Campus Redundancy Models

Version 1.0

Note

ARUBA networks

Copyright

Open Source Code

Certain Aruba products include Open Source software code developed by third parties, including software code subject to the GNU General Public License ("GPL"), GNU Lesser General Public License ("LGPL"), or other Open Source Licenses. The Open Source code used can be found at this site:

http://www.arubanetworks.com/open_source

Legal Notice

ARUBA DISCLAIMS ANY AND ALL OTHER REPRESENTATIONS AND WARRANTIES, WEATHER EXPRESS, IMPLIED, OR STATUTORY, INCLUDING WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, NONINFRINGEMENT, ACCURACY AND QUET ENJOYMENT. IN NO EVENT SHALL THE AGGREGATE LIABILITY OF ARUBA EXCEED THE AMOUNTS ACUTALLY PAID TO ARUBA UNDER ANY APPLICABLE WRITTEN AGREEMENT OR FOR ARUBA PRODUCTS OR SERVICES PURSHASED DIRECTLY FROM ARUBA, WHICHEVER IS LESS.

Warning and Disclaimer

This guide is designed to provide information about wireless networking, which includes Aruba Network products. Though Aruba uses commercially reasonable efforts to ensure the accuracy of the specifications contained in this document, this guide and the information in it is provided on an "as is" basis. Aruba assumes no liability or responsibility for any errors or omissions.

ARUBA DISCLAIMS ANY AND ALL OTHER REPRESENTATIONS AND WARRANTIES, WHETHER EXPRESSED, IMPLIED, OR STATUTORY, INCLUDING WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, NONINFRINGEMENT, ACCURACY, AND QUIET ENJOYMENT. IN NO EVENT SHALL THE AGGREGATE LIABILITY OF ARUBA EXCEED THE AMOUNTS ACTUALLY PAID TO ARUBA UNDER ANY APPLICABLE WRITTEN AGREEMENT OR FOR ARUBA PRODUCTS OR SERVICES PURCHASED DIRECTLY FROM ARUBA, WHICHEVER IS LESS.

Aruba Networks reserves the right to change, modify, transfer, or otherwise revise this publication and the product specifications without notice.

ARUBA networks

www.arubanetworks.com

1344 Crossman Avenue
Sunnyvale, California 94089

Phone: 408.227.4500
Fax 408.227.4550

Table of Contents

Chapter 1: About the Validated Reference Design Series

The Aruba Validated Reference Designs (VRDs) and application notes are a collection of technology deployment guides that include descriptions of Aruba technology, recommendations for product selections, network design decisions, configuration procedures, and best practices for deployment. Together these guides comprise a reference model for understanding Aruba technology and network designs for common customer deployment scenarios.

Each Aruba VRD and application note contains designs that are constructed in a lab environment and thoroughly tested by Aruba engineers. Our partners and customers use these proven designs to rapidly deploy Aruba solutions in production with the assurance that they will perform and scale as expected. Each guide in the series fits into one of five categories defined in the following manner:

- **Fundamentals:** These essential technology guides cover a broad range of Aruba products including mobility controllers, access points (APs), and site surveys. This is the starting point for new engineers to get familiar with Aruba products. All guides in the other categories build on the information in the fundamental guides. The guide that you are reading is part of the fundamental series.
- **Campus:** These guides cover designs for large campuses networks, including enterprise and education. Typically these networks are carpeted space deployments with hundreds of devices spread over several buildings.
- **Distributed Enterprise:** These guides focus on deployments that cover a wider geography, typically with smaller user and device counts. These deployments include K-12 schools, retail chains, branch offices, and remote workers.
- **Outdoor:** These guides cover large scale outdoor networks. These networks can include deployments such as metro-mesh, video surveillance, rail yards, point-to-point mesh links, and shipping facilities.
- **Network Services:** Guides in this section cover the operation of Aruba products including ClearPass, AirWave, and APAS. The solutions delivered by these guides apply equally to campus, distributed enterprise, and outdoor networks.

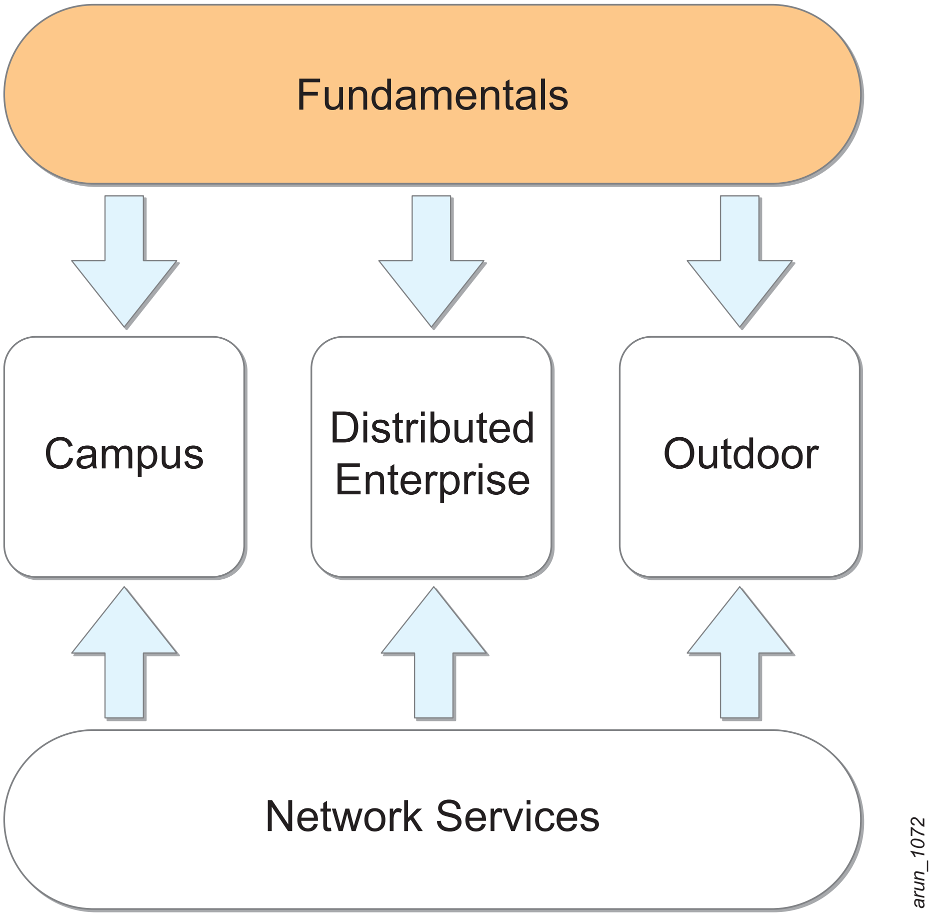

Figure 1 Aruba technology deployment guides

Chapter 2: Introduction

As the WLAN moves from a convenience network to mission-critical, the need for network availability also increases. Redundancy in a mobility controller system is designed to keep the APs functioning during an outage. To determine the network design, organizations must decide between the cost of building redundancy layers and the risk of the network being unavailable. In some cases, multiple types of redundancy are possible, and organizations must gauge their tolerance for risk given the pros and cons of each redundancy model.

In centralized Aruba WLAN deployments, the mobility controller is the heart of the network. The controller operates as a stand-alone master, or in a master-local cluster. Aruba provides several redundancy models for deploying mobility controllers. Each of these options, including the choice to forgo redundancy, must be understood so that the correct choice can be made for each deployment model.

