

UniVirStor Manager v2.15.18 Release Notes

UVS Manager 2.14.18 Release Notes	2
What's New in Ceph Octopus	2
What's new in UVS Manager v2.15.18	2
Upgrade Guide	3
Upgrade from Version 2.14.17	3
Prerequisite	3
If a node fails to upgrade	4
Step1: Upgrade the First Monitor Node	4
Step 2. Upgrade other monitors by UVS manager	6
Step 3. Upgrade OSD nodes.	7
Step 4. Upgrade MDS nodes	7
Step 5: Check the software version after upgrade.	7
Step 6: Upgrade the RADOS Gateway.	8
Step 7: Upgrade the external iSCSI Gateway.	8
Upgrade Note	9

UVS Manager 2.14.18 Release Notes

We are delighted to announce the release of UVS manager v2.15.18. v2.15.18 is our first release for Ceph Octopus. The new shipments of Mars 400 will use this version as their default release.

We modified a significant portion of UVS codes to introduce the Ceph new management and orchestration tool in this release. The latest release enhances many user experiences on UVS manager. As the Ceph Octopus requires Python 3 support, we also upgrade the CentOS to version 8.

Deploying a new Ceph cluster or upgrade to v2.15.18 on Mars 400 will automatically run Ceph MON, OSD, and MDS in containers.

Using the earlier version, customers can keep using the Ceph Nautilus or upgrade it to Ceph Octopus. We also released a UVS version 2.14.17-3 to fix some bugs in the UVS manager and upgrade the Ceph to version 14.2.16 for customers who want to stay in the Ceph Nautilus. Please refer to the 2.14.7-3 release notes. Ambedded will continue the support of Ceph Nautilus on the Mars 400 appliance.

If your Mars 400 cluster runs Ceph Nautilus, we provide the automatic updater to upgrade your Ceph cluster to Octopus without interrupting your storage service. The updater upgrades the Ceph, the Linux kernel, and CentOS. The updater will help you to update your cluster node by node, and the update will not interrupt your service.

What's New in Ceph Octopus

- All Ceph daemons are now running as containers.
- New Ceph dashboard
- PG autoscaling: You can allow the cluster to either make recommendations or automatically tune PGs based on how you use the cluster by enabling *pg-autoscaling*.
- Partial object recovery replaces log-based recovery to improve the speed of recovery.
- Snapshot trimming improvement eliminates metadata in the OSD map that (previously) would grow with cluster age.
- Bluestore performance improvement
- Support new RBD client cache policy write-around as the default setting. The write-around policy differs from the write-back policy in that it does not attempt to service read requests from the cache, unlike the write-back policy, and is, therefore, faster for high-performance write workloads.
- Snapshot-based RBD mirroring reduces the I/O overhead of journal-based snapshots. Snapshot-based mirroring supports kernel RBD that journal-based mirroring doesn't.
- RGW supports object lock for Write-Once-Read-Many(WORM) data protection.

What's new in UVS Manager v2.15.18

- **Ceph Version:** 15.2.8 (Octopus)
- **Operating System**

- Upgrade Linux to CentOS 8
- Kernel upgrade to 5.4.92
- **Deployment & Management:**
 - Use cephadm and ceph orchestrator as the deployment and management tools
 - Modify the UVS manager to meet the requirement of deploying ceph on containers
 - Modify the UVS manager to adopt the change of Ceph Octopus.
 - Modify the console RESTORE function to accelerate the container deployment.
- **NTP Server:** When deploying the new cluster, the UVS manager will automatically use the node as the default NTP server when deploying a new cluster. This can avoid the error if the user forgets to configure the NTP server. Users can change the NTP server after cluster deployment.
- **Network Configuration:**
 - Add configuration option for separating the Ceph public and in-cluster network.
 - Replace the Linux bonding module with Open vSwitch SLB bonding to provide a better stable network.
- **OSD:** Add an option for deploy OSD to use the SATA HDD/SSD as the bluestore storage
- **MON:** Users can deploy multiple Ceph monitors in parallel to shorten the deployment time.
- **RADOS Gateway**
 - Simplify the deployment of new RGW
 - Unify the stand-alone and multisite RGW. All new deployed RGW will be multi-site ready.
- **Ceph File System:** Simplify the deployment of the Ceph file system. Now users can create the CephFS, pools, and metadata server (MDS) on one user interface.
- **Demo license:** A demo license will build in the new Mars 400 shipments.
- **Bug Fixed:**
 - Fix wrong setting of CephFS max_mds after the UVS manager upgrade (Backport to v2.14.17-3)
 - Use log rotation to prevent the logs from over-consuming the storage.

