



**FIRST**®

EINSTEIN INVESTIGATION REPORT

**TABLE OF CONTENTS**

Introduction .....3

    Einstein Match Summary .....4

    Field Setup and Troubleshooting .....5

    Root Cause Analysis .....5

Testing Summary .....6

    Post-Championship Field Testing .....6

        Test 1: Verify Basic Field Operation .....6

        Test 2: Verify Basic Robot Control .....6

        Test 3: Verify Einstein Team Numbers .....6

        Test 4: Verify Einstein Match Sequence .....6

        Test 5: Verify Actual Robot Control .....7

        Test 6: Verify Einstein Matches with Robots .....7

    Wireless Network Security Testing .....7

        Failed Client Authentication .....7

        Denial-Of-Service Attacks .....10

    Einstein Weekend Testing .....11

        Tethered Practice Matches .....11

        Individual On-Field Practice Matches .....11

        Einstein Match Sequence .....12

        Individual Robot Testing .....13

        Lighting Test .....18

        Playing Field Signal Strength Measurements .....18

        Failed Client Authentication Test .....18

Root Cause Conclusions .....20

Next Steps .....22

Appendix A .....24

Appendix B .....25

## INTRODUCTION

In the final rounds of the 2012 *FIRST* Robotics Competition (FRC) Championship in St. Louis, twelve teams, representing the winning alliances from four divisions, competed on the Einstein field to determine the alliance that would be crowned 2012 World Champions. During eight matches, eight of the twelve final teams encountered situations where their robots failed to respond to command signals from the driver station. Loss or disruption in wireless communications and power supply issues on the robot are two examples of issues that may lead to a robot not responding to command signals.

After the Einstein events, *FIRST* headquarters opened an investigation to better understand the root cause(s) of these failures to respond. This report summarizes the testing plan that was conducted along with accompanying results and conclusions. This report also details next steps that *FIRST* headquarters has planned.

To investigate the issues on Einstein, all 12 teams were asked to send their robot and one to two representatives to New Hampshire to participate in testing weekend. All 12 teams accepted the invitation and actively participated. In addition to the team representatives and FRC staff, 18 outside experts and key volunteers were in attendance. A rigorous testing protocol was established with the group of subject matter experts and a series of tests were run; attempting to simulate and reproduce the Einstein issues and ultimately determine any corresponding root causes.

**All 18 subject matter experts and key volunteers and the 12 Einstein teams who participated in the *FIRST* testing weekend have reviewed this report and endorse the Root Cause Conclusions section.**

Key volunteers and subject matter experts who attended testing weekend included:

Eugene Baik (Qualcomm)  
Russ Beavis (DEKA)  
Dave Blumberg (DEKA)  
Bryan Finseth (DEKA)  
Chris Fultz (Rolls Royce)  
Jose Graziani (octoScope)  
Joe Hershberger (National Instruments)  
Jacob Komar (Einstein FTA)  
Paul Malmsten (WPI)  
Greg McKaskle (National Instruments)  
Brad Miller (WPI)  
Al Skierkiewicz (Chief Inspector)  
Liz Smith (Einstein FTA)  
John Toebe (Cisco)  
Eric VanWyk (National Instruments)  
Anurag Vardhan (Qualcomm)  
Derek White (Oracle)  
Bill Wike (Qualcomm Atheros)

In addition to the rigorous testing that was completed, additional information surfaced that is included in this report. Along with the electrical and systems problems that may be expected to occur occasionally in a high-tech competition, there was an intentional act of interference during the Einstein matches.

We encourage you to email feedback on this report to [2012frcfeedback@usfirst.org](mailto:2012frcfeedback@usfirst.org). While we are not able to respond to every email sent, we promise to read each piece of communication, and your ideas will be taken into account as we continue to move *FIRST* forward.

Thank you for your continued patience and support.

## EINSTEIN MATCH SUMMARY

The following table lists the command response failures observed during the Einstein matches, and notes about the issues as observed by field staff at the time. Specific details about the programming language, robot radio version, etc. for each of the robots from the Einstein matches are provided in Appendix A.

Match	Team	Station	Start and Duration of Command Response Failures	Notes
Semi-Final 1-1	118	Red 2	0:04 into hybrid period for remainder of the match	Could ping robot radio from field computers, driver station still showed live camera feed
Semi-Final 2-1	1114	Red 1	final 0:46 of match	
	4334	Red 2	0:05 into hybrid period for remainder of match	
	2056	Red 3	0:44 into teleoperated period for 0:09, final 0:24 of match	
Semi-Final 1-1 Replay	118	Red 2	0:04 into hybrid period for remainder of the match	Could ping robot radio from field computers, driver station still showed live camera feed
Semi-Final 2-1 Replay	1114	Red 1	final 0:16 of match	
	2056	Red 3	0:06 into teleoperated period for 1:33	Start of command response failures seemed to coincide with collision between robot and blue driver wall
Semi-Final 1-2	118	Red 3	0:05 into hybrid period for remainder of the match	Could ping robot radio from field computers, driver station still showed live camera feed
Semi-Final 2-2	2056	Red 1	final 1:30 of match	Used a spare radio (provided by FRC), radio replied to network pings
	1114	Red 2	0:48 into teleoperated period for 0:17	
Final 1	233	Blue 2	0:08 into hybrid period for 0:02, 0:27 into teleoperated period for 0:01, 1:01 into teleoperated for 0:22	Driver station software indicated connection with robot radio, but not cRIO controller
	987	Blue 3	0:57 into teleoperated for 0:06	
Final 2	180	Red 1	1:09 into teleoperated period for 0:44	
	25	Red 2	1:29 into teleoperated period for 0:02	
	987	Blue 3	final 0:19 of match	

## FIELD SETUP AND TROUBLESHOOTING

The Einstein field electronics were set up and tested at night on Friday, April 29<sup>th</sup>, in St. Louis. The mechanical elements of the Einstein field came from a spare and unused field, so that these elements would look brand new for the final matches of the event. The Field Management System (FMS), which includes all electronics and cabling on the playing field, used on Einstein came from another field that was used successfully at four FRC Regionals:

- Chesapeake Regional, Baltimore, MD
- Virginia Regional, Richmond, VA
- Midwest Regional, Chicago, IL
- 10,000 Lakes Regional, Minneapolis, MN

All standard pre-event testing was done on the Einstein FMS, including testing of the scoring counters, E-stop buttons, player station lights and signs, player station Ethernet cables, and the field's wireless access point. All system components passed these tests. The number of active wireless access points was also checked to confirm that the D-Link 1522 radios would not have problems linking with the field access point. A total of 16 access points were detected across the 2.4GHz and 5GHz frequency bands. Because this number was well below 60, and remained at this level for the duration of all Einstein matches, FRC staff had no information to indicate it would be helpful to implement the Code Bondé radios as described in Team Update 2012-01-20. These radios were intended for use only in wireless environments with an excessive number of access points.

During the Einstein matches, members of the FRC staff, as well as external control system development team members, were working on identifying the causes of the command response failures. After observing and monitoring the command response failures during the first two matches, FRC staff members came to the conclusion that there may have been a problem with the field access point. The decision was then made to replay both Semi-Final 1-1 and 2-1, and swap out the field's wireless access point with the unit from the Galileo field. No other field components were swapped out, since a failure in these other devices would exhibit different symptoms and cause a breakdown in communications between FMS and the driver station computers. Communications between FMS and team driver station computers were observed to be solid throughout each match.

## ROOT CAUSE ANALYSIS

As part of the investigation, Ted McKain, formerly the Technical and Quality Director for Rolls Royce North America, currently retired, and a consultant to Rolls Royce, came to FRC headquarters to work with the *FIRST* staff. With his support, a complete list of possible causes for command response failures was compiled into a fishbone diagram. Using log data from the Einstein matches, notes from *FIRST* staff members and volunteers, and information provided by the twelve Einstein teams, many of these causes were ruled out. Of the remaining possible causes, *FIRST* staff, with the assistance of Mr. McKain, identified which causes were testable based on which conditions could be reasonably replicated. For each item that was not going to be tested, a risk assessment was made to determine the likelihood that a root cause could be missed by not testing that item. As testing was completed, the document was updated to identify whether the items under test were confirmed to be a root cause or not, or if they still remained as a potential cause. The final document can be found in Appendix B.

