Cluster Administration

Nutanix Certified Professional - Multicloud Infrastructure

KNOWLEDGE CHEATSHEET

OPERATE THE CLUSTER

Inventory Discover versions → **Precheck** Health and capacity → **Update** Orchestrate changes → **Validate** Review results

01 Cluster administration

Prism and cluster services

- **Prism Element (PE)** manages one cluster. **Prism Central (PC)** centralizes registered clusters, categories, policies and cross-cluster performance views. Use a PC Viewer role for datacenter-wide read-only visibility.
- The **Prism service** supplies the management interface, nCLI and REST API. **Stargate** handles storage I/O; **Cassandra** stores distributed metadata; **Curator** performs background storage scans and maintenance.
- **Genesis** manages service startup and supervision. **ZooKeeper** provides distributed coordination; **Zeus** manages cluster configuration. A healthy management page alone does not establish storage or host health.
- **NTP** synchronizes clocks for logs, alerts and authentication. A **time zone** changes displayed local time; it does not correct clock drift. Configure DNS, NTP and the cluster time zone during initial setup.

LCM inventory and update sources

- **Life Cycle Manager (LCM)** inventories software and firmware, evaluates compatibility and orchestrates updates. Inventory refreshes component versions and can automatically refresh the LCM framework from its configured source.
- For an **inventory failure**, check source URL, DNS, proxy, firewall and reachability. Match the firewall rule to the configured protocol: HTTP normally uses **80**; HTTPS normally uses **443**.
- A **dark site** uses supported offline bundles. Stage the appropriate bundle on a reachable internal web server, point LCM to the local source and run inventory before choosing updates.
- Run **health and compatibility prechecks** before updates. Keep **NCC** current, resolve failed checks and confirm evacuation capacity. AHV and some firmware updates require host reboots; review each component's impact.

Expansion and maintenance

- Before **cluster expansion**, verify supported hardware, compatible AOS/AHV versions, unique addresses and discovery connectivity. Check management VLANs, allowed traffic and uplink configuration when a new node is not discovered.
- **Maintenance mode** evacuates eligible VMs before servicing a host. Check spare CPU and memory, placement rules and nonmigratable devices before starting; a host-local dependency can prevent evacuation.

Cluster Administration / Storage & Data Protection

Nutanix Certified Professional - Multicloud Infrastructure

KNOWLEDGE CHEATSHEET

OPERATE THE CLUSTER

Inventory Discover versions → **Precheck** Health and capacity → **Update** Orchestrate changes → **Validate** Review results

- For a normal **three-node cluster**, add a replacement node before removing an old node. Removal must preserve the supported cluster size, replica placement and enough capacity to evacuate its data.
- An **agent VM** is tied to a host and is handled differently from an ordinary migratable VM during maintenance. **GPU passthrough** and other host-bound devices also constrain live migration.

Authentication, roles and SSH

- **Authentication** verifies identity; **role mapping** grants access. Configuring AD or LDAP does not by itself authorize every directory user. Map appropriate users or groups to roles and resource scopes.
- **PC RBAC** combines allowed actions with scoped entities, including cluster or category-based access. **PE roles** govern a single cluster. A self-service project groups delegated resources; it is separate from an AHV host permission scope.
- For directory sign-in failures, inspect role mapping, group membership, account lock or expiry, password-change requirements, DNS and time synchronization. Retain a controlled local administrator account for recovery.
- Replace default **CVM, AHV host and IPMI/BMC passwords** separately. **Cluster Lockdown** restricts password-based SSH access to CVMs and AHV hosts and supports administrator public keys; validate key access before relying on it.

