

White Paper |



The Next Step in the Evolution of Wireless Mesh Networking

August 2010

ARUBA[®]
networks

Table of Contents

The Next Step in the Evolution of Wireless Mesh Networking

Meeting today's wireless mesh network requirements	2
The modern wireless mesh network is here	3
Solving scalability challenges with AWR	3
Enforcing quality of service for wireless	8
Authentication and encryption	9
MobileMatrix high-speed roaming	9
HQ-quality video with Active Video Transport.....	11
Conclusion	13

1 Meeting today's wireless mesh network requirements

Wireless mesh networks allows outdoor environments to be interconnected without any wires and with the security and reliability of a wired network. They solve a wide range of communications challenges across different outdoor environments, making them well suited for public safety, emergency response, oil rigs, video surveillance, large-scale events and transportation hubs.

An all-wireless network mesh brings the convenience of easy installation and lower deployment costs. Wireless mesh networks must meet the same standards for scalable capacity, reliability and security as do their wireless LAN (WLAN) counterparts.

Issues with throughput, quality and security in a wireless mesh network have largely been resolved, but scalable capacity remains an obstacle with some vendor solutions. Many wireless mesh solutions simply cannot scale without compromising performance, quality of service (QoS) or availability across multiple hops in a mesh infrastructure. And in always-on, mission-critical communications environments, that's simply not acceptable.

The scalable capacity challenge is rooted in the nature of wireless networking. The inherent inefficiencies of sharing the radio frequency (RF) spectrum are a primary contributor to scalability issues. A second cause is Layer 2 switching or bridging. The design of the link-layer wireless protocols used in switching or bridging have adverse consequences for wireless mesh scalability, flexibility and performance.

Wired Ethernet has resolved the scalable capacity challenge. First, significantly greater bandwidth is available with wired networks than with wireless, so the impact is smaller. And second, wired Ethernet uses IP routing.

The shared nature of RF means that the available bandwidth on a wireless network will be less than on a wired network, but the scalable capacity challenge can be met by using true network-layer routing on the wireless mesh network.

Until now, true network-layer routing has been considered too complex and costly to be practical. But Aruba Networks has made a significant investment in its adaptive Layer 3 wireless mesh routing protocol to advance the state-of-the-art networking in three significant ways:

- Network-layer routing. Aruba provides efficient network-layer routing that is designed specifically for the wireless mesh. With network-layer routing, wireless mesh networks can deliver the scalability, throughput and low-latency across multiple mesh hops and over large geographic areas to meet the demands of delay-sensitive applications.
- High-speed mobility. Aruba's high-speed roaming capabilities – both within a single IP domain and across multiple domains – integrate IP routing with wireless link-level access to support seamless roaming.
- Support for high-quality voice and video. Aruba's innovative traffic-shaping technology ensures the delivery of high-definition video and high-quality voice by enforcing QoS and bandwidth management.

Aruba's innovation, combined with advances in multi-radio backhaul, quality of service, security and manageability, give Aruba's wireless mesh network solution an unprecedented level of network intelligence along with a compelling total cost of ownership (TCO).



Figure 1: Aruba supports the convergence of voice, video and data applications onto an intelligent and scalable wireless mesh network that delivers superior performance, seamless roaming, enhanced quality and carrier-grade security and reliability.

2 The modern wireless mesh network is here

Early IEEE 802.11 wireless mesh networks encountered scalability problems caused by link-level protocols, such as carrier-sense multiple-access with collision avoidance (CSMA/CA). Wired Ethernet uses carrier-sense multiple-access with collision detection (CSMA/CD), and the ability to detect Ethernet collisions – which cannot be done with RF signals – makes the protocol more efficient.

In addition, Ethernet has several orders of magnitude more bandwidth to solve this challenge. But in an all-wireless environment, there is far less bandwidth than a wired network and the CSMA/CA protocol imposes capacity limitations, especially in single-radio access points (APs).

Subsequent wireless mesh solutions used separate radios for access and backhaul to mitigate the effects. Some products use multiple radios for backhaul and directional antennas to minimize self-interference, which can dramatically increase the number of collisions in large networks. Nevertheless, CSMA/CA still has a detrimental effect on throughput, and that impact accumulates with each mesh hop and ultimately restricts the available bandwidth across the network.

