
# Overview of Redundant Systems





atlantic.net

Secure Cloud Services
Managed & Compliant Infrastructure

888-618-DATA (3282)
sales@atlantic.net
www.atlantic.net

# Table of Contents

**What Is a Redundant System?** 3
**Types of Redundant Systems** 3
**Examples of Redundant Software Services** 4
Hyper-V Replica 5
Hyper-V Clustering 5
HAProxy 5
Heartbeat 5
**Examples of Redundant Hardware Services** 6
RAID 6
**Networking Redundancy** 8
First Hop Redundancy Protocols (FHRP) 8
Virtual Router Redundancy Protocol (VRRP) 8
Hot Standby Router Protocol (HSRP) 8
Gateway Load Balancing Protocol (GLBP) 9
**Data Center Redundancy** 9
**Sources** 10

The purpose of this white paper is to explain redundancy in terms of computing, networking, and hosting. We will provide real-world examples of redundant technology solutions to illustrate what redundancy is and how it works.

Atlantic.Net has created multiple hosting environments, including a durable cloud platform, high-speed VPS hosting, HIPAA compliant infrastructure, and managed private cloud hosting. All of our systems are built with redundancy as a primary driving factor of the design process.

In everyday English, redundancy may have a negative connotation; something redundant is usually not needed or considered superfluous. However, in a cloud hosting environment, redundancy can mean the difference between seamless system availability and unwanted or unexpected downtime.

## What Is a Redundant System?

A redundant system will provide failover or load balancing support to protect a live system in the event of an unexpected failure. In the case of power, mechanical, or software failure, a redundant system will have a duplicate component or platform to fall back to. In general, any component of a system with a single point of failure can be seen as a risk to production services.

Power or mechanical systems have simpler fall back strategies requiring the mere presence of another of the same type of service; software failovers usually require extra configuration on the host system or a master or gateway.

Redundancy capabilities are recommended for any business-critical system, but particularly for systems that have a significant impact during downtime. Some businesses may keep all of their critical customer information in a database; therefore, for business continuity purposes, protecting that database with redundancy will protect the data integrity in the event of a catastrophic failure.

## Types of Redundant Systems

A redundant system consists of at least two systems that are interconnected and designed for the same purpose. There are many different types of redundant system configurations available, and different implementations of the system provide unique approaches to how to keep a system up at all times.

Not all servers need to be configured with redundancy; rather, only the most critical should be considered. We highly recommended detailed risk assessment to understand what servers are in scope and the maximum amount of downtime your servers can handle. Use this assessment to determine an RTO (Recovery Time Objective) and

RPO (Recovery Point Objective) strategy. RTO is the maximum amount of acceptable downtime. This can range from 5 seconds to 24 hours. The RPO is the point in time from which you require your data; for example, your business can function with a maximum loss of 24 hours worth of data.

Here are a few popular examples[1]:

- **Active-Inactive/Hot-Cold** – When one component of a system is the active system and another is inactive or shut down. The inactive component is activated only when the currently running component fails or undergoes maintenance.
- **Active-Active/Hot-Hot** – When both systems are live and making connections. This is most commonly known as clustering. Usually, the device in front of both machines will determine how to split incoming traffic.
- **Active-Standby/Hot-Warm** – When both systems are on, but only one is making connections. The second system is meant to periodically receive updates or backups from the primary system. In the event of a failure, the system on standby takes the primary role until the initial system can be recovered.

Each type has its own pros and cons.

- **Active-Inactive/Hot-Cold** systems can provide a simple redundant platform, but any failover will result in users seeing an older version of the system.
- **Active-Active/Hot-Hot** will require a constant update of both systems, either manually or through a separate service, to ensure that all users can use either system. This approach can heavily reduce the active load on a service you're providing to customers.
- **Active-Standby/Hot-Warm** will provide the failover capabilities of hot-cold with a more up-to-date copy of your active system on the failover, but it does not provide any load easing.

Other forms of multiple node redundancy are available that allow for greater redundancy and robust load balancing solutions. At that point, you'll have a high-availability cluster, also known as an HA cluster.

