



oVirt Changed-Block Tracking API

Developer's perspective

Marcin Kubacki
Chief Software Architect

September/2021



Agenda

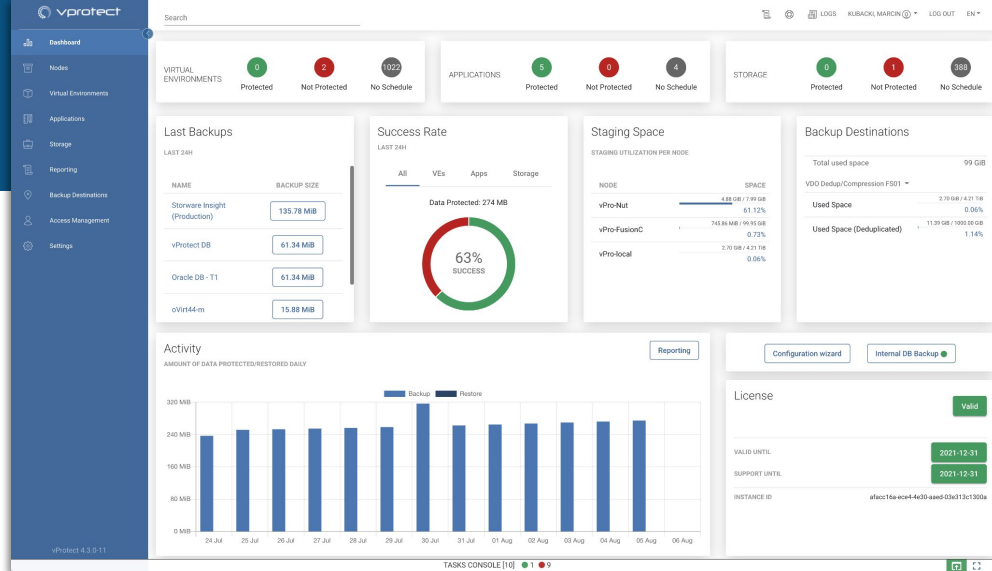
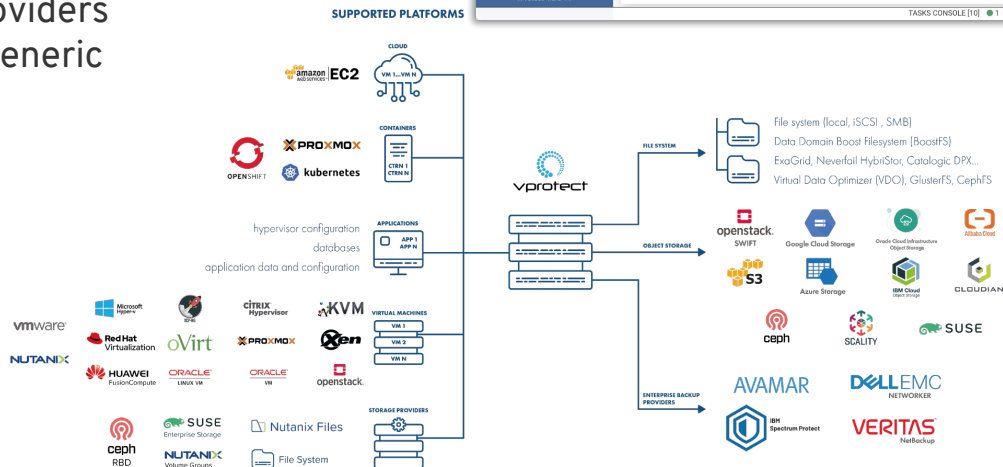
- Agent-less data protection for oVirt
 - vProtect overview
- Few strategies to backup VMs in oVirt
- oVirt CBT APIs and SDK
- Demo

