Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: June 22, 2019

E. Grossman, Ed.
DOLBY

December 19, 2018

Deterministic Networking Use Cases
draft-ietf-detnet-use-cases-20

Abstract

This draft presents use cases from diverse industries which have in common a need for "deterministic flows". "Deterministic" in this context means that such flows provide guaranteed bandwidth, bounded latency, and other properties germane to the transport of time-sensitive data. These use cases differ notably in their network topologies and specific desired behavior, providing as a group broad industry context for DetNet. For each use case, this document will identify the use case, identify representative solutions used today, and describe potential improvements that DetNet can enable. The Use Case Common Themes section then extracts and enumerates the set of common properties implied by these use cases.

Table of Contents

1. Introduction

This draft documents use cases in diverse industries which require deterministic flows over multi-hop paths. DetNet flows can be established from either a Layer 2 or Layer 3 (IP) interface, and such flows can co-exist on an IP network with best-effort traffic. DetNet also provides for highly reliable flows through provision for redundant paths.

The DetNet Use Cases explicitly do not suggest any specific design for DetNet architecture or protocols; these are topics of other DetNet drafts.

The DetNet use cases as originally submitted explicitly were not considered by the DetNet Working Group to be concrete requirements; The DetNet Working Group and Design Team considered these use cases,

identifying which elements of them could be feasibly implemented within the charter of DetNet, and as a result certain of the originally submitted use cases (or elements of them) have been be moved to the Use Cases Explicitly Out of Scope for DetNet section.

The DetNet Use Cases document provide context regarding DetNet design decisions. It also serves a long-lived purpose of helping those learning (or new to) DetNet to understand the types of applications that can be supported by DetNet. It also allow those WG contributors who are users to ensure that their concerns are addressed by the WG; for them this document both covers their contribution and provides a long term reference to the problems they expect to be served by the technology, both in the short term deliverables and as the technology evolves in the future.

The DetNet Use Cases document has served as a "yardstick" against which proposed DetNet designs can be measured, answering the question "to what extent does a proposed design satisfy these various use cases?"

The Use Case industries covered are professional audio, electrical utilities, building automation systems, wireless for industrial applications, cellular radio, industrial machine-to-machine, mining, private blockchain, and network slicing. For each use case the following questions are answered:

- What is the use case?
- How is it addressed today?
- How should it be addressed in the future?
- What should the IETF deliver to enable this use case?

The level of detail in each use case is intended to be sufficient to express the relevant elements of the use case, but not greater than that.

DetNet does not directly address clock distribution or time synchronization; these are considered to be part of the overall design and implementation of a time-sensitive network, using existing (or future) time-specific protocols (such as [IEEE8021AS] and/or [RFC5905]).

2. Pro Audio and Video

2.1. Use Case Description

The professional audio and video industry ("ProAV") includes:

- Music and film content creation
- Broadcast
- Cinema
- Live sound
- Public address, media and emergency systems at large venues (airports, stadiums, churches, theme parks).

These industries have already transitioned audio and video signals from analog to digital. However, the digital interconnect systems remain primarily point-to-point with a single (or small number of) signals per link, interconnected with purpose-built hardware.

These industries are now transitioning to packet-based infrastructure to reduce cost, increase routing flexibility, and integrate with existing IT infrastructure.

Today ProAV applications have no way to establish deterministic flows from a standards-based Layer 3 (IP) interface, which is a fundamental limitation to the use cases described here. Today deterministic flows can be created within standards-based layer 2 LANs (e.g. using IEEE 802.1 AVB) however these are not routable via IP and thus are not effective for distribution over wider areas (for example broadcast events that span wide geographical areas).

It would be highly desirable if such flows could be routed over the open Internet, however solutions with more limited scope (e.g. enterprise networks) would still provide a substantial improvement.

The following sections describe specific ProAV use cases.

2.1.1. Uninterrupted Stream Playback

Transmitting audio and video streams for live playback is unlike common file transfer because uninterrupted stream playback in the presence of network errors cannot be achieved by re-trying the transmission; by the time the missing or corrupt packet has been identified it is too late to execute a re-try operation. Buffering can be used to provide enough delay to allow time for one or more

retries, however this is not an effective solution in applications where large delays (latencies) are not acceptable (as discussed below).

Streams with guaranteed bandwidth can eliminate congestion on the network as a cause of transmission errors that would lead to playback interruption. Use of redundant paths can further mitigate transmission errors to provide greater stream reliability.

Additional techniques such as forward error correction can also be used to improve stream reliability.

2.1.2. Synchronized Stream Playback

Latency in this context is the time between when a signal is initially sent over a stream and when it is received. A common example in ProAV is time-synchronizing audio and video when they take separate paths through the playback system. In this case the latency of both the audio and video streams must be bounded and consistent if the sound is to remain matched to the movement in the video. A common tolerance for audio/video sync is one NTSC video frame (about 33ms) and to maintain the audience perception of correct lip sync the latency needs to be consistent within some reasonable tolerance, for example 10%.

A common architecture for synchronizing multiple streams that have different paths through the network (and thus potentially different latencies) is to enable measurement of the latency of each path, and have the data sinks (for example speakers) delay (buffer) all packets on all but the slowest path. Each packet of each stream is assigned a presentation time which is based on the longest required delay. This implies that all sinks must maintain a common time reference of sufficient accuracy, which can be achieved by any of various techniques.

This type of architecture is commonly implemented using a central controller that determines path delays and arbitrates buffering delays.

2.1.3. Sound Reinforcement

Consider the latency (delay) from when a person speaks into a microphone to when their voice emerges from the speaker. If this delay is longer than about 10-15 milliseconds it is noticeable and can make a sound reinforcement system unusable (see slide 6 of [SRP_LATENCY]). (If you have ever tried to speak in the presence of a delayed echo of your voice you may know this experience).

Note that the 15ms latency bound includes all parts of the signal path, not just the network, so the network latency must be significantly less than 15ms.

In some cases local performers must perform in synchrony with a remote broadcast. In such cases the latencies of the broadcast stream and the local performer must be adjusted to match each other, with a worst case of one video frame (33ms for NTSC video).

In cases where audio phase is a consideration, for example beam-forming using multiple speakers, latency can be in the 10 microsecond range (1 audio sample at 96kHz).

2.1.4. Secure Transmission

2.1.4.1. Safety

Professional audio systems can include amplifiers that are capable of generating hundreds or thousands of watts of audio power which if used incorrectly can cause hearing damage to those in the vicinity. Apart from the usual care required by the systems operators to prevent such incidents, the network traffic that controls these devices must be secured (as with any sensitive application traffic).

2.2. Pro Audio Today

Some proprietary systems have been created which enable deterministic streams at Layer 3 however they are "engineered networks" which require careful configuration to operate, often require that the system be over-provisioned, and it is implied that all devices on the network voluntarily play by the rules of that network. To enable these industries to successfully transition to an interoperable multi-vendor packet-based infrastructure requires effective open standards, and establishing relevant IETF standards is a crucial factor.

2.3. Pro Audio Future

2.3.1. Layer 3 Interconnecting Layer 2 Islands

It would be valuable to enable IP to connect multiple Layer 2 LANs.

As an example, ESPN constructed a state-of-the-art 194,000 sq ft, $125 million broadcast studio called DC2. The DC2 network is capable of handling 46 Tbps of throughput with 60,000 simultaneous signals. Inside the facility are 1,100 miles of fiber feeding four audio control rooms (see [ESPN_DC2]).

In designing DC2 they replaced as much point-to-point technology as they could with packet-based technology. They constructed seven individual studios using layer 2 LANS (using IEEE 802.1 AVB) that were entirely effective at routing audio within the LANs. However to interconnect these layer 2 LAN islands together they ended up using dedicated paths in a custom SDN (Software Defined Networking) router because there is no standards-based routing solution available.

2.3.2. High Reliability Stream Paths

On-air and other live media streams are often backed up with redundant links that seamlessly act to deliver the content when the primary link fails for any reason. In point-to-point systems this is provided by an additional point-to-point link; the analogous requirement in a packet-based system is to provide an alternate path through the network such that no individual link can bring down the system.

2.3.3. Integration of Reserved Streams into IT Networks

A commonly cited goal of moving to a packet based media infrastructure is that costs can be reduced by using off the shelf, commodity network hardware. In addition, economy of scale can be realized by combining media infrastructure with IT infrastructure. In keeping with these goals, stream reservation technology should be compatible with existing protocols, and not compromise use of the network for best-effort (non-time-sensitive) traffic.

2.3.4. Use of Unused Reservations by Best-Effort Traffic

In cases where stream bandwidth is reserved but not currently used (or is under-utilized) that bandwidth must be available to best-effort (i.e. non-time-sensitive) traffic. For example a single stream may be nailed up (reserved) for specific media content that needs to be presented at different times of the day, ensuring timely delivery of that content, yet in between those times the full bandwidth of the network can be utilized for best-effort tasks such as file transfers.

This also addresses a concern of IT network administrators that are considering adding reserved bandwidth traffic to their networks that "users will reserve large quantities of bandwidth and then never un-reserve it even though they are not using it, and soon the network will have no bandwidth left".

2.3.5. Traffic Segregation

Sink devices may be low cost devices with limited processing power. In order to not overwhelm the CPUs in these devices it is important to limit the amount of traffic that these devices must process.

As an example, consider the use of individual seat speakers in a cinema. These speakers are typically required to be cost reduced since the quantities in a single theater can reach hundreds of seats. Discovery protocols alone in a one thousand seat theater can generate enough broadcast traffic to overwhelm a low powered CPU. Thus an installation like this will benefit greatly from some type of traffic segregation that can define groups of seats to reduce traffic within each group. All seats in the theater must still be able to communicate with a central controller.

There are many techniques that can be used to support this feature including (but not limited to) the following examples.

2.3.5.1. Packet Forwarding Rules, VLANs and Subnets

Packet forwarding rules can be used to eliminate some extraneous streaming traffic from reaching potentially low powered sink devices, however there may be other types of broadcast traffic that should be eliminated using other means for example VLANs or IP subnets.

2.3.5.2. Multicast Addressing (IPv4 and IPv6)

Multicast addressing is commonly used to keep bandwidth utilization of shared links to a minimum.

Because of the MAC Address forwarding nature of Layer 2 bridges it is important that a multicast MAC address is only associated with one stream. This will prevent reservations from forwarding packets from one stream down a path that has no interested sinks simply because there is another stream on that same path that shares the same multicast MAC address.

Since each multicast MAC Address can represent 32 different IPv4 multicast addresses there must be a process put in place to make sure this does not occur. Requiring use of IPv6 address can achieve this, however due to their continued prevalence, solutions that are effective for IPv4 installations are also desirable.

2.3.6. Latency Optimization by a Central Controller

A central network controller might also perform optimizations based on the individual path delays, for example sinks that are closer to the source can inform the controller that they can accept greater latency since they will be buffering packets to match presentation times of farther away sinks. The controller might then move a stream reservation on a short path to a longer path in order to free up bandwidth for other critical streams on that short path. See slides 3-5 of [SRP_LATENCY].

Additional optimization can be achieved in cases where sinks have differing latency requirements, for example in a live outdoor concert the speaker sinks have stricter latency requirements than the recording hardware sinks. See slide 7 of [SRP_LATENCY].

2.3.7. Reduced Device Cost Due To Reduced Buffer Memory

Device cost can be reduced in a system with guaranteed reservations with a small bounded latency due to the reduced requirements for buffering (i.e. memory) on sink devices. For example, a theme park might broadcast a live event across the globe via a layer 3 protocol; in such cases the size of the buffers required is proportional to the latency bounds and jitter caused by delivery, which depends on the worst case segment of the end-to-end network path. For example on todays open internet the latency is typically unacceptable for audio and video streaming without many seconds of buffering. In such scenarios a single gateway device at the local network that receives the feed from the remote site would provide the expensive buffering required to mask the latency and jitter issues associated with long distance delivery. Sink devices in the local location would have no additional buffering requirements, and thus no additional costs, beyond those required for delivery of local content. The sink device would be receiving the identical packets as those sent by the source and would be unaware that there were any latency or jitter issues along the path.

2.4. Pro Audio Asks

- Layer 3 routing on top of AVB (and/or other high QoS networks)
- Content delivery with bounded, lowest possible latency
- IntServ and DiffServ integration with AVB (where practical)
- Single network for A/V and IT traffic
- Standards-based, interoperable, multi-vendor

- IT department friendly

- Enterprise-wide networks (e.g. size of San Francisco but not the whole Internet (yet...))

3. Electrical Utilities

3.1. Use Case Description

Many systems that an electrical utility deploys today rely on high availability and deterministic behavior of the underlying networks. Presented here are use cases in Transmission, Generation and Distribution, including key timing and reliability metrics. In addition, security issues and industry trends which affect the architecture of next generation utility networks are discussed.

3.1.1. Transmission Use Cases

3.1.1.1. Protection

Protection means not only the protection of human operators but also the protection of the electrical equipment and the preservation of the stability and frequency of the grid. If a fault occurs in the transmission or distribution of electricity then severe damage can occur to human operators, electrical equipment and the grid itself, leading to blackouts.

Communication links in conjunction with protection relays are used to selectively isolate faults on high voltage lines, transformers, reactors and other important electrical equipment. The role of the teleprotection system is to selectively disconnect a faulty part by transferring command signals within the shortest possible time.

3.1.1.1.1. Key Criteria

The key criteria for measuring teleprotection performance are command transmission time, dependability and security. These criteria are defined by the IEC standard 60834 as follows:

- Transmission time (Speed): The time between the moment where state changes at the transmitter input and the moment of the corresponding change at the receiver output, including propagation delay. Overall operating time for a teleprotection system includes the time for initiating the command at the transmitting end, the propagation delay over the network (including equipments) and the selection and decision time at the receiving end, including any additional delay due to a noisy environment.

- Dependability: The ability to issue and receive valid commands in the presence of interference and/or noise, by minimizing the probability of missing command (PMC). Dependability targets are typically set for a specific bit error rate (BER) level.

- Security: The ability to prevent false tripping due to a noisy environment, by minimizing the probability of unwanted commands (PUC). Security targets are also set for a specific bit error rate (BER) level.

Additional elements of the teleprotection system that impact its performance include:

- Network bandwidth

- Failure recovery capacity (aka resiliency)

3.1.1.1.2. Fault Detection and Clearance Timing

Most power line equipment can tolerate short circuits or faults for up to approximately five power cycles before sustaining irreversible damage or affecting other segments in the network. This translates to total fault clearance time of 100ms. As a safety precaution, however, actual operation time of protection systems is limited to 70- 80 percent of this period, including fault recognition time, command transmission time and line breaker switching time.

Some system components, such as large electromechanical switches, require particularly long time to operate and take up the majority of the total clearance time, leaving only a 10ms window for the telecommunications part of the protection scheme, independent of the distance to travel. Given the sensitivity of the issue, new networks impose requirements that are even more stringent: IEC standard 61850 limits the transfer time for protection messages to 1/4 - 1/2 cycle or 4 - 8ms (for 60Hz lines) for the most critical messages.

3.1.1.1.3. Symmetric Channel Delay

Teleprotection channels which are differential must be synchronous, which means that any delays on the transmit and receive paths must match each other. Teleprotection systems ideally support zero asymmetric delay; typical legacy relays can tolerate delay discrepancies of up to 750us.

Some tools available for lowering delay variation below this threshold are:

- For legacy systems using Time Division Multiplexing (TDM), jitter buffers at the multiplexers on each end of the line can be used to offset delay variation by queuing sent and received packets. The length of the queues must balance the need to regulate the rate of transmission with the need to limit overall delay, as larger buffers result in increased latency.

- For jitter-prone IP packet networks, traffic management tools can ensure that the teleprotection signals receive the highest transmission priority to minimize jitter.

- Standard packet-based synchronization technologies, such as 1588-2008 Precision Time Protocol (PTP) and Synchronous Ethernet (Sync-E), can help keep networks stable by maintaining a highly accurate clock source on the various network devices.

3.1.1.1.4. Teleprotection Network Requirements (IEC 61850)

The following table captures the main network metrics as based on the IEC 61850 standard.

Teleprotection Requirement	Attribute
One way maximum delay	4-10 ms
Asymetric delay required	Yes
Maximum jitter	less than 250 us (750 us for legacy IED)
Topology	Point to point, point to Multi-point
Availability	99.9999
precise timing required	Yes
Recovery time on node failure	less than 50ms - hitless
performance management	Yes, Mandatory
Redundancy	Yes
Packet loss	0.1% to 1%

Table 1: Teleprotection network requirements

3.1.1.1.5. Inter-Trip Protection scheme

"Inter-tripping" is the signal-controlled tripping of a circuit breaker to complete the isolation of a circuit or piece of apparatus in concert with the tripping of other circuit breakers.

Inter-Trip protection Requirement	Attribute
One way maximum delay	5 ms
Asymetric delay required	No
Maximum jitter	Not critical
Topology	Point to point, point to Multi-point
Bandwidth	64 Kbps
Availability	99.9999
precise timing required	Yes
Recovery time on node failure	less than 50ms - hitless
performance management	Yes, Mandatory
Redundancy	Yes
Packet loss	0.1%

Table 2: Inter-Trip protection network requirements

3.1.1.1.6. Current Differential Protection Scheme

Current differential protection is commonly used for line protection, and is typical for protecting parallel circuits. At both end of the lines the current is measured by the differential relays, and both relays will trip the circuit breaker if the current going into the line does not equal the current going out of the line. This type of protection scheme assumes some form of communications being present between the relays at both end of the line, to allow both relays to compare measured current values. Line differential protection schemes assume a very low telecommunications delay between both relays, often as low as 5ms. Moreover, as those systems are often not time-synchronized, they also assume symmetric telecommunications paths with constant delay, which allows comparing current measurement values taken at the exact same time.

