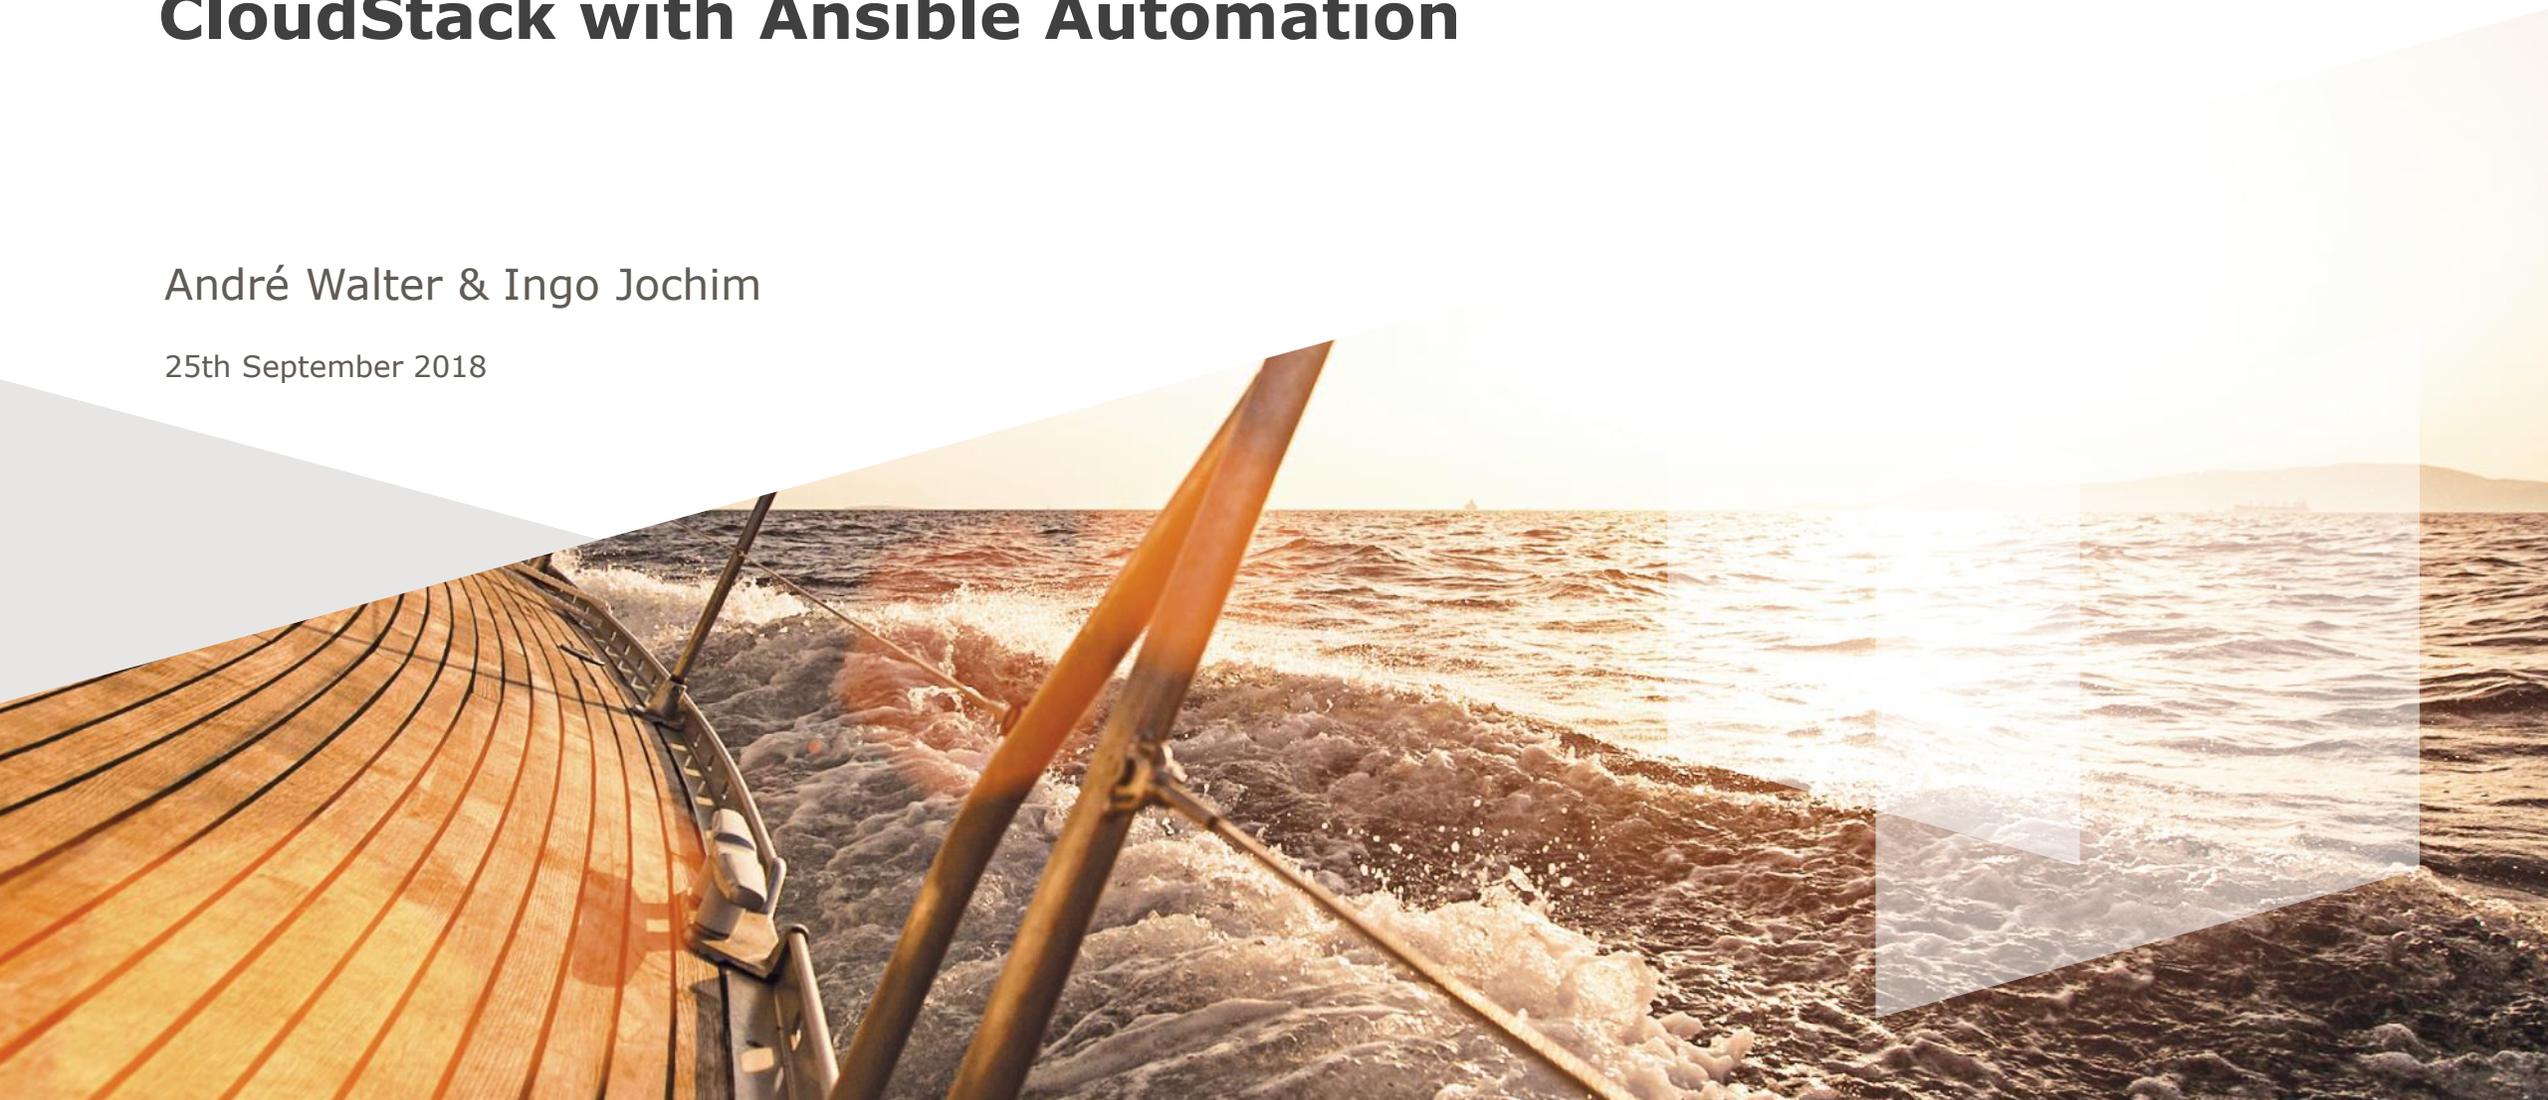# CloudStack with Ansible Automation

André Walter & Ingo Jochim

25th September 2018

Agenda

1. Talk André
2. Talk Ingo

About André

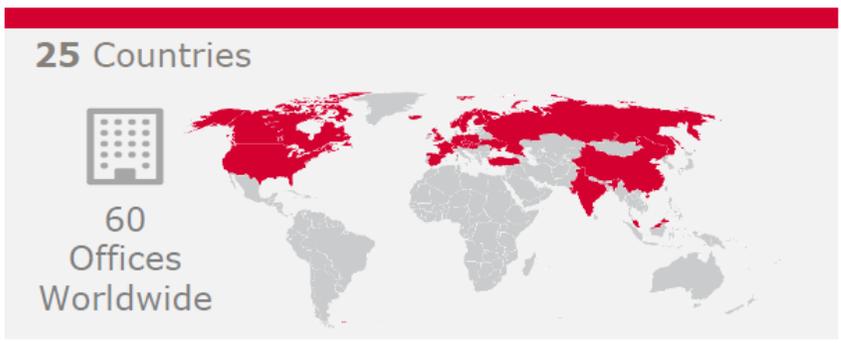