

opentext™

ESM HAについて

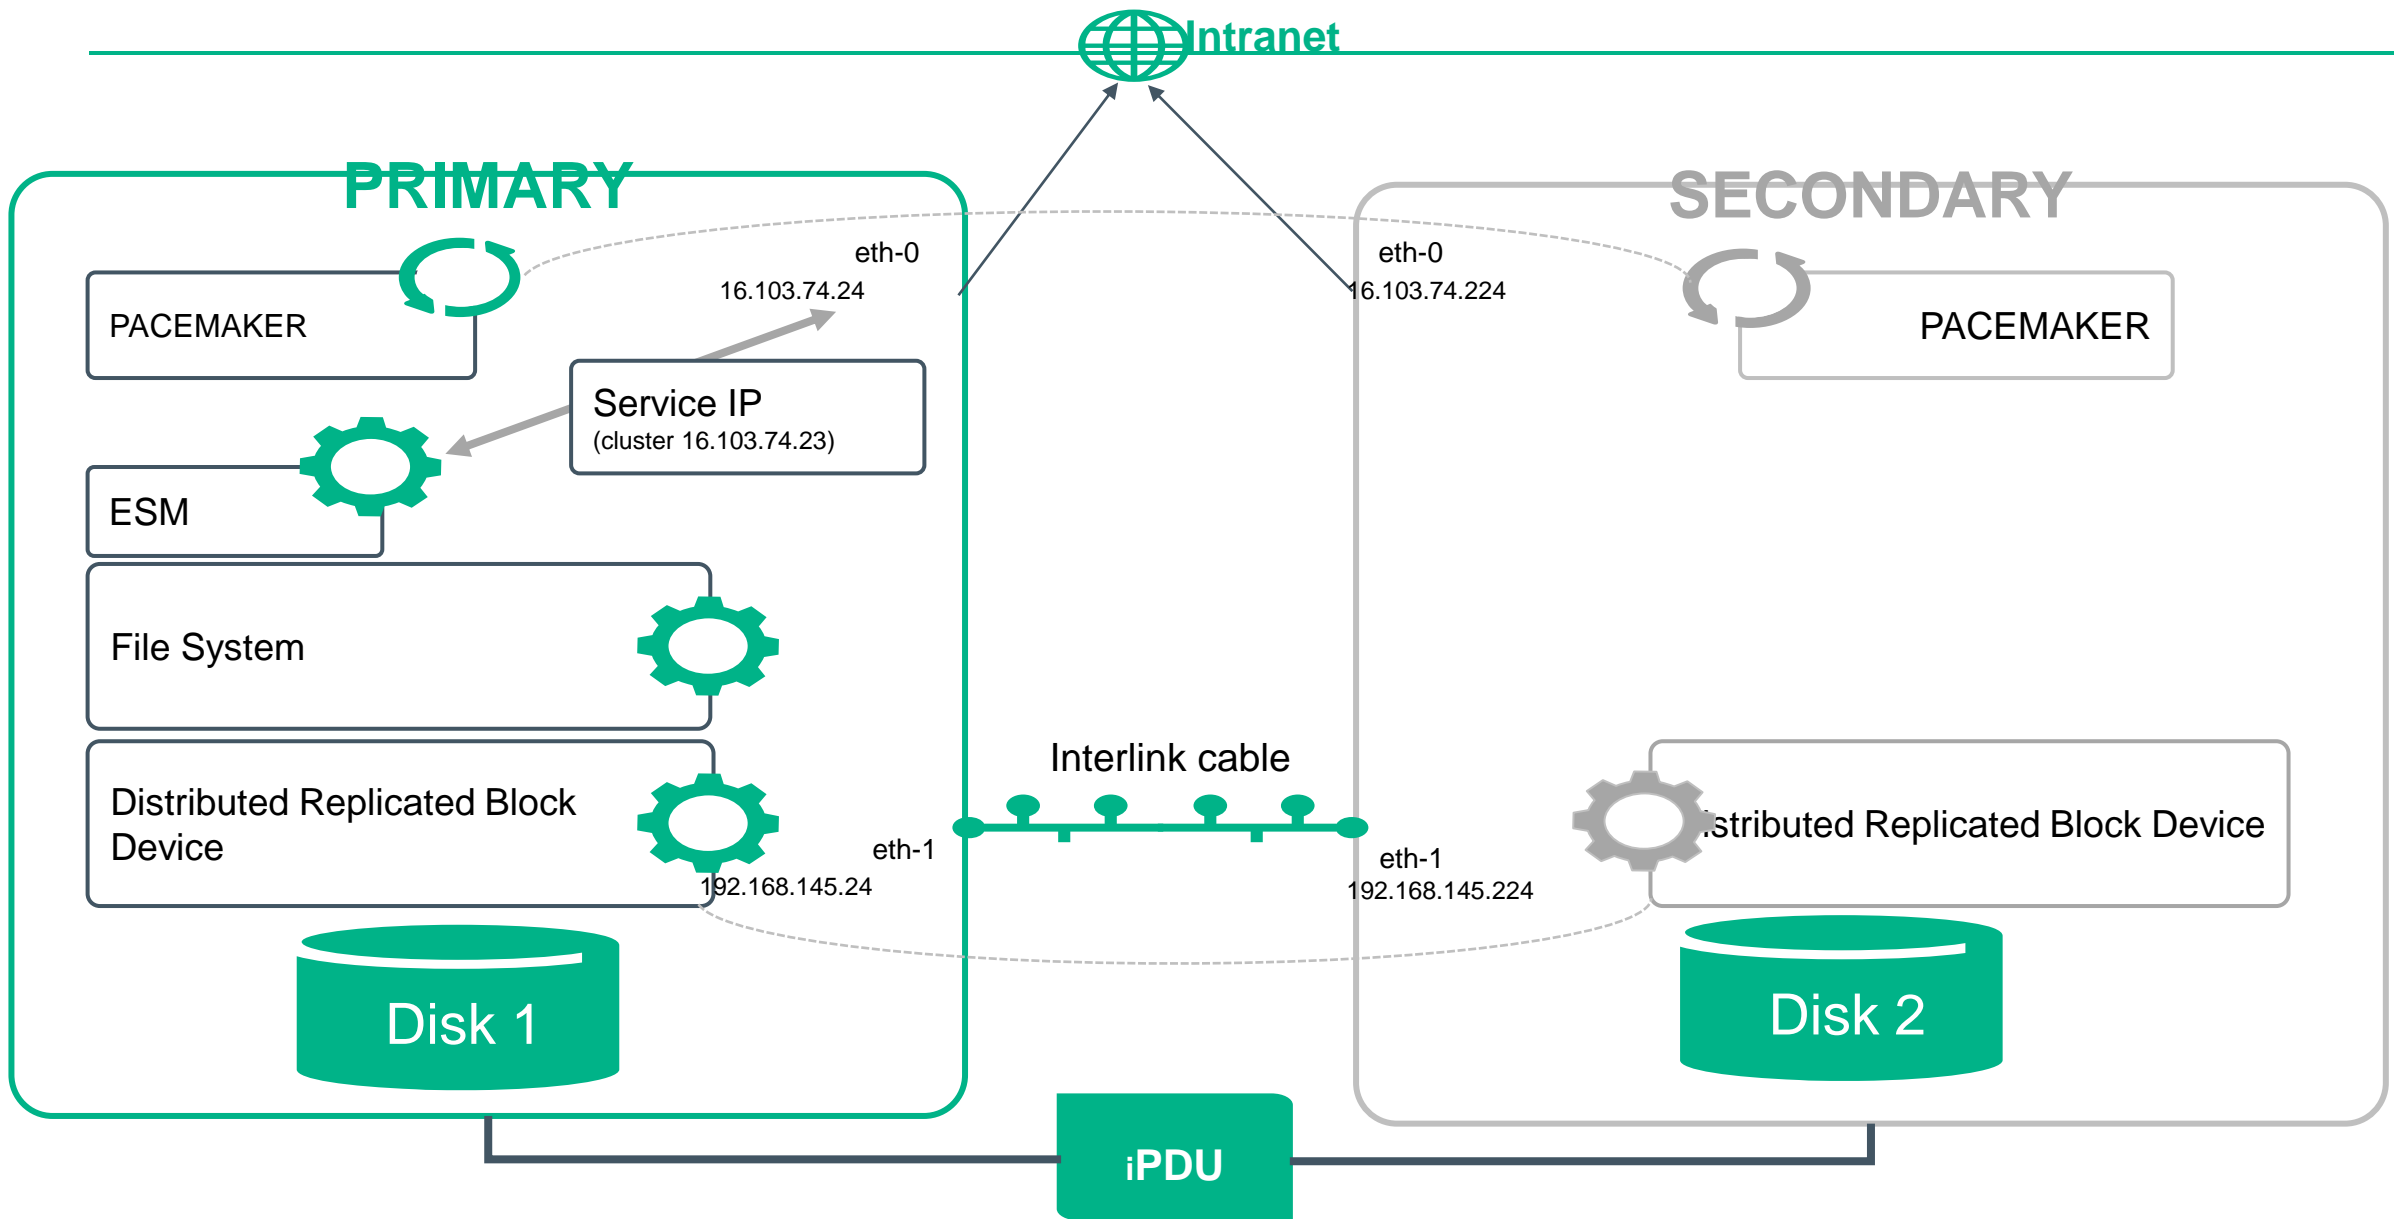
2023/03

| ArcSight Support

Agenda

- **ESM HA Component Overview and Terminology**
 - ESM HA Terminology
 - Hardware Requirements