Current Differential protection Requirement	Attribute
One way maximum delay	5 ms
Asymetric delay Required	Yes
Maximum jitter	less than 250 us (750us for legacy IED)
Topology	Point to point, point to Multi-point
Bandwidth	64 Kbps
Availability	99.9999
precise timing required	Yes
Recovery time on node failure	less than 50ms - hitless
performance management	Yes, Mandatory
Redundancy	Yes
Packet loss	0.1%

Table 3: Current Differential Protection metrics

3.1.1.1.7. Distance Protection Scheme

Distance (Impedance Relay) protection scheme is based on voltage and current measurements. The network metrics are similar (but not identical to) Current Differential protection.

Distance protection Requirement	Attribute
One way maximum delay	5 ms
Asymetric delay Required	No
Maximum jitter	Not critical
Topology	Point to point, point to Multi-point
Bandwidth	64 Kbps
Availability	99.9999
precise timing required	Yes
Recovery time on node failure	less than 50ms - hitless
performance management	Yes, Mandatory
Redundancy	Yes
Packet loss	0.1%

Table 4: Distance Protection requirements

3.1.1.1.8. Inter-Substation Protection Signaling

This use case describes the exchange of Sampled Value and/or GOOSE (Generic Object Oriented Substation Events) message between Intelligent Electronic Devices (IED) in two substations for protection and tripping coordination. The two IEDs are in a master-slave mode.

The Current Transformer or Voltage Transformer (CT/VT) in one substation sends the sampled analog voltage or current value to the Merging Unit (MU) over hard wire. The MU sends the time-synchronized 61850-9-2 sampled values to the slave IED. The slave IED forwards the information to the Master IED in the other substation. The master IED makes the determination (for example based on sampled value differentials) to send a trip command to the originating IED. Once the slave IED/Relay receives the GOOSE trip for breaker tripping, it opens the breaker. It then sends a confirmation message back to the master. All data exchanges between IEDs are either through Sampled Value and/or GOOSE messages.

Inter-Substation protection Requirement	Attribute
One way maximum delay	5 ms
Asymetric delay Required	No
Maximum jitter	Not critical
Topology	Point to point, point to Multi-point
Bandwidth	64 Kbps
Availability	99.9999
precise timing required	Yes
Recovery time on node failure	less than 50ms - hitless
performance management	Yes, Mandatory
Redundancy	Yes
Packet loss	1%

Table 5: Inter-Substation Protection requirements

3.1.1.2. Intra-Substation Process Bus Communications

This use case describes the data flow from the CT/VT to the IEDs in the substation via the MU. The CT/VT in the substation send the analog voltage or current values to the MU over hard wire. The MU converts the analog values into digital format (typically time-synchronized Sampled Values as specified by IEC 61850-9-2) and sends them to the IEDs in the substation. The GPS Master Clock can send

1PPS or IRIG-B format to the MU through a serial port or IEEE 1588 protocol via a network. Process bus communication using 61850 simplifies connectivity within the substation and removes the requirement for multiple serial connections and removes the slow serial bus architectures that are typically used. This also ensures increased flexibility and increased speed with the use of multicast messaging between multiple devices.

Intra-Substation protection Requirement	Attribute
One way maximum delay	5 ms
Asymetric delay Required	No
Maximum jitter	Not critical
Topology	Point to point, point to Multi-point
Bandwidth	64 Kbps
Availability	99.9999
precise timing required	Yes
Recovery time on Node failure	less than 50ms - hitless
performance management	Yes, Mandatory
Redundancy	Yes - No
Packet loss	0.1%

Table 6: Intra-Substation Protection requirements

3.1.1.3. Wide Area Monitoring and Control Systems

The application of synchrophasor measurement data from Phasor Measurement Units (PMU) to Wide Area Monitoring and Control Systems promises to provide important new capabilities for improving system stability. Access to PMU data enables more timely situational awareness over larger portions of the grid than what has been possible historically with normal SCADA (Supervisory Control and Data Acquisition) data. Handling the volume and real-time nature of synchrophasor data presents unique challenges for existing application architectures. Wide Area management System (WAMS) makes it possible for the condition of the bulk power system to be observed and understood in real-time so that protective, preventative, or corrective action can be taken. Because of the very high sampling rate of measurements and the strict requirement for time synchronization of the samples, WAMS has stringent telecommunications requirements in an IP network that are captured in the following table:

WAMS Requirement	Attribute
One way maximum delay	50 ms
Asymetric delay Required	No
Maximum jitter	Not critical
Topology	Point to point, point to Multi-point, Multi-point to Multi-point
Bandwidth	100 Kbps
Availability	99.9999
precise timing required	Yes
Recovery time on Node failure	less than 50ms - hitless
performance management	Yes, Mandatory
Redundancy	Yes
Packet loss	1%
Consecutive Packet Loss	At least 1 packet per application cycle must be received.

Table 7: WAMS Special Communication Requirements

3.1.1.4. IEC 61850 WAN engineering guidelines requirement classification

The IEC (International Electrotechnical Commission) has published a Technical Report which offers guidelines on how to define and deploy Wide Area Networks for the interconnections of electric substations, generation plants and SCADA operation centers. The IEC 61850-90-12 is providing a classification of WAN communication requirements into 4 classes. Table 8 summarizes these requirements:

WAN Requirement	Class WA	Class WB	Class WC	Class WD
Application field	EHV (Extra High Voltage)	HV (High Voltage)	MV (Medium Voltage)	General purpose
Latency	5 ms	10 ms	100 ms	> 100 ms
Jitter	10 us	100 us	1 ms	10 ms
Latency Asymetry	100 us	1 ms	10 ms	100 ms
Time Accuracy	1 us	10 us	100 us	10 to 100 ms
Bit Error rate	10-7 to 10-6	10-5 to 10-4	10-3	
Unavailability	10-7 to 10-6	10-5 to 10-4	10-3	
Recovery delay	Zero	50 ms	5 s	50 s
Cyber security	extremely high	High	Medium	Medium

Table 8: 61850-90-12 Communication Requirements; Courtesy of IEC

3.1.2. Generation Use Case

Energy generation systems are complex infrastructures that require control of both the generated power and the generation infrastructure.

3.1.2.1. Control of the Generated Power

The electrical power generation frequency must be maintained within a very narrow band. Deviations from the acceptable frequency range are detected and the required signals are sent to the power plants for frequency regulation.

Automatic Generation Control (AGC) is a system for adjusting the power output of generators at different power plants, in response to changes in the load.

FCAG (Frequency Control Automatic Generation) Requirement	Attribute
One way maximum delay	500 ms
Asymetric delay Required	No
Maximum jitter	Not critical
Topology	Point to point
Bandwidth	20 Kbps
Availability	99.999
precise timing required	Yes
Recovery time on Node failure	N/A
performance management	Yes, Mandatory
Redundancy	Yes
Packet loss	1%

Table 9: FCAG Communication Requirements

3.1.2.2. Control of the Generation Infrastructure

The control of the generation infrastructure combines requirements from industrial automation systems and energy generation systems. This section considers the use case of the control of the generation infrastructure of a wind turbine.

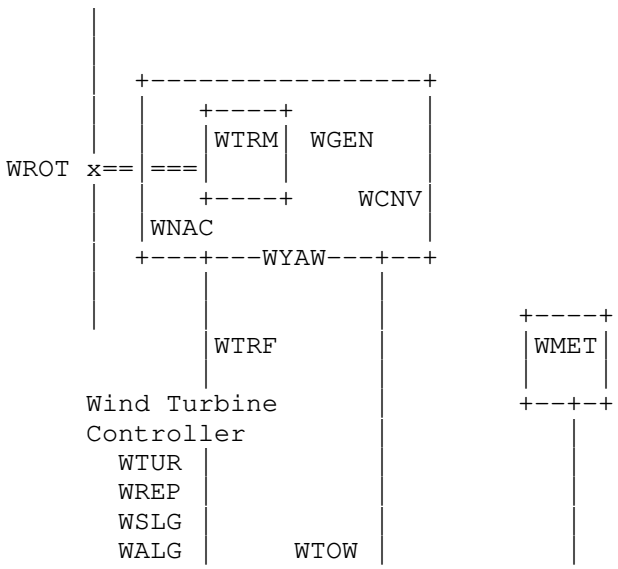

Figure 1: Wind Turbine Control Network

Figure 1 presents the subsystems that operate a wind turbine. These subsystems include

- WROT (Rotor Control)
- WNAC (Nacelle Control) (nacelle: housing containing the generator)
- WTRM (Transmission Control)
- WGEN (Generator)
- WYAW (Yaw Controller) (of the tower head)
- WCNV (In-Turbine Power Converter)
- WMET (External Meteorological Station providing real time information to the controllers of the tower)

Traffic characteristics relevant for the network planning and dimensioning process in a wind turbine scenario are listed below. The values in this section are based mainly on the relevant references [Ahm14] and [Spe09]. Each logical node (Figure 1) is a part of the metering network and produces analog measurements and status information which must comply with their respective data rate constraints.

Subsystem	Sensor Count	Analog Sample Count	Data Rate (bytes/sec)	Status Sample Count	Data rate (bytes/sec)
WROT	14	9	642	5	10
WTRM	18	10	2828	8	16
WGEN	14	12	73764	2	4
WCNV	14	12	74060	2	4
WTRF	12	5	73740	2	4
WNAC	12	9	112	3	6
WYAW	7	8	220	4	8
WTOW	4	1	8	3	6
WMET	7	7	228	-	-

Table 10: Wind Turbine Data Rate Constraints

Quality of Service (QoS) constraints for different services are presented in Table 11. These constraints are defined by IEEE 1646 standard [IEEE1646] and IEC 61400 standard [IEC61400].

Service	Latency	Reliability	Packet Loss Rate
Analogue measure	16 ms	99.99%	< 10-6
Status information	16 ms	99.99%	< 10-6
Protection traffic	4 ms	100.00%	< 10-9
Reporting and logging	1 s	99.99%	< 10-6
Video surveillance	1 s	99.00%	No specific requirement
Internet connection	60 min	99.00%	No specific requirement
Control traffic	16 ms	100.00%	< 10-9
Data polling	16 ms	99.99%	< 10-6

Table 11: Wind Turbine Reliability and Latency Constraints

3.1.2.2.1. Intra-Domain Network Considerations

A wind turbine is composed of a large set of subsystems including sensors and actuators which require time-critical operation. The reliability and latency constraints of these different subsystems is shown in Table 11. These subsystems are connected to an intra-domain network which is used to monitor and control the operation of the turbine and connect it to the SCADA subsystems. The different

components are interconnected using fiber optics, industrial buses, industrial Ethernet, EtherCat, or a combination of them. Industrial signaling and control protocols such as Modbus, Profibus, Profinet and EtherCat are used directly on top of the Layer 2 transport or encapsulated over TCP/IP.

The Data collected from the sensors and condition monitoring systems is multiplexed onto fiber cables for transmission to the base of the tower, and to remote control centers. The turbine controller continuously monitors the condition of the wind turbine and collects statistics on its operation. This controller also manages a large number of switches, hydraulic pumps, valves, and motors within the wind turbine.

There is usually a controller both at the bottom of the tower and in the nacelle. The communication between these two controllers usually takes place using fiber optics instead of copper links. Sometimes, a third controller is installed in the hub of the rotor and manages the pitch of the blades. That unit usually communicates with the nacelle unit using serial communications.

3.1.2.2.2. Inter-Domain network considerations

A remote control center belonging to a grid operator regulates the power output, enables remote actuation, and monitors the health of one or more wind parks in tandem. It connects to the local control center in a wind park over the Internet (Figure 2) via firewalls at both ends. The AS path between the local control center and the Wind Park typically involves several ISPs at different tiers. For example, a remote control center in Denmark can regulate a wind park in Greece over the normal public AS path between the two locations.

The remote control center is part of the SCADA system, setting the desired power output to the wind park and reading back the result once the new power output level has been set. Traffic between the remote control center and the wind park typically consists of protocols like IEC 60870-5-104 [IEC-60870-5-104], OPC XML-DA [OPCXML], Modbus [MODBUS], and SNMP [RFC3411]. At the time of this writing, traffic flows between the wind farm and the remote control center are best effort. QoS requirements are not strict, so no SLAs or service provisioning mechanisms (e.g., VPN) are employed. In case of events like equipment failure, tolerance for alarm delay is on the order of minutes, due to redundant systems already in place.

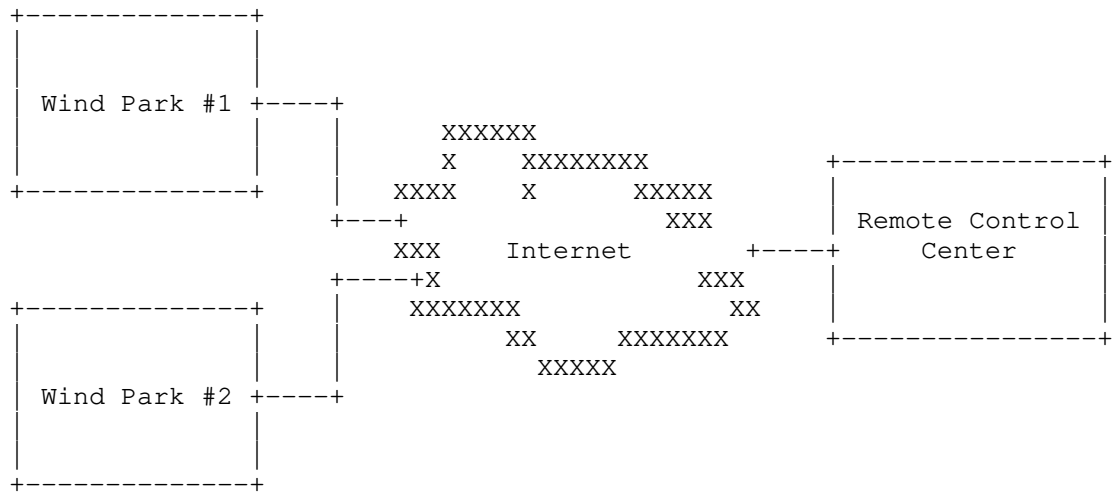

Figure 2: Wind Turbine Control via Internet

Future use cases will require bounded latency, bounded jitter and extraordinary low packet loss for inter-domain traffic flows due to the softwarization and virtualization of core wind farm equipment (e.g. switches, firewalls and SCADA server components). These factors will create opportunities for service providers to install new services and dynamically manage them from remote locations. For example, to enable fail-over of a local SCADA server, a SCADA server in another wind farm site (under the administrative control of the same operator) could be utilized temporarily (Figure 3). In that case local traffic would be forwarded to the remote SCADA server and existing intra-domain QoS and timing parameters would have to be met for inter-domain traffic flows.

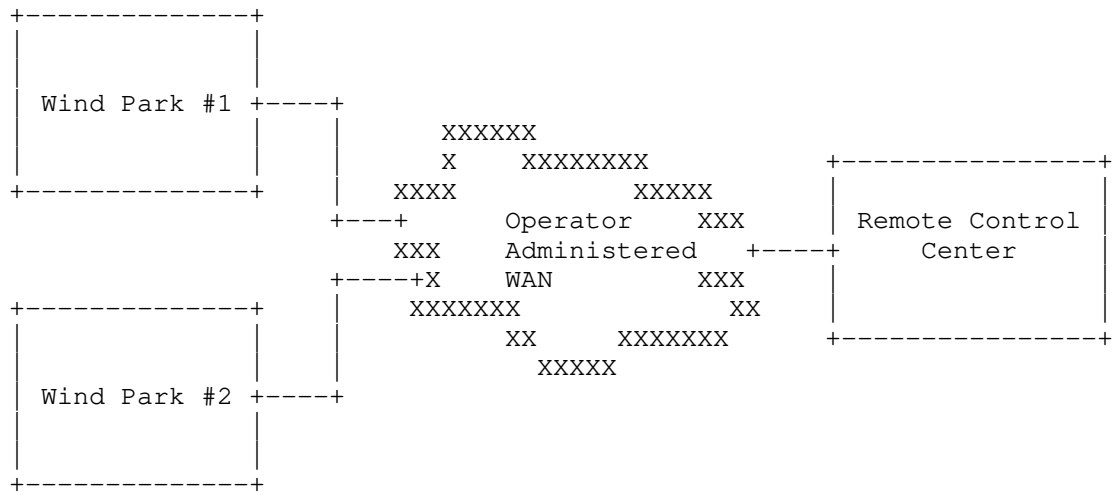

Figure 3: Wind Turbine Control via Operator Administered WAN

3.1.3. Distribution use case

3.1.3.1. Fault Location Isolation and Service Restoration (FLISR)

Fault Location, Isolation, and Service Restoration (FLISR) refers to the ability to automatically locate the fault, isolate the fault, and restore service in the distribution network. This will likely be the first widespread application of distributed intelligence in the grid.

Static power switch status (open/closed) in the network dictates the power flow to secondary substations. Reconfiguring the network in the event of a fault is typically done manually on site to energize/ de-energize alternate paths. Automating the operation of substation switchgear allows the flow of power to be altered automatically under fault conditions.

