ASK THE EXPERT: EXTREME HIGH AVAILABILITY WITH POSTGRESQL

9 June, 2022

EDB™

OUR SPEAKERS

HOST

Kevin Li
Director Sales
Engineer, EMEA

SPEAKER

Michael Willer
Senior Sales Engineer,
EMEA

EDB™

AGENDA

- Why extreme HA?
- Solutions for HA
- How EDB Delivers extreme HA
- Architectures
- Failover demo

EDB™

WHY EXTREME HA?

HIGH AVAILABILITY

High availability (HA) is a characteristic of a system, which aims to ensure an agreed level of operational performance, usually uptime, for a higher than normal period.

Key principles:

- Eliminate single point of failure
- Reliable crossover
- Detection of failures

Ref: https://en.wikipedia.org/wiki/High_availability

EDB™

SCHEDULED/UNSCHEDULED DOWNTIME

- **Scheduled/planned downtime** is a result of maintenance that is disruptive to system operation and usually cannot be avoided with a currently installed system design.
  - It include patches to system software that require a reboot or system configuration changes that only take effect upon a reboot.

- **Unscheduled/Unplanned downtime** is the result of downtime events due to some physical failures/events, such as hardware or software failure or environmental anomaly.
  - For example, power outages, failed CPU or RAM components (or possibly other hardware components failure), network failure, security breaches, or various applications, middleware, and operating system failures result in Unplanned outage/Unscheduled downtime.

EDB™

IMPACT OF DOWNTIME

Loss of Revenue

When the application is unavailable - you lose money from existing customers.

Disruption of Business

When the application is unavailable, you lose customers.

Cost of workarounds, increased employee costs.

Brand impact

Loss of confidence in the Brand.

Loss of market-share.

Human impact (Health, Life)

Health-care applications, Fire brigade, Telco, Police, …

EDB

LOSS OF REVENUE CALCULATION (EXAMPLE)

Loss of revenue per hour of downtime	$1,000,000					
Average uptime per year	99%	99.50%	99.90%	99.95%	99.99%	99.999%
Downtime per year (minutes)	5256	2628	525.6	262.8	52.56	5.256
Downtime per year (hours)	87.6	43.8	8.76	4.38	0.876	0.0876
Total loss of revenue per year	**$87,600,000**	**$43,800,000**	**$8,760,000**	**$4,380,000**	**$876,000**	**$87,600**

EXTREME HIGH AVAILABILITY

Key requirements:

- 24/7
- No downtime (scheduled or not)
- Geo-redundancy
- Fast automated failover
- No data loss (or as little possible)

EDB™

SOLUTIONS FOR HA

1 PRIMARY / x STANDBY ARCHITECTURES

Automatically detect failures

- Monitor database health - detect failure and takes action
- Promote Standby
- (optional) Reconfigure connection routing
- (optional) Avoid “split brain” scenario

Examples

- EDB Failover Manager (EFM)
- RepMgr
- Patroni

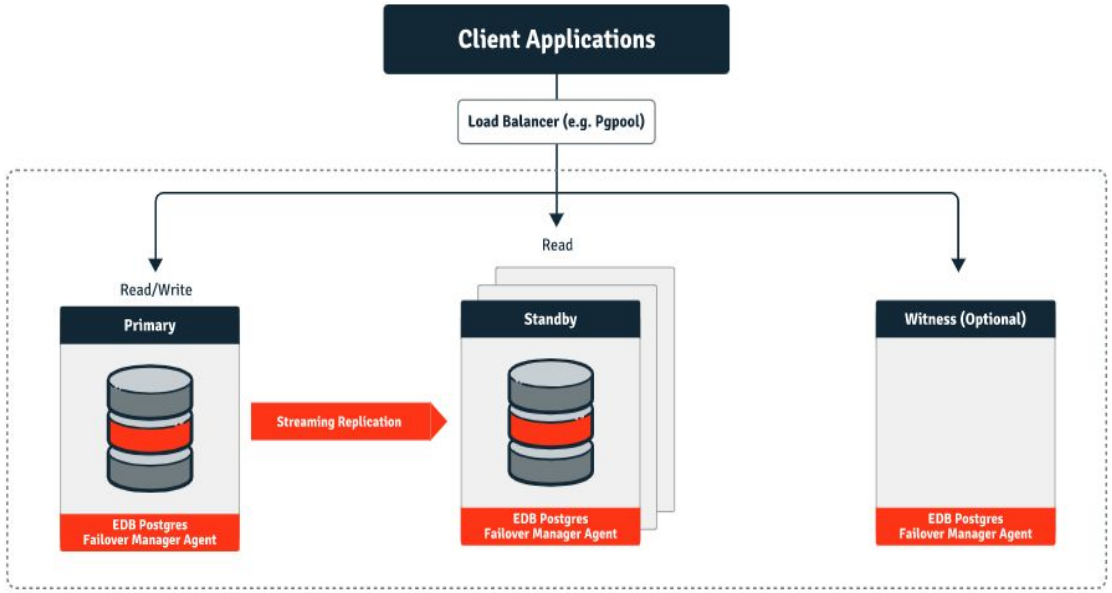

ACTIVE-ACTIVE ENVIRONMENTS

Multi-master replication enabling highly available and geographically distributed Postgres clusters

- Logical replication of data and schema enabled via standard PostgreSQL extension
- Data consistency options that span from immediate to eventual consistency
- Robust tooling to manage conflicts, monitor performance, and validate consistency
- Deploy natively to cloud, virtual, or bare metal environments

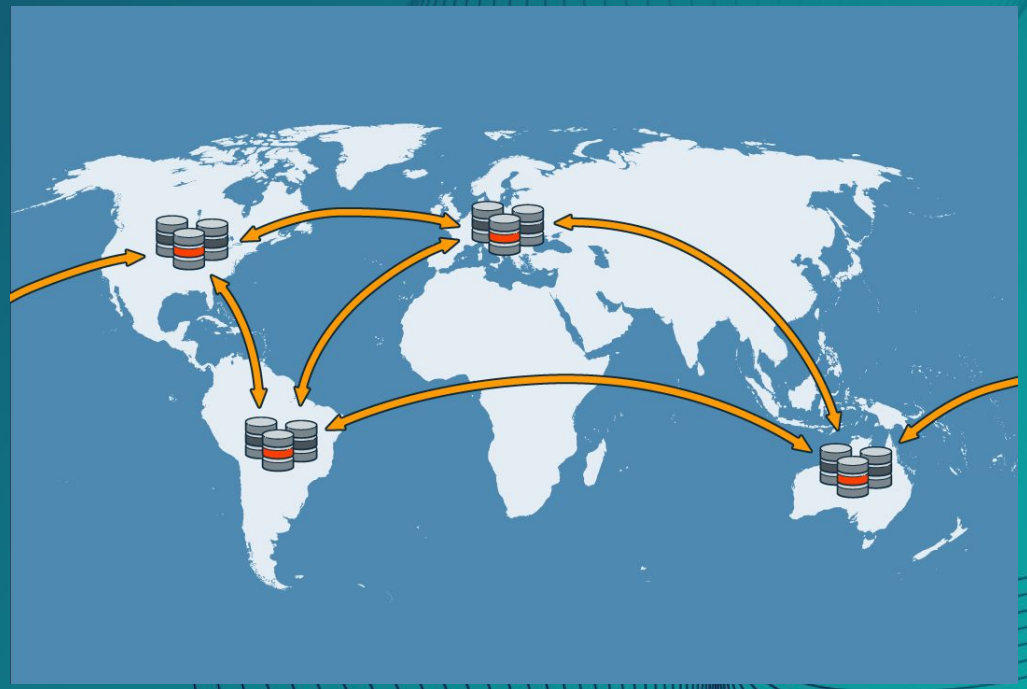

EDB™

HOW EDB DELIVERS EXTREME HA

EDB

EDB Postgres Distributed - Always-ON Platinum

- Multi-master cluster
- Mesh architecture to minimize latency between nodes
- Raft consensus layer
- Integrated with other services
  - Pooling, backup, proxy
- Multiple possible architectures
  - Logical standbys
  - Subscriber-only nodes
  - Witness nodes
- Cloud, on-premises, or hybrid

