LOFAR-WAN ADD

D.H.P. Maat, R.B. Gloudemans

Verified:			
Name	Signature	Date	Rev.nr.

Accepted:		
Work Package Manager	System Engineering Manager	Program Manager
D.H.P. Maat ………………………………….. *date:*	A. Gunst ………………………………….. *date:*	J.M. Reitsma ………………………………….. *date:*

Distribution list:

Group:	For Information:
ASTRON D.H.P. Maat R. Gloudemans T. Grit K. Stuurwold J.M. Reitsma A.W. Gunst	ASTRON

Document history:

Revision	Date	Section	Page(s)	Modification
1.0	2007-15-3		-	Creation

Table of contents:

1 Introduction ... 5
2 Architecture description ... 6
2.1 Introduction ... 6
2.2 LOFAR data transport ... 6
3 Monitoring and Control ... 11
3.1 Sensor Control ... 11
3.2 Transient detection mode ... 11
4 Partner sensors ... 12
5 E-LOFAR ... 13
6 System Requirement Specification compliance ... 14

1 Introduction

The function of the LOFAR Wide Area Network is to transport data between the LOFAR stations and the central processor site. The main part of the LOFAR data that is transported throughout the network is observation data. In addition, a small part of the LOFAR data stream concerns MAC related data transport, partner data and active equipment management data. The role of the Wide Area Network with respect to the transport of the observation data is depicted in Figure 1.

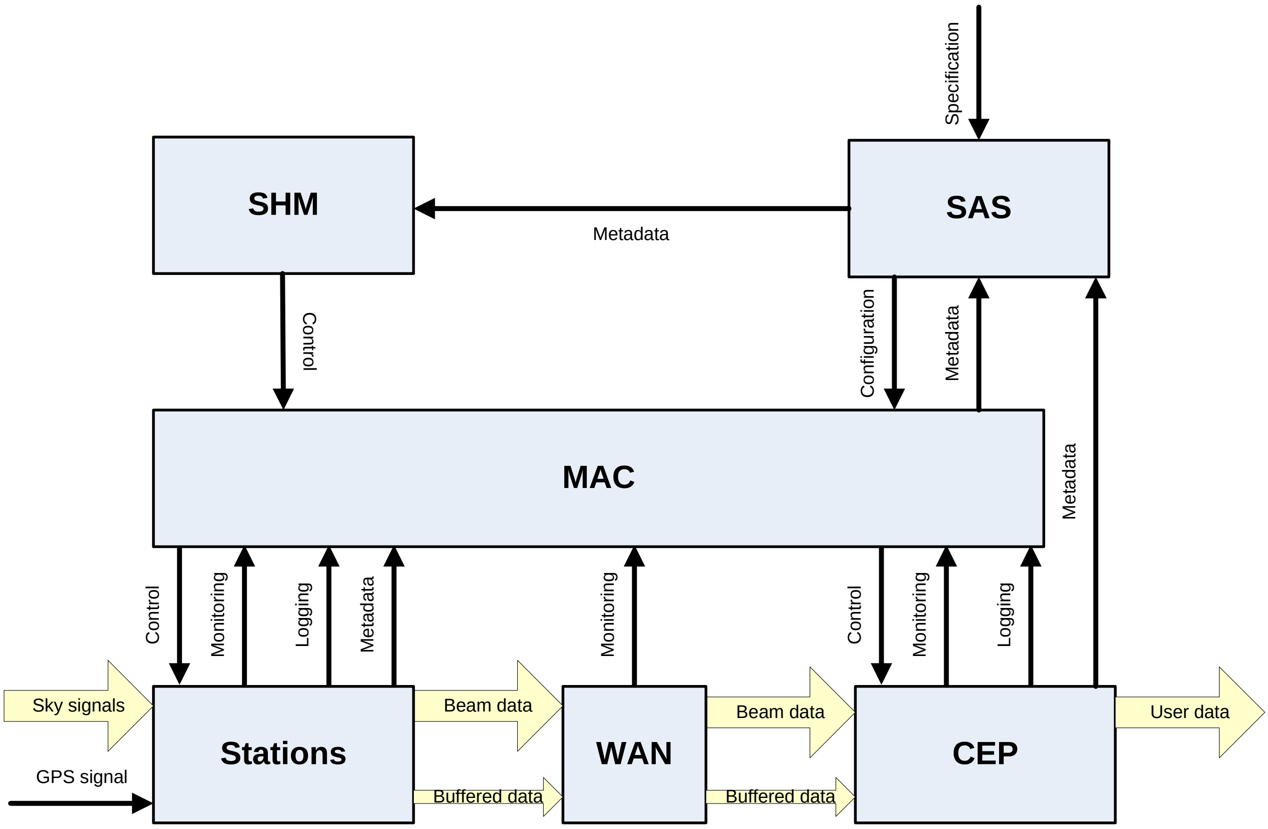

Figure 1 The role of the Wide Area Network in the observation data transport system

