ixia

Enterprise Master Test Plan

IxVeriWave

Contents

Introduction: Lifecycle Management of Enterprise Class Wireless Networks

Managing the lifecycle of enterprise-class wireless networks requires the ability to quickly and effectively execute a wide range of disparate tasks. The figure below outlines the significant steps of this lifecycle, but while the graphic shows an orderly progression from one step to the next, the reality is that many of these steps overlap, occur simultaneously, or happen out of order.

Those managing wireless networks understand the frustration points of living the lifecycle. While there is not one tool that does it all, the IxVeriWave family of products is designed with this lifecycle in mind and shares measurement technology across the family – making Ixia's solution the only one on the market that extends from the lab to the desktop, scaling from one to thousands of clients.

Scope

This document is divided into chapters, each of which describes the best-practice approach for that step. Users of this document can turn to the chapter pertaining to the task immediately in front of them and get to work. As steps are completed, users are rewarded with a solid foundation of data and test configurations ready for use and providing leverage across other steps of the lifecycle. Learning to use IxVeriWave tools to select client devices, for example, sets the user up perfectly for verifying new driver releases for wireless NICs, verifying new client devices, or assessing the performance of a given access point (AP).

Some steps, like selecting infrastructure equipment, happen only occasionally. Other steps such as responding to user complaints, happen daily. But best-in-class network management touches on all the steps at some point, and enables significant increases in operational efficiency through the reuse of data and configurations across the steps.

Chapter 1: Selecting Infrastructure Equipment

Selecting the most suitable equipment for your network infrastructure should be more than reading the data sheets, comparing prices, and getting the sales pitch from the competing vendors. The cost of the communications infrastructure (access points, controllers, switches, and routers), will most likely be the single largest investment in your wireless networks, so choosing wisely is critical. However, making the wrong choice may have an interesting side effect – the long-term cost of supporting a network that is ill-equipped to handle user needs may exceed the initial investment. Therefore, utmost care must be given to this choice.

The process of selecting wireless network equipment should include the step of subjecting samples of that equipment to stressful and customer-specific loads, including the mix of traffic, number and types of client devices and deterministic measurement of quality of experience (QoE) delivered to each client device. The scale of the test needs to exceed the requirement of the environment so that margin-to-failure is observed, and establish the amount of capacity the network solution provides for extending into the future. Measurements need to be repeatable, quantitative, and controllable.

Typically called "bake-offs" or "shootouts," subjecting candidate network solutions to such tests allows the consumer to observe:

- *The ability of the network to meet design goals*
- *What's required to properly configure the network to run at scale (time and complexity)*
- *Stability of the network under stress*
- *Performance of the network under stress and at scale with measurements appropriate to type of data*

Benefits

By performing the testing described in this section, the following benefits are derived:

- *Numerical, quantitative comparison of wireless network alternatives*
- *Performance expressed in end-user Quality of Experience (QoE) metrics*
- *Test configurations ready for re-use*
- *Foundation of performance data for isolating network issues from client issues*

Overview

This testing is done in three basic steps:

1. Perform benchmark testing of throughput, TCP Goodput, latency, roaming, client capacity, and rate vs. range.
2. Perform QoE testing with varying client count/client mix and stressful traffic loads of mixed type, appropriate to the end user environment.
3. Perform testing with actual client devices.

Steps one and two can only be performed with equipment that is proven to be capable of exceeding the capability of the network under test, otherwise the process is using the network to measure the stressing equipment and not the other way around.

The equipment must be capable of running at IEEE 802.11 and 802.3 line rate, and present each AP involved with the required client count and behavior. For this reason, the only known method of performing this testing in an effective, economical, and repeatable fashion is with test equipment designed for the task. Step three is performed with actual client devices that are important to the end-user environment, but done at scale and with companion/competing traffic as will be encountered in the deployed environment.

Test Strategy

The system under test (SUT) includes the controller, switch, and access points from the candidate vendor. The Test System is:

- *For Steps 1 and 2 above: IxVeriWave Ethernet and Wireless WaveBlades, use the Benchmarking and WaveQoE Applications. Connect the APs to WaveBlades using RF cables, and use the test applications to perform testing. A representative test setup is shown below.*

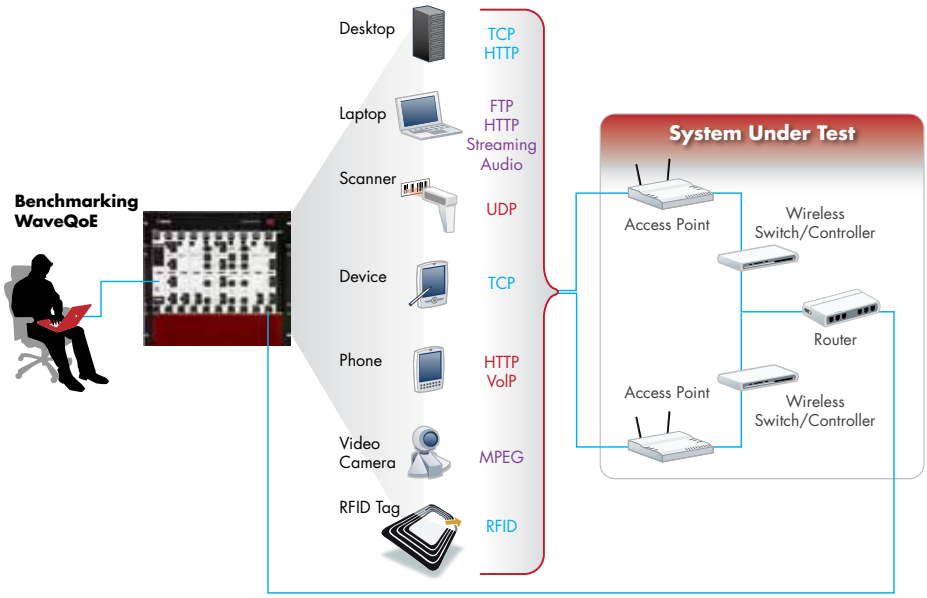

- *For Step 3 above: The IxVeriWave Ethernet WaveBlade or the portable EF1101 Ethernet Test Appliance, using the WaveInsite and WaveDeploy applications. A representative test setup is shown below.*

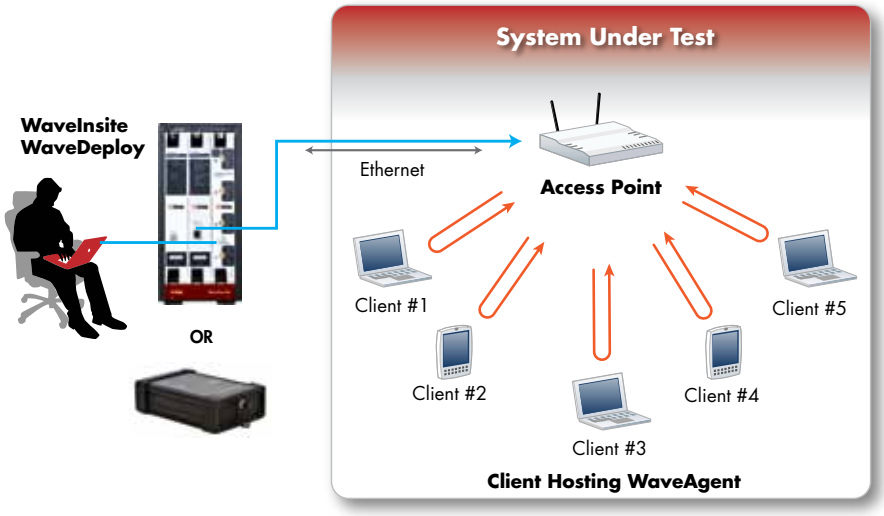

Directly connecting the test equipment to the SUT allows measurement of the SUT's ability to support performance, independent of the variable of RF/radio performance. If first measurements are taken in open air and do not meet expectations, the problem of differentiating between RF-related causes and other non-RF network problems arises. Best practice starts with establishing performance metrics without the RF link being a variable in the test, including putting the candidate access points into RF enclosures. If measurements obtained in this cabled environment do not meet expectations, performance will only get worse once the network is deployed and actual client devices are involved.

1. Benchmark Tests 1 AP per Radio

Maximum Client Capacity

- All connected clients are serviced with data
- What this test reveals:
  - » Maximum functional connected client count

»

UDP Throughput

- Default frame lengths (64, 88, 128, 256, 512, 1024, 1280, 1518)
- Upstream, downstream and bidirectional
- Open and Secure
- 5, 20, and 64 clients
- What this test reveals:
  - » Maximum forwarding rate at zero loss in each direction for different size frames

- Shows the maximum packets/second the system can handle (short frame performance)
- 802.11n downstream aggregation efficiency
- Ability for AP to get access to the medium in the presence of many clients
- Performance curve of traffic vs. client count
- Penalty for security encryption
- TCP Goodput

- Default MSS settings (536, 1460)
- Upstream, downstream and bidirectional
- 5, 20 and 64 clients
- What this test reveals:
  - Maximum TCP goodput rate
  - TCP performance curve of traffic vs. client count

Latency

- 1000-byte packets, UDP
- Upstream, downstream and bidirectional
- 5, 20 and 64 clients
- 50% of forwarding rate figure from throughput testing
- What this test reveals:
  - Latency as a function of client count
  - Latency profile in each direction

Roaming Delay

- Involve multiple APs
- Open and secure
- Involve appreciable number of clients 100-200
- What this test reveals:
  - Ability for SUT to handle roaming clients
  - Packet loss, roam time, dropped roams

2. WaveQoE Test: 2 Radios per AP (as applicable), 2-4 APs

Mixed Client/Traffic Nominal Performance Verification

- Establish voice, video, and data traffic representative of enterprise
- Spread traffic across multiple clients, incorporate legacy and 802.11n mixture
- Load both radios of an AP simultaneously

- Verify Service Level Agreements (SLAs) met for all traffic types and clients
- What this test reveals:
  - » Issues with traffic prioritization
  - » Performance with heterogeneous traffic and client mix

Mixed Client/Traffic Maximum Performance Verification

- Repeat above test while increasing total load incrementally
- Continue increasing load until first SLA failures are observed
- What this test reveals:
  - » Correct traffic gets prioritized in overload condition
  - » Actual maximum transport capability of client-specified traffic mix Performance with Evolving Client/Traffic Mix
- Repeat above test, moving client population to what is expected at end of network life
- Repeat above test, adjusting traffic mix to reflect usage trend over network life
- What this test reveals:
  - » How much margin is available in network solution to accommodate change/growth
  - » How will network perform under special-case conditions

3. Open-Air Client Testing

Key Client Device Performance Testing with WaveInsite

- Establish maximum forwarding rate performance available with key client devices
- Measure upstream, downstream UDP and TCP forwarding rate
- Measure close to AP, then in distance increments of 10 feet up to maximum AP spacing
- What this test reveals:
  - » Initial view of Client/AP interaction under stressful load
  - » Maximum performance available from client device
  - » Performance degradation with distance, each direction
  - »

Scale

This testing approach is easily scaled to fit the requirements of the enterprise. The following table is a guide to hardware needs for different size organizations.

Small Office — Single, stand-alone AP, no independent controller	• Steps 1 and 2: WT20 with (1) WBW2000 MIMO blade, (1) WBE1101 Ethernet Blade. Tabletop testing, 500 client capability, test one radio at a time • Step 3: EF1101 or WT20 with (1) WBE1101 Ethernet Blade. Open air testing
Small / Medium Office – Small scale controller with up to 4 APs, or distributed controller architecture	• Steps 1 and 2: WT20 with (1) WBW1104N SISO blade (4 ports per blade, individual radios), (1) WBW2000 MIMO blade, (1) WBE1101 Ethernet Blade, Tabletop testing, 2000 total clients. Test all four APs at once, one radio at a time, or test both radios per AP on two APs. • Step 3: EF1101 or WT20 with (1) WBE1101 Ethernet Blade. Open Air Testing
Office Building – One or more controller, multiple APs	• Steps 1 and 2: WT90 with (2) WBW2000 MIMO blades, (1) WBW1104N SISO blade, (1) WBE1104 Ethernet blade (4 ports). Tabletop or rack-mount testing, 3000 client capability, test many combinations of radios and APs. • Step 3: EF1101 or WT20 with (1) WBE1101 Ethernet Blade. Open air testing.
Campus – Multiple controllers, multiple APs	• Steps 1 and 2: WT90 with (2) WBW2000 MIMO blades, (5) WBW1104N SISO blades, (1) WBE1104 Ethernet blade (4 ports). Tabletop or rack-mounting testing, 11000 client capability, load multiple APs and controllers simultaneously • Step 3: EF1101 or WT20 with(1) WBE1101 Ethernet Blade. Open air testing.

