Ný kynslóð gagnagrunna

HANA nú

Applicon

UT MESSAN 2013

SAP HANA

Hvað er HANA

Forsendur fyrir HANA

Af hverju HANA

Hvar er HANA notað

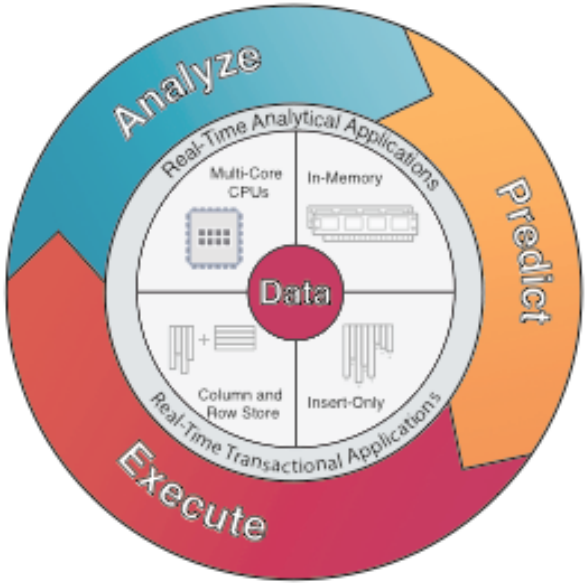

Applicon

Hvað er HANA

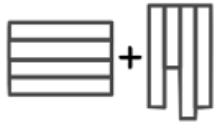

Gangagrunnur – en mjög sérstækur

- Gögn geymd í minni – RAM er ~10.000 hraðari en diskar
- Gögn geymd í dálkum – Mikil hraðaaukning og samþjöppun
- Gögn líka geymd í röðum
- Samhliðavinnsla

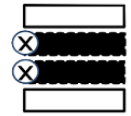

Transactional og analytical gagnagrunnur á einum stað

- Einfaldar kerfishögun, gagnalíkön, forritun + færri villur
- Breyttar kröfur til vöruhúss gagna - (pre)aggregates, indexar, stjörnuskema, tuning, gagnaflutningi
- Rauntíma OLAP á OLTP gagnahögun
- Vinnsla og útreikningur framkvæmdur í gagnagrunninum

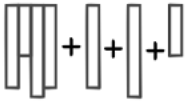

Dæmi

- 1.000 Terabyte grunnur gefur fyrirhafnarlaust 1 sek svartíma
- „Lítill" SQL grunnur gefur 6 sek svartíma – eftir mikla aðkomu sérfræðinga og 1 klst indexing vinnslu

Applicon

Hvað er HANA – hagkvæmni

Ný högun mun lækka kostnað umtalsvert

- Sparar gagnaflutning og kostnaðarsama vinnslu á gögnum
- Öryggismál og aðgangsstýringar á einum stað
- Lægri vélbúnaðarkostnaður

- Verðþrýstingur á núverandi gagnagrunnsframleiðendur

Forsendur fyrir HANA

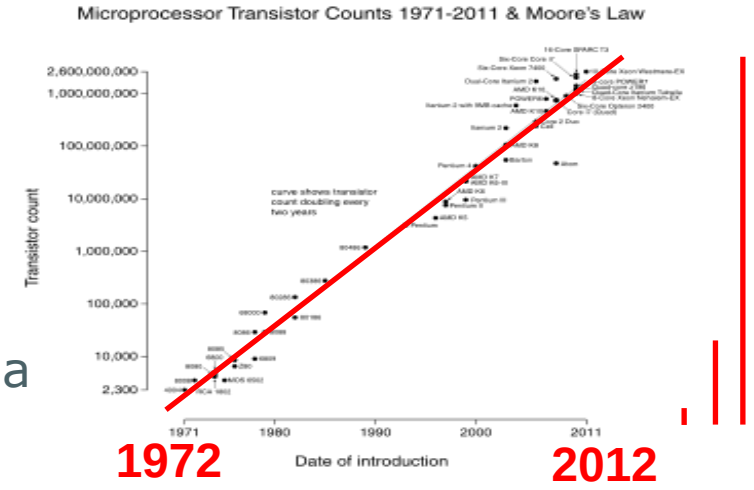

Moor's Law

- Afköst tvöfaldast á tveggja ára fresti
- Gagnagrunnar hafa ekki þróast samhliða

