

UVS Manager v2.14.16 Release Notes

Major Changes from v2.12.15	2
Bugs Fixed	3
New network bonding mode option:	3
Mars 400 bonding mode 2 (Node network configured by UVS v2.12.15 and earlier versions)	4
Mars 400 with bonding mode 6 (Node Network Configured by UVS v2.14.16)	4
Feature Highlights	4
Upgrade Guide	9
Upgrade from Version 2.12.15 to 2.14.16	9
Step1: Upgrade the First Monitor Node	10
Step 2: Upload the Temporary UVS License	11
Step 3: Upgrade Other Monitor and OSD	11
Step 4: Upgrade the MDS	12
Step 5: Check the software version after upgrade.	12
Step 6: Upgrade the external RADOS Gateway	13
Step 7: Request a new UVS manager license from Ambedded	13
Appendixes	15
Upgrade from UVS version before v2.12.15	15
Upgrading From Version Before 2.0.13	15
Upgrading From Version 2.0.14 to 2.12.15	15

Change Logs

Document Version	Released Date	Description	Remarks
V1.1	23-March-2021	1. Add Upgrade Guide Step 1-1: add patch to avoid timeout while updating OSD.	---

We are glad to announce the release of UVS manager v2.14.16. This is our first release for Ceph Nautilus v14.2.5. There are a lot of changes to meet the Ceph Nautilus release. This UVS manager release adds many new features and enhancements. Please read this release note carefully before you upgrade the Ceph cluster.

Major Changes from v2.12.15

- Upgrade Ceph to Nautilus. Some awesome features are listed as follows. Please refer to the [Ceph Nautilus Release Notes](#).
 - New Ceph dashboard
 - You can decrease the number of placement groups at any time, and the cluster can automatically tune the PG count.
 - New erasure code plugin CLAY is available for reducing the recovery time.
- **Kernel Change**
 - [To solve the issue of I/O will occasionally halt for 1 min due to the TCP read, kernel update to new version 4.12.14. Nodes will be rebooted during the upgrade process.](#)
- **UVS manager new features**
 - Deployment**
 - Add OSD with encryption
 - Parallely deploy multiple OSD to reduce the time for deployment
 - Safely remove OSD by moving OSD to trash before destroying
 - Control LEDs of storage device location for finding the devices easily
 - Change host network MTU (Maximum Transmission Unit) by UVS manager
 - RADOS Gateway**
 - Support SSL encryption on RADOS gateways (https)
 - Users are able to select the CRUSH rule while creating the RADOS Gateway
 - Automatically create all pools while creating RADOS gateway. This helps in planning the number of placement groups.
 - Ceph File System**
 - Automatically deploy standby MDS(Metadata Server) on Monitor nodes after creating active MDS.
 - Add MDS fail-back features for easily fail-back a fixed MDS to active mode.
 - CRUSH Map**
 - Users can edit and rename the CRUSH types to meet the data center infrastructure with up to ten levels of CRUSH map hierarchy.
 - CRUSH rule and Erasure code profile support selecting the device class
- **Other Enhancements**
 - Use ETCD distributed database to synchronize the data between multiple UVS manager services and reduce the web page loading time. Improve the UVS Manager Node Page load time to be much faster
 - Avoid UVS manager notification, users, and login being visible while the user is editing the password
 - Enable the failover of UVS notification
 - Add Generate and Download Diagnostic Log
 - Automate firmware upgrade with the flexibility of grouping multiple nodes to upgrade sequentially or parallely. The upgrade procedure is customized for each UVS version upgrade. Allow the user to use characters other than a-z as UVS user name.
 - Allow the user to use special characters other than a-z and up to 128 characters as the password
 - Improve the Ansible Execution Performance
 - Add Disk Diagnostics to admin Console

- Change the node ethernet bonding mode to mode 6 balance-alb
- Allow user to delete old update the upd and rpm files
- Upgrade PHP from version 5 to PHP7 to get better performance.

