INSTANA

Learning from the visible past to accelerate the observable future

What can chocolate cake teach us about Observability

What is Observability

Depends on who you ask …

“If you are observable then **data** is available so I can **analyze** and **understand** you.”

Paraphrase from **Observability vs. Monitoring,**
https://www.instana.com/blog/observability-vs-monitoring/

And this needs to be done in **production** …

Challenging Trends

- **Production is the validation environment**
- **Increasing speed everywhere**: dev, build, deploy, troubleshoot, ..
- **Embrace complexity**: complex software manages more complex software
- **Just in time structure:** FaaS
- **Infrastructure is less visible**

Important Questions and Considerations

Why bake a chocolate cake?	What problems and challenges to tackle?
What are the best baking techniques?	Can we identify some best practices?
What can we learn from past cakes?	What can the past teach us about how to proceed?

A Taste of the Problem Space

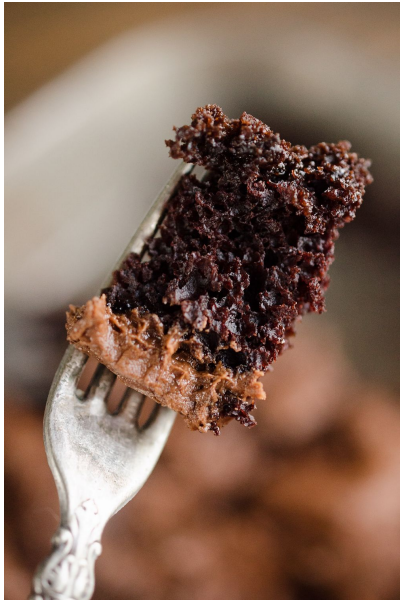

Who	What	Example Problem Statement
Business owners	User journeys (workflows, SLI/SLO)	What services are involved in this user journey?
Developers	Services (business logic, endpoints)	What business goals are impacted by these services? What h/w or infrastructure is impacting this service?
SRE / Operators	Infrastructure (h/w, schedulers, frameworks, libraries, virtualization)	If I retire this h/w, what services are impacted?

