

Feature Overview and Configuration Guide

Autonomous Management Framework™ Plus (AMF Plus)

AMF PLUS



AlliedWare Plus™
OPERATING SYSTEM

AlliedTelesis.com

C613-22136-00-REV C

Acknowledgments

This product includes software developed by the University of California, Berkeley and its contributors. Copyright ©1982, 1986, 1990, 1991, 1993 The Regents of the University of California. All rights reserved.

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. For information about this see www.openssl.org/.
Copyright (c) 1998-2019 The OpenSSL Project
Copyright (c) 1995-1998 Eric A. Young, Tim J. Hudson
All rights reserved.

For the full list of acknowledgments, and respective copyright notices, run the **show version** command on your device.

This product includes software licensed under the under v2 and v3 of the GNU General Public License, available from: www.gnu.org/licenses/gpl2.html and www.gnu.org/licenses/gpl.html respectively.

Source code for all GPL licensed software in this product can be obtained from the Allied Telesis GPL Code Download Center at: www.alliedtelesis.com/support/default.aspx

Allied Telesis is committed to meeting the requirements of the open source licenses including the GNU General Public License (GPL) and will make all required source code available.

If you would like a copy of the GPL source code contained in Allied Telesis products, please send us a request by registered mail including a check for US\$15 to cover production and shipping costs and a CD with the GPL code will be mailed to you.

GPL Code Request
Allied Telesis Labs (Ltd)
PO Box 8011
Christchurch
New Zealand

Allied Telesis, Inc. reserves the right to make changes in specifications and other information contained in this document without prior written notice. The information provided herein is subject to change without notice. In no event shall Allied Telesis, Inc. be liable for any incidental, special, indirect, or consequential damages whatsoever, including but not limited to lost profits, arising out of or related to this manual or the information contained herein, even if Allied Telesis, Inc. has been advised of, known, or should have known, the possibility of such damages.

Allied Telesis, AlliedWare Plus, Allied Telesis Management Framework, EPSRing, SwitchBlade, VCStack and VCStack Plus are trademarks or registered trademarks in the United States and elsewhere of Allied Telesis, Inc. Additional brands, names and products mentioned herein may be trademarks of their respective companies.

©2023 Allied Telesis Inc. All rights reserved. No part of this publication may be reproduced without prior written permission from Allied Telesis, Inc.

Contents

Acknowledgments.....	2
Contents	3
Introduction to AMF and AMF Plus	8
The difference between AMF and AMF Plus.....	9
Products and software versions that apply to this guide.....	10
What is an AMF Plus network?	12
Key benefits of AMF Plus	13
Elements of AMF Plus	14
Applications that use AMF Plus	21
Vista Manager EX™.....	21
AMF Security Controller and AMF Application Proxy	23
AMF Plus Subscription Licenses.....	24
AMF Plus starter license.....	24
Managing AMF Plus licenses	24
Automatically obtaining and activating licenses	27
Combining AMF and AMF Plus licensing.....	30
The Concept of AMF Plus Areas	32
Configuring an AMF Plus controller	32
Connections from AMF Plus controllers to the other areas	33
Example - Configuring a multi-area network.....	35
Area links on AR-series Eth ports.....	38
Areas with 120-300 nodes	38
Controllers with 301-1000 areas	39
Example AMF Plus Configuration	40
Configuring a simple stand-alone area	40
User account management.....	45
NTP and AMF Plus	46

Special Considerations when Using LACP Aggregations as AMF Plus Links	48
Sharing AMF Plus links with other network operations.....	49
Reserved IP address range	50
AMF Plus on VCStacks	50
AMF Plus links on AR-series Eth interfaces	51
AMF Plus interaction with QoS and ACLs.....	51
AMF Plus interaction with STP on AR-series devices.....	52
Renaming your AMF Plus network.....	52
Using the AMF Plus Menu in Vista Manager	53
Dynamic Connection.....	55
Internet Breakout.....	61
Auto Traffic Shaping	64
Application Priority	66
Security	71
Health Monitoring.....	75
Networks	86
Smart ACL.....	87
Intent-based QoS.....	94
Configuring AMF Plus Nodes: the Unified CLI.....	130
Working-sets	130
Local working-set.....	131
Creating a working-set.....	131
Working-set groups.....	131
Executing commands on working-sets.....	133
Interactive commands.....	136
Copying files between nodes	136
Node Provisioning.....	137
When to use node provisioning.....	137
Provisioning multiple device-types on the same node.....	138
Creating a new provisioned node.....	138
Configuring adjacent nodes	139
Connecting a provisioned node to an AMF Plus network.....	142

Node Recovery	144
Automatic node recovery	144
Replacing a device with a similar model	147
Restoring a node to a “clean” state	148
Recommended procedure when replacing a device using automatic node recovery	149
Recovering devices that have subscription licenses	150
Recovering devices that have release licenses.....	151
Fixing a failed node recovery: AMF Plus safe configuration	152
Recovering a node manually	154
Recovering a VCStack — SBx908 GEN2.....	155
Recovering a VCStack — other switches	157
Recovering AMF Plus devices with special links	158
Recovering and Provisioning Isolated Nodes	160
Recovering a TQ5403 series access point.....	165
Recovering a different access point (not TQ5403 series).....	168
AMF Plus guest nodes	175
Overview.....	175
Not all guest nodes are equal.....	175
AMF Plus guest discovery.....	176
AMF Plus functionality supported by AMF Plus guests	177
AMF Plus guest configuration	178
AMF Plus guest node show commands	183
AMF Plus support for ONVIF Profile Q devices.....	187
AMF Plus Backups	191
Backups by different types of nodes	191
Which files are backed up?	195
Backup destinations.....	196
Controlling the backup behaviour of controller and master nodes	198
Scheduling backups.....	198
Performing a manual backup	199
Backups on chassis or VCStacks running as AMF Plus controllers or masters	200
Forcing all master nodes in an area to perform a backup.....	201

Backing up to remote servers	203
Multiple backup destinations	209
AMF Plus Tunneling (Virtual-links)	211
Secure virtual-links	213
Configuring a virtual-link	214
Prioritizing the tunneled traffic.....	220
Virtual cross-links	223
Firmware Auto Upgrade	224
Advantages of reboot-rolling upgrade	224
Disadvantages of reboot-rolling upgrade.....	225
Advantages of distribute firmware upgrade	225
Support for AMF Plus Network Upgrades	225
Summary of the AMF Plus upgrade process	226
Detailed explanation of the AMF Plus upgrade process	226
Example 1 - Performing a reboot-rolling upgrade.....	229
Example 2 - AMF Plus distribute firmware upgrade	231
AMF Plus Security.....	233
Default security level	233
AMF Plus link management.....	233
Increasing AMF Plus security	234
AMF Plus restricted-login	235
AMF Plus Secure Mode.....	236
AMF Plus Cloud	246
Introduction	246
AMF Plus Cloud Documentation	247
What is AMF Plus virtualization?	247
Licensing	247
Multiple Tenants on AMF Plus Cloud.....	248
Introduction	248
Feature overview	248
Licensing	250

Configuration example for private cloud installation	251
Configuration example for public cloud installations	260
Using AMF Plus in EPSR Rings	269
Down-links and cross-links when adding AMF Plus to an EPSR ring	269
Dual-ring EPSR network with a common segment between two transit nodes	273

Introduction to AMF and AMF Plus

AMF is the Allied Telesis Autonomous Management Framework™, and AMF Plus is an expanded version of AMF. Both AMF and AMF Plus are a suite of features that combine to simplify network management across all supported network equipment from the core to the edge. They also integrate with Vista Manager, our graphical monitoring and management platform.

AMF and AMF Plus provide simplified device recovery and firmware upgrade management. Their primary function is to reduce the management and maintenance overhead on a network, while improving its responsiveness in handling equipment failures. They:

- enable an entire network to be managed as a single virtual device from any node within the network (excepting guest nodes). Configuration changes can be simultaneously made on multiple devices, and new devices can easily be assimilated into the network.
- can easily be overlaid on top of an existing network without changing its physical topology. They will determine the optimal logical topology for their own control plane.

The features of AMF and AMF Plus enable network engineers to lower network operating costs by reducing the complexity of network management and automating many routine tasks.

This guide provides a conceptual introduction to AMF and AMF Plus, together with their benefits, and presents configuration guidelines that explain their practical application in real networks.

This guide does not include Vista Manager's AMF Plus features. See the [Vista Manager EX Release Notes](#) and [User Guide](#) for details about those features.

The difference between AMF and AMF Plus

On the AlliedWare Plus command line, AMF and AMF Plus are identical. The difference between them is in Vista Manager, where AMF Plus includes additional AMF Plus intent-based networking features. The following figures show this difference.

Figure 1: Differences between AMF and AMF Plus

AMF feature	AMF	AMF PLUS	Vista Manager feature	AMF	AMF PLUS
Centralized management	✓	✓	Base features	✓	✓
Auto recovery	✓	✓	Wireless manager	✓	✓
Auto provisioning	✓	✓	Third-party device manager	✓	✓
Auto upgrade	✓	✓	Centralized management	✓	✓
Auto backup	✓	✓	Allied Intent-based Orchestrator menu	✓	
			Dynamic connections	✓	✓
			Internet breakout	✓	✓
			Auto traffic shaping	✓	✓
			Application priority	✓	✓
			AMF Plus menu		✓
			Health monitoring		✓
			Smart ACLs		✓
			Intent based QoS		✓

Vista Manager is purchased separately to AMF/AMF Plus
AMF Plus menu is included with your AMF Plus license and Vista Manager purchase

Products and software versions that apply to this guide

This guide applies to AlliedWare Plus™ products running version **5.5.2-2.3** and later. This includes the AMF Plus Cloud product.

AMF Plus works if devices are running versions earlier than 5.5.2-2.3, provided that master and controllers are running 5.5.2-2.3 or later. However, this guide assumes that all devices in your AMF Plus network have been upgraded to 5.5.2-2.3 and have AMF Plus licenses.

AMF and AMF Plus master and controller nodes require a feature license. For information about which products can act as AMF Plus master or controller nodes and the available licenses, see the [AlliedWare Plus Datasheet](#) and the [AMF Plus Datasheet](#).

AMF Plus was introduced in AlliedWare Plus firmware version 5.5.2-2.3. Versions 5.5.2-1.x and earlier can only use AMF licenses, whereas versions 5.5.2-2.x and above can use both AMF and AMF Plus licenses. Only AMF Plus licenses are now available for purchase.

Both AMF and AMF Plus provide the same AlliedWare Plus-based automation functions. However, AMF licenses do not include Vista Manager's AMF Plus networking features. To access these features, you need to upgrade from AMF to AMF Plus.

Note: For information about upgrading, see the [Vista Manager 3.10.x release note](#).

You only need to change to AMF Plus licenses if you want to manage more nodes or want to use the AMF Plus menu. Existing AMF licenses remain valid but cannot access the AMF Plus menu, and will instead see the Allied Intent-based Orchestrator (AIO) menu.

Because only AMF Plus licenses are now available for purchase, the rest of this guide refers only to AMF Plus and the AMF Plus menu.

Vista Manager EX AMF Plus requirements and licensing

The following requirements are needed to run AMF Plus:

- AlliedWare Plus firmware version 5.5.2-2.3 or later running on AMF masters and controllers.
- AMF Plus license for AMF masters and controllers.
- Vista Manager EX version 3.10.1 or later.

Note that AMF Plus devices must first be configured in the CLI before they can run with Vista Manager EX. For information on how to configure AMF Plus devices to run with Vista Manager EX, see "[Configuring AMF Plus to communicate with Vista Manager EX](#)" on page 21.

How many AMF Plus licenses do I need?

One AMF Plus license manages up to 10 nodes:

- If your network has 75 nodes, then 8 licenses are required.
- A license is available for either a 1 or 5 year period.
- The license code name is **AT-SW-APM10-xYR**

See the [AMF Plus datasheet](#) for full licensing details.

Can I mix AMF and AMF Plus licenses in Vista Manager?

Yes - it is possible for an AMF Master/Controller to have a combination of both AMF and AMF Plus node/area licenses.

However, in order to use the AMF Plus menu in Vista Manager, you must have **only** AMF Plus licenses.

If there are any AMF masters with any AMF node licenses or any AMF controllers with AMF area licenses, then Vista Manager will not display the AMF Plus menu, and will instead display the Allied Intent-Based Orchestrator (AIO) menu.

Both AMF and AMF Plus node/area licenses will count towards the total number of AMF nodes/areas available.

When is the AMF Plus menu visible in Vista Manager EX?

The AMF Plus menu is visible in Vista Manager EX when you **only** have AMF Plus licenses

The AMF Plus menu replaces the AIO menu in Vista Manager when all the AMF Masters and AMF Controllers have:

- An AMF Plus Controller/Master license on all Master and Controllers

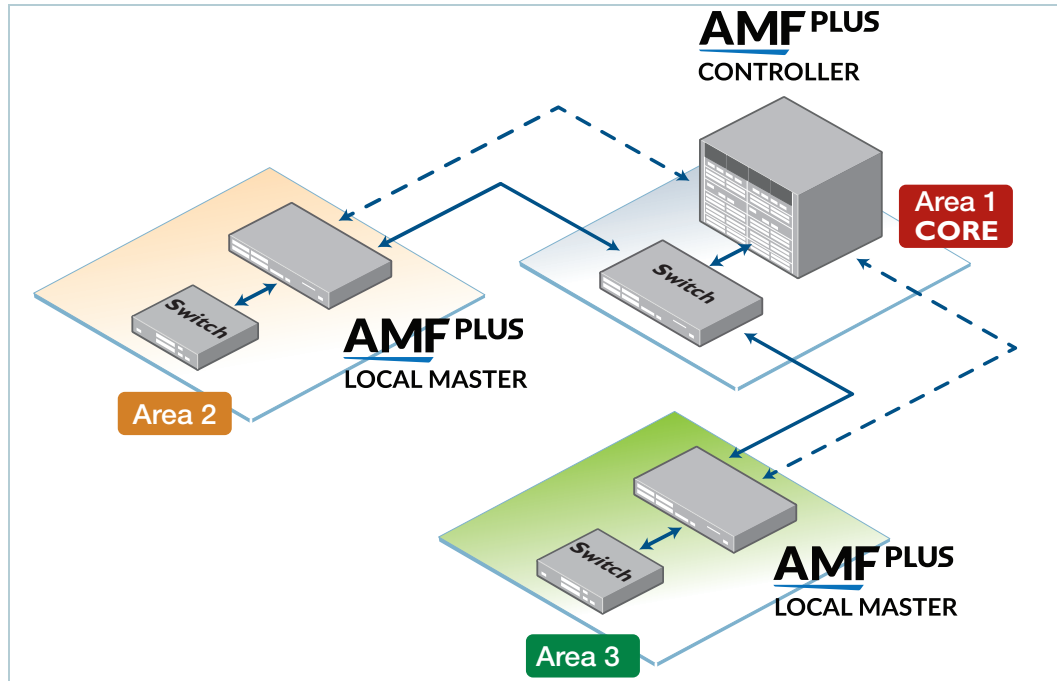
AND

- No AMF Controller/Master licenses applied or AMF Controller/Master licenses are disabled with the **atmf amfplus-license-only** command.

What is an AMF Plus network?

Before considering the detail of the various elements that operate together, it is worth taking a high-level overview of the structure of an **AMF Plus network**, in order to provide a visual of the context in which the elements exist.

Figure 2: AMF network overview



The primary structural entity in an AMF Plus network is the **AMF Plus area**.

Each area can consist of up to 300 network nodes, referred to as **AMF Plus member nodes**. These are coordinated by one or more units, known as **AMF Plus masters**. A single stand-alone AMF Plus area, consisting of its master(s) and member nodes, is a viable AMF Plus network and can provide a thorough range of AMF Plus functionality.

To scale up to larger networks, AMF Plus can operate across multiple AMF Plus areas. To operate a multi-area network requires a further level of hierarchy. This requirement to communicate with multiple AMF Plus areas is met by the introduction of an **AMF Plus controller**. A controller acts as a master for the masters in each of the individual AMF Plus areas.

A controller can connect to up to 60 AMF Plus areas. However, it can only connect to one area at a time. That is, the controller can connect to any one of its client masters, and perform management activities via that master, but cannot perform management activities on multiple masters simultaneously.

So, an AMF Plus network can be considered to be composed of two realms: the first realm is the management plane within each individual AMF Plus area (intra-area realm), and the second realm is the aggregation of individual areas (inter-area realm) into a larger management network. The aggregated management network is managed by a controller or multiple controllers. With area aggregation and multiple controllers, AMF Plus managed networks can grow up to 18,000 nodes.

Key benefits of AMF Plus

The key benefits of AMF Plus include its unified Command Line Interface (CLI), simple configuration backup and recovery process, smart provisioning of new network nodes, and time-saving rolling firmware upgrade.

Unified command-line

The conventional means of configuring and controlling AlliedWare Plus™ devices is to use their text-based Command Line Interface (CLI). In existing networks, this CLI is available via a serial console port and also via remote login sessions such as SSH.

AMF Plus extends this capability from managing either a single device, or a VCStack™ of devices, through to managing a whole network all from a single (unified) CLI session. Using the unified CLI, a network administrator can nominate either all nodes, or a subset of nodes, within the AMF Plus network to comprise an entity known as a **working-set**. Commands can then be executed concurrently across all network nodes within this working-set as if they were a single unit. Any existing configuration or diagnostic actions can thus be applied to multiple devices using a single command sequence. This reduces maintenance costs and configuration complexity, while still retaining complete flexibility in network design and control. For more information, see "[Configuring AMF Plus Nodes: the Unified CLI](#)" on page 130.

Vista Manager's AMF Plus menu

Vista Manager EX contains an AMF Plus menu where you can access various features and configure your network in different ways from its GUI. Features that you can use with AMF Plus include the AMF Plus menu, Health Monitoring, Smart ACL, and Intent-Based QoS. Features in the AMF Plus menu help you to monitor and manage your network efficiently. Note that if you do not have an AMF Plus license, the AMF Plus menu will not be visible and you will instead see the Allied Inten-based Orchestrator menu.

AMF Plus remote login

The AMF Plus remote login feature allows a user logged on to an AMF Plus node to connect to any other AMF Plus node. They can then run commands on that node as if they were local to that node.

In AMF Plus secure mode remote login to other AMF Plus nodes will only be allowed from an AMF Plus master node. AMF Plus member nodes will not be able to use the AMF Plus remote login feature to connect to other nodes in the network, including to AMF Plus master nodes.

Configuration backup and recovery

AMF Plus master nodes automatically backup the complete configuration information for all their member nodes, including boot configuration, firmware, licenses, and user scripts.

If an AMF Plus member node should fail, the AMF Plus process will automatically recognize and reconfigure an unconfigured replacement (standby unit), completely recreating the stored configuration of the failed unit into the replacement unit. The new unit will then reboot and resume service, without any need for user intervention beyond physical hardware replacement and cable

connection. In this way AMF Plus provides a complete zero-touch recovery solution. For more information, see "[AMF Plus Backups](#)" on page 191. Similarly, AMF Plus controller nodes can backup the master nodes of the AMF Plus areas under their control, to provide automatic recovery of failed masters.

Auto upgrade

Installing firmware upgrades on a production network is typically an infrequent but sensitive and labor-intensive process. AMF Plus is able to roll out upgrades to a user-selected subset of nodes. All that needs to be entered is the target group of nodes, and the location where the new firmware is stored; AMF Plus will then take care of the rest. Nodes are upgraded in a serial fashion, with each node being tested before continuing on to upgrade the next node.

If an upgrade fails on a particular node, the upgrade process is automatically terminated and that node will revert to its previous firmware version. In this way firmware updates are almost completely hands-free, whilst also providing confidence that a bad update will not result in loss of service. For more information, see "[Firmware Auto Upgrade](#)" on page 224.

Node provisioning

It is generally undesirable to have unconfigured devices connected to the network. Node provisioning enables you to preconfigure a port ready to accept and automatically configure a "clean" (as new) device for connection at a later date. This is achieved by storing the future node's configuration in the master node's backup files ready to be loaded to the new device when connected. For more information, see "[Node Provisioning](#)" on page 137.

Elements of AMF Plus

This section contains a description of the elements that make up AMF Plus and what each term means.

AMF Plus network

An **AMF Plus network** comprises a set of networked devices that contain embedded network management intelligence. These devices collaborate together, under the management of master and/or controller devices, to automate and expedite a number of network management activities.

Because there is an inherent limit to the number of devices that can fully collaborate together, network scalability is maintained by partitioning the network into an number of semi-independent managed regions called **AMF Plus areas**.

Network name

In order to provide the capability for networks to interconnect, an **AMF Plus network name** is necessary that can identify the AMF Plus network to which any given node belongs. It follows therefore, that all nodes within a single AMF Plus network must be configured with the same AMF Plus network name. Note that in an AMF Plus network consisting of multiple areas, all the member nodes in all the AMF Plus areas must be configured with the same AMF Plus network name.

Although each autonomous AMF Plus network has a finite size of 60 areas, data transfer may occur between devices residing in different autonomous AMF Plus networks. In this situation, we would want the user data (called **data plane information**) to pass between these networks. However, we do not want to pass the information that manages the internal operation of each individual AMF Plus network (called **control plane information**). The existence of different network names helps to ensure that there is no exchange of control plane traffic between autonomous networks.

AMF Plus nodes

Five types of nodes exist within an AMF Plus network: controller, master, member, gateway and AMF Plus guest nodes. Any of these, except the AMF Plus guest node, can comprise either a single switch, or a VCStack.

Controller node An **AMF Plus controller node** sits at the highest level of hierarchy in an AMF Plus network. A node is designated as a controller by the command **atmf controller**, see "[Configuring an AMF Plus controller](#)" on page 32.

The main functions performed by an AMF Plus controller are listed below:

- backing up the master nodes in the AMF Plus areas under its control. This can be on a scheduled basis, and/or on demand.
- recovering the master nodes within the AMF Plus areas under its control.
- running commands simultaneously on multiple nodes within the AMF Plus areas under its control (all the nodes that run the commands simultaneously must be within the same AMF Plus area).
- operating as the master node for its own local AMF Plus area.

Only one AMF Plus area in the AMF Plus network may contain controller nodes. Up to eight controller nodes can be created in this AMF Plus area, which forms the hub of a star topology, and they (the controllers) will operate independently of each other. We recommend that you have at least 2 AMF Plus controllers per network for redundancy purposes.

Master node **Master nodes** are user defined by configuration. They then form the core domain of the AMF Plus area. Aspects of master node functionality include:

- performing file system backups of all nodes in the AMF Plus area.
- acting as a file server for firmware and configuration for the member nodes in its area.
- providing an essential component for the formation of an AMF Plus network. That is, an AMF Plus network cannot exist without the existence of at least one master node.
- managing the membership of all nodes.
- recovering master or member nodes within the area.

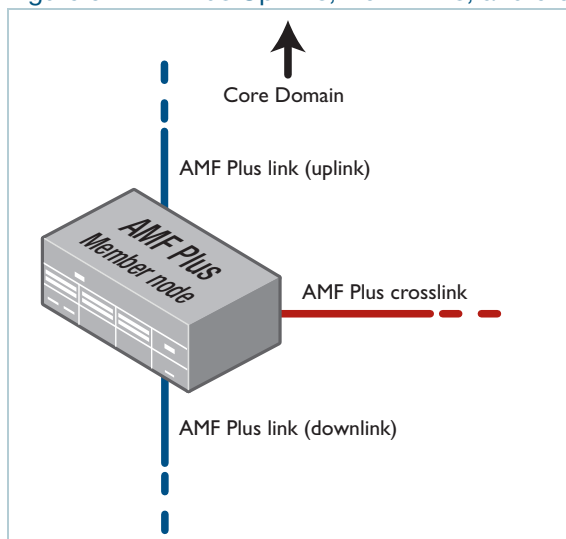
When more than one master node exists in an area, their operation is completely independent and unsynchronized. All master nodes within an AMF Plus area must reside within the same AMF Plus domain. We recommend that you have at least 2 masters per area for redundancy purposes.

- Member node** **Member nodes** are referred to simply as nodes. The maximum number of nodes in an AMF Plus area depends on which product is used as an AMF Plus master, and can be up to 300 nodes.
- Gateway node** **Gateway nodes** exist at the end of an AMF Plus area link and are referred to as the gateway nodes for their area. There are no special requirements on gateway nodes, they may be the controller or master node in their area, or they could be just member nodes.
- Edge node** **Edge nodes** are CentreCOM series switches that can only be used as edge switches in an AMF Plus network. The full management power and convenience of AMF Plus is available on these switches, but they can only link to one other AMF Plus node. They cannot form cross-links.
- Guest node** **Guest nodes** are devices that either do not run the AlliedWare Plus operating system or run a version that does not support AMF Plus. Guest functionality provides limited participation in an AMF Plus network. Guest devices do not require any operating system modifications or have AMF Plus software loaded onto them.
- Parent node** A **parent node** is an AMF Plus node that also directly connects to a specific AMF Plus node. For example, if an access point is connected to an AMF Plus network, then the node to which it directly connects is the access point's parent node.

Node interconnection

Nodes can connect either horizontally using **cross-links**, or vertically using **uplinks/downlinks**. This is shown in the illustration below:

Figure 3: AMF Plus Uplinks, Downlinks, and cross-links



AMF Plus links of either type are used to pass AMF Plus management traffic between nodes; however, they can also carry other network traffic.

- Cross-links** Cross-links are used to connect AMF Plus nodes to other AMF Plus nodes within what is termed an **AMF Plus domain**. Configuring an interface as an AMF Plus cross-link will automatically put its port into trunk mode. A cross-link can be a single link, a static aggregator, or a dynamic (LACP) aggregator. AMF Plus master nodes must be connected to each other using AMF Plus cross-links to ensure they are part of the uppermost domain level.

Up/down links Up/down links and virtual links interconnect domains in a vertical hierarchy, with the highest domain being the **core domain**. In effect, they form a tree of interconnected AMF Plus domains. The tree of interconnected AMF Plus domains must be loop-free, so there should never be rings formed by only up/downlinks.

In other words: Within each domain, cross-links between AMF Plus nodes define those nodes as siblings within the same domain. You can form rings by combining cross-links with up/down links and/or virtual links, as long as each AMF Plus domain links upwards to only a single parent domain. Each domain may link downwards to multiple child domains.

AMF Plus domains

Every AMF Plus node belongs to an **AMF Plus domain**. Domains can comprise of a single node or multiple nodes. AMF Plus master nodes are included in the highest domain level within an AMF Plus area, also known as the core domain, and all other domains are rooted in this domain.

As previously mentioned, domains are determined by AMF Plus cross-links. All nodes connected via cross-links form part of the same domain, and nodes connected via up/down links will be part of either higher or lower level domains.

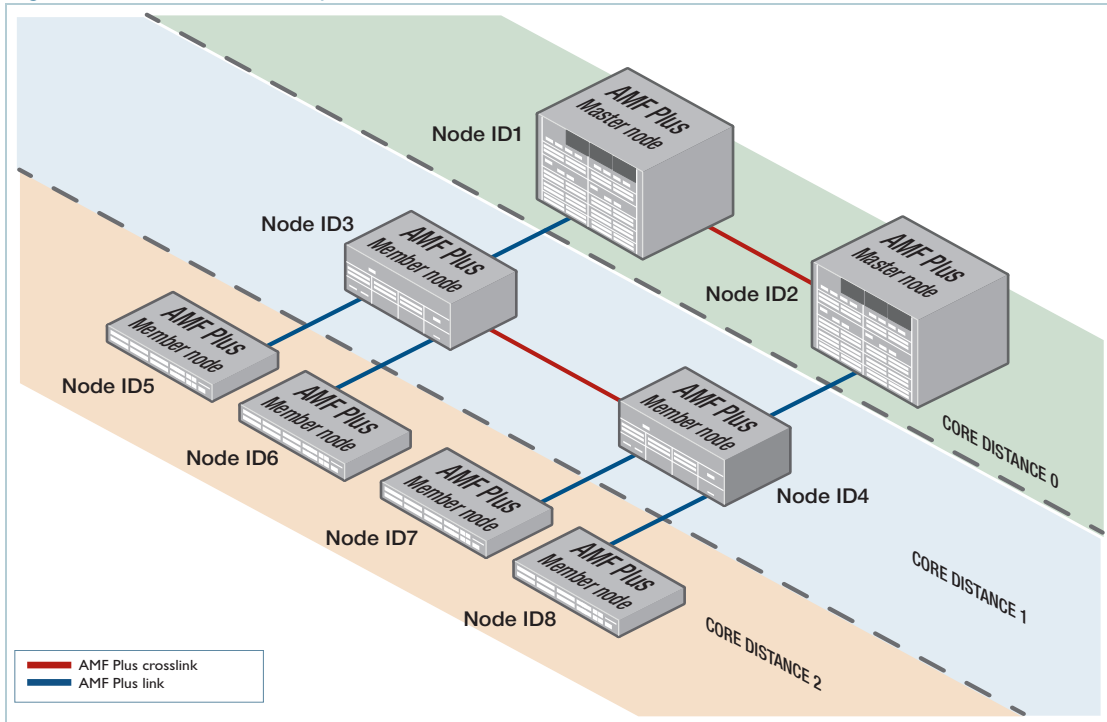
Connections between nodes that are in different domains are deemed to be vertical (because they connect from one layer to another), and connections between nodes in the same domain are deemed to be horizontal.

Note: Nodes within a domain must be connected in either a chain or ring topology. This means that a maximum of **two** cross-links should be configured on any single node.

The advantage of an AMF Plus domain is that two links from a domain to a single higher level domain will provide redundant AMF Plus links. We recommend you only connect each domain to one higher level domain, though you can connect each domain to multiple lower level domains. We also recommend that you set a maximum number of **12** nodes per domain.


Hop count The vertical distance of a domain from its core domain is known as its hop count. Figure 4 below shows the relationship between nodes, domains, and core distance (hop count). The core domain has a core distance (hop count) of 0, and the maximum recommended core distance in an AMF Plus area is eight.

Figure 4: Core distance hop counts between domains



If an AMF node has multiple vertical uplinks, the core distance for each uplink must be the same. For example, in the following diagram, the thick purple line shows an attempted virtual link. This virtual link is invalid, because node 5 would have a core distance of 2 through its uplinks and 1 through the virtual link. Configurations like this are not supported, because they lead to unpredictable behavior in the AMF Plus network.

Figure 5: An example of conflicting hop counts



AMF Plus areas

AMF Plus is a highly scalable framework, designed to unify the management of very large networks. The inherent value of AMF Plus is the capability to embed management intelligence into network nodes. This enables them to work cooperatively to automate network management tasks. As a result, devices within an AMF Plus region of operation need to maintain a reasonable degree of knowledge of all other devices in that region.

When operating on the scale of thousands of nodes, it is necessary to apply some structure to the Framework, by dividing it into separated operating regions. This way, strong integration can be maintained between nodes within a region, but the coupling between nodes in different regions can be considerably reduced. This is achieved in AMF Plus by dividing a network into regions known as **AMF Plus areas**.

Conceptually, an AMF Plus area consists of a series of domains, arranged in layers, with the core domain (the domain containing the master(s)) at the top. Each AMF Plus area consists of one or more master nodes, and a set of member nodes. The masters and members within an area operate in a unified fashion, but have no interaction with masters or members of other regions.

Coordinating the AMF Plus network as a whole are up to eight controller nodes, each of which can communicate with the master nodes in other areas. All the controller nodes may be configured either to communicate with the masters in **all** other areas, or in order to spread the load across the controllers, different controllers could be configured to communicate with the masters within selected sets of areas.

The area that contains the controller(s) is called the **core area**. The controllers are not necessarily the master nodes of their own local area. Configuring a node to be a controller is independent from configuring a node to be a master. So, the master node(s) of the core area can be quite separate from the controller(s) within that area. Or the controller(s) that exist in that area could also be configured to be master(s).

Virtual-links

It is simple to form an AMF Plus link between two AMF Plus nodes when they are directly connected to each other. However, a framework that relies on all member nodes being directly connected to each other is rather limited in scope. It is far better if the framework can extend across regions in which AMF Plus is not active. For example, it is desirable for the framework to extend between sites that communicate with each other via the Internet, or to be able to hop over a section of non-AMF Plus-capable equipment within a site.

These sorts of non-contiguous connections within an AMF Plus network are made possible by the use of **virtual-links**.

Virtual-links are achieved by encapsulating AMF protocol packets within IP wrappers (L2TPv3 encapsulation, to be exact), so that they can be transported across any arbitrary path that consists of IP forwarding devices.

Any AMF Plus node, except a guest node, can terminate a virtual-link. Virtual-links can be created between:

- member nodes
- member nodes and master nodes
- master nodes and controller nodes (actually, connections between controller nodes and nodes in other AMF Plus areas have a special status and are named area links).

The details of creating and optimizing virtual-links are described in the section "[AMF Plus Tunneling \(Virtual-links\)](#)" on page 211.

Area links

The links between different areas are termed **AMF Plus area links**. These links may be just normal direct AMF Plus links (i.e. AMF Plus links between directly connected devices) or they may be virtual-links.

The devices at each end of an area link are referred to as the gateway nodes for their area. There are no special requirements on gateway nodes. They could be the controller or master node in their area, or they could be just a standard member node.

The main restriction on area links is that they must run between the core area (the area that contains the controller(s)) and another area. It is not possible to have an area link between two non-core AMF Plus areas.

The details of configuring area links are described in the section "[Connections from AMF Plus controllers to the other areas](#)" on page 33.

Loop-free data plane

AMF Plus ensures that its own data plane (i.e. the AMF Plus VLANs) is kept loop-free. However, it does not manage the network data plane (i.e. the paths defined by data VLANs configured on the device). It is therefore important that the data plane configured in the network is kept loop free. RSTP is enabled by default for this, but see "[AMF Plus interaction with STP on AR-series devices](#)" on page 52 for important information about those devices.

Applications that use AMF Plus

Vista Manager EX™

Vista Manager EX is a graphical network monitoring and management tool for AMF Plus networks. It automatically creates a complete topology map from an AMF Plus network of switches, firewalls and wireless access points, showing areas and multiple levels of connected nodes and devices. Vista Manager EX facilitates simple management of many, or all, network devices from a dashboard that gives you a central overview of your network. From the dashboard you can monitor up-to-date network status, and take action to resolve any network problems.

Configuring AMF Plus to communicate with Vista Manager EX

AMF Plus devices must first be configured in the CLI before they can run with Vista Manager EX.

Perform the following steps on your AMF Plus network to allow it to communicate with Vista Manager EX.

Step 1: Activate the HyperText Transfer Protocol (HTTP) service.

Enable the HTTP service on all AMF Plus nodes, including all AMF Plus masters and controllers, using the following commands:

```
awplus#configure terminal
awplus(config)#service http
```

You can use an AMF Plus working set command to configure this option on all AMF Plus devices in an area:

```
awplus#atmf working-set group all
AMF[10]#configure terminal
AMF[10](config)#service http
```

Step 2: Allow Vista Manager EX to discover the AMF Plus network topology.

Run the following command on your AMF Plus controller (if you have one in your network) and all AMF Plus masters to allow Vista Manager EX to discovery your AMF Plus network:

```
awplus#configure terminal
awplus(config)#atmf topology-gui enable
```

You can use an AMF Plus working set command to configure this option on all controllers and masters in an area:

```
awplus#atmf working-set group controller, master
AMF[2]#configure terminal
AMF[2](config)#atmf topology-gui enable
```

Step 3: Configure the AMF Plus log event host.

If the AMF Plus controller, or AMF Plus master, you intend to register with Vista Manager EX is configured to send event notifications to Vista Manager EX, then Vista Manager EX will display them on its dashboard and event log page. This command need only be run on the AMF Plus controller/master registered with Vista Manager EX.

```
awplus(config)#log event-host <ip-address> atmf-topology-event
```

Additional Information

For information on installing and using the Vista Manager EX, see the [Vista Manager EX Installation and User Guide](#).

AMF Security Controller and AMF Application Proxy

The Allied Telesis AMF Security (AMF-Sec) Controller and AMF Application Proxy work with selected firewalls to provide additional protection to AMF nodes from malware or virus attacks.

The AMF-Sec Controller is management software for Allied Telesis devices. It is part of the Software-defined Networking (SDN) solution, which is a network architecture for controlling network traffic from a central controller. It simplifies network management by removing management tasks and decisions from individual devices or device stacks, and centralizing them. The AMF-Sec Controller and Allied Telesis devices communicate over a network pathway referred to as the control plane. The control plane can be based on either the OpenFlow protocol or the AMF Application Proxy.

This feature is available on AMF version 5.4.7-2.x or newer.

Configuring AMF to communicate with your AMF-Sec Controller

Perform the following steps on your AMF network to allow it to communicate with the AMF-Sec Controller.

Step 1: Activate the HyperText Transfer Protocol (HTTP) service.

Enable the HTTP service on all AMF masters, using the following commands:

```
awplus#configure terminal
awplus(config)#service http
```

You can use an AMF working set command to configure this option on all controllers and masters in an area:

```
awplus#atmf working-set group master
AMF[2]#configure terminal
AMF[2](config)#service http
```

Step 2: Activate the AMF application proxy service.

Run this command on all AMF nodes, including all AMF masters and controllers:

```
awplus(config)#service atmf-application-proxy
```

You can use an AMF working set command to configure this option on all AMF nodes in an area:

```
awplus#atmf working-set group master
AMF[2]#configure terminal
AMF[2](config)#service atmf-application-proxy
```

Additional Information

For information on installing and using the AMF-Sec Controller with AMF Application Proxy see the [AMF Security Controller documentation](#).

AMF Plus Subscription Licenses

AMF and AMF Plus master and controller nodes require a subscription license. For information about which products can act as AMF Plus master or controller nodes and the available licenses, see the [AlliedWare Plus Datasheet](#) and the [AMF Plus Datasheet](#).

AMF Plus was introduced in AlliedWare Plus firmware version 5.5.2-2.3. Versions 5.5.2-1.x and earlier can only use AMF licenses, whereas versions 5.5.2-2.x and above can use both AMF and AMF Plus licenses. Only AMF Plus licenses are now available for purchase.

Because only AMF Plus licenses are now available for purchase, the rest of this chapter refers only to AMF Plus licenses. However, the features and procedures are the same for AMF licenses.

AMF Plus starter license

All AMF Plus master capable devices come with a free 3-node starter license. The built-in starter license lets you try AMF Plus before investing in a more comprehensive licensing option.

Managing AMF Plus licenses

To subscribe to AMF Plus and manage your licenses, use the following steps:

Step 1: Obtain the serial number for your AMF Plus master and/or controller devices

Subscription licenses are tied to the serial number of the device.

Use the **show system serialnumber** command to display the serial number:

```
awplus# show system serialnumber
A05050G144700002
```

Step 2: Obtain the subscription license

To purchase a subscription license, contact your authorized Allied Telesis representative. You will need to supply the device serial number.

Step 3: Download the subscription license

Subscription licenses are contained in a Capability Response File (CRF). You can download the CRF from the [Allied Telesis Download Center](#) by logging into your account.

Once you have reached the **Download Center Homepage**, you can locate your device type by clicking **Search Devices** from the **Devices** menu on the left. You can select your specific device by clicking the serial number from the **Serial Number** list.

From the **View Device** page, you can download a CRF by clicking the **Download Capability Response** link. CRFs are saved as **.bin** files.

Step 4: Load the subscription license onto the device

After you have downloaded your CRF, you can transfer it onto the device's Flash storage by any preferred method. For example, you can use the **copy** command to copy the CRF file from a USB device to your Flash storage:

```
awplus#copy usb flash
```

Output 1: Example from the **copy usb flash** command

```
awplus#copy usb flash
Enter source path with file name[:A05050G144700002.bin
Copying...
Successful operation
```

Step 5: Activate the subscription license

Display the filename of the CRF in Flash storage, by using the following command:

```
awplus#dir *.bin
```

Then activate it by using the following command:

```
awplus#license update <CRF-filename>
```

This command copies license entitlements from the CRF into the device's internal encrypted license library. You can then safely delete the CRF from the device. For this command to successfully activate the license, the CRF must be valid and be tied to the serial number of the device.

Step 6: Verify your CRF activation

You can verify the license by using the following command:

```
awplus#show license external
```

This displays the license name, the serial number of the device, and the license's valid dates.

Updating subscription licenses

If a subscription license expires, the device immediately reverts to the 3-node AMF Plus Starter license. Warning messages will be printed in the device log 28 days, 21 days, 14 days, 7 days, and 1 day prior to a license expiring. The Allied Telesis Download Center will also send you an email reminder prior to your license expiring.

To renew your license, contact your Allied Telesis representative. You can use the command **show license external** to confirm the serial number of the device. After renewing the license, follow steps 3-6 above to download and activate it.

Subscription licenses on VCStacks

If you are licensing a VCStack, you only need to purchase a license for one member of the stack. This does not need to be the VCStack master.

To load the license onto the stack, follow the steps above on the stack master. The software checks that the CRF is valid for one of the stack members and applies the license entitlement to all members of the stack. The command **show license external stored** shows which stack member is the source of the license entitlement.

Output 2: Example from the **show license external stored** command

```
awplus#show license external stored

Feature entitlements sourced from license file on local flash:

Stack member 1, serial A04435H101200015
No valid entitlements found

Stack member 2, serial C20YB7309

AMF Master

    Start date:                25 Jan 2023 00:00
    Expiry date:               19 Jan 2024 23:59
    Maximum nodes:             10

Stack member 3, serial B04435H101200015
No valid entitlements found
```

If you need to modify the license, for example to extend the date or change the number of nodes under management, make sure you modify the license for the same device as the original license. Do not create a new license for a different stack member instead.

If a device leaves the stack

If the device that is the source of the license entitlement leaves the stack, then:

- a warning message alerts you to this event. The message displays on the console, is logged, and appears in the **show license external** output.
- the remaining members of the stack retain their entitlement and continue to operate as an AMF Plus controller/master without any disruption in service.
- if the remaining partial stack reboots, it loses access to the license when it restarts.

If you need to permanently replace the device that is the source of the license entitlement, you can transfer the license to another stack member. To do this:

1. On the [Allied Telesis Download Center](#), transfer the license to the other stack member's serial number.
2. Follow steps 3-4 above to transfer the CRF to the stack member.
3. Force the stack to re-synchronize its license entitlement by using the command:

```
awplus#license redistribute
```

Multiple copies of a license on a stack

As said above, you only need to purchase a single license for multiple stack members, and therefore you only need to activate **one** CRF for the whole stack.

However, if you activate multiple CRFs for the same feature on the stack, the stack will obtain its license entitlements from the device with the **lowest** stack-ID. Note that stack-ID is the only factor that determines which license is used; factors such as license expiry date are not checked.

This means that it is possible (but not recommended) to have multiple CRFs for the same feature, where those CRFs have different expiry dates or support a different number of nodes. In that situation, it is possible for the stack to obtain the wrong license entitlements. If the stack obtains the wrong license entitlements:

- enter the **license redistribute** command.
- if that does not resolve the issue, then renumber the stack members so that the device with the preferred license entitlements has the lowest stack-ID amongst the devices that have any license installed, and reboot the renumbered devices. Once the stack has fully re-formed, if licenses are still not as desired, enter the **license redistribute** command again.

Automatically obtaining and activating licenses

Software version 5.4.6-2.x introduced simplified installation of licenses. Simply run the following command:

```
awplus#license update online
```

When the command **license update online** is entered, the device will:

1. Connect to the Download Center
2. Check if new or changed licenses are available for the device, keyed to the device's serial number
3. For each such license it finds, download and install the license.

Note that AlliedWare Plus devices do not automatically connect to the Download Center and check whether licenses are available. They only check when you run the **license update online** command.

On VCStacks, running **license update online** updates all stack members. Each stack member individually checks for licenses on the Download Center and installs any that are found.

On SBx8100 systems, running **license update online** updates all CFCs that are present, including all CFCs on both chassis in a stack. Each CFC individually checks for licenses on the Download Center and installs any that are found.

Firewall rules

Subscription licensing originating from firewall

Most firewalls block all traffic by default, so in order for the 'license update online' command to function correctly, you may need to configure your firewall to allow outbound DNS lookups and HTTPS connections. The following figure shows a recommended example configuration for an Allied Telesis AR-Series firewall, when the WAN interface to the Internet is configured as a ppp0 interface and the subscription licensing is being performed from the firewall itself.

```
zone public
network wan
  ip subnet 0.0.0.0/0 interface ppp0
  host ppp0
  ip address dynamic interface ppp0

firewall
rule 10 permit https from public.wan.ppp0 to public.wan
rule 20 permit dns from public.wan.ppp0 to public.wan
protect
```

These rules permit DNS and HTTPS packets to any destination IP address, if:

- the source IP address of the packets is the IP address of the ppp0 interface, and
- the packets are egressing the firewall via interface ppp0.

DNS packets are permitted so that the device can look up the address of the Download Center. HTTPS packets are permitted so the secure communication session with the Download Center can proceed.

The rule uses a subnet of 0.0.0.0/0 to match on any destination IP address.

The "from" part of the rule uses "public.wan.ppp0" because the firewall itself is originating the connection to the Download Center, rather than allowing traffic to flow through it. The traffic that is involved in the connection to the Download Center originates from the IP address of the PPP interface.

Subscription licensing through the firewall

AlliedWare Plus devices configured with features such as AMF and OpenFlow also use subscription-based licensing. These devices could be located within a private firewall zone, accessing the subscription service located in the Internet, via the AR-Series firewall.

In order to allow access to the subscription licensing services from a private zone to the Internet, firewall permit rules need to be created.

```
zone private
network lan
ip subnet 10.1.1.0/24 interface vlan1

zone public
network wan
ip subnet 0.0.0.0/0 interface ppp0
host ppp0
ip address dynamic interface ppp0

firewall
rule 30 permit https from private to public
rule 40 permit dns from private to public.wan
protect
```

These rules permit DNS and HTTPS packets to any destination IP address, to allow devices located within a private zone to access subscription-based services located on the Internet through the AR-Series firewall. If the firewall is also performing NAT, then corresponding NAT-based masquerade rules for HTTPS and DNS will also need to be configured. For more information about firewall and NAT rules, see the [Firewall and Network Address Translation \(NAT\) Feature Overview and Configuration Guide](#).

Verifying the update

The update process normally takes approximately 5 seconds.

If the console does not respond for 10 or more seconds after typing the command, a network, routing or firewall configuration error is probably preventing the connection from establishing. If this happens, you can abort the command by pressing Ctrl-C, or wait for the command to time out after 30 seconds.

If the connection to the Download Centers fails and times out, an error message will be generated on the CLI to indicate the problem. If you abort the command, no error message is displayed.

If the update is successful, the device will produce log messages to say which features have had their licensing state updated (activated, deactivated, or expiration/count changed). If the command completes successfully but there are no licenses available for the device, or no change in the licenses already on the device, no log messages will be produced.

You can also use the **show license external** command to confirm which licenses are active on the device after the update has been applied.

Combining AMF and AMF Plus licensing

From AlliedWare Plus version 5.5.2-2.3 onwards, AMF licenses are no longer available to buy. Instead, AMF Plus licenses are available. Existing AMF licenses are still valid, so you only need to change to AMF Plus licenses if you:

- need to manage more nodes, or
- want to use the new AMF Plus menu in Vista Manager EX.

The following sections describe how to change AMF masters to AMF Plus. The same steps apply to AMF controllers.

Managing more nodes

If you want to add more nodes, and you do not want to use AMF Plus in Vista Manager EX, then simply buy an AMF Plus license for the additional nodes and install it on your AMF Plus master. You do not need to change the existing AMF license.

For example, if you have a license for 10 AMF nodes and want to add another 10 nodes, buy and install an AMF Plus license for the extra 10 nodes.

Using AMF Plus in Vista Manager EX

If you want to use AMF Plus in Vista Manager EX, then you need to change your licensing to AMF Plus completely. To do this:

Step 1: Obtain enough AMF Plus licenses to cover all the nodes in your AMF Plus network.

This will be the existing AMF nodes and any additional nodes you need.

Step 2: Upgrade to 5.5.2-2.3

Upgrade the AMF master to 5.5.2-2.3 or later. If possible, upgrade all other nodes too, especially any that have downstream nodes attached to them.

Step 3: Install the AMF Plus license

Follow the steps above to install the AMF Plus license. Alternatively, you can do this in Vista Manager EX, through the **Asset Management** page.

Step 4: Change to the AMF Plus license

Change to using AMF Plus instead of AMF. To do this, either enter the following command on your master:

```
awplus(config)#atmf amfplus-license-only
```

or in Vista Manager EX, go to **System Management > Licenses** and enable the **AMF Plus Forced** button:

AMF Licenses

i AMF Plus functionality is only available when all AMF Controllers and Masters in the network have an active AMF Plus license without any active AMF licenses, or the AMF licenses have been negated by enabling AMF Plus Forced.

⚠ Changing these values will save the current running configuration to the device.

Device Name	AMF Plus Forced
AMF-Cloud-Master M	<input checked="" type="checkbox"/>

The Concept of AMF Plus Areas

AMF Plus is a highly scalable framework, designed to unify the management of very large networks.

The inherent value of AMF Plus is the capability to embed management intelligence into network nodes. This enables them to work cooperatively to automate network management tasks. As a result, devices within an AMF Plus region of operation need to maintain a reasonable degree of knowledge of all other devices in that region. When operating on the scale of thousands of nodes, it is necessary to apply some structure to the Framework, and divide it into separate operating regions. This maintains strong integration between nodes within a region while reducing the coupling between nodes in different regions.

This is achieved in AMF Plus by dividing a network into regions known as AMF Plus areas.

Each AMF Plus area consists of one or more master nodes, and a set of member nodes. The masters and members within an area operate in a unified fashion, but have no interaction with masters or members of other regions.

Coordinating the AMF Plus network as a whole are up to eight controller nodes, each of which can communicate with the master nodes in other areas. All the controller nodes may be configured either to communicate with the masters in **all** other areas, or in order to spread the load across the controllers, different controllers could be configured to communicate with the masters within selected sets of areas.

The area that contains the controller(s) is called the **core area**. The controllers are not necessarily the master nodes of their own local area. Configuring a node to be a controller is independent from configuring a node to be a master. The master node(s) of the core area can be quite separate from the controller(s) within that area. Or the controller(s) that exist in that area could also be configured to be master(s).

Configuring an AMF Plus controller

To set up a node as an AMF Plus controller it first needs to have an AMF Plus controller license installed. Then it can be configured it as a controller by using the command:

```
atmf controller
```

The area to which it belongs needs to be given a name and an ID number:

```
atmf area <area-name> id <1-4094> [local]
```

The parameter **local** indicates that this command is specifying the name and ID number of the area in which the controller resides.

The controller needs to be informed of the identities of the areas it is controlling, and to be given passwords for authenticating its communication to masters in those areas.

The following commands need to be configured on the controller for each of the areas it controls:

```
atmf area <area-name> id <1-4094>
atmf area <area-name> password [8] <password>
```

The corresponding password also needs to be configured on master nodes and gateway nodes (see below for an explanation of gateway nodes) in each of those areas.

In addition the controller will need configuration relating to the backing up of master nodes in the controlled areas, as described in the section ["Controlling the backup behaviour of controller and master nodes"](#) on page 198.

Connections from AMF Plus controllers to the other areas

For an AMF Plus controller to communicate with the other areas that comprise the network, it needs AMF Plus links into those areas. The core area has a link to each of the other areas in the network. The other areas do not have links to each other, so the only inter-area links are those from the core area to the other areas.

The links from the core area to another area are referred to as **area links**. These links may be direct connections between neighboring nodes, or they may be virtual-links; either is quite valid. The end points of the area links can be any nodes within the two areas that are being connected. There is no requirement that the area link terminate on a controller or master node, and similarly there is no rule that the area link can't terminate on a controller or master node.

The devices at each end of the area link are referred to as **gateway nodes**, as they constitute the "gateways" into their respective areas.

The configuration required on a gateway node is:

1. The identity of the area that the node belongs to.

```
atmf area <area-name> id <1-4094> local
```

2. If the gateway is **not** in the core area, then it needs to be configured with the password for the area in which it resides.

If the gateway is in the core area, then it needs to be configured with the passwords for any areas to which it will be forming area links.

```
atmf area <area-name> password <password>
```

3. The names and ID number of any areas to which the gateway will be forming area links. If the gateway is not in the core area, then it will only be forming an area link to the core area, and therefore only needs the name and ID number of the core area.

A gateway node in the core area may be forming area links to multiple remote areas, and will need to be configured with the names and ID numbers of all those areas.

```
atmf area <area-name> id <1-4094>
```

4. The area link definition(s)

If an area link is a link between directly connected neighbors, then the area link is simply configured on the interface that connects to the neighbor in the other area.

```
interface portx.y.z
switchport atmf-arealink remote-area <area-name> vlan <2-4094>
```

The VLAN that is configured on the area link is a VLAN that must be dedicated to the area link, and not used for other purposes.

If an area link is a virtual-link, then the link is defined like a normal virtual-link, as described in ["AMF Plus Tunneling \(Virtual-links\)" on page 211](#), except that an extra **remote-area** parameter is appended to the command, to indicate that the far end of the virtual-link is in another area.

```
atmf virtual-link id <1-4094> ip <a.b.c.d> remote-id <1-4094> remote-ip
<a.b.c.d> remote-area <area-name>
```

Configuring master nodes in a multi-area network

The master nodes in the core and non-core areas need to be aware of which area they are in, so that they will correctly validate connections from the controller(s). The master nodes need to be configured with the identity of the area they belong to, and the password for that area. In addition master nodes need to be configured with the identity of the controller's area.

```
atmf area <area-name> id <1-4094> local
atmf area <area-name> password <password>
atmf area <controller-area-name> id <controller-area-id>
```

Connecting from the controller to another area

A key benefit of having controller nodes is that they can be used to carry out management tasks in any of their controlled areas. From one place - a login on the controller - a network manager can operate on any node in the whole multi-area network.

This is achieved by the controller connecting to a master in any of its controlled areas, it can then carry out any activities that this master can perform, such as kicking off a rolling reboot in its area, or executing commands on a working-set within that area, or provisioning new nodes with that area.

The command that connects the controller to a remote master is:

```
atmf select-area <area-name>
```

Once this command has been entered, you are effectively in control of the master node in the specified area, and can execute any commands that could be executed in that master. To relinquish control of that remote master node, enter one of the following commands:

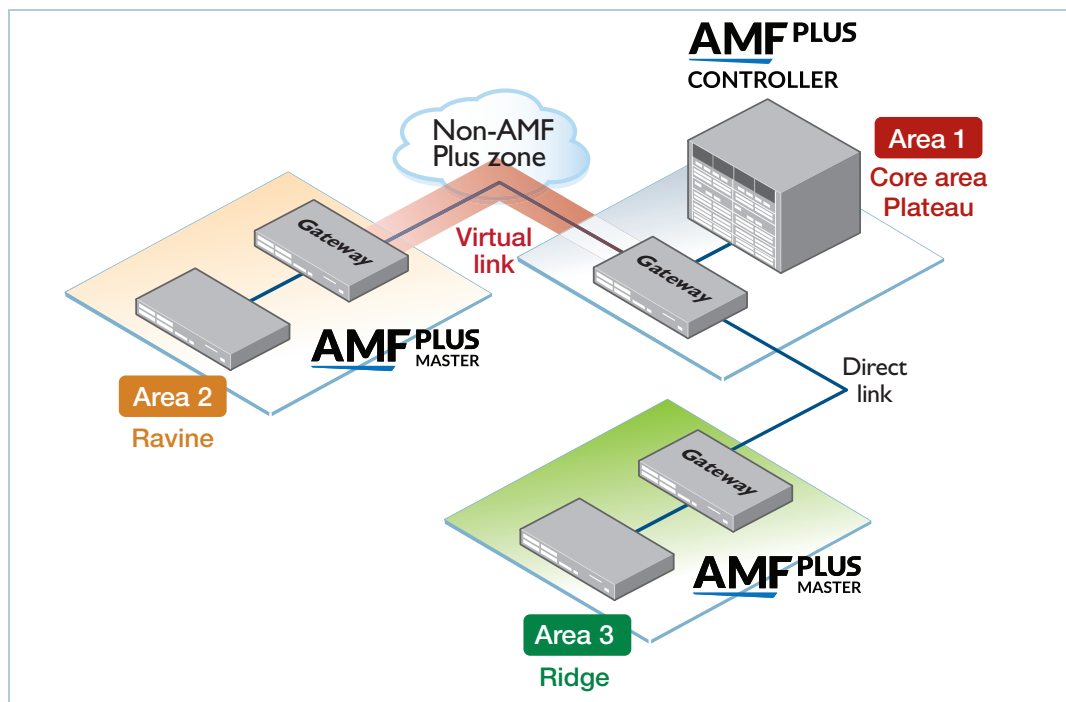
```
no atmf select-area
```

or

```
atmf select-area local
```

Note that a controller can only connect to one remote master at a time.

Example - Configuring a multi-area network



The AMF Plus and related configurations for the six nodes illustrated in this diagram are:

Controller/master in the core area

```

Hostname Highpoint
atmf network-name Terrain
atmf controller
atmf master
atmf area Plateau id 1 local
atmf area Ravine id 2
atmf area Ravine password 8 ElphMJTGVBWuCXcv4xtM19cBE+1wWa/KGtPmEmosAI4=
atmf area Ridge id 3
atmf area Ridge password 8 9nSHUJAdV6mHEygAhpCutXmywVgwAHAE6e4U42e1158=
atmf management vlan 4000
atmf backup area-masters enable
atmf backup server id 1 10.37.74.1 username root path /tftpboot/
backups_from_on_highlander
atmf backup 14:30 frequency 4

interface port1.0.9
  switchport atmf-link
  switchport mode trunk
  switchport trunk native vlan none

```

Gateway in the core Area

```

Hostname Highgate
vlan database
  vlan 10
atmf network-name Terrain
atmf area Plateau id 1 local
atmf area Ravine id 2
atmf area Ravine password 8 ElphMJTGVBWuCXcv4xtM19cBE+1wWa/KGtPmEmosAI4=
atmf area Ridge id 3
atmf area Ridge password 8 9nSHUJAdV6mHEygAhpCutXmywVgwAHAE6e4U42e1158=

atmf virtual-link id 12 ip 154.23.17.9 remote-id 21 remote-ip 92.48.201.10
remote-area Ravine

interface port1.0.9
  switchport atmf-link
  switchport mode trunk
  switchport trunk native vlan none

int port1.0.10
  switchport mode trunk
  switchport trunk allowed vlan add 10

int port1.0.11
  switchport atmf-arealink remote-area Ridge vlan 20
  switchport mode trunk
  switchport trunk native vlan none

int vlan10
  ip address 154.23.17.9

```

Master in Area 2

```

Hostname Rapids
atmf network-name Terrain
atmf master
atmf area Ravine id 2 local
atmf area Ravine password 8 ElphMJTGVBWuCXcv4xtM19cBE+1wWa/KGtPmEmosAI4=
atmf area Plateau id 1
atmf management vlan 4000

interface port1.0.9
  switchport atmf-link
  switchport mode trunk
  switchport trunk native vlan none

```

Gateway in Area 2

```

Hostname Moraine
vlan database
  vlan 10
atmf network-name Terrain
atmf area Ravine id 2 local
atmf area Ravine password 8 E1phMJTGVBWuCXcv4xtM19cBE+1wWa/KGtPmEmosAI4=
atmf area Plateau id 1

atmf virtual-link id 21 ip 92.48.201.10 remote-id 12 remote-ip 154.23.17.9
remote-area Plateau

interface port1.0.9
  switchport atmf-link
  switchport mode trunk
  switchport trunk native vlan none

int port1.0.10
  switchport mode trunk
  switchport trunk allowed vlan add 10

int vlan10
  ip address 92.48.201.10

```

Master in Area 3

```

Hostname Peak
atmf network-name Terrain
atmf master
atmf area Ridge id 3 local
atmf area Ridge password 8 9nSHUJAdV6mHEygAhpCutXmywVgwAHAE6e4U42e1158=
atmf area Plateau id 1
atmf management vlan 4000

atmf backup server id 1 192.168.231.54 username climber path /home/
climber/node_backups
atmf backup redundancy enable
atmf backup 16:30 frequency 2

interface port1.0.9
  switchport atmf-link
  switchport mode trunk
  switchport trunk native vlan none

```

Gateway in Area 3

```

Hostname Saddle
atmf network-name Terrain
atmf area Ridge id 3 local
atmf area Ridge password 8 9nSHUJAdV6mHEygAhpCutXmywVgwAHAE6e4U42e1158=
atmf area Plateau id 1

interface port1.0.9
switchport atmf-link
switchport mode trunk
switchport trunk native vlan none

int port1.0.11
switchport atmf-arealink remote-area Plateau vlan 20
switchport mode trunk
switchport trunk native vlan none

```

Area links on AR-series Eth ports

AMF Plus area links are supported over an AR-series device's Eth interfaces. To use this feature your AMF Plus network must be in AMF Plus secure mode.

Use the **atmf-arealink** command on an Eth interface to configure it as an AMF Plus area link. For example, to configure the Eth1 interface as an AMF Plus area link to the 'Auckland' area on VLAN 6, use the following commands:

```

awplus# configure terminal
awplus(config)# interface eth1
awplus(config-if)# atmf-arealink remote-area Auckland vlan 6

```

Areas with 120-300 nodes

If your network includes AMF Plus areas over 120 nodes in size, you must enable restricted-login, by using the commands:

```

node1#configure terminal
node1(config)#atmf restricted-login

```

When restricted-login is enabled, only the AMF Plus master nodes are able to create working sets. Rolling-reboot is also only available from master nodes, because it uses working sets.

120-300 node AMF Plus areas are only available if every node in your network is running version 5.4.6-2.1 or later.

Controllers with 301-1000 areas

When using AMF Plus Cloud as the controller of 301 to 1000 areas, the following restrictions apply:

- Multiple controllers are not supported.
- The controller must be the only master in its area.

Step 2: Set the AMF Plus network name

The network name must be the same on all nodes in all areas within the AMF Plus network. More precisely, within the same **autonomous AMF Plus network**. For more information see ["Network name" on page 14](#).

```
AMF_Master(config)#atmf network-name atmf1
```

Step 3: Configure the switch to be the AMF Plus master

```
AMF_Master(config)#atmf master
```

- On standalone devices one master license is required per device. A SBx8100 chassis with dual CFC controller cards fitted is treated as a single device, and so only a single AMF Plus master license is required. If the SBx8100 chassis has two CFC cards installed, install the license on one CFC card. It will be automatically copied over to the other CFC card.
- On VCStacks, only a single AMF Plus master license is required per VCStack. Install this license onto any stack member.