Return on Investment

Network managers who have invested in this approach for making network infrastructure purchasing decisions enjoy the following returns:

- Strong basis of comparison data to aid in detailed discussions with network vendors
- Ability to directly measure and validate vendor claims when software or configuration changes are expected to deliver performance improvements
- Ability to quickly test and verify suspected faulty APs during deployment phase
- Ability to present repeatable and quantifiable evidence to both network and client vendors then communicating issues for action
- Capacity figures, measured at scale with relevant traffic and client mix, for use when planning the network topology and AP placement
- Acceptable load levels and SLA settings for use in client testing, overall site assessment, and troubleshooting

Chapter 2: Select Client Devices, Determine, Interoperability and Coexistence

Client devices such as smart phones, tablet-PCs, wireless-scanners, and laptops are the interface between the user and the network and have a significant effect on the experience the user will observe. Client devices change rapidly with new devices or versions being introduced every few months, and new software builds constantly being downloaded.

IT managers typically have less control over client devices, yet they represent the most critical link to end user Quality of Experience (QoE). The fact that each client device is manufactured by a different vendor, and all need to coexist and interoperate with the chosen infrastructure makes them the weakest link in the overall network.

Although the process of client selection, interoperability verification, and coexistence assessment can be viewed as three separate steps, it is best addressed as one comprehensive test that addresses all aspects of client behavior. After all, a "good" client that interoperates well with the APs but interferes with other clients is as disruptive as one that doesn't work well with the network infrastructure. The process of selecting wireless client devices should subject each of them to loads and traffic types as can be expected in the "production network," and include deterministic measurement of the QoE delivered.

The testing should be performed first in an environment that is RF-clean (such as a shielded RF room) to establish baseline client performance, and then in a variety of physical locations with respect to the placement of access points and in the presence of other client devices to ensure performance. The test equipment needs to exceed the capability of the applications that will run on the client device so that margin is observed.

Subjecting candidate devices to such tests allows the user to observe:

- Absolute best performance that can be expected from a client device
- Stability of the client device under stress
- Performance of the device under stress, at scale with measurements appropriate to type of data
- Simultaneous operation with the network and in the vicinity of other client devices
- Performance at various locations in an environment
- Whether the client device connects and passes the required amount of traffic at all required locations, without fail
- Whether the client and network manage the connection and traffic forwarding duties correctly as the client device moves through the physical environment

Benefits

By performing the testing described in this section, the following benefits are derived:

- Numerical, quantitative comparison of wireless client devices
- Performance expressed in end-user experience metrics
- Test configurations ready for re-use
- Foundation of performance data for isolating client issues from network issues
- Confirmation of NIC and network configuration variables

Overview

This testing is done in 4 steps:

1. Perform benchmark testing of forwarding rate, goodput, and throughput with different traffic types. Establish QoE metrics that are achievable and acceptable.
2. Repeat testing at various distances and orientations characterizing performance versus location applicable to the deployment environment.
   Add other client devices, and check for client interaction.
3. Perform load testing at a known favorable network location, close to an access point. Include ecosystem traffic and other client devices to provide realistic environment.
4. Repeat testing at challenging physical locations. Include behaviors that actual users will subject the client device to, such as physical movement of the devices.

Steps 1 and 2 above are performed with WaveInsite with WaveAgent installed on the client device. The Ethernet source/sink of traffic on the network side is either a WBE1101 WaveBlade or an EF1101 hand-held Traffic Generator Analyzer. Traffic performance measurements made with WaveInsite while the client is in an RF isolated (or at least clean) environment provide the ability to establish best-case performance for the client device, and set reasonable traffic loads and metrics for later testing in the deployed case. Assuming this testing is being done on a network that has already been benchmarked for performance per Chapter 1, the user can then be assured that the test equipment and network combination exceed the capability of the client device, and therefore measurements can be attributed to the client device.

Steps 3 and 4 above are performed with WaveDeploy. The WF1101 can be used to create ecosystem clients and traffic, and can also be used as a high-performance packet capture device during troubleshooting.

WaveInsite is also very useful during troubleshooting. Preferably, WaveAgent is installed on the client device and all other companion devices, and the Ethernet source/sink of traffic on the network side is an EF1101 hand-held Traffic Generator Analyzer.

Test Strategy

The device under test (DUT) is the client product such as a smart phone, netbook, or laptop. The test system is WaveInsite and WaveDeploy. Other client devices with which the DUT needs to be compatible are included to verify co-existence. All testing is done in open air, as shown below.

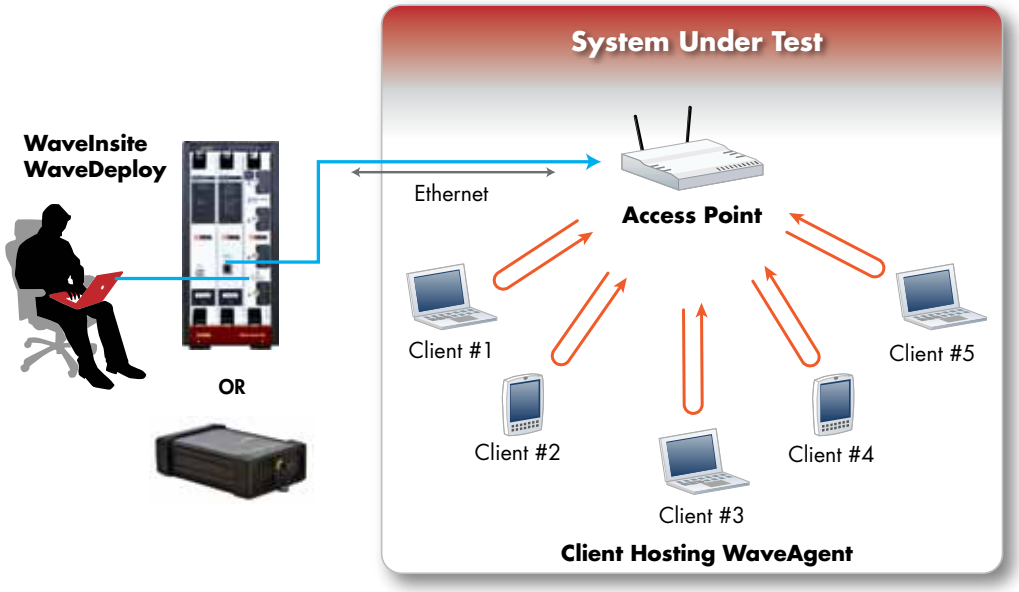

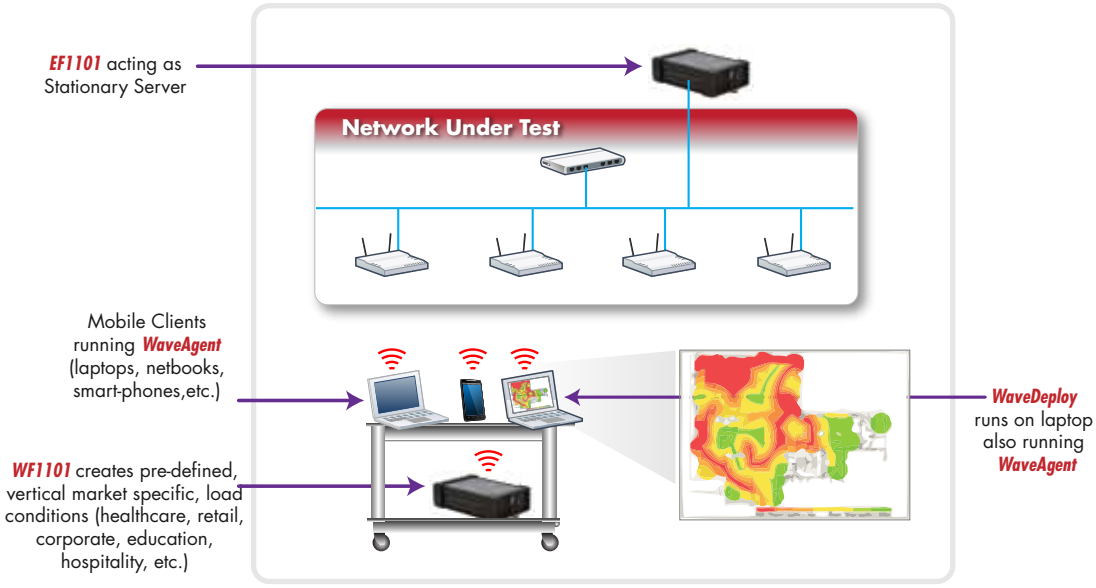

1. WaveInsite Testing (Controlled-RF Environment) – 1 AP, 1 Client Device

Maximum UPD Forwarding Rate

- Upstream, downstream, bidirectional
- Multiple frame lengths (64, 512, 1024, 1518)
- Open and Secure
- Power Save
- What this test reveals:
  - » Shows the maximum packets/second the DUT can handle (short frame performance)
  - » Shows maximum forwarding rate in each direction
  - » Penalty for security encryption
  - » 802.11n aggregation efficiency

TCP Goodput

- Default MSS settings (536, 1460)
- 1, 2, and 4 streams per flow type
- Upstream and downstream
- What this test reveals:
  - » Ability to sync downstream TCP flows, per session
  - » Ability to secure TCP flows

Voice / Video / Data Mixed Flow

- HTTP downstream, Voice bi-directional, UDP upstream and downstream
- Create flows that reflect applications
- Upstream and downstream
- What this test reveals:
  - » Maximum performance for client NIC to carry mix of data with appropriate QoS

Rate vs. Range

- Run at 10 ft distance increments from AP
- Perform rate vs. range test while both client and AP are under stress
- Traverse range until traffic stops
- What this test reveals:
  - » QoS under stress
  - » Quality of NIC receiver and transmitter
  - » Effectiveness of rate adaptation algorithm in client device

2. WaveInsite Coexistence Testing –1 AP, 2 or More Client Devices

Application Performance with Other Clients Present

- Enable one or more additional client devices on the same AP and BSSID as the DUT. For special-purpose devices, run actual customer application (e.g. patient monitor)
- Repeat tests in 1) above and assess effect on DUT and other devices
- What this test reveals:
  - » Client device's response to other co-existing traffic on same band/AP
  - » Network configuration requirements specific to the client device

3. WaveDeploy Testing (Optimum Location)

Single-Client Coexistence and Speed Test

- Use Ixia's "WLAN Site Assessment: Best Practices for Pre- and Post-deployment Verification" guidelines for Target Load settings and SLAs
- Select all traffic types for coexistence and speed test traffic that are of interest
- Client device configured per IT guidelines (security, power save, etc)
- Make multiple measurements from same location to demonstrate repeatability
- What this test reveals:
  - » Verifies client device/network combination can deliver required performance under stress
  - » If Target Load or SLA settings need to be adjusted
  - » If there are hidden issues with optimum location
  - » Whether unexpected variation between repeated measurements is cause for investigation
  - » Whether, if performance cannot be achieved, adding other clients and competing traffic will make things worse

Multiple-Client Coexistence and Speed Test (DUT Plus Other Relevant Client Devices)

