STORAGE FABRIC

Fast Lane

Storage Capabilities

Storage Constructs

Reads and Writes

Failures

CAPABILITIES

Capabilities listed

- Intelligent Tiering
- Data Locality
- Automatic Disk Balancing
- Data Path Redundancy

INTELLIGENT TIERING

Fast Lane

- Continuously monitors data access patterns
  - Optimizes placement on HDD or SSD tier
  - No Administrative intervention
  - Maximum performance for hot data and random I/O
  - Maximum capacity for cold data and sequential I/O

DATA LOCALITY

Fast Lane

- Store the data of a VM to on the node that runs the VM
  - This prevents I/O to go through the network
  - Optimizes performance
  - Data follows the relocation of a VM based on read patterns
    - when the VM is moved
    - during an HA event

AUTOMATIC DISK BALANCING

Fast Lane

- Allows for different workloads and different node types
  - When storage utilization on a node crosses a threshold
  - Movement of data always between to same storage tier
  - Different storage capacities of nodes are taken into account (compute-heavy vs. storage-heavy nodes)

DATA PATH REDUNDANCY

Fast Lane

- In the case of an CVM unavailability
  - Requests are rerouted to a healthy CVM on another node
  - Failover is fully transparent to hypervisor and applications

STORAGE POOLS

- A collection of different drive types spanning multiple nodes
- CVMs communicate and create a single storage pool (clusterwide)