Step 4: Configure the data VLANs

```
AMF_Master(config)#vlan database
AMF_Master(config-vlan)#vlan 2
AMF_Master(config-vlan)#vlan 3
AMF_Master(config-vlan)#exit
```

Step 5: Configure ports as AMF plus links

```
AMF_Master(config)#interface port1.1.1-1.1.2
AMF_Master(config-if)#switchport atmf-link
```

Step 6: Configure data VLANs on AMF Plus links as required

```
AMF_Master(config-if)#switchport trunk allowed vlan add 2-3
AMF_Master(config-if)#switchport trunk native vlan none
```

Step 7: Save the configuration

```
AMF_Master#copy running-config startup-config
Building configuration...[OK]
```

Configuring Member1

The configuration of each of the member nodes does not differ vastly. However, the set of ports used for AMF Plus links is not the same on all the members. Member1 and Member2 both use one set of ports, while Member3 and Member4 use another set. So, we will look first at the configuration required on Member1 and Member2, and then consider the configuration used on Member3 and Member4.

Configure Member1

Step 1: Set the host name

Host names are used as the node name for AMF Plus nodes therefore they MUST BE UNIQUE within the AMF Plus area. Each of the member nodes needs to be given a different hostname, e.g. Member1, Member2.

```
awplus#configure terminal
awplus(config)#hostname Member1
```

Step 2: Set the AMF Plus network name

The AMF Plus network name must be the same on all nodes on all AMF Plus areas within the AMF Plus network.

```
Member1(config)#atmf network-name atmf1
```

Step 3: Configure the data VLANs

```
Member1(config)#vlan database
Member1(config-vlan)#vlan 2-3
```

Step 4: Configure ports as AMF Plus links

```
Member1(config)#interface port1.1.1,port1.1.3
Member1(config-if)#switchport atmf-link
```

Step 5: Configure data VLANs on the AMF Plus links as required

```
Member1(config-if)#switchport trunk allowed vlan add 2-3
Member1(config-if)#switchport trunk native vlan none
```

Step 6: Configure an AMF Plus cross-link

AMF Plus links and cross-links do not need to be configured with data VLANs and can be used solely to provide redundant links in the AMF Plus management VLAN.

```
Member1(config)#interface port1.1.2
Member1(config-if)#switchport atmf-crosslink
Member1(config-if)#switchport trunk allowed vlan add 2-3
Member1(config-if)#switchport trunk native vlan none
```

Step 7: Save the configuration

```
Member1#copy running-config startup-config
```

Configure Member2

Because members 1 and 2 have the same port configuration, you can repeat the steps used to configure Member1 but set the hostname to be Member2.

Configure Member3

Step 1: Set the host name

Host names are used as the node name for AMF Plus nodes and MUST BE UNIQUE within the AMF Plus area. So, each of the member nodes needs to be given a different hostname, e.g. Member3, Member4.

```
awplus#configure terminal
awplus(config)#hostname Member3
```

Step 2: Set the AMF Plus network name

The AMF Plus network name must be the same on all nodes in all areas within the AMF Plus network.

```
Member3(config)#atmf network-name atmf1
```

Step 3: Configure the data VLANs

```
Member3(config)#vlan database
Member3(config-vlan)#vlan 2-3
```

Step 4: Configure ports as AMF Plus links

```
Member3(config)#interface port1.0.1
Member3(config-if)#switchport atmf-link
```

Step 5: Configure data VLANs on the AMF-links as required

```
Member3(config-if)#switchport trunk allowed vlan add 2-3
Member3(config-if)#switchport trunk native vlan none
```

Step 6: Save the configuration

```
Member3#copy running-config startup-config
```

Configure Member4

Because members 3 and 4 have the same port configuration, you can repeat the steps used to configure Member3 but set the hostname to be Member4.

Verifying the AMF Plus Network

To check that all nodes have joined the AMF Plus network use the **show atmf** command with the **summary** parameter. You can run this command from any node in an AMF Plus network.

Output 3: Checking AMF Plus configuration using the show atmf summary command

```
AMF_Master#show atmf summary
ATMF Summary Information:
ATMF Status           : Enabled
Network Name          : atmfl
Node Name              : AMF_Master
Role                   : Master
Current ATMF Nodes   : 5
AMF_Master#
```

The **Current ATMF Nodes** field in the output above shows that all 5 nodes have joined the AMF Plus network.

Use the **show atmf nodes** command to check information on individual nodes:

Output 4: Output from the show atmf nodes command

```
AMF_Master#show atmf nodes
Node Information:
 * = Local device
  SC = Switch Configuration:
    C = Chassis   S = Stackable   N = Standalone
Node      Device      ATMF      Node
Name      Type          Master     SC  Parent  Depth
-----
* AMF_Master  AT-SBx81CFC960    Y         CS  none    0
Member1     AT-SBx908 GEN2    N         S   AMF_Master  1
Member2     AT-SBx908 GEN2    N         S   AMF_Master  1
Member4     x230-18GT         N         S   Member2    2
Member3     x230-18GT         N         S   Member2    2
Current ATMF node count 5
```

The **Parent** field in the output above refers to the parent domain and not the upstream device. In the example output above, Member2 is the domain controller for the parent domain for Member3 and Member4.

Use the **show atmf links** command to check information on individual AMF Plus links:

Output 5: Checking output with the show atmf links command

```
switch1# show atmf links
```

ATMF Links Brief:

Local Port	Link Type	Port Status	ATMF State	Adjacent Node	Adjacent Ifindex	Link State
sa1	Crosslink	Up	TwoWay	Building_1	4501	Forwarding
1.1.1	Downlink	Up	Full	Bld1_Floor_1	5001	Forwarding
1.1.2	Downlink	Up	Full	Bld1_Floor_2	5003	Forwarding
1.1.3	Downlink	Up	Full	Bld2_Floor_1	6101	Forwarding
1.1.4	Crosslink	Down	Init	*switch3		Blocking

* = provisioned

User account management

The default **username** for an AlliedWare Plus login is **manager** and the default **password** is **friend**. Users should change this password on all their nodes to provide login security.

AMF Plus also supports remote login authentication via either RADIUS or TACACS+.

Because AMF Plus's goal is to provide a uniform management plane across the whole network, if you are using the local user database we recommend you create the same user accounts on all the nodes in the network. In reality, though, it is not essential to have the same accounts on all the nodes. Users can remote login from one node to a second node even if they are logged into the first node with a user account that does not exist on the second node, provided that restricted-login is disabled and the user account on the first node has privilege level 15.

If login authentication via RADIUS or TACACS+ is configured, a user is remotely authenticated when they first log into a device on the AMF Plus network. Thereafter AMF Plus uses a key exchange mechanism. This means the default AAA authentication method group of **local** is compatible with AMF Plus. For AMF Plus to work correctly using RADIUS or TACACS+, ensure the AAA method includes **local** as a backup method group, for example:

```
aaa authentication default group radius local
```

NTP and AMF Plus

AMF Plus uses NTP to synchronize the system clocks across nodes within the network. All AMF Plus nodes automatically receive time from the AMF Plus master's NTP server. For this to operate you need to configure at least one external NTP server on each AMF Plus master in your network to ensure accurate logging, and consistent timestamps between all AMF Plus nodes. Configuration of three or more NTP servers is considered best practice. Configured servers do not need to be the same between AMF Plus masters. One option is to use the pool of NTP servers provided by the NTP Pool Project (www.pool.ntp.org).

In some networks the AMF Plus masters may not have a path to an external NTP server. This may be due to the AMF Plus masters and core of the network being locked down with no Internet access. If this is the case a local NTP server, or AMF Plus node which does have Internet access, can be configured as the desired NTP server.

When you have multiple AMF Plus masters, the masters will act as NTP peers of each other and all other nodes will use the masters as NTP servers. This happens automatically; you do not have to configure it.

The primary function of NTP within an AMF Plus network is to ensure that date stamps on backups are consistent across member nodes. In a network that has multiple AMF Plus master nodes, it is particularly important to ensure that node recovery is performed with the most up-to-date backup. It is a good idea to set the **time zone** to be the same on all AMF Plus nodes.

Configuring NTP on the AMF Plus network

On all AMF Plus masters, you should configure **three** external NTP servers. If this is not possible, because the masters are not connected to the Internet, then at least one node connected to the Internet should be configured with NTP. The masters can then be configured to use these nodes, with Internet access, as their AMF Plus server.

Note: AMF Plus masters act as NTP peers of each other; all other nodes use the masters as NTP servers. This happens automatically, so you do not have to configure it.

For example:

```
awplus(config)# ntp server 1.pool.ntp.org
awplus(config)# ntp server 2.pool.ntp.org
awplus(config)# ntp server 3.pool.ntp.org
```

You can check that nodes have synchronized with the NTP server using the **show ntp status** command, for example:

Output 6: Output from the **show ntp status** command

```
awplus#show ntp status
associd=0 status=0618 leap_none, sync_ntp, 1 event, no_sys_peer,
system peer:      10.37.109.1:123
system peer mode: client
leap indicator:   00
stratum:          4
log2 precision:   -18
root delay:       32.810
root dispersion:  159.658
reference ID:     10.37.109.1
reference time:   db5f5f4e.94ac8ebe Thu, Feb 18 2023 10:10:22.580
system jitter:    0.482072
clock jitter:     0.366
clock wander:     0.247
broadcast delay:  0.000
symm. auth. delay: 0.000
```

Special Considerations when Using LACP Aggregations as AMF Plus Links

Using LACP aggregations as AMF Plus links requires specific default behavior on the part of AMF Plus nodes.

AMF Plus requires that a completely unconfigured node will successfully form an AMF Plus connection and become integrated into the network when attached to an AMF Plus network.

If the unconfigured node is attached to the network by an LACP aggregation it must be possible for the unconfigured node to form an LACP aggregation.

By default AMF Plus nodes need to recognize when the connected ports on a neighbor device are dynamically (LACP) aggregated, and then to negotiate an aggregated link with that neighbor's ports.

Specific functionality is available in AlliedWare Plus to support this default behavior. It is called **LACP global passive mode**.

LACP global passive mode

AlliedWare Plus devices can self-configure LACP channel-groups dynamically when they are connected to another device that has LACP channel-groups configured with active mode.

When a device starts with factory default configuration (or the start-up configuration file is missing), LACP global passive mode is automatically turned on. This is useful if you want to attach a new device to an existing LACP configured network, because the newly added device will then automatically form LACP channel-groups.

This feature can be turned on or off by the following CLI commands in global configuration mode:

- `lacp global-passive-mode enable`
- `no lacp global-passive-mode enable`

The current configuration setting is displayed by using the command **show running-config**.

Dynamically learned LACP channel-groups behave the same as manually configured ones (that are configured by the **channel-group** command). The only exception is that dynamically learned LACP channel-groups are not displayed in the running configuration. Currently known—both dynamically created and manually configured—LACP channel-groups are displayed by entering the following commands:

- `show etherchannel`
- `show etherchannel detail`

Dynamically learned LACP channel-groups

A dynamically learned LACP channel-group will be removed from the port, in any of the following situations:

- LACP global passive mode is turned off
- the port is removed (hot-swapped out)
- the port is down
- the **no channel-group** command is executed on that port.

A dynamically learned LACP channel-group will become a normal, manually configured, LACP channel-group and appear in the running configuration, in any of the following situations if you:

- add any configuration in Interface Configuration mode of the aggregation or any member of the aggregation
- execute the **channel-group** command in any member of the aggregation, or
- add a new port to the aggregation.

Mixed LACP configuration (manual and dynamic)

When LACP global passive mode is turned on, we do **not** recommend using a mixed configuration in an LACP channel-group. A mixed configuration is one where some links are manually configured and others are dynamically learned in the same channel-group. Global passive mode is turned on using the **lACP global-passive-mode enable** command.

Sharing AMF Plus links with other network operations

AMF Plus links have special significance within the AMF Plus network. They are the links used to carry the AMF Plus management and control traffic flows. Moreover, the AMF Plus software includes its own algorithm for ensuring loop-free operation of the AMF Plus management VLANs that run over AMF Plus links.

However, despite the special significance of AMF Plus links, they are not used exclusively for AMF Plus communication, and are also able to participate in other aspects of the operation of the network. Specifically, they can also carry data VLANs, and therefore transport all manner of user data that is being exchanged within the network.

However, although AMF Plus does ensure loop-free operation of the AMF Plus management VLANs that operate over its AMF Plus links, it does not provide the same service to the data VLANs (including the native vlan if present) that may also be configured to use these links. Users are, therefore, responsible for protecting their data VLANs - either by explicitly avoiding VLAN loops by configuring EPSR, or by using the spanning tree protocol. In this respect the following should be noted:

- AMF Plus coexists with spanning tree, AMF Plus, so spanning tree will operate on AMF Plus links without adversely affecting the operation of the AMF Plus management VLANs.
- There is no restriction regarding the use of EPSR with AMF Plus. EPSR rings can coexist on ports that are also configured with AMF Plus links. See ["Using AMF Plus in EPSR Rings" on page 269](#) for information on supported EPSR topologies.

Reserved IP address range

Some of the AMF Plus-related communication that occurs between AMF Plus nodes is in the form of IP traffic. A class-B subnet is reserved for the use of this AMF Plus-related IP communication.

By default, the reserved range is the subnet 172.31.0.0/16. Addresses in this subnet must be reserved for AMF Plus and should be used for no other purpose. AMF Plus actually further divides this subnet into two /17 subnets, used for different purposes:

- 172.31.0.0/17 assigned to the AMF Plus management VLAN
- 172.31.128.0/17 assigned to the AMF Plus domain VLAN

It is possible to change the subnet used by AMF Plus, using the command:

```
atmf management subnet <a.b.0.0>
```

This command assumes that the subnet being allocated has a /16 netmask. AMF Plus will automatically further subdivide the allocated subnet into two /17 subnets:

- a.b.0.0/17 assigned to the AMF Plus management VLAN
- a.b.128.0/17 assigned to the AMF Plus domain VLAN

The new management subnet will not become effective until all members of the AMF Plus network have been updated and all units rebooted.

To return the subnet to the default 172.31.0.0/16, use the command:

```
no atmf management subnet
```

AMF Plus on VCStacks

If any VCStacks are included as AMF Plus nodes, the VCS virtual MAC feature should be enabled to ensure correct operation of the AMF Plus network. If the VCStack is running as an AMF Plus master node and is required to backup member nodes, then removable storage media must be installed in all stack members.

AMF Plus links on AR-series Eth interfaces

Up/down links and area links are supported over an AR-series device's Eth interface. This enables you to provision and recover AMF Plus nodes over these interfaces.

To use this feature your AMF Plus network must be in AMF Plus secure mode.

Use the **atmf-link** and **atmf-arealink** commands on an Eth interface to configure it as an AMF Plus link. For example, to configure an up/down link on Eth1 port, use the following commands:

```
awplus# configure terminal
awplus(config)# interface eth1
awplus(config-if)# atmf-link
```

If you run the command and AMF Plus secure mode is not enabled, you will see the following error message:

```
Node_1(config)#int eth1
Node_1(config-if)#atmf-link
% Cannot configure eth1 because atmf secure-mode is not enabled.
```

By default AMF Plus recovery is disabled on these links. Enable recovery by running the **atmf recovery over-eth** command in privilege exec mode.

```
awplus# atmf recovery over-eth
```

This setting persists even after restoring a device to a “clean” state with the **erase factory-default** or **atmf cleanup** command.

AMF Plus interaction with QoS and ACLs

It's important that ACL and QoS rules do not block the following traffic types:

- any traffic on VLANs 4091 and 4092 as they are the default AMF control VLANs.
- any traffic on the subnet which is reserved for AMF Plus management purposes - by default 172.31.0.0/16
- packets with protocol type 0xfbae.
- BPDUs that use the MAC address 0180.c200.002e.
- any IPv6 addresses in the range FD00:4154:4D46::/48 as these are used for inter-area communication.

With AMF Plus enabled the number of ACLs available on the x230, x550, x930, XS900MX, x950, and SBx908 GEN2 switches decreases by 1. If you are not using AMF Plus, you can disable AMF Plus to reclaim the additional ACL.

AMF Plus interaction with STP on AR-series devices

On AR-Series devices, if you use STP at the same time as AMF Plus, you may find that AMF Plus downlinks/cross-links on blocked STP ports remain in a blocking state (either a state of Up OneWay Blocking or Up RequestReset Blocking).

When (or if) STP unblocks the port, the AMF Plus link will fully synchronize and correctly forward AMF Plus traffic in both directions, but there may be a delay of a few seconds before it does so. This could interrupt any AMF Plus working-set operation that is in progress at the time.

Note that this potential for blocking by STP does not occur with AMF Plus virtual-links. Therefore, on these devices, we recommend connecting the device to the AMF Plus network via a virtual-link.

Renaming your AMF Plus network

To rename the AMF Plus network, use the command:

```
node_1(config)# atmf network-name <new-name>
```

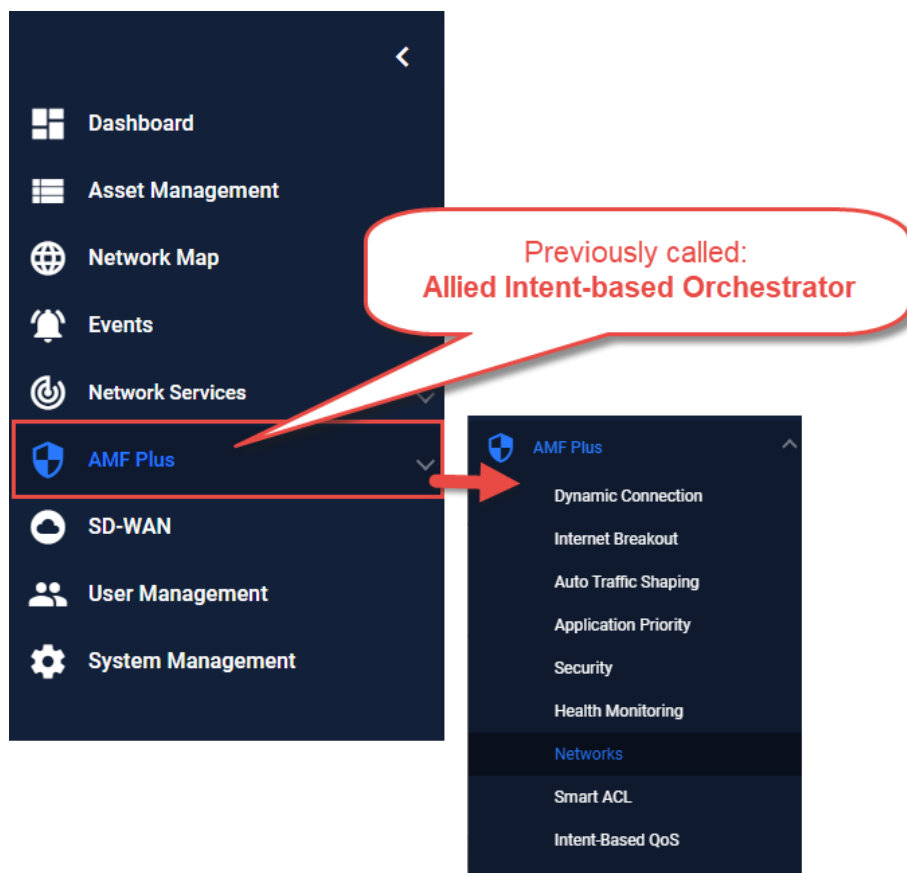
You do not need to reboot your device after changing the network name.

Using the AMF Plus Menu in Vista Manager

Allied Telesis Autonomous Management Framework™ Plus (AMF Plus) replaces the Allied Intent-based Orchestrator (AIO) menu when you have an AMF Plus license. AMF Plus provides network optimization, automation, management, and visualization. The uniquely designed intent-based configuration, reporting, and map facilities of Vista Manager EX make these powerful tools simple to configure, initiate, and manage. AMF Plus offers automation of branch security and WAN bandwidth management.

AMF Plus devices must first be configured in the CLI before they can run with Vista Manager EX. For information on how to configure your devices in the CLI, see "[Configuring AMF Plus to communicate with Vista Manager EX](#)" on page 21.

The **AMF Plus** menu is located in the left-hand menu.



Note: AMF licenses do not include Vista Manager's AMF Plus features. To access these features, you need to upgrade from AMF to AMF Plus.

For the AMF Plus feature to become fully available to you and for all menu items to be activated, install the feature license. The AMF Plus license is not part of the base Vista Manager EX license, but a trial is included in the 90-day trial license.

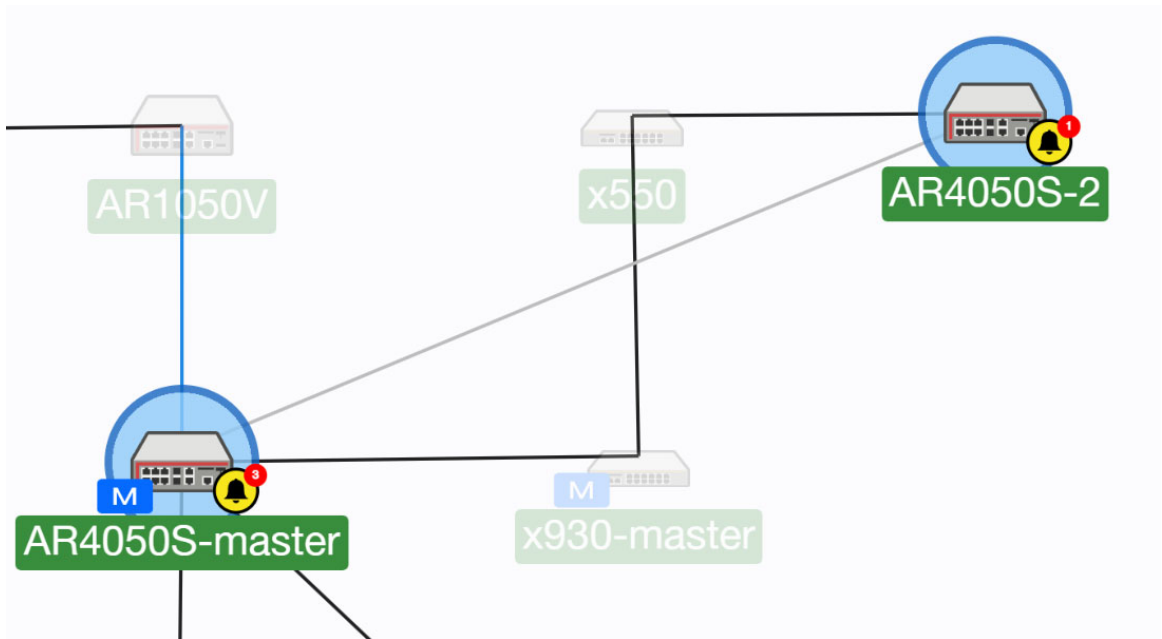
If you want to continue using AMF Plus after the 90-day trial license expires, you need to install a feature license for it. Contact your authorized [Allied Telesis salesperson](#) for assistance.

For more information about Vista Manager, see the [Vista Manager EX User Guide](#).

The AMF Plus menu is made up of several tools to help you manage your network:

- ["Dynamic Connection" on page 55](#)
- ["Internet Breakout" on page 61](#)
- ["Auto Traffic Shaping" on page 64](#)
- ["Application Priority" on page 66](#)
- ["Security" on page 71](#)
- ["Health Monitoring" on page 75](#)
- ["Networks" on page 86](#)
- ["Smart ACL" on page 87](#)
- ["Intent-based QoS" on page 94](#)

Dynamic Connection



The Dynamic Connection map lets you click-and-drag to create new VPN tunnels between the AR-Series devices (firewalls or routers) at different locations across your WAN.

There are two types of tunnels you can create on the Dynamic Connection map:

- Point-to-point tunnels require a source device and destination device.
- Point-to-multipoint tunnels require a source device and multiple destination devices.

Note: The AR1050V does not support Dynamic Connection.

For this feature to be fully functional, apart from installing the AIO license, you must also have either administrator access or write permission on a device.

To create a tunnel, both devices must be part of the AMF network, support GRE tunnels and be running firmware version AlliedWare Plus 5.5.0-2.x or later.

You cannot create multiple tunnels with the same source and destination interface pair (e.g. eth1). Split up the interface if you wish to create more than one tunnel, for example, split eth ports into sub-interfaces. You may create another tunnel with the same source interface as long as the destinations are on different devices.

All tunnels are encrypted with IPSec to secure your WAN traffic. Each tunnel will have a different crypto key with a unique name.

Distributed tunnel routing

When you create a tunnel, you can choose to distribute routes to additional devices in order to create a return routing path.

You will see a list of subnets to choose from, with these subnets being accessible from the device. However, not all networks and devices at the tunnel destination are used to form new primary routes. The list of destinations are pre-filtered.

The following types of networks and hosts are allowed:

- connected by static routes
- directly connected to the end router (direct routes)
- routed through a dynamic routing protocol

Example: When a tunnel is created from (A) to (B), (A) will distribute networks and hosts (X) to (B). However, that does not necessarily mean (X) can reach (B), so networks on (B) are allowed to be distributed to add as routes on (X).

If a tunnel is deleted, all static routes associated with the nexthops of that tunnel will also be deleted. However, manual routes can still be added from the pull-down menus.

Administrative distances are added to static routes; static routes with the same default administrative distance (zero) to the same destination is not supported. When a route is shared, Vista Manager adds a 1 to its distance. Therefore, a direct connection route with a default distance of 0 will have a distance of 1 when added to a destination device's route table.

For this feature to be fully supported, AlliedWare Plus version 5.5.1-2.1 or later is required.

Distribute Routes Example

Settings for the source end of tunnel (Auckland)

1. The route to **Auckland (1.0.0.0/8)** is selected in the "Distribute Routes" input. This route is added to the route table of the Christchurch device, allowing traffic to go from Auckland to Christchurch.
2. A tunnel between Auckland and Waimate already exists, so the route to **Waimate (2.0.0.0/8)**, is an option in the "Distribute Routes" input. This route is added to the route table of the Christchurch device, allowing traffic to go from Waimate to Christchurch.
3. Nothing is needed in the "Distribute routes to devices" input because the selected routes are automatically distributed to the destination end of the tunnel (Christchurch).

Settings for the destination end of tunnel (Christchurch)

4. The route to **Christchurch (3.0.0.0/8)** is selected in the “Distribute Routes” input. This route is added to the route table of the Auckland device, allowing traffic to go from Christchurch to Auckland.
5. Because the route to Waimate is added to the route table of the Christchurch device, there is now an option to distribute a route to Christchurch on the Waimate device. This route is added to Waimate, allowing traffic to go from Christchurch to Waimate.

Note: It is mandatory to choose a route. Vista Manager is unable to prevent loops from being created as all forwarding paths in the network are not known. Some WAN-facing interfaces will not be included in the list of routing destinations, as this could form routing loops caused by networks beyond the immediate control of the user.

Feature limitations

There are some feature limitations to take note of:

- Because this is adding static routing, there may be potential for routing loops. The risk of causing such loops cannot be eliminated.
- Entity subnets will not be filtered out if they overlap or are duplicated with other subnets. It is up to the user to create valid entities.
- Changes made to subnets and entities after the tunnel has been created will not be automatically deleted; routes on the devices will not be updated. Users will have to make these changes on the tunnels and devices if they make changes to subnet and entities.
- IPv6 routes are supported as static routes, but are not supported as distributed subnets. The IP version of static routes must match the IP version of the tunnel IP address.
- mGRE tunnels use GRE-based protocols and are therefore stateless. Static routes on mGRE will not be re-routed automatically if a hub-to-spoke tunnel link goes down.

How to create a Point-to-point tunnel

To create a Point-to-point tunnel, do the following:

1. Use the pencil icon to draw a line between devices (firewalls/routers) at the two locations you wish to connect with a new VPN tunnel.
2. Next, set up tunnel options. Select tunnel mode.
3. Select an interface for the tunnel to be on.
4. Vista Manager EX generates the tunnel interface IP addresses. The subnet prefix is /30.

Note: If you choose your own IP address, it must be in the same subnet and must not be used on another interface on those devices.

5. Enter a description name for the tunnel.
6. Configure tunnel routing.

Note: The options here are default or static. You may enter IP addresses for each end of the tunnel by selecting static routing.

7. Repeat steps 3-6 to set up tunnel options for the destination device.
8. Click **Check connectivity**. There should be a ping from source interface to destination interface if there is a connection.
9. Click **Create** when complete.

1

2

3

4

5

6

7

8

9

Create Point to Point Tunnel

Tunnel protection type is IPSec

Tunnel Mode
GRE

AR4050S
Tunnel Interface *
eth1 ()

Tunnel IP Address *
172.28.0.1 /30

Tunnel Description
AKL-source

Tunnel Routing
No Routing

AR2010V
Tunnel Interface *
eth2 ()

Tunnel IP Address *
172.28.0.2 /30

Tunnel Description
NSN-destination

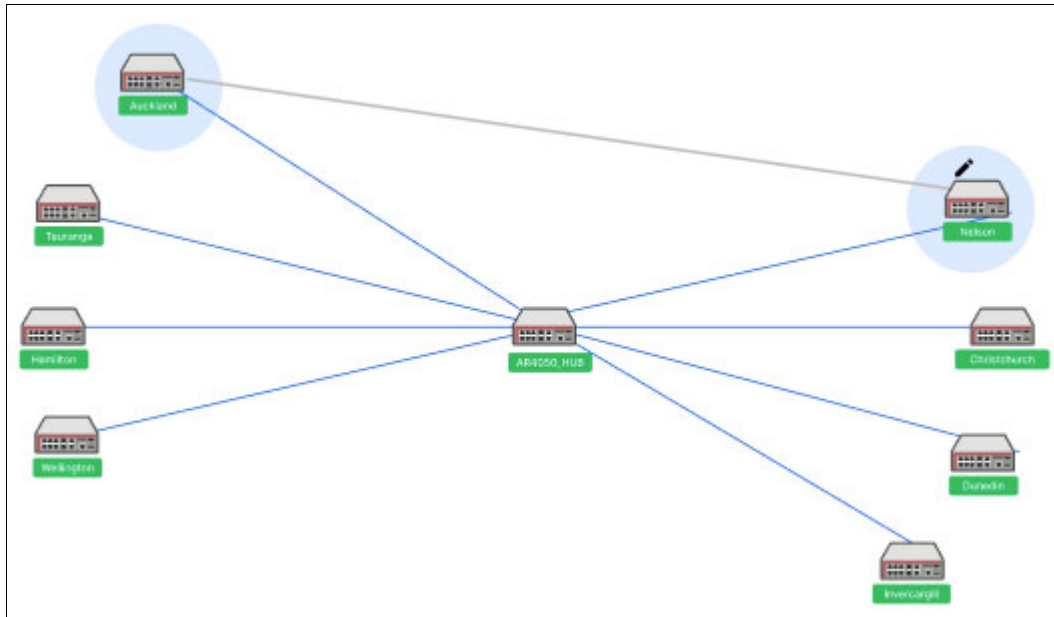
Tunnel Routing
No Routing

Check connectivity

Cancel Create

How to create a Point-to-multipoint tunnel

To create a Point-to-multipoint tunnel, do the following



1. Click on the pencil icon and select point-to-multipoint tunnel.
2. Use the pencil icon to first select a tunnel hub. This is usually a head office router.
3. Next, select spokes one by one. These should be your branch offices.
4. Perform Option 1 steps 2-6 to set up tunnel options.

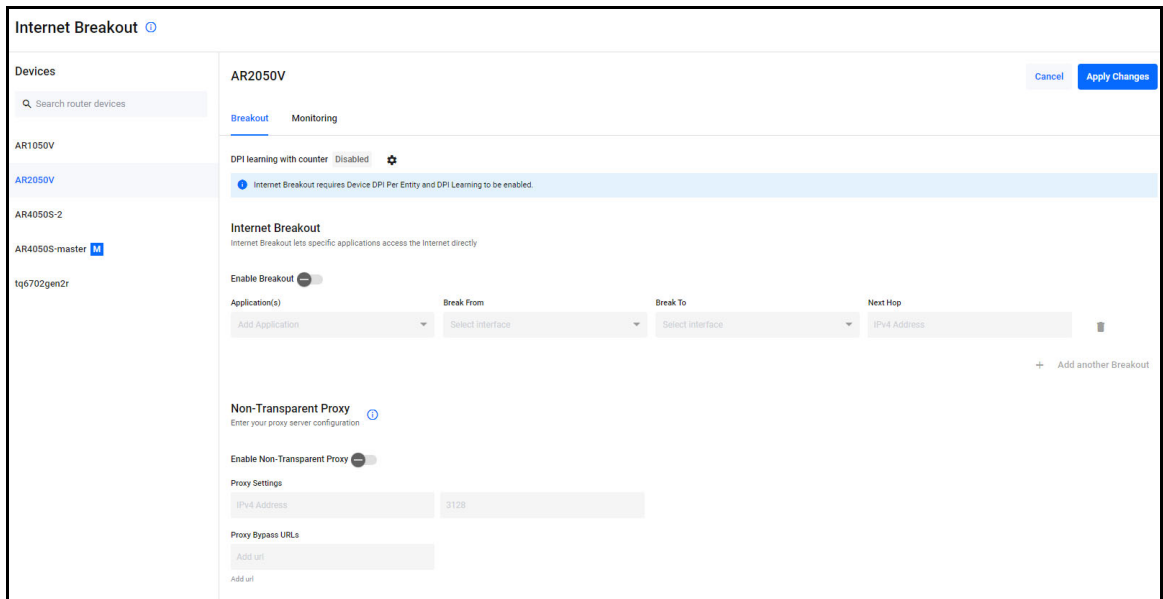
Note: In version 3.5.0, adding a static route to the hub of a multipoint tunnel is not supported.

5. Repeat for all your spokes (branch offices).
6. Click **Check connectivity**.
7. Click **Create** when complete.

Note: For multipoint tunnels, a hub of multipoint tunnel cannot share the same interface (with the same IP address) as a GRE point-to-point tunnel.

Note: Connectivity is not needed for the new tunnel configuration to be created, although the tunnel will not be fully formed until there is a connection.

Internet Breakout



Internet Breakout lets specific applications being used at branch office locations, access the Internet directly, rather than going via the head office. This improves the performance of cloud-based applications (e.g. Office 365) and reduces traffic volumes on VPN connections between branch offices and the head office.

Before configuring, start by identifying the types of applications you may want to allow direct Internet access.

Note that this feature requires AR-series devices to run AlliedWare Plus 5.5.0-2.1 or later.

- Internet Breakout requires Device DPI Per Entity and DPI Learning to be enabled.
- Enabling this feature reduces router throughput.

Caution: Any traffic that bypasses security processing may reduce security and threat protection at the local branch office. Carefully consider the potential consequences of giving direct Internet access to a type of traffic, and whether additional local or cloud-based security needs to be implemented to protect Internet Breakout traffic and the branch office.

- Internet Breakout needs to classify applications for sending direct to the Internet. It does this most effectively when it can read both incoming and outgoing traffic on the interface that was/is sending those applications to the head office. For IPSec protected tunnels, this requires a feature called tunnel security reprocessing. Vista Manager does not enable tunnel security reprocessing because it reduces router performance.

- To enable tunnel security reprocessing, enter the following commands on the router's CLI:

```
enable
conf t
tunnel security-reprocessing
```

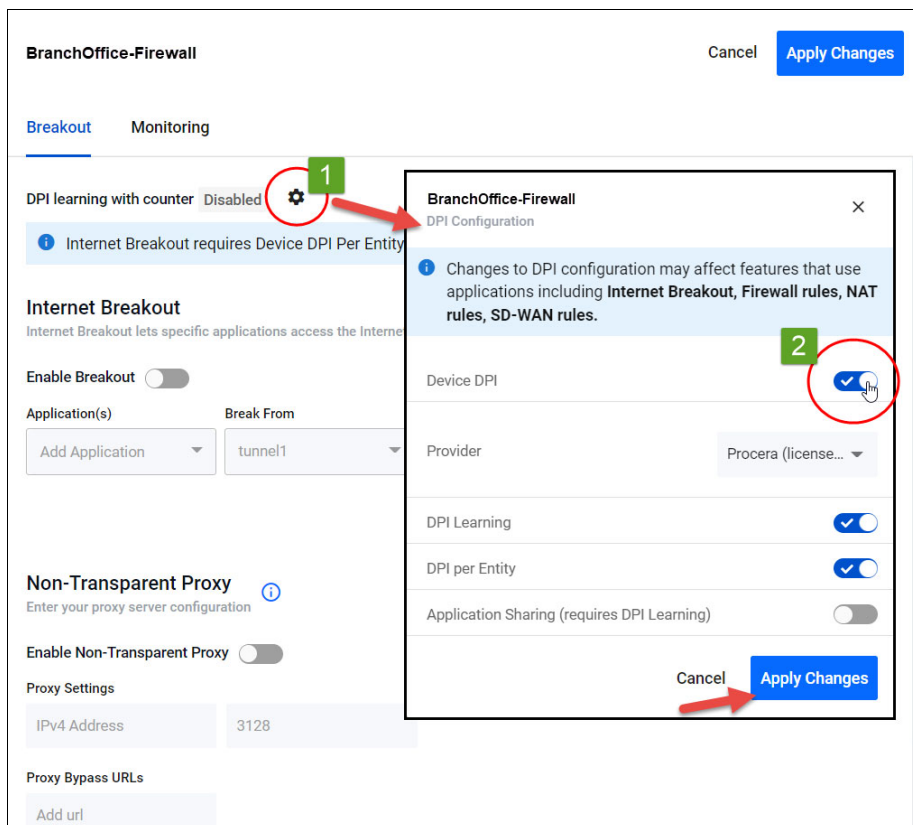
How to Enable Internet Breakout

Step 1. Enable Internet Breakout and specify the traffic path for applications

By default, Internet Breakout inputs are disabled until **Breakout** or **Non-Transparent Proxy** is enabled. If invalid options are selected followed by disabling Internet Breakout, these options will be removed. This prevents saving a disabled but invalid configuration.

Use the **Internet Breakout > Breakout** tab and select a device:

1. Click the **Settings** icon.
2. Enable **Device DPI** (save any changes).



3. Select **Enable Breakout**
4. Add **Applications** to the Breakout List, for example, Office365, Google, Youtube, etc.
5. Select the interfaces to **Break from** and **Break to**.
 - To add another break from or break to interface, click + **Add another Breakout** and repeat steps 4, 5, and 6.
6. Enter the **Next Hop** address (optional). If 'tunnel' is selected as 'break from', then next hop is disabled.
7. Enable and configure the **Non-Transparent Proxy** settings (optional).
8. Click **Apply Changes**.

BranchOffice-Firewall Cancel **Apply Changes**

8

Breakout Monitoring

DPI learning with counter Running: Server

Internet Breakout

Internet Breakout lets specific applications access the Internet directly

Enable Breakout **3**

Application(s)	Break From 5	Break To	Next Hop 6
<input type="text" value=""/> google <input type="button" value="X"/> 4 office <input type="button" value="X"/> youtube <input type="button" value="X"/>	tunnel1 <input type="button" value="X"/>	eth1 <input type="button" value="X"/> eth1 eth2 tunnel1 tunnel10 tunnel2	10.0.0.2 <input type="button" value="X"/>

[+ Add another Breakout](#)

Non-Transparent Proxy

Enter your proxy server configuration

Enable Non-Transparent Proxy **7**

Proxy Settings

12.22.34.45	3128
-------------	------

Proxy Bypass URLs

Add url

Add url


- www.google.com
- www.microsoft.com

Step 2. Monitor the breakout

Use the Internet Breakout > Monitoring tab

Two charts are available here:

- The pie chart shows the top 5 breakout applications. Clicking applications on the vertical legend adds/removes them to/from the chart.
- The line graph shows breakout traffic over a set period of time. Clicking applications on the horizontal legend or using the drop-down list adds/removes them to/from the graph.



Auto Traffic Shaping

Auto Traffic Shaping dynamically adjusts the maximum transit capacity of remote locations (spoke tunnels) to not exceed the receive capacity of the central site (hub). This is termed the **maximum Rx bandwidth** of the hub.

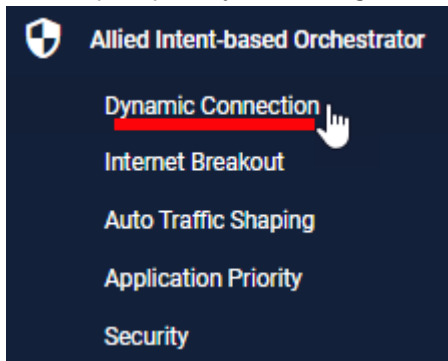
To allocate this bandwidth optimally, we recommend you also deploy Application Priority profiles on each spoke tunnel.

To manage traffic, an algorithm uses current spoke tunnel traffic rates, and any configured application priority settings, across all spoke tunnels to fairly allocate bandwidth. Spoke tunnels have a guaranteed transmit bandwidth. This equals the sum of the CIRs (committed information rate) plus system bandwidth defaulted to 5%.

Prerequisite step: Configure tunnels between spokes and hubs.

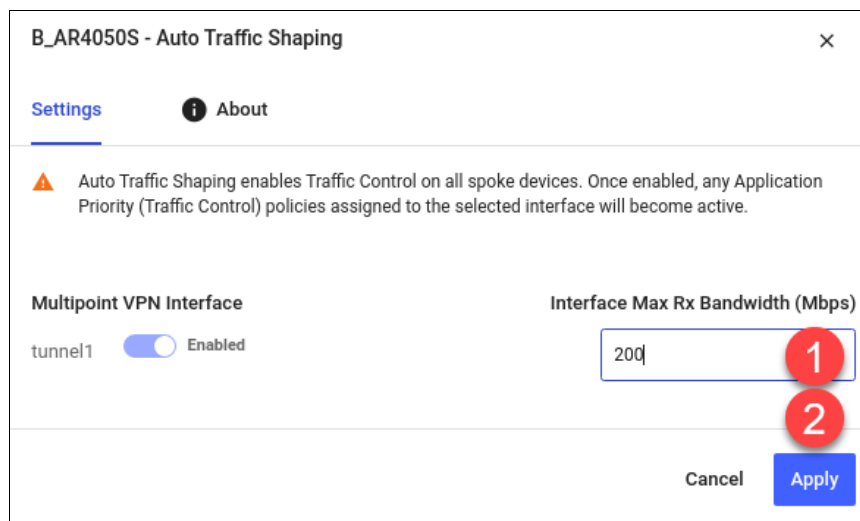
Use the Allied Intent-based Orchestrator (or AMF Plus) > Dynamic Connection feature

This step requires you to navigate away from Auto Traffic Shaping.

**Step 1: Configure the Interface Max Rx Bandwidth value.**

Use the Auto Traffic Shaping > Settings tab or button


1. Enter the maximum bandwidth a hub can handle. The algorithm calculates and applies optimal traffic shaping based on this number.
2. Click **Apply**.



Step 2: Monitor hub utilization and traffic loss.

Use the **Auto Traffic Shaping > Monitoring** tab

View charts in the Monitoring page, where you can use filters to specify what traffic is shown.



Application Priority

You can use Application Priority to choose specific applications and prioritize or deprioritize them. This ensures your most important business traffic is prioritized for transmission between locations across your WAN. Vista Manager EX provides 3 priority classes:

1. **Critical Services**
2. **Daily Operations**
3. **Non-Essential**

You can assign different applications to each priority class, save the assignment in a policy, and deploy the policy on the AR-Series device (firewall or router) at each location in your WAN.

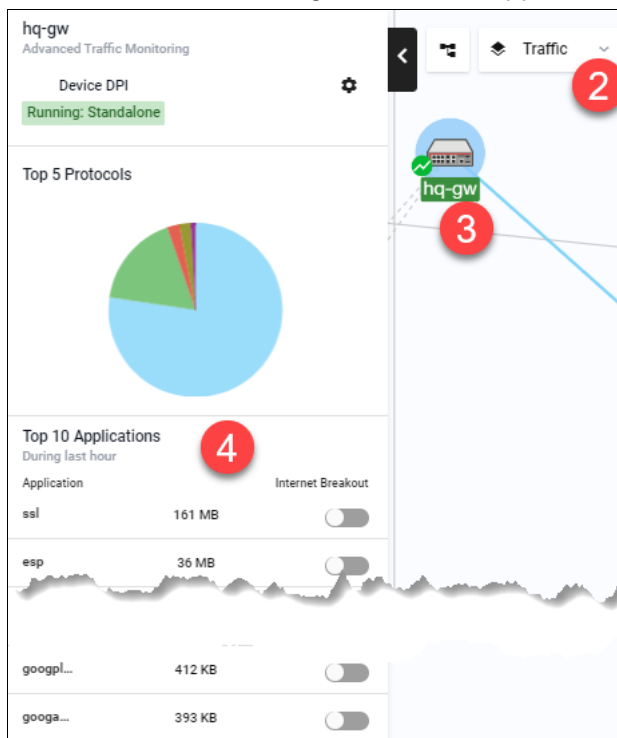
A policy is the overall title for a set of rules and priorities. It also defines the type of algorithm for how it calculates the priority of traffic. Traffic for any unassigned applications set in the rules will fall into the **Default** policy class. The default class is not directly visible when creating a policy, but you can see the traffic matching the default class (either in throughput or packet loss) in the Monitoring graphs.

This feature lets you view any existing Application Priority policies, and shows throughput and packet loss graphs for devices that have a policy deployed on them. You can also see how much guaranteed bandwidth each class has and how much shared bandwidth remains. When the network is congested, use the slider or advanced option to set bandwidth requirements to ensure smooth application traffic.

Vista Manager EX application usage data lets you better prioritize applications. When creating or deploying policies, you can analyze current traffic present on a device, which helps you assign applications into the most appropriate priority classes for a policy.

How to check application usage on a device

1. Navigate to the Network Map.
2. Select **Traffic** mode from the drop-down list.
3. Select the device you want to check. A blue circle appears around it.
4. Examine the traffic usage data, which appears in the left-hand panel.



How to create an Application Priority policy

1. To create a policy, you may:
 - Right-click on the device in Traffic mode of the Network Map and select: **Application Priority > Add Policy**, or
 - Navigate to the Application Priority menu item and click **+Add Policy**, or
 - Navigate to the Application Priority menu item, and **clone** an existing policy by clicking the 3 dots for that policy, in the Action column.

All of the above approaches open the **Add Policy** page.

2. Next, type in a policy name. For example, **Branch-Office**.
3. Select an Application Provider. By default, Built-in is selected. If you have bought an Advanced Firewall licence for your AR-Series UTM firewall, select Procera instead, which enables a much larger application list to work with.
4. On the right-hand panel, choose a Category of applications. For example, **Remote Access**. A list of applications will appear.
5. Assign appropriate classes to the relevant applications. You can use the Assigned Class filter at any point to see what applications you have assigned to a class. When you assign a class, it appears accordingly on the policy classes on the left.
6. Here, you may adjust the bandwidth for each class. To do this, either **move the slider** or **enable advanced bandwidth adjustment** to type in the percentage. Percentages will be converted to Mbps values when deployed to device. If the advanced option is used after the slider, any manually-set values are automatically replaced by the slider pre-sets.

Note: The reserved percentage of guaranteed bandwidth for system traffic is displayed here. The 5% value is based on the default value that traffic control sets on a device. The actual value may vary depending on what device(s) the user deploys the policy on to. Vista Manager will just show **5% as the system bandwidth**.
7. If you have accessed the page via the Network Map, click **Save and Deploy**. Otherwise, click **Save**.

← Add Policy
Cancel **Save**

Policy Name * 2

Give your policy a descriptive title E.g. AppCategory_RequiredMbps

Application Provider 3

Built-in Procera (license required)

Application Priority
Assign applications to give priority over lower classes and default traffic

1 Critical Services
Applications critical to business operations.
Include services whose disruption would result in a high cost. E.g. Database & Backup.

Applications Clear All

citrix rdp ssh

2 Daily Operations
Applications used in day-to-day business operations.
E.g. File sharing.

Applications Clear All

teamviewer telnet

3 Non-Essential
Applications that are commonly used but not essential to business operations.
E.g. Social media.

Applications Clear All

pcanywhere

Guaranteed Bandwidth
Control minimum bandwidth requirements to ensure application traffic when network is congested.
Percentages will be converted to Mbps values when deployed to device.

25%

25%
⚙️

<input type="radio"/>	System		5 %
<input checked="" type="radio"/>	Critical Services	- +	4 %
<input checked="" type="radio"/>	Daily Operations	- +	10 %
<input type="radio"/>	Non-Essential	- +	6 %
<input checked="" type="radio"/>	Shared Bandwidth		75 %

6

Assign Applications 7

Category: Remote Access ▾ Assigned class: All ▾

Search applications

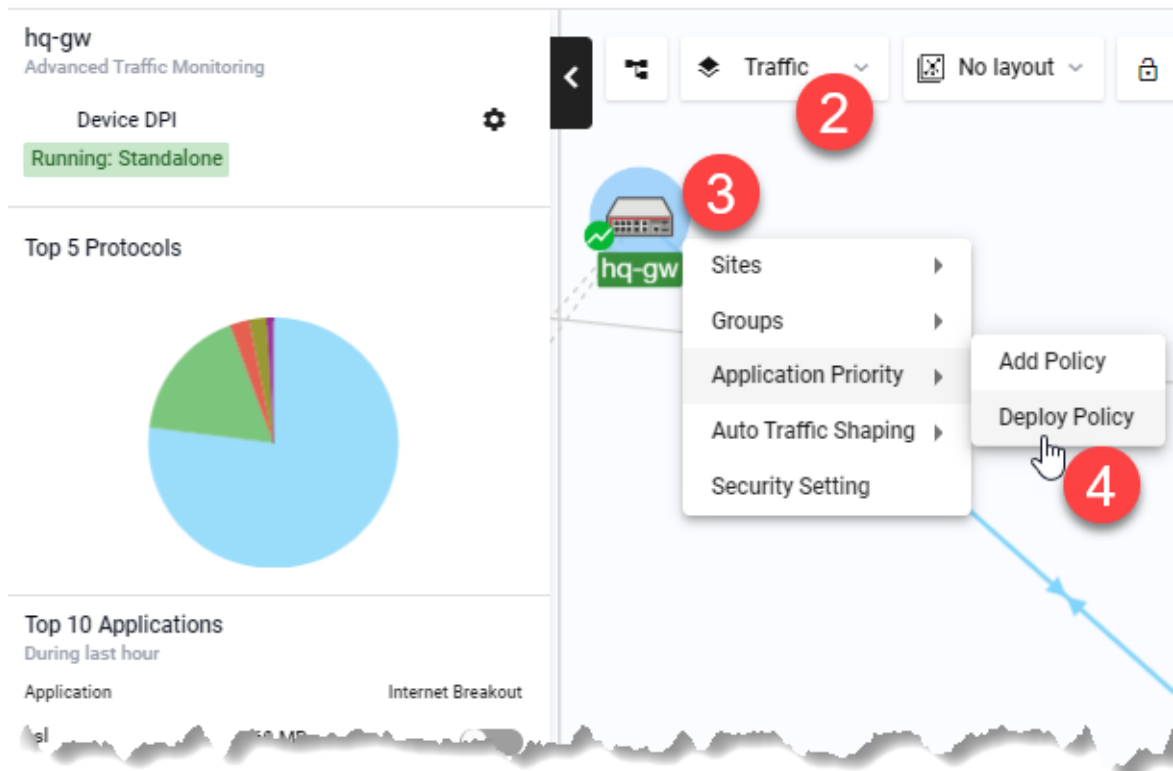
Applications	Assigned to class
citrix	Critical Services ▾
pcanywhere	Non-Essential ▾
rdp	Critical Services ▾
ssh	Critical Services ▾
teamviewer	Daily Operations ▾
telnet	Daily Operations ▾
vnc	Unassigned ▾

1 to 7 of 7 << < Page 1 of 1 > >>

How to deploy a policy

First, create a policy in the section above.

1. Navigate to the Network Map.
2. Select **Traffic** mode from the drop-down list.
3. Select the device you want to deploy a policy to. A blue circle appears around it.
4. Right-click on the device, select Application Priority > **Deploy Policy**.



5. Select a policy to deploy.
6. Specify a **Source Entity** to match traffic against.
7. Specify an interface for **Destination Entity**.
8. Define a **maximum bandwidth**. This places a cap on the virtual bandwidth.

9. Click **Deploy Policy** when complete.

← Deploy Policy Cancel **Deploy Policy**

hq-gw
Advanced Traffic Monitoring

Device DPI
Running: Standalone

Top 5 Protocols

Top 10 Applications
During last hour

Application	Internet Breakout
ssl	160 MB <input type="checkbox"/>
esp	36 MB <input type="checkbox"/>
icmp	5 MB <input type="checkbox"/>
udp	2 MB <input type="checkbox"/>
eth	1 MB <input type="checkbox"/>
facebo...	967 KB <input type="checkbox"/>
supercel	500 KB <input type="checkbox"/>
tcp	470 KB <input type="checkbox"/>
googpl...	449 KB <input type="checkbox"/>
googa...	393 KB <input type="checkbox"/>

Device: hq-gw

Source Entity: VM_Zone.BREAKOUT_TO

Destination Entity: VM_Zone.BREAKOUT_FROM (tunnel1)

Select Policy: test-deploy

Destination Max Bandwidth (Mbps): 100

test-deploy

- Critical Services**
Applications critical to business operations.
Include services whose disruption would result in a high cost. E.g. Database & Backup.
Applications
- Daily Operations**
Applications used in day-to-day business operations.
E.g. File sharing.
Applications
circuit
- Non-Essential**
Applications that are commonly used but not essential to business operations.
E.g. Social media.
Applications

Security

The security feature lets you configure the web control and IP reputation features on the UTM firewalls at a number of locations simultaneously, for centralized and simplified management.

- **Web control** offers an easy way to monitor and control the types of websites viewed by employees.
- **IP reputation** blocks employee access to websites that are known source of spam, viruses and other malicious activity, to protect your network against security threats.

The overall security feature allows you to enable recommended security settings for a group of UTM firewall devices based on an industry type and security strength. This simplifies the process as there is no need to manually choose website or reputation categories for each device.

Note: For this feature to be fully functional, you may need to do additional configuration in the device GUI. Internet access and domain name lookup are required. Enable the ATL Live update server in order to download and check for IP reputation or web control updates.

Step 1: Enable security features and select industry settings.

Use the Security > General tab

1. Enable **IP Reputation** and **Web Control** for the desired device group(s).
2. Select a Security provider (e.g. Digital Arts).
3. Select the industry type (e.g. High School).
4. Set a time to check for updates.
5. Select the desired security strength for the industry (e.g. Medium).
6. Click **Apply Changes**.

CanterburyHigh 6 [Apply Changes](#)

General Advanced IP Reputation Advanced Web Control Monitoring

Security Features

IP Reputation

Web Control 1

Web-control (License Required)

Digital Arts Webroot 2

Industry

High School x 3

Check for updates every

24 hours 4

Industry Strength

Low Medium High 5

Step 2: Edit advanced IP reputation settings if required.**Use the Security > Advanced IP Reputation tab**

1. Click **Edit Industry Settings**.

CanterburyHigh 3 **Apply Changes**

General **Advanced IP Reputation** Advanced Web Control Monitoring

Industry Strength Check for updates every

Low **Medium** High 24 hours

Industries 1 **Edit Industry Settings**

Custom

Infant School View Details

Elementary School View Details

Junior High School View Details

High School View Details

2. **Permit, alert, or deny** a reputation category action as needed.
A warning appears if one or more devices in the group have had different IP reputation or Web Control settings already applied via another group.

IP Reputation	×
Abused TLD	Permit Alert Deny
Bitcoin Related	Permit Alert Deny
Blackhole	Permit Alert Deny
Bot	Permit Alert Deny

3. Click **Apply Changes**. This changes the industry type to Custom.

Step 3: Edit advanced Web Control settings if required.

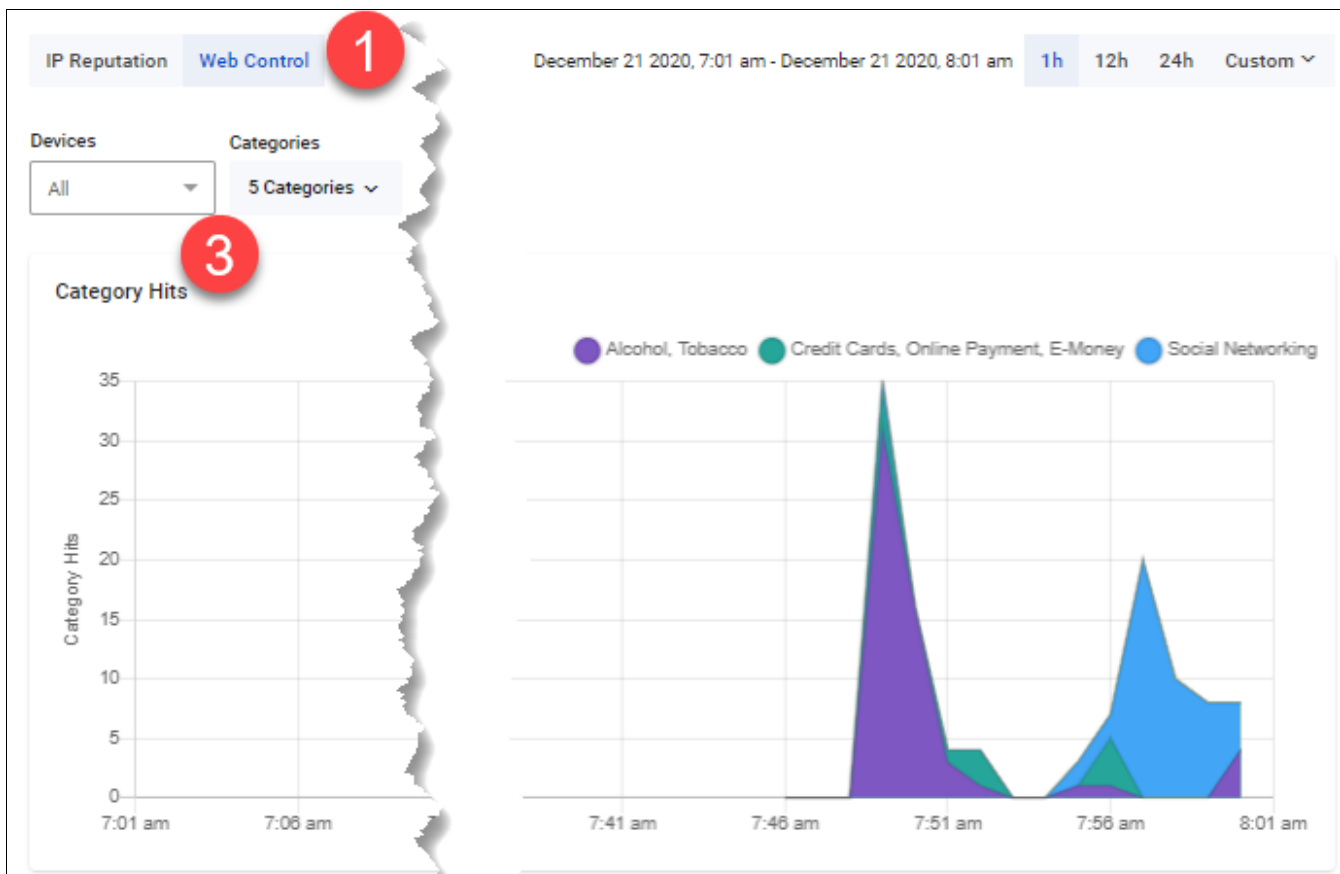
Use the Security > Advanced Web Control tab

1. Click **Edit Industry Settings**.
2. **Permit** or **deny** website categories as needed.
A warning appears if one or more devices in the group have had different IP reputation or Web Control settings already applied via another group.
3. Click **Apply Changes**. This changes the industry type to Custom.

Step 4: Monitor Web Control and IP Reputation performance.

Use the Security > Monitoring tab

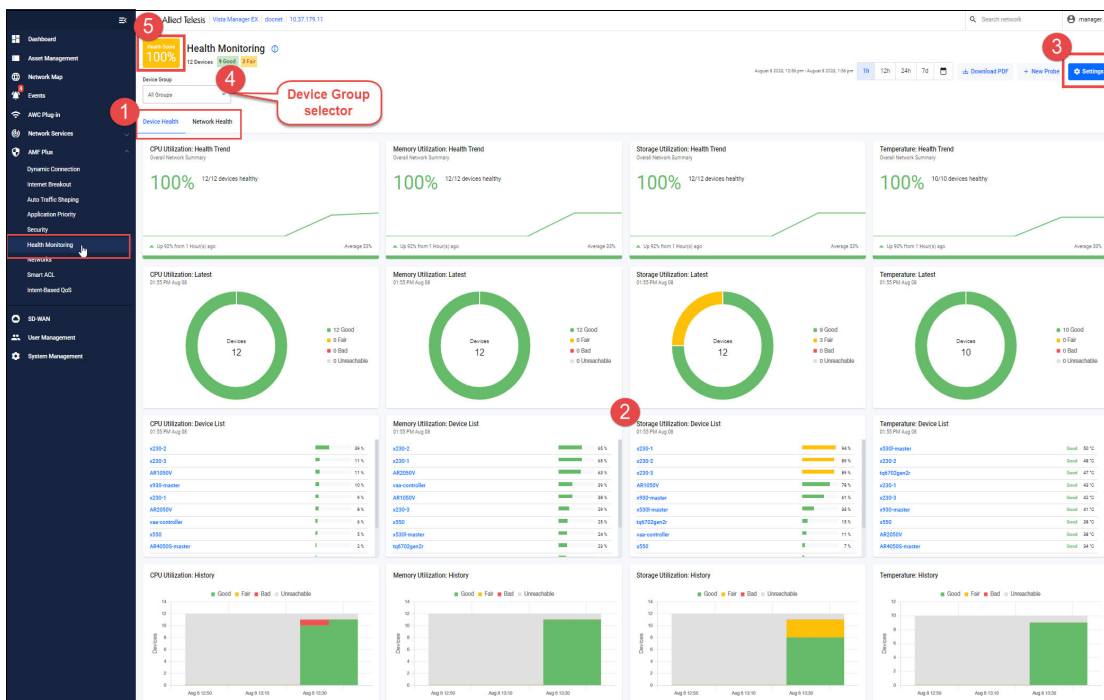
1. Click on the **IP Reputation** and **Web Control** buttons to view respective graphs.
2. For IP Reputation, click the drop-down list to select the UTM firewall device from a specific location to view.
3. For Web Control, click the drop-down list to select the UTM firewall device from a specific location to view. Clicking Categories on the legend or drop-down list lets you add/remove categories to view on the graph.



Health Monitoring

You can use the Health Monitoring feature to view a summary of the state of your network's health as part of AMF Plus. Understanding network health indicators enables you to investigate, analyze, and improve the overall health of your network quickly. Such indicators include CPU utilization, storage, temperature, and memory usage.


- There are two key indicators (tabs):
 - Device Health:** hardware state (CPU, memory, storage, temperature)
 - Network Health:** traffic and related errors (traffic health, interface counters)
- Each key indicator has the following widget types: Health Trend, Latest, Device List, History, and Probes (for Network Health only).
- Use the **Settings** panel of the dashboard to configure the key indicator threshold values. These thresholds apply to the entire network and are used to determine the health status of devices and interfaces.
 - Use the **Device Group** selector to display a group. The default is All Groups (i.e. all known devices). Auto-generated and user generated groups are located under the **Asset Management** menu.



