Gremlin OSC

OSC Korea & Demo

김대성 과장, DevOps Engineer, OSC Korea

Break Things On Purpose

Inject something harmful to
build an immunity.

Gremlin

항체 주입을 통해 면역체계 형성

Microservices

"경쟁 우위 확보를 위해 빠른 혁신 불가피"

Speed of innovation is a competitive advantage

DevOps

"The time has come for IT to face the actual complexity required to really drive its own transformation"

– McKinsey

Cloud

속도는 상응하는 대가를 요구함

Velocity Comes at a Price

Application/Service는 급격한 속도로 복잡해지고 있음 : **운영 난이도 증가**

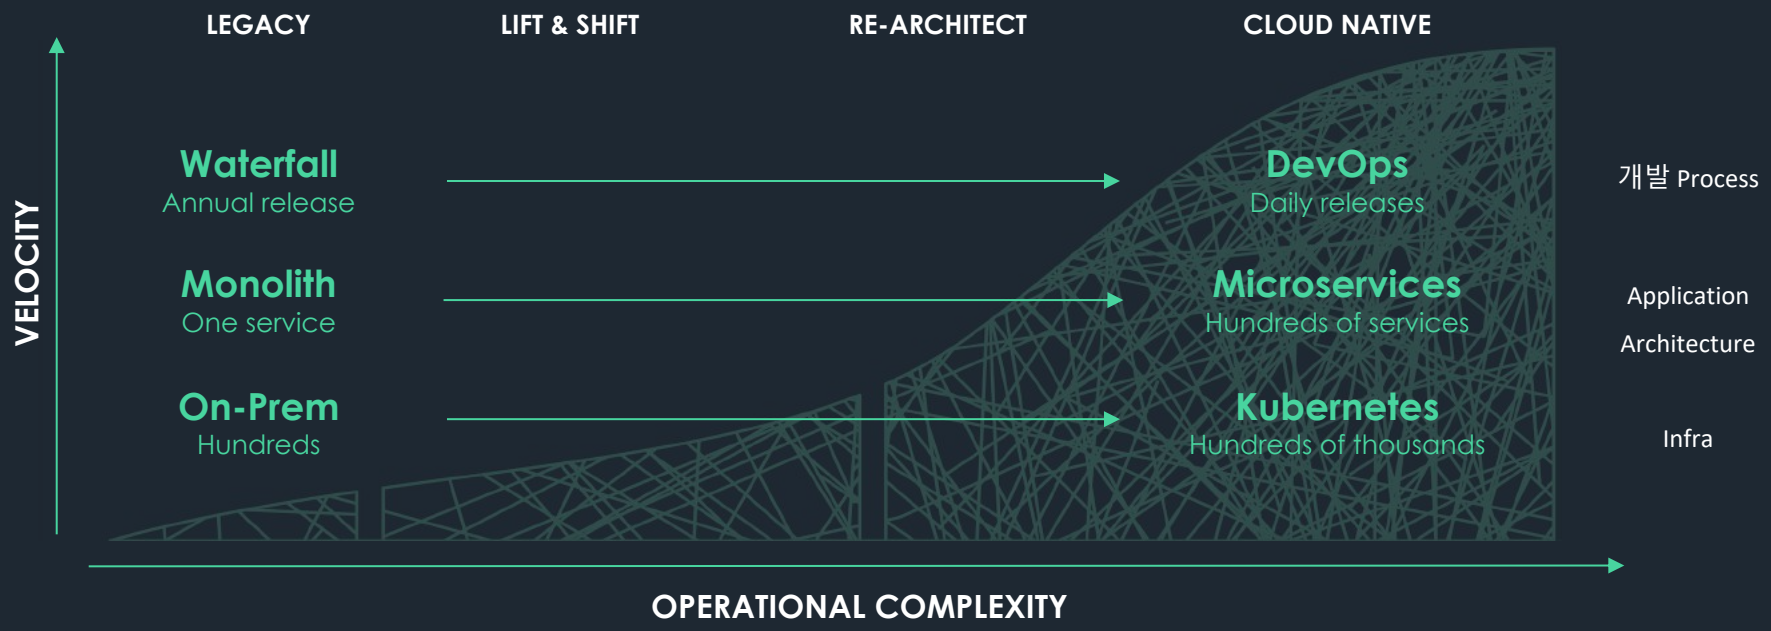

Container기반의 Service를 제공하는 Global 기업 예제

Gremlin OSC

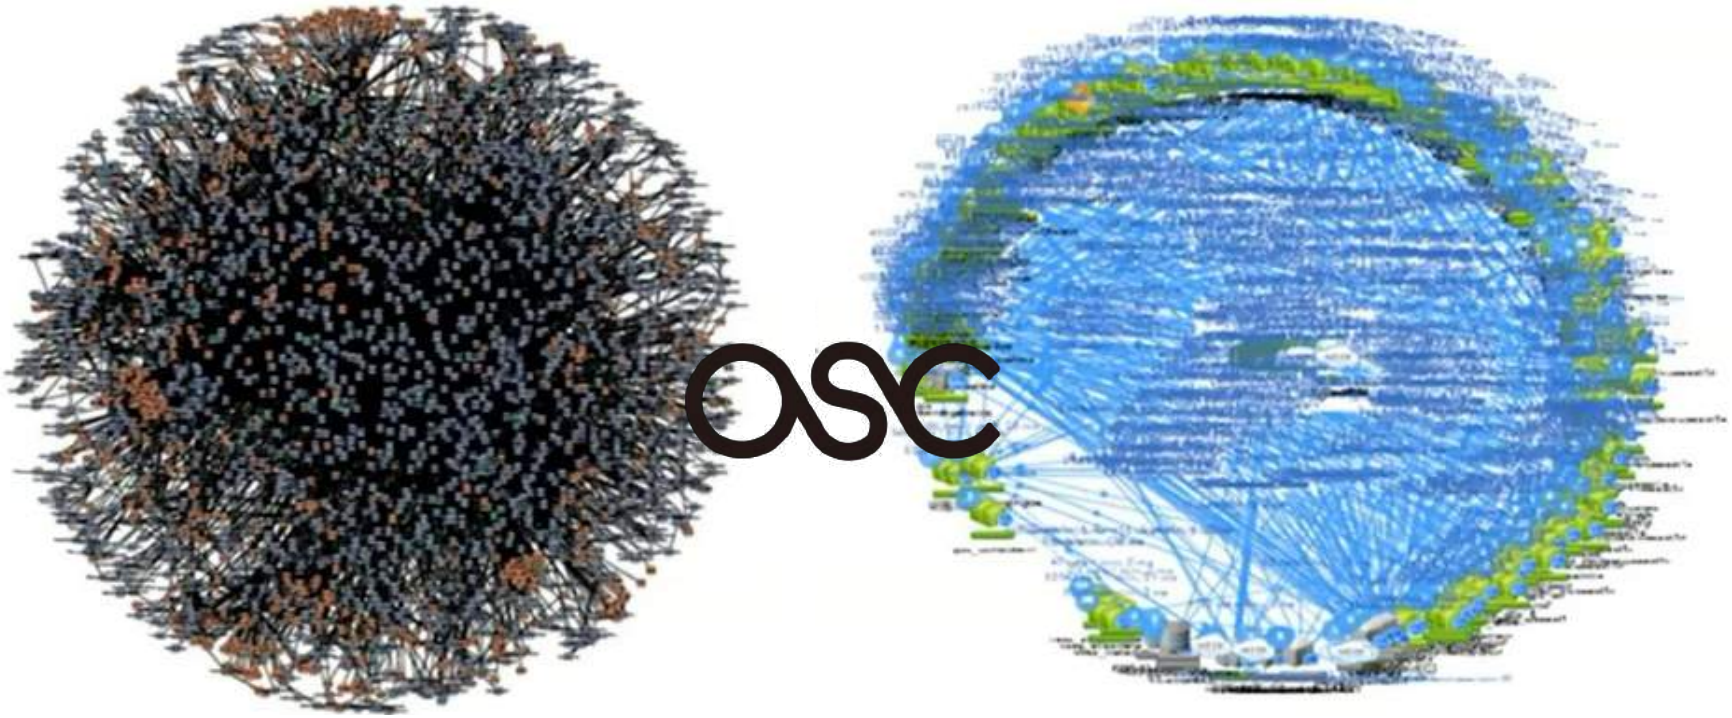