## TESTING SUMMARY

### POST-CHAMPIONSHIP FIELD TESTING

Immediately after the Championship, the Einstein FMS was set up at *FIRST* headquarters for testing. The overall goal of the field tests outlined below was to determine if any of the command response failures could be reproduced. Testing was done using driver station computers (E09 and E11 models), robot simulators (D-Link 1522 robot radios connected to netbooks running a Java application that would simulate a cRIO), and robot control system test platforms (a mix of boxes on wheels and bench-top configurations) built by FRC Engineering. In the following tests 1-6, the field access point from the Galileo field was used. The original access point used on Einstein was not included in these tests because the issues observed in the replays of SF1-1 and SF2-1 were similar to when these matches were played originally, indicating that the causes of the issues were likely systemic and not tied to just a single failure of the field access point.

---

#### TEST 1: VERIFY BASIC FIELD OPERATION

The purpose of this test was to identify any basic system-level problems with the Einstein FMS. A test match was run without any driver stations or robots connected to the FMS. The Einstein FMS, including all field input and output devices, showed no errors during the test.

---

#### TEST 2: VERIFY BASIC ROBOT CONTROL

The purpose of this test was to identify any station-specific problems with the Einstein FMS by running test matches with a single driver station and robot simulator connected. A total of six matches were run, with the driver station and robot simulator in a different station for each match. All six matches ran successfully with no issues identified.

---

#### TEST 3: VERIFY EINSTEIN TEAM NUMBERS

The purpose of this test was to identify any issues with the Einstein FMS related to the identification numbers of the twelve teams that competed on the Einstein field. The SSID's and IP addressing used on the field network is based on the six teams participating on the field. The format is 10.TE.AM.XXX for the various devices used by each team. Incorrect configuration of these settings will prevent a team's driver station and robot from being able to link on the playing field. Six matches were run using the teams from Semi-Final 1 (16, 25, 118, 180, 548, and 2194) represented by FRC Engineering driver stations and robot simulators. For each match, the six teams were rotated one position from the station they were in during the previous match, such that each team was tested in each station. This process was then repeated using the six teams from Semi-Final 2 (207, 233, 987, 1114, 2056, and 4334). No issues were found during the test.

---

#### TEST 4: VERIFY EINSTEIN MATCH SEQUENCE

The purpose of this test was to identify if there were any problems with the Einstein FMS related to the actual sequence of matches that made up the Einstein tournament at the Championship. The Einstein match sequence was played through using FRC Engineering driver stations and robot simulators. The Semi-Final matches were played with the teams in the same stations as they were during the actual Einstein matches. Since the FMS

software automatically populates the teams for the Finals matches at the conclusion of the Semi-Finals, the Finals were played with teams in different stations than the actual Finals in St. Louis. To compensate, an additional six test matches were run after using the same strategy as test #3 above, but using the six teams from the Finals. No issues were discovered during this test.

---

#### TEST 5: VERIFY ACTUAL ROBOT CONTROL

The purpose of this test was to identify issues with the Einstein FMS's ability to interface with actual robot control systems. Three test matches were run using 6 robot test platforms built by FRC Engineering. During the matches, two of the robots exhibited momentary losses of communication with their driver stations. Further investigations into these losses revealed a faulty D-Link 1522 robot radio on one robot, and a problem with the user code on the other robot. An additional test match was run without any communication problems after fixing these two issues. Otherwise, no issues were found during this test.

---

#### TEST 6: VERIFY EINSTEIN MATCHES WITH ROBOTS

This test was a combination of tests #4 and #5 to verify that the Einstein FMS was capable of playing the Einstein match sequence using robot control system test platforms (a mix of boxes on wheels and bench-top configurations) built by FRC Engineering instead of the robot simulators used in test #4. No issues were found during this test.

---

### WIRELESS NETWORK SECURITY TESTING

---

#### FAILED CLIENT AUTHENTICATION

During the Post-Championship field testing an attempt was made to connect to one of the team SSIDs set up on the field network, but the WPA key was entered incorrectly. This was observed to sever the communication between the driver station and robot associated with that SSID. After further testing, a link between failed authentication attempts and a disruption of the communication between a driver station and robot was confirmed, though not every failed authentication attempt resulted in a communication disruption. After the communication had been severed, communication was not reestablished automatically. Additional authentication attempts, whether they succeeded or failed, sometimes reestablished the connection between the driver station and robot. This initial testing was done using an iPad 2 as the client failing authentication.

---

#### ALTERNATIVE ROBOT RADIO TESTING

Testing was then done with different client bridge configurations to determine susceptibility to the failed client authentication vulnerability. The bridges tested were:

- D-Link DAP-1522 Revision A (firmware versions 1.21, 1.31, 1.40, 1.41b)
- D-Link DAP-1522 Revision B (firmware versions 2.00, 2.01, 2.03b)
- Linksys WGA600N
- Linksys WES610N
- D-Link DIR-825 (DD-WRT version 17201).

Only the D-Link DAP-1522 Revision A was affected by the failed authentication attempts, all other tested radios showed no disruption to the communication between the driver station and robot. All tested firmware versions on the DAP-1522 Revision A were found to be affected by this issue.

---

## ALTERNATIVE FIELD ACCESS POINT TESTING

All prior failed client authentication testing was done on the field access point in the configuration used on Einstein, a Cisco Aironet 1252 Access Point (AIR-AP1252AG-x-K9) running firmware version 12.4(25d)JA connected to the rest of the field hardware in the typical configuration. Testing with other alternative access point configurations included:

- Cisco 1252 with 12.4(25d)JA firmware standalone (i.e without the field network)
- Cisco 1252 with 12.4(10b)JA firmware on the field network
- Cisco 1252 with 12.4(10b)JA firmware standalone
- DAP-1522 in AP mode standalone
- Linksys WRT610N standalone

The only Access Point configuration affected by the failed authentication attempts was the Cisco Aironet 1252 access point running firmware version 12.4(25d)JA, regardless of whether it was operated on its own or connected to the field network.

The version of the firmware found to exhibit the behavior was installed in all field access points for events occurring in Week 4 and later. The decision to upgrade the firmware used in the field access points was made based on a recommendation from Cisco to address an issue where the access points were rebooting during the programming sequence that occurs prior to each match. Firmware version 12.4(25d)JA was tested through a large number of simulated matches in order to verify that it would address the issue with the AP rebooting while being programmed, and was tested with robots at *FIRST* to verify proper operation. Attempting to authenticate an additional client, either successfully or unsuccessfully, was not tested as part of this validation. The exhibited behavior stands in sharp contrast to what is expected from an 802.11n system. Wireless experts consulted as part of this investigation were not surprised to find that this was not tested as part of the validation test suite. Prior to Week 4 of the 2012 season the 12.4(10b)JA firmware was used in all FRC field access points.

