



*Fast Reoptimization with only a few Changes:*  
Enhancing Tactical Traffic Engineering with Segment Routing  
Midpoint Optimization

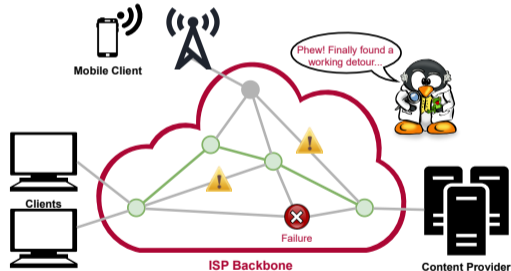
**Alexander Brundiers**, Timmy Schüller, Nils Aschenbruck

GI/ITG KuVS FG “Network Softwarization”  
April 03–04, 2025

# Quick Introduction - What is Traffic Engineering?

Traffic Engineering in one (simplified) sentence:

Control the paths of traffic flows in your network to achieve various objectives.



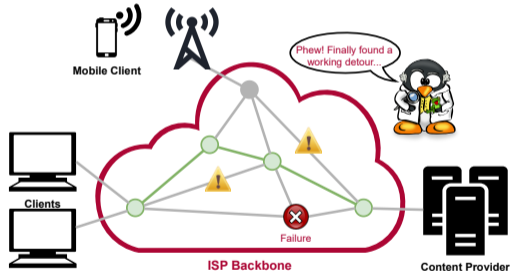
# Quick Introduction - What is Traffic Engineering?

Traffic Engineering in one (simplified) sentence:

Control the paths of traffic flows in your network to achieve various objectives.

Possible Use Cases or Objectives:

- Reduce load of highly utilized links
- Reduce energy consumption (GreenTE)
- Deal with failures & traffic changes
- ...




# Quick Introduction - What is Traffic Engineering?

Traffic Engineering in one (simplified) sentence:

Control the paths of traffic flows in your network to achieve various objectives.

Possible Use Cases or Objectives:

- Reduce load of highly utilized links
- Reduce energy consumption (GreenTE)
- Deal with failures & traffic changes
- ...




Implemented via *Metric-Tuning*, *MPLS (with RSVP-TE)*, *Segment Routing*, ...

## Segment Routing (SR) in a Nutshell:

Control a packet's path by defining **interim destinations/waypoints**:

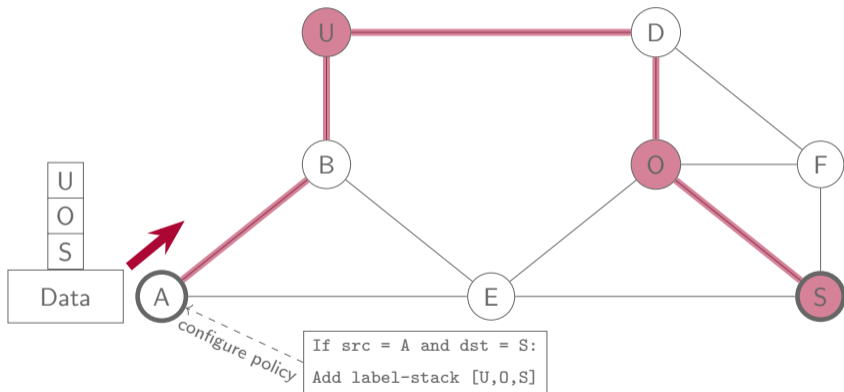
# Segment Routing (SR) in a Nutshell:

Control a packet's path by defining **interim destinations/waypoints**:



# Segment Routing (SR) in a Nutshell:

Control a packet's path by defining **interim destinations/waypoints**:



Terminology:

**SR Policy:** "Rule" determining which segments to add to a packet

# Common Operational Challenges: *“Mom, the Internet is broken!!!”*



Hardware Failures &  
Traffic Changes




# Common Operational Challenges: *“Mom, the Internet is broken!!!”*



Hardware Failures &  
Traffic Changes



LOADING



Service Deterioration  
or Disruption


# Common Operational Challenges: *“Mom, the Internet is broken!!!”*



Hardware Failures &  
Traffic Changes



LOADING



Service Deterioration  
or Disruption



Unhappy Customers

Such issues can be addressed with SR TE!

Such issues can be addressed with SR TE!

## A) Proactive Protection

Idea: “**Hedge** networks against said events so they don’t become a problem.”