FLISR can be managed centrally from a Distribution Management System (DMS) or executed locally through distributed control via intelligent switches and fault sensors.

FLISR Requirement	Attribute
One way maximum delay	80 ms
Asymetric delay Required	No
Maximum jitter	40 ms
Topology	Point to point, point to Multi-point, Multi-point to Multi-point
Bandwidth	64 Kbps
Availability	99.9999
precise timing required	Yes
Recovery time on Node failure	Depends on customer impact
performance management	Yes, Mandatory
Redundancy	Yes
Packet loss	0.1%

Table 12: FLISR Communication Requirements

3.2. Electrical Utilities Today

Many utilities still rely on complex environments formed of multiple application-specific proprietary networks, including TDM networks.

In this kind of environment there is no mixing of OT and IT applications on the same network, and information is siloed between operational areas.

Specific calibration of the full chain is required, which is costly.

This kind of environment prevents utility operations from realizing the operational efficiency benefits, visibility, and functional integration of operational information across grid applications and data networks.

In addition, there are many security-related issues as discussed in the following section.

3.2.1. Security Current Practices and Limitations

Grid monitoring and control devices are already targets for cyber attacks, and legacy telecommunications protocols have many intrinsic network-related vulnerabilities. For example, DNP3, Modbus,

PROFIBUS/PROFINET, and other protocols are designed around a common paradigm of request and respond. Each protocol is designed for a master device such as an HMI (Human Machine Interface) system to send commands to subordinate slave devices to retrieve data (reading inputs) or control (writing to outputs). Because many of these protocols lack authentication, encryption, or other basic security measures, they are prone to network-based attacks, allowing a malicious actor or attacker to utilize the request-and-respond system as a mechanism for command-and-control like functionality. Specific security concerns common to most industrial control, including utility telecommunication protocols include the following:

- Network or transport errors (e.g. malformed packets or excessive latency) can cause protocol failure.
- Protocol commands may be available that are capable of forcing slave devices into inoperable states, including powering-off devices, forcing them into a listen-only state, disabling alarming.
- Protocol commands may be available that are capable of restarting communications and otherwise interrupting processes.
- Protocol commands may be available that are capable of clearing, erasing, or resetting diagnostic information such as counters and diagnostic registers.
- Protocol commands may be available that are capable of requesting sensitive information about the controllers, their configurations, or other need-to-know information.
- Most protocols are application layer protocols transported over TCP; therefore it is easy to transport commands over non-standard ports or inject commands into authorized traffic flows.
- Protocol commands may be available that are capable of broadcasting messages to many devices at once (i.e. a potential DoS).
- Protocol commands may be available to query the device network to obtain defined points and their values (i.e. a configuration scan).
- Protocol commands may be available that will list all available function codes (i.e. a function scan).

These inherent vulnerabilities, along with increasing connectivity between IT an OT networks, make network-based attacks very feasible.

Simple injection of malicious protocol commands provides control over the target process. Altering legitimate protocol traffic can also alter information about a process and disrupt the legitimate controls that are in place over that process. A man-in-the-middle attack could provide both control over a process and misrepresentation of data back to operator consoles.

3.3. Electrical Utilities Future

The business and technology trends that are sweeping the utility industry will drastically transform the utility business from the way it has been for many decades. At the core of many of these changes is a drive to modernize the electrical grid with an integrated telecommunications infrastructure. However, interoperability concerns, legacy networks, disparate tools, and stringent security requirements all add complexity to the grid transformation. Given the range and diversity of the requirements that should be addressed by the next generation telecommunications infrastructure, utilities need to adopt a holistic architectural approach to integrate the electrical grid with digital telecommunications across the entire power delivery chain.

The key to modernizing grid telecommunications is to provide a common, adaptable, multi-service network infrastructure for the entire utility organization. Such a network serves as the platform for current capabilities while enabling future expansion of the network to accommodate new applications and services.

To meet this diverse set of requirements, both today and in the future, the next generation utility telecommunnications network will be based on open-standards-based IP architecture. An end-to-end IP architecture takes advantage of nearly three decades of IP technology development, facilitating interoperability and device management across disparate networks and devices, as it has been already demonstrated in many mission-critical and highly secure networks.

IPv6 is seen as a future telecommunications technology for the Smart Grid; the IEC (International Electrotechnical Commission) and different National Committees have mandated a specific adhoc group (AHG8) to define the migration strategy to IPv6 for all the IEC TC57 power automation standards. The AHG8 has finalised the work on the migration strategy and the following Technical Report has been issued: IEC TR 62357-200:2015: Guidelines for migration from Internet Protocol version 4 (IPv4) to Internet Protocol version 6 (IPv6).

Cloud-based SCADA systems will control and monitor the critical and non-critical subsystems of generation systems, for example wind farms.

3.3.1. Migration to Packet-Switched Network

Throughout the world, utilities are increasingly planning for a future based on smart grid applications requiring advanced telecommunications systems. Many of these applications utilize packet connectivity for communicating information and control signals across the utility's Wide Area Network (WAN), made possible by technologies such as multiprotocol label switching (MPLS). The data that traverses the utility WAN includes:

- Grid monitoring, control, and protection data
- Non-control grid data (e.g. asset data for condition-based monitoring)
- Physical safety and security data (e.g. voice and video)
- Remote worker access to corporate applications (voice, maps, schematics, etc.)
- Field area network backhaul for smart metering, and distribution grid management
- Enterprise traffic (email, collaboration tools, business applications)

WANs support this wide variety of traffic to and from substations, the transmission and distribution grid, generation sites, between control centers, and between work locations and data centers. To maintain this rapidly expanding set of applications, many utilities are taking steps to evolve present time-division multiplexing (TDM) based and frame relay infrastructures to packet systems. Packet-based networks are designed to provide greater functionalities and higher levels of service for applications, while continuing to deliver reliability and deterministic (real-time) traffic support.

3.3.2. Telecommunications Trends

These general telecommunications topics are in addition to the use cases that have been addressed so far. These include both current and future telecommunications related topics that should be factored into the network architecture and design.

3.3.2.1. General Telecommunications Requirements

- IP Connectivity everywhere
- Monitoring services everywhere and from different remote centers

- Move services to a virtual data center
- Unify access to applications / information from the corporate network
- Unify services
- Unified Communications Solutions
- Mix of fiber and microwave technologies - obsolescence of SONET/ SDH or TDM
- Standardize grid telecommunications protocol to opened standard to ensure interoperability
- Reliable Telecommunications for Transmission and Distribution Substations
- IEEE 1588 time synchronization Client / Server Capabilities
- Integration of Multicast Design
- QoS Requirements Mapping
- Enable Future Network Expansion
- Substation Network Resilience
- Fast Convergence Design
- Scalable Headend Design
- Define Service Level Agreements (SLA) and Enable SLA Monitoring
- Integration of 3G/4G Technologies and future technologies
- Ethernet Connectivity for Station Bus Architecture
- Ethernet Connectivity for Process Bus Architecture
- Protection, teleprotection and PMU (Phaser Measurement Unit) on IP

3.3.2.2. Specific Network topologies of Smart Grid Applications

Utilities often have very large private telecommunications networks. It covers an entire territory / country. The main purpose of the network, until now, has been to support transmission network monitoring, control, and automation, remote control of generation

sites, and providing FCAPS (Fault, Configuration, Accounting, Performance, Security) services from centralized network operation centers.

Going forward, one network will support operation and maintenance of electrical networks (generation, transmission, and distribution), voice and data services for ten of thousands of employees and for exchange with neighboring interconnections, and administrative services. To meet those requirements, utility may deploy several physical networks leveraging different technologies across the country: an optical network and a microwave network for instance. Each protection and automatism system between two points has two telecommunications circuits, one on each network. Path diversity between two substations is key. Regardless of the event type (hurricane, ice storm, etc.), one path needs to stay available so the system can still operate.

In the optical network, signals are transmitted over more than tens of thousands of circuits using fiber optic links, microwave and telephone cables. This network is the nervous system of the utility's power transmission operations. The optical network represents ten of thousands of km of cable deployed along the power lines, with individual runs as long as 280 km.

3.3.2.3. Precision Time Protocol

Some utilities do not use GPS clocks in generation substations. One of the main reasons is that some of the generation plants are 30 to 50 meters deep under ground and the GPS signal can be weak and unreliable. Instead, atomic clocks are used. Clocks are synchronized amongst each other. Rubidium clocks provide clock and 1ms timestamps for IRIG-B.

Some companies plan to transition to the Precision Time Protocol (PTP, [IEEE1588]), distributing the synchronization signal over the IP/MPLS network. PTP provides a mechanism for synchronizing the clocks of participating nodes to a high degree of accuracy and precision.

PTP operates based on the following assumptions:

It is assumed that the network eliminates cyclic forwarding of PTP messages within each communication path (e.g. by using a spanning tree protocol).

PTP is tolerant of an occasional missed message, duplicated message, or message that arrived out of order. However, PTP assumes that such impairments are relatively rare.

PTP was designed assuming a multicast communication model, however PTP also supports a unicast communication model as long as the behavior of the protocol is preserved.

Like all message-based time transfer protocols, PTP time accuracy is degraded by delay asymmetry in the paths taken by event messages. Asymmetry is not detectable by PTP, however, if such delays are known a priori, PTP can correct for asymmetry.

IEC 61850 defines the use of IEC/IEEE 61850-9-3:2016. The title is: Precision time protocol profile for power utility automation. It is based on Annex B/IEC 62439 which offers the support of redundant attachment of clocks to Parallel Redundancy Protocol (PRP) and High-availability Seamless Redundancy (HSR) networks.

3.3.3. Security Trends in Utility Networks

Although advanced telecommunications networks can assist in transforming the energy industry by playing a critical role in maintaining high levels of reliability, performance, and manageability, they also introduce the need for an integrated security infrastructure. Many of the technologies being deployed to support smart grid projects such as smart meters and sensors can increase the vulnerability of the grid to attack. Top security concerns for utilities migrating to an intelligent smart grid telecommunications platform center on the following trends:

- Integration of distributed energy resources
- Proliferation of digital devices to enable management, automation, protection, and control
- Regulatory mandates to comply with standards for critical infrastructure protection
- Migration to new systems for outage management, distribution automation, condition-based maintenance, load forecasting, and smart metering
- Demand for new levels of customer service and energy management

This development of a diverse set of networks to support the integration of microgrids, open-access energy competition, and the use of network-controlled devices is driving the need for a converged security infrastructure for all participants in the smart grid, including utilities, energy service providers, large commercial and industrial, as well as residential customers. Securing the assets of electric power delivery systems (from the control center to the

substation, to the feeders and down to customer meters) requires an end-to-end security infrastructure that protects the myriad of telecommunications assets used to operate, monitor, and control power flow and measurement.

"Cyber security" refers to all the security issues in automation and telecommunications that affect any functions related to the operation of the electric power systems. Specifically, it involves the concepts of:

- Integrity : data cannot be altered undetectably
- Authenticity (data origin authentication): the telecommunications parties involved must be validated as genuine
- Authorization : only requests and commands from the authorized users can be accepted by the system
- Confidentiality : data must not be accessible to any unauthenticated users

When designing and deploying new smart grid devices and telecommunications systems, it is imperative to understand the various impacts of these new components under a variety of attack situations on the power grid. Consequences of a cyber attack on the grid telecommunications network can be catastrophic. This is why security for smart grid is not just an ad hoc feature or product, it’s a complete framework integrating both physical and Cyber security requirements and covering the entire smart grid networks from generation to distribution. Security has therefore become one of the main foundations of the utility telecom network architecture and must be considered at every layer with a defense-in-depth approach. Migrating to IP based protocols is key to address these challenges for two reasons:

- IP enables a rich set of features and capabilities to enhance the security posture
- IP is based on open standards, which allows interoperability between different vendors and products, driving down the costs associated with implementing security solutions in OT networks.

Securing OT (Operation technology) telecommunications over packet-switched IP networks follow the same principles that are foundational for securing the IT infrastructure, i.e., consideration must be given to enforcing electronic access control for both person-to-machine and machine-to-machine communications, and providing the appropriate

levels of data privacy, device and platform integrity, and threat detection and mitigation.

3.4. Electrical Utilities Asks

- Mixed L2 and L3 topologies
- Deterministic behavior
- Bounded latency and jitter
- Tight feedback intervals
- High availability, low recovery time
- Redundancy, low packet loss
- Precise timing
- Centralized computing of deterministic paths
- Distributed configuration may also be useful

4. Building Automation Systems

4.1. Use Case Description

A Building Automation System (BAS) manages equipment and sensors in a building for improving residents' comfort, reducing energy consumption, and responding to failures and emergencies. For example, the BAS measures the temperature of a room using sensors and then controls the HVAC (heating, ventilating, and air conditioning) to maintain a set temperature and minimize energy consumption.

A BAS primarily performs the following functions:

- Periodically measures states of devices, for example humidity and illuminance of rooms, open/close state of doors, FAN speed, etc.
- Stores the measured data.
- Provides the measured data to BAS systems and operators.
- Generates alarms for abnormal state of devices.
- Controls devices (e.g. turn off room lights at 10:00 PM).

4.2. Building Automation Systems Today

4.2.1. BAS Architecture

A typical BAS architecture of today is shown in Figure 4.

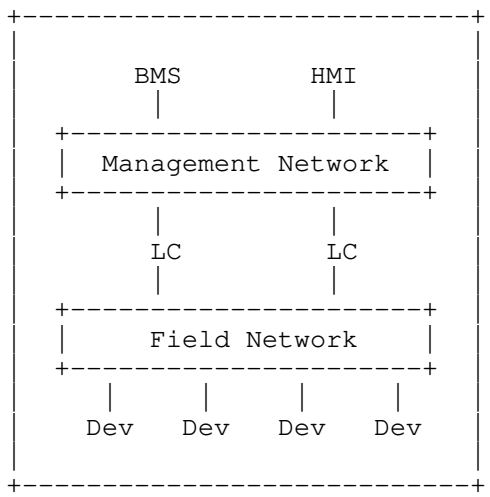

BMS := Building Management Server
HMI := Human Machine Interface
LC := Local Controller

Figure 4: BAS architecture

There are typically two layers of network in a BAS. The upper one is called the Management Network and the lower one is called the Field Network. In management networks an IP-based communication protocol is used, while in field networks non-IP based communication protocols ("field protocols") are mainly used. Field networks have specific timing requirements, whereas management networks can be best-effort.

A Human Machine Interface (HMI) is typically a desktop PC used by operators to monitor and display device states, send device control commands to Local Controllers (LCs), and configure building schedules (for example "turn off all room lights in the building at 10:00 PM").

A Building Management Server (BMS) performs the following operations.

- Collect and store device states from LCs at regular intervals.

- Send control values to LCs according to a building schedule.

- Send an alarm signal to operators if it detects abnormal devices states.

The BMS and HMI communicate with LCs via IP-based "management protocols" (see standards [bacnetip], [knx]).

A LC is typically a Programmable Logic Controller (PLC) which is connected to several tens or hundreds of devices using "field protocols". An LC performs the following kinds of operations:

- Measure device states and provide the information to BMS or HMI.

- Send control values to devices, unilaterally or as part of a feedback control loop.

There are many field protocols used at the time of this writing; some are standards-based and others are proprietary (see standards [lontalk], [modbus], [profibus] and [flnet]). The result is that BASs have multiple MAC/PHY modules and interfaces. This makes BASs more expensive, slower to develop, and can result in "vendor lock-in" with multiple types of management applications.

4.2.2. BAS Deployment Model

An example BAS for medium or large buildings is shown in Figure 5. The physical layout spans multiple floors, and there is a monitoring room where the BAS management entities are located. Each floor will have one or more LCs depending upon the number of devices connected to the field network.

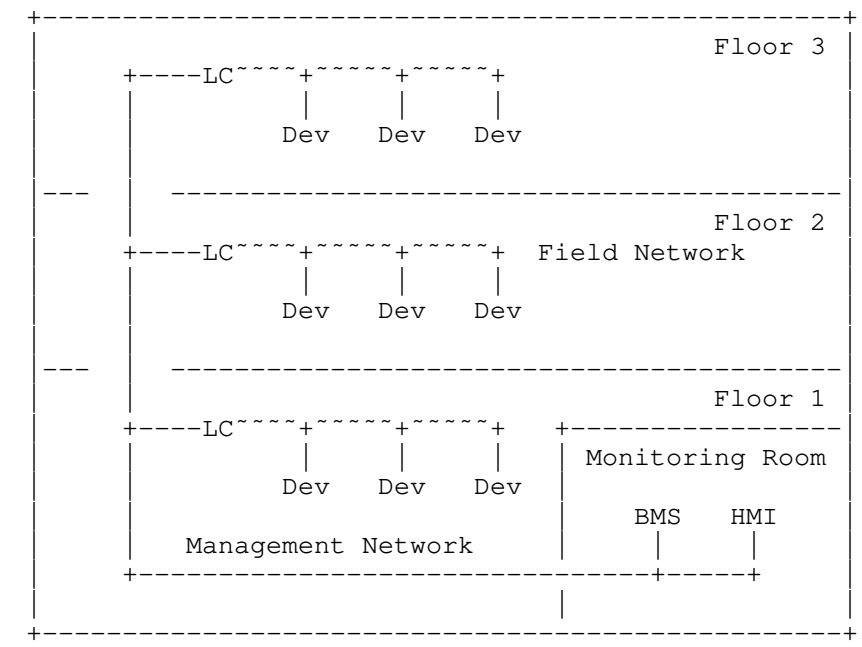

Figure 5: BAS Deployment model for Medium/Large Buildings

Each LC is connected to the monitoring room via the Management network, and the management functions are performed within the building. In most cases, fast Ethernet (e.g. 100BASE-T) is used for the management network. Since the management network is non-realtime, use of Ethernet without quality of service is sufficient for today's deployment.

