

Wireless LAN Controller (WLC) Configuration Best Practices

Document ID: 82463

Introduction

Prerequisites

- Requirements

- Components Used

- Conventions

Best Practices

- Wireless/RF

- Network Connectivity

- Network Design

- Mobility

- Security

- General Administration

- How to Transfer the WLC Crash File from the WLC CLI to the TFTP Server

NetPro Discussion Forums – Featured Conversations

Related Information

Introduction

This document offers short configuration tips that cover several Wireless Unified Infrastructure issues commonly seen in the Technical Assistance Center (TAC).

The objective is to provide important notes that you can apply on most network implementations in order to minimize possible problems.

Note: Not all networks are equal, therefore some tips might not be applicable on your installation. Always verify them before you perform any changes on a live network.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Knowledge of how to configure the Wireless LAN Controller (WLC) and Lightweight Access Point (LAP) for basic operation
- Basic knowledge of Lightweight Access Point Protocol (LWAPP) and wireless security methods

Components Used

The information in this document is based on these software and hardware versions:

- Cisco 2000 / 2100 / 4400 Series WLC that runs firmware 3.2 or 4.0
- LWAPP based Access Points, series 1230, 1240, 1130, 10x0 and 1500

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Conventions

Refer to Cisco Technical Tips Conventions for more information on document conventions.

Best Practices

Wireless/RF

These are the best practices for wireless/radio frequency (RF):

- For any wireless deployment, always do a proper site survey to insure proper quality of service for your wireless clients. The requirements for voice or location deployments are more strict than for data services. Auto RF might help on channel and power settings management, but it cannot correct a bad RF design.
- The site survey must be done with devices that match the power and propagation behavior of the devices to be used on the real network. For example, do not use a 350 802.11B radio with omni antenna to study coverage if the final network uses 1240 dual radios for 802.11A and G.
- In the same related idea, limit the number of service set identifiers (SSIDs) configured at the controller. Based on your access point model, you can have configured 8 or 16 simultaneous SSIDs, but as each WLAN/SSID needs separated probe responses, and beaconing, the RF pollution increases as more SSIDs are added. The results are that some smaller wireless stations like PDA, WiFi Phones and barcode scanners cannot cope with a high number of basic SSID (BSSID) information. This results in lockups, reloads or association failures. Also the more SSIDs, the more beaconing needed, so less RF space is available for real data transmits.
- For RF environments that are clear spaces, like factories where there are access points in a large space without walls, it might be necessary to adjust the Transmit Power Threshold from the default of -65 dBm, to a lower value like -76 dBm. This allows you to lower the co-channel interference (number of BSSID heard from a wireless client in a given moment). The best value is dependant on each site environmental characteristics, so it should be evaluated carefully with a site survey.

Power Transmit Threshold This value, expressed in dBm, is the cut-off signal level at which the Transmit Power Control (TPC) algorithm adjusts the power levels downward, such that this value is the strength at which the third strongest neighbor of an AP is heard.

- Some 802.11 client software might encounter difficulties if it hears more than a certain fixed number of BSSIDs (for example, 24 or 32 BSSIDs.) When you reduce the transmit power threshold and hence the average AP transmit level, you can reduce the number of BSSIDs that such clients hear.
- Do not enable aggressive load balancing unless the network has available a high density of access points in the area, and never if there is voice over wireless. If you enable this feature with access points spaced to far away from each other, it might confuse the roaming algorithm of some clients, and induce coverage holes in some cases. In the latest software versions, this feature is disabled by default.

Network Connectivity

These are the best practices for network connectivity:

- Do not use spanning tree on controllers.

For most topologies, Spanning Tree Protocol (STP) that runs in the controller is not needed. STP is disabled by default.

For non-Cisco switches, it is also recommended that you also disable STP on a per port basis.

Use this command in order to verify:

```
Cisco Controller) >show spanningtree switch

STP Specification..... IEEE 802.1D
STP Base MAC Address..... 00:18:B9:EA:5E:60
Spanning Tree Algorithm..... Disable
STP Bridge Priority..... 32768
STP Bridge Max. Age (seconds)..... 20
STP Bridge Hello Time (seconds)..... 2
STP Bridge Forward Delay (seconds)..... 15

(Cisco Controller) >
```

- Although most of controller configuration is applied "on the fly", it is good idea to reload controllers after you change the following configuration settings:

- ◆ Management address
- ◆ SNMP configuration This is very important if you use older software.
- For all trunk ports that connect to the controllers, filter out the VLANs that are not in use.

For example in Cisco IOS® switches, if the management interface is on VLAN 20, plus VLAN 40 and 50 are used for two different WLANs, use this configuration command at the switch side:

```
switchport trunk allowed vlans 20,40,50
```

- Do not configure the service port with an overlapping subnet to the management interface.
- Do not leave an interface with a 0.0.0.0 address, for example an unconfigured service port. It might affect DHCP handling in the controller.