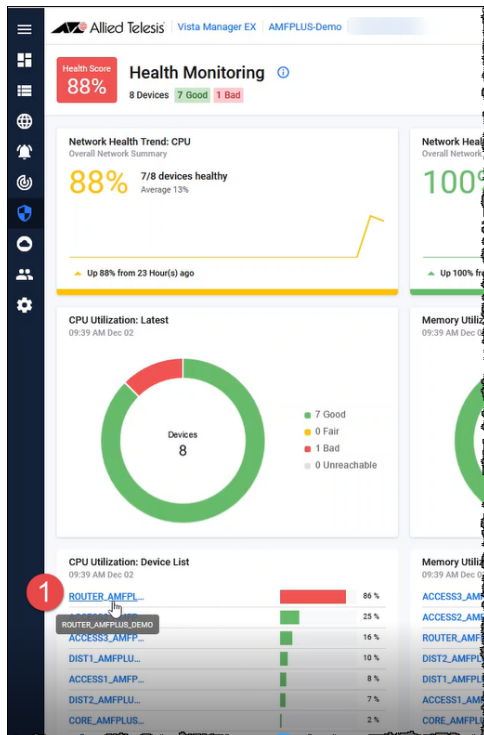
- The **Health Score** (located at top left corner) is a percentage based on how many devices are healthy in the network. The state of each device is selected based on the worst state of any of the gathered statistics. Result charts are color coded for easy understanding of device status: Green = Good, Yellow=Fair, Red=Bad, and Grey= Unreachable.

Device Health tab

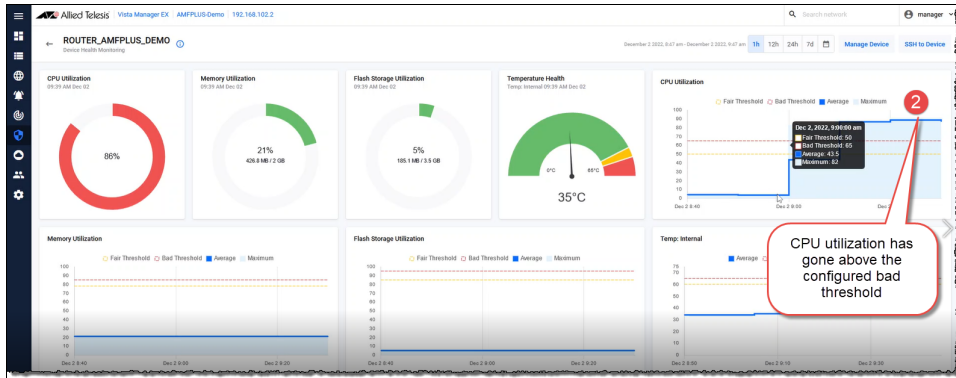
In the example below, you can see a Health Score of 88%. There are 8 devices in this network, but there's a CPU issue with one of them. The bad device is highlighted in red.



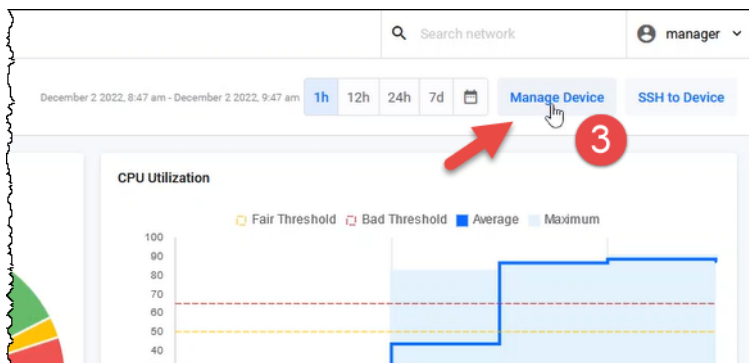
1. Click on the 'bad' device name to investigate further.



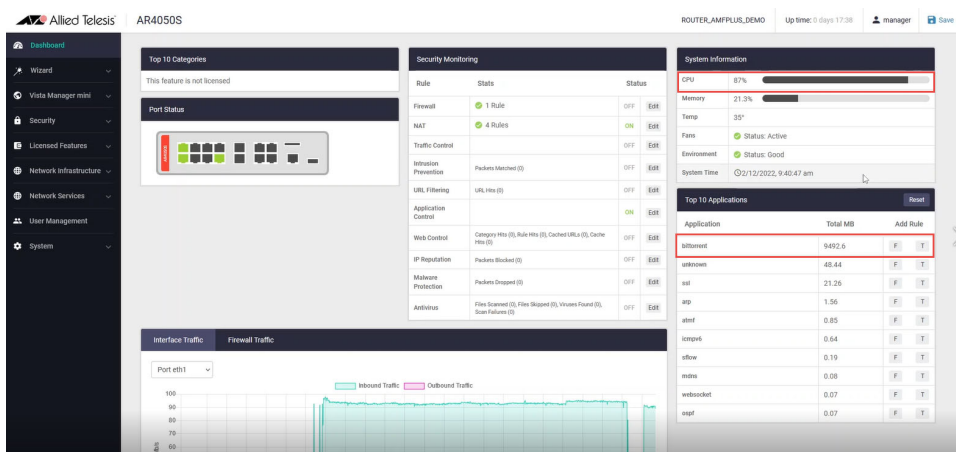
- Drilling down confirms that around 9am CPU utilization rose above the configured band threshold.



- To further diagnose the issue, click on **Manage Device** to open the device's GUI.



For this example device (AR4050S), the system information indicates a very high CPU usage and the applications show the bittorrent traffic increasing quite rapidly...which is likely to be the cause of the high CPU utilization.



- At this point you may decide to disallow the bittorrent traffic by adding a firewall rule.

Application	Total MB	Add Rule	
bittorrent	9526.82	F	T
unknown	48.49	F	T
ssl	21.26	F	T
arp	1.57	F	T
atmf	0.85	F	T
icmpv6	0.64	F	T
sflow	0.19	F	T
mdns	0.08	F	T
websocket	0.07	F	T
ospf	0.07	F	T

- Configure the firewall rules.

New Firewall Rule

Action: Deny

Application: bittorrent

From: ANY

To: ANY

- Turn on the firewall.

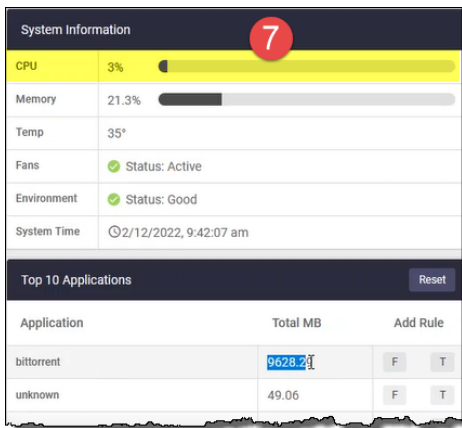
AR4050S ROUTER_AMFPLUS_DEMO Up time: 0 days 17:39 manager Save

Firewall ON

2 Rules

Action	Application	From	To	Errors
Deny	bittorrent	ANY	ANY	
Permit	any	ANY	ANY	

- Go back to the Dashboard and check the CPU percentage.



FAQs

- What devices are monitored, and can you select the devices that will be monitored?
 - All AlliedWare Plus devices are automatically added to Health Monitoring. Devices cannot be added or deleted manually.
- How often is a device polled?
 - Polling occurs every 5 minutes.
- How much historical data is stored?
 - 7 days.

Note: You can export a PDF of the selected tab from the Health Monitoring page by clicking the **Download PDF** button. The PDF scales to match what you see in Vista Manager when you zoom in or out in your browser window.

Network Health tab

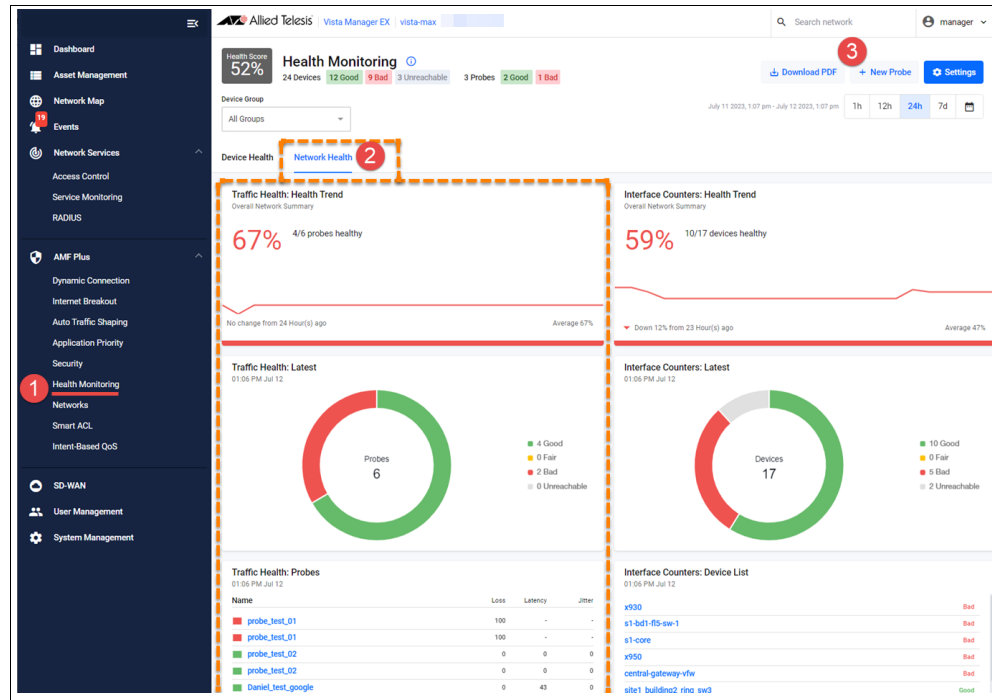
The Health Monitoring feature allows you to create Link Monitoring (Linkmon) Probes and will report statistics (latency, jitter and packet loss) on these probes, representing the health of the traffic.

For example, you could use the probes to monitor links:

- from a company's router to the Internet, to ensure it is operating and at an acceptable level.
- inside a company's LAN from the core switch to a highly used backup server to check latency.
- to a remote office used for video broadcast to check the jitter.
- between the core switches of two remote offices.

To use this feature:

1. Go to **AMF Plus > Health Monitoring**
2. Click **Network Health**.
 - The following widgets are visible: Health Trend, Latest, Probes, and History.
3. To add a probe, click **+ New Probe**




4. The **New Traffic Health Probe** window opens.
 - Select the **Probe Type** - ICMP Echo or HTTP GET.
 - Select a **Source** device - the drop down box lists all linkmon capable devices in the network.
 - Type in a **Destination** - for ICMP echo probes, this is either an IP address or FQDN. For HTTP GET probes, this can only be an FQDN.
 - Enter an **Interval** in seconds. For ICMP probes the default is 1 and for HTTP GET probes the default is 30.
 - Set the **Thresholds** for packet loss, jitter, and latency.

Health Monitoring Polling and Interface Counters Polling

To view updated information on Interface Counters in the Network Health tab, you must first enable both **Health Monitoring Polling** and **Interface Counters Health Monitoring Polling** in the settings. This is because Interface Counters Health Monitoring Polling only affects the interface counters (not the entire network health tab, just the right column).

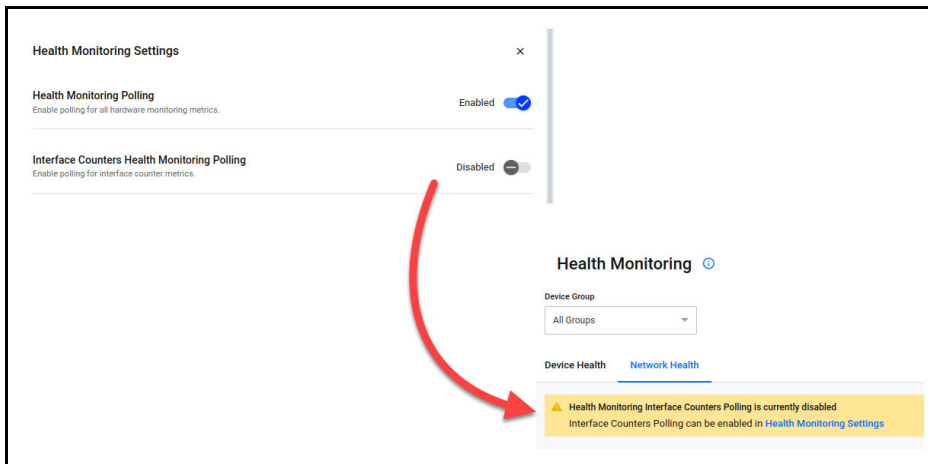
To access the Health Monitoring settings, click the **Settings** button on the top right corner of the Health Monitoring page:



You must first enable Health Monitoring Polling before you can enable Interface Counters Health Monitoring Polling.

Interface Counter Polling is a subset of Health Monitoring Polling and cannot be enabled otherwise.

A warning will appear on the Network Health tab of Health Monitoring if **Interface Counters Health Monitoring Polling** is disabled, but Health Monitor polling is enabled:



The interface counters dashboard includes the following widgets:

- Health Trend
- Latest
- Device List
- History

Use the dashboard to easily monitor:

- Interface errors occurring across the entire network or on a particular device.
- Overall health status and specific health metrics of interfaces.
- Detailed explanations of the errors, enabling you to effectively diagnose and resolve any issues that arise.

You can configure the threshold values for each interface error type in the main settings panel of the dashboard. These thresholds apply to the entire network and are used to determine the health status of devices and interfaces.

Health Monitoring third-party devices

From version 3.10.3 onwards, the Health Monitoring page incorporates information on third-party devices (i.e., non-AMF devices). Prior to this, only network devices belonging to AMF were subject to monitoring. Using the SNMP Plugin, the Health Monitoring dashboard displays the latest stats for all devices with an IP address.

The SNMP plugin collects detailed information and statistics from network devices, and utilizes a Standard MIB compiler to generate charts based on MIB values. See [Table 1 on page 83](#) for the supported MIB information.

If the SNMP device can provide such information, the Health Monitoring dashboard will display CPU, RAM, Storage, and Temperature statistics.

Table 1: Supported MIB information

VENDOR NAME	SNMP AGENT	MEMORY SIZE	STORAGE SIZE	CPU LOAD	SENSOR
Standard MIB	SNMP Service (Windows) Net-SNMP package (Linux)	HOST-RESOURCE-MIB::hrStorageTable (hrStotageType=hrStorag eRam)	HOST-RESOURCE-MIB::hrStorageTable (hrStorageType=hrStora geFixedDisk or hrStorageRamDisk)	HOST-RESOURCE-MIB::hrProcessor Load	ENTITY-SENSOR-MIB::entPhySens orTableENTITY-MIB::entPhysical Table

Activating Health Monitoring statistics for non-AMF devices

First you need to add the SNMP Plug-in, then add and configure an SNMP network.

1. To add the SNMP plugin:

- Go to **System Management > Plugins**
- Click **+ Add Plugin**
- Enter the Server URL - `https://<ip-address>:6443/NetManager`
- Verify the connection and **Save**.

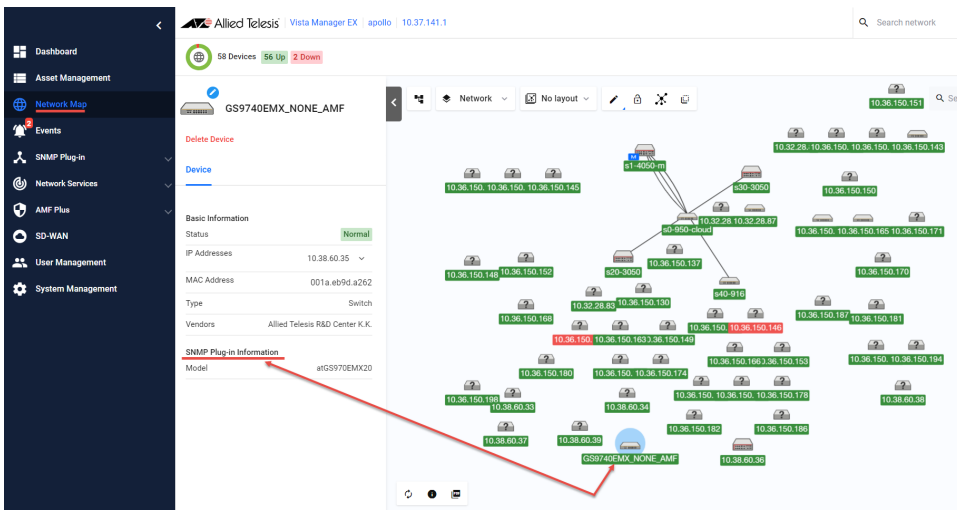
2. To add an SNMP network.

- Go to **SNMP Plug-in > Network Tree**
- Create a subnet.
- After a few minutes the SNMP Plugin automatically discovers available devices under the specified subnet.
- After auto discovery is complete, a list of devices is shown on the Network Tree.

3. Check the Network Map to find the devices that were discovered by the SNMP plugin.

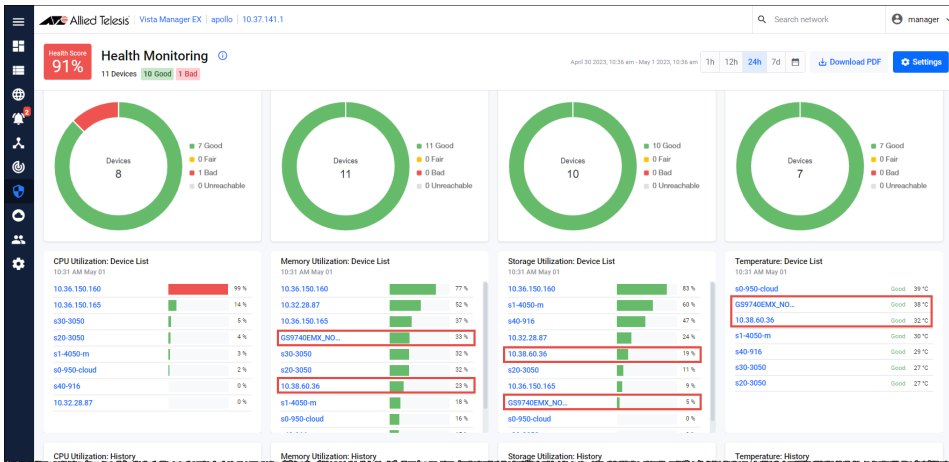
- Go to **Network Map**

- Click on a device to see its SNMP Plugin information.



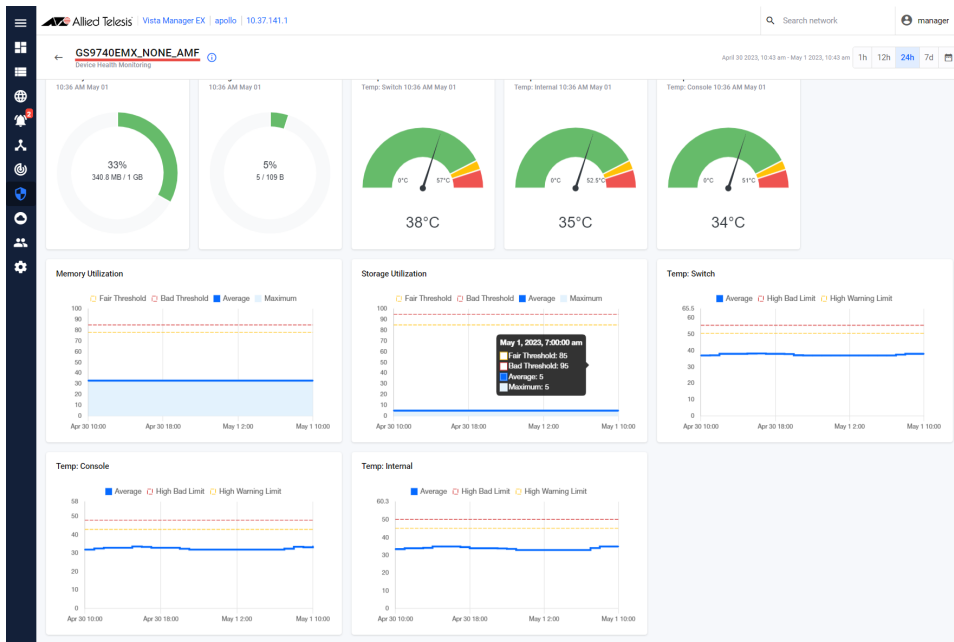
- To obtain the latest statistics for third-party or non-AMF devices discovered by the SNMP plugin, check the Health Monitoring dashboard or the Device Specific dashboard.

- Go to **AMF Plus > Health Monitoring**
- The Health Monitoring page includes all valid third-party or non-AMF devices in the summary charts.



- Click on a specific device to see its details.

- For example, the image below shows the details for 'GS970EMX_NONE_AMF'.

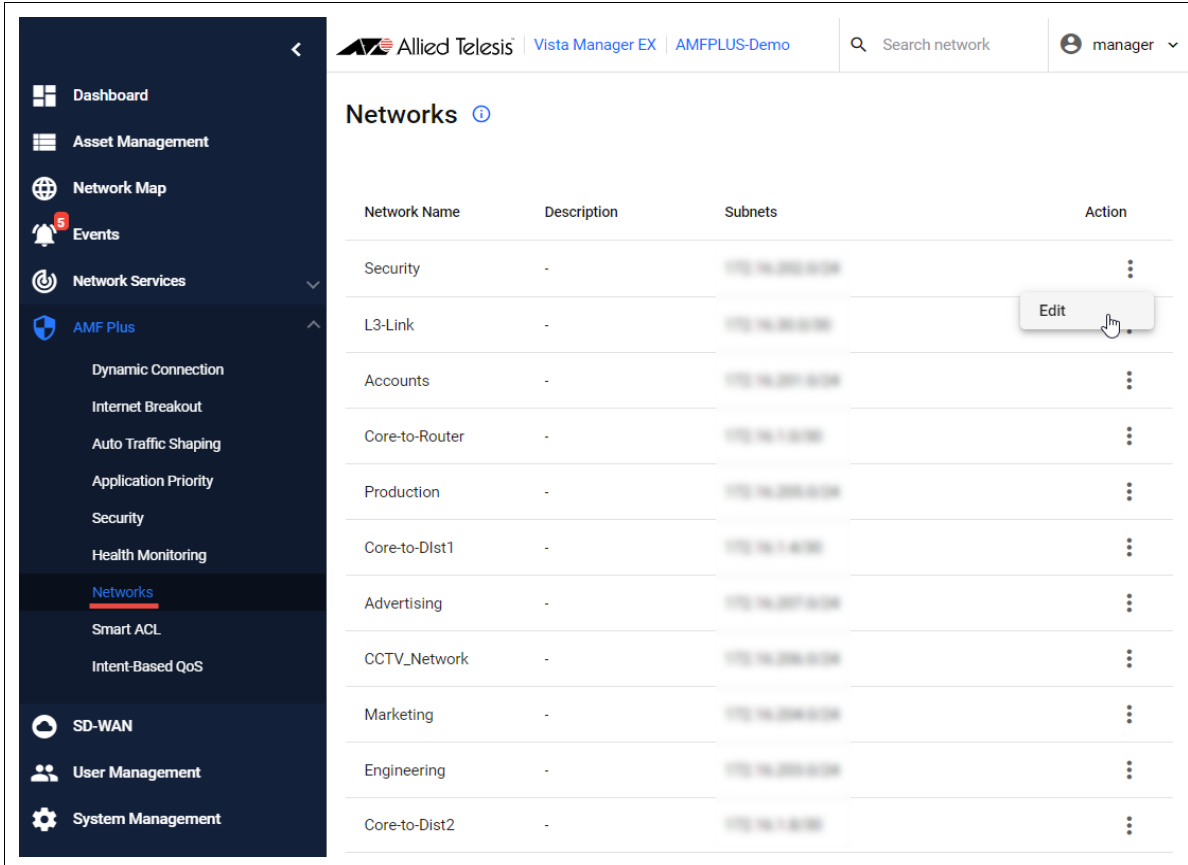


Please note that you can only find third-party devices on the Health Monitoring dashboard if the correct MIB is supported on the third party device.

Networks

Vista Manager EX defines a Network as an IP subnet attached to a VLAN. For example, subnet 192.168.1.0/24 is associated with VLAN1. As part of the Smart ACL feature, network entries are automatically imported and maintained by Vista Manager EX via the attached network devices.

Each network is given a default **Network Name** and **Description** which you can **Edit** to suit your needs.



Network Name	Description	Subnets	Action
Security	-	192.16.200.0/24	⋮
L3-Link	-	192.16.200.0/24	Edit
Accounts	-	192.16.200.0/24	⋮
Core-to-Router	-	192.16.1.0/24	⋮
Production	-	192.16.200.0/24	⋮
Core-to-Dist1	-	192.16.1.0/24	⋮
Advertising	-	192.16.200.0/24	⋮
CCTV_Network	-	192.16.200.0/24	⋮
Marketing	-	192.16.200.0/24	⋮
Engineering	-	192.16.200.0/24	⋮
Core-to-Dist2	-	192.16.1.0/24	⋮

NETWORKS - FIELD	DESCRIPTION
Network Name	These are auto-generated in sequence Network-1, Network-2...Network-n, but you can rename them via the Edit Action.
Description	The network description, for example: VLAN100. Use the Edit Action to add or change a description.
Subnets	The subnet IP address, these are auto-generated and derived from the attached networks. But, networks can only be added via the CLI, i.e. configuring a VLAN with subnet(s).
Action	Use the Action menu to edit the network name and description.

Smart ACL

The Smart ACL tool allows you to manage ACLs across devices in the network. ACLs provide traffic flow control and decide which types of traffic are forwarded or blocked.

You can (could) use Smart ACL to control the resources that clients access in the network. For example, you might want to stop marketing clients from being able to see a security client's CCTV video stream and also stop the security clients from accessing marketing videos.

There are three parts to the Smart ACL tool:

1. **Networks:** VLANs configured with an IP subnet.
2. **Policies:** Access List filters (rules) used to control network traffic.
3. **Policy Matrix:** A display of:
 - currently configured source and destination networks
 - policy status - configured, active, and hits on the ACL policy

The objective of Smart ACL is to allow you to apply policies between networks - to control traffic from a source network going to a destination network.

The screenshot displays the Smart ACL interface with the following components:

- Header:** Allied Telesis Vista Manager EX
- Section:** Smart ACL
- Navigation:** Policy Matrix (selected) and Policies
- Legend:**
 - Default (grey circle)
 - Policy Configured (dark grey circle)
 - Policy Active (light blue circle)
 - Policy Hit (dark blue circle)
- Grid:** A 5x5 matrix with Source (Network-1 to Network-5) on the y-axis and Destination (Network-1 to Network-5) on the x-axis. All cells are currently empty.
- Actions:** 'Edit Policies' links are located to the right of each row in the grid.

Getting started with Smart ACLs

You need to do some initial configuration before you can use the Smart ACL tool. The initial configuration ensures that the Policy Matrix shows the current active policies.

In brief, you first configure a network and optionally assign it a meaningful name, then create an ACL policy and apply it to the network. Let's look at each step in more detail:

1. Configure a network.

- Use the **CLI** to configure a network on your AlliedWare Plus device.

For example:

```
vlan database
vlan 100

interface port1.0.5
switchport mode trunk
switchport trunk allowed vlan add 100

interface vlan100
ip address 172.16.2.1/24
```

2. Assign a meaningful name to the network (optional).

- Go to **AMF Plus > Networks**
- By default, networks are auto-generated in sequence Network-1, Network-2...Network-n, but you can change the default name to a more meaningful one by using the **Edit** action. You can also add a useful **Description** to the **Network Name**.

Networks ⓘ

Network Name	Description	Subnets	Action
Network-1	VLAN200 - has these subn...	172.16.200.0/24	⋮ Edit
Network-2	VLAN100 - has these subn...	172.16.100.0/24	⋮
Network-3	VLAN1037 - has these sub...	172.16.1037.0/24	⋮
Network-4	VLAN1034 - has these sub...	172.16.1034.0/24	⋮
Network-5	VLAN2 - has these subnets.	172.16.2.0/24	⋮

Networks

Network Name

Security

Description

VLAN200-has these subnets

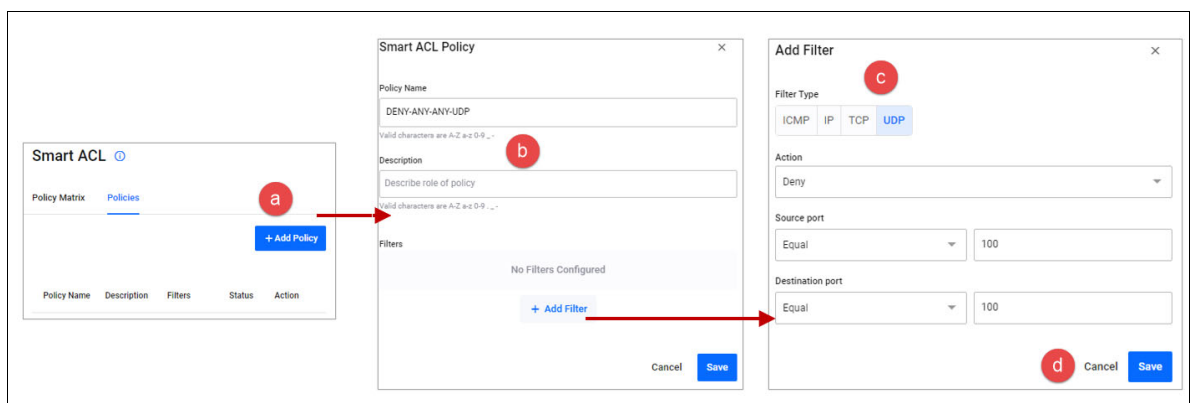
3. Create an ACL Policy

- Go to **AMF Plus > Smart ACL**
- The **Policy Matrix** displays all currently configured networks. In the example below, there are 5 networks configured with default names.



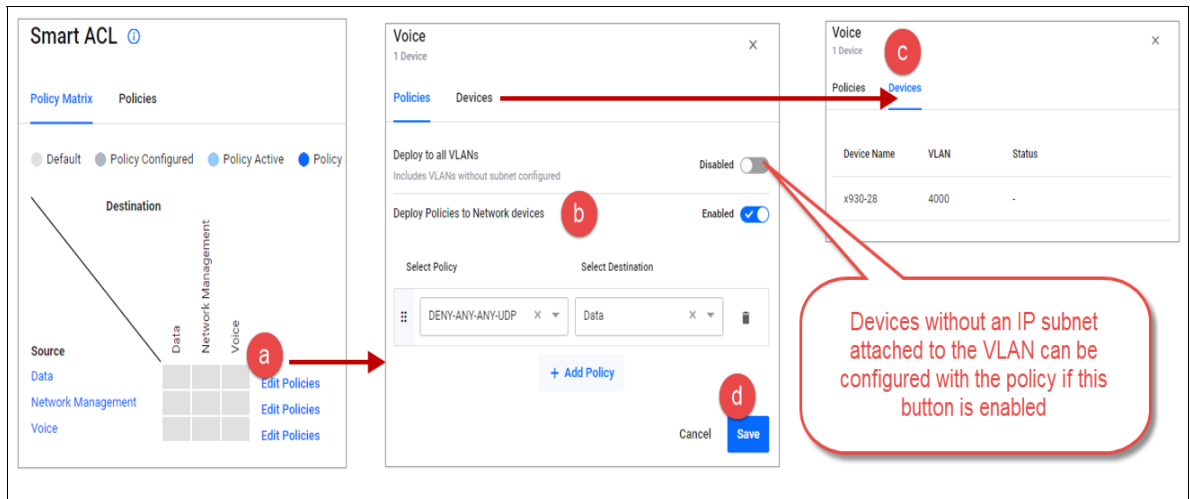
4. Select the **Policies** tab, then:

- a. Click **+Add Policy**
- b. Enter a **Policy Name** and **Description**
- c. Click **+ Add Filter** - set the **Action** and **Filter Type**
- d. Click **Save**.



In the example above, an ACL policy called DENY-ANY-ANY-UDP has an action of DENY if the packet matches UDP source port =100 and destination port =100.


5. Back in the **Policy Matrix** tab, apply a policy to a network.
 - a. Select **Edit Policies**
 - b. Configure as required - i.e. select a policy and destination
 - c. Check the policy is applied to the correct device(s)
 - d. Click **Save**



In the example above:

- The ACL policy DENY-ANY-ANY-UDP is applied to packets from the Voice network going to the Data network.
- The Devices tab shows all the devices that the policy will be applied to. In this case, only the device x930-28 will be configured with the ACL policy.

Now you can see the policy is active from source network **Voice** to destination network **Data**.

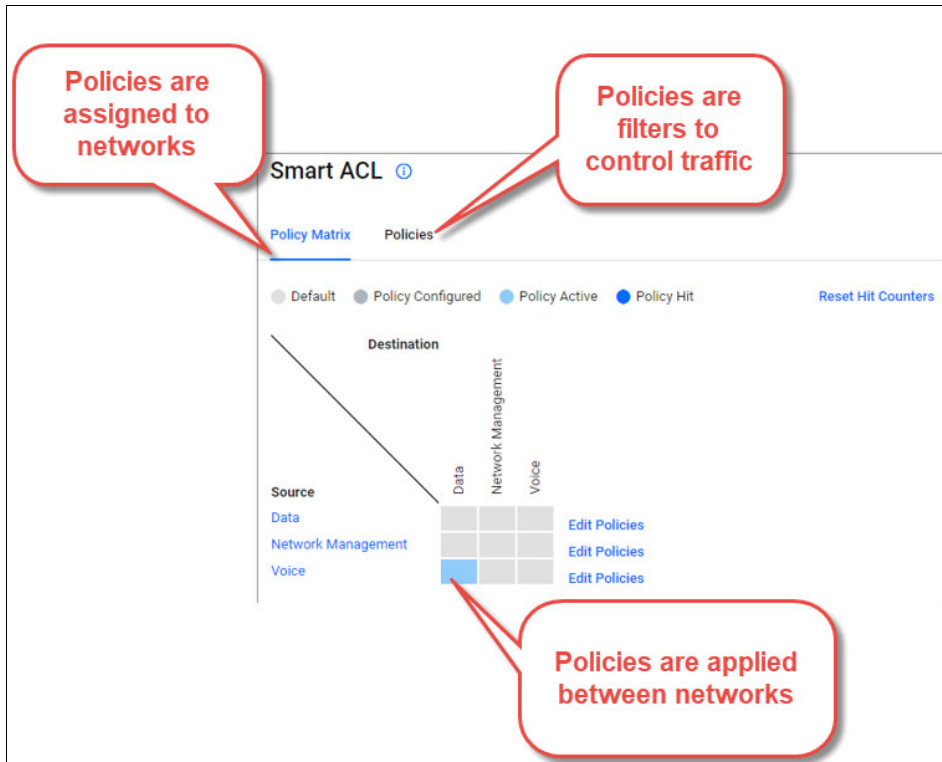


This completes the initial configuration.

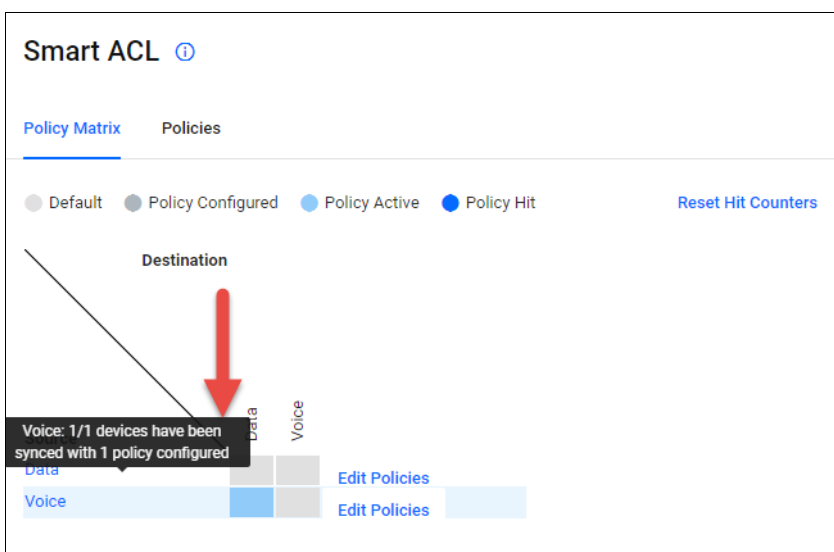
Understanding the Smart ACL Policy Matrix and its operation

The Smart ACL tool makes configuring complex ACLs on networks easier. It allows you to create, edit, view, and delete ACL policies. ACL policy changes are synced and applied by Vista Manager EX automatically to VLANs using **per-VLAN ACLs**.

Once the initial configuration is complete, the Policy Matrix is set up with the configured networks. In the example below, you can see an active policy from source network **Voice** to destination network **Data**.



You can hover your mouse over a network name to see how many of the devices in that network have been synced with the ACL configuration for the policy.




The benefit of this is every time a new device is added as part of the network and this subnet, the values will increase and the new device will automatically receive the policy.

What are per-VLAN ACLs?

Per-VLAN ACLs filter traffic as it **ingresses** a VLAN.

Per-VLAN ACL rules are applied to **all** ports on which the VLAN is active. This means they are applied to all ports that are access ports in the VLAN, all trunk ports that allow packets tagged for the VLAN, and all trunk ports whose native VLAN is this VLAN.



Can Smart ACL configure other types of ACLs, for example an interface ACL?

Smart ACL only supports per-VLAN ACLs, and only applies when traffic is going from one subnet to another subnet.

What actual configuration is applied to the device?

Take the example used in ["Getting started with Smart ACLs"](#) on page 88:

- An ACL policy called DENY-ANY-ANY-UDP has an action of DENY if the packet matches UDP source port =100 and destination port=100.
- This policy is applied to traffic from the Voice network (V4000, 172.16.0.0/16) going to the Data network (V1, 10.37.62/27).

```
! acl-group matching the Data subnet 10.37.62.64/27
acl-group ip address VISTA_V4_1
  ip 10.37.62.64/27

! Deny traffic matching source IP = any, UDP source port = 100, and
destination IP = Data subnet, UDP destination = 100.
access-list hardware VISTA_V4_source2_destination1_policy1
  deny udp any eq 100 host-group VISTA_V4_1 eq 100

! Apply access-list to access-map
vlan access-map VISTA_ACCESS_MAP_source2
  match access-group VISTA_V4_source2_destination1_policy1

! Attach access-map to VLAN 4000
vlan filter VISTA_ACCESS_MAP_source2 vlan-list 4000 input
```

What commands can I use to view the Smart ACL configuration?

Use the following commands to view the Smart ACL configuration:

```
show acl ip address
show access-list
show vlan access-map
show vlan filter
```

To view the hit counters, use the command:

```
show access-list counters
```

Intent-based QoS

Quality of Service (QoS) is a way to prioritize network traffic to ensure that the most important traffic gets through the network with minimal delay or interference.

QoS is a complicated feature with many configuration options and different ways to configure the feature. To configure QoS on a network, you will typically follow these steps:

- Identify the types of traffic that are important and need to be prioritized, such as voice or video traffic.
- Assign each type of traffic a priority level based on its importance. This is typically done using a QoS tagging system.
- Configure your network devices (routers, switches, etc.) to recognize the QoS tags and prioritize traffic accordingly.
- Set bandwidth limits or rate limits on non-priority traffic to prevent it from interfering with the prioritized traffic.

By configuring QoS on your network, you can ensure that critical applications like voice and video are given priority over less important traffic, leading to better network performance and user experience.

From Vista Manager EX version 3.10.1 onwards, you can use **Intent-based QoS** to easily manage and troubleshoot a basic QoS configuration on your network as part of **AMF Plus**.

The benefits of Intent-based QoS

In a congested network where packets are being dropped, it is quite difficult to determine where the drops are occurring. A network could consist of numerous devices, each with a number of ports with egress queues. Detecting drops on one of the queues, on one of those ports, on one of those devices is challenging. Intent-based QoS helps you troubleshoot and visualize the performance of egress queues and manage their settings.

You can:

- Visualise egress queues across the entire network and for individual devices:
 - Drops
 - Throughput
- Modify egress queue settings:
 - Strict priority – queue egress limits
 - Weighted Round Robin – queue weightings

Getting started


First you need to **manually** apply a default QoS configuration VISTA_DEFAULT_POLICY to all switches in your network. Please see "[Configuring the Vista Manager EX default policy](#)" on page 109 for guidelines and some complete configuration examples.

This default configuration sets up 2 priority queues and 6 weighted round robin (WRR) queues. The strict priority queues have an egress rate limit applied, and the WRR queues each have a weighting applied. The configuration also defines a mapping of DSCP fields to QoS queue based on industry standards. This mapping cannot be changed via Vista Manager EX.

New ports configured with the default QoS policy are added to the list of polled ports. Likewise, removed ports with the default QoS policy are deleted from the list of polled ports.

Once the default configuration has been applied on the network, the **Intent-Based QoS** dashboard shows the state of the network in regards to the QoS queues.

Each of the eight QoS queues has a label based loosely on what sort of traffic is expected on the queue. For example, the highest priority queue, QoS queue 7 has the label 'Voice' as this queue will be used for VoIP traffic. Queue 6 has the label 'Video' as this queue will be used for a variety of video services, and so on.



You can see in the diagram above that the **Streaming** queue is experiencing queue drops. Using the dashboard, you can investigate further to see when and on which device drops are occurring.

Using the Dashboards

You can adjust the rate limit or weighting of a problematic queue on the entire network by using a simple graphical tool - the Intent-Based QoS dashboard.

In fact, there are three interlinked dashboards:

1. **Intent-based QoS** - displays egress queue details across the network. Data is aggregated from all ports on all devices in the network.
 - Click on a device name to open the device dashboard.
2. **Device** - displays egress queue details from a single device. Data is aggregated from all ports on the device.
 - Click on a port name to open the port dashboard.
3. **Port** - displays queue drops and throughput from a port.

The three dashboards allow you manage QoS configurations on your network. You can use them to drill-down from a wide-angle view of the network traffic, select a device, and then select a port on that device.

Navigating the dashboards

The Intent-Based QoS dashboard shows queue details for the entire network. Data is aggregated from all ports configured with the VISTA_DEFAULT_POLICY. Vista Manager EX scans the network for any ports configured with the default Vista QoS policy. Every five minutes Vista will poll these ports for queue drops and queue throughput (transmitted bytes). Intent-Based QoS presents the data in dashboards:

The layout is similar for all three dashboards. The Queue status ribbon run along the top, with specific queue details and historic charts below.

The screenshot shows the 'Intent-Based QoS' dashboard. At the top, there is a 'Queue status ribbon' with eight queues: 7 Voice (0), 6 Video (0), 5 Network Management (0), 4 Streaming (0), 3 Transactional (0), 2 Bulk Data (0), 1 Scavenger (1.7 B), and 0 Best Effort (0). The 'Scavenger' queue is highlighted in red. Callouts include: 'Queue status ribbon: Click on a queue to see queue details', 'Queue details: Details for a particular queue', 'Historic charts: Cumulative data from all queues over the selected time period', 'Time period: March 21 2023, 10:38 am - March 28 2023, 10:38 am', and 'Access to: - Queue configuration - Monitoring thresholds'. Below the ribbon, the 'Queue 7: Voice' details are shown, including bandwidth limit (50%), throughput (1.09 Mbps), and queue drops (0). Two charts are displayed: 'Throughput During last 7 days' and 'Queue Drops During last 7 days'. A table on the right shows 'Queue Drops by Device' for various devices like ACCESS1_AMFPLUS_DE...

The Queue status ribbon displays the drops and the status of each queue.

This close-up shows the 'Queue status ribbon' with eight queues: 7 Voice (0), 6 Video (0), 5 Network Management (0), 4 Streaming (0), 3 Transactional (0), 2 Bulk Data (0), 1 Scavenger (1.7 B), and 0 Best Effort (0). Callouts explain: 'Queue number' and 'Queue label' pointing to the queue identifiers; 'Click on a queue to display its details' pointing to the ribbon; 'Cumulative drops on the queue network-wide in the last 5 minutes' pointing to the '0' values; and 'The queue color changes based on threshold settings: - Green - below the Fair threshold, - Yellow - between Fair and bad, - Red - above the Bad threshold' pointing to the red 'Scavenger' queue.

Drops should be investigated, especially on higher numbered queues, as it could be an indication that congestion is occurring in the network, and potentially impacting on user experience.

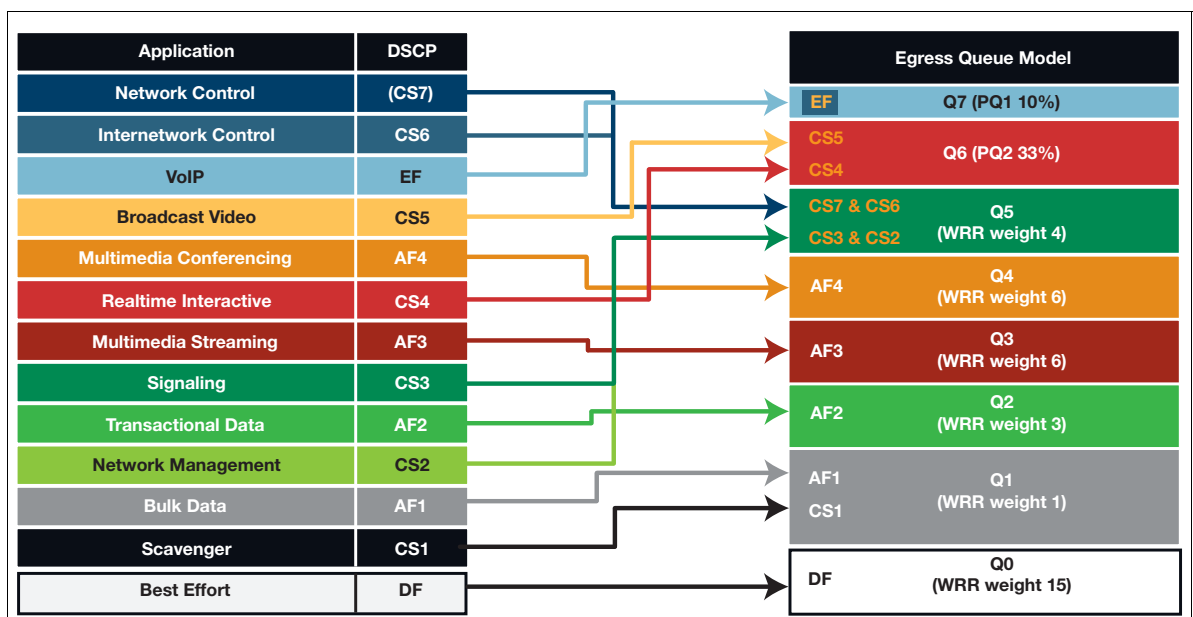
Where do the dashboard queue names come from?

Each of the eight queues has a label describing what sort of traffic is expected on the queue.

QUEUE	LABEL	DESCRIPTION
7	Voice	Traffic requiring minimum loss, latency, and jitter, such as VoIP telephony.
6	Video	Traffic requiring low loss, latency, and jitter, such as video-conferencing.
5	Network Management	Traffic protected with a minimum bandwidth guarantee such as SNMP, NTP, and Syslog.
4	Streaming	Highly interactive traffic, such as instant messaging and Telnet
3	Transactional	Low response time traffic where users wait for transactions to finish, such as SAP and Oracle.
2	Bulk Data	Low interaction, not drop sensitive traffic, such as FTP, E-Mail and Backup Operations.
1	Scavenger	Business-irrelevant traffic, such as Gaming and Peer-to-Peer Media Sharing.
0	Best Effort	Traffic not requiring differentiated treatment.

We recommended traffic is place into the correct queues, but there is no strict requirement. For example, there is nothing stopping you from putting Voice traffic into the Streaming queue. However, the labels in Intent-Based QoS cannot be changed.

It is ultimately up to you how you want to bind RFC4594 traffic classes to egress queues, however the bindings denoted in the following diagram are recommended.



Queue details

Click on a queue in the Queue Status ribbon to see its details. In the example below, queue 7 Voice is selected and its details displayed underneath.

Queue 7: Voice
Traffic requiring minimum loss, latency and jitter, such as VoIP Telephony

Queue Drops
During last 7 days

Queue Drops by Device
During last 7 days

Device	Total Throughput	Queue Drops
ACCESS1_AMFPLUS_DE...	55.82 GB	0
CORE_AMFPLUS_DEMO	0 B	0
DIST1_AMFPLUS_DEMO	0 B	0
DIST2_AMFPLUS_DEMO	0 B	0
ACCESS2_AMFPLUS_DE...	0 B	Unsupported
ACCESS3_AMFPLUS_DE...	0 B	Unsupported

Annotations:

- Queue number, label, description
- Queue egress rate limit % setting
- Aggregated throughput (in bps) for the queue across the network. Helps identify possible congestion and issues in the network.
- Aggregated drops for the queue from all ports across the network over the last 5 minutes.
- Queue Drops displays drops over a time period. This allows you to determine if there is a persistent drop issue or just a one off or intermittent spike in drops
- Fair/Bad threshold setting
- Link to the Device Dashboard of each device. Sorted by largest queue drops. If drops are equal, sorted by largest throughput.
- QoS enabled
- Number of drops

Historic charts

The historic charts display past details for all queues across all devices on the network. In the left chart below you can see the throughput per queue aggregated from ports over the selected time period, in this case the last 7 days. In the right chart, you can see drops per queue aggregated from ports over the same time period.

Throughput
During last 7 days


Queue Drops
During last 7 days

Annotations:

- QoS enabled
- QoS enabled
- Drops broken down by queue

Configuring the queue settings

Use the **Settings** button to access the **Intent-Based QoS Settings** window.



 A screenshot of the "Intent-Based QoS Settings" window. The window has a dark blue sidebar on the left with a menu including "Dashboard", "Asset Management", "Network Map", "Events", "Network Services", "Access Control", "Service Monitoring", "RADIUS", "AMF Plus", "Dynamic Connection", "Internet Breakout", "Auto Traffic Shaping", "Application Priority", "Security", "Health Monitoring", "Networks", "Smart ACL", "SD-WAN", "User Management", and "System Management". The "AMF Plus" section is expanded, and "Intent-Based QoS" is selected. The main content area shows "Intent-Based QoS" with a summary of queues: "7 Voice", "6 Video", "5 Network Manag...", and "4 Streami...". Below this is "Queue 7: Voice" with details like "Bandwidth Limit: 10%", "Throughput: 0 bps", and "Queue Drops: 0". A "Queue Drops" graph is also visible. On the right, the "Monitoring Thresholds" tab is active, showing a table of settings for various queues. The table has columns for "Queue", "Fair", and "Bad" thresholds, with "Good" also implied. Each row has a "Fair" and "Bad" slider and a "Good" checkbox. The "Good" checkboxes are all checked. At the bottom right of the window are "Cancel" and "Save" buttons.

Queue	Fair	Bad	Good
7 Voice	1	1	<input checked="" type="checkbox"/>
6 Video	1	1	<input checked="" type="checkbox"/>
5 Network Management	1	1	<input checked="" type="checkbox"/>
4 Streaming	1	1	<input checked="" type="checkbox"/>
3 Transactional	1	1	<input checked="" type="checkbox"/>
2 Bulk Data	1	1	<input checked="" type="checkbox"/>

This is where you set or change queue parameters for Strict Priority egress limits and WRR queue weightings. Any changes you make are pushed out to all devices configured with the QoS policy named: VISTA_DEFAULT_POLICY.

If you don't see a screen resembling the one shown above, where no data appears in the dashboard, it indicates that the default values have not been configured. You will need to use the CLI to configure these settings. This is described in the section "[Configuring the Vista Manager EX default policy](#)" on page 109.

There are three setting tabs: **Monitoring Thresholds**, **Queue Configuration** and **Auto Queue Configuration**:

Monitoring Thresholds tab

The **Monitoring Thresholds** tab lets you change the drop threshold for queues. By default, the Fair and Red thresholds are set to 1 drop. This means that if there are ≥ 1 drops, the queue display will show as red.

Thresholds define the queue state display in the dashboard (Good, Fair, Bad). The thresholds are based on the acceptable number of dropped packets per queue during a 5-minute time period. You can set threshold values for each queue individually.

Intent-Based QoS Settings [Close]

Monitoring Thresholds | Queue Configuration | Auto Queue Configuration

QoS Polling Enabled [Toggle]

Queue Monitoring Thresholds
Thresholds define Queue state in dashboard (Good, Fair, Bad). Adjust the amount of queue drops acceptable per queue during 5 min time period.

Queue	Fair	Bad	Include
7 Voice Traffic requiring minimum loss, latency and jitter, such as VoIP Telephony	1	1	[Toggle]
6 Video Traffic requiring low loss, latency and jitter, such as Videoconferencing	1	1	[Toggle]
5 Network Management Traffic protected with a minimum bandwidth guarantee such as SNMP, NTP and Syslog	1	1	[Toggle]
4 Streaming Highly interactive traffic, such as instant messaging and Telnet	1	1	[Toggle]
3 Transactional Low response time traffic where users wait for transactions to finish, such as SAP and Oracle	1	1	[Toggle]
2 Bulk Data Low interaction, not drop sensitive traffic, such as FTP, E-Mail and	1	1	[Toggle]

[Cancel] [Save]

Thresholds are applied network-wide and cannot be set on a single device or port. If the QoS configuration is different between device ports, a warning message is displayed in the Intent-Based QoS dashboard.

Queue Configuration tab

The **Queue Configuration** tab lets you set queue parameters for: Strict Priority (egress rate limiting) and WRR - weighting.

Strict Priority queue settings

Use the Strict Priority queues for traffic requiring minimum loss, latency, and jitter, such as VoIP and video conferencing.

Priority queues send their packets first, in order of priority. This means that queue 7 sends packets until it is empty (or reaches its bandwidth limit), then queue 6 sends packets. Queues 0-5 only get to send packets when both queues of 6 and 7 are empty or have reached their bandwidth limit. If you don't restrict the queue bandwidths, the highest priority queues could stop the other queues from getting any bandwidth on particularly busy interfaces.

The settings allow you to limit the bandwidth used on the interface on each device, by setting the egress rate limit % value of the queue. This allows you to reserve bandwidth on the interfaces of the devices for other lower priority queues, which stops the highest priority queues from ever using all the bandwidth.

- You cannot set the egress rate limit to zero, because this is the same as disabling traffic flow from the interface.
- You can set the egress rate limit to 100%. This means that the queue in question will use as much of the egress bandwidth as it can, up to the capacity of the interface.
- If you set both Strict Priority queues to an egress rate limit of 100%, then the higher priority queue (7 - Voice) will use as much capacity as it needs. The capacity that queue 7 does not use is available for the lower priority queue (6 - Video), which will use as much of that remaining capacity as it needs.
- If you try to set the total egress rate limit to over 100%, Vista Manager gives you a warning, because this will allow the Strict Priority queues to "starve" the WRR queues if the strict priority queues' traffic demand uses all the interface bandwidth.

Weighted Round Robin (WRR) queue settings

Use the Weighted Round Robin (WRR) queues for:

- Network Management, Streaming, Transactional, Bulk Data, Scavenger, and Best Effort.

Weighted Round Robin (WRR)
 Weighted Round Robin (WRR) WRR queues are configured using a weight of 1-15. The higher the weight the more frames are sent. Queue values will be converted to absolute throughput value using the actual speed of the interface on each device. [More info](#)

Queue	Weight 1 - 15
5 Network Management Traffic protected with a minimum bandwidth guarantee such as SNMP, NTP and Syslog	4
4 Streaming Highly interactive traffic, such as instant messaging and Telnet	6
3 Transactional Low response time traffic where users wait for transactions to finish, such as SAP and Oracle	6
2 Bulk Data Low interaction, not drop sensitive traffic, such as FTP, E-Mail and Backup Operations	3
1 Scavenger Entertainment traffic, such as Gaming and Peer-to-Peer Media	1

WRR queues have a weight of 1-15. The weights are relative to each other and work as ratios. When the egress interface is congested, a greater proportion of traffic is sent over queues with a higher relative weighting.

For example, a queue configured with a weighting value of 15 will send 15 times as much traffic as a queue configured with a weighting of 1 when the egress interface is congested. Likewise, if all queues are configured with a value of 15, all the queues will send the same amount of traffic. It is the relative difference that matters, so setting all queues to 15 is the same as setting all queues to 1.

Intent-Based QoS

Configuration mismatch
 Some devices are out of sync with Vista Manager Settings. Save settings to apply configuration to all devices.

7 Voice	6 Video	5 Network Management	4 Streaming	3 T
0	0	0	0	0

Auto Queue Configuration tab

Autonomous Queue Configuration empowers you to automate QoS configuration changes across the network in response to changing traffic flows. If any queue's number of egress drops exceeds the 'bad' monitoring thresholds, the Auto Queues Configuration feature will adjust the resource allocated to that queue.

You can choose how frequently you want the autonomous configuration to run. You can also provide upper and lower bounds for resources automatically allocated to each queue.

Each time the Auto Queue Configuration feature changes the QoS configuration, you will receive a message in the event log specifying which queues changed, their previous values, and their new values.

The priority queue optimisation algorithm

Each time the Autonomous Queue Configuration algorithm runs, Intent-based QoS inspects the egress queue drops for each priority queue in turn. These queue drops include all drops for the queue across the whole network.

If the queue is in a 'bad' state, Intent-based QoS increases the egress rate for the queue by 3%. It applies the new configuration to every QoS capable device in the network.

It determines a 'bad' state by summing all egress drops for the queue within each 5-minute monitoring period, since the algorithm last ran. This gives a number for each 5-minute period, called the "summed drops". If the summed drops in any of the 5-minute periods exceed the 'bad' monitoring threshold, then the queue is considered to be in a bad state. You can change the monitoring threshold for each queue on the Monitoring Thresholds tab.

Intent-based QoS can automatically increase each queue up to the queue's maximum egress rate.

The weighted round robin optimisation algorithm

Each time the Autonomous Queue Configuration algorithm runs, Intent-based QoS inspects the egress queue drops for each weighted round robin queue in turn. These queue drops include all drops for the queue across the whole network.

If the queue is in a 'bad' state, Intent-based QoS donates one unit of weight from a good queue to the bad queue. The good queue is called the 'donor' queue. The algorithm attempts to perform this transfer for all queues that are in a bad state.

It determines a 'bad' state by summing all egress drops for the queue within each 5-minute monitoring period since the algorithm last ran. If the summed drops in any of these 5-minute windows exceeds the 'bad' monitoring threshold, the queue is considered to be in a bad state. You can change the monitoring threshold for each queue on the Monitoring Thresholds tab.

The algorithm selects the 'donor' queue according to the following criteria:

- It must be in a good state, which means it has not exceeded the queue drops specified in its monitoring threshold.
- Its current weight must be greater than the minimum weight.
- Its current weight must be lower than the maximum weight.
- Among all the queues that satisfy the above criteria, the algorithm chooses the queue with the largest baseline surplus. The baseline surplus is the queue's initial weight minus the queue's current weight.

The lowest priority queue which satisfies all these conditions is chosen as the donor queue. If no donor queue can be found, no weight will be transferred between the queues.

- If the recipient queue's current weight is already equal to its maximum weight, no weight will be donated to it.
- If the algorithm decides the weight must be changed, it applies the new configuration to every QoS capable device in the network.

You can limit the extent of the automatic weight changes by specifying min and max weights. Auto Queue Configuration will keep the queue weightings within those min and max weights.


Port congestion

When a port receives more traffic than it can transmit, it buffers the traffic until the traffic can be sent. If the buffer becomes full and cannot buffer any more packets, any new incoming packets will be dropped, this is known as tail drop. This can cause two issues:

- **packet delay** - the packet in the buffer is delayed until the port is ready to send it.
- **packet drops** - the packet is dropped and lost forever.

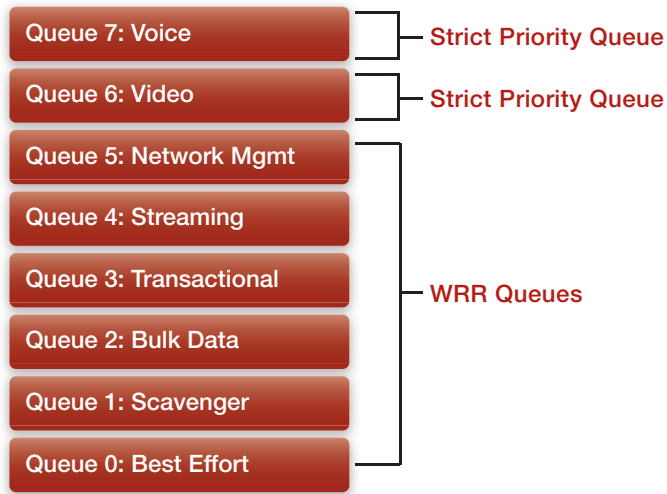
The transmitting device may choose to resend the lost packet, but this could take some time, because it has to detect the packet has been lost.

Delays and drops result in network degradation, and for some applications can cause serious problems. For example, voice traffic is sensitive to packet loss, so excessive loss will cause a deterioration of voice quality.



Egress queue modelling

Vista Manager's Intent-based QoS uses two strict priority and six WRR queues. The QoS queue types, Strict priority and WRR are described in more detail next.

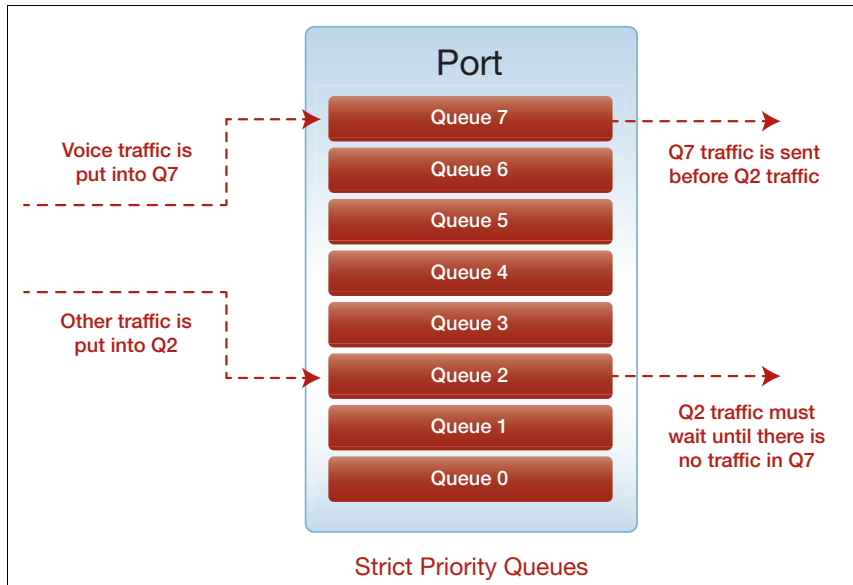


QoS egress queue types


Egress queues help with application performance by allocating a preference to outgoing traffic. For example, voice traffic could be given a high priority so it will be sent before other types of traffic.