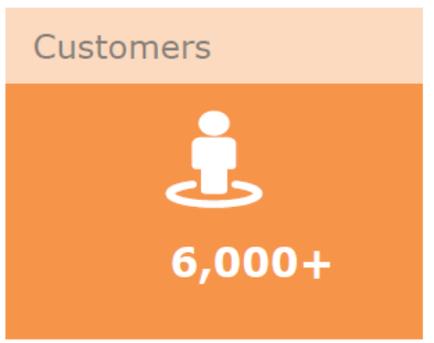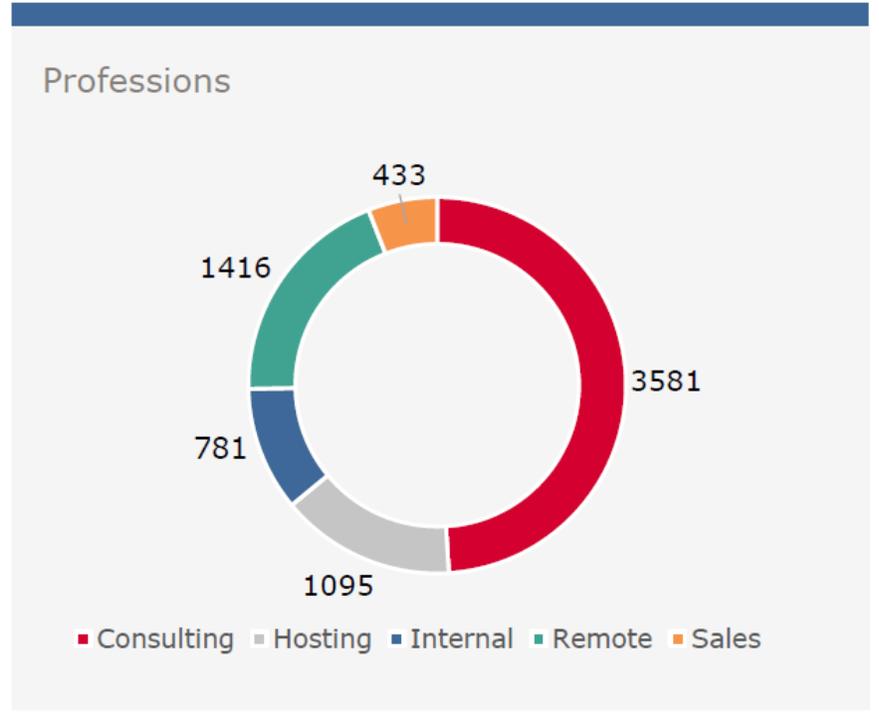- Head of itelligence GMS Cloud Infrastructure Services
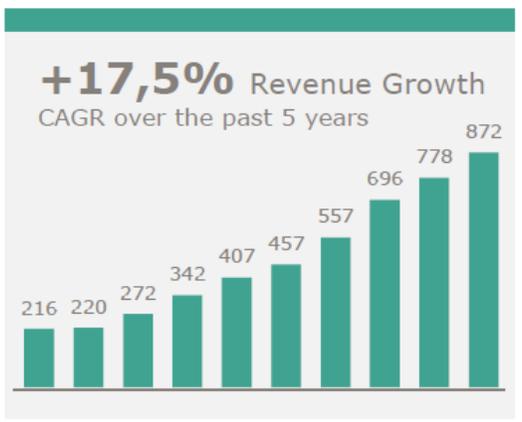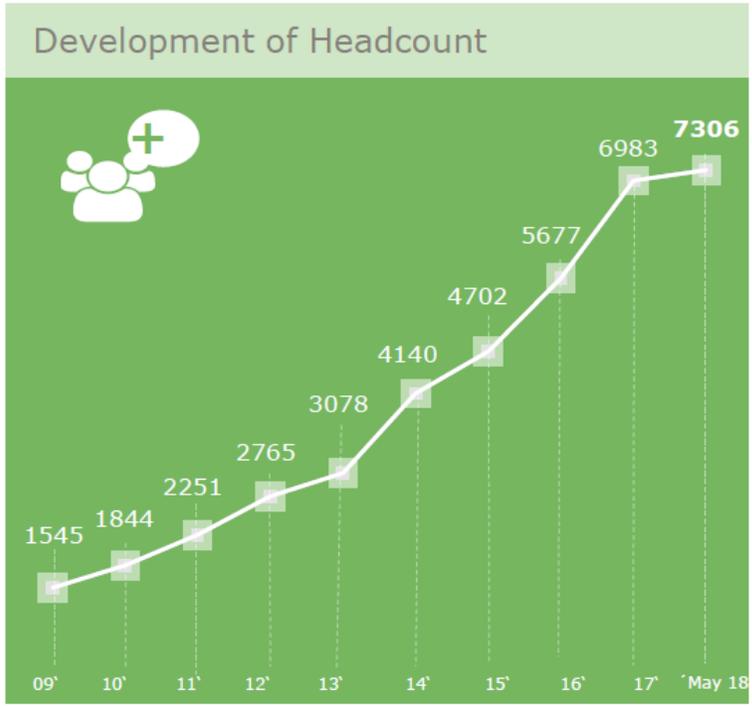
- 　　　@meradioc