In the field network a variety of physical interfaces such as RS232C and RS485 are used, which have specific timing requirements. Thus if a field network is to be replaced with an Ethernet or wireless network, such networks must support time-critical deterministic flows.

In Figure 6, another deployment model is presented in which the management system is hosted remotely. This is becoming popular for small office and residential buildings in which a standalone monitoring system is not cost-effective.

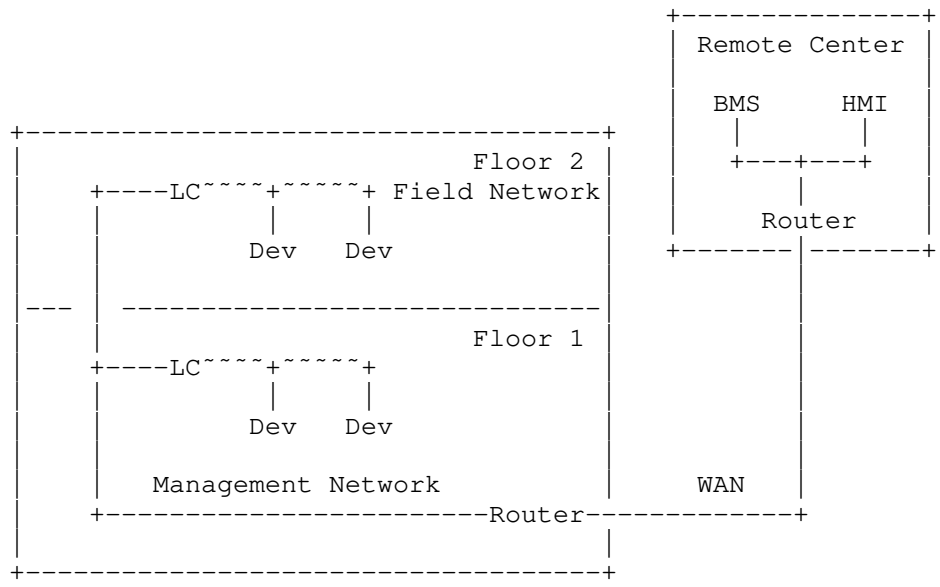

Figure 6: Deployment model for Small Buildings

Some interoperability is possible today in the Management Network, but not in today's field networks due to their non-IP-based design.

4.2.3. Use Cases for Field Networks

Below are use cases for Environmental Monitoring, Fire Detection, and Feedback Control, and their implications for field network performance.

4.2.3.1. Environmental Monitoring

The BMS polls each LC at a maximum measurement interval of 100ms (for example to draw a historical chart of 1 second granularity with a 10x sampling interval) and then performs the operations as specified by the operator. Each LC needs to measure each of its several hundred sensors once per measurement interval. Latency is not critical in this scenario as long as all sensor values are completed in the measurement interval. Availability is expected to be 99.999 %.

4.2.3.2. Fire Detection

On detection of a fire, the BMS must stop the HVAC, close the fire shutters, turn on the fire sprinklers, send an alarm, etc. There are typically ~10s of sensors per LC that BMS needs to manage. In this

scenario the measurement interval is 10-50ms, the communication delay is 10ms, and the availability must be 99.9999 %.

4.2.3.3. Feedback Control

BAS systems utilize feedback control in various ways; the most time-critial is control of DC motors, which require a short feedback interval (1-5ms) with low communication delay (10ms) and jitter (1ms). The feedback interval depends on the characteristics of the device and a target quality of control value. There are typically ~10s of such devices per LC.

Communication delay is expected to be less than 10ms, jitter less than 1ms while the availability must be 99.9999% .

4.2.4. Security Considerations

When BAS field networks were developed it was assumed that the field networks would always be physically isolated from external networks and therefore security was not a concern. In today's world many BASs are managed remotely and are thus connected to shared IP networks and so security is definitely a concern, yet security features are not available in the majority of BAS field network deployments .

The management network, being an IP-based network, has the protocols available to enable network security, but in practice many BAS systems do not implement even the available security features such as device authentication or encryption for data in transit.

4.3. BAS Future

In the future more fine-grained environmental monitoring and lower energy consumption will emerge which will require more sensors and devices, thus requiring larger and more complex building networks.

Building networks will be connected to or converged with other networks (Enterprise network, Home network, and Internet).

Therefore better facilities for network management, control, reliability and security are critical in order to improve resident and operator convenience and comfort. For example the ability to monitor and control building devices via the internet would enable (for example) control of room lights or HVAC from a resident's desktop PC or phone application.

4.4. BAS Asks

The community would like to see an interoperable protocol specification that can satisfy the timing, security, availability and QoS constraints described above, such that the resulting converged network can replace the disparate field networks. Ideally this connectivity could extend to the open Internet.

This would imply an architecture that can guarantee

- Low communication delays (from <10ms to 100ms in a network of several hundred devices)
- Low jitter (< 1 ms)
- Tight feedback intervals (1ms – 10ms)
- High network availability (up to 99.9999%)
- Availability of network data in disaster scenario
- Authentication between management and field devices (both local and remote)
- Integrity and data origin authentication of communication data between field and management devices
- Confidentiality of data when communicated to a remote device

5. Wireless for Industrial Applications

5.1. Use Case Description

Wireless networks are useful for industrial applications, for example when portable, fast-moving or rotating objects are involved, and for the resource-constrained devices found in the Internet of Things (IoT).

Such network-connected sensors, actuators, control loops (etc.) typically require that the underlying network support real-time quality of service (QoS), as well as specific classes of other network properties such as reliability, redundancy, and security.

These networks may also contain very large numbers of devices, for example for factories, "big data" acquisition, and the IoT. Given the large numbers of devices installed, and the potential pervasiveness of the IoT, this is a huge and very cost-sensitive

market such that small cost reductions can save large amounts of money.

5.1.1. Network Convergence using 6TiSCH

Some wireless network technologies support real-time QoS, and are thus useful for these kinds of networks, but others do not.

This use case focuses on one specific wireless network technology which provides the required deterministic QoS, which is "IPv6 over the TSCH mode of IEEE 802.15.4e" (6TiSCH, where TSCH stands for "Time-Slotted Channel Hopping", see [I-D.ietf-6tisch-architecture], [IEEE802154], [IEEE802154e], and [RFC7554]).

There are other deterministic wireless busses and networks available today, however they are imcompatible with each other, and incompatible with IP traffic (for example [ISA100], [WirelessHART]).

Thus the primary goal of this use case is to apply 6TiSCH as a converged IP- and standards-based wireless network for industrial applications, i.e. to replace multiple proprietary and/or incompatible wireless networking and wireless network management standards.

5.1.2. Common Protocol Development for 6TiSCH

Today there are a number of protocols required by 6TiSCH which are still in development, and a second intent of this use case is to highlight the ways in which these "missing" protocols share goals in common with DetNet. Thus it is possible that some of the protocol technology developed for DetNet will also be applicable to 6TiSCH.

These protocol goals are identified here, along with their relationship to DetNet. It is likely that ultimately the resulting protocols will not be identical, but will share design principles which contribute to the eficiency of enabling both DetNet and 6TiSCH.

One such commonality is that although at a different time scale, in both TSN [IEEE802.1TSNTG] and TSCH a packet crosses the network from node to node follows a precise schedule, as a train that leaves intermediate stations at precise times along its path. This kind of operation reduces collisions, saves energy, and enables engineering the network for deterministic properties.

Another commonality is remote monitoring and scheduling management of a TSCH network by a Path Computation Element (PCE) and Network Management Entity (NME). The PCE/NME manage timeslots and device resources in a manner that minimizes the interaction with and the

load placed on resource-constrained devices. For example, a tiny IoT device may have just enough buffers to store one or a few IPv6 packets, and will have limited bandwidth between peers such that it can maintain only a small amount of peer information, and will not be able to store many packets waiting to be forwarded. It is advantageous then for it to only be required to carry out the specific behavior assigned to it by the PCE/NME (as opposed to maintaining its own IP stack, for example).

It is possible that there will be some peer-to-peer communication, for example the PCE may communicate only indirectly with some devices in order to enable hierarchical configuration of the system.

6TiSCH depends on [PCE] and [I-D.ietf-detnet-architecture].

6TiSCH also depends on the fact that DetNet will maintain consistency with [IEEE802.1TSNTG].

5.2. Wireless Industrial Today

Today industrial wireless is accomplished using multiple deterministic wireless networks which are incompatible with each other and with IP traffic.

6TiSCH is not yet fully specified, so it cannot be used in today's applications.

5.3. Wireless Industrial Future

5.3.1. Unified Wireless Network and Management

DetNet and 6TiSCH together can enable converged transport of deterministic and best-effort traffic flows between real-time industrial devices and wide area networks via IP routing. A high level view of a basic such network is shown in Figure 7.

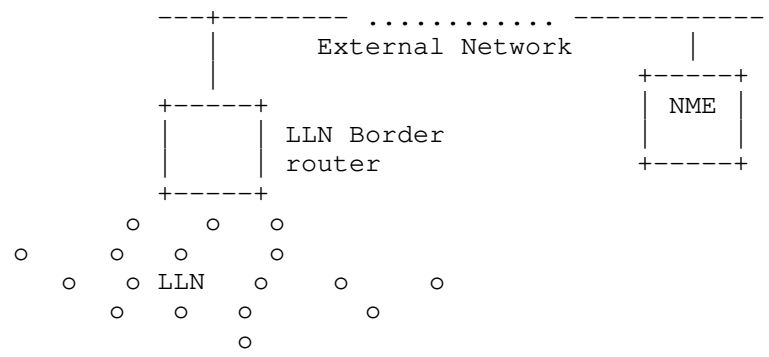

Figure 7: Basic 6TiSCH Network

Figure 8 shows a backbone router federating multiple synchronized 6TiSCH subnets into a single subnet connected to the external network.

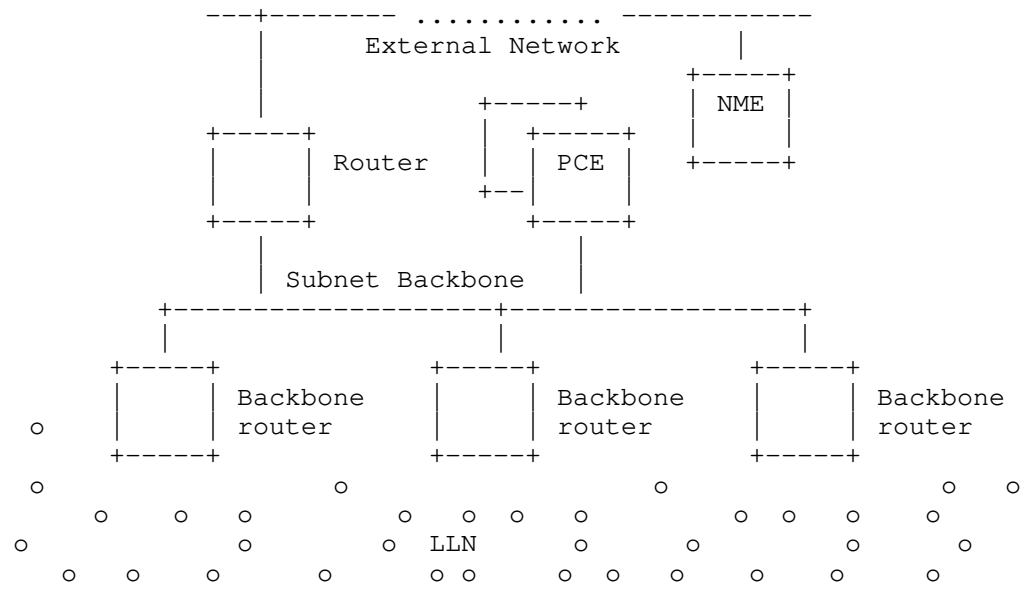

Figure 8: Extended 6TiSCH Network

The backbone router must ensure end-to-end deterministic behavior between the LLN and the backbone. This should be accomplished in conformance with the work done in [I-D.ietf-detnet-architecture] with respect to Layer-3 aspects of deterministic networks that span multiple Layer-2 domains.

The PCE must compute a deterministic path end-to-end across the TSCH network and IEEE802.1 TSN Ethernet backbone, and DetNet protocols are expected to enable end-to-end deterministic forwarding.

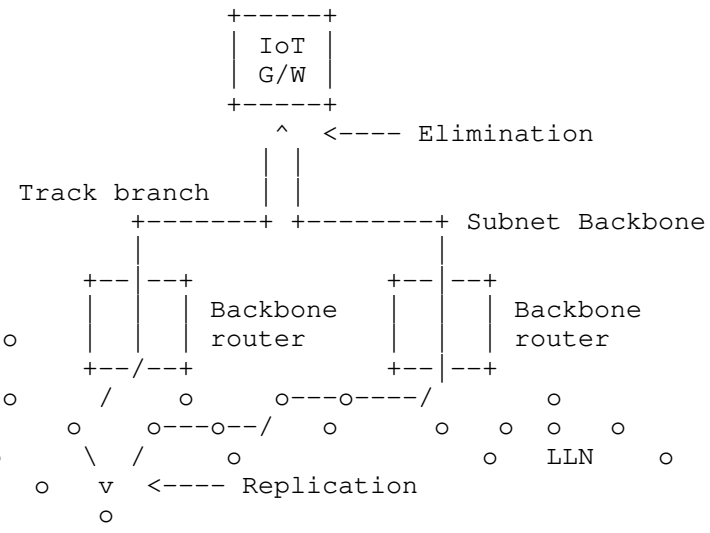

Figure 9: 6TiSCH Network with PRE

5.3.1.1. PCE and 6TiSCH ARQ Retries

6TiSCH uses the IEEE802.15.4 Automatic Repeat-reQuest (ARQ) mechanism to provide higher reliability of packet delivery. ARQ is related to packet replication and elimination because there are two independent paths for packets to arrive at the destination, and if an expected packed does not arrive on one path then it checks for the packet on the second path.

Although to date this mechanism is only used by wireless networks, this may be a technique that would be appropriate for DetNet and so aspects of the enabling protocol could be co-developed.

For example, in Figure 9, a Track is laid out from a field device in a 6TiSCH network to an IoT gateway that is located on a IEEE802.1 TSN backbone.

In ARQ the Replication function in the field device sends a copy of each packet over two different branches, and the PCE schedules each hop of both branches so that the two copies arrive in due time at the gateway. In case of a loss on one branch, hopefully the other copy

of the packet still arrives within the allocated time. If two copies make it to the IoT gateway, the Elimination function in the gateway ignores the extra packet and presents only one copy to upper layers.

At each 6TiSCH hop along the Track, the PCE may schedule more than one timeSlot for a packet, so as to support Layer-2 retries (ARQ).

In deployments at the time of this writing, a TSCH Track does not necessarily support PRE but is systematically multi-path. This means that a Track is scheduled so as to ensure that each hop has at least two forwarding solutions, and the forwarding decision is to try the preferred one and use the other in case of Layer-2 transmission failure as detected by ARQ.

5.3.2. Schedule Management by a PCE

A common feature of 6TiSCH and DetNet is the action of a PCE to configure paths through the network. Specifically, what is needed is a protocol and data model that the PCE will use to get/set the relevant configuration from/to the devices, as well as perform operations on the devices. This protocol should be developed by DetNet with consideration for its reuse by 6TiSCH. The remainder of this section provides a bit more context from the 6TiSCH side.

5.3.2.1. PCE Commands and 6TiSCH CoAP Requests

The 6TiSCH device does not expect to place the request for bandwidth between itself and another device in the network. Rather, an operation control system invoked through a human interface specifies the required traffic specification and the end nodes (in terms of latency and reliability). Based on this information, the PCE must compute a path between the end nodes and provision the network with per-flow state that describes the per-hop operation for a given packet, the corresponding timeslots, and the flow identification that enables recognizing that a certain packet belongs to a certain path, etc.

For a static configuration that serves a certain purpose for a long period of time, it is expected that a node will be provisioned in one shot with a full schedule, which incorporates the aggregation of its behavior for multiple paths. 6TiSCH expects that the programing of the schedule will be done over COAP as discussed in [I-D.ietf-6tisch-coap].

6TiSCH expects that the PCE commands will be mapped back and forth into CoAP by a gateway function at the edge of the 6TiSCH network. For instance, it is possible that a mapping entity on the backbone transforms a non-CoAP protocol such as PCEP into the RESTful

interfaces that the 6TiSCH devices support. This architecture will be refined to comply with DetNet [I-D.ietf-detnet-architecture] when the work is formalized. Related information about 6TiSCH can be found at [I-D.ietf-6tisch-6top-interface] and RPL [RFC6550].

A protocol may be used to update the state in the devices during runtime, for example if it appears that a path through the network has ceased to perform as expected, but in 6TiSCH that flow was not designed and no protocol was selected. DetNet should define the appropriate end-to-end protocols to be used in that case. The implication is that these state updates take place once the system is configured and running, i.e. they are not limited to the initial communication of the configuration of the system.

A "slotFrame" is the base object that a PCE would manipulate to program a schedule into an LLN node ([I-D.ietf-6tisch-architecture]).