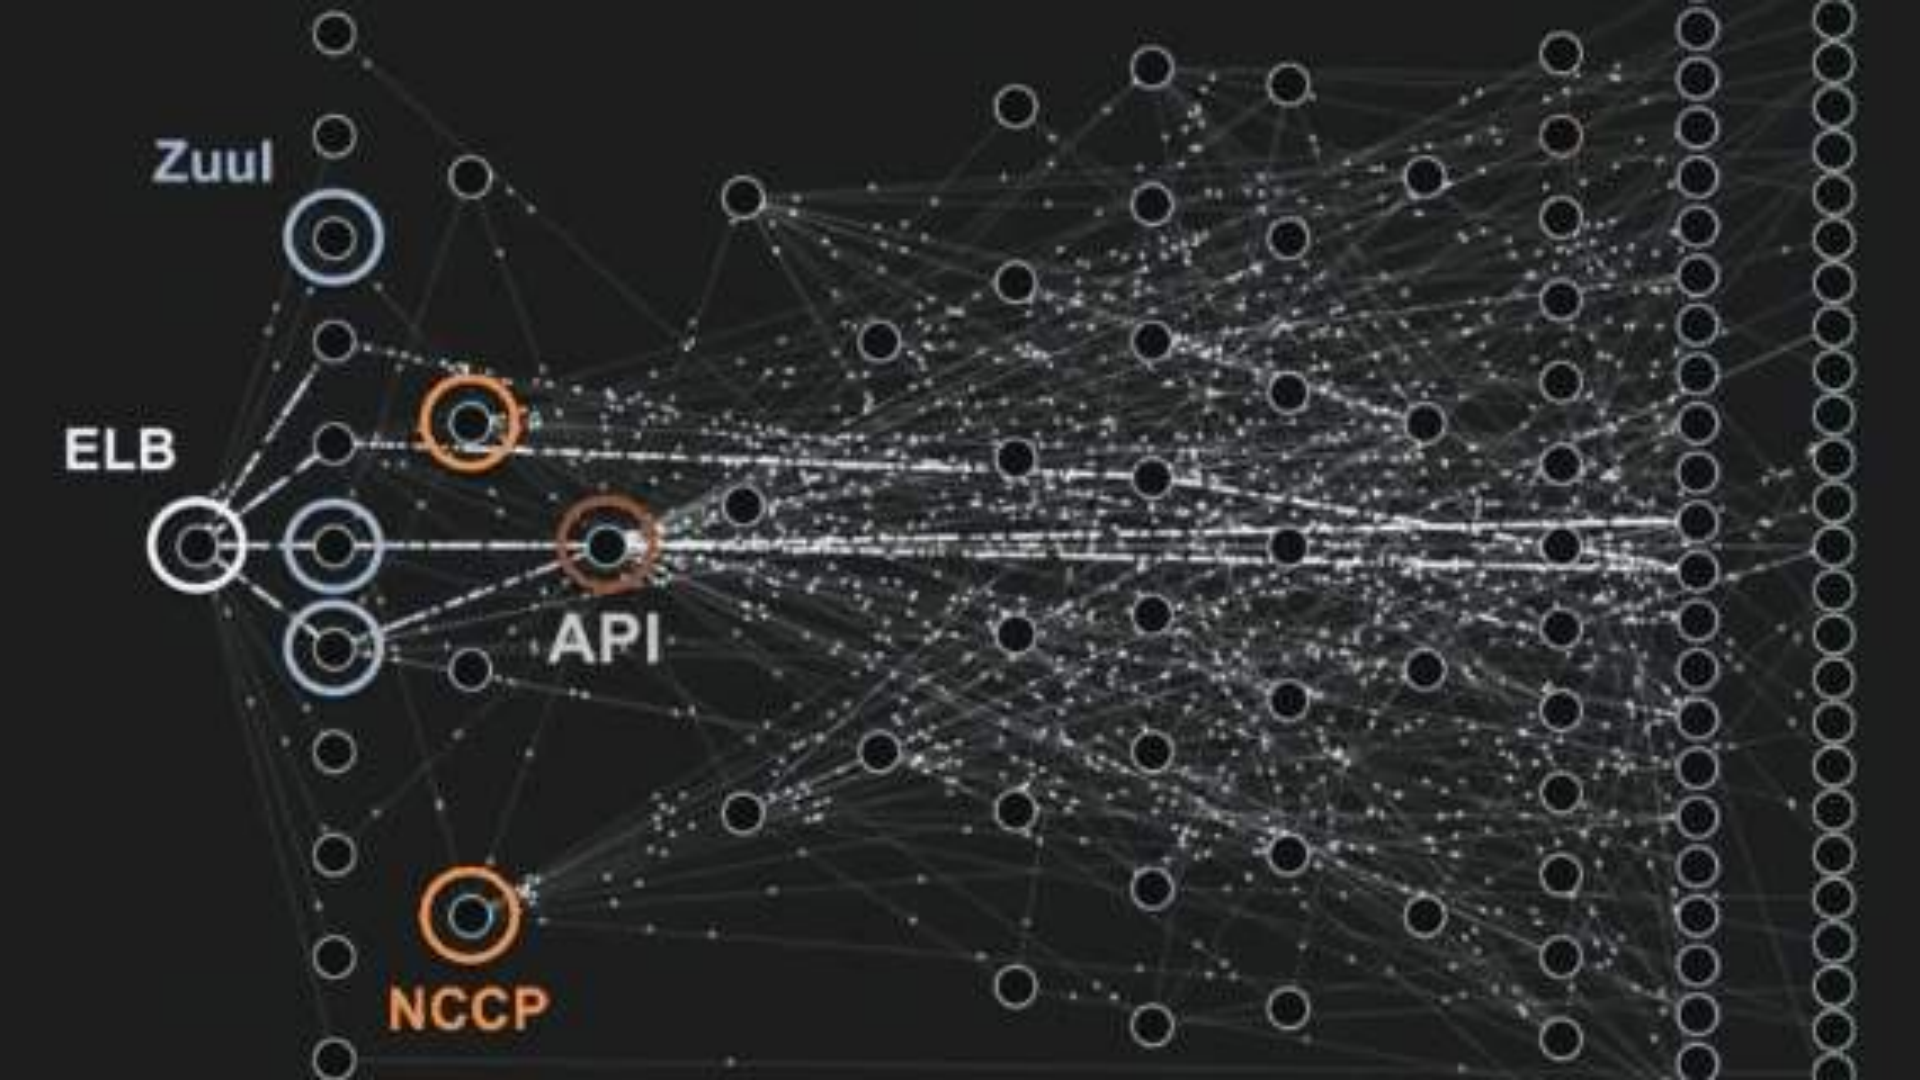
Zuul
ELB
API
NCCP

Dependency Map - Instana View

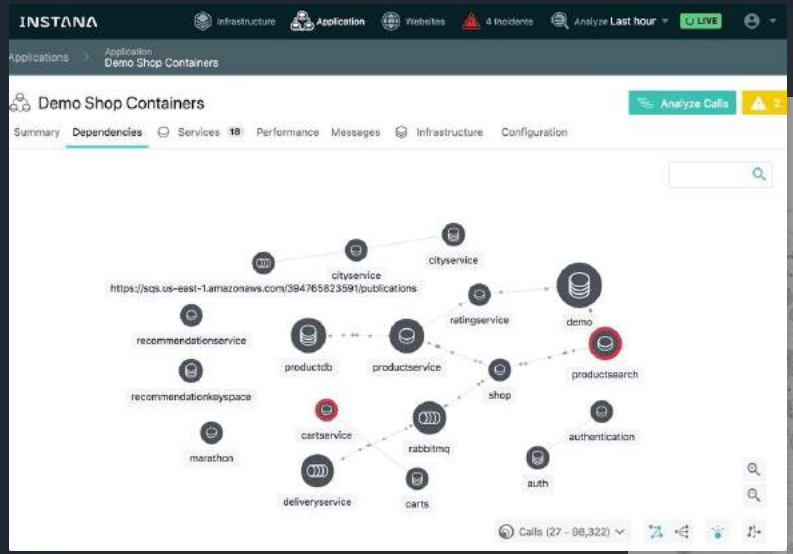

SIT / QA 환경

STG / PRD 환경

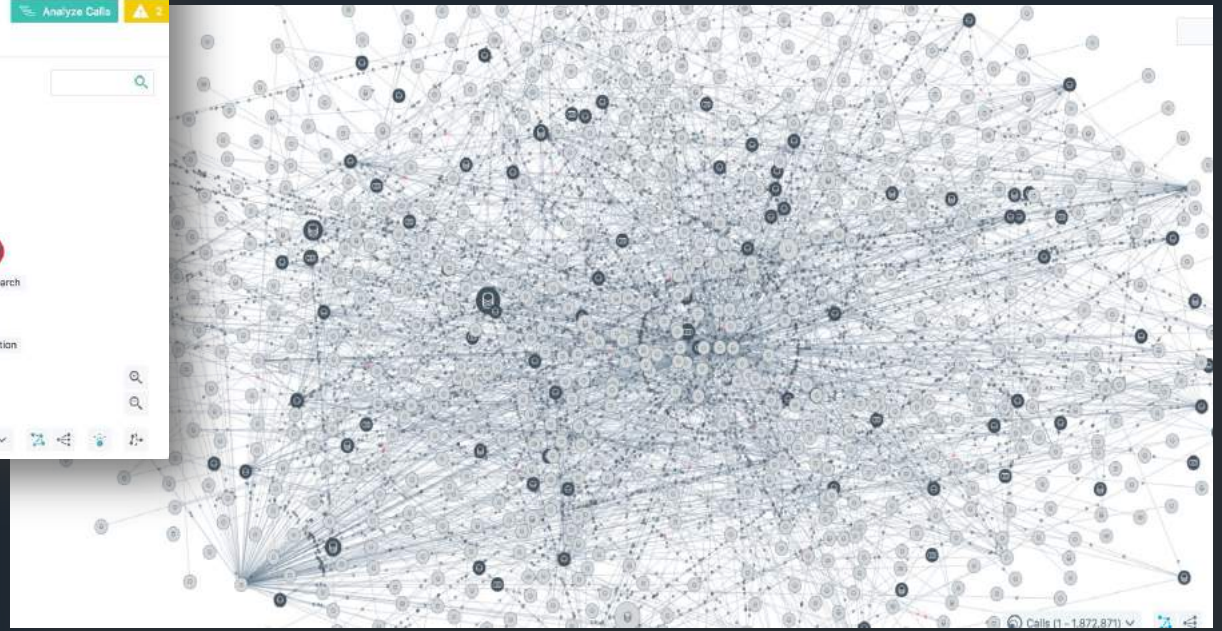

Infrastructure Map - Instana View

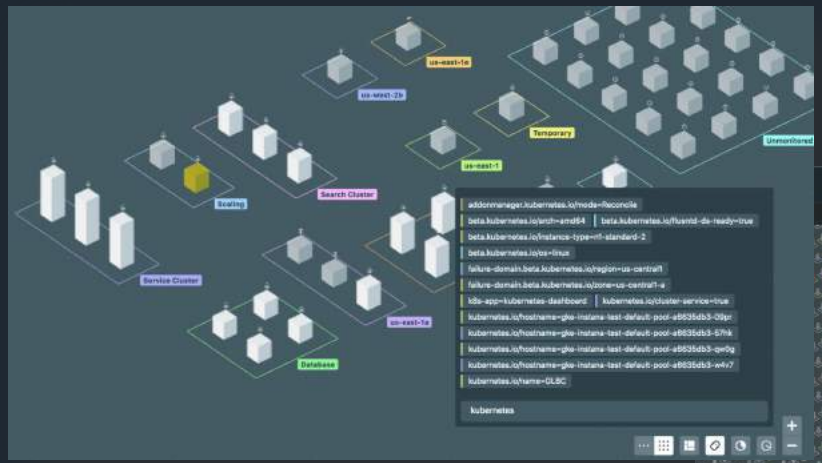

SIT / QA 환경

STG / PRD 환경

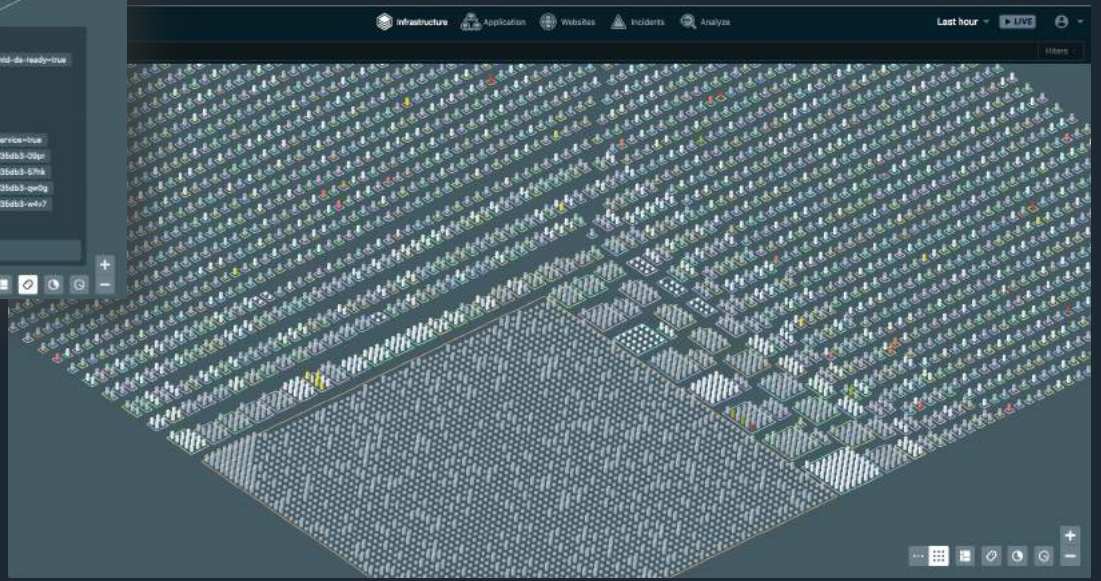

빠른 혁신의 이면 - 장애사례 3가지

Black Friday 장애

Forbes

기술 Issue로 인해 수조원대 손실이 발생한 유통업체

12.01.16

금융 장애

Newsweek

City Bank Web-site, Mobile Service 장애

2.28.19

항공사 장애

npr

System 장애로 인한 항공일정 지연

4.1.19

The Reliability Gap

"잦은 배포에도 품질수준을 유지하기 위한 방안 필요"