---

## ALTERNATIVE SOURCE TESTING

The failed client authentication issue was tested and confirmed with the following 802.11n/5GHz-capable client devices as the source of the communications disruption:

- HP Elitebook 8540w (Windows XP)
- HP Elitebook 8540w (Windows 7)
- MacBook Pro (OSX 10.6.8)
- iPad2 (iOS 5.1)
- iPad gen3 (iOS 5.1)
- Samsung Galaxy Nexus (Android 4.0.4)



Each of the above client devices was further evaluated to determine if they attempt to re-associate with the AP after the initial failed attempt, and how often the re-association attempts occur. The results of these evaluations are listed below:

- HP Elitebook 8540w (Windows XP)
  - Repeat interval 5-6 seconds, however communication drops were not observed every 5-6 seconds. In some tests the link would remain up for 20-30 seconds, then drop for 10-15 seconds. In other tests the link would drop immediately and come back minutes later. A test which caused a repeatable condition was not found.
- HP Elitebook 8540w (Windows 7)
  - Repeat interval 75 seconds, however communication drops were not observed every 75 seconds. In some tests the link would drop immediately and come back 75 seconds later. In other tests the link would drop at the first 75 second repeat interval and then would not come back. A test which caused a repeatable condition was not found.
- MacBook Pro (OSX 10.6)
  - Would not attempt to re-associate
- iPad2 (iOS 5.1)
  - Would not attempt to re-associate
- iPad gen3 (iOS 5.1)
  - Would not attempt to re-associate
- Samsung Galaxy Nexus (Android 4.0.4)
  - Test results show that this phone repeats connection attempts only once 20 seconds after an unsuccessful association. This single attempt at re-association was the only recorded attempt within a 10 minute window after the original attempt to connect to the network.
  - It was also observed that it is possible to cause communication drops between the driver station and the robot that last for less than 20 seconds. This can be done by manually selecting one of the field access point SSID's and rapidly issuing the "Connect" command. Testing achieved communication drop intervals down to 3 seconds. When this test was done with multiple robots and driver stations connected to the field network it was observed that the phone refreshes all authentication attempts at the same time. For example:
    1. Initiate client authentication attempt on SSID 1 and wait 20 seconds to observe the second attempt.
    2. Initiate client authentication attempt on SSID 2. The Galaxy Nexus will also attempt to authenticate with SSID 1 again. The phone will attempt to connect with both SSID's on the repeat interval 20 seconds later.

Each of these authentication attempts has the potential to cause working communication to drop and a dropped connection to be reestablished between the driver station and the robot. Repeated attempts to connect to multiple SSID's can result in robots that are drivable and robots that are not over the course of the match.

Over the course of these tests, FRC Engineering was able to determine how to identify a failed client authentication through the log data recorded in the field access point. However, the configuration of the field access points used during the 2012 FRC competitions, including the matches on Einstein, is such that log data is not retained when the access point is powered off.

---

## FAILED CLIENT AUTHENTICATION ON EINSTEIN

While the Einstein matches were in progress, an individual was observed near the field using a cell phone in an apparent attempt to access the field WiFi network. This individual had attempted to engage field personnel in discussions while the field personnel were troubleshooting other issues. This individual was asked to put away the cell phone, and complied. Later, the individual was observed using the cell phone again, and at that point, before the last two Einstein matches were played, was asked to leave the field area, and did so.

After Championship, this individual came forward wishing to share knowledge regarding the failed client authentication issue. The individual claimed to have attempted to connect to the network associated with Team 2056 during Semi-Final 2-1 and observed that this attempt corresponded with the robot losing communication. This individual reported using a Samsung Galaxy Nexus mobile phone in this activity. The match in which this incident occurred, Semi-Final 2-1, was one of the two matches that were replayed after changing the field access point. This individual denied that any attempt was made to access the network of any other robot in that match, or to access the network of any robot during any other match.

FIRST also received reports from witnesses who were near the field during the Einstein matches of someone utilizing a cell phone during the original Semi-Final 2-1, the replay of Semi-Final 2-1, and Semi-Final 2-2. This individual was observed to be pulling up a screen which contained the team numbers of the six teams currently competing, selecting a team, and then rapidly typing in text. In one case, 1114 was seen to have been selected specifically, although the witness to this event does not recall in which match this took place; the replay of Semi-Final 2-1, or Semi-Final 2-2. Witness memories on other specific team selections are less certain. The individual using the Galaxy Nexus phone was observed to be rapidly repeating this process until shortly before the end of each match. The description of this individual provided by these witnesses in these cases fits the individual who came forward as described in the paragraph above.

---

## DENIAL-OF-SERVICE ATTACKS

The most common type of Denial of Service attack that can be performed on a wireless network is known as a “Death Flood”. This involves the attacker sending a series of de-authentication packets to the client, and optionally the access point. By sending a continuous stream of these packets the client can be kept off the network indefinitely. The AirTight system in use on all FRC fields is a wireless intrusion prevention system that is designed to detect these types of attacks, as well as other Wi-Fi events. The purpose of the testing described below was to identify the effect of this attack on an FRC network and the ability of the AirTight system to detect and flag such an attack.

A set of software tools was used to generate Denial of Service attacks on the field network. The first test performed was to verify that this attack could sever the communication between a single driver station and robot by using a broadcast attack against all clients connected to a particular SSID. This attack was successfully used to disrupt communications for an entire match. The AirTight system successfully detected and flagged this test as a “Denial of Service Attack” event. Further testing revealed that it is possible to disrupt communications for any period of time over the course of a match.

The next tests involved targeting the attack at the specific client (the robot radio) and throttling the rate at which the tool sent the de-authentication packets. These attacks were also successful at disrupting the communication, but were able to elude detection by the AirTight sensor. Further investigation into the thresholds for classifying an event of this type in the AirTight system revealed a minimum of 8 directed or 4 broadcast de-authentication

packets per second for a minimum of 90 seconds must be detected for the AirTight system to flag and classify a denial of service event. *FIRST* has determined that these thresholds provide insufficient protection for the FRC wireless use case.

At this time there is no evidence suggesting that an attack of this nature was used against any robot on Einstein, though it cannot be completely ruled out as a cause of any communication disruption based on the log data available.

## EINSTEIN WEEKEND TESTING

Representatives from the 12 Einstein teams, FRC Staff, and the 18 outside experts and key volunteers listed in the Introduction of this report convened in New Hampshire for a weekend of detailed tests in an effort to reproduce the issues seen on Einstein and determine the root cause of each issue.

Each team uncrated their own robot and initial inspections were performed to confirm that no damage occurred to the robots during shipment to New Hampshire.

### TETHERED PRACTICE MATCHES

The first test of the weekend involved running each team through tethered practice matches using the “Practice Match” feature included in the FRC driver station software. The goal of this test was to determine if any of the issues seen on Einstein could be reproduced off of the field. These tests were done immediately after uncrating the robots.

Team 2056’s robot was the only machine to exhibit connection issues during the tethered matches. During the first match, the communications between the driver station software and the robot were lost for the first 7 seconds in the hybrid period. A similar communications loss was observed during the 2<sup>nd</sup> tethered match, but was only one second in duration. It was noted that during the communications loss, the indicator lights on the robot radio seemed to indicate that the radio was rebooting. A 3<sup>rd</sup> match was run, during which communications were lost three times. Each of these losses coincided with the robot radio rebooting.

### INDIVIDUAL ON-FIELD PRACTICE MATCHES

The next test was to run each team individually on the Einstein field in a practice match to again attempt to reproduce any of the issues seen on Einstein in St. Louis. The issues encountered during these tests are listed below.

#### 180

During the match, connection was lost with the robot after it had slammed into one of the bridges on the field. An immediate check of the robot electrical system revealed that one side of the Anderson battery connection was loose. A team representative indicated that both halves of the connection are usually zip-tied to the robot frame for each match, but this had only been done for one half before the test. After zip-tying the other half of the connection to the frame, another match was run without any connection losses. It was agreed that the battery connection was likely the source of the command response failure in this particular test, but not in Finals match 2 on Einstein as a failure of the battery connector as seen in this test would not allow control to be re-established.

---

## 2056

Upon colliding with the barrier in the center of the field while attempting to cross over it, the robot radio on the 2056 robot rebooted, causing a command response failure for 48 seconds. When trying again to cross the barrier, the radio lost power completely, causing a command response failure for the remainder of the match.

Investigation of the power wiring for the radio revealed a broken wire on an Anderson Power Lock connector that connected the wires coming out of DC-DC converter to the wires going to the D-Link radio itself. Team representatives replaced the Anderson Power Lock units with new connectors from the FRC spare parts case and ran another match without issue.

The command response failures during this match were observed to be similar to the failures seen by 2056 in the replay of Semi-Final 2-1. During that match, 2056 was observed to have lost communication around the same time that the robot collided with the blue alliance driver wall. While it is not known for certain if the connector wire was broken during that match, it is a possible root cause of the command response failure.

