

# Geographic Logical Volume Manager (a203869)

—  
Antony (Red) Steel  
Belisama Pte. Ltd.



**IBM Champion**

noun \ 'champ-pé-on \

They're experts. They're leaders.

IBM Champion

IBM

`#include <std_disclaimer.h>`

These notes have been prepared by an Australian, so beware of unusual spelling and pronunciation.

All comments regarding futures are probably nothing more than the imagination of the speaker and are IBM Confidential till after GA.

Thanks to:  
Steven Finnes  
Ravi Shankar  
Shawn Bodily

# Agenda

- Background and plans
- Concepts
  - Examine the operation and history of the Geographic Logical Volume Manager:
    - Synchronous mode
    - Asynchronous mode.
- Operation
  - Primary site failure.
  - Secondary site failure.
  - the differences in configuration between sync and async modes.
- Planning
  - Design and sizing
- Monitoring
- Configuration
- (Backup slides with detailed examples of configuration)



# Background

- Replication over IP for AIX started with GeoRM / HAGEo.
- Standalone GLVM provided as part of AIX 5.3 in October 2005.
  - AIX filesets
    - Remote Physical Volume Client      glvm.rpv.client
    - Remote Physical Volume Server      glvm.rpv.server
    - Geographic LVM Utilities              glvm.rpv.util
- HACMP/XD GLVM initially released in April 2005 on HACMP/XD 5.2 (to replace HACMP/XD HAGEO).
- Initial release characteristics:
  - Single data mirroring network only.
  - Synchronous mirroring.
  - Supports file systems or raw logical volumes.
  - GMVG could not be an enhanced concurrent mode vg.
- GLVM 5.4
  - Multilink support – multiple networks type XD\_data (up to 4) – improved throughput and reliability.
  - Support for enhanced concurrent mode volume groups, but not concurrent access from both sites.
- GLVM 5.5
  - Asynchronous support (so GeoRM removed from availability).
  - Uses mirror pools (feature of AIX 6.1 TL2).



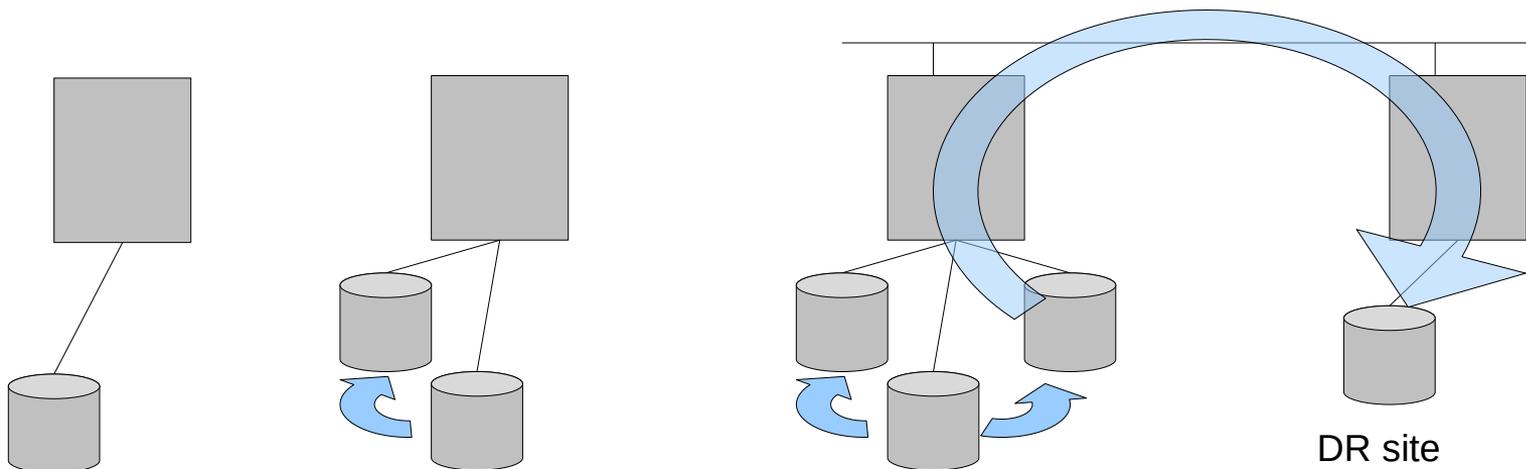
# GLVM Stand alone or with PowerHA SystemMirror EE

- GLVM can be used stand alone, but need careful management as GLVM itself has no concept of sites or status of the remote components. Operation can be scripted, but care must be taken as relatively easy to corrupt or loose data.
- Operation with PowerHA SystemMirror Extended Edition:
  - PowerHA provides the framework to monitor and manage the operation of GLVM to control the replication of data over IP network(s) and adds:
    - A configuration assistant with enhancements to convert existing Volume Groups to GMVGs and include into the appropriate resource group.
    - Delete and rollback features added.
    - Both synchronous and asynchronous modes are supported in the PowerHA GUI.
    - The following networks are included:
      - XD\_data network(s) – Up to 4 supported, with data sent round-robin over the available networks.
      - XD\_ip network for the RSCT-based keepalives.
    - LVM preferred read management to manage read from the local site copy.
- PowerHA can ensure that GLVM data is available at only one site, if running standalone you are responsible for managing this and avoiding concurrent access to data at both sites.

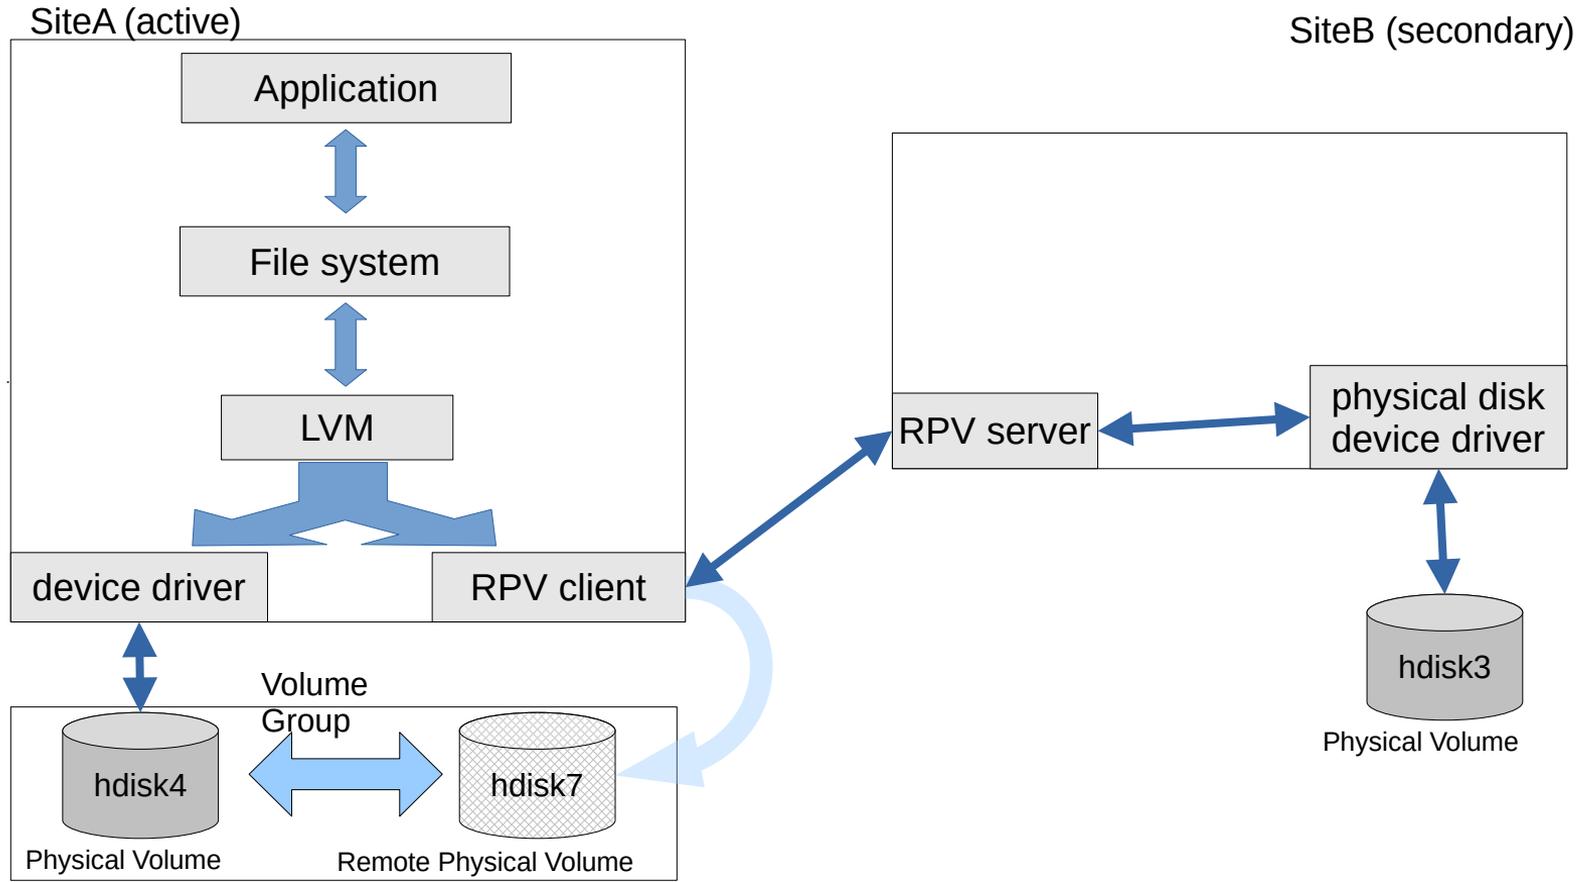


# Concepts

- Why? Provides solutions for:
  - Non IBM storage, Only have IP between sites, Applications with no replication, Need to move or replicate data between sites....
- Geographic Logical Volume Manager (GLVM) builds on the AIX LVM – adding the capability to mirror to a remote machine over a TCP/IP network.
- Advantage over GeoRM is that it is built into / part of the AIX LVM.
- Now asynchronous support and multiple networks are supported.



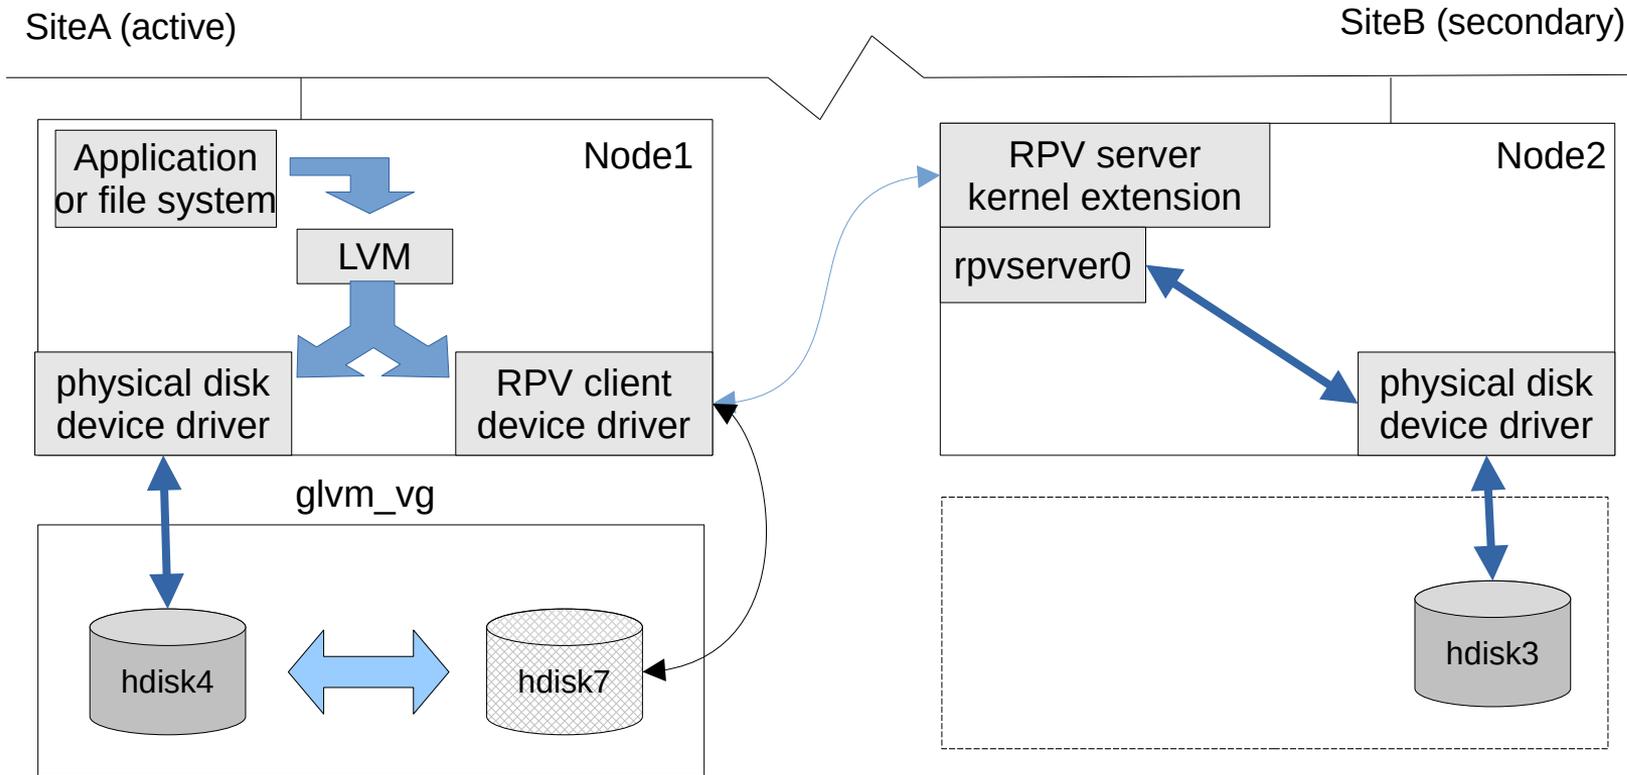
# Concepts – I/O path



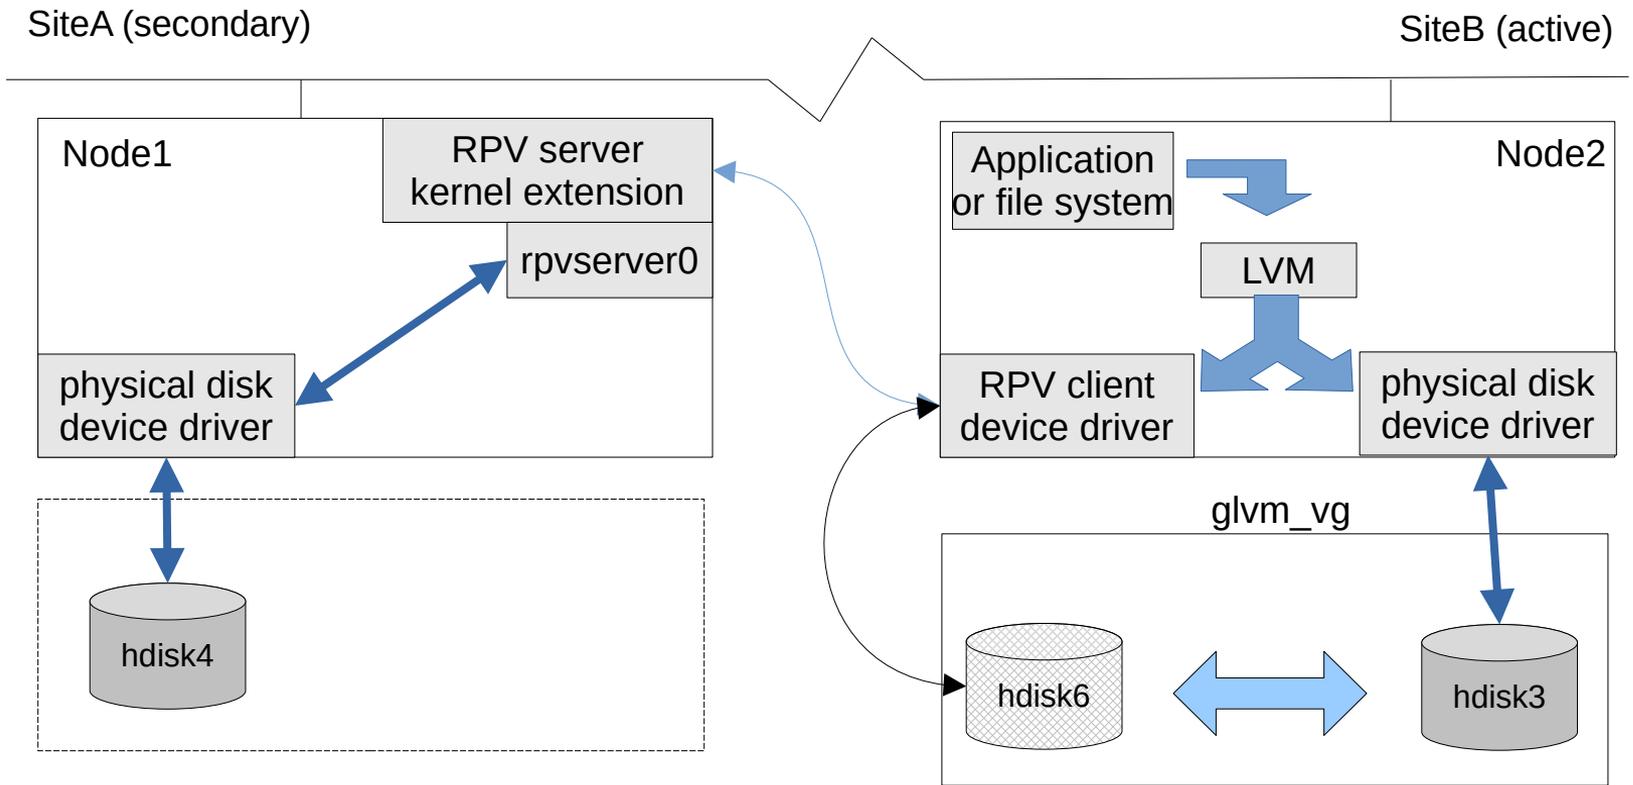
# Concepts - terminology

- GLVM
  - Geographic logical volume manager
- Remote Physical Volume (RPV)
  - Pseudo device driver that provides access to the remote disks as if they were locally attached.
  - Remote system connected via tcp/ip.
  - Distance limited by application latency/performance requirements
- RPV Server
  - kernel extension of the RPV device driver
  - One on remote machine for each replicated PV
  - Each one points to local PV
  - named rpvserverN
- RPV Client
  - The pseudo device driver (hdiskN)
  - LVM sees as hdisk
  - Definition includes local and remote server addresses and the timeout
- GMVG
  - Geographically mirrored volume group

