

Next Generation Computing

Enterprise Computing requirements are changing to address next generation technologies like AI/ML, BlockChain, Analytics, Multisite clouds, edge computing etc. Enterprises need computing where data resides and pushing the enormous data to public data centers will be an impossible task as Data generated at the On-Premises rapidly growing at very high phase. Computing infrastructure has to move to On-Premises or to Edge to address the huge computing requirements. Deploying cloud at On-Premises demands more IT forces and it forced the innovation of Hyper converged Infrastructure. It simplifies multiple aspects of IT such as setting up, monitoring, provisioning and backup functions of IT infrastructure. Today you must support an even wider set of application deployment models virtually anywhere the data resides:

- ❖ **Traditional Computing** Requirements requires datacenter type of setup to address its requirements
- ❖ **Big data Analytics** needs massive amount of computing and acceleration. There are different types acceleration setups are available currently in the market like GPU, FPGA, TPU etc. New IT infrastructure should support these requirements.
- ❖ **Cloud Native Architecture** needs Virtual Machines and Containers to deploy micro services.
- ❖ **Distributed Applications** needs computing and storage where data and user resides.

The Enterprises spend lots of money and resources on building turnkey infrastructure that unifies the servers, storage, and virtualization along with end-to-end systems management and operations management capabilities on legacy infrastructure. But the results of these implementation activities never gives the desired effects as it is the time to change the way we think while reorganizing the Data Centre resources. Also Enterprise IT today is looking for ways to implement on-premises cloud and services with the speed and operational efficiency of public cloud services such as Amazon Web Services (AWS), Microsoft Azure, and Google Cloud. Hyper Converged infrastructure address all above mentioned requirements of next generation and traditional computing as well as cloud requirements of Enterprises. It simplifies deployment, provisioning, monitoring of the solution.

What is Hyper Converged Infrastructure (HCI)?

- ❖ HCI follows distributed server, storage, network, and other Data Centre resources.
- ❖ HCI combines locally attached resources to a compute and distribute these resources using intelligent software.
- ❖ HCI replace legacy expensive infrastructure consisting of separate servers, storage networks, and storage arrays

What are the advantages of HCI?

- ❖ HCI improves the efficiency of silos
- ❖ HCI removes the complexities of legacy infrastructure
- ❖ HCI accelerate the scale of deployment with ease
- ❖ HCI remove painful management of legacy datacenter infrastructure
- ❖ HCI virtualization endpoint security of virtual machines (VM) and hosts
- ❖ HCI introduces zero data loss scheme in the Data Centre storage domain
- ❖ HCI improves service availability of Data Centre components

HCI allows enterprises to deploy infrastructure in minutes and shift the focus to applications that power the business. Acceleron delivers a comprehensive enterprise cloud platform that bridges the wide gap between traditional infrastructure and public cloud services.

1. Acceleron HCI Solution

Acceleron delivers a web-scale, hyper converged infrastructure solution custom made for virtualization and cloud deployments. This solution address the scale, resilience, and economic benefits of web-scale architecture to the enterprise through the Acceleron Enterprise Cloud Platform, which combines three components—Acceleron Lambda, Acceleron Graviton and Acceleron Nebula.

❖ **Scalability**

Acceleron HCI scales both horizontally and vertically to meet enterprise demands. All the HCI nodes available in the markets can increase their storage, Compute and networking capability by adding another HCI node. But Acceleron HCI node not only scales horizontally, but also it can scales vertically by increasing computing, storage and network resources inside a HCI node. Acceleron HCI can scale up to 64 HCI nodes and any nodes in HCI can share the storage resources from all nodes in the cluster. Acceleron HCI Intelligently distributes data across all nodes and uniform capacity utilization across all nodes at all time is maintained. HCI can independently scale storage and compute as and when needed without any downtime.

❖ **Different Hypervisors and workloads**

Acceleron HCI node is designed to address different type of work loads by enabling computing cloud components like Intel Servers, GPUs, and FPGAs etc. Acceleron HCI supports multiple industry standard hypervisors, such as KVM, RHEL, VMware, Hyper-V etc. Acceleron HCI has been architected as a **Hypervisor Agnostic Solution** to address various enterprise customers' hypervisor requirements.

❖ **High Availability**

Acceleron HCI software is designed to meet the high-availability standards and data protection that are expected from enterprise-class arrays. It is designed to be ready for next generation applications and data demands. Native data protection features found in HCI Software is essential to all infrastructures as data protection is very important. Acceleron HCI architecture allows you to deploy your infrastructure with simplified management and independently scaled compute and storage nodes, further it ensures that resources are deployed when and where they required. By deploying Acceleron HCI, IT organizations can transform their data center, driving operational efficiencies and reducing costs.

2. How to Deploy Acceleron HCI Solution

Acceleron HCI cluster is a completely software defined cluster. Each node in a cluster runs a hypervisor (VMware vSphere, Microsoft Hyper-V, KVM and RHEL), and the Acceleron Graviton and Nebula software. The Acceleron Lambda is a web based software for MDC410 hardware and Chassis management tool, it is only required when HCI is built on MDC410.

