



Vendor Neutrality Drivers and Hindrances-Optical Spectrum as a Service in Disaggregated and Open Networks

Downloaded from: <https://research.chalmers.se>, 2026-02-28 06:29 UTC

Citation for the original published paper (version of record):

Kaeval, K., Kerm, H., Järvelill, T. et al (2025). Vendor Neutrality Drivers and Hindrances-Optical Spectrum as a Service in Disaggregated and Open Networks. 2025 European Conference on Optical Communications (ECOC). <http://dx.doi.org/10.1109/ECOC66593.2025.11263188>

N.B. When citing this work, cite the original published paper.

Vendor Neutrality Drivers and Hindrances - Optical Spectrum as a Service in Disaggregated and Open Networks

Kaida Kaeval⁽¹⁾, Hendrik Johann Kerm⁽¹⁾, Torm Järvelill⁽¹⁾, Carlos Natalino⁽²⁾, Jasper Müller⁽³⁾

⁽¹⁾ Tallinn University of Technology, Tallinn, Estonia, kaida.kaeval@taltech.ee

⁽²⁾ Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden

⁽³⁾ Adtran, Martinsried, Germany

Abstract Network disaggregation over the past decade has led to the emergence of Optical Spectrum as a Service (OSaaS). This paper discusses how Telemetry as a Service (TaaS) can help OSaaS operation in disaggregated optical networks and discusses the data communication network requirements for TaaS platform. ©2025 The Author(s)

Introduction

The idea of providing Optical Spectrum as a Service (OSaaS) to monetize spectral resources in lightly populated Open Line Systems (OLS) has been around for over a decade [1]. OSaaS provides end-user a continuous spectral region in the Open Line System (OLS), compared to access to a single wavelength provided by alien wavelength services. If the same wavelength region is available in multiple adjacent network domains, the signals can travel thousands of kilometres, enabling new multi-domain network architectures, as described in [2]. This unlocks several benefits, such as avoiding expensive transceiver investments which are otherwise required for regenerations at domain borders, and implementation of end-to-end low-margin networking or signal pre-emphasis.

However, creation, management, and maintenance of such end-to-end services is complicated, due to often limited access to underlying OLS domains. While the inter-domain OSaaS signal parameters can be controlled locally at domain borders, using a domain interface device (NeDID) [3] or a software defined networking solution [4], large portions of third-party light can introduce power redistribution, anomalies and failures in the adjacent OLS domains. To detect and compensate for anomalies and evolving failures of the underlying OLSs, telemetry data can be collected. However, in multi-operator environments, such data is often considered business-critical. This delays the implementation of OSaaS in live networks, as the raw telemetry data collected from the OSaaS end-user transceivers is often not sufficient for fast anomaly detection or end-to-end performance optimizations [5]. To overcome this, a new framework comprising data-sovereign features for secure telemetry sharing and machine-learning (ML) empowerment in the networks has been proposed [6].

In essence, the services of a telemetry

sharing platform can be called Telemetry as a Service (TaaS). To distinguish between different service levels compared to a single-operator telemetry collection, the responsibility of the telemetry data availability alongside with various value-added data streams, such as labels, conceptualized datasets, or knowledge graphs complementing raw data can be defined under service-specific Quality-of-Service (QoS) agreements, as described in Fig. 1.

In order to assure telemetry data representativeness and timeliness for ML-based anomaly detection and performance adjustments in the OSaaS use-case, the collection rate of the raw telemetry data from the OLSs should be optimized. For example, assuring 99.999% parameter validation or service availability calculation per day requires telemetry collection with below-second interval [7]. The number of Network Elements (NEs), internal modules in the NE, number and rate of parameters collected, determine the capacity and reliability demands on the data link between the NEs and telemetry agent. In conventional telemetry collection from domains under a single operator management, the Data Communication Network (DCN) channels between the NEs and telemetry agent are often implemented using the capacity of an internal Optical Supervisory Channel (OSC) of the OLS and their performance can be controlled and adjusted regularly. However, various network topologies and TaaS use cases may

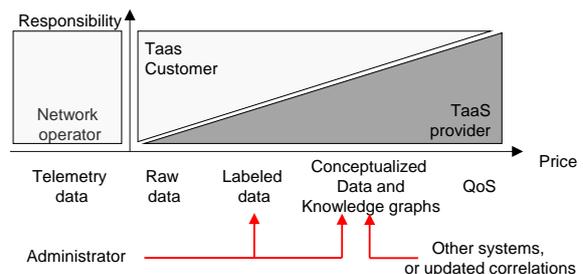


Fig. 1: Responsibility boundaries and value-added data streams in the Telemetry as a Service concept.