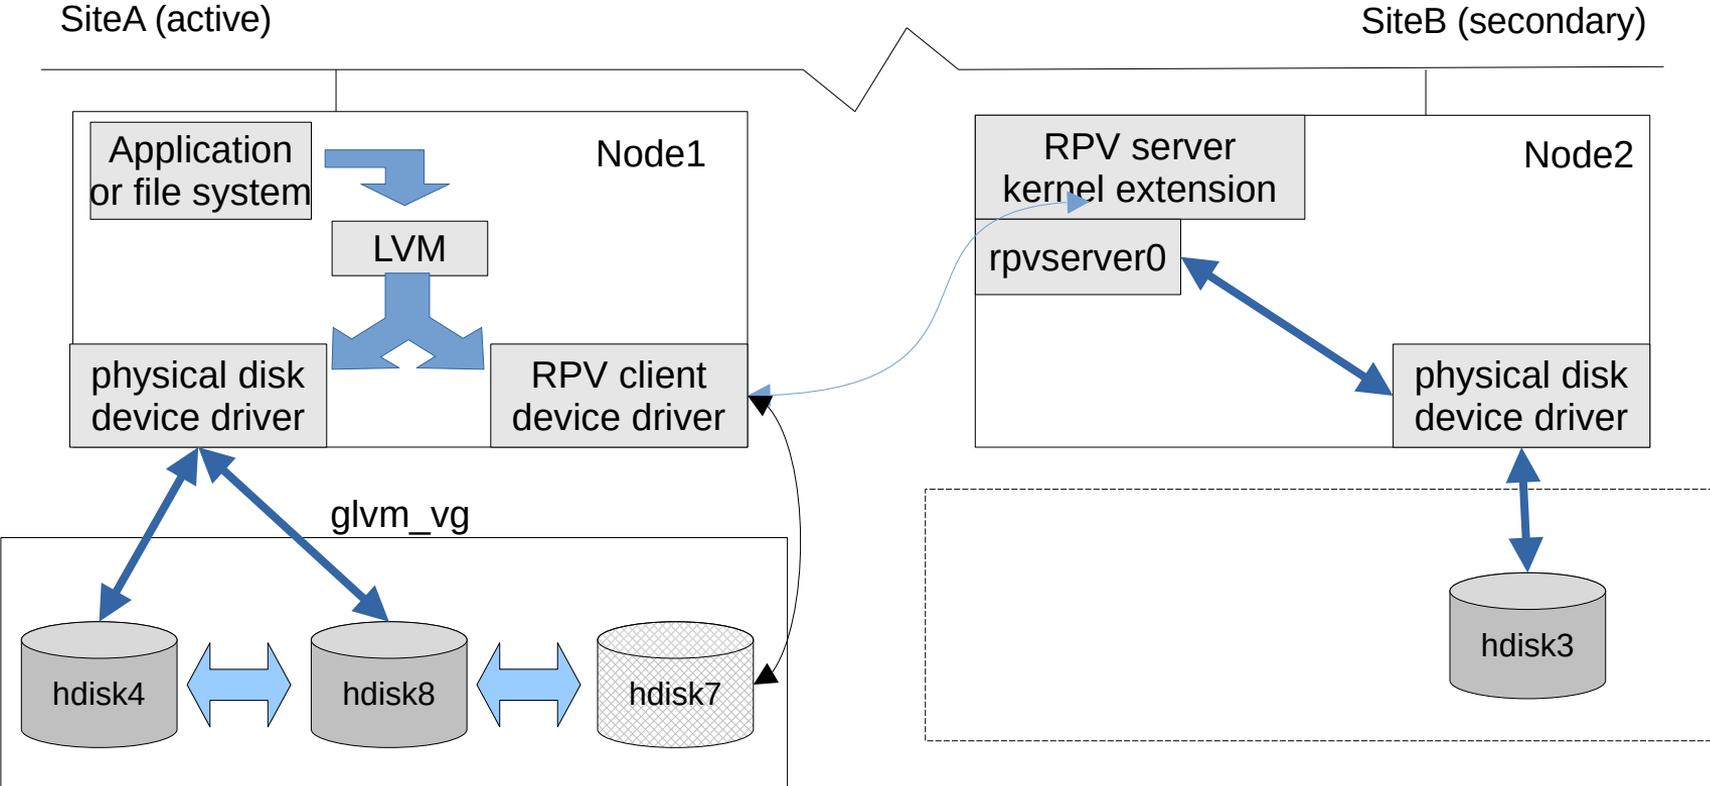
# Concepts – Site A active



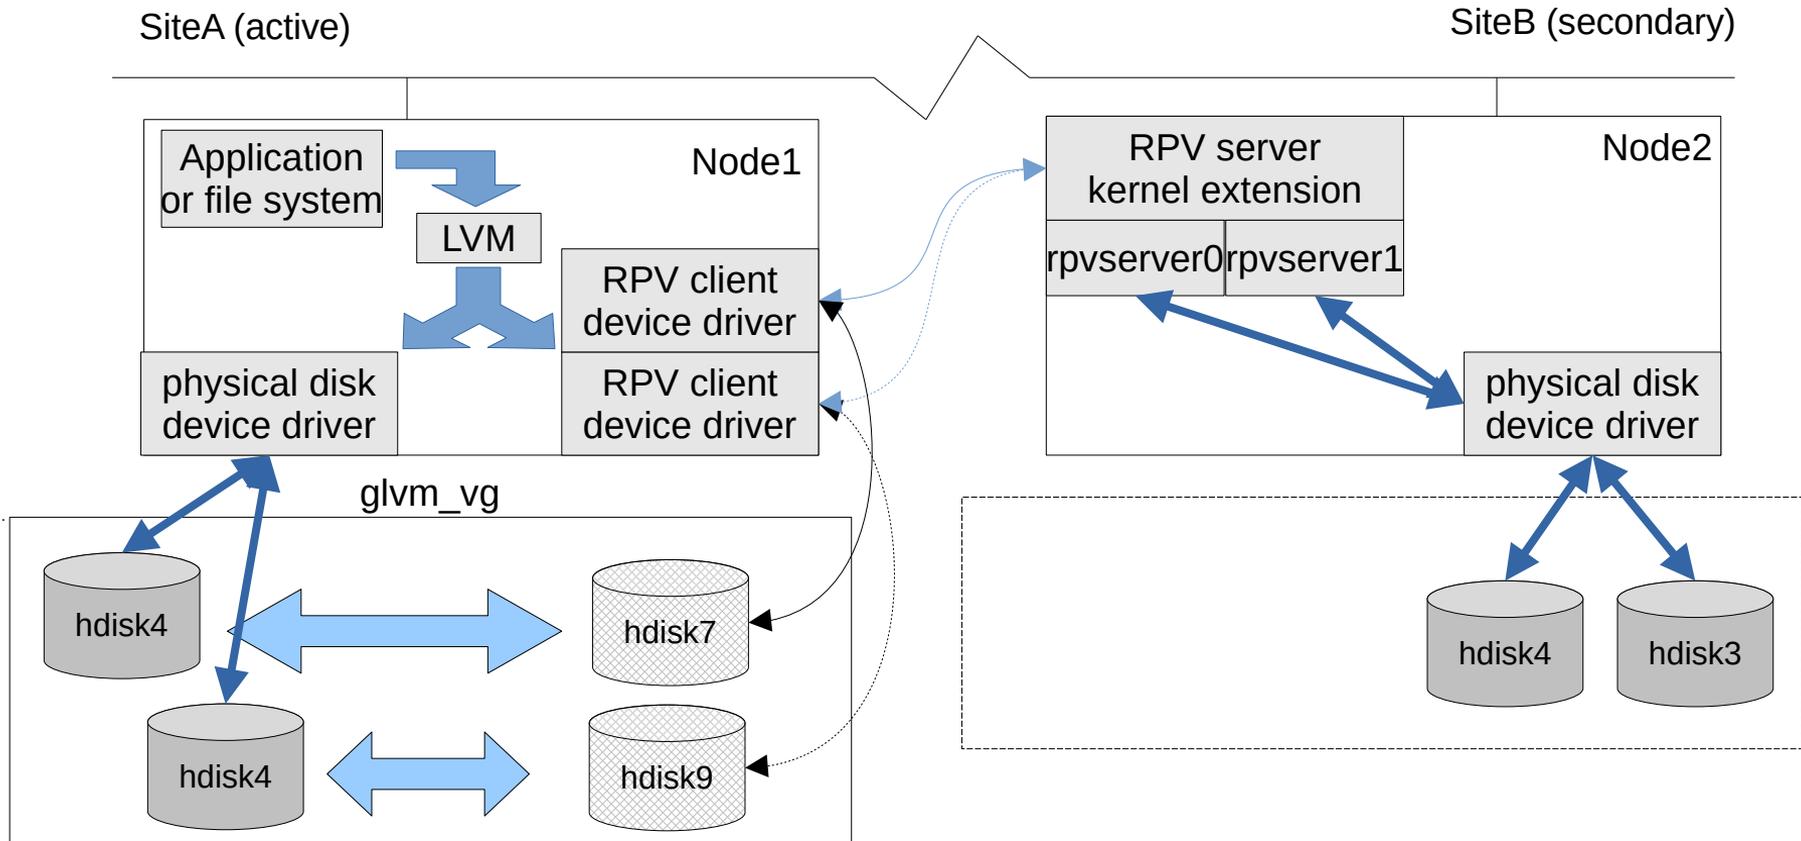
# Concepts – Changed to Site B active



# Concepts – Two copies of mirror at Site A

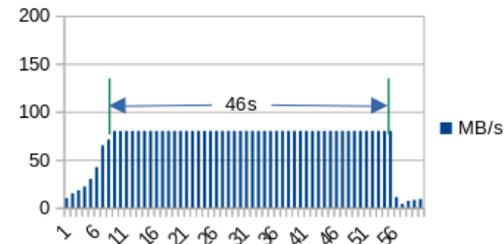
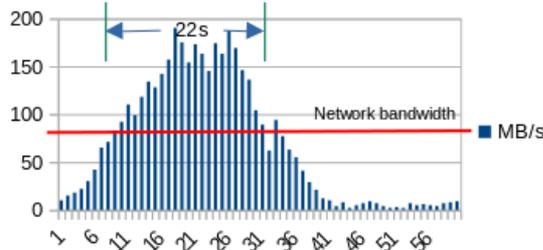


# Concepts – Two mirrored LUNs



# Limitations of synchronous GLVM

- How synchronous mode may impact your application:
  - Time to write to remote PV has greatest impact
    - depends on network latency and bandwidth
  - Bandwidth limiting factor when the volume of data is excessive
    - Network must be sufficient to meet peaks in I/O to guarantee an acceptable response time
    - If peaks rare, excessive network cost required and rarely used.
  - Latency also important
    - Even if bandwidth is sufficient, latency adds to every I/O operation
    - Limits GLVM practically to under roughly 100km and may require costly low latency equipment
- Benefits of Asynchronous GLVM
  - Network delay removed as an issue
  - I/O done returns to application after local write completes



# Introduction to asynchronous mode

- The following changes introduced to support asynchronous mode.
  - AIX 6.1 Mirror Pools required for asynchronous mode (a good idea for synchronous!).
  - AIX LVM has been enhanced to handle asynchronous mirroring:
    - New commands `chmp` and `lsmp`.
    - Changes to `varonvg` and `varyoffvg`.
    - Support for `aio_cache` logical volume.
    - New functionality to handle asynchronous writes, included failure of deferred write.
  - RPV device driver changed to process write requests asynchronously:
    - Includes a new async I/O cache to record asynchronous RPV writes on the client side.
    - Expand GLVM utilities for Asynchronous GMVG fast path.
    - New path for mirror pool based GMVG (including the asynchronous GMVGs).
    - Retain path for legacy GMVGs.
  - Enhance PowerHA SystemMirror Enterprise Edition to support asynchronous GLVM.
    - New code for configuration and recovery.
  - Maintain asynchronous statistics for the RPV device driver.
    - New options in the `rpvstat` command.



# Asynchronous mode

- Benefits of asynchronous GLVM:
  - Writes cached at active site, remote writes complete in the background.
    - Hides both latency and bandwidth induced delay in replication.
    - Cache must be non-volatile to recover node crash.
    - Choosing asynchronous mode implies there will be a data loss in event of a failure. The size of the cache controls the maximum amount of data that will potentially be lost.
    - When cache full, all I/Os wait until space is free in the cache.
      - Hence the cache must be big enough, but not too big.
  - Can maintain acceptable performance through peaks in I/O spreading out the network load.
  - If customer can afford some data loss in the event of a disaster, this offers many benefits
    - Greater distances (protect against larger disasters)
    - Smaller bandwidth required for same workload
    - More flexible and suits more customer environments than synchronous GLVM.
- Drawbacks of asynchronous GLVM
  - More complex to plan and administer, particularly recovery where the issue of data divergence may have to be addressed.
  - Data will be loss if production site fails
  - Extra steps taken to ensure that there is a copy of each mirrored partition at each site. Overcome with use of mirror pools – required for asynchronous mode, recommended for synchronous.