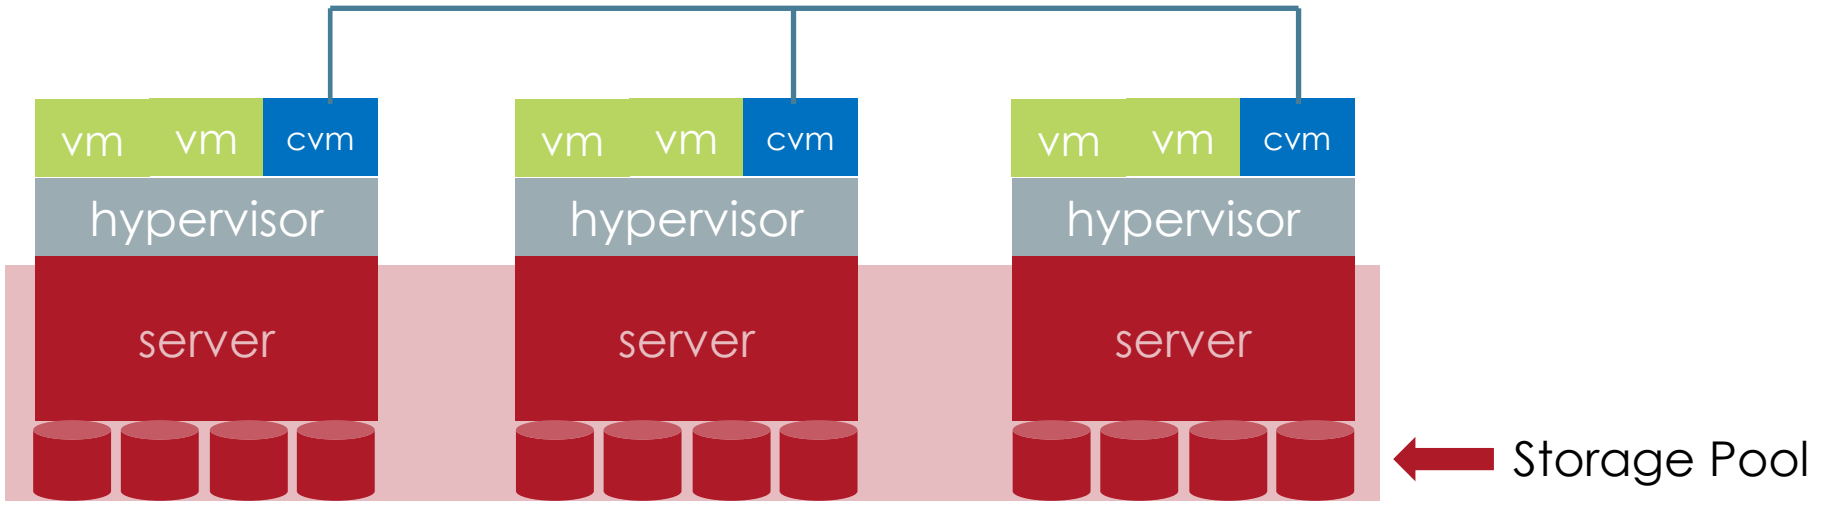

STORAGE CONTAINER

- A container is logical segmentation of the storage pool
- A container contains a group virtual disks (vdisks)
- Thin provisioned
- Can be compared to a '*Vmware datastore*'

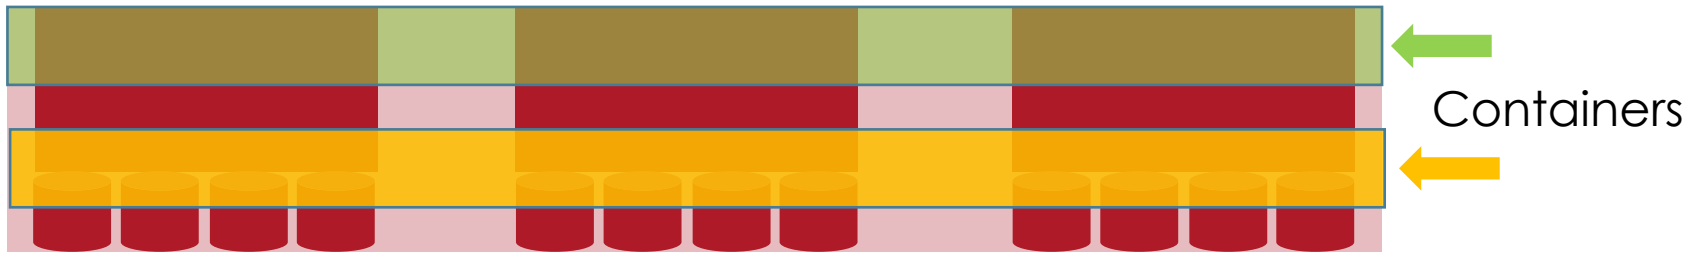

WRITING IO

- Data is written to the local node the VM is running on
- With **replication factor 2**, data will be written to two locations
- In the example the original copy is written to the first node
  The second copy is written to another node in the cluster

