WIRELESS LAN RELIABILITY IN THE UNWIRED ENTERPRISE

Assuring Enterprise-Wide Mobility with a Nonstop Wireless LAN

Copyright © 2011, Juniper Networks, Inc.

Table of Contents

Executive Summary . . . 3
Introduction . . . 3
Signal Integrity . . . 4
Predictive RF Modeling . . . 4
Locating Rogue Access Points and Sources of RF Interference . . . 5
System Reliability . . . 5
Controller Virtualization Benefits . . . 7
Resource Availability . . . 8
Maintaining Reliable Quality of Service (QoS) . . . 8
Seamless Mobility . . . 9
Holistic Manageability . . . 10
Unified Wireless LAN Infrastructure and Mobility Services Management . . . 11
Conclusion . . . 11
About Juniper Networks . . . 12

Table of Figures

Figure 1: Predictive RF modeling coverage map . . . 5
Figure 2: Hot standby versus virtualization . . . 6
Figure 3: Virtual controller cluster operation . . . 6
Figure 4: Hitless failover . . . 7
Figure 5: Access point rebalancing . . . 7
Figure 6: Seamless mobility campus-wide . . . 9
Figure 7: Centralized versus distributed forwarding . . . 10

Copyright © 2011, Juniper Networks, Inc.

Executive Summary

The unwired enterprise has long been a goal in most organizations, and the advent of 802.11n now makes achieving that goal a practical reality. 802.11n and the proliferation of mobile devices are driving the migration from wired to wireless networks throughout the enterprise. But are wireless LANs up to the challenge? Just as wireless security was the biggest barrier to enterprise-wide adoption a few years back, a lack of continuous wireless connectivity is now preventing the unwired enterprise from becoming a ubiquitous reality. This white paper examines what it takes to maximize availability in the unwired enterprise, and explores the five elements needed to make the wireless LAN as dependable as wired Ethernet.

Introduction

Wireless LANs used to be considered a convenience, existing merely to service applications that required mobility to be effective. Those applications still exist, of course, and still require mobility. Several are even now mission critical, such as medical record access at the point of care, production line monitoring, mobile point-of-sale kiosks, along with many others.

With IEEE 802.11n fully standardized, industry pundits are now predicting a full-fledged wired to wireless migration, accompanied by the erosion of wired Ethernet access layer port sales. Dell'Oro Group's Q4 2010 Ethernet Vendor Report found that the global LAN switch port market is already flat, and is even declining in some geographies (excluding data center switch ports of 10GbE and above). Ethernet's future growth in the enterprise now exists, instead, mostly in the backbone and the data center.

> *"802.11n puts pervasive mobility on the fast track. IT professionals should start thinking now about how they will deploy, maintain, and benefit from an all wireless LAN."*
>
> —Paul DeBeasi, research director, Burton Group

> *"The war is over. Any new Wi-Fi deployments or retrofits need to be 802.11n-compatible. Enterprises will increasingly look to wireless technology as a viable alternative to recabling."*
>
> —Mike Jude, research analyst, Nemertes Research

The enhanced productivity and many other benefits of mobility have always been quite attractive, and the proliferation of mobile devices like tablets and feature-rich smartphones now makes the unwired enterprise even more compelling. One needs look no further than one's own mobile phone use to realize the advantages of anywhere, anytime voice communications. Voice over wireless LAN (VoWLAN) lets users roam throughout facilities and across campuses without the recurring cost of cell phone air time, while significantly reducing voicemail messages and phone tag. Just as it makes sense to untether voice, so too does it make sense to untether data. Having an unwired enterprise also lowers costs by requiring less wiring and equipment, and effectively eliminating user moves, adds and changes. And it is more environmentally-friendly owing to its smaller footprint and lower power consumption.

Before wireless access can fully replace wired connectivity, however, it must be as mission critical as the many applications it supports. Just as wireless security was the biggest barrier to enterprise-wide adoption a few years ago, a lack of continuous wireless connectivity is now preventing the unwired enterprise from becoming a reality. This white paper explores the five elements needed to make the wireless LAN as dependable as wired Ethernet:

- Signal integrity
- System reliability
- Resource availability
- Seamless mobility
- Holistic manageability

Copyright © 2011, Juniper Networks, Inc.

Failure is simply no longer an option with today's unwired enterprise, where users depend on full mobility with constant, seamless access, consistent levels of performance, and few or no disruptions. "Failure" can take many forms from a user's perspective, ranging from a VoWLAN call being dropped to a widespread outage. The former can be tolerated occasionally; the latter simply cannot—ever.

