

# Make IT

2022

## OLVM Incremental Backups

Speaker:

- **Urh Srečnik** <[urh.srecnik@abakus.si](mailto:urh.srecnik@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

# APPM

Abakus Plus  
Performance  
Monitor



- For Oracle Database Standard Edition
- Made by DBAs for DBAs
- Temporal performance comparison
- Resource allocation optimization
- Database performance tracking
- Performance bottleneck optimization

**One Indicator to Rule Them All!**

Speaker: **Boris Oblak**

Friday, 3. june, 9:50 – 10:35

Hall: James Cook

# Backup server

supports Oracle Databases and OLVM VMs



- **Backup**  
takes no time
- **Recovery**  
data recovery is almost instant
- **Disk space**  
backed up data takes up minimal amount of disk space
- **Availability**  
data is always available and always in view
- **Security**  
backed up data can not be deleted without support personnel intervention
- **Alternative uses**  
BI analysis / reporting / DB upgrade verification / R&D testing / seamless business continuation

# Customers

Gorenjska  Banka

 **GENERALI**  
Zavarovalnica

 Ljubljana Airport

 **EKDIS**  
Ekspresno. Ekonomično.

  
REPUBLIKA SLOVENIJA  
MINISTRSTVO ZA OBRAMBO

 **NOVA BANKA**

 **MILENIJUM OSIGURANJE**

 KONTROLA ZRAČNEGA PROMETA SLOVENIJE

 **Iskra**

 **Hidria**

 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**

 **TSS** PEST MANAGEMENT SOLUTIONS

**MAGNETIK d.o.o.**  
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)

# Linux KVM

- <https://www.linux-kvm.org/>
- *KVM (for Kernel-based Virtual Machine) is a full virtualization solution for Linux on x86 hardware containing virtualization extensions (Intel VT or AMD-V). It consists of a loadable kernel module, `kvm.ko`, that provides the core virtualization infrastructure and a processor specific module, `kvm-intel.ko` or `kvm-amd.ko`.*

# QEMU

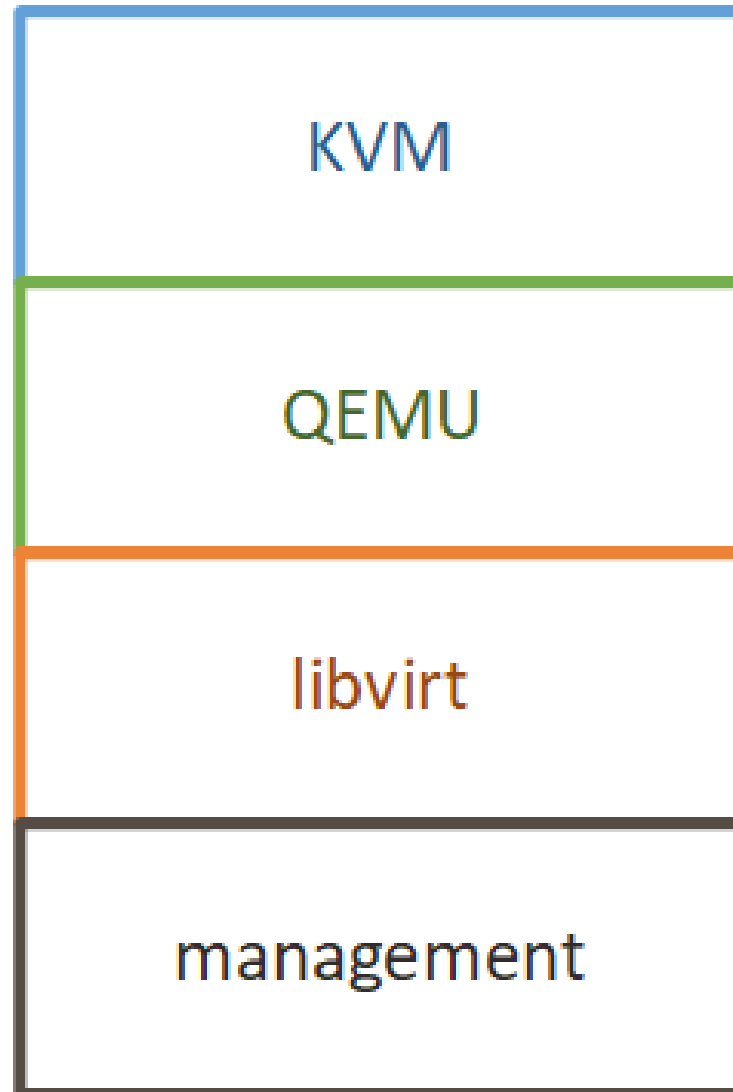
- A generic and open source machine emulator and virtualizer
- *Run operating systems for any machine, on any supported architecture*
- <https://www.qemu.org/>



# QEMU/KVM Management Solutions

- Based on libvirt:
  - oVirt
  - RHEV
  - OLVM
  - virt-manager
- **Not** based on libvirt:
  - Proxmox

# »Layers«



# Storage-level solutions?

- You can snapshot, backup, replicate, migrate, ...
  - From **storage** (SAN, CEPH, ...)

