

Make IT

2022

Moderna infrastruktura za Oracle

Speakers:

- **Urh Srečnik** <urh.srecnik@abakus.si>
- **Sergej Rožman** <sergej.rozman@abakus.si>
- **Boris Oblak** <boris.oblak@abakus.si>



ORACLE®

Gold
Partner

Abakus Plus d.o.o.

- Infrastructure Team

- Services

- OS & NET admin
- DBA, Programming

- Applications

- **APPM**
- **Backup Server & Deja Vu**
- Arbiter

- Development Team

- Enterprise Applications

- Document Management

- Newspaper Distribution

- Flight Information System

Customers

Gorenjska Banka

GENERALI
Zavarovalnica

Ljubljana Airport

EKDIS
Ekspresno. Ekonomično.

REPUBLIKA SLOVENIJA
MINISTRSTVO ZA OBRAMBO

NOVA
BANKA

MILENIJUM
OSIGURANJE

KONTROLA
ZRAČNEGA
PROMETA
SLOVENIJE

Iskra

Hotria

Mestna občina
Ljubljana

LON

triglav

ANDRITZ

jata emona
LJUBLJANA

UNIVERZITETNA PSIHIATRIČNA
KLINIKA LJUBLJANA
University Psychiatric Clinic Ljubljana

skbbanka
otp group

SAVARe

MERKUR

TRELLEBORG

SODO
SISTEMSKI OPERATER
DISTRIBUCIJSKEGA OMREŽJA Z
ELEKTRIČNO ENERGIJO

BANKA
SLOVENIJE

PRVA

MAGNETIK d.o.o.
TSS PEST MANAGEMENT SOLUTIONS

Trelleborg Slovenija, d.o.o.

NLB Vita
Življenjska zavarovalnica

Mercator

GOODYEAR DUNLOP
SAVA TIRES

MM
KARTON

studio ritem

Blubit
TikO
TOVARNA KOVINSKE OPREME

ZAVOD ZA
ŠPORT RS
PLANICA

PH
Primorska
hranilnica

CENTROSINERGIJA

PANTEON
GROUP

Lonia

PRONET
CHOOSE THE FUTURE

hit alpinea
Kranjska Gora

SAVA
HOTELS & RESORTS

LASERLINE

ORACLE

ROS d.o.o.

NFOTRANS

PARK
POSTOJSKA
JAMA



ADRIA ANKARAN
HOTEL & RESORT

OLVM

ORACLE Linux Virtualization Manager

Dashboard

Compute

Network

Storage

Administration

Events

Last Updated 5/4/2022, 2:05:19 PM GMT+2

1 Data Centers 1	2 Clusters N/A	3 Hosts 3	1 Data Storage Domains 1	8 Virtual Machines 8	16 Events 1 1 14
---------------------	-------------------	--------------	-----------------------------	-------------------------	---------------------

Global Utilization

CPU	Memory	Storage
99% Available of 100% Virtual resources - Committed: 41%, Allocated: 41%	606.1 Available of 754.3 GiB Virtual resources - Committed: 21%, Allocated: 21%	28.0 Available of 38 GiB Virtual resources - Committed: 0%, Allocated: 0%
1% Used	148.3 GiB Used	10.0 GiB Used

Cluster Utilization

CPU	Memory


Storage Utilization

Storage

Legend: > 90% (Red), 75-90% (Orange), 65-75% (Yellow), < 65% (Light Blue)

oVirt

upstream, community supported



Red Hat
Virtualization

ORACLE

Linux

Virtualization Manager

more stable with commercial support

oVirt OPEN VIRTUALIZATION MANAGER

Dashboard | Last Updated 9/19/2017, 4:11:43 PM GMT+3

1 Data Centers 1	1 Clusters N/A	2 Hosts 2	3 Data Storage Domains 3	3 Virtual Machines 2 1	18 Events 2 7 9
---------------------	-------------------	--------------	-----------------------------	-----------------------------	------------------------

Global Utilization

CPU 98% Available of 100% Over commit: 25% (allocated 75%)	Memory 2.2 Available of 3.6 GiB Over commit: 14% (allocated 35%)	Storage 0.2 Available of 0.2 TiB Over commit: 1% (allocated 2%)
---	---	--

ORACLE Linux Virtualization Manager

Dashboard | Last Updated 5/4/2022, 2:05:19 PM GMT+2

1 Data Centers 1	2 Clusters N/A	3 Hosts 3	1 Data Storage Domains 1	8 Virtual Machines 8	16 Events 1 1 14
---------------------	-------------------	--------------	-----------------------------	-------------------------	-------------------------