# Introduction to AIX Mirror Pools

- Mirror pools:
  - Introduced in AIX 6.1 TL2 for Scalable Volume Groups.
  - Are just a collection of disks within a Volume Group. Each disk can be in only 1 mirror pool.
  - Each logical volume copy can be assigned to a mirror pool, therefore only partitions from disks in that pool will be allocated.
  - Characteristics can be changed, but will not affect existing allocations. Use a reorgvg to ensure all allocated partitions meet the pool restrictions.

```

Change Characteristics of a Physical Volume

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

* Physical volume NAME          [Entry Fields]
Allow physical partition ALLOCATION?  yes          +
Physical volume STATE           active       +
Set hotspare characteristics     n           +
Set Mirror Pool                [site-a]      +
Change Mirror Pool Name          []           +
Remove From Mirror Pool

F1=Help      F2=Refresh      F3=Cancel      F4=List
F5=Reset     F6=Command    F7=Edit       F8=Image
F9=Shell    F10=Exit       Enter=Do
  
```

```

Add a Logical Volume

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

[Entry Fields]
[TOP]
Logical volume NAME          [ulv11]
* VOLUME GROUP name          datavg
* Number of LOGICAL PARTITIONS [1]          #
PHYSICAL VOLUME names       [hdisk0 hdisk4]  +
Logical volume TYPE          [jfs2log]      +
POSITION on physical volume  middle      +
RANGE of physical volumes    minimum      +
MAXIMUM NUMBER of PHYSICAL VOLUMES to use for allocation [ ]          #
Number of COPIES of each logical partition                2          +
Mirror Write Consistency?    passive        +
Allocate each logical partition copy on a SEPARATE physical volume? superstrict +
...
Mirror Pool for First Copy          site-a      +
Mirror Pool for Second Copy       site-b      +
Mirror Pool for Third Copy
Infinite Retry Option        no              +

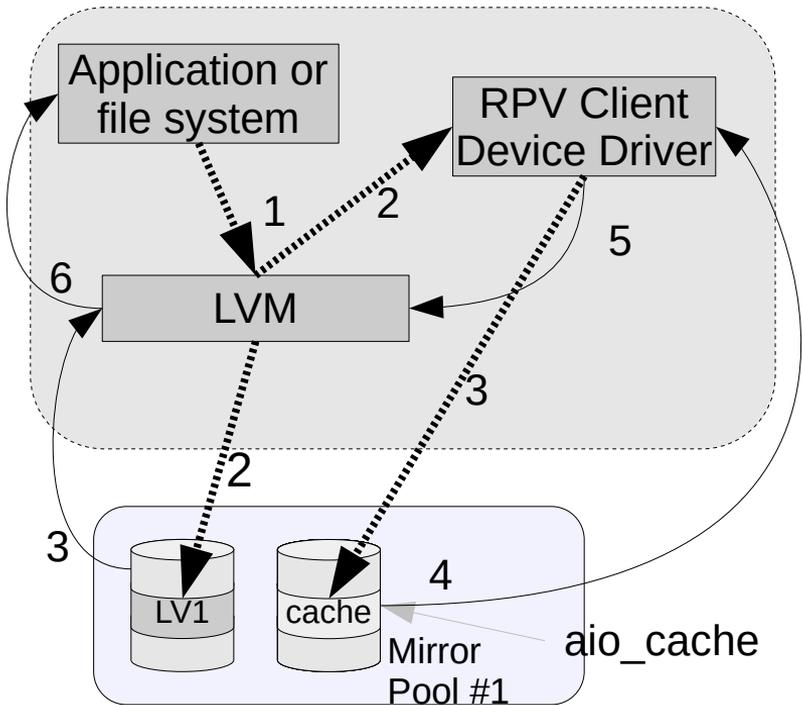
F1=Help      F2=Refresh      F3=Cancel      F4=List
F5=Reset     F6=Command    F7=Edit       F8=Image
F9=Shell    F10=Exit       Enter=Do
  
```



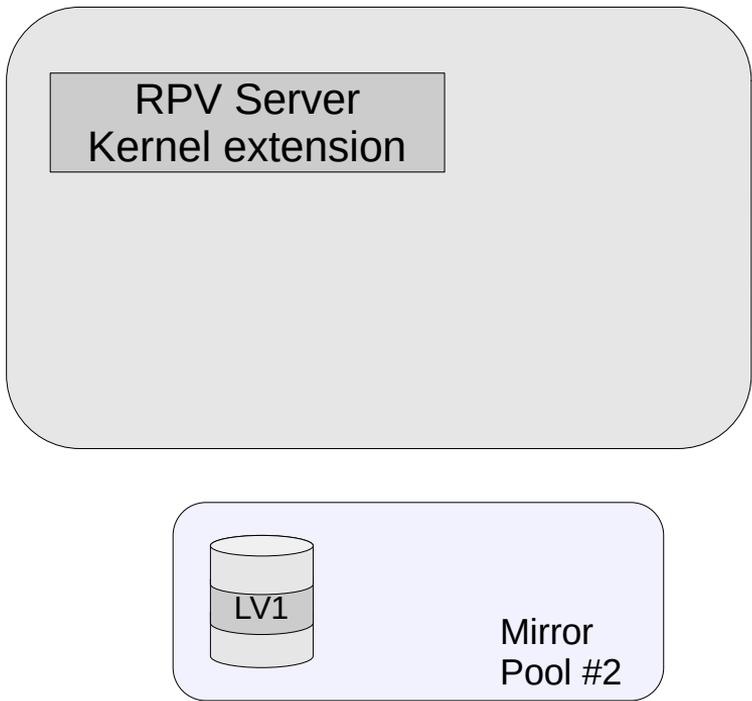


# Async mode GLVM (cont) – local write

Production



DR

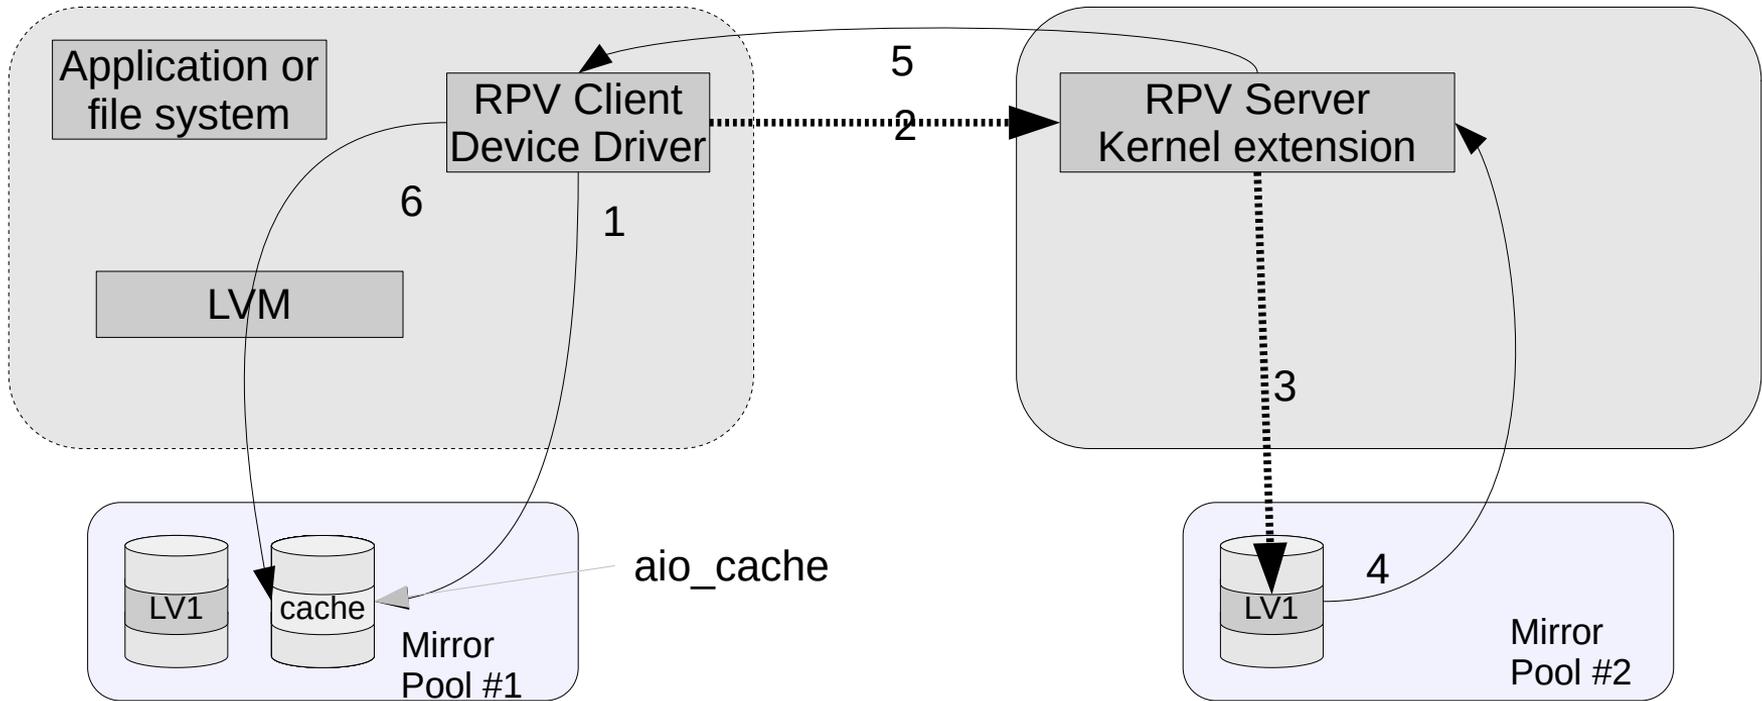


# Async mode GLVM (cont) – remote write

(at sometime later)

Production

DR



## Async mode GLVM (cont)

- Summary
  - Build on AIX mirror pools:
    - A convenient way to group disks in a VG with local and remote disks are in different pools.
    - LVM commands allocate Physical Partitions from different pools, so each site has a complete copy.
    - Works with synchronous as well as asynchronous (this is a property of the Mirror pool).
    - You can have mixed synchronous and asynchronous Mirror Pools in one VG so you can mirror sync one way and async back.
    - Each asynchronous Mirror Pool requires a aio\_cache LV, for the RPV to cache remote writes.
    - The aio\_cache LV cannot be mirrored across sites and there is only one per pool.
    - The aio\_cache LV that is responsible for the writes to Mirror Pool#1 actually belongs to Mirror Pool#2.
    - The RPV Device Driver informs the LVM that the write is complete when it is stored in the the aio\_cache.
    - If a node crashes with an asynchronous VG active, then on recovery, the contents of the aio\_cache can be sent to the remote site, ie no greater data loss that you would expect with normal VGs.
  - After creation, the management is similar to that of the synchronous GMVGs
    - varyoffvg takes the VG offline, However it waits till the aio\_cache is empty before completing (so could take a while if the cache contains large amounts of data).
    - Complexity not apparent until the possibility of data divergence raises it's head.



- Summary

- Logical volumes must be super strict with mirror pools defined for each copy for asynchronous.
- Volume groups must be scalable (can be enhanced concurrent mode, but will not support concurrent access across sites)
- rootvg and paging spaces not supported.
- Cannot snapshot the Volume Group.
- Bad-block relocation should be turned off for each logical volume.
- Each copy of logical volume must be allocated to a pool.
- Only one asynchronous mirror pool across 2 sites is supported.



## Async mode GLVM (cont)

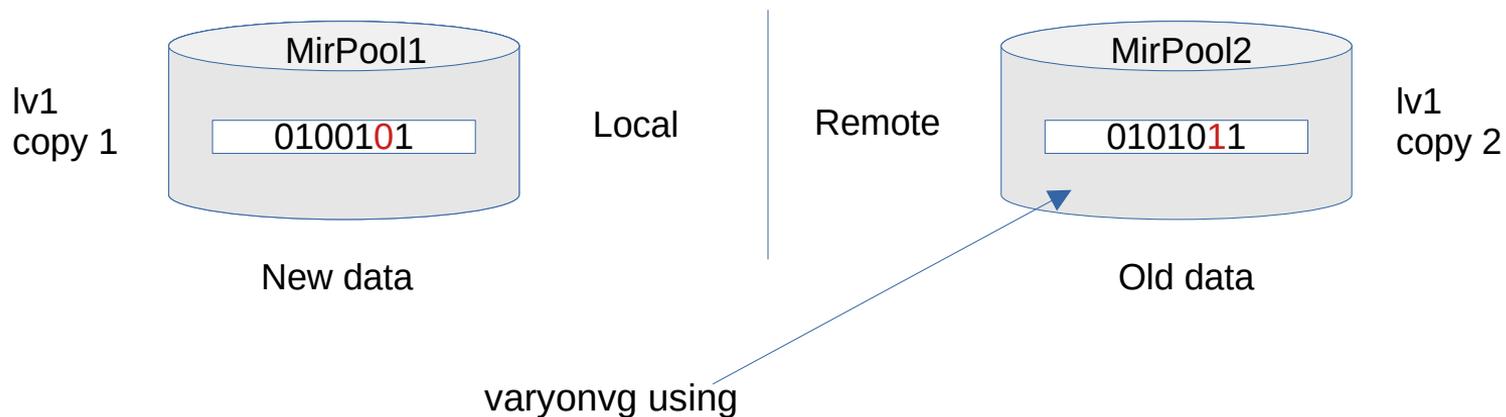
- Examples of the changes to varyonvg for mirror pools
  - Modified to handle recovery for async mirroring – specifically data divergence
    - varyonvg -k loc | rem
    - To specify which copy of the data to keep
      - loc – keep the data from the local mirror pool
      - rem – keep the data from the remote mirror pool



- If there are stale partitions in the copy specified with the -k flag (local or remote), then the varyon will fail. The new flag (-o) will force and allow you to use the copy which may have partitions that are marked stale.
- This flag can only be used with the -k flag

# Async mode GLVM (cont)

- Examples of the changes to varyonvg for mirror pools (cont)
  - If you attempt to varyon a VG, and the system thinks that there is data in the remote cache which is inaccessible, then the user is given a warning. There is a new flag (-d) to force the varyon in this case – using potential back level data.



## Async mode GLVM (cont)

- Details of the addition flags for varyonvg command (-d; -k; -o)
  - -d:

Allows data divergence. The -d flag only takes effect when you try to bring the volume group online while the cache at the opposite site might contain unmirrored data updates and that cache is not accessible. If the varyonvg command detects that you might use back-level data and you do not specify the -d flag, the command fails with a warning.
  - -k loc | rem

Keeps data from the local mirror copy or remote mirror copy. You can specify the following attributes with the -k flag:
    - loc Retains the local mirror copy data. Local means local physical volumes and not primary site
    - rem Retains the remote mirror copy data. Remote means remote physical volumes and not remote site.
- -o

Allows using data from partitions that are stale in the copy you selected but fresh in the other copy. The varyonvg command fails if you specify the -k flag to preserve either local copy or remote copy in the data divergence case and the varyonvg command cannot preserve the complete copy because some partitions are not fresh in the local or remote copy that you selected. You can override the failure by specifying the -o flag to use data from partitions that are stale in the copy that you selected but fresh in the other copy. The -o flag is only valid with the -k flag.



## Async mode GLVM (cont)

- Changes to varyonvg for mirror pools (cont)
  - New exit codes
    - 30 – is returned if user did not specify “-d” option and it detects that there is an IO cache at the remote site and it is not accessible because the site is down. In this case varyonvg can not allow activation because updates in VG will cause the data divergence. Hence user intervention is required to take the decision whether to allow it or not (with -d flag).
    - 31 - varyonvg will return this exit code if user did not specify “-k” option and it detects that data divergence has occurred. In the data divergence case if “-k” option is not specified then varyonvg does not know which data copy to preserve hence user intervention is required to take this decision (-k loc | rem).
- Changes to varyoffvg
  - The aim is for varyoffvg is to de-activate the VG cleanly with all data consistently mirrored to the remote site.
    - Drain all outstanding I/Os in the aio\_cache LV. This may have an impact on the performance of the varyoffvg command.
    - A persistent message if the aio\_cache is not emptied.
  - Administration options for asynchronous GLVM
    - Not recommended during times of data divergence
      - May force a full resync of all data



## Async mode GLVM (cont)

- Other LVM Changes related to asynchronous GLVM
  - `chlv`
    - Cannot change `aio_cache` LV type if in use by a asynchronous mirror pool
    - Once asynchronous mirroring configured, you cannot change bad block reallocation for LV
  - `chpv`
    - Initially `chpv -v r` (marks a pv unavailable) not supported for async VG
    - `varyonvg` may mark PVs as unavailable – in this case `chpv -v a` will not work, `varyonvg` must do the recovery
  - `chvg` options with asynchronous volume groups has been limited, you cannot
    - Turn off mirror pool strictness
    - Turn on auto varyon
    - Turn on bad block relocation
  - `mk1v` options
    - Paging type LV not supported in GMVG
    - `aio_cache` LV type only supported in Scalable VG



## Async mode GLVM (cont)

- New command – change mirror pool (chmp)
  - Change a pool mode and set it's properties:

```
chmp -A [-c aio_cache_lv] [-h HW_Mark] -m MP_Name VG_Name
chmp -h HW_Mark -m MP_Name VG_Name
chmp -S [-f] -m MP_Name VG_Name
```

    - A Configures a mirror pool for asynchronous mirroring.
    - c aio\_cache type logical volume name
    - f Forces a mirror pool from async mirroring to sync mirroring.
    - h The high water mark as a percent of I/O cache size (default 100%, range 1-100).
    - m Specifies the mirror pool name.
    - S Changes a mirror pool from async mirroring to sync mirroring.
  - For example convert async mirroring to sync
    - Volume group must be active and mirror pool must (obviously) async and all disks active. There is a -f flag to force.

```
# chmp -S -m MyPool datavg
```



## Managing mirror pools

- New command – list mirror pools (lsmp)

```
lsmp [-L] [-A] [-n descriptorpv] [-m mirrorpoolname] vgname:
```

- A Displays information specific to asynchronous mirroring, if it exists.
- L Specifies no waiting to obtain a lock on the volume group. Note: If the volume group is being changed, using the -L flag gives unreliable data.
- m mirrorpoolname Displays mirror pool information only for the mirror pool that is specified by the mirrorpoolname parameter. If you do not specify the flag, information for all mirror pools that belong to the volume group is displayed.
- n descriptorpv Accesses information from the descriptor area that is specified by the descriptorpv variable. The volume group does not need to be active when you use the -n flag.

