IJECBE (2026), **4**, **1**, 17–49
Received (29 April 2026) / Revised (22 May 2026)
Accepted (25 May 2026) / Published (30 March 2026)
https://doi.org/10.62146/ijecbe.v4i1.239
https://ijecbe.ui.ac.id
ISSN 3026-5258

International Journal of Electrical, Computer and Biomedical Engineering

RESEARCH ARTICLE

Delivery-Aware Hybrid Throughput Forecasting for Targeted Capacity Planning and Proactive Congestion Avoidance in Microwave Transmission Networks

Gede Deny Marthafani and Catur Apriono*

Department of Electrical Engineering, Faculty of Engineering, Universitas Indonesia, Depok, Indonesia
*Corresponding author. Email: catur@eng.ui.ac.id

Abstract

Abstract-Microwave transmission networks operate under strict physical-layer capacity constraints, where excessive utilization may reduce service reliability and increase congestion risk. In practical operations, capacity expansion decisions are both time-sensitive and location-sensitive, requiring operators to determine not only when an upgrade should be executed, but also which links require intervention. However, many existing forecasting approaches remain focused on numerical prediction accuracy and do not explicitly support delivery-aware and targeted planning decisions. This study proposes a delivery-aware hybrid smoothing–residual forecasting framework for link-level microwave throughput prediction using hourly Network Management System (NMS) data from 110 links observed over a six-month period. The framework separates the throughput series into a smoothed structural component and a residual component, where primary forecasting is performed on the structural signal, and residual error is modeled using LightGBM. The final forecast is reconstructed by combining both components and evaluated across multiple horizons using WMAPE, MAE, and RMSE. Experimental results show that CatBoost with LightGBM residual correction provides the most consistent performance across all forecasting horizons. Beyond improving prediction accuracy, the proposed framework enables delivery-aware and targeted capacity planning by estimating threshold-crossing timing and identifying high-risk links. This supports proactive congestion avoidance, more efficient upgrade prioritization, and improved resource allocation in microwave transmission networks.

Keywords: Microwave transmission; throughput forecasting; hybrid forecasting; LightGBM residual correction; multi-horizon prediction; zero-congestion planning

1. Introduction

Microwave transmission remains a critical Layer-1 infrastructure for mobile backhaul networks, particularly in areas where fiber deployment is constrained by geography, cost, or implementation limitations. In many practical network environments, microwave links continue to serve as a flexible and cost-efficient solution for connecting base stations to aggregation nodes. Compared with full fiber-based transport, microwave systems often offer faster deployment and lower initial investment. However, microwave networks operate under strict physical-layer capacity constraints, where achievable throughput is limited by channel bandwidth, modulation scheme, coding efficiency, spectral utilization, and propagation conditions [1]-[2].

As mobile traffic continues to grow, microwave capacity planning becomes increasingly challenging. In operational practice, transmission teams commonly monitor utilization thresholds to identify links that may require expansion. In this study, 80% utilization is treated as the primary congestion-planning threshold, since traffic approaching this level indicates that the available transmission headroom is becoming limited and the risk of congestion is increasing. Once a link is projected to cross this threshold, corrective actions such as bandwidth expansion, channel upgrade, rerouting, swap, or fiberization may be required [3]–[2]. The challenge is that these actions require lead time for procurement, material readiness, implementation, integration, and Ready-for-Service (RFS). As a result, late planning decisions may cause a link to remain above the safe utilization boundary before corrective action is completed.

In practical transmission operations, expansion decisions are not only time-sensitive but also location-sensitive. Operators must determine not only when an upgrade should be initiated, but also which specific links should be prioritized. Inaccurate timing may lead to delayed upgrades and prolonged congestion, whereas inaccurate link selection may result in inefficient capital deployment, where expansion is executed on links that are not yet critical while truly high-risk links remain underserved. Therefore, forecasting should not be viewed only as a numerical prediction task. Instead, it should function as a decision-support mechanism for timely and targeted capacity planning.

Under this perspective, forecasting becomes relevant not only for estimating future throughput, but also for identifying whether and when a link is likely to move toward the 80% utilization threshold. When implemented at the link level, forecasting can provide two important planning signals simultaneously: temporal guidance for determining upgrade timing and spatial guidance for identifying which links require intervention. This dual role is particularly important in microwave backhaul planning, where congestion risk is highly link-specific and operational resources must be allocated selectively.

Traffic forecasting in telecommunication systems has been widely studied using statistical, machine-learning, and neural-network-based approaches [4]-[5]. Classical models such as SARIMA are effective for capturing linear temporal dependence and recurring seasonality [6], [7], whereas machine-learning models such as CatBoost, LightGBM, and MLP are more flexible in learning nonlinear and heterogeneous traffic behavior [8], [5], [9]–[10]. However, many previous studies directly model the original traffic series without explicitly separating stable structural patterns from short-term irregular fluctuations.

For hourly microwave throughput data, this distinction is important because the observed series often contains both a smoother recurring component and irregular behavior caused by burst traffic, local disturbances, and residual fluctuation. Modeling both components in a single stage may reduce forecast robustness, particularly when traffic behavior differs across links and regions. In addition, many forecasting studies emphasize numerical error metrics such as MAE, RMSE, and WMAPE, but do not explicitly connect the forecast to threshold-based planning decisions, delivery timing, or targeted upgrade prioritization. As a result, forecasting accuracy is often evaluated in isolation from its operational usefulness.

To address these limitations, this study proposes a delivery-aware hybrid smoothing-residual forecasting framework for link-level microwave throughput prediction. The approach first extracts a smoothed structural component from the throughput series to represent the dominant temporal signal, and then models the remaining residual error separately using LightGBM as the residual learner. The final forecast is obtained by combining the structural and residual predictions. Beyond improving numerical accuracy, the framework is designed to support delivery-aware and targeted planning by estimating threshold-crossing timing and identifying links that require earlier intervention.

The proposed framework is evaluated using hourly Network Management System (NMS) data from 110 microwave links, consisting of 68 links in Central Java and 42 links in Bali-Lombok, observed from 1 September 2025 to 26 February 2026. Forecasting performance is assessed over 1-day, 3-day, 7-day, and 14-day horizons using WMAPE, MAE, and RMSE, and is further interpreted against the 80% utilization threshold to support congestion-risk assessment and upgrade planning. Although formal validation is limited to 14-day horizons, the framework is intended to operate in a continuous rolling-update mode, where forecasts are regenerated as new hourly observations become available.

This study contributes in three main ways. First, it proposes a link-level microwave throughput forecasting framework that explicitly reflects physical-layer capacity constraints and practical congestion thresholds. Second, it introduces a hybrid smoothing-residual modeling strategy that separates structural traffic behavior from irregular fluctuation to improve forecasting robustness. Third, it extends the role of forecasting from numerical prediction to operational decision support by linking threshold-crossing forecasts with delivery timing and targeted upgrade prioritization. In this way, the proposed framework supports proactive congestion avoidance, more effective resource allocation, and delivery-aware zero-congestion planning in microwave transmission networks.

The remainder of this paper is organized as follows. Section 2 reviews related work on microwave backhaul constraints, traffic forecasting, and hybrid forecasting models. Section 3 describes the proposed methodology. Section 4 presents the experimental results and discussion. Section 5 concludes the paper and outlines future work.

2. TRAFFIC FORECASTING ON TRANSMISSION MICROWAVE

2.1 Microwave Transmission and Backhaul Capacity Constraints

Microwave transmission remains an important component of mobile backhaul networks, particularly in areas where fiber deployment is constrained by geography, infrastructure availability, or cost efficiency [1]–[2]. In many practical deployment environments, microwave links provide a fast and flexible solution for connecting base stations to aggregation nodes [3]–[2]. However, unlike higher-layer packet networks, microwave throughput is bounded by strict physical-layer constraints, including channel bandwidth, modulation profile, coding efficiency, and propagation conditions [1], [3]. As utilization approaches the configured link capacity, the available margin for burst traffic and temporary degradation becomes increasingly limited, thereby increasing the risk of congestion and service instability [1], [11], [12].

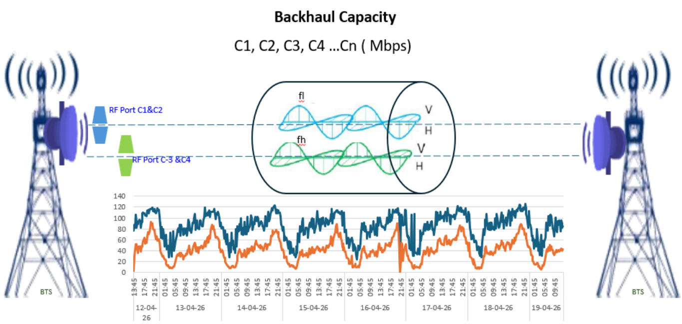

Figure 1. Microwave Backhaul Link Configuration and Observed Hourly Throughput Pattern

Figure 1 presents the relationship between microwave backhaul configuration and observed traffic behavior. The microwave link connects two BTS sites through radio channels that operate under specific frequency, polarization, and bandwidth constraints. These physical-layer parameters determine the maximum achievable throughput of the link. The traffic illustration in the lower part of the figure shows that hourly throughput may fluctuate over time while still maintaining certain recurring temporal patterns. As traffic demand increases, these fluctuations can push link utilization closer to the configured capacity limit. Therefore, the figure provides the motivation for applying link-level throughput forecasting to support proactive capacity planning and congestion avoidance in microwave transmission networks.

For this reason, threshold-based monitoring remains common in transmission operations. In this study, 80% utilization is treated as the primary operational planning threshold. Once a link moves toward this level, the issue is no longer simply one of monitoring, but of deciding how early upgrade preparation must begin. Thus, throughput forecasting becomes directly relevant to operational planning rather than merely to traffic estimation.

2.2 Network Traffic Forecasting and Model Perspectives

Traffic forecasting has long been used in telecommunication systems to anticipate future demand, improve resource allocation, and support network expansion decisions [4]–[5]. Classical models such as SARIMA are effective for capturing linear temporal dependence and recurring seasonality [6], [7]. In contrast, machine-learning approaches such as CatBoost, LightGBM, and MLP are more flexible in learning nonlinear and heterogeneous traffic behavior [8], [5], [9]–[10].

This distinction is important for link-level microwave forecasting. Hourly microwave throughput often contains both a smoother recurring temporal structure and irregular local fluctuation. Under such conditions, SARIMA remains relevant as a structured statistical baseline, while boosting-based models and neural networks provide alternative perspectives for capturing nonlinear behavior. In particular, LightGBM is suitable for residual correction because it is efficient and effective for structured nonlinear regression, whereas CatBoost is theoretically attractive for the structural forecasting stage due to its ability to model nonlinear interactions across heterogeneous traffic patterns [9], [13], [10].

