Continue

Unisphere for vmax 9.0 installation guide

Dell EMC VMAX drives, VMAXISCSIDriver, and VMAXFCDriver support the use of Dell EMC VMAX storage arrays with the Cinder Block Storage project. Both provide equivalent functions and differ only to support their host e-mail methods. Drivers communicate with the back-end VMAX storage management software to perform batches. They use the HTTP Requests library to communicate with an Unisphere for a VMAX instance using the RESTAPI interface at the back end to perform VMAX storage operations. Note Workload support has been removed in PowerMax OS 5978. A fix queens is combined to support a PowerMax OS 5878 upgrade to a VMAX All Flash, so all functionality created from the existing OpenStack Instances and Volumes VMAX All Flash storage will work as before. Before you can upgrade a VMAX All Flash operating system, contact the OpenStack distributor to make sure that the fix is included. Dell EMC VMAX Cinder driver supports VMAX-3 hybrid series and VMAX All-Flash arrays. To run the Dell EMC VMAX Cinder driver, array operating system software, Solutions Enabler 8.4.0.x series (minimum 8.4.0.7) and Unisphere (minimum 8.4.0.0.15) are required for VMAX 8.4.0.x series. Solutions Enabler 9.0.x for PowerMax 9.0.x and Unisphere had a post attribute to support PowerMax OS in PowerMax and All Flash series. You can download Solutions Enabler and Unisphere from Dell EMC's support website (login required). For the VMAX 8.4.0 Setup Guide in support.emc.com, see the Solutions Enabler 8.4.0 Setup and Configuration Guide and Unisphere. There are five Dell EMC Software Packages sold with VMAX Hybrid arrays: Base Suite Advanced Suite Local Replication Suite Remote Replication Suite Total Productivity Suite Dell EMC VMAX Cinder driver requires Advanced Suite and Local Replication Suite or Total Productivity Pack (which includes Advanced Suite and Local Replication Suite). The Remote Replication Package is also required to use the VMAX Remote Replication function. For full functionality, including SRDF for VMAX All Flash, it is necessary to add fx package or F package plus SRDF a la carte. The storage system also requires an Unisphere for VMAX (SMC) eLicense. Each is individually licensed. For more information about how to obtain the relevant license(s), the following reference is eLicensing Support. To activate your authorizations and receive your VMAX license files, visit the Service Center at , as indicated in the License Authorization Code (LAC) letter emailed to you. For help with missing or incorrect rights after activation (otherwise, expected functionality is unavailable because it is not licensed), contact your EMC account representative or authorized resen agent. for help with an error License files through Solutions Enabler and contact Dell EMC Customer Support Center. If a LAC letter is missing or you want more instructions on how to activate your licenses through the Online Support site, contact EMC's worldwide Licensing team at licensing@emc.com or search: North America, Latin America, APJK, Australia, New Zealand: SVC4EMC (800-782-4362) and follow the voice prompts. EMEA: +353 (0) 21 4879862 and follow the voice commands. VMAX drivers support these operations: Volume creation, list, delete, attach and delete volume snapshots Copy an image with a batch clone Create a volume of a unit (Host and storage-supported volume transition) Create a volume from the global volume group and delete the global volume group Create a general volume group (add and delete volumes) Create a Global volume group from source, SFSS/SFD SRDF/A and S THERDF Metro Quality of Service (QoS) Management and unmanaged volumes and snapshots VMAX drivers also support the following features: Dynamic target iSCSI IP address iSCSI multi-path support Oversubscription Service Level support SnapVX support Compression support (All Flash only) CHAP Authentication Not VMAX All Flash array Solutions Enabler 8.3.0.11 or later by compression associated with the default Service associated compression level. This means that volumes added to newly created storage groups will be compressed. Note Since the release of PowerMax in May 2018, Unisphere 9.0.0.6 for PowerMax may support new PowerMax features such as Duplication of Replicated volumes and Online Device Expansion. Volume Restore from Snapshot is also supported in all VMAX All Flash and PowerMax arrays. Download and install download solutions enabler .com VMAX Driver Integration, support.emc. You can install Solutions Enabler on a non-OpenStack host. Supported platforms include different flavors of Windows, Red Hat and SUSE Linux. Solutions Enabler can be installed on a physical server or as a Virtual Device (VMware ESX server VM). Also, starting with HYPERMAX OS Q3 2015, you can manage VMAX3 arrays using the Embedded Management (eManagement) container application. For more information, see the Solutions Enabler 8.4.0 Setup guide at support.emc.com. Note Before you can use VMAX drivers, you need to explore storage arrays. For more information, follow the instructions .com Solutions Enabler 8.4.0 Setup and Configuration Guide at support.emc. support.emc.com unisphere and install it. Unisphere can be installed in local, remote, or embedded configurations – that is, running Solutions Enabler on the same server; On a server connected to the Solutions Enabler server; or using eManagement application (including Solution Enabler and Unisphere for VMAX). See the Unisphere for the VMAX 8.4.0 Installation Guide at support.emc.com. Configure Block Storage in Cinder.conf Note For security and back-end uniformity, the use of the XML file for VMAX back-end configuration has been depreciation-ready in Queens. While xml file usage is still supported, an alert will be issued about its upcoming depreciation. VMAX parameter cinder.conf parameter Default Required Description RestServerIp san_ip Yes IP address Unisphere server RestServerPort san_rest_port 8443 Port Unisphere server RestUserName san_login 'admin' Yes Username Unisphere server RestPassword san_password Yes Password Unisphere server Array vmax_array No Yes Unique VMAX array serial number SRP vmax_srp No Yes Name storage resource pool PortGroups vmax_port_groups None Yes Name VMAX port group(s) SSLVerify driver_ssl_cert_verify driver_ssl_cert_path False None No my_unisphere_host.pem Note san_rest_path is 8443 by default but can be changed if necessary. The default default is considered for the purposes