The scale of redundancy has different levels, as seen in Figure 2. A completely redundant network is more resilient and costs more than the lower levels of the scale:

- Having a completely redundant network
- Having redundancy for aggregation-level mobility controllers (AP level)
- Having redundancy between a set of mobility controllers (small deployments)
- Having no redundancy at all

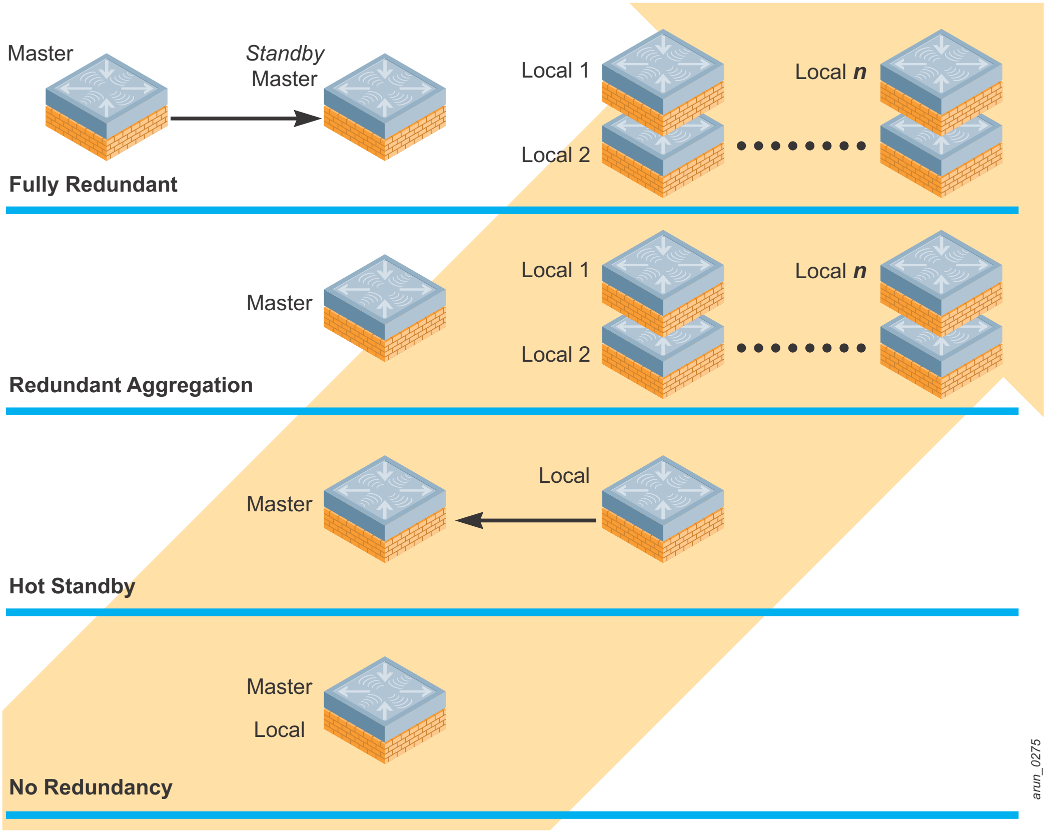

Figure 2 Scale of redundancy for mobility controllers

As a network moves up the scale, the cost and complexity increases. At the same time, the chance of the network being unusable due to a network outage decreases. This guide discusses redundancy at each level and what the consequences are of running a network without redundancy.

Appendix A provides failover testing times with the Aruba system. We have done extensive testing with real clients and simulators in an effort to decrease our failover times. It has become very apparent that client behaviors during an outage vary widely, and most of the outage time is related to clients reassociating to the network.

Additional Reference Material

- This guide assumes a working knowledge of Aruba products. Recommended reading for this guide is the *Aruba Mobility Controllers VRD.* This guide and others in the series are available for free at http://www.arubanetworks.com/vrd.
- The complete suite of Aruba technical documentation is available for download from the Aruba support site. These documents present complete, detailed feature and functionality explanations outside the scope of the VRD series. The Aruba support site is located at: https://support.arubanetworks.com/.
  This site requires a user login and is for current Aruba customers with support contracts.

Chapter 3: Master Redundancy

The master mobility controller is the control plane in a centralized WLAN. The master controller handles initial AP boot up in Layer 3 deployments, policy configuration and push to the local mobility controllers, local database access, and services such as security coordination and location. Additionally, if CPsec is enabled on the network, the master is responsible for certificate generation.

Master Active-Standby

To achieve high availability of the master mobility controller uses master redundancy (see Figure 3). In this scenario, two controllers are used at the management layer: one controller is configured as an active master and one is configured as a standby master. The two masters operate in a warm standby redundancy model. One master is the active primary, and the second is a standby that receives updates from the master about the state of the network.

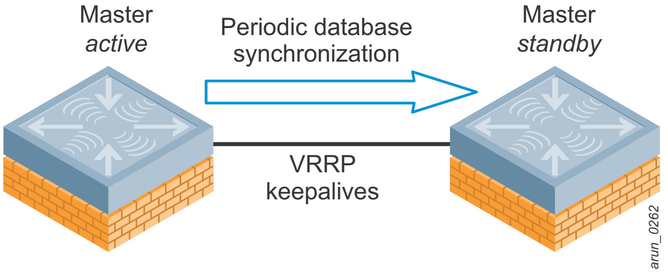

Figure 3 Master redundancy using VRRP and database synchronization

The two masters synchronize databases at a designated interval, typically every 30 minutes. The two controllers run a VRRP instance between them. The VRRP virtual IP (VIP) address is used by the local mobility controllers, mobility access switches (MASs), and APs that attempt to discover a mobility controller. The VIP address is also used for network administration.

When the primary master becomes unreachable for the timeout period, the backup master promotes itself to be the primary master and uses the VIP address. All traffic from locals and APs to the master automatically switches to the new primary as seen in Figure 4.

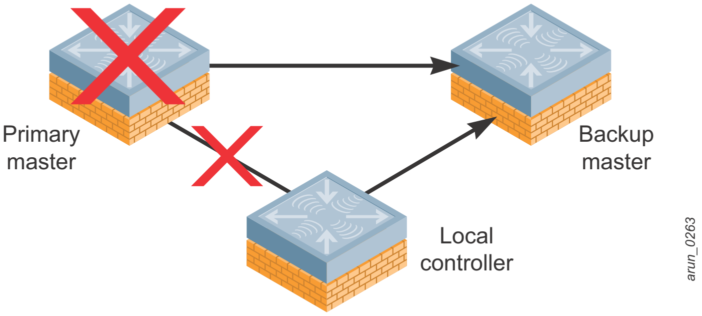

Figure 4 Master redundancy failure scenario for the local mobility controller

While VRRP pre-emption is supported, Aruba does not recommend enabling it in the master redundancy model. If preemption is disabled and a failover occurs, the new primary remains the primary even when the original master comes back online. The new primary does not revert to a backup unless an administrator forces it to do so. Disabling preemption prevents the master from “flapping” between two controllers and it allows the administrator to investigate the cause of the outage.