2.3 Model Characteristics for Throughput Forecasting

2.3.1 SARIMA

Seasonal Autoregressive Integrated Moving Average (SARIMA) is one of the most established models in time-series forecasting and remains useful when the observed series shows relatively stable temporal dependence and recurring seasonality [6], [7]. It works well when future values can be described by a structured combination of autoregressive, differencing, and seasonal components.

That said, SARIMA is less flexible when the data become heterogeneous, partially non-stationary, or affected by nonlinear local fluctuation. This limitation is relevant for hourly microwave throughput, where recurring structure often coexists with burst traffic, irregular recovery behavior, and local disturbances. Under such conditions, SARIMA remains a meaningful baseline, but it may not fully capture the complexity of link-level traffic behavior.

2.3.2 LightGBM

LightGBM is a gradient-boosting model designed for efficient regression and classification on structured data [9], [10]. Its main advantage lies in its ability to capture nonlinear relationships and feature interactions while remaining computationally efficient. This makes it particularly attractive in forecasting settings where the target variable depends on complex combinations of lagged features, temporal descriptors, or contextual inputs.

In the present study, LightGBM is not treated mainly as a competing primary forecaster. Instead, it plays a more specific role as the residual learner. This is theoretically reasonable because the residual component is expected to contain the part of the signal that remains nonlinear, irregular, and difficult to explain through a simpler structural forecast. Under that setup, LightGBM provides a practical way to improve the final prediction without making the hybrid architecture unnecessarily complex.

2.3.3 CatBoost

CatBoost is also a boosting-based model and is widely recognized for its strong performance in structured regression problems [13]. It is particularly useful when the data contain nonlinear interactions and heterogeneous patterns that are difficult to represent through a single global linear structure. For link-level microwave forecasting, this matters because different links may behave quite differently even when they belong to the same network region.

Within the proposed framework, CatBoost is used as one of the candidate primary forecasting models for the smoothed structural series. This makes theoretical sense because, once the short-term noise has been reduced through smoothing, the remaining structural signal may still be too complex for a linear seasonal model but still sufficiently organized to be learned effectively by a boosting algorithm.

2.3.4 Multilayer Perceptron (MLP)

Multilayer Perceptron (MLP) is a neural-network-based model that can learn nonlinear input–output relationships [14]. In forecasting, MLP is often used as a general nonlinear benchmark because it is flexible and can model complex patterns without requiring explicit assumptions about functional form.

However, that flexibility comes with trade-offs. MLP is often more sensitive to parameter tuning, training configuration, and data consistency than boosting-based models. In operational forecasting problems where the data are link-specific, the training window is limited, and the model is expected to behave consistently across many cases, that sensitivity may affect robustness. For this reason, MLP is included here as a useful comparison model rather than as the main residual-correction method.

2.4 Hybrid Smoothing–Residual Forecasting

Forecasting the raw traffic series in a single stage is common, but it becomes less reliable when stable temporal structure and irregular fluctuation coexist in the same signal [15], [16], [17]. This is often the case with hourly microwave throughput, where daily or weekly regularity may still be visible, while burst traffic and local residual deviation remain significant.

Hybrid forecasting provides a natural solution to this problem by separating the dominant structural component from the remaining residual behavior [15], [16], [17]. The structural component is forecast first, and the remaining error is then modeled separately. This approach is particularly relevant in the present study because the throughput series is assumed to contain both a learnable smoothed signal and a nonlinear residual component. This forms the theoretical basis for combining a primary forecasting model with LightGBM-based residual correction.

2.5 Research Gap and Positioning of This Study

Despite the growing literature on network traffic forecasting, several important gaps remain when the problem is viewed from the perspective of practical microwave transmission planning.

First, many existing studies focus on aggregate traffic prediction or general network-wide forecasting, whereas real transmission expansion is executed at the individual-link level [4]–[16], [18]. In operational environments, each microwave link has its own configured capacity, traffic growth pattern, utilization exposure, and upgrade urgency. As a result, aggregate prediction is not sufficient for transmission planning that requires link-specific intervention decisions.

Second, prior studies commonly assess forecasting quality mainly through numerical error metrics such as MAE, RMSE, and WMAPE [15], [16], [6], [7], [17]. Although these metrics are important, they do not fully represent operational usefulness. In microwave planning, the key decision question is not only whether the predicted throughput is close to the actual value, but also whether the forecast can indicate when utilization is likely to approach or cross the operational threshold that triggers upgrade preparation.

Third, many forecasting approaches are developed as prediction tools only, without clear integration with delivery timing and execution feasibility. In practice, transmission upgrades require procurement lead time, material readiness, implementation, integration, and operational safety margins. Therefore, a forecast becomes operationally meaningful only if it supports sufficiently early action before congestion persists.

Fourth, prior studies rarely address the issue of targeted upgrade allocation. In real transmission planning, the problem is not only when to expand capacity, but also where intervention should be prioritized. Inaccurate prioritization may result in inefficient capital deployment, where upgrades are performed on links that are not yet critical while truly high-risk links remain underserved. This indicates that forecasting should support both temporal decision accuracy and spatial decision accuracy.

Fifth, many studies directly forecast the original traffic series without explicitly separating stable structural behavior from short-term irregular fluctuation [15], [16], [17]. For hourly microwave throughput data, this distinction is important because the observed series often combines recurring temporal structure with burst traffic, local disturbances, and residual irregularity. Under such conditions, single-stage forecasting may reduce robustness, especially across heterogeneous links and regions.

Against this background, the present study is positioned as a delivery-aware hybrid smoothing-residual forecasting framework for link-level microwave throughput prediction. The framework operates at the individual-link level, separates structural behavior from residual fluctuation, and applies LightGBM as a fixed residual learner across all hybrid configurations. More importantly, the framework extends forecasting from numerical prediction into decision support by linking threshold-crossing estimation with delivery timing and targeted upgrade prioritization. In this way, the study contributes not only to forecasting methodology, but also to practical zero-congestion planning in microwave transmission networks.

Table 1 highlights the conceptual shift from accuracy-oriented forecasting toward a delivery-aware and business-oriented planning framework. Beyond numerical prediction, the proposed approach enables timely and targeted upgrade decisions, supports more efficient capital allocation, and contributes to congestion avoidance and service stability. These operational improvements may further support better customer experience and indirectly reduce churn risk, thereby enhancing the overall effectiveness of microwave network planning.

Table 1. From Forecasting Accuracy to Delivery-Aware and Business-Oriented Planning

Dimension	Existing Studies	This Study
Forecasting Scope	Aggregate / network-level	Link-level forecasting
Objective	Prediction accuracy	Decision-oriented forecasting
Threshold Awareness	Not explicitly considered	80% threshold-based planning
Delivery Integration	Not included	Delivery-aware (PO timing)
Upgrade Targeting	Not addressed	Targeted link prioritization
Traffic Modeling	Direct raw series modeling	Smoothing + residual decomposition
Residual Handling	Implicit	Explicit (LightGBM correction)
Planning Outcome	Analytical insight	Actionable planning framework
Resource Allocation	Not discussed	More efficient CAPEX allocation
Operational Impact	Not explicitly linked	Congestion avoidance & QoS stability
Business Implication	Not considered	Potential support for service continuity and business protection

Based on these identified gaps, the next section presents the proposed methodology, including data preprocessing, traffic-pattern categorization, hybrid throughput forecasting, threshold-crossing analysis, and delivery-feasibility interpretation.

3. Methodology

This study proposes a delivery-aware hybrid smoothing–residual forecasting framework for link-level microwave throughput prediction under physical-layer capacity constraints. The methodology is designed to estimate future throughput and to determine whether and when a link is likely to approach or exceed the 80% utilization threshold, enabling early identification of congestion risk for transmission planning [3], [2].

The framework is structured as a multi-stage workflow that integrates data preparation, hybrid forecasting, and planning-oriented interpretation. At the data preparation stage, hourly NMS throughput data are processed to ensure temporal consistency and analytical reliability. Traffic pattern categorization is also applied to capture heterogeneous link behavior and provide contextual insight into forecasting conditions across different links and regions [15], [19],[20] [18], [6].

The forecasting stage is based on decomposition of the observed throughput series into a smoothed structural component and a residual component. The structural component captures dominant temporal behavior, while the residual component represents irregular and nonlinear fluctuations. This design follows hybrid forecasting principles, where the primary signal is modeled first and the remaining error is learned separately to improve robustness [15], [16], [17]. Primary forecasting is performed on the structural signal using SARIMA, LightGBM, CatBoost, and MLP, while the residual component is modeled using LightGBM as a fixed error-correction layer [6]–[10]. The final forecast is reconstructed by combining both components.

To support planning, the forecast output is evaluated against the 80% utilization threshold to estimate threshold-crossing timing. This estimate is then integrated with deterministic delivery lead-time modeling to assess upgrade feasibility and determine the latest feasible PO release date [3], [2]. Through this integration, the methodology links forecasting outputs with practical decision-making in microwave transmission planning. The complete workflow of the proposed framework is illustrated in Figure 2.

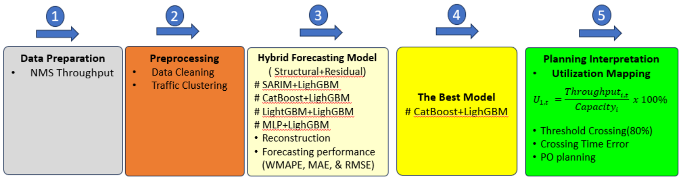

Figure 2. Proposed Delivery-Aware Hybrid Forecasting Framework

Figure 2 presents the proposed delivery-aware hybrid forecasting framework. The process begins with hourly NMS throughput data, followed by preprocessing, data cleaning, and traffic clustering. Several hybrid forecasting configurations are then evaluated, including SARIMA + Residual LightGBM, CatBoost + Residual LightGBM, LightGBM + Residual LightGBM, and MLP + Residual LightGBM, with performance measured using WMAPE, MAE, and RMSE. Based on the evaluation results, CatBoost + Residual LightGBM is selected as the best-performing model and is used for planning interpretation. The selected model output is mapped into utilization using configured link capacity and further analyzed against the 80% threshold to determine threshold-crossing events, crossing-time error, and PO planning. Thus, the framework translates forecasting results into actionable capacity-planning decisions for proactive congestion avoidance.

