IceWarp

HIGH AVAILABILITY SETUP

WWW.ICEWARP.COM

FAILOVER CLUSTERED SETUP
ACTIVE-PASSIVE WITH DB CLUSTER

1. Frontend load balancer/s

- Can be just one to avoid a single point of failure, there are two in the setup
- For failover send 100% of the traffic to the main server (IW 1)
- Monitors TCP sockets and in general network availability of the machine
- May monitor the availability of the services (command > response)
- In case of failure of IW 1 all traffic is redirected to IW 2
- The application load balancer should support both route and NAT to support SMTP right

Any load balancer capable of the functionality mentioned above is suitable for the setup. The suggested solution for this setup is ***Keepalived for Linux.***

2. IceWarp servers (IW 1, IW 2.. IWn)

- Servers are synchronised by sharing the database and config/mail folders
- Servers are identical and host all services
- Database connection may be set to database load balancer 1
- Backup connection to database load balancer 2 or Keepalived virtual IP used as a target IP for both IceWarp servers which is always assigned to master
- All servers are attached to Storage

3. Database load balancer

- Provides a single point of connection for the application services
- Can be just one to avoid a single point of failure, there are two in the setup
- Receives request from IceWarp server and forwards them to database cluster nodes

The only solution that works as expected is either ***Percona ProxySQL or MariaDB MaxScale****.*

4. Database cluster

- Consists of three (or more) servers in the cluster
- 3 or 5 nodes for single geo.located setups, 2,4 + garbd for multi geo.located setups
- Servers are replicated using SQL server mechanisms (wsrep)
- Optional is DR slave replication to a separate server (requirements are half of the server in a cluster)

We recommend using either ***Percona XtraDB Cluster (PXC) or MariaDB in Galera cluster****.*

5. Storage

- Storage is connected to all servers
- The recommended file system is NFS4, alternatively, you may use FC
- The dedicated network connection for NFS is strongly recommended, of the speed at least 10Gb/s
- Low latency dedicated interconnects like Mellanox RDMA/InfiniBand are strongly recommended (dedicated VLAN in a virtualized environment)

- Storage is replicated to a backup storage
- We do recommend using ZFS snapshots (other mechanisms have significant latency and may cause data loss when storage goes down) (zrepl daemon)
- Some manufacturers such as NetApp have built-in replication in the storages, so the storage provides this
- Recommended architecture is TIER III storage