- Use "Best Practices" guidelines for other client devices
- Select all traffic types for coexistence and speed test traffic that are of interest for ALL client devices
- All client devices configured per IT guidelines
- Make multiple measurements from the same location to demonstrate repeatability
- What this test reveals:
  - » Client-to-client interaction; if speed test results are inconsistent with coexistence results
  - » If Target Load or SLA settings need to be adjusted for multiple-client case
  - » Whether unexpected variation between repeated measurements is cause for investigation

Multiple Client Coexistence and Speed Test, Plus Ecosystem Traffic (if required)

- To increase client count and amount of traffic in the collision domain, use WF1101
- Use WaveDeploy Expert to create ecosystem of up to 64 clients per WF1101 with constant background traffic (multiple WF1101s can be used for very high client count)
- Repeat Multiple-Client Coexistence and Speed Tests from above
- Make multiple measurements from the same location to demonstrate repeatability
- What this test reveals:
  - » Confirms interoperability under highest client count condition with stressful traffic load
  - » If Target Load or SLA settings need to be adjusted for multiple-client/ ecosystem case
  - » Whether unexpected variation between repeated measurements is cause for investigation
  - » Client-to-client interaction; if speed test results are inconsistent with coexistence results
  - » If Target Load or SLA settings need to be adjusted for multiple-client case

4. WaveDeploy Testing – Challenging Locations

Radio Power-Up at Challenging Location

- Use multiple client coexistence and speed test, plus ecosystem traffic test configuration from above
- Adjust Target Loads and SLAs downward if locations warrant a decrease in expected performance
- Turn on client devices and ecosystem traffic clients at the challenging location
- Make multiple measurements from the same location to demonstrate repeatability
- If performance is not as expected, repeat with ecosystem traffic removed, then other client devices removed to diagnose
- What this test reveals:
  - » Interoperability demonstrated under worst-case condition when DUT and other clients are first connected at that location
  - » Rate adaptation, radio performance, airtime fairness, AP selection algorithm performance at the edge of the network
  - » Appropriate Target Loads and SLA settings at the edges of the network

Radio Power-Up at Challenging Location

- Test loads and SLA will be identical to prior test
- Turn on ecosystem traffic at the challenging location
- Turn on all client devices at another location so that all connect to an AP other than the preferred one at the target location
- Make one measurement, then move all clients to the challenging location

- Make multiple measurements at the challenging location
- If performance is not as expected, repeat with ecosystem traffic removed, then other client devices removed to diagnose
- What this test reveals:
  - » If clients stick to a non-optimum access point as they migrate to the challenging location
  - » Through repeated measurements, it is possible to observe how quickly the client device(s) change their connection to the correct AP, and what degree of disruption occurs during the process
  - » Provides an excellent test method for evaluating "roaming aggressiveness" settings
  - » If client does not resolve connection in a reasonable time, packet captures with the WF1101 reveal the characteristics of the poor link (phy rate, retries, rate adaptation churning)

Scale

For smaller organizations that have relatively small client counts, performing these tests with WaveInsite, WaveDeploy Pro and a small number of additional client devices may be sufficient. For larger organizations or those with a large variety of client devices, use WaveInsite and WaveDeploy Expert.

This is preferred for two reasons: 1) the ability to use the WF1101 for ecosystem traffic and packet capture, and 2) the use of the EF1101 as the Ethernet traffic source/sink. The preferred approach is summarized in the following table:

Small Office — Single, stand-alone AP, no independent controller	• WaveInsite & WaveDeploy Pro
Small / Medium Office – Small scale controller with up to 4 APs, or distributed controller architecture	• WaveInsite & WaveDeploy Expert
Office Building – One or more controller, multiple APs	• WaveInsite & WaveDeploy Expert • Use multiple WF1101s if more than 64 ecosystem clients
Campus – Multiple controllers, multiple APs	• WaveInsite & WaveDeploy Expert • Use multiple WF101s if more than 64 ecosystem clients

Return on Investment

Network managers who have invested in this approach for determining client behavior and performance, network/client interoperability, and client coexistence, enjoy the following returns:

- Strong basis of comparison data to aid in detailed discussions with client and network vendors
- Ability to directly measure and validate vendor claims when driver, software or configuration changes are expected to deliver performance improvements
- Confirmation of the NIC configuration settings for client devices, and ability to verify for each driver revision
- Ability to present repeatable and quantifiable evidence to both network and client vendors when communicating issues for action
- Capacity figures, measured on clients alone and in the presence of other client devices, for use when planning the network topology and AP placement
- Acceptable load levels and Service Level Agreement settings for use in site assessment and troubleshooting

Chapter 3: Set Service Level Goals

When realistic service goals are established for your network, user expectations can be managed, capacity can be allocated, and the amount of trouble-tickets minimized. The risk of not testing before network turn-on and not setting SLAs is severe: users will complain, support costs will increase, and adoption by users will decline making it hard to demonstrate ROI.

Establishing Target Loads and Service Level Agreement (SLA) settings for WaveDeploy ensures that conducting a site survey with a known set of client devices results will be indicative of where corrective action needs to be taken. By setting the Target Loads and SLAs in accordance with the traffic and performance levels supporting business needs, any location that shows as "green" can be trusted to deliver the required performance. Any location that shows as "red" can be assumed to be a problem spot that must be resolved.

While there are default loads and SLA settings available within Ixia's WaveDeploy solution, personalizing them to meet the needs of a specific organization eliminates the potential for "false alarms" if the defaults are set too high, or "missed alarms" if set too low. To get satisfactory performance, it is required that all elements of 802.11 transmissions are operating at expected and satisfactory levels. These elements are:

- RF signal power (RSSI) at every location
- Demonstrated transport rates at every location (Phy rate)
- Clearly preferable AP at every location, with sufficient separation from other APs and networks (Co-channel Interference)
- No interference from other client devices (including those clients outside your control)
- Sufficient data transport capability by network and client device (Target Loads and SLA)

If there are problems with any of the above attributes, end user experience is negatively affected. These elements are highly interactive and confusing to untangle. Often, the observable symptom is poor traffic delivery performance, but the cause is co-channel interference or clients that are connected to the wrong AP (resulting in extremely poor transport rates and a high number of retries).

The key point is that the only way to ensure performance is to run stressful traffic loads at locations of interest, and observe the end-user quality-of-experience measures. If everything is in order, good traffic QoE figures will result. If any one thing is amiss, the QoE measurements will reflect this. In other words, it is not best practice to measure one ingredient or two ingredients of performance, and make claims based on assumptions.

Testing is best done in two phases. The first phase is to test at stress levels in a location favorable to the network, with all client devices involved simultaneously. The second phase is to move the client devices to a worst-case physical location and repeat testing. The test equipment used must exceed the capability of the applications that will run on the client device, so that margin is observed. Measurements need to be repeatable, quantitative and contollable.

Benefits

By performing the testing described in this section, the following benefits are derived:

- Demonstration of sustainable performance under real-life deployment conditions
- Performance expressed in end-user experience metrics
- Test configurations ready for deployment testing and site assessments
- WaveDeploy Mapping accurately pinpoints where corrective action is required

Overview

Testing is done in two steps:

1. Perform load testing with WaveDeploy at a known favorable network location, close to an access point. Include ecosystem traffic and other client devices to provide realistic environment. Adjust loads and SLAs as described in the Best Practices Guide.
2. Repeat testing at challenging physical locations. Verify that applying the Target Loads and achieving the SLA corresponds to the level of traffic a typically network user will require for satisfaction. Both steps are best performed with WaveDeploy Expert. The WF1101 can be used to create ecosystem clients and traffic, or can be used as a high-performance packet capture device during troubleshooting.

WaveInsite is also very useful during troubleshooting. Preferably, WaveAgent is installed on the client device and all other companion devices, and the Ethernet source/sink of traffic on the network side is an EF1101.

Test Strategy

The device under test (DUT) is the combination of the network and client devices. The test system is WaveDeploy Expert. If specific loads need to be applied to the DUT for long periods of time for troubleshooting, WaveInsite is used. All testing is done on a deployed network, in open air, as shown below:

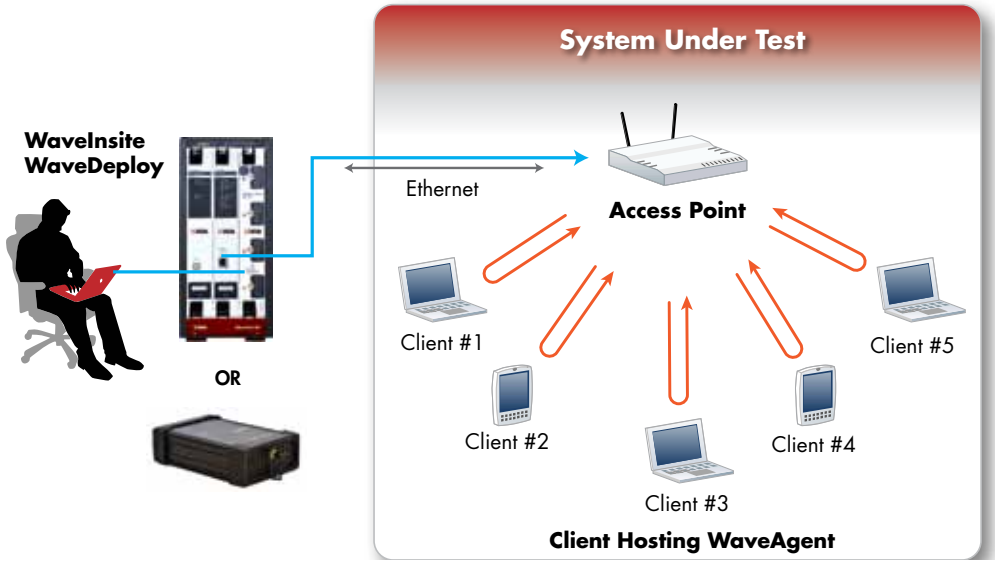

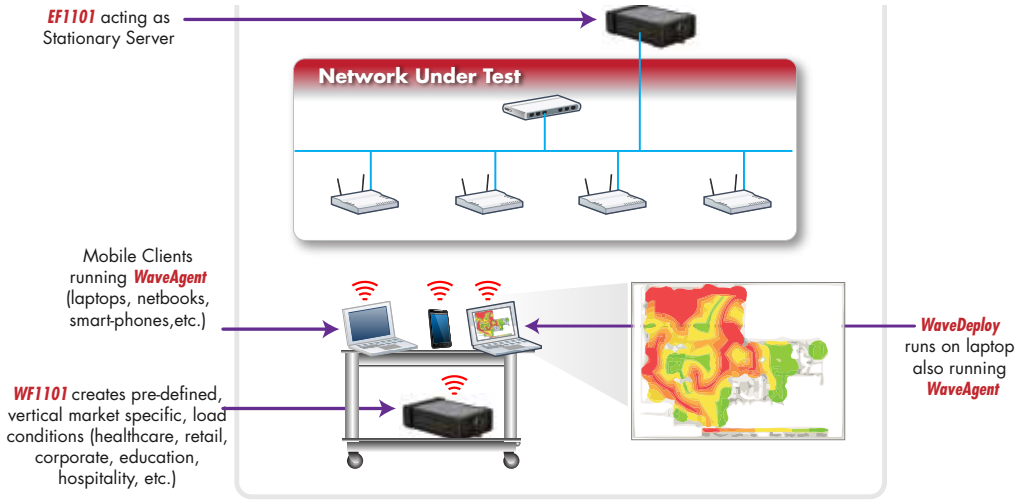

1. WaveDeploy Testing (Optimum Location)

Representative Clients, Coexistence and Speed Test

- Use Ixia's "WLAN Site Assessment Best Practices for Pre- and Post-Deployment Verification" guidelines for Target Load settings and SLAs
- Adjust Target Loads and SLAs to represent the requirements of the enterprise
- Select all traffic types for coexistence and speed test traffic that are of interest
- All client devices configured per IT guidelines (security, power save, etc.)
- Make multiple measurements from same location to demonstrate repeatability
- Verify that all results meet their SLA, or make adjustments as necessary to get to "all green"
- What this test reveals:
  - » Configuration issues with clients or network, if target loads and SLA cannot be sustained repeatedly
  - » What margin exists between required performance and maximum observed performance
  - » If there are hidden issues with optimum location
  - » Unexpected variation between repeated measurements is cause for investigation
  - » If clients interact with each other, discrepancies will appear between coexistence tests and speed tests