## Managing mirror pools (cont)

- **lsmg example**

- To display all mirror pool information for a volume group, enter the following command:

```
# lsmg -A datavg
  VOLUME GROUP:  datavg      Mirror Pool Super Strict: yes
  MIRROR POOL:   SydPool    Mirroring Mode:           SYNC
  MIRROR POOL:   MelPool    Mirroring Mode:           SYNC
```

- To display all mirror pool information for a volume group and include asynchronous mirroring information in the output, enter the following command:

```
# lsmg -A datavg
  VOLUME GROUP:  datavg      Mirror Pool Super Strict:  yes

  MIRROR POOL:   SydPool    Mirroring Mode:           ASYNC
  ASYNC MIRROR STATE:inactive ASYNC CACHE LV:          datacache_lv1
  ASYNC CACHE VALID:  yes    ASYNC CACHE EMPTY:      yes
  ASYNC CACHE HWM:   100    ASYNC DATA DIVERGED:   no

  MIRROR POOL:   MelPool    Mirroring Mode:           ASYNC
  ASYNC MIRROR STATE: active  ASYNC CACHE LV:          datacache_lv2
  ASYNC CACHE VALID:  yes    ASYNC CACHE EMPTY:      no
  ASYNC CACHE HWM:   100    ASYNC DATA DIVERGED:   no
```



## Managing mirror pools (cont)

- lsvg changes
  - To display all mirror pool information for a logical volume.

```
# lsvg -m glvm_vg
Logical Volume      Copy 1          Copy 2          Copy 3
glvmlv01            glvm1          glvm2          None
glvmlv02            glvm1          glvm2          None
glvm2_cache        glvm1          None           None
glvm1_cache        glvm2          None           None
```

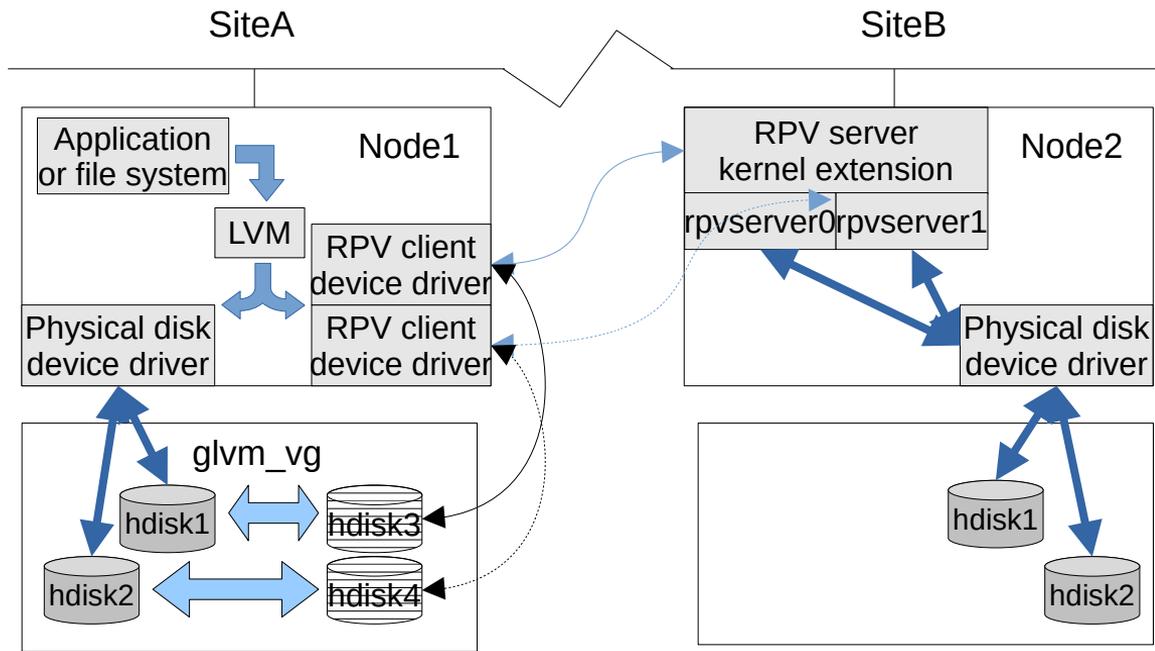


# GLVM in action



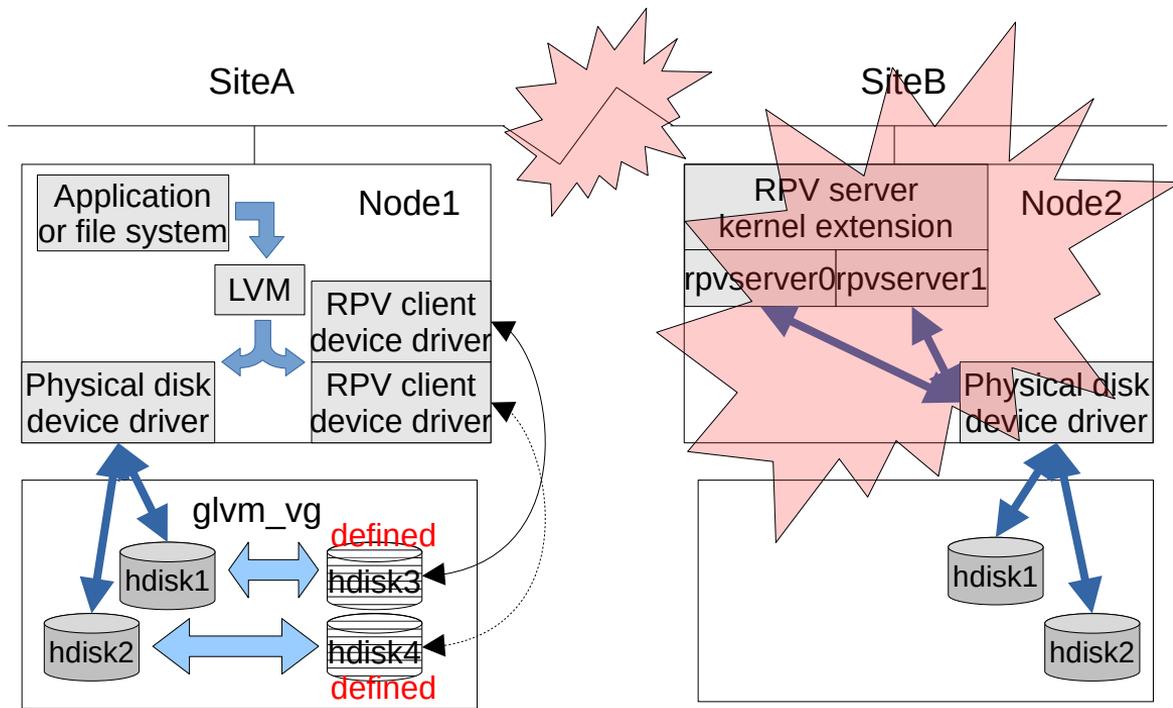
# Operation with a simple example

- Example configuration, 2 sites, GMVG with 2 hdisks / site.



# Operation with a simple example (cont)

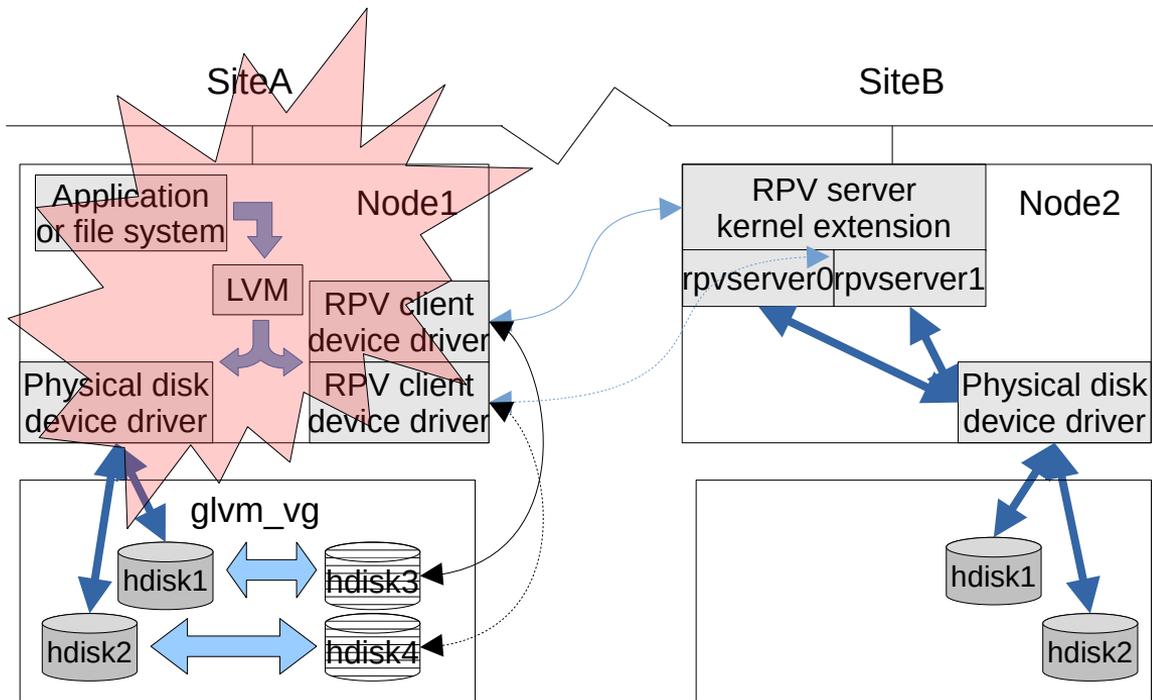
- Situation where either network(s) or remote node fails.





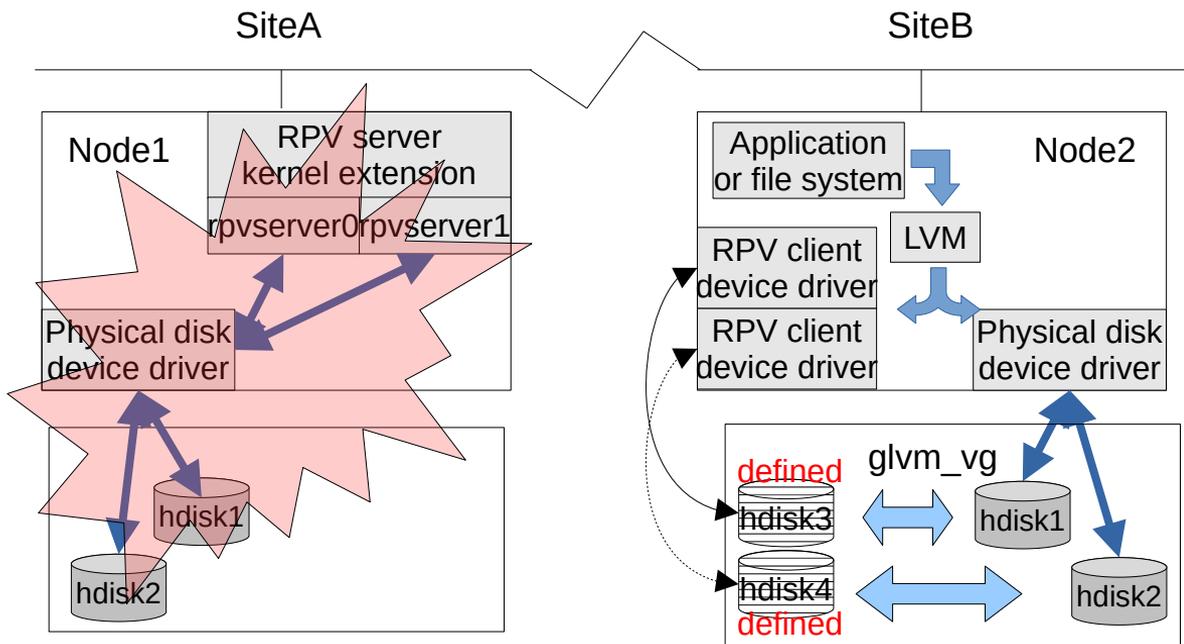
# Operation active site failing

- Previous configuration with active node failing.



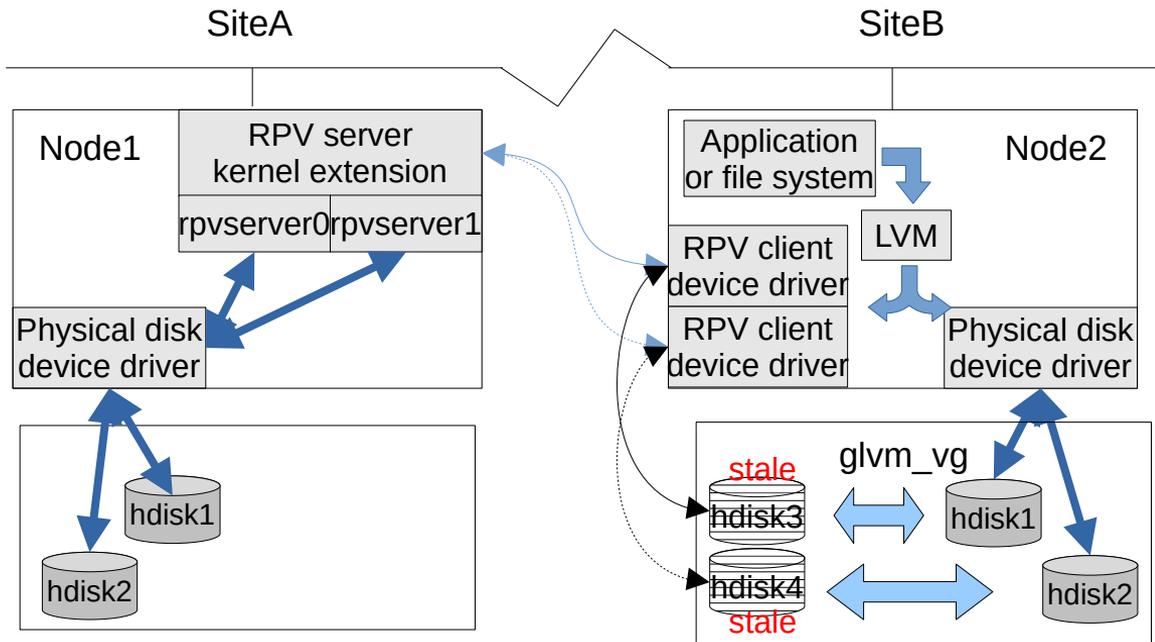
# Operation active site failing (cont)

- Application / file systems active on Site B.



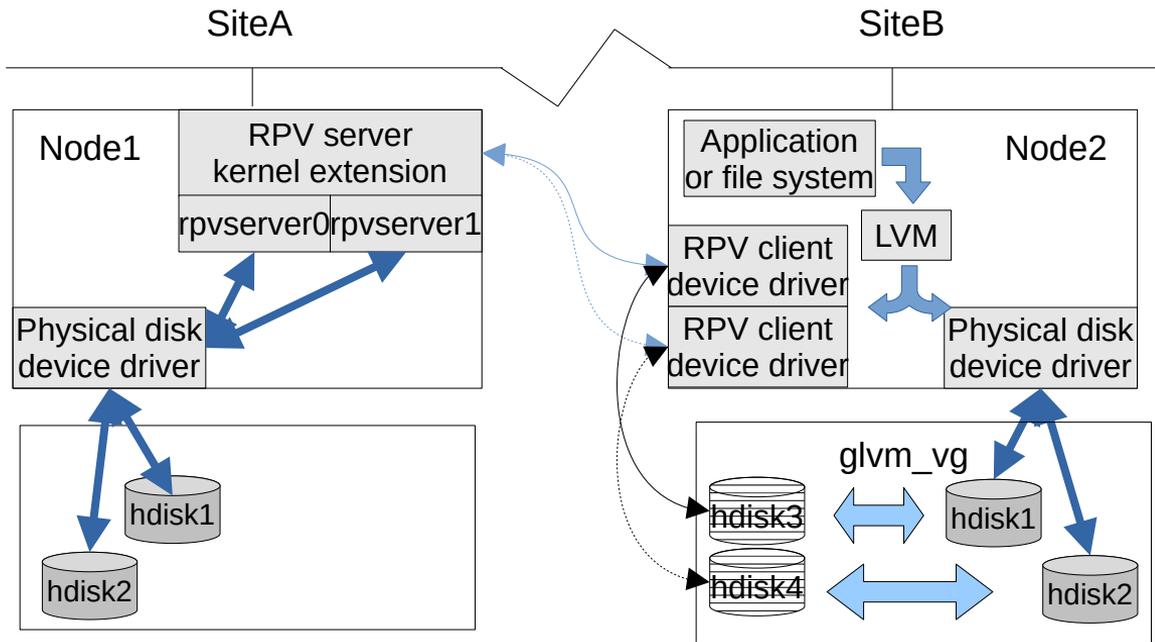
# Operation active site failing (cont)

- Site A recovers and replication started.