HANA þróun hefst 2006

- Lítil kerfi til að byrja með
- Vélbúnaðarþróun þó fyrirsjáanleg

HANA hönnunakrafa frá upphafi

- Engar samtölur í gagnagrunni
- SQL skipanir og samtölur útbúnar í rauntíma

Applicon

Af hverju HANA

5 víddir í nútíma greiningarkerfum

- Breidd — mikið gagnamagn
- Dýpt — flóknar fyrirspurnir
- Hraði — svartími og sveigjanleiki
- Einfaldleiki — einfaldur rekstur
- Rauntími — glóðvolg gögn

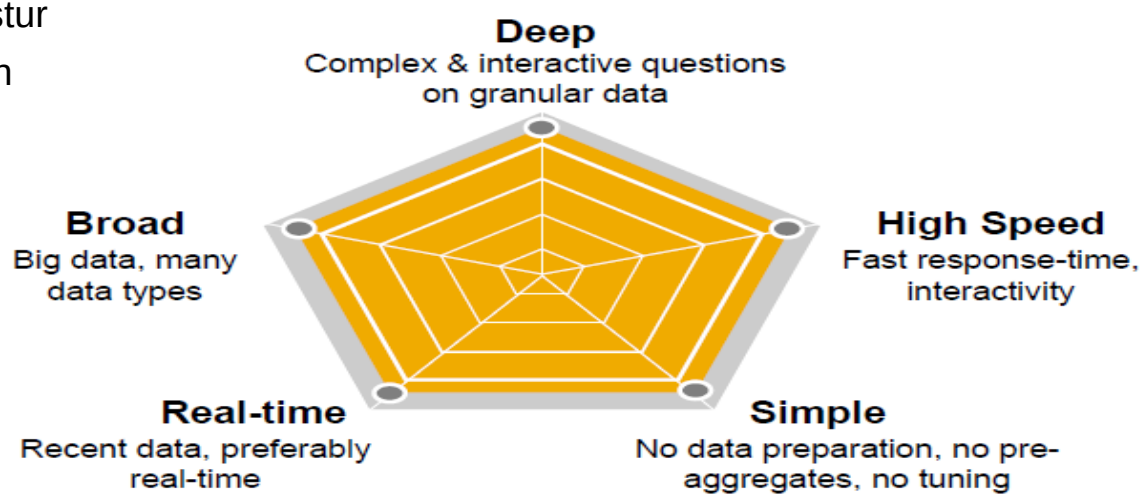

Af hverju HANA

Hraðari og fullkomnari….

- Vinnsla í venjulegum forritum
- Greiningarkerfi
- Áætlanagerð
- Spálíkön
- Hugbúnaðarþróun
- …
- Ákvarðanataka
- Stuðningur við snjalltæki - 2 sek krafa

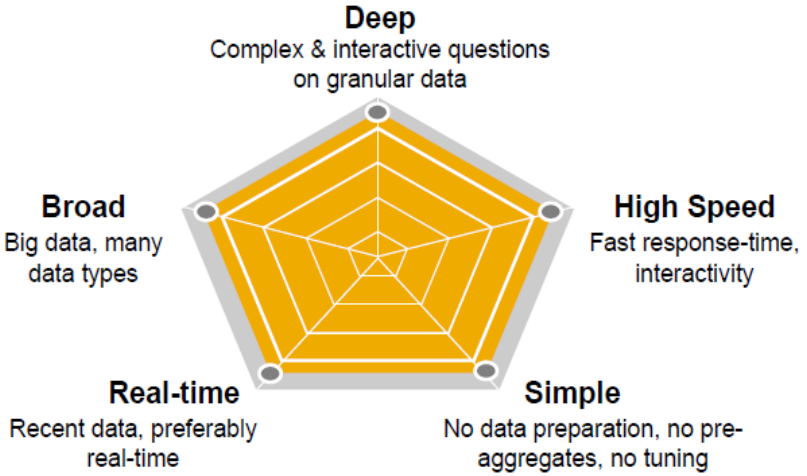

Upplýsingar í rauntíma

- Svörin strax – ekki á morgun

Applicon

Hvað ef HANA væri flugvél

Á meðan HANA flýgur á milli San Fransico og Tokyo þá.....