Representative Clients, Coexistence and Speed Test, Plus Ecosystem Traffic (if required)

- To increase client count and amount of traffic in the collision domain, use WF1101

- Use WaveDeploy Expert to create ecosystem of up to 64 clients per WF1101 with constant background traffic (multiple WF1101s can be used for very high client count)
- Repeat Coexistence and Speed Tests from above
- Make multiple measurements from the same location to demonstrate repeatability
- Adjust Target Loads and SLA downward as necessary to get to "all green"
- What this test reveals:
  - » Confirms interoperability under highest client count condition with stressful traffic load
  - » If Target Load or SLA settings need to be adjusted for multiple-client/ ecosystem case
  - » Unexpected variation between repeated measurements is cause for investigation

2. WaveDeploy Testing (Challenging Locations)

Representative Clients, Coexistence and Speed Test

- Repeat similar test from above, check that the results against SLA are still "all green"
- Adjust Target Loads and SLAs downward if locations warrant a decrease in expected performance
- Make multiple measurements from the same location to demonstrate repeatability
- If performance is not as expected, repeat with ecosystem traffic removed, then other client devices removed to diagnose
- What this test reveals:
  - » Interoperability demonstrated under worst-case condition when clients are at that location
  - » Rate adaptation, radio performance, airtime fairness, AP selection algorithm performance at the edge of the network
  - » Appropriate Target Loads and SLA settings at the edges of the network

Representative Clients, Coexistence, Speed Test Plus Ecosystem Traffic

- Test loads and SLA will be identical to prior test
- Turn on ecosystem traffic at the challenging location
- Repeat above test and ensure SLAs are all "all green"
- Make multiple measurements from the same location to demonstrate repeatability
- What this test reveals:
  - » Target Loads and SLA set correctly for "all green" performance under worst-case conditions
  - » Any results obtained in site assessment or diagnostic testing that do not meet SLA is indicative of a problem that needs to be resolved
  - » Lack of repeatability for a given measurement indicates a marginal performance issue

Scale

All organizations should conduct this test with the complement of client devices that are used in the enterprise, with traffic loads on each type of device that represents all clients of that type within the radius of influence for that measurement. For situations where large client count is required, WaveDeploy Expert with WF1101 should be used in addition to the actual client devices. The preferred approach is summarized in the table below:

Small Office — Single, stand-alone AP, no independent controller	• WaveDeploy Pro. At least one of every type of client device included in the test.
Small / Medium Office – Small scale controller with up to 4 APs, or distributed controller architecture	• WaveDeploy Expert. At least one of every type of client device included in the test. Use EF1101 and WaveInsite for advanced troubleshooting.
Office Building – One or more controller, multiple APs	• WaveDeploy Expert. At least one of every type of client device included in the test. Use EF1101 and WaveInsite for advanced troubleshooting. Use Multiple WF1101s if more than 64 ecosystem clients.
Campus – Multiple controllers, multiple APs	• WaveDeploy Expert. At least one of every type of client device included in the test. Use EF1101 and WaveInsite for advanced troubleshooting. Use multiple WF1101s if more than 64 ecosystem clients.

Return on Investment

Network managers who have invested in this approach for setting Target Loads and SLAs enjoy the following returns:

- Proven settings with relevant client devices for WaveDeploy
- Ability to quickly repeat test in diagnosis of a trouble ticket
- Confirmation of the NIC configuration settings for client devices, and ability to verify for each driver revision
- Ability to present repeatable and quantifiable evidence to both network and client vendors when communicating issues for action
- Capacity figures measured at on clients alone and in the presence of other client devices for use when planning the network topology and AP placement
- Simultaneous operation with the network and in the vicinity of other client devices
- The client and network manage the connection and traffic forwarding duties correctly as the client device moves through the physical environment
- Rapid and accurate visual analysis of WaveDeploy result maps

Chapter 4: Plan and Deploy the Network

When it comes to network planning and deployment, every network vendor/installer will have recommendations and planning tool to assist in that process. This section is not intended to compete with those tools or recommendations, but rather work in concert with them.

Simply put, this chapter is about performing a limited number of quick measurements to confirm several things as the planning and deployment is being done. During planning, there will be questions about how to best cover certain locations or situations, and proposed channel plans, AP placement or network configurations will need to be resolved. Being able to quickly propose and conduct experiments to answer key questions of performance enables the vendor/installer to make decisions confidently based on relevant data.

During deployment, best practice is to verify every access point, controller, and switch as it is placed into service. Being able to run a significant level of traffic quickly at each AP to verify that all connections and the device itself are performing as required helps to eliminate call-backs. If the AP is able to put out good beacons, but has trouble sustaining line-rate 802.11n traffic, the best time to find this out is while the ladder is right there.

This testing is straightforward. For deployment testing, simply run WaveDeploy with at least one client device at a location near the given access point, and verify that all results come up meeting SLA settings. If the activities outlined in the section above (Set Service Level Goals) have been completed, the user already knows what performance to expect out of each AP, and it becomes a simple matter of verifying that each newly placed AP works like all the others. If not, find the issue before moving on. Keep in mind that the issue could be a neighboring network or client devices, not the newly-placed AP.

For testing during planning, the testing is still simple. Use traditional network prototyping techniques, such as placing APs on temporary stands and varying their location, and use WaveDeploy to make the measurements. Vary the number of clients, or use ecosystem traffic where very high client counts are required.

Experiments can quickly be done to check:

- The impact on traffic delivery of placing another AP on the same channel nearby
- The impact legacy clients at 1 and 2 Mbps rates will have on overall system performance
- How many APs will be required to support a meeting room environment when multiple wireless clients are all depending on the network for voice and/or video
- The impact placing an AP in a particular location has on performance (important when evaluating performance vs. esthetic tradeoffs)

All of the above (and more) are simply different instances of WaveDeploy testing, where the experiment requirement can range from a single client, to multiple simultaneous clients, to simultaneous clients plus ecosystem traffic provided by a WF1101.

Benefits

By performing the testing described in this section, the following benefits are derived:

- Rapid and data-driven decisions while working out planning issues
- Verification of all network equipment while the installer equipment/personnel is onsite
- Objective data to provide to network vendors when faulty equipment is identified

Overview

For testing done in support of network planning, the objective is to translate the situation into an experiment. The network vendor will have proposals for addressing the situation; typically the difficult part of the equation is coming up with client devices in sufficient number, variety, and measurement capability to quantify the effectiveness of the proposal.

WaveDeploy provides the mechanism to manage the traffic generation and analysis on the client devices, plus harvesting the performance data, and a low-touch means to add many tens of diverse clients and traffic flows. This makes conducting virtually any "what-if" test within minutes very simple.

For testing done in support of deployment, the objective is to run the same traffic loads and SLAs that will later be used for site assessment, but to do so at a few key locations while the network is being deployed. For advanced situations where advanced diagnostics and continuous traffic are required, the EF1101 with WaveInsite is used.

Test Strategy

The device under test (DUT) is the network being deployed. The test system is WaveDeploy Expert. If specific loads need to be applied to the DUT for long periods of time for troubleshooting, WaveInsite is used. All testing is done on a deployed network, in open air, as shown below.

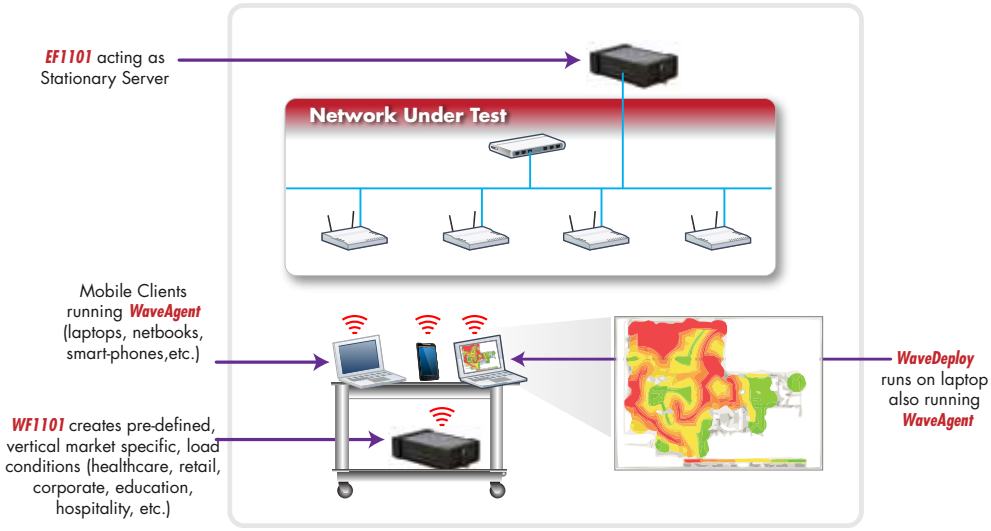

1. WaveDeploy Testing (Planning Phase)

Representative Clients, Coexistence and Speed Test with Ecosystem Traffic (if required)

- Use "Best Practices" guidelines or results of Chapter 4 for target load settings and SLAs
- Add clients and/or ecosystem to match the needs of the experiment
- Select all traffic types for coexistence and speed test traffic that are of interest
- All client devices configured per IT guidelines (security, power save, etc.)
- Use WaveDeploy results analysis to derive answers
- Use WaveInsite to investigate specific network/client performance issues
- What this test reveals:
  - » Key configuration and placement questions answered with repeatable data

2. WaveDeploy Testing (Deployment Phase)

Representative Clients, Coexistence, and Speed Test, with Ecosystem Traffic (if required)

- Use "Best Practices" guidelines or results of Chapter 4 for target load settings and SLAs
- Verify under each access point as it is turned on that measurements are "all green" at a distance of about 20 feet from the AP
- Repeat measurements if there appears to be any variation or SLA failure
- If not "all green" examine WaveDeploy results to determine cause
- What this test reveals:
  - » Faulty access points
  - » Sources of interference
  - » Deployment issues not anticipated in the planning phase

Scale

For the planning phase, the scale of the test should match the complexity of the experiment needing to be performed. For situations where large client count is required, WaveDeploy Expert with WF1101 should be used in addition to the actual client devices. For the deployment phase, it is typical to conduct testing using a single client device has been demonstrated to be capable of high performance. The preferred approach is summarized in the following table.

Small Office –Single, stand-alone AP, no independent controller	• WaveDeploy Pro. Use a laptop capable of supporting the highest data rates required by the enterprise
Small/Medium Office – Small scale controller with up to 4 APs, or distributed controller architecture	• WaveDeploy Expert. Use a laptop capable of supporting the highest data rates required by the enterprise
Office Building – One or more controllers, multiple APs	• WaveDeploy Expert. Use a laptop capable of supporting the highest data rates required by the enterprise. Use EF1101 and WaveInsite for advanced troubleshooting. Use multiple WF1101s if more than 64 ecosystem clients.
Campus – Multiple controllers, multiple APs	• WaveDeploy Expert. Use a laptop capable of supporting the highest data rates required by the enterprise. Use EF1101 and WaveInsite for advanced troubleshooting. Use multiple WF1101s if more than 64 ecosystem clients.

Return on Investment

Network managers who have invested in this approach for planning and deploying the network enjoy the following returns:

- Network planning decisions proven to meet requirements, minimizing the need for "additional touches" and continued tuning in challenging spots
- Find faulty network equipment before the site assessment phase
- Ability to present repeatable and quantifiable evidence to both network and client vendors when communicating issues for action

Chapter 5: Adjust and Verify Deployed Network

Performing a site assessment using a tool capable of running and measuring multiple streams of traffic to/from multiple client devices is the one and only means of characterizing the experience users will have on the network at any location in the enterprise. Performance cannot be estimated, calculated, or simulated. There are no "safe" assumptions –companies must either put their client devices and network under traffic stress and prove performance, or they will answer trouble tickets and chase assumptions later.

