

HiPerWA: High Performance Wireless Analytics

A survey into the challenges of constructing, managing, debugging and optimizing an in-production large-scale enterprise network

Bjørn Ludvig Langaas Johansen

INF-3981: Master thesis in Computer Science — June 2016

To the brave people who dare search for the moving asm line...

“The best computer is a man, and it’s the only one that can be mass-produced
by unskilled labor.”
–Wernher von Braun

“The most likely way for the world to be destroyed, most experts agree, is by
accident. That’s where we come in; we’re computer professionals. We cause
accidents.”
–Nathaniel S. Borenstein

Abstract

Enterprise wireless networks are becoming larger and larger, with more and more users and devices connecting to the networks. This requires the infrastructure to be closely monitored and adjusted to ensure an optimum experience for users and their devices. Existing proprietary solutions exist, but are either costly, resource demanding, rigid or may not be able to deliver the functionality that is required to operate a modern wireless infrastructure efficiently.

In this project, a closer look has been taken at some of the key metrics from a wireless infrastructure, and what they can tell about the health and state of the infrastructure. These metrics has been collected by a modular, customizable implementation designed to be extensible and capable of delivering customized metrics and analytics that helps in the day to day operation of the infrastructure.

As an example, information collected has been used to show how interference in high-density deployments of 2.4GHz 802.11 radios can be reduced by deactivating some of the radios to ensure a better 2.4GHz environment, and better client experience.

Acknowledgements

First I want to thank my advisor, John Markus Bjørndalen for many interesting discussions, both on-topic and off-topic, and for believing in my (sometimes) vague ideas or notions.

I would also thank my colleagues at the IT-department at UiT for their valuable input, support and encouragement during my studies. A special thanks to my colleague Anders Baardsgaard, with whom I have shared both an office, and many interesting discussions on the topic of wireless networks and protocols.

A special thanks to my friends, fellow classmates, and all fellow students throughout my studies. The memories of late nights, early mornings, heated discussions, hardcore procrastination, and the satisfaction of discovering new things with friends will never be forgotten.

"It was the best of times, ~~it was the worst of times~~" [1, p. 3]

Finally, I would like to thank my family, my mother and father, my sister, and my two brothers, for encouraging, supporting and believing in me. I would also thank my two cats, Turbo and Pusi for helping me to relax, and teaching me the value of a "thinking powernap", and when it is needed.

Contents

Abstract	iii
Acknowledgements	v
List of Figures	xi
List of Abbreviations	xiii
1 Introduction	1
1.1 History	3
1.1.1 UiT	3
1.1.2 Previous work	5
1.2 Motivation	5
1.3 Problem description	6
1.3.1 Interference issues	8
1.4 Challenges	10
1.4.1 Data mining	10
1.4.2 Potential bottlenecks	11
1.4.3 Client device support limitations	11
1.4.4 Legal limitations	11
2 Background information	13
2.1 Enterprise wireless networks	13
2.1.1 Common Architectures	14
2.1.2 Access Points	16
2.1.3 RRM	16
2.1.4 CAPWAP	17
2.2 Radio Resource management	18
2.2.1 DCA	18
2.2.2 TPC	20
2.2.3 Air Quality	20
2.2.4 Other measurements	21
2.3 802.11	21
2.3.1 Protocol	21

2.3.2	Interference	22
2.3.3	Noise	23
2.3.4	Signal-Noise Ratio	24
2.3.5	Protocol Impact	24
2.3.6	Problem illustration	25
2.4	Eduroam	26
3	Goals	27
3.1	Improvements over the current system	27
3.1.1	System integration and visualization	28
3.1.2	Automatic (pre) provisioning	28
3.1.3	Monitoring	29
3.2	Interference Reduction	29
3.3	2.4GHz radio shutoff	29
4	Data collection and storage	31
4.1	Development and Experimental Setup	31
4.2	SNMP	34
4.3	Data format and namespace	34
4.4	Storage	35
4.5	Extensions	35
5	System design and adaptations	37
5.1	SNMP-interaction	37
5.2	Access Point model	38
5.2.1	Accesspoints (plural)	38
5.2.2	Accesspoint (singular)	38
5.3	Controller model	39
5.4	Analytics and correlations	39
5.5	Optimizations	39
5.5.1	Poor SNR conditions	40
5.5.2	Detection of failed radios	40
5.5.3	Detection of failed subsystems in access points	41
5.6	(Semi-)Automatic actions	42
5.7	Alerts	42
6	General monitoring	45
6.1	CAPWAP health	45
6.1.1	Detecting and classifying access point teardowns	46
6.1.2	Controller caused teardowns	47
6.2	Controller health	48
7	Interference reduction	49
7.1	Requirements	49

7.2	Method	49
7.3	External Interference adaptations	50
7.4	Other efforts	51
7.5	Further development	51
8	Testing	53
8.1	Performance	53
8.1.1	Data collection	53
8.1.2	Data presentation	54
8.1.3	Analytics	54
8.2	Use cases	54
8.3	Correctness	55
9	Results and observations	57
9.1	Interference reduction	57
9.2	SNR observations	58
9.3	Data size impact	58
10	End products	61
10.1	SNMP abstraction layer	61
10.2	Data collection and storage code	61
10.2.1	Wireless controller	61
10.2.2	Access points	62
10.3	HTTP GUI	62
10.3.1	Dashboards	64
10.4	Command line provisioning tool	66
10.5	Unfinished products	66
10.5.1	Alert module	66
10.5.2	Self-service portal	66
10.5.3	Other projects	67
11	Discussion	69
11.1	Improvements over current system	69
11.1.1	Responsiveness	69
11.2	Interference reduction and effects	70
11.2.1	Interference usage	70
11.3	Adverse effects	70
11.4	Alternate measurements	71
11.5	Client station collection	72
11.6	Limitations	72
11.7	Related work	72
11.8	Future work	73
11.8.1	Personalized SSID system	73
11.8.2	Digital exam monitoring system	73

11.8.3 CleanAir management	74
11.9 Future plans	74
11.10 Evaluation	75
11.11 Future perspective	75
12 Conclusion	77
Bibliography	79

List of Figures

1.1	Access points with violating interference levels at UiT(as of 2016-05-25)	9
2.1	Clients using the wireless infrastructure at UiT(as of 2015-05-27)	14
2.2	Co-channel interference between two wireless cells, with a single client in the middle	22
2.3	Signal strength(RSSI), Noise and the Signal-to-Noise Ratio(SNR)	24
8.1	UTF-8 characters correctly collected from controller and displayed in the HTTP GUI view at UiT	55
9.1	Access points with violating interference levels at UiT(as of 2016-05-31)	57
10.1	List of some access points at UiT	63
10.2	Showing access point tf-2410-01-rw	63
10.3	List of some controllers at UiT	64
10.4	Main dashboard view at UiT	65
10.5	Alta dashboard view at UiT	65

List of Abbreviations

API Application Programming Interface

ASM Assembly

ASN.1 Abstract Syntax Notation One

CAPWAP Control And Provisioning of Wireless Access Points

CCI Co-channel Interference

CPI Cisco Prime Infrastructure

CSMA/CA Carrier Sense Multiple Access / Collision Avoidance

CTS Clear-To-Send

DCA Dynamic Channel Allocation/Assignment

DFS Dynamic Frequency Selection

DHCP Dynamic Host Configuration Protocol

DSSS Direct Sequence Spread Spectrum

DTLS Datagram Transport Layer Security

GUI Graphical User Interface

HA-SSO High Availability with Stateful Switchover

IEEE Institute of Electrical and Electronics Engineers

IoE Internet of Everything

- IoT** Internet of Things
- ISM** Industrial, Scientific and Medical
- JSON** JavaScript Object Notation
- LLDP** Link Layer Discovery Protocol
- MAC** Media Access Control
- MIB** Management Information Base
- NAT** Network Address Translation
- NAV** Network Administration Visualized
- OID** Object Identifier
- RRM** Radio Resource Management
- RSSI** Received signal strength indication
- RTS** Request-To-Send
- RTT** Round-trip time
- SNMP** Simple Network Management Protocol
- SSID** Service Set Identifier
- TPC** Transmitter Power Control
- UiT** University of Tromsø
- VLAN** Virtual Local Area Network
- VoLTE** Voice over LTE
- VoWiFi** Voice over WiFi
- WCS** Wireless Control System
- XML** Extensible Markup Language



Introduction

In this project, a few of the important factors in large scale wireless computer network operation has been investigated to take a closer look on what can be done to ensure optimal operation, with a reasonable amount of efficiency so that a rational operational expense(OpEx) can be achieved. Every large organization which have locations where employees, customers, members, students, associates or partners gather, should deliver fast, reliable, available and easy-to-use wireless internet access.

In a connected world, where everyone and everything is being connected to the Internet, locations where large amounts of people gather, the average home or small business solution will not scale, neither in infrastructural design and architecture, in protocol and resource allocation, nor in administrative scale. Therefore, large scale enterprise grade wireless networks become more and more common, and the coverage and service level expected by users become greater and greater. However, the step from small-scale wireless network installations based purely on acceptable coverage, to high-density, seamless, and reliable wireless networks often carry the need for huge investments in physical and logical infrastructure. Everything from structured media wiring and increased electrical installations, to more advanced wireless radios called access points and centralized controllers that handle the increase in traffic along with the need for much more advanced management of the radio spectrum resources available, is required. This process is often accompanied by a very steep learning curve riddled with hidden traps and pitfalls that even seasoned network engineers never could have anticipated.

To manage an ever growing infrastructure, a management and monitoring system capable of delivering relevant information to the engineers in a timely fashion is critical. Information about the radio environment around each access point must be collected and aggregated to detect potential problematic conditions as soon as they occur, and present them to engineers or operators within short time. The challenge of this is the fact that the amount of information needed to keep a complete view of the infrastructure, grows almost exponentially with the size of the infrastructure. This means that the design and requirements set forth for the management system must consider performance, while also ensuring that the cost of investment is as low as possible. A previously reasonable price of USD \$40 per access point, suddenly becomes an substantial investment when the size of the infrastructure grows from 10's of devices to 1000's of devices. Further, as the coverage area increases, so does the amount of users, and proportionally the amount of support cases. It is therefore important that common problems like failed access points, local power failures or potential problems with the radio environment be detected as soon as possible to avoid the first line of support being overrun by users. To ensure this, an efficient management system, capable of collecting the relevant information in real-time, or near real-time is important.

As such, it has as a part of this project and thesis been proposed if it is possible, using standard libraries, to implement a basis for a management system capable of scaling to the size required by an organization like UiT, and to see if the information gathered by the initial stage of development can be used in everyday management of the wireless infrastructure.

The basis for this project and thesis is the experiences had, challenges faced, lessons learned, mistakes done, successes achieved and notes taken during two years of experience doing large scale wireless network operation, design and expansion at UiT. Both the project and thesis are based on practical experience solving real-world problems and tasks through (sometimes) time consuming debugging, tedious manual operations and interesting dives into how a truly diverse jungle of different client devices and industry standard components work together. From this, several observations have been made with regard to what is important to consider, monitor and act upon when designing, operating and debugging a wireless network that should work for most users, all of the time.

As UiT is a large governmental institution in in Northern-Norway, it is also somewhat tasked with a regional responsibility with regards educating the future workforce of the region, advancing research areas and ensuring a strong local presence above the arctic circle. As such, it also acts as a regional authority when it comes to delivering services to the rest of the public sector and large public happenings where downtime or bad public experience might not be as

acceptable as it may be in a home environment. Examples of events where University of Tromsø (UiT) has had a strong presence with its wireless network infrastructure is the 2014 Chess Olympiad¹, and the annual Forskningstorget at Forskningsdagene² in Tromsø, in addition to delivering wireless coverage in public arenas like Kultursalen³ in Alta.

1.1 History

1.1.1 UiT

In the last years wireless networks have become more and more important in the everyday life of users of mobile and portable technology. As mobile devices become smarter and smarter, and continuously fulfill more and more uses in the lives of its users, the need for high-performance, reliable and available wireless networks become greater. Cellular network providers expand their networks with more advanced technologies to deliver higher speed, lower latency and more power efficient data transfer while on the move. Similarly, large public and private institutions see the need to expand their own wireless networks within their premises to deliver its employees, visitors, customers and associates with the ability to connect their devices to a local network capable of not only offloading the cellular network and reducing telecommunication costs, but also providing fast, reliable access to the Internet, along with the same reliable and fast access to local services like printing, media sharing, file access, backup and restricted on-premises services.

For UiT, this has meant that the traditional wireless network built over the last 10 years, based purely on priority of areas, and borderline acceptable coverage, has all but been replaced by a state-of-the-art, high performance wireless network. Using the newest technologies, and focusing on delivering excellent coverage(signal strength) in all areas where this is expected, along with radio-technical capacity improvements in high and very-high density areas like auditoriums, public areas, libraries, study halls, social and cultural arenas. In addition to the areas owned and/or rented/leased by UiT, the student welfare organization has hired UiT to begin deployment of wireless network in its student dormitories, apartments and other areas.

This expansion has meant that in few years, the wireless network administered by UiT has grown from approximately 60-70 wireless access points deliver-

1. <https://chess24.com/en/olympiad2014>

2. <http://www.forskningsdagene.no/>

3. <http://www.altakultursal.no/>

ing spots of wireless coverage, to approximately 2200 access point delivering full area coverage, with significantly higher theoretical throughput, and much higher client/user count capacity.

In addition to infrastructural improvements and expansions, the areas administered by UiT has grown significantly, both in the amount of area, and the geographical distances. In 2013, UiT and Finnmark University College merged to form UiT The Arctic University of Norway (usually shortened to UiT). This meant that UiT no longer had students and employees in the immediate vicinity of Tromsø, but also in Alta, Kirkenes and Hammerfest. From a management perspective, this makes it harder to effectively manage the infrastructure in place at each location, as the distances between them become too great to just stop by to check something. In addition, it is not viable to have local representatives from the IT-department on the smaller locations. This created the need for management tools that should give the administrators located elsewhere an effective and relevant overview of the local conditions on each location.

At that point, UiT currently used Cisco Wireless Control System (WCS) to monitor and gather information from the current wireless infrastructure. This system was primarily based on the principles of wired network monitoring and management⁴, and focused on rudimentary aspects like access point reachability, frequency and transmit power, client count and controller status. During the fall of 2013 and early 2014, the next generation of network management tools, Cisco Prime Infrastructure (CPI)⁵ was being launched, and subsequently deployed at UiT⁶.

This new tool focused on merging the wired and wireless network management into one tool, capable of both with the goal of looking at all technologies as "one network". Additionally, a new feature gradually being introduced were the northbound APIs in this new product, making it possible to develop own systems based on the information potentially available through this API. However, this soon turned out to be an unfinished product, with many features lacking, or not working as intended, which meant that most of the management still had to be done directly on the respective wireless controllers, with the additional work this introduces with regards to consistency between controllers. As the subsequent versions of CPI was launched, these introduced more and more working features, while maintaining the "one network view". Even with the first version, the virtual machine requirements [2, Table 1] [3, Table 1] were substantial, and made this virtual machine appliance the most resource demanding in

4. At that point known as Cisco NCS (Network control system)

5. <http://www.cisco.com/c/en/us/products/cloud-systems-management/prime-infrastructure/index.html>

6. as of version 1.3

the UiT virtual machine infrastructure, with a provisioned capacity of 16 non-hyperthreaded cores, 24GiB of RAM and 1200GiB of high-speed, low-latency disk storage.