Agent-less data protection for oVirt

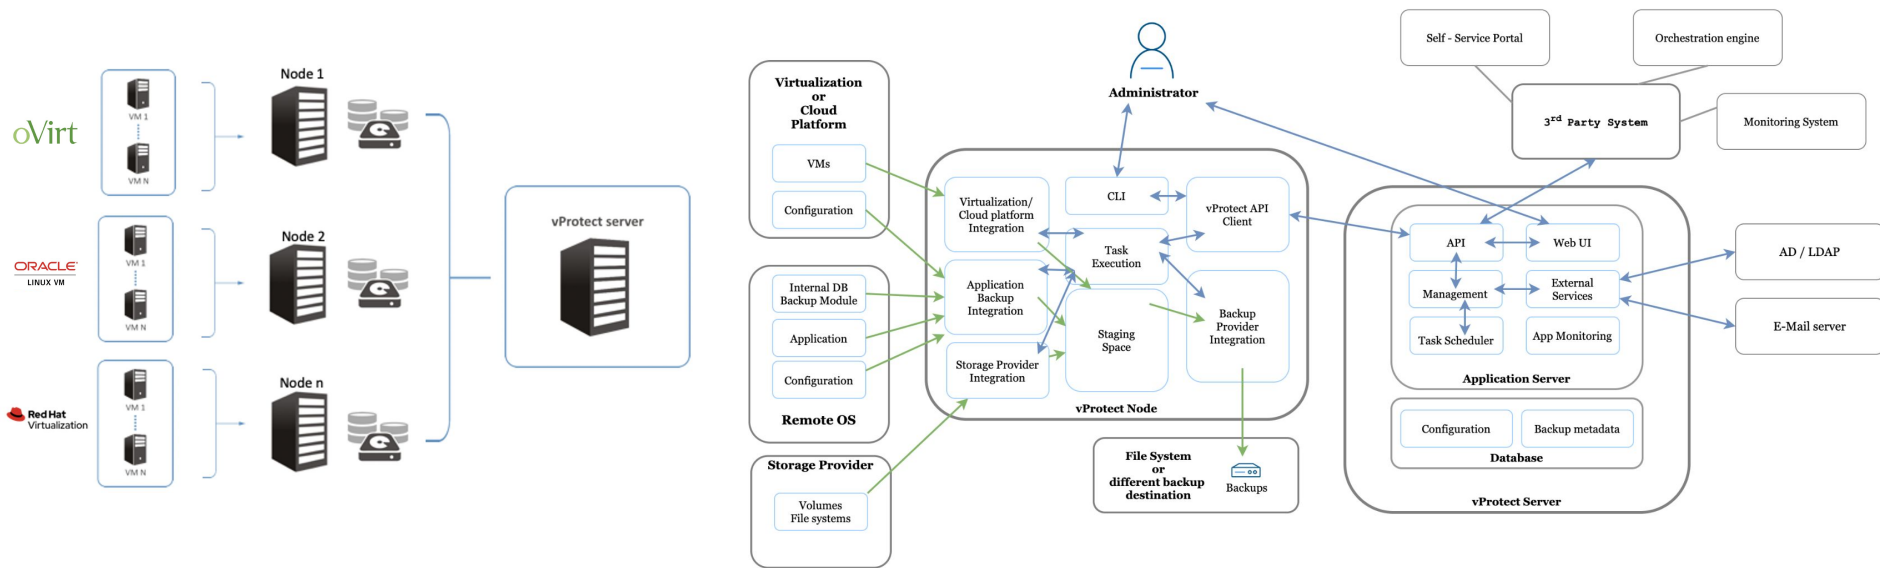
vProtect overview

Backup solution for multiple virtualization platforms including oVirt-based, such as RHV or OLVM:

- Agent-less backups
 - incremental with CBT
 - synthetic
- File-level restore
- Snapshot management
- Multiple backup providers
- Application-level, generic backup mechanism
- Storage providers