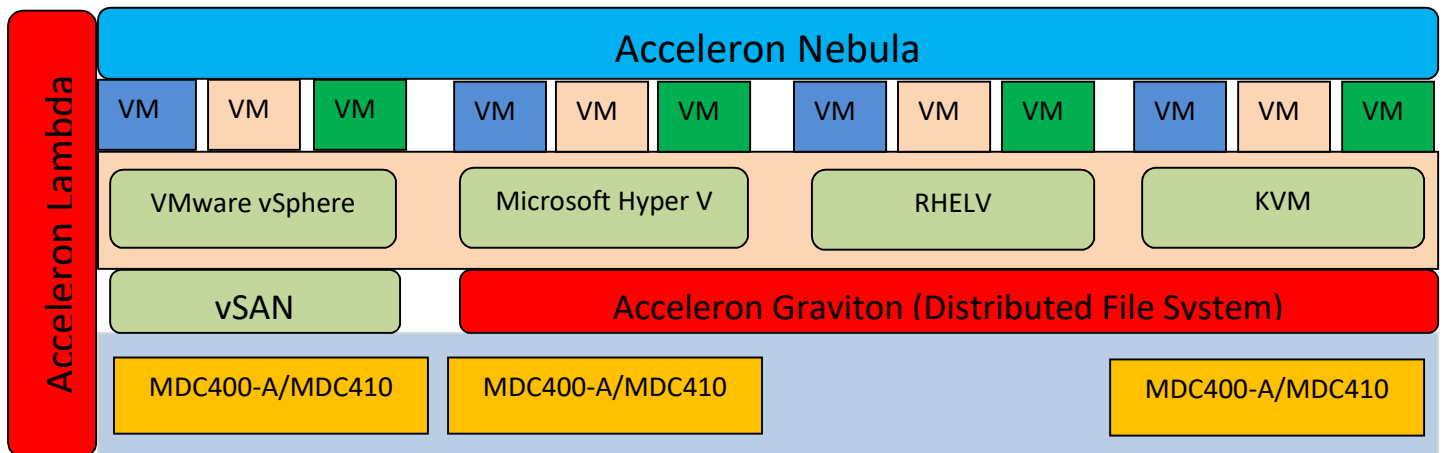
Acceleron HCI software can be deployed on following hardware

- ❖ MDC400A is a generic 2U HCI node for both compute and storage. It can also be configured for both HCI Compute only and HCI storage only solutions (Acceleron Graviton and Nebula is required).
- ❖ MDC410 hyper converged appliance, it can build HCI solutions in a box with more than 200 Intel Xeon cores and it can also supports additional 600 TB storage. MDC410 has built-in two low latency switch fabric for connecting HCI nodes inside appliance (Acceleron Lambda, Graviton and Nebula is required).
- ❖ On any industry standard server hardware from standard vendors (Acceleron Graviton and Nebula is required).

2.1. Graviton:

Graviton is a distributed storage fabric with software defined storage features and it runs on the each node of the cluster. Graviton consists of storage Cluster monitor, cluster manager, Object storage Demon and Metadata Server. For high availability and redundancy of storage nodes, these components will be running on all HCI nodes. Graviton store data as objects within logical storage pools and it enables the Storage Cluster to scale, rebalance, and recover dynamically. It can be configured for replication and erasure coding or both. It enables Storage High Availability features and multiplatform storage sharing capability. Graviton enables inter-operability of storage sharing with third party HCI platforms. Graviton’s iSCSI gateway functionalities are enabled for volumes block services which provides iSCSI access to applications that require direct access to block storage. This can be non-virtualized systems, or virtual machines with specific requirements.

- ❖ Graviton enables the Enterprise storage services for applications, eliminating the need for separate solutions from vendors such as NetApp, EMC, and HP
- ❖ Graviton enables a comprehensive set of capabilities for storage performance acceleration, data reduction, data protection, and much more
- ❖ Graviton support for VMware® vSphere, Microsoft® Hyper-V, KVM, and RHEL V.



The Graviton, Distributed Storage Fabric simplifies storage and data management for virtual environments. Graviton supports different types of storages such as SATA, SAS, SSD and NVMe. It supports different combinations of hard drives and SSDs storage across a HCI cluster and exporting it as a data store to the virtualization layer as iSCSI, NFS, and SMB shares, Graviton eliminates the need for SAN and NAS solutions. The Graviton keeps multiple replicas of VM storage intelligently on HCI nodes. It also supports erasure coding which gives data protections similar to RAID 5/6 configurations. Graviton implemented

2.2. Nebula:

Acceleron Management tool is integrated with Acceleron Nebula, it is formally called **Acceleron HCI**. Acceleron Nebula is a distributed instance of cloud suite and it is the heart of HCI software and it enables all VM-centric policies and operation of HCI. Complete virtualization solution included with Acceleron Nebula at no additional hidden charge. It enables Enterprise virtualization for compute, storage and network elements in the HCI cluster. VM creation, Storage Volume creation, network creation etc are major functionalities of Nebula software and it also orchestrates storage provisioning related features like snapshot, backup and cloning etc.

Nebula enables some Enterprise critical features such as Node High availability and VM live migration. Nebula also has HCI management features which simplify and streamline common workflows and eliminating the need for separate management solutions for servers, storage networks, storage, and virtualization. This suite of features enables application and business requirements to be met within the hyper-converged infrastructure, with no need to rely on other third-party tools. In other words, you can build a complete datacenter on Acceleron HCI.

Virtual networks are created by tenants and administrators to provide networking capability between VMs managed by HCI compute. Neutron, which is a part of Nebula, is a network management service that exposes an extensible set of APIs for creating and managing virtual networks.

Nebula supports SNMP-based monitoring using third-party SNMP tools, certain important metrics and details are included as SNMP monitoring. User can set SMTP-based email alerting for certain parameters from Nebula dashboard.

Nebula supports multiple clusters as different projects. Admin can add multiple HCI clusters under same management utility. This feature is given as part of Acceleron HCI dashboard.