Experiences with version 1.x and 2.x deemed slow and unresponsive, while still drawing substantial amounts of the provisioned resources. With the release of CPI version 3.0, some of the responsiveness aspects improved significantly. As it were, UiT were the first Cisco customer in northern Europe to deploy version 3.0, mere hours after being release, due to the need for management support for the newly acquired 8540 wireless controllers.

On 2016-01-01, UiT merged with Narvik University College and Harstad University College, to incorporate them into UiT The Arctic University of Norway. As a part of the merge, the existing wireless networks on both campuses, along with associated minor locations were, or is to to be replaced, with a total amount of $\approx 350 - 400$ access points on campus(excluding The Arctic Student Welfare Organization of Norway). As a part of the merge, it was decided that the IT-departments in Narvik and Harstad should not be responsible for managing the wireless network, only the physical mounting of access points and wiring in their respective locations. The management of the wireless networks at UiT is mainly to be done from Tromsø, with some local management in Alta.

1.1.2 Previous work

During the fall of 2015, a capstone project was completed to take a closer look on the current wireless management system, how this could be improved, and some of the metrics that currently weren't supported, but was available. This capstone project laid the grounds for this master thesis, and was called HiPerWA(High Performance Wireless Analytics), and focused mostly on how and what information and statistics that could be collected to make informed management decisions for the wireless network infrastructure.

1.2 Motivation

The motivation behind this project was rooted in the problems, challenges and needs experienced in everyday work running the wireless infrastructure at UiT, with experiences gained throughout the process of expanding the wireless network, both in capacity, and geographically.

During the process of expanding the wireless infrastructure, several aspects of management, provisioning and maintenance have required large amounts

of manual interventions. These tasks have often been related to correlating data from different sources, comparing data or doing seemingly menial work that should be operated. It is therefore in the interest to comprise some of the experiences and knowledge into a system that not only covers the essential metrics typically found in a management system, but also supports advanced analytics, adjustments and decision making.

A well-designed system can be expanded upon to support functionality to automate or aid processes like provisioning, zero-touch deployment and assurance of return on investment.

1.3 Problem description

The current management and monitoring system in use at UiT is slow to collect data, very little responsive on most platforms, and is lacking in its ability to be customized to account for the multi-campus model at UiT and the variety of configurations in the infrastructure. In addition, the virtual appliance running the current system is extremely resource demanding for something that at first glance seems like a lightweight task. The average collection interval of access point information in CPI is approximately 15 minutes, depending on system backlog, load and periodically running processes. This means that information about the status of each access point, client count, radio metrics and controller status potentially is delayed 15 minutes. In real time and production critical applications and situations, this is in many cases considered unacceptable. As an educational institution, tasks like digital exams, digital education through interactive and responsive lectures and the ability for students to rely on the connectivity provided by the network is important. Therefore it is important for administrators to be able to quickly detect or be alerted situations of interest and be able to act upon them as quickly as possible.

As the current system aims on providing every tool for managing the network infrastructure, the user experience is quite poor. This is primarily due to the lack of responsiveness, where pages with information uses a very long time to load, and Extensible Markup Language (XML) and JavaScript Object Notation (JSON) Application Programming Interfaces (APIs) have an experienced response time of 400-1200 milliseconds, making it hard to develop external services that can be used without another layer of caching or delay. This problem can temporarily be remedied through total system reboots, which ranges from 30 to 50 minutes in time, from the reboot is initiated, to the command line, HTTP GUI and GUI is available. After a reboot, most subsystems seem more responsive and the consumed host system resources decrease dramatically for a while. This however is not permanent, and over the course of a couple of

days, the lack of responsiveness and resource consumption is back to its normal high.

Another aspect where there is a dire need for improvement is the lack of available or easy-to-use data APIs from the existing system[4]. As the architecture and layout of the existing system is designed to cater to every aspect of network management, from authentication and accounting, to automatic configuration archiving and auditing[2], the useful information pertaining to wireless technologies and monitoring is either lacking or very hard to aggregate without extensive polling through pagination and sorting, based on the global namespace holding all device configurations. As the API is an integrated part of the HTTP GUI, it is also haunted by the same responsiveness issues and latency of information gathering.

A similar, open-source system, Network Administration Visualized (NAV)[5] is already in use at UiT and most other educational and health institutions in Norway⁷. This system, with its sub-services is responsible for monitoring and reporting for the entire IT-infrastructure and a range of building automation at UiT. It is in the interest of the people responsible for the wireless infrastructure at UiT, along with the incident team at the IT-department to be able to check the status of the wireless infrastructure alongside the other critical parts for its operation. The current version of NAV does not contain any features to monitor a wireless infrastructure apart from individual ping-response measurements for devices added. As UiT administer approximately 2200 access points, it is not viable to set up monitoring of each of these access points the same way one monitors the wireless controllers. Additionally, as NAV primarily is designed for basic Simple Network Management Protocol (SNMP) monitoring of network equipment, it lacks the support for monitoring specific services of wireless controllers.

Implementing a module for NAV using the existing APIs from the existing system has been attempted, but efforts were abandoned due to the high latency and slow responsiveness. Until now, there has not been a suitable library for adding this functionality directly into NAV, as most libraries are developed by the manufacturers of the wireless equipment, and is integrated into their own management systems (like the current system in use at UiT), and is either proprietary or not compatible with third party systems.

One of the most resource consuming tasks when deploying wireless networks is provisioning of the access points. The existing system supports basic deployment templates where static options can be set on a selection of access points, but this often requires manual administrator intervention to set specific fields

7. See <https://nav.uninett.no/wiki/navusers> for some of the users

like location, name and custom interface configuration. As access points are mounted by contractors, interns, local representatives from the IT department or other personnel not familiar with the normal operation of the access points, access points may not be connected correctly, and therefore may not connect to the wireless controller at all. As the wireless controller has no relation to whether or not it is missing an access point, there is no way of knowing if an access point has failed or has not been connected yet without an overlay system. The current system does not support batch-based provisioning, where lists of access point can be checked against the set of already provisioned, and therefore tracked access points. This means that it is up to the administrator to ensure that all access points have been provisioned, and put into service.

1.3.1 Interference issues

One of the challenges one faces when building high-density wireless networks, is the fact that there is a limited number of non-overlapping frequency slots(channels) available. This means that when the number of wireless access points increase beyond a certain threshold, depending on the environment, one risks having more than one access point on the same channel, and within normal reception range of each other. This means that they then can interfere with each other(see Subsection 2.3.2 for explanation of interference), and has to account for each others traffic when trying to send traffic of its own.

This means that the overall throughput of the wireless network in an area actually may decrease when additional access points are added. To avoid this, modern access points and clients are capable of adjusting their transmit power dynamically during communication, to ensure that their signal only reaches as far as the intended destination with acceptable strength, thereby trying to avoid to cause unnecessary interference. However, as the 2.4GHz Industrial, Scientific and Medical (ISM) spectrum in in most parts of the world only have 3 non-overlapping channels[6], the probability for interference is much higher. Additionally, the reach of 2.4GHz signals is much further than the same original strength 5GHz signal[7, 2.9][8]. When deploying access points in high density areas, this means that even at the lowest transmit power available, the access point is audible for other access points and clients nearby, while obtaining a more adapted signal strength on 5GHz, while also avoiding interference due to a larger amount of available channels.

At UiT this problem is noticeable in many areas, and preemptive measures have been done in the largest auditoriums, where 2.4GHz radios have been deactivated in all but 3 access points, so that there is no possibility for channel overlap. In other areas, interference due to high density deployments are still present, with an average ≈ 100 access points reporting interference levels higher than

the recommended maximum(20%) threshold, as seen in Figure 1.1. The figure also shows that that the problem consists of both persistent interference in the form of static access point-to-access point interference, but also interference generated when there is client traffic. The inter access point interference is often generated by 802.11 beacons sent regularly from access point, announcing available networks(Service Set Identifiers (SSIDs))

To alleviate this, and contribute to an overall better experience for users, clients are encouraged to connect to the 5GHz radios, both actively through mechanisms like Band Select⁸, or more passively by generally delivering a stronger signal strength on 5GHz due to less limiting of transmit power to avoid interference. This seem to work, with 5GHz usage reaching an average of $\approx 66\%$ during mid-day peaks, and nearly 90% in auditoriums⁹. Still, the amount

To decrease the amount of interference seen in the infrastructure, some of the 2.4GHz radius must be turned off, to decrease the density of active radios and free potential air time. To ensure an acceptable client experience, interference hotspots where the center a cluster of access points must be found, so that maximum interference reduction can be achieved, with minimal cost to the overall capacity. As noted above, the interference both consists of persistent interference present through both the night and day, and a varying degree of interference due to varying amount of client traffic, with a low point during the night. It is believed that if the persistent interference can be decreased, there should also be increased available air time and the intermittent interference should also decrease.




Figure 1.1: Access points with violating interference levels at UiT(as of 2016-05-25)

Source: Bjørn Johansen, Base Services/Network, IIA, UiT

8. https://documentation.meraki.com/MR/Radio_Settings/Band_Steering_Overview

9. Measurements done using the code developed in this project

1.4 Challenges

1.4.1 Data mining

One of the major challenges when constructing a system for collection of information from industry leading manufacturers is the availability of reliable, up-to-date documentation of where to find the relevant information, and how it should be harvested.

As most enterprise network equipment today make system information available through SNMP[9], it is necessary to acquire documentation for the interesting and relevant SNMP Object Identifier (OID) trees. For Cisco systems, this currently consists of 1149 public¹⁰ text files that describe the publicly available interfaces that a Cisco device may respond to over SNMP. Some of these interfaces have been added through company acquisitions done over the last 15 years, like the Airespace interfaces[10]. This means that to find one specific data point, substantial search efforts may have to be done to find the correct source of the information.

Additionally, there are proprietary or undocumented interfaces that either have been deprecated, not documented, internal for Cisco usage or that should not be known to customers. These interfaces should still be available, but may pose challenging to find the definitions for and where the root of their tree is located.

Data formats

As SNMP was designed over 20 years ago, it comes with a range of potential limitations. One relevant limitation in Northern Europe is the support for Nordic letters and symbols, that may be supported in the GUI parts of the devices, but may pose a challenge when being encoded into SNMP Abstract Syntax Notation One (ASN.1) format.

Additionally, to account for quirks in the implementation of both SNMP libraries, and the acsnmp implementations in the devices themselves, it may be required to insert some kind of sanity checks and format conversion and washing before using the collected data fields in analysis, in exported information or presentation to users.

10. <ftp://ftp.cisco.com/pub/mibs/v1/>

1.4.2 Potential bottlenecks

As this kind of potentially large-scale data collection from the infrastructure at UiT has not been attempted before, there may be systematic bottlenecks that may become apparent when the collection of information starts. Previous experience with collection of information from switches and printers at UiT has shown that some devices, or specific software or firmware versions may be vulnerable to large scale information gathering, or that specific branches of acsnmp trees may lead to infinitely deep recursion to find the end leaf. This has then caused the system load of these devices to reach levels where it has impacted its primary tasks like packet forwarding or printing, which is unacceptable for critical infrastructure like a wireless network.

To ensure the risk of adverse effects on the production infrastructure, several tests has been performed to see if the infrastructure is affected by collection without any limiting. Thus far, no notable effects has been noted.

1.4.3 Client device support limitations

One of the largest barriers to completely ending the use of the 2.4GHz frequency band for regular, client 802.11 traffic, is the fact that there are still relatively popular client devices being developed and sold that does not support a pure 5GHz infrastructure. Devices like Apple Watch[11, p. 37], Raspberry Pi 3 does not have a 5GHz radio. This means that a 2.4GHz infrastructure still has to be present to serve the needs of these devices, and a complete abandonment is not possible for the foreseeable future.

1.4.4 Legal limitations

When it comes to collection and storage of information, especially information about personal devices, one has to take special care that this information is not disseminated or spread to third parties. Additionally, information that can be used to identify persons, organizations or other parties should not be stored longer than absolutely necessary.

To achieve this, information about end user devices should as far as it is possible for the goals in this project, not be collected. If it is necessary to collect this kind of information, it should happen in dialogue with the legal department at UiT, and the information should be limited to the bare necessities.

/2

Background information

2.1 Enterprise wireless networks

An introduction to some of the fundamentals for enterprise wireless networks can be found in [12, 2.4]

In comparison to most home and smaller office wireless networks, large scale or enterprise grade wireless networks consists of multiple wireless access points that contribute to a singular logical network. This means that clients connected to the wireless network may move from access point without the user needing to intervene, and without reconfiguring settings to account for talking to another physical access point. The act of moving from access point, while still maintaining the connection and state is called roaming.

In addition to consisting of multiple access points, enterprise scale wireless networks must support a much larger amount of clients, in the scale of 100 – 100000 clients. This means that the physical infrastructure, along with the management tools used to administer the network must support the same amount of clients, along with its proportionally large amount of traffic. At UiT, a theoretical maximum of approximately 118 000 clients and 10000 access points are supported, with a current access point count of ≈ 2200 , and an average 3200 clients connected at any time, with peaks of ≈ 9000 clients. This capacity has been geographically distributed across the major campuses in Tromsø, Alta, Harstad and Narvik.

When a wireless infrastructure reaches this size, measures are usually taken to ensure that the infrastructure is working as intended, and without systemic faults the entire time, with very high availability.

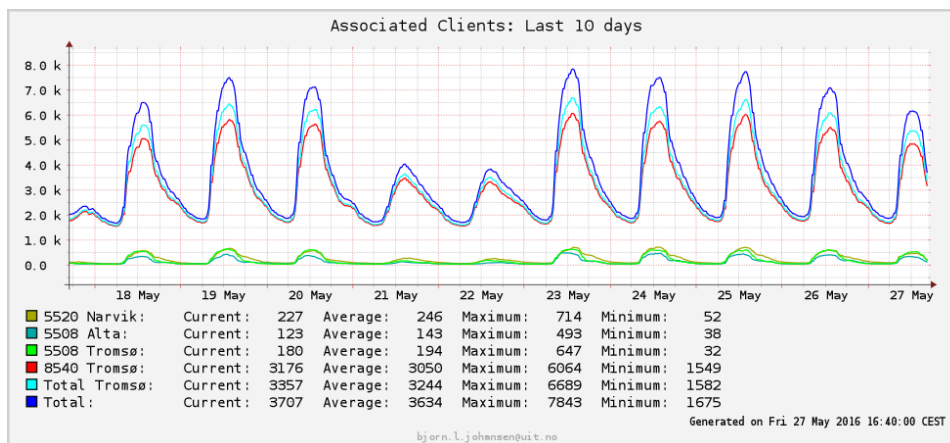


Figure 2.1: Clients using the wireless infrastructure at UiT (as of 2015-05-27)

Source: Bjørn Johansen, Base services/Network, ITA, UiT

2.1.1 Common Architectures

Most large enterprise scale wireless networks today consists architectures with central controllers and wireless access points carefully placed across the areas where one wants wireless coverage. The architecture with simple access points with little or no pre-configuration required simplifies deployment significantly over traditional access points where each access point would have to be configured with at least some kind of authentication for remote management before being deployed. Instead, lightweight access points are connected to the network and through information like Option 43¹, DNS-search domain or broadcast would discover one or more wireless controllers to control it. In this design, one or more central controllers form a mobility domain, where clients seamlessly can roam between access points, which in turn is connected to one of the controllers. When roaming between access points, the authentication state, along with other parameters are kept, and in most cases the same IP-address is kept, with minimal noticeable packet loss incurred. This means that the higher level protocols and application for the most part is oblivious of the roaming. Roaming within a single mobility domain also means that a client may move its traffic between controllers in relation to which access point it is currently talking to, with the client being oblivious to

1. <http://www.cisco.com/c/en/us/support/docs/wireless-mobility/wireless-lan-wlan/97066-dhcp-option-43-00.html>

this fact, and invisibly from the user.