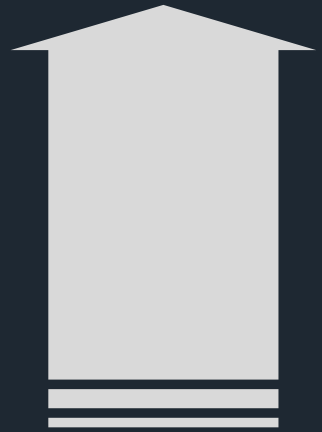

가용성 vs. 변경 속도

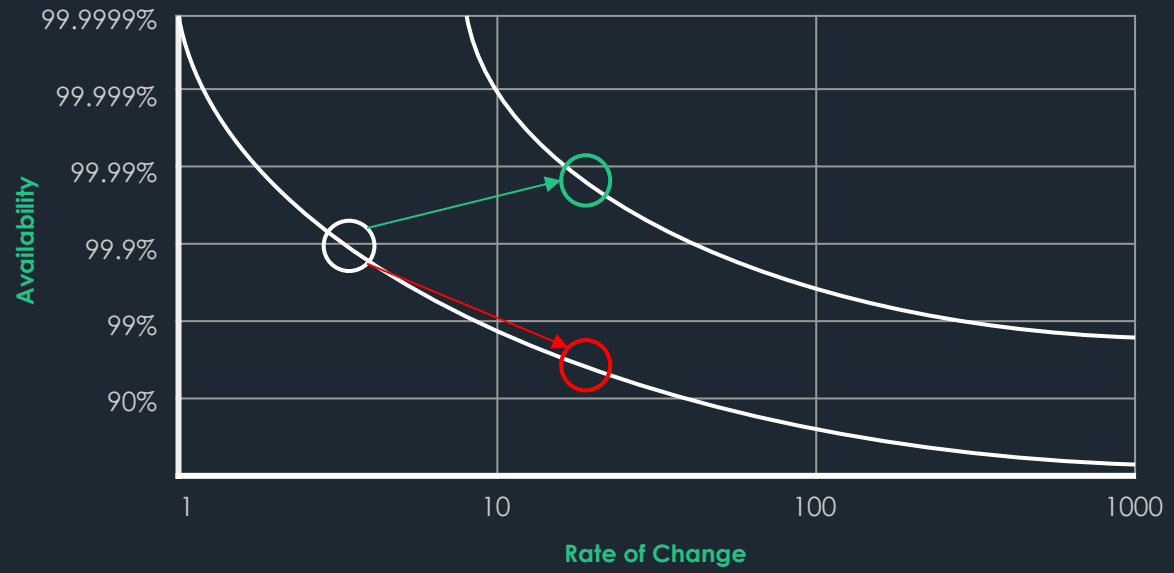

"*Change introduces new forms of failure that are difficult to see before the fact....*"

– Richard Cook, How Complex Systems Fail

Traditional Testing
is not enough
anymore.

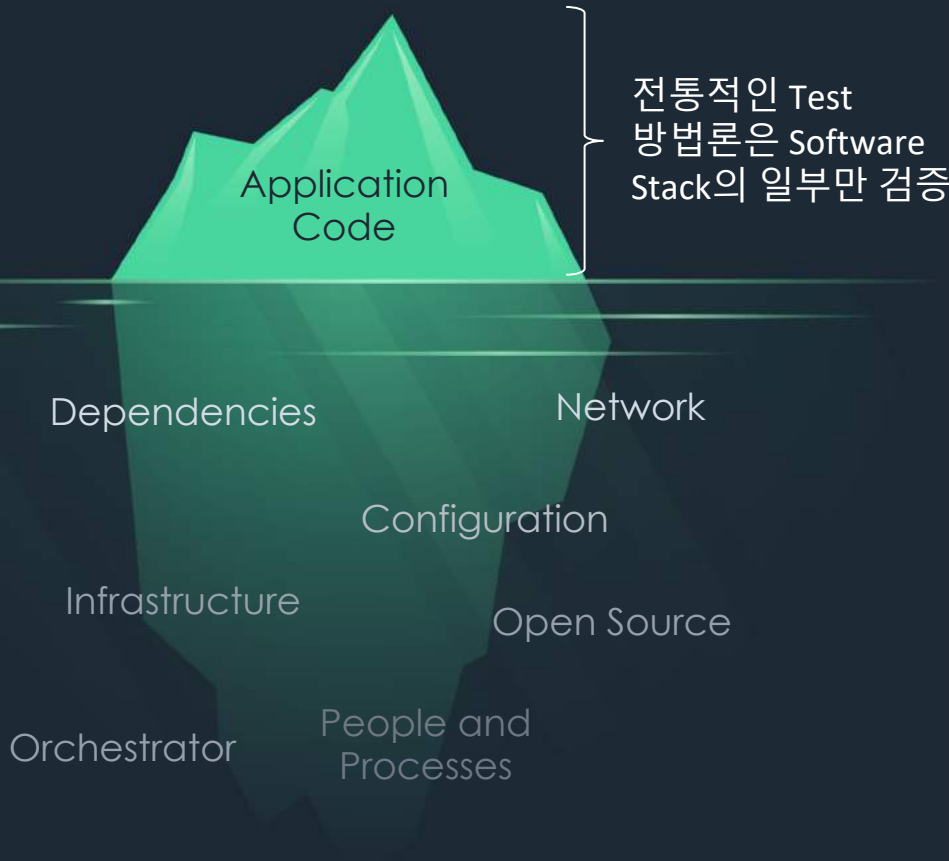

> Chaos Engineering은 System이 어떻게 반응하는지 확인하기 위해 제어 가능한 수준의 장애를 인위적으로 발생시켜 System의 취약점을 찾아내는 것

Chaos Engineering

Thoughtful, controlled experiments designed to reveal the weakness in our systems.

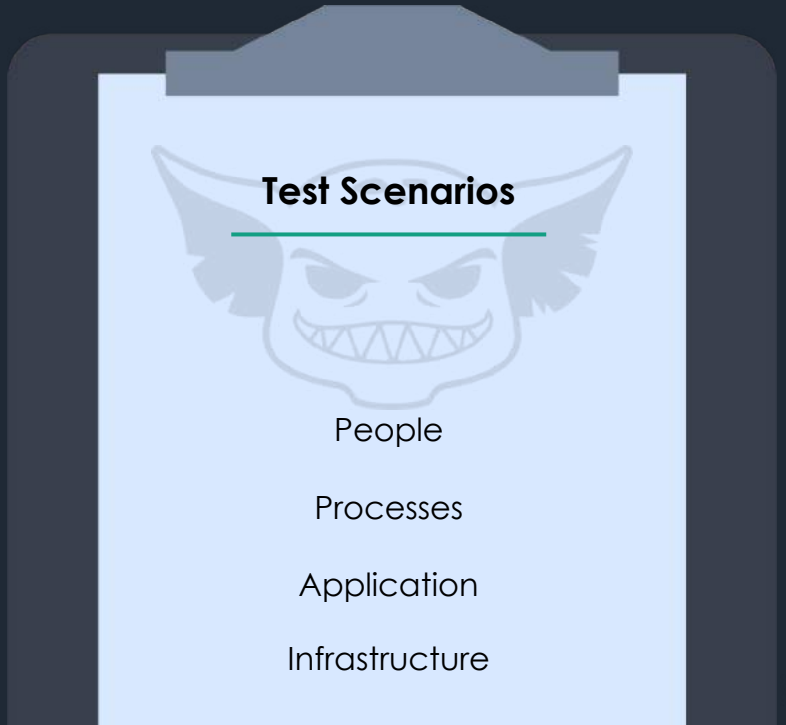

Progressively test your system to isolate problems and mitigate risk

실험 대상을 점진적으로 확대

Chaos Engineering은 선도적으로 Issue를 찾아내며, System의 취약점을 사전에 확인하여 장애 대응력을 향상시킴

기대효과

장애 요소 선제적 조치

Engineers **proactively** test to find and fix issues and limit the impact of failures

장애 대응력 향상

And are more effective when they work **reactively** during an incident

운영환경의 Monitoring & Detection 방식을 세부적으로 튜닝하여 위험요소 감소

> *Chaos Engineering changes the exercise from one of guessing, to one of staging and observing. It helps mitigate risk from emergent failure paths that we literally have no other way to discover.*

Workiva **Matt Simons**
Sr. Engineering Manager

Benefits of Chaos Engineering

장애감소

장애 10배 감소

Dropbox

장애 해결시간 감소

MTTD (Mean Time to Detect)
수시간에서 5분으로 감소

Backcountry

빠른 배포

4배 빠른 K8S Migration

UNDER ARMOUR

성공적인 런치

GameDay 한번으로 35개의 High Priority 버그를 발견하여 수정

media dpg

Chaos Engineering completes DevOps

" 분산System을 안정적으로 운영하는 유일한 방법은 Chaos Engineering을 도입하는 것 "
- Gartner

> *By 2023, 40% of organizations will implement chaos engineering practices as part of DevOps initiatives,* ***reducing unplanned downtime by 20%.***
>
> –Gartner

"2023년까지 40%의 조직이 Chaos Engineering을 DevOps의 일환으로 도입하여 Downtime을 20% 감소시킬 것" - Gartner

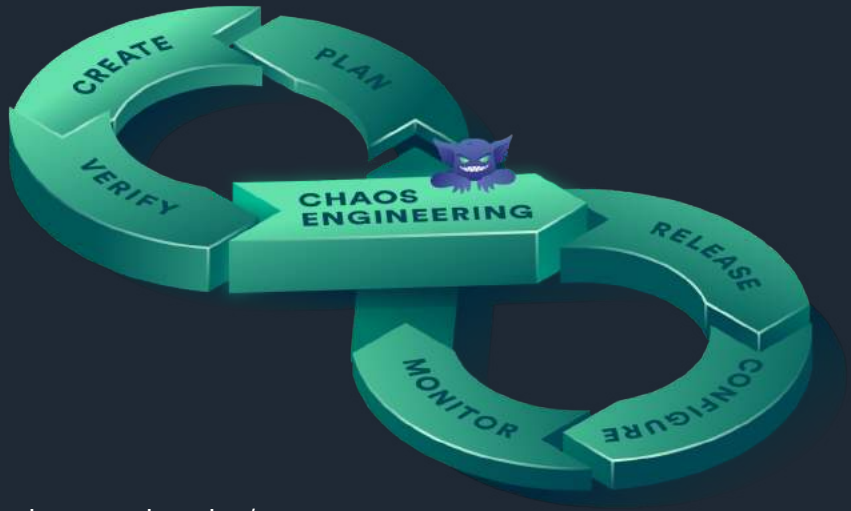

https://www.gartner.com/smarterwithgartner/the-io-leaders-guide-to-chaos-engineering/