The PCE should read energy data from devices and compute paths that will implement policies on how energy in devices is consumed, for instance to ensure that the spent energy does not exceeded the available energy over a period of time. Note: this statement implies that an extensible protocol for communicating device info to the PCE and enabling the PCE to act on it will be part of the DetNet architecture, however for subnets with specific protocols (e.g. CoAP) a gateway may be required.

6TiSCH devices can discover their neighbors over the radio using a mechanism such as beacons, but even though the neighbor information is available in the 6TiSCH interface data model, 6TiSCH does not describe a protocol to proactively push the neighborhood information to a PCE. DetNet should define such a protocol; one possible design alternative is that it could operate over CoAP, alternatively it could be converted to/from CoAP by a gateway. Such a protocol could carry multiple metrics, for example similar to those used for RPL operations [RFC6551]

5.3.2.2. 6TiSCH IP Interface

"6top" ([I-D.wang-6tisch-6top-sublayer]) is a logical link control sitting between the IP layer and the TSCH MAC layer which provides the link abstraction that is required for IP operations. The 6top data model and management interfaces are further discussed in [I-D.ietf-6tisch-6top-interface] and [I-D.ietf-6tisch-coap].

An IP packet that is sent along a 6TiSCH path uses the Differentiated Services Per-Hop-Behavior Group called Deterministic Forwarding, as described in [I-D.svshah-tsvwg-deterministic-forwarding].

5.3.3. 6TiSCH Security Considerations

On top of the classical requirements for protection of control signaling, it must be noted that 6TiSCH networks operate on limited resources that can be depleted rapidly in a DoS attack on the system, for instance by placing a rogue device in the network, or by obtaining management control and setting up unexpected additional paths.

5.4. Wireless Industrial Asks

6TiSCH depends on DetNet to define:

- Configuration (state) and operations for deterministic paths
- End-to-end protocols for deterministic forwarding (tagging, IP)
- Protocol for packet replication and elimination

6. Cellular Radio

6.1. Use Case Description

This use case describes the application of deterministic networking in the context of cellular telecom transport networks. Important elements include time synchronization, clock distribution, and ways of establishing time-sensitive streams for both Layer-2 and Layer-3 user plane traffic.

6.1.1. Network Architecture

Figure 10 illustrates a 3GPP-defined cellular network architecture typical at the time of this writing, which includes "Fronthaul", "Midhaul" and "Backhaul" network segments. The "Fronthaul" is the network connecting base stations (baseband processing units) to the remote radio heads (antennas). The "Midhaul" is the network inter-connecting base stations (or small cell sites). The "Backhaul" is the network or links connecting the radio base station sites to the network controller/gateway sites (i.e. the core of the 3GPP cellular network).

In Figure 10 "eNB" ("E-UTRAN Node B") is the hardware that is connected to the mobile phone network which communicates directly with mobile handsets ([TS36300]).

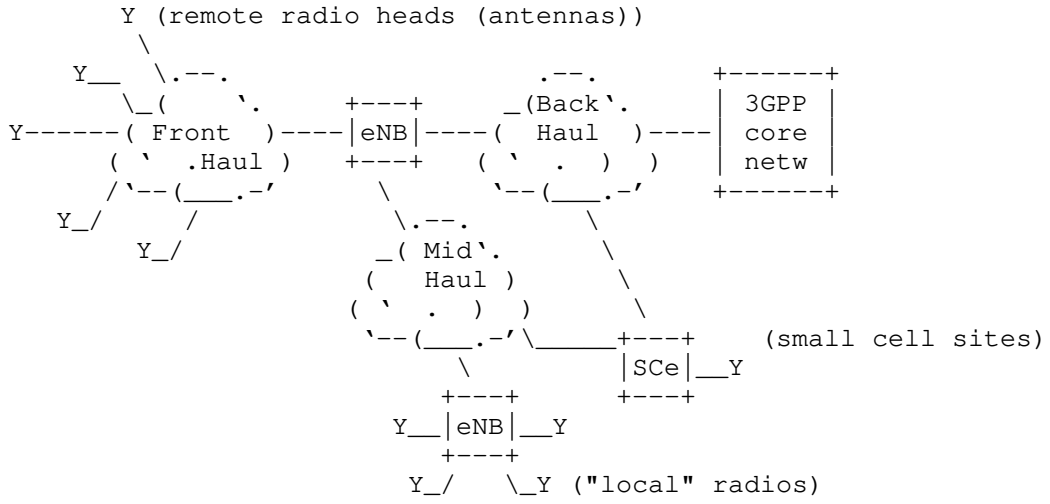

Figure 10: Generic 3GPP-based Cellular Network Architecture

6.1.2. Delay Constraints

The available processing time for Fronthaul networking overhead is limited to the available time after the baseband processing of the radio frame has completed. For example in Long Term Evolution (LTE) radio, processing of a radio frame is allocated 3ms but typically the processing uses most of it, allowing only a small fraction to be used by the Fronthaul network (e.g. up to 250us one-way delay, though the existing spec ([NGMN-fronth]) supports delay only up to 100us). This ultimately determines the distance the remote radio heads can be located from the base stations (e.g., 100us equals roughly 20 km of optical fiber-based transport). Allocation options of the available time budget between processing and transport are under heavy discussions in the mobile industry.

For packet-based transport the allocated transport time (e.g. CPRI would allow for 100us delay [CPRI]) is consumed by all nodes and buffering between the remote radio head and the baseband processing unit, plus the distance-incurred delay.

The baseband processing time and the available "delay budget" for the fronthaul is likely to change in the forthcoming "5G" due to reduced radio round trip times and other architectural and service requirements [NGMN].

The transport time budget, as noted above, places limitations on the distance that remote radio heads can be located from base stations (i.e. the link length). In the above analysis, the entire transport

time budget is assumed to be available for link propagation delay. However the transport time budget can be broken down into three components: scheduling /queueing delay, transmission delay, and link propagation delay. Using today's Fronthaul networking technology, the queuing, scheduling and transmission components might become the dominant factors in the total transport time rather than the link propagation delay. This is especially true in cases where the Fronthaul link is relatively short and it is shared among multiple Fronthaul flows, for example in indoor and small cell networks, massive MIMO antenna networks, and split Fronthaul architectures.

DetNet technology can improve this application by controlling and reducing the time required for the queuing, scheduling and transmission operations by properly assigning the network resources, thus leaving more of the transport time budget available for link propagation, and thus enabling longer link lengths. However, link length is usually a given parameter and is not a controllable network parameter, since RRH and BBU sights are usually located in predetermined locations. However, the number of antennas in an RRH sight might increase for example by adding more antennas, increasing the MIMO capability of the network or support of massive MIMO. This means increasing the number of the fronthaul flows sharing the same fronthaul link. DetNet can now control the bandwidth assignment of the fronthaul link and the scheduling of fronthaul packets over this link and provide adequate buffer provisioning for each flow to reduce the packet loss rate.

Another way in which DetNet technology can aid Fronthaul networks is by providing effective isolation from best-effort (and other classes of) traffic, which can arise as a result of network slicing in 5G networks where Fronthaul traffic generated in different network slices might have differing performance requirements. DetNet technology can also dynamically control the bandwidth assignment, scheduling and packet forwarding decisions and the buffer provisioning of the Fronthaul flows to guarantee the end-to-end delay of the Fronthaul packets and minimize the packet loss rate.

[METIS] documents the fundamental challenges as well as overall technical goals of the future 5G mobile and wireless system as the starting point. These future systems should support much higher data volumes and rates and significantly lower end-to-end latency for 100x more connected devices (at similar cost and energy consumption levels as today's system).

For Midhaul connections, delay constraints are driven by Inter-Site radio functions like Coordinated Multipoint Processing (CoMP, see [CoMP]). CoMP reception and transmission is a framework in which multiple geographically distributed antenna nodes cooperate to

improve the performance of the users served in the common cooperation area. The design principal of CoMP is to extend single-cell to multi-UE (User Equipment) transmission to a multi-cell-to-multi-UEs transmission by base station cooperation.

CoMP has delay-sensitive performance parameters, which are "midhaul latency" and "CSI (Channel State Information) reporting and accuracy". The essential feature of CoMP is signaling between eNBs, so Midhaul latency is the dominating limitation of CoMP performance. Generally, CoMP can benefit from coordinated scheduling (either distributed or centralized) of different cells if the signaling delay between eNBs is within 1-10ms. This delay requirement is both rigid and absolute because any uncertainty in delay will degrade the performance significantly.

Inter-site CoMP is one of the key requirements for 5G and is also a goal for 4.5G network architecture.

6.1.3. Time Synchronization Constraints

Fronthaul time synchronization requirements are given by [TS25104], [TS36104], [TS36211], and [TS36133]. These can be summarized for the 3GPP LTE-based networks as:

Delay Accuracy:
: +-8ns (i.e. +-1/32 Tc, where Tc is the UMTS Chip time of 1/3.84 MHz) resulting in a round trip accuracy of +-16ns. The value is this low to meet the 3GPP Timing Alignment Error (TAE) measurement requirements. Note: performance guarantees of low nanosecond values such as these are considered to be below the DetNet layer - it is assumed that the underlying implementation, e.g. the hardware, will provide sufficient support (e.g. buffering) to enable this level of accuracy. These values are maintained in the use case to give an indication of the overall application.

Timing Alignment Error:
: Timing Alignment Error (TAE) is problematic to Fronthaul networks and must be minimized. If the transport network cannot guarantee low enough TAE then additional buffering has to be introduced at the edges of the network to buffer out the jitter. Buffering is not desirable as it reduces the total available delay budget. Packet Delay Variation (PDV) requirements can be derived from TAE for packet based Fronthaul networks.

* For multiple input multiple output (MIMO) or TX diversity transmissions, at each carrier frequency, TAE shall not exceed 65 ns (i.e. 1/4 Tc).

* For intra-band contiguous carrier aggregation, with or without MIMO or TX diversity, TAE shall not exceed 130 ns (i.e. 1/2 Tc).

* For intra-band non-contiguous carrier aggregation, with or without MIMO or TX diversity, TAE shall not exceed 260 ns (i.e. one Tc).

* For inter-band carrier aggregation, with or without MIMO or TX diversity, TAE shall not exceed 260 ns.

Transport link contribution to radio frequency error:
+-2 PPB. This value is considered to be "available" for the Fronthaul link out of the total 50 PPB budget reserved for the radio interface. Note: the reason that the transport link contributes to radio frequency error is as follows. At the time of this writing, Fronthaul communication is from the radio unit to remote radio head directly. The remote radio head is essentially a passive device (without buffering etc.) The transport drives the antenna directly by feeding it with samples and everything the transport adds will be introduced to radio as-is. So if the transport causes additional frequency error that shows immediately on the radio as well. Note: performance guarantees of low nanosecond values such as these are considered to be below the DetNet layer - it is assumed that the underlying implementation, e.g. the hardware, will provide sufficient support to enable this level of performance. These values are maintained in the use case to give an indication of the overall application.

The above listed time synchronization requirements are difficult to meet with point-to-point connected networks, and more difficult when the network includes multiple hops. It is expected that networks must include buffering at the ends of the connections as imposed by the jitter requirements, since trying to meet the jitter requirements in every intermediate node is likely to be too costly. However, every measure to reduce jitter and delay on the path makes it easier to meet the end-to-end requirements.

In order to meet the timing requirements both senders and receivers must remain time synchronized, demanding very accurate clock distribution, for example support for IEEE 1588 transparent clocks or boundary clocks in every intermediate node.

In cellular networks from the LTE radio era onward, phase synchronization is needed in addition to frequency synchronization ([TS36300], [TS23401]). Time constraints are also important due to their impact on packet loss. If a packet is delivered too late, then the packet may be dropped by the host.

6.1.4. Transport Loss Constraints

Fronthaul and Midhaul networks assume almost error-free transport. Errors can result in a reset of the radio interfaces, which can cause reduced throughput or broken radio connectivity for mobile customers.

For packetized Fronthaul and Midhaul connections packet loss may be caused by BER, congestion, or network failure scenarios. Different fronthaul functional splits are being considered by 3GPP, requiring strict frame loss ratio (FLR) guarantees. As one example (referring to the legacy CPRI split which is option 8 in 3GPP) lower layers splits may imply an FLR of less than 10E-7 for data traffic and less than 10E-6 for control and management traffic.

Many of the tools available for eliminating packet loss for Fronthaul and Midhaul networks have serious challenges, for example retransmitting lost packets and/or using forward error correction (FEC) to circumvent bit errors is practically impossible due to the additional delay incurred. Using redundant streams for better guarantees for delivery is also practically impossible in many cases due to high bandwidth requirements of Fronthaul and Midhaul networks. Protection switching is also a candidate but at the time of this writing, available technologies for the path switch are too slow to avoid reset of mobile interfaces.

Fronthaul links are assumed to be symmetric, and all Fronthaul streams (i.e. those carrying radio data) have equal priority and cannot delay or pre-empt each other. This implies that the network must guarantee that each time-sensitive flow meets their schedule.

6.1.5. Security Considerations

Establishing time-sensitive streams in the network entails reserving networking resources for long periods of time. It is important that these reservation requests be authenticated to prevent malicious reservation attempts from hostile nodes (or accidental misconfiguration). This is particularly important in the case where the reservation requests span administrative domains. Furthermore, the reservation information itself should be digitally signed to reduce the risk of a legitimate node pushing a stale or hostile configuration into another networking node.

Note: This is considered important for the security policy of the network, but does not affect the core DetNet architecture and design.

6.2. Cellular Radio Networks Today

6.2.1. Fronthaul

Today's Fronthaul networks typically consist of:

- Dedicated point-to-point fiber connection is common
- Proprietary protocols and framings
- Custom equipment and no real networking

At the time of this writing, solutions for Fronthaul are direct optical cables or Wavelength-Division Multiplexing (WDM) connections.

6.2.2. Midhaul and Backhaul

Today's Midhaul and Backhaul networks typically consist of:

- Mostly normal IP networks, MPLS-TP, etc.
- Clock distribution and sync using 1588 and SyncE

Telecommunication networks in the Mid- and Backhaul are already heading towards transport networks where precise time synchronization support is one of the basic building blocks. While the transport networks themselves have practically transitioned to all-IP packet-based networks to meet the bandwidth and cost requirements, highly accurate clock distribution has become a challenge.

In the past, Mid- and Backhaul connections were typically based on Time Division Multiplexing (TDM-based) and provided frequency synchronization capabilities as a part of the transport media. Alternatively other technologies such as Global Positioning System (GPS) or Synchronous Ethernet (SyncE) are used [SyncE].

Both Ethernet and IP/MPLS [RFC3031] (and PseudoWires (PWE) [RFC3985] for legacy transport support) have become popular tools to build and manage new all-IP Radio Access Networks (RANs) [I-D.kh-spring-ip-ran-use-case]. Although various timing and synchronization optimizations have already been proposed and implemented including 1588 PTP enhancements [I-D.ietf-tictoc-1588overmpls] and [RFC8169], these solution are not necessarily sufficient for the forthcoming RAN architectures nor do

they guarantee the more stringent time-synchronization requirements such as [CPRI].

There are also existing solutions for TDM over IP such as [RFC4553], [RFC5086], and [RFC5087], as well as TDM over Ethernet transports such as [MEF8].

6.3. Cellular Radio Networks Future

Future Cellular Radio Networks will be based on a mix of different xHaul networks (xHaul = front-, mid- and backhaul), and future transport networks should be able to support all of them simultaneously. It is already envisioned today that:

- Not all "cellular radio network" traffic will be IP, for example some will remain at Layer 2 (e.g. Ethernet based). DetNet solutions must address all traffic types (Layer 2, Layer 3) with the same tools and allow their transport simultaneously.

- All forms of xHaul networks will need some form of DetNet solutions. For example with the advent of 5G some Backhaul traffic will also have DetNet requirements, for example traffic belonging to time-critical 5G applications.

- Different splits of the functionality run on the base stations and the on-site units could co-exist on the same Fronthaul and Backhaul network.

Future Cellular Radio networks should contain the following:

- Unified standards-based transport protocols and standard networking equipment that can make use of underlying deterministic link-layer services

- Unified and standards-based network management systems and protocols in all parts of the network (including Fronthaul)

New radio access network deployment models and architectures may require time- sensitive networking services with strict requirements on other parts of the network that previously were not considered to be packetized at all. Time and synchronization support are already topical for Backhaul and Midhaul packet networks [MEF22.1.1] and are becoming a real issue for Fronthaul networks also. Specifically in Fronthaul networks the timing and synchronization requirements can be extreme for packet based technologies, for example, on the order of sub +-20 ns packet delay variation (PDV) and frequency accuracy of +0.002 PPM [Fronthaul].

The actual transport protocols and/or solutions to establish required transport "circuits" (pinned-down paths) for Fronthaul traffic are still undefined. Those are likely to include (but are not limited to) solutions directly over Ethernet, over IP, and using MPLS/ PseudoWire transport.

Interesting and important work for time-sensitive networking has been done for Ethernet [TSNTG], which specifies the use of IEEE 1588 time precision protocol (PTP) [IEEE1588] in the context of IEEE 802.1D and IEEE 802.1Q. [IEEE8021AS] specifies a Layer 2 time synchronizing service, and other specifications such as IEEE 1722 [IEEE1722] specify Ethernet-based Layer-2 transport for time-sensitive streams.

