

Internet Engineering Task Force (IETF)
Request for Comments: 8578
Category: Informational
ISSN: 2070-1721

E. Grossman, Ed.
DOLBY
May 2019

Deterministic Networking Use Cases

Abstract

This document presents use cases for diverse industries that 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 Deterministic Networking (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.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8578>.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	6
2. Pro Audio and Video	7
2.1. Use Case Description	7
2.1.1. Uninterrupted Stream Playback	8
2.1.2. Synchronized Stream Playback	9
2.1.3. Sound Reinforcement	9
2.1.4. Secure Transmission	10
2.1.4.1. Safety	10
2.2. Pro Audio Today	10
2.3. Pro Audio in the Future	10
2.3.1. Layer 3 Interconnecting Layer 2 Islands	10
2.3.2. High-Reliability Stream Paths	11
2.3.3. Integration of Reserved Streams into IT Networks ..	11
2.3.4. Use of Unused Reservations by Best-Effort Traffic ..	11
2.3.5. Traffic Segregation	11
2.3.5.1. Packet-Forwarding Rules, VLANs, and Subnets	12
2.3.5.2. Multicast Addressing (IPv4 and IPv6)	12
2.3.6. Latency Optimization by a Central Controller	12
2.3.7. Reduced Device Costs due to Reduced Buffer Memory ..	13
2.4. Pro Audio Requests to the IETF	13
3. Electrical Utilities	14
3.1. Use Case Description	14
3.1.1. Transmission Use Cases	14
3.1.1.1. Protection	14
3.1.1.2. Intra-substation Process Bus Communications	21
3.1.1.3. Wide-Area Monitoring and Control Systems ..	23
3.1.1.4. WAN Engineering Guidelines Requirement Classification	25

3.1.2. Generation Use Case	26
3.1.2.1. Control of the Generated Power	26
3.1.2.2. Control of the Generation Infrastructure ..	27
3.1.3. Distribution Use Case	32
3.1.3.1. Fault Location, Isolation, and Service Restoration (FLISR)	32
3.2. Electrical Utilities Today	33
3.2.1. Current Security Practices and Their Limitations ...	34
3.3. Electrical Utilities in the Future	35
3.3.1. Migration to Packet-Switched Networks	36
3.3.2. Telecommunications Trends	37
3.3.2.1. General Telecommunications Requirements ...	37
3.3.2.2. Specific Network Topologies of Smart-Grid Applications	38
3.3.2.3. Precision Time Protocol	38
3.3.3. Security Trends in Utility Networks	39
3.4. Electrical Utilities Requests to the IETF	41
4. Building Automation Systems (BASs)	41
4.1. Use Case Description	41
4.2. BASs Today	42
4.2.1. BAS Architecture	42
4.2.2. BAS Deployment Model	44
4.2.3. Use Cases for Field Networks	45
4.2.3.1. Environmental Monitoring	45
4.2.3.2. Fire Detection	46
4.2.3.3. Feedback Control	46
4.2.4. BAS Security Considerations	46
4.3. BASs in the Future	46
4.4. BAS Requests to the IETF	47
5. Wireless for Industrial Applications	47
5.1. Use Case Description	47
5.1.1. Network Convergence Using 6TiSCH	48
5.1.2. Common Protocol Development for 6TiSCH	48
5.2. Wireless Industrial Today	49
5.3. Wireless Industrial in the Future	49
5.3.1. Unified Wireless Networks and Management	49
5.3.1.1. PCE and 6TiSCH ARQ Retries	51
5.3.2. Schedule Management by a PCE	52
5.3.2.1. PCE Commands and 6TiSCH CoAP Requests	52
5.3.2.2. 6TiSCH IP Interface	54
5.3.3. 6TiSCH Security Considerations	54
5.4. Wireless Industrial Requests to the IETF	54

6. Cellular Radio	54
6.1. Use Case Description	54
6.1.1. Network Architecture	54
6.1.2. Delay Constraints	55
6.1.3. Time-Synchronization Constraints	57
6.1.4. Transport-Loss Constraints	59
6.1.5. Cellular Radio Network Security Considerations	60
6.2. Cellular Radio Networks Today	60
6.2.1. Fronthaul	60
6.2.2. Midhaul and Backhaul	60
6.3. Cellular Radio Networks in the Future	61
6.4. Cellular Radio Networks Requests to the IETF	64
7. Industrial Machine to Machine (M2M)	64
7.1. Use Case Description	64
7.2. Industrial M2M Communications Today	66
7.2.1. Transport Parameters	66
7.2.2. Stream Creation and Destruction	67
7.3. Industrial M2M in the Future	67
7.4. Industrial M2M Requests to the IETF	67
8. Mining Industry	68
8.1. Use Case Description	68
8.2. Mining Industry Today	68
8.3. Mining Industry in the Future	69
8.4. Mining Industry Requests to the IETF	70
9. Private Blockchain	70
9.1. Use Case Description	70
9.1.1. Blockchain Operation	71
9.1.2. Blockchain Network Architecture	71
9.1.3. Blockchain Security Considerations	72
9.2. Private Blockchain Today	72
9.3. Private Blockchain in the Future	72
9.4. Private Blockchain Requests to the IETF	72
10. Network Slicing	73
10.1. Use Case Description	73
10.2. DetNet Applied to Network Slicing	73
10.2.1. Resource Isolation across Slices	73
10.2.2. Deterministic Services within Slices	74
10.3. A Network Slicing Use Case Example - 5G Bearer Network ...	74
10.4. Non-5G Applications of Network Slicing	75
10.5. Limitations of DetNet in Network Slicing	75
10.6. Network Slicing Today and in the Future	75
10.7. Network Slicing Requests to the IETF	75
11. Use Case Common Themes	76
11.1. Unified, Standards-Based Networks	76
11.1.1. Extensions to Ethernet	76
11.1.2. Centrally Administered Networks	76
11.1.3. Standardized Data-Flow Information Models	76

11.1.4. Layer 2 and Layer 3 Integration	76
11.1.5. IPv4 Considerations	76
11.1.6. Guaranteed End-to-End Delivery	77
11.1.7. Replacement for Multiple Proprietary Deterministic Networks	77
11.1.8. Mix of Deterministic and Best-Effort Traffic	77
11.1.9. Unused Reserved Bandwidth to Be Available to Best-Effort Traffic	77
11.1.10. Lower-Cost, Multi-Vendor Solutions	77
11.2. Scalable Size	78
11.2.1. Scalable Number of Flows	78
11.3. Scalable Timing Parameters and Accuracy	78
11.3.1. Bounded Latency	78
11.3.2. Low Latency	78
11.3.3. Bounded Jitter (Latency Variation)	79
11.3.4. Symmetrical Path Delays	79
11.4. High Reliability and Availability	79
11.5. Security	79
11.6. Deterministic Flows	79
12. Security Considerations	80
13. IANA Considerations	80
14. Informative References	80
Appendix A. Use Cases Explicitly Out of Scope for DetNet	90
A.1. DetNet Scope Limitations	90
A.2. Internet-Based Applications	90
A.2.1. Use Case Description	91
A.2.1.1. Media Content Delivery	91
A.2.1.2. Online Gaming	91
A.2.1.3. Virtual Reality	91
A.2.2. Internet-Based Applications Today	91
A.2.3. Internet-Based Applications in the Future	91
A.2.4. Internet-Based Applications Requests to the IETF	92
A.3. Pro Audio and Video - Digital Rights Management (DRM)	92
A.4. Pro Audio and Video - Link Aggregation	92
A.5. Pro Audio and Video - Deterministic Time to Establish Streaming	93
Acknowledgments	93
Contributors	95
Author's Address	97

1. Introduction

This memo documents use cases for diverse industries that require deterministic flows over multi-hop paths. Deterministic Networking (DetNet) flows can be established from either a Layer 2 or Layer 3 (IP) interface, and such flows can coexist 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 for other DetNet documents.

The DetNet use cases, as originally submitted, explicitly were not considered by the DetNet Working Group (WG) to be concrete requirements. The DetNet WG and Design Team considered these use cases, identifying which of their elements could be feasibly implemented within the charter of DetNet; as a result, certain originally submitted use cases (or elements thereof) were moved to Appendix A ("Use Cases Explicitly Out of Scope for DetNet") of this document.

This document provides context regarding DetNet design decisions. It also serves a long-lived purpose of helping those learning (or new to) DetNet understand the types of applications that can be supported by DetNet. It also allows those WG contributors who are users to ensure that their concerns are addressed by the WG; for them, this document (1) covers their contributions and (2) provides a long-term reference regarding the problems that they expect will be served by the technology, in terms of the short-term deliverables and also as the technology evolves in the future.

This 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 industries covered by the use cases in this document are

- o professional audio and video (Section 2)
- o electrical utilities (Section 3)
- o building automation systems (BASs) (Section 4)
- o wireless for industrial applications (Section 5)
- o cellular radio (Section 6)

- o industrial machine to machine (M2M) (Section 7)
- o mining (Section 8)
- o private blockchain (Section 9)
- o network slicing (Section 10)

For each use case, the following questions are answered:

- o What is the use case?
- o How is it addressed today?
- o How should it be addressed in the future?
- o 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 no more 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 [IEEE-8021AS] and/or [RFC5905]).

Section 11 enumerates the set of common properties implied by these use cases.

2. Pro Audio and Video

2.1. Use Case Description

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

- o Music and film content creation
- o Broadcast
- o Cinema
- o Live sound
- o Public address, media, and emergency systems at large venues (e.g., 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 signal or a small number of signals per link, interconnected with purpose-built hardware.

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

Today, ProAV applications have no way to establish deterministic flows from a standards-based Layer 3 (IP) interface; this is a fundamental limitation of the use cases described here. Today, deterministic flows can be created within standards-based Layer 2 LANs (e.g., using IEEE 802.1 TSN ("TSN" stands for "Time-Sensitive Networking")); however, these flows 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 of more-limited scope (e.g., enterprise networks) would still provide substantial improvements.

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 in that uninterrupted stream playback in the presence of network errors cannot be achieved by retrying the transmission; by the time the missing or corrupt packet has been identified, it is too late to execute a retry 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. The use of redundant paths can further mitigate transmission errors and thereby provide greater stream reliability.