A Well Stocked Kitchen

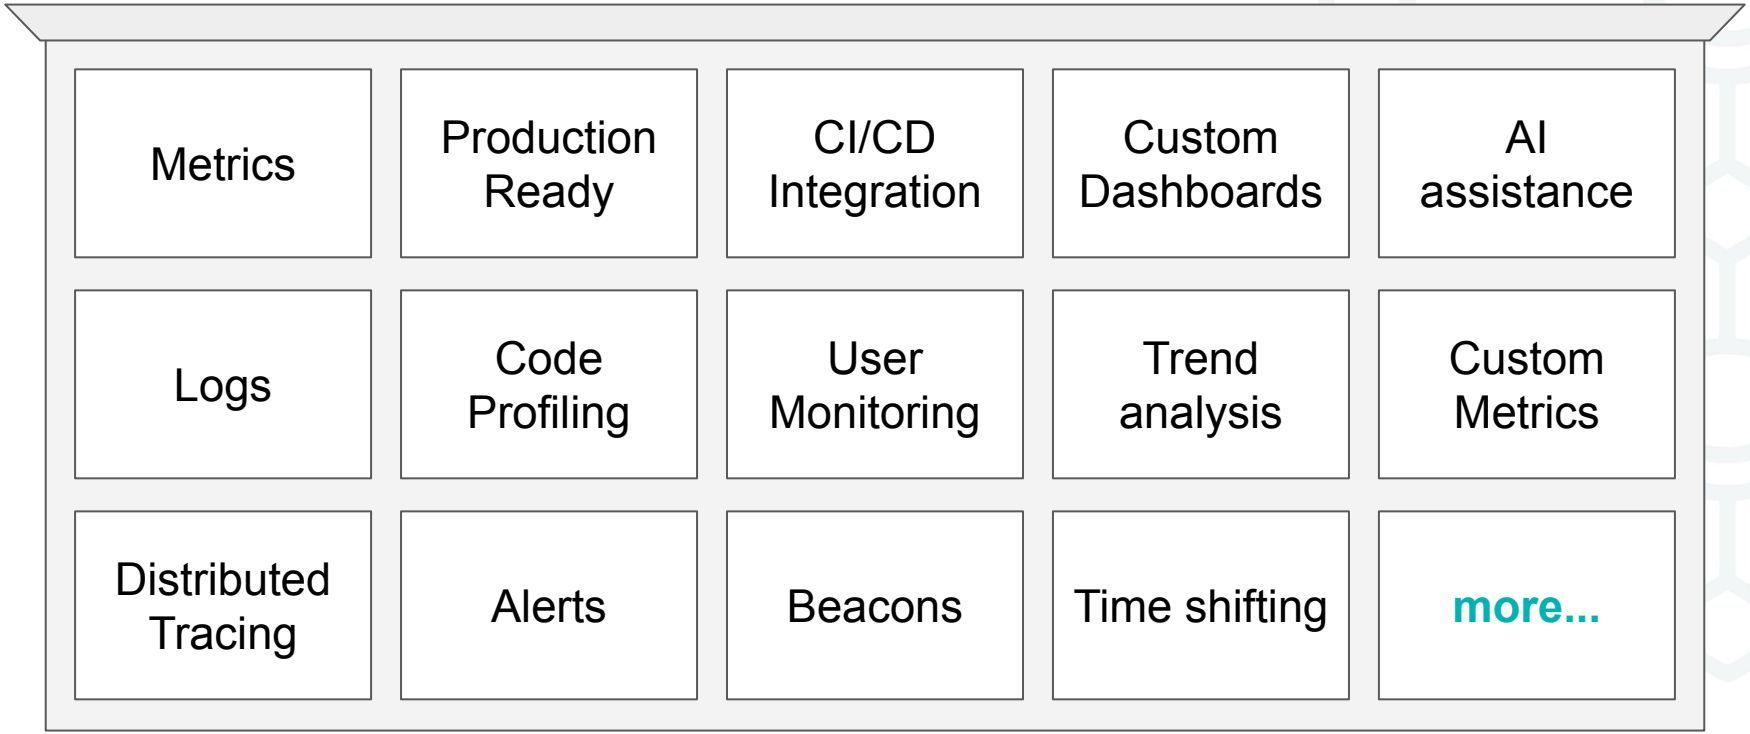

Ingredients or Chocolate Cake First?

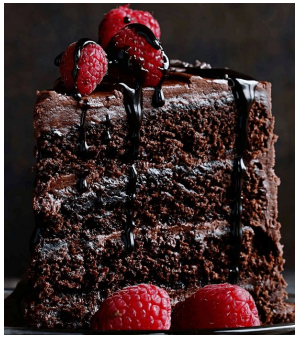

Ingredients or Observability First?

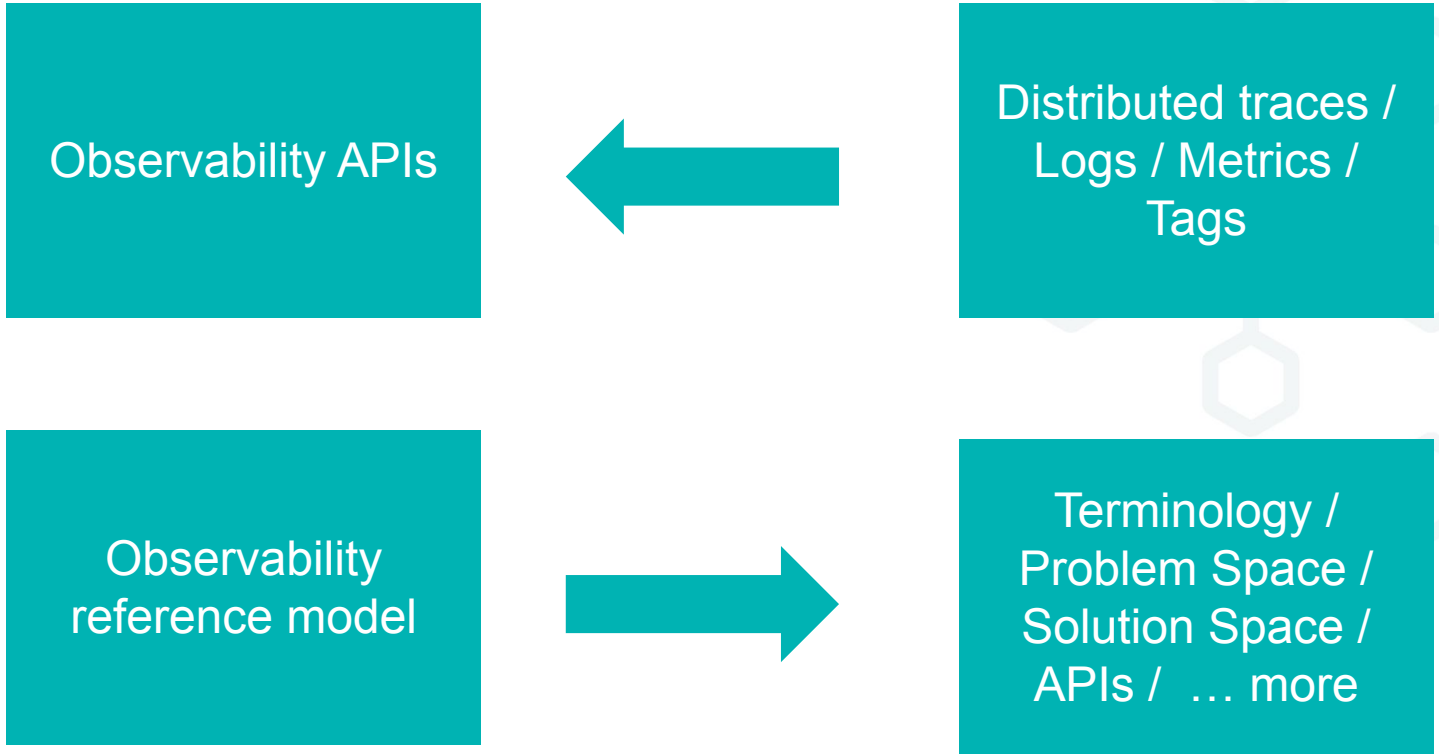

Why a Reference Model

- Promote understanding of the problem space and not specific solutions
- Technology agnostic
- Terminology, entities, and relationships
- Only clarify "things within an environment" or a problem space