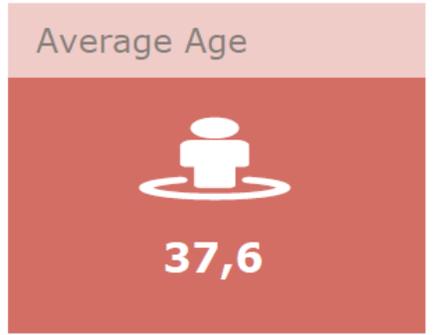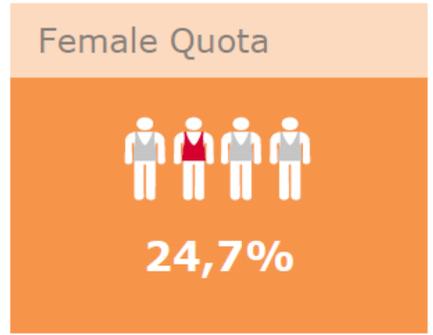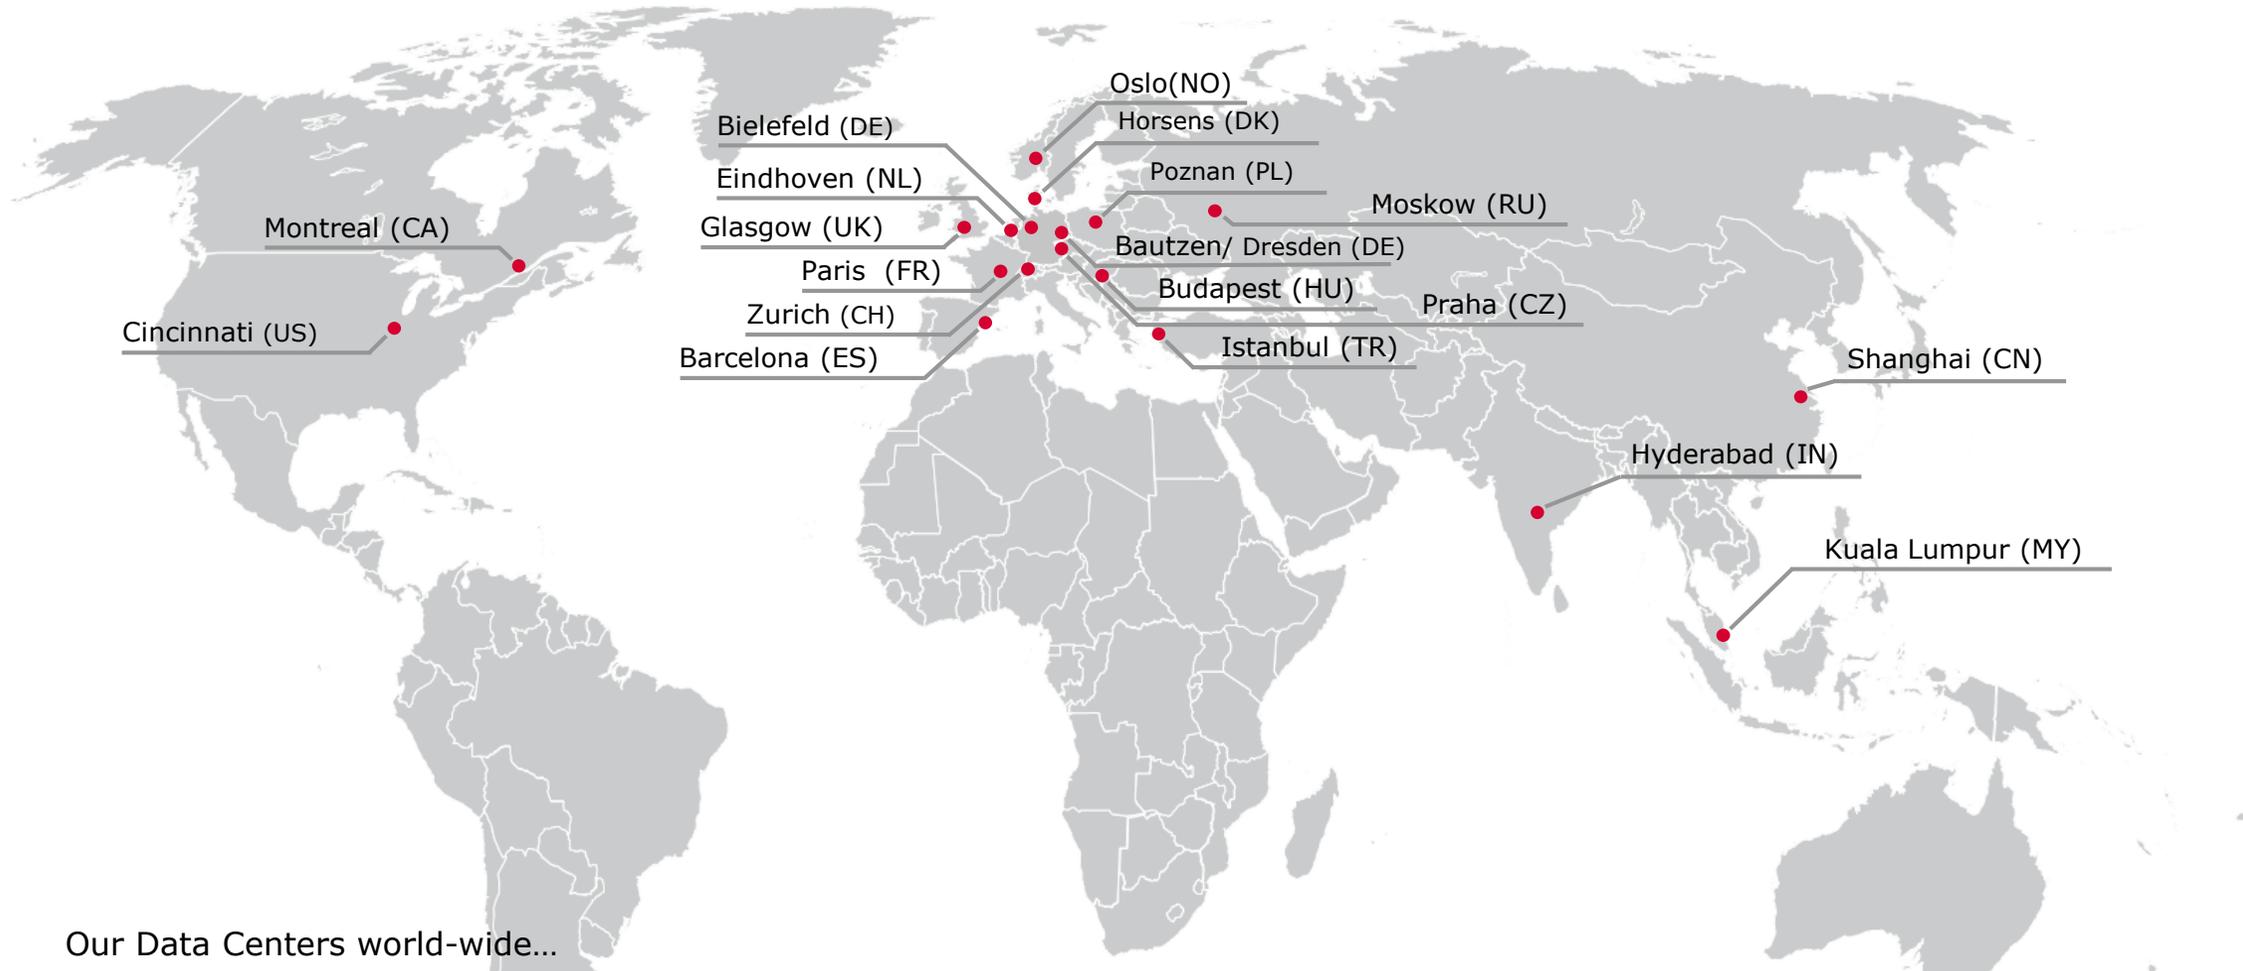- 　　　https://www.linkedin.com/in/andre-walter-87215a35/