However even these Ethernet TSN features may not be sufficient for Fronthaul traffic. Therefore, having specific profiles that take the requirements of Fronthaul into account is desirable [IEEE8021CM].

New promising work seeks to enable the transport of time-sensitive fronthaul streams in Ethernet bridged networks [IEEE8021CM]. Analogous to IEEE 1722 there is an ongoing standardization effort to define the Layer-2 transport encapsulation format for transporting radio over Ethernet (RoE) in the IEEE 1904.3 Task Force [IEEE19143].

As mentioned in Section 6.1.2, 5G communications will provide one of the most challenging cases for delay sensitive networking. In order to meet the challenges of ultra-low latency and ultra-high throughput, 3GPP has studied various "functional splits" for 5G, i.e., physical decomposition of the gNodeB base station and deployment of its functional blocks in different locations [TR38801].

These splits are numbered from split option 1 (Dual Connectivity, a split in which the radio resource control is centralized and other radio stack layers are in distributed units) to split option 8 (a PHY-RF split in which RF functionality is in a distributed unit and the rest of the radio stack is in the centralized unit), with each intermediate split having its own data rate and delay requirements. Packetized versions of different splits have been proposed including eCPRI [eCPRI] and RoE (as previously noted). Both provide Ethernet encapsulations, and eCPRI is also capable of IP encapsulation.

All-IP RANs and xHaul networks would benefit from time synchronization and time-sensitive transport services. Although Ethernet appears to be the unifying technology for the transport, there is still a disconnect providing Layer 3 services. The protocol stack typically has a number of layers below the Ethernet Layer 2 that shows up to the Layer 3 IP transport. It is not uncommon that on top of the lowest layer (optical) transport there is the first layer of Ethernet followed one or more layers of MPLS, PseudoWires

and/or other tunneling protocols finally carrying the Ethernet layer visible to the user plane IP traffic.

While there are existing technologies to establish circuits through the routed and switched networks (especially in MPLS/PWE space), there is still no way to signal the time synchronization and time-sensitive stream requirements/reservations for Layer-3 flows in a way that addresses the entire transport stack, including the Ethernet layers that need to be configured.

Furthermore, not all "user plane" traffic will be IP. Therefore, the same solution also must address the use cases where the user plane traffic is a different layer, for example Ethernet frames.

There is existing work describing the problem statement [I-D.ietf-detnet-problem-statement] and the architecture [I-D.ietf-detnet-architecture] for deterministic networking (DetNet) that targets solutions for time-sensitive (IP/transport) streams with deterministic properties over Ethernet-based switched networks.

6.4. Cellular Radio Networks Asks

A standard for data plane transport specification which is:

- Unified among all xHauls (meaning that different flows with diverse DetNet requirements can coexist in the same network and traverse the same nodes without interfering with each other)

- Deployed in a highly deterministic network environment

- Capable of supporting multiple functional splits simultaneously, including existing Backhaul and CPRI Fronthaul and potentially new modes as defined for example in 3GPP; these goals can be supported by the existing DetNet Use Case Common Themes, notably "Mix of Deterministic and Best-Effort Traffic", "Bounded Latency", "Low Latency", "Symmetrical Path Delays", and "Deterministic Flows".

- Capable of supporting Network Slicing and Multi-tenancy; these goals can be supported by the same DetNet themes noted above.

- Capable of transporting both in-band and out-band control traffic (OAM info, ...).

- Deployable over multiple data link technologies (e.g., IEEE 802.3, mmWave, etc.).

A standard for data flow information models that are:

- Aware of the time sensitivity and constraints of the target networking environment

- Aware of underlying deterministic networking services (e.g., on the Ethernet layer)

7. Industrial Machine to Machine (M2M)

7.1. Use Case Description

Industrial Automation in general refers to automation of manufacturing, quality control and material processing. This "machine to machine" (M2M) use case considers machine units in a plant floor which periodically exchange data with upstream or downstream machine modules and/or a supervisory controller within a local area network.

The actors of M2M communication are Programmable Logic Controllers (PLCs). Communication between PLCs and between PLCs and the supervisory PLC (S-PLC) is achieved via critical control/data streams Figure 11.

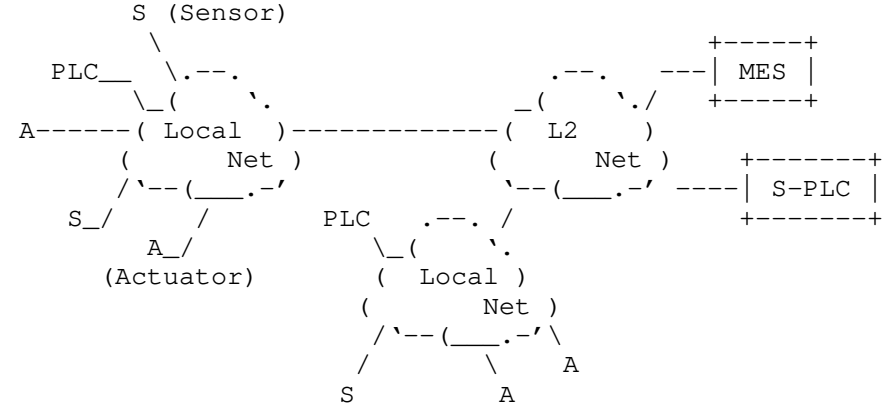

Figure 11: Current Generic Industrial M2M Network Architecture

This use case focuses on PLC-related communications; communication to Manufacturing-Execution-Systems (MESs) are not addressed.

This use case covers only critical control/data streams; non-critical traffic between industrial automation applications (such as communication of state, configuration, set-up, and database communication) are adequately served by prioritizing techniques available at the time of this writing. Such traffic can use up to

80% of the total bandwidth required. There is also a subset of non-time-critical traffic that must be reliable even though it is not time-sensitive.

In this use case the primary need for deterministic networking is to provide end-to-end delivery of M2M messages within specific timing constraints, for example in closed loop automation control. Today this level of determinism is provided by proprietary networking technologies. In addition, standard networking technologies are used to connect the local network to remote industrial automation sites, e.g. over an enterprise or metro network which also carries other types of traffic. Therefore, flows that should be forwarded with deterministic guarantees need to be sustained regardless of the amount of other flows in those networks.

7.2. Industrial M2M Communication Today

Today, proprietary networks fulfill the needed timing and availability for M2M networks.

The network topologies used today by industrial automation are similar to those used by telecom networks: Daisy Chain, Ring, Hub and Spoke, and Comb (a subset of Daisy Chain).

PLC-related control/data streams are transmitted periodically and carry either a pre-configured payload or a payload configured during runtime.

Some industrial applications require time synchronization at the end nodes. For such time-coordinated PLCs, accuracy of 1 microsecond is required. Even in the case of "non-time-coordinated" PLCs time sync may be needed e.g. for timestamping of sensor data.

Industrial network scenarios require advanced security solutions. At the time of this writing, many industrial production networks are physically separated. Preventing critical flows from being leaked outside a domain is handled by filtering policies that are typically enforced in firewalls.

7.2.1. Transport Parameters

The Cycle Time defines the frequency of message(s) between industrial actors. The Cycle Time is application dependent, in the range of 1ms - 100ms for critical control/data streams.

Because industrial applications assume deterministic transport for critical Control-Data-Stream parameters (instead of defining latency and delay variation parameters) it is sufficient to fulfill the upper

bound of latency (maximum latency). The underlying networking infrastructure must ensure a maximum end-to-end delivery time of messages in the range of 100 microseconds to 50 milliseconds depending on the control loop application.

The bandwidth requirements of control/data streams are usually calculated directly from the bytes-per-cycle parameter of the control loop. For PLC-to-PLC communication one can expect 2 - 32 streams with packet size in the range of 100 - 700 bytes. For S-PLC to PLCs the number of streams is higher - up to 256 streams. Usually no more than 20% of available bandwidth is used for critical control/data streams. In today's networks 1Gbps links are commonly used.

Most PLC control loops are rather tolerant of packet loss, however critical control/data streams accept no more than 1 packet loss per consecutive communication cycle (i.e. if a packet gets lost in cycle "n", then the next cycle ("n+1") must be lossless). After two or more consecutive packet losses the network may be considered to be "down" by the Application.

As network downtime may impact the whole production system the required network availability is rather high (99.999%).

Based on the above parameters some form of redundancy will be required for M2M communications, however any individual solution depends on several parameters including cycle time, delivery time, etc.

7.2.2. Stream Creation and Destruction

In an industrial environment, critical control/data streams are created rather infrequently, on the order of ~10 times per day / week / month. Most of these critical control/data streams get created at machine startup, however flexibility is also needed during runtime, for example when adding or removing a machine. Going forward as production systems become more flexible, there will be a significant increase in the rate at which streams are created, changed and destroyed.

7.3. Industrial M2M Future

We foresee a converged IP-standards-based network with deterministic properties that can satisfy the timing, security and reliability constraints described above. Today's proprietary networks could then be interfaced to such a network via gateways or, in the case of new installations, devices could be connected directly to the converged network.

For this use case time synchronization accuracy on the order of 1us is expected.

7.4. Industrial M2M Asks

- Converged IP-based network
- Deterministic behavior (bounded latency and jitter)
- High availability (presumably through redundancy) (99.999 %)
- Low message delivery time (100us - 50ms)
- Low packet loss (with bounded number of consecutive lost packets)
- Security (e.g. prevent critical flows from being leaked between physically separated networks)

8. Mining Industry

8.1. Use Case Description

The mining industry is highly dependent on networks to monitor and control their systems both in open-pit and underground extraction, transport and refining processes. In order to reduce risks and increase operational efficiency in mining operations, a number of processes have migrated the operators from the extraction site to remote control and monitoring.

In the case of open pit mining, autonomous trucks are used to transport the raw materials from the open pit to the refining factory where the final product (e.g. Copper) is obtained. Although the operation is autonomous, the tracks are remotely monitored from a central facility.

In pit mines, the monitoring of the tailings or mine dumps is critical in order to minimize environmental pollution. In the past, monitoring has been conducted through manual inspection of pre-installed dataloggers. Cabling is not usually exploited in such scenarios due to the cost and complex deployment requirements. At the time of this writing, wireless technologies are being employed to monitor these cases permanently. Slopes are also monitored in order to anticipate possible mine collapse. Due to the unstable terrain, cable maintenance is costly and complex and hence wireless technologies are employed.

In the underground monitoring case, autonomous vehicles with extraction tools travel autonomously through the tunnels, but their

operational tasks (such as excavation, stone breaking and transport) are controlled remotely from a central facility. This generates video and feedback upstream traffic plus downstream actuator control traffic.

8.2. Mining Industry Today

At the time of this writing, the mining industry uses a packet switched architecture supported by high speed ethernet. However in order to achieve the delay and packet loss requirements the network bandwidth is overestimated, thus providing very low efficiency in terms of resource usage.

QoS is implemented at the Routers to separate video, management, monitoring and process control traffic for each stream.

Since mobility is involved in this process, the connection between the backbone and the mobile devices (e.g. trucks, trains and excavators) is solved using a wireless link. These links are based on 802.11 for open-pit mining and "leaky feeder" communications for underground mining. (A "leaky feeder" communication system consists of a coaxial cable run along tunnels which emits and receives radio waves, functioning as an extended antenna. The cable is "leaky" in that it has gaps or slots in its outer conductor to allow the radio signal to leak into or out of the cable along its entire length.)

Lately in pit mines the use of LPWAN technologies has been extended: Tailings, slopes and mine dumps are monitored by battery-powered dataloggers that make use of robust long range radio technologies. Reliability is usually ensured through retransmissions at L2. Gateways or concentrators act as bridges forwarding the data to the backbone ethernet network. Deterministic requirements are biased towards reliability rather than latency as events are slowly triggered or can be anticipated in advance.

At the mineral processing stage, conveyor belts and refining processes are controlled by a SCADA system, which provides the in-factory delay-constrained networking requirements.

At the time of this writing, voice communications are served by a redundant trunking infrastructure, independent from data networks.

8.3. Mining Industry Future

Mining operations and management are converging towards a combination of autonomous operation and teleoperation of transport and extraction machines. This means that video, audio, monitoring and process

control traffic will increase dramatically. Ideally, all activities on the mine will rely on network infrastructure.

Wireless for open-pit mining is already a reality with LPWAN technologies and it is expected to evolve to more advanced LPWAN technologies such as those based on LTE to increase last hop reliability or novel LPWAN flavours with deterministic access.

One area in which DetNet can improve this use case is in the wired networks that make up the "backbone network" of the system, which connect together many wireless access points (APs). The mobile machines (which are connected to the network via wireless) transition from one AP to the next as they move about. A deterministic, reliable, low latency backbone can enable these transitions to be more reliable.

Connections which extend all the way from the base stations to the machinery via a mix of wired and wireless hops would also be beneficial, for example to improve remote control responsiveness of digging machines. However to guarantee deterministic performance of a DetNet, the end-to-end underlying network must be deterministic. Thus for this use case if a deterministic wireless transport is integrated with a wire-based DetNet network, it could create the desired wired plus wireless end-to-end deterministic network.

8.4. Mining Industry Asks

- Improved bandwidth efficiency
- Very low delay to enable machine teleoperation
- Dedicated bandwidth usage for high resolution video streams
- Predictable delay to enable realtime monitoring
- Potential to construct a unified DetNet network over a combination of wired and deterministic wireless links

9. Private Blockchain

9.1. Use Case Description

Blockchain was created with bitcoin as a ’public’ blockchain on the open Internet, however blockchain has also spread far beyond its original host into various industries such as smart manufacturing, logistics, security, legal rights and others. In these industries blockchain runs in designated and carefully managed networks in which

deterministic networking requirements could be addressed by DetNet. Such implementations are referred to as 'private' blockchain.

The sole distinction between public and private blockchain is defined by who is allowed to participate in the network, execute the consensus protocol, and maintain the shared ledger.

Today's networks treat the traffic from blockchain on a best-effort basis, but blockchain operation could be made much more efficient if deterministic networking services were available to minimize latency and packet loss in the network.

9.1.1. Blockchain Operation

A 'block' runs as a container of a batch of primary items such as transactions, property records etc. The blocks are chained in such a way that the hash of the previous block works as the pointer to the header of the new block. Confirmation of each block requires a consensus mechanism. When an item arrives at a blockchain node, the latter broadcasts this item to the rest of the nodes which receive and verify it and put it in the ongoing block. The block confirmation process begins as the number of items reaches the predefined block capacity, at which time the node broadcasts its proved block to the rest of the nodes, to be verified and chained. The result is that block N+1 of each chain transitively vouches for blocks N and before of that chain.

9.1.2. Blockchain Network Architecture

Blockchain node communication and coordination is achieved mainly through frequent point-to-multi-point communication, however persistent point-to-point connections are used to transport both the items and the blocks to the other nodes. For example, consider the following implementation.

When a node is initiated, it first requests the other nodes' address from a specific entity such as DNS, then it creates persistent connections each of with other nodes. If a node confirms an item, it sends the item to the other nodes via these persistent connections.

As a new block in a node is completed and is proven by the surrounding nodes, it propagates towards its neighbor nodes. When node A receives a block, it verifies it, then sends an invite message to its neighbor B. Neighbor B checks to see if the designated block is available, and responds to A if it is unavailable, then A sends the complete block to B. B repeats the process (as done by A above) to start the next round of block propagation.

The challenge of blockchain network operation is not overall data rates, since the volume from both block and item stays between hundreds of bytes to a couple of megabytes per second, but is in transporting the blocks with minimum latency to maximize efficiency of the blockchain consensus process. The efficiency of differing implementations of the consensus process may be affected to a differing degree by the latency (and variation of latency) of the network.

9.1.3. Security Considerations

Security is crucial to blockchain applications, and at the time of this writing, blockchain systems address security issues mainly at the application level, where cryptography as well as hash-based consensus play a leading role in preventing both double-spending and malicious service attacks. However, there is concern that in the proposed use case of a private blockchain network which is dependent on deterministic properties, the network could be vulnerable to delays and other specific attacks against determinism which could interrupt service.

9.2. Private Blockchain Today

Today private blockchain runs in L2 or L3 VPN, in general without guaranteed determinism. The industry players are starting to realize that improving determinism in their blockchain networks could improve the performance of their service, but as of today these goals are not being met.

9.3. Private Blockchain Future

Blockchain system performance can be greatly improved through deterministic networking service primarily because it would accelerate the consensus process. It would be valuable to be able to design a private blockchain network with the following properties:

- Transport of point-to-multi-point traffic in a coordinated network architecture rather than at the application layer (which typically uses point-to-point connections)
- Guaranteed transport latency
- Reduced packet loss (to the point where packet retransmission-incurred delay would be negligible.)

9.4. Private Blockchain Asks

- Layer 2 and Layer 3 multicast of blockchain traffic
- Item and block delivery with bounded, low latency and negligible packet loss
- Coexistence in a single network of blockchain and IT traffic.
- Ability to scale the network by distributing the centralized control of the network across multiple control entities.

10. Network Slicing

10.1. Use Case Description

Network Slicing divides one physical network infrastructure into multiple logical networks. Each slice, corresponding to a logical network, uses resources and network functions independently from each other. Network Slicing provides flexibility of resource allocation and service quality customization.

Future services will demand network performance with a wide variety of characteristics such as high data rate, low latency, low loss rate, security and many other parameters. Ideally every service would have its own physical network satisfying its particular performance requirements, however that would be prohibitively expensive. Network Slicing can provide a customized slice for a single service, and multiple slices can share the same physical network. This method can optimize the performance for the service at lower cost, and the flexibility of setting up and release the slices also allows the user to allocate the network resources dynamically.

