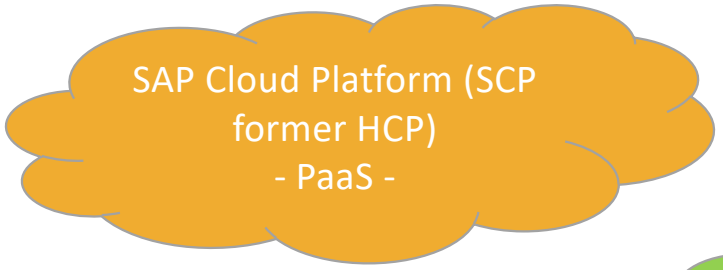

SAP HANA Cloud

- as Gateway to your sources -

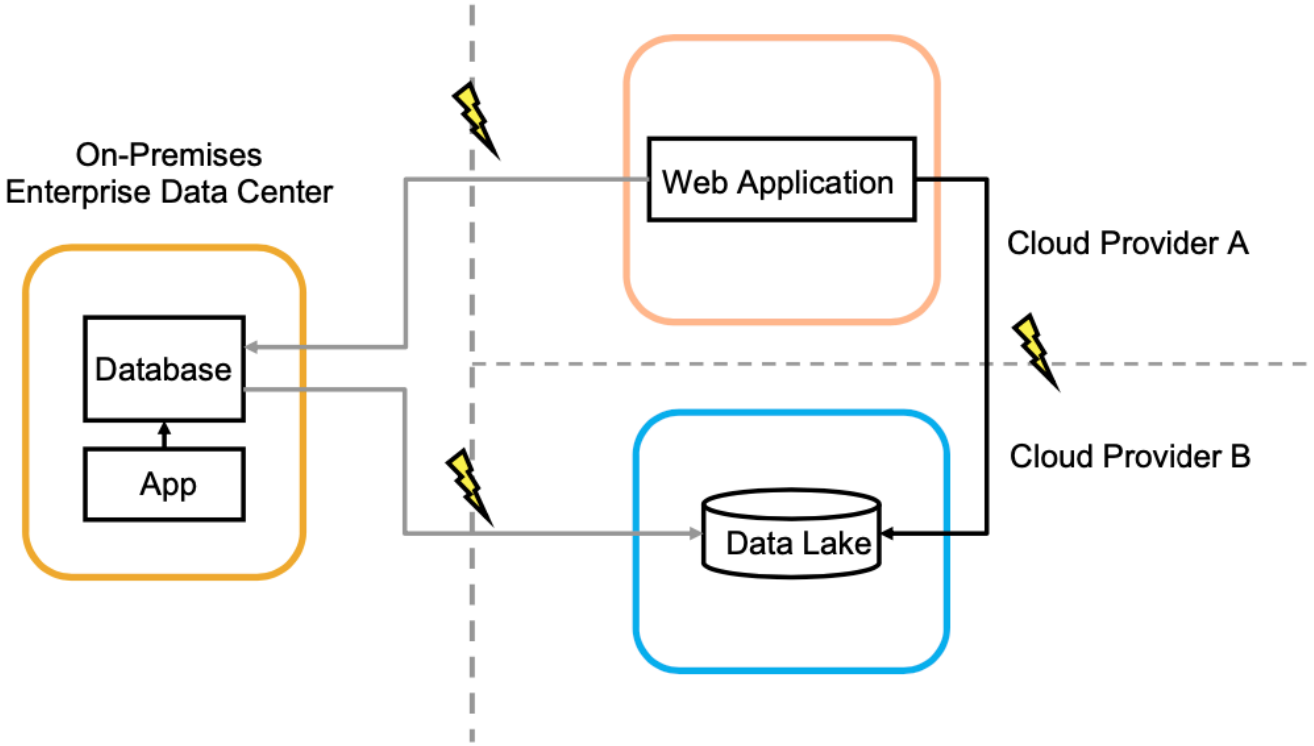

- How to Set Up?
- How To Connect?
- Performance?
- Secure Access?

"74% of enterprises describe their strategy as hybrid / multi-cloud today".

Source: D. Bartoletti, Forrester Research

SAP HANA Cloud

- Overview -

- Announced at SAP TechEd 2019 (replay)
- Trial will be usable in the next weeks
- SAP HANA Cloud is a multi-cloud database as a service (DBaaS)
- Central Store / Data Fabric as Gateway to all you sources
- How SAP HANA Cloud Solves Your Biggest Challenges
  - Expanding Volumes of Data and High Costs of Storage
  - Limited Connection Management with Cloud Sources (access to all data, no matter the source, through a single gateway)
  - In-memory Storage is Cost Prohibitive
  - Not Enough Computing Power or Storage