Bugs Fixed

- Create Multisite RGW will create extra pools if there is a standalone RGW exist. Correct the issue to use the same groups of pools.
- CephFS active MDS needs another standby MDS for failover. UVS shall automatically create standby MDS when the user creates the active MDS.
- Firmware Update fail if there is UNREACHABLE Ceph cluster's node
- UVS manager display date/time in the NTP page always displays UTC+0. It shall use the same time zone as the NTP server set in the system.
- Node Admin Console feature "Change Password" shall change the admin password instead of changing the root password.
- When the user opens the RGW user management page, the UVS manager shall not trigger the creation of the rgw pool.
- UVS manager shall show notification for duplicate inputs when the user creates multiple OSD or CRUSH map buckets at a time.

New network bonding mode option:

Since this UVS version, the admin console offers new bonding mode 6 balance-alb beside the old default mode 2 balance-xor.


Mode 6 uses adaptive load balancing and support failover. It does not require any special switch support. Mode 6 is simpler for switch configuration, and it is stabler compared to mode 2.

You don't have to make any change for your in production Mars 400 cluster.

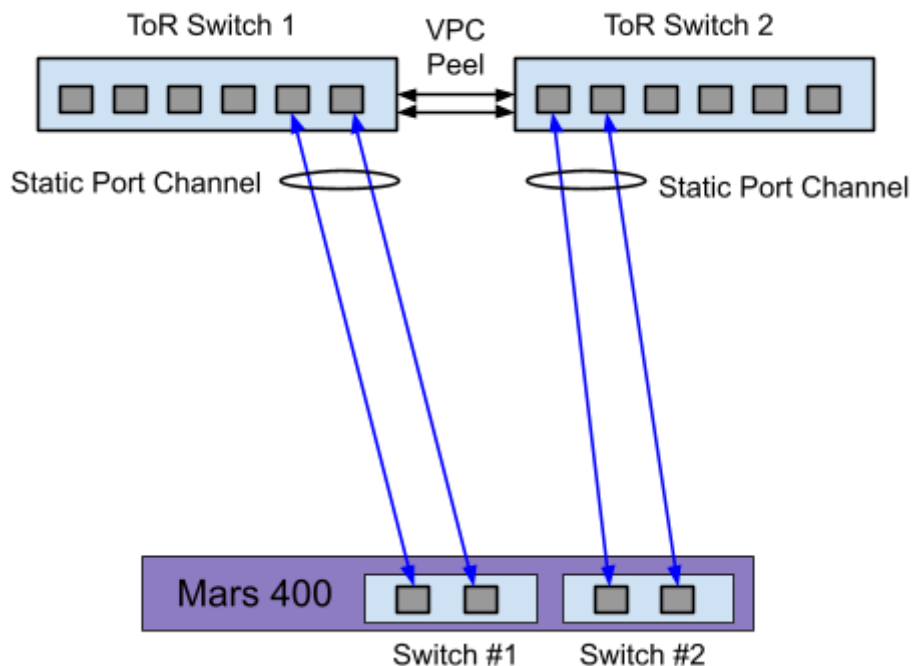
We recommend you to use bonding mode 6 for deploying new Mars 400 appliance. Mode 2 is still available for use.

The configuration of the network port on top of rack switch for Mar 400 with bonding mode 2 and mode 6 is different. Please refer to the following diagrams.

Mars 400 bonding **mode 2** (Node network configured by UVS v2.12.15 and earlier versions)



Mars 400 with bonding **mode 6** (Node Network Configured by UVS v2.14.16)



Feature Highlights

- **Using the ETCD distributed database to improve the UVS manager performance.**
Before v2.14.16, the UVS manager did not have a high availability database to keep the latest status of the Ceph cluster in the background. When the user switches the UVS manager to a new page, the UVS manager must query much information from all nodes. If the scale of the Ceph cluster becomes big, some pages in the UVS manager take much time to load. Introducing the ETCD database can automatically update the cluster status data. This new feature shortens the time to load a page a lot. The ETCD keeps the UVS manager status data

synchronized between all Monitor nodes. Users can get the same information no matter he uses the UVS manager on which Monitor node. Also, the notification service on the UVS manager is high-available now because of the utilizing ETCD.