02 Storage and data protection

Data path and locality

- A **Controller VM (CVM)** on each node provides distributed AOS storage services. **Stargate** characterizes writes: bursty random I/O can use the **OpLog**; sustained random or sequential writes can go directly to the **Extent Store**.
- The **OpLog** is a persistent low-latency write buffer, not the final location of every write. The **Extent Store** holds persistent extents. Replication protects acknowledged data according to the configured replication factor.
- **Data locality** favors serving a VM's data from its current node, reducing remote reads. Migration can initially require remote reads while locality is reestablished; local access does not eliminate replica traffic.
- **Shadow Clones** cache a shared, read-mostly base vDisk near consuming VMs, useful for linked-clone VDI. **Flash Mode** favors flash residency for selected data; plan for flash capacity, workload size and failure headroom.

Storage & Data Protection

Nutanix Certified Professional - Multicloud Infrastructure

KNOWLEDGE CHEATSHEET

PROTECT DATA AND SERVICE

- **Replication** — Data copies
- **HA reserve** — Restart capacity
- **Snapshots** — Recovery points
- **Remote site** — Disaster recovery

Containers and space efficiency

- A **storage container** is a logical policy boundary within a storage pool. **Advertised capacity** controls the capacity presented for that container. Physical free space and rebuild headroom still depend on the shared cluster.
- **Inline compression** reduces data size as it enters persistent storage. **Deduplication** eliminates repeated content; cache deduplication and capacity savings address different storage tiers. **Thin provisioning** allocates space as data is written.
- **Erasure coding (EC-X)** encodes eligible write-cold data with parity to reduce redundancy overhead. It suits archive-like data better than continuously overwritten hot data and has separate node and failure-domain requirements.
- **Category-based storage policies** apply supported data-service settings to matching VMs. Check policy association, replication-factor compatibility and application state; a new setting need not produce immediate savings on existing data.

Replication and failure domains

- **RF2** keeps two data copies and supports one failure; **RF3** keeps three and supports two, subject to healthy placement and capacity. A standard RF3 cluster needs at least **five nodes**. **RF1** provides no redundant data copy.
- **Node, block and rack awareness** place replicas across independent failure domains. A block is a chassis grouping. Count independent domains, not just nodes, when planning block or rack fault tolerance.
- **Replication factor** protects stored data. **HA reservation** preserves compute capacity to restart VMs after a host failure. **Anti-affinity** separates selected VMs. Each addresses a different failure consequence.
- After a disk or node failure, **Curator scans** identify affected data and coordinate restoration of redundancy. Review **data resiliency**, **Reserve Capacity** and storage alerts together; changing a warning threshold does not restore rebuild headroom.

Replication and recovery

- A **snapshot** preserves a recovery point. A **protection domain** groups protected entities and replication schedules. **RPO** measures acceptable data loss; **RTO** measures acceptable service recovery time.
- **Asynchronous replication** sends scheduled recovery points to a remote site. **NearSync** uses frequent lightweight snapshots for a lower RPO. Replication bandwidth must sustain the workload's changed-data rate.
- If **NearSync** repeatedly falls back, investigate WAN bandwidth, latency, backlog and source/target health. A short configured schedule cannot overcome a transfer path that cannot keep up.
- **Metro Availability** uses synchronous replication between eligible sites. Validate supported hypervisors, network latency and site configuration. A **Witness** supplies an independent failure-decision point; it is not a backup or data replica.

Virtual Machines & Images

Nutanix Certified Professional - Multicloud Infrastructure

KNOWLEDGE CHEATSHEET

PREPARE AND PLACE WORKLOADS

Image — Prepared guest → **Policy** — Eligible clusters → **Deploy** — Drivers and sizing → **Schedule** — Hosts and HA

03 Virtual machines and images

Image service and placement

- The **Image Service** stores ISO installation media and disk images. Mount an ISO as virtual optical media; clone a prepared disk image into a VM disk. Verify image type, transfer status and cluster availability.
- Manage **PC-owned images** through Prism Central. An image visible in Prism Element can have limited local edit actions because its management authority is PC; also check the user's permissions.
- An **Image Placement Policy** associates image categories with cluster categories. Matching categories determine automatic image distribution; verify both sides when an image is absent from a target cluster.
- **Hard enforcement** restricts policy-managed image placement to eligible clusters. **Soft enforcement** distributes to selected clusters while permitting placement elsewhere. Use hard enforcement when residency must be confined.