of this documentation, so that the label does not appear in any of the following cinder.conf statements. Note VMAX PortGroups must be preconfigured to expose array-managed volumes. Port groups can be provided in cinder.conf or specified as an extra type of spec storage in the volume type: portgroupname. The latter provides the user with more control. When a dynamic masking view is created by the VMAX driver, if the port group specified as an additional specification does not exist, the port group is randomly selected from the PortGroup list to evenly distribute loads to the set of supplied groups. Note The Service Level and workload can be added to cinder.conf when the back end is in the default state and there is no associated volume type. This is not a recommended configuration because it is very restrictive. VMAX parameter cinder.conf parameter Default Required ServiceLevel vmax_service_level No Workload vmax_workload No Configuration in Cinder.conf Block Storage /etc/cinder/cinder.conf: enabled_backends = CONF_GROUP_ISCSI, CONF_GROUP_FC [CONF_GROUP_ISCSI] volume_driver = cinder.volume.drivers.dell_emc.vmax.iscsi.VMAXISCSIDriver volume_backend_name = VMAX_ISCSI_DIAMOND vmax_port_groups = [OS-ISCSI-PG] san_ip = 10.10.10.10.10.10 san_login = my_username san_password = my_password vmax_array = vmax_service_level = 000123456789 vmax_srp = SRP_1 [CONF_GROUP_FC] volume_driver = cinder.volume.drivers.dell_emc.vmax.fc.VMAXFCDrive in volume_backend_name VMAX_FC_DIAMOND vmax_port_groups = san_ip 10.10.10.10 san_login = my_username san_password = my_password vmax_array = 000123456789 vmax_srp = SRP_1 example, two back-end configuration groups are enabled: CONF_GROUP_ISCSI and CONF_GROUP_FC. Each configuration group there is a section unique parameters for connections, volume_backend_name and data to be used. After Cinder.conf is updated, create Volume Types, openStack commands must be volume_backend_names create and associate openStack volume types with declared resources. Additionally, each volume type must be specified by an associated volume, which is an additional specification that indicates the combination of service level/workload to pool_name type. Storage type:There is also the option to assign a port group to a volume type by setting the portgroupname additional specification. Note It is possible to create as many volume types as the number of combinations of Service Level and Workload (available) for volume provisioning. pool_name must be set and is an additional property with a format: <ServiceLevel>+<Workload><SRP>+ +<Array id=>. This cinder get-pools can be obtained from output - detail. Create $openstack volume type VMAX_ISCSI_SILVER_OLTP $ openstack volume type set --property volume_backend_name=ISCSI_backend \ --property pool_name=Silver+OLTP+SRP_1+000123456789 \- -property storage type:portgroupname=OS-PG2 \ $VMAX_ISCSI_SILVER_OLTP create openstack volume type $VMAX_FC_DIAMOND_DSS openstack volume type set --property volume_backend_name=FC_backend \ --property pool_name=Diamond+DSS+SRP_1+000123456789 \ --property storage type:portgroupname=OS-PG1 \ VMAX_FC_DIAMOND_DSS By ising these commands, the Block Storage volume type is associated with VMAX_ISCSI_SILVER_OLTP ISCSI_backend, Silver Service Level, and an OLTP workload. VMAX_FC_DIAMOND_DSS type FC_backend associated with the Diamond Service Level and DSS workload. ServiceLevel manages basic storage to ensure expected performance. Setting ServiceLevel to Yok means that fast non-managed storage groups (storage groups that are not associated with any service level) will be created. If ServiceLevel does not exist, the Workload must be Destroyed. When the openstack volume set --property pool_name=None+None+SRP_1+000123456789 Workload is added, the latency interval decreases due to the information added. Setting workload to No means that the latency interval will be the widest range for the Service Level type. Please note that you cannot set up a Workload without a Service Level. The openstack volume type set --feature pool_name=Diamond+None+SRP_1+000123456789 Note VMAX Hybrid Optimized supports Diamond, Platinum, Gold, Silver, Bronze, and NONE service levels. VMAX supports All Flash, Diamond and None. Both support DSS_REP, DSS, OLTP_REP, OLTP, and None workloads. Seamless upgrades from an SMI-S-based drive to a RESTAPI-based drive follow the above setup instructions are supported by several exceptions: Live migration functionality does not work on older volumes that have already been added/are in use. These volumes must first be separated and re-sewn using RESTAPI-based</Array> </SRP> </ServiceLevel> openssl s_client -showcerts \ -connect my_unisphere_host:8443 /dev/null \ | openssl x509 -outform PEM > my_unisphere_host.pem Where my_unisphere_host is the name of the single-instance instance in hostname and my_unisphere_host.pem file. Add this path driver_ssl_cert_path cinder.conf under back-end stanza driver_ssl_cert_path = /path/to/my_unisphere_host.pem, or follow these steps: LINK TO IST (if step 2 is complete): Copy the pem file to the system certificate directory: #cp my_unisphere_host /us r/share/ca-certificates/ca_cert.crt Optional: Update the CA certificate database with the following commands: #sudo dpkg-ca-certificates Not Note Check that the new ca_cert.crt will activate it by selecting prompt in the dialog box. If activation is not enabled, use the up and down keys to select it and the space key to enable or disable it. Make sure that the path defined in #sudo update-ca-certificates driver_ssl_cert_verify cinder.conf backend stanza or step 1 is set to True. Note Problem SSLError(hostname does not match 'xx.xx' Solution Control these requests and dependencies are up-to-date: $ sudo pip installation requests --confirm that the upgrade SSL cert was created using command: $openssl s_client -showcerts -{connect my_unisphere_host:{port/dev/null|opens -x500 Form PEM >{ Verify using the cert_name}.pem Cert command: $openssl s_client -connect {ip_address}:{port} -CAfile {cert_name}.pem -verify 9 Install ipaddress if requests are considered up-to-date and the certificate is validated but the host error still occurs, install the host error: $ sudo pip install ipaddress Zone Manager is required when there is a fabric array. This is necessary for larger configurations, where pre-zoning will be very complex and open zoning will raise safety concerns. Make sure that the Iscsi-initiator-utils package is installed on all Process nodes. Note You can only ping VMAX iSCSI destination ports when there is a valid masking view. An insertion process creates this masking