Even where throughput degradation has been minimized, Layer 2 switching – with its flat topology and reliance upon the spanning-tree protocol – still places severe limitations on wireless mesh network scalability. Switched Ethernet networks also suffer from this limitation, despite the advent of rapid and multiple spanning trees.

The only way to eliminate these final limitations to wireless mesh scalability is to use end-to-end network-layer wireless routing. Routing is what makes the Internet scalable and resilient, and Aruba's Adaptive Wireless Routing™ (AWR™) protocol brings the same network intelligence to an all-wireless network infrastructure.

3 Solving scalability challenges with AWR

To overcome the scalable capacity and performance limitations inherent in a large, flat link-level network, wireless mesh vendors typically use a wired infrastructure to support many APs.

In an all-wireless network, however, inter-node communications are also wireless – which is why most wireless mesh vendors claim to use some form of “routing” among all mesh nodes. But such “routing” is really nothing more than an extension of the spanning-tree protocol designed for Layer 2 networks.

Most wireless mesh vendors do not use Layer 3 routing because of its high cost and complexity. Regularly exchanging route information can consume precious wireless bandwidth. In large networks, routing tables can grow correspondingly large, requiring a considerable amount of memory in the routers.

Constantly updating routes and making real-time packet forwarding decisions demands substantial computational resources. Because network-layer routing can undermine the price/performance of a vendor's product, most vendors choose instead to "route" at the link level.

Aruba combines the efficiency of distance-vector routing with significant enhancements for wireless link-state awareness to create the industry's first network routing solution purpose-built for scaling wireless mesh networks. AWR is true, network-layer routing that is designed for link-level LANs, both wireless and wired.




Figure 2: Aruba's Adaptive Wireless Routing overcomes the limitations imposed by Layer 2 networks to make all-wireless network infrastructures more intelligent, scalable, flexible and resilient.

As a dynamic, distributed routing protocol purpose-built for wireless mesh networks, AWR forwards packets at near nominal wireless data rates across multiple hops with a per-hop latency of less than two milliseconds. AWR is robust and remarkably efficient, which allows it to be implemented cost-effectively on Aruba's multi-radio wireless mesh routers.

AWR uses the proven distance-vector routing protocol (DVRP). DVRP is well suited for wireless mesh networks because it has low message overhead and low computational complexity. With DVRP, all routers periodically exchange route information with neighboring routers via user datagram protocol (UDP) messages. Any change in a neighboring route is reflected in a router's own route table, which is communicated to its neighbors.

Aruba employs three enhancements to optimize DVRP for wireless mesh networks:

- **Awareness of link-state quality.** Incorporating link-state awareness allows Aruba to create a quality metric that is specific to wireless communications. Indicators for assessing current link quality include the link's data rate, received signal-strength Indication (RSSI) and external interference.

AWR makes all packet forwarding decisions based on a combination of the link-state quality and DVRP parameters. The result is superior resiliency and throughput performance with optimal utilization of all available paths.

- **Mobility with seamless roaming.** AWR supports wireless mobility with seamless roaming capabilities within a single IP domain and across multiple domains.

By incorporating information about current client connections in all route tables, cross-IP subnet roaming facilitates direct and rapid handoff of RF communications from one router to another across the wireless mesh. This design yields a more efficient network traffic flow and makes it possible to support applications that require seamless, real-time roaming.

- **Routing around link failures.** A local-repair mechanism prevents loops from forming in route tables. While these loops normally do not occur when a router fails, they can occur when the physical link fails between adjacent routers.

While link failures are rare in wired networks, wireless mesh networks are susceptible to sources of RF interference that can cause a link to degrade substantially or fail altogether. When a link failure occurs, the affected routers exchange information to discover an alternate path.




Figure 3: Routing is what makes the Internet so scalable and resilient, and Adaptive Wireless Routing brings the same network intelligence to wireless mesh. Any router can immediately route around any link failure.

Aruba's enhancements make DVRP effective and efficient in a wireless mesh network. Peak performance is maintained across geographically dispersed networks because of updates to routing tables. Intelligent, dynamic routing eliminates the need for manual intervention to optimize overall traffic flows.

AWR makes it possible to balance the load among all available gateways to the Internet or other external networks. In fact, the entire Aruba wireless mesh constantly balances the total load to optimize traffic flow, even under adverse conditions with high levels of RF interference.

