Testing Triple Play Services and Security in Enterprise Networks

Testing scenarios, caveats and issues encountered in testing converged or "triple play" (voice, video, data) services and security in converged enterprise-class deployment scenarios

Introduction

With the maturation of IEEE 802.11x, voice over IP (VoIP), and sophisticated Ethernet bridging protocols, enterprises are deploying mixed wired and wireless networks to support a complex combination of "triple play" voice and video applications within a converged network infrastructure. The University of New Hampshire InterOperability Laboratory (UNH-IOL) organized and hosted the first multi-vendor test event specifically designed to begin validating the interoperability, performance and scalability of such networks. The UNH-IOL's first "Enterprise Services and Security" event, conducted March 21-25, 2005, incorporated hardware and software products from 13 suppliers. The open-industry event tested triple play as it would actually be deployed in a large corporate LAN serving several offices on a distributed campus.

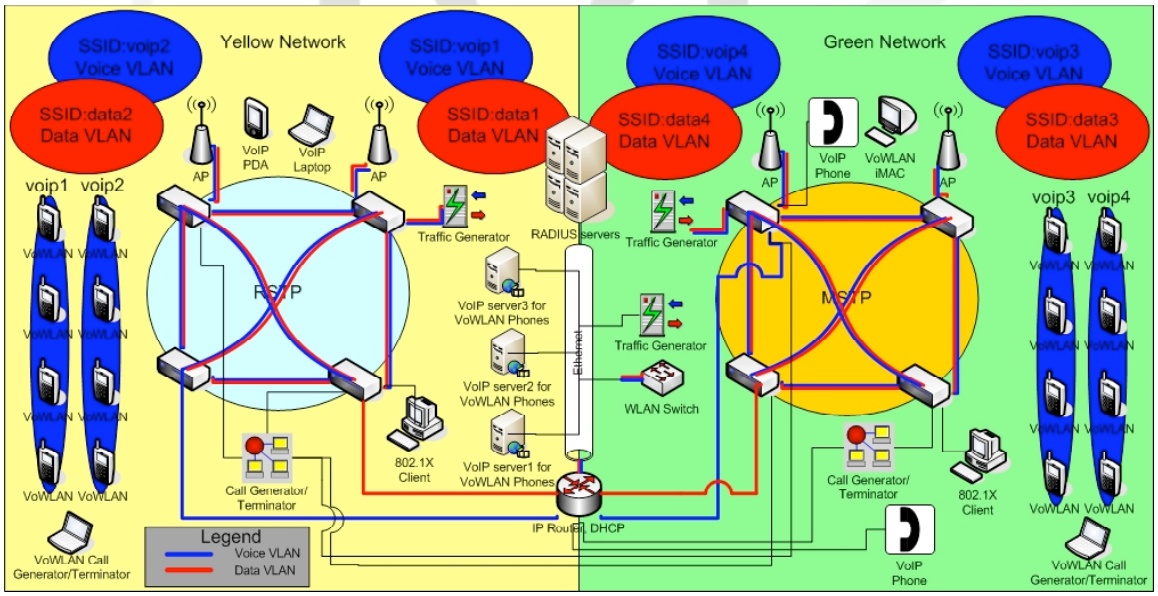

Figure 1 – Enterprise network topology

This deployment-style, multi-vendor network duplicated the mixed traffic characteristic of larger companies converging VoIP, video and wireless onto their existing data networks. The test bed included VoIP, VoIP-over-wireless, both software

and hardware voice clients, session initiation protocol (SIP) proxies, PDAs, Ethernet bridges, call generation/termination equipment, voice-quality analyzers, routers, wireless access points, laptop clients, streaming media servers and a mix of security protocols. The test was also the first time that three different UNH-IOL testing groups – the laboratory's Bridge Functions, Wireless and VoX Consortiums – were involved in a single event.

The following companies participated in the plugfest:

The overall event consisted of bridging, wireless and VoIP testing performed in two phases: small network interoperability and conformance testing; and build-out of an enterprise network to provide the underlying infrastructure for triple-play and wireless services. This white paper presents a basic overview of the technologies tested during the various portions of the event as well as a generalized view of the testing results. Network technologies tested included, VLANs, Spanning Tree(s), and wired and wireless security such as WEP, WPA (TKIP/AES), and 802.1X. Across this mesh of underlying protocols, testers ran voice and video services across the network.