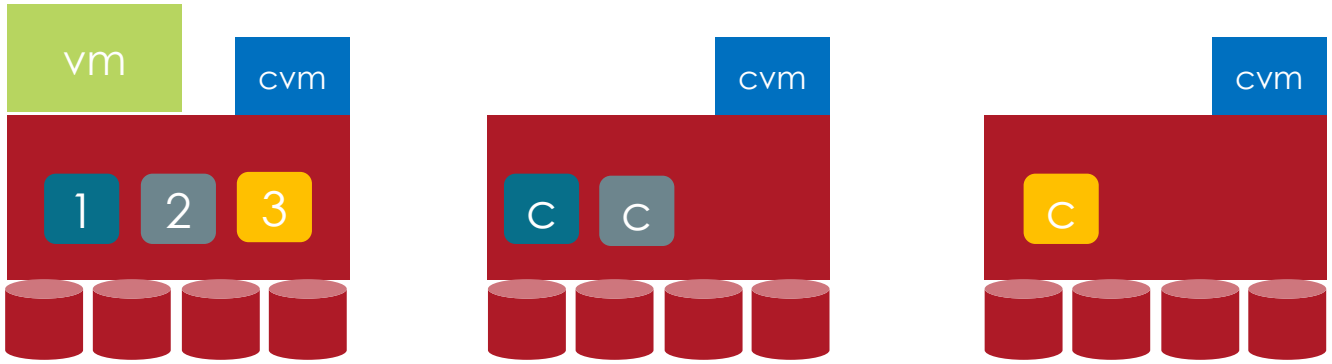

READING IO

- Data is read from the node on which the VM resides

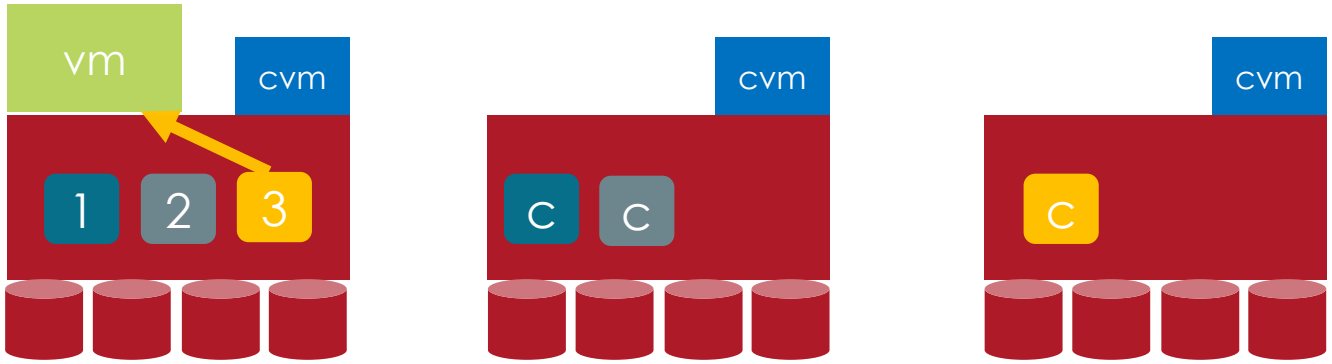

READING IO

- When data is read from another node the data will then be copied
- In this example the vm has migrated to the second node and the data was on the first node

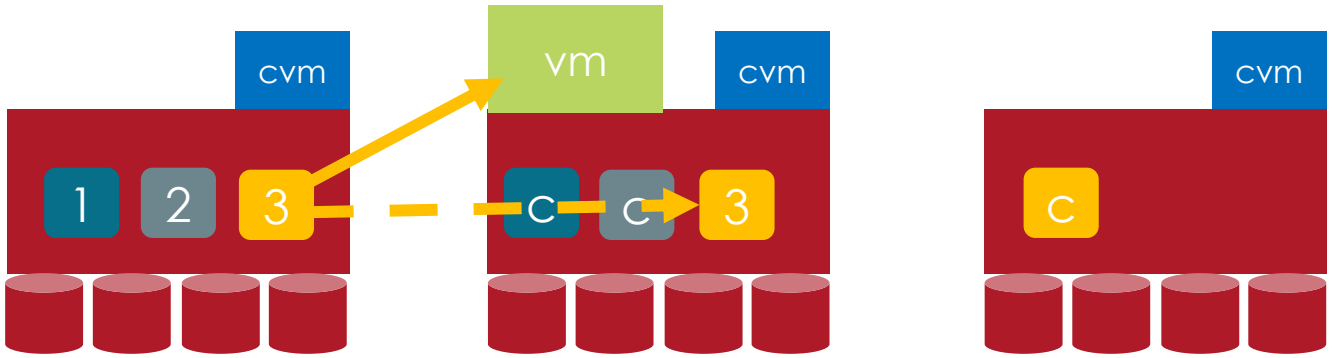

STORAGE CONSTRUCTS LISTED

- Storage Pool
- Storage Container
- vDisk
- Volume Group

STORAGE CONSTRUCTS LISTED

- vBlock
- Extent
- Extent Group
- OpLog

BOOT DRIVE FAILURE

- CVM will fail
- Other VMs are unaffected
- Data Path Redundancy will redirects the path to another CVM

DATA DRIVE FAILURE

Fast Lane

- Single drive failure will not result in data loss
- Cluster create a new replica of the data
- With replication factor 2, losing a second drive can result in data loss