AWR supports multicasting to provide more efficient bandwidth utilization for applications such as video, which broadcast to multiple destinations. By making all packet-forwarding decisions in a cross-layer fashion, AWR can deliver superior service quality to all applications.

Wireless mesh networks from Aruba also support open shortest-path first (OSPF) routing. OSPF is a hierarchical, interior gateway protocol that uses link state to create an optimal hierarchy of paths among routers in a network.

The OSPF autonomous system boundary router (ASBR) function is utilized so that external routers can include the wireless mesh in their route tables. Backbone and non-backbone areas in OSPF are supported

OSPF periodically publishes link-state advertisements (LSAs), which are required by external gateway routers. Support for these OSPF features enables network operators to optimize routing across multiple autonomous systems – wired and wireless – in any single network domain.

AWR is compatible with Wi-Fi standards at both the data-link and physical layers, even where Aruba has made additional enhancements to further optimize routing in wireless mesh network infrastructures.



Figure 4: Aruba's intelligent network routing adds tremendous flexibility to the wireless mesh and eases the integration with other networks, such as the Internet, an enterprise network or a service provider's infrastructure.

AWR overcomes the internetworking limitations imposed by spanning tree on a purely Layer 2 network. To overcome the performance issues caused by access contention, Aruba supports multiple independent, software-configurable radios in a single router.

Each can operate in either AP mode or backhaul mode in the unlicensed 2.4-GHz or 5-GHz bands or in the licensed 4.9-GHz public safety band. AWR's network intelligence takes full advantage of this powerful multi-radio, multi-path topology in its route creation and packet-forwarding decisions.

Aruba made several enhancements to the basic IEEE 802.11a/g standards that optimize bandwidth utilization at the link level. Based on Super A/G, which supports 90 Mbps maximum theoretical throughput, these enhancements add a special data rate control capability to take full advantage of static channel binding, fast-frame formation and hardware compression features available in the wireless chipset.

The Super A/G enhancements are:

- **Data rate control.** Data rate control establishes an intelligent tradeoff between transmission speed and the error rate caused by factors like interference and low signal-to-noise ratios.

On some wireless links, a higher data rate often experiences a disproportionately higher error rate, and the resulting TCP retransmissions effectively reduce the overall throughput. Aruba's data-rate control feature takes into account acceptable levels of packet loss for UDP applications, and can be applied to different services, such as video, voice and Internet access.

- **Static channel binding.** Static channel binding enables two standard 20-MHz channels to be bound together to operate effectively as a single channel with twice the throughput capacity.
- This capability is analogous the inverse multiplexing technology used to bind multiple T1/E1 lines or the link aggregation groups (LAGs) used to bind multiple Ethernet switch ports into a unified connection that delivers higher bandwidth.
- **Fast frame formation.** Fast frame formation overcomes the inefficiency caused by the mismatch in frame sizes between wired Ethernet (1,500 bytes) and wireless (4,096 bytes). Fast frame formation eliminates overhead by combining contiguous Ethernet frames into a single wireless packet with a common header.
- **Hardware compression.** To boost application performance, Aruba uses an efficient, lossless compression algorithm to reduce the size of packets with redundant payload content. The extent of the performance gain from hardware compression and fast frame formation depends on the content and size of the packet payload.

Because Super A/G enhancements operate only in the backhaul links across the wireless mesh, the AP function in Aruba wireless mesh routers is fully compliant with the 802.11 a/b/g standards and interoperates with all Wi-Fi clients.

Layer 1, the physical layer, is where the AWR protocol takes into account the RF characteristics that affect link quality, including RSSI and external interference. To optimize these characteristics and maximize backhaul throughput and transmission range in outdoor deployments, Aruba supports omni-directional and directional antennas.

Directional antennas are preferred for backhaul because of their ability to minimize self-interference. By contrast, the AP functionality in wireless mesh router would normally use an omni-directional antenna, although directional coverage may be beneficial in some situations.

To further optimize performance at Layer 1, Aruba implements two additional enhancements:

- **Radio frequency management.** Aruba's Radio Frequency Management (RFM) module automatically scans all frequencies and channels to discover all neighboring routers within range. The Aruba RFM module constantly assesses link quality and monitors link status. AWR uses this information to optimize route creation and make packet forwarding decisions.
- **Automated interference detection and avoidance.** Aruba's Automated Interference Detection and Avoidance (AIDA) module determines when it is advantageous to change channels for any AP or backhaul radio that experiences high levels of RF interference. The Aruba AIDA module monitors all channels, including unused channels, in order to perform this service.