This is how you verify:

```
(Cisco Controller) >show interface summary
```

Interface Name	Port	Vlan Id	IP Address	Type	Ap Mgr
ap-manager	LAG	15	192.168.15.66	Static	Yes
example	LAG	30	0.0.0.0	Dynamic	No
management	LAG	15	192.168.15.65	Static	No
service-port	N/A	N/A	10.48.76.65	Static	No
test	LAG	50	192.168.50.65	Dynamic	No
virtual	N/A	N/A	1.1.1.1	Static	No

- Do not use LAG unless all ports of the controller have the same Layer 2 configuration on the switch side. For example, avoid filtering some VLANs in one port, and not the others.
- When you use LAG, the controller relies on the switch for the load balancing decisions on traffic that comes from the network. It expects that traffic that belongs to an AP (LWAPP or network to wireless user) always enters on the same port. Use only ip-src or ip-src ip-dst load balancing options in the switch EtherChannel configuration. Some switch models might use unsupported load balancing mechanisms by default, so it is important to verify.

This is how to verify the EtherChannel load balancing mechanism:

```
switch#show etherchannel load-balance
EtherChannel Load-Balancing Configuration:
src-dst-ip
```

EtherChannel Load-Balancing Addresses Used Per-Protocol:
Non-IP: Source XOR Destination MAC address
IPv4: Source XOR Destination IP address
IPv6: Source XOR Destination IP address

This is how to change the switch configuration (IOS):

```
switch(config)#port-channel load-balance src-dst-ip
```

- Do not configure a LAG connection that spans across multiple switches. When you use LAG, it must be with all ports that belong to the same EtherChannel that goes to the same physical switch. It is possible to bypass this limitation using specific scenarios with 3750 cross-stack EtherChannel functionality only.
- If you do *not* use a LAG topology, you should always create an AP-manager per physical port, for redundancy and scalability.
- Never configure a backup port for an AP-manager interface, even if it is allowed in older software versions. The redundancy is provided by the multiple AP-manager interfaces as mentioned earlier in this document.

Network Design

These are the best practices for network design:

- Limit the number of access points per VLAN. A good number is around 30 to 60 and depends on network characteristics. This helps to minimize reassociation problems in case of network failure.
- In relation to the first tip, do not put more than 20 access points in the same VLAN with the management interface of the controller. It is possible that due to a high number of broadcast messages generated by the access points, some discovery messages are dropped and this results in a slower access point joining processes.
- Per design, most of the CPU initiated traffic is sent from the management address in the controller. For example, SNMP traps, RADIUS authentication requests, and so forth.

The exception to this rule is DHCP related traffic, which is sent from the interface related to the WLAN settings, for controller software version 4.0 and later. For example, if a WLAN uses a dynamic interface, the DHCP request is forwarded using this Layer 3 address.

This is important to take into account when you configure firewall policies or design the network topology. It is important to avoid configuring a dynamic interface in the same sub network as a server that has to be reachable by the controller CPU, for example a RADIUS server, as it might cause asymmetric routing issues.

Mobility

These are the best practices for mobility:

- All controllers in a mobility group should have the same IP address for a virtual interface, for example 1.1.1.1. This is important for roaming.

This is how to verify:

```
(Cisco Controller) >show interface summary
```

Interface Name	Port	Vlan Id	IP Address	Type	Ap Mgr
-----	-----	-----	-----	-----	-----

ap-manager	LAG	15	192.168.15.66	Static	Yes
management	LAG	15	192.168.15.65	Static	No
service-port	N/A	N/A	10.48.76.65	Static	No
test	LAG	50	192.168.50.65	Dynamic	No
virtual	N/A	N/A	1.1.1.1	Static	No

- The virtual gateway address should be "not routable" inside your network infrastructure. It is only intended to be reachable for a wireless client when connected to a controller, never from a wired connection.
- Do not create unnecessarily large mobility groups. A mobility group should only have all controllers that have access points in the area where a client can physically roam, for example all controllers with access points in a building. If you have a scenario where several buildings are separated, they should be broken into several mobility groups. This saves memory and CPU, as controllers do not need to keep large lists of valid clients, rogues and access points inside the group, which would not interact anyway.

Keep in mind that WLC redundancy is achieved through the mobility groups. So it might be necessary in some situations to increase the mobility group size, including additional controllers for redundancy (N+1 topology for example).

- In scenarios where there is more than one controller in a mobility group, it is normal to see some rogue access point alerts about our own access points in the network after a controller reload. This happens due to the time it takes to update the access point, client and rogue lists between mobility group members.
- The DHCP Required option in WLAN settings allows you to force clients to do a DHCP address request/renew every time they associate to the WLAN before they are allowed to send or receive other traffic to the network. From a security standpoint, this allows for a more strict control of IP addresses in use, but also might have affects in the total time for roaming before traffic is allowed to pass again.