Bridging Introduction

Ethernet bridging is edging out competing technologies in distributed networks because it can deliver the same data services, such as voice, video and data, at a lower cost and with reduced complexity. Industrial control systems have traditionally used proprietary data communications systems to manage their networks, while service provider networks employed various transport technologies, and enterprise networks used routers. Today, managers of large-scale wireless networks are replacing the “fat” APs, which perform a relatively heavy amount of network data processing, with “light” APs managed by an Ethernet bridge. In fact, recent advances in Ethernet bridging allow a properly configured bridged Ethernet network to achieve nearly the same results as any of the aforementioned solutions at a reduced cost. In many cases a bridged Ethernet network is easier to deploy, configure and maintain than the alternative(s).

Increased focus on Ethernet bridging since the ratification of several key standards in the late 1990s has rapidly matured the technology. IEEE Working Group 802.1 has produced a set of bridging protocols that meets the needs of a wide range of applications. Virtual Local Area Networks (VLANs) separate bridged Ethernet networks into numerous virtual networks, increasing scalability and deployment options. Spanning Tree Protocol (STP) ensures that a stable topology will arise from a randomly connected network with multiple data paths. Port-Based Network Authentication (IEEE 802.1X)

and other bridge security protocols extend the security umbrella to “Layer 2”, while still other protocols address specific network problems.

VLANs enable as many as 4,094 distinct broadcast domains within a single physical LAN, each broadcast domain a “virtual” LAN (VLAN). Separate broadcast domains reduce network overhead by containing broadcast Media Access Control (MAC) frames within a smaller LAN. Each VLAN is assigned a unique identifier, or VLAN ID. Every MAC frame transmitted on the LAN is associated with a specific VLAN, whether a tagged frame or an untagged frame. Tagged VLAN frames have the added benefit of providing a means to implement quality of service (QoS). MAC frames are not allowed to travel between VLANs, although they may be routed between different VLANs representing IP subnets. VLANs can be used to reduce the number of routers needed in a network by creating separate IP subnets in a flat, bridged network.

Within a STP-enabled LAN, a root bridge is elected, and redundant links are pruned until the active topology is a simply and fully connected tree. STP allows multiple redundant data paths to be deployed and automatically manages the use of these paths, eliminating data loops. Currently four fully interoperable versions of STP exist:

- IEEE Std 802.1D™-1998 - legacy STP
- IEEE Std 802.1w™-2001 - initial Rapid STP (RSTP) version
- IEEE Std 802.1D™-2004 - revised RSTP version; IEEE 802.1w-2001 was modified and rolled into the 2004 update of 802.1D; as a result an 802.1w RSTP bridge behaves differently than an 802.1D RSTP bridge
- IEEE Std 802.1Q™-2003 - Multiple STP (MSTP), also known as 802.1s, which was the name of the amendment to 802.1Q

Legacy STP allows for topology convergence times in the 30 to 60 second range. RSTP greatly reduced this convergence time by exchanging additional information in the Protocol Data Units (PDUs) and eliminating reliance on timers, which are restricted to enforcing worst-case delays. MSTP allows administrators to manage numerous direct data paths in the network by creating multiple instances of RSTP within the same physical LAN. In such networks, the administrator can configure each MST instance with a particular set of VLANs. This helps ensure a more optimal use of available bandwidth and allows for new and unique topologies.

Network administrators can limit and monitor device access to bridged Ethernet networks using IEEE Std 802.1X™-2001 (Port-Based Network Access Control) with a RADIUS backend. IEEE 802.1AE™ (MACsec) and IEEE 802.1af™ (KEYsec) are currently in development and will further enhance the security of such networks. Port-Based Network Access Control, often referred to as “.1X” (pronounced “dot one X”), enables authentication of devices and isolation of unauthorized devices from the network. A device (the supplicant) connecting to the network via a bridge (the authenticator) participates in an authentication process. The authenticator requests a set of supplicant credentials ranging from a certificate to a username/password and contacts the authentication server to verify the credentials. To allow for network access customization

the authentication server, generally referred to as the RADIUS server, has the option of pushing down a specific configuration to the supplicant, with attributes ranging from a VLAN ID to various vendor-defined attributes. MACsec, short for MAC security, provides a means of secure communication between devices on a LAN. KEYsec defines a framework for security key management and a method for establishing secure associations used by MACsec and 802.1X.

Bridging Testing Overview

Interoperability and conformance testing in the bridging portion of the ESS test event verified protocol functionality and identified several issues prior to deployment of the enterprise network. This initial phase focused on the testing and verification of VLANs, STP, RSTP, Filtering Database and 802.1X. Pre-testing of pre-release network devices early on identified issues in proper device configuration and network topology. Clearing these up eased the network deployment phase and improved support of triple play and wireless services.

Plugfest participants, with assistance from consortium employees, executed test plans provided by the UNH-IOL Bridge Functions Consortium for all interoperability and conformance assessment (www.iol.unh.edu/testsuites/bfc). The interoperability test plans placed the device under test (DUT) in a small test network consisting of between two and six bridges and numerous test stations with which to exercise the protocol mechanisms. The conformance test plans verified state machine and timer functionality by directly connecting the DUT to several test stations, transmitting a series of predefined frames, and noting the response.