Komust DW notendur 5-10x hægar án HANA

Boston

Bank of America fór 56x hægar án HANA
* Afkomugreining á 100 milljónum færslna

Las Vegas

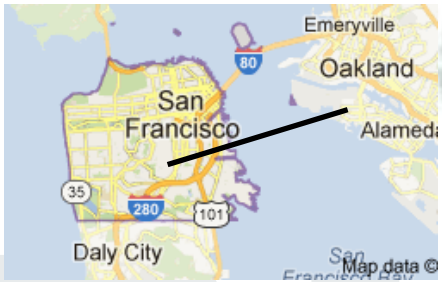

Margir með 1000x hraðabreytingu eftir HANA
* sjá dæmi á saphana/com

Oakland

Applicon

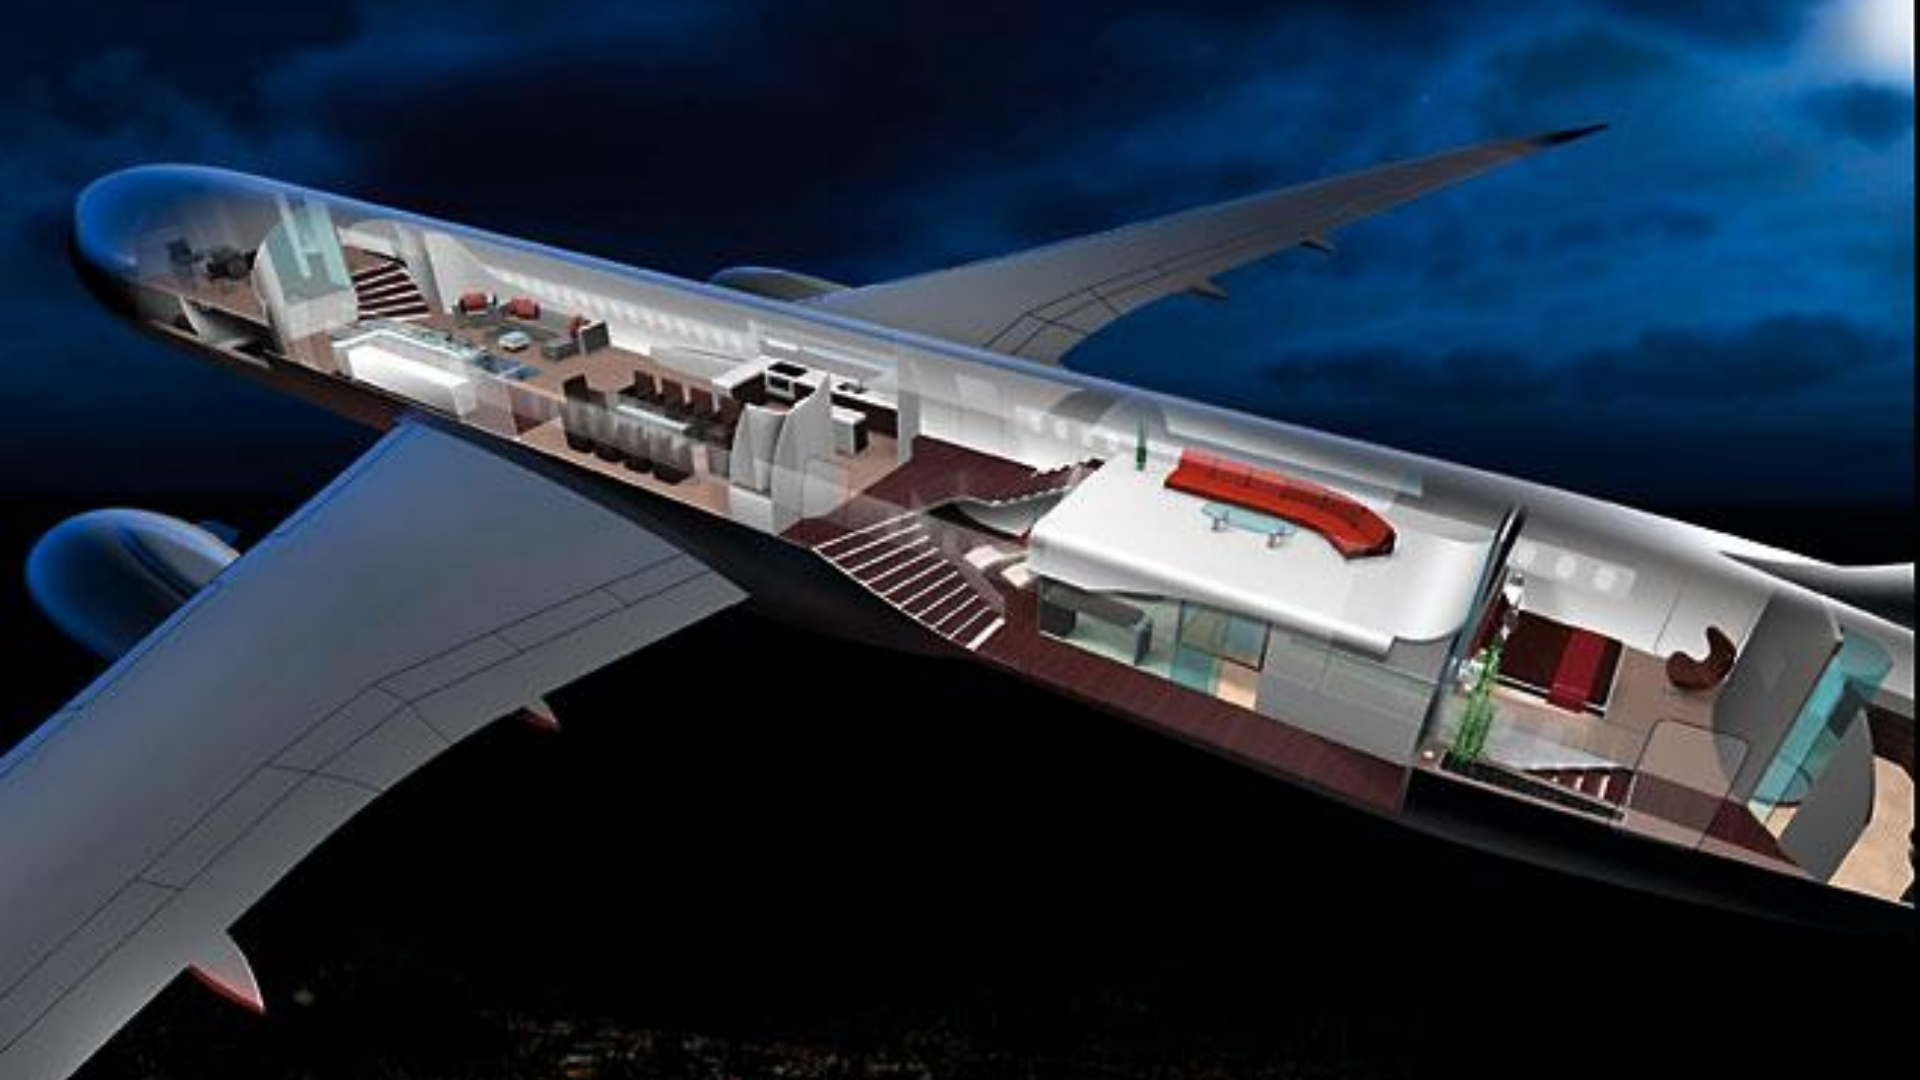

Af hverju HANA - Hverjir fljúga hraðar

18 mánuðir >800 millionUSD í tekjur ~1.000 viðskiptavinir >65.000 notendur

MEDTRONICS

60x Faster Processing Queries

10x Data Compression
From 1.5 TB to 150 GB

250x Better Complaint Analysis
(Long Text Data)

MITSUI KNOWLEDGE INDUSTRY

408.000x Faster Than Traditional Disk-Based Systems in Technical PoC

216x Faster DNA Analysis
Results - From 2-3 Days to 20 Minutes

CHARITÉ

1.000x Faster Tumor Data Analyzed in Seconds Instead of Hours

2-10 sec For Report Execution

Hvar er HANA notað

Almennur gagnagrunnur fyrir
- Hvaða gögn sem er
- Nýþróun - svo sem genarannsóknir, vefleikir, PFM, bankar, fleira
- Gögn eða kerfi í skýinu

Fyrir SAP kerfi
- BW, BPC, B1, ByD, CRM + **ERP on HANA**
- Þekktur innleiðingakostnaður – Rapid Deployment Solution
  - Profitability Analysis, Spend Analytics, BW migration ....

Nýjar SAP lausnir: Smart Meter Analytics, Shopper Insight....

Applicon

Dæmi - Gögn geymd í dálkum?