Global Utilization

CPU 99% Available of 100% Virtual resources - Committed: 41%, Allocated: 41%	Memory 606.1 Available of 754.3 GiB Virtual resources - Committed: 21%, Allocated: 21%	Storage 28.0 Available of 38 GiB Virtual resources - Committed: 0%, Allocated: 0%
---	---	--

RED HAT VIRTUALIZATION

Dashboard | Last Updated 11/24/2017, 9:31:32 AM GMT+10

1 Data Centers 1	1 Clusters N/A	1 Hosts 1	1 Data Storage Domains 1	1 Virtual Machines 1	8 Events 8
---------------------	-------------------	--------------	-----------------------------	-------------------------	---------------

Global Utilization

CPU 99% Available of 100% Over commit: 0% (allocated 25%)	Memory 11.0 Available of 11.6 GiB Over commit: 0% (allocated 9%)	Storage 0.1 Available of 0.2 TiB Over commit: 0% (allocated 5%)
--	---	--

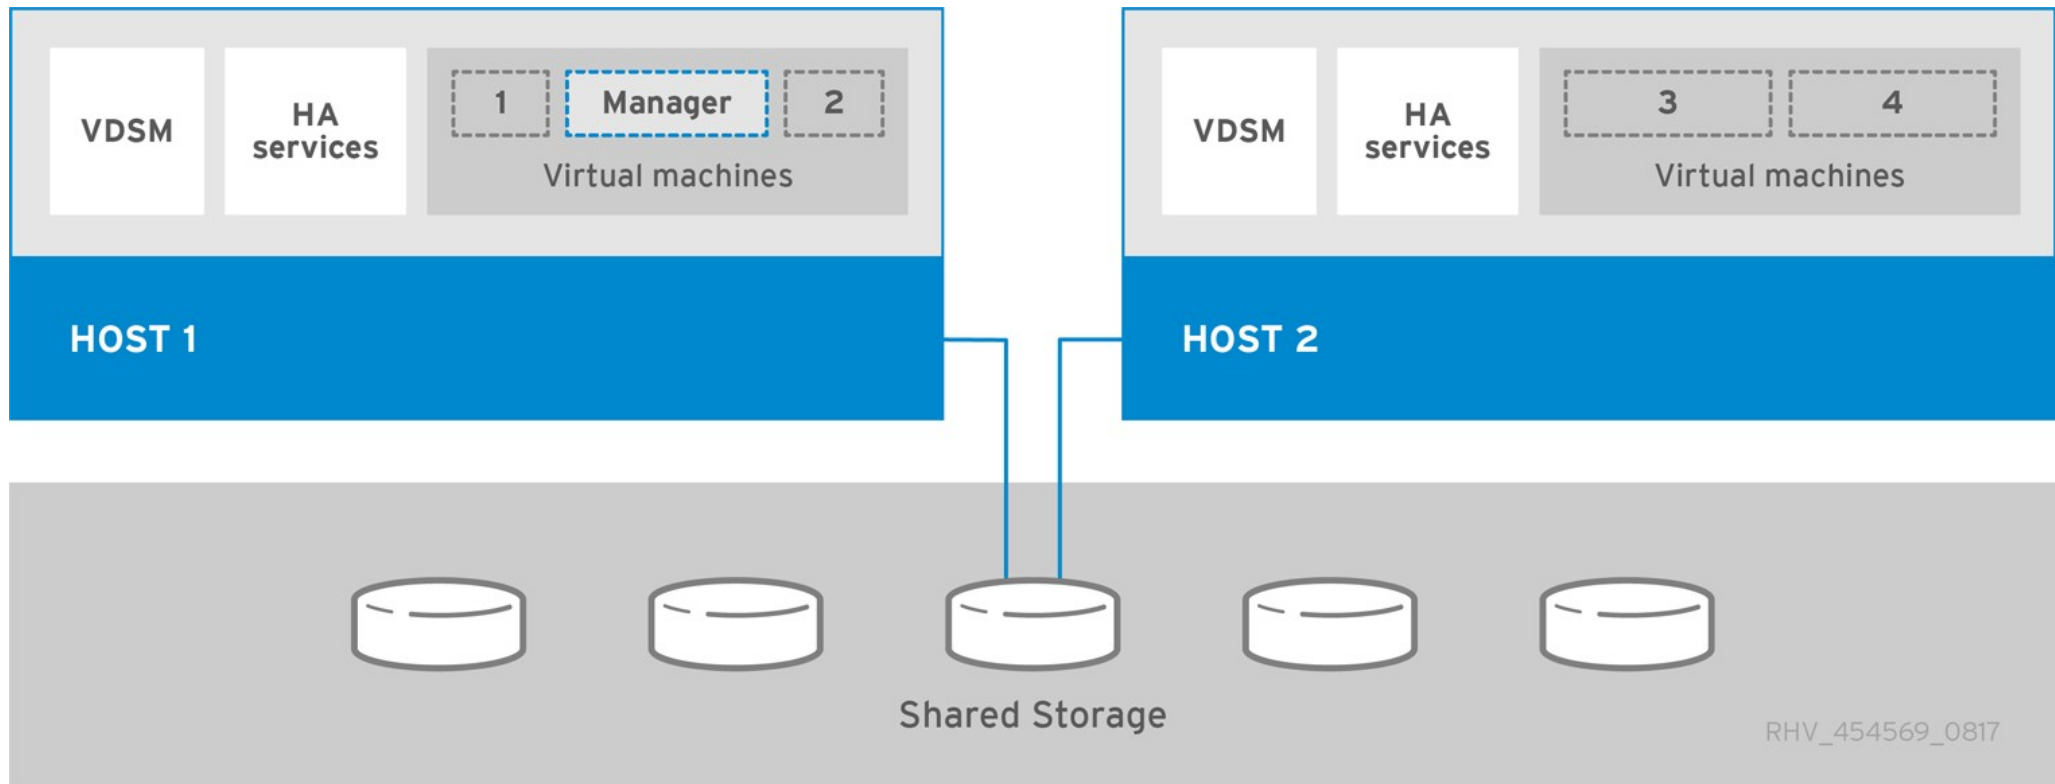


image from:

https://www.ovirt.org/documentation/installing_ovirt_as_a_self-hosted_engine_using_the_command_line/

Built On

- **KVM** - <https://www.linux-kvm.org/>
Linux kernel module capable of full virtualization support for Linux on x86 hardware containing virtualization extensions (Intel VT or AMD-V)
- **QEMU** - <https://www.qemu.org/>
a generic and open source machine emulator and virtualizer
- **libvirt** - <https://libvirt.org/>
virtualization API and toolkit to manage virtualization platforms

OLVM: Inkrementalni backupi

Speaker:

Urh Srečnik

Podjetje

Abakus plus d.o.o.

Petek, 3. junij

 12:35 - 13:20

Težavnost: Basic

Dvorana: James Cook

Področje: Database

OLVM Licensing


- Free to use/install
- ***Optional*** commercial support by Oracle
- Approved hard partitioning technologies by Oracle:
 - OLVM
 - OVM (end of life, »Extended Support« still available)
 - Oracle Solaris Zones
 - Source:
<https://www.oracle.com/assets/partitioning-070609.pdf>

Other Virtualizations?

- Oracle Database licences requires you to license all the *physical* CPUs in all the nodes in the vm cluster.
- It *may* be impossible to use OLVM hard partitioning for Standard Edition Databases anyway
 - Because each chip in multi-chip processor is considered as a socket

Oracle License Management Services (LMS) is the established authority on Oracle licensing policy. Licensing information in this presentation is for educational purposes only and may even be incorrect.

VE Evolution





Modern Infrastructure for Oracle (and beyond)

mag. Sergej Rožman

sergej.rozman@abakus.si





KVM – Kernel-based Virtual Machine



- since 2006 by Qumranet, acquired by Red Hat in 2008
- clouds based on KVM: Openstack, IBM Cloud, Rackspace, Google, ...
- virtualizations based on KVM: Red Hat Enterprise Virtualization, Oracle Virtualization, Proxmox VE, oVirt, ...





KVM – Kernel-based Virtual Machine

Run a VM

```
# qemu-system-x86_64 -enable-kvm -hda /tmp/test.raw  
VNC server running on 127.0.0.1:5900
```

- initial parameters of VM are specified in a command line
- default graphical console is at VNC localhost:5900 (or specified in command line)
- VM management goes through UNIX socket (JSON format) or QEMU/Monitor (Ctrl-Alt-2)

Manage a VM – example: eject CD/DVD

```
# find and connect to right socket
```

```
{ "execute": "eject", "arguments": { "device": "ide1-cd0" } }  
{"return": {}}
```





libvirt



- since 2005 by Red Hat
<https://libvirt.org/>
- is a toolkit to manage virtualization platforms
- XML configuration

Manage a VM – example: eject CD/DVD

```
# virsh change-media VMname hdc --eject  
succeeded to complete action eject on media
```