SAP HANA Cloud

- Overview -

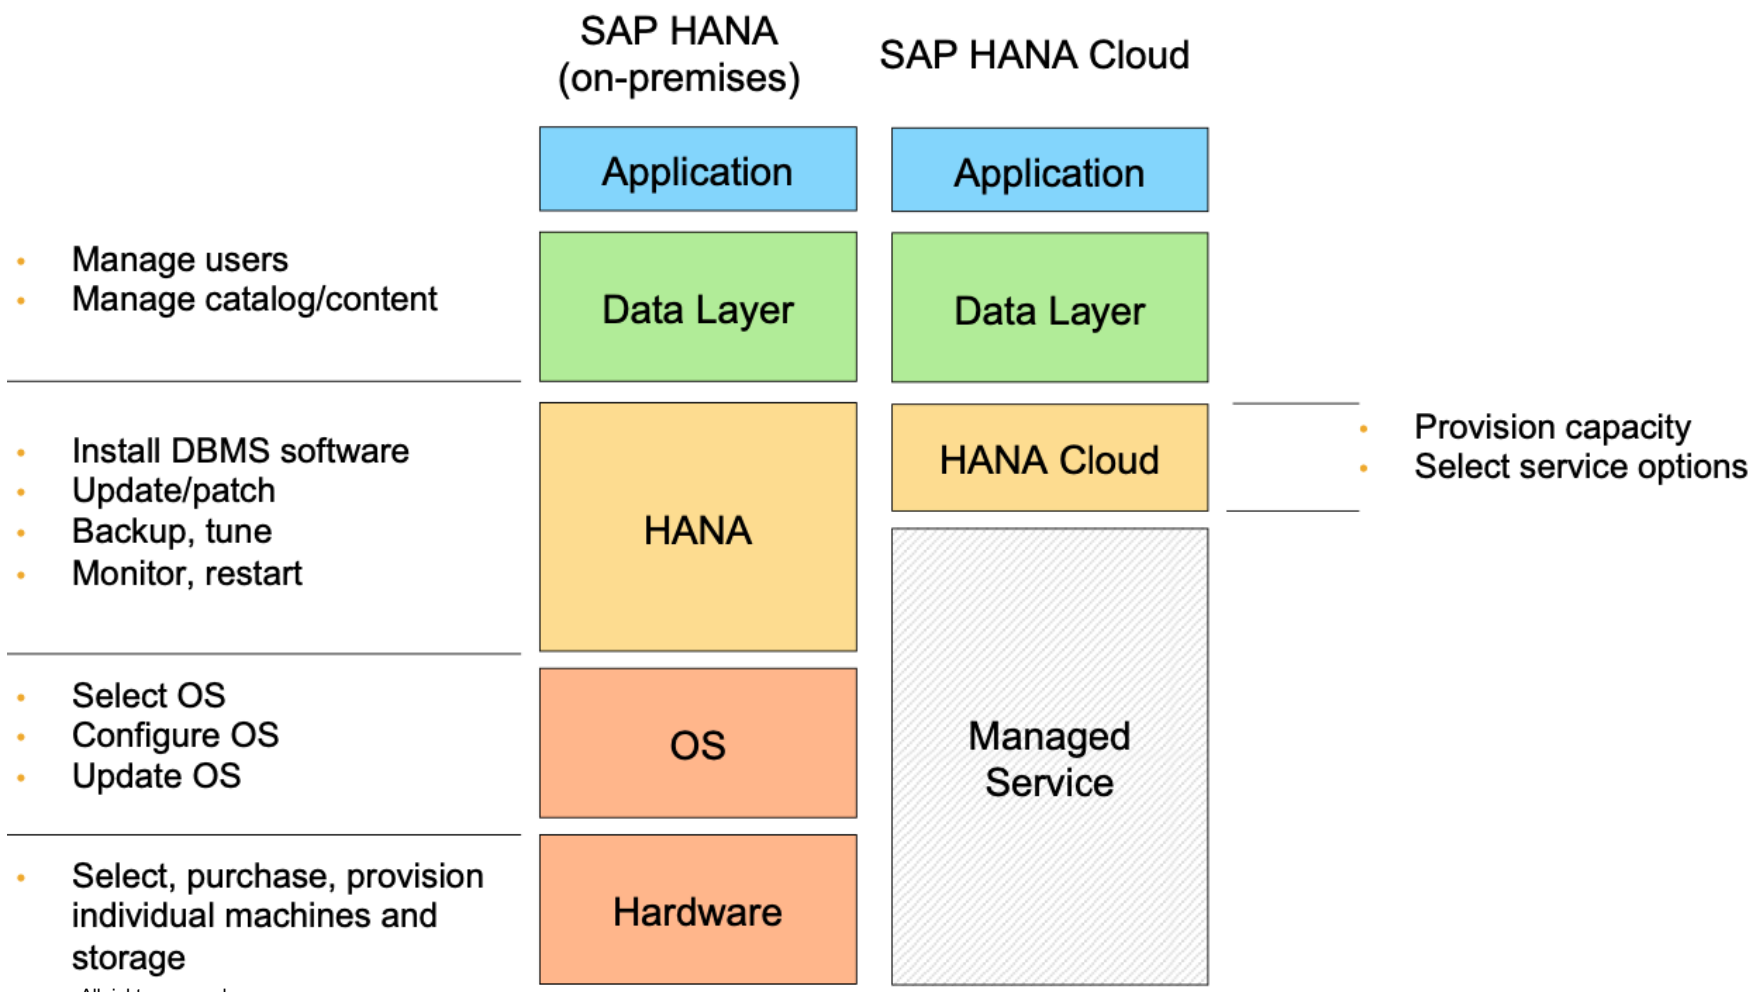

SAP HANA Cloud

- Overview -

SAP HANA Cloud: Shared responsibility for security

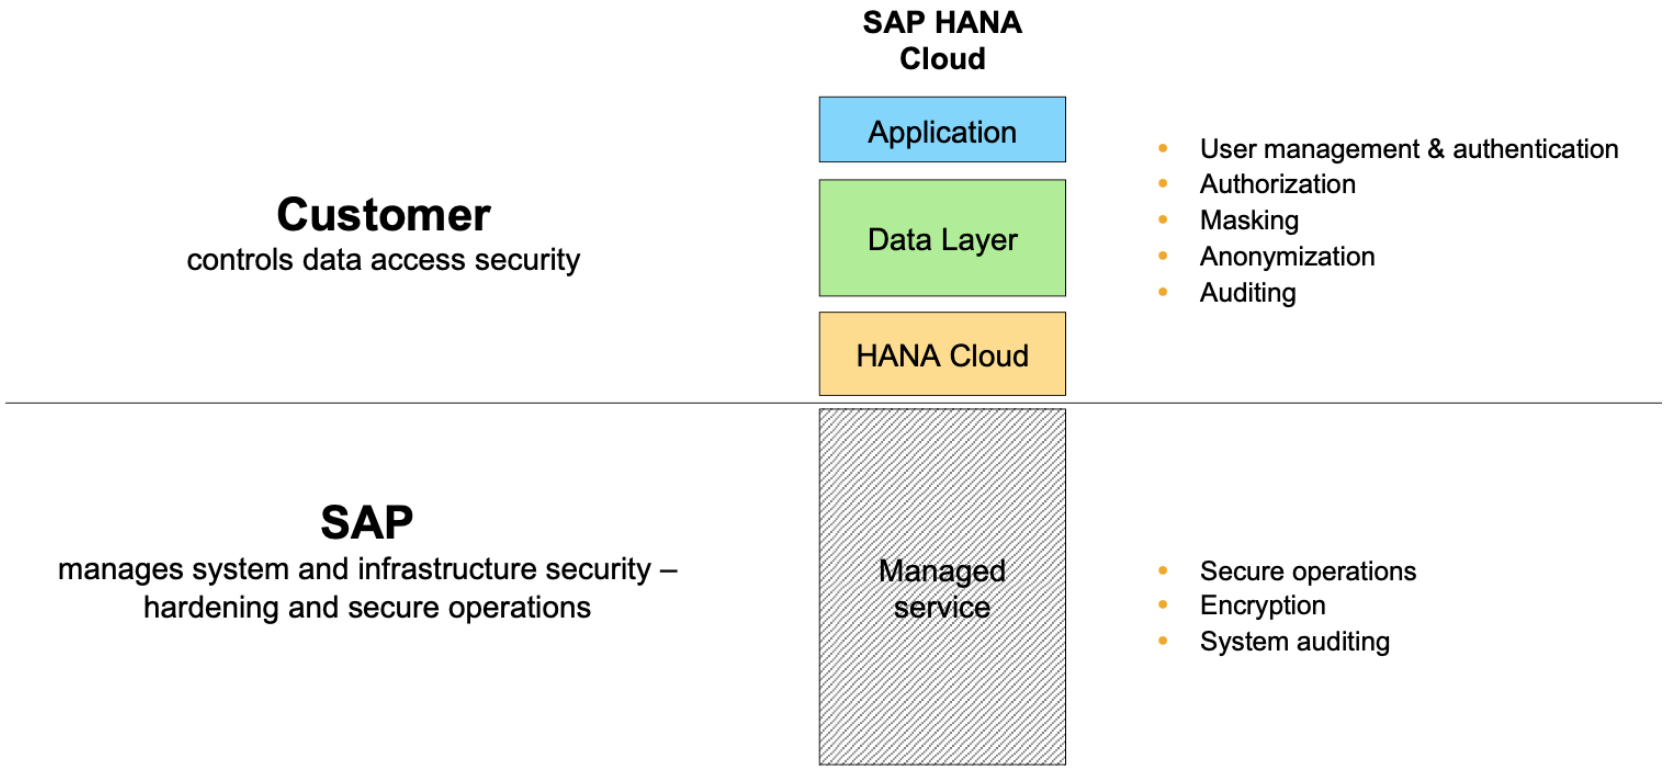

SAP HANA Cloud

- Overview -

What is the difference between SAP HANA as a Service and SAP HANA Cloud?

SAP HANA Cloud is a new cloud-native implementation of SAP HANA, while SAP HANA as a Service is an SAP HANA on-premise version running in a cloud infrastructure.

How does SAP HANA Cloud relate to the SAP Cloud Platform (SCP)?

SAP Cloud Platform is a platform as a service (Paas) for the development, integration, and operation of cloud services. SAP HANA Cloud is a database as a service (DBaaS) that is offered, provisioned, and managed by the SAP Cloud Platform.

Source: https://saphanacloudservices.com/hana-cloud/quick-answers/

SAP HANA Cloud

- Overview -

Is SAP HANA Cloud replacing SAP HANA as a Service?

Yes. SAP HANA Cloud is the successor of SAP HANA as a Service. However, we will continue supporting SAP HANA as a Service for current customers in accordance with existing contracts. SAP is fully committed to assisting SAP HANA as a Service customers during their transition to SAP HANA Cloud with both the technical support and the license migration.

What role will SAP BW have with SAP HANA Cloud and BW? Can it use SAP HANA Cloud storage or integrate with SAP HANA Cloud?

Yes, SAP BW will be able to access SAP HANA Cloud and the data lake as a federated source.

Source: https://saphanacloudservices.com/hana-cloud/quick-answers/

SAP HANA Cloud

- Overview -

Who manages or is the admin for the HANA database in SAP HANA Cloud?

SAP HANA Cloud is a DBaaS that manages parts of your instance, like the complete hardware setup and other cloud-native aspects.

The users have one admin to start with and can create a list of customizable roles and permissions for their own organization.

Are there any concrete performance metrics for performance at different layers? How could users confirm that they aren't experiencing contention on infrastructure?

You will have access to the SAP HANA Cockpit, a communications platform as a service (CPaas) within SAP HANA Cloud. Here, you can monitor various performance metrics, like CPU, memory, response time, disk I/O, and network to confirm contention on infrastructure.

Source: https://saphanacloudservices.com/hana-cloud/quick-answers/

SAP HANA Cloud

- Overview -

What is the availability of SAP HANA Cloud on hyperscalers?

SAP HANA Cloud will be available on Azure and AWS first in the first half of 2020. Availability in GCP and AliCloud are planned in late 2020.

Are there any plans to use AI/machine learning to automate scaling?

Yes. There are plans to use machine learning in SAP HANA Cloud for automated operations in future.