“Community”

- A **reference model** is an abstract framework or domain-specific ontology consisting of an interlinked set of clearly defined concepts … to encourage clear communication.
- it represents a complete set. Reference models are often illustrated as a set of concepts with some indication of the relationships between the concepts.

https://en.wikipedia.org/wiki/Reference_model

Past Reference Models

"CCS: Calculus of Communicating Systems"
Milner, 1980

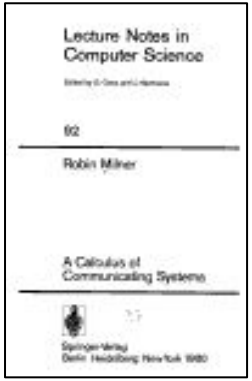

Business user journeys, workflows

"Communicating Sequential Processes"
Hoare, 1986

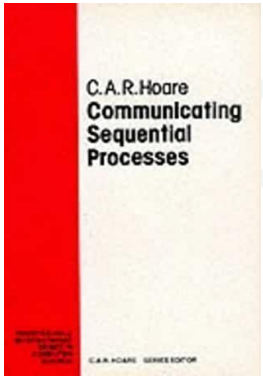

Services,
(business logic, endpoints)

"Queueing Systems"
Kleinrock, 1975

Infrastructure, schedulers

INSTANA

Three Dimensions (3D) of Observability

Business user journeys
(workflows,SLI / SLO)

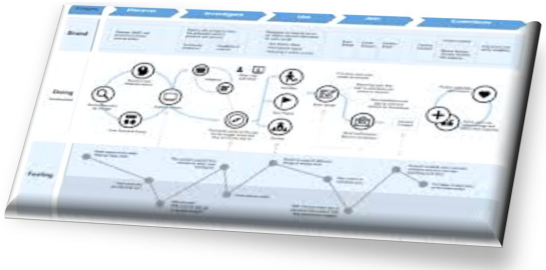

Services,
(business logic, endpoints)

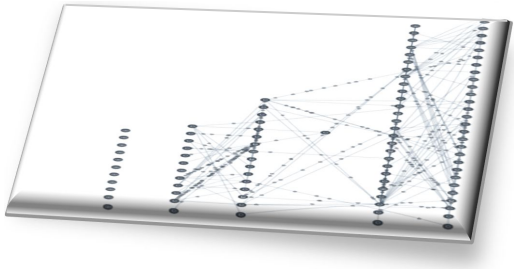

Infrastructure
(h/w, schedulers frameworks, libraries)

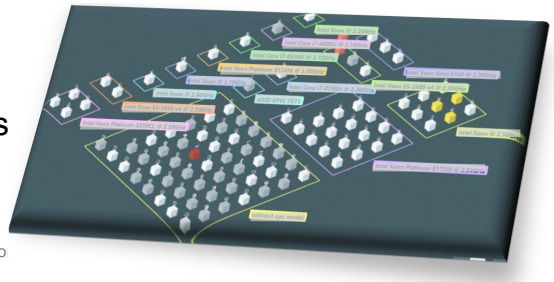

Three Dimensions (3D) of Observability

Business user journeys (workflows,SLI / SLO)

Services, (business logic, endpoints)

Infrastructure (h/w, schedulers frameworks, libraries)

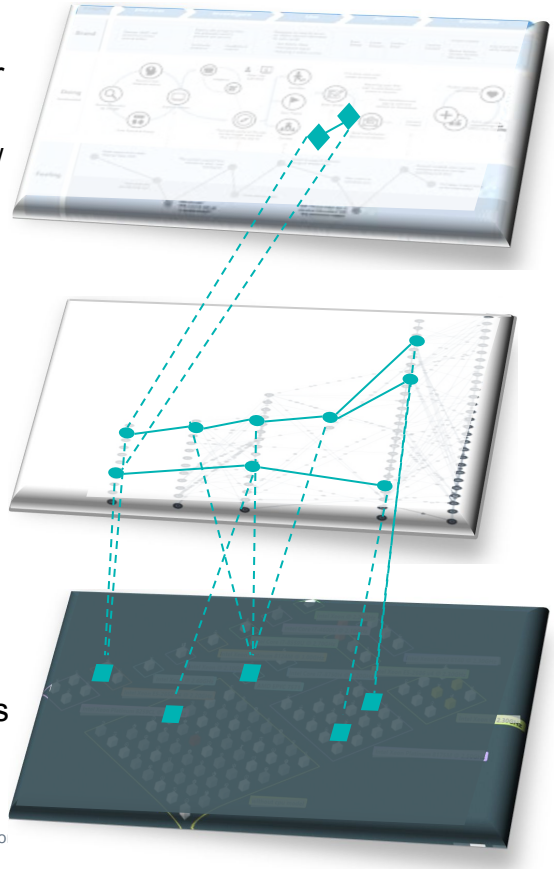

- What services are involved in this user journey?
- What h/w or infrastructure is impacting this service?

