B2B Digital Dependency: A Thinking Artifact

Yu-Jou Ting · Independent Follow-up · 2025

CONTEXT

During an on-site strategy exercise at SFR, teams were given limited time to respond to a live B2B segmentation problem. Due to time constraints, most outputs remained surface-level.

This document is an independent follow-up, built afterward without access to SFR internal data, and never submitted.

What it demonstrates: problem reframing under ambiguity, hypothesis-driven segmentation design, and anticipation of operational constraints.

SCOPE NOTE

This work focuses on strategic framing and testable hypotheses, not full implementation planning.

EXECUTIVE FRAMING

The focus here is on what already exists but has not been framed correctly. The objective is to transition from reactive coverage (serving based on visible size) to predictive partnership (serving based on digital criticality).

The proposals are designed as No-Regret moves. They rely on existing datasets (DTM library) and non-intrusive external signals. Execution is envisioned as a series of low-stakes POC sprints to validate hypotheses before any systemic overhaul.

IF IMPLEMENTED

Decision relevance if validated:

If validated through targeted POCs, this framework would enable:

Earlier identification of high-growth B2B clients before revenue signals appear in usage data.

More efficient allocation of SLA resources based on business criticality rather than headcount size.

Increased retention among multi-site and data-dependent clients through targeted service design.

Initial validation would measure:

Retention uplift among segmented high-dependency clients.

Upsell conversion for premium products (SD-WAN, premium SLA).

Churn reduction in high-dependency segments identified through digital intensity scoring.

STRUCTURAL DIAGNOSIS

The Core Tension SFR Is Facing

SFR faces a structural tension between local optimization and strategic momentum. Each unit is rationally managed. The company as a whole cannot advance. This is a framing problem, not a performance problem.

Three observations underpin this:

1. SFR has long been positioned in B2B as a low-cost alternative rather than a professional partner. The root issue is the absence of trust data. Clients do not believe SFR can deliver high-quality SLA maintenance. The goal is to convert existing, invisible reliability into a visible commercial advantage.

2. Single-site companies are cost-sensitive and will switch for cheaper contracts. Multi-site companies need stability, will pay a premium for it, and once a pricing framework is established across their system, they rarely migrate.

3. Given the high dimensionality and sparsity of telecom data, Feature Importance Ranking should precede model development. The non-linear relationship between Recovery Time and Churn should be tested before defaulting to correlation analysis. A small-scale POC (5,000 samples from DTM_client_features_library) is the appropriate starting point.

THREE STRATEGIC DIRECTIONS

Three strategic directions (from exploratory to data-driven)

A. Growth Intent Mining

If a company is hiring Cloud Engineers or Data Scientists, its future telecom demand will increase substantially regardless of current headcount. Monitoring public hiring signals provides a non-intrusive, legally defensible indicator of digital growth intent.

The goal is to reframe the telecom contract from a utility cost into a dependency-building growth tool. Early offers (cloud credits, security consultations) create switching costs that compound over time. Proactive Transparency Reports (real-time SLA dashboards shared with multi-site clients) shift brand perception from cheap to accountable.

> *Note: all external signal collection must use publicly available information and non-intrusive indicators, subject to legal and compliance validation.*

B. SLA Elasticity Segmentation

Clients who have historically generated compensation claims following outages have already revealed that their business depends on network stability. These clients are most likely to pay a premium for guaranteed SLA, even if small by headcount.

A grey-zone test in high Recovery Time demand areas (a Golden 1-Hour Repair add-on) can validate retention impact before any system change.

Critical: Premium SLA for small clients must be Logical VIP, not Physical VIP. Provide automatic 4G/5G backup dongles rather than faster repair crew dispatch. This achieves 99.9% uptime without competing with large-client repair priorities.

C. Digital Intensity Scoring

Replace headcount-based segmentation with a Digital Intensity metric: Total Data Usage (Invoice) divided by Workforce Size (Profile). Companies scoring above the 80th percentile are flagged as Data-Hungry Clients.

Workforce size data may be unreliable. If accuracy is low, the model risks collapse. Proxy indicators (active SIM count, or office square footage where available) should be tested in parallel.

A/B Test: select 100 Data-Hungry clients currently using only basic SIM plans. Route them to a dedicated account manager (not automated voice) for SD-WAN upsell. Measure conversion and retention delta.

ANTICIPATING OBJECTIONS

Known weaknesses in the proposed framework

On Digital Intensity: the data circularity trap

Challenge: the metric captures only clients who already use a lot of data. It misses clients whose growth is constrained by SFR own bandwidth limitations. This produces Selection Bias.

Fix: add Traffic Growth Velocity as a variable. A small company with flat headcount but 20% monthly traffic growth is a hidden high-value client that usage-volume metrics alone would not surface.

On LLM-Assisted Labeling: scalability and cost

Challenge: GPT-4 API costs are high, web scraping is slow, anti-scraping mechanisms exist, and a one-time 5,000-sample run does not constitute a production pipeline.

Fix: Tiered Labeling. Apply LLM deep-labeling only to high-traffic or high-change clients. Use keyword-rule filters for the rest.

On SLA Segmentation: operational contradiction

Challenge: during a region-wide outage, a repair manager must choose between a 500-person factory and a 10-person VIP consultancy. If the factory is served first, the Premium SLA becomes a liability.

Fix: addressed via Logical VIP (automatic backup hardware). The commitment is fulfilled without competing for physical repair resources.

On Single vs Multi-site: geographic ambiguity

Challenge: many companies have multiple connections at the same address. This appears as Multi-site in the data but is not. Upselling VPN to these clients will produce confusion and waste Sales capacity.

Fix: require Address Zip Code Count validation. Only flag a SIREN as Multi-site strategic target when it has two or more distinct postal codes.

WHAT THIS DOCUMENT IS AND IS NOT

An honest account

This is a working document, produced independently after real field exposure.

The value lies in how the problem is framed, the hypotheses proposed, and the decisions these enable.

The framing, segmentation logic, and objection anticipation are developed. The financial modeling and systemic validation remain open questions.

Certain directions, particularly around Feature-wise Modulation informed by technical exchange with the Numberly data science team, are promising but carry real operational risks: explainability gaps, noise-to-signal challenges, and maintenance cost. These risks are flagged, with fixes proposed where possible.

The recommendations are designed as No-Regret moves: low-commitment tests that generate useful signal regardless of outcome. Every POC is structured to answer: what do we confirm if it works, what do we rule out if it fails, and what is the regret cost to the organization either way.

STRATEGIC INTENT

The purpose of this work is to enable better decisions under uncertainty.

Yu-Jou Ting · 2025 · Independent analysis, not affiliated with SFR