“Preparation is the key to success!”

— A. G. Bell




Such issues can be addressed with SR TE!

## A) Proactive Protection

Idea: “**Hedge** networks against said events so they don’t become a problem.”

“Preparation is the key to success!”

— A. G. Bell




Schüller et al., “Failure Resiliency with only a few Tunnels – Enabling Segment Routing for Traffic Engineering”, IEEE/ACM ToN, 2021

## B) Reactive Restoration (aka “Tactical TE”)

Idea: “Quickly **reconfigure** networks to mitigate the critical event.”

“Gotta go fast!”

— Sonic



Gay et al., “Expect the Unexpected: Sub-second Optimization for Segment Routing”, IEEE INFOCOM, 2017

Such issues can be addressed with SR TE!

## A) Proactive Protection

Idea: “**Hedge** networks against said events so they don’t become a problem.”

“Preparation is the key to success!”

— A. G. Bell




Schüller et al., “Failure Resiliency with only a few Tunnels – Enabling Segment Routing for Traffic Engineering”, IEEE/ACM ToN, 2021

## B) Reactive Restoration (aka “Tactical TE”)

Idea: “Quickly **reconfigure** networks to mitigate the critical event.”

“Gotta go fast!”

— Sonic



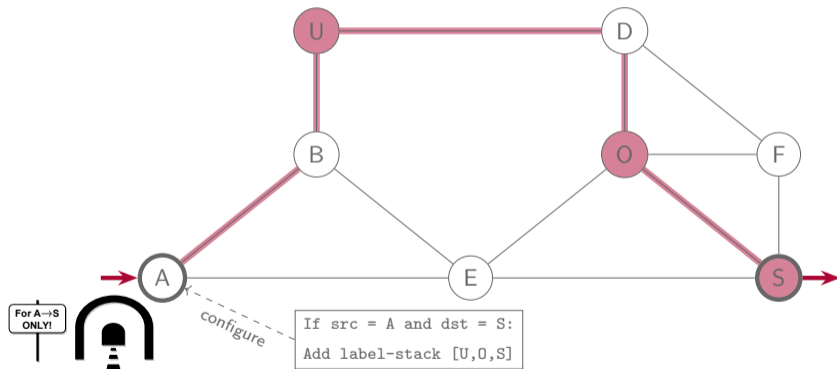
Gay et al., “Expect the Unexpected: Sub-second Optimization for Segment Routing”, IEEE INFOCOM, 2017

Limitation: Existing approaches all rely on **end-to-end** SR!

SR Policy = *“E2E tunnel for a **single** traffic demand”*

Limitation: Existing approaches all rely on **end-to-end** SR!

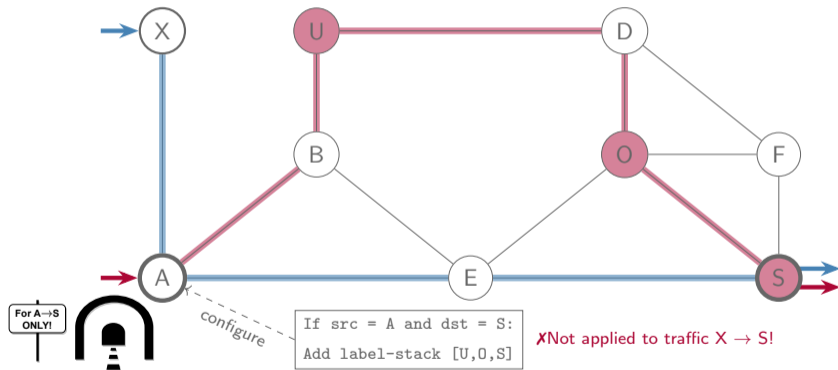
SR Policy = "E2E tunnel for a **single** traffic demand"





Limitation: Existing approaches all rely on **end-to-end** SR!

SR Policy = "E2E tunnel for a **single** traffic demand"



What's the problem with that?

## What's the problem with that?

E2E SR requires a **dedicated policy for each traffic demand** to be detoured.

→ TE solutions with **hundreds or thousands of policies!**

## What's the problem with that?

E2E SR requires a **dedicated policy for each traffic demand** to be detoured.

→ TE solutions with **hundreds or thousands of policies!**

→ Hard to understand & verify!