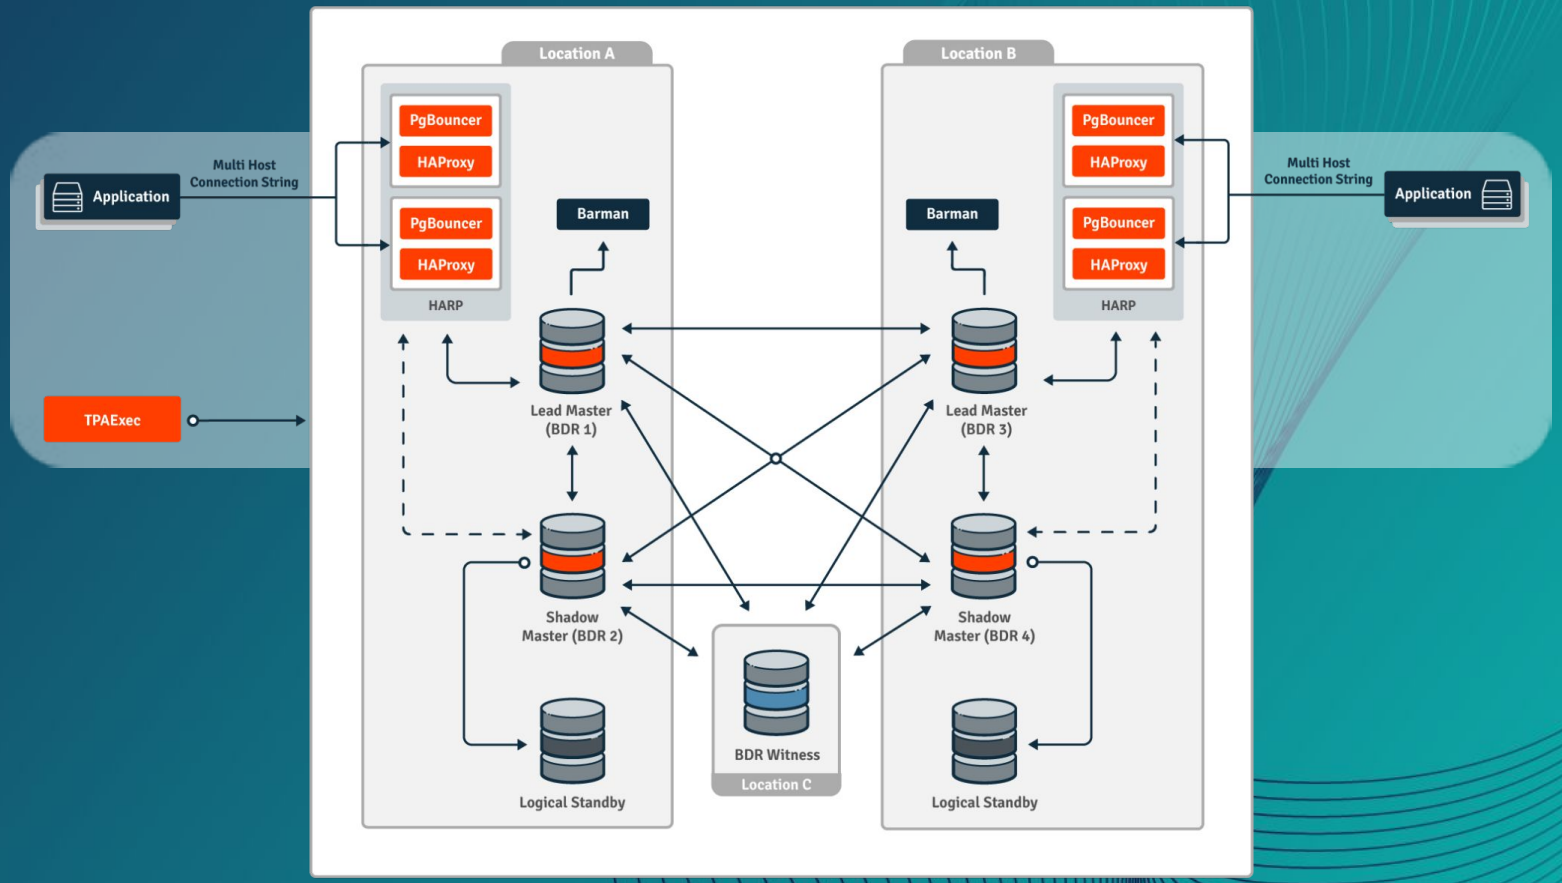

BDR Always On Platinum Architecture

ARCHITECTURES

EDB

MULTIPLE ARCHITECTURES TO FIT YOUR HA NEEDS

Four standard options support a range of requirements from basic to most advanced

Single Active Location

Always On Bronze (single active location)

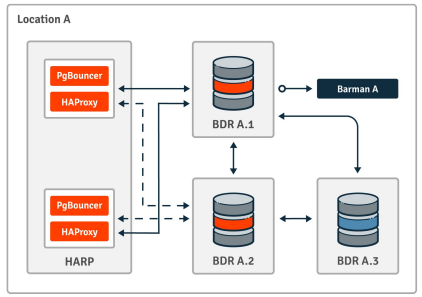

Always On Silver (single active location, backup in DR location)

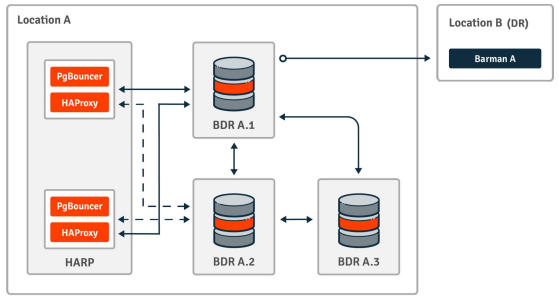

Maintains HA on node failure

Two Location, Active-Active Architectures

Always On Gold (two active locations)

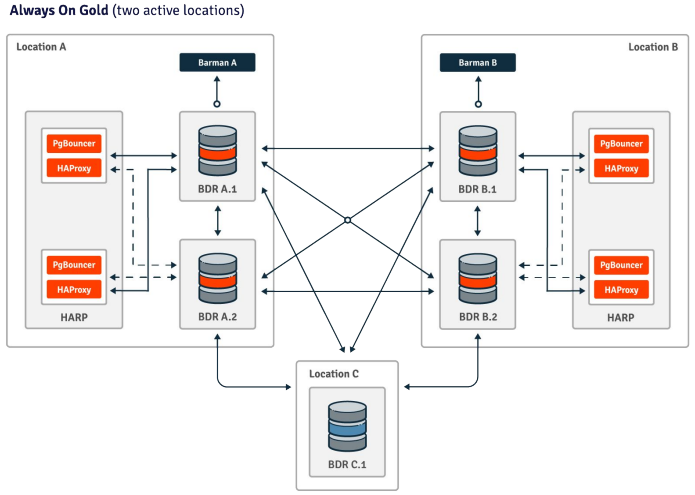

For deployments where quick provisioning of replacement nodes is possible to maintain local HA in case of node failure, includes optional witness node in third location

Always On Platinum (two active locations)

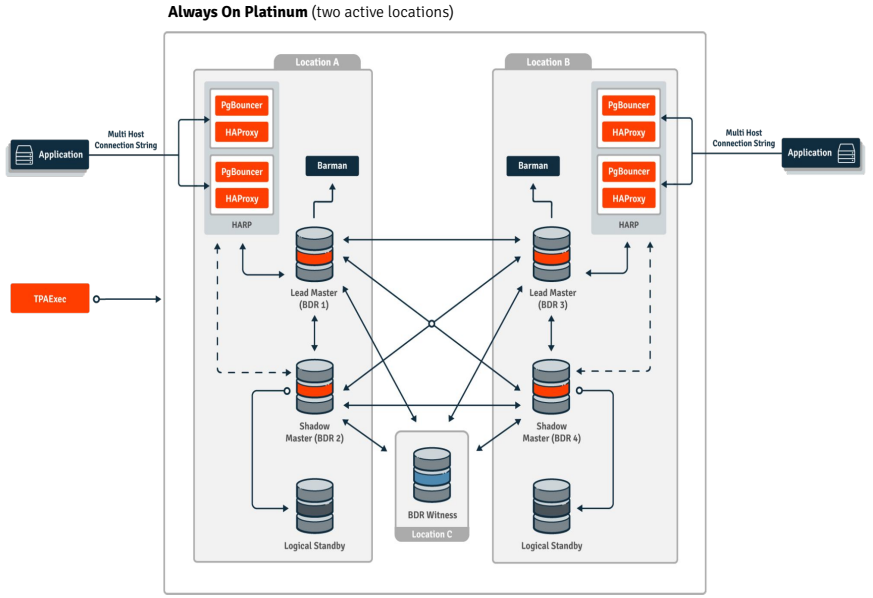

For on-premises deployments where hot standby is required to maintain local HA in case of hardware failure, includes optional witness node in third location

HOW TO CHOOSE THE RIGHT ARCHITECTURE

	Always On - Bronze	Always On - Silver	Always On - Gold	Always On - Platinum
Hardware failure protection	Yes	Yes	Yes	Yes
Location failure protection	No (unless Barman is moved offsite)	Yes - Recovery from backup	Yes - instant failover to fully functional site	Yes - instant failover to fully functional site
Failover to DR or Full second location	DR (if Barman is located offsite); NA otherwise	DR	Full second location	Full second location
Zero downtime upgrade	Yes	Yes	Yes	Yes
Support of AZs in public/ private cloud	Yes	Yes	Yes	Yes
Fast local restoration of high availability after device failure	No; time to restore HA: (1) VM prov + (2) approx 60 min/500GB	Yes; three local BDR nodes allow to maintain HA after device failure	No; time to restore HA: (1) VM prov + (2) approx 60 min/500GB	Yes; logical standbys can quickly be promoted to full BDR nodes
Cross data center network traffic	No	Backup traffic only	Full replication traffic	Full replication traffic
BDR license cost (price per core on each node)	2 BDR nodes	3 BDR nodes	4 BDR nodes	4 BDR nodes 2 logical standbys