---

## 4334

During their wireless practice match, 4334 was observed to have packet trip times greater than 100 milliseconds throughout the match. Further investigation revealed that two instances of the WPI Smart Dashboard program were running on the team's driver station computer. Running a 2<sup>nd</sup> test match with only one instance of the dashboard running reduced the packet trip times to a more reasonable, though still elevated, level.

---

## EINSTEIN MATCH SEQUENCE

The next step in the testing was to play the Einstein matches in sequence, with teams in the same stations they were in during the real matches in St. Louis. The goal was to reproduce any of the issues seen during the Einstein tournament that had not yet been reproduced in earlier testing.

During the first match of the test (replication of Semifinal 1-1) Team 180's robot lost communication during hybrid mode. The orange "bridge mode" light on the robot radio was observed to be off. The robot lost communication 4 seconds before the end of hybrid and regained communication with 44 seconds remaining in teleop, an outage of 80 seconds. Shortly before the robot reconnected to the field the "bridge light" on the robot radio was observed to begin blinking again and the robot relinked to the field and remained connected for the remainder of the match.

In investigating this failure after the completion of the test it was discovered that the robot radio power supply was connected to one of the unregulated ports on the power distribution board and not to the regulated 12V radio power supply port. If the robot suffered from a large enough dip in power provided by the main robot battery, it could have caused the radio to lose power and reboot. This is a possible root cause of the command response failure seen by 180 during Final 2 on Einstein.

No other issues were observed.

---

## INDIVIDUAL ROBOT TESTING

To identify any other explanations for the issues seen on Einstein, a number of additional tests were conducted with the robots that suffered from command response failures during the tournament. These tests were unique to each robot, and based on initial feedback about robot construction and programming provided to FRC staff and the experts from the teams in the weeks leading up to the Einstein testing weekend activities. The details and results of these tests are listed below.

---

### 25

Examining the driver station and FMS log files from Einstein Finals match 2, it was observed that during the period of time where team 25 suffered a command response failure all 5 robots that were connected to the field (team 180 had no communication at the time) showed increased trip times and lost packets.

Investigation of team 25's robot and dashboard code showed no indicators suggesting that the robot or dashboard code or configuration may have been related to this network event. Team 25's robot had no issues during any of the tests performed over the weekend.

Observation of the wireless network during testing, including primary metrics such as channel utilization and packet error rate in addition to secondary metrics such as robot trip time and packet loss, revealed a few possible causes of the network disruption, including anomalous behavior by the Smart Dashboard and another robot losing communication. It also remains possible that an external RF interference event, such as a lightning strike, caused the disruption. Any one of these scenarios could have been the root cause for team 25's 2 second loss of communication during Final 2.

---

### 118

Prior to arriving for the Einstein testing, Team 118 identified a potential problem with their code where it would become stuck in a loop when entering hybrid mode if the robot's gyro sensor is unplugged. Testing the robot using the Practice Match mode of the driver station with the gyro unplugged resulted in a loss of communication to the robot after a few seconds. Specifically, the code performs a reset of the gyro accumulator then uses a while-loop to check that the reset has been completed by looping until the gyro value is detected to be less than a small tolerance value. There is a finite amount of time, between when the gyro accumulator is reset and when the while-loop reads the value, during which the accumulator continues to measure the gyro. Typically this additional accumulation would be a very small value resulting from measurement noise. Anything that causes this additional accumulation to be greater than the tolerance of the check will prevent the loop from terminating.

Several possible causes for the unexpected gyro signal were explored. The most likely culprit is a faulty crimp in the connector between the gyro and the analog inputs creating an intermittent connection. When disconnected, it was interpreted as an extremely rapid counter-clockwise rotation. This triggered the initialization routine bug.

Two alternate theories were also tested and proven to be unlikely:

- A gyro's null response (motionless output voltage) will vary with temperature. This is automatically compensated for by calibrating the null response when the gyro first becomes active in the system. However, self-heating or ambient heating may continue to drift the null, causing the gyro to report that it is slowly spinning in one direction. This can be reproduced with a heat gun and/or freeze spray. It was

hypothesized that lighting differences or schedule differences exacerbated this, but we found the effect to be less than the user code tolerance value. Additionally, power cycling the robot would have resolved this potential cause.

- The gyro’s mounting holes are isolated from its ground plane, but it is possible to short to it with large or aggressive mounting hardware. This would lead to a chassis isolation fault, which in turn would direct noise through that subsystem. The gyro hardware was found to be well mounted with no isolation fault.

By itself this issue should cause the code to remain caught in the loop indefinitely, but should not result in a loss of communication. Communication with the driver station should still be maintained and a “Reboot cRIO” command issued using the driver station should be able to reboot the system. The first thing investigated as a potential cause of the command response failure was thread priorities. The FRC communication thread should run with a higher priority than any team code so that it always receives CPU time to process incoming driver station packets. It was determined that all threads were running with appropriate priorities.

During further investigation of this issue it was revealed that packets from the on-board BeagleBone, used for vision processing, were filling up the cRIO’s network buffer causing all subsequent driver station control packets to be rejected. This condition occurred because the portion of the code to read the packets from the BeagleBone was not being reached when the code was stuck in the loop related to the gyro. The network stack for the VxWorks 6.3 operating system used on the cRIO uses a single buffer to store incoming packets from all open sockets. By not reading the BeagleBone packets out of this buffer, they accumulated until the buffer was full leaving no space for the driver station packets, including the command to reboot the cRIO. It was observed during testing that it took approximately 4 seconds for this buffer to fill and cause communication loss; this timing corresponds with the observations and logs from the matches on Einstein.

Testing done during the weekend was able to clearly replicate the symptoms seen by 118 during their three matches on Einstein. The filling of the communications buffer with packets from the BeagleBone, caused by the user code failing to exit the gyro reset loop, was determined to be the root cause of the command response failures 118 suffered in St. Louis.

---

## 180

The team 180 robot showed 100% cRIO CPU usage at all times. This was traced to having empty DisabledContinuous, TeleopContinuous, and AutonomousContinuous in the code. This resulted in those methods being called as fast as possible, utilizing any excess available CPU time. It is unlikely that this behavior would cause a complete command response failure.

---

## 233

In St. Louis, after experiencing connection issues in Einstein Finals Match 1, Team 233 tethered to their robot and was able to connect to the Axis camera, but not to the cRIO. After replacing the Ethernet cable between the cRIO and the robot radio with one provided by one of the Einstein FIRST Technical Advisors they were able to again connect to the cRIO. They played in Finals Match 2 with no communication drops; though they did show increased trip times and packet loss during the same time span as teams 16, 25, 207, and 987.

The first test performed on the 233 robot was an examination of the Ethernet ports on the D-Link 1522 robot radio and on the cRIO. All ports were found to be free of debris and operated normally. All ports retained the Ethernet

cable firmly, none allowed the connector to be withdrawn enough to break contact without depressing the tab on the connector.

The minimum connection drop from an Ethernet cable removal and reinsertion was tested next. The minimum time observed for the connection to be reestablished was one second. The replacement cable used in Einstein Finals 2 had been returned to the FTA at the conclusion of the matches in St. Louis, but using a new cable the robot showed no connection issues during any of the tests performed over the weekend.

The reboot time of the cRIO running the Einstein code was also tested. This reboot was measured to take a minimum of 24 seconds when tethered. Based on this, it is unlikely that the 22 second drop experienced by 233 in Finals match 1 was caused by a cRIO reboot as it is shorter than the minimum reboot time.

Team 233 also had an unusual pattern of cRIO CPU Usage appearing in their driver station logs from all matches throughout the Championship. After talking with the team this was discovered to be a result of a number of “printf” statements in the code which were triggered periodically.

At the conclusion of testing, the seemingly-faulty Ethernet cable used during Finals 1 was determined to be the likely cause of the command response failures for 233. However, without the chance to test the cable itself, this conclusion could not be completely confirmed.

