Timeless Timing Attacks:

Exploiting Concurrency to Leak Secrets over Remote Connections

Tom Van Goethem, Christina Pöpper,
Wouter Joosen, Mathy Vanhoef

DistriNet

جامعة نيويورك أبوظبي
NYU ABU DHABI

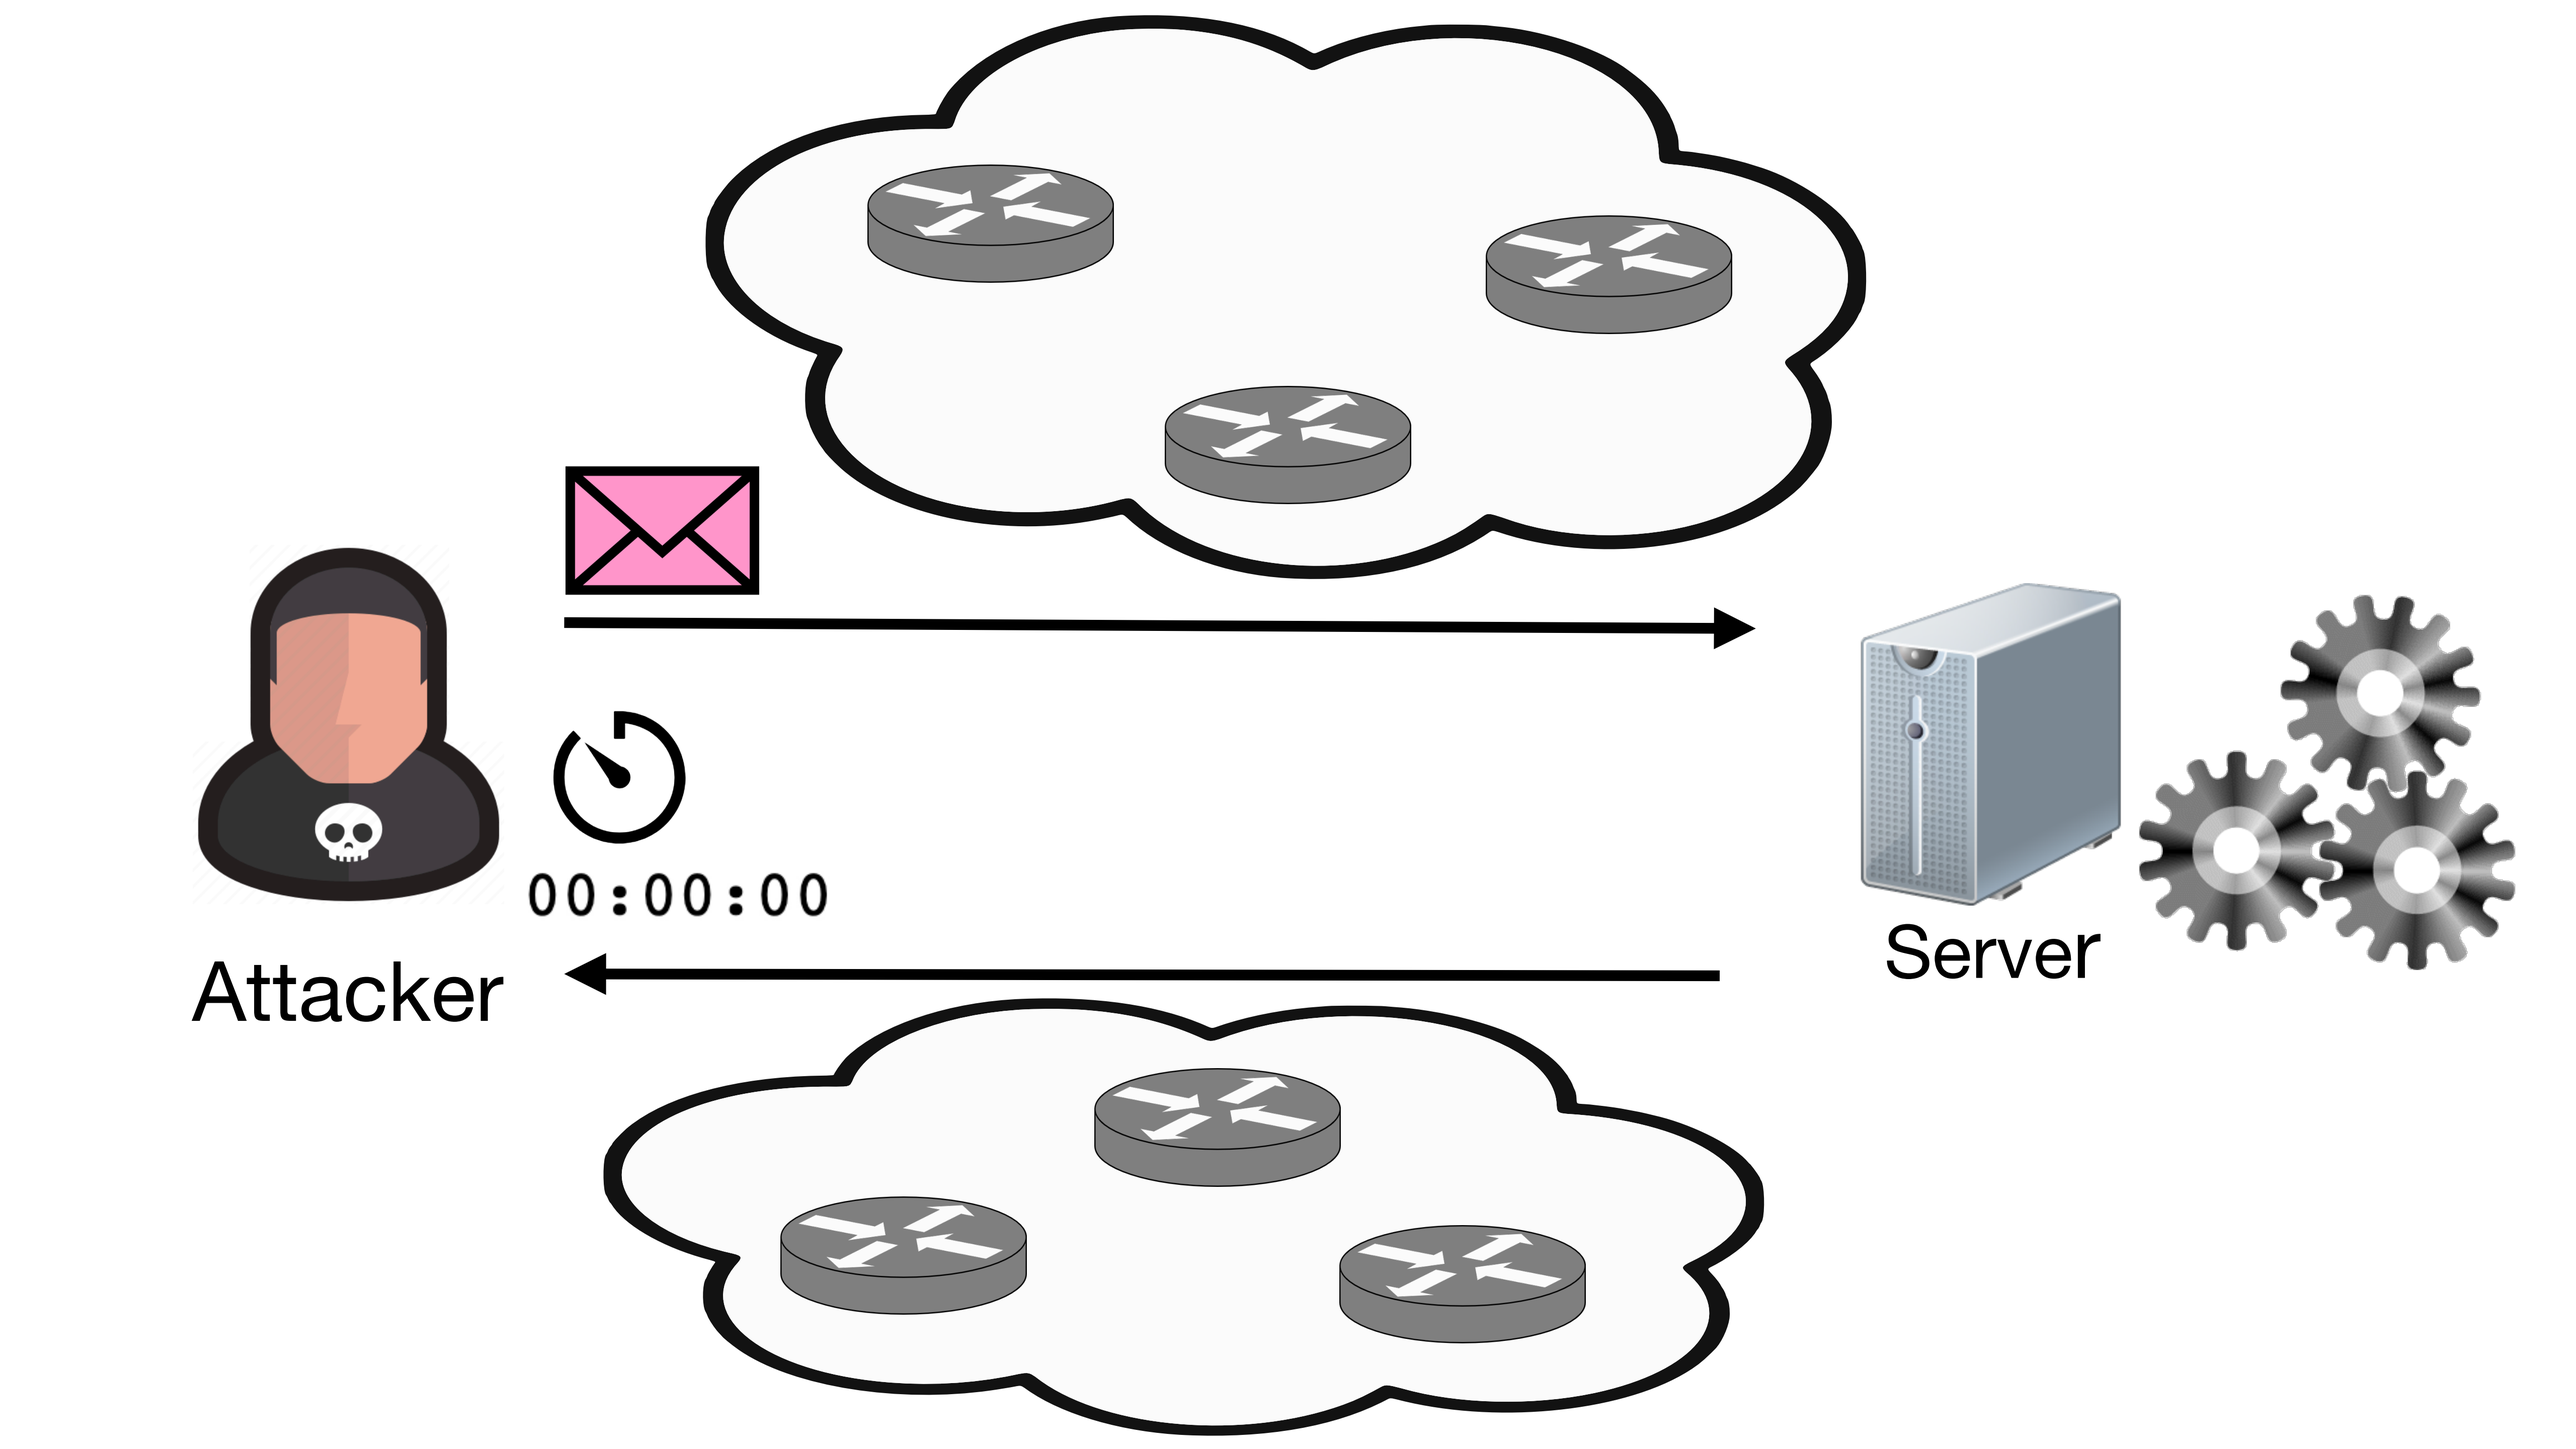
00:00:00
Attacker
Server

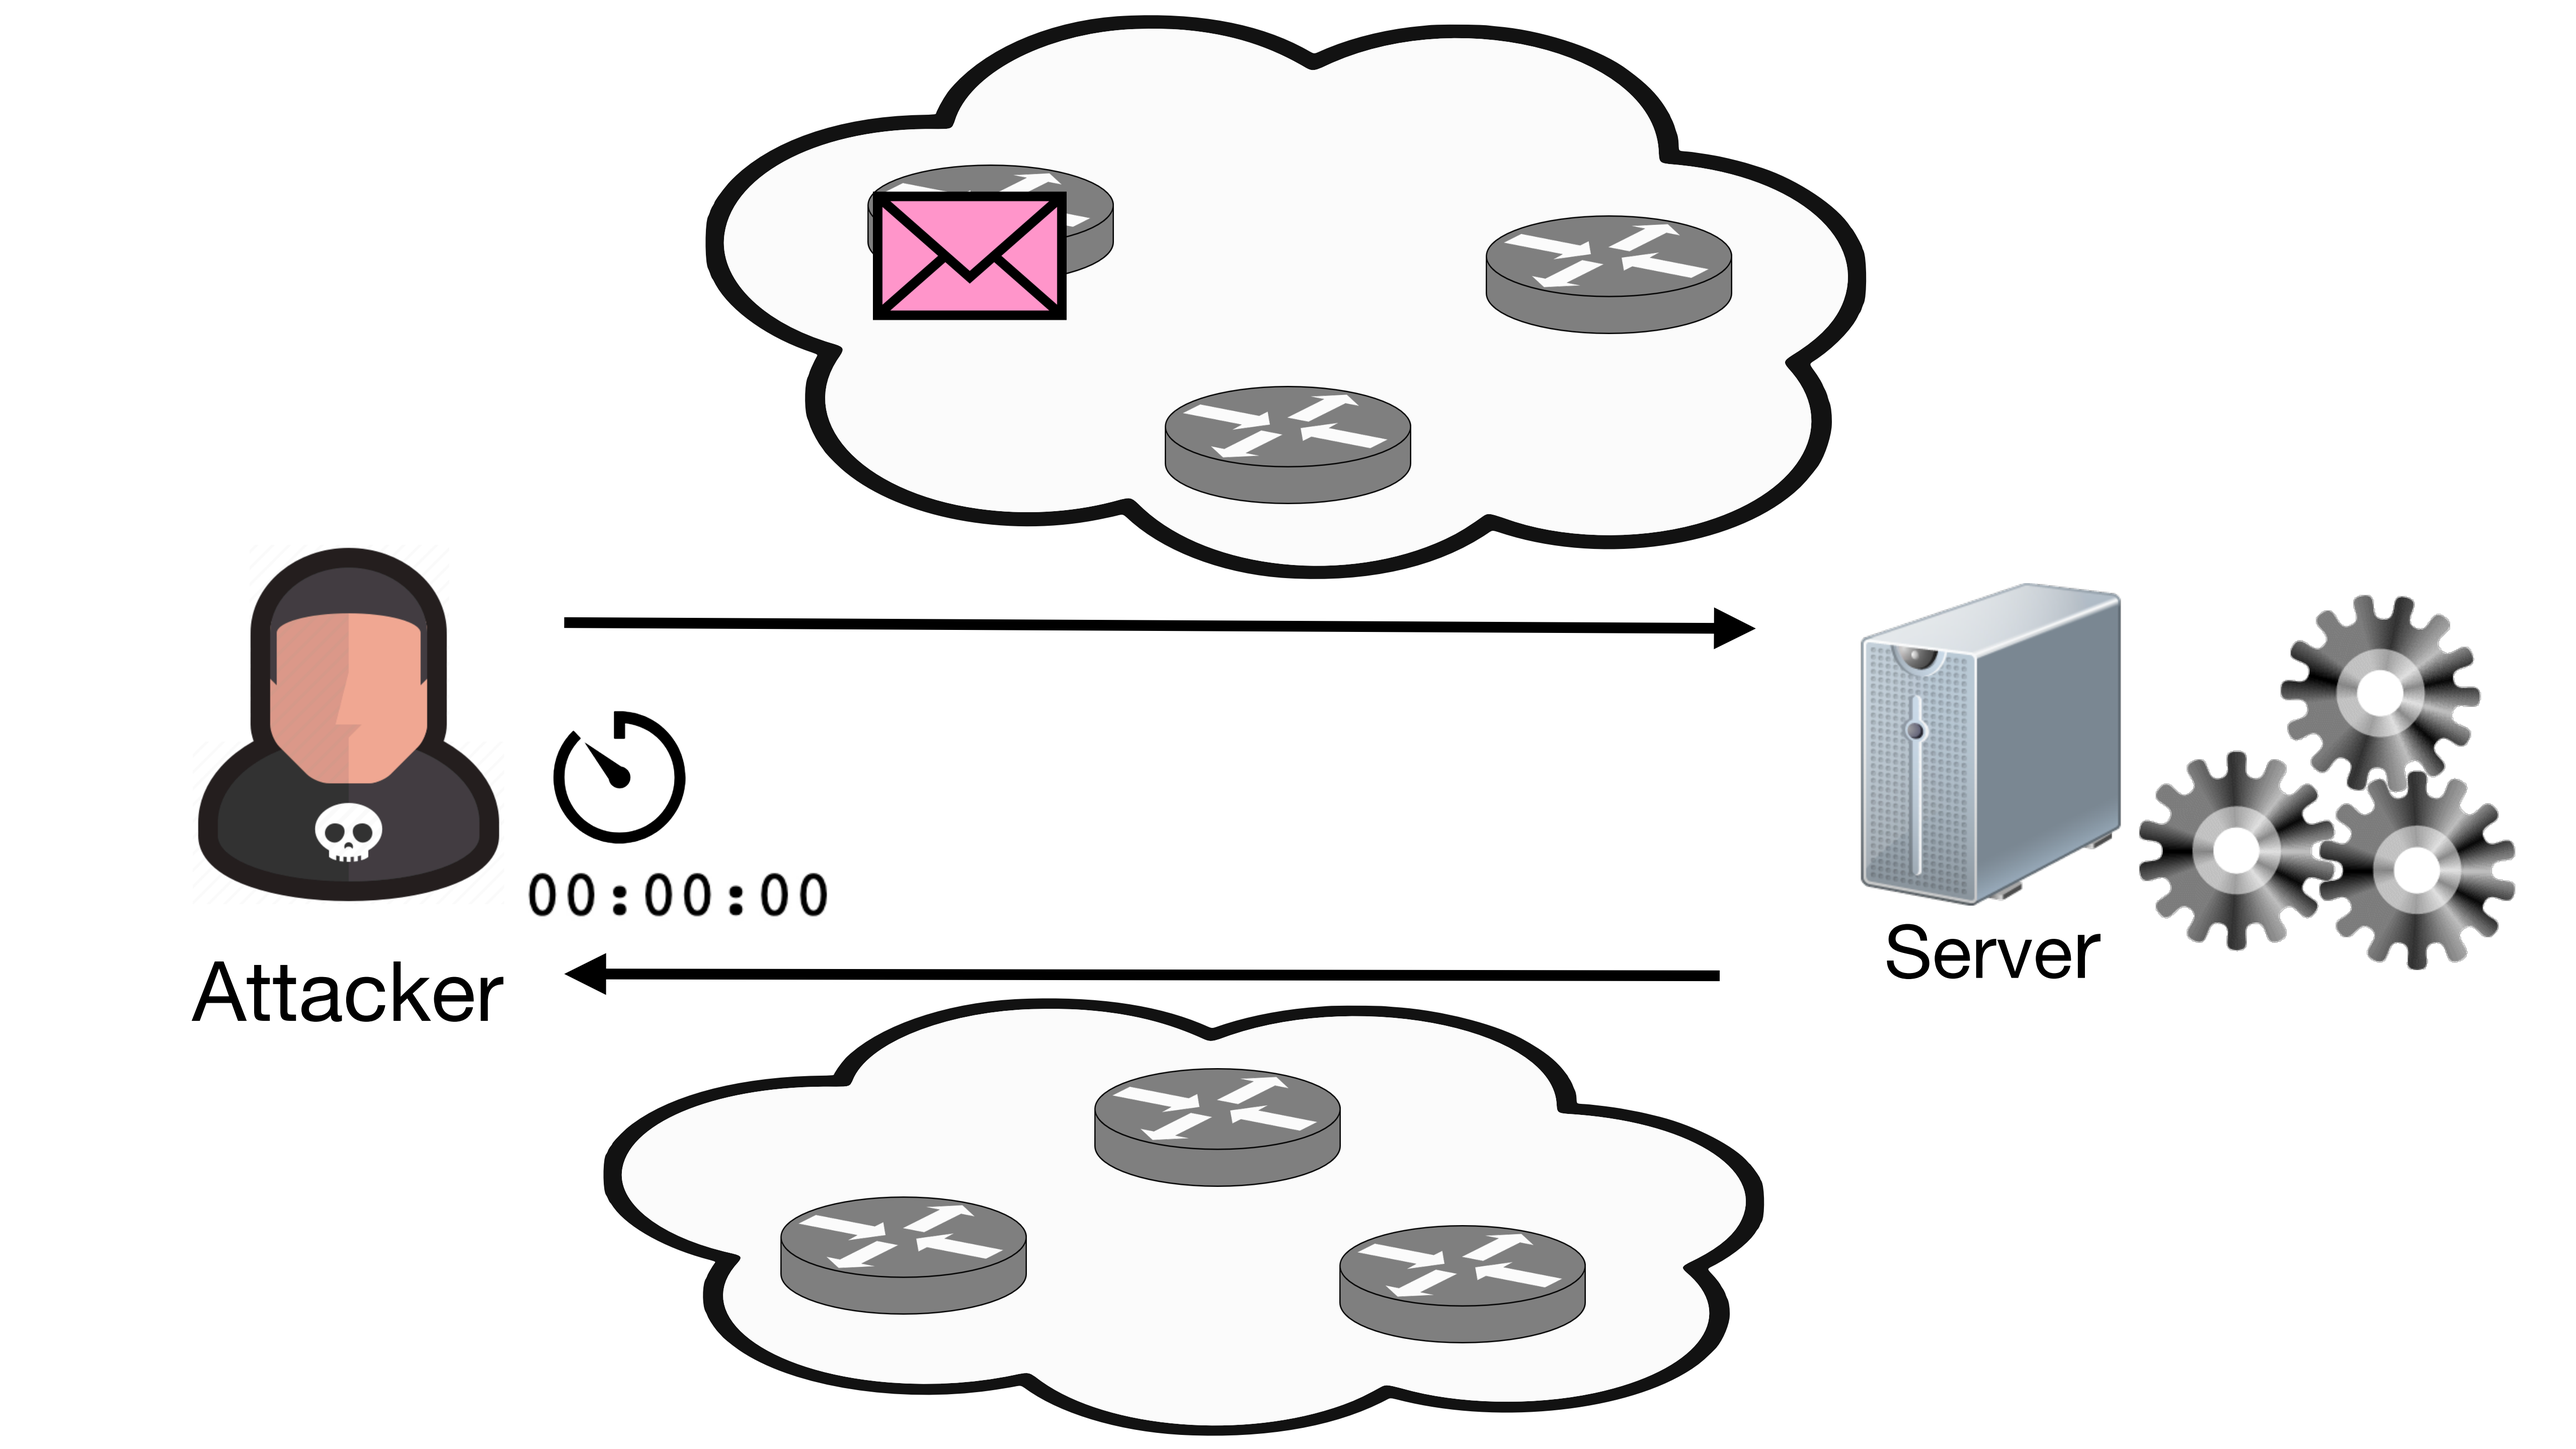
00:00:00
Attacker
Server