3.1 Data Preprocessing

The dataset used in this study consists of hourly microwave throughput measurements collected from the Network Management System (NMS) of a mobile network operator. For each microwave link, the dataset includes observed throughput, configured link capacity, and recorded utilization. However, forecasting in this study is performed in the throughput domain, while utilization is used for planning interpretation. The observation period spans from 1 September 2025 to 26 February 2026 and covers 110 microwave links, consisting of 68 links in Central Java and 42 links in Bali–Lombok. Since the data are recorded hourly, each link provides 24 observations per day, enabling the analysis of intraday and weekly traffic behavior. Link-level forecasting is adopted because each microwave link may operate under different bandwidth, modulation profile, and capacity limits, making aggregate modeling less suitable for operational planning [1]–[11], [4].

Before forecasting, several preprocessing steps are performed to ensure temporal consistency and analytical validity. First, all observations are aligned to uniform hourly timestamps. Second, missing data caused by temporary monitoring interruptions are handled to preserve time-series continuity. Third, abnormal spikes and extreme anomalies are filtered to reduce their adverse impact on clustering, smoothing, and forecasting stages.

3.2 Traffic Pattern Categorization

Microwave traffic behavior may vary substantially across service areas and regions. Prior studies show that network traffic often exhibits heterogeneous temporal patterns, including regular seasonality, burst dynamics, sudden local drops, and irregular fluctuation [15], [16], [18], [21], [22] [8], [23]. Such heterogeneity is relevant because links with different temporal behavior may also differ in forecasting difficulty and congestion-risk evolution.

To capture this heterogeneity, clustering-based traffic pattern categorization is conducted before forecasting. Each link is represented using summary descriptors derived from the hourly throughput series, including average throughput, volatility, daily autocorrelation, weekly autocorrelation, midday variation, recovery behavior, and peak-shift tendency. These descriptors are chosen because they describe the temporal morphology of link behavior rather than merely its average magnitude [15], [16], [18], [6].

Based on centroid interpretation, the links are grouped into four traffic-pattern categories: Cluster-0, Cluster-1, Cluster-2, and Cluster-3. Cluster-0 represents Healthy Seasonality, where the traffic pattern shows stable repetition and strong daily or weekly structure. Cluster-1 represents Midday Sink, where the traffic still preserves daily structure but contains a noticeable reduction during certain daytime periods. Cluster-2 represents Broken Recovery, which is characterized by irregular drops followed by unstable recovery behavior. Cluster-3 represents Shifted Seasonal Pattern, where the peak timing and seasonal structure tend to shift across time. Figures 3–6 illustrate representative traffic patterns for each cluster.

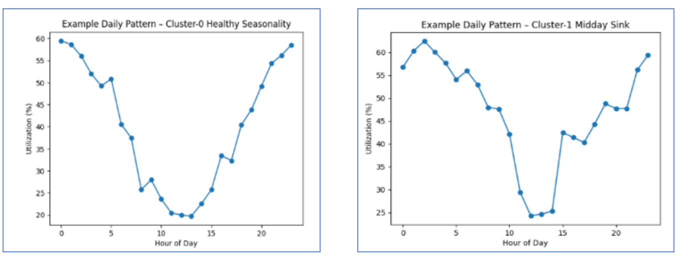

Figure 3. Representative Traffic Pattern of Cluster-0: Healthy Seasonality

Figure 3 shows the representative traffic pattern of Cluster-0, which is characterized by relatively stable and repeated daily or weekly traffic behavior. This pattern indicates that the link has strong seasonality and lower volatility, making it more suitable for forecasting. From a planning perspective, links in this cluster are generally easier to monitor because their future behavior tends to follow a more predictable structure.

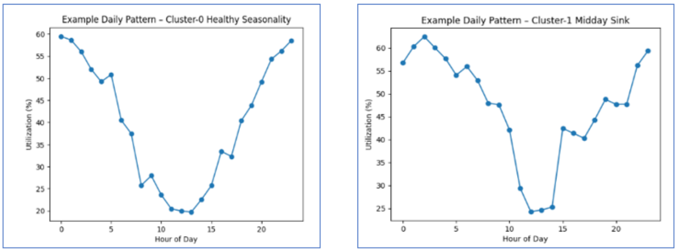

Figure 4. Representative Traffic Pattern of Cluster-1: Midday Sink

Figure 4 presents the representative pattern of Cluster–1, where the traffic still shows daily structure but contains a noticeable traffic reduction during certain daytime periods. This midday sink behavior indicates moderate variability in the traffic profile. Although the pattern remains forecastable, the presence of local drops may increase forecasting difficulty compared with Cluster–0.

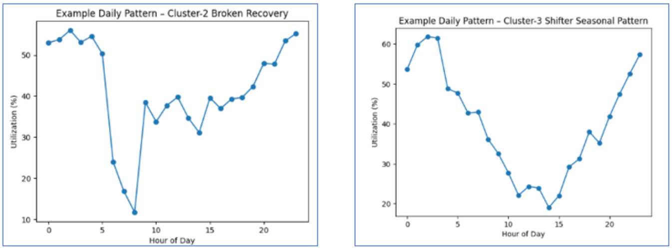

Figure 5. Representative Traffic Pattern of Cluster-2: Broken Recovery

Figure 5 illustrates the traffic pattern of Cluster–2, which is marked by irregular traffic drops followed by recovery behavior. This pattern reflects a more unstable traffic condition, where the link does not consistently return to the same seasonal structure after fluctuation. As a result, forecasting for this cluster becomes more challenging, and residual correction becomes important to capture irregular deviations.

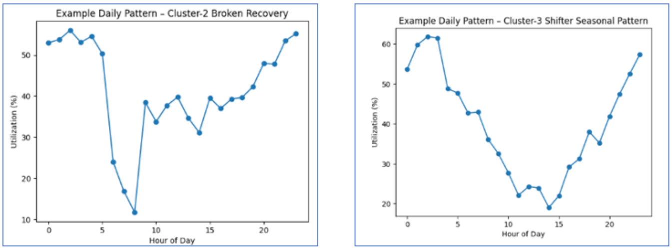

Figure 6. Representative Traffic Pattern of Cluster-3: Shifted Seasonal Pattern

Figure 6 shows the representative traffic behavior of Cluster-3, where the seasonal pattern tends to shift over time. This cluster has higher volatility and less stable peak timing compared with the other clusters. Operationally, links in this cluster require closer monitoring because the shifting traffic behavior may reduce forecast stability and increase the risk of prediction error.

Table 2. Data Clustering Forecast Risk

Cluster	Pattern	Operational Interpretation	Utilization Characteristics	Seasonality Indicator	Behavioral Indicator
Cluster-0	Healthy Seasonality	Stable traffic growth	low volatility	Strong daily and weekly autocorrelation (ACF(24) > 0.60, ACF(168) >0.50)	Stable peak hour, peak \shift std <2
Cluster-1	Midday Sink	Behavioral traffic pattern	moderate volatility	Daily seasonality present (ACF(24) >0.50)	Significant midday traffic drop
Cluster-2	Broken Recovery	Irregular traffic fluctuation	high volatility	Weak seasonality (ACF(24) < 0.40)	Frequent traffic drops followed by recovery
Cluster-3	Shifter Seasonal	Potential operational anomaly	Very High volatility	Moderate seasonality ACF(24) with inconsistent autocorrelation pattern	Peak-hour shifting across weeks

Table 2 presents the four traffic-pattern clusters used to characterize heterogeneous microwave link behavior. Each cluster is described based on its operational interpretation, utilization behavior, seasonality indicator, and temporal characteristics. This classification helps explain why some links are easier to forecast than others, since stable seasonal links tend to produce more reliable forecasts, while volatile or shifting patterns introduce higher forecasting difficulty.

This categorization does not replace independent per-link forecasting; instead, it improves interpretability by describing the traffic environment in which each model is expected to operate. In this sense, clustering serves as a contextual analytical layer that explains heterogeneous forecasting performance across links and regions [15], [16], [6], [7], [17].

3.3 *Structural Smoothing of Throughput Series*

Since forecasting is performed in the throughput domain, the observed throughput series is decomposed into structural and residual components in order to reduce short-term fluctuation and improve forecasting stability. This is consistent with general time-series principles stating that observed traffic may contain a dominant temporal signal together with irregular deviations [15], [16], [6], [7], [17]. Let $Y_t$denote the observed throughput at time t. The series is decomposed as

$$Y_t = S_t + R_t \tag{1}$$

where S_t denotes the structural component and R_t denotes the residual component. The structural component captures the smoother long-term and seasonal behavior that forms the dominant traffic pattern, while the residual component represents irregular fluctuation that remains after structural extraction [15], [16], [6], [7], [17].

In this study, exponential smoothing is applied to estimate the structural signal:

$$S_{i,t} = \alpha Y_{i,t} + (1 - \alpha)S_{i,t-1} \tag{2}$$

Where: $S_{i,t}$= Smoothed structural component $Y_{i,t}$= Observed throughput value α = smoothing parameter ($0 < \alpha < 1$)

A higher value of α makes the smoothed series more responsive to recent observations, while a lower value of α produces a smoother structural component by giving more weight to past observations.

To improve reproducibility and avoid manual parameter selection, the smoothing parameter α was determined using an Optuna-based optimization process implemented in the Python forecasting program. The optimization process was conducted by considering the predefined traffic-pattern clusters, so that the smoothing configuration could reflect the structural behavior of each cluster. Optuna first recommended the smoothing span for each cluster based on the observed throughput behavior. The corresponding smoothing parameter was then derived using the standard exponential smoothing relationship:

$$\alpha = \frac{2}{Span + 1} \tag{3}$$

where span represents the number of data points used to define the smoothing responsiveness. Based on the Optuna optimization results, the recommended α values for Cluster-0, Cluster-1, Cluster-2, and Cluster-3 were 0.67, 0.50, 0.50, and 0.67, respectively. These values correspond to short smoothing spans and indicate that the structural extraction was designed to remain responsive to hourly microwave traffic dynamics. Cluster-0 and Cluster-3 used a higher α, indicating a more responsive smoothing configuration, while Cluster-1 and Cluster-2 used a moderate α, indicating a slightly smoother extraction of the structural component.

Table 3 shows the smoothing parameter selection obtained from the Optuna-based optimization process for each traffic cluster. The recommended alpha values indicate that different traffic patterns require different levels of smoothing responsiveness. Cluster-0 and Cluster-3 use higher alpha values to preserve responsive structural movement, while Cluster-1 and Cluster-2 use moderate smoothing to reduce local fluctuation while maintaining the dominant traffic trend.

