

# Distributed applications: A challenge for systems, networks, and application development

---

Holger Karl  
Paderborn University




Computer Networks Group  
Universität Paderborn

- Developing and deploying applications: Together
  - Buzzword: Continuous software engineering, CI/CD
- Deploying: Applications and infrastructure
  - Buzzword: Infrastructure as code
- Education issues

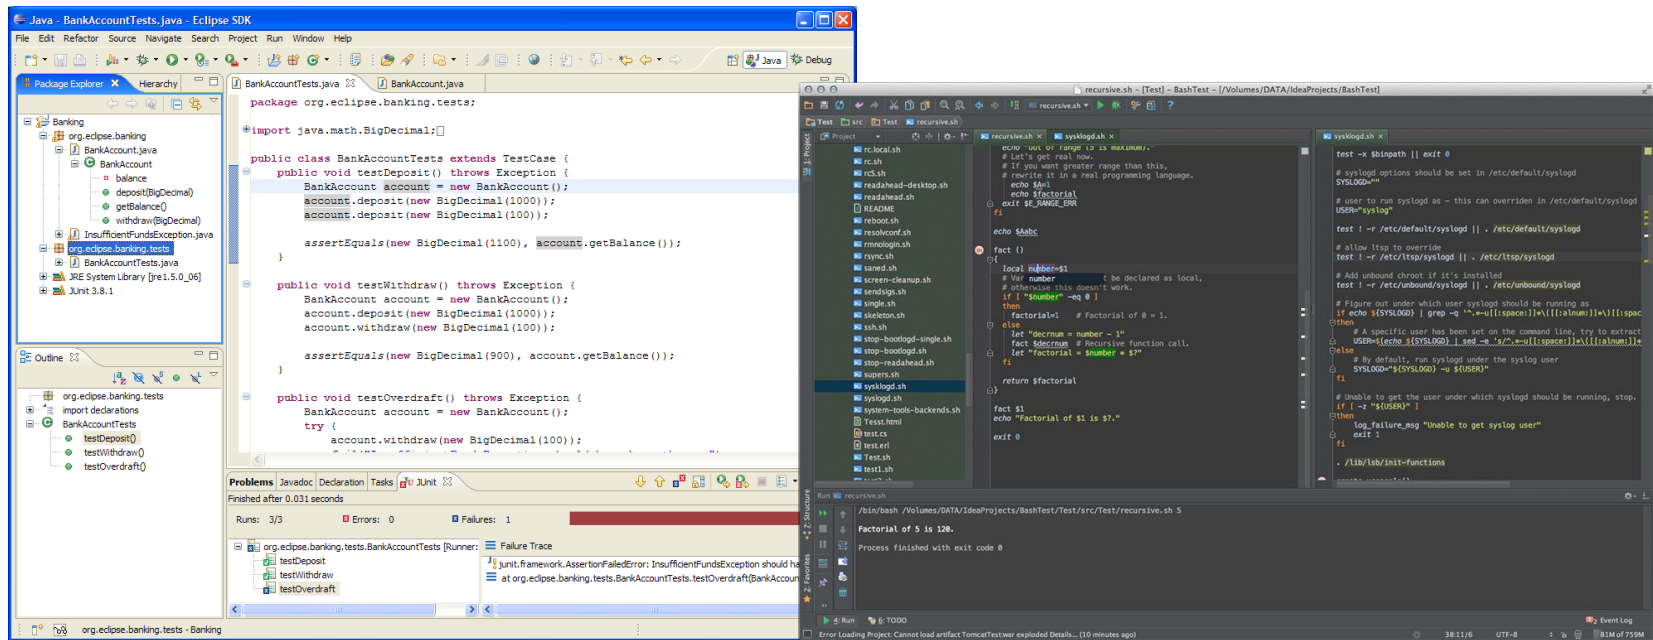


# Key ingredients



# Developing old school

- Monolithic IDEs producing monolithic code



Code

# Deploying?

---



# Deploying, old school

---

```
$ # download
$ wget http://www.developer.com/app.tgz
$ tar xfz app.tgz
$ # compile and build
$ cd app
$ ~/app> ./configure
$ ~/app> make
$ ~/app> make install
$ # run
$ ~/app> ./app
```

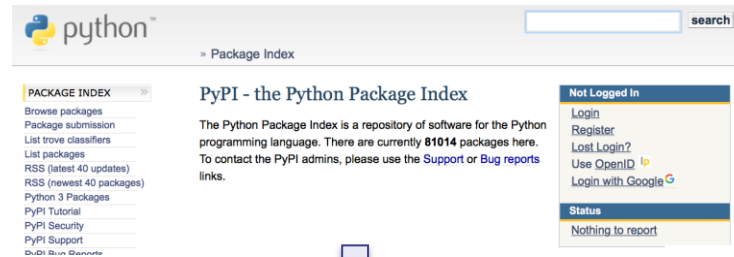
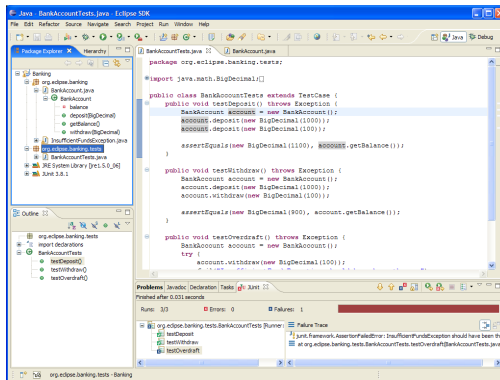
- Great for single application, single machine



# Developing, ca. 2017

- Own code, plus tons of modules, libraries, **microservices**,

...



Code



Code



Code



Code




# Deploying using virtual machines

```
$ # download description
$ wget http://www.developer.com/app.tgz
$ tar xfz app.tgz
$ cd app
$ # edit configuration:
$ vi ansible-playbook.yml
$ # download actual image, configure it
$ vagrant --provision up
$ # tell DNS where to find the app:
$ curl --request POST www.mydns.com -d {'fdqn': myapp.com, 'ip':
1.2.3.4}
```

Customize a ready-made  
image to local needs


E.g., network addresses

- Great for simple applications
- Even inside, e.g., VMWare clusters





# Setup after deployment



# Deploying complex applications – Manually

- Anything real will need multiple machines

```
$ wget http://www.developer.com/app.tgz
$ tar xfz app.tgz
$ # a frontend server:
$ cd app/frontend
$ # complex editing of network configuration
$ app/frontend> vi ansible-playbook.yml
$ app/frontend> vagrant --provision up
$
$ # a backend server:
$ cd ../app/backend
$ # complex editing of network configuration
$ app/backend> vi ansible-playbook.yml
$ app/backend> vagrant --provision up
$ # tell DNS where to find the app:
$ curl --request POST www.mydns.com -d {'fdqn': myapp.com, 'ip':
frontend_ip}
```


Tell frontend how  
to talk to backend:  
IP, ports, keys, ...

Tell backend which  
frontend to accept:  
IP, ports, keys, ...

