Technical Guide

Allied Telesis™

Feature Overview and Configuration Guide

Autonomous Management Framework™ (AMF)

AMF™

AlliedWare Plus™
OPERATING SYSTEM

alliedtelesis.com

C613-22047-00 REV T

Acknowledgments

This product includes software developed by the University of California, Berkeley and its contributors.

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. For information about this see www.openssl.org/.

This product includes software licensed under the GNU General Public License available from: www.gnu.org/licenses/gpl2.html

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.

Contents

Introduction to AMF

The Allied Telesis Autonomous Management Framework (AMF) is a suite of features that combine to simplify network management across all supported network equipment from the core to the edge.

AMF provides simplified device recovery and firmware upgrade management. The primary function of AMF is to reduce the management and maintenance overhead on a network, while improving its responsiveness in handling equipment failures.

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

AMF's features 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, together with its benefits, and presents configuration guidelines that explain the practical application of AMF in real networks.

Products and software versions that apply to this guide

This guide applies to AlliedWare Plus™ products running version **5.4.5-0.x** and later**.**

AMF master and controller nodes require a feature license. For information about which products can act as AMF master or controller nodes and the available licenses, see the AlliedWare Plus Datasheet and the Autonomous Management Framework Datasheet.

The following features are supported since the following software versions:

- AMF Guestnode - 5.4.6-0.x
- AMF Agent for integrating x600 Series Switches - 5.4.6-1.x
- AMF subscription licenses - 5.4.6-1.x
- Up to 300 nodes in a single area, if the AMF master is an AT-SBx81CFC960 or the Virtual AMF Appliance (VAA) - 5.4.6-2.x
- User authentication via a RADIUS or TACACS+ server - 5.4.6-2.x
- Multiple tenants on a Virtual AMF Appliance (VAA) - 5.4.7-0.x
- AMF secure mode - 5.4.7-0.3
- Deploying the VAA on Microsoft Azure - 5.4.7-1.x
- Auto-recovery of isolated nodes - 5.4.7-2.x
- Auto-recovery support on GS900MX/MPX series switches' VCStack ports - 5.4.7-2.x
- Multiple tenants on a Virtual AMF Appliance (VAA) installed on a public cloud - 5.4.7-2.x
- Recovery progress indicator added to UTM firewalls and VPN routers - 5.4.8-0.x
- Subscription license recovery with grace period - 5.4.8-0.x
- Recovery file maintenance - 5.4.8-0.x
- Secure virtual-links - 5.4.9-0.x
- Dynamic tunnel addresses for virtual-links - 5.4.9-0.x
- Node recovery using an equivalent devices - 5.4.9-0.x
- Node provisioning using several device types - 5.4.9-0.x
- AMF guest node support for ONVIF Profile Q devices - 5.4.9-2.x
- Support for AMF links on AR-series Eth interfaces - 5.5.0-1.x

Overview of an AMF 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 network**, in order to provide a picture of the context within which the elements exist.

Figure 1: AMF network overview

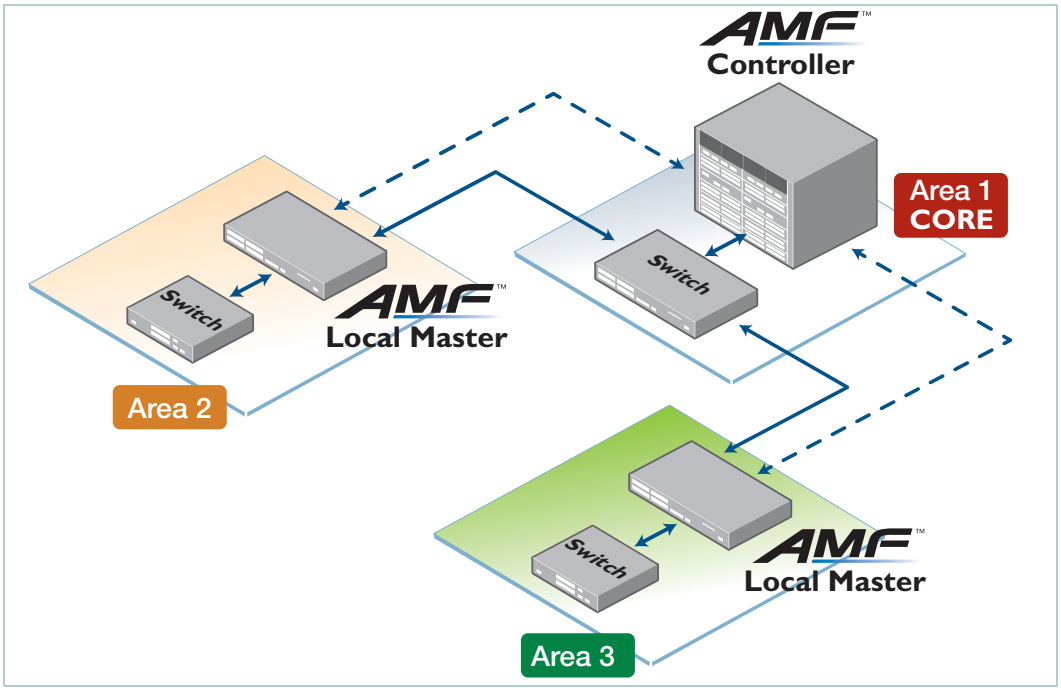

The primary structural entity in an AMF network is the **AMF area**.

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

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

A controller can connect to up to 60 AMF 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 network can be considered to be composed of two realms: the first realm is the management plane within each individual AMF 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 managed networks can grow up to 18,000 nodes.

Key benefits of AMF

The key benefits of AMF 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 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 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 Nodes: the Unified CLI" on page 53.

AMF remote login

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

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

Configuration backup and recovery

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

If an AMF member node should fail, the AMF 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 provides a complete zero-touch recovery solution. For more information, see "AMF Backups" on page 60. Similarly, AMF controller nodes can backup the master nodes of the AMF 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 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 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 108.

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 117.

Elements of AMF

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

AMF network

An **AMF 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 areas**.

Network name

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

Although each autonomous AMF network has a finite size of 60 areas, data transfer may occur between devices residing in different autonomous AMF 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 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 nodes

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

Controller node

An **AMF controller node** sits at the highest level of hierarchy in an AMF network. A node is designated as a controller by the command **atmf controller**, see "Configuring an AMF controller" on page 46.

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

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

Only one AMF area in the AMF network may contain controller nodes. Up to eight controller nodes can be created in this AMF 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 controllers per network for redundancy purposes.

Master node

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

- performing file system backups of all nodes in the AMF area.
- acting as a file server for firmware and configuration for the member nodes in its AMF area.
- providing an essential component for the formation of an AMF network. That is, an AMF 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 AMF area.

When more than one AMF master node exists in an AMF area, their operation is completely independent and unsynchronized. All master nodes within an AMF area must reside within the same AMF 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 area depends on which product is used as an AMF master, and can be up to 300 nodes.

Gateway node

Gateway nodes exist at the end of an AMF 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 network. The full management power and convenience of AMF is available on these switches, but they can only link to one other AMF node. They cannot form cross-links.

AMF guest node

AMF guest nodes are devices that either do not run the AlliedWare Plus operating system or run a version that does not support AMF. AMF guest functionality provides limited participation in an AMF network. AMF guest devices do not require any operating system modifications or have AMF software loaded onto them.

Parent node

A **parent node** is an AMF node that also directly connects to a specific AMF node. For example, if an access point is connected to an AMF network, then the AMF 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 2: AMF Uplinks, Downlinks, and cross-links

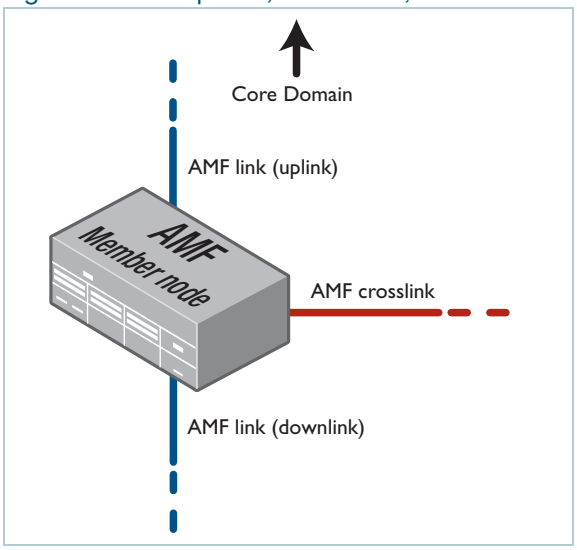

AMF links of either type are used to pass AMF management traffic between nodes; however, they can also carry other network traffic.

Cross-links

Cross-links are used to connect AMF nodes to other AMF nodes within what is termed an **AMF domain**. Configuring an interface as an AMF 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 master nodes must be connected to each other using AMF 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 domains. The tree of interconnected AMF 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 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 domain links upwards to only a single parent domain. Each domain may link downwards to multiple child domains.

AMF domains

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

As previously mentioned, AMF domains are determined by AMF cross-links. All nodes connected via AMF cross-links form part of the same domain, and nodes connected via up/down AMF 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 domain is that two links from a domain to a single higher level domain will provide redundant AMF links. We recommend that an AMF domain only be connected to a single higher level domain, though it may be connected 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 3 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 area is eight.

Figure 3: Core distance hop counts between domains

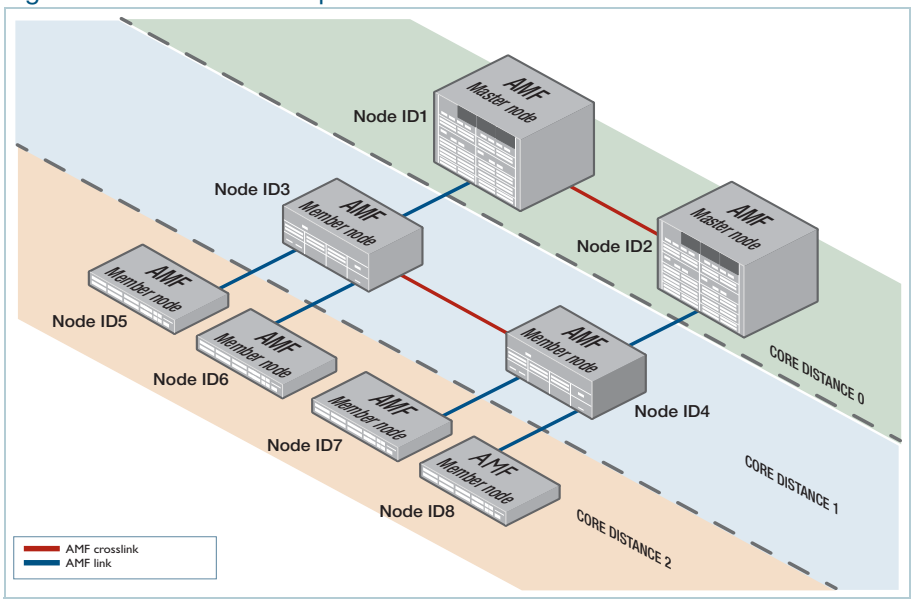

AMF areas

AMF is a highly scalable framework, designed to unify the management of very large networks. The inherent value of AMF 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 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 by dividing a network into regions known as **AMF areas**.

Conceptually, an AMF 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 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 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 link between two AMF 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 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-capable equipment within a site.

These sorts of non-contiguous connections within an AMF 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 node, except an AMF guestnode, can terminate a virtual-link. Virtual-links can be created between AMF nodal types such as:

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

The details of creating and optimizing virtual-links are described in the section"AMF Tunneling (Virtual-links)" on page 36.

Area links

The links between different AMF areas are termed **AMF area links**. These links may be just normal direct AMF links (i.e. AMF 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 AMF area. It is not possible to have an area link between two non-core AMF areas.

The details of configuring AMF area links are described in the section "Connections from AMF controllers to the other areas" on page 47.

Loop-free data plane

Currently (AlliedWare Plus software version 5.4.6-1.4), AMF ensures that its own data plane (i.e. the AMF VLANs) is kept loop-free. AMF however 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.

Basic AMF Configuration

This section presents a series of examples explaining how to configure an AMF network.

Example - Configuring a simple stand-alone area

The following configuration example uses a simplified network to explain the basic steps required to configure AMF.

Figure 4: Simple AMF single master example

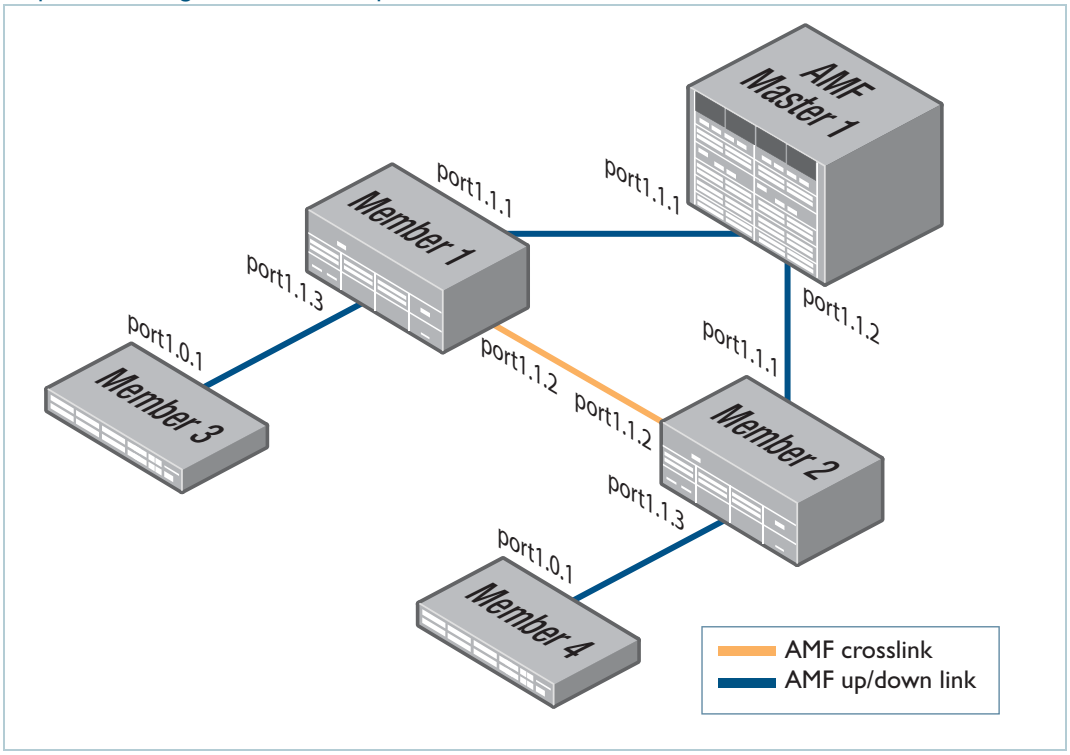

Configuring the AMF master node

Step 1: **Set the host name**

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

Note: Host names are used as the AMF node name and **must be unique** within the AMF area.

Step 2: **Set the AMF network name**

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

Note: The AMF network name must be the same on all nodes in all areas within the AMF network. More precisely, within the same **autonomous AMF network**. For more information see "Network name" on page 11.

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

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

- On standalone devices one master license is required per device. In the case of a SBx8100 chassis with dual CFC controller cards fitted - this is still considered a single device, and so only a single AMF master license is required. If the SBx8100 has two CFC cards installed, then the licence can be installed on one CFC card, and will be automatically copied over to the other CFC card.
- On VCStacks, prior to software version 5.4.5, an AMF master license was required for each VCStack member. From 5.4.5 onwards, the licensing has been relaxed, and only a single AMF master license is required per VCStack. However, this license needs to be installed on each 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-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-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 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**

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

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

Step 2: Set the AMF network name

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

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

Step 3: Configure the data VLANs

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

Step 4: Configure ports as AMF links

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

Step 5: Configure data VLANs on the AMF 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 cross-link

```
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
```

Note: AMF 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 management VLAN.

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

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

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

Step 2: **Set the AMF network name**

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

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

Step 3: **Configure the data VLANs**

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

Step 4: **Configure ports as AMF-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 Network

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

Output 1: Checking AMF configuration using the show atmf summary command

```
AMF_Master#show atmf summary
ATMF Summary Information:
ATMF Status            : Enabled
Network Name           : atmf1
Node Name              : AMF_Master
Role                   : Master
Current ATMF Nodes     : 5
AMF_Master#
```

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

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

Output 2: 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          x510-52GTX           N        S   Member2          2
  Member3          x510-52GTX           N        S   Member2          2
 Current ATMF node count 5
```

Note: 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 links:

Output 3: Checking output with the show atmf links command

```
switch1# show atmf links

ATMF  Links Brief:

Local     Link       Port     ATMF          Adjacent            Adjacent    Link
Port      Type       Status   State         Node                Ifindex     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.

Prior to version 5.4.6-2 any user account used for managing an AMF network had to exist in the local user database of all devices in the AMF network. This is no longer a requirement and AMF now supports remote login authentication via either RADIUS or TACACS+.

Because AMF'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 network. Thereafter AMF uses a key exchange mechanism. This means the default AAA authentication method group of **local** is compatible with AMF. For AMF 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

AMF uses NTP to synchronize the system clocks across nodes within the network. All AMF nodes automatically receive time from the AMF master's NTP server. For this to operate you need to configure at least one external NTP server on each AMF master in your network to ensure accurate logging, and consistent timestamps between all AMF nodes. Configuration of three or more NTP servers is considered best practice. Configured servers do not need to be the same between AMF 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 masters may not have a path to an external NTP server. This may be due to the AMF masters and core of the network being locked down with no Internet access. If this is the case a local NTP server, or AMF node which does have Internet access, can be configured as the desired NTP server.

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

The primary function of NTP within an AMF network is to ensure that date stamps on backups are consistent across member nodes. In an AMF network that has multiple AMF 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 nodes.

Configuring NTP on the AMF network

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

Note: AMF masters act as NTP peers of each other, all other nodes use the AMF masters as NTP servers. This happens automatically; 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 4: 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, Aug 18 2016 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 Links

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

AMF requires that a completely unconfigured node will successfully form an AMF connection and become integrated into the network when attached to an AMF 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 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—by using the **lacp global-passive-mode enable** command—we do **not** recommend using a mixed configuration in an LACP channel-group; i.e. where some links are manually configured and others are dynamically learned in the same channel-group.

Sharing AMF links with other network operations

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

However, despite the special significance of AMF links, they are not used exclusively for AMF 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 does ensure loop-free operation of the AMF management VLANs that operate over its AMF 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 either

configuring EPSR, or by using the spanning tree protocol. In this respect the following should be noted:

- AMF coexists with spanning tree - so spanning tree will operate on AMF links without adversely affecting the operation of the AMF management VLANs.
- There is no restriction regarding the use of EPSR with AMF. EPSR rings can coexist on ports that are also configured with AMF links. See "Using AMF in EPSR Rings" on page 179 for information on supported ESPR topologies.

Reserved IP address range

Some of the AMF-related communication that occurs between AMF nodes is in the form of IP traffic. A class-B subnet is reserved for the use of this AMF-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 and should be used for no other purpose. AMF actually further divides this subnet into two /17 subnets, used for different purposes:

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

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

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

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

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

The new management subnet will not become effective until all members of the AMF 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 on VCStacks

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

AMF links on AR-series Eth interfaces

From AlliedWare Plus version 5.5.0-1.1 onwards, AMF up/down links and AMF area links are supported over an AR-series device's Eth interface. This enables you to provision and recovery AMF nodes over these interfaces.

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

Use the **atmf-link** and **atmf-arealink** commands on an Eth interface to configure it as an AMF link. For example, to configure an AMF 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 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 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 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 management purposes - by default 172.31.0.0/16
- packets with protocol type 0xfbae.
- BPDU packets 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 enabled the number of ACLs available on the x230, x310, x510, x550, x930, GS900MX/MPX, XS900MX, x950, and SBx908 GEN2 switches decreases by 1. If you are not using AMF you can disable AMF to reclaim the additional ACL.

AMF interaction with STP on IE200 and AR-series devices

On AR-Series devices and IE200 series switches, if you use STP at the same time as AMF, you may find that AMF 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 link will fully synchronize and correctly forward AMF traffic in both directions, but there may be a delay of a few seconds before it does so. This could interrupt any AMF working-set operation that is in progress at the time.

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

Renaming your AMF network

To rename the AMF network, use the command:

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

In software versions prior to 5.4.6-1.x, if you renamed the AMF network on an AlliedWare Plus device, we recommended you reboot your device. Version 5.4.6-1.x **removed** the need to reboot.

AMF Subscription Licenses

To use AMF, you must have an AMF subscription license (since version **5.4.6-1.x**). To see the available licenses, check your device's datasheet, which is available at alliedtelesis.com.

AMF starter license

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

Managing AMF licenses

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

Step 1: **Obtain the serial number for your AMF 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 5: 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 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 6: 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 Apr 2016 00:00
    Expiry date:                  19 Apr 2017 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 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.

Node licensing prior to 5.4.6-1.x

If your network is running an earlier version than 5.4.6-1.x, you need to purchase and install an AMF feature license. Contact your authorized Allied Telesis sales representative for more information.