This cluster-based alpha selection improves the representativeness of the smoothed structural signal under heterogeneous microwave traffic patterns. Links with stable seasonality and shifting seasonal behavior require a more responsive smoothing configuration so that the structural component can still follow relevant hourly dynamics. In contrast, links with midday sink or broken recovery characteristics require moderate smoothing to reduce short-term irregular fluctuation while preserving the dominant structural trend. As a result, the smoothed structural component becomes more suit-

Table 3. Optuna-Based Smoothing Parameter Selection by Traffic Cluster

Traffic Cluster	Traffic Pattern	Recommended Span	Derived $\alpha = \frac{2}{Span+1}$	Interpretation
Cluster-0	Healthy Seasonality	2	0.67	More responsive smoothing to preserve regular traffic dynamics
Cluster-1	Midday Sink	3	0.5	Moderate smoothing to reduce local fluctuation while preserving structural behavior
Cluster-2	Broken Recovery	3	0.5	Moderate smoothing to balance irregular recovery behavior and structural trend
Cluster-3	Shifted Seasonal Pattern	2	0.67	More responsive smoothing to capture shifting seasonal behavior

able for subsequent primary forecasting, while the remaining irregular fluctuation is handled separately through residual modeling.

3.4 Residual Component Extraction

After the structural smoothing stage, the residual component is obtained by subtracting the smoothed structural signal from the original throughput series. The residual term is defined as:

$$e_{i,t} = Y_{it} - S_{i,t} \tag{4}$$

where:
$e_{i,t}$ represents irregular traffic fluctuations not captured by the structural model. Residual modeling is important because the forecasting problem in microwave networks is not purely seasonal. Even when the dominant daily and weekly structure is stable, short-term irregular behavior can materially affect the time at which utilization crosses the operational threshold. Therefore, modeling the residual component separately allows the proposed hybrid framework to capture nonlinear variation that would otherwise reduce the planning reliability of the forecast [15], [16], [10], [17].

3.5 Primary Forecasting Models and Residual Learning

All forecasting models in this study are applied to the throughput series, ensuring consistency in the modeling domain. To evaluate forecasting robustness across different model families, this study uses four candidate primary forecasting models for the smoothed structural series:

- SARIMA, to capture temporal dependence and seasonality [6], [7]
- LightGBM, for nonlinear gradient-boosting regression [9], [10]
- CatBoost, for robust boosting-based prediction on structured data [13]
- Multi-Layer Perceptron (MLP), as a neural-network baseline [14]

The purpose of these candidate models is to forecast the smoothed structural series, not the raw series directly. This design follows the hybrid forecasting principle that the primary stage should first model the dominant predictable structure before residual correction is applied [15], [16], [17].

For the residual component, In this study, LightGBM is used as the sole residual learner because the residual stage is designed as a consistent error-correction layer rather than as a separate model-comparison task. After the primary model captures the smoothed structural component, the remaining residual mainly contains irregular and nonlinear error patterns. LightGBM is selected for this role due to its strong regression capability, computational efficiency, and suitability for modeling structured nonlinear residual behavior. Keeping the residual learner fixed also ensures methodological consistency across hybrid configurations and allows the comparison to remain focused on the forecasting capability of the primary models. Accordingly, the hybrid configurations evaluated are:

- Configuration 1: SARIMA + Residual LightGBM
- Configuration 2: LightGBM + Residual LightGBM
- Configuration 3: CatBoost + Residual LightGBM
- Configuration 4: MLP + Residual LightGBM

The final forecast is reconstructed as:

$$\hat{Y}_t = \hat{S}_t + \hat{R}_t \tag{5}$$

where $\hat{S}_t$ is the predicted structural component and $\hat{R}_t$ is the predicted residual component.

3.6 Multi-Horizon Forecasting Setup

Forecasting experiments are conducted over four operational horizons 1-day, 3-day, 7-day and 14-day. These horizons are selected because they are operationally relevant for early warning and near-term transmission planning. Although formal model validation in this study is limited to these horizons, the proposed framework is intended to be used in a continuous rolling-update manner. In practice, forecasts can be regenerated periodically as new hourly observations become available, allowing the framework to support longer-term planning through repeated multi-horizon updates rather than relying on a single long direct forecast.

3.7 Computing Environment and Experimental Implementation

All preprocessing, clustering, smoothing, forecasting, and residual-learning experiments in this study were implemented in a Python-based computing environment. The main libraries used include pandas and numpy for data handling, scikit-learn for preprocessing and clustering support, statsmodels for SARIMA modeling, and LightGBM and CatBoost for machine-learning-based forecasting and residual correction. This implementation setup is sufficient to support practical link-level throughput forecasting using hourly NMS data across multiple forecasting horizons.

Table 4. Environment Experiment

Item	Description
Programming Language	Python 3.11.4
Data Processing	Pandas Numpy
Statistical Model	Statsmodels Pmdarima
Machine Learning	Scikit-learn CatBoost LightGBM
Forecast Scope	110 links, hourly data
Forecast Horizon	1-day; 3-day; 7-day; 14-day
CPU Utilization	8 Core
RAM	32 GB
OS	Windows 11

Table 4 summarizes the computing environment and software libraries used in the experimental implementation. The forecasting workflow was implemented using Python with supporting libraries for data processing, statistical modeling, machine learning, and residual correction. This environment provides sufficient computational capability to process hourly NMS data from 110 microwave links across multiple forecasting horizons.

3.8 Forecast Performance Evaluation

The predictive performance of the proposed framework is evaluated using three standard error metrics: Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and Weighted Mean Absolute Percentage Error (WMAPE), which are widely used in forecasting literature [15], [6], [7].

$$MAE = \frac{1}{n}\sum_{t}^{n} |Yt - \hat{Y}t| \; (unit\ Mbps) \tag{6}$$

$$RMSE = \sqrt{\frac{1}{n}\sum_{t=1}^{n}(Yt - \hat{Y}t)^2 \; (unit\ Mbps)} \tag{7}$$

$$WMAPE = \frac{100}{n}\sum_{t=1}^{n} |\frac{Y_t - \hat{y}_t}{Y_t}| \; (\tag{8}$$

Where Y_t is the actual throughput value, Y_t is the predicted throughput value, and is the number of observations in the evaluation period. These metrics quantify numerical prediction accuracy, but in this study they are complemented by threshold-oriented evaluation to better reflect planning usefulness [15], [16], [6].

3.9 Threshold-Crossing Analysis

After the final throughput forecast is obtained, predicted throughput is mapped into utilization relative to the configured link capacity. In this study, utilization is not treated as the primary forecasting target, but as an operational indicator used to interpret congestion risk and planning thresholds.

Throughput values are converted into utilization values relative to configured microwave capacity so that traffic intensity can be interpreted in operationally meaningful terms [15], [8], [24], [6], [7], [17].

The utilization for link i at time t is defined as:

$$U_{1,t} = \frac{Throughtput_{i,t}}{Capacity_i} x100 \tag{9}$$

Where:

$Throughput_{i,t}$ is measured hourly traffic, $Capacity_i$ the configured Layer-1 capacity determined by bandwidth and modulation scheme. Each link operates under different bandwidth and modulation configurations; therefore, modeling is performed independently per link.

For planning interpretation, this study adopts a delivery-aware utilization policy designed to eliminate congestion risks through proactive transmission planning. The proposed policy introduces an 80% utilization threshold as a hard operational boundary for upgrade completion. Table 5 presents the proposed utilization categories under the zero-congestion planning framework.

Table 5. Planning Management Stage

Planning Action	Utilization Range (%)
Risk-Critical-1	100
Risk-Critical-2	90-100
Risk-High	80-90
Normal-Medium High	70 – 80
Normal-Medium	50 – 70
Normal-Low	<50

Table 5 defines the utilization-based planning stages used to interpret congestion risk. Utilization below 50% is considered low-risk, while higher utilization ranges require closer monitoring and earlier planning action. In particular, the 80–90% range is treated as a high-risk delivery zone, where upgrade preparation should already be completed or accelerated to prevent the link from entering a more critical congestion condition.

By shifting the operational threshold from reactive monitoring to proactive planning, this policy enables network operators to initiate transmission upgrades before congestion occurs. This approach significantly reduces the risk of emergency upgrades and improves network stability under rapidly growing traffic demand.

Finally, the predicted threshold-crossing time is integrated with deterministic upgrade lead-time modeling to evaluate the feasibility of transmission expansion. Microwave links are considered to approach congestion risk when utilization exceeds 80% of configured capacity. The predicted threshold-crossing time for link is defined as:

$$t_i^* = Min\left\{t : \hat{U}_{i,t} \geq 80\%\right\} \tag{10}$$

where t_i^* represents the earliest predicted time when utilization exceeds the operational threshold. To measure forecasting stability, the crossing-time prediction error is defined as:

$$E_{cross} = \hat{t}_1 - t_1 \tag{11}$$

Where:

- E_{cross} = Crossing Time Error

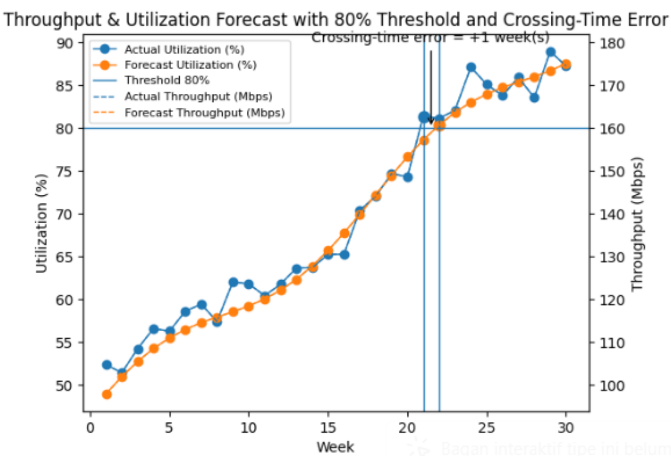

Figure 7. Illustration of Crossing-Time Error Between Predicted and Actual Threshold-Crossing Events

Figure 7 illustrates the gap between the predicted and actual threshold-crossing times. This figure explains how E_{cross} is calculated and interpreted in relation to the 80% utilization threshold. A smaller crossing-time error indicates that the model can estimate the congestion-risk timing more accurately, which is important for determining when upgrade preparation or PO release should be initiated.

To clarify the direction of the crossing-time prediction error, the sign of E_{cross} is interpreted as follows. A negative value of E_{cross} indicates that the predicted threshold-crossing time occurs earlier than the actual threshold-crossing time. This means that the model provides an early warning before the actual crossing event occurs.