Interoperability testing assayed each device against a sample of devices from the UNH-IOL Bridge Functions Consortium test bed, consisting of 40+ products representing a variety of network system vendors. Testing against devices from five to ten different system vendors provides a fairly high degree of confidence in the DUT's interoperability. Conformance testing allowed for further investigation of issues observed during the interoperability testing phase.

Wireless Introduction

Recent innovations in wireless technologies, such as offloading processing to network-edge WLAN switches, have helped clear the way for large-scale enterprise-class wireless network deployments. However, despite improved maintenance and administration tools, switched WLAN networks are vulnerable to issues arising from the converged network environment and the integration of services such as VoIP. Due to the high demand of VoIP services in professional settings, it is imperative to enterprise adoption that sensitive VoIP streams and data can coexist within a WLAN switch system. Network users expect network services to work over wireless LANs just as they do over wired LANs, and this will hold true for triple play services over the wireless medium as administrators add voice and video into a primarily data driven network that already has many caveats to deployment and administration. Expectations are especially high for triple play services, as end users expect service quality equal to that enjoyed by

traditional telephone and video systems.

The ESS event tested traditional WLAN systems as well as WLAN switched systems incorporating enhanced solutions from several manufacturers.

Wireless Testing Overview

The test network topology combined traditional wireless access points (APs) and newer WLAN switch systems in an environment duplicating conditions in the field. The primary goal was to determine how diverse vendors' enhanced WLAN switch systems performed in the converged network setting compared with legacy AP solutions that have been tested in the past.

Much of the testing focused on session and application awareness features that optimize wireless networks and improve coexistence between VoIP and wireless devices. Session-aware WLAN switches separate VoIP traffic and regular wireless data to prevent VoIP phones from accessing data streams and data devices from accessing VoIP streams. Application-aware functionality allows the WLAN switch to monitor all TCP/UDP ports in a wireless network or networks. The application-aware WLAN switch enables and disables certain TCP/UDP ports as required. This enhances security and greatly decreases the chance of unwanted access through open ports.

Session-aware functionality introduces additional features to WLAN Switch implementations. Chief among these is priority load balancing. Normally, wireless data and VoIP traffic would be kept on separate APs/SSIDs for ease of administration. However, there are times when VoIP and wireless devices must coexist on the same AP/SSID; this is referred to as a converged network because VoIP and data traffic share a common infrastructure. To preserve voice quality in converged networks, VoIP traffic must be given priority over other traffic types. WLANs with Session-aware functionality distinguish high priority VoIP traffic from regular wireless traffic, thus reducing or eliminating VoIP quality degradation across the network.

The testing also examined an additional benefit of session awareness, enhanced radio frequency (RF) scanning. RF scanning shuts down the light APs in a WLAN switch implementation for a brief period of time to scan the air for clearer channels and/or rogue devices. However, because even a temporary shutdown causes a noticeable drop in voice quality, some implementations disable RF scanning when integrating VoIP devices. Disabling RF scanning can compromise network efficiency and security. Enhanced RF scanning solutions use session-aware functionality to detect light APs which are engaged in VoIP data transmission. APs involved in VoIP servies are left on while other APs are used for RF scanning. Thus VoIP traffic can travel the network without compromising maintenance and security.

Obviously, VoIP traffic streams are particularly sensitive to interference and disruption; to compensate, many devices have settings for the maximum number of VoIP devices allowed and the maximum number of VoIP calls allowed. The group test examined a

wireless call load balancing feature that automatically forces a VoIP device to roam to another AP on the WLAN switch if either maximum (devices or calls) is reached.

In addition, manufacturers have developed client blacklisting systems that help further protect sensitive VoIP traffic streams. Products use this functionality to blacklist any client that violates the policy of a certain AP or SSID. The group test verified that if the blacklisting system is designed to only allow VoIP traffic on a specific AP or SSID, any client attempting to send wireless data would be blacklisted and not allowed to access the network.

VoIP Testing Scenarios

The VoIP tests performed were divided into four sections:

- call performance and call load capacity
- voice aware features
- voice QoS and
- voice security.

Each section consisted of multiple tests that verify the different aspects of each solution.

- The Call Performance and Call Load Capacity section consist of three tests. The first test is a Call Capacity test, which tests the maximum number of calls allowed per AP for both the SIP and SVP protocols. The second and third sections test Call Performance and verify if the WLAN switch system prioritized VoIP traffic when both VoIP and wireless data were transmitted through the system. In the first Call Performance test, the voice servers and light APs were plugged directly into the WLAN switch. In the second Call Performance test, they were connected with the WLAN switch via a bridged/routed network.