Depending on the size and geographical expanse of the network infrastructure, one or more mobility domain form the entire wireless infrastructure of an organization. As an organization like UiT may have multiple locations around a country or even on multiple continents, the network latency between locations may cause communication between a location and a central location to be costly. In this case, a separate mobility domain is advisable, as client information will in most cases only be useful at the location, and live roaming between mobility domains are very rear. Additionally, it is preferable for the client traffic to and from local and internet resources to enter into the normal network infrastructure as soon as possible in the topology, to avoid additional, unnecessary round trip time and extra network hops.

With both single and multiple mobility domains, a client may have a primary anchor controller assigned, to ensure an authoritative source for authentication and statistics collection. This anchor controller may be responsible for coordinating roaming between controllers or mobility domains, and may facilitate the authentication and association of a client in a foreign mobility domain, effectively simplifying the configuration and operation of multiple mobility domains. The anchor controller can also help the client to obtain and use an IP-address from its home campus, which helps the client to reach resources within its own home base campus.

This also makes it possible to build redundant, high availability networks, where the access points can be moved between controllers, and even mobility domains if necessary. Additionally, it is possible to include geo-separated clustered controllers which acts as a single device, while being situated in different geographical locations.

An alternate design used by smaller organizations with many smaller locations, or organization without the capacity or infrastructure to run centralized designs is a combination of mesh and cloud based infrastructures. These installations consists of access points which for the most part act as their own miniature controllers. In such a mesh network, a single access point acts as the controller for its nearby access points, and in case of the primary controller access point failing, another of the access points in the mesh takes seamlessly over. This mesh network may in turn be administered remotely using a cloud service, with the network administrator never having to touch the access point.

2.1.2 Access Points

Access points in an enterprise wireless network usually act as the local bridge between the wired infrastructure and the wireless domain. The main task of the access point is to be responsible for the wireless communication, with 802.11 protocol operation, encryption, and local radio resource management.

Depending on the architecture, the access point either forms an encrypted tunnel to its controller to transport management and data traffic back to the controller, where the client reaches the rest of the network and the internet, or may exit some or all of the traffic locally at the local network of the location. This makes it possible to place access points in adverse or hostile environments.

The local radio parameters of the access points are adjustable with regard to frequency/channel and transmit power, which is adjustable on-the-fly to adapt to a changing radio environment and clients. Depending on the access point model, it is also capable of doing local environment analytics to record and identify noise, detect interference or other radio resource management tasks. The information gathered may either be aggregated and acted upon locally, or forwarded to a controller or the rest of the mesh infrastructure for further usage.

When an access point is connected to a wireless controller, it can be provisioned with the correct host name, its location in addition to which controller it should regard as its primary controller and strive to remain connected to. Some access points also support bridging of additional wired interfaces so that a single uplink can be shared between wired clients and the access point. In this case, each wired interface of the access point has to be provisioned with regard to VLAN tagging and whether or not it can be used for client traffic. Some access points also support supplying power to other equipment over the wire, in example IP-phones or local sensors or small displays.

2.1.3 RRM

One important function to have in an enterprise wireless network, regardless of size is the mechanisms that can be called Radio Resource Management (RRM). RRM consists of gathering live and historical information about the environment around each access point, and through analysis, statistics and aggregation make the best informed decisions about the frequencies and transmit powers in use. This can be to avoid interference between access points in the same infrastructure, avoid intermittent noise or disruptions, or more cleverly assign the different frequencies in (very) high density areas where the amount

of non-overlapping channels are too few, and random assignment cannot be used.

2.1.4 CAPWAP

The Control And Provisioning of Wireless Access Points (CAPWAP) protocol[13][14] is a standardised, extendable networking protocol designed for intercommunication between lightweight access points(sometimes referred to as *wireless termination point*) and wireless controllers(sometimes referred to as *access controller*). Its primary goal is to serve as a vendor-independent, omni-present protocol that handles all aspects of communication between an lightweight access point and a wireless controller infrastructure.

Upon powering up, an access point queries the network by either using information served to it using Dynamic Host Configuration Protocol (DHCP), via DNS or by broadcasting a discovery packet[13, 4.5],[15, 5.3] to the local broadcast domain(physical/Virtual Local Area Network (VLAN)) to collect a list of viable wireless controllers. When a list of controllers has been gathered, the access point tries to join the controller with the highest preference set. If the access point has been configured a specific controller previously, or it has been in association with one of the controllers in the list before, this controller will be configured.

When the access point has associated with a controller, it establishes a Datagram Transport Layer Security (DTLS)[16][17] CAPWAP tunnel with the controller. The tunnel is usually split into two main parts, a control path and a data path, where the data path in practice can be multiple layer-2 tunnels(VLANs). Depending on the capabilities of the access point, parts or whole of the tunnel is encrypted to ensure that privileged configuration information and potentially traffic is not subject to man-in-the-middle attacks between the access point and the controller. When the tunnel has been established, the access point is configured and provisioned with the parameters the controller requires to be present for the access point to be in compliance. This can range from specifying regulatory domain and allowed channels, security parameters, secondary and tertiary controllers, to the individual SSIDs it should broadcast and where client traffic should be directed.

During normal operation, the wireless controller and the access point are in regular contact to relay management and control information. For access points capable of collecting information about the radio environment or do local analysis or security tasks, this information is also sent back to the controller with regular intervals[18].

The use of CAPWAP tunnels enable access points to be positioned anywhere in the network, even at locations geographically remote, as long as it is possible for the access point to reach the wireless controller either via IP-routing or through the same broadcast domain. As the CAPWAP tunnel is capable of transporting layer-2 traffic, clients can be assigned IP-adresses from the same IP-pool, even though they may be located thousand of kilometers apart. This makes it possible for employees, students or associates to bring a personal access point with them while traveling or to their home, and they will not only be able to use the same wireless networks regardless of where they are, but also to obtain an internal IP-address and have their traffic securely transported back to their home institution, where it can exit into the local network to reach local resources, or exit to the internet. As the CAPWAP protocol is designed as an extensible protocol, additional functionality like transport of wired traffic[19], making it possible for employees to have a transparent, on-the-go VPN-service with them.

If the CAPWAP tunnel between the wireless controller and the access point breaks down, the connectivity for control traffic, and potentially data traffic is lost. This means that, depending on the configuration of the access point, SSIDs may be deactivated and clients disconnected, or the access point may work in stand-alone(FlexConnect² mode). For completely lightweight access points, it is therefore important that the CAPWAP tunnel is kept alive, to ensure an acceptable client/user experience.

2.2 Radio Resource management

2.2.1 DCA

One of the most useful, and important features in modern radio resource management for wireless networks, is the possibility for Dynamic Channel Allocation/Assignment (DCA). DCA is the ability for a wireless network infrastructure to change its broadcasting(send and receive) frequency setting during operation, based on various parameters. These parameters may consist of regulatory or regional restrictions on the use of certain frequencies with regard to other users of the same frequency, transmitted power, actual duty cycle or utilization or some local considerations that may intermittently or continuously be present.

As the use of frequency bands is regulated by individual entities on different

2. http://www.cisco.com/c/en/us/td/docs/wireless/controller/8-0/configuration-guide/b_cg80/b_cg80_chapter_010001000.html

countries and regions, the available spectrum usable for wireless networks is different in the different parts of the world, with the possibility of some frequencies being available in Europe (and Norway), while not available (legally) in the US. While this may not present a problem for small organizations with few or one locations, this may become a challenge for larger institutions like universities or large corporations which may have offices or campuses in different regulatory domains, or may have employees, students or researchers travelling across borders between the domains. This means that there may be a need for the DCA mechanisms to take the location of the access point (radio) when doing the assignment of the channel. Another related consideration is the restriction of transmitted power on certain frequencies on some areas, where the allowed maximum or average transmitted power may be different depending on whether or not the radio is positioned inside a building, or outdoors, as this may impact the dispersment of the signal, with typically much lower dispersment inside buildings, and very low leakage out of the building.

Further, some frequencies may carry usage restrictions with regards to other users of the same frequency spectrum. The most common example of this is radars used for weather observations, aviation, naval and marine purposes, and military applications. These radars usually operate within certain frequencies in the 5GHz spectrum which in some regulatory domains (including Europe/Norway) also may be used for wireless networks, among other uses. This spectrum originally allocated for use with radars is referred to as the Dynamic Frequency Selection (DFS) channels [20]. To avoid interfering with radar systems, and avoid disruption of the wireless network, certain mechanisms have been put in place to ensure proper operation. As radars can be identified by wireless radios by their distinctive bursts of energy, quick measures can be taken to minimize mutual interference, and future disruption. Upon detection of a radar signal, an access point will change its operating channel as soon as possible, and if possible inform its associated clients (via 802.11h) that a radar signal has been detected and that a change of frequency has been initiated. This enables the clients to change to the new frequency at the same time as the access point, minimizing the disruption. As the energy contained in some radar bursts may be much higher than the maximum allowed energy of a wireless network station, this also protects the radio hardware of both clients and the access point from damage. In addition to the DFS/DCA mechanism, there are also signal strength restrictions on some of the frequencies that may be used by radars, to avoid disruption of low-power radars that may not be detected by the access points themselves.

The most important use of DCA is in noise and interference avoidance. As the frequency spectrum allocated for wireless networks is unlicensed spectrum, wireless networks are not the only users of this spectrum. This means that

in a large infrastructure, there is bound to be a number of devices that does not conform to the 802.11 specification that uses the same frequencies as the wireless network. This means that it may be hard or even impossible for a wireless network to coexist on the same frequency as a non-802.11 device. The solution for this is therefore often for the 802.11 capable equipment (access points and clients) to change its operating channel to avoid the noise. This is done by the use of DCA, where the access point may listen periodically on all channels to determine the optimal channel to use in its location.

When it comes to interference avoidance, the same principles as for noise avoidance applies, however it is much more likely for several 802.11 networks to coexist on the same channel/frequency without major impact on performance and operation. However, when it is possible it is still recommended and practiced that DCA ensures that a unused or little used channel is chosen. For access points in the same wireless infrastructure, the final channel assignment may be a result of a computed optimum channel plan for an area, taking into account several access points to ensure a best-fit with the regard to all DCA parameters.

2.2.2 TPC

Transmitter Power Control (TPC) is the ability for wireless stations to adjust their transmit power depending on their environment, local or regulatory restriction, or other factors that may impact coverage, reliability and experienced or generated interference.

Depending on the architecture of wireless access points, most enterprise wireless access points are manufactured with 4-8 predefined transmit power levels ranging from -3dBm to 17dBm. The intention of this is for wireless network operators or RRM systems to be able to dynamically change the transmit power of access points to ensure that the best coverage is given for the intended area, while also ensuring minimal interference with surrounding access points. This also enables access points to deliver the same coverage/range on both 2.4GHz and 5GHz and avoid clients unnecessarily sticking to a single access point, and for network designers to a degree use the same radio coverage patterns for both frequencies.

2.2.3 Air Quality

Air Quality within wireless networks is a term used to describe the quality of a certain frequency range. Not unlike ordinary air quality this is an indicator of the amount of clean air in a sample, compared to the amount of pollutants. The

scale of air quality is an inverse scale, with near perfect quality being 100(%), and "undiluted" pollution being 0.

When calculating the air quality, the algorithm starts with a base of 100, and for each pollutant the proportional degree is subtracted. The degree of pollution from a pollutant is a product of its consumed air resource(time or frequency range width), and signal strength, from which signifies the overall impact. In other words, a Bluetooth signal with high air time usage, but very low signal strength does not have a large impact, as the signal is not heard by a significant part of the wireless cell.

2.2.4 Other measurements

One of the potentially available measurements that may tell something about a wireless cell alone or a collection of cells in a wireless network infrastructure, is the amount of clients that is heard with a low or poor SNR-ratio.

Not all access points in use in a enterprise network infrastructure may have the capability or the hardware for doing advanced detection and classification of noise. For these access points, looking at the SNR-ratio of individual, associated clients may tell something about the environment, without explicitly looking at the noise level.

In Cisco-based wireless network infrastructures, it is possible on some controller models to collect an integer number per access point with the number of associated clients which fall below the threshold for poor SNR-ratio. If the client at the same time has an acceptable Received signal strength indication (RSSI), it indicates an environment with much noise, and that there should be taken automatic or manual measures.

2.3 802.11

2.3.1 Protocol

The fundamentals of the 802.11 protocol, its mechanisms and how has been covered in [12, 2.2]

2.3.2 Interference

In 802.11 protocols interference is considered all radio energy that can be successfully demodulated as valid 802.11 communication, but is not intended for, or pertaining to the particular wireless cell it is being detected on. This can be neighboring wireless cells using the same channel, or a wireless station transmitting without pertaining to the otherwise strictly followed protocol. It is important to note that in 802.11 wireless networks, interference usually does not distort or alter the signals sent between the members of the wireless cell, it just inhibits their ability to initiate own traffic due to the Carrier Sense Multiple Access / Collision Avoidance (CSMA/CA) mechanisms in play, or occupies the exclusive transmission or reception equipment(antenna and radio).

Interference between wireless cells is often referred to as Co-channel Interference (CCI) and is interference in a station-to-station(client station to client station, or client station to access point) relationship(cell) on a specific frequency or channel. This occurs when two or more cells can detect each other on the same channel, due to limited number of available channels, high transmit power, or too close proximity. The result of this is that the CSMA/CA mechanisms are rightfully, but excessively triggered, as the cells more or less contend for the same air time and need to take each other into account before transmitting, in addition to "ignoring" demodulated. This may lead to lower overall throughput and in some cases connectivity issues between stations, as high CCI may cause the stations to be unable to contact each other within the necessary interval of time.




Figure 2.2: Co-channel interference between two wireless cells, with a single client in the middle

Source: Wireless LAN Design Guide for High Density Client Environments in Higher Education[20]

2.3.3 Noise

Noise within 802.11 wireless networks is defined as any detectable energy that cannot be decoded as valid 802.11 protocol communication, foreign or pertaining to one's own cell. As such, any source of radio energy with a frequency within the frequency range of the channel of a cell is considered noise. Depending on the data rate, channel width, sub-carriers in use, and the bandwidth of the noise source, the impact may vary.