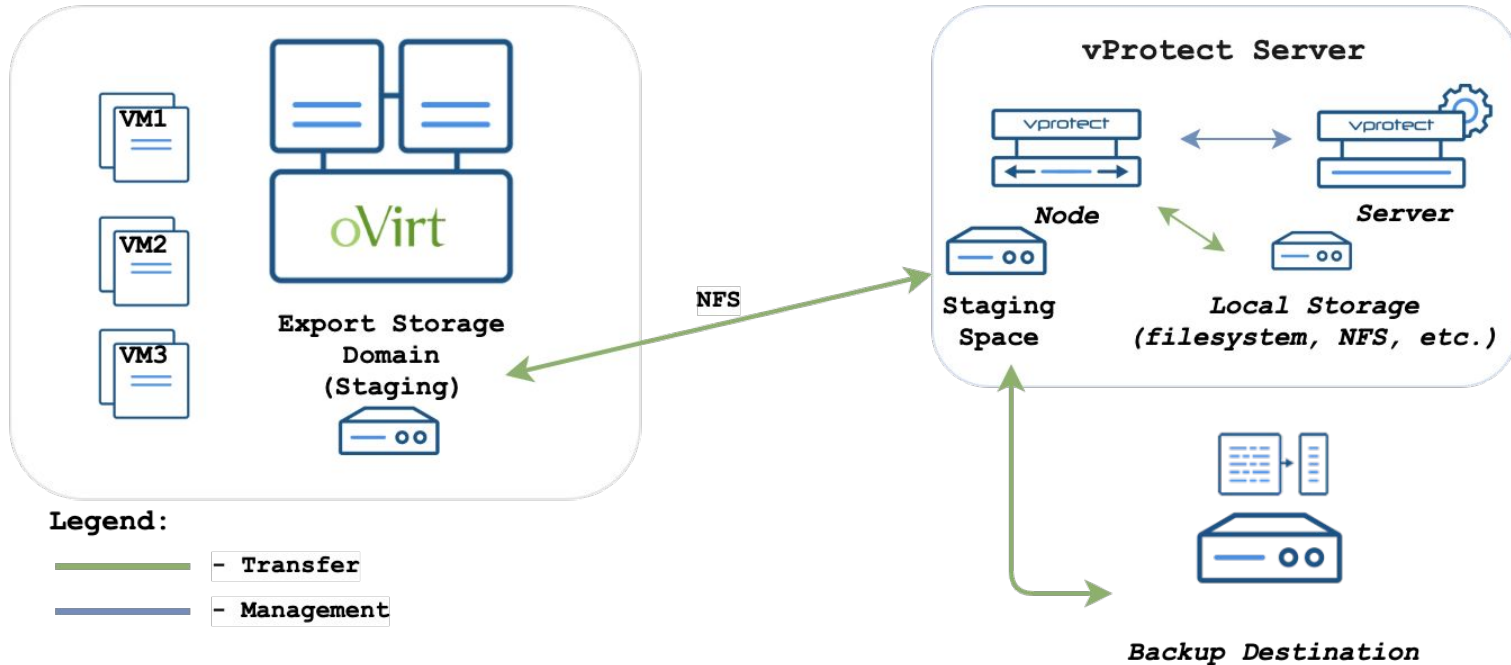


vProtect architecture



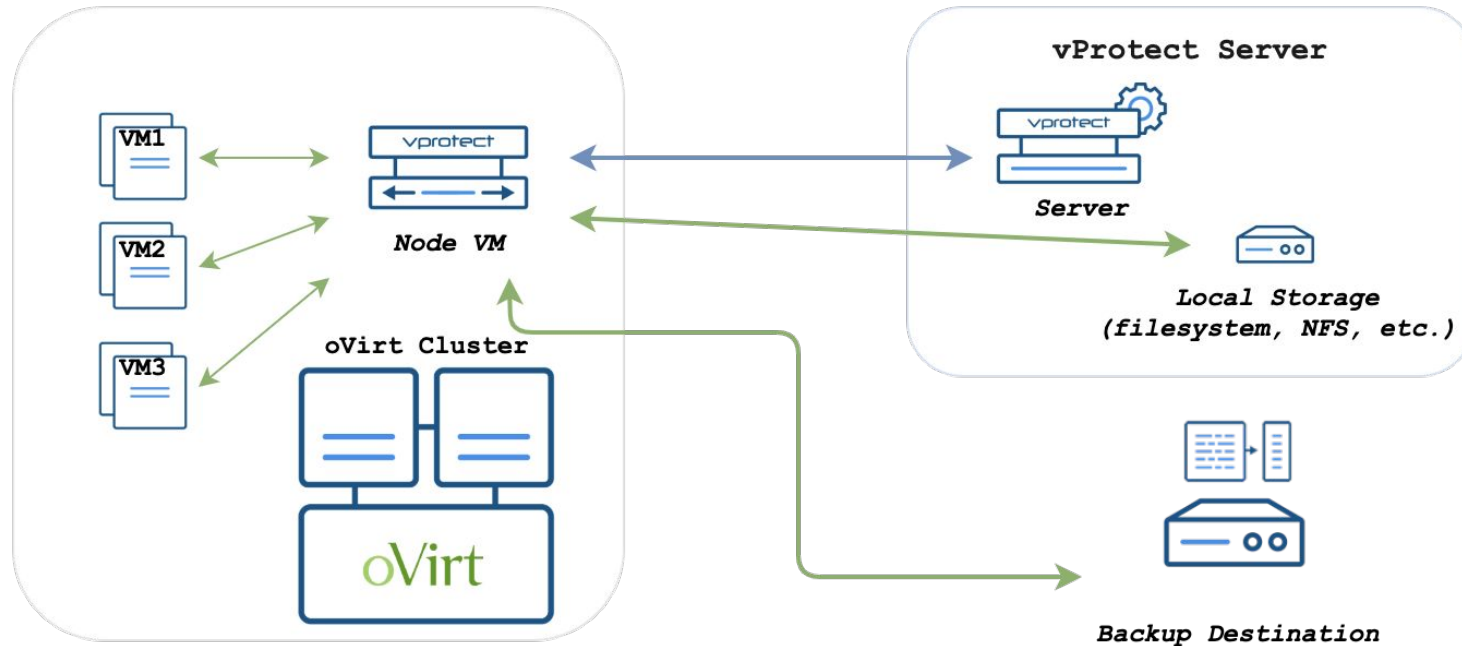
Backup strategies

Export storage domain (v3 API)