The problem with answering trouble tickets without employing a solid testing methodology is two-fold. The obvious problem is the amount of time it takes to react to each occurrence – with the operative word here being "react." The not-so-obvious problem is that chasing problems where there is a lack of foundational data means that picking apart the problem is all done empirically and without a point of reference. Changes are made to network configurations that can actually have unintended negative consequences, but don't show up until another time, when a new set of stresses are put on the network.

The testing required in this chapter is actually quite simple and straightforward. It requires a bit of pre-work in the form of setting up an Ethernet "fixed client" to anchor the measurements, and configuration of the proper target loads and SLAs within Ixia's WaveDeploy prior to performing the site assessment. This process is described in more detail in Chapter 3. Then, it is a matter of walking the portion of the enterprise (all or part), making the measurements, and reviewing the results. The time investment is similar to that of RF mapping traditionally.

It is important to note that it is not a requirement that an entire physical enterprise be measured in one assessment. Often users will start by doing an assessment with traffic using WaveDeploy in the section of a building that is problematic, resolving client and/or network issues there first. This has the effect of allowing the user to build a foundation of data a little at a time, and then merge assessments together at a later date.

Benefits

By performing the testing described in this section, the following benefits are derived:

- Performance maps for traffic delivery by type, RF signal power, phyrate, co-channel interference
- Performance data on relevant client devices and drivers
- Foundation of data for solving problems in the future
- Demonstration that network can carry required level of traffic
- Simple test configuration to reuse when network configurations are changed

Overview

Perform a site assessment using WaveDeploy, and include relevant client devices in the assessment running target loads of traffic that correspond to the amount of data the enterprise requires. With a bit of attention to setting target loads and SLAs, the resulting map clearly shows the locations where end-user performance will not meet expectations and further work is required. For environments requiring very high client counts, use WaveDeploy Expert with one or more WF1101s to provide ecosystem traffic.

The resulting performance maps establish not only signal levels and PHY rates that are experienced by actual client devices, but demonstrate traffic performance in both directions for voice, video, and data for relevant clients individually and together. For situations where advanced diagnostics and continuous traffic are required, WaveInsite is used.

Test Strategy

The device under test (DUT) is the network being deployed. The test system is WaveDeploy Expert. If specific loads need to be applied to the DUT for long periods of time for troubleshooting, WaveInsite is used. All testing is done on a deployed network, in open air, as shown below.

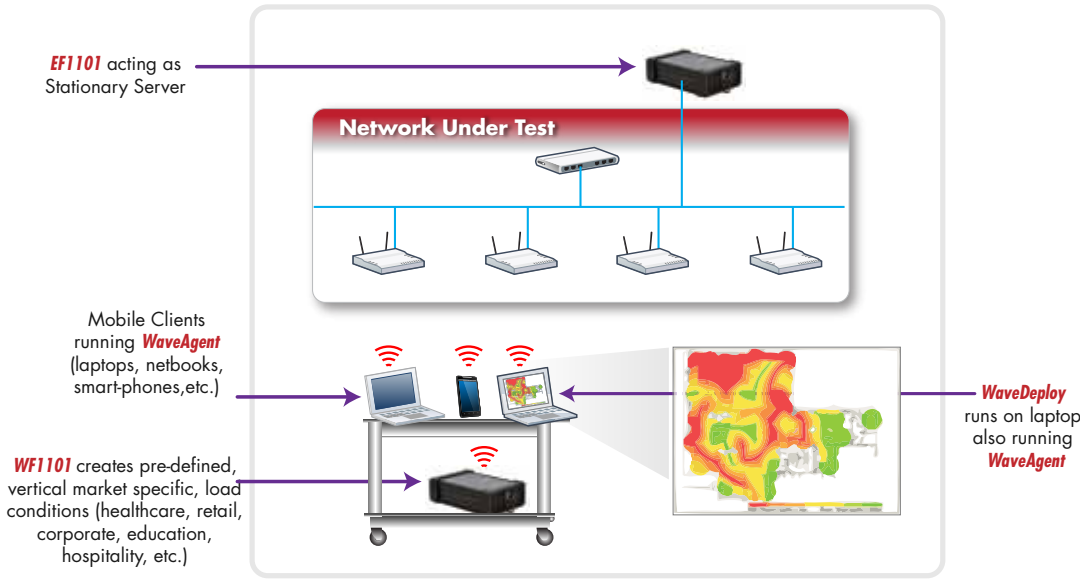

The testing approach has two basic components. The first is to perform baseline testing to ensure that the network configuration is correct, and that the target loads and SLAs for each client device included are set correctly. This is done at a favorable location: short distance from the access point, no source of interference, known good network devices and clients. Following the instructions in Chapter 4 provides the configuration information necessary for setting target loads and SLAs. The second component, once the baseline is established, is to perform a series of "walk around" assessments.

It is best practice to include actual client devices that are representative of those the end-users have, and include the mixture of clients the enterprise supports. For example, if an enterprise supports 802.11n laptops, 802.11g laptops, and smart phones, it is best practice to include at least one of each in the baseline testing and "walk around" assessments, with each configured for the target loads and SLAs appropriate to their capability.

1. WaveDeploy Testing (Baseline)

Representative Clients, Coexistence, and Speed Test With Ecosystem Traffic (if required)

Use "Best Practices" guidelines or results of Chapter 4 for target load settings and SLAs

- Measure at favorable location and ensure SLAs are met in all performance categories
- Select all traffic types for coexistence and speed test traffic that are of interest
- All client devices configured per IT guidelines (security, power save, etc)
- Use WaveDeploy results analysis to derive answers
- Use WaveInsite to investigate specific network/client performance issues
- What this test reveals:
  - » Configuration of network is correct to support traffic mix
  - » Configuration of client devices is correct to support end user experience
  - » Baseline level of performance

2. WaveDeploy Testing (Walk-Around)

Representative Clients, Coexistence, and Speed Test with Ecosystem Traffic (if required)

- Use target load and SLAs from baselining
- Load site map into WaveDeploy
- Perform "walk around" site assessment
- If not "all green" examine WaveDeploy results to determine cause
- What this test reveals:
  - » Locations where network/client fails to deliver required performance level
  - » Roaming issues with network and/or client –stickiness, failed roams
  - » Problems with automatic radio management algorithm/configurations
  - » Presence of other networks or clients that will be problematic

Scale

Small offices typically conduct this testing with a single laptop, or a laptop plus key client devices. Medium offices typically use a combination of laptops, netbooks and smart phones. Larger enterprises will add WaveDeploy Expert and increase client count using the WF1101 for ecosystem traffic in certain locations. The preferred approach is summarized in the table below:

Small Office –Single, stand-alone AP, no independent controller	• WaveDeploy Pro. Use representative mix of client devices
Small/Medium Office – Small scale controller with up to 4 APs, or distributed controller architecture	• WaveDeploy Pro. Use representative mix of client devices
Office Building – One or more controllers, multiple APs	• WaveDeploy Expert. Use representative mix of client devices. Use EF1101 and WaveInsite for advanced troubleshooting. Use multiple WF1101s if more than 64 ecosystem clients
Campus – Multiple controllers, multiple APs	• WaveDeploy Expert. Use representative mix of client devices. Use EF1101 and WaveInsite for advanced troubleshooting. Use multiple WF1101s if more than 64 ecosystem clients

Return on Investment

Network managers who have invested in this approach to verifying the deployed network enjoy the following returns:

- Ability to base "sign-off" of deployment on repeatable, quantitative measurements done on relevant client devices
- Foundation of data, by location, that provides the baseline measurements to compare to when responding to trouble tickets
- Ability to present repeatable and quantifiable evidence to both network and client vendors when communicating issues for action
- Ability to more quickly separate client issues from network issues in the future
- Solid baseline data against which to compare when network configuration changes or firmware upgrades are performed

Chapter 6: Respond to and Remediate User Issues

Many organizations spend a majority of their time and effort dealing with problems. The goal of this chapter is to lay the groundwork for enabling an enterprise to become less driven by reaction to trouble tickets, and more equipped to get ahead of the user population and prevent issues.

The basis of best practices for remediating customer issues is always rooted in repeatable measurements and prior verification. If the person resolving the issue is armed with data upon which they can rely, they can quickly pinpoint which element is not operating correctly.

Looking at a specific example makes the point clear. Say a trouble ticket comes in that states, "Wireless network appears very slow at Jane Taylor's desk." Best practices would dictate that the network management team already knows what maker of laptop Jane has, what NIC and driver is installed, and what version of OS is running on that laptop.

It is also known where Jane's desk is located and which access points serve that area. Best practices would also dictate that a WaveDeploy site assessment was done of that area, and the data can be pulled up to show performance maps at that location. Finally, a performance test of Jane's laptop would have been completed before deployment so the team knows what the network and the client device should be capable of.

It now becomes a simple matter of seeing what has changed. If it's the network, a short retest with WaveDeploy at that site will reveal this in minutes. If it is the wireless NIC in the laptop, a short WaveDeploy test will reveal the problem equally fast. If the network and wireless NIC demonstrate that the delivery of traffic to Jane's laptop isn't the issue, the focus can quickly be shifted to the laptop itself.

Benefits

By performing the testing described in this section, the following benefits are derived:

- A method of rapidly isolating network from client issues
- The ability to determine whether the problem is really the wireless or not
- The ability to conduct a performance test remotely
- The means to confirm the corrective actions

Overview

Remediation can be reduced to a rapid three-step process:

1. Test the location in question against previous WaveDeploy results, using a known good client device. If those results are bad, something has changed with the network or wireless environment.

2. Retest individual APs with WaveDeploy to determine whether new interferes are present, AP radios are functioning, or if a configuration change has resulted in either too much power or too little power being emitted.
3. If the test with known good client devices shows good results, the problem can be traced to either new or changed client devices (wireless NIC, driver, other interfering processes). Test the new or changed client devices with WaveInsite for interoperability and performance.

Test Strategy

The testing approach is designed to isolate the network from the client device as quickly as possible. The first step is to check the network and the environment and compare to expected results. The most effective approach is to take a client device of known performance to the location in question, and run WaveDeploy from a previous configuration so as to simply repeat a measurement that was taken during site assessment.

If the measurement can't be repeated, WaveDeploy analysis maps point to the problem. If the network and wireless environment check out, run the same test on the client device and see if the results can be duplicated. Satisfactory results mean that communication from the IP stack inside the client device, all the way through the wireless NIC and access point, are working as expected.

If the test with known good client devices points to the network, then the system under test (SUT) is the network and the test system WaveDeploy Pro. All testing is done on a deployed network in open air as shown below.

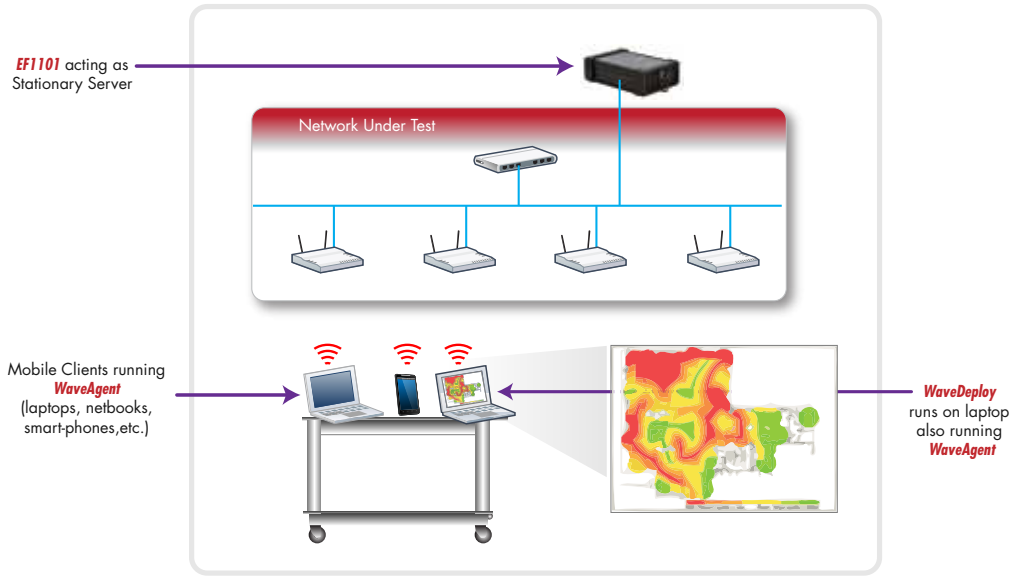

If the test with known good client devices points to the clients, then the system under test (SUT) is the client(s). Use the WaveTest Ethernet WaveBlade or the portable EF1101 Ethernet Test Appliance using the WaveInsite application. A representative test setup is shown below.