HCI Management tool has server and client architecture and every computing node and storage nodes has client portion of HCI management tool. Management tool grows as the number of cluster nodes increases and it is built into each and every node.

2.3. Lambda:

It is a web based Hardware Management tool, which configures the MDC410 HCI cluster and monitors the alerts, temperature, vibration, various logs etc. Various applications such as HCI, VDI etc can be launched from this web tool. Lambda has built-in dashboard to view desktop of all servers and nodes in HCI cluster. Firmware upgrades of all components of HCI are managed through Lambda. Lambda is the “block” management tool for Acceleron MDC410. MDC410 is a cluster server in a box which has minimum 6 x Dual Socket servers / 12 x Single Socket servers in a box with two numbers of low latency switch fabric.

2.4. Infrastructure Resilience

The Acceleron HCI is designed to be fault tolerant, with no single points of failure or performance bottlenecks. This is achieved through multiple features such as storage redundancy, high availability of HCI nodes, error correction codes built-in distributed storage fabric and data path redundancy.

2.4.1. Storage Redundancy

Acceleron HCI has mechanism for Metadata protection for all offered nodes within the cluster so as to provide high availability and no single point of failure. MetaData Server (MDS) which provides a more traditional style of filesystem based on POSIX standards that translate into objects stored in the OSD pool. This is typically where a non-Linux platform can implement client support for Graviton. This can be shared via CIFS and NFS to non-Graviton and non-Linux based systems including Windows. Graviton requires all of its servers to be able to see each other directly in the cluster. So this filesystem would also be the point where external systems would be able to see the content without having direct access to the Graviton Cluster. **Meta Data Server implemented with Reed-Solomon code for Data Integrity Checks.**

Graviton storage container is configured with a replication factor (k) of two or three or four or more based on the initial configuration of replication factor. Acceleron achieve replication in m storage disks k copies. RF=2 ensures that two copies of data are maintained at all times, allowing the cluster to survive the failure of a single node or drive. When RF is set to 3 (RF=3), three copies of the data are maintained in a cluster, providing resilience against two simultaneous failures. RF=2 is considered the best practice in most scenarios.

Graviton has Cache Cluster support which speeds up low performance storages (NLSAS/SATA) access. Normally cache is populated with SAS SSD disk or VMe SSD disks.

2.4.2. Storage Optimization

Graviton support different type of storage optimization methods to make efficient use of storage available in the cluster.

2.4.2.1. Deduplication

Data deduplication process that eliminates excessive copies of data and significantly decreases storage capacity requirements. Deduplication can be run as an inline process as the data is being written into the storage system and it can also run as a background process to eliminate duplicates after the data is written to disk

When deduplication is enabled, data is fingerprinted on ingest using a SHA-1 hash. Deduplication operations are software driven and leverage the hardware-assist capabilities of the Intel chipset for the SHA-1 fingerprint generation. Because SHA-1 is a strong hash, deduplication is performed based on a fingerprint match

2.4.2.2. Compression

Data compression can be done in-line and as post process after writing into the storage disk. Post process compress understands the random data access patterns and improves the compress performance.

2.4.2.3. Erasure Coding

Acceleron HCI implements the innovative data protection scheme similar to RAID across multiple storage nodes. It reduces the capacity cost compared to replication factor based data protection scheme. This EC scheme is enabled for a cluster. EC mode which is similar to RAID 5 can give available storage as 75% of the total storage in the cluster. There are different modes are possible, but it will enable as part of initial configuration. Dynamic assignment of EC coding is not allowed. User has to decide on Replication factor and EC scheme during the deployment of the cluster.

EC scheme encodes a strip of data blocks on different nodes and calculates parity. In the event of a disk or node failure, parity is used to calculate any missing data blocks. Graviton uses an extent group as the data block, and each data block in a strip must be on a different node and belong to a different vDisk. The number of data and parity blocks in a strip is configured based on the desired number of failures to withstand.

2.4.3. Data Protection and Backup

Acceleron HCI gives data protection and backup features for Enterprise storage cloud. Acceleron HCI defines set of RTO (Recovery Time Objective) and RPO (Recovery Point Objective) guidelines to help the customers to customize their backup requirements of different applications.

2.4.3.1. Local data Snapshots

Graviton can create unlimited local snapshots with VM and application level consistency and recover data instantly to meet a wide range of backup and data protection requirements. Metadata based snapshots require minimal disk overhead and ensure high performance recovery. Graviton uses VM-centric snapshots to provide production-level data protection without sacrificing performance.

2.4.3.2. Remote Backup and Disaster Recovery

Organizations prepare for everything from natural disasters to cyber-attacks with disaster recovery plans that detail a process to resume mission-critical functions quickly and without major losses in revenues or business operations.

Acceleron HCI Disaster Recovery and Remote Back are based on incremental backup which is executed at defined intervals. Those incremental backup files are stored on to a remote data centre or a backup facility. Acceleron replication logic supports one-one, one- many and many - many replication formats based on customer requirements. Additionally, Acceleron has integrated other third party tools like **Veeam** etc also as part of its solution.

Acceleron Nebula interface displays a simplified view of all local and remote snapshots, which allow the administrators to restore a VM from a snapshot with a single click. In case of disaster, administrators can launch a failed VM to the secondary datacenter.

2.4.3.3. DR and DC

Nebula supports multiple clusters, one can be DC cluster and other one can be the DR Cluster. Once DC fail completely or has to shut down completely, then DR can bring all VMs and storage alive with Acceleron HCI support.