Additional techniques, such as Forward Error Correction (FEC), 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 stream and the video stream must be bounded and consistent if the sound is to remain matched to the movement in the video. A common tolerance for audio/video synchronization is one National Television System Committee (NTSC) video frame (about 33 ms); to maintain the audience's 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) enables measurement of the latency of each path and has 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 that 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 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) between the time when a person speaks into a microphone and when their voice emerges from the speaker. If this delay is longer than about 10-15 ms, 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 might be familiar with this experience.)

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

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 (33 ms for NTSC video).

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

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. If used incorrectly, such amplifiers 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 that enable deterministic streams at Layer 3; however, they are "engineered networks" that require careful configuration to operate and often require that the system be over-provisioned. Also, 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. Establishing relevant IETF standards is a crucial factor.

2.3. Pro Audio in the 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 "Digital Center 2" (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 TSN) 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 redundancy 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 should not compromise the 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 underutilized), 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 unreserve 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 1,000-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 via 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 Layer 2 bridges by design forward Media Access Control (MAC) addresses, it is important that a multicast MAC address only be 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.

In other words, since each multicast MAC address can represent 32 different IPv4 multicast addresses, there must be a process in place to make sure that any given multicast MAC address is only associated with exactly one IPv4 multicast address. Requiring the use of IPv6 addresses could help in this regard, due to the much larger address range of IPv6; however, due to the continued prevalence of IPv4 installations, solutions that are effective for IPv4 installations would be practical in many more use cases.

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 sinks that are farther away. 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, at 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 Costs due to Reduced Buffer Memory

Device costs 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 defined by the worst-case latency and jitter values of the worst-case segment of the end-to-end network path. For example, on today's 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 packets identical to those sent by the source and would be unaware of any latency or jitter issues along the path.

2.4. Pro Audio Requests to the IETF

- o Layer 3 routing on top of Audio Video Bridging (AVB) (and/or other high-QoS (Quality of Service) networks)
- o Content delivery with bounded, lowest possible latency
- o IntServ and DiffServ integration with AVB (where practical)
- o Single network for A/V and IT traffic
- o Standards-based, interoperable, multi-vendor solutions
- o IT-department-friendly networks
- o Enterprise-wide networks (e.g., the 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 for transmission, generation, and distribution, including key timing and reliability metrics. In addition, security issues and industry trends that 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 International Electrotechnical Commission (IEC) Standard 60834 [IEC-60834] as follows:

- o Transmission time (speed): The time between the moment when a state change occurs at the transmitter input and the moment of the corresponding change at the receiver output, including propagation delay. The overall operating time for a teleprotection system is the sum of (1) the time required to initiate the command at the transmitting end, (2) the propagation delay over the network (including equipment), and (3) the time required to make the necessary selections and decisions at the receiving end, including any additional delay due to a noisy environment.

- o Dependability: The ability to issue and receive valid commands in the presence of interference and/or noise, by minimizing the Probability of Missing Commands (PMC). Dependability targets are typically set for a specific Bit Error Rate (BER) level.
- o 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 BER level.

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

- o Network bandwidth
- o 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 a total fault clearance time of 100 ms. As a safety precaution, however, the actual operation time of protection systems is limited to 70-80% of this period, including fault recognition time, command transmission time, and line breaker switching time.

Some system components, such as large electromechanical switches, require a particularly long time to operate and take up the majority of the total clearance time, leaving only a 10 ms window for the telecommunications part of the protection scheme, independent of the distance of travel. Given the sensitivity of the issue, new networks impose requirements that are even more stringent: IEC Standard 61850-5:2013 [IEC-61850-5:2013] limits the transfer time for protection messages to 1/4-1/2 cycle or 4-8 ms (for 60 Hz lines) for messages considered the most critical.

3.1.1.1.3. Symmetric Channel Delay

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

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

- o 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.
- o For jitter-prone IP networks, traffic management tools can ensure that the teleprotection signals receive the highest transmission priority to minimize jitter.
- o Standard packet-based synchronization technologies, such as the IEEE 1588-2008 Precision Time Protocol (PTP) [IEEE-1588] and synchronous Ethernet (syncE) [syncE], 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

Table 1 captures the main network metrics. (These metrics are based on IEC Standard 61850-5:2013 [IEC-61850-5:2013].)

Teleprotection Requirement	Attribute
One-way maximum delay	4-10 ms
Asymmetric delay required	Yes
Maximum jitter	Less than 250 us (750 us for legacy IEDs)
Topology	Point to point, point to multipoint
Availability	99.9999%
Precise timing required	Yes
Recovery time on node failure	Less than 50 ms - 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
Asymmetric delay required	No
Maximum jitter	Not critical
Topology	Point to point, point to multipoint
Bandwidth	64 kbps
Availability	99.9999%
Precise timing required	Yes
Recovery time on node failure	Less than 50 ms - 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 typically used to protect parallel circuits. At both ends of the lines, the current is measured by the differential relays; 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 that some form of communication is present between the relays at both ends of the line, to allow both relays to compare measured current values. Line differential protection schemes assume that the telecommunications delay between both relays is very low -- often as low as 5 ms. Moreover, as those systems are

often not time-synchronized, they also assume that the delay over symmetric telecommunications paths is constant; this allows the comparison of current measurement values taken at exactly the same time.

Current Differential Protection Requirement	Attribute
One-way maximum delay	5 ms
Asymmetric delay required	Yes
Maximum jitter	Less than 250 us (750 us for legacy IEDs)
Topology	Point to point, point to multipoint
Bandwidth	64 kbps
Availability	99.9999%
Precise timing required	Yes
Recovery time on node failure	Less than 50 ms - 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

The distance (impedance relay) protection scheme is based on voltage and current measurements. The network metrics are similar (but not identical) to the metrics for current differential protection.

Distance Protection Requirement	Attribute
One-way maximum delay	5 ms
Asymmetric delay required	No
Maximum jitter	Not critical
Topology	Point to point, point to multipoint
Bandwidth	64 kbps
Availability	99.9999%
Precise timing required	Yes
Recovery time on node failure	Less than 50 ms - 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 values and/or GOOSE (Generic Object Oriented Substation Events) messages between Intelligent Electronic Devices (IEDs) in two substations for protection and tripping coordination. The two IEDs are in 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 sampled values (as specified by IEC 61850-9-2:2011 [IEC-61850-9-2:2011]) 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 message containing the command to trip the breaker, it opens the breaker. It then sends a confirmation message back to the master. All data exchanges between IEDs are through sampled values and/or GOOSE messages.

Inter-substation Protection Requirement	Attribute
One-way maximum delay	5 ms
Asymmetric delay required	No
Maximum jitter	Not critical
Topology	Point to point, point to multipoint
Bandwidth	64 kbps
Availability	99.9999%
Precise timing required	Yes
Recovery time on node failure	Less than 50 ms - 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 sends 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 61850-9-2:2011 [IEC-61850-9-2:2011]) and sends them to the IEDs in the substation. The Global Positioning System (GPS) Master Clock can send 1PPS or IRIG-B format to the MU through a serial port or IEEE 1588 protocol

via a network. 1PPS (One Pulse Per Second) is an electrical signal that has a width of less than 1 second and a sharply rising or abruptly falling edge that accurately repeats once per second. 1PPS signals are output by radio beacons, frequency standards, other types of precision oscillators, and some GPS receivers. IRIG (Inter-Range Instrumentation Group) time codes are standard formats for transferring timing information. Atomic frequency standards and GPS receivers designed for precision timing are often equipped with an IRIG output. Process bus communication using IEC 61850-9-2:2011 [IEC-61850-9-2:2011] simplifies connectivity within the substation, 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
Asymmetric delay required	No
Maximum jitter	Not critical
Topology	Point to point, point to multipoint
Bandwidth	64 kbps
Availability	99.9999%
Precise timing required	Yes
Recovery time on node failure	Less than 50 ms - hitless
Performance management	Yes; mandatory
Redundancy	Yes or 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 (PMUs) 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 the real-time nature of synchrophasor data presents unique challenges for existing application architectures. The 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, the WAMS has stringent telecommunications requirements in an IP network, as captured in Table 7:

WAMS Requirement	Attribute
One-way maximum delay	50 ms
Asymmetric delay required	No
Maximum jitter	Not critical
Topology	Point to point, point to multipoint, multipoint to multipoint
Bandwidth	100 kbps
Availability	99.9999%
Precise timing required	Yes
Recovery time on node failure	Less than 50 ms - hitless
Performance management	Yes; mandatory
Redundancy	Yes
Packet loss	1%
Consecutive packet loss	At least one packet per application cycle must be received.

Table 7: WAMS Special Communication Requirements

3.1.1.4. WAN Engineering Guidelines Requirement Classification

The IEC has published a technical report (TR) that offers guidelines on how to define and deploy Wide-Area Networks (WANs) for the interconnection of electric substations, generation plants, and SCADA operation centers. IEC TR 61850-90-12:2015 [IEC-61850-90-12:2015] provides four classes of WAN communication requirements, as summarized in Table 8:

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 asymmetry	100 us	1 ms	10 ms	100 ms
Time accuracy	1 us	10 us	100 us	10 to 100 ms
BER	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
Cybersecurity	Extremely high	High	Medium	Medium

Table 8: Communication Requirements (Courtesy of IEC TR 61850-90-12:2015)

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
Asymmetric 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 describes the use case for control of the generation infrastructure of a wind turbine.

Figure 1 presents the subsystems that operate a wind turbine.

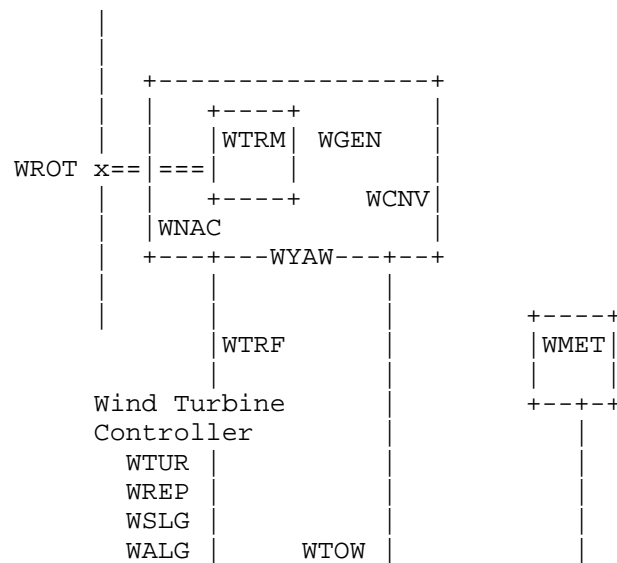


Figure 1: Wind Turbine Control Network

The subsystems shown in Figure 1 include the following:

- o WROT (rotor control)
- o WNAC (nacelle control) (nacelle: housing containing the generator)
- o WTRM (transmission control)
- o WGEN (generator)
- o WYAW (yaw controller) (of the tower head)
- o WCNV (in-turbine power converter)
- o WTRF (wind turbine transformer information)

- o WMET (external meteorological station providing real-time information to the tower's controllers)
- o WTUR (wind turbine general information)
- o WREP (wind turbine report information)
- o WSLG (wind turbine state log information)
- o WALG (wind turbine analog log information)
- o WTOW (wind turbine tower information)

Traffic characteristics relevant to 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 that must comply with their respective data-rate constraints.