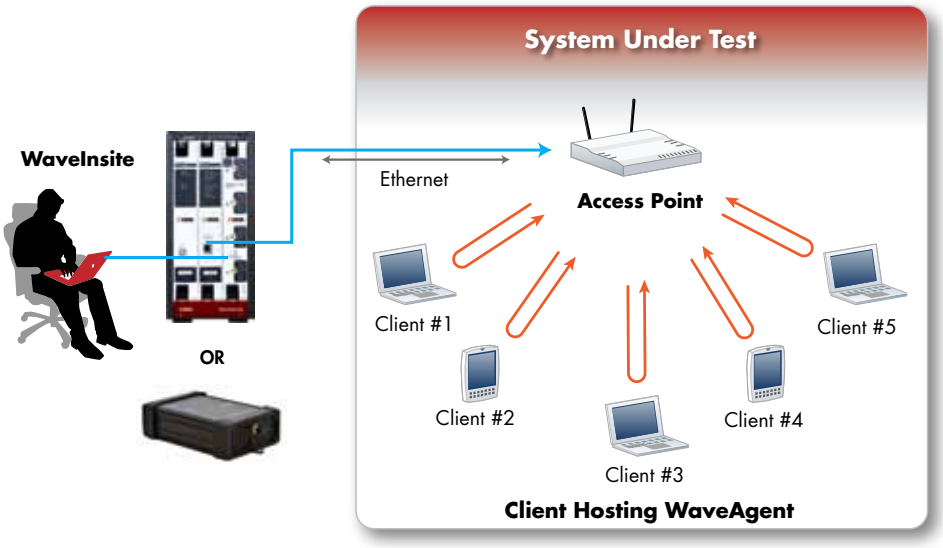

1. WaveDeploy Testing of the Local Network and Environment

Using Known Good Client

- Run stressful load, include coexistence and speed test types
- Use same target loads and SLAs from site assessment, if available
- Measure at the location where the problem is reported
- Compare to previous results and check for differences
- Examine details of results that are below SLAs
- Correct issues and repeat
- What this test reveals:
  - » Malfunctioning access points (low phy rate, low traffic performance, poor RSSI)
  - » Environmental issues (new sources of co-channel interference, interfering clients
  - » Changes in the channel utilization mapping
  - » Ability to quickly confirm whether changes made are helping

2. WaveInsite Testing of Client Device (Local or Remote)

If Testing Above Shows Network Performance is Good

- Load WaveAgent onto client device
- Test maximum upstream, downstream UDP, and TCP forwarding rate performance to clients and compare with results obtained with known good clients
- Measure close to AP, then in distance increments of 10 feet, up to maximum AP spacing
- What this test reveals:
  - » Client/AP interaction under stressful load as compared to known good client
  - » Expected performance from client device as compared to known good client
  - » Performance degradation with distance, each direction as compared to known good client
  - » Issues with the NIC driver or client device
  - » Network interoperability issues specific to a given client device
  - » Ability to quickly confirm if changes made are helping

For environments in which the client is remote to the test facility (e.g. satellite office, retail store, etc.) this is a fast step for measuring the network, environment, and performance as seen by the suspect client device.

Scale

Small offices typically conduct this testing with a single known good laptop. Medium offices will use the same approach, but can combine other client devices as necessary. Larger enterprises will add WaveDeploy Expert and increase client count using the WF1101 for ecosystem traffic in certain locations. The preferred approach is summarized in the following table.

Small Office –Single, stand-alone AP, no independent controller	• WaveInsite & WaveDeploy Pro
Small/Medium Office – Small scale controller with up to 4 APs, or distributed controller architecture	• WaveInsite & WaveDeploy Expert
Office Building – One or more controllers, multiple APs	• WaveInsite & WaveDeploy • Use multiple WF1101s if more than 64 ecosystem clients
Campus – Multiple controllers, multiple APs	• WaveInsite & WaveDeploy Expert • Use multiple WF1101s if more than 64 ecosystem clients

Return on Investment

Network managers who have invested in this approach to remediation enjoy the following returns:

- Faster resolution times: the ability to put focus on where the problem truly lies
- Lower cost: the ability to address some situations without deploying personnel to the location
- Ability to present repeatable and quantifiable evidence to both network and client vendors when communicating issues for action
- Ability to more quickly separate client issues from network issues

Chapter 7: Verify Software Updates and Configuration Changes

When new software is supplied for the network or client devices, the best approach is to repeat previous testing to ensure the performance and functionality has not degraded, and validate expected improvements in performance. The same is true of network configuration changes since wireless network configuration is extremely complex, and it is quite possible to make changes that bring about unintended and unobvious consequences.

The preferred approach when verifying changes is to have conducted the baseline testing described in Chapters 1 and Chapter 2 so that configurations of the test applications and the resulting data against which to compare are established. For the network, the best practice is to put the controller and access points under truly stressful conditions, with high client counts and significant traffic loads. Running benchmark performance tests of throughput, latency, maximum client capacity, rate vs. range, roaming tests, and Quality of Experience tests at scale yields the only "apples to apples" comparison from which to judge the upgrade or change.

For clients, the best practice is to subject the devices to a stressful environment to ensure performance. Forwarding rates, delivery of mixed traffic, coexistence with other client devices, and roaming behavior are all important elements to confirm. The key point to keep in mind is the high cost of deploying a software upgrade, configuration change, or NIC driver change that has an issue, because it will be found by and affect the efficiency of the population of users.

Benefits

By performing the testing described in this section, the following benefits are derived:

- Defects are uncovered and identified prior to deployment
- Cllaims regarding improvements and compatibility from network and client vendors can be verified
- A basis of performance data against which configuration tradeoffs can be measured can be easily maintained

Overview

Since software upgrades may be applied to the network, the clients, or both, it is best to change only one variable at a time, testing either network changes while clients are held constant, or vice versa. Begin by testing the location in question against previous WaveDeploy results using a known good client device. If those results are bad, something has changed within the network or wireless environment:

1. For network software changes, the test must be performed with equipment that is proven to be capable of exceeding the capability of the network under test, otherwise the process is using the network to measure the stressing equipment and not the other way around. The equipment must be capable of running at IEEE 802.11 and 802.3 line rate, and present each AP involved with the required client count and behavior. For this reason, the only known method of performing this testing in an effective, economical, and repeatable fashion is with test equipment designed for the task.

2. For client software changes, perform benchmark and deployment testing of client devices to check driver/OS upgrades and configuration changes. This is done using WaveInsite and WaveDeploy.

Test Strategy

When confirming network side changes, the SUT is the wireless network controller and access points. The test system is the WaveTest Ethernet and Wireless WaveBlades, using the Benchmarking and WaveQoE applications. Connect the APs to WaveBlades using RF cables, and use the test applications to perform testing. A representative test setup is shown below.

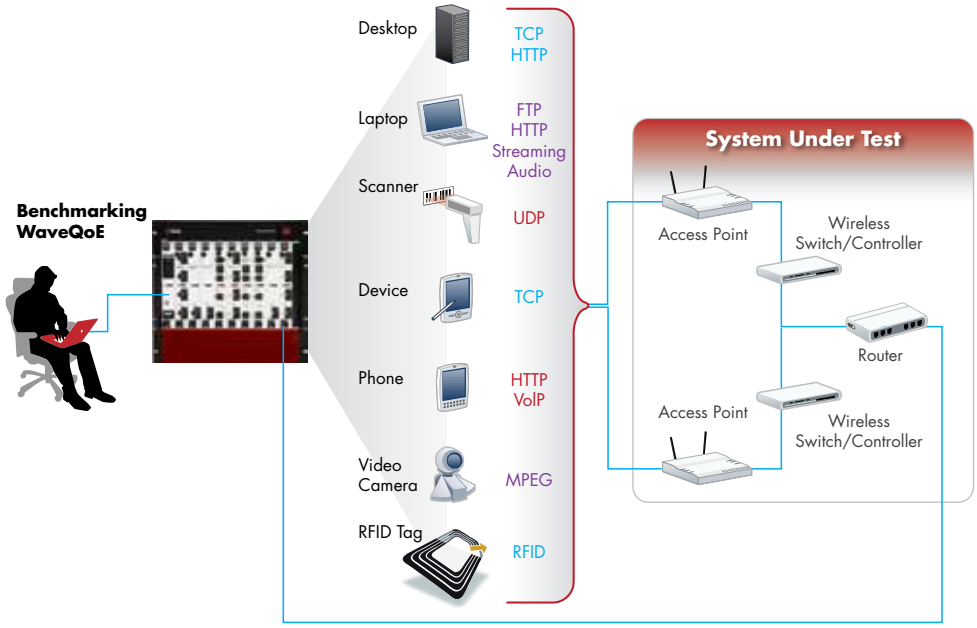

Refer to chapter 1 for detailed instructions.

When confirming client side changes the device under test (DUT) is the client product, such as a smart phone, netbook, or laptop. The test system is WaveInsite and WaveDeploy. Other client devices with which the DUT needs to be compatible, are included to verify co-existence. All testing is done in open air, as shown below.

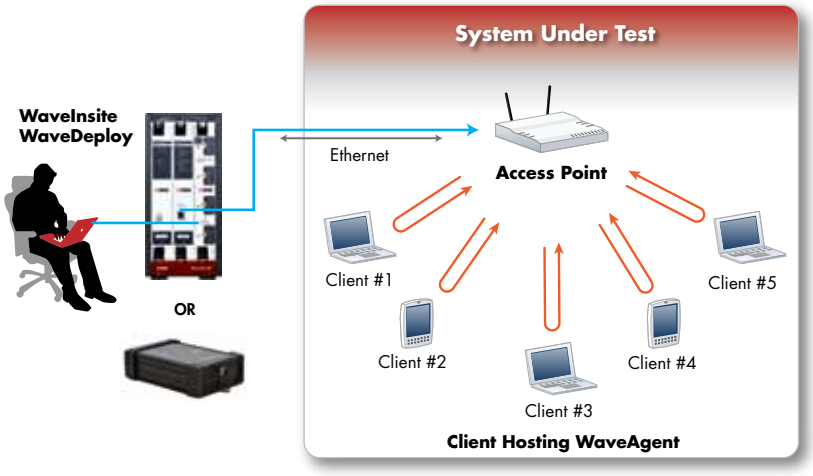

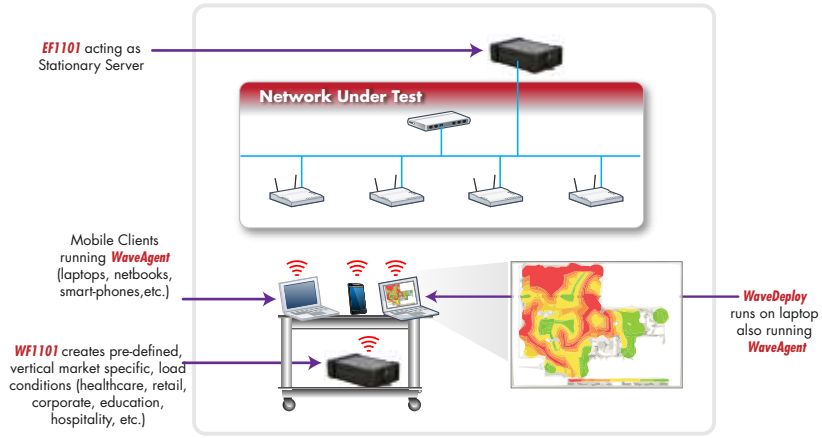

Return on Investment

Network managers who have invested in this approach for verifying software updates and configuration changes enjoy the following returns:

- Eliminate surprises after upgrades
- Ability to drive concerns back to network and client vendors with quantifiable data prior to “going live”
- Ability to preview the effects of configuration changes and new features turned on in a network
- Ability to present repeatable and quantifiable evidence to both network and client vendors when communicating issues for action

Chapter 8: Verify New Client Devices

Client devices change rapidly with new devices or versions being introduced every few months, and new software builds constantly being downloaded. As an IT manager, you typically have less control over client devices, yet they represent the most critical link to end-user quality of experience.