DEMO

TERMINAL LAYOUT

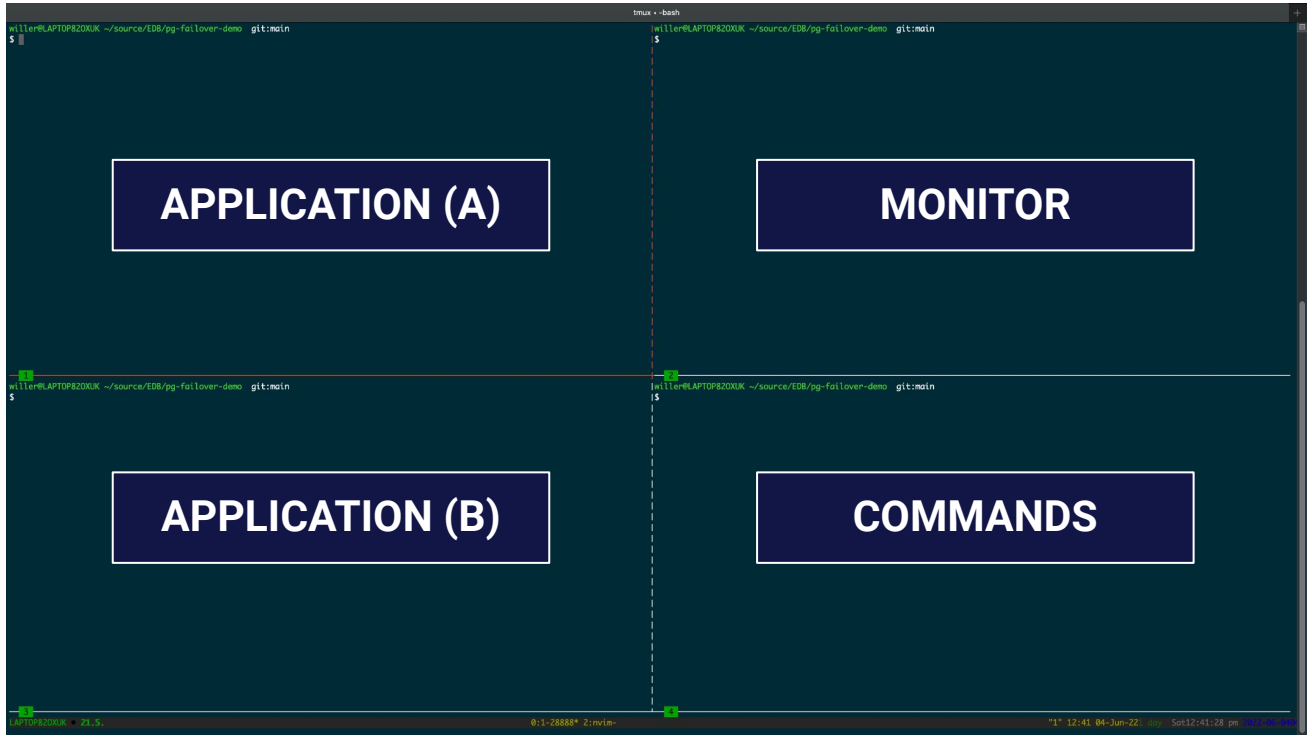

HARP MANAGER

Ensures database is running
elects lead

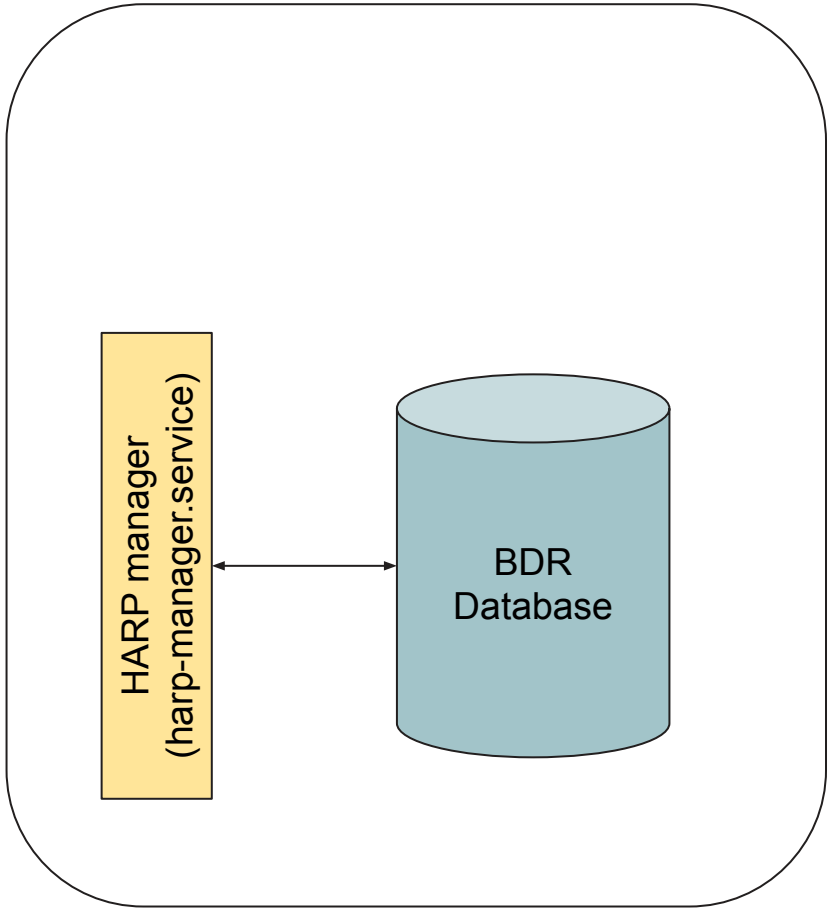

HARP PROXY

Directs traffic to the lead database