- Getting slightly less great



# Setup after deployment



# Scaling out?

- Load increases:  
need another frontend

Start second  
frontend

Integrate it  
with backend

```
$ # copy frontend configuration:
$ cp -r app/frontend app/frontend2
$ cd app/frontend2
$ app/frontend2> # complex editing of network configuration
$ app/frontend> vi ansible-playbook.yml
$ # bring up a new frontend instance
$ app/frontend> vagrant --provision up
$
$ # tell backend to accept new frontend
$ cd ../app/backend
$ aap/backend> vagrant ssh
[backend] $ vi backend.config # add the second frontend
server
[backend] $ ./backend restart # trigger reloading of
config
```



# Scaling out?

- But wait, there's more:

Distribute traffic to two frontends:  
Loadbalancer needed

Firewall needs to be told to allow traffic


*firewall*

DNS should point to loadbalancer

```
$ wget http://www.loadbalancer.com/lb.tgz
$ tar xzf lb.tgz
$ cd lb
$ # complex editing of lb configuration
$ lb> vagrant --provision up
$ cd ..
$ ssh firewall # reconfigure
$ # tell DNS where to find the app:
$ curl --request POST www.mydns.com -d {'fdqn': myapp.com, 'ip': lb_ip}
```

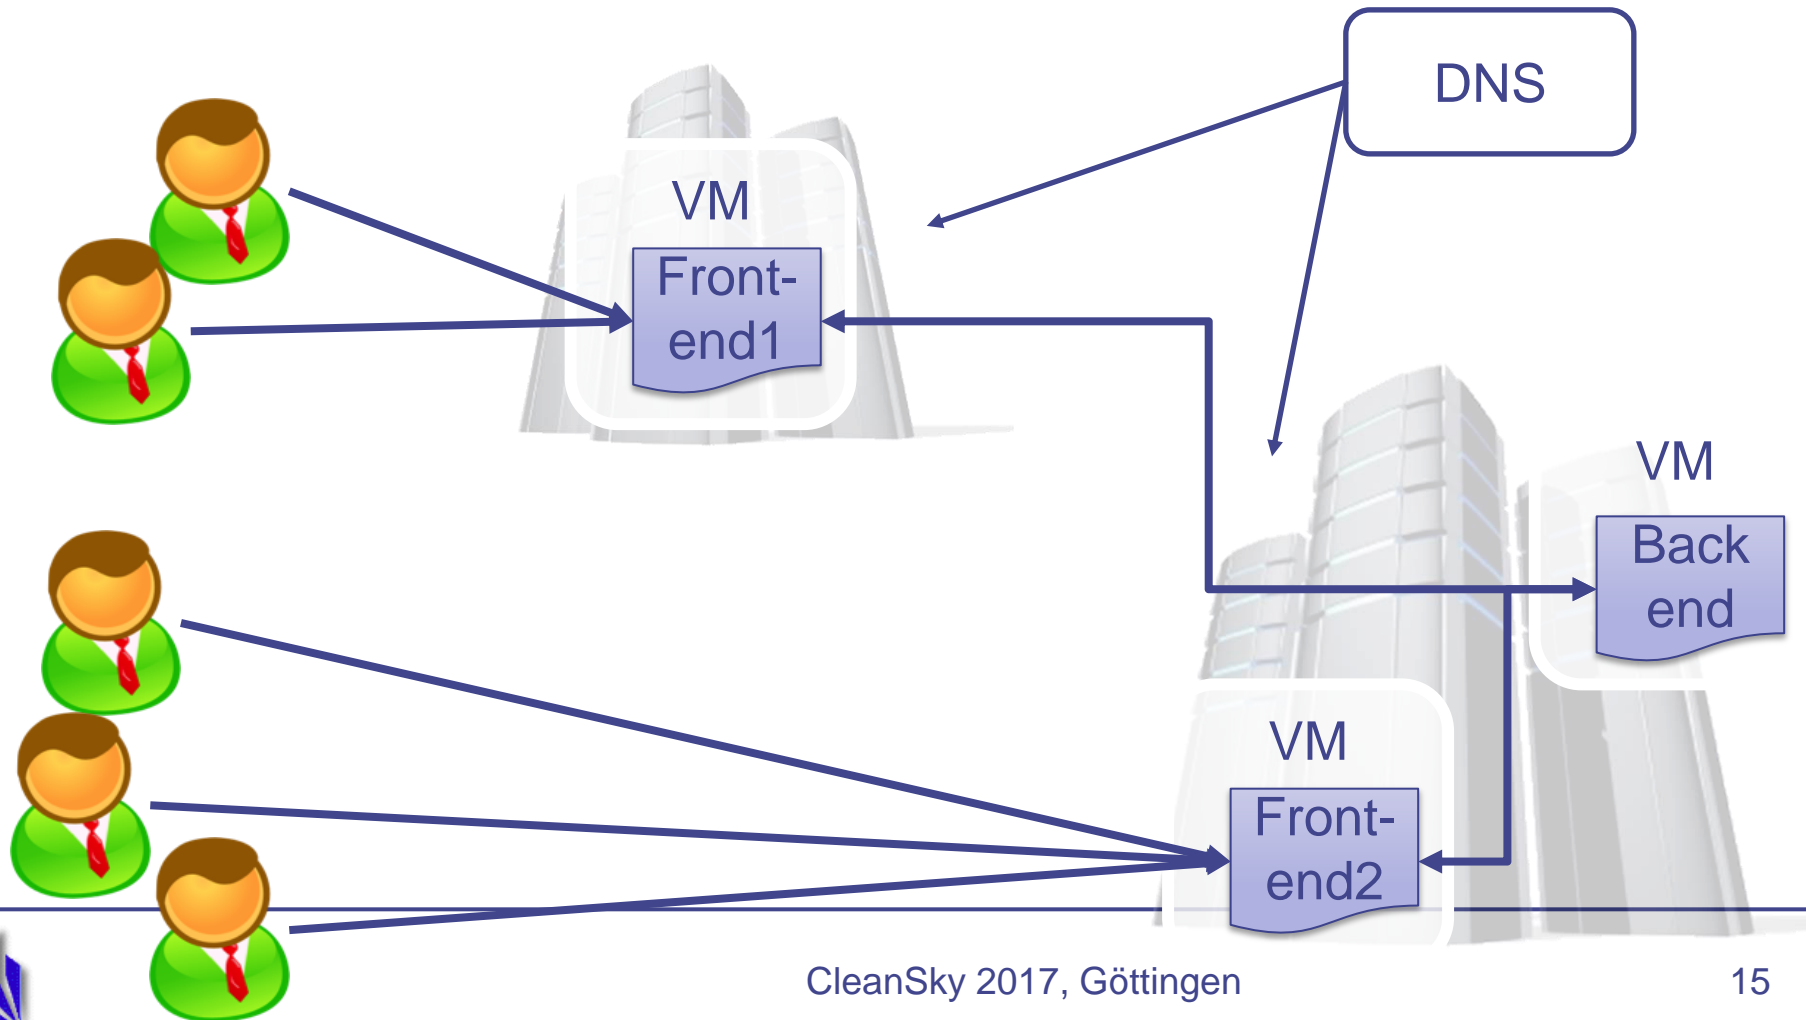


# Setup after deployment



# Frontends at different sites?

- Wouldn't this make more sense?



# Deploying frontends at remote sites

```
$ # configure frontend2 to run elsewhere
$ cd app/frontend2
$ vagrant halt
$ # put IP of other site into configuration
$ vagrant up
$
$ # drop loadbalancer
$ # reconfigure firewall at both sites
$ # reconfigure DNS
$ # reconfigure backend
```

- Definitely not great!