Chaos Engineering Diagram (from Chaos Engineering Slack Community)

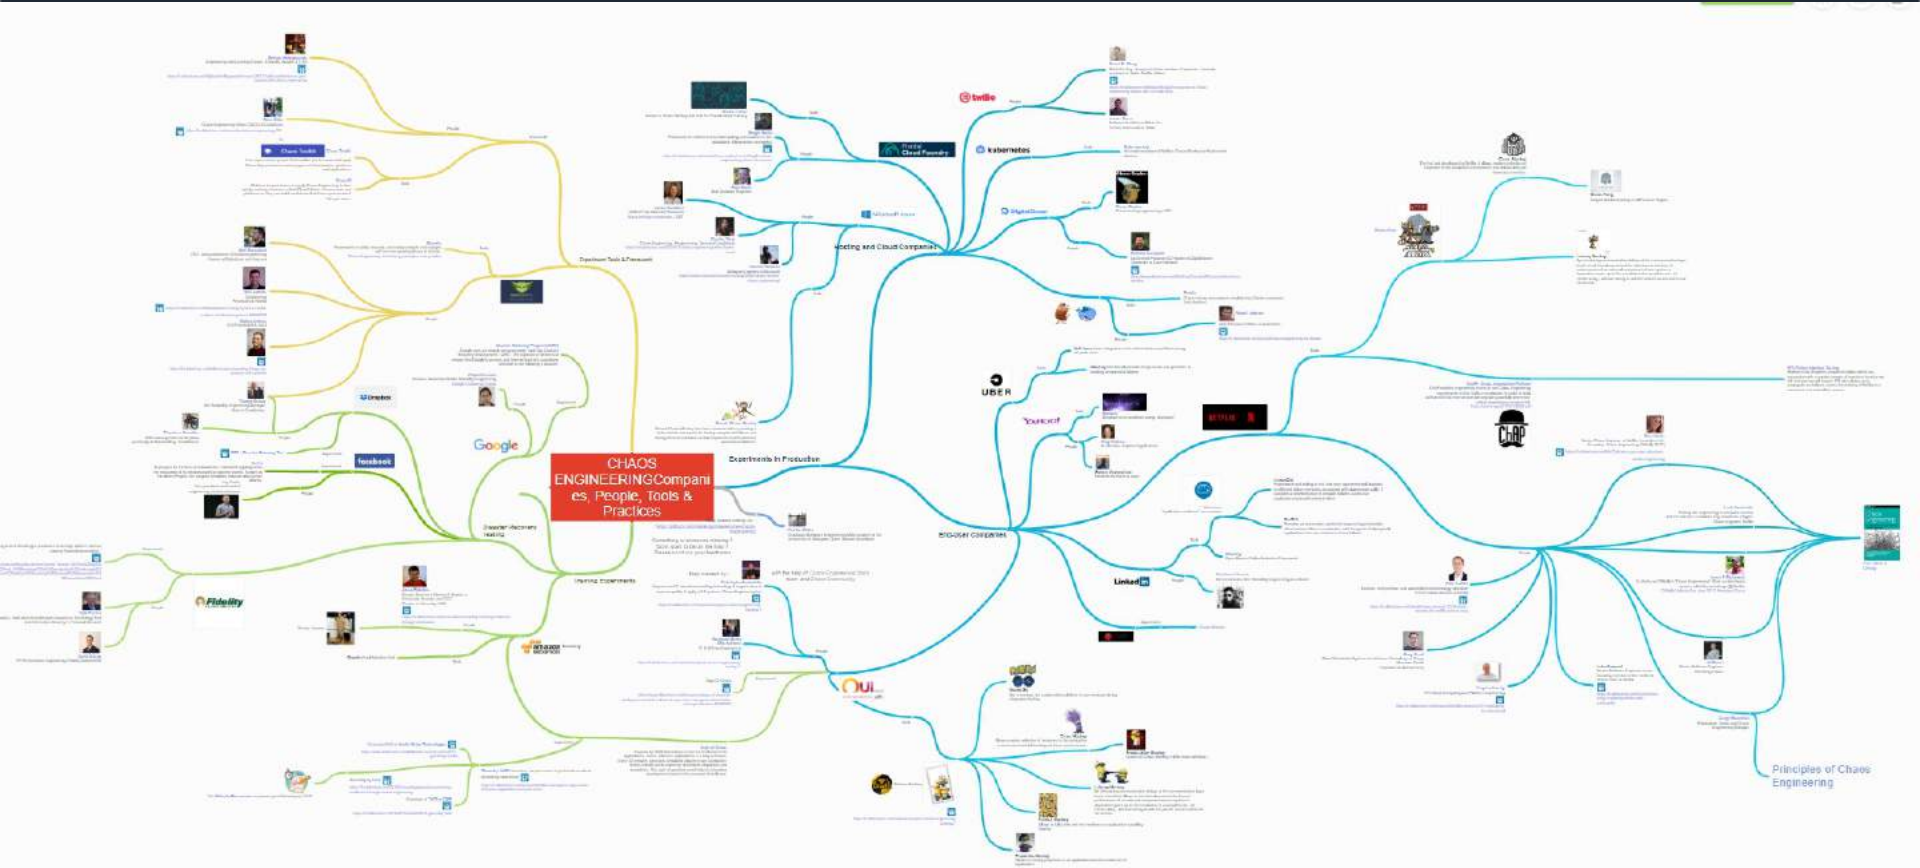

https://coggle.it/diagram/WiKceGDAwgABrmyv/t/chaos-engineeringcompanies%2C-people%2C-tools-practices/0a2d4968c94723e48e1256e67df51d0f4217027143924b23517832f53c536e62

Chaos Engineering Platform

Gremlin

Gremlin은 Amazon 및 Netflix에서 Chaos Engineering을 선도하던 전문가들에 의해 설계됨

Kolton Andrus
CO-FOUNDER & CEO

Chaos Engineer & Call Leader
Netflix & Amazon

Matthew Fornaciari
CO-FOUNDER & CTO

Senior Platform Engineer
Salesforce & Amazon

We've done this before.

NETFLIX amazon salesforce

주요 고객사

Trusted By Teams Worldwide

Across Retail, Financial Services, High Tech, Media, Healthcare, and more

Why Gremlin?

"Failure as a Service"

복잡한 System의 장애요소를 찾아내는 경험과 전문성에 기반한 Solution

Simple

단순한 Interface와 정리된 API 문서 제공

Safe

안전한 장애 유발 및 해소, Roll Back 지원

Secure

SOC II 인증
RBAC, MFA, SSO

Comprehensive

다양한 Stack Layer에 대한 Attack 지원

Gremlin Architecture & Use Cases

Expedia | TARGET | UNDER ARMOUR | NIKE | Workiva | qualtrics

Cloud 전환 | Team 교육 | Cloud Native 전환 | Monitoring/Tool 검증 | 의존성/상관관계 검증 | DR 검증

Gremlin

Simple Interface
UI | CLI | API

Orchestration
Schedule | Scenarios

Safety
Magnitude | Blast Radius | Halt Button

Security
MFA | SAML | Google SSO | RBAC | Access Logs

Attacks
Network | Resources | State | Application

VMware | AWS | Azure | Google | Docker | K8S | Linux | Windows | Serverless

Gremlin 구성도

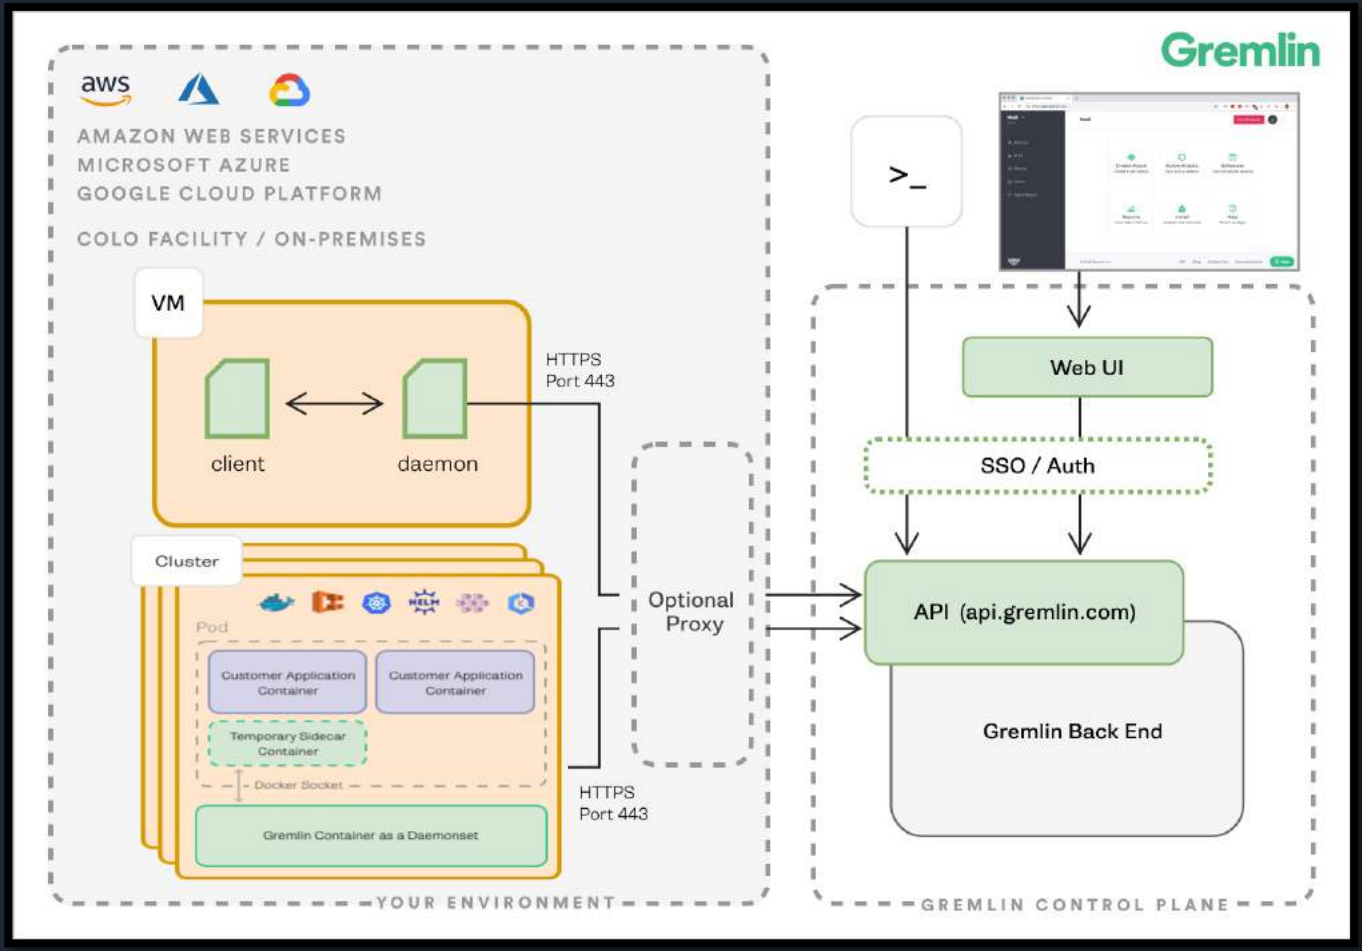
Gremlin
aws
AMAZON WEB SERVICES
MICROSOFT AZURE
GOOGLE CLOUD PLATFORM
COLO FACILITY / ON-PREMISES
VM
client
daemon
HTTPS
Port 443
Cluster
Pod
Customer Application Container
Customer Application Container
Temporary Sidecar Container
Docker Socket
Gremlin Container as a Daemonset
HTTPS
Port 443
Optional Proxy
>_
Web UI
SSO / Auth
API (api.gremlin.com)
Gremlin Back End
YOUR ENVIRONMENT
GREMLIN CONTROL PLANE

