

CPSC 436C Cloud Computing

Winter 2025 Term 1 (September 11, 2025)

Tony Mason (fsgeek@cs.ubc.ca), Lecturer




Logistics

Tony's Office Hours:

- **Monday September 15 @ 16:00-17:00 (Discord)**
- **One more TBA**
- **Random discord/twitch streaming**

TA Office Hours: TBA

Project 1 released



Quick Check-in

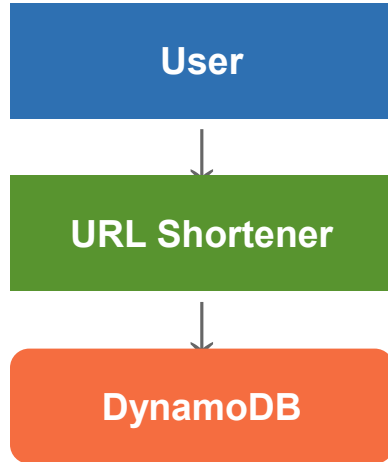
Show of hands: Who got their AI environment working?

Quick share: Most interesting setup challenge?

Remember: This isn't pass/fail - we're learning together



Link Shortener (Simple)



Generate a short name

- Hash value
- Collision – use next available value
- Bitly: $6^{26} = 56,800,235,584$



Need:

- Database
- Index against URI and short name
- Web front-end

We Built It!

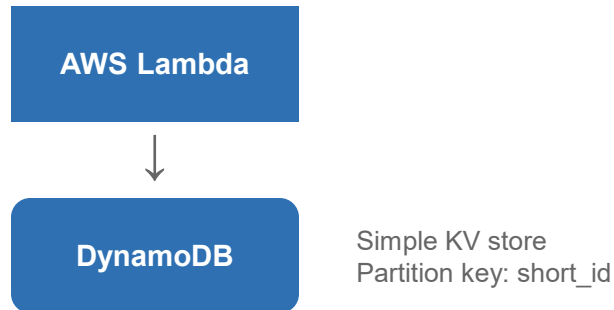
The absolute minimum:

```
# Lambda handler
def shorten(url):
    key = hash(url)[:6]
    table.put_item(
        Key={'id': key},
        Item={'url': url}
    )
    return key
```

Services needed:

- Lambda (1 function)
- DynamoDB (1 table, key-value)
- API Gateway (2 endpoints)

Monthly cost: \$0 (free tier)

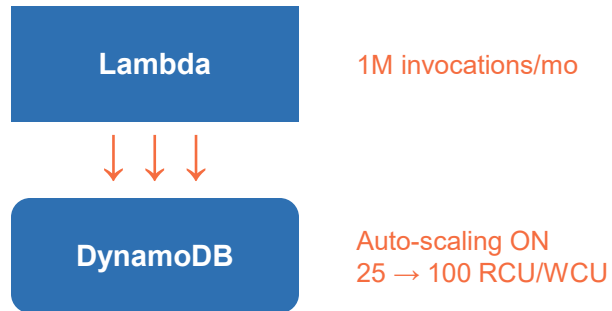


People Love It!

Viral growth hits:

Week 1: 100 URLs/day
Week 2: 1,000 URLs/day
Week 4: 10,000 URLs/day
Week 8: 50,000 URLs/day

*DynamoDB auto-scales!
(But not for free...)*



The First Bill

AWS Cost Explorer

Lambda:	\$12.47
1.2M requests	
DynamoDB:	\$28.93
100 RCU/WCU avg	
API Gateway:	\$5.83
1.2M calls	
Total:	\$47.23

Breaking free tier limits:

Lambda:

- Free: 1M requests
- Used: 1.2M requests
- Overage: 200K @ \$0.20/1M

DynamoDB:

- Free: 25 capacity units
- Used: 100 capacity units
- Overage: 75 @ \$0.00065/hour


"How do we pay for this?"

The Meeting

Three options on the table:

1. Shut down
 - ✗ We have 50K daily users!
2. Charge \$5/month
 - ✗ Bit.ly is free
3. Add analytics & monetize
 - ✓ "We already have the data..."

Current architecture



"Just Analytics"

"We need to understand traffic"

New requirements:

- Count clicks per URL
- Track click timestamps
- Basic geographic data

Problem: DynamoDB can't do SQL queries efficiently!

**Solution: Add RDS PostgreSQL
+\$15/month (db.t3.micro)**

Lambda

RDS PostgreSQL

CloudWatch

DynamoDB

Lambda now writes to both:

- DynamoDB for redirects
- RDS for analytics

New tables:

- clicks (timestamp, short_id, ip)
- analytics (url, count, last_click)

Marketing Wants In

"Show us conversion funnels!"

New requirements:


- User sessions (cookies)
- Referrer tracking
- Campaign attribution
- Real-time dashboards

New services needed:

- ElastiCache (sessions)
- Kinesis (real-time)
- QuickSight (dashboards)

Revenue: \$500/month

Costs: \$125/month



Enterprise Features

"We need user accounts!"


Enterprise requirements:

- User authentication
- Custom domains
- API access
- Bulk operations
- Data exports

New services:

- Cognito (auth)
- Route 53 (domains)
- CloudFront (CDN)
- S3 (exports)
- SES (emails)

Revenue: \$5,000/month



12 AWS services
\$800/month infrastructure

The Data Product

"We're a data company now"

Selling to enterprises:

- Real-time click feeds
- Behavioral analytics
- ML predictions
- Industry benchmarks

New infrastructure:

- EMR (big data)
- SageMaker (ML)
- Redshift (warehouse)
- Athena (queries)
- Glue (ETL)

Revenue: \$50K/month

Costs: \$8K/month

Original purpose: 2% of infrastructure

EMR + Kinesis + Redshift

SageMaker + Comprehend + Forecast

Athena + QuickSight + Glue

20+ AWS services
100TB data/month
10M user profiles

What Did We Build?

