



NetStream Primo

Technical Description

Revision: 2.01
Update: August 2016

This document contains information that is proprietary to Netronics Technologies Inc.

No part of this publication may be reproduced, modified, or distributed without prior written authorization of Netronics Communications LLC.

This document is provided as is, without warranty of any kind.

Statement of Conditions

The information contained in this document is subject to change without notice.

Netronics shall not be liable for errors contained herein or for incidental or consequential damage in connection with the furnishing, performance, or use of this document or equipment supplied with it.

Information to User

Any changes or modifications of equipment not expressly approved by the manufacturer could void the user's authority to operate the equipment and the warranty for such equipment.

Copyright © 2016 by Netronics. All rights reserved.

Regulatory Compliance

General Note

This system has achieved Type Approval in various countries around the world. This means that the system has been tested against various local technical regulations and found to comply. The frequency bands in which the system operates may be “unlicensed” and in these bands, the system can be used provided it does not cause interference.

FCC - Compliance

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generate, uses and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.

Consult the dealer or an experienced radio/TV technician for help.

Changes or modifications to this equipment not expressly approved by the party responsible for compliance could void the user's authority to operate the equipment.



Warning

It is the responsibility of the installer to ensure that when using the outdoor antenna kits in the United States (or where **FCC** rules apply), only those antennas certified with the product are used. The use of any antenna other than those certified with the product is expressly forbidden by **FCC** rules 47 CFR part 15.204.



Warning

It is the responsibility of the installer to ensure that when configuring the radio in the United States (or where **FCC** rules apply), the Tx power is set according to the values for which the product is certified. The use of Tx power values other than those, for which the product is certified, is expressly forbidden by **FCC** rules 47 CFR part 15.204.



Caution

Outdoor units and antennas should be installed ONLY by experienced installation professionals who are familiar with local building and safety codes and, wherever applicable, are licensed by the appropriate government regulatory authorities. Failure to do so may void the product warranty and may expose the end user or the service provider to legal and financial liabilities. Resellers or distributors of this equipment are not liable for injury, damage or violation of regulations associated with the installation of outdoor units or antennas. The installer should configure the output power level of antennas according to country regulations and antenna type.



Warning

Where Outdoor units are configurable by software to Tx power values other than those for which the product is certified, it is the responsibility of the Professional Installer to restrict the Tx power to the certified limits.

This product was tested with special accessories - indoor unit (IDU or PoE), FTP CAT 5e shielded cable with sealing gasket, 12 AWG grounding cable - which must be used with the unit to insure compliance.

Indoor Units comply with part 15 of the FCC rules. Operation is subject to the following two conditions:

- These devices may not cause harmful interference.
- These devices must accept any interference received, including interference that may cause undesired operation.

Canadian Emission Requirements for Indoor Units

This Class B digital apparatus complies with Canadian ICES-003.
Cet appareil numérique de la classe B est conforme à la norme NMB-003 du Canada.

China MII

Operation of the equipment is only allowed under China MII 5.8 GHz band regulation configuration with EIRP limited to 33 dBm (2 Watt).

India WPC

Operation of the equipment is only allowed under India WPC GSR-38 for 5.8GHz band regulation configuration.

Unregulated

In countries where the radio is not regulated the equipment can be operated in any regulation configuration, best results will be obtained using Universal regulation configuration.

Safety Practices

Applicable requirements of National Electrical Code (NEC), NFPA 70; and the National Electrical Safety Code, ANSI/IEEE C2, must be considered during installation.



A Primary Protector is not required to protect the exposed wiring as long as the exposed wiring length is limited to less than or equal to 140 feet, and instructions are provided to avoid exposure of wiring to accidental contact with lightning and power conductors in accordance with NEC Sections 725-54 (c) and 800-30.

In all other cases, an appropriate Listed Primary Protector must be provided. Refer to Articles 800 and 810 of the NEC for details.

For protection of ODU against direct lightning strikes, appropriate requirements of NFPA 780 should be considered in addition to NEC.

For Canada, appropriate requirements of the CEC 22.1 including Section 60 and additional requirements of CAN/CSA-B72 must be considered as applicable.

Table of Contents

1. Synonyms and Acronyms.....	11
2. Introduction.....	13
2.1. Product Overview.....	14
2.2. Unique NetStream Primo Feature Set.....	15
2.3. System Configurations	16
2.4. New Features in Version C8.3	17
3. NetStream Primo Hardware Description	18
3.1. NetStream Primo Unit Description	19
3.1.1. Hardware Architecture.....	19
3.1.2. Interfaces.....	20
3.1.3. Management Connection for 1+1 HSB Configurations	21
3.2. PoE Injector	23
3.2.1. PoE Injector Interfaces	23
4. Activation Keys	25
4.1. Working with Activation Keys	26
4.2. Demo Mode.....	27
4.3. Activation Key-Enabled Features	28
5. Feature Description.....	32
5.1. Innovative Techniques to Boost Capacity and Reduce Latency.....	33
5.1.1. Capacity Summary.....	33
5.1.2. Header De-Duplication.....	34
5.1.3. Frame Cut-Through	36
5.1.4. Adaptive Coding Modulation (ACM)	37
5.1.5. External Protection.....	41
5.1.6. ATPC	43
5.1.7. Radio Signal Quality PMs.....	44
5.1.8. Radio Utilization PMs	44
5.2. Ethernet Features.....	45
5.2.1. Ethernet Services Overview	45
5.2.2. NetStream Primo's Ethernet Capabilities	60
5.2.3. Supported Standards.....	61
5.2.4. Ethernet Service Model.....	61
5.2.5. Ethernet Interfaces.....	77
5.2.6. Quality of Service (QoS)	86
5.2.7. Global Switch Configuration.....	111
5.2.8. Automatic State Propagation.....	112
5.2.9. Adaptive Bandwidth Notification (EOAM)	113
5.2.10. Network Resiliency.....	115
5.2.11. OAM	120
5.3. Synchronization.....	124
5.3.1. NetStream Primo Synchronization Solution.....	124
5.3.2. Available Synchronization Interfaces	124
5.3.3. Synchronous Ethernet (SyncE)	125
5.3.4. IEEE-1588v2 PTP Optimized Transport	125

5.3.5.	SSM Support and Loop Prevention	131
5.4.	AES-256 Payload Encryption	133
5.4.1.	AES Benefits.....	133
5.4.2.	NetStream Primo AES Implementation.....	133
5.4.3.	AES Interoperability	134
6.	NetStream Primo Management	135
6.1.	Management Overview.....	136
6.2.	Automatic Network Topology Discovery with LLDP Protocol	138
6.3.	Management Communication Channels and Protocols.....	139
6.4.	Web-Based Element Management System (Web EMS).....	141
6.5.	Command Line Interface (CLI).....	142
6.6.	Configuration Management.....	143
6.7.	Software Management	144
6.7.1.	Backup Software Version	144
6.8.	IPv6 Support	145
6.9.	In-Band Management	146
6.10.	Local Management.....	147
6.11.	Alarms	148
6.11.1.	Configurable BER Threshold for Alarms and Traps	148
6.11.2.	Alarms Editing	148
6.12.	NTP Support	149
6.13.	UTC Support	150
6.14.	System Security Features	151
6.14.1.	Netronics' Layered Security Concept	151
6.14.2.	Defenses in Management Communication Channels	151
6.14.3.	Defenses in User and System Authentication Procedures.....	152
6.14.4.	Secure Communication Channels	155
6.14.5.	Security Log	157
7.	Standards and Certifications	160
7.1.	Supported Ethernet Standards.....	161
7.2.	MEF Certifications for Ethernet Services	162
8.	Specifications	163
8.1.	General Radio Specifications.....	164
8.2.	Frequency Accuracy	165
8.3.	Radio Capacity Specifications.....	166
8.3.1.	3.5 MHz Channel Bandwidth (ACCP) (MRMC 1523)	166
8.3.2.	7 MHz Channel Bandwidth (ACCP) (MRMC 1508)	166
8.3.3.	14MHz Channel Bandwidth (ACCP) (MRMC 1509)	167
8.3.4.	28 MHz Channel Bandwidth (ACCP) (MRMC 1504)	167
8.3.5.	28 MHz Channel Bandwidth (ACAP) (MRMC 1505)	168
8.3.6.	40 MHz Channel Bandwidth (ACCP) (MRMC 1507)	168
8.3.7.	56 MHz Channel Bandwidth (ACCP) (MRMC 1502)	169
8.3.8.	56 MHz Channel Bandwidth (ACAP) (MRMC 1506)	169
8.3.9.	80 MHz Channel Bandwidth (ACCP) (MRMC 1501)	170
8.4.	Transmit Power Specifications	171
8.5.	Receiver Threshold Specifications.....	172
8.5.1.	Overload Thresholds	176

8.6.	Frequency Bands	177
8.7.	Mediation Device Losses	186
8.8.	Ethernet Latency Specifications	187
8.8.1.	Latency – 3.5 MHz Channel Bandwidth.....	187
8.8.2.	Latency – 7 MHz Channel Bandwidth.....	187
8.8.3.	Latency – 14 MHz Channel Bandwidth	187
8.8.4.	Latency – 28 MHz Channel Bandwidth	188
8.8.5.	Latency – 40 MHz Channel Bandwidth	188
8.8.6.	Latency – 56 MHz Channel Bandwidth	189
8.8.7.	Latency – 80 MHz Channel Bandwidth	189
8.9.	Interface Specifications	190
8.9.1.	Ethernet Interface Specifications	190
8.10.	Carrier Ethernet Functionality.....	191
8.11.	Synchronization Functionality	192
8.12.	Network Management, Diagnostics, Status, and Alarms.....	193
8.13.	Mechanical Specifications	194
8.14.	Standards Compliance.....	195
8.15.	Environmental Specifications.....	196
8.16.	Antenna Specifications.....	197
8.17.	Power Input Specifications.....	198
8.18.	Power Consumption Specifications.....	199
8.19.	Power Connection Options	200
8.20.	PoE Injector Specifications	201
8.20.1.	Power Input.....	201
8.20.2.	Environmental	201
8.20.3.	Standards Compliance.....	201
8.20.4.	Mechanical	201
8.21.	Cable Specifications.....	202
8.21.1.	Outdoor Ethernet Cable Specifications.....	202
8.21.2.	Outdoor DC Cable Specifications	203

List of Figures

NetStream Primo Rear View (Left) and Front View (Right).....	19
Cable Gland Construction	19
NetStream Primo Block Diagram	20
NetStream Primo Interfaces.....	20
Protection Signaling Cable Pinouts	22
PoE Injector	23
PoE Injector Ports.....	24
Header De-Duplication Potential Throughput Savings per Layer	35
Propagation Delay with and without Frame Cut-Through.....	36
Frame Cut-Through	37
Frame Cut-Through	37
Adaptive Coding and Modulation with 11 Working Points.....	38
NetStream Primo ACM with Adaptive Power Contrasted to Other ACM Implementations ..	40
1+1 HSB Protection	41
Internal and Local Management	42
Basic Ethernet Service Model	46
Ethernet Virtual Connection (EVC).....	46
Point to Point EVC	47
Multipoint to Multipoint EVC.....	47
Rooted Multipoint EVC.....	48
MEF Ethernet Services Definition Framework	49
E-Line Service Type Using Point-to-Point EVC.....	51
EPL Application Example	51
EVPL Application Example.....	52
E-LAN Service Type Using Multipoint-to-Multipoint EVC	52
Adding a Site Using an E-Line service.....	53
Adding a Site Using an E-LAN service.....	53
MEF Ethernet Private LAN Example	54
MEF Ethernet Virtual Private LAN Example	55
E-Tree Service Type Using Rooted-Multipoint EVC.....	55
E-Tree Service Type Using Multiple Roots	56
MEF Ethernet Private Tree Example	57
Ethernet Virtual Private Tree Example.....	58
Mobile Backhaul Reference Model	58
Packet Service Core Building Blocks.....	59
NetStream Primo Services Model	62
NetStream Primo Services Core	63
NetStream Primo Services Flow	64
Point-to-Point Service	65
Multipoint Service	65
Management Service	67
Management Service and its Service Points	69
SAPs and SNPs	70
Pipe Service Points	70
SAP, SNP and Pipe Service Points in a Microwave Network	71
Service Path Relationship on Point-to-Point Service Path	75
Physical and Logical Interfaces.....	77
Grouped Interfaces as a Single Logical Interface on Ingress Side	78

Grouped Interfaces as a Single Logical Interface on Egress Side	78
Relationship of Logical Interfaces to the Switching Fabric	82
QoS Block Diagram	87
Standard QoS and H-QoS Comparison	88
Hierarchical Classification	89
Classification Method Priorities	90
Ingress Policing Model	94
NetStream Primo Queue Manager	97
Synchronized Packet Loss.....	98
Random Packet Loss with Increased Capacity Utilization Using WRED.....	99
WRED Profile Curve.....	100
Detailed H-QoS Diagram	103
Scheduling Mechanism for a Single Service Bundle.....	106
Network Topology with NetStream Primo Units and Third-Party Equipment.....	114
ABN Entity	114
G.8032 Ring in Idle (Normal) State.....	117
G.8032 Ring in Protecting State	117
Load Balancing Example in G.8032 Ring	118
NetStream Primo End-to-End Service Management	121
SOAM Maintenance Entities (Example)	122
Ethernet Line Interface Loopback – Application Examples.....	123
IEEE-1588v2 PTP Optimized Transport – General Architecture.....	126
Calculating the Propagation Delay for PTP Packets	127
Transparent Clock – General Architecture.....	130
Transparent Clock Delay Compensation	131
AES-256 Encrypted Link.....	133
Integrated NetStream Primo Management Tools	137
Security Solution Architecture Concept.....	151

1. Synonyms and Acronyms

ACAP	Adjacent Channel Alternate Polarization
ACCP	Adjacent Channel Co-Polarization
ACM	Adaptive Coding and Modulation
AES	Advanced Encryption Standard
AIS	Alarm Indication Signal
ATPC	Automatic Tx Power Control
BER	Bit Error Ratio
BLSR	Bidirectional Line Switch Ring
BPDU	Bridge Protocol Data Units
BWA	Broadband Wireless Access
CBS	Committed Burst Size
CE	Customer Equipment
CET	Carrier-Ethernet Transport
CFM	Connectivity Fault Management
CIR	Committed Information Rate
CLI	Command Line Interface
CoS	Class of Service
CSF	Client Signal Failure
DA	Destination Address
DSCP	Differentiated Service Code Point
EBS	Excess Burst Size
EIR	Excess Information Rate
EPL	Ethernet Private Line
EVPL	Ethernet Virtual Private Line
EVC	Ethernet Virtual Connection
FEC	Forward Error Correction
FTP (SFTP)	File Transfer Protocol (Secured File Transfer Protocol)
GbE	Gigabit Ethernet
GMT	Greenwich Mean Time
HTTP (HTTPS)	Hypertext Transfer Protocol (Secured HTTP)
LAN	Local area network
LOC	Loss of Carrier
LOF	Loss Of Frame
LOS	Loss of Signal
LTE	Long-Term Evolution
MEN	Metro Ethernet Network
MPLS	Multiprotocol Label Switching
MRU	Maximum Receive Unit
MSE	Mean Square Error
MSP	Multiplex Section Protection

MSTP	Multiple Spanning Tree Protocol
MTU	Maximum Transmit Capability
NMS	Network Management System
NTP	Network Time Protocol
OAM	Operation Administration & Maintenance (Protocols)
PBB-TE	Provider Backbone Bridge Traffic Engineering
PBS	Peak Burst Rate
PDV	Packed Delay Variation
PIR	Peak Information Rate
PM	Performance Monitoring
PN	Provider Network (Port)
PTP	Precision Timing-Protocol
QoE	Quality of-Experience
QoS	Quality of Service
RBAC	Role-Based Access Control
RDI	Reverse Defect Indication
RMON	Ethernet Statistics
RSL	Received Signal Level
RSTP	Rapid Spanning Tree Protocol
SAP	Service Access Point
SFTP	Secure FTP
SLA	Service level agreements
SNMP	Simple Network Management Protocol
SNP	Service Network Point
SNR	Signal-to-Noise Ratio
SNTP	Simple Network Time Protocol
SP	Service Point
STP	Spanning Tree Protocol
SSH	Secured Shell (Protocol)
SSM	Synchronization Status Messages
SyncE	Synchronous Ethernet
TC	Traffic Class
TOS	Type of Service
UNI	User Network Interface
UTC	Coordinated Universal Time
VC	Virtual Containers
Web EMS	Web-Based Element Management System
WG	Wave guide
WFQ	Weighted Fair Queue
WRED	Weighted Random Early Detection

2.1. Product Overview

Delivering ultra-high capacity in any licensed band (6-42 GHz), NetStream Primo provides a new level of wireless capacity in current and emerging backhaul scenarios. NetStream Primo combines a compact, all-outdoor design, QPSK to 2048 QAM modulation, and an advanced feature set for Carrier Ethernet Transport that includes a sophisticated Ethernet services engine, cutting-edge header de-duplication techniques, frame cut-through, and more, to provide a compact and efficient solution for a wide variety of cost-effective deployment scenarios including macrocell backhaul, small-cell aggregation, and emerging fronthaul applications.

NetStream Primo is easily and quickly deployable compared with fiber, enabling operators to achieve faster time to new revenue streams, lower total cost of ownership, and long-term peace of mind.

NetStream Primo is an integral part of the NS Primo/Diplo family of high-capacity wireless backhaul products. Together, the NS Primo/Diplo product series provides a wide variety of backhaul solutions that can be used separately or combined to form integrated backhaul networks or network segments.

The NS Primo/Diplo series “pay-as-you-go” activation key model further enables operators to build for the future by adding capacity and functionality over time to meet the needs of network growth without requiring additional hardware.

2.2. Unique NetStream Primo Feature Set

The following table summarizes the basic NetStream Primo feature set.

NetStream Primo Feature Set

Extended Modulation Range	ACM 4-2048 QAM (11 ACM points)
Frequency Bands	6-42 GHz
Wide Range of Channels	3.5, 7, 14, 28, 40, 56, 80 MHz
Power over Ethernet (PoE)	Proprietary
Small Form Factor	(H)230mm x (W)233mm x (D)98mm
Antennas	Netronics proprietary NetStream Primo/Diplo ODU interface Direct and remote mount – standard flange
Durable All-Outdoor System	IP66-compliant

2.3. System Configurations

NetStream Primo is designed to support the following site configurations:

- 1+0 Direct/Remote Mount
- 1+1 HSB Direct/Remote Mount
- 2+0 Direct/Remote Mount

2.4. New Features in Version C8.3

The following table lists the features that have been added in NetStream OS version C8.3, and indicates where further information can be found on the new features in this manual and where configuration instructions can be found in the NetStream Primo User Manual.

New Features in Versions C8.3

Feature	Further Information	Configuration Instructions in the User's Guide
AES Radio Encryption	AES-256 Payload Encryption on page 133	Section 5.5, Configuring AES-256 Payload Encryption

3. NetStream Primo Hardware Description

This chapter describes the NetStream Primo and its components, interfaces, and mediation devices.

This chapter includes:

- NetStream Primo Unit Description
- PoE Injector


3.1. NetStream Primo Unit Description

NetStream Primo features an all-outdoor architecture consisting of a single unit directly mounted on the antenna.




The equipment is type approved and labeled according to EU Directive 1999/5/EC (R&TTE).

NetStream Primo Rear View (Left) and Front View (Right)



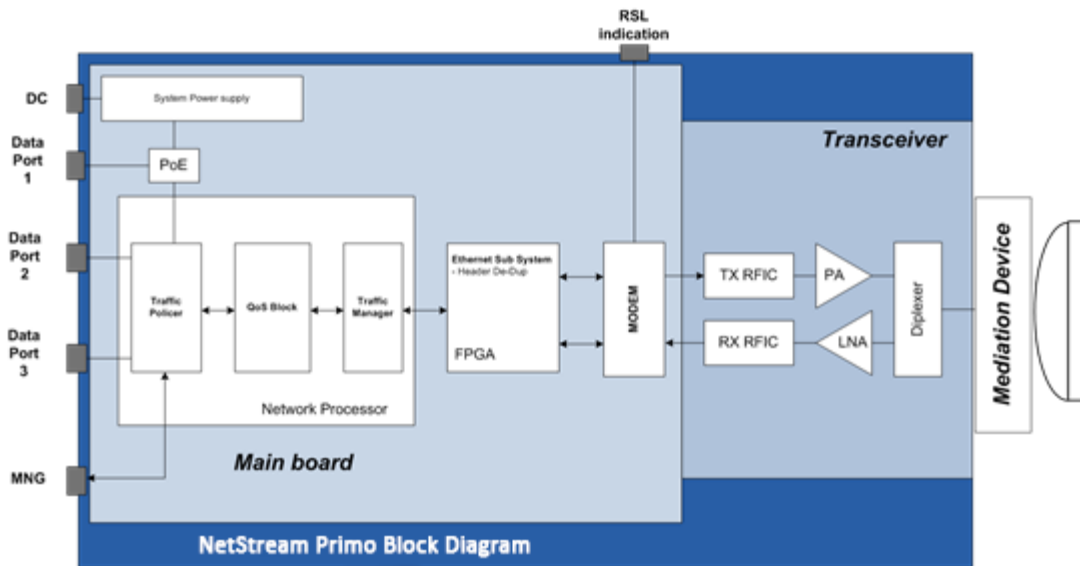
Cable Gland Construction



3.1.1. Hardware Architecture

The following diagram presents a detailed block diagram of the NetStream Primo.

NetStream Primo Block Diagram




The NetStream Primo combines full system capabilities with a very compact form-fit. The all outdoor system architecture is designed around Netronics' IP core components.

For a detailed description of the system interfaces, refer to *Interfaces* on page 20.

3.1.2. Interfaces

NetStream Primo Interfaces



- Data Port 1 for GbE traffic:
 - Electric: 10/100/1000Base-T. Supports PoE.
 - Optical: 1000Base-X (optional)

- Data Port 2 for GbE traffic:
 - Electric 10/100/1000Base-T
 - Optical: 1000Base-X (optional)
- Data Port 3 for GbE traffic
 - Electric: 10/100/1000Base-T
 - Optical: 1000Base-X (optional)

**Note**

For more details, refer to *Interface Specifications* on page 190.

- Power interface (-48VDC)
- Management Port: 10/100Base-T
- 1 RF Interface – Standard interface per frequency band
- RSL interface: BNC connector
- Grounding screw

3.1.3. Management Connection for 1+1 HSB Configurations

In HSB protection configurations, two Y-splitter cables and a special signaling cable must be used to connect the management ports (MGT/PROT) of the two NetStream Primo units and provide management access to each unit.

When Out-of-Band management is used, a splitter is required to connect the management ports to local management and to each other.

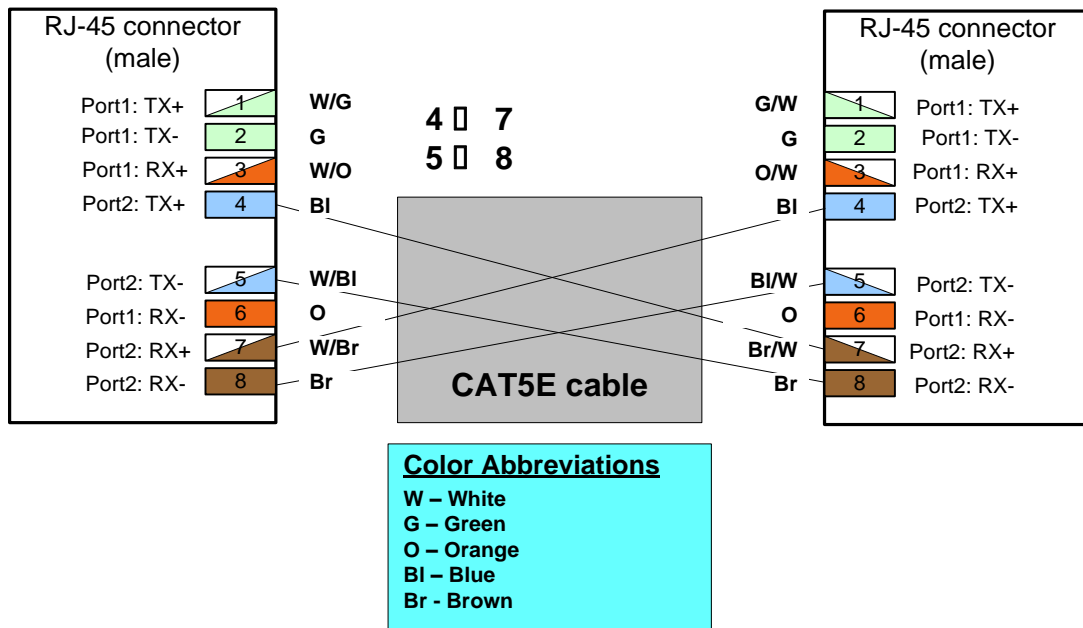
The Protection signaling cables are available pre-assembled from Netronics in various lengths, but users can also prepare them in the field.

The following sections explain how to prepare and connect these cables.

3.1.3.1. Preparing a Protection Signaling Cable

The Protection signaling cables require the following pinouts.

Protection Signaling Cable Pinouts



Other than the pinout connection described above, the cable should be repaired according to the cable preparation procedure described in the NetStream Primo Installation Guide.

3.1.3.2. Connecting the Protection Splitters and Protection Signaling Cable

Each splitter has three ports:

- System plug (“Sys”) – The system plug should be connected to the NetStream Primo’s management port.
- Management port (“Mng”) – A standard CAT5E cable should be connected to the splitter’s management port in order to utilize out-of-band (external) management.



Even for systems that use in-band management, initial configuration of an HSB protection configuration must be performed manually using out-of-band management.

- Protection signaling port (“Prot”) – A Protection signaling cross cable, as described above, should be connected between this port and the other “Prot” port of the second splitter on the mate NetStream Primo unit.


3.2. PoE Injector

The PoE injector box is designed to offer a single cable solution for connecting both data and the DC power supply to the NetStream Primo system.

To do so, the PoE injector combines 48VDC input and GbE signals via a standard CAT5E cable using a proprietary Netronics design.

The PoE injector can be ordered with a DC feed protection and with +24VDC support, as well as EMC surge protection for both indoor and outdoor installation options. It can be mounted on poles, walls, or inside racks.

PoE Injector




Two models of the PoE Injector are available:

- **PoE_Inj_AO_2DC_24V_48V** – Includes two DC power ports with power input ranges of $\pm(18-60)V$ each.
- **PoE_Inj_AO** – Includes one DC power port (DC Power Port #1), with a power input range of $\pm(40-60)V$.

3.2.1. PoE Injector Interfaces

- DC Power Port 1 $\pm(18-60)V$ or $\pm(40-60)V$
- DC Power Port 2 $\pm(18-60)V$ (Optional)
- GbE Data Port supporting 10/100/1000Base-T
- Power-Over-Ethernet (PoE) Port
- Grounding screw

PoE Injector Ports



4. Activation Keys

This chapter describes NetStream Primo's activation key model. NetStream Primo offers a pay as-you-grow concept in which future capacity growth and additional functionality can be enabled with activation keys. For purposes of the activation keys, each NetStream Primo unit is considered a distinct device. Each device contains a single activation key cipher.

This chapter includes:

- Working with Activation Keys
- Demo Mode
- Activation Key-Enabled Features

4.1. Working with Activation Keys

Netronics provides a web-based system for managing activation keys. This system enables authorized users to generate activation keys, which are generated per device serial number.

In order to upgrade an activation key, the activation key must be entered into the NetStream Primo. The system checks and implements the new activation key, enabling access to new capacities and/or features.

In the event that the activation-key-enabled capacity and feature set is exceeded, an Activation Key Violation alarm occurs and the Web EMS displays a yellow background and an activation key violation warning. After a 48-hour grace period, all other alarms are hidden until the capacity and features in use are brought within the activation key's capacity and feature set.

4.2. Demo Mode

The system can be used in demo mode, which enables all features for 60 days. Demo mode expires 60 days from the time it was activated, at which time the most recent valid activation key cipher goes into effect. The 60-day period is only counted when the system is powered up. 10 days before demo mode expires, an alarm is raised indicating to the user that demo mode is about to expire.

**Note**

Demo mode does not include AES radio encryption functionality unless a valid AES activation key has been applied for at least one carrier when demo mode is activated.

4.3. Activation Key-Enabled Features

The default (base) activation key provides each carrier with a capacity of 10 Mbps. In addition, the default activation key provides:

- A single management service.
- Unlimited Smart Pipe (L1) services.
- A single 1 x GbE port for traffic.
- Full QoS with basic queue buffer management (fixed queues with 1 Mbit buffer size limit, tail-drop only).
- LAG
- No synchronization



As described in more detail below, a CET Node activation key allows all CET service/EVC types including Smart Pipe, Point-to-Point, and Multipoint for all services, as well as an additional GbE traffic port for a total of 2 x GbE traffic ports.

As your network expands and additional functionality is desired, activation keys can be purchased for the features described in the following table.

Activation Key Types

Marketing Model	Description	For Addition Information
Capacity Refer to Capacity Activation Keys on page 30.	Enables you to increase your system’s radio capacity in gradual steps by upgrading your capacity activation key. Without a capacity activation key, each NetStream Primo unit has a capacity of 10 Mbps. Activation-key-enabled capacity is available from 50 Mbps to 650 Mbps.	Capacity Summary
NS DIPLO/PRIMO-SL-ACM	Enables the use of Adaptive Coding and Modulation (ACM) scripts.	Adaptive Coding Modulation (ACM)
NS DIPLO/PRIMO-SL-Header-DeDuplication	Enables the use of Header De-Duplication, which can be configured to operate at L2 through L4.	Header De-Duplication
NS DIPLO/PRIMO-SL-Encryption-AES256	Enables the use of AES-256 encryption for full radio payload encryption. Note that: If no AES activation key is configured for the unit and the user attempts to enable AES on a radio carrier, in addition to an Activation Key Violation alarm the feature will remain inactive and no encryption will be performed. After entering an AES activation key, the user must reset the unit before AES can be activated. Unit reset is only necessary for the first AES activation key. If AES activation keys are acquired later for additional radio carriers, unit reset is not necessary.	AES-256 Payload Encryption
NS DIPLO/PRIMO-SL-GE-Port	Enables the use of an Ethernet port in GbE mode (10/100/1000baseT or 1000baseX).. An activation key is required for each traffic port that is used on the device. Any of these activation keys can be installed multiple times with dynamic allocation inside the unit. Note: Two Ethernet ports are enabled in FE mode (10/100baseT) by default without requiring any activation key.	Interfaces
Refer to Edge CET Node Activation Keys on page 30.	Enables Carrier Ethernet Transport (CET) and a number of Ethernet services (EVCs), depending on the type of CET Node activation key: Edge CET Node – Up to 8 EVCs. Aggregation Level 1 CET Node – Up to 64 EVCs. A CET Node activation key also enables the following: A GbE traffic port in addition to the port provided by the default activation key, for a total of 2 GbE traffic ports. Network resiliency (MSTP/RSTP) for all services. Full QoS for all services including basic queue buffer management (fixed queues buffer size limit, tail-drop only) and eight queues per port, no H-QoS.	Ethernet Service Model Quality of Service (QoS)
NS DIPLO/PRIMO-SL-Network-Resiliency	Network Resiliency	Enables G.8032 for improving network resiliency.
NS DIPLO/PRIMO-SL-H-QoS	Enables H-QoS. This activation key is required to add service-bundles with dedicated queues to interfaces. Without this activation key, only the default eight queues per port are supported.	Quality of Service (QoS)
NS DIPLO/PRIMO-SL-Enh-Packet-Buffer	Enables configurable (non-default) queue buffer size limit for Green and Yellow frames. Also enables WRED. The default queue buffer size limit is 1Mbits for Green frames and 0.5 Mbits for Yellow frames.	Quality of Service (QoS)
NS DIPLO/PRIMO-SL-Sync-Unit	Enables the G.8262 synchronization unit. This activation key is required in order to provide end-to-end synchronization distribution on the physical layer. This activation key is also required to use Synchronous Ethernet (SyncE).	Synchronization
NS DIPLO/PRIMO-SL-IEEE-1588-TC	Enables IEEE-1588 Transparent Clock support. ¹	IEEE-1588v2 PTP Optimized Transport
NS DIPLO/PRIMO-SL-Frame-Cut-Through	Enables Frame Cut-Through.	Frame Cut-Through
NS DIPLO/PRIMO-SL-Secure-Management	Enables secure management protocols (SSH, HTTPS, SFTP, SNMPv3, and RADIUS).	Secure Communication Channels
NS DIPLO/PRIMO-SL-Eth-OAM-FM	Enables Connectivity Fault Management (FM) per Y.1731/ 802.1ag (CET mode only). ²	Connectivity Fault Management (FM)
NS DIPLO/PRIMO-SL-Eth-OAM-PM	Enables performance monitoring pursuant to Y.1731 (CET mode only). ³	

¹ IEEE-1588 Transparent Clock is planned for future release.

² FM support is planned for future release.

Capacity Activation Keys

Marketing Model	Description
NS DIPLO/PRIMO-SL-Capacity-50M	NS DIPLO/PRIMO SL - Capacity 50M
NS DIPLO/PRIMO-SL-Capacity-100M	NS DIPLO/PRIMO SL - Capacity 100M
NS DIPLO/PRIMO-SL-Capacity-150M	NS DIPLO/PRIMO SL - Capacity 150M
NS DIPLO/PRIMO-SL-Capacity-200M	NS DIPLO/PRIMO SL - Capacity 200M
NS DIPLO/PRIMO-SL-Capacity-225M	NS DIPLO/PRIMO SL - Capacity 225M
NS DIPLO/PRIMO-SL-Capacity-250M	NS DIPLO/PRIMO SL - Capacity 250M
NS DIPLO/PRIMO-SL-Capacity-300M	NS DIPLO/PRIMO SL - Capacity 300M
NS DIPLO/PRIMO-SL-Capacity-350M	NS DIPLO/PRIMO SL - Capacity 350M
NS DIPLO/PRIMO-SL-Capacity-400M	NS DIPLO/PRIMO SL - Capacity 400M
NS DIPLO/PRIMO-SL-Capacity-450M	NS DIPLO/PRIMO SL - Capacity 450M
NS DIPLO/PRIMO-SL-Capacity-500M	NS DIPLO/PRIMO SL - Capacity 500M
NS DIPLO/PRIMO-SL-Capacity-650M	NS DIPLO/PRIMO SL - Capacity 650M
NS DIPLO/PRIMO-SL-Upg-25M-50M	NS DIPLO/PRIMO SL - Upg 25M - 50M
NS DIPLO/PRIMO-SL-Upg-50M-100M	NS DIPLO/PRIMO SL - Upg 50M - 100M
NS DIPLO/PRIMO-SL-Upg-100M-150M	NS DIPLO/PRIMO SL - Upg 100M - 150M
NS DIPLO/PRIMO-SL-Upg-150M-200M	NS DIPLO/PRIMO SL - Upg 150M - 200M
NS DIPLO/PRIMO-SL-Upg-200M-225M	NS DIPLO/PRIMO SL - Upg 200M - 225M
NS DIPLO/PRIMO-SL-Upg-225M-250M	NS DIPLO/PRIMO SL - Upg 225M - 250M
NS DIPLO/PRIMO-SL-Upg-250M-300M	NS DIPLO/PRIMO SL - Upg 250M - 300M
NS DIPLO/PRIMO-SL-Upg-300M-350M	NS DIPLO/PRIMO SL - Upg 300M - 350M
NS DIPLO/PRIMO-SL-Upg-350M-400M	NS DIPLO/PRIMO SL - Upg 350M - 400M
NS DIPLO/PRIMO-SL-Upg-400M-450M	NS DIPLO/PRIMO SL - Upg 400M - 450M
NS DIPLO/PRIMO-SL-Upg-450M-500M	NS DIPLO/PRIMO SL - Upg 450M - 500M
NS DIPLO/PRIMO-SL-Upg-500M-650M	NS DIPLO/PRIMO SL - Upg 500M - 650M

Edge CET Node Activation Keys

Marketing Model	# of Bundled GbE Ports for User Traffic	Management Service	# of Pipe (L1) Ethernet Services	# of CET (L2) Ethernet Services
Default (No Activation Key)	1	Yes	Unlimited	-
NS DIPLO/PRIMO-SL-Edge-CET-Node	2	Yes	Unlimited	8
NS DIPLO/PRIMO-SL-Agg-Lvl-1-CET-Node	2	Yes	Unlimited	64

If a CET activation key is not generated on the NS DIPLO/PRIMO device upon initial configuration, the device uses by default a base smart pipe activation key (SL-0311-0). If the operator later wants to upgrade from the base smart pipe activation key to a CET activation key, the customer must use a CET upgrade activation key. The following table lists the CET upgrade activation keys:

³ PM support is planned for future release.

Edge CET Note Upgrade Activation Keys

Marketing Model	Upgrade From	Upgrade To
NS DIPLO/PRIMO-SL-Upg-Pipe/Edge-CET	NG Smart Pipe Activation Key (SL-0311-0)	NS DIPLO/PRIMO-SL-Edge-CET-Node (SL-0312-0)
NS DIPLO/PRIMO-SL-Upg-Edge/Agg-Lvl-1	NS DIPLO/PRIMO-SL-Edge-CET-Node (SL-0312-0)	NS DIPLO/PRIMO-SL-Agg-Lvl-1-CET-Node (SL-0313-0)

5. Feature Description

This chapter describes the main NetStream Primo features. The feature descriptions are divided into the categories listed below.

This chapter includes:

- Innovative Techniques to Boost Capacity and Reduce Latency
- Ethernet Features
- Synchronization

5.1. Innovative Techniques to Boost Capacity and Reduce Latency

NetStream Primo utilizes Netronics' innovative technology to provide a high-capacity low-latency solution. The total switching capacity of NetStream Primo is 5 Gbps or 3.125 mpps, whichever capacity limit is reached first. NetStream Primo also utilizes established Netronics technology to provide low latency, representing a 50% latency reduction for Ethernet services compared to the industry benchmark for wireless backhaul.

NetStream Primo's Header De-Duplication option enables NetStream Primo to boost capacity and provide operators with efficient spectrum utilization, with no disruption of traffic and no addition of latency.

Another of Netronics' innovative features is Frame Cut-Through, which provides unique delay and delay-variation control for delay-sensitive services. Frame Cut-Through enables high-priority frames to bypass lower priority frames even when the lower-priority frames have already begun to be transmitted. Once the high-priority frames are transmitted, transmission of the lower-priority frames is resumed with no capacity loss and no re-transmission required.

Netronics was the first to introduce hitless and errorless Adaptive Coding Modulation (ACM) to provide dynamic adjustment of the radio's modulation to account for up-to-the-minute changes in fading conditions. NetStream Primo utilizes Netronics' advanced ACM technology, and extends it to the range of QPSK to 2048 QAM.

This section includes:

- Capacity Summary
- Header De-Duplication
- Frame Cut-Through
- Adaptive Coding Modulation (ACM)
- External Protection
- ATPC
- Radio Signal Quality PMs
- Radio Utilization PMs

5.1.1. Capacity Summary

The total switching capacity of NetStream Primo is 5 Gbps or 3.125 mpps, whichever capacity limit is reached first.

A NetStream Primo unit can provide the following radio capacity:

- **Supported Channels** – 3.5/7/14/28/40/56/80 MHz channels
- **All licensed bands** – L6, U6, 7, 8, 10, 11, 13, 15, 18, 23, 26, 28, 32, 38, 42 GHz
- **High Modulation** – QPSK to 2048 QAM

For additional information:

- Radio Capacity Specifications


5.1.2. Header De-Duplication

NetStream Primo offers the option of Header De-Duplication, enabling operators to significantly improve Ethernet throughput over the radio link without affecting user traffic. Header De-Duplication can be configured to operate on various layers of the protocol stack, saving bandwidth by reducing unnecessary header overhead. Header De-duplication is also sometimes known as header compression.



Without Header De-Duplication, NetStream Primo still removes the IFG and Preamble fields. This mechanism operates automatically even if Header De-Duplication is not selected by the user.

Header De-Duplication



Header De-Duplication identifies traffic flows and replaces the header fields with a "flow ID". This is done using a sophisticated algorithm that learns unique flows by looking for repeating frame headers in the traffic stream over the radio link and compressing them. The principle underlying this feature is that frame headers in today's networks use a long protocol stack that contains a significant amount of redundant information.