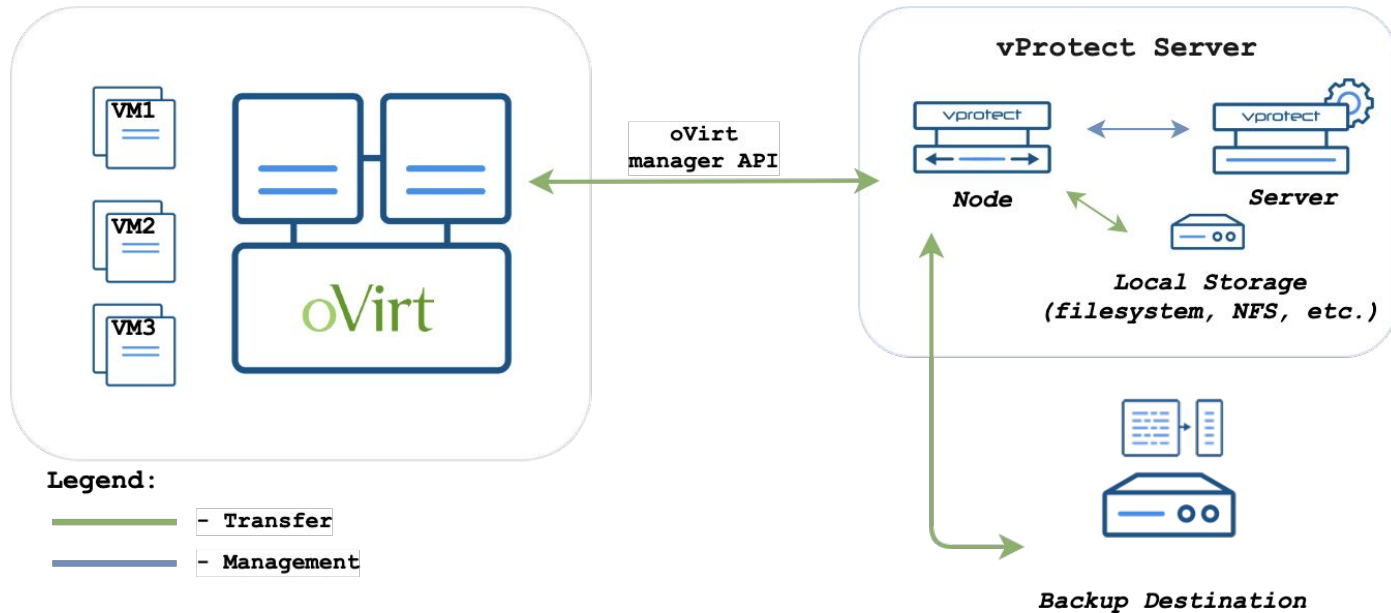
- Filesystem, NFS, deduplication appliances
- Object storage on-premise or via cloud providers
- Enterprise backup providers (Dell EMC, IBM, Veritas)