- **Integrate CLAY erasure code to reduce OSD fail recovery time by 60%**

Ceph's default replication level provides excellent protection against data loss by storing three copies. However, storing 3 copies of data increases the cost of hardware as well as power consumption and cooling.

Erasure code offers a solution similar to RAID 6, which consume less raw capacity and provide the same level of data protection as replica 3. The drawback of erasure code is taking a much longer time to recover data stored in failed disks. Shortening the time to recover is a big motivation to use erasure code.


In this release, we introduce the new Clay erasure code plugin. A benchmark test of recovery time on replica 3, Jerasure, and Clay erasure code pools shows Clay code improve 62% of recovery time compare to Jerasure code.

The disadvantage of the traditional erasure codes is their long recovering time. Clay (Coupled Layer) code offers a simplified construction for decoding/repair for Ceph.


Clay code recovers OSD fail with much less time compared to Jerasure Code.

The following diagrams are the performance and recovery tests on 21x OSD with k=4, m=2, s=5. The performance test tool is RADOS bench.

Recovery Time



4K-READ & 4K-WRITE



4M-READ & 4M-WRITE



- **Parallelly create OSD for shortening the time for deployment**
Deploying multiple OSD by UVS manager become much faster. The new release of the UVS manager deploys multiple OSDs parallelly instead of sequentially. For example, deploying 20 OSD takes only 21 minutes now. It took 100 minutes in the past.
- **OSD encryption:** Users have the option to encrypt OSD while deploying new OSDs. The OSD encryption utilizes the dm-crypt of Linux kernel.

Add OSD

IP Address*: (One Per Line)

OSD Encryption*:

Yes

Journal on SSD*:

Yes

Select Chassis*:

mars400a

Add

- Allow user to define CRUSH type name.**

Ceph uses the CRUSH map and rule to define the data placement of a pool as a hierarchy. We use the CRUSH types to define the levels of the infrastructure of a hierarchy. The default CRUSH types are root, data center, room, row, pos, pdu, rack, chassis, and host.

However, the administrator may want to use CRUSH types other than the default types. For example, racks are located in cages; cages are located on floors, floors in a building, and buildings in a data center. Customize the types has to decompile and recompile the CRUSH map.

Before v2.12.15, users can create the CRUSH map with types of rack and chassis. Since v2.14.16, users can customize the CRUSH types easily in the web user interface with up to 10 levels of hierarchy. You can also rename the default buck types to any name suit for your real infrastructure.

CrushMap Tree

+ Create Bucket
Move Bucket
Crush F

root
region
zone
datacenter
room
pod
pdu
row
rack
chassis

Expand All
Collapse All


+ 11 default

Rename Bucket Type

Current Type Name	Modified Type Name
Type 11 : root:	root
Type 10 : region:	region
Type 9 : zone:	zone
Type 8 : datacenter:	datacenter
Type 7 : room:	layerd
Type 6 : pod:	layerc
Type 5 : pdu:	layerb
Type 4 : row:	layera
Type 3 : rack:	rack
Type 2 : chassis:	chassis
Type 1 : host:	host
Type 0 : osd:	osd

Update
Close

Example:



- **LED control on UVS manager**

Locate the failed disk drive shall be very careful. Mis-replace a faulty drive could make the situation even worse. In the next release of the UVS manager, the administrator could use the UVS manager to blink the LED of the chassis, and the LED beside the storage devices. This makes the replacement job easily and reduces the possibility of a human-made mistake.