# What do we need?

---

- Get rid of manual configuration via console
- **Describing** configuration of a composed application
  - Its actual application-oriented components
  - Support components like firewalls, ...
  - Scaling properties
- **Deploying** such a composed application
- **Configuring** underlying infrastructure



# Issues



# Describing complex applications – Templates

---

- **Templates** describe components in an application
- Needs
  - Which components
  - How interconnected
  - How to scale
  - Where to place



# Template examples

---

- OpenStack: HEAT Orchestration Template (HOT)
  - HEAT: OpenStack's orchestration plugin
  - Composition, interconnected, limited scaling properties, single site
- Docker Swarm
  - Container-based (instead of VMs)
  - Relatively simple setup, easy generation
- Google Kubernetes



# Docker/Kubernetes example: Guestbook application


```
1  apiVersion: extensions/v1beta1
2  kind: Service
3  metadata:
4    name: redis-master
5    labels:
6      app: redis-master
7      tier: backend
8      role: master
9  spec:
10    ports:
11      # the port on which we will access the cluster
12      - port: 6379
13        targetPort: 6379
14    selector:
15      app: redis-master
16      tier: backend
17      role: master
18    ---
19  apiVersion: extensions/v1beta1
20  kind: Deployment
21  metadata:
22    name: redis-master
23    # these labels can be applied automatically
24    # from the labels in the pod template
25    # labels:
26    #   app: redis-master
27    #   role: master
28    #   tier: backend
29  spec:
30    # this replicates the pod
31    # modify it according to your case
32    replicas: 1
33    # selector can be applied automatically
34    # from the labels in the pod template if not set
35    # selector:
36    #   matchLabels:
37    #     app: guestbook
38    #     role: master
39    #     tier: backend
40    template:
41      metadata:
42        labels:
43          app: redis-master
44          role: master
45          tier: backend
46      spec:
47        containers:
48          - name: redis-master
49            image: gcr.io/google_samples/gb-redis:v1
50            resources:
51              requests:
52                cpu: 100m
53                memory: 100Mi
54            ports:
55              - containerPort: 6379
56    ---
57  apiVersion: extensions/v1beta1
58  kind: Deployment
59  metadata:
60    name: redis-slave
61    # these labels can be applied automatically
62    # from the labels in the pod template
63    # labels:
64    #   app: redis-slave
65    #   role: slave
66    #   tier: backend
67  spec:
68    # this replicates the pod
69    # modify it according to your case
70    replicas: 3
71    # selector can be applied automatically
72    # from the labels in the pod template if not set
73    # selector:
74    #   matchLabels:
75    #     app: guestbook
76    #     role: slave
77    #     tier: backend
78    template:
79      metadata:
80        labels:
81          app: redis-slave
82          role: slave
83          tier: backend
84      spec:
85        containers:
86          - name: redis-slave
87            image: gcr.io/google_samples/gb-redis:v1
88            resources:
89              requests:
90                cpu: 100m
91                memory: 100Mi
92            ports:
93              - containerPort: 6379
94    ---
95  apiVersion: v1
96  kind: Service
97  metadata:
98    name: frontend
99    # these labels can be applied automatically
100    # from the labels in the pod template if not set
101    # labels:
102    #   app: guestbook
103    #   tier: frontend
104  spec:
105    # this replicates the value is default
106    # modify it according to your case
107    replicas: 3
108    # selector can be applied automatically
109    # from the labels in the pod template if not set
110    # selector:
111    #   matchLabels:
112    #     app: guestbook
113    #     tier: frontend
114    template:
115      metadata:
116        labels:
117          app: guestbook
118          tier: frontend
119      spec:
120        containers:
121          - name: php-redis
122            image: gcr.io/google_samples/gb-frontend:v3
123            resources:
124              requests:
125                cpu: 100m
126                memory: 100Mi
127            env:
128              - name: GET_HOSTS_FROM
129                value: dns
130              # If your cluster config does not include a dns service, then to
131              # instead access environment variables to find service host
132              # info, comment out the 'value: dns' line above, and uncomment the
133              # line below.
134              # value: env
135            ports:
136              - containerPort: 80
```



<https://github.com/kubernetes/kubernetes/blob/release-1.2/examples/guestbook/all-in-one/guestbook-all-in-one.yaml>

# Templates, as they should be

- **Quantitative** annotations
  - From profiling
- Scaling options
- ...



# Services?

- “Ordinary” application boxes?

Front-  
end

Back-  
end

Data-  
base

- “Network” functions?

Fire-  
wall

Load-  
balancer

- What’s the key, pivotal difference?



# Claim: None!

---


- There is no inherent difference between
  - microservices
  - network functions
- Differences are a historical, ecosystem-based coincidence
  - E.g., separate orchestrators, ...





# Mid-term goal: Merge NFV and (D)CC

---




# Crossing from VNFs to SDN?

---

- Example: Load balancer
  - Given as a VNF
  - Or as SDN rules
- Decide which one?
- Automatically convert?
- Open!

