

# IBM InfoSphere Managed Data License

4/29/19

IBM InfoSphere Virtual Data Pipeline (VDP) licensing is based on source application data under management. This is referred to as Managed Data License (MDL). This document explains:

[How is MDL Measured?](#) on page 1

[Factors that Influence MDL Calculations](#) on page 3

[Frequently Asked Questions](#) on page 3

## How is MDL Measured?

IBM InfoSphere measures MDL based on the actual application size at the front-end. If the application under management reports the volume size for data under management, IBM InfoSphere takes into account the volume size reported (for example, the MDL calculation for VMware will be consistent with the reported size of the VM in vCenter). IBM InfoSphere MDL can be independently verified from the application.

The unit of measurement is TB. In other words, IBM InfoSphere measures usage based on the actual size of the applications it is managing. If you manage 10 TB of Oracle data spread across multiple databases, the IBM InfoSphere MDL reports 10 TB of MDL use. For more information on calculations, see [Frequently Asked Questions](#) on page 3.

IBM InfoSphere MDL measurement takes into consideration only the size of the data under management based on the last successful copy. It does *not* take into account:

- The frequency of data capture
- Where the copies are stored or how many copies are made
- The data change rate
- The retention period or how much storage is occupied by copies

IBM InfoSphere VDP - Global Manager (IVGM), Resiliency Director (RD), and Object Recovery for Exchange and SharePoint licensing are in addition to the base MDL licensing.

This section details how MDL is computed for various use cases.

- [MDL Computation for VMware](#)
- [MDL Computation for Microsoft Hyper-V](#)
- [MDL Computation for Applications Stored Out-of-Band](#)
- [MDL Computation for Databases](#)
- [MDL Computation for OnVault Images of Unprotected Applications](#)

[Factors that Influence MDL Calculations](#) and [Frequently Asked Questions](#) on page 3 address additional scenarios.

## MDL Computation for VMware

When the entire VMware VM is managed by InfoSphere VDP Appliances (all volumes), the MDL count is the total allocated size of all managed volumes. If the volumes are thin-provisioned, then thin provision values are used. If the volumes are thick-provisioned, then the total allocated size is used.

IBM InfoSphere takes into account the size of the managed volumes reported by VMware vCenter. If a VM has been configured with an 8 TB thin-provisioned volume, and it has been allocated 2 TB, then 2 TB is counted towards MDL use. When the allocated size increases, the MDL count will increase from the next successful copy after the increase in size.

This method of measurement applies all VMware backup scenarios. When an IBM InfoSphere VDP Connector is installed inside a VM and the application data within the VM is captured using the VDP Connector, the MDL calculations are based on out-of-band connector based model described in [MDL Computation for Applications Stored Out-of-Band](#).

## MDL Computation for Microsoft Hyper-V

MDL measurement is consistent with MDL measurement for out-of-band applications using the VDP Connector, because the Hyper-V solution leverages the VDP Connector on the Hyper-V server. The size of .VHD and .VHDX files under protection will be counted towards MDL. This will be the actual size of the file on the disk.

## MDL Computation for Applications Stored Out-of-Band

In out-of-band scenarios, where the VDP Connector is used, the actual size of the application is taken into account. If the managed application is 5 TB in actual size on a 7 TB volume, IBM InfoSphere MDL count will be 5 TB.

## MDL Computation for Databases

For Oracle, Exchange, and SQL Server applications, only the databases protected are counted towards MDL. Log files are not taken into account:

- Oracle: The allocated size of the database files under protection are counted towards MDL. The allocated size includes data files and control files.
- Microsoft SQL Server: The total size of all the database files, including .MDF, .LDF and .NDF files under protection are counted towards MDL. Log files (.TRN) are not counted towards MDL.
- Microsoft Exchange: The total size of the EDB files and the log files under protection are counted towards MDL.

## MDL Computation for OnVault Images of Unprotected Applications

OnVault images are only counted toward MDL while the source application is protected. If the source application becomes unprotected and if all snapshot and dedup images are expired, then the OnVault image is not counted for MDL.