The figure shows that the WAN has interfaces with the Station, CEP and MAC sections. The WAN will handle the data transport for both the stations in the Netherlands and the E-LOFAR stations that are located abroad. The first part of the Wide Area Network consists out of (leased/owned) managed dark fibre and active communication equipment. The E-LOFAR stations will be connected with the use of bandwidth/wavelength services that are provided by NRENs (National Research and Education Networks) or commercial parties.
In addition to the LOFAR related data, the network will also be used for the transportation of data that is produced by the sensor systems of the LOFAR partners, like the KNMI and University of Wageningen. The data from these sensors also needs to be routed over the LOFAR network. This has to be done in such a way that other running experiments will not be influenced.

In this document the design of both the national LOFAR and the E-LOFAR data transport systems is described.

2 Architecture description

2.1 Introduction

In this chapter the design and configuration of the LOFAR data transport network is described. Since the approach for the Dutch LOFAR and the E-LOFAR network are very different, both are treated in different sections.

2.2 LOFAR data transport

2.2.1 Network topology and fibre-optic cable configuration

In Figure 2 a schematic topology of the Dutch LOFAR-WAN is depicted. Three parts can be distinguished:

(1) The LOFAR 'arms' (green), along which the remote LOFAR antenna stations are positioned. The central hub of these links is located near Exloo (concentrator node or CN), near the LOFAR central core. Distances from the antenna stations along the spiral arms to the core range up to 80 kilometres. The data rate per remote station is 2Gb/s.

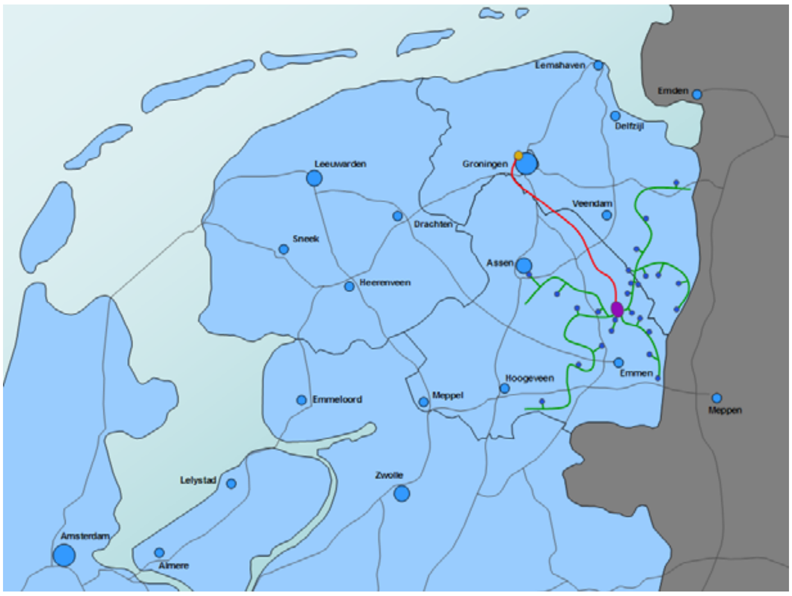

Figure 2 Generalised topology of the Dutch LOFAR-WAN (actual remote station sites and cable routes are not shown)

(2) The LOFAR core (purple) in which the central field antenna groups are placed. These core stations generate at first 2Gb/s. Later on the data rate will be increased to 12Gb/s to accommodate for more beams in the core. The active equipment that is to be installed will be capable of handling the latter data rate. The data generated in the core stations is transported via short distance (< 2 km) links to the earlier mentioned communication hub (CN) near the core.

ASTRON

(3) LOFAR backbone (red), which is used for the transportation of the remote and core station data from the central hub (CN) near Exloo to the communication node at the Central Processing Facility (CEP node or CEPN). Data rates up to ~ 700 Gb/s need to be transported via this link, which has a length of 65 km.