In contrast, a positive value of E_{cross} indicates that the predicted threshold-crossing time occurs later than the actual threshold-crossing time, meaning that the model detects the crossing after the actual event. From an operational planning perspective, early prediction may provide additional preparation time, while late prediction may increase the risk of delayed intervention. Therefore, both the sign and the absolute magnitude of E_{cross} are useful for interpreting planning reliability.

Table 6. Interpretation of Crossing-Time Error E_{cross}

Ecross Value	Meaning	Operational Implication
$E_{cross} < 0$	The predicted crossing occurs earlier than the actual crossing	Safer for planning because it provides an early warning
$E_{cross} = 0$	The predicted crossing occurs exactly at the actual crossing time	Highly accurate crossing-time prediction
$E_{cross} > 0$	The predicted crossing occurs later than the actual crossing	More risky because the actual threshold crossing has already occurred

As shown in Table 6, both the sign and magnitude of E_{cross} are important for operational interpretation. A negative value is generally safer from a planning perspective because it provides an earlier warning, whereas a positive value indicates a late prediction and may increase the risk of delayed intervention. while the sign of E_{cross} indicates whether the model prediction is early or late.

3.10 Delivery Feasibility Analysis

The final methodological stage integrates the predicted threshold-crossing time with deterministic upgrade lead time in order to evaluate execution feasibility. In practical transmission operations, an accurate forecast is useful only if it provides sufficient time for procurement, installation, integration, and service activation before the link enters the high-risk utilization zone [3], [2].

$$t_{Po,i} = t_{cross} - \left(L_{Po} + L_{impl} + B_{uffer}\right) - E_{cross} \tag{12}$$

Where:

- t_{cross} = Date of crossing time on threshold 80%
- L_{Po} = lead time PO release till material readiness
- L_{impl} = lead time implementation
- B_{uffer} = safety buffer
- $t_{Po,i}$ = latest PO release date

The latest feasible PO release date is defined as the predicted threshold-crossing date minus the total delivery lead time. In this study, the total delivery lead time consists of procurement lead time, material readiness, implementation and integration duration, and an operational safety buffer. This deterministic formulation is adopted as a baseline planning assumption to simplify execution-feasibility analysis and to provide an interpretable decision rule for transmission upgrade planning.

Although this study uses deterministic lead-time modeling as a baseline planning assumption, real-world delivery lead time may vary due to procurement approval duration, vendor material availability, logistics constraints, field readiness, and integration complexity. Therefore, the safety buffer B_{uffer} can be interpreted not only as a fixed operational margin, but also as an adjustable parameter to absorb delivery uncertainty. In future implementation, B_{uffer} may be defined dynamically using historical procurement performance, vendor SLA reliability, material readiness status, or probabilistic lead-time distributions. This allows the proposed framework to remain compatible with both deterministic baseline planning and more realistic stochastic delivery conditions.

4. Data Analysis and Experimental Results

4.1 Dataset Description

The dataset used in this study consists of hourly microwave throughput measurements collected from the Network Management System (NMS) of a mobile network operator. The observation period spans from 1 September 2025 to 26 February 2026 and includes 110 microwave links, consisting of 68 links in Central Java (CJ) and 42 links in Bali–Lombok (BL). Each link is observed at hourly granularity, allowing the forecasting framework to capture daily and short-term utilization dynamics relevant to operational transmission planning.

In this study, the forecasting framework is evaluated under a link-level setting, where each microwave link is modeled independently according to its observed throughput behavior and configured capacity. This approach is important because microwave congestion risk is inherently link-specific. Different links may exhibit different utilization growth, volatility, seasonal behavior, and threshold-crossing tendencies, even within the same region.

Table 7. Region and Cluster Traffic

Region	Cluster Data				Total
	0	1	2	3	
Bali-Lombok(BL)	19	9	9	5	42
Central Java(CJ)	12	29	22	5	68
Grand Total	**31**	**38**	**31**	**10**	**110**

Table 7 summarizes the distribution of the 110 microwave links across regions and traffic clusters. Central Java contributes 68 links, while Bali–Lombok contributes 42 links. The table also shows that Central Java contains more Cluster-1 and Cluster-2 links, whereas Bali–Lombok has a higher proportion of Cluster-0 links. This indicates that the two regions have different traffic morphologies, making the dataset suitable for testing forecasting robustness under heterogeneous operating conditions.

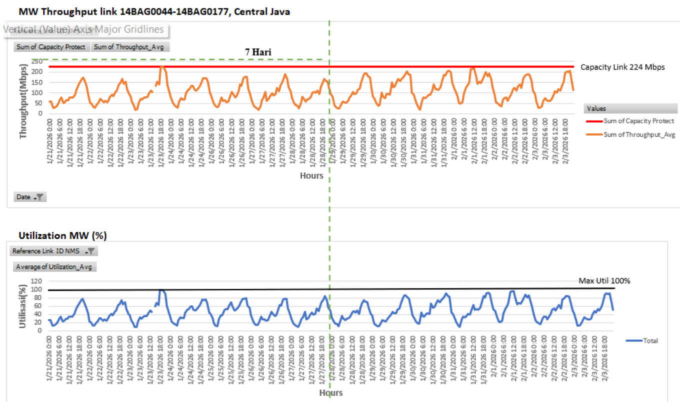

Figure 8. Example of Hourly NMS Throughput Data Used for Forecasting

Figure 8 presents an example of hourly throughput data collected from the NMS. The figure shows that microwave traffic contains both recurring temporal patterns and short-term fluctuations. This characteristic supports the need for a hybrid approach, where the structural component is separated from irregular residual behavior before forecasting is performed.

4.2 Hybrid Forecasting and Horizon-Wise Forecasting Performance

The proposed hybrid forecasting framework combines primary forecasting on the smoothed structural throughput series with residual correction using LightGBM. Four hybrid configurations are evaluated: SARIMA + Residual LightGBM, LightGBM + Residual LightGBM, CatBoost + Residual LightGBM, and MLP + Residual LightGBM.

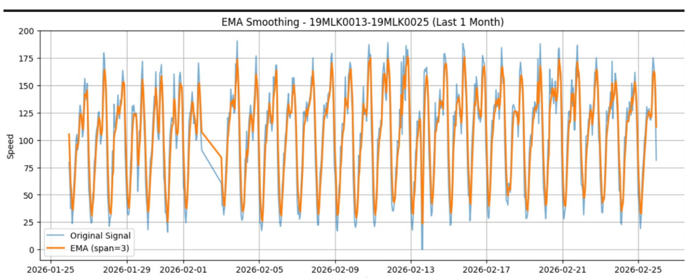

Figure 9. Smoothed Structural Component of the Throughput Series

Figure 9 shows the smoothed structural component obtained from the original throughput series. The smoothing process reduces short-term noise while preserving the dominant traffic trend and seasonal behavior. This smoothed signal is used as the primary input for the forecasting models so that the models can focus on the more stable and learnable component of the traffic series.

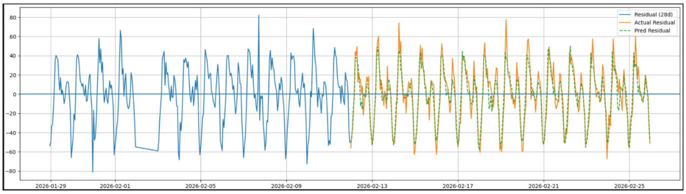

Figure 10. Residual Component After Structural Smoothing

Figure 10 presents the residual component obtained after subtracting the smoothed structural signal from the original throughput series. The residual represents irregular fluctuations that are not captured by the structural component. In the proposed framework, this residual is modeled using LightGBM to improve the final forecast by correcting remaining nonlinear and irregular errors.

Performance is measured using WMAPE, MAE, and RMSE over four operational forecasting horizons: 1-day, 3-day, 7-day, and 14-day. The overall summary results show that CatBoost + Residual LightGBM provides the most consistent forecasting performance across all horizons. Averaged across the full dataset, CatBoost achieves WMAPE = 13.63% for 1-day forecasting, 14.98% for 3-day forecasting, 15.77% for 7-day forecasting, and 15.80% for 14-day forecasting. It also produces the lowest overall MAE and RMSE among the evaluated models for all four horizons. By comparison, LightGBM + Residual LightGBM is the second-best performer and remains highly competitive, while MLP + Residual LightGBM shows moderate accuracy. SARIMA + Residual LightGBM consistently produces the highest error, particularly for medium and longer forecasting horizons.

Table 8. General Forecast Performance Horizons Result

Hybri Model Forecasting	Link	1 Day			3 Days			7 Days			14 Days		
		WMAPE (%)	MAE (Mbps)	RMSE (Mbps)	WMAPE (%)	MAE (Mbps)	RMSE (Mbps)	WMAPE (%)	MAE (Mbps)	RMSE (Mbps)	WMAPE (%)	MAE (Mbps)	RMSE (Mbps)
CatBoost+LightGBM	110	13.6	16.7	22.2	14.7	18.0	24.6	15.8	19.5	27.0	15.8	20.0	27.1
LightGBM+LighGBM	110	14.2	17.3	22.5	15.1	18.4	24.9	16.0	19.8	27.4	16.1	20.5	28.0
MLP+LightGBM	110	15.5	20.5	26.2	16.6	21.9	28.9	17.6	23.3	31.0	17.5	23.2	30.7
SARIMA+LightGBM	110	16.3	20.1	25.7	18.3	22.8	29.7	19.8	24.8	32.4	22.6	28.4	36.1

Table 8 presents the overall forecasting performance across all 110 microwave links. The results are aggregated over the full dataset and therefore provide the main benchmark for comparing the evaluated hybrid forecasting configurations across 1-day, 3-day, 7-day, and 14-day horizons. Overall, CatBoost + Residual LightGBM achieves the most consistent performance, followed by LightGBM + Residual LightGBM, while SARIMA produces higher errors, particularly for longer forecasting horizons.

These results indicate that, within the proposed smoothing–residual framework, boosting-based primary forecasting is better suited for capturing the structural behavior of microwave throughput than the statistical SARIMA approach. Although

SARIMA remains useful as a classical baseline, the observed microwave traffic patterns in this dataset appear to be better represented by nonlinear learning models, especially when heterogeneous link behavior and residual irregularity are present.

Figures 11–14, the actual curve represents the observed hourly throughput, while the forecast curve represents the predicted throughput generated by each hybrid forecasting configuration over the 14-day horizon.