Header De-Duplication can be customized for optimal benefit according to network usage. The user can determine the layer or layers on which Header De-Duplication operates, with the following options available:

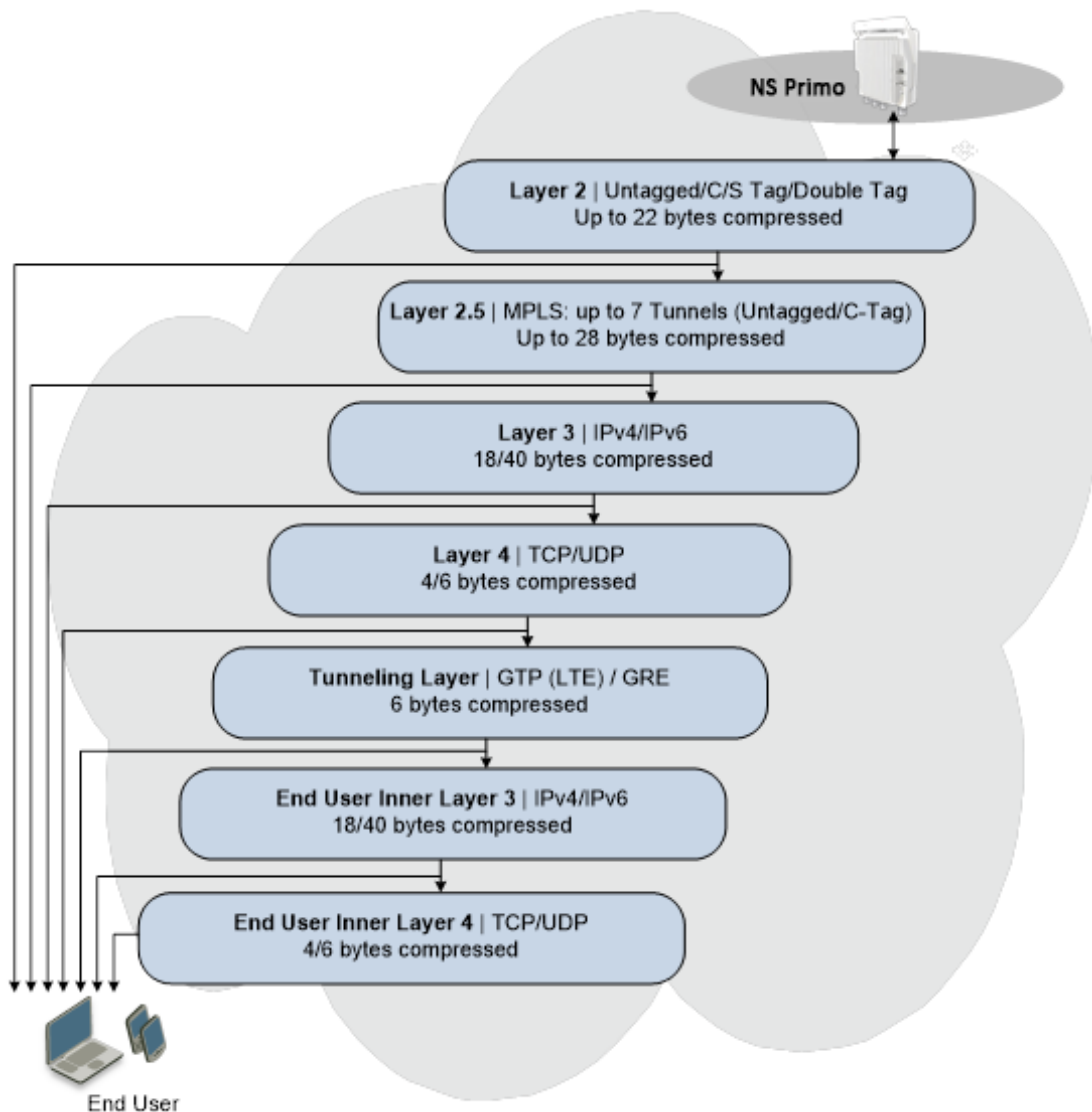
- Layer2 – Header De-Duplication operates on the Ethernet level.
- MPLS – Header De-Duplication operates on the Ethernet and MPLS levels.
- Layer3 – Header De-Duplication operates on the Ethernet and IP levels.
- Layer4 – Header De-Duplication operates on all supported layers up to Layer 4.
- Tunnel – Header De-Duplication operates on Layer 2, Layer 3, and on the Tunnel layer for packets carrying GTP or GRE frames.

- Tunnel-Layer3 – Header De-Duplication operates on Layer 2, Layer 3, and on the Tunnel and T-3 layers for packets carrying GTP or GRE frames.
- Tunnel-Layer4 – Header De-Duplication operates on Layer 2, Layer 3, and on the Tunnel, T-3, and T-4 layers for packets carrying GTP or GRE frames.

Operators must balance the depth of De-Duplication against the number of flows in order to ensure maximum efficiency. Up to 256 concurrent flows are supported.

The following graphic illustrates how Header De-Duplication can save up to 148 bytes per frame.

Header De-Duplication Potential Throughput Savings per Layer



Depending on the packet size and network topology, Header De-Duplication can increase capacity by up to:

- 50% (256 byte packets)
- 25% (512 byte packets)
- 8% (1518 byte packets)

5.1.2.1. Header De-Duplication Counters

In order to help operators optimize Header De-Duplication, NetStream Primo provides counters when Header De-Duplication is enabled. These counters include real-time information, such as the number of currently active flows and the number of flows by specific flow type. This information can be used by operators to monitor network usage and capacity, and optimize the Header De-Duplication settings. By monitoring the effectiveness of the de-duplication settings, the operator can adjust these settings to ensure that the network achieves the highest possible effective throughput.

5.1.3. Frame Cut-Through

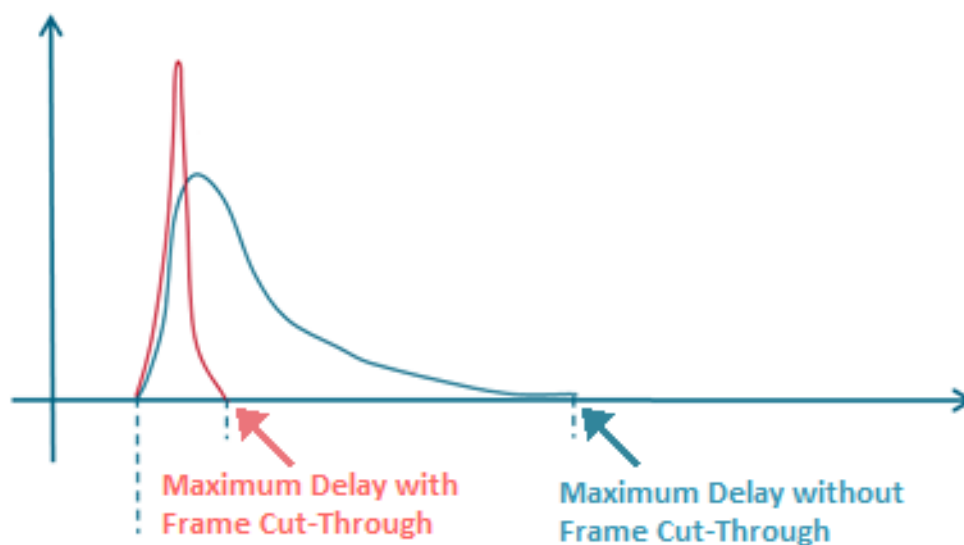
Related topics:

- Ethernet Latency Specifications
- Egress Scheduling

Frame Cut-Through is a unique and innovative feature that ensures low latency for delay-sensitive services, such as CES, VoIP, and control protocols. With Frame Cut-Through, high-priority frames are pushed ahead of lower priority frames, even if transmission of the lower priority frames has already begun. Once the high priority frame has been transmitted, transmission of the lower priority frame is resumed with no capacity loss and no re-transmission required. This provides operators with:

- Immunity to head-of-line blocking effects – key for transporting high-priority, delay-sensitive traffic.
- Reduced delay-variation and maximum-delay over the link:
 - Improved QoE for VoIP and other streaming applications.
 - Expedited delivery of critical control frames.

Propagation Delay with and without Frame Cut-Through



5.1.3.1. Frame Cut-Through Basic Operation

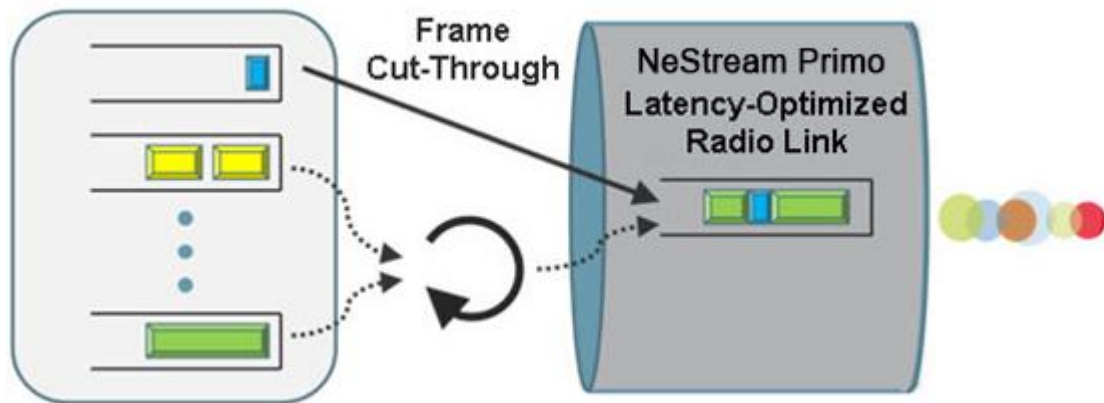
Using Frame Cut-Through, frames assigned to high priority queues can pre-empt frames already in transmission over the radio from other queues. Transmission of the pre-empted frames is resumed after the cut-through with no capacity loss or re-transmission required. This feature provides services that are sensitive to delay and delay variation, such as VoIP, with true transparency to lower priority services, by enabling the transmission of a high priority, low-delay traffic stream.

Frame Cut-Through



When enabled, Frame Cut-Through applies to all high priority frames, i.e., all frames that are classified to a CoS queue with 4th (highest) priority.

Frame Cut-Through



5.1.4. Adaptive Coding Modulation (ACM)

Related topics:

- Quality of Service (QoS)

NetStream Primo employs full-range dynamic ACM. NetStream Primo’s ACM mechanism copes with 100 dB per second fading in order to ensure high transmission quality. NetStream Primo’s ACM mechanism is designed to work with NetStream Primo’s QoS mechanism to ensure that high priority voice and data frames are never dropped, thus maintaining even the most stringent service level agreements (SLAs).

The hitless and errorless functionality of NetStream Primo’s ACM has another major advantage in that it ensures that TCP/IP sessions do not time-out. Without ACM, even interruptions as short as 50 milliseconds can lead to timeout of TCP/IP sessions, which are followed by a drastic throughput decrease while these sessions recover.


5.1.4.1. Eleven Working Points

NetStream Primo implements ACM with 11 available working points, as shown in the following table:

ACM Working Points (Profiles)

Working Point (Profile)	Modulation
Profile 0	QPSK
Profile 1	8 PSK
Profile 2	16 QAM
Profile 3	32 QAM
Profile 4	64 QAM
Profile 5	128 QAM
Profile 6	256 QAM
Profile 7	512 QAM
Profile 8	1024 QAM (Strong FEC)
Profile 9	1024 QAM (Light FEC)
Profile 10	2048 QAM

Adaptive Coding and Modulation with 11 Working Points



5.1.4.2. Hitless and Errorless Step-by Step Adjustments

ACM works as follows. Assuming a system configured for 128 QAM with ~170 Mbps capacity over a 28 MHz channel, when the receive signal Bit Error Ratio (BER) level reaches a predetermined threshold, the system pre-emptively switches to 64 QAM and the throughput is stepped down to ~140 Mbps. This is an errorless, virtually instantaneous switch. The system continues to operate at 64 QAM until the fading condition either intensifies or disappears. If the fade intensifies, another switch takes the system down to 32 QAM. If, on the other hand, the weather condition improves, the modulation is switched back to the next higher step (e.g., 128 QAM) and so on, step by step. The switching continues automatically and as quickly as needed, and can reach all the way down to QPSK during extreme conditions.

5.1.4.3. ACM Radio Scripts

An ACM radio script is constructed of a set of profiles. Each profile is defined by a modulation order (QAM) and coding rate, and defines the profile's capacity (bps). When an ACM script is activated, the system automatically chooses which profile to use according to the channel fading conditions.

The ACM TX profile can be different from the ACM RX profile.

The ACM TX profile is determined by remote RX MSE performance. The RX end is the one that initiates an ACM profile upgrade or downgrade. When MSE improves above a predefined threshold, RX generates a request to the remote TX to upgrade its profile. If MSE degrades below a predefined threshold, RX generates a request to the remote TX to downgrade its profile.

ACM profiles are decreased or increased in an errorless operation, without affecting traffic.

ACM scripts can be activated in one of two modes:

- **Fixed Mode.** In this mode, the user can select the specific profile from all available profiles in the script. The selected profile is the only profile that will be valid, and the ACM engine will be forced to be OFF. This mode can be chosen without an ACM activation key.
- **Adaptive Mode.** In this mode, the ACM engine is running, which means that the radio adapts its profile according to the channel fading conditions. Adaptive mode requires an ACM activation key.

The user can define a maximum profile. For example, if the user selects a maximum profile of 9, the system will not climb above the profile 9, even if channel fading conditions allow it.

5.1.4.4. ACM Benefits

The advantages of NetStream Primo's dynamic ACM include:

- Maximized spectrum usage
- Increased capacity over a given bandwidth
- 11 modulation/coding work points (~3 db system gain for each point change)
- Hitless and errorless modulation/coding changes, based on signal quality
- An integrated QoS mechanism that enables intelligent congestion management to ensure that high priority traffic is not affected during link fading

5.1.4.5. ACM and Built-In QoS

NetStream Primo's ACM mechanism is designed to work with NetStream Primo's QoS mechanism to ensure that high priority voice and data frames are never dropped, thus maintaining even the most stringent SLAs. Since QoS provides priority support for different classes of service, according to a wide range of criteria, you can configure NetStream Primo to discard only low priority frames as conditions deteriorate.

If you want to rely on an external switch’s QoS, ACM can work with the switch via the flow control mechanism supported in the radio.

5.1.4.6. ACM with Adaptive Transmit Power

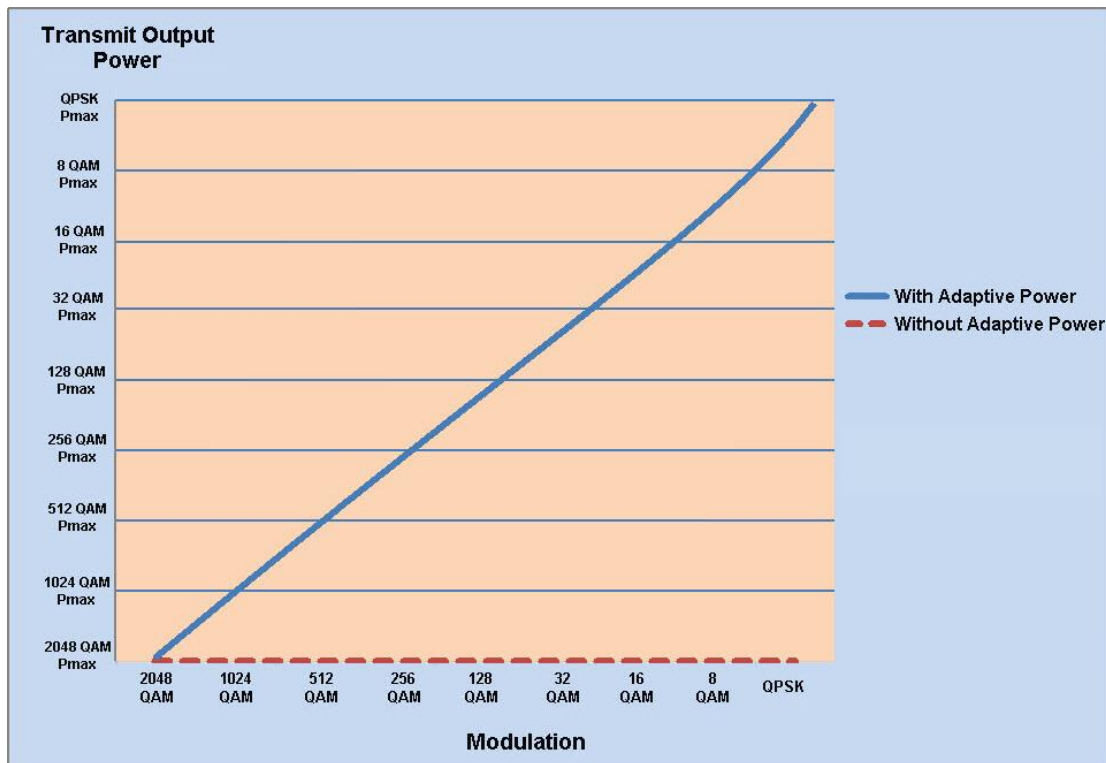
This feature requires:

- ACM script

When planning ACM-based radio links, the radio planner attempts to apply the lowest transmit power that will perform satisfactorily at the highest level of modulation. During fade conditions requiring a modulation drop, most radio systems cannot increase transmit power to compensate for the signal degradation, resulting in a deeper reduction in capacity. NetStream Primo is capable of adjusting power on the fly, and optimizing the available capacity at every modulation point.

The following figure contrasts the transmit output power achieved by using ACM with Adaptive Power to the transmit output power at a fixed power level, over an 18-23 GHz link. This figure shows how without Adaptive Transmit Power, operators that want to use ACM to benefit from high levels of modulation (e.g., 2048 QAM) must settle for low system gain, in this case, 16 dB, for all the other modulations as well. In contrast, with NetStream Primo’s Adaptive Transmit Power feature, operators can automatically adjust power levels, achieving the extra system gain that is required to maintain optimal throughput levels under all conditions.

NetStream Primo ACM with Adaptive Power Contrasted to Other ACM Implementations

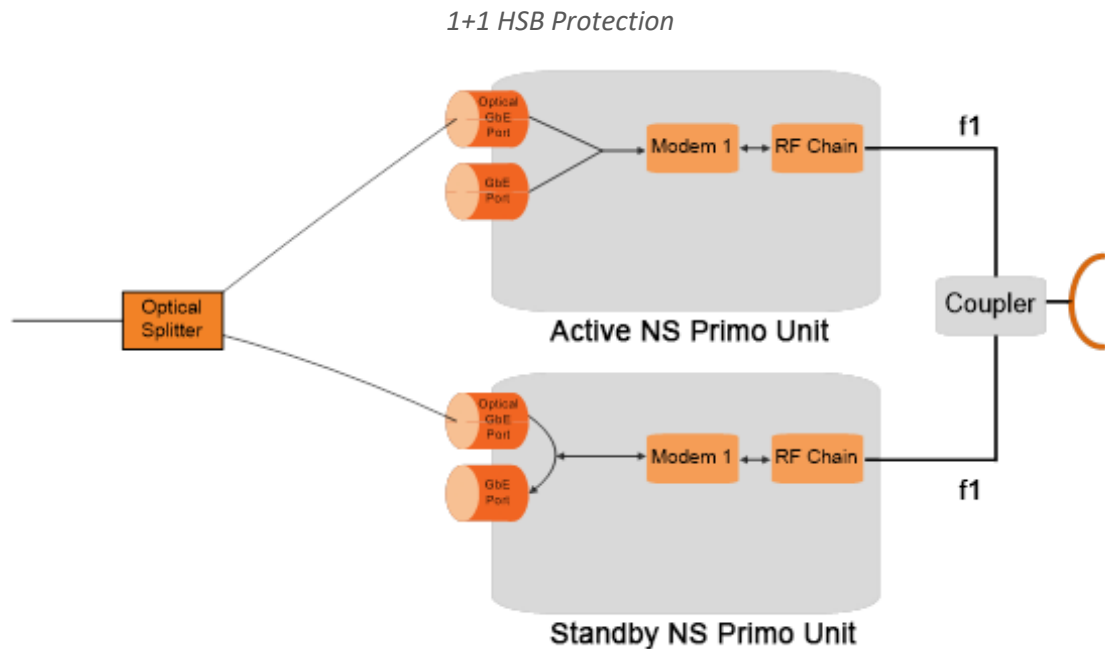


5.1.5. External Protection

1+1 HSB protection utilizes two NetStream Primo units with a single antenna, to provide hardware redundancy for Ethernet traffic. One NetStream Primo operates in active mode and the other operates in standby mode. If a protection switchover occurs, the roles are switched. The active unit goes into standby mode and the standby unit goes into active mode.

The standby unit is managed by the active unit. The standby unit’s transmitter is muted, but the standby unit’s receiver is kept on in order to monitor the link. However, the received signal is terminated at the switch level.

One GbE port on each NetStream Primo is connected to an optical splitter. Both ports on each NetStream Primo belong to a LAG, with 100% distribution to the port connected to the optical splitter on each NetStream Primo unit. Traffic must be routed to an optical GbE port on each NetStream Primo unit. No protection forwarding cable is required.



In a 1+1 HSB configuration, each NetStream Primo monitors its own radio. If the active NetStream Primo detects a radio failure, it initiates a switchover to the standby NetStream Primo.

5.1.5.1. Management for External Protection

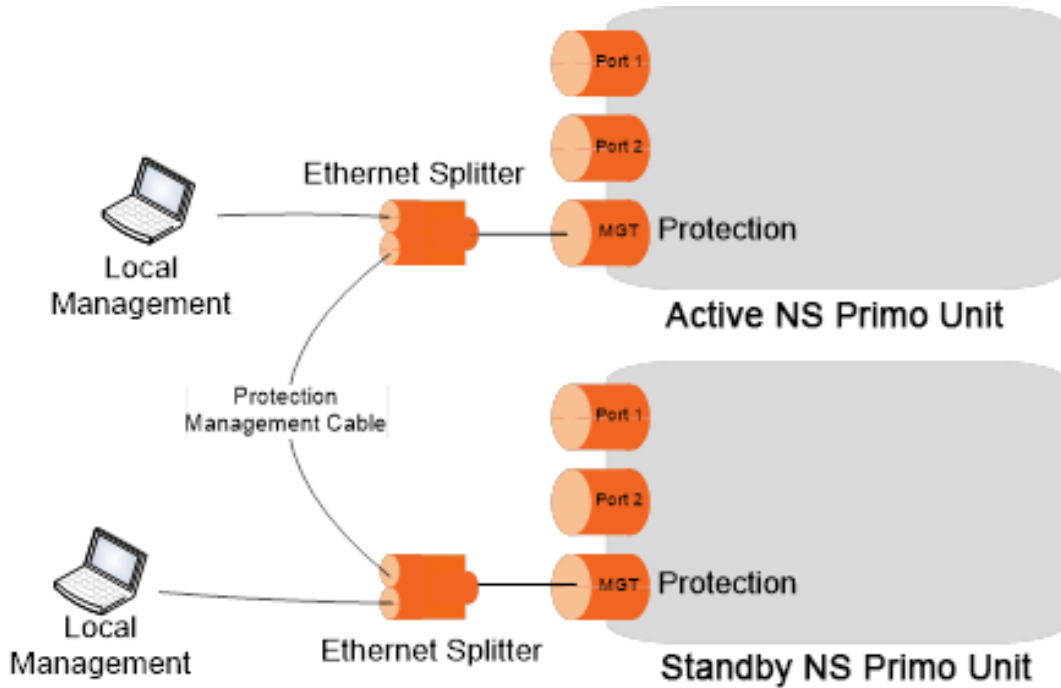
In an external protection configuration, the standby unit is managed via the active unit. A protection cable connects the two NetStream Primo units via their management ports. This cable is used for internal management. By placing an Ethernet splitter on the protection port, the user can add another cable for local management (for a detailed description, refer to *Management Connection for 1+1 HSB Configurations* on page 21). A single IP address is used for both NetStream Primo units, to ensure that management is not lost in the event of switchover.



Note

If in-band management is used, no splitter is necessary.

Internal and Local Management



The active and standby units must have the same configuration. The configuration of the active unit can be manually copied to the standby unit. Upon copying, both units are automatically reset. Therefore, it is important to ensure that the units are fully and properly configured when the system is initially brought into service.



Note

Dynamic and hitless copy-to-mate functionality is planned for future release.

5.1.5.2. Switchover

In the event of switchover, the standby unit becomes the active unit and the active unit becomes the standby unit. Switchover takes less than 50 msec.

The following events trigger switchover for HSB protection according to their priority, with the highest priority triggers listed first:

- 1 Loss of active unit
- 2 Lockout
- 3 Radio/Ethernet interface failure
- 4 Manual switch

5.1.6. ATPC

ATPC is a closed-loop mechanism by which each carrier changes the transmitted signal power according to the indication received across the link, in order to achieve a desired RSL on the other side of the link.

ATPC enables the transmitter to operate at less than maximum power for most of the time. When fading conditions occur, transmit power is increased as needed until the maximum is reached.

The ATPC mechanism has several potential advantages, including less transmitter power consumption and longer amplifier component life, thereby reducing overall system cost.

ATPC is frequently used as a means to mitigate frequency interference issues with the environment, thus allowing new radio links to be easily coordinated in frequency congested areas.

The Power Consumption Saving mode enables the system to adjust the power automatically to reduce the power used, when possible.

5.1.6.1. ATPC Override Timer



Note

ATPC Override Timer is planned for future release.

Without ATPC, if loss of frame occurs the system automatically increases its transmit power to the configured maximum. This may cause a higher level of interference with other systems until the failure is corrected.

In order to minimize this interference, some regulators require a timer mechanism which will be manually overridden when the failure is fixed. The underlying principle is that the system should start a timer from the moment maximum power has been reached. If the timer expires, ATPC is overridden and the system transmits at a pre-determined power level until the user manually re-establishes ATPC and the system works normally again.

The user can configure the following parameters:

- **Override timeout** (0 to disable the feature): The amount of time the timer counts from the moment the system transmits at the maximum configured power.
- **Override transmission power**: The power that will be transmitted if ATPC is overridden because of timeout.

The user can also display the current countdown value.

When the system enters into the override state, ATPC is automatically disabled and the system transmits at the pre-determined override power. An alarm is raised in this situation.

The only way to go back to normal operation is to manually cancel the override. When doing so, users should be sure that the problem has been corrected; otherwise, ATPC may be overridden again.

5.1.7. Radio Signal Quality PMs

NetStream Primo supports the following radio signal quality PMs. For each of these PM types, users can display the minimum and maximum values, per radio, for every 15 minute interval. Users can also define thresholds and display the number of seconds during which the radio was not within the defined threshold.

- RSL (users can define two RSL thresholds)
- TSL
- MSE
- XPI

Users can also display BER PMs and define thresholds for Excessive BER and Signal Degrade BER. Alarms are issued if these thresholds are exceeded. See *Configurable BER Threshold for Alarms and Traps* on page 148.

5.1.8. Radio Utilization PMs

NetStream Primo supports the following counters, as well as additional PMs based on these counters:

- Radio Traffic Utilization – Measures the percentage of radio capacity utilization, and used to generate the following PMs for every 15-minute interval:
 - Peak Utilization (%)
 - Average Utilization (%)
 - Over-Threshold Utilization (seconds). The utilization threshold can be defined by the user (0-100%).
- Radio Traffic Throughput – Measures the total effective Layer 2 traffic sent through the radio (Mbps), and used to generate the following PMs for every 15-minute interval:
 - Peak Throughput
 - Average Throughput
 - Over-Threshold Utilization (seconds). The threshold is defined as 0.
- Radio Traffic Capacity – Measures the total L1 bandwidth (payload plus overheads) sent through the radio (Mbps), and used to generate the following PMs for every 15-minute interval:
 - Peak Capacity
 - Average Capacity
 - Over-Threshold Utilization (seconds). The threshold is defined as 0.
- Frame Error Rate – Measures the frame error rate (%), and used to generate Frame Error Rate PMs for every 15-minute interval.

5.2. Ethernet Features

NetStream Primo features a service-oriented Ethernet switching fabric that provides a total switching capacity of up to 5 Gbps or 3.125 Mpps. NetStream Primo has an electrical GbE interface that supports PoE, and two SFP interfaces, as well as an FE interface for management. The second SFP interface can also be used for data sharing.

NetStream Primo's service-oriented Ethernet paradigm enables operators to configure VLAN definition and translation, CoS, and security on a service, service-point, and interface level.

NetStream Primo provides personalized and granular QoS that enables operators to customize traffic management parameters per customer, application, service type, or in any other way that reflects the operator's business and network requirements.

This section includes:

- Ethernet Services Overview
- Carrier-grade service resiliency (G.8032)
- NetStream Primo's Ethernet Capabilities
- Supported Standards
- Ethernet Service Model
- Ethernet Interfaces
- Quality of Service (QoS)
- Global Switch Configuration
- Automatic State Propagation
- Adaptive Bandwidth Notification (EOAM)
- Network Resiliency
- OAM


5.2.1. Ethernet Services Overview

The NetStream Primo services model is premised on supporting the standard MEF services (MEF 6, 10), and builds upon this support by the use of very high granularity and flexibility. Operationally, the NetStream Primo Ethernet services model is designed to offer a rich feature set combined with simple and user-friendly configuration, enabling users to plan, activate, and maintain any packet-based network scenario.

This section first describes the basic Ethernet services model as it is defined by the MEF, then goes on to provide a basic overview of NetStream Primo's Ethernet services implementation.

The following figure illustrates the basic MEF Ethernet services model.

Basic Ethernet Service Model



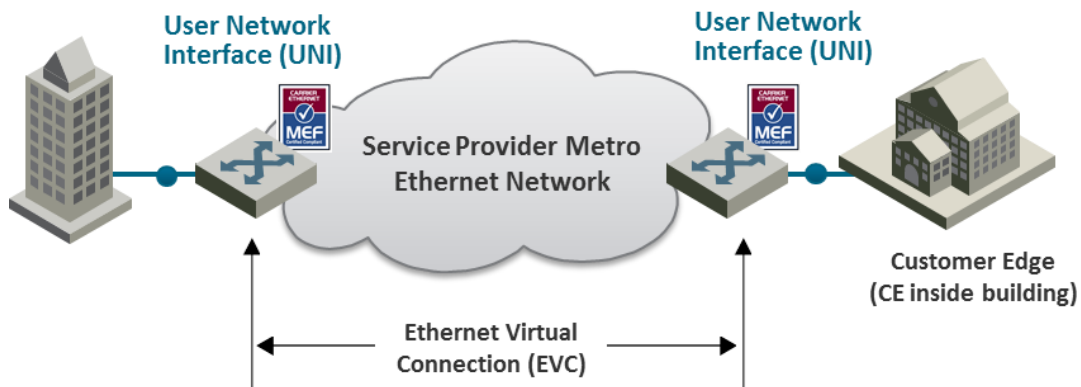
In this illustration, the Ethernet service is conveyed by the Metro Ethernet Network (MEN) provider. Customer Equipment (CE) is connected to the network at the User Network Interface (UNI) using a standard Ethernet interface (10/100 Mbps, 1 Gbps). The CE may be a router, bridge/switch, or host (end system). A NI is defined as the demarcation point between the customer (subscriber) and provider network, with a standard IEEE 802.3 Ethernet PHY and MAC.

The services are defined from the point of view of the network’s subscribers (users). Ethernet services can be supported over a variety of transport technologies and protocols in the MEN, such as SDH/SONET, Ethernet, ATM, MPLS, and GFP. However, from the user’s perspective, the network connection at the user side of the UNI is only Ethernet.

5.2.1.1. EVC

Subscriber services extend from UNI to UNI. Connectivity between UNIs is defined as an Ethernet Virtual Connection (EVC), as shown in the following figure.

Ethernet Virtual Connection (EVC)



An EVC is defined by the MEF as an association of two or more UNIs that limits the exchange of service frames to UNIs in the Ethernet Virtual Connection. The EVC perform two main functions:

- Connects two or more customer sites (UNIs), enabling the transfer of Ethernet frames between them.
- Prevents data transfer involving customer sites that are not part of the same EVC. This feature enables the EVC to maintain a secure and private data channel.

A single UNI can support multiple EVCs via the Service Multiplexing attribute. An ingress service frame that is mapped to the EVC can be delivered to one or more of the UNIs in the EVC, other than the ingress UNI. It is vital to avoid delivery back to the ingress UNI, and to avoid delivery to a UNI that does not belong to the EVC. An EVC is always bi-directional in the sense that ingress service frames can originate at any UNI in an EVC.


Service frames must be delivered with the same Ethernet MAC address and frame structure that they had upon ingress to the service. In other words, the frame must be unchanged from source to destination, in contrast to routing in which headers are discarded. Based on these characteristics, an EVC can be used to form a Layer 2 private line or Virtual Private Network (VPN).

One or more VLANs can be mapped (bundled) to a single EVC.

The MEF has defined three types of EVCs:


- 1 **Point to Point EVC** – Each EVC contains exactly two UNIs. The following figure shows two point-to-point EVCs connecting one site to two other sites.

Point to Point EVC




- 2 **Multipoint (Multipoint-to-Multipoint) EVC** – Each EVC contains two or more UNIs. In the figure below, three sites belong to a single Multipoint EVC and can forward Ethernet frames to each other.

Multipoint to Multipoint EVC



- 3 **Rooted Multipoint EVC (Point-to-Multipoint)** – Each EVC contains one or more UNIs, with one or more UNIs defined as Roots, and the others defined as Leaves. The Roots can forward frames to the Leaves. Leaves can only forward frames to the Roots, but not to other Leaves.

Rooted Multipoint EVC



In the NetStream Primo, an EVC is defined by either a VLAN or by Layer 1 connectivity (Pipe Mode).

5.2.1.2. Bandwidth Profile

The bandwidth profile (BW profile) is a set of traffic parameters that define the maximum limits of the customer's traffic.

At ingress, the bandwidth profile limits the traffic transmitted into the network:

- Each service frame is checked against the profile for compliance with the profile.
- Bandwidth profiles can be defined separately for each UNI (MEF 10.2).
- Service frames that comply with the bandwidth profile are forwarded.
- Service frames that do not comply with the bandwidth profile are dropped at the ingress interface.

The MEF has defined the following three bandwidth profile service attributes:

- Ingress BW profile per ingress UNI
- Ingress BW profile per EVC
- Ingress BW profile per CoS identifier

The BW profile service attribute consists of four traffic parameters:

- CIR (Committed Information Rate)
- CBS (Committed Burst Size)
- EIR (Excess Information Rate)
- EBS (Excess Burst Size)

Bandwidth profiles can be applied per UNI, per EVC at the UNI, or per CoS identifier for a specified EVC at the UNI.

The Color of the service frame is used to determine its bandwidth profile. If the service frame complies with the CIR and EIR defined in the bandwidth profile, it is marked Green. In this case, the average and maximum service frame rates are less than or equal to the CIR and CBS, respectively.

If the service frame does not comply with the CIR defined in the bandwidth profile, but does comply with the EIR and EBS, it is marked Yellow. In this case, the average service frame rate is greater than the CIR but less than the EIR, and the maximum service frame size is less than the EBS.

If the service frame fails to comply with both the CIR and the EIR defined in the bandwidth profile, it is marked Red and discarded.

In the NetStream Primo, bandwidth profiles are constructed using a full standardized TrTCM policer mechanism.

5.2.1.3. Ethernet Services Definitions

The MEF provides a model for defining Ethernet services. The purpose of the MEF model is to help subscribers better understand the variations among different types of Ethernet services. NetStream Primo supports a variety of service types defined by the MEF. All of these service types share some common attributes, but there are also differences as explained below.

Ethernet service types are generic constructs used to create a broad range of services. Each Ethernet service type has a set of Ethernet service attributes that define the characteristics of the service. These Ethernet service attributes in turn are associated with a set of parameters that provide various options for the various service attributes.

MEF Ethernet Services Definition Framework



The MEF defines three generic Ethernet service type constructs, including their associated service attributes and parameters:

- Ethernet Line (E-Line)
- Ethernet LAN (E-LAN)
- Ethernet Tree (E-Tree)

Multiple Ethernet services are defined for each of the three generic Ethernet service types. These services are differentiated by the method for service identification used at the UNIs. Services using All-to-One Bundling UNIs (port-based) are referred to as “Private” services, while services using Service Multiplexed (VLAN-based) UNIs are referred to as “Virtual Private” services. This relationship is shown in the following table.

MEF-Defined Ethernet Service Types

Service Type	Port Based (All to One Bundling)	VLAN-BASED (EVC identified by VLAN ID)
E-Line (Point-to-Point EVC)	Ethernet Private Line (EPL)	Ethernet Virtual Private Line (EVPL)
E-LAN (Multipoint-to-Multipoint EVC)	Ethernet Private LAN (EP-LAN)	Ethernet Virtual Private LAN (EVP-LAN)
E-Tree (Rooted Multipoint EVC)	Ethernet Private Tree (EP-Tree)	Ethernet Virtual Private Tree (EVP-Tree)

All-to-One Bundling refers to a UNI attribute in which all Customer Edge VLAN IDs (CE-VLAN IDs) entering the service via the UNI are associated with a single EVC. Bundling refers to a UNI attribute in which more than one CE-VLAN ID can be associated with an EVC.


To fully specify an Ethernet service, additional service attributes must be defined in addition to the UNI and EVC service attributes. These service attributes can be grouped under the following categories:

- Ethernet physical interfaces
- Traffic parameters
- Performance parameters
- Class of service
- Service frame delivery
- VLAN tag support
- Service multiplexing
- Bundling
- Security filters

5.2.1.4. E-Line Service

The Ethernet line service (E-Line service) provides a point-to-point Ethernet Virtual Connection (EVC) between two UNIs. The E-Line service type can be used to create a broad range of Ethernet point-to-point services and to maintain the necessary connectivity. In its simplest form, an E-Line service type can provide symmetrical bandwidth for data sent in either direction with no performance assurances, e.g., best effort service between two FE UNIs. In more sophisticated forms, an E-Line service type can provide connectivity between two UNIs with different line rates and can be defined with performance assurances such as CIR with an associated CBS, EIR with an associated EBS, delay, delay variation, loss, and availability for a given Class of Service (CoS) instance. Service multiplexing can occur at one or both UNIs in the EVC. For example, more than one point-to-point EVC can be offered on the same physical port at one or both of the UNIs.


E-Line Service Type Using Point-to-Point EVC



5.2.1.5. Ethernet Private Line Service

An Ethernet Private Line (EPL) service is specified using an E-Line Service type. An EPL service uses a point-to-point EVC between two UNIs and provides a high degree of transparency for service frames between the UNIs that it interconnects such that the service frame’s header and payload are identical at both the source and destination UNI when the service frame is delivered (L1 service). A dedicated UNI (physical interface) is used for the service and service multiplexing is not allowed. All service frames are mapped to a single EVC at the UNI. In cases where the EVC speed is less than the UNI speed, the CE is expected to shape traffic to the ingress bandwidth profile of the service to prevent the traffic from being discarded by the service. The EPL is a port-based service, with a single EVC across dedicated UNIs providing site-to-site connectivity. EPL is the most popular Ethernet service type due to its simplicity, and is used in diverse applications such as replacing a TDM private line.

EPL Application Example



5.2.1.6. Ethernet Virtual Private Line Service

An Ethernet Virtual Private Line (EVPL) is created using an E-Line service type. An EVPL can be used to create services similar to EPL services. However, several characteristics differ between EPL and EVPL services.

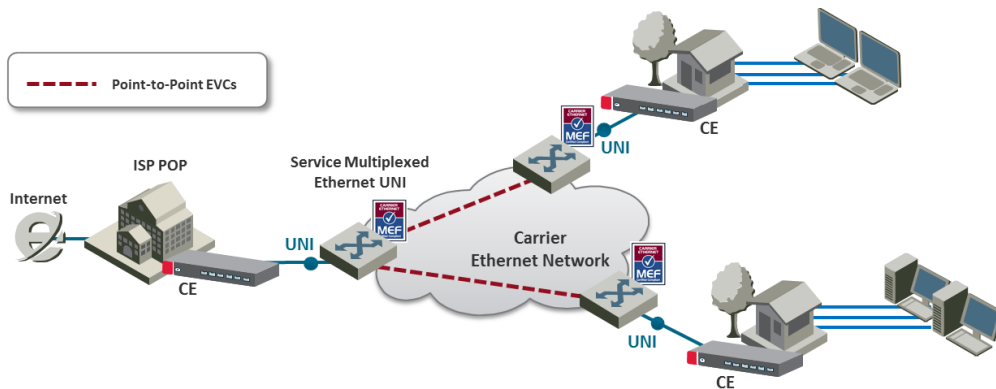
First, an EVPL provides for service multiplexing at the UNI, which means it enables multiple EVCs to be delivered to customer premises over a single physical

connection (UNI). In contrast, an EPL only enables a single service to be delivered over a single physical connection.

Second, the degree of transparency for service frames is lower in an EVPL than in an EPL.

Since service multiplexing is permitted in EVPL services, some service frames may be sent to one EVC while others may be sent to other EVCs. EVPL services can be used to replace Frame Relay and ATM L2 VPN services, in order to deliver higher bandwidth, end-to-end services.

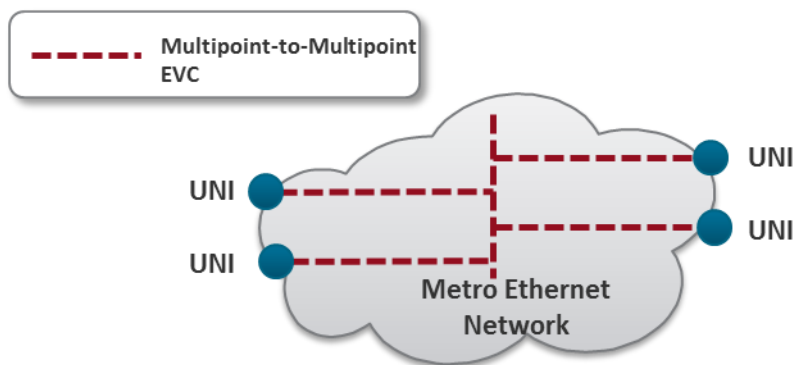
EVPL Application Example



5.2.1.7. E-LAN Service

The E-LAN service type is based on Multipoint to Multipoint EVCs, and provides multipoint connectivity by connecting two or more UNIs. Each site (UNI) is connected to a multipoint EVC, and customer frames sent from one UNI can be received at one or more UNIs. If additional sites are added, they can be connected to the same multipoint EVC, simplifying the service activation process. Logically, from the point of view of a customer using an E-LAN service, the MEN can be viewed as a LAN.

E-LAN Service Type Using Multipoint-to-Multipoint EVC



The E-LAN service type can be used to create a broad range of services. In its basic form, an E-LAN service can provide a best effort service with no performance assurances between the UNIs. In more sophisticated forms, an E-LAN service type can be defined with performance assurances such as CIR with an associated CBS, EIR


with an associated EBS, delay, delay variation, loss, and availability for a given CoS instance.

For an E-LAN service type, service multiplexing may occur at none, one, or more than one of the UNIs in the EVC. For example, an E-LAN service type (Multipoint-to-Multipoint EVC) and an E-Line service type (Point-to-Point EVC) can be service multiplexed at the same UNI. In such a case, the E-LAN service type can be used to interconnect other customer sites while the E-Line service type is used to connect to the Internet, with both services offered via service multiplexing at the same UNI.

E-LAN services can simplify the interconnection among a large number of sites, in comparison to hub/mesh topologies implemented using point-to-point networking technologies such as Frame Relay and ATM.


For example, consider a point-to-point network configuration implemented using E-Line services. If a new site (UNI) is added, it is necessary to add a new, separate EVC to all of the other sites in order to enable the new UNI to communicate with the other UNIs, as shown in the following figure.

Adding a Site Using an E-Line service



In contrast, when using an E-LAN service, it is only necessary to add the new UNI to the multipoint EVC. No additional EVCs are required, since the E-LAN service uses a multipoint to multipoint EVC that enables the new UNI to communicate with each of the others UNIs. Only one EVC is required to achieve multi-site connectivity, as shown in the following figure.