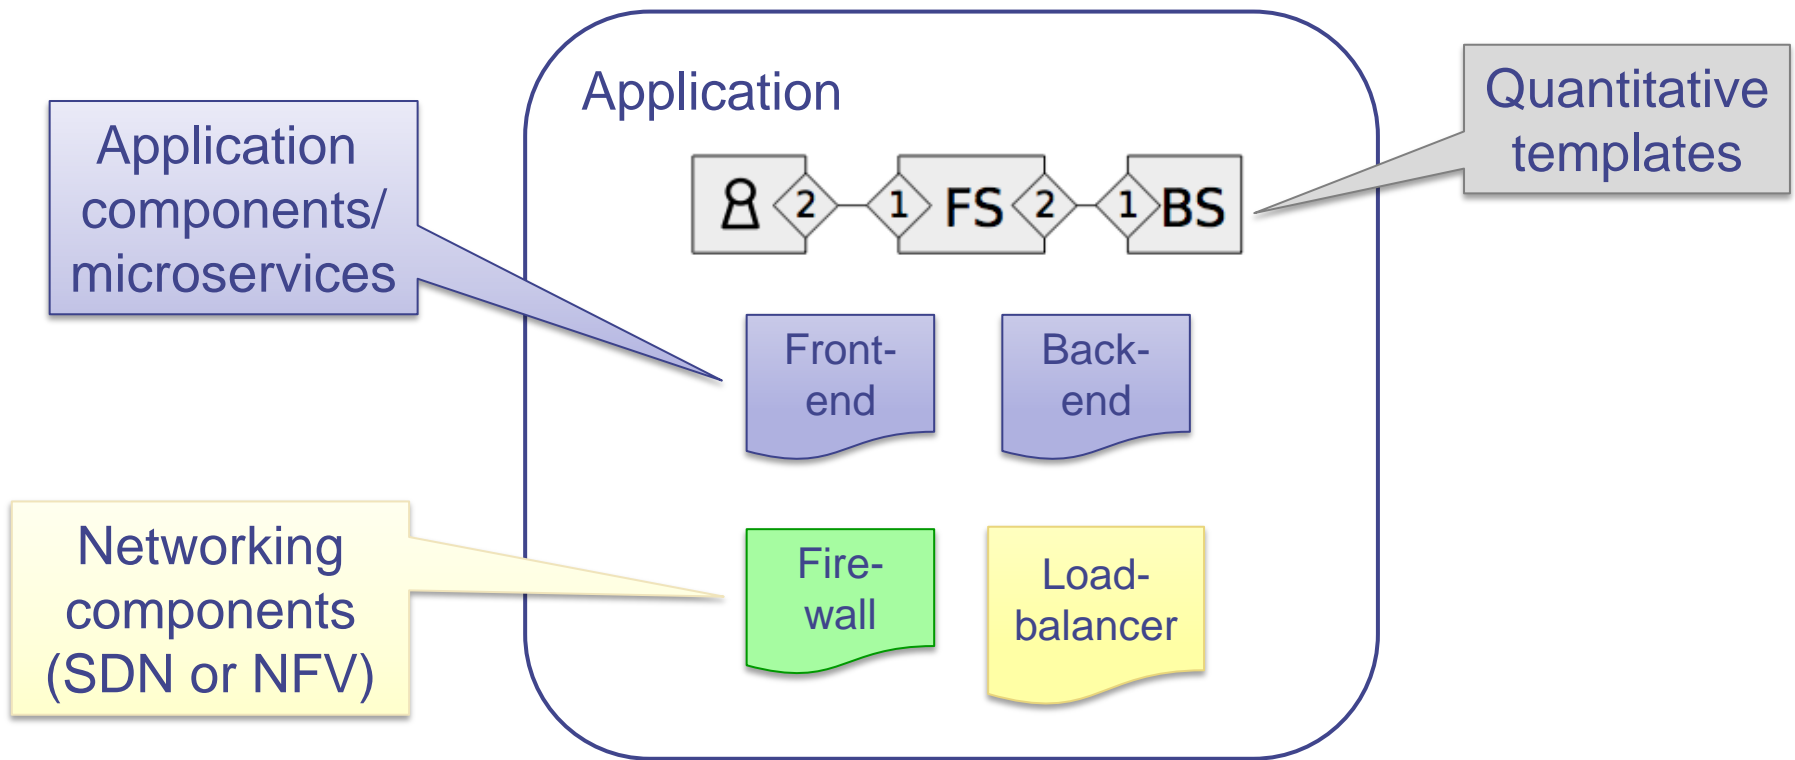


## Mid-term goal: Merge NFV, DCC, SDN





# Applications become more complex

- **Application =  
Actual application + networking**




Developer knows best!

# Example template: Quantitative scaling



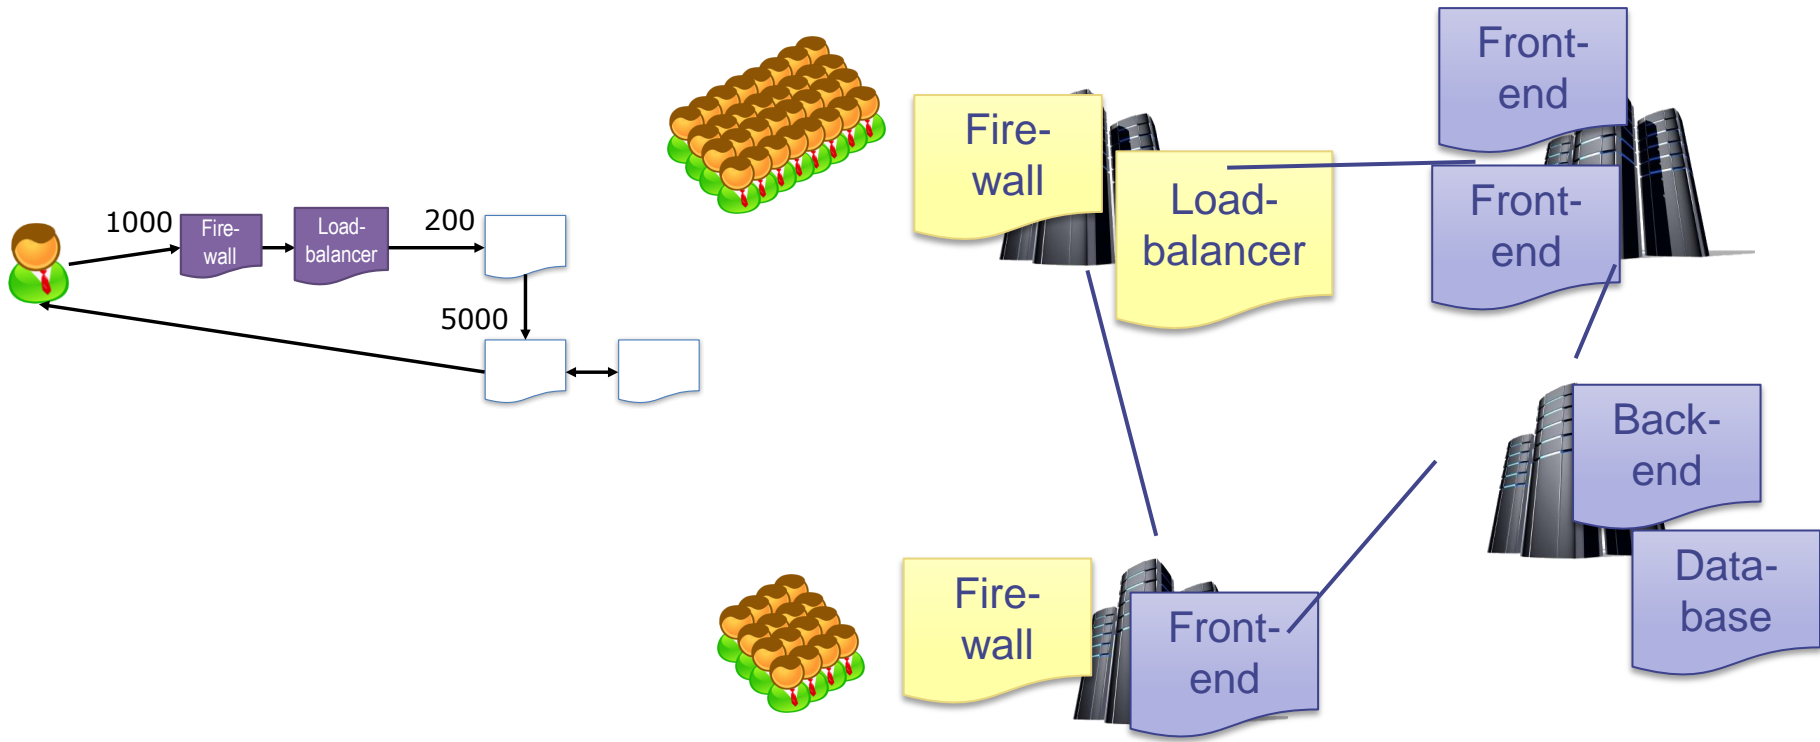
M. Keller, C. Roberts, H. Karl: Template Embedding: Using Application Architecture to Allocate Resources in Distributed Clouds, UCC 2014

# Example template: Reordering



S. Mehraghdam, M. Keller, H. Karl: Specifying and Placing Chains of Virtual Network Functions, CloudNet 2014

# Issues





# Deploying?

- Scaling?
- Placement?
- Lifecycle management?