Master – No Redundancy

If the master fails without a backup, the following services stop working:

- **AP boot:** During the AP boot cycle, the AP must discover and connect to a provisioning mobility controller. In almost all deployments this is the master mobility controller, because that mobility controller typically is not serving APs and is able to be a single source for AP provisioning. It is also far easier to configure either DNS lookup or a single DHCP option to find a single mobility controller than to manage multiple lookups or scopes. It is also possible to use Layer 2 discovery mechanisms to find a local mobility controller, but this is not realistic in larger deployments. If the master is unreachable, in certain cases the APs may not be able to reboot until the master is restored or their boot process is modified:
  - In situations where DHCP option 43 is used, the APs are unable to boot until a new master is in place or the DHCP scope option is modified to point at either a new master or to a local.
  - If DNS is used to locate the master, the APs are down unless a second IP
  - address is also returned in the DNS response that points to a local. Note that this configuration results in a protracted outage, and the local must have AP capacity to bring up and then redirect the APs as they fail to the backup DNS response. This outage is longer in duration, with each AP taking approximately 4-5 minutes to fail to the backup. Depending on the AP capacity on the backup, several attempts may be needed before the AP is able to connect and be redirected properly.
  - APs that rely on Aruba Discovery Protocol (ADP) continue to operate as long as a local is capable of answering their ADP request. These APs require Layer 2 connectivity to the local for ADP to function.
  - In all cases, APs that are currently operating continue to do so in the event that the master becomes unreachable until they are rebooted or power cycled.
- **Local policy configuration:** Configuration, done either on the master or AirWave, requires that the master is operational to push configurations to the locals. If the master is not available, changes to the network policy configuration are not possible unless each mobility controller is modified manually, though local configuration at the IP level is possible.
- **Local database access is lost:** If the master becomes unreachable, guest access using the local database, as well as when roaming between locals when machine authentication is enabled, is lost.
- **Monitoring, heat maps, and location:** If AirWave is not present in the network, centralized network monitoring, heat map generation, and location services all are down.
- **Valid AP table:** When the master is down, the valid AP table is no longer available for updates. The locals continue to function with cached data until that ages out. After that time, other APs in the network are seen as “unknown” instead of valid, interfering, or rogue. When this occurs,

Adaptive Radio Management (ARM) increases power to the edge APs on both sides in an attempt to increase coverage and work around the now unknown AP. At AP border areas, overlapping channels and power lead to increased interference.

- **RFProtect coordination:** When the master is down, RFProtect security loses its coordination capabilities between locals. Any new APs that show up are classified as “unknown,” which prevents automatic containment from functioning. Existing data remains until it ages out, and then all of the APs begin to be reclassified as “unknown.” If protection of valid stations is enabled, clients are prevented from joining any AP that is not valid, which after some time will be all APs that the mobility controller can see that are not directly attached.
- **AP white lists:** The two varieties of white lists are the campus AP (CAP) and remote AP (RAP) white lists. For the CAP white list, all mobility controllers share a copy of the white list, but without the master, they lose the capability to synchronize the lists. The RAP white list must be exported to the local manually to ensure that operations continue, but no additional APs can be authorized while the master is unreachable.
- **CPsec:** Failure to have a backup for CPsec results in the same failures as a master mobility controller, with an additional problem. If the master physically must be replaced, as soon as it is brought online, the entire network goes back through the recertification list. In addition, the AP white list must be rebuilt.

Chapter 4: Local Redundancy

Three models of local redundancy are available. Each type of local redundancy is appropriate in a particular scenario, and sometimes they operate together. In each redundancy method, the goal is to provide the AP with a location where it can establish connectivity in the event of a mobility controller failure. The two methods for doing this are the Virtual Router Redundancy Protocol (VRRP) and the local management switch and backup local management switch (LMS / BLMS). These methods can be combined to provide local and data center redundancy across the three available deployment models. Table 1 describes the features and timings of each method.

Table 1 VRRP and LMS / BLMS Feature Comparison

Feature	VRRP	LMS / BLMS
Layer 2 or Layer 3 Operation	Operates at Layer 2	Operates at Layer 3
AP Reconnection	The AP radios rebootstrap on failover after the heartbeat times out. This takes about 10 seconds.	The AP radios rebootstrap after the heartbeat times out. The AP attempts to reestablish a connection to the primary LMS before failing to the backup LMS. This is takes approximately 1 min.

VRRP tends to be faster than LMS redundancy, but it only works at Layer 2. Aruba recommends running VRRP wherever possible, and reserving LMS redundancy where Layer 2 adjacency is not available, such as between data centers.

Active-Active (1:1)

In the Aruba active-active redundancy model, two locals share a set of APs, divide the load, and act as a backup for each other. Aruba recommends the active-active method of deploying redundant locals whenever they are Layer 2 adjacent. When two controllers operate together, they must run two instances of VRRP with each controller acting as the primary for one instance and backup for the other as shown in Figure 5.

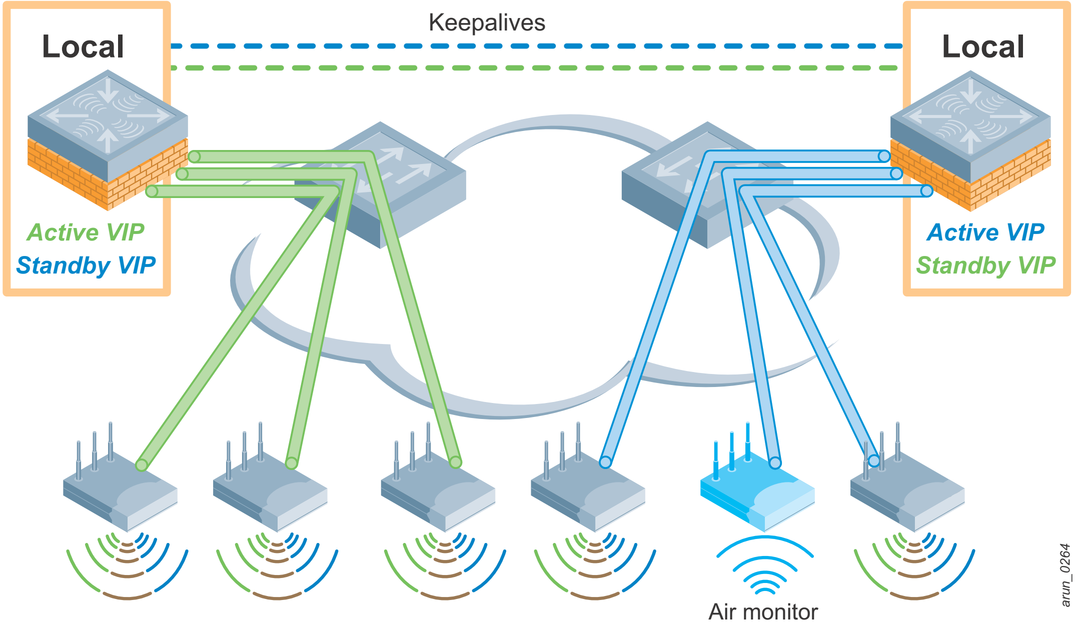

Figure 5 Active-active redundancy, both mobility controllers reachable

The controllers each terminate half of the APs in this redundancy cluster. The APs are configured in two different AP groups, each with a different VIP as the LMS IP address for that AP group. When one active local controller becomes unreachable, as in Figure 6, APs that are connected to the unreachable controller fail over to the second local. That controller now terminates all of the APs in the redundancy cluster. Therefore each controller must have sufficient processing power and licenses to accommodate all of the APs that are served by the entire cluster.

