

Redundancy Design Best Practices

Revision 1.1
June 13, 2005



Aruba Networks
1322 Crossman Ave
Sunnyvale, CA 94089
+1 408 227 4500
<http://www.arubanetworks.com>

Table of Contents

Introduction	3
Relevant concepts	3
Master-local switches/controllers	3
LMS-IP/BACKUP-LMS-IP	3
VRRP	3
Redundancy solutions	3
1:1 LMS redundancy	3
Active-Standby model	4
Active-Active model	6
1:1 master switch redundancy	10
1:N/N:1 switch redundancy	13
Standby switch backing up multiple local switch(es)	13
Active -Active N to 1 Redundancy	15
Summary	17

Introduction

This document covers some of the basic concepts and requirements of inter-switch redundancy in a network using the Aruba solution.

Relevant concepts

Master-local switches/controllers

The Aruba network consists of one or more Aruba mobility controllers - each controller managing a set of access points. A GRE tunnel is established between each access point and the Aruba controller managing the access point. The controller then provides switching, encryption/decryption, authentication, and access control functionality to the clients associated with the access point.

In a network consisting of multiple controllers, one of the controllers is denoted the "Aruba Master". This Master controller provides a central point of all configurations such as WLAN and security parameters. It also provides the ability to manage and troubleshoot the network from a central location. All other controllers are referred to as "Local mobility controllers" or "Local mobility switches (LMS)".

LMS-IP/BACKUP-LMS-IP

As described above, each access point is managed by an Aruba mobility controller/switch. This switch is then called the "LMS" (Local Mobility Switch) for this access point and the IP address used by the access point to connect to is referred to as the "LMS-IP".

It is also possible to specify the IP address of a different switch that the access point can connect to if it is unable to connect (or loses its connection) to the "LMS". This IP address is referred to as the "Backup-LMS-IP".

VRRP

The Virtual Router Redundancy Protocol (VRRP) is designed to eliminate the single point of failure inherent in switching systems in the static default routed environment. VRRP specifies an election protocol that dynamically assigns responsibility for a virtual router to one of the VRRP routers on a LAN. The VRRP router controlling the IP addresses associated with a virtual router is called the Master, and forwards packets sent to these IP addresses. The election process provides dynamic fail-over to a redundant switch/controller for the forwarding responsibility should the Master become unavailable.

The Aruba solution uses VRRP protocol to elect a "VRRP Master" between a pair of switches. The VRRP Master assumes control of the virtual IP address and this IP address is used by the respective access points and local switches/controllers to connect to the "currently active switch in the redundant pair". This is described in greater detail in the sections below for the different redundancy scenarios.

Redundancy solutions

1:1 LMS redundancy

In this solution a single Aruba switch/controller backs up another Aruba switch. In other words, when an Aruba switch fails (or is disconnected from the network) the access points that are connected to this switch fail over to the second Aruba switch (backup switch). This is illustrated in the figure below.




Fig. 1 Local switch redundancy

There are two options of deploying a pair of redundant switches for local switch redundancy - active-standby and active-active.

Active-Standby model

In this model the backup switch does not serve any access points or clients unless there is a failure.

An example of a topology that uses this model of redundancy is shown below.




Fig. 2 Example topology for active-standby local redundancy

The following is the configuration needed for the above topology.

Local 1 configuration

```
vlan 2
interface vlan 2
    ip address 10.200.168.2 255.255.255.0
!
interface gigabitethernet 2/24
    switchport access vlan 2
    switchport mode access
    trusted
!
vrrp 2
    vlan 2
    ip address 10.200.168.254
    no shutdown
!
```

Local 2 configuration

```
vlan 2
interface vlan 2
    ip address 10.200.168.3 255.255.255.0
!
interface gigabitethernet 2/24
    switchport access vlan 2
    switchport mode access
    trusted
!
vrrp 2
    vlan 2
    ip address 10.200.168.254
    no shutdown
!
```

Master configuration

```
ap location 1.0.0
    lms-ip 10.200.168.254
!
```

How it works

In the model shown above, the pair of redundant switches (local 1 and local 2 in the example) is connected over a layer 2 network on at least one of the VLANs. VRRP is configured to run on this VLAN and elects a “VRRP Master” switch between the two switches. This elected VRRP Master switch assumes control over the virtual IP address. All APs with location codes that indicate that they are located in “building 1” (location codes 1.x.y) are configured to connect to the same virtual IP address.