- When to start, stop, migrate state of a component

- Where to put which component?
- Virtual network embedding – NP-hard
- Embedding **scalable** templates: generalization
  - Still solvable with good heuristics



# One-size-fits-all solution?

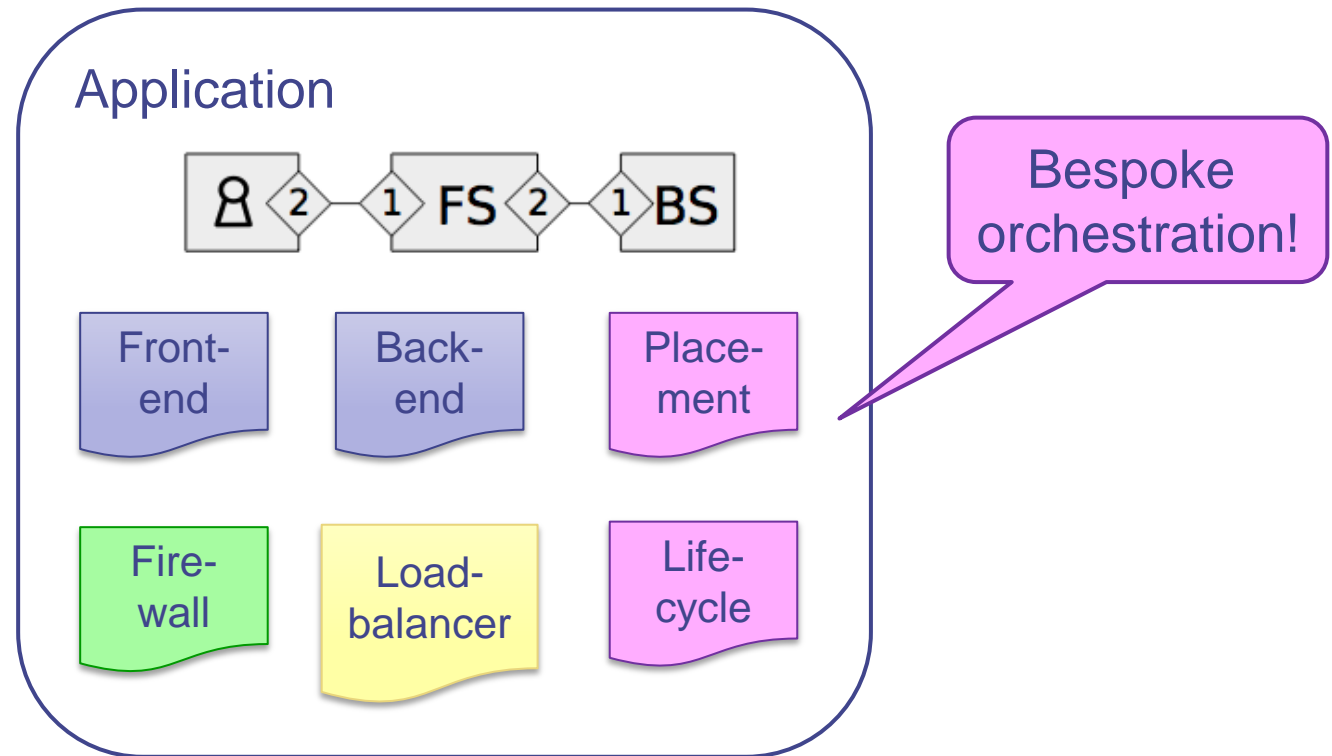
---

- Use a single, “perfect” deployment scheme?
  - Independent of application?
- Remember: Developer knows best!
  - SotA: Some parameterization exists, but limited in flexibility
    - E.g., Kubernetes
- No! Need bespoke deployment, orchestration schemes!



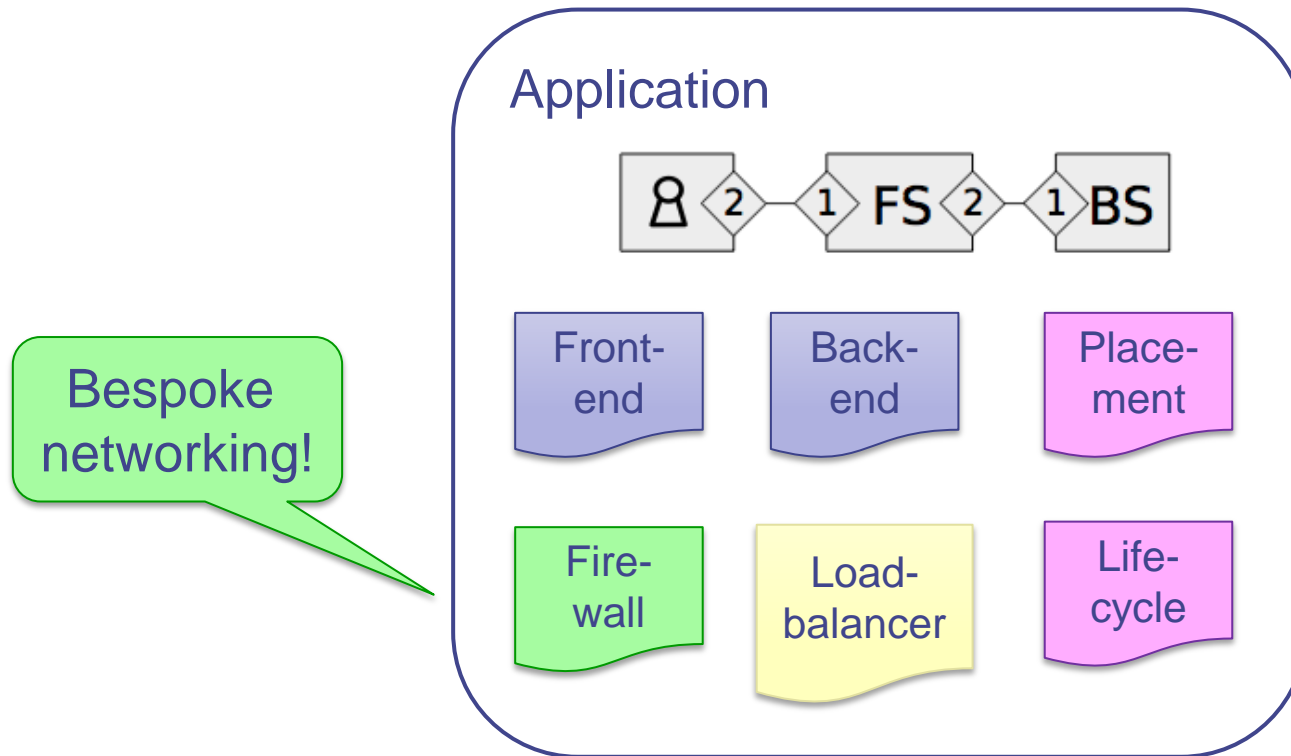
# Developer knows best!

- **Application =  
Orchestration + actual application + networking**
  - UNIFY, SONATA, ...