Three Dimensions (3D) of Observability

Business user journeys (workflows,SLI / SLO)

Services, (business logic, endpoints)

Infrastructure (h/w, schedulers frameworks, libraries)

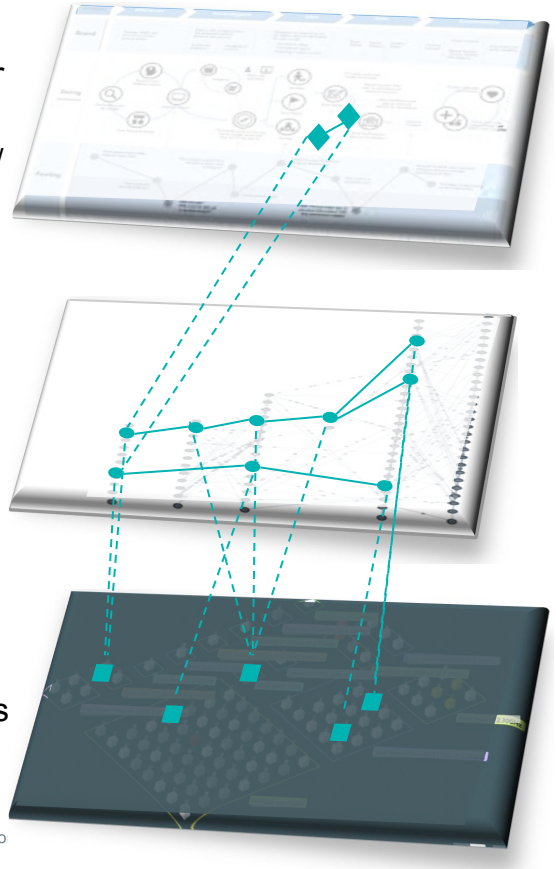

- What services are involved in this user journey?
- What h/w or infrastructure is impacting this service?

Three Dimensions (3D) of Observability

Business user journeys (workflows,SLI / SLO)

Services, (business logic, endpoints)

Infrastructure (h/w, schedulers frameworks, libraries)

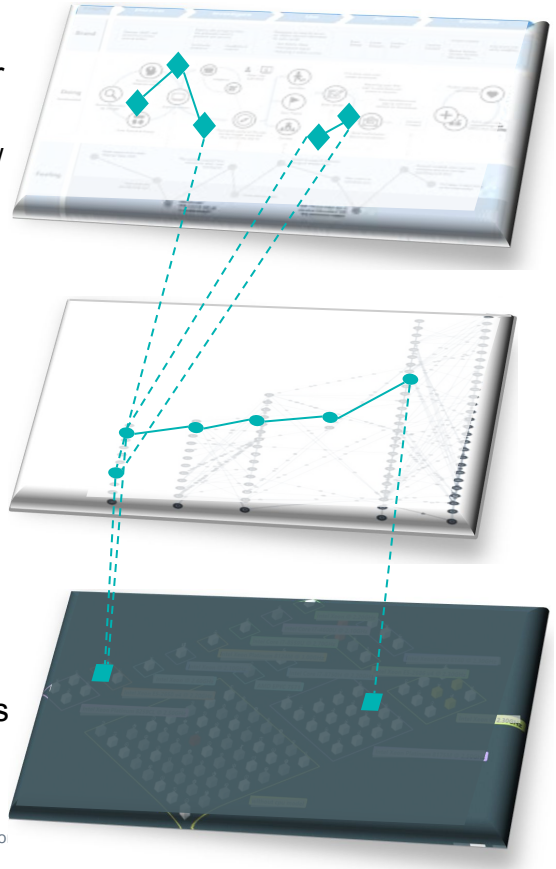

- What business goals are impacted by these services?
- If I retire this h/w what services are impacted?

Three Dimensions (3D) of Observability

Business user journeys
(workflows,SLI / SLO)

Services,
(business logic, endpoints)

Infrastructure
(h/w, schedulers frameworks, libraries)

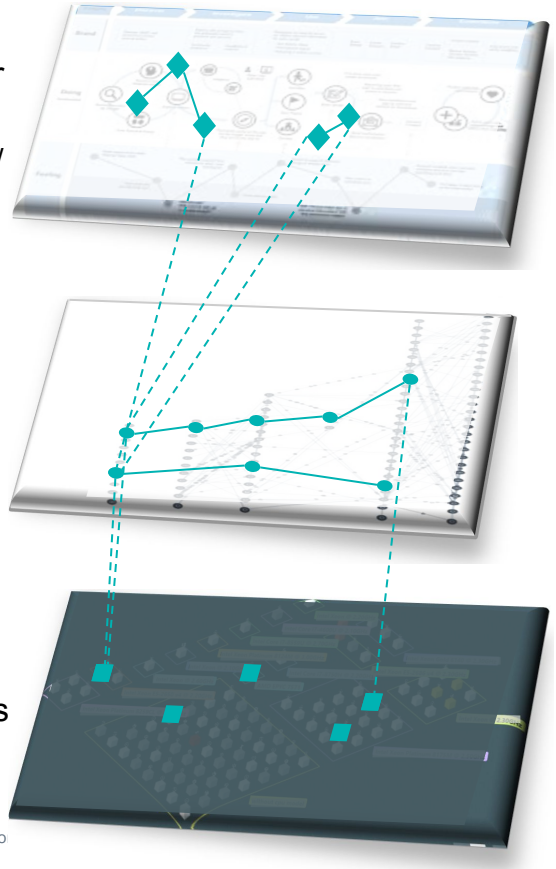

- What business goals are impacted by these services?