Subsystem	Sensor Count	Analog Sample Count	Data Rate (bytes/s)	Status Sample Count	Data Rate (bytes/s)
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

QoS constraints for different services are presented in Table 11. These constraints are defined by IEEE Standard 1646 [IEEE-1646] and IEC Standard 61400 Part 25 [IEC-61400-25].

Service	Latency	Reliability	Packet Loss Rate
Analog measurement	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 that require time-critical operation. The reliability and latency constraints of these different subsystems are shown in Table 11. These subsystems are connected to an intra-domain network that 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 [EtherCAT], or a combination thereof. Industrial signaling and control protocols such as Modbus [MODBUS], PROFIBUS [PROFIBUS], PROFINET [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 at the bottom of the tower and also 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 Autonomous System (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.




Figure 2: Wind Turbine Control via Internet

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-Data Access

(XML-DA) [OPCXML], Modbus [MODBUS], and SNMP [RFC3411]. At the time of this writing, traffic flows between the remote control center and the wind park are best effort. QoS requirements are not strict, so no Service Level Agreements (SLAs) or service-provisioning mechanisms (e.g., VPNs) are employed. In the case of such events as equipment failure, tolerance for alarm delay is on the order of minutes, due to redundant systems already in place.

Future use cases will require bounded latency, bounded jitter, and extraordinarily low packet loss for inter-domain traffic flows due to the softwarization and virtualization of core wind-park 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 failover of a local SCADA server, a SCADA server in another wind-park 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.

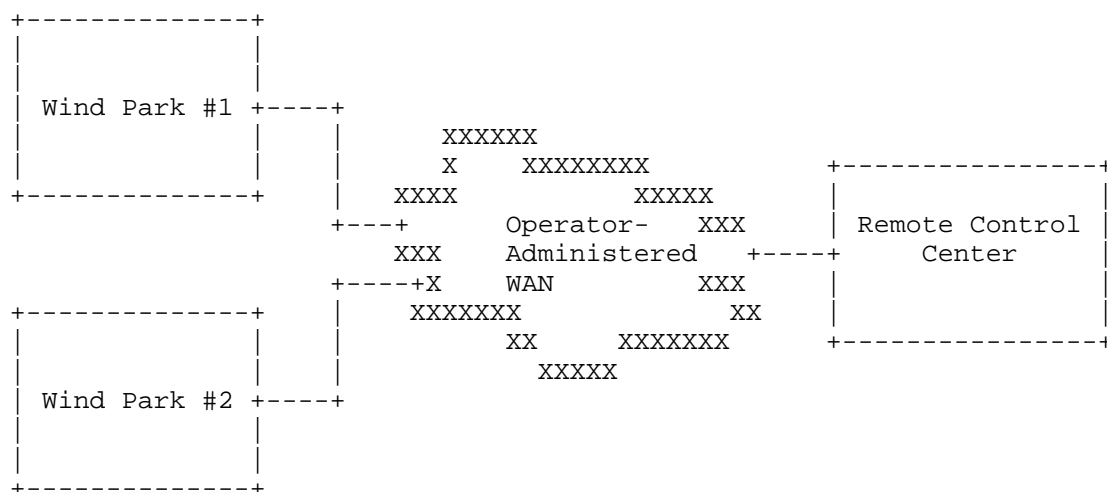


Figure 3: Wind Turbine Control via Operator-Administered WAN

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.

The 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
Asymmetric delay required	No
Maximum jitter	40 ms
Topology	Point to point, point to multipoint, multipoint to multipoint
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 consisting of multiple application-specific proprietary networks, including TDM networks.

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

Specific calibration of the full chain is required; this is costly.

This kind of environment prevents utility operations from realizing 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. Current Security Practices and Their Limitations

Grid-monitoring and control devices are already targets for cyber attacks, and legacy telecommunications protocols have many intrinsic network-related vulnerabilities. For example, the Distributed Network Protocol (DNP3) [IEEE-1815], 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 perform data retrieval (reading inputs) or control functions (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 functionality similar to command and control. Specific security concerns common to most industrial-control protocols (including utility telecommunications protocols) include the following:

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

- o Protocol commands may be available that will query the device network to obtain defined points and their values (i.e., perform a configuration scan).
- o Protocol commands may be available that will list all available function codes (i.e., perform a function scan).

These inherent vulnerabilities, along with increasing connectivity between IT and OT networks, make network-based attacks very feasible. By injecting malicious protocol commands, an attacker could take 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 result in (1) improper control over a process and (2) misrepresentation of data that is sent back to operator consoles.

3.3. Electrical Utilities in the 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's 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 telecommunications network will be based on an 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 has already been demonstrated in many mission-critical and highly secure networks.

IPv6 is seen as a future telecommunications technology for the smart grid; the IEC and different national committees have mandated a specific ad hoc group (AHG8) to define the strategy for migration to IPv6 for all the IEC Technical Committee 57 (TC 57) power automation standards. The AHG8 has finalized its work on the migration strategy, and IEC TR 62357-200:2015 [IEC-62357-200:2015] has been issued.

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

3.3.1. Migration to Packet-Switched Networks

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 WAN, made possible by technologies such as Multiprotocol Label Switching (MPLS). The data that traverses the utility WAN includes:

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

WANs support this wide variety of traffic to and from substations, the transmission and distribution grid, and 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 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 provided 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

- o IP connectivity everywhere
- o Monitoring services everywhere, and from different remote centers
- o Moving services to a virtual data center
- o Unified access to applications/information from the corporate network
- o Unified services
- o Unified communications solutions
- o Mix of fiber and microwave technologies - obsolescence of the Synchronous Optical Network / Synchronous Digital Hierarchy (SONET/SDH) or TDM
- o Standardizing grid telecommunications protocols to open standards, to ensure interoperability
- o Reliable telecommunications for transmission and distribution substations
- o IEEE 1588 time-synchronization client/server capabilities
- o Integration of multicast design
- o Mapping of QoS requirements
- o Enabling future network expansion
- o Substation network resilience
- o Fast convergence design
- o Scalable headend design
- o Defining SLAs and enabling SLA monitoring

- o Integration of 3G/4G technologies and future technologies
- o Ethernet connectivity for station bus architecture
- o Ethernet connectivity for process bus architecture
- o Protection, teleprotection, and PMUs on IP

3.3.2.2. Specific Network Topologies of Smart-Grid Applications

Utilities often have very large private telecommunications networks that can cover an entire territory/country. Until now, the main purposes of these networks have been to (1) support transmission network monitoring, control, and automation, (2) support remote control of generation sites, and (3) provide FCAPS (Fault, Configuration, Accounting, Performance, and Security) services from centralized network operation centers.

Going forward, one network will support the operation and maintenance of electrical networks (generation, transmission, and distribution), voice and data services for tens of thousands of employees and for exchanges with neighboring interconnections, and administrative services. To meet those requirements, a utility may deploy several physical networks leveraging different technologies across the country -- for instance, an optical network and a microwave network. Each protection and automation 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 links, and telephone cables. This network is the nervous system of the utility's power transmission operations. The optical network represents tens of thousands of kilometers 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 underground 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 1 ms timestamps for IRIG-B.

Some companies plan to transition to PTP [IEEE-1588], 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:

- o The network eliminates cyclic forwarding of PTP messages within each communication path (e.g., by using a spanning tree protocol).
- o 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.
- o As designed, PTP expects a multicast communication model; however, PTP also supports a unicast communication model as long as the behavior of the protocol is preserved.
- o Like all message-based time transfer protocols, PTP time accuracy is degraded by delay asymmetry in the paths taken by event messages. PTP cannot detect asymmetry, but if such delays are known a priori, time values can be adjusted to correct for asymmetry.

The use of PTP for power automation is defined in IEC/IEEE 61850-9-3:2016 [IEC-IEEE-61850-9-3:2016]. It is based on Annex B of IEC 62439-3:2016 [IEC-62439-3:2016], 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:

- o Integration of distributed energy resources
- o Proliferation of digital devices to enable management, automation, protection, and control

- o Regulatory mandates to comply with standards for critical infrastructure protection
- o Migration to new systems for outage management, distribution automation, condition-based maintenance, load forecasting, and smart metering
- o 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 customers, and 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.

"Cybersecurity" 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:

- o Integrity: data cannot be altered undetectably
- o Authenticity (data origin authentication): the telecommunications parties involved must be validated as genuine
- o Authorization: only requests and commands from authorized users can be accepted by the system
- o 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. The consequences of a cyber attack on the grid telecommunications network can be catastrophic. This is why security for the smart grid is not just an ad hoc feature or product; it's a complete framework integrating both physical and cybersecurity 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 addressing these challenges for two reasons:

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

Securing OT telecommunications over packet-switched IP networks follows the same principles that are foundational for securing the IT infrastructure, i.e., consideration must be given to (1) enforcing electronic access control for both person-to-machine and machine-to-machine communications and (2) providing the appropriate levels of data privacy, device and platform integrity, and threat detection and mitigation.

3.4. Electrical Utilities Requests to the IETF

- o Mixed Layer 2 and Layer 3 topologies
- o Deterministic behavior
- o Bounded latency and jitter
- o Tight feedback intervals
- o High availability, low recovery time
- o Redundancy, low packet loss
- o Precise timing
- o Centralized computing of deterministic paths
- o Distributed configuration (may also be useful)

4. Building Automation Systems (BASs)

4.1. Use Case Description

A 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:

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

4.2. BASs Today

4.2.1. BAS Architecture

A typical present-day BAS architecture is shown in Figure 4.




Figure 4: BAS Architecture

There are typically two layers of a network in a BAS. The upper layer is called the management network, and the lower layer 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.

An 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.