The combined effect of these wireless mesh network enhancements – all fully integrated in Layers 1-3 – is industry-leading intelligence and scalable performance. Aruba delivers significantly higher throughput performance compared to multi-radio solutions that rely solely on Layer 2 switching at the link layer. Just as important, the increased throughput is sustained independently from source to destination across multiple hops in the wireless mesh network.

3.1 Enforcing quality of service for wireless

Aruba's Adaptive Wireless Routing and other enhancements deliver the high-throughput, low-latency performance that is fundamental to providing exceptional QoS in mesh networks that support voice, video and data.

Provisions for path and packet forwarding optimization, load balancing, interference avoidance, and multi-radio backhaul with no throughput degradation across multiple hops all help to sustain peak levels of performance.

To deliver satisfactory QoS requires another layer of intelligence – one that optimizes the total throughput performance among all applications. Aruba provides this intelligence at Layer 3 and Layer 2 by integrating support for differentiated services (DiffServ), IEEE 802.11e and IEEE 802.1Q virtual LANs (VLANs).

- **DiffServ.** DiffServ is the preferred way to control and enforce QoS in a routed network. It provides a way to classify and manage network traffic to enhance QoS for mission-critical applications, while still maintaining acceptable levels of QoS for all other applications.

Aruba's implementation takes full advantage of DiffServ's ability to create a hierarchy of categories that enables network operators to minimize latency and guarantee throughput for specific applications by provisioning a preferential allocation of bandwidth. Other less-critical applications, like email and file transfers, are then allocated bandwidth on an as-available, best-effort basis.

- **802.11e.** 802.11e defines QoS enhancements specifically for Layer 2 WLANs. In addition to 802.11e, Aruba also supports Wi-Fi Multimedia (WMM) and Wireless Multimedia Extensions (WME) established by the Wi-Fi Alliance. Aruba's implementation allocates a separate queue for each of four flow categories -- voice, video, best-effort and background.

Multiple queues optimize the allocation of bandwidth based on an organization's established service level policies. Voice is usually assigned the highest priority because of its extraordinary sensitivity to latency. With Aruba, the priorities established by DiffServ at Layer 3 can be mapped into these Layer 2 flow categories, which create greater flexibility for the network operator.

- **VLANs.** 802.1Q allows network operators to create logical partitions, or VLANs, on a Layer 2 network. Segmentation creates an additional layer of security and can also help manage QoS. Using VLANs to isolate applications is particularly important for Wi-Fi-enabled devices that do not support 802.11e.

Aruba's implementation allows a VLAN ID to be associated with a WLAN service set identifier (SSID). The association extends VLAN QoS priorities to users and applications assigned to different SSIDs. The VLAN typically gets its priority treatment by mapping its ID to a designated differentiated-services code point (DSCP) at Layer 3.




Figure 5: Aruba employs multiple methods of enforcing QoS – including DiffServ at Layer 3, 802.11e and VLANs – to support the use of voice, video and data on a wireless mesh.

3.2 Authentication and encryption

Aruba's wireless mesh solution provides strong security beginning with the connection between clients and APs and extending across the backhaul from source to destination. Security provisions are also top-to-bottom – from user access control at the application layer to encryption of traffic during transmission at the physical layer. These same provisions secure all communications required to manage the wireless mesh network infrastructure and to exchange routing information among all routers.

Aruba's network-layer intelligence affords superior compatibility with other network security measures. These include network address translation (NAT), network access control (NAC) and IPsec virtual private networks (VPNs) as well as separate access control lists (ACLs) or authentication, authorization and accounting (AAA) systems, such as RADIUS.

For basic NAC, Aruba supports IEEE 802.1X authentication, which is based on the extensible authentication protocol (EAP). Client devices must first authenticate with the wireless mesh router AP before being granted access to the network. In addition to RADIUS, Aruba supports extensions to EAP, including X.509 digital certificates, EAP-subscriber identity module (SIM), EAP-transport layer security (TLS) and EAP-tunneled transport layer security (TTLS).