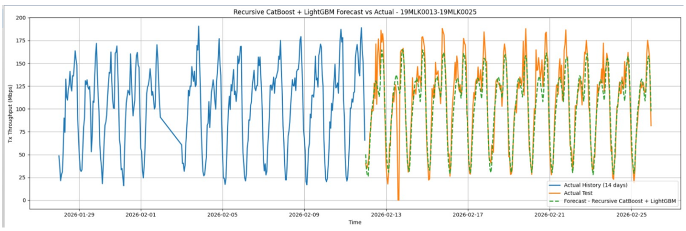

Figure 11. Two-Week Forecasting Result Using CatBoost + Residual LightGBM

Figure 11 shows 14-day forecasting result produced by CatBoost combined with LightGBM residual correction. The forecast curve follows the actual throughput pattern more consistently, indicating that CatBoost is effective in learning the smoothed structural behavior of the link. This result supports the overall finding that CatBoost + Residual LightGBM provides the most stable performance across forecasting horizons.

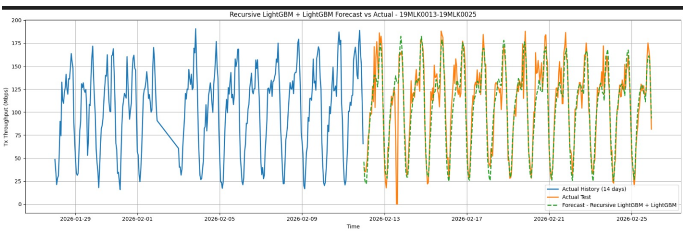

Figure 12. Two-Week Forecasting Result Using LightGBM + Residual LightGBM

Figure 12 presents the two-week forecasting result using LightGBM as the primary model with LightGBM residual correction. The model is able to capture the general traffic direction and performs competitively with CatBoost. However, compared with CatBoost, LightGBM may show slightly less stable tracking in some local fluctuations, which explains why it becomes the second-best model in the overall comparison.

Figure 13 shows the two-week forecasting result using MLP with LightGBM residual correction. The MLP model is able to learn nonlinear relationships, but its prediction may be more sensitive to data variability, training configuration, and short-term fluctuation. This explains why MLP produces moderate forecasting performance compared with the boosting-based models.

Figure 14 presents the two-week forecasting result using SARIMA with LightGBM residual correction. SARIMA captures the general temporal structure but is less flexible in representing nonlinear and heterogeneous link-level traffic behavior. As a

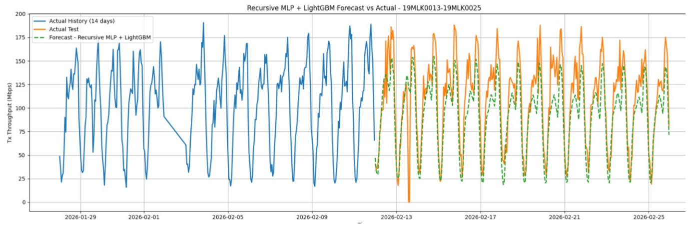

Figure 13. Two-Week Forecasting Result Using MLP + Residual LightGBM

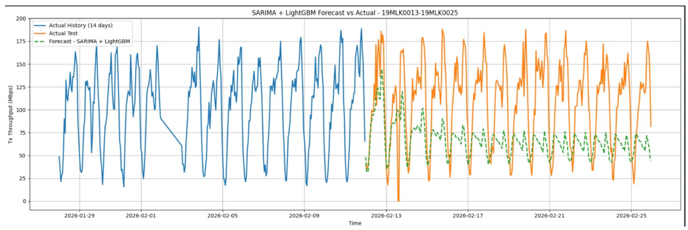

Figure 14. Two-Week Forecasting Result Using SARIMA + Residual LightGBM

result, its forecast tends to be less accurate than the boosting-based models, especially when the throughput series contains irregular fluctuation and local pattern changes.

4.3 Regional Forecast Performance Comparison

To evaluate the robustness of the proposed framework under different operational environments, forecasting performance is analyzed separately for the two study regions, namely Central Java (CJ) and Bali-Lombok (BL). These regions differ in traffic composition, operational context, and link-level temporal characteristics, making regional comparison important for assessing the generalizability of the proposed forecasting framework. The summary results show that the same overall model ranking remains visible across both regions, where CatBoost + Residual LightGBM consistently provides the strongest performance, followed by LightGBM + Residual LightGBM, while MLP + Residual LightGBM remains moderate and SARIMA + Residual LightGBM produces the largest error.

In Central Java, the forecasting framework benefits from a larger number of observed links and a traffic environment that still retains identifiable recurring demand structure for many links. Under this setting, CatBoost demonstrates strong performance in learning the smoothed structural signal, while the residual LightGBM layer improves the final forecast by correcting remaining nonlinear deviations. The results indicate that the hybrid framework remains stable even when the region contains a mixture of relatively regular links and more volatile operational traffic behavior. This suggests that the proposed approach is suitable for mature microwave environments where traffic growth is not fully uniform across all links.

In Bali-Lombok, the summary results also show that CatBoost remains the strongest primary forecasting model within the hybrid framework. However, the regional pattern suggests a different traffic environment, where some links may be influenced by more heterogeneous short-term fluctuation. Even under such conditions, the ranking between models remains broadly consistent, which indicates that the proposed smoothing-residual architecture is sufficiently flexible to handle regional variation. This is an important finding because it shows that the forecasting framework is not only effective in one dominant network environment, but can also maintain useful predictive stability across different traffic settings.

Table 9. Region Forecast Performance Horizons Result

Region	Hybri Model Forecasting	Link	1 Day			3 Days			7 Days			14 Days		
			WMAPE (%)	MAE (Mbps)	RMSE (Mbps)	WMAPE (%)	MAE (Mbps)	RMSE (Mbps)	WMAPE (%)	MAE (Mbps)	RMSE (Mbps)	WMAPE (%)	MAE (Mbps)	RMSE (Mbps)
Bali-Lombok	CatBoost+LightGBM	42	14.20	18.7	24.7	13.7	19.2	25.8	13.4	18.0	24.8	14.1	19.7	25.8
Bali-Lombok	LightGBM+LighGBM	42	14.30	18.7	24.7	15.9	22.7	29.7	13.9	18.9	25.8	14.2	20.4	27.0
Bali-Lombok	MLP+LightGBM	42	15.90	20.8	26.9	22.0	30.9	38.3	16.6	23.8	30.8	22.1	29.8	37.2
Bali-Lombok	SARIMA+LightGBM	42	14.90	21.4	28.0	14.6	22.2	28.8	15.8	24.3	31.4	17.1	26.2	33.2
Central Java	CatBoost+LightGBM	68	13.30	15.4	20.7	15.8	18.9	24.5	17.3	20.5	28.3	16.9	20.2	27.9
Central Java	LightGBM+LighGBM	68	14.10	16.4	21.2	15.7	20.0	25.5	17.3	20.4	28.3	17.2	20.6	28.5
Central Java	MLP+LightGBM	68	16.60	19.7	25.0	22.9	29.3	35.6	21.8	25.4	33.4	23.0	27.5	35.4
Central Java	SARIMA+LightGBM	68	15.90	20.0	25.1	19.9	25.0	31.5	18.7	22.6	30.8	17.7	21.4	29.2

Table 9 compares the forecasting performance between Central Java and Bali-Lombok. The results show that the overall model ranking remains broadly consistent across both regions, with CatBoost + Residual LightGBM providing the strongest performance in most horizons.

From a practical perspective, the regional comparison confirms that the main benefit of the proposed framework is not tied to a single geographical traffic profile. Instead, its strength lies in its ability to preserve stable forecasting performance across multiple operating conditions while still maintaining the same hybrid logic: structural forecasting on the smoothed series and nonlinear residual correction using LightGBM. This regional consistency strengthens the operational relevance of the proposed framework for broader microwave planning deployment.

4.4 Forecast Performance by Traffic Cluster

Beyond regional comparison, the forecasting results are also interpreted according to the four traffic-pattern clusters obtained during the clustering stage, namely Cluster-0 (Healthy Seasonality), Cluster-1 (Midday Sink), Cluster-2 (Broken Recovery), and Cluster-3 (Shifted Seasonal Pattern). Since these clusters represent different temporal morphology of link behavior, cluster-wise analysis is useful for understanding how traffic regularity and volatility affect forecasting difficulty. The uploaded summary shows that performance differences across clusters are meaningful, and that model behavior is not identical under all traffic conditions.

For Cluster-0, forecasting performance is generally better and more stable because the links in this group exhibit stronger temporal regularity and healthier seasonality. In this type of traffic environment, the dominant structural component can be captured more effectively by the primary forecasting model after smoothing. As a result, both CatBoost and LightGBM perform strongly, while residual correction mainly acts as a refinement layer rather than a major compensating mechanism. This indicates that links with clearer seasonality are naturally more favorable for multi-horizon forecasting.

For Cluster-1, the forecasting framework still performs well, although the presence of moderate volatility makes the task more challenging than Cluster-0. Nevertheless, because Cluster-1 retains identifiable daily structure, the proposed hybrid design remains effective in separating stable traffic behavior from residual fluctuation. In this cluster, CatBoost continues to show superior performance, which is consistent with the expectation that a robust boosting model can learn moderately complex structural behavior better than purely statistical alternatives.

For Cluster-2, forecasting becomes more demanding due to higher volatility and weaker seasonality. Even so, the summary results indicate that the CatBoost-based hybrid framework remains the strongest configuration overall. This suggests that although Cluster-2 is operationally more unstable than Cluster-1, the smoothing stage successfully extracts a structural signal that is still learnable, while the residual LightGBM layer helps absorb irregular error components that would otherwise degrade forecast quality. In other words, the hybrid architecture is particularly useful under this type of partially disrupted traffic behavior.

For Cluster-3, forecasting difficulty is the highest among all groups. This is expected because links in this cluster exhibit very high volatility, unstable shifting patterns, and weaker repeatability of temporal structure. Under these conditions, all models experience larger forecast error, and the remaining uncertainty after smoothing is also more difficult to correct. Operationally, this means that Cluster-3 links should be interpreted not only as a difficult forecasting group, but also as a group requiring closer operational attention, validation, and potentially earlier preventive review. Overall, the cluster-wise results confirm that forecasting difficulty is strongly influenced by traffic-pattern morphology, and that the proposed hybrid design provides its greatest value when structural regularity is present but incomplete.

Table 10 presents the forecasting performance across the four traffic-pattern clusters. The results show that forecasting accuracy is influenced by the temporal characteristics of each cluster. Links with stable seasonality generally produce more reliable forecasts, while links with irregular recovery or shifting seasonal behavior are more difficult to predict. This confirms that traffic-pattern morphology plays an important role in determining forecasting difficulty and model robustness.