- o Collects and stores device states from LCs at regular intervals.
- o Sends control values to LCs according to a building schedule.
- o Sends an alarm signal to operators if it detects abnormal device states.

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

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

- o Measures device states and provides the information to a BMS or HMI.
- o Sends control values to devices, unilaterally or as part of a feedback control loop.

At the time of this writing, many field protocols are in use; some are standards-based protocols, and others are proprietary (see standards [LonTalk], [MODBUS], [PROFIBUS], and [FL-net]). The result is that BASs have multiple MAC/PHY modules and interfaces. This makes BASs more expensive and 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 includes a monitoring room where the BAS management entities are located. Each floor will have one or more LCs, depending on 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 not a real-time network, the use of Ethernet without QoS is sufficient for today's deployments.

Many physical interfaces used in field networks have specific timing requirements -- for example, RS232C and RS485. Thus, if a field network is to be replaced with an Ethernet or wireless network, such networks must support time-critical deterministic flows.

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

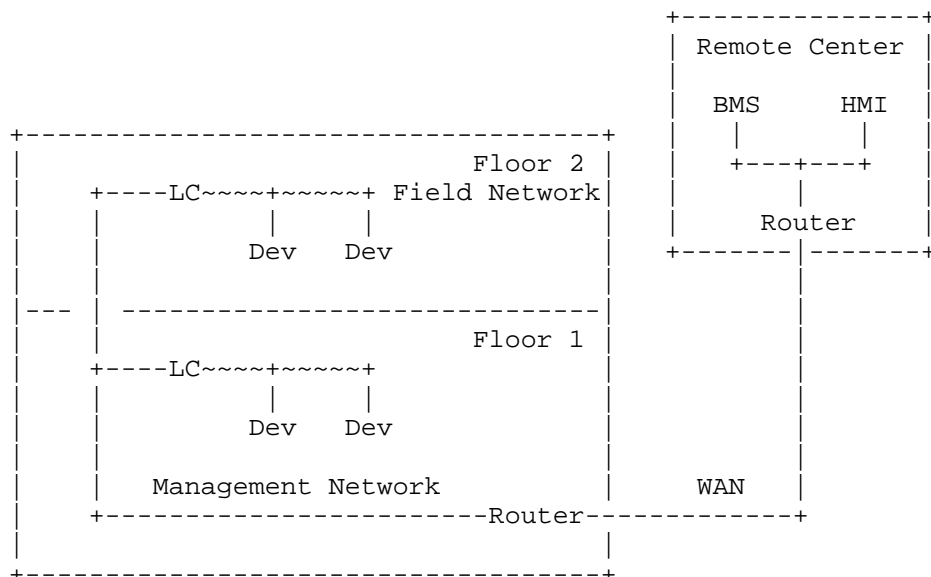


Figure 6: Deployment Model for Small Buildings

Some interoperability is possible in today's management networks but is not possible 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 100 ms (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 value measurements are completed within 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 tens of fire sensors per LC that the BMS needs to manage. In this scenario, the measurement interval is 10-50 ms, the communication delay is 10 ms, and the availability must be 99.9999%.

4.2.3.3. Feedback Control

BASs utilize feedback control in various ways; the most time-critical is control of DC motors, which require a short feedback interval (1-5 ms) with low communication delay (10 ms) and jitter (1 ms). The feedback interval depends on the characteristics of the device and on the requirements for the control values. There are typically tens of feedback sensors per LC.

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

4.2.4. BAS Security Considerations

When BAS field networks were developed, it was assumed that the field networks would always be physically isolated from external networks; therefore, security was not a concern. In today's world, many BASs are managed remotely and are thus connected to shared IP networks; therefore, security is a definite concern. Note, however, that security features are not currently 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 BASs do not implement even such available security features as device authentication or encryption for data in transit.

4.3. BASs in the Future

In the future, lower energy consumption and environmental monitoring that is more fine-grained will emerge; these 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 networks, home networks, and the 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 Requests to the IETF

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

- o Low communication delays (from <10 ms to 100 ms in a network of several hundred devices)
- o Low jitter (<1 ms)
- o Tight feedback intervals (1-10 ms)
- o High network availability (up to 99.9999%)
- o Availability of network data in disaster scenarios
- o Authentication between management devices and field devices (both local and remote)
- o Integrity and data origin authentication of communication data between management devices and field devices
- o 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, (1) when portable, fast-moving, or rotating objects are involved and (2) 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 QoS, as well as such specific network properties 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 that provides the required deterministic QoS: "IPv6 over the TSCH mode of IEEE 802.15.4e" (6TiSCH, where "TSCH" stands for "Time-Slotted Channel Hopping"; see [Arch-for-6TiSCH], [IEEE-802154], and [RFC7554]).

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

Thus, the primary goal of this use case is to apply 6TiSCH as a converged IP-based 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 that are still in development. Another goal 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 that contribute to the efficiency of enabling both DetNet and 6TiSCH.

One such commonality is that -- although on a different time scale -- in both TSN [IEEE-8021TSNTG] and TSCH, a packet that crosses the network from node to node follows a precise schedule, as does a train that leaves intermediate stations at precise times along its path. This kind of operation reduces collisions, saves energy, and enables engineering of 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 and 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; it will have limited bandwidth between peers such that it can maintain only a small amount of peer information, and it will not be able to store many packets waiting to be forwarded. It is advantageous, then, for the IoT device to only be required to carry out the specific behavior assigned to it by the PCE and 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 [DetNet-Arch].

6TiSCH also depends on the fact that DetNet will maintain consistency with [IEEE-8021TSNTG].

5.2. Wireless Industrial Today


Today, industrial wireless technology ("wireless industrial") is accomplished using multiple deterministic wireless networks that 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 in the Future

5.3.1. Unified Wireless Networks and Management

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



LLN: Low-Power and Lossy Network

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 [DetNet-Arch] 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 IEEE 802.1 TSN Ethernet backbone, and DetNet protocols are expected to enable end-to-end deterministic forwarding.

5.3.1.1. PCE and 6TiSCH ARQ Retries

6TiSCH uses the Automatic Repeat reQuest (ARQ) mechanism [IEEE-802154] to provide higher reliability of packet delivery. ARQ is related to Packet Replication and Elimination (PRE) because there are two independent paths for packets to arrive at the destination. If an expected packet 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 technique might be appropriate for DetNet, and aspects of the enabling protocol could therefore 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 an IEEE 802.1 TSN backbone.

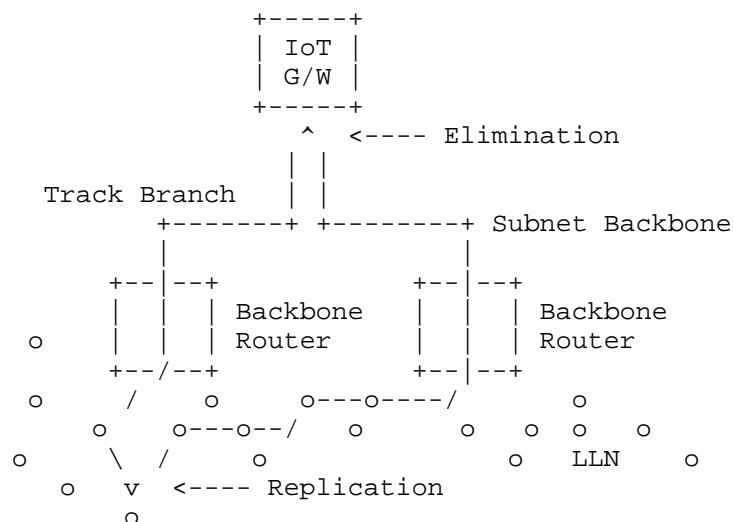


Figure 9: 6TiSCH Network with PRE

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 the case of a loss on one branch, one hopes that the other copy of the packet will still arrive 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).

At the time of this writing, a deployment's TSCH track does not necessarily support PRE but is systematically multipath. This means that a track is scheduled so as to ensure that each hop has at least two forwarding solutions. The forwarding decision will be to try the preferred solution and use the other solution in the 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 actions taken by a PCE when configuring 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. Specifically, both DetNet and 6TiSCH need to develop a protocol (and associated data model) that the PCE can use to (1) get/set the relevant configuration from/to the devices and (2) perform operations on the devices. These could be initially developed by DetNet, with consideration for their 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 traffic requirements 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, 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, i.e., a schedule that defines the behavior

of the node with respect to all data flows through that node. 6TiSCH expects that the programming of the schedule will be done over the Constrained Application Protocol (CoAP) as discussed in [CoAP-6TiSCH].

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 the Path Computation Element Communication Protocol (PCEP) into the RESTful interfaces that the 6TiSCH devices support. This architecture will be refined to comply with DetNet [DetNet-Arch] when the work is formalized. Related information about 6TiSCH can be found in [Interface-6TiSCH-6top] and [RFC6550] ("RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks").

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 [Arch-for-6TiSCH].

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 exceed the available energy over a period of time. Note that this statement implies that an extensible protocol for communicating device information 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 neighbor 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, metrics similar to those used for RPL operations [RFC6551].

5.3.2.2. 6TiSCH IP Interface

Protocol translation between the TSCH MAC layer and IP is accomplished via the "6top" sublayer [Sublayer-6TiSCH-6top]. The 6top data model and management interfaces are further discussed in [Interface-6TiSCH-6top] and [CoAP-6TiSCH].

An IP packet that is sent along a 6TiSCH path uses a differentiated services Per-Hop Behavior Group (PHB) called "deterministic forwarding", as described in [Det-Fwd-PHB].

5.3.3. 6TiSCH Security Considerations

In addition to 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 Requests to the IETF

6TiSCH depends on DetNet to define:

- o Configuration (state) and operations for deterministic paths
- o End-to-end protocols for deterministic forwarding (tagging, IP)
- o A protocol for PRE

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 to establish 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. The architecture includes "Fronthaul", "Midhaul", and "Backhaul" network segments. The "Fronthaul" is the network connecting base stations (Baseband Units (BBUs)) to the Remote Radio Heads (RRHs) (also referred to here as "antennas"). The "Midhaul" is the network that interconnects 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).