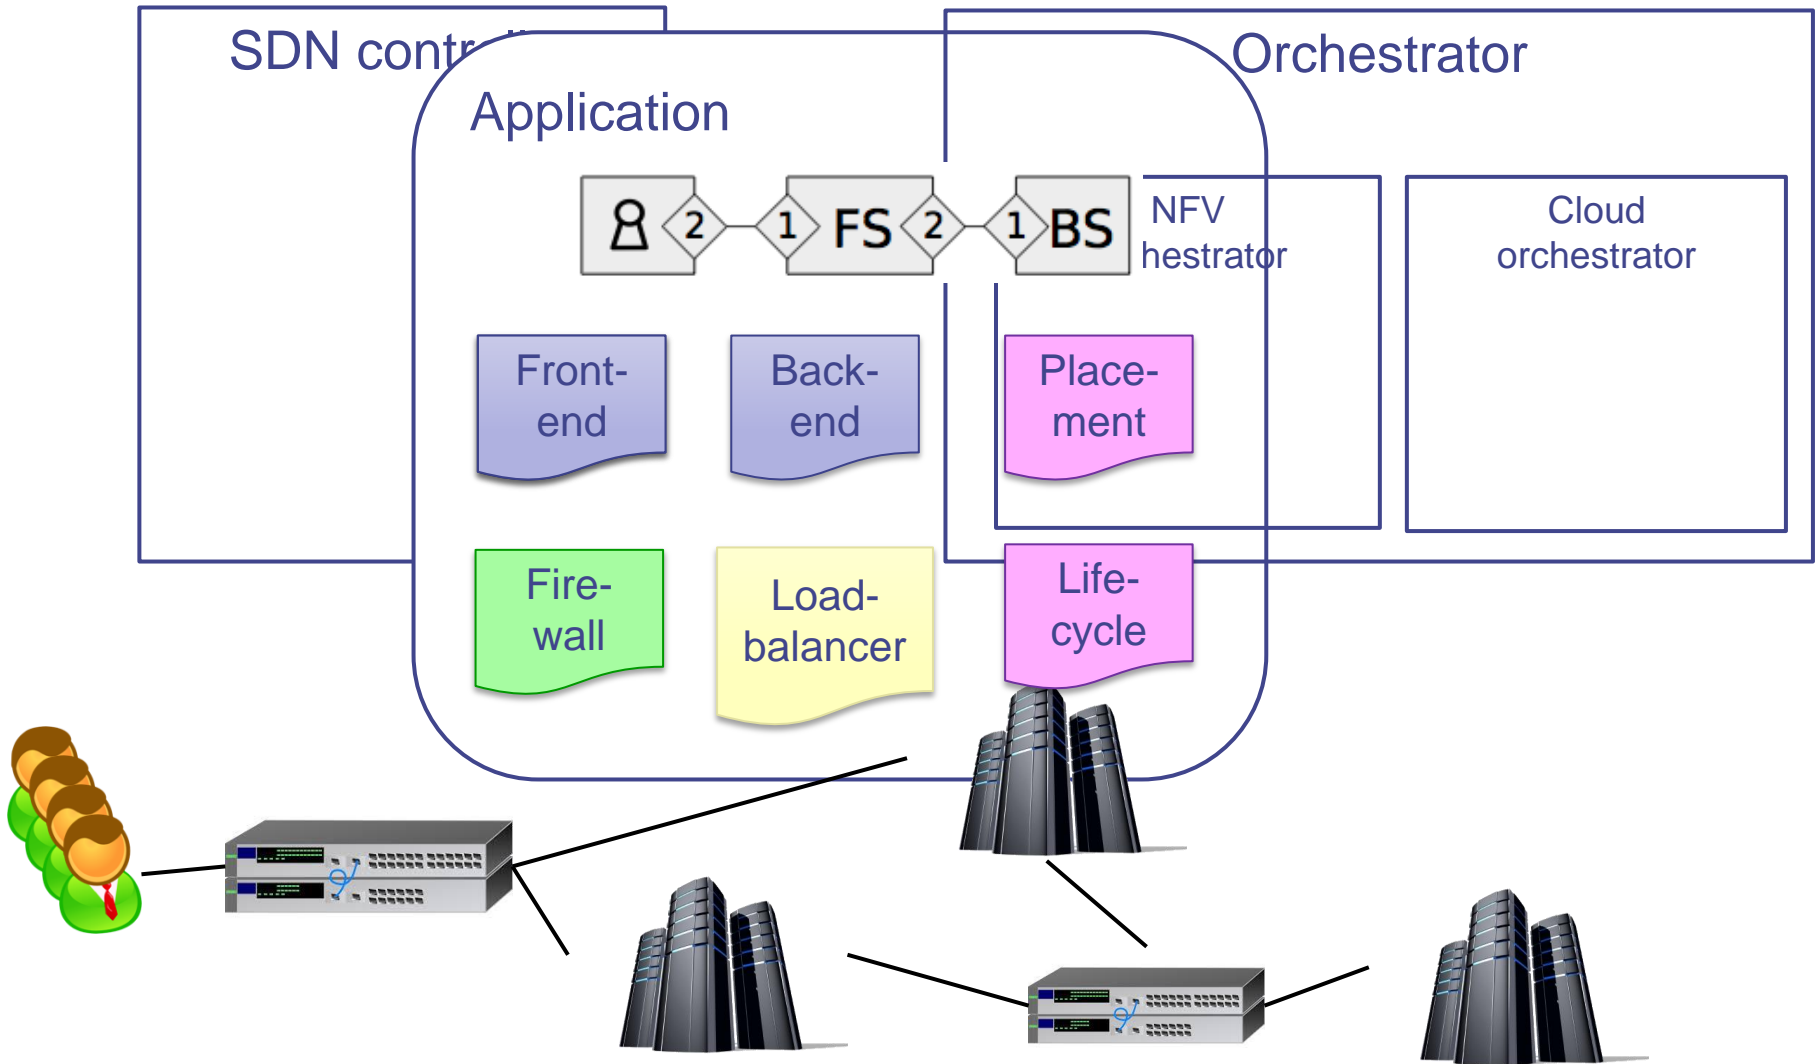


# Reality check: Networking

- Networking: SDN controller applications
- Create network topology for application
- Developer support? Legacy?



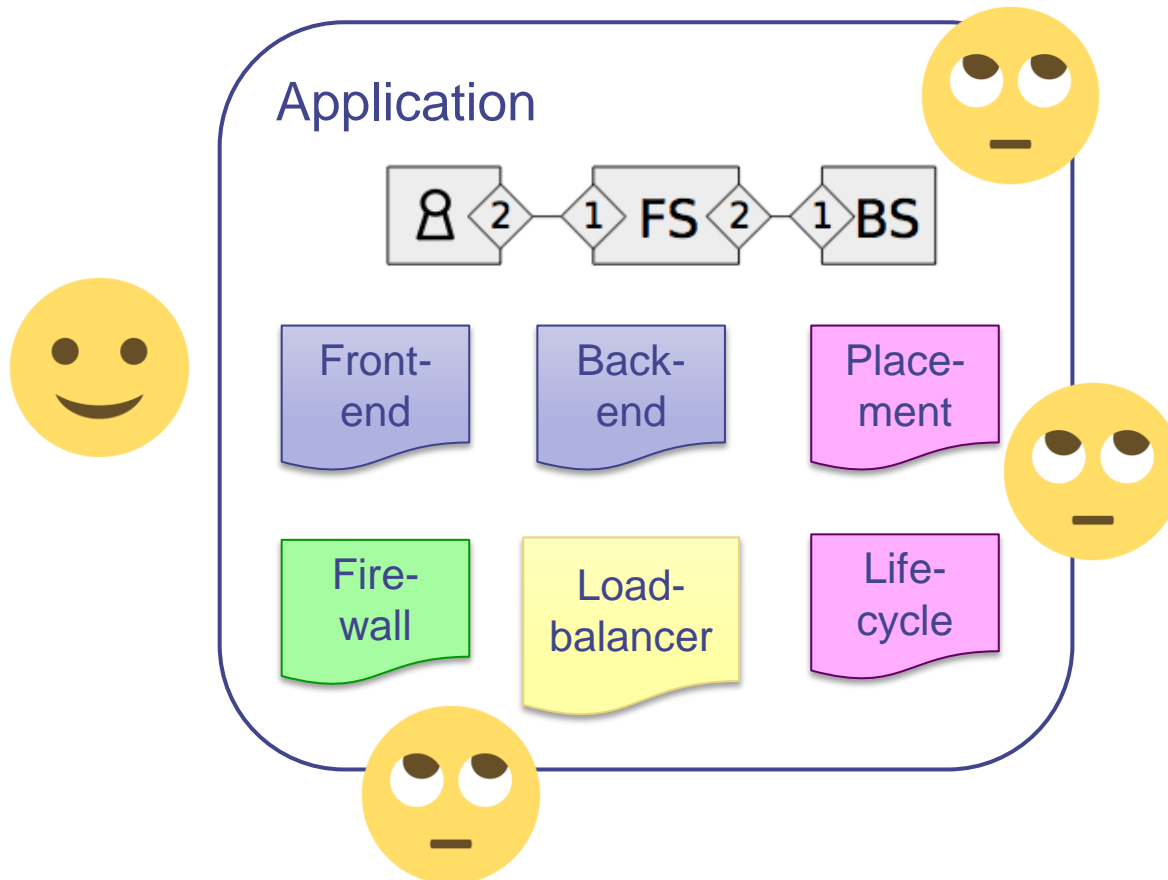
# High-level architecture



Developer knows best?