Consider this: If a wired LAN access port fails, one user is affected. If that same port backhauls a wireless access point, dozens of users are typically affected—a number comparable to the failure of an entire Ethernet access switch. And if a WLAN controller fails, hundreds if not thousands of users might be affected. When properly implemented, the five elements covered here eliminate widespread outages and minimize the occurrence of isolated outages, resulting in a nonstop wireless experience.

Signal Integrity

"Clean air" is such a wholesome-sounding concept. Who wouldn't want to breathe pure, clean air? But there are two problems with this concept in the context of wireless LANs. The first is that there is no such thing as "clean air" in the unlicensed RF spectrum. With so many wireless devices running so many different applications, there will always be the potential for RF interference. The second fallacy is that the situation is hopeless without buying some new, sophisticated spectral analysis tool and upgrading to the latest generation of premium priced access points. WLAN solutions have long had effective techniques for detecting and mitigating against RF interference. From powerful planning tools that employ 3D modeling during the installation to features like rogue access point detection, the air is made sufficiently clean for dependable operation of wireless LANs today.

What will change in the future is the ability to have more sophisticated spectral analysis capabilities integrated directly into the access points. This capability depends on certain features being incorporated into the RF chipsets that made their debut in the middle of 2010, and vendors are now beginning to take advantage of these features in their network management systems. Over time, the built-in spectral analysis capabilities of these new RF chipsets will enable quite sophisticated capabilities for improving performance, supporting enhanced service levels, assisting with troubleshooting efforts, and implementing new mobility services. Of critical importance will be how well vendors are able to integrate spectral analysis and its many potential applications with other mobility services, especially those involving real-time location services and advanced security provisions. For example, spectral analysis should be able to reveal the presence of a rogue access point, but will it be able to identify the location precisely enough to quickly remove this threat to security?

In addition, IT managers should realize that these additional spectral analysis capabilities do not require dedicated sensors. The chipsets make it possible for a properly designed access point to function simultaneously in both spectrum scanning mode and client services mode. Sufficient spectral information can be gathered with a one in four, or one in five sensor-to-client services ratio.

An upgrade to these new access points purely to gain their hardware-based spectral analysis capabilities will be very difficult to cost-justify, however, because it will require replacing all (or most) access points. If such a "rip and replace" exercise is part of a migration to 802.11n, then the additional cost for the spectral analysis-enabled chipset should be negligible. But for organizations that have already implemented 802.11n, existing provisions for detecting and mitigating against RF interference are sufficient for achieving dependable LAN connectivity. And because these existing provisions are implemented entirely in software, without any dependence on the newer chipsets, enhancements that improve capabilities over time are possible as part of routine updates.

Predictive RF Modeling

A prudently planned and properly configured network is far more dependable than one that is neither. This is why deploying a high availability wireless LAN begins with good RF modeling tools. Predictive modeling ensures successful operation of device-level resiliency features, such as auto-RF tuning to automatically increase the power of access points to compensate for a neighboring access point's outage. Although this RF power adjustment occurs automatically when a neighboring access point fails, it is important to have each and every access point placed properly to fill in any such voids. Planning tools with predictive modeling use what-if RF coverage scenarios to ensure good signal quality and performance at all times, in all locations, and under all conditions.

Copyright © 2011, Juniper Networks, Inc.

The figure below shows how a predictive planning tool can present graphical what-if analysis views of RF coverage in various layouts and failure scenarios. After importing an AutoCAD drawing, or a JPEG or GIF drawing of the facility, network planners are able to perform predictive analysis in three dimensions to see how coverage holds up under various scenarios for both indoor and outdoor wireless LAN deployments.