view. Masking views are dynamically created by VMAX FC and iSCSI drivers using the following naming conventions. [protocol] is either I for connected volumes via Fibre Channel for attached volumes over iSCSI or F. OS-[shortHostName]-[protocol]-[portgroup_name]-A launcher group is created or re-used for each host attached to VMAX volumes using MV Drivers (by attachment type). All known types of initiators for this hosts included in the group. With each new insertion volume operation, the VMAX driver receives initiators (WWNs or IQNs) from OpenStack and adds or updates the contents of the Launcher Group as needed. The names are in the following format. [protocol] is either I for connected volumes via Fibre Channel for attached volumes over iSCSI or F. OS-[shortHostName]-[protocol]-IG Not Hosts that adds OPENStack managed VMAX storage cannot also be installed with the same VMAX storage that is not managed by OpenStack. VMAX-arrayed FA ports to be used in a new masking view are retrieved from the port group provided as extra specs of the audio type or selected from the list provided in the Dell EMC configuration file. As volumes are attached to a host, they are added to an existing storage group (if any), or a new storage group is created and a volume is added. Storage groups common created from a pool connected to a single host through a single connection type (iSCSI or FC). [protocol] is either I for connected volumes via Fibre Channel for attached volumes over iSCSI or F. The VMAX cinder driver uses cascading storage groups - a parent storage group associated with masking view that includes child storage groups for each configured SRP/slo/workload/compression enabled or disabled/compression enabled or disabled combination. VMAX All Flash and Hash Parent storage group: OS-[shortHostName]-[protocol]-[portgroup_name]-SG Child storage group: OS-[shortHostName]-[SRP]-[ServiceLevel/Workload]-[portgroup_name]-CD-RE Not CD and RE are set only if compression is explicitly disabled or replication is explicitly enabled. See compression and replication sections below. By default, intervals and re-attempts have 3 seconds and 200 re-attempts, respectively. These determine how long (ranges) and how many times (re-attempts) a user wants to wait for a single Rest call, 3*200=600seconds. Depending on the use, this cinder.conf may need to be overridden by the user. For example, if performance is a factor, the interval should be reduced to check the job status more frequently, and if multiple concurrent provisioning requests are granted, re-attempts should be increased so that calls do not run out of time. In the following example, the driver checks the status of the job every 3 seconds. It will continue to check 200 re-attempts before time is out. Add the following lines to the VMAX back end in Cinder.conf: [CONF_GROUP_ISCSI] volume_driver = VMAX_ISCSI_DIAMOND vmax_port_groups = [OS-ISCSI-PG] san_ip = 10.1.10.10.10 san_login = my_username san_password = my_password vmax_array = 000123456789 vmax_srp = SRP_1 range = 3 repee times = 200 Quality of Service (QoS) has traditionally been associated with network bandwidth usage. Network limitations on specific networks in terms of bandwidth usage for clients. This ensures that they provide a cost-based tiered level of service. Nova/cinder QoS offers similar functionality based on unit type setting limits for host storage bandwidth per service offer. Each volume type depends on specific QoS attributes, some of which are specific to each storage vendor. In hypervisor, QoS limits the following: Total bytes/s, bytes/s arrows, bytes/s Limit by IOPS - Total IOPS/s, reading IOPS/s, IOPS/s in Cürse, IOPS/s QoS application writing in Cürse, hypervisor (front end), storage sub-system (back end) or both. This section discusses the QoS limits applied by the VMAX back end and hypervisor front end instead of each other or only by the rear end (Vendor Special). The VMAX driver offers support for the total byte/s limit and the IOPS's Total IOPS/s limit. The VMAX driver supports the following attributes total_iops_sec-end/rear-end agnostic system - Maximum IOPs (in I/Os per second). Valid values range from 100 IO/S to 100,000 IO/s total_bytes_sec - Maximum bandwidth (throughput) per second bytes. Valid values range from 1048576 bytes (1MB) to 104857600000 bytes (100,000MB), the VMAX driver offers the following attribute, which depends on the VMAX-specific vendor and the set total_iops_sec. Dynamic Distribution - Provides/disables dynamic distribution of host I/O limits. Possible values include: Always - Provides full dynamic deployment mode. When enabled, configured host I/O boundaries are dynamically distributed across configured ports, allowing the boundaries of each port to be adjusted to fluctuating demand. OnFailure - Provides port error. When enabled, fractions of the configured host I/O limits available for a configured port are set according to the number of ports currently online. Never - This feature is disabled (Default). Prerequisites - VMAX Host I/O Limit (IO/sn) - Unlimited Set Dynamic Deployment - N/A Prerequisites - Block Storage (cinder) back end (storage group)) Key Value total_iops_sec 500 total_bytes_sec 104857600 (100MB) DistributionType Always create Qos Properties with the above required values: $ openstack volume qos create --consumer background end \ -- property total_iops_sec=5 000 \ --property total_bytes_sec=104857600 \ --property DistributionType=Always \ my_qos Common QoS properties with specified volume type: $ openstack volume qos association my_qos my_volume_type Create volume with the volume type mentioned above: $ openstack volume generation -- size 1 --type my_volume_type my_volume Result - VMAX (storage group) Host I/O Limit (MB/Sn) - 100 Host Limit (IO/Sn) - 500 Sets Dynamic Distribution - Always Result - Block Storage (cinder) Volume Volume against volume type and QoS is applied with the above parameters. Prerequisites - VMAX Host I/O Limit (IO/Sec) - 2000 Host I/O Limit (IO/Sec) - 2000 Set Dynamic Deployment - Always Prerequisites - Block Storage (cinder) back end (storage group)) Value total_iops_sec 500 total_bytes_sec 104857600 (100MB) DistributionType OnFailure Never Prerequisite - Block Storage (cinder) back end (storage group), total_iops_sec 500 total_bytes_sec 500 total_bytes_sec=1048576 00 \ -- feature DistributionType=Always \ my_qos Associate QoS properties with the specified volume type: $ openstack volume qos association my_qos my_volume_type Create volume with the volume type mentioned above: $ openstack volume generation with the volume type mentioned above: $ openstack volume generation -- size 1 --type my_volume my_volume Add volume created in 3 steps to an instance $ openstack server my_volume my_instance Result - VMAX (storage group) Host I/O Limit (MB/Sn) - 100 Host I/O Limit (IO/Sn) - 500 Sets Dynamic Distribution - Always Result - Block Storage (cinder) Volume is created by volume type and QoS is applied with the above parameters. Conclusion - Hypervisor (nova) Libvirt contains an extra XML flag <disk> in the section called iotune, which is responsible for the rate limitation. To verify this, first get the OS-EXT-SRV-ATTR:instance_name value of the server instance, instance-00000003. $ openstack server show <serverid>

