Red Hat Summit

Connect

AI Chronicles

Notes & Lessons Learned from Real World Projects

Fatih E. Nar | Distinguished Chief Architect,

Red Hat CTO Office

The Quick Check

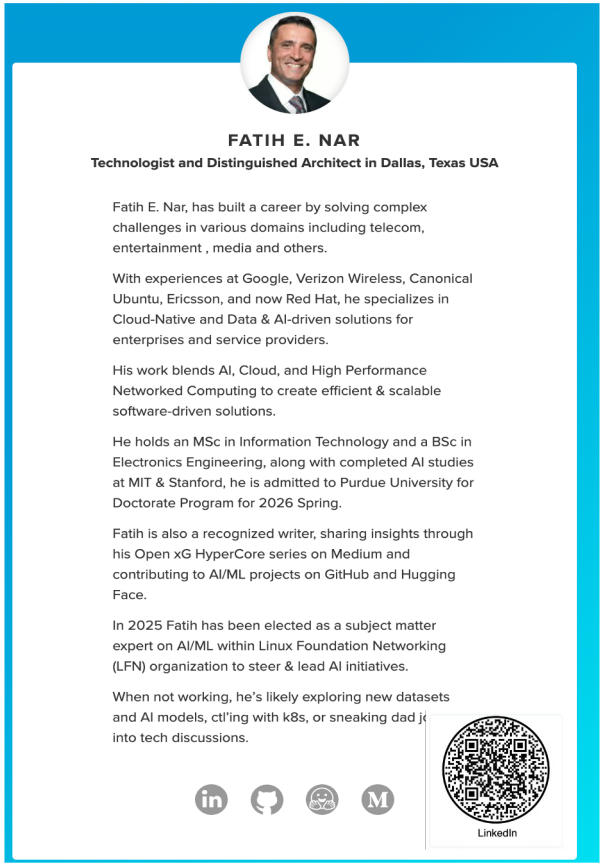

FATIH E. NAR

Technologist and Distinguished Architect in Dallas, Texas USA

Fatih E. Nar, has built a career by solving complex challenges in various domains including telecom, entertainment , media and others.

With experiences at Google, Verizon Wireless, Canonical Ubuntu, Ericsson, and now Red Hat, he specializes in Cloud-Native and Data & AI-driven solutions for enterprises and service providers.

His work blends AI, Cloud, and High Performance Networked Computing to create efficient & scalable software-driven solutions.

He holds an MSc in Information Technology and a BSc in Electronics Engineering, along with completed AI studies at MIT & Stanford, he is admitted to Purdue University for Doctorate Program for 2026 Spring.

Fatih is also a recognized writer, sharing insights through his Open xG HyperCore series on Medium and contributing to AI/ML projects on GitHub and Hugging Face.

In 2025 Fatih has been elected as a subject matter expert on AI/ML within Linux Foundation Networking (LFN) organization to steer & lead AI initiatives.

When not working, he's likely exploring new datasets and AI models, ctl'ing with k8s, or sneaking dad j into tech discussions.

LinkedIn

"No Gimmicks - No Hype Pumping True Applied AI for Enterprise"