Dashboard

Scenarios

Attack

Schedule

Client

보고서

문서

API

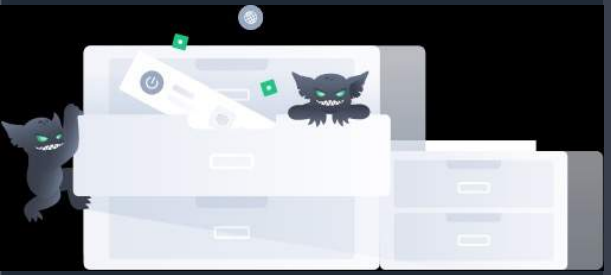

Gremlin

Dashboard

Halt All Attacks

ENTERPRISE Client Version: 2.15.2 - 2.15.3

SE

JK TI

+ Invite your team members

Welcome back, Thomas!

Link your attacks together

Create Scenario

Want to recreate a failure?

View Recommended Scenarios

Quickly run an attack

Create Attack

Targets

Provisioned: 50/day

View All

Last Attacked: 5 hours ago

ip-172-31-61-227.ap-northeast-2.compute.internal

ACTIVE

Attack

Last Attacked: 5 hours ago

REVOKED 172.31.10.86

REVOKED 172.31.13.165

Total Attacks Run

This week

Latency 17

Blackhole 9

CPU 8

Recent Scenarios

Latency
Last run: 7/16/2020

CPU
Last run: 6/22/2020

Memory
Last run: 6/22/2020

Latency
Last run: 6/4/2020

Recent Attacks

Latency
7/17/2020 1:28 pm
SUCCESSFUL

Latency
7/17/2020 1:25 pm
SUCCESSFUL

CPU
7/17/2020 11:19 am
SUCCESSFUL

CPU
7/17/2020 11:16 am
SUCCESSFUL

지원

Attack Setup 절차

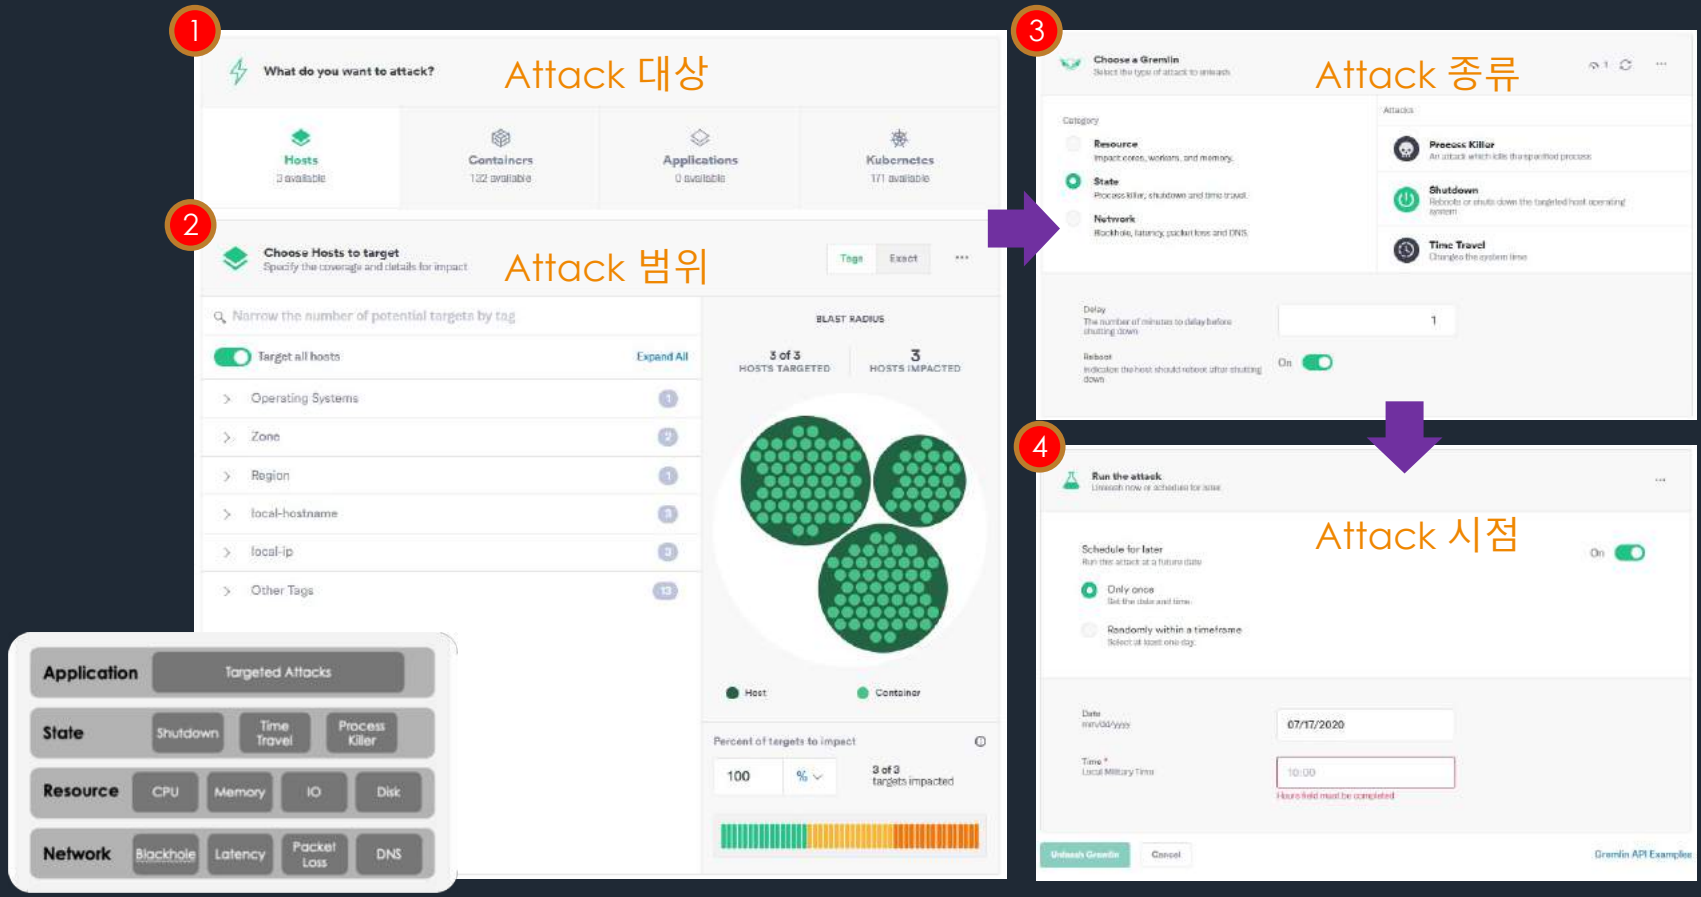

Scenarios 지원

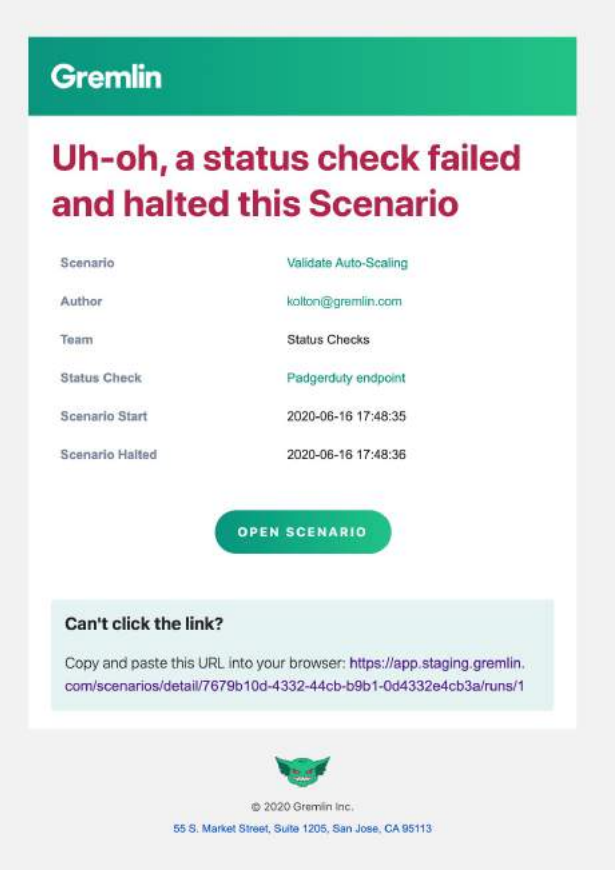

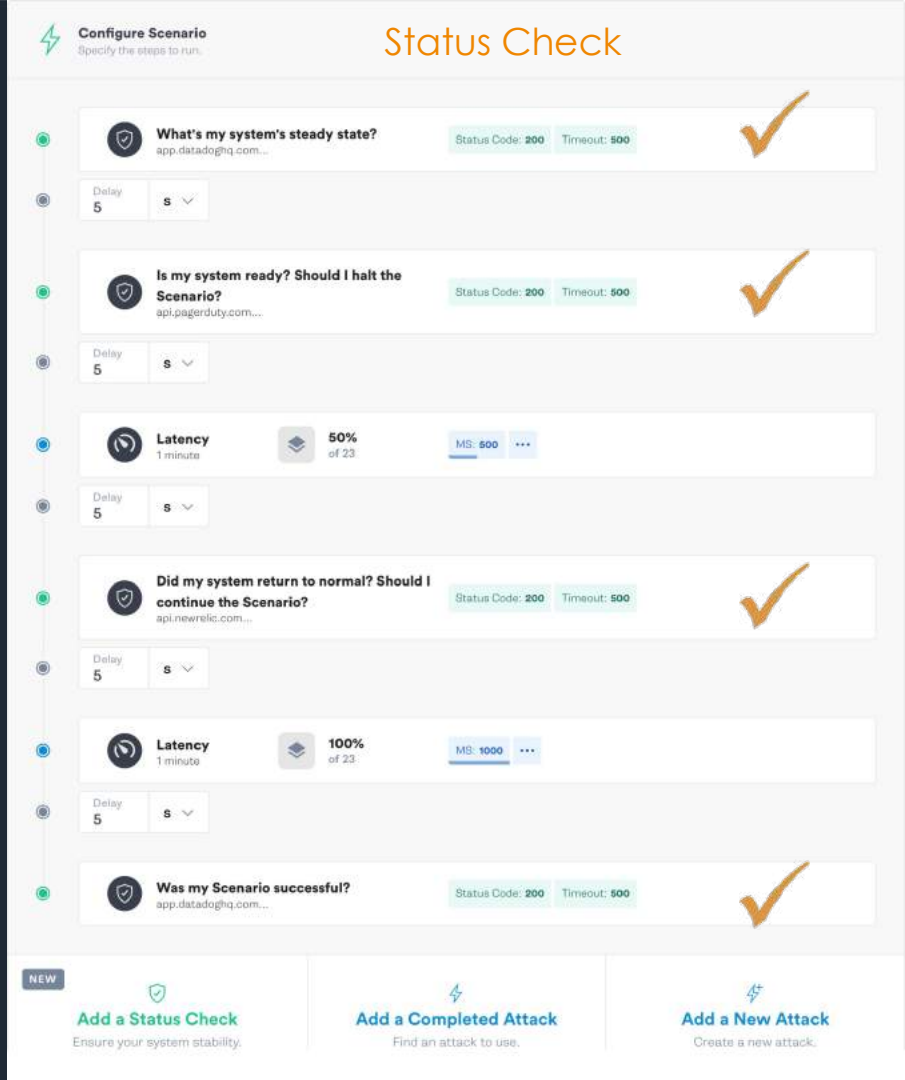

Application Level Fault Injection (ALFI)

Request 기반 Metadata 활용

Customer ID, Device ID, Country등을 이용해 **Attack Radius**를 제어

System 접근이 어려운 경우

AWS Lambda등 **Serverless** 환경에서 Application에 장애 삽입 메커니즘 구현 (e.g. Run in the JVM)

응용 사례

특정 Customer ID에 대한 적용되는 Attack을 생성하여 Monitoring 하여 다른 사용자는 영향을 주지 않도록 구성

일부 특정 Endpoint에만 장애를 유발하고 나머지 Endpoint는 정상인 환경 구성

Always-on 장애 환경을 관리하에 있는 특정 Device에 대해서만 구성하여 정기적으로 장애영향 분석

Application 수준의 장애 주입 설정

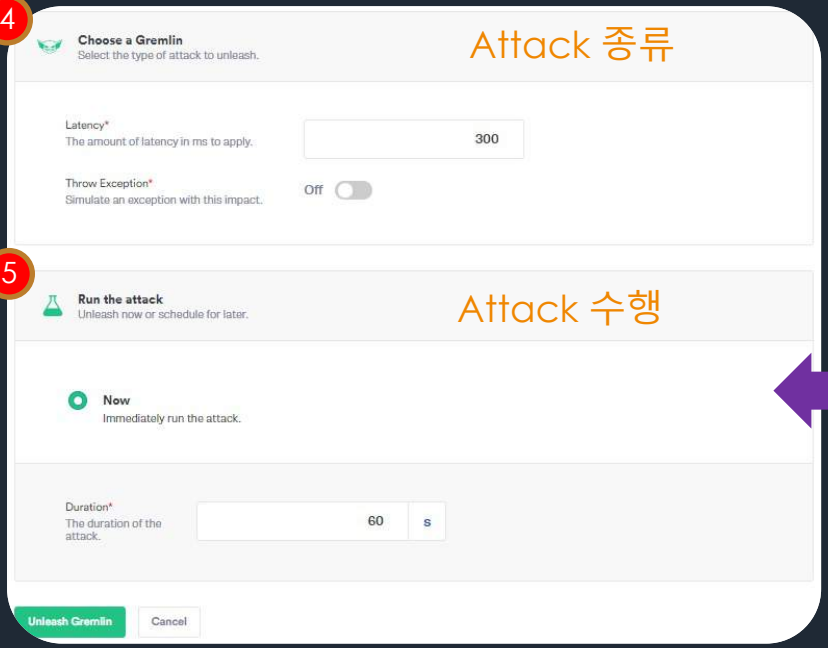

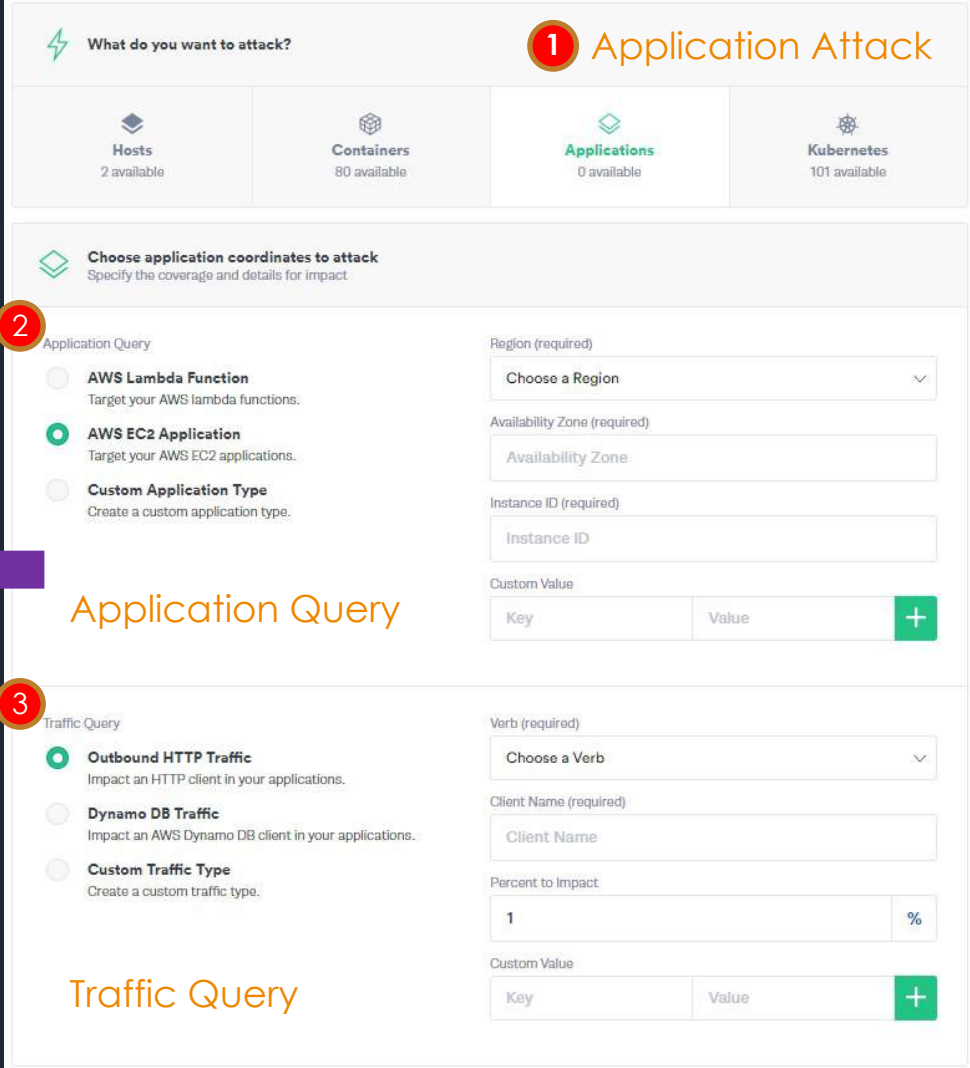

Gremlin Pricing Plan