Upgrade Guide

Upgrade from Version 2.14.17

Please read this upgrade note carefully before you start to upgrade your Mars 400 appliance software.

There are three major software upgrades.

1. Linux: Due to the Ceph version Octopus (v15.2.x) Python 3 module dependencies, we have to upgrade the current CentOS 7 and kernel to CentOS 8 and kernel 5.4.92.
2. Upgrade Ceph to version 15.2.5. Ceph daemon monitor, OSD and MDS, will upgrade to containers.
3. The new UVS Manager version 2.15.18

Prerequisite

1. Download two update files from Ambedded Help Center.
 - a. ceph-mars400.tar.gz
 - b. 2.15-18.upd
2. Please make sure that your cluster is in HEALTHY status.

We made a video to demonstrate the upgrade procedure. ([The YouTube video is HERE](#))

The steps to completely upgrade the software are

1. Upgrade the first Monitor node by node admin console.
2. Upgrade the rest of the monitor nodes by UVS manager.
3. Upgrade the OSD nodes.
4. Upgrade MDS nodes if you have MDS nodes in use.
5. Check the UVS software versions.
6. Upgrade the external RADOS gateways.
7. Upgrade the external iSCSI gateways.
8. Finally, please upgrade all of your ceph clients.
9. *** It is not necessary to keep the UVS manager open all the time during the update. The updater works in the background. You can close the UVS manager and reopen it to check the result later. ***

If a node fails to upgrade

During the upgrade, the UVS updater will reboot the nodes twice to replace the kernel and CentOS. The time for updating a node is about 10 minutes or longer, depending on the cluster's workload while updating.

The updater will do a rolling upgrade one node by one node. You shall not experience a service stop during the software upgrade. If you meet a failure while upgrading a node, the upgrade program will stop upgrading if the node's upgrade fails. You can skip the failed node and keep upgrading other nodes.

Redo the upgrade on the node that failed to upgrade. If you still get an upgrade failure on the node after repeating the update, please contact Ambedded to get the support.

Step1: Upgrade the First Monitor Node

1. Copy two update files (ceph-mars400.tar.gz and 2.15-18.upd) to one of your monitor node's directory **/var/www/html/mars200update**.
2. Log in to the BMC of the Mars 400 that hosts the monitor node.
3. Connect the node console by entering "C" and the node number.
4. Login that monitor node with user admin (default pass is admin) to execute Update. Enter the IP address of this node. **(Figure 1)**
5. Wait for the updater to complete the upgrade on this node. The update will take totally around 12 minutes.
6. The first stage is to upgrade the kernel, CentOS and Ceph. When you see the console output as Figure 2, you can log in with the root user to check the Ceph version. (Command: **ceph versions**) (Figure 3)
7. It will take several minutes to convert the Ceph daemons on this node to Ceph containers. You can use the command "**podman ps**" to check if we have converted daemons to containers. Check the Ceph version of Ceph services running on the node with the command: **ceph versions**. (Figure 3)
8. If you see all the containers are up and running, the upgrade is complete.

Figure 1.