This can use any combination of the previously noted redundancy solutions with maximum flexibility in the approach or amount of redundancy needed. HA clusters can also be set up across multiple physical locations to allow for availability up to the internet backbone level.[2]

## Examples of Redundant Software Services

Short of low resource availability, there is very little reason to not have proprietary replication or redundant services set up in a

virtual environment; thus, many such services are available by default in most virtualization systems. All of our cloud services have replication available, a feature that allows us to replicate any server from one node to another, whether they are in the same data center or separate data center regions.

## Hyper-V Replica

Hyper-V Replica is a form of hot-warm redundancy. A primary virtual machine is created on one physical host and accepts incoming connections. When enabling replication, the virtual hard disks of the new machine are transferred to a separate physical Hyper-V host. This host then configures a VM on itself that replicates on a user-defined schedule to ensure that the most recent image of the active server is taken. Additional checkpoints points can be kept as well.[3] Hyper-V private hosting with managed services is provided by Atlantic.Net with this feature baked in; contact our team for further information.

## Hyper-V Clustering

Hyper-V is also capable of clustering through a connection to other Hyper-V hosts. VMs on any Hyper-V host can be clustered together on that singular host to provide redundancy on a local level through virtual networking.

Microsoft Network Load Balancing (NLB) can be used to create a single resource made up of multiple hosts that share the same information to provide a simple point of access for file sharing. Since this is only capped by the amount of resources you have available, you can theoretically set up multiple hosts with multiple VMs for maximum redundancy, which would also allow you to perform maintenance on individual VMs without sacrificing service or resource availability.[4] Hyper-V private hosting with managed services is provided by Atlantic.Net with this feature baked in; contact our team for further information.

## HAProxy

Aside from Hyper-V, a gateway device such as a firewall can be used for failover or load balancing services. For example, Atlantic.Net can provide pfSense with High Availability Proxy, also known as HAProxy.

HAProxy will act as a load balancer, a proxy, or a simple hot-warm high availability solution for TCP and HTTP-based applications.[5] HAProxy is a very popular, Linux-based open-source solution used by some of the most visited sites in the world.[6]

## Heartbeat

Heartbeat is a service available on most distributions of Linux that is used to determine whether nodes in a cluster are still up or responsive. It's very simple to set up and provides failover capabilities to any system

working over TCP.

The developers of Heartbeat also recommend other cluster resource managers which start or stop services based on whether a particular host is down. Heartbeat has this included, but other managers are available. Due to Heartbeat's simplicity, it is highly customizable.[7] Cloud Hosting platforms provided by Atlantic.Net already have this feature baked in, and we can assist you with implementing Heartbeat on your own private Linux distribution, if needed.

## Examples of Redundant Hardware Services

The best part about redundant hardware is its simplicity. While software services may require excessive configuration and are possibly quite sensitive, the hardware is usually very simple to set up and incredibly durable. The first example we will look at is the widely used RAID technology.

### RAID

RAID stands for **R**edundant **A**rray of **I**ndependent **D**isks (or **R**edundant **A**rray of **I**nexpensive **D**isks depending on how long you've been using it) and has multiple levels used either for data protection or increased disk I/O.

RAID can either be set up via a software or hardware controller. The controller has the software and configuration necessary to manage the RAID disks. The configuration can be exported to different systems with little to no additional configuration.

RAID can be set up in a few different ways to provide a good balance of both of its qualities:

- **RAID 0** – This is essentially no redundancy. No disks on the system share data through mirroring, but all data is striped across each disk providing increased read/write speed. Each drive can still use the storage provided to it at its fullest, meaning the more drives you add to a RAID 0 the more space you'll have.

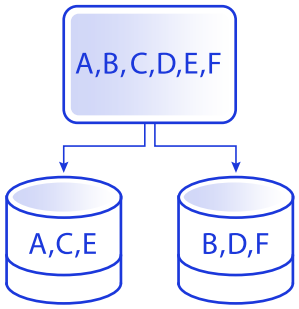


- **RAID 1** – A basic form of mirroring providing excellent redundancy at the cost of space. In a two-drive system, a complete copy of the data on one drive is written to the other. This redundancy is enhanced with each drive added. Since all data must be mirrored across all drives, total space on the system will be limited to just the space of the smallest drive in the system.