+ Create New MON

Service Restart

Node Reboot

Search

<input type="checkbox"/>	Rank	Node Name	IP Address	Status	M.2 SSD SMART
<input type="checkbox"/>	mon.0	mycluster121-796d9	↑192.168.1.121	Up	<div></div> <div></div> <div></div>
<input type="checkbox"/>	mon.1	mycluster122-796d9	↑192.168.1.122	Up	<div></div> <div></div> <div></div>

Showing 1 to 2 of 2 rows

OSD

+ Create New OSD

Service Restart

Node Reboot

Remove OSD

Move To Recycle

Search

<input type="checkbox"/>	Rank			Class	Node Name	IP Address	PG	Status	Data SMART	MetaData SMART
<input type="checkbox"/>	osd.4			hdd	mycluster123-14d45	↑192.168.1.123	341	Up	<div></div> <div></div> <div></div>	<div></div> <div></div> <div></div>
<input type="checkbox"/>	osd.5			hdd	mycluster124-14d45	↑192.168.1.124	46	Up	<div></div> <div></div> <div></div>	<div></div> <div></div> <div></div>
<input type="checkbox"/>	osd.0			hdd	mycluster125-14d45	↑192.168.1.125	373	Up	<div></div> <div></div> <div></div>	<div></div> <div></div> <div></div>
<input type="checkbox"/>	osd.2			ssd	mycluster127-2bd0d	↑192.168.1.127	138	Up	<div></div> <div></div> <div></div>	<div></div> <div></div> <div></div>
<input type="checkbox"/>	osd.3			hdd	mycluster128-2bd0d	↑192.168.1.128	430	Up	<div></div> <div></div> <div></div>	<div></div> <div></div> <div></div>

- Enabling user to download the log through UVS manager for troubleshooting**
 It is quite often to collect necessary log and system status while supporting customers to find out the issues. Instead of manually collect logs and check Ceph status, Ambedded implements a new feature that can automatically collect pieces of information need for troubleshooting. No more remote access and manually collect data are required.

Upgrade Guide

Upgrade from Version 2.12.15 to 2.14.16

There four major software upgrades from v2.12.15 to 2.14.16. Please read this upgrade notes carefully before you start to upgrade your Mars 400 appliance software.

- Upgrade PHP version from PHP5 to PHP7:** PHP is the major software language for UVS manager front end web-based user interface. We have to upgrade one of the UVS managers running on the Monitor node before you can proceed with other upgrades.
- UVS manager:** Upgrade from version 15 to version 16 with many new features and bugs fixed.
- Upgrade the Ceph:** Version from Luminous 12.2.11 to Nautilus 14.2.5. The UVS manager upgrade assistant will automatically restart the Ceph daemons after push the upgrade package to the specified nodes.
- Update Linux kernel:** For the OS stability. We've upgraded the Linux kernel to 4.12.14. Hence, the UVS nodes have to reboot for loading the new kernel.