For the case where the SBx8100 is the AMF master, only one AMF master or controller license is required, even if two CFCs (Controller Fabric Cards) are installed. If the SBx8100 has two CFC cards

installed, then the license can be installed on one CFC card, and will be automatically copied over to the other CFC card.

For the case where another AlliedWare Plus switch is the AMF master and the AMF master is a VCStack, an AMF master license is required on all devices in the stack. A stack will form successfully even if the master license is present on some stack members, but not others. However, if a master failover event occurs, and the new master unit in the VCStack happens to be one that does not have the AMF master license installed, then the stack will cease to operate as an AMF master. So, it is highly advisable to ensure that before a VCStack goes live as an AMF master, the AMF master license is installed on all stack members. See the Licensing Feature Overview and Configuration Guide for details.

When more than one AMF master node exists in an AMF area, it is important to know that these master nodes operate completely independently of each other, and there is no synchronization between them. For redundancy, an AMF area can have multiple master nodes, each acting as a master for the area. However, because there is no synchronization of status or data files between the masters, the behavior of a master node is not affected by the presence of other master nodes.

AMF Tunneling (Virtual-links)

AMF 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 tunnel is shown in Figure 5 on page 36. Switches 1 and 2 encapsulate the Layer 2 AMF 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) perform a conventional routing function, reading the IP addresses of the tunneled packets and forwarding them to their destination.

Once connected through the tunnel, the remote AMF members will have the same AMF capabilities as a directly connected AMF member.

Figure 5: AMF virtual-link

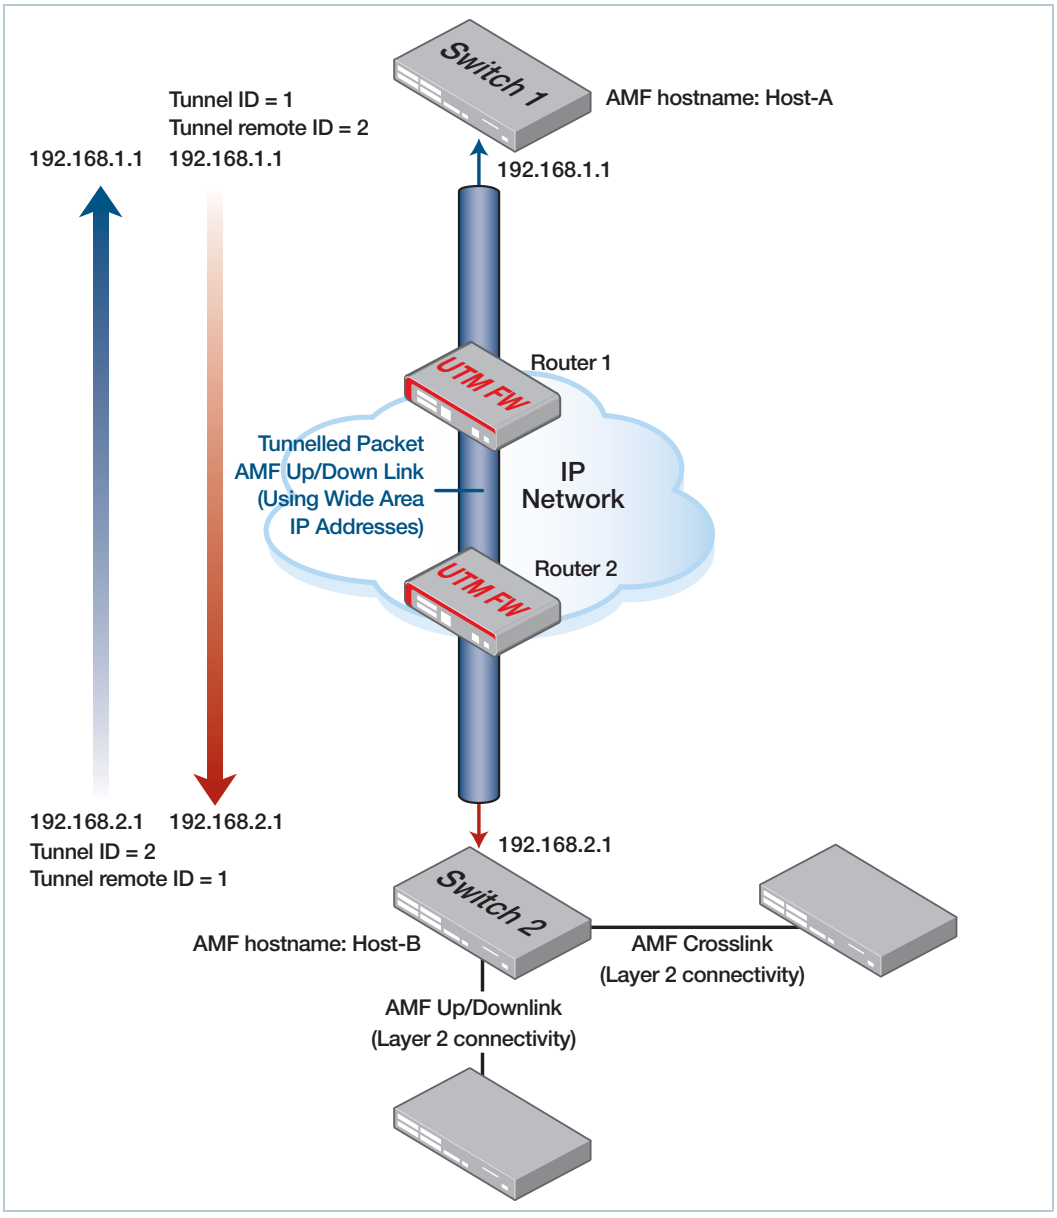

Configuring a virtual-link

Configuration involves specifying the following:

- local tunnel ID
- local IP address
- remote tunnel ID
- remote IP address

The Layer 2 tunnel, created by the command **atmf virtual-link id ip remote-id remote-ip**, enables a local AMF session to appear to pass transparently across a Wide Area Network (WAN) such as the Internet.

The addresses configured as the local and remote tunnel IP addresses 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 involve using some type of managed WAN service such as a site-to-site VPN. Tunnels are only supported using IPv4.

A reciprocal configuration is also required on the corresponding remote device. The 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. So, the tunnel with ID 10 on the switch at one end of the virtual-link could connect to the tunnel with ID 40 on the switch at the other end of the link.

The tunneled link may pass through external (non AMF capable) routers in order to provide wide area network connectivity. However, in this configuration, these devices perform conventional IP forwarding of the tunneled packets based on the content of the IP headers in their IP encapsulation. The protocol tunneling function is accomplished by the AMF nodes at each end of the virtual-link.

Note that the requirement to preconfigure the local IP address and tunnel ID on a device located at the far end of an AMF virtual-link tunnel, means that zero touch device replacement of a remote device that terminates the tunnel cannot be achieved by delivering backed up files from a AMF 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 member nodes" on page 60

Since version 5.4.7-2.x, you can recover isolated nodes at the far end of an AMF virtual-link tunnel by using the method described in the section "Auto-recovery and Provisioning of Isolated Nodes" on page 103.

Example

Use the following commands to create the tunnel shown in figure "AMF virtual-link" on page 36.

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
```

Caution On an IP interface that is carrying AMF virtual-link traffic, do not set the MTU (Maximum Transmission Unit) to less than its default value of 1500 bytes.

Secure virtual-links

From AMF version 5.4.9-0.1 onwards, you can create secure AMF virtual-links by encapsulating the L2TPv3 frames of the virtual-link with IPsec. This feature is only supported on some AlliedWare Plus devices, see "Secure virtual-links limitations" on page 38.

Secure virtual-links make it possible for your AMF 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 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 members:

Switch devices support a maximum of 20 downstream AMF nodes when using a secure virtual-link as an uplink.

Secure virtual-links behind NAT:

When there are two or more AMF members behind a shared NAT device, only one of the members will be able to use secure virtual-links.

AMF Multi-tenant environment:

An AMF Multi-tenant environment supports a maximum cumulative total of 1200 secure virtual-links across all AMF containers.

Supported devices

Secure virtual-links are only supported on the devices listed in the following table. There is also a limit to the number of links these devices support.

Table 1: Supported number of protected virtual-links

DEVICE	VIRTUAL-LINK LIMIT
AMF Cloud / VAA	300
AR4050S AR3050S AR2050V AR2010V	60
x220 x230/x230L x310 x510/x510L IX5-28GPX	2

Example Configuration

Step 1: **Configure an AMF 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
```

Virtual-links with dynamic IP addresses

From version 5.4.9-0.1 onwards, you can configure an AMF virtual-link by using a dynamic local and/or a dynamic remote tunnel address.

Configuring a virtual-link with a dynamic local IP address

If an interface acquires its IP address dynamically then the local side of the tunnel can be specified by using the interface's name instead of using its IP address. When using a dynamic local address the remote address of the other side of the virtual-link must be configured with either:

- the IP address of the NAT device the dynamically configured interface is behind, or
- 0.0.0.0, if the virtual-link is configured as a secure virtual-link, see "Configuring a virtual-link with a dynamic remote IP address" on page 41).

Figure 6: Virtual-link with one of the devices behind a NAT device with a static IP address.

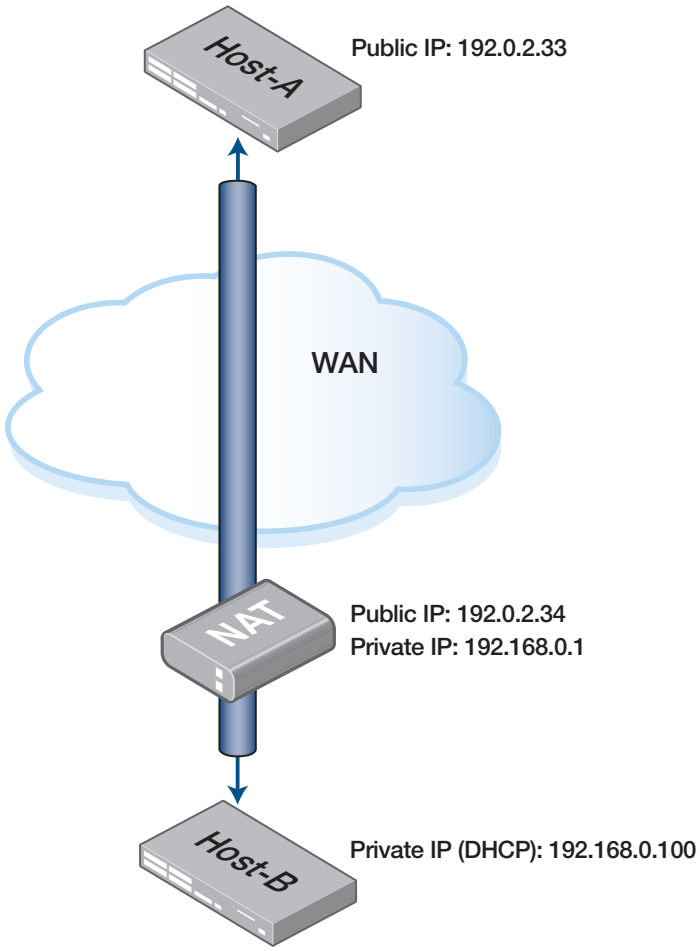

For example, if you wish to create a virtual -link between Host A and Host B where Host B is behind a NAT device with a known static public IP address, use the following commands:

On Host-A

```
Host-A(config)# atmf virtual-link id 1 ip 192.168.0.33 remote-id 2 remote-ip 192.0.2.34
```

On Host-B

```
Host-B(config)# atmf virtual-link id 2 interface port1.0.1 remote-id 1 remote-ip 192.0.2.33
```

Configuring a virtual-link with a dynamic remote IP address

Using a dynamic address for the remote side of an AMF virtual-link is only available on secure AMF virtual-links. When the IP address of the remote side of a secure AMF virtual-link is unknown you configure an AMF virtual-link by specifying a dynamic address for the remote side. This is done by setting the remote-ip parameter to **0.0.0.0**.

Figure 7: Secure virtual-link with one of the devices behind a NAT device with a dynamic IP address.

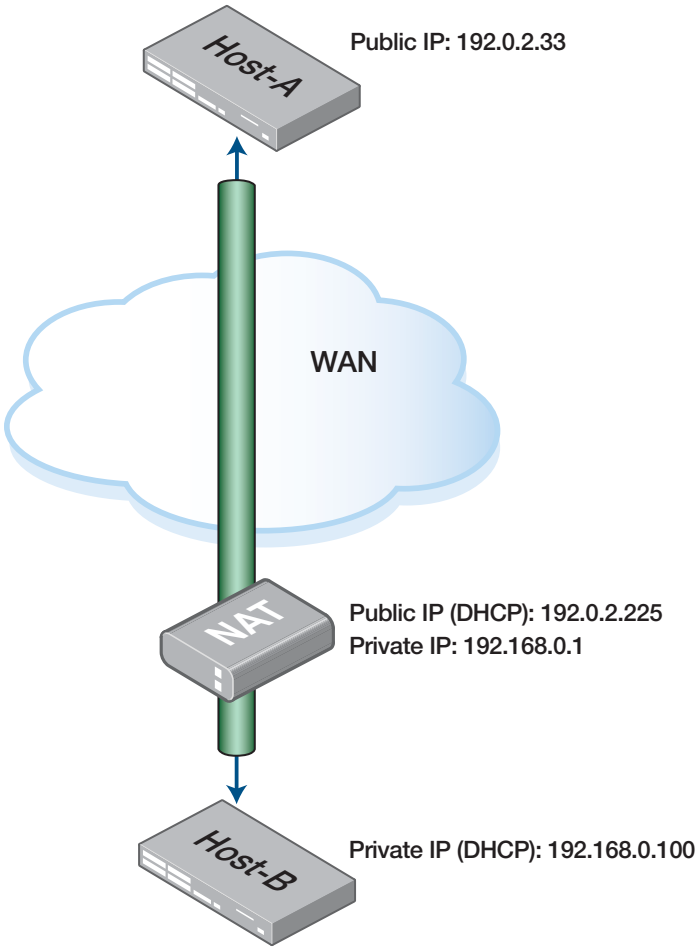

For example, if you wish to create a virtual-link between Host A and Host B where Host B is behind a NAT device with an unknown public IP address, use the following commands:

On Host-A

```
Host-A(config)# atmf virtual-link id 1 ip 192.168.0.33 remote-id 2 remote-ip 0.0.0.0
Host-A(config)# atmf virtual-link id 1 protection ipsec key secure-password
```

On Host-B

```
Host-B(config)# atmf virtual-link id 2 interface port1.0.1 remote-id 1 remote-ip 192.168.1.1
Host-B(config)# atmf virtual-link id 2 protection ipsec key secure-password 192.0.2.33
```

Note: A dynamic remote address cannot be used on both sides of a secure AMF virtual-link.

Prioritizing the tunneled traffic

The traffic that is carried in the tunnel includes AMF 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 "AMF virtual-link" on page 36, 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 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 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 a VAA (Virtual AMF Appliance). In this situation a single virtual cross-link can be created to enable a master residing on a VAA to share the AMF master role with an AMF master running on a physical device.

Note: Creating virtual cross-links between container masters (in a multi-tenant VAA installation) and physical masters is not supported.

The Concept of AMF Areas

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

The inherent value of AMF 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 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 by dividing a network into regions known as AMF areas.

Each AMF 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 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 controller

To set up a node as an AMF controller it first needs to have an AMF 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 <1-126> [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 <1-126>
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 67.

Connections from AMF controllers to the other areas

For an AMF controller to communicate with the other areas that comprise the network, it needs AMF 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 <1-126> 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 <1-126>
```

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 Tunneling (Virtual-links)" on page 36, 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 <1-126> 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

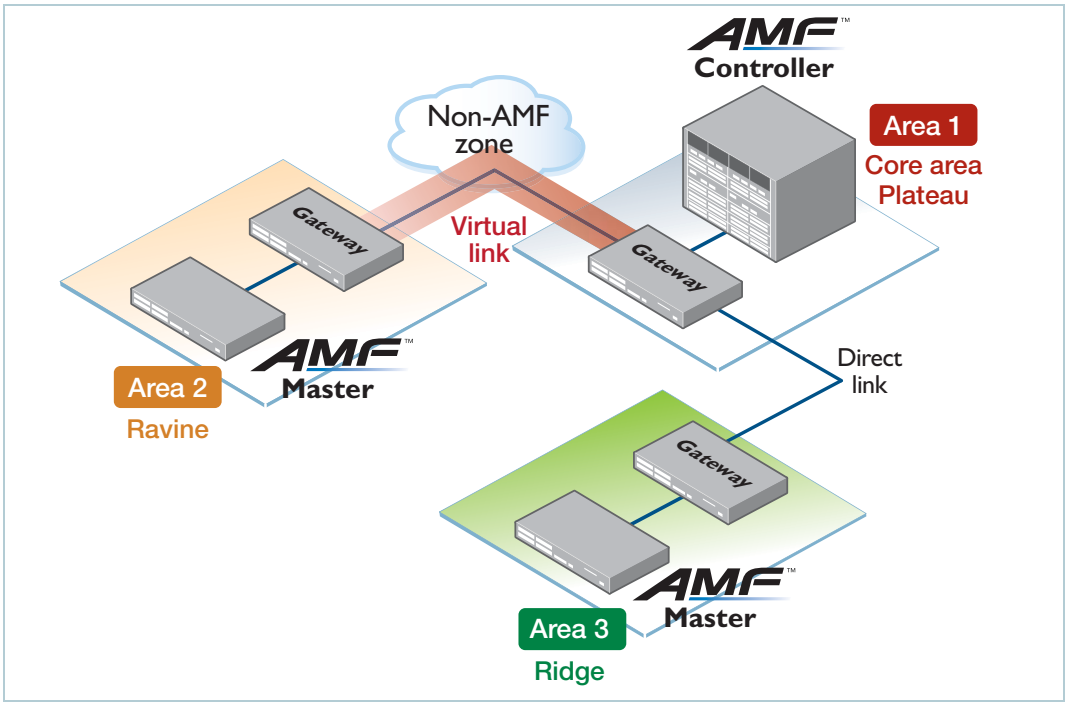

The AMF and related configurations for the six nodes illustrated in this diagram are:

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 E1phMJTGVBWuCXcv4xtM19cBE+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 E1phMJTGVBWuCXcv4xtM19cBE+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 E1phMJTGVBWuCXcv4xtM19cBE+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

From AlliedWare Plus version 5.5.0-1.1 onwards, AMF area links are supported over an AR-series device's Eth interfaces.

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

Use the **atmf-arealink** command on an Eth interface to configure it as an AMF area link. For example, to configure the Eth1 interface as an AMF 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 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 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 areas are only available if every node in your network is running version 5.4.6-2.1 or later.

Configuring AMF Nodes: the Unified CLI

The unified CLI is a central component of AMF. It provides you with a configuration and display interface that can control a selected collection of nodes, or the entire AMF 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 network (unless **restricted-login** is enabled). All the members of a working-set must reside in the same AMF area. It is **not** possible for a working-set to span areas. Figure 8 below shows a number of switches that comprise a working-set.

Figure 8: AMF working-set

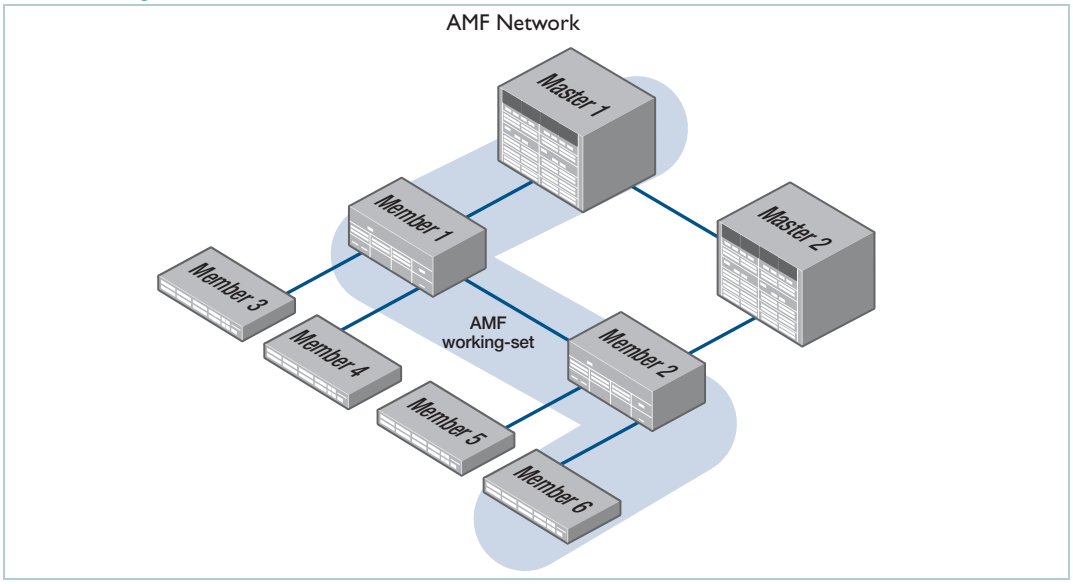

Note: For security reasons **restricted-login** limits the action of working-sets on the AMF network. See "AMF restricted-login" on page 126 for more information.

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 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 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 atmf1 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 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 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 network:

1. **All**—all nodes within the AMF 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 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 area, use the command **show atmf group** with the **automatic** parameter:

Output 7: show atmf group members automatic

```
x908_VCS_1#show atmf group members automatic

Retrieving Automatic groups from:
x510_1 Master x908_VCS_2 x908_VCS_1

ATMF Group membership

Automatic        Total
Groups           Members   Members
poe              1         Master
x510             1         x510_1
SBx8100          1         Master
x900             2         x908_VCS_2 x908_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 x510_1, x510_2 and all nodes in the group named x900, use the following command:

```
x908_VCS_1# atmf working-set x510_1,x510_2 group x900
x510_1, x510_2, x908_VCS_1, x908_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 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 area with the command **show atmf group members user-defined**:

Output 8: show atmf group members user-defined

```
x908_VCS_1#show atmf group members user-defined

Retrieving Automatic groups from:
x510_1 Master1, Master2, x908_VCS_2 x908_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 9: 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

===========
x908_VCS_1:
===========

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

===========
x908_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 10: 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, x908_VCS_1, x908_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 networking using the copy command and adding the AMF node name suffixed with ".atmf" to the path.