We wanted to:

Shorten URLs

2 AWS services
\$0/month
10 lines of code

We built:

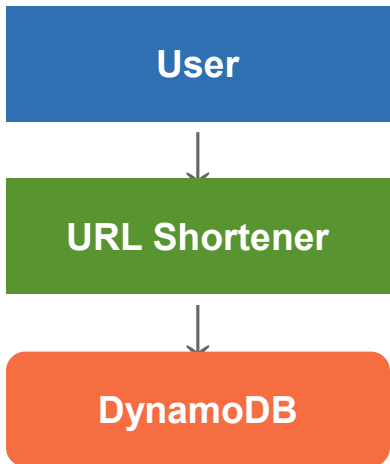
Surveillance infrastructure

25+ AWS services
\$8,000/month
100K+ lines of code

*The cloud services made it easy.
The business model made it inevitable.*

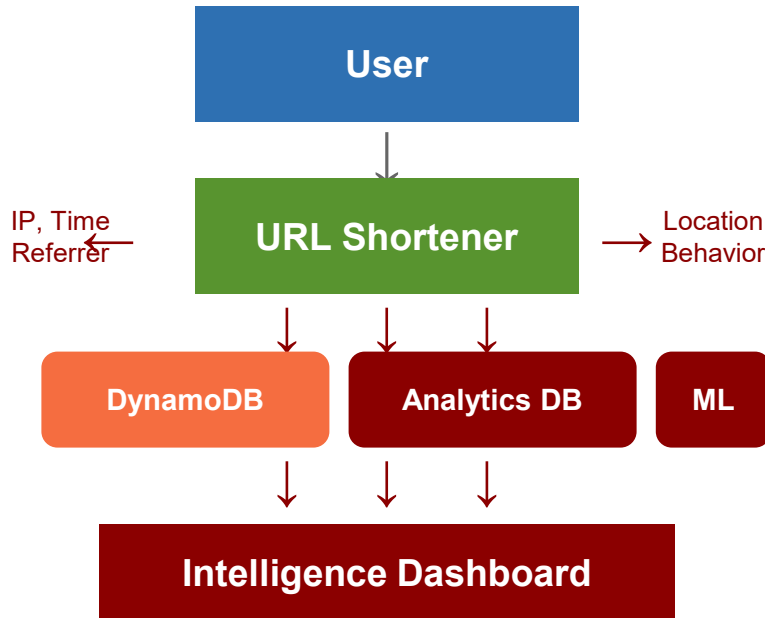
Architecture Evolution: Tool → Surveillance

Tuesday: "Simple" URL Shortener



Clean, simple, functional

Thursday: Surveillance Infrastructure



"Same functionality, surveillance infrastructure"

Questions You Should Ask

Before building anything:

- What data does this collect?
- Who has access to it?
- What decisions will it drive?
- Who bears the cost? Who reaps the benefit?
- Would I want to use this on myself?



From Recognition to Understanding

Data collection is not inherently exploitative

- How will it be used?
- How do we preserve privacy?
- How do we minimize sensitive data exposure?
- How do we ensure transparency?
- What are we building?
 - What does it do?
 - What does it *not* do?



Applying Tuesday's Lesson

What we learned:

- Vague queries = useless responses
- Constraints *improve* outputs
- AI **does not** understand implications: you do



ML Model Selection

Bad: "What's the best ML model?"

[Generic, useless response]

Better: "I need a model for [task] with [constraints]..."

[Improved, specific response]

Best: "...What are the surveillance implications of each option?"

[Critical analysis of data collection and privacy]



Project 1: Scenario

Your "startup" needs a prediction API

- Current: 100 requests/day
- Viral event: 50,000 requests/hour
- Constraint: Free tier only



The subtext: You're building infrastructure that will outlive your job

Project Requirements: Technical



- Design document with architecture
- Implementation (git repo)
- AI collaboration documentation
- **NEW: Privacy impact assessment**
- Due: September 23 11:00am PT

Project Requirements: Data Usage



- What data are you collecting?
- How could this be misused?
- What safeguards did you include?
- Document your ethical considerations

Project Focus

Before diving into code:

1. Understand your model's data requirements
2. Map the information flow
3. Identify potential misuse cases
4. THEN start architecture

Resources: Discord #project-1-help | Office hours for technical AND ethical questions



The Path Forward



- You're learning to build powerful systems
- With power comes responsibility (yes, the Spider-Man quote)
- Technical excellence includes ethical consideration
- Document everything - especially your concerns

Tuesday Goal

Preparation (not graded, but important):

- Initial problem analysis (1 page max)
- Include: technical approach AND ethical considerations
- Come ready to discuss trade-offs



Questions & Support

- Technical questions?
- Ethical concerns?
- Setup issues?



Remember: Questioning the assignment IS part of the assignment

Capstone Project Ideas: Coming Soon

I have several of these will put them into a Github repo.



Discussion

Too much lecture... Let's talk.



100 SECONDS OF





THE UNIVERSITY OF BRITISH COLUMBIA

THE UNIVERSITY OF BRITISH COLUMBIA