

Multi-Agent Digital Twin for SLA Management in Coherent P2MP Metro-Access Networks

Original

Multi-Agent Digital Twin for SLA Management in Coherent P2MP Metro-Access Networks / Zeb, S., Dipto, I.C., Masood, M.U., Costa, N., Pedro, J., Napoli, A., Curri, V.. - (2026). (17th International Conference on Network of the Future (NoF 2026) Roma (Ita) 29 September - 02 October 2026).

Availability:

This version is available at: 11583/3015758 since: 2026-09-21T10:46:39Z

Publisher:

IEEE

Published

DOI:

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Multi-Agent Digital Twin for SLA Management in Coherent P2MP Metro-Access Networks

Sanwal Zeb*, Imran Chowdhury Dipto*, Muhammad Umar Masood*, Nelson Costa[†],
João Pedro[‡], Antonio Napoli[‡], Vittorio Curri*

*DET, Politecnico di Torino, Torino, Italy

[†]Nokia Optical Networks, Carnaxide, Portugal [‡]Nokia Optical Networks, Munich, Germany

Email: sanwal.zeb@polito.it

Abstract—Dynamic multi-service traffic in coherent point-to-multipoint (P2MP) metro-access networks makes proactive service level agreement (SLA) management increasingly challenging. Existing AI-enabled network management solutions primarily support reactive monitoring and provide limited capability for anticipating future SLA degradation. This demonstration presents an AI-native digital twin for predictive SLA management that integrates traffic forecasting, SLA-state prediction, recovery recommendation, and a large language model (LLM)-assisted conversational interface into a unified decision-making framework. By predicting future enterprise, radio access network (RAN), and passive optical network (PON) traffic demands, the framework identifies potential SLA violations, estimates network risk, and recommends proactive mitigation actions before service degradation occurs. Through natural-language interactions, operators can query future network conditions, understand the causes of predicted SLA violations, and receive explainable operational recommendations. The demonstration highlights the potential of AI-native digital twins to enable proactive, intent-driven, and explainable management of next-generation coherent P2MP metro-access networks.

Index Terms—Multi-Agent, Digital Twin, Point-to-multipoint (P2MP), Metro-Access, Service-Level-Agreement (SLA)

I. INTRODUCTION

The rapid growth of mobile, enterprise, and residential connectivity services is transforming the operational landscape of metro-access networks. These heterogeneous services exhibit diverse traffic characteristics while sharing a common backend transport network. Ensuring service-level agreement (SLA) compliance across such multi-service environments has therefore become a critical challenge for network operators. Coherent point-to-multipoint (P2MP) transmission based on digital subcarrier multiplexing (DSCM) has emerged as a promising solution for supporting flexible bandwidth allocation and scalable resource sharing in converged metro-access infrastructures [1, 2]. However, the increasing traffic variability and service diversity introduce additional operational complexity, motivating the need for intelligent and predictive network management mechanisms.

Network digital twins have emerged as an effective approach for monitoring and analyzing complex communication infrastructures [3, 4]. They provide virtual representations of physical networks and support data-driven operational insights. In parallel, machine learning techniques have been widely adopted for traffic forecasting, performance optimization, and quality-of-transmission (QoT) estimation in optical

networks [5]. Recent advances in artificial intelligence and agentic systems further allow intelligent agents to reason, collaborate, and assist operators through natural-language interactions [6]. Together, these developments create new opportunities to transform digital twins from monitoring platforms into intelligent systems for proactive network management.

Despite these advances, most existing solutions remain largely reactive and focus on observing current network conditions rather than anticipating future service degradations. They provide limited support for predicting SLA violations and recommending mitigation actions before network performance is affected. This challenge is particularly significant in metro-access networks, where multiple services share common optical resources and exhibit highly dynamic traffic patterns. Consequently, there is a growing need for AI-native systems that can forecast future network states, assess SLA risks, and provide explainable assistance for proactive network management.

In this demonstration, we present an AI-native digital twin for predictive SLA management in coherent P2MP metro-access networks. The proposed framework integrates traffic forecasting, SLA-state prediction, recovery recommendation, and LLM-assisted operator interaction into a unified decision-making pipeline. An XGBoost model predicts future RAN, enterprise, and PON traffic demands, while a Random Forest classifier identifies normal, degraded, and failure-prone network states. Based on these predictions, the framework generates proactive mitigation recommendations and translates them into explainable natural-language responses through an LLM-assisted conversational interface. The demonstration illustrates how predictive analytics and AI can jointly enable proactive, and intent-driven SLA management.