- The voice aware features section also consists of three tests. The first test verifies that VoIP Aware RF scanning does not compromise VoIP quality. Three different setups are tested: a baseline control test, where VoIP quality is tested with RF scanning off; VoIP quality is then measured with regular RF scanning on; and lastly, the VoIP quality is tested with Voice Aware RF scanning on. The second test in this section is the Call Admission Control test. This verified the WLAN switch system's ability to detect VoIP devices connected to the network, whether or not they are engaged in a call (referred to as "on hook" if not engaged and "off hook" if engaged), and to balance and control the number of calls per AP based on the number of active voice calls. The last portion of this testing is the Voice Application Awareness testing. This checks that the WLAN switch can both successfully distinguish the different VoIP protocols (SIP, SVP) and open the ports that are needed, as necessary.

- The third section tests the Voice QoS solutions. The first test verifies the prioritized queuing of VoIP data. This test ensures that the VoIP traffic is

prioritized higher than wireless data transmitted at the same time. VoIP prioritization is verified in the wireless medium, as compared to the Call Performance tests, which only test the wired medium. The second test of this section is the Traffic Tagging test, which ensures that the WLAN switch tags packets (for QoS purposes) to and from the AP. Both 802.1p and DSCP tags are tested. The final test is a Bandwidth Control test; this test verifies that the data traffic is limited to a certain throughput amount, so that VoIP traffic is guaranteed the bandwidth it requires.

- The fourth section is designed to test the Voice Security solutions. The first test in this section is the Limiting Network Access for Voice Users test. This ensures that VoIP devices on the voice networks are not allowed to access anything but the voice relevant portions of the network. The second and final test in this section was the Blacklisting Client test. This test ensures that any device on the VoIP network that violates the voice policies is blacklisted from the network and kept there until cleared by an administrator.

No 802.1X enabled wireless VoIP handsets were on hand during the test event to test the 802.1X security functions and features under the conditions and tests listed above. Future testing will examine the security features of wireless VoIP handsets.

Results:

Bridged network results

No major interoperability issues were noted during pre-deployment testing and enterprise network build-out. Minor interoperability issues related to network convergence and feature operation were observed on pre-release devices brought to the test event by attending vendors seeking R&D-style feedback. Devices which exhibited interoperability issues with the possibility of affecting network performance were excluded from the enterprise network and scheduled for further conformance testing.

With the goal to support triple-play services, the enterprise network effectively deployed VLANs, RSTP, MSTP, 802.1X and Link Aggregation with no observed interoperability issues. VLANs partitioned the two network traffic types: VoIP and data. RSTP and MSTP eliminated data loops by pruning the core bridged network's mesh topology to a fully and simply connected network. Test scenarios demonstrated the ability of 802.1X to limit access to the network resources by requiring authorized credentials from the network device.

Each bridged network was connected to the central router via two 1 Gbps uplinks, one for each VLAN (VoIP and data). One bridged network deployed a single instance of RSTP, the other ran MSTP configured with a VoIP instance and data instance.

Proof-of-concept testing demonstrated the operation of 802.1X across a Spanning Tree enabled bridged Ethernet network, and through a router to the WAN. 802.1X operates in Layer 2 between the supplicant and authenticator. RADIUS, the authenticating backend, operates at Layer 3 between the authenticator and RADIUS server, allowing an authenticator to authenticate with a RADIUS server outside its subnet. Participants authenticated a variety of supplicants from Apple, Funk, Meetinghouse and Microsoft with RADIUS servers from Cisco, FreeRADIUS, Funk, Infoblox, Meetinghouse and Microsoft.

Wireless network results

The wireless network environment poses specific vulnerabilities for voice and video applications. Testing of triple play has been extensive in broadband and other wired transport environments, but there has been relatively little public multi-vendor testing of these technologies in converged wired and wireless LANs. Among other findings, the results detailed below show that it is absolutely necessary for systems to be voice aware and to perform the various functions detailed in the tests successfully without disruption of the voice call setup, teardown, or even the voice quality of the calls.

For example, voice-aware functionality tests proved that, consistent with their intended purpose, enhanced RF scanning features did not have any noticeable effect on voice quality. R-Factor values recorded only varied by a value of 0.6 within each run, and the results were consistent with or without the enhanced RF scanning features enabled. By contrast, with typical RF scanning features turned on (but not voice sensitive, i.e. not enhanced RF scanning) no R-Factor results could be recorded whatsoever, as the calls would drop when the AP went into RF scan mode. This shows that a particular implementation's feature set must be completely and fully integrated with, and aware of the services running on the system. For instance, if the system knows that there are no voice clients, services or sessions running on an AP, then a typical RF scanning function may be used without detriment. However, when such services are present, enhanced RF scanning features must be used in order to avoid any disruptive effects on the voice services.