+----------------------------------+--+ | Field | Value | +----------------------------------+--+ | OS-DCF:diskConfig | AUTO | | OS-EXT-AZ:availability_zone | nova | | OS-EXT-SRV-ATTR:host | myhost | | OS-EXT-SRV-ATTR:hypervisor_hostname | myhost | | OS-EXT-SRV-ATTR:instance_name | instance-00000003 | | OS-EXT-STS:power_state | Running | | OS-EXT-STS:task_state | Not there | | Os-EXT-STS:vm_state | active | | Os-SRV-USG:terminated_at | Not there | | accessIPv4 | | | accessIPv6 | | | custom=fd21:99c2:73f3:0:f816:3eff:febe:30ed, 10.0.0.3 | | config_drive | | created | 2017-11-02T08:15:34Z | | flavor | m1.tiny (1) | | hostId | e7b8312581f9fbb8508587d45c0b6fb4dc86102c632ed1f3a6ad9d42 | | id | 0ef0ff4c-dbda-4dc7-b8ed-45d27fc2f31db | image | cirros-0.3.5-x86_64-disk (b7c220f5-2408-4296-9e58-fc5a41cb7e9d) | key_name | myhostname | name | myhostname | | progress | 0 | project_id | bae4b97a0d8b4c28a5add483981e5db | features | | | security_groups | name='default' | | status | EVENT | | updated | 2017-11-02T08:15:42Z | | user_id | 7bccf456740546799a7e20457f13c38b | volumes_attached | | +----------------------------------+--+

<serverid><disk/disk>We then run the following command using the OS-EXT-SRV-ATTR:instance_name taken above. $ virsh dumpxml example-00000003 | grep -1 total_bytes_sec|total_iops_sec The output of the command contains the following xml. The start and end tag include <disk> . <iotune> <total_bytes_sec></total_bytes_sec> <total_bytes_sec>500</total_bytes_sec> </iotune> Prerequisite - VMAX Host I/O Limit (MB/Sn) - 100 Host I/O Limit (IO/Sn) - 500 Sets Dynamic Deployment - Always Prerequisites - Block Storage (cinder) back end (storage group)) Value total_iops_sec 500 total_bytes_sec 104857600 (100MB) DistributionType OnFailure Create QoS specification with the above prerequisite values: $ openstack volume qos create --consumer back end \ -- property total_iops_sec=1048576 00 \ --property DistributionType=Always \ my_qos Associate QoS specifications with the specified volume type: $ openstack volume qos association my_qos my_volume Creating volume with the volume type mentioned above - $ openstack volume generation --size 1 --type my_volume_type my_volume Result - VMAX (storage group) Host I/O Limit (MB/S) - 100 Host I/O Limit (ICSec) - Unlimited Set Dynamic Distribution - N/A Prerequisites - Block Storage (cinder) back end (storage group)+Key Value Distribution Display Always create QoS specifications with the above prerequisite values: $ openstack volume qos create --consumer back end \ --property DistributionType=Always \ my_qos Associate QoS with the specified volume type: $ openstack volume qosos association my_qos my_volume_type Create volume with the volume type mentioned above : $ openstack volume generation -- size 1 --type my_volume_type my_volume Outcome - VMAX (storage group) Host I/O Limit (MB/Sn) - Unlimited Host I/O Limit (IO/Sn) - Unlimited Set Dynamic Distribution - N/A Result - Block Storage (cinder) It is created by volume type and there is no QoS exchange. Install open iscsi on all nodes on your system Do not install EMC PowerPath as native multipath software cannot exist with multipath tools Ubuntu must have all nova computing nodes installed: #apt-get multipath-tools #multipath-tools #multipath modules install multipath-tools #multipath modules install sysfsutils sg3-utils #file system programs apt #-get scsis #SCSI openSUSE and SUSE Linux Enterprise Server tools: #zipper install multipath-tools #multipath modules #zipper install sysfsutils sg3-benefits #file system utilities #zipper install scsitools #SCSI tools Red Hat Enterprise Linux and CentOS: #yum install #ensure iSCSI yükli # cihaz-mapper-multipath #multipath modülleri yüklemek # </disk> <disk> Installing system utilities with sysfsutils sg3-utils #file page=pre-spc3-83 --whitelisted --device=/devin path_selector round-robin 0 path_checker tour features 0 hardware, handler 0 prio const r_weight uniform no_path_retry 6 rr_min_io 1000 rr_min_io_rq 1 } } You may need to restart the host or restart the iSCSI and multipath services after installing the MPIO tools. Ubuntu day: #service on-iscsi re-service #service multipath-tools openSUSE, SUSE Linux Enterprise Server, Red Hat Enterprise Linux and CentOS: #systemctl restart's open-iscsi tools $ lsblk Name MAJ:MIN RM SIZE RO TYPE MOUNTPOINT sda 8:0 0 1G 0 disk ... 3600009700001960701868533030303030303235 (dm-6) 252:6 0 1G 0 mpath vda 253:0 1T 0 disk Computer (nova) node, Add the following flag to the [libvirt] section of /etc/nova/nova-cpu.conf: volume_use_multipath = True on cinder controller node, iSCSI MPIO can be set globally in [DEFAULT], or after change can be changed in the VMAX back-end cluster /etc/cinder/cinder.conf: use_multipath_for_image_xfer = True Restart nova-compute and cinder-volume services. Create a 3GB VMAX volume, confirm that you have multiple initiators on the compute node for I/O. Create an instance from an image outside of local LVM storage or from VMAX storage, for example, connect 3GB volume from a launchable volume to the new sample: #multipath -ll mpath1000196700053153303030303039 (dm-6) 252:6 0 3G 00 disk . 3600009700001960700531533030303030383039 (dm-6) 252:6 0 3G 0 mpath sdb 8:16 0 3G 0 disk ... 3600009700001967001868533030303038309 (dm-6) 252:6 0 3G 0 mpath sdc 8:16 0 3G 0 disk . 