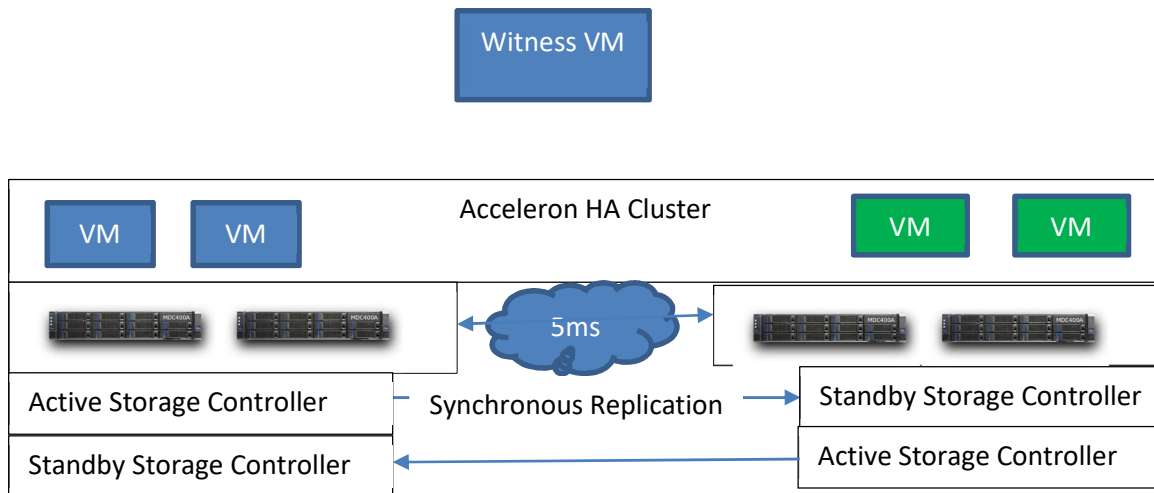
For critical workloads requiring zero RPO, and near-zero RTO, Acceleron offers DR and DC solution which has High availability for server, storage and network, which ensure continuous data availability across separate sites within an edge network. Data is written synchronously to both sites, so it is always available to applications in the event a site fails or needs to undergo maintenance. Virtual machines can be non-disruptively migrated between sites for planned maintenance events or other needs.

2.4.3.4. Acceleron Metro Availability Overview

Acceleron Metro Availability is part of a suite of data protection features offered on the Acceleron HCI platform for DR-DC. Metro Availability is a continuous availability solution, provides a global file system namespace across “stretched” between Acceleron clusters. Synchronous storage replication mechanism supports the independent Acceleron clusters, using the “protection domain” and “remote site” constructs. The platform enables synchronous replication at Cluster level and replicates all VMs and files stored in that storage controller synchronously to another Acceleron cluster. Each protection domain maps to one container.

Administrators can create multiple protection domains to enable different policies, including bidirectional replication, where each Acceleron cluster replicates synchronously to one or more clusters. Containers have two primary roles while enabled for Metro Availability: active and standby. As shown in the following figure, active storage controller replicate data synchronously to standby storage controller. Administrators create a protection domain by specifying an active storage controller on the local cluster and a standby storage controller with the same name on the remote cluster. The active and standby storage controller mount to their respective hypervisor hosts using the same datastore name, which effectively spans the datastore across both clusters and sites.

With a datastore stretched across both Acceleron clusters, you can create a single hypervisor cluster and use common clustering features, like Live Migration and High Availability, to manage the environment. Hosts presenting standby storage controller can run VMs targeted for the standby storage controller; however, standby storage controller are not available for direct VM traffic. The system forwards all I/O targeted for a standby storage controller to the active site.



The Acceleron HCI platform supports Metro Availability in conjunction with other data management features, including compression, deduplication, and tiering. Metro Availability also allows you to enable compression for the synchronous replication traffic between the Acceleron clusters. When you configure the remote site, enabling replication traffic compression reduces the total bandwidth required to maintain the synchronous relationship

❖ Active Storage Controller

All local write operations, including both random and sequential writes, go through the operation log of the HCI Controller with Metro Availability enabled. When local writes occur, data replicates locally based on the container's replication factor. In parallel, the system sends remote writes to the operation log of a HCI Controller in the cluster maintaining the standby storage controller. The standby storage controller's replication factor determines the number of writes required to complete the remote write operation against the peer cluster. After all replication factor writes are processed for both the active and standby storage controller, the VM receives acknowledgement that the write is complete.

❖ Standby Storage Controller

Standby storage controller receives synchronous write updates from the active storage controller. Standby storage controller mount to the appropriate hypervisor instances in the Acceleron cluster, so you can run VMs locally to the hypervisor nodes that own the storage controller and datastore. While the standby storage controller appears available, it does not process VM I/O directly. The system forwards all VM read or write I/O targeted for a standby storage controller to a HCI controller in the cluster that owns the active storage controller before returning data or acknowledgement to the VM.

❖ **Metro Availability Witness**

The Metro Availability witness allows Acceleron HCI clusters to automate management of storage and replication states, including the promotion of standby protection domains, so VMs can automatically fail over during site failures. The witness is a VM you can install in any server or PC. It resides in a site and failure domain separate from the Acceleron HCI clusters forming the Metro relationship, enabling failover decisions that avoid conditions that could lead to a split-brain scenario, where storage is active on both sites. Acceleron HCI clusters register with the witness, and each protection domain represents a unique relationship. This design permits multiple sets of clusters to register with one witness and supports both active and standby protection domains.

❖ **Metro Availability Replication States**

When you first enable Metro Availability, a full copy of the data residing in the active storage controller replicates between the clusters. Acceleron HCI uses a snapshot for this initial full copy and maintains it as a reference point. The relationship is in a synchronizing state until the copy completes. Any writes that occur following the initial snapshot replicate synchronously. Once fully synchronized, the protection domain enters an “enabled” state, which indicates that the standby storage controller is consistent.