## Factors that Influence MDL Calculations

Here are some important factors that influence MDL calculations in Out-of-Band scenarios:

- **Compressed Volumes:** When volume is compression enabled, the MDL calculations count the post-compression values. For example, if a 2 TB volume has 2.5 TB of data that is compressed into 1.8 TB, MDL count will be 1.8 TB, not 2.5 TB.
- **Windows Optimized Volumes:** For Windows optimized volumes, VDP rehydrates the volume for backup, and the MDL count will be the rehydrated value. For example, if a 1TB Windows optimized volume contains 800GB data, which when rehydrated for backup ends up as 1.1 TB, the MDL will be 1.1 TB.
- **Block sizes:** The block size of the staging disk is taken into account for MDL measurement. If the source volume's block size and the staging disk block size match, then the MDL values will exactly match the source volume. If the block size used on the staging disk is different from the source volume, then there will be a minor difference because the MDL calculation is done on the staging disk.

---

**Note:** Note: Since ZFS volumes use EXT staging disks; ZFS Compression and Dedup do not affect the MDL calculations, which will reflect the full data size.

---

- **Consistency Groups:** The MDL count for a consistency group will be the sum of all application sizes in the consistency group. Applications are measured individually and summed.
- **Reparse points and Linked lists:** These have no impact on MDL.

## Frequently Asked Questions

1. Q: What is the unit of measurement for an IBM InfoSphere MDL?  
A: IBM InfoSphere MDL unit of measurement is TB. The calculations are based on binary (base 2) and not decimal (base 10). In other words, IBM InfoSphere counts 1024 GB as 1 TB, not 1000 GB as 1 TB.
2. Q: Does the Dedup ratio impact my MDL?  
A: The Dedup ratios achieved by the system do not affect the MDL consumption, as the MDL count is based on the front-end application size and not the back-end storage consumption.
3. Q: We just purchased a new InfoSphere VDP Appliance to replicate from our primary appliance for DR purposes. Do I need additional MDL licenses to replicate between InfoSphere VDP Appliances?  
A: IBM InfoSphere MDL is based on front-end application size. It does not take into account how many copies are retained or where they are retained. Adding an appliance for DR purposes *only* will not change MDL consumption. Any applications managed locally on the DR site or VMware datastores replicated between InfoSphere VDP Appliances will impact MDL consumption.
4. Q: If my file system application has 3 TB of data, and I use Prune Paths and Exclude lists to eliminate 1 TB of files from management, does IBM InfoSphere count 3 TB as managed capacity or 2 TB as managed capacity?  
A: The MDL for file system applications is based on the actual amount of data managed, 2 TB in this example.
5. Q: How often does the IBM InfoSphere MDL count get updated?  
A: IBM InfoSphere MDL values are computed and updated once a day, at 3:00 AM (local time on the appliance). The capture operation should be successfully completed before 3:00 AM for it to be included in MDL measurement.
6. Q: My MDL count for my application seems to be lower than what it reported yesterday, why is that?  
A: IBM InfoSphere MDL count is based on the most recent successful copy, not the largest recoverable copy. Applications shrink or expand over a period of time (irrespective of the change rates involved). When the application size shrinks, it is reflected on the MDL count the next day.