*Neprekinjeno poslovanje na odprtokodni način*

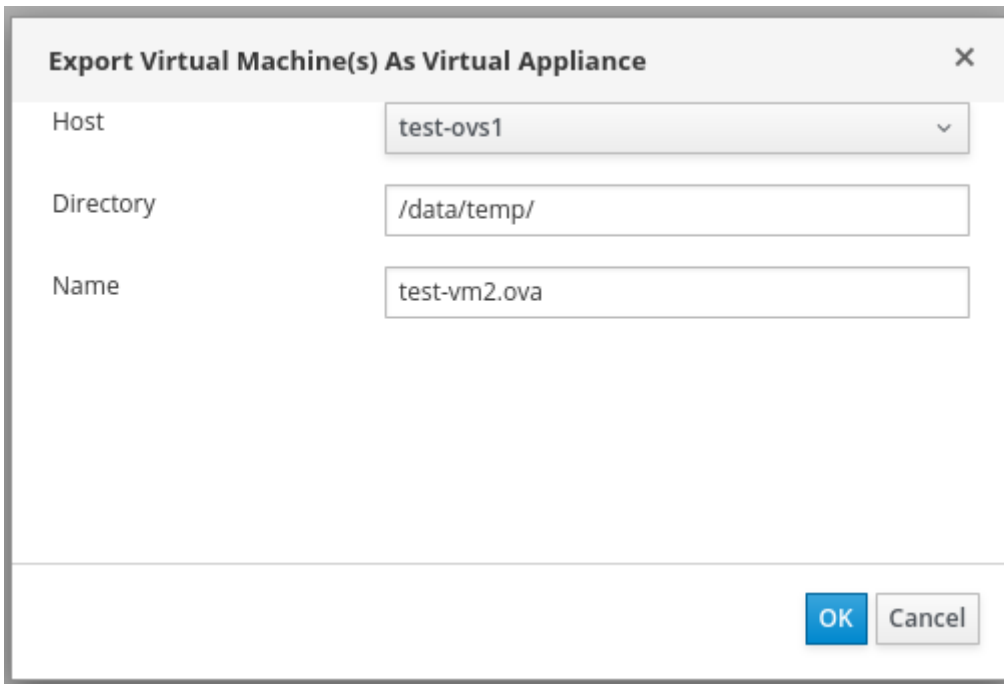
Speaker: **Sergej Rožman**

Thursday, 2. june, 13:40 – 14:25

Hall: James Cook

# Export?

- Export to Export Domain
- Export as OVA
- Also, Storage Domains can be detached/attached



The screenshot shows a dialog box titled "Export Virtual Machine(s) As Virtual Appliance" with a close button (X) in the top right corner. The dialog contains three input fields: "Host" with a dropdown menu showing "test-ovs1", "Directory" with a text box containing "/data/temp/", and "Name" with a text box containing "test-vm2.ova". At the bottom right, there are two buttons: "OK" (highlighted in blue) and "Cancel".

**All those actions  
require VM to be  
shut down.**

# OLVM: What is the Backup and Restore Solution for VMs running in OLVM (Doc ID 2677830.1)

- VM configuration backup can be taken by following oVirt guide under Section: '**Backing Up and Restoring Virtual Machines Using the Backup and Restore API**'
- *Currently there is no official support from Oracle on the backup solution.*

# oVirt Backup API

- The Backup and Restore API
- Incremental Backup and Restore API
- [https://ovirt.org/documentation/administration\\_guide/#sect-Backing\\_Up\\_and\\_Restoring\\_Virtual\\_Machines\\_Using\\_the\\_Backup\\_and\\_Restore\\_API](https://ovirt.org/documentation/administration_guide/#sect-Backing_Up_and_Restoring_Virtual_Machines_Using_the_Backup_and_Restore_API)

# oVirt REST API

```
$ curl --insecure \  
  --user "admin@internal:PASSWORD" \  
  --header "Accept: application/xml" \  
  'https://hostname.fqdn.com/ovirt-engine/api/vms/<VMID>
```

```
<vm ... id="695b29bc-073f-4ae3-b3ca-d4a173567b23">  
  <name>vm-name</name>  
  <description>test description</description>  
  <cpu>  
    <architecture>x86_64</architecture>  
    <topology>  
      <cores>1</cores>  
      <sockets>16</sockets>  
      <threads>1</threads>  
    </topology>  
  </cpu>  
  ...  
</vm>
```