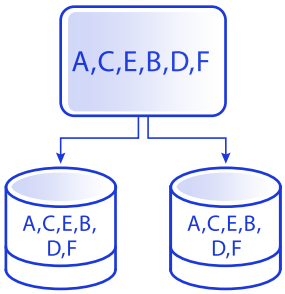


- **RAID 5 –** This form of RAID is usually used to increase read speed and reliability. In this case, stripes are placed about each drive in the system, with the minimum being 3 drives. At the same time, an extra block of error-correcting data is placed about each drive in a technique called parity. This checks whether data is changed when transferring from one drive to another.[8] This also provides a minimal form of redundancy since 1 of these drives can fail and the system can still run. The more drives added to this type of RAID setup, the more your read speed increases. With minimum redundancy and striping across all drives, the total amount of space in this setup is equal to the size of your logical RAID volume times the number of drives you use, minus one. For example, if you have 5 500 GB drives in a RAID 5, you would have 2000 GB usable, or 2 TB (500 *(5-1)=2000).

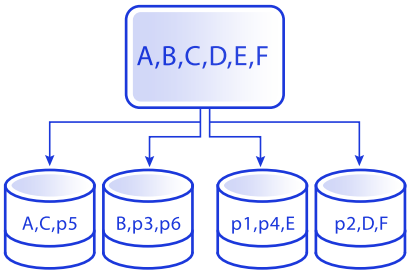


- **RAID 10** – This is a combination of RAID 1 and RAID 0. In this case, all data is striped across each device with blocks of data also being mirrored across the entirety of the striped system. For example, in a 4 drive RAID 10 system 2 500 GB drives may have the same data, but not all the data needed for the system to work properly. 2 other drives' data would be required. Think of each RAID 1 system as a single drive, and each of those systems placed into a RAID 0 array. In this setup, performance can be drastically increased as in RAID 0, with some redundancy still in place with the mirroring. Up to half of the drives in the system can fail before the system crashes, but as with any redundant array, it's best to replace drives as soon as possible. Atlantic.Net uses RAID 10 for all SSD Cloud VPS Storage.

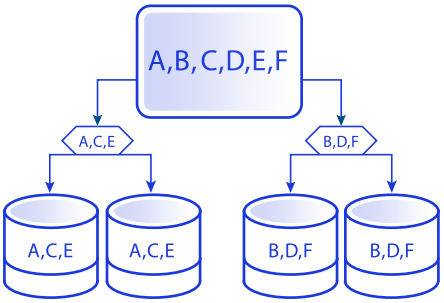


For added protection, the RAID controllers are protected by battery backup units that power the ROM chips used to save the configuration in memory in case of power loss, etc. A BBU will provide power to a RAID array that's part of a powered down system for a small amount of time, allowing the content of a RAID controller's cache to

stay intact. This can be a lifesaver if the information is constantly being fed into your RAID array and any downtime could cause data corruption.[9]

So, your physical system and the services within can be constructed redundantly rather adequately. But what about your connection to any part of your system? As in, your direct internet connection to your system as a whole?

# Networking Redundancy

## First Hop Redundancy Protocols (FHRP)

In contrast to dynamic gateway discovery protocols, static gateways allow for straightforward hops between the client and their appropriate gateway, but this creates a single point of failure – namely the gateway itself.

To prevent or reduce the impact of gateway failure, FHRPs were created. They provide redundant gateways a fallback, or offer load-balancing for high traffic systems, along with redundancy. These protocols include VRRP, HSRP, and GLBP.[10]

## Virtual Router Redundancy Protocol (VRRP)

VRRP is a form of redundancy used for routers that requires at least two physically separate routers connected via either Ethernet or optical fiber connections. In this situation, a 'virtual router' containing static routes is created and shared between each system.

One system is considered the 'master' and another the 'backup'. When the master fails, the backup takes over as the next master. This can be set up with multiple backups for extra redundancy. The concept is very similar to Heartbeat in that the backup systems will check to see if the master is available. Once it does not receive a response, after a predetermined amount of time the backup will assume control of the virtual switch and accept connections for all requests coming in for the default IP configured for the master switch.[11] [12]

## Hot Standby Router Protocol (HSRP)

HSRP is like VRRP; however, in this scenario, the configured virtual switch isn't a 'switch', but rather a logical group of multiple routers. The IP of the group is an IP not assigned to a physical host. Instead, the group is assigned an IP and one of the routers is determined to be the 'active' router.