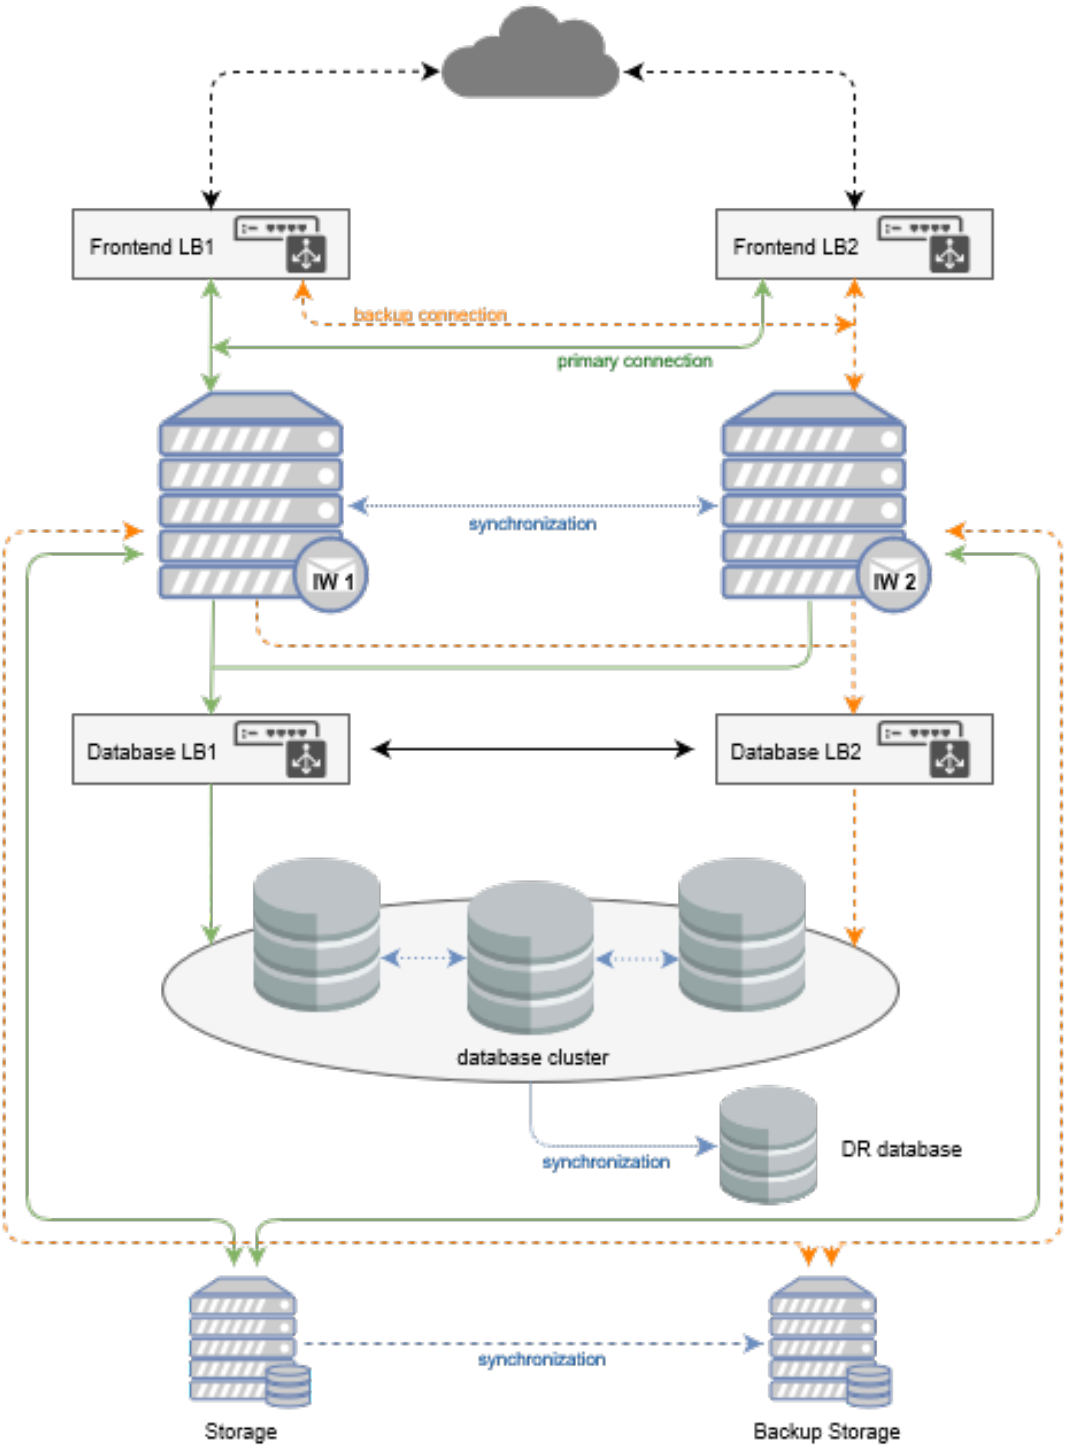

FAILOVER SETUP
ACTIVE-PASSIVE WITH DB REPLICATION

1. Frontend load balancer/s

- Can be just one to avoid a single point of failure, there are two in the setup
- For failover send 100% of the traffic to the main server (IW 1)
- Monitors TCP sockets and in general network availability of the machine
- May monitor the availability of the services (command > response)
- In case of failure of IW 1 all traffic is redirected to IW 2

Any load balancer capable of the functionality mentioned above is suitable for the setup. The suggested solution for this setup is ***Keepalived for Linux.***

2. IceWarp servers (IW 1, IW 2.. IWn)

- Servers are identical and host all services
- The database connection is set to a database that is accessed by one server only
- Both servers are attached to shared storage
- There is always a single access

3. Database servers

- Can run on the same machine or separately
- Servers are replicated using SQL server mechanisms