Adding a Site Using an E-LAN service

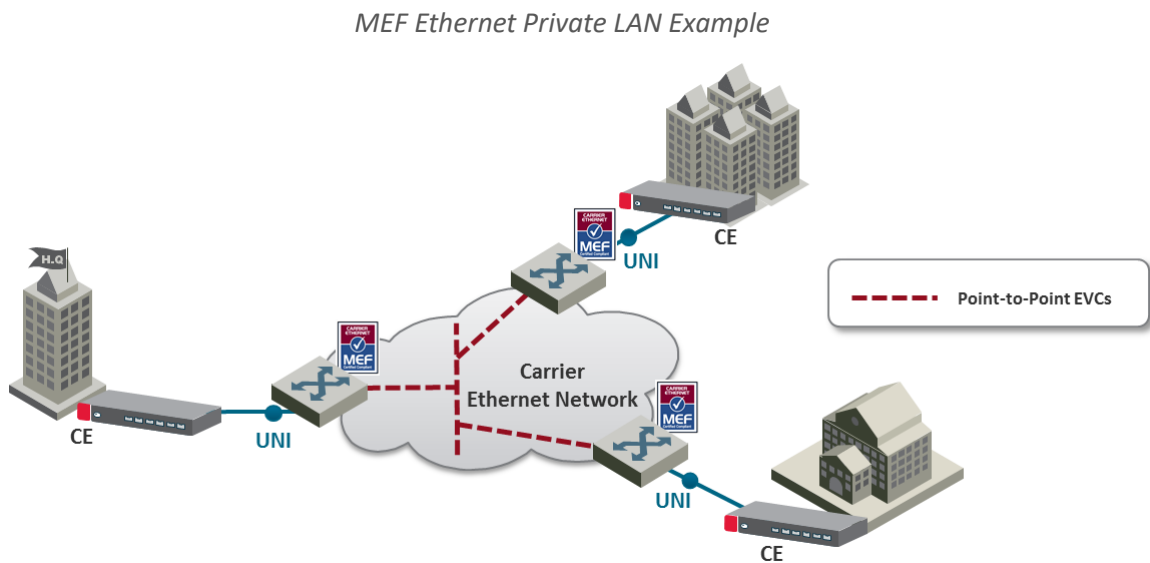


The E-LAN service type can be used to create a broad range of services, such as private LAN and virtual private LAN services.

5.2.1.8. Ethernet Private LAN Service

It is often desirable to interconnect multiple sites using a Local Area Network (LAN) protocol model and have equivalent performance and access to resources such as servers and storage. Customers commonly require a highly transparent service that connects multiple UNIs. The Ethernet Private LAN (EP-LAN) service is defined with this in mind, using the E-LAN service type. The EP-LAN is a Layer 2 service in which each UNI is dedicated to the EP-LAN service. A typical use case for EP-LAN services is Transparent LAN.

The following figure shows an example of an EP-LAN service in which the service is defined to provide Customer Edge VLAN (CE-VLAN) tag preservation and tunneling for key Layer 2 control protocols. Customers can use this service to configure VLANs across the sites without the need to coordinate with the service provider. Each interface is configured for All-to-One Bundling, which enables the EP-LAN service to support CE-VLAN ID preservation. In addition, EP-LAN supports CE-VLAN CoS preservation.



5.2.1.9. Ethernet Virtual Private LAN Service

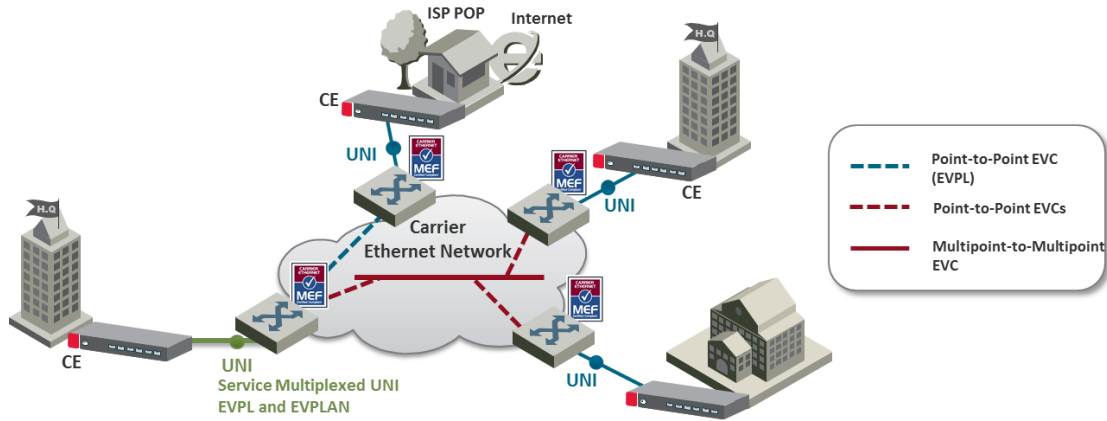
Customers often use an E-LAN service type to connect their UNIs in an MEN, while at the same time accessing other services from one or more of those UNIs. For example, a customer might want to access a public or private IP service from a UNI at the customer site that is also used to provide E-LAN service among the customer's several metro locations. The Ethernet Virtual Private LAN (EVP-LAN) service is defined to address this need. EVP-LAN is actually a combination of EVPL and E-LAN.

Bundling can be used on the UNIs in the Multipoint-to-Multipoint EVC, but is not mandatory. As such, CE-VLAN tag preservation and tunneling of certain Layer 2 control protocols may or may not be provided. Service multiplexing is allowed on

each UNI. A typical use case would be to provide Internet access a corporate VPN via one UNI.

The following figure provides an example of an EVP-LAN service.

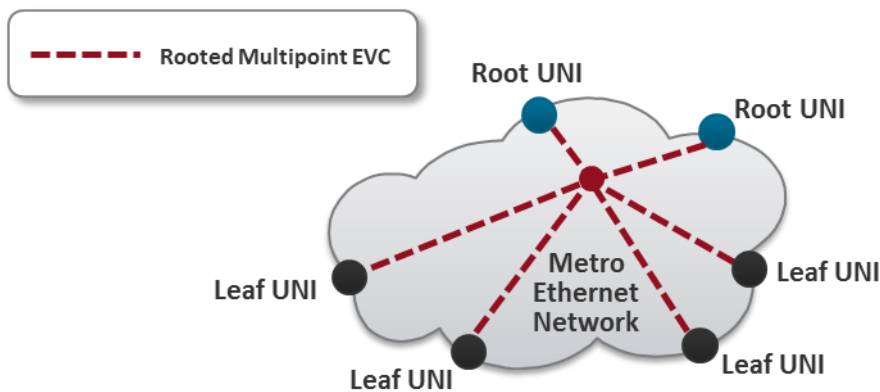
MEF Ethernet Virtual Private LAN Example



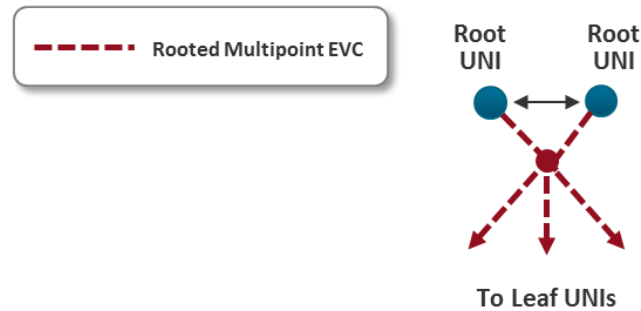
5.2.1.10. E-Tree Service

The E-Tree service type is an Ethernet service type that is based on Rooted-Multipoint EVCs. In its basic form, an E-Tree service can provide a single Root for multiple Leaf UNIs. Each Leaf UNI can exchange data with only the Root UNI. A service frame sent from one Leaf UNI cannot be delivered to another Leaf UNI. This service can be particularly useful for Internet access, and video-over-IP applications such as multicast/broadcast packet video. One or more CoS values can be associated with an E-Tree service.

E-Tree Service Type Using Rooted-Multipoint EVC



Two or more Root UNIs can be supported in advanced forms of the E-Tree service type. In this scenario, each Leaf UNI can exchange data only with the Root UNIs. The Root UNIs can communicate with each other. Redundant access to the Root can also be provided, effectively allowing for enhanced service reliability and flexibility.

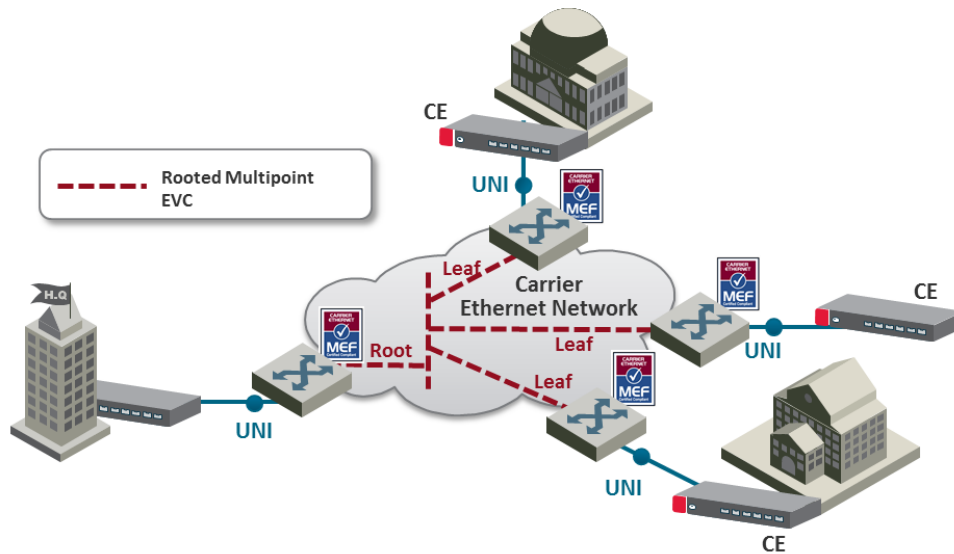
E-Tree Service Type Using Multiple Roots

Service multiplexing is optional and may occur on any combination of UNIs in the EVC. For example, an E-Tree service type using a Rooted-Multipoint EVC, and an E-Line service type using a Point-to-Point EVC, can be service multiplexed on the same UNI. In this example, the E-Tree service type can be used to support a specific application at the Subscriber UNI, e.g., ISP access to redundant PoPs (multiple Roots at ISP PoPs), while the E-Line Service type is used to connect to another enterprise site with a Point-to-Point EVC.

5.2.1.11. Ethernet Private Tree Service

The Ethernet Private Tree service (EP-Tree) is designed to supply the flexibility for configuring multiple sites so that the services are distributed from a centralized site, or from a few centralized sites. In this setup, the centralized site or sites are designed as Roots, while the remaining sites are designated as Leaves. CE-VLAN tags are preserved and key Layer 2 control protocols are tunneled. The advantage of such a configuration is that the customer can configure VLANs across its sites without the need to coordinate with the service provider. Each interface is configured for All-to-One Bundling, which means that EP-Tree services support CE-VLAN ID preservation. EP-Tree also supports CE-VLAN CoS preservation. EP-Tree requires dedication of the UNIs to the single EP-Tree service.

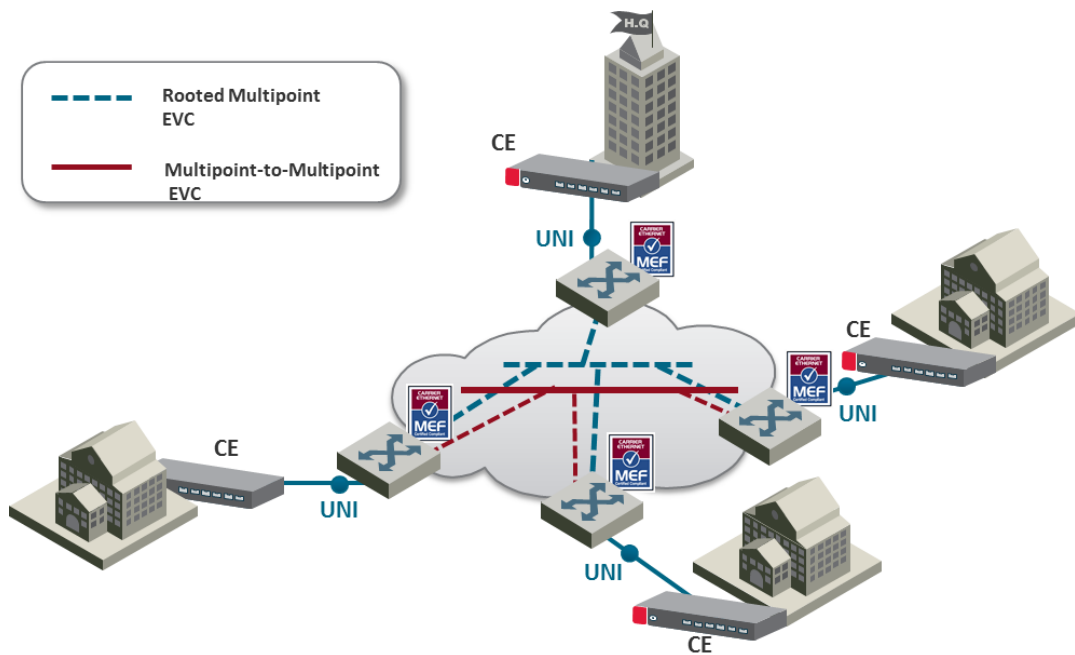
The following figure provides an example of an EP-Tree service.

MEF Ethernet Private Tree Example

5.2.1.12. Ethernet Virtual Private Tree Service

In order to access several applications and services from well-defined access points (Root), the UNIs are attached to the service in a Rooted Multipoint connection. Customer UNIs can also support other services, such as EVPL and EVP-LAN services. An EVP-Tree service is used in such cases. Bundling can be used on the UNIs in the Rooted Multipoint EVC, but it is not mandatory. As such, CE-VLAN tag preservation and tunneling of certain Layer 2 Control Protocols may or may not be provided. EVP-Tree enables each UNI to support multiple services. A good example would be a customer that has an EVP-LAN service providing data connectivity among three UNIs, while using an EVP-Tree service to provide video broadcast from a video hub location. The following figure provides an example of a Virtual Private Tree service.


Ethernet Virtual Private Tree Example



NetStream Primo enables network connectivity for **Mobile Backhaul** cellular infrastructure, fixed networks, private networks and enterprises.

Mobile Backhaul refers to the network between the Base Station sites and the Network Controller/Gateway sites for all generations of mobile technologies. Mobile equipment and networks with ETH service layer functions can support MEF Carrier Ethernet services using the service attributes defined by the MEF.

Mobile Backhaul Reference Model



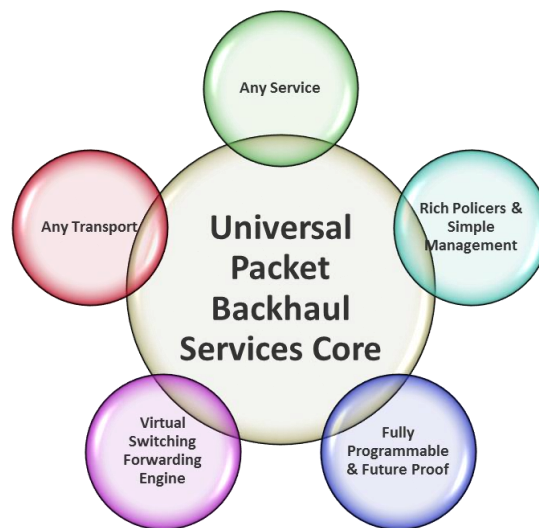
The NetStream Primo services concept is purpose built to support the standard MEF services for mobile backhaul (MEF 22, mobile backhaul implementation agreement), as an addition to the baseline definition of MEF Services (MEF 6) using service attributes (as well as in MEF 10). E-Line, E-LAN and E-Tree services are well defined as the standard services.

5.2.1.13. NetStream Primo Universal Packet Backhaul Services Core

NetStream Primo addresses the customer demand for multiple services of any of the aforementioned types (EPL, EVPL, EP –LAN, EVP-LAN, EP-Tree, and EVP-Tree) through its rich service model capabilities and flexible integrated switch application. Additional Layer 1 point-based services are supported as well, as explained in more detail below.

Services support in the mobile backhaul environment is provided using the NetStream Primo services core, which is structured around the building blocks shown in the figure below. NetStream Primo provides rich and secure packet backhaul services over any transport type with unified, simple, and error-free operation.

Packet Service Core Building Blocks



5.2.1.14. Any Service

- Ethernet services (EVCs)
 - E-Line (Point-to-Point)
 - E-LAN (Multipoint)
 - E-Tree (Point-to-Multipoint)⁴
- Port based (Smart Pipe) services

5.2.1.15. Any Transport

- Native Ethernet (802.1Q/Q-in-Q)
- Any topology and any mix of radio and fiber interfaces
- Seamless interworking with any optical network (NG-SDH, packet optical transport, IP/MPLS service/VPN routers)

⁴ E-Tree services are planned for future release.

5.2.1.16. Virtual Switching/Forwarding Engine

- Clear distinction between user facing service interfaces (UNI) and intra-network interfaces
- Fully flexible C-VLAN and S-VLAN encapsulation (classification/preservation/ translation)
- Improved security/isolation without limiting C-VLAN reuse by different customers
- Per-service MAC learning with 128K MAC addresses support

5.2.1.17. Fully Programmable and Future-Proof

- Network-processor-based services core
- Ready today to support emerging and future standards and networking protocols

5.2.1.18. Rich Policies and Tools with Unified and Simplified Management

- Personalized QoS (H-QoS)⁵
- Superb service OAM (FM, PM)⁶
- Carrier-grade service resiliency (G.8032)⁷

5.2.2. NetStream Primo's Ethernet Capabilities

NetStream Primo is built upon a service-based paradigm that provides rich and secure frame backhaul services over any type of transport, with unified, simple, and error-free operation. NetStream Primo's services core includes a rich set of tools that includes:

- Service-based Quality of Service (QoS).
- Service OAM, including granular PMs, and service activation.
- Carrier-grade service resiliency using G.8032⁸

The following are NetStream Primo's main Carrier Ethernet transport features. This rich feature set provides a future-proof architecture to support backhaul evolution for emerging services.

- Up to 64 services
- Up to 32 service points per service
- All service types:⁹
 - Multipoint (E-LAN)

⁵ H-QoS support is planned for future release.

⁶ PM support is planned for future release.

⁷ G.8032 support is planned for future release.

⁸ G.8032 support is planned for future release.

⁹ Point-to-Multipoint service support is planned for future release.

- Point-to-Point (E-Line)
- Point-to-Multipoint (E-Tree)
- Smart Pipe
- Management
- 128K MAC learning table, with separate learning per service (including limiters)
- Flexible transport and encapsulation via 802.1q and 802.1ad (Q-in-Q), with tag manipulation possible at ingress and egress
- High precision, flexible frame synchronization solution combining SyncE and 1588v2
- Hierarchical QoS with 2K service level queues, deep buffering, hierarchical scheduling via WFQ and Strict priority, and shaping at each level
- 1K hierarchical two-rate three-Color policers
 - Port based – Unicast, Multicast, Broadcast, Ethertype
 - Service-based
 - CoS-based
- Up to four link aggregation groups (LAG)
 - Hashing based on L2, L3, MPLS, and L4
- Enhanced <50msec network level resiliency (G.8032) for ring/mesh support

5.2.3. Supported Standards

NetStream Primo is fully MEF-9 and MEF-14 certified for all Carrier Ethernet services. For a full list of standards and certifications supported by NetStream Primo, refer to the following sections:


- Supported Ethernet Standards
- MEF Certifications for Ethernet Services

5.2.4. Ethernet Service Model

NetStream Primo's service-oriented Ethernet paradigm is based on Carrier-Ethernet Transport (CET), and provides a highly flexible and granular switching fabric for Ethernet services.

NetStream Primo's virtual switching/forwarding engine is based on a clear distinction between user-facing service interfaces and intra-network service interfaces. User-facing interfaces (UNIs) are configured as Service Access Points (SAPs), while intra-network interfaces (E-NNIs or NNIs) are configured as Service Network Points (SNPs).

NetStream Primo Services Model



The NetStream Primo services core provides for fully flexible C-VLAN and S-VLAN encapsulation, with a full range of classification, preservation, and translation options available. Service security and isolation is provided without limiting the C-VLAN reuse capabilities of different customers.

Users can define up to 64 services on a single NetStream Primo. Each service constitutes a virtual bridge that defines the connectivity and behavior among the network element interfaces for the specific virtual bridge. In addition to user-defined services, NetStream Primo contains a pre-defined management service (Service ID 257). If needed, users can activate the management service and use it for in-band management.

To define a service, the user must configure virtual connections among the interfaces that belong to the service. This is done by configuring service points (SPs) on these interfaces.


A service can hold up to 32 service points. A service point is a logical entity attached to a physical or logical interface. Service points define the movement of frames through the service. Each service point includes both ingress and egress attributes.



Management services can hold up to 30 SPs.

The following figure illustrates the NetStream Primo services model, with traffic entering and leaving the network element. NetStream Primo’s switching fabric is designed to provide a high degree of flexibility in the definition of services and the treatment of data flows as they pass through the switching fabric.

NetStream Primo Services Core




5.2.4.1. Frame Classification to Service Points and Services

Each arriving frame is classified to a specific service point, based on a key that consists of:

- The Interface ID of the interface through which the frame entered the NetStream Primo.
- The frame’s C-VLAN and/or S-VLAN tags.

If the classification mechanism finds a match between the key of the arriving frame and a specific service point, the frame is associated to the specific service to which the service point belongs. That service point is called the ingress service point for the frame, and the other service points in the service are optional egress service points for the frame. The frame is then forwarded from the ingress service point to an egress service point by means of flooding or dynamic address learning in the specific service. Services include a MAC entry table of up to 131,072 entries, with a global aging timer and a maximum learning limiter that are configurable per-service.

NetStream Primo Services Flow



5.2.4.2. Service Types

NetStream Primo supports the following service types:

- Point-to-Point Service (P2P)
- MultiPoint Service (MP)
- Management Service
- Point-to-Multipoint Service (E-Tree)



Note


Support for E-Tree services is planned for future release.

5.2.4.3. Point to Point Service (P2P)

Point-to-point services are used to provide connectivity between two interfaces of the network element. When traffic ingresses via one side of the service, it is immediately directed to the other side according to ingress and egress tunneling rules. This type of service contains exactly two service points and does not require MAC address-based learning or forwarding. Since the route is clear, the traffic is tunneled from one side of the service to the other and vice versa.

The following figure illustrates a P2P service.

Point-to-Point Service




P2P services provide the building blocks for network services such as E-Line EVC (EPL and EVPL EVCs) and port-based services (Smart Pipe).

5.2.4.4. Multipoint Service (MP)

Multipoint services are used to provide connectivity between two or more service points. When traffic ingresses via one service point, it is directed to one of the service points in the service, other than the ingress service point, according to ingress and egress tunneling rules, and based on the learning and forwarding mechanism. If the destination MAC address is not known by the learning and forwarding mechanism, the arriving frame is flooded to all the other service points in the service except the ingress service point.

The following figure illustrates a Multipoint service.

Multipoint Service



Multipoint services provide the building blocks for network services such as E-LAN EVCs (EP-LAN and EVP-LAN EVCs), and for E-Line EVCs (EPL and EVPL EVCs) in which

only two service points are active. In such a case, the user can disable MAC address learning in the service points to conserve system resources.

5.2.4.5. Learning and Forwarding Mechanism

NetStream Primo can learn up to 131,072 Ethernet source MAC addresses. NetStream Primo performs learning per service in order to enable the use of 64 virtual bridges in the network element. If necessary due to security issues or resource limitations, users can limit the size of the MAC forwarding table. The maximum size of the MAC forwarding table is configurable per service in granularity of 16 entries.

When a frame arrives via a specific service point, the learning mechanism checks the MAC forwarding table for the service to which the service point belongs to determine whether that MAC address is known to the service. If the MAC address is not found, the learning mechanism adds it to the table under the specific service.

In parallel with the learning process, the forwarding mechanism searches the service's MAC forwarding table for the frame's destination MAC address. If a match is found, the frame is forwarded to the service point associated with the MAC address. If not, the frame is flooded to all service points in the service.

The following table illustrates the operation of the learning and forwarding mechanism.

Ethernet Services Learning and Forwarding

MAC Forwarding Table			
Input Key for learning / forwarding (search) operation		Result	Entry type
Service ID	MAC address	Service Point	
13	00:34:67:3a:aa:10	15	dynamic
13	00:0a:25:33:22:12	31	dynamic
28	00:0a:25:11:12:55	31	static
55	00:0a:25:33:22:12	15	dynamic
55	00:c3:20:57:14:89	31	dynamic
55	00:0a:25:11:12:55	31	dynamic

In addition to the dynamic learning mechanism, users can add static MAC addresses for static routing in each service. These user entries are not considered when determining the maximum size of the MAC forwarding table.

Users can manually clear all the dynamic entries from the MAC forwarding table. Users can also delete static entries per service.

The system also provides an automatic flush process. An entry is erased from the table as a result of:

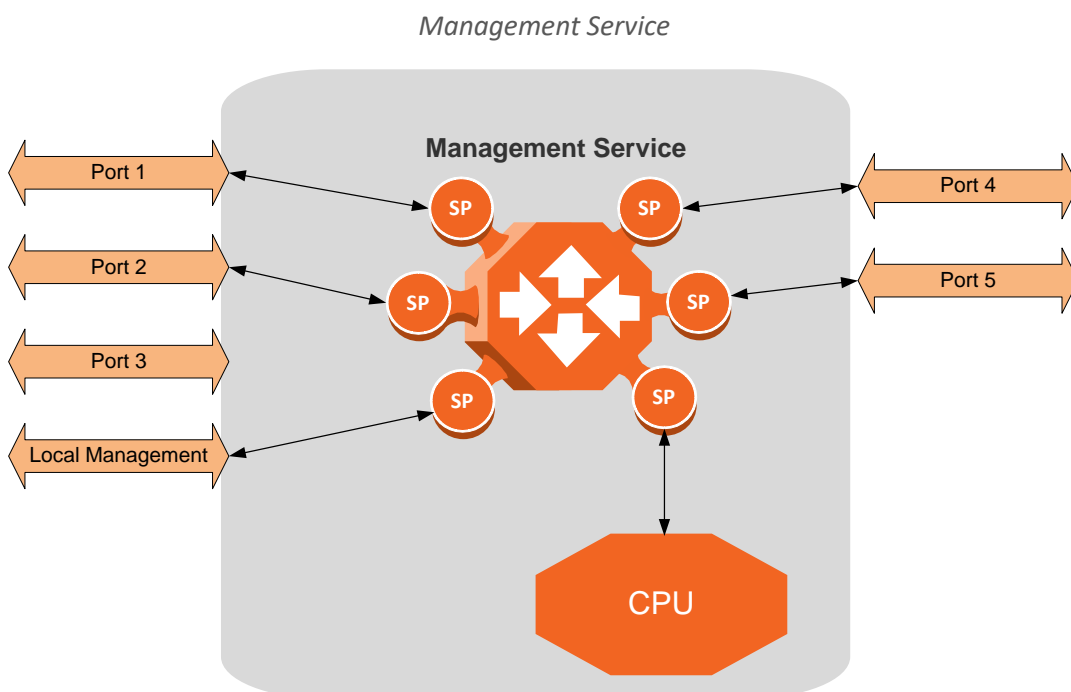
- The global aging time expires for the entry.
- Loss of carrier occurs on the interface with which the entry is associated.
- Resiliency protocols, such as MSTP or G.8032.

5.2.4.6. Management Service (MNG)

The management service connects the local management port, the network element host CPU, and the traffic ports into a single service. The management service is pre-defined in the system, with Service ID 257. The pre-defined management service has a single service point that connects the service to the network element host CPU and the management port. To configure in-band management over multiple network elements, the user must connect the management service to the network by adding a service point on an interface that provides the required network connectivity.

Users can modify the attributes of the management service, but cannot delete it. The CPU service point is read-only and cannot be modified. The local management port is also connected to the service, but its service point is not visible to users. The management port is enabled by default and cannot be disabled.

The following figure illustrates a management service.



Management services can provide building blocks for network services such as E-LAN EVCs (EP-LAN and EVP-LAN), as well as E-Line EVCs (EPL and EVPL EVCs) in which only two service points are active.

5.2.4.7. Service Attributes

NetStream Primo services have the following attributes:

- **Service ID** – A unique ID that identifies the service. The user must select the Service ID upon creating the service. The Service ID cannot be edited after the service has been created. Service ID 257 is reserved for the pre-defined Management service.
- **Service Type** – Determines the specific functionality that will be provided for Ethernet traffic using the service. For example, a Point-to-Point service

provides traffic forwarding between two service points, with no need to learn a service topology based on source and destination MAC addresses. A Multipoint service enables operators to create an E-LAN service that includes several service points.

- **Service Admin Mode** – Defines whether or not the service is functional, i.e., able to receive and transmit traffic. When the Service Admin Mode is set to Operational, the service is fully functional. When the Service Admin Mode is set to Reserved, the service occupies system resources but is unable to transmit and receive data.
- **EVC-ID** – The Ethernet Virtual Connection ID (end-to-end). This parameter does not affect the network element's behavior, but is used by the NMS for topology management.
- **EVC Description** – The Ethernet Virtual Connection description. This parameter does not affect the network element's behavior, but is used by the NMS for topology management.
- **Maximum Dynamic MAC Address Learning per Service** – Defines the maximum number of dynamic Ethernet MAC address that the service can learn. This parameter is configured with a granularity of 16, and only applies to dynamic, not static, MAC addresses.
- **Static MAC Address Configuration** – Users can add static entries to the MAC forwarding table. The global aging time does not apply to static entries, and they are not counted with respect to the Maximum Dynamic MAC Address Learning. It is the responsibility of the user not to use all the 131,072 entries in the table if the user also wants to utilize dynamic MAC address learning.
- **CoS Mode** – Defines whether the service inherits ingress classification decisions made at previous stages or overwrites previous decisions and uses the default CoS defined for the service. For more details on NetStream Primo's hierarchical classification mechanism, refer to *Classification* on page 88.
- **Default CoS** – The default CoS value at the service level. If the CoS Mode is set to overwrite previous classification decisions, this is the CoS value used for frames entering the service.
- **xSTP Instance (0-46, 4095)** – The spanning tree instance ID to which the service belongs. The service can be a traffic engineering service (instance ID 4095) or can be managed by the xSTP engines of the network element.

5.2.4.8. Service Points


Service points are logical entities attached to the interfaces that make up the service. Service points define the movement of frames through the service. Without service points, a service is simply a virtual bridge with no ingress or egress interfaces.

NetStream Primo supports several types of service points:

- **Management (MNG) Service Point** – Only used for management services. The following figure shows a management service used for in-band management among four network elements in a ring. In this example, each

service contains three MNG service points, two for East-West management connectivity in the ring, and one serving as the network gateway.


Management Service and its Service Points



- **Service Access Point (SAP) Service Point** – An SAP is equivalent to a UNI in MEF terminology and defines the connection of the user network with its access points. SAPs are used for Point-to-Point and Multipoint traffic services.
- **Service Network Point (SNP) Service Point** – An SNP is equivalent to an NNI or E-NNI in MEF terminology and defines the connection between the network elements in the user network. SNPs are used for Point-to-Point and Multipoint traffic services.

The following figure shows four network elements in ring. An MP Service with three service points provides the connectivity over the network. The SNPs provide the connectivity among the network elements in the user network while the SAPs provide the access points for the network.

SAPs and SNPs



- **Pipe Service Point** – Used to create traffic connectivity between two points in a port-based manner (Smart Pipe). In other words, all the traffic from one port passes to the other port. Pipe service points are used in Point-to-Point services.

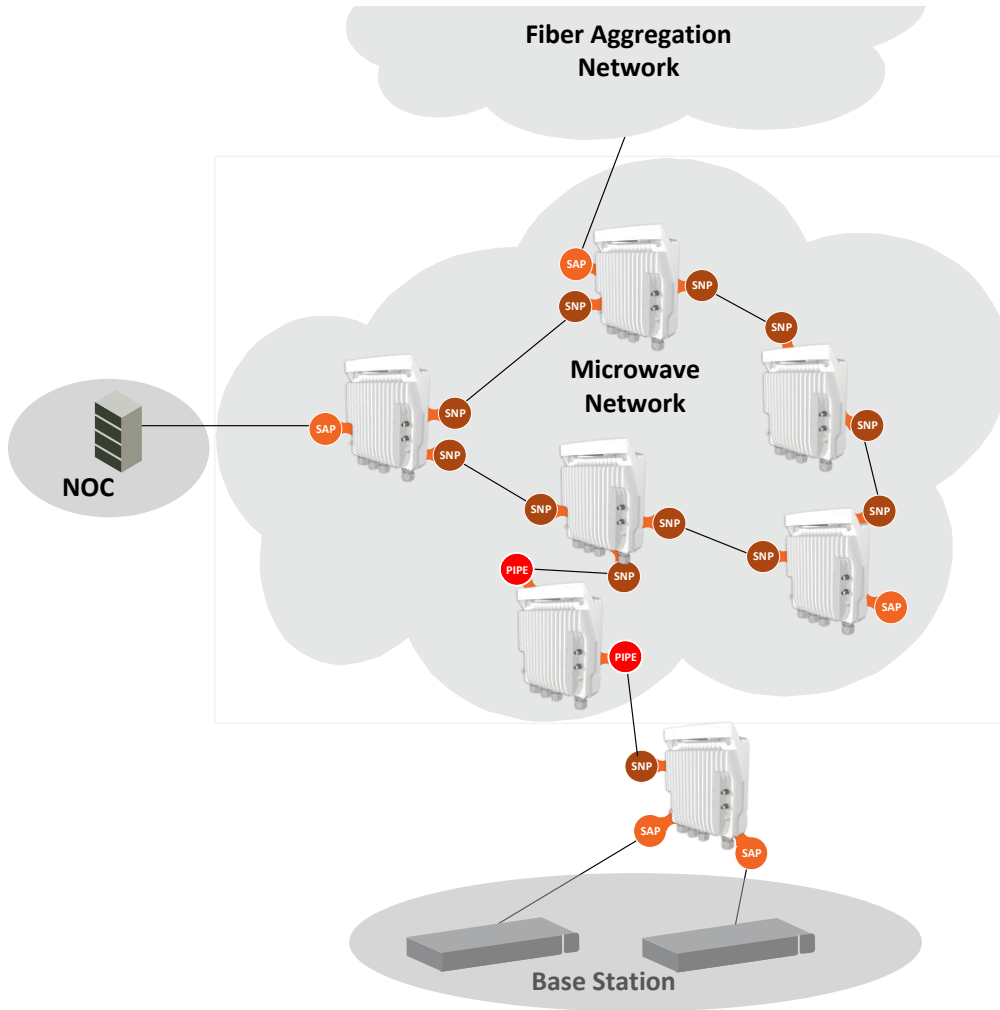
The following figure shows a Point-to-Point service with Pipe service points that create a Smart Pipe between Port 1 of the network element on the left and Port 2 of the network element on the right.

Pipe Service Points



The following figure shows the usage of SAP, SNP and Pipe service points in a microwave network. The SNPs are used for interconnection between the network elements while the SAPs provide the access points for the network. A Smart Pipe is also used, to provide connectivity between elements that require port-based connectivity.

SAP, SNP and Pipe Service Points in a Microwave Network



The following table summarizes the service point types available per service type.

Service Point Types per Service Type

		Service point type			
		MNG	SAP	SNP	Pipe
Service Type	Management	Yes	No	No	No
	Point-to-Point	No	Yes	Yes	Yes
	Multipoint	No	Yes	Yes	No

5.2.4.9. Service Point Classification

As explained above, service points connect the service to the network element interfaces. It is crucial that the network element have a means to classify incoming frames to the proper service point. This classification process is implemented by means of a parsing encapsulation rule for the interface associated with the service point. This rule is called the Attached Interface Type, and is based on a three part key consisting of:

- The Interface ID of the interface through which the frame entered.
- The frame's C-VLAN and/or S-VLAN tags.

The Attached Interface Type provides a definitive mapping of each arriving frame to a specific service point in a specific service. Since more than one service point may be associated with a single interface, frames are assigned to the earliest defined service point in case of conflict.

5.2.4.10. SAP Classification

SAPs can be used with the following Attached Interface Types:

- **All to one** – All C-VLANs and untagged frames that enter the interface are classified to the same service point.
- **Dot1q** – A single C-VLAN is classified to the service point.
- **QinQ** – A single S-VLAN and C-VLAN combination is classified to the service point.
- **Bundle C-Tag** – A set of multiple C-VLANs are classified to the service point.
- **Bundle S-Tag** – A single S-VLAN and a set of multiple C-VLANs are classified to the service point.

5.2.4.11. SNP classification

SNPs can be used with the following Attached Interface Types:

- **Dot1q** – A single C-VLAN is classified to the service point.
- **S-Tag** – A single S-VLAN is classified to the service point.

5.2.4.12. PIPE classification

Pipe service points can be used with the following Attached Interface Types:

- **Dot1q** – All C-VLANs and untagged frames that enter the interface are classified to the same service point.
- **S-Tag** – All S-VLANs and untagged frames that enter the interface are classified to the same service point.

5.2.4.13. MNG classification

Management service points can be used with the following Attached Interface Types:

- **Dot1q** – A single C-VLAN is classified to the service point.
- **S-Tag** – A single S-VLAN is classified to the service point.
- **QinQ** – A single S-VLAN and C-VLAN combination is classified into the service point.

The following table shows which service point types can co-exist on the same interface.

Service Point Types that can Co-Exist on the Same Interface

	MNG SP	SAP SP	SNP SP	Pipe SP
MNG SP	Only one MNG SP is allowed per interface.	Yes	Yes	Yes
SAP SP	Yes	Yes	No	No
SNP SP	Yes	No	Yes	No
PIPE SP	Yes	No	No	Only one Pipe SP is allowed per interface.

The following table shows in more detail which service point – Attached Interface Type combinations can co-exist on the same interface.

Service Point Type-Attached Interface Type Combinations that can Co-Exist on the Same Interface

	SP Type	SAP					SNP		Pipe		MNG		
SP Type	Attached Interface Type	802.1q	Bundle C-Tag	Bundle S-Tag	All to One	QinQ	802.1q	S-Tag	802.1q	S-Tag	802.1q	QinQ	S-Tag
SAP	802.1q	Yes	Yes	No	No	No	No	No	Only for P2P Service	No	Yes	No	No
	Bundle C-Tag	Yes	Yes	No	No	No	No	No	Only for P2P Service	No	Yes	No	No
	Bundle S-Tag	No	No	Yes	No	Yes	No	No	No	No	No	Yes	No
	All to One	No	No	No	Only 1 All to One SP Per Interface	No	No	No	No	No	No	No	No
	QinQ	No	No	Yes	No	Yes	No	No	No	No	No	Yes	No
SNP	802.1q	No	No	No	No	No	Yes	No	Only for P2P Service	No	Yes	No	No
	S-Tag	No	No	No	No	No	No	Yes	No	Only for P2P Service	No	No	Yes
Pipe	802.1q	Only for P2P Service	Only for P2P Service	No	No	No	Only for P2P Service	No	Only one Pipe SP Per Interface	No	Yes	No	No
	S-Tag	No	No	No	No	No	No	Only for P2P Service	No	Only one Pipe SP Per Interface	No	No	Yes
MNG	802.1q	Yes	Yes	No	No	No	Yes	No	Yes	No	No	No	No
	QinQ	No	No	Yes	No	Yes	No	No	No	No	No	No	No
	S-Tag	No	No	No	No	No	No	Yes	No	Yes	No	No	No


5.2.4.14. Service Point Attributes

As described above, traffic ingresses and egresses the service via service points. The service point attributes are divided into two types:

- **Ingress Attributes** – Define how frames are handled upon ingress, e.g., policing and MAC address learning.
- **Egress Attributes** – Define how frames are handled upon egress, e.g., preservation of the ingress CoS value upon egress, VLAN swapping.

The following figure shows the ingress and egress path relationship on a point-to-point service path. When traffic arrives via port 1, the system handles it using service point 1 ingress attributes then forwards it to service point 2 and handles it using the SP2 egress attributes:

Service Path Relationship on Point-to-Point Service Path



Service points have the following attributes:

5.2.4.15. General Service Point Attributes

- **Service Point ID** – Users can define up to 32 service points per service, except for management services which are limited to 30 service points in addition to the pre-defined management system service point.
- **Service Point Name** – A descriptive name, which can be up to 20 characters.
- **Service Point Type** – The type of service point, as described above.
- **S-VLAN Encapsulation** – The S-VLAN ID associated with the service point.
- **C-VLAN Encapsulation** – The C-VLAN ID associated with the service point.
- **Attached C VLAN** – For service points with an Attached Interface Type of Bundle C-Tag, this attribute is used to create a list of C-VLANs associated with the service point.
- **Attached S-VLAN** – For service points with an Attached Interface Type of Bundle S-Tag, this attribute is used to create a list of S-VLANs associated with the service point.

5.2.4.16. Ingress Service Point Attributes

The ingress attributes are attributes that operate upon frames when they ingress via the service point.

- **Attached Interface Type** – The interface type to which the service point is attached, as described above. Permitted values depend on the service point type.
- **Learning Administration** – Enables or disables MAC address learning for traffic that ingresses via the service point. This option enables users to enable or disable MAC address learning for specific service points.
- **Allow Broadcast** – Determines whether to allow frames to ingress the service via the service point when the frame has a broadcast destination MAC address.
- **Allow Flooding** – Determines whether incoming frames with unknown MAC addresses are forwarded to other service points via flooding.
- **CoS Mode** – Determines whether the service point preserves the CoS decision made at the interface level, overwrites the CoS with the default CoS for the service point.
- **Default CoS** – The service point CoS. If the CoS Mode is set to overwrite the CoS decision made at the interface level, this is the CoS value assigned to frames that ingress the service point.
- **Token Bucket Profile** – This attribute can be used to attach a rate meter profile to the service point. Permitted values are 1– 250.
- **CoS Token Bucket Profile** – This attribute can be used to attach a rate meter profile to the service point at the CoS level. Users can define a rate meter for each of the eight CoS values of the service point. Permitted values are 1– 250 for CoS 0–7.
- **CoS Token Bucket Admin** – Enables or disables the rate meter at the service point CoS level.