The steps to completely upgrade the software are

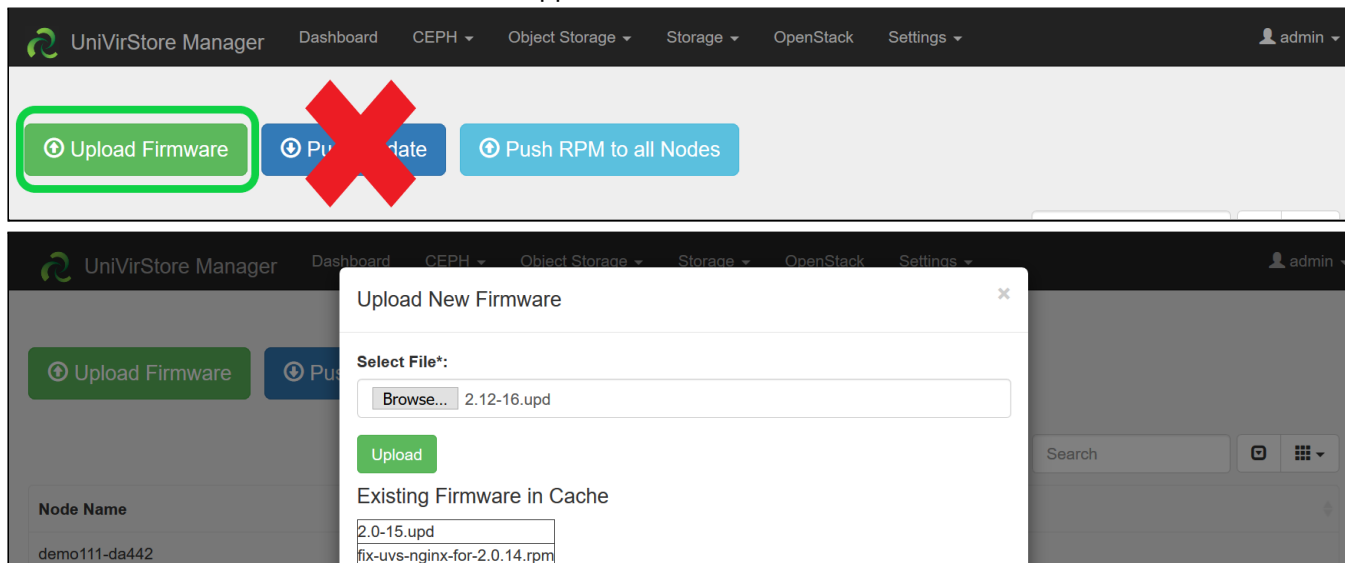
- Download the update files: 2.12-6.upd, updates.php & the temporary license
- Upgrade the first monitor node by login node admin console. The first monitor node means the monitor node you are going to use for the rest of upgrade steps by UVS manager.
- copy updates.php file to the first updated monitor node.
- Upgrade rest of monitors
- Upgrade OSD nodes
- Upgrade MDS nodes if you have MDS nodes in use.

7. Upgrade the external RADOS gateways
8. Finally, upgrade your Ceph clients

Please follow the following sections for details of each step. Also, we put a video about the upgradations on youtube, here is the link <https://youtu.be/h9KvTco6cJM>

Step1: Upgrade the First Monitor Node

1. Open your UVS manager on a browser.
2. For upgrading to UVS v2.14-16, we have to update the PHP version from version 5 to version 7. Please upload the "2.12-16.upd" package to this monitor node by using the UVS "Upload Firmware" feature.
3. After you upload the upd file, logout and close the UVS manager. PLEASE DON'T PUSH UPDATE BY UVS MANAGER AT THIS MOMENT.
4. Login admin console of the monitor node that you just used for uploading the upd file, by ssh admin@IP_ADDRESS on a terminal app.



5. Use the update function 6 and enter the IP address of this node to update the firmware. It may take 5 minutes to complete the upgrade.

```
Tue, 24 Dec 2019 03:29:18 +0000

MARS400 Version 2.12-15
Node: demo111-da442 192.168.1.111
Node Number(Board ID): 0

Main Menu
  Q. 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
Enter Option : 6
Update Server Name/IP? :127.0.0.1_
```

```

87150K ..... 99% 83.4M 0s
87200K ..... 99% 88.9M 0s
87250K ..... 99% 80.0M 0s
87300K ..... 99% 96.3M 0s
87350K ..... 99% 83.3M 0s
87400K ..... 99% 92.4M 0s
87450K ..... 99% 93.8M 0s
87500K ..... 99% 95.3M 0s
87550K ..... 100% 68.7M=5.5s

2019-12-24 03:30:06 (15.6 MB/s) - '/tmp/2.12-16.upd' saved [89670893/89670893]

Exit status : 0
Done
    Decrypting Update...Done
    Extracting Update...Done
    Running Pre Execution Task...Done
    Updating Files...Done
    Running Post Execution Task...[ 3757.664236] systemd: 61 output lines su
ppressed due to ratelimiting
Done
    Updated Successfully to Version 2.12-16...Done
    Final Cleanup...Done
Press enter key to continue....