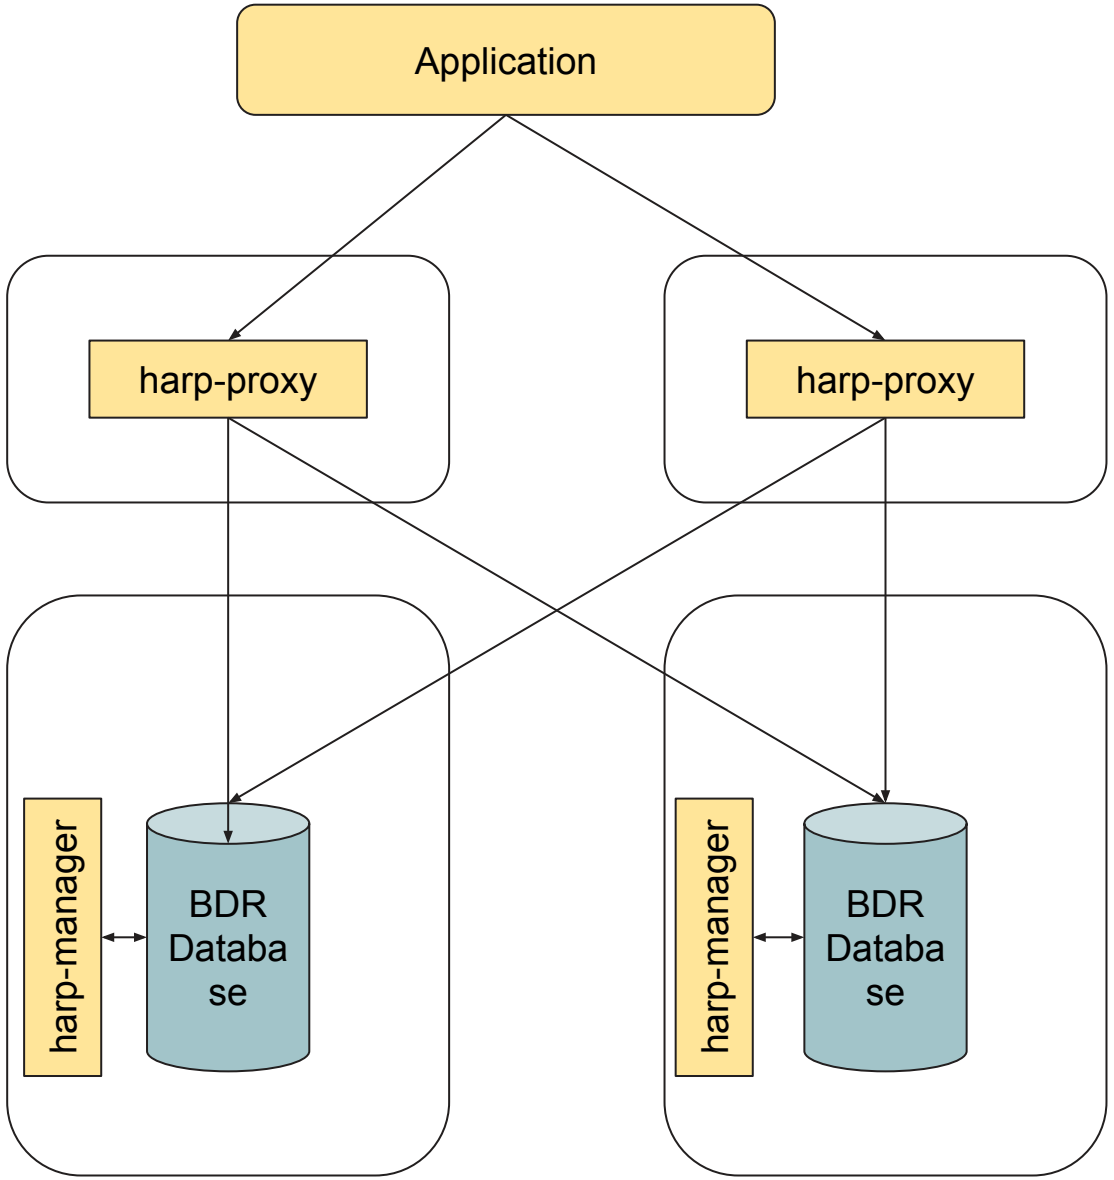

EDB POSTGRES DISTRIBUTED

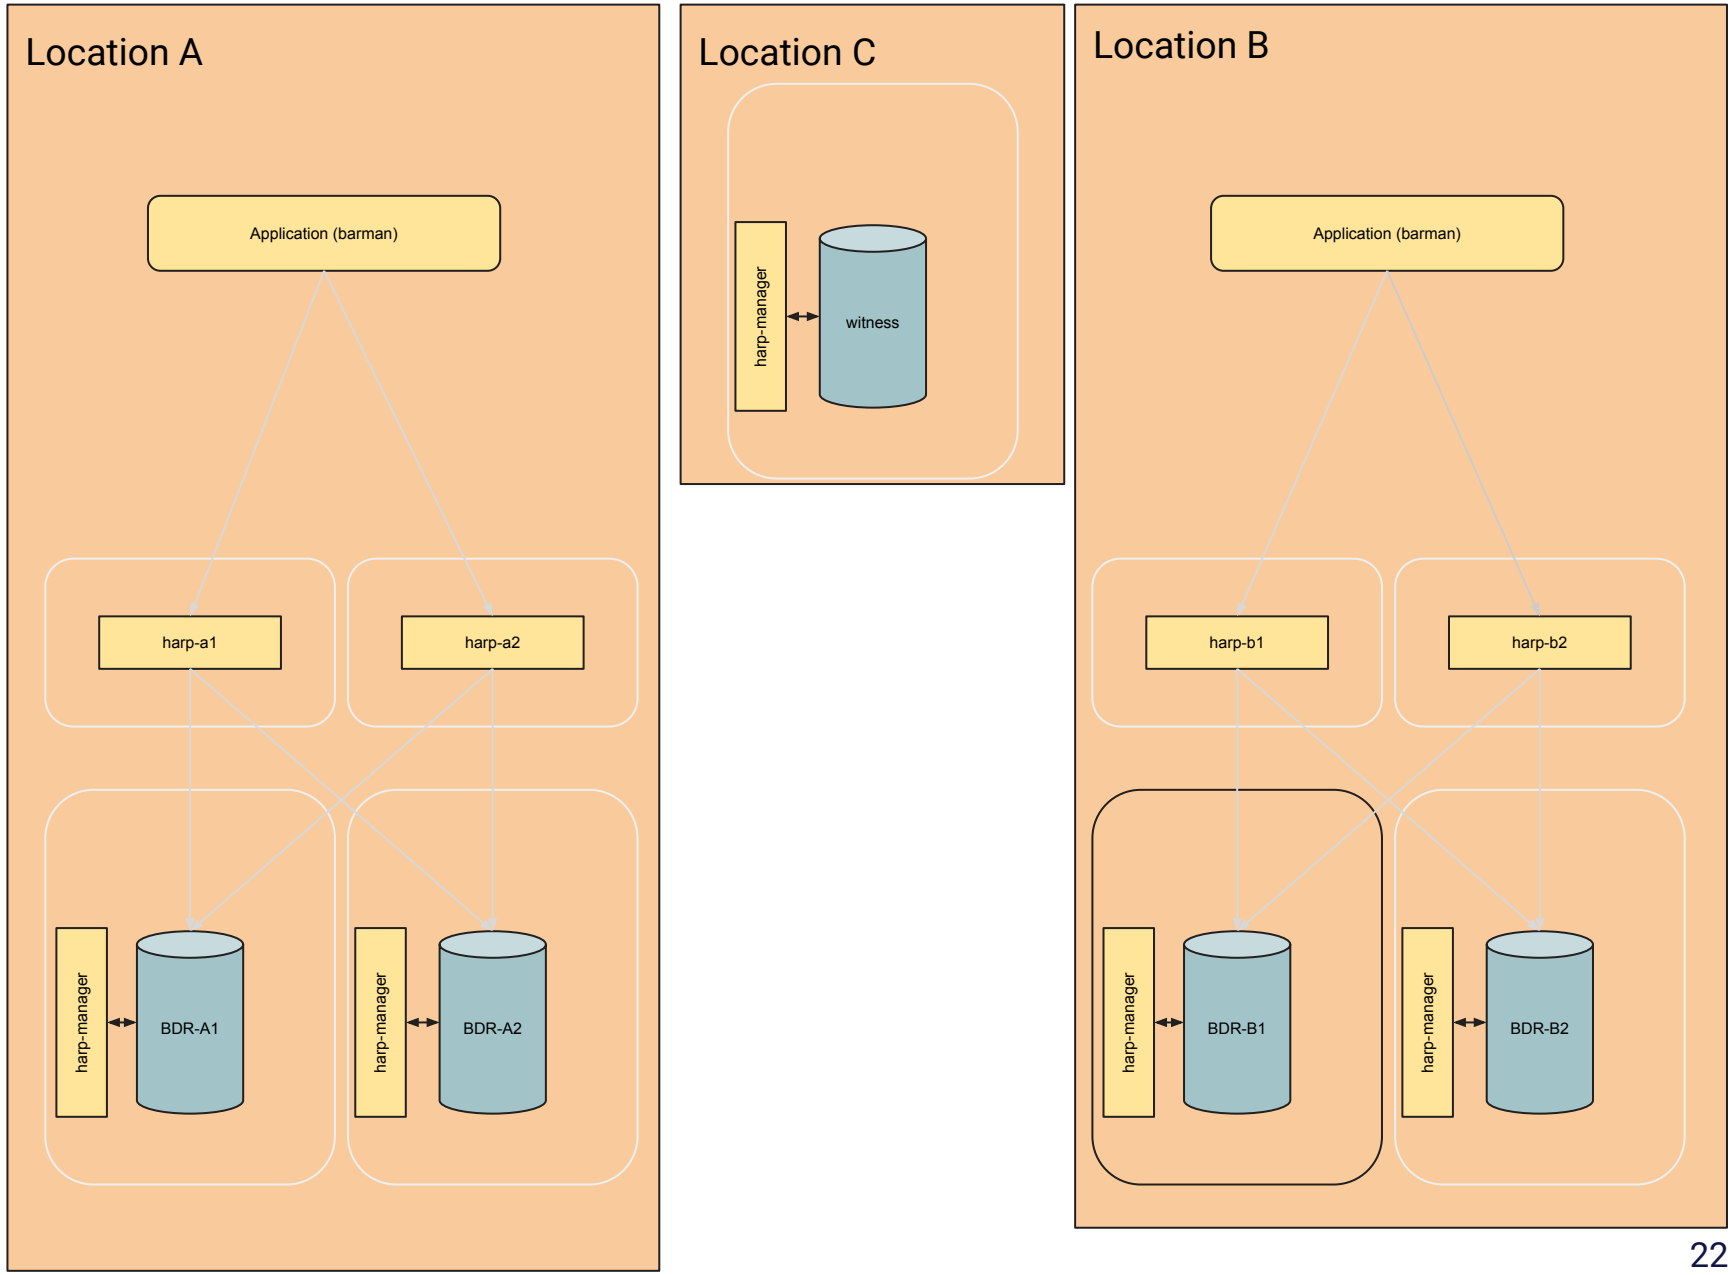

EDB POSTGRES DISTRIBUTED

Multiplex app connections

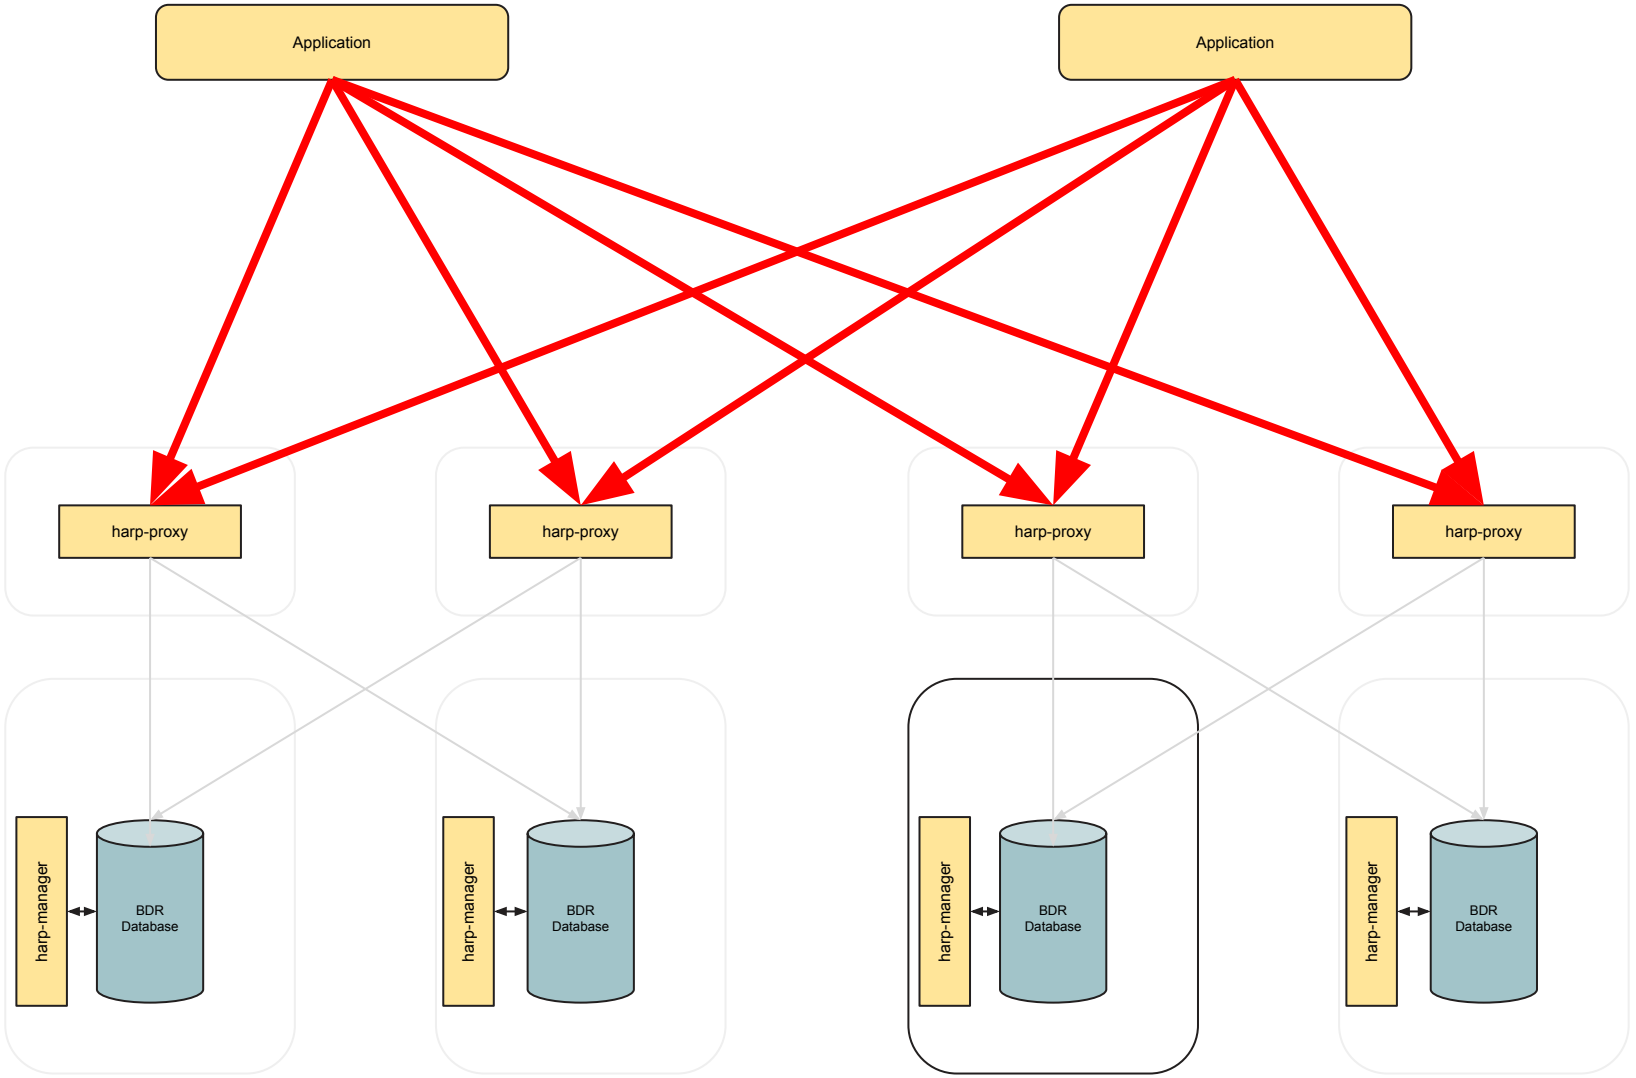

CURRENT