The fact that each client device is manufactured by a different vendor and all need to coexist and interoperate with the chosen infrastructure makes them the weak link in the overall network. A proactive approach to managing the population of wireless devices in the enterprise means performing a verification of the device as it is deployed, and eliminating the effects of changes before the unit gets to the end user. To be efficient, this test has to be quick and easy to conduct within minutes.

Benefits

By performing the testing described in this section, the following benefits are derived:

- Uncover and identify defects and before deployment
- Catalog the performance of each unit by serial number to enable detection of degradation over time
- Maintain a basis of performance data against which configuration tradeoffs can be measured

Overview

The process for conducting this test is to install WaveAgent on the candidate device, and run a WaveDeploy test that includes that client device to check for performance that matches test results for that model. This test is typically conducted at a known location that has been previously characterized for performance and lack of interference.

Test Strategy

The device under test (DUT) is the client product, such as a smart phone, netbook, or laptop. The test system is WaveInsite and WaveDeploy. Other client devices with which the DUT needs to be compatible, are included to verify co-existence. All testing is done in open air, as shown below.

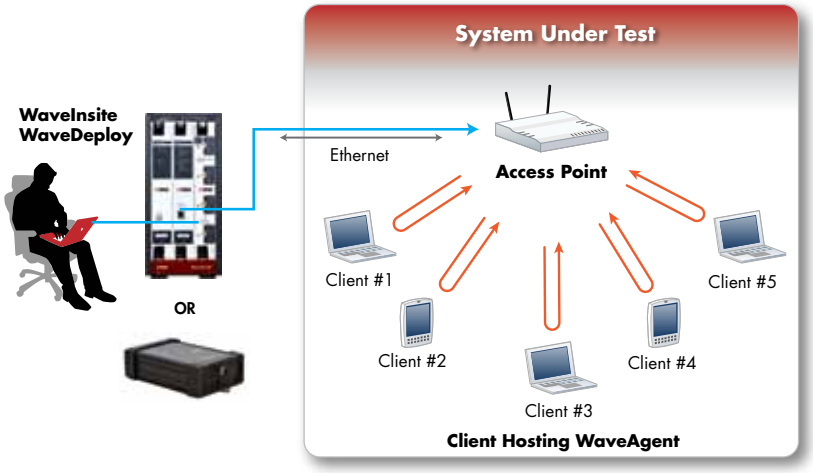

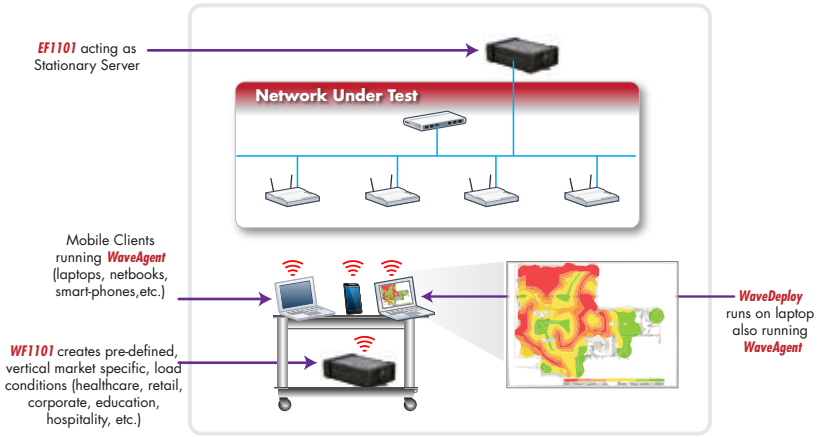

1. WaveInsite Testing (Controlled-RF Environment) –1 AP, 1 Client Device

Maximum UPD Forwarding Rate

- Upstream, downstream, bidirectional
- Multiple frame lengths (64, 512, 1024, 1518)
- Open and Secure
- Power Save
- What this test reveals:
  - » Maximum packets/second the DUT can handle (short frame performance)
  - » Maximum forwarding rate in each direction
  - » Penalty for security encryption
  - » 802.11n aggregation efficiency

TCP Goodput

- Default MSS settings (536, 1460)
- 1, 2, and 4 streams per flow type
- Upstream and downstream
- What this test reveals:
  - » Ability to sink downstream TCP flows, per session
  - » Ability to source TCP flows

Voice / Video / Data Mixed Flow

- HTTP downstream, Voice bidirectional, UDP upstream, and downstream
- Create flows that reflect applications
- Upstream and downstream
- What this test reveals:
  - » Maximum performance for client NIC to carry mix of data with appropriate QoS

Rate vs. Range

- Run at ten-foot distance increments from AP
- Perform rate vs. range test while both client and AP are under stress
- Traverse range until traffic stops
- What this test reveals:
  - » QoS under stress
  - » Quality of NIC receiver and transmitter
  - » Effectiveness of rate adaptation algorithm in client device

2. WaveInsite Co-Existence Testing –1 AP, 2 or More Client Devices

Application Performance with Other Clients Present

- Enable one or more additional client devices on the same AP and BSSID as the DUT. For special-purpose devices, run actual customer application (e.g. patient monitor)
- Repeat tests in 1) above and assess effect on DUT and other devices
- What this test reveals:
  - » Client device's response to other coexisting traffic on same band/AP
  - » Network configuration requirements specific to the client device

3. WaveDeploy Testing - Optimum Location

Single-Client Coexistence and Speed Test

- Use "WLAN Site Assessment: Best Practices for Pre and Post Deployment Verification" guidelines for target load settings and SLAs
- Select all traffic types for coexistence and speed test traffic that are of interest
- Client device configured per IT guidelines (security, power save, etc)
- Make multiple measurements from same location to demonstrate repeatability
- What this test reveals:
  - » Verifies client device/network combination can deliver required performance under stress
  - » Whether target load or SLA settings need to be adjusted
  - » Hidden issues with optimum location
  - » Whether unexpected variation between repeated measurements is cause for investigation
  - » Whether, if performance cannot be achieved, adding other clients and competing traffic will make things worse

Multiple-Client Coexistence and Speed Test (DUT Plus Other Relevant Client Devices)

- Use "Best Practices" guidelines for other client devices
- Select all traffic types for coexistence and speed test traffic that are of interest for ALL client devices
- All client devices configured per IT guidelines
- Make multiple measurements from the same location to demonstrate repeatability
- What this test reveals:
  - » Client-to-client interaction – if speed test results are inconsistent with coexistence results
  - » If target load or SLA settings need to be adjusted for multiple-client case
  - » Unexpected variation between repeated measurements is cause for investigation

Multiple Client Coexistence and Speed Test Plus Ecosystem Traffic (if required)

- To increase client count and amount of traffic in the collision domain, use WF1101
- Use WaveDeploy Expert to create ecosystem of up to 64 clients per WF1101 with constant background traffic (multiple WF1101s can be used for very high client count)
- Repeat Multi-client Coexistence and Speed Tests from above
- Make multiple measurements from the same location to demonstrate repeatability

- What this test reveals:
  - » Confirms interoperability under highest client count condition with stressful traffic load
  - » Whether target load or SLA settings need to be adjusted for multiple-client/ecosystem case
  - » Unexpected variation between repeated measurements is cause for investigation

4. WaveDeploy Testing - Challenging Locations

Radio Power-Up at Challenging Location

- Use multi-client coexistence and speed test, plus ecosystem traffic test configuration from above
- Adjust target loads and SLAs downward if locations warrant a decrease in expected performance
- Turn on client devices and ecosystem traffic clients at the challenging location
- Make multiple measurements from the same location to demonstrate repeatability
- If performance is not as expected, repeat with ecosystem traffic removed, then other client devices removed to diagnose
- What this test reveals:
  - » Interoperability demonstrated under worst-case condition when DUT and other clients are first connected at that location
  - » Rate adaptation, radio performance, airtime fairness, AP selection algorithm performance at the edge of the network
  - » Appropriate target loads and SLA settings at the edges of the network

Radio Power-Up, Then Migrate to Challenging Location

- Test loads and SLA will be identical to prior test
- Turn on ecosystem traffic at the challenging location
- Turn on all client devices at another location, so that all connect to an access point other than the preferred one at the target location
- Make one measurement, then move all clients to the challenging location
- Make multiple measurements at the challenging location
- If performance is not as expected, repeat with ecosystem traffic removed, then other client devices removed to diagnose
- What this test reveals:
  - » If clients stick to a non-optimum access point as they migrate to the challenging location
  - » How quickly the client device(s) change their connection to the correct AP, and what degree of disruption occurs during the process

- Effectiveness of “roaming aggressiveness” settings
- Characteristics of the poor link (phy rate, retries, rate adaptation churning) if a client does not resolve connection within a reasonable time

Chapter 9: Plan for Network Upgrades and New Services

In the course of operating an enterprise network, the process of continuously expanding the network and the services offered is a normal occurrence. The convenience of wireless coupled with the proliferation of client devices and applications make a powerful incubator for innovation, leading to a constant stream of requests to network managers.

The response can range from simply turning on features of the existing network equipment, to upgrading the technology and capability of the network altogether. Planning for upgrades and new services is about projecting the organization and its appetite for capacity into the future, and being equipped to assess the ability of the existing network to meet the need.

If capacity is exceeded, the planning activity then lays the groundwork for establishing performance criteria when considering new network equipment. The typical questions considered are:

- What equipment is required to see if the existing network can support the expanding need? How can we verify expanded operation quickly, efficiently and economically?
- Will the existing network handle the expanded functionality and scale required? Is it a matter of adding access points, or is the job bigger than that?
- If more capability/capacity is required, how do we solidify the requirements so that the decision-making process for the next investment can be data-driven?

Benefits

By performing the testing described in this section, the following benefits are derived:

- Effective means of performing "what if" testing on existing networks, at scale
- Performance data expressed in end-user experience metrics
- Decisions based on performance metrics rather than estimates and projections, reducing risk to the organization

Overview

This testing typically takes one of two forms: smaller scale testing aimed primarily at assessing the network's ability to provide new or denser services, and larger scale testing aimed at planning for network upgrades. An example of providing new services would be an organization investigating the possibility of bringing VoIP phones into the workplace, and wishing to see how the existing network can support the current data users while a new population of prioritized traffic and client devices is introduced.

Important attributes of QoS, roaming performance, and how many VoIP clients can be handled per AP become key questions. Unfortunately, what has become standard practice in addressing such questions is not best practice. An all too common response to investigating is to bring a small sample of the target devices into the environment, and do an empirical check with uncontrolled methodology. Basically, if one laptop appears to be able to download some video from the internet while a VoIP phone holds up a decent call nearby, that becomes the "test."