2.4.4. HCI Network

Acceleron HCI network is a distributed switch (DRS) and each node in the HCI network clients often deploys several processes across a number of nodes. These processes interact with each other and other HCI services. The main process of the HCI Networking service is neutron-server, a Python daemon that exposes the HCI Networking API and passes tenant requests to a suite of plug-ins for additional processing. Acceleron HCI management tool provides visibility of network infrastructure. HCI network support L2&L3 routing features.

The HCI Networking components are:

❖ **Server Component**

This service runs on the network node to service the Networking API and its extensions. It also enforces the network model and IP addressing of each port. The server requires indirect access to a persistent database. This is accomplished through plugins, which communicate with the database using AMQP (Advanced Message Queuing Protocol).

❖ **plugin agent**

Runs on each compute node to manage local virtual switch (vswitch) configuration. The plug-in that you use determine which agents run. This service requires message queue access and depends on the plugin used.

❖ **DHCP agent**

It provides DHCP services to tenant networks. This agent is the same across all plug-ins and is responsible for maintaining DHCP configuration. Dhcp-agent requires message queue access.

❖ **L3 agent**

It provides L3/NAT forwarding for external network access of VMs on tenant networks. It requires message queue access.

❖ **Network provider services (SDN server/services).**

It provides additional networking services to tenant networks. These SDN services may interact with server, and plugin-agents through communication channels such as REST APIs.

2.4.5. WAN Optimizations for storage cluster

Acceleron HCI uses WAN optimization techniques like Deduplication, Compression, local caching, QoS etc. Deduplication eliminates the transfer of redundant data across the storage cluster network by sending references instead of the actual data. Similarly compression relies on data patterns that can be represented more efficiently. Essentially compression techniques similar to ZIP, RAR, ARJ etc. are applied on-the-fly to data passing through hardware (or virtual machine). Acceleron Graviton supports these two features and made available to storage clusters. Local caching is enable using graviton on cache pool which enhances the Read and Write latency behavior of the cluster.

2.4.6. Security

Acceleron Nebula combines features such as two-factor authentication and data-at-rest encryption. Authentication is an integral part of any real world HCI deployment and so careful thought should be given to this aspect of system design.

Acceleron HCI Identity service supports multiple methods of authentication, including user name & password, LDAP, and external authentication methods. Upon successful authentication, The Identity service provides the user with an authorization token used for subsequent service requests.

LDAP server management is configuration file based feature. Admin User can update details of LDAP server like IP address user name etc by editing this File.

2.4.6.1. Two-factor Authentication

Identity services employs multi-factor authentication for network access to privileged user accounts. The Identity service supports external authentication services through the Apache web server that can provide this functionality. Servers may also enforce client-side authentication using certificates.

2.4.6.2. Encryption at rest

Acceleron HCI support encryption of data in three key areas - local storage for Virtual Machine instances, block devices and Object Storage block devices. The objective of this feature is to mitigate the risk of data compromise in the event that disks or full servers are removed from data center deployments.

Encryption of underlying block devices is performed using dm-crypt with LUKS; key management is provided by Vault, which provides secure encrypted storage of the keys used for each block device with automatic sealing of secrets in the event of reboot/restart of services.

2.4.7. Hypervisor

Acceleron HCI uses **Hypervisor Agnostic solution** which supports various industry standard virtualization frameworks. It gives flexibility in choosing customer workload specific virtualization solutions based on their requirements. Acceleron HCI can support **Acceleron Native Hypervisor, VMware vSphere, Microsoft Hyper-V, Linux KVM** and **RHELV** on its devices.

2.4.7.1. Acceleron Native Hypervisor (ANH)-KVM based

ANH is an enhanced Ubuntu Linux KVM implementation which has integrated with various Enterprise applications. ANH being the Linux flavor it has the inherent capability to run all Linux VMs, Windows VMs and FreeBSD VMs. Acceleron Labs' applications like Acceleron VDI, Acceleron Connect (Video Conferencing solution) etc are some of the applications which can work seamlessly on ANH. ANH is customized for Acceleron HCI, taking advantage of the intelligent storage services provided by Nebula, Graviton and Lambda. Nebula is optimized for use with server virtualization; it provides full data resiliency and data services such as snapshots, clones, and provisioning operations at VM granularity. With power of ANH, Acceleron HCI delivers compute services, storage service, SDN and high availability.

There are three main components in ANH are following

❖ **KVM-kmod**

KVM kernel module.

❖ **Libvirtd**

An API, daemon, and management tool for managing KVM and QEMU. Communication between ANH and KVM and QEMU occurs through libvirtd.

❖ **Qemu-KVM**

A machine emulator and virtualizer that runs in user space for every VM (domain). Hardware-assisted virtualization and VMs run on ANH uses Qemu-KVM.

2.4.7.2. ANH VM Network Security

ANH VM security includes east-west firewalls, or micro segmentation, allowing administrator to easily manage network isolation and VM level network policy. Each VM running on ANH can be provisioned with Firewall rules to allow access of it from outside world. Connectivity of VM to outside network can be achieved through **Floating IP** address assignment to native VMs.

2.4.7.3. ANH VM Data Protection

Each VM running on ANH is protected using local snapshots as well as replication to a remote site. ANH has full access to all the data protection capabilities of Acceleron HCI described in this document. Snapshot and back procedure need extra amount of storage in the cluster as well as in the remote location. Both

Snapshot and backup are executed on customer configuration. Periodic backup configurations such as RPO and RTO needs to be appropriately configured based on your backup storage capability.