FREE

FOR GETTING STARTED

Familiarize yourself with Chaos Engineering by safely experimenting with a limited Blast Radius.

- 1 Team
- 1 Installed Agent
- Unlimited Attacks per Target
- 2 Attacked Targets per month
- Community support

STARTER

FOR SINGLE APPLICATIONS

First time users running GameDays on a single application. Unlocks all experiment types for one Team.

- 1 Team
- 5 Installed Agents
- Unlimited Attacks per Target
- 5 Attacked Targets per month
- Community support

PRO

FOR SMALL ORGANIZATIONS

For small organizations with a limited number of Teams and applications. Adds security controls and compliance.

- 3 Teams
- Unlimited Agents
- Unlimited Attacks per Target
- 30 Attacked Targets per week
- Basic 8-8, M-F support

ENTERPRISE

FOR MID-LARGE ORGANIZATIONS

Mid and large organizations with multiple Teams running and automating experiments. Includes security controls and production support.

- 5 Teams
- Unlimited Agents
- Unlimited Attacks per Target
- 50 Attacked Targets per day
- 2 Large scale experiments
- Dedicated 24/7 support

Add-on packages

STARTER

- 5 additional Attacked Targets per month

1 Add-on maximum

PRO

- 1 additional Team
- 10 additional Attacked Targets per week

ENTERPRISE

- 1 additional Team
- 10 additional Attacked Targets per day

Note) Application Attack, SAML, RBAC, Audit Trail은 Pro & Enterprise Version에서만 지원

Feature Comparison

	FREE	STARTER	PRO	ENTERPRISE
Linux/Windows	✓	✓	✓	✓
Containers	✓	✓	✓	✓
Kubernetes	✓	✓	✓	✓
Application/Serverless	✗	✗	✓	✓
State Attack	1	3	3	3
Resource Attack	1	4	4	4
Network Attack	1	4	4	4
Large Scale Exp.	✗	✗	✗	2
SAML	✗	✗	✓	✓
RBAC	✗	✗	✓	✓
Audit Trail	✗	✗	✓	✓