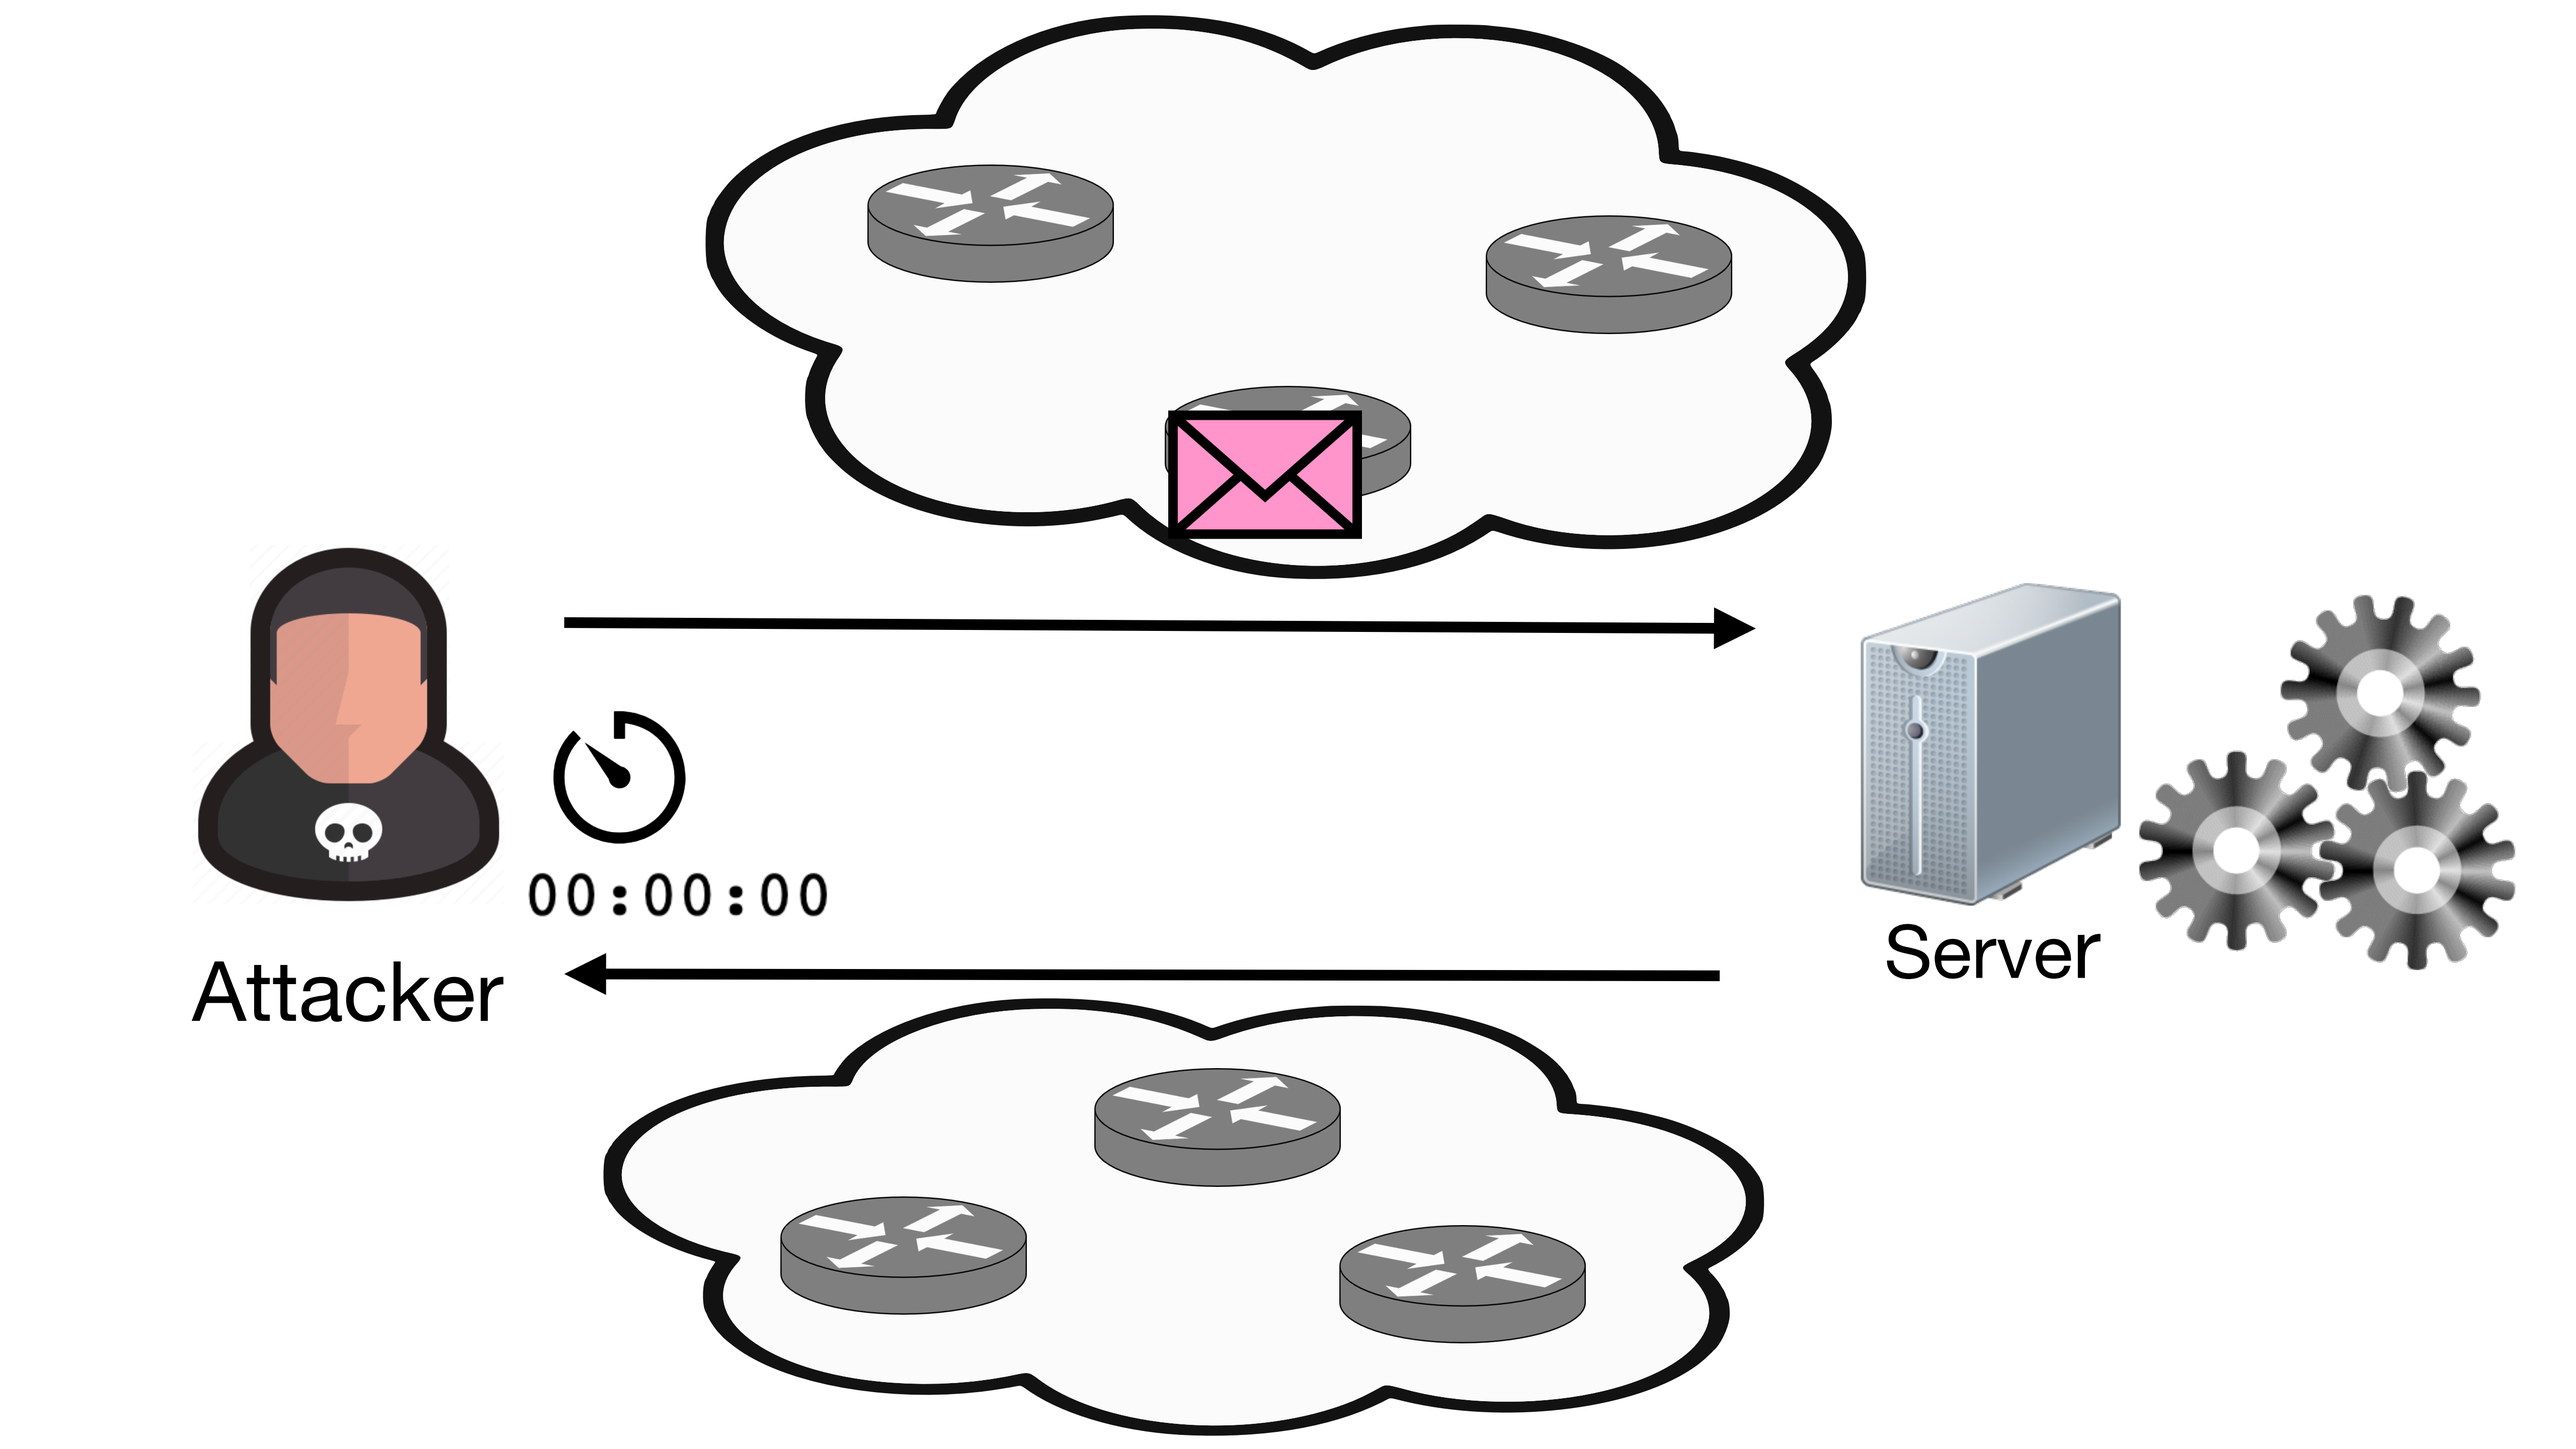
00:00:00
Attacker
Server

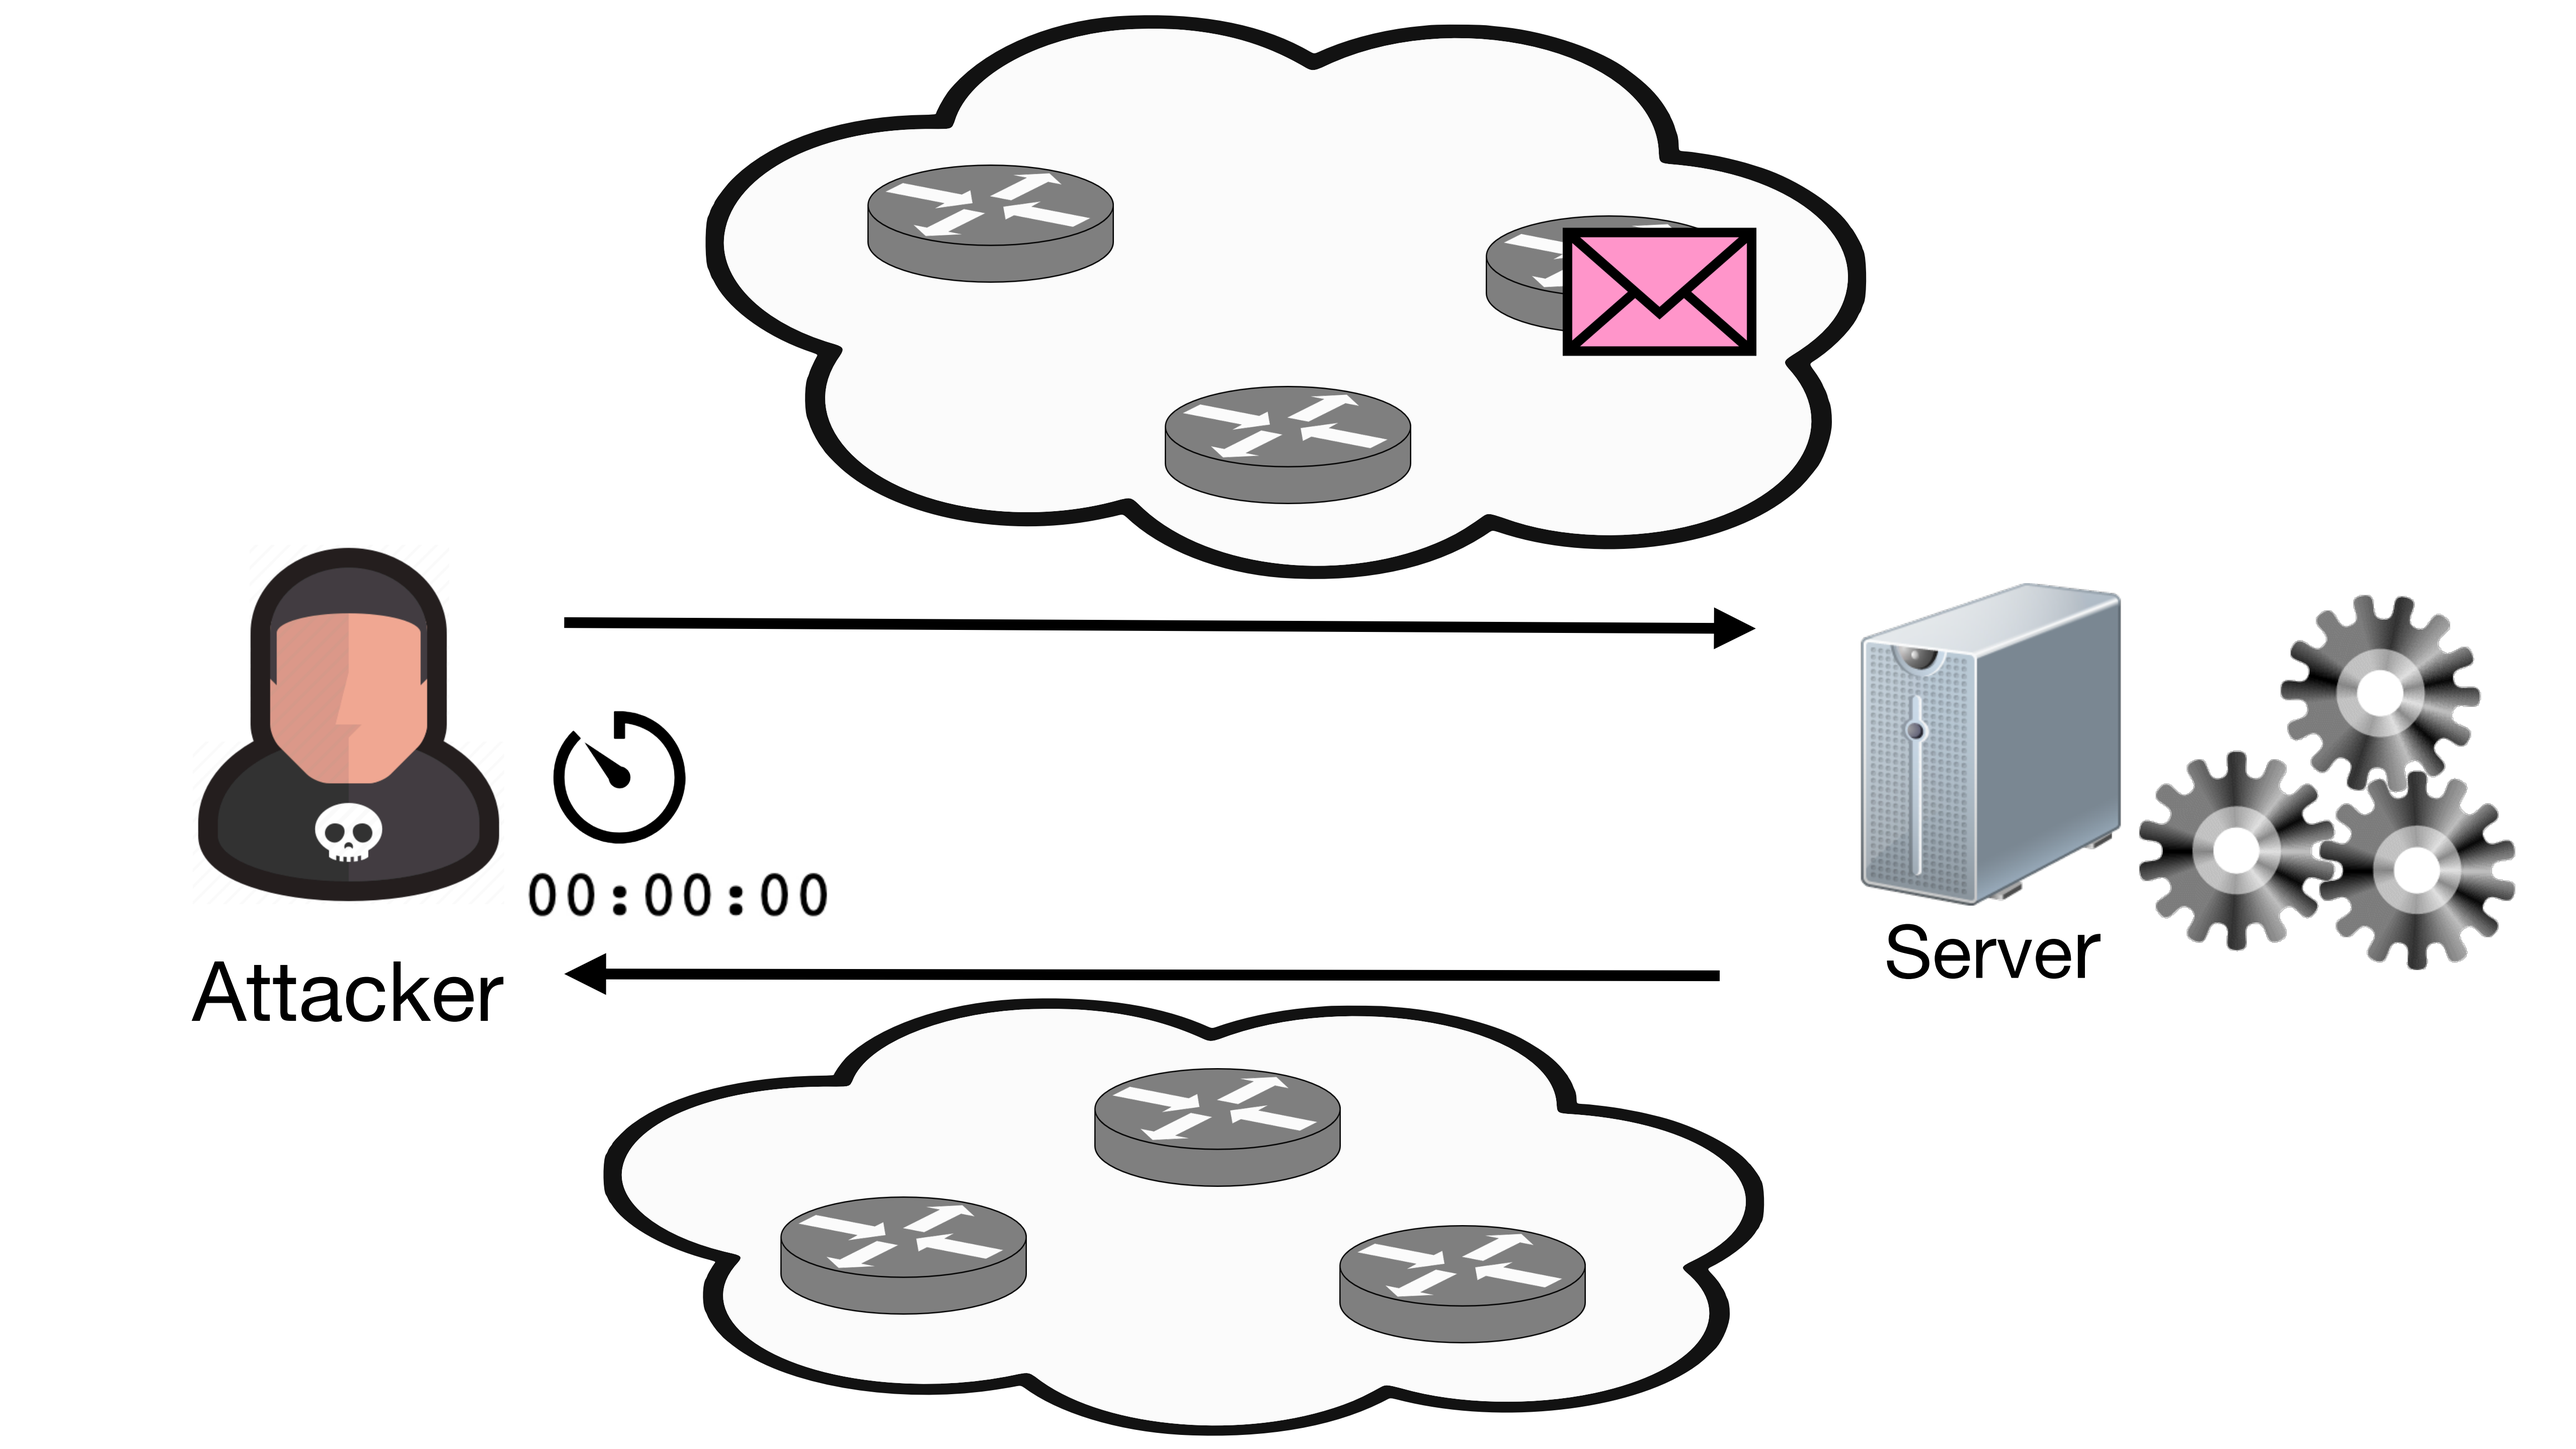
00:00:00
Attacker
Server

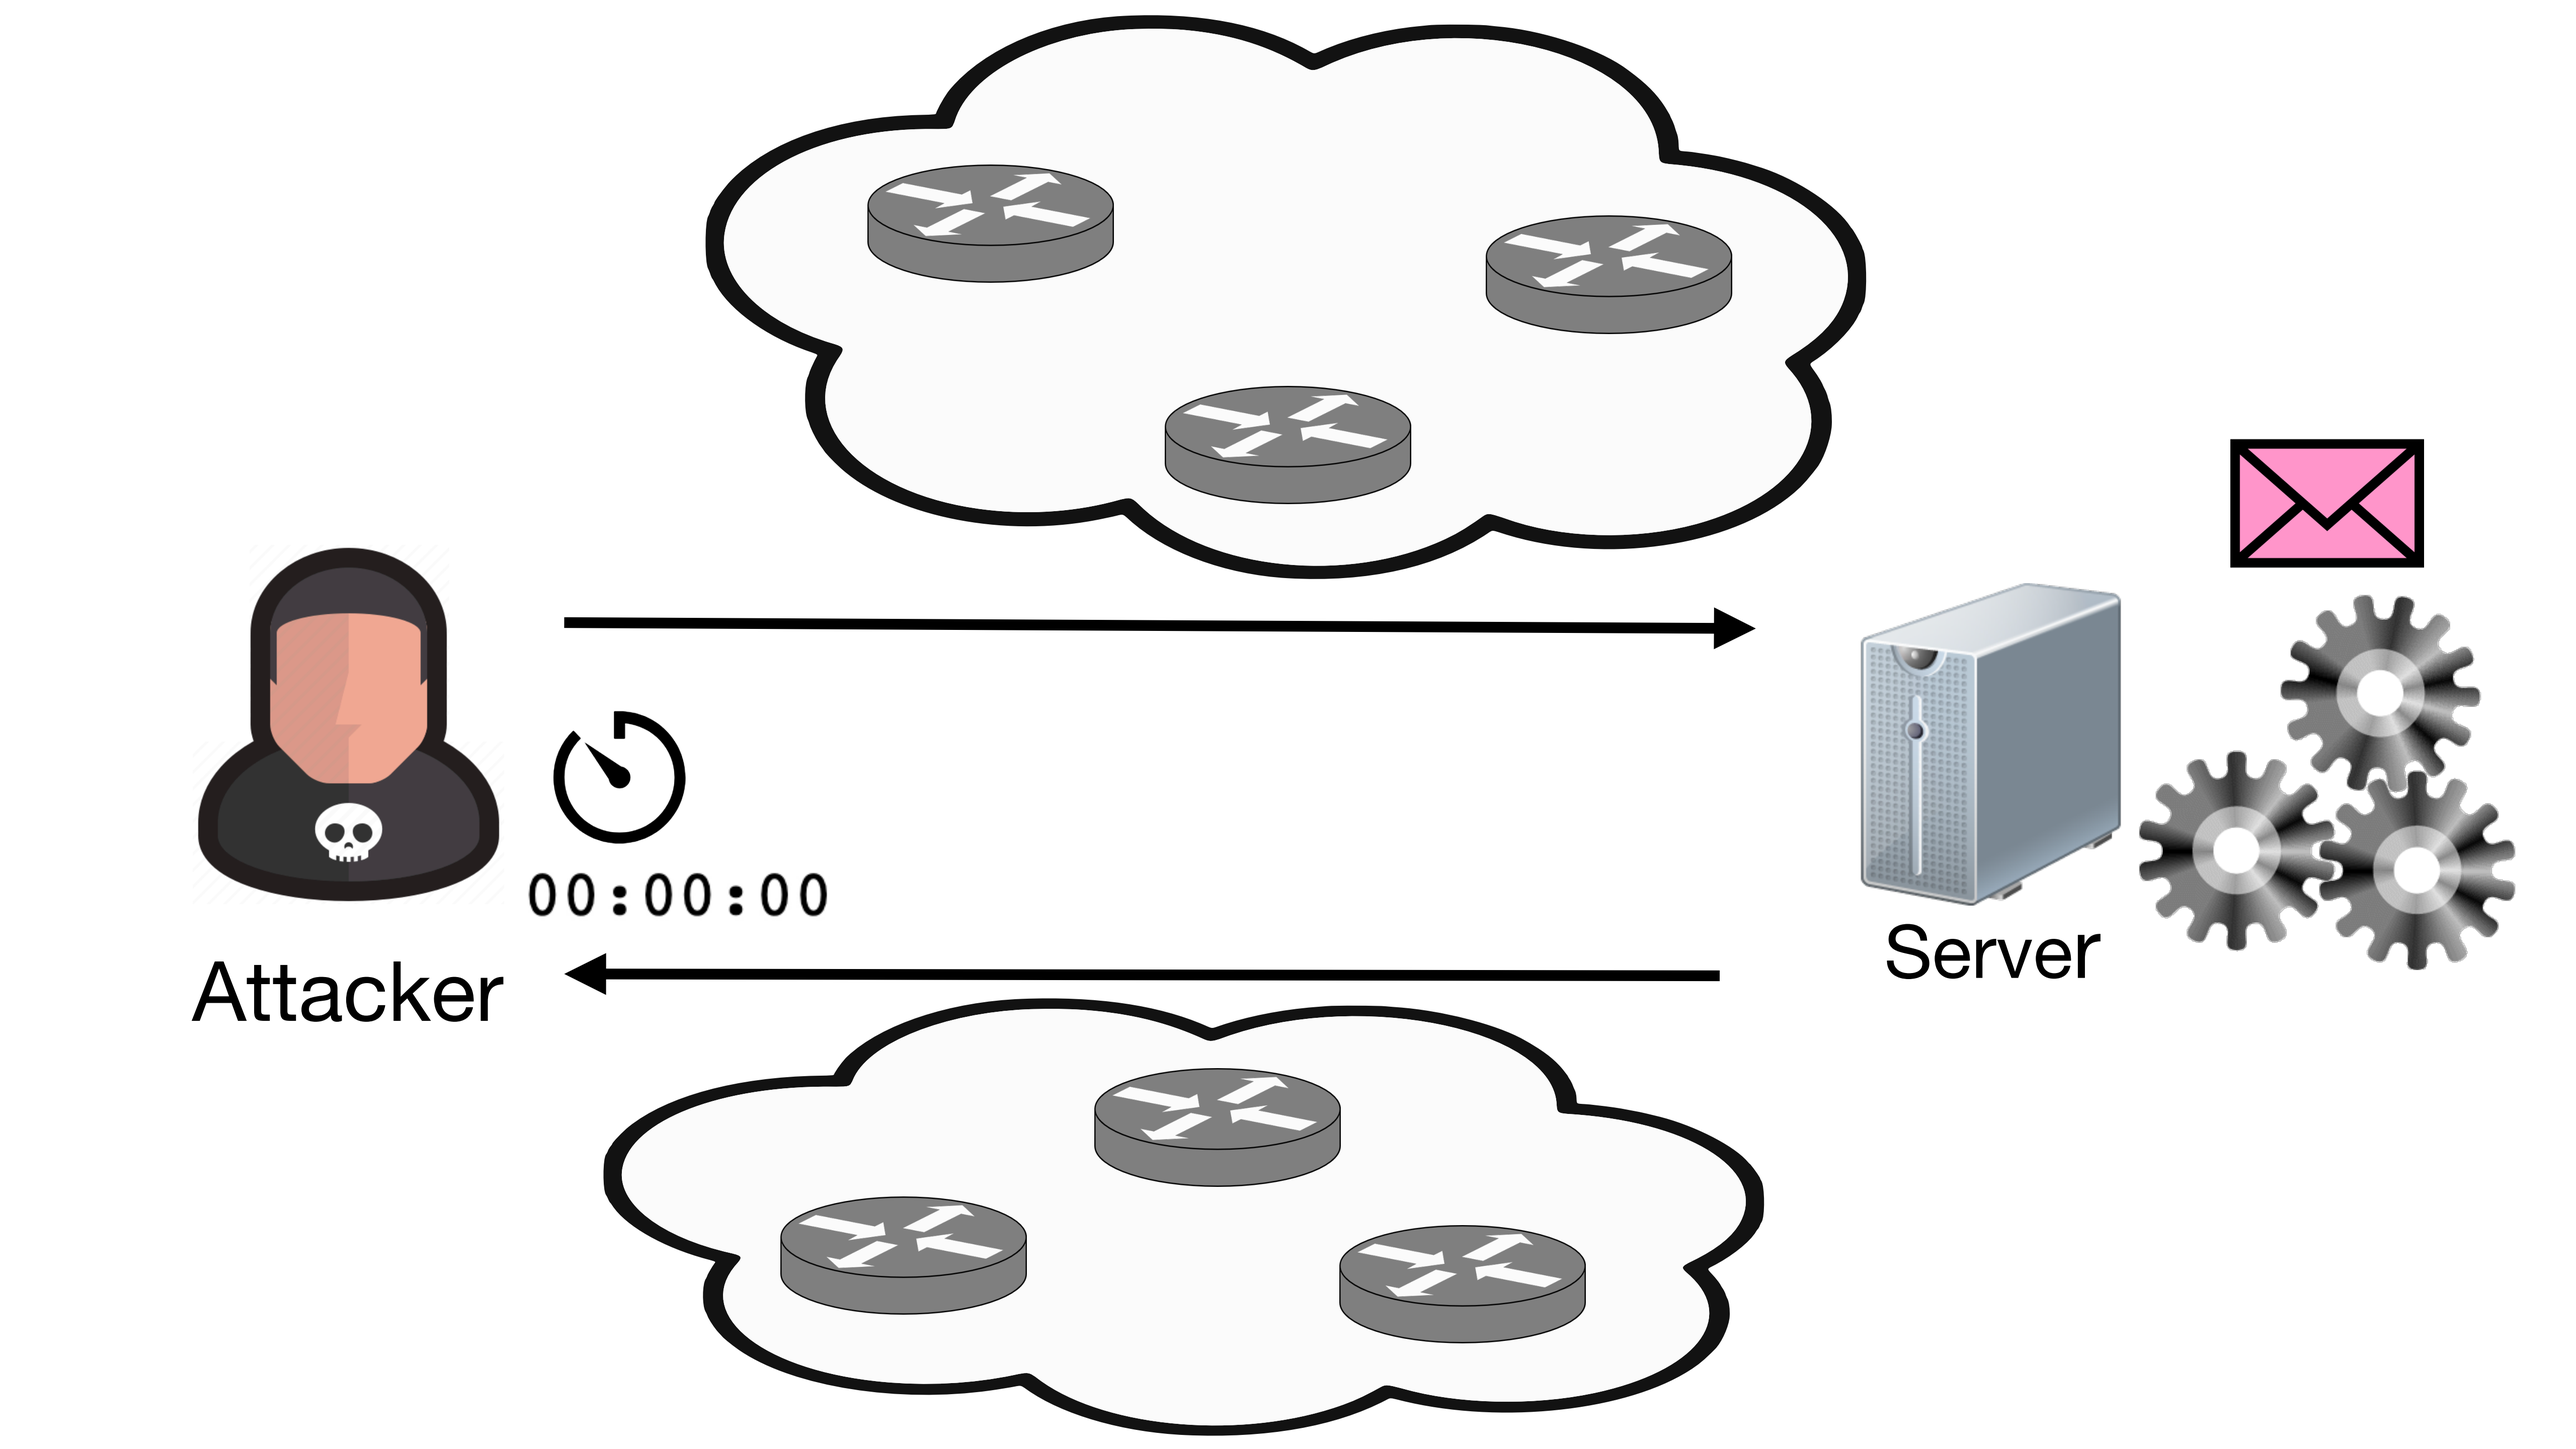
00:00:00
Attacker
Server

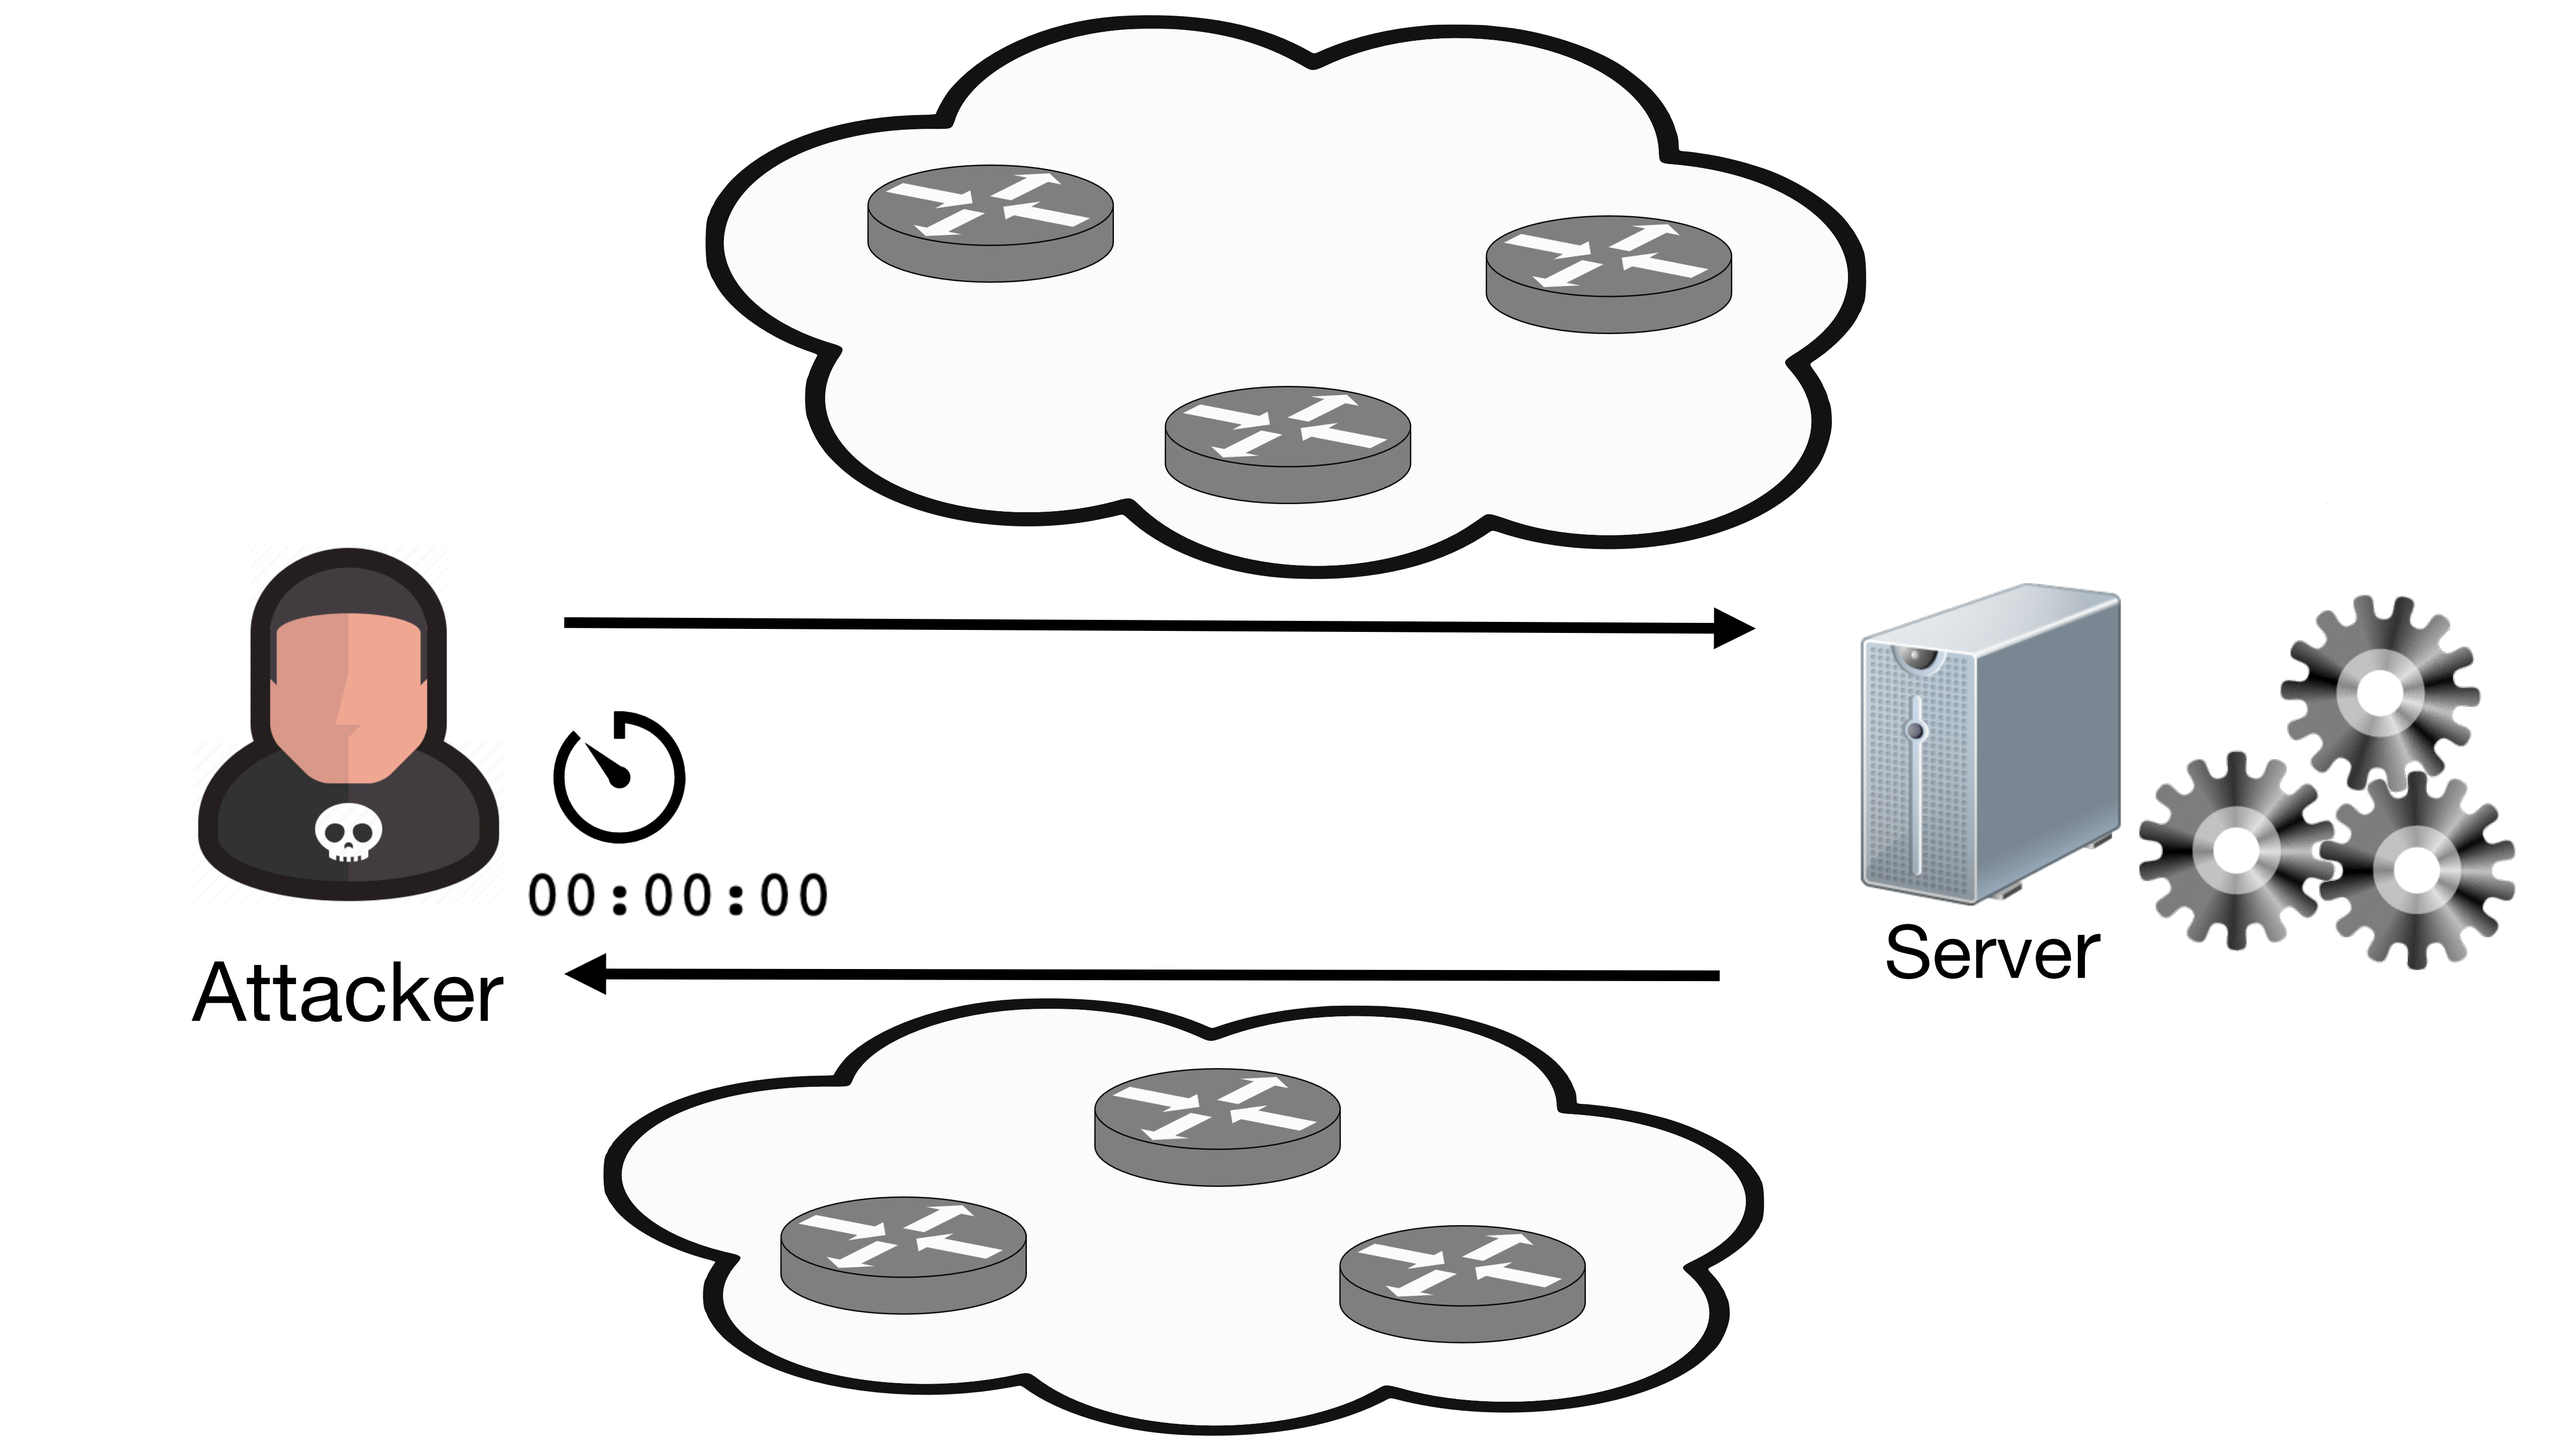
00:00:00
Attacker
Server

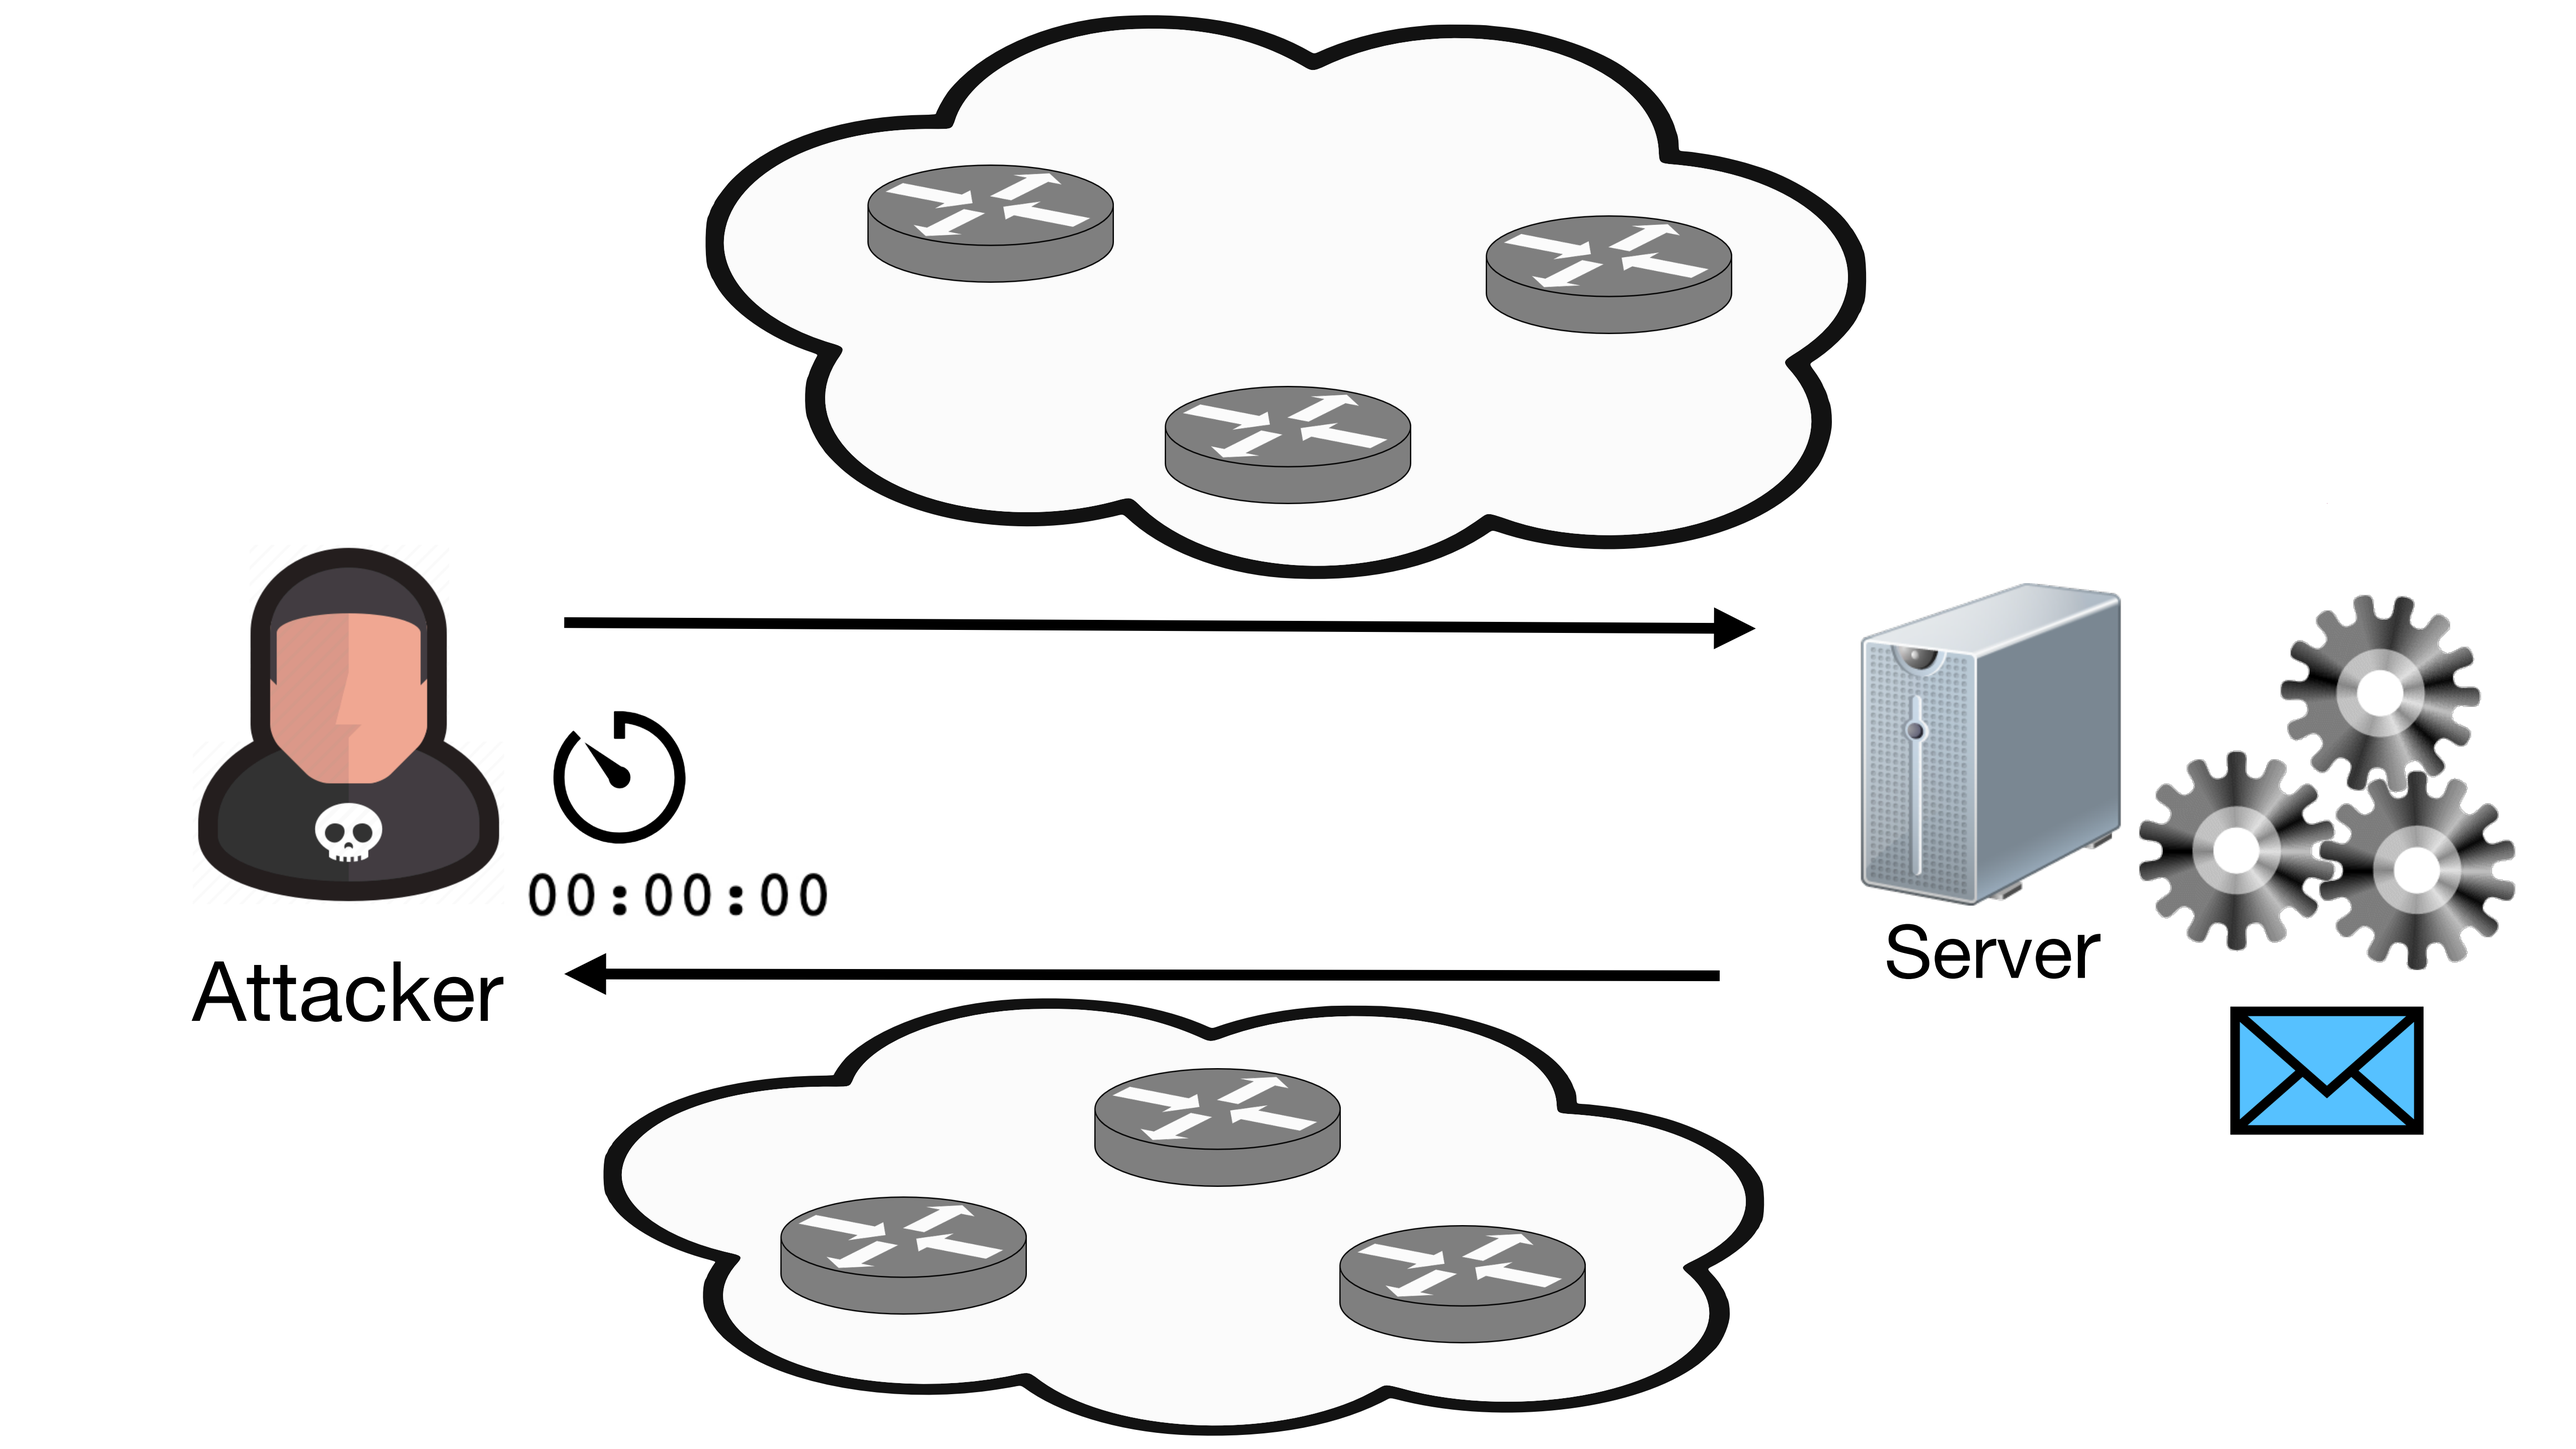
00:00:00
Attacker
Server

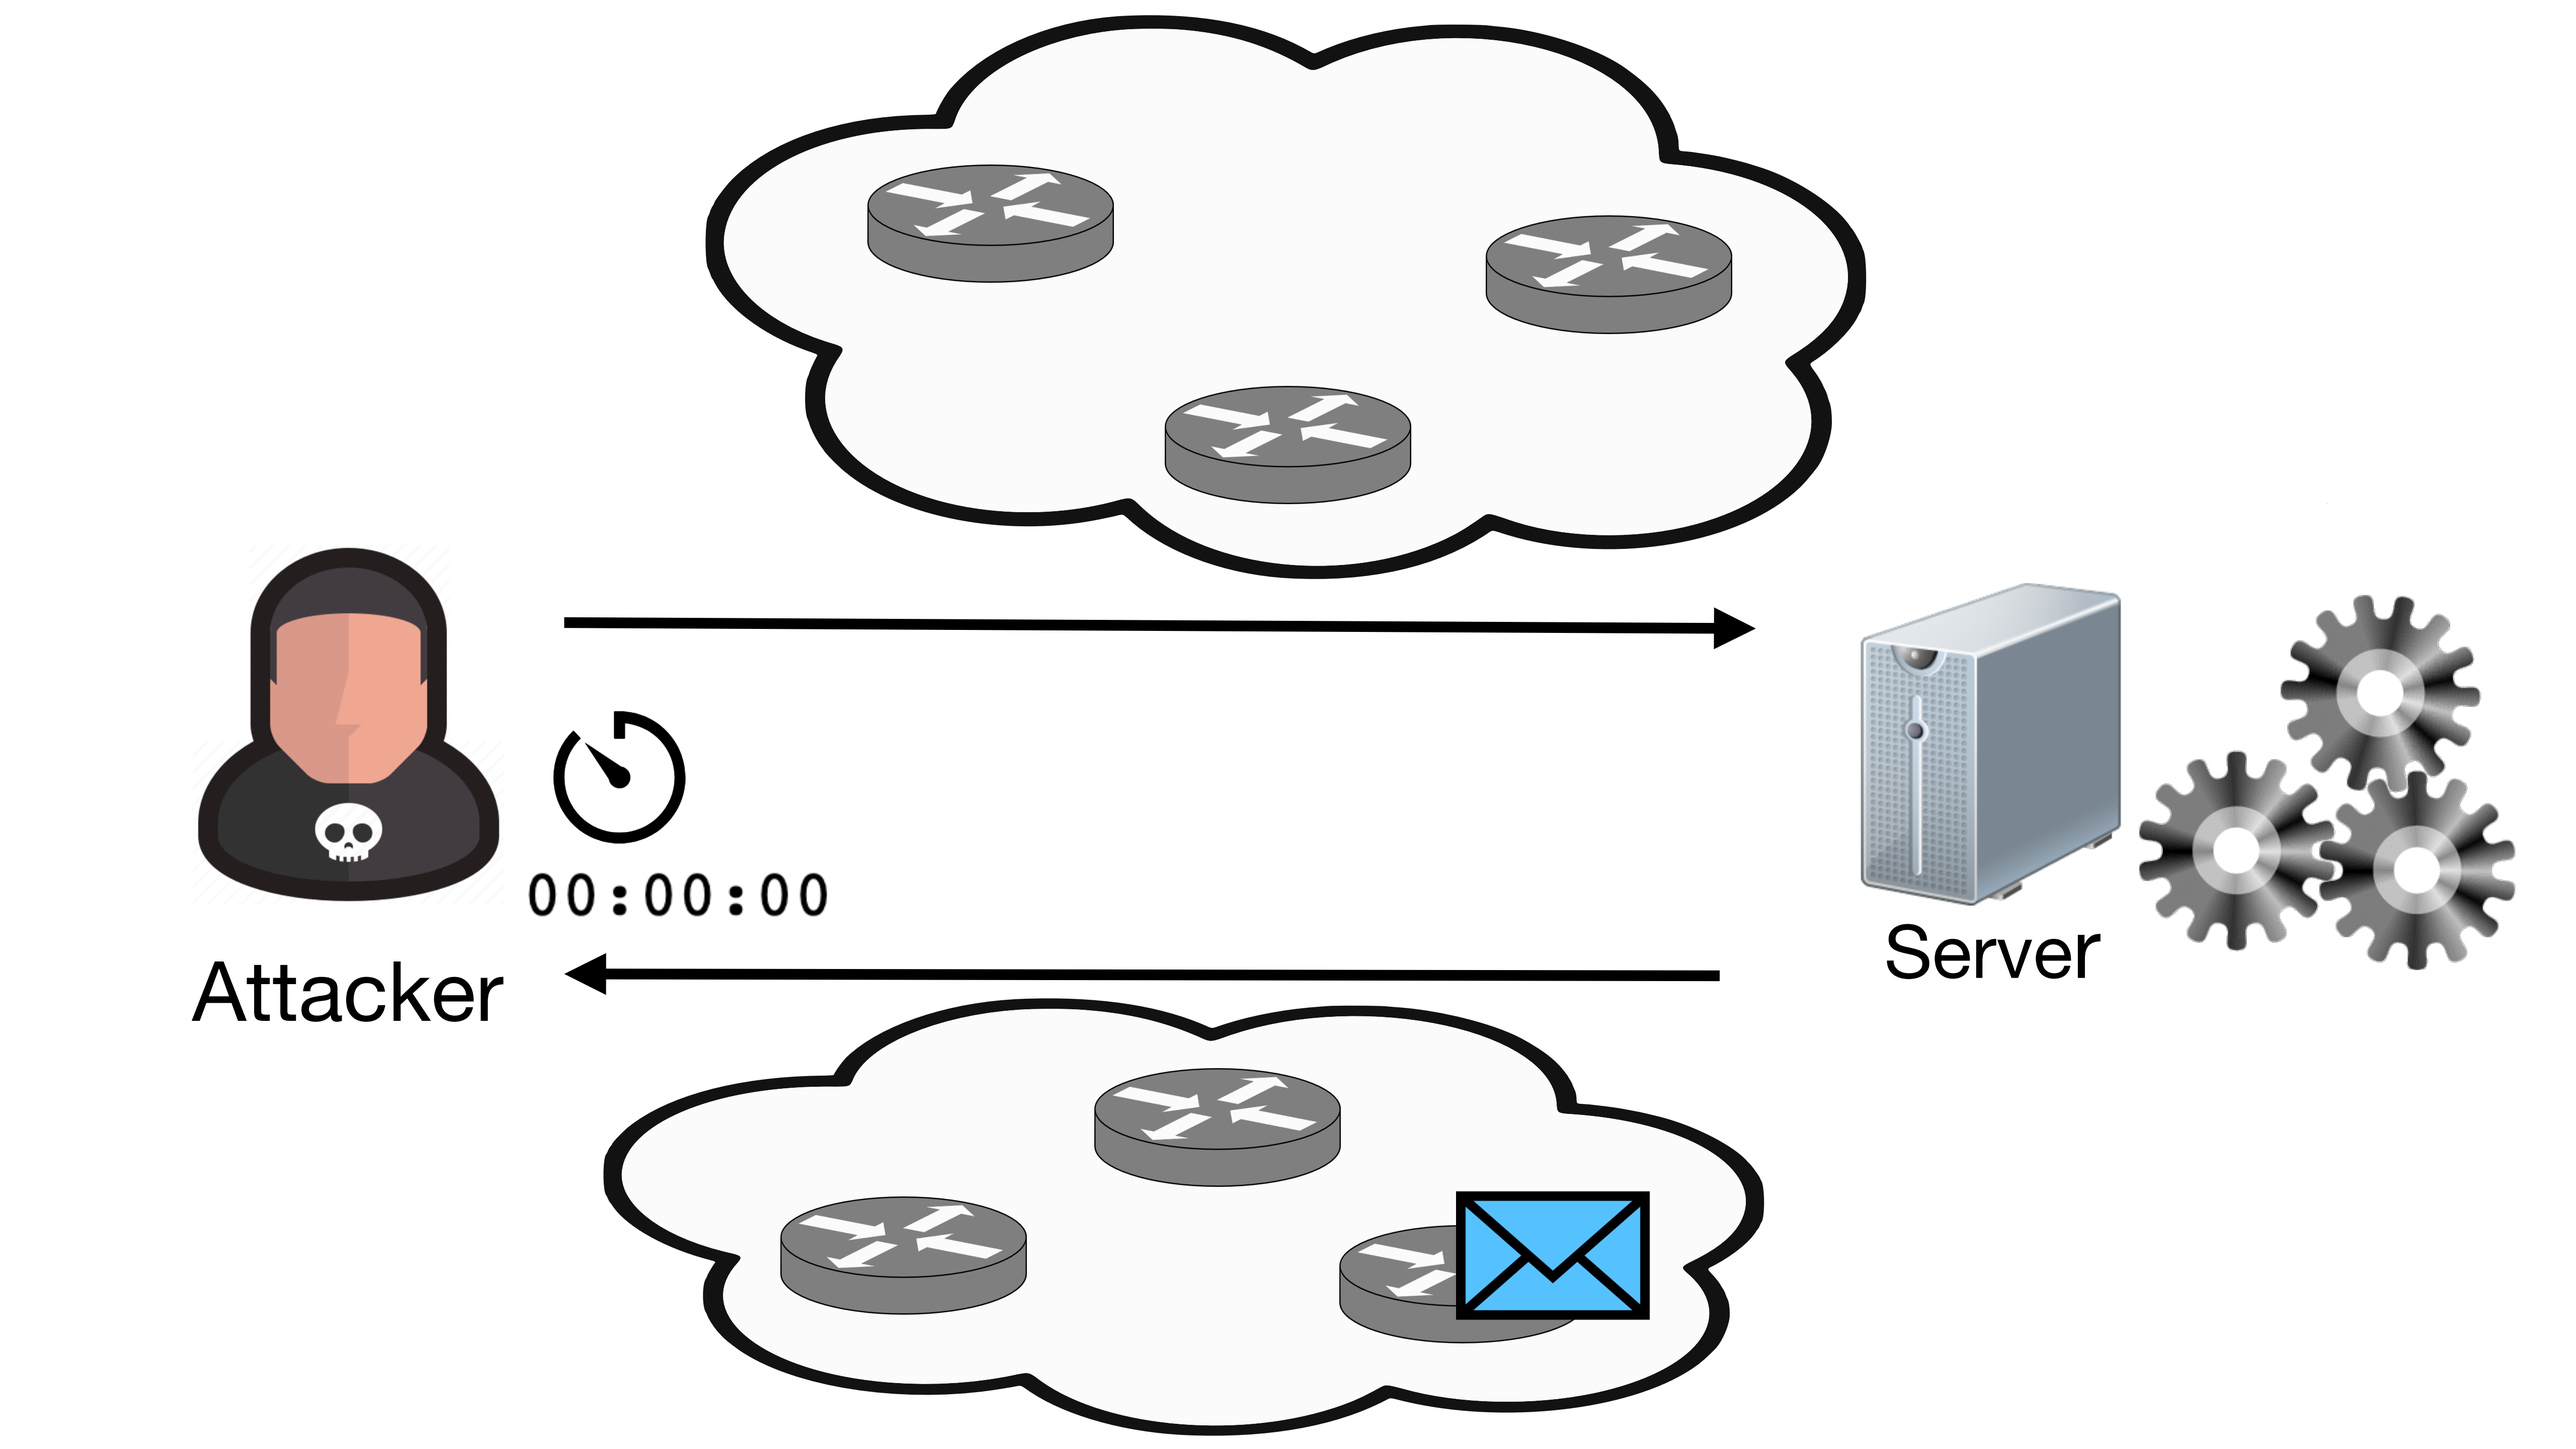