ACLs provide additional controlling and filtering of traffic on the Aruba mesh based on an authenticated user's source IP or MAC address and the application's destination address. Additional filtering is available at the MAC-address level to ensure that only valid MAC addresses gain access to the network.

To secure communications between clients and APs, Aruba supports Wi-Fi protected access (WPA and WPA2), which uses the temporal key integrity protocol (TKIP) to change keys dynamically at regular intervals.

Aruba also supports the advanced encryption standard (AES) to encrypt both AP and backhaul traffic in a wireless mesh router. AES affords considerably stronger encryption with larger key sizes of 128 bits, 192 bits and 256 bits. WPA2 supports AES for the stronger encryption.

VLANs can be used to isolate traffic from different groups of users or applications as an added security measure. With Aruba, an SSID can be associated with a VLAN ID, and each VLAN can have its own security policy pertaining to access control, authentication and encryption.

Aruba allows some or all SSIDs to be hidden to prevent their detection by the auto-scanning feature in most Wi-Fi clients. In this instance, eligible users must provide a valid SSID separately in order to gain network access.

4 MobileMatrix high-speed roaming

Aruba's MobileMatrix™ provides the ability to roam seamlessly, potentially at very high speeds, throughout the Aruba wireless mesh infrastructure. The integration of the network and link layers in AWR provides the cross-IP subnet roaming capability needed to allow clients to move from wireless mesh router AP to AP in less than 50 milliseconds while maintain session persistence and keeping the same IP address. Fast roaming maintains a continuous application connection, which is critical for latency-sensitive applications like voice and video.

The IEEE 802.11 standard does not specify a robust, interoperable mechanism for roaming. In fact, it requires each client have only a single connection. There is also no provision for an AP to identify clients within its range, which places the burden on the clients to detect available connections and initiate a roaming request.

The inter-AP protocol (IAPP) specified in IEEE 802.11f provides a means for nomadic movement between APs. But the delay involved is generally longer than required to support real-time voice and video applications. Plus, IAPP is problematic for connections that use WEP, WPA or WPA2 security. Without any enhancements, IAPP is suitable only for data applications that are insensitive to latency and require no security.

Roaming in Wi-Fi networks is possible at Layer 3 with the Mobile IP standard described in IETF RFC3344. Optional for IPv4 and required for IPv6, Mobile IP enables packets to be forwarded in tunnels from a system with a fixed IP address

to mobile devices that roam among multiple networks. These mobile devices can roam across multiple subnets, where it becomes necessary to assign a new and different IP address.

Significantly, Mobile IP was designed to be transparent to both the mobile node itself and the correspondent node at the remote end, which may be either mobile or stationary. But Mobile IP is very complicated and is not widely adopted.

Aruba's MobileMatrix leverages the capabilities of IAPP and adopts a simplified version of Mobile IP to make roaming fast and seamless without high overhead. MobileMatrix maintains full interoperability with ordinary Wi-Fi clients. It does not require any special software on servers, clients or internetworking systems external to the wireless mesh.

MobileMatrix roaming begins by inheriting the standard MAC layer trigger mechanism initiated by the Wi-Fi client. But IAPP only supports roaming within a single IP subnet, which is problematic for IP applications. In fact, if a client roams with IAPP to an AP on a different IP subnet, its IP address can no longer be used in a current session. Consequently, it will require a new IP address and re-initiate a new session.

To support cross-IP subnet roaming, MobileMatrix uses a special gateway function, which is analogous to Mobile IP's home agent, to recognize that a client IP address is now using a different AP. MobileMatrix immediately updates the routing tables in AWR to route packets via the new AP.

While MobileMatrix and Mobile IP take a similar approach to roaming, the important difference is speed. The ability to complete the transition in less than 50 milliseconds gives MobileMatrix seamless session persistence for virtually any IP application, including voice.

The 50-millisecond transition period is from beginning to end for both the client and the wireless mesh network. Wi-Fi clients constantly scan for available wireless mesh APs. When a stronger signal is detected, the client can initiate an IAPP roaming request.

In fewer than 50 milliseconds, MobileMatrix recognizes the request and initiates an update to the route tables in AWR, and then propagates the changes to affected routers in the wireless mesh while the client simultaneously re-associates with the new wireless mesh AP.