# What's the problem with that?

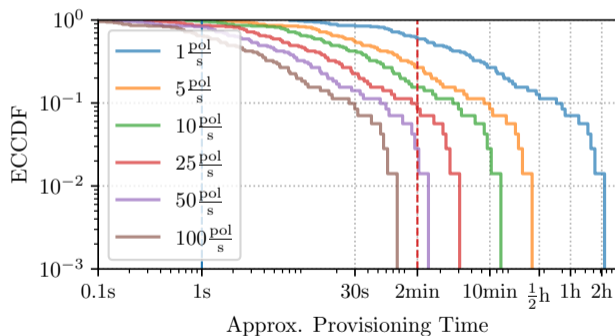
E2E SR requires a **dedicated policy** for each traffic demand to be detoured.

→ TE solutions with **hundreds or thousands of policies!**

→ Hard to understand & verify!



→ Slow roll-out: (i.e. multiple minutes)



# What's the problem with that?


E2E SR requires a **dedicated policy** for each traffic demand to be detoured.

→ TE solutions with **hundreds or thousands of policies!**

→ Hard to understand & verify!



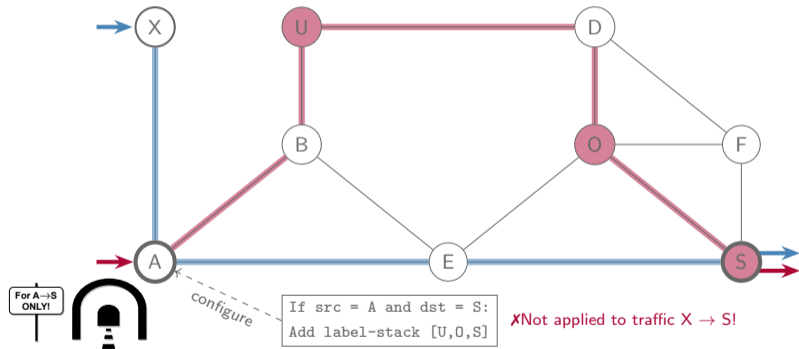
→ Slow roll-out: (i.e. multiple minutes)



→ **Limitations** regarding the practical usability of reactive restoration!

# How to overcome these issues?

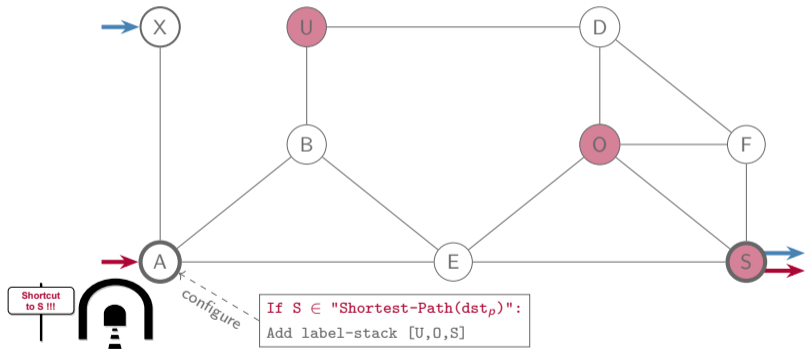
Problem: Dedicated "demand-bound" policies



# How to overcome these issues? Use SR Midpoint Optimization!

Problem: ~~Dedicated “demand-bound” policies~~

→ More flexible steering mechanisms (i.e. IGP Shortcut)