GUI Management

- 44 management utilities
according to https://www.linux-kvm.org/page/Management_Tools
- Oracle Linux Virtualization Manager is not one of them, oVirt is





Red Hat Virtual Machine Manager



- since 2008 by Red Hat
<https://virt-manager.org/>
- Red Hat Virtual Machine Manager, also known as virt-manager, is a **desktop** virtual machine manager/monitor
- just tool, not an environment





virt-manager

File Virtual Machine View Send Key

Overview OS information Performance CPUs Memory Boot Options VirtIO Disk 1 NIC :b9:6c:87 NIC :76:e9:f3 Tablet Mouse Keyboard Display Spice Serial 1 Channel qemu-ga Channel spice Video Virtio Controller USB 0 Controller SATA 0 Controller PCIe 0 Controller VirtIO Serial 0 USB Redirector 1 USB Redirector 2 RNG /dev/urandom

+ Add Hardware

Details XML

CPU usage

55 %

Memory usage

1522 MiB of 4096 MiB

Disk I/O

0 KiB/s read 20 KiB/s write

Network I/O

1 KiB/s in 0 KiB/s out

Prekliči Uveljavi



Proxmox VE



- since 2008 by Proxmox Server Solutions GmbH, Austria
<https://www.proxmox.com/en/proxmox-ve>
- web-based graphical interface management
- simple and intuitive use
- robust, not easy to break
- **well integrated with all modern hardware and storage systems**
- can be all-in-one solution



Proxmox VE



Browser address bar: <https://kvm1:8006/#v1:0=qemu%2F10083:4:.....>

Page title: PROXMOX Virtual Environment 7.1-8

Server View: Virtual Machine 10083 (artemis) on node 'kvm1'

Summary for VM 10083 (artemis):

- Status: running
- HA State: started, Group: none
- Node: kvm1
- CPU usage: 13.51% of 8 CPU(s)
- Memory usage: 96.45% (15.43 GiB of 16.00 GiB)
- Bootdisk size: 180.00 GiB
- IPs: No Guest Agent configured

Notes: tukaj gor teče IPA in APPM. Pa tudi abakus docker repository, čeprav mislim, da se ta ne uporablja več.

CPU usage (Hourly average):

Memory usage (Total and RAM usage):

Network traffic (netin and netout):

Disk IO (diskread and diskwrite):

Tasks:

Start Time	End Time	Node	User name	Description	Status
May 25 06:22:03		kvm1	root@pam	VM/CT 50108 - Console	
May 25 02:47:44	May 25 02:47:45	kvm2	root@pam	VM 99007 - Stop	OK