Disk attachment



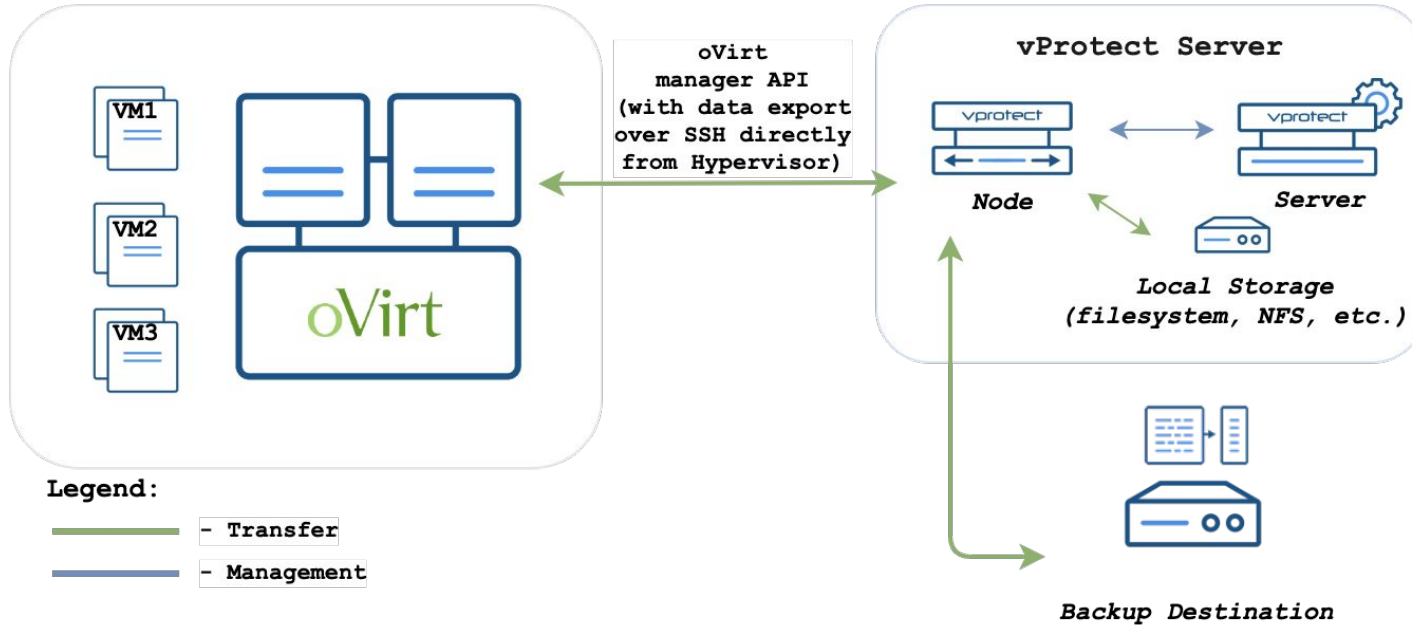
- Filesystem, NFS, deduplication appliances
- Object storage on-premise or via cloud providers
- Enterprise backup providers (Dell EMC, IBM, Veritas)