MobileMatrix uses four methods to achieve fast, cross-IP subnet roaming:

- *The Global MobileMatrix Service* process maintains all mobile user information required by the other MobileMatrix processes.
- The *Access Point MobileMatrix* is responsible for initiating and completing the roaming requests on behalf of clients and within the wireless mesh infrastructure. It is equivalent to the Mobile IP foreign agent, and runs on all routers.
- Any router that serves as a gateway to an external network uses two additional processes that are equivalent to the Mobile IP home agent. The *Local MobileMatrix Service* maintains the mobile user information required by the local gateway. It is an extension of the *Global MobileMatrix Service*.
- The *MobileMatrix Traffic Gateway* is the companion process in gateway routers. It is responsible for establishing the route to the mobile client's current AP, which includes advertising any new route to the AWR protocol.



Figure 6: With MobileMatrix, users are free to move about the entire wireless mesh without reinitiating active sessions because the seamless transition from one AP to another is completed in less than 50 milliseconds.

5 HQ-quality video with Active Video Transport

Aruba's Active Video Transport™ (AVT™) traffic-shaping system delivers progressive, non-interlaced HD-quality video at up to 30 frames per second. With AVT, users perceive a significant improvement in quality, while behind the scenes AVT makes intelligent tradeoffs between latency and impairments to video quality.

Fully appreciating AVT's ability to improve video performance requires understanding the causes and effects of the most common impairments to video quality – packet loss, packet reordering and packet jitter.

Voice and video traffic is normally transmitted in wired and wireless networks using the connectionless UDP protocol. Real-time applications like voice normally cannot benefit from the retransmission feature of the connection-oriented TCP protocol. When packets are lost or corrupted in a UDP data stream, they are simply lost and are never recovered.

Uncompressed video signals are tolerant of moderate packet loss, but any amount of packet loss for compressed video signals becomes noticeable – often in annoying ways. And because wireless mesh networks have limited bandwidth, some form of compression is typically used.

Traffic loss in wireless mesh networks can be significant due to periodic congestion. Dropped packets or transmission errors that corrupt packets can occur due to a link data rate being too high, external RF noise or interference, antenna misalignment, moving obstacles, multi-path fading, user mobility, or a low or variable RSSI.

Consequently, packets often arrive at their destination in a different order than they are sent from the source. TCP reorders these packets into their original sequence, but UDP does not. With UDP applications, packets are consumed in the order they are received.

With compressed digital video, the effect of packet reordering can be worse than the equivalent amount of packet loss because an out-of-sequence packet disrupts the decoding process.

For this reason, video equipment is often designed to simply drop the packet, or with high-end systems, build in some delay in the decoder to create a brief window of opportunity for reordering out-of-sequence packets.

The amount of packet reordering in a routed network is normally poses less of a problem than packet loss, and is usually caused by a change in the end-to-end route during the session. Routes generally change for two reasons – link failures that require rerouting or link congestion that triggers load balancing along new paths.

Without some provision in the video decoder, jitter causes quality to degrade with noticeable pixilation or blurred images. The same delay built into sophisticated video equipment to create a window of opportunity for packet reordering also facilitates the removal of jitter from the incoming packet stream.

The sources of packet jitter in a wireless mesh network include variations in delay at the source, variable link data rates along the path, changing traffic conditions in QoS queues, changes in end-to-end routes, the non-deterministic effects of the CSMA/CA protocol, and roaming.

Compression algorithms commonly used in digital video applications are stateful. In stateful communications, the arriving bit stream is used to make changes in the existing image rather than construct a new image during each frame interval.

Stateful compression algorithms have the advantage of being highly efficient, which is desirable in a wireless mesh network. However, they have the disadvantage of not being tolerant of packet loss, and often require a disruptive resynchronization between the encoder and decoder when packet loss, reordering and jitter become severe enough.

As stated earlier, Aruba's AVT traffic-shaping system delivers high-definition video by making an intelligent tradeoff between latency and the impairments to video quality. The increased latency required to compensate for packet loss, reordering and jitter is imperceptible to users. What is perceptible is the significant improvement in video quality.

AVT uses four technologies to deliver cinema-quality 30-frame progressive video across the wireless mesh network – deep packet inspection, MAC protocol optimization, an in-network retransmission protocol, and adaptive video jitter removal.