There are two types of egress queues available, strict priority and weighted round robin:

- **Strict priority** - traffic in a higher queue is sent before traffic in a lower queue. The lowest queue is queue 0 and the highest is queue 7.



- **Weighted round robin** - queues are given a weighting. When the egress interface is congested, the specified weightings act as relative ratios to each other. For example:
 - If Q2 weight = 1 and Q5 weight = 15, then Q5 will send 15 times as much traffic as Q2.
 - If Q2 weight = 15 and Q5 weight = 15, then Q2 and Q5 will send the same amount of traffic.

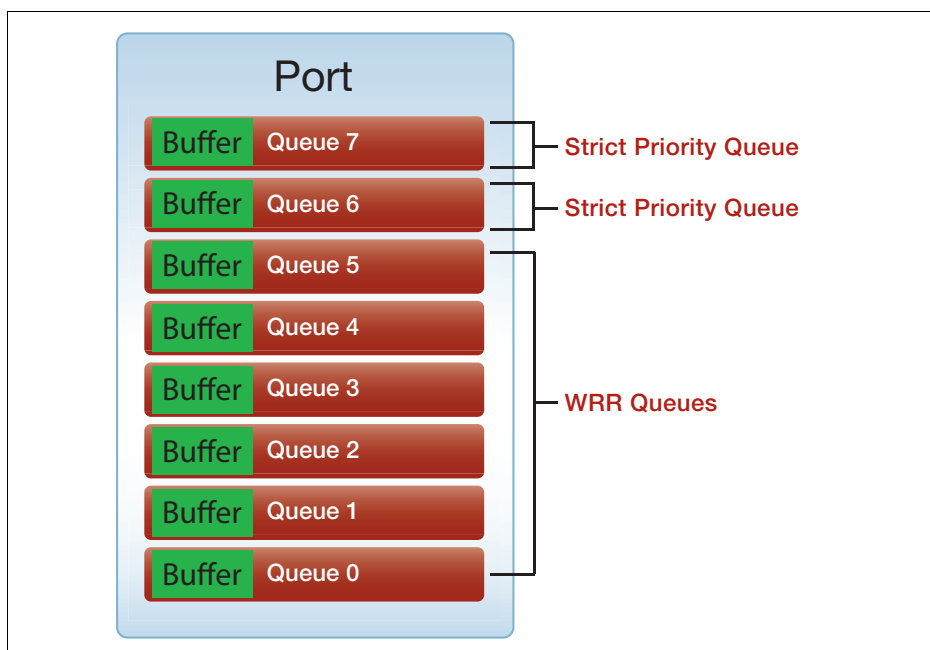


Advantages and disadvantages of WRR and Strict Priority queues

The main advantage of strict priority queues is that they ensure that drop sensitive traffic can be forwarded without loss. The difficulty with strict priority queues is that they can lead to starvation of traffic on lower priority queues.

The main advantage of WRR queues is that they ensure that at least some traffic on all queues in a WRR group is sent when congestion occurs, making full starvation of lower priority queues impossible. The difficulty with WRR queues is that some degree of packet-loss occurs on all queues when under congestion, which is problematic for applications sensitive to packet-loss.

Vista Manager EX uses two strict priority queues with egress-rate-limiting and six WRR queues. This ensures forwarding of drop sensitive traffic, while also ensuring that starvation doesn't occur on the lower priority queues.



The initial default configuration ensures that:

- Packets are marked and put into an appropriate queue.
- Queues types are set and configured with the appropriate weight and bandwidth settings:
 - Strict Priority queues for the high priority traffic (queues 7 Voice and 6 Video).
 - Weighted Round Robin (WRR) queues for the lower priority queues (all other queues).
- Interfaces are configured for QoS.

Configuring the Vista Manager EX default policy

The initial manual configuration includes: enabling QoS on devices, creating a default policy 'VISTA_DEFAULT_POLICY', and applying the default policy to interfaces.

- For platforms:

1. Enable QoS on all devices to be managed by Intent-based QoS:

```
mls qos enable
```

2. Create a QoS policy VISTA_DEFAULT_POLICY and apply it to all ports that you want Intent-based QoS to monitor and manage:

```
policy-map VISTA_DEFAULT_POLICY
  trust dscp
  class default
```

3. Set default policy queue weights:

For platforms: x8100, x220 and GS980M, x530 and GS980MX, x320 and GS980EM:

On these platforms the weights can be configured to any multiple of 17 that you choose (between 17 and 255). The reason for this is that AlliedWare Plus platforms, aside from the ones listed above, only support weightings between 1 and 15.

For Vista Manager EX to support both platform weightings - i.e. some between 1 and 15 and some between 17 and 255 in a single network, the number of possible weightings on platforms which support 17 and 255 has to be reduced to 15 possible combinations. $255/17=15$, hence why these platforms must be configured as with a weighting which is a multiple of 17.

If the weight is not a multiple of 17, then when the configuration is updated by Vista Manager, it will be updated to a multiple of 17.

- Set the scheduler to configure the WRR queue weights.

```
mls qos scheduler-set 1 wrr-queue group 1 weight 255 queue 0
mls qos scheduler-set 1 wrr-queue group 1 weight 11 queue 1
mls qos scheduler-set 1 wrr-queue group 1 weight 40 queue 2
mls qos scheduler-set 1 wrr-queue group 1 weight 104 queue 3
mls qos scheduler-set 1 wrr-queue group 1 weight 104 queue 4
mls qos scheduler-set 1 wrr-queue group 1 weight 70 queue 5
```

- The QoS policy must then be applied to each interface that will use QoS.

In addition to this the queue weights and egress rate **limits** must be set on each queue. The egress-rate limit can be set to whatever values you choose. Here they are set to 333m and 100m on a 1Gig link, this is equivalent to 33% and 10% of the total bandwidth of this interface.

```
interface port1.0.1
  service-policy input VISTA_DEFAULT_POLICY
  strict-priority-queue egress-rate-limit 333m queues 6
  strict-priority-queue egress-rate-limit 100m queues 7
  mls qos scheduler-set 1
```

■ **For other platforms:**

For other platforms, the configuration is slightly different. Instead of having a scheduler-set the weights, they are applied individually to each interface. The egress-rate limit can be set to whatever values you choose. Here they are set to 333m and 100m on a 1Gig link, this is equivalent to 33% and 10% of the total bandwidth of this interface. The percentage values must be consistent across the entire network. If queue 7 is set to the equivalent of 10% on one interface, then it must be the same percentage for all other interfaces.

```
interface port1.0.1
service-policy input VISTA_DEFAULT_POLICY
wrr-queue weight 15 queues 0
wrr-queue weight 1 queues 1
wrr-queue weight 3 queues 2
wrr-queue weight 6 queues 3
wrr-queue weight 6 queues 4
wrr-queue weight 4 queues 5
strict-priority-queue egress-rate-limit 333m queues 6
strict-priority-queue egress-rate-limit 100m queues 7
```

QoS mapping traffic to the right queue

For platforms: x8100, x220 and GS980M, x530 and GS980MX, x320, and GS980E:

Additionally, you will need to ensure the right traffic ends up on the right queues. Here are two possible ways of doing this, but it's entirely up to you how this is done.

1: Mapping from DSCP values to queues

One way to achieve this is with the following configuration that uses the existing DSCP value on each packet to map the packet into the specified queue.

```
mls qos map premark-dscp 8 to new-queue 1
mls qos map premark-dscp 10 to new-queue 2
mls qos map premark-dscp 12 to new-queue 2
mls qos map premark-dscp 14 to new-queue 2
mls qos map premark-dscp 16 to new-queue 5
mls qos map premark-dscp 18 to new-queue 3
mls qos map premark-dscp 20 to new-queue 3
mls qos map premark-dscp 22 to new-queue 3
mls qos map premark-dscp 24 to new-queue 5
mls qos map premark-dscp 26 to new-queue 4
mls qos map premark-dscp 28 to new-queue 4
mls qos map premark-dscp 30 to new-queue 4
mls qos map premark-dscp 32 to new-queue 6
mls qos map premark-dscp 34 to new-queue 6
mls qos map premark-dscp 36 to new-queue 6
mls qos map premark-dscp 38 to new-queue 6
mls qos map premark-dscp 40 to new-queue 6
mls qos map premark-dscp 46 to new-queue 7
mls qos map premark-dscp 48 to new-queue 5
mls qos map premark-dscp 56 to new-queue 5
```

2: Mapping from CoS to DSCP to queue:

Alternatively, if CoS is being used then it can first be mapped to a DSCP value on the edge of the network, and then on the internal parts of the network, the previous configuration can be used.

To map the CoS values to DSCP values the following configuration can be used, the VISTA_DEFAULT_POLICY will then need to be applied to each interface as described above.

```
mls qos map premark-dscp 8 to new-queue 1
mls qos map premark-dscp 10 to new-queue 2
mls qos map premark-dscp 12 to new-queue 2
mls qos map premark-dscp 14 to new-queue 2
mls qos map premark-dscp 16 to new-queue 5
mls qos map premark-dscp 18 to new-queue 3
mls qos map premark-dscp 20 to new-queue 3
mls qos map premark-dscp 22 to new-queue 3
mls qos map premark-dscp 24 to new-queue 5
mls qos map premark-dscp 26 to new-queue 4
mls qos map premark-dscp 28 to new-queue 4
mls qos map premark-dscp 30 to new-queue 4
mls qos map premark-dscp 32 to new-queue 6
mls qos map premark-dscp 34 to new-queue 6
mls qos map premark-dscp 36 to new-queue 6
mls qos map premark-dscp 38 to new-queue 6
mls qos map premark-dscp 40 to new-queue 6
mls qos map premark-dscp 46 to new-queue 7
mls qos map premark-dscp 48 to new-queue 5
mls qos map premark-dscp 56 to new-queue 5
mls qos scheduler-set 1 wrr-queue group 1 weight 255 queues 0
mls qos scheduler-set 1 wrr-queue group 1 weight 11 queues 1
mls qos scheduler-set 1 wrr-queue group 1 weight 40 queues 2
mls qos scheduler-set 1 wrr-queue group 1 weight 104 queues 3
mls qos scheduler-set 1 wrr-queue group 1 weight 104 queues 4
mls qos scheduler-set 1 wrr-queue group 1 weight 70 queues 5
class-map COS-DSCP_TRANSLATE_7
  match cos 7

class-map COS-DSCP_TRANSLATE_6
  match cos 6

class-map COS-DSCP_TRANSLATE_5
  match cos 5

class-map COS-DSCP_TRANSLATE_4
  match cos 4

class-map COS-DSCP_TRANSLATE_3
  match cos 3

class-map COS-DSCP_TRANSLATE_2
  match cos 2

class-map COS-DSCP_TRANSLATE_1
  match cos 1

policy-map VISTA_DEFAULT_POLICY
  trust dscp
  class default
  class COS-DSCP_TRANSLATE_7
    set dscp 56
    set queue 5
  class COS-DSCP_TRANSLATE_6
    set dscp 48
```

```

set queue 5
class COS-DSCP_TRANSLATE_5
set dscp 46
set queue 7
class COS-DSCP_TRANSLATE_4
set dscp 34
set queue 6
class COS-DSCP_TRANSLATE_3
set dscp 26
set queue 4
class COS-DSCP_TRANSLATE_2
set dscp 18
set queue 3
class COS-DSCP_TRANSLATE_1
set dscp 10
set queue 2

```

QoS mapping traffic to the right queue - For other platforms

1: Mapping from DSCP values to queues

```

mls qos map cos-queue 0 to 0
mls qos map cos-queue 1 to 1
mls qos map cos-queue 2 to 2
mls qos enable
mls qos map premark-dscp 8 to new-cos 1
mls qos map premark-dscp 10 to new-cos 2
mls qos map premark-dscp 12 to new-cos 2
mls qos map premark-dscp 14 to new-cos 2
mls qos map premark-dscp 16 to new-cos 5
mls qos map premark-dscp 18 to new-cos 3
mls qos map premark-dscp 20 to new-cos 3
mls qos map premark-dscp 22 to new-cos 3
mls qos map premark-dscp 24 to new-cos 5
mls qos map premark-dscp 26 to new-cos 4
mls qos map premark-dscp 28 to new-cos 4
mls qos map premark-dscp 30 to new-cos 4
mls qos map premark-dscp 32 to new-cos 6
mls qos map premark-dscp 34 to new-cos 6
mls qos map premark-dscp 36 to new-cos 6
mls qos map premark-dscp 38 to new-cos 6
mls qos map premark-dscp 40 to new-cos 6
mls qos map premark-dscp 46 to new-cos 7
mls qos map premark-dscp 48 to new-cos 5
mls qos map premark-dscp 56 to new-cos 5

```

2: Mapping from DSCP values to queues

```

mls qos map cos-queue 0 to 0
mls qos map cos-queue 1 to 1
mls qos map cos-queue 2 to 2
mls qos enable
mls qos map premark-dscp 8 to new-cos 1
mls qos map premark-dscp 10 to new-cos 2
mls qos map premark-dscp 12 to new-cos 2
mls qos map premark-dscp 14 to new-cos 2
mls qos map premark-dscp 16 to new-cos 5
mls qos map premark-dscp 18 to new-cos 3
mls qos map premark-dscp 20 to new-cos 3
mls qos map premark-dscp 22 to new-cos 3
mls qos map premark-dscp 24 to new-cos 5
mls qos map premark-dscp 26 to new-cos 4
mls qos map premark-dscp 28 to new-cos 4

```



```
mls qos map premark-dscp 30 to new-cos 4
mls qos map premark-dscp 32 to new-cos 6
mls qos map premark-dscp 34 to new-cos 6
mls qos map premark-dscp 36 to new-cos 6
mls qos map premark-dscp 38 to new-cos 6
mls qos map premark-dscp 40 to new-cos 6
mls qos map premark-dscp 46 to new-cos 7
mls qos map premark-dscp 48 to new-cos 5
mls qos map premark-dscp 56 to new-cos 5
!
class-map COS_7
  match cos 7
!
class-map COS_6
  match cos 6
!
class-map COS_5
  match cos 5
!
class-map COS_4
  match cos 4
!
class-map COS_3
  match cos 3
!
class-map COS_2
  match cos 2
!
class-map COS_1
  match cos 1
!
class-map EF
  match dscp 46
!
class-map CS7
  match dscp 56
!
class-map CS6
  match dscp 48
!
class-map CS5
  match dscp 40
!
class-map CS4
  match dscp 32
!
class-map CS3
  match dscp 24
!
class-map CS2
  match dscp 16
!
class-map CS1
  match dscp 8
!
class-map AF41
  match dscp 34
!
class-map AF42
  match dscp 36
!
class-map AF43
  match dscp 38
!
```

```

class-map AF31
  match dscp 26
!
class-map AF32
  match dscp 28
!
class-map AF33
  match dscp 30
!
class-map AF21
  match dscp 18
!
class-map AF22
  match dscp 20
!
class-map AF23
  match dscp 22
!
class-map AF11
  match dscp 10
!
class-map AF12
  match dscp 12
!
class-map AF13
  match dscp 14
!
policy-map VISTA_DEFAULT_POLICY
  class default
    remark new-cos 0 internal
  class COS_7
    remark new-cos 5 internal
    remark-map to new-dscp 56
  class COS_6
    remark new-cos 5 internal
    remark-map to new-dscp 48
  class COS_5
    remark new-cos 7 internal
    remark-map to new-dscp 46
  class COS_4
    remark new-cos 6 internal
    remark-map to new-dscp 34
  class COS_3
    remark new-cos 4 internal
    remark-map to new-dscp 26
  class COS_2
    remark new-cos 3 internal
    remark-map to new-dscp 18
  class COS_1
    remark new-cos 2 internal
    remark-map to new-dscp 10
  class EF
    remark new-cos 7 internal
  class CS7
    remark new-cos 5 internal
  class CS6
    remark new-cos 5 internal
  class CS3
    remark new-cos 5 internal
  class CS2
    remark new-cos 5 internal
  class CS5
    remark new-cos 6 internal
  class CS4

```

```
    remark new-cos 6 internal
class AF41
    remark new-cos 6 internal
class AF42
    remark new-cos 6 internal
class AF43
    remark new-cos 6 internal
class AF31
    remark new-cos 4 internal
class AF32
    remark new-cos 4 internal
class AF33
    remark new-cos 4 internal
class AF21
    remark new-cos 3 internal
class AF22
    remark new-cos 3 internal
class AF23
    remark new-cos 3 internal
class AF11
    remark new-cos 2 internal
class AF12
    remark new-cos 2 internal
class AF13
    remark new-cos 2 internal
class CS1
    remark new-cos 1 internal
```

Complete configuration example - for the x220 and x230 series switches

You could use the following configuration on an access switch. The configuration for distribution and core switches would largely be identical, except that the configured egress-rate-limiting would occur on all ports, not just on uplinks.

x220 - Access CoS to DSCP

```
x230#show run
!
service password-encryption
!
hostname x230
!
no banner motd
!
username manager privilege 15 password 8 $1$bJoVec4D$JwOJGPr7YqoExA0GVasdE0
!
!
no service ssh
!
service telnet
!
service http
!
no clock timezone
!
snmp-server
!
!
!
aaa authentication enable default local
aaa authentication login default local
!
!
!
!
!
ip domain-lookup
!
!
!
no service dhcp-server
!
spanning-tree mode rstp
!
service power-inline
no lacp global-passive-mode enable
!
mls qos map cos-queue 0 to 0
mls qos map cos-queue 1 to 1
mls qos map cos-queue 2 to 2
mls qos enable
mls qos map premark-dscp 8 to new-cos 1
mls qos map premark-dscp 10 to new-cos 2
mls qos map premark-dscp 12 to new-cos 2
mls qos map premark-dscp 14 to new-cos 2
mls qos map premark-dscp 16 to new-cos 5
mls qos map premark-dscp 18 to new-cos 3
mls qos map premark-dscp 20 to new-cos 3
mls qos map premark-dscp 22 to new-cos 3
```

```
mls qos map premark-dscp 24 to new-cos 5
mls qos map premark-dscp 26 to new-cos 4
mls qos map premark-dscp 28 to new-cos 4
mls qos map premark-dscp 30 to new-cos 4
mls qos map premark-dscp 32 to new-cos 6
mls qos map premark-dscp 34 to new-cos 6
mls qos map premark-dscp 36 to new-cos 6
mls qos map premark-dscp 38 to new-cos 6
mls qos map premark-dscp 40 to new-cos 6
mls qos map premark-dscp 46 to new-cos 7
mls qos map premark-dscp 48 to new-cos 5
mls qos map premark-dscp 56 to new-cos 5
!
class-map COS_7
  match cos 7
!
class-map COS_6
  match cos 6
!
class-map COS_5
  match cos 5
!
class-map COS_4
  match cos 4
!
class-map COS_3
  match cos 3
!
class-map COS_2
  match cos 2
!
class-map COS_1
  match cos 1
!
class-map EF
  match dscp 46
!
class-map CS7
  match dscp 56
!
class-map CS6
  match dscp 48
!
class-map CS5
  match dscp 40
!
class-map CS4
  match dscp 32
!
class-map CS3
  match dscp 24
!
class-map CS2
  match dscp 16
!
class-map CS1
  match dscp 8
!
class-map AF41
  match dscp 34
!
class-map AF42
  match dscp 36
!
```

```

class-map AF43
  match dscp 38
!
class-map AF31
  match dscp 26
!
class-map AF32
  match dscp 28
!
class-map AF33
  match dscp 30
!
class-map AF21
  match dscp 18
!
class-map AF22
  match dscp 20
!
class-map AF23
  match dscp 22
!
class-map AF11
  match dscp 10
!
class-map AF12
  match dscp 12
!
class-map AF13
  match dscp 14
!
policy-map VISTA_DEFAULT_POLICY
  class default
    remark new-cos 0 internal
  class COS_7
    remark new-cos 5 internal
    remark-map to new-dscp 56
  class COS_6
    remark new-cos 5 internal
    remark-map to new-dscp 48
  class COS_5
    remark new-cos 7 internal
    remark-map to new-dscp 46
  class COS_4
    remark new-cos 6 internal
    remark-map to new-dscp 34
  class COS_3
    remark new-cos 4 internal
    remark-map to new-dscp 26
  class COS_2
    remark new-cos 3 internal
    remark-map to new-dscp 18
  class COS_1
    remark new-cos 2 internal
    remark-map to new-dscp 10
  class EF
    remark new-cos 7 internal
  class CS7
    remark new-cos 5 internal
  class CS6
    remark new-cos 5 internal
  class CS3
    remark new-cos 5 internal
  class CS2
    remark new-cos 5 internal

```

```

class CS5
  remark new-cos 6 internal
class CS4
  remark new-cos 6 internal
class AF41
  remark new-cos 6 internal
class AF42
  remark new-cos 6 internal
class AF43
  remark new-cos 6 internal
class AF31
  remark new-cos 4 internal
class AF32
  remark new-cos 4 internal
class AF33
  remark new-cos 4 internal
class AF21
  remark new-cos 3 internal
class AF22
  remark new-cos 3 internal
class AF23
  remark new-cos 3 internal
class AF11
  remark new-cos 2 internal
class AF12
  remark new-cos 2 internal
class AF13
  remark new-cos 2 internal
class CS1
  remark new-cos 1 internal
!
!
interface port1.0.1-1.0.16
  switchport
  switchport mode access
  service-policy input VISTA_DEFAULT_POLICY
!
interface port1.0.17-1.0.18
  switchport
  switchport mode access
  service-policy input VISTA_DEFAULT_POLICY
  wrr-queue weight 15 queues 0
  wrr-queue weight 1 queues 1
  wrr-queue weight 3 queues 2
  wrr-queue weight 6 queues 3
  wrr-queue weight 6 queues 4
  wrr-queue weight 4 queues 5
  wrr-queue egress-rate-limit 333m queues 6
  wrr-queue egress-rate-limit 100m queues 7
!
!
line con 0
line vty 0 4
!
end

```

x220 Access basic

```

x220#show run
!
service password-encryption
!
hostname x220

```

```

!
no banner motd
!
username manager privilege 15 password 8 $1$bJoVec4D$JwOJGPr7YqoExA0GVasdE0
!
!
no service ssh
!
service telnet
!
service http
!
no clock timezone
!
snmp-server
!
!
!
aaa authentication enable default local
aaa authentication login default local
!
!
!
!
!
ip domain-lookup
!
!
spanning-tree mode rstp
!
service power-inline
lACP global-passive-mode enable
!
mls qos enable
mls qos map premark-dscp 8 to new-queue 1
mls qos map premark-dscp 10 to new-queue 2
mls qos map premark-dscp 12 to new-queue 2
mls qos map premark-dscp 14 to new-queue 2
mls qos map premark-dscp 16 to new-queue 5
mls qos map premark-dscp 18 to new-queue 3
mls qos map premark-dscp 20 to new-queue 3
mls qos map premark-dscp 22 to new-queue 3
mls qos map premark-dscp 24 to new-queue 5
mls qos map premark-dscp 26 to new-queue 4
mls qos map premark-dscp 28 to new-queue 4
mls qos map premark-dscp 30 to new-queue 4
mls qos map premark-dscp 32 to new-queue 6
mls qos map premark-dscp 34 to new-queue 6
mls qos map premark-dscp 36 to new-queue 6
mls qos map premark-dscp 38 to new-queue 6
mls qos map premark-dscp 40 to new-queue 6
mls qos map premark-dscp 46 to new-queue 7
mls qos map premark-dscp 48 to new-queue 5
mls qos map premark-dscp 56 to new-queue 5
mls qos scheduler-set 1 wrr-queue group 1 weight 255 queues 0
mls qos scheduler-set 1 wrr-queue group 1 weight 11 queues 1
mls qos scheduler-set 1 wrr-queue group 1 weight 40 queues 2
mls qos scheduler-set 1 wrr-queue group 1 weight 104 queues 3
mls qos scheduler-set 1 wrr-queue group 1 weight 104 queues 4
mls qos scheduler-set 1 wrr-queue group 1 weight 70 queues 5
!
!
policy-map VISTA_DEFAULT_POLICY
  trust dscp

```



```

class default
!
interface port1.0.1-1.0.47
  switchport
  switchport mode access
  service-policy input VISTA_DEFAULT_POLICY
!
interface port1.0.48-1.0.50
  switchport
  switchport mode access
  service-policy input VISTA_DEFAULT_POLICY
  wrr-queue egress-rate-limit 333m queues 6
  wrr-queue egress-rate-limit 100m queues 7
  mls qos scheduler-set 1
!
interface port1.0.51-1.0.52
  switchport
  switchport mode access
!
line con 0
line vty 0 4
!
end

```

x220 Distribution or Core basic

```

x220#show run
!
service password-encryption
!
hostname x220
!
no banner motd
!
username manager privilege 15 password 8 $1$bJoVec4D$JwOJGPr7YqoExA0GVasdE0
!
!
no service ssh
!
service telnet
!
service http
!
no clock timezone
!
snmp-server
!
!
!
aaa authentication enable default local
aaa authentication login default local
!
!
!
!
!
ip domain-lookup
!
!
spanning-tree mode rstp
!

```

```

service power-inline
lACP global-passive-mode enable
!
mls qos enable
mls qos map premark-dscp 8 to new-queue 1
mls qos map premark-dscp 10 to new-queue 2
mls qos map premark-dscp 12 to new-queue 2
mls qos map premark-dscp 14 to new-queue 2
mls qos map premark-dscp 16 to new-queue 5
mls qos map premark-dscp 18 to new-queue 3
mls qos map premark-dscp 20 to new-queue 3
mls qos map premark-dscp 22 to new-queue 3
mls qos map premark-dscp 24 to new-queue 5
mls qos map premark-dscp 26 to new-queue 4
mls qos map premark-dscp 28 to new-queue 4
mls qos map premark-dscp 30 to new-queue 4
mls qos map premark-dscp 32 to new-queue 6
mls qos map premark-dscp 34 to new-queue 6
mls qos map premark-dscp 36 to new-queue 6
mls qos map premark-dscp 38 to new-queue 6
mls qos map premark-dscp 40 to new-queue 6
mls qos map premark-dscp 46 to new-queue 7
mls qos map premark-dscp 48 to new-queue 5
mls qos map premark-dscp 56 to new-queue 5
mls qos scheduler-set 1 wrr-queue group 1 weight 255 queues 0
mls qos scheduler-set 1 wrr-queue group 1 weight 11 queues 1
mls qos scheduler-set 1 wrr-queue group 1 weight 40 queues 2
mls qos scheduler-set 1 wrr-queue group 1 weight 104 queues 3
mls qos scheduler-set 1 wrr-queue group 1 weight 104 queues 4
mls qos scheduler-set 1 wrr-queue group 1 weight 70 queues 5
!
!
policy-map VISTA_DEFAULT_POLICY
  trust dscp
  class default
!
interface port1.0.1-1.0.50
  switchport
  switchport mode access
  service-policy input VISTA_DEFAULT_POLICY
  wrr-queue egress-rate-limit 333m queues 6
  wrr-queue egress-rate-limit 100m queues 7
  mls qos scheduler-set 1
!
interface port1.0.51-1.0.52
  switchport
  switchport mode access
!
line con 0
line vty 0 4
!
end

```

x230 Access CoS to DSCP

```

x230#show run
!
service password-encryption
!
hostname x230
!
no banner motd

```

```

!
username manager privilege 15 password 8 $1$bJoVec4D$JwOJGPr7YqoExA0GVasdE0
!
!
no service ssh
!
service telnet
!
service http
!
no clock timezone
!
snmp-server
!
!
!
aaa authentication enable default local
aaa authentication login default local
!
!
!
!
!
ip domain-lookup
!
!
!
no service dhcp-server
!
spanning-tree mode rstp
!
service power-inline
no lacp global-passive-mode enable
!
mls qos map cos-queue 0 to 0
mls qos map cos-queue 1 to 1
mls qos map cos-queue 2 to 2
mls qos enable
mls qos map premark-dscp 8 to new-cos 1
mls qos map premark-dscp 10 to new-cos 2
mls qos map premark-dscp 12 to new-cos 2
mls qos map premark-dscp 14 to new-cos 2
mls qos map premark-dscp 16 to new-cos 5
mls qos map premark-dscp 18 to new-cos 3
mls qos map premark-dscp 20 to new-cos 3
mls qos map premark-dscp 22 to new-cos 3
mls qos map premark-dscp 24 to new-cos 5
mls qos map premark-dscp 26 to new-cos 4
mls qos map premark-dscp 28 to new-cos 4
mls qos map premark-dscp 30 to new-cos 4
mls qos map premark-dscp 32 to new-cos 6
mls qos map premark-dscp 34 to new-cos 6
mls qos map premark-dscp 36 to new-cos 6
mls qos map premark-dscp 38 to new-cos 6
mls qos map premark-dscp 40 to new-cos 6
mls qos map premark-dscp 46 to new-cos 7
mls qos map premark-dscp 48 to new-cos 5
mls qos map premark-dscp 56 to new-cos 5
!
class-map COS_7
  match cos 7
!
class-map COS_6
  match cos 6

```

```
!  
class-map COS_5  
  match cos 5  
!  
class-map COS_4  
  match cos 4  
!  
class-map COS_3  
  match cos 3  
!  
class-map COS_2  
  match cos 2  
!  
class-map COS_1  
  match cos 1  
!  
class-map EF  
  match dscp 46  
!  
class-map CS7  
  match dscp 56  
!  
class-map CS6  
  match dscp 48  
!  
class-map CS5  
  match dscp 40  
!  
class-map CS4  
  match dscp 32  
!  
class-map CS3  
  match dscp 24  
!  
class-map CS2  
  match dscp 16  
!  
class-map CS1  
  match dscp 8  
!  
class-map AF41  
  match dscp 34  
!  
class-map AF42  
  match dscp 36  
!  
class-map AF43  
  match dscp 38  
!  
class-map AF31  
  match dscp 26  
!  
class-map AF32  
  match dscp 28  
!  
class-map AF33  
  match dscp 30  
!  
class-map AF21  
  match dscp 18  
!  
class-map AF22  
  match dscp 20  
!
```

```

class-map AF23
  match dscp 22
!
class-map AF11
  match dscp 10
!
class-map AF12
  match dscp 12
!
class-map AF13
  match dscp 14
!
policy-map VISTA_DEFAULT_POLICY_DOWNLINK
  class default
    remark new-cos 0 internal
  class COS_7
    remark new-cos 5 internal
    remark-map to new-dscp 56
  class COS_6
    remark new-cos 5 internal
    remark-map to new-dscp 48
  class COS_5
    remark new-cos 7 internal
    remark-map to new-dscp 46
  class COS_4
    remark new-cos 6 internal
    remark-map to new-dscp 34
  class COS_3
    remark new-cos 4 internal
    remark-map to new-dscp 26
  class COS_2
    remark new-cos 3 internal
    remark-map to new-dscp 18
  class COS_1
    remark new-cos 2 internal
    remark-map to new-dscp 10
  class EF
    remark new-cos 7 internal
  class CS7
    remark new-cos 5 internal
  class CS6
    remark new-cos 5 internal
  class CS3
    remark new-cos 5 internal
  class CS2
    remark new-cos 5 internal
  class CS5
    remark new-cos 6 internal
  class CS4
    remark new-cos 6 internal
  class AF41
    remark new-cos 6 internal
  class AF42
    remark new-cos 6 internal
  class AF43
    remark new-cos 6 internal
  class AF31
    remark new-cos 4 internal
  class AF32
    remark new-cos 4 internal
  class AF33
    remark new-cos 4 internal
  class AF21
    remark new-cos 3 internal

```

```

class AF22
  remark new-cos 3 internal
class AF23
  remark new-cos 3 internal
class AF11
  remark new-cos 2 internal
class AF12
  remark new-cos 2 internal
class AF13
  remark new-cos 2 internal
class CS1
  remark new-cos 1 internal
!
policy-map VISTA_DEFAULT_POLICY_UPLINK
  trust dscp
  class default
!
interface port1.0.1-1.0.16
  switchport
  switchport mode access
  service-policy input VISTA_DEFAULT_POLICY_DOWNLINK
!
interface port1.0.17-1.0.18
  switchport
  switchport mode access
  service-policy input VISTA_DEFAULT_POLICY_UPLINK
  wrr-queue weight 15 queues 0
  wrr-queue weight 1 queues 1
  wrr-queue weight 3 queues 2
  wrr-queue weight 6 queues 3
  wrr-queue weight 6 queues 4
  wrr-queue weight 4 queues 5
  wrr-queue egress-rate-limit 333m queues 6
  wrr-queue egress-rate-limit 100m queues 7
!!
line con 0
line vty 0 4
!
end

```

x230 Access basic

```

!
service password-encryption
!
hostname x230
!
no banner motd
!
username manager privilege 15 password 8 $1$bJoVec4D$JwOJGPr7YqoExA0GVasdE0
!
!
no service ssh
!
service telnet
!
service http
!
no clock timezone
!
snmp-server
!
aaa authentication enable default local
aaa authentication login default local

```

```

!
ip domain-lookup
!
!
!
no service dhcp-server
!
spanning-tree mode rstp
!
service power-inline
no lacp global-passive-mode enable
!
mls qos map cos-queue 0 to 0
mls qos map cos-queue 1 to 1
mls qos map cos-queue 2 to 2
mls qos enable
mls qos map premark-dscp 8 to new-cos 1
mls qos map premark-dscp 10 to new-cos 2
mls qos map premark-dscp 12 to new-cos 2
mls qos map premark-dscp 14 to new-cos 2
mls qos map premark-dscp 16 to new-cos 5
mls qos map premark-dscp 18 to new-cos 3
mls qos map premark-dscp 20 to new-cos 3
mls qos map premark-dscp 22 to new-cos 3
mls qos map premark-dscp 24 to new-cos 5
mls qos map premark-dscp 26 to new-cos 4
mls qos map premark-dscp 28 to new-cos 4
mls qos map premark-dscp 30 to new-cos 4
mls qos map premark-dscp 32 to new-cos 6
mls qos map premark-dscp 34 to new-cos 6
mls qos map premark-dscp 36 to new-cos 6
mls qos map premark-dscp 38 to new-cos 6
mls qos map premark-dscp 40 to new-cos 6
mls qos map premark-dscp 46 to new-cos 7
mls qos map premark-dscp 48 to new-cos 5
mls qos map premark-dscp 56 to new-cos 5
!
policy-map VISTA_DEFAULT_POLICY
  trust dscp
  class default
!
interface port1.0.1-1.0.16
  switchport
  switchport mode access
  service-policy input VISTA_DEFAULT_POLICY
!
interface port1.0.17-1.0.18
  switchport
  switchport mode access
  service-policy input VISTA_DEFAULT_POLICY
  wrr-queue weight 15 queues 0
  wrr-queue weight 1 queues 1
  wrr-queue weight 3 queues 2
  wrr-queue weight 6 queues 3
  wrr-queue weight 6 queues 4
  wrr-queue weight 4 queues 5
  wrr-queue egress-rate-limit 333m queues 6
  wrr-queue egress-rate-limit 100m queues 7
!
line con 0
line vty 0 4
!
end

```

x230 Distribution or Core basic

```

!
service password-encryption
!
hostname x230
!
no banner motd
!
username manager privilege 15 password 8 $1$bJoVec4D$JwOJGPr7YqoExA0GVasdE0
!
!
no service ssh
!
service telnet
!
service http
!
no clock timezone
!
snmp-server
!
!
!
aaa authentication enable default local
aaa authentication login default local
!
!
!
!
ip domain-lookup
!
!
!
no service dhcp-server
!
spanning-tree mode rstp
!
service power-inline
no lacp global-passive-mode enable
!
mls qos map cos-queue 0 to 0
mls qos map cos-queue 1 to 1
mls qos map cos-queue 2 to 2
mls qos enable
mls qos map premark-dscp 8 to new-cos 1
mls qos map premark-dscp 10 to new-cos 2
mls qos map premark-dscp 12 to new-cos 2
mls qos map premark-dscp 14 to new-cos 2
mls qos map premark-dscp 16 to new-cos 5
mls qos map premark-dscp 18 to new-cos 3
mls qos map premark-dscp 20 to new-cos 3
mls qos map premark-dscp 22 to new-cos 3
mls qos map premark-dscp 24 to new-cos 5
mls qos map premark-dscp 26 to new-cos 4
mls qos map premark-dscp 28 to new-cos 4
mls qos map premark-dscp 30 to new-cos 4
mls qos map premark-dscp 32 to new-cos 6
mls qos map premark-dscp 34 to new-cos 6
mls qos map premark-dscp 36 to new-cos 6
mls qos map premark-dscp 38 to new-cos 6
mls qos map premark-dscp 40 to new-cos 6
mls qos map premark-dscp 46 to new-cos 7
mls qos map premark-dscp 48 to new-cos 5

```



```
mls qos map premark-dscp 56 to new-cos 5
!
policy-map VISTA_DEFAULT_POLICY
  trust dscp
  class default
!
interface port1.0.1-1.0.18
  switchport
  switchport mode access
  service-policy input VISTA_DEFAULT_POLICY
  wrr-queue weight 15 queues 0
  wrr-queue weight 1 queues 1
  wrr-queue weight 3 queues 2
  wrr-queue weight 6 queues 3
  wrr-queue weight 6 queues 4
  wrr-queue weight 4 queues 5
  wrr-queue egress-rate-limit 333m queues 6
  wrr-queue egress-rate-limit 100m queues 7
!
line con 0
line vty 0 4
!
end
```

Configuring AMF Plus Nodes: the Unified CLI


The unified CLI is a central component of AMF Plus. It provides you with a configuration and display interface that can control a selected collection of nodes, or the entire AMF Plus area, from a single point. This control is provided through the **atmf working-set** command.

Working-sets

Conceptually a **working-set** is a collection of switches that can be configured centrally as if they are a single device. A working-set may comprise a predefined group that has been automatically created based on some common set of physical attributes such as switch type etc., or it may be an arbitrary set of devices created by a network administrator to simplify configuration.

Specifying or selecting a working-set allows a single CLI command to be executed on all nodes within the selected working-set, from a single device. A working-set can be defined, selected and configured from any node within an AMF Plus network (unless **restricted-login** is enabled). All the members of a working-set must reside in the same AMF Plus area. It is **not** possible for a working-set to span areas. The following figure shows a number of switches that comprise a working-set.

Figure 1: AMF Plus working-set



Note: For security reasons **restricted-login** limits the action of working-sets on the AMF Plus network. See "[AMF Plus restricted-login](#)" on page 235 for more information.

Note: In secure mode **restricted-login** is on by default and cannot be turned off.

Local working-set

By default, when you first log into a node that is part of an AMF Plus network, you are implicitly placed into the working-set group **local**, a working-set that contains only the local node. In this instance the CLI prompt when you log in will be either:

- the host name, if one has been assigned, or
- in the case of a new node in safe mode, a host name based on its MAC address followed by the usual prompt (> or #)

```
Node1> enable
Node1#
```

Creating a working-set

To create a working-set, use the command **atmf working-set** followed by a comma separated list of the nodes you want to comprise the working-set.

Whenever you select a working-set containing any nodes other than the local device, the CLI prompt will display the AMF Plus network name (set using the **atmf network-name** command) followed by the number of nodes in the working-set. From the example below, **atmf1[2]** is a 2 node working-set on the atm1 network.

```
Node1# atmf working-set Node1,Node2
Node1,Node2
Working set join
atmf1[2]#
```

Once you have joined the working-set, any command that you type in will be sent to all the members of the working-set.

To return to controlling just the local device from any other working-set, use the command: **atmf working-set group local**.

Working-set groups

AMF Plus includes the ability to have working-set groups, so that it is not always necessary to use a comma separated list to specify a working-set. AMF Plus working-set groups can be split into two types:

- automatic
- user-defined

Automatic working-set groups

There are three automatic working-set groups that will exist on every AMF Plus network:

1. **All**—all nodes within the AMF Plus area.
2. **Current**—the current working-set of nodes. This group is useful for adding additional nodes to the current working-set.
3. **Local**—the local device.

In any AMF Plus area there will also be a number of other automatic working-set groups that are dependent on the platform types which exist within the area.

To see the platform- dependent automatic working-set groups that exist within the AMF Plus area, use the command **show atmf group** with the **automatic** parameter:

Output 1: show atmf group members automatic

```
x930_VCS_1#show atmf group members automatic

Retrieving Automatic groups from:
x530_1 Master x930_VCS_2 x930_VCS_1

ATMF Group membership

Automatic      Total
Groups         Members  Members
poe            1        Master
x530           1        x530_1
SBx8100       1        Master
x950           2        x930_VCS_2 x930_VCS_1
```

To select a working-set group use the **atmf working-set** command with the **group** parameter, followed by the group name. You can specify a single group, a comma-separated list of groups, or a comma-separated list of individual nodes followed by a comma-separated list of groups. For example, to create a working set made up of x530_1, x530_2 and all nodes in the group named x930, use the following command:

```
x930_VCS_1# atmf working-set x530_1,x530_2 group x930
x530_1, x530_2, x930_VCS_1, x930_VCS_2
Working set join
atmf1[4]#
```

- If you specify a partially invalid working-set node list or group list, only the valid nodes or groups will join the working-set.
- If you specify a completely invalid working-set, you will create a working-set containing no nodes. The switch will generate a warning message to alert you that the current working-set is empty:

```
atmf1[3]# atmf working-set group x511
% Warning - working set is now empty
atmf1[0]#
```

User-defined working-set groups

In addition to the automatic working-set groups, you can create user-defined groups for arbitrary sets of nodes that you wish to group together, for example, all AMF Plus master nodes.

To create a user-defined working-set group:

1. Create a working-set containing the desired nodes.
2. Having joined the working-set, then in global configuration mode use the command **atmf group**.

```
Master# atmf working-set Master1,Master2
Master1,Master2
Working set join
atmf1[2]# conf t
atmf1[2]# atmf group new-group-name
```

You can see all user-defined working-set groups that exist on the AMF Plus area with the command **show atmf group members user-defined**:

Output 2: `show atmf group members user-defined`

```
x930_VCS_1#show atmf group members user-defined

Retrieving Automatic groups from:
x530_1 Master1, Master2, x930_VCS_2 x930_VCS_1

ATMF Group membership

User-defined      Total
Groups           Members  Members
-----
Masters          2       Master1 Master2

Master#
```

Executing commands on working-sets

Executing commands on a working-set of nodes is very similar to executing commands on a single AlliedWare Plus device.

When a command is executed that is valid for all nodes within the working-set, the output is displayed for each of the nodes separately. However, output will be grouped when it is the same for more than one node.

Here is an example output of the **show arp** command run from a working-set:

Output 3: show arp command output

```

atmf1[4]#show arp
=====
Master:
=====

  IP Address      MAC Address      Interface      Port      Type
  172.31.0.1      eccd.6d7d.a542   ATMF           sa1       dynamic
  172.31.0.3      0000.cd2b.0329   ATMF           sa1       dynamic
  172.31.0.10     0000.cd37.0163   ATMF           sa1       dynamic

=====
x510_1:
=====

  IP Address      MAC Address      Interface      Port      Type
  172.31.0.2      eccd.6d03.10f9   ATMF           sa4       dynamic

=====
x930_VCS_1:
=====

  IP Address      MAC Address      Interface      Port      Type
  172.31.0.2      0000.cd37.1050   ATMF           sa1       dynamic

=====
x930_VCS_2:
=====

  IP Address      MAC Address      Interface      Port      Type
  172.31.0.2      0000.cd37.1050   ATMF           sa3       dynamic

atmf1[4]#

```

Invalid working-set commands

Some commands can only be executed on particular nodes within the working-set. Initially the command will be attempted on all nodes within the working-set. However, on any node for which the command is invalid, the command execution will fail and the output displayed will indicate the nodes on which the command succeeded and nodes on which the command failed.

In the example below, output is displayed from the **show card** command run from a working-set that is only a valid command for the SBx8100 series switches.

Output 4: SBx8100 Series show card command output

```

atmf1[4]# show card
=====
Master:
=====

Slot Card Type          State
-----
1    AT-SBx81GP24       Online
2    AT-SBx81GP24       Online
3    AT-SBx81GP24       Online
4    AT-SBx81XS6        Online
5    AT-SBx81CFC400     Online (Active)
6    -                  -
7    -                  -
8    -                  -
9    -                  -
10   -                  -
11   -                  -
12   -                  -
-----

=====
x510_1, x930_VCS_1, x930_VCS_2:
=====
% Invalid input detected at '^' marker.

```

Sub-configuration limitations for some nodes in a working-set

There will be some instances where a sub-configuration mode is only valid for some of the nodes in the working-set. One example of this would be when entering interface configuration mode for a port that exists on some members of the working-set and not on others. For example:

```

atmf1[4]# conf t
atmf1[4](config)# int port1.1.1
% Can't find interface port1.1.1
atmf1[4:2](config-if)# conf t

```

In the example above the interface **port1.1.1** exists on two of the nodes in the working-set, but does not exist on nodes “Master” or “x510_1”. The interface configuration mode fails for these nodes, and a warning message is output to indicate this.

Inside the square brackets, the first number indicates the total number of nodes in the working-set, and the second number indicates the number of nodes in the sub-configuration mode that has been entered. Any configuration commands configured in this mode will only be executed on the nodes that successfully entered the sub-configuration mode. Entering **exit** while in this mode will return to global configuration mode for all nodes within the working-set:

```

atmf1[4:2](config-if)# exit
atmf1[4](config)# (config)#

```

Interactive commands

It is inappropriate to execute **interactive** commands simultaneously across multiple nodes within a working-set. These commands can only be executed on the local node working-set or on a working-set with a single member.

When any interactive commands are entered from within a working-set they will give an error:

```
atmf1[4]# ping 4.2.2.1
% Working set must contain only single node for this command
```

Interactive commands include:

- ping
- mtrace/mstat
- traceroute
- boot system
- boot configuration-file
- banner login
- tcpdump
- edit
- copy*
- mail
- move
- terminal monitor

Copying files between nodes

You can copy files between nodes in an AMF Plus networking using the copy command and adding the AMF node name suffixed with “.atmf” to the path.

For example to copy a file named “x550-5.5.2-2.3.rel” from the current directory on the AMF Plus master to flash on an AMF Plus node named “node1” use the command:

```
master# copy x550-5.5.2-2.3.rel node1.atmf/flash:
```


Node Provisioning

You can preconfigure (provision) a port for a future node before that node is physically added to the network. A provisioned node can be created as a new unique entity, or can be cloned using the backup data from an existing node. When you connect the new node to the provisioned port in the AMF Plus network, its configuration is loaded from the information stored in the backup media.

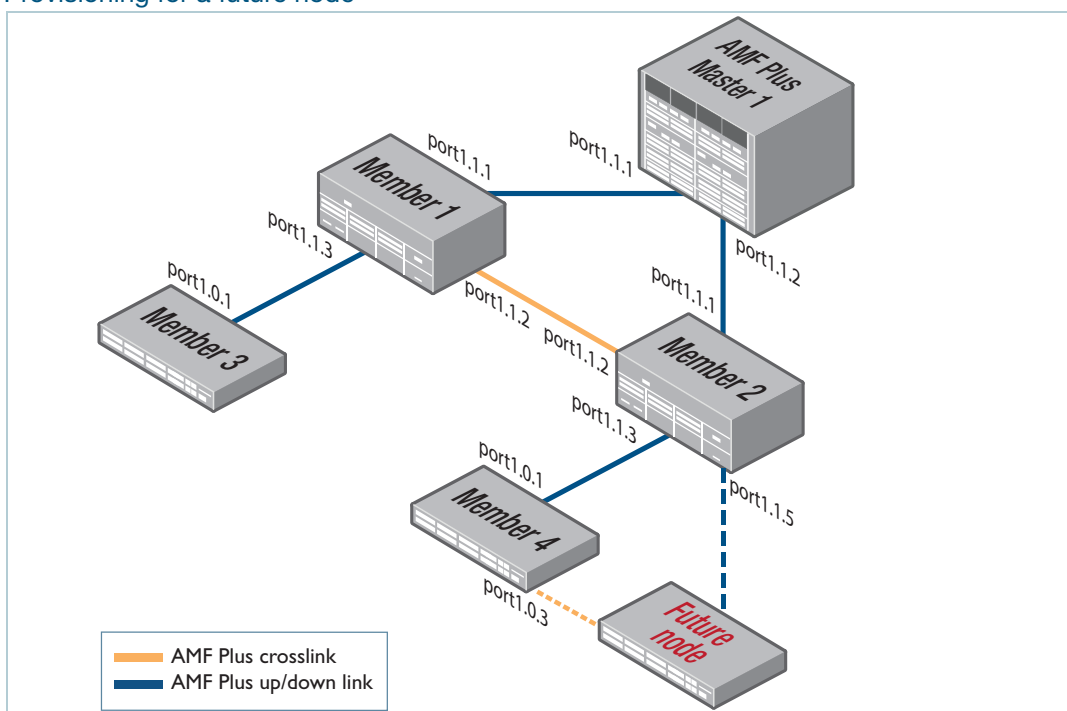
AMF Plus commands are used to create and configure a provisioned node and to specify the port(s) that the node is expected to appear on.

When to use node provisioning

Node provisioning can be used in the following instances:

- For future extension of the AMF Plus network. You can preconfigure future AMF Plus network nodes via the **atmf provision node** commands. The following figure illustrates the position of a future, provisioned node. Port 1.1.5 on Member 2 and port 1.0.3 on Member 4 are configured to expect the future node.

Figure 2: Provisioning for a future node



- For replacing an existing node with a node that has a **different platform** (e.g. replacing an x510 switch with an x530), and/or with a **different host name**: using the **atmf provision node** commands, you can configure the ports on adjacent nodes to accept a replacement member.

If you are replacing an existing node that has the same platform and host name, refer to "[Node Recovery](#)" on page 144. In this case, node provisioning is not necessary. Use node recovery instead.

Provisioning multiple device-types on the same node

From AlliedWare Plus version 5.4.9-0.1 onwards you can provision a node with multiple device-type backups. When a device is then attached to the network, AMF Plus uses its device-type to find the correct configuration to use. For example, if you create an x510 and an x530 provisioning configuration for a node called "future-node", then if either an x510 or an x530 is attached to that node, the appropriate configuration will be used.

Creating a new provisioned node

Node provisioning is effectively the process of creating a backup file-set on a master node that can be loaded onto a provisioned node some time in the future. This file-set is created just as if the provisioned node really existed and was connected to the network. Typically these comprise configuration, operating system, and license files etc.

AMF Plus stores these configuration files for the provisioned node on the master node's backup media or a remote backup server. These files are automatically loaded onto the new node's Flash when it is introduced to the network, in the same way as backed up files are loaded to the replacement for an existing node.

You can preconfigure nodes either by **creating** a new directory, or by **cloning** an existing node (see [Table 3 on page 141](#) and [Table 2 on page 140](#)).

In brief, the operations of the two methods are:

1. Using the command **create**.

This command creates an "empty" directory ready to hold release and configuration files for use on a future node. You then need to copy configuration and release files from an existing device into the new directory. After this use the **configure** commands to set the boot release and configuration files.

A convenient way to do this is to use the commands:

```
awplus# atmf provision node <nodename> [device <device-type>]
awplus(atmf-provision)# copy flash:<release-file> ./<release-file>
awplus(atmf-provision)# copy flash:<config-file> ./<config-file>
awplus(atmf-provision)# configure boot system <release-file>
awplus(atmf-provision)# configure boot config <config-file>
```

where *<nodename>* is the hostname used for the provisioned node and *<device-type>* is an optional parameter for specifying which model device the configuration is for.

Alternatively, you can create the configuration files by using the text editor to edit a configuration script:

- Into the file, enter commands similar to those described in "[Example AMF Plus Configuration](#)" on page 40.
- Copy the newly created configuration file into the directory that has been created for holding files for this future node. This procedure is described in [Table 3 on page 141](#).

2. Using the command **atmf provision node clone**.

This command creates a new directory and copies most settings and files from another backed up or provisioned node, referred to as the 'donor node'. You can make additional changes manually to these files, if needed.

We recommend that you select the donor node to have a configuration as close as possible to that needed on the new node, and for it to contain the same number of ports, or have the same expansion modules (XEMs or LIF cards) installed in the same bays. This limits the number of manual changes required to replicate the configuration of the new node.

It is convenient to set the working directory to be the directory, or backup media, in which those files reside when editing or creating files for provisioning. This saves you from having to type out the full path to the nodes backup location. The following command sets the working directory to be the storage directory for a given provisioned node:

```
awplus# atmf provision node <nodename> [device <device-type>]
awplus(atmf-provision)# locate
```

where *<nodename>* is the hostname used for the provisioned node and *<device-type>* is an optional parameter for specifying which model device the configuration is for.

Configuring adjacent nodes

You need to configure the AMF Plus links and cross-links on the adjacent node before the new node is connected. Later, when the provisioned node is introduced to the AMF Plus network, the adjacent node(s) will recognize it and the new node will automatically join the AMF Plus network.

If you plan to **replace** an existing AMF Plus node with one that has a **different host name**, use the **atmf provision** command to configure the adjacent node to expect the new node in the future. This command is used to configure all AMF Plus links and cross-links to the new node (excluding virtual-links). The command is entered in port configuration mode for the port to which the provisioned node will be connected. It effectively informs the node that a provisioned node, with a specified name, will be connected to that port.

If you plan to **extend** your AMF Plus network via ports that have not been used before, you must first fully configure the ports beforehand. Such configuration includes using the **atmf provision** command and other commands, some of which are shown in the following tables.

Carry out the procedures outlined in [Table 2 on page 140](#) if you want to achieve the following situations:

- **clone** a provisioned node.
- configure the existing node(s) that the provisioned node will eventually connect to.

Carry out the procedures outlined in [Table 3 on page 141](#) if you want to achieve the following situations:

- **create** a provisioned node.
- configure the existing node(s) that the provisioned node will (eventually) connect to.

Table 2: Procedure for cloning a provisioned node and configuring its adjacent nodes

1. Enter Privileged Exec mode	AMF_Master1>enable
2. Enter AMF Plus provisioning mode and set the name to “future_node” and optionally specify the configuration is for an x530.	AMF_Master1#atmf provision node future_node OR AMF_Master1#atmf provision node future_node device x530
3. This command clones the settings from member_3 to future_node	AMF_Master1(atmf-provision)#clone member_3 If further changes are required, edit the configuration, as explained in step 6 in Table 3 above. Note that it is essential to give the provisioned node a unique hostname.
4. On adjacent node(s), configure the port(s) that will be connected to the provisioned node. In this example, port1.0.3 on member_4 is being configured as an AMF Plus link and to expect the provisioned node future_node	AMF_Master1#atmf working-set member_4 member_4#configure terminal member_4(config)#interface port1.0.3 member_4(config-if)#switchport atmf-link member_4(config-if)#switchport trunk native vlan none member_4(config-if)#atmf provision future_node member_4(config-if)#exit member_4(config)#exit member_4#atmf working-set group local AMF_Master1# Note that AMF Plus links and cross-links do not need to be configured with data VLANs and can be used solely to provide AMF Plus management VLAN redundancy. Step 4 can be repeated to configure the ports on other adjacent nodes to expect the provisioned node.

Table 3: Procedure for creating a provisioned node and configuring its adjacent node(s)

1. Enter Privileged Exec mode	Member_4>enable
2. Enter AMF Plus provisioning mode and set the name to “future_node” and optionally specify the configuration is for an x530.	Member_4#atmf provision node future_node OR Member_4#atmf provision node future_node device x530
3. This command sets up an empty directory on the backup media for use with the provisioned node.	Member_4 (atmf-provision) #create
4. Copy and set release file	<p>To copy a release file from Member4’s Flash into the future_node directory, and set that release file to load onto future_node when it first boots up, enter the following commands:</p> <pre>Member_4 (atmf-provision) #locate Member_4 (atmf-provision) #copy flash:member4.rel ./ future_node.rel Member_4 (atmf-provision) #configure boot system future_node.rel OR Member_4 (atmf-provision) #locate Member_4 (atmf-provision) #copy current-software member4.rel ./future_node.rel Member_4 (atmf-provision) #configure boot system future_node.rel</pre> <p>For information on downloading AlliedWare Plus release files see the Download Centre at www.alliedtelesis.com/support/software.</p>
5. Copy and set the configuration file	<p>To copy a configuration file named current.cfg from Member4’s Flash into the future_node directory, and set that configuration file to load onto future_node when it first boots up, enter the following commands:</p> <pre>Member_4 (atmf-provision) #locate Member_4 (atmf-provision) #copy flash:current.cfg ./ future_node.cfg Member_4 (atmf-provision) #configure boot config future_node.cfg</pre> <p>For information on configuring a switch for AMF Plus see "Example AMF Plus Configuration" on page 40.</p>
6. Edit configuration file if necessary	<p>Note that it is essential to give the provisioned node a unique hostname. For information on configuring a switch for AMF Plus see "Example AMF Plus Configuration" on page 40.</p>

Table 3: Procedure for creating a provisioned node and configuring its adjacent node(s) (continued)

<p>7. Copy and set license file</p>	<p>To copy a license certificate named member_4.txt from member4's Flash into the future_node directory, and set that license certificate to load onto future_node when it first boots up, enter the following commands:</p> <pre>Member_4(atmf-provision)#locate Member_4(atmf-provision)#copy flash:member_4.txt ./ future_node.txt Member_4(atmf-provision)#license-cert future_node.txt</pre>
<p>8. On adjacent node(s), configure the port(s) that will be connected to the provisioned node. In this example, port1.0.3 on member4 is being configured as an AMF Plus link and to expect the provisioned node future_node</p>	<pre>Member_4#configure terminal Member_4(config)#interface port1.0.3 Member_4(config-if)#switchport atmf-link Member_4(config-if)#switchport trunk native vlan none Member_4(config-if)#atmf provision future_node Member_4(config-if)#exit Member_4(config)#exit Member_4#atmf working-set group local</pre> <p>Note that AMF Plus links and cross-links do not need to be configured with data VLANs and can be used solely to provide AMF Plus management VLAN redundancy.</p> <p>Step 8 can be repeated to configure the ports on other adjacent nodes that will be connected to the provisioned node.</p>

Connecting a provisioned node to an AMF Plus network

When you add the new node to the AMF Plus network, its settings and files are automatically downloaded from the master node's backup media, or a remote backup server, to the new node's Flash.

All you need to do is cable the new clean device in to the network. The switch's port LEDs will flash to show that its settings are being loaded. Progressive strobing of all the port LEDs indicates that a recovery is underway. For more information on the node recovery LEDs see "[Recovery progress indication](#)" on page 149.

The following example shows the expected output when a provisioned node named **future_node** joins the AMF Plus network to replace a node called **member_5**.

```
21:57:35 awplus ATMF[999]: ATMF network detected
21:57:35 awplus ATMF[999]: ATMF safe config applied (forwarding disabled)
21:57:45 awplus ATMF[999]: Shutting down all non ATMF ports.
21:57:45 awplus ATMF[999]: member_5 has left. 0 member in total.
21:57:45 x530-2 ATMF[999]: future_node has joined. 1 member in total.
21:57:45 x530-2 ATMF[999]: Automatic node recovery started
21:57:45 x530-2 ATMF[999]: Attempting to recover as future_node
21:57:46 x530-2 ATMF[999]: Checking master node availability
21:57:52 x530-2 ATMF[999]: AMF_Master1 has joined. 2 members in total
21:57:54 x530-2 ATMF[999]: member_1 has joined. 3 members in total.
21:57:56 x530-2 ATMF[999]: member_2 has joined. 4 members in total.
21:58:00 x530-2 ATMF[999]: member_3 has joined. 5 members in total.
21:58:03 x530-2 ATMF[999]: member_4 has joined. 6 members in total.
21:58:04 x530-2 ATMFFSR[6779]: Retrieving recovery data from master node
AMF_Master1
21:58:34 x530-2 ATMFFSR[6779]: Licence installed from certificate.
21:58:35 x530-2 ATMFFSR[6779]: File recovery from master node succeeded.
Node will now reboot
```

Node Recovery

Automatic node recovery

AMF Plus allows you to replace a failed node with another device and let AMF Plus automatically load the appropriate configuration, operating system, licenses, and other files onto the replacement device.

For this to work, the replacement device must have no configuration file. This means it must be either:

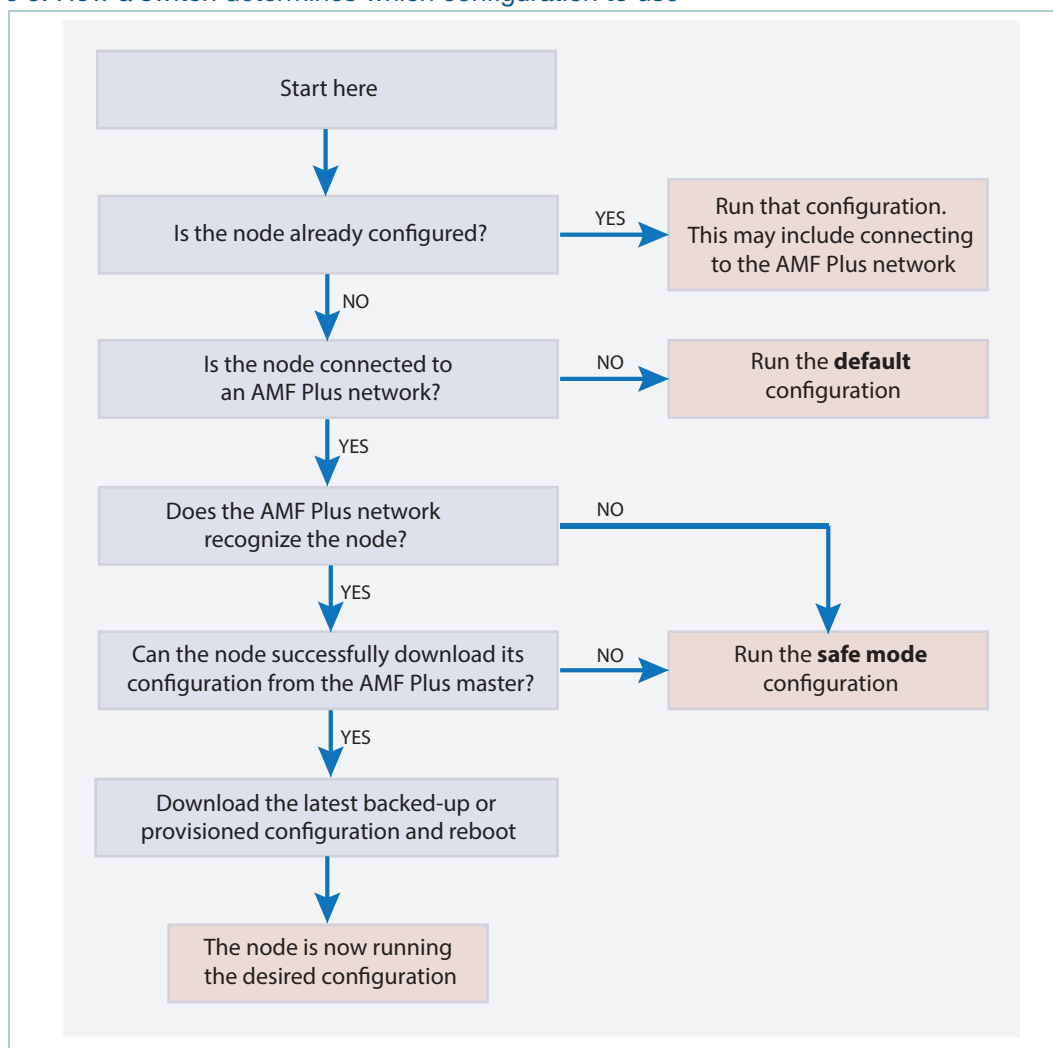
- a factory-new device, or
- a used device that has been returned to a “clean” state (see ["Restoring a node to a “clean” state" on page 148](#))

To replace a failed device with a new device that has either a different platform or a different node name you need to provision the network to expect the new device. See ["Node Provisioning" on page 137](#).