If local 1 is elected the “VRRP Master” in VLAN 2, all the access points form their GRE tunnels with this switch. In this scenario all access points and clients are now handled by local 1. In the event of “local 1” failing or losing its connectivity, VRRP elects local 2 as the new “Master” and local 2 assumes control over the virtual IP address. The access points attempt to connect to the same IP address again and now form the GRE tunnels with local 2. Thus all access points fail over to the backup switch. On recovery, it becomes the “Backup” by following the VRRP procedure and is now in the “standby” mode.

Active-Active model

In this model both the Aruba switches are serving access points and clients in the normal mode of operation. Each switch acts as a backup for the access points and clients on the other switch. This places some restrictions on the load that can be placed on each switch (when both are active) to ensure that each of the switches can still serve the total number of access points and users in a failure scenario.

An example of a topology that uses this model of redundancy is shown below.

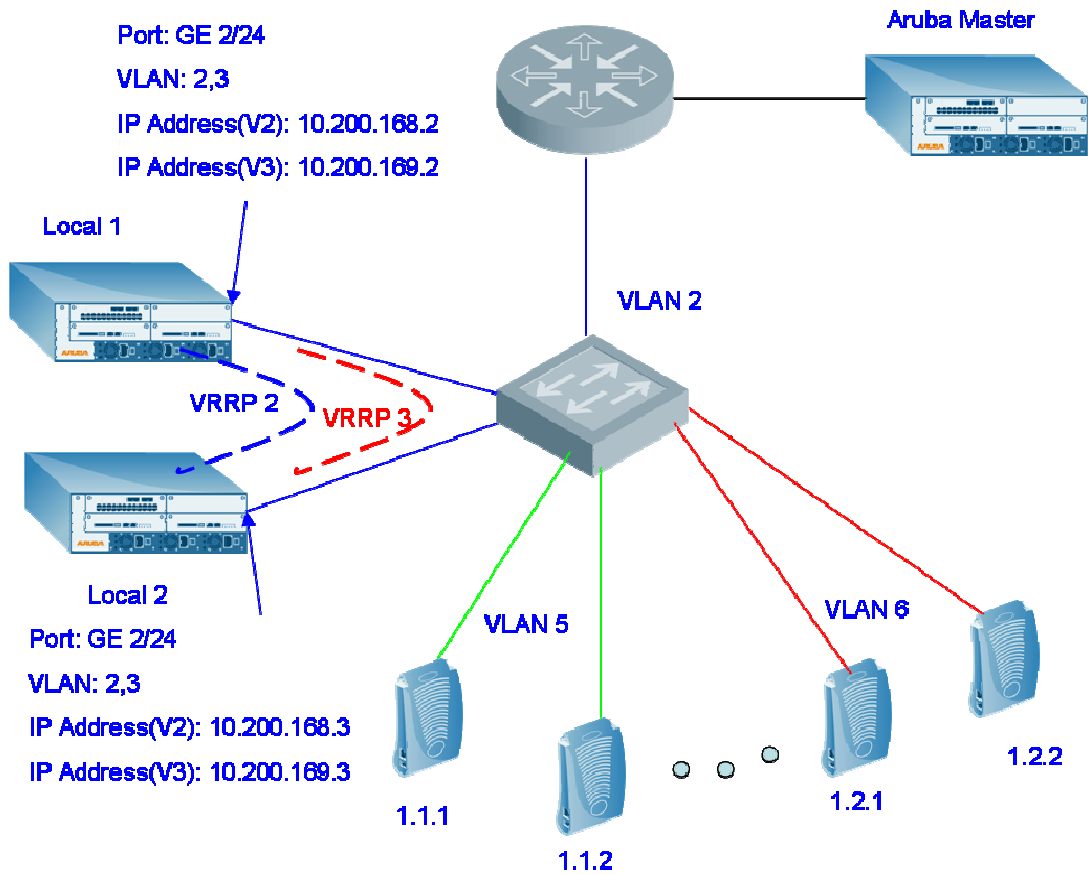


Fig. 3 Example topology for active-active local redundancy

The following is the configuration needed for the above topology.