Detailed below are some of the results collected during the recent test event. They correspond to the testing sections described above in the Wireless Testing Overview.

Results from Wireless Network Section 1: Call Performance and Call Load Capacity Tests

The UNH-IOL was unable to perform this testing to the fullest capacity at the ESS event due to the large number of phones required. The laboratory plans to test full load performance and capacity testing at a future date.

Results from Wireless Network Section 2: Voice Aware Feature Tests

The Device Under Test (DUT) first underwent the Voice Aware RF Management tests. The DUT had 14 phones connected to it, and they engaged in 7 calls. For each part of this test, the wireless testing device measured an "R-Value" for each phone call. An R-Value is the quantitative measurement of VoIP quality in a network, and is one of the

preferred measurements of VoIP quality. An R-Value of 75 or above is considered toll quality or better. The first part of the test ran with RF scanning disabled on the DUT. This provided a control value that represented the DUT's performance when unencumbered. The following table provides the values recorded during this test.

Test Label	1st Test Run	2nd Test Run
R-Value of VoIP flow 1	77.9	77.3
R-Value of VoIP flow 2	77.9	77.8
R-Value of VoIP flow 3	78.0	77.9
R-Value of VoIP flow 4	No Data	No Data
R-Value of VoIP flow 5	77.8	77.8
R-Value of VoIP flow 6	77.9	77.8
R-Value of VoIP flow 7	77.9	77.9

Table 1 – Calculated R-Values w/ RF scanning disabled

The second part of this test ran with the RF scanning enabled on the DUT but with no Voice sensitive features. This showed a significant drop in voice quality, which caused broken voice output; hence no voice quality metrics were captured.

The final portion of this test ran with enhanced RF scanning enabled on the DUT. As evidenced by the table below, enhanced RF scanning features successfully preserved the integrity of the voice traffic stream. The R-Values are comparable to the first test without RF scanning; in this test however, the enhanced RF scanning provided maintenance and security without sacrificing VoIP quality.

Test Label	1st Test Run	2nd Test Run
R-Value of VoIP flow 1	77.9	77.9
R-Value of VoIP flow 2	77.8	77.9
R-Value of VoIP flow 3	77.8	77.9
R-Value of VoIP flow 4	No Data	No Data
R-Value of VoIP flow 5	77.9	77.8
R-Value of VoIP flow 6	77.3	77.8
R-Value of VoIP flow 7	77.5	77.9

Table 2 – Calculated R-Values w/ RF scanning enabled

The Call Admission Control test was the next test performed in this section. Call admission control and Call load balancing were enabled on the DUT. The test began with 12 wireless VoIP handsets connected to a light AP (AP 1) controlled by the DUT. Another light AP (AP 2) with no VoIP devices connected to it was brought up with the same settings. This test used the SVP protocol, with a maximum number of 4 SVP calls allowed. The following output from the DUT's command line interface (CLI) displayed the status of the phones at the beginning of the test.

```
VoIP Client Table (Flags: C = call in progress)
-----------------
MAC Address       Location   BSSID              ESSID VLAN  Protocol  Flags
-----------       --------   -----              ----- ----  --------  -----
```

```
00:90:7a:02:4a:09   AP 1     00:0b:86:c0:f9:60  voip   1    svp
00:90:7a:02:15:d2   AP 1     00:0b:86:c0:f9:60  voip   1    svp
00:90:7a:02:36:42   AP 1     00:0b:86:c0:f9:60  voip   1     svp
00:90:7a:02:38:a5   AP 1     00:0b:86:c0:f9:60  voip   1     svp
00:90:7a:02:49:af   AP 1     00:0b:86:c0:f9:60  voip   1     svp
00:90:7a:02:19:31   AP 1     00:0b:86:c0:f9:60  voip   1     svp
00:90:7a:02:4a:03   AP 1     00:0b:86:c0:f9:60  voip   1     svp
00:90:7a:02:49:f5   AP 1     00:0b:86:c0:f9:60  voip   1     svp
00:90:7a:02:2f:85   AP 1     00:0b:86:c0:f9:60  voip   1     svp
00:90:7a:02:2b:82   AP 1     00:0b:86:c0:f9:60  voip   1     svp
00:90:7a:02:0e:d6   AP 1     00:0b:86:c0:f9:60  voip   1     svp
00:90:7a:02:24:b8   AP 1     00:0b:86:c0:f9:60  voip   1     svp
Number of Clients:12
```