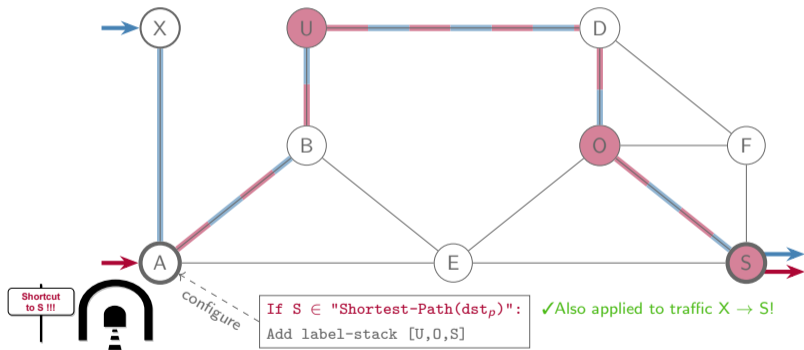




# How to overcome these issues? Use SR Midpoint Optimization!

Problem: ~~Dedicated~~ "demand-bound" policies

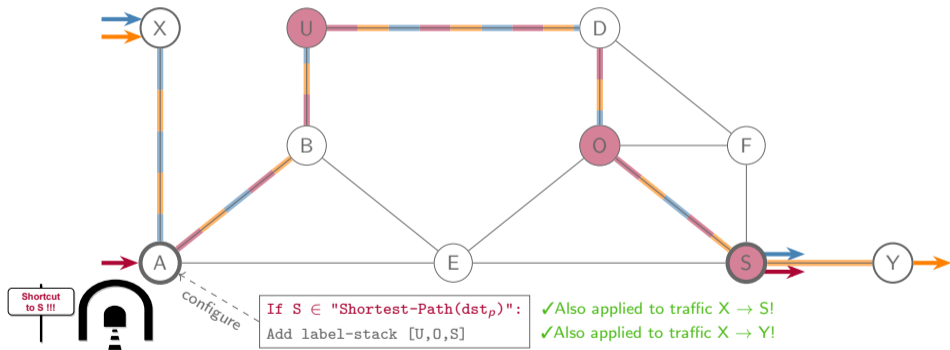
→ More flexible steering mechanisms (i.e. IGP Shortcut)



# How to overcome these issues? Use SR Midpoint Optimization!

Problem: ~~Dedicated~~ "demand-bound" policies


→ More flexible steering mechanisms (i.e. IGP Shortcut)



# How to overcome these issues? Use SR Midpoint Optimization!

Problem: ~~Dedicated~~ "demand-bound" policies

→ More flexible steering mechanisms (i.e. IGP Shortcut)

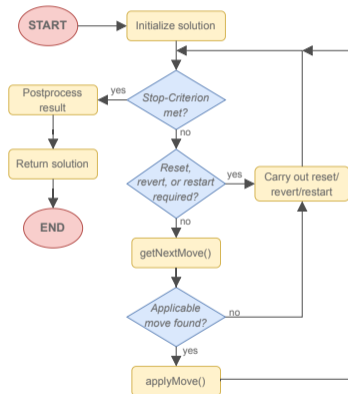



→ Detouring of **multiple demands** with only **a single policy!**

How to compute suitable solutions fast?

# How to compute suitable solutions fast?

## New algorithm: **Midpoint Optimization Local Search (MOLS)**




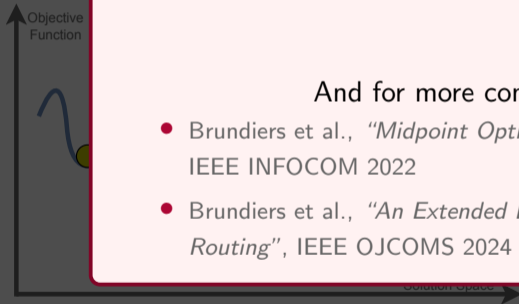
# How to compute suitable solutions fast?

New algorithm: **Midpoint Optimization Local Search (MOLS)**

For more details, see our paper!

And for more context on SR MO, see:

- Brundiens et al., *"Midpoint Optimization for Segment Routing"*, IEEE INFOCOM 2022
- Brundiens et al., *"An Extended Look at Midpoint Optimization for Segment Routing"*, IEEE OJCOMS 2024



# Evaluation Setup: (simulation based)



Topology Information

Source	Destination	Traffic
RTM	LAX	42
OS	RTM	5
...	...	...

Traffic Matrix

Input

Further Optional  
Requirements & Constraints

$$\begin{aligned} \min & \theta \\ \text{s.t.} & \sum_{p \in P_d} x_p^d = 1 \quad \forall d \in D \\ & \sum_{d \in D} \sum_{p \in P_d} \delta_p(e) x_p^d t_d \leq \theta c_e \quad \forall e \in E \\ & x_p^d \geq 0 \quad \forall d \in D \\ & \quad \forall p \in P \end{aligned}$$