3600009700001960700531533030303030383039 (dm-6) 252:6 0 3G 0 mpath vda This VMAX back-end into one-way initiator supports CHAP authentication functionality. With CHAP one-way authentication, the storage array expects to receive a valid credential and CHAP secret in response to the host challenge and exchange during the initial connection negotiation. When challenged, the host provides chap credentials and CHAP secrets to the storage array. Storage array, this ID stored in information in the host launcher group (IG) in iSCSI and CHAP searches for his secret When positive authentication occurs, the storage array sends an acceptance message to the host. However, if the storage array cannot find any records of the credential/hidden pair, it sends a rejection message and the connection is closed. Assumptions, Constraints, and Prerera terms: Host initiator IQN is required along with the credentials that the host launcher will use to log on to the storage sequence. The same credentials should be used on a multi-node system that connects to the same array. Enable one-way CHAP authentication for the iSCSI initiator in the storage sequence using SYMCLI. The template and example are shown below. For this installation purpose, the credentials used/hidden iqn.1991-05.com.company.lcseb130 # symaccess -sid -iscsi enable chap | disabled | chap set chap -cred -secret #symaccess will be my_username/my_password Settings and Configuration# Set up and restart cinder in the cinder.conf VMAX back group using the following parameters. Chap_use_chap_auth True Yes chap_password my_password The configuration options required for Yes chap_password my_password Value Yes chap_username my_username Image_volume_cache_enabled VMAX_ISCSI_DIAMOND Yes chap_password = True value_driver = zero volume_driver = cinder.volume.drivers.dell_emc.vmax.iscsi.VMAXISCSIDriver volume_backend_name = VMAX_ISCSI_DIAMOND san_ip = 10.10.10.10 san_login = my_u4v_username san_password SRP_1 vmax_array = my_u4v_password vmax_srp = 0001234456789 vmax_port_groups = [OS-ISCSI-PG] use_chap_auth = Correct chap_username = my_username chap_password = using my_password using symcli Enable CHAP authentication for a host launcher as described above, but cinder.conf use_chap_auth, chap_username't chap_password the phone. Create a bootable volume. openstack volume creation --size 1 \ --image image_name>\ <volume_type> -- type \test Boot instance named test_server using the volume created above: create an openstack server --volume test \ --flavor m1.small \ --nic net-id=private \ test_server Verify that the audio operation succeeded, but the boot instance fails when CHAP authentication fails. Update Cinder.conf with a chap_username set use_chap_auth correct and correct credentials chap_password correct credentials. Run the Openstack server again confirm that the boot instance process is running correctly and that the volume has been reached. You have successfully performed both volume and boot instance operations and confirm that the user has access to the volume. In the entire Flash series, the creation of a storage group has a compressed attribute by default. Setting compression to a storage group does not mean that all devices will be compressed immediately. This means that all incoming writes will be considered for compression. This feature applies to all Flash 250F, 450F, 850F and PowerMax 2000, 8000 arrays. Note Since the release of PowerMax in May 2018, Unisphere 9.0.0.6 for PowerMax may support Replication. Compression and Replication go hand in hand, so if Compression is enabled, there is also deduplication. Compression disabling is also enabled by default. Replication is available only in the PowerMax directory. Use case 1 - Create a new volume type called compression deactivalization, VMAX_COMPRESSION_DISABLED, allocation, and deletion volume. volume_backend_name an extra spec on the storage. Set a new extra spec storage type:disablecompression = True. Create a new volume. Check Unisphere or symcli to see if the volume is in the storage group<srp>OS--<servicelevel>:-CD-SG and compression is disabled in that storage group<workload>. Plug the volume into an example. Check Unisphere or symcli to see if the volume is in the storage group<srp>OS--<servicelevel>:-CD-SG and compression is disabled in that storage group<workload>. Delete the volume. If this operating system is the final volume-<srp><servicelevel><workload>:CD-SG storage group, it must also be deleted. Repeat use case steps 1-4. Create a new volume type. For example, VMAX_COMPRESSION_ENABLED. Set a new extra spec compression as storage type:disablecompression = False or DO NOT set this extra spec. VMAX_COMPRESSION_DISABLED from VMAX_COMPRESSION_ENABLED to The New One. Check Unisphere or symcli to see if the volume is in the storage group OS-- SG<srp><servicelevel><workload>and compression is enabled in that storage group. Note Extra spec storage type:If the disable compression is set to a hash, compression is considered because VMAX3 is not a hybrid property. Configure the source and destination arrays, configure an SRDF group between the source and destination arrays selected for the volume. The source array must correspond to the cinder.conf 'vmax_array' entry (or input in the VMAX XML file for older I:Array>installations). Select both director and ports for SRDF emulsing. Network topology is important when selecting admin endpoints unutmayın. Desteklenen modlar Synchronous, Asynchronous ve Metro'dur. Not Kaynak ve hedef diziler aynı Unisphere sunucusu tarafından yönetilmiyorsa (diğer bir deyişle, hedef dizi uzaktan sunucuya bağlanır - örneğin, gömülü yönetim kullanıyorsanız),</Array>< <workload>; <servicelevel> <srp>; <workload>; <servicelevel> <srp>; <workload>; <servicelevel> <srp>; <portgroup> <servicelevel>; <srp>; <host-management>: <workload>; <servicelevel> <srp>; <srp> </srp>; In a complete disaster scenario (i.e. the primary array is completely lost and the entire connection is gone) event, the Unisphere server will no longer be able to communicate with the target array. In this scenario, volumes automatically fail on the target array, but administrator intervention will either be required; configure the destination (remote) array locally to the current Unisphere server (if it is a stand-alone server), or enter the details of a second Unisphere server into the locally connected cinder.conf for example, the Unisphere server, the embedded management of the destination directory), and then restart the cinder volume service. Enable replication in /etc/cinder/cinder.conf. To enable replication functionality on the VMAX cize drive, it is required to create a replication volume type. For