

🔰 www.atix.de

# Container, Orchestration, Service Mesh

Automation for your Container Cluster



www.linuxtage.at

Andy Wirtz

27. April 2019

Andy:

Andy Wirtz

- ► IT Consultant at ATIX AG, Germany
- automation of data centers
- deployment of cloud native services
- expertise in Docker, Kubernetes, Istio

Contact:

- phone: +49 (0)89 452 35 38-248
- mail: wirtz@atix.de
- www.xing.com/profile/Andy\_Wirtz2
- www.linkedin.com/in/andy-wirtz







India





#atix #grazerlinuxtage2019 #container #orchestration #servicemesh



How is it possible that my requested service:

- gets as fast and as many new features?
- is always and instantly available?
- looks sometimes diffrent for me than for my mate?



How is it possible that my requested service:

- gets as fast and as many new features?
- is always and instantly available?
- looks sometimes diffrent for me than for my mate?



www.atix.de

How is it possible that my requested service:

- gets as fast and as many new features?
- is always and instantly available?
- Iooks sometimes diffrent for me than for my mate?

























## Wish list



### End user:

how is it possible that my requested service gets as fast and as many new features?

#### DevOps – I want:

- N deployments per day
- independence of my deployments from:
  - hardware
  - operation System
  - other applications



## Wish list



End user:

how is it possible that my requested service gets as fast and as many new features?

DevOps – I want:

- N deployments per day
- independence of my deployments from:
  - hardware
  - operation System
  - other applications



Limit 1



Software monolith:

- many interlocked processes per application
- difficult to:
  - implement new features
  - get new versions running



Idea 1



Microservices:

- one process per application
- communication via simple, well-defined APIs
- faster and easier feature implementation





Dependencies:

- shared use of libraries
- intransparent and error-prone
- deployment dependent on:
  - hardware
  - operation system
  - other applications



Idea 2



#### Container:

- application plus needed libraries
- ► isolation
- ► independence



## New possibilities

Velocity:

- for development
- for testing
- for deployment

Isolation:

- portability
- no dependency conflicts







![](_page_15_Picture_0.jpeg)

![](_page_15_Picture_1.jpeg)

![](_page_15_Picture_2.jpeg)

![](_page_15_Picture_3.jpeg)

![](_page_15_Picture_4.jpeg)

![](_page_15_Picture_6.jpeg)

![](_page_16_Picture_1.jpeg)

www.atix.de

#### End user:

how is it possible that my requested service is always and instantly available?

#### DevOps – I want:

- to use over 80% of my hardware at any given time
- my services up and running at any given time

![](_page_16_Figure_7.jpeg)

![](_page_17_Picture_1.jpeg)

www.atix.de

End user:

how is it possible that my requested service is always and instantly available?

DevOps – I want:

- to use over 80% of my hardware at any given time
- my services up and running at any given time

![](_page_17_Figure_7.jpeg)

Limit 1

![](_page_18_Picture_1.jpeg)

Static deployments:

- limit resource usage of containers is predefined
- hosts with significant overhead
- deployment independent of service requests
- manual scaling as necessary

![](_page_18_Figure_7.jpeg)

Idea 1

![](_page_19_Picture_1.jpeg)

Abstraction of resources:

- collection of host to one cloud
- deployment of pods (containers) in this cluster

Orchestrator:

- dynamic scheduling of pods
- defining the desired state
- automatic adjustment of the current state if necessary
- automatic scaling as necessary

![](_page_19_Figure_10.jpeg)

Limit 2

![](_page_20_Picture_1.jpeg)

Updates with downtime:

- all existing pods are killed
- new ones are created
- the service is down for a small amount of time

![](_page_20_Figure_6.jpeg)

Idea 2

![](_page_21_Picture_1.jpeg)

Rolling updates:

- new pods are created
- they are checked for readyness
- old wones are killed
- zero downtime

![](_page_21_Figure_7.jpeg)

# New possibilities

## Efficiency:

- hardware abstraction
- multiple services per cloud

Automation:

- scaling dependent on the requests
- some ready to use pods in stock
- self-healing
- updates without downtime
- standardization and reproducibility

![](_page_22_Figure_11.jpeg)

![](_page_22_Figure_12.jpeg)

![](_page_22_Picture_13.jpeg)

![](_page_23_Picture_0.jpeg)

![](_page_23_Picture_1.jpeg)

![](_page_23_Picture_2.jpeg)

![](_page_23_Picture_3.jpeg)

![](_page_23_Picture_4.jpeg)

![](_page_23_Picture_6.jpeg)

## Wish list

![](_page_24_Picture_1.jpeg)

End user:

how is it possible that my requested service looks sometimes diffrent for me than for my mate?