Note: Fill the UPD File URL with the IP address of the current node.

```
Node: node21-1072 (Static)
Node Number(Board ID): 0

Main Menu
  0. Quit
  1. Shutdown
  2. Reboot
  3. Change Password
  4. Network Configuration
  5. Change Time Zone
  6. Update
  7. Network Diagnostics
  8. CEPH Diagnostics
  9. Default Restore Check
  d. Disk Diagnostics
[Enter Option : 6
Input Update Server IP or
[UPD File URL (http:// or https:// ) :192.168.1.111
UPD File Name ? (default is 2.15.19.upd): 2.15-18.upd
```

Figure 2.

```

MARS400
LVS Manager

*****
*
*   This system is for the use of authorized users only.
*
*
*****
Date: 10:37:10 Tue Mar 2 2021
Kernel 5.4.100-MARS400 on an aarch64
Last Restore return Pass

[node111-1072 login: root
[Password:
```

Figure 3.

```
[root@node111-1072 ~]# podman ceph versions
Error: unrecognized command 'podman ceph'
Try 'podman --help' for more information.
[root@node111-1072 ~]# ceph versions
{
  "mon": {
    "ceph version 14.2.12 (2f3caa3b8b3d5c5f2719a1e9d8e7deea5ae1a5c6) nautilus (stable)": 2,
    "ceph version 15.2.5 (2c93eff00150f0cc5f106a559557a58d3d7b6f1f) octopus (stable)": 1
  },
  "mgr": {
    "ceph version 14.2.12 (2f3caa3b8b3d5c5f2719a1e9d8e7deea5ae1a5c6) nautilus (stable)": 2,
    "ceph version 15.2.5 (2c93eff00150f0cc5f106a559557a58d3d7b6f1f) octopus (stable)": 1
  },
  "osd": {
    "ceph version 14.2.12 (2f3caa3b8b3d5c5f2719a1e9d8e7deea5ae1a5c6) nautilus (stable)": 4
  },
  "mds": {
    "ceph version 15.2.5 (2c93eff00150f0cc5f106a559557a58d3d7b6f1f) octopus (stable)": 1
  },
  "rgw": {
    "ceph version 15.2.5 (2c93eff00150f0cc5f106a559557a58d3d7b6f1f) octopus (stable)": 1
  },
  "overall": {
    "ceph version 14.2.12 (2f3caa3b8b3d5c5f2719a1e9d8e7deea5ae1a5c6) nautilus (stable)": 8,
    "ceph version 15.2.5 (2c93eff00150f0cc5f106a559557a58d3d7b6f1f) octopus (stable)": 4
  }
}
[root@node111-1072 ~]# podman ps
CONTAINER ID   IMAGE                                COMMAND                                CREATED        STATUS        PORTS        NAMES
aad007a8da3f   docker.io/ceph/ceph:v15             -n mon,node111-10... 7 minutes ago  Up 7 minutes ago        ceph-46473832-f350-4ab9-8c5a-9d850b612fab
-mon,node111-1072
3fe128e87756   docker.io/ceph/ceph:v15             -n mgr,node111-10... 7 minutes ago  Up 7 minutes ago        ceph-46473832-f350-4ab9-8c5a-9d850b612fab
-mgr,node111-1072
0e57805dd0db   docker.io/ceph/ceph:v15             -n mds,node111-10... 7 minutes ago  Up 7 minutes ago        ceph-46473832-f350-4ab9-8c5a-9d850b612fab
-mds,node111-1072
[root@node111-1072 ~]#
```

Step 2. Upgrade other monitors by UVS manager

1. Open the browser to login to the New UVS manager on the first monitor already upgraded to v2.15.18.
2. Navigate to the Settings → Firmware Update page. Select the monitor nodes that need updating.
3. Click **Push Update** to upgrade the monitors. (Figure 4)
4. Select the rolling update and select the 2.15-18.upd file.
5. Wait for the upgrade to complete. Every monitor will take 12 minutes to complete the upgrade. (Figure 5)

Figure 4.