For example to copy a file named "x550-5.4.8-0.1.rel" from the current directory on the AMF master to flash on an AMF node named "node1" use the command:

```
master# copy x550-5.4.8-0.1.rel node1.atmf/flash:
```

AMF Backups

AMF backups are a valuable part of AMF network operation. Backups ensure that appropriate devices within the AMF network have copies of other devices' information and files. This means that if a node fails, the AMF 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 area.

They are the mechanism by which AMF master nodes update their records of their AMF area, and so have all the information and files required to enable automatic node recovery. By default, AMF master nodes are configured to perform automatically scheduled backups of the entire AMF area once per day at 3:00 a.m. AMF 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 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 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 uplink to its higher-domain neighbor.

If a device has a direct AMF uplink to its higher-domain neighbor and needs to be replaced; all the replacing device needs to do is send AMF 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 devices between itself and the other end of the virtual-link. So, if it sends out AMF messages towards the master, they will simply be dropped by the non-AMF devices in between.

Figure 9: Backups by member nodes

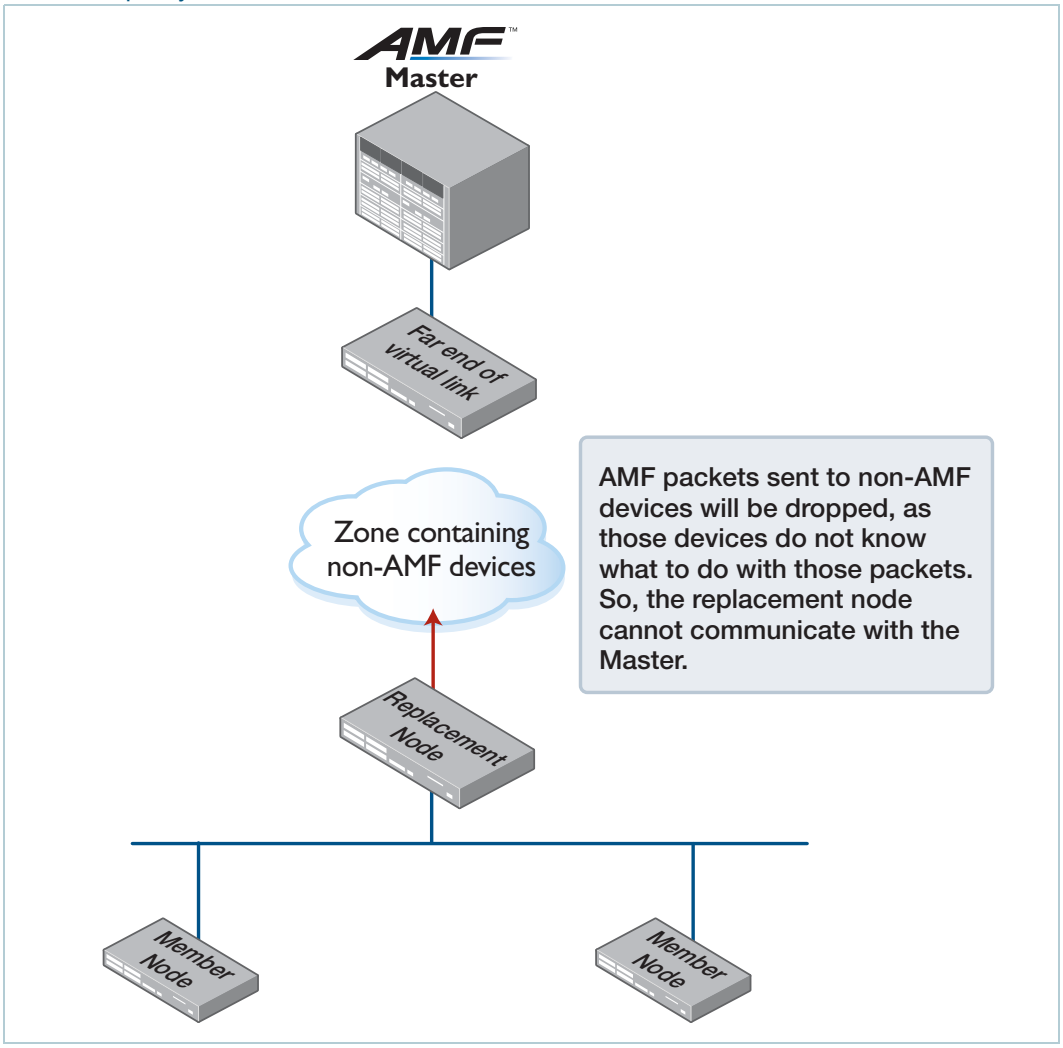

An alternative solution is required to enable recovery in this situation. The solution is to enable the member node neighbors of the virtual-link endpoint node to perform limited backups of the virtual-link endpoint node.

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 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 virtual-link becomes operational and the recovering node can now connect to the area master.

Figure 10: Recovery step 1

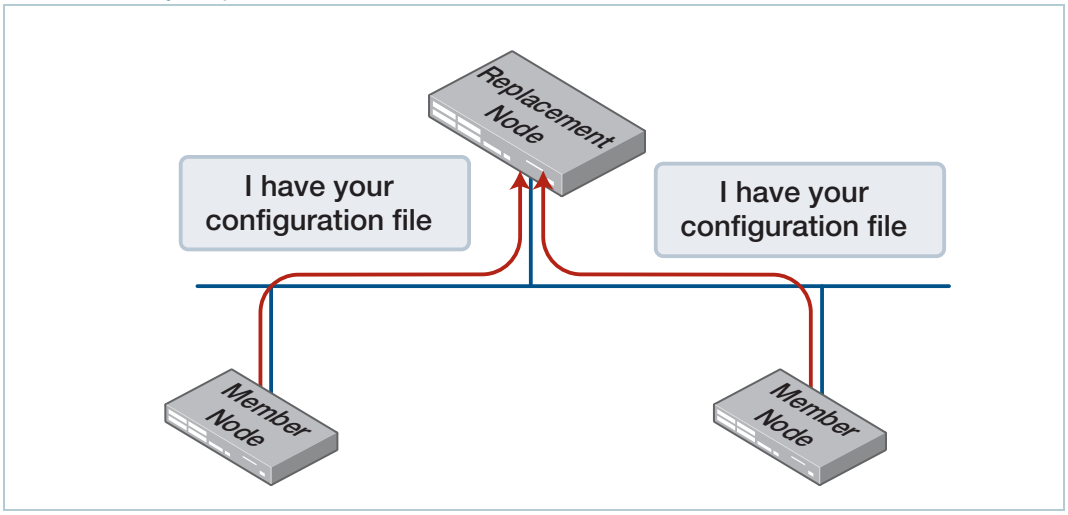

Figure 11: Recovery step 2

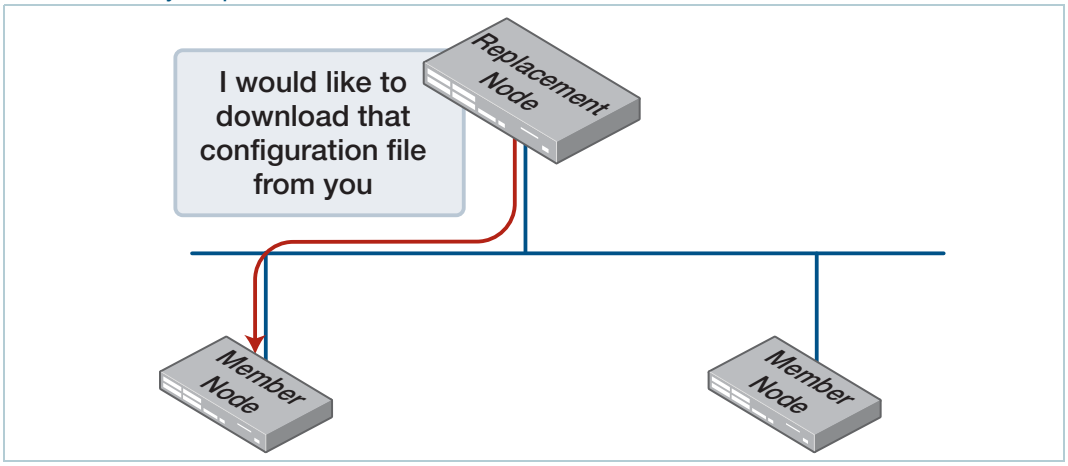

Figure 12: Recovery step 3

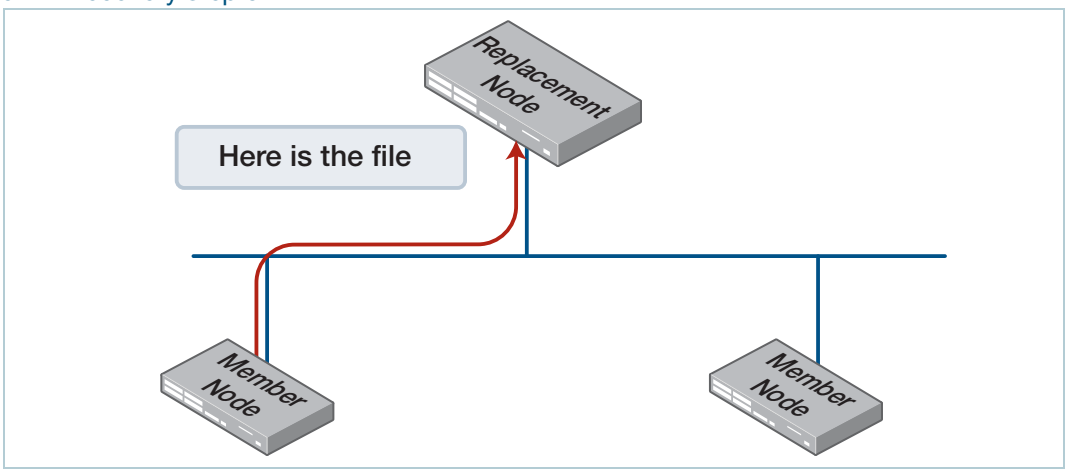

The recovering node then reboots and comes up again running the configuration that includes the definition of the virtual-link. It re-establishes the virtual-link, and communicates with the master to complete the recovery.

Figure 13: Recovery step 4: Re-establishing the virtual-link

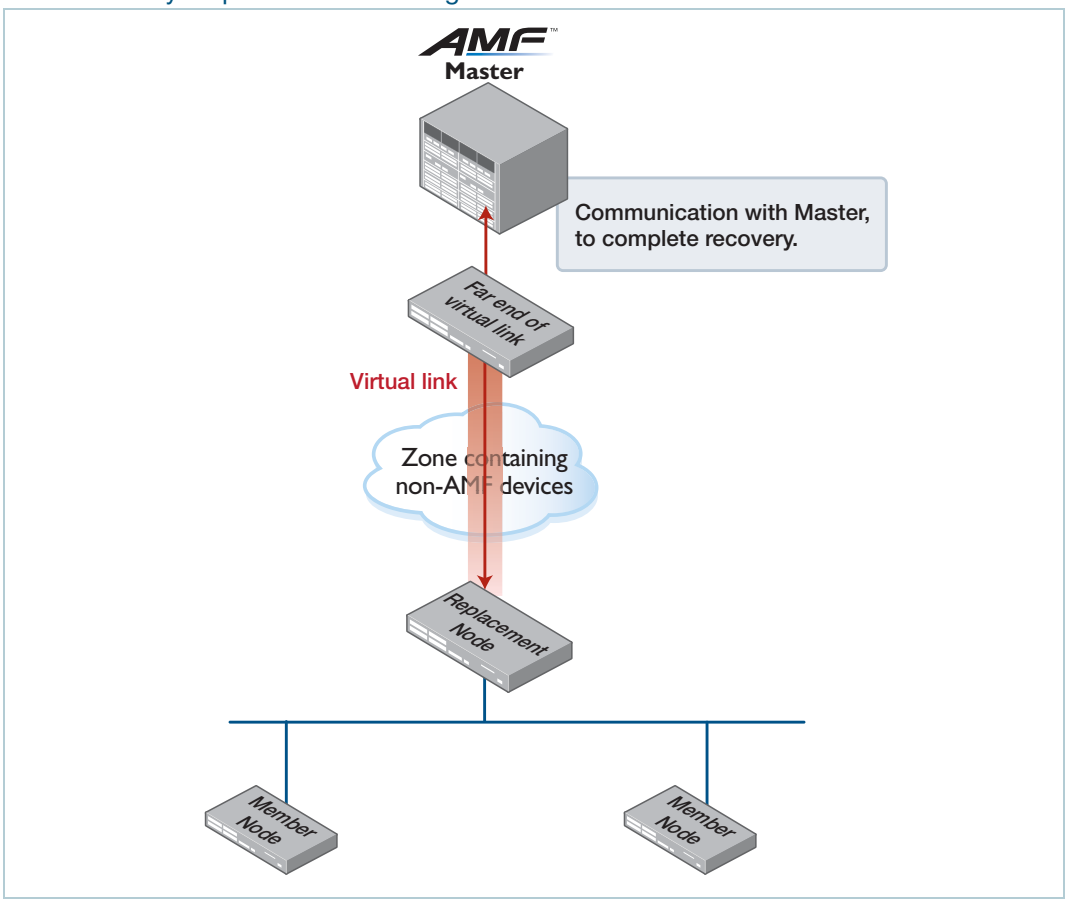

Events that trigger the pushing of the configuration file to neighbor nodes

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 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 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 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 Areas.

These are discussed below in the section "Multiple backup destinations" on page 78.

If an AMF master is storing backup data on removable media then:

- if the AMF master is a SBx8100 system with dual CFC controllers, both CFCs should have removable media installed.
- if the AMF 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 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 area.

  You can store other data on the storage device as long as you make sure that enough space is reserved for future AMF backups.

  AMF 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 11: Output showing backup media space warning

```
master1#show atmf backup

Scheduled Backup ...... Disabled
  Schedule ............ 1 per day starting at 12:45
  Next Backup Time .... 25 May 2014 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 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 backup files or file system, we recommend that the following procedure be followed before rebooting or removing any removable storage media from an AMF 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 12: 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 .... 25 May 2014 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 network as a provisioned node, we recommend that you perform a manual backup from the AMF master.

To perform a manual backup of the entire AMF area, on the AMF 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 node by running the following commands on the AMF master:

```
master1# atmf backup now <node-name>
master1(config)# atmf backup enable
master1(config)# exit
```

To check the status of the AMF backup, use the **show atmf backup** command. The “Date”, “Time”, and “On Media” details update once the backup for that node is finished.

Output 13: 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 Dec 2013 03:00
Backup Media .......... USB (Total 3692.6MB, Free 1782.7MB)
Current Action ........ Doing manual backup
  Started ............. 13 Dec 2012 05:20
  Current Node ........ Member1
Backup Redundancy ..... Disabled

-------------------------------------------------------------------------------
Node Name                Date         Time      In ATMF  On Media  Status
-------------------------------------------------------------------------------
AMF_Master               13 Dec 2012  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 14: 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 .... 13 Dec 2013 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 Dec 2013  05:20:16  Yes      Yes       Good
Member1                       13 Dec 2013  05:20:27  Yes      Yes       Good
Member2                       13 Dec 2013  05:20:40  Yes      Yes       Good
Member3                       13 Dec 2013  05:20:52  Yes      Yes       Good
Member4                       13 Dec 2013  05:21:08  Yes      Yes       Good
```

Note that the file system used by the AMF 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 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 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 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 72.

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 master nodes in an AMF 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 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 master nodes in an area, first manually select all AMF 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 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 backup node.
- A public user key for root from the AMF 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 15: Importing a host key from remote file server

```
x908(config)#crypto key pubkey-chain knownhosts ip 10.37.165.65 rsa

10.37.165.65 ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAABAQD4LkpjHYXHrFU6/
0vS2PTsIkvKh8j0JkIwIMiBHAxJVadED73G6iq4C6Lk
NB+M7BZOohF/
ow0Bhp4Jr8XK0MfM09gfriRHNNQlGsbfPugDKXnJpFMU88Yu3MaSrkGULgtD7B0MXBEw05H
PNjo1RCR9KI3Z3GFGM1TJy8T/6
xikczyaxbhqfUeqtpMgDMzRhieqIdpl7Umg4fJxhMDSha8af0HrpRpntsw23+h5IUX9Sw+p
G9F1zxczncM1BsKQ579iYA0Ek+pWiFlxK2lziO
86oIkYrlcsnHmcYKjrO/9GIlSFSAm6v2bBnXMh6wzcpl0A+6TAU4Bp9c7WNq4KlU0x
Are you sure you want to add this public key (yes/no)? yes
x908(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 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 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 AAAAB3NzaC1yc2EAAAADAQABAAAAgQDFRNJeSqqMmMesU3wF6RM/
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 controller/master

Once the keys have been exchanged, the AMF 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 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 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 78, 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 78, 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.

Note: 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 16: Output from the **show atmf backup** command - configuration of two remote backup file servers

```
x908#show atmf backup
  Scheduled Backup ...... Enabled
    Schedule ............ 24 per day starting at 14:25
    Next Backup Time .... 19 May 2014 11:25
  Backup Bandwidth ...... Unlimited
  Backup Media .......... FILE SERVER (Total 503837.5MB, Free 186818.0MB)
  Server Config .........
    Synchronization ..... Synchronized
      Last Run .......... 19 May 2014 11:09:50
     1 .................. Configured   (Mounted)
      Host .............. 10.36.150.54
      Username .......... user_1
      Path .............. temp/x908_1
      Port .............. -
   * 2 .................. Configured   (Mounted, Primary)
      Host .............. tb165.test.com
      Username .......... user_2
      Path .............. temp/x908_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
  ------------------------------------------------------------------------
  x210a                    19 May 2014  11:09:37  Yes      Yes       Good
                           19 May 2014  11:09:46  2        1         Good
  x930a                    19 May 2014  11:09:17  Yes      Yes       Good
                           19 May 2014  11:09:19  2        1         Good
  x930b                    19 May 2014  11:09:49  Yes      Yes       Good
                           19 May 2014  11:09:49  2        1         Good
  x930c                    19 May 2014  11:09:20  Yes      Yes       Good
                           19 May 2014  11:09:20  2        1         Good
  x930d                    19 May 2014  11:09:19  Yes      Yes       Good
                           19 May 2014  11:09:19  2        1         Good
  x908                     19 May 2014  11:09:49  Yes      Yes       Good
                           19 May 2014  11:09:50  2        1         Good
  x908stk                  19 May 2014  11:09:47  Yes      Yes       Good
                           19 May 2014  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 17: **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 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**.

Node Recovery

Automatic node recovery

AMF allows you to replace a failed node with another device and let AMF 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 85)

Note: If your recovering device is a GS900MX/MPX series switch where the VCStack port has been used as an AMF link see "Recovering a GS900MX/MPX series switch" on page 86

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 117.

Note: From AlliedWare Plus version 5.4.9-0.1 onwards, 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 an equivalent model" on page 82).

When a switch boots up, it follows the process shown in the flowchart of Figure 14 to determine what configuration to use. This flowchart indicates when automatic node recovery is successful.