Source: https://saphanacloudservices.com/hana-cloud/quick-answers/

SAP HANA Cloud

- Overview -

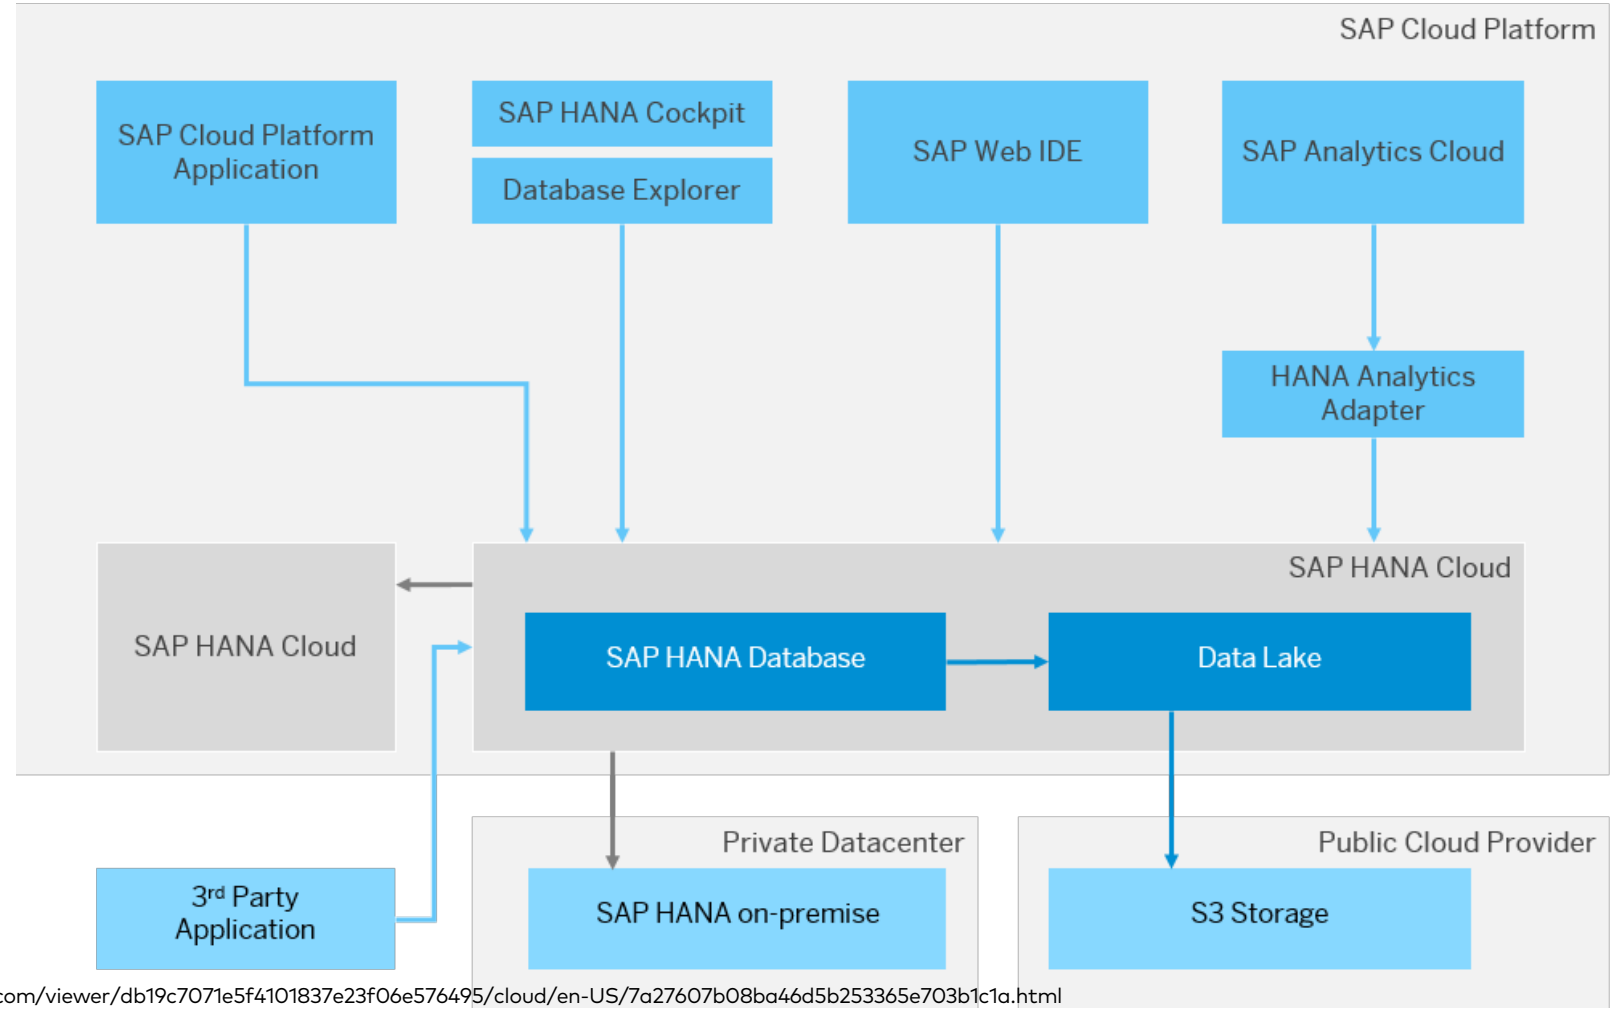

Source : https://help.sap.com/viewer/db19c7071e5f4101837e23f06e576495/cloud/en-US/7a27607b08ba46d5b253365e703b1c1a.html

SAP HANA Cloud

- Overview -

- Releases
  - A new version of SAP HANA Cloud is released **every three months**
  - Every version can be used for **up to six months**
  - Afterwards, an SAP HANA Cloud instance **is automatically upgraded** to the latest available version.
  - **submit a service request** to upgrade to the next available version

SAP HANA C

- Overview -

Features are not supported as part of SAP HANA Cloud

- As SAP HANA Cloud is a managed cloud service, not all ALTER SYSTEM commands are supported
- Anonymization with CalcViews (supported by native SQL views instead)
- BO Explorer including SQL Extensions
- Capture and Replay
- Delta Log
- DWF
- Dynamic Tiering
- EML (aka TensorFlow integration)
- Flexible Table
- SAP HANA Accelerator for ASE (A4A)
- SAP HANA CDS
- SAP HANA extended application services, classic model (XS classic)
- Live Cache
- MDC
- MDX
- Non-unique Inverted Hash Index
- R Integration
- Smart Data Quality (SDQ)
- Streaming Analytics
- Temporary Row Tables
- Text Analysis
- Text Mining
- Time Series Tables (SQL functions are still available)
- Views (JOIN / OLAP)
- No direct access to SAP HANA Cloud OS level
- Previous deprecated features in SAP HANA Platform have been removed from SAP HANA Cloud
- For recoverability in SAP HANA Cloud, disabling log is no longer supported
- SAP Host Agent has been removed from SAP HANA Cloud

For details please refer to the SAP HANA Cloud Feature Compatibility documentation.