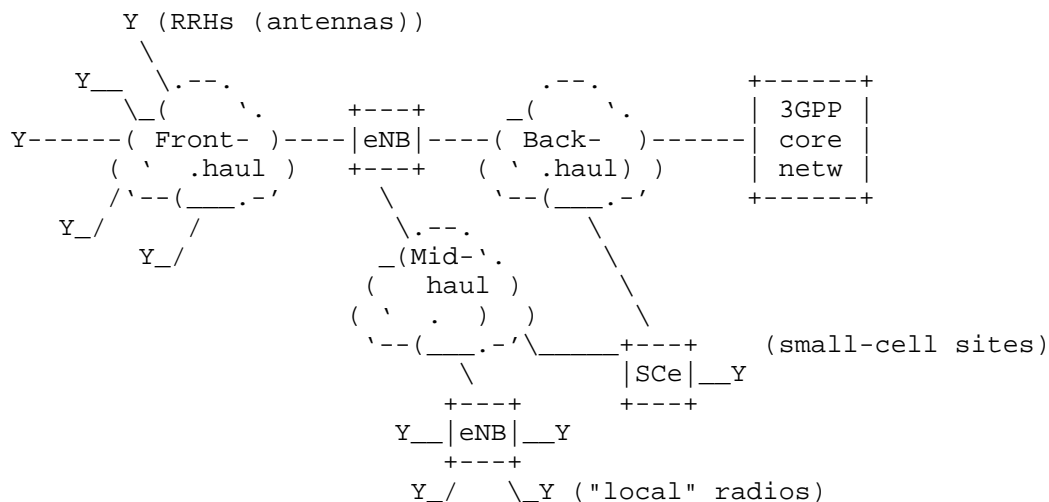


Figure 10: Generic 3GPP-Based Cellular Network Architecture

In Figure 10, "eNB" ("E-UTRAN Node B") is the hardware that is connected to the mobile phone network and enables the mobile phone network to communicate with mobile handsets [TS36300]. ("E-UTRAN" stands for "Evolved Universal Terrestrial Radio Access Network".)

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, 3 ms is allocated for the processing of a radio frame, but typically the baseband processing uses most of it, allowing only a small fraction to be used by the Fronthaul network. In this example, out of 3 ms, the maximum time allocated to the Fronthaul network for one-way delay is 250 us, and the existing specification [NGMN-Fronth] specifies a maximum delay of only 100 us. This ultimately determines the distance the RRHs can be located from the base stations (e.g., 100 us equals roughly 20 km of optical fiber-based transport). Allocation options regarding the available time budget between processing and transport are currently undergoing heavy discussion in the mobile industry.

For packet-based transport, the allocated transport time between the RRH and the BBU is consumed by node processing, buffering, and distance-incurred delay. An example of the allocated transport time is 100 us (from the Common Public Radio Interface [CPRI]).

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 RRHs can be located from base stations (i.e., the link length). In the above analysis, it is assumed that the entire transport time budget is available for link propagation delay. However, the transport time budget can be broken down into three components: scheduling/queuing 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 is shared among multiple Fronthaul flows -- for example, in indoor and small-cell networks, massive Multiple Input Multiple Output (MIMO) antenna networks, and split Fronthaul architectures.

DetNet technology can improve Fronthaul networks by controlling and reducing the time required for the queuing, scheduling, and transmission operations by properly assigning network resources, thus (1) leaving more of the transport time budget available for link propagation and (2) enabling longer link lengths. However, link length is usually a predetermined parameter and is not a controllable network parameter, since RRH and BBU sites are usually located in predetermined locations. However, the number of antennas in an RRH site might increase -- for example, by adding more antennas, increasing the MIMO capability of the network, or adding support for massive MIMO. This means increasing the number of 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 can 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 between flows -- for example, between flows originating in different slices within a network-sliced 5G network. Note, however, that this isolation applies to DetNet flows for which resources have been preallocated, i.e., it does not apply to best-effort flows within a DetNet. DetNet technology can also dynamically control the bandwidth-assignment, scheduling, and

packet-forwarding decisions, as well as 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 systems 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 cost and energy-consumption levels similar to today's systems).

For Midhaul connections, delay constraints are driven by inter-site radio functions such as Coordinated Multi-Point (CoMP) processing (see [CoMP]). CoMP reception and transmission constitute a framework in which multiple geographically distributed antenna nodes cooperate to improve performance for the users served in the common cooperation area. The design principle of CoMP is to extend single-cell-to-multi-UE (User Equipment) transmission to a multi-cell-to-multi-UE transmission via cooperation among base stations.

CoMP has delay-sensitive performance parameters: "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 1-10 ms. This delay requirement is both rigid and absolute, because any uncertainty in delay will degrade performance significantly.

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

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:

+/-8 ns (i.e., $\pm 1/32 T_c$, where T_c is the Universal Mobile Telecommunications System (UMTS) Chip time of $1/3.84$ MHz), resulting in a round-trip accuracy of +/-16 ns. The value is this low in order to meet the 3GPP Timing Alignment Error (TAE) measurement requirements. Note that 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.

TAE:

TAE is problematic for Fronthaul networks and must be minimized. If the transport network cannot guarantee TAE levels that are low enough, 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 measurements for packet-based Fronthaul networks.

- * For MIMO or TX diversity transmissions, at each carrier frequency, TAE measurements shall not exceed 65 ns (i.e., $1/4 T_c$).
- * For intra-band contiguous carrier aggregation, with or without MIMO or TX diversity, TAE measurements shall not exceed 130 ns (i.e., $1/2 T_c$).
- * For intra-band non-contiguous carrier aggregation, with or without MIMO or TX diversity, TAE measurements shall not exceed 260 ns (i.e., $1 T_c$).
- * For inter-band carrier aggregation, with or without MIMO or TX diversity, TAE measurements shall not exceed 260 ns.

Transport link contribution to radio frequency errors:

+/-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 that the transport link contributes to radio frequency errors for the following reason: at the time of this writing, Fronthaul communication is direct communication from the radio unit to the RRH. The RRH is essentially a passive device (e.g., without buffering). The transport drives the antenna directly by feeding it with samples, and everything the transport adds will be introduced to the radio "as is". So, if the transport causes any additional frequency errors, the errors will show up immediately on the radio as well. Note that 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 are more difficult to meet 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 that transport is almost error free. Errors can cause a reset of the radio interfaces, in turn causing 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-layer splits may imply an FLR of less than 10^{-7} for data traffic and less than 10^{-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 or using FEC to circumvent bit errors (or both) is practically impossible, due to the additional delay incurred. Using redundant streams for better guarantees of delivery is also practically impossible in many cases, due to high bandwidth requirements for 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 a reset of mobile interfaces.

It is assumed that Fronthaul links are symmetric. All Fronthaul streams (i.e., those carrying radio data) have equal priority and cannot delay or preempt each other.

All of this implies that it is up to the network to guarantee that each time-sensitive flow meets its schedule.

6.1.5. Cellular Radio Network 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:

- o Dedicated point-to-point fiber connection (common)
- o Proprietary protocols and framings
- o 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:

- o Mostly normal IP networks, MPLS-TP, etc.
- o Clock distribution and synchronization using IEEE 1588 and syncE

Telecommunications networks in the Midhaul and Backhaul are already heading towards transport networks where precise time-synchronization support is one of the basic building blocks. In order to meet

bandwidth and cost requirements, most transport networks have already transitioned to all-IP packet-based networks; however, highly accurate clock distribution has become a challenge.

In the past, Midhaul and Backhaul connections were typically based on TDM and provided frequency-synchronization capabilities as a part of the transport media. More recently, other technologies such as GPS or syncE [syncE] have been used.

Ethernet, IP/MPLS [RFC3031], and pseudowires (as described in [RFC3985] ("Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture") for legacy transport support)) have become popular tools for building and managing new all-IP Radio Access Networks (RANs) [SR-IP-RAN-Use-Case]. Although various timing and synchronization optimizations have already been proposed and implemented, including PTP enhancements [IEEE-1588] (see also [Timing-over-MPLS] and [RFC8169]), these solutions are not necessarily sufficient for the forthcoming RAN architectures, nor do they guarantee the more stringent time-synchronization requirements such as [CPRI].

Existing solutions for TDM over IP include those discussed in [RFC4553], [RFC5086], and [RFC5087]; [MEF8] addresses TDM over Ethernet transports.

6.3. Cellular Radio Networks in the Future

Future cellular radio networks will be based on a mix of different xHaul networks (xHaul = Fronthaul, Midhaul, and Backhaul), and future transport networks should be able to support all of them simultaneously. It is already envisioned today that:

- o 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 and Layer 3) with the same tools and allow their transport simultaneously.
- o All types of xHaul networks will need some types 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).
- o Different functional splits between the base stations and the on-site units could coexist on the same Fronthaul and Backhaul network.

Future cellular radio networks should contain the following:

- o Unified standards-based transport protocols and standard networking equipment that can make use of underlying deterministic link-layer services
- o Unified and standards-based network management systems and protocols in all parts of the network (including Fronthaul)

New RAN deployment models and architectures may require TSN 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 also becoming a real issue for Fronthaul networks. Specifically, in Fronthaul networks, the timing and synchronization requirements can be extreme for packet-based technologies -- for example, on the order of a PDV of ± 20 ns or less and frequency accuracy of ± 0.002 PPM [Fronthaul].

The actual transport protocols and/or solutions for establishing required transport "circuits" (pinned-down paths) for Fronthaul traffic are still undefined. Those protocols 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 TSN has been done for Ethernet [IEEE-8021TSNTG]; this work specifies the use of PTP [IEEE-1588] in the context of IEEE 802.1D and IEEE 802.1Q. [IEEE-8021AS] specifies a Layer 2 time-synchronizing service, and other specifications such as IEEE 1722 [IEEE-1722] 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 Fronthaul requirements into account is desirable [IEEE-8021CM].

New promising work seeks to enable the transport of time-sensitive Fronthaul streams in Ethernet bridged networks [IEEE-8021CM]. Analogous to IEEE 1722, standardization efforts in the IEEE 1914.3 Task Force [IEEE-19143] to define the Layer 2 transport encapsulation format for transporting Radio over Ethernet (RoE) are ongoing.

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 5G "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 enhanced CPRI (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 when it comes to providing Layer 3 services. The protocol stack typically has a number of layers below Ethernet Layer 2 that might be "visible" to Layer 3. In a fairly common scenario, on top of the lowest-layer (optical) transport is the first (lowest) Ethernet layer, then one or more layers of MPLS, pseudowires, and/or other tunneling protocols, and finally one or more Ethernet layers that are visible to Layer 3.