In the WAN configuration as described above, the fibre-optic cables of the individual stations can share the cable trenches in an optimal way, resulting in the most easy network deployment and the lowest deployment costs. In addition, by collecting all data at a point near the central field, pre-processing can be performed near the central field, before all data is transferred to the central processor at Groningen.
Since the availability requirement for LOFAR is relatively low (95%), when compared to commercial data communication networks, the use of redundancy is not needed in the network.

2.2.2 Fibre-optic cabling

Throughout the LOFAR system single mode fibre is employed of type G652-C/D. This is a standard single mode fibre with low water peak (low attenuation between 1300 nm and 1600 nm). By employing this type of fibre the low cost wavelength division multiplexing technique CWDM can be applied.

2.2.2.1 Fibre count: remote stations

The remote stations will be connected to the CN via leased lines. The running fibre cable tenders show that the lowest link deployment costs are obtained by minimising the amount of fibres that run from each station to the CN. As such each remote station is equipped with only two fibre pairs. One pair (with three CWDM data links) is used for observation data, MAC related data and partner data. The switch management traffic is transferred via the second pair. By separating the data links from the switch management link, a data link failure due to a wrong configuration of this link can still be solved remotely.

2.2.2.2 Fibre count: core stations

The fibres in the LOFAR core will be installed and owned by LOFAR. Since the costs of installing a cable are mostly determined by the digging costs (in the case of relatively low fibre counts), the amount of fibres per antenna station does not have to be as small as possible. As such, three fibre pairs per station will be used, which is an amount of fibre pairs per station that equals the amount of 1GbE links per station that are initially needed. Later on it is likely that two 10GbE links and one switch management link will be employed per core station.

2.2.2.3 Fibre count: backbone

The amount of fibres in the backbone should be sufficient to transport all LOFAR data. The maximum amount of data that LOFAR currently can handle is equal to the input data rate of the BlueGene/L which is about 720 Gb/s. At the time the backbone needed to be installed, only 1GbE CWDM technology was capable of providing the required data rates per fibres at an acceptable cost level. The optical loss of the link and the power budget of the optical transceiver modules were such that about 10 1GbE CWDM channels could be transported from CN to CEP per fibre pair. To be capable of providing the 720 Gb/s, a cable with 72 fibre pairs is installed at the link CN – CEP.

2.2.3 Protocols and active equipment configuration

The chosen network protocols for the transport of the observation data are Ethernet and UDP. Since huge amounts of data need to be transported via the LOFAR-WAN, special care needs to be taken to minimise the network overhead. Because of this, jumbo frames (9kB size) will be used on the Ethernet level instead of regular Ethernet (1.5kB) packets.
No internal routing is needed in the LOFAR WAN and packets can be sent directly to the destination at CEP using the MAC (Medium Access Control) address. Since the LOFAR network is fairly simple, this task can be fulfilled by layer 2+ switches. In addition to the observation data, all other data will also be transferred via the switches. The latter data will use regular TCP/IP connections.

At each station all data streams (local and external) are handled by a relatively simple (edge) routing switch (stack). At the CN and CEPN more complex switching frames will be needed to handle the data streams.

A good method to handle the large amount of data streams is the usage of Virtual LAN's (VLAN). These can be defined at the network appliance level. A VLAN can separate one physical network into several virtual ones. Network traffic can be routed between VLAN's, but the VLAN can also be totally isolated from other VLAN's. VLAN's can also be used to split the LOFAR network into separate networks for the different sensor types.
VLAN's have one downside; they increase the network overhead between network appliances, but not between network endpoints. The extra overhead is just a few bytes per packet. When jumbo Ethernet frames are used (9kB packets) this extra overhead is less than 0.1% of each packet.