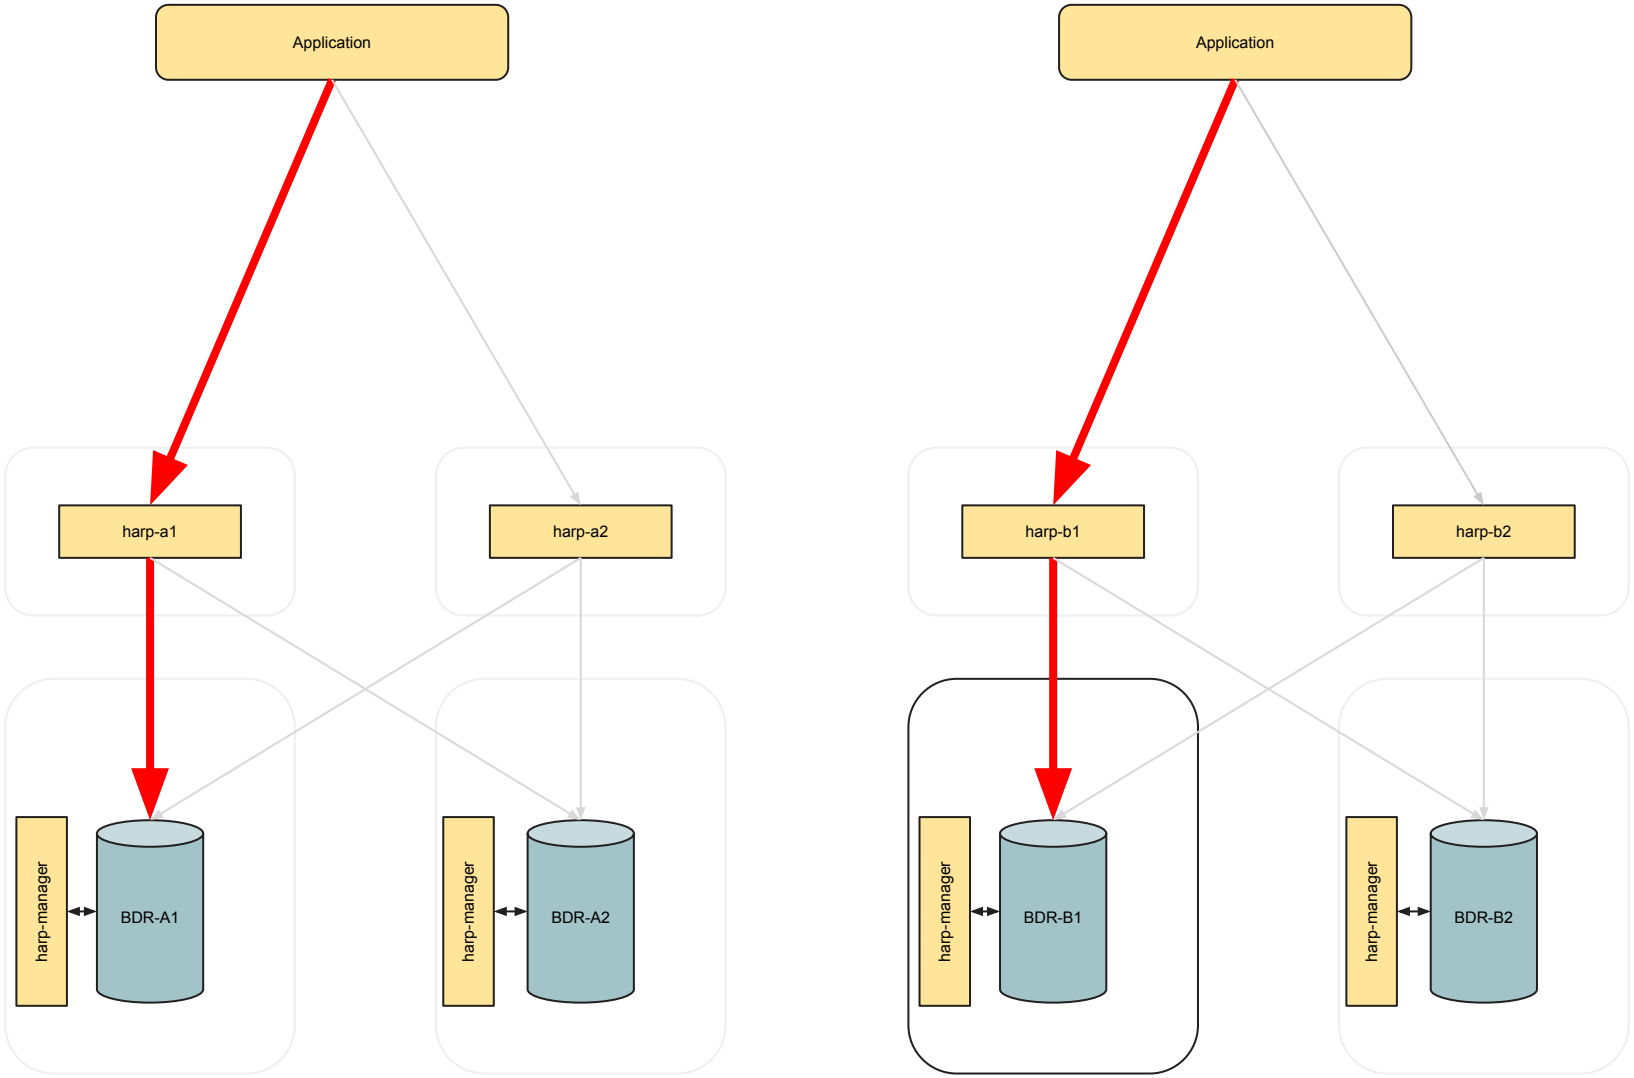
Application
harp-a1
harp-a2
harp-manager
BDR-A1
harp-manager
BDR-A2
Application
harp-b1
harp-b2
harp-manager
BDR-B1
harp-manager
BDR-B2

BDR-A1 FAILS

BDR-B2 FAILS (FOR LATER)