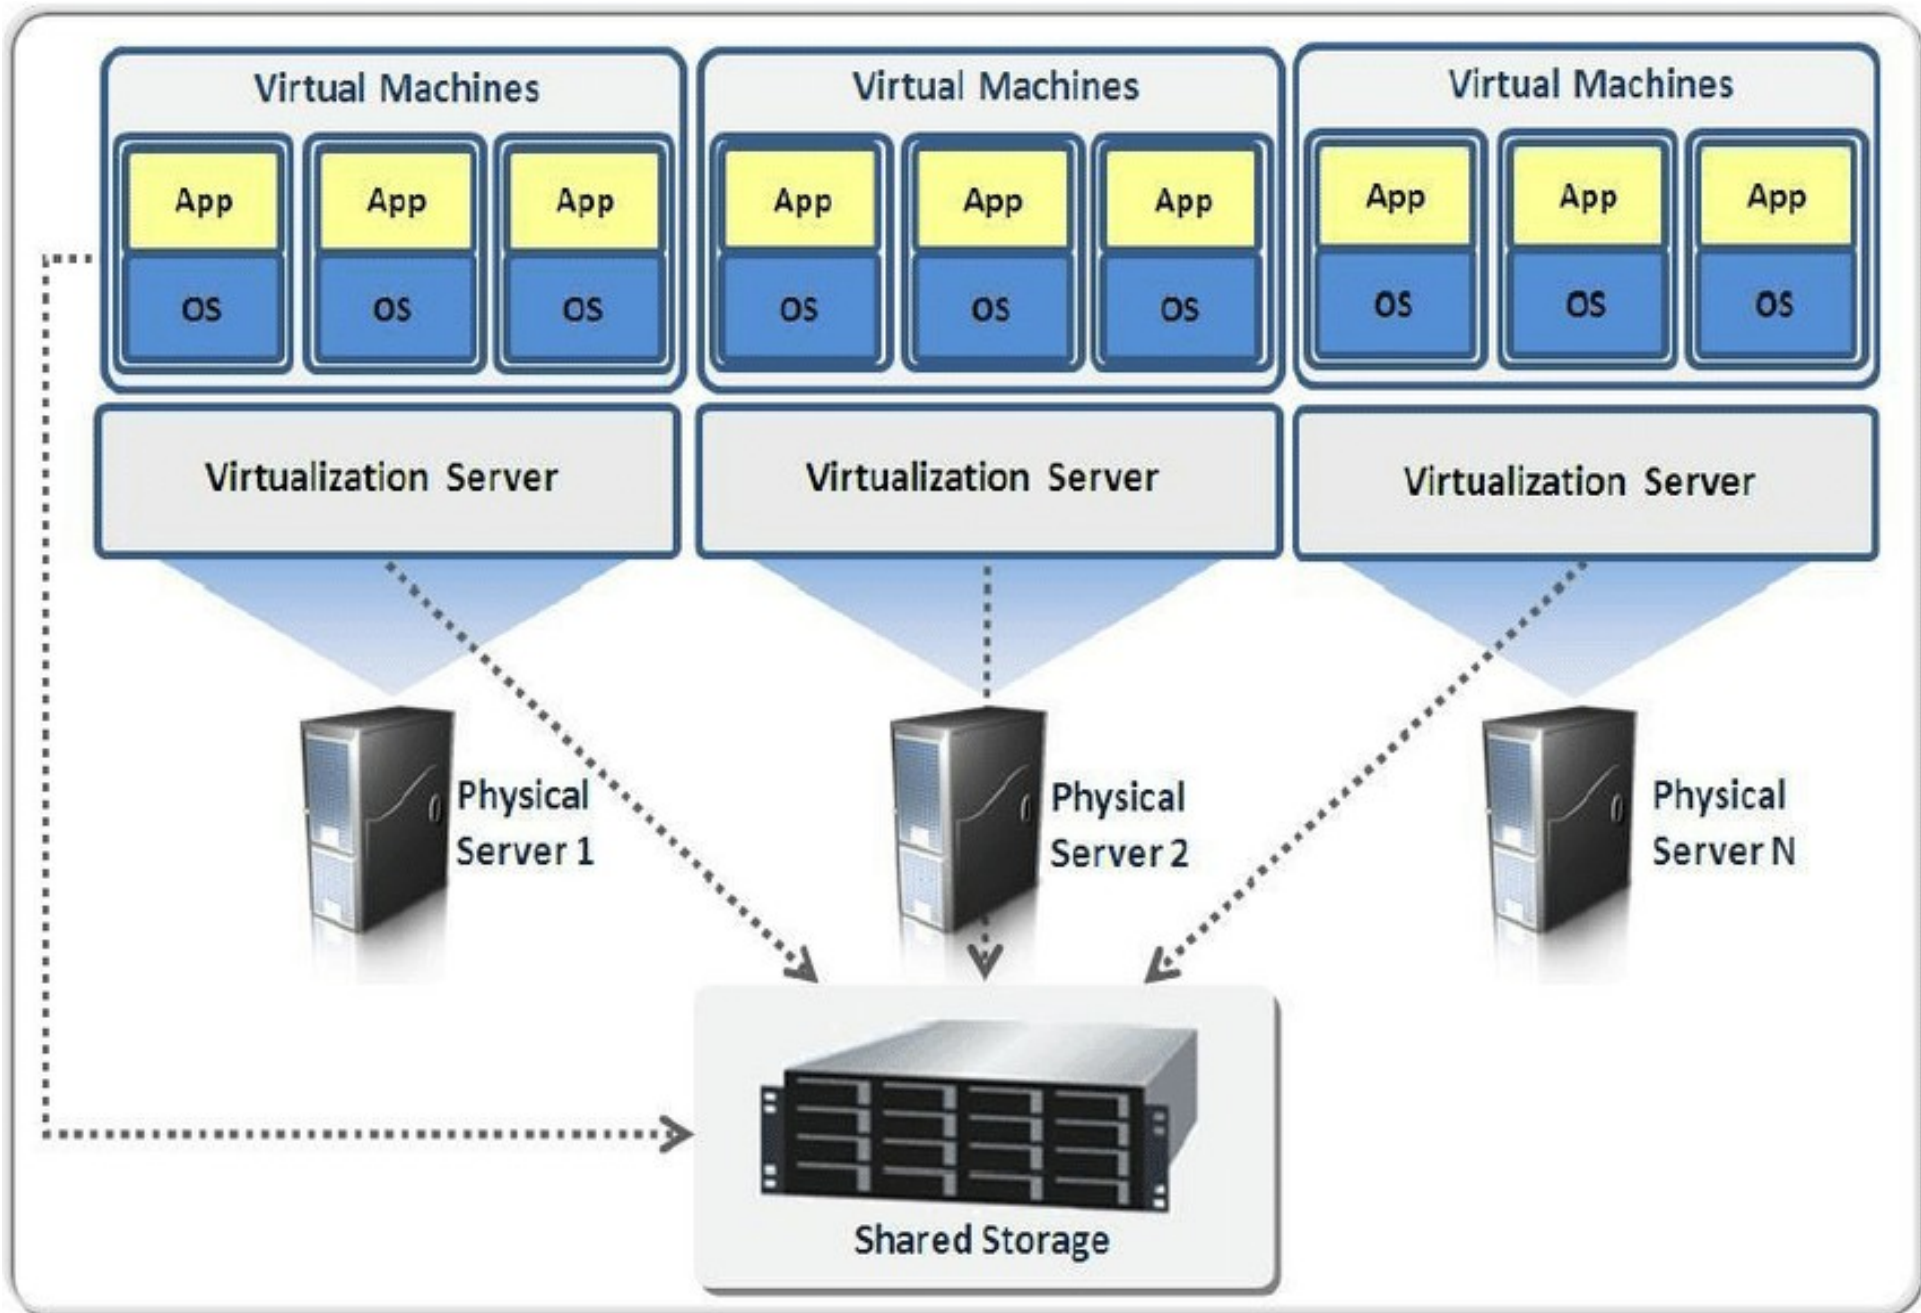


Modern Infrastructure

Boris will now tell us about the experience with the Oracle database in virtual and non-virtual environments.