Agenda - André

1. About itelligence and our focus

2. Typical use case (SAP HANA)

3. Resulting requirements

4. Why Apache CloudStack

itelligence Worldwide - Numbers

Female Quota

24,7%

Average Age

37,6

Development of Headcount

1545
1844
2251
2765
3078
4140
4702
5677
6983
7306

09' 10' 11' 12' 13' 14' 15' 16' 17' ´May 18

Employees

7.300+

+17,5% Revenue Growth
CAGR over the past 5 years

216 220 272 342 407 457 557 696 778 872

Professions

433
1416
3581
781
1095

■ Consulting ▪ Hosting ■ Internal ▪ Remote ▪ Sales

Customers

6,000+

25 Countries

60 Offices Worldwide

itelligence Managed Services – Service Centers Worldwide

Oslo(NO)
Bielefeld (DE)
Horsens (DK)
Eindhoven (NL)
Poznan (PL)
Montreal (CA)
Glasgow (UK)
Moskow (RU)
Paris (FR)
Bautzen/ Dresden (DE)
Cincinnati (US)
Zurich (CH)
Budapest (HU)
Praha (CZ)
Barcelona (ES)
Istanbul (TR)
Shanghai (CN)
Hyderabad (IN)
Kuala Lumpur (MY)

Our Data Centers world-wide…

Bielefeld

Kuala Lumpur

Bautzen

Poznan

Cincinnati

GMS Cloud Strategy | Change of Customer Application Landscape

„We manage your (SAP) Cloud!"