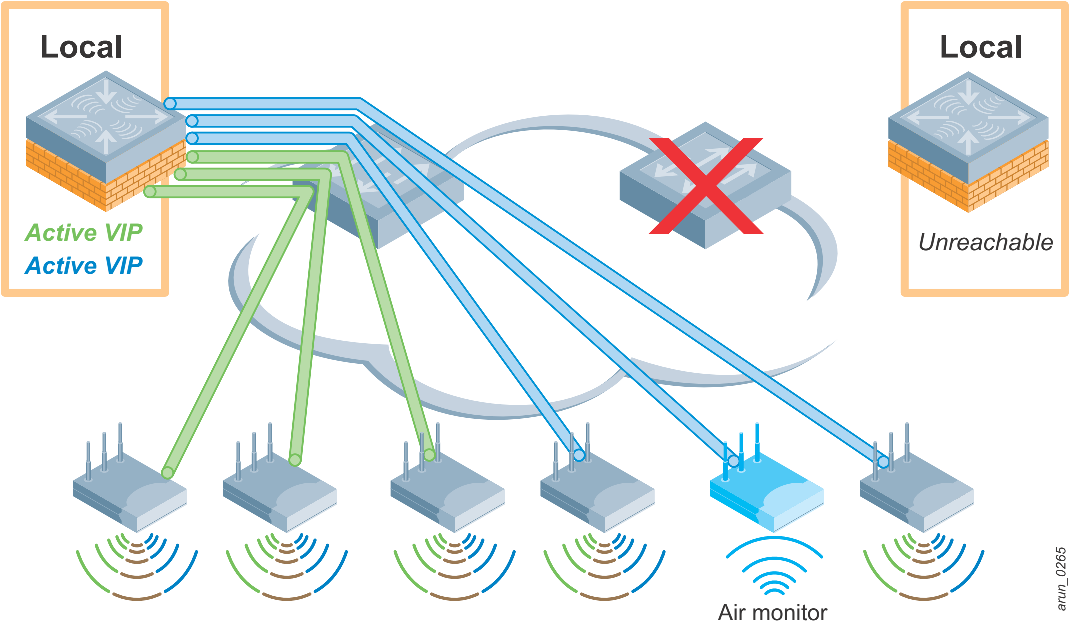

Figure 6 Active-active redundancy, mobility controller unreachable

In this model, preemption should be disabled so that APs are not forced to fail back to the original primary when it comes back online. APs will not fail back, so this model requires that the mobility controller be sized appropriately to carry the entire planned failover AP capacity for an extended period of time.

NOTE

When determining the AP load for active-active, give some thought (from a capacity standpoint) to what will happen to the backup controller when the APs fail over. If each mobility controller is at 50% of total capacity, when a failure occurs, the mobility controller that the APs fail over to will now be at 100% capacity. This leaves no room for future expansion of the system. Aruba recommends that each mobility controller be planned to run at 40% capacity (80% of total) to allow for future expansion.

“Salt-and-Pepper” Deployments

Active-active designs should not be confused with or considered synonymous with so called “Salt-and-Pepper” (SNP) designs. In an SNP design, APs are interspersed in the same RF coverage area, such that every other AP goes to a different mobility controller. Aruba strongly recommends against this design as it can have negative impacts on ARM operation and increases Layer 3 roaming events.

Part of what ARM does is to adjust AP channel and power settings to balance that AP with surrounding APs. In an SNP each APs is surrounded by an AP that is homed to a different controller but is still valid. ARM instances expect only to encounter this situation at the boundary of a controllers RF domain. Channel changes will not be coordinated between the APs in the same RF neighborhood. Aruba recommends each building be considered a single RF neighborhood for AP termination.

Layer 3 roaming is impacted because each time a device moves to a new AP, it either moves to a foreign agent, which causes a tunnel to be constructed and authentication to occur, or back to its home agent. This type of roaming is more disruptive to the client as opposed to roaming across APs on the same controller.

Aruba recommends against deploying in a SNP model.

Active-Standby (1+1)

The active-standby model also has two controllers, but in this case, one controller sits idle while the primary controller supports the full load of APs and users (see Figure 7).

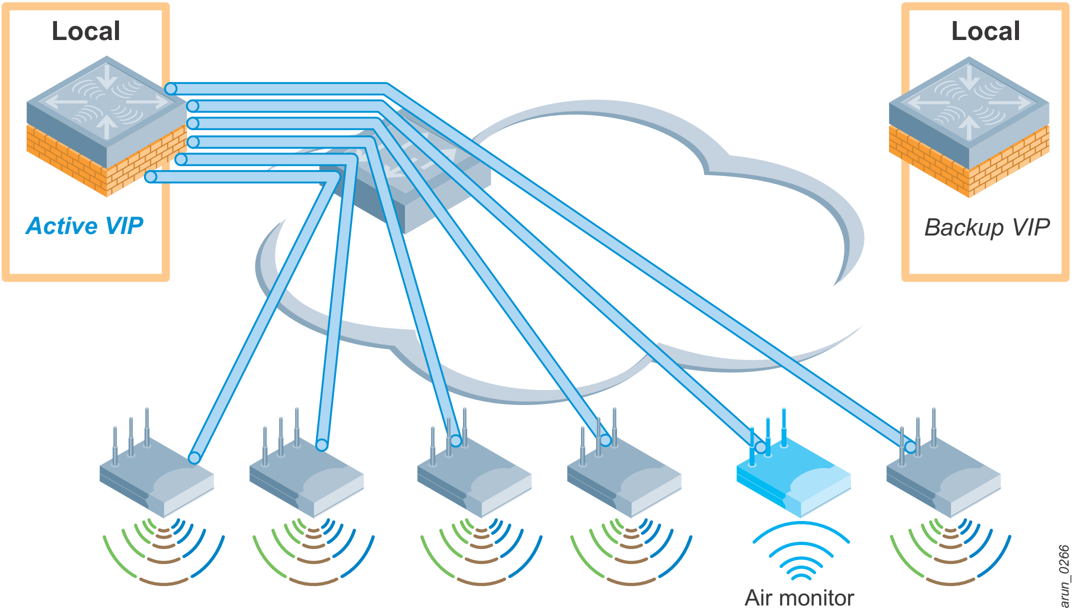

Figure 7 Active-standby redundancy, primary mobility controller is reachable

When a failure occurs in the active-standby model, all of the APs and users must fail over to the backup controller. This model has a larger failure domain and will have some increased latency because the full load of APs must fail over to the backup controller and users re-authenticate as shown in Figure 8. This form of redundancy typically uses the LMS and backup LMS configuration for the AP group because the controllers are usually in separate data centers. Alternatively, a single VRRP instance could be run between the two controllers, and all APs for the pair would terminate against this VRRP IP address.

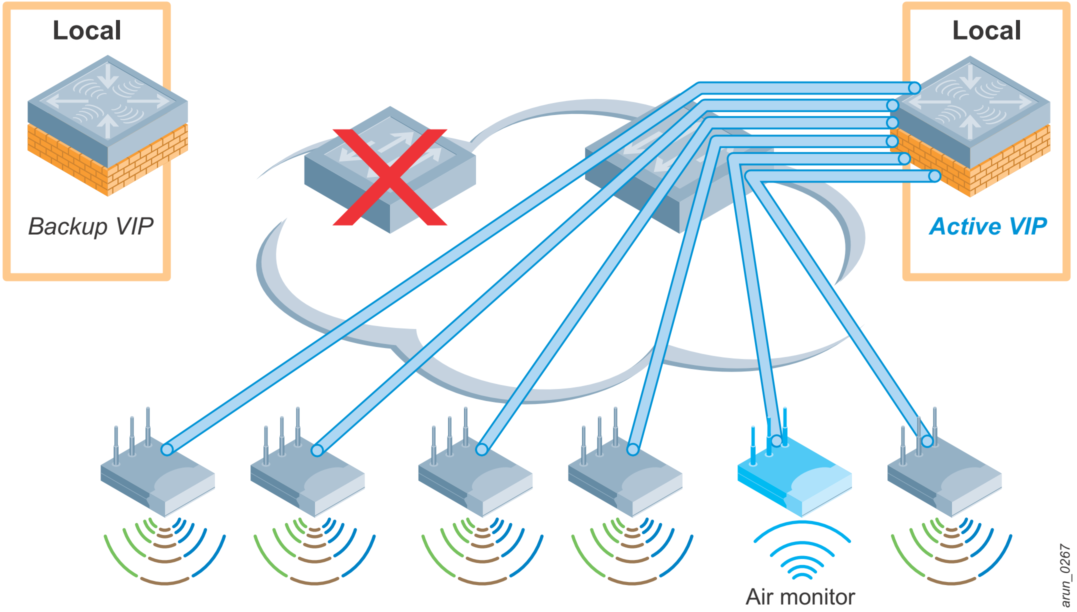

Figure 8 Active-standby controller, primary mobility controller is unreachable

As with active-active, when the active local becomes unreachable, all of the APs that are connected to the unreachable controller fail over to the standby local. That controller carries the full AP load duration of the outage. Therefore each controller must have sufficient processing power and licenses to accommodate all of the APs served by the entire cluster.