# The Backup and Restore API

1. Create a snapshot of the virtual machine to back up
2. Retrieve the OVF definition (xml) from snap
3. Get the snapshot ID
4. Identify the disk ID of the snapshot
5. Attach the snapshot to a **backup virtual machine** as an active disk attachment
6. Use the **backup software on the backup virtual machine** to back up the data on the snapshot disk
7. Remove the snapshot disk attachment
8. Delete the snapshot



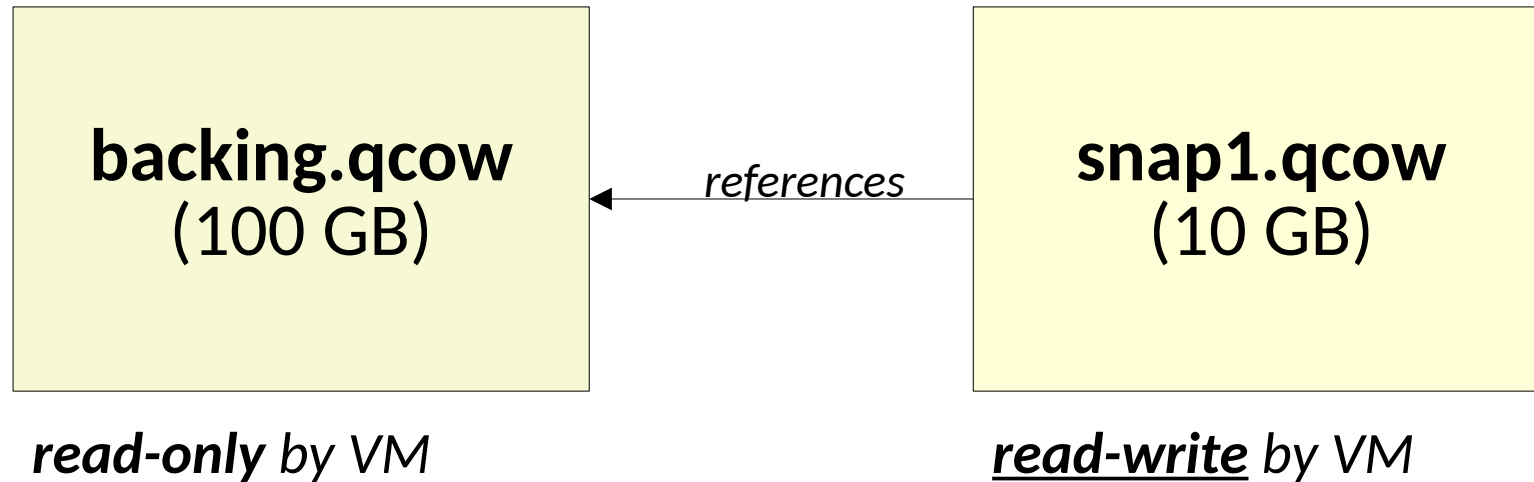
# Incremental Backup and Restore API

1. Find virtual machine disks that should be included in the backup.
2. The backup application starts a:
  1. **full backup** or
  2. **incremental backup.**
3. Wait for engine to prepare the virtual machine for backup.
4. Create an **image transfer object** for every disk in backup
5. Get a list of changed blocks from ovirt-imageio for every image transfer.
6. Download changed blocks in RAW format
7. Finalize all image transfers and backup

## (cont.) Full Backup

*When taking a [full] backup of a running virtual machine, the process **creates a scratch disk** on the same storage domain as the disk being backed up. The backup process creates this disk to enable new data to be written to the running virtual machine during the backup. It is automatically deleted when the backup finishes.*

# QCOW Snapshots



```
# qemu-img info snap1.qcow
image: /path/to/snap1.qcow
file format: qcow2
virtual size: 15 GiB (16106127360 bytes)
disk size: 192 KiB
cluster_size: 65536
backing file: backing.qcow (actual path: /path/to/backing.qcow)
backing file format: qcow2
...
```

# QMP Protocol

- <https://wiki.qemu.org/Documentation/QMP>  
*The QEMU Machine Protocol (QMP) is a JSON-based protocol which allows applications to control a QEMU instance.*

```
$ telnet localhost 4444
Trying 127.0.0.1...
Connected to localhost.
Escape character is '^]'.
{
  "QMP": {
    "version": {
      "qemu": {
        "micro": 0,
        "minor": 0,
        "major": 3
      },
      "package": "v3.0.0"
    },
    "capabilities": [
      "oob"
    ]
  }
}
```

# QEMU Dirty Bitmaps

- <https://www.qemu.org/docs/master/interop/bitmaps.html>
- *Dirty Bitmaps are in-memory objects that track writes to block devices. They can be used in conjunction with various block job operations to perform incremental or differential backup regimens.*

# Abakus Method

## (using QEMU Dirty Bitmaps)

- Works with **any managment layer** as we only communicate with QEMU via QMP.
- Can backup **any format** (qcow, raw, direct lun, ...)
  - Even when snapshotting is not available
- Uses **incremental backups** using dirty bitmap index.