The only solution that works as expected is ***MariaDB****.*

4. Storage

- Storage is connected to both servers
  The recommended file system is NFS4; alternatively, you may use FC
- The dedicated network connection for NFS is strongly recommended, of the speed at least 10Gb/s
- Storage is single or replicated to a backup storage
- When replicated, we do recommend using ZFS snapshots (other mechanisms have significant latency and may cause data loss when storage goes down).
- Some manufacturers such as NetApp have built-in replication in the storages, so the storage provides this.
- Recommended architecture is TIER III storage.
- When single, this setup requires only one active branch that serves all the users. When the backup branch is taking over, the direction of synchronisation must be changed manually, and main server becomes a spare until next failure.

Replicated storage

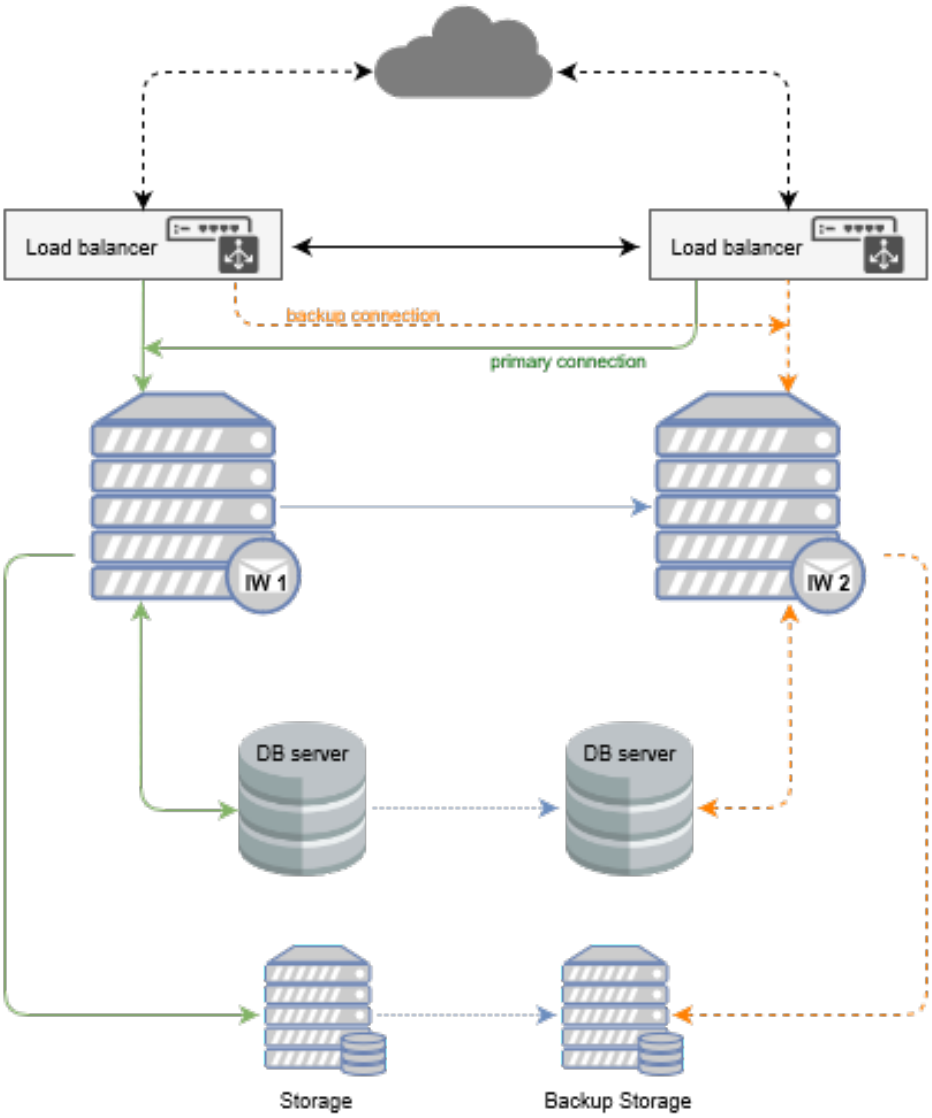

Single storage

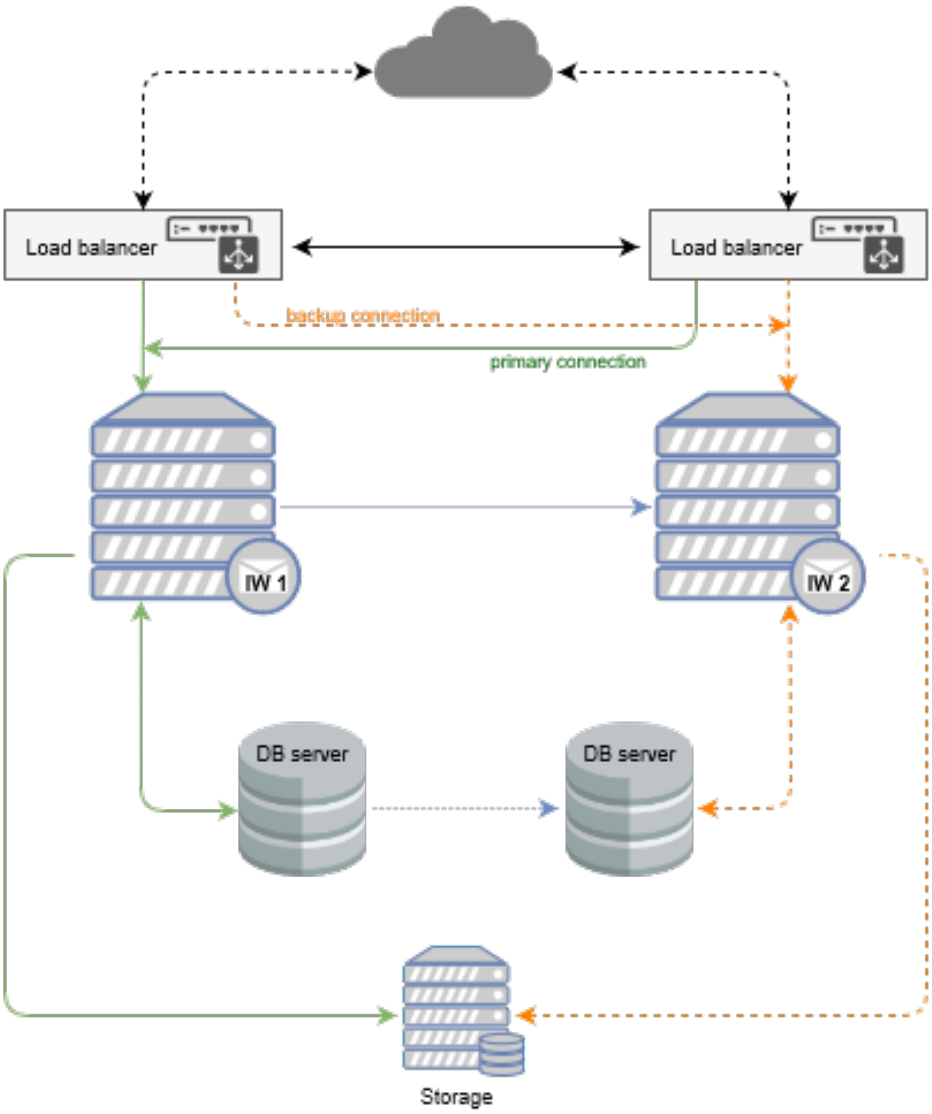

SERVER SIZING OPTIONS

Users	Standard setup	High Availability Setup	
1000	1x Production server (IceWarp, Database, Storage on the same server)	2x Production server (IceWarp) 3x Galera MySQL server (Database) + 2x VM load balancer 2x Storage server (optionally Backup storage server)	Note: Galera can be virtualized Production + SQL servers: 16GB RAM, 4core CPU Storage: less HDD (e.g. 5), 32GB RAM