Figure 14: How a switch determines which configuration to use

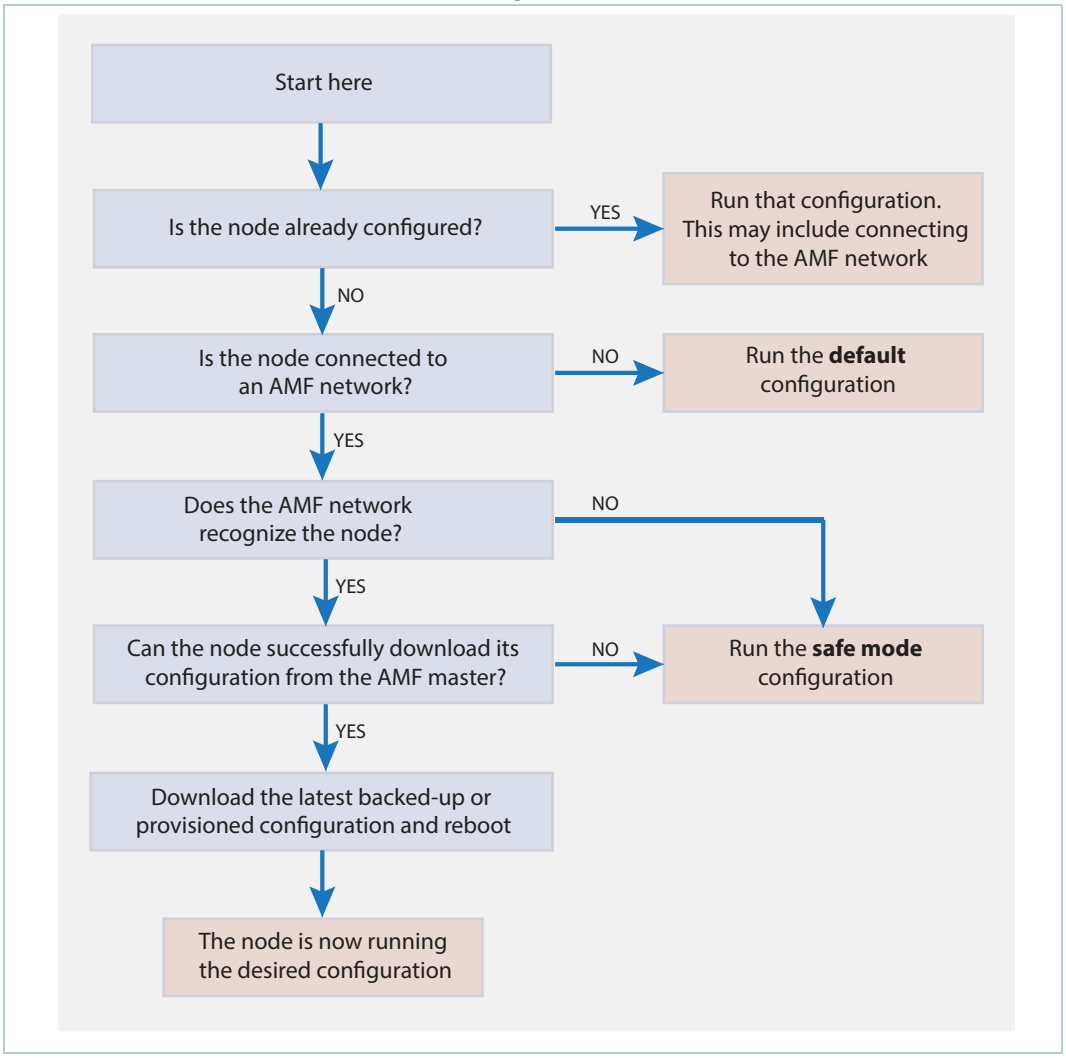

How does the recovering node work out which files to download?

A key step in the flowchart above is "**Does the AMF network recognize the node?**" This is the step where the power of AMF 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 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 15: How a recovering node works out which files to download

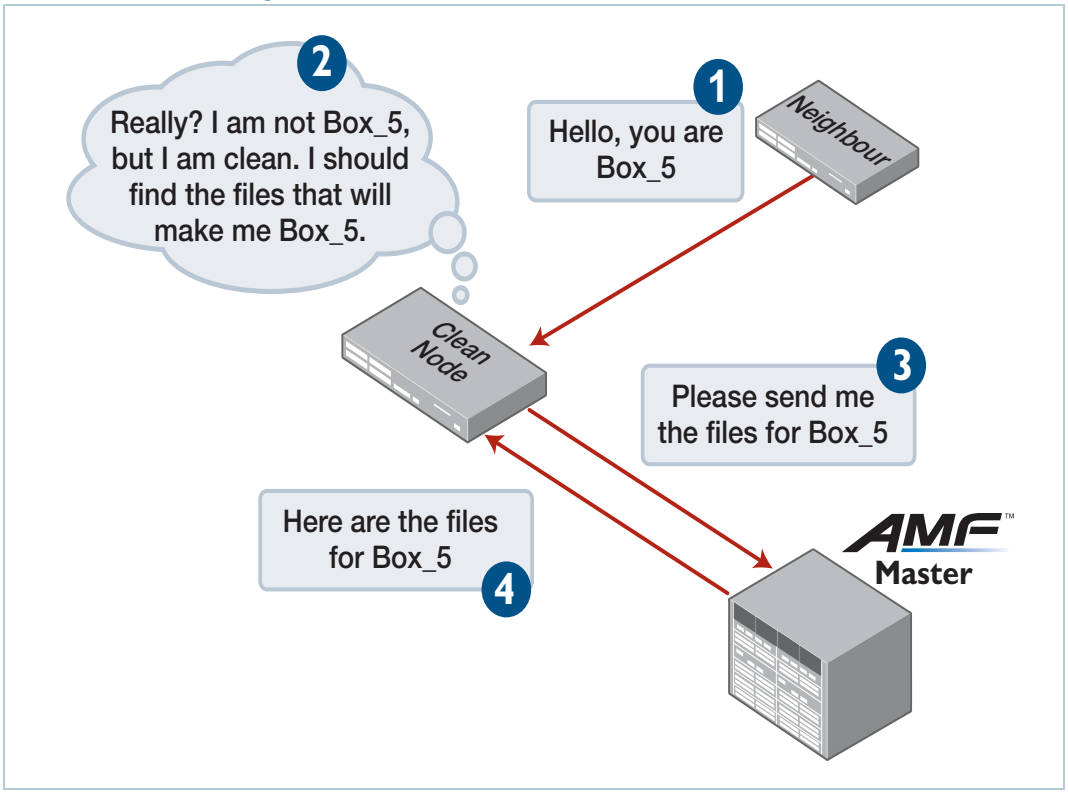

Replacing a device with an equivalent model

From AlliedWare Plus version 5.4.9-0.1 onwards, automatic recovery will work on AlliedWare Plus devices replaced with an equivalent model device as per the table below:

Table 2: Permitted device substitutions

Original device type	Replacement device type	Supported from version*
GS900MX/MPX	GS980MX	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 needs to be running on the AMF master and the node's neighbor for the recovery to work.

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 16: Example log output showing release file error

```
...
07:02:52 x510 atmffsd: /flash/x510-5.4.9-0.1.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:x530-5.4.9-
0.1.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.

Recommended procedure when replacing a device using automatic node recovery

The following procedure is recommended when replacing a device using AMF automatic node recovery.

1. Ensure that the replacement device is in a 'clean' state, see "Restoring a node to a "clean" state" on page 85
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 can successfully auto-recover the broken device.

A special note on LAGs

If the device to be replaced was connected to the AMF network using aggregators, it is important to follow the procedure above. Failure to do this may result in the AMF 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.

Points to note about automatic 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. 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

Figure 17: 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 x510_1 ATMF[863]: Automatic node recovery started
23:03:26 x510_1 ATMF[863]: Attempting to recover as x510_1
23:03:26 x510_1 ATMF[863]: Checking master node availability
23:03:32 x510_1 ATMF[863]: Master has joined. 2 members in total.
23:03:32 x510_1 ATMF[863]: x908_VCS_2 has joined. 3 members in total.
23:03:32 x510_1 ATMF[863]: x908_VCS_1 has joined. 4 members in total.
23:03:37 x510_1 ATMFFSR[2950]: Retrieving recovery data from master
node Master
23:05:18 x510_1 ATMFFSR[2950]: File recovery from master node
succeeded. Node will now reboot
Flushing file system buffers...
Unmounting any remaining filesystems...
Restarting system.
```

Recovery progress indication

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 recovery progress indication is active.

Note: This feature is not available on the x8100 series switches.

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 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. 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

The device is then rebooted to put it into a clean state. The device can then be used for automatic node recovery.

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 network.

Caution

If the **atmf cleanup** command is run on a node that is connected to an AMF network, and the AMF 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 master backup. For this reason it is recommended that you disconnect the node from the AMF network before the **atmf cleanup** command is executed.

Recovering a GS900MX/MPX series switch

From AlliedWare Plus version 5.4.7-2.1 onwards, you can, with some preparation, auto-recover GS900MX/MPX Series switches where a VCStack port has been used as an AMF link.

This is achieved by either:

- disabling stacking on the device from the CLI, before using it in a recovery situation, or
- using a USB storage device to provide an autoboot configuration that disables stacking.

Disable stacking on the device using the CLI

Log in to the device and run the following commands:

```
awplus>enable
awplus#configure terminal
awplus(config)#no stack 1 enable
```

```
Warning: this will disable the stacking hardware on member-1.
Are you sure you want to continue? (y/n): y
Warning: A reboot without saving device configuration leaves dual use
ports in network mode.
Warning: To allow AMF auto-recovery with dual use ports do not save the
config.
% stackports cannot be reconfigured until after the config has been saved
and the device has been rebooted
23:59:44 awplus VCS[674]: Deactivating Stacking Ports on stack member 1
```

```
awplus(config)#exit
awplus#atmf cleanup
```

```
This command will erase all NVS, all flash contents except for
the boot release, and any license files, and then reboot the switch.
Proceed ? (y/n):y
13:05:17 a3-gs900 IMISH[4465]: Switch cleaned, rebooting at request of
manager
```

The device is now ready for use in a recovery or provisioning situation.

Using a USB storage device to disable stacking

An alternative to the CLI method is to make use of a USB storage device and the AlliedWare Plus **autoboot** feature.

The **autoboot** feature allows a factory-new device to load configuration files, and a release image, from a USB storage device, the first time the device boots.

Step 1: **Prepare your USB storage device**

Create the following three files and copy them to the root directory of your USB storage device.

autoboot.txt

```
;Autoboot Creation, Local Time Thu, 14 Sep 2017 09:50:31
;

[AlliedWare Plus]
Copy_from_external_media_enabled=yes
Boot_Config=default.cfg
```

default.cfg

```
no stack 1 enable
trigger 1
type reboot
script 1 usb:/atmf_clean.sh
```

atmf_clean.sh

```
#!/bin/ash

#Activating findme so user knows script is about to run.
echo -e "enable\nfindme" | imish -l

sleep 20
#Creating special file so default startup knows to ignore stacking
configuration.
touch .configs/.stk_hw_disabled
echo -e "enable\n atmf clean\n y\n" | imish -l
```

Step 2: **Connect the USB storage device, with required files, to a factory-new GS900MX/MPX.**

Step 3: **Power on the GS900MX/MPX.**

Step 4: **The device boots from flash and detects the USB storage device**

The device will boot from flash, then detect the USB storage device. If you have access to the console the following information is displayed:

```
awplus Autoboot: Restoring configuration from usb:/default.cfg to
flash:default.cfg, Please wait ..
awplus Autoboot: Configuration successfully restored
awplus Autoboot: Autoboot restore successful, rebooting device.
```

This shows the device has read the configuration from the USB device and is about to reboot.

Step 5: **Device reboots and applies the new configuration file.**

The new configuration will run the **atmf_clean.sh** script on the USB device. The script disables stacking, does an **atmf cleanup**, and executes the **findme** command, which causes all the LEDs on the switch to flash.

Step 6: **Wait for the flashing LEDs.**

The **findme** command causes the device's port LEDs to flash. This is a visual cue to the person installing the switch that the script is running.

Step 7: **Remove the USB storage device**

Once the LEDs stop flashing it is safe to remove the USB storage device.

Step 8: **Device reboots and initiates auto-recovery procedure.**

Note: If the USB storage device is not removed from the switch in step 7 it will be stuck in a reboot loop. If this happens, wait until after you see the flashing LEDs again (Step 6) before removing the USB device.

Manual node recovery

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 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 safe configuration mode applied (see "AMF safe configuration procedures" on page 91). After investigating the failure and taking remedial action you may choose to 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 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 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 network and function correctly.

Output 18: Example output showing manual recovery

```
amf1#atmf recover x510_1 Master
This command will erase ALL flash contents.  Continue node recovery?
(y/n)y
Manual node recovery successfully initiated
x510_1#23:15:32 x510_1 ATMFFSR[8477]: Retrieving recovery data from
master node Master
23:17:17 x510_1 ATMFFSR[8477]: Manual node recovery completed
x510_1#
```

Node recovery on VCStacks

Node recovery on VCStacks that are part of an AMF 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.

In the extremely unlikely situation of needing to replace an entire VCStack within an AMF network, you can use AMF 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 85).

It is recommended that at least one AMF link per VCStack is configured on either a static or LACP aggregator that contains ports across all stack members. This ensures that if a stack member fails, there is no risk that the VCStack will leave the AMF Network or Area.

The procedure for recovering an entire stack is as follows:

- Connect a clean device to the AMF network, and power it on. The connections into the AMF network should be between the appropriately configured AMF links on the neighboring node, and the ports configured as AMF links in the backup of the failed node’s configuration, (i.e. the ports configured as AMF links on the failed node).
- The AMF 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.
- 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.
- 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.

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.

Recovery of devices with subscription licenses

Subscription licenses are keyed to the serial number of a device. From release 5.4.8-0.1 onwards, if you replace a device, then AMF recovery automatically transfers these licenses to the replacement device.

AlliedWare Plus grants these subscription licenses a 28-day grace period. During this grace period, subscription features operate as normal. This gives you time to register the licenses to the new device's serial number.

For this feature to work:

- your AMF master/controller must be running 5.4.8-0.1 or newer
- the replaced (failed) device must have been running 5.4.8-0.1 or newer
- AMF must have backed up the replaced device, either automatically, or manually, while it was running 5.4.8-0.1 or newer.

The **show license external** command displays the following message:

```
awplus#show license external

NOTICE: This device has undergone ATMF recovery and is currently using
        licenses registered to serial number Axxxxxxxxxxxxxxx.
        The grace period ends at 00:52:28 11 Apr 2018. To ensure
        continued operations, please contact Allied Telesis Customer
        support to have the license entitlements transferred to this
        device using serial Axxxxxxxxxxxxxxx
...
```

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.

Recovery of devices with release licenses

SwitchBlade x908 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. 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(s) for the replacement device has/have been obtained, these can then be installed on the new device using the **license** command. To obtain release license(s), 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.

AMF safe configuration

If AMF automatic node recovery fails, AMF 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 safe configuration operation

A log message is generated whenever AMF 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 safe configuration procedures

The procedures for AMF 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 management VLANs or any other VLANs that are detected on the AMF 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 links or cross-links are shut down. These links and cross-links are detected by AMF 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 19: **show vlan brief command** output - for a device in AMF 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 20: **show running-config** output for a device in AMF 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 safe configuration

If your node has had AMF safe configuration applied, you can use normal CLI configuration commands to modify the running-configuration to whatever configuration is required.

The example below shows a device being returned from AMF safe configuration mode to having its default VLAN and port settings applied. Note that in this example a 24-port interface card has been used.

```
awplus# configure terminal
awplus(config)# interface port1.0.1-port1.0.24
awplus(config)# interface port1.1.1-port1.1.24
awplus(config-if)# switchport trunk native vlan 1
awplus(config-if)# switchport trunk allowed vlan remove 4090
awplus(config-if)# switchport mode access
% port1.1.5 has ATMF link configured so its mode cannot be changed
% 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 network, AMF link and cross-link settings should not be changed. In the example above you can see that port1.0.5 is an automatically configured AMF link.

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.

Recovery of AMF devices with special links

AMF recovery files are created for nodes with special links. Special links include:

- virtual links,
- area links terminating on an AMF master, and
- area virtual links terminating on an AMF master.

An AMF 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 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 21: **show atmf recovery-file** output

```
node1#show atmf recovery-file

AlliedWare Plus (TM) 5.4.7 12/17/17 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 2018  19:21:19
node2                       18 Feb 2018  19:21:19

=======
node2:
=======

ATMF Recovery File Info: Special Link Present
Location                    Date         Time
USB storage device          20 Dec 2017  18:59:06
node1                       18 Feb 2018  19:20:55
node2                       18 Feb 2018  19:20:55
```

From version 5.4.8-0.2, 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 22: **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.

How to recover a TQ AP Guest

When correctly configured as AMF Guest Nodes, Allied Telesis TQ-series access points offer the additional advantage that their configuration files can be included in the automated backups performed by AMF. While zero-touch recovery is not available on AMF Guest nodes, TQ-series access points included in the AMF backup regime can have their configuration manually recovered. This is useful when replacing a failed TQ device. The following example shows the configuration, backup, and restore of a TQ access point using AMF.

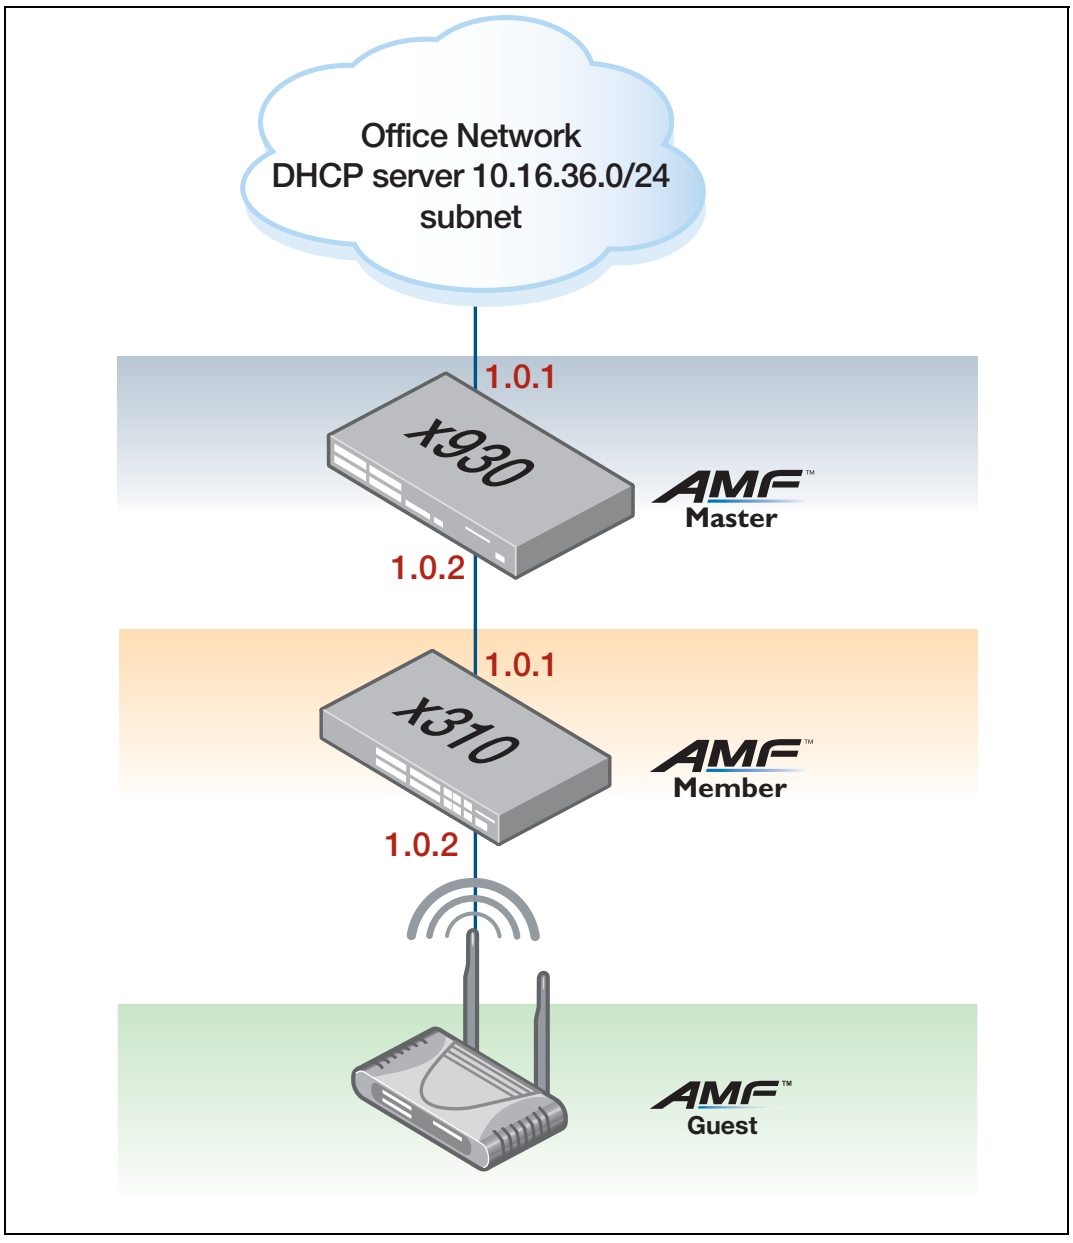

Annotated x930 (AMF Master) configuration