Figure 1 – *The four VoIP devices placed phone calls after this setup was verified. This was the maximum number of calls allowed per AP, and therefore the DUT was expected to correctly balance the system, by moving all of the "on hook" (not engaged in phone calls) devices to a neighboring light AP (AP 2).*

```
VoIP Client Table (Flags: C = call in progress)
-----------------
MAC Address       Location  BSSID              ESSID VLAN   Protocol   Flags
-----------       --------  -----              ----- ----   --------   -----
00:90:7a:02:4a:09   AP 2    00:0b:86:c0:f9:60   voip   1     svp
00:90:7a:02:15:d2   AP 2    00:0b:86:c0:f9:60   voip   1     svp
00:90:7a:02:36:42   AP 2    00:0b:86:c0:f9:60   voip   1     svp
00:90:7a:02:38:a5   AP 2    00:0b:86:c0:f9:60   voip   1     svp
00:90:7a:02:49:af   AP 2    00:0b:86:c0:f9:60   voip   1     svp
00:90:7a:02:19:31   AP 2    00:0b:86:c0:f9:60   voip   1     svp
00:90:7a:02:4a:03   AP 2    00:0b:86:c0:f9:60   voip   1     svp
00:90:7a:02:49:f5   AP 2    00:0b:86:c0:f9:60   voip   1     svp
00:90:7a:02:2f:85   AP 1    00:0b:86:c0:f9:60   voip   1     svp         C
00:90:7a:02:2b:82   AP 1    00:0b:86:c0:f9:60   voip   1     svp         C
00:90:7a:02:0e:d6   AP 1    00:0b:86:c0:f9:60   voip   1     svp         C
00:90:7a:02:24:b8   AP 1    00:0b:86:c0:f9:60   voip   1     svp         C
Number of Clients:12
```

Figure 2 – *The DUT correctly detected the call status of a connected VoIP device, as evidenced by the CLI output. The DUT proactively moved all "on hook" VoIP devices to AP2, which correctly balanced out the system when the maximum number of calls was reached.*

The Voice Application Awareness test, the final test of this section, assayed voice application awareness enabled on the AP WLAN systems. This solution ensured that specified ports were only opened when VoIP calls of various protocols were engaged (SIP, SVP, Cisco SCCP). All other ports should have remained closed. Successful phone calls made using the SIP, SVP, and SCCP protocols verified that this solution worked properly. The AP WLAN systems properly detected and opened ports for specific VoIP traffic protocols, thereby passing the test.

Results from Wireless Section 3: Voice Quality of Service Tests

The third section of tests focused on the ability of wireless APs to prioritize voice traffic. The prioritized queuing test, the first test in this section, tested the wireless AP with prioritized queuing enabled such that VoIP traffic had a higher priority than data

traffic. Eight devices connected to the wireless AP established four VoIP phone calls, while a wireless testing device transmitted data traffic through the AP at 4 Mbps. The VoIP calls demonstrated good voice quality, but minimal data traffic from the wireless testing devices was observed. This indicated that, for the most part, the AP properly prioritized the VoIP traffic higher than the data traffic. It should be noted that as a consequence of the prioritized queuing, the service represented by the wireless data traffic was impaired.

The Traffic Tagging test checked the ability of the AP to tag packets in Differentiated Services Code Point (DSCP) format and transmit them to the LAN. The test ran with the traffic tagging option enabled on the AP, and an Ethernet trace was taken from the LAN. This trace verified that the AP tagged the packets as intended, so that a DSCP-enabled LAN could process the QoS information contained in the AP's traffic.

The Bandwidth Control test, the last test in this section, tested the AP with the Bandwidth Control option configured to allow only 1 Mbps of throughput from the wireless testing device. This limited the amount of wireless data and reserved the medium for throughput-sensitive VoIP traffic. A wireless testing device transmitted data at 4 Mbps through the AP for a specified amount of time. At the end of the test, the wireless test tool verified that only 1 Mbps of wireless throughput was received. This verified that the APs correctly limited bandwidth for wireless data users and effectively preserved the quality of the VoIP traffic streams.

The results from this QoS testing demonstrated the features commonly supported by wireless APs in relation to VoIP network traffic. Integration of these features into an overall QoS network architecture would improve the networks' ability to prioritize traffic types from end-to-end.

Results from Wireless Section 4: Voice Security Tests

The final section tested the VoIP network's security features. The first test ensured the ability of the AP WLAN system to limit the access for voice devices. This test ran with a feature enabled that secures data networks from relatively unsecured voice networks. This feature works by limiting the access for devices on a voice SSID to only voice ports and other voice devices. A wireless data client established a connection to the limited voice SSID, and all ICMP traffic sourced to a non-VoIP device failed. The AP WLAN system denied all attempted access to a non-VoIP device or port, verifying the functionality of the limited voice network access feature.