- If I retire this h/w what services are impacted?

Three Dimensions (3D) of Observability

Business user journeys (workflows,SLI / SLO)

Services, (business logic, endpoints)

Infrastructure (h/w, schedulers frameworks, libraries)

- What user journeys (business KPIs) have priority?
- How to map logical services to actual services?
- What number of of replicas?
- Where to place the replicas?

Three Dimensions (3D+1) of Observability

Business user journeys
(workflows,SLI / SLO)

Services,
(business logic, endpoints)

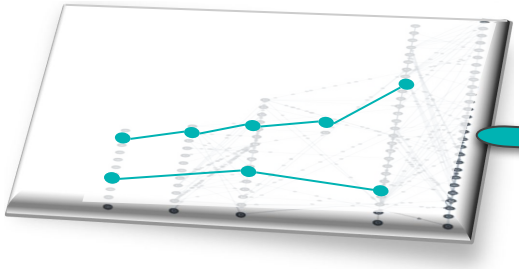

Infrastructure
(h/w, schedulers frameworks, libraries)

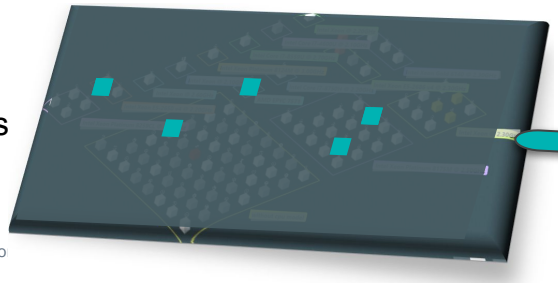

Platform
(global scheduler, orchestration)

INSTANA

One Possible Formalism

Business user journeys (workflows,SLI / SLO)

Services, (business logic, endpoints)

Infrastructure (h/w, schedulers frameworks, libraries)

Platform (global scheduler, orchestration)

One Possible Formalism

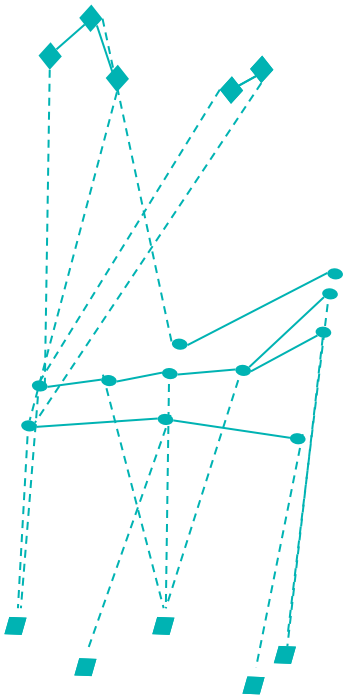

- An algorithmic attributed graph grammar
  - Nodes (objects) are typed with attributes (tag+value)
  - Edges (relationships) are typed and have meaning
  - Sub-graphs have semantics
  - Reason or simplify using graph rewrite rules

- Technology agnostic
  - Can be implemented in various ways: Graph DB, SQL, noSQL, etc.

The Root Cause Triage Recipe

Business user journeys
(workflows,SLI / SLO)

Services,
(business logic, endpoints)

Infrastructure
(h/w, schedulers frameworks, libraries)

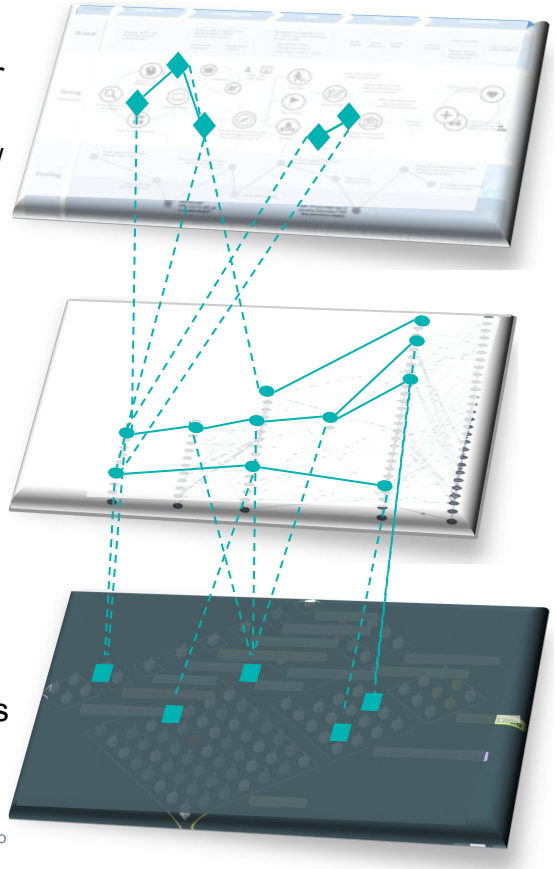

The root cause triage recipe:

1. Contrast when things were good
2. Compare with a healthy object of the same type
3. Examine the upstream or downstream objects
4. Check recent config changes
5. Review a similar issue in the past
6. Is it second order interference
7. Was a limit reached
8. Bug somewhere
9. At a dead end so go back several objects, goto 1
10. Move to a different dimension and goto 1

Slicing the Cake

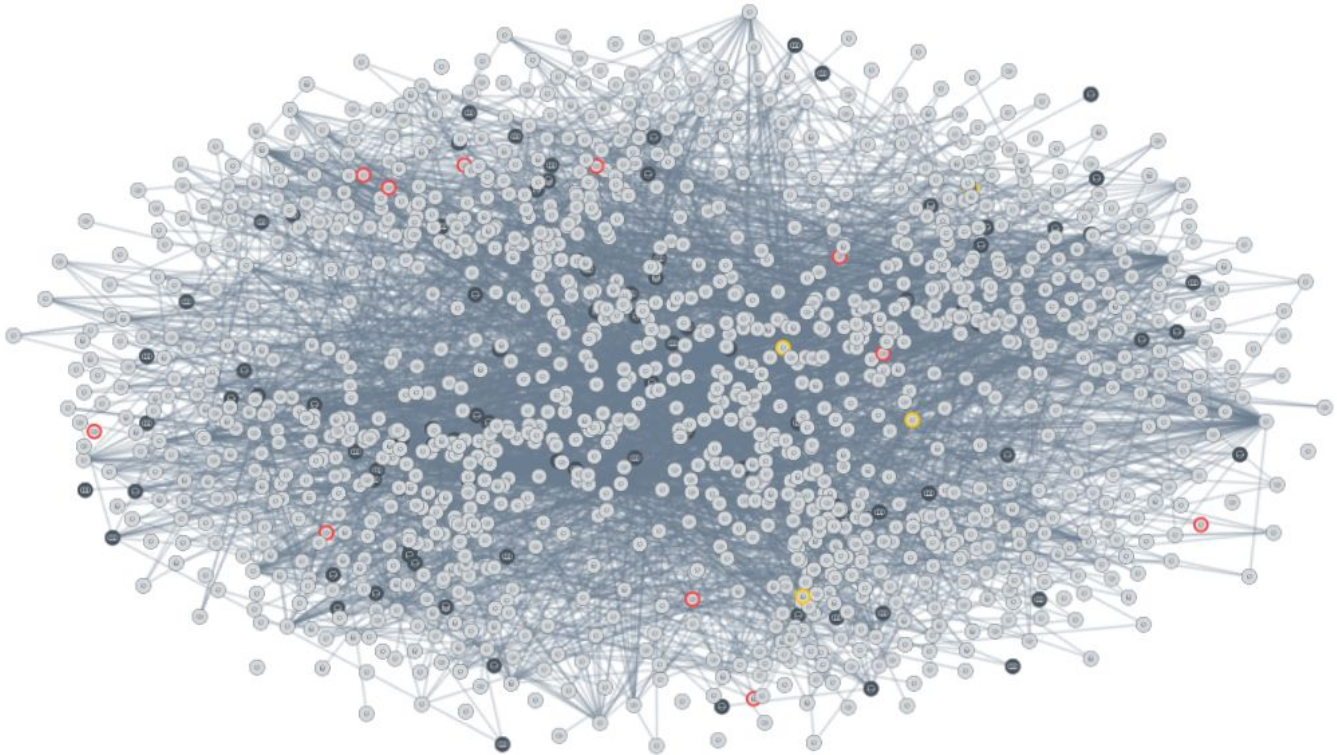

Is a Big Slice Too Much?