```


Step 1-1:

1. Copy the updates.php file to the path **ceph-mon-node: /var/www/html/updates.php**
`$ scp updates.php ${ceph-mon-node}:/var/www/html/updates.php`
2. Change the owner to the "updates.php" from root to nginx
`$ chown nginx:nginx /var/www/html/updates.php`

Step 2: Upload the Temporary UVS License

Because the PHP of new UVS uses PHP7, the old UVS license for PHP5 can not work for the new UVS manager. You have to upgrade the UVS license too. Please upload a temporary license provided by Ambedded. The temporary license is valid until 2020 Jun. 30th. **You have to request a new license from Ambedded for UVS version 2.14.16 and future versions.**

1. Open the UVS manager on this node and navigate to the "Firmware Update" page.
2. Upload the temporary license. After checking the license, you shall see the first monitor node has been upgraded to UVS v2.14-16.



Step 3: Upgrade Other Monitor and OSD

After you upload the temporary license, you could use the **"push update"** function and update the remaining nodes by following the section [Firmware Update](#). We recommend pushing all the monitor nodes first at the same time. **Then rolling push the OSD nodes.** After Ceph monitor & OSD nodes have been upgraded, please push the MDS & remaining nodes at the same time.

1. Select the nodes that you want to update the firmware. You can enter the role name in the search box to filter the nodes you want to push update. For example, enter mon will show only monitor nodes on the list.

UniVirStore Manager Dashboard CEPH Object Storage Storage OpenStack Settings admin

Upload Firmware Push Update Push RPM

mon

<input checked="" type="checkbox"/>	Node Name	IP Address	Role	UVS Version	Version
<input type="checkbox"/>	demo111-09d7b	192.168.1.111	mon mgr	2.14-16.20200213	CEPH: 14.2.5 KERNEL: 4.12.14
<input checked="" type="checkbox"/>	demo112-09d7b	192.168.1.112	mon mgr	2.12-15.20190722	CEPH: 12.2.11 KERNEL: 4.19.52-MARS400
<input checked="" type="checkbox"/>	demo113-1be75	192.168.1.113	mon mgr	2.12-15.20190722	CEPH: 12.2.11 KERNEL: 4.19.52-MARS400

Showing 1 to 3 of 3 rows

2. Choose the push mode, parallel push or rolling push.

Push UPD to Selected Nodes

Select Push Mode*:

Rolling Push (Push One By One)

--Please Select Push Mode--

Rolling Push (Push One By One)

Step 4: Upgrade the MDS

If you have MDS in use, you shall also upgrade the MDS nodes after the OSD. Select the MDS nodes and push upgrade by Parallel Push. After the upgrade, UVS will also deploy standby MDS on all monitor nodes.

Step 5: Check the software version after upgrade.

After you complete the upgrade and reboot. Please check the firmware version of each node by navigating to the Settings -> Firmware Update.

The correct software versions after this upgrade shall be

1. UVS: 2.14.16
2. Ceph : 14.2.5
3. Kernel: 4.12.14

UniVirStore Manager
Dashboard
CEPH
Object Storage
Storage
OpenStack
Settings
admin