00:00:00
Attacker
Server

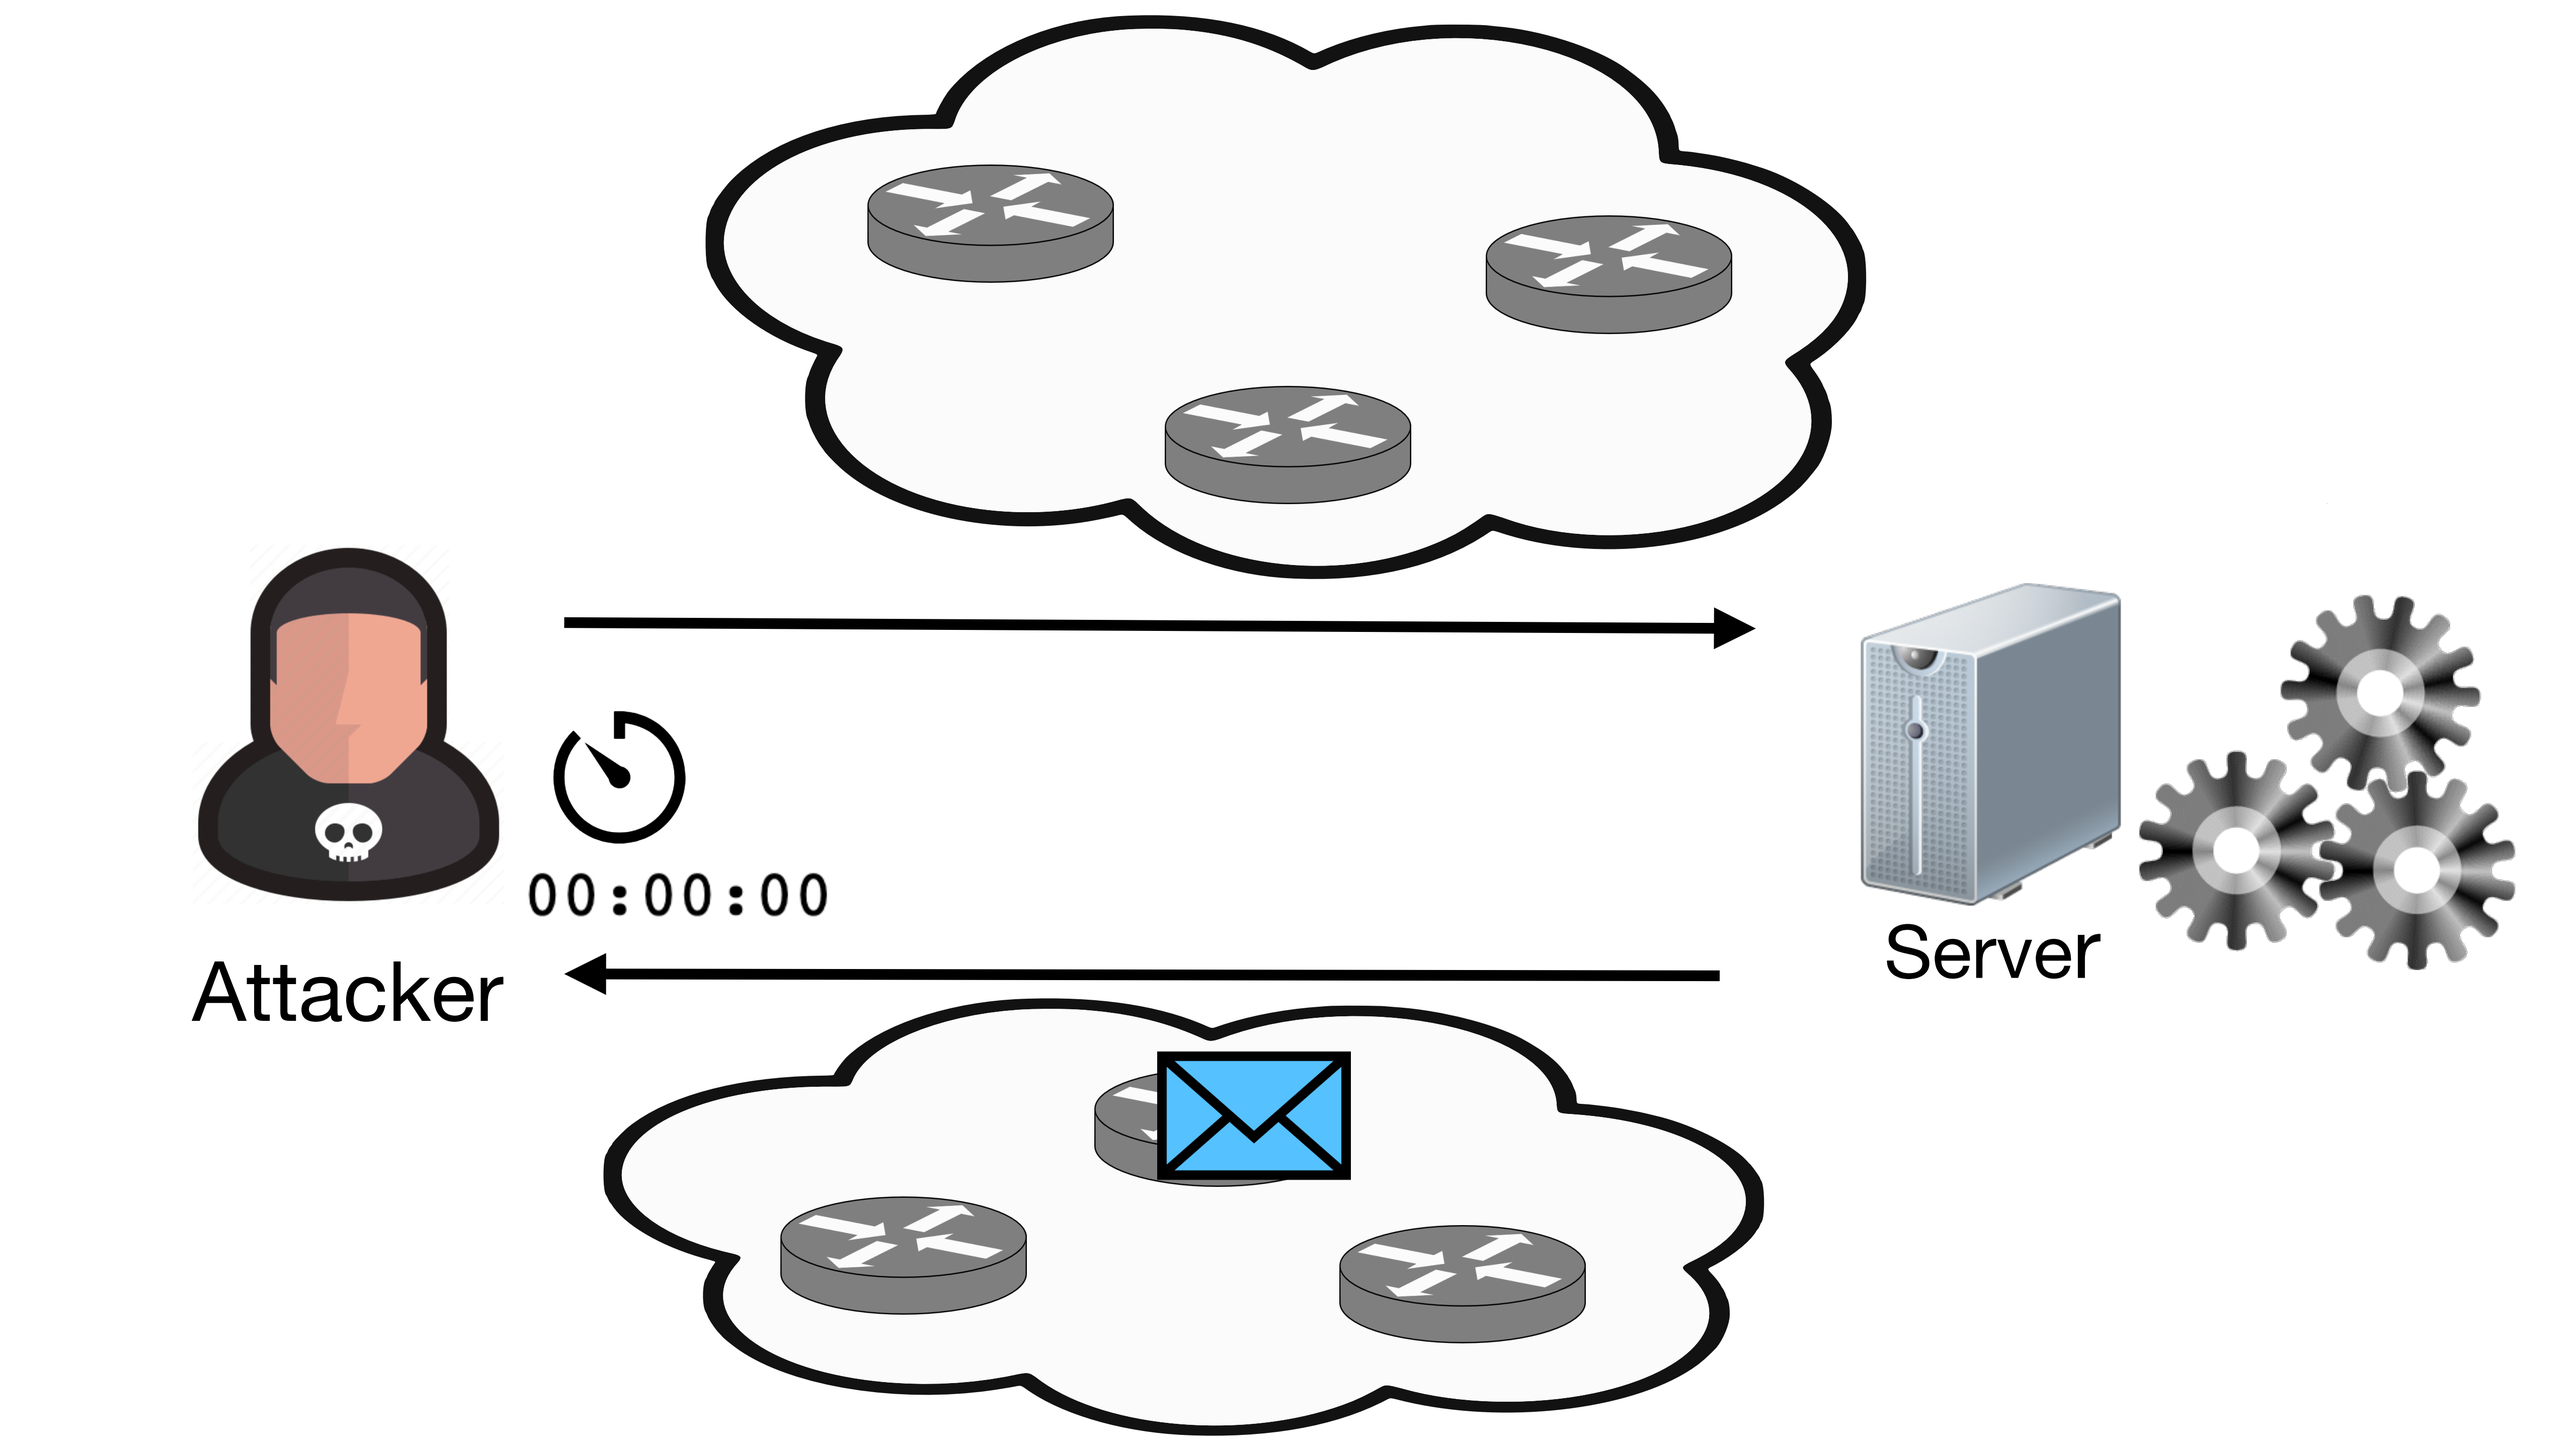
00:00:00
Attacker
Server

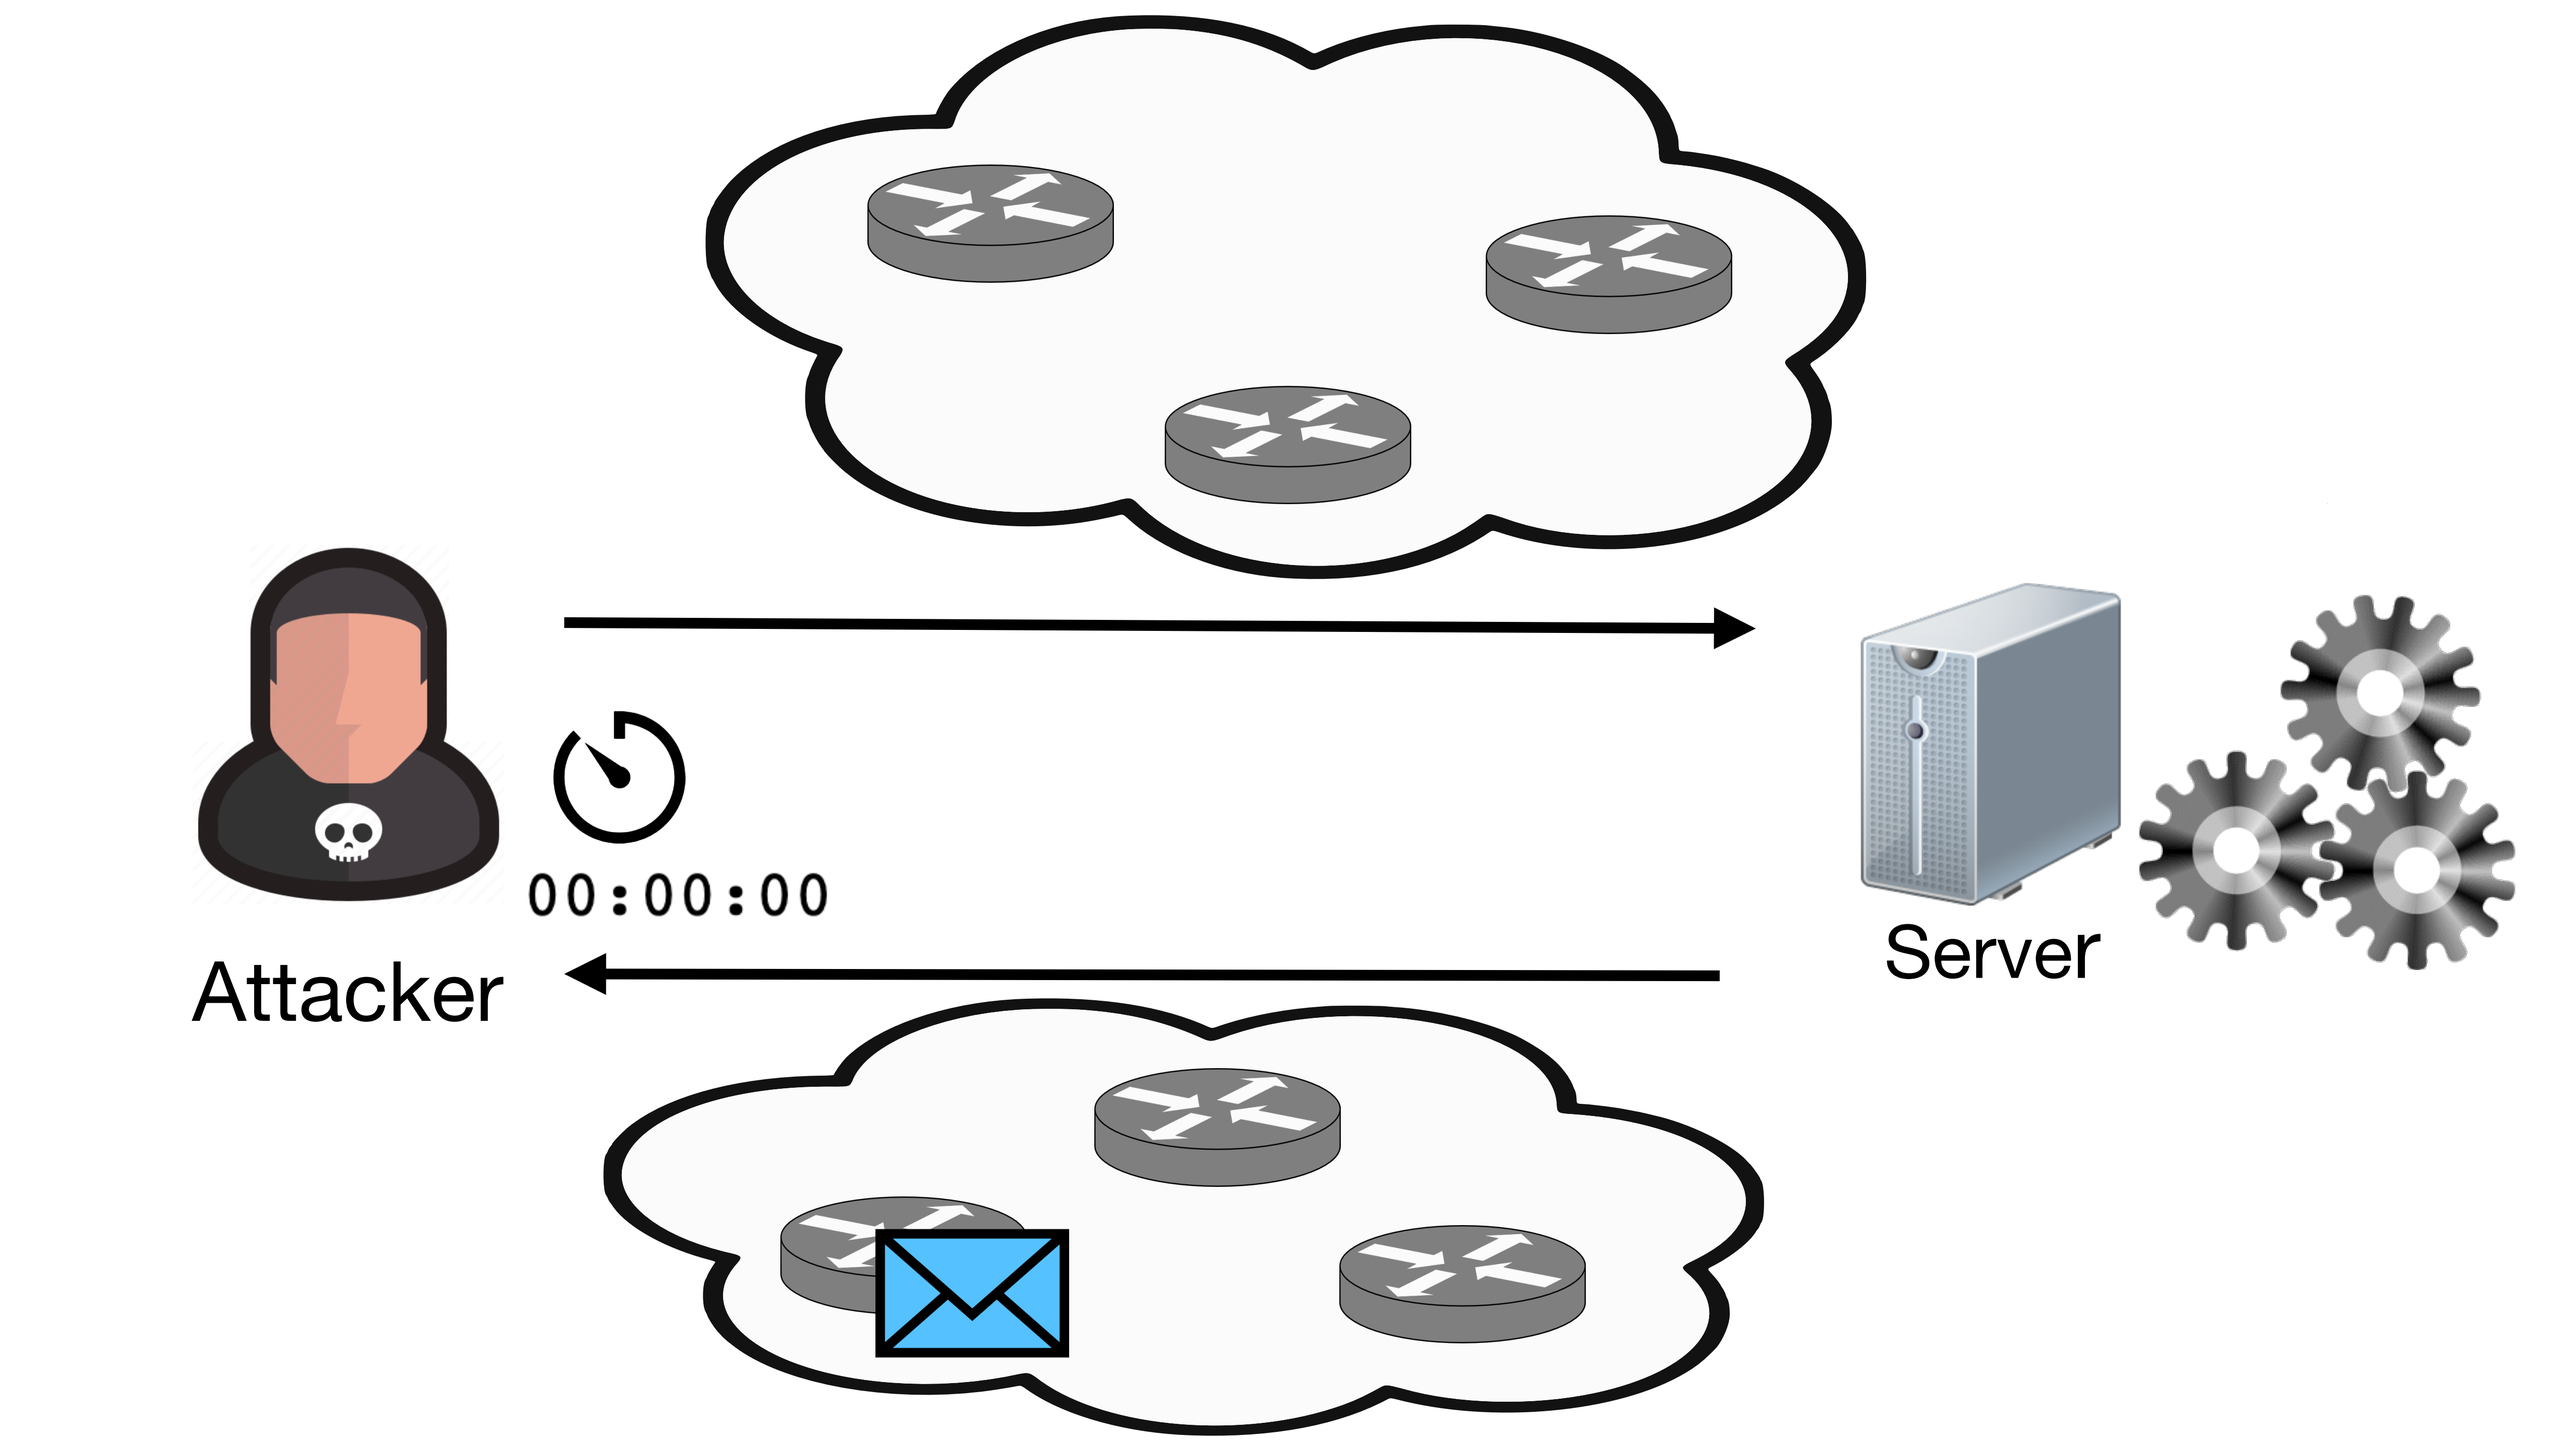
00:00:00
Attacker
Server

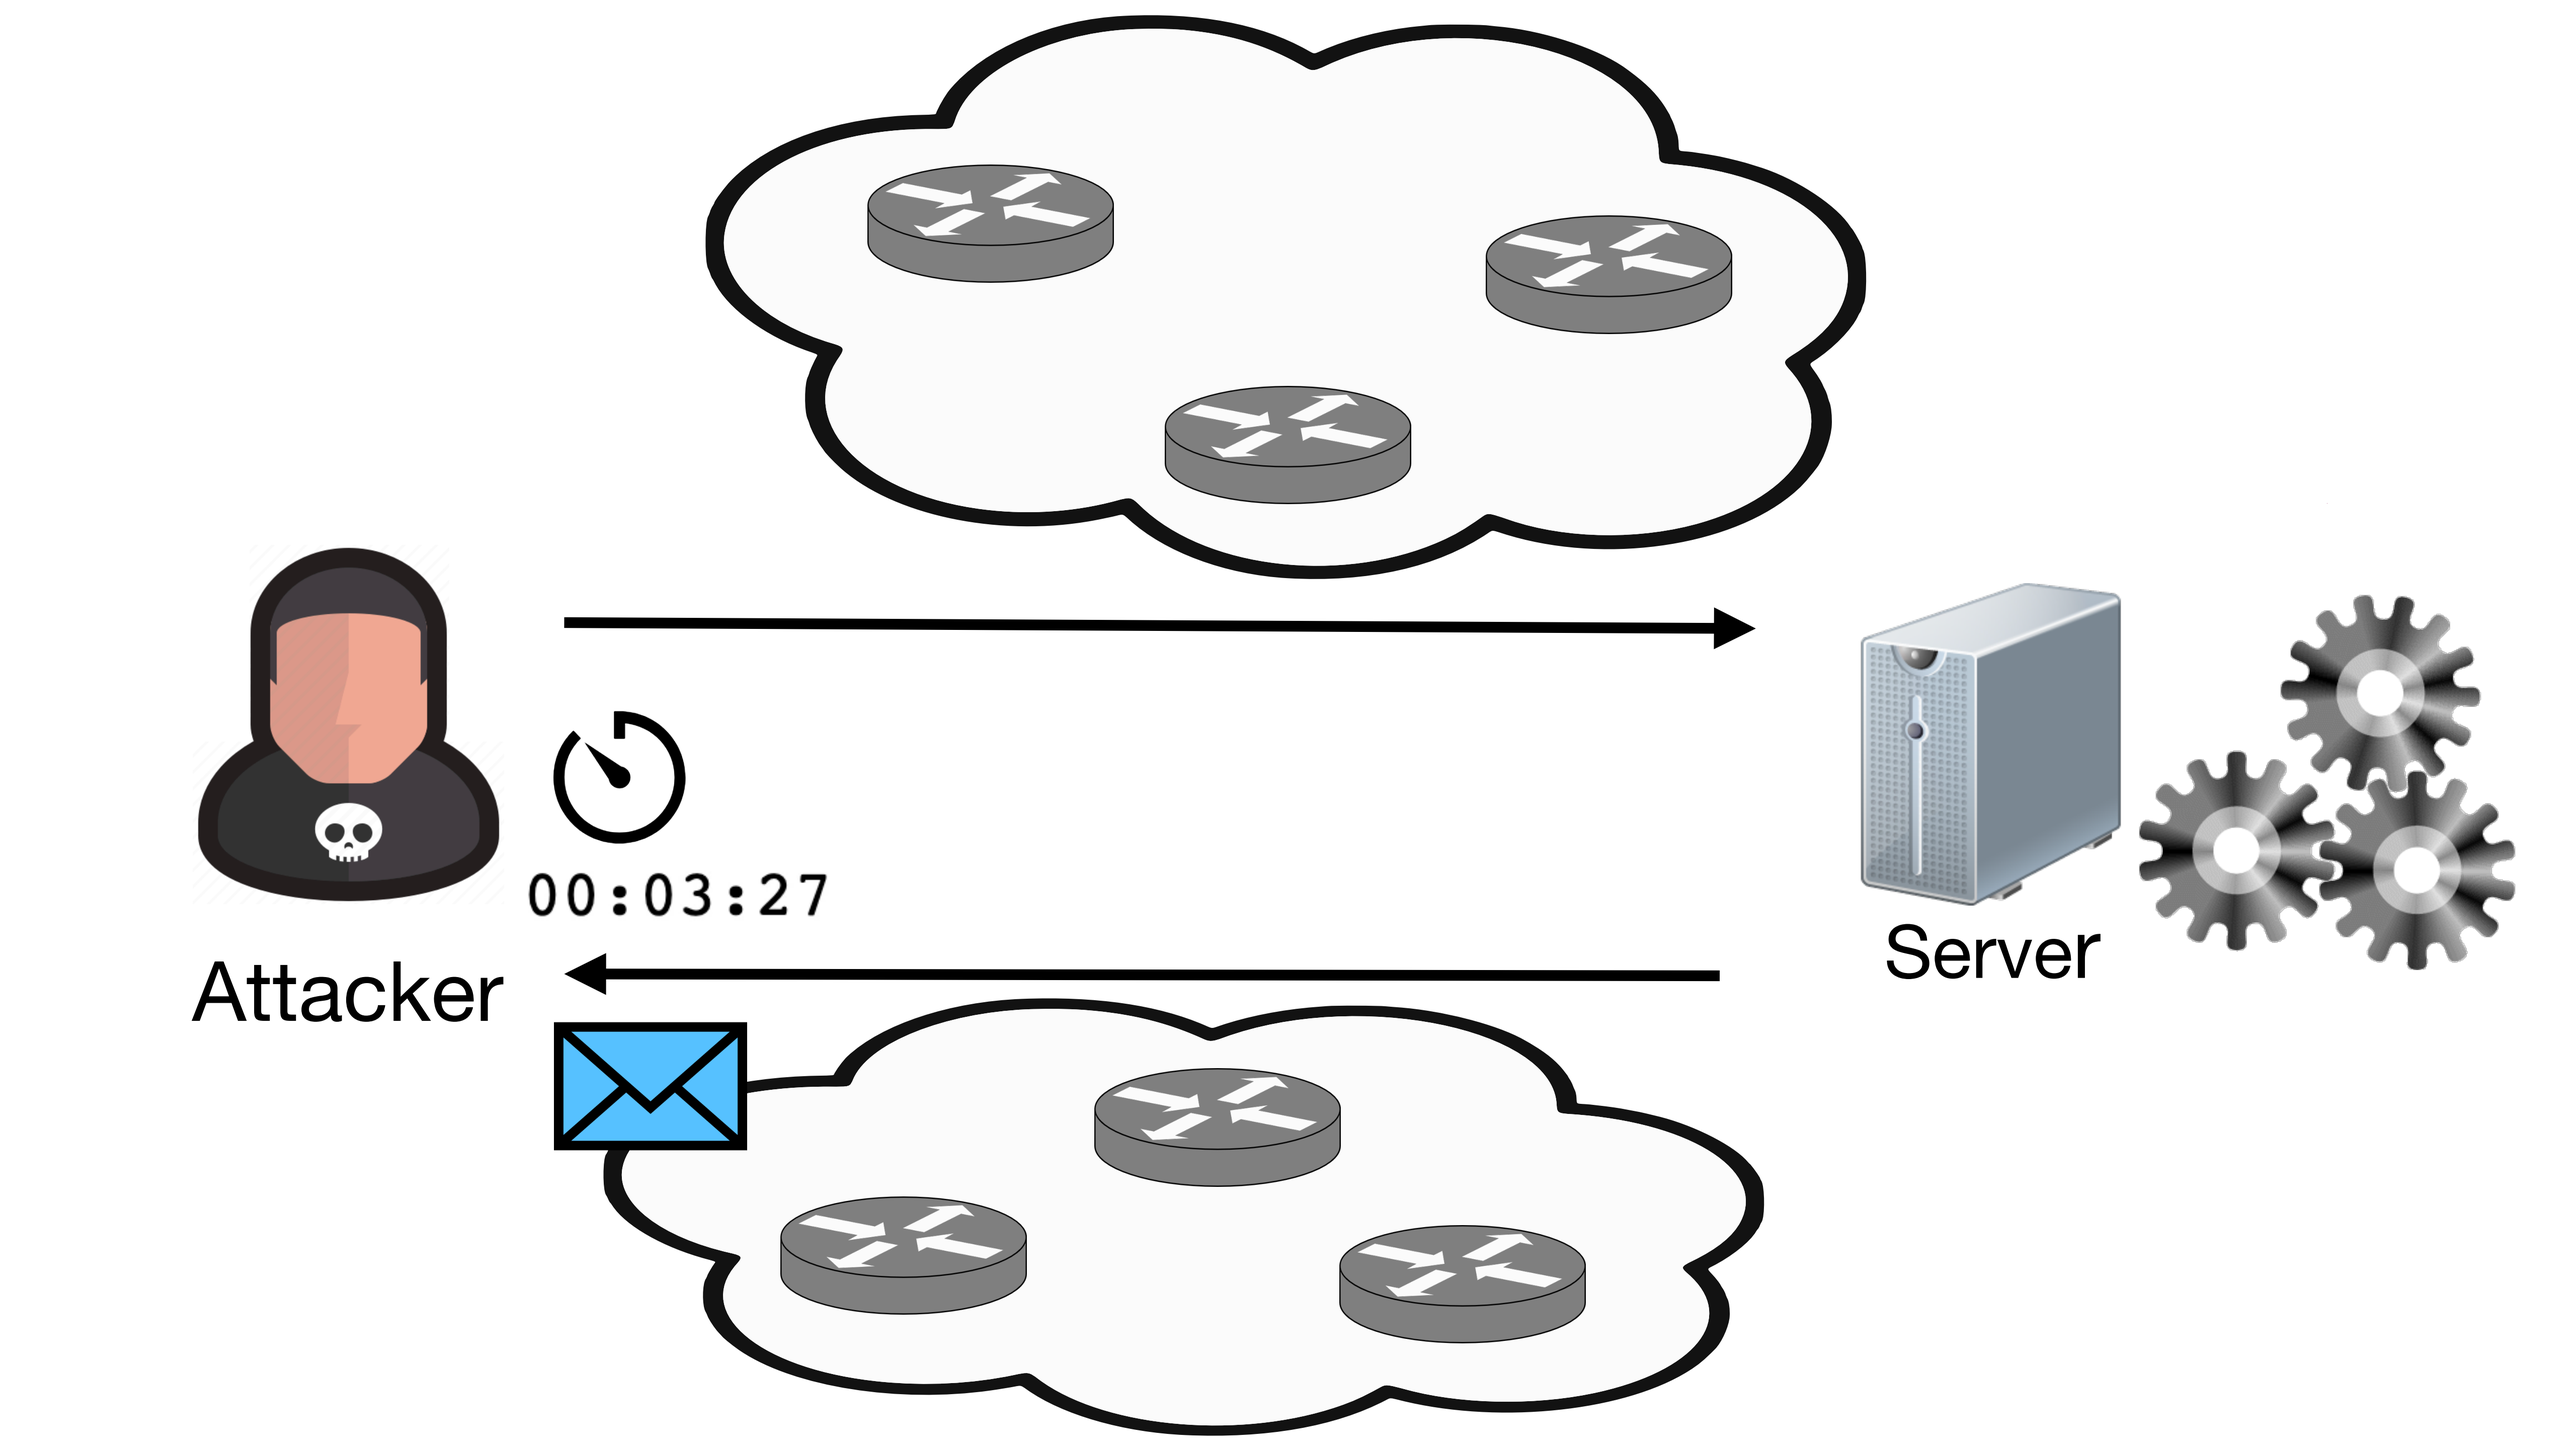
00:03:27
Attacker
Server

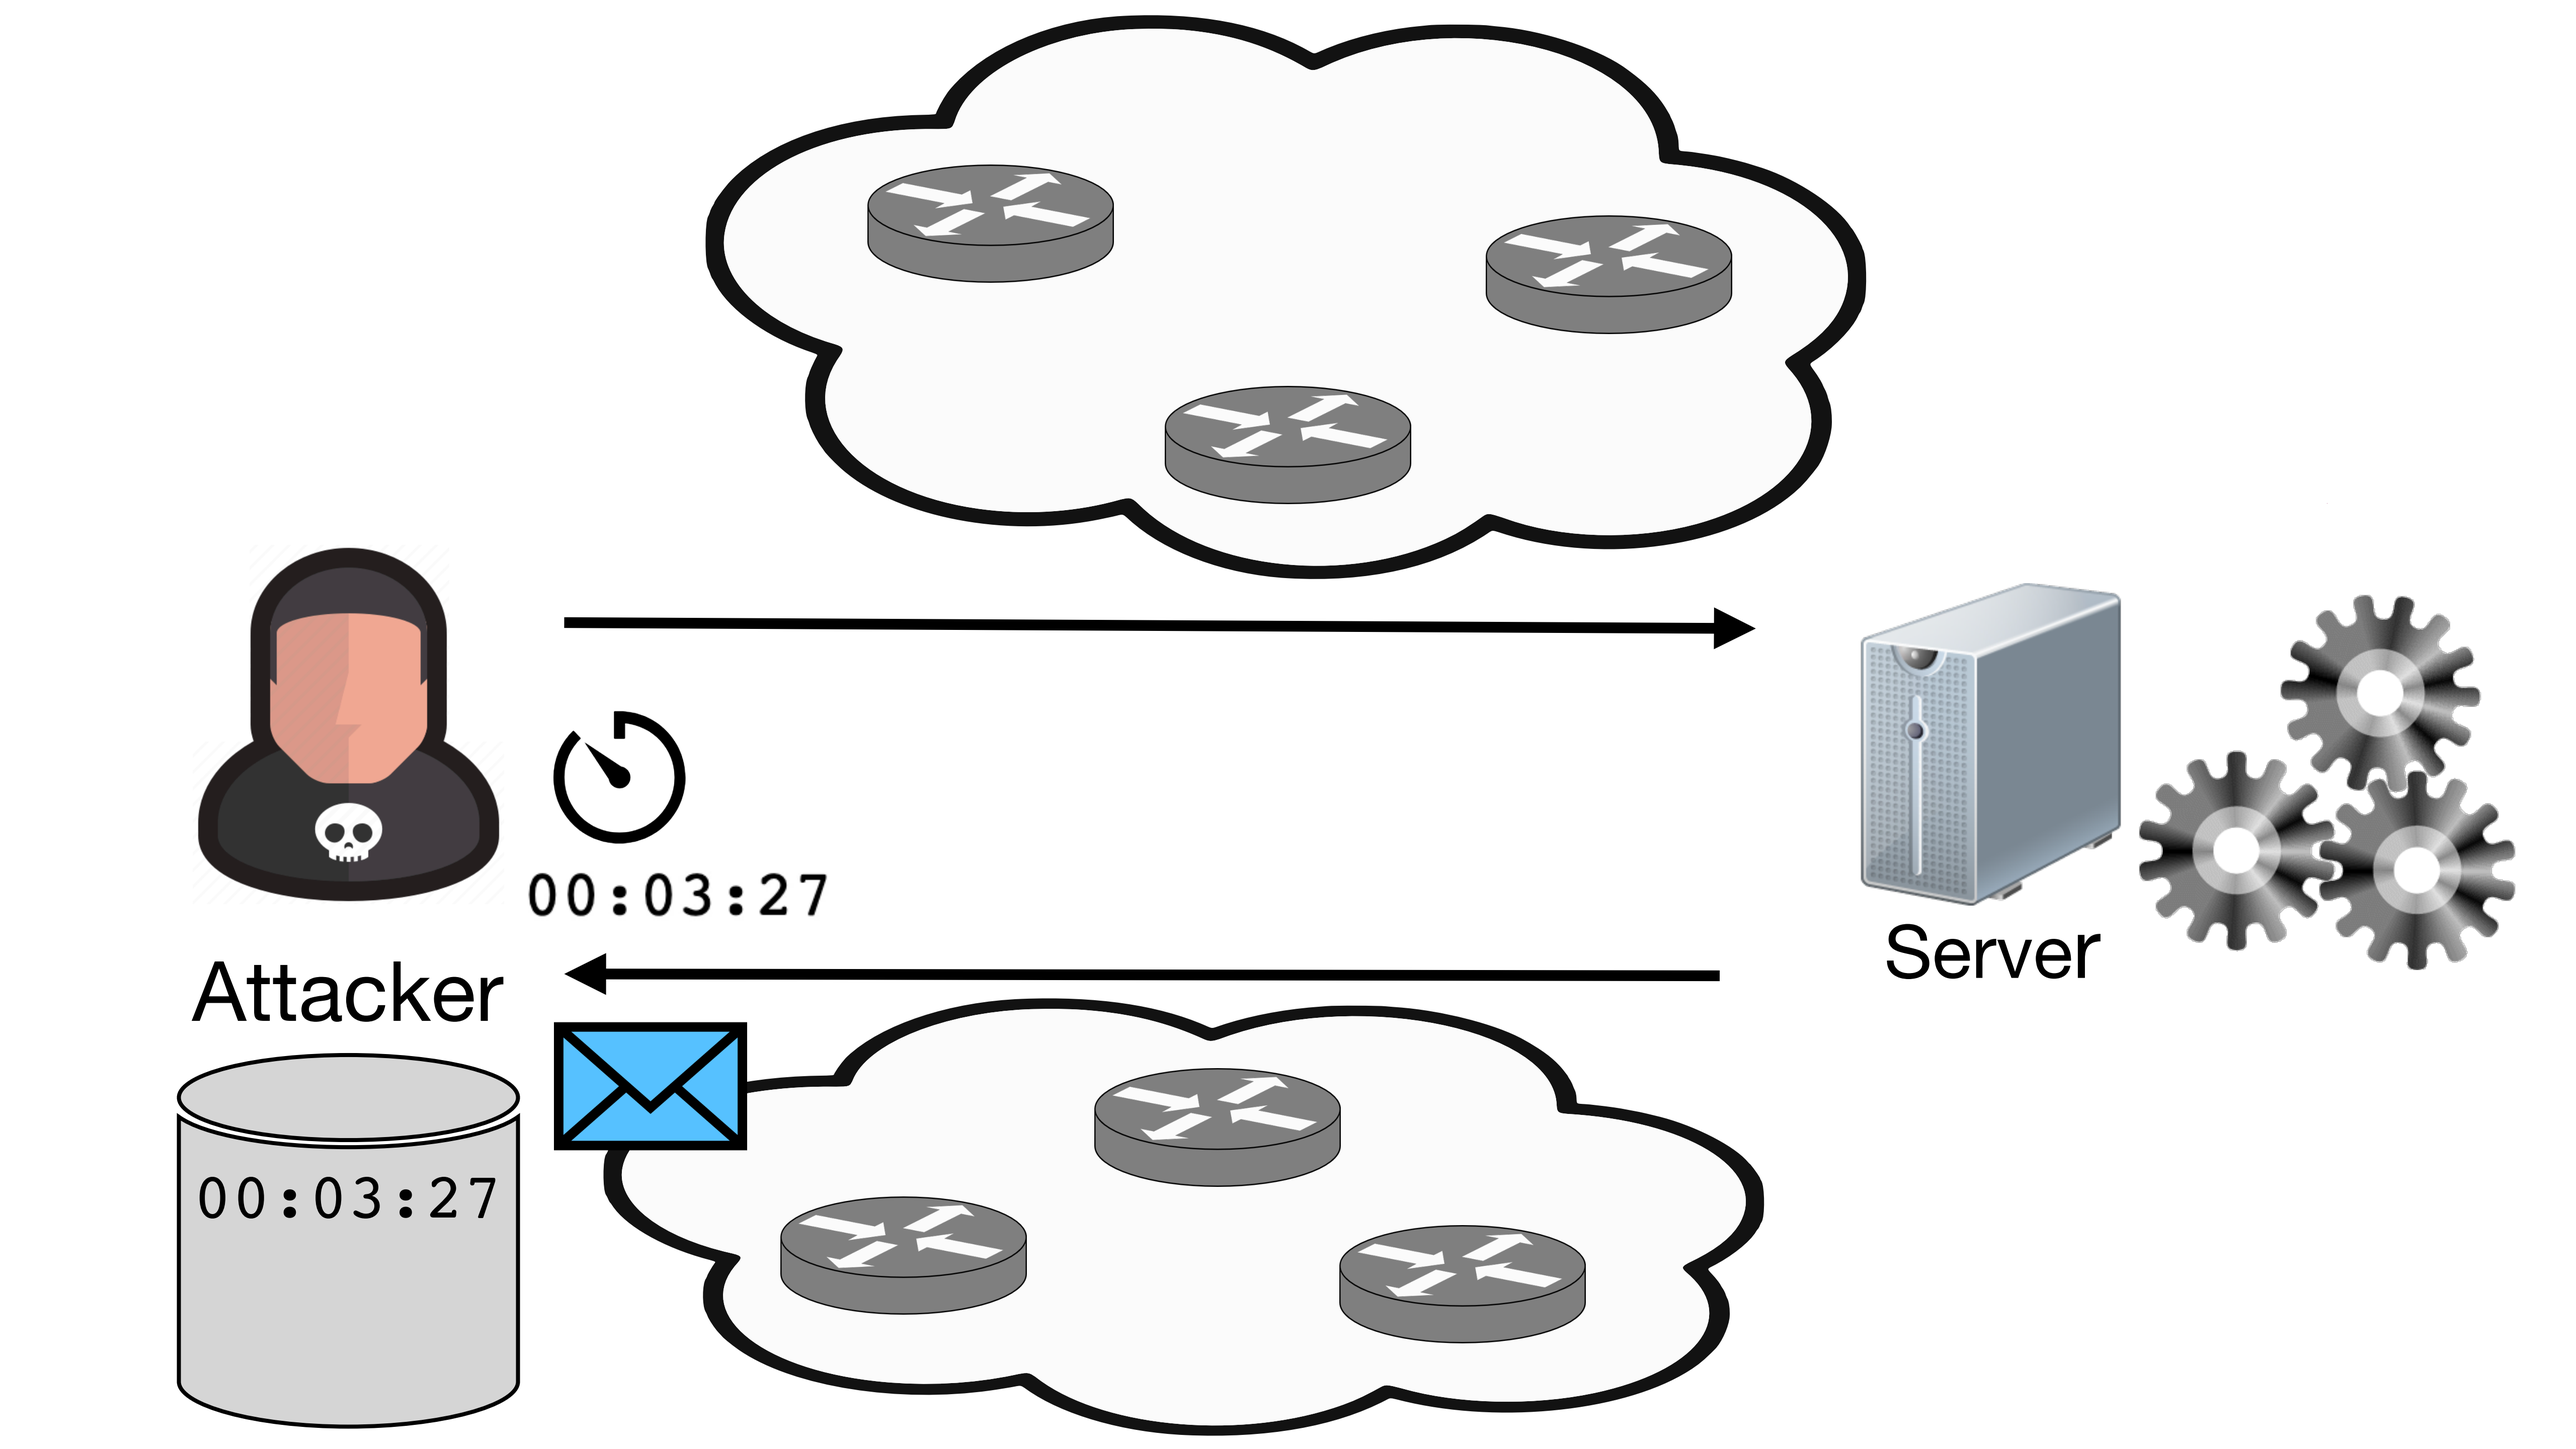
00:03:27
Server
Attacker
00:03:27

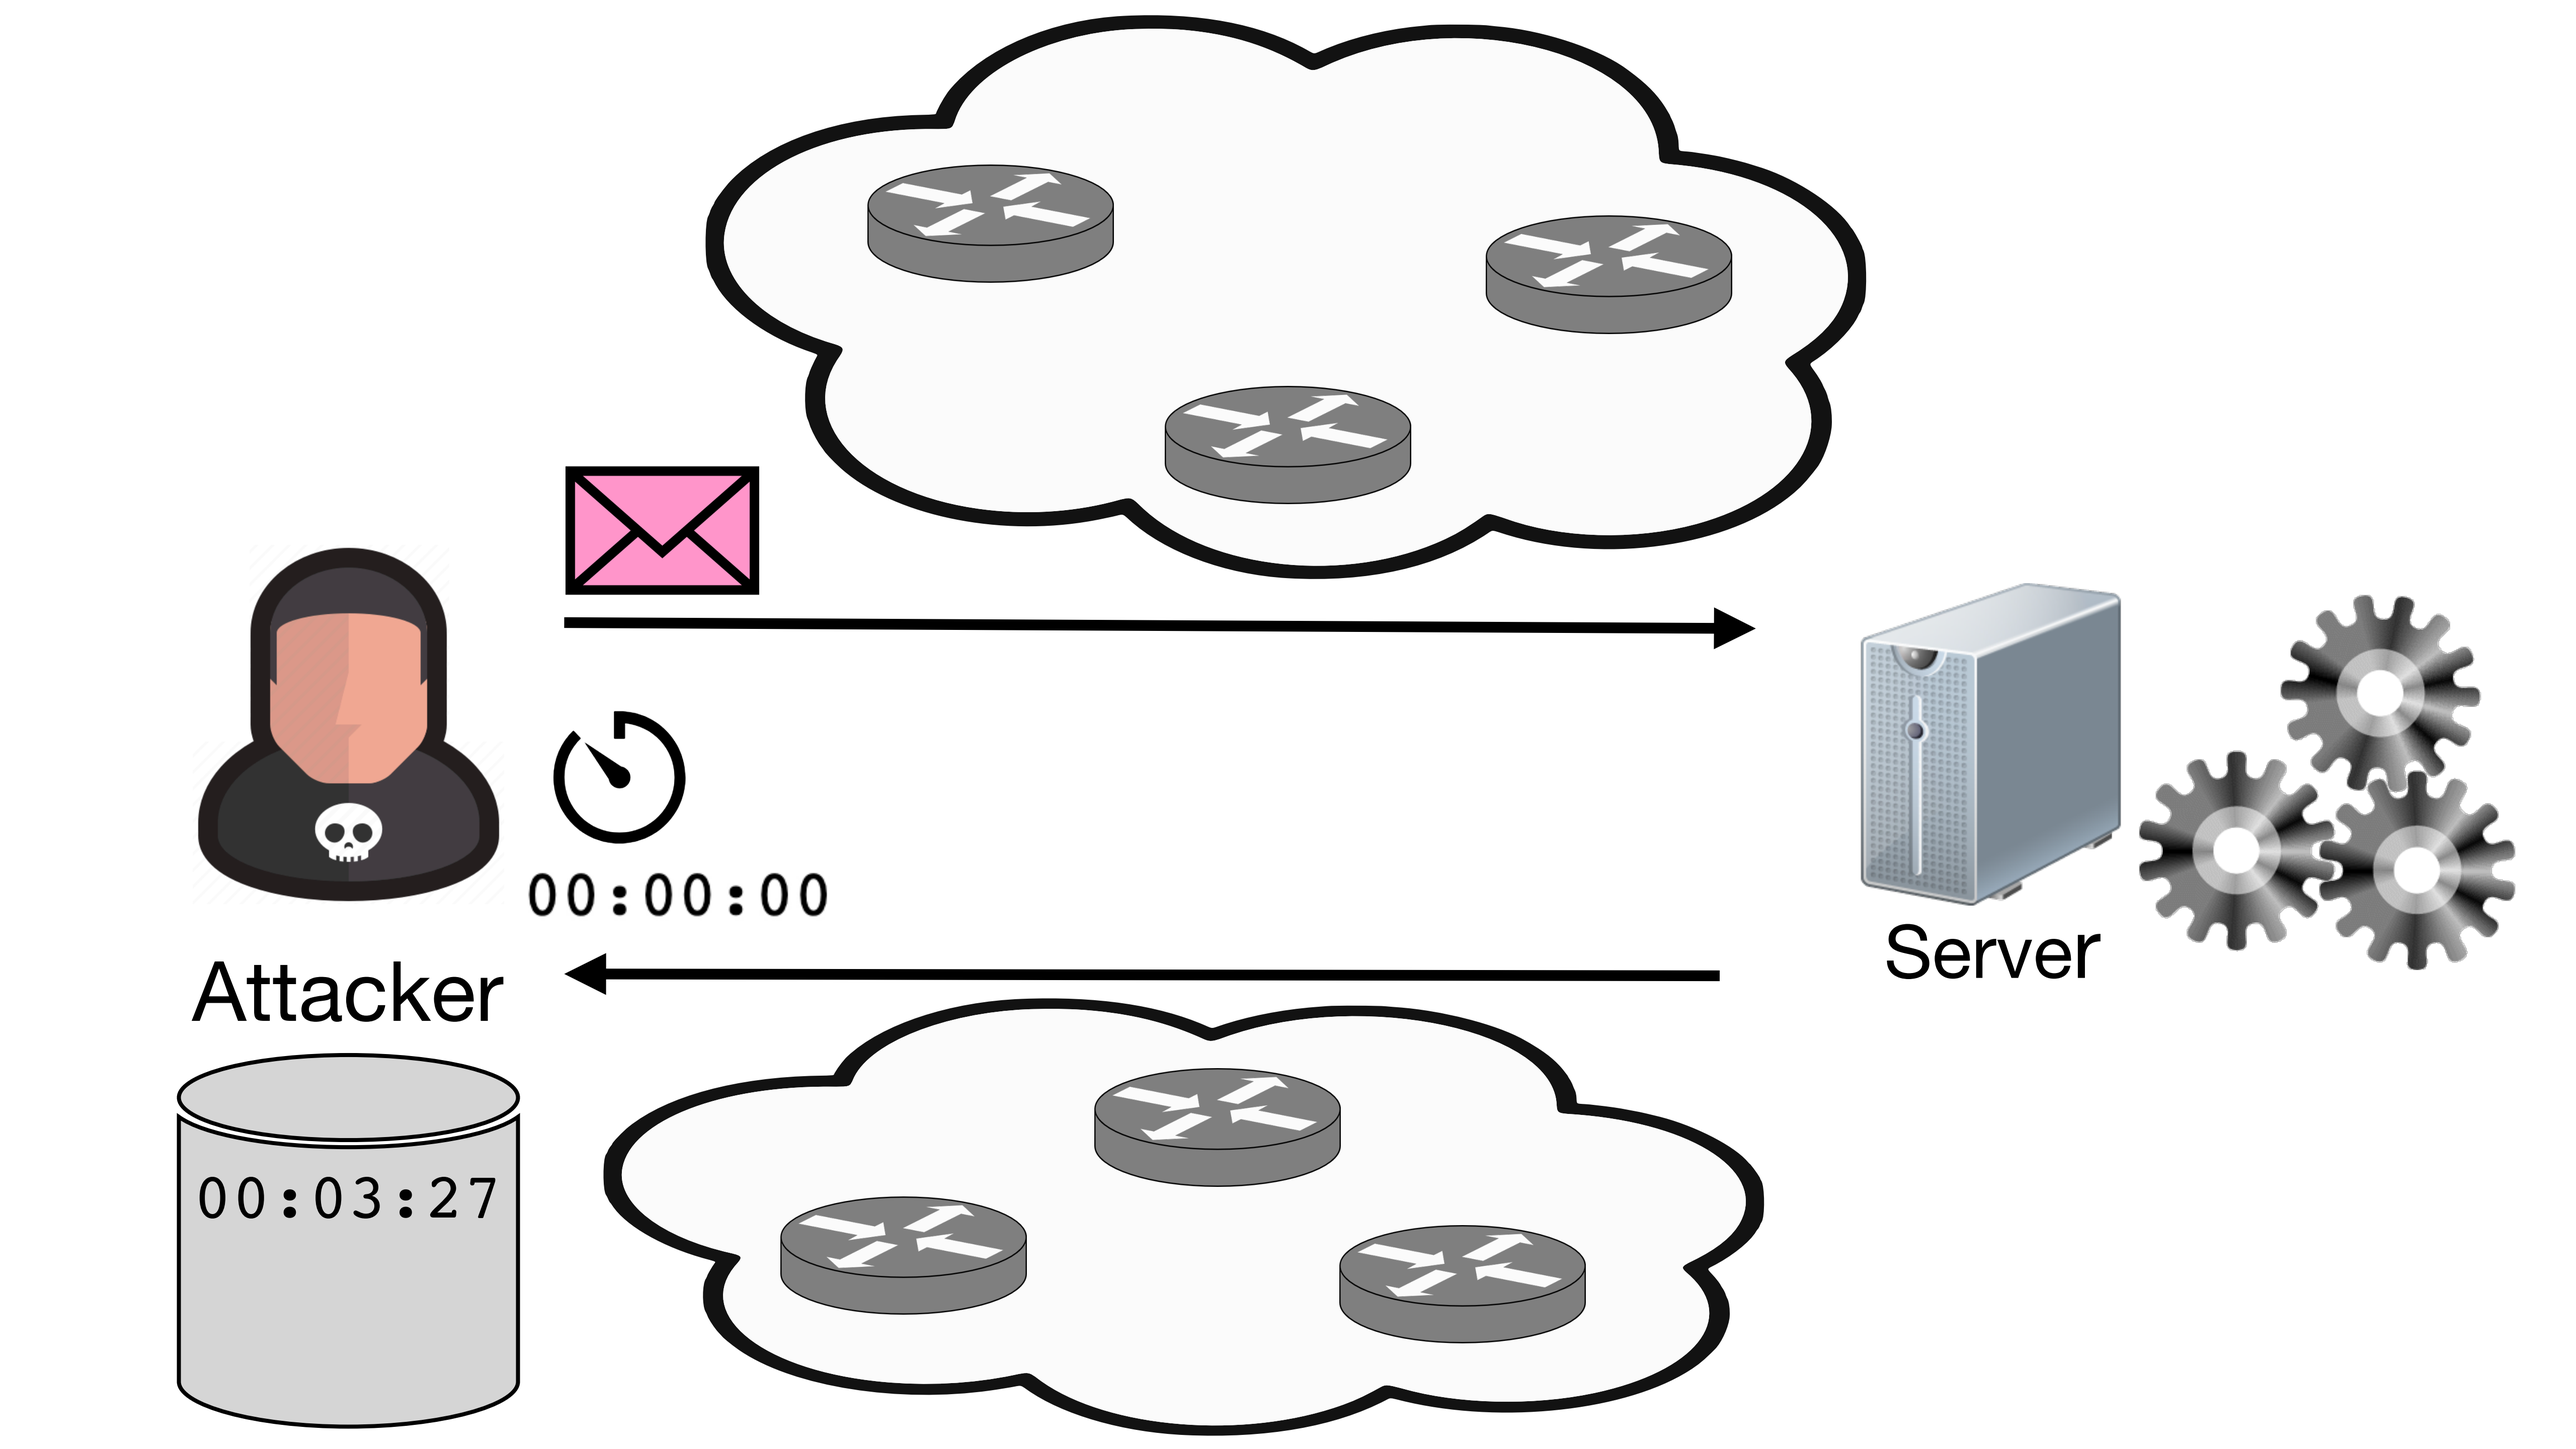