2.4.8. Virtual Machine Management

VM management on Acceleron HCI focuses on creation, updates, deletion, data protection, and monitoring of VMs and their resources. These Acceleron HCI services and features are all available via the Nebula interface. Acceleron HCI's management dashboard gives all functionalities to manage the HCI clusters. Acceleron HCI management console can be access through web GUI which can be access remotely or locally using an html browser. It provides detailed matrix about Instants, vCPU, RAM, Volumes, Volume snapshots, Volume storage, Network details, security groups etc. It help the user to monitor the health of all HCI subsystems.

2.4.8.1. Thin provisioning

Thin provisioning gives the appearance of having more physical resources than are actually available. If a system always has enough resource to simultaneously support all of the virtualized resources, then it is not thin provisioned. It is not advisable thin provision real memory in a server. Each task acts as if it has real real resources allocated. The sum of the allocated virtual resources assigned to tasks typically exceeds the total of real real resources. Acceleron HCI thin provisioning is a configuration file based feature and Admin user can edit this file while setting up the HCI cluster.

2.4.8.2. VM Provisioning

Nebula displays a list of all VMs in an HCI cluster along with configuration, resource usage, and performance details for each VM. Administrators can create VMs and execute operations on selected VMs, including power on or off, power cycle, reset, shutdown, reboot, snapshots and clones, migration, pause, update, delete, Attach Floating IP, launch instant and launch a remote console.

2.4.8.3. VM Image Management

Virtual machine images contain a virtual disk that holds a bootable operating system on it. Disk images provide templates for virtual machine file systems. The Image service controls image storage and management.

Instances are the individual virtual machines that run on physical compute nodes inside the cloud. Users can launch any number of instances from the same image. Each launched instance runs from a copy of the base image. Any changes made to the instance do not affect the base image. Snapshots capture the state of an instances running disk. Users can create a snapshot, and build a new image based on these snapshots. The Compute service controls instance, image, and snapshot storage and management.

When you launch an instance, you must choose a flavor, which represents a set of virtual resources. Flavors define virtual CPU number, RAM amount available, and ephemeral disks size. Users must select from the set of available flavors defined on their cloud. Acceleron HCI provides a number of predefined

flavors that you can edit or add to the existing ones. You can add and remove additional resources from running instances, such as persistent volume storage, or public IP addresses.

The image management service can store the virtual disk files that are used to create fully functioning VMs or operating system installation media as an .iso file that you can mount to provide a fresh operating system install experience. Incorporated into Nebula, the image service can import and convert existing virtual disk formats, including .raw, .vhd, .vmdk, .vdi and .qcow2. Additionally for supporting hybrid cloud with Amazon image service has .aki, .ari, .ami format support.

2.4.8.4. VM Allocation and Resource sharing

Acceleron HCI can intelligently place virtual machines on nodes in a cluster based on the computes and storage resources availability. VM scheduler will take care of various aspects of the VM and resources before allocating VM in the host. Following filter are considered while deciding the host node for VM allocation

- ❖ High Availability Zone
- ❖ Compute capability
- ❖ Image properties
- ❖ Server Group Affinity
- ❖ Server Group Anti-Affinity

Affinity controls will have the ability to govern where VMs should run. VM-host affinity ties a VM to a host or group of hosts, so that the VM only runs on that host or group. Antiaffinity allows you designate VMs that should not run on the same hosts.

2.4.8.5. Cross-Hypervisor Migration

Acceleron HCI has built-in capability to take snapshot and back up to another cluster of ESXi or AHV etc. This feature is used to cross-hypervisor migration of VMs from other clusters to Acceleron HCI. You can create one or more protection zone on the source cluster and set the Acceleron HCI cluster as the target remote cluster. Then, snapshot VMs on the source ESXi/AHV cluster and replicate them to the Acceleron HCI cluster, where you can restore them and bring them online as Acceleron HCI VMs.

2.4.8.6. Automated High Availability

Virtual machine High availability (HA) is an enterprise feature, which allows VM seamlessly migrate to another host or block upon service outage of host or block. If a host fails, the VMs previously running on that host restart on healthy nodes throughout the cluster. There are multiple HA configuration options available to account for different cluster scenarios:

- ❖ By default, all Acceleron HCI clusters provide auto HA (best effort), even when you haven't configured the cluster for HA. Auto HA works without reserving any resources and does not enforce admission control, so the capacity may not be sufficient to start all the VMs from the failed host.

- ❖ You can also configure an Acceleron HCI cluster for HA with resource reservation to guarantee that the resources required to restart VMs are always available. Acceleron HCI offers two modes of resource reservation: host reservations and segment reservations. Clusters with uniform host configurations (for example, RAM on each node) use host reservation, while clusters with heterogeneous configurations use segment reservation.

- ❖ **Host reservations**

This method reserves an entire host for failover protection. Nebula selects the least-used host in the cluster as a reserve node and migrates all VMs on that node to other nodes in the cluster so that the reserve node's full capacity is available for VM failover. Nebula determines the number of failover hosts needed to match the number of failures the cluster tolerates for the configured replication factor.

- ❖ **Segment reservations**

Segment reservation divides the cluster into fixed-size segments of CPU and memory. Each segment corresponds to the largest VM to be restarted after a host failure. The scheduler, also taking into account the number of host failures the system can tolerate, implements admission control to ensure that there are enough resources reserved to restart VMs if any host in the cluster fails.