Optimization Model/Algorithm

Output

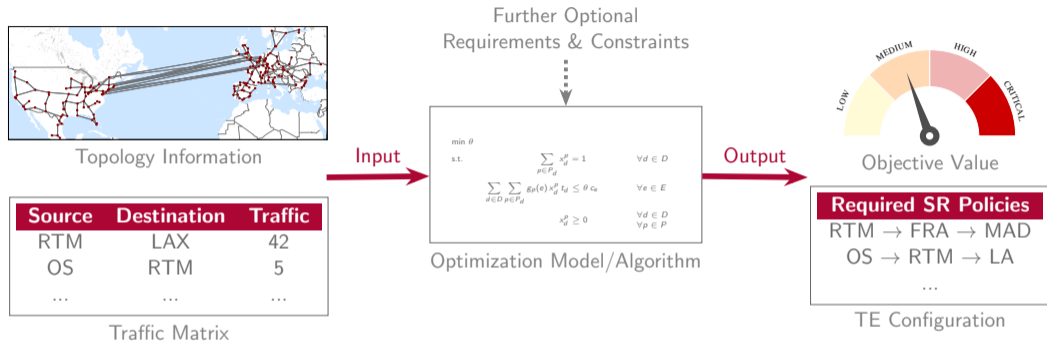


Objective Value

**Required SR Policies**  
RTM → FRA → MAD  
OS → RTM → LA  
...

TE Configuration

# Evaluation Setup: (simulation based)



## Data Sources:

- Semi-Artificial *Repetita/Topology Zoo* data
- Real-world data from a **Tier-1 ISP backbone**




MOLS allows to handle failure scenarios exceptionally well!

MOLS allows to handle failure scenarios exceptionally well!

Our algorithm resolves **over 99%** of failure scenarios

MOLS allows to handle failure scenarios exceptionally well (and fast)!

Our algorithm resolves **over 99%** of failure scenarios, mostly in **sub-second** fashion!



The use of MO also greatly improves provisioning times!

The use of MO also greatly improves provisioning times!

MOLS: Around **97% reduction** of policy numbers compared to E2E SR!


→ often just a **low double-digit** number

# The use of MO also greatly improves provisioning times!

**MOLS: Around 97% reduction of policy numbers compared to E2E SR!**

→ often just a **low double-digit** number

→ **Much faster roll-out:** (generally < 1min)





# The use of MO also greatly improves provisioning times!

MOLS: Around **97% reduction** of policy numbers compared to E2E SR!

→ often just a **low double-digit** number

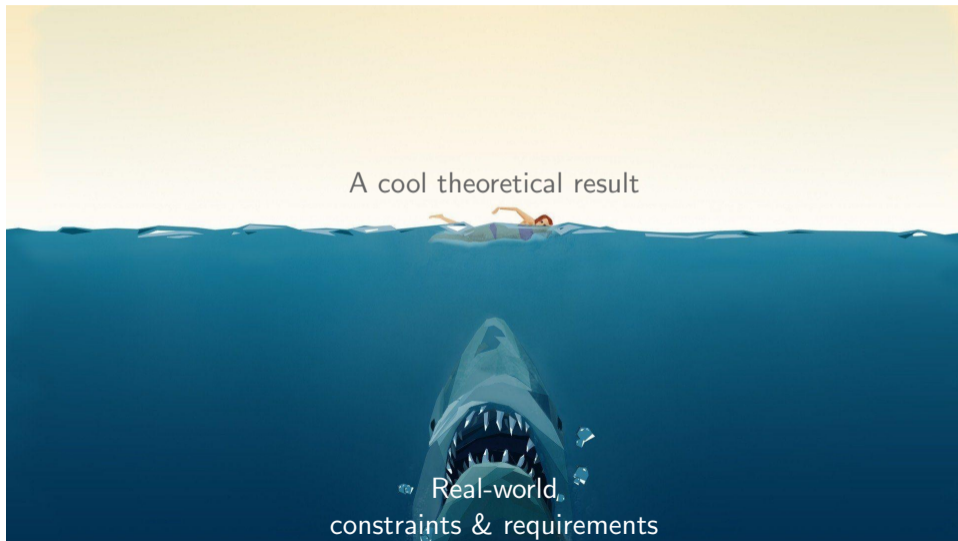
→ **Much faster roll-out:** (generally < 1min)



Fast computation, verification & provisioning

→ **Greatly improved "Time-to-Repair"!**

# Bridging the Gap: Theory → Practice





# Important Constraints & Requirements

- Policy Numbers
- Routing Loops

- MSD Limits
- Traffic Splitting

- Latency Bounds
- ...



# Important Constraints & Requirements

- Policy Numbers

- MSD Limits

- Latency Bounds

**Our approach considers ALL these requirements!**

**Optimization Algorithm**  
(+ Real-world Requirements)

+

**Technical Implementations**  
(e.g., Juniper or Cisco)

=

**Practically usable  
results!**

Hey, I can *actually*  
use this stuff!