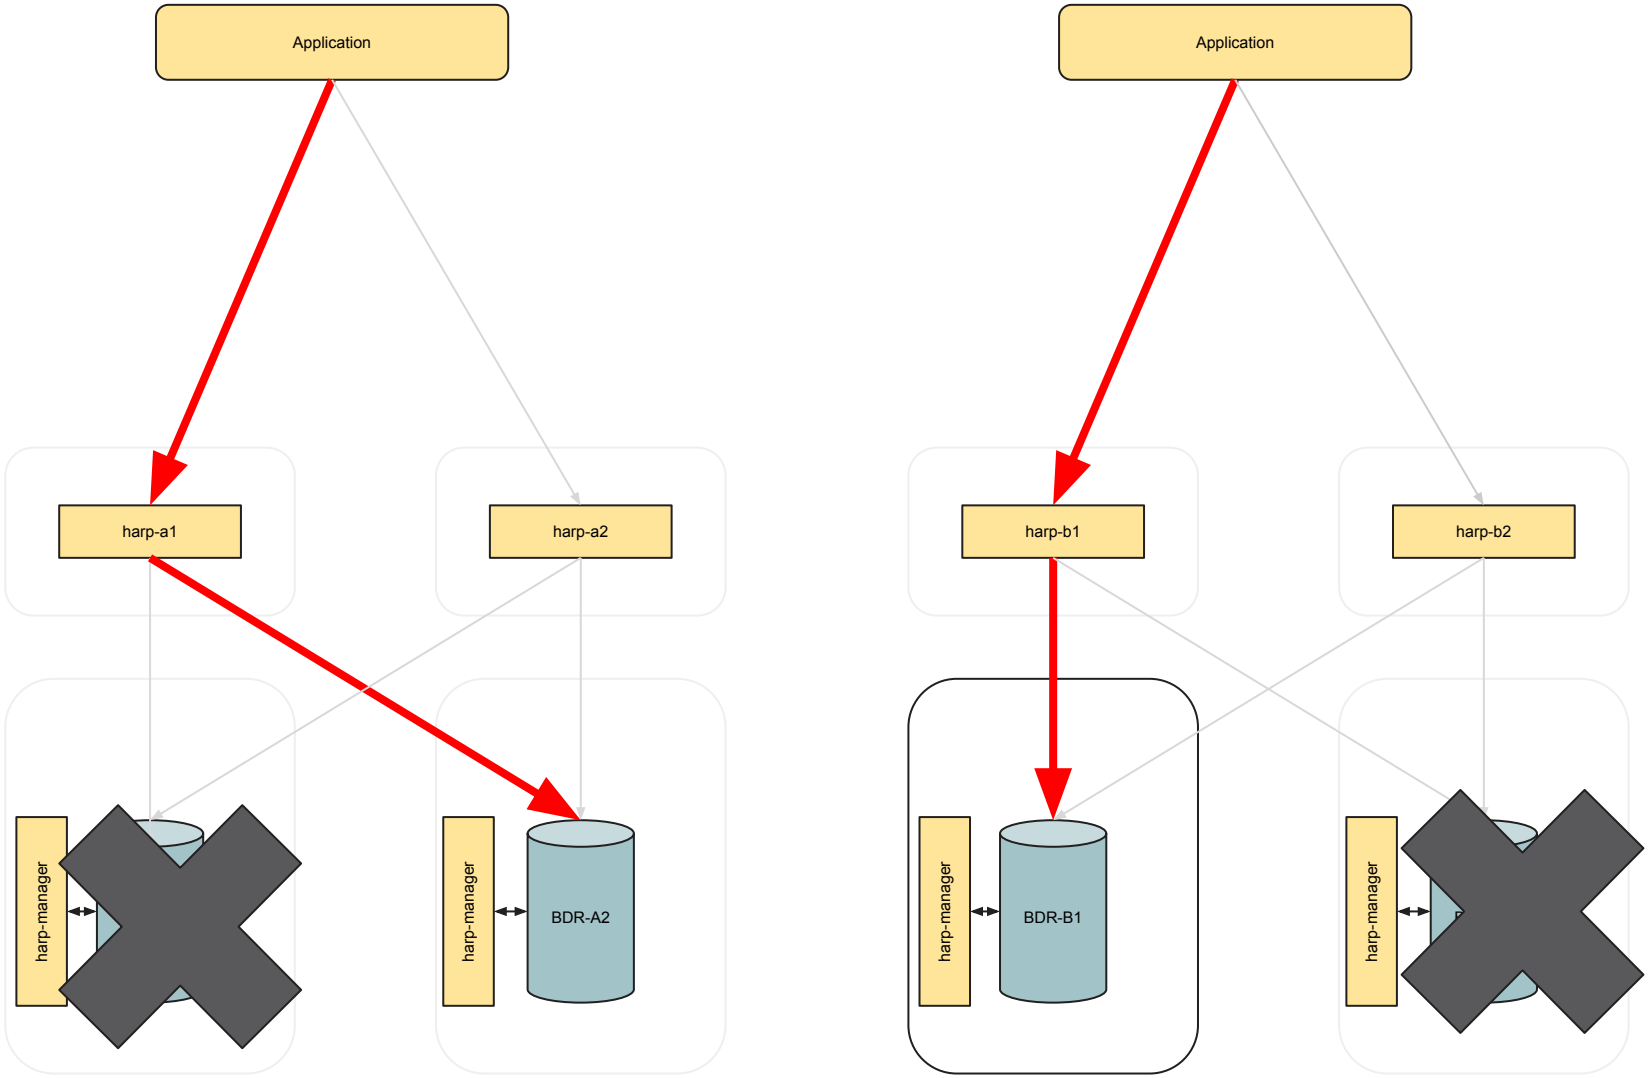

BACK TO DEFAULT

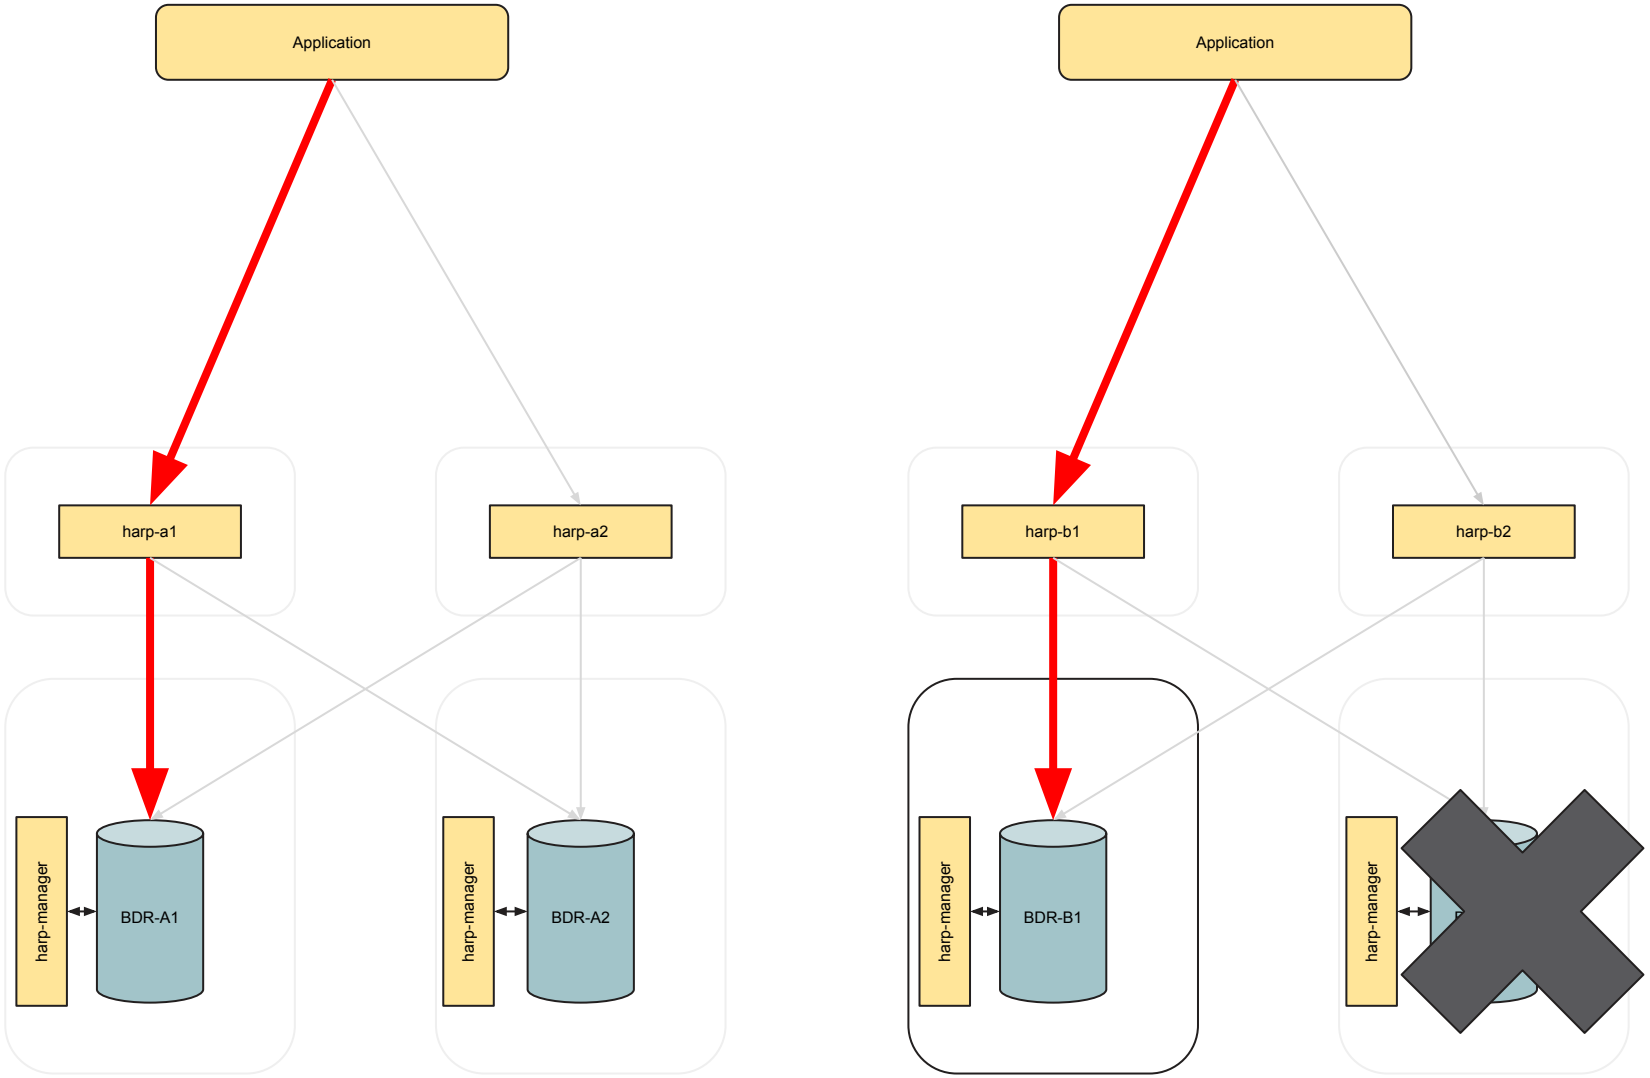

HARP PROXY FAILS (harp-a1)