this type of unit, the corresponding back-end quartet in cinder.conf must replication_device parameter that is later selected. This parameter defines a single replication target array and takes the form of a list of key value pairs. enabled_backends = VMAX_FC_REPLICATION [VMAX_FC_REPLICATION] volume_driver = cinder.volume.drivers.dell_emc.vmax.fc. VMAXFCDriver san_ip = 10.10.10.10 san_login = my_u4v_username san_password = my_u4v_password vmax_array = SRP_1 vmax_array = 000123456789 vmax_port_groups = [OS-OS-FC-PG] use_chap_auth = True chap_username = my_password chap_password = my_password VMAX_FC_REPLICATION replication_device = target_device_id 000197811111, rdf_group_label, remote_port_groups-failover-pg, remote_pool:SRP_1, rdf_group_label: 28_11_07, allow_extend:False, mode:Metro, metro_use_bias:False, allow_delete_metro:False ... Note:: Key value pairs replication_device "1000" must be in the same row (separated by comma). These are displayed in the reserved rows above to be ready. target_device_id is a unique VMAX array serial number of the target array. For full failed functionality, values must be discovered and managed by the same U4V server. remote_port_groups must prepare for the back end in the event of a port that this port group contains all FC or all iSCSI port groups (for a specific back end) in accordance with the configured drive (iSCSI or FC). remote_pool is the unique pool name of the given target directory. rdf_group_label name of a VMAX SRDF group that is preconfigured between the source and destination arrays. allow_extend flag that allows the extension of replicated volumes. To extend a volume in an SRDF relationship, it must first break that relationship, independently expand both source and target volumes, and then re-establish the replication relationship. If not explicitly set, this flag is used as False by default. Note as SRDF should be carried out, caution should be taken when performing this operation. If absolutely necessary, only one pair of resources and targets should be extended at a time. Note Currently it is not possible to expand SRDF/Metro protected volumes. mode is the required replication mode. Options are 'Synchronous', 'Asynchronous', and 'Metro'. This defaults to 'Synchronous' metro_use_bias parameter when using 'bias' protection is used instead of Witness. This false is the default. allow_delete_metro is a flag that indicates whether subway devices have been deleted. All Metro devices in an RDF group must be managed together, so the entire group must be suspended first to delete one of the pairings. Therefore, we want this flag to be clearly set. This flag is against False by default. Note Service Level and Workload: An attempt will be made to create a storage group with the same combination of service level and workload as the primary in the target series. However, if this combination is not available on the remote array, then the source device is Hybrid, the target array is an All Flash, and the Entire Flash incompatible service level, such as Bronze, is configured), no service level is applied. Note VMAX naps drives can support a single replication target per back end, we do not support Concurrent SRDF or Cascaded SRDF. Make sure that there is only one replication_device input per back-end stanza. Create a volume type with replication enabled. After the VMAX back-end input in cinder.conf is replication_device parameter, a corresponding volume type must be replication_enabled the property set is specified. See Above setup VMAX drivers for details. # openstack volume type set -- property replication_enabled=<is> True \ VMAX_FC_REPLICATION Most features are supported except for: Replication Group operations are available only for volumes in Synchronous mode. Storage-supported rewrite operations on replication-enabled VMAX volumes (from a non-replicated type to a replicated type, and vice versa. For example, switching to another service level/workload combination) is not supported. Currently, it is not possible to expand SRDF/Metro protected volumes. If a larger volume size is required for an SRDF/Metro protected volumes, this can be achieved by cloning the original volume and choosing a larger size for the new cloned volume. Image volume cache functionality is supported by setting image_volume_cache_enabled = True), but one of two actions must be taken when creating a cached volume: The first boot volume generated on a back end (which will trigger the cached volume to be created) must be the smallest required size. For example, if the least size disk to hold the image is 5GB, create the initial boot volume at 5 GB. Alternatively, allow_extend:<is><>replication_device parameter is set to True (please note that it is not possible to extend SRDF/Metro protected volumes). This is because the initial boot volume is created in the least size required for the desired image, and then expanded to the specified size of the user. In order to extend a Cinder.conf, for example, if the primary directory is upgraded, the administrator can use the failover-host method. #cinder failover-host host@VMAX_FC_REPLICATION If the primary array becomes reusable, it can start a rollback using the same command, and by specifying --backend_id default: #cinder failover-host cinder_host@VMAX_FC_REPLICATION --backend_id default Note failover and Failback operations are not in the current Metro configurations. Asynchronous and metro volumes in an RDF session, that is, belonging to an SRDF group, must be managed together for RDF operations (although it is an option exempt from consistency to create and release pairs in an Async group). To facilitate this management, we are creating an internal RDF management storage group at the back end. It is very important for the correct management that the volumes in this storage group are directly corresponding to the volumes in the RDF group. Therefore, the RDF group specified in cinder.conf is mandatory for exclusive use by this cürnbest back end. SRDF/Metro is a high availability solution. RdF works by masking both sides of the relationship to the host and presenting all routes to the host, appearing to point to a single device. To do this, there must be a lot of software running to manage writing to multiple paths. Audio rewriting with storage-supported migration is now supported for VMAX3 arrays. Cinder requires that a volume cannot be rewritten between the back ends for storage-supported migration. To use storage-supported