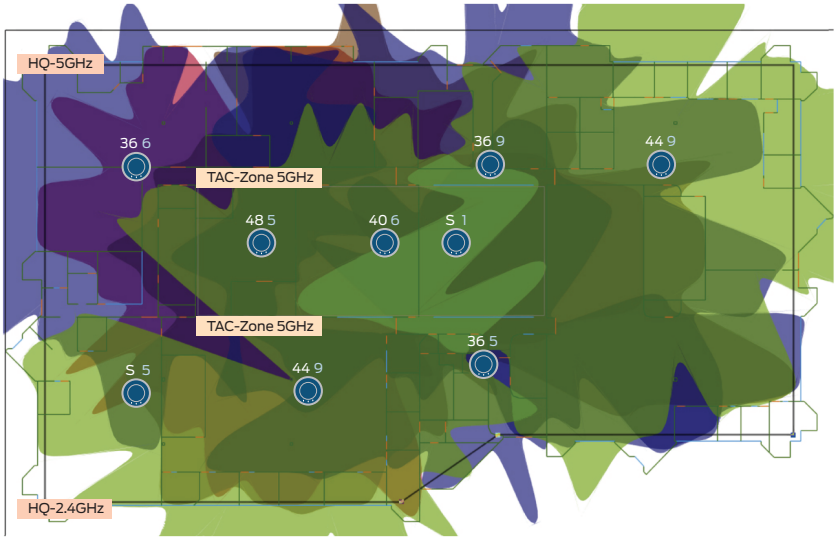

Figure 1: Predictive RF modeling coverage map

Locating Rogue Access Points and Sources of RF Interference

Most wireless LAN solutions are able to detect the presence of RF interference, whatever the cause. But very few are able to pinpoint the location with real precision. In some situations, the source and its location may be both benign and obvious, such as a microwave oven in a lunch room. But in other situations, the problem could be far more serious, such as a rogue access point being backhauled by a third-generation/fourth-generation (3G/4G) cellular service on a tablet or smartphone. In this situation, the rogue is bypassing enterprise security provisions and may be causing interference for the enterprise wireless LAN.

The most accurate means for detecting device (and, therefore, user) locations is with a solution that supports real-time location services. Real-time location systems (RTLS) employ triangulation or other means to precisely pinpoint a wireless device's location. Although the degree of precision can vary among RTLS solutions, all are far more precise than determining location based solely on the access point being used (which could be on another floor). There are numerous applications for RTLS, but in the context of signal integrity, its value exists in the ability to quickly and precisely locate the source of any RF interference. In the unwired enterprise, the RTLS solution should also be as effective outdoors throughout a campus setting as it is indoors.

System Reliability

The best way to make the wireless LAN, as a system, immune to widespread outages is to implement redundancy in all critical elements, especially at the centralized controllers. The traditional method for deploying redundant controllers has been to install a "hot standby" or secondary controller that takes over when the primary fails. Virtual Router Redundancy Protocol (VRRP) is often used in this configuration. For it to work seamlessly, the standby must constantly track the state of the primary (or primaries), including all active sessions and their security associations. The secondary controller must also have some means of detecting a failure in the primary (or primaries). This may all work well enough, but there is a major problem with the approach: It becomes very expensive and exponentially more complex as the network scales to accommodate more users.

A far superior way to implement active/active redundancy is by virtualizing all of the controllers. Virtualization has become increasingly popular for both servers and network-attached storage based on both its cost saving and reliability enhancing advantages. The same approach and advantages now make sense for the mission critical wireless LAN.

Copyright © 2011, Juniper Networks, Inc.

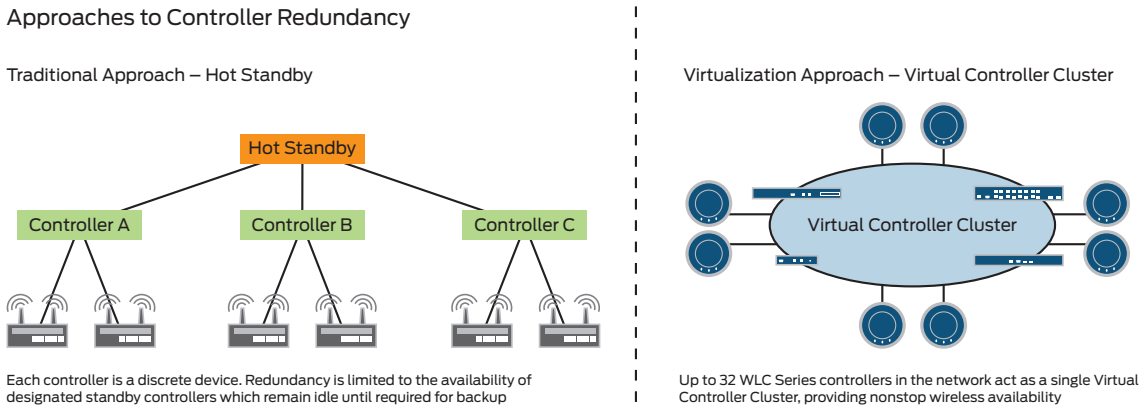

Figure 2: Hot standby versus virtualization

Maximum reliability is achieved with controller virtualization that supports either N:1 or N:N automatic failover for controllers, as well as for the association of access points to controllers. While the N:N configuration is more difficult to engineer because it requires a distributed architecture at its foundation, those solutions that do possess this capability achieve far superior redundancy for active sessions. The reason is that N:N is a form of many-to-many redundancy where all controllers act as backups for one another. This requires that each carry a copy of the configuration of all other controllers at all times. With this holistic, self-healing approach, all controllers are in constant operation, and if one should fail, one or more of the others instantaneously takes over control of the failed controller's access points and current session. The result is exactly what users want—a highly resilient wireless LAN.