Many-to-One (N+1)

The many-to-one model typically is used in remote networks where branch offices have local mobility controllers but redundancy on site is not feasible. The store controllers are typically smaller models with a limited numbers of APs, and a much larger controller is deployed as the +1 in the data center as seen in Figure 9. This model requires that a secure connection is established between the sites that is independent of the mobility controllers, and that the connection should have high bandwidth and low latency.

It is possible to use N+1 on the campus as well, but here consideration should be given to the ratio and likelihood that sections of the campus might become unreachable, which would cause a multiple controller failover.

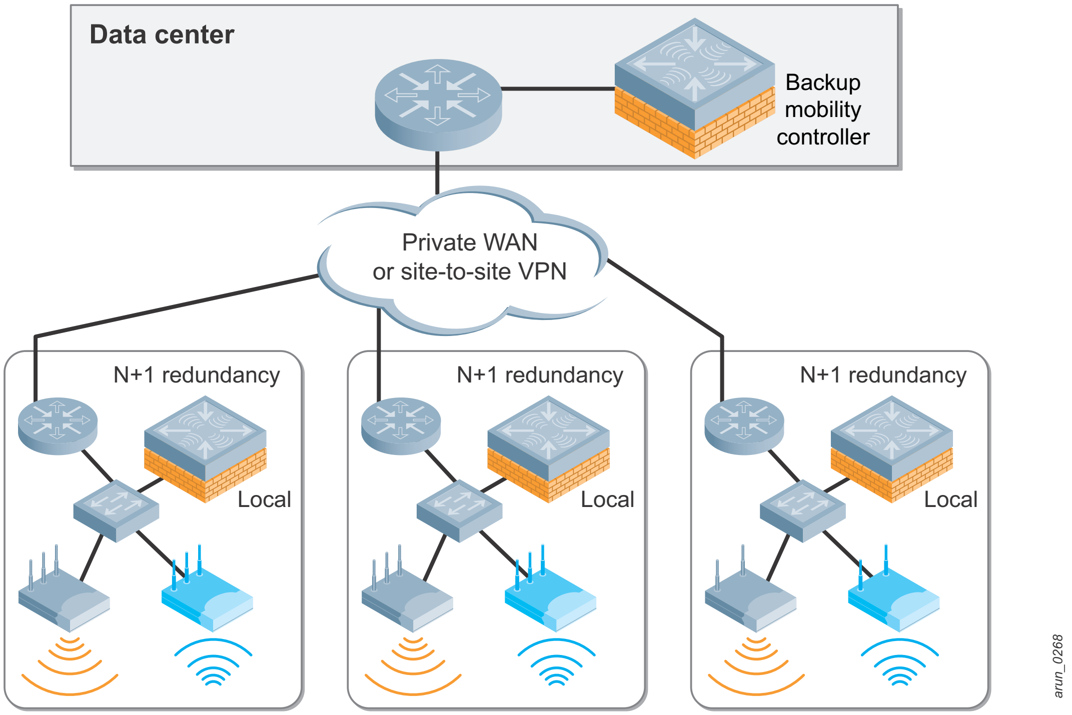

Figure 9 N+1 redundancy, local active

When the local at the remote site fails, the APs fail back to the backup LMS that is configured for that purpose, just as in the active-standby scenario (see Figure 10).

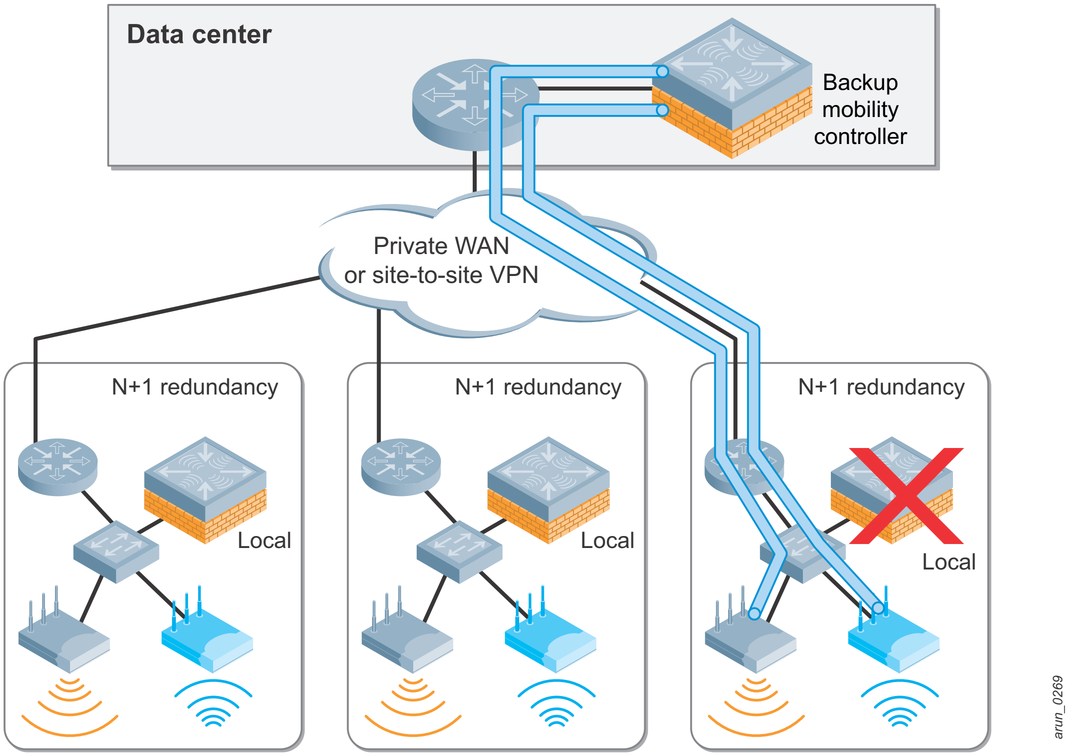

Figure 10 N+1 redundancy, local failed, AP connected across the WAN

The difference in the N+1 scenario is that this failure is typically across a WAN link, and the backup controller should be large enough to handle multiple site failures at the same time. Though a typical small site might have a handful of APs on a smaller mobility controller, the central site must have a much larger mobility controller with increased licensing to handle the expected number of failures of locals.

In typical campus designs, only a single failure is anticipated, but some organizations require more resiliency against failure of multiple sites. Common cases include retail stores, where more than a single store may have an outage at any one time due to the sheer number of sites and the fact that the controller may be in user-accessible space.

Aruba strongly recommends that preemption be enabled in this scenario. Due to the limited capacity of the redundant mobility controller and the possible delay introduced by failing over to a remote site, it is recommended that APs be moved back to their original mobility controller as soon as service is restored. In a campus deployment where the backup is often equal in capacity to the local, preemption is even more critical. If a second controller becomes unreachable, the backup controller typically does not have the capacity to accept the additional APs.