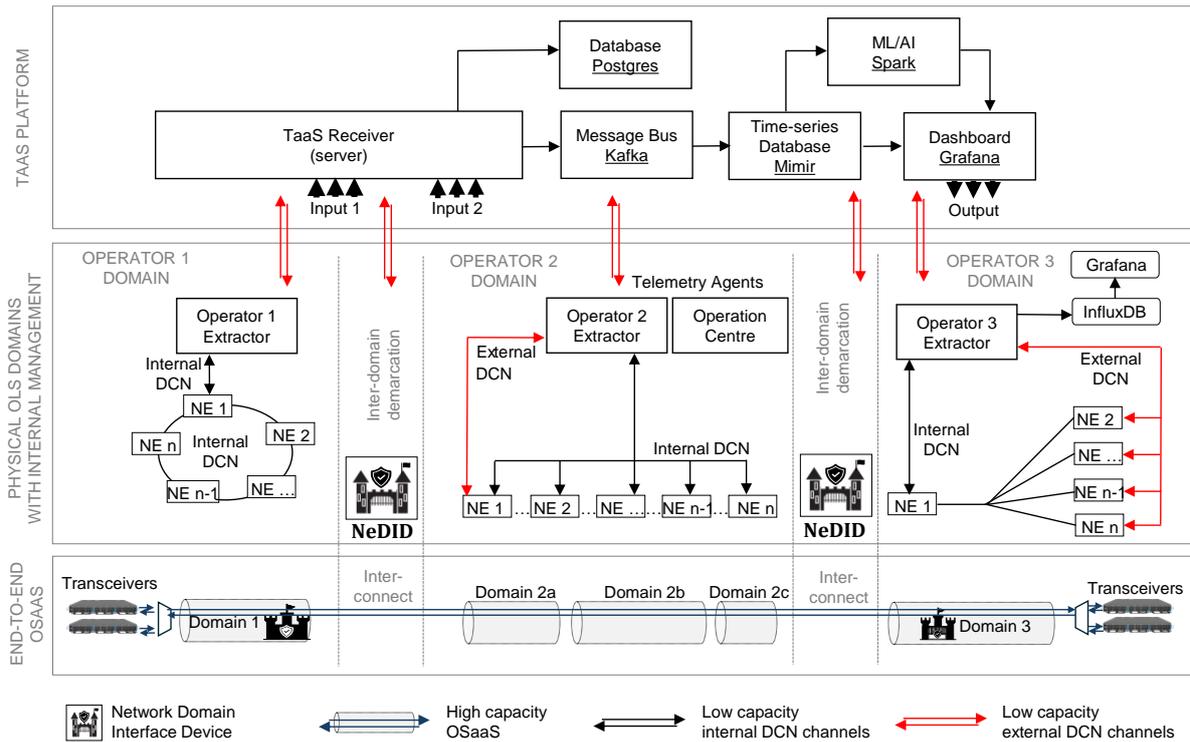


Fig. 2 Telemetry as a Service platform components with underlying network domains carrying Optical Spectrum as a Service

require an external data channel between the NEs, domain telemetry agent (extractor) or TaaS receiver, which often lack monitoring and control functions. These channels, illustrated with a red line in Fig. 2, are independent from the underlying physical OLSs, as they may be implemented using VPNs, dedicated Ethernet, or even dark fibre infrastructure – whichever is available or more cost efficient.