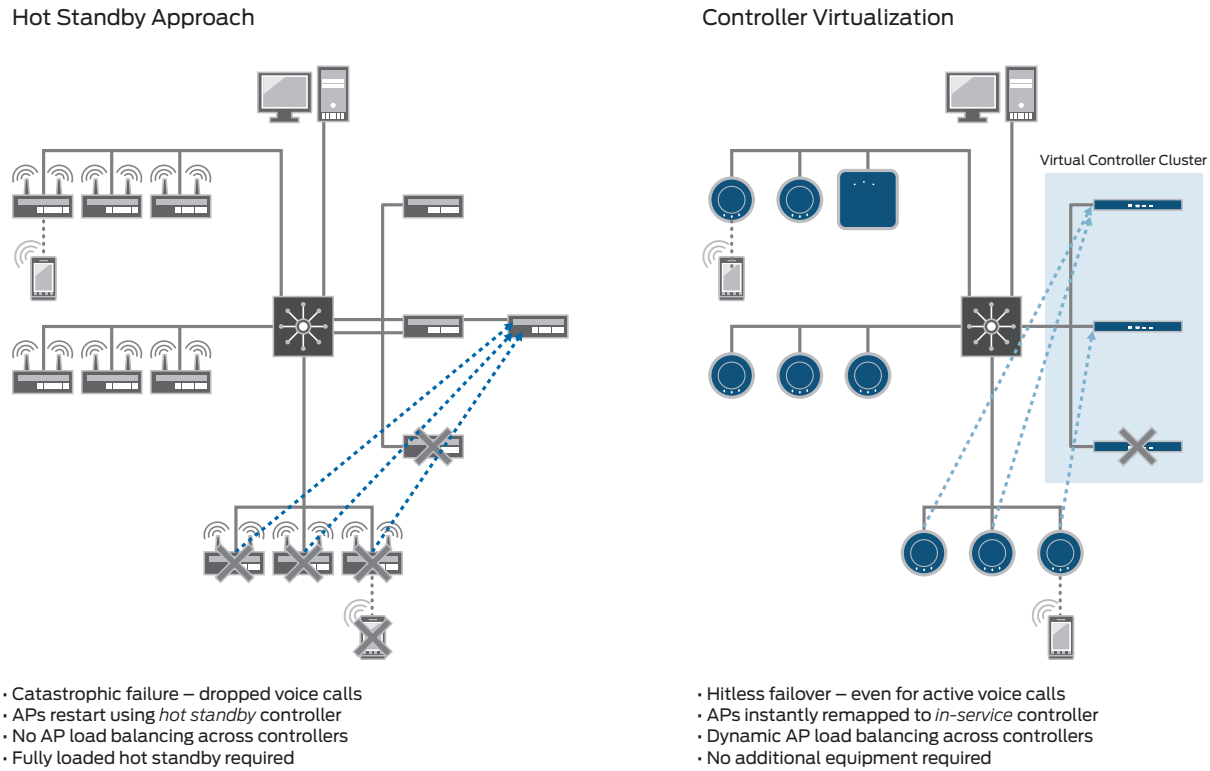

Figure 3: Virtual controller cluster operation

There are two other high availability provisions that normally accompany virtualized controllers. The first is use of redundant, hot-swappable power supplies in each WLAN controller. This prevents some failures of, and the urgent need to repair, an individual controller. The second is load balancing and/or redundancy for the authentication, authorization and accounting servers. AAA servers are mission critical resources that provide network access control, and a complete failure in this separate system would effectively "disable" the wireless LAN by preventing users from being able to log into the network.

Copyright © 2011, Juniper Networks, Inc.

Controller Virtualization Benefits

Just as with virtualized servers, virtualizing the WLAN controllers affords some significant operational benefits without increasing management complexity. The reason is this—the entire cluster of controllers is normally managed as a single, cohesive unit. In other words, only one "master" controller in the virtual cluster is used to configure and manage the entire configuration of controllers, regardless of their number. The single virtual cluster also normally shares a common configuration database, which is then propagated throughout the entire cluster, thereby eliminating the cost and complexity of managing and maintaining different configurations on each controller.