In comparison with Interference sources that can be measured to a given signal strength (RSSI) at a given point in the cell, the signal strength of noise sources is much harder to quantify. This is because the energy pattern, duty-cycle and band-width of noise sources vary, and therefore is hard to compare to each other. Further, noise sources may not be present at all times, or may follow irregular transmit patterns that makes them hard to quantify. Therefore, noise sources have traditionally been compounded, and only been considered as a part of the environment noise-floor (see Subsection 2.3.4), which describes the average (over time and/or width of energy signature) level of energy in the environment that is considered noise.

In recent years, commercial manufacturers have begun delivering products that employ advanced, digital signal processing chips to detect, analyse and classify noise sources to be able to anticipate how the source may behave, and to avoid being noteworthy affected. Based on what kind of noise source it is, how the signal is shaped (width and power across the width) and the received signal strength (RSSI), an Air Quality rating is made.

Noise sources can be divided into two main categories, active transmitters, and passive transmitters. Active transmitters are devices that intentionally use the same frequency spectrum for its own use case, which may be everything from surveillance systems and building management systems, cell phones and competing data transmission systems, to microwave ovens and induction chargers. Passive transmitters include all kind of transmitters that transmit in the same frequency spectrum, but without it being intentional or designed for it. This may be due to poorly executed electrical design, cheap manufacturing, or accidental design flaws. Typical units in this category include transformers, Christmas lights, high voltage contactors or high frequency AC-voltage lines or wiring with externally inducted signals. One of the most famous examples of accidental design flaws is the case of USB 3.0 data transmission and its effect on Bluetooth and 802.11 wireless cards[21].

2.3.4 Signal-Noise Ratio

To describe the quantified relation between the useful signal and environmental noise at a given point in a wireless cell, Signal-to-Noise Ratio(SNR) is used. As the noise in a area is compounded to form a total view, called the noise floor, one can measure the distance between the signal strength of useful signal and the noise floor. The distance is a measurement of the proportional relationship between the signal and the average noise level. With the use of a logarithmic scale like dB, the difference in power is expressed as a multiple of 10 per 10 dB. Figure 2.3 shows how the distance between the actual signal and the noise floor can be illustrated. If the distance between the noise floor and the signal




Figure 2.3: Signal strength(RSSI), Noise and the Signal-to-Noise Ratio(SNR)

Source: Wireless fundamentals: Signal-to-Noise Ratio (SNR) and wireless signal strength[22]

is too low(usually less than 25db[11, p. 37]), it becomes less and less likely that the modem in 802.11 stations are capable of demodulating the signal and decode the data without error.

2.3.5 Protocol Impact

As 802.11 is a CSMA/CA-based protocol, both interference and noise may, and will cause the total utilization of the available air time to decrease. As each transmitter participating in a cell must consider both interference and noise, and potentially wait before transmitting its data or management frames, the total throughput of the cell exponentially decreases as the amount of interference and noise increases. This is due to the fact that multiple members of the cell

may be affected, which may cause large amounts of traffic to be queued across the cell. When an available airtime slot opens, multiple parties may want to send their traffic, which causes additional protocol communication (Request-To-Send (RTS)/Clear-To-Send (CTS)) to be used. Further, retransmissions may be invoked from either the 802.11 members, or higher layer protocols due to the delay incurred, additionally increasing the amount of traffic. For some members of the cell closer to the interference or noise source, starvation may be observed.

If high amounts of noise is experienced, or repeated transmissions are corrupted, members of the cell may rate-shift to a lower data rate to ensure the transmission is received successfully. This causes proportionally higher usage of the available air time. In example, a station originally transferring data with 150Mbps, will use 3-4 times longer time to transmit the same amount of data if it has to rate-shift down to 54Mbps, due to fragmentation and overhead, essentially decreasing the overall goodput of the cell.

For higher layer protocols and users, this is usually noticed as experienced latency, with minor interference or noise present causing up to 10 times higher round trip time (RTT). This may not have noticeable impact on regular data transmission, as this can be, and is often masked with large segment window sizes and parallel downloading of data. However, when it comes to voice-based traffic, real-time streaming video or typical online games, latency, and especially varying latency causes large problems. A sudden 50-100 millisecond spike in round-trip time may be the deciding factor when dealing with first-person shooter games, or a similar degree of increase in latency may cause additional buffering, or may induce stuttering in the video or audio while streaming or performing voice or video calls. This becomes even more relevant when cellular carriers[23] start deploying services like Voice over WiFi (VoWiFi) and Voice over LTE (VoLTE), which is especially relevant for larger organizations with users located indoors. In this case, cellular operator and customers are reliant upon stable wireless networks with even latency.

It is therefore not the amount of data or the throughput which quantifies the quality of a wireless network, but rather its ability to transfer data efficiently without much delay and jitter.

2.3.6 Problem illustration

The problem with the 802.11 protocol when it comes to large scale, dense deployments is that the 802.11 is regarded as an extremely polite protocol. Due to its CSMA/CA[24] protocol, an 802.11 station must go through extensive collision avoidance protocols before sending data across the wireless medium.

As the frequency band between two or more stations is shared, the probability for other traffic to be present is large. To avoid collision, a 802.11 station will therefore wait for the air to be free before sending its data, and if the transmission is interrupted by another transmitter, the broadcast is halted, and retried after a little, random length halt period.

This creates a range of challenges in high density or very high density deployments, where the probability of two or more wireless cells sharing the same channel is large, and mechanisms to be able to share the same frequency band must be put in place.

2.4 Eduroam

UiT is a member of the Eduroam-providers organization, and offers all eduroam users in the world internet access through its wireless network infrastructure. All employees, students and associated personnel at UiT have access to eduroam, which is the main wireless network offered to users at UiT.

The way eduroam works, is that each user is identified by their username, suffixed with their membership organization, in example `username@uit.no`, which signifies the user `username` belonging at the institution `uit.no`. Each member institution is associated to a regional or national network registry which is responsible for authentication routing to and from the institution. As such, external users visiting `uit.no` will try to log on to the SSID at UiT, and their authentication request is forwarded to the regional or national hub for forwarding to the external user's home institution or regional hub.

The eduroam network is described in [25], but does not give any pointers towards which frequency bands and data rates to support, or not support.

/3

Goals

3.1 Improvements over the current system

The primary initiative for this project was the need for a new base framework capable of collecting arbitrary pieces of information from the wireless infrastructure at UiT, and present or store it for further usage by other systems, integrated or associated. Initial focus was on looking closer on the information being collected by previous efforts[12] and investigate if this information could be collected more efficiently, mainly with regards to information from access points.

An important requirement was that standardized libraries to be used as much as possible, to ensure easy platform portability. This was to ensure that the developed system framework would be compatible with other potential systems¹, in addition to the resulting source code being compatible to include in future efforts with the same goal.

As seen from the existing system in use at UiT, it was important that the resulting system was trivial to deploy and redeploy using minimal customization and effort. To ensure this, the system was to be designed to run as a standalone file structure, only utilizing local libraries and storage engines capable of trivial data export should the system have to be moved. This would also mean that the storage of system data would be possible externally from the system host

1. Mainly NAV[5]

itself.

The main goal of this project was to identify aspects of the existing system, in addition to new requested features that was needed or could deliver some useful functionality, and try to implement them in a way that could scale and deliver actionable information, within reasonable resource usage.

3.1.1 System integration and visualization

As mentioned, it is advantageous to build the resulting library or system in a way that is possible to include or couple to existing systems such as NAV. Additionally, to showcase some of the information collected, and make it easier for other parts of the management staff to keep an eye on the infrastructure, a rudimentary, but well designed graphical user interface should be developed. Along with this, it is possible to incorporate the necessary components for additional system integrations through interfaces like a RESTful[26, Chapter 5] and WHOIS[27].

3.1.2 Automatic (pre) provisioning

One of the more useful features that becomes possible with a more customized system, is the ability to create automatically provisioning software that can parse the already standardized forms in use at UiT for installing new structured media, to also provision access points as soon as they are installed. As the forms used contains not only the name of the outlet where access points are connected, but also the switch port where the access point is connected, along with the mac address of the access point, the switch port can be configured to the access point, the access point can be configured with correct name, and a location from the form.

In student housing, where access point with additional wired interfaces are used, these interfaces can be configured as soon as the access point is connected to the controller, and the access point can be added to the specific access point group for the student welfare organization to only deliver the specific SSID available for residents. In the future, services like personalized SSIDs may become available(see Subsection 11.8.1), where additional configuration may be necessary both on the access point, the controller and the surrounding infrastructure for this to work as intended.

3.1.3 Monitoring

When information has been collected, it should be possible to implement trivial procedural actions that can be run to periodically monitor and report the state of the infrastructure.

3.2 Interference Reduction

To utilize the new infrastructure installed at UiT, it is desirable to be able to reduce persistent interference as much as possible, both to gain a better radio environment for clients, to avoid unnecessary alerts and warnings and to ensure that the infrastructure remains robust and capable of supporting real-time and latency-sensitive applications.

To do this, it is a goal to create a mechanism to identify centers or hubs within clusters of high interference, and to take measures to break up or remove the cluster hotspots. By doing this, the impact of channel sharing or channel re-usage, with co-channel interference being the result, should be minimized.

By removing the larger interference hotspots, a step is taken towards removing the areas where multiple wireless cells overlap to create what is called the "Don't want"-zone[28].

3.3 2.4GHz radio shutoff

A side-effect of the efforts to remove co-channel interference, mainly on the 2.4GHz spectrum, is that this paves the way for the first steps towards removing the 2.4GHz spectrum as a user-occupied radio network. As more and more clients choose 5GHz SSIDs due to better RSSI-conditions, the usage of 2.4GHz will most likely decrease, even with the improved interference levels.

As such, it is a goal to investigate whether or not a discontinuation of 2.4GHz as a user-occupied network is possible in the near future. This can then be used in the upcoming efforts at UiT to enforce 5GHz as the primary wireless network standard, with a lesser-quality, limited service for legacy clients on 2.4GHz.

The plan, if this succeeds, is to make the 2.4GHz spectrum available for devices participating in typical Internet of Things (IoT)/Internet of Everything (IOE)

[29] devices that due to power constraints, or cost of production does not support nor contain a 5GHz 802.11 radio. As these devices does not require large quantities of data to be transferred, and may not need nor want to continuously transfer data, or require low latency, the 2.4GHz 802.11 network can be used for the uplink or the backhaul for their communication. This essentially creates the possibility for a two-tier network infrastructure, where minimal effort can be put into optimizing what many call a lost cause, and focusing on optimizing and planning for a 5GHz and future 60GHz[30] infrastructure.

/4

Data collection and storage

4.1 Development and Experimental Setup

The platform used for developing the source code, and running tests and experiments for this thesis, was the actual wireless infrastructure in use at UiT, spread across four major campuses, with access points spread across Northern Norway from Bodø in south-west, to Kirkenes in north-east. The following wireless infrastructure hardware is present at, or connected to the different campuses:

- Tromsø

Wireless controllers

- * 2× Cisco 8540 Wireless controller (in High Availability with Stateful Switchover (HA-SSO) redundancy setup)
- * 1× Cisco 5508 Wireless controller
- * 1× Cisco 2504 Wireless controller

– Access Points

- * Cisco 1131AG

- * Cisco 1142N
- * Cisco 1242AG
- * Cisco 1262N
- * Cisco 702W
- * Cisco 2702i
- * Cisco 3502i
- * Cisco 3602i
- * Cisco 3702i
- Alta
 - Wireless controllers
 - * 2× Cisco 5508 Wireless controllers(in HA-SSO redundancy setup)
 - Access Points
 - * Cisco 1131AG
 - * Cisco 1142N
 - * Cisco 2702i
 - * Cisco 2702e
 - * Cisco 3602i
 - * Cisco 3702i
 - *
- Narvik
 - Wireless controllers
 - * 2× Cisco 5520 Wireless controllers(in HA-SSO redundancy

setup)

- Access points
 - * Cisco 1142N
 - * Cisco 2702i
- Harstad
 - Wireless controllers
 - * 2× Cisco 5520 Wireless controllers(in HA-SSO redundancy setup)
 - Access points
 - * Cisco 2702i
 - * Various Trapeze/Juniper access points scheduled for decommissioning(not used in this project)

A total of 10 wireless controllers(6 logical units) and 2129 access points were used for the development and testing process of this project. All wireless controllers are connected to a 1 or 10 Gigabit redundant network connection, depending on controller hardware support.

To run the source code developed, a HP Z400 Workstation with an Intel Xeon W3550 CPU with 4 physical and 8 logical(HyperThreaded) cores running at 3.06GHz and 24GiB of DDR3 1066MHz RAM was used. As host operating system, Arch Linux¹ with Linux 4.4.1-2-ARCH and standard python 3.5.1 with standard libraries was used. For acsnmp communication, and some value encoding/decoding, pysnmp² 4.3.1 was used.

The computer was located in the primary datacenter of UiT in Tromsø, with average network round-trip time to all controllers of less than 6 milliseconds.

The current management system in use, Cisco Prime Infrastructure version 3.0 was located in the same datacenter, but in a virtualized environment powered by a [INSERT HERE]

1. <https://www.archlinux.org/>

2. <http://pysnmp.sourceforge.net/>

4.2 SNMP

Information from each controller, and by extension every access point was collected using acsnmp version 2c, which enables acsnmp libraries to request batches of information from a acsnmp engine(device). Due to security concerns, collection of information over acsnmp, or changing controller parameters over acsnmp is restricted to specific IP-addresses or subnets to avoid unauthorized access. As such, the machine used for collection was placed on a privileged network alongside the existing management system.

While some information can be collected and stored directly from acsnmp, other information is stored in ASN.1³ format. This means that not all fields are automatically translated by the acsnmp-libraries used, or may not be possible for the library do decode due to proprietary data formats in the underlying data. For this purpose, some wrapper functions were developed to try to "clean" and reencode received data before storing it.

4.3 Data format and namespace

The storage format used to store the information collected is JSON⁴ as this closely resembles the dictionary datatype found in python, and enables a structured, document based storage structure, where parts of the data could be stored as sub-values of a parent document. It was important in the early stages of development that the number, type and naming of fields were flexible, as different controllers, access points and other types of data stored may not have fixed fields, but may be dependent on model, type or availability of information

Further, for persistent storage **MongoDB**⁵ was used to store the information as documents. As the layout of documents stored in **MongoDB** is based on binary JSON(BSON), it closely resembles the layout of dictionaries in python, and bindings between the two are almost seamless. This enables very fast searching, projections and collection aggregations. As each access point, controller, or in the future client is stored as individual documents, the flexibility in what data

3. <http://www.itu.int/en/ITU-T/asn1/Pages/introduction.aspx>

4. <http://www.json.org/>

5. <https://www.mongodb.com/>

is stored and its naming is kept. When doing mining on the data collected, it is possible to make copies of the relevant documents, and write the resulting information back to the same document collection, or to a separate document collection.

The naming of field within the database, in the source code and by references are sourced from the various acsnmp Management Information Bases (MIBs) used to find the information. These names are potentially Cisco proprietary names, and may be changed at a later stage to account for heterogeneous infrastructures. The acsnmp MIBs are usually organized in deep trees consisting of lists or leafs, but when storing the relevant information in documents, a more shallow document tree structure has been used. It is assumed that the field names are somewhat unique within the scope of this system, and that it at the current stage in development does not pose a threat.

The main namespace for information storage has been set to `hiperwa`, with collections under this namespace for access points(`aps`), controllers(`wlcs`). Additional collection names are created and destroyed as needed. Future expansion with a permanent `clients`, a `rouges`, a `noise` and a `analytics` is planned.