Fig. 2 explains the 3 layers for the end-to-end OSaaS, traversing three adjacent physical OLS domains, interconnected with NeDID devices. The bottom layer presents the end-to-end OSaaS service layer. For OSaaS, the telemetry data can be usually retrieved from the transceivers, while the access to OLS telemetry is restricted or absent. The second layer portrays the physical OLS layer with NEs, each composed of multiple modules, providing numerous parameters as a telemetry output. Each parameter can be sent in a separate telemetry message, or packed into larger messages, including multiple parameters per module, or per NE. Further considering the data representativeness requirements for various network anomalies and faults, the increased collection and upload rate can create a burden on DCN channels, and possible delay and packet loss may impact the QoS of the TaaS. While this can be easily addressed for internal DCN channels in single-operator networks, understanding the impact would help to adjust the required parameters for the external DCN channels. The last layer on Fig. 2 presents the

TaaS components. The telemetry (input 1) and possible sensing data (input 2) from the operators and individual devices from the interconnected domains are collected in the TaaS receiver, which stores the data into a time-series database. For this, various open-source database implementations, like Mimir, can be used. Then, the data can be directly presented on dashboards, using Grafana, or first ran through the ML agent, and then displayed. This work analyses the performance of the telemetry data transfer between operator-controlled extractors and TaaS receiver, focusing on throughput of telemetry messages under external DCN imperfections such as delay, packet loss and limited channel bandwidth [8].

Experimental setup

In order to simulate TaaS systems, the system was realized as a full-stack pipeline complete with extractor (client), receiver (server), and time-series database backend. The extractor generates synthetic telemetry messages from real-world data templates and sends them to the receiver over a direct LAN connection, providing 810 Mbits/s capacity [8]. The receiver processes incoming data, transforms it into Mimir-compatible format, and forwards it to the time-series database. The goal of this setup is to measure the performance of gRPC telemetry streams under controlled and impaired network conditions such as transmission latency and packet loss, simulated by the Clumsy tool [9].

Each transmitted telemetry message contained at least one telemetry parameter, consisting of metadata describing the NE and the module, and the measured value. With each test, 1,000 telemetry messages were transmitted between extractor and receiver, reporting critical performance metrics such as message throughput (messages/sec) and total transmission time.

Results

To evaluate how bundling multiple telemetry parameters into a single gRPC telemetry message impacts the streaming performance, first, the number of parameters was varied per message over constant channel conditions. The test scenarios ranged from messages containing just 1 parameter up to 1250 parameters. For each configuration, the throughput was measured in messages per second. Fig. 3, left, illustrates the findings with system's highest throughput at 168.58 messages/s when transmitting messages with one parameter only, and 13.04 messages/s when 1250 parameters per message were sent. We account this to additional computational and serialization overhead for larger messages.

To examine the robustness against unforeseen network disruptions, the extractor was fixed to send 50 and 200 parameters per telemetry message, and transmission delay and packet loss were simulated using Clumsy. Fig. 3, middle, demonstrates the average throughput in messages/s for transmission delay ranging from 1 to 1500 ms, plotted for 50 and 200 parameters per message. While larger telemetry messages are more impacted by the transmission delay in general, the relative difference equals out with growing transmission delay.

To evaluate the effects of packet loss, loss rates of 1%, 5%, 10%, and 15% were introduced by Clumsy. Similarly to the delay use case, the extractor was fixed to send 50 and 200 parameters per message. Fig. 3, right, shows throughput degradation with increased loss rate. Furthermore, the impact was worse for 200 parameter messages, differing up to 4 times from the residual throughput of smaller packets.

Conclusions

Intense efforts towards network disaggregation and openness over the last decade have allowed the emergence of new service models like Optical Spectrum as a Service. The benefits from such service models are specifically interesting in multi-operator environments. However, this service model also introduces new requirements for telemetry collection due to its fundamental role in the end-to-end performance estimations and machine-learning-based anomaly detection in open, disaggregated network environments. In addition to secure and data sovereign raw telemetry sharing, as proposed in [6], operators expect the telemetry data upload rates to be representative of possible network anomalies [7]. When availability of the raw data, labelling, context and knowledge graphs are guaranteed with certain quality of service, Telemetry as a Service can be provided. However, the new requirements on TaaS data impose a stress on the communication links transmitting telemetry and sensing data between data extractors and receiver at TaaS platform. This stress is specifically impactful, when operating under limited capacity, delay, or packet loss conditions.