5.2.4.17. Egress Service Point Attributes

The egress attributes are attributes that operate upon frames egressing via the service point.

- **C-VLAN ID Egress Preservation** – If enabled, C-VLAN frames egressing the service point retain the same C-VLAN ID they had when they entered the service.
- **C-VLAN CoS Egress Preservation** – If enabled, the C-VLAN CoS value of frames egressing the service point is the same as the value when the frame entered the service.
- **S-VLAN CoS egress Preservation** – If enabled, the S-VLAN CoS value of frames egressing the service point is the same as the value when the frame entered the service.
- **Marking** – Marking refers to the ability to overwrite the outgoing priority bits and Color of the outer VLAN of the egress frame, either the C-VLAN or

the S-VLAN. If marking is enabled, the service point overwrites the outgoing priority bits and Color of the outer VLAN of the egress frame. Marking mode is only relevant if either the outer frame is S-VLAN and S-VLAN CoS preservation is disabled, or the outer frame is C-VLAN and C-VLAN CoS preservation is disabled. When marking is enabled and active, marking is performed according to global mapping tables that map the 802.1p-UP bits and the DEI or CFI bit to a defined CoS and Color value.

- **Service Bundle ID** – This attribute can be used to assign one of the available service bundles from the H-QoS hierarchy queues to the service point. This enables users to personalize the QoS egress path. For details, refer to *Standard QoS and Hierarchical QoS (H-QoS)* on page 101.

5.2.5. Ethernet Interfaces

The NetStream Primo switching fabric distinguishes between physical interfaces and logical interfaces. Physical and logical interfaces serve different purposes in the switching fabric.

The concept of a physical interface refers to the physical characteristics of the interface, such as speed, duplex, auto-negotiation, master/slave, and standard RMON statistics.

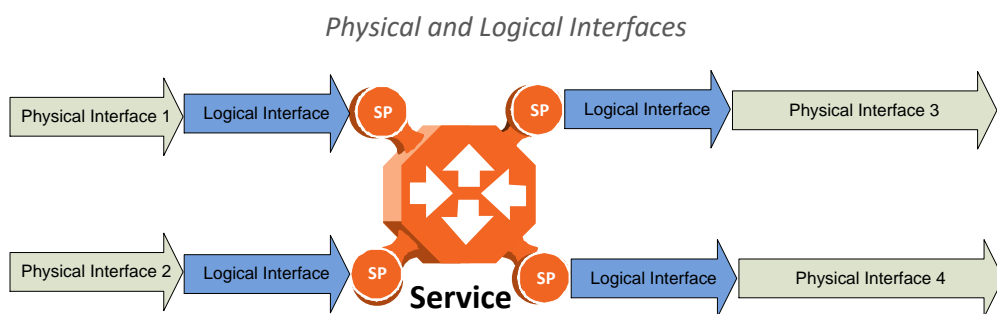
A logical interface can consist of a single physical interface or a group of physical interfaces that share the same function. Examples of the latter are protection groups and link aggregation groups. Switching and QoS functionality are implemented on the logical interface level.

It is important to understand that the NetStream Primo switching fabric regards all traffic interfaces as regular physical interfaces, distinguished only by the media type the interface uses, e.g., RJ-45, SFP, or Radio.

From the user’s point of view, the creation of the logical interface is simultaneous with the creation of the physical interface. For example, when the user enables a radio interface, both the physical and the logical radio interface come into being at the same time.

Once the interface is created, the user configures both the physical and the logical interface. In other words, the user configures the same interface on two levels, the physical level and the logical level.

The following figure shows physical and logical interfaces in a one-to-one relationship in which each physical interface is connected to a single logical interface, without grouping.





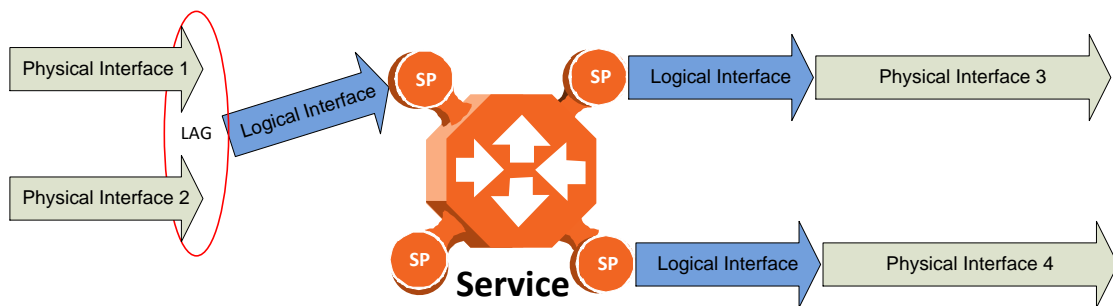
Note

For simplicity only, this figure represents a uni-directional rather than a bi-directional traffic flow.

The next figure illustrates the grouping of two or more physical interfaces into a logical interface, a link aggregation group (LAG) in this example. The two physical interfaces on the ingress side send traffic into a single logical interface. The user configures each physical interface separately, and configures the logical interface as a single logical entity. For example, the user might configure each physical interface to 100 mbps, full duplex, with auto-negotiation off. On the group level, the user might limit the group to a rate of 200 mbps by configuring the rate meter on the logical interface level.

When physical interfaces are grouped into a logical interface, NetStream Primo also shows standard RMON statistics for the logical interface, i.e., for the group. This information enables users to determine the cumulative statistics for the group, rather than having to examine the statistics for each interface individually.

Grouped Interfaces as a Single Logical Interface on Ingress Side




Note

For simplicity only, this figure represents a uni-directional rather than a bi-directional traffic flow.

The following figure shows the logical interface at the egress side. In this case, the user can configure the egress traffic characteristics, such as scheduling, for the group as a whole as part of the logical interface attributes.

Grouped Interfaces as a Single Logical Interface on Egress Side



**Note**

For simplicity only, this figure represents a uni-directional rather than a bi-directional traffic flow.

5.2.5.1. Physical Interfaces

The physical interfaces refer to the real traffic ports (layer 1) that are connected to the network. The Media Type attribute defines the Layer 1 physical traffic interface type, which can be:

- Radio interface
- RJ-45 or SFP Ethernet interface.

5.2.5.2. Physical Interface Attributes

The following physical interface parameters can be configured by users:

- **Admin** – Enables or disables the physical interface. This attribute is set via the Interface Manager section of the Web EMS.
- **Auto Negotiation** – Enables or disables auto-negotiation on the physical interface. Auto Negotiation is always off for radio and SFP interfaces.
- **Speed and Duplex** – The physical interface speed and duplex mode. Permitted values are:
 - **Ethernet RJ-45 interfaces:** 10Mbps HD, 10Mbps FD, 100Mbps HD, 100Mbps FD, and 1000Mbps FD.
 - **Ethernet SFP interfaces:** Only 1000FD is supported
 - **Radio interfaces:** The parameter is read-only and set by the system to 1000FD.
- **Flow Control** – The physical port flow control capability. Permitted values are: Symmetrical Pause and/or Asymmetrical Pause. This parameter is only relevant in Full Duplex mode.¹⁰
- **IFG** – The physical port Inter-frame gap. Although users can modify the IFG field length, it is strongly recommended not to modify the default value of 12 bytes without a thorough understanding of how the modification will impact traffic. Permitted values are 6 to 15 bytes.
- **Preamble** – The physical port preamble value. Although users can modify the preamble field length, it is strongly recommended not to modify the default values of 8 bytes without a thorough understanding of how the modification will impact traffic. Permitted values are 6 to 15 bytes.
- **Interface description** – A text description of the interface, up to 40 characters.

The following read-only physical interface status parameters can be viewed by users:

¹⁰ This functionality is planned for future release.

- **Operational State** – The operational state of the physical interface (Up or Down).
- **Actual Speed and Duplex** – The actual speed and duplex value for the Ethernet link as agreed by the two sides of the link after the auto negotiation process.
- **Actual Flow Control State** – The actual flow control state values for the Ethernet link as agreed by the two sides after the auto negotiation process.
- **Actual Physical Mode** (only relevant for RJ-45 interfaces) – The actual physical mode (master or slave) for the Ethernet link, as agreed by the two sides after the auto negotiation process.

5.2.5.3. Ethernet Statistics

The NetStream Primo platform stores and displays statistics in accordance with RMON and RMON2 standards.

Users can display various peak TX and RX rates (in seconds) and average TX and RX rates (in seconds), both in bytes and in packets, for each measured time interval. Users can also display the number of seconds in the interval during which TX and RX rates exceeded the configured threshold.

The following transmit statistic counters are available:

- Transmitted bytes (not including preamble) in good or bad frames. Low 32 bits.
- Transmitted bytes (not including preamble) in good or bad frames. High 32 bits.
- Transmitted frames (good or bad)
- Multicast frames (good only)
- Broadcast frames (good only)
- Control frames transmitted
- Pause control frame transmitted
- FCS error frames
- Frame length error
- Oversized frames – frames with length > 1518 bytes (1522 bytes for VLAN-tagged frames) without errors
- Undersized frames (good only)
- Fragments frames (undersized bad)
- Jabber frames – frames with length > 1518 bytes (1522 for VLAN-tagged frames) with errors
- Frames with length 64 bytes, good or bad
- Frames with length 65-127 bytes, good or bad
- Frames with length 128-255 bytes, good or bad

- Frames with length 256-511 bytes, good or bad
- Frames with length 512-1023 bytes, good or bad.
- Frames with length 1024-1518 bytes, good or bad
- Frames with length 1519-1522 bytes, good or bad

The following receive statistic counters are available:

- Received bytes (not including preamble) in good or bad frames. Low 32 bits.
- Received bytes (not including preamble) in good or bad frames. High 32 bits.
- Received frames (good or bad)
- Multicast frames (good only)
- Broadcast frames (good only)
- Control frames received
- Pause control frame received
- FCS error frames
- Frame length error
- Code error
- Counts oversized frames – frames with length > 1518 bytes (1522 bytes for VLAN-tagged frames) without errors *and* frames with length > MAX_LEN without errors
- Undersized frames (good only)
- Fragments frames (undersized bad)
- Counts jabber frames – frames with length > 1518 bytes (1522 for VLAN-tagged frames) with errors
- Frames with length 64 bytes, good or bad
- Frames with length 65-127 bytes, good or bad
- Frames with length 128-255 bytes, good or bad
- Frames with length 256-511 bytes, good or bad
- Frames with length 512-1023 bytes, good or bad
- Frames with length 1024-1518 bytes, good or bad
- VLAN-tagged frames with length 1519-1522 bytes, good or bad
- Frames with length > MAX_LEN without errors
- Frames with length > MAX_LEN with errors

5.2.5.4. Logical Interfaces


A logical interface consists of one or more physical interfaces that share the same traffic ingress and egress characteristics. From the user's point of view, it is more convenient to define interface behavior for the group as a whole than for each

individual physical interface that makes up the group. Therefore, classification, QoS, and resiliency attributes are configured and implemented on the logical interface level, in contrast to attributes such as interface speed and duplex mode, which are configured on the physical interface level.

It is important to understand that the user relates to logical interfaces in the same way in both a one-to-one scenario in which a single physical interface corresponds to a single logical interface, and a grouping scenario such as a link aggregation group or a radio protection group, in which several physical interfaces correspond to a single logical interface.

The following figure illustrates the relationship of a LAG group to the switching fabric. From the point of view of the user configuring the logical interface attributes, the fact that there are two Ethernet interfaces is not relevant. The user configures and manages the logical interface just as if it represented a single Ethernet interface.

Relationship of Logical Interfaces to the Switching Fabric



5.2.5.5. Logical Interface Attributes

The following logical interface attributes can be configured by users:

5.2.5.6. General Attributes

- **Traffic Flow Administration** – Enables traffic via the logical interface. This attribute is useful when the user groups several physical interfaces into a single logical interface. The user can enable or disable traffic to the group using this parameter.

5.2.5.7. Ingress Path Classification at Logical Interface Level

These attributes represent part of the hierarchical classification mechanism, in which the logical interface is the lowest point in the hierarchy.

- **VLAN ID** – Users can specify a specific CoS and Color for a specific VLAN ID. In the case of double-tagged frames, the match must be with the frame's outer VLAN. Permitted values are CoS 0 to 7 and Color Green or Yellow per VLAN ID. This is the highest classification priority on the logical interface level, and overwrites any other classification criteria at the logical interface level.

- **802.1p Trust Mode** – When this attribute is set to Trust mode and the arriving packet is 802.1Q or 802.1AD, the interface performs QoS and Color classification according to user-configurable tables for 802.1q UP bit (C-VLAN frames) or 802.1AD UP bit (S-VLAN frames) to CoS and Color classification.
- **IP DSCP Trust Mode** – When this attribute is set to Trust mode and the arriving packet has IP priority bits, the interface performs QoS and Color classification according to a user-configurable DSCP bit to CoS and Color classification table. 802.1p classification has priority over DSCP Trust Mode, so that if a match is found on the 802.1p level, DSCP bits are not considered.
- **MPLS Trust Mode** – When this attribute is set to Trust mode and the arriving packet has MPLS EXP priority bits, the interface performs QoS and Color classification according to a user-configurable MPLS EXP bit to CoS and Color classification table. Both 802.1p and DSCP classification have priority over MPLS Trust Mode, so that if a match is found on either the 802.1p or DSCP levels, MPLS bits are not considered.
- **Default CoS** – The default CoS value for frames passing through the interface. This value can be overwritten on the service point and service level. The Color is assumed to be Green.

For more information about classification at the logical interface level, refer to *Logical Interface-Level Classification* on page 89.

5.2.5.8. Ingress Path Rate Meters at Logical Interface Level

- **Unicast Traffic Rate Meter Admin** – Enables or disables the unicast rate meter (policer) on the logical interface.
- **Unicast Traffic Rate Meter Profile** – Associates the rate meter (policer) with a specific rate meter (policer) profile.
- **Multicast Traffic Rate Meter Admin** – Enables or disables the multicast rate meter (policer) on the logical interface.
- **Multicast Traffic Rate Meter Profile** – Associates the rate meter (policer) with a specific rate meter (policer) profile.
- **Broadcast Traffic Rate Meter Admin** – Enables or disables the broadcast rate meter (policer) on the logical interface.
- **Broadcast Traffic Rate Meter Profile** – Associates the rate meter (policer) with a specific rate meter (policer) profile.
- **Ethertype 1 Rate Meter Admin** – Enables or disables the Ethertype 1 rate meter (policer) on the logical interface.
- **Ethertype 1 Rate Meter Profile** – Associates the rate meter (policer) with a specific rate meter (policer) profile.
- **Ethertype 1 Value** – The Ethertype value to which the user wants to apply this rate meter (policer). The field length is 4 nibbles (for example, 0x0806 - ARP).

- **Ethertype 2 Rate Meter Admin** – Enables or disables the Ethertype 2 rate meter (policer) on the logical interface.
- **Ethertype 2 Rate Meter Profile** – Associates the rate meter (policer) with a specific rate meter (policer) profile.
- **Ethertype 2 Value** – The Ethertype value to which the user wants to apply the rate meter (policer). The field length is 4 nibbles (for example, 0x0806 - ARP).
- **Ethertype 3 Rate Meter Admin** – Enables or disables the Ethertype 3 rate meter (policer) on the logical interface.
- **Ethertype 3 Rate Meter Profile** – Associates the rate meter (policer) with a specific rate meter (policer) profile.
- **Ethertype 3 Value** – The Ethertype value to which the user wants to apply the rate meter (policer). The field length is 4 nibbles (for example, 0x0806 - ARP).
- **Inline Compensation** – The logical interface's ingress compensation value. The rate meter (policer) attached to the logical interface uses this value to compensate for Layer 1 non-effective traffic bytes.

5.2.5.9. Egress Path Shapers at Logical Interface Level

- **Logical Port Shaper Profile** – Users can assign a single leaky bucket shaper to each interface. The shaper on the interface level stops traffic from the interface if a specific user-defined peak information rate (PIR) has been exceeded.¹¹
- **Outline Compensation** – The logical interface's egress compensation value. Any shaper attached to this interface, in any layer, uses this value to compensate for Layer 1 non-effective traffic bytes. Permitted values are even numbers between 0 and 26 bytes. The default value is 0 bytes.

5.2.5.10. Egress Path Scheduler at Logical Interface Level

- **Logical Interface Priority Profile** – This attribute is used to attach an egress scheduling priority profile to the logical interface.
- **Logical Port WFQ Profile** – This attribute is used to attach an egress scheduling WFQ profile to the logical interface. The WFQ profile provides a means of allocating traffic among queues with the same priority.

The following read-only logical interface status parameters can be viewed by users:

- **Traffic Flow Operational Status** – Indicates whether or not the logical interface is currently functional.

¹¹ This attribute is reserved for future use. The current release supports traffic shaping per queue and per service bundle, which provides the equivalent of shaping per logical interface.

5.2.5.11. Logical Interface Statistics

5.2.5.12. RMON Statistics at Logical Interface Level

As discussed in *Ethernet Statistics* on page 80, if the logical interface represents a group, such as a LAG or a 1+1 HSB pair, the NetStream Primo platform stores and displays RMON and RMON2 statistics for the logical interface.

5.2.5.13. Rate Meter (Policer) Statistics at Logical Interface Level

For the rate meter (policer) at the logical interface level, users can view the following statistics counters:

- Green Frames
- Green Bytes
- Yellow Frames
- Yellow Bytes
- Red Frames
- Red Bytes

**Note**

Rate meter (policer) counters are 64 bits wide.

5.2.5.14. Link Aggregation Groups (LAG)

Link aggregation (LAG) enables users to group several physical interfaces into a single logical interface bound to a single MAC address. This logical interface is known as a LAG group. Traffic sent to the interfaces in a LAG group is distributed by means of a load balancing function. NetStream Primo uses a distribution function of up to Layer 4 in order to generate the most efficient distribution among the LAG physical ports, taking into account:

- MAC DA and MAC SA
- IP DA and IP SA
- C-VLAN
- S-VLAN
- Layer 3 Protocol Field
- UDP/TCP Source Port and Destination Port
- MPLS Label

LAG can be used to provide redundancy for Ethernet interfaces, both on the same NetStream Primo unit (line protection) and on separate units (line protection and equipment protection). LAGs can also be used to provide redundancy for radio links. LAG can also be used to aggregate several interfaces in order to create a wider (aggregate) Ethernet link. For example, LAG can be used to create a 3 Gbps channel by grouping the three Ethernet interfaces to a single LAG.

Up to four LAG groups can be created.

LAG groups can include interfaces with the following constraints:

- Only physical interfaces (including radio interfaces), not logical interfaces, can belong to a LAG group.
- Interfaces can only be added to the LAG group if no services or service points are attached to the interface.
- Any classification rules defined for the interface are overridden by the classification rules defined for the LAG group.
- When removing an interface from a LAG group, the removed interface is assigned the default interface values.

NetStream Primo enables users to select the LAG members without limitations, such as interface speed and interface type. Proper configuration of a LAG group is the responsibility of the user.

5.2.6. Quality of Service (QoS)

Related topics:

- Ethernet Service Model
- In-Band Management

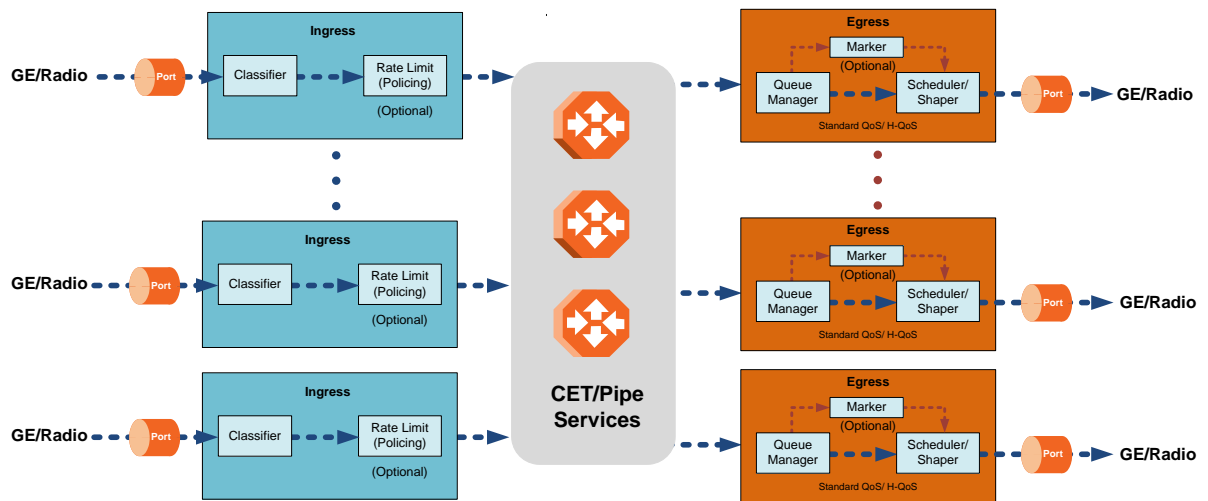
Quality of Service (QoS) deals with the way frames are handled within the switching fabric. QoS is required in order to deal with many different network scenarios, such as traffic congestion, packet availability, and delay restrictions.

NetStream Primo's personalized QoS enables operators to handle a wide and diverse range of scenarios. NetStream Primo's smart QoS mechanism operates from the frame's ingress into the switching fabric until the moment the frame egresses via the destination port.

QoS capability is very important due to the diverse topologies that exist in today's network scenarios. These can include, for example, streams from two different ports that egress via single port, or a port-to-port connection that holds hundreds of services. In each topology, a customized approach to handling QoS will provide the best results.

The figure below shows the basic flow of NetStream Primo's QoS mechanism. Traffic ingresses (left to right) via the Ethernet or radio interfaces, on the "ingress path." Based on the services model, the system determines how to route the traffic. Traffic is then directed to the most appropriate output queue via the "egress path."

QoS Block Diagram



The ingress path consists of the following QoS building blocks:

- **Ingress Classifier** – A hierarchical mechanism that deals with ingress traffic on three different levels: interface, service point, and service. The classifier determines the exact traffic stream and associates it with the appropriate service. It also calculates an ingress frame CoS and Color. CoS and Color classification can be performed on three levels, according to the user's configuration.
- **Ingress Rate Metering** – A hierarchical mechanism that deals with ingress traffic on three different levels: interface, service point, and service point CoS. The rate metering mechanism enables the system to measure the incoming frame rate on different levels using a TrTCM standard MEF rate meter, and to determine whether to modify the color calculated during the classification stage.

The egress path consists of the following QoS building blocks:

- **Queue Manager** – This is the mechanism responsible for managing the transmission queues, utilizing smart WRED per queue and per packet color (Green or Yellow).
- **Scheduling and Shaping** – A hierarchical mechanism that is responsible for scheduling the transmission of frames from the transmission queues, based on priority among queues, Weighted Fair Queuing (WFQ) in bytes per each transmission queue, and eligibility to transmit based on required shaping on several different levels (per queue, per service bundle, and per port).
- **Marker** – This mechanism provides the ability to modify priority bits in frames based on the calculated CoS and Color.

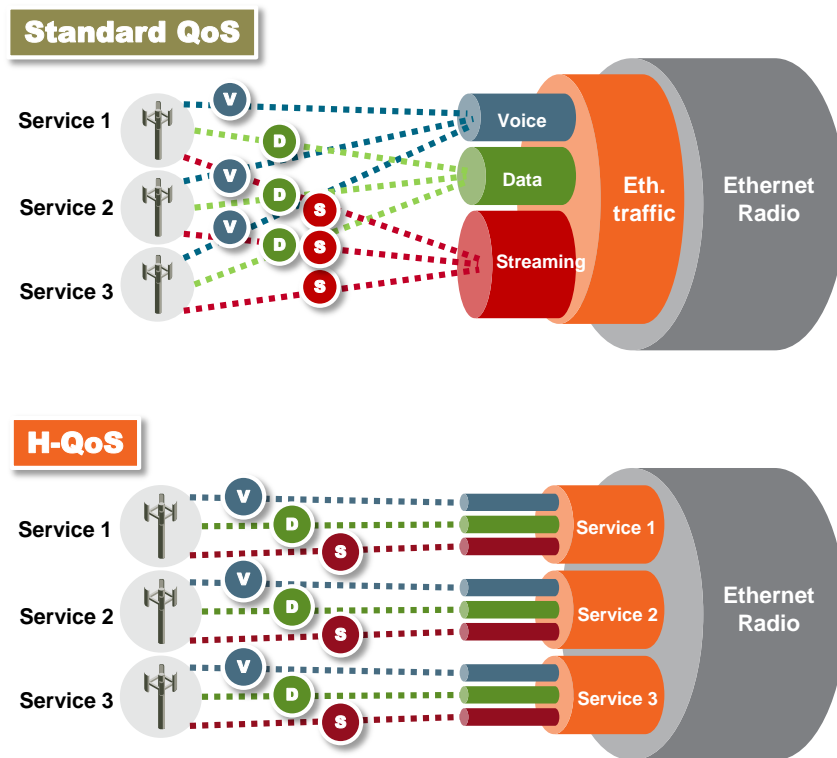
The following two modes of operation are available on the egress path:

- **Standard QoS** – This mode provides eight transmission queues per port.
- **Hierarchical QoS (H-QoS)** – In this mode, users can associate services from the service model to configurable groups of eight transmission queues

(service bundles), from a total 2K queues. In H-QoS mode, NetStream Primo performs QoS in a hierarchical manner in which the egress path is managed on three levels: ports, service bundles, and specific queues. This enables users to fully distinguish between streams, therefore providing a true SLA to customers.

The following figure illustrates the difference between how standard QoS and H-QoS handle traffic:

Standard QoS and H-QoS Comparison



5.2.6.1. QoS on the Ingress Path

5.2.6.2. Classification

NetStream Primo supports a hierarchical classification mechanism. The classification mechanism examines incoming frames and determines their CoS and Color. The benefit of hierarchical classification is that it provides the ability to “zoom in” or “zoom out”, enabling classification at higher or lower levels of the hierarchy. The nature of each traffic stream defines which level of the hierarchical classifier to apply, or whether to use several levels of the classification hierarchy in parallel.

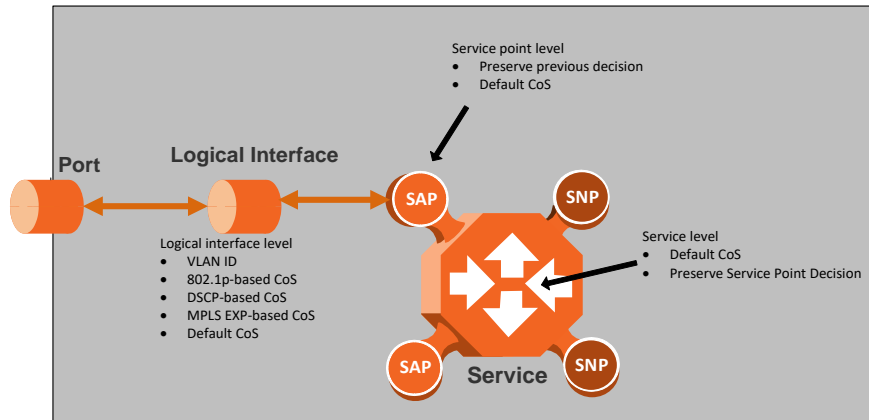
The hierarchical classifier consists of the following levels:

- Logical interface-level classification
- Service point-level classification
- Service level classification

The following figure illustrates the hierarchical classification model. In this figure, traffic enters the system via the port depicted on the left and enters the service via

the SAP depicted on the upper left of the service. The classification can take place at the logical interface level, the service point level, and/or the service level.

Hierarchical Classification



5.2.6.3. Logical Interface-Level Classification

Logical interface-level classification enables users to configure classification on a single interface or on a number of interfaces grouped together, such as a LAG group. The classifier at the logical interface level supports the following classification methods, listed from highest to lowest priority. A higher level classification method supersedes a lower level classification method:

- VLAN ID
- 802.1p bits.
- DSCP bits.
- MPLS EXP field.
- Default CoS

NetStream Primo performs the classification on each frame ingressing the system via the logical interface. Classification is performed step by step from the highest priority to the lowest priority classification method. Once a match is found, the classifier determines the CoS and Color decision for the frame for the logical interface-level.


For example, if the frame is an untagged IP Ethernet frame, a match will not be found until the third priority level (DSCP priority bits). The CoS and Color values defined for the frame's DSCP priority bits will be applied to the frame.

Users can disable some of these classification methods by configuring them as un-trusted. For example, if 802.1p classification is configured as un-trusted for a specific interface, the classification mechanism does not perform classification by VLAN UP bits. This is useful, for example, if the required classification is based on DSCP priority bits.

If no match is found at the logical interface level, the default CoS is applied to incoming frames at this level. In this case, the Color of the frame is assumed to be Green.

The following figure illustrates the hierarchy of priorities among classification methods, from highest (on the left) to lowest (on the right) priority.

Classification Method Priorities



Interface-level classification is configured as part of the logical interface configuration. For details, refer to *Ingress Path Classification at Logical Interface Level* on page 82.

The following tables show the default values for logical interface-level classification. The key values for these tables are the priority bits of the respective frame encapsulation layers (VLAN, IP, and MPLS), while the key results are the CoS and Colors calculated for incoming frames. These results are user-configurable, but it is recommended that only advanced users should modify the default values.

C-VLAN 802.1 UP and CFI Default Mapping to CoS and Color

802.1 UP	CFI	CoS (configurable)	Color (configurable)
0	0	0	Green
0	1	0	Yellow
1	0	1	Green
1	1	1	Yellow
2	0	2	Green
2	1	2	Yellow
3	0	3	Green
3	1	3	Yellow
4	0	4	Green
4	1	4	Yellow
5	0	5	Green
5	1	5	Yellow
6	0	6	Green
6	1	6	Yellow
7	0	7	Green
7	1	7	Yellow

S-VLAN 802.1 UP and DEI Default Mapping to CoS and Color

802.1 UP	DEI	CoS (Configurable)	Color (Configurable)
0	0	0	Green

802.1 UP	DEI	CoS (Configurable)	Color (Configurable)
0	1	0	Yellow
1	0	1	Green
1	1	1	Yellow
2	0	2	Green
2	1	2	Yellow
3	0	3	Green
3	1	3	Yellow
4	0	4	Green
4	1	4	Yellow
5	0	5	Green
5	1	5	Yellow
6	0	6	Green
6	1	6	Yellow
7	0	7	Green
7	1	7	Yellow

DSCP Default Mapping to CoS and Color

DSCP	DSCP (bin)	Description	CoS (Configurable)	Color (Configurable)
0 (default)	000000	BE (CS0)	0	Green
10	001010	AF11	1	Green
12	001100	AF12	1	Yellow
14	001110	AF13	1	Yellow
18	010010	AF21	2	Green
20	010100	AF22	2	Yellow
22	010110	AF23	2	Yellow
26	011010	AF31	3	Green
28	011100	AF32	3	Yellow
30	011110	AF33	3	Yellow
34	100010	AF41	4	Green
36	100100	AF42	4	Yellow
38	100110	AF43	4	Yellow
46	101110	EF	7	Green
8	001000	CS1	1	Green
16	010000	CS2	2	Green
24	011000	CS3	3	Green
32	100000	CS4	4	Green
40	101000	CS5	5	Green
48	110000	CS6	6	Green
51	110011	DSCP_51	6	Green
52	110100	DSCP_52	6	Green

DSCP	DSCP (bin)	Description	CoS (Configurable)	Color (Configurable)
54	110110	DSCP_54	6	Green
56	111000	CS7	7	Green

Default value is CoS equal best effort and Color equal Green.

MPLS EXP Default Mapping to CoS and Color

MPLS EXP bits	CoS (configurable)	Color (configurable)
0	0	Yellow
1	1	Green
2	2	Yellow
3	3	Green
4	4	Yellow
5	5	Green
6	6	Green
7	7	Green

5.2.6.4. Service Point-Level Classification

Classification at the service point level enables users to give special treatment, in higher resolution, to specific traffic flows using a single interface to which the service point is attached. The following classification modes are supported at the service point level. Users can configure these modes by means of the service point CoS mode.

- Preserve previous CoS decision (logical interface level)
- Default service point CoS

If the service point CoS mode is configured to preserve previous CoS decision, the CoS and Color are taken from the classification decision at the logical interface level. If the service point CoS mode is configured to default service point CoS mode, the CoS is taken from the service point's default CoS, and the Color is Green.

5.2.6.5. Service-Level Classification

Classification at the service level enables users to provide special treatment to an entire service. For example, the user might decide that all frames in a management service should be assigned a specific CoS regardless of the ingress port. The following classification modes are supported at the service level:

- Preserve previous CoS decision (service point level)
- Default CoS

If the service CoS mode is configured to preserve previous CoS decision, frames passing through the service are given the CoS and Color that was assigned at the service point level. If the service CoS mode is configured to default CoS mode, the CoS is taken from the service's default CoS, and the Color is Green.

5.2.6.6. Rate Meter (Policing)

NetStream Primo's TrTCM rate meter mechanism complies with MEF 10.2, and is based on a dual leaky bucket mechanism. The TrTCM rate meter can change a frame's CoS settings based on CIR/EIR+CBS/EBS, which makes the rate meter mechanism a key tool for implementing bandwidth profiles and enabling operators to meet strict SLA requirements.

The NetStream Primo hierarchical rate metering mechanism is part of the QoS performed on the ingress path, and consists of the following levels:

- Logical interface-level rate meter
- Service point-level rate meter¹²
- Service point CoS-level rate meter¹³

MEF 10.2 is the de-facto standard for SLA definitions, and NetStream Primo's QoS implementation provides the granularity necessary to implement service-oriented solutions.

Hierarchical rate metering enables users to define rate meter policing for incoming traffic at any resolution point, from the interface level to the service point level, and even at the level of a specific CoS within a specific service point. This option enables users to customize a set of eight policers for a variety of traffic flows within a single service point in a service.

Another important function of rate metering is to protect resources in the network element from malicious users sending traffic at an unexpectedly high rate. To prevent this, the rate meter can cut off traffic from a user that passes the expected ingress rate.

TrTCM rate meters use a leaky bucket mechanism to determine whether frames are marked Green, Yellow, or Red. Frames within the Committed Information Rate (CIR) or Committed Burst Size (CBS) are marked Green. Frames within the Excess Information Rate (EIR) or Excess Burst Size (EBS) are marked Yellow. Frames that do not fall within the CIR/CBS+EIR/EBS are marked Red and dropped, without being sent any further.

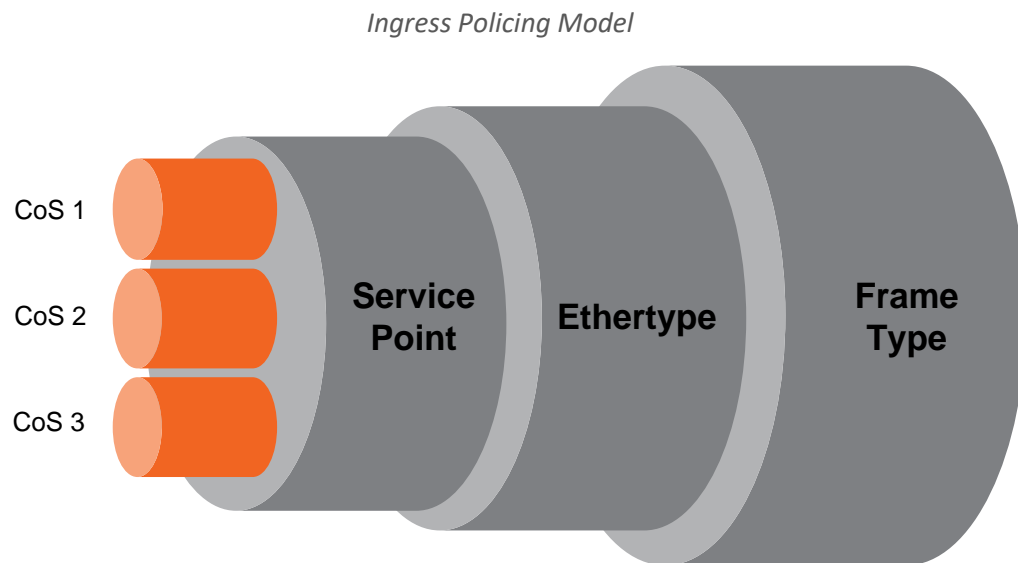
NetStream Primo provides up to 1024 user-defined TrTCM rate meters. The rate meters implement a bandwidth profile, based on CIR/EIR, CBS/EBS, Color Mode (CM), and Coupling flag (CF). Up to 250 different profiles can be configured.

Ingress rate meters operate at three levels:

- Logical Interface:
 - Per frame type (unicast, multicast, and broadcast)
 - Per frame ethertype
- Per Service Point
- Per Service Point CoS

¹² Service point-level rate metering is planned for future release.

¹³ Service point and CoS-level rate metering is planned for future release.



At each level (logical interface, service point, and service point + CoS), users can attach and activate a rate meter profile. Users must create the profile first, then attach it to the interface, service point, or service point + CoS.

5.2.6.7. Global Rate Meter Profiles

Users can define up to 250 rate meter user profiles. The following parameters can be defined for each profile:

- **Committed Information Rate (CIR)** – Frames within the defined CIR are marked Green and passed through the QoS module. Frames that exceed the CIR rate are marked Yellow. The CIR defines the average rate in bits/s of Service Frames up to which the network delivers service frames and meets the performance objectives. Permitted values are 0 to 1 Gbps, with a minimum granularity of 32Kbps.
- **Committed Burst Size (CBS)** – Frames within the defined CBS are marked Green and passed through the QoS module. This limits the maximum number of bytes available for a burst of service frames in order to ensure that traffic conforms to the CIR. Permitted values are 2 to 128 Kbytes, with a minimum granularity of 2 Kbytes.
- **Excess Information Rate (EIR)** – Frames within the defined EIR are marked Yellow and processed according to network availability. Frames beyond the combined CIR and EIR are marked Red and dropped by the policer. Permitted values are 0 to 1 Gbps, with a minimum granularity of 32 Kbps.
- **Excess Burst Size (EBS)** – Frames within the defined EBS are marked Yellow and processed according to network availability. Frames beyond the combined CBS and EBS are marked Red and dropped by the policer. Permitted values are 2 to 128 Kbytes, with a minimum granularity of 2 Kbytes.
- **Color Mode** – Color mode can be enabled (Color aware) or disabled (Color blind). In Color aware mode, all frames that ingress with a CFI/DEI field set to 1 (Yellow) are treated as EIR frames, even if credits remain in the CIR

bucket. In Color blind mode, all ingress frames are treated first as Green frames regardless of CFI/DEI value, then as Yellow frames (when there is no credit in the Green bucket). A Color-blind policer discards any previous Color decisions.

- **Coupling Flag** – If the coupling flag between the Green and Yellow buckets is enabled, then if the Green bucket reaches the maximum CBS value the remaining credits are sent to the Yellow bucket up to the maximum value of the Yellow bucket.

The following parameter is neither a profile parameter, nor specifically a rate meter parameter, but rather, is a logical interface parameter. For more information about logical interfaces, refer to *Logical Interfaces* on page 81.

- **Line Compensation** – A rate meter can measure CIR and EIR at Layer 1 or Layer 2 rates. Layer 1 capacity is equal to Layer 2 capacity plus 20 additional bytes for each frame due to the preamble and Inter Frame Gap (IFG). In most cases, the preamble and IFG equals 20 bytes, but other values are also possible. Line compensation defines the number of bytes to be added to each frame for purposes of CIR and EIR calculation. When Line Compensation is 20, the rate meter operates as Layer 1. When Line Compensation is 0, the rate meter operates as Layer 2. This parameter is very important to users that want to distinguish between Layer 1 and Layer 2 traffic. For example, 1 Gbps of traffic at Layer 1 is equal to ~760 Mbps if the frame size is 64 bytes, but ~986 Mbps if the frame size is 1500 bytes. This demonstrates that counting at Layer 2 is not always fair in comparison to counting at Layer 1, that is, the physical level.

5.2.6.8. Rate Metering (Policing) at the Logical Interface Level

Rate metering at the logical interface level supports the following:

- Unicast rate meter
- Multicast rate meter
- Broadcast rate meter
- User defined Ethertype 1 rate meter
- User defined Ethertype 2 rate meter
- User defined Ethertype 3 rate meter

For each rate meter, the following statistics are available:

- Green Frames (64 bits)
- Green Bytes (64 bits)
- Yellow Frames (64 bits)
- Yellow Bytes (64 bits)
- Red Frames (64 bits)
- Red Bytes (64 bits)

5.2.6.9. Rate Metering (Policing) at the Service Point Level

Users can define a single rate meter on each service point, up to a total number of 1024 rate meters per network element at the service point and CoS per service point levels.

The following statistics are available for each service point rate meter:

- Green Frames (64 bits)
- Green Bytes (64 bits)
- Yellow Frames (64 bits)
- Yellow Bytes (64 bits)
- Red Frames (64 bits)
- Red Bytes (64 bits)

5.2.6.10. Rate Metering (Policing) at the Service Point + CoS Level

Users can define a single rate meter for each CoS on a specific service point, up to a total number of 1024 rate meters per network element at the service point and CoS per service point levels.

The following statistics are available for each service point + CoS rate meter:

- Green Frames (64 bits)
- Green Bytes (64 bits)
- Yellow Frames (64 bits)
- Yellow Bytes (64 bits)
- Red Frames (64 bits)
- Red Bytes (64 bits)

5.2.6.11. QoS on the Egress Path

5.2.6.12. Queue Manager

The queue manager (QM) is responsible for managing the output transmission queues. NetStream Primo supports up to 2K service-level transmission queues, with configurable buffer size. Users can specify the buffer size of each queue independently. The total amount of memory dedicated to the queue buffers is 2 Gigabits.