Some consideration should be given to the ratio of primary to backup controllers. If the backup mobility controller is the same model and scale as all of the locals that it is backing up, a single local can become unreachable in the network and the network still operates properly. If a second local becomes unreachable, all APs that exceed the capacity of the backup mobility controller are listed as unlicensed and do not operate. It is recommended that, where possible, the backup have the ability to terminate multiple locals in the event that multiple mobility controllers go off line.

Local – No Redundancy

If a local becomes unreachable and has no backups configured for the APs, all APs that are assigned to that mobility controller go down and no users can connect. Any AMs that are associated to the controller are also down, which eliminates the capability to scan for threats and contain rogue devices. This situation continues until the APs are reprovisioned and assigned to another mobility controller or the original or replacement local becomes reachable again.

Comparison of Local Redundancy Models

Table 2 summarizes the pros and cons of each redundancy model, which allows network managers to make the proper redundancy decision for their network.

Table 2 Comparison of Redundancy Models

Redundancy Type	Pro	Con
Active-Active (1:1)	● The failure domain is smaller, because fewer APs must fail over in the event of an outage. ● The outage duration is smaller, because fewer APs will take less time to recover, typically about half as long as failing over a fully loaded mobility controller. ● All mobility controllers are in use at all times. ● Each mobility controller has a reduced load.	● More expensive than N+1, because all mobility controllers must be licensed to handle the full complement of APs in the failure domain. Aruba recommends that this load be planned to 80% of the maximum capacity of each mobility controller ● Twice as many mobility controllers are required vs. no redundancy.
Active-Standby (1+1)	● If APs fail to the backup controller, essentially nothing has changed in the network except where the APs and users are hosted.	● Has the same cost structure as the active-active redundancy model, with two sets of mobility controllers and two sets of licenses. ● The failure domain is larger, all APs must fail to the backup mobility controller, which typically takes twice as long as active-active. ● The outage duration will be longer, because more APs must be recovered.
Many-to-One (N+1)	● In this cost-optimized model, fewer redundant mobility controllers are required, and they need only be licensed and scaled to handle the maximum number of failed mobility controllers. ● Typically only one redundant mobility controller is deployed.	● Multiple failures can overwhelm the redundant mobility controller, which causes a network down scenario. ● Preemption must be enabled to clear APs back to the primary mobility controller as soon as it is recovered, which results in a second unplanned outage.

Aruba recommends using active-active redundancy wherever possible. Active-active provides the fastest recovery time in the event of a network outage with the least disruption to the end user. Aruba also recommends in all models that mobility controllers not be loaded past the 80% mark. This load level helps increase the stability of the network during prolonged outages and allows for future growth of the network.

Chapter 5: Data Center Redundancy

The data center of an organization may experience an outage where all local mobility controllers at a particular site are offline but the network continues to operate. The APs can fail over to a redundant set of mobility controllers in another location as seen in Figure 11. The redundant controllers can be either in the same data center but connected by discrete power and data connections, or in a remote data center that is reachable by a private WAN or IPsec link.

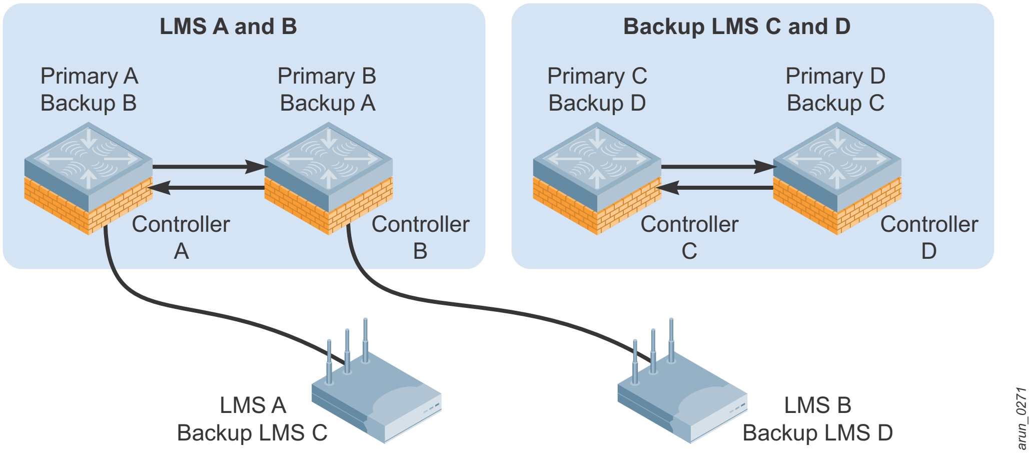

Figure 11 Active-active plus LMS and standby backup LMS

When the data center is at a remote site, consider the link between sites. The primary concerns are latency, overall bandwidth, and security. Latency affects authentication, such as 802.1X, and voice calls. Overall bandwidth needs to consider AP control traffic and user traffic. Finally, the connection should be secure between the sites, especially if decrypt tunnel is in use.

Data center redundancy consists of two to four total controllers, and up to four instances of VRRP. The APs can be set up either to split between two of the mobility controllers (active-active) with a pair in hot standby, or spread evenly across all four mobility controllers. In this model, the APs are set up so that they operate on the VIP address of their primary pair of mobility controllers (Figure 12), and their backup is one of the two VIP addresses on the second pair of mobility controllers (Figure 13).

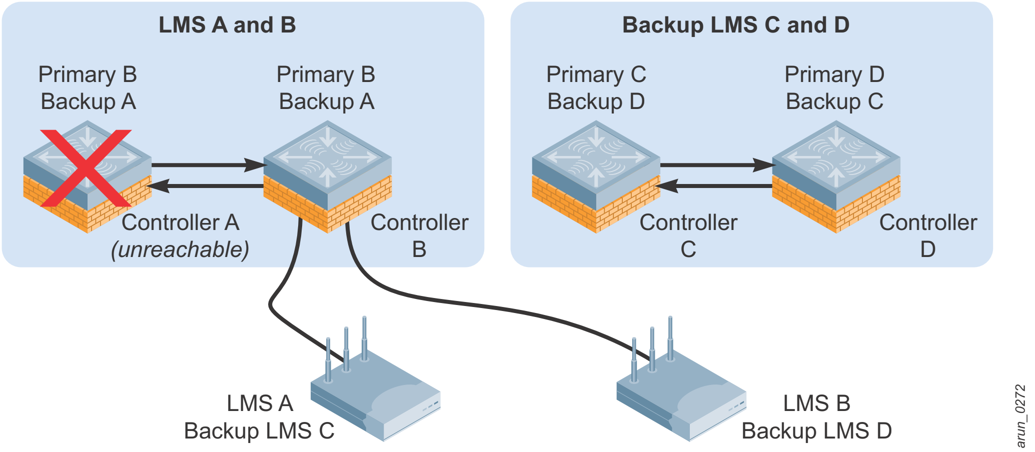

Figure 12 ***Failure of single primary mobility controllers in active-active with LMS and backup LMS***