4.5 Threshold-Crossing Accuracy and Delivery-Aware Planning

After evaluating the forecasting performance across the overall dataset, regions, and traffic clusters, this section demonstrates how the proposed framework can be translated into operational planning decisions. The analysis focuses on a representative high-risk microwave link to examine whether the forecast can accurately estimate when utilization approaches or crosses the 80% threshold. This threshold-crossing information is then used as the basis for delivery-aware planning, including crossing-time error evaluation and PO release scheduling.

Table 10. Cluster Forecast Performance Horizons Result

Cluster	Hybri Model Forecasting	Link	1 Day			3 Days			7 Days			14 Days		
			WMAPE (%)	MAE (Mbps	RMSE (Mbps)	WMAPE (%)	MAE (Mbps)	RMSE (Mbps)	WMAP E	MAE (Mbps)	RMSE (Mbps)	WMAPE (%)	MAE (Mbps)	RMSE (Mbps
0	CatBoost+LightGBM	31	13.4	15.25	19.43	15	17.16	22.18	14.3	16.27	21.96	17	19.58	25.38
0	LightGBM+LighGBM	31	13.6	15.23	19.27	15.1	18.49	23.16	15.1	17.23	23.12	17	19.67	25.63
0	MLP+LightGBM	31	15.5	16.92	20.76	21.4	25.12	29.55	17.7	20.03	26.27	26.1	29.27	35.75
0	SARIMA+LightGBM	31	13.9	15.57	19.68	17.3	20.26	25.45	17.6	20.69	27.15	18.8	22.57	28.97
1	CatBoost+LightGBM	38	13.7	15.63	22.05	15	17.35	23.36	16.8	18.42	25.39	15.4	17.57	24.27
1	LightGBM+LighGBM	38	14.3	16.23	22.33	15.9	19.71	26.83	17.1	18.81	25.74	15.6	17.85	24.54
1	MLP+LightGBM	38	16.2	18.67	24.69	22.4	27.72	35.14	19.1	20.91	27.87	20.3	23.27	30.08
1	SARIMA+LightGBM	38	16.9	20.75	27.77	17.1	19.77	26.11	18.2	21.12	28.42	16.7	19.07	25.83
2	CatBoost+LightGBM	31	13.1	17.49	22.97	14.5	19.16	25.17	16.3	22.69	31.48	14.9	20.62	29.04
2	LightGBM+LighGBM	31	14.3	19.02	24.03	15.9	23.1	28.87	16	22.22	31.4	15.4	21.09	29.63
2	MLP+LightGBM	31	17	23.04	29.03	23.5	34.21	41.32	20.8	28.4	37.24	21	29.01	37.98
2	SARIMA+LightGBM	31	14.8	21.05	26.18	19	28	35.05	16.6	23.14	32.12	15.9	22.27	30.95
3	CatBoost+LightGBM	10	15.8	22.3	28.84	16.6	30.72	39.77	14.8	24.07	34.63	16	28.41	37.57
3	LightGBM+LighGBM	10	15.5	22.3	28.57	17.2	27.07	34.33	14.5	24.34	34.06	16.9	31.71	42.89
3	MLP+LightGBM	10	17.1	26.62	34.68	23.6	39.51	49.36	26.1	42.72	53.51	25.8	43.15	53.75
3	SARIMA+LightGBM	10	17.7	33.36	40.77	19.3	38.52	48.11	18.3	39.69	49.69	21	44.02	54.03

By using the best-performing model, CatBoost + Residual LightGBM, the forecasted throughput is converted into utilization relative to the configured link capacity. This allows the model output to be interpreted not only as a numerical prediction, but also as a practical planning signal for identifying congestion risk, determining upgrade urgency, and estimating the latest feasible procurement window before persistent congestion occurs.

4.5.1 Crossing Time Error

To demonstrate the operational value of the proposed forecasting framework, this study examines the threshold-crossing behavior of one representative high-risk link, 14DMK0004-14SMG0428, which belongs to Cluster-0, is located in Central Java, and is categorized as High in the congestion-level overview. This case is selected to illustrate how multi-horizon forecasting can support practical planning decisions when a link operates near or above the 80% utilization threshold.

This case does not represent an initial threshold-crossing event. The selected link had already exhibited utilization behavior near or above the 80% threshold before the forecasting window. Accordingly, the purpose of this case study is not to demonstrate initial congestion discovery, but rather to evaluate how accurately the proposed framework can track and estimate subsequent threshold-crossing behavior for a high-risk link operating close to the congestion boundary. Under this setting, threshold-crossing accuracy remains operationally meaningful because the planning objective is to prevent persistent or recurring congestion through timely upgrade preparation.

Figure 15 presents the throughput forecast and the corresponding utilization trajectory of link 14DMK0004–14SMG0428. The lower panel shows the relationship between the forecast utilization path and the 80% threshold, allowing the predicted threshold-crossing timing to be identified visually. This analysis is important because the practical objective is not only to minimize numerical forecasting error, but also to determine whether a high-risk link is likely to remain in, return to, or move further into the congestion-planning zone that requires earlier corrective action.

Table 11. Forecasting Performance of Representative High-Risk Link 14DMK0004-14SMG0428

Link_id	Cluster	Region	Hybri Model Forecasting	1 Day			3 Days			7 Days			14 Days		
				WMAPE (%)	MAE (Mbps)	RMSE (Mbps)	WMAPE (%)	MAE (Mbps)	RMSE (Mbps)	WMAPE (%)	MAE (Mbps)	RMSE (Mbps)	WMAPE (%)	MAE (Mbps)	RMSE (Mbps)
14DMK0004-14SMG0428	0	CJ	CatBoost+LightGBM	10.0	11.6	16.7	11.7	13.6	18.5	13.4	17.0	26.5	16.4	20.9	28.9
14DMK0004-14SMG0428	0	CJ	LightGBM+LighGBM	11.8	13.7	18.4	13.1	16.6	22.1	14.9	18.8	28.2	18.2	23.1	31.0
14DMK0004-14SMG0428	0	CJ	MLP+LightGBM	15.0	17.4	21.9	20.7	25.9	31.1	15.9	20.1	27.2	17.3	22.0	29.5
14DMK0004-14SMG0428	0	CJ	SARIMA+LightGBM	12.6	14.7	18.4	36.3	42.4	50.3	16.3	20.6	29.0	20.7	26.3	33.9

Table 11 presents the forecasting performance of the representative high-risk link 14DMK0004–14SMG0428 across multiple horizons. The results show that CatBoost + Residual LightGBM provides the lowest error and the most stable performance for this link. This finding supports its use in the subsequent threshold-crossing and PO release analysis, where reliable forecasting is required to estimate congestion timing and support proactive upgrade planning.

For this representative link, CatBoost provides the best forecasting performance across all validated horizons. At the 1-day horizon, CatBoost achieves WMAPE 10.0%, MAE 11.603, and RMSE 16.662. At the 3-day horizon, it records WMAPE 11.7%, MAE 13.610, and RMSE 18.48. At the 7-day horizon, CatBoost remains the strongest model with WMAPE 13.4%, MAE 16.96, and RMSE 26.53. At the 14-day horizon, it still outperforms the other models with WMAPE 16.4%, MAE 20.88, and RMSE 28.90. These results confirm that CatBoost provides the most stable forecast trajectory for this high-risk link across all evaluated forecasting windows.

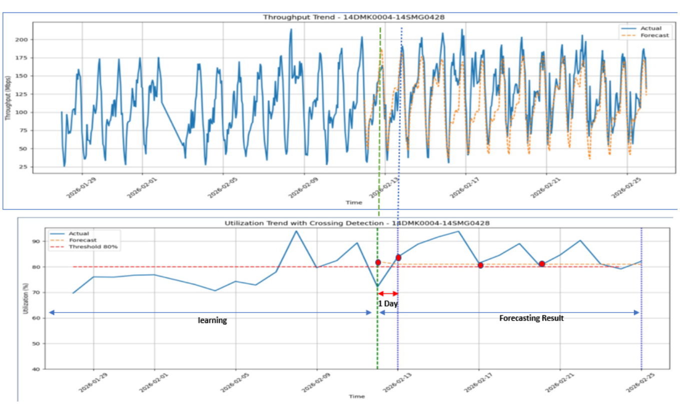

Figure 15. Threshold-Crossing Time Error for Link 14DMK0004–14SMG0428

Figure 15 illustrates the threshold-crossing behavior of the representative high-risk link 14DMK0004–14SMG0428. The figure compares the predicted utilization trajectory with the actual crossing event around the 80% utilization threshold. The result shows a crossing-time error of 1 day, indicating that the proposed framework can provide a reliable early-warning signal for planning PO release and capacity upgrade actions.

4.5.2 Planning PO Release

From a practical planning perspective, threshold-crossing accuracy is more actionable than throughput prediction alone because the key operational question is not only how much traffic will increase, but when corrective action must begin. In this context, the predicted threshold-crossing date can be used as the reference point for backward planning of procurement and implementation activities. Therefore, PO release should be initiated before the predicted threshold-crossing date by considering the total delivery lead time, including procurement processing, material readiness, implementation, integration, and operational safety buffer. As a result, the forecast provides not only an indication that a link is approaching congestion, but also a practical basis for determining the latest feasible window for PO release so that the upgrade can be completed before the link remains persistently above the 80% utilization threshold.

For the representative link 14DMK0004–14SMG0428, the predicted threshold-crossing date is 12 February 2026, while the actual crossing occurs on 13 February 2026, resulting in a 1-day crossing-time error. Under the normal planning assumption, the total delivery lead time consists of 14 days for PO release to material readiness, 7 days for implementation and integration, and 3 days of safety buffer, giving a total of 24 days. Based on this assumption, the latest baseline PO release date is 19 January 2026.

$$t_{Po,i} = t_{cross} - (L_{Po} + L_{impl} + B_{uffer}) - E_{cross}$$

Where:

- t_{cross} Predicted crossing threshold 80% on 12 Feb 2026
- L_{Po} lead time PO release include material readiness time: assume 14 days
- L_{impl} Implementation, integration and RFS: 7 Days
- B_{uffer} safety buffer: 3 days
- E_{cross} Crossing-time error allowance: 1 Day
- $t_{Po,i}$ Expected PO released date

For the representative link 14DMK0004–14SMG0428, the planning assumptions are as follows:

$$L_{Total} = L_{Po} + L_{impl} + B_{uffer} = 14 + 7 + 3 = 24\ Days$$

After including the 1-day crossing-time error allowance, the safer PO release date becomes:

$$t_{Po,i} = 12\ Feb\ 2026 - 24 - 1$$
$$t_{Po,i} = 18\ Jan\ 2026$$

To account for the 1-day crossing-time error, a safer planning date can be set 1 day earlier, resulting in an adjusted PO release date of 18 January 2026. This means that, from an operational perspective, the PO should ideally be released between 18 and 19 January 2026, depending on whether the operator applies baseline planning or a more conservative buffer.