Thank you





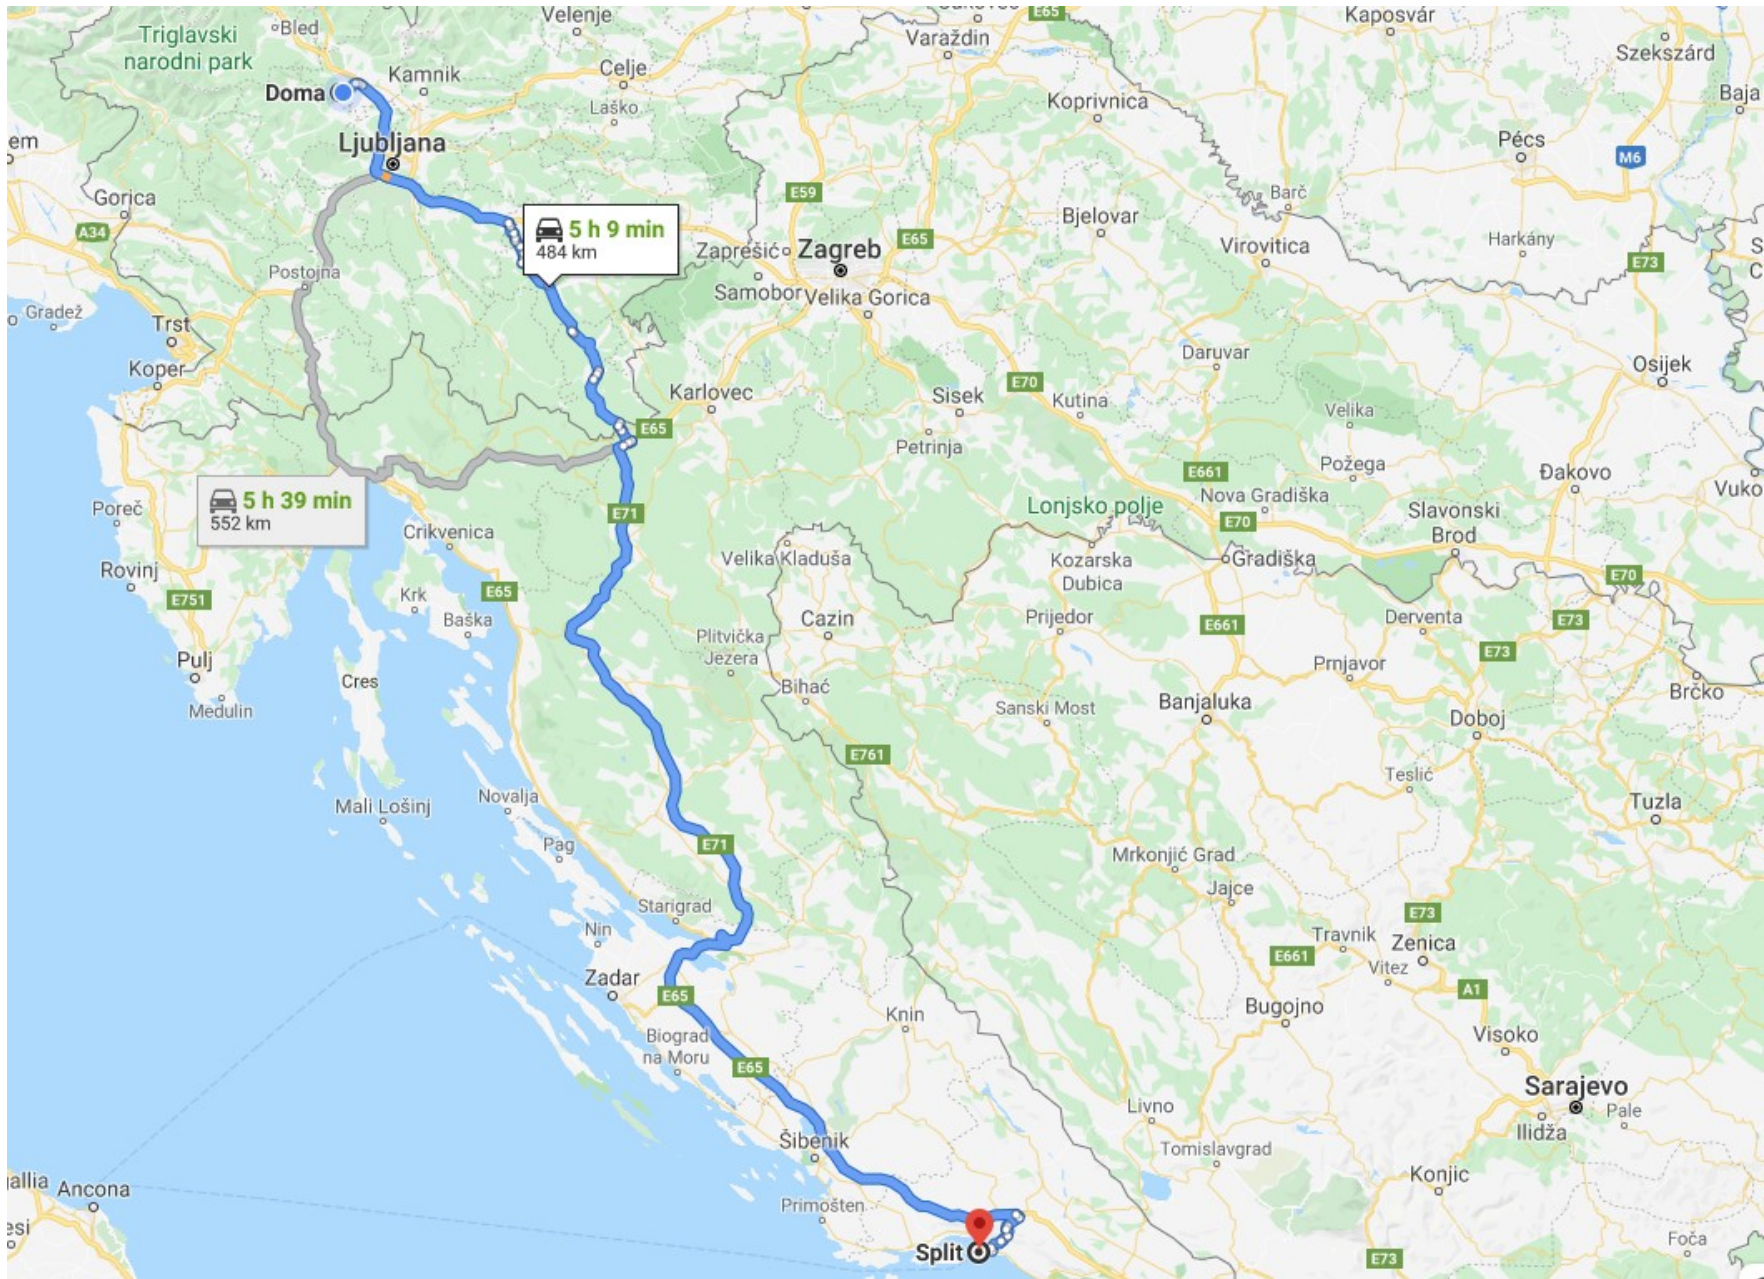
Make IT

2022

How is database performance
doing today?