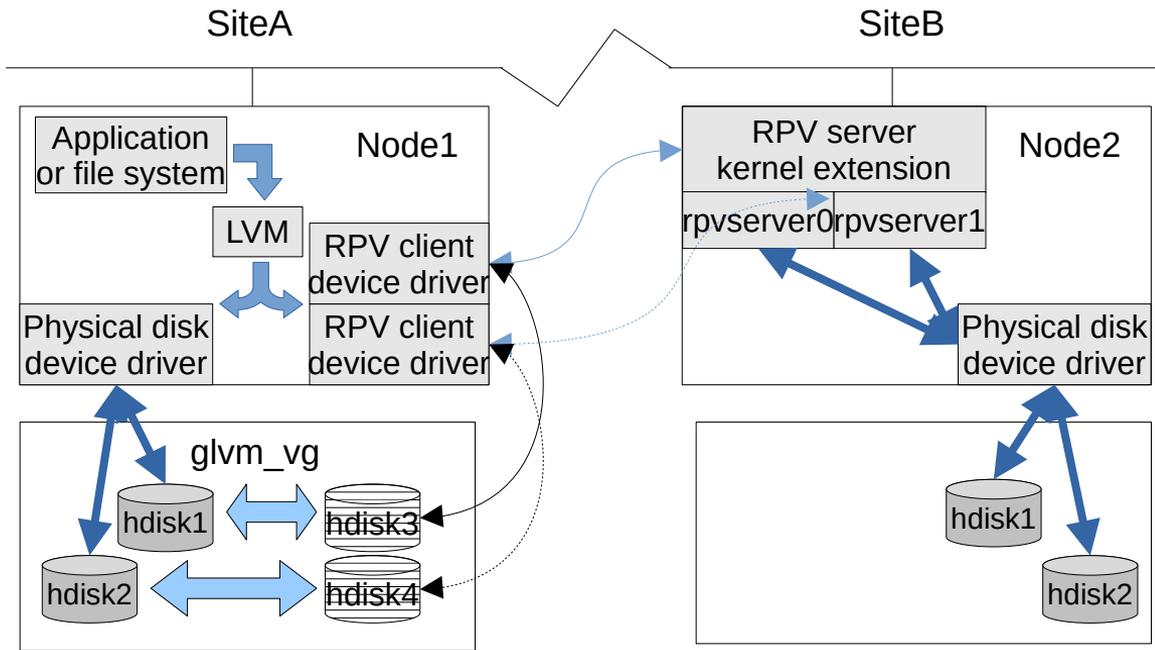
# Operation active site failing (cont)

- Sites in sync



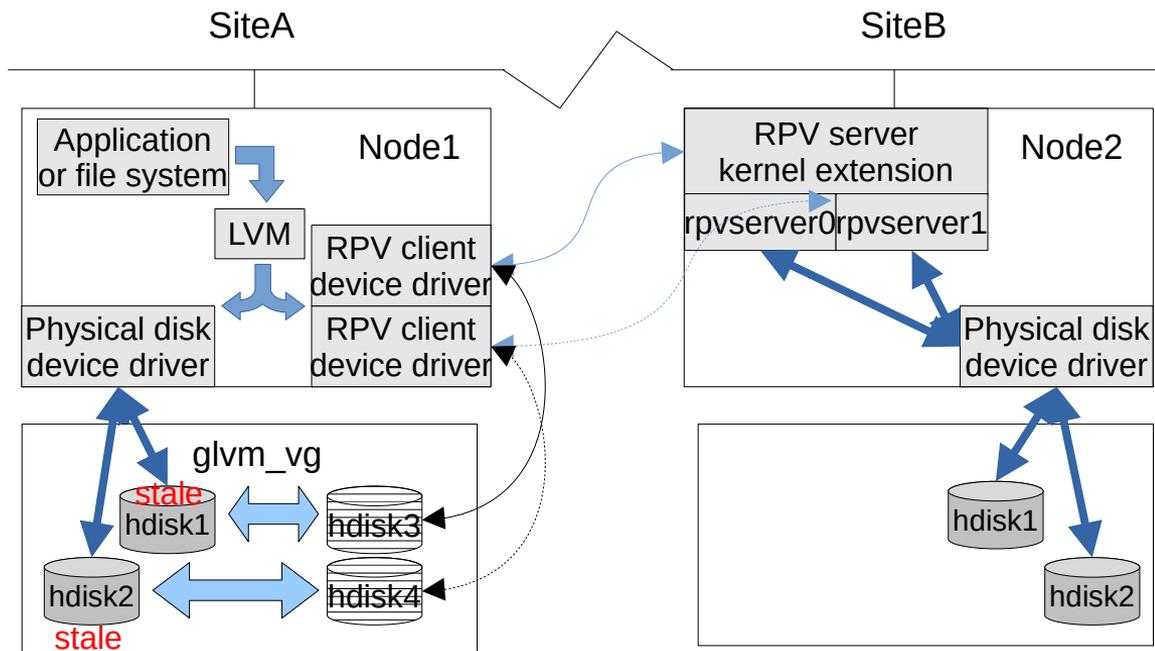
# Operation with a simple example

- Application moved back to Site A after synchronisation complete.



## Operation with a simple example

- Application moved back to Site A before synchronisation completed.
- Note: there will be a performance impact:
  - Synchronisation still running
  - Access a stale partition for I/O must be completed at the remote site.
- Note: the implication if there is more than one copy at one site, as GLVM will send two I/O operations per partition (discuss later)



# Planning



# Planning

- Know your environment

Avoid concurrent access!

- Application
  - I/O load (both peak and average) and I/O profile
  - if an existing application, use tools: gmdsizing\*; lvmstat; iostat; nmon
- Network
  - Bandwidth; latency; redundancy
  - tuning (no options)  
(There will be growth!)
- Storage
- Performance

```
# /usr/sbin/gmd/gmdsizing -i 1 -t 24h -V -v datavg
=====
Disk      Reads      Writes
hdisk4    0           272
. . . . .
=====
Disk      Reads      Writes
hdisk4    0           304
=====
=====
Disk      block      total      minimum      maximum
hdisk4    size      read  write  read  write  read  write
. . . . .
```

```
# iostat hdisk0 hdisk2 5 17280
tty:  tin      tout      avg-cpu:  % user % sys % idle % iowait physc % entc
      6.0      660.0
                    1.6  10.0  88.4  0.0  0.1  13.6

Disks:      % tm_act      Kbps      tps      Kb_read  Kb_wrtn
hdisk2      20.0      1792.0      14.0      0        1792
hdisk0      0.0        0.0        0.0        0         0
. . . . .
```

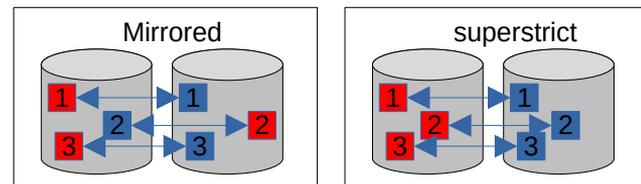
\* a old tool that is part of PowerHA samples



- The following configuration options should be considered

- Network configuration
- LVM Issues

- Physical partition size
- Number of copies
- LV allocation policy – superstrict and mirror pools to ensure copies on each site is consistent
- Mirror write consistency – recommend set to passive
  - Active – penalty on write; Off – syncvg run if not clean shutdown; passive no write penalty, passive sync
- bad block relocation should be off
- LV scheduling policies – Parallel
- Preferred read discussed later
- Write verification (off default)
- Quorum planning
  - Quorum “on” is good for data availability ( $\frac{1}{2} + 1$  VGDA's to activate and keep available)
  - Quorum “off” is good for data integrity (all disks available to activate, 1 to keep available)
  - Quorum off is recommended (recall that PowerHA SystemMirror EE does checking prior to forced activation).



# LVM semi aware of nature of RPV Client

- LVM not aware of the geographic nature of the physical volume
  - Sees each geographic physical volume as:
    - Slow (network latency)
    - Less reliable (network loss)
  - Does not coalesce synchronisation across the network
- LVM supports preferred read allows you to set pool. The PreferredRead variable can be set to a value ranging from 0 to 3 (0 disables; 1-3 is the pool number). If using GLVM stand-alone, this must be configured when you start at each site (`chlv -R # lv_name`).

```
# lslv glvmlv01
LOGICAL VOLUME:      glvmlv01          VOLUME GROUP:      glvm_vg
LV IDENTIFIER:      00c8d23000004b0000000017a5c413a99.1 PERMISSION:        read/write
VG STATE:           active/complete    LV STATE:           opened/syncd
TYPE:               jfs2log            WRITE VERIFY:       off
MAX LPs:            512                 PP SIZE:            16 megabyte(s)
COPIES:             2                   SCHED POLICY:       parallel
. . . . .
INFINITE RETRY:     no                  PREFERRED READ:    1
DEVICESUBTYPE:     DS_LVZ
COPY 1 MIRROR POOL: glvm1
COPY 2 MIRROR POOL: glvm2
COPY 3 MIRROR POOL: None
ENCRYPTION:         no
```

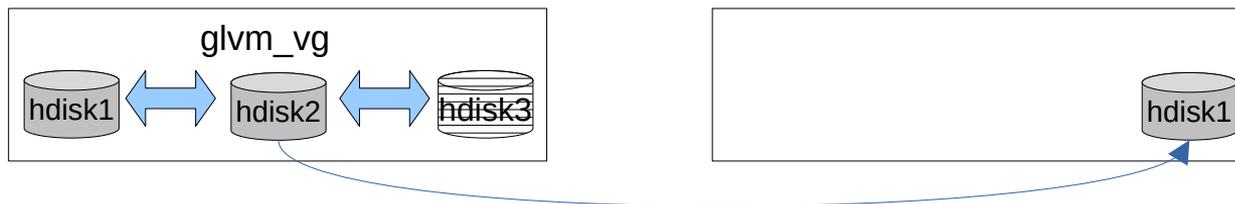
```
chlv -R 1 glvmlv01
```

PowerHA  
will manage

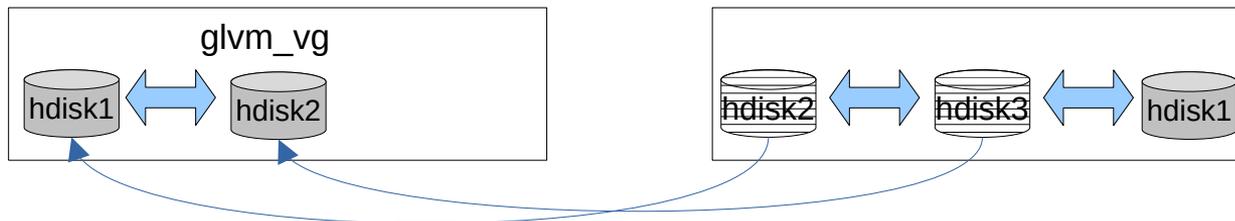


## Network does not coalesce updates

- From previous example with 2 copies at one site
  - From primary site, there will be one network transfer for each I/O



- But when site changed, each I/O results in twice the network traffic



- This also has implications when to move site after a failure in asynchronous mode, when network bandwidth may be crucial to the speed of recovery.

# LVM semi aware of nature of RPV Client

- PowerHA adds ability to set preferred read by volume group and points to new setting for Physical Volumes - Storage location (default, flashstorage, pool1 etc):
  - roundrobin – this is the default LVM preferred read policy. The LVM determines which copy to read.
  - favourcopy – manual states that this option if you want to read from the Flash storage irrespective of the resource group location
  - siteaffinity - Select this option if you want to read from the local storage path that is based on the resource group location.

### Change/Show Characteristics of a Physical Volume

Type or select values in entry fields.  
Press Enter AFTER making all desired changes.

```

[TOP]                [Entry Fields]
* Physical Volume Name      hdisk0
* Volume Group Name         datavg
* Node List                 pcha1,pcha2
  Resource Group            myapp
* Reference node           pcha1
. . . .
Current Mirror Pool        sitea
Set Mirror Pool            []
Change Mirror Pool Name    []
Remove from Mirror Pool
Storage location           default
  
```

### Change/Show characteristics of a Volume Group

Type or select values in entry fields.  
Press Enter AFTER making all desired changes.

```

[Entry Fields]
* VOLUME GROUP name      datavg
  Resource Group Name    myapp
  Node Names             pcha1,pcha2
* Activate volume group AUTOMATICALLY  no
. . . .
LVM Preferred Read      roundrobin
  
```

default  
flashstorage  
pool1 name  
pool2 name

roundrobin  
favorcopy  
siteaffinity



- Replace aio\_cache LV
  - If cache LV needs to be replaced, perform the following steps:
    - Change the mirror pool from async to sync
      - `chmp -S -m MelPool datavg`
    - Remove the current aio\_cache LV in mirror pool SydPool
      - `rmlv datacache_lv1`
    - Create a new aio\_cache LV in mirror pool SydPool
      - `mklv -t aio_cache -p copy1=SydPool -y n_datacache_lv1 datavg 1`
    - Set up async mirroring for mirror pool MelPool using the new aio\_cache LV
      - `chmp -A -c n_datacache_lv1 -h 90 -m MelPool datavg`
- Setting preferred read
  - use `chlv -R N LV_name`, where N is number of the local mirror pool.

# Maintenance tasks (cont)

- Change modes

- `smit glvm_utils` → Geographically Mirrored Volume Groups
  - Manage Geographically Mirrored Volume Groups with Mirror Pools
  - Configure Mirroring Properties of a Mirror Pool

Configure Mirroring Properties of a Mirror Pool

Move cursor to desired item and press Enter.

List Asynchronous Mirroring Information for Mirror Pools  
▲ Convert to Asynchronous Mirroring for a Mirror Pool  
Change Asynchronous Mirroring Attributes for a Mirror Pool  
Convert to Synchronous Mirroring for a Mirror Pool

- Command line

- To async: `chmp -A [ -c aiocachelv ] [ -h highwatermark ] -m mirrorpool vg`
- To sync: `chmp -S [-f ] -m mirrorpool vg`



- RPV tuning
  - rpvutil command to set:
    - rpv\_net\_monitor to set if the RPV client performs network failure detection and attempts to resume after recovery. 1=enabled, 0=disabled (default).
    - compression to control if data is compressed before being sent to server (using POWER NX842 acceleration unit if present). 1=enabled, 0=disabled (default).
    - io\_grp\_latency to set the maximum expected delay in milliseconds, before receiving the IO acknowledgement for a mirror pool that is configured in asynchronous mode. By default GLVM waits 10ms before forming group, but can be reduced to improve response time, but may increase CPU usage.