Additionally, this might affect some client implementations which do not do a DHCP renew until the lease time expires. For example, Cisco 7920 or 7921 phones might have voice problems while they roam if this option is enabled, as the controller does not allow voice or signaling traffic to pass until the DHCP phase is completed. Some third-party printer servers might also be affected. In general, it is a good idea not to use this option if the WLAN has non-Windows clients. This is because the more strict controls might induce connectivity issues, based on how the DHCP client side is implemented. This is how you verify:

```
(Cisco Controller) >show wlan 1
```

```
WLAN Identifier..... 1
Profile Name..... 4400
Network Name (SSID)..... 4400
Status..... Enabled
MAC Filtering..... Disabled
Broadcast SSID..... Enabled
AAA Policy Override..... Disabled
Number of Active Clients..... 0
Exclusionlist Timeout..... 60 seconds
Session Timeout..... 1800 seconds
Interface..... management
WLAN ACL..... unconfigured
DHCP Server..... Default
DHCP Address Assignment Required..... Disabled
Quality of Service..... Silver (best effort)
WMM..... Disabled
CCX - AironetIe Support..... Enabled
CCX - Gratuitous ProbeResponse (GPR)..... Disabled
Dot11-Phone Mode (7920)..... Disabled
Wired Protocol..... None
```

Security

These are the best practices for security:

- It is good idea to change the RADIUS timeout to 5 seconds. The default of 2 seconds is acceptable for a fast RADIUS failover, but probably not enough for Extensible Authentication Protocol–Transport Layer Security (EAP–TLS) authentication, or if the RADIUS server has to contact external databases (Active Directory, NAC, SQL, and so forth).

This is how to verify:

```
(Cisco Controller) >show radius summary
Vendor Id Backward Compatibility..... Disabled
Credentials Caching..... Disabled
Call Station Id Type..... IP Address
Administrative Authentication via RADIUS... Enabled
Aggressive Failover..... Disabled
Keywrap..... DisabledAuthentication Servers
```

!--- This portion of code has been wrapped to several lines due to spatial concerns.

Idx	Type	Server Address	Port	State	Tout	RFC3576
1	N	10.48.76.50	1812	Enabled	2	Enabled

```
IPSec -AuthMode/Phase1/Group/Lifetime/Auth/Encr
-----
Disabled - none/unknown/group-0/0 none/none
```

This is how to configure:

```
config radius auth retransmit-timeout 1 5
```

- Check on the SNMPv3 default user. By default, the controller comes with a username that should be disabled or changed.

This is how to verify:

```
(Cisco Controller) >show snmpv3user

SNMP v3 User SNMP v3 User Name AccessMode Authentication Encryption
-----
default Read/Write HMAC-MD5 CBC-DES
```

This is how to configure:

```
config snmp v3user delete default
config snmp v3user create nondefault rw hmacsha des authkey encrkey
```

Keep in mind that your SNMP settings must match between the controller and the Wireless Control System (WCS). Also, you should use an encryption and hash keys that match your security policies.

- When you perform web authentication with an external authentication page, do not use a server which has a web server and a proxy server at the same time to host the login page. The controller allows HTTP traffic from the wireless client to the server before authentication is complete. This allows the client to navigate using the proxy service present on the server.

- In the controllers, the default timeout for the EAP Identity request is 1 second, which is not enough for some situations like One Time Passwords, or Smart Card implementations, where the user is prompted to write a PIN or password before the wireless client can answer the identity request. In autonomous access points, the default is 30 seconds, so this should be taken into account while you migrate autonomous to infrastructure wireless networks.

This is how to change:

```
config advanced eap identity-request-timeout 30
```

- On aggressive environments, a helpful feature is to enable access point authentication with a threshold of 2. This permits both to detect possible impersonation and minimize false positive detections.

This is how to configure:

```
config wps ap-authentication enable
config wps ap-authentication threshold 2
```

- In relation to the previous tip, Management Frame Protection (MFP) can also be used to authenticate all 802.11 management traffic detected between nearby access points in the wireless infrastructure. Take into consideration that some common third party wireless cards have problems in their driver implementation that do not handle correctly the extra information elements added by MFP. Make sure you use the latest drivers from your card manufacturer before you test and use MFP.
- NTP is very important for several features. It is mandatory to use NTP synchronization on controllers if you use any of these features: Location, SNMPv3, access point authentication, or MFP.