volume rewriting, follow these steps: To pass a volume from one Service Level or Workload combination to another, use Rewrite the application to another using the migration policy. The target audio type must VMAX plug-in on the audio unit. In source masking view, the volume is moved from the FAST storage group with the parent storage group. The volume is added to the FAST storage group in the target parent storage group of the target masking view. At this point, the volume belongs to two storage groups. One of two things happens: If the connection to the target example succeeds, the volume is removed from the FAST storage group in the source masking view and deletes the original storage group that contains another volume. If the connection to the target instance fails, the volume is removed from the target storage group that is removing the storage group if it is empty. The volume was returned to the original storage group. For more information, please see the following: Note OpenStack osio uses an open standard for messaging group and groups of them that protect the host path, the volume can remain 'live'. During live migration, the following steps are performed by the VMAX plug-in on the audio unit: In source masking view, the volume is moved from the FAST storage group with the parent storage group. The volume is added to the FAST storage group in the target parent storage group of the target masking view. The volume is removed from the FAST storage group in the source masking view and deletes the original storage group that contains another volume. If the connection to the target example fails, the volume is removed from the target storage group that is removing the storage group if it is empty. The volume was returned to the original storage group. For more information, please see the following: Note OpenStack osio uses an open standard for messaging intermediate software as AMQP. This messaging intermediate software (RPC messaging system) provides OpenStack services that work on multiple servers to talk to each other. By default, the RPC messaging client is set to time out after 60 seconds, meaning that if any operation you perform takes longer than 60 seconds to complete, the process times out and the ERROR message MessageTimetime Timeout: Timed out waitingwaiting to message ID [message_id] In this case, cinder.conf and nova.conf increase the rpc_response_timeout flag value to all Cinders and Novas and restart services. To change this value, which will depend entirely on its environment, you may only need to increase it slightly, or you may need a little more time than usual if your environment is under heavy network load. Fine-tuning is required here, change the value, and reset services. At least rpc_response_timeout to 240, but this will need to be raised if high concurrency is a factor. This should also be enough for all cinder backup commands. NOVA-INST-DIR/instances/ (for example, /opt/stack/data/nova/instances) must be installed by shared storage. Make sure nova-inst-dir (set to state_path in nova.conf file) is the same on all computers. Configure your DNS or/etc/hosts and make sure it is consistent on all hosts. The host computers are able to perform name resolution with each other. As a test, use the ping command to ping each host from each other. $ping HostA$ping HostB$ping HostC Export NOVA-INST-DIR/samples from HostA and make sure it is readable by HostB and HostC Compute user. For more information, please refer to the relevant operating system documentation: for example enable the 'play/search' bit on your shared directory so that qemu can use images within directories on all computing hosts. $ chmod o+x NOVA-INST-DIR/instances Note If you are switching from process to controller, make sure that it runs two steps above on the controller node to export the sample directory. For our use case shown below, we have three hosts with the names HostA, HostB, and HostC. HostB and HostC are process nodes, while HostA is the process node. The following have also been used in live migration. 2 gb boot volume using Cirros image. The sample 2048 is created using the above 2gb volume with a flavor m1.small using RAM, 20GB Disk and 1 VCPU. Create a bootable volume. openstack volume creation --volume size 2 \ --image cirros-0.3.5-x86_64-disk \ --volume_lm_1 start an instance by using the volume created above. $ openstack server create --volume lm_1 \ --flavor m1.small \ --nic --nic --security group default \ --availability zone nova:HostB \ server_lm_1 HostB has an example created by running: $ openstack server lm_1 file | grep hypervisor_hostname\\instance_name | OS-EXT-SRV-ATTR:hypervisor_hostname | HostB | OS-EXT-SRV-ATTR:instance_name | instance-00000006 Confirm, Through virsh, instance_name steps are returned in step 3 (instance-00000006), created using an example in HostB: $ virsh list --migrate the migration instance from all Id Name State --- 1 instance-00000006 Running HostB to HostA to HostC likes: $ openstack server migration --live HostA \ server_lm_1 Command run in 3 steps above. The hypervisor must run the Command in Host A on Host A to verify that it was created, for example, through virsh. Managing volumes in OpenStack is the process by which a volume on a storage device is transferred to OpenStack for use in an OpenStack environment. In order for a volume to be said according to OpenStack, the following prerequisites must be met: Cinder is not part of the Masking View that has volume in the managed pool Volume is not part of an SRDF relationship Volume TDEV (thin device) Volume is set to FBA emulation Volume must be in an entire gb instance. Volume 5.5GB is not a valid target because it cannot target a snapvx for a volume to said in openstack, the same Storage Resource Pool (SRP) must reside as the back end configured for use in openstack. Specifying the pool correctly can be entered manually as in the same format: Pool format: <service_level>+ + Pool instance<workload_type><srp><array_id>; 1: Diamond+DSS+SRP_1+11111111Hav example 2: Diamond+SRP_1+11111111111Havuz values service_level The service level of the volume to be managed the workload of the volume to be managed Storage Resource Pool (12 digit numeric) configured for use by the back-end array_id Can now manage the volume if the back-end array_id matches the service_level that your pool name defined, this CLI command is possible with cinder management. If a service_level is not bootable, leave the boot