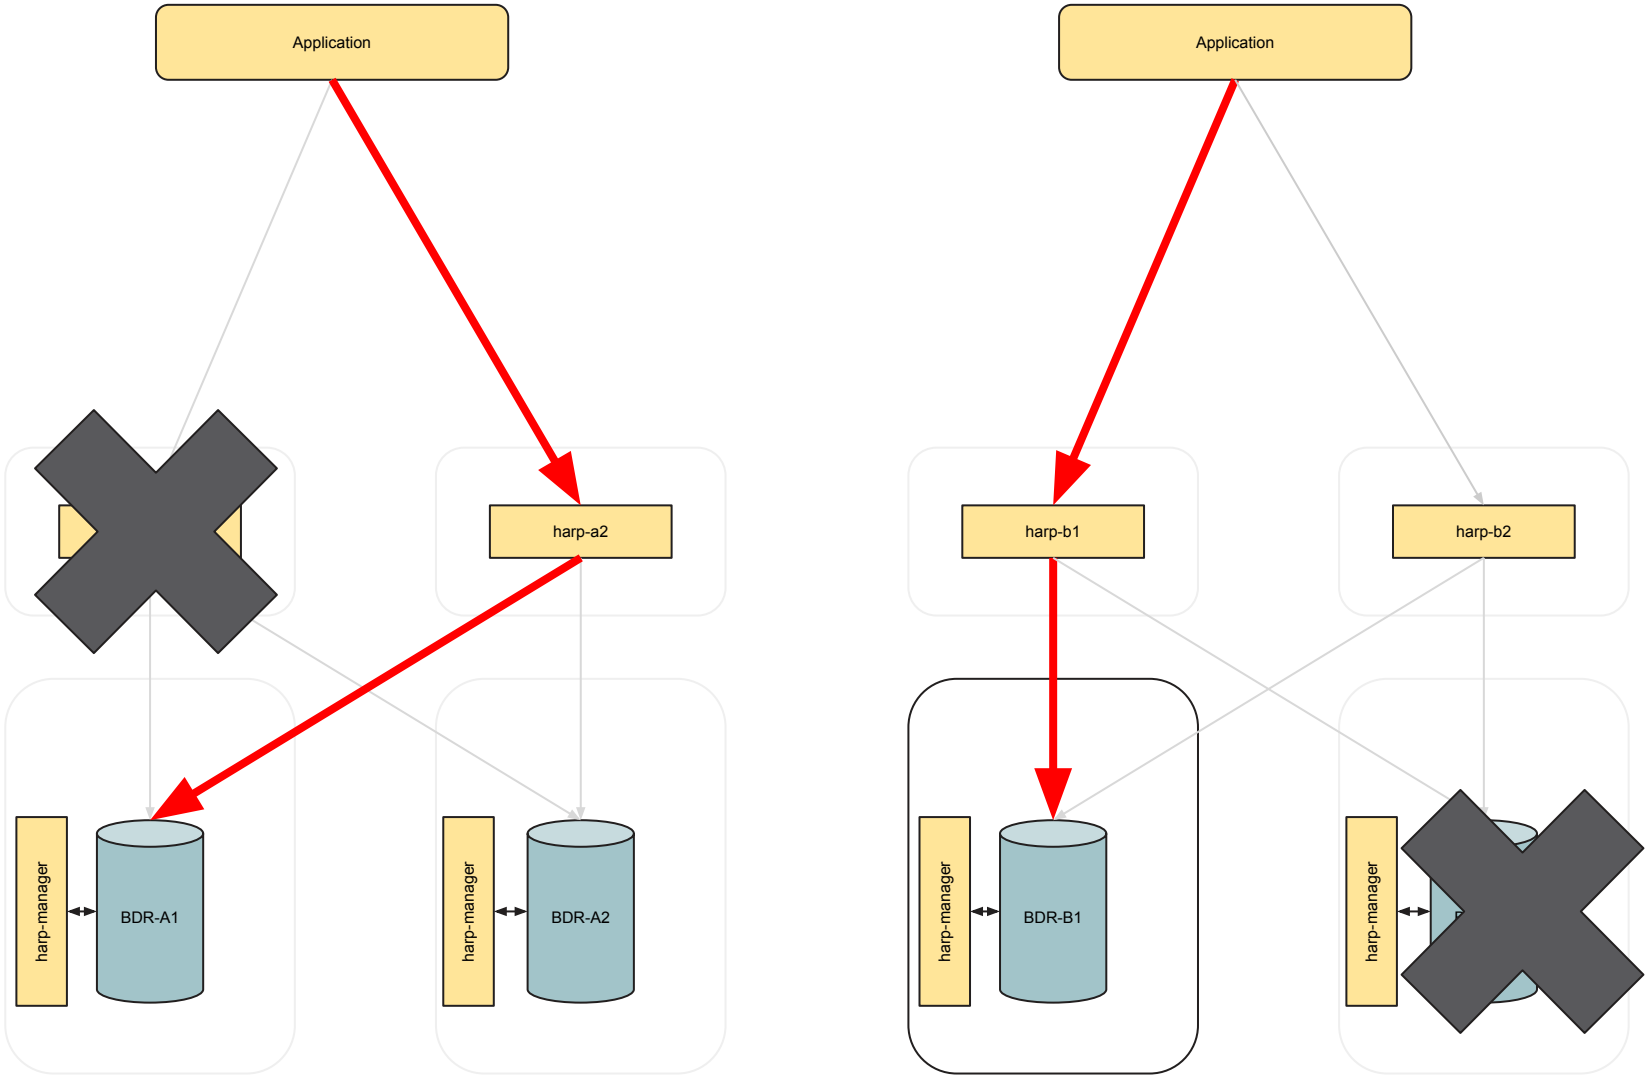

BACK TO DEFAULT

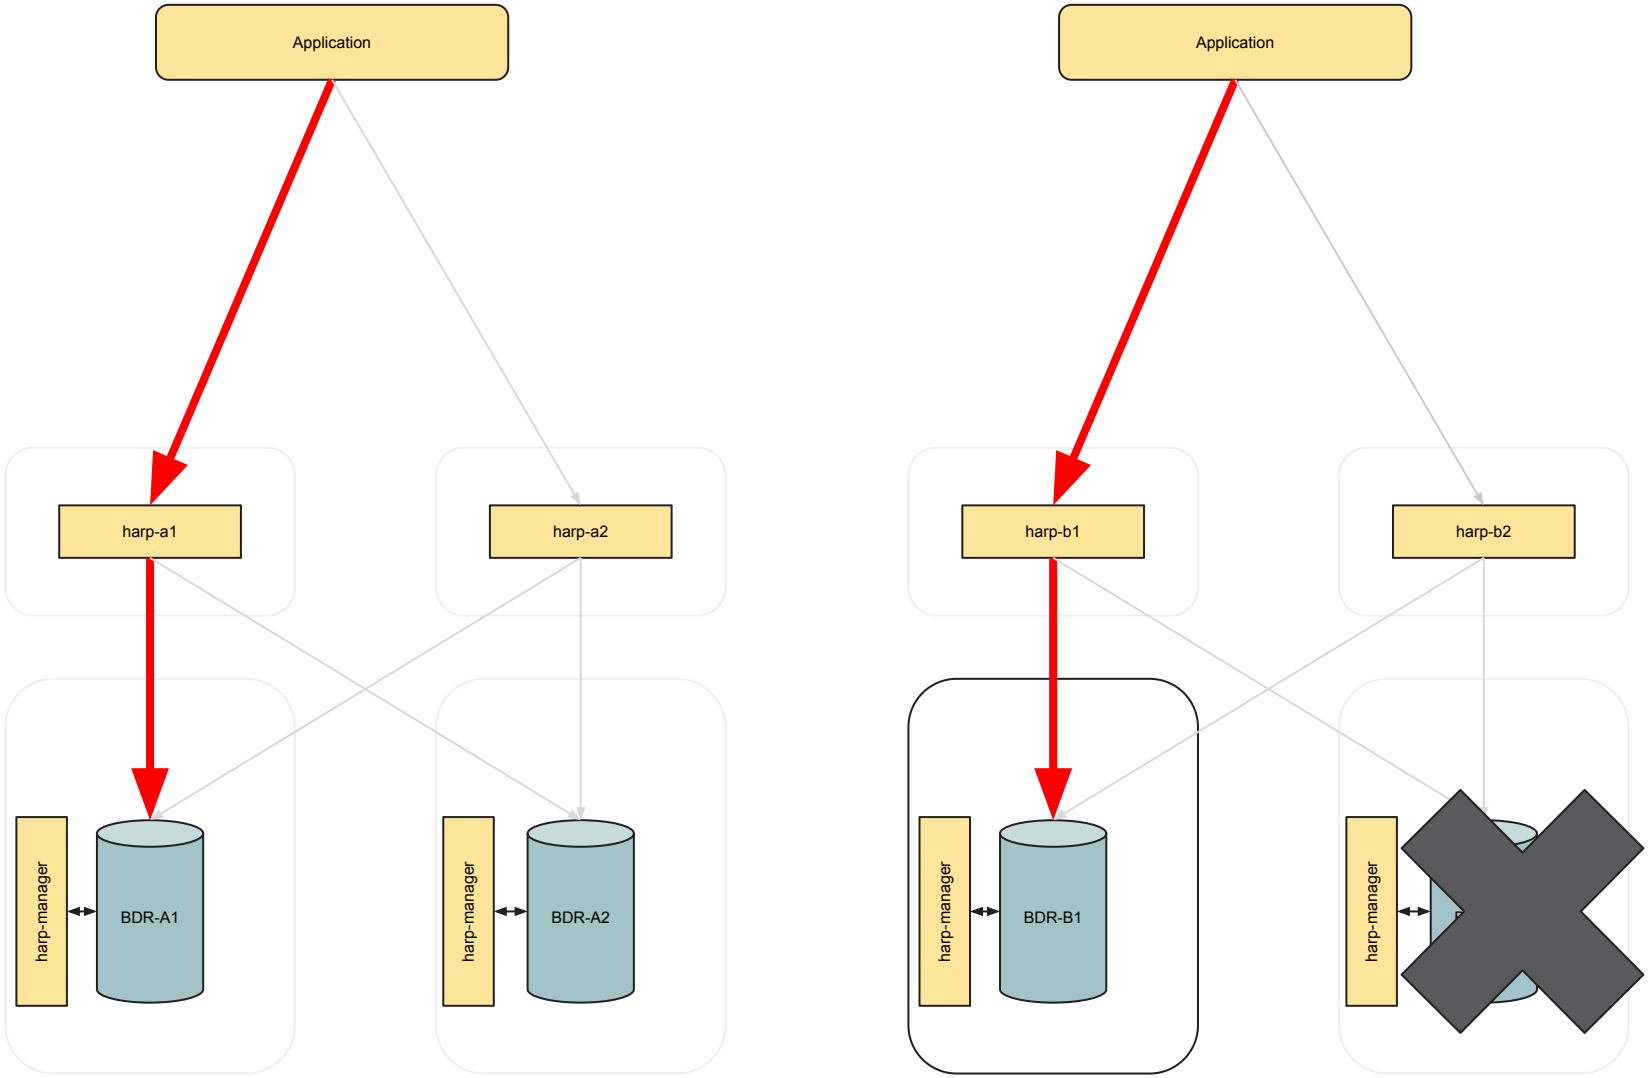

RESTORE BDR-B2

Automatic resync