Distance and time

- I drove my car to Split for 4 hours.
- How fast did it take?



Distance and time

- I drove my car to Split for 4 hours.
- How fast did it take?
- Google maps: 484 km
- Put this together:
 - $484 \text{ km} / 4 \text{ h} = 121 \text{ km/h}$

Test case

- I drove my car to Split and it takes 6 hours.
- $484 / 6 = 80,67 \text{ km/h}$
- $80,67 < 121$:
 - it takes »too long«;
 - something went wrong;
 - don't know what,
 - ... but it indicates I need something to do

Test case

- I drove my car to Split and it takes 6 hours.
- $484 / 6 = 80,67 \text{ km/h}$
- $80,67 < 121$:
 - it takes »too long«;
 - something went wrong;
 - don't know what,
 - ... but it indicates I need something to do:
 - new car?
 - new road?
 - new route?

Indicator km/h

- The indicator does not tell what is wrong or what to do.
- It tells us about relationship between distance and time.
- We may want to investigate further.

Work and Time



1 worker 12 wheelbarrows in 60 min

Work and Time

- 1 worker (workload).
- 12 wheelbarrows (work_done).
 - **unit of work**: wheelbarrow.
- 60 min (time).
- $\text{Work_time} = \text{workload} * \text{time} = 1 * 60 = 60 \text{ min.}$
- $\text{Response Time} = \text{work_time} / \text{work_done} = 60 \text{ min} / 12 = 5 \text{ min} / \text{wheelbarrow.}$

Work and Time



2 workers 24 wheelbarrows in 60 min

Work and Time

- 2 workers (workload).
- 24 wheelbarrows (work_done).
- 60 minutes (time).
- $\text{Work_time} = \text{workload} * \text{time} = 2 * 60 = 120 \text{ min.}$
- $\text{Response Time} = \text{work_time} / \text{work_done} = 120 \text{ min} / 24 = 5 \text{ min / wheelbarrow.}$

Response Time

- rate of **work_done** and **work_time** is linear.
- as work_done increases so does the work_time.
- **Response Time is constant.**

Response Time

- rate of **work_done** and **work_time** is linear.
- as work_done increases so does the work_time.
- **Response Time is constant.**
- until ...

Response Time

- rate of **work_done** and **work_time** is linear.
- as work_done increases so does the work_time.
- **Response Time is constant.**
- until ...
- the system get's too busy.

Work and Time



4 workers
(waiting)
24 wheelbarrows
in 60 min

Work and Time

- 4 workers (workload).
- 24 wheelbarrows (work_done).
- 60 minutes (time).
- $\text{Work_time} = \text{workload} * \text{time} = 4 * 60 = 240 \text{ min.}$
- $\text{Response Time} = \text{work_time} / \text{work_done} = 240 \text{ min} / 24 = 10 \text{ min} / \text{wheelbarrow.}$

Unit of Work: LIO (logical IO)

- LIO processing is the number-one bottleneck for many business processes.
- LIO consumes two of system's most expensive resources: CPU and latches.

<https://method-r.com/wp-content/uploads/2017/07/Why-You-Should-Focus-on-LIOs-Instead-of-PIOs.pdf>
<https://blog.pythian.com/do-you-know-if-your-database-slow/>
<https://blog.orapub.com/20181204/do-direct-path-reads-count-as-logical-reads.html>

Work and Time

	Workers	Oracle Server
workload	4 workers	active sessions
unit of work	wheelbarrow	LIO*
work done	24 wheelbarrows	LIOs
wall time	60 min	
work time	workload * time	DB time
response time	work time / work done	

Oracle Server Calculation

- workload = DB time / wall time.
- **response time** = DB time / LIO (ms/LIO).
- throughput: LIO / wall time (LIO/ms).
- **Response time is hardware dependent.**

Oracle Server Calculation

- workload = DB time / wall time.
- **response time** = DB time / LIO (ms/LIO).
- throughput: LIO / wall time (LIO/ms).
- **Response time is hardware dependent.**
- Calculate response time **baseline** when database performs »acceptable«.
 - **Carve it in stone.**

How is database performance
doing today?

Good, normal or bad?

How is database performance doing today?

Good, normal or bad?


current response time



baseline response time

Useful cases

- Virtual environment.
- Evolving application.
- New application installed.
- Database upgrade.
- ...



Make IT

2022

- 9:50 Boris Oblak
One indicator to rule them all
(James Cook)
- 12:35 Urh Srečnik
OLVM: Inkrementalni backupi (dirty bitmaps)
(James Cook)



<http://www.abakus.si/>