```
!
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 x310 configuration

```
!
service password-encryption
!
hostname x310
!
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 x310-50
!
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

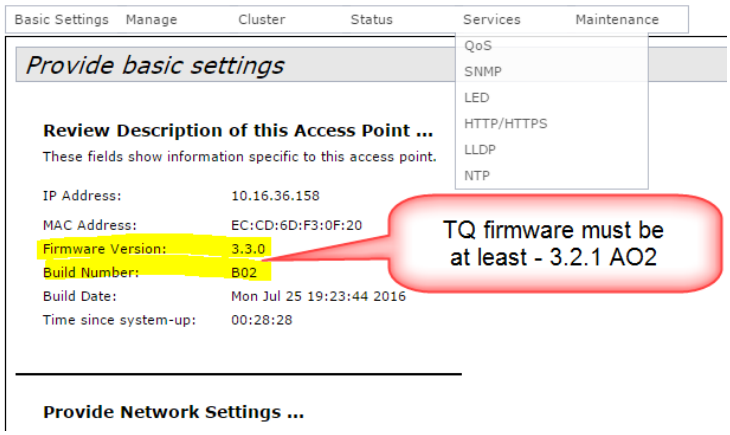

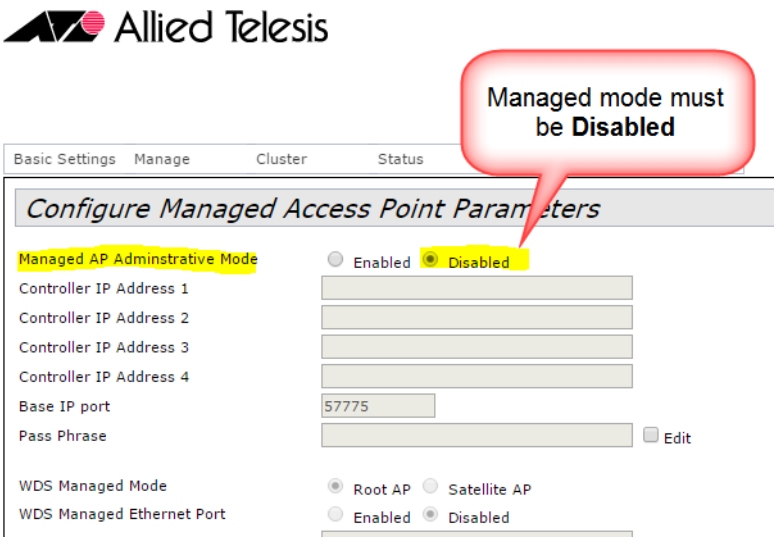

Some sample show commands

Output 23: **show atmf nodes all** output (only available on AMF 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               3930-52GTX       Y     S      none             0
  x310               x310-50FT        N     S      x930             1
  x310-1.0.2         AT-TQ4600        N     G      x310             -

 Current ATMF node count 3 (guests 1)
```

Output 24: **show atmf guests** output (only available on AMF masters and controllers)

```
x930#show atmf guests

  Guest Information:

  Device           Device           Parent       Guest      IP/IPv6
  Name             Type             Node         Port       Address
-------------------------------------------------------------------------
  x310-1.0.2       AT-TQ4600        x310         1.0.2      10.16.36.158

  Current ATMF guest node count 1
```

Output 25: **show atmf guests details** output (only available on AMF masters and controllers)

```
x930#show atmf guests detail

ATMF Guest Node Information:

Node Name                       : x310
Port Name                       : port1.0.2
Ifindex                         : 5002
Guest Description               : x310-1.0.2
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 26: **show atmf backup guests** output (only available on AMF 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 Oct 2016 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
-------------------------------------------------------------------------
x310              port1.0.2           06 Oct 2016   13:42:12   Good
```

Output 27: **show atmf links** output (available on all AMF nodes)

```
x310#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     x310-1.0.2   -          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
│   │   └── x310
│   ├── logs
│   │   └── x310
│   ├── nodes
│   │   └── x310
│   │       └── port1.0.2
│   │           └── config.xml
├── info
├── logs
└── nodes
    ├── x930
    │   └── flash
    └── x310
        └── flash
```

How to manually recover a failed TQ AP guest.

In this example, we have already replaced the failed device with another TQ-4600.

```
x310#debug atmf error
x310#terminal monitor
x310#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 x310 IMISH[6105]: [manager@ttyS0]atmf recover guest port1.0.2
15:07:12 x310 ATMFFSR[22152]: Attempting to ping master (x930) (attempt=1)
15:07:13 x310 ATMFFSR[22152]: Automatic node recovery ping -c 1 -w 1
172.31.1.12 was successful
15:07:13 x310 ATMFFSR[22152]: Node recovery pinged master (x930) reached
15:07:13 x310 ATMFFSR[22152]: Node recovery found 1 potential nodes with
recovery services
15:07:13 x310 ATMFFSR[22152]: Interrogating node x930 for backup
15:07:13 x310 ATMFFSR[22152]: Selected master recovery node x930
15:07:13 x310 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/x310/
port1.0.2 /tmp/.atmf_guest/recover
15:07:18 x310 DHCPSN[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 x310 NSM[602]: Port down notification received for port1.0.2

15:07:20 x310 NSM[602]: Port up notification received for port1.0.2
15:07:42 x310 MSTP[876]: CIST port1.0.2 now forwarding, propagating TC to
other ports
15:07:44 x310 DHCPSN[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
x310#
```

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.

Auto-recovery and Provisioning of Isolated Nodes

Version 5.4.7-2 adds support for the auto-recovery and provisioning of isolated nodes.

An isolated node is an AMF member that is only connected to the rest of the AMF network via a virtual-link. As there is no physical connection to a neighboring AMF device, isolated nodes cannot identify their location in the AMF network via traditional means. Instead, it is identified using an **identity token** stored on the AMF master. This token is saved to the AMF 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 master must be accessible to the isolated AMF node. This is achieved by using DHCP to send the Uniform Resource Identifier (URI) of the AMF master to the recovering node. The URI is either the IP address or DNS entry of the AMF master when running as a recovery server. This URI must be both resolvable and reachable from the recovering node.

- After an AMF member is backed up the identity token of the recovering node is stored in a database on the AMF master
- When the AMF 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 AMF 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 AMF 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 18: Isolated node auto-recovery example

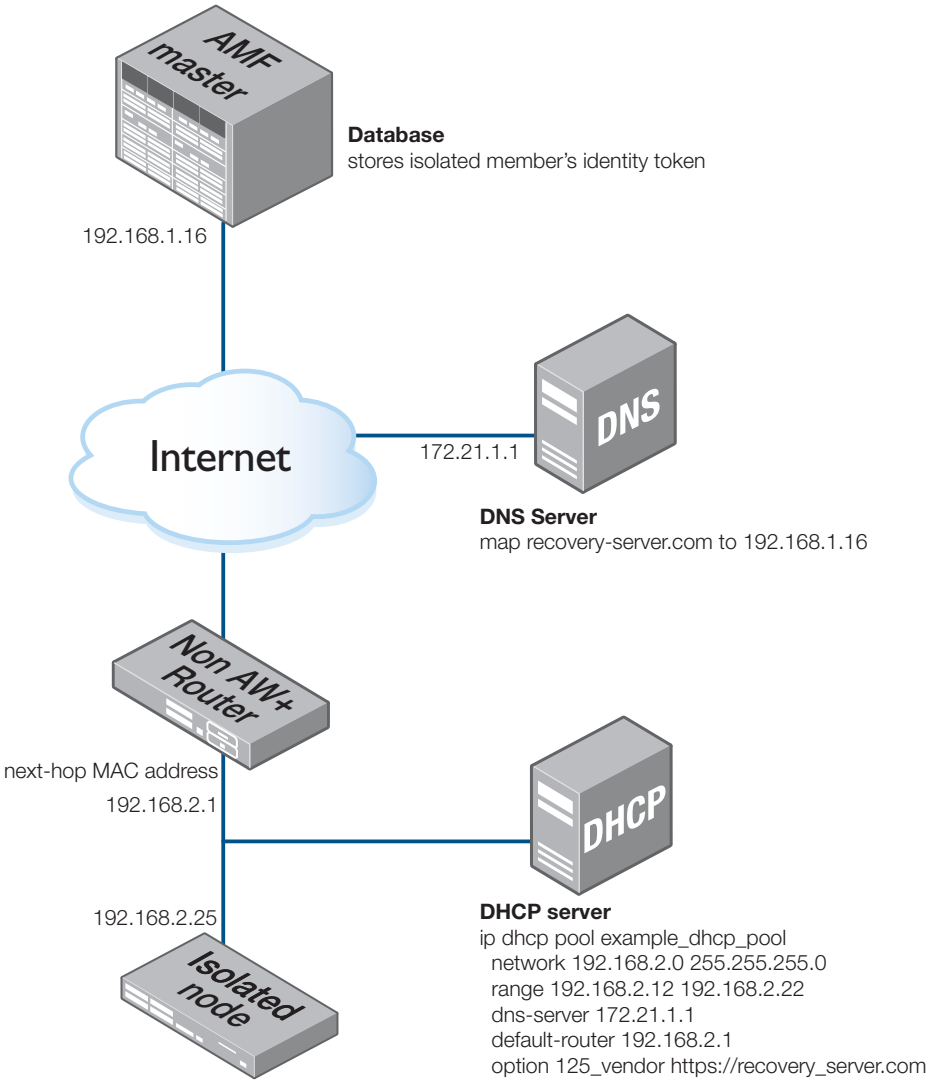

Note: All AMF traffic running over a public network should be encrypted by configuring a VPN between the AMF network and any remote sites.

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 master that does not depend on the recovering device.
- The AMF 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 AMF 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 node is backed up once the new device has been installed.

DHCP Server

There are three alternatives for supplying the AMF master's URI to the recovering node:

- Configure the DHCP server to send the AMF master's FQDN using option 125.
- Configure the DHCP server to send the AMF 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 28: Sample DHCP configuration supplying the AMF master's FDQN 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 29: Sample DHCP configuration supplying the AMF 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 30: 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 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 AMF master's IP address instead of its FQDN.

Configuring AMF

In order to handle recovery requests from isolated nodes, the **recovery-server** must be enabled on the AMF master. Use the following command on the AMF master to enable **recovery-server**.

```
awplus# atmf recovery-server
```

Replace an existing node

Once **recovery-server** is enabled on an AMF network, the next time an isolated node is backed up its **identity token** will be stored in the AMF 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 85
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 can successfully auto-recover the broken 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 117 for general information on how to provision a node.

Note: As an isolated node has no AMF 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 command:

```
awplus# atmf provision node my-x930 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 command:

```
awplus# atmf provision node my-x930 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 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 Tunneling (Virtual-links)" on page 36 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 85, and power up. If the network has been prepared correctly the new device will now contact the AMF master and provision itself automatically.

Firmware Auto Upgrade

There are two recommended methods for performing a firmware upgrade on an AMF area:

1. AMF reboot-rolling upgrade
2. AMF distribute firmware upgrade

The reboot-rolling firmware upgrade feature allows nodes within an AMF 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 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.

AMF reboot-rolling feature is supported from firmware version 5.4.3 and the AMF distribute firmware feature is supported from firmware version 5.4.4. Reboot-rolling and distribute firmware upgrades are not supported on the GS900MX series switches.

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 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 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 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 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 Network Upgrades

Allied Telesis strongly recommends that all nodes in an AMF 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 on newer firmware versions is compatible with AMF on older firmware versions.

AMF 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 upgrade process

There are a number of steps to follow in order to upgrade nodes within an **area** using AMF. 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 area.
3. Copy the releases to the location you intend to use for the upgrade.
4. Decide which AMF 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 network upgrade using your selected method. The AMF upgrade can be initiated from any node in the AMF area. It does not have to be initiated from the master.

Detailed explanation of the AMF 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 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 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 area**.

Once you have decided to upgrade your AMF 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 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 upgrade can be initiated from any node in the AMF network.

4. **Decide which AMF upgrade method is most suitable for your network.**

The AMF 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 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 distribute firmware method so that all nodes can 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 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 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 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               SBx81CFC400-5.4.7-0.1.rel  Release ready
core2               SBx81CFC400-5.4.7-0.1.rel  Release ready
distribution1       x930-5.4.7-0.1.rel         Release ready
distribution2       x930-5.4.7-0.1.rel         Release ready
edge1               x510-5.4.7-0.1.rel         Insufficient space
edge2               x510-5.4.7-0.1.rel         Release ready

Continue upgrading releases ? (y/n):
```

If the AMF 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 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 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 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 area, first select all nodes using the default **working-set group all** command:

```
SBx8100#atmf working-set group all
===============================================
SBx8100, SBx908-VCS1, SBx908-VCS2, x510_1, x510_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 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.

```
csg_vcf[5]#atmf reboot-rolling usb:/rel/*.rel
Retrieving data from SBx8100
Retrieving data from SBx908-VCS2
Retrieving data from x510_1
Retrieving data from x510_2
Retrieving data from SBx908-VCS1

ATMF Rolling Reboot Nodes:

                    Timeout
Node Name           (Minutes)  New Release File           Status
-----------------------------------------------------------------------
x510_2               9          x510-main-20121203-1.rel   Release ready
x510_1               6          x510-main-20121203-1.rel   Release ready
SBx908-VCS1          9          x900-main-20121203-1.rel   Release ready
SBx908-VCS2          9          x900-main-20121203-1.rel   Release ready
SBx8100             11         SBx81CFC400-main-20121203  Release ready
                                   -1.rel
% The controlling node (SBSBx8100) 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    : x510-main-20121203-1.rel to x510_2
Updating Release   : x510-main-20121203-1.rel information on x510_2
=================================================================
ATMF Rolling Reboot: Rebooting x510_2
=================================================================
02:11:32 SBx8100 ATMF[1973]: x510_2 has left. 4 members in total.

% x510_2 has left the working-set
02:13:30 SBx8100 ATMF[1973]: x510_2 has joined. 5 members in total.
Reboot of x510_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 area.

```
================================================================
Copying Release    : x510-main-20121203-1.rel to x510_1
Updating Release   : x510-main-20121203-1.rel information on x510_1
================================================================
ATMF Rolling Reboot: Rebooting x510_1
================================================================
02:14:13 SBx8100 ATMF[1973]: x510_1 has left. 4 members in total.

% x510_1 has left the working-set
02:15:53 SBx8100 ATMF[1973]: x510_1 has joined. 5 members in total.
Reboot of x510_1 has completed

================================================
Copying Release    : x900-main-20121203-1.rel to SBx908-VCS1
Updating Release   : x900-main-20121203-1.rel information on SBx908-VCS1
================================================================
ATMF Rolling Reboot: Rebooting SBx908-VCS1
================================================================
02:19:02 SBx8100 ATMF[1973]: x510_1 has left. 4 members in total.
02:19:02 SBx8100 ATMF[1973]: SBx908-VCS1 has left. 3 members in total.

% SBx908-VCS1 has left the working-set
02:20:48 SBx8100 ATMF[1973]: SBx908-VCS1 has joined. 4 members in total.
Reboot of SBx908-VCS1 has completed
02:20:51 SBx8100 ATMF[1973]: x510_1 has joined. 5 members in total.
================================================================
Copying Release    : x900-main-20121203-1.rel to SBx908-VCS2
Updating Release   : x900-main-20121203-1.rel information on SBx908-VCS2
================================================================
ATMF Rolling Reboot: Rebooting SBx908-VCS2
================================================================
02:21:54 SBx8100 ATMF[1973]: x510_2 has left. 4 members in total.
02:21:54 SBx8100 ATMF[1973]: SBx908-VCS2 has left. 3 members in total.

% SBx908-VCS2 has left the working-set
02:23:35 SBx8100 ATMF[1973]: SBx908-VCS2 has joined. 4 members in total.
Reboot of SBx908-VCS2 has completed
================================================================
Copying Release    : SBx81CFC400-main-20121203-1.rel to SBSBx8100
02:23:39 SBx8100 ATMF[1973]: x510_2 has joined. 5 members in total.
Updating Release   : SBx81CFC400-main-20121203-1.rel information on
SBx8100
================================================================
ATMF Rolling Reboot: Rebooting SBSBx8100
================================================================
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 distribute firmware upgrade

To perform an AMF 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 x510,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 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               x510-5.4.7-0.1.rel         Release ready
edge1               x510-5.4.7-0.1.rel         Release ready
distribution2       x930-5.4.7-0.1.rel         Release ready
distribution1       x930-5.4.7-0.1.rel         Release ready
core2               SBx81CFC400-5.4.7-0.1.rel  Release ready
Continue upgrading releases ? (y/n): y
===================================================================
Copying Release    : x510-5.4.7-0.1.rel to edge2
Updating Release   : x510-5.4.7-0.1.rel information on edge2
===================================================================
Copying Release    : x510-5.4.7-0.1.rel to edge1
Updating Release   : x510-5.4.7-0.1.rel information on edge1
===================================================================
Copying Release    : x930-5.4.7-0.1.rel to distribution2
Updating Release   : x930-5.4.7-0.1.rel information on distribution2
===================================================================
Copying Release    : x930-5.4.7-0.1.rel to distribution1
Updating Release   : x930-5.4.7-0.1.rel information on distribution1
===================================================================
Copying Release    : SBx81CFC400-5.4.7-0.1.rel to core2
Updating Release   : SBx81CFC400-5.4.7-0.1.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
```

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 network, its configuration is loaded from the information stored in the backup media.

AMF 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 network. You can preconfigure future AMF network nodes via the **atmf provision node** commands. The following figure illustrates the position of a future, provisioned node. Port1.1.5 on Member 2 and port1.0.3 on Member 4 are configured to expect the future node.

Figure 19: Provisioning for a future node

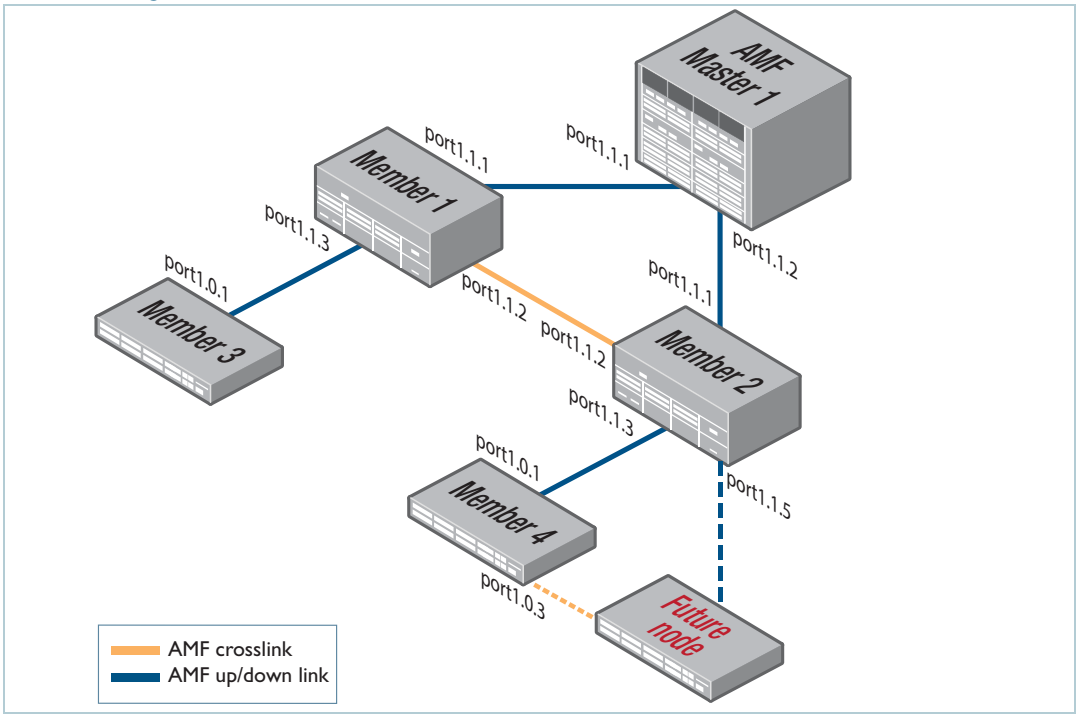

- For replacing an existing node with a node that has a **different platform** (e.g. replacing an x310 switch with an x510), and/or with a **different host name**. Using the **atmf provision node** commands you can configure the ports on adjacent nodes to accept a replacement AMF member.

If you are replacing an existing node that has the same platform and host name, refer to "Node Recovery" on page 80. In this case, node provisioning is not necessary, and node recovery will suffice.

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 uses its device-type to find the correct configuration to use. For example you can create an x510 and an x530 provisioning configuration for a node called “future-node” and if either an x510 or an x530 is attached to that node the appropriate configuration will be used.

Notes on provisioning prior to AlliedWare Plus version 5.4.9-0.1

The provisioning examples in this document are for release 5.4.9-0.1 or newer. Prior to this you were not able to specify a device-type. Also, there was no `atmf-provision` mode on the CLI, instead all the provisioning commands need to be prefixed with **atmf provision node *<nodename>***.

For example (5.4.9-0.1 or newer syntax):

```
awplus# atmf provision node <nodename>
awplus(atmf-provision)# create
awplus(atmf-provision)# locate
awplus(atmf-provision)# copy flash:<config-file> ./<config-file>
```

becomes the following in AMF versions prior to 5.4.9-0.1:

```
awplus# atmf provision node <nodename> create
awplus# atmf provision node <nodename> locate
awplus# copy flash:<config-file> ./<config-file>
```

If you are using an AMF version prior to 5.4.9-0.1, you will need to make the necessary adjustments to the examples in Table 3 on page 121 and Table 4 on page 122.

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 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 121 and Table 4 on page 122).

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 "Basic AMF Configuration" on page 17.
   - 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 121.

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

If you plan to **replace** an existing AMF 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 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 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 3 on page 121 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.

Carry out the procedures outlined in Table 4 on page 122 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.

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 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	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:
	
	`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`
	
	For information on downloading AlliedWare Plus release files see the Download Centre at www.alliedtelesis.com/support

5. Copy and set the configuration file	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:
	
	`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`
	
	For information on configuring a switch for AMF see "Basic AMF Configuration" on page 17.

6. Edit configuration file if necessary	Note that it is essential to give the provisioned node a unique hostname. For information on configuring a switch for AMF see "Basic AMF Configuration" on page 17.

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

7. Copy and set license file	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:
	
	`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`

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 link and to expect the provisioned node future_node	`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`
	
	Note that AMF links and cross-links do not need to be configured with data VLANs and can be used solely to provide AMF management VLAN redundancy.
	
	Step 8 can be repeated to configure the ports on other adjacent nodes that will be connected to the provisioned node.

Table 4: Procedure for cloning a provisioned node and configuring its adjacent nodes

1. Enter Privileged Exec mode	`AMF_Master1>enable`

2. Enter AMF 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.

Table 4: Procedure for cloning a provisioned node and configuring its adjacent nodes

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 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 links and cross-links do not need to be configured with data VLANs and can be used solely to provide AMF management VLAN redundancy.
	Step 4 can be repeated to configure the ports on other adjacent nodes to expect the provisioned node.

Connecting a provisioned node to an AMF network

When you add the new node to the AMF 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 84.

The following example shows the expected output when a provisioned node named ***future_node*** joins the AMF 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 x510-2 ATMF[999]: future_node has joined. 1 member in total.
21:57:45 x510-2 ATMF[999]: Automatic node recovery started
21:57:45 x510-2 ATMF[999]: Attempting to recover as future_node
21:57:46 x510-2 ATMF[999]: Checking master node availability
21:57:52 x510-2 ATMF[999]: AMF_Master1 has joined. 2 members in total
21:57:54 x510-2 ATMF[999]: member_1 has joined. 3 members in total.
21:57:56 x510-2 ATMF[999]: member_2 has joined. 4 members in total.
21:58:00 x510-2 ATMF[999]: member_3 has joined. 5 members in total.
21:58:03 x510-2 ATMF[999]: member_4 has joined. 6 members in total.
21:58:04 x510-2 ATMFFSR[6779]: Retrieving recovery data from master node
AMF_Master1
21:58:34 x510-2 ATMFFSR[6779]: Licence installed from certificate.
21:58:35 x510-2 ATMFFSR[6779]: File recovery from master node succeeded.
Node will now reboot
```

AMF Security

AMF 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 nodes are in a physically-isolated location, no AMF virtual links go over insecure paths, and you have complete trust in all privileged users on all the AMF nodes.

By default, AMF includes the following security features:

- AMF operates on a closed physical network and only exchanges AMF messages across links that have been configured as AMF links
- The AMF protocol is not IP-based, which means that it does not listen to connection requests over the Internet. AMF networks are not subject to remote access
- AMF creates a virtual L2 (Layer 2) management network, which is secure because the device blocks packets from external networks from entering the AMF L2 management network.

This means that attackers can only compromise an AMF network if they have physical access to it (unless it includes virtual links over insecure paths – see “Protecting AMF virtual links”).

However, any privileged user on any AMF node can configure any other AMF 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 auto-recovery and provisioning allow auto-configuration of new devices, but AMF is not affected by the reported vulnerabilities. AMF is not susceptible to attack by remote Internet hosts because the AMF protocol, by design, is only available to link partners.

AMF link management

You should only configure a link as an AMF link if it specifically connects two AMF nodes together.

If you do this, attackers can only inject packets into an AMF 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 nodes in a physically-secure location, and/or by using secure mode.

Increasing AMF security

There are three other things you can do to increase AMF security:

- configure AMF “restricted login”
- protect any AMF virtual-links that are over insecure paths
- enable AMF “secure mode”.

The following sections summarize these options.

Restricted login

With restricted login, only privileged users on the AMF master can use working-sets and automatic connections to other AMF nodes. To maximize the benefit of restricted login, the AMF master should be in a physically-secure location.

See "AMF restricted-login" on page 126 for configuration information.

Protecting AMF virtual-links

AMF virtual-links connect non-adjacent nodes by tunneling AMF 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 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 Tunneling (Virtual-links)" on page 36 for configuration information.

Note: Secure mode encrypts AMF packets. A VPN or IPsec encapsulation is, therefore, not necessary for protecting AMF 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.

Note: IPsec encapsulating is only available from AMF version 5.4.9-0.1, see "Secure virtual-links" on page 38.

Secure Mode

For the highest level of security within an AMF network, you can enable AMF “secure mode”. With secure mode enabled:

- AlliedWare Plus encrypts all AMF packets and uses certificates to verify the identity of each node in the AMF network
- Restricted login is automatically enabled and can’t be disabled
- A node can only join the AMF network if it has been authorized by a privileged user on the AMF master. This makes it impossible for an attacker to connect a device without your knowledge.

See "AMF Secure Mode" on page 127 for important details and configuration information.

AMF restricted-login

By default, users who are logged into any node on an AMF network are able to manage any other node by using either working-sets or an AMF 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 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 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 areas with more than 120 nodes, you must enable restricted-login.

AMF Secure Mode

Introduction

The AMF secure mode feature improves the security of the AMF network by reducing the risk of your network being compromised through unauthorized access to the AMF network. It achieves this by:

- Adding an authorization mechanism before allowing an AMF member to join an AMF network.
- Encrypting all AMF packets sent between AMF nodes.
- Additional logging, which enables network administrators to monitor attempts to gain unauthorized access to the AMF network.

AMF secure mode is optional and enabled from the command line interface. When running in secure mode the controllers and masters in the AMF network form a group of certification authorities. A node may only join a secure AMF network once authorized by a master or controller. When enabled, all devices in the AMF network must be running in secure mode, unsecured devices will not be able to join a secure AMF network.

Note: When an AMF network is running in AMF 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 master. This feature cannot be disabled independently of secure mode. See "AMF restricted-login" on page 126 for more information.

Licensing

AMF secure mode does not require a special license, and is enabled in the base license from release 5.4.7-0.3 onwards. Note that AMF secure mode cannot be enabled if the AMF master only has an AMF starter license.

Requirements

An AMF area operating in secure mode is limited to 126 AMF devices, this includes AMF masters and member nodes.

If an AMF 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 network. Running some AMF nodes in secure mode and some in non-secure mode is not possible.

If secure mode is enabled on an AMF network containing an AMF controller, i.e. the AMF network contains multiple areas, the controller and master of the local area must be on the same device. All AMF masters must be authorized by the AMF controller before they can join the AMF network.

Additionally, for an area master to join the AMF controller network, the controller must also be authorized by the area master.

Note: In secure mode the manual AMF recovery feature is disabled, therefore backup and recovery of AMF guest nodes is not supported.

Recommendation for multiple AMF masters

A secure AMF network with a single AMF master is vulnerable to disruption if the master is lost. A replacement master would need to re-authorize all AMF 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 network has two AMF masters.

A viable alternative to running multiple masters is to have the AMF 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 master nodes must reside in the same AMF domain and are required to be directly connected using AMF cross-links. In order to meet this requirement for AMF masters running on a VAA (Virtual AMF Appliance), a single virtual cross-link can be created using the **atmf virtual-crosslink** command. This enables a master residing on a VAA to share the AMF master role with an AMF 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 container masters, therefore if multiple tenants on a single VAA host are configured for secure mode, only a single AMF master is supported per area. In this scenario it is expected that redundancy will be provided through the virtualization hypervisor.

Enabling AMF secure mode

AMF secure mode is enabled on an AMF 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 member and is the recommended way of enabling secure mode on new or existing AMF networks. Run the following command on an AMF master to enable secure mode on an entire AMF network.

```
awplus# atmf secure-mode enable-all
```

Individual AMF 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 network they must be authorized on one of the AMF masters. Once nodes have been authorized on a master, the authorized state information is propagated through the network to all other nodes. AMF masters and controllers act as a certification authority and issue the AMF 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 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 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 network.

AMF backups when running in secure mode

When an AMF network is running with secure mode enabled, AMF auto-recovery will only function if the AMF 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 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 auto-recovery continues to work until before the scheduled backup runs.

Note: A message is displayed when an AMF network changes to or from secure mode recommending that the administrator initiates an immediate AMF backup.

Note: Guest node backups are disabled when secure mode is enabled.

AMF 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 nodes to a working-set. In a non-secure mode AMF network this feature is optional via configuration. In a secured AMF network the restricted-login feature is required, and will automatically be enabled when secure mode is enabled. This means AMF working-sets can only be created and used by a user logged on to an AMF master. If secure mode is disabled on an AMF network restricted login will remain enabled and will need to be disabled independently, if this is desired.

AMF remote login

In a non-secure mode AMF network a user logged in to an AMF node can use the AMF remote login feature to connect to any other AMF node. In AMF secure mode remote login to other AMF nodes will only be allowed from an AMF master node. AMF member nodes will not be able to use the AMF remote login feature to connect to other nodes in the network, including to AMF master nodes.

AMF auto-recovery in secure mode

The pre-requisites for the successful auto-recovery of an AMF member while the network is in secure mode are:

- The clean node must be pre-loaded with software that supports secure mode; i.e. AlliedWare Plus version 5.4.7-0.3 or later.
- The software on the AMF 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 members connected to the AMF network via virtual-links. For these nodes auto-recovery from external media should be used.

AMF secure mode pre-authorization

It is possible for an administrator to allow an AMF 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 Master with the **atmf authorize provision** command. This is done by either specifying the MAC address of the replacement device, or by specifying the neighouring node hostname and interface, i.e. the AMF 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 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 Jun 2017 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 Jun 2017 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 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 85.

Enabling secure mode on an existing AMF network

Executing the **atmf secure-mode enable-all** command on an AMF master node will reconfigure an existing AMF network into secure mode.

When the command is executed, all the AMF members leave the AMF network and rejoin in secure mode with each member being automatically authorized. There is a small disruption to the AMF 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 secure mode on an AMF master automatically enables the AMF restricted login feature if this has not already been enabled.

Note: This process saves the **running-config** on each device in the AMF 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 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 node to a secure mode network

An AMF 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 master nodes.

All AMF master nodes will receive the authorization request from the new member node. Once the node is authorized on any AMF master device this information will then be propagated to all other AMF masters in the area.

If a new AMF node, with secure mode enabled, is connected and attempts to join the AMF network a message is displayed on the AMF 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 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 network

If you decide to remove secure mode from an AMF network this can be done by executing the command **no atmf secure-mode enable-all** on an AMF 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 network.

Verifying secure mode on an AMF 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 master to display a list of all AMF 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         x510L-28GT       master         port1.0.3
  area_1_node_3         x510L-28GT       master         sa1
  area_1_node_4         x310-26FT        master         port1.0.1
```

Using the **show atmf authorization current** command from an AMF master to view all devices in the AMF network that are currently authorized to join the AMF 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 May 2017
  master                master                15 May 2017
  area_1_node_5         master                16 May 2017
  area_1_node_6         master                16 May 2017
```

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 network, but has not yet been authorized by an AMF 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 secure mode network

The **show atmf secure-mode audit** command displays a list of security recommendations for securing your AMF 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:

-----------------------------------------------------------------------
2017 Feb 2 00:59:25 user.notice node1 ATMF[848]: Sec_Audit - ATMF Secure
Mode is enabled.
2017 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     16/02/2017 09:28:22   Member3
```

AMF Guestnode

Overview

The AMF guestnode feature provides basic AMF functionality to non-AMF capable devices. These AMF “guests” are devices that either do not run the AlliedWare Plus operating system or run a version that does not support AMF. Essentially, any device that has either an IPv4 or IPv6 address can become an AMF guest.

AMF nodes can recognize the presence of an AMF guest either statically, or dynamically if it uses a protocol such as DHCP or LLDP. Once recognized, the AMF node is then able to provide a limited level of support to these devices.

A requirement of AMF guests is that they each connect directly to their own port on an AMF node with no intermediate devices.

Note: From AlliedWare Plus version 5.4.7, AMF guestnodes are not supported on ports using the OpenFlow protocol.

AMF parent nodes are aware, from the AMF guestnode configuration commands, which of their ports are connected to guest nodes. If configured to be dynamic, the AMF node then listens on these ports for DHCP or LLDP frames, in order to obtain information about any of its attached AMF guest devices.

Having discovered information about an AMF guest, or by having its information configured statically, the AMF node transmits this information to its local AMF master(s). The AMF guest information is then accessible to users or management tools that access the AMF network information database.

Information about AMF guests is not distributed to the whole AMF 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 guests, such as the TQ wireless access points, AMF has a more detailed knowledge of the management interface. AMF is, therefore, able to perform specific actions on these devices such as making backups of the configuration, or performing a recover on a replacement units.

Other guest nodes, however, offer very little information that can be obtained by the AMF parent. This is because they do not transmit frames that provide useful AMF guest information. For these guest node types information needs to be statically configured on the port of the AMF node that directly connects to the AMF guest.

AMF guest discovery

As stated in the previous section, AMF guests can be configured to be either static or dynamic depending on how the AMF node detects their existence. If the guest node uses DHCP to get an IP address dynamically then the AMF parent node can use DHCP snooping to learn the address of the guest node. This process is termed **dynamic IP discovery**.

If the AMF guest does not use DHCP to obtain an IP address, then the IPv4, or IPv6, address of the AMF guest must be manually configured on the port of the parent AMF node. This process is termed **static IP discovery**.

Once an AMF node is aware of an AMF guest's IP address, it can use this to interrogate the device for its MAC address.

Dynamic guest nodes

When AMF guests are configured for dynamic discovery, the switchport on the AMF parent is configured to use DHCP snooping to discover the AMF guest's IP and MAC addresses. Additionally, if the AMF 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 AMF node in order for dynamic discovery to work.

Static AMF guests

AMF guests configured for static IP discovery may not support protocols such as DHCP or LLDP and, therefore, offer the AMF 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 node that connects to the AMF guest.

Please note the following:

- If an AMF 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.
- An AMF 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 AMF guests with downstream device.

AMF functionality supported by AMF guests

The degree to which AMF management capabilities are extended to an individual AMF guest depends on the nature of the AMF guest as defined by its model type.

By default, the minimum management functionality provided to an AMF guest is basic topology reporting. At a minimum, AMF will track and report an AMF guest's network activity recorded at the link level. Whatever information is obtained about the AMF guest and its operational status will be stored in an AMF device database and can be viewed by users.

For some AMF guest types, however, device-specific support capability is included in AMF. If the AMF network discovers that one of these advanced-support devices is connected, then it enables the user to carry out more management operations on that AMF guest.

An AMF 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 (from AlliedWare Plus version 5.4.9-2.x).
- 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 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 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 backup of TQ AP guest devices will only occur with TQ firmware release 3.2.1.a02 and later.

Login fallback

From version 5.5.0-1.1 onwards, you can enable login fallback on TQ guest nodes using the command **login-fallback enable**. This feature allows AMF 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 144.

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. 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. 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 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 guest configuration

There are two components of AMF 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 guest devices. Each guest-class is a profile that can be assigned to multiple AMF guests.

Guest-classes are modal templates that can be applied to selected AMF 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 an AMF 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 either static or dynamic. 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.

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 AMF 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 an AMF guest configuration to a switchport of an AMF node, then associates the AMF guest with a guest-class profile. This will determine the method AMF uses to interrogate the AMF guest.

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
```

Note: An IPv4 address is required on the VLAN interface of the parent AMF node in order for dynamic discovery to work.

Configure a TQ wireless access point with static discovery.

Step 1: Create a guest-class named TQ5403.

```
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 as 'manager' and the password as 'tq54-guestpass' for the TQ5403.

```
Parent-Node1(config-atmf-guest)# username manager password tq54-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 TQ5403 ip
192.168.10.1
```

The guest-class specified in the **atmf-guestlink** command must have already been created using the **atmf guest-class** command. If the guest-class name does not correspond to an predefined guest-class then the command will return an error. Also note that names of guest-classes are case-sensitive.

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 DHCP snooping see the LLDP Feature Overview and Configuration Guide

Step 1: **Enable LLDP on the AMF 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
```

helpSet 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
```

Replacing a TQ guest node with login fallback

From version 5.5.0-1.1 onwards, you can enable login fallback on TQ guest nodes using the command **login-fallback enable**.

Login fallback means; if a guest node’s saved username and password fail, AMF 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 an AMF guest node. AMF can then start the AMF 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 AT-TQ5k, use the commands:

```
node1#configuration terminal
node1(config)#atmf guest-class AT-TQ5k
node1(config-atmf-guest)#username testuser password testpass
node1(config-atmf-guest)#login-fallback enable
node1(config-atmf-guest)#end
node1#
```

AMF guestnode show commands

Guest node show commands can be executed on either AMF masters or AMF controllers. When executed on an AMF master the command displays guest node information for the local AMF area only, and when executed on an AMF controller the command displays output for all AMF areas. The **show** command syntax is different on AMF masters and AMF controllers.

show atmf guests

This command displays a list of guests that an AMF network has discovered. It needs to be run on an AMF master.

Output 46: 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 20: Parameter Descriptions from the **show atmf guests** command

PARAMETER	DESCRIPTION
Device Name	The name assigned for this device within the AMF network. It could be a name that is discovered from the device, or failing that, a name that is auto-assigned by AMF. 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 member name of the AMF 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 master.

Output 47: 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 21: Parameter Descriptions from the **show atmf guests detail** command

PARAMETER	DESCRIPTION
Node Name	The AMF 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 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 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.

Figure 21: Parameter Descriptions from the **show atmf guests detail** command (continued)

PARAMETER	DESCRIPTION
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 48: 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 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 guestnode 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 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 guest backup feature.
- restored manually using the AMF 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 parent node. This should include the username, password, and port information for connecting to the device via HTTP.
3. Create an AMF guest-link using the with the configured ONVIF guest-class.

Figure 22: 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 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 parent node. This should include the username, password, and port information for connecting to the device via HTTP.
5. Create an AMF 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 23: 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, and if that ONVIF device supports system backup and restore, AMF 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 a**tmf 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 parent node to see a list of all locally configured ONVIF guest nodes.

Output 49: 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 parent node to see details about the configured guest nodes.

Output 50: 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 master.

Output 51: 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 Oct 2019 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 Oct 2019   03:00:13   Good
x930              port1.0.2           22 Oct 2019   03:00:01   Good
```

AMF Support for x600 Series Switches

x600 Series Switched went End of Life before AMF was included in the AlliedWare Plus software. This mean that these switches cannot be full AMF members. However, AMF networks running version 5.4.6-1.x or later can be more seamlessly integrated with AlliedWare Plus x600 Series switches, as long as the x600 Series switch is running version 5.4.2-3.14 or later.

The x600 Series switch must be directly connected to an AMF node that is running 5.4.6-1.x or later. The x600 Series switch provides the following information to the AMF node that it is connected to:

- MAC address of the port connected to the AMF node
- IPv4 address
- IPv6 address
- name/type of the device (Allied Telesis x600)
- name of the current firmware
- version of the current firmware
- configuration name

Earlier software versions made most of this information available to AMF from x600 Series switches, but it was necessary to configure the x600 as an AMF guestnode (so it needed to be configured with DHCP and/or LLDP). With version 5.4.2-3.16 or later, as soon the x600 is connected to an appropriately configured port of an AMF node, it is immediately integrated into the AMF network.

To configure the new functionality, use the following steps:

Step 1: **Upgrade the software version on the x600 Series switch**

The x600 Series switch must be running version 5.4.2-3.16 or later.

Step 2: **Configure the link to the x600 Series switch**

On the AMF node to which the x600 Series switch is connected, configure the link to the x600, using the command:

```
node_1(config-if)# switchport atmf-agentlink
```

Step 3: **Monitor the x600 Series switch**

On the AMF node to which the x600 Series switch is connected, you can see the details of the x600 by running the following command:

```
node_1# show atmf links guest detail
```

Virtual AMF Appliance/AMF Cloud

Introduction

The Virtual AMF Appliance (VAA), also known as AMF Cloud, allows your AMF master and/or controller to be run on a virtual appliance, rather than running on an physical Allied Telesis device. This means you get all the functionality of integrated hardware-based management, with the added advantages of private or public cloud access and flexibility.