parameter out of this command or optional. OpenStack also determines the size of the valve when managed without specifying the volume size. Command format: $ cinder manage --name volume_name>volume-type id> <array_id> <device_id>Command Example: $cinder manage --name image --volume-type volume VMAX_ISCSI_DIAMOND \ --availability-zone nova demo@VMAX_ISCSI_DIAMOND#Diamond+SRP_1+11111111111111111111111111111d8 After the above command is run, the volume will be available in the same way as the other command<id>/identifier> <host> <h/--bootable> <lt/av_zone> <lt/vmax_vol_type> <lt/array_id> <lt/device_id> <lt/service_level> <lt/service_level> VMAX volume. An unmanaged volume with an 'OS-' pren over the identifier name cannot be managed in OpenStack because it is a keyword reserved for managed volumes. If the identifier name has this prenation, an exception is thrown by the VMAX driver in a management process. Although it is not possible to manage volumes that are part of an SRDF relationship to OpenStack, this is done by having the VMAX volume type with replication enabled during the process of managing the process (see <a0> </a0>. Once managed, replication is enabled for this volume. Not managing a volume is not the same as when deleting a volume. When a volume is deleted from OpenStack, VMAX is also deleted. The volume can also be managed back to OpenStack, but remains for further use in VMAX. The unmanaged is performed using the Cinder unmanage CLI command: Command format: $cinder unmanage <volume_name olume id> Command example: $cinder unmanage vmax_test_vol OpenStack once unmanaged, the volume can still be retrieved using its device ID or OpenStack volume. In unisphere you will also notice that the 'OS-' pren if the prenir has been removed, this sound is now another visual indicator managed by OpenStack. If the resource volume already exists in Cinder, users can manage VMAX SnapVX snapshots to OpenStack. Similarly, users will be able to manage OpenStack snapshots to remove them from Cinder but keep storage backpoints. Setup, restrictions, and requirements: No additional settings or configurations are required to support this functionality. SnapVX functionality support is required in VMAX to manage/not manage snapshots. Managing/Managing Snapshots in OpenStack Cinder is currently only supported through Cinder CLI commands. In Cinder, it is possible to manage or manage only one snapshot at a time. It is possible to manage VMAX SnapVX snapshots to OpenStack, the volume of resources received from the snapshot already exists and is managed by OpenStack Cinder. The resource volume may have been created in OpenStack Cinder, or it may have been managed from OpenStack Cinder as well. Removed the restriction on managing SnapVX resource volumes with support for managing SnapVX snapshots in OpenStack Queens. Note It is not possible to manage target volumes linked to OpenStack SnapVX or existing volumes in a replication session. Requirements/Restrictions: SnapVX resource volume must be found in Cinder and managed by Cinder. SnapVX snapshot name should not start <lt/volume_name>Os-. SnapVX snapshot source volume should not be in a failed state. Why snapVX snapshot management is allowed only if the snapshot does not have linked target units. Command Structure: Define your SnapVX snapshot for management in VMAX, pay attention to the name. Note the device ID, make sure that the resource volume is managed into OpenStack Cinder. Using Cinder CLI, use the following command structure to manage snapshots to OpenStack Cinder: $ cinder snapshot-manage --id type resource name <lt/name> <lt/identifier> <lt/metadata><lt/key=value>; <lt/key=value>; ...] <lt/device_id> <lt/identifier>Positional arguments: <lt/device_id>; VMAX device id <lt/identifier>: Name of existing snapshot Optional arguments: --name <lt/name>: Snapshot name (Default=N/A) --description <lt/description>: Snapshot description (Default=N/A) --metadata [<lt/key=value>; <lt/key=value>; [...]] Metadata key=value pairs (Default=No) Example: $ cinder snapshot-manage --name SnapshotManaged1 --description Managed Queens Feb18 \ 0021A Where is VMAXSnapshot: After managing the SnapVX snapshot, the name snapshotmanaged of the snapshot description will be managed Queens February18. The source volume device ID is 0021A. At the end of VMAX is the instant name VMAXSnapshot. Conclusion: After the process of management of Snapshot is complete, the SnapVX snapshot at the back end of the VMAX is predetermined by the letters OS-VMAXSnapshot, leaving a snapshot in this example called OS-VMAXSnapshot. The associated snapshot managed by Cinder will be available for use as the same name SnapshotManaged. The unmanaged snapshot in Cinder is the process by which the snapshot is removed from Cinder and no longer managed by Cinder, but still at the storage back end. OpenStack Cinder follows this behavior of unmanaging a SnapVX snapshot, so after Cinder removes a VMAX SnapVX snapshot management, the snapshot is removed from OpenStack but is still available for VMAX back-end use. Requirements/Constraints: SnapVX resource volume should not fail on a case-by-case condition. Command Structure: Identify the SnapVX snapshot from OpenStack cinder that you don't want to manage, note the snapshot name or ID as specified by Cinder. Use the following command structure to manage SnapVX snapshot using Cinder CLI: $ cinder snapshot-unmanage <lt/snapshot>Positional arguments: <lt/snapshot>; Cinder snapshot name or ID. Example: After cinder snapvx management unstall, the snapshot on the VMAX back end is no longer managed in OpenStack icin işletim sistemi öneki kaldırılmış olacaktır. Yukarıdaki örnekte, OpenStack'ten yönetilmekten sonraki anlık görüntü, depolama ını arka ucunda VMAXSnapshot olarak adlandırılacaktır. Not Bu özellik sadece Unisphere REST'ten edinilebilir<lt/snapshot>. <lt/snapshot> <lt/key=value>; <lt/key=value> <lt/description> <lt/name>; <lt/name>; restore from a snapshot or return a volume to a snapshot, restore data from the snapshot to the volume, meaning the volume is overwritten with the time data in the snapshot. Unseation is not allowed when the volume and snapshot are not the same size, it only works with the latest snapshot assigned to the volume. Note This feature is only available from Unisphere REST 9.0.0.6, and the directory must be used in HyperMaxOS 5978 or later. Device expansion was designed to support RDF devices and devices with snapVX sessions. Session.