Source https://launchpad.support.sap.com/#/notes/2868742/E

SAP HANA Cloud

- Overview -

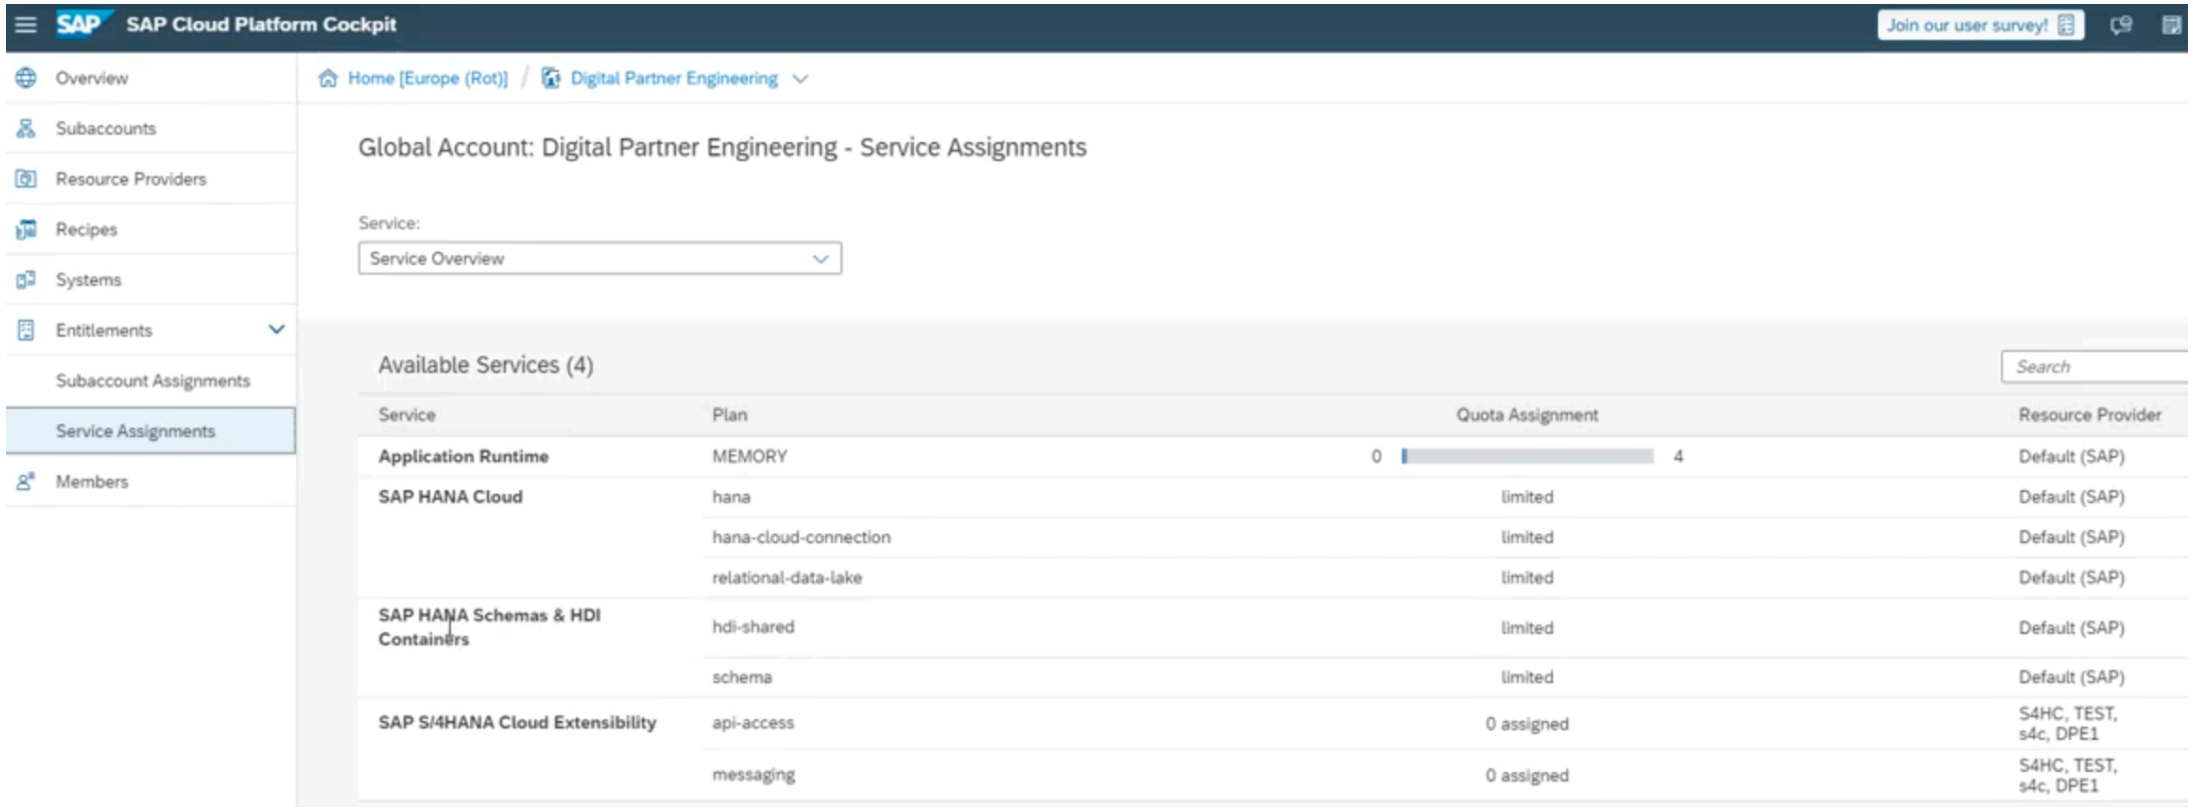

SAP HANA Cloud

- Overview -

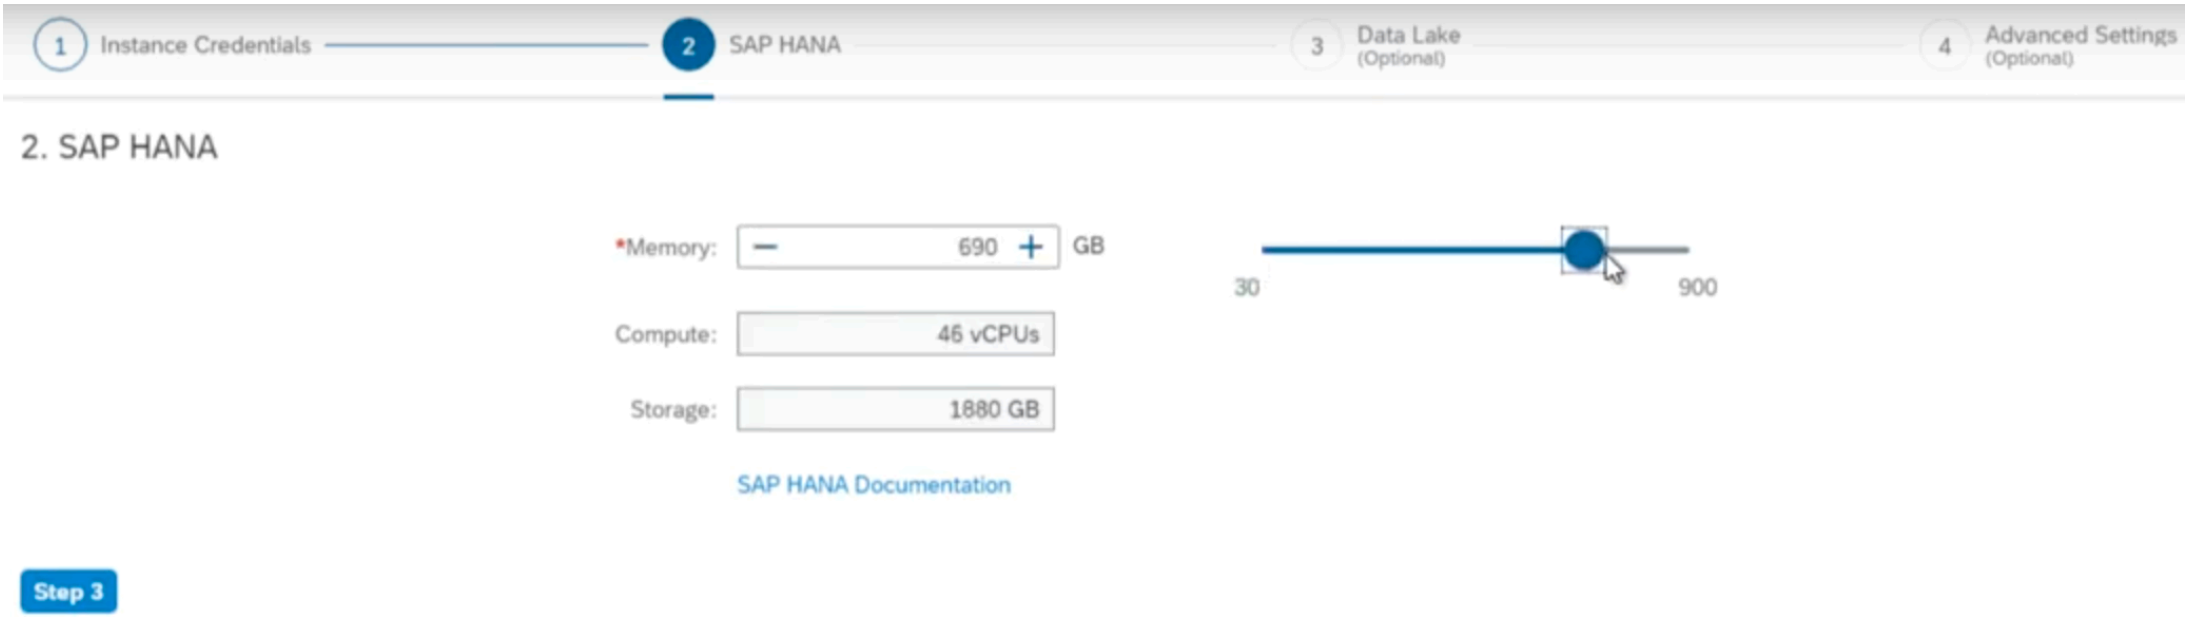

SAP HANA Cloud

- Overview -

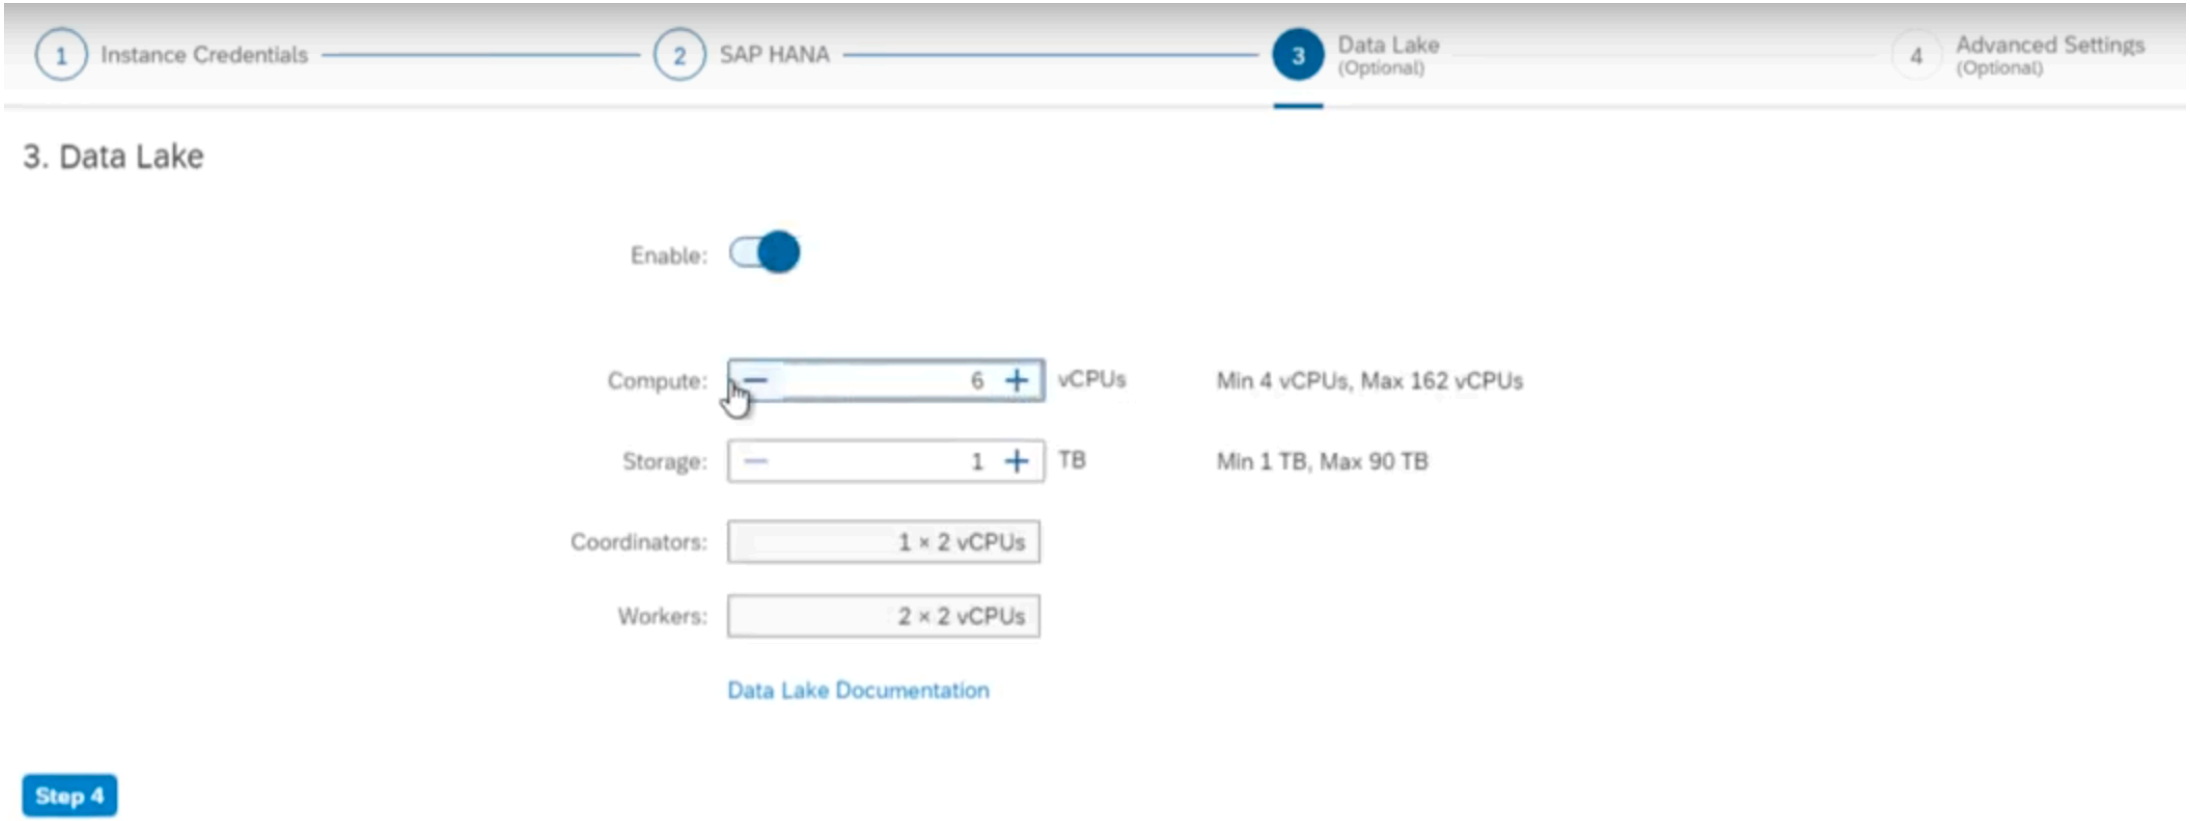

SAP HA

- Overview -

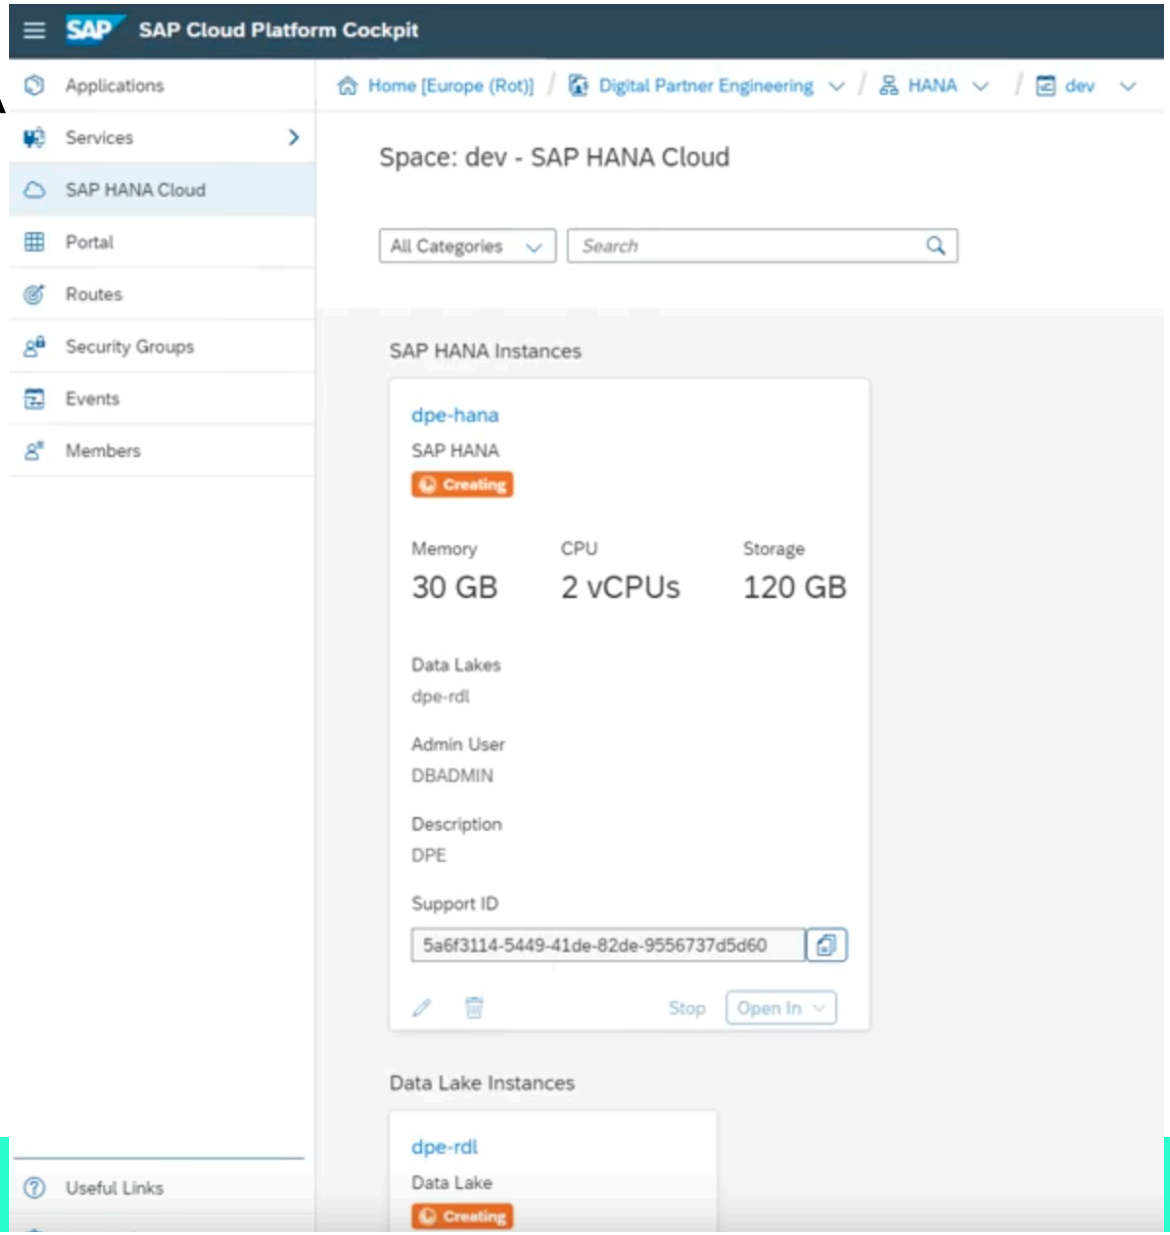

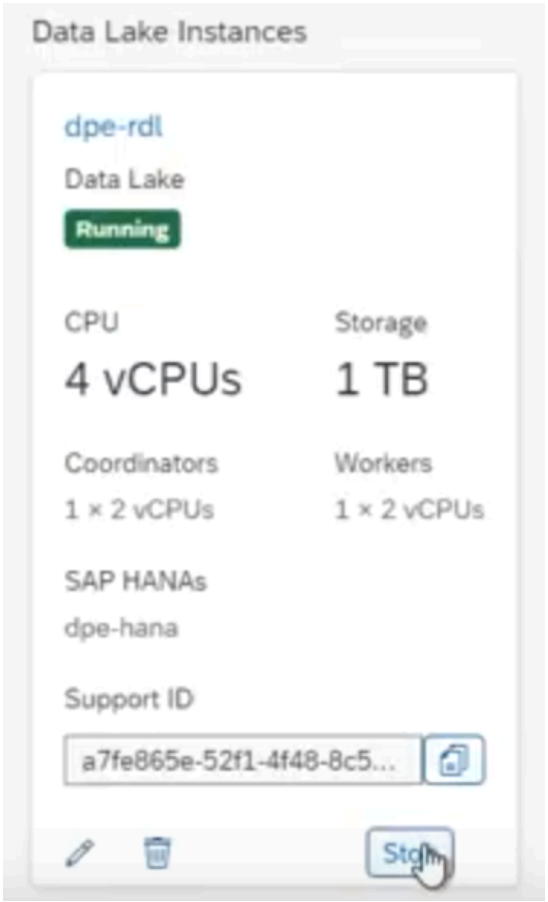

SAP HANA Cloud

- Overview -

HANA Cloud Migration Paths – Overview

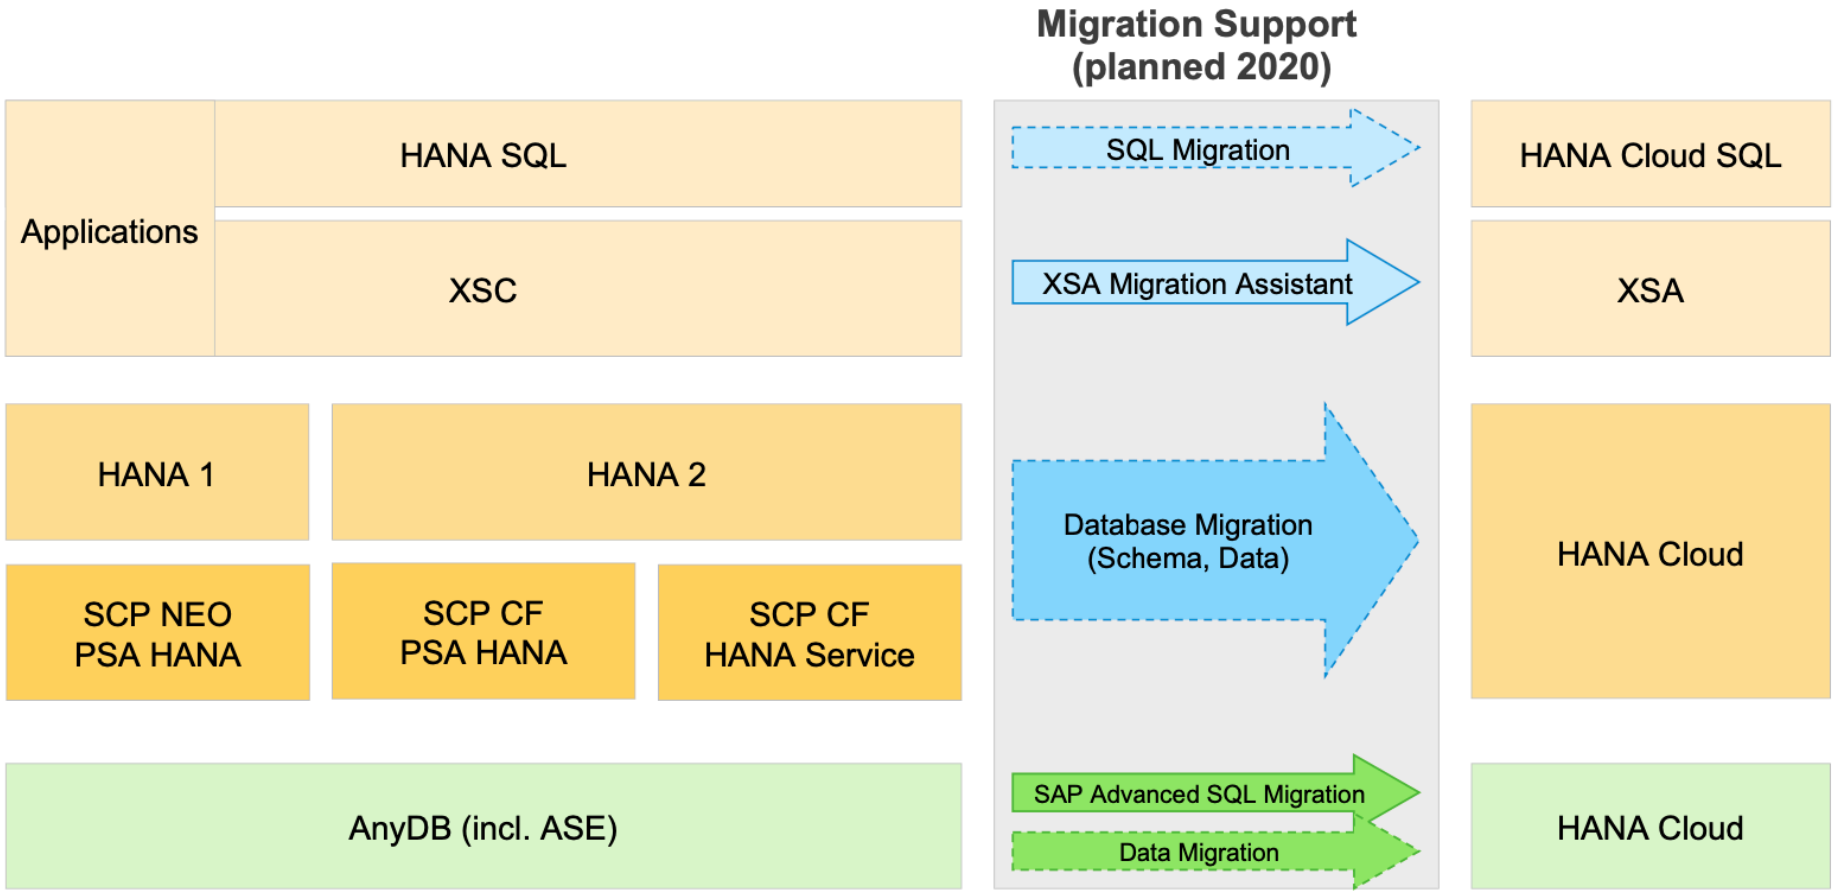

SAP HANA Cloud

- Overview -

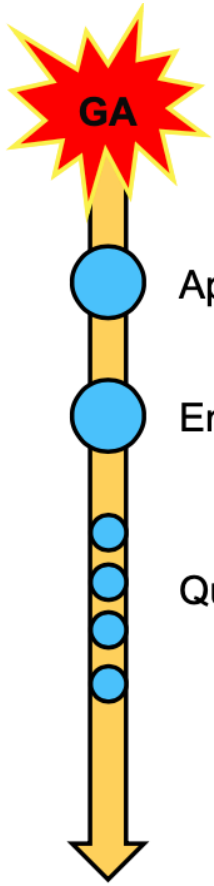

March 27th: HANA Cloud available on Azure in DC EU-West and US-West

April 2nd: HANA Cloud available on AWS in DC EU-Central and US-East

End Q2: HANA Cloud Trial availability; HANA Cloud availability in more regions

Quarterly Releases, Continuous Innovations

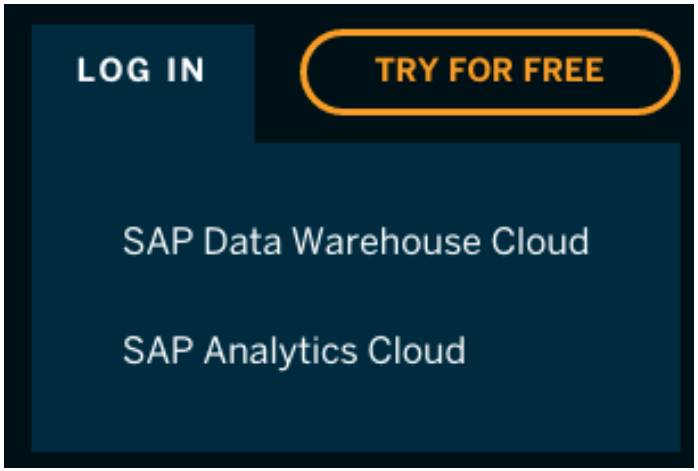

SAP HANA Cloud

- Overview -

SAP HANA Cloud Trial

If you are interested in taking part in the trial, your first step is to sign up here for more information. Once the trial is available, you will be informed via email and you can immediately sign up and start using SAP HANA Cloud.

Here is some basic information about the upcoming trial:

- Your trial account to test SAP HANA Cloud is currently planned for a 30-day duration.
- You will be given a chance to renew your trial period before having to either purchase the product or stop using it.
- The configuration of your trial instance of SAP HANA Cloud will be typically 30GB of memory, 2vCPUs, and 120GB of storage.