- As many customers are planning to use public or hybrid cloud services, combination of cloud services is key for success

- Use of cloud infrastructure capabilities of NTT Group

- itelligence partnership with **aws** as the leading Hyperscaler

- High quality Technical and Application Management Services are foundation for smart Cloud Management
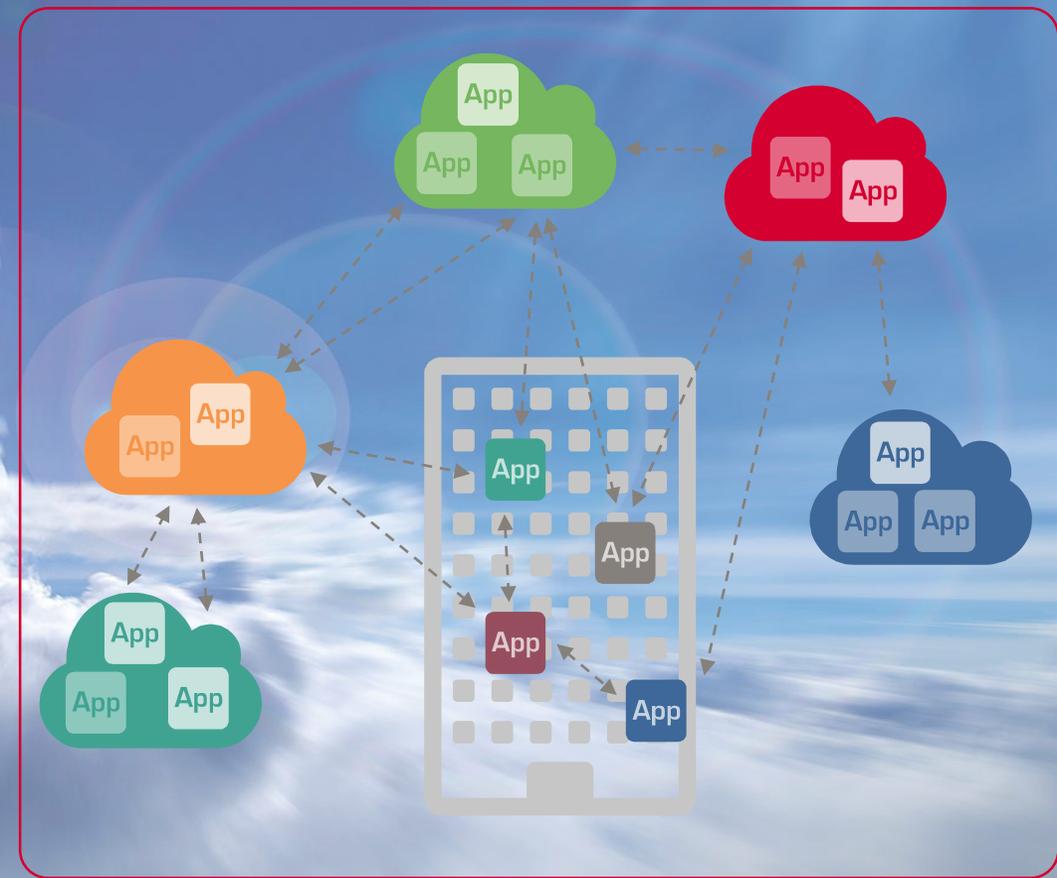
Use Case - SAP Hana

The SAP HANA database is a hybrid, relational in-memory database that combines row-based, column-based, and object-based database technologies and has been optimized for the use of parallel processing functionality of modern multi-core and CPU architectures.
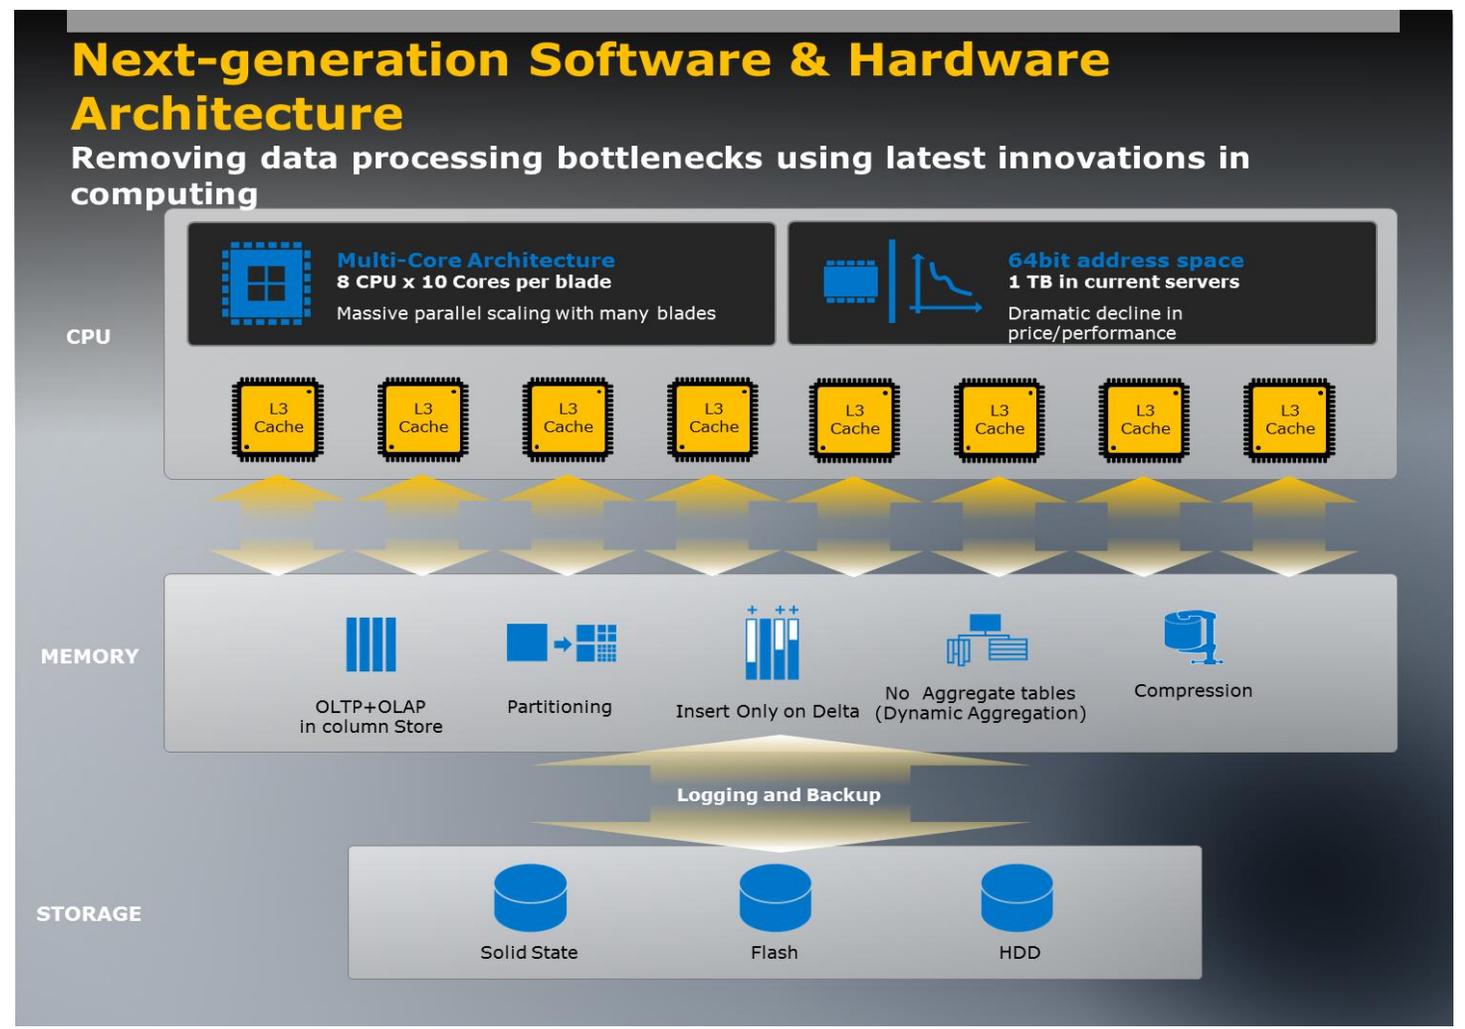
This architecture enables SAP applications to benefit from current hardware technologies.

