



Missing the Metaphor for Self-Driving Databases

HPTS, September 16, 2024

Johann Schleier-Smith
CEO & Founder, CrystalDB

www.crystaldb.cloud

@jssmith

Oracle 2018: Autonomous Database



Source: [Super Color Digital](#)

Carnegie Mellon 2017: Self-Driving Database

Self-Driving Database Management Systems

Andrew Pavlo, Gustavo Angulo, Joy Arulraj, Haibin Lin, Jiexi Lin, Lin Ma, Prashanth Menon
Todd C. Mowry, Matthew Perron, Ian Quah, Siddharth Santurkar, Anthony Tomasic
Skye Toor, Dana Van Aken, Ziqi Wang, Yingjun Wu*, Ran Xian, Tieying Zhang
Carnegie Mellon University, *National University of Singapore

ABSTRACT

In the last two decades, both researchers and vendors have built advisory tools to assist database administrators (DBAs) in various aspects of system tuning and physical design. Most of this previous work, however, is incomplete because they still require humans to make the final decisions about any changes to the database and are

Much of the previous work on self-tuning systems is focused on standalone tools that target only a single aspect of the database. For example, some tools are able to choose the best logical or physical design of a database [16], such as indexes [30, 17, 58], partitioning schemes [6, 44], data organization [7], or materialized views [5]. Other tools are able to select the tuning parameters for

Google 2024: Gemini in Databases



Source: [SkyTel](#)

Microsoft 1999

Self-Tuning Technology in Microsoft SQL Server

Surajit Chaudhuri, Eric Christensen, Goetz Graefe, Vivek Narasayya, Michael Zwilling
Microsoft Corporation
One Microsoft Way, Redmond, WA, 98052
Email Contact: surajitc@microsoft.com

1 Introduction

Today's databases require database administrators (DBAs) who are responsible for performance tuning. However, as usage of databases becomes pervasive, it is important that the databases are able to automatically tune themselves to application needs and hardware capabilities rather than require external intervention by DBAs



Technical Insight



Fundamental challenge

- Workloads are valuable
- Workloads are complex

Will automation ever take over the job?