ESM HA Terminology



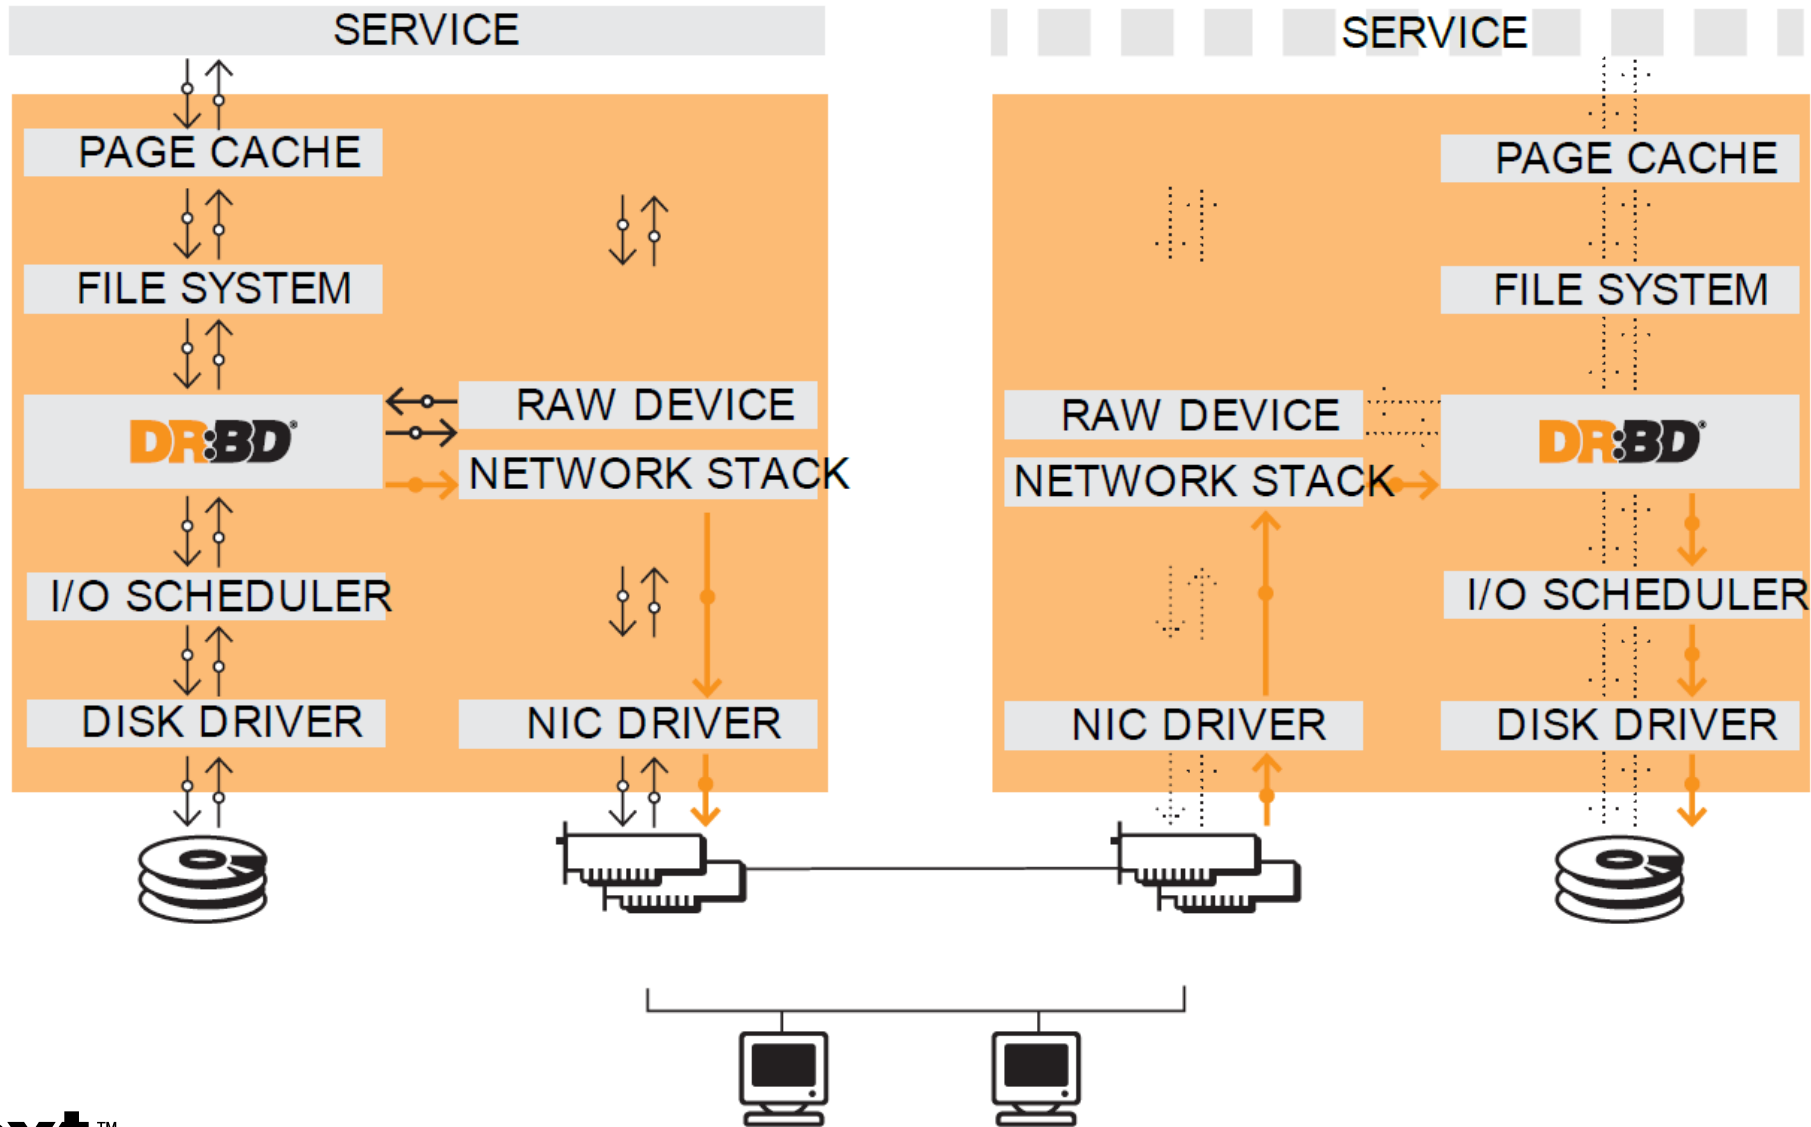
ESM HA Terminology

- **HA:** High Availability or Failover Cluster Solution for ESM
- **Nodes:** Servers part of the Clustering Solution (Maximum 2)
- **Passive/Failover:** Only Cluster configuration currently supported
- **HDD:** Each server/node must have it's own hard drive.
- **Service IP:** Shared IP address used by the Cluster solution. Never directly configured on any server.
- **Node IP:** IP address directly configured on each Node of the cluster.
- **Service Hostname:** Shared hostname for the cluster. Used for console login and Command Center Access
- **Primary Hostname:** Hostname assigned to the primary server/node of the cluster
- **Secondary Hostname:** Hostname assigned to the secondary server/node of the cluster
- **Primary Cable IP:** IP address configured on the interconnect network interface of the primary server/node
- **Secondary Cable IP:** IP address configured on the interconnect network interface of the secondary server/node

ESM HA Terminology Contd.

- **Shared Disk:** Mount point or partition where ESM directories will be located on both servers. Each server must have its own disk