The Acceleron HCI controller restarts the VMs on the healthy hosts. The Acceleron HCI controller tracks host health by monitoring connections to the all cluster hosts. If the Acceleron HCI controller becomes partitioned or isolated, or if it fails, the healthy portion of the cluster elects a new Acceleron HCI controller.

2.4.9. GPU Support

A graphics processing unit (GPU) is the hardware or software that displays graphical content to Monitors. In non-virtualized Servers, a GPU is either a physical card or built directly into the CPU hardware, while GPU functions in the virtualized world have been software driven and consumed additional CPU cycles. With Modern PCIe interconnect, GPU can be virtualized with PCIe virtual functions on certain GPUs (AMD GPUs) and NVIDIA follows a different architecture for virtualized GPU. With advent of next generation applications as well as 3D tools, more and more enterprises using hardware GPU in the virtual world. Servers with installed physical GPU cards and expose virtual GPU to guest VMs using pass through or vGPU mode.

2.4.9.1. GPU Pass through

The GPU cards installed in server nodes for virtualization in GPU pass through mode, Acceleron HCI can pass a GPU through to a VM, allowing the VM to own that whole physical GPU device without any GPU virtualization. This is similar to configuring a node with one or more GPUs directly attaching them to a physical host/node. The Pass through mode helps the guest VMs to acquire one or more GPUs to them. Acceleron HCI currently supports NVIDIA Grid cards for GPU pass through; refer to our product documentation for the current list of supported devices.

With pass through, you can also use GPUs for offloading computational workloads—a more specialized situation than the typical graphical use cases. GPU compute scenarios assign one or more GPUs for a VM to use for processing.

2.4.9.2. vGPU

The passthrough method works well for a smaller number of VMs requiring larger amounts of GPU resources, workloads such as VDI often have different requirements. VDI workloads typically have a much larger number of VMs that need varying amounts of GPU resources based on application types and usage. NVIDIA Grid GPU cards contain 1–4 GPU on each physical PCI card, and each

physical host can support installing one or two cards. This capability allows for up to 8 GPU in each node to meet density requirements driven by VDI workloads. For maximum flexibility, vGPU mode allows you to slice each GPU into smaller segments that you can virtually assign to VMs.

vGPU profiles allow you to assign different levels of resources to VMs, so each of them can utilize a defined maximum number of displays and quality of resolution and Acceleron HCI support NVIDIA Grid cards for vGPU.

2.4.10. Acceleron HCI automation framework and REST API

Nebula support REST-API framework for automating HCI operations such as admin Logic API, Application catalog APIs, Container Service APIs, Backup Service APIs, Block storage API etc.

Admin Logic API supports functions such as Authentication, HTTP Status Codes, Service Discovery, Administrative Endpoints, DelegateAPI Endpoints etc

Application Logic API supports functions such as Getting information from the VM, like a config that is generated during the deployment, VM rebooting and Scaling

Backup Service API supports functions such as Backups allow users to record their backups and metadata information about those backups and when backups were taken. It holds the backup information. Jobs allow users to schedule and execute backup jobs on one node or more. Jobs are stored through the api in the database. Every job contains set of actions that carry out the backup job.

Block Storage APIs supports functions to create an environment with multiple-storage back ends and the API spawns Block Storage volume back ends and keys them from a unique queue

3. Acceleron Hyper Converged Infrastructure Software

HCI Software Specification brief		
1	Guest Operating Systems	Below operating systems are supported Windows10/11 Windows Server 12/16/19/22 Red Hat SUSE Ubuntu CentOS Solaris x86
2	Storage	Supports protocols NFS, ISCSI, SMB and FC. Has support for external storage products (NAS, ISCSI & FC) from leading vendors. Supports abstraction of external storage (SAN and NAS) devices making them HCI software aware. Supports thin/thick provisioning
3	VM Migration	Live Migration of Virtual Machines Migration of Virtual Machines between different generations of CPUs Migration of Virtual Machines without shared storage. Live Migration of Virtual Machines across storage arrays (NFS, ISCSI).
4	Networking	Support for NIC teaming for load sharing and redundancy Supports creation of virtual switches that connect virtual Machines
5	Backup and Recovery	Support for taking point-in-time snapshots of the virtual machines. Support for backup and recovery
6	Power Management	Support for Industry Standard Advanced Configuration and Power Interface (ACPI) specifications to save power in manual mode.
7	High Availability	Supports migration of VM from failed host to redundant host when problem occurs. Supports one or more nodes failures for both storage and compute nodes based on configuration.

8	Monitoring	Supports VM level monitoring of Storage CPU RAM Network
9	Over Commit	Supports CPU over commit Supports RAM over commit
10	Replication	Supports in-built Replication with RF=2,3,4, etc capability Supports VM data replication DC-DR bi-directional replication with any-any Mesh network data centers
11	Resource Allocation	Supports for adding VCPU to Virtual Machine Supports for adding RAM to Virtual Machine Supports for adding Disk space to Virtual Machine
12	Encryption	Supports for VM-level encryption which protects unauthorized data access both at-rest and in-motion
13	Hypervisor	Hypervisor Agnostic Solution VMware vSphere Microsoft Hyper V
14	Provisioning	Support for VM installation, configuration and upgrade on multiple physical servers. Support for VM flavor creation
15	Hybrid Cloud Support	Support drivers for managing VM and Storage in Amazon Cloud
16	End point Protection	Integrated Bit Defender "Gravity" Software as part of Acceleron HCI. User should get separate license to use the end point protection from Bit Defender.