5000	1x Production server (IceWarp, Storage) 1x MySQL server (Database)	2x Production server (IceWarp) 3x Galera MySQL server (Database) + 2x VM load balancer 2x Storage server (optionally Backup storage server)	Note: Galera can be virtualized Production + SQL servers: 32GB RAM, 4core CPUs Storage: less HDD (e.g. 10), 64GB RAM
10000	1x Production server (IceWarp, Storage) 1x MySQL server (Database)	2x Production server (IceWarp) 3x Galera MySQL server (Database) + 2x VM load balancer 2x Storage server (optionally Backup storage server) 10GbE network	Note: Galera can be virtualized Production + SQL servers: 64GB RAM, 6core CPUs Storage: less HDD (e.g. 15), 128GB RAM
15000	1x Production server (IceWarp, Storage) 1x MySQL server (Database)	2x Production server (IceWarp) 3x Galera MySQL server (Database) + 2x VM load balancer 2x Storage server (optionally Backup storage server) 10GbE network	Production + SQL servers: 64GB RAM, 6core CPUs Storage: less HDD (e.g. 20), 128GB RAM
25000	1x Production server (IceWarp) 1x MySQL server (Database) 1x Storage server	2x Production server (IceWarp) 3x Galera MySQL server (Database) + 2x load balancer 2x Storage server (optionally Backup storage server) 10GbE network	Production + SQL servers: 64GB RAM, 8core CPUs Storage: less HDD (e.g. 25), 128GB RAM
IceWarp Cluster		Nx Production server (IceWarp) 3x Galera MySQL server (Database) + 2x load balancer 2x Storage server 1x Backup storage server 2x Control server 1x Management unit 2x 10GbE Switch	

HW CONFIGURATION EXAMPLE

CONTROL SERVER
1x Intel Xeon E5-2603v4 - 1,7GHz@6,4GT 15MB cache, 6core, 85W,LGA2011 1
4x 16GB 2400MHz DDR4 ECC Registered 2R×4, LP(31mm), Samsung (M393A2G40EB1-CRC) 4