Note: It is possible to replace some AlliedWare Plus devices with an equivalent model and still make use of automatic recovery (see ["Replacing a device with a similar model" on page 147](#)).

When a switch boots up, it follows the process shown in the flowchart of [Figure 3](#) to determine what configuration to use. This flowchart indicates when automatic node recovery is successful.

Figure 3: How a switch determines which configuration to use



How does the recovering node work out which files to download?

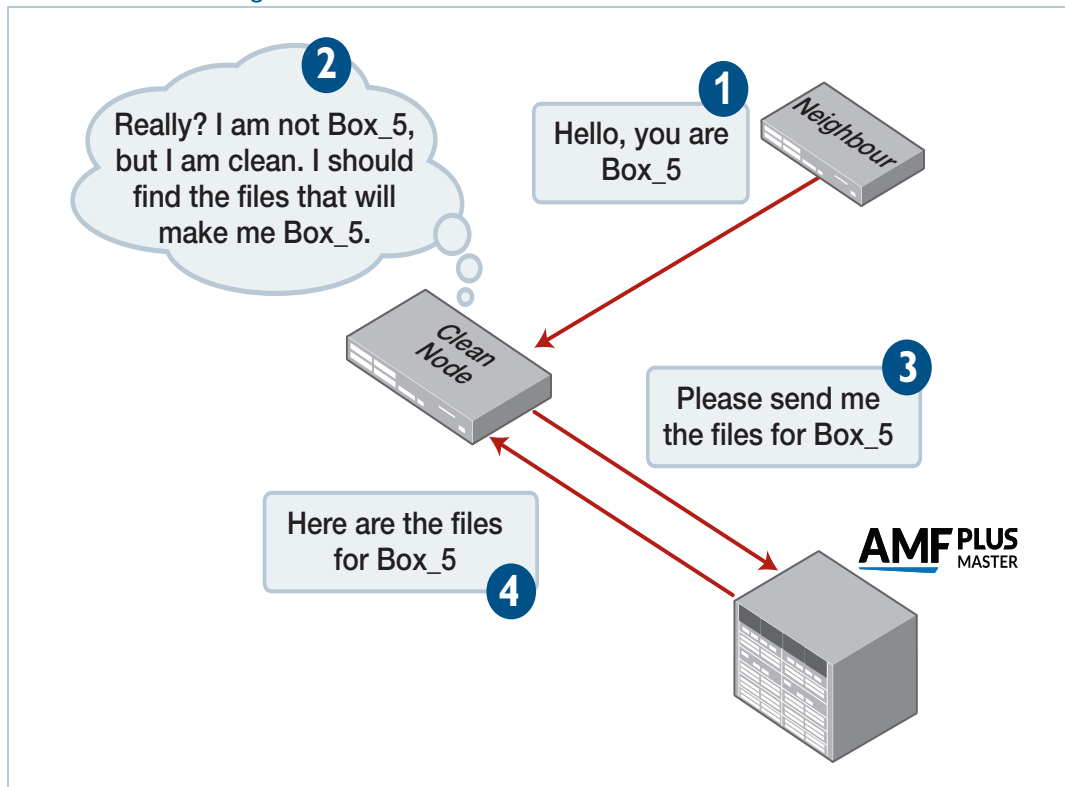
A key step in the flowchart above is “**Does the AMF Plus network recognize the node?**” This is the step where the power of AMF Plus comes into play.

When a node disappears and a new, clean, node is added in its place, the following algorithm is used:

1. The neighbors will send AMF Plus messages with the host name of the previously operating node.
2. The new node will query all known masters to find which has the most recent back-up for that host name given by the adjacent neighbors.
3. The new node will restore all the relevant files that relate to the hostname in question, from the master it has chosen.

A simplified representation of the process is as follows:

Figure 4: How a recovering node works out which files to download



Replacing a device with a similar model

Automatic recovery works on AlliedWare Plus devices replaced with an equivalent device as shown in the table below:

Table 4: Permitted device substitutions

Original device	Replacement device	Supported from version ^a
x310 Series	x530 Series and x530L Series	5.5.2-0.1
FS980M Series	GS970EMX Series and GS980MX Series	5.5.2-0.1
IE200 Series	IE340 Series and IE340L Series	5.5.2-0.1
IE300 Series	IE340 Series and IE340L Series	5.5.2-0.1
IE510-28GSX	x530 Series and x530L Series	5.5.2-0.1
GS900MX/MPX Series	GS980M/52PS or GS980M/52	5.5.1-1.1 ^b
GS900MX/MPX Series	GS980MX Series	5.4.9-2.1
IX5-28GPX	x530 Series	5.4.9-0.1
x210 Series	x230 Series	5.4.9-0.1
x510 Series	x530 Series	5.4.9-0.1
x610 Series	x530 Series	5.4.9-0.1
x900 Series	x930 Series or x950 Series	5.4.9-0.1
x930 Series	x950 Series	5.4.9-0.1

- The supported version (or later) needs to be running on the node's neighbor.
- Also supported in the most recent maintenance versions for 5.4.9-2.x, 5.5.0-2.x and 5.5.1-0.x

Boot release file

When a device is replaced with an equivalent model, the replacement device receives the configuration files from the node's backup but retains its own boot release file. For example, if you replace an x510 series device with an x530 series device, then the x510's configuration files will be used, but the recovering node will use the x530's release file. This throws an error in the log file that is expected and can be ignored.

Figure 5: Example log output showing release file error

```

...
07:02:52 x510 atmffsd: /flash/x510-5.5.3-2.3.rel is not a valid release file
(Release not intended for this device)
07:02:52 x510 atmffsd: No valid boot system found
07:02:52 x510 atmffsd: Restoring original firmware to flash:x530-5.5.3-2.3.rel
...

```

As this release file is the one that was shipped with the device, we recommend that you update this release file to match the version of AlliedWare Plus you are currently running in your network.

Restoring a node to a “clean” state

When replacing a failed device, your replacement device should be one of the following types, in order for AMF Plus automatic node recovery to work:

- a factory-new device
- a used device that has been returned to a “clean” state

A clean device is one that has had its previous configuration components removed. A process of cleaning is required when replacing a failed device with one that, although in working condition, has been used previously and still retains components of its previous configuration.

If you keep on-site spares, store them with clean configurations and current releases. When you upgrade your network to a new AlliedWare Plus version, we recommend that you upgrade your spare devices too.

To clean up a previously used device, use the **atmf cleanup** command.

Caution

If the **atmf cleanup** command is run on a node that is connected to an AMF Plus network, and the AMF Plus master for that network is currently running a backup, there is a small risk that the files for the node being cleaned will be erased from the AMF Plus master backup. For this reason it is recommended that you disconnect the node from the AMF Plus network before the **atmf cleanup** command is executed.

This command erases all data from NVS and Flash, apart from the following:

- the boot release file (a .rel file) and its release setting file
- license files
- the latest GUI release file

Any other user files that remain in Flash will be overwritten during the automatic recovery process. If there are any files stored in the Flash of the replacement device that need to be retained, back these files up prior to installing the device into the AMF Plus network.

The device is then rebooted to put it into a clean state. The device can then be used for automatic node recovery.

Recommended procedure when replacing a device using automatic node recovery

This section describes the basic procedure for node recovery. If your node is one of the following, see later sections for important details:

- ["Recovering a VCStack — SBx908 GEN2" on page 155](#)
- ["Recovering AMF Plus devices with special links" on page 158](#)
- ["Recovering and Provisioning Isolated Nodes" on page 160](#)
- ["Recovering a different access point \(not TQ5403 series\)" on page 168](#)

Basic node recovery

1. Ensure that the replacement device is in a 'clean' state. See ["Restoring a node to a "clean" state" on page 148](#)
2. Power down the device to be replaced.
3. With the replacement device powered down, move the port connections from the broken device to the replacement device, being careful to ensure the ports are connected in exactly the same way as they were connected to the broken device.
4. Power up the replacement device.

Following these recommendations will ensure AMF Plus can successfully auto-recover the broken device.

LAGs If the device to be replaced was connected to the AMF Plus network using aggregators, it is important to follow the procedure above. Failure to do this may result in the AMF Plus network failing to auto-recover the device. Failure to auto-recover will leave the device in safe mode. If this happens, the device may be successfully auto-recovered by ensuring all of the connections are in place then powering the replacement device down, then back up. This will re-start the auto-recovery process.

Recovery progress indication

There are two ways to tell recovery progress: from the device's LEDs and from log messages.

LED progress indicator

This is a visual feature that uses front-panel LEDs to display the recovery status during automatic recovery. This feature uses two distinct flash patterns to indicate the following states:

RECOVERY STATE	LED INDICATION (GREEN)
Recovery in progress	Progressive strobing of all port LEDs.
Recovery failure	All port LEDs simultaneously flashing on and off.

If the recovery fails, the LEDs will keep flashing until you turn off the failure-alert indicator. To do this, use the command **atmf recover led-off**. This command will return the port LEDs to their normal running state.

If an automatic recovery fails, you need to determine the cause of the failure, and take appropriate action.

Note that the **findme**, **findme trigger**, and **ecofriendly** LED features cannot be used while AMF Plus recovery progress indication is active.

Note: This feature is not available on the x8100 series switches.

Log message

A log message will appear on the console or other VTY session indicating when recovery has finished (whether successfully or with errors). This message can also be found by viewing the log with the **show log** command

Example log output showing automatic node recovery

```
23:03:15 awplus ATMF[863]: ATMF network detected
23:03:15 awplus ATMF[863]: ATMF safe config applied (forwarding
disabled)
23:03:25 awplus ATMF[863]: Shutting down all non ATMF ports
23:03:26 x530_1 ATMF[863]: Automatic node recovery started
23:03:26 x530_1 ATMF[863]: Attempting to recover as x530_1
23:03:26 x530_1 ATMF[863]: Checking master node availability
23:03:32 x530_1 ATMF[863]: Master has joined. 2 members in total.
23:03:32 x530_1 ATMF[863]: x950_VCS_2 has joined. 3 members in total.
23:03:32 x530_1 ATMF[863]: x950_VCS_1 has joined. 4 members in total.
23:03:37 x530_1 ATMFFSR[2950]: Retrieving recovery data from master
node Master
23:05:18 x530_1 ATMFFSR[2950]: File recovery from master node
succeeded. Node will now reboot
Flushing file system buffers...
Unmounting any remaining filesystems...
Restarting system.
```

Other points to note about node recovery

- Automatic node recovery is not intended to restore multiple nodes simultaneously. If multiple nodes have failed, you must recover them one at a time.
- Do not make any changes to the device's configuration while a node recovery is underway.

Recovering devices that have subscription licenses

Subscription licenses are keyed to the serial number of a device. If you replace a device, then AMF Plus recovery automatically transfers these licenses to the replacement device.

After node recovery, any licenses from the failed device will be stored and will operate on the new device for a 28-day grace period. During this grace period, the new device will generate a log

message once per day. This log message tells you how many days are left until the license will stop working and advises you to contact the Allied Telesis support team. The support team can transfer the license from the old device's serial number to the new device's serial number. The output of the **show license external** command displays the same message.

Also, if there were any expired licenses on the failed device, they will be stored for the 28-day grace period too. You can use the **show license external stored** command to see any expired licenses. The following output shows an example of an expired license.

Example of **show license external** and **show license external stored** for an expired license

```
awplus#show license external

NOTICE: This device has undergone ATMF recovery and is currently using
        licenses registered to serial number A05049G151700005.
        The grace period ends at 00:00:00 18 Jan 2023. To ensure
        continued operations, please contact Allied Telesis Customer
        support to have the license entitlements transferred to this
        device using serial A05050G144800006

Features with installed entitlements:

No active entitlements found.

awplus#show license external stored

Feature entitlements stored on this unit (A05050G144800006):

AMF Master

        Start date:                14 Jan 2023 00:00
        Expiry date:                15 Jan 2024 23:59
        AMF nodes:                  10
```

Contact your authorized Allied Telesis representative to transfer your licenses from the old to the new serial number. If you don't do this before the end of the grace period, AlliedWare Plus disables the subscription licenses, and you will no longer have access to the subscription features.

Recovering devices that have release licenses

SwitchBlade x908 GEN2 and SwitchBlade x8100 switches need to have release licenses installed. These release licenses are keyed to the MAC address of the device. If an existing SwitchBlade is removed from the network, and replaced by a new one, the new unit will not have the same MAC address as the unit that was removed. As a result, the release license that was backed up from the original unit will not apply to the new unit.

The procedure to deal with this situation is:

1. Even though the new unit will not have a valid release license initially, the lack of relevant release licenses will not block a firmware upgrade via AMF Plus. The software will be successfully

installed on the new unit, and it will run in unlicensed mode (which means that it will not be able to upgrade to yet further software versions).

2. Following the successful recovery of a failed device and once the release license for the replacement device has been obtained, you can install it on the new device using the **license** command. To obtain release licenses, contact Allied Telesis support.

Note that backed up feature licenses (unless they are tied to a physical entity such as a device serial number or MAC address) will automatically be installed into the replacement unit - see "[Recovering devices that have subscription licenses](#)" on page 150.

Fixing a failed node recovery: AMF Plus safe configuration

If AMF Plus automatic node recovery fails, AMF Plus contains a safety net feature that puts the replacement node into a safe configuration state. This is to prevent an unconfigured device from joining the network and creating loops.

Detecting AMF Plus safe configuration operation

A log message is generated whenever AMF Plus safe configuration is applied. This message will appear in the log some time after the startup sequence. This message will also be displayed on the console or any connected VTY session.

AMF Plus safe configuration procedures

The procedures for AMF Plus safe configuration are shown below:

- A special VLAN is created in the disabled state and given the name **atmf_node_recovery_safe_vlan**. The VID of this VLAN is determined dynamically to ensure that it does not conflict with either of the AMF Plus management VLANs or any other VLANs that are detected on the AMF Plus network.
- All ports are removed from their default VLAN membership (VLAN 1).
- All ports are set as tagged members of the safe VLAN.
- Additionally, all ports that are not AMF Plus links or cross-links are shut down. These links and cross-links are detected by AMF Plus and added to the dynamic configuration. This is done to ensure correct behavior of static aggregators and Layer 3 protocols configured on neighboring devices.

Output 5: **show vlan brief** command output - for a device in AMF Plus safe configuration mode

```
awplus#show vlan brief


VLAN ID  Name          Type      State  Member ports  (u)-Untagged, (t)-Tagged
=====  =====
1         default       STATIC    ACTIVE
4090     atmf_node_recovery_safe_vlan
                STATIC    SUSPEND  port1.1.1(t)  port1.1.2(t)  port1.1.3(t)
                port1.1.4(t)  port1.1.5(t)  port1.1.6(t)
                port1.1.7(t)  port1.1.8(t)  port1.1.9(t)
                port1.1.10(t) port1.1.11(t)
                port1.1.12(t) port1.1.13(t)
                port1.1.14(t) port1.1.15(t)
                port1.1.16(t) port1.1.17(t)
                port1.1.18(t) port1.1.19(t)
                port1.1.20(t) port1.1.21(t)
                port1.1.22(t) port1.1.23(t)
                port1.1.24(t)
```

Output 6: **show running-config** output for a device in AMF Plus safe configuration mode

```
awplus#show running-config
...
!
vlan database
  vlan 4090 name atmf_node_recovery_safe_vlan
  vlan 4090 state disable
!
interface port1.0.1-1.0.4
  shutdown
  switchport
  switchport mode trunk
  switchport trunk allowed vlan add 4090
  switchport trunk native vlan none
!
interface port1.0.5
  switchport
  switchport atmf-link
  switchport mode trunk
  switchport trunk allowed vlan add 4090
  switchport trunk native vlan none
!
interface port1.0.6-1.0.24
  shutdown
  switchport
  switchport mode trunk
  switchport trunk allowed vlan add 4090
  switchport trunk native vlan none
!
...
```

Undoing an AMF Plus safe configuration

If your node has had AMF Plus safe configuration applied, you can use normal CLI configuration commands to modify the running-configuration to whatever configuration is required.

Caution  No changes should be made to the device's configuration while a node recovery is underway. A log message will appear on the console or other logged in session indicating when recovery has finished (whether successfully or with errors). This message can also be found by viewing the log after running the show log command.

The example below shows a device being returned from AMF Plus safe configuration mode to having its default VLAN and port settings applied. Note that in this example a 24-port switch has been used.

```
awplus# configure terminal
awplus(config)# interface port1.0.1-port1.0.24
awplus(config-if)# switchport trunk native vlan 1
awplus(config-if)# switchport trunk allowed vlan remove 4090
awplus(config-if)# switchport mode access
% port1.0.5 has ATMF link configured so its mode cannot be changed
awplus(config-if)# no shutdown
awplus(config-if)# exit
awplus(config-if)# vlan database
awplus(config-if)# no vlan 4090
awplus(config-if)# end
```

In order to retain connectivity to the AMF Plus network, AMF Plus link and cross-link settings should not be changed. In the example above you can see that port1.0.5 is an automatically configured link.

Recovering a node manually

There are certain situations where automatic recovery may fail. Automatic recovery has been designed to be cautious in its approach to recovering nodes for reasons such as:

- the backup stored on the AMF Plus master does not have a “good” status
- the replacement device has a version of the AlliedWare Plus operating system installed on it that too old to be compatible with Alliedware Plus version running on the neighbor or the master.

When these situations occur, automatic node recovery will fail.

In this failed state, the replacement device will have the AMF Plus safe configuration mode applied (see ["AMF Plus safe configuration procedures" on page 152](#)). After investigating the failure and taking remedial action you can initiate a manual node recovery. To do this, enter the following command:

```
amf1# atmf recover {<node_name>} {<master_node_name>}
```

where:

- **node_name** is the host name of the device being recovered.
- **master_node_name** is the host name of the AMF Plus master that contains the backup you want to use for the recovery.

The manual recovery command will bypass the usual checks performed by automatic node recovery. Make sure that the backup configuration stored on the specified AMF Plus master is correct before you execute the command.

If you attempt to manually recover a node with the backup file of a node from a **different platform**, the release file from the backup will be incompatible and won't be copied to the replacement device. Instead, the existing release on the replacement device will be used, in order to ensure the device can join the AMF Plus network and function correctly.

Output 7: Example output showing manual recovery

```
amf1#atmf recover x530_1 Master
This command will erase ALL flash contents. Continue node recovery?
(y/n)y
Manual node recovery successfully initiated
x530_1#23:15:32 x530_1 ATMFFSR[8477]: Retrieving recovery data from
master node Master
23:17:17 x530_1 ATMFFSR[8477]: Manual node recovery completed
x530_1#
```

Recovering a VCStack — SBx908 GEN2

This procedure describes the recovery of a stacked device that is also the AMF Plus Master. It uses a combination of the AMF Plus node recovery and VCStack recovery mechanisms.

In summary, you need to:

1. Prepare the replacement device.
2. Use AMF Plus recovery to recover the replacement device as the AMF Plus Master as a standalone device (not connected to any other stack members).
3. Use VCStack to recover the other stack members.

This example describes a two-device stack. For stacks with more devices, repeat step 3 as necessary.

This procedure assumes the failed AMF Plus Master had an external media device connected to it, such as a USB stick, and that this external media device contains a recent backup of the AMF Plus Master or the whole AMF Plus network. You may also need the following information:

- The hostname of the device to be replaced
- The AMF Plus network name

- The VLAN and subnet used for stacking, if these have been changed from their default settings.

Step 1: Prepare the replacement devices

1. Connect the stacking cables to the device which will be the Stack Master (which we will call SM1) and connect a console cable to this device. Make sure no other cables are connected to this device and that the other ends of the stacking cables are not connected.
2. Connect a console cable to the other stack member (which we will call sm2).
3. Connect the external media device with the AMF Plus backup to SM1.
4. Run the command **atmf cleanup** on both SM1 and sm2.

Step 2: Use AMF Plus to recover the replacement device as a standalone unit

This will happen automatically if the replacement device finds a recovery file on the external media. This will be the case if an AMF Plus virtual-link had been connected to the device. Otherwise, do the following:

5. On SM1, configure the AMF Plus network name, and hostname of 'atmf master'.
6. If the failed device used a non-standard VLAN and subnet for stacking, configure these.
7. Run the command **atmf recover** on SM1.
8. Once the recovery has succeeded, reboot SM1.

Step 3: Use VCStack to recover the other stack member

9. On sm2, configure stacking and an appropriate stack ID, for example using the commands:

```
awplus(config)# stack enable
awplus(config)# stack 1 renumber 2
```

10. Save the configuration on sm2 then reboot.
11. Configure stacking ports on sm2, for example using the commands:

```
awplus(config)# interface port2.7.1-2.7.5
awplus(config-if)# stackport
```

12. Save the config on sm2 then reboot.
13. While rebooting, connect the stacking cables from SM1 to sm2. Then sm2 will join as a stack member.
14. Connect all of the other network cables.

Recovering a VCStack — other switches

Node recovery on VCStacks that are part of an AMF Plus network is slightly different to node recovery of standalone devices.

This is because VCStack has its own node recovery mechanism that has different requirements to AMF Plus.

In the extremely unlikely situation of needing to replace an entire VCStack within an AMF Plus network, you can use AMF Plus automatic node recovery to first recover Stack ID 1, which will become the VCStack master.

The replacement device for this must be a clean unit, (see ["Restoring a node to a “clean” state" on page 148](#)).

We recommend you configure at least one AMF Plus link per VCStack on an aggregator that contains ports across all stack members. This makes sure the stack stays in the AMF Plus network or area if a stack member fails.

The procedure for recovering an entire stack is as follows:

1. Connect a clean device to the AMF Plus network, and power it on. The connections into the AMF Plus network should be between the appropriately configured AMF Plus links on the neighboring node, and the ports configured as AMF Plus links in the backup of the failed node's configuration, (i.e. the ports configured as AMF Plus links on the failed node).
2. The AMF Plus network should detect the replacement device and begin the automatic node recovery process. Wait until this process completes, then check that the replacement device has come up correctly as VCStack ID 1, and that the configuration is correct.
3. Configure the next replacement device as VCStack ID 2. Ensure it is installed with a compatible release and the same set of licenses that exist on ID 1. Connect the VCStack cables and power it on. For details of how to add devices to a stack, see the [VCStack Feature Overview and Configuration Guide](#).
4. VCStack ID 1 should detect ID 2 and synchronize the configuration and firmware release. Once this has completed, check that the VCStack has formed correctly, and then connect the remaining network connections.
5. For any additional VCStack members, repeat the last two steps, ensuring that the VCStack ID is set to the next sequential value for each additional device that is added to the VCStack.

Recovering AMF Plus devices with special links

AMF Plus recovery files are created for nodes with special links. Special links include:

- virtual links,
- area links terminating on an AMF Plus master, and
- area virtual links terminating on an AMF Plus master.

An AMF Plus node with one of these special links pushes its startup configuration to its neighbors and to any attached external media. It then fetches and applies this configuration at recovery time. This configuration enables it to contact the AMF Plus master and initiate a recovery.

Recovery files can be out of date if:

- a node's neighbor is off line when changes are made to its configuration, or
- a node no longer contains a special link.

Use the **show atmf recovery-file** command to see the status of a node's recovery files and the dates they were created.

Output 8: show atmf recovery-file output

```

node1#show atmf recovery-file

AlliedWare Plus (TM) 5.5.2-2 12/17/22 19:43:41

=====
node1, node2:
=====

Working set join

=====
node1:
=====

ATMF Recovery File Info: Special Link Present
Location          Date           Time
USB storage device Media Not Found
node1             18 Feb 2023   19:21:19
node2             18 Feb 2023   19:21:19

=====
node2:
=====

ATMF Recovery File Info: Special Link Present
Location          Date           Time
USB storage device 20 Dec 2022   18:59:06
node1             18 Feb 2023   19:20:55
node2             18 Feb 2023   19:20:55

```

Use the **clear atmf recovery-file** command to delete all of a node's recovery files. It deletes the recovery files stored on:

- the local node,
- neighbor nodes, and
- external media (USB or SD card).

If AlliedWare Plus detects that a node contains a special link then the following message is displayed:

Output 9: **clear atmf recovery-file** output

```
node1#clear atmf recovery-file
% Warning: ATMF recovery files have been removed.
ATMF recovery may fail. Please save running-configuration.
```

Note: This command deletes all of a node's recovery files. If a node still has a special link you must save the node's running configuration after running the clear command. Saving the running configuration creates new recovery files on the node's neighbors and on any attached external media.

Recovering and Provisioning Isolated Nodes

An isolated node is an AMF Plus member that is only connected to the rest of the AMF Plus network via a virtual-link. As there is no physical connection to a neighboring AMF Plus device, isolated nodes cannot identify their location in the AMF Plus network via traditional means. Instead, it is identified using an **identity token** stored on the AMF Plus master. This token is saved to the AMF Plus master when the node is backed up and is created using the MAC address of the next-hop on the isolated node's virtual-link interface.

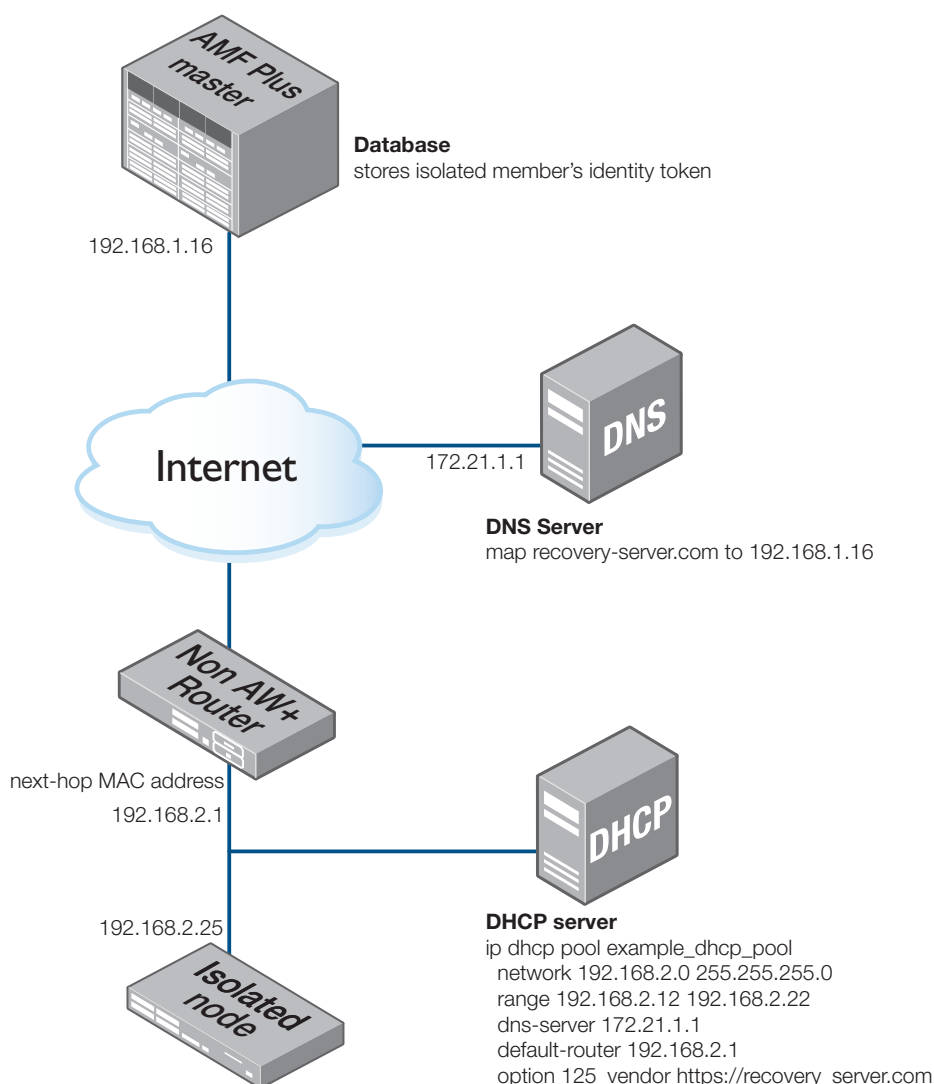
Note: When provisioning a new device, it is possible to optionally specify the new device's serial number, instead of the next-hop MAC address, as the identity token.

How auto-recover works for isolated nodes

In order to initiate a recovery the AMF Plus master must be accessible to the isolated AMF Plus node. This is achieved by using DHCP to send the Uniform Resource Identifier (URI) of the master to the recovering node. The URI is either the IP address or DNS entry of the master when running as a recovery server. This URI must be both resolvable and reachable from the recovering node.

- After a member is backed up the identity token of the recovering node is stored in a database on the master
- When the member is replaced with a clean unit, this recovering node will send a DHCP discovery request with special option 125.
- The DHCP server may respond to this request with a DHCP offer, which includes special option 125 containing the URI of the master.
- If the DHCP server does not send a URI, then the default address (<https://amfrecovery.alliedtelesis.com>) is used.
- The recovering node accesses the URI supplied by the DHCP server using its **identity token** as a parameter.
- The master will use the identity token to identify the recovering node and send the device its configuration, which is then applied to the running configuration.
- Using this configuration the recovering node will establish a virtual-link to the master and auto-recovery will begin.

Figure 6: Isolated node auto-recovery example



Note: All AMF Plus traffic running over a public network should be encrypted by configuring a VPN between the AMF Plus network and any remote sites.

Preparing your network

Before configuring auto-recovery or provisioning of an isolated node ensure the following:

- There is a route from the remote network to the AMF Plus master that does not depend on the recovering device.
- The master is reachable from the remote network using the information provided by the DHCP server.
- If the DHCP server is not configured with option 125, a DNS server must be available to resolve the default address (<https://amfrecovery.alliedtelesis.com>) to the IP address of the master.
- The virtual-link interface of the isolated device should be configured with a statically assigned address that is outside the range served by the DHCP server, but within the same subnet.

- The virtual-link next-hop address must also be in the same IP subnet as that provided by the DHCP server.
- If the port between the neighbor device and the recovering node is trunked, the virtual-link must be on the **native vlan** that exists on the neighbor device. If it isn't then the DHCP server will not see the DHCP discover packets.
- Only the recovery of single isolated nodes is supported. Multiple isolated nodes, sharing the same next-hop device, create ambiguous identity tokens, unless they are connected to a device that supports separate MAC addresses per interface.

Note: The MAC address of the next-hop device is critical to the recovery of a device as it is used to identify the recovering device. For this reason, if the next-hop device is replaced it is vital the AMF Plus node is backed up once the new device has been installed.

DHCP Server

There are three alternatives for supplying the AMF Plus master's URI to the recovering node:

- Configure the DHCP server to send the master's FQDN using option 125.
- Configure the DHCP server to send the master's IP address using option 125
- Use the default URI, <https://amfrecovery.alliedtelesis.com>.

The DHCP server must be configured on a device on your network that is reachable by the recovering node.

Sample DHCP pool configurations for an AlliedWare+ device are shown below.

Output 10: Sample DHCP configuration supplying the AMF Plus master's FQDN via option 125

```
ip dhcp option 125 name 125_vendor ascii
!
ip dhcp pool example_dhcp_pool
 network 192.168.2.0 255.255.255.0
 range 192.168.2.12 192.168.2.22
 dns-server 172.21.1.1
 default-router 192.168.2.1
 option 125_vendor https://recovery_server.com
!
service dhcp-server
!
```

Note: The recovery request adds a path to the base URI (/api/v1/atmf/pre_recovery_request). It is important, therefore, not put anything after the FQDN portion of the URI. If you do the URI will be used verbatim. In this example https://recovery_server.com is valid, while https://recovery_server.com/ is not.

Output 11: Sample DHCP configuration supplying the AMF Plus master's IP address via option 125.

```
ip dhcp option 125 name 125_vendor ascii
!
ip dhcp pool example_dhcp_pool
network 192.168.2.0 255.255.255.0
range 192.168.2.12 192.168.2.22
dns-server 172.21.1.1
default-router 192.168.2.1
option 125_vendor https://192.168.1.16
!
service dhcp-server
!
```

Output 12: Sample DHCP configuration using the default URI.

```
!
ip dhcp pool example_dhcp_pool
network 192.168.2.0 255.255.255.0
range 192.168.2.12 192.168.2.22
default-router 192.168.2.1
!
service dhcp-server
!
```

DNS Server The DNS server maps the AMF Plus master's FQDN to its IP address. This will either be the custom URI specified using option 125 on the DHCP server or the default URI, <https://amfrecovery.alliedtelesis.com>.

A DNS server is not required if the DHCP server sends a URI containing the master's IP address instead of its FQDN.

Configuring AMF Plus

In order to handle recovery requests from isolated nodes, the **recovery-server** must be enabled on the AMF Plus master. Use the following command on the master to enable **recovery-server**.

```
awplus(config)# atmf recovery-server
```

Replace an existing node

Once **recovery-server** is enabled on an AMF Plus network, the next time an isolated node is backed up its **identity token** will be stored in the AMF Plus master's database. Should the device fail, it can then be replaced using the following procedure:

1. Ensure that the replacement device is in a 'clean' state, see "[Restoring a node to a "clean" state](#)" on page 148
2. Power down the device to be replaced.
3. With the replacement device powered down, move the port connections from the broken device to the replacement device, being careful to ensure the ports are connected in exactly the same way as they were connected to the broken device.
4. Power up the replacement device.

Provision a new node To deploy a new device to a remote location do the following:

Step 1: Provision a node.

See "[Node Provisioning](#)" on page 137 for general information on how to provision a node.

Note: As an isolated node has no AMF Plus neighbors there is no need to configure an adjacent node to identify it.

Step 2: Create an identity token for the device that will be provisioned

You can create this by either specifying its next-hop MAC address or by specifying the serial number of the replacement device. The advantage of using the next-hop MAC address is that any device, regardless of its serial number, can be added to the network but using the serial number maybe preferential if it is not easy to find the next-hop MAC address.

To create a identity token for a device named "my-x930" with serial number "A10064A172100008" use the following commands:

```
awplus# atmf provision node my-x930
awplus(atmf-provision)# identity serial-number A10064A172100008 prefix
192.168.2.25/24
```

To create a identity token for a device named "my-x930" with next-hop MAC address "e01a.ea2a.70e9" use the following commands:

```
awplus# atmf provision node my-x930
awplus(atmf-provision)# identity mac-address e01a.ea2a.70e9 prefix
192.168.2.25/24
```

Note: The prefix is the IP address and subnet mask of the virtual-link interface on the isolated node.

Step 3: Add a virtual-link

Create the local side of the virtual-link before powering up the provisioned node. In our example the following command needs to be run on the AMF Plus master:

```
awplus(config)# atmf virtual-link id 1 ip 192.168.1.16 remote-id 2 remote-
ip 192.168.2.25
```

Note: This command needs to be run on the device terminating the local side of the virtual-link.

See "[AMF Plus Tunneling \(Virtual-links\)](#)" on page 211 for more information on setting up virtual-links.

Step 4: Add a new device to the network

At the remote site cable in a new device that is in a 'clean' state, see "[Restoring a node to a "clean" state](#)" on page 148, and power up. If the network has been prepared correctly the new device will now contact the AMF Plus master and provision itself automatically.

Recovering a TQ5403 series access point

From AlliedWare Plus version 5.5.3-0.1 onwards, it is possible to use AMF Plus auto-recovery to recover TQ5403, TQ5403e and TQm5403 access points (APs). If an AP fails, you can simply replace it with a new or factory-reset identical model AP. AMF Plus will use the original AP's backed-up configuration file and a stored firmware file to recover the AP. This is supported from TQ5403 series firmware version 6.0.3-0.1 onwards.

Preparing for auto-recovery

Auto-recovery is available for APs that are attached to any AlliedWare Plus product that has switch ports. To set up an AlliedWare Plus device so that auto-recovery can happen for an AP attached to the device, first make sure that the following apply:

- AMF backup is enabled (it is enabled by default)
- The AMF master has suitable external media installed for saving the backups
- Backups of the AP's configuration are happening successfully
- A VLAN has been configured on the port that the AP is connected to on the AlliedWare Plus device. The VLAN must have an IP address in the same subnet as the AP's IP address
- The AP's IP address can be statically configured or assigned by DHCP. If you want to use the AlliedWare Plus device as a DHCP server for the AP, see the [DHCP Feature Overview and Configuration Guide](#) for details about configuring DHCP pools.

Then do the following steps:

Step 1: Create a guest class on the AlliedWare Plus device that the AP is connected to in the AMF Plus network

For example, create a guest class called TQ5403. In this example, the AP is connected to an x230 series switch:

```
x230(config)# atmf guest-class TQ5403
x230(config-atmf-guest)# modeltype tq
x230(config-atmf-guest)# username manager password tq54-guestpass
```

Step 2: Specify the discovery method

Set the discovery method to 'agent':

```
x230(config-atmf-guest)# discovery agent
```

Step 3: Specify login fallback on the guest class

Login fallback means that if a guest node's saved username and password fail, AMF Plus will try to connect to the node using the factory default username and password (manager/friend). Login fallback is disabled by default and is only valid for model type **tq**.

Enable login fall back on the guest-class TQ5403:

```
x230(config-atmf-guest)# login-fallback enable
```

Step 4: Store the access point's firmware file in an accessible location and specify that location in the guest class

AMF Plus cannot back up AP firmware files (only configuration files), so you need to store the firmware file somewhere accessible and provide the AlliedWare Plus device with its location. Note:

- You can specify the file's directory or its filename. Specifying a directory is recommended because that makes it easier to keep the firmware file up to date.
- The following protocols are supported: http, https, tftp, usb, and card.
- Do not change the filenames of stored firmware files.

For example, store the firmware file in the top level of a USB stick, insert the USB stick into the AlliedWare Plus device, and specify that on the x230 Series switch:

```
x230(config-atmf-guest)# firmware-url usb:
x230(config-atmf-guest)# exit
```

Step 5: Create a guest link between the AP and AlliedWare Plus device

The guest link needs to be configured on an access switchport or the native VLAN of a trunk switchport. If the AP has a static IP address, specify that address as well.

For example, configure a static guest link connected on port 1.0.2 on the x230 Series switch, for an AP with a static address of 192.168.2.1:

```
x230(config)# interface port1.0.2
x230(config-if)# switchport atmf-guestlink class TQ5403 ip 192.168.2.1
```

If the AP's address is dynamically assigned, simply remove the IP address from the guest link command:

```
x230(config-if)# switchport atmf-guestlink class TQ5403
```

Recovering a failed AP

If your TQ5403 series AP fails, auto-recovery is very easy.

Step 1: Make sure the replacement AP is in a factory-default state

If the AP is not new, use the **Maintenance > Configuration** menu in its web-based management interface to return it to the factory defaults. See the [User Guide](#) for more information.

Step 2: Make sure the replacement AP is running recent enough firmware


The replacement AP needs to be running firmware version 6.0.3-0.1 or later. If necessary, use the **Maintenance > Update** menu to upgrade it.

Step 3: Replace the failed AP with the replacement AP

Cable up the replacement AP in the same way as the failed AP, and power it on. AMF Plus will automatically recover it.

Recovering a different access point (not TQ5403 series)

When TQ series APs are configured as AMF Plus guest nodes, their configuration files can be included in the automated backups performed by AMF Plus. While zero-touch recovery is not available on AMF Plus guest nodes - except for TQ5403 series APs - you can manually recover the configurations of other TQ-series APs included in the AMF Plus backup regime. This is useful when replacing a failed TQ device. The following example shows the configuration, backup, and restore of such a TQ AP using AMF Plus.



Replacing a TQ guest node with login fallback

You can enable login fallback on TQ guest nodes using the command **login-fallback enable**.

Login fallback means that if a guest node's saved username and password fail, AMF Plus will try to connect to the node using the factory default username and password (manager/friend). When a

new TQ replaces an existing TQ, this allows the new TQ to be discovered and managed as a guest node. AMF Plus can then start the guest node recovery procedure.

Login fallback is disabled by default and is only valid for model type **tq**.

Example:

To enable login fall back on the guest-class TQ6602, use the commands:

```
node1# configuration terminal
node1(config)# atmf guest-class TQ6602
node1(config-atmf-guest)# username example-user password example-pass
node1(config-atmf-guest)# login-fallback enable
```

Annotated x930 (AMF Plus Master) configuration

```
!
service password-encryption
!
hostname x930
!
no banner motd
!
username manager privilege 15 password 8 $1$bJoVec4D$JwOJGPr7YqoExA0GVasdE0
!
no service ssh
!
platform hwfilter-size ipv4-limited-ipv6
!
service telnet
!
service http
!
no clock timezone
!
snmp-server
!
aaa authentication enable default local
aaa authentication login default local
!
stack virtual-chassis-id 2132
!
atmf network-name MyNet
atmf master
atmf backup bandwidth 250
!
ip domain-lookup
!
no service dhcp-server
!
no ip multicast-routing
!
spanning-tree mode rstp
!
lacp global-passive-mode enable
!
switch 1 provision x930-52
!
interface port1.0.1
switchport
```

```

    switchport mode access
    !
interface port1.0.2
    switchport
    switchport atmf-link
    switchport mode trunk
    !
interface port1.0.3-1.0.50
    switchport
    switchport mode access
    !
interface vlan1
    ip address dhcp
    !
line con 0
    exec-timeout 0 0
line vty 0 4
    !

```

Annotated x230 configuration

```

!
service password-encryption
!
hostname x230
!
no banner motd
!
username manager privilege 15 password 8 $1$bJoVec4D$JwOJGPr7YqoExA0GVasdE0
!
no service ssh
!
platform hwfilter-size ipv4-limited-ipv6
!
service telnet
!
service http
!
no clock timezone
!
snmp-server
!
aaa authentication enable default local
aaa authentication login default local
!
stack virtual-chassis-id 1865
!
atmf network-name MyNet
!
atmf guest-class TQAP (With guest-class, dynamic discovery is the default)
modeltype tq
username manager password 8 9gzfX8VlZFw5ZKDsepGsKAm7LodYtmhJo5aK3vak7h8=
! (A username and password is mandatory when using the 'tq' modeltype)
ip domain-lookup
no service dhcp-server
!
no ip multicast-routing
!
service dhcp-snooping (Important; so we can detect when guests leave)
ip dhcp snooping delete-by-linkdown
!
spanning-tree mode rstp
!

```

```

lACP global-passive-mode enable

!
switch 1 provision x230-28
!
interface port1.0.1
  switchport
  switchport atmf-link

  switchport mode trunk
  ip dhcp snooping trust
  (Basic DHCP, it identifies the uplink port we expect DHCP requests to go to)
!
interface port1.0.2
  switchport
  switchport atmf-guestlink class TQAP
  (This is the port that we will connect the guest device to)

  switchport mode access
!
interface port1.0.3-1.0.50
  switchport
  switchport mode access

!
interface vlan1
  ip dhcp snooping (We want to snoop on this interface for guests)
  ip address dhcp (The interface on the guest parent interface MUST have an IPv4
address)
!
line con 0
  exec-timeout 0 0
line vty 0 4
!
end

```

TQ4600 Information



Basic Settings	Manage	Cluster	Status	Services	Maintenance
<i>Provide basic settings</i>				QoS	
Review Description of this Access Point ... These fields show information specific to this access point.				SNMP	
IP Address:	10.16.36.158			LED	
MAC Address:	EC:CD:6D:F3:0F:20			HTTP/HTTPS	
Firmware Version:	3.3.0			LLDP	
Build Number:	B02			NTP	
Build Date:	Mon Jul 25 19:23:44 2016				
Time since system-up:	00:28:28				
Provide Network Settings ...					

TQ firmware must be at least - 3.2.1 AO2



Managed mode must be Disabled

Basic Settings Manage Cluster Status

Configure Managed Access Point Parameters

Managed AP Administrative Mode Enabled Disabled

Controller IP Address 1

Controller IP Address 2

Controller IP Address 3

Controller IP Address 4

Base IP port

Pass Phrase Edit

WDS Managed Mode Root AP Satellite AP

WDS Managed Ethernet Port Enabled Disabled

Some sample show commands

Output 13: **show atmf nodes all** output (only available on AMF Plus masters and controllers)

```
x930#show atmf nodes all

Node and Guest Information:

* = Local device

SC = Switch Configuration:
C = Chassis   S = Stackable   N = Standalone G = Guest

Node/Guest      Device           ATMF           Parent          Node
Name            Type             Master  SC      Domain          Depth
-----
* x930           x930-52GTX      Y         S      none            0
  x230           x230-28GT       N         S      x930            1
  x230-1.0.2     AT-TQ4600       N         G      x230            -

Current ATMF node count 3 (guests 1)
```

Output 14: **show atmf guests** output (only available on AMF Plus masters and controllers)

```
x930#show atmf guests

Guest Information:

Device           Device           Parent           Guest           IP/IPv6
Name             Type             Node             Port            Address
-----
x230-1.0.2       AT-TQ4600       x230             1.0.2           10.16.36.158

Current ATMF guest node count 1
```

Output 15: `show atmf guests details` output (only available on AMF Plus masters and controllers)

```
x930#show atmf guests detail

ATMF Guest Node Information:

Node Name           : x230
Port Name           : port1.0.2
Ifindex             : 5002
Guest Description   : x230-1.0.2
Device Type         : AT-TQ4600
Backup Supported    : Yes
MAC Address         : eccd.6df3.0f20
IP Address          : 10.16.36.158
IPv6 Address        : Not Set
HTTP Port           : 0
Firmware Version    : 3.3.0 B02
```

Output 16: `show atmf backup guests` output (only available on AMF Plus masters and controllers)

```
x930#atmf backup guests now
x930#show atmf backup guests
Guest Backup ..... Enabled
  Schedule ..... 1 per day starting at 03:00
  Next Backup Time .... 07 Feb 2023 03:00
Backup Bandwidth ..... 250 KBps
Backup Media ..... SD (Total 1875.7MB, Free 1513.2MB)
Current Action ..... Idle
  Started ..... -
  Current Node ..... -
Backup Redundancy .... Disabled

Parent Node Name  Port Name          Date           Time           Status
-----
x230              port1.0.2          06 Feb 2023   13:42:12      Good
```

Output 17: `show atmf links` output (available on all AMF Plus nodes)

```
x230#show atmf links

ATMF Link Brief Information:


Local  Link      Link   ATMF   Adjacent  Adjacent  Link
Port   Type       Status State   Node/Area  Ifindex   State
-----
1.0.1  Uplink     Up     Full   x930      5002      Forwarding
1.0.2  Guestlink  Up     Full   x230-1.0.2 -         Active

* = Provisioned.
```

Structure of the backup media for guest node backups

The following diagram shows the structure of the backup media for guest node backups with the guest config backup file highlighted:

```
admin@amf-backup-server:~$ tree /media/sf_server-share/atmf/MyNet
/media/sf_server-share/atmf/MyNet
├── areas
├── guests
│   ├── info
│   └── x230
```



How to manually recover a failed TQ AP guest.

In this example, we have already replaced the failed device with another TQ-4600.

```

x230#debug atmf error
x230#terminal monitor
x230#atmf recover guest port1.0.2 (This command only takes a few seconds to complete.
Note that it must be run on the guest parent node. If the restore succeeds the AP
will be restarted automatically.)
15:07:12 x230 IMISH[6105]: [manager@ttyS0]atmf recover guest port1.0.2
15:07:12 x230 ATMFFSR[22152]: Attempting to ping master (x930) (attempt=1)
15:07:13 x230 ATMFFSR[22152]: Automatic node recovery ping -c 1 -w 1
172.31.1.12 was successful
15:07:13 x230 ATMFFSR[22152]: Node recovery pinged master (x930) reached
15:07:13 x230 ATMFFSR[22152]: Node recovery found 1 potential nodes with
recovery services
15:07:13 x230 ATMFFSR[22152]: Interrogating node x930 for backup
15:07:13 x230 ATMFFSR[22152]: Selected master recovery node x930
15:07:13 x230 ATMFFSR[22152]: RSYNC COMMAND: rsync -rtDqhW --modify-
window=1 --stats --itemize-changes --timeout=30 --bwlimit=250 --log-file=/
var/log/atmf_guest_recovery 172.31.1.12::EXMEDIA/MyNet/guests/nodes/x230/
port1.0.2 /tmp/.atmf_guest/recover
15:07:18 x230 DHCPSPN[15262]: Binding Delete: 10.16.36.158, chaddr
eccd.6df3.0f20, vlan1, port1.0.2, Server 10.16.36.254, Type Dynamic (with
691065 seconds remaining)
15:07:18 x230 NSM[602]: Port down notification received for port1.0.2

15:07:20 x230 NSM[602]: Port up notification received for port1.0.2
15:07:42 x230 MSTP[876]: CIST port1.0.2 now forwarding, propagating TC to
other ports
15:07:44 x230 DHCPSPN[15262]: Binding Add: 10.16.36.158, chaddr
eccd.6df3.0f20, vlan1, port1.0.2, Server 10.16.36.254, Type Dynamic,
Expires in 691200 seconds
x230#

```

Note that the text shown in pale grey above is only displayed if the terminal monitor is activated with the command **debug atmf error** enabled.

AMF Plus guest nodes

Overview

The AMF Plus guest node feature provides basic AMF Plus functionality to non-AMF Plus capable devices. These AMF Plus “guests” are devices that either do not run the AlliedWare Plus operating system or run a version that does not support AMF Plus. Essentially, any device that has either an IPv4 or IPv6 address can become an AMF Plus guest.

AMF Plus nodes can recognize the presence of an AMF Plus guest either statically, or dynamically if it uses a protocol such as DHCP or LLDP. Once recognized, the AMF Plus node is then able to provide a limited level of support to these devices.

A requirement of AMF Plus guests is that they each connect directly to their own port on an AMF Plus node with no intermediate devices.

Note: AMF Plus guest nodes are not supported on ports using the OpenFlow protocol.

AMF Plus parent nodes are aware, from the AMF Plus guest node configuration commands, which of their ports are connected to guest nodes. If configured to be dynamic, the AMF Plus node then listens on these ports for DHCP or LLDP frames, in order to obtain information about any of its attached AMF Plus guest devices.

Having discovered information about an AMF Plus guest, or by having its information configured statically, the AMF Plus node transmits this information to its local AMF Plus master(s). The AMF Plus guest information is then accessible to users or management tools that access the AMF Plus network information database.

Information about AMF Plus guests is not distributed to the whole AMF Plus network, it is only retained on the directly connected (parent) node, and the parent nodes master(s). This avoids consuming potentially large amounts of memory on the many nodes that may exist within a network.

Not all guest nodes are equal

For some types of AMF Plus guests, such as the TQ wireless access points, AMF Plus has a more detailed knowledge of the management interface. AMF Plus is, therefore, able to perform specific actions on these devices such as making backups of the configuration, or performing recovery on a replacement unit.

Other guest nodes, however, offer very little information that can be obtained by the AMF Plus parent. This is because they do not transmit frames that provide useful AMF Plus guest information. For these guest node types, information needs to be statically configured on the port of the AMF Plus node that directly connects to the AMF Plus guest.

AMF Plus guest discovery

As stated in the previous section, AMF Plus guests can be configured to be either static or dynamic depending on how the AMF Plus node detects their existence. If the guest node uses DHCP to get an IP address dynamically, then the AMF Plus parent node can use DHCP snooping to learn the address of the guest node. This process is termed **dynamic IP discovery**.

If the AMF Plus guest does not use DHCP to obtain an IP address, then the IPv4, or IPv6, address of the AMF Plus guest must be manually configured on the port of the parent AMF Plus node. This process is termed **static IP discovery**.

Once an AMF Plus node is aware of an AMF Plus guest's IP address, it can use this to interrogate the device for its MAC address.

Dynamic guest nodes

When AMF Plus guests are configured for dynamic discovery, the switchport on the AMF Plus parent is configured to use DHCP snooping to discover the guest's IP and MAC addresses. Additionally, if the guest supports LLDP and sends out LLDP packets advertising information about itself, this information can be used to discover the device's MAC address, and other details.

On devices that support LLDP-MED, this information originates from that provided by the equipment supplier. For other devices, information is retrieved from the descriptor string within the LLDP protocol.

Note that dynamic discovery is only available when using IPv4. For IPv6 networking, static IP discovery must be used. Also note that an IPv4 address is required on the VLAN interface on the parent node in order for dynamic discovery to work.

Static AMF Plus guests

AMF Plus guests configured for static IP discovery may not support protocols such as DHCP or LLDP and, therefore, offer the AMF Plus parent node no information to discover their IP address. For these nodes, information such as IP address must be statically configured on the port of the AMF Plus node that connects to the guest.

Please note the following:

- If a guest is statically configured but is also using DHCP and/or LLDP and there is a difference between the static and dynamic addresses, the static address will be used.
- A guest that has downstream devices with their own IP addresses assigned by DHCP - such as a router or a wireless access point - can only use dynamic discovery if LLDP discovery mode is chosen and the guest device is configured to send LLDP information. DHCP snooping cannot be used to dynamically discover guests with downstream device.

AMF Plus functionality supported by AMF Plus guests

The degree to which AMF Plus management capabilities are extended to an individual AMF Plus guest depends on the nature of the guest as defined by its model type.

By default, the minimum management functionality provided to a guest is basic topology reporting. At a minimum, AMF Plus will track and report a guest's network activity recorded at the link level. Whatever information is obtained about the guest and its operational status will be stored in an AMF Plus device database and can be viewed by users.

For some guest types, however, device-specific support capability is included in AMF Plus. If the AMF Plus network discovers that one of these advanced-support devices is connected, then it enables the user to carry out more management operations on that guest.

A guest model type is defined using the **modeltype** command, and can be one of the following:

- TQ access points
- Alliedware devices
- AlliedWare Plus devices
- ONVIF devices
- Other devices.

TQ access points

This model type applies to the Allied Telesis TQ series of wireless access points. Management support for TQ guest nodes extends to operational status monitoring at the link and network levels. AMF Plus Master nodes will also automatically or manually backup configuration files for TQ guest nodes, and allow for manual configuration recovery. This allows for easy reconfiguration of replacement TQ guest nodes.

Please note the following:

- AMF Plus Guest backup of TQ APs may fail if the APs are running in managed mode. If this happens, log in to the TQ and disable managed mode, then reboot.
- Guest backup of TQ APs may fail if there is a user logged on to the AP GUI.
- Reliable AMF Plus backup of TQ AP guest devices will only occur with TQ firmware release 3.2.1.a02 and later.

Login fallback

You can enable login fallback on TQ guest nodes using the command **login-fallback enable**. This feature allows AMF Plus to manage the TQ guest node using the factory default username/password if the saved username/password combination fails. For configuration details, see ["Replacing a TQ guest node with login fallback" on page 168](#).

Auto-recovery of TQ5403 series APs

From version 5.5.3-0.1 onwards, AMF Plus can auto-recover a failed TQ5403 series AP. For configuration details, see ["Recovering a TQ5403 series access point" on page 165](#).

AlliedWare devices

This model type offers additional support for Allied Telesis devices that are running the legacy AlliedWare operating system, which does not support AMF Plus. This support is limited to operational status monitoring at the link and network levels.

AlliedWare Plus devices

This model type offers support for devices that are running older versions of the Allied Telesis AlliedWare Plus operating system, which are not capable of running AMF Plus. This support is limited to operational status monitoring at the link and network levels.

ONVIF devices

This model type supports ONVIF (Open Network Video Interface Forum) Profile Q devices and offers AMF Plus system backup and manual recovery capabilities as well as operational status monitoring at the link and network levels.

Other devices

This model type is intended for devices that do not fit into any of the above categories. Support is limited to operational status monitoring at the link and network levels.

AMF Plus guest configuration

There are two components of AMF Plus guest configuration:

- guest-class configuration, and
- guest-link configuration

Guest-class configuration

Guest-classes are used to set up the general parameters for a class of AMF Plus guest devices. Each guest-class is a profile that can be assigned to multiple guests.

Guest-classes are modal templates that can be applied to selected guest types. Once you have created a guest-class, you can select it by entering its mode. From here, you can then configure a further set of operational settings specifically for this guest-class. These settings can then be applied to a guest-link by running the **switchport atmf-guestlink** command.

The following settings can be configured from within each guest-class mode:

- Discovery method** This can be static, dynamic, or (for TQ5403 series APs) agent. If unconfigured, the command will apply its default setting of dynamic. AR-series devices do not support DHCP snooping, therefore, dynamic guest-class discovery is only supported using LLDP on these products. If you want to configure auto-recovery on TQ5403 series APs, use the agent discovery type (and see ["Recovering a TQ5403 series access point" on page 165](#) for other necessary configuration).
- HTTP enable** This parameter is used to enable GUI access to a guest node. When **http-enable** is configured the port number is set to the default of 80. If the guest node is using a different port for HTTP, you can configure this using the port number attribute.
- Model type** This parameter can be set to one of *alliedware*, *aw+*, *onvif*, *tq*, or *other*. If unconfigured it will apply the default of *other*.
- User name and password** The username and password for devices supporting extended guest functionality. This is the username and password used to login and manage the device via HTTP.

Configuration method

The method applied in the following example assigns a guest configuration to a switchport of an AMF Plus node, then associates the guest with a guest-class profile. This will determine the method AMF Plus uses to interrogate the guest.

An IPv4 address is required on the VLAN interface of the parent AMF Plus node in order for dynamic discovery to work.

The following steps are used to define a guest-class:

Step 1: Define a guest-class name and enter the configuration mode for that guest-class.

Create a guest-class named 'Camera' and enter the guest-class configuration mode for that guest-class.

```
Parent-Node1(config)# atmf guest-class Camera
Parent-Node1(config-atmf-guest)#
```

Step 2: Configure parameters for the new class.

Select the model type for the guest-class to be a 'other'.

```
Parent-Node1(config-atmf-guest)# modeltype other
```

Set the username as 'manager' and the password as 'guestpass'.

```
Parent-Node1(config-atmf-guest)# username manager password guestpass
```

Step 3: Configure port on a parent node to know that it is connected to a guest node.

Select port2.1.1 to configure.

```
Parent-Node1(config)# interface port2.1.1
```

Configure the port to be a guest-link.

```
Parent-Node1(config-if)# switchport atmf-guestlink class Camera
```

Configure a TQ wireless access point with static discovery.

Step 1: Create a guest-class named TQ6602.

You must create the guest-class first, so that you can use it later in the **atmf-guestlink** command. Note that names of guest-classes are case-sensitive.

```
Parent-Node1(config)# atmf guest-class TQ6602
Parent-Node1(config-atmf-guest)#
```

Step 2: Configure parameters for the new class.

Select the model type for the guest-class to be a 'tq'.

```
Parent-Node1(config-atmf-guest)# modeltype tq
```

Set the username as 'manager' and the password as 'tq62-guestpass' for the TQ6602.

```
Parent-Node1(config-atmf-guest)# username manager password tq62-guestpass
```

Set the device for static discovery.

```
Parent-Node1(config-atmf-guest)# discovery static
```

Step 3: Configure the port on the parent node that the guest node is connected to.

Select port2.1.1 to configure.

```
Parent-Node1(config)# interface port2.1.1
```

Configure the port to be a guest-link.

```
Parent-Node1(config-if)# switchport atmf-guestlink class TQ6602 ip
192.168.10.1
```

Configure a TQ5403 wireless access point with agent discovery.

Step 1: Create a guest-class named TQ5403.

You must create the guest-class first, so that you can use it later in the **atmf-guestlink** command. Note that names of guest-classes are case-sensitive.

```
Parent-Node1(config)# atmf guest-class TQ5403
Parent-Node1(config-atmf-guest)#
```

Step 2: Configure parameters for the new class.

Select the model type for the guest-class to be a 'tq'.

```
Parent-Node1(config-atmf-guest)# modeltype tq
```

Set the username and password, for example, username of 'manager' and password of 'tq54-guestpass'.

```
Parent-Node1(config-atmf-guest)# username manager password tq54-guestpass
```

Set the device to use agent discovery.

```
Parent-Node1(config-atmf-guest)# discovery agent
```

Step 3: Configure the port on the parent node that the guest node is connected to.

Select port2.1.1 to configure.

```
Parent-Node1(config)# interface port2.1.1
```

Configure the port to be a guest-link. If the AP has a static IP address of 192.168.10.1, use the command:

```
Parent-Node1(config-if)# switchport atmf-guestlink class TQ5403 ip
192.168.10.1
```

If the AP has a dynamically-assigned address, use the command:

```
Parent-Node1(config-if)# switchport atmf-guestlink class TQ5403
```

Configure a camera guest node with dynamic DHCP discovery.

For more information on DHCP snooping see the [DHCP Snooping Feature Overview and Configuration Guide](#)

Step 1: Enable the DHCP snooping service.

```
Parent-Node1(config)# service dhcp-snooping
```

Step 2: Configure a VLAN with an IP address and allow DHCP snooping on it.

```
Parent-Node1(config)# interface vlan2
```

Enable DHCP snooping on this VLAN.

```
Parent-Node1(config-if)# ip dhcp snooping
```

Assign an IP address to the VLAN.

```
Parent-Node1(config-if)# ip address 192.168.2.1/24
```

Step 3: Create a guest-class named Camera.

```
Parent-Node1(config)# atmf guest-class Camera
```

```
Parent-Node1(config-atmf-guest)#
```

Step 4: Configure parameters for the new class.

Select the modeltype for the guest-class to be a 'onvif'.

```
Parent-Node1(config-atmf-guest)# modeltype onvif
```

Set the username as 'admin' and the password as 'secret' for the class 'Camera'.

```
Parent-Node1(config-atmf-guest)# username admin password secret
```

Enable http GUI access for the guest node.

```
Parent-Node1(config-atmf-guest)# http-enable
```

Step 5: Configure the port on the parent node that the guest node is connected to.

Select port1.0.2 to configure.

```
Parent-Node1(config)# interface port1.0.2
```

Assign the port to VLAN 2.

```
Parent-Node1(config)# switchport access vlan 2
```

Configure the port to be a guest-link.

```
Parent-Node1(config-if)# switchport atmf-guestlink class Camera
```

Configure a guest node with LLDP discovery.