DevOps – I want:

- to visualize the traffic between my services
- to deny access between some services
- to encrypt the communication between my services
- to let some end users test a new version of my service

![](_page_24_Figure_10.jpeg)

![](_page_24_Figure_11.jpeg)

## Wish list

![](_page_25_Picture_1.jpeg)

End user:

how is it possible that my requested service looks sometimes diffrent for me than for my mate?

DevOps – I want:

- to visualize the traffic between my services
- to deny access between some services
- to encrypt the communication between my services
- to let some end users test a new version of my service

#atix #grazerlinuxtage2019 #container #orchestration #servicemesh

![](_page_25_Figure_10.jpeg)

![](_page_26_Picture_1.jpeg)

🔰 www.atix.de

Big cloud:

- possibly millions of containers in one cluster
- losing track of the connections and the communication

![](_page_26_Figure_5.jpeg)

![](_page_27_Picture_1.jpeg)

Service Mesh:

- data plane with proxies as side car
- control plane with pilot, citadel and mixer

Proxies:

- take over all network communication
- send telemetry, logs, metrics

Mixer:

- collects telemetry, logs, metrics
- has adapters for the respective backends

![](_page_27_Figure_12.jpeg)

![](_page_27_Figure_13.jpeg)

![](_page_28_Picture_1.jpeg)

🔰 www.atix.de

Permissions:

- every container is allowed to reach every other container
- role based access control is possible
- no restrictions with respect to labels, attributes, ip-adresses

![](_page_28_Figure_6.jpeg)

![](_page_29_Picture_1.jpeg)

#### Proxies:

- call Mixer before each request
- perform precondition checks from cache

Mixer:

- has adapter for policy backends
- performs precondition checks

![](_page_29_Figure_8.jpeg)

![](_page_30_Picture_1.jpeg)

Network security:

- different attributes from different software defined networks
- often times no encryption of traffic

![](_page_30_Figure_5.jpeg)

# Idea 3: Security features

![](_page_31_Picture_1.jpeg)

🔰 www.atix.de

Proxies:

- get certificates, keys and secure naming
- perform mutual TLS communication

Citadel:

- creates certificate and key pairs
- stores them as Kubernetes secrets
- automatically rotates them

Pilot:

 generates secure naming information

passes it to proxies
 #atix #grazerlinuxtage2019 #container #orchestration #servicemesh

![](_page_31_Figure_12.jpeg)

![](_page_32_Picture_1.jpeg)

Service discovery:

- round robin load balancing mode
- no routing rules

![](_page_32_Figure_5.jpeg)

![](_page_33_Picture_1.jpeg)

Proxies:

- get service discovery and traffic rules
- perform service discovery
- load balancing modes: round robin, random, and weighted least request

Pilot:

- manages and configures all proxy instances
- is responsible for the lifecycle of the proxies

![](_page_33_Figure_9.jpeg)

# visualizing metricscollecting logs

- collecting trace spans
- visualizing the mesh

### Policy enforcement:

Telemetry capturing:

- dynamically limit traffic to service
- modify request headers and routing
- control access to service

![](_page_34_Figure_7.jpeg)

# New possibilities

collecting, querying and

![](_page_34_Figure_9.jpeg)

![](_page_34_Picture_10.jpeg)

## New possibilities

Telemetry capturing:

- collecting, querying and visualizing metrics
- collecting logs
- collecting trace spans
- visualizing the mesh

Policy enforcement:

- dynamically limit traffic to service
- modify request headers and routing
- control access to service

![](_page_35_Figure_12.jpeg)

![](_page_35_Picture_13.jpeg)

# New possibilities

![](_page_36_Picture_1.jpeg)

🔰 www.atix.de

Security features:

- using authentication and authorization
- provisioning identity
- using mutual TLS