Unlike the other use cases presented here, Network Slicing is not a specific application that depends on specific deterministic properties; rather it is introduced as an area of networking to which DetNet might be applicable.

10.2. DetNet Applied to Network Slicing

10.2.1. Resource Isolation Across Slices

One of the requirements discussed for Network Slicing is the "hard" separation of various users' deterministic performance. That is, it should be impossible for activity, lack of activity, or changes in activity of one or more users to have any appreciable effect on the deterministic performance parameters of any other slices. Typical techniques used today, which share a physical network among users, do

not offer this level of isolation. DetNet can supply point-to-point or point-to-multipoint paths that offer bandwidth and latency guarantees to a user that cannot be affected by other users' data traffic. Thus DetNet is a powerful tool when latency and reliability are required in Network Slicing.

10.2.2. Deterministic Services Within Slices

Slices may need to provide services with DetNet-type performance guarantees, however note that a system can be implemented to provide such services in more than one way. For example the slice itself might be implemented using DetNet, and thus the slice can provide service guarantees and isolation to its users without any particular DetNet awareness on the part of the users' applications. Alternatively, a "non-DetNet-aware" slice may host an application that itself implements DetNet services and thus can enjoy similar service guarantees.

10.3. A Network Slicing Use Case Example - 5G Bearer Network

Network Slicing is a core feature of 5G defined in 3GPP, which is under development at the time of this writing [TR38501]. A network slice in a mobile network is a complete logical network including Radio Access Network (RAN) and Core Network (CN). It provides telecommunication services and network capabilities, which may vary from slice to slice. A 5G bearer network is a typical use case of Network Slicing; for example consider three 5G service scenarios: eMMB, URLLC, and mMTC.

- eMBB (Enhanced Mobile Broadband) focuses on services characterized by high data rates, such as high definition videos, virtual reality, augmented reality, and fixed mobile convergence.

- URLLC (Ultra-Reliable and Low Latency Communications) focuses on latency-sensitive services, such as self-driving vehicles, remote surgery, or drone control.

- mMTC (massive Machine Type Communications) focuses on services that have high requirements for connection density, such as those typical for smart city and smart agriculture use cases.

A 5G bearer network could use DetNet to provide hard resource isolation across slices and within the slice. For example consider Slice-A and Slice-B, with DetNet used to transit services URLLC-A and URLLC-B over them. Without DetNet, URLLC-A and URLLC-B would compete for bandwidth resource, and latency and reliability would not be guaranteed. With DetNet, URLLC-A and URLLC-B have separate bandwidth

reservation and there is no resource conflict between them, as though they were in different logical networks.

10.4. Non-5G Applications of Network Slicing

Although operation of services not related to 5G is not part of the 5G Network Slicing definition and scope, Network Slicing is likely to become a preferred approach to providing various services across a shared physical infrastructure. Examples include providing electrical utilities services and pro audio services via slices. Use cases like these could become more common once the work for the 5G core network evolves to include wired as well as wireless access.

10.5. Limitations of DetNet in Network Slicing

DetNet cannot cover every Network Slicing use case. One issue is that DetNet is a point-to-point or point-to-multipoint technology, however Network Slicing ultimately needs multi-point to multi-point guarantees. Another issue is that the number of flows that can be carried by DetNet is limited by DetNet scalability; flow aggregation and queuing management modification may help address this. Additional work and discussion are needed to address these topics.

10.6. Network Slicing Today and Future

Network Slicing has the promise to satisfy many requirements of future network deployment scenarios, but it is still a collection of ideas and analysis, without a specific technical solution. DetNet is one of various technologies that have potential to be used in Network Slicing, along with for example Flex-E and Segment Routing. For more information please see the IETF99 Network Slicing BOF session agenda and materials.

10.7. Network Slicing Asks

- Isolation from other flows through Queuing Management
- Service Quality Customization and Guarantee
- Security

11. Use Case Common Themes

This section summarizes the expected properties of a DetNet network, based on the use cases as described in this draft.

11.1. Unified, standards-based network

11.1.1. Extensions to Ethernet

A DetNet network is not "a new kind of network" - it based on extensions to existing Ethernet standards, including elements of IEEE 802.1 AVB/TSN and related standards. Presumably it will be possible to run DetNet over other underlying transports besides Ethernet, but Ethernet is explicitly supported.

11.1.2. Centrally Administered

In general a DetNet network is not expected to be "plug and play" - it is expected that there is some centralized network configuration and control system. Such a system may be in a single central location, or it maybe distributed across multiple control entities that function together as a unified control system for the network. However, the ability to "hot swap" components (e.g. due to malfunction) is similar enough to "plug and play" that this kind of behavior may be expected in DetNet networks, depending on the implementation.

11.1.3. Standardized Data Flow Information Models

Data Flow Information Models to be used with DetNet networks are to be specified by DetNet.

11.1.4. L2 and L3 Integration

A DetNet network is intended to integrate between Layer 2 (bridged) network(s) (e.g. AVB/TSN LAN) and Layer 3 (routed) network(s) (e.g. using IP-based protocols). One example of this is "making AVB/TSN-type deterministic performance available from Layer 3 applications, e.g. using RTP". Another example is "connecting two AVB/TSN LANs ("islands") together through a standard router".

11.1.5. Consideration for IPv4

This Use Cases draft explicitly does not specify any particular implementation or protocol, however it has been observed that various of the use cases described (and their associated industries) are explicitly based on IPv4 (as opposed to IPv6) and it is not considered practical to expect them to migrate to IPv6 in order to use DetNet. Thus the expectation is that even if not every feature of DetNet is available in an IPv4 context, at least some of the significant benefits (such as guaranteed end-to-end delivery and low latency) are expected to be available.

11.1.6. Guaranteed End-to-End Delivery

Packets in a DetNet flow are guaranteed not to be dropped by the network due to congestion. However, the network may drop packets for intended reasons, e.g. per security measures. Similarly best-effort traffic on a DetNet is subject to being dropped (as on a non-DetNet IP network). Also note that this guarantee applies to the actions of DetNet protocol software, and does not provide any guarantee against lower level errors such as media errors or checksum errors.

11.1.7. Replacement for Multiple Proprietary Deterministic Networks

There are many proprietary non-interoperable deterministic Ethernet-based networks available; DetNet is intended to provide an open-standards-based alternative to such networks.

11.1.8. Mix of Deterministic and Best-Effort Traffic

DetNet is intended to support coexistance of time-sensitive operational (OT) traffic and information (IT) traffic on the same ("unified") network.

11.1.9. Unused Reserved BW to be Available to Best-Effort Traffic

If bandwidth reservations are made for a stream but the associated bandwidth is not used at any point in time, that bandwidth is made available on the network for best-effort traffic. If the owner of the reserved stream then starts transmitting again, the bandwidth is no longer available for best-effort traffic, on a moment-to-moment basis. Note that such "temporarily available" bandwidth is not available for time-sensitive traffic, which must have its own reservation.

11.1.10. Lower Cost, Multi-Vendor Solutions

The DetNet network specifications are intended to enable an ecosystem in which multiple vendors can create interoperable products, thus promoting device diversity and potentially higher numbers of each device manufactured, promoting cost reduction and cost competition among vendors. The intent is that DetNet networks should be able to be created at lower cost and with greater diversity of available devices than existing proprietary networks.

11.2. Scalable Size

DetNet networks range in size from very small, e.g. inside a single industrial machine, to very large, for example a Utility Grid network spanning a whole country, and involving many "hops" over various

kinds of links for example radio repeaters, microwave linkes, fiber optic links, etc.. However recall that the scope of DetNet is confined to networks that are centrally administered, and explicitly excludes unbounded decentralized networks such as the Internet.

11.2.1. Scalable Number of Flows

The number of flows in a given network application can potentially be large, and can potentially grow faster than the number of nodes and hops. So the network should provide a sufficient (perhaps configurable) maximum number of flows for any given application.

11.3. Scalable Timing Parameters and Accuracy

11.3.1. Bounded Latency

The DetNet Data Flow Information Model is expected to provide means to configure the network that include parameters for querying network path latency, requesting bounded latency for a given stream, requesting worst case maximum and/or minimum latency for a given path or stream, and so on. It is an expected case that the network may not be able to provide a given requested service level, and if so the network control system should reply that the requested services is not available (as opposed to accepting the parameter but then not delivering the desired behavior).

11.3.2. Low Latency

Applications may require "extremely low latency" however depending on the application these may mean very different latency values; for example "low latency" across a Utility grid network is on a different time scale than "low latency" in a motor control loop in a small machine. The intent is that the mechanisms for specifying desired latency include wide ranges, and that architecturally there is nothing to prevent arbirtrarily low latencies from being implemented in a given network.

11.3.3. Bounded Jitter (Latency Variation)

As with the other Latency-related elements noted above, parameters should be available to determine or request the allowed variation in latency.

11.3.4. Symmetrical Path Delays

Some applications would like to specify that the transit delay time values be equal for both the transmit and return paths.

11.4. High Reliability and Availability

Reliablity is of critical importance to many DetNet applications, in which consequences of failure can be extraordinarily high in terms of cost and even human life. DetNet based systems are expected to be implemented with essentially arbitrarily high availability (for example 99.9999% up time, or even 12 nines). The intent is that the DetNet designs should not make any assumptions about the level of reliability and availability that may be required of a given system, and should define parameters for communicating these kinds of metrics within the network.

A strategy used by DetNet for providing such extraordinarily high levels of reliability is to provide redundant paths that can be seamlessly switched between, while maintaining the required performance of that system.

11.5. Security

Security is of critical importance to many DetNet applications. A DetNet network must be able to be made secure against devices failures, attackers, misbehaving devices, and so on. In a DetNet network the data traffic is expected to be be time-sensitive, thus in addition to arriving with the data content as intended, the data must also arrive at the expected time. This may present "new" security challenges to implementers, and must be addressed accordingly. There are other security implications, including (but not limited to) the change in attack surface presented by packet replication and elimination.

11.6. Deterministic Flows

Reserved bandwidth data flows must be isolated from each other and from best-effort traffic, so that even if the network is saturated with best-effort (and/or reserved bandwidth) traffic, the configured flows are not adversely affected.

12. Security Considerations

This document covers a number of representative applications and network scenarios that are expected to make use of DetNet technologies. Each of the potential DetNet uses cases will have security considerations from both the use-specific and DetNet technology perspectives. While some use-specific security considerations are discussed above, a more comprehensive discussion of such considerations is captured in DetNet Security Considerations [I-D.ietf-detnet-security]. Readers are encouraged to review this

document to gain a more complete understanding of DetNet related security considerations.

13. Contributors

RFC7322 limits the number of authors listed on the front page of a draft to a maximum of 5, far fewer than the 20 individuals below who made important contributions to this draft. The editor wishes to thank and acknowledge each of the following authors for contributing text to this draft. See also Section 14.

Craig Gunther (Harman International)
10653 South River Front Parkway, South Jordan,UT 84095
phone +1 801 568-7675, email craig.gunther@harman.com

Pascal Thubert (Cisco Systems, Inc)
Building D, 45 Allee des Ormes - BP1200, MOUGINS
Sophia Antipolis 06254 FRANCE
phone +33 497 23 26 34, email pthubert@cisco.com

Patrick Wetterwald (Cisco Systems)
45 Allees des Ormes, Mougins, 06250 FRANCE
phone +33 4 97 23 26 36, email pwetterw@cisco.com

Jean Raymond (Hydro-Quebec)
1500 University, Montreal, H3A3S7, Canada
phone +1 514 840 3000, email raymond.jean@hydro.qc.ca

Jouni Korhonen (Broadcom Corporation)
3151 Zanker Road, San Jose, 95134, CA, USA
email jouni.nospam@gmail.com

Yu Kaneko (Toshiba)
1 Komukai-Toshiba-cho, Saiwai-ku, Kasasaki-shi, Kanagawa, Japan
email yu1.kaneko@toshiba.co.jp

Subir Das (Vencore Labs)
150 Mount Airy Road, Basking Ridge, New Jersey, 07920, USA
email sdas@appcomsci.com

Balazs Varga (Ericsson)
Konyves Kalman krt. 11/B, Budapest, Hungary, 1097
email balazs.a.varga@ericsson.com

Janos Farkas (Ericsson)
Konyves Kalman krt. 11/B, Budapest, Hungary, 1097
email janos.farkas@ericsson.com

Franz-Josef Goetz (Siemens)
Gleiwitzerstr. 555, Nurnberg, Germany, 90475
email franz-josef.goetz@siemens.com

Juergen Schmitt (Siemens)
Gleiwitzerstr. 555, Nurnberg, Germany, 90475
email juergen.jues.schmitt@siemens.com

Xavier Vilajosana (Worldsensing)
483 Arago, Barcelona, Catalonia, 08013, Spain
email xvilajosana@worldsensing.com

Toktam Mahmoodi (King's College London)
Strand, London WC2R 2LS, United Kingdom
email toktam.mahmoodi@kcl.ac.uk

Spiros Spirou (Intracom Telecom)
19.7 km Markopoulou Ave., Peania, Attiki, 19002, Greece
email spiros.spirou@gmail.com

Petra Vizarreta (Technical University of Munich)
Maxvorstadt, ArcisstraBe 21, Munich, 80333, Germany
email petra.stojsavljevic@tum.de

Daniel Huang (ZTE Corporation, Inc.)
No. 50 Software Avenue, Nanjing, Jiangsu, 210012, P.R. China
email huang.guangping@zte.com.cn

Xuesong Geng (Huawei Technologies)
email gengxuesong@huawei.com

Diego Dujovne (Universidad Diego Portales)
email diego.dujovne@mail.udp.cl

Maik Seewald (Cisco Systems)
email maseewal@cisco.com

14. Acknowledgments

14.1. Pro Audio

This section was derived from draft-gunther-detnet-proaudio-req-01.

The editors would like to acknowledge the help of the following individuals and the companies they represent:

Jeff Koftinoff, Meyer Sound

Jouni Korhonen, Associate Technical Director, Broadcom

Pascal Thubert, CTAO, Cisco

Kieran Tyrrell, Sienda New Media Technologies GmbH

14.2. Utility Telecom

This section was derived from draft-wetterwald-detnet-utilities-reqs-02.

Faramarz Maghsoodlou, Ph. D. IoT Connected Industries and Energy Practice Cisco

Pascal Thubert, CTAO Cisco

The wind power generation use case has been extracted from the study of Wind Farms conducted within the 5GPPP Virtuwind Project. The project is funded by the European Union's Horizon 2020 research and innovation programme under grant agreement No 671648 (VirtuWind).

14.3. Building Automation Systems

This section was derived from draft-bas-usecase-detnet-00.

14.4. Wireless for Industrial Applications

This section was derived from draft-thubert-6tisch-4detnet-01.

This specification derives from the 6TiSCH architecture, which is the result of multiple interactions, in particular during the 6TiSCH (bi)Weekly Interim call, relayed through the 6TiSCH mailing list at the IETF.

The authors wish to thank: Kris Pister, Thomas Watteyne, Xavier Vilajosana, Qin Wang, Tom Phinney, Robert Assimiti, Michael Richardson, Zhuo Chen, Malisa Vucinic, Alfredo Grieco, Martin Turon, Dominique Barthel, Elvis Vogli, Guillaume Gaillard, Herman Storey, Maria Rita Palattella, Nicola Accettura, Patrick Wetterwald, Pouria Zand, Raghuram Sudhaakar, and Shitanshu Shah for their participation and various contributions.

14.5. Cellular Radio

This section was derived from draft-korhonen-detnet-telreq-00.

14.6. Industrial Machine to Machine (M2M)

The authors would like to thank Feng Chen and Marcel Kiessling for their comments and suggestions.

14.7. Internet Applications and CoMP

This section was derived from draft-zha-detnet-use-case-00 by Yiyong Zha.

This document has benefited from reviews, suggestions, comments and proposed text provided by the following members, listed in alphabetical order: Jing Huang, Junru Lin, Lehong Niu and Oilver Huang.

14.8. Network Slicing

This section was written by Xuesong Geng, who would like to acknowledge Norm Finn and Mach Chen for their useful comments.

14.9. Mining

This section was written by Diego Dujovne in conjunction with Xavier Vilasojana.

14.10. Private Blockchain

This section was written by Daniel Huang.

15. IANA Considerations

This memo includes no requests from IANA.

16. Informative References

[Ahm14] Ahmed, M. and R. Kim, "Communication network architectures for smart-wind power farms.", Energies, p. 3900-3921. , June 2014.

[bacnetip] ASHRAE, "Annex J to ANSI/ASHRAE 135-1995 - BACnet/IP", January 1999.

[CoMP] NGMN Alliance, "RAN EVOLUTION PROJECT COMP EVALUATION AND ENHANCEMENT", NGMN Alliance NGMN_RANEV_D3_CoMP_Evaluation_and_Enhancement_v2.0, March 2015, <https://www.ngmn.org/uploads/media/NGMN_RANEV_D3_CoMP_Evaluation_and_Enhancement_v2.0.pdf>.

[CONTENT_PROTECTION]
Olsen, D., "1722a Content Protection", 2012, <http://grouper.ieee.org/groups/1722/contributions/2012/avtp_dolsen_1722a_content_protection.pdf>.