Local 1 configuration

```
vlan 2
interface vlan 2
    ip address 10.200.168.2 255.255.255.0
!

vlan 3
interface vlan 3
    ip address 10.200.169.2 255.255.255.0
!

interface gigabitethernet 2/24
    switchport trunk allowed vlan 2,3
    switchport mode trunk
    trusted
!

vrrp 2
    vlan 2
    ip address 10.200.168.254
    priority 110
    preempt
    no shutdown
!

vrrp 3
    vlan 3
    ip address 10.200.169.254
    priority 100
    preempt
    no shutdown
!
```


Local 2 configuration

```
vlan 2
interface vlan 2
    ip address 10.200.168.3 255.255.255.0
!

vlan 3
interface vlan 3
    ip address 10.200.169.3 255.255.255.0
!

interface gigabitethernet 2/24
    switchport trunk allowed vlan 2,3
    switchport mode trunk
    trusted
!

vrrp 2
    vlan 2
    ip address 10.200.168.254
    priority 100
    preempt
    no shutdown
!

vrrp 3
    vlan 3
    ip address 10.200.169.254
    priority 110
    preempt
    no shutdown
!
```

Master configuration

```
ap location 1.1.0
    lms-ip 10.200.168.254
!
ap location 1.2.0
    lms-ip 10.200.169.254
!
```

How it works

In the above topology, each of the local switches has two VLANs (2 and 3 in the example shown above) and VRRP running on each of these VLANs. By giving each switch a higher priority in one of the VLANs respectively (and by enabling preemption), the administrator can configure the system to always elect the same switch as the “VRRP Master” for each instance when both switches are up and running.

When both switches are up, local 1 becomes the “VRRP Master” for VLAN 2 and thus assumes control of the virtual IP 10.200.168.254 while local 2 becomes the “VRRP Master” for VLAN 3 and thus assumes control for the virtual IP 10.200.169.254. On the “Aruba Master” switch, the APs with location codes 1.1.x are configured to connect to the IP address 10.200.168.254 while the APs with location codes 1.2.x are configured to connect to the IP address 10.200.169.254. Thus when both the switches are up, all APs 1.1.x connect to local 1 while all APs 1.2.x connect to local 2.

In the event of a failure of a switch, the VRRP election process causes the other switch to become the VRRP master for both the VLANs. For example, if local 1 fails, local 2 becomes the VRRP master for VLAN 2 and takes control of the virtual IP address 10.200.168.254. Thus all APs 1.1.x now fail over to local 2. When local 1 regains connectivity, it preempts local 2 and becomes the VRRP master again, thus causing the APs 1.1.x to return to local 1.

1:1 master switch redundancy

In an Aruba network, one of the switches is configured to be a “master” switch. This switch acts as a central controller and configuration point. It is important to note that the master switch does NOT act as a “single point of failure” for the network. In the absence of the master switch, the network can still serve the clients as before. The only functionality that is lost is the ability to change global configuration such as security parameters and monitor the Aruba network.

The master is used to make global configuration changes and keep a view of the global network such as global user events, RF events, and IDS events. If required, an Aruba switch can be configured to back up the master switch and take over the responsibilities of the master switch in the event of a failure. This is referred to as “Master redundancy” in an Aruba network.

The example below shows a pair of Aruba switches configured a pair of redundant master switches.




Fig. 4 Master Redundancy Topology

The configuration for the above shown topology is described below.