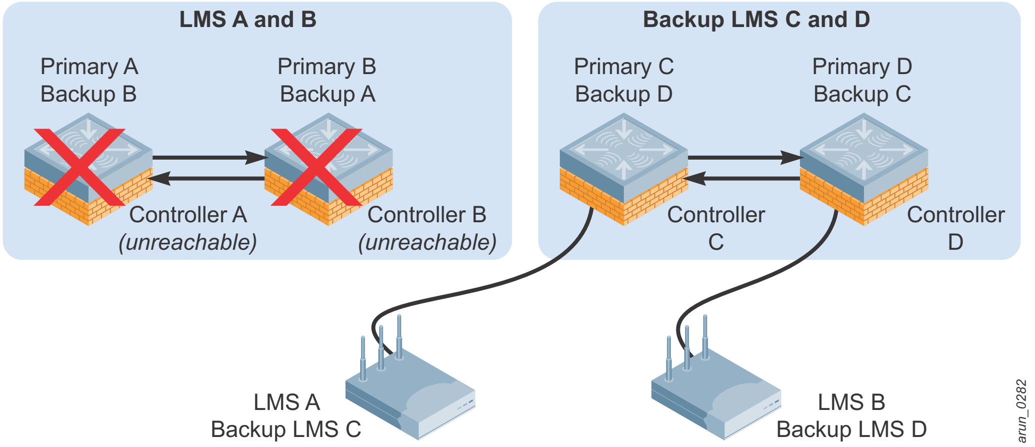

Figure 13 ***Failure of primary the primary data center in active-active with LMS and backup LMS***

In a failure scenario, the failure of one mobility controller in a pair results in typical active-active failover. If the second mobility controller in the pair fails, the APs fail over to their backup pair of controllers and split between the two VIP instances. In either deployment model, all four mobility controllers must be licensed and capable of supporting the full AP load.

Figure 14 shows scenario 1:

1. 412 APs were split across two active M3 mobility controllers (206 APs each), with each group active on one of the two VRRP instances in the first pair of locals and the two VRRP instances in the second pair standing by to receive APs.
2. When the local A fails, the APs move to the active backup local B. This change results in 412 APs on the backup local B, and 0 APs on each local (C and D) in the second cluster.
3. When local B fails, the 412 APs from the failed cluster distribute themselves evenly across the two locals C and D that are still active in the second cluster. This change results in 206 APs on each local.
4. If local C fails, all 412 APs become active on the remaining local D.
5. As a result, each mobility controller must be licensed to support all 412 APs if three of the other mobility controllers become unreachable.

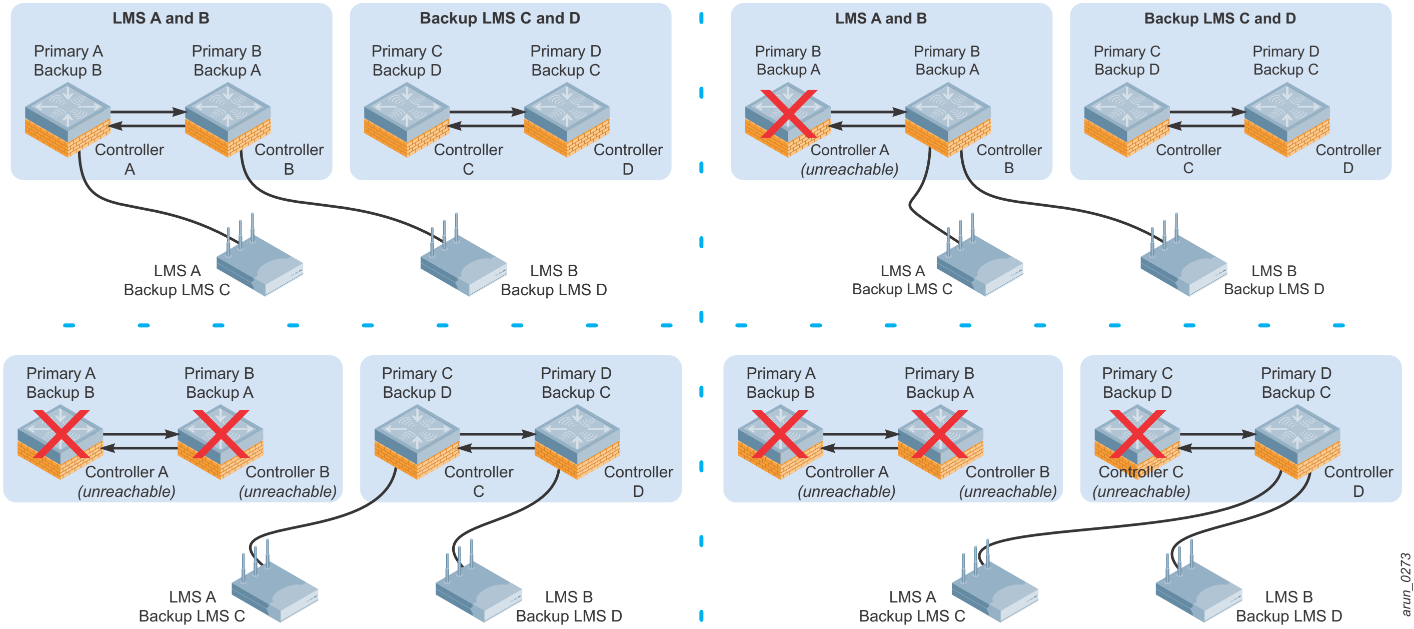

Figure 14 Failure series, active-active with LMS and standby backup LMS

Figure 15 shows scenario 2:

1. 412 APs were split across four active M3 mobility controllers (103 APs each), and each group was active on one VRRP.
2. When local A fails, the APs move to the active backup local B. This change results in 206 APs on the backup local B, and 103 APs on each local (C and D) in the second cluster.
3. When the local B fails, the 206 APs from the failed cluster distribute themselves evenly across the locals C and D that are still active in the second cluster. This change results in 206 APs on each mobility controller.
4. If local C fails, all 412 APs become active on the remaining local D.
5. As a result, each mobility controller must be licensed to support all 412 APs if three of the other mobility controllers become unreachable.

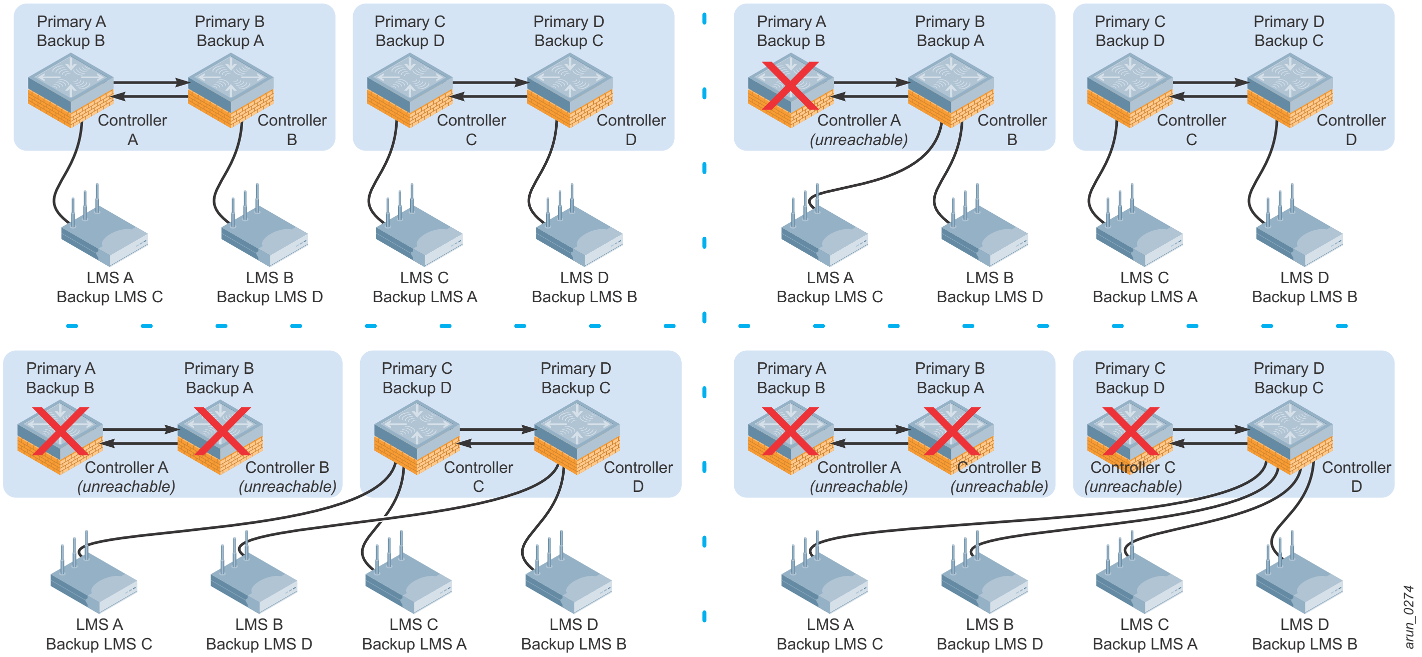

Figure 15 ***Failure series, active-active with LMS and backup LMS also in use***