7. Q: I am managing a 4 TB Oracle database. The Oracle database has a 10% daily change rate, but the size of the database is always 4 TB. What will be my MDL on any given day?
- A: IBM InfoSphere MDL measurement is based on the size of the last successful copy. In the above example, the managed application size will be 4 TB, so the MDL calculation will be 4 TB. Unless the size of the application changes, the change rate does not directly impact the MDL calculation.
8. Q: If I have an SQL Server application running on a VMware VM, and I manage the SQL Server application using VDP Connector and the entire VM using VMware VADP, will my SQL database be counted on top of the VM?
- A: IBM InfoSphere MDL counts VMware VM separately and the SQL database separately, which can lead to double counting. However, customers usually manage only the OS volume of the VM and not the entire VM, when managing the applications residing on the VMs separately. This effectively eliminates double counting.
9. Q: I manage Microsoft SQL and Oracle applications using VDP. Do you count only the database size or do you include the log files in MDL measurement?
- A: IBM InfoSphere counts only the managed database files that are needed for a consistent database backup towards MDL measurement. IBM InfoSphere does not count log files towards MDL measurement.
10. Q: I'm no longer actively managing an application that was backed up daily for over a year. When will IBM InfoSphere MDL measurement stop including this application?
- A: IBM InfoSphere MDL measurement is based on the last successful copy of the application available. IBM InfoSphere MDL will consider the application for MDL measurement until all the copies under management have expired. In other words, as long as there's a recoverable image, IBM InfoSphere will count it towards MDL. This includes orphan images as well (images that are retained per the SLA, whose source application has been deleted from VDP management).
11. Q: I leverage VDP for my Test/Dev. As part of the Test/Dev we use external tools for data masking and take new snapshots after data masking and present them to the test/dev teams, who also take snapshots periodically. How do changes in workflows like data masking affect my MDL count?
- A: Use of LiveClone and Snapshots manually or via workflow has no effect on MDL. If you rediscover these snapshots as new applications and apply an SLA to them, then they will be counted towards MDL use.
12. Q: Will my operations be disrupted if I reach the limits of my purchased MDL?
- A: IBM InfoSphere recognizes the critical nature of its products operation in customers' environments. There is a built-in grace operation to ensure that jobs are not stopped when MDL utilization extends beyond the purchased MDL. You should monitor your MDL usage regularly and contact IBM InfoSphere when your usage approaches the purchased MDL.
13. Q: Do I need licenses to access/recover/restore my data?
- A: VDP does not check for the presence of MDL licenses during access/restore/recovery. IBM InfoSphere licensing does not come in the way of customer, when the customer tries to access or restore their data.
14. Q: I would like to chargeback my customers using a different scheme. Can I get the required information to generate my chargeback reports?
- A: The Report Manager provides a rich set of reports that you can leverage for chargeback. If you have needs that are not met with the Report Manager, please contact your IBM InfoSphere Representative.
15. Q: How would a Cloud Service Provider running a large public IaaS do chargeback and bill his customer? Should he use MDL or some other chargeback model?
- A: The CSP could use MDL based model along with some level of margin applied or choose to chargeback based on backend storage capacity of the applications with the dedup ratio applied.
16. Q: I have an application with 3 months retention. I no longer need to protect this application, so I no longer actively protect it. When will the MDL consumed by this application be released?

A: Any application that has a recoverable image, either under active protection or inactive, will be counted towards MDL use. There are two ways the MDL is released, at the end of 3 months when all recoverable images for the application have expired, or by manually expiring all images for the specific application.

17. Q: What happens if I accidentally delete an application and then re-discover it again?

A: When an application is deleted and rediscovered, the newly discovered application is considered as a brand new application for MDL calculations. If recoverable images of the deleted application remain in the system, those images count towards MDL.

18. Q: How do I verify my VMware MDL usage?

A: IBM InfoSphere MDL calculations for VMware are consistent with the VMware vCenter reported size for that VM.

`du -h *.vmdk` output from the appropriate VM folder on the datastore should match the MDL count.

---

**Note:** When snapshots external to IBM InfoSphere are found on the VM, the allocated size of the VM will be taken for MDL measurement until the snapshots are deleted, as the external snapshots can artificially inflate the MDL calculations.

---

19. Q: How do I verify my Oracle MDL usage?

A: IBM InfoSphere MDL calculations for Oracle are based on the allocated size for the database. Here is a sample query to verify the Oracle database size.

```
select (d.total + c.total) total from (select sum(bytes) total from v$datafile) d,  
(select sum(block_size*file_size_blks) total from v$controlfile) c;
```

Then subtract the following: `select sum(bytes) free from dba_free_space;`

20. Q: How do I verify my MDL usage for File Systems?

A: IBM InfoSphere MDL calculations for file system based applications:

Windows Used File System size reported by DiskManager

Linux Used File System size reported by `df -k`

AIX Used File System size reported by `df -k`

Solaris Used File System size reported by `df -k`

HP-UX Used File System size reported by `bdf`

---

**Note:** Start paths and prune paths ARE respected. Therefore, the MDL could be lower than DiskManager or `df -k` output, and additional scripting would be required to get the actual MDL measurement.

---