00:00:00
Attacker
00:03:27
Server

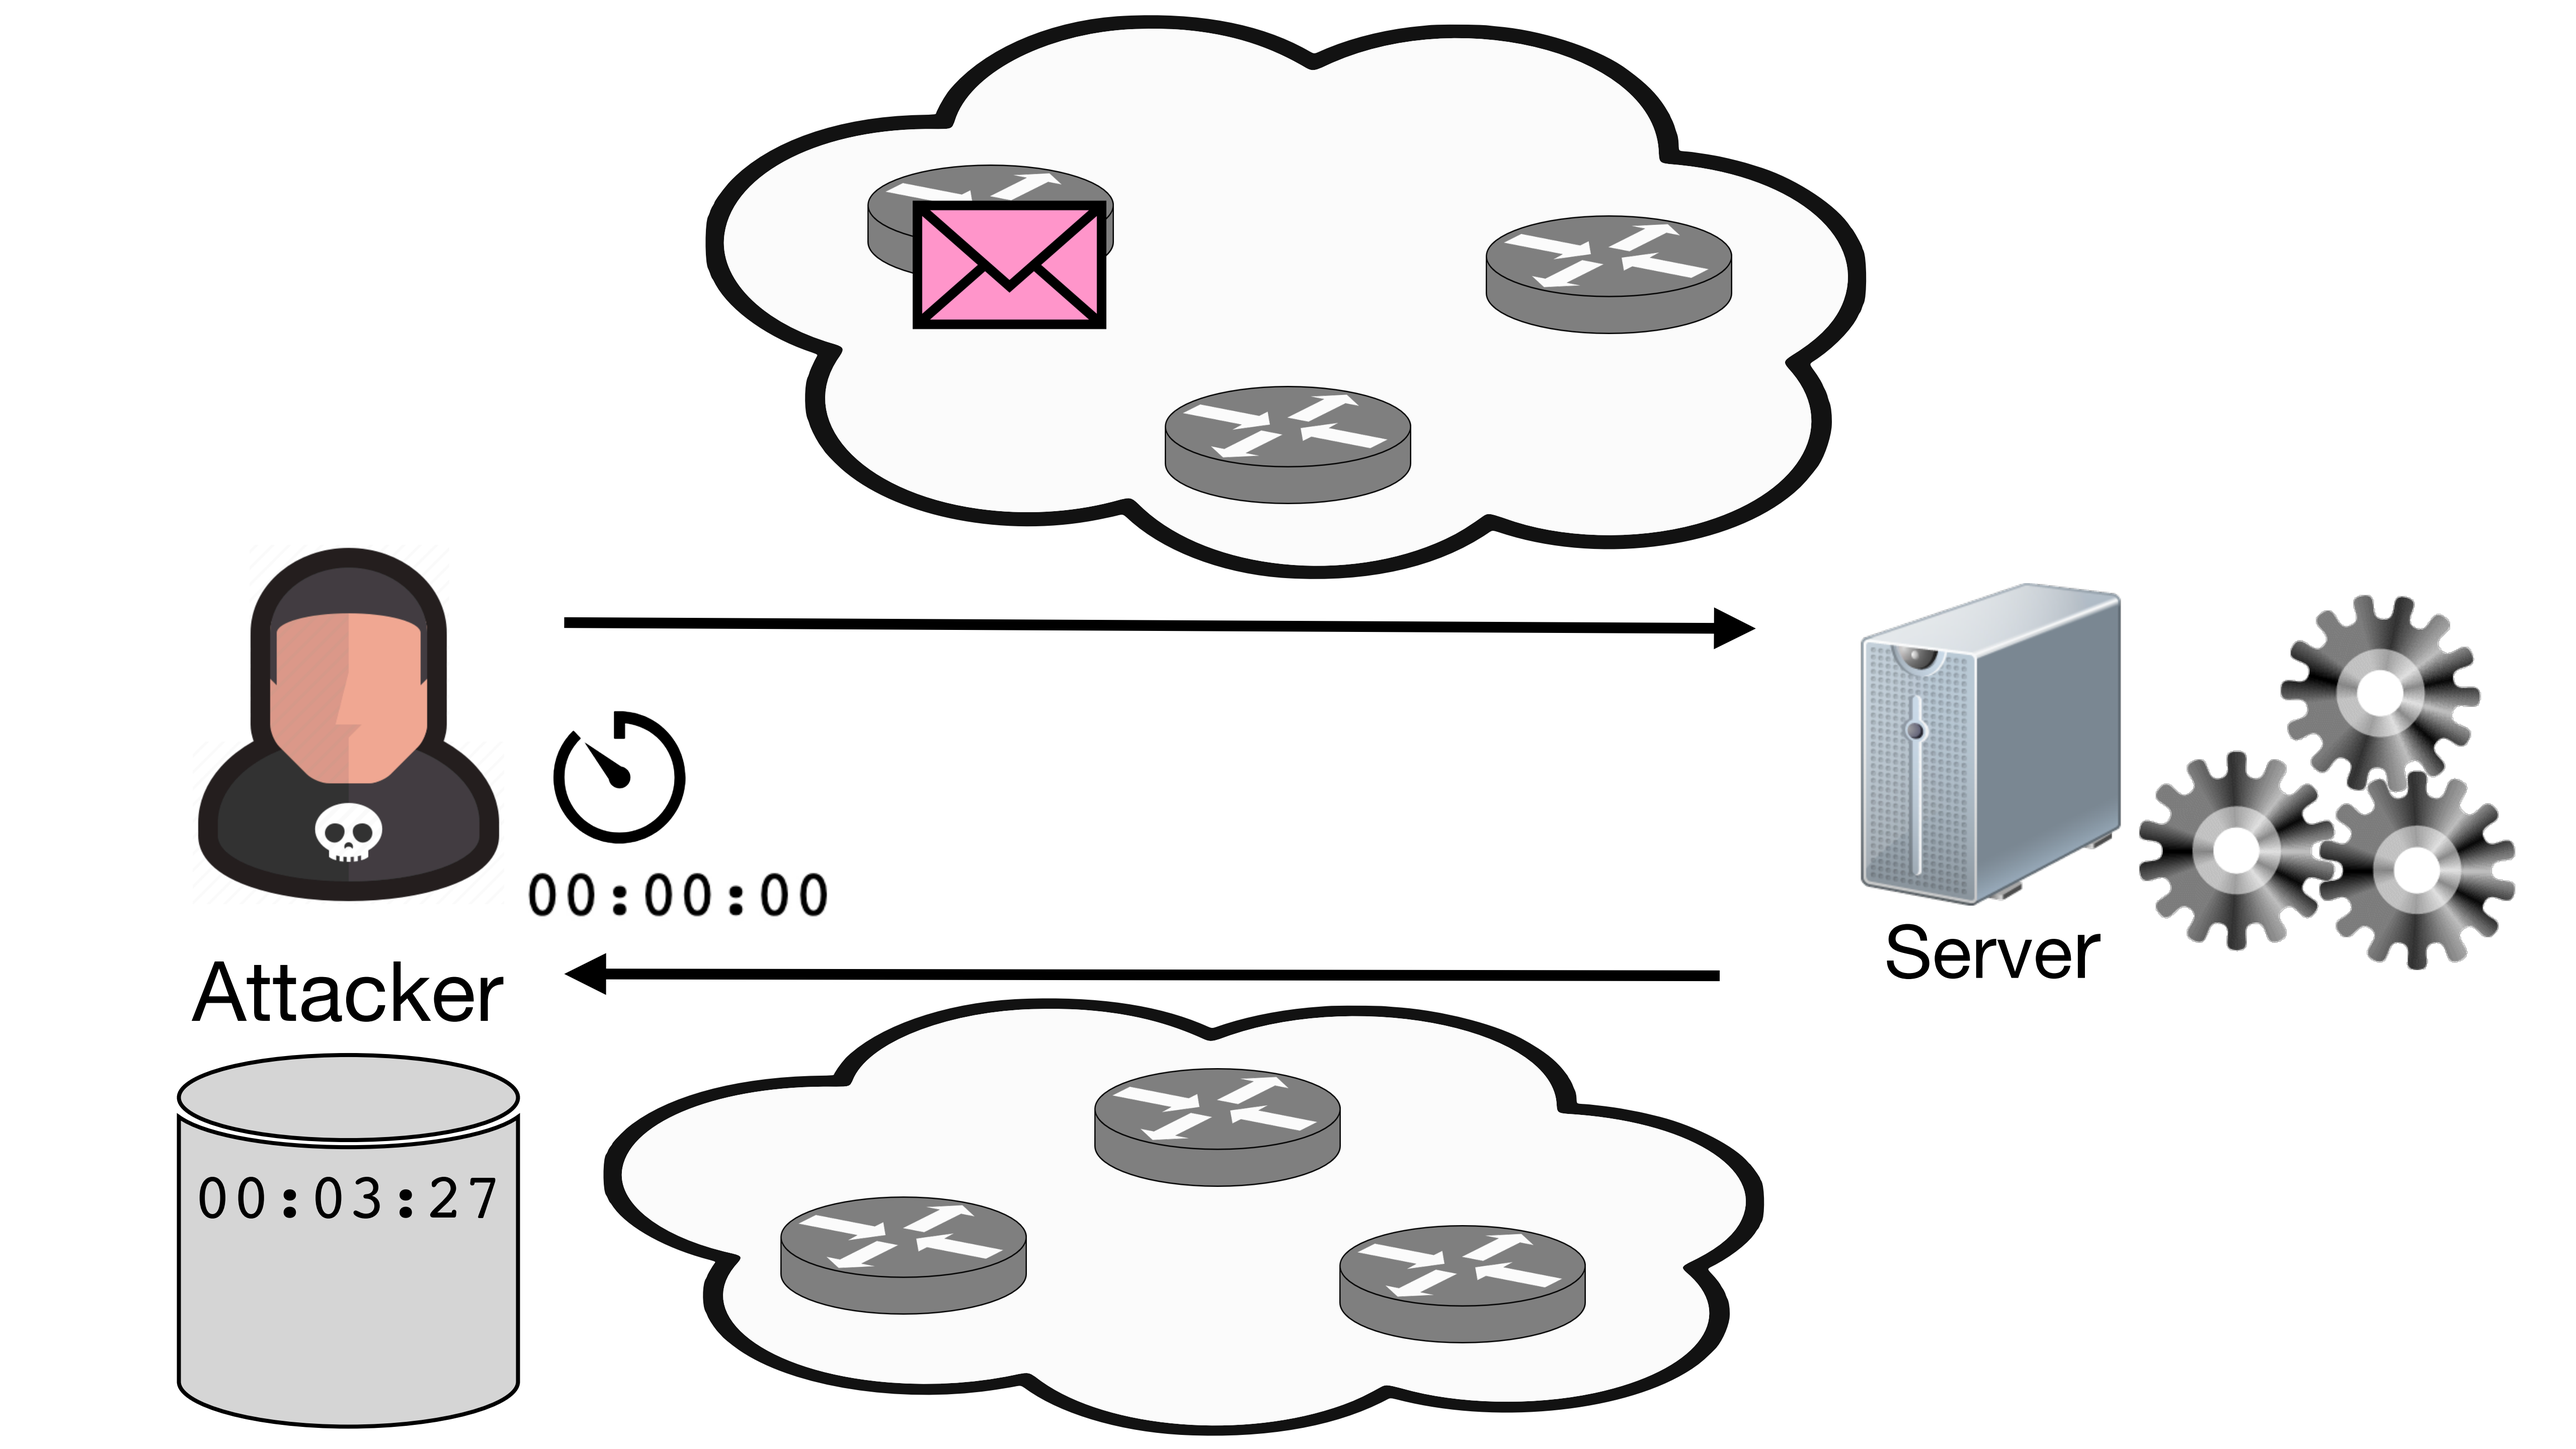
00:00:00
Server
Attacker
00:03:27

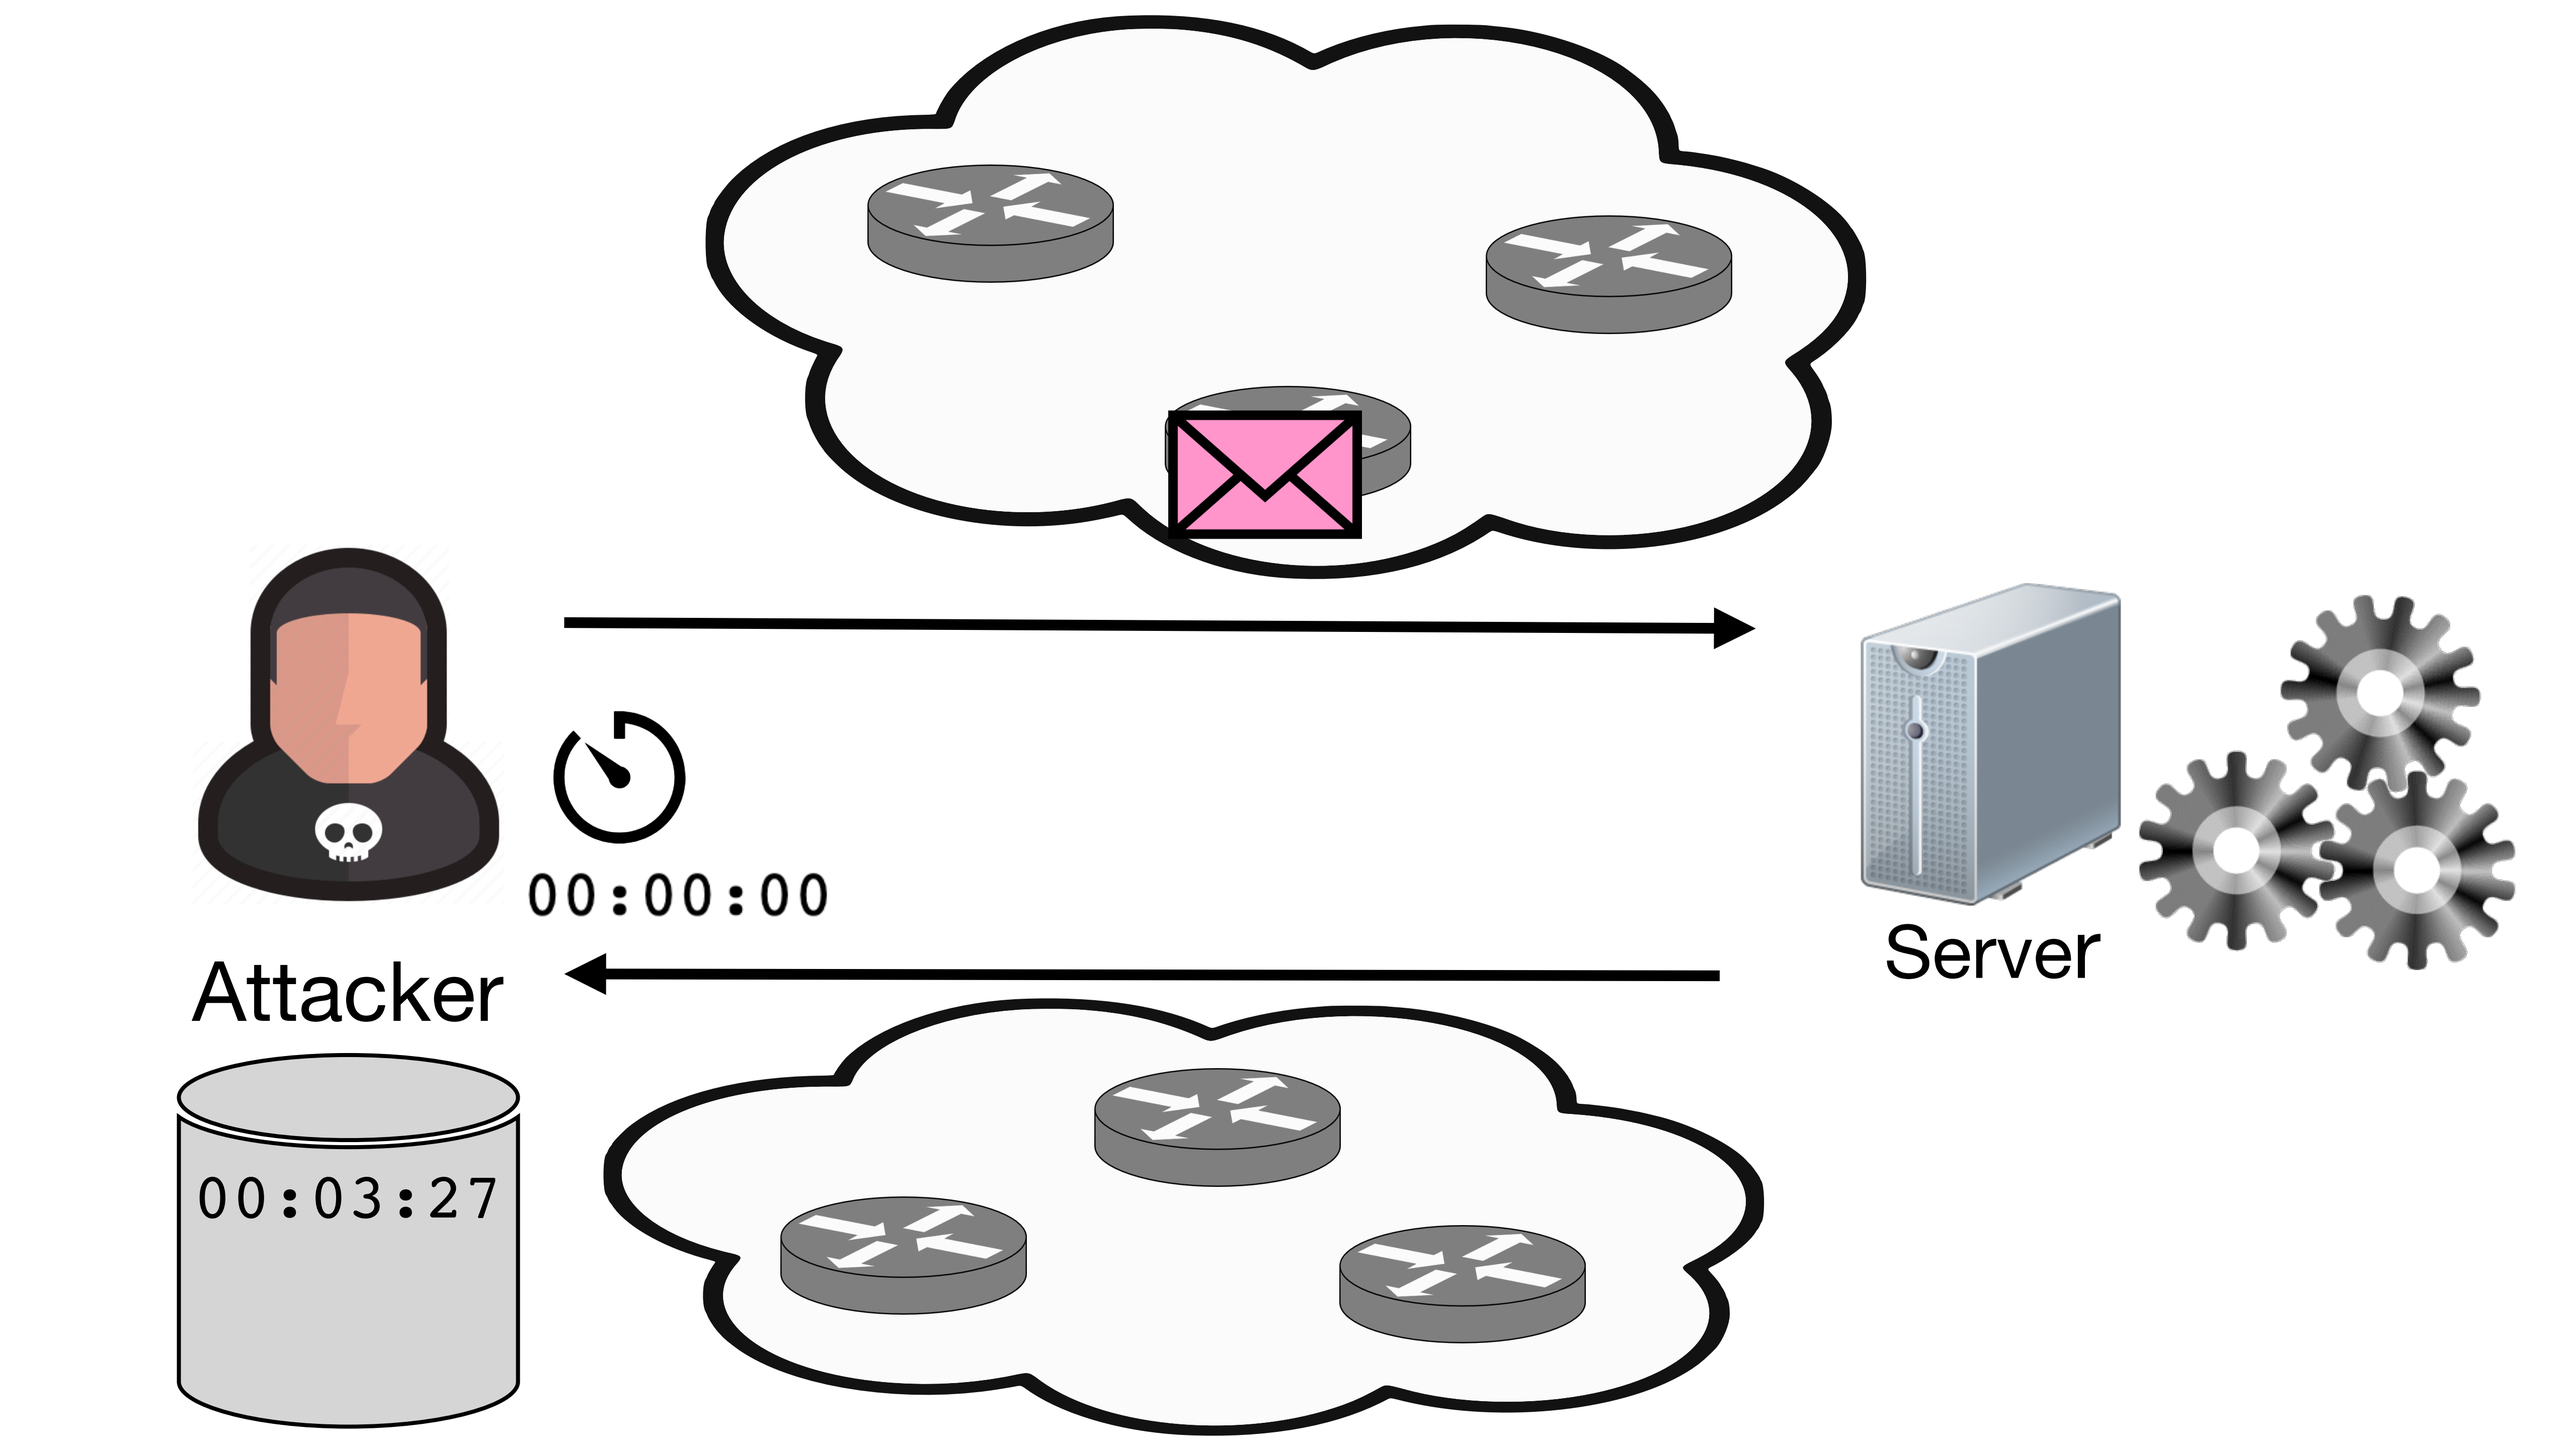
00:00:00
Attacker
Server
00:03:27

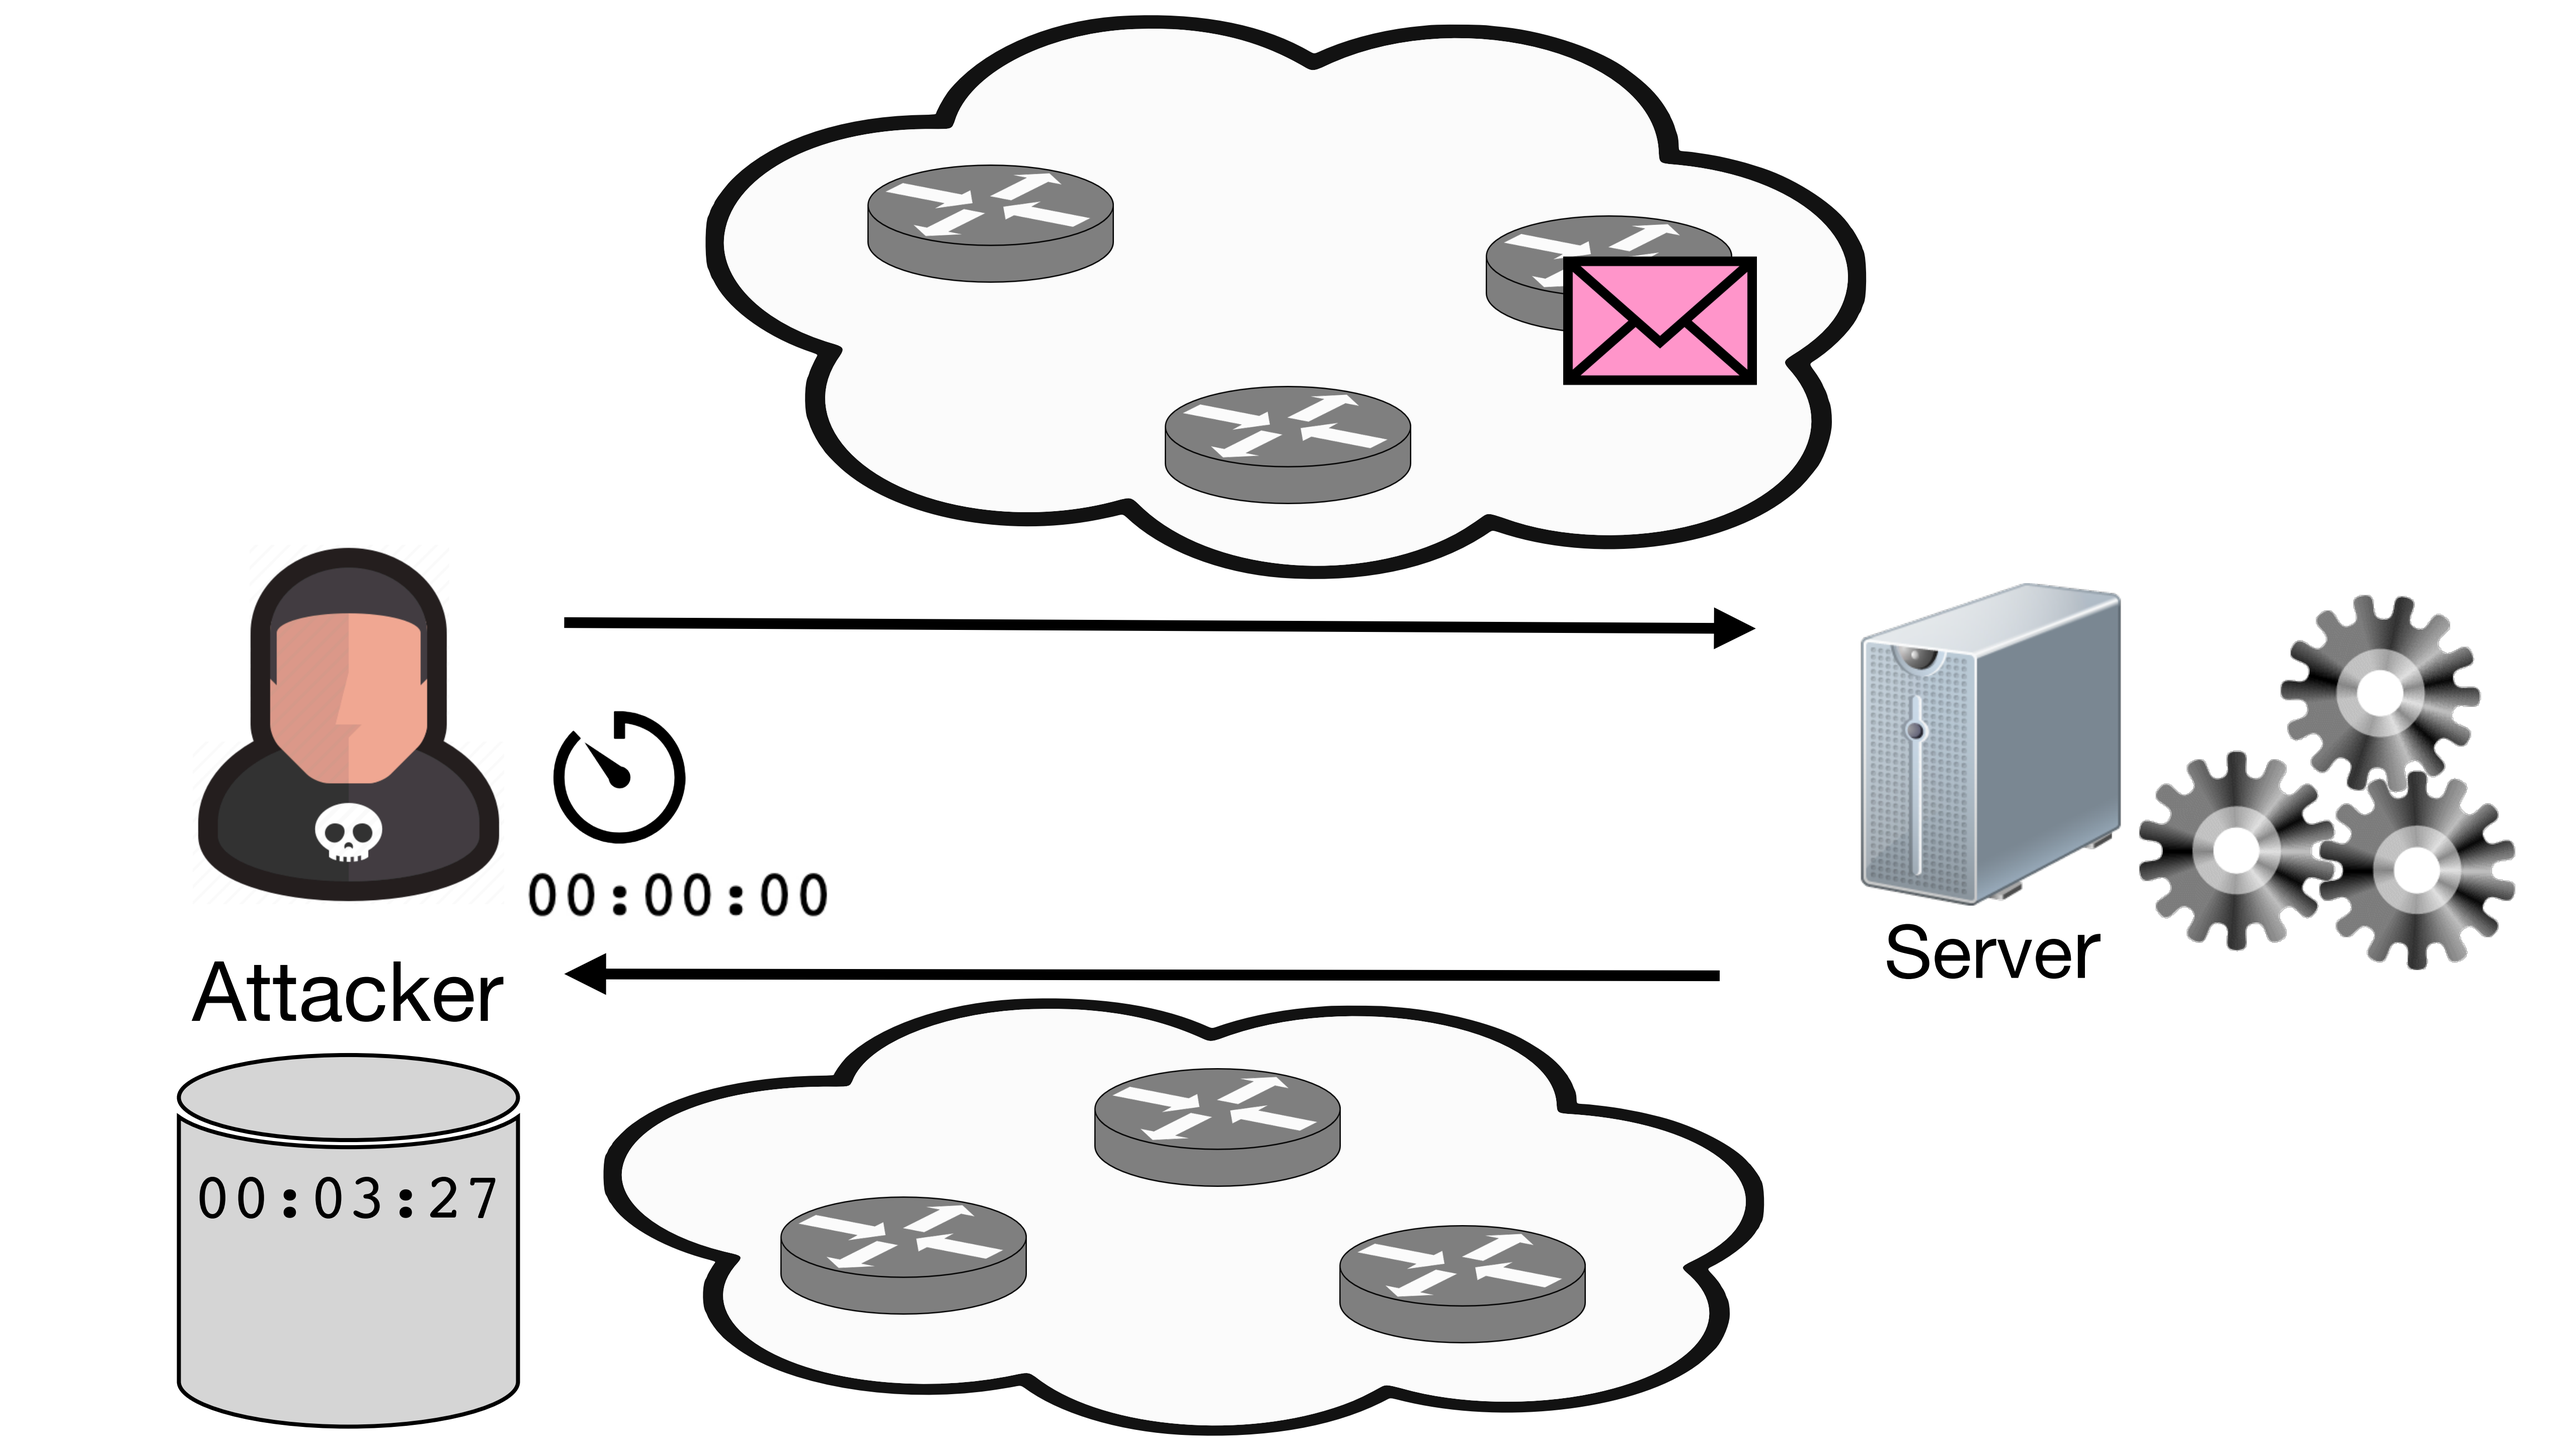
00:00:00
Attacker
Server
00:03:27

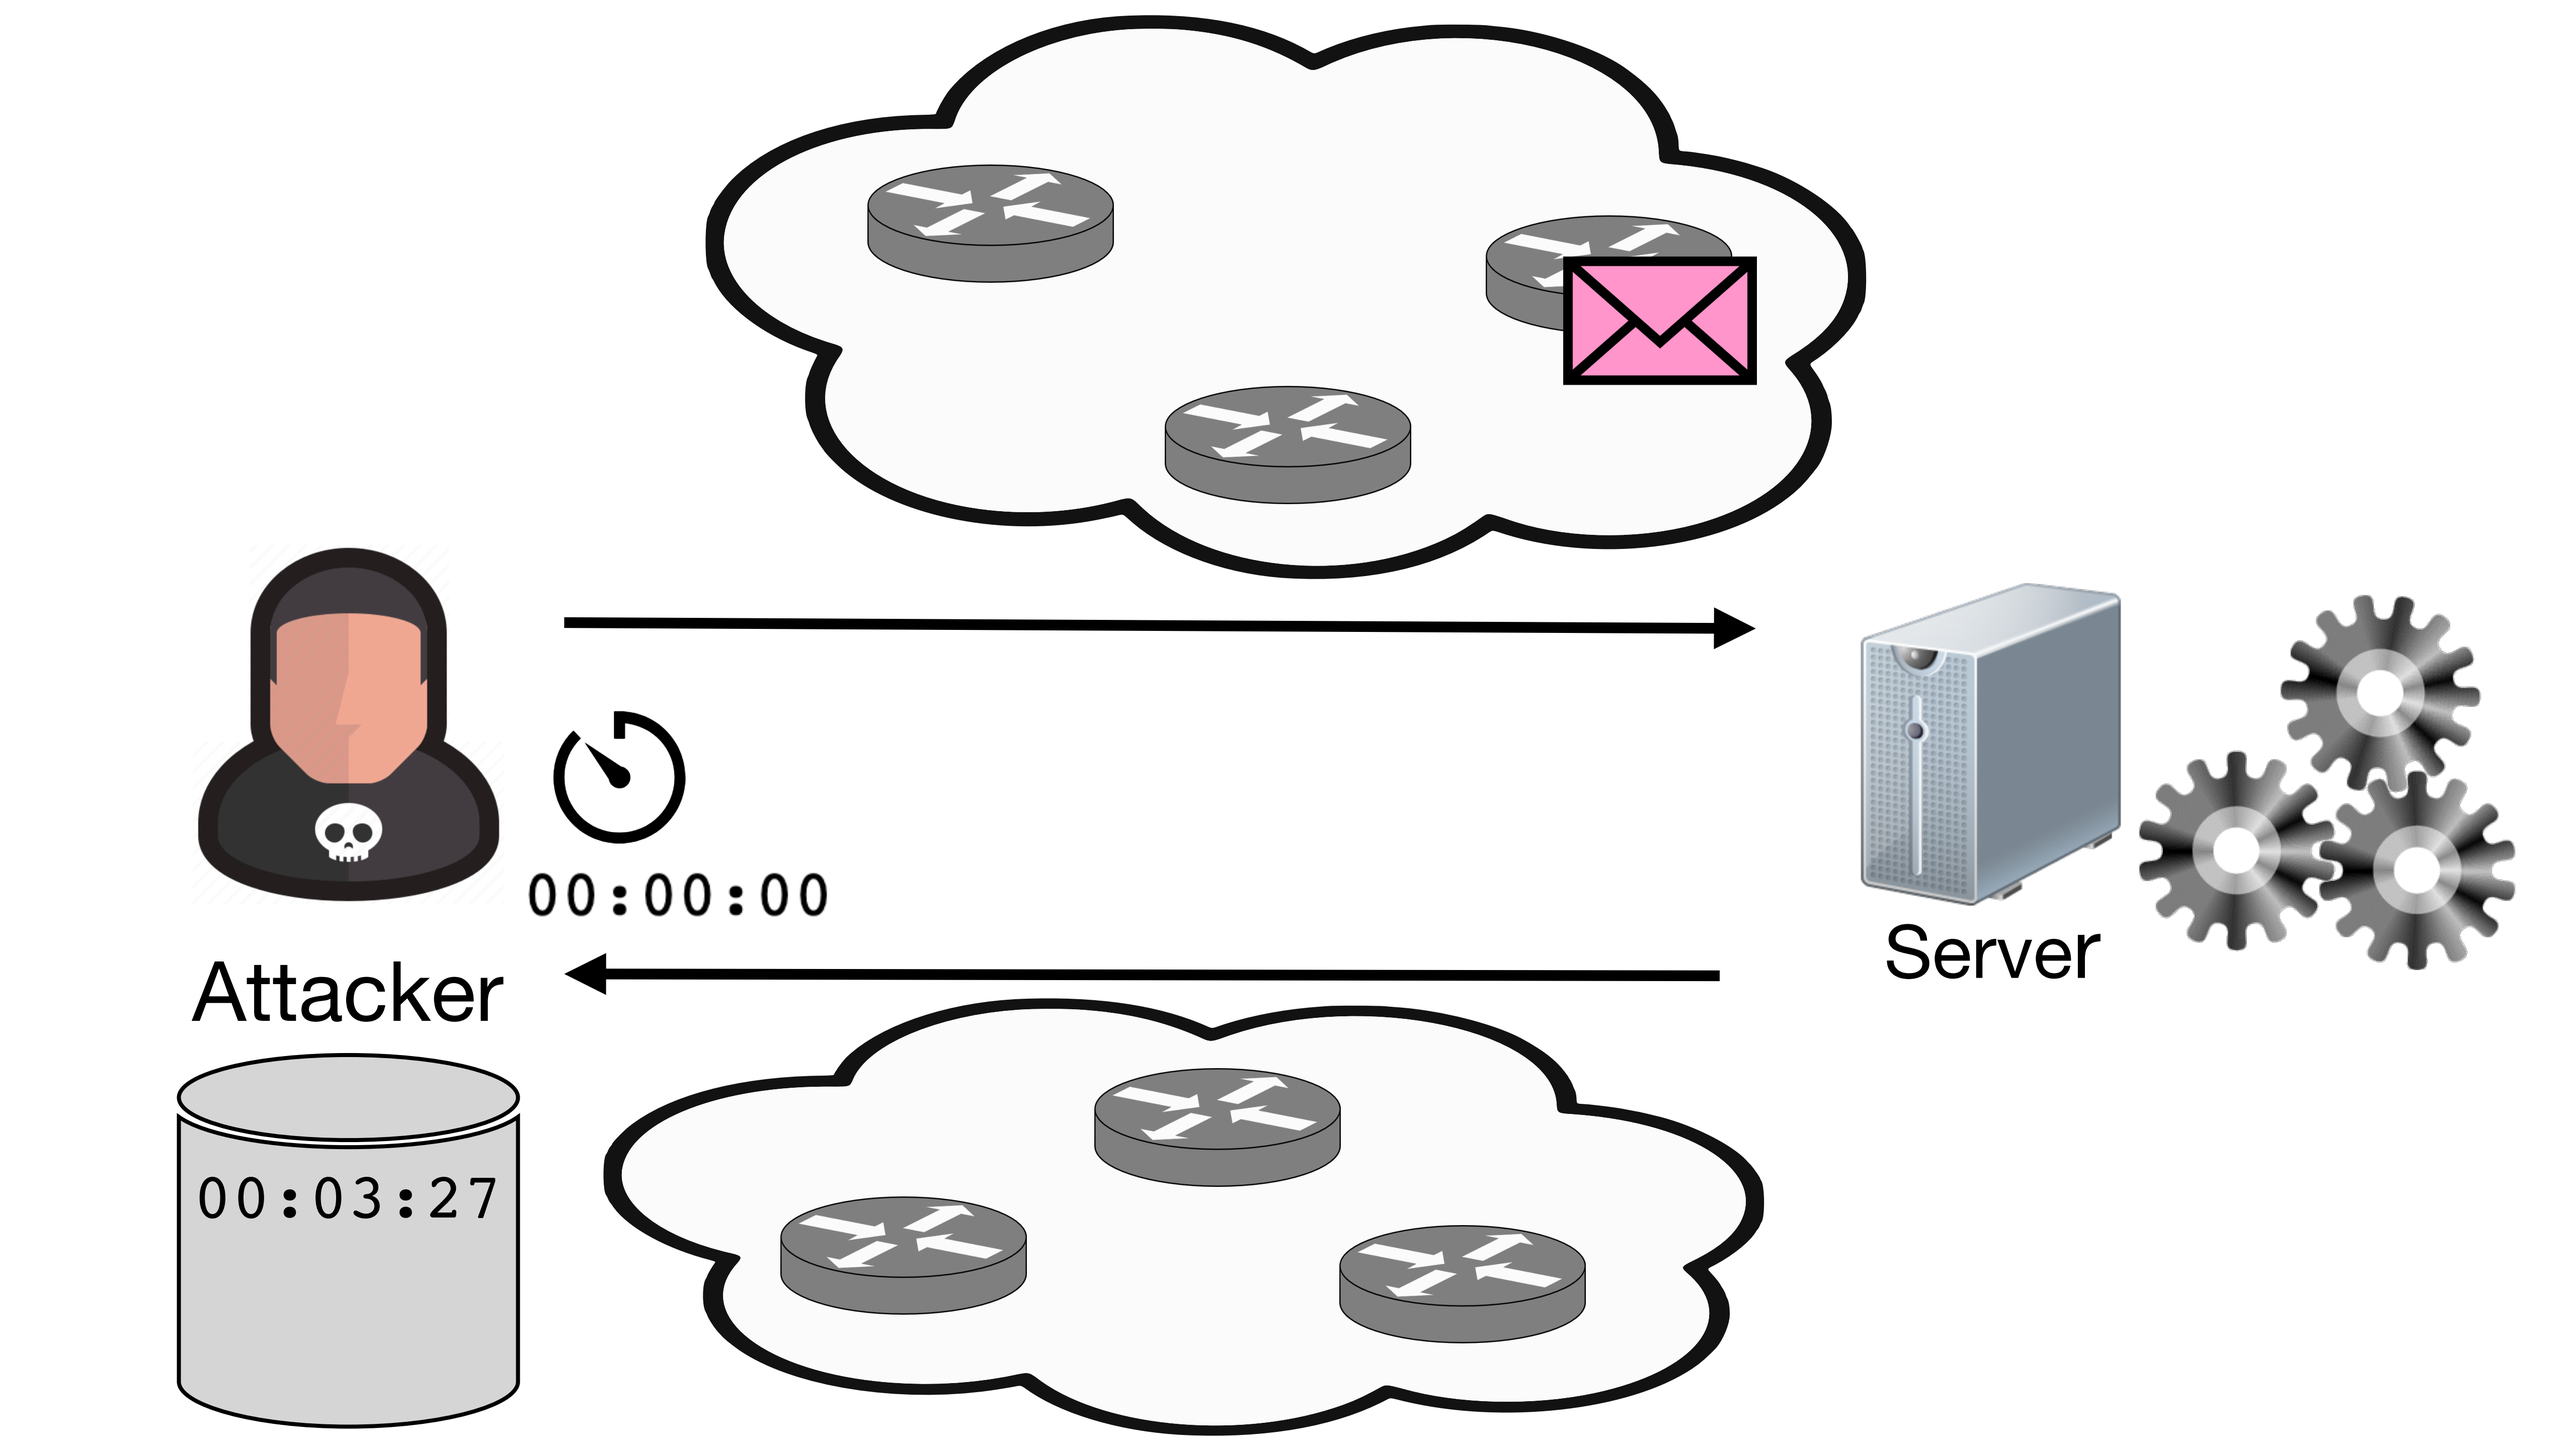
00:00:00
Attacker
Server
00:03:27