Guest preparation and sizing

- Windows setup may need the **Nutanix VirtIO SCSI driver** to discover virtual SCSI disks. Attach the VirtIO driver ISO and load the storage driver; a network, balloon or serial driver will not expose the boot disk.
- Install **cloud-init** in a Linux image for deployment-time hostname, users, SSH keys and scripts. **Sysprep** generalizes Windows images. **Nutanix Guest Tools (NGT)** provides guest integration rather than replacing storage drivers.
- Align large VM CPU and memory allocations with physical **NUMA** boundaries. **vNUMA** exposes a useful topology to a NUMA-aware guest; avoid allocating more vCPUs than the workload can use.
- **Memory overcommit** can improve density for suitable workloads but increases pressure and paging risk. Right-size oversized VMs first and preserve HA capacity. Multiple virtual disks can expose storage parallelism when the application uses them.

Placement, HA and scheduling

- **VM-VM anti-affinity** separates related VMs across hosts, such as members of an application cluster. **VM-host affinity** restricts eligible hosts and can reduce evacuation or failover choices.
- **AHV HA** restarts eligible VMs after host failure. Enable **HA reservation** when restart capacity must be protected; enough replica copies alone do not provide CPU and memory for every restarted VM.

Virtual Machines & Images / Networking & Traffic Security

Nutanix Certified Professional - Multicloud Infrastructure

KNOWLEDGE CHEATSHEET

TRACE THE NETWORK PATH

VM NIC Guest network → **OVS** VLAN and bridge → **Bond** Physical uplinks → **Fabric** Trunks and LACP

- **Acropolis Dynamic Scheduling (ADS)** can migrate VMs to reduce host resource contention. Inspect task history, host demand and placement constraints when a VM moves unexpectedly. Migration does not cure a guest's undersizing.
- Before rolling updates, check affinity rules, host-bound devices, free capacity and VM migration readiness. Treat constraints as workload requirements; validate availability before changing them to simplify maintenance.

04 Networking and traffic security

Virtual switches and VLANs

- AHV networking uses **Open vSwitch (OVS)**. The default **vs0** virtual switch manages the **br0** bridge across hosts; **br0-up** is the usual default uplink bond. Manage supported switch settings consistently through Prism.
- A **tagged VLAN** must be allowed on every relevant upstream trunk. Same-host communication can succeed while cross-host communication fails if the physical fabric does not carry that VLAN.
- AHV **VLAN 0** denotes untagged traffic on a VLAN-backed network. Match the switch's native VLAN; entering its numeric VLAN as a tag can send traffic differently from the expected untagged path.
- A **managed network** provides Nutanix IP address management and DHCP from a configured pool. An **unmanaged network** relies on external address assignment. Verify subnet, gateway, pool and guest DHCP configuration.

Uplink bonds and failover

- **Active-Backup** uses one uplink at a time with standby failover. **Active-Active with MAC pinning (balance-slb)** distributes VM NICs across uplinks without switch-side link aggregation; one VM NIC remains limited to one uplink's rate.
- **Active-Active (balance-tcp)** uses **LACP** and distributes flows across uplinks, allowing aggregate traffic from one VM NIC to use several links. A single flow remains limited by its selected link and hashing behavior.
- Use compatible NICs of the **same speed and vendor/model** within a bond and redundant physical paths. Validate identical guest VLAN allowances on both switch ports; management traffic surviving a failover does not prove guest connectivity.
- A physical NIC already assigned to **vs0** cannot also be assigned to another virtual switch. Check current membership even for disconnected NICs. Reconfigure uplinks carefully because host and CVM reachability use the physical fabric.