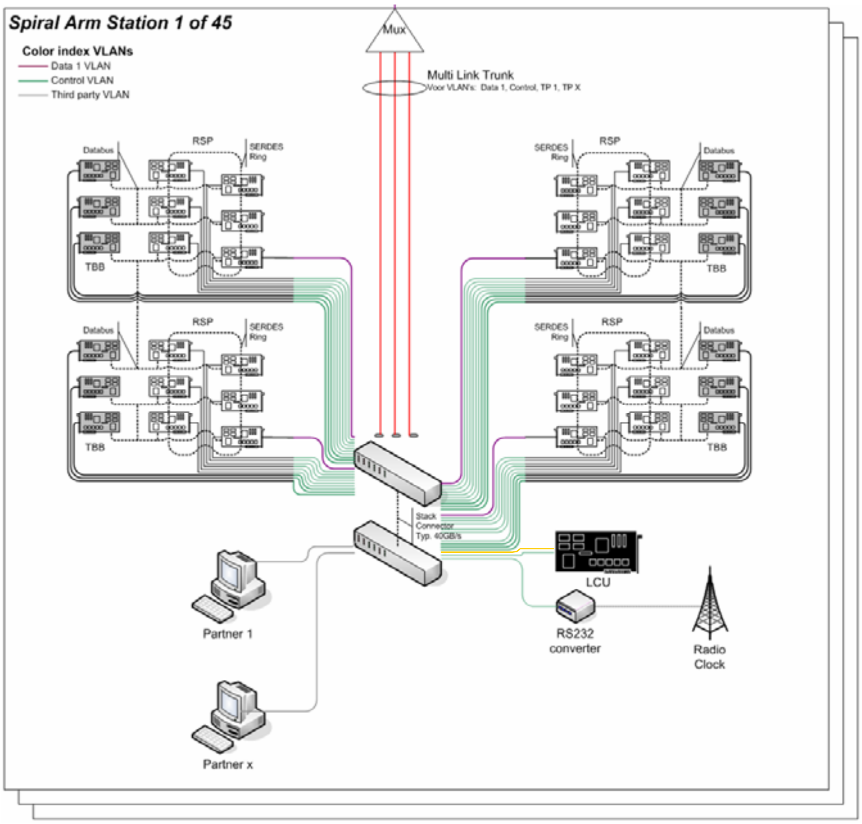

Figure 3 Schematic overview over the station network architecture

Using the this approach, the data streams in each station are divided into a number of VLAN's. A schematic VLAN network architecture in the LOFAR station is given in Figure 3.

- VLAN1: The observation data from the four RSP board outputs are placed in a VLAN that ranges up to the input cluster of the CEP (purple).
- VLAN2: the LCU can communicate with the CCU at the CEP via a MAC VLAN (yellow).
- VLAN3: is a local VLAN via which the LCU communicates with all the systems at the station (green).
- VLAN4: This VLAN will handle the switch control traffic.
- VLAN5 .. N: are the VLAN's of the different LOFAR partners. Each partner has a single VLAN (grey).

All the relevant VLAN's are transferred to the CEP location via a multi link trunk (red). The switch control via the separate fibre pair is connected to the serial port of the switch via a console switch. The shown network architecture will be employed at all stations. For the core station upgrade, mentioned in Section 2.2.1, no VLAN changes are needed.

2.2.4 Communication technology

The type of the communication technology that is to be installed throughout the network will be determined by the companies that respond to a WAN active equipment tender. As such, details on the communication technology can not be provided in this document.

2.2.4.1 Switches

The only information that is currently available about the switches that are going to be employed is related to the switch requirements that will be given in the active equipment tender document. For the WAN switches the requirements are:

- Support for the MAC (Medium Access Control) address based traffic
- Support for VLAN's
- Support for 9k byte jumbo frames
- Support for UDP traffic
- Support for multicast (needed for the communication in the MAC VLAN)
- Support for QOS (Quality Of Service is needed for control of the partner bandwidths)
- Controllable by SNMP (Simple Network Management Protocol for interfacing to MAC)
- Sufficient backplane speed
- Sufficient 1000BASE-T and SFP/XFP ports.

Switches that comply with the mentioned requirements are widely available.

2.2.4.2 Tx/Rx – optics and WDM (de)multiplexers

In the LOFAR switches pluggable optical transceivers will be used. Many different pluggable optics standards are available. LOFAR will use SFP and XFP pluggable optics, since these standards are expected to have a longer support than all other standards. For 1 GbE links single mode SFP's are available with transmission distances between 10 and 120 km. These SFP types are also available for CWDM systems with up to 16 wavelength channels. SFP-based DWDM (Dense WDM) technology is currently also appearing on the market. For (10GbE) XFP's the situation is more or less the same, although the pricing level of XFP technology is still higher. For transmission of 10GbE traffic via the employed G652-C/D fibre beyond ~60 km, dispersion compensation is required. Currently both fibre based and (XFP-based) electronic dispersion compensation is available. The electronic dispersion compensation is preferred over the fibre based dispersion compensation due to the relatively high optical losses in the latter.