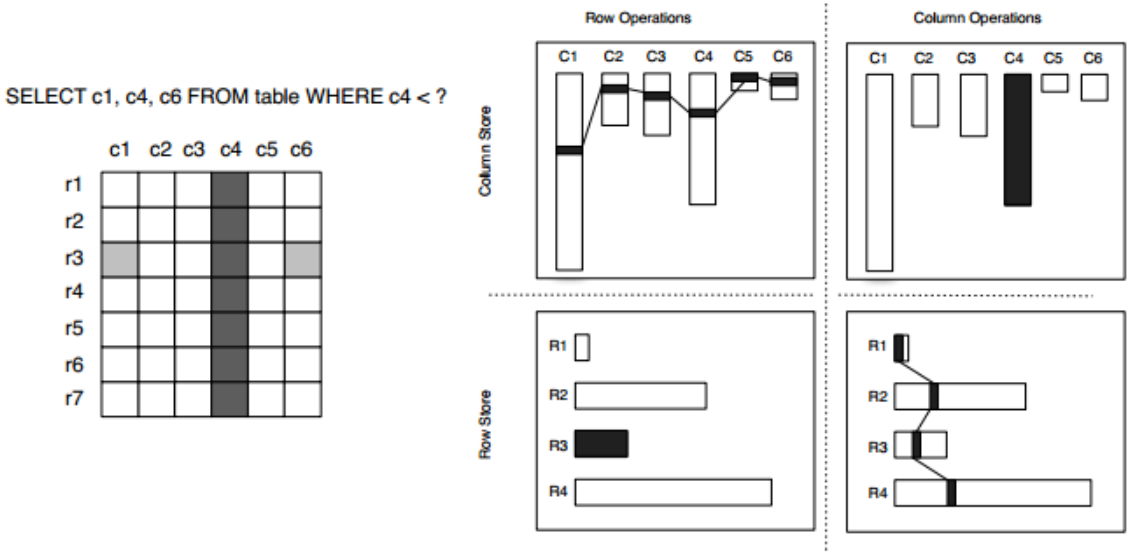

Tafla með starfsmannaupplýsingum geymd í dálkum

- Hvað ef lína á við starfsmann og „dálkur 4“ geymir launatölur?
- Spurningum um heildarlaun er fljótsvarað
- Birting á öllum gildum starfsmanns verður þó hægvirkari

Dæmi – USA eftir póstnúmerum

State	City	Zip Code	Education expenditure	Age	1 – Male 2 – Female	Income	# Wallmart Stores
CA	SF	10042	100	25	1	20000	2
CA	SF	32042	120	41	2	30000	3
MA	CA	23042	110	56	1	40000	0
MA	CA	42322	150	73	1	1000000	2
TX	DA	32032	100	55	2	2000000	1
TX	DA	3423	90	32	1	5600000	0

Viltu prófa? - Dæmi fyrir Chrome með 70 milljón línur

Fleiri dæmi á www.saphana.com

Applicon

Samantekt

In-Memory Data Management for Enterprise Applications

Status Quo

- Systems are **traditionally separated** into **transactional** (OLTP) and **analytical** (OLAP) data management systems
- This separation has many drawbacks
  - OLAP system does not have the latest data and relies on pre-fabricated data
  - Cost-intensive ETL process to sync both systems
  - Fortune 500 companies need a team of 20 people to manage the redundancy

The Vision

- Build an enterprise system **combining OLTP** and **OLAP** in one single database
- This approach enables **real-time analytics** and a **simplified** application and database architecture
- Business questions can be answered in **less than one second**

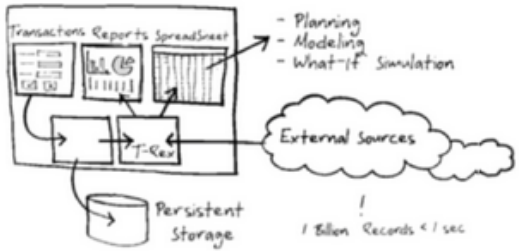

- Recent trends in hardware support this vision
  - Massively **increasing main-memory capacity** at a **lower TCO** (up to 4 TB/server)
  - **Multi-core CPU** architectures (up to 64 cores)

Technology

- **Column orientation** instead of row orientation

Id	Name	Salary
1	Miller	4.000
2	Jones	3.000

row orientation: 1, Miller, 4.000; 2, Jones, 3.000;

column orientation: 1, 2; Miller, Jones; 4.000, 3.000;

- **Compression** of business data saves memory and speeds up processing (up to a factor 10)
- **Partitioning** allows for massive parallelization

Research Results

- **Validated** in-memory column-oriented database technology **with real customer systems**
- Removed secondary indexes, pre-calculated and materialized sum tables, reduced complexity in application
- **Improved** dunning run **from 20 min down to 1 s** (1200-times faster than before)
- Work in progress
  - Parallelizing **planning processes**
  - Augmenting **Available-to-Promise** with real-time analytics and flexible order fulfillment
  - Provisioning **multi-tenant** analytics through a **cloud-based infrastructure**
  - Combining **row** and **column** storage in a **hybrid** system

Research Group of Prof. Plattner/Dr. Zeier: Enterprise Platform and Integration Concepts (EPIC)
Anja Bog, Martin Grund, Jens Krüger, Peng Liu, Stephan Müller, Jan Schaffner, Christian Tinnefeld
hana-epic@hpi.uni-potsdam.de

Hasso Plattner Institut
IT Systems Engineering | Universität Potsdam