Here are the three other benefits of virtualization that pertain to high availability:

- **Hitless failover**—As shown in the figure below, if one controller fails, the access points currently under its control automatically use another. The failover happens automatically and quickly because each and every access point has an active dual-homed connection to multiple controllers, with only one designated as the primary. In this way, sessions and data flows continue seamlessly without interruption.

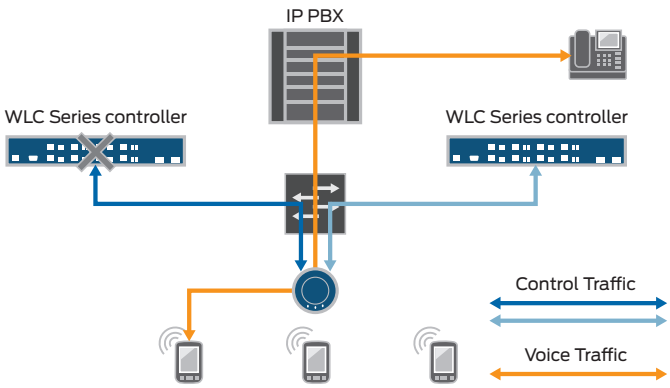

Figure 4: Hitless failover

- **Zero downtime maintenance**—Just as in a failover scenario, when any controller requires servicing, its access points are automatically reassigned in real-time across the virtual cluster, again without interruption. As a result, moves, adds, and changes, along with periodic maintenance tasks that require powering down or resetting a controller, will not affect any services or users.
- **Expansion without disruption**—Because the entire load is balanced evenly among all of the controllers in the virtualized cluster, new controllers can be added to scale overall capacity incrementally and seamlessly. The new controller automatically inherits its configuration, and the total load is then rebalanced accordingly, again without any service interruption, as shown in the figure below.

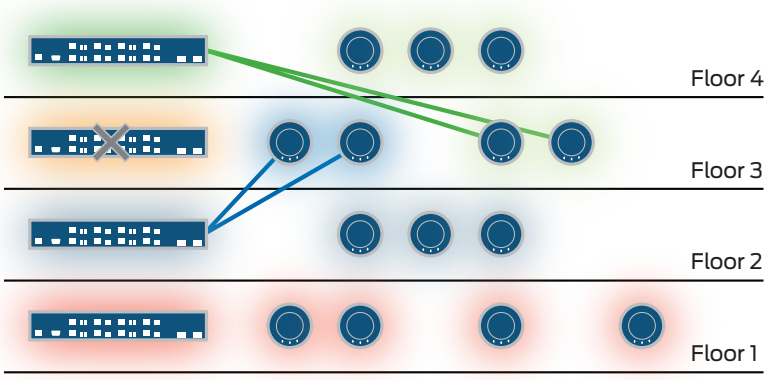

Figure 5: Access point rebalancing

Copyright © 2011, Juniper Networks, Inc.

Resource Availability

Controller virtualization assures high availability for these centralized, system-wide elements of the wireless LAN. But it does nothing to increase the reliability of the remainder of the wireless LAN infrastructure that is distributed throughout the premises. Here are some additional provisions to help minimize or eliminate failures in the "ports, paths, and power" at the edge of the network:

- Deploying access points within close enough range of one another so that if any one of them fails, the surrounding ones automatically adjust their RF signal strength to compensate for the void (see Predictive RF Modeling in the Signal Integrity section above)
- Balancing the load among access points to minimize the number of users affected by a failure and the subsequent failover, when the access point will either shed its load or need to acquire more users
- If the network is fully aware of a user's location, active sessions can be directed to roam to a suitable nearby access point during routine maintenance.
- Use of access points capable of detecting changes that affect attenuation (the movement of furniture, for example) and automatically recalibrating themselves and/or altering their transmit power levels to adjust for these changes
- Employing distributed security enforcement to ensure that roaming occurs fast enough to maintain continuity of sessions and security associations as users move from one access point to another, especially for VoWLAN calls when even a momentary disruption may seem like a network failure to the user (More details are available in the next section on Seamless Mobility.)
- Creating a mesh topology wherever wireless backhaul is needed, which is normally the case outdoors where the access points are untethered
- Utilizing resilient Layer 2 switching protocols, such as Spanning Tree, Rapid Spanning Tree, and IEEE 802.3ad/ EtherChannel, to ensure reliable connectivity in the wired Ethernet infrastructure

Maintaining Reliable Quality of Service (QoS)