Upload Firmware
Push Update
Push RPM

Search

<input type="checkbox"/>	Node Name	IP Address	Role	UVS Version	Version
<input type="checkbox"/>	demo111-09d7b	192.168.1.111	mon mgr	2.14-16.20200213	CEPH: 14.2.5 KERNEL: 4.12.14
<input type="checkbox"/>	demo112-09d7b	192.168.1.112	mon mgr	2.14-16.20200213	CEPH: 14.2.5 KERNEL: 4.12.14
<input type="checkbox"/>	demo113-1be75	192.168.1.113	mon mgr	2.14-16.20200213	CEPH: 14.2.5 KERNEL: 4.12.14
<input type="checkbox"/>	demo114-6a5bf	192.168.1.114	osd	2.14-16.20200213	CEPH: 14.2.5 KERNEL: 4.12.14
<input type="checkbox"/>	demo115-9ce91	192.168.1.115	osd	2.14-16.20200213	CEPH: 14.2.5 KERNEL: 4.12.14
<input type="checkbox"/>	demo116-a55f1	192.168.1.116	osd	2.14-16.20200213	CEPH: 14.2.5 KERNEL: 4.12.14
<input type="checkbox"/>	demo117-fd9b1	192.168.1.117	osd	2.14-16.20200213	CEPH: 14.2.5 KERNEL: 4.12.14

Showing 1 to 7 of 7 rows

© 2018 Ambedded Technology Co. Ltd.; Taiwan. All Rights Reserved | UniVirStore Manager Version 2.14-16.20200213 for aarch64 | PHP Version 7.2.10 for aarch64

Step 6: Upgrade the external RADOS Gateway

During the firmware update of monitor nodes, we've updated the internal RADOS Gateway daemons collocated with monitor nodes.

If you have external RADOS gateways, you can upgrade the external RGW one by one by deleting and re-deploying them. The UVS manager "Delete" external gateway feature purges only the ceph information on RGW container; no RADOS data will be deleted.

Please refer to the section "[Creating RADOS Gateways on x86 servers](#)" for details about pulling the new docker image for Ceph Nautilus.

While re-creating RGW, you have options for non-SSL/SSL support. As the original RGW pools already exist, the new features of selecting CRUSH rules and the number of PG will be ignored while creating new RGW, as shown in Figure A & B. Once you re-create the new external gateway, users can start to access data via RGW.

Step 7: Request a new UVS manager license from Ambedded

Remember to request a new license from Ambedded as the temporary license is valid only before 2020 Jun. 30th.

Create External Rados Gateway

External RGW IP Address*:

192.168.1.245

Select Crush Rule*:

replicated_rule [Default]

Select RGW Bucket Data Pool PG*:

128

Figure A

SSL Encrytion Option*:

☒ No SSL Encryption

☐ Self Signed Certification

☐ Upload CA & SSL Key

Create

Close

Figure B

Appendixes

Upgrade from UVS version before v2.12.15


Upgrading From Version Before 2.0.13

If your UVS Manager version is 2.0.13 or older, please upgrade UVS to the next version of yours' in order before you upgrade to 2.14.16. For example, if your current version is 2.0.11, you have to upgrade version like 2.0.11 >> 2.0.12 >> 2.0.13 >> 2.0.14 >> 2.12.15 >> 2.14.16

Upgrading From Version 2.0.14 to 2.12.15

UVS Manager v2.12.15 is the next version of v2.0.14. In v2.12.15, we upgraded the **Ceph version** from 12.2.7 to 12.2.11 & **kernel version** from 4.16.9 to 4.19.52. It's a huge upgrading and please read the following instructions carefully.

1. First, you should **push the RPM patch** for v2.0.14 before you upgrade with the 2.0-15.upd file. Go to "**UVS Manager**" >> "**Settings**" >> "**Push RPM to all Nodes**" >> upload "**fix-uvn-nginx-for-2.0.14.rpm**"



2. Upload "**2.0-15.upd**" and push update by following the section Firmware Update. Once UVS has been upgraded, the UVS version should display 2.12-15, not 2.0.15.