Although there exist technologies for establishing circuits through the routed and switched networks (especially in the 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 solution in question also must address the use cases where the user-plane traffic is on a different layer (for example, Ethernet frames).

6.4. Cellular Radio Networks Requests to the IETF

A standard for data-plane transport specifications that is:

- o 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)
- o Deployed in a highly deterministic network environment
- o 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" (Section 11); of special note are Sections 11.1.8 ("Mix of Deterministic and Best-Effort Traffic"), 11.3.1 ("Bounded Latency"), 11.3.2 ("Low Latency"), 11.3.4 ("Symmetrical Path Delays"), and 11.6 ("Deterministic Flows")
- o Capable of supporting network slicing and multi-tenancy; these goals can be supported by the same DetNet themes noted above
- o Capable of transporting both in-band and out-of-band control traffic (e.g., Operations, Administration, and Maintenance (OAM) information)
- o Deployable over multiple data-link technologies (e.g., IEEE 802.3, mmWave)

A standard for data-flow information models that is:

- o Aware of the time sensitivity and constraints of the target networking environment
- o 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 M2M use case focuses on machine units on a plant floor that periodically exchange data with upstream or downstream machine modules and/or a supervisory controller within a LAN.

PLCs are the "actors" in M2M communications. Communication between PLCs, and between PLCs and the supervisory PLC (S-PLC), is achieved via critical control/data streams (Figure 11).

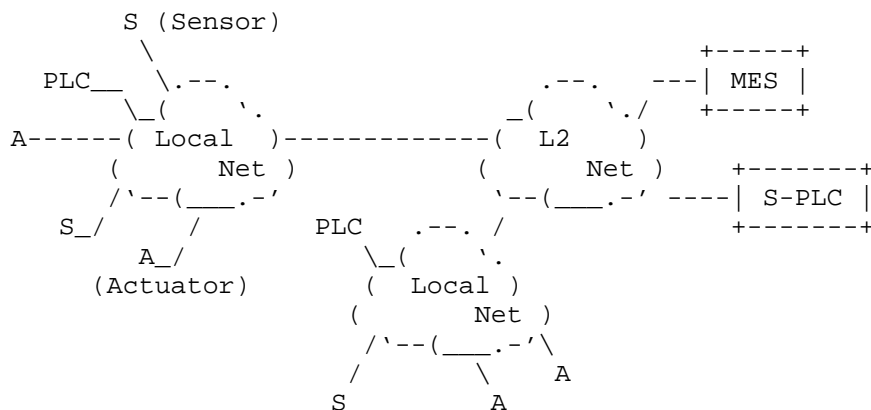


Figure 11: Current Generic Industrial M2M Network Architecture

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, setup, and database communication) is 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, deterministic networking is primarily needed 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 that 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 Communications 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 preconfigured 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 us is required. Even in the case of "non-time-coordinated" PLCs, time synchronization 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. Filtering policies that are typically enforced in firewalls are used to prevent critical flows from being leaked outside a domain.

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 1-100 ms for critical control/data streams.

Because industrial applications assume that deterministic transport will be used for critical control-data-stream parameters (instead of having to define latency and delay-variation parameters), it is sufficient to fulfill requirements regarding the upper bound of latency (maximum latency). The underlying networking infrastructure must ensure a maximum end-to-end message delivery time in the range of 100 us to 50 ms, 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 sizes in the range of 100-700 bytes. For S-PLC-to-PLC communication, 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, 1 Gbps links are commonly used.

Most PLC control loops are rather tolerant of packet loss; however, critical control/data streams accept a loss of no more than one packet per consecutive communication cycle (i.e., if a packet gets lost in cycle "n", then the next cycle ("n+1") must be lossless). After the loss of two or more consecutive packets, 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 and delivery time.

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. As production systems become more flexible going forward, there will be a significant increase in the rate at which streams are created, changed, and destroyed.

7.3. Industrial M2M in the 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; alternatively, 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 1 us is expected.

7.4. Industrial M2M Requests to the IETF

- o Converged IP-based network
- o Deterministic behavior (bounded latency and jitter)
- o High availability (presumably through redundancy) (99.999%)
- o Low message delivery time (100 us to 50 ms)

- o Low packet loss (with a bounded number of consecutive lost packets)
- o Security (e.g., preventing 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, in both open-pit and underground extraction as well as in transport and refining processes. In order to reduce risks and increase operational efficiency in mining operations, the location of operators has been relocated (as much as possible) from the extraction site to remote control and monitoring sites.

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 trucks 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 was conducted through manual inspection of preinstalled dataloggers. Cabling is not typically used in such scenarios, due to its high 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; hence, wireless technologies are employed.

In the case of underground monitoring, autonomous vehicles with extraction tools travel independently through the tunnels, but their operational tasks (such as excavation, stone-breaking, and transport) are controlled remotely from a central facility. This generates upstream video and feedback 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 comply with requirements regarding delay and packet loss, the network bandwidth is overestimated. This results in 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 connections between the backbone and the mobile devices (e.g., trucks, trains, and excavators) are implemented using a wireless link. These links are based on IEEE 802.11 [IEEE-80211] for open-pit mining and "leaky feeder" communications for underground mining. (A "leaky feeder" communication system consists of a coaxial cable, run along tunnels, that 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 Low-Power WAN (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 Layer 2. 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 triggered slowly or can be anticipated in advance.

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

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

8.3. Mining Industry in the 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 at the mine will rely on network infrastructure.

Wireless for open-pit mining is already a reality with LPWAN technologies; it is expected to evolve to more-advanced LPWAN technologies, such as those based on LTE, to increase last-hop reliability or novel LPWAN flavors 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. These networks connect many wireless Access Points (APs) together. 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 that 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 the responsiveness of digging machines to remote control. 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 Requests to the IETF

- o Improved bandwidth efficiency
- o Very low delay, to enable machine teleoperation
- o Dedicated bandwidth usage for high-resolution video streams
- o Predictable delay, to enable real-time monitoring
- o Potential for constructing 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 manage 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.

9.1.1.1. Blockchain Operation

A "block" runs as a container of a batch of primary items (e.g., transactions, property records). 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 it, 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 previous of that chain.

9.1.1.2. Blockchain Network Architecture

Blockchain node communication and coordination are achieved mainly through frequent point-to-multipoint 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' addresses from a specific entity, such as DNS. The node then creates persistent connections with each of the 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 and 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; A then sends the complete block to B. B repeats the process (as was done by A) to start the next round of block propagation.

The challenge of blockchain network operation is not overall data rates, since the volume from both the block and the item stays between hundreds of bytes and a couple of megabytes per second; rather, the challenge is in transporting the blocks with minimum latency to maximize the 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. Blockchain Security Considerations

Security is crucial to blockchain applications; 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 for a private blockchain network that is dependent on deterministic properties the network could be vulnerable to delays and other specific attacks against determinism, as these delays and attacks could interrupt service.

9.2. Private Blockchain Today

Today, private blockchain runs in Layer 2 or Layer 3 VPNs, generally 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 at present these goals are not being met.

9.3. Private Blockchain in the Future

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

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

9.4. Private Blockchain Requests to the IETF

- o Layer 2 and Layer 3 multicast of blockchain traffic
- o Item and block delivery with bounded, low latency and negligible packet loss
- o Coexistence of blockchain and IT traffic in a single network
- o 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, which corresponds 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 performance for the service at lower cost, and the flexibility of setting up and releasing the slices also allows the user to allocate 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 a user bandwidth and latency guarantees that cannot be affected by other users' data traffic. Thus, DetNet is a powerful tool when reliability and low latency are required in network slicing.

10.2.2. Deterministic Services within Slices

Slices may need to provide services with DetNet-type performance guarantees; note, however, 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 as defined in 3GPP. The system architecture for 5G is under development at the time of this writing [TS23501]. A network slice in a mobile network is a complete logical network, including RANs and Core Networks (CNs). It provides telecommunications services and network capabilities, which may vary from slice to slice. A 5G bearer network is a typical use case for network slicing; for example, consider three 5G service scenarios: eMBB, URLLC, and mMTC.

- o eMBB (Enhanced Mobile Broadband) focuses on services characterized by high data rates, such as high-definition video, Virtual Reality (VR), augmented reality, and fixed mobile convergence.
- o URLLC (Ultra-Reliable and Low Latency Communications) focuses on latency-sensitive services, such as self-driving vehicles, remote surgery, or drone control.
- o mMTC (massive Machine Type Communications) focuses on services that have high connection-density requirements, such as those typically used in smart-city and smart-agriculture scenarios.

A 5G bearer network could use DetNet to provide hard resource isolation across slices and within a given 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 resources, and latency and reliability requirements would not be guaranteed. With DetNet, URLLC-A and URLLC-B have separate bandwidth reservations; there is no resource conflict between them, as though they were in different physical networks.

10.4. Non-5G Applications of Network Slicing

Although the 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 for providing various services across a shared physical infrastructure. Examples include providing services for electrical utilities and pro audio via slices. Use cases like these could become more common once the work for the 5G CN 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 multipoint-to-multipoint 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 issue. Additional work and discussion are needed to address these topics.

10.6. Network Slicing Today and in the Future

Network slicing has promise in terms of satisfying many requirements of future network deployment scenarios, but it is still a collection of ideas and analyses without a specific technical solution. DetNet is one of various technologies that could potentially be used in network slicing, along with, for example, Flex-E and segment routing. For more information, please see the IETF 99 Network Slicing BoF session agenda and materials as provided in [IETF99-netslicing-BoF].

10.7. Network Slicing Requests to the IETF

- o Isolation from other flows through queuing management
- o Service quality customization and guarantees
- o 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 document.

11.1. Unified, Standards-Based Networks

11.1.1. Extensions to Ethernet

A DetNet network is not "a new kind of network" -- it is based on extensions to existing Ethernet standards, including elements of IEEE 802.1 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 Networks

In general, a DetNet network is not expected to be "plug and play"; rather, some type of centralized network configuration and control system is expected. Such a system may be in a single central location, or it may be 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. Layer 2 and Layer 3 Integration

A DetNet network is intended to integrate between Layer 2 (bridged) network(s) (e.g., an 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. IPv4 Considerations

This document explicitly does not specify any particular implementation or protocol; however, it has been observed that various use cases (and their associated industries) described herein are explicitly based on IPv4 (as opposed to IPv6), and it is not considered practical to expect such implementations 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) will be available.

11.1.1.6. 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 actions taken by DetNet protocol software and does not provide any guarantee against lower-level errors such as media errors or checksum errors.