CWDM (de)multiplex technology is commercially available in a number of system configurations, ranging from complete CWDM Tx/RX systems to standalone (passive) (de)mux components. The systems can handle up to 16 wavelength channels and can be used for both 1GbE and 10GbE traffic. DWDM (de)multiplexers are currently being introduced at the market. Their pricing is still relatively high.

3 Monitoring and Control

3.1 Sensor Control

The LOFAR monitoring and control by the MAC system will take place using SNMP. Network equipment not only needs to be able to pass these SNMP commands in the MAC VLAN, but also needs to be controlled itself using SNMP. This means that the SNMP commands with which the switches can be controlled and monitored should as close to the industrial standard as possible. Relevant information about the status of the switches is stored in a Management Information Base. A description of this information data base needs to provided by the switch supplier.

3.2 Transient detection mode

In order to detect transients it should be possible to switch antenna observation modes and/or dump the cache (from the TBB boards) with one command. To do this, two mechanisms are available:

- If only a single MAC VLAN is used, a SNMP command may be broadcast.
- If the MAC connections are distributed over several VLAN's, Multicast should be used.

4 Partner sensors

Apart from LOFAR data transport, the network will also be used by the LOFAR partners for the transportation of their sensor data. The network should be built in such a way that the individual data streams do not affect each other. This has implications in two areas:

- Bandwidth. A sensor may never generate more data than is allotted to that sensor and the full bandwidth allotted to that sensor should always be available. To obtain this, the Ethernet Quality of Service (QoS) protocol can be used, but it is probably better and simpler to limit the network interfaces of a switch at a certain speed. For most switches this is a per port configurable parameter. The downside is that possible network speeds are 10, 100 or 1000 Mbit/s. QoS offers finer grained control, but also adds to the network overhead. Most important in this is that it generates load on the CPU of a switch, thus decreasing the maximum data speed a switch can handle.
- The partners should never be able to access the sensor systems of the other partners. This can be accomplished using the VLAN protocol. Another option may be Virtual Private Networks (VPN). The advantage of the VPN is that the network operator cannot influence the data stream, but it also requires much more management. Best solution is to use both. But since the VLAN makes the network safe enough for the network operator, usage of VPN should be left to the sensor owners. (They can do this by installing a VPN endpoint at the sensor location)

5 E-LOFAR

Apart from the stations in the Netherlands, LOFAR stations will also be placed in other countries in Europe (E-LOFAR), like Germany, Great Britain, France, Sweden and Poland. The data connection between these stations will be realised with the use of bandwidth/wavelength services that are to be provided by commercial parties or NRENs. For the transfer of the data, dedicated links will be used which will only contain the traffic of the E-LOFAR stations. In this way a safe transfer of the data can be realised, with a well predictable delay. The dedicated links are configured in such a way, that the E-LOFAR data only needs to be provided to one or more interfaces at a switch (owned by the bandwidth/wavelength supplier) at the E-LOFAR station side, to have it available at an output interface of a switch (owned by the bandwidth/wavelength supplier) at the CEP site in Groningen. The NRENs throughout Europe are currently deploying so-called lightpath technology that is capable of providing this type of data transport.
By making use of this type of data transfer technology, the E-LOFAR station configuration can be identical to the configuration of the LOFAR stations in the Netherlands. Since the E-LOFAR partners also want to make use of the locally generated observation, a data splitter (PC) needs to be placed in the data path, such that the data is transferred both to the CEP in Groningen and to a local system.

A schematic layout of the local E-LOFAR data transmission system is given in Figure 4.