[CPRI] CPRI Cooperation, "Common Public Radio Interface (CPRI); Interface Specification", CPRI Specification V6.1, July 2014, <http://www.cpri.info/downloads/CPRI_v_6_1_2014-07-01.pdf>.

[DCI] Digital Cinema Initiatives, LLC, "DCI Specification, Version 1.2", 2012, <http://www.dcimovies.com/>.

[eCPRI] IEEE Standards Association, "Common Public Radio Interface, "Common Public Radio Interface: eCPRI Interface Specification V1.0", 2017, <http://www.cpri.info/>.

[ESPN_DC2]
Daley, D., "ESPN's DC2 Scales AVB Large", 2014, <http://sportsvideo.org/main/blog/2014/06/espns-dc2-scales-avb-large>.

[flnet] Japan Electrical Manufacturers Association, "JEMA 1479 - English Edition", September 2012.

[Fronthaul]
Chen, D. and T. Mustala, "Ethernet Fronthaul Considerations", IEEE 1904.3, February 2015, <http://www.ieee1904.org/3/meeting_archive/2015/02/tf3_1502_che n_1a.pdf>.

[I-D.ietf-6tisch-6top-interface]
Wang, Q. and X. Vilajosana, "6TiSCH Operation Sublayer (6top) Interface", draft-ietf-6tisch-6top-interface-04 (work in progress), July 2015.

[I-D.ietf-6tisch-architecture]
Thubert, P., "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4", draft-ietf-6tisch-architecture-19 (work in progress), December 2018.

[I-D.ietf-6tisch-coap]
Sudhaakar, R. and P. Zand, "6TiSCH Resource Management and Interaction using CoAP", draft-ietf-6tisch-coap-03 (work in progress), March 2015.

[I-D.ietf-detnet-architecture]
Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", draft-ietf-detnet-architecture-09 (work in progress), October 2018.

[I-D.ietf-detnet-problem-statement]
Finn, N. and P. Thubert, "Deterministic Networking Problem Statement", draft-ietf-detnet-problem-statement-08 (work in progress), December 2018.

[I-D.ietf-detnet-security]
Mizrahi, T., Grossman, E., Hacker, A., Das, S., Dowdell, J., Austad, H., Stanton, K., and N. Finn, "Deterministic Networking (DetNet) Security Considerations", draft-ietf-detnet-security-03 (work in progress), October 2018.

[I-D.ietf-tictoc-1588overmpls]
Davari, S., Oren, A., Bhatia, M., Roberts, P., and L. Montini, "Transporting Timing messages over MPLS Networks", draft-ietf-tictoc-1588overmpls-07 (work in progress), October 2015.

[I-D.kh-spring-ip-ran-use-case]
Khasnabish, B., hu, f., and L. Contreras, "Segment Routing in IP RAN use case", draft-kh-spring-ip-ran-use-case-02 (work in progress), November 2014.

[I-D.svshah-tsvwg-deterministic-forwarding]
Shah, S. and P. Thubert, "Deterministic Forwarding PHB", draft-svshah-tsvwg-deterministic-forwarding-04 (work in progress), August 2015.

[I-D.wang-6tisch-6top-sublayer]
Wang, Q. and X. Vilajosana, "6TiSCH Operation Sublayer (6top)", draft-wang-6tisch-6top-sublayer-04 (work in progress), November 2015.

[IEC-60870-5-104]
International Electrotechnical Commission, "International Standard IEC 60870-5-104: Network access for IEC 60870-5-101 using standard transport profiles", June 2006.

[IEC61400]
"International standard 61400-25: Communications for monitoring and control of wind power plants", June 2013.

[IEEE1588]
IEEE, "IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems", IEEE Std 1588-2008, 2008, <http://standards.ieee.org/findstds/standard/1588-2008.html>.

[IEEE1646]
"Communication Delivery Time Performance Requirements for Electric Power Substation Automation", IEEE Standard 1646-2004 , Apr 2004.

[IEEE1722]
IEEE, "1722-2011 - IEEE Standard for Layer 2 Transport Protocol for Time Sensitive Applications in a Bridged Local Area Network", IEEE Std 1722-2011, 2011, <http://standards.ieee.org/findstds/standard/1722-2011.html>.

[IEEE19143]
IEEE Standards Association, "P1914.3/D3.1 Draft Standard for Radio over Ethernet Encapsulations and Mappings", IEEE 1914.3, 2018, <https://standards.ieee.org/develop/project/1914.3.html>.

[IEEE802.1TSNTG]
IEEE Standards Association, "IEEE 802.1 Time-Sensitive Networks Task Group", March 2013, <http://www.ieee802.org/1/pages/avbridges.html>.

[IEEE802154]
IEEE standard for Information Technology, "IEEE std. 802.15.4, Part. 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks".

[IEEE802154e]
IEEE standard for Information Technology, "IEEE standard for Information Technology, IEEE std. 802.15.4, Part. 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks, June 2011 as amended by IEEE std. 802.15.4e, Part. 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs) Amendment 1: MAC sublayer", April 2012.

[IEEE8021AS]
IEEE, "Timing and Synchronizations (IEEE 802.1AS-2011)", IEEE 802.1AS-2001, 2011, <http://standards.ieee.org/getIEEE802/download/802.1AS-2011.pdf>.

[IEEE8021CM]
Farkas, J., "Time-Sensitive Networking for Fronthaul", Unapproved PAR, PAR for a New IEEE Standard; IEEE P802.1CM, April 2015, <http://www.ieee802.org/1/files/public/docs2015/new-P802-1CM-dr aft-PAR-0515-v02.pdf>.

[ISA100] ISA/ANSI, "ISA100, Wireless Systems for Automation", <https://www.isa.org/isa100/>.

[knx] KNX Association, "ISO/IEC 14543-3 - KNX", November 2006.

[lontalk] ECHELON, "LonTalk(R) Protocol Specification Version 3.0", 1994.

[MEF22.1.1]
MEF, "Mobile Backhaul Phase 2 Amendment 1 -- Small Cells", MEF 22.1.1, July 2014, <http://www.mef.net/Assets/Technical_Specifications/PDF/MEF_22.1.1.pdf>.

[MEF8] MEF, "Implementation Agreement for the Emulation of PDH Circuits over Metro Ethernet Networks", MEF 8, October 2004, <https://www.mef.net/Assets/Technical_Specifications/PDF/MEF_8.pdf>.

[METIS] METIS, "Scenarios, requirements and KPIs for 5G mobile and wireless system", ICT-317669-METIS/D1.1 ICT-317669-METIS/D1.1, April 2013, <https://www.metis2020.com/wp-content/uploads/deliverables/METIS_D1.1_v1.pdf>.

[modbus] Modbus Organization, "MODBUS APPLICATION PROTOCOL SPECIFICATION V1.1b", December 2006.

[MODBUS] Modbus Organization, Inc., "MODBUS Application Protocol Specification", Apr 2012.

[NGMN] NGMN Alliance, "5G White Paper", NGMN 5G White Paper v1.0, February 2015, <https://www.ngmn.org/uploads/media/NGMN_5G_White_Paper_V1_0.pdf>.

[NGMN-fronth]
NGMN Alliance, "Fronthaul Requirements for C-RAN", March 2015, <https://www.ngmn.org/uploads/media/NGMN_RANEV_D1_C-RAN_Fronthaul_Requirements_v1.0.pdf>.

[OPCXML] OPC Foundation, "OPC XML-Data Access Specification", Dec 2004.

[PCE] IETF, "Path Computation Element", <https://datatracker.ietf.org/doc/charter-ietf-pce/>.

[profibus]
IEC, "IEC 61158 Type 3 - Profibus DP", January 2001.

[RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, DOI 10.17487/RFC3031, January 2001, <https://www.rfc-editor.org/info/rfc3031>.

[RFC3411] Harrington, D., Presuhn, R., and B. Wijnen, "An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks", STD 62, RFC 3411, DOI 10.17487/RFC3411, December 2002, <https://www.rfc-editor.org/info/rfc3411>.

[RFC3985] Bryant, S., Ed. and P. Pate, Ed., "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", RFC 3985, DOI 10.17487/RFC3985, March 2005, <https://www.rfc-editor.org/info/rfc3985>.

[RFC4553] Vainshtein, A., Ed. and YJ. Stein, Ed., "Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)", RFC 4553, DOI 10.17487/RFC4553, June 2006, <https://www.rfc-editor.org/info/rfc4553>.

[RFC5086] Vainshtein, A., Ed., Sasson, I., Metz, E., Frost, T., and P. Pate, "Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)", RFC 5086, DOI 10.17487/RFC5086, December 2007, <https://www.rfc-editor.org/info/rfc5086>.

[RFC5087] Stein, Y(J)., Shashoua, R., Insler, R., and M. Anavi, "Time Division Multiplexing over IP (TDMoIP)", RFC 5087, DOI 10.17487/RFC5087, December 2007, <https://www.rfc-editor.org/info/rfc5087>.

[RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, DOI 10.17487/RFC5905, June 2010, <https://www.rfc-editor.org/info/rfc5905>.

[RFC6550] Winter, T., Ed., Thubert, P., Ed., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, DOI 10.17487/RFC6550, March 2012, <https://www.rfc-editor.org/info/rfc6550>.

[RFC6551] Vasseur, JP., Ed., Kim, M., Ed., Pister, K., Dejean, N., and D. Barthel, "Routing Metrics Used for Path Calculation in Low-Power and Lossy Networks", RFC 6551, DOI 10.17487/RFC6551, March 2012, <https://www.rfc-editor.org/info/rfc6551>.

[RFC7554] Watteyne, T., Ed., Palattella, M., and L. Grieco, "Using IEEE 802.15.4e Time-Slotted Channel Hopping (TSCH) in the Internet of Things (IoT): Problem Statement", RFC 7554, DOI 10.17487/RFC7554, May 2015, <https://www.rfc-editor.org/info/rfc7554>.

[RFC8169] Mirsky, G., Ruffini, S., Gray, E., Drake, J., Bryant, S., and A. Vainshtein, "Residence Time Measurement in MPLS Networks", RFC 8169, DOI 10.17487/RFC8169, May 2017, <https://www.rfc-editor.org/info/rfc8169>.

[Spe09] Sperotto, A., Sadre, R., Vliet, F., and A. Pras, "A First Look into SCADA Network Traffic", IP Operations and Management, p. 518-521. , June 2009.

[SRP_LATENCY] Gunther, C., "Specifying SRP Latency", 2014, <http://www.ieee802.org/1/files/public/docs2014/cc-cgunther-acceptable-latency-0314-v01.pdf>.

[SyncE] ITU-T, "G.8261 : Timing and synchronization aspects in packet networks", Recommendation G.8261, August 2013, <http://www.itu.int/rec/T-REC-G.8261>.

[TR38501] 3GPP, "3GPP TS 38.501, Technical Specification System Architecture for the 5G System (Release 15)", 2017, <https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3144>.

[TR38801] 3GPP, "3GPP TR 38.801, Technical Specification Group Radio Access Network; Study on new radio access technology: Radio access architecture and interfaces (Release 14)", 2017, <https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3056>.

[TS23401] 3GPP, "General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access", 3GPP TS 23.401 10.10.0, March 2013.

[TS25104] 3GPP, "Base Station (BS) radio transmission and reception (FDD)", 3GPP TS 25.104 3.14.0, March 2007.

[TS36104] 3GPP, "Evolved Universal Terrestrial Radio Access (E-UTRA); Base Station (BS) radio transmission and reception", 3GPP TS 36.104 10.11.0, July 2013.

[TS36133] 3GPP, "Evolved Universal Terrestrial Radio Access (E-UTRA); Requirements for support of radio resource management", 3GPP TS 36.133 12.7.0, April 2015.

[TS36211] 3GPP, "Evolved Universal Terrestrial Radio Access (E-UTRA); Physical channels and modulation", 3GPP TS 36.211 10.7.0, March 2013.

[TS36300] 3GPP, "Evolved Universal Terrestrial Radio Access (E-UTRA) and Evolved Universal Terrestrial Radio Access Network (E-UTRAN); Overall description; Stage 2", 3GPP TS 36.300 10.11.0, September 2013.

[TSNTG] IEEE Standards Association, "IEEE 802.1 Time-Sensitive Networks Task Group", 2013, <http://www.IEEE802.org/1/pages/avbridges.html>.

[WirelessHART] www.hartcomm.org, "Industrial Communication Networks - Wireless Communication Network and Communication Profiles - WirelessHART - IEC 62591", 2010.

Appendix A. Use Cases Explicitly Out of Scope for DetNet

This section contains use case text that has been determined to be outside of the scope of the present DetNet work.

A.1. DetNet Scope Limitations

The scope of DetNet is deliberately limited to specific use cases that are consistent with the WG charter, subject to the interpretation of the WG. At the time the DetNet Use Cases were solicited and provided by the authors the scope of DetNet was not clearly defined, and as that clarity has emerged, certain of the use cases have been determined to be outside the scope of the present DetNet work. Such text has been moved into this section to clarify that these use cases will not be supported by the DetNet work.

The text in this section was moved here based on the following "exclusion" principles. Or, as an alternative to moving all such text to this section, some draft text has been modified in situ to reflect these same principles.

The following principles have been established to clarify the scope of the present DetNet work.

- The scope of network addressed by DetNet is limited to networks that can be centrally controlled, i.e. an "enterprise" aka "corporate" network. This explicitly excludes "the open Internet".

- Maintaining synchronized time across a DetNet network is crucial to its operation, however DetNet assumes that time is to be maintained using other means, for example (but not limited to) Precision Time Protocol ([IEEE1588]). A use case may state the accuracy and reliability that it expects from the DetNet network as part of a whole system, however it is understood that such timing properties are not guaranteed by DetNet itself. At the time of this writing it is an open question as to whether DetNet protocols will include a way for an application to communicate such timing expectations to the network, and if so whether they would be expected to materially affect the performance they would receive from the network as a result.

A.2. Internet-based Applications

There are many applications that communicate over the open Internet that could benefit from guaranteed delivery and bounded latency. However as noted above, all such applications when run over the open Internet are out of scope for DetNet. These same applications may be in-scope when run in constrained environments, i.e. within a centrally controlled DetNet network. The following are some examples of such applications.

A.2.1. Use Case Description

A.2.1.1. Media Content Delivery

Media content delivery continues to be an important use of the Internet, yet users often experience poor quality audio and video due to the delay and jitter inherent in today's Internet.

A.2.1.2. Online Gaming

Online gaming is a significant part of the gaming market, however latency can degrade the end user experience. For example "First Person Shooter" games are highly delay-sensitive.

A.2.1.3. Virtual Reality

Virtual reality has many commercial applications including real estate presentations, remote medical procedures, and so on. Low latency is critical to interacting with the virtual world because perceptual delays can cause motion sickness.

A.2.2. Internet-Based Applications Today

Internet service today is by definition "best-effort", with no guarantees on delivery or bandwidth.

A.2.3. Internet-Based Applications Future

An Internet from which one can play a video without glitches and play games without lag.

For online gaming, the maximum round-trip delay can be 100ms and stricter for FPS gaming which can be 10-50ms. Transport delay is the dominate part with a 5-20ms budget.

For VR, 1-10ms maximum delay is needed and total network budget is 1-5ms if doing remote VR.

Flow identification can be used for gaming and VR, i.e. it can recognize a critical flow and provide appropriate latency bounds.

A.2.4. Internet-Based Applications Asks

- Unified control and management protocols to handle time-critical data flow
- Application-aware flow filtering mechanism to recognize the timing critical flow without doing 5-tuple matching

- o Unified control plane to provide low latency service on Layer-3 without changing the data plane

- o OAM system and protocols which can help to provide E2E-delay sensitive service provisioning

A.3. Pro Audio and Video - Digital Rights Management (DRM)

This section was moved here because this is considered a Link layer topic, not direct responsibility of DetNet.

Digital Rights Management (DRM) is very important to the audio and video industries. Any time protected content is introduced into a network there are DRM concerns that must be maintained (see [CONTENT_PROTECTION]). Many aspects of DRM are outside the scope of network technology, however there are cases when a secure link supporting authentication and encryption is required by content owners to carry their audio or video content when it is outside their own secure environment (for example see [DCI]).

As an example, two techniques are Digital Transmission Content Protection (DTCP) and High-Bandwidth Digital Content Protection (HDCP). HDCP content is not approved for retransmission within any other type of DRM, while DTCP may be retransmitted under HDCP. Therefore if the source of a stream is outside of the network and it uses HDCP protection it is only allowed to be placed on the network with that same HDCP protection.

A.4. Pro Audio and Video - Link Aggregation

Note: The term "Link Aggregation" is used here as defined by the text in the following paragraph, i.e. not following a more common Network Industry definition.

For transmitting streams that require more bandwidth than a single link in the target network can support, link aggregation is a technique for combining (aggregating) the bandwidth available on multiple physical links to create a single logical link of the required bandwidth. However, if aggregation is to be used, the network controller (or equivalent) must be able to determine the maximum latency of any path through the aggregate link.

A.5. Pro Audio and Video - Deterministic Time to Establish Streaming

The DetNet Working Group has decided that guidelines for establishing a deterministic time to establish stream startup are not within scope of DetNet. If bounded timing of establishing or re-establish streams

is required in a given use case, it is up to the application/system to achieve this.

Author's Address

Ethan Grossman (editor)
Dolby Laboratories, Inc.
1275 Market Street
San Francisco, CA 94103
USA

Phone: +1 415 645 4726
Email: ethan.grossman@dolby.com
URI: http://www.dolby.com