20+ years in SWE & Telecom & Cloud Eng

Real Observability & AI implementations in production

Open source datasets, models, and code

Focus: **Business outcomes**, not just technical achievements

Ironed Framework

6-phase methodology adapted from data mining best practices

1. Business Understanding

Define objectives & AI goals

2. Data Exploration

Harvest, assess volume/velocity/variety

3. Information Engineering

Select, cleanse, ETL pipelines

4. AI Modeling

Build, test multiple techniques

5. QA & Tuning

Evaluate results, review process

6. Production

MLOps, observe, measure impact

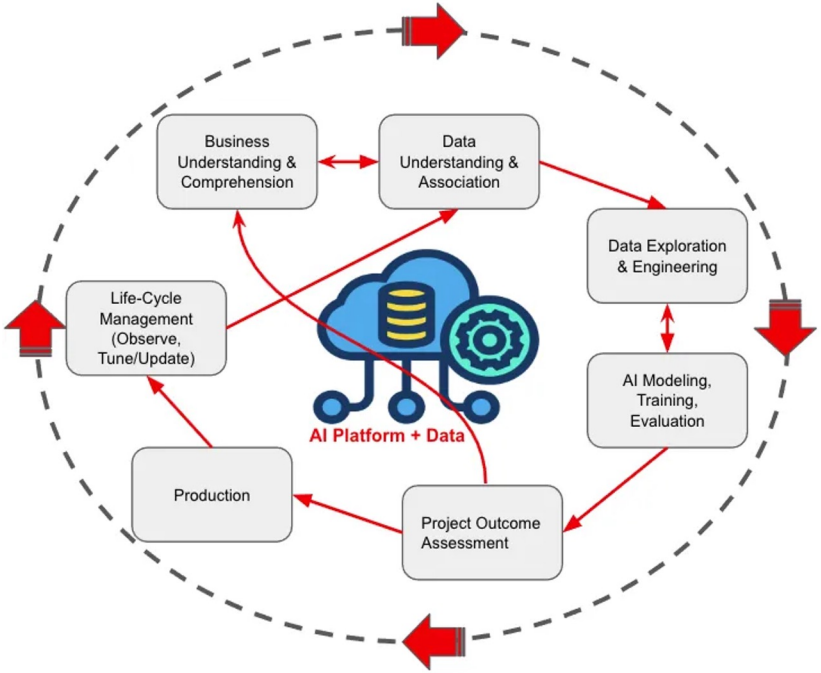

Without this discipline, AI projects fail to meet business objectives

Strategy #1: The Foundational Platform

The #1 Barrier to AI Success

81% of database workloads still on-premises

80% of organizations use multiple clouds

Data fragmented across systems = **Data silos**

Apps Distributed across clouds = **Platform silos**

- Creating unified AI governance framework is key

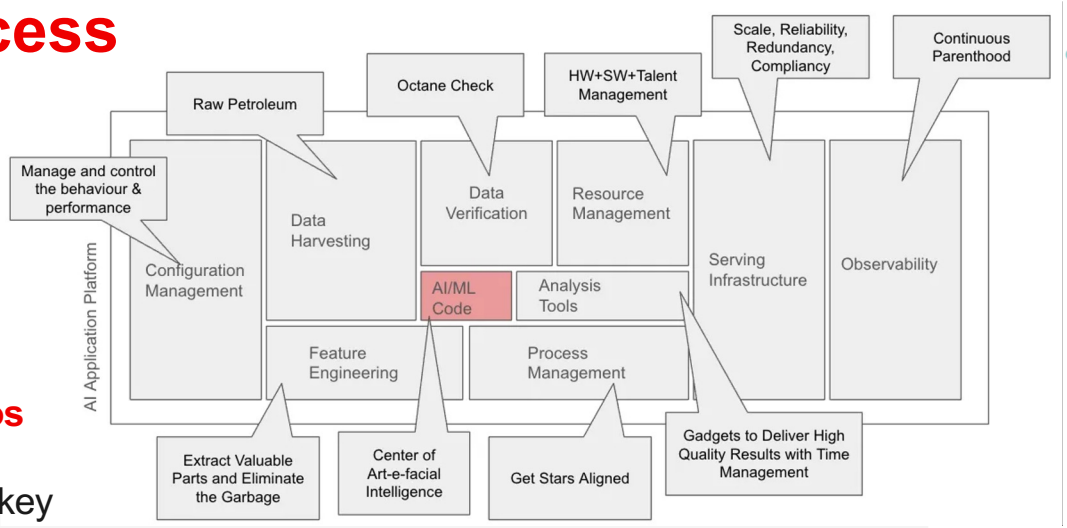

> ***"It's difficult to think about what my data is, where it is, and how I utilize all data types stored for 10-15 years."***
>
> — VP of Cloud Operations, US Tier 1 Service Provider