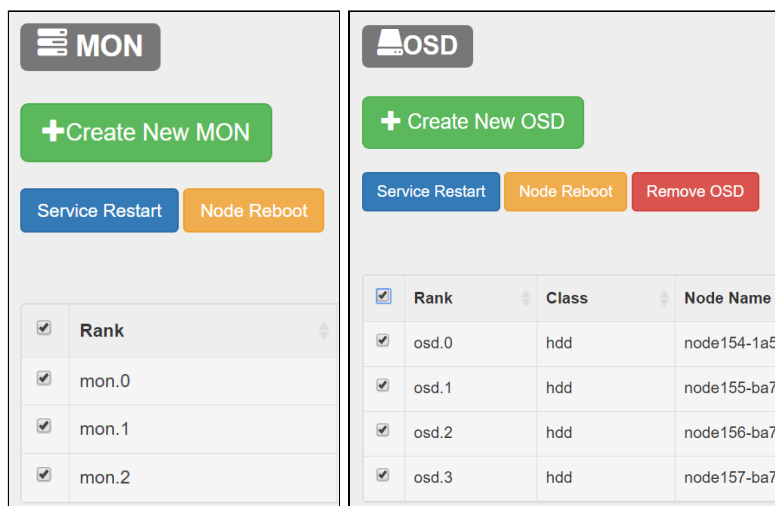
© 2018 Ambedded Technology Co. Ltd.; Taiwan. All Rights Reserved | UniVirStore Manager Version 2.12-15.20190722 for aarch64

3. After upgrading, you should **rolling reboot** all of the MON, OSD & MDS nodes in order to apply it. The sequence of the reboot is reboot MON first, then reboot OSD, lastly, reboot the MDS nodes. Please use the UVS manager node page to make the rolling reboot on MON and OSD nodes. The rolling reboot will reboot nodes one by one. After the first MON reboot completely and cluster become healthy again, the UVS manager will reboot the next MON and so on until all MONs are rebooted. The reboot of each monitor shall take around 1 to 2 minutes. During the reboot, the cluster will be temporarily unhealthy. If you have more than 3 MONs, the rolling reboot will not halt the client services.

Warning: While MON rebooting, **internal services** such as **iSCSI & RadosGW would be disconnected** for a few minutes. Please make sure your application data are stored and synchronized.

OSD rolling reboot is similar to the MON reboot. UVS manager will reboot the OSD one after the other. The duration of reboot one OSD till the cluster becomes healthy will be about five minutes. If you have more than 3 OSDs, the rolling reboot will not halt your client services. Please follow the steps in section 3.1 & 3.2. to reboot MON, OSD, and MDS.

- 3.1. For MON & OSD nodes, choose nodes by clicking the checkbox beside Rank and click "**Node Reboot**". Please complete the reboot of MON nodes before reboot the OSD node.



- 3.2. For MDS nodes, please follow this section [Access node console through BMC](#), login to MDS node as “**admin**” and choose “**2. Reboot**” in the main menu.

```

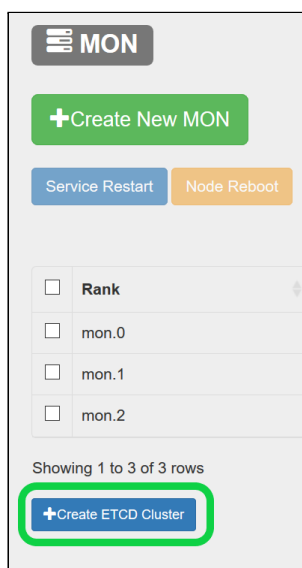
Wed, 24 Jul 2019 02:41:54 +0000

MARS400 Version 2.0-14
Node: node148-4f41d 192.168.1.148
Node Number(Board ID): 7

Main Menu
  Q. Quit
  1. Shutdown
  2. Reboot
  3. Change Password
  4. Network Configuration
  5. Change Time Zone
  6. Update
  7. Network Diagnostics
  8. CEPH Diagnostics
Enter Option : 

```

4. Since UVS Manager v2.12.15, ETCD cluster is used to accelerate UI operations. Please go to the “**Node**” page, and you’ll see the “**+Create ETCD Cluster**” button below MON Nodes. After you reboot all the Ceph cluster nodes, click it and install the ETCD Cluster.



5. After you complete the upgrade and reboot. Please check the firmware version of each node by navigating to the **Settings -> Firmware Update**.