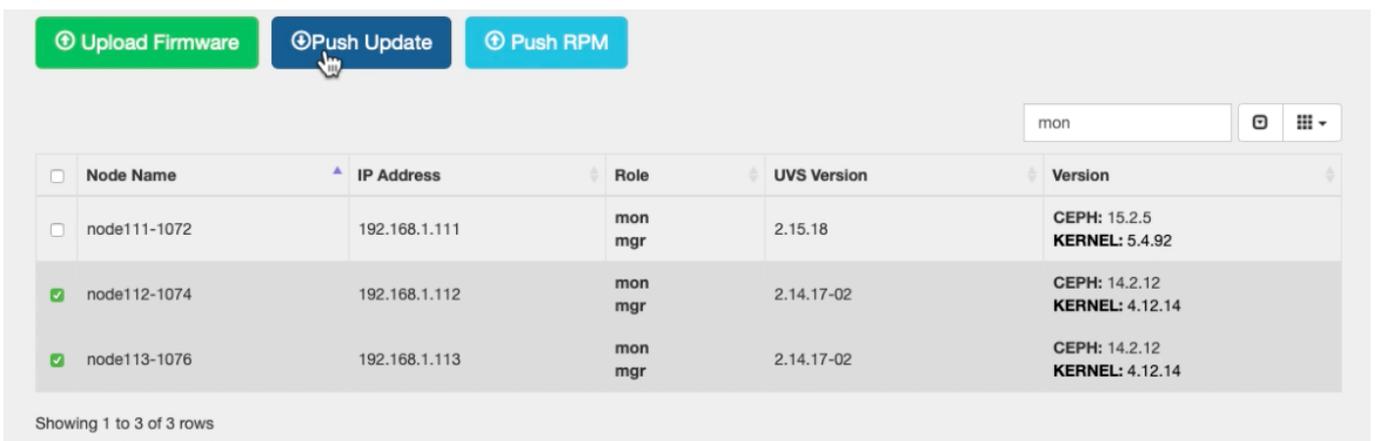


Figure 5.

```

PLAY [Restart ETCD Refresh UVS Daemon] *****

TASK [Gathering Facts] *****
fatal: [node112-1074]: FAILED! => {"ansible_facts": {}, "changed": false, "failed_modules": {"setup": {"failed": true, "module_stderr": "Warning: Permanently added 'node112-1074,192.168.1.112' (ECDSA) to the list of known hosts.\r\n/bin/sh: /usr/bin/python: No such file or directory\n", "module_stdout": "", "msg": "The module failed to execute correctly, you probably need to set the interpreter.\nSee stdout/stderr for the exact error", "rc": 127}}, "msg": "The following modules failed to execute: setup\n"}
fatal: [node113-1076]: FAILED! => {"ansible_facts": {}, "changed": false, "failed_modules": {"setup": {"failed": true, "module_stderr": "Warning: Permanently added 'node113-1076,192.168.1.113' (ECDSA) to the list of known hosts.\r\n/bin/sh: /usr/bin/python: No such file or directory\n", "module_stdout": "", "msg": "The module failed to execute correctly, you probably need to set the interpreter.\nSee stdout/stderr for the exact error", "rc": 127}}, "msg": "The following modules failed to execute: setup\n"}
ok: [node111-1072]

TASK [Reload Daemon] *****
changed: [node111-1072]

TASK [restart services] *****
changed: [node111-1072] => (item=etcdrefreshuvs)

PLAY RECAP *****
node111-1072      : ok=3    changed=2    unreachable=0    failed=0    skipped=0    rescued=0    ignored=0
node112-1074      : ok=0    changed=0    unreachable=0    failed=1    skipped=0    rescued=0    ignored=0
node113-1076      : ok=0    changed=0    unreachable=0    failed=1    skipped=0    rescued=0    ignored=0

Exit status : 2
Check if OSDs are all upgraded...
Check if MONs are all upgraded...
All MONs are upgraded to Octopus...
Check if MDSs are all upgraded...

Some Nodes are not Upgraded ...System will be Stable Until ALL Nodes are upgra

```

[Back](#)

Step 3. Upgrade OSD nodes.