# Monitoring and statistics



# Display configuration

- Command `lsglvm` to display configuration and check the configuration

```
# lsglvm
#Volume Group   Logical Volume   RPV         PVID         Site
glvm_vg         glvm1_cache     hdisk2     00c8cf4057f2d781   glvm2
glvm_vg         glvmlv01        hdisk2     00c8cf4057f2d781   glvm2
glvm_vg         glvmlv02        hdisk2     00c8cf4057f2d781   glvm2
```

```
# lsglvm -p
glvm_vg: (Asynchronously mirrored)
# Logical Volume  RPV         PVID         Site         Mirror Pool
glvm1_cache     hdisk2     00c8cf4057f2d781   glvm2       glvm2
glvmlv01        hdisk2     00c8cf4057f2d781   glvm2       glvm2
glvmlv02        hdisk2     00c8cf4057f2d781   glvm2       glvm2
```

```
# lsglvm -m
# Table of All Physical Volumes in all Geographic Logical Volumes
# Site         Copy Physical Volumes
glvm_vg
glvmlv01
glvm1         PV1 hdisk1
glvm2         PV2 hdisk2
glvmlv02
glvm1         PV1 hdisk1
glvm2         PV2 hdisk2
glvm2_cache
glvm1         PV1 hdisk1
glvm1_cache
glvm2         PV1 hdisk2
. . . .
```

```
lsglvm -c
Checking Volume Group glvm_vg
# Site         Copy Physical Volumes
#glvm1         PV1 hdisk1
glvm2         PV2 hdisk2
Checking Logical Volume glvmlv01
Checking Logical Volume glvmlv02
Checking Logical Volume glvm2_cache
Checking Logical Volume glvml_cache
```



# Synchronous statistics

- Continuous display of statistics
  - -c for count
  - -i for interval
  - -d for just delta

```
# rpvstat -c3 -i3 -d
```

```
Remote Physical Volume Statistics:
```

RPV Client	Comp Reads cx Pend Reads	Comp Writes Pend Writes	Comp KRead Pend KRead	Comp KWrite Pend KWrite	Errors
hdisk2	1	0 121107	0 12	0 54058148	0 5124

```
Remote Physical Volume Statistics:
```

RPV Client	Comp Reads cx Pend Reads	Comp Writes Pend Writes	Comp KRead Pend KRead	Comp KWrite Pend KWrite	Errors
hdisk2	1	+0 +320	+0 19	+149136 9220	+0

```
. . . . .
```



# Synchronous statistics (cont)

- Display statistics by network
  - Disk and network statistics

```
# rpvstat -n
```

```
Remote Physical Volume Statistics:
```

RPV Client	cx	Comp Reads Pend Reads	Comp Writes Pend Writes	Comp KRead Pend KRead	Comp KWrite Pend KWrite	Errors
hdisk2	1	45 0	13039 0	780 0	1655205 0	0
192.168.200.78	Y	45 0	13039 0	780 0	1655205 0	0

- Just network

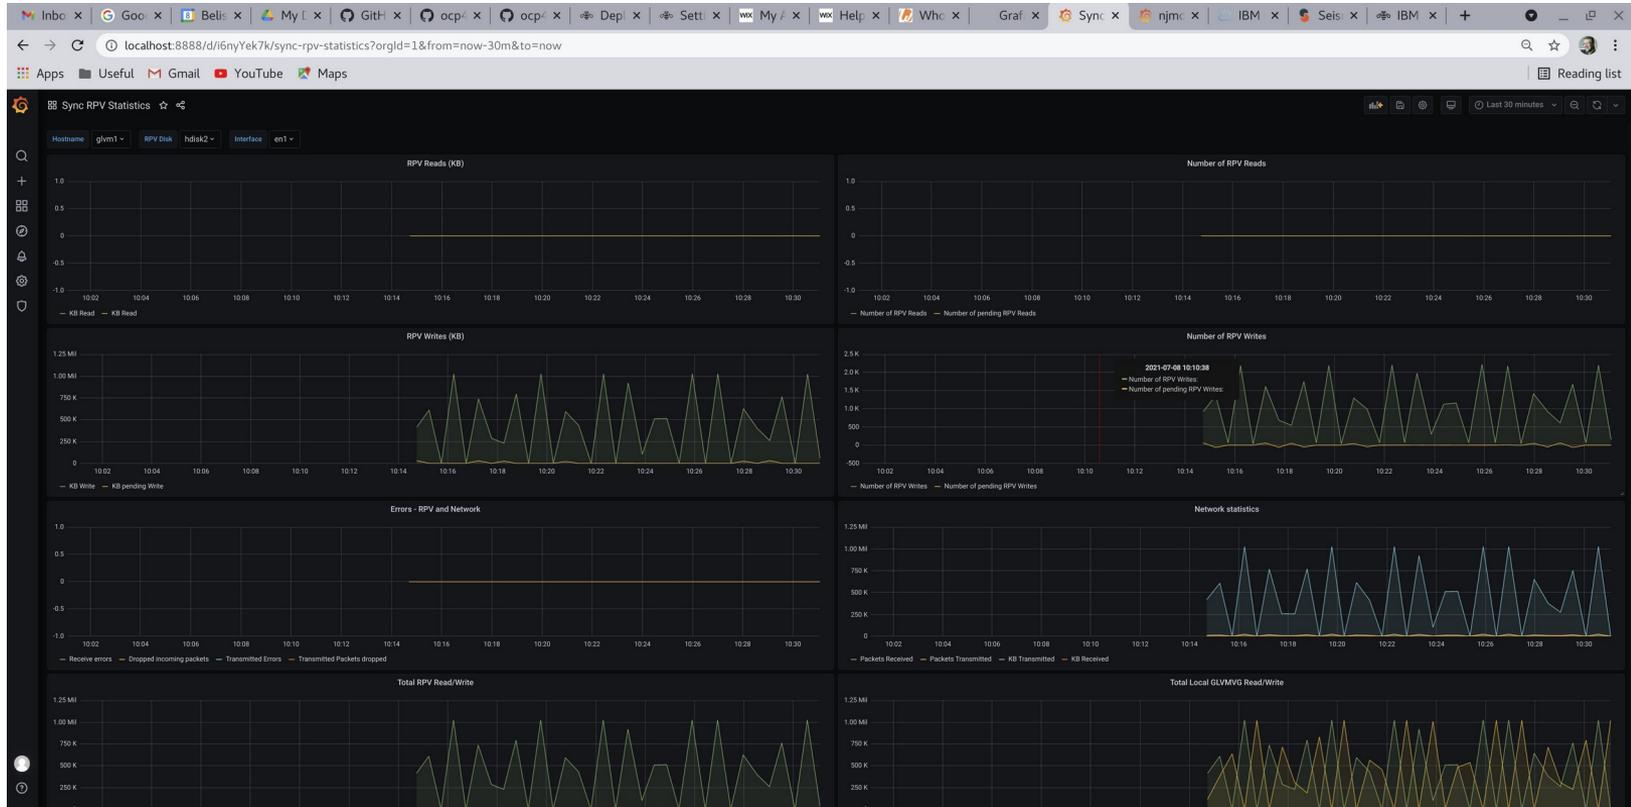
```
# rpvstat -N
```

```
Remote Physical Volume Statistics:
```

RPV Client Network	Comp Reads Pend Reads	Comp Writes Pend Writes	Comp KRead Pend KRead	Comp KWrite Pend KWrite	Errors KB/sec
192.168.200.78	45 0	13040 0	780 0	1655209 0	0 -

# Synchronous statistics (cont)

- Load statistics into influxDB and present with Grafana (all on AIX now)



# Asynchronous statistics

- Display statistics
  - Asynchronous statistics

```
# rpvstat -A

Remote Physical Volume Statistics:

      Completd  Completed  Cached  Cached  Pending  Pending
      Async    Async      Async  Async    Async    Async
RPV Client  ax Writes  KB Writes  Writes  KB Writes  Writes  KB Writes
-----
hdisk2      A      178      70664    55      27652    4      2048
```

- Cache statistics

```
# rpvstat -C

Remote Physical Volume Statistics:

      Total Async  Max  Pending  Total  Max  Cache Free
      ax Writes    Cache Cache  Cache  Cache  Cache Free
      ax Writes    Util % Writes  Wait % Wait  Space KB
-----
glvm_vg      A      1587  99.99  39  14.43  6  13058
```



# Asynchronous statistics (cont)

- Cache details

```
# rpvstat -G

Remote Physical Volume Statistics:

GMVG name ..... glvm_vg
AIO total commit time (ms) ..... 183576
Number of committed groups ..... 546
Total committed AIO data (KB) ..... 2041105
Average group commit time (ms) ..... 336
AIO data committed per sec (KB) ..... 11000
AIO total complete time (ms) ..... 305749
Number of completed groups ..... 537
Total completed AIO data (KB) ..... 2008071
Average group complete time (ms) ..... 569
AIO data completed per sec (KB) ..... 6000
Number of groups read ..... 107
Total AIO data read (KB) ..... 9573
Total AIO cache read time (ms) ..... 2845478
Average group read time (ms) ..... 26593
AIO data read per sec (KB) ..... 0
Number of groups formed ..... 547
Total group formation time (ms) ..... 5174
Average group formation time (ms) ..... 9
Number of cache fulls detected ..... 84
Total cache usage time (ms) ..... 989930
Total wait time for cache availability (ms) .. 18890
Total AIO write data in transit (KB) ..... 0
```

# Asynchronous statistics (cont)

- Load statistics into influxDB and present with Grafana (all on AIX now)



# GMVG statistics

- GMVG statistics
  - GMVG details

```
# gmvstat
GMVG Name          PVs  RPVs  Tot Vols  St Vols  Total PPs  Stale PPs  Sync
-----
glvm_vg            1    1     2         0       2550       0         100%
```

- With RPV statistics

```
gmvstat -t -r
Geographically Mirrored Volume Group Information          01:23:06 AM 13 Aug 2021
-----
glvm1
glvm1
GMVG Name          PVs  RPVs  Tot Vols  St Vols  Total PPs  Stale PPs  Sync
-----
glvm_vg            1    1     2         0       2550       0         100%

Remote Physical Volume Statistics:

RPV Client          cx  Comp Reads  Comp Writes  Comp KRead  Comp KWrite  Errors
-----
hdisk2              1      48        21987        781        5716693      0
                   0          0          0          0          0
```



# Configuration



- Configuring GLVM between two sites.
  - Configure sites on both nodes
    - `smit glvm_utils` → Remote Physical Volume Servers
      - Remote Physical Volume Server Site Name Configuration
    - `/usr/sbin/rpvsitename -a sitename`
  - Configure RPV Server(s) (configure no to start automatically, but yes to start now) on node2
    - `smit glvm_utils` -> Remote Physical Volume Servers -> Add Remote Physical Volume Servers
      - > Select the local physical volume (from name and pvid listed)
    - `/usr/sbin/mkdev -c rpvserver -s rpvserver -t rpvstype -a rpv_pvid=00c8d230573f3c0f`
      - a client\_addr='192.168.200.78' -a auto\_online='n'
  - Configure RPV Client(s) (using Server IP address, the hdisk there and timeout) on node1
    - `smit glvm_utils` -> Remote Physical Volume Clients -> Select if your mirroring network uses IPv6
      - > Add the RPV Server IP address
    - `/usr/sbin/mkdev -c disk -s remote_disk -t rpvclient -a pvid 00c8cf40 573f7392`
      - a server\_addr='192.168.200.78' -a local\_addr='192.168.200.138' -a io\_timeout='180'
  - Create the GMVG on node1
    - Select the local hdisk(s) and the RPV Client(s) they will be mirrored to. Select superstrict.



## Configuration (cont)

- Configuring GLVM between two sites (cont).
  - Configure disks
    - Turn of bad block relocation
      - `chvg -b n glvm_vg`
    - Add disks at each site to the local mirror pool
      - `chpv -p glvm1 hdisk1; chpv -p glvm2 hdisk2`
  - Create logical volumes
    - Configure passive mirror write consistency, set superstrict allocation policy, and mirror pool for each copy.
    - If using asynchronous mode, configure `aio_cache` logical volume
  - Create file system(s) as required
  - Configure the mirror pool mode (synchronous or asynchronous) with `chmp`
- Stop the current configuration and create the opposite pair of RPV Server(s) and Client(s)
  - unmount the file system(s), `varyoffvg` the volume group(s)
  - stop the `rpvclient(s)` on node1 - `rmdev -l hdiskN`
  - stop the `rpvserver(s)` on node2 – `rmdev -l rpvserverN`
  - Create the RPV Server(s) on node1 and RPV Client(s) on node2



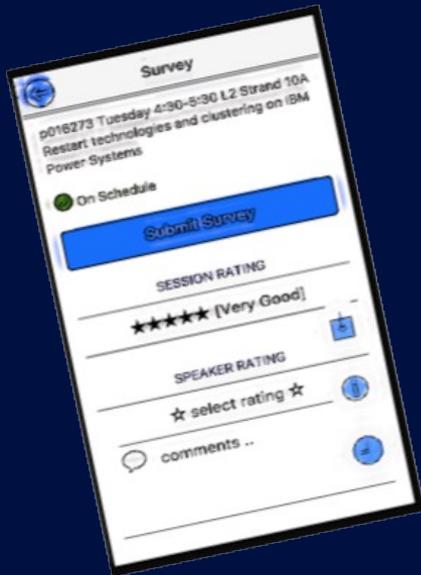
## References

- Asynchronous Geographic Logical Volume Mirroring (GLVM)
- Best Practices for Cloud deployments



# Thank you a203869 – Geographic Logical Volume Manager

Antony (Red) Steel  
antony.steel@belisama.com.sg



**Please don't forget to complete the session evaluation!**

# Notices and disclaimers

© 2021 International Business Machines Corporation. No part of this document may be reproduced or transmitted in any form without written permission from IBM.

## **U.S. Government Users Restricted Rights — use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.**

Information in these presentations (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. **This document is distributed “as is” without any warranty, either express or implied. In no event, shall IBM be liable for any damage arising from the use of this information, including but not limited to, loss of data, business interruption, loss of profit or loss of opportunity.** IBM products and services are warranted per the terms and conditions of the agreements under which they are provided.

IBM products are manufactured from new parts or new and used parts. In some cases, a product may not be new and may have been previously installed. Regardless, our warranty terms apply.”

**Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.**

Performance data contained herein was generally obtained in a controlled, isolated environments. Customer examples are presented as illustrations of how those

customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.

References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.

Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM. All materials and discussions are provided for informational purposes only, and are neither intended to, nor shall constitute legal or other guidance or advice to any individual participant or their specific situation.

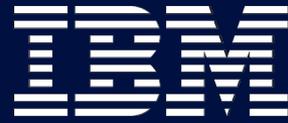
It is the customer's responsibility to insure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer's business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer follows any law.



# Notices and disclaimers

- Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products about this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. IBM does not warrant the quality of any third-party products, or the ability of any such third-party products to interoperate with IBM's products. **IBM expressly disclaims all warranties, expressed or implied, including but not limited to, the implied warranties of merchantability and fitness for a purpose.**
- The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents, copyrights, trademarks or other intellectual property right.
- IBM, the IBM logo, ibm.com and [names of other referenced IBM products and services used in the presentation] are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at: [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml)





# Backup slides



# Create environment in IBM Cloud

**Power Systems Virtual Server-p9** Active Add tags

**Virtual server instances**

Create a new instance for Power Systems Virtual Server-p9

Virtual Server instances are container software services, all running services can be created in your resource list.

Virtual server instances

SSH keys

Storage volumes

Boot images

Subnets

Number of instances: 3

Instance name: gpm

Capacity: 3

Capacity constraint:  Numerical profile  Numerical profile

VM profile: DR

SSH key:  Create SSH key

Boot image:  IBM POWER9 HV22

Operating system:  IBM AIX

Profile:  IBM P12

Core type:  Dedicated  Shared uncapped  Shared capped

Core:  3

Memory (GB):  4

Attached storage volumes (optional)

Name	Size	Type	Shareable	Number of volumes
No volumes attached				

Network interfaces

Public network:

Name	IP address	IP range	CIDR
ibm_public_network	170.148.200.2	170.148.200.0/24	170.148.200.0/24

**Summary** United States of A...

VM type: IBM POWER9 HV22

vCPUs: 3

RAM: 4K

Network interface

Storage volume: 100 GB

Total estimated cost: \$10.17/mo

IBM Cloud logo



# Site configuration

```
glvm1/#!/# lspv
hdisk0          00fa00d6b552f41b          rootvg          active
hdisk1          00c8d23057b60c26          None

Host file:
192.168.138.138          glvm1
192.168.200.138        glvm1-priv
e80::4f:10ff:fe11:9111  glvmv6-1-priv
```

```
glvm2/#!/# lspv
hdisk0          00fa00d6b552f41b          rootvg          active
hdisk1          00c8cf4057f2d781          None

Host file:
192.168.138.140        glvm2
192.168.200.78         glvm2-priv
e80::4f:10ff:fe11:9112 glvmv6-2-priv
```

## System Storage Management (Physical &amp; Logical Storage)

Move cursor to desired item and press Enter.

Logical Volume Manager

File Systems

Files & Directories

System Backup Manager

**Geographic Logical Volume Manager Utilities**

F1=Help

F2=Refresh

F3=Cancel

F8=Image

F9=Shell

F10=Exit

Enter=Do



## Geographic Logical Volume Manager Utilities

Move cursor to desired item and press Enter.

Geographically Mirrored Volume Groups  
Geographically Mirrored Logical Volumes  
Remote Physical Volume Clients  
Remote Physical Volume Servers  
Status Monitors

F1=Help  
F9=Shell

F2=Refresh  
F10=Exit

F3=Cancel  
Enter=Do

F8=Image



## Remote Physical Volume Servers

Move cursor to desired item and press Enter.