To provide a broader operational view, three lead-time scenarios can also be considered. Under the Fast scenario, the latest PO release date is 29 January 2026. Under the Normal scenario, it is 19 January 2026. Under the Conservative scenario, it becomes 7 January 2026. These results show that the same threshold-crossing forecast can be directly translated into different procurement deadlines depending on delivery assumptions, thereby supporting delivery-aware expansion planning and proactive congestion prevention.

Table 12. Procurement Planning Scenarios Based on Predicted Threshold-Crossing Date

Scenario	Predicted Crossing Date	Total Lead Time	Latest PO Release Date	Planning Interpretation
Fast	12-Feb-26	14 days	29-Jan-26	feasible with short procurement cycle
Normal	12-Feb-26	24 days	19-Jan-26	recommended baseline planning
Conservative	12-Feb-26	36 days	7-Jan-26	requires earlier decision

Table 12 translates the predicted threshold-crossing date into practical PO release scenarios. By applying different delivery lead-time assumptions, the same forecast can generate different procurement deadlines under fast, normal, and conservative planning conditions. This demonstrates how the proposed forecasting framework can be connected directly to delivery-aware planning, allowing operators to determine when procurement should be initiated to complete upgrades before congestion becomes persistent.

Overall, this simulation shows that the same threshold-crossing forecast can be directly translated into procurement timing recommendations depending on the assumed delivery lead time. This confirms that the proposed forecasting framework is not only useful for predicting utilization behavior, but also for supporting delivery-aware expansion planning and proactive congestion prevention in microwave transmission networks.

4.6 Practical Interpretation for Continuous Rolling Forecasting

Although the formal experimental validation in this study is limited to 1-day, 3-day, 7-day, and 14-days horizons, the operational interpretation of the framework is broader. The proposed method is not intended to function as a single one-time predictor for a very long horizon in one execution. Instead, it is designed to be used in a continuous rolling-update mode, where the model is rerun periodically as new hourly NMS observations become available. This operational interpretation is fully consistent with the forecasting logic and planning orientation described in the earlier sections.

Under this rolling setup, short- and medium-horizon forecasts become an ongoing early-warning engine rather than isolated predictions. Each new forecast update refines the expected utilization path, revises the estimated threshold-crossing risk, and improves the timing basis for upgrade decisions. In this sense, the value of the framework does not depend on claiming direct validated accuracy for a very long single-run horizon. Instead, its value lies in continuously monitoring whether a link is progressively moving toward the 80% threshold, and whether the remaining lead time is still sufficient for corrective action.

This interpretation is particularly important for practical transmission planning because network traffic is inherently dynamic. Forecasts generated today may need to be revised tomorrow as new traffic behavior emerges, especially for links in higher-volatility clusters. A rolling-update mechanism is therefore more realistic and more operationally defensible than relying on a single fixed long-horizon forecast. It also aligns better with real engineering workflows, where planning decisions are repeatedly updated in response to current utilization, procurement status, and field readiness.

Accordingly, the practical contribution of the proposed framework can be understood as follows: it provides a repeatable multi-horizon forecasting mechanism that supports continuous congestion-risk monitoring, iterative upgrade prioritization, and delivery-aware zero-congestion planning over longer operational windows. This makes the framework suitable not only for experimental forecasting evaluation, but also for routine transmission planning deployment in environments where data are updated continuously and decisions must be revised over time.

5. Conclusion

This study proposes a delivery-aware hybrid smoothing–residual forecasting framework for link-level microwave throughput prediction under physical-layer capacity constraints. The framework is designed not only to improve forecasting accuracy, but also to support proactive transmission planning by identifying whether and when a link is moving toward the operational 80% utilization threshold. To achieve this, the observed throughput series is separated into a smoothed structural component and a residual component, where the structural signal is forecast by the primary model and the remaining error is corrected using LightGBM as a fixed residual learner.

The experimental evaluation uses hourly Network Management System (NMS) data from 110 microwave links, consisting of 68 links in Central Java and 42 links in Bali–Lombok, observed from 1 September 2025 to 26 February 2026. Forecasting performance is assessed over 1-day, 3-day, 7-day, and 14-days horizons using WMAPE, MAE, and RMSE. The results show that the proposed hybrid architecture is effective for handling heterogeneous microwave traffic behavior. Among the evaluated configurations, CatBoost + Residual LightGBM provides the best and most consistent forecasting performance across all validated horizons, achieving the strongest result at the 1-day horizon with WMAPE 13.6%, MAE 16.65, and RMSE 22.19, while remaining the most stable configuration for longer horizons.

From an operational perspective, the main contribution of this study lies in translating forecast outputs into threshold-oriented planning signals. By relating the predicted utilization trajectory to the 80% congestion-planning threshold, the framework supports early warning, upgrade prioritization, and PO release planning under delivery lead-time constraints. The case study on link 14DMK0004–14SMG0428 further shows that the proposed method can estimate threshold-crossing timing with a 1-day crossing-time error at the 14-day horizon, indicating that the framework is sufficiently reliable for near-term expansion planning and proactive congestion prevention.

A limitation of this study is that formal model validation is conducted only up to a 14-day horizon using approximately six months of hourly observations. Therefore, this study does not claim direct validated forecasting accuracy for very long horizons in a single run. Instead, the proposed method is intended to be deployed in a continuous rolling-update mode, where forecasts are regenerated periodically as new hourly data become available. Under this operational setting, repeated short- and medium-horizon forecasts can provide continuous congestion-risk monitoring and support longer-term transmission planning through iterative forecast refinement.

Future research may extend this framework in several directions. First, more advanced deep learning models such as Long Short-Term Memory (LSTM) may be explored to capture longer temporal dependency and more complex nonlinear traffic behavior, particularly for highly volatile links. However, such models would require higher computational resources, more intensive parameter tuning, and longer training duration than the current boosting-based approach. Second, future studies may incorporate longer historical observation windows and additional exogenous variables, such as weather effects, special-event traffic, or region-specific operational factors. Third, tighter integration between threshold-crossing prediction, procurement timing, and automated transmission planning dashboards may further enhance the practical value of the proposed framework for operational deployment.

References

[1] ITU-R. *Recommendation ITU-R P.530-19: Propagation Data and Prediction Methods Required for the Design of Terrestrial Line-of-Sight Systems.* Tech. rep. International Telecommunication Union, Sept. 2025.

[2] Telecom Regulatory Authority of India. *Recommendations on Assignment of the Microwave Spectrum for Backhaul.* Tech. rep. Dec. 2025.

[3] Huawei Technologies. *Microwave Industry White Paper 2025.* Tech. rep. Nov. 2025.

[4] G. A. Christian, I. P. Wijaya, and R. F. Sari. “Network Traffic Prediction of Mobile Backhaul Capacity Using Time Series Forecasting”. In: *Proceedings of the International Seminar on Intelligent Technology and Its Applications (ISITIA).* 2021. DOI: 10.1109/ISITIA52817.2021.9502256.

[5] J. Kim et al. “A Deep Dive into AI-Based Network Traffic Prediction”. In: *Applied Sciences* (2025). DOI: 10.3390/app16010367.

[6] Rob J. Hyndman and George Athanasopoulos. *Forecasting: Principles and Practice.* 3rd ed. OTexts, 2021. URL: https://otexts.com/fpp3/.

[7] George E. P. Box et al. *Time Series Analysis: Forecasting and Control.* 5th ed. Wiley, 2015.

[8] M. Abbasi et al. “Deep Learning for Network Traffic Monitoring and Analysis: A Review”. In: *Computer Communications* (2021). DOI: 10.1016/j.comcom.2021.01.021.

[9] Guolin Ke et al. “LightGBM: A Highly Efficient Gradient Boosting Decision Tree”. In: *Advances in Neural Information Processing Systems*. 2017.

[10] Tianqi Chen and Carlos Guestrin. “XGBoost: A Scalable Tree Boosting System”. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2016. DOI: 10.1145/2939672.2939785.

[11] Ericsson. *Microwave Outlook 2025: Near Fiber Split, AI and 2x Capacity*. Tech. rep. Oct. 2025.

[12] Huawei Technologies. *Microwave Industry White Paper 2024*. Tech. rep. 2024.

[13] Liudmila Prokhorenkova et al. “CatBoost: Unbiased Boosting With Categorical Features”. In: *Advances in Neural Information Processing Systems*. 2018.

[14] Simon Haykin. *Neural Networks and Learning Machines*. 3rd ed. Pearson, 2009.

[15] G. O. Ferreira et al. “Forecasting Network Traffic: A Survey and Tutorial With Open-Source Comparative Evaluation”. In: *IEEE Access* 11 (2023). DOI: 10.1109/ACCESS.2023.3236261.

[16] M. Di Mauro et al. “Hybrid Learning Strategies for Multivariate Time Series Cellular Traffic Forecasting”. In: *Computer Networks* (2024). DOI: 10.1016/j.comnet.2024.110286.

[17] G. P. Zhang. “Time Series Forecasting Using a Hybrid ARIMA and Neural Network Model”. In: *Neurocomputing* 50 (2003), pp. 159–175.

[18] Wu et al. “Deep Learning for Mobile Traffic Forecasting: A Survey”. In: *IEEE Transactions on Network and Service Management* (2019). DOI: 10.1109/TNSM.2019.2929152.

[19] S. Li et al. “Deep Learning Based Prediction of Traffic Peaks in Mobile Networks”. In: *Computer Networks* (2024). DOI: 10.1016/j.comnet.2023.110167.

[20] X. Liu et al. “A Meta-Learning-Based Framework for Cellular Traffic Forecasting”. In: *Applied Sciences* (2025). DOI: 10.3390/app152111616.

[21] Y. Peng et al. “Network Traffic Prediction With Attention-Based Spatial-Temporal Modeling”. In: *Computer Networks* (2024). DOI: 10.1016/j.comnet.2024.110296.

[22] Z. Jin et al. “A Mobility-Aware Network Traffic Prediction Model”. In: *Computer Networks* (2023). DOI: 10.1016/j.comnet.2023.109981.

[23] E. Lykakis et al. “Data Traffic Prediction for 5G and Beyond”. In: *Electronics* (2025). DOI: 10.3390/electronics14234611.

[24] S. Saha et al. “Overcoming Data Limitations in Internet Traffic Forecasting”. In: *Computer Communications* (2025). DOI: 10.1016/j.comcom.2025.108280.