- If you already use other services in SAP Cloud Platform, those will not be affected or limited in any way by your participation in the SAP HANA Cloud Trial.

SAP HANA Cloud

- Overview -

- Webinars:
  - https://saphanacloudservices.com/hana-cloud/webinars/

JUN 4
8:00 AM PST /
5:00 PM CET

BEST PRACTICES

Migration to SAP HANA Cloud

In this session, we would like to give an overview of the possible migration paths, and the planned migration tools and procedures to support customers in their transition to SAP HANA Cloud. We would also like to collect our customers' opinions and suggestions on different ideas to better understand our customers' demands and needs.

MAY 13
8:00 AM PST /
5:00 PM CET

GETTING STARTED

Getting Started with SAP HANA Cloud

In this session, we would like to show you how to get started with SAP HANA Cloud, our new offering in SAP Cloud Platform. We will show you how you can provision an instance and how you can start working with that instance. In addition, we would like to collect your feedback to better understand your needs.

MAY 21
8:00 AM PST /
5:00 PM CET

DEEP DIVE

Extend your On-Premise Systems with SAP HANA Cloud

This session will provide a more detailed look on how you can extend your existing on-premise systems with SAP HANA Cloud capabilities.

MAY 5
8:00 AM PST /
5:00 PM CET

GETTING STARTED

Overview of SAP HANA Cloud's Innovative Data Lake