Master configuration

```
vlan 10
interface vlan 10
    ip address 10.200.175.1 255.255.255.0
!

interface gigabitethernet 2/24
    switchport access vlan 10
    switchport mode access
    trusted
!
vrrp 10
    vlan 10
    ip address 10.200.175.254
    priority 110
    preempt
    tracking master-up-time 30 add 20
    no shutdown
!
master-redundancy
    master-vrrp 10
    peer-ip-address 10.200.175.2
!
database synchronize period 60
```

Backup Master configuration

```
vlan 10
interface vlan 10
    ip address 10.200.175.2 255.255.255.0
!

interface gigabitethernet 2/24
    switchport access vlan 10
    switchport mode access
    trusted
!
vrrp 10
    vlan 10
    ip address 10.200.175.254
    priority 110
    preempt
    tracking master-up-time 30 add 20
    no shutdown
!
master-redundancy
    master-vrrp 10
    peer-ip-address 10.200.175.1
!
database synchronize period 60
```

Local configuration

masterip 10.200.175.254

How it works

In the topology shown above, VRRP is configured on a VLAN between a pair of Aruba switches. By configuring a higher priority for one of the switches (the “preferred master”) and enabling preemption, the administrator can ensure that this switch is always elected the master switch when both switches are up and in normal operation. The local switches now point to the virtual IP on this VRRP instance as their master IP address; thereby ensuring that the master functionality is always present on the network as long as one of the two switches is alive.

VRRP tracking

The above example adds a new factor in the configuration - VRRP tracking. The VRRP tracking mechanism provides the ability to “track” a particular variable and adjust the priority accordingly. In this case the variable being tracked is the “duration that the switch has been Master for this VRRP instance”. Therefore the command “**tracking master-up-time 30 add 20**” means that a priority of 20 should be added to the configured priority if the switch stays as the Master for this VRRP instance for 20 minutes or more. The Aruba Master switch maintains state about the network and various events occurring in the network. These can include RF parameters such as channel and power allocation for the access points, IDS security events such as interfering/rogue APs detected/classified, security vulnerabilities observed and other wireless intrusions detected. This state is synchronized between the pair of Aruba master switches periodically. When a master switch fails, the backup switch might not have this database updated. Thus depending on how long the original master switch takes to recover, it might or might not be useful to revert to the original master switch. This flexibility is provided by using the tracking mechanism described above.

Master-redundancy configuration

In the master-redundancy configuration, the administrator needs to specify the VRRP instance that will be used to elect the Master for the Aruba network. Also the IP address of the peer switch for this VRRP instance is specified. This IP address can be the loopback IP address (or the switch IP) or the interface IP address.

Database Synchronization

The command “**database synchronize period 60**” determines how often the two switches synchronize their databases. It is recommended for system performance that this operation should not be performed at a frequency greater than every 30 minutes.

1:N/N:1 switch redundancy

There are scenarios where it is more suitable to have a single switch backing up multiple switches or having a set of switches backing up a single switch. The following sections describe redundancy models that address these scenarios. In addition, these models are also extremely useful in scenarios where it is not possible to place all redundant pair of switches in a Layer 2 network (a restriction that is present in the previously mentioned models due to the VRRP protocol constraints).

Standby switch backing up multiple local switch(es)

This model is illustrated by the following figure.