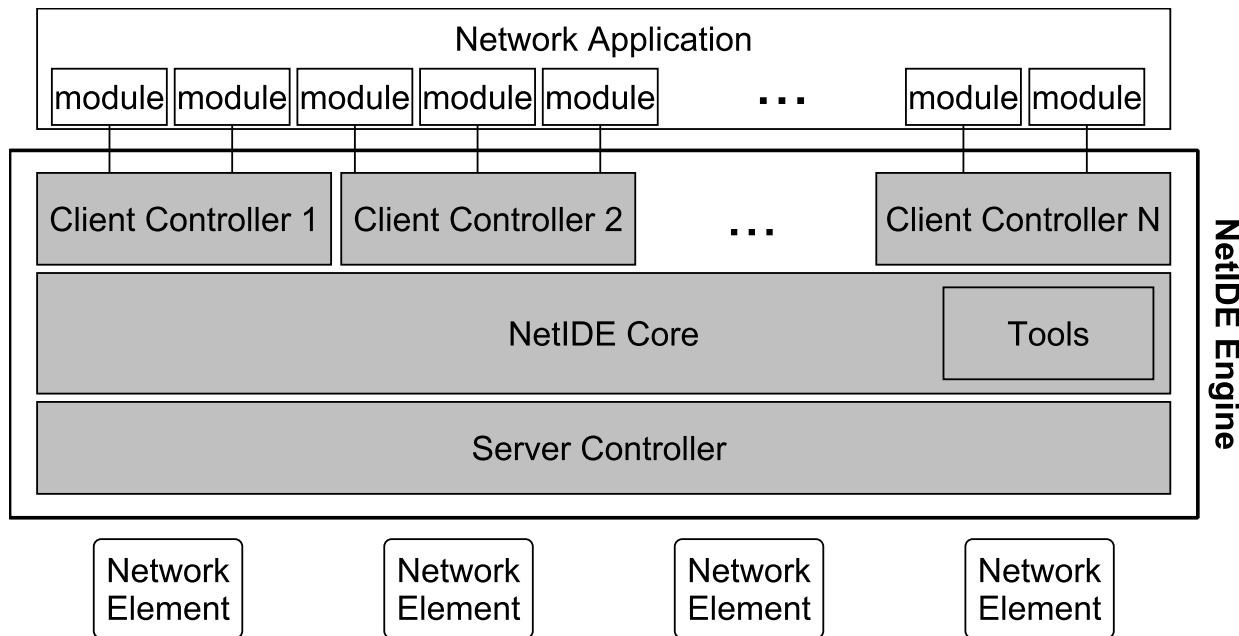


# Developer knows best?



# Reality check: Develop SDN network applications

- NetIDE: Bring SDN closer to developers
- Supporting legacy applications on controller frameworks



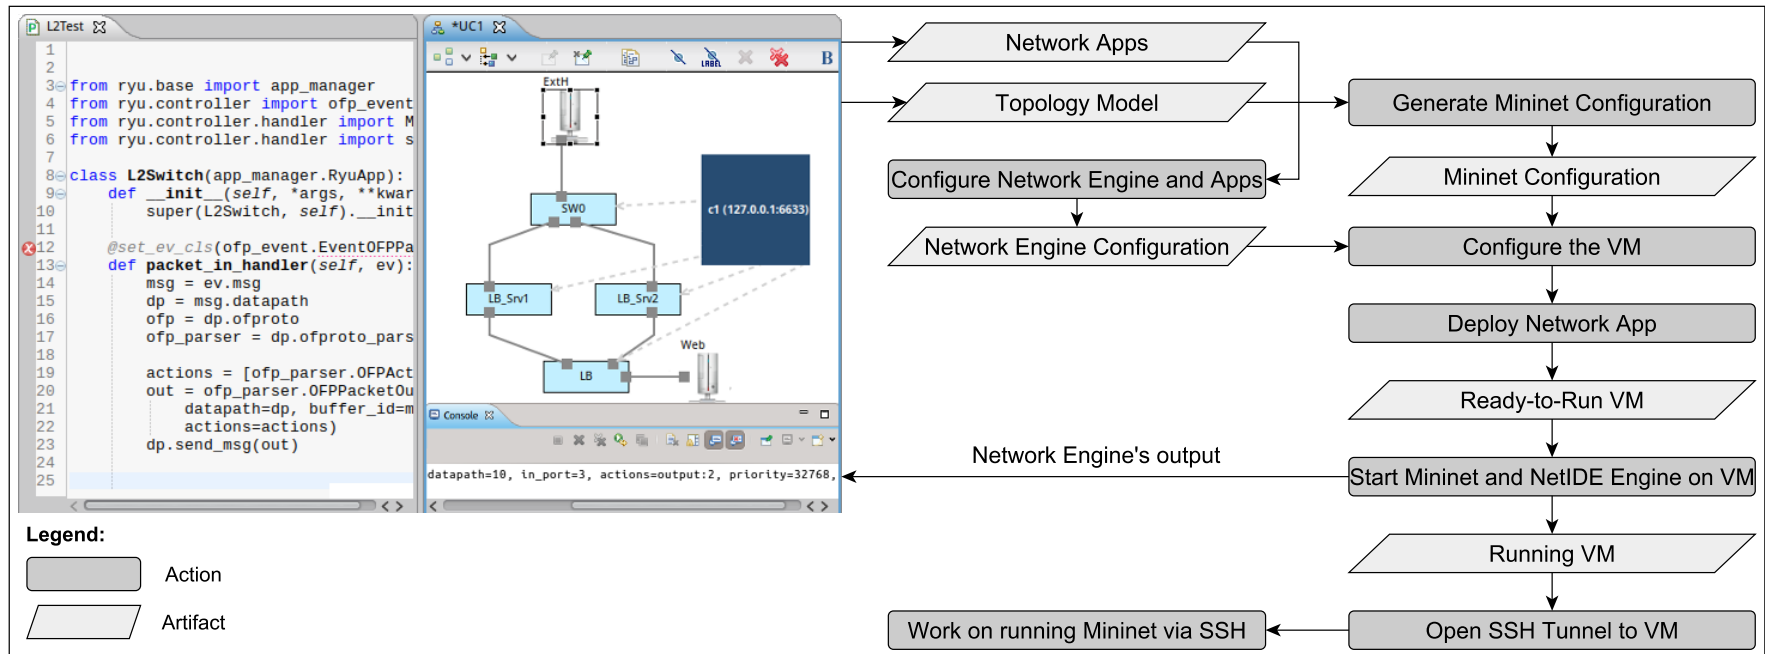
<http://www.netide.eu>



P. A. Aranda Gutiérrez et al., NetIDE: All-in-one framework for next generation, composed SDN applications, NetSoft Demo, 2016



# NetIDE: Ease developer's tasks by SDK



P. A. Aranda Gutiérrez et al., NetIDE: All-in-one framework for next generation, composed SDN applications, NetSoft Demo, 2016

# Reality check: Templates

---

- Ideal developer writes perfect templates
- Real developer?
- Option: semi-automatic generation?
  - Based on semantic understanding of components