We developed SAP HANA Cloud with four key characteristics in mind: unlimited, trusted, simple, and open. One of the fundamental features of SAP HANA Cloud is Data Lakes. Learn how Data Lakes can meet the need of a large storage and low cost data solution. Join Jason Hinsperger, SAP's senior product manager, for a comprehensive overview of data lakes in HANA Cloud and discover SAP HANA Cloud's powerful capabilities.

- Blogs
  - https://saphanacloudservices.com/hana-cloud/resources/?resource_type=blog

SAP HANA Cloud

- Overview -

- Summary
  - cloud-native data management platform
  - Successor of HANA as a Service
  - Gateway for on-prem and cloud systems
  - Connects different cloud providers and apps
  - **Not** S/C/BW/4HANA ready
  - There are plenty new acronyms:
    - DBasS (no typos here – „DBas**s**“)
    - CPasS

SAP HANA Cloud

- Overview -

- Details:
  - Onboarding Guide
    - https://saphanacloudservices.com/hana-cloud/onboarding-guide/
  - Blog by Denys van Kempen
    - https://blogs.sap.com/2020/03/28/getting-started-with-sap-hana-cloud/
  - SAP Notes 2868742 - Differences between SAP HANA Cloud and SAP HANA Platform for SQL, SQLScript and SAP HDI (SAP HANA Deployment Infrastructure)
    - https://launchpad.support.sap.com/#/notes/2868742/E
  - SAP HANA Cloud Feature Compatibility
    - https://help.sap.com/viewer/3c53bc7b58934a9795b6dd8c7e28cf05/cloud/en-US/e131e792973348d1ac072590fe3d137c.html

Architecture

- Indexes-

Architecture

Indexes

Braucht die HANA zwingend Indizes?

Benötigen Indizes eine Persistenz?