CVM FAILURE

- Address becomes unavailable
- Data traffic is redirected transparantly
- Self-healing will transfer traffic back to original CVM when it is back

HOST FAILURE

- HA-protected VMs are restarted on other nodes
- Stargate will start migrating extents
- Curator instructs Stargate to create a new replicas

STORAGE CONTAINERS

- Three default storage containers
  - NutanixManagementShare
  - SelfServiceContainer
  - Default-container

UPDATING STORAGE CONTAINERS

Fast Lane

- Some conditions when updating a Storage Container
  - Cannot be renamed through the Update Storage Container dialog
  - Cannot be renamed if it contains vdisks
  - Replication factor cannot be changed when updating the Storage Container

RESERVED AND ADVERTISED CAPACITY

- By default a Storage Container can use all capacity in a Storage Pool
  - Reserved Capacity reserves a minimum amount of space
  - Advertized Capacity sets the maximum amount of space
  - Advertized Capacity should be larger than Reserved Capacity

RESERVED AND ADVERTISED CAPACITY

- By default a Storage Container can use all capacity in a Storage Pool
  - Reserved Capacity reserves a minimum amount of space
  - Advertized Capacity sets the maximum amount of space
  - Advertized Capacity should be larger than Reserved Capacity

STORAGE OPTIMIZATION FEATURES

Fast Lane

- Compression
- Deduplication
- Erasure Coding

COMPRESSION

- Inline and post-process
- Inline compression is enabled by default
- Large I/O and sequential I/O is not treated the same as random data

DEDUPLICATION

Fast Lane

- Dupliate blocks are consolidated
- Not enabled by default
- Cache deduplication
- Capacity deduplication
- Additional RAM needed for CVM

ERASURE CODING

Fast Lane

- Increases usable disk space without compromising resilience
- Stripes individual blocks across nodes
- Data blocks and parity blocks
- Most suitable for cold data, archives and backups