---

## 987

Testing on team 987’s robot was initially focused on the network anomalies that had been detected during the wireless testing and match testing. During this testing a number of network traffic spikes occurred which appeared to be coming from the robot or driver station of team 987. Further testing of the robot pointed to the Smart Dashboard as a potential source of these traffic spikes. A data capture of one of the spikes with the robot tethered and using the Practice Match mode of the driver station software revealed that the network traffic spike appeared to be composed of a large number of 1 byte packets erroneously being sent by the Smart Dashboard client on the driver station PC. In an effort to identify the cause of these spikes the team dramatically increased the amount of data they were logging using the Smart Dashboard. Further testing was unable to reproduce or pinpoint what was causing the Smart Dashboard to send these packets. Similar network traffic patterns were also observed on other robots using the Smart Dashboard.

As part of the Smart Dashboard investigation, the team also used a feature of their dashboard that sent data from the dashboard to the robot. It was discovered that using this feature, combined with the increase in the Smart Dashboard data being sent by the robot, could cause a condition where the robot code would lock-up. The code was then reverted to the state used on Einstein and it was shown that the same condition could be repeated, but required many more invocations of the feature which sent data from the dashboard to the robot. The cause of this lock-up was eventually traced to an issue with the NetworkTables C++ library on the robot which was crashing while in possession of a semaphore lock. The crash was caused by a bad pointer delete in the write thread of the NetworkTables code. When the code reached a point where it attempted to acquire this locked semaphore it would freeze waiting for the semaphore to be released.

The specific feature the team was invoking to reliably reproduce this issue was not a feature they used on Einstein, and based on the controls involved in using the feature it is highly unlikely it was triggered accidentally. As a result of this, it was determined that this bug was unlikely to be the root cause of 987’s command response failures on Einstein.

---

## 1114

In all of the matches on Einstein, 1114 exhibited 100% cRIO CPU usage. As such, individual testing on the robot started with checking the code and underlying frameworks for tasks that could cause such high usage. The user code was found to be executing `Java Thread.yield()`, instead of `Thread.sleep()`, in the vision processing thread. At first, this was the likely suspect of the high CPU usage, but the cRIO still indicated 100% usage even after replacing the `Thread.yield()` with `Thread.sleep()`. It wasn't until the cRIO vision processing was turned off that the CPU usage dropped below 100% (to 60%). It was recommended to the team that they turn down the camera frame rate to avoid 100% CPU usage in the future. While this may have impacted the performance of the cRIO, it is unlikely that it was the source of 1114's command response failures on Einstein.

A number of other software checks were performed:

- Checked NetConsole tool for too many `println()` messages
- Check VxWorks cRIO operating system priority levels
- Checked network traffic with Wireshark for packet floods
- Checked for Java garbage collection overhead
- Checked for memory leaks

No issues were found during these tests that could explain the connection problems seen on Einstein by 1114.

After the software testing was completed, FRC staff members and on-site experts decided to examine the robot's electrical system. During this check, it was discovered that the positive terminal on the main battery connector was not fully inserted. The team reported that just prior to the matches on Einstein in St. Louis they noticed a failure in the original positive wire between the battery connector and main robot circuit breaker. After discovering this failure, the wire and main breaker were both replaced. During the testing weekend, personnel were unable to fully insert or remove the terminal, suggesting the terminal or connector body may have been defective. Closer inspection of the terminal showed none of the typical marks that result from normal use. It is possible that this could have caused intermittent power losses during the matches on Einstein in St. Louis. The 1114 robot did not suffer from any command response failures during the matches prior to this discovery, but the terminal was pushed as far into the connector as it would go before playing any additional matches.

Later in the inspection of the electrical system the main breaker's "off" button was found to be abnormally sensitive; a light tap on the button could cause a brief brown-out condition in the electronics. Tapping the off button harder (but not with enough force to fully actuate the button) was found to reboot the camera, and on occasion, the D-Link robot radio and cRIO robot controller. To determine if this sensitivity was a potential source of the issues seen on Einstein, the robot was brought to the field for more testing. The team representatives drove the robot aggressively throughout a test match, purposefully slamming into walls, bridges, and barriers at high speeds in an attempt to trigger the button on the main breaker. At no time during the on-field testing was the team able to trigger the button and cause a brown-out. A combination of reports from 1114, 2056, and reviews of video footage from 1114's matches on Einstein also confirm that the robot signal light remained on and flashing for the duration of all three matches in which 1114 participated, indicating that the cRIO never lost power. Based on these results, it is unlikely that the sensitive main breaker would have caused command response failures during the actual Einstein matches.



---

## 2056

In Semi-Final match 2-2 on Einstein, team 2056's robot failed to respond to commands for the final 90 seconds of the match. During this time, the robot radio still replied to network pings sent from the field computers, indicating that the radio was still connected to the wireless network.<sup>1</sup> Typically, this symptom is indicative of a power loss to the cRIO, or a broken Ethernet connection between the radio and cRIO. Individual testing was done on the 2056 robot to identify any explanation for this symptom.

Testing was conducted to identify any broken or unreliable Ethernet ports on the radio or cRIO. Each port was tested one-by-one, physically straining the port via a connected Ethernet cable. The robot was also driven aggressively on the field without encountering any issues. These tests indicated that a broken or faulty Ethernet port was not likely a cause of 2056's command response failures on Einstein.

To verify if a cRIO power loss could have been a root cause, the boot-up time of 2056's cRIO was measured. This time was found to approximately 25 seconds. Based on this data, a loss of cRIO power could explain the command response failure seen in Semi-Final 2-2, but only if the power loss was permanent. Because this scenario would result in the robot signal light (RSL) turning off, and the RSL on 2056's robot was observed to remain blinking during the match, a permanent loss of cRIO power could not have been the cause of the command response failure in Semi-Final 2-2.

---

## 4334

Throughout the weekend testing and sporadically through all of their driver station logs from the FRC Championship, Team 4334 showed trip time measurements higher than those typically experienced by the other robots on the field at the same time. They also consistently showed 100% CPU usage on their cRIO. The initial focus of the testing to identify the source of this anomaly was the Smart Dashboard.

First, the robot was run using the Practice Match feature of the driver station with the Smart Dashboard client on the driver station PC closed. This test showed no observable change in behavior. The robot was then tested with a number of the Smart Dashboard calls removed from the robot code. This test still did not alleviate the 100% CPU usage condition on the cRIO.

The complete robot program was then examined and two issues were found which contributed to the 100% cRIO CPU usage condition. The first thing that was discovered was that the code implemented the DisabledContinuous, TeleopContinuous, and AutonomousContinuous methods with no Wait or Sleep calls. Each of these methods is called as fast as possible by the code framework which will result in 100% CPU usage when the robot is in the mode corresponding to the method. The other issue that was found was the use of while-loops in the hybrid code with no Wait or Sleep call, again allowing them to run as fast as possible. Addressing these two issues resulted in the trip time dropping substantially to match the other teams on the field, which was confirmed by running both a match with just 4334 on the field and a match with a full complement of 6 robots on the field.

It is unlikely that the code issues found on 4334's robot relating to the 100% cRIO CPU usage would be able to cause a complete command response failure for the duration experienced in the initial playing of Semifinal 2-1.

---

<sup>1</sup> Network pings are sent using packets that conform to the ICMP protocol, which is different from the UDP protocol used for the control and status packet sent between the driver station and the robot. Further detail on FMS and the way that communication is handled between the driver station and robot will be included in the FMS white paper as detailed in the Next Steps section at the end of this document.

---

## LIGHTING TEST

One of the differences between the Einstein field and all four of the division fields was the lighting conditions. The purpose of this test was to investigate whether that difference could cause control or connection issues.

Lighting was brought in to replicate the brightness and approximate positioning of the lighting used on the Einstein field. The teams who were using vision processing of any kind on Einstein were identified and tested under these lighting conditions. This list included teams 16, 118, 233, 987 and 1114.

No command response failures were observed during these tests and no latency or lag in controls were reported. It was observed that the quality of the tracking for a number of the teams did seem to decrease under the Einstein lighting compared to other tests using the facility's fluorescent lighting. The facility fluorescent lighting conditions were similar to the lighting used on the division fields.