Recap of key findings/contributions:

By utilizing the **SR Midpoint Optimization** concept, our MOLS algorithm ...

## Recap of key findings/contributions:

By utilizing the **SR Midpoint Optimization** concept, our MOLS algorithm ...

... enables **sub-second** computation of solutions **on-par or better** than E2E SR,

## Recap of key findings/contributions:

By utilizing the **SR Midpoint Optimization** concept, our MOLS algorithm ...

... enables **sub-second** computation of solutions **on-par or better** than E2E SR,

... while considerably **reducing policy numbers** (up to 97%),


## Recap of key findings/contributions:

By utilizing the **SR Midpoint Optimization** concept, our MOLS algorithm ...

... enables **sub-second** computation of solutions **on-par or better** than E2E SR,

... while considerably **reducing policy numbers** (up to 97%),

... which translates to substantially **improved provisioning times** and, thus, a faster treatment of critical events,



## Recap of key findings/contributions:


By utilizing the **SR Midpoint Optimization** concept, our MOLS algorithm ...

... enables **sub-second** computation of solutions **on-par or better** than E2E SR,

... while considerably **reducing policy numbers** (up to 97%),

... which translates to substantially **improved provisioning times** and, thus, a faster treatment of critical events,

... and all that while even respecting **real-world constraints!**



# Recap of key findings/contributions:


By utilizing the **SR Midpoint Optimization** concept, our MOLS algorithm ...

... enables **sub-second** computation of solutions **on-par or better** than E2E SR,

... while considerably **reducing** pol

That's a **major improvement**,  
over the current State-of-the-Art!

... which translates to sub-second provisioning times  
and, thus, a faster tree



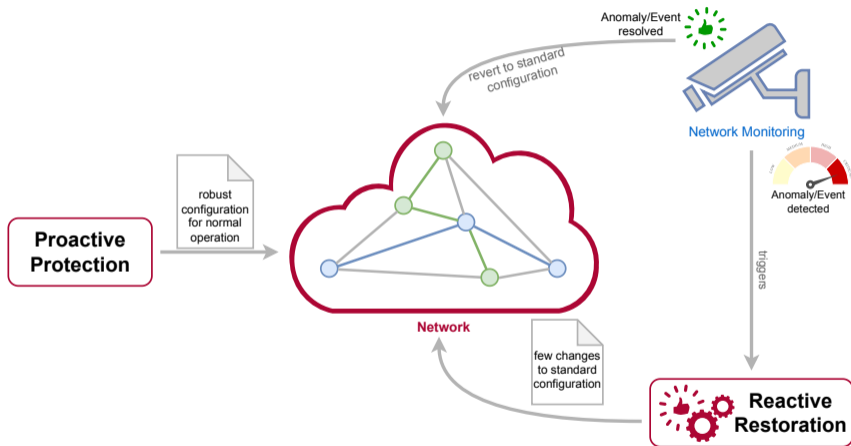
... and all that while even respecting **real-world constraints!**






# A sneak-peak into (potential) future work:

Combine MOLS with MO-based proactive protection approaches!




A sneak-peak into (potential) future work:

Combine MOLS with MO-based proactive protection approaches!



**The Final Takeaway:**

**SR MO enables some pretty cool stuff!  
Consider using it!**



A sneak-peak into (potential) future work:

## The Final Takeaway:

**SR MO enables some pretty cool stuff!**  
**Consider using it!**



Time's up? Find me at  
[sys.cs.uos.de/brundiars](https://sys.cs.uos.de/brundiars)  
or contact me directly:  
[brundiars@uos.de](mailto:brundiars@uos.de)



## Appendix & Backup-Slides

# Image Sources/Credits

- The “scientist tux” (penguin) on [Slide 2](#) is taken from:  
M. Barbieri, 2010, “*Tux version of scientist Lazzaro Spallanzani*”, Wikimedia Commons.  
online: [https://commons.wikimedia.org/wiki/File:Tux\\_Spallanzani.svg](https://commons.wikimedia.org/wiki/File:Tux_Spallanzani.svg)  
License: Massimo Barbieri, CC BY-SA 3.0 <https://creativecommons.org/licenses/by-sa/3.0>, via Wikimedia Commons
- Some of the illustrations on [Slide 4](#) were created using the *AI Image Generator of DeepAI* (<https://deepai.org/machine-learning-model/text2img>)