Disk Image Transfer API (4.2)



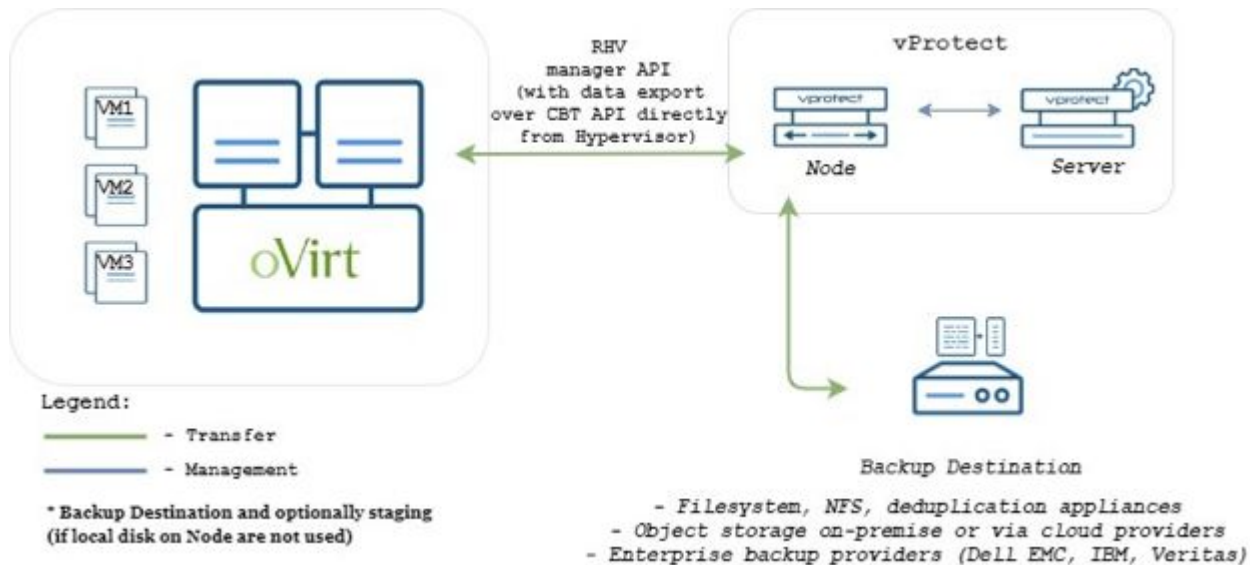
- Filesystem, NFS, deduplication appliances
- Object storage on-premise or via cloud providers
- Enterprise backup providers (Dell EMC, IBM, Veritas)

SSH Transfer



- Filesystem, NFS, deduplication appliances
- Object storage on-premise or via cloud providers
- Enterprise backup providers (Dell EMC, IBM, Veritas)

Changed-Block Tracking (CBT)