MDC400A Hardware Specifications	
Form factor	<ul style="list-style-type: none"> • 1U/2U/4U, 19" rack mount • Dimensions (WxDxH) – 438 x 658 x 87 /174(mm)
Processor subsystem	<ul style="list-style-type: none"> • Intel Processors <ul style="list-style-type: none"> ○ 2 x Intel® Xeon® 2nd and 3rd Generation processors upto 265W TDP ○ 16 DIMM slots (For 2nd Generation Cascade Lake processor) ○ 32 DIMM Slots (For 3rd Generation Ice lake processor) ○ Support DDR4 2133/2400/2666/2933/3200MHz RDIMM/LR DIMM, for up to 64GB per DIMM. ○ Intel® C621/C621A chipset • AMD Processors <ul style="list-style-type: none"> ○ Upto AMD Processor 64Cores / 128T, 256M Cache, 225W ○ 32 DIMM Slots ○ Support DDR4 2133/2400/2666/2933/3200MHz RDIMM/LR DIMM, for up to 64GB per DIMM.
GPU	<ul style="list-style-type: none"> • 4 x NVIDIA GPU supported (MDC400A-G Model)
Storage Type	<ul style="list-style-type: none"> • SATA 3.0 • SAS HDD • SATA SSD • SAS SSD • NVMe SSD
Security	<ul style="list-style-type: none"> • Trusted Platform Module 2.0 v2and v3 • Silicon-based Hardware Root of Trust • Cryptographically signed image upload • Signed firmware updates • Secure alerting • Automatic BIOS recovery • Rapid OS recovery • System Secure Erase • System Lockdown
Certifications /Compliances (OS)	<ul style="list-style-type: none"> • Windows Server • Red Hat Linux • Suse Linux • Ubuntu • Cent OS
Certifications /Compliances	<ul style="list-style-type: none"> • VMWare • HyperV • Red Hat virtualization

Hypervisors or Virtualization	<ul style="list-style-type: none"> • Xen • KVM
Server Management	<ul style="list-style-type: none"> • BMC web based GUI interface to manage and monitor the sever resources. • HTML based management framework allows to manage servers in physical, local and remote environments while operating in-band • BMC network ports to manage the server out-of-band without a systems management software agent. • BMC software plugins to monitor and analysis hardware performance parameters, alerting, telemetry streaming features etc. • HTML based monitoring, driver updates & configuration, power monitoring & capping, RAID management, Out-of-band hardware & firmware inventory • Zero-touch auto configuration to auto deploy a baseline server configuration profile • Power Management parameter database storage for 5days. • Dedicated remote management port and support IPv4 and IPv6. • System health and systems components (CPU, RAM, HD, FANS, Power Supplies, HBA's, NICs) health monitoring support. • System hardware configuration and license restore and archiver support. • PXE based automated hardware configuration and Operating System deployment to multiple servers • Virtual IO management / stateless computing • BMC software support for integration with 3rd party management tools such as vCentre, System Center, BMC, Ansible and support for Redfish API
Hardware Certification available	<ul style="list-style-type: none"> • CE,FCC,RHoS
USB	<ul style="list-style-type: none"> • 4x USB 4.0 on the front
Front Control	<ul style="list-style-type: none"> • Power button • System reset button
Front I/O Ports	<ul style="list-style-type: none"> • 4 x USB 3.0
Visual Indicators	<ul style="list-style-type: none"> • Power • UID • LAN activity • HDD status
Rear Panel	<ul style="list-style-type: none"> • 4 x RJ45 Network interfaces (10/100/1000 Base-T)

	<ul style="list-style-type: none"> • 1 x 100Gbps Eth SFP28 / HDR 100 (optional) • 4 x 10Gbps SFP+/Base-T • 1 x IPMI interface (10/100/1000 Base-T) • 1 x VGA • 1 x DB-9 (serial port)
Expansion slots	<ul style="list-style-type: none"> • 4 x PCIe 4.0 x16, Low Profile (It will be used for system with Ethernet Interface, Infiniband Adapter cards, SAS Interface support) • 2 x PCIe 4.0 x8, Low Profile
Cache Drive	<ul style="list-style-type: none"> • Upto 4 x 1.6TB SAS SSD with 3DWPD/10DWPD
Drive bays	<ul style="list-style-type: none"> • 8 x 3.5" SAS/SATA/SSD or 8 x 2.5" SATA/SAS/SSD disks or 12 x 25"/3.5" SATA/SAS/SSD disks or 24 x 25"/3.5" ATA/SAS/SSD disks or 36 x 25"/3.5" ATA/SAS/SSD disks (configured at purchase) • 2x NVMe SSD U.2 Disks • Boot DISK 2 x M.2 or 2 x 960GB SAS SSD
Power	<ul style="list-style-type: none"> • 2 (1+1) CRPS (80+ Platinum) or 2 (1+1) (80+ Titanium) Redundant hot swappable PS– Configured at purchase • 4U support upto 2000W • 2U support upto 2000W • 1U support upto 1400W
Cooling	<ul style="list-style-type: none"> • 80 x 25/38mm internal redundant fans
Temperature	<ul style="list-style-type: none"> • Operating: 10°C to 35°C (50°F to 95°F) • Non-operating: -40°C to 70°C (-40°F to 158°F)
Weight	<ul style="list-style-type: none"> • 15Kg
Operating System Supported	<ul style="list-style-type: none"> • Microsoft Windows – Server 2012 and above versions • CentOS 6.8/7.4 • SUSE Enterprise Linux Server • Ubuntu 16.04/18.04 • Redhat Enterprise Linux • VMware EXi
Warranty options	<ul style="list-style-type: none"> • 1 years • 3 Years • 5 Years