1x 10× serialATA 3.0 6Gb/s RAID 0/1/5/10 1
1x network card 2× GbE
1x Aspeed AST2400 with 16MB VRAM 1
2x Intel SSD DC S3100 Series 240GB SATA3 6Gbps 2,5" 54/4,4kIOPS 5DWPD 2

1x IPMI 2.0 module with KVM-over-LAN
1x SC113MTQ-330C 1U ATX 8SFF,slimCD,330W (80+ GOLD)
1x RR1U-E16 (1U RC) - PCI-E16 (H8SMI/X8STi/X8DTL/H8QM+/X9SCI-LN4F)
1x SNK-P0057PS Passive 1U heatsink pro 2P LGA2011-3 BLADE + X10SRL-F 1

STORAGE SERVER

8x SFF-8643 (SAS-HD) -> SFF-8643 (SAS-HD), 80cm cabel
2x Intel Xeon E5-2620v3 - 2,4GHz@8,0GT 15MB cache, 6core,HT, 85W,LGA2011
1x X10DRH-iT 2S-R3,PCI-E16(g3),6PCI-E8(g3),2×10GbE-T, 8sATA3,16DDR4-2400,IPMI
8x 32GB 2133MHz DDR4 ECC Registered 2R×4, LP(31mm), Samsung
2x LSI SAS9300-8i(3008) SAS3HBA(JBOD) 2×8643,exp:1024HD,PCI-E8 g3,MD2,SGL
36x 8TB Hitachi Ultrastar 7K4000 - 7200rpm, SAS2, 512n, 64MB, 3,5" 16
2x 1TB Seagate Constellation ES.3 -7200rpm, SAS2, 512n, 128MB, 3,5" 2
2x Hitachi Ultrastar SSD800MH.B 400GB SAS3 12Gbps 2,5" 145/100kIOPS, 25DWPD 2
2x Samsung SSD SM863 960GB SATA3 6Gbps 2,5" 97/26kIOPS 3,6DWPD SED
1x ConnectX-3 Pro EN MCX312B-XCCT - Dual Port 10GbE (SFP+), PCI-E8(g3)
1x SC847BE2C-R1K28LPB 4U eATX13 24+12 SAS3 (dual SAS3 exp.),2SFF, rPS 1280W (80+ PLATINUM), LP
2x reduction 3,5"->2,5" for SC813/815/818/825/829/832/836/842/848/932/933/733
1x Additional box for 2×2,5" hotswap do SC826B/SC216B/846X/417B at the back side including backplane)
2x SNK-P0048AP4 Active 2U for 1P/2P LGA2011 (52dBA,8500rpm,4pin) + (CAS BKT-0048L-C32 for H8 boards)