Its 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 24: VAA deployed online with Amazon Web Services

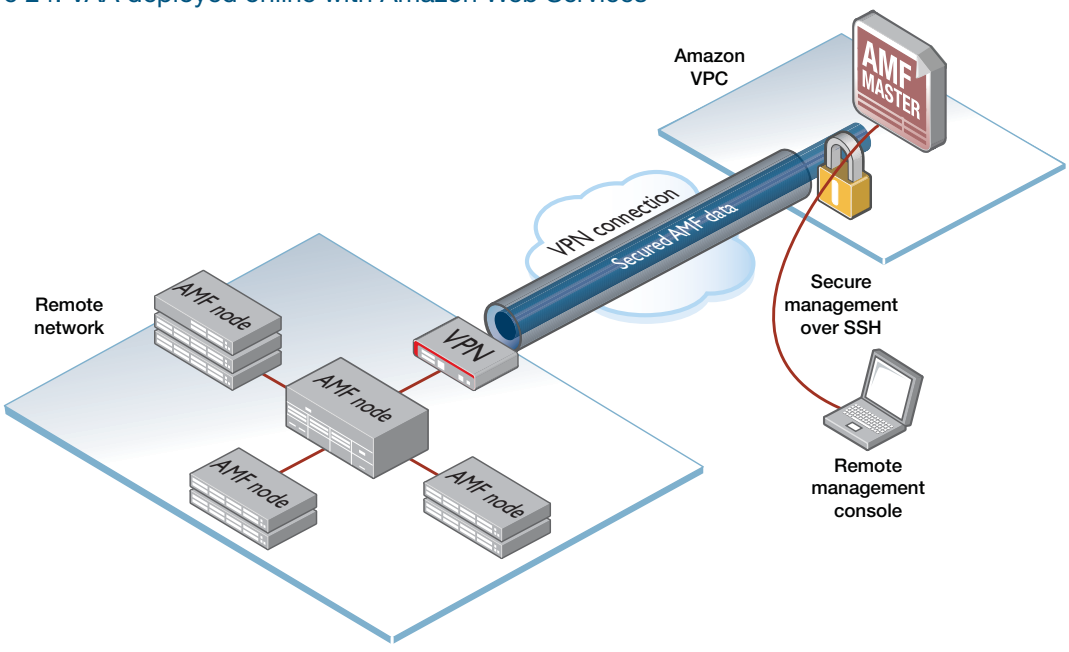

For for information on installing the Virtual AMF Appliance see the Install Guide: Virtual AMF Appliance (VAA) for AMF Cloud.

What is AMF virtualization?

The VAA is a variant of AlliedWare Plus that is supplied as an ISO image and is loaded on 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 controller or master just as you would a physical device.

The VAA supports a minimum level of functionality required to support AMF, therefore the VAA does not support standard L2 switching and L3 routing. However, all AMF 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 the VAA, and triggers created that will run these scripts.

The VAA connects to other AMF devices (both physical and virtual) using AMF virtual-links.

Licensing

VAA licensing is subscription-based and depends on the size of the network under management. You will need to consider:

- how many AMF masters throughout the network are linked to an AMF controller, and
- how many nodes in each AMF area are linked to that area's AMF master.

Each VAA acting as an AMF controller or AMF master will need its own unique license file that is based on the unique serial number of the VAA. This license defines the number and type of nodes allowed throughout the AMF network.

See the Autonomous Management Framework Datasheet for information on AMF Cloud master and controller licensing.

See the Install Guide: Virtual AMF Appliance (VAA) for AMF Cloud for information on obtaining and installing these licenses.

Multiple Tenants on AMF Cloud

Introduction

Running multiple tenants on a single Virtual AMF Appliance (VAA) provides an efficient way to configure and control different service networks via a centralized virtual machine. It allows one VAA, also known as AMF Cloud, to act as both the AMF controller and the AMF masters for up to 300 AMF areas. Each VAA AMF master runs in its own virtual AlliedWare Plus environment known as an AMF container.

An AMF container is an isolated instance of AlliedWare Plus with its own network interfaces, configuration, and file system. The features available inside an AMF container are a sub-set of the features available on the host VAA, this allows it to function as a uniquely identifiable AMF master. From the host VAA each AMF area can be managed using either Vista Manager or the command line interface (CLI).

The tenants in each AMF area could be branch offices of a single organization, or separate customers managed by a single service provider. Hosting multiple tenants on a single VAA could also be used where a service provider provisions an AMF container for a client tenant, but the tenant manages their own AMF area. This is possible because each area is isolated from all other areas through the use of AMF containers, with each tenant only able to view and control their own area.

The key advantage of hosting multiple tenants on a single VAA over a traditional AMF installation is that each area does not need a device with the capabilities to act as an AMF master. For information on installing the Virtual AMF Appliance see the Install Guide: Virtual AMF Appliance (VAA) for AMF Cloud.

Feature overview

The VAA host provides the CLI commands necessary to create and configure AMF containers. The VAA host functions as the AMF controller for the local area, while remote AMF areas are then defined and assigned to each AMF container.

Once the containers are configured with an area-link and enabled, they are automatically set-up as AMF masters with their own AMF area configuration. The administrator can then configure AMF virtual-links inside these AMF containers, which connect to physical AMF members in the remote AMF area.

Each AMF container has two virtual interfaces, eth0 and eth1:

- The eth0 interface is used to connect the container to the VAA 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 master within each container to establish AMF virtual-links with other AMF devices in its area.

Figure 25: AMF container architecture

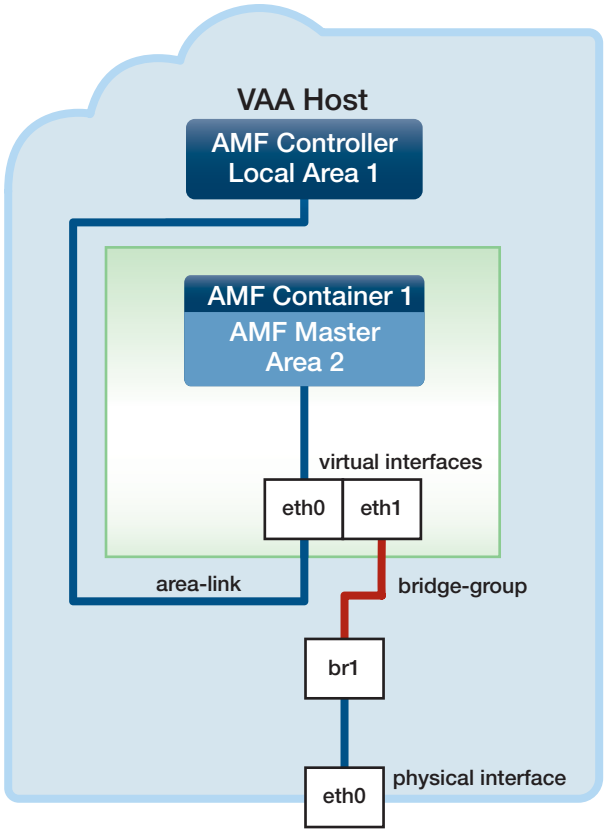

- In the case of a VAA hosted on an public cloud service, the virtual-links are carried from the remote AMF area over an IPSec protected L2TPv3 tunnel.

Each AMF container in the VAA 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 containers do not support NTP. Therefore, NTP should be configured on the host VAA, and every AMF node directly connected to the AMF container

Licensing

This feature is included in the Virtual AMF Appliance software and does not require a special license. It does, however, need the following:

- A VAA AMF controller license on the host VAA for the number of areas the controller will support (up to 300).

  The AMF controller license allows you to create an AMF master on the controller node if there is no other AMF master in the area.

- A VAA AMF master license on each AMF container for the number of nodes the master will support (up to 300).

  With no VAA master license installed, an AMF container supports only a single node, i.e. the container's AMF master.

See the Autonomous Management Framework Datasheet for more information on AMF Cloud master and controller licensing.

Configuration example for private cloud installation

Complete the following steps to configure multiple tenants on a single VAA host:

1. “Configure AMF on the VAA host”
2. “Create an AMF container on the VAA host”
3. “Configure the AMF container”
4. “Configuring the bridge connecting the container to the Hypervisor Ethernet interface”
5. “Assigning the bridge group to a container on the VAA”
6. “Configuring the remote AlliedWare Plus device”
7. “Verifying the AMF container”
8. “Managing an AMF area”

The diagram below shows a VAA virtual machine configured with a local area and three AMF containers. The left-hand container illustrates the AMF virtual-link between the AMF master in container “vac-wlg-1” and a physical switch at the remote site within AMF 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 virtual-link. The IP address of the eth1 container interface must have direct IP connectivity with the remote device terminating the AMF virtual-link. Typically this would be achieved using a site-to-site VPN.

Figure 26: AMF multiple tenant configuration on private cloud example

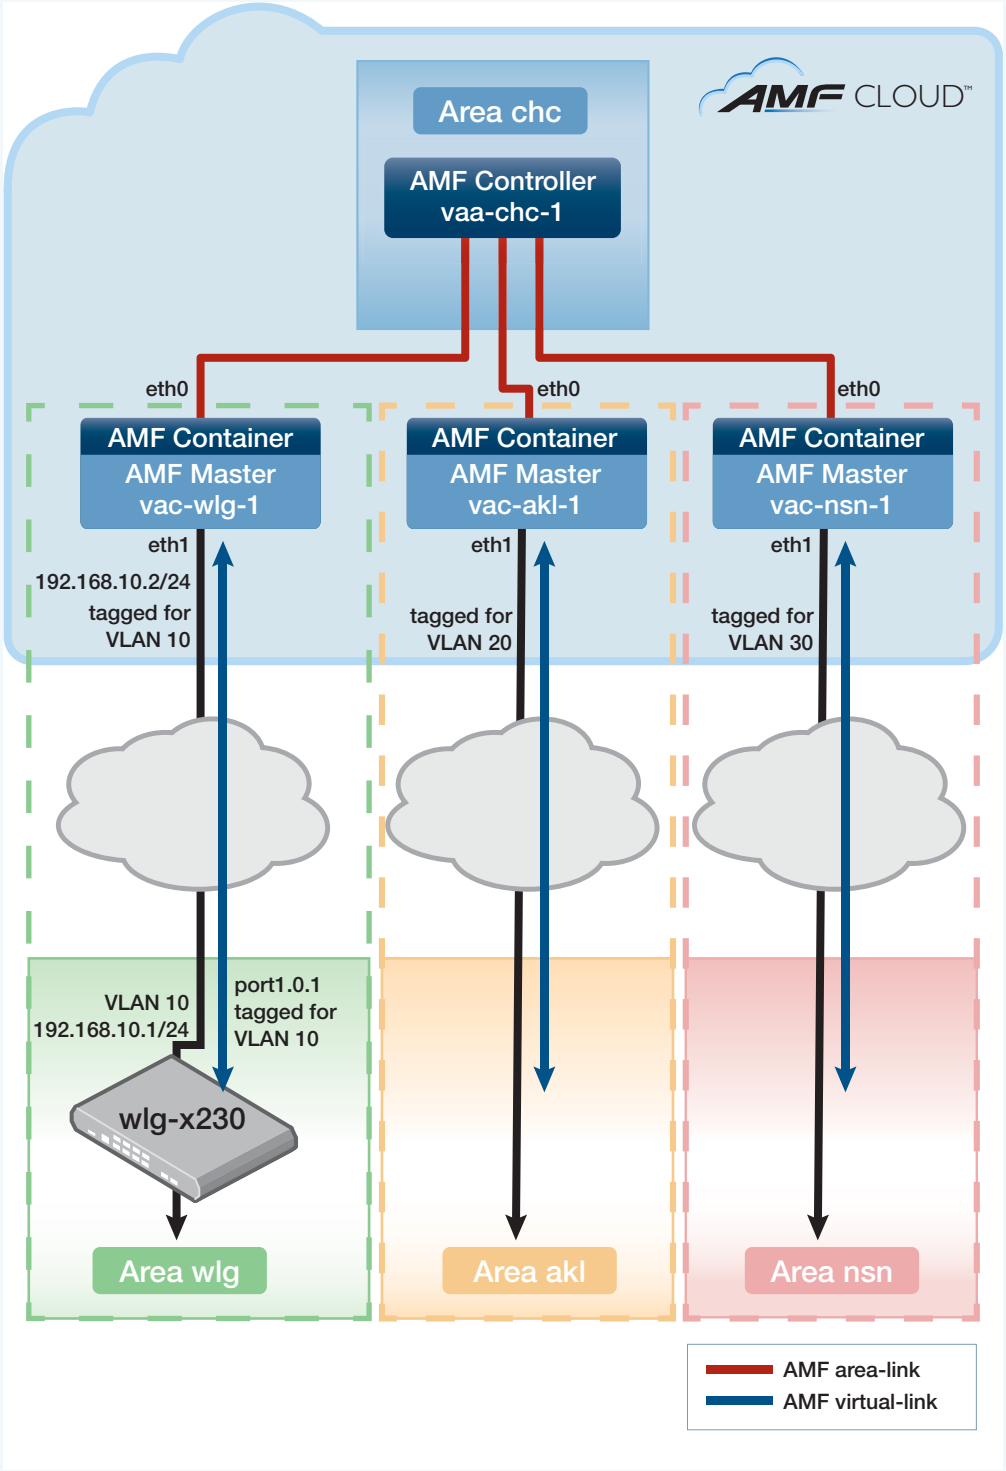

Configure AMF on the VAA host

First create the AMF areas, one for the local area which contains the AMF controller, and one for each AMF container. In this example we will create a local area on the VAA 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 VAA hostname and AMF network name**

```
awplus#configure terminal
awplus(config)#hostname vaa-chc-1
vaa-chc-1(config)#atmf network-name atlnz
```

Step 2: **Configure the VAA as AMF controller and master**

```
vaa-chc-1(config)#atmf master
vaa-chc-1(config)#atmf controller
```

Note: If the VAA is the only AMF node in its area it must be configured as both an AMF master and an AMF controller. If the local area contains other AMF nodes, one of the other nodes in that area can be configured as an AMF master but there must be an AMF master somewhere in the AMF controller's area.

Step 3: **Configure the local AMF area**

This area, "chc", will contain the local AMF controller and local AMF 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 container on the VAA host

Next we create an AMF container named "vac-wlg-1" and, inside this container, configure an **area-link** to area "wlg":

Step 5: **Create the AMF 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 container by the VAA host. This assigns the container to the area "wlg" and configures it as an AMF 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 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 container

Connect to the AMF container “vac-wlg-1”, login, and add an IP address to the eth1 interface. Eth1 is the AMF container interface that will connect to the physical AlliedWare Plus devices within the container’s remote area.

Step 8: Connect to the AMF 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 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 container’s configuration

If you have finished configuring the AMF 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 VAA to the container “vac-wlg-1” you need to exit to the host VAA.

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 host VAA**

```
vac-wlg-1(config)#exit
vac-wlg-1#
Type <Ctrl+a q> to exit the container and return to the AMF 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 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 VAA

The bridge group created on the VAA 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 controller's configuration

If you have finished configuring the AMF 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 VAA. Each AMF container's configuration must be saved from within that container.

Configuring the remote AlliedWare Plus device

Configure an AMF virtual-link from the AMF member "wlg-x230" in area "wlg" to the AMF 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 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 virtual-link from the remote device to the AMF 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 container

You can check the state and resource utilization of an AMF container with the **show atmf container** command.

Output 52: 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 container.

Output 53: 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 containers running on a VAA host use the **show atmf container detail** command.

Output 54: 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 container running on a VAA 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 55: 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 area

The virtual AMF master, inside the AMF container, can now be used to manage its AMF area in exactly the same way as if it was a physical AMF master.

Configuration example for public cloud installations

Complete the following steps to configure multiple tenants on a single VAA host:

1. “Configure AMF on the VAA host”
2. “Create an AMF container on the VAA host”
3. “Configuring L2TPv3 tunnel with IPsec encryption on VAA host”
4. “Configure the AMF container”
5. “Assigning the bridge group to a container on the VAA”
6. “Configuring the remote AlliedWare Plus device”
7. “Verifying the AMF container”
8. “Managing an AMF area”

The diagram below shows a VAA virtual machine configured with a local area and three AMF containers. The left-hand container illustrates the AMF virtual-link between the AMF master in container “vac-wlg-1” and a physical switch at the remote site within AMF 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 virtual-link. The IP address of the eth1 container interface must have direct IP connectivity with the remote device terminating the AMF virtual-link. Typically this would be achieved using a site-to-site VPN.

Figure 27: AMF multiple tenant configuration on public cloud example

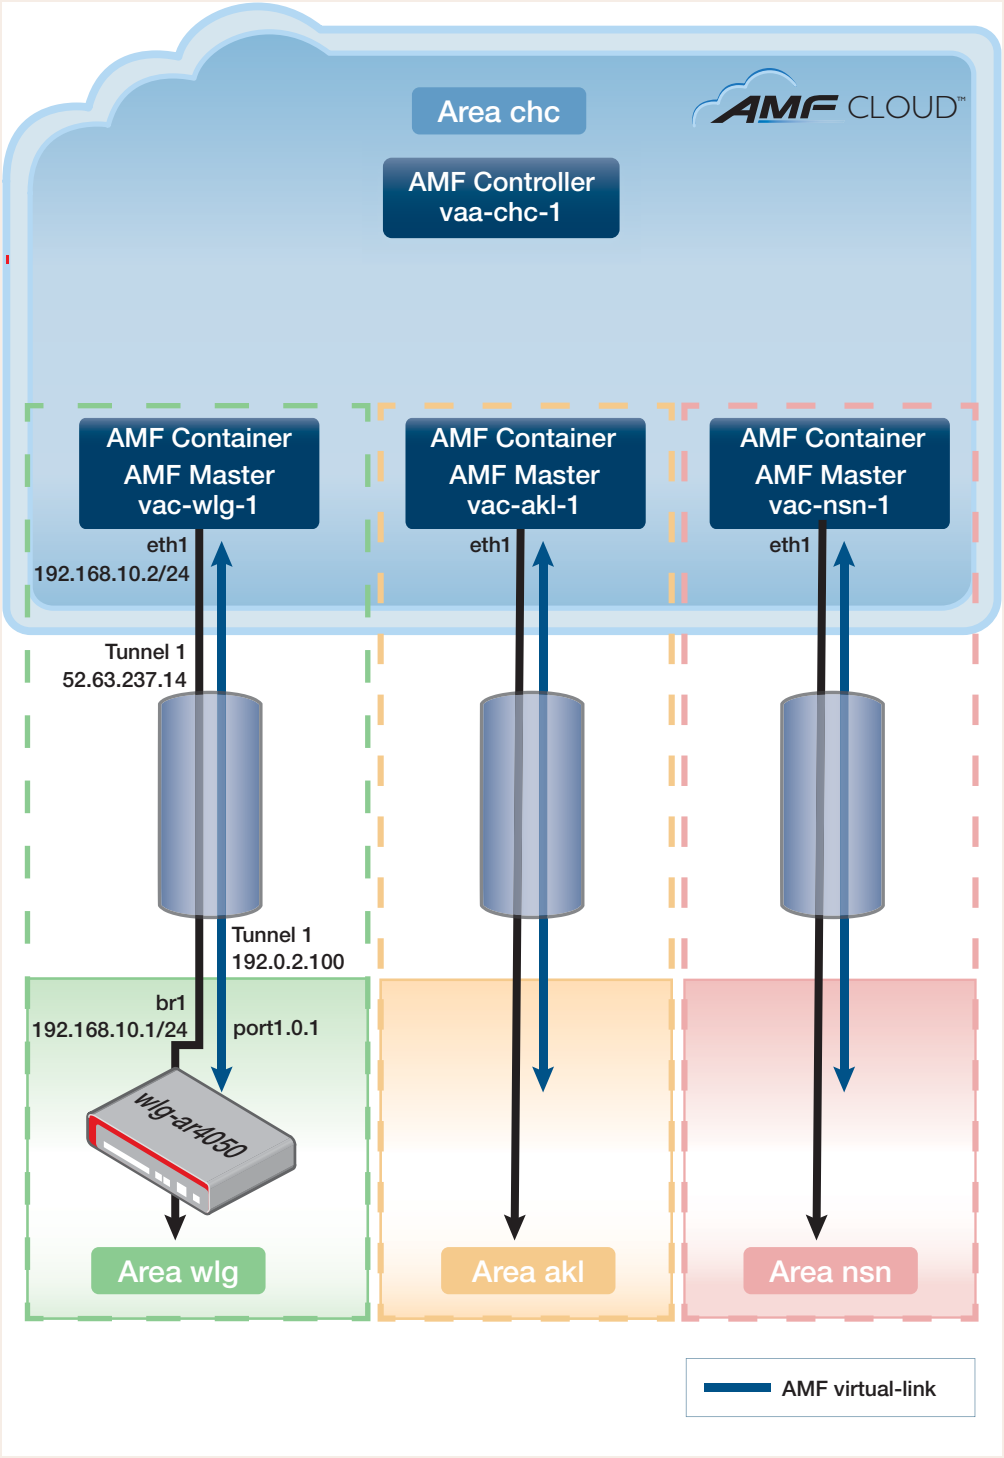

Configure AMF on the VAA host

First create the AMF areas, one for the local area which contains the AMF controller, and one for each AMF container. In this example we will create a local area on the VAA 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 VAA hostname and AMF network name**

```
awplus#configure terminal
awplus(config)#hostname vaa-chc-1
vaa-chc-1(config)#atmf network-name atlnz
```

Step 2: **Configure the VAA as AMF controller and master**

```
vaa-chc-1(config)#atmf master
vaa-chc-1(config)#atmf controller
```

Note: If the VAA is the only AMF node in its area it must be configured as both an AMF master and an AMF controller. If the local area contains other AMF nodes, one of the other nodes in that area can be configured as an AMF master but there must be an AMF master somewhere in the AMF controller's area.

Step 3: **Configure the local AMF area**

This area, "chc", will contain the local AMF controller and local AMF 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 container on the VAA host

Next we create an AMF container named "vac-wlg-1" and, inside this container, configure an **area-link** to area "wlg":

Step 5: **Create the AMF 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 container by the VAA host. This assigns the container to the area "wlg" and configures it as an AMF 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 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 VAA host

Add an L2TPv3 tunnel with IPsec encryption to the VAA 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 VAA 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 container

Connect to the AMF container “vac-wlg-1”, login, and add an IP address to the eth1 interface. Eth1 is the AMF container interface that will connect to the physical AlliedWare Plus devices within the container’s remote area.

Step 12: **Connect to the AMF 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 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 container's configuration

If you have finished configuring the AMF 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 host VAA.

Step 16: Exit to the host VAA

```
vac-wlg-1(config)#exit
vac-wlg-1#
Type <Ctrl+a q> to exit the container and return to the AMF 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 VAA

The bridge group created on the VAA 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 controller’s configuration

If you have finished configuring the AMF 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 VAA. Each AMF container’s configuration must be saved from within that container.