**Remote Physical Volume Server Site Name Configuration**

List All Remote Physical Volume Servers

Add Remote Physical Volume Servers

Change / Show a Remote Physical Volume Server

Change Multiple Remote Physical Volume Servers

Remove Remote Physical Volume Servers

Configure Defined Remote Physical Volume Servers

F1=Help

F2=Refresh

F3=Cancel

F8=Image

F9=Shell

F10=Exit

Enter=Do



# Create Server on glvm2 (cont)

## Remote Physical Volume Server Site Name Configuration

Move cursor to desired item and press Enter.

**Define / Change / Show Remote Physical Volume Server Site Name**

Remove Remote Physical Volume Server Site Name

F1=Help

F2=Refresh

F3=Cancel

F8=Image

F9=Shell

F10=Exit

Enter=Do



# Create Server on glvm2 (cont)

Define / Change / Show Remote Physical Volume Server Site Name

Type or select values in entry fields.  
Press Enter AFTER making all desired changes.

\* Remote Physical Volume Server Site Name [Entry Fields]  
[glvm2]

F1=Help  
F5=Reset  
F9=Shell

F2=Refresh  
F6=Command  
F10=Exit

F3=Cancel  
F7=Edit  
Enter=Do

F4=List  
F8=Image

# Create RPV server on glvm2 (cont) – add RPV server

## Remote Physical Volume Servers

Move cursor to desired item and press Enter.

Remote Physical Volume Server Site Name Configuration

List All Remote Physical Volume Servers

**Add Remote Physical Volume Servers**

Change / Show a Remote Physical Volume Server

Change Multiple Remote Physical Volume Servers

Rem+-----+

Con|

Physical Volume Identifiers

| Move cursor to desired item and press F7.

| ONE OR MORE items can be selected.

| Press Enter AFTER making all selections.

| # Physical Volume                    Physical Volume Identifier

| # -----

| **hdisk1                    00c8cf4057f2d781**

| F1=Help                                F2=Refresh                            F3=Cancel

| F7=Select                              F8=Image                              F10=Exit

| F1=He| Enter=Do                        /=Find                                n=Find Next

F9=Sh+-----+

# Create RPV server on glvm2 (cont) – add RPV server (cont)

## Add Remote Physical Volume Servers

Type or select values in entry fields.  
Press Enter AFTER making all desired changes.

	[Entry Fields]	
Physical Volume Identifiers	00c8cf4057f2d781	
* Remote Physical Volume Client Internet Address	[192.168.200.138]	+
Configure Automatically at System Restart?	[no]	+
Start New Devices Immediately?	[yes]	+

F1=Help  
F5=Reset  
F9=Shell

F2=Refresh  
F6=Command  
F10=Exit

F3=Cancel  
F7=Edit  
Enter=Do

F4=List  
F8=Image



# Create RPV server on glvm2 (cont) – add RPV server (cont)

## COMMAND STATUS

Command: OK                    stdout: yes                    stderr: no

Before command completion, additional instructions may appear below.

rpvserver0 Available

F1=Help

F2=Refresh

F3=Cancel

F6=Command

F8=Image

F9=Shell

F10=Exit

/=Find

n=Find Next

# Create RPV client on glvm1

## Add Remote Physical Volume Clients

Type or select a value for the entry field.  
Press Enter AFTER making all desired changes.

[Entry Fields]

\* Does data mirroring network use Internet Protocol Version 6 (IPv6)?    no    +

F1=Help

F2=Refresh

F3=Cancel

F4=List

F5=Reset

F6=Command

F7=Edit

F8=Image

F9=Shell

F10=Exit

Enter=Do

# Create RPV client on glvm1 (cont)

## Add Remote Physical Volume Clients

Type or select a value for the entry field.  
Press Enter AFTER making all desired changes.

\* Remote Physical Volume Server Internet Address      [Entry Fields]  
   [192.168.200.78]      +

F1=Help

F2=Refresh

F3=Cancel

F4=List

F5=Reset

F6=Command

F7=Edit

F8=Image

F9=Shell

F10=Exit

Enter=Do

## Create RPV client on glvm1 (cont)

## Add Remote Physical Volume Clients

Type or select a value for the entry field.  
Press Enter AFTER making all desired changes.

```

                                     [Entry Fields]
* Remote Physical Volume Server Internet Address   [192.168.200.78]   +
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Remote Physical Volume Local Internet Address                                     |
|                                     |                                     |                                     |
| Move cursor to desired item and press Enter. |                                     |                                     |
|                                     |                                     |                                     |
|   192.168.138.138   glvm1           |                                     |                                     |
|   192.168.200.138  glvm1-priv      |                                     |                                     |
|                                     |                                     |                                     |
| F1=Help           F2=Refresh         F3=Cancel   |                                     |                                     |
F1| F8=Image        F10=Exit           Enter=Do    |                                     |                                     |
F5| /=Find          n=Find Next        |                                     |                                     |
F9+-----+-----+-----+-----+-----+-----+

```



## Create RPV client on glvm1 (cont)

## Add Remote Physical Volume Clients

```

Ty+-----+
Pr|           Remote Physical Volume Server Disks
|
| Move cursor to desired item and press F7.
* |   ONE OR MORE items can be selected.
| | Press Enter AFTER making all selections.
|
| # These remote physical volumes are available
| # at site glvm2:
| #
| # (The physical volume names are as they are known
| # on the host named
| # glvm-2)
| #
| # Physical Volume           Physical Volume Identifier
| # -----
| # hdisk1                    00c8cf4057f2d7810000000000000000
|
| F1=Help           F2=Refresh           F3=Cancel
F1| F7=Select       F8=Image             F10=Exit
F5| Enter=Do        /=Find               n=Find Next
F9+-----+

```



## Create RPV client on glvm1 (cont)

## Add Remote Physical Volume Clients

Type or select values in entry fields.  
Press Enter AFTER making all desired changes.

	[Entry Fields]	
Remote Physical Volume Server Internet Address	192.168.200.78	
Remote Physical Volume Local Internet Address	192.168.200.138	
Physical Volume Identifiers	00c8cf4057f2d78100000>	
I/O Timeout Interval (Seconds)	[10]	#
Start New Devices Immediately?	[yes]	+

F1=Help

F2=Refresh

F3=Cancel

F4=List

F5=Reset

F6=Command

F7=Edit

F8=Image

F9=Shell

F10=Exit

Enter=Do



# Create RPV client on glvm1 (cont)

## COMMAND STATUS

Command: OK                    stdout: yes                    stderr: no

Before command completion, additional instructions may appear below.

hdisk2 Available

# lspv			
hdisk0	00fa00d6b552f41b	rootvg	active
hdisk1	00c8d23057b60c26	None	
hdisk2	00c8cf4057f2d781	None	

F1=Help                    F2=Refresh                    F3=Cancel                    F6=Command  
 F8=Image                    F9=Shell                    F10=Exit                    /=Find  
 n=Find Next



# Create a scalable volume group

## Add a Scalable Volume Group

Type or select values in entry fields.  
 Press Enter AFTER making all desired changes.

```

                                [Entry Fields]
VOLUME GROUP name                [glvm_vg]
Physical partition SIZE in megabytes
* PHYSICAL VOLUME names           [hdisk1 hdisk2]
Force the creation of a volume group?    no
Activate volume group AUTOMATICALLY    no
    at system restart?
Volume Group MAJOR NUMBER           []
Create VG Concurrent Capable?        no
Max PPs per VG in units of 1024      32
Max Logical Volumes                  256
Enable Strict Mirror Pools           Superstrict
Infinite Retry Option                no
    
```

```

F1=Help          F2=Refresh      F3=Cancel      F4=List
F5=Reset         F6=Command      F7=Edit        F8=Image
F9=Shell         F10=Exit        Enter=Do
    
```



## Turn off bad block relocation

```
# chvg -b n glvm_vg
```

stop bad block relocation

```
# lsvg glvm_vg
```

VOLUME GROUP:	glvm_vg	VG IDENTIFIER:	00c8d23000004b000000017a5c413a99
VG STATE:	active	PP SIZE:	16 megabyte(s)
VG PERMISSION:	read/write	TOTAL PPs:	1275 (20400 megabytes)
MAX LVs:	256	FREE PPs:	1275 (20400 megabytes)
LVs:	0	USED PPs:	0 (0 megabytes)
OPEN LVs:	0	QUORUM:	2 (Enabled)
TOTAL PVs:	1	VG DESCRIPTORS:	2
STALE PVs:	0	STALE PPs:	0
ACTIVE PVs:	1	AUTO ON:	no
MAX PPs per VG:	32768	MAX PVs:	1024
LTG size (Dynamic):	512 kilobyte(s)	AUTO SYNC:	no
HOT SPARE:	no	<b>BB POLICY:</b>	<b>non-relocatable</b>
MIRROR POOL STRICT:	super	INFINITE RETRY:	no
PV RESTRICTION:	none	CRITICAL VG:	no
DISK BLOCK SIZE:	512	CRITICAL PVs:	no
FS SYNC OPTION:	no		
ENCRYPTION:	no		



# Put local disk into pool glvm1

```
# chpv -p glvm1 hdisk1
```

put local disk in mirror pool glvm1

```
lspv hdisk1
```

```
PHYSICAL VOLUME:      hdisk1                VOLUME GROUP:      glvm_vg
PV IDENTIFIER:        00c8d23057b60c26  VG IDENTIFIER      00c8d23000004b000000017a5c413a99
PV STATE:             active
STALE PARTITIONS:    0                ALLOCATABLE:       yes
PP SIZE:              16 megabyte(s)          LOGICAL VOLUMES:   0
TOTAL PPs:           1275 (20400 megabytes)    VG DESCRIPTORS:    2
FREE PPs:            1275 (20400 megabytes)    HOT SPARE:         no
USED PPs:            0 (0 megabytes)           MAX REQUEST:       512 kilobytes
FREE DISTRIBUTION:   255..255..255..255..255
USED DISTRIBUTION:   00..00..00..00..00
MIRROR POOL:       glvm1
```



## Put remote disk into pool glvm2

```
# chpv -p glvm2 hdisk2
```

```
put local disk in mirror pool glvm2
```

```
lspv hdisk2
```

```
PHYSICAL VOLUME:      hdisk2                VOLUME GROUP:      glvm_vg
PV IDENTIFIER:        00c8cf4057f2d781  VG IDENTIFIER      00c8d23000004b000000017a5c413a99
PV STATE:             active
STALE PARTITIONS:    0                ALLOCATABLE:       yes
PP SIZE:              16 megabyte(s)         LOGICAL VOLUMES:   0
TOTAL PPs:           1275 (20400 megabytes)  VG DESCRIPTORS:    2
FREE PPs:            1275 (20400 megabytes)  HOT SPARE:         no
USED PPs:            0 (0 megabytes)         MAX REQUEST:       512 kilobytes
FREE DISTRIBUTION:   255..255..255..255..255
USED DISTRIBUTION:   00..00..00..00..00
MIRROR POOL:       glvm2
```

# Display pools

```
# lsmc -A glvm_vg
VOLUME GROUP:      glvm_vg          Mirror Pool Super Strict: yes

MIRROR POOL:       glvm1            Mirroring Mode:           SYNC
MIRROR POOL:       glvm2            Mirroring Mode:           SYNC
```

# Create the logical volumes

## Add a Logical Volume

Type or select values in entry fields.  
Press Enter AFTER making all desired changes.

[TOP]	[Entry Fields]	
Logical volume NAME	[glvmlv01]	
* VOLUME GROUP name	glvm_vg	
* Number of LOGICAL PARTITIONS	[1]	#
PHYSICAL VOLUME names	[hdisk1 hdisk2]	+
Logical volume TYPE	[jfs2log]	+
POSITION on physical volume	middle	+
RANGE of physical volumes	minimum	+
MAXIMUM NUMBER of PHYSICAL VOLUMES to use for allocation	[ ]	#
Number of COPIES of each logical partition	2	+
Mirror Write Consistency?	<b>passive</b>	+
Allocate each logical partition copy on a SEPARATE physical volume?	<b>superstrict</b>	+
RELOCATE the logical volume during reorganization?	yes	+
Logical volume LABEL	[glvm_jfslog]	
MAXIMUM NUMBER of LOGICAL PARTITIONS	[512]	#
Enable BAD BLOCK relocation?	no	+
SCHEDULING POLICY for writing/reading logical partition copies	parallel	+
Enable WRITE VERIFY?	no	+
File containing ALLOCATION MAP	[ ]	
Stripe Size?	[Not Striped]	+
Serialize IO?	no	+
Mirror Pool for First Copy	<b>glvm1</b>	+
Mirror Pool for Second Copy	<b>glvm2</b>	+
Mirror Pool for Third Copy		+
Infinite Retry Option	no	+



# Create the logical volumes (cont)

## Add a Logical Volume

Type or select values in entry fields.  
Press Enter AFTER making all desired changes.

[TOP]	[Entry Fields]	
Logical volume NAME	[glvmlv02]	
* VOLUME GROUP name	glvm_vg	
* Number of LOGICAL PARTITIONS	[100]	#
PHYSICAL VOLUME names	[hdisk1 hdisk2]	+
Logical volume TYPE	[jfs2]	+
POSITION on physical volume	middle	+
RANGE of physical volumes	minimum	+
MAXIMUM NUMBER of PHYSICAL VOLUMES to use for allocation	[ ]	#
Number of COPIES of each logical partition	2	+
Mirror Write Consistency?	<b>passive</b>	+
Allocate each logical partition copy on a SEPARATE physical volume?	<b>superstrict</b>	+
RELOCATE the logical volume during reorganization?	yes	+
Logical volume LABEL	[glvm_data]	
MAXIMUM NUMBER of LOGICAL PARTITIONS	[512]	#
Enable BAD BLOCK relocation?	no	+
SCHEDULING POLICY for writing/reading logical partition copies	parallel	+
Enable WRITE VERIFY?	no	+
File containing ALLOCATION MAP	[ ]	
Stripe Size?	[Not Striped]	+
Serialize IO?	no	+
Mirror Pool for First Copy	<b>glvm1</b>	+
Mirror Pool for Second Copy	<b>glvm2</b>	+
Mirror Pool for Third Copy		+
Infinite Retry Option	no	+



# Create the logical volumes (cont)

## Add a Logical Volume

Type or select values in entry fields.  
Press Enter AFTER making all desired changes.

[TOP]	[Entry Fields]	
Logical volume NAME	[glvm1_cache]	
* VOLUME GROUP name	glvm_vg	
* Number of LOGICAL PARTITIONS	[4]	#
PHYSICAL VOLUME names	[hdisk1]	+
Logical volume TYPE	[ <b>aio_cache</b> ]	+
POSITION on physical volume	middle	+
RANGE of physical volumes	minimum	+
MAXIMUM NUMBER of PHYSICAL VOLUMES to use for allocation	[ ]	#
Number of COPIES of each logical partition	2	+
Mirror Write Consistency?	<b>passive</b>	+
Allocate each logical partition copy on a SEPARATE physical volume?	yes	+
RELOCATE the logical volume during reorganization?	yes	+
Logical volume LABEL	[glvm1-cache]	
MAXIMUM NUMBER of LOGICAL PARTITIONS	[512]	#
Enable BAD BLOCK relocation?	no	+
SCHEDULING POLICY for writing/reading logical partition copies	parallel	+
Enable WRITE VERIFY?	no	+
File containing ALLOCATION MAP	[ ]	
Stripe Size?	[Not Striped]	+
Serialize IO?	no	+
Mirror Pool for First Copy	<b>glvm1</b>	+
Mirror Pool for Second Copy		+
Mirror Pool for Third Copy		+
Infinite Retry Option	no	+



# Create the logical volumes (cont)

## Add a Logical Volume

Type or select values in entry fields.  
Press Enter AFTER making all desired changes.