**ROAD
WORK
AHEAD**

JAGUAR
WAYMO

Lyons & Lyon Market
Liquor Wine Beer ATM
Groceries 415-775-4480

LYONS


5WDA993

9GRC786

2811EV2

Typical Autonomous Vehicle (AV) Architecture

Perception



Planning



Control

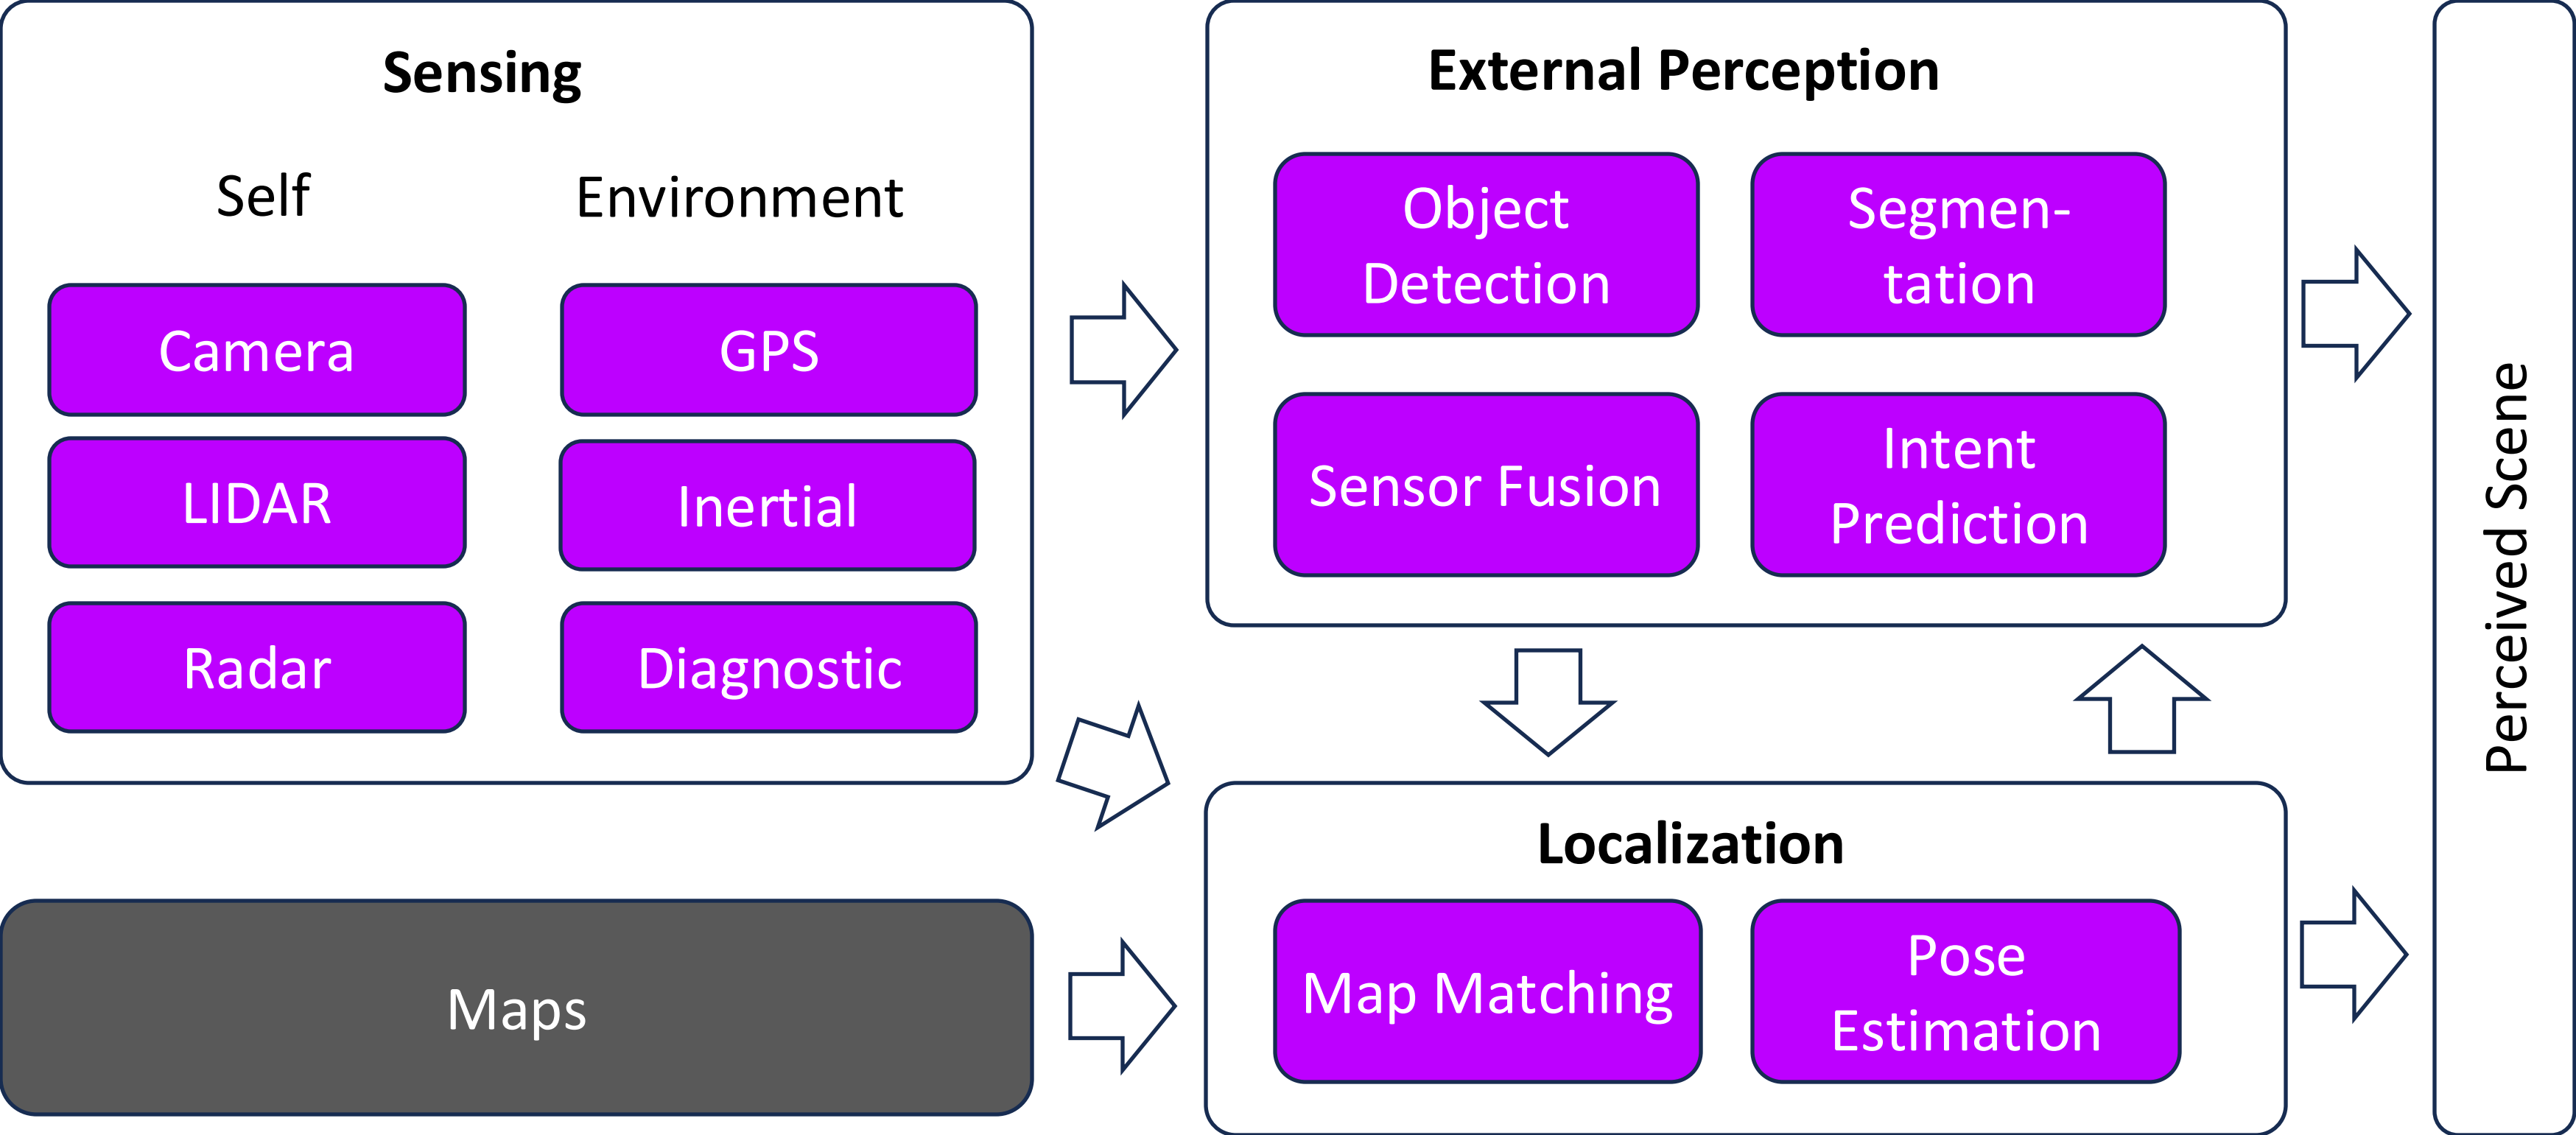









Autonomous Vehicle Perception



Typical Autonomous Vehicle (AV) Architecture

Perception



Planning



Control



Good ML Reliability

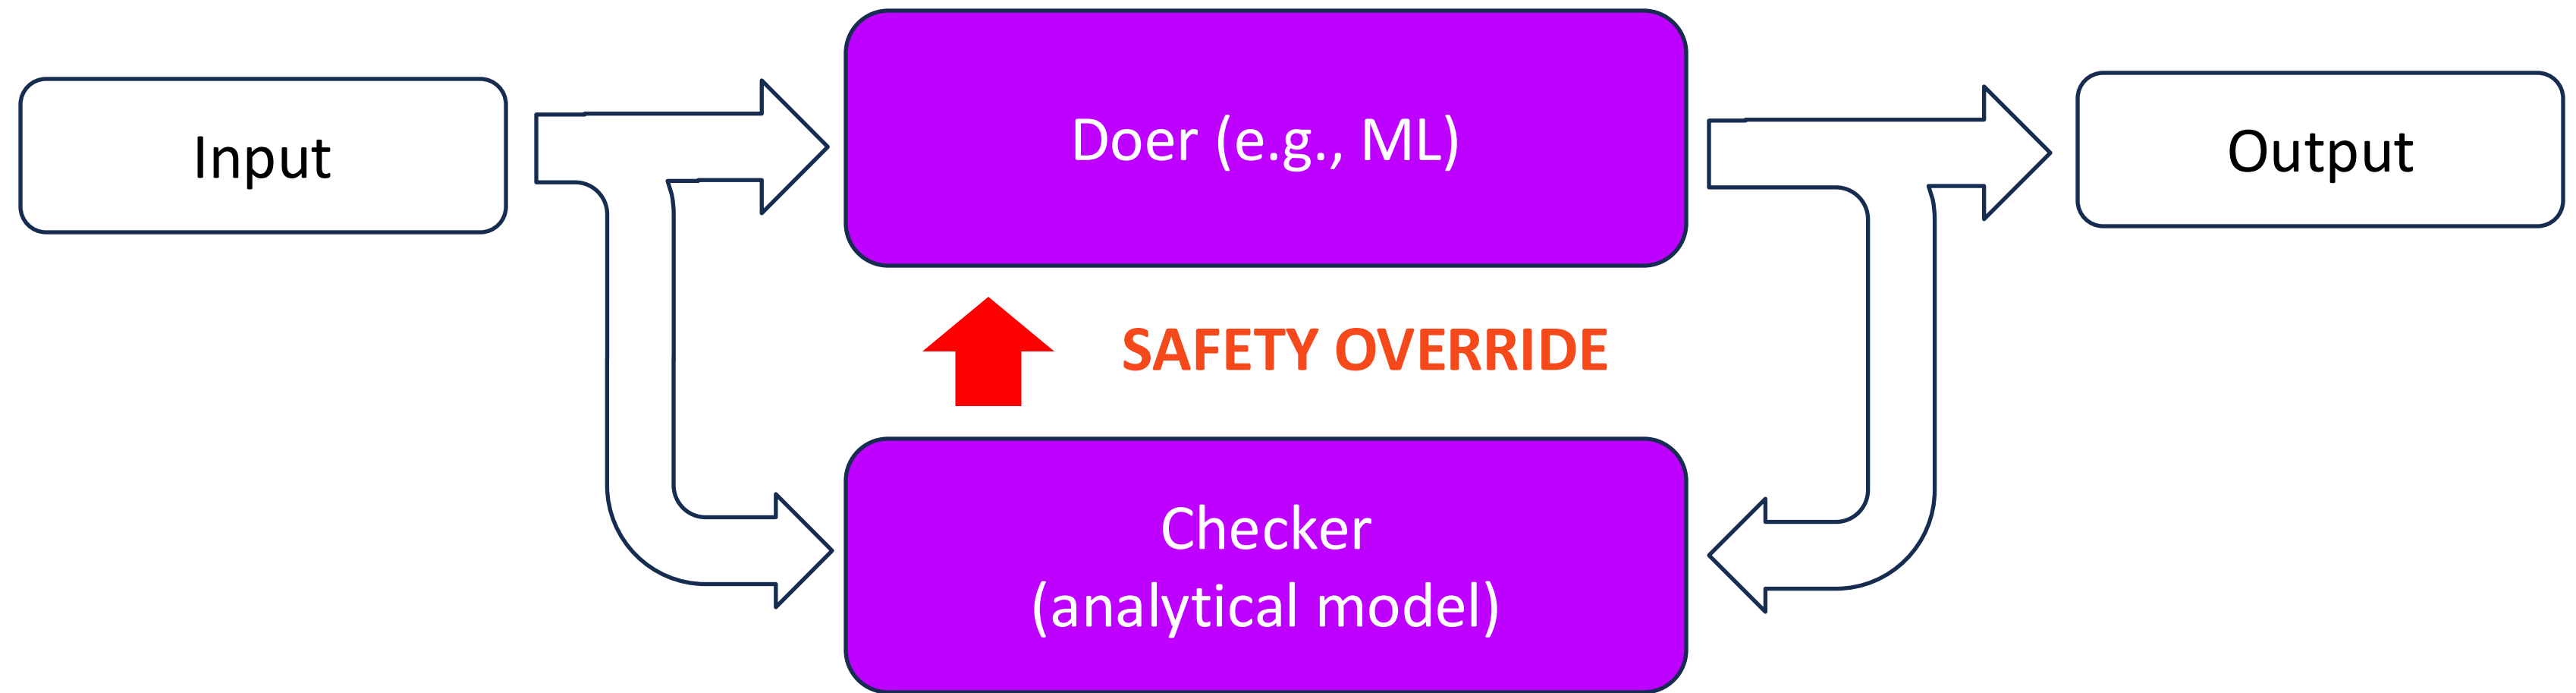