Configuring the remote AlliedWare Plus device

Configure an AMF virtual-link from the AMF member “wlg-ar4050” in area “wlg” to the AMF 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 VAA 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 VAA 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 VAA 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 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 virtual-link from the remote device to the AMF 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 container

You can check the state and resource utilization of an AMF container with the **show atmf container** command.

Output 56: 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 container.

Output 57: 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 containers running on a VAA host use the **show atmf container detail** command.

Output 58: 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 container running on a VAA 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”.

Applications that use AMF

Vista Manager EX™

Vista Manager EX is a graphical network monitoring and management tool for AMF networks. It automatically creates a complete topology map from an AMF 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.

This feature is available on AMF versions 5.4.7-0.1 or newer.

Configuring AMF to communicate with Vista Manager EX

Perform the following steps on your AMF 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 nodes, including all AMF masters and controllers, 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 AMF 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 network topology.**

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

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

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

Step 3: **Configure the AMF log event host.**

If the AMF controller, or AMF 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 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 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.

Using AMF in EPSR Rings

In this chapter we look at how to appropriately select whether to configure AMF links or AMF 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 to an EPSR ring

When AMF is added to an existing EPSR ring, a common mistake is to mark **all** ring links as AMF down-links. An example of this is shown below in Figure 28.

As you can see, each IE300 device has two routes back to the AMF 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 28: Common error - all ring links marked as AMF down-links

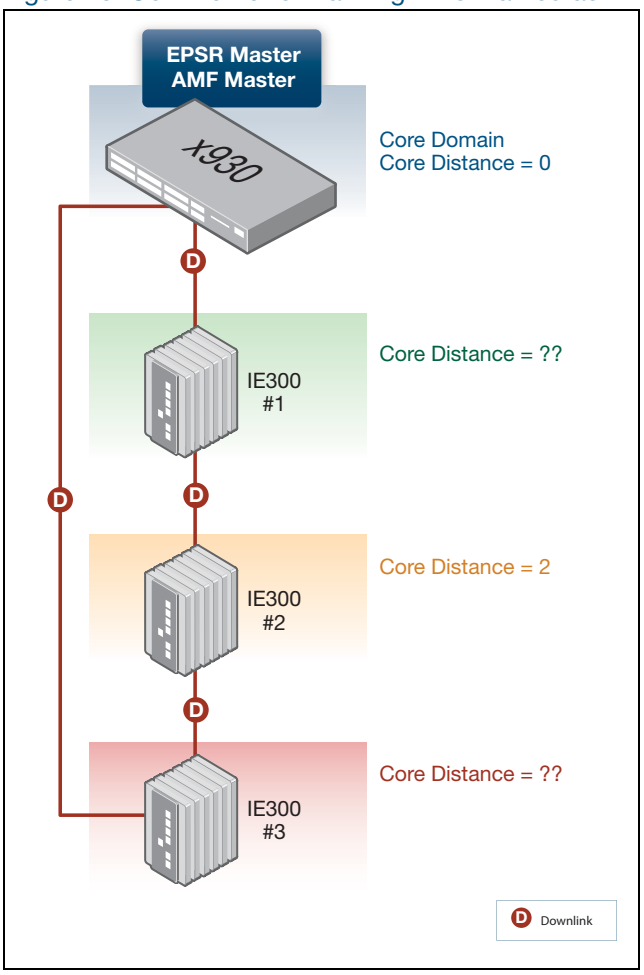

How can you configure this example above to avoid the core distance issue? Let us now consider some good and flawed alternative solutions.

Solution 1

If the EPSR Master does not have existing amf cross-links to other nodes, the easiest solution to the problem in Figure 28, is to mark all EPSR ring links as AMF cross-links; thereby placing all EPSR nodes within the same AMF 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 Domain.

Solution 2

You can see in the diagram below, that both EPSR Master ring links are AMF down-links, but all other ring links are made AMF cross-links.

Figure 29: Mark EPSR Master links as down-links and all other ring links as cross-links

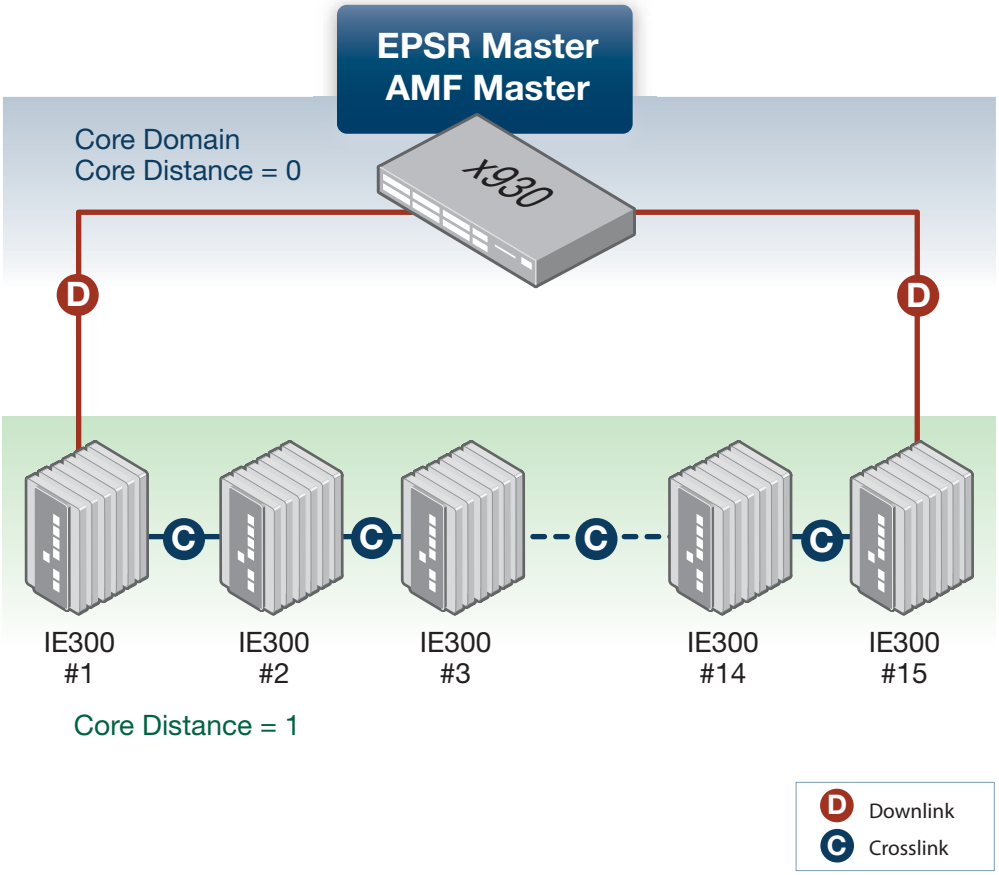

Solution 3

Building on the diagram above in Figure 29, what if your EPSR ring contains 24 nodes? How can you avoid exceeding the recommended maximum of 12 nodes per AMF domain?

You can do this by placing some of the transit nodes in the same AMF domain as the EPSR master, and some in a child domain.

Figure 30: Place some transit nodes into a child domain

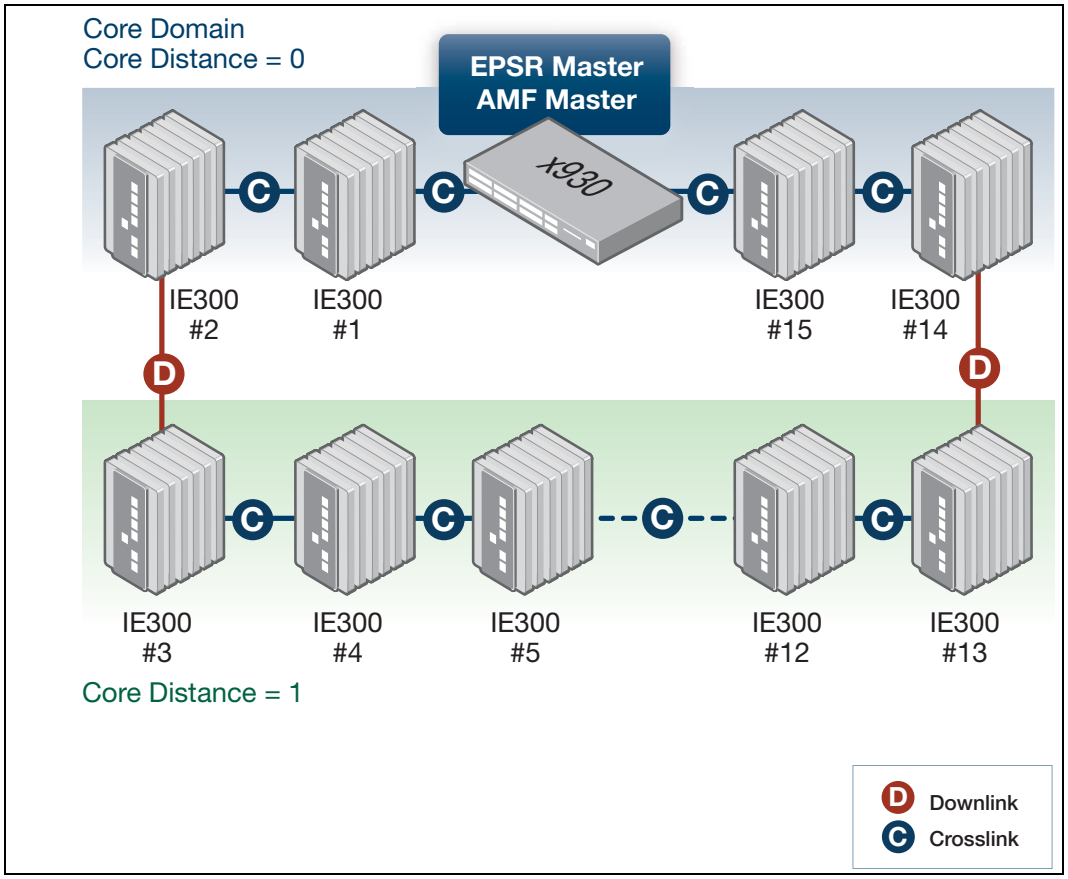

Solution 4

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 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 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 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 domain at a core distance of 2. The two uplinks (AMF down-links) from this domain each connect to different parent domains. This breaks the rule that an AMF domain can only have a **single** parent domain.

Figure 31: EPSR ring containing more than 24 nodes

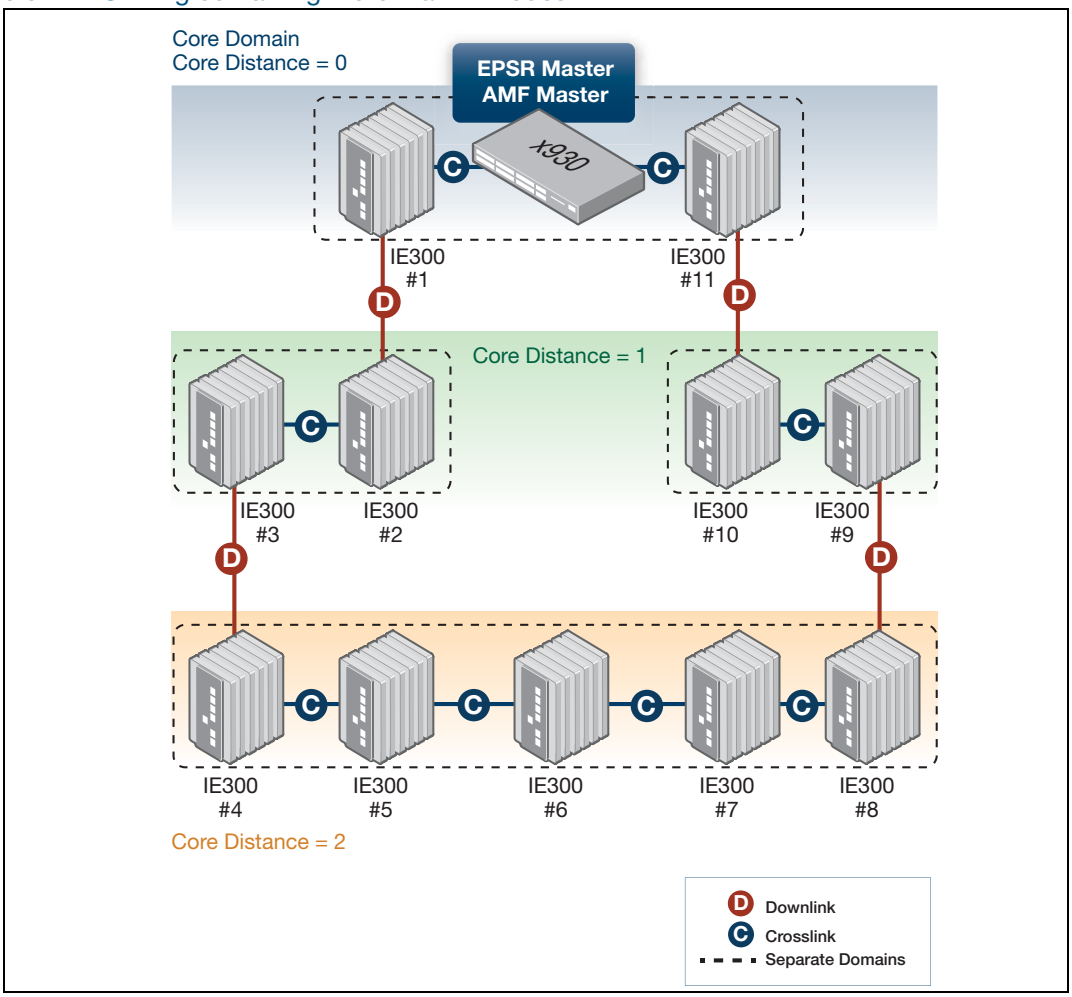

A possible workaround to this problem above is to remove the AMF down-link between nodes 8 and 9. But the obvious downside is that if there is a physical break in the EPSR ring, then the AMF master may lose connectivity with some nodes.

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 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 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 master as well as the EPSR master.

Here are three variations to consider:

1. Core domain ring with down-links to chain

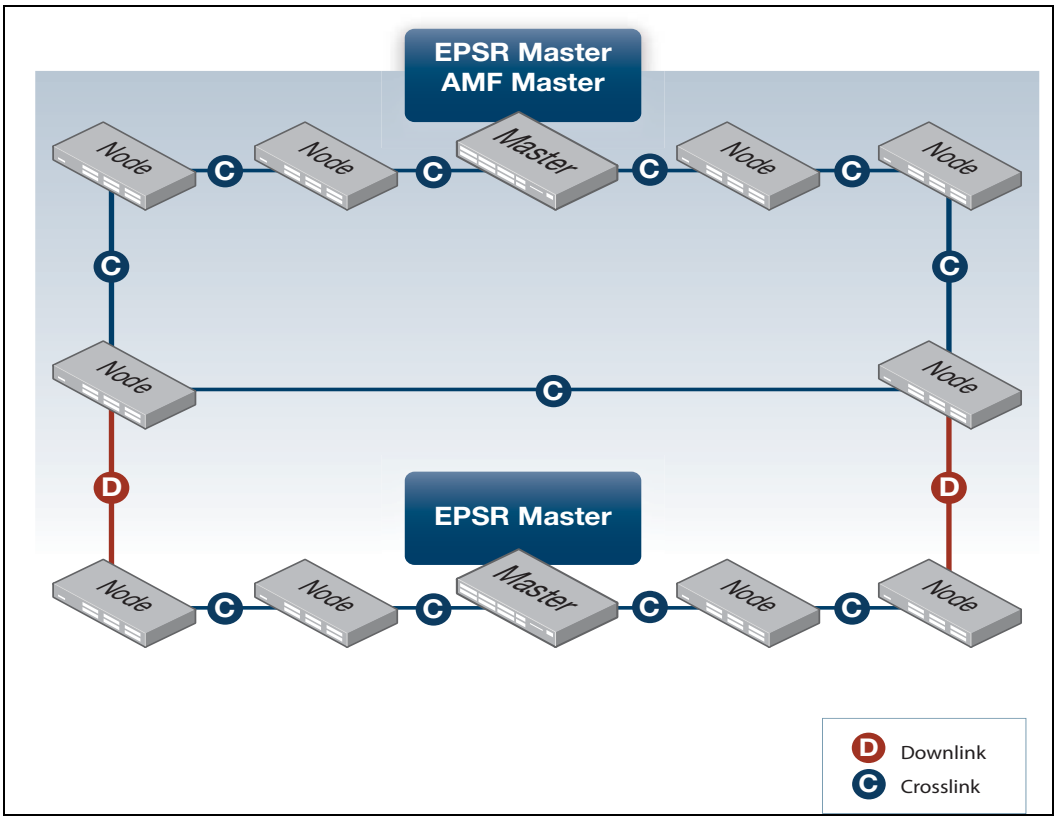

2. Core domain chain with down-links to ring

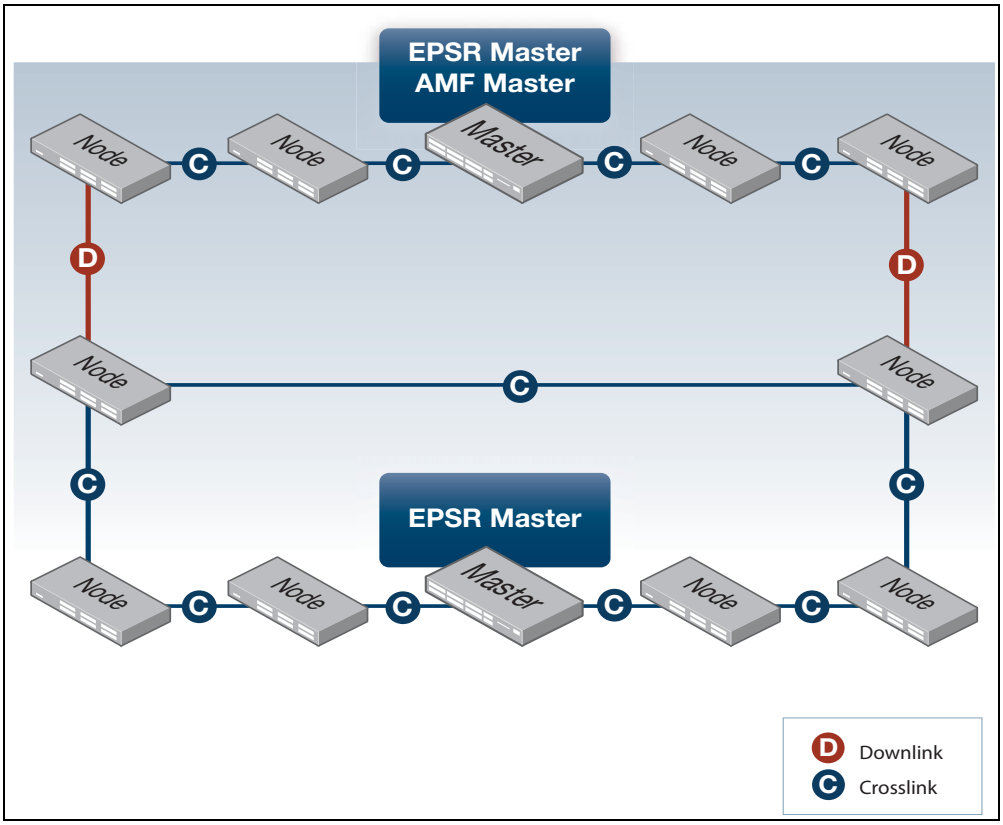

3. Core domain chain with down-links to chain domains