Architecture

Indexes

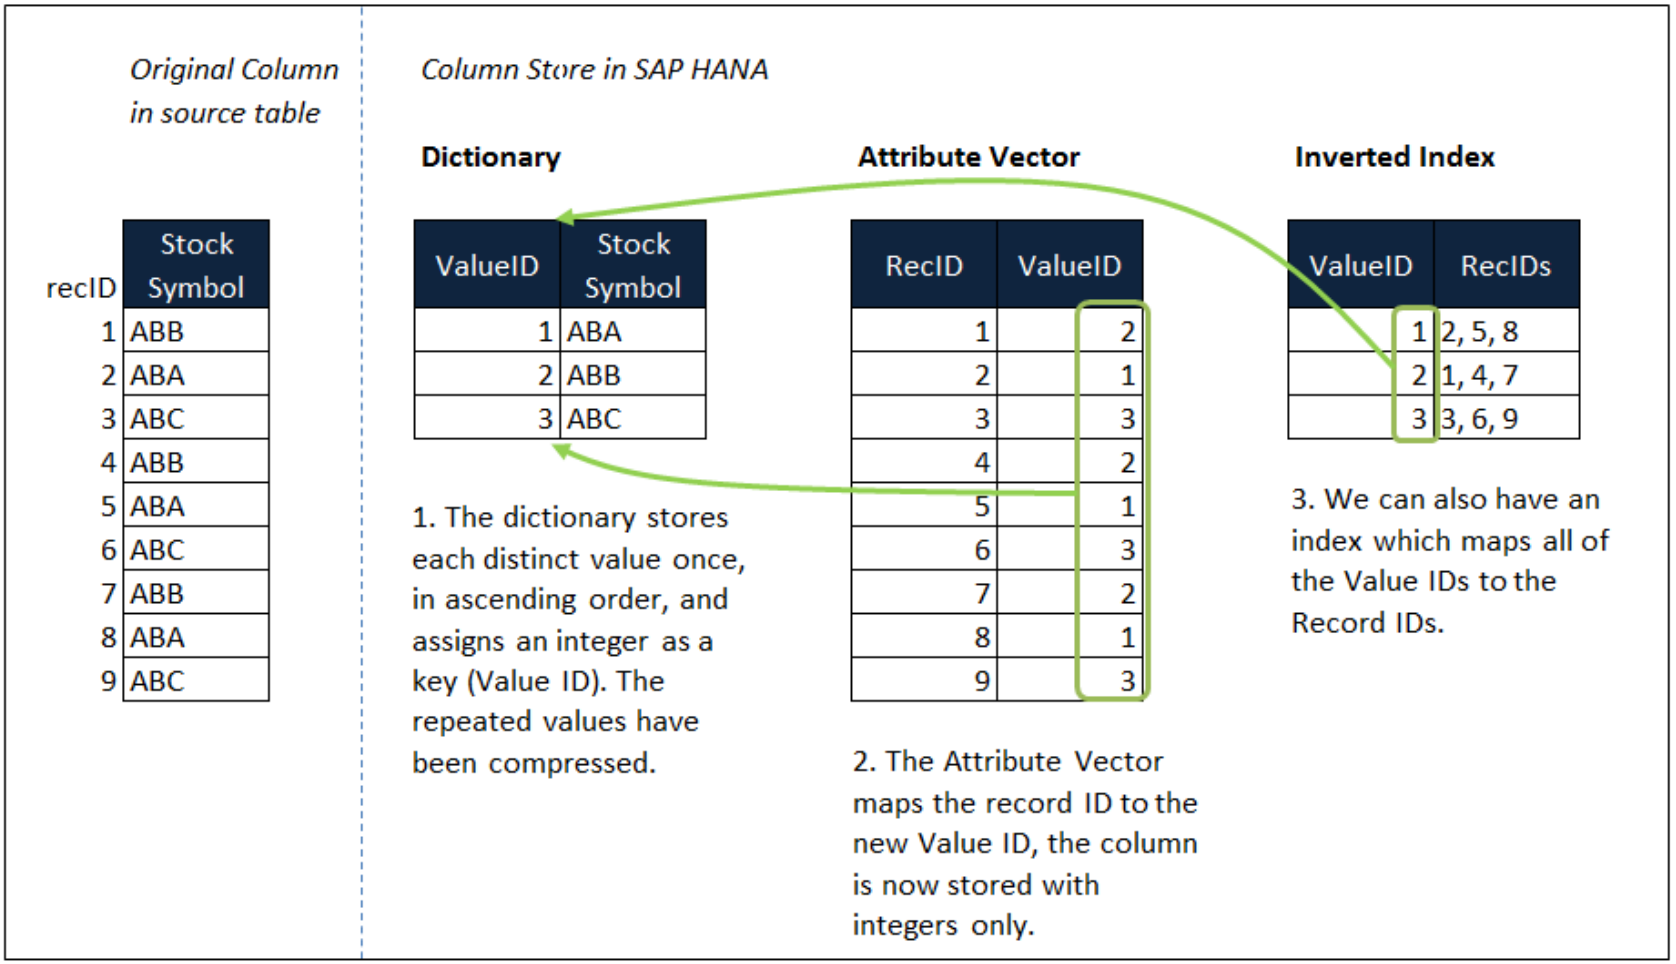

Architecture

Indexes

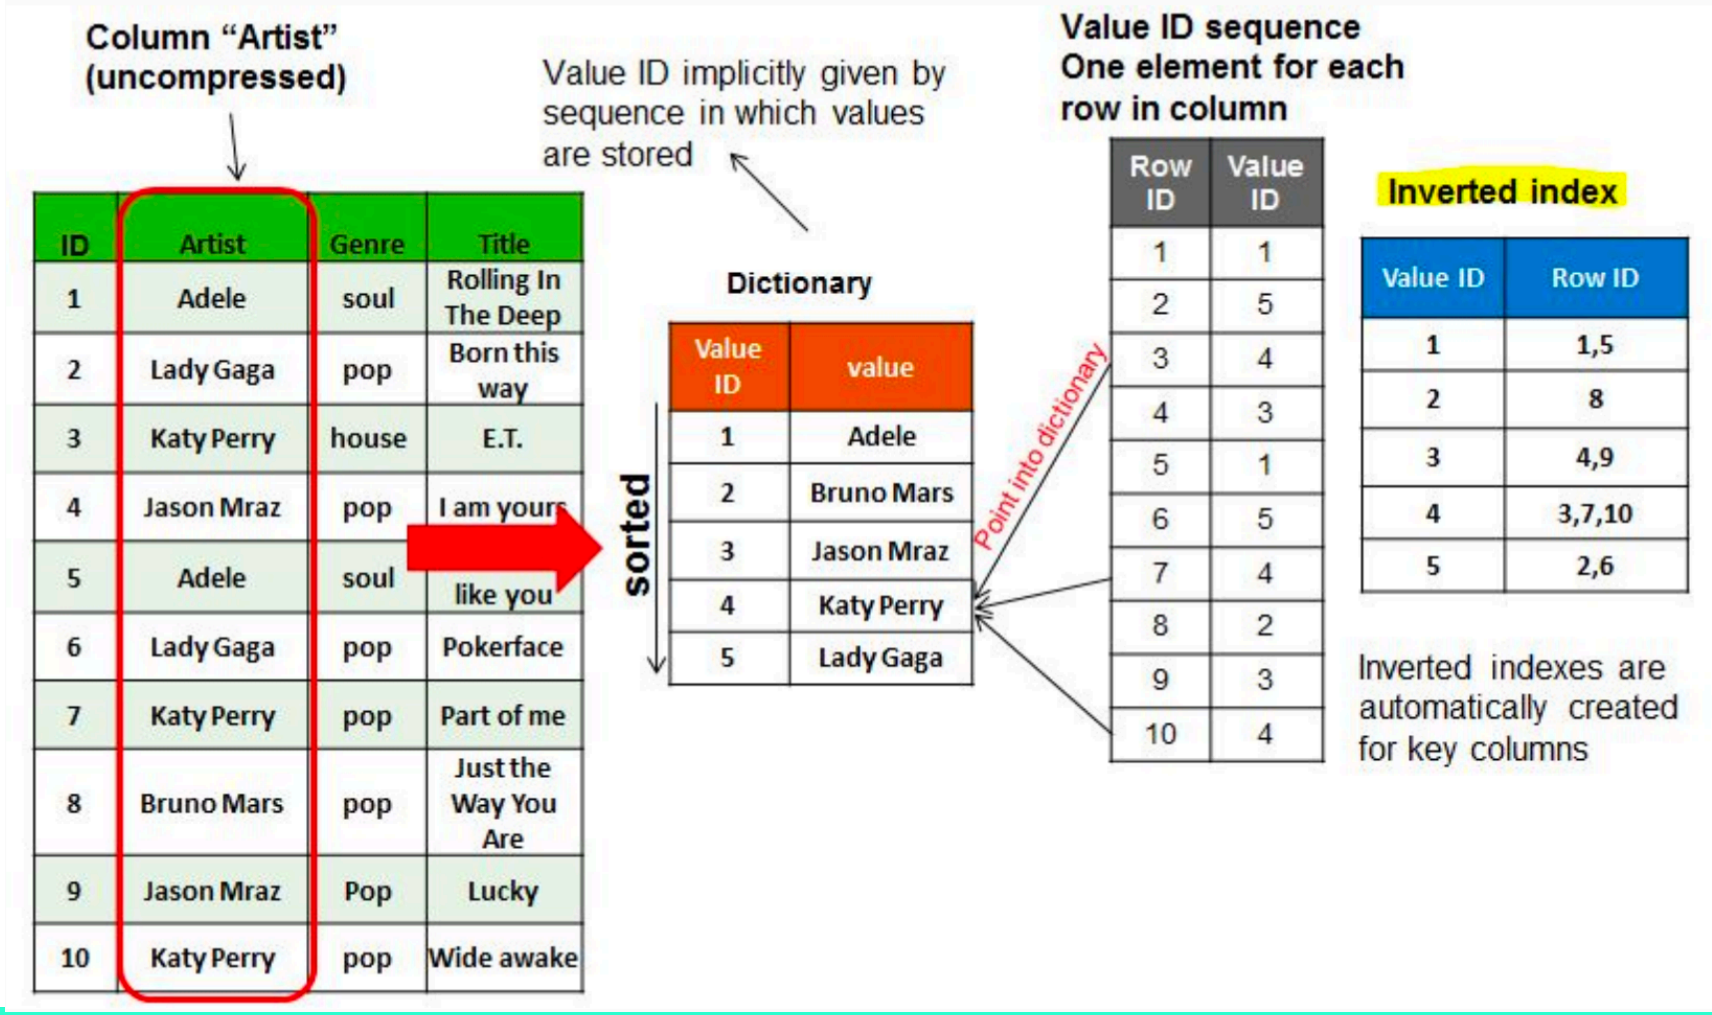

ID	Artist	Genre	Title
1	Adele	soul	Rolling In The Deep
2	Lady Gaga	pop	Born this way
3	Katy Perry	house	E.T.
4	Jason Mraz	pop	I am yours
5	Adele	soul	like you
6	Lady Gaga	pop	Pokerface
7	Katy Perry	pop	Part of me
8	Bruno Mars	pop	Just the Way You Are
9	Jason Mraz	Pop	Lucky
10	Katy Perry	pop	Wide awake

Value ID	value
1	Adele
2	Bruno Mars
3	Jason Mraz
4	Katy Perry
5	Lady Gaga

Row ID	Value ID
1	1
2	5
3	4
4	3
5	1
6	5
7	4
8	2
9	3
10	4

Value ID	Row ID
1	1,5
2	8
3	4,9
4	3,7,10
5	2,6

Architecture

Indexes

Store	Index type	SAP Note	Details	Creation command
Row store	BTREE [UNIQUE]		B*tree index on row store table	`CREATE [UNIQUE] BTREE INDEX ...`
Row store	CPBTREE [UNIQUE]	2112604	B*tree index with compressed prefix on row store table (default row store index type)	`CREATE [UNIQUE] [CPBTREE] INDEX ...`
Column store	FULLTEXT	2800008	Fulltext index Be aware that alternatively to explicit fulltext indexes also implicit fulltext indexes with reduced functionality can be created when column types BINTEXT, SHORTTEXT or TEXT are used. See the blog SAP HANA fulltext indexes for more details.	`CREATE FULLTEXT INDEX ...`

Architecture

Indexes

Column store	FULLTEXT TEXT ANALYSIS	2800008	These indexes are fulltext indexes (see above) with activated text analysis feature. In this case an implicit text analysis table with naming convention $TA_<index_name> is created.	`CREATE FULLTEXT INDEX ... TEXT ANALYSIS ...`
Column store	INVERTED HASH [UNIQUE]	2109355	Inverted hash index, more memory efficient alternative to inverted value indexes Maps column dictionary value IDs to row IDs, no B*tree structure	`CREATE [UNIQUE] INVERTED HASH INDEX...`
Column store	INVERTED INDIVIDUAL [UNIQUE]	2600076	Inverted individual index, much more memory efficient alternative to unique inverted value indexes (>= SAP HANA 2.0 SPS 03) Check uniqueness based on individual columns, no concat attribute required	`CREATE [UNIQUE] INVERTED INDIVIDUAL INDEX ...`
Column store	INVERTED VALUE [UNIQUE]		Inverted value index, default column store index type that maps value IDs of dictionary to row IDs of column Maps column dictionary value IDs to row IDs, no B*tree structure	`CREATE [UNIQUE] [INVERTED VALUE] INDEX ...`

Architecture

Indexes

Index type	Column indexed	Index size	Related compression types	Description
BLOCK	X	> 0	CLUSTERED INDIRECT LINEAR RLE RLE SPARSE (>= Rev. 122.03)	Block indexes are used for compressed columns with a few exceptions. The underlying data structure is block based which results in smaller memory footprints than FULL indexes.
BLOCK	X	0	PREFIXED SPARSE (<= Rev. 122.02)	Inverted indexes aren't possible on PREFIXED columns and SPARSE columns (SAP HANA <= 1.00.122.02), so even if an index is explicitly created, it technically doesn't exist. A manual recompression will make sure that a compatible compression type is chosen and an index is created (see SAP Note 2112604).
FULL	X	> 0	DEFAULT	Full indexes are used for columns without advanced compression. It assigns each value in the column a list of its positions. The memory consumption can be significant, particularly in case of many distinct values.

Architecture

Indexes

FULL	X	> 0	CLUSTERED
			INDIRECT
			LINEAR RLE
			PREFIXED
			RLE
			SPARSE
				Starting with SAP HANA 1.00.122.10 and 2.00.002.01 FULL indexes are possible on columns with advanced compression if certain data distribution characteristics are met and a real index is defined on the column (not just an automatically created internal index). The transition from BLOCK to FULL exclusively happens during optimize compression (SAP Note 2112604).
				
				In case of trouble (e.g. join related performance overhead as described in SAP Note 2516807) this feature can be deactived with the following setting:
				
				`indexserver.ini -> [optimize_compression] -> singleindex_consider_for_compressed_columns = false`
				
				In order to adjust columns that already have a FULL index, a forced compression needs to be triggered after the parameter was adjusted:
				
				`UPDATE "<table_name>" WITH PARAMETERS ('OPTIMIZE_COMPRESSION' = 'FORCE')`

NONE		0		A column without an inverted index has index type NONE.