[TOP]	[Entry Fields]	
Logical volume NAME	[glvm2_cache]	
* VOLUME GROUP name	glvm_vg	
* Number of LOGICAL PARTITIONS	[4]	#
PHYSICAL VOLUME names	[hdisk2]	+
Logical volume TYPE	[aio_cache]	+
POSITION on physical volume	middle	+
RANGE of physical volumes	minimum	+
MAXIMUM NUMBER of PHYSICAL VOLUMES to use for allocation	[ ]	#
Number of COPIES of each logical partition	2	+
Mirror Write Consistency?	<b>passive</b>	+
Allocate each logical partition copy on a SEPARATE physical volume?	yes	+
RELOCATE the logical volume during reorganization?	yes	+
Logical volume LABEL	[glvm2-cache]	
MAXIMUM NUMBER of LOGICAL PARTITIONS	[512]	#
Enable BAD BLOCK relocation?	no	+
SCHEDULING POLICY for writing/reading logical partition copies	parallel	+
Enable WRITE VERIFY?	no	+
File containing ALLOCATION MAP	[ ]	
Stripe Size?	[Not Striped]	+
Serialize IO?	no	+
Mirror Pool for First Copy	<b>glvm2</b>	+
Mirror Pool for Second Copy		+
Mirror Pool for Third Copy		+
Infinite Retry Option	no	+



## Create the file systems

- Create a file system using the jfs2log lv and the jfs2 lv.
- Confirm that the file system mounts

```
# df
Filesystem      512-blocks      Free %Used    Iused %Iused Mounted on
/dev/hd4        196608          51776   74%      2991   33% /
/dev/hd2        4915200         46016  100%     42322  82% /usr
/dev/hd9var     393216           0  100%     1055   72% /var
/dev/hd3        458752          448896   3%        41     1% /tmp
/dev/hd1        65536           64792   2%         8     1% /home
/dev/hd11admin  262144          261384   1%         5     1% /admin
/proc           -                -     -         -     - /proc
/dev/hd10opt    786432          85088  90%     11471  52% /opt
/dev/livedump   524288          523552   1%         4     1% /var/adm/ras/livedump
/dev/repo00     15794176        80424  100%     2130   19% /usr/sys/inst.images
/dev/ulv01      7471104         3725000  51%        68     1% /home/red
/dev/ulv04      3080192         1275536  59%       121     1% /var/lib/influxdb
/ahafs          -                -     -         35     1% /aha
/dev/glvmlv02   3276800         3003120   9%        16     1% /data
```



# Stop current configuration

- Stop application(s)
- unmount the file system(s)
- Stop the RPV Client(s)

```
glvm1://:# rmdev -l hdisk2  
hdisk2 Defined
```

- Stop the RPV Server(s)

```
glvm2://:# rmdev -l rpvserver0  
rpvserver0 Defined
```

## Remote Physical Volume Server Site Name Configuration

Move cursor to desired item and press Enter.

**Define / Change / Show Remote Physical Volume Server Site Name**

Remove Remote Physical Volume Server Site Name

F1=Help

F2=Refresh

F3=Cancel

F8=Image

F9=Shell

F10=Exit

Enter=Do



# Create RPV server on glvm1 (cont)

Define / Change / Show Remote Physical Volume Server Site Name

Type or select values in entry fields.  
Press Enter AFTER making all desired changes.

\* Remote Physical Volume Server Site Name [Entry Fields]  
[glvm1]

F1=Help

F2=Refresh

F3=Cancel

F4=List

F5=Reset

F6=Command

F7=Edit

F8=Image

F9=Shell

F10=Exit

Enter=Do

# Create RPV server on glvm1 (cont) – add RPV server (cont)

## Remote Physical Volume Servers

Move cursor to desired item and press Enter.

```
Remote Physical Volume Server Site Name Configuration
List All Remote Physical Volume Servers
Add Remote Physical Volume Servers
Change / Show a Remote Physical Volume Server
Change Multiple Remote Physical Volume Servers
```

### Physical Volume Identifiers

Move cursor to desired item and press F7.

ONE OR MORE items can be selected.

Press Enter AFTER making all selections.

# Physical Volume	Physical Volume Identifier
-------------------	----------------------------

#	-----
---	-------

<b>hdisk1</b>	<b>00c8d23057b60c26</b>
---------------	-------------------------

F1=Help

F2=Refresh

F3=Cancel

F7=Select

F8=Image

F10=Exit

F1| Enter=Do

/=Find

n=Find Next

F9+



# Create RPV server on glvm1 (cont) – add RPV server (cont)

## Add Remote Physical Volume Servers

Type or select values in entry fields.

Press Enter AFTER making all desired changes.

	[Entry Fields]	
Physical Volume Identifiers	00c8d23057b60c26	
* Remote Physical Volume Client Internet Address	[192.168.200.78]	+
Configure Automatically at System Restart?	[no]	+
Start New Devices Immediately?	[yes]	+

F1=Help

F2=Refresh

F3=Cancel

F4=List

F5=Reset

F6=Command

F7=Edit

F8=Image

F9=Shell

F10=Exit

Enter=Do



# Create RPV server on glvm1 (cont) – add RPV server (cont)

## COMMAND STATUS

Command: OK                    stdout: yes                    stderr: no

Before command completion, additional instructions may appear below.

rpvserver0 Available

F1=Help

F2=Refresh

F3=Cancel

F6=Command

F8=Image

F9=Shell

F10=Exit

/=Find

n=Find Next

# Create RPV client on glvm2

## Add Remote Physical Volume Clients

Type or select a value for the entry field.  
Press Enter AFTER making all desired changes.

[Entry Fields]

\* Does data mirroring network use Internet Protocol Version 6 (IPv6)?      no      +

F1=Help

F2=Refresh

F3=Cancel

F4=List

F5=Reset

F6=Command

F7=Edit

F8=Image

F9=Shell

F10=Exit

Enter=Do



# Create RPV client on glvm2 (cont)

## Add Remote Physical Volume Clients

Type or select a value for the entry field.  
Press Enter AFTER making all desired changes.

```

                                     [Entry Fields]
* Remote Physical Volume Server Internet Address  [192.168.200.138]      +
  
```

F1=Help

F2=Refresh

F3=Cancel

F4=List

F5=Reset

F6=Command

F7=Edit

F8=Image

F9=Shell

F10=Exit

Enter=Do



# Create RPV client on glvm2 (cont)

## Add Remote Physical Volume Clients

Type or select a value for the entry field.  
Press Enter AFTER making all desired changes.

```

                                     [Entry Fields]
* Remote Physical Volume Server Internet Address   [192.168.200.138]   +

+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     |                                     |
|           Remote Physical Volume Local Internet Address           |
| Move cursor to desired item and press Enter.                     |
| 192.168.138.140  glvm2                                             |
| 192.168.200.78  glvm2-priv                                       |
|                                                                     |
| F1=Help           F2=Refresh           F3=Cancel                   |
F1=He| F8=Image     F10=Exit             Enter=Do                    |
F5=Re| /=Find       n=Find Next          |                             |
F9=Sh+-----+-----+-----+-----+-----+-----+-----+

```



# Create RPV client on glvm2 (cont)

## Add Remote Physical Volume Clients

```

Type +-----+
Press|          Remote Physical Volume Server Disks          |
|
| Move cursor to desired item and press F7.                    |
| * Rem | ONE OR MORE items can be selected.                    |
|       | Press Enter AFTER making all selections.              |
|
| # These remote physical volumes are available                |
| # at site glvm1:                                             |
| #                                                            |
| # (The physical volume names are as they are known          |
| # on the host named                                         |
| # glvm-1)                                                    |
| #                                                            |
| # Physical Volume      Physical Volume Identifier          |
| # -----|-----|
| # hdisk1              00c8d23057b60c260000000000000000
|
| F1=Help                F2=Refresh                F3=Cancel
F1=He| F7=Select          F8=Image                 F10=Exit
F5=Re| Enter=Do           /=Find                    n=Find Next
F9=Sh+-----+
    
```



# Create RPV client on glvm2 (cont)

## Add Remote Physical Volume Clients

Type or select values in entry fields.  
Press Enter AFTER making all desired changes.

```

                                     [Entry Fields]
Remote Physical Volume Server Internet Address 192.168.200.138
Remote Physical Volume Local Internet Address  192.168.200.78
Physical Volume Identifiers                   00c8d23057b60c2600000000000000>
I/O Timeout Interval (Seconds)                [10] #
Start New Devices Immediately?                 [yes] +

```

F1=Help  
F5=Reset  
F9=Shell

F2=Refresh  
F6=Command  
F10=Exit

F3=Cancel  
F7=Edit  
Enter=Do

F4=List  
F8=Image



# Create RPV client on glvm2 (cont)

## COMMAND STATUS

Command: OK                    stdout: yes                    stderr: no

Before command completion, additional instructions may appear below.

hdisk2 Available



# lspv			
hdisk0	00fa00d6b552f41b	rootvg	active
hdisk1	00c8cf4057f2d781	None	
hdisk2	00c8d23057b60c26	None	

F1=Help  
F8=Image  
n=Find Next

F2=Refresh  
F9=Shell

F3=Cancel  
F10=Exit

F6=Command  
/=Find



## Check Volume Group

- Import volume group from the local disk in the GMVG
- mount the file system(s) and confirm data

# Check the configuration

## Geographically Mirrored Volume Groups

Move cursor to desired item and press Enter.

List Geographic Volume Group Information

Manage Geographically Mirrored Volume Groups with Mirror Pools

Manage Legacy Geographically Mirrored Volume Groups

Replace a Remote Physical Volume with Another on the same Site

**Verify Mirror Copy Site Locations for a Volume Group**

F1=Help  
F9=Shell

F2=Refresh  
F10=Exit

F3=Cancel  
Enter=Do

F8=Image



# Verify Mirror Copy Site Locations for a Volume Group

## COMMAND STATUS

Command: OK                    stdout: yes                    stderr: no

Before command completion, additional instructions may appear below.

```
Checking Volume Group glvm_vg
# Site      Copy Physical Volumes
#glvm1     PV1 hdisk1
glvm2     PV2 hdisk2
Checking Logical Volume glvmlv01
Checking Logical Volume glvmlv02
Checking Logical Volume glvm1_cache
Checking Logical Volume glvm2_cache
```

F1=Help  
F8=Image  
n=Find Next

F2=Refresh  
F9=Shell

F3=Cancel  
F10=Exit

F6=Command  
/=Find



## Verify Mirror Copy Site Locations for a Volume Group (bad example)

## COMMAND STATUS

Command: OK                    stdout: yes                    stderr: no

Before command completion, additional instructions may appear below.

Checking Volume Group glvm\_vg

# Site            Copy Physical Volumes

#glvm1            PV1 hdisk1

glvm2            PV2 hdisk2

Checking Logical Volume glvmlv01

# Warning: The Logical Volume glvmlv01 allocation policy is y,

# it must be s, for SuperStrict. Use the chlv -s s -u <upperbound> glvmlv01

# command to correct this before proceeding.

Checking Logical Volume glvmlv02

# Warning: The Logical Volume glvmlv02 allocation policy is y,

# it must be s, for SuperStrict. Use the chlv -s s -u <upperbound> glvmlv02

# command to correct this before proceeding.

Checking Logical Volume glvm1\_cache

# Warning: The Logical Volume glvm1\_cache allocation policy is y,

# it must be s, for SuperStrict. Use the chlv -s s -u <upperbound> glvm1\_cache

# command to correct this before proceeding.

Checking Logical Volume glvm2\_cache

# Warning: The Logical Volume glvm2\_cache allocation policy is y,

# it must be s, for SuperStrict. Use the chlv -s s -u <upperbound> glvm2\_cache

# command to correct this before proceeding.

F1=Help

F2=Refresh

F3=Cancel

F6=Command

F8=Image

F9=Shell

F10=Exit

/=Find

n=Find Next



# Move back to glvm1

```
glvm2://:# rmdev -l hdisk2
hdisk2 Defined
glvm2://:# mkdev -l rpvserver0
rpvserver0 Available
```

```
glvm1://:# rmdev -l rpvserver0
rpvserver0 Defined
glvm1://:# mkdev -l hdisk2
hdisk2 Available
glvm1://:# lspv
hdisk0          00fa00d6b552f41b          rootvg          active
hdisk1          00c8d23057b60c26          None
hdisk2          00c8cf4057f2d781          None
```

## Geographically Mirrored Volume Groups

Move cursor to desired item and press Enter.

List Geographic Volume Group Information

**Manage Geographically Mirrored Volume Groups with Mirror Pools**

Manage Legacy Geographically Mirrored Volume Groups

Replace a Remote Physical Volume with Another on the same Site

Verify Mirror Copy Site Locations for a Volume Group

F1=Help  
F9=Shell

F2=Refresh  
F10=Exit

F3=Cancel  
Enter=Do

F8=Image



## Manage Geographically Mirrored Volume Groups with Mirror Pools

Move cursor to desired item and press Enter.

Add Remote Physical Volumes to a Mirror Pool  
Remove Remote Physical Volumes from a Mirror Pool  
Add a Remote Site Mirror Copy to a Mirror Pool  
Remove a Remote Site Mirror Copy from a Mirror Pool  
**Configure Mirroring Properties of a Mirror Pool**  
Activate an Asynchronously Mirrored Volume Group

F1=Help  
F9=Shell

F2=Refresh  
F10=Exit

F3=Cancel  
Enter=Do

F8=Image



# Converting to asynchronous mode (cont)

## Configure Mirroring Properties of a Mirror Pool

Move cursor to desired item and press Enter.

List Asynchronous Mirroring Information for Mirror Pools

**Convert to Asynchronous Mirroring for a Mirror Pool**

Change Asynchronous Mirroring Attributes for a Mirror Pool

Convert to Synchronous Mirroring for a Mirror Pool

```

+-----+
|                                     |
|                               Select a VOLUME GROUP name |
|                                     |
| Move cursor to desired item and press Enter. |
|                                     |
|      glvm_vg |
|                                     |
| F1=Help           F2=Refresh           F3=Cancel |
| F8=Image          F10=Exit             Enter=Do  |
| /=Find            n=Find Next          |
|                                     |
+-----+

```

F1=Help  
F9=Shell





# Converting to asynchronous mode (cont)

## Convert to Asynchronous Mirroring for a Mirror Pool

Type or select values in entry fields.  
 Press Enter AFTER making all desired changes.

	[Entry Fields]	
* VOLUME GROUP Name	glvm_vg	
* MIRROR POOL Name	glvm1	
Logical Volume for I/O Cache	[ ]	+
I/O Cache High Water Mark Value	[100]	#

Logical Volume for I/O Cache

Move cursor to desired item and press Enter.

**glvm1\_cache**

F1=Help	F2=Refresh	F3=Cancel
F5=Reset	F10=Exit	Enter=Do
F9=Shell	/=Find	n=Find Next

F1=Help  
 F5=Reset  
 F9=Shell



# Converting to asynchronous mode (cont)

## Convert to Asynchronous Mirroring for a Mirror Pool

Type or select values in entry fields.  
Press Enter AFTER making all desired changes.

- \* VOLUME GROUP Name
- \* MIRROR POOL Name
- Logical Volume for I/O Cache
- I/O Cache High Water Mark Value

```
[Entry Fields]
glvm_vg
glvm1
[glvm1_cache]          +
[075]                   #
```

Specifies the I/O-cache high watermark.  
The value is the percent of I/O cache size.  
The default value is 100%.

F1=Help  
F5=Reset  
F9=Shell

F2=Refresh  
F6=Command  
F10=Exit

F3=Cancel  
F7=Edit  
Enter=Do

F4=List  
F8=Image



## Converting to asynchronous mode (cont)

- Or from the command line:

```
/usr/sbin/chmp -A -m'glvm1' -c'glvm1_cache' -h'75' glvm_vg  
/usr/sbin/chmp -A -m'glvm2' -c'glvm2_cache' -h'75' glvm_vg
```

# Show asynchronous configuration

```

lsmpl glvm_vg
VOLUME GROUP:          glvm_vg          Mirror Pool Super Strict: yes

MIRROR POOL:          glvm1            Mirroring Mode:          ASYNC
MIRROR POOL:          glvm2            Mirroring Mode:          ASYNC

glvm1:/:# lsmpl -AL glvm_vg
VOLUME GROUP:          glvm_vg          Mirror Pool Super Strict: yes

MIRROR POOL:          glvm1            Mirroring Mode:          ASYNC
ASYNC MIRROR STATE:  inactive          ASYNC CACHE LV:          glvm1_cache
ASYNC CACHE VALID:   yes                ASYNC CACHE EMPTY:       yes
ASYNC CACHE HWM:     75                  ASYNC DATA DIVERGED:    no

MIRROR POOL:          glvm2            Mirroring Mode:          ASYNC
ASYNC MIRROR STATE:  active             ASYNC CACHE LV:          glvm2_cache
ASYNC CACHE VALID:   yes                ASYNC CACHE EMPTY:       no
ASYNC CACHE HWM:     75                  ASYNC DATA DIVERGED:    no

```



# Example of simple script

- Scripts

- Example of scripts to start and stop RPV Server and Client (start the remote server first)

## Start

```

if [ $# -ne 1 ]; then
    echo "try entering client|server"
    exit 1
fi
case $1
in
    client) echo "Starting Client"
            lspv
            mkdev -l hdisk2
            lspv
            varyonvg glvm_vg
            mount /data
            ;;
    server) echo "Starting Server"
            mkdev -l rpvserver0
            ;;
esac

```

## Stop

```

if [ $# -ne 1 ]; then
    echo "try entering client|server"
    exit 1
fi
case $1
in
    client) echo "Stopping Client"
            umount /data
            varyoffvg glvm_vg
            rmdev -l hdisk2
            lspv
            ;;
    server) echo "Stopping Server"
            rmdev -l rpvserver0
            ;;
esac

```