In an always-on wireless LAN, the network must continue functioning well across potentially wildly fluctuating bandwidth demands caused by the many different users and applications. Although such performance issues do not constitute an availability problem per se, they can be (and are) perceived by users as network failures. For example, a minor and temporary increase in congestion could cause latency and jitter to increase enough to disrupt VoWLAN calls. The six measures listed here will all help improve reliability by minimizing the periods when users experience performance, or connection problems, or both:

- Bandwidth allocation—This provision limits the actual bandwidth a user, user group, or particular service set identifier (SSID) may consume, and is usually expressed as a percentage or in Kbps.
- Band steering—This provision moves as many laptops as possible to the 5 GHz band to better balance the load between the two radios in Multiple Input/Multiple Output (MIMO) access points. The effect is to leave sufficient availability for devices operating exclusively in the 2.4 GHz band, which is the case for most tablets and VoWLAN phones that use the lower frequency to conserve power.
- Dynamic call access control—Dynamic CAC recognizes the always-on nature of the many VoWLAN-enabled devices, and only considers those with current VoIP sessions in its call count, thereby enabling more actual calls to go through.
- Dynamic authorization (RFC 3576)—This flexible capability automatically corrects and prevents potential network abuse and meltdowns. For example, by tracking the cumulative bandwidth utilization of guests, any exceeding a certain traffic threshold during critical business hours can be throttled back to a lower allocation.
- IEEE 802.11e media access control (MAC) protocol, Wi-Fi Multimedia (WMM), and Traffic Specification (TSPEC)—These QoS functions are particularly valuable for voice traffic, but are still applicable to other applications. WMM provides negotiated priority access, while TSPEC provides call access control by preventing new calls during peak demand in order to prevent wireless bandwidth from being over allocated.
- Unscheduled Automatic Power Save Delivery (U-APSD) and Wi-Fi Multimedia Power Save (WMM-PS)—These power conservation methods are designed to improve battery performance, which helps prevent a dead or dying battery from causing connectivity problems.

Copyright © 2011, Juniper Networks, Inc.

Seamless Mobility

The inability to roam seamlessly among access points must be considered a network failure in the unwired enterprise. Lost sessions, dropped calls, or the need to reauthenticate at any time are all symptoms of an inability to roam. While detailed treatment of this important topic is beyond the intent and scope of this white paper, at least some coverage is warranted.

Before discussing the specific capabilities required, it is important to note a limitation of the 802.11i standard for "fast roaming." Most wireless LAN vendors support this standard, but it only supports roaming between access points on the same controller and, therefore, does not work across multiple controllers. Because virtual controller clusters function as a single controller, this configuration is capable of supporting fast roaming enterprise-wide.

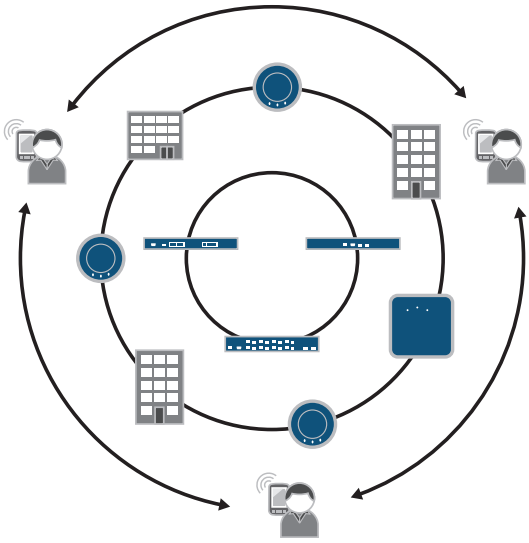

Figure 6: Seamless mobility campus-wide

Centralizing the WLAN controllers is a common practice and even a "best practice" now. But without certain capabilities in the entire wireless LAN infrastructure, the centralized controllers can cause roaming problems. One such capability is the distributed enforcement of security provisions. This can be achieved by propagating security credentials among controllers, thereby enabling each controller's access points to instantly recognize existing authenticated and authorized users who roam into range. Among distributed enforcement's many advantages is its ability to deliver the fast, secure roaming (in less than 100 milliseconds) needed to avoid disrupting VoWLAN calls, even when crossing major network boundaries such as roaming across controllers, or moving from indoors to outside.

A related capability involves the endpoints of encrypted communications. The encryption function provides the privacy enforcement behind authentication and authorization. Without encryption, network security could be compromised because it would be too easy to discover and duplicate another user's identity. Wireless LAN solutions that rely on centralized controllers as the endpoints for encryption and decryption will invariably cause roaming problems. The centralized approach traces its roots to the early days of wireless LANs when the ill-fated Wired Equivalent Privacy (WEP) protocol was broken. With the new, improved Wi-Fi Protected Access (WPA and WPA2) protocols replacing WEP, along with the cryptographic capabilities now built into access points, the distributed encryption model is increasingly preferred based on its superior scalability, security, and reliability.