COPY EXAMPLE

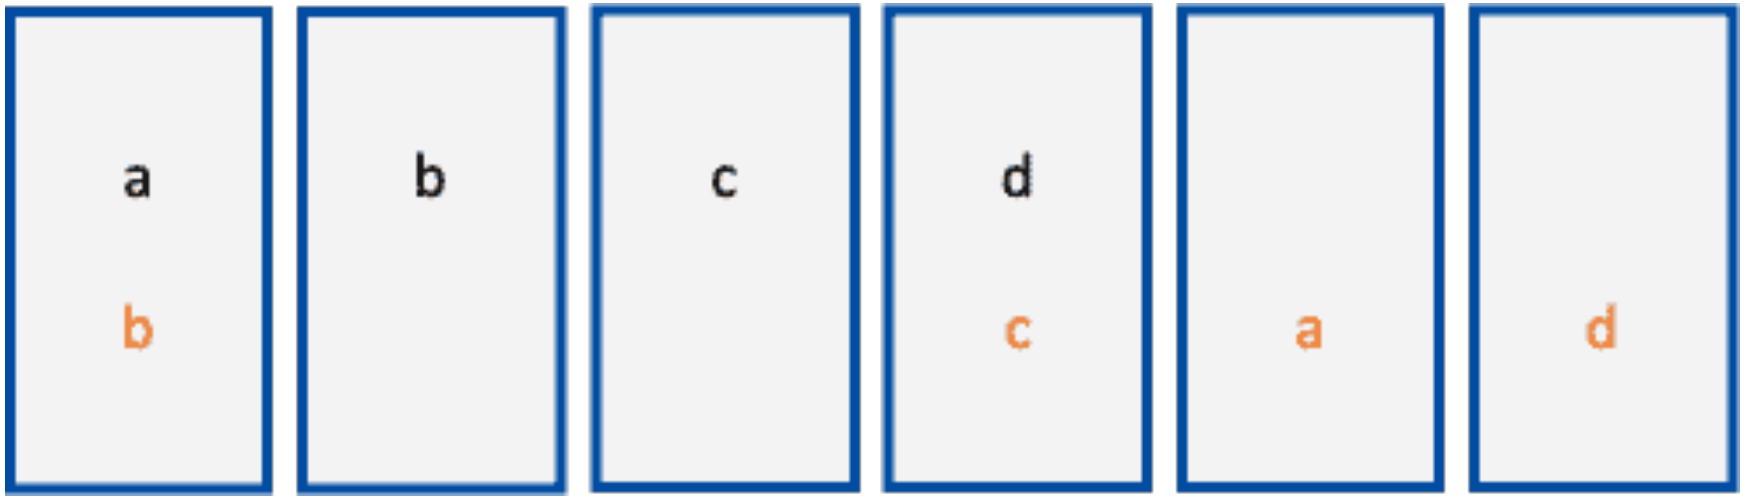

- 6-node cluster
- Redundancy factor 2

ERASURE CODING EXAMPLE

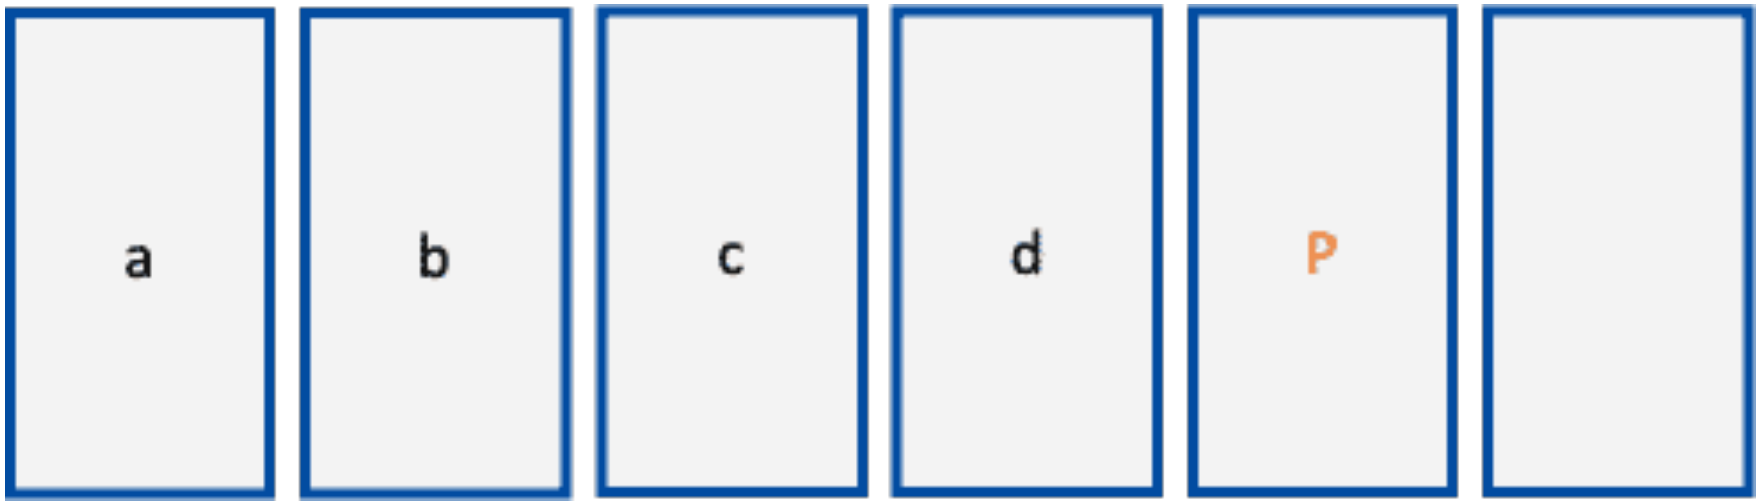

- Parity is calculated
- Copies are removed and replaced by Parity block

SUMMARY

- Dataprotection
- Reads and writes
- Failover and optimalization