The SAP HANA database is a central component of SAP products that help customers make their processes more efficient, flexible, and profitable.

Resulting Requirements

- Often physical servers.
- No containers.
- All data are in memory.
- Machines with lot's of RAM.
- Huge workloads.
- Complex backup and disaster recovery scenarios.

Next-generation Software & Hardware Architecture — Removing data processing bottlenecks using latest innovations in computing

Resulting Requirements - Private Cloud

- Complete control over tools and technology.

- Full data control.

- SLAs important.

- Special use cases.
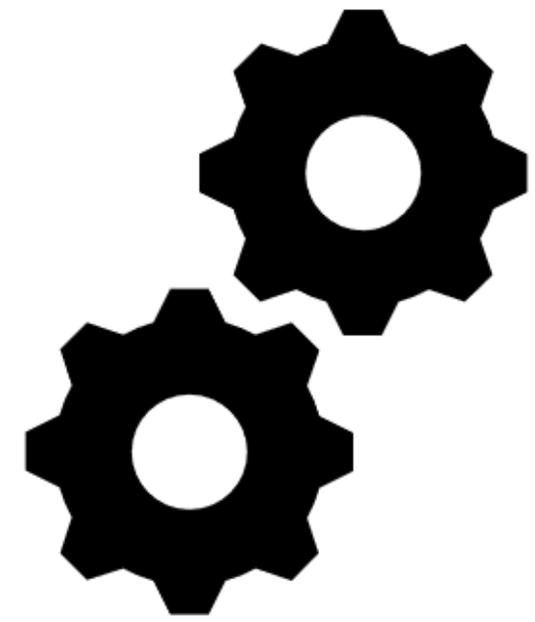
What we do with our cloud?

- Mostly running SAP systems / IaaS.

- Production / Training / Demo / Development.
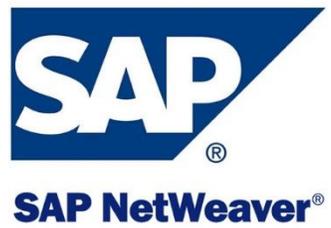
Why Apache CloudStack?

- Our requirements:
  - Multi tenancy.
  - OpenSource.
  - Powerful API.
  - Easy to use.
  - Frequent updates.
  - Good community.
  - Multi hypervisor support.

#cloudstackworks – Our Experience, What We Like?

- Made good experience because #cloudstackworks.

- Easy to upgrade.

- Stable, simple, easy to maintain.

- Great community.

- Free, OpenSource.

We ♡

Questions?

About Ingo

- Head of CIS - Cloud Implementation

- @ingojochim

- https://www.linkedin.com/in/ingo-jochim-b30b8590/

- CloudStack user since 2014

Agenda - Ingo

1. Big picture of our cloud.

2. CloudStack THE central component.

3. Customer portal.

4. Portal - How it works?

5. Ansible daemon.

6. Response in user interface.

7. Dynamical user interface.

8. Automation host - network setup.

9. Scalable setup.

10. Application installation automation.

Big Picture of Our Cloud
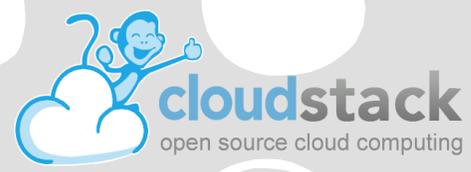
Talk Reference

- **CloudStack + Ceph**   (Sebastian Bretschneider - Monday, 24th Sep, 15:40)

  Abstract:

  This is about our way from NFS to CEPH as primary Storage in CloudStack. Disadvantages about NFS as primary Storage - recommendations for CEPH - what can go wrong. Stories of a Sysadmin.

  https://apachecon.dukecon.org/acna/2018/#/scheduledEvent/d3b061c4a120801e8

- **Billing with CloudStack**   (Alexander Stock - Monday, 24th Sep, 16:40)

  Abstract:

  In this talk you will see which tools we use to generate billing reports for cloudstack and what were the difficulties which we discovered during the implementation.

  https://apachecon.dukecon.org/acna/2018/#/scheduledEvent/dcf32b078a6d4f21d

CloudStack The central component

- Used via API.

- We use it for:

  - Compute.

  - Block storage.

  - Network.

  - Usage recording.