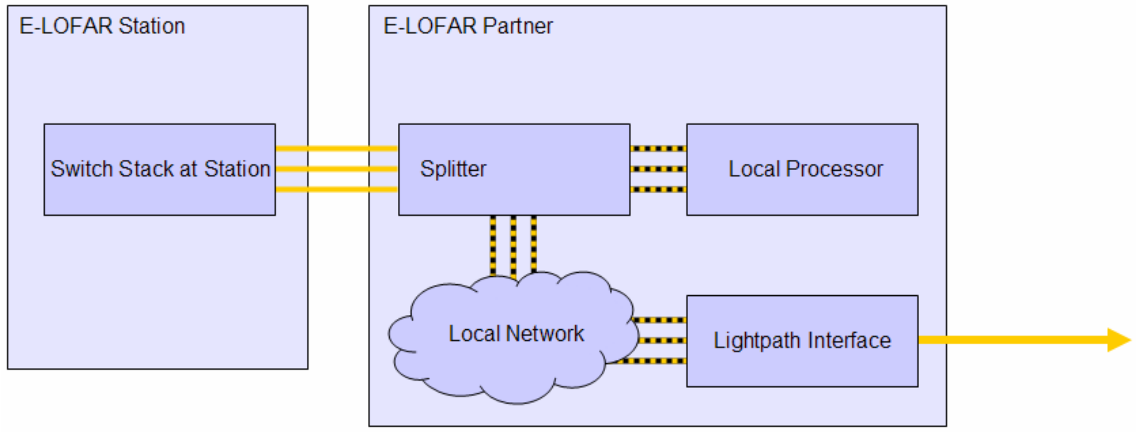

Figure 4 Schematic layout of the data transmission system at the E-LOFAR site

The splitter, local network equipment and lightpath interface needs to be capable of handling the LOFAR data as specified in Chapter 2. At the CEP location in Groningen the lightpath interfaces of all the E-LOFAR lightpaths are present. From these interfaces the E-LOFAR data will be transferred to the input cluster of the CEP system.

6 System Requirement Specification compliance

Within the System Requirement Specification a number of requirements for the WAN are given. In this chapter the compliance with these requirements is treated (NC: not compliant, TBD: to be determined).

LO-3.09.4 Data Transport function

01	LOFAR shall provide for transport of measurement data, monitoring data and control data.	OK	
02	LOFAR shall provide for transport of data from all Remote Stations to the Central Processor.	OK	
03	LOFAR shall provide for transport of data from the Virtual Core (substations) to the Central Processing facilities.	OK	
04	LOFAR shall be capable of transporting data autonomously (this implies that the data transport network shall be able to operate independently of other subsystems).	OK	
05	The data transport network shall incorporate its own autonomous network management and control system.	OK	
06	The data transport shall be protected against any environmental influences as specified in section 5.01.	NC	(see LO-5.01)
07	The data transport network shall incorporate the ability to detect and identify failing network equipment (hardware and software).	OK	
08	The data transport shall be protected against (mechanical) damages caused by inadvertent human or animal activities.	NC	The network will be deployed according to widely applied standards. This implies that only a partly protection is provided.
09	The data transport network shall be protected against unauthorized use.	OK	
10	The data transport shall not be interrupted or corrupted during maintenance activities, related to the data transport.	NC	During network maintenance (parts of) the network will not be available for the transport of data.
11	LOFAR shall be able to transport all monitor and control data between the LOFAR subsystems that it connects.	OK	
12	LOFAR shall provide for a redundant (back-up) network to transport all control data and consolidated monitor data between the LOFAR central processor (CEP) and the remote and virtual core stations.	TBD	The mentioned redundancy is relatively costly. The benefits of the redundancy vs. the costs need to be considered.
13	The monitoring and control data shall be transported via a single full duplex network.	OK	
14	The data transport network connecting the LOFAR stations the central processing facility shall be deployed according t the regulations applicable to general data networks.	OK	
15	It shall be possible to reuse the LOFAR data transport network for other purposes after the lifetime of LOFAR.	OK	
16	The data transport network shall preserve the time ordering of each input data flow (this implies that all output data from a station keeps the correct relative time order).	OK	

ASTRON

LO-5.01 Environmental Requirements (only relevant issues are mentioned)

LO-5.01.02 Site and infrastructure requirements

01	Buildings or parts of buildings containing central processing equipment and operator areas shall have a climatic conditioning system which can control the temperature within the range of 18 ºC to 23 ºC and the humidity within the range of 50 % to 70 % independent of weather conditions.	OK	
06	LOFAR central processing and operating facilities shall be provided with remotely video and audio monitoring equipment.	TBD	The mentioned feature needs a relatively costly data link. The benefits of the feature vs. the costs need to be considered.
07	LOFAR stations shall be provided with remotely operated video and audio monitoring equipment.	TBD	The mentioned feature needs a relatively costly data link. The benefits of the feature vs. the costs need to be considered.
08	LOFAR station shall be provided with a fixed telephone line to call with the operators in the central processing room	TBD	The mentioned feature needs a relatively costly link. The benefits of the feature vs. the costs need to be considered.