

SQL SERVER CHAOS ENGINEERING

ANDREW PRUSKI

SQL SERVER DBA & MICROSOFT DATA PLATFORM MVP

 @dbafromthecold

 dbafromthecold@gmail.com

 www.dbafromthecold.com

 github.com/dbafromthecold

SESSION AIM

How can Chaos Engineering be applied to SQL Server?

AGENDA

Identifying SQL Server weaknesses
Running a Chaos Engineering experiment
Testing SQL Server on Kubernetes

IDENTIFYING WEAKNESSES

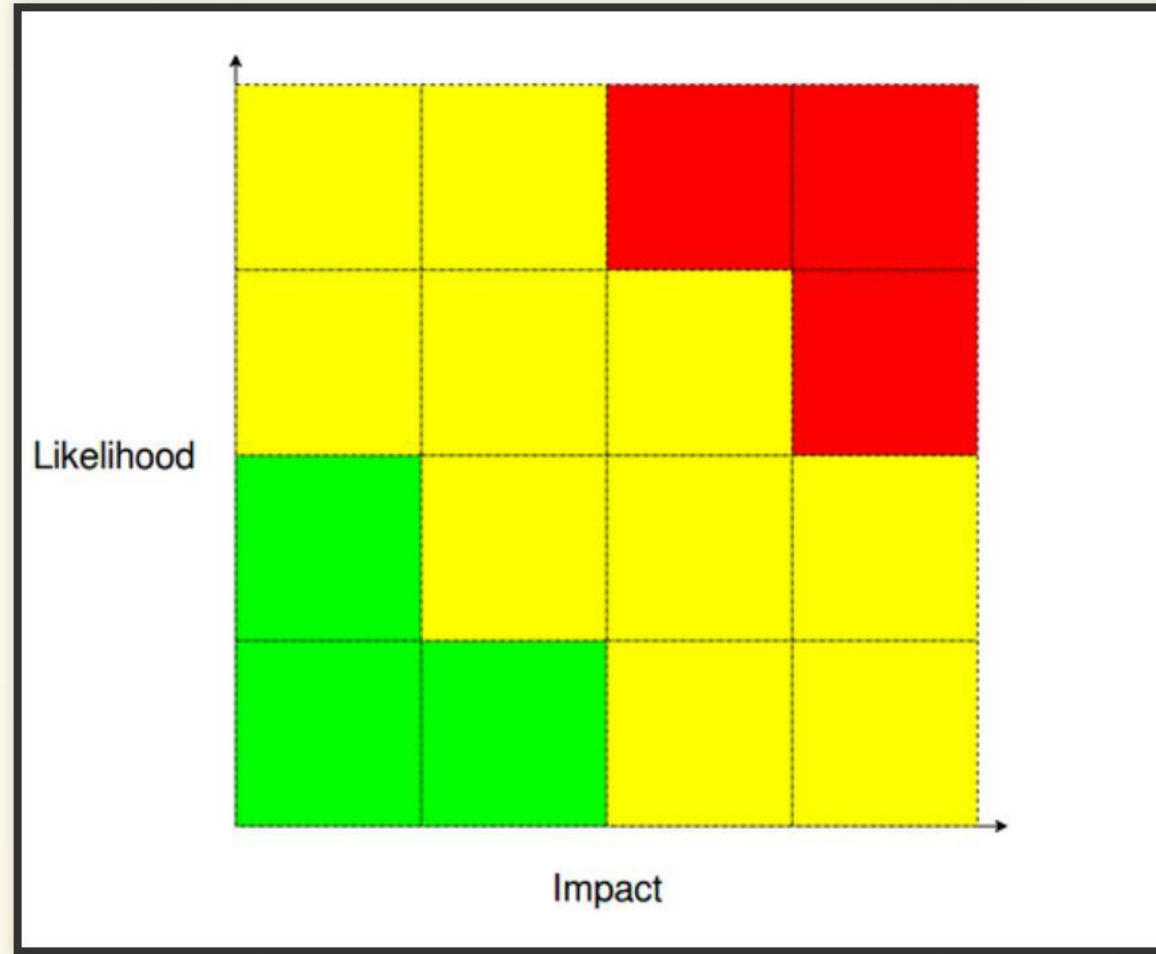
PAST INCIDENT ANALYSIS

How has the system failed previously?

What technologies/strategies are now in place?

What was learnt from those previous failures?

LIKELIHOOD-IMPACT MAP



DEFINING AN EXPERIMENT

Which failure has the highest likelihood?

Which failure has the highest impact?

What will you gain from testing that failure?

Is there anything else that can be tested?

RUNNING AN EXPERIMENT

WHAT FAILURE ARE WE GOING TO
TEST?

WHAT FAILURE ARE WE GOING TO TEST?

- What happens if the primary node in an Availability Group cluster fails?

DEFINING THE EXPERIMENT

DEFINING THE EXPERIMENT

- Hypothesis

DEFINING THE EXPERIMENT

- Hypothesis
 - The listener of the availability group should remain online

DEFINING THE EXPERIMENT

- Hypothesis
 - The listener of the availability group should remain online
- Method

DEFINING THE EXPERIMENT

- Hypothesis
 - The listener of the availability group should remain online
- Method
 - Stop the SQL database engine service on the primary node

DEFINING THE EXPERIMENT

- Hypothesis
 - The listener of the availability group should remain online
- Method
 - Stop the SQL database engine service on the primary node
- Rollback

DEFINING THE EXPERIMENT

- Hypothesis
 - The listener of the availability group should remain online
- Method
 - Stop the SQL database engine service on the primary node
- Rollback
 - Restart the SQL database engine service on the primary node

DEMO

SQL SERVER RUNNING ON KUBERNETES

KUBEINVADERS



RESOURCES

<https://github.com/dbafromthecold/SqlServerChaosEngineering>

<https://github.com/lucky-sideburn/KubeInvaders>