The Blacklisting Clients test, the second and final test of the section, ensures that devices on a voice network that attempted to access non-voice devices or use non-voice data protocols are placed on a blacklist, wherein all network access is prohibited. First, a wireless data client connects to the voice SSID where blacklisting is enabled. Second, the wireless client attempts to access the data network via a web browser. The AP WLAN system correctly denied all access to the data network. The AP WLAN system also added the wireless data client to the Blacklist. This is demonstrated by the AP WLAN system CLI output seen below. The device highlighted in bold letters is the data client.

MAC Address	Location	BSSID	ESSID	Role	Age(d:h:m)
00:0d:28:2e:8f:9d	AP1	00:0b:86:c0:fa:60	voip	cisco	00:01:07
00:90:7a:02:15:d2	AP1	00:0b:86:c0:fa:60	voip	phones	00:01:06
00:0e:38:50:26:ff	AP1	00:0b:86:c0:fa:60	voip	cisco	00:01:07
00:90:7a:02:4a:03	AP1	00:0b:86:c0:fa:60	voip	phones	00:01:07
00:90:7a:02:49:af	AP1	00:0b:86:c0:fa:60	voip	phones	00:01:07
00:90:7a:02:36:42	AP1	00:0b:86:c0:fa:60	voip	phones	00:01:07
00:90:7a:02:38:a5	AP1	00:0b:86:c0:fa:60	voip	phones	00:01:07
00:90:7a:02:24:b8	AP1	00:0b:86:c0:fa:60	voip	phones	00:01:07
00:90:7a:02:19:31	AP1	00:0b:86:c0:fa:60	voip	phones	00:01:07
00:90:7a:02:2b:82	AP1	00:0b:86:c0:fa:60	voip	phones	00:01:07
00:90:7a:02:4a:09	AP1	00:0b:86:c0:fa:60	voip	phones	00:01:07
00:90:7a:02:49:f5	AP1	00:0b:86:c0:fa:60	voip	phones	00:01:07
00:90:7a:02:2f:85	AP1	00:0b:86:c0:fa:60	voip	phones	00:01:07
00:90:7a:02:0e:d6	AP1	00:0b:86:c0:fa:60	voip	phones	00:01:07
00-05-4E-44-20-3B	**AP1**	**00:0b:86:c0:fa:60**	**voip**	**phones**	**00:00:10**

Figure 3 – ***The wireless client appeared on the blacklist after it attempted to access the data network.***

DoS STA List

STA	reason	block-time(sec)
00:05:4E:44:20:3B	**session-blacklist**	**60**

Figure 4 – *The DUT denied all attempts to communicate on the network after the wireless data client was blacklisted, until an Administrator manually removed the client from the blacklist.*

In summary, prioritized queuing, bandwidth control, and voice sensitive RF scanning all help to protect sensitive VoIP streams and preserve the voice quality. Assigning roles and limiting access provides security for data networks that may contain private information; the Blacklisting feature provides additional protection from devices that attempted to break these rules. The results of this test event indicate that while there are obstacles to integrating the two network mediums, solutions exist to overcome them.

VoIP Network Specific Results

The most striking results of the voice portion of the testing involved findings related to the scalability of wireless APs. The test results effectively doubled the number of high-quality VoIP calls previously recorded on a single AP, from seven to between 14 and 16. A single AP was found to scale up to 24 VoIP calls, but not without seriously impairing the quality of service. At one extreme, the testing loaded the network with 24,000 calls, achieving a steady state rate greater than 12,000 calls, with 10 calls per wireless access point and 11,990 calls over the wired network. VoIP test solutions in a 17-hour period pushed 1.7 terabytes of voice traffic over the network, both wired and wireless.

Several call generation and termination systems were present at the test event. All test systems were capable of measuring both intrusive (PESQ, PSQM+) and non-intrusive (R-factor, MOS scores) voice quality measurements. One of the goals of the test

event was to setup the test networks with many data flows on the data VLAN portions of the network and simultaneously use the voice VLAN portion of the network to its fullest capacity. This is useful for several reasons, including the ability to characterize the network's call capacity, another is to determine the effect on voice quality when the underlying network is operating at or above its design limits or under specific extreme conditions.

Among the specific issues, some wireless systems were found to be able to connect erroneously to the data side, resulting in dropped calls, where others employed mechanisms to control which network they were associated with.