For more information on LLDP see the [LLDP Feature Overview and Configuration Guide](#)

Step 1: Enable LLDP on the AMF Plus parent node.

```
Parent-Node1(config)# lldp run
```

Step 2: Create a guest-class for the guest node.

```
Parent-Node1(config)# atmf guest-class LLDPsample
Parent-Node1(config-atmf-guest)#
```

Step 3: Configure parameters for the new class (these are device dependent).

Select the model type for the guest-class to be a 'other'.

```
Parent-Node1(config-atmf-guest)# modeltype other
```

Set the username as 'admin' and the password as 'secret' for the class 'LLDPsample'.

```
Parent-Node1(config-atmf-guest)# username admin password secret
```

Step 4: Configure the port on the parent node that the guest node is connected to.

Select port1.0.5 to configure.

```
Parent-Node1(config)# interface port1.0.5
```

Allow LLDP TLVs to be transmitted by the port.

```
Parent-Node1(config)# lldp tlv-select all
```

Configure the port to be a guest-link.

```
Parent-Node1(config-if)# switchport atmf-guestlink class LLDPsample
```

AMF Plus guest node show commands

Guest node show commands can be executed on either AMF Plus masters or AMF Plus controllers. When executed on a master, the command displays guest node information for the local AMF Plus area only, and when executed on a controller the command displays output for all AMF Plus areas. The **show** command syntax is different on masters and controllers.

show atmf guests

This command displays a list of guests that an AMF Plus network has discovered. It needs to be run on an AMF Plus master.

Output 18: Example output from the `show atmf guests` command

```

master#show atmf guests

  Guest Information:

  Device          Device          Parent          Guest          IP/IPv6
  Name            Type            Node            Port            Address
  -----
  master-2.1.1    AR415S          master          2.1.1          192.168.2.10
  master-2.1.2    AT-9924T        master          2.1.2          192.168.1.10
  master-2.1.4    AT-TQ4600       master          2.1.4          192.168.1.12

  Current ATMF guest node count 3

```

Figure 7: Parameter Descriptions from the `show atmf guests` command

PARAMETER	DESCRIPTION
Device Name	The name assigned for this device within the AMF Plus network. It could be a name that is discovered from the device, or failing that, a name that is auto-assigned by AMF Plus. The auto-assigned name consists of <parent node name>-<attached port number>
Device Type	This is the product name of the guest node and is discovered from the device. If no device Type can be discovered, then the model name configured on the guest-class assigned to the connected port is used.
Parent Node	The AMF Plus member name of the member that directly connects to the guest node.
Guest Port	The port on the parent node that directly connects to the guest node.
IP/IPv6 Address	The address discovered from the node, or statically configured on the parent node's attached port.

show atmf guests detail

This command displays the details of each discovered guest node. It needs to be run on an AMF Plus master.

Output 19: Example output from the `show atmf guests detail` command

```

master#show atmf guests detail

ATMF Guest Node Information:

Node Name           : master
Port Name           : port2.1.1
Ifindex             : 6101
Guest Description   : master-2.1.1
Device Type         : AR415S
Backup Supported    : No
MAC Address         : 0000.cd1d.b114
IP Address          : 192.168.2.10
IPv6 Address        : Not Set
HTTP Port           : 0
Firmware Version    : 2.9.2-09

```



```

Node Name           : master
Port Name          : port2.1.4
Ifindex            : 6104
Guest Description   : master-2.1.4
Device Type        : AT-TQ3200
Backup Supported    : Yes
MAC Address        : 001a.eb85.fd60
IP Address         : 192.168.1.12
IPv6 Address       : Not Set
HTTP Port          : 0
Firmware Version   : 3.2.1 A02

```

Figure 8: Parameter Descriptions from the **show atmfguests detail** command

PARAMETER	DESCRIPTION
Node Name	The AMF Plus device that directly connects to the guest node, i.e. the parent node.
Port Name	The port on the parent node that directly connects to the guest node.
Ifindex	An internal index number that maps to the port number of the guest's parent node.
Guest Description	Either a description that a user has manually entered by using the description command in interface mode for the interface, or a default description consisting of the AMF Plus parent node name plus the port number that connects it to the guest.
Device Type	A device type name that is extracted from the device.
Backup Supported	Indicates whether AMF Plus is able to backup this device.
MAC Address	The MAC address of the guest node.
IP Address	The IP address of the guest node.
IPv6 Address	The IPv6 address of the guest node.
HTTP Port	The HTTP port enables you to specify a port when enabling HTTP to allow a URL for the HTTP user interface of a guest node. This is determined by the http-enable <port> command.
Firmware Version	If available, the firmware level that the guest is running, as extracted from the guest node.

show atmf links guests

This command shows the details of guest-links on the parent node.

Output 20: Output from the **show atmf links guests** command

```
Parent-Node1#show atmf links guest

Guest Link Information:

DC = Discovery configuration
S = static D = dynamic

Local      Guest      Model      MAC      IP / IPv6
Port       Class      Type       DC Address Address
-----
2.1.1     alliedware Alliedware S 0000.cd1d.b114 192.168.2.10
2.1.2     awdyn      Alliedware D 0000.cd24.023a 192.168.1.10
2.1.4     TQ         TQ         S 001a.eb85.fd60 192.168.1.12
2.1.6     FireSensor Other      S -                2001:af34:93::fe9
2.1.7     -          Other      D -                -
```

Parameter descriptions from the **show atmf links guest** command

PARAMETER	DESCRIPTION
Local Port	The local port on which the guest-link is configured.
Guest Class	The guest-class that has been configured on the local port.
Model Type	The model that has been defined in that associated guest-class.
DC	The device discovery method, S=static and D=dynamic.
MAC Address	The MAC address that has been discovered for the guest node.
IP/IPv6 address	The IP4 or IPv6 of the guest node. This may be either discovered or statically entered.

AMF Plus support for ONVIF Profile Q devices

ONVIF (Open Network Video Interface Forum) Profile Q devices are IP-based network devices (for example network cameras, network switches, or network monitors), which can be discovered and configured by a Profile Q client (or manager). For more information on ONVIF and Profile Q devices, see the ONVIF web site <http://onvif.org>.

ONVIF devices are managed as an AMF Plus guest node by configuring the guest-class with **modeltype onvif**. The ONVIF nodes can be configured with either a static IP address or a dynamic address using DHCP snooping. Communication between the AMF Plus parent node and the ONVIF guest node relies on HTTP.

Once configured an ONVIF guest node can be:

- backed up manually or automatically using the AMF Plus guest backup feature.
- restored manually using the AMF Plus guest recovery feature.

Note that only one ONVIF device is supported per port.

Configure an ONVIF guest node with a static IP address

To configure an ONVIF guest node with a static IP address, do the following:

1. Consult the ONVIF device's user manual to identify the username, password, and IP address of the device. If necessary use the device management application, or GUI provided by the device, to modify the user account and network setup.
2. Configure an ONVIF guest-class on the AMF Plus parent node. This should include the username, password, and port information for connecting to the device via HTTP.
3. Create an AMF Plus guest-link using the with the configured ONVIF guest-class.

Figure 9: Configuration for an ONVIF guest node with a static IP address.

```
atmf guest-class camera
modeltype onvif
username admin password 8 ntAaRMi+IXGOg/SDpjeebGrublIw8BtPw4bwbhGTQhE=
http-enable
discovery static
interface port1.0.1
switchport
switchport mode access
switchport atmf-guestlink class camera ip 192.168.1.66
```

Configure an ONVIF device with a dynamic IP address

To configure an ONVIF guest node with a dynamic IP address, do the following:

1. Consult the ONVIF device's user manual to identify the username and password. If necessary use the device management application, or GUI provided by the device, to modify the user account and network setup. Ensure the device is configured to get its IP address via DHCP.
2. Enable DHCP snooping on the AMF Plus parent node.
3. Create a VLAN for the ONVIF guest node, assign an IP address range to this VLAN, and enable DHCP snooping on it.
4. Configure an ONVIF guest-class on the AMF Plus parent node. This should include the username, password, and port information for connecting to the device via HTTP.
5. Create an AMF Plus guest-link using the configured ONVIF guest-class.
6. Assign the port that the guest-node is attached to, to the VLAN you created before.

Figure 10: Configuration for an ONVIF guest node with a dynamic IP address.

```

service dhcp-snooping
interface vlan2
  ip address 192.168.2.1/24
  ip dhcp snooping
atmf guest-class dynonvif
  modeltype onvif
  username admin password 8 ntAaRmi+IXGOg/SDpjeebGrublIw8BtPw4bwbhGTQhE=
  http-enable
interface port1.0.2
  switchport
  switchport mode access
  switchport access vlan 2
  switchport atmf-guestlink class dynonvif

```

Backing up an ONVIF device

Once an ONVIF device is detected by AMF Plus, and if that ONVIF device supports system backup and restore, AMF Plus automatic and manual backup features will be available for that node.

Note that the **show atmf links guest detail** command shows whether an ONVIF guest node supports system backup and restore in the 'Backup Supported' field.

The following commands are available:

- **(no) atmf backup guests enable** to enable or disable automatic guest node backup.
- **atmf backup guests now** to start a manual backup.
- **show atmf backup guests** to check the status of guest node backups.

The ONVIF system backup file is considered as binary data as the file format is vendor-specific. The contents are simply uploaded to the ONVIF device when a recovery is executed.

Restoring an ONVIF device

An ONVIF device's configuration can be restored manually from the previous system backup. To do this:

- Connect the new ONVIF device to same port the old device was attached to.
- Make sure the new device is the same, or at least compatible, with the old device.
- Use the **atmf recover guest <guest-port>** command to start the recovery process.
- Once the ONVIF device is successfully restored it may be necessary to reboot the device.

The ONVIF system backup file is considered as binary data as the file format is vendor-specific. The contents are simply uploaded to the ONVIF device when a recovery is executed.

Show commands for ONVIF guest nodes

Run the **show atmf links guest** command on the AMF Plus parent node to see a list of all locally configured ONVIF guest nodes.

Output 21: Example output from the **show atmf links guest** command

```
Parent-Node1# show atmf links guest

Guest Link Information:

DC = Discovery configuration
S = static D = dynamic

Local      Guest      Model      MAC      IP / IPv6
Port      Class     Type       DC Address Address
-----
port1.0.1 camera    ONVIF      S 001f.553d.65ba 192.168.1.66
port1.0.2 dynamic  ONVIF      D 001f.5541.cd9d 192.168.2.68

Total number of guest links configured 2
```

Run the **show atmf links guest detail** command on the AMF Plus parent node to see details about the configured guest nodes.

Output 22: Example output from the **show atmf links guest details** command

```
Parent-Node1# show atmf links guest detail

Detailed Guest Link Information:

Interface                : port1.0.1
  Link State              : Full
  Class Name              : camera
  Model Type              : ONVIF
  Discovery Method        : Static
  IP Address              : 192.168.1.66
  Username                : admin
  Node State              : Full
  Backup Supported        : No
  MAC address             : 001f.553d.65ba
  Device Type             : HEW4PER2B
  Description             : x930-1.0.1
  Serial Number           : C057100245
  Firmware Version        : 1.000.HW01.2 build: 2019-05-14
  HTTP port               : 80

Interface                : port1.0.2
  Link State              : Full
...
```

To check the backup status of the ONVIF guest nodes run the **show atmf backup guest** command on the AMF Plus master.

Output 23: Example output from the **show atmf backup guest** command

```
master# show atmf backup guest
Guest Backup ..... Enabled
  Schedule ..... 1 per day starting at 03:00
  Next Backup Time .... 23 Feb 2023 03:00
Backup Bandwidth ..... Unlimited
Backup Media ..... SD (Total 3668.4MB, Free 2817.5MB)
Current Action ..... Idle
  Started ..... -
  Current Node ..... -
Backup Redundancy ..... Disabled
Parent Node Name  Port Name      Date           Time           Status
-----
x930              port1.0.1     22 Feb 2023   03:00:13      Good
x930              port1.0.2     22 Feb 2023   03:00:01      Good
```

AMF Plus Backups

AMF Plus backups are a valuable part of AMF Plus network operation. Backups ensure that appropriate devices within the AMF Plus network have copies of other devices' information and files. This means that if a node fails, the AMF Plus network can automatically configure its replacement.

Backups by different types of nodes

Backups can be performed by either a master, controller, or member node. The operational details of backups differs between the three types of nodes.

Backups by master nodes

The backups performed by master nodes are a fundamental part of their contribution to the operation of an AMF Plus area.

They are the mechanism by which AMF Plus master nodes update their records of their AMF Plus area, and so have all the information and files required to enable automatic node recovery. By default, AMF Plus master nodes are configured to perform automatically scheduled backups of the entire AMF Plus area once per day at 3:00 a.m. AMF Plus masters can store their backups either on **remote file servers** or on **removable media** such as USB sticks or SD cards. These backup files can be used in the recovery of a failed node.

It is also possible to initiate a manual backup of the AMF Plus network from a master node.

Backups by controller nodes

By default, controller nodes do not perform backups. However, they can be configured to backup the master nodes of all their controlled areas on a regular scheduled basis. Alternatively they can be used to initiate a backup on a specified of area's master nodes immediately.

Controller nodes backup only the master nodes in their controlled areas. They do not backup member nodes.

Backups by member nodes

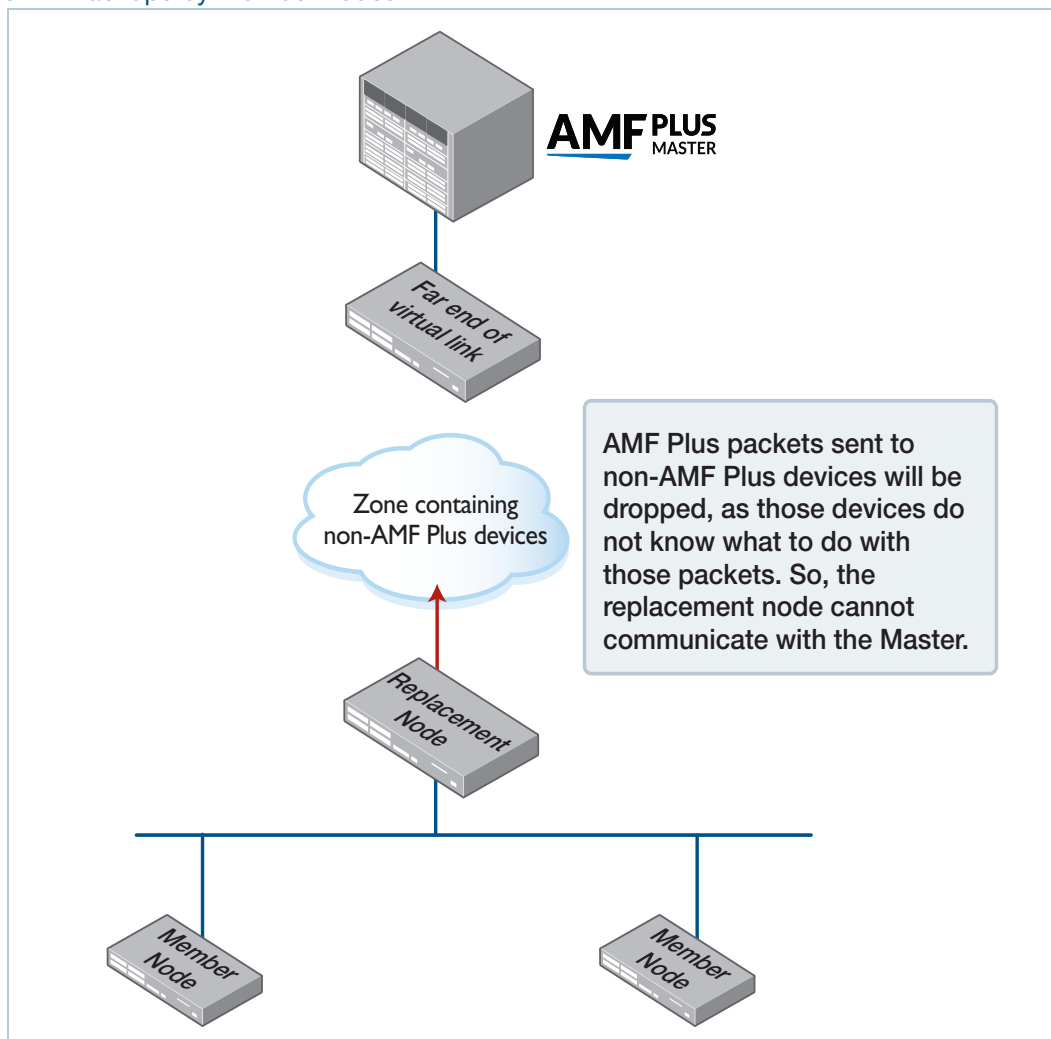
The act of a member node performing a limited backup of another member node is required in order to deal with a specific circumstance - namely when a portion of an AMF Plus network is located at the end of a virtual-link.

The automatic recovery of a unit that terminates a virtual-link, and accesses its master via that virtual-link, is less straightforward than the recovery of a unit that has a direct AMF Plus uplink to its higher-domain neighbor.

If a device has a direct AMF Plus uplink to its higher-domain neighbor and needs to be replaced; all the replacing device needs to do is send AMF Plus messages from its uplink port to its neighbor. Its neighbor responds with information about the device being replaced. The replacing device then requests the master to provide it with the files necessary to fully adopt the role of the replaced device.

However, if a unit is at the end of a virtual-link, there will be one or more non-AMF Plus devices between itself and the other end of the virtual-link. So, if it sends out AMF Plus messages towards the master, they will simply be dropped by the non-AMF Plus devices in between.

Figure 11: Backups by member nodes



An alternative solution is required to enable recovery in this situation. The solution is to enable the member node neighbors of the virtual-link endpoint node to perform limited backups of the virtual-link endpoint node.

The node at the end of the virtual-link pushes its startup configuration to its adjacent neighbors. At the time of recovery, the replacement unit then fetches the startup configuration back and applies the configuration. This provides it with the configuration required to establish the virtual-link, and thereby make contact with the master node, to obtain the rest of the files it needs to complete a full recovery.

The AMF Plus messages that the adjacent nodes send to the replacement node indicate that the sending node has the configuration file that the replacement node needs.

If a recovering node detects a neighbor that indicates it has the required configuration file, it can then download and apply the configuration file from that neighbor. This restores the recovering node to its prefailure configuration.

With the original configuration restored the AMF Plus virtual-link becomes operational and the recovering node can now connect to the area master.

Figure 12: Recovery step 1

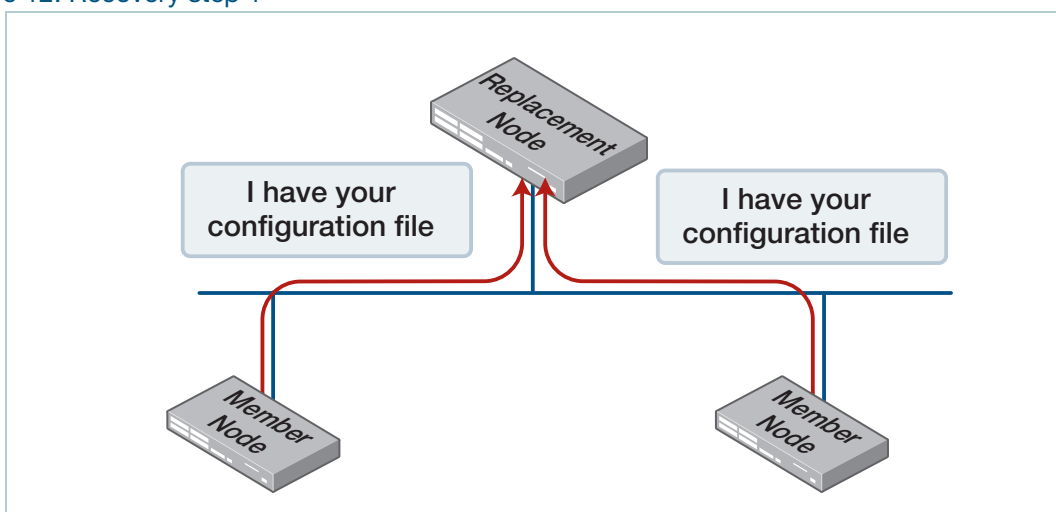


Figure 13: Recovery step 2

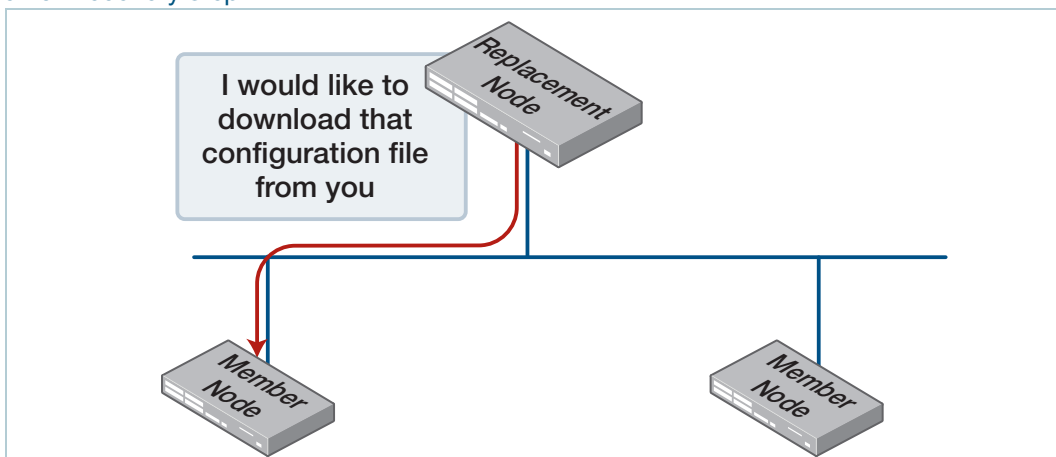

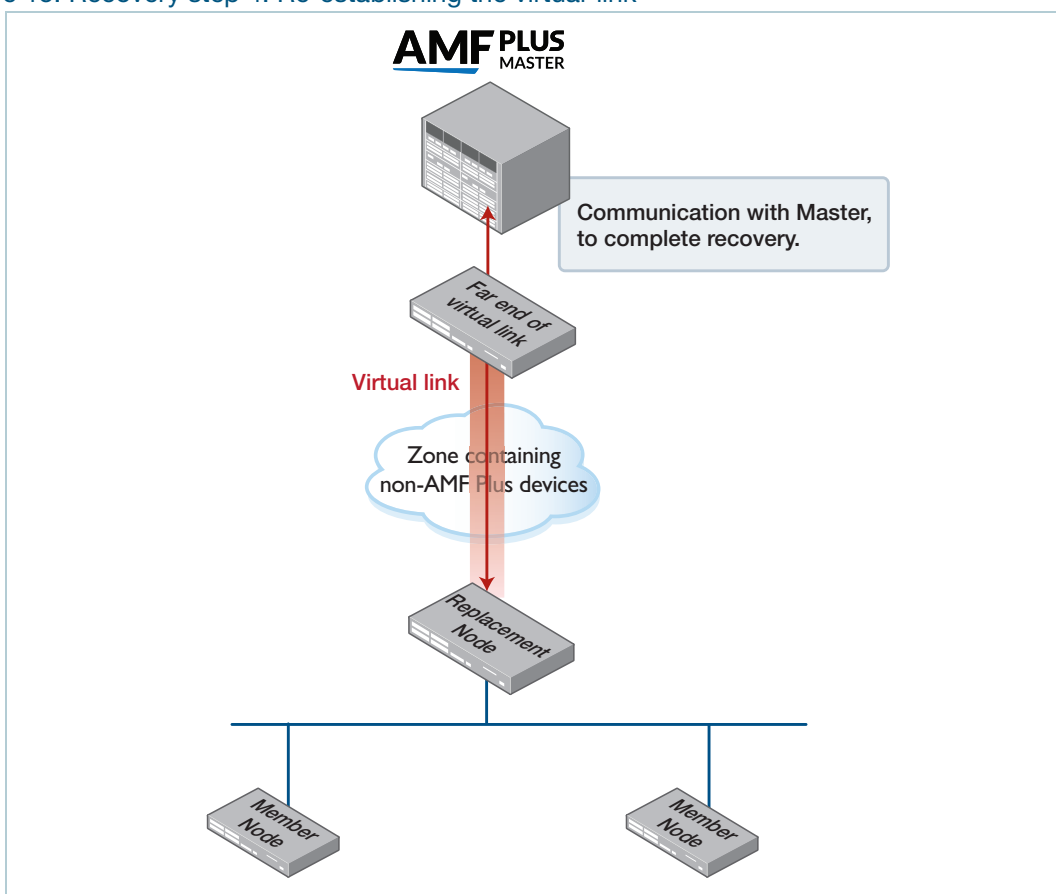


Figure 14: Recovery step 3



The recovering node then reboots and comes up again running the configuration that includes the definition of the virtual-link. It re-establishes the virtual-link, and communicates with the master to complete the recovery.

Figure 15: Recovery step 4: Re-establishing the virtual-link



Events that trigger the pushing of the configuration file to neighbor nodes

On appropriate occasions, the node at the end of the virtual-link needs to push its configuration file to its neighbors.

These appropriate occasions include when:

- the running configuration on the node is saved to startup configuration.
- the node detects that a new AMF Plus neighbor has been directly connected to it.
- the virtual-link on the node goes to the "FULL" state.
- a gateway node connected to the virtual-link changes state from non-master to master.

It is important to note that the node terminating the virtual-link does not push its configuration file to its directly connected neighbors at regular scheduled intervals, but only when one of the above events occurs. Therefore, if you want to ensure that the configuration file of the endpoint of a virtual-link has been backed up to its neighbors, then you should simply copy the running config to startup config. This will also force the backup to the neighbors to occur.

Which files are backed up?

The full backup process (i.e. the backup performed by a master or controller node) backs up most of the files on a node, including its configuration, AlliedWare Plus images, license files, and more.

The sorts of files that are not backed up are:

- stacking configuration
- coredump, exception log, tech support etc.
- DHCP Snooping database
- history and file editor state
- random number seed
- signature files for security services, because they go out of date—on start-up the firewall will check for, and if necessary, download the latest file-set
- password files—on start-up the password files will be regenerated from the running-config.

Backup destinations

Controller and master nodes can save backups either on a separate file server, or on removable media, such as a USB stick or SD card, installed locally.

The backups can be saved to whichever type of removable media that the controller or master supports - USB or SD card. We recommend using the ext3 or ext4 filesystem on external media that are used for AMF Plus backups.

When backing up to remote file servers, up to two servers can be specified as backup destinations for any given master or controller node.

A good level of resilience can be provided by sending backups to both removable media and remote file servers at the same time as enabled by the command **atmf backup redundancy**.

When a master or controller is utilizing more than one backup destination at once (either two remote file servers or remote file servers and removable media) there are extra rules and commands that relate to this multi-destination mode of operation.

Note: If multiple AMF Plus Areas are configured to use the same backup server, with the same path a conflict could arise.

For example, If there is a device called Router1 in both Areas 1 and 2, the backups will both go to the same directory called `.../nodes/Router1` and overwrite each other. To avoid this conflict, it is recommended that you configure different paths for different AMF Plus Areas.

These are discussed below in the section "[Multiple backup destinations](#)" on page 209.

If an AMF Plus master is storing backup data on removable media then:

- if the AMF Plus master is a SBx8100 system with dual CFC controllers, both CFCs should have removable media installed.
- if the AMF Plus master is a VCStack, all stack members should have removable media installed.
- the removable media installed must have sufficient capacity to hold all of the relevant files stored in the Flash on every node in the AMF Plus area, including other master nodes. Files that are backed up include all configuration files, release files, and scripts, but do not include files like core dumps, exception logs, or technical support files.

Typically a 4 GB SD card or USB storage device would hold backups for a 40 node AMF Plus area.

You can store other data on the storage device as long as you make sure that enough space is reserved for future AMF Plus backups.

AMF Plus requires up to 128 MB backup space for SBx8100 nodes and up to 64 MB backup space for other nodes. The output from the **show atmf backup** command will provide warnings if free space on the backup media falls below a safe level.

Output 24: Output showing backup media space warning

```

master1#show atmf backup

Scheduled Backup ..... Disabled
  Schedule ..... 1 per day starting at 12:45
  Next Backup Time .... 14 Feb 2023 12:45
Backup Media ..... SD (Total 3827.0MB, Free 7.1MB)
                               WARNING: Space on backup media is below 64MB
Current Action ..... Idle
  Started ..... -
  Current Node ..... -

```

Safe removal of removable storage media

Removing removable storage media, or rebooting the controller or master node, while an AMF Plus backup is underway could potentially cause corruption to files in the backup. Although files damaged as a result of mishandling backup media will be replaced during the next backup cycle, if the file system on the media becomes damaged, it may require reformatting before being usable again. To avoid any damage to the AMF Plus backup files or file system, we recommend that the following procedure be followed before rebooting or removing any removable storage media from an AMF Plus master or controller:

1. Disable backups to prevent a scheduled backup from occurring while the card is being removed.
2. Terminate any backup already in process.
3. Verify that it is safe to remove the media by checking that backups are disabled and that there are no backups currently in progress.

Output 25: Example of the safe storage media removal procedure

```

master1#conf t

master1(config)#no atmf backup enable
master1(config)#exit
master1#atmf backup stop
master1#show atmf backup

Scheduled Backup ..... Disabled
  Schedule ..... 1 per day starting at 12:45
  Next Backup Time .... 14 Feb 2023 12:45
Backup Media ..... SD (Total 3827.0MB, Free 3257.1MB)
Current Action ..... Idle
  Started ..... -
  Current Node ..... -

```

Once the media has been reinstalled, ensure that the backup scheduler is re-enabled.

Controlling the backup behaviour of controller and master nodes

By default, master nodes will perform a backup of their whole area every day at 3:00 a.m. The default behavior for controller nodes is not to perform any backups at all.

On master nodes, the performing of backups can be turned on and off with the command:

```
(no) atmf backup enable
```

Once backups have been enabled, you can trigger an immediate backup and/or schedule regular backups, as described in the next section.

On controller nodes, backups can be either enabled or disabled by using the command:

```
(no) atmf backup area-masters enable
```

Scheduling backups

Scheduling backups, on both controllers and masters, can be configured with the command:

```
atmf backup {default|<hh:mm> frequency <1-24>}
```

For example, to schedule three backups per day, with the first backup at 7:20 a.m., the command is:

```
atmf backup 07:20 frequency 3
```

Triggering immediate backups

On a master node, the backup of all nodes in its area, or a specified node, can be kicked off immediately with the command:

```
atmf backup now [<nodename>]
```

On a controller, the backup of the masters in all of its controlled areas, or all or one of the masters in a particular area, can be kicked off immediately with the command:

```
atmf backup area-masters now [area <area-name>|area <area-name> node <node-name>]
```

Stopping a backup

To stop a backup that is currently in progress, use the command:

```
atmf backup stop
```

Deleting backed up files

On a controller, you can delete the backup files for a given master with the command:

```
atmf backup area-masters delete area <area-name> node <node-name>
```

On a master, you can delete the backup files for a given node with the command:

```
atmf backup delete <node-name>
```

Performing a manual backup

Whenever a new device is physically added to the AMF Plus network as a provisioned node, we recommend that you perform a manual backup from the AMF Plus master.

To perform a manual backup of the entire AMF Plus area, on the AMF Plus master enter the command **atmf backup now**:

```
master1# atmf backup now
master1(config)# atmf backup enable
master1(config)# exit
```

You can perform a manual backup of a single AMF Plus node by running the following commands on the AMF Plus master:

```
master1# atmf backup now <node-name>
master1(config)# atmf backup enable
master1(config)# exit
```

To check the status of the AMF Plus backup, use the **show atmf backup** command. The “Date”, “Time”, and “On Media” details update once the backup for that node is finished.

Output 26: Example output from the **show atmf backup** command entered during a backup

```
AMF_Master#show atmf backup
Scheduled Backup ..... Enabled
  Schedule ..... 1 per day starting at 03:00
  Next Backup Time .... 14 Feb 2023 03:00
Backup Media ..... USB (Total 3692.6MB, Free 1782.7MB)
Current Action ..... Doing manual backup
  Started ..... 14 Feb 2023 05:20
  Current Node ..... Member1
Backup Redundancy ..... Disabled
```

Node Name	Date	Time	In ATMF	On Media	Status
AMF_Master	14 Feb 2023	05:20:16	Yes	Yes	Good
Member1	-	-	Yes	Yes	In Progress
Member2	-	-	Yes	No	-
Member3	-	-	Yes	No	-
Member4	-	-	Yes	No	-

Below is example output from the **show atmf backup** command entered after the backup has completed.

Output 27: Example output from the **show atmf backup** command entered after backup was completed

```
AMF_Master#show atmf backup
Scheduled Backup ..... Enabled
  Schedule ..... 1 per day starting at 03:00
  Next Backup Time .... 14 Feb 2023 03:00
Backup Media ..... USB (Total 3692.6MB, Free 1651.1MB)
Current Action ..... Idle
  Started ..... -
  Current Node ..... -
Backup Redundancy ..... Disabled
```

```
-----
Node Name                Date           Time           In ATMF  On Media  Status
-----
ATMF_Master              13 Feb 2023   05:20:16      Yes      Yes      Good
Member1                  13 Feb 2023   05:20:27      Yes      Yes      Good
Member2                  13 Feb 2023   05:20:40      Yes      Yes      Good
Member3                  13 Feb 2023   05:20:52      Yes      Yes      Good
Member4                  13 Feb 2023   05:21:08      Yes      Yes      Good
```

Note that the file system used by the AMF Plus backup will not backup files that have the same name but different case (e.g. “test.txt” and “TEST.txt”). Only **one** of these files will be stored in the backup. For this reason we recommend that all files on a node be given unique file names.

Backups on chassis or VCStacks running as AMF Plus controllers or masters

This section is only applicable in configurations that are **not** using remote backup servers.

When a chassis is operating as a controller or master node, AMF Plus backups will only occur on the removable media of the CFC that is the chassis master. Therefore, in the event of a CFC failure, the new master CFC will have no access to this backup information.

The same applies to the case of a VCStack performing backups. Only the Stack master will be carrying out backups.

To avoid this situation, you can either configure a remote backup file server or use trigger scripts to automatically perform a manual backup of the AMF Plus network following a failover event. This section provides some example trigger scripts to automatically apply a manual backup. To apply the remote file server solution see "[Backing up to remote servers](#)" on page 203.

Example 1 This example uses a manual backup activation script called **triggered-atmfbackup.scp**. When activated, this script applies the following commands to initiate a network backup:

```
enable
wait 180
atmf backup now
```

When a CFC failure event occurs, the trigger **type chassis active-CFC-fail** will activate. The following example shows how the above scripted steps can be automatically applied if this event occurs.

This example shows a trigger script configuration for the **SBx8100**:

```
master1# conf t
master1(config)# trigger 1
master1(config-trigger)# type chassis active-CFC-fail
master1(config-trigger)# script 1 triggered-atmfbackup.scp
```

To explain the sequence; if there is a failure of a CFC that is operating as a chassis master, trigger 1, which is associated with the trigger **type chassis active-CFC-fail**, will activate.

This process runs the script **triggered-atmfbackup.scp**, which will then execute the command to carry out an atmf backup immediately.

Example 2 In the event of a VCS master failure, the trigger **type stack master-fail** will activate. The following example shows how the above scripted steps can be automatically applied if this event occurs.

This example shows a trigger script configuration that can operate when a stack master node fails:

```
Master1# conf t
Master1(config)# trigger 1
Master1(config-trigger)# type stack master-fail
Master1(config-trigger)# script 1 triggered-atmfbackup.scp
```

To explain the sequence; if there is a failure of a node that is operating as a stack master, trigger 1, which is associated with the trigger **type stack master-fail**, will activate.

This process runs the script **triggered-atmfbackup.scp**.

Forcing all master nodes in an area to perform a backup

If there are multiple AMF Plus master nodes in an AMF Plus area, you may also want to use a trigger script or perform a manual backup by **all** master nodes after a failover event, so that all backups are up to date.

Create an AMF Plus working-set group that contains all master nodes in an area, then use the **atmf working-set** command in the trigger script to execute the manual backup on all nodes within the working-set.

To create a working-set containing all AMF Plus master nodes in an area, first manually select all AMF Plus masters using the **atmf working-set** command:

```
Master1# atmf working-set Master1,Master2
NetworkName[2]# conf t
NetworkName[2](config)# trigger 1
```

This command displays an output screen similar to the one shown below:

```
=====
Master1, Master2
=====

Working set join

ATMF1[2]#
```

Enter configuration commands, one per line. End with CNTL/Z

```
ATMF1{2}# conf t
ATMF1[2](config)# trigger 1
ATMF1[2](config-trigger)# type stack master-fail
ATMF1[2](config-trigger)# script 1 triggered-atmfbackup.scp
```

Enter configuration commands, one per line. End with CNTL/Z:

```
ATMF1{2}# conf t
ATMF1[2](config)# trigger 2
ATMF1[2](config-trigger)#type chassis active-CFC-fail
ATMF1[2](config-trigger)#script 1 triggered-atmfbackup.scp
```

Next, create a user-defined working-set group containing the nodes in the current working-set using the **atmf group (membership)** command:

```
atmf1[2]# conf t
atmf1[2](config) atmf group AMF_masters
```

You could also carry out a manual backup by all the masters in the area by using the commands:

```
atmf working-set group AMF_masters
atmf backup now
```

Backing up to remote servers

Because the connection to the remote server(s) must be secure, there are a few steps to setting up the Remote Server backup.

Setting up SSH keys with the file server

To enable AMF Plus Remote backup, SSH keys need to be setup on the file server.

The following points and processes apply:

1. For the File Server

- Any modern Linux server can be used.
- The server destination file system should support file permissions. Note, that the FAT32 file system is **not** supported.
- Default OpenSSH versions will work without modifying the SSH settings.
- An OpenSSH server daemon must be running with its default settings.
- The ability to add known SSH keys to a user.
- There must be enough hard drive space on the server (at worst: the sum of all Flash space on all devices in the network).
- Users must be given write access to the folder that is the root of the backup folders.

2. For the Public Key

For security reasons, the exchange of user and host keys must be done manually by the user using existing crypto commands. This can be done either before or after configuring the remote file server.

- A public host key from the remote file server needs to be imported to the AMF Plus backup node.
- A public user key for root from the AMF Plus backup node needs to be installed on the remote file server.

The process for this is described in the next two sections.

Importing a host key from the remote file server

The following process imports an **RSA** host key from the remote file server with the IP address of **10.37.165.65**.

Output 28: Importing a host key from remote file server

```
x950(config)#crypto key pubkey-chain knownhosts ip 10.37.165.65 rsa

10.37.165.65 ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQD4LknpjHYXhrFU6/
0vS2PTsIkvKh8j0JkIwIMiBHaxJVadED73G6iq4C6Lk
NB+M7BZOohF/
ow0Bhp4Jr8XK0Mfm09gfriRHNNQlGsbfpugDKXnJpFMU88Yu3MaSrkgULgtD7B0MXBEw05H
PNjo1RCR9KI3Z3GFGM1TJy8T/6
xikczyaxbhqfUeqtpMgDMzRhieqIdpl7Umg4fJxhMDSHa8af0HrpRpntsw23+h5IUX9Sw+p
G9F1zxczncM1BsKQ579iYA0Ek+pWiFlxK2lziO
86oIkYr1csnHmcYKjrO/9GI1SFSAm6v2bBnXMh6wzcp10A+6TAU4Bp9c7WNq4K1U0x
Are you sure you want to add this public key (yes/no)? yes
x950(config)#
```

Caution When you respond to the question:



“Are you sure you want to add this public key (yes/no)?“

You must type in the **complete word**. For example **‘yes’** (not just the letter ‘y’).

This process implements a trusted relationship between the server and the AMF Plus master. It is therefore the users’ responsibility to verify that the public key is being imported from the correct host, and not a substitute host using the same host name or IP address.

Exporting a public user key from the AMF Plus backup node for the remote file server

1. Generate a public key for root on the AlliedWare Plus Switch:

```
x930a(config)#crypto key generate userkey root rsa
Generating user key for root (1024 bits rsa)
This may take a while. Please wait...
x930a(config)#exit
x930a#
```

2. Create a file containing the public key:

```
x930a#show crypto key userkey root rsa > root-rsa.pub
```

3. Confirm that the file exists and looks right.

```
x930a#show file root-rsa.pub
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQgQDFRNJeSqqMmMesU3wF6RM/
F2rieTwU5ILWzstlizVF
Sz1qLPYb5VaRL8s7pMTElP5aAjIs2dS041a/
0WseVt1qYSUlw9kJzcyR26yy+bUzt0L5DspQq+uZkpmv
KiTE3BQFvw6RospJh+36fT301ywUPfJiRVyigdyfbch9YG6s4Q==
```

4. At this stage, transfer the key via TFTP or via removable media. There is no special requirement to keep this key secure.

```
x930a#copy root-rsa.pub tftp://10.37.165.33/
Enter destination file name[root-rsa.pub]:
Copying...
Successful operation
x930a#
```

5. Now on the external server, append the public key to your authorized_keys file.

```
amf-admin@atmf23:~$ cat /tftpboot/root-rsa.pub >> ~/.ssh/authorized_keys
```

Linux SSH daemon directory permissions

The Linux SSH daemon file and directory permissions need to be correctly configured.

To set the file permissions on the ssh and authorized_keys directories, use the following commands:

```
amf-admin@tb165 ~ $ chmod 700 ~/.ssh
```

```
amf-admin@tb165 ~ $ chmod 600 ~/.ssh/authorized_keys
```

File/user permissions on the remote file server

The default users/groups and permissions bits should be set to 'user' in the folder that is the root of the backup folders. These defaults should be the same for all directories above the one supplied. The easiest way to achieve this is to log in as the supplied user to the server, make the directories required, then not to touch the user/group or permissions on the folders/files created.

Security on the remote file server

The user will need to create a directory on the remote file server to receive the backed-up files and directories. The user should limit the permissions on this directory so as to keep these files as secure as possible. Note that the default permissions will allow group access to this directory:

```
amf-admin@atmf23:~$ mkdir network_backups
amf-admin@atmf23:~$ ls -l
total 56
drwxr-xr-x  4 amf-admin  4096 Dec 19 14:33 atmf
drwxr-xr-x  2 amf-admin  4096 May  1 17:04 network_backups
drwxr-xr-x 17 amf-admin 32768 May  1 10:23 release_tarballs
drwxr-xr-x 17 amf-admin 12288 Apr 29 14:43 scripts
drwxr-xr-x  4 amf-admin  4096 Apr 30 13:24 temp
```

- Typically this will be access for the user only.

```
amf-admin@atmf23:~$ chmod 700 network_backups/
amf-admin@atmf23:~$ ls -l
total 56
drwxr-xr-x  4 samh      stdept  4096 Dec 19 14:33 atmf
drwx-----  2 samh      stdept  4096 May  1 17:04 network_backups
drwxr-xr-x 17 samh      stdept 32768 May  1 10:23 release_tarballs
drwxr-xr-x 17 samh      stdept 12288 Apr 29 14:43 scripts
drwxr-xr-x  4 samh      stdept  4096 Apr 30 13:24 temp
amf-admin@atmf23:~$
```

- Access to backups is now restricted to the user **samh**.

Configuring the backup server on the AMF Plus controller/master

Once the keys have been exchanged, the AMF Plus controller or master can be configured to use the server for backups.

The command to carry out this configuration is:

```
atmf backup server id {1|2} <hostlocation> username <username>
[path <path>|port <1-65535>]
```

Where:

id just provides a method for referring to the two different servers if two backup servers have been configured.

hostlocation is the IPv4 or IPv6 address of the server.

username is the name that the controller or master will use when it connects to the server to write or read backups.

path is the directory path to the folder where the server stores the backups. By default, this is the home directory of the user specified by the username parameter.

port is the TCP port used for connections to the server. By default, port 22 is used.

For the example being worked through here, the command would be:

```
atmf backup server id 1 10.37.165.65 username samh path /home/samh/  
network_backups
```

Each AMF Plus controller or master supports a maximum of two remote file servers. The remote backup file servers are mounted on the controller or master's file system using SSH and appear as folders.

Configuring a backup to a remote server

After you have configured the servers you can check the backup media, location, log details, and server status using the **show atmf backup** command. You can also manually synchronize the contents of an active server and other configured servers if required.

The following steps describe how to set up two backup servers:

1. Use the command **atmf backup server** for backup server 1.

This command configures a remote file server(s) as the destination for AMF Plus backups.

Configuration of a remote server will switch the backup process to using remote server functionality and disable any further backup to removable media. Use the **no** variant of this command to remove the destination servers and revert to backing up to removable media.

Note that if no servers are configured, the backup will go to removable media. If no servers are configured and no removable media exists, no backup will occur. As described below in the section "[Multiple backup destinations](#)" on page 209, it is actually possible to override the disabling of the backup to removable media when remote servers have been configured, and thereby backup to both removable media and remote file servers.

2. Repeat step (1) for backup server 2.

You should now have two file servers configured to backup your network.

As described below in the section "[Multiple backup destinations](#)" on page 209, it is actually possible to override the disabling of the backup to removable media when remote servers have been configured, and thereby backup to both removable media and remote file servers.

3. Use the command **atmf backup now** to force a manual backup of your network.

This step is optional. Alternatively you could wait until the next scheduled back occurs.

4. Use the command **show atmf backup**.

If you forced a manual backup, you will probably want to display the location and state of each configured file server. The display from this command also shows diagnostic results that test connectivity to each server by using the optional **server-status** parameter.

In the example shown below, output from the **show atmf backup** command displays the configuration of two remote backup file servers.

Output 29: Output from the **show atmf backup** command - configuration of two remote backup file servers

```
x950#show atmf backup
Scheduled Backup ..... Enabled
  Schedule ..... 24 per day starting at 14:25
  Next Backup Time .... 14 Feb 2023 11:25
Backup Bandwidth ..... Unlimited
Backup Media ..... FILE SERVER (Total 503837.5MB, Free 186818.0MB)
Server Config .....
  Synchronization ..... Synchronized
  Last Run ..... 14 Feb 2023 11:09:50
  1 ..... Configured (Mounted)
  Host ..... 10.36.150.54
  Username ..... user_1
  Path ..... temp/x950_1
  Port ..... -
  * 2 ..... Configured (Mounted, Primary)
  Host ..... tbl65.test.com
  Username ..... user_2
  Path ..... temp/x950_2
  Port ..... -
Current Action ..... Idle
  Started ..... -
  Current Node ..... -
```

Node Name	Date	Time	In ATMF	On Media	Status
Synchronization	Date	Time	From Id	To Id	Status
x230a	14 Feb 2023	11:09:37	Yes	Yes	Good
	14 Feb 2023	11:09:46	2	1	Good
x930a	14 Feb 2023	11:09:17	Yes	Yes	Good
	14 Feb 2023	11:09:19	2	1	Good
x930b	14 Feb 2023	11:09:49	Yes	Yes	Good
	14 Feb 2023	11:09:49	2	1	Good
x930c	14 Feb 2023	11:09:20	Yes	Yes	Good
	14 Feb 2023	11:09:20	2	1	Good
x930d	14 Feb 2023	11:09:19	Yes	Yes	Good
	14 Feb 2023	11:09:19	2	1	Good
x950	14 Feb 2023	11:09:49	Yes	Yes	Good
	14 Feb 2023	11:09:50	2	1	Good
x950stk	14 Feb 2023	11:09:47	Yes	Yes	Good
	14 Feb 2023	11:09:48	Yes	Yes	Good

You can use the **show atmf backup** command with the parameter **server-status** to display the results of the diagnostics that test connectivity to each server:

Output 30: **show atmf backup** command showing diagnostic test results from each server

```
Master1#sh atmf backup server-status
  Id  Last Check  State
-----
  1   186 s      File server ready
  2    1 s      SSH no route to host
```

Multiple backup destinations

For resilience, AMF Plus enables a master or controller node to store backups in multiple locations. The backups can go to up to two remote file servers, and to removable media at the same time.

Backing up to two remote file servers

When a master or controller node is configured with two remote file servers, it will store backups on both file servers. However, one of the file servers will automatically be assigned the role of primary server. When a backup is performed, the following sequence of events occurs:

- The master or controller backs up required data to the primary remote server.
- When the backup is complete, the master or controller synchronizes the backed-up files from the primary server over to the other (the backup) server.

The identity of the primary server can be seen from the output of the command:

```
show atmf backup
```

One of the servers will be labeled Primary, as shown below:

```
* 2 ..... Configured (Mounted, Primary)
```

At any time, you can force the servers to synchronize by using the command:

```
atmf backup synchronize
```

For example, if the backup server has been off line for a while, and during that time, a backup has occurred, then the primary server will have more recent versions of the backed-up files than the backup server. So, when that backup server is brought back on line, its data can be brought up to date by running this manual synchronization process.

Backing up to remote file server(s) and removable media

If a master or controller has been configured with one or two remote file servers for backups then the default behavior is no longer to send backups to removable media.

However, if removable media is present in the unit, and you want to send backups to this media in addition to the remote file server(s), then this functionality can be implemented by using the command:

```
atmf backup redundancy enable
```

When this has been enabled, the rules are:

- If remote file servers are configured and accessible, then the primary backup destination will always be one of the remote file servers.
- When a backup to the primary remote server is complete, the backup is first synchronized to the other remote file server (if a second remote server has been configured, and is accessible) and then to the removable media.
- The remote file server(s)—if available—is always the preferred location for retrieving backups for a recovery. The removable media will only be used for delivering files for a recovery if no remote file servers are accessible.
- The command **atmf backup synchronize** will synchronize the backed-up files between all backup destinations—the remote file server(s) and the removable media.
- If the removable media has been absent for a while, and a new item of removable media is installed into the controller/master node, the backed-up files on the remote file server(s) will not be automatically synchronized over to the removable media. This synchronization must be initiated manually, using the command **atmf backup synchronize**.

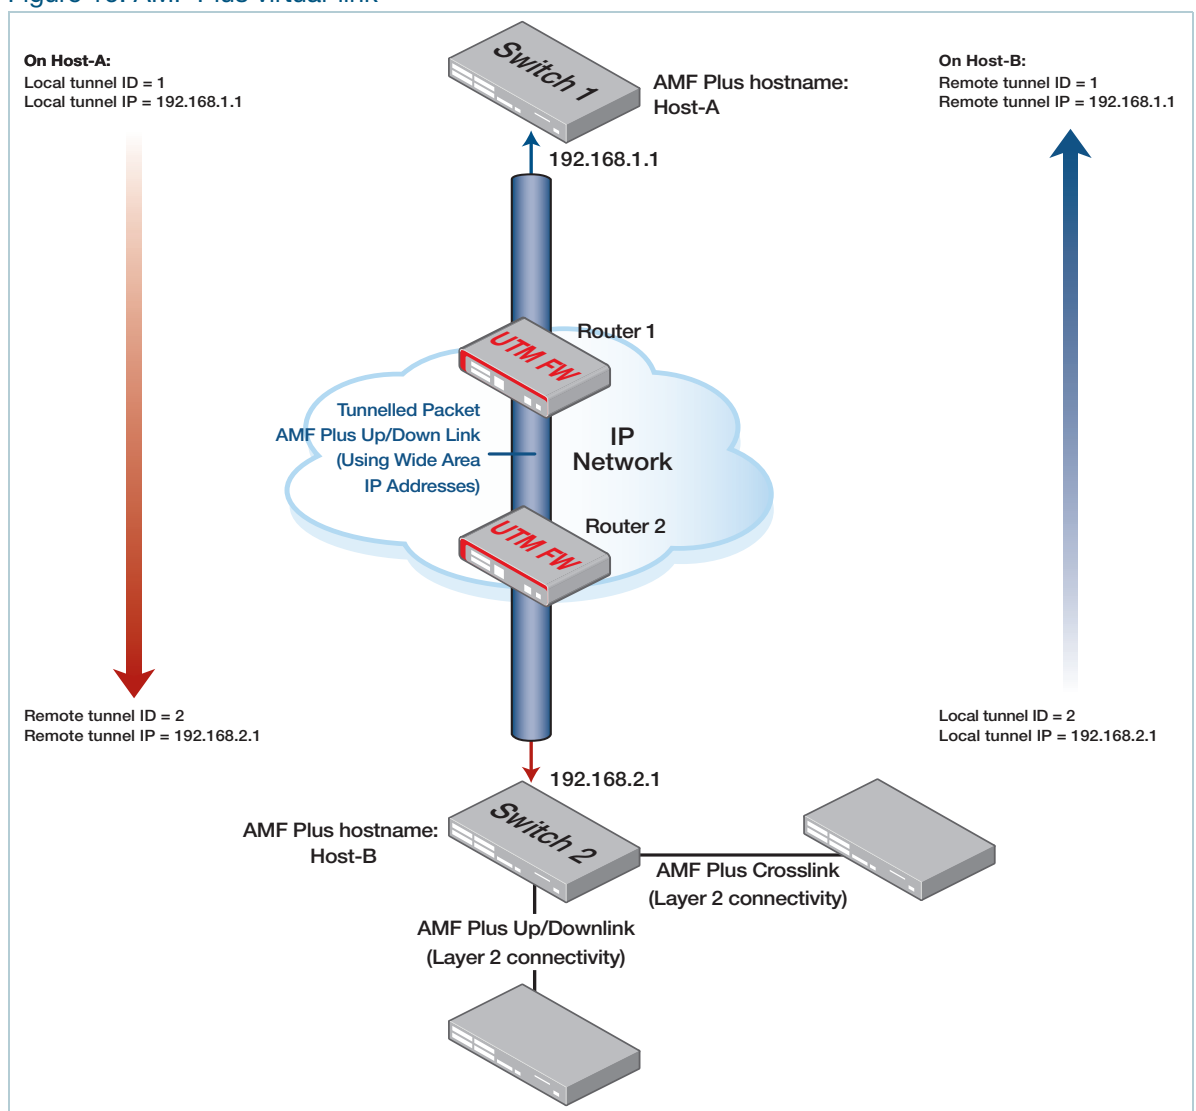
AMF Plus Tunneling (Virtual-links)

AMF Plus tunneling enables you to extend your local uplinks and downlinks across a wide area network. The tunneled data is wrapped in a Layer 3 IP packet for transmission across a wide area IP network. A simple AMF Plus tunnel is shown in Figure 16 below. Switches 1 and 2 encapsulate the Layer 2 AMF Plus uplink and downlink data and wrap this inside a Layer 3 IP packet to enable it to traverse an IP Network. Routers 1 and 2 (and any other routers within the cloud) route traffic conventionally, reading the IP addresses of the tunneled packets and forwarding them to their destination.

The two ends of the tunnel must have IP connectivity to each other. If the tunnel is configured to connect a head office and branch office over the Internet, typically this would use some type of managed WAN service such as a site-to-site VPN. Tunnels are only supported using IPv4.

Once connected through the tunnel, the remote AMF Plus members will have the same AMF Plus capabilities as a directly connected AMF Plus member.

Figure 16: AMF Plus virtual-link



You need to create matching local and remote configurations on the devices at each end of the tunnel. Each local tunnel ID must be unique to the device on which it is configured. The local and remote tunnel ID numbers do not have to be the same value. For example, a tunnel with ID 10 on the switch at one end of the virtual-link could connect to a tunnel with ID 40 on the switch at the other end of the link.

The tunneled link may pass through external (non AMF Plus-capable) routers in order to provide wide area network connectivity. These devices forward the traffic conventionally. The protocol tunneling function is accomplished by the AMF Plus nodes at each end of the virtual-link.

Different AlliedWare Plus devices support different numbers of virtual links, as shown in the following table:

Table 5: Supported number of virtual-links

DEVICE	VIRTUAL-LINK LIMIT
AMF Plus Cloud AR4000S Cloud 10G UTM Firewall AR4050S AR3050S AR2050V AR2010V	1000
AR1050V	100
SBx81CFC960 SBx908 GEN2	60
All other x-series and IE-series devices that support AMF Plus	32

Effect of virtual links on zero-touch recovery

Note that zero touch device replacement of a remote device that terminates the tunnel cannot be achieved by delivering backed-up files from an AMF Plus master that is located in the vicinity of the local end of the tunnel. This is because the master cannot deliver the files to the replacement unit until the link is up, but the link cannot form until the replacement unit has its config files. Another mechanism is used for backing up the configs on virtual tunnel end-points. This is described in the section "[Backups by master nodes](#)" on page 191.

You can recover isolated nodes at the far end of an AMF Plus virtual-link tunnel by using the method described in the section "[Recovering and Provisioning Isolated Nodes](#)" on page 160.

Secure virtual-links

You can create secure AMF Plus virtual-links by encapsulating the L2TPv3 frames of the virtual-link with IPsec. Secure virtual-links make it possible for your AMF Plus data to securely traverse a wide area IP network without the need to create a secure VPN tunnel.

IPsec provides the following security services to the AMF Plus virtual-link:

- data origin authentication, i.e. it identifies who sent the data.
- confidentiality (encryption), this ensures the data cannot be intercepted and read.
- integrity (authentication) - ensures the data has not been changed en-route.
- replay protects - by detecting packets received more than once, this helps protect against denial of service attacks.

Secure virtual-links limitations

The following limitations need to be considered when creating secure virtual-links.

Total number of downstream AMF Plus members:

Switch devices support a maximum of 20 downstream AMF Plus nodes when using a secure virtual-link as an uplink.

Secure virtual-links behind NAT:

When there are two or more AMF Plus members behind a shared NAT device, only one of the members will be able to use secure virtual-links.

AMF Plus Multi-tenant environment:

An AMF Plus Multi-tenant environment supports a maximum cumulative total of 1200 secure virtual-links across all AMF Plus containers.

Supported devices

Not all AlliedWare Plus devices support secure virtual links, as shown in the following table:

Table 6: Supported number of secure virtual-links

DEVICE	SECURE VIRTUAL-LINK LIMIT
AMF Cloud	300
10G UTM Firewall AR4050S AR3050S AR2050V AR2010V	60
x220 Series x230 and x230L Series x330 Series	2

Example Configuration

Step 1: Configure an AMF Plus virtual-link as normal.

```
Host-A(config)# atmf virtual-link id 1 ip 192.168.1.1 remote-id 2 remote-ip
192.168.2.1
```

Step 2: Apply protection to the virtual-link

```
Host-A(config)# atmf virtual-link id 1 protection ipsec key secure-
password
```

Step 3: Repeat these steps on the other side of the link

```
Host-B(config)# atmf virtual-link id 2 ip 192.168.2.1 remote-id 1 remote-ip
192.168.1.1
```

```
Host-B(config)# atmf virtual-link id 2 protection ipsec key secure-
password
```

Configuring a virtual-link

You can configure virtual links when your link end-points include static or dynamic IP addresses. This section shows examples for each of:

"1. Local and remote addresses of the link are statically configured" on page 215

"2. Remote address of the link is identified by DNS" on page 216

"3. Local address of the link is identified by interface" on page 218

"4. Remote address of the link is identified via IPsec" on page 219

These options are not mutually exclusive. If you use an interface to identify the local address of the link, you can use any of the available methods to identify the remote address (a static IP address, a hostname, or a dynamic IP address via IPsec).

Caution On an IP interface that is carrying AMF Plus virtual-link traffic, do not set the MTU (Maximum Transmission Unit) to less than its default value of 1500 bytes.

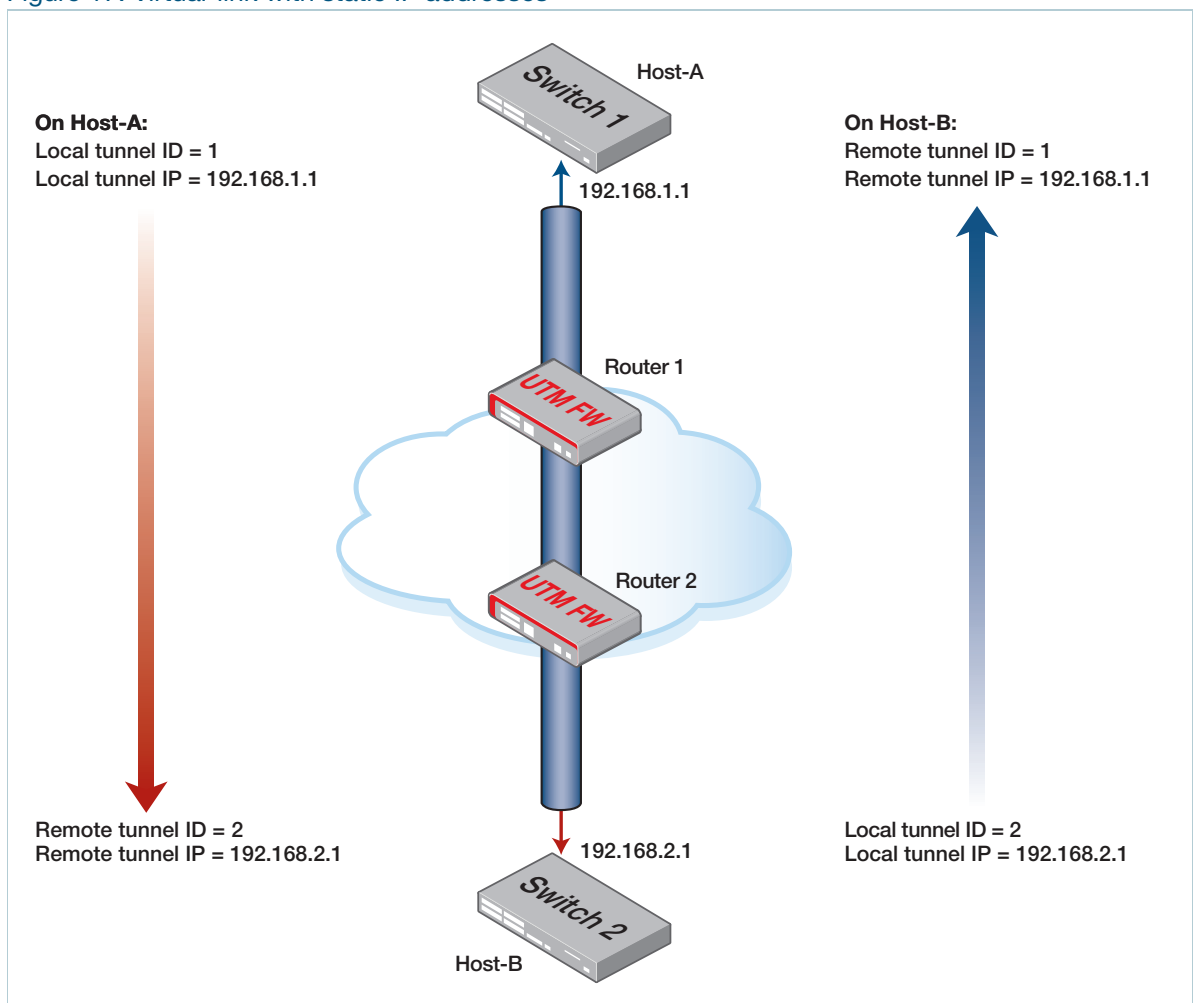


1. Local and remote addresses of the link are statically configured

This example creates a virtual-link between Host-A and Host-B where both hosts have known static IP addresses. Configuration involves specifying the virtual-link parameters in the table below.