Plan별 적용 예

Service A

STARTER

DEV/STG

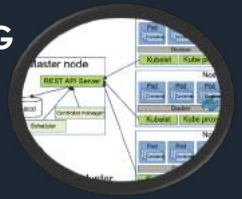

PRD

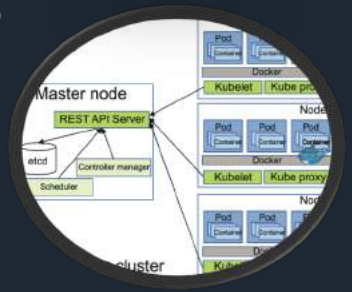

Service B

DEV/STG

PRD

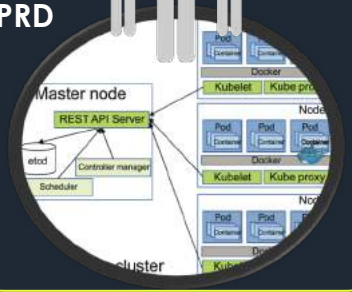

PRO

ENTERPRISE

Service C

DEV/STG

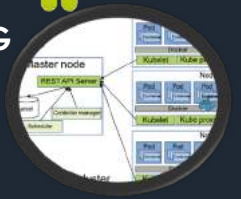

PRD

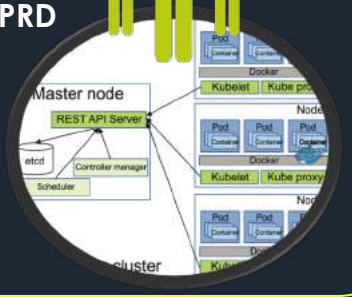

Pricing FAQs

- License 정책 : 최소 계약기간 12개월
- Plan 변경 : Upgrade 가능
- Team별 최대 사용자 : 제한 없음
- Target 당 Attack 수 제한 : Attack 횟수 및 Attack Type 에 대한 제한 없음
- Attacked Target 산정 : 첫번째 Attack이 Target에 대해 이루어지면 Attacked Target으로 기록됨
- Attacked Target 수 제한 : Plan에 따라 월별/주별/일별 Unique한 Target 수 한정
- Team간 Attacked Target 수 조정 : 계약기간내 변경 불가 (Enterprise Plan만 분기별 재조정 가능)
- ALFI 산정 : Query 하나당 하나의 Attacked Target으로 산정 (Query내 Traffic Query는 산정 하지 않음)

> *The best way to avoid failure is to fail constantly*

> *Chaos doesn't cause problems. It reveals them*

Gremlin OSC

OSC Korea Demo

OSC Korea는
디지털 트랜스포메이션
전문기업입니다.

고객에게 혁신적이고 이상적인
모던 아키텍처와 최적화된 기술을
제공합니다.

한국 리눅스 재단 운영

- 국내 오픈소스 생태계 활성화 지원
- CNCF, Cloud Foundry, Hyperledger, LFEdge, LFAI, LFPH 등 후원

MSA(Microservice) 전문기업

- 오픈소스 기반 아키텍처 컨설팅 및 구축
- 다양한 MSA 방법론에 기반한 Modern Architecture 설계 지원

글로벌 솔루션 공급

- 디지털 트랜스포메이션에 필요한 글로벌 솔루션 발굴
- 고객 요구사항 및 시장 환경에 부합하는 최적의 솔루션 통합 제공

Gremlin OSC

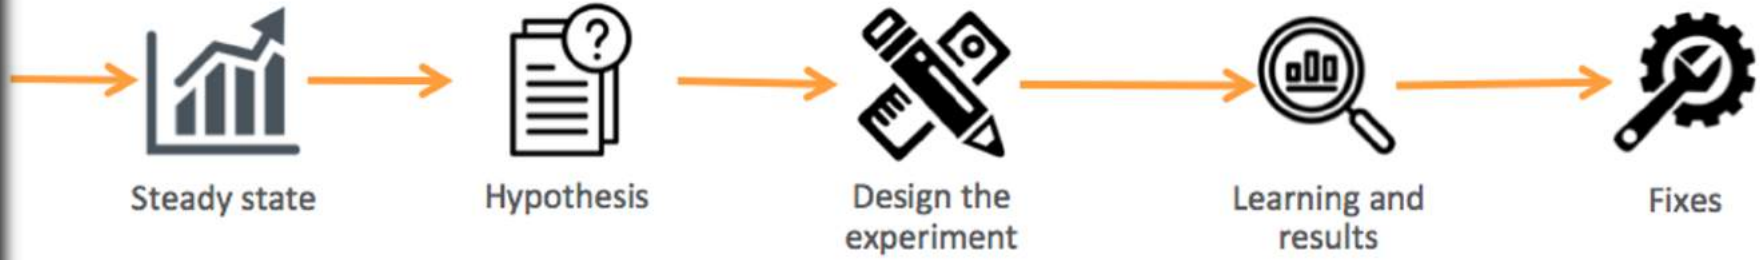
Steady state
Hypothesis
Design the experiment
Learning and results
Fixes

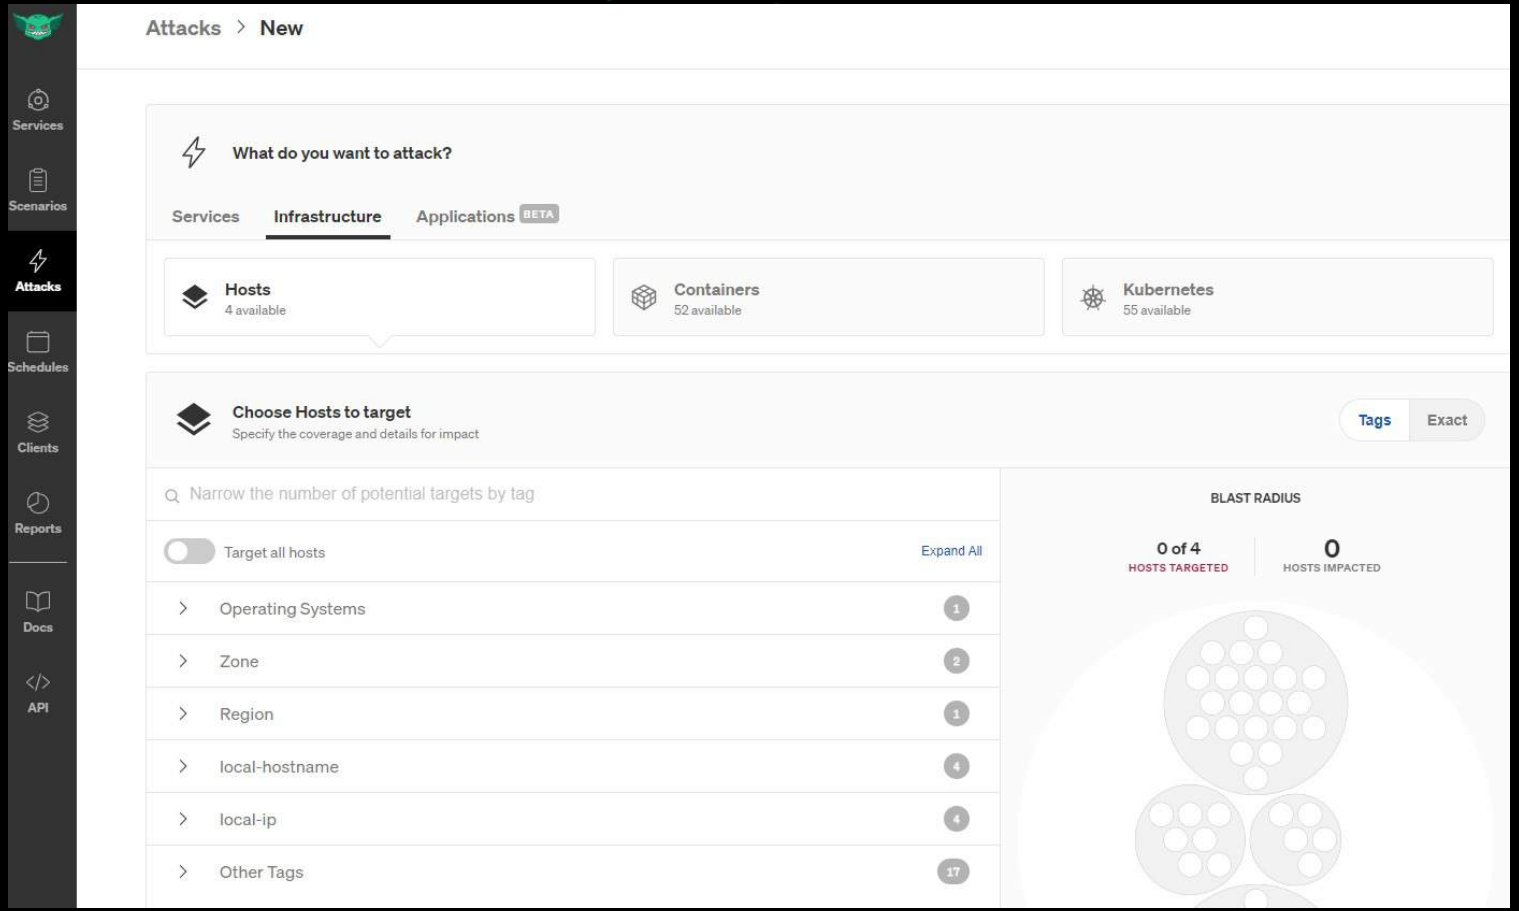
Services
Scenarios
Attacks
Schedules
Clients
Reports
Docs
API
Attacks > New
What do you want to attack?
Services
Infrastructure
Applications BETA
Hosts
4 available
Containers
52 available
Kubernetes
55 available
Choose Hosts to target
Specify the coverage and details for impact
Tags
Exact
Narrow the number of potential targets by tag
Target all hosts
Expand All
Operating Systems 1
Zone 2
Region 1
local-hostname 4
local-ip 4
Other Tags 17
BLAST RADIUS
0 of 4
HOSTS TARGETED
0
HOSTS IMPACTED

Gremlin OSC

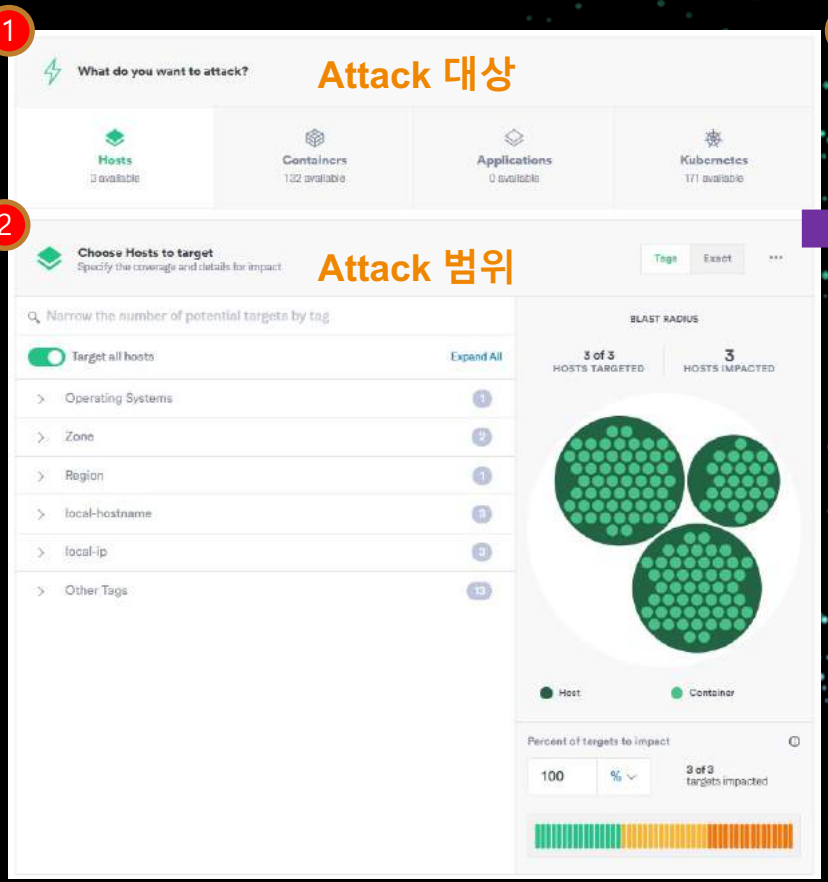
1
What do you want to attack?
Attack 대상
Hosts
3 available
Containers
122 available
Applications
0 available
Kubernetes
171 available
2
Choose Hosts to target
Specify the coverage and details for impact
Attack 범위
Tags
Exact
Narrow the number of potential targets by tag
Target all hosts
Expand All
Operating Systems
Zone
Region
local-hostname
local-ip
Other Tags
BLAST RADIUS
3 of 3
HOSTS TARGETED
3
HOSTS IMPACTED
Host
Container
Percent of targets to impact
100
%
3 of 3
targets impacted