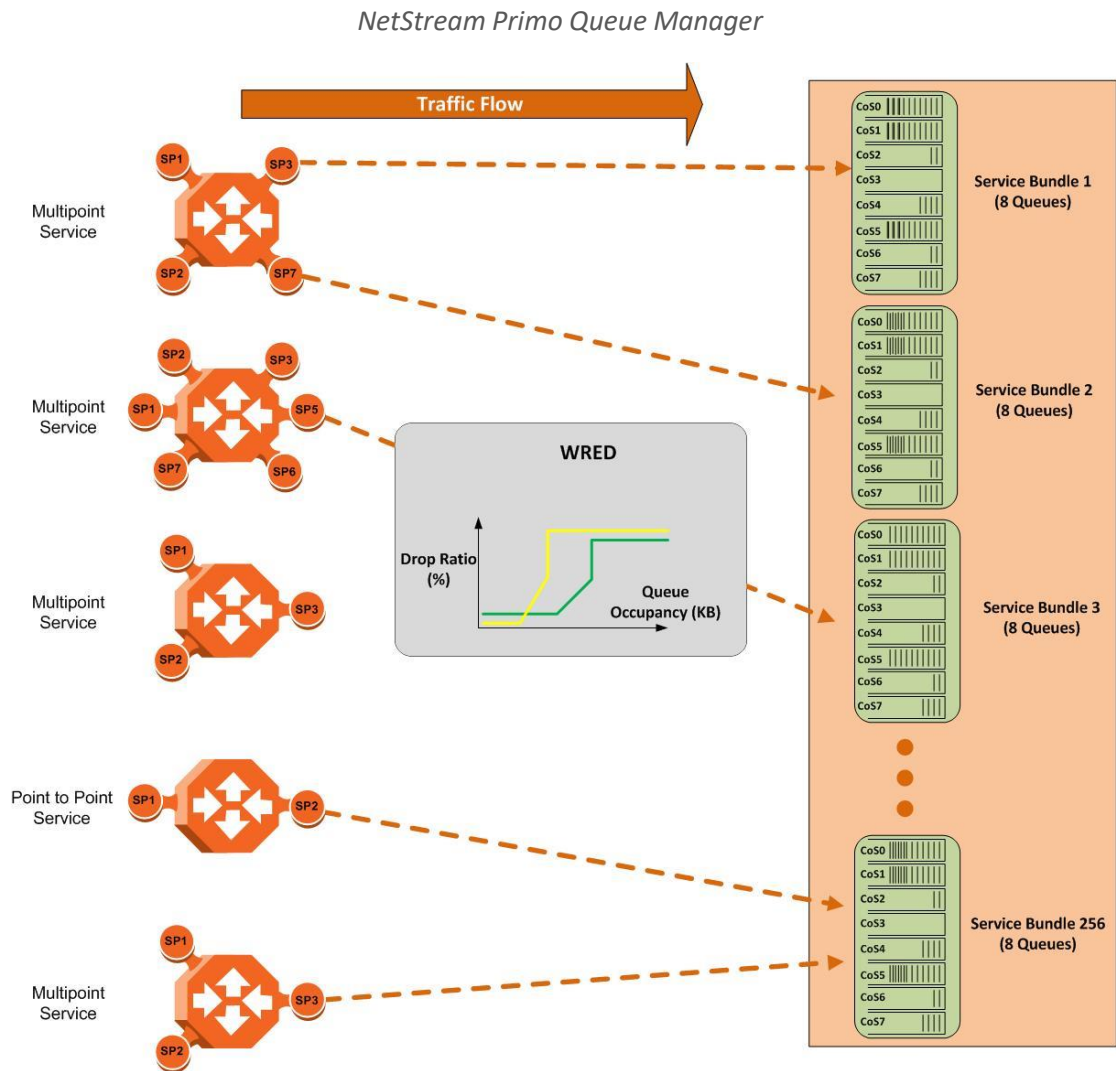
The following considerations should be taken into account in determining the proper buffer size:

- **Latency considerations** – If low latency is required (users would rather drop frames in the queue than increase latency) small buffer sizes are preferable.
- **Throughput immunity to fast bursts** – When traffic is characterized by fast bursts, it is recommended to increase the buffer sizes to prevent packet loss. Of course, this comes at the cost of a possible increase in latency.

Users can configure burst size as a tradeoff between latency and immunity to bursts, according to the application requirements.

The 2K queues are ordered in groups of eight queues. These eight queues correspond to CoS values, from 0 to 7; in other words, eight priority queues.

The following figure depicts the queue manager. Physically, the queue manager is located between the ingress path and the egress path.



In the figure above, traffic is passing from left to right. The traffic passing from the ingress path is routed to the correct egress destination interfaces via the egress service points. As part of the assignment of the service points to the interfaces, users define the group of eight queues through which traffic is to be transmitted out of the service point. This is part of the service point egress configuration.

After the traffic is tunneled from the ingress service points to the egress service points, it is aggregated into one of the eight queues associated with the specific service point. The exact queue is determined by the CoS calculated by the ingress path. For example, if the calculated CoS is 6, the traffic is sent to queue 6, and so on.

Before assigning traffic to the appropriate queue, the system makes a determination whether to forward or drop the traffic using a WRED algorithm with a

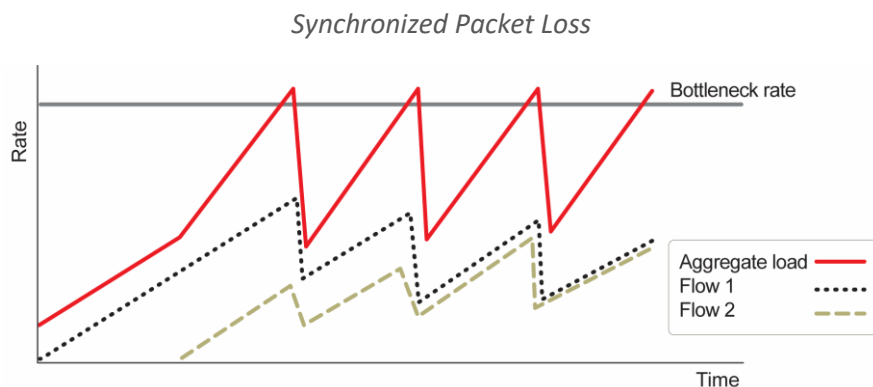
predefined green and yellow curve for the desired queue. This operation is integrated with the queue occupancy level.

The 2K queues share a single memory of 2 Gbits. NetStream Primo enables users to define a specific size for each queue which is different from the default size. Moreover, users can create an over-subscription scenario among the queues for when the buffer size of the aggregate queues is lower than the total memory allocated to all the queues. In doing this, the user must understand both the benefits and the potential hazards, namely, that if a lack of buffer space occurs, the queue manager will drop incoming frames without applying the usual priority rules among frames.

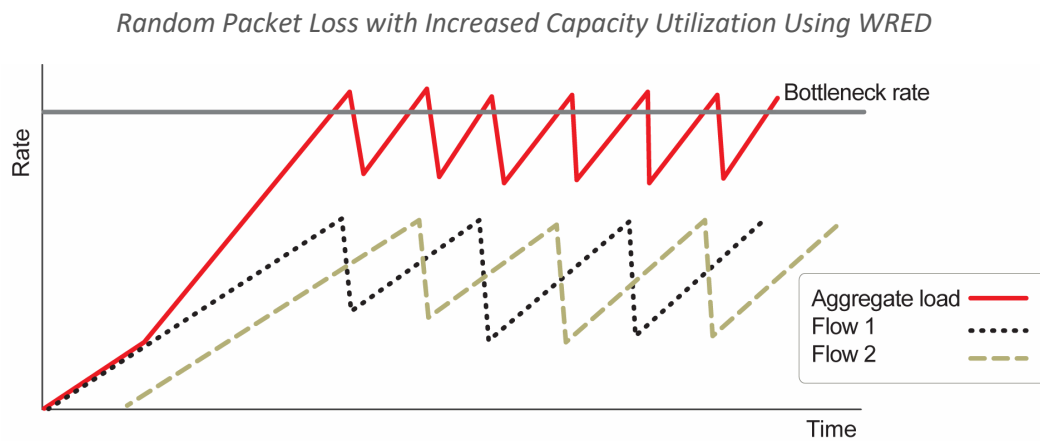
The queue size is defined by the WRED profile that is associated with the queue. For more details, refer to *WRED* on page 98.

5.2.6.13. WRED

The Weighted Random Early Detection (WRED) mechanism can increase capacity utilization of TCP traffic by eliminating the phenomenon of global synchronization. Global synchronization occurs when TCP flows sharing bottleneck conditions receive loss indications at around the same time. This can result in periods during which link bandwidth utilization drops significantly as a consequence of simultaneous falling to a "slow start" of all the TCP flows. The following figure demonstrates the behavior of two TCP flows over time without WRED.



WRED eliminates the occurrence of traffic congestion peaks by restraining the transmission rate of the TCP flows. Each queue occupancy level is monitored by the WRED mechanism and randomly selected frames are dropped before the queue becomes overcrowded. Each TCP flow recognizes a frame loss and restrains its transmission rate (basically by reducing the window size). Since the frames are dropped randomly, statistically each time another flow has to restrain its transmission rate as a result of frame loss (before the real congestion occurs). In this way, the overall aggregated load on the radio link remains stable while the transmission rate of each individual flow continues to fluctuate similarly. The following figure demonstrates the transmission rate of two TCP flows and the aggregated load over time when WRED is enabled.



When queue occupancy goes up, this means that the ingress path rate (the TCP stream that is ingressing the switch) is higher than the egress path rate. This difference in rates should be fixed in order to reduce packet drops and to reach the maximal media utilization, since NetStream Primo will not egress packets to the media at a rate which is higher than the media is able to transmit.

To deal with this, NetStream Primo enables users to define up to 30 WRED profiles. Each profile contains a Green traffic curve and a Yellow traffic curve. These curves describe the probability of randomly dropping frames as a function of queue occupancy. In addition, using different curves for Yellow packets and Green packets enables users to enforce the rule that Yellow packets be dropped before Green packets when there is congestion.

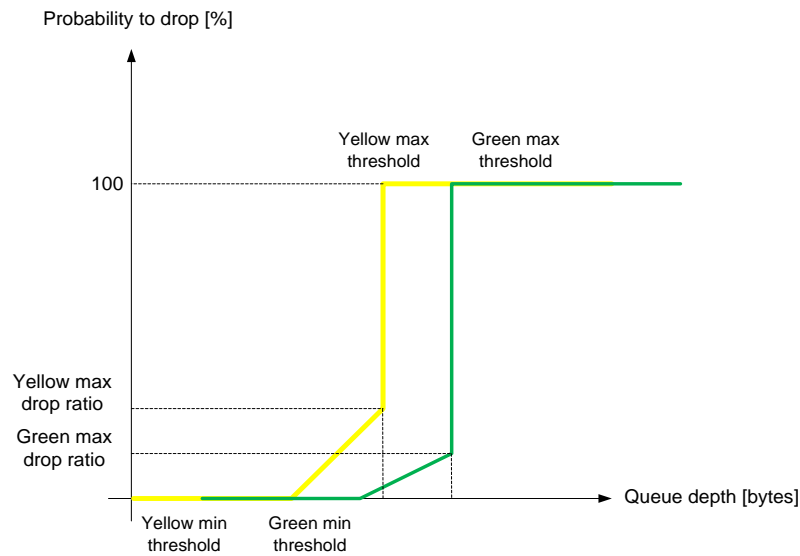
NetStream Primo also includes a pre-defined read-only WRED profile that defines a tail-drop curve. This profile is assigned profile number 31, and is configured with the following values:

- 100% Yellow traffic drop after 16kbytes occupancy.
- 100% Green traffic drop after 32kbytes occupancy.
- Yellow maximum drop is 100%
- Green maximum drop is 100%

A WRED profile can be assigned to each queue. The WRED profile assigned to the queue determines whether or not to drop incoming packets according to the occupancy of the queue. Basically, as queue occupancy grows, the probability of dropping each incoming frame increases as well. As a consequence, statistically more TCP flows will be restrained before traffic congestion occurs.

The following figure provides an example of a WRED profile.

WRED Profile Curve



The tail-drop profile, Profile 31, is the default profile for each queue. A tail drop curve is useful for reducing the effective queue size, such as when low latency must be guaranteed.

5.2.6.14. Global WRED Profile Configuration

NetStream Primo supports 30 user-configurable WRED profiles and one pre-defined (default) profile. The following are the WRED profile attributes:

- **Green Minimum Threshold** – Permitted values are 0 Kbytes to 8 Mbytes, with granularity of 8 Kbytes.
- **Green Maximum Threshold** – Permitted values are 0 Kbytes to 8 Mbytes, with granularity of 8 Kbytes.
- **Green-Maximum Drop** – Permitted values are 1% to 100%, with 1% drop granularity.
- **Yellow Minimum Threshold** – Permitted values are 0 Kbytes to 8 Mbytes, with granularity of 8 Kbytes.
- **Yellow Maximum Threshold** – Permitted values are 0 Kbytes to 8 Mbytes, with granularity of 8 Kbytes.
- **Yellow Maximum Drop** – Permitted values are 1% to 100%, with 1% drop granularity.



K is equal to 1024.

Users can enter any value within the permitted range. Based on the value entered by the user, the software automatically rounds off the setting

according to the granularity. If the user enters a value below the lowest granular value (except 0), the software adjusts the setting to the minimum.

For each curve, frames are passed on and not dropped up to the minimum Green and Yellow thresholds. From this point, WRED performs a pseudo-random drop with a ratio based on the curve up to the maximum Green and Yellow thresholds. Beyond this point, 100% of frames with the applicable Color are dropped.

The system automatically assigns the default “tail drop” WRED profile (Profile ID 31) to every queue. Users can change the WRED profile per queue based on the application served by the queue.

5.2.6.15. Standard QoS and Hierarchical QoS (H-QoS)

In a standard QoS mechanism, egress data is mapped to a single egress interface. This single interface supports up to eight priority queues, which correspond to the CoS of the data. Since all traffic for the interface egresses via these queues, there is no way to distinguish between different services and traffic streams within the same priority.

The figure below shows three services, each with three distinct types of traffic streams:

- Voice – high priority
- Data – medium priority
- Streaming – lower priority

While the benefits of QoS on the egress path can be applied to the aggregate streams, without H-QoS they will not be able to distinguish between the various services included in these aggregate streams. Moreover, different behavior among the different traffic streams that constitute the aggregate stream can cause unpredictable behavior between the streams. For example, in a situation in which one traffic stream can transmit 50 Mbps in a shaped manner while another can transmit 50 Mbits in a burst, frames may be dropped in an unexpected way due to a lack of space in the queue resulting from a long burst.

Hierarchical QoS (H-QoS) solves this problem by enabling users to create a real egress tunnel for each stream, or for a group of streams that are bundled together. This enables the system to fully perform H-QoS with a top-down resolution, and to fully control the required SLA for each stream.

5.2.6.16. H-QoS Hierarchy

The egress path hierarchy is based on the following levels:

- Queue level
- Service bundle level
- Logical interface level

The queue level represents the physical priority queues. This level holds 2K queues. Each eight queues are bundled and represent eight CoS priority levels. One or more service points can be attached to a specific bundle, and the traffic from the service

point to one of the eight queues is based on the CoS that was calculated on the ingress path.

**Note**

With standard QoS, all services are assigned to a single default service bundle.

The service bundle level represents the groups of eight priority queues. Every eight queues are managed as a single service bundle.

The interface level represents the physical port through which traffic from the specified service point egresses.

The following summarizes the egress path hierarchy:

- Up to 5 physical interfaces
- One service bundle per interface in standard QoS / 32 service bundles per interface in H-QoS.
- Eight queues per service bundle

5.2.6.17. H-QoS on the Interface Level

Users can assign a single leaky bucket shaper to each interface. The shaper on the interface level stops traffic from the interface if a specific user-defined peak information rate (PIR) has been exceeded.

In addition, users can configure scheduling rules for the priority queues, as follows:

- Scheduling (serve) priorities among the eight priority queues.
- Weighted Fair Queuing (WFQ) among queues with the same priority.

**Note**

The system assigns the rules for all service bundles under the interface.

RMON counters are valid on the interface level.

5.2.6.18. H-QoS on the Service Bundle Level

Users can assign a dual leaky bucket shaper to each service bundle. On the service bundle level, the shaper changes the scheduling priority if traffic via the service bundle is above the user-defined CIR and below the PIR. If traffic is above the PIR, the scheduler stops transmission for the service bundle.

Service bundle traffic counters are valid on this level.

**Note**

With standard QoS, users assign the egress traffic to a single service bundle (Service Bundle ID 1).

5.2.6.19. H-QoS on the Queue Level

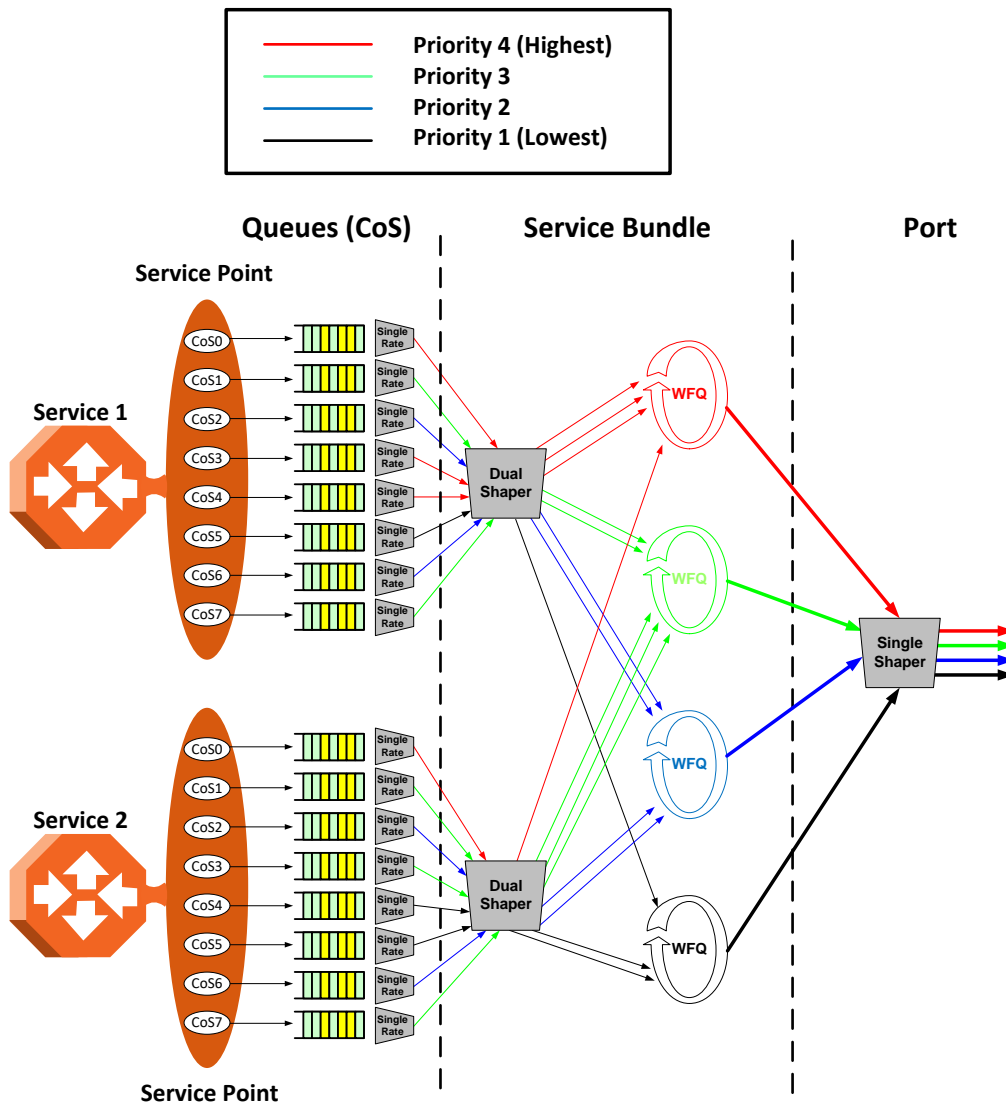
The egress service point points to a specific service bundle. Depending on the user application, the user can connect either a single service point or multiple service points to a service bundle. Usually, if multiple service points are connected to a service bundle, the service points will share the same traffic type and characteristics. Mapping to the eight priority queues is based on the CoS calculated on the ingress path, before any marking operation, which only changes the egress CoS and Color.

Users can assign a single leaky bucket to each queue. The shaper on the queue level stops traffic from leaving the queue if a specific user-defined PIR has been exceeded.

Traffic counters are valid on this level.

The following figure provides a detailed depiction of the H-QoS levels.

Detailed H-QoS Diagram



5.2.6.20. H- QoS Mode

As discussed above, users can select whether to work in Standard QoS mode or H-QoS mode.

- If the user configured all the egress service points to transmit traffic via a single service bundle, the operational mode is Standard QoS. In this mode, only Service Bundle 1 is active and there are eight output transmission queues.
- If the user configured the egress service points to transmit traffic via multiple service bundles, the operational mode is H-QoS. H-QoS mode enables users to fully distinguish among the streams and to achieve SLA per service.

5.2.6.21. Shaping on the Egress Path

Egress shaping determines the traffic profile for each queue. NetStream Primo performs egress shaping on the following three levels:

- Queue level – Single leaky bucket shaping.
- Service Bundle level – Dual leaky bucket shaping
- Interface level – Single leaky bucket shaping

5.2.6.22. Queue Shapers

Users can configure up to 31 single leaky bucket shaper profiles. The CIR value can be set to the following values:

- 16,000 – 32,000,000 bps – granularity of 16,000 bps
- 32,000,000 – 131,008,000 bps – granularity of 64,000 bps



Users can enter any value within the permitted range. Based on the value entered by the user, the software automatically rounds off the setting according to the granularity. If the user enters a value below the lowest granular value (except 0), the software adjusts the setting to the minimum.

Users can attach one of the configured queue shaper profiles to each priority queue. If no profile is attached to the queue, no egress shaping is performed on that queue.

5.2.6.23. Service Bundle Shapers

Users can configure up to 255 dual leaky bucket shaper profiles. The profiles can be configured as follows:

- Valid CIR values are:
 - 0 – 32,000,000 bps – granularity of 16,000 bps
 - 32,000,000 – 1,000,000,000 bps – granularity of 64,000 bps
- Valid PIR values are:

- 16,000 – 32,000,000 bps – granularity of 16,000 bps
- 32,000,000 – 1,000,000,000 bps – granularity of 64,000 bps

**Note**

Users can enter any value within the permitted range. Based on the value entered by the user, the software automatically rounds off the setting according to the granularity. If the user enters a value below the lowest granular value (except 0), the software adjusts the setting to the minimum.

Users can attach one of the configured service bundle shaper profiles to each service bundle. If no profile is attached to the service bundle, no egress shaping is performed on that service bundle.

5.2.6.24. Interface Shapers

Users can configure up to 31 single leaky bucket shaper profiles. The CIR can be set to the following values:

- 0 – 8,192,000 bps – granularity of 32,000 bps
- 8,192,000 – 32,768,000 bps – granularity of 128,000 bps
- 32,768,000 – 131,072,000 bps – granularity of 512,000 bps
- 131,072,000 – 999,424,000 bps – granularity of 8,192,000 bps

**Note**

Users can enter any value within the permitted range. Based on the value entered by the user, the software automatically rounds off the setting according to the granularity. If the user enters a value below the value (except 0), the software adjusts the setting to the minimum.

Users can attach one of the configured interface shaper profiles to each interface. If no profile is attached to the interface, no egress shaping is performed on that interface.

5.2.6.25. Line Compensation for Shaping

Users can configure a line compensation value for all the shapers under a specific logical interface. For more information, refer to *Global Rate Meter Profiles* on page 94.

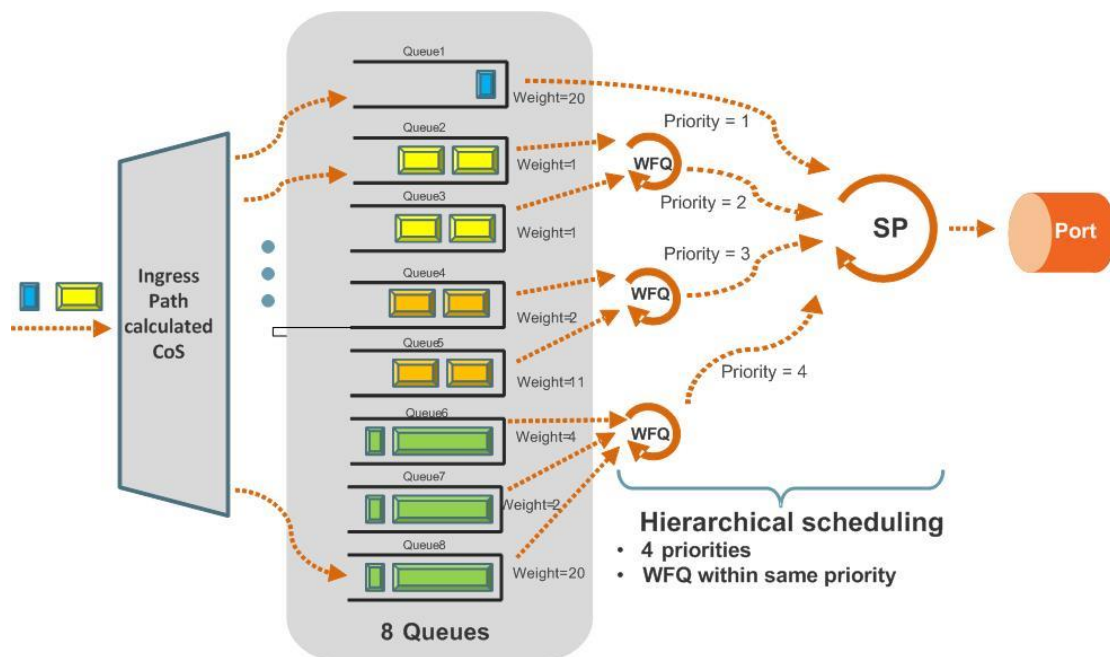
5.2.6.26. Egress Scheduling

Egress scheduling is responsible for transmission from the priority queues. NetStream Primo uses a unique algorithm with a hierarchical scheduling model over the three levels of the egress path that enables compliance with SLA requirements. The scheduler scans all the queues over all the service bundles, per interface, and determines which queue is ready to transmit. If more than one queue is ready to transmit, the scheduler determines which queue transmits first based on:

- **Queue Priority** – A queue with higher priority is served before lower-priority queues.
- **Weighted Fair Queuing (WFQ)** – If two or more queues have the same priority and are ready to transmit, the scheduler transmits frames from the queues based on a WFQ algorithm that determines the ratio of frames per queue based on a predefined weight assigned to each queue.

The following figure shows the scheduling mechanism for a single service bundle (equivalent to Standard QoS). When a user assigns traffic to more than a single service bundle (H-QoS mode), multiple instances of this model (up to 32 per port) are valid.

Scheduling Mechanism for a Single Service Bundle



5.2.6.27. Interface Priority

The profile defines the exact order for serving the eight priority queues in a single service bundle. When the user attaches a profile to an interface, all the service bundles under the interface inherit the profile.

The priority mechanism distinguishes between two states of the service bundle:

- Green State – Committed state
- Yellow state – Best effort state

Green State refers to any time when the *service bundle total rate* is below the user-defined CIR. Yellow State refers to any time when the *service bundle total rate* is above the user-defined CIR but below the PIR.

User can define up to four Green priority profiles, from 4 (highest) to 1 (lowest). An additional four Yellow priority profiles are defined automatically.

The following table provides a sample of an interface priority profile. This profile is also used as the default interface priority profile.

QoS Priority Profile Example

Profile ID (1-9)			
CoS	Green Priority (user defined)	Yellow Priority (read only)	Description
0	1	1	Best Effort
1	2	1	Data Service 4
2	2	1	Data Service 3
3	2	1	Data Service 2
4	2	1	Data Service 1
5	3	1	Real Time 2 (Video with large buffer)
6	3	1	Real Time 1 (Video with small buffer)
7	4	4	Management (Sync, PDUs, etc.)

When the service bundle state is Green (committed state), the service bundle priorities are as defined in the Green Priority column. When the service bundle state is Yellow (best effort state), the service bundle priorities are system-defined priorities shown in the Yellow Priority column.

**Note**

CoS 7 is always marked with the highest priority, no matter what the service bundle state is, since it is assumed that only high priority traffic will be tunneled via CoS 7.

The system supports up to nine interface priority profiles. Profiles 1 to 8 are defined by the user, while profile 9 is the pre-defined read-only default interface priority profile.

The following interface priority profile parameters can be configured by users:

- **Profile ID** – Profile ID number. Permitted values are 1 to 8.
- **CoS 0 Priority** – CoS 0 queue priority, from 4 (highest) to 1 (lowest).
- **CoS 0 Description** – CoS 0 user description field, up to 20 characters.
- **CoS 1 Priority** – CoS 1 queue priority, from 4 (highest) to 1 (lowest).
- **CoS 1 Description** – CoS 1 user description field, up to 20 characters.
- **CoS 2 Priority** – CoS 2 queue priority, from 4 (highest) to 1 (lowest).
- **CoS 2 Description** – CoS 2 user description field, up to 20 characters.
- **CoS 3 Priority** – CoS 3 queue priority, from 4 (highest) to 1 (lowest).
- **CoS 3 Description** – CoS 3 user description field, up to 20 characters.
- **CoS 4 Priority** – CoS 4 queue priority, from 4 (highest) to 1 (lowest).
- **CoS 4 Description** – CoS 4 user description field, up to 20 characters.
- **CoS 5 Priority** – CoS 5 queue priority, from 4 (highest) to 1 (lowest).
- **CoS 5 Description** – CoS 5 user description field, up to 20 characters.
- **CoS 6 Priority** – CoS 6 queue priority, from 4 (highest) to 1 (lowest).

- **CoS 6 Description** – CoS 6 user description field, up to 20 characters.
- **CoS 7 Priority** – CoS 7 queue priority, from 4 (highest) to 1 (lowest).
- **CoS 7 Description** – CoS 7 user description field, up to 20 characters.

Users can attach one of the configured interface priority profiles to each interface. By default, the interface is assigned Profile ID 9, the pre-defined system profile.

5.2.6.28. Weighted Fair Queuing (WFQ)

As described above, the scheduler serves the queues based on their priority, but when two or more queues have data to transmit and their priority is the same, the scheduler uses Weighted Fair Queuing (WFQ) to determine the priorities within each priority. WFQ defines the transmission ratio, in bytes, between the queues. All the service bundles under the interface inherit the WFQ profile attached to the interface.

The system supports up to six WFQ interface profiles. Profile ID 1 is a pre-defined read-only profile, and is used as the default profile. Profiles 2 to 6 are user-defined profiles.

The following table provides an example of a WFQ profile.

WFQ Profile Example

Profile ID (1-7)		
CoS	Queue Weight (Green)	Queue Weight (Yellow – not visible to users)
0	20	20
1	20	20
2	20	20
3	20	20
4	20	20
5	20	20
6	20	20
7	20	20

For each CoS, the user can define;

- **Profile ID** – Profile ID number. Permitted values are 2 to 6.
- **Weight** – Transmission quota in bytes. Permitted values are 1 to 20.

Users can attach one of the configured interface WFQ profiles to each interface. By default, the interface is assigned Profile ID 1, the pre-defined system profile.

5.2.6.29. Egress Statistics

5.2.6.30. Queue-Level Statistics

NetStream Primo supports the following counters per queue at the queue level:

- Transmitted Green Packet (64 bits counter)
- Transmitted Green Bytes (64 bits counter)

- Transmitted Green Bits per Second (32 bits counter)
- Dropped Green Packets (64 bits counter)
- Dropped Green Bytes (64 bits counter)
- Transmitted Yellow Packets (64 bits counter)
- Transmitted Yellow Bytes (64 bits counter)
- Transmitted Yellow Bits per Second (32 bits counter)
- Dropped Yellow Packets (64 bits counter)
- Dropped Yellow Bytes (64 bits counter)

5.2.6.31. Service Bundle-Level Statistics

NetStream Primo supports the following counters per service bundle at the service bundle level:

- Transmitted Green Packets (64 bits counter)
- Transmitted Green Bytes (64 bits counter)
- Transmitted Green Bits per Second (32 bits counter)
- Dropped Green Packets (64 bits counter)
- Dropped Green Bytes (64 bits counter)
- Transmitted Yellow Packets (64 bits counter)
- Transmitted Yellow Bytes (64 bits counter)
- Transmitted Yellow Bits per Second (32 bits counter)
- Dropped Yellow Packets (64 bits counter)
- Dropped Yellow Bytes (64 bits counter)

5.2.6.32. Interface-Level Statistics

For information on statistics at the interface level, refer to *Ethernet Statistics* on page 80.

5.2.6.33. Marker

Marking refers to the ability to overwrite the outgoing priority bits and Color of the outer VLAN of the egress frame. Marking mode is only applied if the outer frame is S-VLAN and S-VLAN CoS preservation is disabled, or if the outer frame is C-VLAN and C-VLAN CoS preservation is disabled. If outer VLAN preservation is enabled for the relevant outer VLAN, the egress CoS and Color are the same as the CoS and Color of the frame when it ingress into the switching fabric.

Marking is performed according to a global table that maps CoS and Color values to the 802.1p-UP bits and the DEI or CFI bits. If Marking is enabled on a service point, the CoS and Color of frames egressing the service via that service point are overwritten according to this global mapping table.

If marking and CoS preservation for the relevant outer VLAN are both disabled, marking is applied according to the Green frame values in the global marking table. When marking is performed, the following global tables are used by the marker to decide which CoS and Color to use as the egress CoS and Color bits.

802.1q UP Marking Table (C-VLAN)

CoS	Color	802.1q UP (Configurable)	CFI Color (Configurable)
0	Green	0	0
0	Yellow	0	1
1	Green	1	0
1	Yellow	1	1
2	Green	2	0
2	Yellow	2	1
3	Green	3	0
3	Yellow	3	1
4	Green	4	0
4	Yellow	4	1
5	Green	5	0
5	Yellow	5	1
6	Green	6	0
6	Yellow	6	1
7	Green	7	0
7	Yellow	7	1

802.1ad UP Marking Table (S-VLAN)

CoS	Color	802.1ad UP (configurable)	DEI Color (configurable)
0	Green	0	0
0	Yellow	0	1
1	Green	1	0
1	Yellow	1	1
2	Green	2	0
2	Yellow	2	1
3	Green	3	0
3	Yellow	3	1
4	Green	4	0
4	Yellow	4	1
5	Green	5	0
5	Yellow	5	1
6	Green	6	0
6	Yellow	6	1

CoS	Color	802.1ad UP (configurable)	DEI Color (configurable)
7	Green	7	0
7	Yellow	7	1

The keys for these tables are the CoS and Color. The results are the 802.1q/802.1ad UP and CFI/DEI bits, which are user-configurable. It is strongly recommended that the default values not be changed except by advanced users.

5.2.6.34. Standard QoS and Hierarchical QoS (H-QoS) Summary

The following table summarizes and compares the capabilities of standard QoS and H-QoS.

Summary and Comparison of Standard QoS and H-QoS

Capability	Standard QoS	Hierarchical QoS
Number of transmission queues per port	8	256
Number of service bundles	1 (always service bundle id equal 1)	32
WRED	Per queue (two curves – for green traffic and for yellow traffic via the queue)	Per queue (two curves – for green traffic and for yellow traffic via the queue)
Shaping at queue level	Single leaky bucket	Single leaky bucket
Shaping at service bundle level	Dual leaky bucket	Dual leaky bucket
Shaping at port level	Single leaky bucket (this level is not relevant since it is recommended to use service bundle level with dual leaky bucket)	Single leaky bucket
Transmission queues priority	Per queue priority (4 priorities).	Per queue priority (4 priorities). All service bundles for a specific port inherit the 8-queues priority settings.
Weighted fair Queue (WFQ)	Queue level (between queues)	Queue level (between queues) Service Bundle level (between service bundles)
Marker	Supported	Supported
Statistics	Queue level (8 queues) Service bundle level (1 service bundle) Port level	Queue level (256 queues) Service bundle level (32 service bundles) Port level

5.2.7. Global Switch Configuration

The following parameters are configured globally for the NetStream Primo switch:

- **S-VLAN Ethertype** –Defines the ethertype recognized by the system as the S-VLAN ethertype. NetStream Primo supports the following S-VLAN ethertypes:
 - 0x8100
 - 0x88A8 (default)
 - 0x9100
 - 0x9200

- **C-VLAN Ethertype** – Defines the ethertype recognized by the system as the C-VLAN ethertype. NetStream Primo supports 0x8100 as the C-VLAN ethertype.
- **MRU** – The maximum segment size defines the maximum receive unit (MRU) capability and the maximum transmit capability (MTU) of the system. Users can configure a global MRU for the system. Permitted values are 64 bytes to 9612 bytes.

5.2.8. Automatic State Propagation

Related topics:

- Network Resiliency
- External Protection
- Link Aggregation Groups (LAG)

Automatic State Propagation (ASP) enables propagation of radio failures back to the Ethernet port. You can also configure ASP to close the Ethernet port based on a radio failure at the remote carrier. ASP improves the recovery performance of resiliency protocols.



Note

It is recommended to configure both ends of the link to the same ASP configuration.

5.2.8.1. Automatic State Propagation Operation

ASP is configured as pairs of interfaces. Each ASP pair includes a Monitored Interface and a Controlled Interface. The Monitored Interface is a radio interface. The Controlled Interface is an Ethernet interface. Only one ASP pair can be configured per radio interface, and only one ASP pair can be configured per Ethernet interface.

The following events in the Monitored Interface trigger ASP:

- Radio LOF
- Radio Excessive BER

The user can also configure the ASP pair so that Radio LOF, Radio Excessive BER, or loss of the Ethernet connection at the remote side of the link will also trigger ASP.

When a triggering event takes place:

- If the Controlled Interface is an electrical GbE port, the port is closed.
- If the Controlled Interface is an optical GbE port, the port is muted.

The Controlled Interface remains closed or muted until all triggering events are cleared.

In addition, when a local triggering event takes place, the ASP mechanism sends an indication to the remote side of the link. Even when no triggering event has taken place, the ASP mechanism sends periodic update messages indicating that no triggering event has taken place.

5.2.8.2. Automatic State Propagation and Protection

When the Controlled Interface is part of a 1+1 HSB protection configuration, a port shutdown message is only sent to the remote side of the link if both of the protected interfaces are shut down.

In a 1+1 HSB configuration using Multi-Unit LAG mode, in which two Ethernet interfaces on each unit belong to a static LAG, an ASP triggering event only shuts down the external user port.

When the Monitored interface is part of a 1+1 HSB configuration, ASP is only triggered if both interfaces fail.

Closing an Ethernet port because of ASP does not trigger a protection switch.

5.2.8.3. Preventing Loss of In-Band Management

If the link uses in-band management, shutting down the Ethernet port can cause loss of management access to the unit. To prevent this, users can configure ASP to operate in Client Signal Failure (CSF) mode. In CSF mode, the ASP mechanism does not physically shut down the Controlled Interface when ASP is triggered. Instead, the ASP mechanism sends a failure indication message (a CSF message). The CSF message is used to propagate the failure indication to external equipment.

CSF mode is particularly useful when the NetStream Primo unit is an element in the following network topologies:

- Ring or mesh network topology.
- An NS DIPLO/PRIMON connected to a NetStream Primo unit being utilized as a pipe via an Ethernet interface (back-to-back on the same site).¹⁴
- Payload traffic is spanned by G.8032 in the network.
- In-band management is spanned by MSTP in the network.
- A NetStream Primo unit being utilized as a pipe is running one MSTP instance for spanning in-band management.



Note

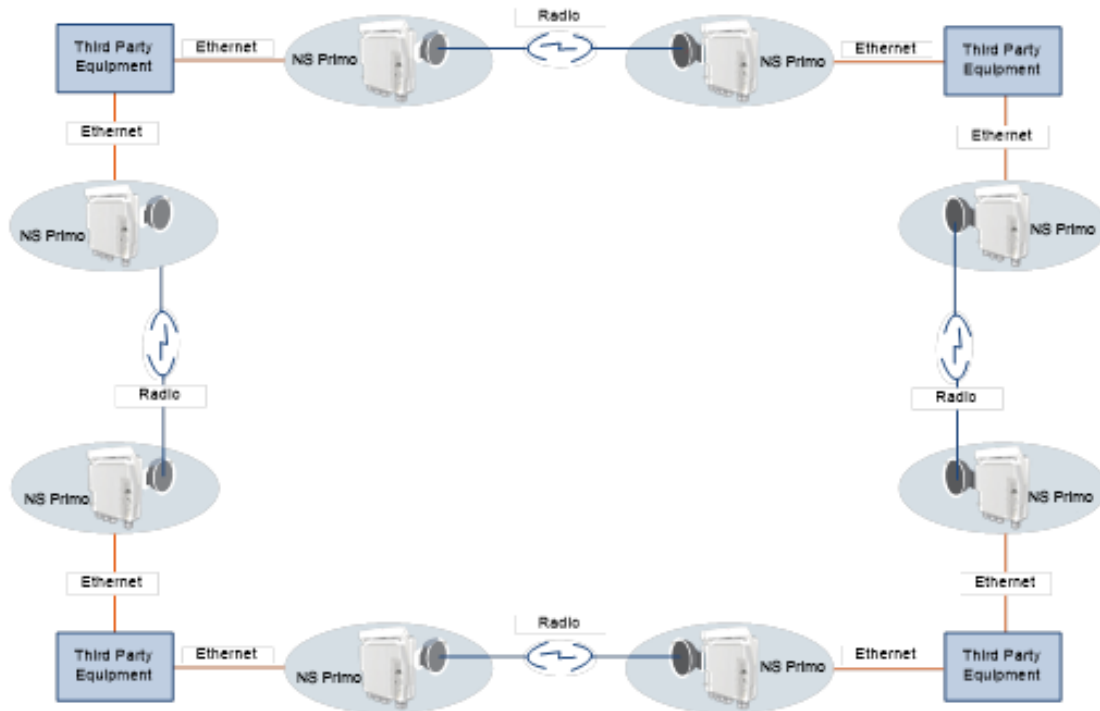
CSF mode is planned for future release.