---

## PLAYING FIELD SIGNAL STRENGTH MEASUREMENTS

Measurements were taken on the playing field to determine the RF power level of the wireless signal transmitted by the field access point. The following was measured using an RF power meter:

- -34dBm directly in front of the field access point.
- -44dBm when standing in the far corners of the playing field; directly in front of Player Station Blue 1 and Red 3.
- Power levels dropped by ~10dB (-44dBm to -54dBm) when several people stood between the field access point and the playing field.

An additional two tests were performed to evaluate the signal strength in extreme cases. Both field access points on Einstein had all antennas installed in the correct location.

- With no antennas installed on the field access point, the measured signal levels dropped by ~30dB (-64dBm to -74dBm)
- With the 2.4GHz antennas installed on the 5GHz terminals of the field access point the signal levels dropped by ~10dB (-44dBm to -54dBm).

These signal levels are within acceptable limits for the D-Link 1522 radio, however actual signal levels received by the D-Link 1522 may be further reduced due to radio placement on the robot, RF interference, and obstructions caused by other robots.

---

## FAILED CLIENT AUTHENTICATION TEST

The goal of this test was to test the effects of failed client authentication on other robots during a match. Based on information received by FRC in the weeks leading up to the testing weekend, the testing was focused on 2056 and their Einstein matches.

The first test was to run Semi-Final 2-1 and attempt to stop communication between 2056's driver station software and their cRIO robot controller. FRC staff members used a 3<sup>rd</sup> generation Apple iPad to successfully cause 2056 to lose communication by attempting to connect to the wireless network for that team and entering an incorrect passkey. The other five teams on the field reported no difference in controls or communications with

their robots during the match. A second match was run identical to the first, but with a narrower wireless channel to increase the accuracy of network packet capture tools being used to analyze robot communications. This second match produced results identical to the first.

As stated earlier, team 2056 suffered a command response failure for the final 90 seconds of the Semi-Final 2-2. During those last 90 seconds their robot radio was still responsive to network ping requests. In an attempt to reproduce this symptom, Semi-Final 2-2 was run as an additional failed client authentication test. To replicate the match, representatives from 2056 swapped out their team's robot radio for the same spare radio they used during Semi-Final 2-2 in St. Louis. FRC staff members were again able to successfully cause a command response failure for 2056, during which the spare robot radio was indeed responsive to network ping requests sent by field computers. This behavior matched the symptoms seen during Semi-Final 2-2 in St. Louis, indicating that failed client authentication is a possible explanation for the command response failure seen by 2056 during that match.

## ROOT CAUSE CONCLUSIONS

The table below lists the root-cause conclusions for each command response failure as reached through the Einstein investigation. Each root cause has a probability listed to indicate the level of confidence that the listed root cause was the source of the command response failure. These probability levels are:

- **Confirmed:** conclusion is well supported by extensive testing and data, and there is little or no chance that something else could have caused the command response failure.
- **Likely:** the conclusion is supported by extensive testing and data, but there remains a small possibility that something else caused the command response failure.
- **Plausible:** the conclusion is supported by some testing and limited data. The listed conclusion is the strongest explanation for the command response failure, but there is reasonable possibility that something else could be the cause.

Match	Team	Conclusion	Probability
Semi-Final 1-1	118	cRIO user code failed to exit gyro reset loop, custom electronics filled communications buffer preventing cRIO from receiving control packets from driver station	Confirmed
Semi-Final 2-1	1114	Failed Client Authentication	Plausible
	4334	Failed Client Authentication	Plausible
	2056	Failed Client Authentication	Confirmed
Semi-Final 1-1 Replay	118	Same as Semi-Final 1-1 above	Confirmed
Semi-Final 2-1 Replay	1114	Failed Client Authentication	Likely
	2056	Failed Client Authentication	Plausible
		Robot radio reboot due to wiring problem	Plausible
Semi-Final 1-2	118	Same as Semi-Final 1-1 above	Confirmed
Semi-Final 2-2	2056	Failed Client Authentication	Likely
	1114	Failed Client Authentication	Likely
Final 1	233	Faulty Ethernet cable between robot D-Link radio and cRIO	Likely
	987	Failed Client Authentication	Plausible
Final 2	180	Robot D-Link radio reboot due to power dip	Likely
	25	Field network congestion	Likely
	987	Failed Client Authentication	Plausible

Detailed summaries of how the conclusions above were reached are listed below by team:

- **25** – Thorough inspection of the robot electrical system and software revealed no faulty components or code. At the time of the brief command response failure, all teams on the field were observed to have higher packet round trip times than normal, suggesting RF interference or general field network congestion.
- **118** – This chain of events was replicated during the Einstein testing weekend and caused symptoms identical to those seen in all three of 118's matches on Einstein.
- **180** – The robot radio was found to be drawing its electrical power from an unregulated port on the power distribution board; brief but low dips in the main battery voltage could cause the radio to reboot. The duration of the command response failure is approximately equal to the time it takes the robot radio to reboot after a power loss. The team reported attempting to use the "Reboot cRIO" feature of the driver

station during the communication outage based on experience with a code failure at a prior event. A code failure is unlikely to be the source of the command response failures on Einstein as no such failure was reproduced during the testing.

- **233** – The team was unable to connect to the cRIO when tethered through the radio off the field after the match, indicating a faulty Ethernet cable between the cRIO and radio. The cable was replaced, and no other command response failures were experienced. The duration of the command response failures for this robot were too short to be attributed to a cRIO or robot radio rebooting, and Ethernet ports on the robot were found to be clear of debris.
- **987** – While testing of this robot revealed programming issues that could cause higher than normal trip times, or the cRIO controller user code to lock up, none of these issues were found to cause a command response failure. The duration of these losses were too short to be attributed to a cRIO or robot radio reboot, but fit well with the symptoms of failed client authentication.
- **1114** – While the positive terminal on the main battery connector was not fully inserted, and the main breaker had a sensitive “off” switch, rigorous and aggressive testing of the robot on the field did not result in any command response failures. Because of this, failed client authentication is a stronger explanation of the command response failures, especially in light of the witness report of this activity. Additionally, the cRIO on the robot was tested to determine how long it takes to reboot after a power interruption. The command response failure in Semi-Final 2-2 was shorter in duration than the measured reboot time, indicating that a loss of power to the cRIO was not the cause.
- **2056** – There was a confirmed report of failed client authentication as the cause of the command response failure for this robot in the original Semi-Final 2-1. In the replay of Semi-Final 2-1, the D-Link robot radio was unresponsive to network pings, suggesting that the radio may have lost power. During initial Einstein weekend testing, the radio power wiring was found to be faulty. Further testing later in the weekend revealed that this robot radio exhibited the same unresponsive behavior when failed client authentication was used. Therefore, either scenario could explain this command response failure. During the command response failure in Semi-Final 2-2, the replacement radio used on the robot was responsive to network pings and the robot signal light remained blinking for the entire match. Testing during Einstein weekend revealed that the radio was also responsive when failed client authentication was used. While it is possible that the cRIO lost Ethernet connection to the radio, thorough inspection of the robot electrical system did not turn up any faulty connections or components that could cause this.
- **4334** – Thorough examination and testing of this robot did not reveal any potential explanations for the command response failure. A complete inspection of the robot electrical system did not reveal any loose connections or damaged components that could cause complete power loss to the radio or cRIO. No issues found in the robot code could result in the code completely crashing for the duration of the match. Failed client authentication is the strongest remaining explanation.

## NEXT STEPS

Based on the results of this investigation and the input of outside experts after observing the system during testing, FRC is planning on investigating a number of improvements for the 2013 season and beyond.

A number of the improvements FRC had already planned to make were reinforced by the events of Einstein:

- Improved Robot Diagnostics – Improvements will be made to the robot diagnostic information available to both teams and volunteers. This will likely include changes to the information stored in the driver station logs, improvements to the driver station log file viewer, and changes to the FMS logs and Field Monitor. At minimum, information about the status of the robot radio (ping status) will be added to these logs. Additional documentation on the DS logs and how to use them as a troubleshooting and debugging tool will be created. Additional methods or tools to assist teams in debugging robot code issues, particularly on the competition field, will be investigated including changes to the NetConsole, additions of easy-to-use interfaces for logging to disk, and documentation on ways to use dashboards and any new features to debug robot behavior.
- FMS White Paper – FRC will be producing a white paper which describes how the Field Management System operates. This will include details on the topology of the system, components used and the communications paths between the various components.