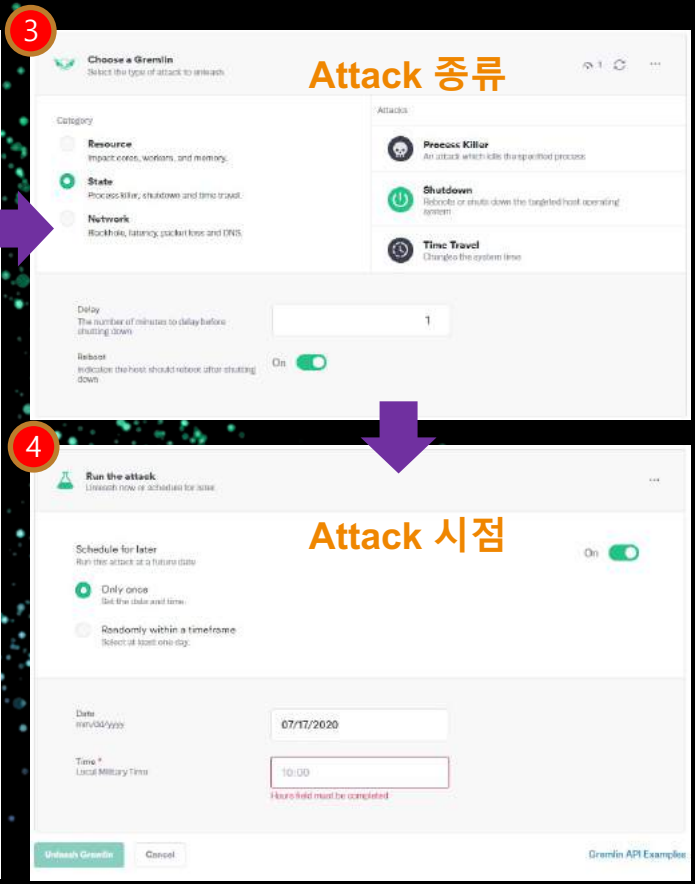
3
Choose a Gremlin
Attack 종류
Category
Resource
State
Network
Attacks
Process Killer
Shutdown
Time Travel
Delay
1
Reboot
On
4
Run the attack
Attack 시점
Schedule for later
On
Only once
Randomly within a timeframe
Date
07/17/2020
Time *
10:00
Unleash Gremlin
Cancel
Gremlin API Examples

Gremlin OSC

리소스

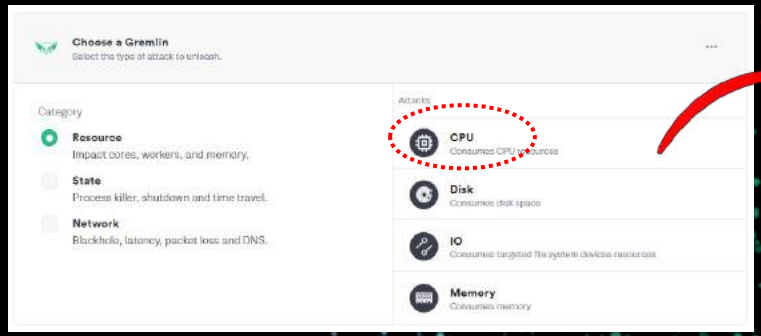

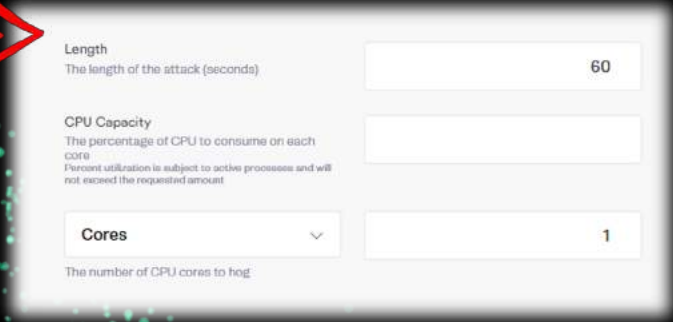

상태

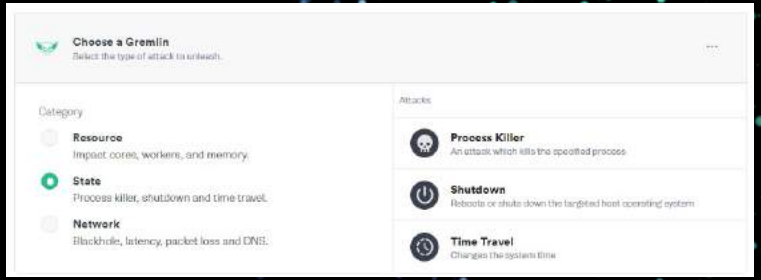

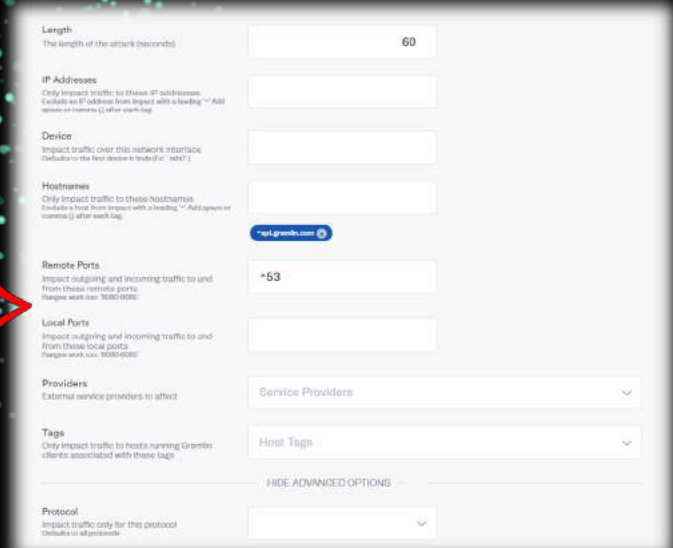

네트워크

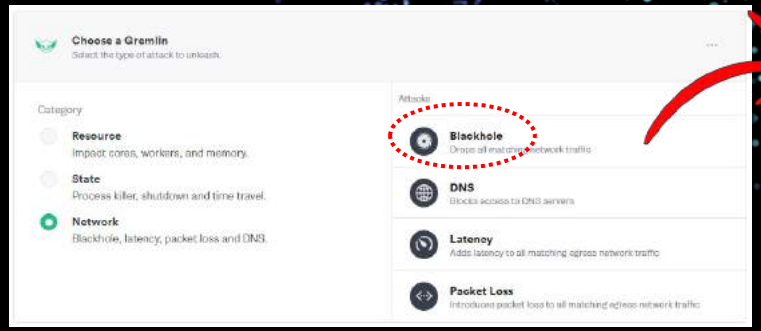

Hipster Shop Anatomy (Online Boutique)

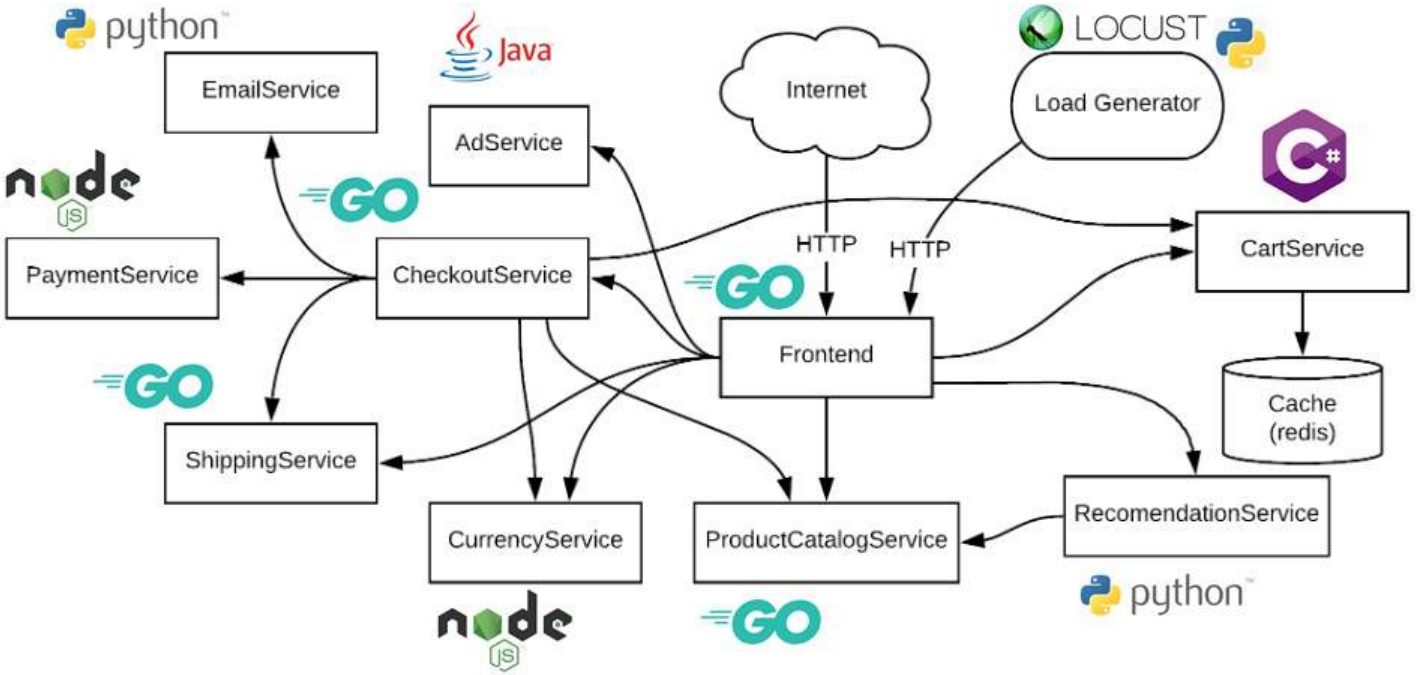

https://github.com/GoogleCloudPlatform/microservices-demo

Gremlin OSC

감사합니다.