4.4 Storage

As mentioned above, the storage of information has been done in **MongoDB**. Within the database engine, a `hiperwa` namespace has been created.

This makes it possible to store information on a different machine than the actual collection of data, without having to consider file storage. Additionally, it is planned to implement data-level redundancy and backup on the database level, with potential horizontal capacity expansion if necessary.

4.5 Extensions

If permitted, a port of the source code developed earlier[12] is to be done to be able to collect client information and store this in the same dataset as access point and controller information. At the current stage, this information is still very strictly regulated, and subject to periodical pruning to avoid having the

potential of collecting to track individual users. The data retention capabilities⁶ in MongoDB makes it a prime candidate for ensuring data about clients and users is not kept for longer than allowed. By setting a time-to-live value on whole or parts of this kind of data, the privacy of users can be secured.

6. <https://docs.mongodb.com/v3.0/core/index-ttl/>

/5

System design and adaptations

5.1 SNMP-interaction

To be able to communicate with components in the infrastructure in a reasonable way, and for the individual components in the system to be agnostic to the low-level data format of the information provided, a `acsntp-middleware-library` has been implemented. The main purpose of the library was to provide the traditional **get**, **set** and **walk** semantics of traditional `acsntp`, but at the same time hide idiosyncrasies in the components and language `acsntp` library from the rest of the modules.

Additionally, some internal and some externally visible tool functions was added to be able to use the internal `acsntp` information to convert the returned data to a sane format compatible with use in other contexts. This makes the rest of the system agnostic to how the data is transported to and from the components in the infrastructure, and in theory, it is possible to replace `acsntp` with another management protocol at a later time, with the implementation of a new middleware library that glues the new data source to the same **get**, **set** and **walk** semantics exposed in the current system.

5.2 Access Point model

5.2.1 Accesspoints (plural)

To account for how one usually systematically handles a set of access points, sets of access points have been defined as a separate interface(class). The primary intention of this is to provide some of the essential methods/actions used on access points as a wrapper method. One of the features of this is that parallelizable or otherwise optimizable set operations can be optimized internally, without exposing this to external callers of these methods.

From the `Accesspoints` interface, searching within the stored access points is possible, and will return a list of access points that match. The search dictionary is currently directly passed to the database engine, and should be a subset of the structure of an access point instance represented as a JSON document.

To be able to create a new access point entry from one or more controllers, a preliminary search/walk must be done to collect a minimum of information required to be able to create a new access point entry. The information needed is in most cases the system MAC address of the access point, which uniquely identifies the access point, in addition to being used as an index in `acsnmp` subtrees from the controller. From the point the access point is present in the database, it is possible to request and store more information about the access point from its controller. As access points may move between controllers outside the control of this system, a regular poll for new access points must be done, and is for all intents and purposes not any different than polling for completely new access points on a controller. As the system MAC-address of the access point does not change, the database is updated with the new controller address for the access point, and information about the access point is collected from the new controller at the next call to the update method.

5.2.2 Accesspoint (singular)

Each access point is handled as a separate instance of the `Accesspoint` class, and is intended as the interface used by all users of the library.

The class handles both creation of an access point entry, deletion and collection of information. The layout is designed to match the logical layout of how a wireless access point can be viewed, with static information like MAC-addresses, serial number, model and number of radios, dynamical information like uptime, name, location, and current primary controller. Each radio is also stored as a sub-entry of an access point, and information is collected from the controller based on the index of the access point, and the radio slot number.

Depending on the access point model, an access point may also have a number of wired(LAN) interfaces. These are treated the same way as the radios, and is collected based on the index of(acsnmp suffix) of the access point, and interface number.

5.3 Controller model

As controllers at its core are simple devices, a flat model with very little functionality is needed. At the current stage, few fields compared to access points are necessary to collect from controllers, and due to this fact, both static and dynamic information is collected every time.

In the future, if deemed necessary, more functionality may be implemented for controllers. Examples of potential actions may be to migrate all associated access points to other controllers, where capacity allows, so that the controller in question can be serviced or put into maintenance.

5.4 Analytics and correlations

An important design decision that sets the new system design apart from the existing solution is the fact that data is being collected and stored in a way that enables out-of-bound analysis and mining, without interfering with the process of collecting the data.

As the current data storage engine in use is document based, it enables the developer to easily make copies of the relevant portions of a document collection to do analysis, without risking the data being changed or removed during analysis. It is also possible to do changes to the copies, or remove document parts to decrease the amount of data shuffled around, which makes it easy to scale the analysis to much larger amounts of data sources(access points)

5.5 Optimizations

To ensure that the time used by collection of information is as short as possible, the design of the system allows for certain parts to be run in parallel. In both the library for access points and the controllers, mass collection of information is done with individual threads or processes(depending on OS mapping) for each controller or access point. This makes it possible to collect information for

a large amount of objects at the same time, without having to incur the same network overhead costs compounded for each object if it the same collection was done in serial. As such it is also in theory possible to do collection of multiple machines at the same time, if the controllers allow collection from multiple IP-addresses or subnets at the same time. Depending on the controller capacity, it should also be possible to run multiple instances of the devised system, for individual and collection for different purposes.

5.5.1 Poor SNR conditions

To detect poor SNR conditions, one can look at the raw noise metrics collected by access points every 180 - 300 seconds[31, p. 11]. However, this information is not available for all access points depending on access point model or software version. As such, different methods may be used to indirectly collect, or infer the noise in an area.

The most low hanging fruit is to collect and compare the number of clients which are classified within the poor SNR range by access points themselves. As this information is received from the modems in access points, the information about noise level/floor may not be directly available, but reported as a statistic from the modem, as a relation to the RSSI.

5.5.2 Detection of failed radios

To detect whether or not a radio has failed to during operation, failed to come up after reboot/power loss, or may be failing, there are several methods that can be deployed.

A typical indicator of a radio that may have a failed modem or transmitter is when an access point that have been operational for some time, but suddenly have one or more radios reported as operationally down, but administratively up. As most access points have two radios, one can deduce a potential transmitter failure or modem crash through observing one the one being still operational. When a down radio is detected, a primary test to determine if it is permanently disabled, or the modem may have a crashed can be determined by toggling the administrative status of the radio, from enabled, to disabled and back again. This resets the modem, and should in theory bring the radio back up if it is possible.

If not, there are one key indicator that can be used to determine if the radio transmitter or the modem has failed. If the status of the radio remains administratively disabled after attempting to toggle it, the radio is being kept down

due to other reasons. This can be due to lack of enough power from the power supply, due to changed regulatory conditions or other reasons that typically requires manual intervention.

5.5.3 Detection of failed subsystems in access points

Detecting faulty subsystems in access points may be much harder, and may not be possible directly through established or standardized interfaces like acsnmp. Examples of failed subsystems may be broken or missing certificate stores, failing or faulty flash storage or missing radio or modem interfaces.

To detect missing radio interfaces through acsnmp, there are several methods. Usually, an access point is equipped with a fixed number of radios, and may have an expansion slot, in which the access point will report if there is a valid radio preset. If one tries to collect information from one of these slots, three outcomes may arise. If the radio is functioning as designed, information from the radio can be obtained through SNMP, and depending on if the radio is enabled, the information can be used, or considered stale. On the other hand, if only parts of the information is possible to retrieve, this may indicate that parts of the radio has failed, or may not be operational, and the information should be considered stale. Lastly, if an error occurs, or the controller reports that there is no information at the given OID, it can be assumed that the radio is no longer actively connected to the access point control plane.

Another approach is to investigate the system logs from the controllers, to see if there are any access points that present invalid radio types or send invalid or malformed CAPWAP join requests or control messages. This may be an indicator of malfunctioning or corrupt hardware or failing flash storage.

If collection can be extended to the individual switches where access points are connected, it is also possible to detect corrupt flash storage by examining the Institute of Electrical and Electronics Engineers (IEEE) 802.3af/at power information, or Link Layer Discovery Protocol (LLDP) information supplied by access points. Access points with failed flash storage which are unable to boot their onboard operating system, will also be unable to identify themselves to the switch supplying power, and will appear as an IEEE power device, instead of an access point with a specific model name and power requirement.

5.6 (Semi-)Automatic actions

Based on information gathered from this system, some automatic, or semi-automatic adjustments are possible.

When an access point is found to be missing, an automatic task can be devised that may query the switch where the access point was connected, to investigate whether or not the access point is still connected, if it draws power, and if it identifies itself as an access point. If the switch cannot be reached, it can be assumed that it may be a local power failure in the area, the switch has failed or the area has lost uplink connectivity, all of which may be classified as "not an wireless problem".

To account for access points moving between controllers, an automatic task to search for a specific access point on all known controllers can be implemented. This reduces the perceived down-time of the access point from several collection cycles until a full poll of the controllers are done, to mere seconds, depending on the amount of controllers, and other workload.

Another useful task may be automatic redundancy failover to ensure that the primary controller in a HA-SSO setup is the active controller when possible, to avoid prolonged periods of time where licences are temporarily inherited, and to avoid data only available from the primary controller being lost.

Further, as outlined in Section 3.2, it may be possible to generate lists of access points that due to interference reduction reasons, may have their 2.4GHz radio turned off or disabled. However actions that may impact user experience, or cause temporary or permanent connectivity loss is advisable to be done semi-automatically, or partly manually. This can be solved by generating finished setups for these actions, and have an operator approve, deny or limit the extent of the operation, to ensure the best client experience.

5.7 Alerts

When information about access points and controllers has been collected, it is simple to generate automated alert based on state or thresholds. In example, to generate alerts for all access points that no longer is associated with a controller, one simply queries the database for access points that no longer are present. In future developments, it is planned to have individual alert generation as part of the periodic collection process of each access point. This way, access points that are polled very often can have their alerts actively pushed to the attention of an operator or another system, while less important access points that are

polled more seldom, may not generate alerts quite as fast.

With collection of information from controllers, example alerts like sanity checks that ensure that a redundancy failover notice can be sent to the operator, licence threshold warnings and early temperature warnings can be issued. Also here is it planned to make alert generation a part of the collection process instead of alerts being generated on the fly when polling the database for irregular conditions.

/6

General monitoring

6.1 CAPWAP health

During early deployment of access points for the Student Welfare Organization in 2014, substantial problems with maintaining stable CAPWAP tunnels were experienced. Due to lacks in the management system in use, the issue went largely unnoticed for a long time, as the collection interval of the management system was too long to effectively notice missing access points from the controllers. From when the first problems were reported, it still took a lot of time and effort to pinpoint the cause of the client problems in the areas affected. When discovering that the problem described by users was caused by, or indicated by the teardown of CAPWAP tunnels due to loss in connectivity, it became apparent that this could be a potential indicator not only to similar problems, but also a key point on any wireless debugging checklist.

The cause was later traced to a bug in the energy-saving protocols in use in modern switches and access points¹. This caused the access points to briefly lose contact with the switches they were attached to and powered from, and due to long reconnection time from spanning-tree protocols on the switch ports configured for access points, the CAPWAP tunnels were torn down after timing out. This not only caused packet loss for associated clients before the tunnel was torn down, but also disassociated all clients when the tunnel were torn down. A temporary fix was created to alleviate the capwap tunnel teardown

1. <https://quickview.cloudapps.cisco.com/quickview/bug/CSCus35889>

by forcing short convergence times for the spanning-tree, which enabled the CAPWAP tunnel to survive in most cases.

As mentioned in Subsection 2.1.4, maintaining the CAPWAP tunnel is important for client experience. To be able to detect potentially silent or hidden faults or issues in the infrastructure between the access point and the controller, or even issues with the access point itself, it is therefore to often, and regularly monitor the uptime of the CAPWAP tunnel. Fortunately this number is available for collection through acsnmp from the controllers used in this project.

6.1.1 Detecting and classifying access point teardowns

Detecting current torn down CAPWAP tunnels is different from platform to platform. For the controllers used in this project, no explicit way can be used to count or collect the currently down CAPWAP tunnels, as the controllers in practice have no relationship with access points not currently associated to the controller. This means that missing access points must be detected through periodical polling of the controller. The amount of, or type of data being queried is irrelevant, as the indicating factor is whether or not the controller is able to respond with information or not. If the controller is unable to supply the information, this means that the access point is not associated with the controller, and the CAPWAP tunnel has been torn down. In other wireless controller models², a traditional 802.3-compliant interface is present on the controller from the first time the access point was associated, and remains present until the controller is rebooted.

To determine the cause of the teardown, one has to look at information from the access point itself if it has reestablished its CAPWAP tunnel, or draw conclusions from other access points (found through RRM neighborships). If the access point in question is present, information and difference in the CAPWAP and access point uptime can tell the operator whether or not the loss of connectivity was caused by powerloss or the access point being disconnected from the network outlet. If the uptime of the access point and CAPWAP tunnel is similar it can be deduced that the access point was rebooted. Depending on whether or not nearby access points show the similar statistics, it can be determined whether or not the access point, switch port or cabling may be compromised or may be failing. If all access point in the area or connected to the same switch show the same statistics, it may have been caused due to a local power loss to the area.

2. like Cisco 3650, 3850 and 5760

Does the CAPWAP uptime greatly differ from the access point, it may indicate that the network connectivity was lost somewhere between the access point and the controller. Using the same logic as above, if the same difference or CAPWAP uptime can be observed on other access points, it is fair to assume that the connectivity loss was to the entire area. If the CAPWAP uptime only differ on a subset of the access points, it may indicate failing wiring, or potential tampering.

By keeping statistics of the average CAPWAP tunnel uptime before uptime drops to zero again can help administrators identify potentially problematic access points that struggle to keep a tunnel up for extended periods of time.

It has also been observed cases where the internal certificate on access points have expired due to an age limit of 10 years on the manufacturer installed certificate, which in turn caused the access points to not be able to establish a CAPWAP tunnel to the controller. An age check of access points can therefore be implemented to warn administrators to either replace the access points, or add manual overrides for this age check. Other causes for CAPWAP failure may be broken cryptographic chips or flash memory which causes invalid CAPWAP messages to be formulated.

6.1.2 Controller caused teardowns

A CAPWAP tunnel may also fail due to faults in the controller hardware or software. However these causes may be significantly harder to detect, especially if it may only affect certain access points or a small subset at a time. The best known indicator to determine whether or not a controller may have silently failed or may experience hidden capacity issues is to monitor the average CAPWAP uptime of all its access points, in addition to monitoring the number of associated access points (and thereby open CAPWAP tunnels) closely.

During the spring of 2015, a hidden capacity issue was discovered in the Cisco 5760 Next Generation Wireless Controller. While doing regular maintenance on other wireless controllers, 800 access points were gracefully moved to the 5760 controller, and started to establish CAPWAP tunnels. As the number of access points grew at a slow pace, gradual CAPWAP tunnel failures were observed, and access points started disassociating. By logging into the console of the controller, and running operating system diagnostics, it was determined that the cryptographic capacity of the hardware was insufficient for maintaining more than 450-600 simultaneous CAPWAP tunnels, and that the CPU of the controller was busy trying to compensate. Similar observations was later done

at NTNU³ and UiT⁴. Current recommendations for this controller now dictate a maximum of 600 concurrent access points on this controller, with a previous datasheet specification of 1000.