**SFB901**



ON - THE - FLY COMPUTING

<https://sfb901.uni-paderborn.de>



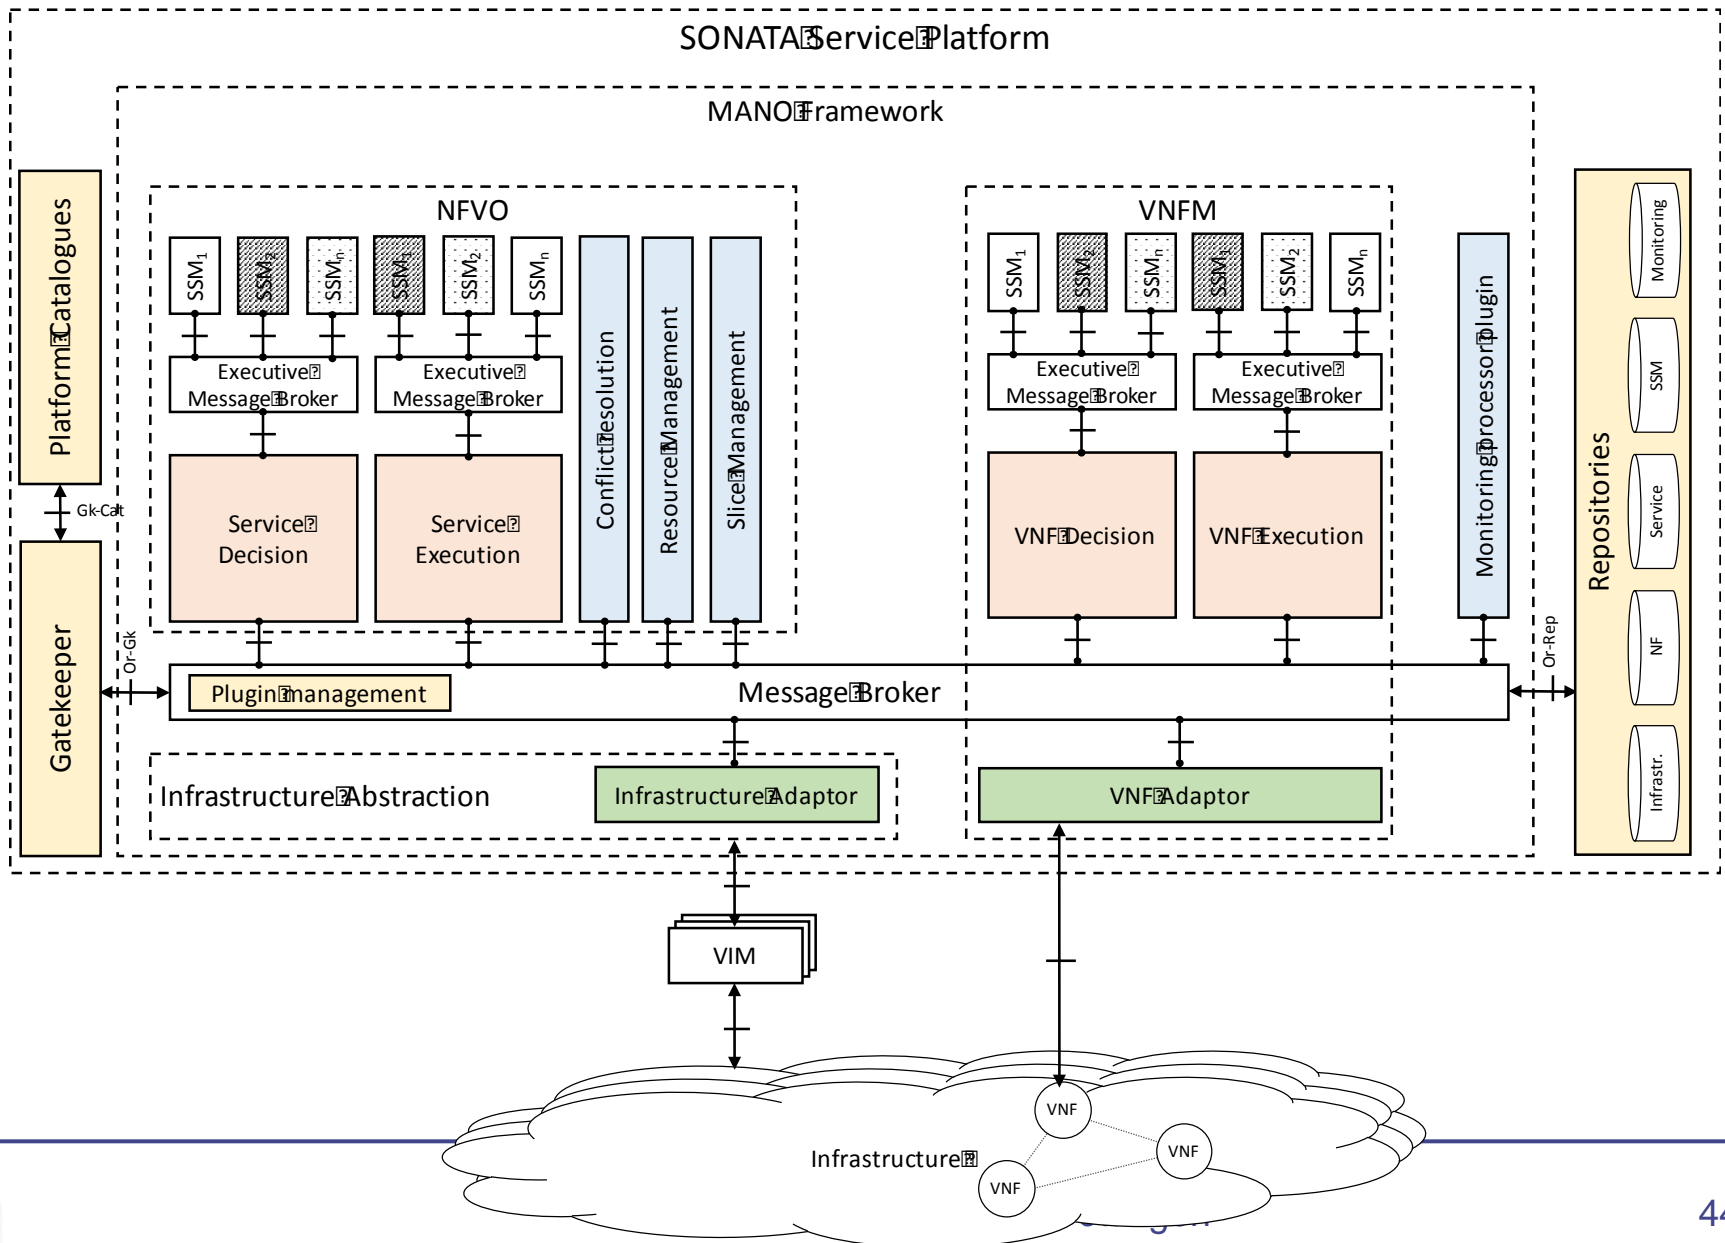
# Reality check: Orchestration with 3rd party modules?

- Run code from outside application inside an orchestration platform?
  - Security, isolation, conflicts, ...?
- Sure – as microservices!
  - SONATA (5G PPP): microservices-based orchestration



<http://sonata-nfv.eu>

# SONATA: Architecture overview



# Challenge: Education!

---

- Teaching curricula: Integrate networking, distributed systems, virtualization
- Continuous education



# Conclusion

---

- NFV and (D)CC need to merge; SDN integrated
- Developers and Ops need to understand that
  - DevOps for complex, distributed applications
- Research & education task