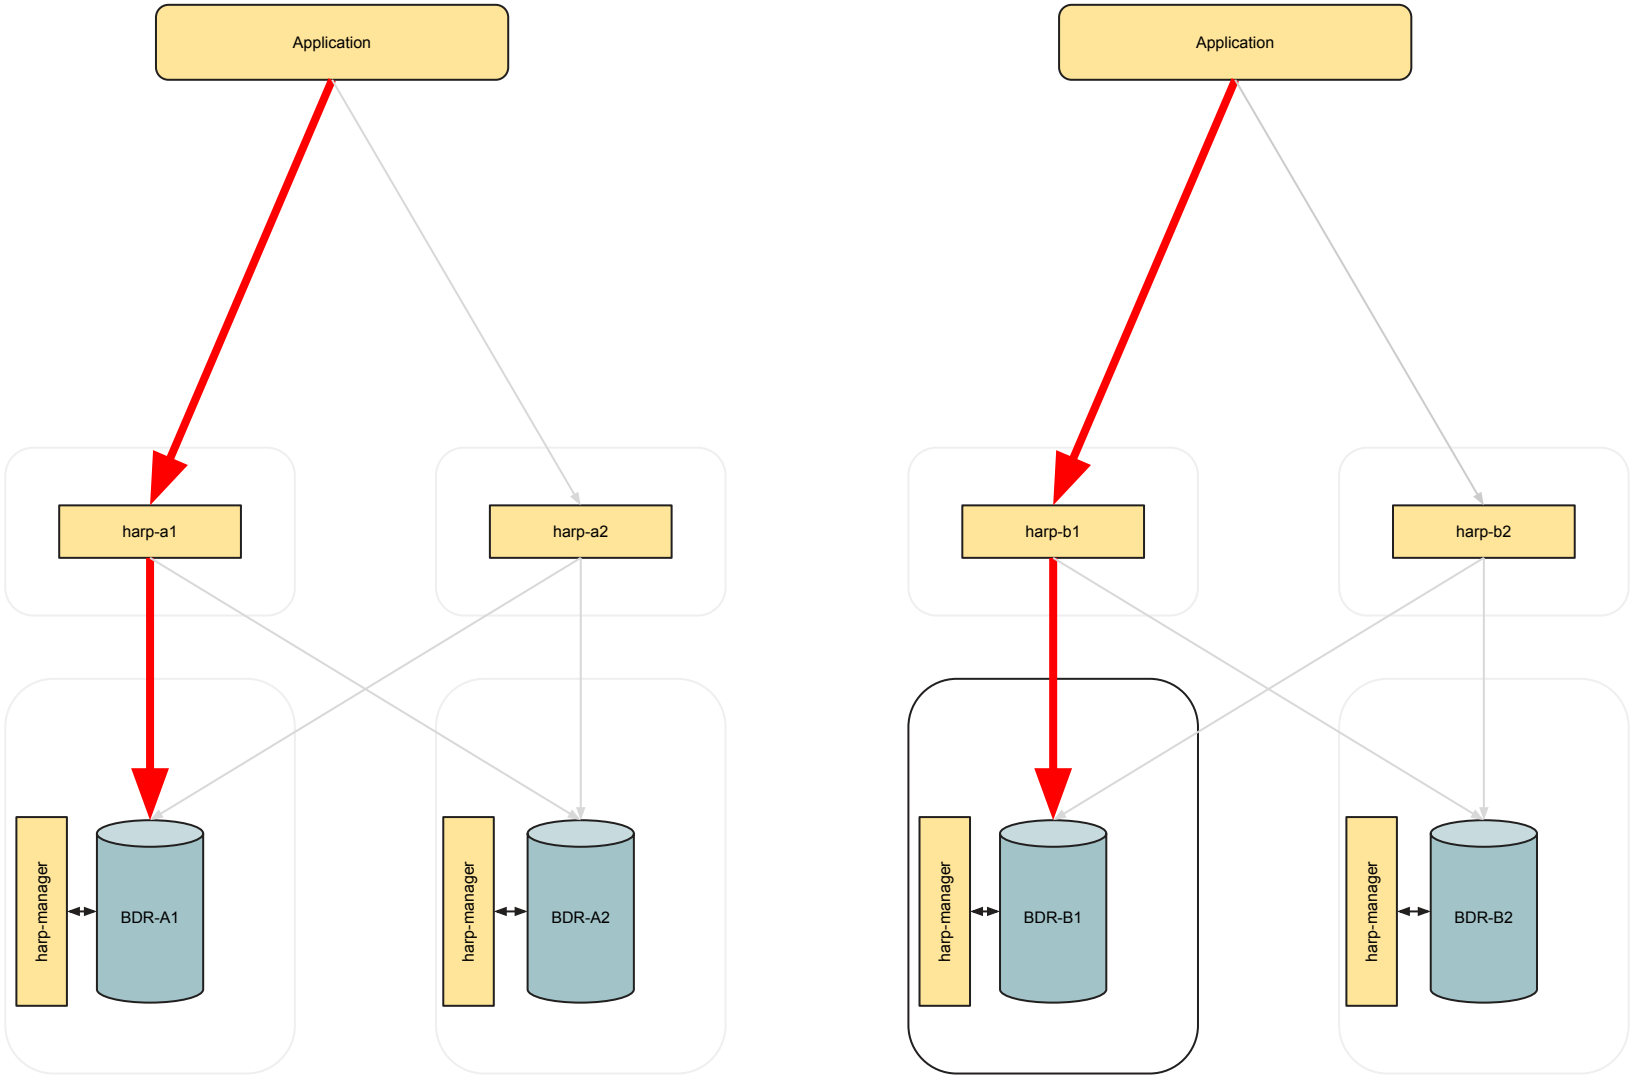

THANK YOU FOR ATTENDING!

Attend our next session, Ask the Expert: Perform at Your Best with Postgres

EDB™