Of the three test solutions at the event, each varied in the specific features, but the biggest differentiation was call capacity. There is a conundrum in the voice testing space, in that "one call" requires two endpoints, i.e. a caller and a callee. So, each call typically is counted as two actual calls. This introduces a certain ambiguity into gauging the capacity of various systems. One call is two endpoints communicating via a bi-directional RTP stream. One system as configured was able to originate 120 calls and terminate another 120 calls, thus making 240 concurrent calls; another was able to handle a maximum of 144 on each side, thus peaking at 288 total calls (sometimes still referred to in TDM terms, aka "channels"). However the highest capacity system available at the test event exceeded both systems as previously noted, and could handle 12,000 originating and 12,000 terminating channels, for 24,000 calls.

By attempting to run all available channels from all voice test systems, the maximum capacity of the underlying wired network was discovered to be not a static number. Instead, it fluctuated between approximately 12,000 and 21,500 connected endpoints, (6,000 to 10,250 concurrent RTP streams). These numbers varied based on the network conditions, and seemed to depend on the active packet paths within each network.

Video Testing results

Of all three services tested, video proved the most fallible. However, the bottleneck did not appear to be the network itself so much as limitations to the computational power and speed of the clients employed. Video quality appeared to depend upon the specific *combination* of transport mechanism (wired and wireless) and client, rather than either one in isolation.

Video testing during the event used several actual video streams concurrently. All the streams sourced from the same device, and video server. The video server streamed a variety of files using IPv4/IPv6 multicast, thus providing efficient delivery of video streams to the various clients from a known multicast source, and allowing for much lower bandwidth consumption when compared to multiple unicast video streams to specific clients. This also allowed for the flexibility to have client stations move from place to place throughout the network and still be able to receive the video stream without having to modify any settings (i.e. the video sources came from various port

numbers from a single multicast IP address). With this configuration the client may always receive the video content regardless of IP address.

The software used for both the client side and the server side of the video streams was Video LAN Client (VLC) a free piece of software specifically designed for this purpose. Both Linux and Windows client machines were used on both wired and wireless connections to the network. Some of the test machines were older laptops with 802.11b cards and limited RAM and Processor speeds. With these systems the limitations of the underpowered systems and the limited network speed (only 802.11b, 11Mbps best case) showed that only a single video stream could be received, and the video playback was less than optimal due to dropped/lost packets (due to a variety of conditions) and their associated retransmissions at the Media Access Layer (802.11 MAC). It is important to note that the IP/UDP video stream running on the 802.11 links would not, at the IP/UDP layer attempt retransmissions of lost or dropped frames, however it was the 802.11 MAC layer which did attempt to retransmit said frames.

In the wired Ethernet scenario(s), since full duplex switched environments were used, lost or dropped packets were very unlikely and collisions should never occur. The observed behavior for the wired network was superior. Where the wireless setup had trouble maintaining good quality video playback of one stream, the wired setup reliably received and played three to four simultaneous video streams (albeit the memory and processors were severely taxed on these older systems). Demonstrating both voice and video performance on less-than-optimal hardware configurations is key when testing triple play services, as rarely would every system attempting to use the triple play services be of sufficient specifications to guarantee consistent behavior.

In addition, newer and more powerful laptop systems were used with faster wireless LAN cards that were compatible with 802.11g and 802.11a standards. In scenarios in which the laptops had more than adequate memory and processor speed, they easily maintained up to four simultaneous DVD-quality video streams using either wired or wireless connections.

Summary

Few large enterprise networks, most of which employ some combination of wired and wireless infrastructures, will completely replace their existing systems to accommodate demand for voice and video services. Rather, administrators will gradually replace and build out their networks, adding the new services to their existing systems. Continued testing of triple play services in enterprise-class networks should help raise awareness of the complexities and caveats specific to the sensitive, high-demand services of voice and video in converged network environments.

The first UNH-IOL Enterprise Services and Security test identified diverse issues inherent to running VoIP and wireless devices in converged enterprise LANs; however these issues are not without solutions. In some cases, solutions implemented to improve performance in homogenous wireless data networks can potentially have an adverse effect on the quality of triple play services in the converged wired and wireless LAN.

Examples of these enhancements include roaming, client blacklisting and detection of rogue APs. In environments such as triple play, where packet arrival and overall QoS are paramount, these will be persistent concerns. Manufacturers are addressing these issues, and they should not prove to be impediments in the long term.

However, additional testing is needed to assess QoS in true real-world networks. This test did not assess the mapping of quality of service from the wireless to the wired side and vice versa. Most deployment topologies will run voice calls over a wireless AP through some kind of network cloud to a SIP server on the other side. Prioritization will need to be end-to-end on such a network, otherwise the QoS on the AP will break down in the cloud and voice and video quality will degrade.

Overall, the test results indicate a positive outlook for enterprises looking to combine the power of triple play services with the versatility of a converged network. However, the tests revealed an assortment of technical caveats and issues specific to such deployments. These issues and other issues will need to be addressed as existing networks add new services and emerging protocols to ever more complex network topologies.