A standby router is ready to take any connections should the active router go down. All routers besides the active and standby are all listening to determine its place in line. HSRP is a Cisco proprietary protocol and has very few, minor differences to VRRP such as their default timers determining when to failover. HSRP has been around a bit longer and is more well-known compared to VRRP.[12]

## Gateway Load Balancing Protocol (GLBP)

GLBP's main advantage over HSRP and VRRP is its ability to load balance on top of providing redundancy to a gateway with little to no extra configuration. Much like HSRP and VRRP, GLBP will create a group between physical routers and determine an Active Virtual Gateway, or AVG.

A virtual IP not currently used by any of the routers in the group is assigned to the AVG. The AVG then distributes virtual MAC addresses among the rest of the routers in the group. Each backup router is now considered an Active Virtual Forwarder, or AVF.

ARP requests sent to the AVG will provide a different virtual MAC address to the client sending the request. At that point, traffic from that client to the virtual IP of the group forwards to the router whose virtual MAC address they received, allowing each router to still be used instead of sitting idly by.

In the event of a failure of the AVG, priority-based election takes place, just like in HSRP and VRRP, and the next backup takes its place, distributing virtual MAC addresses as normal. The other routers still retain the virtual MAC address provided by the original AVG and things continue as normal. In the event of a failure of one of the AVFs, the AVG will prevent routing traffic to its virtual MAC address.

Just like HSRP, GLBP is a Cisco proprietary form of FHRP.

## Data Center Redundancy

In addition to redundancy measures for your personal servers or routers, data centers are designed to be resilient to system failure. Data centers fall under tiers defined by the Uptime Institute to provide fault tolerance for failure of any mechanical or service failure, allowing for as much uptime as possible.

There are four tiers, each building upon one another to provide high availability to all clients within a data center:

- **Tier I – Basic Capacity:** This requires space for an IT group for data center operations, an uninterruptible power supply (UPS) that monitors and filters power usage and dedicated cooling equipment that is constantly running 24/7. This also includes a power generator in the case of electrical power failure.
- **Tier II – Redundant Capacity Components:** Everything that Tier I provides, plus redundant power and cooling to the facility. This can include extra UPS units or extra generators.
- **Tier III – Concurrently Maintainable:** Everything Tier II provides, plus extra equipment in place to prevent any need for shutdowns for equipment replacement or maintenance. At this tier,

redundant power and cooling are applied directly to all technical equipment, and the equipment itself is configured for redundancy or seamless failover.

- **Tier IV – Fault Tolerance:** Everything Tier III provides, plus uninterrupted service at the provider level. While a data center may have electricity or water provided by a city or state provider, a secondary line of each service utilized by the data center is required. This includes the ISP as well. In the event of a failure at any section leading up to client equipment, there is a backup plan in place ready for a seamless transition.

## Sources


[1] *Cold/Warm/Hot Server: http://searchwindowsserver.techtarget.com/definition/cold-warm-hot-server*

[2] *High Availability Clustering: https://www.mulesoft.com/resources/esb/high-availability-cluster*

[3] *Hyper-V Replica: https://technet.microsoft.com/en-us/library/jj134172(v=ws.11).aspx*

[4] *Hyper-V and High Availability: https://technet.microsoft.com/en-us/library/hh127064.aspx*

[5] *HAProxy Description: http://www.haproxy.org/#desc*

[6] *HAProxy – They use it!: http://www.haproxy.org/they-use-it.html*

[7] *Heartbeat: http://www.linux-ha.org/wiki/Main_Page*

[8] *RAID Definition: http://searchstorage.techtarget.com/definition/RAID*

[9] *RAID Battery Backup Units: https://www.thomas-krenn.com/en/wiki/Battery_Backup_Unit_(BBU/BBM)_Maintenance_for_RAID_Controllers*

[10] *High-Availability – VRRP, HSRP, GLBP: http://www.freeccnastudyguide.com/study-guides/ccna/ch14/vrrp-hsrp-glbp/*

[11] *Understanding VRRP: http://www.juniper.net/techpubs/en_US/junos/topics/concept/vrrp-overview-ha.html*

[12] *Configuring VRRP: http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipap-p_fhrp/configuration/15-mt/fhp-15-mt-book/fhp-vrrp.html*



**Ready to get started with securing your cloud services?**

Contact Atlantic.Net's cloud security experts today about how we can help secure your cloud infrastructure! Reach out to our sales team at **888-618-DATA (3282)** or email us at **sales@atlantic.net**!