oVirt CBT APIs

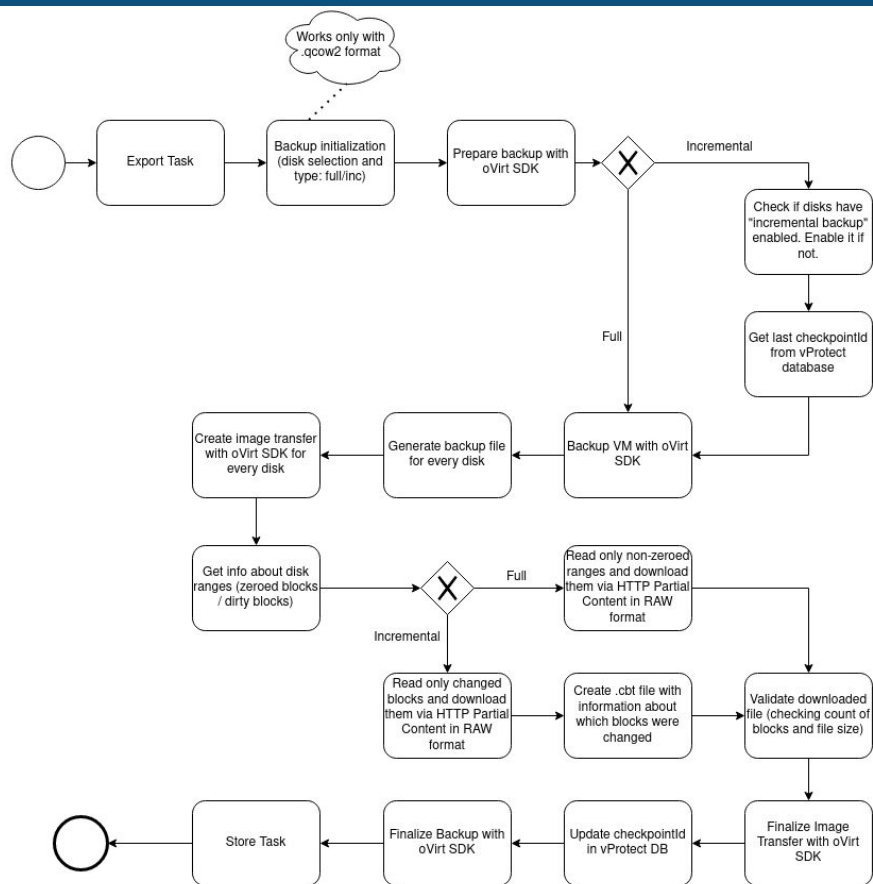
Changed-Block Tracking

- CBT API allows to query for changes that occurred in the VM disk
- Snapshots have no longer to be kept in the virtualization platform for future backups
- Behind the scenes, oVirt, or more accurately KVM with QEMU is going to keep track of what has changed using dirty-block maps.
- Snapshots can occupy a significant amount of space and affect performance - no wonder why CBT was a long-awaited feature.
- Whenever vProtect performs a backup, it needs to record checkpoint ID
 - an ID that is later used to perform the next backup to refer to changes that happened since this particular point in time
 - it also includes information about zeroed blocks.
- Disk Image Transfer APIs still used as a base
 - Data export directly from hypervisor - transfer from the Image I/O service that runs on hypervisor
- In the past (before oVirt 4.4.5) VMs that were down could not be backed up using CBT method
 - fallback to regular Disk-Image Transfer