6.2 Controller health

With a geographically distributed controller architecture, it is important to be able to monitor the health of the controllers without having to be physically present, nor having to drill down into menus to find relevant metrics. Information like current temperature, licence level, amount of clients connected through the controller, along with indication whether or not the controller high-availability setup has failed over to the secondary hardware controller is crucial for the operators to have.

To avoid silent licence issues, it is important to monitor the number of available access point licences on the controllers in question. As access point may be automatically or manually failed over to one or more controllers, the available free licences may vary from time to time, or it may be necessary to move licences to other controllers depending on necessity. It is also interesting to monitor licence usage to see if a reasonable percentage of the licences are used, or lay unused due to inefficient distribution.

Additionally, the way high-availability setup work, two or more identical physical controllers may work in lock-step as one logical controller. If the primary controller fails, the secondary controller takes over in the course of microseconds, and maintains all state through what is called a high-availability stateful switchover(HA-SSO). This can also be induced manually during hardware maintenance on the active, primary controller. The challenge of this setup is that due to licensing restrictions, the secondary or tertiary controller can only hold(inherit) the licences for 90 days. If a controller failover goes unnoticed for more than 90 days, parts of the infrastructure may stop serving clients, or access points may be unable to associate to the controller after CAPWAP teardown. Additionally, in the interest of this project, there are also a subset of SNMP OIDs that will not resolve when queried to the secondary controller, even though it is currently in the possession of the shared IP-address.

3. Norwegian University of Science and Technology

4. University of Oslo



Interference reduction

7.1 Requirements

To accomplish a noticeable, and over time, significant reduction in interference, a series of different information aspects must be covered to make the efforts as effective and precise as possible. As not all areas are as equally affected, a primary set of access points or areas can be devised by gathering all access points which report an interference threshold violation. From these, it is possible to find their closest neighboring access points, that also may candidates for the interference reduction process. It is also important to note that the intention of potentially deactivating radios in the interest of reducing interference should not be confused with future intentions of decommissioning the 2.4GHz spectrum as a user-spectrum, and that the signal strength and coverage should not fall below recommended levels. As a significant part of devices being sold to consumers still only support 2.4GHz wireless networks, areas with a high number of devices still using this spectrum may not benefit from deactivating radios, even if the amount of interference decreases.

7.2 Method

As mentioned above, to identify areas where the interference is above the recommended threshold, one fetches all the access points which reports an interference threshold violation `bsnAPIInterferenceProfileState == 0`. These

access points form the primary candidates for deactivation of their radios. Further, if information about the nearest (RX) neighbors of each access points is available, these form secondary candidates for deactivation. If access points are aptly named, with building code and room encoded into the name (like `tf-2410-01-rw`) it is also possible to include access points within the same room into the pool. This can be access points within the same auditorium, or other large venues (Kultursalen in Alta is a prime candidate for this, with 11 access points using directional antennas).

From this information, it is possible to map the relationship between the access points in an area as a graph, with edges between the access points on the same channel, creating "islands" for each building, room or area. By identifying the central hubs or hotspots in the graph, potential high-gain, low-effort candidates can be marked for deactivation, with the result being the hub breaks up in part or in whole, and may reform smaller groups which may later be reevaluated for further deactivation. As a final safeguard, access points which at the point of running this mechanism, or if available, historically have had a high number of 2.4GHz clients in the recent past, may not be advisable to deactivate. This can be weighed against the information about neighbors, which may indicate the coverage in the area can be considered sufficient, and that deactivation still is advisable.

7.3 External Interference adaptations

As UiT and most large organizations have wireless infrastructure in places where other parties also have wireless infrastructure, it is not always as easy as deactivating some of ones own radios. In these cases it is important to monitor the interference on all viable channels, to be able to adjust ones own infrastructure to the potentially changing radio environments, as more and more foreign radios may appear. It may be more appealing to run several of ones own access points on the same channels and incur some interference, than to share multiple channels with other parties, and having to consider their own infrastructure usage, and their potentially ill advised RRM planning. Other times, all channels are approximately equally crowded, and in these cases, it may be advisable to make technological shifts towards a pure 5GHz infrastructure, where possible, and cost allowing. This may not only decrease support and management for users, but also improve work performance.

In some areas, policy enforcement may be necessary. One example of this is enforcement of no private wireless networks on student dormitories where existing infrastructure run by UiT or the student welfare organization is present. As wireless access points/routers manufactured for the home marked, with

high speed and long range in mind does not contain the necessary RRM capabilities, it may cause serious amounts of interference in coalition with other similar access points in the area, making the experience for both users of the private wireless network, as well as the sanctioned wireless networks, worse. To counter this, it is possible to collect information about where the potentially unsanctioned access point is positioned based on information about signal strength and which access points that can see it as an interference source and as a rouge access point. The collection and aggregation of rouge access point data is one of the planned features for the future extension of the developed system.

7.4 Other efforts

It has historically been attempted several different efforts to decrease overall interference and amount of busy air time. One method which has been employed by most educational and public institutions is to deactivate and disable 802.11b rates and support in their infrastructure. As 802.11b rates range from 1 Megabit per second to 11 Megabits per second, the time spent transmitting per unit of data is much higher than for more modern data rates.

At UiT, 802.11b rates were removed as part of the supported rates in the infrastructure during 2014, and the minimum required data rate for communication raised to 24 megabits per second. This enables much more efficient use of the available air time, and forces "lazy" clients that prefer lower data rates to use the shared resource much more efficiently. Further, by disabling 802.11b as a standard, safeguards and adaptations that are in place in the newer standards (802.11g and 802.11n) can be removed, contributing to a more efficient communication for the much more modern infrastructure components. As 802.11b also support Direct Sequence Spread Spectrum (DSSS), which occupies 22MHz with its channels, the gains received by removing the 2MHz overlap to the nearest non-overlapping channel can be significant, in addition to contributing to potentially more channels being considered "available".

7.5 Further development

The logical next step in the efforts to reduce interference is to not only address the areas where interference levels surpass the defined threshold level, but also to try to reduce the average interference level overall. Further investigations into interference reduction in the 5GHz band may be the next major technological step, to ensure that the main wireless information highway is as congestion free

as possible. With the advancement and development in deployment density of wireless access points, the same interference issues may arise for the 5GHz spectrum in few years. To counter this issue, the first access points[32] with 5GHz micro/pico-cell design, capable of creating very small(5-15 meters) cells are becoming available. This enables channel reuse without the same degree of interference.

/ 8

Testing

8.1 Performance

8.1.1 Data collection

To test the performance of the collection of data, a simple collection program(`collect.py`) was developed early on, which have been used during the entire development period to test new parts of code, in addition to testing in-production large scale collection of data. The program pulls all stored access points from the storage engine as an access point collection, and proceeds to call the method for collection of dynamically changing data fields for all access points.

This spawns one process per access point, which proceeds to ask the controller currently in charge of the access point for the dynamic information. In practice, this spawns thousands of processes which all spend most of their time waiting to establish a connection and ask the controller for its data.

To see if the amount of information has anything to do with the overall running time of a single collection cycle, the amount fields to be collected, was adjusted. This was to determine if the amount of data, or the sheer number of individual, non-neighboring fields would impact overall collection time, and thereby impact the freshness of the data being collected.

The collection program has also been used to continuously collect data for

analytics for the entirety of the development cycle, both to be able to operate on fresh data, to collect new interesting fields as they are found or required, and to test the longevity of the code to see if any new or unexpected faults in hardware or source code can cause the system to crash or output incorrect information to its callers.

8.1.2 Data presentation

To test the performance of the data presentation (ie. Fetching of information from storage), ab¹ along with the mongo interactive utility was used. For actual graphical performance testing Mozilla Firefox 44.0.2 with developer console was used.

The goal of testing the performance of the data presentation system, was to ensure that it would perform as intended with acceptable response time to its users, even under heavy load on the decoupled information collection system.

8.1.3 Analytics

As the total amount of data present in the active data set cannot be considered a dataset large enough to test throughput or large scale performance, it is hard to quantify the performance of the analytics themselves.

However, as the underlying database inherently support concurrent read and writes through read and write concerns, which describe how consistent the data need to be, the scalability and performance potential is large compared to previous implementations[12]. While earlier implementation could cause seemingly random latency spikes on data retrieval due to table locking during insertion, the current implementation display no such symptom, and analytics like interference reduction, access point information aggregation and correlation execute in the course of milliseconds, with a complete GUI load completing within 20 milliseconds.

8.2 Use cases

The current implementation of the collection and analysis library has been used in development implementations of dashboards that display the status

1. ApacheBench: <https://httpd.apache.org/docs/2.4/programs/ab.html>

of the wireless infrastructure on each respective campus. Further, command line tools for for automatic provisioning of Cisco 702W access points is under development, for use in the student housing complexes in Tromsø and Alta. This tool will use information about access points, switches and their associated controller to automatically name, configure and test the access point shortly after it being mounted, to ensure an acceptable experience from the first moment. As the system provides an abstraction for external code to use, the same tool can be used to provision other access points, such as the upcoming Cisco 1810W and Cisco OEAP1810 access points planned for deployment during July-October.

8.3 Correctness

As the data being collected and presented in the developed system is information collected from a real-world environment, parts of the data may be hard to verify its correctness, as the nature of the information may be fluctuating, or changing from one collection to the next. The information provided by controllers however, has painstakingly been verified every step of the way, to ensure that the correct fields are collected, and that the information contained within, is the information desired.

To ensure that conversions between ASN.1 format and Unicode/UTF-8 format in python and MongoDB worked correctly, and was able to decode/convert a wide range of characters, a series of manual tests were done. Conversion of common, but odd UTF-8 characters were tested by inserting a selection of characters into text string fields on controllers, and then collected and converted and displayed using the newly implemented system. The tests showed no problems with the current implementation with regard to conversion of UTF-8 characters, even though the underlying protocol(SNMP) does not explicitly support this character set.



Figure 8.1: UTF-8 characters correctly collected from controller and displayed in the HTTP GUI view at UiT

To ensure the correctness of the implementation, the newest version of the implementation was run in production from the time it was first possible, to the completion of this thesis. During this time, the implementation correctly handled failing access points, failing controllers, network problems and problems with the underlying storage due to disk failure. Even though the console was left open, the implementation was able to run without supervision for the entirety of the period.

/9

Results and observations

9.1 Interference reduction

The most prominent result of this project is the result of the collection of data from access points, and the correlation done to pinpoint access points which contributes to the increase of 802.11-classified interference, and thereby contributing to the overall decrease in performance. As shown in Figure 9.1, the amount of access point with interference levels violating the threshold level of 20% has decreased significantly. The majority of the decrease can be accounted to the deactivation of access points in student dormitories.

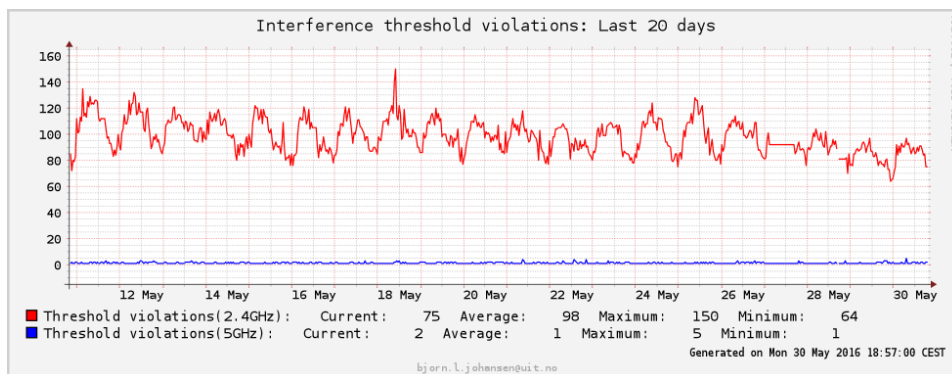


Figure 9.1: Access points with violating interference levels at UiT(as of 2016-05-31)

Source: Bjørn Johansen, Base Services/Network, IIA, UiT

It should also be noted that at least one of the access points that experience interference threshold violations in the 5GHz band seem to have a faulty radio or radio modem, which causes it to misinterpret signals as 802.11 traffic, and reports a constant 100% interference level. This is valuable information to use in further development, as this may be a suitable indicator of failing or faulty access points. The access point in question is at the time of writing, approximately 11 years old, and therefore scheduled for a quick replacement. In the mean time, the advisable thing to do with these kinds of faulty radios is to deactivate them, to ensure that the access point does not cause problems with clients trying to associate, or problems for surrounding access points through the 802.11 CSMA/CA-mechanisms with unnecessary RTS/CTS control messages.

9.2 SNR observations

As was somewhat expected, the majority of the access points with clients suffering from poor signal-to-noise ratio, were access points in areas with few access points, and large distances between access points. This means that the average client in these areas had low RSSI, which was supported by the information about clients collected from the access points. Further, when clients suffered from poor SNR, but still having acceptable RSSI within the specified requirements for high-density deployments, the poor SNR was found to be caused by actual noise sources.

The most prominent example being a new surveillance camera, which rendered channel 1 on the 2.4GHz spectrum completely useless for 802.11 traffic within 100 meters. Due to a high density of access points in the area in question, two access points were still forced to use this channel, and had serious problems with maintaining a stable, low-latency connection to clients. In this case, the 2.4GHz radios of these access points were subsequently deactivated, so that the access points would only announce SSIDs and associate with clients on 5GHz channels.

9.3 Data size impact

To determine if the time used in a complete collection cycle of all access points were dependent on the amount of data or fields scheduled for collection, the amount of fields to collect were varied in different testing scenarios. This was additionally tested along the development process, where more and more fields were added, and in some cases more advanced processing directives applied

to the data.

What was observed was little to no impact on the total time used during sequential and semi-sequential collection of data, with varying collection time seemingly dependent on controller load, and test machine load, much more than the amount of fields being collected. This was further verified by timing and profiling the collection of a single access point, where the majority of the time for collection spent in OS and library overhead with regards to allocating an available socket to establish the connection, and busy-waiting for the result from the controller.

This was also reflected in the system load when collecting all access points in parallel, with the system spawning an excess of 2100 simultaneous processes, with minimal actual CPU load, and most of the processes resorting to a waiting state. Minimal impact on the collection time, most time is used establishing connection and local/remote busy wait. The system load average still reflected the vast amount of processes present on a single machine, but still remained fully responsive, and trivial command line CPU benchmarks seemed unaffected by the seemingly high system load.

/10

End products

The following products and works can be drawn from the practical efforts of this project and thesis.

10.1 SNMP abstraction layer

A *SNMP* wrapper library for conversion of various kinds of data from a wireless and wired infrastructure has been developed and tested. This forms the foundation for all collection of data in the project, and at its current stage in development, fulfills the necessities of the rest of the system. Over time, the internals of this library is to be replaced by the low-level parts of the *pysnmp* library in efforts to further speed up collection, and make it possible to schedule data collection asynchronously from the rest of the control flow of an continuously collecting system.

10.2 Data collection and storage code

10.2.1 Wireless controller

A skeleton framework class responsible for collecting information from wireless controllers has been implemented. The framework class exhibits the presently

needed and intended functionality necessary for the collection, but is meant to be extended and divided into classes that inherits the main class, to be able to support and adjust functionality to different controller models, and in the future different manufacturers.