11.1.1.7. 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.

11.1.1.8. Mix of Deterministic and Best-Effort Traffic

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

11.1.1.9. Unused Reserved Bandwidth 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; this occurs 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.

11.1.1.10. 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. In other words, vendors should be able to create DetNet networks 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 (e.g., a utility-grid network spanning a whole country and involving many "hops" over various kinds of links -- for example, radio repeaters, microwave links, or fiber optic links). However, recall that the scope of DetNet is confined to networks that are centrally administered and thereby 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

DetNet data-flow information models are 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 expected that the network may not be able to provide a given requested service level; if this is indeed the case, the network control system should reply that the requested services are not available (as opposed to accepting the parameter but then not delivering the desired behavior).

11.3.2. Low Latency

Various applications may state that they require "extremely low latency"; however, depending on the application, "extremely low" may imply very different latency bounds. For example, "low latency" across a utility-grid network is a different order of magnitude of latency values compared to "low latency" in a motor control loop in a small machine. It is intended that the mechanisms for specifying desired latency include wide ranges and that architecturally there is nothing to prevent arbitrarily 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 that can determine or request permitted variations in latency should be available.

11.3.4. Symmetrical Path Delays

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

11.4. High Reliability and Availability

Reliability is of critical importance to many DetNet applications, because the 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% uptime (where 99.9999 means "six nines") or even 12 nines. 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 so that a system can seamlessly switch between the paths while maintaining its required level of performance.

11.5. Security

Security is of critical importance to many DetNet applications. A DetNet network must have the ability to be made secure against device failures, attackers, misbehaving devices, and so on. In a DetNet network, the data traffic is expected to 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 PRE.

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 use cases will have security considerations from both the use-specific perspective and the DetNet technology perspective. While some use-specific security considerations are discussed above, a more comprehensive discussion of such considerations is captured in [DetNet-Security] ("Deterministic Networking (DetNet) Security Considerations"). Readers are encouraged to review [DetNet-Security] to gain a more complete understanding of DetNet-related security considerations.

13. IANA Considerations

This document has no IANA actions.

14. Informative References

- [Ahm14] Ahmed, M. and R. Kim, "Communication Network Architectures for Smart-Wind Power Farms", *Energies* 2014, pp. 3900-3921, DOI 10.3390/en7063900, June 2014.
- [Arch-for-6TiSCH]
Thubert, P., Ed., "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4", *Work in Progress*, draft-ietf-6tisch-architecture-20, March 2019.
- [BACnet-IP]
ASHRAE, "Annex J to ANSI/ASHRAE 135-1995 - BACnet/IP", January 1999, <<http://www.bacnet.org/Addenda/Add-1995-135a.pdf>>.
- [BAS-DetNet]
Kaneko, Y. and S. Das, "Building Automation Use Cases and Requirements for Deterministic Networking", *Work in Progress*, draft-bas-usecase-detnet-00, October 2015.
- [CoAP-6TiSCH]
Sudhaakar, R., Ed. and P. Zand, "6TiSCH Resource Management and Interaction using CoAP", *Work in Progress*, draft-ietf-6tisch-coap-03, March 2015.
- [CoMP]
NGMN Alliance, "RAN EVOLUTION PROJECT COMP EVALUATION AND ENHANCEMENT", VERSION 2.0, NGMN Alliance, March 2015, <https://www.ngmn.org/fileadmin/user_upload/NGMN_RANEV_D3_CoMP_Evaluation_and_Enhancement_v2.0.pdf>.

- [Content_Protection]
Olsen, D., "1722a Content Protection", April 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 1.3", June 2018, <<https://www.dcinovies.com/>>.
- [Det-Fwd-PHB]
Shah, S. and P. Thubert, "Deterministic Forwarding PHB",
Work in Progress,
draft-svshah-tsvwg-deterministic-forwarding-04,
August 2015.
- [DetNet-6TiSCH]
Thubert, P., Ed., "6TiSCH requirements for DetNet", Work
in Progress, draft-thubert-6tisch-4detnet-01, June 2015.
- [DetNet-Arch]
Finn, N., Thubert, P., Varga, B., and J. Farkas,
"Deterministic Networking Architecture", Work in Progress,
draft-ietf-detnet-architecture-13, May 2019.
- [DetNet-Audio-Reqs]
Gunther, C., Ed. and E. Grossman, Ed., "Deterministic
Networking Professional Audio Requirements", Work in
Progress, draft-gunther-detnet-proaudio-req-01,
March 2015.
- [DetNet-Mobile]
Zha, Y., "Deterministic Networking Use Case in Mobile
Network", Work in Progress, draft-zha-detnet-use-case-00,
July 2015.
- [DetNet-RAN]
Korhonen, J., "Deterministic networking for radio
access networks", Work in Progress,
draft-korhonen-detnet-telreq-00, May 2015.

[DetNet-Security]

Mizrahi, T., Grossman, E., Ed., Hacker, A., Das, S., Dowdell, J., Austad, H., Stanton, K., and N. Finn, "Deterministic Networking (DetNet) Security Considerations", Work in Progress, draft-ietf-detnet-security-04, March 2019.

[DetNet-Util-Reqs]

Wetterwald, P. and J. Raymond, "Deterministic Networking Utilities requirements", Work in Progress, draft-wetterwald-detnet-utilities-reqs-02, June 2015.

[eCPRI]

IEEE Standards Association, "Common Public Radio Interface: eCPRI Interface Specification V1.2", June 2018, <<http://www.cpri.info/>>.

[ESPN_DC2]

Daley, D., "ESPN's DC2 Scales AVB Large", SVG News, June 2014, <<https://sportsvideo.org/main/blog/2014/06/espns-dc2-scales-avb-large>>.

[EtherCAT]

"EtherCAT Technology Group", <<https://www.ethercat.org/default.htm>>.

[FL-net]

Japan Electrical Manufacturers Association, "JEMA 1479 - English Edition", September 2012, <[https://www.jema-net.or.jp/Japanese/standard/opcn/pdf/JEM_1479e\(20120927\).pdf](https://www.jema-net.or.jp/Japanese/standard/opcn/pdf/JEM_1479e(20120927).pdf)>.

[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_chen_1.pdf>.

[IEC-60834]

International Electrotechnical Commission, "Teleprotection equipment of power systems - Performance and testing", IEC 60834, October 1999.

[IEC-60870-5-104]

International Electrotechnical Commission, "Telecontrol equipment and systems - Part 5-104: Transmission protocols - Network access for IEC 60870-5-101 using standard transport profiles", IEC 60870-5-104, June 2006.

[IEC-61400-25]

International Electrotechnical Commission, "Communications for monitoring and control of wind power plants", IEC 61400-25, June 2013.

[IEC-61850-5:2013]

International Electrotechnical Commission, "Communication networks and systems for power utility automation - Part 5: Communication requirements for functions and device models", IEC 61850-5, January 2013.

[IEC-61850-9-2:2011]

International Electrotechnical Commission, "Communication networks and systems for power utility automation - Part 9-2: Specific communication service mapping (SCSM) - Sampled values over ISO/IEC 8802-3", IEC 61850-9-2, September 2011.

[IEC-61850-90-12:2015]

International Electrotechnical Commission, "Communication networks and systems for power utility automation - Part 90-12: Wide area network engineering guidelines", IEC TR 61850-90-12, July 2015.

[IEC-62357-200:2015]

International Electrotechnical Commission, "Power systems management and associated information exchange - Part 200: Guidelines for migration from Internet Protocol version 4 (IPv4) to Internet Protocol version 6 (IPv6)", IEC 62357-200:2015, July 2015.

[IEC-62439-3:2016]

International Electrotechnical Commission, "Industrial communication networks - High availability automation networks - Part 3: Parallel Redundancy Protocol (PRP) and High-availability Seamless Redundancy (HSR)", March 2016.

[IEC-IEEE-61850-9-3:2016]

International Electrotechnical Commission, "Communication networks and systems for power utility automation - Part 9-3: Precision time protocol profile for power utility automation", IEC 61850-9-3, May 2016.

[IEEE-1588]

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

- [IEEE-1646]
IEEE, "IEEE Standard Communication Delivery Time Performance Requirements for Electric Power Substation Automation", IEEE Standard 1646,
<<https://standards.ieee.org/standard/1646-2004.html>>.
- [IEEE-1722]
IEEE, "IEEE Standard for a Transport Protocol for Time-Sensitive Applications in Bridged Local Area Networks", IEEE Standard 1722,
<<https://standards.ieee.org/findstds/standard/1722-2016.html>>.
- [IEEE-1815]
IEEE Standards Association, "IEEE Standard for Electric Power Systems Communications-Distributed Network Protocol (DNP3)", IEEE Standard 1815, <<https://ieeexplore.ieee.org/servlet/opac?punumber=6327576>>.
- [IEEE-19143]
IEEE Standards Association, "IEEE Standard for Radio over Ethernet Encapsulations and Mappings", IEEE 1914.3,
<<https://standards.ieee.org/develop/project/1914.3.html>>.
- [IEEE-80211]
IEEE Standard for Information technology, "IEEE Std. 802.11, Telecommunications and information exchange between systems--Local and metropolitan area networks--Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications",
<https://standards.ieee.org/standard/802_11-2016.html>.
- [IEEE-802154]
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 (WPANs)",
<https://standards.ieee.org/standard/802_15_4-2015.html>.
- [IEEE-8021AS]
IEEE, "IEEE Standard for Local and Metropolitan Area Networks - Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks", IEEE 802.1AS,
<<http://www.ieee802.org/1/pages/802.1as.html>>.