Strategy #2: Right Business Problems

Technology-Fits Where Business Needed To Thrive

Business Driven Architectures

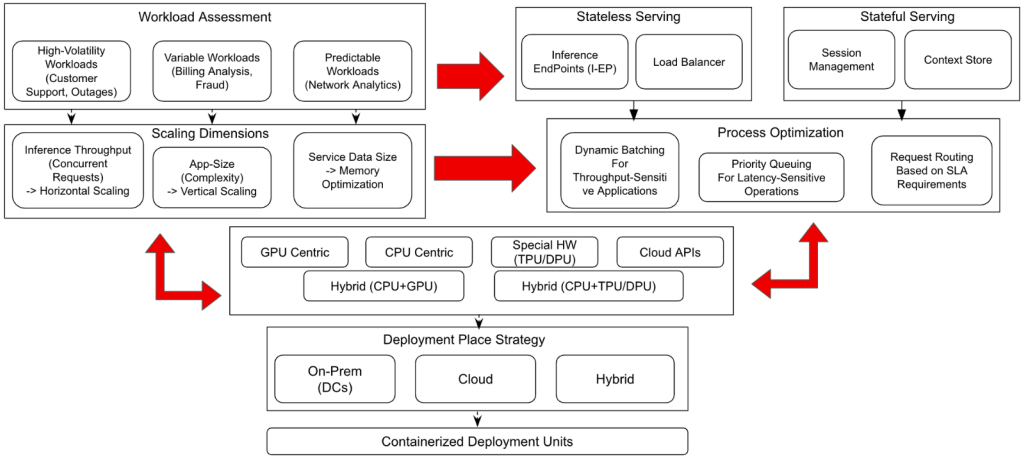

Wrong: "Let's use AI because it's cool"

Right: "Reduce customer churn by 15%"

Start with **business problem**, not technology!

Define **success metrics** upfront

Real Example: US Telco Device Retailer

Problem: Poor product search • Solution: Hybrid search with GenAI Recommendation Engine

Impact: 20+% improvement

Telco-AIX: 15+ Real Implementations

Open source AI use cases for telecommunications

💰 Revenue

Revenue Assurance, Fraud Detection, Churn Prediction

⚙️ Operations

Service Assurance, 5G RAN, Root Cause Analysis

🔒 Security

IoT Security, SecOps-AI, DDoS Detection

🌍 Sustainability

Energy Efficiency, Smart Grid, Carbon Optimization

🤖 Agentic

Intent Classification, Autonomous Networks

📊 Models

All datasets & models on HuggingFace

Various Telco AI Usecases & Experiments

- Readme
- MIT license
- Activity
- Custom properties
- 183 stars
- 32 watching
- 80 forks

github.com/open-experiments/Telco-AIX

Use Case #1: Revenue Assurance

The Business Problem

- Revenue leakage from billing errors
- Fraudulent transactions
- Manual audits costly & slow

Data Engineering

- Transaction logs & CDRs
- 20+ engineered features
- Dataset on HuggingFace

The Approach

- Balanced Random Forest
- Neural Network Transformers

Business Impact

85%+ accuracy

Cost savings: Reduced revenue leakage

	Call_Duration	Data_Usage	Sms_Count	Roaming_Indicator	MobileWallet_Use	Plan_Type	Cost	Cellular_Location_Distance	Last_Time_Pin_Used	Avg_Call_Duration	Avg_Data_Usage	Fraud
0	4.692681	539.146554	1	0	0	postpaid	71.603238	3.629675	46.011773	1.460109	336.312984	0
1	30.101214	247.225104	4	0	0	postpaid	-3.794503	3.654629	47.339394	30.817472	150.969590	0
2	13.167457	117.971674	3	0	0	postpaid	4.581320	2.506765	8.653167	13.554912	79.394244	0
3	9.129426	411.883231	4	0	0	postpaid	24.955166	0.098861	18.875911	7.990501	317.191998	1
4	1.696249	1134.432099	4	0	0	prepaid	43.517320	4.204570	30.711223	0.159457	1073.585260	0

Use Case #2: Service Assurance

The Business Problem

- Reactive network issue handling
- Customer satisfaction at risk
- High MTTR

Data Sources

- Platform metrics (Prometheus)
- Application telemetry (OTel)
- Network KPIs

The Approach

- Neural Networks (PyTorch)
- Predict latency & NPS
- OpenTelemetry data

Business Impact

40% MTTR reduction

Proactive detection • Customer retention

```
(app-root) (app-root) curl -X POST http://0.0.0.0:5000/predict -H "Content-Type: application
/json" -d '{
     "input_data": "Cell Availability: 70.0, MTTR: 4.0, Throughput: 50.0, Latency: 1000, Pac
ket Loss Rate: 0.96, Call Drop Rate: 0.97, Handover Success Rate: 77.0, Data Usage: 7000, Us
er Count: 9000, Signal Strength: -80, Jitter: 600, Connection Setup Success Rate: 83.0, Secu
rity Incidents: 5, Authentication Failures: 5, Temperature: 10, Humidity: 70, Weather: Rain,
 Issue Reported: Yes, City: New York, State: NY, Zip: 10001",
     "question": "What would be the Fault Occurrence Rate?"
   }'
{"answer":"37.073938081897"}
(app-root) (app-root)
```