PARAMETER	ON HOST-A	ON HOST-B
id	A number to identify Host-A's end of the link	A number to identify Host-B's end of the link
ip	Host-A's static IP address	Host-B's static IP address
remote-id	A number to identify Host-B's end of the link	A number to identify Host-A's end of the link
remote-ip	Host-B's static IP address	Host-A's static address

Figure 17: Virtual-link with static IP addresses



Use the following commands to create the tunnel above.

On Host-A Host-A(config)# atmf virtual-link id 1 ip 192.168.1.1 remote-id 2 remote-ip 192.168.2.1

On Host-B Host-B(config)# atmf virtual-link id 2 ip 192.168.2.1 remote-id 1 remote-ip 192.168.1.1

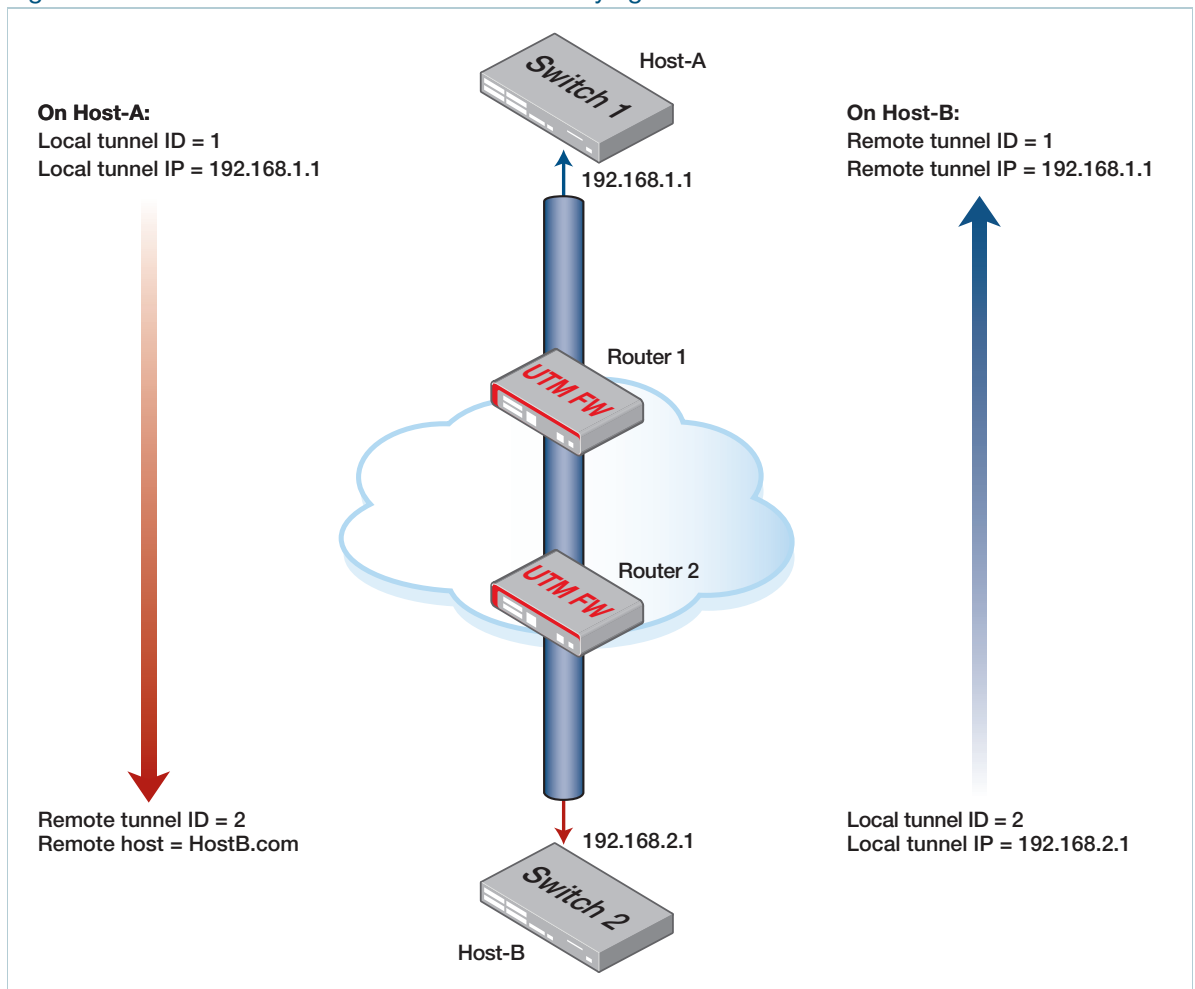
2. Remote address of the link is identified by DNS

This example creates a virtual-link between Host-A and Host-B where the IP address at Host-B's end of the link is unknown but can be resolved via DNS. To identify Host-B's end of the link, use a domain name as remote host on Host-A. That domain name can be the domain name of Host-B or of a NAT device in front of Host-B.

Configuration involves specifying the virtual-link parameters in the table below. You also need to make sure that Host-A can use DNS to resolve the host name.

PARAMETER	ON HOST-A	ON HOST-B
id	A number to identify Host-A's end of the link	A number to identify Host-B's end of the link
ip	Host-A's static IP address	Host-B's static IP address
remote-id	A number to identify Host-B's end of the link	A number to identify Host-A's end of the link
remote-ip		Host-A's static address
remote-host	The domain name of Host-B's end of the link	

Figure 18: Virtual-link with a domain name identifying one end of the link



Use the following commands to create the tunnel above.

On Host-A Host-A(config)# ip name-server 10.10.10.2

Host-A(config)# ip domain-lookup

Host-A(config)# atmf virtual-link id 1 ip 192.168.1.1 remote-id 2
remote-host HostB.com

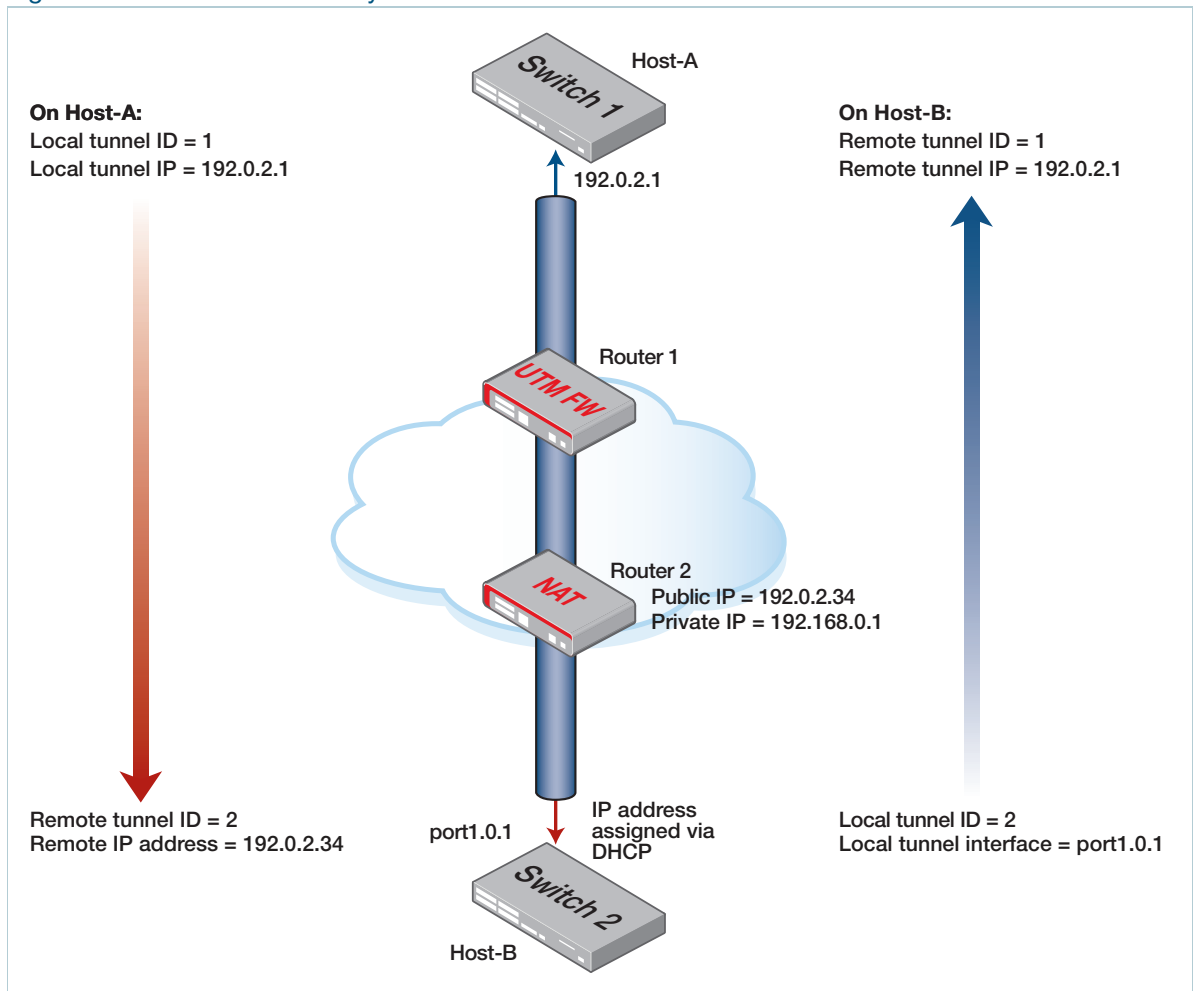
On Host-B Host-B(config)# atmf virtual-link id 2 ip 192.168.2.1 remote-id 1 remote-ip
192.168.1.1

3. Local address of the link is identified by interface

This example creates a virtual-link between Host-A and Host-B where the IP address of Host-B is dynamically-assigned via DHCP. In this example, Host-B is behind a firewall that has a static public IP address. Configuration involves specifying the virtual-link parameters in the table below.

PARAMETER	ON HOST-A	ON HOST-B
id	A number to identify Host-A's end of the link	A number to identify Host-B's end of the link
ip	Host-A's public IP address	
interface		The interface that connects Host-B to the WAN
remote-id	A number to identify Host-B's end of the link	A number to identify Host-A's end of the link
remote-ip	The public IP address of the NAT gateway	Host-A's public IP address

Figure 19: Virtual-link with a dynamic IP address on one end of the link



Use the following commands to create the tunnel above.

On Host-A Host-A(config)# atm virtual-link id 1 ip 192.0.2.1 remote-id 2 remote-ip 192.0.2.34

On Host-B Host-B(config)# atm virtual-link id 2 interface port1.0.1 remote-id 1 remote-ip 192.0.2.1

4. Remote address of the link is identified via IPsec

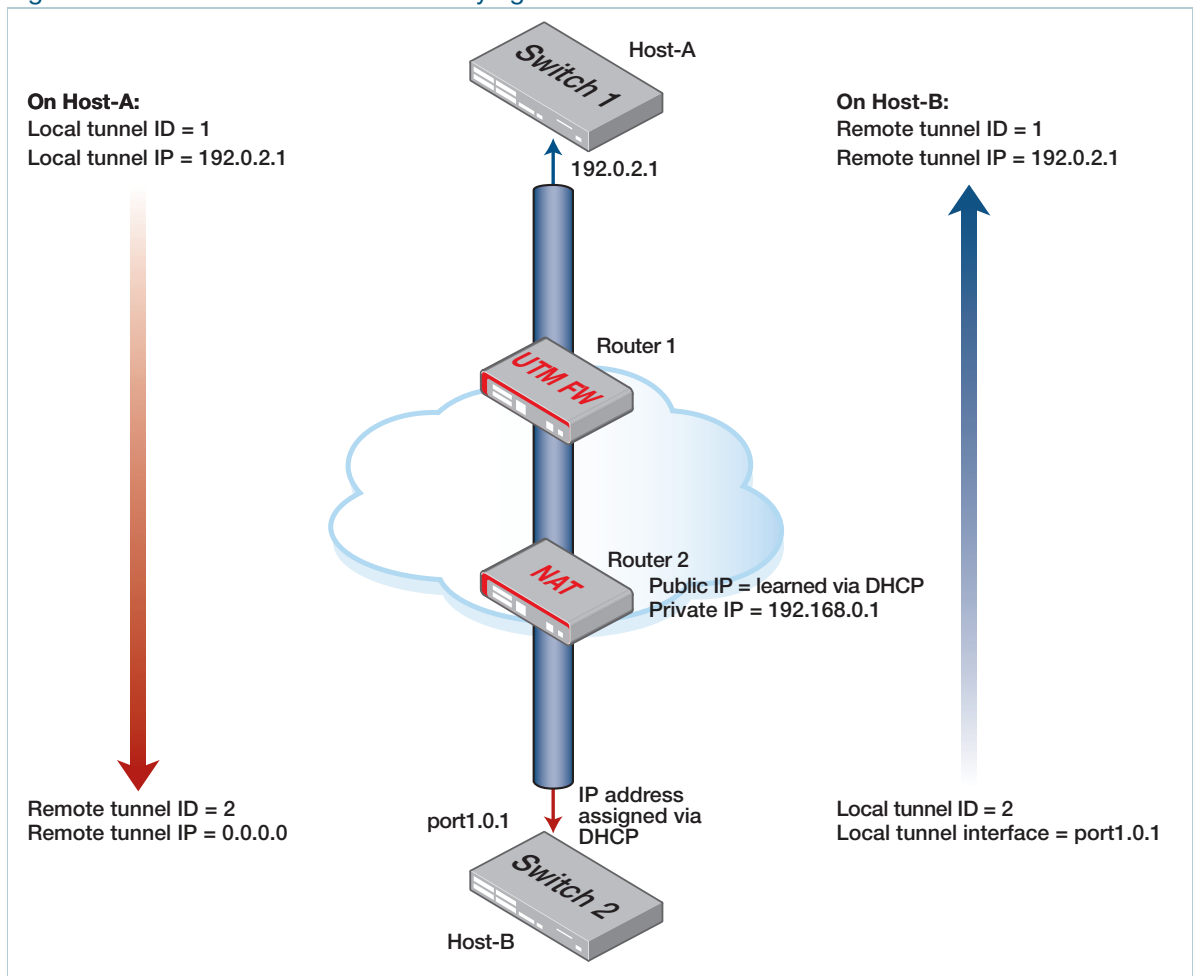
This example creates a virtual-link between Host-A and Host-B where the IP address of Host-B is dynamically-assigned via DHCP. In this example, Host-B is behind a firewall that also has a dynamically-assigned public IP address, so Host-A cannot use the firewall's address to identify the Host-B end of the tunnel.

Configuration involves specifying the virtual-link parameters in the table below. You also need to make the link a secure virtual-link - see "[Secure virtual-links](#)" on page 213.

PARAMETER	ON HOST-A	ON HOST-B
id	A number to identify Host-A's end of the link	A number to identify Host-B's end of the link
ip	Host-A's public IP address	
interface		The interface that connects Host-B to the WAN
remote-id	A number to identify Host-B's end of the link	A number to identify Host-A's end of the link
remote-ip	0.0.0.0	Host-A's public IP address

Note that a dynamic remote address cannot be used on both ends of a secure AMF Plus virtual-link, so you cannot configure a remote IP address of 0.0.0.0 on both Host-A and Host-B.

Figure 20: Virtual-link with IPsec identifying one end of the link



Use the following commands to create the tunnel above.

On Host-A Host-A(config)# atmf virtual-link id 1 ip 192.0.2.1 remote-id 2 remote-ip 0.0.0.0

Host-A(config)# atmf virtual-link id 1 protection ipsec key secret_key

On Host-B Host-B(config)# atmf virtual-link id 2 interface port1.0.1 remote-id 1 remote-ip 192.0.2.1

Host-A(config)# atmf virtual-link id 2 protection ipsec key secret_key

Prioritizing the tunneled traffic

The traffic that is carried in the tunnel includes AMF Plus management information. It is therefore critical that this information is conveyed to the portion of the network that lies at the remote end of the tunnel. Accordingly, the traffic that passes through the tunnel should be given a high QoS priority, so that it will not be lost if other user traffic causes congestion.

There are two key places in the path where we need to consider prioritization of the tunnel traffic:

1. When the tunneled traffic passes through any devices that lie between the tunnel endpoint and the wide area network.
2. When the tunneled traffic arrives at the endpoint device, and needs to go up to its CPU.

Prioritizing tunnel traffic passing through other switches

As the tunnel passes through other switches that lie between the tunnel endpoint and the WAN, the traffic destined for the tunnel will be hardware forwarded by these switches. The tasks that are required to increase the priority of that hardware forwarding are:

- recognise the tunnel destined traffic.
- assign this traffic to a high-priority queue on the egress port.
- insert a priority tag in the VLAN header of the tunnel destined packets.
- insert a priority tag into the IP header of the tunnel destined packets.

Step 1: Recognizing the tunnel traffic

In the illustration in [Figure 16 on page 211](#), the tunnel destined traffic consists of IP packets between 192.168.1.1 and 192.168.2.1.

Consider a switch located between Switch1 and Router1, with its port1.0.1 interface connected to Router1 and its port1.0.2 interface connected to Switch1.

This switch needs to prioritize the tunnel destined traffic in both directions - from Switch1 to Switch2 and from Switch2 to Switch1.

Therefore, it needs to look for traffic with:

- Source IP 192.168.1.1 and Destination IP 192.168.2.1
- Source IP 192.168.2.1 and Destination IP 192.168.1.1

These two sets of traffic can be classified by matching with the following ACLs:

```
access-list hardware vlinkdown
    permit ip 192.168.1.1/32 ip 192.168.2.1/32
access-list hardware vlinkup
    permit ip 192.168.2.1/32 ip 192.168.1.1/32
```

Step 2: Assigning the traffic to a high-priority queue on the egress port

There are a few configuration steps required to achieve this. A classmap needs to be created for each tunnel direction, that matches traffic for that tunnel direction:

```
class-map vlinkup match access-group vlinkup
class-map vlinkdown match access-group vlinkdown
```

A policy map needs to be created for each tunnel direction, that takes the traffic matching the appropriate classmap, and puts that traffic into a high-priority egress queue (in this case, queue 5):

```
policy-map vlinkup
    class default
    class vlinkup
        remark new-cos 5 internal

policy-map vlinkdown
    class default
    class vlinkdown
        remark new-cos 5 internal
```

The appropriate policy map needs to be applied to the port on which the tunnel traffic in the given direction arrives:

```
int port1.0.1
    service-policy input vlinkup
int port1.0.2
    service-policy input vlinkdown
```

Step 3: Insert a priority tag in the VLAN header of tunnel destined packets

This optional action might be useful where additional switches exist between the AMF Plus tunnel endpoint and the WAN and that these switches only support CoS-based prioritization. If the CoS value in the tunnel packets is set to a high value, then typically, these additional switches will be able to give that traffic high priority.

The configuration required to mark the CoS values in these frames is simply to change the lines "remark new-cos 5 internal" in the configuration above to "remark new-cos 5 both". In this way, the frames are sent to a high-priority egress queue on the switch with the QoS config **and** the CoS value in the frames' VLAN header is marked with the value 5.

Step 4: Insert a priority tag in the IP header of tunnel destined packets

Again, this is an optional action. If the devices in the wide-area network are configured to prioritize packets based on the DSCP value in their IP headers, then configuring a high-priority value in DSCP fields of the tunnel packets could be worthwhile. In general, WAN devices are not configured to prioritize based on DSCP, and certainly it cannot be expected on the Internet. But, in the case of a private WAN that is configured with a DSCP prioritization scheme, then the following configuration will be of value:

Add the line:

```
remark-map bandwidth-class green to new-dscp 46
```

to the class configuration in each policy map.

Prioritizing the tunnel traffic to the CPU of the receiving endpoint

When the tunneled traffic arrives at the end-point AMF Plus member, it needs to go up to the CPU of that device to be processed. If a high rate of traffic is arriving at that device, then the link up to the CPU may be oversubscribed, and the tunnel traffic will need to be prioritized to make sure it is not dropped due to the congestion.

The configuration for prioritizing the tunnel traffic up to the CPU is very similar to that for prioritizing the traffic being forwarded in the tunnel. The main difference is that only one direction of traffic (namely, traffic **to** the end-point device) needs to be prioritized.

For example, on Switch1, it is necessary only to match the traffic coming towards it from Switch2:

```
access-list hardware incomingTunnel
    permit ip 192.168.2.1/32 ip 192.168.1.1/32
```

From there, the rest of the configuration is essentially the same as the through-traffic prioritization case:

```
policy-map incomingTunnel
    classmap incomingTunnel
        match access-group incomingTunnel
    class default
```

```
class incomingTunnel
    remark new-cos 5 internal
int port1.0.1
    service-policy input incomingTunnel
```

Virtual cross-links

It is also possible to create virtual cross-links using the **atmf virtual-crosslink** command. These are only supported in a topology making use of an AMF Plus Cloud instance. In this situation a single virtual cross-link can be created to enable a master residing on AMF Plus Cloud to share the AMF Plus master role with an AMF Plus master running on a physical device.

Note: Creating virtual cross-links between container masters (in a multi-tenant AMF Plus Cloud installation) and physical masters is not supported.

Firmware Auto Upgrade

There are two recommended methods for performing a firmware upgrade on an AMF Plus area:

1. AMF Plus reboot-rolling upgrade
2. AMF Plus distribute firmware upgrade

The reboot-rolling firmware upgrade feature allows nodes within an AMF Plus area to be rebooted and upgraded in a rolling sequence in order to minimize downtime and reduce the management overhead. The distribute firmware upgrade feature enables nodes within an AMF Plus area to be upgraded and then rebooted at a later time.

Note: The **atmf reboot-rolling** command can also be used to reboot a set of nodes without upgrading the firmware. Specifying the **force** parameter on the **atmf reboot-rolling** command will result in the upgrade continuing even if an upgrade or reboot fails for a particular node.

How many nodes to update at once

We recommend upgrading a maximum of 42 nodes at once, because both reboot-rolling and distribute firmware upgrades generate large amounts of AMF Plus traffic.

Advantages of reboot-rolling upgrade

The reboot-rolling upgrade offers the following advantages:

- A fully automated upgrade process with report on completion (assuming the initiating node is excluded from the rolling reboot upgrade process).
- The AMF Plus reboot-rolling algorithm automatically selects the appropriate order in which to upgrade the nodes (edge to core).
- The most up-to-date appropriate release is automatically selected from the specified location for each platform (removable media only).
- Multiple nodes can be upgraded and rebooted with a single command.
- AMF Plus checks that each node completes the upgrade successfully and that the node re-joins the network before proceeding to upgrade the next node.
- During the upgrade AMF Plus provides rolling updates for the number of nodes that have been upgraded and rebooted, and the total elapsed time since the upgrade was initiated.
- A report is generated at the end of the reboot-rolling upgrade. (If the initiating node is excluded from the reboot-rolling upgrade working-set).

Disadvantages of reboot-rolling upgrade

The reboot-rolling upgrade does however suffer from the following disadvantages:

- The upgrade operation is not separated from the reboot operation.
- Since only one node is upgraded and rebooted at a time, it could take a long time to upgrade a large network.

Advantages of distribute firmware upgrade

- The most up-to-date appropriate release is automatically selected from the specified location for each platform.
- The upgrade operation is separate from the reboot operation, so a working-set of nodes can be loaded with new firmware, but rebooted at a later convenient time.
- Once the new firmware is distributed to a working-set of nodes, the nodes can be manually rebooted individually, in groups, or all at the same time.

Support for AMF Plus Network Upgrades

Allied Telesis strongly recommends that all nodes in an AMF Plus network run the same firmware version, and that running different versions should be limited to the periods between staged upgrades of the network. Where the network contains End of Life devices, these should run the latest available maintenance release.

However, in general, AMF Plus on newer firmware versions is compatible with AMF Plus on older firmware versions.

AMF Plus upgrade exceptions

There are a few exceptions to the general rule stated above due to known incompatibilities that exist between particular firmware versions. These exceptions are described in detail in the “Important Considerations Before Upgrading” section of the Release Note for each AlliedWare Plus version, which can be downloaded from the [Allied Telesis website](#).

Summary of the AMF Plus upgrade process

There are a number of steps to follow in order to upgrade nodes within an **area** using AMF Plus. These steps are summarized below:

1. Select a group of nodes to be upgraded.
2. Select the new release for each platform to be upgraded in the AMF Plus area.
3. Copy the releases to the location you intend to use for the upgrade.
4. Decide which AMF Plus upgrade method is most suitable for your network.
5. Check that each node to be upgraded (including all members of VCStacks) has enough space in Flash to hold the new release, and is set to boot from Flash.
6. Initiate the AMF Plus network upgrade using your selected method. The AMF Plus upgrade can be initiated from any node in the AMF Plus area. It does not have to be initiated from the master.

Detailed explanation of the AMF Plus upgrade process

This section expands on the steps listed above:

1. **Select a group of nodes that need to be upgraded.**

While we recommend that all nodes in an AMF Plus network are eventually upgraded to the same release, you may decide to perform the upgrade on a selected nodes first. If you elect to use the AMF Plus reboot-rolling upgrade method you may also wish to exclude the node that is controlling the upgrade in order for it to generate a report once the upgrade has completed.

2. **Select the new release for each platform to be upgraded in the AMF Plus area.**

Once you have decided to upgrade your AMF Plus network, you need to look at which product families your network contains and select an appropriate release for each platform. We strongly recommend that the same firmware version is applied to each platform where possible.

3. **Copy the releases to the location you intend to use for the upgrade.**

The supported media locations for AMF Plus network upgrades are:

- Flash
- Removable SD card
- Removable USB storage
- TFTP server
- SCP server
- HTTP server

If you specify either Flash or removable storage media, the newest compatible release for each node will be selected from the stored releases.

Note: Removable storage media must not contain more than 20 releases. More than this and the upgrade will fail and an error message will be generated. If the release file is to be copied from a remote storage location (e.g. via TFTP, HTTP, etc.), then the URL should specify the exact release file-name without using wild card characters.

If you are using removable storage media to store the new firmware releases, this must be installed in the node used to initiate the upgrade. If you are using a remote server as the source location for the new firmware releases, the server must be reachable from the node used to initiate the upgrade. The AMF Plus upgrade can be initiated from any node in the AMF Plus network.

4. **Decide which AMF Plus upgrade method is most suitable for your network.**

The AMF Plus upgrade method that is most suitable for your network will depend on your network topology and your individual requirements. If you prefer an upgrade method that reboots each node one at a time (so that the outage is limited to only what is connected to that node), and then checks that the node successfully re-joins the network before proceeding to upgrade and reboot the next node, then the AMF Plus reboot-rolling method may suit your network. If your network demands a shorter period of interruptions, then it may be more appropriate to use the AMF Plus distribute firmware method so that all nodes can be rebooted at the same time.

5. **Check that each node to be upgraded, including all members of VCStacks, has enough space in Flash to hold the new release and is configured to boot from Flash.**

Once you have decided which AMF Plus upgrade method is most appropriate for your network, the simplest way to determine whether the nodes to be upgraded have enough space in Flash to hold the new firmware version is to create a working-set of the nodes to be upgraded and use the chosen AMF Plus upgrade command on the working-set. This will be either **atmf reboot-rolling location** or **atmf distribute firmware location**, where **location** is one of the supported media locations.

Note: If a remote storage location (e.g. via TFTP, HTTP, etc.), is specified, then this step must be performed separately for each product family, as the exact release file name without using wild card characters must be entered.

These commands will perform several checks on the working-set of nodes to ensure that the upgrade will succeed, including checking that each node is set to boot from Flash, and that they have enough free space, before prompting you to confirm whether you wish to proceed with the upgrade.

Output 31: AMF Plus upgrade showing node with insufficient space

```

core1#atmf working-set group all
=====
core1, core2, distribution1, distribution2, edge1, edge2:
=====

Working set join

AMFname[6]#atmf distribute firmware usb:
Retrieving data from core1
Retrieving data from core2
Retrieving data from distribution1
Retrieving data from distribution2
Retrieving data from edge1
Retrieving data from edge2

ATMF Firmware Upgrade:

Node Name          New Release File          Status
-----
core1               SBx81CFC400-5.5.2-2.3.rel  Release ready
core2               SBx81CFC400-5.5.2-2.3.rel  Release ready
distribution1      x930-5.5.2-2.3.rel        Release ready
distribution2      x930-5.5.2-2.3.rel        Release ready
edge1               x230-5.5.2-2.3.rel        Insufficient space
edge2               x230-5.5.2-2.3.rel        Release ready

Continue upgrading releases ? (y/n):

```

If the AMF Plus upgrade command indicates that there are nodes with insufficient space you can cancel the upgrade operation and free up the necessary space in Flash on those nodes.

6. Initiate the AMF Plus network upgrade using the selected method.

Once you have confirmed that all nodes to be upgraded have enough free space in Flash, you can then initiate the AMF Plus upgrade using the chosen upgrade command. Each node will be updated to boot from the new release and the previous release will be set as the backup release file.

Note: If the AMF Plus distribute firmware method is being used, then the nodes must be rebooted manually to complete the upgrade.

Example 1 - Performing a reboot-rolling upgrade

To perform a reboot-rolling firmware upgrade on all nodes in the AMF Plus area, first select all nodes using the default **working-set group all** command:

```
SBx8100#atmf working-set group all
=====
SBx8100, x950-VCS1, x950-VCS2, x230_1, x230_2:
=====
Working set join
```

Next, using the **atmf reboot-rolling** command specify the path to the release files to which you wish to upgrade the nodes in the AMF Plus network. In this example, the release files are stored on the removable USB storage media installed in the node controlling the reboot-rolling firmware upgrade, in a directory called "rel". Note that because the node controlling the reboot-rolling firmware upgrade is included in the nodes to be upgraded, a message is displayed indicating that no summary will be available on completion.

```
AMFname[5]#atmf reboot-rolling usb:/rel/*.rel
Retrieving data from SBx8100
Retrieving data from x950-VCS2
Retrieving data from x230_1
Retrieving data from x230_2
Retrieving data from x950-VCS1

ATMF Rolling Reboot Nodes:

Node Name                Timeout
                          (Minutes)  New Release File          Status
-----
x230_2                    9           x230-5.5.2-2.3.rel       Release ready
x230_1                    6           x230-5.5.2-2.3.rel       Release ready
x950-VCS1                 9           x950-5.5.2-2.3.rel       Release ready
x950-VCS2                 9           x950-5.5.2-2.3.rel       Release ready
SBx8100                   11          SBx8100-5.5.2-2.3.rel    Release ready

% The controlling node (SBx8100) is included in the
rolling reboot and will be rebooted last.
No summary will be available on completion.
Continue upgrading releases ? (y/n):
=====
Copying Release      : x230-5.5.2-2.3.rel to x230_2
Updating Release     : x230-5.5.2-2.3.rel information on x230_2
=====
ATMF Rolling Reboot: Rebooting x230_2
=====
02:11:32 SBx8100 ATMF[1973]: x230_2 has left. 4 members in total.

% x230_2 has left the working-set
02:13:30 SBx8100 ATMF[1973]: x230_2 has joined. 5 members in total.
Reboot of x230_2 has completed
```

Although no summary report was generated in this particular example, you can refer to the progress messages output to the console to confirm that the upgrades were successful. You can also use the **atmf working-set group all** and the **show boot** commands to confirm the current boot image for each node in the AMF Plus area.

```

=====
Copying Release      : x230-5.5.2-2.3.rel to x230_1
Updating Release    : x230-5.5.2-2.3.rel information on x230_1
=====
ATMF Rolling Reboot: Rebooting x230_1
=====
02:14:13 SBx8100 ATMF[1973]: x230_1 has left. 4 members in total.

% x230_1 has left the working-set
02:15:53 SBx8100 ATMF[1973]: x230_1 has joined. 5 members in total.
Reboot of x230_1 has completed

=====

Copying Release      : x950-5.5.2-2.3.rel to x950-VCS1
Updating Release    : x950-5.5.2-2.3.rel information on x950-VCS1
=====
ATMF Rolling Reboot: Rebooting x950-VCS1
=====
02:19:02 SBx8100 ATMF[1973]: x230_1 has left. 4 members in total.
02:19:02 SBx8100 ATMF[1973]: x950-VCS1 has left. 3 members in total.

% x950-VCS1 has left the working-set
02:20:48 SBx8100 ATMF[1973]: x950-VCS1 has joined. 4 members in total.
Reboot of x950-VCS1 has completed
02:20:51 SBx8100 ATMF[1973]: x230_1 has joined. 5 members in total.
=====
Copying Release      : x950-5.5.2-2.3.rel.rel to x950-VCS2
Updating Release    : x950-5.5.2-2.3.rel information on x950-VCS2
=====
ATMF Rolling Reboot: Rebooting x950-VCS2
=====
02:21:54 SBx8100 ATMF[1973]: x230_2 has left. 4 members in total.
02:21:54 SBx8100 ATMF[1973]: x950-VCS2 has left. 3 members in total.

% x950-VCS2 has left the working-set
02:23:35 SBx8100 ATMF[1973]: x950-VCS2 has joined. 4 members in total.
Reboot of x950-VCS2 has completed
=====
Copying Release      : SBx8100-5.5.2-2.3.rel to SBx8100
02:23:39 SBx8100 ATMF[1973]: x230_2 has joined. 5 members in total.
Updating Release    : SBx8100-5.5.2-2.3.rel information on SBx8100
=====
ATMF Rolling Reboot: Rebooting SBx8100
=====
02:24:07 SBx8100 ATMF: reboot-rolling Rebooting SBx8100 at request of user
manager.

```

Note: Removable storage media must not contain more than 20 releases or the upgrade will not proceed and an error message will be generated. If the release file is to be copied from a remote storage location (e.g. via TFTP, HTTP, etc.), then the URL should specify the exact release filename without using wild card characters.

Example 2 - AMF Plus distribute firmware upgrade

To perform an AMF Plus distribute firmware upgrade:

1. First select the set of nodes you wish to upgrade using the **atmf working-set** command:

```
atmf working-set core2, group x230,x930
=====
core2, distribution1, distribution2, edge1, edge2:
=====

Working set join

AMFname[5]#
```

2. Then, using the **atmf distribute firmware** command, specify the path to the release files to use for the upgrade. In this example the release files are being stored on the removable USB storage media in the controlling node named Core1, which in this instance is excluded from the upgrade.

Note: As previously mentioned, removable storage media must not contain more than 20 releases or the upgrade will not proceed and an error message will be generated. If the release file is to be copied from a remote storage location (e.g. via TFTP, HTTP, etc.), then the URL should specify the exact release file name without using wild card characters.

The AMF Plus distribute firmware process will copy the appropriate firmware release to each node in the working-set, and then configure the nodes to boot from the new release. The previous boot release will, on each node, be automatically configured to be the backup boot release so that any node that fails to load the new release will automatically revert to the old release.

Output 32: Output from the `atmf distribute firmware` command

```

AMFname[5]#atmf distribute firmware usb:
Retrieving data from core2
Retrieving data from distribution1
Retrieving data from distribution2
Retrieving data from edge3
Retrieving data from edge1
Retrieving data from edge2

ATMF Firmware Upgrade:

Node Name                New Release File                Status
-----
edge2                    x230-5.5.2-2.3.rel             Release ready
edge1                    x230-5.5.2-2.3.rel             Release ready
distribution2            x930-5.5.2-2.3.rel             Release ready
distribution1            x930-5.5.2-2.3.rel             Release ready
core2                    SBx81CFC400-5.5.2-2.3.rel      Release ready
Continue upgrading releases ? (y/n): y
=====
Copying Release      : x230-5.5.2-2.3.rel to edge2
Updating Release    : x230-5.5.2-2.3.rel information on edge2
=====
Copying Release      : x230-5.5.2-2.3.rel to edge1
Updating Release    : x230-5.5.2-2.3.rel information on edge1
=====
Copying Release      : x930-5.5.2-2.3.rel to distribution2
Updating Release    : x930-5.5.2-2.3.rel information on distribution2
=====
Copying Release      : x930-5.5.2-2.3.rel to distribution1
Updating Release    : x930-5.5.2-2.3.rel information on distribution1
=====
Copying Release      : SBx81CFC400-5.5.2-2.3.rel to core2
Updating Release    : SBx81CFC400-5.5.2-2.3.rel information on core2
=====
New firmware will not take effect until nodes are rebooted.
=====
AMFname[6]#

```

Once the appropriate firmware release has been distributed to the selected nodes, the nodes can be rebooted at a convenient time, either individually or together to complete the upgrade process.

```

AMFname[6]#rel
% Warning: 6 nodes in total will be rebooted.
reboot system? (y/n): y

```


AMF Plus Security

AMF Plus has been designed to include a number of security features by default, with a focus on providing both security and convenience. However, you can enable extra optional features to maximize security. This chapter describes the default and optional security features.

Default security level

We recommend only using the default security level if all AMF Plus nodes are in a physically-isolated location, no AMF Plus virtual links go over insecure paths, and you have complete trust in all privileged users on all the AMF Plus nodes.

By default, AMF Plus includes the following security features:

- AMF Plus operates on a closed physical network and only exchanges AMF Plus messages across links that have been configured as AMF Plus links
- The AMF Plus protocol is not IP-based, which means that it does not listen to connection requests over the Internet. AMF Plus networks are not subject to remote access
- AMF Plus creates a virtual L2 (Layer 2) management network, which is secure because the device blocks packets from external networks from entering the AMF Plus L2 management network.

This means that attackers can only compromise an AMF Plus network if they have physical access to it (unless it includes virtual links over insecure paths – see "[Protecting AMF Plus virtual-links](#)" on [page 234](#)).

However, any privileged user on any AMF Plus node can configure any other AMF Plus node in the network.

There have been reports of large-scale attacks on third-party devices, which were exploited remotely through their auto-configuration solutions. AMF Plus auto-recovery and provisioning allow auto-configuration of new devices, but AMF Plus is not affected by the reported vulnerabilities. AMF Plus is not susceptible to attack by remote Internet hosts because the AMF Plus protocol, by design, is only available to link partners.

AMF Plus link management

You should only configure a link as an AMF Plus link if it specifically connects two AMF Plus nodes together. If you do this, attackers can only inject packets into an AMF Plus network if they replace one of the actual nodes of the network with another device. An attacker cannot simply connect an extra device into the network. You can prevent an attacker from replacing a node by keeping all AMF Plus nodes in a physically-secure location, and/or by using secure mode.

Increasing AMF Plus security

There are three other things you can do to increase AMF Plus security:

- configure AMF Plus “restricted login”
- protect any AMF Plus virtual-links that are over insecure paths
- enable AMF Plus “secure mode”.

The following sections summarize these options.

Restricted login

With restricted login, only privileged users on the AMF Plus master can use working-sets and automatic connections to other AMF Plus nodes. To maximize the benefit of restricted login, the AMF Plus master should be in a physically-secure location.

See ["AMF Plus restricted-login" on page 235](#) for configuration information.

Protecting AMF Plus virtual-links

AMF Plus virtual-links connect non-adjacent nodes by tunneling AMF Plus traffic over the devices in the path between the nodes. This means virtual-link security depends on the security of the devices between the nodes. If you are not sure that all those devices are secure, you need to protect the virtual-link – especially if it goes over the Internet.

You can protect such AMF Plus virtual-links by either:

- creating a VPN between the parts of the path that you consider insecure, or
- using IPsec to encapsulating the L2TPv3 frames of the virtual-link.

See ["AMF Plus Tunneling \(Virtual-links\)" on page 211](#) for configuration information.

Note: Secure mode encrypts AMF Plus packets. A VPN or IPsec encapsulation is, therefore, not necessary for protecting AMF Plus virtual-links if you are using secure mode. If the same path carries other traffic though, you do need to protect that traffic with a VPN.

Secure Mode

For the highest level of security within an AMF Plus network, you can enable AMF Plus “secure mode”. With secure mode enabled:

- AlliedWare Plus encrypts all AMF Plus packets and uses certificates to verify the identity of each node in the AMF Plus network
- Restricted login is automatically enabled and can't be disabled
- A node can only join the AMF Plus network if it has been authorized by a privileged user on the AMF Plus master. This makes it impossible for an attacker to connect a device without your knowledge.

See "[AMF Plus Secure Mode](#)" on page 236 for important details and configuration information.

AMF Plus restricted-login

By default, users who are logged into any node on an AMF Plus network are able to manage any other node by using either working-sets or an AMF Plus remote login. If the access provided by this feature is too wide, or contravenes network security restrictions, it can be limited by running the **atmf restricted-login** command, which changes the access so that:

- users who are logged into non-master nodes cannot execute any commands that involve working-sets, and
- from non-master nodes, users can use remote-login, but only to login to a user account that is valid on the remote device (via a statically configured account or RADIUS/TACACS+). Users are also required to enter the password for that user account.

The **atmf restricted-login** command will not be saved in the running configuration. It is a network property that can be enabled or disabled from any AMF Plus master. However, the status of restricted-login will be retained over a reboot.

Note that once you have run the command **atmf restricted-login**, certain other commands that utilize the AMF Plus working-set command will operate only on master nodes, such as the **atmf reboot-rolling** and **show atmf group members** commands.

If you have AMF Plus areas with more than 120 nodes, you must enable restricted-login.

AMF Plus Secure Mode

Introduction

The AMF Plus secure mode feature improves the security of the AMF Plus network by reducing the risk of your network being compromised through unauthorized access to the AMF Plus network. It achieves this by:

- Adding an authorization mechanism before allowing an AMF Plus member to join an AMF Plus network.
- Encrypting all AMF Plus packets sent between AMF Plus nodes.
- Additional logging, which enables network administrators to monitor attempts to gain unauthorized access to the AMF Plus network.

AMF Plus secure mode is optional and enabled from the command line interface. When running in secure mode the controllers and masters in the AMF Plus network form a group of certification authorities. A node may only join a secure AMF Plus network once authorized by a master or controller. When enabled, all devices in the AMF Plus network must be running in secure mode, unsecured devices will not be able to join a secure AMF Plus network.

Note: When an AMF Plus network is running in AMF Plus secure mode the **atmf restricted-login** feature is automatically enabled. This restricts the **atmf working-set** command to users that are logged in on an AMF Plus master. This feature cannot be disabled independently of secure mode. See "[AMF Plus restricted-login](#)" on page 235 for more information.

Licensing

AMF Plus secure mode does not require a special license. Note that AMF Plus secure mode cannot be enabled if the AMF Plus master only has an AMF Plus starter license.

Requirements

An AMF Plus area operating in secure mode is limited to 126 AMF Plus devices. This includes AMF masters and member nodes.

If an AMF Plus controller is running in secure mode then all nodes within all areas under the control of that controller must also be running in secure mode. If they are not running in secure mode then they will not be able to join the AMF Plus network. Running some AMF Plus nodes in secure mode and some in non-secure mode is not possible.

If secure mode is enabled on an AMF Plus network containing an AMF Plus controller, i.e. the AMF Plus network contains multiple areas, the controller and master of the local area must be on the same device. All AMF Plus masters must be authorized by the AMF Plus controller before they can

join the AMF Plus network. Additionally, for an area master to join the AMF Plus controller network, the controller must also be authorized by the area master.

Note: In secure mode the manual AMF Plus recovery feature is disabled, therefore backup and recovery of AMF Plus guest nodes is not supported.

Recommendation for multiple AMF Plus masters

A secure AMF Plus network with a single AMF Plus master is vulnerable to disruption if the master is lost. A replacement master would need to re-authorize all AMF Plus nodes in the network and the administrator would need to manually clear the existing certificates on all nodes. For this reason it is recommended that a secure AMF Plus network has two AMF Plus masters.

A viable alternative to running multiple masters is to have the AMF Plus master on a VCStack. The VCStack master will synchronize all security data to the other members of the stack. In the event of a failure of the stack master another stack member will take over with minimal disruption to the network.

All AMF Plus master nodes must reside in the same AMF Plus domain and are required to be directly connected using AMF Plus cross-links. In order to meet this requirement for masters running on AMF Plus Cloud, a single virtual cross-link can be created using the **atmf virtual-crosslink** command. This enables a master residing on AMF Plus Cloud to share the master role with a master running on a physical device.

Note: Nodes within a domain must be connected in either a chain or ring topology. This means that a maximum of **two** cross-links should be configured on any single node.

Virtual cross-links are not supported on AMF Plus container masters, therefore if multiple tenants on a single AMF Plus Cloud host are configured for secure mode, only a single AMF Plus master is supported per area. In this scenario it is expected that redundancy will be provided through the virtualization hypervisor.

Enabling AMF Plus secure mode

AMF Plus secure mode is enabled on an AMF Plus network by entering the **atmf secure-mode enable-all** command in privileged exec mode. This has the effect of running the **atmf secure-mode** command on each AMF Plus member and is the recommended way of enabling secure mode on new or existing AMF Plus networks. Run the following command on an AMF Plus master to enable secure mode on an entire AMF Plus network.

```
awplus# atmf secure-mode enable-all
```

Individual AMF Plus nodes can join an existing secure mode network by executing the **atmf secure-mode** command (in global configuration mode) on that node.

```
awplus(config)# atmf secure-mode
```

Authorizing nodes in secure mode

In order for the member nodes to join a secure mode AMF Plus network they must be authorized on one of the AMF Plus masters. Once nodes have been authorized on a master, the authorized state information is propagated through the network to all other nodes. Masters and controllers act as a certification authority and issue the AMF Plus node with a certificate. By default a node's certificate will last 180 days before expiring. This can be configured to anything between 1 and 365 days using the **atmf secure-mode certificate expiry** command. Alternatively the certificate can be set to never expire with the **infinite** parameter.

When a node's certificate is due to expire, and the node is still an active member of the AMF Plus network, the certificate will be automatically refreshed. The refreshed certificate will be valid for the same number of days as the original certificate. If a node is not currently an active member of the AMF Plus network when the certificate expires, it will not be automatically refreshed and the node will require reauthorization the next time it attempts to join the AMF Plus network.

AMF Plus backups when running in secure mode

When an AMF Plus network is running with secure mode enabled, AMF Plus auto-recovery will only function if the AMF Plus backup was taken while the network was already in secure mode. Similarly, auto-recovery will fail for a non-secure device from a backup taken in secure mode. When secure mode is enabled or disabled on an AMF Plus network, it is recommended you perform a manual backup as soon as possible rather than wait for the automated scheduled backup. This ensures that AMF Plus auto-recovery continues to work until before the scheduled backup runs.

Note: A message is displayed when an AMF Plus network changes to or from secure mode recommending that the administrator initiates an immediate AMF Plus backup.

Note: Guest node backups are disabled when secure mode is enabled.

AMF Plus restricted login

Restricted login is configured on a master node and allows only a privileged user logged into a master node to add other AMF Plus nodes to a working-set. In a non-secure mode AMF Plus network this feature is optional via configuration. In a secured AMF Plus network the restricted-login feature is required, and will automatically be enabled when secure mode is enabled. This means AMF Plus working-sets can only be created and used by a user logged on to an AMF Plus master. If secure mode is disabled on an AMF Plus network restricted login will remain enabled and will need to be disabled independently, if this is desired.

AMF Plus remote login

In a non-secure mode AMF Plus network, a user logged in to an AMF Plus node can use the remote login feature to connect to any other AMF Plus node. In AMF Plus secure mode, remote login to other AMF Plus nodes will only be allowed from an AMF Plus master node. AMF Plus member nodes

will not be able to use the remote login feature to connect to other nodes in the network, including to master nodes.

AMF Plus auto-recovery in secure mode

The pre-requisites for the successful auto-recovery of an AMF Plus member while the network is in secure mode are:

- The software on the AMF Plus backup for the pre-existing device must also support secure mode.
- The pre-existing device that is being replaced must have been configured for secure mode before the device was removed.

In general, recovery will no longer be zero-touch. When a clean device is attached, the recovery process will begin; however, the administrator will need to authorize the new device before it is allowed to complete the recovery process.

Note: Neighbor node auto-recovery is not supported for AMF Plus members connected to the AMF Plus network via virtual-links. For these nodes auto-recovery from external media should be used.

AMF Plus secure mode pre-authorization

It is possible for an administrator to allow an AMF Plus master to pre-authorize a device which can then be used as a replacement device for zero-touch recovery.

- The device is pre-authorized on an AMF Plus Master with the **atmf authorize provision** command. This is done by either specifying the MAC address of the replacement device, or by specifying the neighbouring node hostname and interface, i.e. the AMF Plus node and switch port the replacement device will attach to.
- The default timeout for pre-authorization is 60 minutes but can be extended to up to 6000 minutes (100 hours).
- If a node has been pre-authorized using its MAC address the device with that MAC address can be used to replace any compatible AMF Plus node in the network and will be automatically authorized upon joining the network.
- If pre-authorization has been configured for a specific member node and port combination, a compatible replacement device will only be authorized when connected to the correct member node and port.
- You can check which MAC addresses or node and port combinations have been pre-authorized with the **show atmf authorization provisional** command.

Output 33: .Example specifying a node and port combination

```

master#atmf authorize provision node area_1_node_4 interface port1.0.3
master#show atmf authorization provisional

ATMF Provisional Authorization:

Area - Node Name           Start           Timeout
or MAC Address             Interface       Time            Minutes
-----
area_1_node_4              port1.0.3      22 Feb 2023 07:38:24 60

```

Output 34: .Example specifying a MAC address

```

master#aatmf authorize provision mac 0000.5e00.5e23
master#show atmf authorization provisional

ATMF Provisional Authorization:

Area - Node Name           Start           Timeout
or MAC Address             Interface       Time            Minutes
-----
0000.5e00.5e23            22 Feb 2023 07:38:24 60

```

- Once a MAC address or node and port combination has been pre-authorized, a compatible "clean" device can be connected to the AMF Plus network and it will be automatically authorized by a controller or master and zero-touch recovery can proceed.

Note: A "clean" device is one on which the **atmf cleanup** command has been run, see ["Restoring a node to a "clean" state" on page 148.](#)

Enabling secure mode on an existing AMF Plus network

Executing the **atmf secure-mode enable-all** command on an AMF Plus master node will reconfigure an existing AMF Plus network into secure mode.

When the command is executed, all the AMF Plus members leave the AMF Plus network and rejoin in secure mode with each member being automatically authorized. There is a small disruption to the AMF Plus management network while the members leave and rejoin, however the data-plane traffic on the network will be unaffected by this process and will continue to operate as normal.

Enabling AMF Plus secure mode on an AMF Plus master automatically enables the AMF Plus restricted login feature if this has not already been enabled.

Note: This process saves the **running-config** on each device in the AMF Plus network.

Output 35: . Example output from the **atmf secure-mode enable-all** command

```

master#atmf secure-mode enable-all

Total number of nodes 3
3 nodes support secure-mode

Enable secure-mode across the ATMF network ? (y/n): y

master#02:22:43 AMF_master ATMF[749]: box2 has left. 2 members in total.
02:22:43 master ATMF[749]: box3 has left. 1 member in total.
02:22:47 master ATMF[749]: Node box3 (area:area1) [eccd.6ddc.5e9f] - preauthorized
02:22:47 master ATMF[749]: Node box2 (area:area1) [eccd.6ddc.5eb4] - preauthorized
02:22:51 master ATMF[749]: box2 has joined. 2 members in total.
02:22:58 master ATMF[749]: box3 has joined. 3 members in total.
master#02:23:08 master IMISH[5068]: All 3 compatible nodes have joined the secure
mode network.

```

You can confirm the AMF Plus secure mode status with the **show atmf** command.

Output 36: Example output from the **show atmf** command

```

master#show atmf
ATMF Summary Information:

ATMF Status           : Enabled
Network Name          : Example-ATMF
Node Name              : master
Role                   : Master
Restricted login       : Enabled
Secure Mode            : Enabled
Current ATMF Guests   : 0
Current ATMF Nodes    : 3

```

Adding an AMF Plus node to a secure mode network

An AMF Plus device that is not currently in secure mode can be placed in secure mode by executing the global configuration command **atmf secure-mode** on the node. This will prompt the device to request authorization from the AMF Plus master nodes.

All AMF Plus master nodes will receive the authorization request from the new member node. Once the node is authorized on any AMF Plus master device this information will then be propagated to all other AMF Plus masters in the area.

If a new AMF Plus node, with secure mode enabled, is connected and attempts to join the AMF Plus network a message is displayed on the AMF Plus master indicating that a new node is awaiting authorization.

Output 37: Example log message

```

master>03:06:33 master ATMF[749]: Node area_1_node_1 (area:area1)
[eccd.6db5.1045] - requests authorization

```

The administrator can then authorize this device, allowing it to join the AMF Plus network, using the **atmf authorize** command, either by specifying the node waiting for authorization by name and area,

or by using the **all-pending** parameter which will authorize all nodes currently awaiting authorization.

```
master#atmf authorize area_1_node_1
master#03:10:28 master ATMF[749]: area_1_node_1 has joined. 4 members in
total.
```

Disabling secure mode on an AMF Plus network

If you decide to remove secure mode from an AMF Plus network this can be done by executing the command **no atmf secure-mode enable-all** on an AMF Plus master node.

```
mster#no atmf secure-mode enable-all
% Warning: All security certificates will be deleted.
Disable secure-mode across the ATMF network ? (y/n): y
07:24:42 master IMISH[11133]: Please wait while nodes leave and rejoin
the network with the updated setting.
master#09:24:49 master ATMF[732]: Distribution has left. 2 members in
total.
09:24:49 master ATMF[732]: Edge has left. 1 member in total.
09:24:54 master ATMF[732]: Distribution has joined. 2 members in total.
09:24:54 master ATMF[732]: Edge has joined. 3 members in total.
07:24:56 master IMISH[11133]: All 3 nodes have joined the non-Secure-mode
network.
07:24:56 master IMISH[11133]: The running configuration has been updated
and written on all nodes.
07:24:56 master IMISH[11133]: Please back up all nodes in the network.
```

Note: This process saves the **running-config** on each device in the AMF Plus network.

Verifying secure mode on an AMF Plus network

The **show atmf** command shows whether or not the device currently being managed is in secure mode.

Output 38: Example output from the **show atmf** command

```
master#show atmf
ATMF Summary Information:

ATMF Status           : Enabled
Network Name          : Example-ATMF
Node Name              : master
Role                   : Master
Restricted login       : Enabled
Secure Mode            : Enabled
Current ATMF Guests   : 0
Current ATMF Nodes    : 4
```

Use the **show atmf secure-mode** command to display general information about the current status of secure mode.

Output 39: Example output from the **show atmf secure mode** command

```

master#show atmf secure-mode

ATMF Secure Mode:

Secure Mode Status           : Enabled
Certificate Expiry           : 180 Days
Certificates Total            : 4
Certificates Revoked          : 1
Certificates Rejected         : 0
Certificates Active           : 3

Provisional Authorization    : 0
Pending Requests             : 3

Trusted Master                : master

Key Fingerprint:
  b6:af:95:54:41:f5:e6:2e:17:0f:76:67:c4:02:d5:16:98:d4:cc:5a

```

The above output shows 3 active authorized nodes and three nodes waiting for authorization. Use the **show atmf authorization pending** command on an AMF Plus master to display a list of all AMF Plus nodes waiting for authorization.

Output 40: Example output from the **show atmf authorization pending** command

```

master#show atmf authorization pending

Pending Authorizations:

area1 Requests:
Node Name           Product           Parent Node   Interface
-----
area_1_node_2      x530L-28GTX      master        port1.0.3
area_1_node_3      x530L-28GTX      master        sa1
area_1_node_4      x330-28GTX       master        port1.0.1

```

Using the **show atmf authorization current** command from an AMF Plus master to view all devices in the AMF Plus network that are currently authorized to join the network.

Output 41: Example output from the **show atmf authorization current** command

```

master#show atmf authorization current

area1 Authorized Nodes:
Node Name           Signer           Expires
-----
area_1_node_1      master           16 Feb 2023
master              master           15 Feb 2023
area_1_node_5      master           16 Feb 2023
area_1_node_6      master           16 Feb 2023

```

The **show atmf secure-mode statistics** command shows the total number of valid certificates generated, and provides details on the overall status of certificates within the network. It can be used to confirm whether any invalid certificates have been received on the master from a node, and vice versa. These statistics can be cleared using the **clear atmf secure-mode statistics** command.

Output 42: Example output from the **show atmf secure-mode statistics** command

```

master#show atmf secure-mode statistics
  ATMF Secure Mode Statistics:

  Certificates:
  New ..... 7                Expired ..... 0
  Updated ..... 7            Deleted ..... 0
  Revoked ..... 1           Renewed ..... 2
  Rejected ..... 1         Re-authorized .... 1
  Authorized ..... 0

  Local Certificates:
  Valid ..... 4              Invalid ..... 0

  Certificates Validation:
  Request Valid ..... 2
  Request Invalid ..... 0
  Common Valid ..... 13
  Common Invalid ..... 0
  Issuer Valid ..... 14
  Issuer Invalid ..... 0
  Signature Verified ..... 29
  Signature Invalid ..... 0
  Signature Purpose Invalid ..... 0

  Signatures Signed ..... 12

  Master Certificates:
  Re-issued ..... 3
  Downgraded to member ..... 0

  Public key change ..... 2
  Invalid SA public key ..... 0

```

The **show atmf links** command displays a link state of “OneWaySm” if a device running in secure mode is connected to the AMF Plus network, but has not yet been authorized by an AMF Plus master to join the network.

Output 43: Example output from the **show atmf links** command

```

master#show atmf links

  ATMF Link Brief Information:

  Local      Link      Link      ATMF      Adjacent      Adjacent      Link
  Port      Type      Status    State     Node/Area     Ifindex      State
  -----
  sa1       Crosslink Up        Full      area_1_node_3  4501         Forwarding
  1.0.1     Crosslink Up        Full      area_1_node_4  5001         Blocking
  1.0.3     Downlink  Up        Full      area_1_node_2  5001         Forwarding
  1.0.5     Downlink  Up        OneWaySm  area_1_node_6  0            Blocking

```

Checking for vulnerabilities on an AMF Plus secure mode network

The **show atmf secure-mode audit** command displays a list of security recommendations for securing your AMF Plus network. Items prefaced with **Warning** need to be attended to. In this sample output the default username and password, and telnet, should be disabled.

Output 44: Example output from the **show atmf secure-mode audit** command

```

ATMF Secure Mode Audit:

Warning   : The default username and password is enabled.
Good      : SNMP V1 or V2 is disabled.
Warning   : Telnet server is enabled.
Good      : ATMF is enabled. Secure-Mode is on.
Good      : ATMF Topology-GUI is disabled. No trustpoints configured.

ATMF Secure Mode Log Events:

-----
2023 Feb 2 00:59:25 user.notice node1 ATMF[848]: Sec_Audit - ATMF Secure
Mode is enabled.
2023 Feb 2 01:30:00 user.notice node1 ATMF[848]: Sec_Audit - Established
secure connection to area_1_node_1 on interface vlink1.

```

To identify devices that are connected to a secure mode node that are not in secure mode or are not authorized, use the **show atmf secure-mode audit link** command.

Output 45: Example output from the **show atmf secure-mode audit link** command

```

ATMF Secure Mode Audit Link:

* ATMF links connected to devices which are not authorized or are not in
secure-mode.

Port      Link Type      Discovered          Node/Area Name
-----
vlink1    Downlink       16/02/2023 09:28:22  Member3

```

AMF Plus Cloud

Introduction

AMF Plus Cloud is a virtualized instance of AMF Plus Master and/or Controller running on a virtualization platform. This means you get all the functionality of integrated hardware-based management, with the added advantages of private or public cloud access and flexibility.

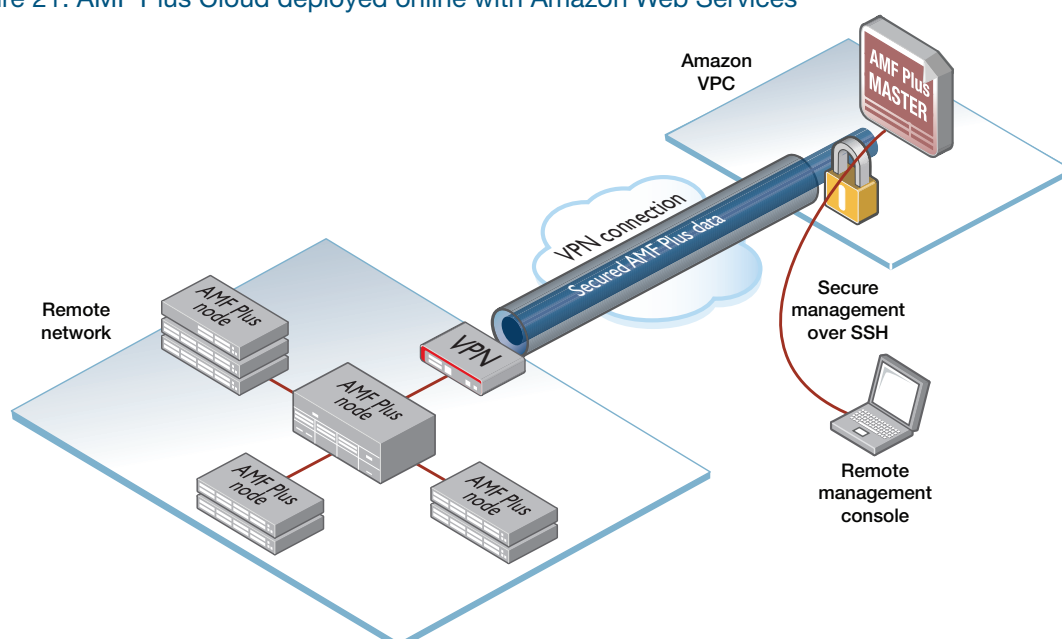
It is supplied either:

- as an ISO image, so you can deploy on one of the supported virtual machine environments,
- preloaded on a Vista Manager Network Appliance (VST- APL),
- or included with Vista Manager Virtual (VST-VRT).

The benefits include:

- Flexible deployment with private or public cloud installation—use your own local server, or deploy fully online with Amazon Web Services or (from version 5.4.7-1.1 onwards) Microsoft Azure.
- Scalable for any size network, with multiple licensing options.
- Lower cost of entry with no dedicated hardware requirements.
- Reduced costs with simple pay-as-you-go licensing.
- Web-based interface for remote network monitoring and management—anywhere, anytime.
- Peace of mind networking with a full back-up stored in the cloud.

Figure 21: AMF Plus Cloud deployed online with Amazon Web Services



AMF Plus Cloud Documentation

For information on deploying AMF Plus Cloud on one of the supported virtualization environments, see the [Virtual AMF Appliance \(VAA\) for AMF Cloud Installation Guide](#). This distribution of AMF Plus Cloud is also referred to as the Virtual AMF Appliance (VAA) in documentation.