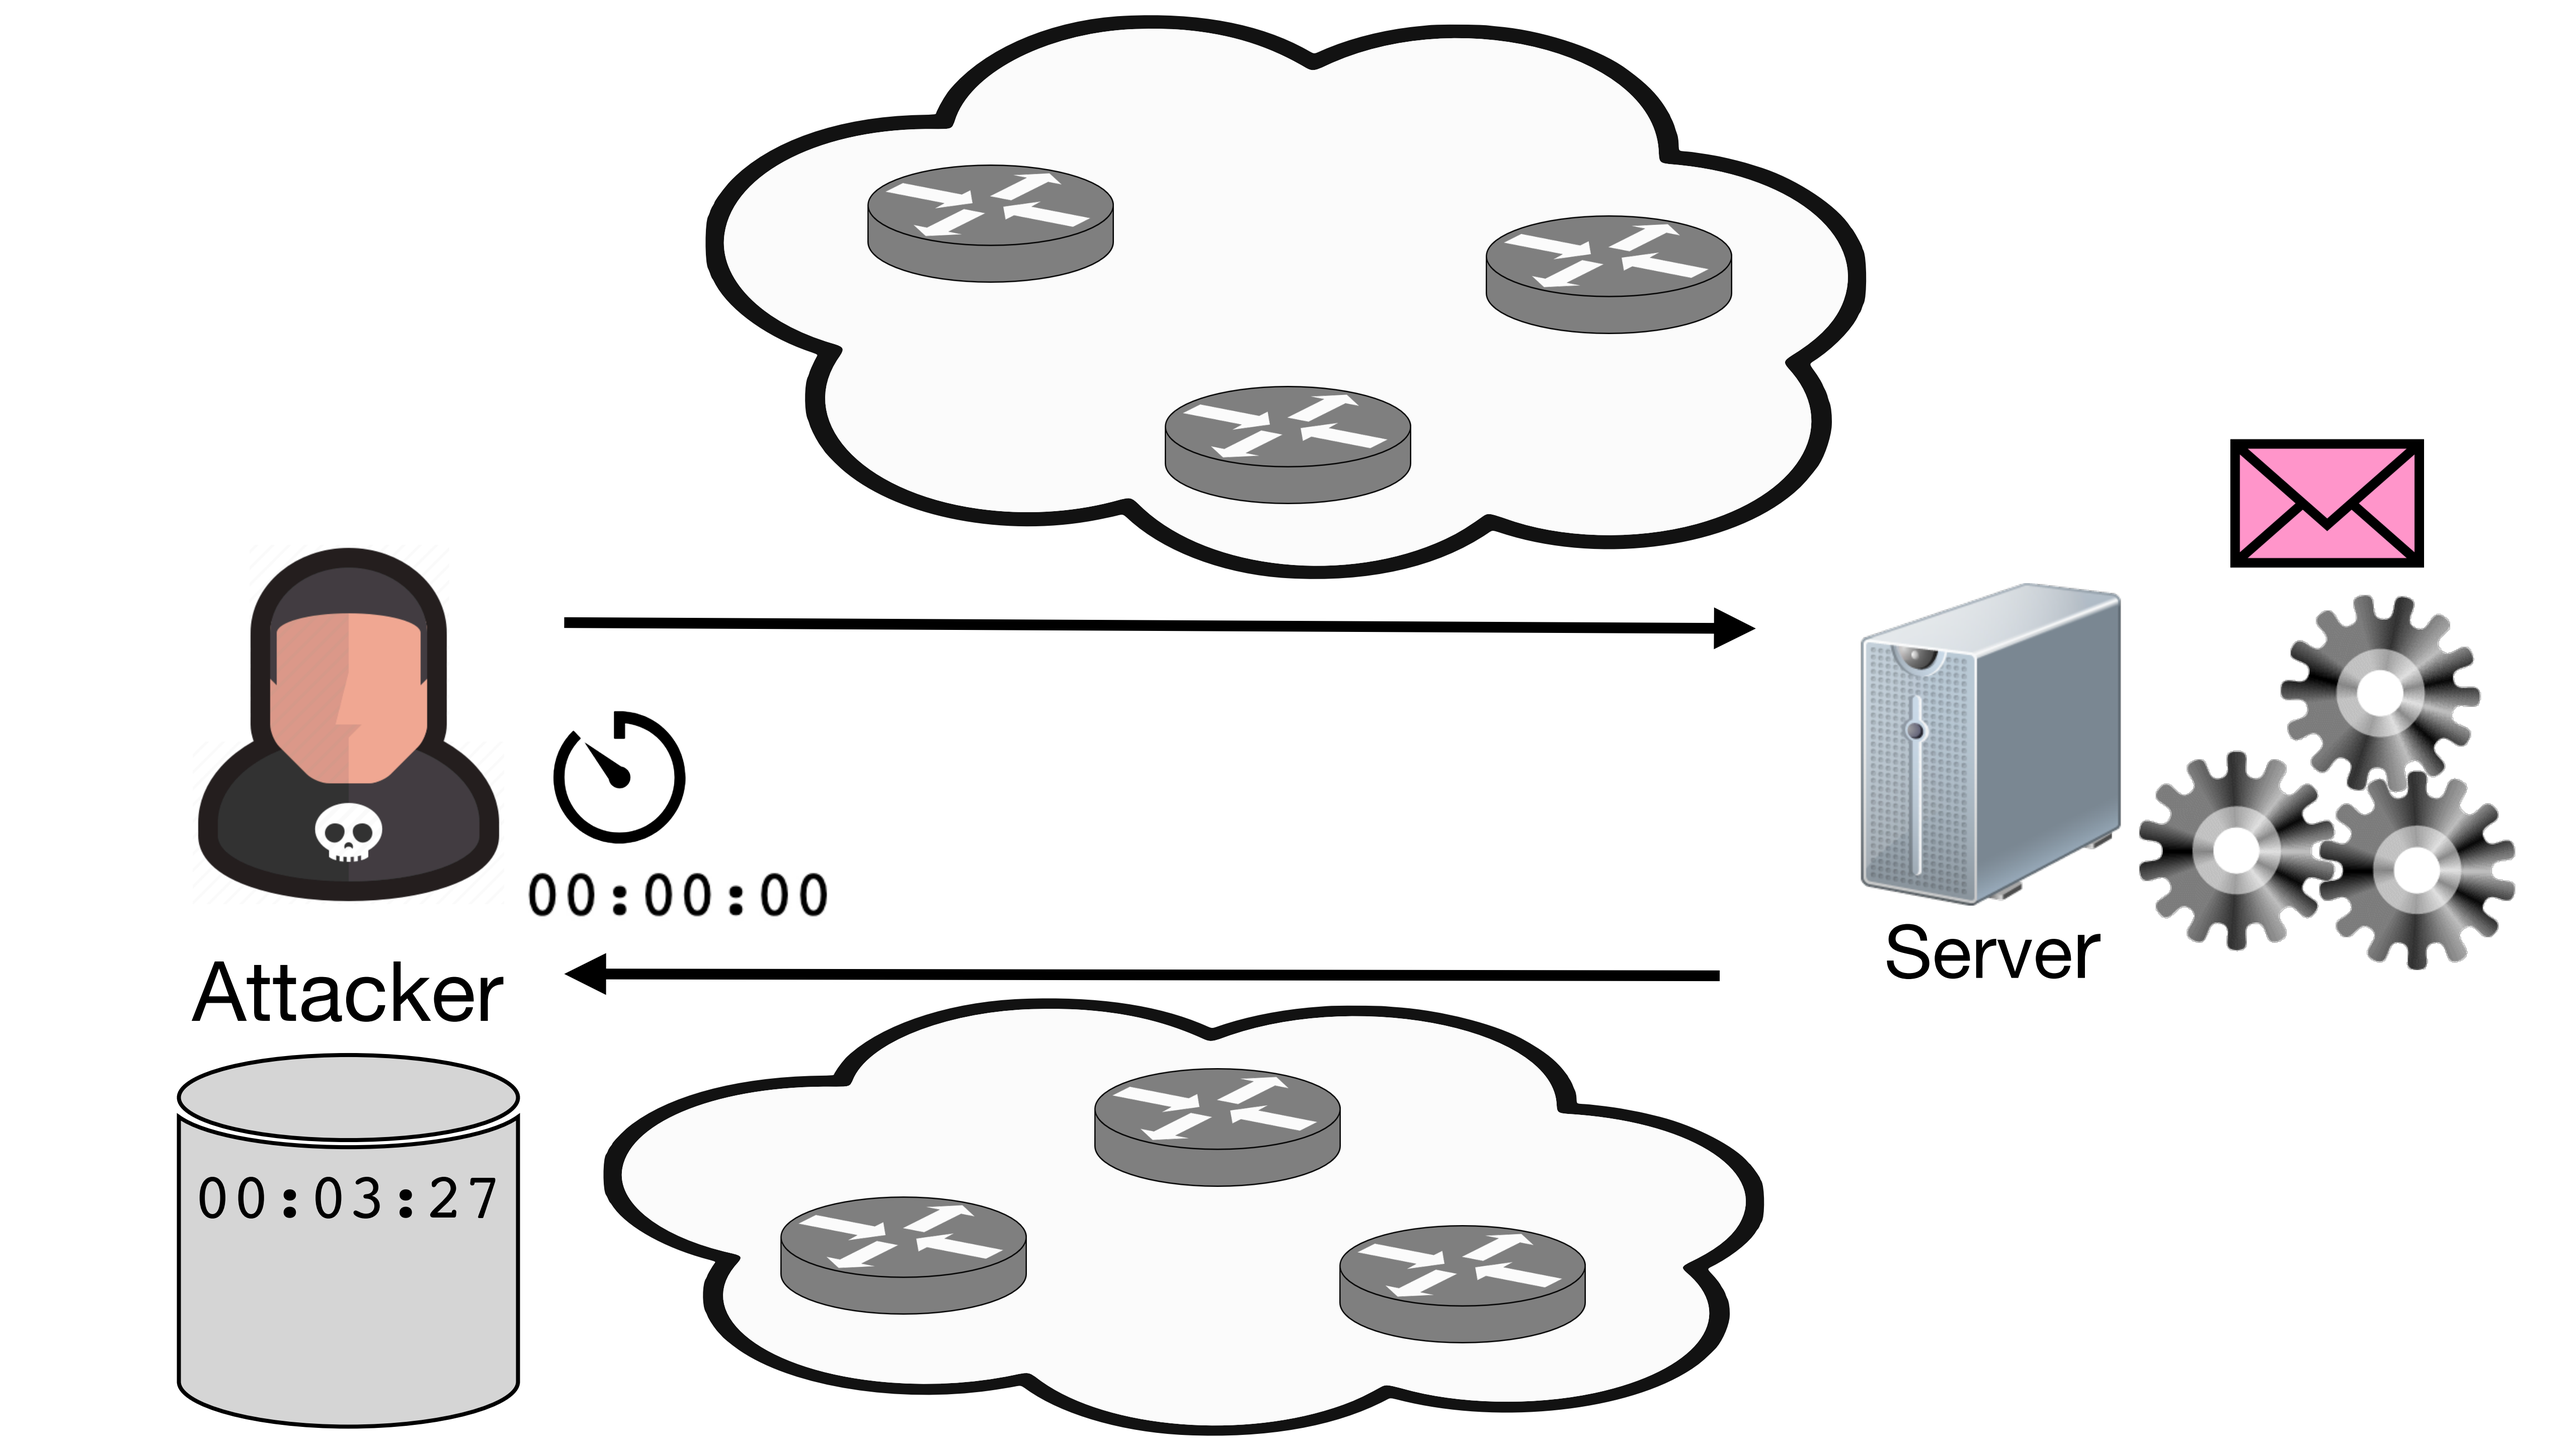
00:00:00
Server
Attacker
00:03:27

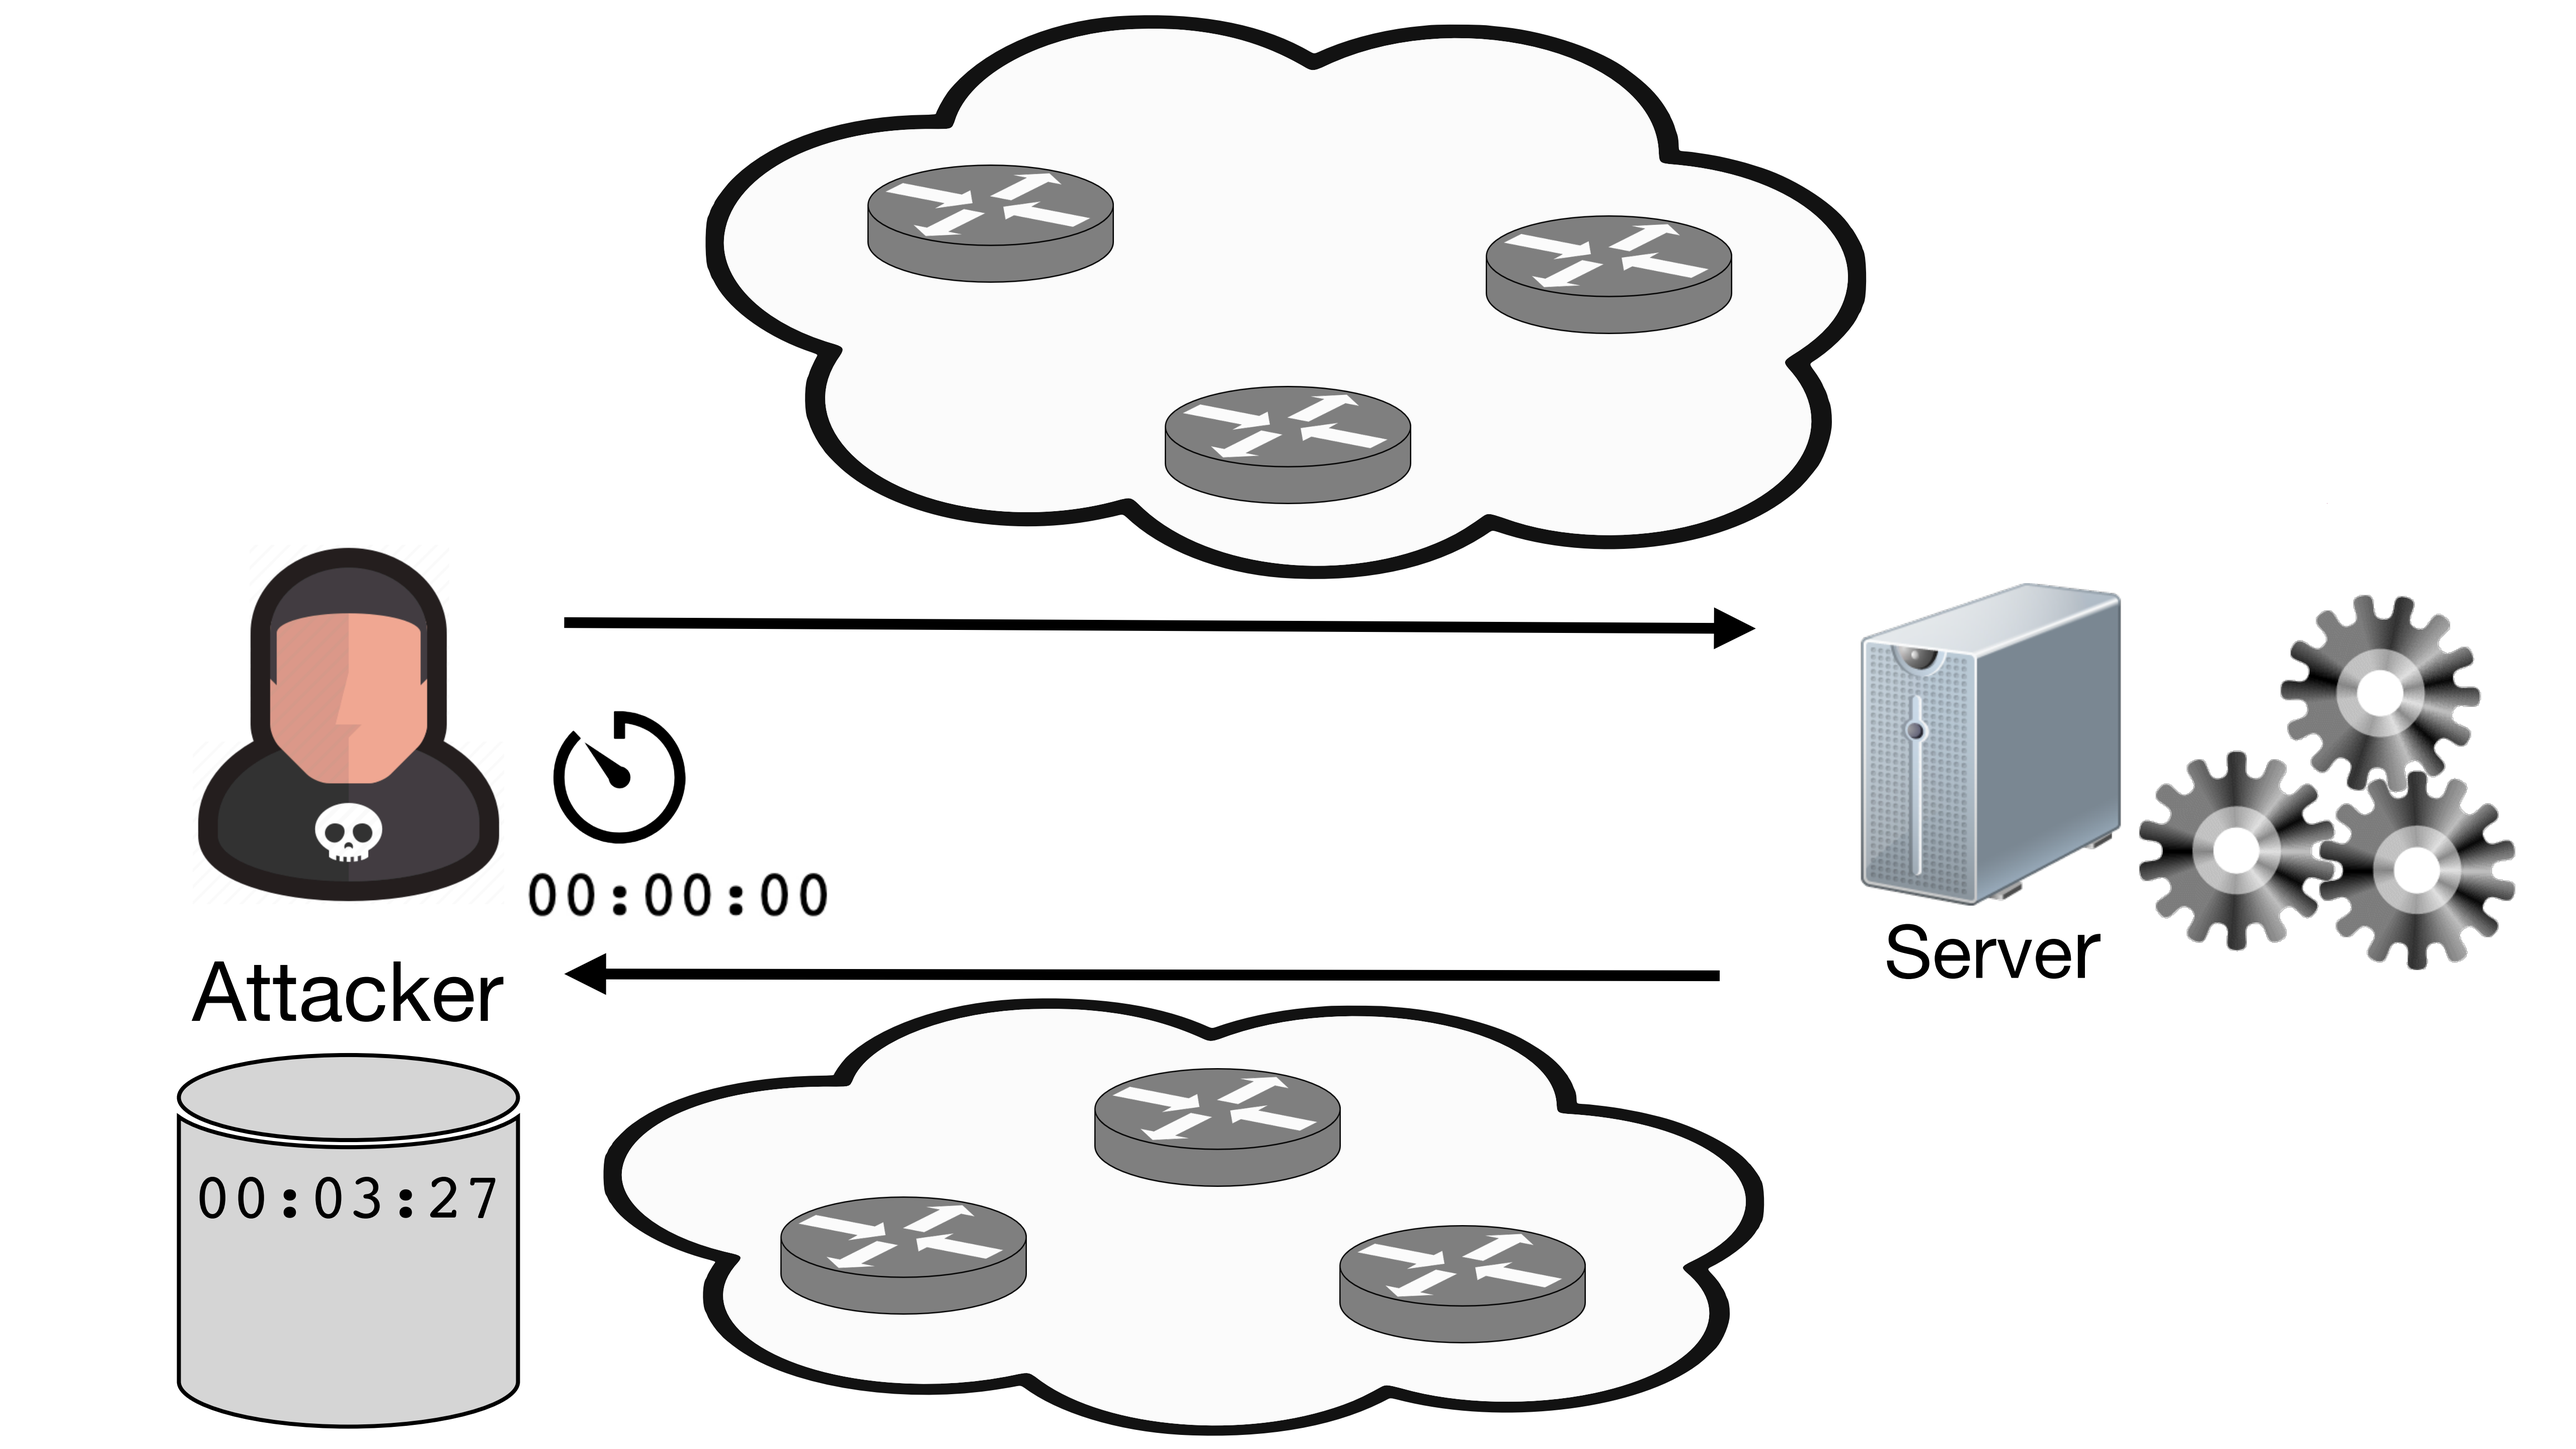
00:00:00
Server
Attacker
00:03:27

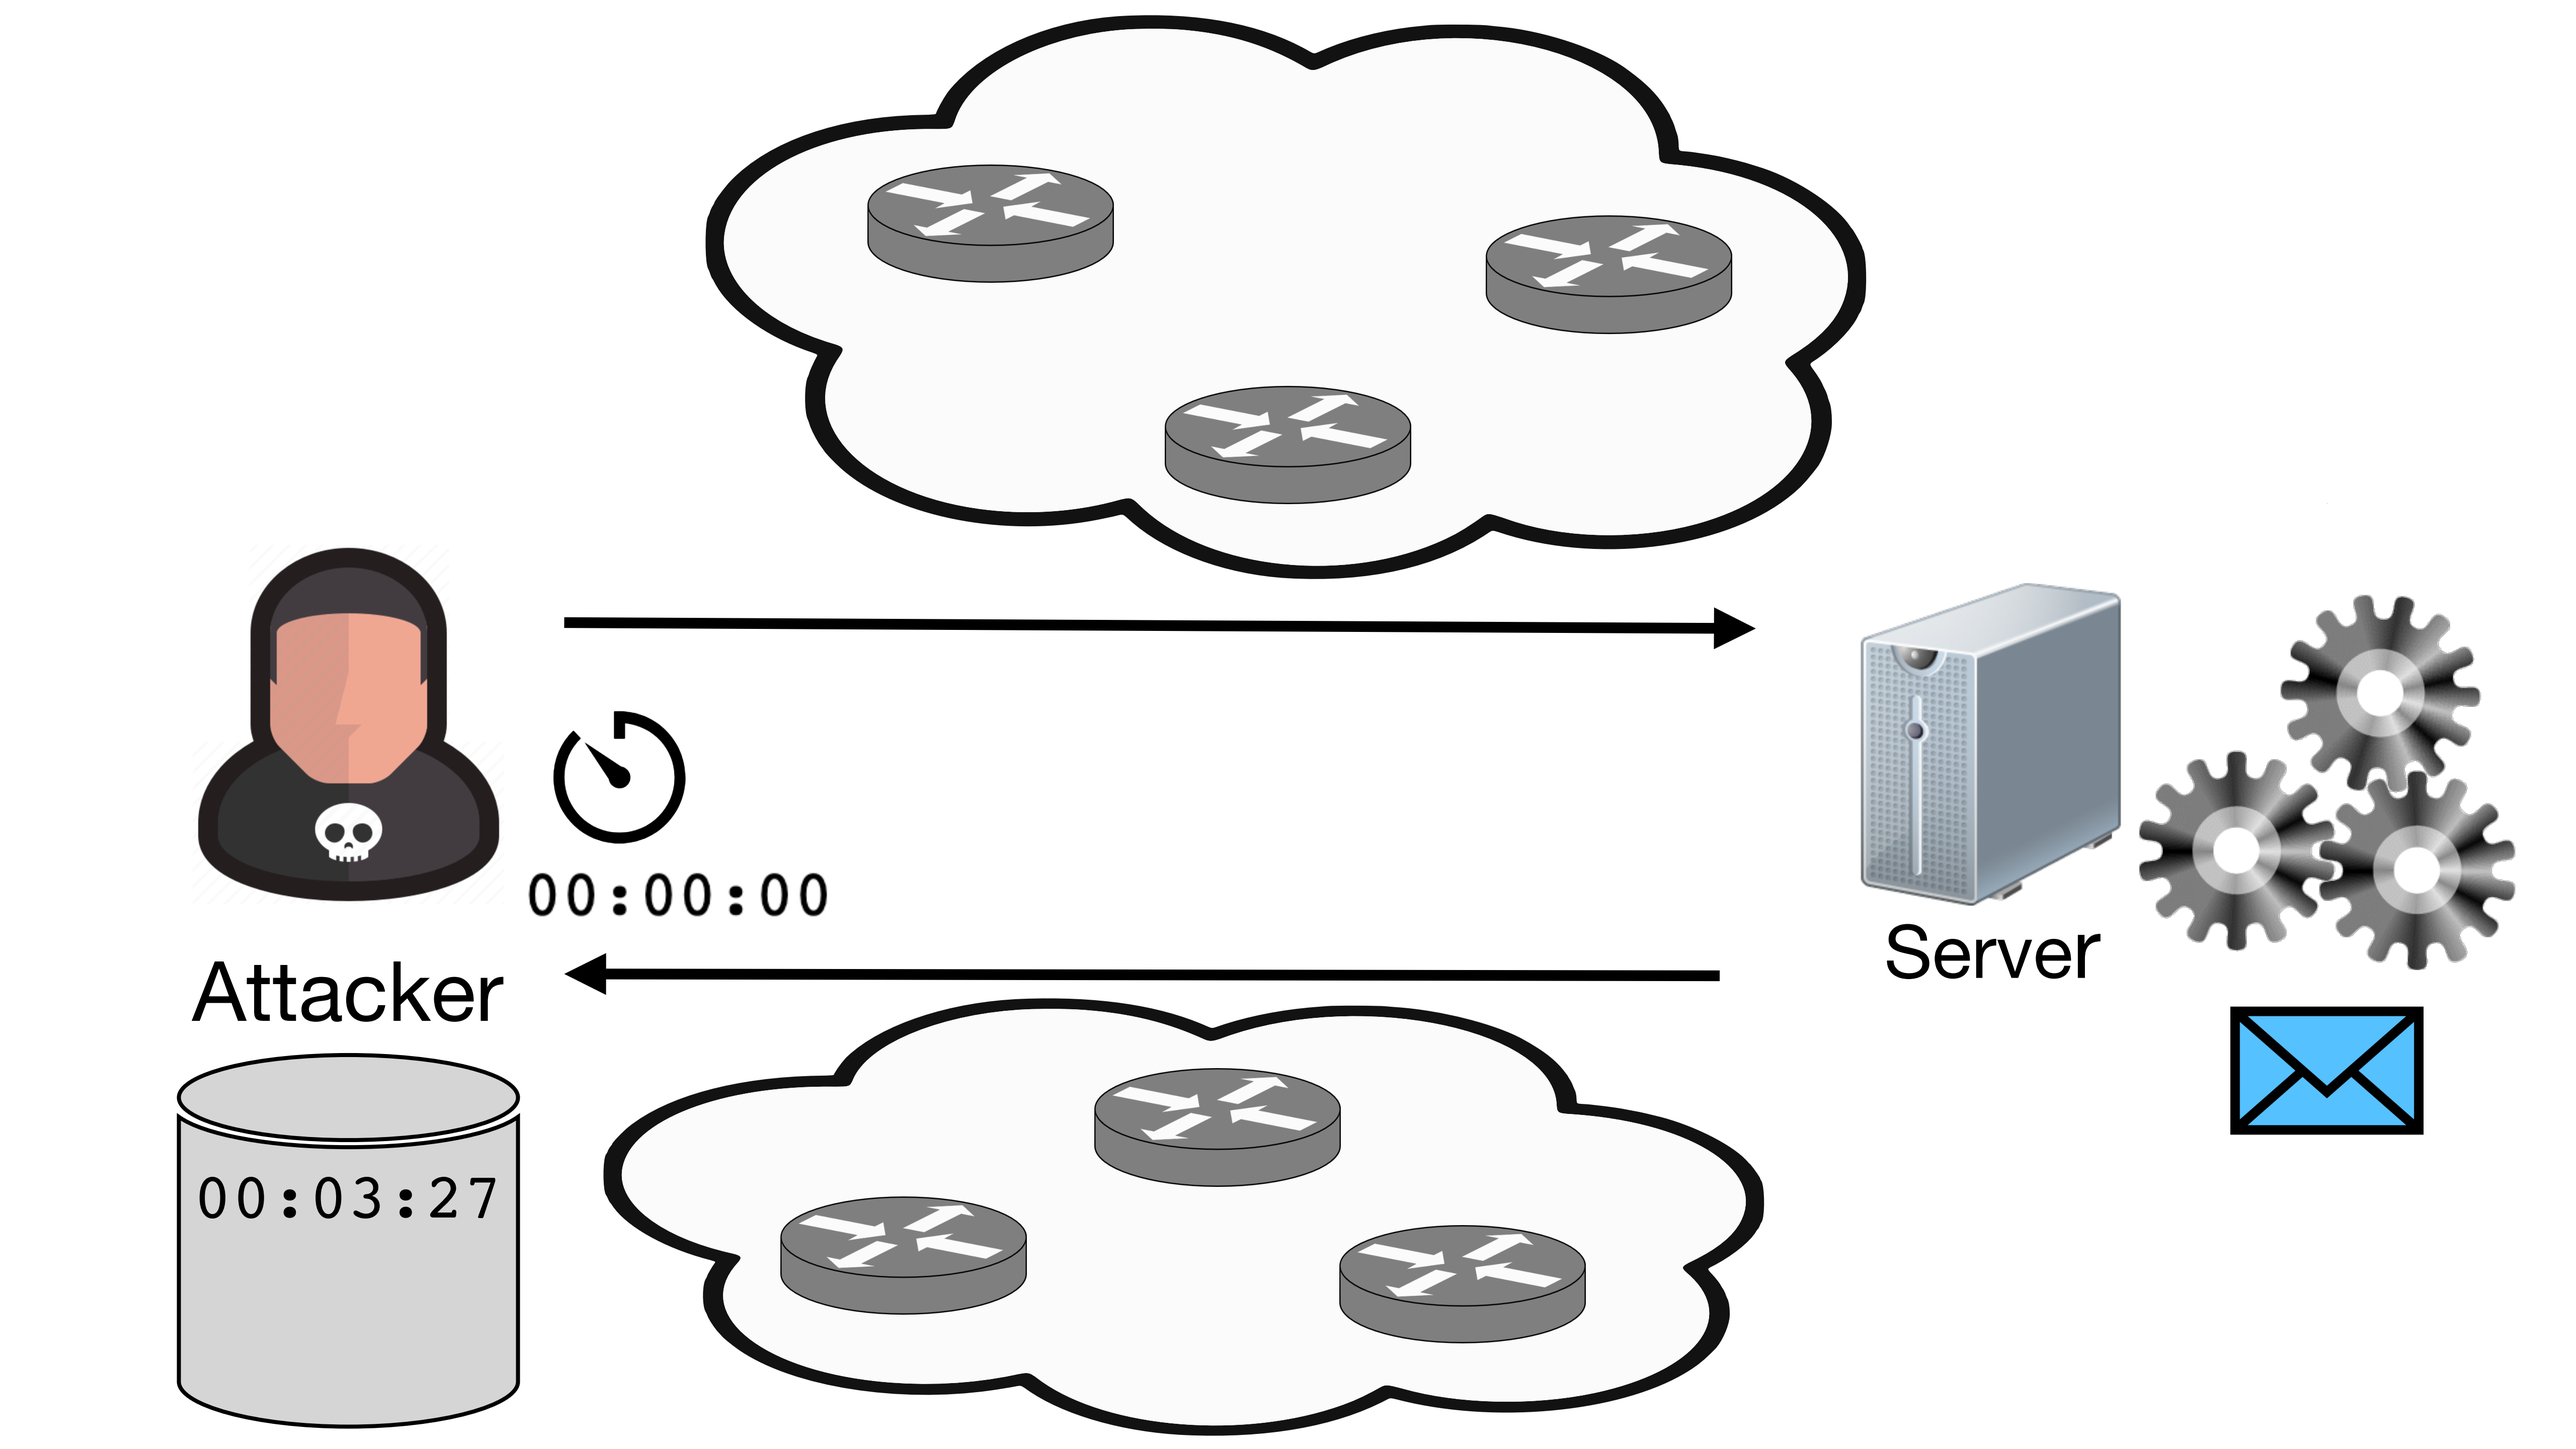
00:00:00
Server
Attacker
00:03:27

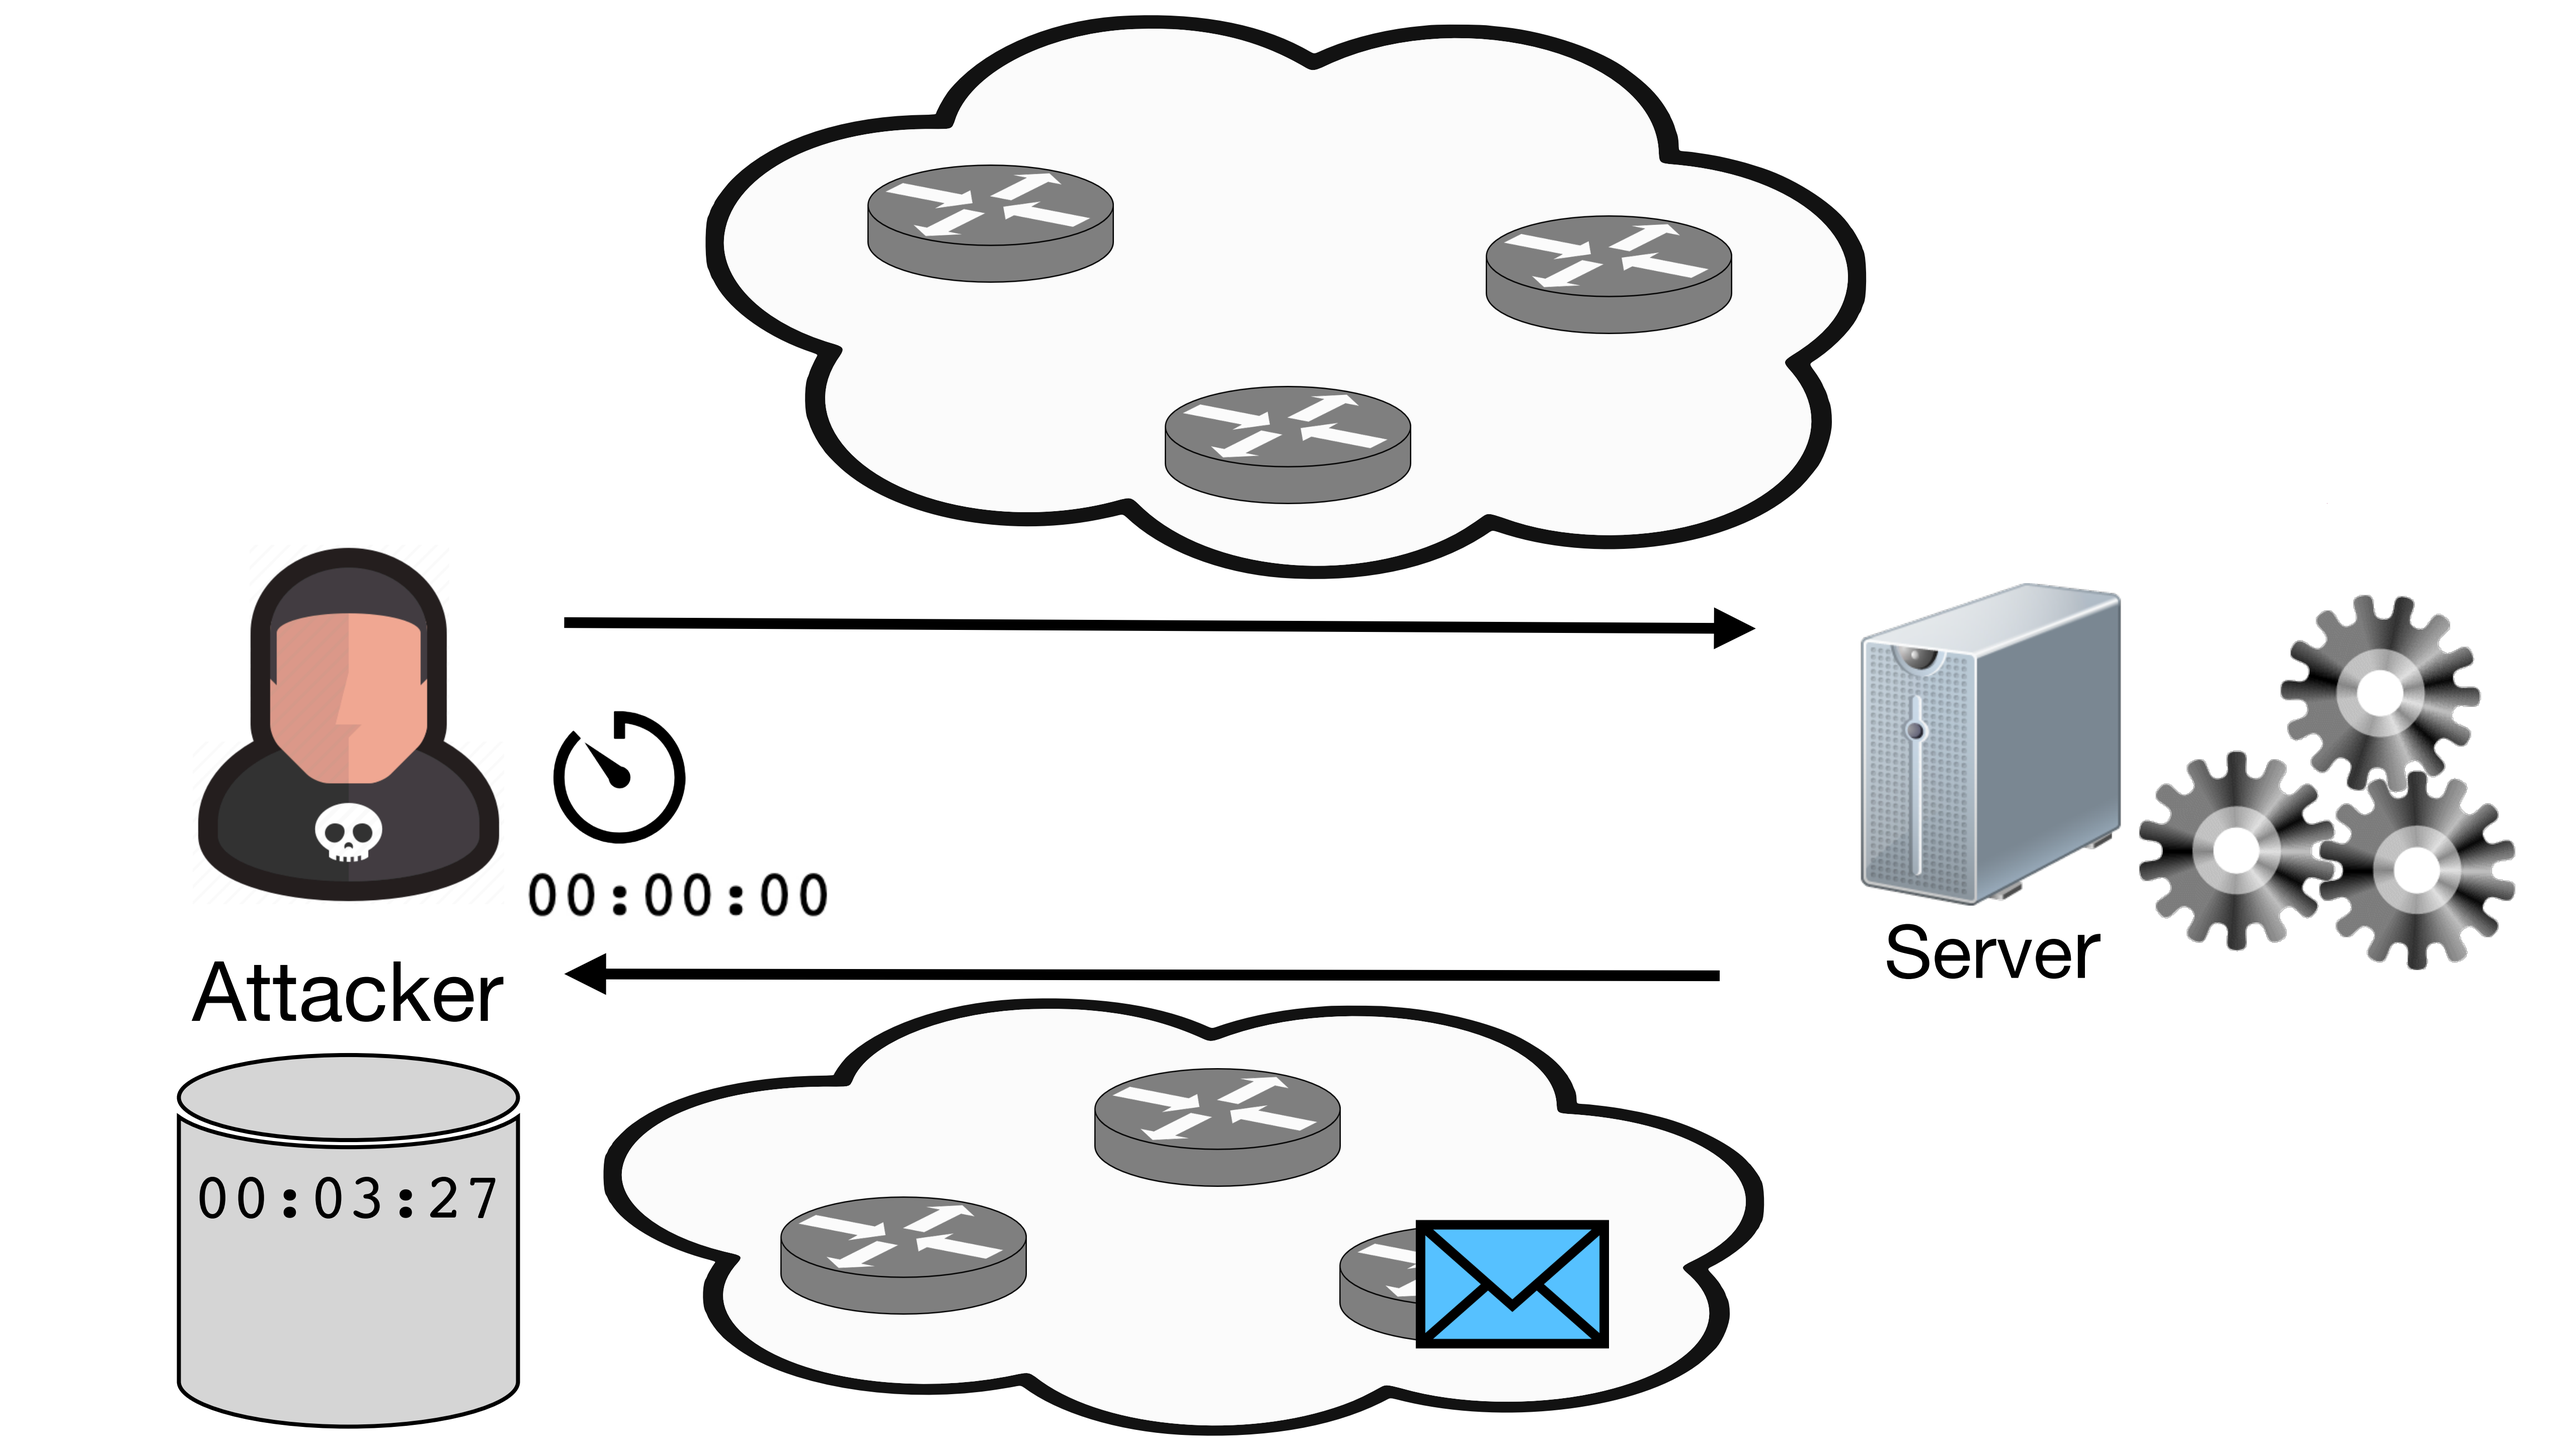
00:00:00
Attacker
Server
00:03:27

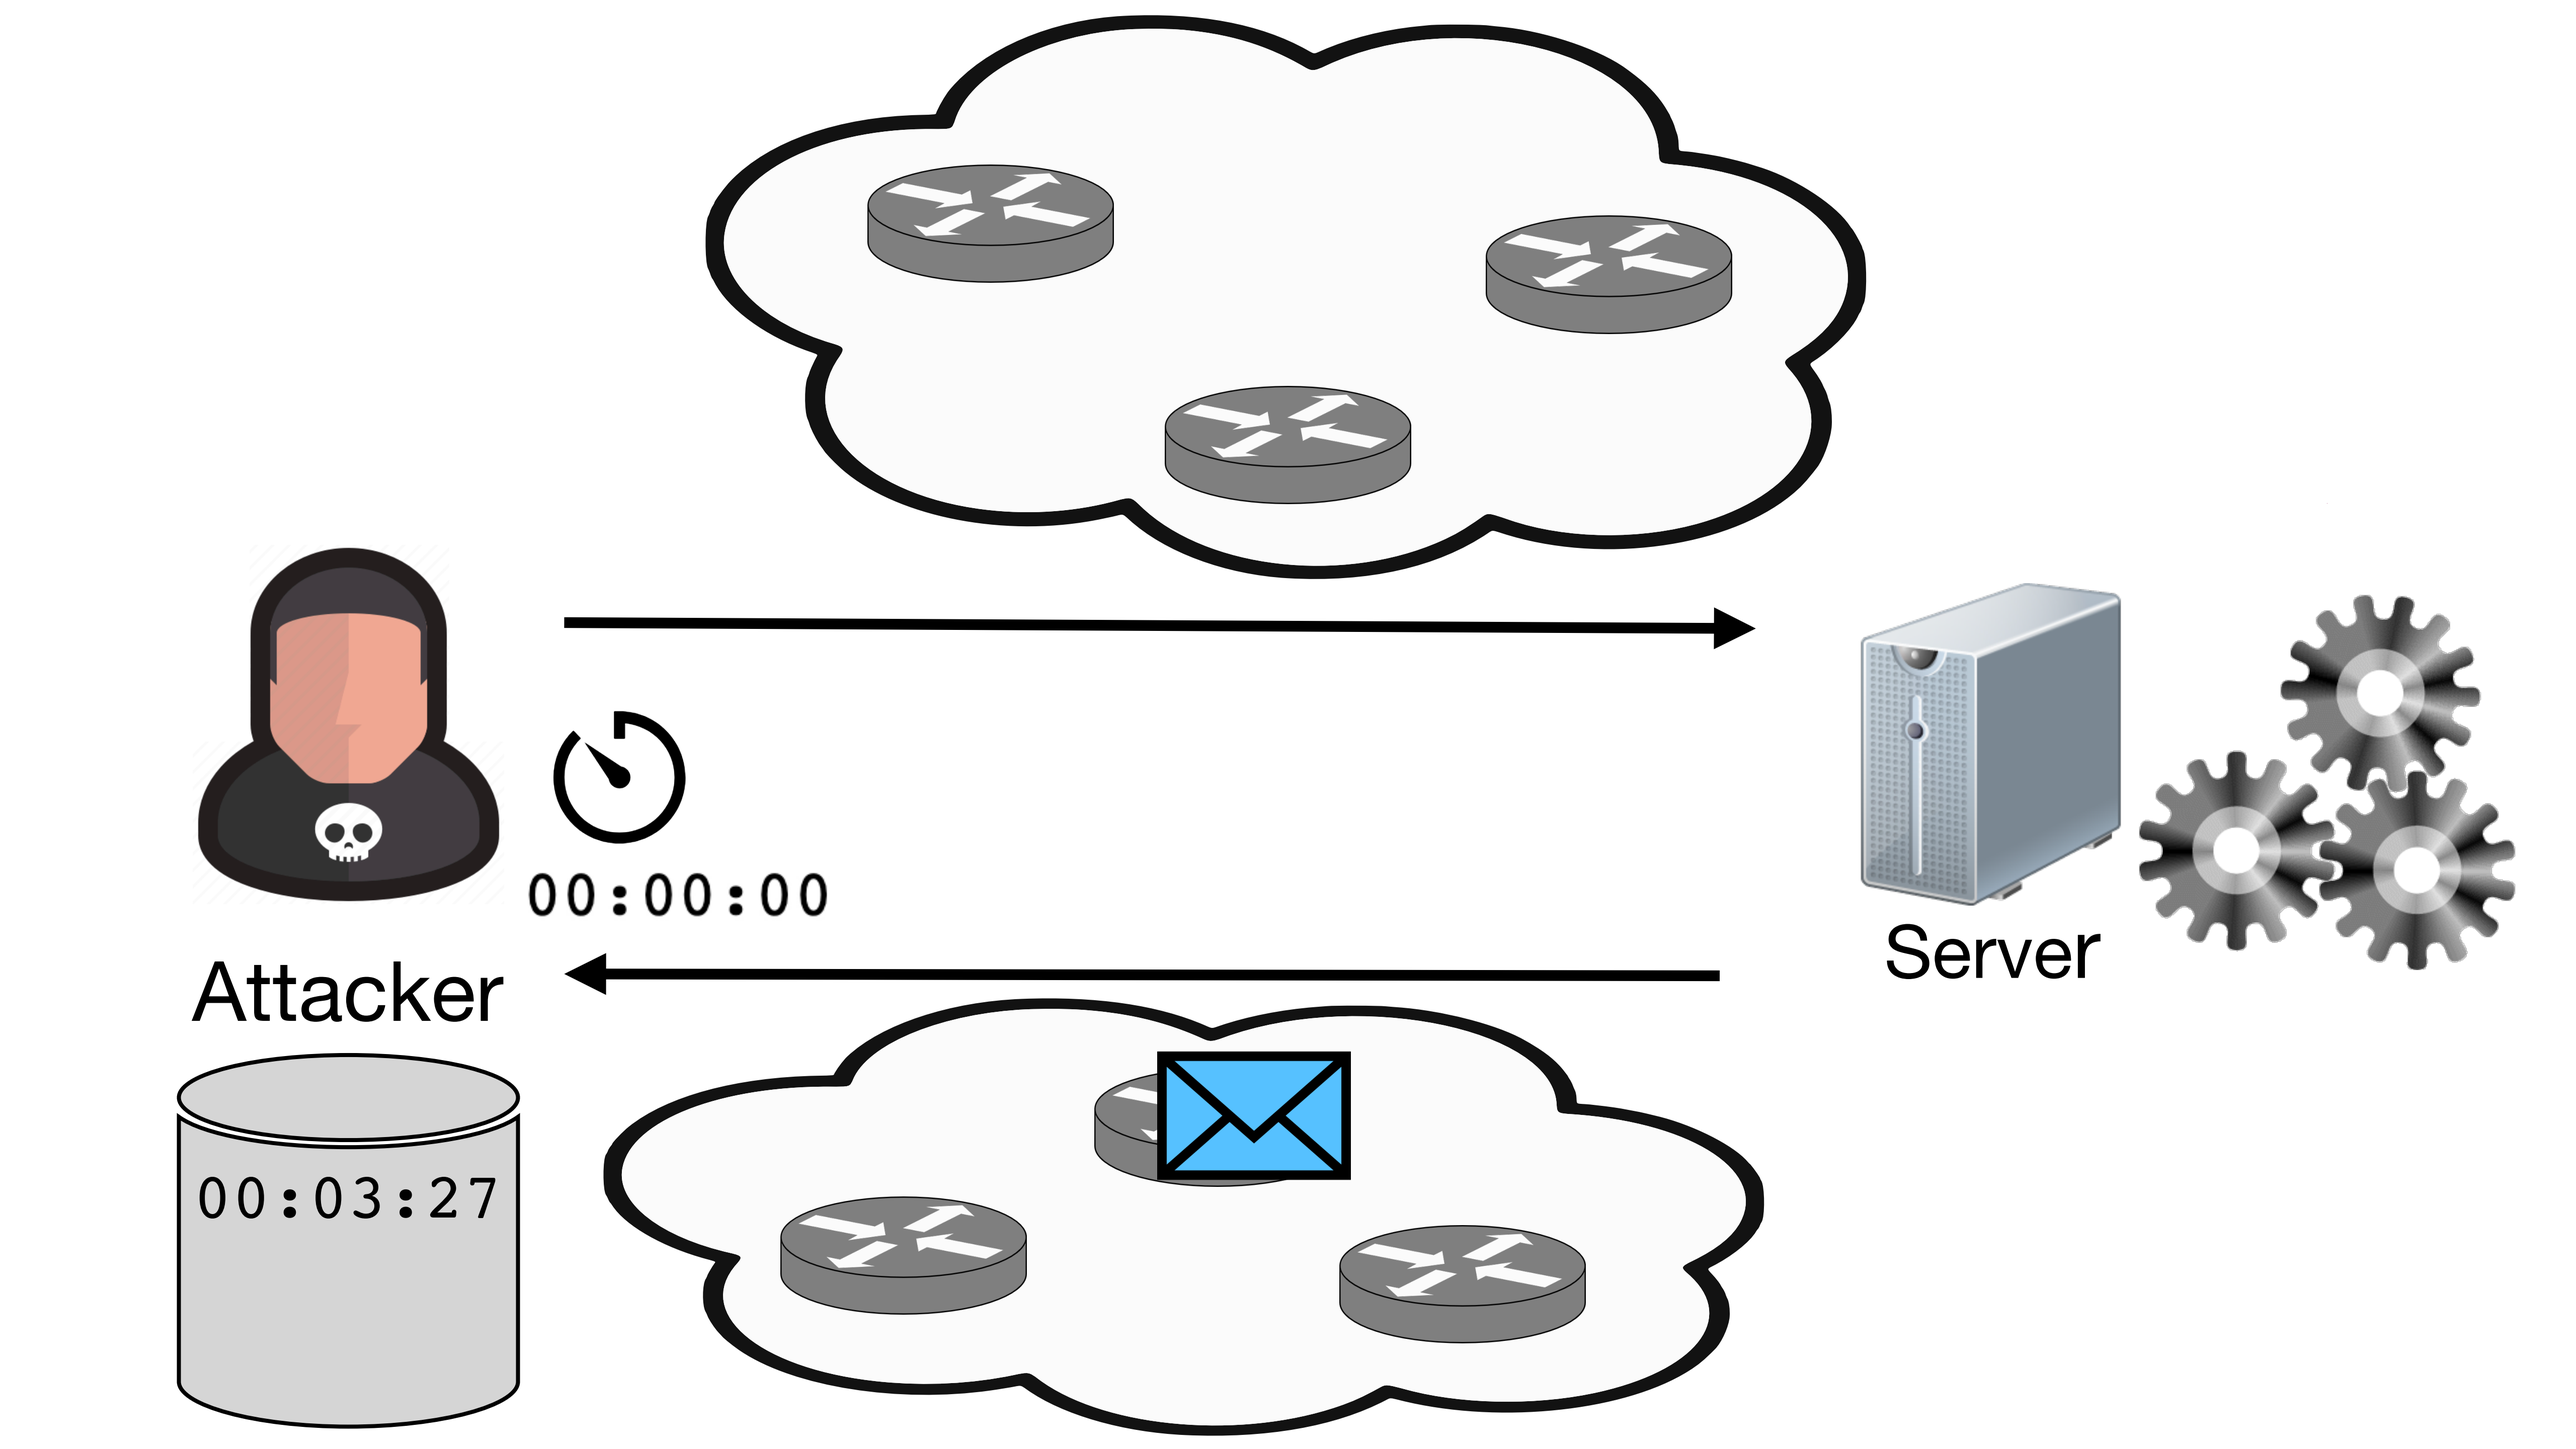
00:00:00
Server