Networking & Traffic Security / Monitoring & Troubleshooting

Nutanix Certified Professional - Multicloud Infrastructure

KNOWLEDGE CHEATSHEET

TRACE THE NETWORK PATH

- **VM NIC** — Guest network
- → **OVS** — VLAN and bridge
- → **Bond** — Physical uplinks
- → **Fabric** — Trunks and LACP

Segmentation and Flow policies

- AOS **network segmentation** separates selected infrastructure traffic, such as management and storage replication. **Flow microsegmentation** applies distributed security policy to workload traffic. VLAN separation alone does not define application access rules.
- **Flow Virtual Networking** supplies VPCs, overlay subnets and network connectivity services. Its network controller uses PC **Microservices Infrastructure**; verify prerequisites before enabling advanced networking.
- An **Application Security Policy** defines permitted application traffic using categories and rules. Use **enforce mode** to block unauthorized communication; monitor mode observes behavior without applying the same blocking protection.
- An **Isolation Environment policy** separates groups of workloads. Choose isolation for group boundaries and application policies for explicit allowed communication. Check categories, rule direction and allowed endpoints before enforcement.

05 Monitoring and troubleshooting

Metrics and bottlenecks

- **VM Storage Controller Bandwidth** measures bytes transferred per unit time. **IOPS** measures operation rate; **latency** measures response time. Match the metric's entity and time window to the affected VM rather than substituting host totals.
- High **CPU ready time** indicates waiting for physical CPU scheduling. Sustained guest memory pressure or **swap rates** suggest memory constraints. Low storage latency does not rule out a CPU or memory bottleneck.
- For remote desktop interruptions, examine **vNIC receive drops**, network paths and guest behavior. For CAD or rendering delays, inspect **GPU usage** and graphics configuration along with CPU, memory and storage metrics.
- Use **dashboard drill-down** and **Analysis charts** to compare affected VMs, hosts and clusters over the same interval. Correlate changes with tasks and events before resizing resources or changing storage settings.

Monitoring & Troubleshooting

Nutanix Certified Professional - Multicloud Infrastructure

KNOWLEDGE CHEATSHEET

INVESTIGATE WITH CONTEXT

Symptom — Affected service → **Metrics** — Entity and interval → **Correlate** — Tasks and events → **Act** — Validate improvement

Efficiency, anomalies and alerts

- A **constrained VM** needs more of a limiting resource; an **over-provisioned VM** has excess allocation. A **bully VM** consumes shared resources and affects neighbors. Use VM Efficiency findings to guide workload-specific right-sizing.
- A **zombie VM** remains powered on with negligible sustained activity; a **dead VM** has remained powered off. Confirm ownership, business schedules and dependencies before reclaiming either VM's resources.
- A **behavioral anomaly** is a deviation from a learned expected range, including unusually low values. A **static threshold** compares against a fixed limit. The predicted band can change with learned patterns.
- A **custom alert policy** combines metric, entity scope, condition and severity. Available warning/critical fields depend on policy type. Review overlapping policy scopes and effective precedence instead of assuming every matching policy applies independently.

Health, logs and automation

- **Nutanix Cluster Check (NCC)** evaluates cluster health. Run checks from the PE Health dashboard or use **ncc health_checks run_all** on a CVM. Inspect failed checks and dependencies before retrying upgrades or maintenance.
- **Remote syslog** needs a destination address, port and transport, plus selected modules and severity. Review **Include Monitor Logs** separately from severity filtering when more monitor messages arrive than expected.
- Configure an **SMTP server** for email delivery. **STARTTLS** upgrades a supported SMTP session to TLS; the port number alone does not enable encryption. Test server reachability, authentication and message delivery.
- A **playbook** connects an event or alert trigger to actions, such as generating or sending a report after host failure. Use **CSV report attachments** when recipients need sortable, filterable VM metrics without console access.