10.2.2 Access points

As for access points, the same skeleton framework has been implemented to collect information about access points and their environment. The main difference between access points and controller with regards to collection of information is that access points requires a binding to a specific controller for the framework to be able to pinpoint where to collect information from.

As a part of the framework developed, a simple autonomous discovery method has been implemented for access point to automatically be added to the database and later be collected using the normal method.

10.3 HTTP GUI

To showcase the information collected in an orderly manner, a web-based GUI based on **Flask**¹ has been designed. Its primary intention was to create a base for how information can be displayed if integrated into systems like **NAV**[5], in addition to creating a potential base on which other interfaces could be created. An example of this is an extension which delivers structured data from the wireless infrastructure via RESTful[26] APIs.

The current version of the GUI is implemented with direct access to the database, but with the further development of the access point and controller library, this will be reimplemented to use their exposed methods and data fields. As such this HTTP GUI is standalone from the collection of data, and is designed to run on a separate server apart from the server collecting the data.

The GUI supports showing access points in lists(Figure 10.1) with relevant information and possibility for drilling down to individual access points(Figure 10.2).

1. <http://flask.pocoo.org/>

Name	Location	Up	Model	Controller	Uptime	CAPWAP uptime
AP-Alta-3702-Sif-adm	Sif-adm	Yes	AIR-CAP3702I-E-K9	alta-wlc5508.infra.uit.no	637236050	637226450
AP-Alta-Forming-2etg	Over himling utentor rom 2123	Yes	AIR-LAP1131AG-E-K9	alta-wlc5508.infra.uit.no	1181421905	576714905
AP-Alta-Haldde-gang	Over himling utentor rom 2107 - 2. etg.	Yes	AIR-LAP1142N-E-K9	alta-wlc5508.infra.uit.no	1182894743	576716243
AP-Alta-Newton-gang	Over himling utentor 1087 - vhagen1	Yes	AIR-LAP1131AG-E-K9	alta-wlc5508.infra.uit.no	1182696582	576715582
AP-Alta-Vinterhagen	Over himling 2.etg vinterhagen - 2610	Yes	AIR-LAP1142N-E-K9	alta-wlc5508.infra.uit.no	1182585556	576716856
AP-prestM200-sw G1/0/25	default location	Yes	AIR-CAP3502I-E-K9	ma-wlc8540.infra.uit.no	803978455	524667955
AP-prestM200-sw G1/0/26	default location	Yes	AIR-CAP3502I-E-K9	ma-wlc8540.infra.uit.no	744985903	524667903
AP-prestM200-sw G1/0/27	default location	Yes	AIR-CAP3502I-E-K9	ma-wlc8540.infra.uit.no	803910779	524666779
AP-prestM200-sw G1/0/28	default location	Yes	AIR-CAP3502I-E-K9	ma-wlc8540.infra.uit.no	804016850	524668250
AP-prestM200-sw G1/0/29	default location	Yes	AIR-CAP3502I-E-K9	ma-wlc8540.infra.uit.no	804113921	524667321
AP0026.9986.c3fc	default location	Yes	AIR-LAP1142N-E-K9	narvik-wlc5520.infra.uit.no	636983484	636971984
AP0462.73b4.627c	default location	No	AIR-CAP2702I-E-K9	harstad-wlc5520.infra.uit.no	310658	301158

Figure 10.1: List of some access points at UiT

tf-2410-01-rw

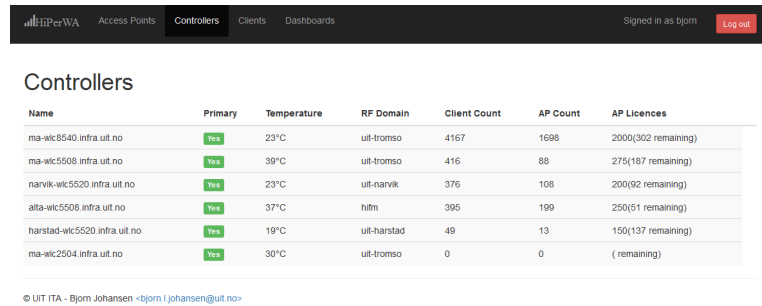
Teorifagbygget 2.410 / EØAΨx

UP X

Info	Radios	
Model: AIR-CAP3702I-E-K9		
Serial: FC21805U0ZQ		
System Mac Address: 18e72877ec20		
Interface Mac Address: 18e728684128		
Controller: ma-wlc8540.infra.uit.no		
AP-group: uit-ita		
	2.4GHz(Slot:0)	5GHz(Slot:1)
	State: Enabled Up	State: Enabled Up
	Channel: 6	Channel: 40
	Power level: 4	Power level: 1
	SSIDs: 6	SSIDs: 7
	Clients: 0	Clients: 5
	Clients with poor SNR: 0	Clients with poor SNR: 0
	Load: OK	Load: OK
	Interference: OK	Interference: OK
	Noise: OK	Noise: OK
	Coverage: OK	Coverage: OK
	Channel Utilization: 23%	Channel Utilization: 1%
	Radio Utilization: RX: 0% TX: 0%	Radio Utilization: RX: 0% TX: 0%

Figure 10.2: Showing access point tf-2410-01-rw

In addition, a quick view of controller status is possible.



Name	Primary	Temperature	RF Domain	Client Count	AP Count	AP Licences
ma-wlc8540.infra.uit.no	Yes	23°C	uit-tromso	4167	1698	2000(302 remaining)
ma-wlc5508.infra.uit.no	Yes	39°C	uit-tromso	416	88	275(187 remaining)
narvik-wlc5520.infra.uit.no	Yes	23°C	uit-narvik	376	108	200(92 remaining)
alta-wlc5508.infra.uit.no	Yes	37°C	hitm	395	199	250(51 remaining)
harstad-wlc5520.infra.uit.no	Yes	19°C	uit-harstad	49	13	150(137 remaining)
ma-wlc2504.infra.uit.no	Yes	30°C	uit-tromso	0	0	(remaining)

© UiT ITA - Bjorn Johansen <bjorn.l.johansen@uit.no>

Figure 10.3: List of some controllers at UiT

10.3.1 Dashboards

A much requested feature, dashboards with critical information about the health of the wireless infrastructure has been created, with rudimentary HTML5, JavaScript and CSS. These are meant to be displayed on monitoring or information displays on each individual campus, and more will be created to give an overview over the infrastructure locally on location where UiT holds its digital exams.

These dashboards could be customized to show controller information and alerts, information and alerts from a specific subset of access points, or information from all sources.

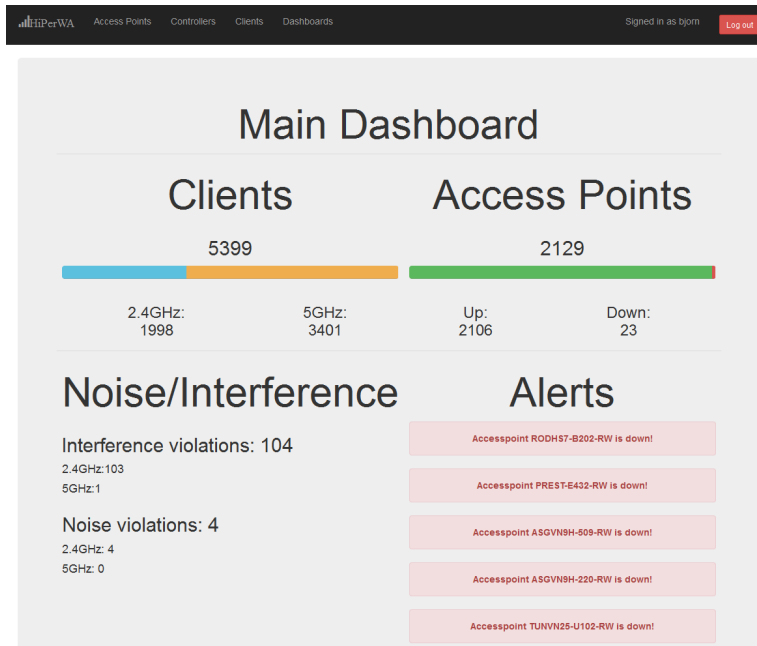
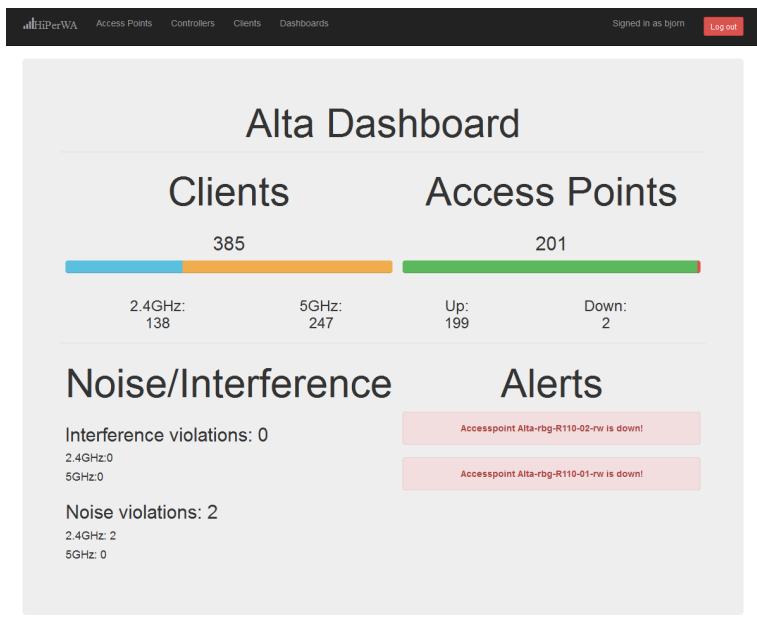


Figure 10.4: Main dashboard view at UiT



© UiT ITA - Bjorn Johansen <bjorn.l.johansen@uit.no>

Figure 10.5: Alta dashboard view at UiT

10.4 Command line provisioning tool

A rudimentary tool using information collected, and the SNMP methods have been implemented to automatically provision access points installed in student housing areas. The tool takes the ethernet Media Access Control (MAC)-address of the access point, and sets the hostname, location, access point group and sets all wired interfaces to a specific student VLAN.

When the new Cisco 1810W and Cisco OEAP1810 access points arrive to market, the tool will be expanded to support provisioning the local ports to CAPWAP tunnels for wired traffic.

10.5 Unfinished products

10.5.1 Alert module

At the current time, alerts and notifications are up to the individual users of the collected data to generate. It is therefore desired create a mechanism to generate and store these at collection-time, in a format that can be used by systems, but with a standardized algorithm for generation.

This would also make it possible to generate e-mail alerts and summaries to be automatically sent out at given intervals or for critical alerts. Extensions for SMS or Syslog should not be hard to implement either. Additionally, small HTML applets for public display can be created to showcase current statistics to the public.

10.5.2 Self-service portal

Another desired feature is a self-service portal for users to log into and check what information is available from the infrastructure about them and their devices. By combining information from controllers and RADIUS about authentication and IP-address, access points with radio and traffic information, potential location services for more accurate positioning, a user could not only see which devices are connected, but where they are, if there is something wrong with their setup and other usefull information.

Further, this could be linked to other planned projects like the personal SSID system described in Subsection 11.8.1.

10.5.3 Other projects

A reimplementaion of the WHOIS-service from earlier[12] was also planned, but due to time constraints this has not been done yet, as the current version is working as intended, and run alongside the new system.

/ 11

Discussion

11.1 Improvements over current system

In the implemented system, several key shortcomings of the existing management system has been addressed. Using already stored data to store static information, and reducing the amount of SNMP requests to the controllers by compounding different fields into one request seem to have reduced the time and resources needed to collect the most important information needed. Further, extensibility and flexibility has been introduced by making it possible to collect custom fields that may not be supported otherwise, in addition to making it possible to dynamically or priority schedule collection of information from access points based on load or the need for updated information from a subset of access points.

It is also easy to extend the functionality of the system, with both collection of new data points and custom functionality to suit the need of users of the system. As the system has been organized into standalone libraries, it is possible to integrate parts of into systems like NAV.

11.1.1 Responsiveness

As the collection of data uses far less resources than the existing system, the responsiveness and effectiveness of use has been greatly improved. While page loads in the existing system took up to 10 seconds to complete, the longest

page loads now take approximately 0.2 seconds, with most of the time spent client side.

11.2 Interference reduction and effects

As seen in Figure 9.1, the reduction in persistent, and partially in intermittent interference threshold violations have decreased, even with just a minor number of deactivated 2.4GHz radios. As this the deactivation of radios currently is manually overseen to avoid any unforeseen or unintentional consequences, the deployment campus wide and to student housing has not been done. Instead, select buildings on campus and select parts of the student housing areas have been analysed, and a total of 23 radios in the 2.4GHz spectrum has been disabled from this. This have resulted in a decrease of peak violations from ≈ 140 to ≈ 100 , minimum number of violations from ≈ 80 to ≈ 60 and average violations over 24 hours of from ≈ 100 to ≈ 80 .

11.2.1 Interference usage

When working with high-density deployments, one use case for interference values can be to see the averaged dispersion of wireless signals in an area, if all access points report approximately the same interference level, or the majority of access points in an building reports interference threshold violations for 2.4GHz, this means that the density of access points is sufficient to be able to cover the area, even with a combination of 2.4GHz and 5GHz, or potentially 5GHz alone. This can be determined by looking at the RX neighborships of the 5GHz radios, along with their TPC settings. If an access point can receive N neighbors, this may indicate that the surrounding areas are covered, and if the TPC setting of the radio is not set to maximum, this means that there is no need to cover more area than already covered.

11.3 Adverse effects

One of the major concerns when starting this project was that efforts undertaken or changes done to the infrastructure during testing would adversely affect the experience of the users of the wireless network at UiT. As the network used for testing not only is used in production for business critical applications like digital exams and live streaming of lectures, but also have a great geographical expanse, changes that could potentially either cause components to fail or loose connectivity have been done with great care during maintenance

windows or in periods of the day with low traffic.

Changes to the radio environments and policies have been done gradually and geographically delimited when possible, so that potential effects could easily be traced back to these changes.

As of this writing, there has not been recorded any lasting adverse effects that has affected the user experience or reputation of the network. It should however be noted that a minor loss of redundancy was observed at one point, due to component failure traced back to actions during testing. This could however be attributed to faults in component software, and this project should not be considered the cause of this. The failure was reported to the manufacturer, which implemented a fix for this scheduled for the next software release. Full redundancy was restored after 66 seconds.

11.4 Alternate measurements

Alternate measurements that can tell something about the health and experience of the wireless infrastructure include Frame Check Sequence values, checksum errors, retransmission count and transmission delay or failures.[33]

Further, by looking at the throughput from clients which actively transmit or receive data, one can determine if the client is sufficiently covered by looking at both the throughput, the number of spatial(simultaneous) streams, RSSI and other factors.