## Traffic management:

- configuring request routing
- traffic shifting
- ► fault injection
- setting up request timeouts
- controlling ingress/egress traffic
- circuit breaking

mirroring #atix #grazerlinuxtage2019 #container #orchestration #servicemesh

![](_page_36_Figure_14.jpeg)

# New possibilities

![](_page_37_Picture_1.jpeg)

Security features:

- using authentication and authorization
- provisioning identity
- using mutual TLS

Traffic management:

- configuring request routing
- traffic shifting
- fault injection
- setting up request timeouts
- controlling ingress/egress traffic
- circuit breaking

mirroring
 #atix #grazerlinuxtage2019 #container #orchestration #servicemesh

![](_page_37_Figure_14.jpeg)

# Summary

![](_page_38_Picture_1.jpeg)

Microservices/containers:

- velocity
- isolation
- Scheduler/Orchestrator
  - ► efficiency
  - automation

Service Mesh

- telemetry capturing
- policy enforcement
- security features
- traffic management

![](_page_38_Picture_13.jpeg)

# Summary

![](_page_39_Picture_1.jpeg)

Microservices/containers:

- velocity
- isolation
- Scheduler/Orchestrator
  - efficiency
  - automation

Service Mesh

- telemetry capturing
- policy enforcement
- security features
- traffic management

![](_page_39_Picture_13.jpeg)

# Summary

![](_page_40_Picture_1.jpeg)

Microservices/containers:

- velocity
- isolation
- Scheduler/Orchestrator
  - efficiency
  - automation

Service Mesh

- telemetry capturing
- policy enforcement
- security features
- traffic management

![](_page_40_Picture_13.jpeg)

docker

cri-o

🔰 www.atix.de

## Cluster automation

![](_page_41_Picture_1.jpeg)

🔰 www.atix.de

Foreman:

- deployment of hosts
- on-premise and in cloud
  atello:
- lifecycle management
- errata management

Ansible:

- configuration management
- automation

![](_page_41_Picture_10.jpeg)

![](_page_41_Picture_11.jpeg)

![](_page_41_Picture_12.jpeg)

## Cluster automation

![](_page_42_Picture_1.jpeg)

Foreman:

- deployment of hosts
- on-premise and in cloud
  Katello:
  - lifecycle management
  - errata management

Ansible:

- configuration management
- automation

![](_page_42_Picture_10.jpeg)

![](_page_42_Picture_11.jpeg)

![](_page_42_Picture_12.jpeg)

![](_page_42_Picture_13.jpeg)

## Cluster automation

![](_page_43_Picture_1.jpeg)

🔰 www.atix.de

Foreman:

- deployment of hosts
- ► on-premise and in cloud Katello:
  - lifecycle management
  - errata management

Ansible:

- configuration management
- automation

![](_page_43_Picture_10.jpeg)

![](_page_43_Picture_11.jpeg)

![](_page_43_Picture_12.jpeg)

## ATIX

![](_page_44_Picture_1.jpeg)

## Orcharhino:

- delivery
- improvements
- quality assurance
- support
- documentation

#### ATIX:

- consulting
- engineering
- support
- ► training

![](_page_44_Picture_13.jpeg)

## ATIX

![](_page_45_Picture_1.jpeg)

## Orcharhino:

- delivery
- improvements
- quality assurance
- support
- documentation

### ATIX:

- consulting
- engineering
- support
- training

![](_page_45_Picture_13.jpeg)

## ATIX booth

![](_page_46_Picture_1.jpeg)

![](_page_46_Picture_2.jpeg)

#atix #grazerlinuxtage2019 #container #orchestration #servicemesh

## **ATIX** presentations

Dr. Sebastian Oehlke:

- Terraform Ein Einblick der Möglichkeiten von Infrastructure-as-a-Code
- ► 4pm,ill

## Manuel Bonk:

- Next Level Ansible
- ▶ 5pm, ill

![](_page_47_Picture_8.jpeg)

![](_page_47_Picture_9.jpeg)

## **ATIX** presentations

Dr. Sebastian Oehlke:

- Terraform Ein Einblick der Möglichkeiten von Infrastructure-as-a-Code
- ▶ 4pm,ill

Manuel Bonk:

- Next Level Ansible
- ► 5pm,ill

![](_page_48_Picture_8.jpeg)

![](_page_48_Picture_9.jpeg)