There are three serious shortcomings to such approaches:

1. Empirical demonstrations do not provide a number of client devices nor traffic levels that represent what would actually be on the network after the new devices are deployed.
2. These approaches require multiple people and devices to be coordinated as a special event, so the logistics and time commitment make the process inefficient, expensive and hard to manage.
3. These approaches are difficult or impractical to repeat. They often involve getting multiple devices manually connected to the network and passing traffic. Subjective quality of experience measurements and the discovery of variables are not anticipated in test planning.

Best practice for evaluating the network's ability to support new services cannot involve assumptions and compromises, but also does not have to be a logistical nightmare. Those addressing these challenges use IxVeriWave equipment and applications to scale the client numbers and traffic loads from within a single application, using purpose-designed hardware to provide relevant client counts and derive repeatable and precise results.

An example of planning for a network upgrade would be the case where an organization is investigating the risks and rewards of moving from an 802.11 legacy network to an 802.11n network, including introducing n-capable client devices into the environment. In this case, the task is one of assessing the existing use of the network so that a current view of network traffic and client population is gained, then specifying what changes to that basis are expected in the future.

Typically there are three cardinal requirements of a new network:

1. It must support the existing use needs of current users immediately; the act of performing a network upgrade must not negatively impact users.

2. Existing problems must be solved immediately. If performance, compatibility, or interoperability issues are limiting the effectiveness of the existing network, the act of installing the new network must solve those problems "out of the box."

3. It must provide sufficient margin for growth as higher-rate client devices are introduced over time. If expanded functionality is a part of the organization's plan, the network must not become a limiting factor.

If the organization believes that adding software features and higher performance access points can meet the expanding need, the approach should be to obtain evaluation versions of the software and hardware, and perform a test at scale. That method is described below.

On the other hand, if the organization believes that the requirements will lead to upgrading the controller, parts of the infrastructure and the access points, the reader is referred to Chapter 1. The task has become one of selecting new infrastructure equipment.

Test Strategy

This section assumes that the new features and functions required of the network have not reached the magnitude where a completely new controller and infrastructure are called for. The system under test (SUT) is the existing controller and switches, with perhaps new software features enabled and higher performance access points.

An excellent example of such a situation is when an enterprise wants to move a subset of the network from legacy to 802.11n APs, and the controller is capable through software upgrade of supporting this move. The best practice for test strategy becomes a straightforward exercise.

Simply put, the team identifies the number of clients, the mix of traffic, and any other pertinent conditions (security, for example) that are called for, and uses either WaveTest or WaveDeploy to test at that scale. Anything less is leaving room for surprises. Deciding whether to use WaveTest or WaveDeploy or both becomes a question of whether the new services need to be proven out on the deployed network over-the-air.

Two cases will be described here, one smaller scale and one larger scale, to illustrate two ends of the spectrum. Network managers can adapt the scale and approach to suit the situation, but the key point being made here is that if risk is introduced anytime the testing approach begins scrimping on numbers – either client count or amount of traffic.

Smaller Scale Example

The first case to be described is a common example: a wireless network already exists at the enterprise, but an important meeting involving a number of people will be taking place on site. All attendees will be bringing their laptops and smart phones, and conducting regular business during the meeting.

The organization wants to be sure of two things:

- That the experience of those in the meeting will be satisfactory (sufficient capacity and performance in the room)
- That the experience of those on the remaining part of the network is not degraded

The most obvious approach would be to go buy or rent the number of laptops and smartphones expected to be in attendance, connect them all to the network, generate traffic, and measure performance to all these devices while also monitoring the performance of the rest of the network. Let's say this event is a shareholder's meeting, and roughly 50 analysts will be in attendance. Let's also say that presently there are two access points that are installed in the meeting room. What is outlined below is an approach that makes this test straightforward, repeatable, and executable by one person.

1. Confirm Client Capacity

WaveDeployExpert (2 WF1101s, 64 wireless client capability each)

- Configure each WF1101 through WaveDeploy for 25 laptop and 25 smartphone clients
- Set target loads to the smallest load level, adjust SLA to match
- Connect both WF1101s to the network and confirm that all clients get connected and pass traffic
- Make "four corner" measurements on representative laptop as mobile client, with minimal traffic
- What this test reveals:
  - » System ability to connect 100 clients
  - » Time it takes to get clients connected
  - » System ability to serve traffic to 100 clients

2. Confirm Load Capacity

WaveDeployExpert (2 WF1101s, 64 wireless client capability each)

- Repeat above test, but set target load to average load to/from each laptop/smartphone
- Set SLAs to the lowest acceptable level
- Connect both WF1101s to the network and confirm all clients get connected and pass traffic
- Make "four corner" measurements on representative laptop as mobile client, with minimal traffic
- Inspect the ecosystem traffic results
- What this test reveals:
  - » System ability to support expected client count and traffic loads
  - » If additional access points will be required because of media saturation
  - » System stability at scale

3. Confirm Total System Capacity

WaveDeployExpert (2 WF1101s, 64 wireless client capability each)

- First, go to other locations in the building and use WaveDeploy to establish best performance measurements at those locations. Establish the level of service available to the rest of the network when the 100 clients and traffic in the target room are not connected.
- Using the configuration from the Load Capacity test above, connect both WF1101s to the network and confirm all clients get connected and pass traffic
- Repeat the WaveDeploy measurements at the other locations in the building.
- Inspect all results
- What this test reveals:
  - » Impact new traffic and clients have on performance at other locations
  - » Whether channel management plans will need to be adjusted to isolate the rest of the environment from the meeting room

The approach above can be modified to suit the scale of the test. Since each WF1101 can generate up to 64 clients and associated traffic on a single channel, and do so at 802.11n line rate, testing a "what if" scenario of 64 clients or less can be accomplished with a single WF1101.

Increasing the number of clients, and/or loading multiple RF channels becomes a simple matter of using additional WF1101s. It can be seen through the above example that the strategy can be scaled, but follows the general model of:

- Verifying that the required number of clients can be connected in a timely manner
- Verifying that the required amount of traffic can be passed, concurrent with the required client count
- Measuring the impact on other parts of the network with a "before and after" approach

Using WaveDeploy, this approach is executable by a single user in a couple of hours, provides documented evidence of the results, and is easily repeatable.

Large Scale Example

The second example will be a large-scale case, also a common situation, where the organization wants to verify that upgrading firmware on an existing controller and deploying a new generation of APs provides increased capacity and performance – without having to completely replace the controller and infrastructure. This would be the case where an influx of 802.11n laptops and smartphones with bandwidth-hungry applications is expected, and the goal is to take an evolutionary rather than revolutionary step in network management.

The organization in this situation wants to observe three things:

- Can the proposed firmware and access point changes satisfy the performance demands within the existing infrastructure?
- Will making these changes disrupt existing users?
- What performance gains are achieved: how far into the future does this take the organization?

The best-practice approach is to perform scaled benchmark tests with WaveTest equipment cabled directly to an access point, followed by an open-air verification. This case differs from the smaller-scale one above, because in the large-scale case the performance improvement needs to be proven at the "whole network" level instead of at a single location.

For this reason alone, the number of client that need to be applied to the network should be in the hundreds, and traffic flows should be elevated to the point where maximums are observed. Obviously this is far outside the range of what can be achieved with a few laptops empirically. What is outlined below addresses this case:

1. WaveQoE Test: 2 Radios per AP (as applicable), 2-4 APs

Mixed Client/Traffic Nominal Performance Verification

- Establish voice, video, and data traffic representative of the enterprise
- Spread traffic across as many clients as the AP can support
- Incorporate legacy and 802.11n clients
- Load both radios of an AP simultaneously
- Verify Service Level Agreements are met for all traffic types and clients
- What this test reveals:
  - » Ability of controller and infrastructure to handle higher load
  - » Performance with heterogeneous traffic and client mix

Mixed Client/Traffic Maximum Performance Verification

- Repeat above test while increasing total load incrementally
- Continue increasing load until first SLA failures are observed
- What this test reveals:
  - » Correct traffic gets prioritized in overload condition
  - » Actual maximum transport capability of client-specified traffic mix, useful for capacity planning

The open-air component of this testing is useful if new AP technology, and therefore client technology, is a part of the growth plan. The essence of this portion of the test is less about scale, and more about observing how new and old client devices and access points will behave together.

2. Open-Air Client Testing

Key Client Device Performance Testing with WaveInsite

- Establish maximum forwarding rate performance available with key client devices
- Measure upstream, downstream UDP, and TCP forwarding rate
- Measure close to AP, then in distance increments of 10 feet, up to maximum AP spacing
- What this test reveals:
  - » Performance of Client/AP interaction under stressful load
  - » Maximum performance available from client device
  - » Performance degradation with distance, each direction

Site Assessment with WaveDeploy

- Network deployed and automated RF management feature active
- Perform site assessment to verify performance to SLA in target area
- Use actual client devices to view performance and traffic delivery fairness
- Use ecosystem traffic to view performance in presence of many clients
- What this test reveals:
  - » Traffic and signal level performance of network as deployed
  - » Result of automated radio/channel management algorithms on SUT
  - » Performance degradation with distance, each direction

WaveInsite testing is very useful in establishing performance expectations between particular client devices and new access points. It can be an eye-opening experience to see just what performance is actually available from systems that are regularly advertised to supply 300 or 450 Mbit transport rates, when realistic traffic and clients mixes are brought into the picture. The WaveDeploy approach is an excellent exercise, especially if a site assessment already exists for the same physical location using the previous technology. Improvements in coverage, if they exist, will be readily recognizable.

Scale for Large-Scale Testing

This testing approach is easily scaled to fit the requirements of the enterprise. The following table offers a guide to hardware needs for different size organizations.

Small Office –Single, stand-alone AP, no independent controller	• WT20 with (1) WBW2000 MIMO Blade, (1) WBE1101 Ethernet Blade. Tabletop testing, 500 client capability, test one radio at a time
Small/Medium Office – Small scale controller with up to 4 APs, or distributed controller architecture	• WT20 with (1) WBW1104N SISO blade (4 ports per blade, individual radios), (1) WBW2000 MIMO blade, (1) WBE1101 Ethernet Blade. Tabletop testing, 2000 total clients. Test all four APs at once, one radio at a time, or test both radios per AP on two APs
Office Building – One or more controllers, multiple APs	• WT90 with (2) WBW2000 MIMO blades, (1) WBW1104N SISO blade, (1) WBE1104 Ethernet blade (4 ports). Tabletop or rack-mount testing, 3000 client capability, test many combinations of radios and APs
Campus – Multiple controllers, multiple APs	• WT90 with (2) WBW2000 MIMO blades, (5) WBW1104N SISO blades, (1) WBE1104 Ethernet blade (4 ports). Tabletop or rack-mount testing, 11000 client capability, load multiple APs and controllers simultaneously.

Return on Investment

Network managers who have invested in this approach in planning for network upgrades and new services enjoy the following returns:

- Fast and efficient means of answering the key questions of risk facing the staff
- Easily repeatable measurements, providing the tools to tune and improve over time
- Strong basis of comparison data to aid in detailed discussions with network vendors
- Ability to directly measure and validate vendor claims when software or configuration changes are expected to deliver performance improvements
- Ability to present repeatable and quantifiable evidence to both network and client vendors then communicating issues for concern
- Capacity figures, measured at scale with relevant traffic and client mix, for use when planning the network topology and AP placement.
- Acceptable load levels and Service Level Agreement settings for use in client testing, overall site assessment and troubleshooting

RECORDS OF REVISIONS

Rev	Date	Description
1	February 2011	Initial Release
2	April 2011	Chapters 6 - 9

Ixia Worldwide Headquarters
26601 Agoura Rd.
Calabasas, CA 91302

(Toll Free North America)
1.877.367.4942

(Outside North America)
+1.818.871.1800

(Fax) 818.871.1805
www.ixiacom.com

Other Ixia Contacts
Info: info@ixiacom.com
Investors: ir@ixiacom.com
Public Relations: pr@ixiacom.com
Renewals: renewals@ixiacom.com
Sales: sales@ixiacom.com
Support: support@ixiacom.com
Training: training@ixiacom.com