Smaller Slices - Easier to Manage

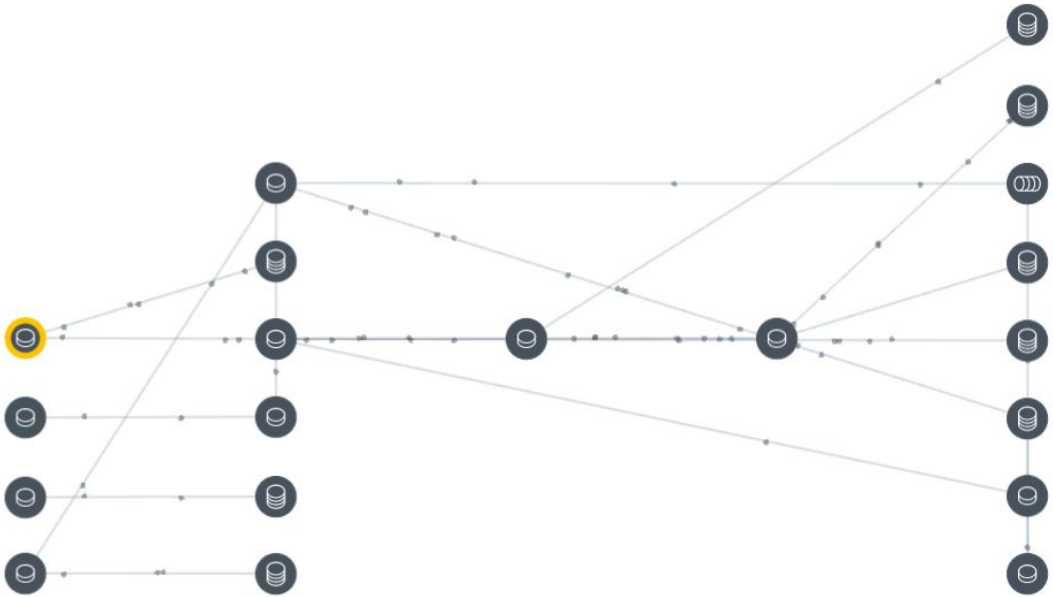

Custom Sized Slices for Everyone

Many ways to slice it …

1. Grouping several services or endpoints
2. K8s namespace or …
3. Location: env, geo, cloud, or hosts
4. Protocol: HTTP, RPC, …
5. By technology: DB, MQ