NONE	X	0		If a table is created with CREATE ... LIKE, implicit inverted indexes on columns of unique indexes aren't copied with SAP HANA <= 122.09, <= 2.00.002.00 and 2.00.010 (SAP Note 2478423). In this case you have to drop and recreate the unique constraint / unique index (**Attention:** Make sure that no concurrent changes are executed on the table, otherwise there is the risk of introducing duplicate keys):
				
				`ALTER TABLE "<table_name>" DROP CONSTRAINT "<unique_constraint_name>"`
				`CREATE UNIQUE INDEX "<index_name>" ON "<table_name>" ("<index_column_1>", ..., "<index_column_N>");`

Architecture

Indexes

You can use *SQL: "HANA_Indexes_InvertedHash_Savings"* from SAP Note 1969700 in order to check for the large indexes that could take advantage of an inverted hash conversion. Column SAVING_GB lists the expected space saving.

SCHEMA_NAME	TABLE_NAME	INDEX_NAME	MAIN_CURR_GB	MAIN_INVHASH_GB	SAVING_GB	CONSTRAINT	INDEX_TYPE
SAPERP	CEALE01	CEALE01~0	64.77	12.18	52.59	PRIMARY KEY	INVERTED VALUE UNIQUE
SAPERP	POC_DB_C_OPER	POC_DB_C_OPER~0	54.97	17.57	37.40	PRIMARY KEY	INVERTED VALUE UNIQUE
SAPERP	POC_DB_C_VALUE	POC_DB_C_VALUE~0	49.44	18.38	31.06	PRIMARY KEY	INVERTED VALUE UNIQUE
SAPERP	POC_DB_C_COMMAND	POC_DB_C_COMMAND~0	47.87	18.97	28.90	PRIMARY KEY	INVERTED VALUE UNIQUE
SAPERP	COSS_BAK	_SYS_TREE_CS_#168305_#0_#P0	25.61	4.47	21.14	PRIMARY KEY	INVERTED VALUE UNIQUE
SAPERP	VRPMA	VRPMA~0	28.11	7.90	20.20	PRIMARY KEY	INVERTED VALUE UNIQUE
SAPERP	VAKPA	VAKPA~0	25.76	6.58	19.17	PRIMARY KEY	INVERTED VALUE UNIQUE
SAPERP	VAPMA	VAPMA~0	24.48	5.54	18.93	PRIMARY KEY	INVERTED VALUE UNIQUE

Architecture

Indexes

No indexes are required.	It is true that SAP HANA can scan large amounts of data even without the existence of an index. Nevertheless there are still situations where indexes are useful: • Support of a primary key or unique key constraint (e.g. primary indexes in SAP ABAP scenarios) • Optimization of central, performance-critical queries • Optimization of queries consuming a lot of CPU resources due to highly parallelized operations • Tables in row store need indexes similarly like tables in classic relational databases For more information on SQL statement optimization and indexes see SAP Notes 2000002 and 2160391.
Indexes are not persisted to disk.	This is not generally true. For example, the main storage of column store concat attributes related to indexes is generally persisted to disk. See SAP Note 2160391 ("Are indexes persisted to disk?") for more details.

Architecture

Indexes

Zwei SAP Systeme mit dem gleichen Dateninhalt

ORACLE

```
SQL> select count(*) from
dba_tables where owner='SAPSR3';

  COUNT(*)
----------
    118.618


SQL> select count(*) from
dba_indexes where
owner='SAPSR3';

  COUNT(*)
----------
    147.323
```

S/4HANA (ohne implizite Indizes)

```
select count(*) from tables
where schema_name='SAPSR3'

 COUNT(*)
----------
141.288

select count(*) from indexes
where schema_name='SAPSR3'

 COUNT(*)
----------
142.131
```

Architecture

Indexes

SQL | Result

```
EXPLAIN PLAN FOR
select * from KB_COLUMN where a=1
```

	OPERATOR_NAME	OPERATOR_DETAILS
1	PROJECT	KB_COLUMN.A, KB_COLUMN.B, KB_COLUMN.C
2	COLUMN TABLE	FILTER CONDITION: KB_COLUMN.A = 1 (DETAIL: ([INDEX LOOK-UP] KB_COLUMN.A = 1))

Architecture

Indexes

Internal column type	Naming convention	Description
CONCAT_ATTRIBUTE	$<col1>$<col2>$...$	Multi column index (concat attribute)
	$uc_<index>$	Unique index
ROWID	$rowid$	Row pointer, unique identifier for a record
TEXT	$_SYS_SHADOW_<type>$	Fulltext index (SAP Note 2800008)
TREX_EXTERNAL_KEY	$trexexternalkey$	Multi column primary key index of type inverted value or inverted hash
TREX_UDIV	$trex_udiv$	Transaction and visibility management
VALID_FROM	$valid_from$	Validity information for records in history tables (begin time)
VALID_TO	$valid_to$	Validity information for records in history tables (end time)

```
select table_name, column_name, compression_type,index_type,
index_loaded from M_CS_COLUMNS where table_name='BALDAT'
```

TABLE_NAME	COLUMN_NAME	COMPRESSION_TYPE	INDEX_TYPE	INDEX_LOADED
BALDAT	MANDANT	RLE	BLOCK	LOADED
BALDAT	RELID	RLE	BLOCK	LOADED
BALDAT	LOG_HANDLE	SPARSE	FULL	LOADED
BALDAT	BLOCK	INDIRECT	FULL	LOADED
BALDAT	SRTF2	CLUSTERED	FULL	LOADED
BALDAT	_DATAAGING	PREFIXED	NONE	NOT APPLICABLE
BALDAT	CLUSTR	INDIRECT	BLOCK	LOADED
BALDAT	CLUSTD	DEFAULT	NONE	NOT APPLICABLE

```
select table_name, column_name, compression_type,index_type,index_loaded
from M_CS_ALL_COLUMNS where table_name='BALDAT'
```

TABLE_NAME	COLUMN_NAME	COMPRESSION_TYPE	INDEX_TYPE	INDEX_LOADED
BALDAT	MANDANT	RLE	BLOCK	LOADED
BALDAT	RELID	RLE	BLOCK	LOADED
BALDAT	LOG_HANDLE	SPARSE	FULL	LOADED
BALDAT	BLOCK	INDIRECT	FULL	LOADED
BALDAT	SRTF2	CLUSTERED	FULL	LOADED
BALDAT	_DATAAGING	PREFIXED	NONE	NOT APPLICABLE
BALDAT	CLUSTR	INDIRECT	BLOCK	LOADED
BALDAT	CLUSTD	DEFAULT	NONE	NOT APPLICABLE
BALDAT	$trex_udiv$	DEFAULT	NONE	NOT APPLICABLE
BALDAT	$rowid$	DEFAULT	FULL	LOADED
BALDAT	$trexexternalkey$	DEFAULT	FULL	LOADED

Quelle: SAP Hinweis: 1986747 - How-To: Analyzing internal Columns in SAP HANA Column Store

Architecture

Indexes

KEY_FIGURE	VALUE_1	VALUE_2	CHANGE_PCT
Begin time	2020/01/21 00:00:00	2020/01/26 06:30:00	
End time	2020/01/26 06:00:00	2020/01/29 10:53:08	
Statement hash	1943857aaf3165f6488ae0c6629a24bd	1943857aaf3165f6488ae0c6629a24bd	
Table types	COLUMN	COLUMN	
Sharing type	GLOBAL	GLOBAL	
Distributed	FALSE	FALSE	
Executions	196163.00	113730.00	-42.02
Records per execution	10.61	8.95	-15.58
Cursor time (ms)	12918856.27	796087.75	-93.83
Cursor time (ms) per execution	65.85	6.99	-89.37
Cursor time (ms) per record	6.20	0.78	-87.40
Execution time (ms)	12879671.91	783514.47	-93.91
Execution time (ms) per execution	65.65	6.88	-89.50
Execution time (ms) per record	6.18	0.76	-87.57
Preparation time (ms)	202.42	322.42	59.28
Preparation time (ms) per execution	0.00	0.00	174.73
Preparation time (ms) per record	0.00	0.00	225.43
Trans. lock time (ms)	0.00	0.00	0.00
Trans. lock time (ms) per execution	0.00	0.00	0.00
Trans. lock time (ms) per record	0.00	0.00	0.00

Architecture

Indexes

KEY_FIGURE	VALUE_1	VALUE_2	CHANGE_PCT
Begin time	2020/01/21 00:00:00	2020/01/26 06:30:00	
End time	2020/01/26 06:00:00	2020/01/29 10:53:08	
Statement hash	1943857aaf2165f6488ae0	1943 6488ae0c6629a24bd	
Table types			
Sharing type			
Distributed			
Executions		0.00	-42.02
Records		95	-15.58
Cursor ti		75	-93.83
Cursor time		.99	-89.37
Cursor time		0.78	-87.40
Execution		783514.47	-93.91
Execution t		6.88	-89.50
Execution tim		0.76	-87.57
Preparation time (ms)		322.42	59.28
Preparation time (ms) per e	00	0.00	174.73
Preparation time (ms) per record	0.00	0.00	225.43
Trans. lock time (ms)	0.00	0.00	0.00
Trans. lock time (ms) per execution	0.00	0.00	0.00
Trans. lock time (ms) per record	0.00	0.00	0.00

Worin unterscheiden sich Cursor und Execution Time?

Architecture

Indexes

Time column type	Description
CURSOR	Contains the overall cursor time, i.e. from start of the execution until having sent the last package to the client; when the client processes the data in multiple fetches, the network and client time during these fetches is included in the cursor time. Preparation time **isn't** included
EXECUTION	Contains the execution time (open + fetch + lock wait + close) on SAP HANA server side, does **not** include preparation time
PREPARATION	Preparation time is linked to the initial parsing and compilation. You can reduce this overhead by using bind variables rather than literals.
LOCK_WAIT	Contains the transaction lock wait time, internal locks are not included

Server + Client

Server

Architecture

Indexes

2160391 - FAQ: SAP HANA Indexes

2600076 - FAQ: SAP HANA Inverted Individual Indexes

2109355 - How-To: Configuring SAP HANA Inverted Hash Indexes

1794297 - Secondary Indexes for S/4HANA and the business suite on HANA