Use Case #3: Churn Prediction

The Business Problem

- Customer acquisition cost >> retention
- Losing customers to competitors
- Need early warning system

The Approach

- Classification models
- Multi-feature analysis
- Behavioral patterns

Key Features

- Service usage patterns
- Support ticket frequency
- Payment behavior
- Contract details

Business Impact

82% accuracy

Revenue protection • Targeted retention

AI Model Selection: When To Use What

Traditional ML

Prediction, classification, forecasting. Higher accuracy, explainable.

GenAI/LLMs

Content generation, conversation, natural language tasks.

Rule-Based

Policy enforcement, compliance. Reliable, transparent.

Neural Networks

Complex patterns, regression, non-linear relationships.

Mixture of Experts

Multi-aspect problems. Combine strengths of multiple models.

Clustering

Segmentation, anomaly detection, unsupervised learning.

⚠️ **Don't replace everything with LLMs!**

Use the right tool. Blend Classic-AI with GenAI.

Strategy #3: Culture of Experimentation

Build Small, Focused Teams

<10 people per team — bigger = communication bottlenecks

"You^AI"

What Works

- Fail fast, iterate quickly
- Define "good" upfront
- Defragment: stakeholders in room
- Automate error detection

What Doesn't Work

- Analysis paralysis
- Perfection before deployment
- Siloed teams
- Technology-first thinking

Strategy #4 & #5: Security + Integration

Security & Governance

- AI adds compliance complexity (GDPR)
- Create unified governance framework
- Make AI part of enterprise fleet

Observability for Auditability

Challenge: Data science wants access; IT enforces compliance. Solution: Collaboration.

System Integration

Augment existing systems

- Legacy models are still useful
- Break down data silos

Classical ML + GenAI = **Best approach**

Reality: Without data integration, AI in silos can't solve enterprise problems.

The Hard Truths: What We Learned

Upskilling is HARD

Telco data scientists are rare! Train internal teams who speak your language.

GPU Economics

GB300 vs RTX A6000: 10x cost for marginal gains. Don't overspend before proving value.

Data Quality

Garbage in, garbage out. Clean data is 80% of the work. Most AI fails due to poor data.

When NOT to Use AI

If Expert Systems Work, Use them!

AI isn't always the answer. Sometimes Excel is sufficient till to a point.

AI Reset is Coming !

The Structural Shift

- Model leadership will rotate every 6-12 months **(no permanent winner)**
- Durable value lives in workflow control, not model superiority
- "Best model" strategy is breaking down -> **prepare for volatility with power-in of open models**.

What's Actually Changing

- **Distribution > Model capability** (who controls/dominates inference wins)
- Open Source Alliance(s) would shift power from Google/OpenAI/Microsoft/Anthropic to Sovereign AI.

Strategic Imperatives for 2026

- **Stop betting on one model vendor** -> architect for multi-model (set of different models) reality
- **Optimize for surfaces** -> where does user intent originate? (voice, email, slack)
- **Focus on workflows, not models** -> become the translator between leadership and tools
- **Build for model churn** -> your primary vendor could pivot, get acquired, or fall behind

Your Edge in the Reset

- **Own proprietary data + workflow knowledge**
- Control specific surfaces (spreadsheets, terminal, calendar)
- **Deliver stable systems from unstable models**
- **Master orchestration, not just prompting**

Top 5 Lessons Learned

- **AI Foundation First** : 80% of success is clean, unified data with consistent enabler platform. Fix data, get core platform ready before buying GPUs.
- **Business Problems, Not Tech** : Start with "What business problem we are solving?" not "How to use AI?"
- **Right Tool for the Job** : Traditional ML often beats GenAI for prediction. Don't replace all with LLMs.
- **Small Teams, Big Impact** : Teams <10 (Two Pizza Team) move faster. Defragment early.
- **Cost-Effective Infrastructure** : Most Desktop GPUs work for most cases. A100/H200/B200/ etc are overkill until you prove a value for +investment.

Red Hat Summit

Connect

Thank you

Join the Telco-AIX Movement


linkedin.com/company/red-hat


facebook.com/redhatinc


youtube.com/user/RedHatVideos


twitter.com/RedHat