This is how to configure:

```
config time ntp server 1 10.1.1.1
```

In order to verify, check for entries like this in your traplog:

```
30 Tue Feb 6 08:12:03 2007 Controller time base status -
Controller is in sync with the central timebase.
```

- When using Protected EAP–Microsoft Challenge Handshake Authentication Protocol version 2 (PEAP–MSCHAPv2), with Microsoft XP SP2, and the wireless card is managed by the Microsoft Wireless Zero Configuration (WZC), you should apply the Microsoft hotfix KB885453 . This prevents several issues on authentication related with PEAP Fast Resume.
- If, for security reasons, the wireless clients should be separated in several sub networks, each one with different security policies, it is good idea to use one or two WLANs (for example, each one has a different Layer 2 encryption policy), together with the AAA–Override feature. This feature allows you to assign per user settings. For example, move the user to either a specific dynamic interface in a separated VLAN, or apply a per user Access Control List.
- Although the controller and access points do support WLAN with SSID using Wi-Fi Protected Access (WPA) and WPA2 simultaneously, it is very common that some wireless client drivers cannot handle complex SSID settings. In general, it is a good idea to keep the security policies simple for any SSID, for example, using one WLAN/SSID with WPA and Temporal Key Integrity Protocol (TKIP), plus a separated one with WPA2 and Advanced Encryption Standard (AES).

General Administration

These are the best practices for General Administration:

- In general, before any upgrade it is a good idea to do a binary backup of the configuration. WLCs support the conversion of older configuration information into new versions, but there is no support

for the reverse process. This applies both to major or minor version changes.

- The use of a newer configuration into an older release might lead to missing settings (access lists, interfaces, and so forth), or on incorrectly working features. If you need to downgrade a controller, it is advisable that after the downgrade, you clear the configuration, configure back the management interface address, and load the binary backup file by TFTP.

How to Transfer the WLC Crash File from the WLC CLI to the TFTP Server

Issue these commands in order to transfer the WLC crash file from the WLC CLI to the TFTP server.

```
transfer upload datatype crashfile
transfer upload serverip <IP address of the TFTP Server>

transfer upload path <Enter directory path>

transfer upload filename <Name of the Crash File>

transfer upload start<yes>
```

Note: When you enter the directory path, "/" usually means the default root directory on the TFTP server.

Here is an example:

```
(Cisco Controller) >debug transfer tftp enable

(Cisco Controller) >debug transfer trace enable

(Cisco Controller) >transfer upload datatype crashfile

(Cisco Controller) >transfer upload filename aire2cra.txt

(Cisco Controller) >transfer upload path /

(Cisco Controller) >transfer upload serverip X.Y.Z.A

(Cisco Controller) >transfer upload start

Mode..... TFTP TFTP Server
IP..... X.Y.Z.A TFTP
Path..... / TFTP
Filename..... aire2cra.txt Data
Type..... Crash File

Are you sure you want to start? (y/N) yes
Thu Dec 29 10:13:17 2005: RESULT_STRING: TFTP Crash File transfer starting.
Thu Dec 29 10:13:17 2005: RESULT_CODE:1

TFTP Crash File transfer starting.
Thu Dec 29 10:13:21 2005: Locking tftp semaphore, pHost=X.Y.Z.A
pFilename=/aire2cra.txt Thu Dec 29 10:13:22 2005:
Semaphore locked, now unlocking,
pHost=X.Y.Z.A pFilename=/aire2cra.txt Thu Dec 29 10:13:22 2005:
Semaphore successfully unlocked,
pHost=X.Y.Z.A pFilename=/aire2cra.txt Thu Dec 29 10:13:22 2005:
tftp rc=0, pHost=X.Y.Z.A pFilename=/aire2cra.txt

pLocalFilename=/mnt/application/bigcrash
```

Thu Dec 29 10:13:22 2005: RESULT_STRING: File transfer operation completed successfully.
Thu Dec 29 10:13:22 2005: RESULT_CODE:11 File transfer operation completed successfully.

NetPro Discussion Forums – Featured Conversations

Networking Professionals Connection is a forum for networking professionals to share questions, suggestions, and information about networking solutions, products, and technologies. The featured links are some of the most recent conversations available in this technology.

NetPro Discussion Forums – Featured Conversations for Wireless
Wireless – Mobility: WLAN Radio Standards
Wireless – Mobility: Security and Network Management
Wireless – Mobility: Getting Started with Wireless
Wireless – Mobility: General

Related Information

- [Lightweight Access Point FAQ](#)
- [Wireless LAN Controller \(WLC\) Troubleshoot FAQ](#)
- [Cisco Wireless LAN Controller Module Q&A](#)
- [Cisco Wireless LAN Controllers Q&A](#)
- [Radio Resource Management under Unified Wireless Networks](#)
- [Wireless Support](#)
- [Wireless LAN \(WLAN\) Technology Support](#)
- [Technical Support & Documentation – Cisco Systems](#)

All contents are Copyright © 2006–2007 Cisco Systems, Inc. All rights reserved. Important Notices and Privacy Statement.

Updated: Jun 12, 2007

Document ID: 82463