The traffic forwarding or switching function should also be distributed among the access points. Centralized forwarding creates an inherent bottleneck, which can lead to disruptive packet losses and delays. With distributed forwarding, traffic flows in the most direct path between the source and the destination. The path may traverse the centralized controllers, but only if necessary. Peer-to-peer traffic, including VoWLAN calls, can often bypass the controllers entirely, and proceed instead from access point to access point via the Ethernet switches. The result is fewer sessions timing out and fewer calls being dropped.

Copyright © 2011, Juniper Networks, Inc.

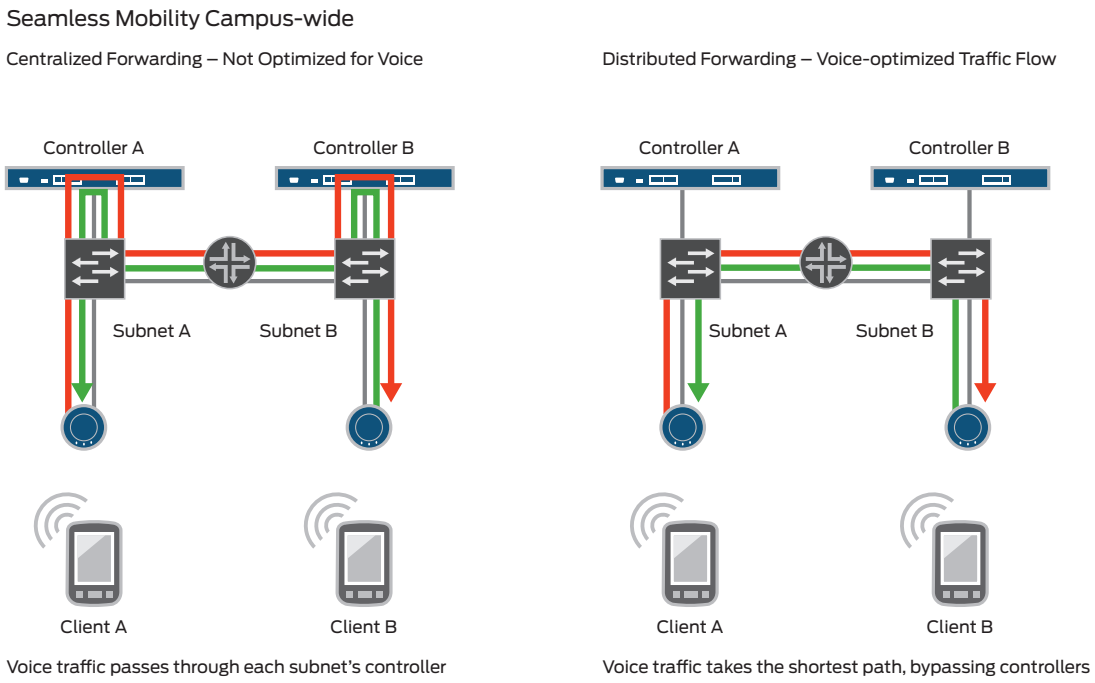

Figure 7: Centralized versus distributed forwarding

Holistic Manageability

All too often, the use of network resiliency provisions in the enterprise, which is essential to achieving always-on wireless connectivity, can unintentionally increase management complexity. The challenge, therefore, is to make the management provisions as simple as they are effective. Such a management model should include:

- A single point of centralized control and a common database for management information
- Access points that can be configured centrally, easily, and in a coordinated fashion, complete with rollback and roll forward capabilities to quickly recover from configuration errors
- In-service moves, adds, and changes that can be made without interrupting network services, potentially during normal business hours
- Awareness of the configuration, operational and capacity status of all controllers for load balancing and in preparation for the possibility of a failover

With controller virtualization, the single point of management is inherent, as all controllers in the cluster are managed as one. This single point of management, known as a seed, can be any designated controller, and a secondary seed can be designated in advance to protect against a primary seed controller failure. The virtual controller cluster then, in turn, treats all access points as a shared resource pool that can be managed centrally.

In the active/standby model, by contrast, it is necessary to manage active and standby controllers separately for either one-to-one or many-to-one configurations. It is also necessary to duplicate the configuration of all active controllers and access points, as well as to update these during moves, adds, and changes. It is often only possible to know the operating behavior of the standby controller when one of the active controllers actually fails. For this reason, the standby controller could also be misconfigured if it is not updated automatically during any changes. As a result, in the event of a failure, session and status information may be lost, requiring users to reauthenticate.