Attacker
00:03:27

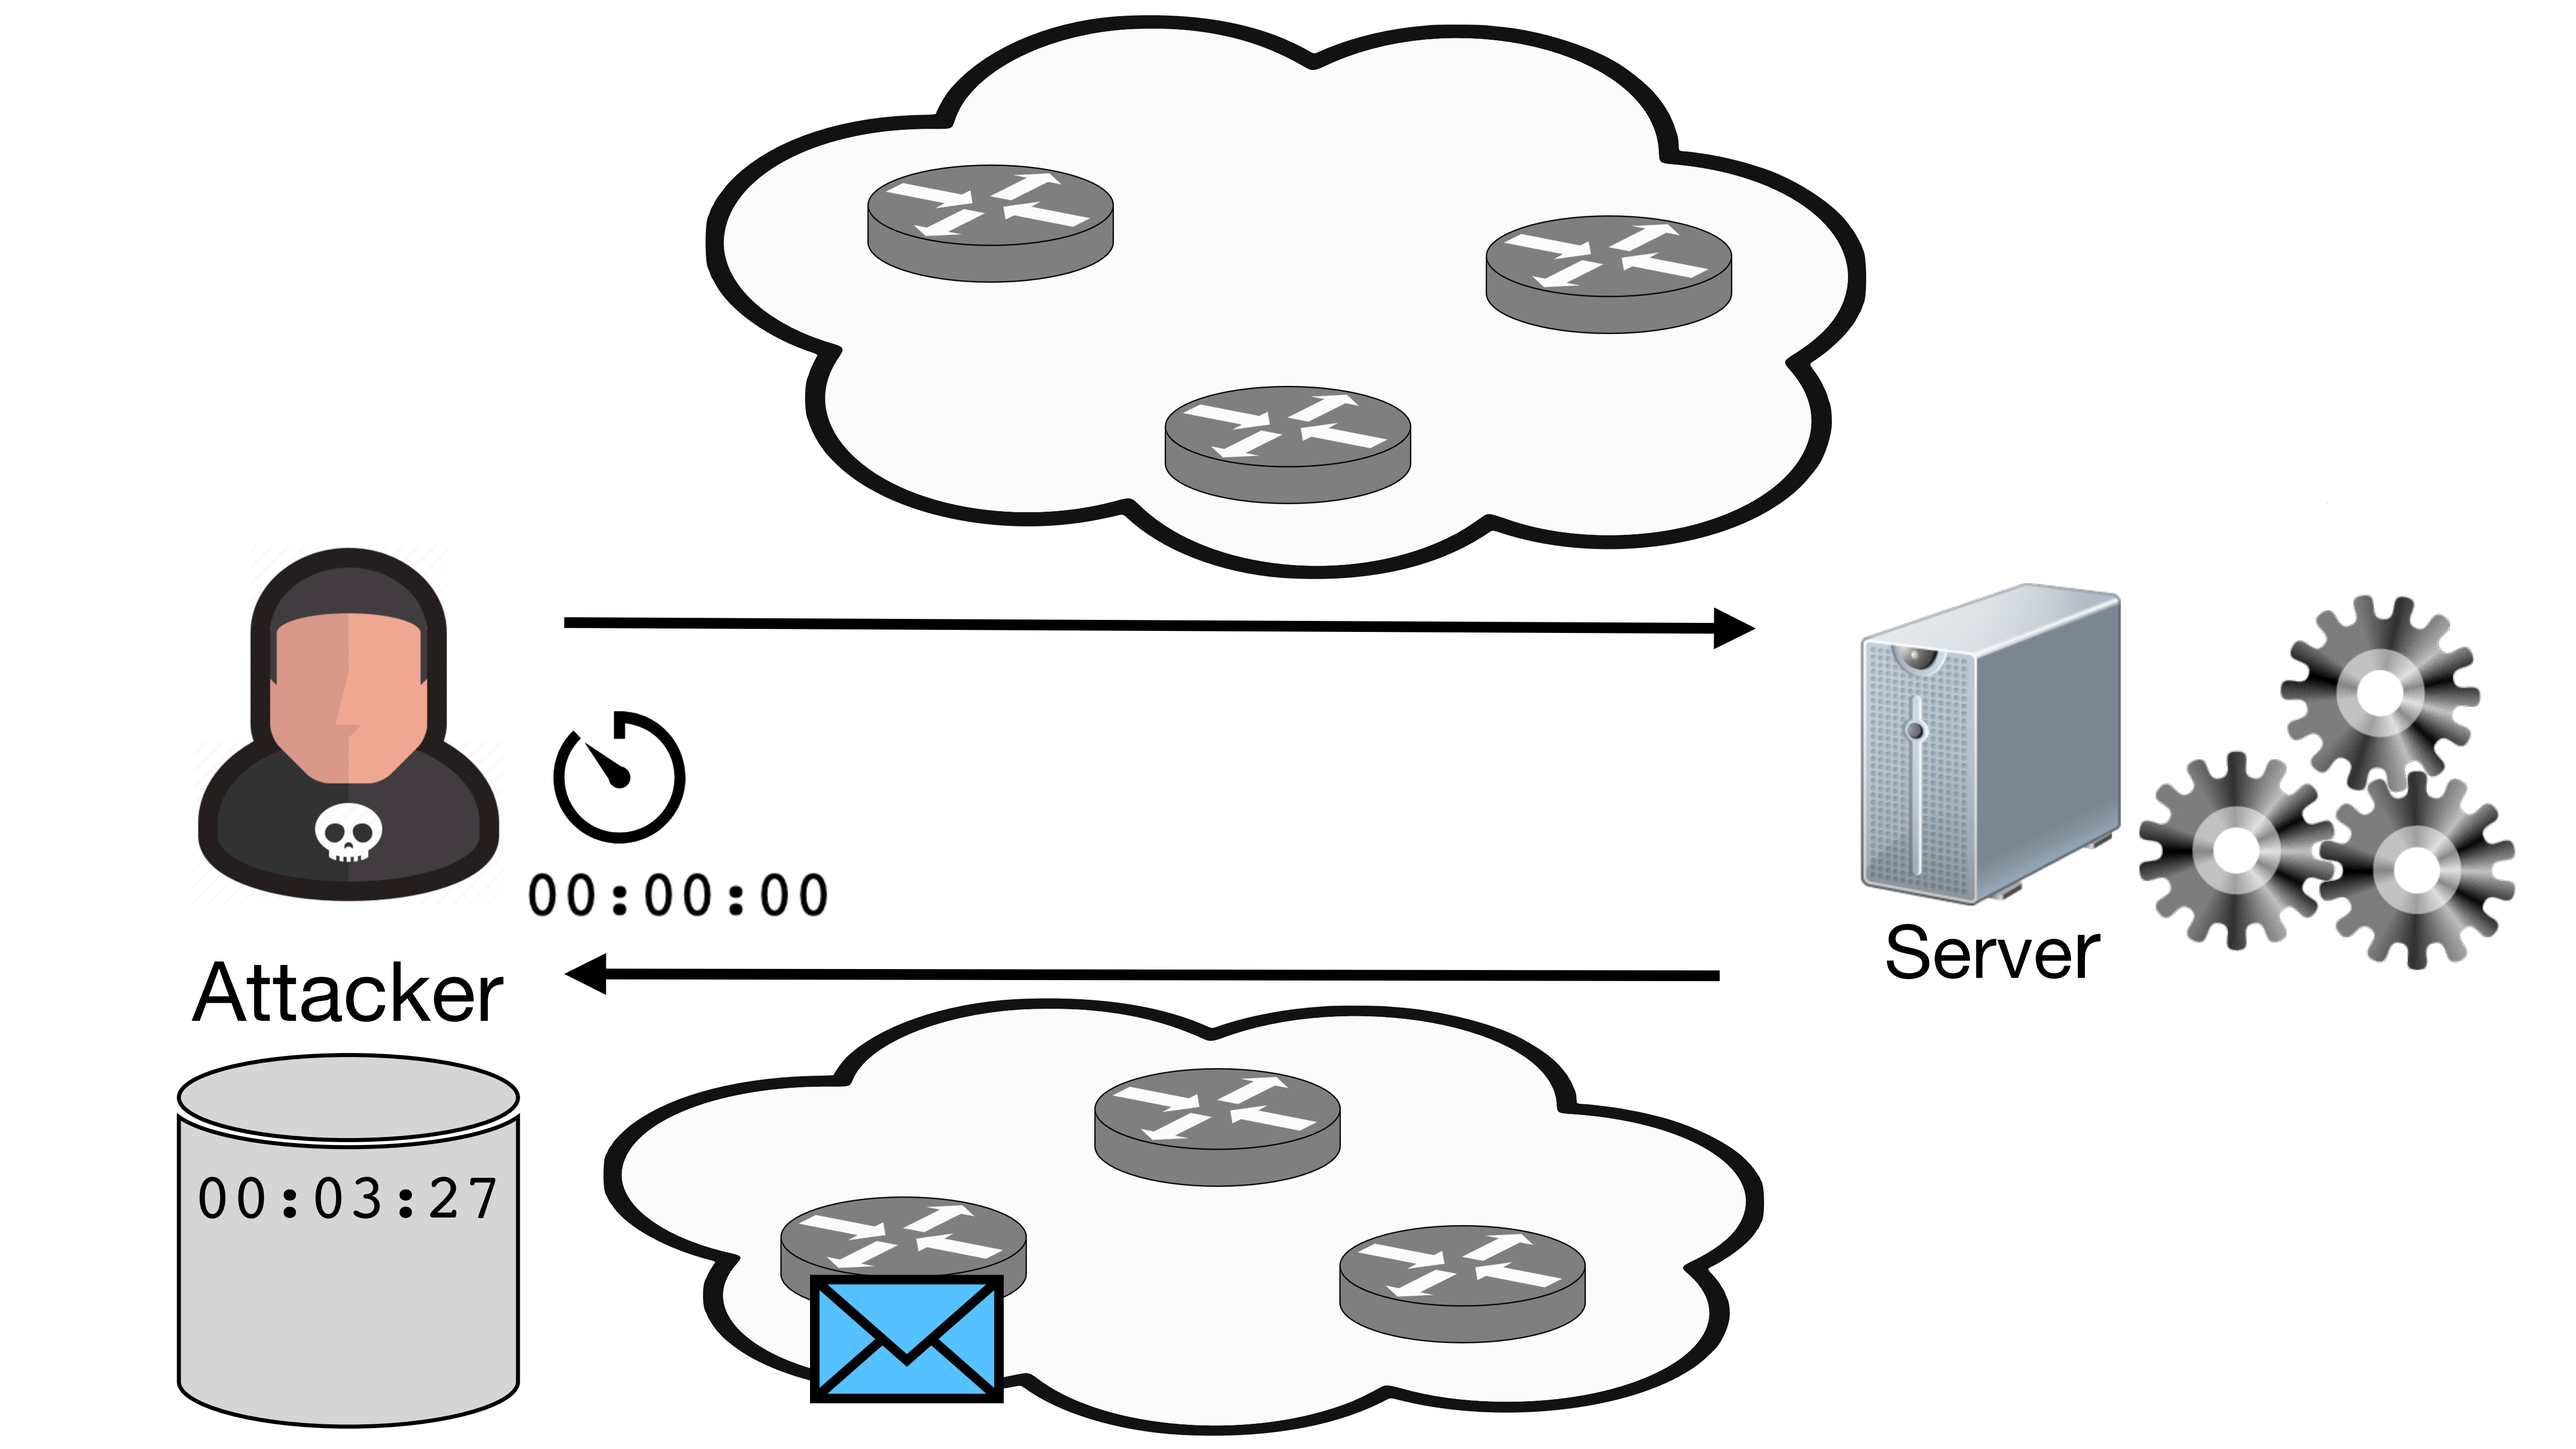
00:00:00
Attacker
Server
00:03:27

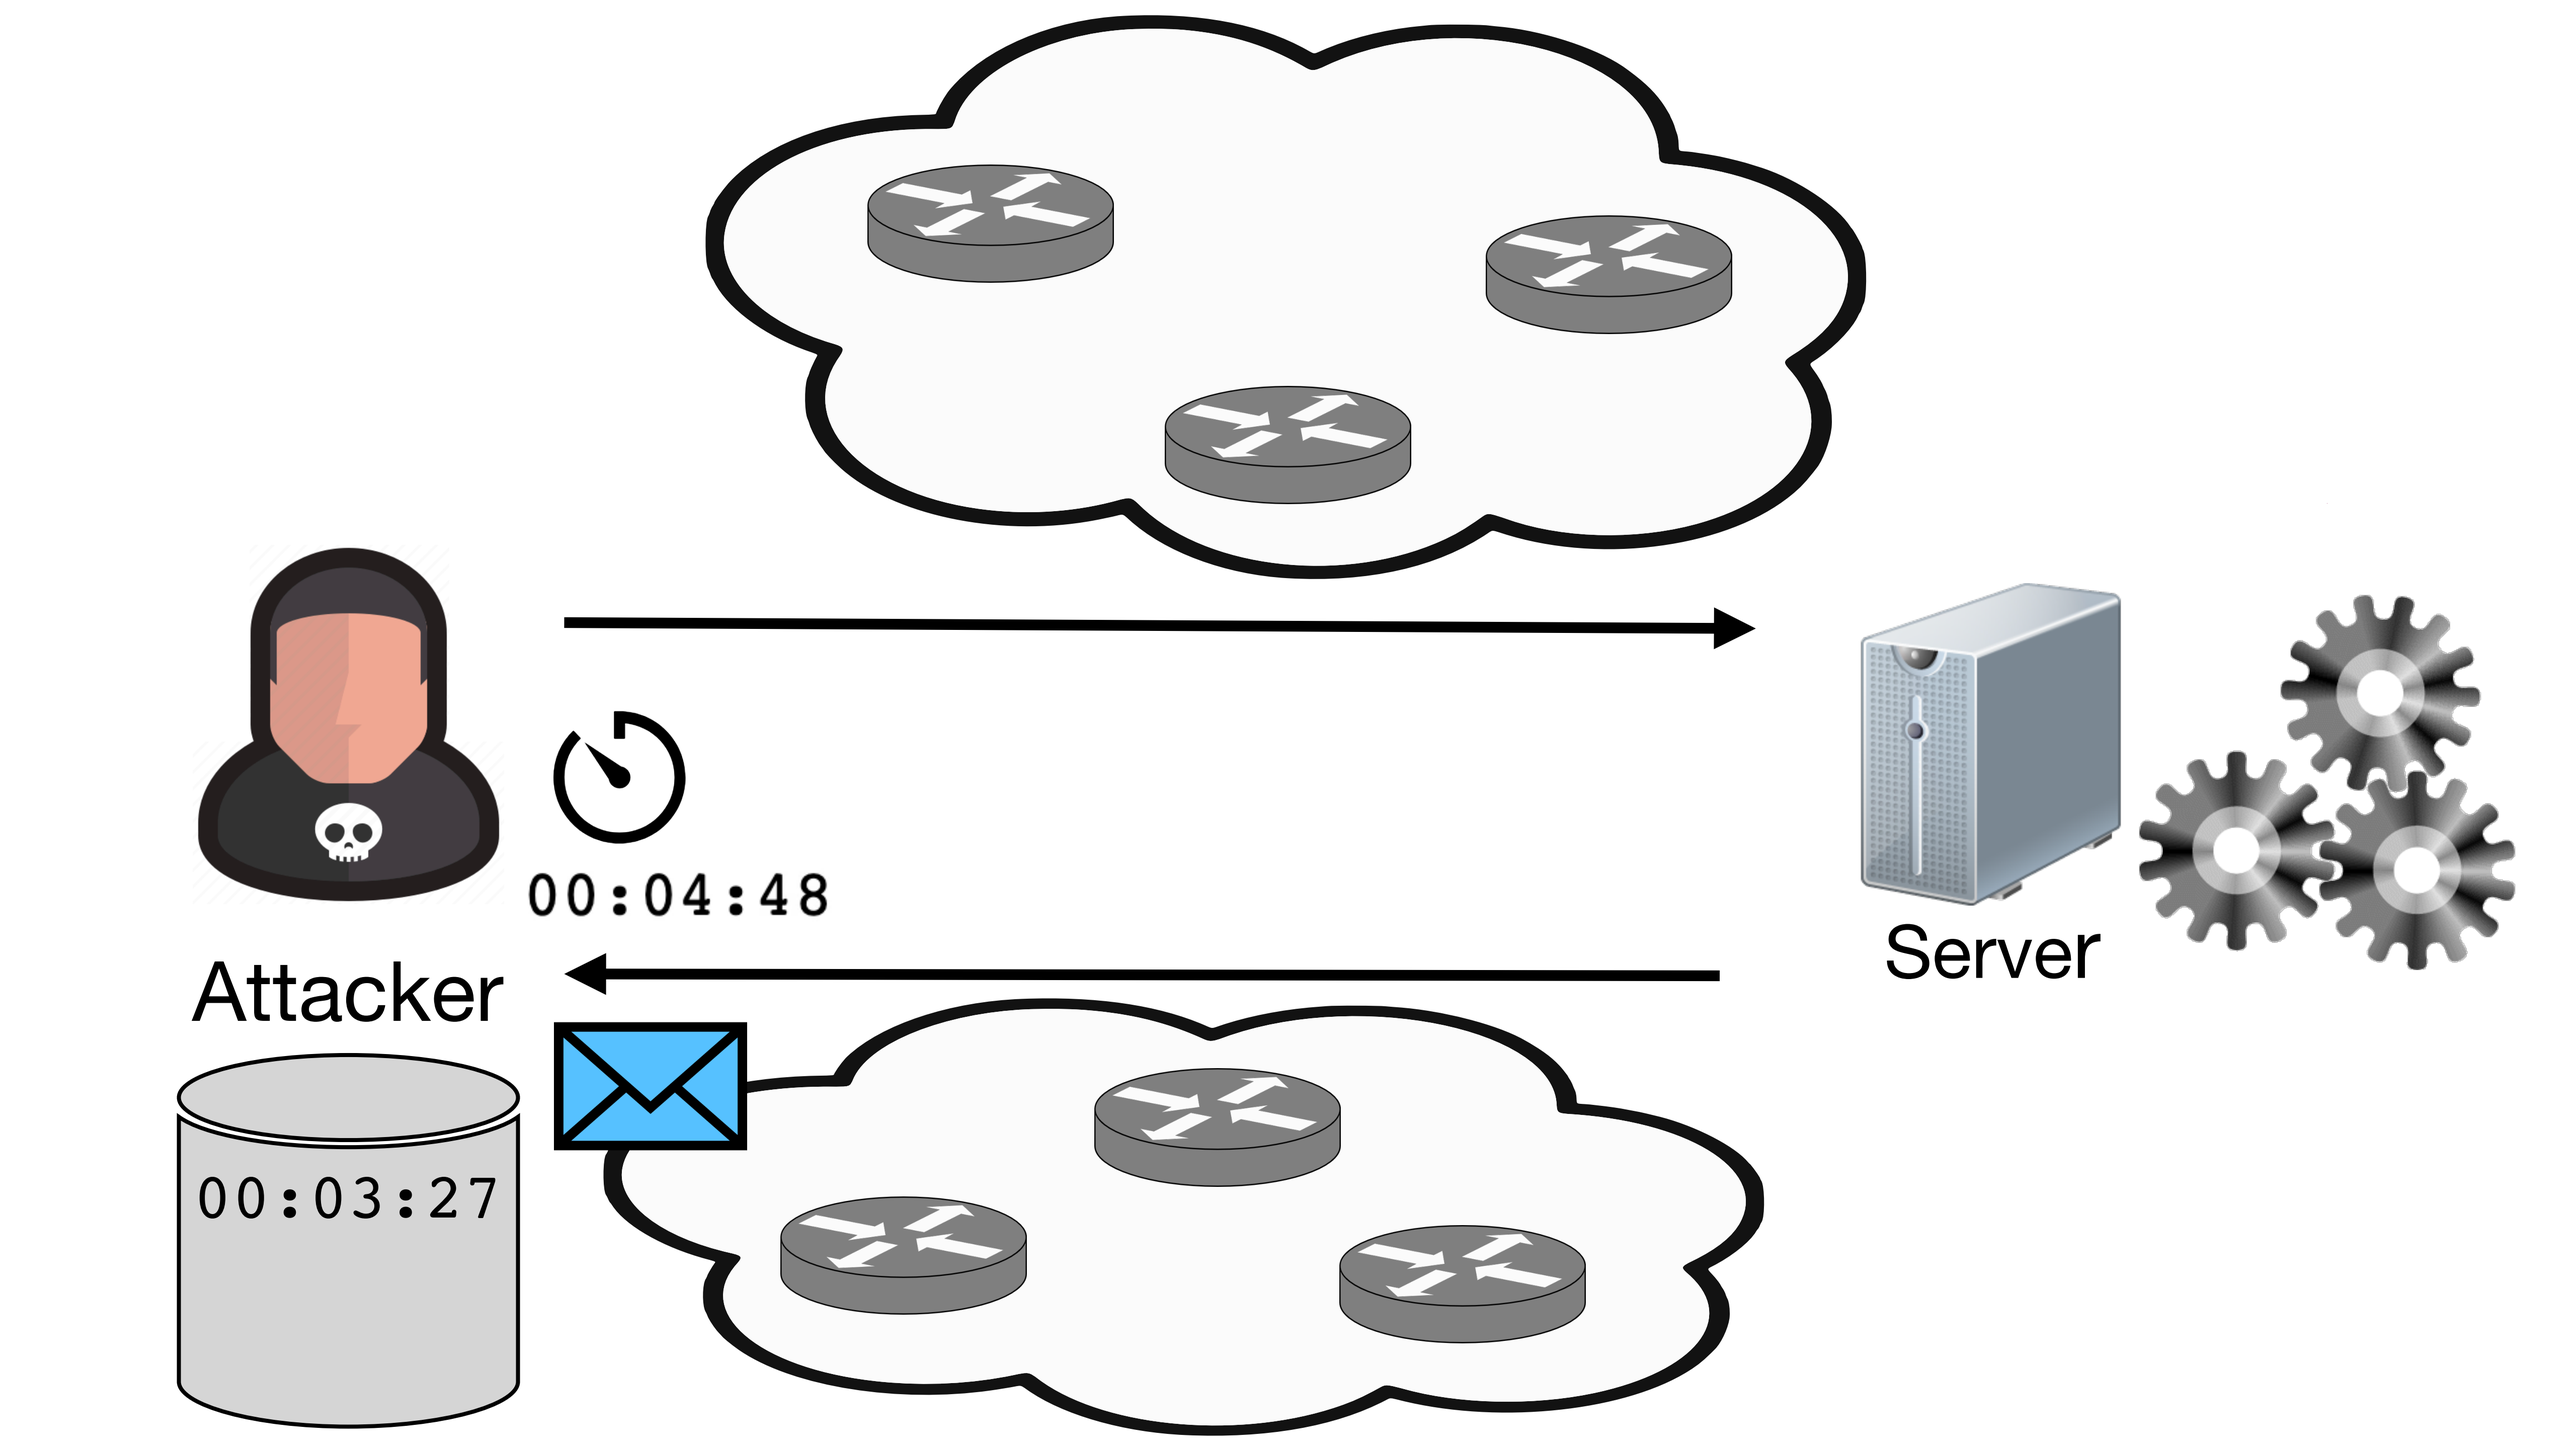
00:04:48
Server
Attacker
00:03:27

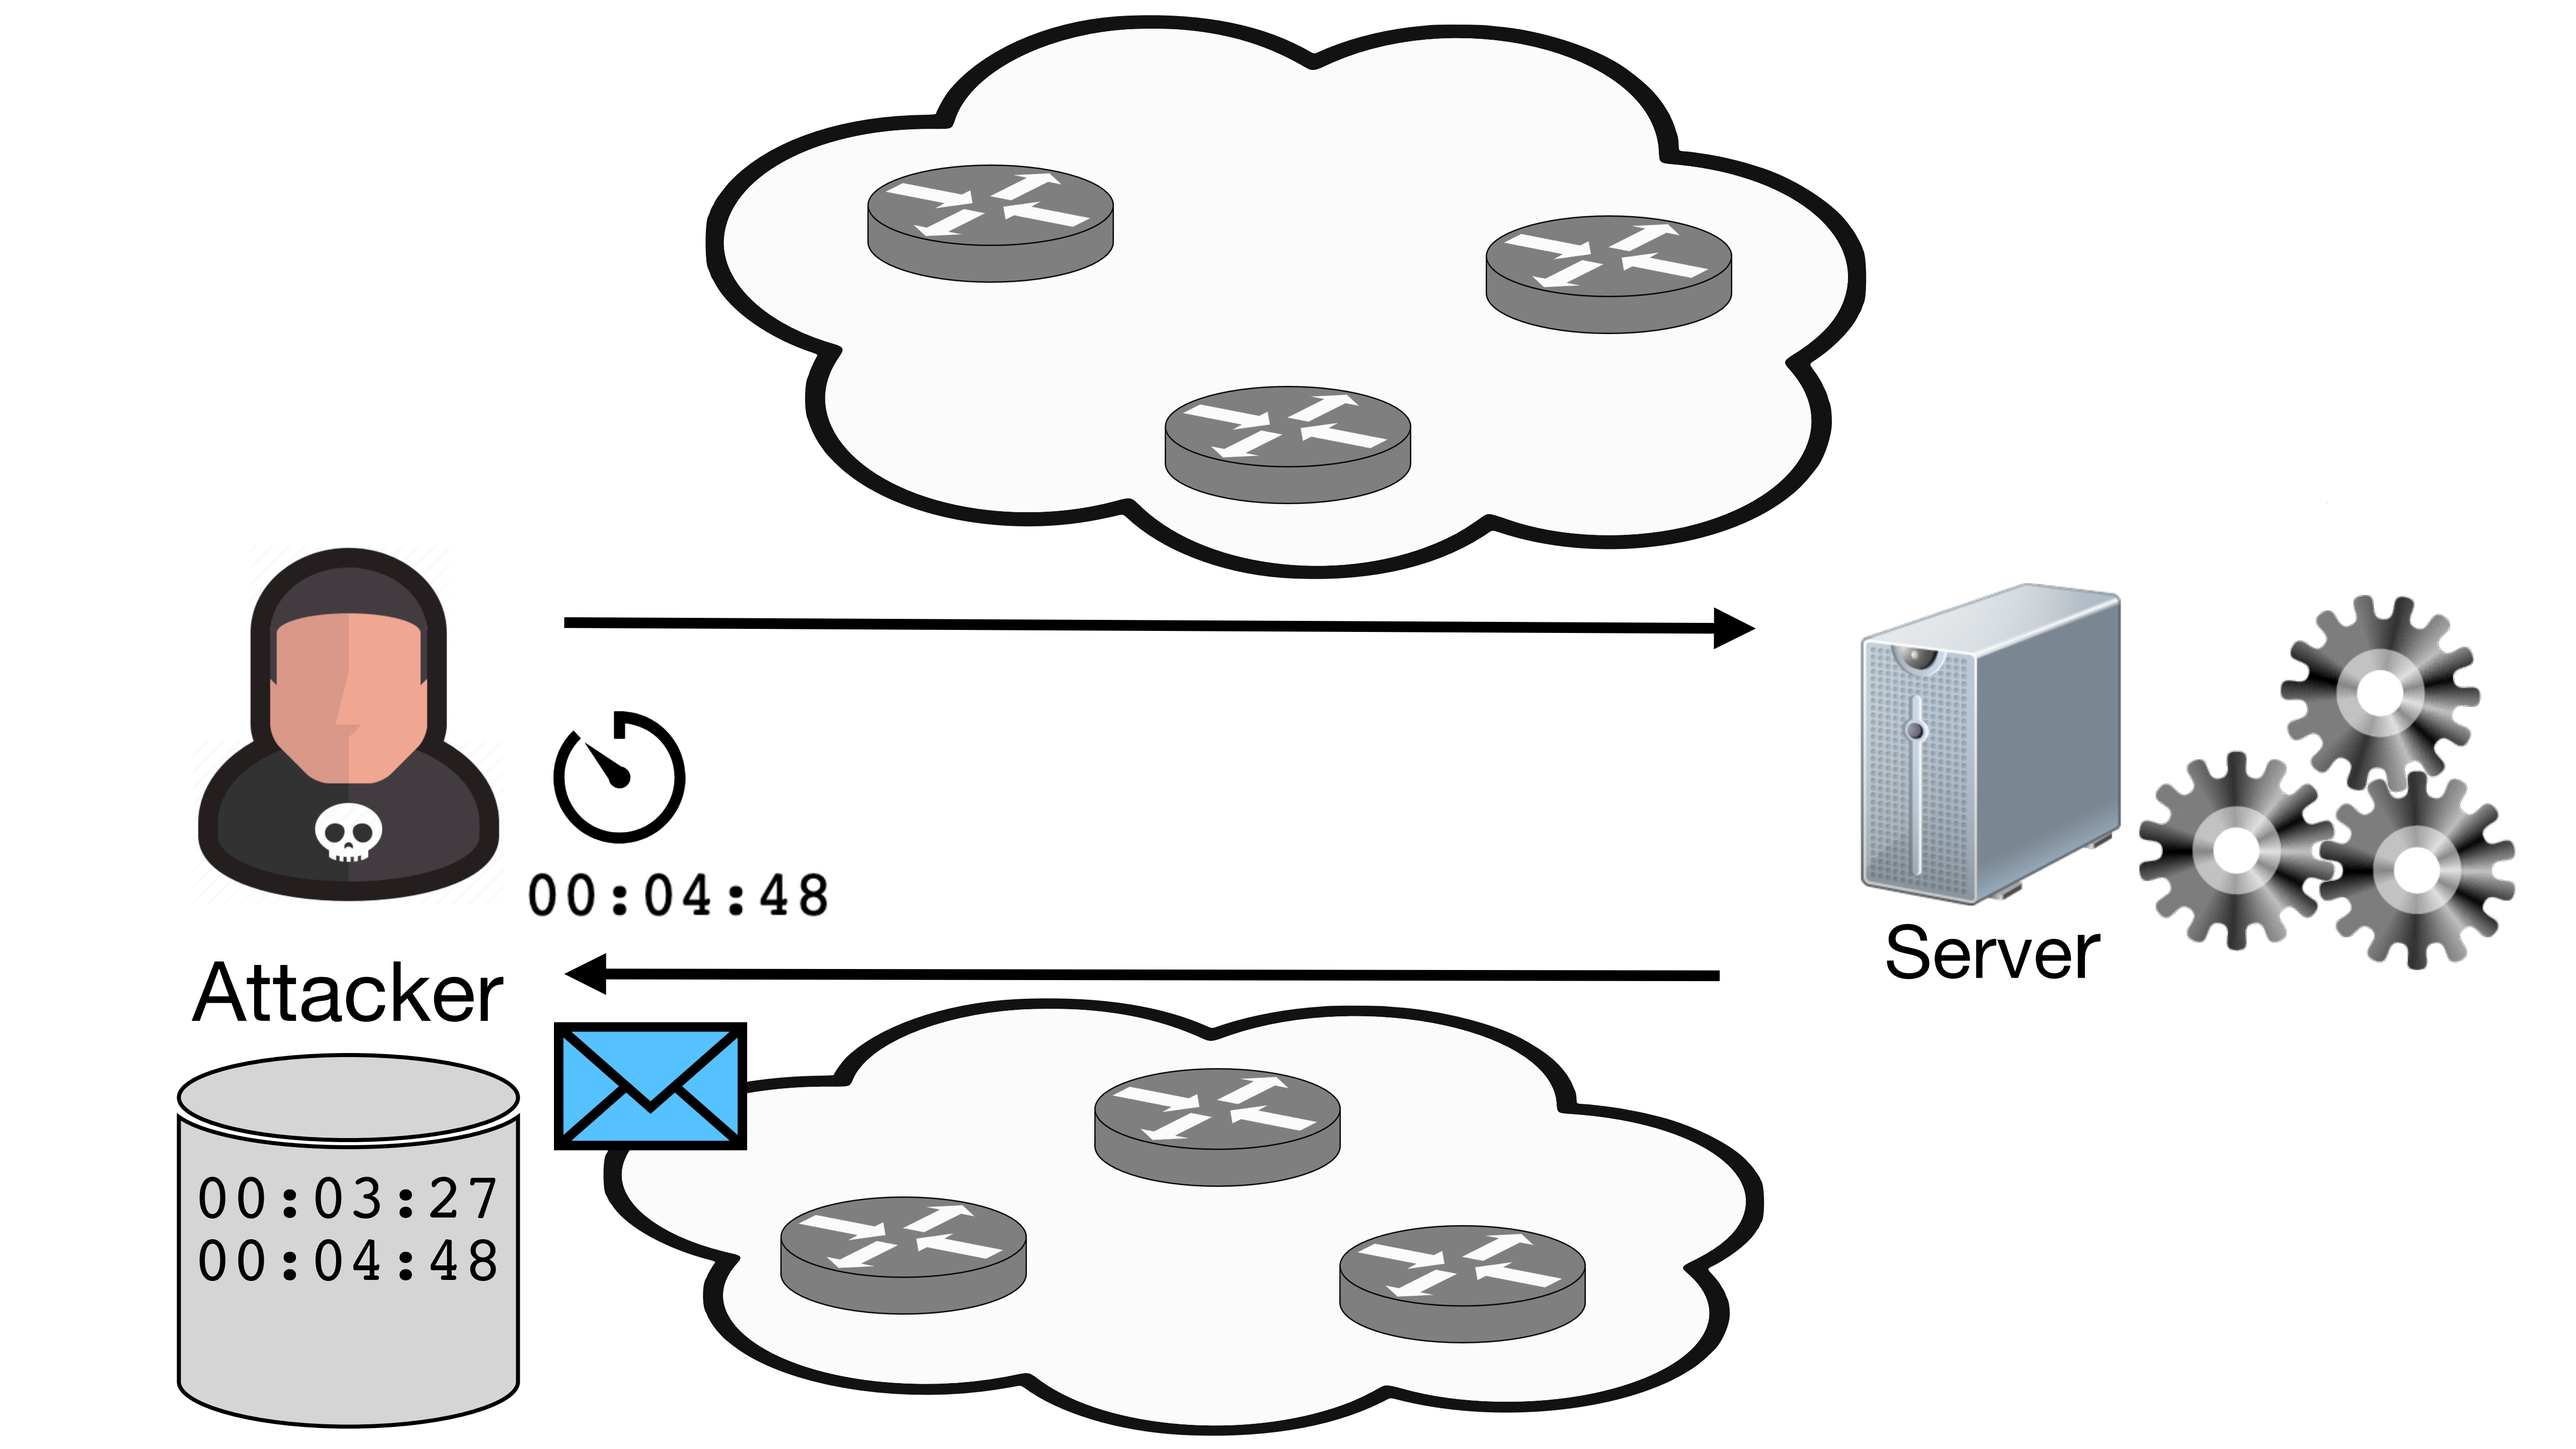
00:04:48
Server
Attacker
00:03:27
00:04:48

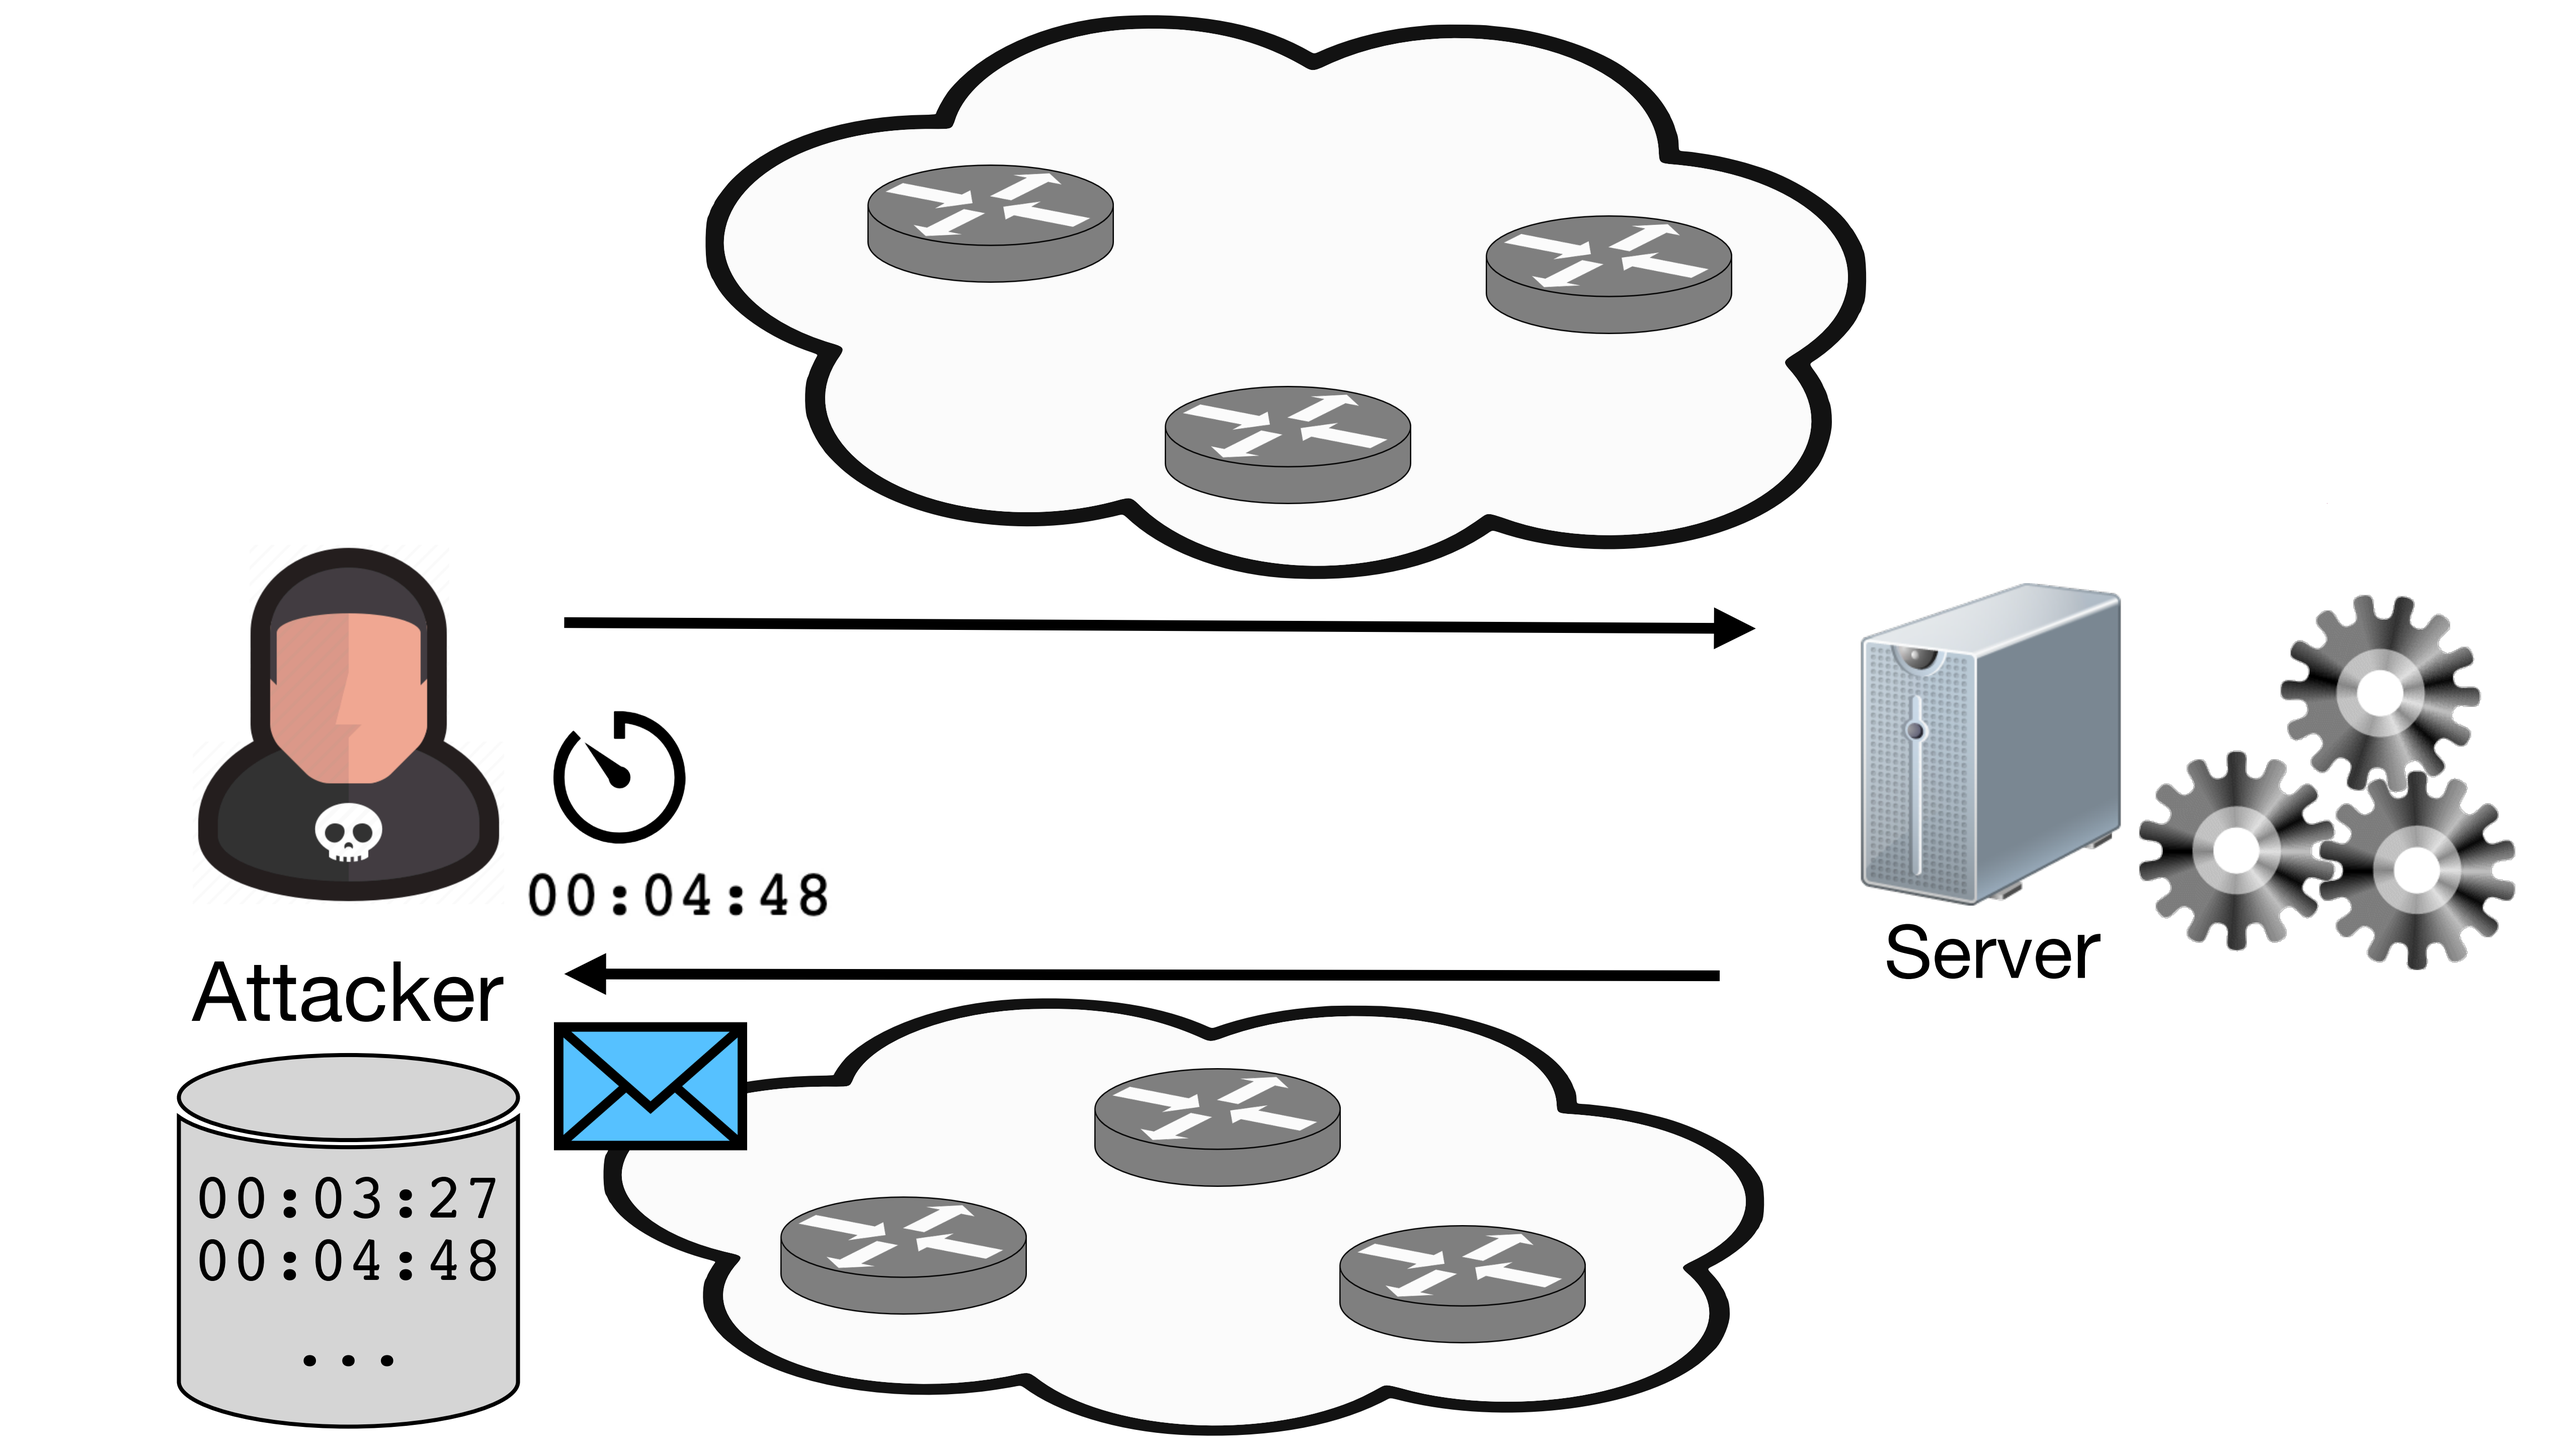
Server
00:04:48
Attacker
00:03:27
00:04:48
...