99%

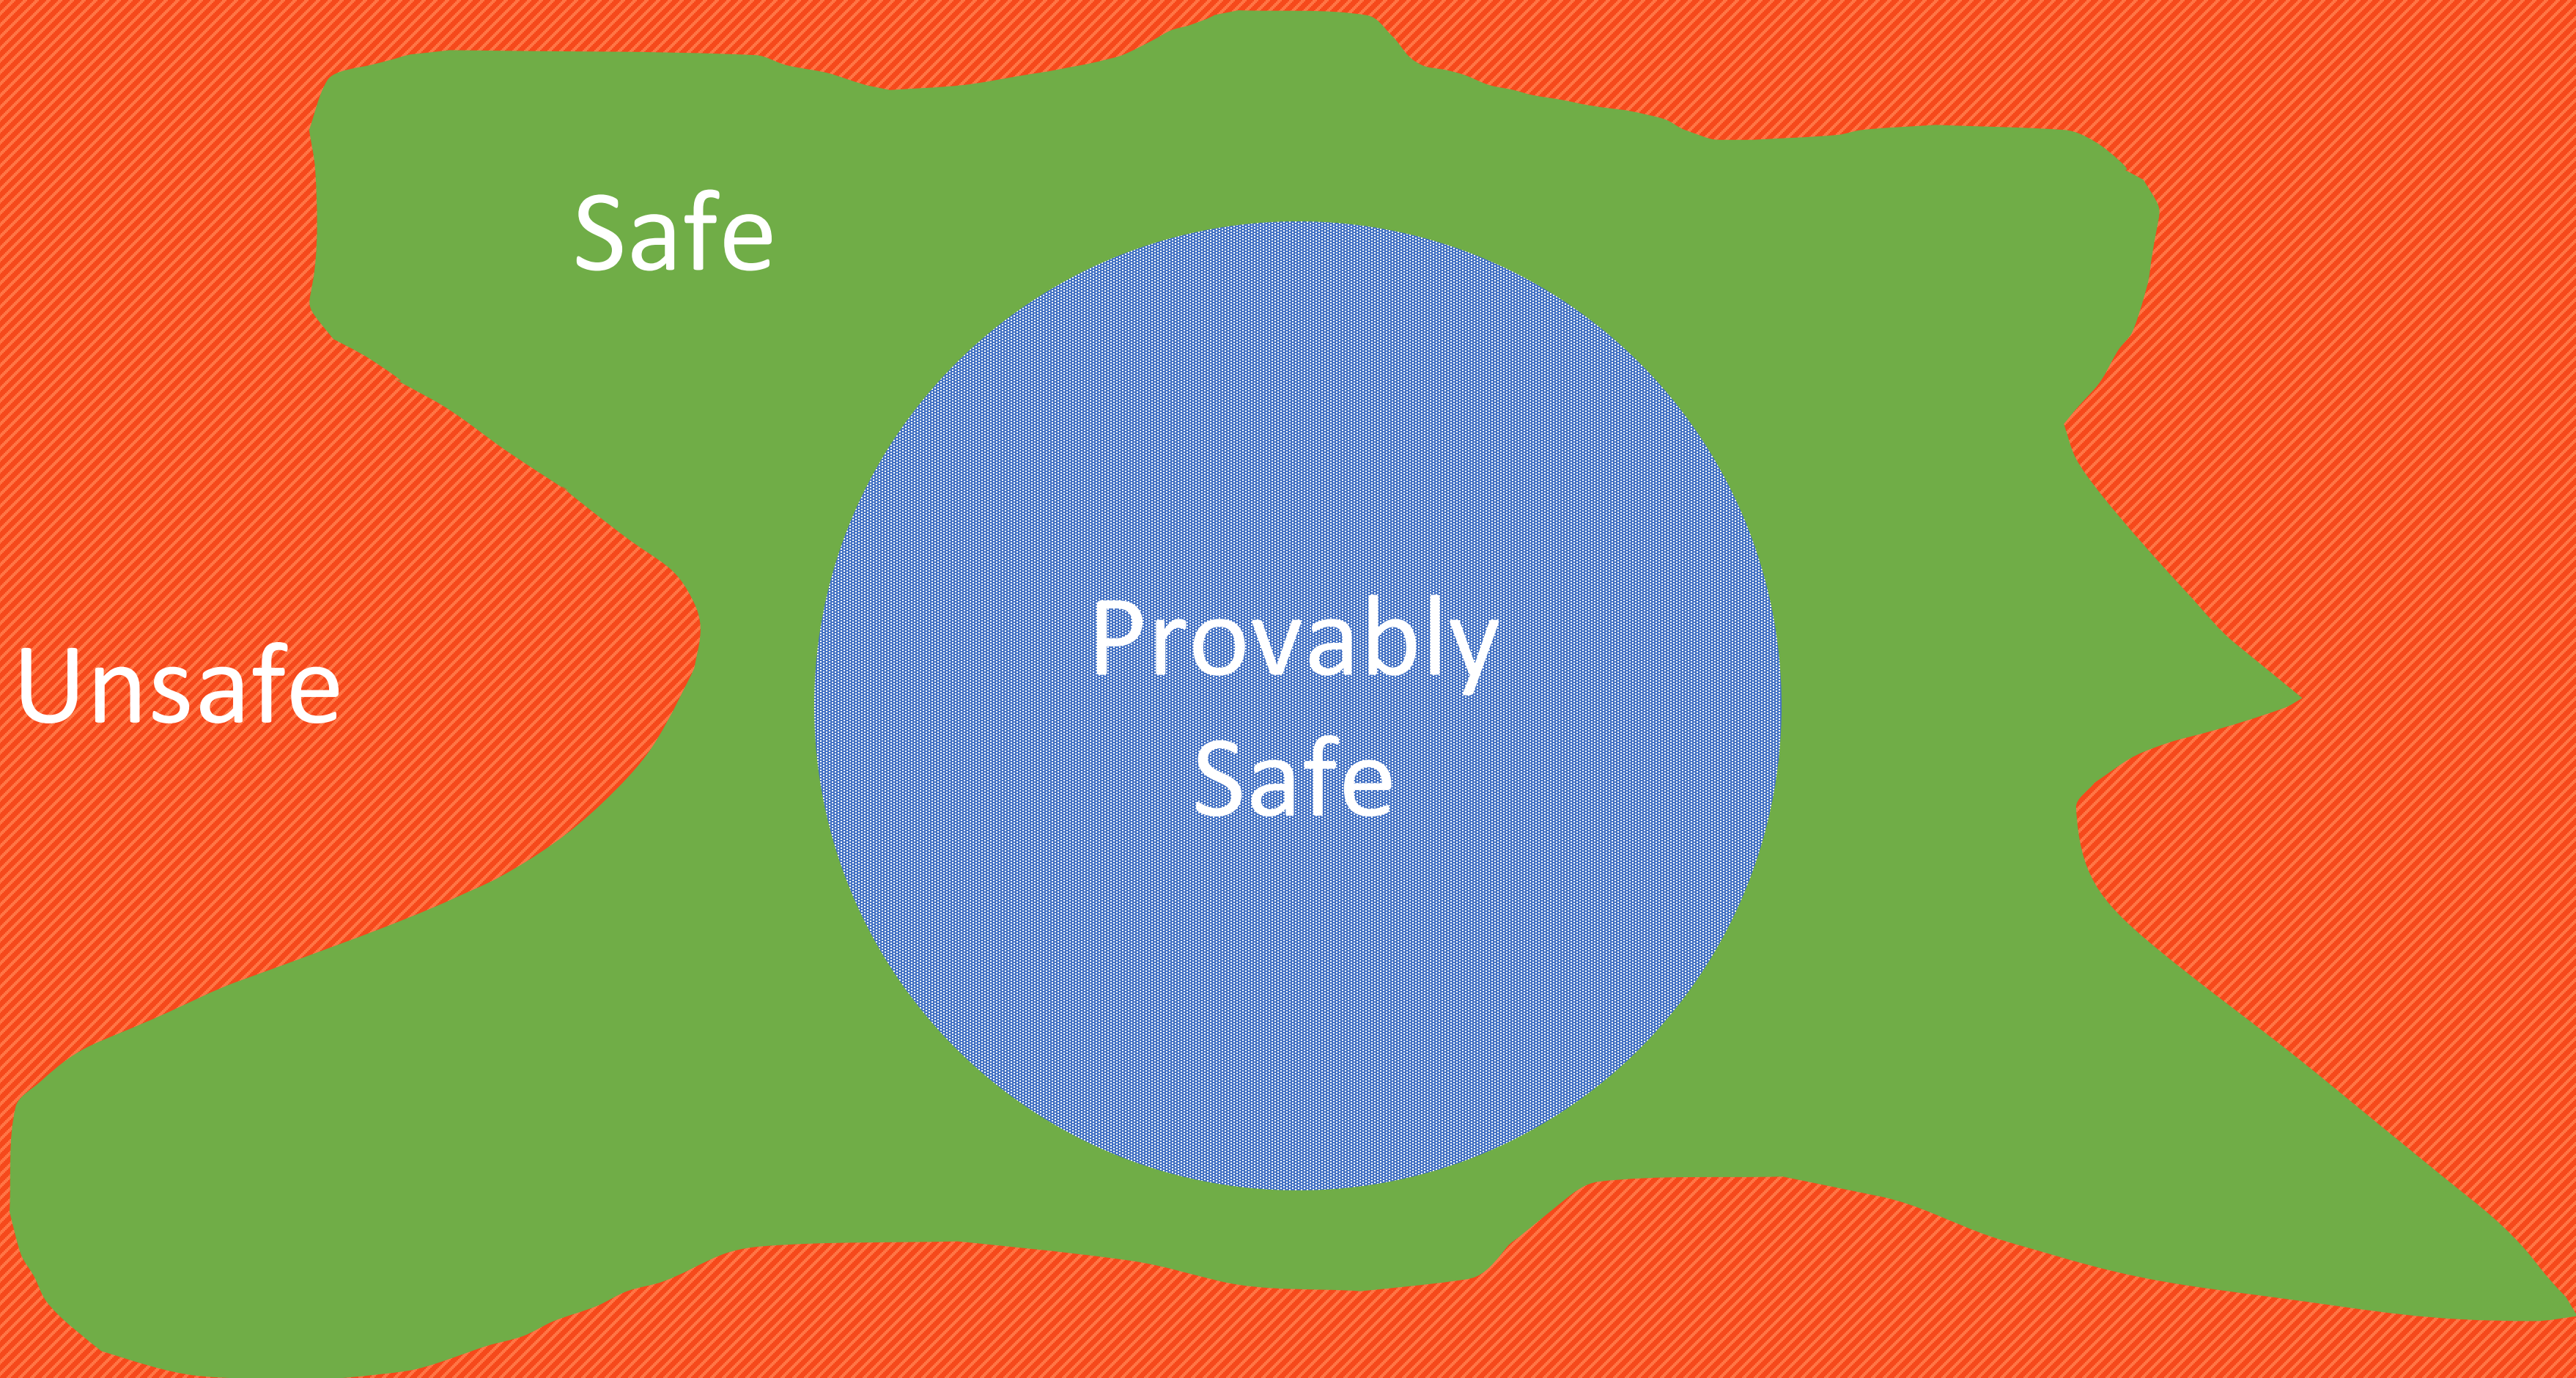


Required AV Reliability

99.99999999999999%

Ensuring a safety envelope






Safe

Unsafe

Provably
Safe

Hazard and Risk Analysis



Safety Cases


Claim: Vehicle stops at stop signs

Argument: Recognizes stop signs and plans accordingly

Evidence: Simulation testing, road testing

Sub-Claims:

- Recognizes stop signs
- Recognizes limit lines
- Adjusts speed to ensure it can stop