BACKUP STORAGE

same as previous, but
2x SSD Hitachi not necessary here
2x Samsung SSD is not necessary here

PRODUCTION SERVER

2x 4×sATA 23+35+45+56cm cabel (pro 813MTQ)
2x Intel Xeon E5-2620v4 - 2,1GHz@8,0GT 20MB cache, 8core,HT, 85W,LGA2011~
1x X10DRW-E 2S-R3,WIO,PCI-E32(g3), 10sATA3,16DDR4-2400,IPMI,bulk
1x ConnectXŽ-3 Pro EN MCX312B-XCCT - Dual Port 10GbE (SFP+), PCI-E8(g3)
8x 32GB 2400MHz DDR4 ECC Registered 2R×4, LP(31mm), Samsung (M393A4K40BB1-CRC)
2x 750GB WD7500BFCX RED RAID Intellipower, sATA3, 16MB, 2,5" 9,5mm, NAS
4x SSD 2,5" 1TB Samsung 850 Pro SATAIII
1x SC113TQ-600W 1U WIO 8SFF,slimCD,600W (80+ PLATINUM)
1x RSC-R1UW-2E16 - 1U WIO->2×PCI-E16g3
2x SNK-P0057PS Passive 1U heatsink for 2P LGA2011-3 BLADE + X10SRL-F

MYSQL GALERA SERVER

similar setup as production server, but
2x Intel Xeon E5-2683v4 - 2,1GHz@9,6GT 40MB cache, 16core,HT,120W,LGA2011~
4x SSD 2,5" 2TB Samsung 850 Pro SATAIII
16x 32GB 2400MHz DDR4 ECC Registered 2R×4, LP(31mm), Samsung

SWITCH (10GB)

Juniper EX4550-32F-AFO/AFI
with EX4550-VC1-128G module
EX-CBL-VCP-1M cabel
license EX4550-AFL (BGP + IS-IS + MPLS)

MANAGEMENT MODULE

Mikrotik RouterBOARD RB3011UiAS-RM

MYSQL GALERA LOAD BALANCER

2x 1U or virtualized

REQUIRED PORTS

Protocol	Type	Source	Source Port	Destination	Destination Port	Notes
*	Inbound	Any	25	Any	25	SMTP
*	Inbound	Any	587	Any	587	SMTP 2nd
*	Inbound	Any	465	Any	465	SMTP SSL
*	Inbound	Any	21	Any	21	FTP
*	Inbound	Any	990	Any	990	FTP SSL
*	Inbound	Any	80	Any	80	HTTP
*	Inbound	Any	32000	Any	32000	HTTP alternative
*	Inbound	Any	32001	Any	32001	HTTPS alternative
*	Inbound	Any	443	Any	443	HTTPS
*	Inbound	Any	110	Any	110	POP3
*	Inbound	Any	995	Any	995	POP3 SSL
*	Inbound	Any	143	Any	143	IMAP
*	Inbound	Any	993	Any	993	IMAP SSL
*	Inbound	Any	5222	Any	5222	IM
*	Inbound	Any	5223	Any	5223	IM SSL
*	Inbound	Any	5229	Any	5229	GROUPWARE
*	Inbound	Any	5060	Any	5060	SIP
*	Inbound	Any	5061	Any	5061	SIP SSL
UDP	Inbound	Any	10000-10255	Any	10000-10255	SIP STREAMS
*	Inbound	Any	1080	Any	1080	SOCKS
*	Inbound	Any	161	Any	161	SNMP
*	Inbound	Any	389	Any	389	LDAP
*	Inbound	Any	636	Any	636	LDAP SSL
*	Inbound	Any	4069	Any	4069	MINGER
*	Inbound	Any	4070	Any	4070	MINGER SSL
*	Inbound	Any	13	Any	13	TIMESYNC
*	Inbound	Any	53	Any	53	DNS
*	Inbound	Any	3306	Any	3306	MYSQL