The Importance of Context in AI x Infra:
Your Infra is a Graph

Roxane Fischer

CEO & Cofounder @ Anyshift.io

Your AI on-call Engineer

Anyshift

I. The Importance of Context in IaC code generation

I.1 Missing Context & Dependencies

Instructions:

“Generate a Terraform configuration for an AWS VPC peering connection”

Anyshift

I.1 Missing Context & Dependencies - by a Dev

```
# Incomplete VPC Peering configuration
resource "aws_vpc_peering_connection" "peer" {
  peer_vpc_id = "vpc-12345678"
  vpc_id      = "vpc-87654321"
  auto_accept = true
  tags = { Name ="Main-to-Peer", team=”infra” }
}
```

Incomplete:

- Hardcoded values

Anyshift

I.1 Missing Context & Dependencies - by Github Copilot

```
# Incomplete VPC Peering configuration
resource "aws_vpc_peering_connection" "peer" {
  peer_vpc_id = aws_vpc.peer_vpc.id
  vpc_id      = aws_vpc.main_vpc.id
  auto_accept = true
  tags = { Name ="Main-to-Peer", team=”infra” }
}
```

Missing Values

Anyshift

I.1 Missing Context & Dependencies - by Github Copilot

```
# Incomplete VPC Peering configuration
resource "aws_vpc_peering_connection" "peer" {
  peer_vpc_id = aws_vpc.peer_vpc.id
  vpc_id      = aws_vpc.main_vpc.id
  auto_accept = true
  tags = { Name ="Main-to-Peer", team=”infra” }
}

resource "aws_vpc" "main_vpc" {…}
```

Complete but non realistic

Anyshift

I.1 Missing Context & Dependencies - by Github Copilot

```
# Incomplete VPC Peering configuration
resource "aws_vpc_peering_connection" "peer" {
  peer_vpc_id = aws_vpc.peer_vpc.id
  vpc_id      = aws_vpc.main_vpc.id
  auto_accept = true
  tags = { Name ="Main-to-Peer", team=”infra” }
}

resource "aws_vpc" "main_vpc" {…}
```

Complete but non realistic

❌ **LLMs cannot do it**

LLMs lack the dependencies needed to fetch the code.

Anyshift

Can Cursor do it?

II.3 Missing Context & Dependencies

```
# Complete
resource "aws_vpc_peering_connection" "peer" {
  peer_vpc_id = aws_vpc.peer_vpc.id
  vpc_id      = data.aws_vpc.main_vpc.id
  auto_accept = true
  tags = { Name ="Main-to-Peer", team=”infra” }
}

resource "aws_vpc" "main_vpc" {…}
```

```
# Data source dependency
data "terraform_remote_state" "vpc" {
  ...
  config = {
    bucket = "acme"
    region = "us-east-1"
    key    = "network.tfstate"
  }
}
```

Complete

Anyshift

I.2 Requirement for Live Cloud Data

```
# Incomplete VPC Peering configuration
resource "aws_vpc_peering_connection" "peer" {
  peer_vpc_id = "vpc-12345678"
  vpc_id      = "vpc-87654321"
  auto_accept = true
  tags = { Name ="Main-to-Peer", team=”infra” }
}
```

Available in the Cloud only

Anyshift

I.2 Requirement for Live Cloud Data

```
# Incomplete VPC Peering configuration
resource "aws_vpc_peering_connection" "peer" {
  peer_vpc_id = "vpc-12345678"
  vpc_id      = "vpc-87654321"
  auto_accept = true
  tags = { Name ="Main-to-Peer", team=”infra” }
}
```

```
# Data source dependency
data "terraform_remote_state" "vpc" {
  ...
  config = {
    bucket = "acme"
    region = "us-east-1"
    key    = "network.tfstate"
  }
}
```

❌ **Cursor cannot do the reconciliation with Cloud Data**

Cursor would need full cloud access to all AWS accounts to know the ARNs and resource IDs.

Anyshift

II. Live Demo

Requirement for Live Data

III. The AI Context for Devops:
Your infra is a graph

III.1 Your Infrastructure is a Graph

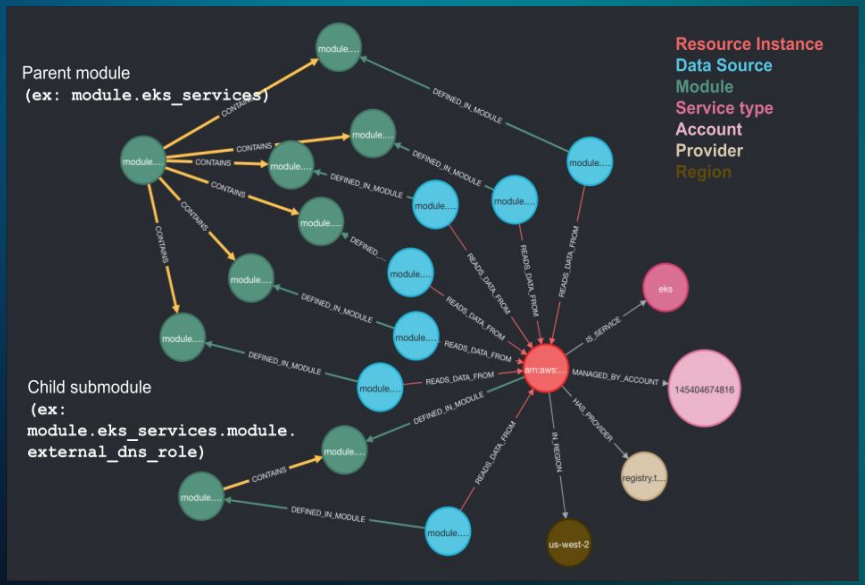

Huge Complexity:

- Multi-accounts
- Multi-cloud
- Ten of thousands of resources
- All the microservices running within

Anyshift

III.1 Your Infrastructure is a Graph

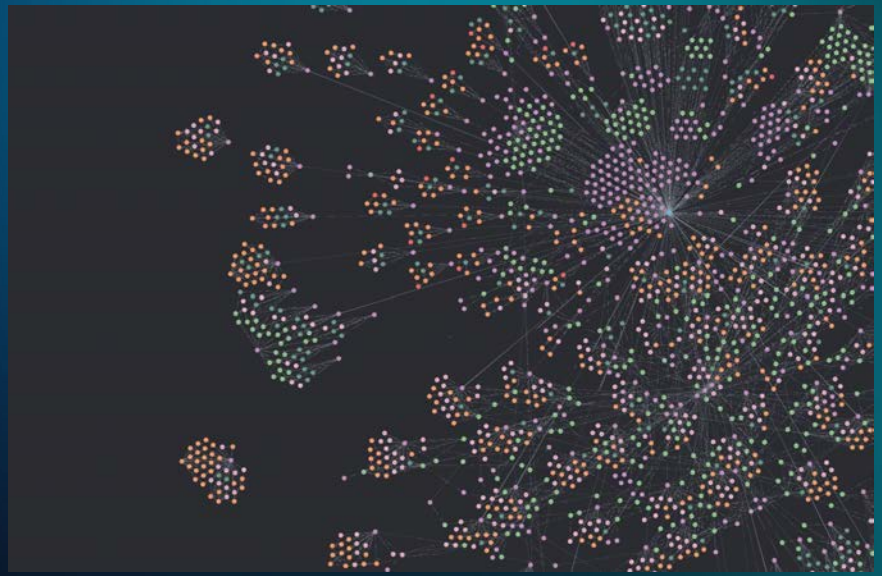

Data Processing:

- Defining meaningful edge weights
- Cloud-to-code relationships

Anyshift

III.2 How to build such graph

1) Build a GraphDB

Nodes + edges built and updated from every infra event
✅ Full graph context, high data quality, minimal LLM drift
❌ Time heavy, hard to maintain.

2) Use MCP Servers

How: Make external calls to MCP servers (e.g AWS)
✅ Easier to build in-house
❌ Fragmented view & Subject to rate limits

Anyshift

III.3 How to build this graph

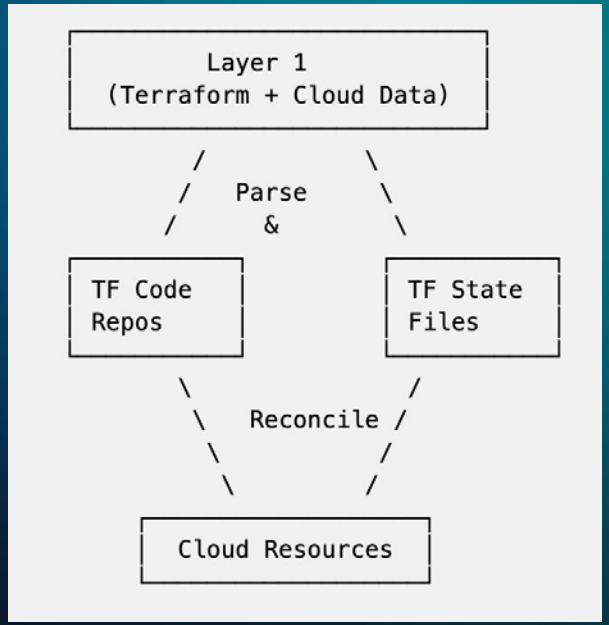

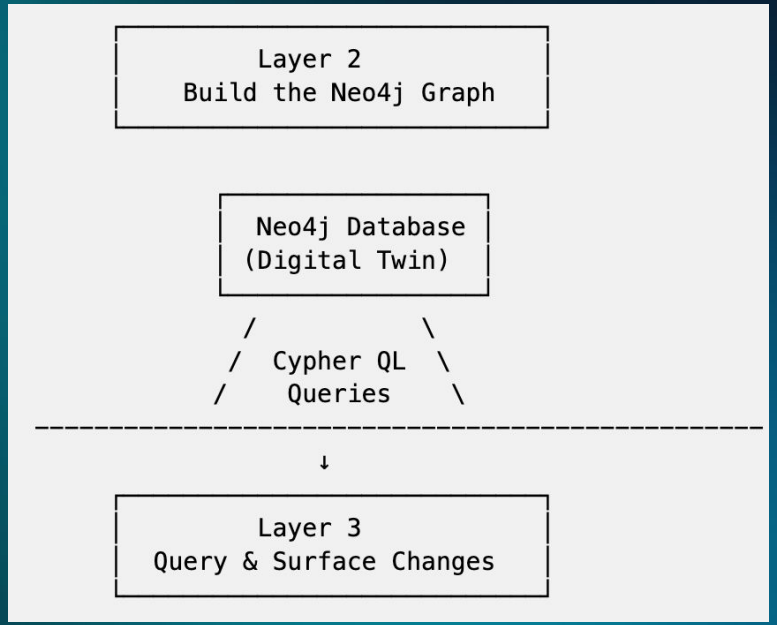

Anyshift

Conclusion

Key Takeaways:

- Context is King: LLMs rely on correct context for accurate IaC
- Your infrastructure is a graph

ROI:

- Offload repetitive tasks from senior SREs
- Accelerate development

Anyshift

Thanks!

Your AI On-Call Engineer

Anyshift