The following changes are planned as a direct response to issues that affected matches on Einstein:

Root Cause	Mitigation
cRIO user code failed to exit gyro reset loop, custom electronics filled communications buffer preventing cRIO from receiving control packets from driver station	Improved Documentation on Advanced Coding Practice Matches
Failed client authentication	Resolve Failed Client Authentication Issue Field Hardware Logging
Robot radio reboot due to wiring problem	Investigate New Robot Radio
Faulty Ethernet cable between robot D-Link radio and cRIO	Improved Robot Diagnostics
Robot D-Link radio reboot due to power dip	Additional emphasis in training and documentation
Field network congestion	Implement Bandwidth Limiting

- Improved Documentation on Advanced Coding – Documentation on advanced code features, particularly threading and networking will be improved.
- Resolve Smart Dashboard Issue – implement a fix to the NetworkTables C++ library used on the robot to prevent crashing while in possession of a semaphore lock.
- Practice Matches - FRC will consider providing time for all Einstein teams to play practice matches on the Einstein field prior to the start of the official matches
- Resolve Failed Client Authentication Issue – There are a number of potential solutions that will be evaluated for the failed client authentication flaw that was discovered including a firmware or settings resolution from Cisco for the field access point, a firmware or settings resolution from D-Link for the DAP-1522 robot radio, and changing robot radios. Additionally, testing for this vulnerability has been added to the testing suite performed whenever changes to the network configuration are being considered. FRC

will revert to firmware version 12.4(10b)JA on the Cisco access point in order to avoid the Failed Client Authentication issue at 2012 off-season events.

- Field Hardware Logging - A system to collect logging information from applicable field hardware (Access Point, Router and Managed Switches) will be implemented on the FMS Server.
- Investigate New Robot Radio – FRC will investigate changing to a new robot radio for the 2013 season. Potential improvements as a result of this change include simplifying the power path by using a 12V radio, improved radio diagnostics and logging, and faster boot time.
- Implement Bandwidth Limiting – Each team should have a defined, fixed bandwidth cap to ensure that the network does not become overloaded. Additional testing will be done to determine the appropriate bandwidth cap for proper operation in all venues including those with crowded wireless environments.
- Camera Calibration – FRC will consider providing time for all Einstein teams to calibrate their cameras on the Einstein field prior to the start of the matches.

Planned changes to the wireless system to increase robustness were confirmed by the feedback from the wireless experts consulted as part of the investigation. These items do not directly address failures that occurred on Einstein but aim to make the wireless network configuration more robust:

- Reduce Robot Data Usage – With a fixed bandwidth cap in place it is important to provide teams with the tools and documentation necessary to work within the cap. Documentation would include information on common camera settings and corresponding data usage rates as well as information on how to determine the data usage of a team’s current setup. Default settings of the camera configuration tool and dashboards will be evaluated to ensure they are using appropriate settings which minimize data usage.
- Quality of Service (QoS) – With a fixed bandwidth cap in place, it becomes critical to prioritize robot control packets over other types of data such as video. QoS can be used to implement this prioritization so that robot control packets will continue to flow even if a team exceeds the bandwidth cap with video or other data.
- Additional changes to the wireless network configuration will be evaluated based on the input and observations of the experts present at the Einstein weekend testing.
- De-Authentication Attack – A number of solutions to the network susceptibility to a “De-Auth Flood” attack will be investigated including working with AirTight to modify thresholds for detection, implementing custom capture and detection of these attacks, and changing radios to a custom solution which is more resistant to these attacks.
- Over-The-Air Monitoring – The feasibility of a system to capture all wireless traffic during matches and analyze that traffic will be evaluated. This data would be used to identify a variety of network attacks, identify poor network conditions by evaluating data rates being used, and determine bandwidth usage of individual robots in order to work with teams that are right up against the bandwidth cap.
- Access Point Location, Orientation and Antennas – A number of possible avenues are being investigated to improve the signal conditions on the playing field. These changes range from relocating and re-orienting the access point to utilizing separate antennas mounted in appropriate locations on the field.
- Utilize Multiple Channels – Currently all robots are run on a single 40MHz wireless channel. Pros and cons of moving to a system that utilizes multiple channels will be evaluated. Many commonly available over-the-air packet capture tools are limited to, or work better on, 20MHz channels.

**APPENDIX A - ROBOT CONFIGURATION INFORMATION**

Division	Team	Language	Camera	Video to DS	Vision	cRIO	CAN	DS	DS 5GHz	Dashboard	Dlink Hardware	Dlink Firmware
Newton	548	Java	2-> 1 used	320x240/30FPS	No	4	Yes-Serial/8 Jags	E09	No	Customized LV Dashboard	A	1.4
	118	C++	2	320x240/17FPS	640x480/17FPS/OnBoard	4	No	E09/Cypress Board	No	Customized LV Dashboard	A	1.4
	2194	Java	1	320x240/30FPS	NA	8	No	E09	No	Default Dashboard	A	1.21
Galileo	16	C++	1	320x240/30FPS	320x240/30FPS/DS	4	Yes-2CAN/1 Jags	Thinkpad T410/ Cypress Board	Yes	Customized SmartDashboard	B	2.0
	25	C++	0	NA	NA	4	No	Acer Aspire One A0722-0473	No	Default Dashboard	B	2.02
	180	C++	1	320x240/30FPS	No	4	Yes-Serial/1 Jag	Acer Aspire One Dual-core A0722-0667	No	SmartDashboard	A	1.21
Archimedes	2056	C++	0	NA	NA	4	No	AcerAspire D257-1814	No	Default Dashboard	A	1.21
	4334	Java	0	NA	NA	4	No	E12	No	SmartDashboard		
	1114	Java	1	No	160x120/cRIO	4	No	ASUS U31SD-AH31	No	SmartDashboard	A	1.21
Curie	987	C++	2	320x240/30FPS x2	Yes	4	No	?		Customized SmartDashboard	A	1.21
	233	C++	1	320x240/30FPS	320x240/cRIO	4	No	Dell	No	Customized LV Dashboard	A	1.4
	207	C++	1	320x240/30FPS	NA	8	No	E09	No	Default Dashboard	B	2.0



**APPENDIX B - ROOT CAUSE ANALYSIS**

Level	Description	Root Cause	Comments
<b>1</b>	<b>Multiple robots inoperable on Einstein 2012</b>		
1.1	Process/Method		
1.1.1	Robot Startup Sequence	No	No issues reproduced by allowing robots to sit before match
1.1.2	Troubleshooting during the event		
1.1.2.1	Field Troubleshooting	Yes	Unable to identify problem leading to multiple/similar failures
1.1.2.2	Robot Troubleshooting	Yes	Unable to identify problem leading to multiple/similar failures
1.1.3	Inspection	Potential	1 Robot radio power wired incorrectly
1.1.4	Radio configuration process at event	No	All robots connected to field to start match
1.1.5	Operational field test at event	No	No problems identified with field hardware
1.1.6	Camera calibration on field	No	No issues occurred in Lighting Test with Einstein robots
1.1.7	FMS startup sequence	No	No evidence at this time
1.1.7.1	Standby time	No	No evidence at this time
1.1.8	Link tests	No	All robots connected to field to start match
1.1.9	Training		
1.1.9.1	Event staff	Potential	Refer to 1.1.2
1.1.9.2	Teams (Driver's Meeting, etc)	No	All robots connected to field to start match
1.1.10	Documentation		
1.1.10.1	Robot Rules	No	No Robot Rules issues identified
1.1.10.2	Control System Manual	No	No evidence at this time
1.1.10.3	Getting Started Guide	No	No evidence at this time
1.1.10.4	LabVIEW Instruction Manual	No	No Einstein robots used LabVIEW
1.1.10.5	C++ Instruction Manual	Yes	Improve documentation on Sockets and Threads