1. Select OSD nodes you want to upgrade.
2. We recommend upgrading OSDs in the same chassis by one push.
3. Click **Push Update**.
4. Wait for the update to complete.
5. Repeat these steps until you upgrade all OSDs.

Step 4. Upgrade MDS nodes

1. If you have MDS in use, please follow the same steps to upgrade.

Step 5: Check the software version after upgrade.

After you complete the upgrade, please check each node's firmware version by navigating to the Settings → Firmware Update.

The correct software versions after this upgrade shall be

1. UVS: 2.15.18
2. Ceph : 15.2.5
3. Kernel: 5.4.92

The screenshot shows the UniVirStore Manager interface. At the top, there are navigation links: Dashboard, CEPH, Ceph FS, Object Storage, Storage, OpenStack, and Settings. Below these are three buttons: Upload Firmware, Push Update, and Push RPM. A search bar is located on the right. The main content is a table with the following data:

Node Name	IP Address	Role	UVS Version	Version
node21-1072	192.168.1.111	mon mds mgr	2.15.18	CEPH: 15.2.5 KERNEL: 5.4.92
node22-1074	192.168.1.112	mon mds mgr	2.15.18	CEPH: 15.2.5 KERNEL: 5.4.92
node23-1076	192.168.1.113	mon mds mgr	2.15.18	CEPH: 15.2.5 KERNEL: 5.4.92
node24-1078	192.168.1.114	osd	2.15.18	CEPH: 15.2.5 KERNEL: 5.4.92
node25-107a	192.168.1.115	osd	2.15.18	CEPH: 15.2.5 KERNEL: 5.4.92
node26-107c	192.168.1.116	osd	2.15.18	CEPH: 15.2.5 KERNEL: 5.4.92
node27-107e	192.168.1.117	osd	2.15.18	CEPH: 15.2.5 KERNEL: 5.4.92
node28-1080	192.168.1.118	mds	2.15.18	CEPH: 15.2.5 KERNEL: 5.4.92

Showing 1 to 8 of 8 rows

Step 6: Upgrade the RADOS Gateway.

If you are using the RGW on Mars 400 nodes, the updater will update the internal RADOS Gateway daemons collocated with monitor nodes. They will work well as they did before upgrading. However, these RGWs are not converted to containers yet. You will still get a warning like the following message. You can ignore these warnings or get help from Ambedded to convert them to containers.

```
HEALTH_WARN 1 stray daemons(s) not managed by cephadm
[WRN] CEPHADM_STRAY_DAEMON: 1 stray daemons(s) not managed by cephadm
stray daemon rgw.node111-1072 on host node111-1072 not managed by cephadm
```

Updating pre-Octopus version RGW to Octopus containers has many considerations that we have to discuss with users. Ambedded offers a free RGW upgrade service to customers who are using Nautilus (UVS 2.14.17).

If you have either internal or external RADOS gateways in use, please contact Ambedded to support you to upgrade your RGW safely.

Step 7: Upgrade the external iSCSI Gateway.

During the firmware update, we've updated the internal iSCSI Gateway daemons collocated with monitor nodes. For the external iSCSI gateways, you could upgrade your Ceph packages manually. You don't have to re-deploy the iSCSI sessions.

- CentOS: # yum update ceph-common
- Debian/Ubuntu: # sudo apt upgrade ceph-common

Upgrade Note

- Upgrade will abort if the M.2 SSD size is less than 64 GB. You need support from Ambedded to make the upgrade for you.
- If the bonding mode is 2, the UVS update will not change the bonding. If the bonding mode is 6, the UVS updater will change bonding to OVS SLB bonding.
- BMC must upgrade to the latest version for supporting configure new bonding.