One method that is under trial, is to look at the traffic itself to see if there are any protocol indicators of latency, retransmission, delay, duplication or other problems. UiT currently route all their guest traffic through an enterprise Network Address Translation (NAT) scheme, which then is capable of tracking individual TCP sessions. By looking on the packets forwarded, it can be determined if the packet is a retransmission, the round trip time of an packet and the subsequent acknowledgement and if the size of the packets decrease due to errors, dropped packets, checksum errors or other potential problems with the underlying(wireless) infrastructure back to the client. Further, as the Round-trip time (RTT) used in TCP considers the entire RTT from end to end, it is possible to look at how much of the RTT is due to local conditions in the wireless infrastructure, and how much is outside the network at UiT.

11.5 Client station collection

As noted in Subsection 1.4.4, collection of information that can be used to identify or track individual persons or entities is not recommended nor wanted, as this may require additional security measures with regards to storage and handling.

Therefore, a choice was made early on to focus the current stage of this project on the infrastructural data available, and make an effort to make up for client specific information through clever reasoning and additional data correlation. An existing service where near real-time client information is available, exists from the previous HiPerWa project, and is still in use in day to day debugging and support. This service was developed as a trial and was exempt from the restrictions mentioned above as it is secured in a way that only a few key employees have access.

11.6 Limitations

The current implementation and design has been tested on wireless controllers, switches and access points from Cisco Systems, Inc. At the current stage in development, it cannot be guaranteed that the same implementation will work on other manufacturers without modification or adjustments.

When it comes to the theories and observations regarding the CAPWAP protocol, these assumptions and conclusions should hold as valid for all manufacturers implementing the protocol in compliance with the specification[15][13][14].

Further, the implementation has only been tested on the Linux operating system, but should in theory work as expected on other platforms which support python 3.5.1 and pysnmp, along with standard libraries and network support.

It should also be noted that minor implementation details have been redacted due to protection of what is considered trade secrets. Their redaction should not impact the operation of the system, nor the outline or conclusion of this thesis.

11.7 Related work

A trial project is under evaluation at University of Oslo, where it is planned to place access points on the nearby subway and bus stations, where only 5GHz

radios are to be enabled. This way, they hope to "prime" clients to prefer 5GHz when proceeding to the campus wireless network, and thereby increasing the use of 5GHz by "encouraging" clients. Deployment of this project, along with results are yet to become a reality, however it may have an actual effect, as most roaming and scanning algorithms try to reuse the same spectrum for further roaming, before switching band. Additionally, protocols like 802.11k[34] and 802.11v[35] would be able to assist, by providing clients with the best access points to roam or connect to.

A similar project[33] as this thesis was done in 2013, in which Cisco 4404 controllers were used, and the goal of this thesis was to look at the effects of more cleverly allocating channels in the 2.4GHz spectrum to decrease failed transmission, failing checksums and retransmissions. This thesis takes a more abstract look at the same problems, but with an approach to decrease interference, and to increase visibility of unknown problems, and at a much larger scale.

11.8 Future work

11.8.1 Personalized SSID system

One of the foremost plans to continue work with the system developed in this project, is to extend the system to support future plans of offering personal SSIDs to students living in housing from the Student Welfare Organization. This would enable each resident to log into a portal where they can manage the access point closest to, or in their place of living, and create, delete or change their personal network.

To achieve this, the current system would have to be extended to not only account for wireless-only infrastructure, but also be able to use information from housing management systems, user databases from UiT, network and subnet assignment procedures at UiT and to a greater extent account for RRM mechanisms to ensure that the interference level is kept at a minimum, even with a substantially greater number of broadcasted SSIDs.

11.8.2 Digital exam monitoring system

One of the features built into the system in this project, was the ability to prioritize collection from specific access points, and schedule these collections more often, or before collection from the rest of the infrastructure. This was done to be able to deliver very fast, responsive information to operations

centers for digital exams at UiT. At these centers, representatives from the IT-department must be able to monitor that the wireless infrastructure is working as intended, and almost preemptively raise alerts to network engineers if something seem to be sub-optimal.

Additionally, to raise the degree of assurance, it is desirable to develop a system which can ensure that specific usernames of candidates are present within the client pool of specific access points, and that there is a fixed amount of devices per user, so that attempts of cheating can be discovered when a user has more than the allowed number of devices due to their cell phone or smart watch is present. Additionally, when monitoring whether or not a specific user is present on the network, this can be used to detect if the user disconnects from the required network, and instead uses a rouge network to communicate with outside parties during the exam.

11.8.3 CleanAir management

It is also planned to look closer at the possibilities in the Cisco CleanAir¹ technology which enables access points to identify noise sources and make smarte decisions on which channels to use, and which to avoid all together. This information is available through SNMP and should be possible to normalize and insert into the developed system. Further, it may be possible to gather raw data from radios to look at how the radio environment is in an area, before deploying more access points or to help other use cases of the 2.4GHz and 5GHz spectrum.

11.9 Future plans

As mentioned in Subsection 1.3.1, it is desirable to avoid usage of the 2.4GHz frequency band to avoid the interference and noise issues of the band. Therefore, plans are currently under development at UiT to gradually move the primary users of the wireless network away from using this band, and over to a pure 5GHz client infrastructure.

To achieve this, several methods are to be deployed. The first step of this process will be to utilize information from this project to see if there are areas at UiT where 2.4GHz usage can be completely abolished through the deactivation of even more radios. This will create the possibility of creating so

1. <http://www.cisco.com/c/en/us/solutions/enterprise-networks/cleanair-technology/index.html>

called greenfield-zones, where new technologies can be tested, to achieve even higher throughput without having to consider supporting old equipment and protocols.

Further, the most used SSIDs at UiT, Eduroam and uit-guest among others will no longer be broadcast on 2.4GHz. This simplifies the management of the infrastructure substantially, as the majority of the traffic sent through these SSIDs is then moved to the 5GHz spectrum, with even higher data rates, more channels, and more advanced coding. Additionally, as the range of a 5GHz signal is substantially shorter than a 2.4GHz signal, the total amount of interference will decrease dramatically, as the amount of air time in use will decrease for the 2.4GHz spectrum, and traffic sent through the air will have a much smaller area of impact.

As previously mentioned, the 2.4GHz spectrum is potentially being retired as a user network at UiT in the coming years, with limited support for legacy equipment for a while. Similar situations were faced when retiring 802.11b as a supported protocol, at which point a total of 6 clients lost connectivity. It is expected that a somewhat larger amount of devices will be affected at this point, but it is believed that this will be for the bettering of the common experience.

11.10 Evaluation

11.11 Future perspective

In the near future, it is believed that major manufacturers[11] will implement their client software with a stronger weighting towards 5GHz usage, and that future devices will support 5GHz as the primary and preferred way of connecting to 802.11 networks.

Further, emerging technologies like HaLow², LoRa³ and similar technologies that primarily will support the growth of the IOT/IOE ecosystems. This means that may become even more important with efficient and flexible management systems for wireless infrastructures, either peer-to-peer or client-station based.

With the adoption of more free spectrum[30] it is still important to manage the limited resource the allocated spectrum represent. As before, the allocated may

2. <http://www.wi-fi.org/discover-wi-fi/wi-fi-halow>

3. <https://www.lora-alliance.org/>

seem endlessly spacious, but when adopted by multiple parties and technologies, it may shrink much faster than expected or foreseen.

/12

Conclusion

A preliminary, extensible and highly parallel library and system for collection of information, metrics and statistics from a enterprise size wireless infrastructure has been implemented. To make use of this information, and to meet a much needed demand, a HTTP-based graphical user interface has been implemented to showcase some of the information being gathered, and to show some of the possible systemic metrics that can be drawn from this.

As the library in some areas just graze the top of the iceberg of what can be found, work on this library and its uses will continue also after this thesis has been written.

Bibliography

- [1] C. Dickens, A Tale of Two Cities. James Nisbet & Company, Limited, 1902.
- [2] Cisco Systems, Inc., “Cisco Prime Infrastructure 3.0 Data Sheet.” <http://www.cisco.com/c/en/us/products/collateral/cloud-systems-management/prime-infrastructure/datasheet-c78-735696.html>, 2015.
- [3] Cisco Systems, Inc., “Cisco Prime Infrastructure 1.3 Data Sheet.” <http://www.cisco.com/c/en/us/products/collateral/cloud-systems-management/prime-infrastructure/datasheet-c78-729877.html>, 2014.
- [4] Cisco Systems, Inc., “Cisco Prime Infrastructure V3.0 API Reference.” <https://developer.cisco.com/site/prime-infrastructure/documents/api-reference/rest-api-v3-0/>.
- [5] Uninett AS, “NAV - Monitor your network with ease.” <https://nav.uninett.no/>.
- [6] Cisco Systems, Inc., “Radio Channel Frequencies.” <http://www.cisco.com/c/en/us/td/docs/routers/access/3200/software/wireless/3200WirelessConfigGuide/RadioChannelFrequencies.pdf>.
- [7] Motorola Solutions, Motorola, Inc., “5Ghz IEEE 802.11a For Interference Avoidance.” http://www.motorolasolutions.com/content/dam/msi/docs/business/_documents/static_files/interference_tb_0809.pdf, 2009.
- [8] Aruba Networks, “RF Basics - Part 1.” http://community.arubanetworks.com/aruba/attachments/aruba/tkb@tkb/121/1/RF-Basics_Part1.pdf, 2007.
- [9] W. Stallings, SNMP, SNMPv2, and CMIP: the practical guide to network-management standards. Addison-Wesley, 1993.
- [10] Airespace, Inc., “AIRESpace-WIRELESS-MIB DEFINITIONS.” <ftp://ftp.airespace.com/ftp/airespace-wireless-mib-definitions.txt>.

cisco.com/pub/mibs/v1/AIRESPACE-WIRELESS-MIB-V1SMI.my, 2010.

- [11] Cisco Systems, Inc. and Apple Inc., “Enterprise Best Practices for Apple Devices on Cisco Wireless LAN.” http://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-2/b_Enterprise_Best_Practices_for_Apple_Devices_on_Cisco_Wireless_LAN.pdf, 2016.
- [12] Bjørn Ludvig Langaas Johansen, “HiPerWA: High Performance Wireless Analytics - An introductory survey into enterprise wireless network analytics,” 2015.
- [13] P. Calhoun, M. Montemurro, and D. Stanley, “Control And Provisioning of Wireless Access Points (CAPWAP) Protocol Specification.” <https://tools.ietf.org/html/rfc5415>, 2009.
- [14] P. Calhoun, M. Montemurro, and D. Stanley, “Control and Provisioning of Wireless Access Points (CAPWAP) Protocol Binding for IEEE 802.11.” <https://tools.ietf.org/html/rfc5416>, 2009.
- [15] P. Calhoun, R. Suri, N. Cam-Winget, M. Williams, S. Hares, B. O’Hara, and S.Kelly, “Lightweight Access Point Protocol.” <https://tools.ietf.org/html/rfc5412>, 2010.
- [16] E. Rescorla and N. Modadugu, “Datagram Transport Layer Security.” <https://tools.ietf.org/html/rfc4367>, 2006.
- [17] E. Rescorla and N. Modadugu, “Datagram Transport Layer Security Version 1.2.” <https://tools.ietf.org/html/rfc6347>, 2012.
- [18] J. Florwick, “Improve Enterprise WLAN Spectrum Quality using Cisco Advanced RF Features.” https://www.ciscolive.com/online/connect/sessionDetail.wv?SESSION_ID=82027, 2015. Video and PDF-slides.
- [19] Cisco Systems, Inc., “Cisco Aironet 1810 Series OfficeExtend Access Points Data Sheet.” <http://www.cisco.com/c/en/us/products/collateral/wireless/aironet-1810-series-officeextend-access-points/datasheet-c78-736868.html>.
- [20] Cisco Systems, Inc., “Wireless LAN Design Guide for High Density Client Environments in Higher Education.” http://www.cisco.com/web/strategy/docs/education/cisco_wlan_design_guide.pdf, 2011.
- [21] Intel Corporation, “USB 3.0 Radio Frequency Interference Impact on 2.4 GHz Wireless Devices.” <http://www.intel.com/content/www/us/en/>

- io/universal-serial-bus/usb3-frequency-interference-paper.html, 2012.
- [22] Cisco Systems, Inc., “Wireless fundamentals: Signal-to-Noise Ratio (SNR) and wireless signal strength.” https://documentation.meraki.com/MR/WiFi_Basics_and_Best_Practices/Wireless_fundamentals%3ASignal-to-Noise_Ratio_%28SNR%29_and_wireless_signal_strength, 2012.
- [23] Magnus Zetterberg, CTO, Telenor Norge, “Ny teknologi for taledekning overallt.” http://www.nkom.no/aktuelt/nyheter/_attachment/20055?_download=true&_ts=1503789014f, 2015.
- [24] M. Oliver and A. Escudero, “Study of different CSMA/CA IEEE 802.11-based implementations.” <http://www.eunice-forum.org/eunice99/027.pdf>, 1999.
- [25] K. Wierenga, S. Winter, and T. Wolniewicz, “The eduroam architecture for network roaming.” <https://tools.ietf.org/html/rfc7593>.
- [26] R. T. Fielding, Architectural styles and the design of network-based software architectures. PhD thesis, University of California, Irvine, 2000.
- [27] L. Daigle and VeriSign, inc., “WHOIS Protocol Specification.” <https://tools.ietf.org/html/rfc3912>. Last visited: 2015-12-06.
- [28] Keith Parsons, “Want, Don’t Want, Don’t Care.” <http://wirelesslanprofessionals.com/wp-content/uploads/2010/01/Want-Dont-Want-Dont-Care.pdf>, 2009.
- [29] Cisco Systems, Inc. and Dave Evans, “The Internet of Everything - How More Relevant and Valuable Connections Will Change the World.” http://www.cisco.com/c/dam/en_us/about/ac79/docs/innov/IoE.pdf, 2012.
- [30] T. Nitsche, C. Cordeiro, A. Flores, E. W. Knightly, E. Perahia, and J. C. Widmer, “IEEE 802.11ad: Directional 60 GHz Communication for Multi-Gbps Wi-Fi.” <http://networks.rice.edu/files/2014/10/11adPaper.pdf>, 2014.
- [31] Cisco Systems, Inc., “Cisco Wireless Controller Online Help, Release 8.1.” <http://www.cisco.com/c/en/us/td/docs/wireless/controller/8-1/olh/wlc-olh-81.pdf>, 2016.
- [32] Cisco Systems, Inc., “Cisco Aironet 3800 Series Access Points Data Sheet.” <http://www.cisco.com/c/en/us/products/collateral/wireless/>

aironet-3800-series-access-points/datasheet-c78-736498.html,
2016.

- [33] E. Mengual Pérez, “Frequency management in a campus-wide wi-fi deployment,” 2013.
- [34] “Ieee standard for information technology– local and metropolitan area networks– specific requirements– part 11: Wireless lan medium access control (mac)and physical layer (phy) specifications amendment 1: Radio resource measurement of wireless lans,” IEEE Std 802.11k-2008 (Amendment to IEEE Std 802.11-2007), pp. 1–244, June 2008.
- [35] D. Stanley and E. Qi, “Status of project ieee 802.11v.” http://grouper.ieee.org/groups/802/11/Reports/tgv_update.htm, 2011.

Appendixes

Included with printed versions of this thesis, there should be a DVD or CD-ROM containing a version of the source code developed in this project, along with digital copies of this report and other relevant materials.

Included with digital versions of this thesis, there should be a link to, or provided a file containing a version of the source code developed in this project, along with digital copies of this report and other relevant materials.