Deep packet inspection identifies and extracts the compression algorithm, video decoding buffer model, video frame-type boundary and video timing used by the packet stream. This information is needed to properly share the traffic through the Aruba wireless mesh network.

The MAC protocol optimization and in-network retransmission protocol work together to minimize and, when necessary, recover from packet loss, respectively. This combined prevent-and-recover approach is especially effective in noisy RF environments where packet loss can occur most often. Finally, the adaptive video jitter removal technology reorders any out-of-sequence packets.

AVT has ingress and egress functions. Video ingress and egress are supported for both gateways and wireless mesh router APs. A single Aruba wireless mesh router provides ingress and egress for different video streams traveling in opposite directions. Because AVT's traffic shaping must work across the entire infrastructure, this technology is implemented in every Aruba wireless mesh router.

The AVT ingress function is responsible for deep packet inspection to identify and characterize video frames. The ingress function applies the special frame format used by the traffic shaping processes and interleaving batches of frames during transmission.

The AVT egress function is responsible for adaptive video jitter removal. It issues in-network retransmission requests when lost and corrupted frames are detected, reorders out-of-sequence frames, and releases the video frames at a consistent jitter-free rate at the end of the playback deadline.

AVT leverages the multicasting capability in AWR to provide concurrent and efficient multi-path transmission of HD-quality video to multiple destinations. Multicasting is especially useful for video surveillance applications that require monitoring and recording at multiple locations or for IPTV applications that broadcast video to multiple viewers.

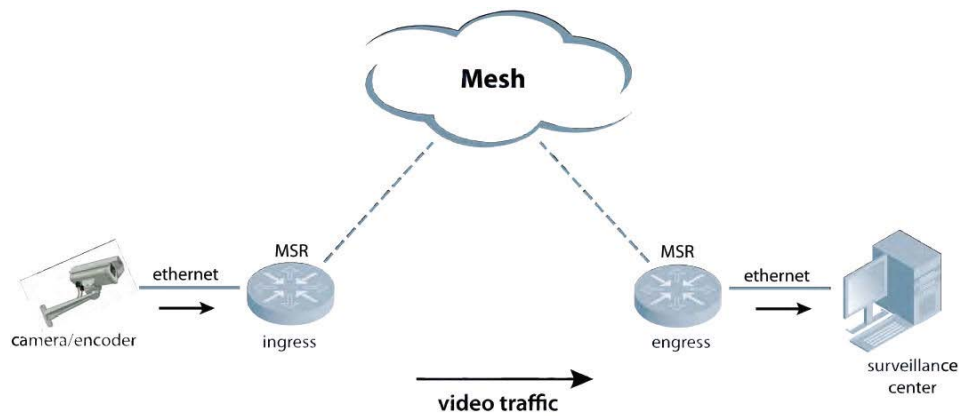


Figure 7: Aruba's Active Video Transport system dramatically improves video quality by shaping the video data stream.

6 Conclusion

Aruba wireless mesh networks bring many advantages to challenging outdoor communications environments that have traditionally been difficult to cover. Quick and easy to install over a wide geographic area while offering the flexibility to be deployed in any outdoor location, wireless mesh networks are an ideal communications solution for public safety, construction, transportation, mining and other industrial applications.

With Aruba's innovative solution for wireless mesh networks, scalable capacity is now achievable. Businesses can deploy a wireless mesh communication infrastructure that supports latency-sensitive video, voice and data applications with the confidence that they can continue to meet requirements for scalability, reliability, security and operational efficiency.

About Aruba Networks

Aruba is the global leader in distributed enterprise networks. Its award-winning portfolio of campus, branch/teleworker, and mobile solutions simplify operations and secure access to all corporate applications and services - regardless of the user's device, location, or network. This dramatically improves productivity and lowers capital and operational costs.

Listed on the NASDAQ and Russell 2000® Index, Aruba is based in Sunnyvale, California, and has operations throughout the Americas, Europe, Middle East, and Asia Pacific regions. To learn more, visit Aruba at <http://www.arubanetworks.com>. For real-time news updates follow Aruba on [Twitter](#), [Facebook](#), or the [Green Island News Blog](#).



1344 Crossman Ave. Sunnyvale, CA 94089-1113
Tel. 408.227.4500 | Fax. 408.227.4550 | 1-866-55-ARUBA
info@arubanetworks.com | <http://www.arubanetworks.com>