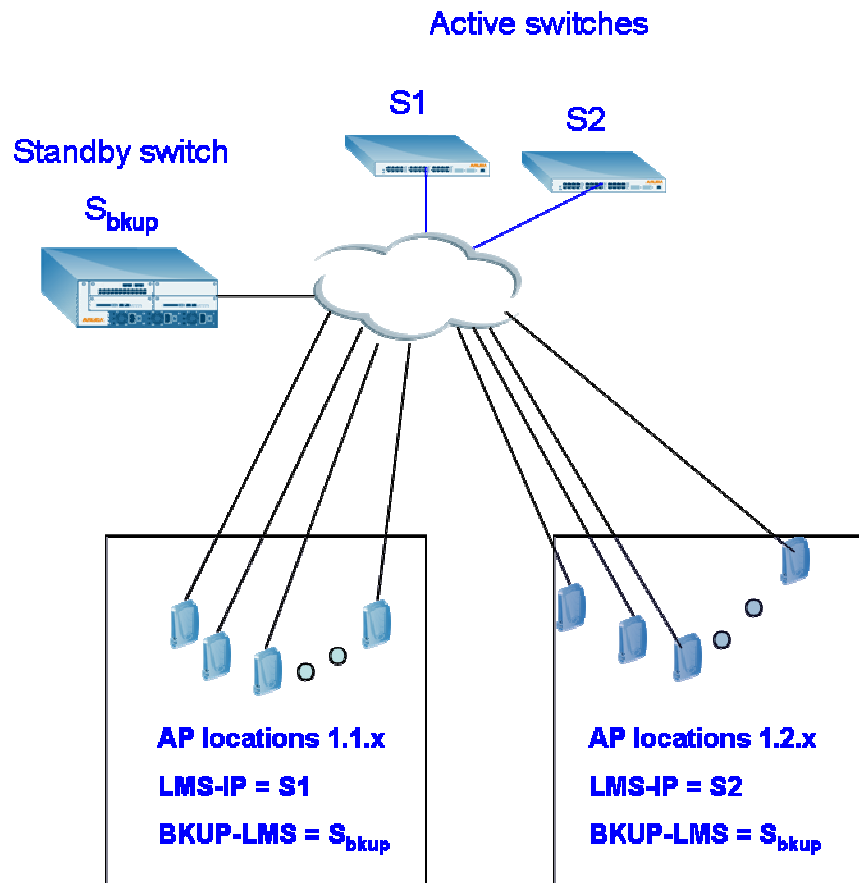


Fig. 5 Standby-switch to backup multiple active switches

The configuration for such a topology implementation is as follows:

Master configuration

```

ap location 1.1.0
  lms-ip <S1>
  bkplms-ip < $S_{bkup}$ >
!
ap location 1.2.0
  lms-ip <S2>
  bkplms-ip < $S_{bkup}$ >
!
```

How it works

In this model the APs are provisioned on the master switch for a “LMS” and “backup-LMS”. When the access point first comes up it attempts to connect to the LMS and establishes the GRE tunnel with this switch. In case of a failure to this switch (which the access point detects

by missing heartbeat messages from the switch), the access points establishes a GRE tunnel with the “backup-LMS” switch and continues its operation through this switch. When the LMS comes back into operation, the administrator can reboot the access points to revert them back to the original LMS if required.

The above example shows how a single switch can be used to backup multiple switches. In the example shows, the switch S_{bkup} is present as a standby (not serving any clients/APs in the normal mode). If either S1 or S2 fail, the APs connected to them will fail over to S_{bkup} .

Note: If the “Standby” switch in the above shown topology is a pair of redundant master switches, the “Backup-LMS-IP” is the virtual IP address on the VRRP instance used for master redundancy.

Active -Active N to 1 Redundancy

This is a variant of the above described topology. In this model, a set of switches are serving access points and are configured in a manner such that if any one of the switches fails, the access points on that switch are load balanced amongst the remaining switches. The following figure illustrates this model.




Fig. 6 Active-active N-1 Redundancy

The configuration for such a topology implementation is as follows.

Master configuration

```
ap location 1.1.0
    lms-ip <S1>
    bkplms-ip <S2>
!
ap location 1.2.0
    lms-ip <S1>
    bkplms-ip <S3>
!
ap location 2.1.0
    lms-ip <S2>
    bkplms-ip <S1>
!
ap location 2.2.0
    lms-ip <S2>
    bkplms-ip <S3>
!
ap location 3.1.0
    lms-ip <S3>
    bkplms-ip <S1>
!
ap location 3.2.0
    lms-ip <S1>
    bkplms-ip <S2>
!
```

How it works

In the configuration above the access points are provisioned with a LMS-IP and Backup-LMS-IP such that the set of access points connecting to a LMS are split into sets with the Backup-LMS being each of the other switches. This is illustrated using the example above.

In the example above, all access points in building 1 (location codes 1.x.y) are configured with the LMS being S1. In addition all access points in building 1 and floor 1 (1.1.x) are configured with the Backup-LMS being S2, while all access points in building 1 and floor 2 (1.2.x) are configured with the Backup-LMS being S3. Thus in the event of a failure of S1 (or S1 losing connectivity), all access points in building 1 and floor 1 fail over to S2 while all access points in building 1 and floor 2 fail over to S3.

This type of redundancy configuration is useful to balance the load of a failed controller amongst multiple active controllers.

Summary

Aruba offers a variety of solutions to provide inter-controller redundancy in a centralized architecture to ensure highly available networks. These solutions can be used by an administrator to implement a highly redundant network depending on the requirements and constraints for each network design.