  - Tagging.
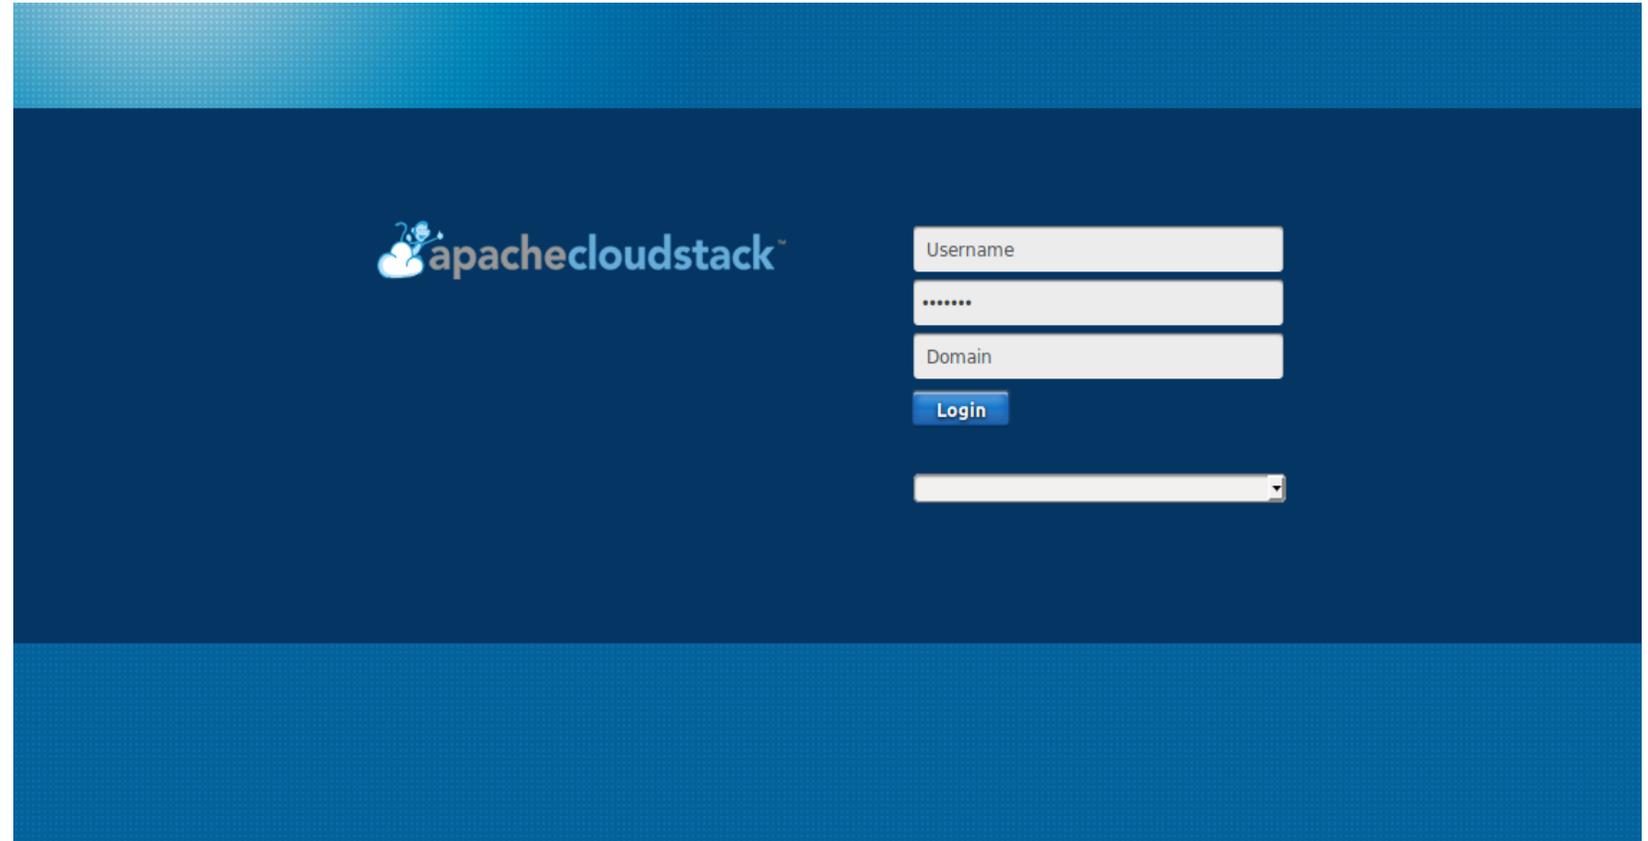
Customer Portal

- Consolidate the varying systems with their individual design and functions.

- Facilitate process automation.

Portal - How it works?

read operations - synchron

write operations - asynchron

Ansible Daemon

Queue

- Self written application.

- Functionality:

  - Queueing jobs.

  - Parallel execution.

  Parallel execution

  - Status response of jobs.

  - Reading automation playbooks (header information).
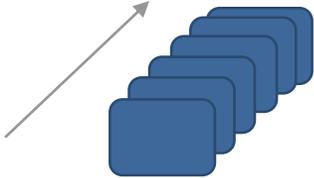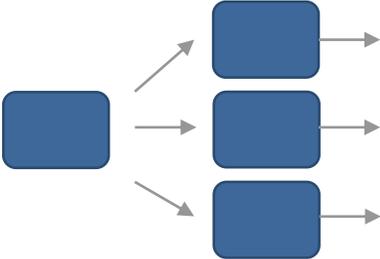
Playbook header