Copyright © 2011, Juniper Networks, Inc.

Unified Wireless LAN Infrastructure and Mobility Services Management

Most mobility services available today are implemented in silos. Silos accelerate the time-to-market, and enable mobility services, such as walled garden guest access, RTLS, VoWLAN, and advanced security options, to be deployed separately as needed, potentially from third-party providers.

But there is a serious problem with implementing mobility services in separate, isolated silos. Because each is installed, configured, and managed separately, the individual services remain completely unaware of one another as they all compete constantly for the same underlying resources. Sometimes the isolation is intentional for legitimate reasons. But normally the isolation is inherent in the fundamental architecture, thereby precluding any possibility of integrated or unified management—now or in the future.

The essential requirements for unifying management of the entire wireless LAN infrastructure and all mobility services are outlined in the following Aberdeen Group best practices for best-in-class enterprises:

- A holistic (versus piecemeal) approach to network architecture, deployment, and ongoing support and management
- Unified network management infrastructure to provide visibility and control over the entire network simultaneously, which reduces redundant control layers and enables support staff to take a proactive approach to maintenance, upgrades, and support
- A consistent policy regarding performance optimization, system upgrades, and maintenance across the entire organization
- Effective enforcement of security and compliance mandates organization-wide

Naturally, such holistic or unified management must leverage the capabilities of the wireless LAN's element management system, as well as the management systems used for the individual mobility services. And in doing so, it must also unite these separate provisions into a coordinated and cohesive set of capabilities for managing the unwired enterprise. For only with such a holistic approach will the all wireless access network be able to deliver reliable connectivity with an acceptable quality of experience for all users.

Conclusion

Because enterprises are becoming so dependent on wireless LANs as the primary access network, and migrating rapidly to an unwired enterprise environment, it bears repeating: Failure is simply no longer an option with today's unwired enterprise. "Failure" can take many forms from a user's perspective, of course, ranging from a VoWLAN call being dropped to a widespread outage. The former can be tolerated occasionally; the latter simply cannot—ever.

For this reason, each element outlined in this white paper is necessary, and only a wireless LAN with all five is sufficient for achieving mission critical reliability on a par with what users have become accustomed to on the wired LAN. Signal integrity is essential to ensuring that users can communicate, without interruption, with access points from any location at any time. Reliability of the centralized controllers is vital, as these systems form the very core of the wireless LAN. It is equally important, however, that the resources at the edge, including the Ethernet switches and wireless access points, deliver consistently high availability to prevent localized outages even under prevailing network outages. Fast seamless roaming is a critical requirement if users are to avoid disruption of individual sessions and calls as they move about office buildings and the enterprise campus. Finally, the management provisions must be holistic to have any chance of enforcing minimal service-level agreements in the face of competing demands.

The proliferation of mobile devices makes the unwired enterprise inevitable. While it may take some organizations some time to cut the cord completely, today's WLAN, however pervasive, will form the foundation for this promising future of anywhere, anytime access. The sooner the wireless LAN can be made as mission critical as the applications it supports, the sooner that day will arrive.

For more information about how your organization can benefit from a nonstop wireless LAN, visit Juniper Networks on the Web at **www.juniper.net/us/en/products-services/wireless**.

Copyright © 2011, Juniper Networks, Inc.

About Juniper Networks

Juniper Networks is in the business of network innovation. From devices to data centers, from consumers to cloud providers, Juniper Networks delivers the software, silicon and systems that transform the experience and economics of networking. The company serves customers and partners worldwide. Additional information can be found at **www.juniper.net**.

Corporate and Sales Headquarters
Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or 408.745.2000
Fax: 408.745.2100
www.juniper.net

APAC Headquarters
Juniper Networks (Hong Kong)
26/F, Cityplaza One
1111 King's Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803

EMEA Headquarters
Juniper Networks Ireland
Airside Business Park
Swords, County Dublin, Ireland
Phone: 35.31.8903.600
EMEA Sales: 00800.4586.4737
Fax: 35.31.8903.601

To purchase Juniper Networks solutions, please contact your Juniper Networks representative at 1-866-298-6428 or authorized reseller.

Copyright 2011 Juniper Networks, Inc. All rights reserved. Juniper Networks, the Juniper Networks logo, Junos, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

2000414-001-EN June 2011 Printed on recycled paper

Copyright © 2011, Juniper Networks, Inc.