This paper has discussed the overall architecture of the open and disaggregated network resources and their connectivity to the TaaS system. We presented experimental results for data channels carrying raw telemetry data between the extractors and the TaaS platform. We have shown the impact of grouping numerous parameters in one telemetry message and identified the impact of delay and packet loss of the DCN channel as the most impactful impairments in TaaS system operation. Future studies on this topic involve channel capacity requirements for co-hosting sensing data transmission alongside with telemetry streams.

Acknowledgements

The work has been funded by the European Commission, grant agreements 10113933 (ECO-eNET) and 101189703 (ICON).

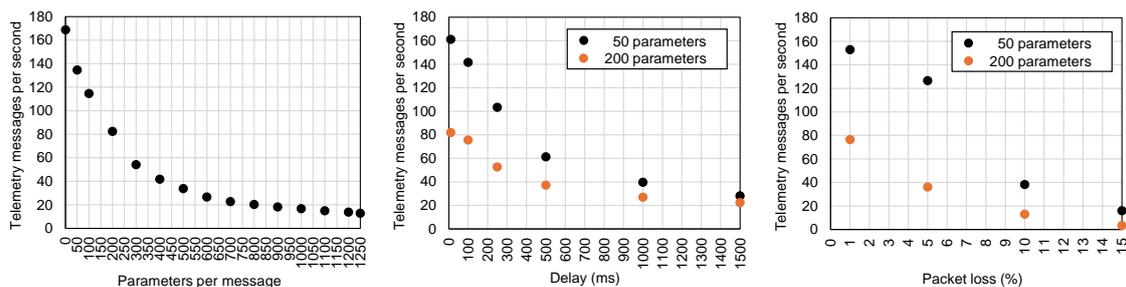


Fig. 3: Telemetry message throughput depending on Left: included parameters, Middle: transmission delay, Right: packet loss

References

- [1] K. Grobe, C. Fürst, A. Autenrieth, and T. Szyrkowicz, "Flexible spectrum-as-a-service," In TERENA Networking Conference (TNC), (2014).
- [2] IOWN Global Forum, "Functional architecture for Multi-domain IOWN Networking," March 2025
- [3] K. Kaeval, K. Grobe and J. -P. Elbers, "Operation of optical spectrum as a service in disaggregated and multi-operator environments [Invited]," in Journal of Optical Communications and Networking, vol. 17, no. 1, pp. A46-A58, January 2025, doi: 10.1364/JOCN.534118.
- [4] Y. Li, W. Mo, S. Zhu, Y. Shen, J. Yu, P. Samadi, K. Bergman, and D.C. Kilper, 2017. "tSDX: enabling impairment-aware all-optical inter-domain exchange," Journal of Lightwave Technology, 36(1), pp.142-154, (2020),
- [5] S. K. Patri et al., "Machine Learning enabled Fault-Detection Algorithms for Optical Spectrum-as-a-Service Users," International Conference on Optical Network Design and Modeling (ONDM), Coimbra, Portugal, 2023, pp. 1-6, doi: 10.23919/ONDM57372.2023.10144858.
- [6] A. Mitrovska, B. Shariati, P. Safari and J. K. Fischer, "Data Governance Framework for Telemetry Sharing," European Conference on Optical Communication (ECOC), Frankfurt, Germany, 2024, pp. 1311-1314.
- [7] K. Kaeval, "End-to-End Link Availability Assessments in an Optical Spectrum as a Service Use-Case," International Conference on Optical Network Design and Modeling (ONDM), Madrid, Spain, 2024, pp. 1-6, doi: 10.23919/ONDM61578.2024.10582649.
- [8] H.J. Kerm, "Development of Telemetry-as-a-Service Platform for Optical Communication Networks and Sensing Grids", 2025, Taltech press.
- [9] Clumsy 0.3 homepage, last accessed 15th of May, 2025, <https://jagt.github.io/clumsy>