- [IEEE-8021CM]
"IEEE Standard for Local and metropolitan area networks - Time-Sensitive Networking for Fronthaul", IEEE Standard 802.1CM,
<https://standards.ieee.org/standard/802_1CM-2018.html>.
- [IEEE-8021TSNTG]
IEEE Standards Association, "IEEE 802.1 Time-Sensitive Networking Task Group",
<<http://www.ieee802.org/1/pages/avbridges.html>>.
- [IETF99-netslicing-BoF]
"Network Slicing (netslicing) BoF", IETF 99, Prague, July 2017, <<https://datatracker.ietf.org/meeting/99/materials/slides-99-netslicing-chairs-netslicing-bof-04>>.
- [Interface-6TiSCH-6top]
Wang, Q., Ed. and X. Vilajosana, "6TiSCH Operation Sublayer (6top) Interface", Work in Progress, draft-ietf-6tisch-6top-interface-04, July 2015.
- [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 Corp., "LonTalk(R) Protocol Specification Version 3.0", 1994, <<http://www.enerlon.com/JobAids/Lontalk%20Protocol%20Spec.pdf>>.
- [MailingList-6TiSCH]
IETF, "6TiSCH Mailing List",
<<https://mailarchive.ietf.org/arch/browse/6tisch/>>.
- [MEF22.1.1]
Metro Ethernet Forum, "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]
Metro Ethernet Forum, "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", Document Number ICT-317669-METIS/D1.1, April 2013, <https://metis2020.com/wp-content/uploads/deliverables/METIS_D1.1_v1.pdf>.
- [MODBUS] Modbus Organization, Inc., "MODBUS Application Protocol Specification", April 2012, <<http://www.modbus.org/specs.php>>.
- [NGMN] NGMN Alliance, "5G White Paper", NGMN 5G White Paper v1.0, February 2015, <https://www.ngmn.org/fileadmin/ngmn/content/downloads/Technical/2015/NGMN_5G_White_Paper_V1_0.pdf>.
- [NGMN-Fronth] NGMN Alliance, "Fronthaul Requirements for C-RAN", March 2015, <https://www.ngmn.org/fileadmin/user_upload/NGMN_RANEV_D1_C-RAN_Fronthaul_Requirements_v1.0.pdf>.
- [OPCXML] OPC Foundation, "OPC Data Access (OPC DA) Specification", <<http://www.opcti.com/opc-da-specification.aspx>>.
- [PCE] IETF, "Path Computation Element", <<https://datatracker.ietf.org/doc/charter-ietf-pce/>>.
- [PROFIBUS] IEC, "PROFIBUS Standard - DP Specification (IEC 61158 Type 3)", <<https://www.profibus.com/>>.
- [PROFINET] "PROFINET Technology", <<https://us.profinet.com/technology/profinet/>>.
- [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] Barbosa, R., Sadre, R., and A. Pras, "A First Look into SCADA Network Traffic", IP Network Operations and Management Symposium, DOI 10.1109/NOMS.2012.6211945, June 2012, <<https://ieeexplore.ieee.org/document/6211945>>.
- [SR-IP-RAN-Use-Case] Khasnabish, B., Hu, F., and L. Contreras, "Segment Routing in IP RAN use case", Work in Progress, draft-kh-spring-ip-ran-use-case-02, November 2014.
- [SRP_LATENCY] Gunther, C., "Specifying SRP Acceptable Latency", March 2014, <<http://www.ieee802.org/1/files/public/docs2014/cc-cgunther-acceptable-latency-0314-v01.pdf>>.
- [Sublayer-6TiSCH-6top] Wang, Q., Ed. and X. Vilajosana, "6TiSCH Operation Sublayer (6top)", Work in Progress, draft-wang-6tisch-6top-sublayer-04, November 2015.
- [syncE] International Telecommunication Union, "Timing and synchronization aspects in packet networks", ITU-T Recommendation G.8261, August 2013, <<https://www.itu.int/rec/T-REC-G.8261>>.
- [Timing-over-MPLS] Davari, S., Oren, A., Bhatia, M., Roberts, P., and L. Montini, "Transporting Timing messages over MPLS Networks", Work in Progress, draft-ietf-tictoc-1588overmpls-07, October 2015.
- [TR38801] 3GPP, "Study on new radio access technology: Radio access architecture and interfaces (Release 14)", 3GPP TR 38.801, April 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 (Release 16)", 3GPP TS 23.401, March 2019, <<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=849>>.
- [TS23501] 3GPP, "System architecture for the 5G System (5GS) (Release 15)", 3GPP TS 23.501, March 2019, <<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3144>>.

- [TS25104] 3GPP, "Base Station (BS) radio transmission and reception (FDD) (Release 16)", 3GPP TS 25.104, January 2019, <<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=1154>>.
- [TS36104] 3GPP, "Evolved Universal Terrestrial Radio Access (E-UTRA); Base Station (BS) radio transmission and reception (Release 16)", 3GPP TS 36.104, January 2019, <<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=2412>>.
- [TS36133] 3GPP, "Evolved Universal Terrestrial Radio Access (E-UTRA); Requirements for support of radio resource management (Release 16)", 3GPP TS 36.133, January 2019, <<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=2420>>.
- [TS36211] 3GPP, "Evolved Universal Terrestrial Radio Access (E-UTRA); Physical channels and modulation (Release 15)", 3GPP TS 36.211, January 2019, <<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=2425>>.
- [TS36300] 3GPP, "Evolved Universal Terrestrial Radio Access (E-UTRA) and Evolved Universal Terrestrial Radio Access Network (E-UTRAN); Overall description; Stage 2 (Release 15)", 3GPP TS 36.300, January 2019, <<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=2430>>.
- [WirelessHART]
International Electrotechnical Commission, "Industrial networks - Wireless communication network and communication profiles - WirelessHART(TM)", IEC 62591:2016, March 2016.

Appendix A. Use Cases Explicitly Out of Scope for DetNet

This appendix contains text regarding use cases that have been determined to be outside 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 that the DetNet use cases were solicited and provided by the authors, the scope of DetNet was not clearly defined. As the scope has been clarified, certain use cases have been determined to be outside the scope of the present DetNet work. Text regarding these use cases was moved to this appendix to clarify that they will not be supported by the DetNet work.

The text was moved to this appendix based on the following "exclusion" principles. Please note that as an alternative to moving all such text to this appendix some 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.

- o The scope of networks 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".
- o Maintaining time synchronization across a DetNet network is crucial to its operation; however, DetNet assumes that time is to be maintained using other means. One example would be PTP [IEEE-1588]. 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, two open questions remain: (1) whether DetNet protocols will include a way for an application to communicate expectations regarding such timing properties to the network and (2) if so, whether those properties would likely have a material effect on network performance 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's experience. For example, "First Person Shooter" (FPS) games are highly delay sensitive.

A.2.1.3. Virtual Reality

VR 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 regarding delivery or bandwidth.

A.2.3. Internet-Based Applications in the Future

One should be able to play Internet videos without glitches and play Internet games without lag.

For online gaming, the desired maximum allowance for round-trip delay is typically 100 ms. However, it may be less for specific types of games; for example, for FPS games, the maximum delay should be 50 ms. Transport delay is the dominant part, with a budget of 5-20 ms.

For VR, a maximum delay of 1-10 ms is needed; if doing remote VR, the total network delay budget is 1-5 ms.

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 Requests to the IETF

- o Unified control and management protocols that handle time-critical data flows
- o An application-aware flow-filtering mechanism that recognizes time-critical flows without doing 5-tuple matching
- o A unified control plane that provides low-latency service on Layer 3 without changing the data plane
- o An OAM system and protocols that can help provide service provisioning that is sensitive to end-to-end delays

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

The following text was moved to this appendix because this information is considered a link-layer topic for which DetNet is not directly responsible.

Digital Rights Management (DRM) is very important to the audio and video industries. Whenever protected content is introduced into a network, there are DRM concerns that must be taken into account (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 such 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 content may be retransmitted under HDCP. Therefore, if the source of a stream is outside of the network and it uses HDCP, it is only allowed to be placed on the network with that same type of protection (i.e., HDCP).

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 that provides

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 WG decided that guidelines for establishing a deterministic time to establish stream startup are not within the scope of DetNet. If the bounded timing for establishing or re-establishing streams is required in a given use case, it is up to the application/system to achieve it.

Acknowledgments

Pro audio (Section 2)

As also acknowledged in [DetNet-Audio-Reqs], the editor 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

Utility telecom (Section 3)

Information regarding utility telecom was derived from [DetNet-Util-Reqs]. As in that document, the following individuals are acknowledged here.

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 parks 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).

Building automation systems (Section 4)

Please see [BAS-DetNet].

Wireless for industrial applications (Section 5)

See [DetNet-6TiSCH].

[DetNet-6TiSCH] 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 [MailingList-6TiSCH].

As also acknowledged in [DetNet-6TiSCH], the editor wishes 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.

Cellular radio (Section 6)

See [DetNet-RAN].

Internet applications and CoMP (Section 6)

As also acknowledged in [DetNet-Mobile], authored by Yiyong Zha, the editor would like to thank the following people for their reviews, suggestions, comments, and proposed text: Jing Huang, Junru Lin, Lehong Niu, and Oliver Huang.

Industrial Machine to Machine (M2M) (Section 7)

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

Mining industry (Section 8)

This text was written by Diego Dujovne, who worked in conjunction with Xavier Vilajosana.

Private blockchain (Section 9)

This text was written by Daniel Huang.

Network slicing (Section 10)

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

Contributors

RFC 7322 ("RFC Style Guide") generally limits the number of authors listed on the front page of a document to five individuals -- far fewer than the 19 individuals listed below, who also made important contributions to this document. The editor wishes to thank and acknowledge each of the following authors for contributing text to this document. See also the Acknowledgments section.

Craig Gunther (Harman International)
10653 South River Front Parkway
South Jordan, UT 84095
United States of America
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 4 97 23 26 34
Email: pthubert@cisco.com

Patrick Wetterwald (Cisco Systems)
45 Allee des Ormes
Mougins 06250
France
Phone: +33 4 97 23 26 36
Email: pwetterw@cisco.com

Jean Raymond (Hydro-Quebec)
1500 University
Montreal, Quebec H3A 3S7
Canada
Phone: +1 514 840 3000
Email: raymond.jean@hydro.qc.ca

Jouni Korhonen (Broadcom Corporation)
3151 Zanker Road
San Jose, CA 95134
United States of America
Email: jouni.nospam@gmail.com

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

Subir Das (Vencore Labs)
150 Mount Airy Road
Basking Ridge, NJ 07920
United States of America
Email: sdas@appcomsci.com

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

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

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

Juergen Schmitt (Siemens)
Gleiwitzerstr. 555
Nurnberg 90475
Germany
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 Vizarrreta (Technical University of Munich)
Maxvorstadt, Arcisstrasse 21
Munich 80333
Germany
Email: petra.stojsavljevic@tum.de

Daniel Huang (ZTE Corporation, Inc.)
No. 50 Software Avenue
Nanjing, Jiangsu 210012
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

Author's Address

Ethan Grossman (editor)
Dolby Laboratories, Inc.
1275 Market Street
San Francisco, CA 94103
United States of America

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