For information on configuring and using AMF Plus Cloud on the Vista Manager Network Appliance (VST-APL), see the [VST-APL-06 and VST-APL-10 Vista Manager Appliances Installation Guide](#) and the [Vista Manager Network Appliance \(VST-APL\) User Guide](#).

For information on configuring and using AMF Plus Cloud on Vista Manager Virtual (VST-APL), see the [Vista Manager Virtual \(VST-VRT\) User Guide](#).

What is AMF Plus virtualization?

The AMF Plus Cloud is a variant of AlliedWare Plus that is supplied as a software image and loaded to a virtual machine at boot up time. Once loaded, the familiar AlliedWare Plus command-line interface (CLI) is available. Use this CLI to configure and manage the virtual AMF Plus controller or master just as you would a physical device.

AMF Plus Cloud supports a minimum level of functionality required to support AMF Plus, therefore the AMF Plus Cloud does not support standard L2 switching and L3 routing. However, all AMF Plus commands, such as **atmf working-set**, **atmf select-area**, and **atmf remote-login**, will work in the same manner as they do on a physical device. Similarly, scripts can be created within AMF Plus Cloud, and triggers created that will run these scripts.

AMF Plus Cloud connects to other AMF Plus devices (both physical and virtual) using virtual-links.

Licensing

AMF Plus Cloud licensing is subscription-based and depends on the size of the network under management. You will need to consider:

- how many AMF Plus masters throughout the network are linked to an AMF Plus controller, and
- how many nodes in each AMF Plus area are linked to that area's AMF Plus master.

Each AMF Plus Cloud instance acting as an AMF Plus controller or AMF Plus master will need its own unique license file that is based on the unique serial number of the AMF Plus Cloud instance. This license defines the number and type of nodes allowed throughout the AMF Plus network.

See the [Autonomous Management Framework Datasheet](#) for information on AMF Plus Cloud master and controller licensing.

Multiple Tenants on AMF Plus Cloud

Introduction

Running multiple tenants on an AMF Plus Cloud instance provides an efficient way to configure and control different service networks via a centralized virtual machine. It allows AMF Plus Cloud to act as both the AMF Plus controller and the AMF Plus masters for up to 300 AMF Plus areas. Each AMF Plus master runs in its own virtual AlliedWare Plus environment known as an AMF Plus container.

An AMF Plus container is an isolated instance of AlliedWare Plus with its own network interfaces, configuration, and file system. The features available inside an AMF Plus container are a sub-set of the features available on the host AMF Plus Cloud instance, this allows it to function as a uniquely identifiable AMF Plus master. From the host, each AMF Plus area can be managed using either Vista Manager or the command line interface (CLI).

The tenants in each AMF Plus area could be branch offices of a single organization, or separate customers managed by a single service provider. Hosting multiple tenants on a single AMF Plus Cloud instance could also be used where a service provider provisions an AMF Plus container for a client tenant, but the tenant manages their own AMF Plus area. This is possible because each area is isolated from all other areas through the use of AMF Plus containers, with each tenant only able to view and control their own area.

The key advantage of hosting multiple tenants on a single AMF Plus Cloud instance over a traditional AMF Plus installation is that each area does not need a device with the capabilities to act as an AMF Plus master. For information on installing AMF Plus Cloud see the [Virtual AMF Appliance \(VAA\) for AMF Cloud Installation Guide](#).

Note: The multiple tenant feature is only supported on AMF Plus Cloud distributed as an ISO image. It is **not** supported on AMF Plus Cloud distributed on the Vista Manager Network Appliance (VST-APL) or with Vista Manager Virtual (VST-VRT).

Feature overview


The AMF Plus Cloud host provides the CLI commands necessary to create and configure AMF Plus containers. It functions as the AMF Plus controller for the local area, while remote AMF Plus areas are then defined and assigned to each AMF Plus container.

Once the containers are configured with an area-link and enabled, they are automatically set-up as AMF Plus masters with their own AMF Plus area configuration. The administrator can then configure AMF Plus virtual-links inside these containers, which connect to physical AMF Plus members in the remote AMF Plus area.

Each AMF Plus container has two virtual interfaces, eth0 and eth1:

- The eth0 interface is used to connect the container to the AMF Plus Clouds' host AMF controller using an AMF area-link.
- The eth1 interface is used to connect the container to the outside world, using a bridged L2 network. This allows the AMF Plus master within each container to establish AMF Plus virtual-links with other AMF Plus devices in its area.

Figure 22: AMF Plus container architecture



- In the case of AMF Plus Cloud hosted on a public cloud service, the virtual-links are carried from the remote AMF Plus area over an IPsec protected L2TPv3 tunnel.

Each AMF Plus container on an AMF Plus Cloud host has its own directory containing the file system for the container. This is the equivalent of the flash file system on a physical AlliedWare Plus device. A container's file system is accessed using the path `"/flash/containers/<container_name>/"`

- Note:** The system clock/time is shared between the host and all containers, so containers are restricted from setting the system clock. This means AMF Plus containers do not support NTP. Therefore, NTP should be configured on the AMF Plus Cloud host, and every AMF Plus node directly connected to the AMF Plus container.

Licensing

This feature is included in the AMF Plus Cloud software and does not require a special license. It does, however, need the following:

- A AMF Plus controller license on the host AMF Plus Cloud instance for the number of areas the controller will support (up to 300).

The AMF Plus controller license allows you to create an AMF Plus master on the controller node if there is no other AMF Plus master in the area.

- A AMF Plus master license on each AMF Plus container for the number of nodes the master will support (up to 300).

With no master license installed, an AMF Plus container supports only a single node, i.e. the container's AMF Plus master.

See the [Autonomous Management Framework Datasheet](#) for more information on AMF Plus Cloud master and controller licensing.

Configuration example for private cloud installation

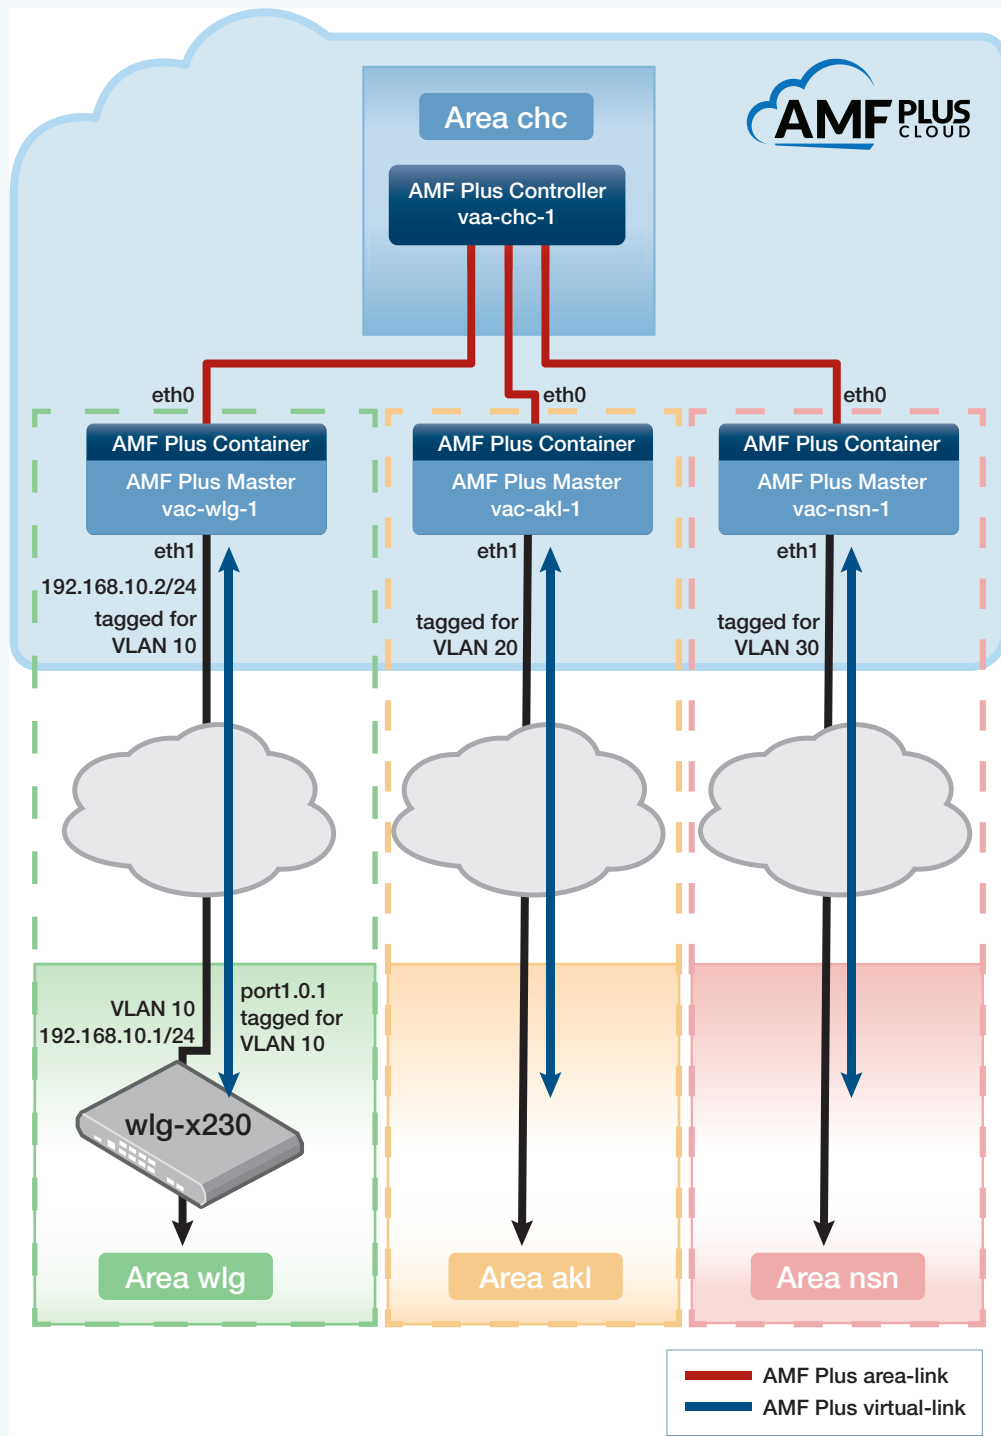
Complete the following steps to configure multiple tenants on a single AMF Plus Cloud instance:

1. “Configure AMF Plus on the host AMF Plus Cloud”
2. “Create an AMF Plus container on the AMF Plus Cloud host”
3. “Configure the AMF Plus container”
4. “Configuring the bridge connecting the container to the Hypervisor Ethernet interface”
5. “Assigning the bridge group to a container on the host AMF Plus Cloud”
6. “Configuring the remote AlliedWare Plus device”
7. “Verifying the AMF Plus container”
8. “Managing an AMF Plus area”

The diagram below shows a virtual machine configured with a local area and three AMF Plus containers. The left-hand container illustrates the AMF Plus virtual-link between the AMF Plus master in container “vac-wlg-1” and a physical switch at the remote site within AMF Plus area “wlg”.

Note: This example assumes that the eth1 IP address of the container is in the same IPv4 subnet as the IP address of the remote switch interface that terminates the AMF Plus virtual-link. The IP address of the eth1 container interface must have direct IP connectivity with the remote device terminating the AMF Plus virtual-link. Typically this would be achieved using a site-to-site VPN.

Figure 23: AMF Plus multiple tenant configuration on private cloud example



Configure AMF Plus on the host AMF Plus Cloud

First create the AMF Plus areas, one for the local area which contains the AMF Plus controller, and one for each AMF Plus container. In this example we will create a local area on the host AMF Plus Cloud named “chc” with an ID of 1, and a remote area named “wlg” with an ID of 2. The area “wlg” will be assigned to a container.

Step 1: Create the AMF Plus Cloud hostname and AMF Plus network name

```
awplus#configure terminal
awplus(config)#hostname vaa-chc-1
vaa-chc-1(config)#atmf network-name atlnz
```

Step 2: Configure the AMF Plus Cloud as AMF Plus controller and master

```
vaa-chc-1(config)#atmf master
vaa-chc-1(config)#atmf controller
```

Note: If the AMF Plus Cloud is the only AMF Plus node in its area, it must be configured as both an AMF Plus master and an AMF Plus controller. If the local area contains other AMF Plus nodes, one of the other nodes in that area can be configured as an AMF Plus master. There must be an AMF Plus master somewhere in the AMF Plus controller's area.

Step 3: Configure the local AMF Plus area

This area, "chc", will contain the local AMF Plus controller and local AMF Plus master:

```
vaa-chc-1(config)#atmf area chc id 1 local
```

Step 4: Create a remote area, and configure it with a password

```
vaa-chc-1(config)#atmf area wlg id 2
vaa-chc-1(config)#atmf area wlg password secret123
```

Create an AMF Plus container on the AMF Plus Cloud host

Next we create an AMF Plus container named "vac-wlg-1" and, inside this container, configure an **area-link** to area "wlg":

Step 5: Create the AMF Plus container

```
vaa-chc-1(config)#atmf container vac-wlg-1
vaa-chc-1(config-atmf-container)#description Wellington
```

Step 6: Configure the area-link

```
vaa-chc-1(config-atmf-container)#area-link wlg
vaa-chc-1(config-atmf-container)#state enable
```

When the **state enable** command is executed, the configuration below is automatically applied to the AMF Plus container by the AMF Plus Cloud host. This assigns the container to the area "wlg" and configures it as an AMF Plus master.

```
atmf network-name atlnz
atmf master
atmf area wlg id 2 local
atmf area wlg password secret123
atmf area chc id 1

interface eth0
 atmf-arealink remote-area chc vlan 4094
```

Once the start-up configuration has been saved from within the AMF Plus container, all further configuration must be added manually.

Step 7: Exit to privilege exec mode

```
vaa-chc-1(config-atmf-container)#exit
vaa-chc-1(config)#exit
vaa-chc-1#
```

Configure the AMF Plus container

Connect to the AMF Plus container “vac-wlg-1”, login, and add an IP address to the eth1 interface. Eth1 is the AMF Plus container interface that will connect to the physical AlliedWare Plus devices within the container’s remote area.

Step 8: Connect to the AMF Plus container and login

```
vaa-chc-1#atmf container login vac-wlg-1
```

```
Connected to tty 1
  Type <Ctrl+a q> to exit the console, <Ctrl+a Ctrl+a> to enter Ctrl+a itself

vac-wlg-1 login: manager
Password: friend

AlliedWare Plus (TM) 5.4.7 02/03/17 08:46:12
vac-wlg-1>
```

```
vac-wlg-1>enable
vac-wlg-1#
```

Step 9: Add an IP address to the eth1 interface

```
vac-wlg-1#configure terminal
vac-wlg-1(config)#interface eth1
vac-wlg-1(config-if)#ip address 192.168.10.2/24
vac-wlg-1(config-if)#exit
```

Step 10: Configure an AMF Plus virtual-link

```
vac-wlg-1(config)#atmf virtual-link id 1 ip 192.168.10.2 remote-id 1
remote-ip 192.168.10.1
```

Step 11: Save the AMF Plus container’s configuration

If you have finished configuring the AMF Plus container, this would be a good time to save its configuration.

```
vac-wlg-1(config)#exit
vac-wlg-1#copy running-config startup-config
```

Configuring the bridge connecting the container to the Hypervisor Ethernet interface

To add the bridge configuration to connect the AMF Plus Cloud host to the container “vac-wlg-1” you need to exit to the AMF Plus Cloud host.

Note: All traffic on bridges must be tagged because “native VLAN” is not supported by the hypervisor virtual switch. If you are using **VMware vSphere Hypervisor** you will need to set “VLAN ID: All (4095)” in the VMware port group settings. This, in effect, tags a port to allow all VLAN IDs to pass through it. This step is not required if you are using the **XenServer** hypervisor.

Step 12: Exit to the AMF Plus Cloud host

```
vac-wlg-1(config)#exit
vac-wlg-1#
```

Type <Ctrl+a q> to exit the container and return to the AMF Plus controller console.

```
vaa-chc-1#
```

Step 13: Configure the bridge

```
vaa-chc-1#configure terminal
vaa-chc-1(config)#bridge 1
vaa-chc-1(config)#interface eth0
vaa-chc-1(config-if)#encapsulation dot1q 10
vaa-chc-1(config-if)#interface eth0.10
vaa-chc-1(config-if)#bridge-group 1
vaa-chc-1(config-if)#exit
vaa-chc-1(config)#
```

Note: Each AMF Plus container must be configured with its own separate VLAN ID and bridge, using the commands above. This ensures isolation of containers such that they may even have duplicate or overlapping IP ranges.

Assigning the bridge group to a container on the host AMF Plus Cloud

The bridge group created on host needs to be assigned to the container “vac-wlg-1”

Step 14: Assign the bridge group to the container

```
vaa-chc-1(config)#atmf container vac-wlg-1
vaa-chc-1(config-atmf-container)#bridge-group 1
```

Step 15: Save the AMF Plus controller's configuration

If you have finished configuring the AMF Plus controller, this would be a good time to save its configuration.

```
vaa-chc-1(config-atmf-container)#exit
vaa-chc-1(config)#exit
vaa-chc-1#copy running-config startup-config
```

Note: This only saves the configuration of the host AMF Plus Cloud. Each AMF Plus container's configuration must be saved from within that container.

Configuring the remote AlliedWare Plus device

Configure an AMF Plus virtual-link from the AMF Plus member "wlg-x230" in area "wlg" to the AMF Plus container "vac-wlg-1".

Step 16: Create a VLAN to use as the VLAN ID for the remote area

```
wlg-x230#configure terminal
wlg-x230(config)#vlan database
wlg-x230(config-vlan)#vlan 10 state enable
```

Step 17: Configure a port for trunk mode and add the VLAN

```
wlg-x230(config)#interface port1.0.1
wlg-x230(config-if)#switchport mode trunk
wlg-x230(config-if)#switchport trunk allowed vlan add 10
```

Step 18: Add an IP address to the VLAN

This IP address must be on the same subnet as the eth1 address of the AMF Plus container for this remote area.

```
wlg-x230(config-if)#interface vlan10
wlg-x230(config-if)#ip address 192.168.10.1/24
```

Step 19: Add an AMF Plus virtual-link from the remote device to the AMF Plus container

```
wlg-x230(config-if)#exit
wlg-x230(config)#atmf virtual-link id 1 ip 192.168.10.1 remote-id 1
remote-ip 192.168.10.2
```

Step 20: Save the remote device's configuration

If you have finished configuring the remote device, this would be a good time to save its configuration.

```
wlg-x230(config)#exit
wlg-x230#copy running-config startup-config
```


Verifying the AMF Plus container

You can check the state and resource utilization of an AMF Plus container with the **show atmf container** command.

Output 46: Example output from the **show atmf container** command

```
vaa-chc-1#show atmf container

ATMF Container Information:

  Container      Area      Bridge  State  Memory  CPU%
-----
  vac-wlg-1     wlg      br1     running 70.3 MB  1.2
  vac-akl-1     ak1      br2     stopped 0 bytes  0.0
  vac-nsn-1     nsn      br3     running 53.2 MB  0.7

Current ATMF Container count: 3
```

This command can also be run for a specific AMF Plus container.

Output 47: Example output from the **show atmf container <container-name>** command

```
vaa-chc-1#show atmf container vac-wlg-1

ATMF Container Information:

  Container      Area      Bridge  State  Memory  CPU%
-----
  vac-wlg-1     wlg      br1     running 70.3 MB  1.2

Current ATMF Container count: 1
```

For more detailed information for all AMF Plus containers running on a AMF Plus Cloud host use the **show atmf container detail** command.

Output 48: Example output from the **show atmf container detail** command

```
vaa-chc-1#show atmf container detail
```

```
ATMF Container Detail Information:
```

```
Name: vac-wlg-1
State: RUNNING
PID: 980
IP: 172.31.0.1
IP: 192.168.0.2
IP: fd00:4154:4d46:3c::1
CPU use: 3.95 seconds
Memory use: 67.07 MiB
KMem use: 0 bytes
Link: vethP31UFA
TX bytes: 166.01 KiB
RX bytes: 141.44 KiB
Total bytes: 307.45 KiB
Link: vethYCT7BB
TX bytes: 674.27 KiB
RX bytes: 698.27 KiB
Total bytes: 1.34 MiB
```

```
Name: vac-akl-1
State: STOPPED
```

```
Name: vac-nsn-1
State: RUNNING
PID: 1086
IP: 172.31.0.1
CPU use: 3.34 seconds
Memory use: 50.82 MiB
KMem use: 0 bytes
Link: veth6LOD7B
TX bytes: 0 bytes
RX bytes: 98.59 KiB
Total bytes: 98.59 KiB
Link: vethJWJ350
TX bytes: 0 bytes
RX bytes: 0 bytes
Total bytes: 0 bytes
```

To show more detailed information for a single AMF Plus container running on a AMF Plus Cloud host use the **show atmf container detail <container-name>** command, where **<container-name>** is the name of the container you wish to examine, for example “vac-wlg-1”.

Output 49: Example output from the `show atmf container detail <container-name>` command

```
vaa-chc-1#show atmf container detail vac-wlg-1
```

```
ATMF Container Detail Information:
```

```
Name: vac-wlg-1  
State: RUNNING  
PID: 980  
IP: 172.31.0.1  
IP: 192.168.0.2  
IP: fd00:4154:4d46:3c::1  
CPU use: 3.95 seconds  
Memory use: 67.07 MiB  
KMem use: 0 bytes  
Link: vethP31UFA  
TX bytes: 166.01 KiB  
RX bytes: 141.44 KiB  
Total bytes: 307.45 KiB  
Link: vethYCT7BB  
TX bytes: 674.27 KiB  
RX bytes: 698.27 KiB  
Total bytes: 1.34 MiB
```

Managing an AMF Plus area

The virtual AMF Plus master, inside the AMF Plus container, can now be used to manage its AMF Plus area in exactly the same way as if it was a physical master.

Configuration example for public cloud installations

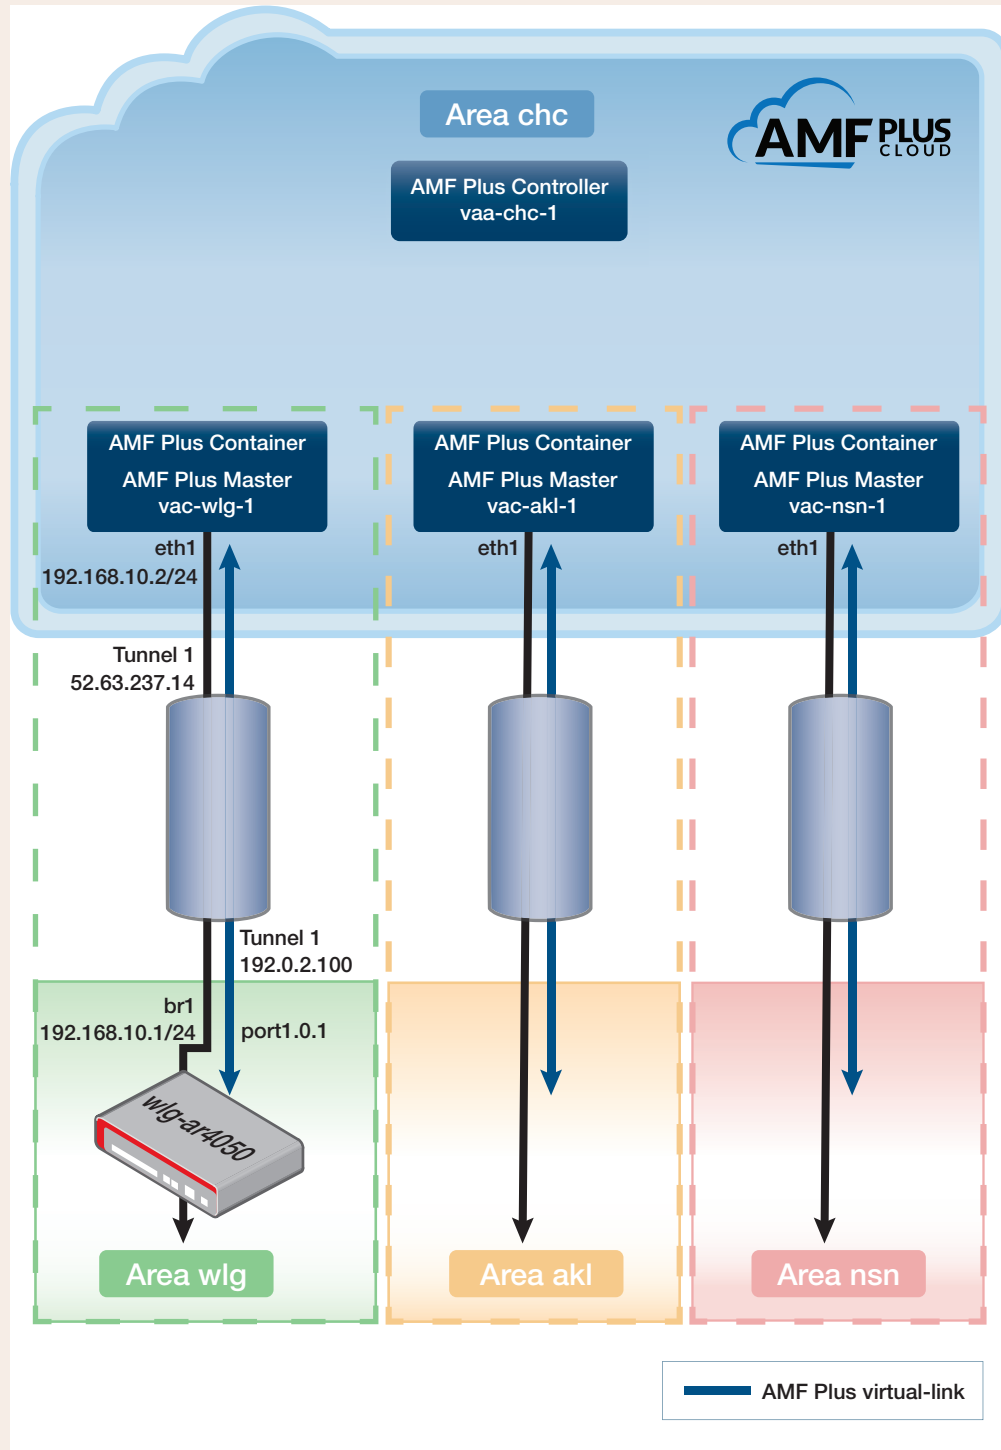
Complete the following steps to configure multiple tenants on a single AMF Plus Cloud host:

1. “Configure AMF Plus on the AMF Plus Cloud host”
2. “Create an AMF Plus container on the AMF Plus Cloud host”
3. “Configuring L2TPv3 tunnel with IPsec encryption on AMF Plus Cloud host”
4. “Configure the AMF Plus container”
5. “Assigning the bridge group to a container on the AMF Plus Cloud host”
6. “Configuring the remote AlliedWare Plus device”
7. “Verifying the AMF Plus container”
8. “Managing an AMF Plus area”

The diagram below shows a AMF Plus Cloud virtual machine configured with a local area and three AMF Plus containers. The left-hand container illustrates the AMF Plus virtual-link between the AMF Plus master in container “vac-wlg-1” and a physical switch at the remote site within AMF Plus area “wlg”.

Note: This example assumes that the eth1 IP address of the container is in the same IPv4 subnet as the IP address of the remote switch interface that terminates the AMF Plus virtual-link. The IP address of the eth1 container interface must have direct IP connectivity with the remote device terminating the AMF Plus virtual-link. Typically this would be achieved using a site-to-site VPN.

Figure 24: AMF Plus multiple tenant configuration on public cloud example



Configure AMF Plus on the AMF Plus Cloud host

First create the AMF Plus areas, one for the local area which contains the AMF Plus controller, and one for each AMF Plus container. In this example we will create a local area on the AMF Plus Cloud host named "chc" with an ID of 1, and a remote area named "wlg" with an ID of 2. The area "wlg" will be assigned to a container.

Step 1: Create the AMF Plus Cloud hostname and AMF Plus network name

```
awplus#configure terminal
awplus(config)#hostname vaa-chc-1
vaa-chc-1(config)#atmf network-name atlnz
```

Step 2: Configure the AMF Plus Cloud as AMF Plus controller and master

```
vaa-chc-1(config)#atmf master
vaa-chc-1(config)#atmf controller
```

Note: If the AMF Plus Cloud is the only AMF Plus node in its area, it must be configured as both an AMF Plus master and an AMF Plus controller. If the local area contains other AMF Plus nodes, one of the other nodes in that area can be configured as a master. There must be an AMF Plus master somewhere in the AMF Plus controller's area.

Step 3: Configure the local AMF Plus area

This area, "chc", will contain the local AMF Plus controller and local AMF Plus master:

```
vaa-chc-1(config)#atmf area chc id 1 local
```

Step 4: Create a remote area, and configure it with a password

```
vaa-chc-1(config)#atmf area wlg id 2
vaa-chc-1(config)#atmf area wlg password secret123
```

Create an AMF Plus container on the AMF Plus Cloud host

Next we create an AMF Plus container named "vac-wlg-1" and, inside this container, configure an **area-link** to area "wlg":

Step 5: Create the AMF Plus container

```
vaa-chc-1(config)#atmf container vac-wlg-1
vaa-chc-1(config-atmf-container)#description Wellington
```

Step 6: Configure the area-link

```
vaa-chc-1(config-atmf-container)#area-link wlg
vaa-chc-1(config-atmf-container)#state enable
```

When the **state enable** command is executed, the configuration below is automatically applied to the AMF Plus container by the AMF Plus Cloud host. This assigns the container to the area "wlg" and configures it as an AMF Plus master.

```
atmf network-name atlnz
atmf master
atmf area wlg id 2 local
atmf area wlg password secret123
atmf area chc id 1

interface eth0
 atmf-arealink remote-area chc vlan 4094
```

Once the start-up configuration has been saved from within the AMF Plus container, all further configuration must be added manually.

Step 7: Exit to privilege configuration mode

```
vaa-chc-1(config-atmf-container)#exit
```

Configuring L2TPv3 tunnel with IPsec encryption on AMF Plus Cloud host

Add an L2TPv3 tunnel with IPsec encryption to the AMF Plus Cloud controller

Step 8: Create the tunnel

```
vac-chc-1#configure terminal
vac-chc-1(config)#interface tunnel1
vac-chc-1(config-if)#mtu 1436
vac-chc-1(config-if)#tunnel protection ipsec
vac-chc-1(config-if)#tunnel mode l2tpv3
```

Step 9: Add the public address of the AMF Plus Cloud controller as the tunnel source

```
vac-chc-1(config-if)#tunnel source 52.63.237.14
vac-chc-1(config-if)#tunnel local id 2
vac-chc-1(config-if)#tunnel local name vac-chc
```

Note: Tunnel local and remote names need to be configured when the devices are behind NAT boundaries.

Step 10: Add the PPP address of the router as the tunnel destination

```
vac-chc-1(config-if)#tunnel destination 192.0.2.100
vac-chc-1(config-if)#tunnel remote id 1
vac-chc-1(config-if)#tunnel remote name wlg-ar4050
vac-chc-1(config-if)#exit
```

Step 11: Create a preshared key for key exchange with the remote end of the tunnel

The hostname used in the key is the same as the tunnel remote name.

```
vac-chc-1(config)#crypto isakmp key tunnelkey hostname wlg-ar4050
vaa-chc-1(config)#exit
vaa-chc-1#
```

Configure the AMF Plus container

Connect to the AMF Plus container “vac-wlg-1”, login, and add an IP address to the eth1 interface. Eth1 is the AMF Plus container interface that will connect to the physical AlliedWare Plus devices within the container’s remote area.

Step 12: Connect to the AMF Plus container and login

```
vaa-chc-1#atmf container login vac-wlg-1
```

```
Connected to tty 1
Type <Ctrl+a q> to exit the console, <Ctrl+a Ctrl+a> to enter Ctrl+a itself

vac-wlg-1 login: manager
Password: friend

AlliedWare Plus (TM) 5.4.7 02/03/17 08:46:12
vac-wlg-1>
```

```
vac-wlg-1>enable
```

```
vac-wlg-1#
```

Step 13: Add an IP address to the eth1 interface

```
vac-wlg-1#configure terminal
```

```
vac-wlg-1(config)#interface eth1
```

```
vac-wlg-1(config-if)#ip address 192.168.10.2/24
```

```
vac-wlg-1(config-if)#exit
```

Step 14: Configure an AMF Plus virtual-link

```
vac-wlg-1(config)#atmf virtual-link id 1 ip 192.168.10.2 remote-id 1
remote-ip 192.168.10.1
```

Step 15: Save the AMF Plus container’s configuration

If you have finished configuring the AMF Plus container, this would be a good time to save its configuration.

```
vac-wlg-1(config)#exit
```

```
vac-wlg-1#copy running-config startup-config
```

Configuring the bridge connecting the container to the L2TPv3 tunnel.

To add the bridge configuration to connect the container “vac-wlg-1” to the tunnel you need to exit to the AMF Plus Cloud host.

Step 16: Exit to the host AMF Plus Cloud

```
vac-wlg-1(config)#exit
```

```
vac-wlg-1#
```

Type <Ctrl+a q> to exit the container and return to the AMF Plus controller console.

```
vaa-chc-1#
```


Step 17: Configure the bridge

```
vaa-chc-1#configure terminal
vaa-chc-1(config)#bridge 1
vaa-chc-1(config)#interface tunnel1
vaa-chc-1(config-if)#bridge-group 1
vaa-chc-1(config-if)#exit
```

Assigning the bridge group to a container on the AMF Plus Cloud host

The bridge group created on the AMF Plus Cloud needs to be assigned to the container “vac-wlg-1”

Step 18: Assign the bridge group to the container

```
vaa-chc-1(config)#atmf container vac-wlg-1
vaa-chc-1(config-atmf-container)#bridge-group 1
```

Step 19: Save the AMF Plus controller’s configuration

If you have finished configuring the AMF Plus controller, this would be a good time to save its configuration.

```
vaa-chc-1(config-atmf-container)#exit
vaa-chc-1(config)#exit
vaa-chc-1#copy running-config startup-config
```

Note: This only saves the configuration of the host AMF Plus Cloud. Each AMF Plus container’s configuration must be saved from within that container.

Configuring the remote AlliedWare Plus device

Configure an AMF Plus virtual-link from the AMF Plus member “wlg-ar4050” in area “wlg” to the AMF Plus container “vac-wlg-1”.

Step 20: Create a VLAN to use as the VLAN ID for the remote area

```
wlg-ar4050#configure terminal
wlg-ar4050(config)#vlan database
wlg-ar4050(config-vlan)#vlan 10 state enable
```

Step 21: Configure access port and add the VLAN

```
wlg-ar4050(config)#interface port1.0.1
wlg-ar4050(config-if)#switchport mode access
wlg-ar4050(config-if)#switchport access vlan 10
wlg-ar4050(config)#exit
```

Configuring the remote end of the L2TPv3 tunnel

Add the remote end of the L2TPv3 tunnel on the AlliedWare Plus device.

Step 22: Create the tunnel

```
wlg-ar4050(config)#interface tunnel1
wlg-ar4050(config-if)#mtu 1436
wlg-ar4050(config-if)#tunnel protection ipsec
wlg-ar4050(config-if)#tunnel mode l2tpv3
```

Step 23: Add the public address of the AMF Plus Cloud controller as the tunnel source

```
wlg-ar4050(config-if)#tunnel source 192.0.2.100
wlg-ar4050(config-if)#tunnel local id 1
wlg-ar4050(config-if)#tunnel local name wlg-ar4050
```

Step 24: Add the public address of the AMF Plus Cloud controller as the tunnel destination

```
wlg-ar4050(config-if)#tunnel destination 52.63.237.14
wlg-ar4050(config-if)#tunnel remote id 2
wlg-ar4050(config-if)#tunnel remote name vaa-chc
wlg-ar4050(config-if)#exit
```

Note: Tunnel local and remote names need to be configured when the devices are behind NAT boundaries.

Step 25: Create a preshared key for key exchange with the AMF Plus Cloud end of the tunnel

The hostname used in the key is the same as the tunnel remote name.

```
wlg-ar4050(config)#crypto isakmp key tunnelkey hostname vaa-chc
```

Configuring the bridge connecting the remote device to the L2TPv3 tunnel.

Step 26: Configure the bridge

```
wlg-ar4050(config)#bridge 1
wlg-ar4050(config)#interface tunnel1
wlg-ar4050(config-if)#bridge-group 1
wlg-ar4050(config-if)#exit
wlg-ar4050(config)#
```

Step 27: Assign the bridge group to VLAN

```
wlg-ar4050(config)#interface port1.0.1
wlg-ar4050(config-if)#bridge-group 1
```

Step 28: Assign an IP address to the bridge

This IP address must be on the same subnet as the eth1 address of the AMF Plus container for this remote area.

```
wlg-ar4050(config-if#interface br1
wlg-ar4050(config-if)#ip address 192.168.10.1/24
```

Step 29: Add an AMF Plus virtual-link from the remote device to the AMF Plus container

```
wlg-ar4050(config-if#exit
wlg-ar4050(config)#atmf virtual-link id 1 ip 192.168.10.1 remote-id 1
remote-ip 192.168.10.2
```

Step 30: Save the device's configuration

If you have finished configuring the remote device, this would be a good time to save its configuration.

```
wlg-ar4050(config)#exit
wlg-ar4050#copy running-config startup-config
```

Verifying the AMF Plus container

You can check the state and resource utilization of an AMF Plus container with the **show atmf container** command.

Output 50: Example output from the **show atmf container** command

```
vaa-chc-1#show atmf container

ATMF Container Information:

  Container          Area          Bridge   State   Memory   CPU%
  -----
  vac-wlg-1         wlg           br1      running 70.3 MB  1.2
  vac-akl-1         akl           br2      stopped 0 bytes  0.0
  vac-nsn-1         nsn           br3      running 53.2 MB  0.7

Current ATMF Container count: 3
```

This command can also be run for a specific AMF Plus container.

Output 51: Example output from the **show atmf container <container-name>** command

```
vaa-chc-1#show atmf container vac-wlg-1

ATMF Container Information:

  Container          Area          Bridge   State   Memory   CPU%
  -----
  vac-wlg-1         wlg           br1      running 70.3 MB  1.2

Current ATMF Container count: 1
```

For more detailed information for all AMF Plus containers running on a AMF Plus Cloud host use the **show atmf container detail** command.

Output 52: Example output from the **show atmf container detail** command

```
vaa-chc-1#show atmf container detail
```

```
ATMF Container Detail Information:
```

```
Name: vac-wlg-1
State: RUNNING
PID: 980
IP: 172.31.0.1
IP: 192.168.0.2
IP: fd00:4154:4d46:3c::1
CPU use: 3.95 seconds
Memory use: 67.07 MiB
KMem use: 0 bytes
Link: vethP31UFA
TX bytes: 166.01 KiB
RX bytes: 141.44 KiB
Total bytes: 307.45 KiB
Link: vethYCT7BB
TX bytes: 674.27 KiB
RX bytes: 698.27 KiB
Total bytes: 1.34 MiB
```

```
Name: vac-akl-1
State: STOPPED
```

```
Name: vac-nsn-1
State: RUNNING
PID: 1086
IP: 172.31.0.1
CPU use: 3.34 seconds
Memory use: 50.82 MiB
KMem use: 0 bytes
Link: veth6LOD7B
TX bytes: 0 bytes
RX bytes: 98.59 KiB
Total bytes: 98.59 KiB
Link: vethJWJ350
TX bytes: 0 bytes
RX bytes: 0 bytes
Total bytes: 0 bytes
```

To show more detailed information for a single AMF Plus container running on a AMF Plus Cloud host use the **show atmf container detail <container-name>** command. Where **<container-name>** is the name of the container you wish to examine, for example “vac-wlg-1”.

Using AMF Plus in EPSR Rings

In this chapter we look at how to appropriately select whether to configure AMF Plus links or AMF Plus Cross-links, firstly in a single-ring EPSR network and then in a dual-ring EPSR network.

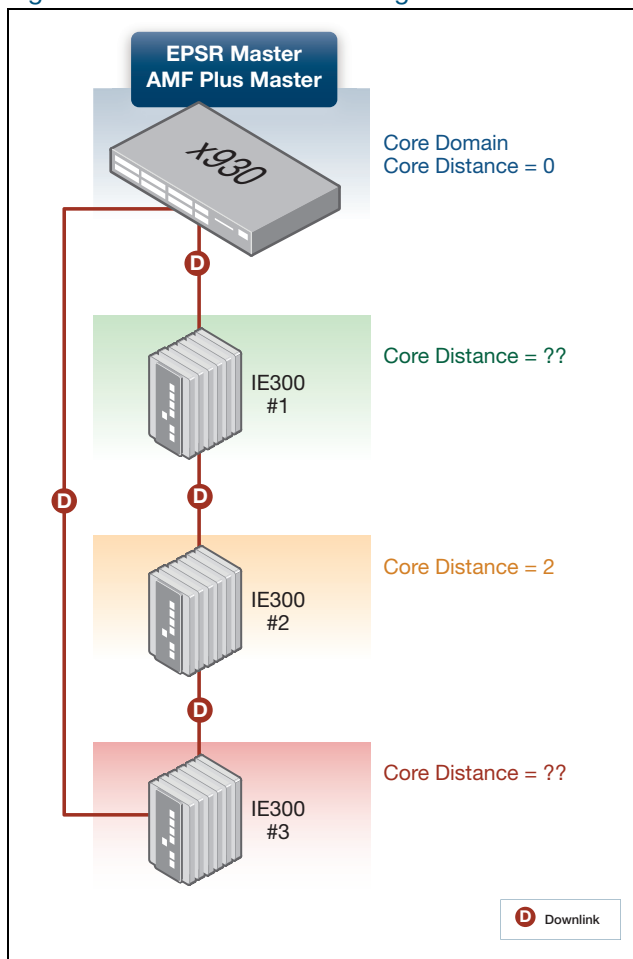
Down-links and cross-links when adding AMF Plus to an EPSR ring

When AMF Plus is added to an existing EPSR ring, a common mistake is to mark **all** ring links as AMF Plus down-links. An example of this is shown below in [Figure 25](#).

As you can see, each IE300 device has two routes back to the AMF Plus Master, and there is a ring of down-links, so what are the respective core distances?

- Does IE300 #1 have a core distance of 1 or 3?
- Does IE300 #3 have a core distance of 1 or 3?

Figure 25: Common error - all ring links marked as AMF Plus down-links



How can you configure this example above to avoid the core distance issue? Let us now consider some good and flawed alternative solutions.

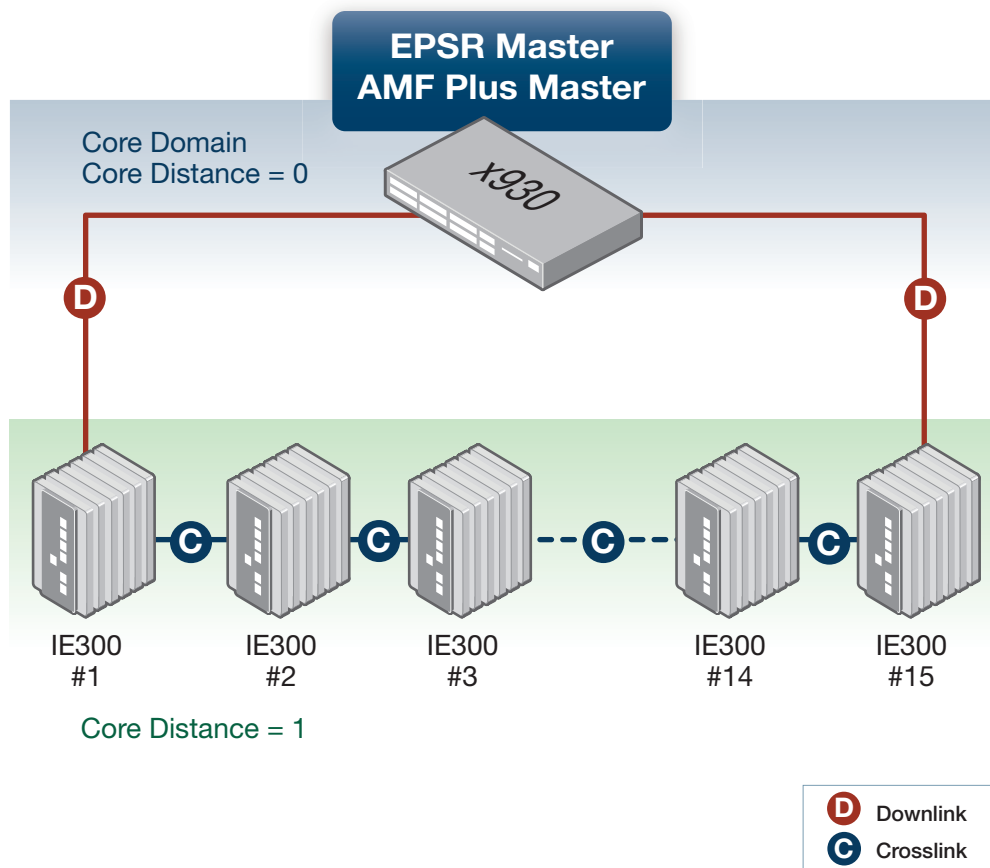
Solution 1

If the EPSR Master does not have existing AMF Plus cross-links to other nodes, the easiest solution to the problem in Figure 25, is to mark all EPSR ring links as AMF Plus cross-links. This places all EPSR nodes within the same AMF Plus Domain. But this is contingent on the EPSR ring not being so large as to exceed the recommended maximum number of 12 nodes in an AMF Plus Domain.

Solution 2

You can see in the diagram below, that both EPSR Master ring links are AMF Plus down-links, but all other ring links are made AMF Plus cross-links.

Figure 26: Mark EPSR Master links as down-links and all other ring links as cross-links

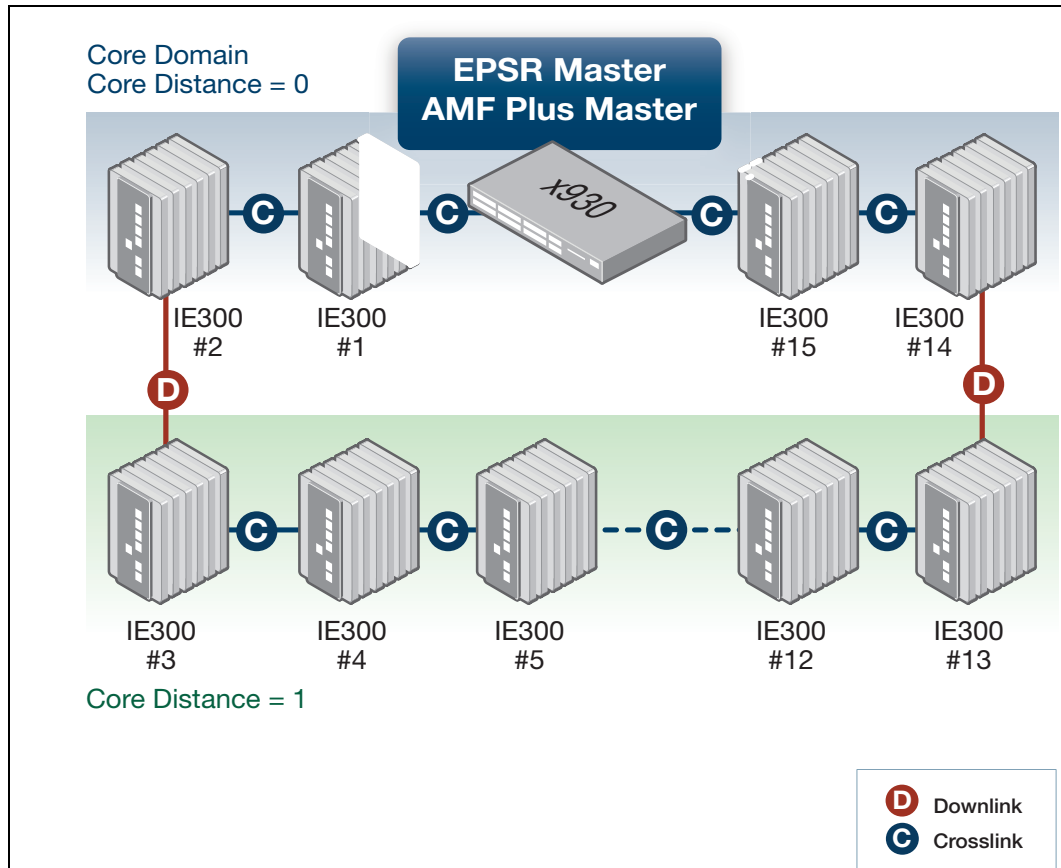


Solution 3

Building on the diagram above in Figure 26, what if your EPSR ring contains 24 nodes? How can you avoid exceeding the recommended maximum of 12 nodes per AMF Plus domain?

You can do this by placing some of the transit nodes in the same AMF Plus domain as the EPSR master, and some in a child domain.

Figure 27: Place some transit nodes into a child domain



Solution 4

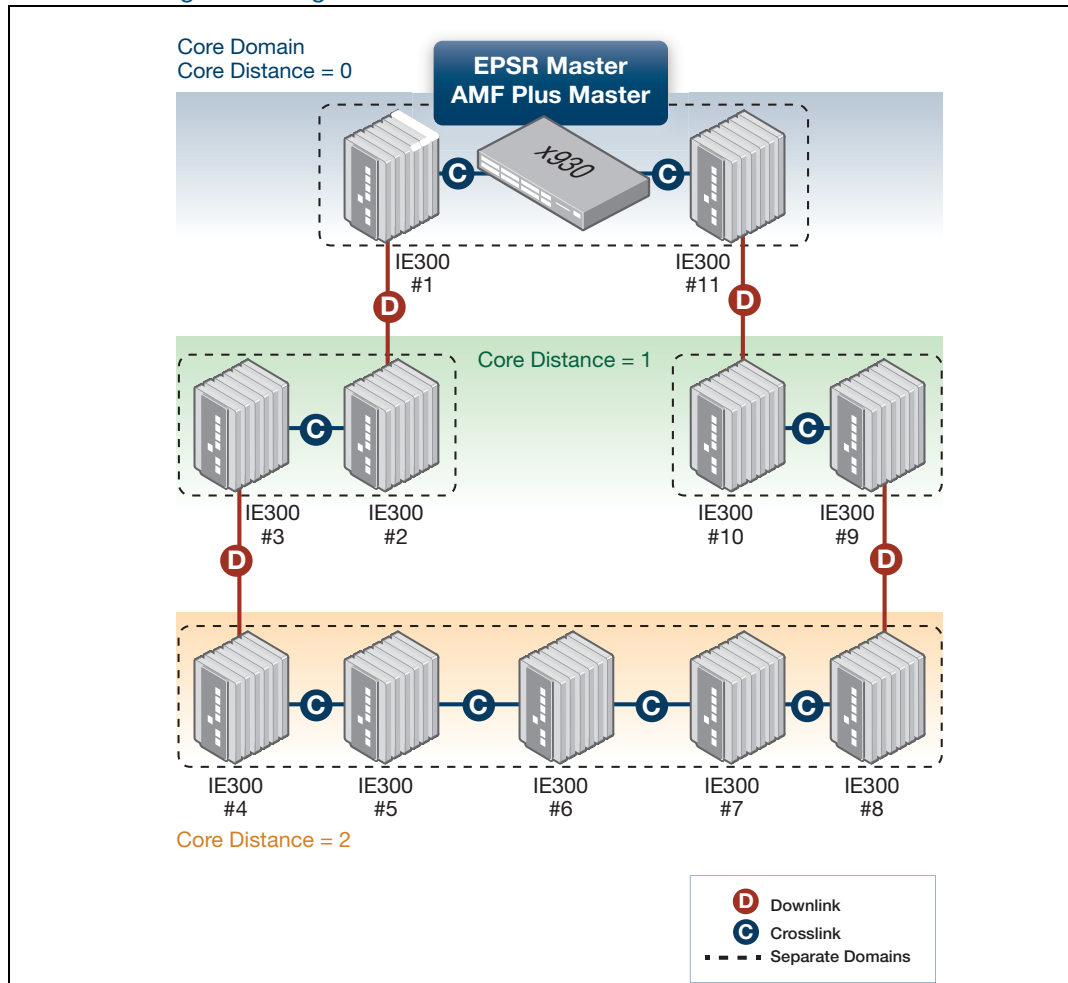
Extending solution 3 above, what if your EPSR ring contains more than 24 nodes? Can you nest to a second domain depth in order to maintain 12 or less nodes in each AMF Plus domain? Yes you can, but there are issues to consider.

Note: While it is always preferable to adhere to the maximum recommended domain limit of 12 nodes, in some networks where an AMF Plus domain only contains devices of sufficient CPU power it may be possible to safely exceed this limit by a small number of nodes. For any questions relating to specific scenarios please contact your authorised Allied Telesis representative.

The diagram below shows an incorrect AMF Plus configuration as a result:

- Transit nodes 2 and 3, and 9 and 10 are all at a core distance of 1 but form two separate domains (indicated by the green dotted lines).
- Transit nodes 4-8 are in the same AMF Plus domain at a core distance of 2. The two uplinks (AMF Plus down-links) from this domain each connect to different parent domains. This breaks the rule that an AMF Plus domain can only have a **single** parent domain.

Figure 28: EPSR ring containing more than 24 nodes



A possible workaround to this problem above is to remove the AMF Plus down-link between nodes 8 and 9. But the obvious downside is that if there is a physical break in the EPSR ring, then the AMF Plus master may lose connectivity with some nodes.

If for example, the link between nodes 5 and 6 fails:

- The EPSR master will unblock its secondary port and network connectivity will be maintained with all nodes in the ring.
- But the AMF Plus master will lose connectivity with nodes 6,7, and 8.

The only way to resolve this problem is to re-factor the EPSR ring to contain fewer nodes so that one of the previous solutions can be used.

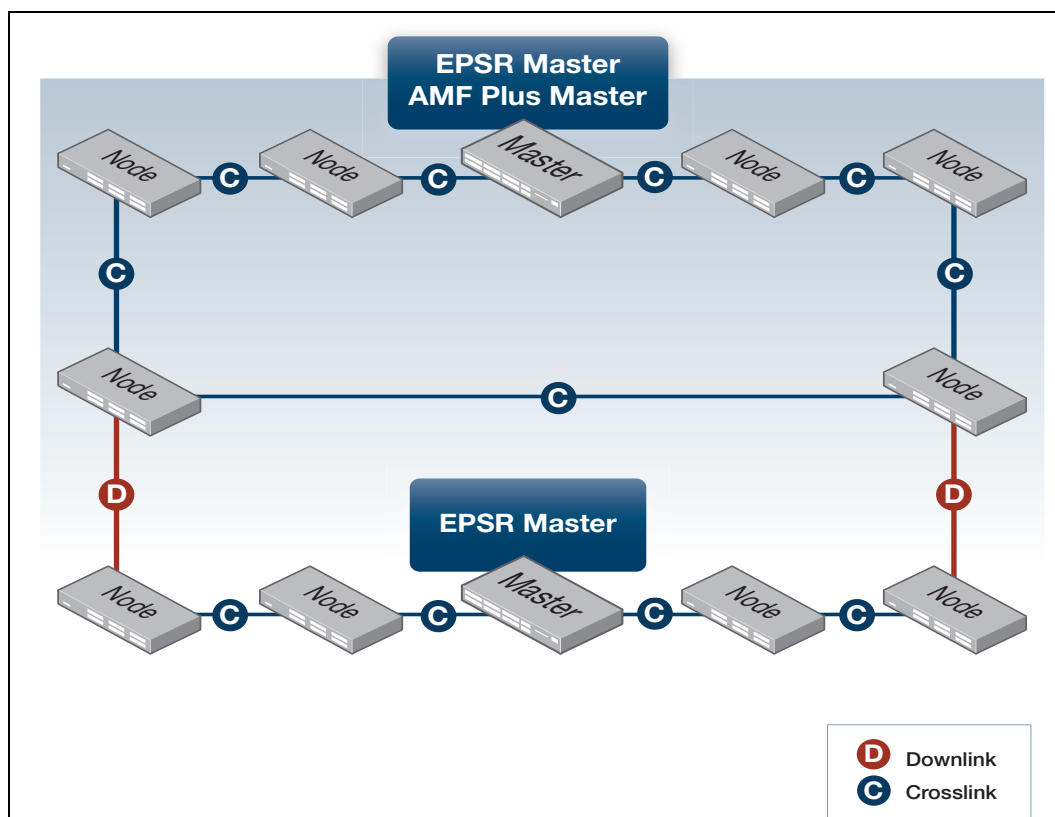
Dual-ring EPSR network with a common segment between two transit nodes

In this section we look at how best to configure AMF Plus in a scenario where you have a dual-ring EPSR network with a common segment between two transit nodes.

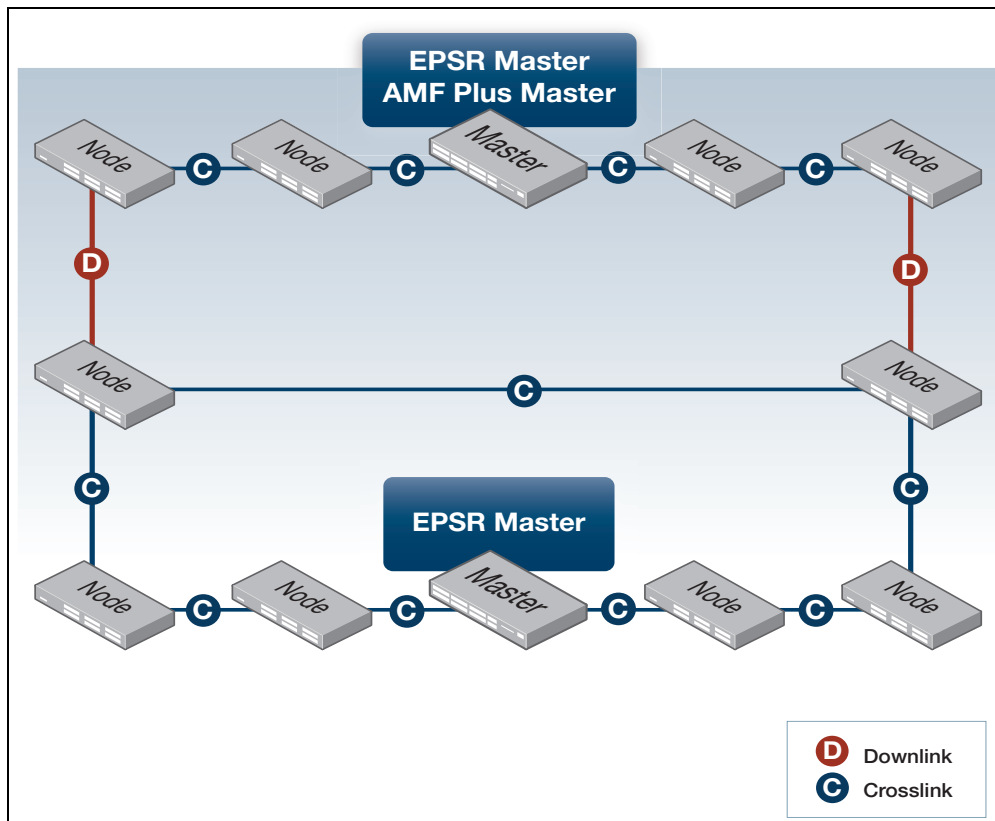
For this sort of topology there are several options. The important point to remember is that the transit nodes with the common segment can only have two cross-links. Note that only the **top** master in these diagrams is the AMF Plus master as well as the EPSR master.

Here are three variations to consider:

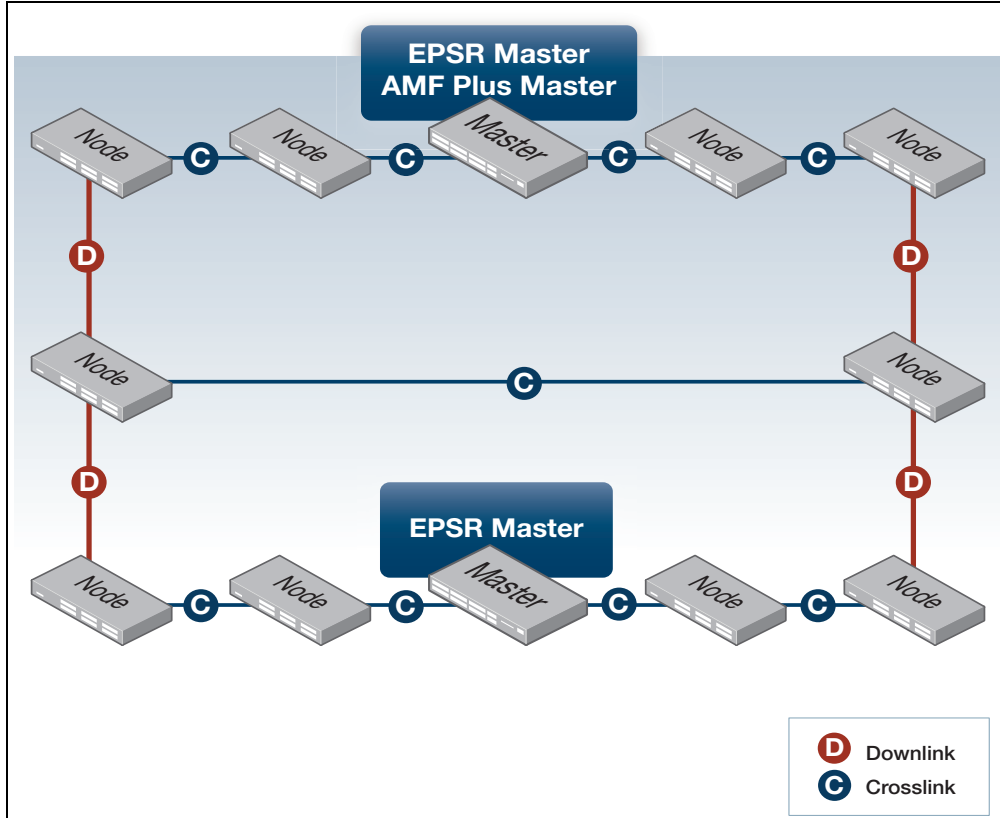
1. Core domain ring with down-links to chain



2. Core domain chain with down-links to ring



3. Core domain chain with down-links to chain domains



C613-22136-00-REV C



NETWORK SMARTER

North America Headquarters | 19800 North Creek Parkway | Suite 100 | Bothell | WA 98011 | USA | T: +1 800 424 4284 | F: +1 425 481 3895

Asia-Pacific Headquarters | 11 Tai Seng Link | Singapore | 534182 | T: +65 6383 3832 | F: +65 6383 3830

EMEA & CSA Operations | Incheonweg 7 | 1437 EK Rozenburg | The Netherlands | T: +31 20 7950020 | F: +31 20 7950021

alliedtelesis.com

© 2023 Allied Telesis, Inc. All rights reserved. Information in this document is subject to change without notice. All company names, logos, and product designs that are trademarks or registered trademarks are the property of their respective owners.