Data Center – No Redundancy

Commonly, data center redundancy is deployed only by organizations with extremely high availability requirements and the ability to have the APs connect through a separate set of infrastructure to the second set of controllers. Each organization must decide on the acceptable level of risk vs. cost around this higher level of redundancy.

Chapter 6: Recommendations for Controller Redundancy

Wireless networks are no longer convenience networks. They are now mission-critical components of the network. As such, they need to be treated like any other mission-critical system. Aruba recommends redundancy at all levels of the system to ensure a highly available network for users.

Table 3 Redundancy Recommendations

Controller	Campus	Branch Office	Remote Access (DMZ)	Data Center
Master	Master redundancy	N/A	Master redundancy	Master redundancy
Local	Active-active redundancy, each mobility controller loaded at 40% of capacity, licensed to 80% of capacity	Active-active redundancy where possible, N+1 redundancy minimum	Active-active redundancy, each mobility controller loaded at 40% of capacity, licensed to 80% of capacity	Active-active redundancy, each mobility controller loaded at 40% of capacity, licensed to 80% of capacity

Appendix A: AP and Client Failover Times

Table 4 and Table 5 show failover times for APs failing over in an active-active deployment. These APs were tested in the Aruba labs using the VeriWave test tools to simulate clients. Aruba has also done extensive over-the-air testing. These tables were developed using two test cases:

- **Test case 1:** The test starts with APs distributed evenly between two mobility controllers. The test ends when all APs have completed their transition from the disconnected mobility controller to the remaining mobility controller.
- **Test case 2:** This test takes the first test case and includes APs and clients failing over between the two mobility controllers. Test case 2 represents a worst case scenario. This test starts with APs and clients evenly distributed between two mobility controllers. The test ends when the last client connection is re-established on the remaining mobility controller. Actual client experience will vary between a few seconds and the maximum value stated.

The client reauthentication rate is affected by a variety of factors outside of the WLAN infrastructure. In this scenario, the clients dominate the resultant failover times. Client reauthentication can vary considerably based on these things:

- **Client supplicant:** Problems with the client supplicant can include slow authentication and noncached credentials.
- **Client driver:** The client NIC card is slow to recognize that the network connection has been interrupted and is again available. The NIC card takes longer than expected to attempt to reconnect to the network.
- **Authentication type:** Different authentication types have different speeds for reauthentication. As an example 802.1X is more involved than an open network, and hence slightly slower.
- **Insufficient AAA infrastructure:** When a large number of clients attempt to reconnect to the network at the same time, the AAA infrastructure, such as RADIUS and LDAP servers, can become overwhelmed. Using a virtual machine for AAA has proven particularly bad when large numbers of clients attempt to re-associate.
- **Insufficient backend infrastructure:** Bottlenecks in the internal infrastructure can lead to longer response times and dropped packets, which create longer authentication times.

NOTE

These numbers are based on active-active redundancy, with half of the APs and users active on each mobility controller. For active-standby or N+1 redundancy, expect that failover times and client authentication can take 25% to 100% longer for each set of numbers. The longer times are caused by the greater number of APs and clients that fail over to the backup controller. For example, in the largest test case, 256 APs need to fail over. In the active-standby model, 512 APs need to fail over.

Simulated Client Tests

Table 4 AP Failover Times

Active-Active				
Test Case 1	CPsec Off	CPsec On	Test Case 2	CPsec On
64 CAP<==>64 CAP	17s	26s	1K User+ 256 CAP<==>1K User + 256 CAP	2m:20s
128 CAP<==>128 CAP	22s	38s	2K User+ 256 CAP<==>2K User + 256 CAP	2m:55s
256 CAP<==>256 CAP	48s	52s	4K User+ 256 CAP<==>4K User + 256 CAP	5m:45

Table 5 RAP Failover Times

RAP Active-Standby		RAP Active-Active	
1K Users + 512 RAP<==>0	3m:00s	1K Users+ 256 RAP<==>1K Users + 256 RAP	2m:10s
1K Users + 1K RAP<==>0	4m:30s	1K Users+ 512 RAP<==>1K Users + 512 RAP	3m:15s

Appendix B: Contacting Aruba Networks

Contacting Aruba Networks

Web Site Support	
Main Site	http://www.arubanetworks.com
Support Site	https://support.arubanetworks.com
Software Licensing Site	https://licensing.arubanetworks.com/login.php
Wireless Security Incident Response Team (WSIRT)	http://www.arubanetworks.com/support/wsirt.php
Support Emails	
Americas and APAC	support@arubanetworks.com
EMEA	emea_support@arubanetworks.com
WSIRT Email Please email details of any security problem found in an Aruba product.	wsirt@arubanetworks.com

Validated Reference Design Contact and User Forum	
Validated Reference Designs	http://www.arubanetworks.com/vrd
VRD Contact Email	referencedesign@arubanetworks.com
AirHeads Online User Forum	http://community.arubanetworks.com

Telephone Support	
Aruba Corporate	+1 (408) 227-4500
FAX	+1 (408) 227-4550
Support	
● United States	+1-800-WI-FI-LAN (800-943-4526)
● Universal Free Phone Service Numbers (UIFN):	
■ Australia	Reach: 1300 4 ARUBA (27822)
■ United States	1 800 9434526 1 650 3856589
■ Canada	1 800 9434526 1 650 3856589
■ United Kingdom	BT: 0 825 494 34526 MCL: 0 825 494 34526

Telephone Support

- Universal Free Phone Service Numbers (UIFN):

Japan	IDC: 10 810 494 34526 * Select fixed phones IDC: 0061 010 812 494 34526 * Any fixed, mobile & payphone KDD: 10 813 494 34526 * Select fixed phones JT: 10 815 494 34526 * Select fixed phones JT: 0041 010 816 494 34526 * Any fixed, mobile & payphone
Korea	DACOM: 2 819 494 34526 KT: 1 820 494 34526 ONSE: 8 821 494 34526
Singapore	Singapore Telecom: 1 822 494 34526
Taiwan (U)	CHT-I: 0 824 494 34526
Belgium	Belgacom: 0 827 494 34526
Israel	Bezeq: 14 807 494 34526 Barack ITC: 13 808 494 34526
Ireland	EIRCOM: 0 806 494 34526
Hong Kong	HKTI: 1 805 494 34526
Germany	Deutsche Telkom: 0 804 494 34526
France	France Telecom: 0 803 494 34526
China (P)	China Telecom South: 0 801 494 34526 China Netcom Group: 0 802 494 34526
Saudi Arabia	800 8445708
UAE	800 04416077
Egypt	2510-0200 8885177267 * within Cairo 02-2510-0200 8885177267 * outside Cairo
India	91 044 66768150