Failure Mode and Effects Analysis

- Potential failure mode
- Potential failure effect
- Severity
- Potential causes
- Expected frequency
- Current controls
- Likelihood of detection
- Computed risk priority
- Recommended action

Process Step	Potential Failure Mode	Potential Failure Effect	SEV ¹	Potential Causes	OCC ²	Current Process Controls	DET ³	RPN ⁴	Action Recommended
What is the step?	In what ways can the step go wrong?	What is the impact on the customer if the failure mode is not prevented or corrected?	How severe is the effect on the customer?	What causes the step to go wrong (i.e., how could the failure mode occur)?	How frequently is the cause likely to occur?	What are the existing controls that either prevent the failure mode from occurring or detect it should it occur?	How probable is detection of the failure mode or its cause?	Risk priority number calculated as SEV x OCC x DET	What are the actions for reducing the occurrence of the cause or for improving its detection? Provide actions on all high RPNs and on severity ratings of 9 or 10.
ATM Pin Authentication	Unauthorized access	• Unauthorized cash withdrawal • Very dissatisfied customer	8	Lost or stolen ATM card	3	Block ATM card after three failed authentication attempts	3	72	
	Authentication failure	Annoyed customer	3	Network failure	5	Install load balancer to distribute work-load across network links	5	75	
Dispense Cash	Cash not disbursed	Dissatisfied customer	7	ATM out of cash	7	Internal alert of low cash in ATM	4	196	Increase minimum cash threshold limit of heavily used ATMs to prevent out-of-cash instances
	Account debited but no cash disbursed	Very dissatisfied customer	8	• Transaction failure • Network issue	3	Install load balancer to distribute work-load across network links	4	96	
	Extra cash dispensed	Bank loses money	8	• Bills stuck to each other • Bills stacked incorrectly	2	Verification while loading cash in ATM	3	48	


Source: isixsigma.com

AutoDBA works with your PostgreSQL, wherever you run it today



AWS, Azure, or GCP
Managed PostgreSQL



On-Prem
PostgreSQL



AutoDBA





Thank You

Johann Schleier-Smith

jssmith@crystal.cloud