# Abakus Method (for each blockdev)

1. (*first-time only*) Create or clear dirty bitmap index
2. (*first-time only*) Backup blockdev contents
3. Create incremental backup from bitmap index
  1. (This also clears the bitmap index)
4. Apply incremental backup to backup image

# Caveat, 1st

- Each management layer require different format of vm definition
  - libvirt: virsh dumpxml
  - proxmox: contents of /etc/pve/qemu-server/<VMID>.conf
  - oVirt: OVF export of its definition from PG database
- We need management-layer specific code to obtain those.



# Caveat, 2nd

- Management layer may »see« disks under different »user friendly« names as QEMU process:
  - Disk name (as seen in management interface)
  - Disk metadata (description, such as »backup=0«)
- We need management-layer specific code to obtain disk names as seen from management GUI.

# Caveat, 3rd (qemu snaps)

- VM sees current snapshot
  - ... and complete chain leading to backing image
- We need code to »resolve« and apply all those snapshots of snapshots of snapshots to the backing image.

# Caveat, 4th

- We need running QEMU process.
- We need to start VM (in paused mode) if it is currently shut down.

# Caveat, 5th

- QEMU Dirty Bitmaps can only persist between reboots if:
  - qcow format is used.
  - clean shutdown was performed (e.g. not »vm kill« :) )
- So, we need either:
  - Use qcow format, or
  - Full backup after each vm boot.
- this is also the reason why some vendors only allow qcow format to be used in incremental strategies.

# Abakus Restore

### Restore Ovirt VM ✕

#### VM Manager

VM Manager  ▼

Password

Login

#### VM Restore

Cluster  ▼

VM Name

Restore VM definition from OVF

Regenerate UUID

Regenerate MAC

Transfer disk images via VM Manager proxy

Disk Image File	Size GB	Restored Disk Name	Storage Domain	Meta	Data
<input type="text" value="olvm-slave_root.full.raw"/>	100	<input type="text" value="olvm-slave_root.fu"/>	vmimg (glusterfs) <span>▼</span>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<input type="text" value="olvm-dejavu_vbsloc.full.raw"/>	150	<input type="text" value="olvm-dejavu_vbslc"/>	vmimg (glusterfs) <span>▼</span>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<input type="text" value="olvm-dejavu_demo.full.raw"/>	15	<input type="text" value="olvm-dejavu_demo"/>	vmimg (glusterfs) <span>▼</span>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Restore

# Abakus Restore

Compute » Virtual Machines

Vms:  ✕ ☆ ▼ 🔍

New Edit Remove ▶ Run ⌚ Suspend ■ Shutdown ↺ Reboot 💻 Console Migrate Create Snapshot ⋮

🔄 ▼ 1 - 4 < >

		Name	Memory	CPU	Network	Status	Uptime	Description
▲	🚗	██████████	9%	12%	0%	Up	135 days	██████████
▲	🚗	██████████-dejavu	2%	0%	0%	Up	20 h	
▼	🚗	██████████-dejavu-T1	--	--	--	Down		Restored using Abakus Backup Server
▲	🚗	██████████	4%	0%	0%	Up	233 days	

Storage » Disks

Disks:  ✕ ☆ ▼ 🔍

New Remove Move Copy Upload ▼ Download ⋮

Disk Type: All Images Direct LUN Cinder Managed Block Content Type: All ▼

🔄 ▼ 1 - 27 < >

Alias	Virtual Size	Status	Type	Description
██████████-dejavu_demo	15 GiB	OK	Image	testni disk za testiranje restore-a
██████████-dejavu_demo.full.raw	15 GiB	OK	Image	Restored using Abakus Backup Server
OVF_STORE	< 1 GiB	OK	Image	OVF_STORE

# Backup server

supports Oracle Databases and OLVM VMs



- **Backup**  
takes no time
- **Recovery**  
data recovery is almost instant
- **Disk space**  
backed up data takes up minimal amount of disk space
- **Availability**  
data is always available and always in view
- **Security**  
backed up data can not be deleted without support personnel intervention
- **Alternative uses**  
BI analysis / reporting / DB upgrade verification / R&D testing / seamless business continuation