Getting a Custom Slice

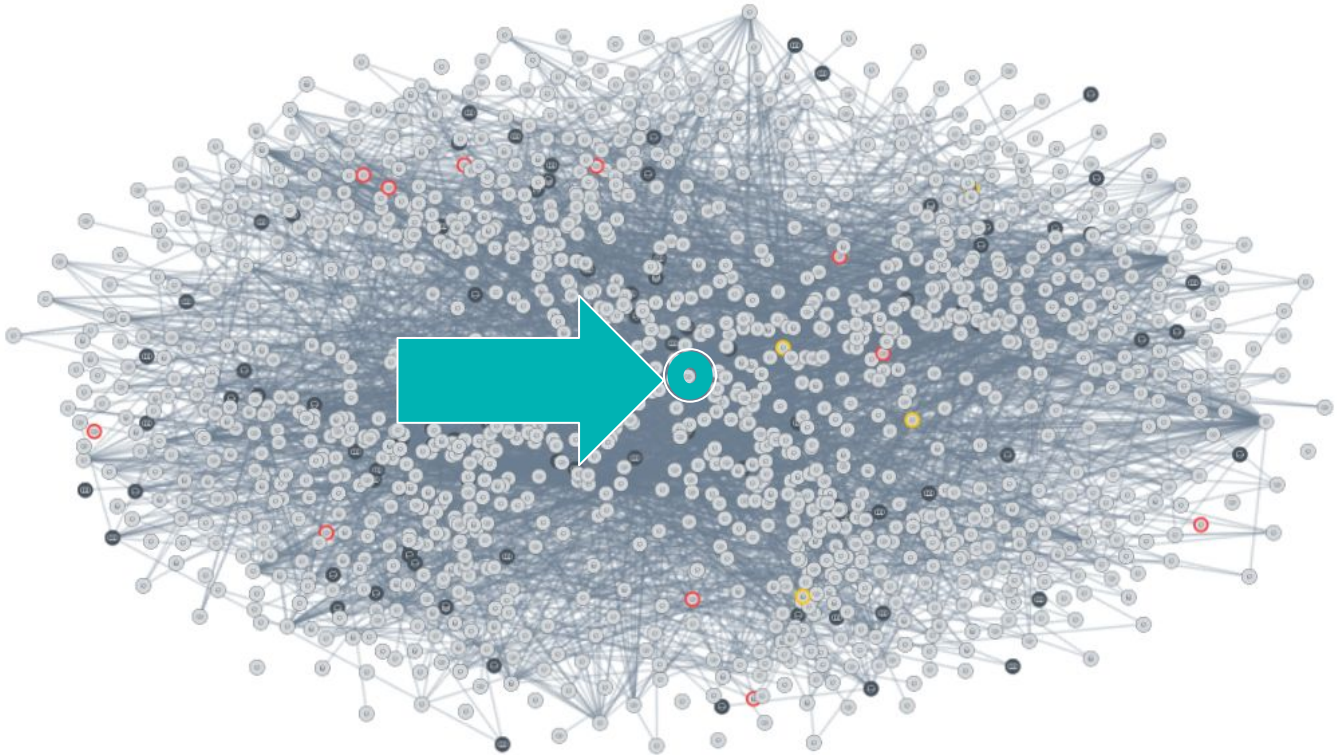

Getting a Custom Slice

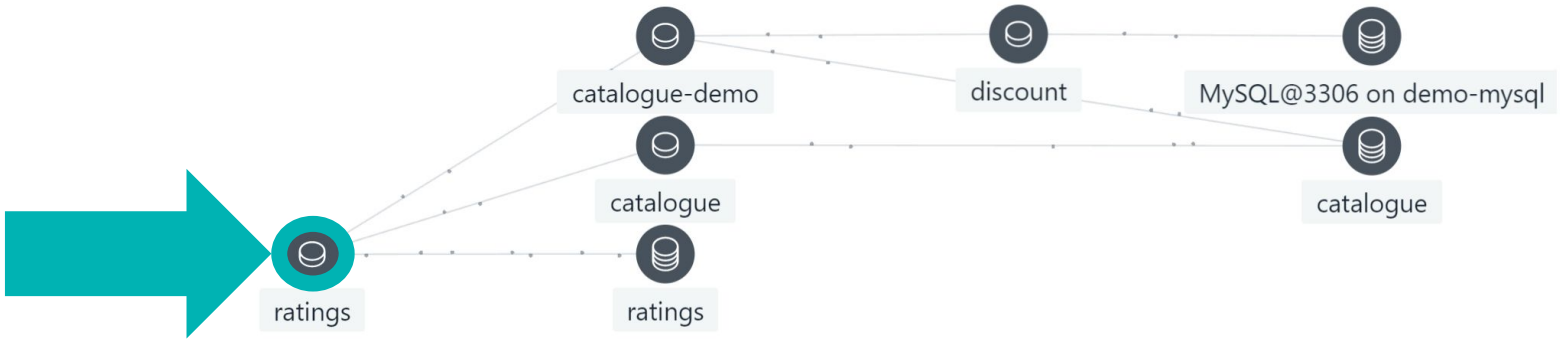

INSTANA

Going from 1 Cake to 1,000 Cakes

Sampling is a way to scale

But sampling impacts the numerator or denominator of the error rate calculation

$$\frac{number\ of\ errors}{number\ of\ successes\ and\ errors}$$

So how to minimize sampling bias or bound the error or … for my needs?

Which sampling strategy to pick?

- No sampling
- Head-based sampling
- Tail-based sampling
- Stratified sampling
- Static sampling
- Dynamic sampling
- Key-based dynamic sampling
- Constant throughput
- Constant throughput per key
- Average sample rate
- Average sample rate with minimum per key
- …

Forming the Mold

- Does the observability community need a reference model to …
  - Define the problem space, terminology, objects, and relationships
  - Use a formalism that provides a semantics and eanbles interoperability
  - Incorporate challenges beyond the API (e.g., scoping, triage workflow, etc.)
- This activity compliments and informs the API development
- If you are interested in taking the next steps Contact me: curtis.hrischuk@instana.com

Q&A

Email me if you are interested in taking the next steps

Curtis.Hrischuk@instana.com

Bonus: Test your observability knowledge by giving our crossword puzzle a try!

https://www.instana.com/crossword

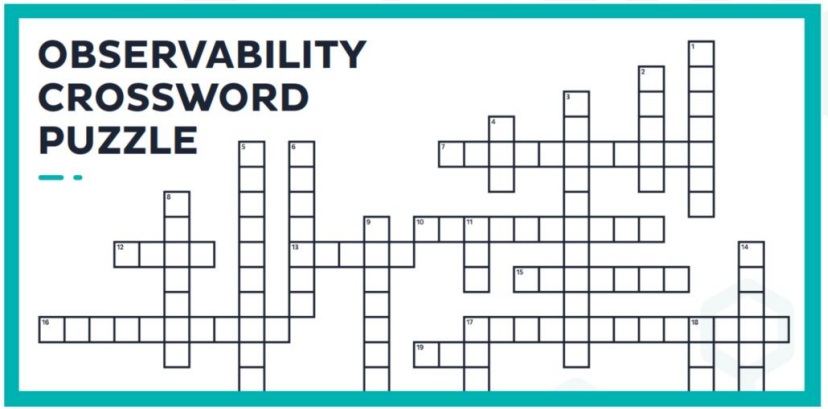

Thank you

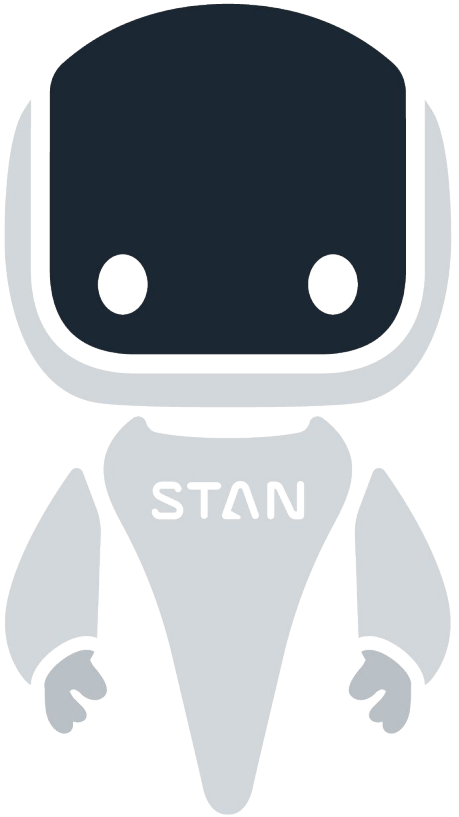

INSTANA

www.instana.com