5.2.9. Adaptive Bandwidth Notification (EOAM)

Adaptive Bandwidth Notification (ABN, also known as EOAM) enables third party applications to learn about bandwidth changes in a radio link when ACM is active. Once ABN is enabled, the radio unit reports bandwidth information to third-party switches.


¹⁴ ASP interoperability among Primo/Diplo units requires that all units be running software version 7.7 or higher.

Network Topology with NetStream Primo Units and Third-Party Equipment



The ABN entity creates a logical relationship between a radio interface, called the Monitored Interface, and an Ethernet interface or a logical group of Ethernet interfaces, called the Control Interface. When bandwidth degrades from the nominal bandwidth value in the Monitored Interface, messages relaying the actual bandwidth values are periodically sent over the Control Interface. A termination message is sent once the bandwidth returns to its nominal level.

ABN Entity



The nominal bandwidth is calculated by the system based on the maximum bandwidth profile.

The ABN entity measures the bandwidth in samples once a change in profile takes place. A weighted average is calculated based on the samples at regular, user-defined intervals to determine whether a bandwidth degradation event has occurred. Bandwidth degradation is reported only if the measured bandwidth remains below the nominal bandwidth at the end of a user-defined hold off period. This prevents the NetStream Primo from reporting bandwidth degradation due to short fading events.

5.2.10. Network Resiliency

NetStream Primo provides carrier-grade service resiliency using the following protocols:

- G.8032 Ethernet Ring Protection Switching (ERPS)
- Multiple Spanning Tree Protocol (MSTP)

These protocols are designed to prevent loops in ring/mesh topologies.



Note

G.8032 and MSTP are planned for future release.

5.2.10.1. G.8032 Ethernet Ring Protection Switching (ERPS)

ERPS, as defined in the G.8032 ITU standard, is currently the most advanced ring protection protocol, providing convergence times of sub-50ms. ERPS prevents loops in an Ethernet ring by guaranteeing that at any time, traffic can flow on all except one link in the ring. This link is called the Ring Protection Link (RPL). Under normal conditions, the RPL is blocked, i.e., not used for traffic. One designated Ethernet Ring Node, the RPL Owner Node, is responsible for blocking traffic at one end of the RPL. When an Ethernet ring failure occurs, the RPL Owner unblocks its end of the RPL, allowing the RPL to be used for traffic. The other Ethernet Ring Node adjacent to the RPL, the RPL Neighbor Node, may also participate in blocking or unblocking its end of the RPL. A number of ERP instances (ERPis) can be created on the same ring.

5.2.10.2. G.8032 ERPS Benefits

ERPS, as the most advanced ring protection protocol, provides the following benefits:

- Provides sub-50ms convergence times.
- Provides service-based granularity for load balancing, based on the ability to configure multiple ERPis on a single physical ring.
- Provides configurable timers to control switching and convergence parameters per ERPis.

5.2.10.3. G.8032 ERPS Operation

The ring protection mechanism utilizes an APS protocol to implement the protection switching actions. Forced and manual protection switches can also be initiated by the user, provided the user-initiated switch has a higher priority than any other local or far-end request.

Ring protection switching is based on the detection of defects in the transport entity of each link in the ring. For purposes of the protection switching process, each transport entity within the protected domain has a state of either Signal Fail (SF) or Non-Failed (OK). R-APS control messages are forwarded by each node in the ring to update the other nodes about the status of the links.

**Note**

An additional state, Signal Degrade (SD), is planned for future release. The SD state is similar to SF, but with lower priority.

Users can configure up to 16 ERPIs. Each ERPI is associated with an Ethernet service defined in the system. This enables operators to define a specific set of G.8032 characteristics for individual services or groups of services within the same physical ring. This includes a set of timers that enables operators to optimize protection switching behavior per ERPI:


- **Wait to Restore (WTR) Timer** – Defines a minimum time the system waits after signal failure is recovered before reverting to idle state.
- **Guard Time** – Prevents unnecessary state changes and loops.
- **Hold-off Time** – Determines the time period from failure detection to response.

Each ERPI maintains a state machine that defines the node's state for purposes of switching and convergence. The state is determined according to events that occur in the ring, such as signal failure and forced or manual switch requests, and their priority. Possible states are:

- Idle
- Protecting
- Forced Switch (FS)
- Manual Switch (MS)
- Pending


As shown in the following figure, in idle (normal) state, R-APS messages pass through all links in the ring, while the RPL is blocked for traffic. The RPL can be on either edge of the ring. R-APS messages are sent every five seconds.

G.8032 Ring in Idle (Normal) State




Once a signal failure is detected, the RPL is unblocked for each ERPI. As shown in the following figure, the ring switches to protecting state. The nodes that detect the failure send periodic SF messages to alert the other nodes in the link of the failure and initiate the protecting state.

G.8032 Ring in Protecting State



The ability to define multiple ERPIs and assign them to different Ethernet services or groups of services enables operators to perform load balancing by configuring a different RPL for each ERPI. The following figure illustrates a ring in which four ERPIs each carry services with 33% capacity in idle state, since each link is designated the RPL, and is therefore idle, for a different ERPI.

Load Balancing Example in G.8032 Ring



5.2.10.4. Multiple Spanning Tree Protocol (MSTP)

MSTP, as defined in IEEE 802.1q, provides full connectivity for frames assigned to any given VLAN throughout a bridged LAN consisting of arbitrarily interconnected bridges.

With MSTP, an independent multiple spanning tree instance (MSTI) is configured for each group of services, and only one path is made available (unblocked) per spanning tree instance. This prevents network loops and provides load balancing capability. It also enables operators to differentiate among Ethernet services by mapping them to different, specific MSTIs. The maximum number of MSTIs is configurable, from 2 to 16.

MSTP is an extension of, and is backwards compatible with, Rapid Spanning Tree Protocol (RSTP).

NetStream Primo supports MSTP according to the following IEEE standards:

- 802.1q
- 802.1ad amendment (Q-in-Q)
- 802.1ah (TE instance)

5.2.10.5. MSTP Benefits

MSTP significantly improves network resiliency in the following ways:

- Prevents data loops by configuring the active topology for each MSTI such that there is never more than a single route between any two points in the network.

- Provides for fault tolerance by automatically reconfiguring the spanning tree topology whenever there is a bridge failure or breakdown in a data path.
- Automatically reconfigures the spanning tree to accommodate addition of bridges and bridge ports to the network, without the formation of transient data loops.
- Enables frames assigned to different services or service groups to follow different data routes within administratively established regions of the network.
- Provides for predictable and reproducible active topology based on management of the MSTP parameters.
- Operates transparently to the end stations.
- Consumes very little bandwidth to establish and maintain MSTIs, constituting a small percentage of the total available bandwidth which is independent of both the total traffic supported by the network and the total number of bridges or LANs in the network.
- Does not require bridges to be individually configured before being added to the network.

5.2.10.6. MSTP Operation

MSTP includes the following elements:

- **MST Region** – A set of physically connected bridges that can be portioned into a set of logical topologies.
- **Internal Spanning Tree (IST)** – Every MST Region runs an IST, which is a special spanning tree instance that disseminates STP topology information for all other MSTIs.
- **CIST Root** – The bridge that has the lowest Bridge ID among all the MST Regions.
- **Common Spanning Tree (CST)** – The single spanning tree calculated by STP, RSTP, and MSTP to connect MST Regions. All bridges and LANs are connected into a single CST.
- **Common Internal Spanning Tree (CIST)** – A collection of the ISTs in each MST Region, and the CST that interconnects the MST regions and individual spanning trees. MSTP connects all bridges and LANs with a single CIST.

MSTP specifies:

- An MST Configuration Identifier that enables each bridge to advertise its configuration for allocating frames with given VIDs to any of a number of MSTIs.
- A priority vector that consists of a bridge identifier and path cost information for the CIST.
- An MSTI priority vector for any given MSTI within each MST Region.

Each bridge selects a CIST priority vector for each port based on the priority vectors and MST Configuration Identifiers received from the other bridges and on an

incremental path cost associated with each receiving port. The resulting priority vectors are such that in a stable network:

- One bridge is selected to be the CIST Root.
- A minimum cost path to the CIST Root is selected for each bridge.
- The CIST Regional Root is identified as the one root per MST Region whose minimum cost path to the root is not through another bridge using the same MST Configuration Identifier.

Based on priority vector comparisons and calculations performed by each bridge for each MSTI, one bridge is independently selected for each MSTI to be the MSTI Regional Root, and a minimum cost path is defined from each bridge or LAN in each MST Region to the MSTI Regional Root.

The following events trigger MSTP re-convergence:

- Addition or removal of a bridge or port.
- A change in the operational state of a port or group (LAG or protection).
- A change in the service to instance mapping.
- A change in the maximum number of MSTIs.
- A change in an MSTI bridge priority, port priority, or port cost.



Note

All except the last of these triggers can cause the entire MSTP to re-converge. The last trigger only affects the modified MSTI.

5.2.10.7. MSTP Interoperability

MSTP in NetStream Primo units is interoperable with:

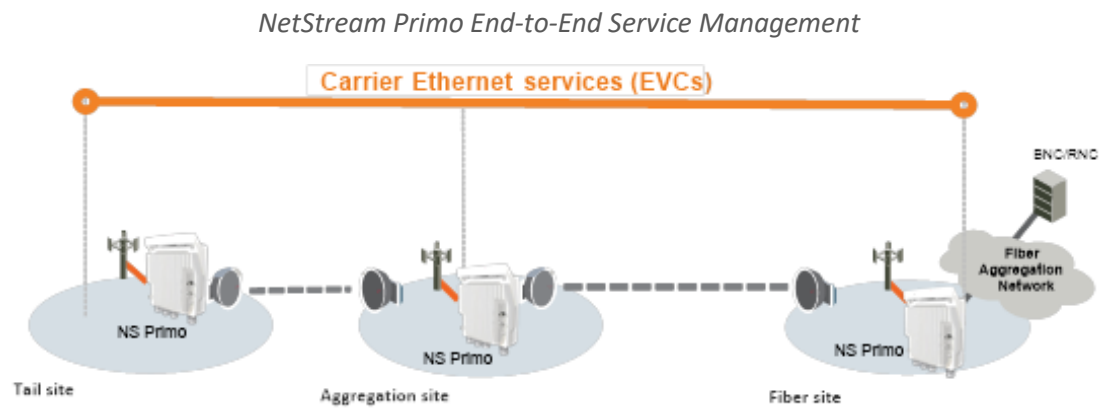
- Old NetAxis units running RSTP.
- Third-party bridges running MSTP.
- Third-party bridges running RSTP

5.2.11. OAM

NetStream Primo provides complete Service Operations Administration and Maintenance (SOAM) functionality at multiple layers, including:

- Fault management status and alarms.
- Maintenance signals, such as AIS, and RDI.
- Maintenance commands, such as loopbacks and Linktrace commands.

NetStream Primo is fully compliant with 802.1ag, G.8013/Y.1731, MEF-17, MEF-20, MEF-30, and MEF-31.



5.2.11.1. Connectivity Fault Management (FM)



Note

This feature is planned for future release.

The IEEE 802.1ag and G.8013/Y.1731 standards and the MEF-17, MEF-20, MEF-30, and MEF-31 specifications define SOAM. SOAM is concerned with detecting, isolating, and reporting connectivity faults spanning networks comprising multiple LANs, including LANs other than IEEE 802.3 media.

IEEE 802.1ag Ethernet FM (Connectivity Fault Management) consists of three protocols that operate together to aid in fault management:

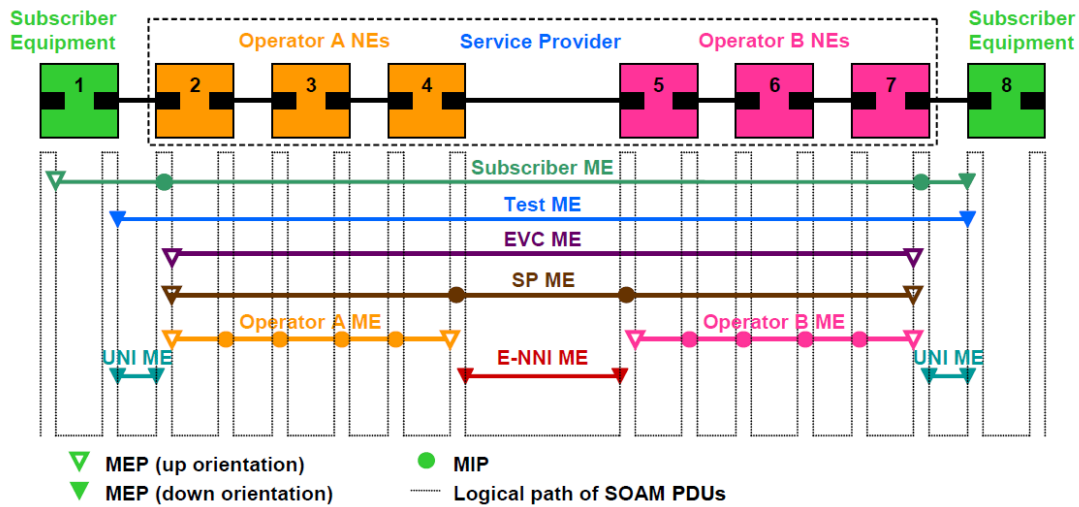
- Continuity check
- Link trace
- Loopback.

NetStream Primo utilizes these protocols to maintain smooth system operation and non-stop data flow.

The following are the basic building blocks of FM:

- Maintenance domains, their constituent maintenance points, and the managed objects required to create and administer them.

SOAM Maintenance Entities (Example)



- Protocols and procedures used by maintenance points to maintain and diagnose connectivity faults within a maintenance domain.
 - CCM (Continuity Check Message): CCM can detect Connectivity Faults (loss of connectivity or failure in the remote MEP).
 - Loopback: LBM/LBR mechanism is an on-demand mechanism. It is used to verify connectivity from any MEP to any certain Maintenance Point in the MA/MEG. A session of loopback messages can include up to 1024 messages with varying intervals ranging from 1 to 60 seconds. Message size can reach jumbo frame size.
 - Linktrace: The LTM/LTR mechanism is an on-demand mechanism. It can detect the route of the data from any MEP to any other MEP in the MA/MEG. It can be used for the following purposes:
 - Adjacent relation retrieval – The ETH-LT function can be used to retrieve the adjacency relationship between an MEP and a remote MEP or MIP. The result of running ETH-LT function is a sequence of MIPs from the source MEP until the target MIP or MEP.
 - Fault localization – The ETH-LT function can be used for fault localization. When a fault occurs, the sequence of MIPs and/or MEP will probably be different from the expected sequence. The difference between the sequences provides information about the fault location.
 - AIS: AIS (defined in G.8013/Y.17310) is the Ethernet alarm indication signal function used to suppress alarms following detection of defect conditions at the server (sub) layer.

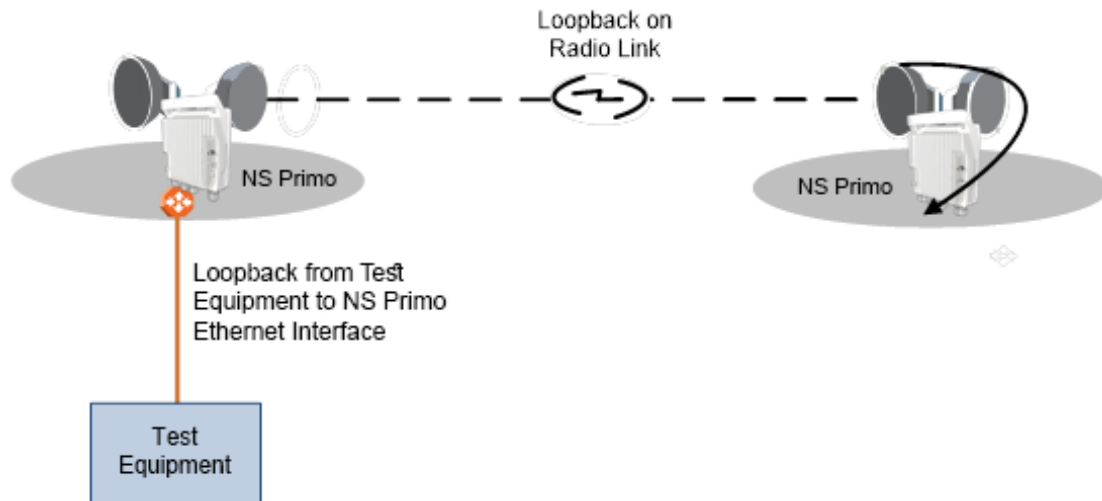
5.2.11.2. Ethernet Line Interface Loopback

NetStream Primo supports loopback testing for its radio interfaces. In addition, the Ethernet Line Interface Loopback feature provides the ability to run loopbacks over the link. When Ethernet loopback is enabled on an interface, the system loops back

all packets ingressing the interface. This enables loopbacks to be performed over the link from other points in the network.

For example, as shown in the figure below, a loopback can be performed from test equipment over the line to an Ethernet interface. A loopback can also be performed from the other side of the radio link.

Ethernet Line Interface Loopback – Application Examples



Ethernet loopbacks can be performed on any logical interface. This includes GbE interfaces, radio interfaces, and LAGS. Ethernet loopbacks cannot be performed on the management interface.

The following parameters can be configured for an Ethernet loopback:

- The interface can be configured to swap DA and SA MAC addresses during the loopback. This prevents Ethernet loops from occurring. It is recommended to enable MAC address swapping if MSTP or LLDP is enabled.
- Ethernet loopback has a configurable duration period of up to 2 ½ hours, but can be disabled manually before the duration period ends. Permanent loopback is not supported.

Ethernet loopbacks can be configured on more than one interface simultaneously. When an Ethernet loopback is active, network resiliency protocols (G.8032 and MSTP) will detect interface failure due to the failure to receive BPDUs.¹⁵

In a system using in-band management, Ethernet loopback activation on the remote side of the link causes loss of management to the remote unit. The duration period of the loopback should take this into account.

¹⁵ G.8032 and MSTP are planned for future release.

5.3. Synchronization

This section describes NetStream Primo's flexible synchronization solution that enables operators to configure a combination of synchronization techniques, based on the operator's network and migration strategy, including:

- PTP optimized transport, supporting IEEE 1588 and NTP, with guaranteed ultra-low PDV and support for ACM and narrow channels.
- Native Sync Distribution, for end-to-end distribution using GbE.
- SyncE PRC Pipe Regenerator mode, providing PRC grade (G.811) performance for pipe ("regenerator") applications.

This section includes:

- NetStream Primo Synchronization Solution
- Available Synchronization Interfaces
- Synchronous Ethernet (SyncE)
- IEEE-1588v2 PTP Optimized Transport
- SSM Support and Loop Prevention

Related topics:

- NTP Support

5.3.1. NetStream Primo Synchronization Solution

Netronics' synchronization solution ensures maximum flexibility by enabling the operator to select any combination of techniques suitable for the operator's network and migration strategy.

- SyncE PRC Pipe Regenerator mode
 - PRC grade (G.811) performance for pipe ("regenerator") applications
- SyncE node¹⁶
- PTP optimized transport
 - Supports a variety of protocols, such as IEEE-1588 and NTP
 - Supports IEEE-1588 Transparent Clock¹⁷
 - Guaranteed ultra-low PDV (<0.015 ms per hop)
 - Unique support for ACM and narrow channels

5.3.2. Available Synchronization Interfaces

Frequency signals can be taken by the system from a number of different interfaces (one reference at a time). The reference frequency may also be conveyed to external equipment through different interfaces.

¹⁶ SyncE node is planned for future release.

¹⁷ IEEE-1588 Transparent Clock is planned for future release.

Synchronization Interface Options

Available interfaces as frequency input (reference sync source)	Available interfaces as frequency output
<ul style="list-style-type: none"> • Radio carrier • GbE Ethernet interfaces 	<ul style="list-style-type: none"> • Radio carrier • GbE Ethernet interfaces

It is possible to configure up to eight synchronization sources in the system. At any given moment, only one of these sources is active; the clock is taken from the active source onto all other appropriately configured interfaces

5.3.3. Synchronous Ethernet (SyncE)

SyncE is standardized in ITU-T G.8261 and G.8262, and refers to a method whereby the frequency is delivered on the physical layer.

5.3.3.1. SyncE PRC Pipe Regenerator Mode

In SyncE PRC pipe regenerator mode, frequency is transported between two GbE interfaces through the radio link.

PRC pipe regenerator mode makes use of the fact that the system is acting as a simple link (so no distribution mechanism is necessary) in order to achieve the following:

- Improved frequency distribution performance, with PRC quality.
- Simplified configuration

In PRC pipe regenerator mode, frequency is taken from the incoming GbE Ethernet signal, and used as a reference for the radio frame. On the receiver side, the radio frame frequency is used as the reference signal for the outgoing Ethernet PHY.

Frequency distribution behaves in a different way for optical and electrical GbE interfaces, because of the way these interfaces are implemented:

- For optical interfaces, separate and independent frequencies are transported in each direction.
- For electrical interfaces, each PHY must act either as clock master or as clock slave in its own link. For this reason, frequency can only be distributed in one direction, determined by the user.

5.3.4. IEEE-1588v2 PTP Optimized Transport



Note

IEEE-1588v2 PTP Optimized Transport is planned for future release.

Precision Timing Protocol (PTP) refers to the distribution of frequency, phase, and absolute time information across an asynchronous frame switched network. PTP can use a variety of protocols to achieve timing distribution, including:

- IEEE-1588
- NTP


- RTP

IEEE-1588 PTP provides both frequency and phase (time) synchronization with the precision that is necessary in packet-switched mobile networks. With IEEE-1588 PTP, clocks distributed throughout the network are synchronized to sub-microsecond accuracy, suitable for mobile networks.

NetStream Primo supports PTP optimized transport, a message-based protocol that can be implemented across packet-based networks. IEEE-1588v2 provides both frequency and phase synchronization, and is designed to provide higher accuracy and precision, to the scale of nanoseconds, and up to 1.5 μ s.


IEEE-1588v2 PTP synchronization is based on a master-slave architecture in which the master and slave exchange PTP packets carrying clock information. The master is connected to a reference clock, and the slave synchronizes itself to the master.

IEEE-1588v2 PTP Optimized Transport – General Architecture



Accurate synchronization requires a determination of the propagation delay for PTP packets. Propagation delay is determined by a series of messages between the master and slave.

Calculating the Propagation Delay for PTP Packets



In this information exchange:

- 1 The master sends a Sync message to the slave and notes the time (t1) the message was sent.
- 2 The slave receives the Sync message and notes the time the message was received (t2).
- 3 The master conveys the t1 timestamp to the slave, in one of the following ways:
 - Embedding the t1 timestamp in the Sync message (requires L1 processing).
 - Embedding the t1 timestamp in a Follow-up message.
- 4 The slave sends a Delay_request message to the master and notes the time the message was sent (t3).
- 5 The master receives the Delay_request message and notes the time the message was received (t4).
- 6 The master conveys the t4 timestamp to the slave by embedding the t4 timestamp in a Delay_response message.

Based on this message exchange, the protocol calculates both the clock offset between the master and slave and the propagation delay, based on the following formulas:

$$\text{Offset} = [(t2 - t1) - (t4 - t3)]/2$$

$$\text{Propagation Delay} = [(t2 - t1) + (t4 - t3)]/2$$

The calculation is based on the assumption that packet delay is constant and that delays are the same in each direction. For information on the factors that may undermine these assumptions and how NetStream Primo's IEEE-1588v2 implementations mitigate these factors, see *Mitigating PDV* on page 128.

5.3.4.1. IEEE-1588v2 Benefits

IEEE-1588v2 provides packet-based synchronization that can transmit both frequency accuracy and phase information. This is essential for LTE applications, and provides a clear advantage over SyncE, which transmits frequency accuracy but not phase information.

Other IEEE-1588v2 benefits include:

- Fractional nanosecond precession.
- Meets strict LTE-A requirements for rigorous frequency and phase timing.
- Hardware time stamping of PTP packets.
- Standard protocol compatible with third-party equipment.
- Short frame and higher message rates.
- Supports unicast as well as multicast.
- Enables smooth transition from unsupported networks.
- Mitigates PDV issues by using Transparent Clock (see *Mitigating PDV* on page 128).
- Minimal consumption of bandwidth and processing power.
- Simple configuration.

5.3.4.2. Mitigating PDV

To get the most out of PTP and minimize PDV, NetStream Primo supports Transparent Clock.

PTP calculates path delay based on the assumption that packet delay is constant and that delays are the same in each direction. Delay variation invalidates this assumption. High PDV in wireless transport for synchronization over packet protocols, such as IEEE-1588, can dramatically affect the quality of the recovered clock. Slow variations are the most harmful, since in most cases it is more difficult for the receiver to average out such variations.

PDV can arise from both packet processing delay variation and radio link delay variation.

Packet processing delay variation can be caused by:

- Queuing Delay – Delay associated with incoming and outgoing packet buffer queuing.
- Head of Line Blocking – Occurs when a high priority frame, such as a frame that contains IEEE-1588 information, is forced to wait until a lower-priority frame that has already started to be transmitted completes its transmission.

- Store and Forward – Used to determine where to send individual packets. Incoming packets are stored in local memory while the MAC address table is searched and the packet's cyclic redundancy field is checked before the packet is sent out on the appropriate port. This process introduces variations in the time latency of packet forwarding due to packet size, flow control, MAC address table searches, and CRC calculations.

Radio link delay variation is caused by the effect of ACM, which enables dynamic modulation changes to accommodate radio path fading, typically due to weather changes. Lowering modulation reduces link capacity, causing traffic to accumulate in the buffers and producing transmission delay.



When bandwidth is reduced due to lowering of the ACM modulation point, it is essential that high priority traffic carrying IEEE-1588 packets be given the highest priority using NetStream Primo's enhanced QoS mechanism, so that this traffic will not be subject to delays or discards.

These factors can combine to produce a minimum and maximum delay, as follows:

- Minimum frame delay can occur when the link operates at a high modulation and no other frame has started transmission when the IEEE-1588 frame is ready for transmission.
- Maximum frame delay can occur when the link is operating at QPSK modulation and a large (e.g., 1518 bytes) frame has just started transmission when the IEEE-1588 frame is ready for transmission.


The worst case PDV is defined as the greatest difference between the minimum and maximum frame delays. The worst case can occur not just in the radio equipment itself but in every switch across the network.

To ensure minimal packet delay variation (PDV), NetStream Primo's synchronization solution includes 1588v2-compliant Transparent Clock. Transparent Clock provides the means to measure and adjust for delay variation, thereby ensuring low PDV.

5.3.4.3. Transparent Clock

NetStream Primo supports End-to-End Transparent Clock, which updates the time-interval correction field for the delay associated with individual packet transfers. End-to-End Transparent Clock is the most appropriate option for the Telecom industry.

A Transparent Clock node resides between a master and a slave node, and updates the timestamps of PTP packets passing from the master to the slave to compensate for delay, enabling the terminating clock in the slave node to remove the delay accrued in the Transparent Clock node. The Transparent Clock node is itself neither a master nor a slave node, but rather, serves as a bridge between master and slave nodes.

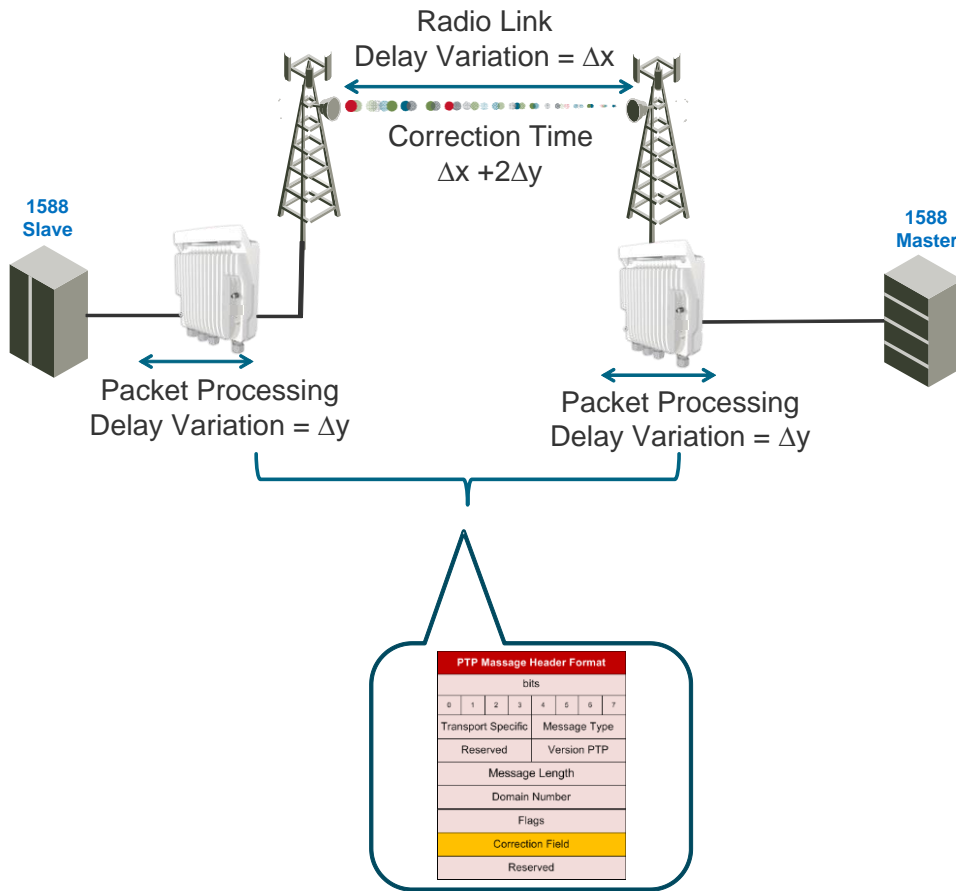
Transparent Clock – General Architecture

NetStream Primo uses 1588v2-compliant Transparent Clock to counter the effects of delay variation. Transparent Clock measures and adjusts for delay variation, enabling the NetStream Primo to guarantee ultra-low PDV.

The Transparent Clock algorithm forwards and adjusts the messages to reflect the residency time associated with the Sync, Follow_Up, and Delay_Request messages as they pass through the device. The delays are inserted in the 64-bit time-interval correction field.

As shown in the figure below, NetStream Primo measures and updates PTP messages based on both the radio link delay variation, and the packet processing delay variation that results from the network processor (switch operation).

Transparent Clock Delay Compensation



5.3.5. SSM Support and Loop Prevention



SSM support is planned for future release.

In order to provide topological resiliency for synchronization transfer, NetStream Primo implements the passing of SSM messages over the radio interfaces. SSM timing in NetStream Primo complies with ITU-T G.781.

In addition, the SSM mechanism provides reference source resiliency, since a network may have more than one source clock.

The following are the principles of operation:

- At all times, each source interface has a “quality status” which is determined as follows:
 - If quality is configured as fixed, then the quality status becomes “failure” upon interface failure (such as LOS, LOC, LOF, etc.).
 - If quality is automatic, then the quality is determined by the received SSMs or becomes “failure” upon interface failure (such as LOS, LOC, LOF, etc.).

- Each unit holds a parameter which indicates the quality of its reference clock. This is the quality of the current synchronization source interface.
- The reference source quality is transmitted through SSM messages to all relevant radio interfaces.
- Each unit determines the current active clock reference source interface:
 - The interface with the highest available quality is selected.
 - From among interfaces with identical quality, the interface with the highest priority is selected.
- In order to prevent loops, an SSM with quality “Do Not Use” is sent towards the active source interface

At any given moment, the system enables users to display:

- The current source interface quality.
- The current received SSM status for every source interface.
- The current node reference source quality.

As a reference, the following are the possible quality values (from highest to lowest):


- AUTOMATIC (available only in interfaces for which SSM support is implemented)
- G.811
- SSU-A
- SSU-B
- G.813/8262 - default
- DO NOT USE
- Failure (cannot be configured by user)

5.4. AES-256 Payload Encryption

NetStream Primo supports AES-256 payload encryption.

The Advanced Encryption Standard (AES) is defined in Federal Information Processing Standard Publication 197 (FIPS 197) for symmetric encryption. AES-256 is widely considered to be secure and efficient and is therefore broadly accepted as the standard for both government and industry applications.

AES-256 Encrypted Link



Note

The AES-256 payload encryption feature is a controlled item under applicable Export Laws. Please contact your Netronics representative to confirm that the encryption feature can be delivered.

5.4.1. AES Benefits

- Provides protection against eavesdropping and man-in-the-middle attacks on the radio
- Full encryption for all radio traffic
- Wire-speed, lowest latency encryption
- Eliminates the need for external encryption devices:
 - Cost effective encryption solution
 - Low Capex and operational costs; fast and simple deployment

5.4.2. NetStream Primo AES Implementation

In NetStream Primo, AES provides full payload encryption for all L1 radio traffic. It also encrypts control data passing through the radio link, such as the Link ID, ATPC data, and SSM messages. AES encryption operates on a point-to-point radio link level. AES is enabled and configured separately for each radio carrier.

NetStream Primo uses a dual-key encryption mechanism for AES.

- The user provides a master key. The master key can also be generated by the system upon user command. The master key is a 32-byte symmetric encryption key. The same master key must be manually configured on both ends of the encrypted link.
- The session key is a 32-byte symmetric encryption key used to encrypt the actual data. Each link uses two session keys, one for each direction.

For each direction, the session key is generated by the transmit side unit and propagated automatically to the other side of the link via a Key Exchange Protocol. The Key Exchange Protocol exchanges session keys by encrypting them with the master key, using the AES-256 encryption algorithm. Session keys are regenerated at user-configured intervals.

AES key generation is completely hitless, and has no effect on ACM operation.

Once AES encryption has been enabled on both sides of the link, the Key Exchange Protocol periodically verifies that both ends of the link have the same master key. If a mismatch is detected, an alarm is raised and traffic transmission is stopped for the mismatched carrier at both sides of the link. The link becomes non-valid and traffic stops being forwarded.

5.4.3. AES Interoperability

NS DIPLO/PRIMO's AES implementation is interoperable among NS DIPLO/PRIMO products that support AES. This means that for all NS DIPLO/PRIMO products that are otherwise interoperable with each other, AES can be used in links between two such products.

6. NetStream Primo Management

This chapter includes:

- Management Overview
- Automatic Network Topology Discovery with LLDP Protocol
- Management Communication Channels and Protocols
- Web-Based Element Management System (Web EMS)
- Command Line Interface (CLI)
- Configuration Management
- Software Management
- IPv6 Support
- In-Band Management
- Local Management
- Alarms
- NTP Support
- UTC Support
- System Security Features

6.1. Management Overview

The Netronics management solution is built on several layers of management:

- NEL – Network Element-level CLI
- EMS – HTTP web-based EMS
- NMS and SML –Netronics NMS platform

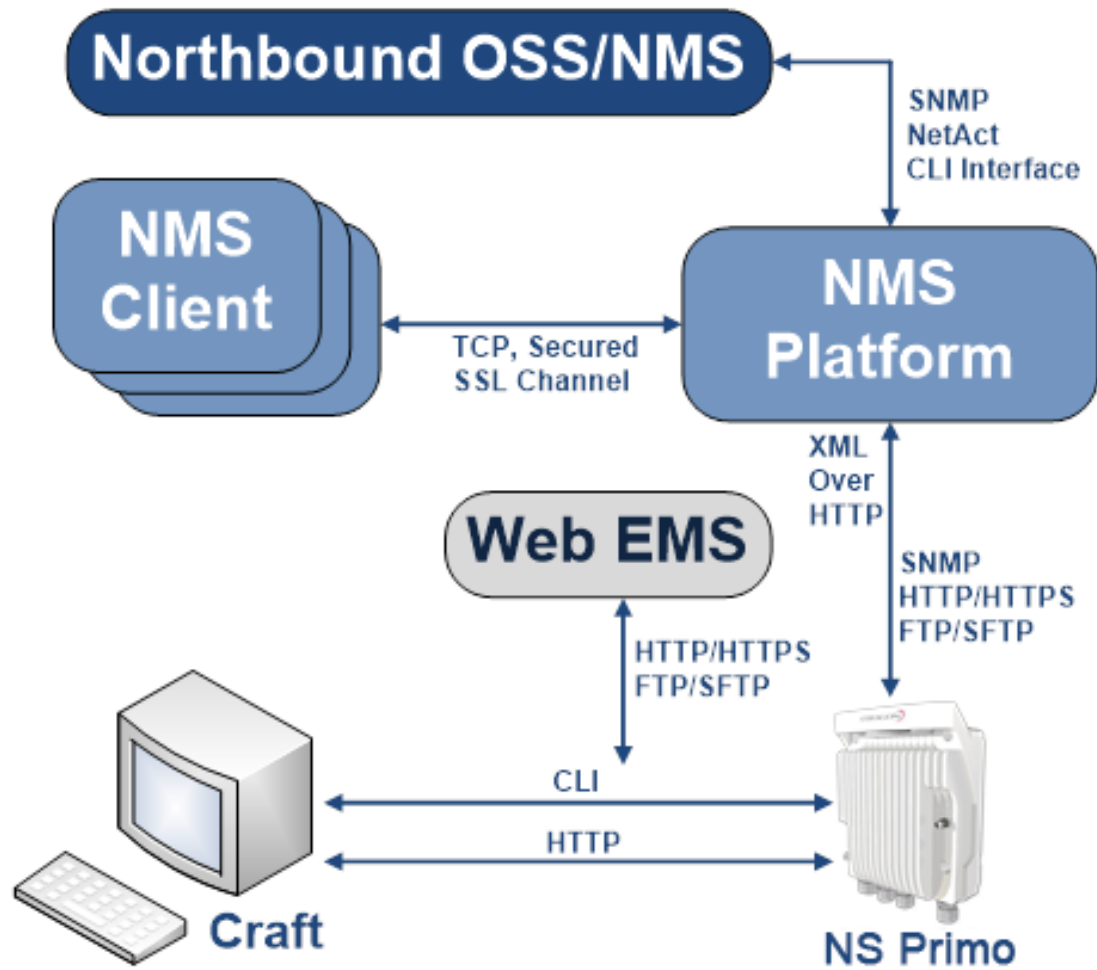
Every Old NetAxis and NS DIPLO/PRIMO network element includes an HTTP web-based element manager that enables the operator to perform element configuration, performance monitoring, remote diagnostics, alarm reports, and more.

In addition, Netronics provides an SNMP v1/v2c/v3 northbound interface on the NetStream Primo.

Netronics offers an NMS solution for providing centralized operation and maintenance capability for the complete range of network elements in a NetStream Primo system.

In addition, management, configuration, and maintenance tasks can be performed directly via the NetStream Primo Command Line Interface (CLI). The CLI can be used to perform configuration operations for NetStream Primo units, as well as to configure several NetStream Primo units in a single batch command.

Integrated NetStream Primo Management Tools



6.2. Automatic Network Topology Discovery with LLDP Protocol

NetStream Primo supports the Link Layer Discovery Protocol (LLDP), a vendor-neutral layer 2 protocol that can be used by a station attached to a specific LAN segment to advertise its identity and capabilities and to receive identity and capacity information from physically adjacent layer 2 peers. NetStream Primo's LLDP implementation is based on the IEEE 802.1AB – 2009 standard.

LLDP provides automatic network connectivity discovery by means of a port identity information exchange between each port and its peer. The port exchanges information with its peer and advertises this information to the NMS managing the unit. This enables the NMS to quickly identify changes to the network topology.

Enabling LLDP on NS DIPLO/PRIMO units enables the NMS to:

- Automatically detect the NS DIPLO/PRIMO unit neighboring the managed NS DIPLO/PRIMO unit, and determine the connectivity state between the two units.
- Automatically detect a third-party switch or router neighboring the managed NS DIPLO/PRIMO unit, and determine the connectivity state between the NS DIPLO/PRIMO unit and the switch or router.

6.3. Management Communication Channels and Protocols

Related Topics:

- Secure Communication Channels

Network Elements can be accessed locally via serial or Ethernet management interfaces, or remotely through the standard Ethernet LAN. The application layer is indifferent to the access channel used.

The NMS can be accessed through its GUI interface application, which may run locally or in a separate platform; it also has an SNMP-based northbound interface to communicate with other management systems.

Dedicated Management Ports

Port number	Protocol	Frame structure	Details
161	SNMP	UDP	Sends SNMP Requests to the network elements
162 Configurable	SNMP (traps)	UDP	Sends SNMP traps forwarding (optional)
25	SMTP (mail)	TCP	Sends the NMS reports and triggers by email (optional)
69	TFTP	UDP	Uploads/ downloads configuration files (optional)
80	HTTP	TCP	Manages devices
443	HTTPS	TCP	Manages devices (optional)
From 21 port to any remote port (>1023)	FTP Control Port	TCP	Downloads software and configuration files. (FTP Server responds to client's control port) (optional)
From Any port (>1023) to any remote port (>1023)	FTP Data Port	TCP	Downloads software and configuration files. The FTP server sends ACKs (and data) to client's data port. Optional FTP server random port range can be limited according to need (i.e., according to the number of parallel configuration uploads).

All remote system management is carried out through standard IP communications. Each NE behaves as a host with a single IP address.

The communications protocol used depends on the management channel being accessed.

As a baseline, these are the protocols in use:

- Standard HTTP for web-based management
- Standard telnet for CLI-based management
- The NMS uses a number of ports and protocols for different functions:

NMS Server Receiving Data Ports

Port number	Protocol	Frame structure	Details
162 Configurable	SNMP (traps)	UDP	Receive SNMP traps from network elements
69	TFTP	UDP	Downloads software and files (optional)
21	FTP Control Port	TCP	Downloads software and configuration files. (FTP client initiates a connection) (optional)

Port number	Protocol	Frame structure	Details
To any port (>1023) from any Port (>1023)	FTP Data Port	TCP	Downloads software and configuration files.(FTP Client initiates data connection to random port specified by server) (optional) FTP Server random port range can be limited according to needed configuration (number of parallel configuration uploads).
9205 Configurable	Propriety	TCP	User Actions Logger server (optional)

Web Sending Data Ports

Port number	Protocol	Frame structure	Details
80	HTTP	TCP	Manages device
443	HTTPS	TCP	Manages device (optional)

Web Receiving Data Ports

Port number	Protocol	Frame structure	Details
21	FTP	TCP	Downloads software files (optional)
Data port	FTP	TCP	Downloads software files (optional)

Additional Management Ports for NetStream Primo

Port number	Protocol	Frame structure	Details
23	telnet	TCP	Remote CLI access (optional)
22	SSH	TCP	Secure remote CLI access (optional)

6.4. Web-Based Element Management System (Web EMS)

The Web EMS is an HTTP web-based element manager that enables the operator to perform configuration operations and obtain statistical and performance information related to the system, including:

- **Configuration Management** – Enables you to view and define configuration data for the NetStream Primo system.
- **Fault Monitoring** – Enables you to view active alarms.
- **Performance Monitoring** – Enables you to view and clear performance monitoring values and counters.
- **Diagnostics and Maintenance** – Enables you to define and perform loopback tests, and software updates.
- **Security Configuration** – Enables you to configure NetStream Primo security features.
- **User Management** – Enables you to define users and user profiles.