II. AI-NATIVE DIGITAL TWIN ARCHITECTURE

We consider a coherent DSCM-based point-to-multipoint (P2MP) metro-access network composed of application, control, and data planes, as illustrated in Fig. 1. A centralized network controller configures the optical resources through northbound and southbound interfaces. In the data plane, a single 400 G coherent hub transceiver serves $N = 4$ access (leaf) nodes through a passive splitter/combiner in a filterless architecture [7]. The hub supports $K = 16$ subcarriers (SCs), each providing 25 Gb/s using 16-QAM modulation, while each leaf node is equipped with a 100 G coherent transceiver using




Fig. 1: Multi-agent framework for SLA management.

$K_s = 4$ SCs [5]. Multi-service traffic comprising RAN, enterprise, and PON is aggregated at each access node and transported over the shared feeder link. The centralized controller dynamically configures SC assignments according to the selected resource allocation policy while ensuring efficient utilization of the available optical capacity.

The proposed AI-native digital twin extends the network digital twin with interconnected AI agents that collectively support predictive SLA management. As illustrated in Fig. 1, the framework processes multi-service traffic through a sequential decision-making pipeline, where each agent consumes the output of the previous stage to progressively transform raw traffic data into actionable operational intelligence.

- **Traffic Forecaster Agent:** It predicts future Enterprise, RAN, and PON traffic using an XGBoost regression trained on historical traffic data. Then the agent calculates the aggregate traffic and reconstructs temporal features needed by the classifier agent. The predicted traffic profiles and reconstructed temporal features are then forwarded to the SLA Prediction Agent for network-state evaluation.
- **SLA Prediction Agent:** It evaluates the predicted traffic using a Random Forest classifier. Based on the patterns of the predicted traffic, it classifies each 5-minute interval into 3 network states: Normal, Degraded, or Failure-prone. The classifier derives an SLA risk score from the prediction confidence; this provides a quantitative measure of the possibility of future service degradation. Then the predicted state and the risk scores are forwarded to the Recovery Recommendation Agent. The predicted network states and associated SLA risk scores are subsequently delivered to the Recovery Recommendation Agent.
- **Recovery Recommendation Agent:** It converts predicted SLA states into network recommendations. It applies pre-defined recovery policies linked with each network state, with mitigation steps. These recommendations provide network operators with actionable guidance before SLA

violations can occur. The generated mitigation recommendations are finally passed to the LLM Assistant Agent for operator-oriented explanation.

- **Operator Assistant Agent:** It serves as the conversational interface of the framework using the Mistral-7B LLM model. It translates the outputs of the previous agents into human-readable explanations to answer operator questions. Then these results are presented to the users using a web interface. This final stage converts structured predictive analytics into explainable natural-language responses that support intuitive operator interaction through the web interface.

III. DEMONSTRATION AND RESULTS

To demonstrate the proposed AI-native digital twin, we evaluate its end-to-end workflow for predictive SLA management. The framework forecasts future multi-service traffic, predicts SLA states, generates recovery recommendations, and delivers intent-driven operational assistance through natural-language interactions.




Fig. 2: Left: XGBoost traffic forecasting. Right: RF network-state classification.

Fig. 2 presents the prediction performance of the proposed framework. The XGBoost forecasting model accurately follows the temporal evolution of aggregate multi-service traffic throughout the day, capturing both gradual traffic variations

Example Queries	SLA Agent (Predictions)	LLM-Assisted Feedback
User Query "Tell me the next SLA violation for tomorrow."	<pre>{ "next_failure": { "time": "20:00", "predicted_total_gbps": 94.43, "risk_score": 0.51 }, "high_risk_intervals": { "20:05", "20:10", "20:15", "20:20" }, "recommended_action": "Backup + Redistribute + Reroute" }</pre>	Tomorrow's network-wide SLA risk occurs at 20:00, with a predicted total of 94.43 Gbps and a risk score of 0.51. The recommended action is to allocate backup capacity, redistribute traffic, and prepare rerouting before the SLA violation. Additional high-risk intervals occur at 20:05, 20:10, 20:15, and 20:20.
User Query "Which traffic service contributes most during the highest-risk interval?"	<pre>{ "highest_risk": { "time": "22:25", "risk_score": 0.89 }, "service_contribution": { "enterprise_gbps": 18.62, "ran_gbps": 45.16, "pon_gbps": 24.18, "total_gbps": 87.95, "dominant_service": "RAN", "dominant_share_percent": 51.3 }, "recommended_action": "Backup + Redistribute + Reroute" }</pre>	At 22:25, the highest-risk interval is predicted with a risk score of 0.89. RAN traffic contributes the most with 45.16 Gbps, corresponding to 51.3% of the total load. The recommended action is to allocate backup capacity, redistribute traffic, and prepare rerouting.
User Query "Generate an operational summary for tomorrow."	<pre>{ "sla_state_distribution_percent": { "normal": 75.0, "degraded": 13.5, "failure_prone": 11.5 }, "peak_risk": { "time": "22:25", "risk_score": 0.89 }, "recommended_action": "Backup + Redistribute + Reroute" }</pre>	Tomorrow's SLA outlook shows that 75% of the intervals will be normal, 13.5% degraded, and 11.5% failure-prone. The peak risk is expected at 22:25 with a risk score of 0.89. The recommended action is to prepare backup capacity, traffic redistribution, and rerouting before the high-risk window.