- **Metadata Volume:** Small raw partition used for the clustering technology to determine the synchronization status of the drives.
- **Connected Hosts:** IP addresses used by the Clustering Software to confirm if a particular node has internet connectivity or not.
- **Preferred Primary:** Server/Node that where ESM services will be running regularly
- **Heartbeat/Pacemaker:** Mechanism used to ensure the nodes on the Cluster are online
- **DRBD:** Synchronization Software implemented for HA

Distributed Replicated Block Device (DRBD) with Linux I/O stack



Distributed Replicated Block Device (DRBD)

- The Distributed Replicated Block Device (DRBD) is a software-based, shared-nothing, replicated
- storage solution mirroring the content of block devices (hard disks, partitions, logical volumes etc.) between hosts.

- DRBD mirrors data:
 - !!!Please don't confuse this with real mirroring where application is writing to both disks at the same time.
 - In DRBD "mirroring" is Not writing to both devices simultaneously.
 - All the time primary disk/node is written first, and after, that updated block is sent to secondary disk/node.
 - in real time.
 - Replication occurs continuously while applications modify the data on the device.
 - transparently.
 - Applications need not be aware that the data is stored on multiple hosts.
 - synchronously or asynchronously.
 - With synchronous mirroring, applications are notified of
 - write completions after the writes have been carried out on all hosts. With asynchronous
 - mirroring, applications are notified of write completions when the writes have completed
 - locally, which usually is before they have propagated to the other hosts.

Split Brain

Split brain is a situation where, due to temporary failure of all network links between cluster nodes, and possibly due to intervention by a cluster management software or human error, both nodes switched to the primary role while disconnected.

This is a potentially harmful state, as it DRBD Features 9 implies that modifications to the data might have been made on either node, without having been replicated to the peer.

Thus, it is likely in this situation that two diverging sets of data have been created, which cannot be trivially merged.

Split Brain

How generation identifiers change

17.2.3.1. Start of a new data generation

When a node loses connection to its peer (either by network failure or manual intervention), DRBD modifies its local generation identifiers in the following manner:

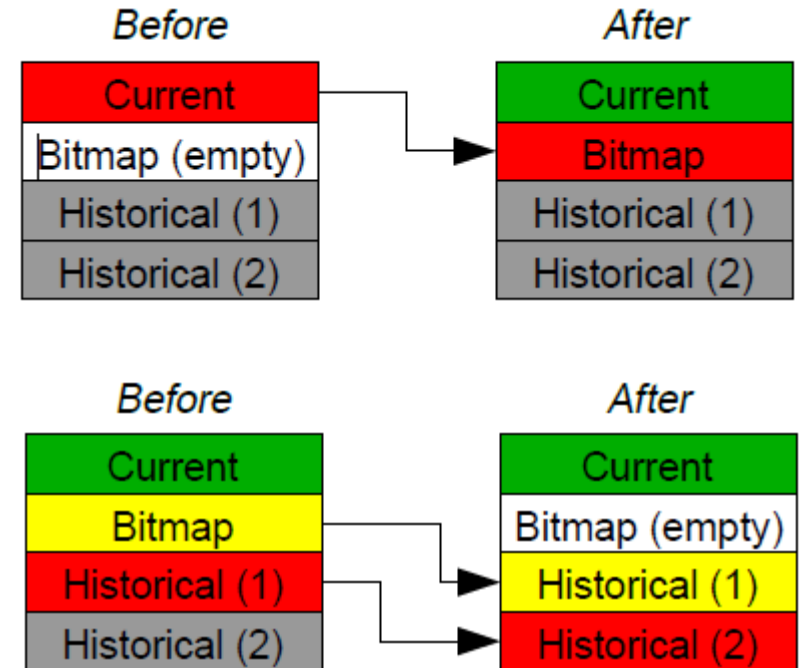
1. A new UUID is created for the new data generation. This becomes the new current UUID for the primary node.
2. The previous UUID now refers to the generation the bitmap is tracking changes against, so it becomes the new bitmap UUID for the primary node.
3. On the secondary node, the GI tuple remains unchanged.

Completion of re-synchronization

When re-synchronization concludes, the following changes are performed:

Bitmap UUIDs match, current UUIDs do not. The local node detects that its current UUID differs from the peer's current UUID, and that the bitmap UUID's match. This is split brain, but one where the data generations have the same parent. This means that DRBD invokes split brain auto-recovery strategies, if configured. Otherwise, DRBD disconnects and waits for manual split brain resolution.

Neither current nor bitmap UUIDs match. The local node detects that its current UUID differs from the peer's current UUID, and that the bitmap UUID's *do not* match. This is split brain with unrelated ancestor generations, thus auto-recovery strategies, even if configured, are moot. DRBD disconnects and waits for manual split brain resolution.



Hardware Requirements

Hardware Requirements

The 2 nodes of the cluster must contain the same hardware specs otherwise they need to have similar values for the following resources

- CPU
- RAM
- HDD
- Ethernet Ports

HDD Partitioning

- HDD configuration on both servers has to be identical
- Do not use default /home partition configuration defined by RHEL or CentOS
- /root, swap and system partitions may use default settings
- Shared disk:
 - Must have it's own partition
 - Will be mounted under /opt directory
 - Recommended to use LVM partition to allow resizing
 - ESM Files will be located here
 - Contents of this disk are synchronized by DRBD software
 - Xfs or ext4 file systems supported

HDD Partitioning Contd.

- Metadata volume
 - Partition used by DRBD to ensure hard drives are properly synchronized
 - Raw partition so it will never have a proper file system
 - Very small size, not bigger than 512MB
 - Recommended to use LVM for such partition
 - Will never have a mount point
 - Will never have a file system configured

Ethernet Ports

- Minimum 2 Ethernet ports
 - Eth0: node IP
 - Eth1: interconnect IP
 - Interconnect IP used by DRBD to synchronize the disks between the nodes
 - Interconnected via crossover cable. Avoid using extra hardware to ensure connectivity
 - Very close proximity between the nodes as Interconnect cable must be crossover
 - 5 IP addresses

The image features the OpenText logo in a bold, white, sans-serif font, centered horizontally. The logo is set against a dark blue background with several glowing, curved lines in a lighter blue color that sweep across the frame from the bottom left towards the top right. The lines have a soft, ethereal glow and vary in thickness, creating a sense of motion and depth. The overall aesthetic is modern and technological.

opentextTM