A Web-Based EMS connection to the NetStream Primo can be opened using an HTTP Browser (Explorer or Mozilla Firefox). The Web EMS uses a graphical interface. Most system configurations and statuses are available via the Web EMS. However, some advanced configuration options are only available via CLI.

The Web EMS shows the actual unit configuration and provides easy access to any interface on the unit.

The Web EMS includes quick link configuration wizards that guide the user, step-by-step, through the creation of 1+0 links with Pipe services.

6.5. Command Line Interface (CLI)

A CLI connection to the NetStream Primo can be opened via telnet. All parameter configurations can be performed via CLI.



Note

Telnet access can be blocked by user configuration.

6.6. Configuration Management

The system configuration file consists of a set of all the configurable system parameters and their current values.

NetStream Primo configuration files can be imported and exported. This enables you to copy the system configuration to multiple NetStream Primo units.

System configuration files consist of a zip file that contains three components:

- A binary configuration file which is used by the system to restore the configuration.
- A text file which enables users to examine the system configuration in a readable format. The file includes the value of all system parameters at the time of creation of the backup file.
- An additional text file which enables users to write CLI scripts in order to make desired changes in the backed-up configuration. This file is executed by the system after restoring the configuration.¹⁸

The system provides three restore points to manage different configuration files. Each restore point contains a single configuration file. Files can be added to restore points by creating backups of the current system state or by importing them from an external server.



Note

In the Web EMS, these restore points are referred to as “file numbers.”

For example, a user may want to use one restore point to keep a last good configuration, another to import changes from an external server, and the third to store the current configuration.

Any of the restore points can be used to apply a configuration file to the system.

The user can determine whether or not to include security-related settings, such as users and user profiles, in the exported configuration file. By default, security settings are included.



Note

The option to enable or disable import and export of security parameters is planned for future release.

¹⁸

The option to edit the backup configuration is planned for future release.

6.7. Software Management

The NetStream Primo software installation and upgrade process includes the following steps:

- **Download** – The files required for the installation or upgrade are downloaded from a remote server.
- **Installation** – The files are installed in the appropriate modules and components of the NetStream Primo.
- **Reset** – The NetStream Primo is restarted in order to boot the new software and firmware versions.

NetStream Primo software and firmware releases are provided in a single bundle that includes software and firmware for all components supported by the system. When the user downloads a software bundle, the system verifies the validity of the bundle. The system also compares the files in the bundle to the files currently installed in the NetStream Primo and its components, so that only files that differ between the new version bundle and the current version in the system are actually downloaded. A message is displayed to the user for each file that is actually downloaded.

**Note**

When downloading an older version, all files in the bundle may be downloaded, including files that are already installed.

Software bundles can be downloaded via FTP, SFTP, HTTP, or HTTPS.

**Note**

Only FTP and SFTP downloads are supported in Release8.2. Support for the other protocols is planned for future release.

After the software download is complete, the user initiates the installation. A timer can be used to perform the installation after a defined time interval. The system performs an automatic reset after the installation.

6.7.1. Backup Software Version

**Note**

Backup software version support is planned for future release.

NetStream Primo maintains a backup copy of the software bundle. In the event that the working software version cannot be found, or the operating system fails to start properly, the system automatically boots from the backup version, and the previously active version becomes the backup version.

Users can also update the backup version manually. The Web EMS includes a field that indicates whether or not the active and backup software versions are identical.

6.8. IPv6 Support

NetStream Primo management communications can use both IPv4 and IPv6. The unit IP address for management can be configured in either or both formats.

Additionally, other management communications can utilize either IPv4 or IPv6. This includes:

- Software file downloads
- Configuration file import and export
- Trap forwarding
- Unit information file export (used primarily for maintenance and troubleshooting)

6.9. In-Band Management

NetStream Primo can optionally be managed In-Band, via its radio and Ethernet interfaces. This method of management eliminates the need for a dedicated management interface. For more information, refer to *Management Service (MNG)* on page 67.

6.10. Local Management

NetStream Primo includes an FE port for local management. This port (MGT) is enabled by default, and cannot be disabled.

6.11. Alarms

6.11.1. Configurable BER Threshold for Alarms and Traps

Users can configure alarm and trap generation in the event of Excessive BER and Signal Degrade BER above user-defined thresholds. Users have the option to configure whether or not excessive BER is propagated as a fault and considered a system event.

6.11.2. Alarms Editing

Users can change the description text (by appending extra text to the existing description) or the severity of any alarm in the system.

This is performed as follows:

- Each alarm in the system is identified by a unique name (see separate list of system alarms and events).
- The user can perform the following operations on any alarm:
 - View current description and severity
 - Define the text to be appended to the description and/or severity
 - Return the alarm to its default values
- The user can also return all alarms and events to their default values.

6.12. NTP Support

Related topics:

- Synchronization

NetStream Primo supports Network Time Protocol (NTP). NTP distributes Coordinated Universal Time (UTC) throughout the system, using a jitter buffer to neutralize the effects of variable latency.

NetStream Primo supports NTPv3 and NTPv4. NTPv4 provides interoperability with NTPv3 and with SNTP.

6.13. UTC Support

NetStream Primo uses the Coordinated Universal Time (UTC) standard for time and date configuration. UTC is a more updated and accurate method of date coordination than the earlier date standard, Greenwich Mean Time (GMT).

Every NetStream Primo unit holds the UTC offset and daylight savings time information for the location of the unit. Each management unit presenting the information (CLI and Web EMS) uses its own UTC offset to present the information in the correct time.

6.14. System Security Features

To guarantee proper performance and availability of a network as well as the data integrity of the traffic, it is imperative to protect it from all potential threats, both internal (misuse by operators and administrators) and external (attacks originating outside the network).

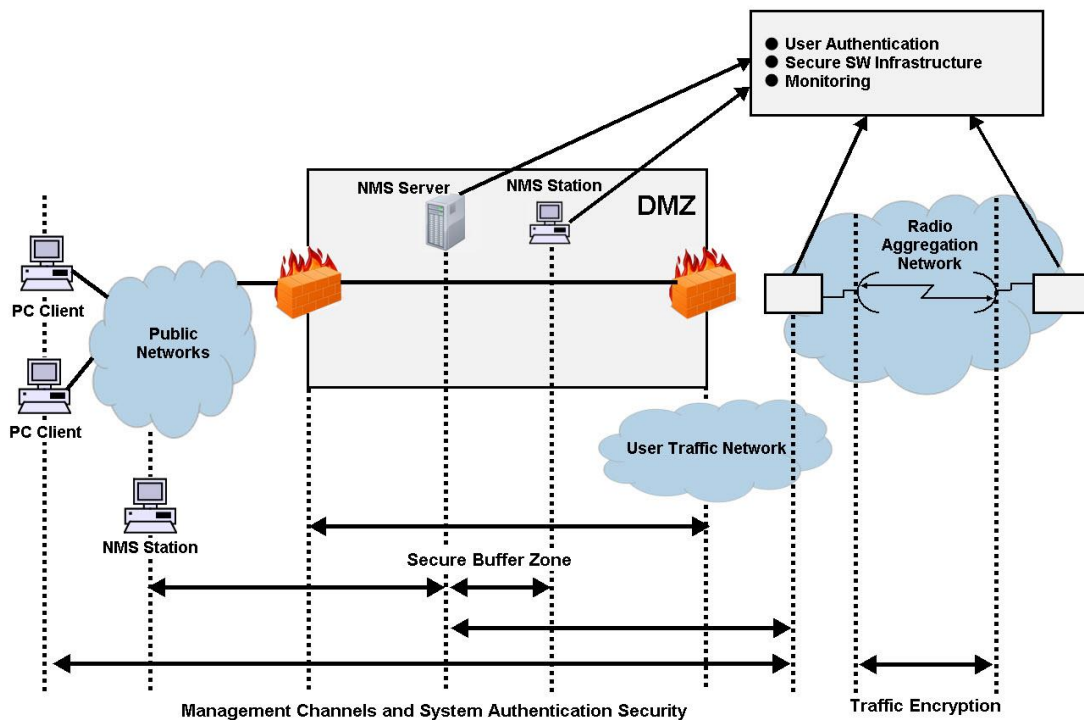
System security is based on making attacks difficult (in the sense that the effort required to carry them out is not worth the possible gain) by putting technical and operational barriers in every layer along the way, from the access outside the network, through the authentication process, up to every data link in the network.

6.14.1. Netronics' Layered Security Concept

Each layer protects against one or more threats. However, it is the combination of them that provides adequate protection to the network. In most cases, no single layer protection provides a complete solution to threats.

The layered security concept is presented in the following figure. Each layer presents the security features and the threats addressed by it. Unless stated otherwise, requirements refer to both network elements and the NMS.

Security Solution Architecture Concept



6.14.2. Defenses in Management Communication Channels

Since network equipment can be managed from any location, it is necessary to protect the communication channels' contents end to end.

These defenses are based on existing and proven cryptographic techniques and libraries, thus providing standard secure means to manage the network, with minimal impact on usability.

They provide defense at any point (including public networks and radio aggregation networks) of communications.

While these features are implemented in Netronics equipment, it is the responsibility of the operator to have the proper capabilities in any external devices used to manage the network.

In addition, inside Netronics networking equipment it is possible to control physical channels used for management. This can greatly help deal with all sorts of DoS attacks.

Operators can use secure channels instead or in addition to the existing management channels:

- SNMPv3 for all SNMP-based protocols for both NEs and NMS
- HTTPS for access to the NE's web server
- SSH-2 for all CLI access SFTP for all software and configuration download between NMS and NEs

All protocols run with secure settings using strong encryption techniques. Unencrypted modes are not allowed, and algorithms used must meet modern and client standards.

Users are allowed to disable all insecure channels.

In the network elements, the bandwidth of physical channels transporting management communications is limited to the appropriate magnitude, in particular, channels carrying management frames to the CPU.

Attack types addressed

- Tempering with management flows
- Management traffic analysis
- Unauthorized software installation
- Attacks on protocols (by providing secrecy and integrity to messages)
- Traffic interfaces eavesdropping (by making it harder to change configuration)
- DoS through flooding

6.14.3. Defenses in User and System Authentication Procedures

6.14.3.1. User Configuration and User Profiles

User configuration is based on the Role-Based Access Control (RBAC) model. According to the RBAC model, permissions to perform certain operations are assigned to specific roles. Users are assigned to particular roles, and through those role assignments acquire the permissions to perform particular system functions.

In the NetStream Primo GUI, these roles are called user profiles. Up to 50 user profiles can be configured. Each profile contains a set of privilege levels per functionality group, and defines the management protocols (access channels) that can be used to access the system by users to whom the user profile is assigned.

The system parameters are divided into the following functional groups:

- Security

- Management
- Radio
- Ethernet
- Synchronization

A user profile defines the permitted access level per functionality group. For each functionality group, the access level is defined separately for read and write operations. The following access levels can be assigned:

- **None** – No access to this functional group.
- **Normal** – The user has access to parameters that require basic knowledge about the functional group.
- **Advance** – The user has access to parameters that require advanced knowledge about the functional group, as well as parameters that have a significant impact on the system as a whole, such as restoring the configuration to factory default settings.

6.14.3.2. User Identification

NetStream Primo supports the following user identification features:

- Configurable inactivity time-out for automatically closing unused management channels
- Optional password strength enforcement. When password strength enforcement is enabled; passwords must comply with the following rules:
 - Password must be at least eight characters long.
 - Password must include at least three of the following categories: lower-case characters, upper-case characters, digits, and special characters.
 - When calculating the number of character categories, upper-case letters used as the first character and digits used as the last character of a password are not counted.
 - The password cannot have been used within the user's previous five passwords.
- Users can be prompted to change passwords after a configurable amount of time (password aging).
- Users can be blocked for a configurable time period after a configurable number of unsuccessful login attempts.
- Users can be configured to expire at a certain date
- Mandatory change of password at first time login can be enabled and disabled upon user configuration. It is enabled by default.

6.14.3.3. Remote Authentication

**Note**

Remote authorization is planned for future release.

Certificate-based strong standard encryption techniques are used for remote authentication. Users may choose to use this feature or not for all secure communication channels.

Since different operators may have different certificate-based authentication policies (for example, issuing its own certificates vs. using an external CA or allowing the NMS system to be a CA), NEs and NMS software provide the tools required for operators to enforce their policy and create certificates according to their established processes.

Server authentication capabilities are provided.

6.14.3.4. RADIUS Support

The RADIUS protocol provides centralized user management services. NetStream Primo supports RADIUS server and provides a RADIUS client for authentication and authorization.

RADIUS can be enabled or disabled. When RADIUS is enabled, a user attempting to log into the system from any access channel (CLI, WEB, NMS) is not authenticated locally. Instead, the user's credentials are sent to a centralized standard RADIUS server which indicates to the NetStream Primo whether the user is known, and which privilege is to be given to the user. RADIUS uses the same user attributes and privileges defined for the user locally.

**Note**

When using RADIUS for user authentication and authorization, the access channel limitations defined per user profile are not applicable. This means that when a user is authorized via RADIUS, the user can access the unit via any available access channel.

RADIUS login works as follows:

- If the RADIUS server is reachable, the system expects authorization to be received from the server:
 - The server sends the appropriate user privilege to the NetStream Primo, or notifies the NetStream Primo that the user was rejected.
 - If rejected, the user will be unable to log in. Otherwise, the user will log in with the appropriate privilege and will continue to operate normally.
- If the RADIUS server is unavailable, the NetStream Primo will attempt to authenticate the user locally, according to the existing list of defined users.



Local login authentication is provided in order to enable users to manage the system in the event that RADIUS server is unavailable. This requires previous definition of users in the system. If the user is only defined in the RADIUS server, the user will be unable to login locally in case the RADIUS server is unavailable.

In order to support NetStream Primo - specific privilege levels, the vendor-specific field is used. Netronics' IANA number for this field is 2281.

The following RADIUS servers are supported:

- FreeRADIUS
- RADIUS on Windows Server (IAS)
 - Windows Server 2008

6.14.4. Secure Communication Channels

NetStream Primo supports a variety of standard encryption protocols and algorithms, as described in the following sections.

6.14.4.1. SSH (Secured Shell)

SSH protocol can be used as a secured alternative to Telnet. In NetStream Primo:

- SSHv2 is supported.
- SSH protocol will always be operational. Admin users can choose whether to disable Telnet protocol, which is enabled by default. Server authentication is based on NetStream Primo's public key.
- RSA and DSA key types are supported.
- Supported Encryptions: aes128-cbc, 3des-cbc, blowfish-cbc, cast128-cbc, arcfour128, arcfour256, arcfour, aes192-cbc, aes256-cbc, aes128-ctr, aes192-ctr, aes256-ctr.
- MAC (Message Authentication Code): SHA-1-96 (MAC length = 96 bits, key length = 160 bit). Supported MAC: hmac-md5, hmac-sha1, hmac-ripemd160, hmac-sha1-96, hmac-md5-96'
- The server authenticates the user based on user name and password. The number of failed authentication attempts is not limited.
- The server timeout for authentication is 10 minutes. This value cannot be changed.

6.14.4.2. HTTPS (Hypertext Transfer Protocol Secure)

HTTPS combines the Hypertext Transfer protocol with the SSL/TLS protocol to provide encrypted communication and secure identification of a network web server. NetStream Primo enables administrators to configure secure access via HTTPS protocol.

6.14.4.3. SFTP (Secure FTP)

SFTP can be used for the following operations:

- Configuration upload and download,
- Uploading unit information
- Uploading a public key
- Downloading certificate files
- Downloading software

6.14.4.4. Creation of Certificate Signing Request (CSR) File

In order to create a digital certificate for the NE, a Certificate Signing Request (CSR) file should be created by the NE. The CSR contains information that will be included in the NE's certificate such as the organization name, common name (domain name), locality, and country. It also contains the public key that will be included in the certificate. Certificate authority (CA) will use the CSR to create the desired certificate for the NE.

While creating the CSR file, the user will be asked to input the following parameters that should be known to the operator who applies the command:

- **Common name** – The identify name of the element in the network (e.g., the IP address). The common name can be a network IP or the FQDN of the element.
- **Organization** – The legal name of the organization.
- **Organizational Unit** - The division of the organization handling the certificate.
- **City/Locality** - The city where the organization is located.
- **State/County/Region** - The state/region where the organization is located.
- **Country** - The two-letter ISO code for the country where the organization is location.
- **Email address** - An email address used to contact the organization.

6.14.4.5. SNMP

NetStream Primo supports SNMP v1, V2c, and v3. The default community string in NMS and the SNMP agent in the embedded SW are disabled. Users are allowed to set community strings for access to network elements.

NetStream Primo supports the following MIBs:

- RFC-1213 (MIB II)
- RMON MIB
- Netronics (proprietary) MIB.

Access to all network elements in a node is provided by making use of the community and context fields in SNMPv1 and SNMPv2c/SNMPv3, respectively.

For additional information:

- NetStream Primo MIB Reference, DOC- 00036524.

6.14.4.6. Server Authentication (SSL / SLLv3)**Note**

SSL and SLLv3 support is planned for future release.

- All protocols making use of SSL (such as HTTPS) use SLLv3 and support X.509 certificates-based server authentication.
- Users with type of “administrator” or above can perform the following server (network element) authentication operations for certificates handling:
 - Generate server key pairs (private + public)
 - Export public key (as a file to a user-specified address)
 - Install third-party certificates
 - The Admin user is responsible for obtaining a valid certificate.
 - Load a server RSA key pair that was generated externally for use by protocols making use of SSL.
- Non-SSL protocols using asymmetric encryption, such as SSH and SFTP, can make use of public-key based authentication.
 - Users can load trusted public keys for this purpose.

6.14.4.7. Encryption**Note**

Support for encryption is planned for future release.

- Encryption algorithms for secure management protocols include:
 - Symmetric key algorithms: 128-bit AES
 - Asymmetric key algorithms: 1024-bit RSA

6.14.5. Security Log

The security log is an internal system file which records all changes performed to any security feature, as well as all security related events.

**Note**

In order to read the security log, the user must upload the log to his or her server.

The security log file has the following attributes:

- The file is of a “cyclic” nature (fixed size, newest events overwrite oldest).
- The log can only be read by users with "admin" or above privilege.
- The contents of the log file are cryptographically protected and digitally signed.
 - In the event of an attempt to modify the file, an alarm will be raised.
- Users may not overwrite, delete, or modify the log file.

The security log records:

- Changes in security configuration
 - Carrying out “security configuration copy to mate”
 - Management channels time-out
 - Password aging time
 - Number of unsuccessful login attempts for user suspension
 - Warning banner change
 - Adding/deleting of users
 - Password changed
 - SNMP enable/disable
 - SNMP version used (v1/v3) change
 - SNMPv3 parameters change
 - Security mode
 - Authentication algorithm
 - User
 - Password
 - SNMPv1 parameters change
 - Read community
 - Write community
 - Trap community for any manager
 - HTTP/HTTPS change
 - FTP/SFTP change
 - Telnet and web interface enable/disable
 - FTP enable/disable
 - Loading certificates
 - RADIUS server
 - Radius enable/disable
 - Remote logging enable/disable (for security and configuration logs)
 - Syslog server address change (for security and configuration logs)

- System clock change
- NTP enable/disable
- Security events
- Successful and unsuccessful login attempts
- N consecutive unsuccessful login attempts (blocking)
- Configuration change failure due to insufficient permissions
- SNMPv3/PV authentication failures
- User logout
- User account expired

For each recorded event the following information is available:

- User ID
 - Communication channel (WEB, terminal, telnet/SSH, SNMP, NMS, etc.)
 - IP address, if applicable
2. Date and time

7. Standards and Certifications

This chapter includes:

- Supported Ethernet Standards
- MEF Certifications for Ethernet Services

7.1. Supported Ethernet Standards

Supported Ethernet Standards

Standard	Description
802.3	10base-T
802.3u	100base-T
802.3ab	1000base-T
802.3z	1000base-X
802.3ac	Ethernet VLANs
802.1Q	Virtual LAN (VLAN)
802.1p	Class of service
802.1ad	Provider bridges (QinQ)
802.3ad	Link aggregation
Auto MDI/MDIX for 1000baseT	
RFC 1349	IPv4 TOS
RFC 2474	IPv4 DSCP
RFC 2460	IPv6 Traffic Classes

7.2. MEF Certifications for Ethernet Services

Supported MEF Specifications

Specification	Description
MEF-2	Requirements and Framework for Ethernet Service Protection
MEF-6.1	Metro Ethernet Services Definitions Phase 2
MEF-8	Implementation Agreement for the Emulation of PDH Circuits over Metro Ethernet Networks
MEF-10.3	Ethernet Services Attributes Phase 3
MEF 22.1	Mobile Backhaul Implementation Agreement Phase 2
MEF-30.1	Service OAM Fault Management Implementation Agreement Phase 2
MEF-35	Service OAM Performance Monitoring Implementation Agreement

MEF Certifications

Certification	Description
CE 2.0	Second generation Carrier Ethernet certification
MEF-18	Abstract Test Suite for Circuit Emulation Services
MEF-9	Abstract Test Suite for Ethernet Services at the UNI. Certified for all service types (EPL, EVPL & E-LAN). This is a first generation certification. It is fully covered as part of CE2.0)
MEF-14	Abstract Test Suite for Traffic Management Phase 1. Certified for all service types (EPL, EVPL & E-LAN). This is a first generation certification. It is fully covered as part of CE2.0)

8. Specifications

This chapter includes:

- General Radio Specifications
- Frequency Accuracy
- Radio Capacity Specifications
- Transmit Power Specifications
- Receiver Threshold Specifications
- Frequency Bands
- Mediation Device Losses
- Ethernet Latency Specifications
- Interface Specifications
- Carrier Ethernet Functionality
- Synchronization Functionality
- Network Management, Diagnostics, Status, and Alarms
- Mechanical Specifications
- Standards Compliance
- Environmental Specifications
- Antenna Specifications
- Power Input Specifications
- Power Consumption Specifications
- Power Connection Options
- PoE Injector Specifications
- Cable Specifications

Related Topics:

- Standards and Certifications



Note

All specifications are subject to change without prior notification.

8.1. General Radio Specifications

Frequency (GHz)	Operating Frequency Range (GHz)	Tx/Rx Spacing (MHz)	Standard
6L,6H	5.85-6.45, 6.4-7.1	252.04, 240, 266, 300, 340, 160, 170, 500	EN 302 217
7,8	7.1-7.9, 7.7-8.5	154, 119, 161, 168, 182, 196, 208, 245, 250, 266, 300,310, 311.32, 500, 530	EN 302 217
10	10.0-10.7	91, 168,350, 550	EN 302 217
11	10.7-11.7	490, 520, 530	EN 302 217
13	12.75-13.3	266	EN 302 217
15	14.4-15.35	315, 420, 475, 644, 490, 728	EN 302 217
18	17.7-19.7	1010, 1120, 1008, 1560	EN 302 217
23	21.2-23.65	1008, 1200, 1232	EN 302 217
24UL	24.0-24.25	Customer-defined	EN 302 217
26	24.2-26.5	800, 1008	EN 302 217
28	27.35-29.5	350, 450, 490, 1008	EN 302 217
32	31.8-33.4	812	EN 302 217
38	37-40	1000, 1260, 700	EN 302 217
42	40.55-43.45	1500	EN 302 217
Standards		ETSI, ITU-R, CEPT	
Frequency Source		Synthesizer	
System Configurations		1+0, 1+1HSB, 2+OSP, DP	
RF Channel Selection		Via EMS/NMS	
Tx Range (Manual/ATPC)		Up to 25dB dynamic range	

8.2. Frequency Accuracy

NetStream Primo provides frequency accuracy of ± 4 ppm¹⁹.

¹⁹ Over temperature.

8.3. Radio Capacity Specifications

Each table in this section includes ranges of capacity specifications according to frame size, with ranges given for no Header De-Duplication, Layer-2 Header De-Duplication, and LTE-optimized Header De-Duplication.



Note

Ethernet capacity depends on average frame size.

The capacity figures for LTE packets encapsulated inside GTP tunnels with IPv4/UDP encapsulation and double VLAN tagging (QinQ). Capacity for IPv6 encapsulation is higher. A Capacity Calculator tool is available for different encapsulations and flow types.

ACAP and ACCP represent compliance with different ETSI mask requirements. ACCP represents compliance with more stringent interference requirements.

8.3.1. 3.5 MHz Channel Bandwidth (ACCP) (MRMC 1523)

Profile	Modulation	Minimum required capacity activation key	Ethernet throughput		
			No Header De-Duplication	L2 Compression	Header De-Duplication
0	QPSK	10	3-4	3-5	4-13
1	16 QAM	10	8-10	8-12	9-32
2	32 QAM	10	11-14	11-17	12-43
3	64 QAM	50	14-17	14-21	15-54
4	128 QAM	50	17-21	17-25	18-65
5	256 QAM	50	19-24	20-29	20-74

8.3.2. 7 MHz Channel Bandwidth (ACCP) (MRMC 1508)

Profile	Modulation	Minimum required capacity activation key	Ethernet throughput		
			No Compression	L2 Compression	Header De-Duplication
0	QPSK	10	8-10	9-13	9-32
1	8 PSK	10	13-16	13-19	13-48
2	16 QAM	50	18-22	18-27	19-69
3	32 QAM	50	24-30	24-36	26-92
4	64 QAM	50	30-37	30-44	32-114
5	128 QAM	50	36-44	36-53	38-137
6	256 QAM	50	42-51	42-61	44-158

Profile	Modulation	Minimum required capacity activation key	Ethernet throughput		
			No Compression	L2 Compression	Header De-Duplication
7	512 QAM	50	45-54	45-66	47-169
8	1024 QAM (Strong FEC)	50	48-58	48-71	50-182
9	1024 QAM (Light FEC)	50	51-62	51-75	53-194

8.3.3. 14MHz Channel Bandwidth (ACCP) (MRMC 1509)

Profile	Modulation	Minimum required capacity activation key	Ethernet throughput		
			No Compression	L2 Compression	Header De-Duplication
0	QPSK	50	19-24	20-29	20-74
1	8 PSK	50	29-36	30-43	31-112
2	16 QAM	50	40-49	41-60	42-153
3	32 QAM	50	53-65	54-79	56-203
4	64 QAM	50	66-80	66-97	69-249
5	128 QAM	100	79-97	80-117	83-301
6	256 QAM	100	90-110	91-134	95-344
7	512 QAM	100	100-122	101-147	105-380
8	1024 QAM (Strong FEC)	100	106-129	106-156	111-402
9	1024 QAM (Light FEC)	100	112-137	113-166	118-426

8.3.4. 28 MHz Channel Bandwidth (ACCP) (MRMC 1504)

Profile	Modulation	Minimum required capacity activation key	Ethernet Throughput		
			No Compression	L2 Compression	Header De-Duplication
0	QPSK	50	40-49	41-59	42-153
1	8 PSK	50	60-74	61-89	63-229
2	16 QAM	100	82-101	83-122	86-313
3	32 QAM	100	108-132	109-160	114-412
4	64 QAM	150	134-163	135-197	140-508
5	128 QAM	150	161-196	162-237	169-612
6	256 QAM	200	183-224	184-270	192-696
7	512 QAM	200	202-247	203-298	212-769
8	1024 QAM (Strong FEC)	225	215-262	216-317	225-817

Profile	Modulation	Minimum required capacity activation key	Ethernet Throughput		
			No Compression	L2 Compression	Header De-Duplication
9	1024 QAM (Light FEC)	225	228-279	230-337	239-868
10	2048 QAM	250	245-299	246-361	257-931

8.3.5. 28 MHz Channel Bandwidth (ACAP) (MRMC 1505)

Profile	Modulation	Minimum required capacity activation key	Ethernet Throughput		
			No Compression	L2 Compression	Header De-Duplication
0	QPSK	50	43-52	43-63	45-162
1	8 PSK	50	62-76	63-92	65-236
2	16 QAM	100	87-107	88-129	92-332
3	32 QAM	100	115-140	116-170	121-437
4	64 QAM	150	141-173	143-209	149-538
5	128 QAM	150	170-208	172-252	179-648
6	256 QAM	200	196-239	197-289	206-745
7	512 QAM	200	209-255	210-308	219-794
8	1024 QAM (Strong FEC)	225	228-278	229-336	239-866
9	1024 QAM (Light FEC)	225	241-295	243-356	253-917
10	2048 QAM	250	263-321	265-389	276-1000

8.3.6. 40 MHz Channel Bandwidth (ACCP) (MRMC 1507)

Profile	Modulation	Minimum required capacity activation key	Ethernet Throughput		
			No Compression	L2 Compression	Header De-Duplication
0	QPSK	50	58-71	58-85	61-220
1	8 PSK	100	86-105	87-127	90-328
2	16 QAM	100	117-143	118-173	123-446
3	32 QAM	150	154-189	156-228	162-588
4	64 QAM	200	190-232	191-280	199-722
5	128 QAM	225	229-280	231-339	241-873
6	256 QAM	250	247-302	249-365	259-939
7	512 QAM	300	270-330	272-399	284-1000
8	1024 QAM (Strong FEC)	300	306-375	309-453	322-1000

Profile	Modulation	Minimum required capacity activation key	Ethernet Throughput		
			No Compression	L2 Compression	Header De-Duplication
9	1024 QAM (Light FEC)	300	325-398	328-481	342-1000
10	2048 QAM	350	352-430	355-520	370-1000

8.3.7. 56 MHz Channel Bandwidth (ACCP) (MRMC 1502)

Profile	Modulation	Minimum required capacity activation key	Ethernet Throughput		
			No Compression	L2 Compression	Header De-Duplication
0	QPSK	100	83-101	83-122	87-314
1	8 PSK	100	123-150	124-182	129-468
2	16 QAM	150	167-205	169-247	176-637
3	32 QAM	225	220-269	222-325	231-838
4	64 QAM	300	270-331	273-400	284-1000
5	128 QAM	300	327-400	329-483	343-1000
6	256 QAM	400	374-457	377-553	393-1000
7	512 QAM	400	406-496	409-600	426-1000
8	1024 QAM (Strong FEC)	450	441-540	445-652	464-1000
9	1024 QAM (Light FEC)	450	469-573	472-693	492-1000
10	2048 QAM	500	508-621	512-751	534-1000

8.3.8. 56 MHz Channel Bandwidth (ACAP) (MRMC 1506)

Profile	Modulation	Minimum required capacity activation key	Ethernet Throughput		
			No Compression	L2 Compression	Header De-Duplication
0	QPSK	100	87-106	88-128	91-331
1	8 PSK	150	127-155	128-187	133-482
2	16 QAM	200	176-215	177-260	185-670
3	32 QAM	225	232-283	233-342	243-881
4	64 QAM	300	284-348	286-420	299-1000
5	128 QAM	350	344-420	346-508	361-1000
6	256 QAM	400	397-485	400-586	416-1000
7	512 QAM	450	426-521	430-630	448-1000
8	1024 QAM (Strong FEC)	450	464-567	467-685	487-1000

Profile	Modulation	Minimum required capacity activation key	Ethernet Throughput		
			No Compression	L2 Compression	Header De-Duplication
9	1024 QAM (Light FEC)	500	493-602	497-728	517-1000
10	2048 QAM	500	534-653	538-789	561-1000

8.3.9. 80 MHz Channel Bandwidth (ACCP) (MRMC 1501)

Profile	Modulation	Minimum required capacity activation key	Ethernet Throughput		
			No Compression	L2 Compression	Header De-Duplication
0	QPSK	100	114-140	115-169	120-435
1	8 PSK	150	162-198	164-240	170-618
2	16 QAM	225	231-283	233-342	243-880
3	32 QAM	300	304-371	306-449	319-1000
4	64 QAM	400	371-454	374-549	390-1000
5	128 QAM	450	439-536	442-649	461-1000
6	256 QAM	500	505-618	509-747	531-1000
7	512 QAM	500	555-679	560-821	583-1000
8	1024 QAM	500	604-738	609-892	634-1000

8.4. Transmit Power Specifications



Nominal TX power is subject to change until the relevant frequency band is formally released. See the frequency rollout plan.

The values listed in this section are typical. Actual values may differ in either direction by up to 1dB.

NetStream Primo Standard Power

Modulation	6 GHz	7 GHz	8 GHz	10-11 GHz	13-15 GHz	18 GHz	23 GHz	24GHz UL ²⁰	26 GHz	28,32,38 GHz	42 GHz
QPSK	25	25	25	23	24	22	20	0	21	18	15
8 PSK	25	25	25	23	24	22	20	0	21	18	15
16 QAM	25	24	24	23	23	21	20	0	20	17	14
32 QAM	24	23	23	22	22	20	20	0	19	16	13
64 QAM	24	23	23	22	22	20	20	0	19	16	13
128 QAM	24	23	23	22	22	20	20	0	19	16	13
256 QAM	24	23	21	22	20	20	18	0	17	14	11
512 QAM	22	21	21	21	20	18	18	0	17	14	11
1024 QAM	22	21	21	20	20	18	17	0	16	13	10
2048 QAM	20	19	19	18	18	16	16	0	15	12	9

NetStream Primo High Power

Modulation	6 GHz	7 GHz	8 GHz	10-11 GHz
QPSK	28	28	28	26
8 PSK	28	28	28	26
16 QAM	28	27	27	26
32 QAM	27	26	26	25
64 QAM	27	26	26	25
128 QAM	27	26	26	25
256 QAM	27	26	24	25
512 QAM	25	24	24	24
1024 QAM	25	24	24	23
2048 QAM	23	22	22	21

²⁰

For 1ft ant or lower.

Customers in countries following EC Directive 2006/771/EC (incl. amendments) must observe the 100mW EIRP obligation by adjusting transmit power according to antenna gain and RF line losses.

8.5. Receiver Threshold Specifications



Note

The values listed in this section are typical. Tolerance range is -1dB/+ 2dB.

Profile	Modulation	Channel Spacing	Frequency (GHz)														
			6	7	8	10	11	13	15	18	23	24 ²¹	26	28-31	32	38	42
0	QPSK	3.5 MHz	-96.5	-96.0	-96.0	-95.5	-96.5	-95.5	-94.5	-96.0	-95.0	-94.5	-94.5	-94.5	-94.0	-94.0	-93.5
1	16 QAM		-90.0	-89.0	-89.0	-89.0	-89.5	-88.5	-88.0	-89.0	-88.0	-87.5	-88.0	-87.5	-87.5	-87.0	-86.5
2	32 QAM		-86.5	-85.5	-85.5	-85.5	-86.0	-85.0	-84.5	-85.5	-84.5	-84.0	-84.5	-84.0	-84.0	-83.5	-83.0
3	64 QAM		-83.0	-82.5	-82.5	-82.0	-83.0	-82.0	-81.0	-82.5	-81.5	-81.0	-81.0	-81.0	-80.5	-80.5	-80.0
4	128 QAM		-79.5	-79.0	-79.0	-78.5	-79.5	-78.5	-77.5	-79.0	-78.0	-77.5	-77.5	-77.5	-77.0	-77.0	-76.5
5	256 QAM		-76.5	-75.5	-75.5	-75.5	-76.5	-75.0	-74.5	-75.5	-75.0	-74.5	-74.5	-74.0	-74.0	-73.5	-73.0
0	QPSK	7 MHz	-93.5	-93.0	-93.0	-92.5	-93.5	-92.5	-91.5	-93.0	-92.0	-91.5	-91.5	-91.5	-91.0	-91.0	-90.5
1	8 PSK		-87.5	-87.0	-87.0	-86.5	-87.5	-86.5	-85.5	-87.0	-86.0	-85.5	-85.5	-85.5	-85.0	-85.0	-84.5
2	16 QAM		-87.0	-86.5	-86.5	-86.0	-87.0	-86.0	-85.0	-86.5	-85.5	-85.0	-85.0	-85.0	-84.5	-84.5	-84.0
3	32 QAM		-83.5	-83.0	-83.0	-82.5	-83.5	-82.5	-81.5	-83.0	-82.0	-81.5	-81.5	-81.5	-81.0	-81.0	-80.5
4	64 QAM		-80.5	-80.0	-80.0	-79.5	-80.5	-79.5	-78.5	-80.0	-79.0	-78.5	-78.5	-78.5	-78.0	-78.0	-77.5
5	128 QAM		-77.5	-76.5	-76.5	-76.5	-77.5	-76.0	-75.5	-76.5	-76.0	-75.5	-75.5	-75.0	-75.0	-74.5	-74.0
6	256 QAM		-74.0	-73.5	-73.5	-73.0	-74.0	-73.0	-72.0	-73.5	-72.5	-72.0	-72.0	-72.0	-71.5	-71.5	-71.0

²¹ Customers in countries following EC Directive 2006/771/EC (incl. amendments) must observe the 100mW EIRP obligation by adjusting transmit power according to antenna gain and RF line losses.