Level	Description	Root Cause	Comments
1.1.10.6	Java Instruction Manual	Yes	Improve documentation on Sockets and Threads
1.2	Environment		
1.2.1	Physical obstruction	No	Nothing observed
1.2.2	Ambient temperature in venue	No	Nothing observed
1.2.3	Humidity	No	Nothing observed
1.2.4	Human congestion (people between AP and robots)	No	Obstruction testing with Einstein robots did not cause disconnects
1.2.5	Light		
1.2.5.1	Luminescence	No	No issues observed in Lighting Test with Einstein robots
1.2.5.2	Heat	No	No issues observed in Lighting Test with Einstein robots
1.2.6	Power from venue to field	Potential	Power fluctuation caused by storm? Should cause failures across all robots; other tests more conclusive
1.2.7	Electrical Interference		
1.2.7.1	Ambient inside venue		
1.2.7.1.1	WiFi	No	Measured, nothing observed
1.2.7.1.2	Other	Potential	Can't measure, but not likely based on feedback from experts
1.2.7.2	Introduced inside venue		
1.2.7.2.1	WiFi	Potential	# of clients was high because of cell phones Can't measure, but not likely based on feedback from experts
1.2.7.2.2	Other	Potential	Phones in audience, tablets, etc. Can't measure, but not likely based on feedback from experts
1.2.7.3	External		
1.2.7.3.1	WiFi	No	Measured, nothing observed, improbable given size of dome

Level	Description	Root Cause	Comments
1.2.7.3.2	Other	Potential	Lightning strikes were abundant around venue
1.2.8	ESD	No	No robot ESD events reported.
1.2.9	Vibration/Shock		
1.2.9.1	Arena vibration	No	Nothing reported/observed
1.2.9.2	Robot to Robot impact	Potential	Nothing observed, but hard to eliminate
1.2.9.3	Robot to field impact	Potential	1 robot hit wall and disconnected at same time
1.3	Management		
1.3.1	Staffing decisions	No	Field crew were all qualified/exemplary
1.3.2	Decision authority	No	All parties were included in decisions, no one made decisions who wasn't qualified
1.3.3	Schedule	Potential	Pressure to troubleshoot quickly, no time allocated for camera calibration/practice matches
1.3.4	Budget	No	No "cut corners" or budget issues identified
1.4	Equipment		
1.4.1	Radio Kiosk	No	Radios were connecting to field, configuration was set properly
1.4.2	Robot		
1.4.2.1	Radio power supply	Yes	Radio power supply plugged into unregulated 12V on 1 robot
1.4.2.2	Sensors	Yes	Issue with 1 robot partially caused by Gyro
1.4.2.3	Analog/Solenoid breakout boards	No	No issues found with Analog or Solenoid breakouts
1.4.2.4	Mechanical systems	No	Nothing observed/reported
1.4.2.5	Camera	No	No reports of camera failure
1.4.2.6	Motors	No	No reports of motor failure
1.4.2.7	Cables/Wiring	Yes	Loose connections found on multiple robots

Level	Description	Root Cause	Comments
1.4.2.8	Battery	No	Nothing alarming on the log, no damage/failure reported
1.4.2.9	Power Distribution Board (PD)	No	No un-identified robot power issues in testing
1.4.2.10	Digital Sidecar (DSC)	No	No reports of issues/failures
1.4.2.11	Radio		
1.4.2.11.1	Hardware	Yes	D-Link DAP 1522 Rev A vulnerable to Failed Client Authentication issue
1.4.2.11.2	Firmware	No	No issues discovered with specific firmware revision
1.4.2.11.3	Orientation	Potential	Signal strength differences based on orientation
1.4.2.11.4	Settings	Yes	WPA 2 security vulnerable to Failed Client Authentication issue
1.4.2.12	Non-KOP electronics	Yes	Issue with 1 robot partially related to BeagleBone processor board
1.4.2.13	Motor controllers/Spikes	No	No CAN or general issues reported
1.4.2.14	cRIO		
1.4.2.14.1	cRIO Software		
1.4.2.14.1.1	vxWorks	Yes	Issue with 1 robot related to single buffer for all network sockets
1.4.2.14.1.2	Usercode/Framework	Yes	Issue with Team 118 related to User Code
1.4.2.14.1.3	FPGA image	No	No issues observed during season
1.4.2.14.2	cRIO Hardware	No	No cRIO damage observed
1.4.2.15	Driver Station (DS)		
1.4.2.15.1	Software		
1.4.2.15.1.1	Dashboard	No	No bandwidth saturation events observed
1.4.2.15.1.2	User code	No	No user code (outside the Dashboard) observed at this level.
1.4.2.15.1.3	DS software	No	No issues observed
1.4.2.15.1.4	Operating System (OS)	No	Nothing observed

Level	Description	Root Cause	Comments
1.4.2.15.2	Hardware		
1.4.2.15.2.1	Joysticks	No	Nothing observed
1.4.2.15.2.2	DS computer	No	Driver Stations maintained communication with FMS
1.4.2.15.2.3	Other team provided input devices	No	Nothing observed
1.4.3	Field		
1.4.3.1	Uninterruptable Power Supply (UPS)	No	Components powered by UPS not a potential Root Cause
1.4.3.2	Power Strips	Potential	Access Point powered by Power Strip
1.4.3.3	A/P		
1.4.3.3.1	Orientation	Potential	No issues in testing, may have greater effect in noisy environment
1.4.3.3.2	Power Supply	No	No issues observed in testing
1.4.3.3.3	Firmware	Yes	Newer firmware allows for FCA vulnerability
1.4.3.4	Estop	No	No estops hit, reported
1.4.3.5	Touch screens	No	No symptoms reported
1.4.3.6	Scoring Hardware and PLC	No	No symptoms reported
1.4.3.7	Classmate power supply	No	No power issues observed
1.4.3.8	Switch (Scorpion case)	No	No issues observed in testing
1.4.3.9	Router	No	No issues observed in testing
1.4.3.10	Cables	No	Collective assessment of the team concurs not an issue. Unable to definitively test DS cables, but can test AP cable.
1.4.3.11	Field indicators	No	Symptoms do not indicate these
1.4.3.12	FTAA computer	No	Symptoms do not indicate these
1.4.3.13	Netbook (pit display)	No	Not in use on Einstein
1.4.3.14	SCCs	No	No issues observed in testing
1.4.3.15	Kinect Stations	No	No issues identified with computers
1.4.3.16	Servers		

Level	Description	Root Cause	Comments
1.4.3.16.1	OS	No	Symptoms do not indicate these
1.4.3.16.2	FMS	No	Symptoms do not indicate these
1.4.3.16.3	Hardware	No	Symptoms do not indicate these
1.5	People		
1.5.1	FIRST HQ Staff		
1.5.1.1	Technical staff provides incorrect/risky information/instruction	Yes	Teams were instructed to configure security at home, feeds a previously unknown vulnerability
1.5.1.2	Admin staff provides incorrect/risky information/instruction	No	Nothing reported/observed
1.5.2	Intentional		
1.5.2.1	WiFi interference	Yes	Connecting to Team SSID with phone
1.5.2.2	Robot damage	No	Nothing reported/observed
1.5.2.3	Field damage	No	Nothing reported/observed
1.5.3	Unintentional		
1.5.3.1	WiFi interference	No	Affected robots did not have Driver Stations with 5GHz enabled
1.5.3.2	Robot damage	No	Nothing reported/observed
1.5.3.3	Field damage	No	Nothing reported/observed
1.5.4	Event staff		
1.5.4.1	Referee	No	Nothing reported/observed
1.5.4.2	Control System Advisor (CSA)	No	Nothing reported/observed
1.5.4.3	Scorekeeper	No	Nothing reported/observed
1.5.4.4	Regional Directors (RD)	No	Nothing reported/observed
1.5.4.5	FIRST Technical Advisor (FTA)	No	Nothing reported/observed
1.5.4.6	MC/Announcer	No	Nothing reported/observed
1.5.4.7	Inspectors	No	Nothing reported/observed