Response in User Interface – Asynchron Jobs

Dynamical User Interface by Playbooks

- Defined in playbook header

  - Category

  - Available playbooks

  - Required values

  - Dependencies

Automation Host and Network Setup

Automation host

Customer network

Customer machines
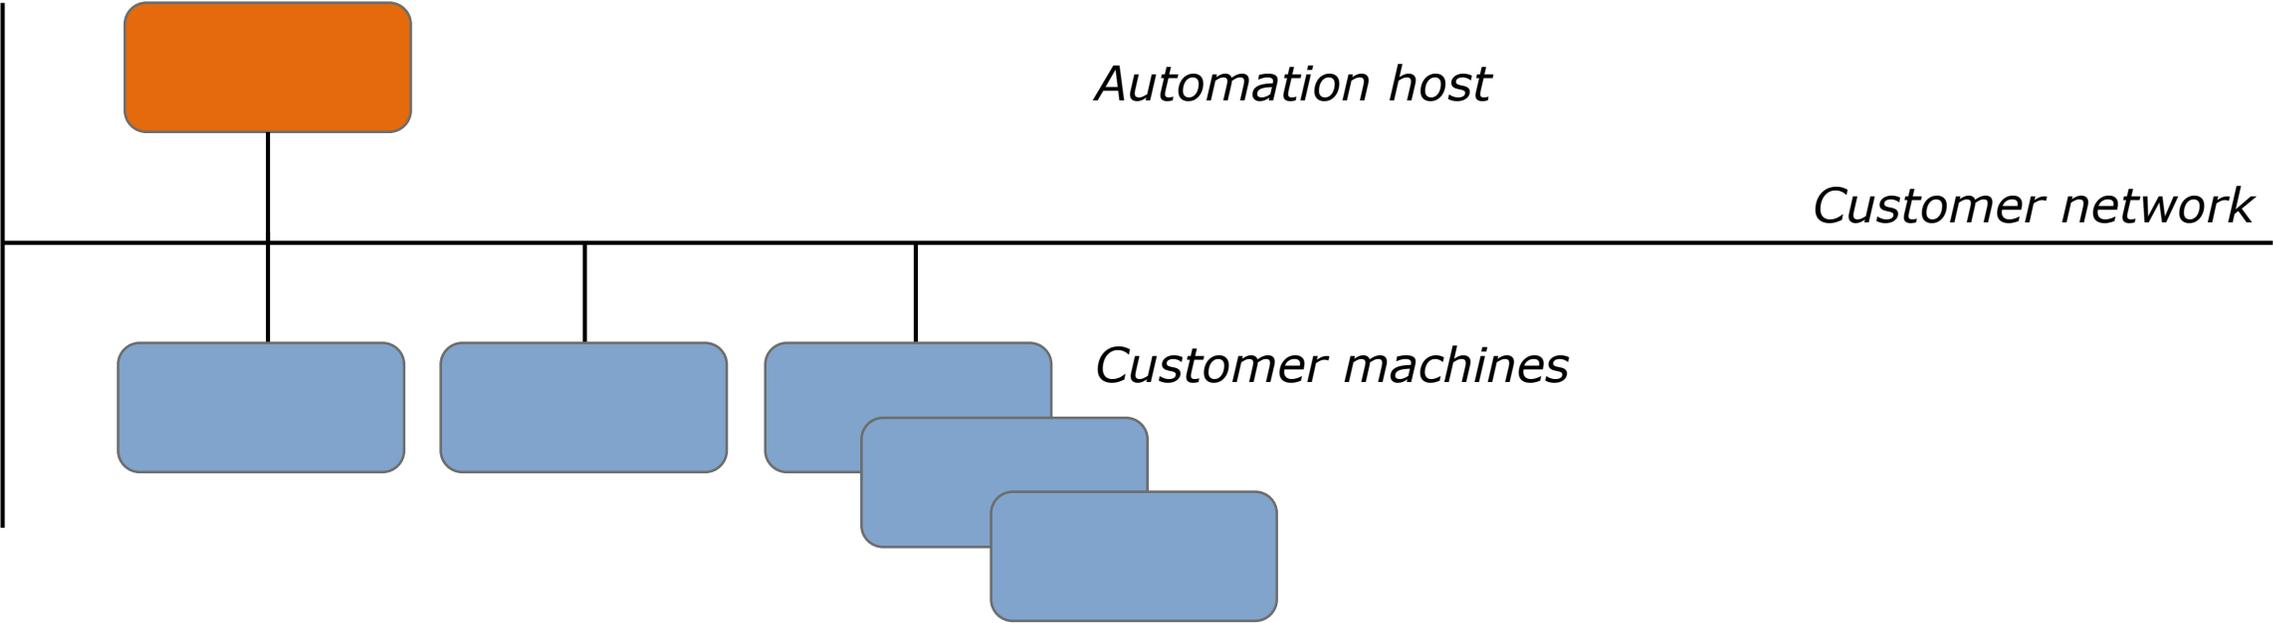
Scalable Setup
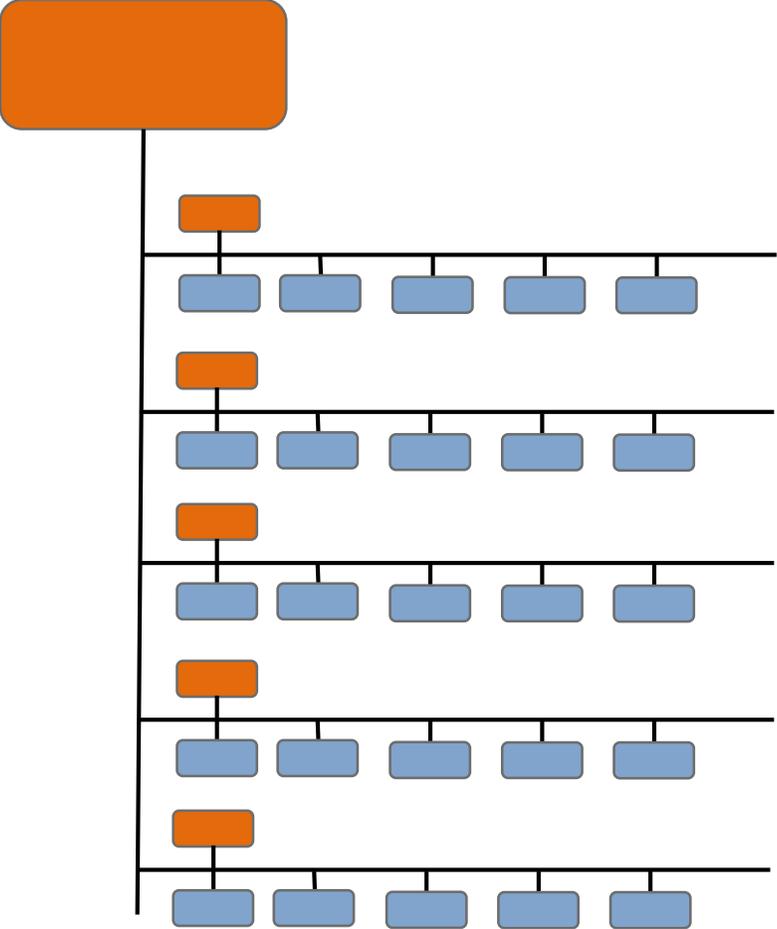
portal, LDAP, monitoring

scalelable: ansible, LDAP, monitoring

Application Installation Automation

Complete automation in a single step

Compute   Storage   Application   Firewall   Monitoring   Backup
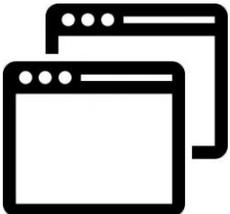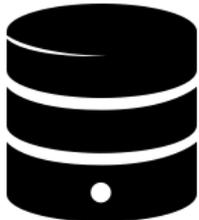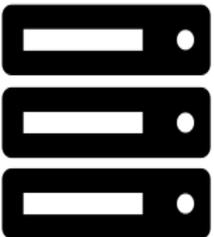
Questions?

Copyright itelligence AG - All rights reserved