Fig. 3: Intent-driven operator interaction.

and peak-demand periods with high fidelity. The accompanying Random Forest classifier successfully categorizes the predicted traffic into *Normal*, *Degraded*, and *Failure-prone* network states, achieving perfect classification for the *Normal* and *Degraded* classes while maintaining reliable detection of *Failure-prone* conditions. Together, these results demonstrate that accurate traffic forecasting enables reliable anticipation of future SLA conditions, providing sufficient time for proactive mitigation before service degradation.

Building upon these predictive results, Fig. 3 demonstrates the interactive operation of the proposed framework. User intents are expressed as natural-language queries and automatically translated into structured SLA-aware analytics, which are subsequently transformed into concise, human-readable operational recommendations. The framework identifies forthcoming SLA violations, pinpoints the dominant traffic contributors responsible for elevated risk, summarizes the predicted network health, and recommends proactive mitigation strategies, including traffic redistribution, backup-capacity allocation, and rerouting. By seamlessly combining predictive analytics with explainable natural-language reasoning, the proposed framework bridges the gap between AI-driven network intelligence and operator decision-making, enabling proactive, transparent, and intent-driven SLA management. Consequently, operators can anticipate future SLA violations, understand their underlying causes, and receive actionable recommendations through an intuitive conversational interface.

IV. CONCLUDING REMARKS

This demonstration showcased an AI-based digital twin that enables predictive SLA management in coherent P2MP metro-access networks. By combining traffic prediction, SLA assessment, automated recovery recommendation, and natural-

language operator interaction, the proposed framework supports proactive identification of service degradation and timely operational decision-making. The demonstrated workflow illustrates how predictive network intelligence can be translated into actionable and explainable operator assistance, reducing reliance on reactive network management. Future work will focus on integrating real-time telemetry, closed-loop control, and autonomous resource optimization to further advance AI-native optical network operation.

ACKNOWLEDGMENTS

Funded by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie project NESTOR under G.A. No. 101119983. This work has also received support from ALLEGRO G.A. (101092766).

REFERENCES

- [1] D. Welch *et al.*, "Point-to-multipoint optical networks using coherent digital subcarriers," *Journal of Lightwave Technology*, vol. 39, no. 16, pp. 5232–5247, 2021.
- [2] D. Welch *et al.*, "Digital subcarrier multiplexing: Enabling software-configurable optical networks," *Journal of Lightwave Technology*, vol. 41, no. 4, pp. 1175–1191, 2023.
- [3] V. Curri, "Gnpy model of the physical layer for open and disaggregated optical networking," *Journal of optical communications and networking*, vol. 14, no. 6, pp. C92–C104, 2022.
- [4] V. Curri, "Digital-twin of physical-layer as enabler for open and disaggregated optical networks," in *2023 International Conference on Optical Network Design and Modeling (ONDM)*. IEEE, 2023, pp. 1–6.
- [5] S. Zeb *et al.*, "AI-driven converged metro-access optical network-as-a-service with point-to-multipoint coherent optics for 6G X-Hauling," *Computer Networks*, vol. 284, p. 112338, 2026.
- [6] I. C. Dipto *et al.*, "LLM-driven optical network-as-a-service (onaas) for next generation metro-access convergence," in *2026 IEEE International Conference on Machine Learning for Communication and Networking (ICMLCN)*, 2026, pp. 1–6.
- [7] S. Zeb *et al.*, "ML-Driven qot estimation for adaptive P2MP optical network-as-a-service in 6G X-Haul," in *2026 IEEE International Conference on Machine Learning for Communication and Networking (ICMLCN)*, 2026, pp. 1–6.