Profile	Modulation	Channel Spacing	Frequency (GHz)															
			6	7	8	10	11	13	15	18	23	24 ²¹	26	28-31	32	38	42	
7	512 QAM	14 MHz	-72.0	-71.5	-71.5	-71.0	-72.0	-71.0	-70.0	-71.5	-70.5	-70.0	-70.0	-70.0	-69.5	-69.5	-69.0	
8	1024 QAM (strong FEC)		-68.5	-68.0	-68.0	-67.5	-68.5	-67.5	-66.5	-68.0	-67.0	-66.5	-66.5	-66.5	-66.0	-66.0	-65.5	
9	1024 QAM (light FEC)		-68.0	-67.0	-67.0	-67.0	-67.5	-66.5	-66.0	-67.0	-66.0	-65.5	-66.0	-65.5	-65.5	-65.0	-64.5	
0	QPSK	14 MHz	-90.5	-90.0	-90.0	-89.5	-90.5	-89.5	-88.5	-90.0	-89.0	-88.5	-88.5	-88.5	-88.0	-88.0	-87.5	
1	8 PSK		-84.5	-84.0	-84.0	-83.5	-84.5	-83.5	-82.5	-84.0	-83.0	-82.5	-82.5	-82.5	-82.0	-82.0	-81.5	
2	16 QAM		-83.5	-83.0	-83.0	-82.5	-83.5	-82.5	-81.5	-83.0	-82.0	-81.5	-81.5	-81.5	-81.0	-81.0	-80.5	
3	32 QAM		-80.5	-79.5	-79.5	-79.5	-80.5	-79.0	-78.5	-79.5	-79.0	-78.5	-78.5	-78.0	-78.0	-77.5	-77.0	
4	64 QAM		-77.5	-76.5	-76.5	-76.5	-77.5	-76.0	-75.5	-76.5	-76.0	-75.5	-75.5	-75.0	-75.0	-74.5	-74.0	
5	128 QAM		-74.0	-73.5	-73.5	-73.0	-74.0	-73.0	-72.0	-73.5	-72.5	-72.0	-72.0	-72.0	-71.5	-71.5	-71.0	
6	256 QAM		-71.5	-70.5	-70.5	-70.5	-71.0	-70.0	-69.5	-70.5	-69.5	-69.0	-69.5	-69.0	-69.0	-68.5	-68.0	
7	512 QAM		-68.5	-68.0	-68.0	-67.5	-68.5	-67.5	-66.5	-68.0	-67.0	-66.5	-66.5	-66.5	-66.0	-66.0	-65.5	
8	1024 QAM (strong FEC)		-65.5	-65.0	-65.0	-64.5	-65.5	-64.5	-63.5	-65.0	-64.0	-63.5	-63.5	-63.5	-63.0	-63.0	-62.5	
9	1024 QAM (light FEC)		-65.0	-64.0	-64.0	-64.0	-65.0	-63.5	-63.0	-64.0	-63.5	-63.0	-63.0	-62.5	-62.5	-62.0	-61.5	
0	QPSK		28 MHz ACCP	-87.5	-87.0	-87.0	-86.5	-87.5	-86.5	-85.5	-87.0	-86.0	-85.5	-85.5	-85.5	-85.0	-85.0	-84.5
1	8 PSK			-83.0	-82.5	-82.5	-82.0	-83.0	-82.0	-81.0	-82.5	-81.5	-81.0	-81.0	-81.0	-80.5	-80.5	-80.0
2	16 QAM	-81.0		-80.5	-80.5	-80.0	-81.0	-79.5	-79.0	-80.5	-79.5	-79.0	-79.0	-79.0	-78.5	-78.0	-78.0	
3	32 QAM	-77.5		-77.0	-77.0	-76.5	-77.5	-76.0	-75.5	-77.0	-76.0	-75.5	-75.5	-75.5	-75.0	-74.5	-74.5	
4	64 QAM	-74.5		-74.0	-74.0	-73.5	-74.5	-73.0	-72.5	-74.0	-73.0	-72.5	-72.5	-72.5	-72.0	-71.5	-71.5	
5	128 QAM	-71.5		-70.5	-70.5	-70.5	-71.0	-70.0	-69.5	-70.5	-69.5	-69.0	-69.5	-69.0	-69.0	-68.5	-68.0	
6	256 QAM	-68.5		-67.5	-67.5	-67.5	-68.0	-67.0	-66.5	-67.5	-66.5	-66.0	-66.5	-66.0	-66.0	-65.5	-65.0	
7	512 QAM	-66.0		-65.0	-65.0	-65.0	-66.0	-64.5	-64.0	-65.0	-64.5	-64.0	-64.0	-63.5	-63.5	-63.0	-62.5	
8	1024 QAM (strong FEC)	-63.0		-62.5	-62.5	-62.0	-63.0	-61.5	-61.0	-62.5	-61.5	-61.0	-61.0	-61.0	-60.5	-60.0	-60.0	

Profile	Modulation	Channel Spacing	Frequency (GHz)														
			6	7	8	10	11	13	15	18	23	24 ²¹	26	28-31	32	38	42
9	1024 QAM (light FEC)	28 MHz ACAP	-62.0	-61.5	-61.5	-61.0	-62.0	-60.5	-60.0	-61.5	-60.5	-60.0	-60.0	-60.0	-59.5	-59.0	-59.0
10	2048 QAM		-58.5	-58.0	-58.0	-57.5	-58.5	-57.0	-56.5	-58.0	-57.0	-56.5	-56.5	-56.5	-56.0	-55.5	-55.5
0	QPSK		-87.5	-87.0	-87.0	-86.5	-87.5	-86.0	-85.5	-87.0	-86.0	-85.5	-85.5	-85.5	-85.0	-84.5	-84.5
1	8 PSK		-82.5	-81.5	-81.5	-81.5	-82.5	-81.0	-80.5	-81.5	-81.0	-80.5	-80.5	-80.0	-80.0	-79.5	-79.0
2	16 QAM		-81.0	-80.0	-80.0	-80.0	-80.5	-79.5	-79.0	-80.0	-79.0	-78.5	-79.0	-78.5	-78.5	-78.0	-77.5
3	32 QAM		-77.0	-76.5	-76.5	-76.0	-77.0	-76.0	-75.0	-76.5	-75.5	-75.0	-75.0	-75.0	-74.5	-74.5	-74.0
4	64 QAM		-74.5	-73.5	-73.5	-73.5	-74.0	-73.0	-72.5	-73.5	-72.5	-72.0	-72.5	-72.0	-72.0	-71.5	-71.0
5	128 QAM		-71.0	-70.5	-70.5	-70.0	-71.0	-70.0	-69.0	-70.5	-69.5	-69.0	-69.0	-69.0	-68.5	-68.5	-68.0
6	256 QAM		-68.0	-67.5	-67.5	-67.0	-68.0	-67.0	-66.0	-67.5	-66.5	-66.0	-66.0	-66.0	-65.5	-65.5	-65.0
7	512 QAM		-66.0	-65.5	-65.5	-65.0	-66.0	-64.5	-64.0	-65.5	-64.5	-64.0	-64.0	-64.0	-63.5	-63.0	-63.0
8	1024 QAM (strong FEC)	-63.0	-62.0	-62.0	-62.0	-62.5	-61.5	-61.0	-62.0	-61.0	-60.5	-61.0	-60.5	-60.5	-60.0	-59.5	
9	1024 QAM (light FEC)	-62.0	-61.0	-61.0	-61.0	-62.0	-60.5	-60.0	-61.0	-60.5	-60.0	-60.0	-59.5	-59.5	-59.0	-58.5	
10	2048 QAM	-58.0	-57.5	-57.5	-57.0	-58.0	-56.5	-56.0	-57.5	-56.5	-56.0	-56.0	-56.0	-55.5	-55.0	-55.0	
0	QPSK	40 MHz	-86.0	-85.5	-85.5	-85.0	-86.0	-85.0	-84.0	-85.5	-84.5	-84.0	-84.0	-84.0	-83.5	-83.5	-83.0
1	8 PSK		-81.0	-80.5	-80.5	-80.0	-81.0	-79.5	-79.0	-80.5	-79.5	-79.0	-79.0	-79.0	-78.5	-78.0	-78.0
2	16 QAM		-79.5	-79.0	-79.0	-78.5	-79.5	-78.0	-77.5	-79.0	-78.0	-77.5	-77.5	-77.5	-77.0	-76.5	-76.5
3	32 QAM		-76.0	-75.0	-75.0	-75.0	-75.5	-74.5	-74.0	-75.0	-74.0	-73.5	-74.0	-73.5	-73.5	-73.0	-72.5
4	64 QAM		-73.0	-72.0	-72.0	-72.0	-73.0	-71.5	-71.0	-72.0	-71.5	-71.0	-71.0	-70.5	-70.5	-70.0	-69.5
5	128 QAM		-70.0	-69.0	-69.0	-69.0	-70.0	-68.5	-68.0	-69.0	-68.5	-68.0	-68.0	-67.5	-67.5	-67.0	-66.5
6	256 QAM		-67.0	-66.0	-66.0	-66.0	-66.5	-65.5	-65.0	-66.0	-65.0	-64.5	-65.0	-64.5	-64.5	-64.0	-63.5
7	512 QAM		-64.0	-63.5	-63.5	-63.0	-64.0	-62.5	-62.0	-63.5	-62.5	-62.0	-62.0	-62.0	-61.5	-61.0	-61.0
8	1024 QAM (strong FEC)	-61.5	-61.0	-61.0	-60.5	-61.5	-60.0	-59.5	-61.0	-60.0	-59.5	-59.5	-59.5	-59.0	-58.5	-58.5	

Profile	Modulation	Channel Spacing	Frequency (GHz)														
			6	7	8	10	11	13	15	18	23	24 ²¹	26	28-31	32	38	42
9	1024 QAM (light FEC)		-60.5	-60.0	-60.0	-59.5	-60.5	-59.5	-58.5	-60.0	-59.0	-58.5	-58.5	-58.5	-58.0	-58.0	-57.5
10	2048 QAM		-58.0	-57.0	-57.0	-57.0	-58.0	-56.5	-56.0	-57.0	-56.5	-56.0	-56.0	-55.5	-55.5	-55.0	-54.5
0	QPSK	56 MHz ACCP	-84.0	-83.5	-83.5	-83.0	-84.0	-83.0	-82.0	-83.5	-82.5	-82.0	-82.0	-82.0	-81.5	-81.5	-81.0
1	8 PSK		-80.0	-79.5	-79.5	-79.0	-80.0	-79.0	-78.0	-79.5	-78.5	-78.0	-78.0	-78.0	-77.5	-77.5	-77.0
2	16 QAM		-77.5	-77.0	-77.0	-76.5	-77.5	-76.5	-75.5	-77.0	-76.0	-75.5	-75.5	-75.5	-75.0	-75.0	-74.5
3	32 QAM		-74.5	-73.5	-73.5	-73.5	-74.0	-73.0	-72.5	-73.5	-72.5	-72.0	-72.5	-72.0	-72.0	-71.5	-71.0
4	64 QAM		-71.0	-70.5	-70.5	-70.0	-71.0	-70.0	-69.0	-70.5	-69.5	-69.0	-69.0	-69.0	-68.5	-68.5	-68.0
5	128 QAM		-68.5	-67.5	-67.5	-67.5	-68.0	-67.0	-66.5	-67.5	-66.5	-66.0	-66.5	-66.0	-66.0	-65.5	-65.0
6	256 QAM		-65.0	-64.5	-64.5	-64.0	-65.0	-64.0	-63.0	-64.5	-63.5	-63.0	-63.0	-63.0	-62.5	-62.5	-62.0
7	512 QAM		-63.0	-62.5	-62.5	-62.0	-63.0	-61.5	-61.0	-62.5	-61.5	-61.0	-61.0	-61.0	-60.5	-60.0	-60.0
8	1024 QAM (strong FEC)		-59.5	-59.0	-59.0	-58.5	-59.5	-58.5	-57.5	-59.0	-58.0	-57.5	-57.5	-57.5	-57.0	-57.0	-56.5
9	1024 QAM (light FEC)		-58.5	-58.0	-58.0	-57.5	-58.5	-57.5	-56.5	-58.0	-57.0	-56.5	-56.5	-56.5	-56.0	-56.0	-55.5
10	2048 QAM		-54.0	-53.5	-53.5	-53.0	-54.0	-53.0	-52.0	-53.5	-52.5	-52.0	-52.0	-52.0	-51.5	-51.5	-51.0
0	QPSK		56 MHz ACAP	-84.5	-84.0	-84.0	-83.5	-84.5	-83.0	-82.5	-84.0	-83.0	-82.5	-82.5	-82.5	-82.0	-81.5
1	8 PSK	-80.0		-79.0	-79.0	-79.0	-79.5	-78.5	-78.0	-79.0	-78.0	-77.5	-78.0	-77.5	-77.5	-77.0	-76.5
2	16 QAM	-77.5		-77.0	-77.0	-76.5	-77.5	-76.0	-75.5	-77.0	-76.0	-75.5	-75.5	-75.5	-75.0	-74.5	-74.5
3	32 QAM	-74.0		-73.0	-73.0	-73.0	-73.5	-72.5	-72.0	-73.0	-72.0	-71.5	-72.0	-71.5	-71.5	-71.0	-70.5
4	64 QAM	-70.5		-70.0	-70.0	-69.5	-70.5	-69.5	-68.5	-70.0	-69.0	-68.5	-68.5	-68.5	-68.0	-68.0	-67.5
5	128 QAM	-68.0		-67.0	-67.0	-67.0	-67.5	-66.5	-66.0	-67.0	-66.0	-65.5	-66.0	-65.5	-65.5	-65.0	-64.5
6	256 QAM	-64.5		-64.0	-64.0	-63.5	-64.5	-63.5	-62.5	-64.0	-63.0	-62.5	-62.5	-62.5	-62.0	-62.0	-61.5
7	512 QAM	-62.5		-62.0	-62.0	-61.5	-62.5	-61.5	-60.5	-62.0	-61.0	-60.5	-60.5	-60.5	-60.0	-60.0	-59.5
8	1024 QAM (strong FEC)	-59.0	-58.5	-58.5	-58.0	-59.0	-58.0	-57.0	-58.5	-57.5	-57.0	-57.0	-57.0	-56.5	-56.5	-56.0	

Profile	Modulation	Channel Spacing	Frequency (GHz)														
			6	7	8	10	11	13	15	18	23	24 ²¹	26	28-31	32	38	42
9	1024 QAM (light FEC)		-58.0	-57.5	-57.5	-57.0	-58.0	-57.0	-56.0	-57.5	-56.5	-56.0	-56.0	-56.0	-55.5	-55.5	-55.0
10	2048 QAM		-55.5	-54.5	-54.5	-54.5	-55.0	-54.0	-53.5	-54.5	-53.5	-53.0	-53.5	-53.0	-53.0	-52.5	-52.0

Profile	Modulation	Channel Spacing	Frequency (GHz)	
			6H	11
0	QPSK	80 MHz	-83.5	-83.0
1	8 PSK		-78.5	-78.5
2	16 QAM		-76.5	-76.5
3	32 QAM		-73.0	-72.5
4	64 QAM		-70.0	-70.0
5	128 QAM		-67.5	-67.0
6	256 QAM		-64.5	-64.5
7	512 QAM		-61.5	-61.5
8	1024 QAM		-59.0	-58.5

8.5.1. Overload Thresholds

- For modulations up to and including 1024 QAM (strong FEC): -20dBm
- For modulations of 1024 (light FEC) and 2048 QAM: -25dBm

8.6. Frequency Bands

Frequency Band	TX Range	RX Range	Tx/Rx Spacing
6L GHz	6332.5-6393	5972-6093	300A
	5972-6093	6332.5-6393	
	6191.5-6306.5	5925.5-6040.5	266A
	5925.5-6040.5	6191.5-6306.5	
	6303.5-6418.5	6037.5-6152.5	
	6037.5-6152.5	6303.5-6418.5	
	6245-6290.5	5939.5-6030.5	260A
	5939.5-6030.5	6245-6290.5	
	6365-6410.5	6059.5-6150.5	
	6059.5-6150.5	6365-6410.5	
	6226.89-6286.865	5914.875-6034.825	252B
	5914.875-6034.825	6226.89-6286.865	
	6345.49-6405.465	6033.475-6153.425	
	6033.475-6153.425	6345.49-6405.465	
	6179.415-6304.015	5927.375-6051.975	252A
	5927.375-6051.975	6179.415-6304.015	
	6238.715-6363.315	5986.675-6111.275	
	5986.675-6111.275	6238.715-6363.315	
	6298.015-6422.615	6045.975-6170.575	
	6045.975-6170.575	6298.015-6422.615	
6235-6290.5	5939.5-6050.5	240A	
5939.5-6050.5	6235-6290.5		
6355-6410.5	6059.5-6170.5		
6059.5-6170.5	6355-6410.5		
6H GHz	6920-7080	6420-6580	500A
	6420-6580	6924.5-7075.5	
	7032.5-7091.5	6692.5-6751.5	340C
	6692.5-6751.5	7032.5-7091.5	
	6764.5-6915.5	6424.5-6575.5	340B
	6424.5-6575.5	6764.5-6915.5	
	6924.5-7075.5	6584.5-6735.5	
	6584.5-6735.5	6924.5-7075.5	
	6781-6939	6441-6599	340A
	6441-6599	6781-6939	
	6941-7099	6601-6759	
	6601-6759	6941-7099	

Frequency Band	TX Range	RX Range	Tx/Rx Spacing
	6707.5-6772.5	6537.5-6612.5	160A
	6537.5-6612.5	6707.5-6772.5	
	6767.5-6832.5	6607.5-6672.5	
	6607.5-6672.5	6767.5-6832.5	
	6827.5-6872.5	6667.5-6712.5	
	6667.5-6712.5	6827.5-6872.5	
7 GHz	7783.5-7898.5	7538.5-7653.5	245A
	7538.5-7653.5	7783.5-7898.5	
	7301.5-7388.5	7105.5-7192.5	196A
	7105.5-7192.5	7301.5-7388.5	
	7357.5-7444.5	7161.5-7248.5	
	7161.5-7248.5	7357.5-7444.5	182A
	7440.5-7499.5	7622.5-7681.5	
	7678.5-7737.5	7496.5-7555.5	
	7496.5-7555.5	7678.5-7737.5	168C
	7580.5-7639.5	7412.5-7471.5	
	7412.5-7471.5	7580.5-7639.5	
	7608.5-7667.5	7440.5-7499.5	
	7440.5-7499.5	7608.5-7667.5	
	7664.5-7723.5	7496.5-7555.5	
	7496.5-7555.5	7664.5-7723.5	168B
	7609.5-7668.5	7441.5-7500.5	
	7441.5-7500.5	7609.5-7668.5	
	7637.5-7696.5	7469.5-7528.5	
	7469.5-7528.5	7637.5-7696.5	
	7693.5-7752.5	7525.5-7584.5	
	7525.5-7584.5	7693.5-7752.5	168A
	7273.5-7332.5	7105.5-7164.5	
	7105.5-7164.5	7273.5-7332.5	
	7301.5-7360.5	7133.5-7192.5	
	7133.5-7192.5	7301.5-7360.5	
	7357.5-7416.5	7189.5-7248.5	
	7189.5-7248.5	7357.5-7416.5	161P
	7280.5-7339.5	7119.5-7178.5	
	7119.5-7178.5	7280.5-7339.5	
	7308.5-7367.5	7147.5-7206.5	
7147.5-7206.5	7308.5-7367.5		
7336.5-7395.5	7175.5-7234.5		
7175.5-7234.5	7336.5-7395.5		
7364.5-7423.5	7203.5-7262.5		

Frequency Band	TX Range	RX Range	Tx/Rx Spacing
	7203.5-7262.5	7364.5-7423.5	
	7597.5-7622.5	7436.5-7461.5	161O
	7436.5-7461.5	7597.5-7622.5	
	7681.5-7706.5	7520.5-7545.5	
	7520.5-7545.5	7681.5-7706.5	
	7587.5-7646.5	7426.5-7485.5	
	7426.5-7485.5	7587.5-7646.5	161M
	7615.5-7674.5	7454.5-7513.5	
	7454.5-7513.5	7615.5-7674.5	
	7643.5-7702.5	7482.5-7541.5	161K
	7482.5-7541.5	7643.5-7702.5	
	7671.5-7730.5	7510.5-7569.5	
	7510.5-7569.5	7671.5-7730.5	
	7580.5-7639.5	7419.5-7478.5	161J
	7419.5-7478.5	7580.5-7639.5	
	7608.5-7667.5	7447.5-7506.5	
	7447.5-7506.5	7608.5-7667.5	
	7664.5-7723.5	7503.5-7562.5	
	7503.5-7562.5	7664.5-7723.5	
	7580.5-7639.5	7419.5-7478.5	161I
	7419.5-7478.5	7580.5-7639.5	
	7608.5-7667.5	7447.5-7506.5	
	7447.5-7506.5	7608.5-7667.5	
	7664.5-7723.5	7503.5-7562.5	
	7503.5-7562.5	7664.5-7723.5	
	7273.5-7353.5	7112.5-7192.5	161F
	7112.5-7192.5	7273.5-7353.5	
	7322.5-7402.5	7161.5-7241.5	
	7161.5-7241.5	7322.5-7402.5	
	7573.5-7653.5	7412.5-7492.5	
	7412.5-7492.5	7573.5-7653.5	
	7622.5-7702.5	7461.5-7541.5	
	7461.5-7541.5	7622.5-7702.5	
	7709-7768	7548-7607	161D
	7548-7607	7709-7768	
	7737-7796	7576-7635	
	7576-7635	7737-7796	
	7765-7824	7604-7663	
	7604-7663	7765-7824	
	7793-7852	7632-7691	

Frequency Band	TX Range	RX Range	Tx/Rx Spacing
	7632-7691	7793-7852	
	7584-7643	7423-7482	161C
	7423-7482	7584-7643	
	7612-7671	7451-7510	
	7451-7510	7612-7671	
	7640-7699	7479-7538	
	7479-7538	7640-7699	
	7668-7727	7507-7566	
	7507-7566	7668-7727	
	7409-7468	7248-7307	
	7248-7307	7409-7468	
	7437-7496	7276-7335	
	7276-7335	7437-7496	
	7465-7524	7304-7363	
	7304-7363	7465-7524	
	7493-7552	7332-7391	
	7332-7391	7493-7552	
	7284-7343	7123-7182	161A
	7123-7182	7284-7343	
	7312-7371	7151-7210	
	7151-7210	7312-7371	
	7340-7399	7179-7238	
	7179-7238	7340-7399	
	7368-7427	7207-7266	
	7207-7266	7368-7427	
	7280.5-7339.5	7126.5-7185.5	154C
	7126.5-7185.5	7280.5-7339.5	
	7308.5-7367.5	7154.5-7213.5	
	7154.5-7213.5	7308.5-7367.5	
	7336.5-7395.5	7182.5-7241.5	
	7182.5-7241.5	7336.5-7395.5	
	7364.5-7423.5	7210.5-7269.5	
	7210.5-7269.5	7364.5-7423.5	
	7594.5-7653.5	7440.5-7499.5	154B
	7440.5-7499.5	7594.5-7653.5	
	7622.5-7681.5	7468.5-7527.5	
	7468.5-7527.5	7622.5-7681.5	
	7678.5-7737.5	7524.5-7583.5	
	7524.5-7583.5	7678.5-7737.5	
	7580.5-7639.5	7426.5-7485.5	154A

Frequency Band	TX Range	RX Range	Tx/Rx Spacing
	7426.5-7485.5	7580.5-7639.5	
	7608.5-7667.5	7454.5-7513.5	
	7454.5-7513.5	7608.5-7667.5	
	7636.5-7695.5	7482.5-7541.5	
	7482.5-7541.5	7636.5-7695.5	
	7664.5-7723.5	7510.5-7569.5	
	7510.5-7569.5	7664.5-7723.5	
8 GHz	8396.5-8455.5	8277.5-8336.5	119A
	8277.5-8336.5	8396.5-8455.5	
	8438.5 – 8497.5	8319.5 – 8378.5	
	8319.5 – 8378.5	8438.5 – 8497.5	
	8274.5-8305.5	7744.5-7775.5	530A
	7744.5-7775.5	8274.5-8305.5	
	8304.5-8395.5	7804.5-7895.5	500A
	7804.5-7895.5	8304.5-8395.5	
	8023-8186.32	7711.68-7875	311C-J
	7711.68-7875	8023-8186.32	
	8028.695-8148.645	7717.375-7837.325	311B
	7717.375-7837.325	8028.695-8148.645	
	8147.295-8267.245	7835.975-7955.925	
	7835.975-7955.925	8147.295-8267.245	
	8043.52-8163.47	7732.2-7852.15	311A
	7732.2-7852.15	8043.52-8163.47	
	8162.12-8282.07	7850.8-7970.75	
	7850.8-7970.75	8162.12-8282.07	
	8212-8302	7902-7992	310D
	7902-7992	8212-8302	
	8240-8330	7930-8020	
	7930-8020	8240-8330	
	8296-8386	7986-8076	
	7986-8076	8296-8386	
8212-8302	7902-7992	310C	
7902-7992	8212-8302		
8240-8330	7930-8020		
7930-8020	8240-8330		
8296-8386	7986-8076		
7986-8076	8296-8386		
8380-8470	8070-8160		
8070-8160	8380-8470		
8408-8498	8098-8188		

Frequency Band	TX Range	RX Range	Tx/Rx Spacing
	8098-8188	8408-8498	
	8039.5-8150.5	7729.5-7840.5	310A
	7729.5-7840.5	8039.5-8150.5	
	8159.5-8270.5	7849.5-7960.5	
	7849.5-7960.5	8159.5-8270.5	
	8024.5-8145.5	7724.5-7845.5	
	7724.5-7845.5	8024.5-8145.5	
	8144.5-8265.5	7844.5-7965.5	
	7844.5-7965.5	8144.5-8265.5	
	8302.5-8389.5	8036.5-8123.5	266C
	8036.5-8123.5	8302.5-8389.5	
	8190.5-8277.5	7924.5-8011.5	266B
	7924.5-8011.5	8190.5-8277.5	
	8176.5-8291.5	7910.5-8025.5	266A
	7910.5-8025.5	8176.5-8291.5	
	8288.5-8403.5	8022.5-8137.5	
	8022.5-8137.5	8288.5-8403.5	
	8226.52-8287.52	7974.5-8035.5	252A
	7974.5-8035.5	8226.52-8287.52	
	8270.5-8349.5	8020.5-8099.5	250A
8016.5-8156.5	7733-7873	283A	
7733-7873	8016.5-8156.5		
8128.5-8268.5	7845-7985		
7845-7985	8128.5-8268.5		
10 GHz			
	10501-10563	10333-10395	168A
	10333-10395	10501-10563	
	10529-10591	10361-10423	
	10361-10423	10529-10591	
	10585-10647	10417-10479	
	10417-10479	10585-10647	
	10501-10647	10151-10297	350A
	10151-10297	10501-10647	
	10498-10652	10148-10302	350B
	10148-10302	10498-10652	
	10561-10707	10011-10157	550A
	10011-10157	10561-10707	
	10701-10847	10151-10297	
10151-10297	10701-10847		
10590-10622	10499-10531	91A	

Frequency Band	TX Range	RX Range	Tx/Rx Spacing
	10499-10531	10590-10622	
	10618-10649	10527-10558	
	10527-10558	10618-10649	
	10646-10677	10555-10586	
	10555-10586	10646-10677	
11 GHz	11425-11725	10915-11207	All
	10915-11207	11425-11725	
	11185-11485	10700-10950	
	10695-10955	11185-11485	
13 GHz	13002-13141	12747-12866	266
	12747-12866	13002-13141	
	13127-13246	12858-12990	
	12858-12990	13127-13246	
	12807-12919	13073-13185	266A
	13073-13185	12807-12919	200
	12700-12775	12900-13000	
	12900-13000	12700-12775	
	12750-12825	12950-13050	
	12950-13050	12750-12825	
	12800-12870	13000-13100	
	13000-13100	12800-12870	
	12850-12925	13050-13150	
	13050-13150	12850-12925	
15 GHz	15110-15348	14620-14858	490
	14620-14858	15110-15348	
	14887-15117	14397-14627	
	14397-14627	14887-15117	
	15144-15341	14500-14697	644
	14500-14697	15144-15341	
	14975-15135	14500-14660	475
	14500-14660	14975-15135	
	15135-15295	14660-14820	
	14660-14820	15135-15295	
	14921-15145	14501-14725	420
	14501-14725	14921-15145	
	15117-15341	14697-14921	
	14697-14921	15117-15341	

Frequency Band	TX Range	RX Range	Tx/Rx Spacing
	14963-15075	14648-14760	315
	14648-14760	14963-15075	
	15047-15159	14732-14844	
	14732-14844	15047-15159	
	15229-15375	14500-14647	728
	14500-14647	15229-15375	
18 GHz	19160-19700	18126-18690	1010
	18126-18690	19160-19700	
	18710-19220	17700-18200	
	17700-18200	18710-19220	
	19260-19700	17700-18140	1560
	17700-18140	19260-19700	
23 GHz	23000-23600	22000-22600	1008
	22000-22600	23000-23600	
	22400-23000	21200-21800	1232 /1200
	21200-21800	22400-23000	
	23000-23600	21800-22400	
	21800-22400	23000-23600	
24UL GHz ²²	24000 - 24250	24000 - 24250	All
26 GHz	25530-26030	24520-25030	1008
	24520-25030	25530-26030	
	25980-26480	24970-25480	
	24970-25480	25980-26480	
	25266-25350	24466-24550	800
	24466-24550	25266-25350	
	25050-25250	24250-24450	
	24250-24450	25050-25250	
28 GHz	28150-28350	27700-27900	450
	27700-27900	28150-28350	
	27950-28150	27500-27700	
	27500-27700	27950-28150	
	28050-28200	27700-27850	350
	27700-27850	28050-28200	

²² Customers in countries following EC Directive 2006/771/EC (incl. amendments) must observe the 100mW EIRP obligation by adjusting transmit power according to antenna gain and RF line losses.

Frequency Band	TX Range	RX Range	Tx/Rx Spacing
	27960-28110	27610-27760	
	27610-27760	27960-28110	
	28090-28315	27600-27825	490
	27600-27825	28090-28315	
	29004-29453	27996-28445	1008
	27996-28445	29004-29453	
	28556-29005	27548-27997	
	27548-27997	28556-29005	
	29100-29125	29225-29250	125
	29225-29250	29100-29125	
31 GHz	31000-31085	31215-31300	175
	31215-31300	31000-31085	
32 GHz	31815-32207	32627-33019	812
	32627-33019	31815-32207	
	32179-32571	32991-33383	
	32991-33383	32179-32571	
38 GHz			1260
	38820-39440	37560-38180	
	37560-38180	38820-39440	
	38316-38936	37045-37676	
	37045-37676	38316-38936	700
	39650-40000	38950-39300	
	38950-39300	39500-40000	
	39300-39650	38600-38950	
	38600-38950	39300-39650	
	37700-38050	37000-37350	
	37000-37350	37700-38050	
	38050-38400	37350-37700	
	37350-37700	38050-38400	
42 GHz			1500
	40550-41278	42050-42778	
	42050-42778	40550-41278	
	41222-41950.5	42722-43450	
	42722-43450	41222-41950.5	

8.7. Mediation Device Losses

Configuration	Interfaces		6-8 GHz	11 GHz	13-15 GHz	18-26 GHz	28-42 GHz
Flex WG	Remote Mount antenna	Added on remote mount configurations	0.5	0.5	1.2	1.5	1.5
1+0	Direct Mount	Integrated antenna	0.2	0.2	0.4	0.5	0.5
OMT	Direct Mount	Integrated antenna	0.3	0.3	0.3	0.5	0.5
Splitter	Direct Mount	Integrated antenna	3.3	3.4	3.4	3.4	3.5

8.8. Ethernet Latency Specifications

8.8.1. Latency – 3.5 MHz Channel Bandwidth

ACM Working Point	Modulation	Frame Size	Latency (μ sec) with GbE Interface					
			64	128	256	512	1024	1518
0	QPSK		1659	1734	1883	2180	2774	3348
1	16 QAM		903	936	1001	1130	1389	1639
2	32 QAM		783	808	856	952	1146	1333
3	64 QAM		704	724	763	842	999	1149
4	128 QAM		650	668	701	767	901	1027
5	256 QAM		587	603	632	692	811	923

8.8.2. Latency – 7 MHz Channel Bandwidth

ACM Working Point	Modulation	Frame Size	Latency (μ sec) with GbE Interface					
			64	128	256	512	1024	1518
0	QPSK		701	763	882	1129	1621	2094
1	8 PSK		516	558	643	810	1144	1466
2	16 QAM		403	432	494	612	852	1081
3	32 QAM		357	381	427	518	702	878
4	64 QAM		326	345	383	459	610	753
5	128 QAM		307	321	353	417	545	666
6	256 QAM		309	324	352	409	522	628
7	512 QAM		346	358	385	437	544	644
8	1024 QAM (strong FEC)		327	340	365	415	516	610
9	1024 QAM (light FEC)		302	314	338	385	481	569

8.8.3. Latency – 14 MHz Channel Bandwidth

ACM Working Point	Modulation	Frame Size	Latency (μ sec) with GbE Interface					
			64	128	256	512	1024	1518
0	QPSK		276	304	360	473	698	913
1	8 PSK		222	242	281	358	512	659
2	16 QAM		178	193	222	280	398	506
3	32 QAM		166	178	201	247	339	424
4	64 QAM		159	168	187	226	303	374
5	128 QAM		191	199	216	248	315	375
6	256 QAM		136	142	158	187	248	302

ACM Working Point	Modulation	Frame Size	Latency (µsec) with GbE Interface					
			64	128	256	512	1024	1518
7	512 QAM		172	179	193	221	277	327
8	1024 QAM (strong FEC)		161	168	182	209	263	310
9	1024 QAM (light FEC)		158	164	177	202	254	299

8.8.4. Latency – 28 MHz Channel Bandwidth

ACM Working Point	Modulation	Frame Size	Latency (µsec) with GbE Interface					
			64	128	256	512	1024	1518
0	QPSK		144	156	183	236	343	446
1	8 PSK		113	122	142	179	256	330
2	16 QAM		98	105	120	148	206	261
3	32 QAM		94	100	112	135	181	225
4	64 QAM		90	95	106	125	165	203
5	128 QAM		84	89	98	115	149	183
6	256 QAM		92	95	104	119	151	181
7	512 QAM		99	103	111	125	155	184
8	1024 QAM (strong FEC)		92	95	103	117	145	173
9	1024 QAM (light FEC)		93	96	104	117	145	171
10	2048 QAM		88	91	99	111	137	162

8.8.5. Latency – 40 MHz Channel Bandwidth

ACM Working Point	Modulation	Frame Size	Latency (µsec) with GbE Interface					
			64	128	256	512	1024	1518
0	QPSK		113	123	144	185	266	346
1	8 PSK		96	103	118	147	205	262
2	16 QAM		81	86	98	120	166	210
3	32 QAM		78	83	92	111	148	184
4	64 QAM		75	79	87	103	135	166
5	128 QAM		71	74	82	96	124	151
6	256 QAM		60	63	71	84	111	137
7	512 QAM		72	75	82	95	120	145
8	1024 QAM (strong FEC)		73	76	82	94	117	141
9	1024 QAM (light FEC)		75	78	84	95	118	140
10	2048 QAM		70	72	78	89	111	132

8.8.6. Latency – 56 MHz Channel Bandwidth

ACM Working Point	Modulation	Frame Size	Latency (μ sec) with GbE Interface					
			64	128	256	512	1024	1518
0	QPSK		85	92	107	138	199	255
1	8 PSK		95	100	112	135	180	222
2	16 QAM		62	67	76	95	132	164
3	32 QAM		59	63	71	87	118	145
4	64 QAM		57	60	67	81	109	133
5	128 QAM		55	58	64	77	103	124
6	256 QAM		55	58	64	76	100	120
7	512 QAM		58	61	67	78	102	121
8	1024 QAM (strong FEC)		55	58	64	75	97	116
9	1024 QAM (light FEC)		56	58	64	75	97	115
10	2048 QAM		53	55	61	72	93	110

8.8.7. Latency – 80 MHz Channel Bandwidth

ACM Working Point	Modulation	Frame Size	Latency (μ sec) with GbE Interface					
			64	128	256	512	1024	1518
0	QPSK		72	78	90	114	163	211
1	8 PSK		58	63	73	91	130	168
2	16 QAM		53	57	65	79	111	141
3	32 QAM		50	54	61	73	101	127
4	64 QAM		48	52	57	69	94	118
5	128 QAM		46	50	55	66	89	111
6	256 QAM		53	56	61	71	93	114
7	512 QAM		50	53	58	67	88	109
8	1024 QAM		47	51	55	65	85	105

8.9. Interface Specifications

8.9.1. Ethernet Interface Specifications

Supported Ethernet Interfaces for Traffic	1 x 10/100/1000Base-T (RJ-45) 2x1000base-X (Optical SFP) or 10/100/1000Base-T (Electrical SFP)
Supported Ethernet Interfaces for Management	10/100 Base-T (RJ-45)
Recommended SFP Types	Optical 1000Base-LX (1310 nm) or SX (850 nm) Note: SFP devices must be of industrial grade (-40°C to +85°C)

The following table lists Netronics-approved SFP devices:

Netronics Marketing Model	Part Number	Item Description
SFP-GE-LX-EXT-TEMP	AO-0097-0	XCVR,SFP,1310nm,1.25Gb, SM,10km,W.DDM,INDUSTRIAL
SFP-GE-SX-EXT-TEMP	AO-0098-0	XCVR,SFP,850nm,MM,1.0625 FC/ 1.25 GBE, INDUSTRIAL
SFP-GE-COPER-EXT-TEMP	AO-0228-0	XCVR,SFP,COOPER 1000BASE-T,RX_LOS DISABLE,INDUSTRIAL TEMP

8.10. Carrier Ethernet Functionality

Latency over the radio link	< 0.15 ms @ 400 Mbps
"Jumbo" Frame Support	Up to 9600 Bytes
General	Enhanced link state propagation Header De-Duplication
Integrated Carrier Ethernet Switch	Switching capacity: 5Gbps / 3.12Mpps Maximum number of Ethernet services: 64 plus one pre-defined management service MAC address learning with 128K MAC addresses 802.1ad provider bridges (QinQ) 802.3ad link aggregation
QoS	Advanced CoS classification and remarking Per interface CoS based packet queuing/buffering (8 queues) Per queue statistics Tail-drop and WRED with CIR/EIR support Flexible scheduling schemes (SP/WFQ/Hierarchical) Per interface and per queue traffic shaping Hierarchical-QoS (H-QoS) – 2K service level queues 2 Gbit packet buffer
Network resiliency	MSTP ERP (G.8032)
OAM	CFM (802.1ag)
Performance Monitoring	Per port Ethernet counters (RMON/RMON2) Radio ACM statistics Enhanced radio Ethernet statistics (Frame Error Rate, Throughput, Capacity, Utilization)
Supported Ethernet/IP Standards	802.3 – 10base-T 802.3u – 100base-T 802.3ab – 1000base-T 802.3z – 1000base-X 802.3ac – Ethernet VLANs 802.1Q – Virtual LAN (VLAN) 802.1p – Class of service 802.1ad – Provider bridges (QinQ) 802.3ad – Link aggregation Auto MDI/MDIX for 1000baseT RFC 1349 – IPv4 TOS RFC 2474 – IPv4 DSCP RFC 2460 – IPv6 Traffic Classes

8.11. Synchronization Functionality

- SyncE
 - SyncE input and output (G.8262)
- IEEE 1588v2 (Precision Time Protocol)
 - Transparent Clock

8.12. Network Management, Diagnostics, Status, and Alarms

Network Management System	Netronics NMS
NMS Interface protocol	SNMPv1/v2c/v3 XML over HTTP/HTTPS toward NMS
Element Management	Web based EMS, CLI
Management Channels & Protocols	HTTP/HTTPS Telnet/SSH-2 FTP/SFTP
Authentication, Authorization & Accounting	User access control X-509 Certificate
Management Interface	Dedicated Ethernet interfaces or in-band in traffic ports
In-Band Management	Support dedicated VLAN for management
TMN	Netronics NMS functions are in accordance with ITU-T recommendations for TMN
RSL Indication	Accurate power reading (dBm) available at NetStream Primo ²³ , and NMS
Performance Monitoring	Integral with onboard memory per ITU-T G.826/G.828

²³ The voltage at the BNC port is 1.XX where XX is the RSL level. For example: 1.59V means an RSL of -59 dBm. Note that the voltage measured at the BNC port is not accurate and should be used only as an aid).

8.13. Mechanical Specifications

Module Dimensions	(H)230mm x (W)233mm x (D)98mm
Module Weight	6 kg
Pole Diameter Range (for Remote Mount Installation)	8.89 cm – 11.43 cm

8.14. Standards Compliance

Specification	Standard
Radio	EN 302 217-2-2
EMC	EN 301 489-1, EN 301 489-4, Class B (Europe) FCC 47 CFR, part 15, class B (US) ICES-003, Class B (Canada) TEC/EMI/TEL-001/01, Class B (India)
Surge	EN61000-4-5, Class 4 (for PWR and ETH1/PoE ports)
Safety	EN 60950-1 IEC 60950-1 UL 60950-1 CSA-C22.2 No.60950-1 EN 60950-22 UL 60950-22 CSA C22.2.60950-22

8.15. Environmental Specifications

- Operating: ETSI EN 300 019-1-4 Class 4.1
 - Temperature range for continuous operating temperature with high reliability:
-33°C to +55°C
 - Temperature range for exceptional temperatures; tested successfully, with limited margins:
-45°C to +60°C
 - Humidity: 5%RH to 100%RH
IEC529 IP66
- Storage: ETSI EN 300 019-1-1 Class 1.2
- Transportation: ETSI EN 300 019-1-2 Class 2.3

8.16. Antenna Specifications

Direct Mount:

Andrew (VHLP), RFS, Xian Putian (WTG), Radio Wave, GD, Shenglu

Remote Mount:

Frequency (GHz)	Waveguide Standard	Waveguide Flange	Antenna Flange
6	WR137	PDR70	UDR70
7/8	WR112	PBR84	UBR84
10/11	WR90	PBR100	UBR100
13	WR75	PBR120	UBR120
15	WR62	PBR140	UBR140
18-26	WR42	PBR220	UBR220
28-38	WR28	PBR320	UBR320
42	WR22	UG383/U	UG383/U

If a different antenna type (CPR flange) is used, a flange adaptor is required. Please contact your Netronics representative for details.

8.17. Power Input Specifications

Standard Input	-48 VDC
DC Input range	-40 to -60 VDC

8.18. Power Consumption Specifications

Maximum Power Consumption	6-11 GHz	13-42 GHz
1+0 Operation	40W	35W
Muted	25W	23W

**Note**

Typical values are 5% less than the values listed above.

8.19. Power Connection Options

Power Source and Range	Data Connection Type	Connection Length	DC Cable Type / Gage
Ext DC -(40.5 ÷ 60)VDC	Optical	≤ 100m	18AWG
		100m ÷ 300m	12AWG
	Electrical	≤ 100m	18AWG
PoE_Inj_AO (All outdoor PoE Injector, -40 ÷ 60VDC)	Electrical	≤ 100m (13 GHz and above) ≤ 75m (6-11 GHz NS Diplo HP)	CAT5e (24AWG)
PoE_Inj_AO_2DC_24V_48V (All outdoor PoE Injector, ±(18 ÷ 60)VDC ²⁴ , DC input redundancy)	Electrical	≤ 100m	CAT5e (24AWG)

²⁴ Optional.

8.20. PoE Injector Specifications

8.20.1. Power Input

Standard Input	-48 or +24VDC (Optional)
DC Input range	$\pm(18^{25}/40.5$ to 60) VDC

8.20.2. Environmental

- Operating: ETSI EN 300 019-1-4 Class 4.1
 - Temperature range for continuous operating temperature with high reliability: **-33°C to +55°C**
 - Temperature range for exceptional temperatures; tested successfully, with limited margins: **-45°C to +60°C**
 - Humidity: 5%RH to 100%RH
IEC529 IP66
- Storage: ETSI EN 300 019-1-1 Class 1.2
- Transportation: ETSI EN 300 019-1-2 Class 2.3

8.20.3. Standards Compliance

Specification	Standard
EMC	EN 301 489-1, EN 301 489-4, Class A (Europe) FCC 47 CFR, part 15, class B (US) ICES-003, Class B (Canada) TEC/EMI/TEL-001/01, Class A (India)
Safety	EN 60950-1 IEC 60950-1 UL 60950-1 CSA-C22.2 No.60950-1 EN 60950-22 UL 60950-22 CSA C22.2.60950-22

8.20.4. Mechanical

Module Dimensions	(H)134mm x (W)190mm x (D)62mm
Module Weight	1kg

²⁵

+18VDC extended range is supported as part of the nominal +24VDC support.

8.21. Cable Specifications

8.21.1. Outdoor Ethernet Cable Specifications

Electrical Requirements	
Cable type	CAT-5e SFUTP, 4 pairs, according to ANSI/TIA/EIA-568-B-2
Wire gage	24 AWG
Stranding	Solid
Voltage rating	70V
Shielding	Braid + Foil
Pinout	<p style="text-align: center;">RJ45,P1 RJ45,P2</p> <p>1 ——— WHITE/GREEN ——— 1</p> <p>2 ——— GREEN ——— 2</p> <p>3 ——— WHITE/ORANGE ——— 3</p> <p>6 ——— ORANGE ——— 6</p> <p>4 ——— BLUE ——— 4</p> <p>5 ——— WHITE/BLUE ——— 5</p> <p>7 ——— WHITE/BROWN ——— 7</p> <p>8 ——— BROWN ——— 8</p> <p>SHELL ——— DRAIN WIRE + SHIELD ——— SHEEL</p>
Mechanical/ Environmental Requirements	
Jacket	PVC, double, UV resistant
Outer diameter	7-10 mm
Operating and Storage temperature range	-40°C - 85°C
Flammability rating	According to UL-1581 VW1, IEC 60332-1
RoHS	According to Directive/2002/95/EC

8.21.2. Outdoor DC Cable Specifications

Electrical Requirements	
Cable type	2 tinned copper wires
Wire gage	18 AWG (for <100m installations) 12 AWG (for >100m installations)
Stranding	stranded
Voltage rating	600V
Spark test	4KV
Dielectric strength	2KV AC min
Mechanical/ Environmental Requirements	
Jacket	PVC, double, UV resistant
Outer diameter	7-10 mm
Operating & Storage temperature range	-40°C - 85°C
Flammability rating	According to UL-1581 VW1, IEC 60332-1
RoHS	According to Directive/2002/95/EC