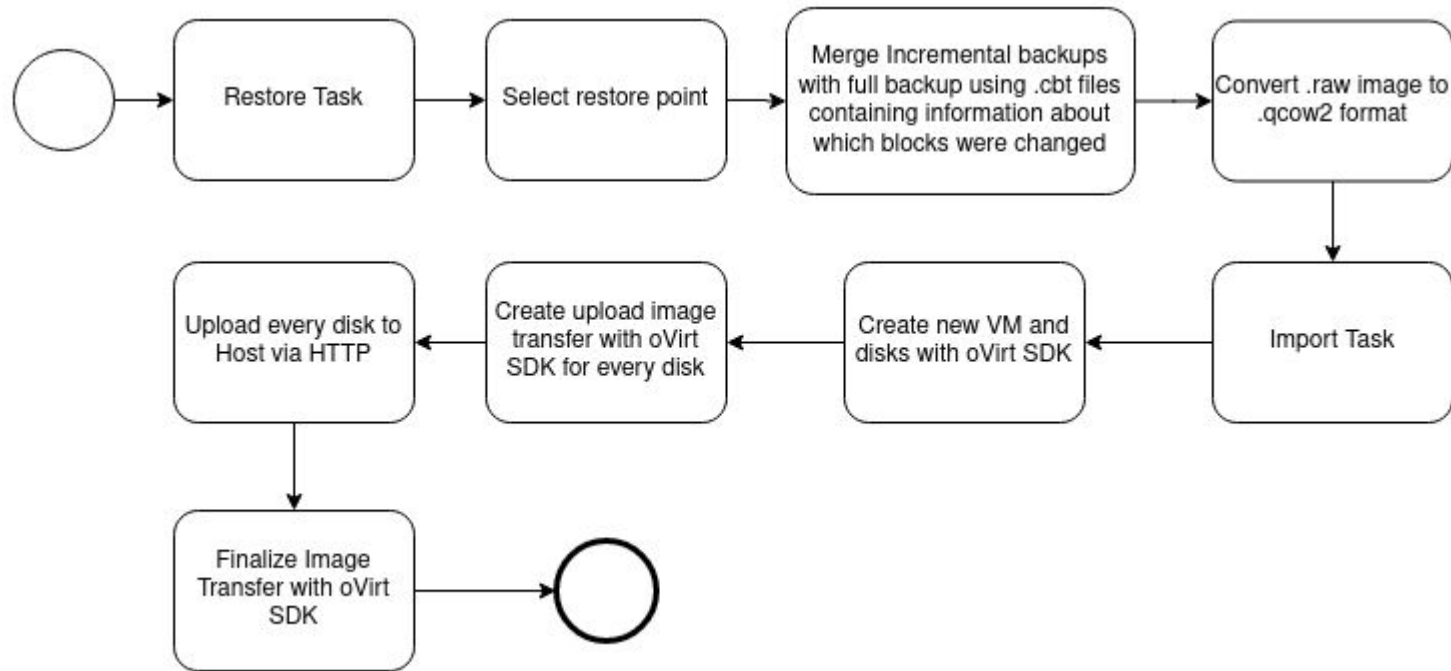
Changed-Block Tracking

- Compared to SSH transfer method:
 - also transfers data from the hypervisor, optionally with netcat,
 - protocol/communication-related aspects have a lower impact on the overall performance
 - multiple HTTP requests for parts of the blocks are generally costly compared to pure netcat transfer
 - but you don't have to read unnecessary blocks, which may significantly shorten the time of the data export phase.

Backup process with CBT



Restore process with CBT



oVirt SDK for Java

- SDKs available for different languages
- You also can invoke APIs directly (HTTP) or with Ansible
- Java example:

```
public Vm getVm(String id) {  
    return conn.systemService().vmsService().get().send().vm();  
}
```

- pom.xml:

```
<dependency>  
    <groupId>org.ovirt.engine.api</groupId>  
    <artifactId>sdk</artifactId>  
    <version>4.2.5</version>  
</dependency>
```

Image Transfer API - changes with CBT API

```
Backup backup = vmBackupsService
    .add()
    .backup(backupBuilder)
    .send()
    .backup();
```

```
ImageTransfer imageTransfer = c.getImageTransfersSvc().addForDisk().imageTransfer(
    imageTransfer()
        .direction(direction)
        .disk(disk)
        .backup(backup)
        .inactivityTimeout(INACTIVITY_TIMEOUT)
        .format(DiskFormat.RAW))
    .send().imageTransfer();
```

A few blockers on the road...

- Initial lack of documentation
 - Hopefully SDKs could help us to reverse-engineer how to process should look like
- oVirt Image I/O - missing ticket-uuid - how to get it?
 - Creating “backup” object in the previous example was the solution - unlike other strategies
- *Parent checkpoint ID does not match the actual leaf checkpoint* during full backups problem:
 - oVirt update helped
- Issue with downloading everything instead of just changed blocks
 - Missing range header in requests
- Problem with backup of powered down VMs
 - Workaround with original Disk Image Transfer API
 - Later solved in oVirt 4.4.5
- VMs going down after backups
 - Seems to be solved in newer oVirt releases

Demo

oVirt

Thank you!

<https://ovirt.org>

<https://www.openvirtualization.pro>
<https://storware.eu>

users@ovirt.org

info@storware.eu



@ovirt

@OpenVirtPro @Storware