	EU	US	Asia
50µs	333	4,492	7,386
20µs	2,926	16,820	-
10µs	23,220	-	-
5µs	-	-	-

Number of requests required to determine timing difference (5-50µs) with 95% accuracy

based on measurements between university network and AWS
imposed maximum: 100,000

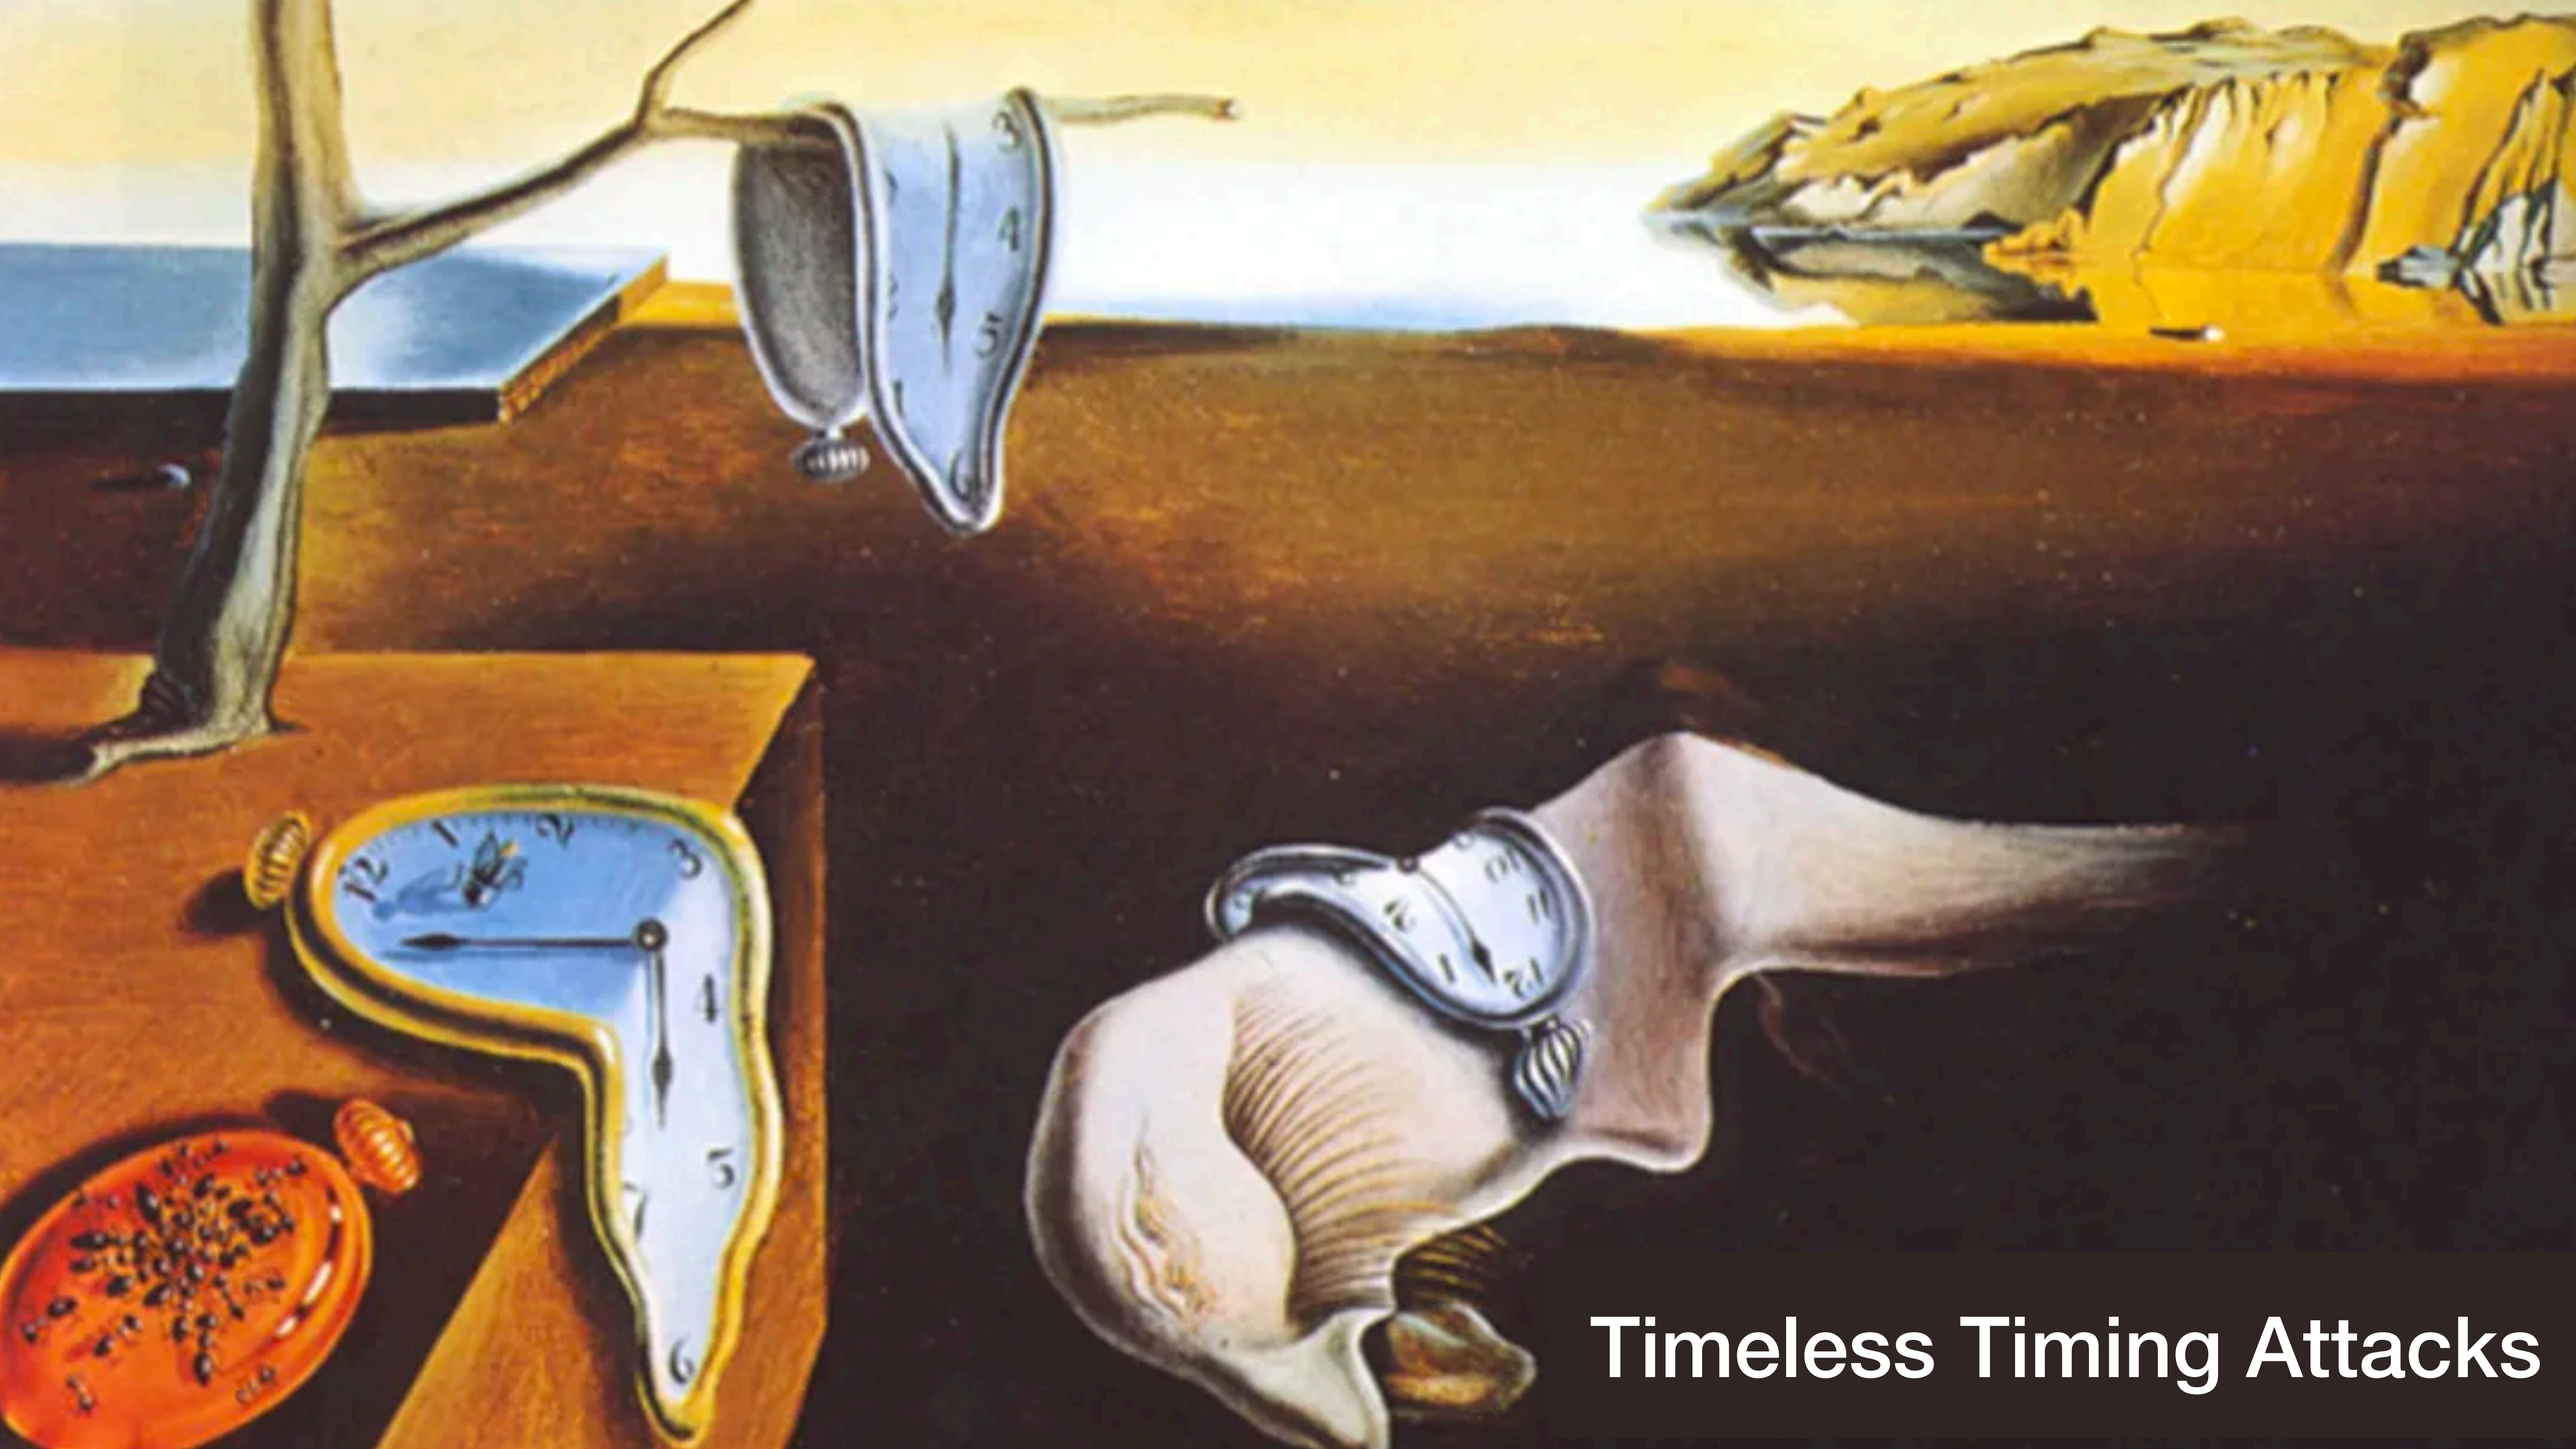
Timeless Timing Attacks

Timeless Timing Attacks

- Absolute response timing is unreliable, as it will always include jitter for every request
- Let’s get rid of the notion of time (hence timeless)
- Instead of relying on sequential timing measurements,
  we **introduce concurrency** and only consider response order
  => no absolute timing measurements
- Timeless timing attacks are unaffected by network jitter

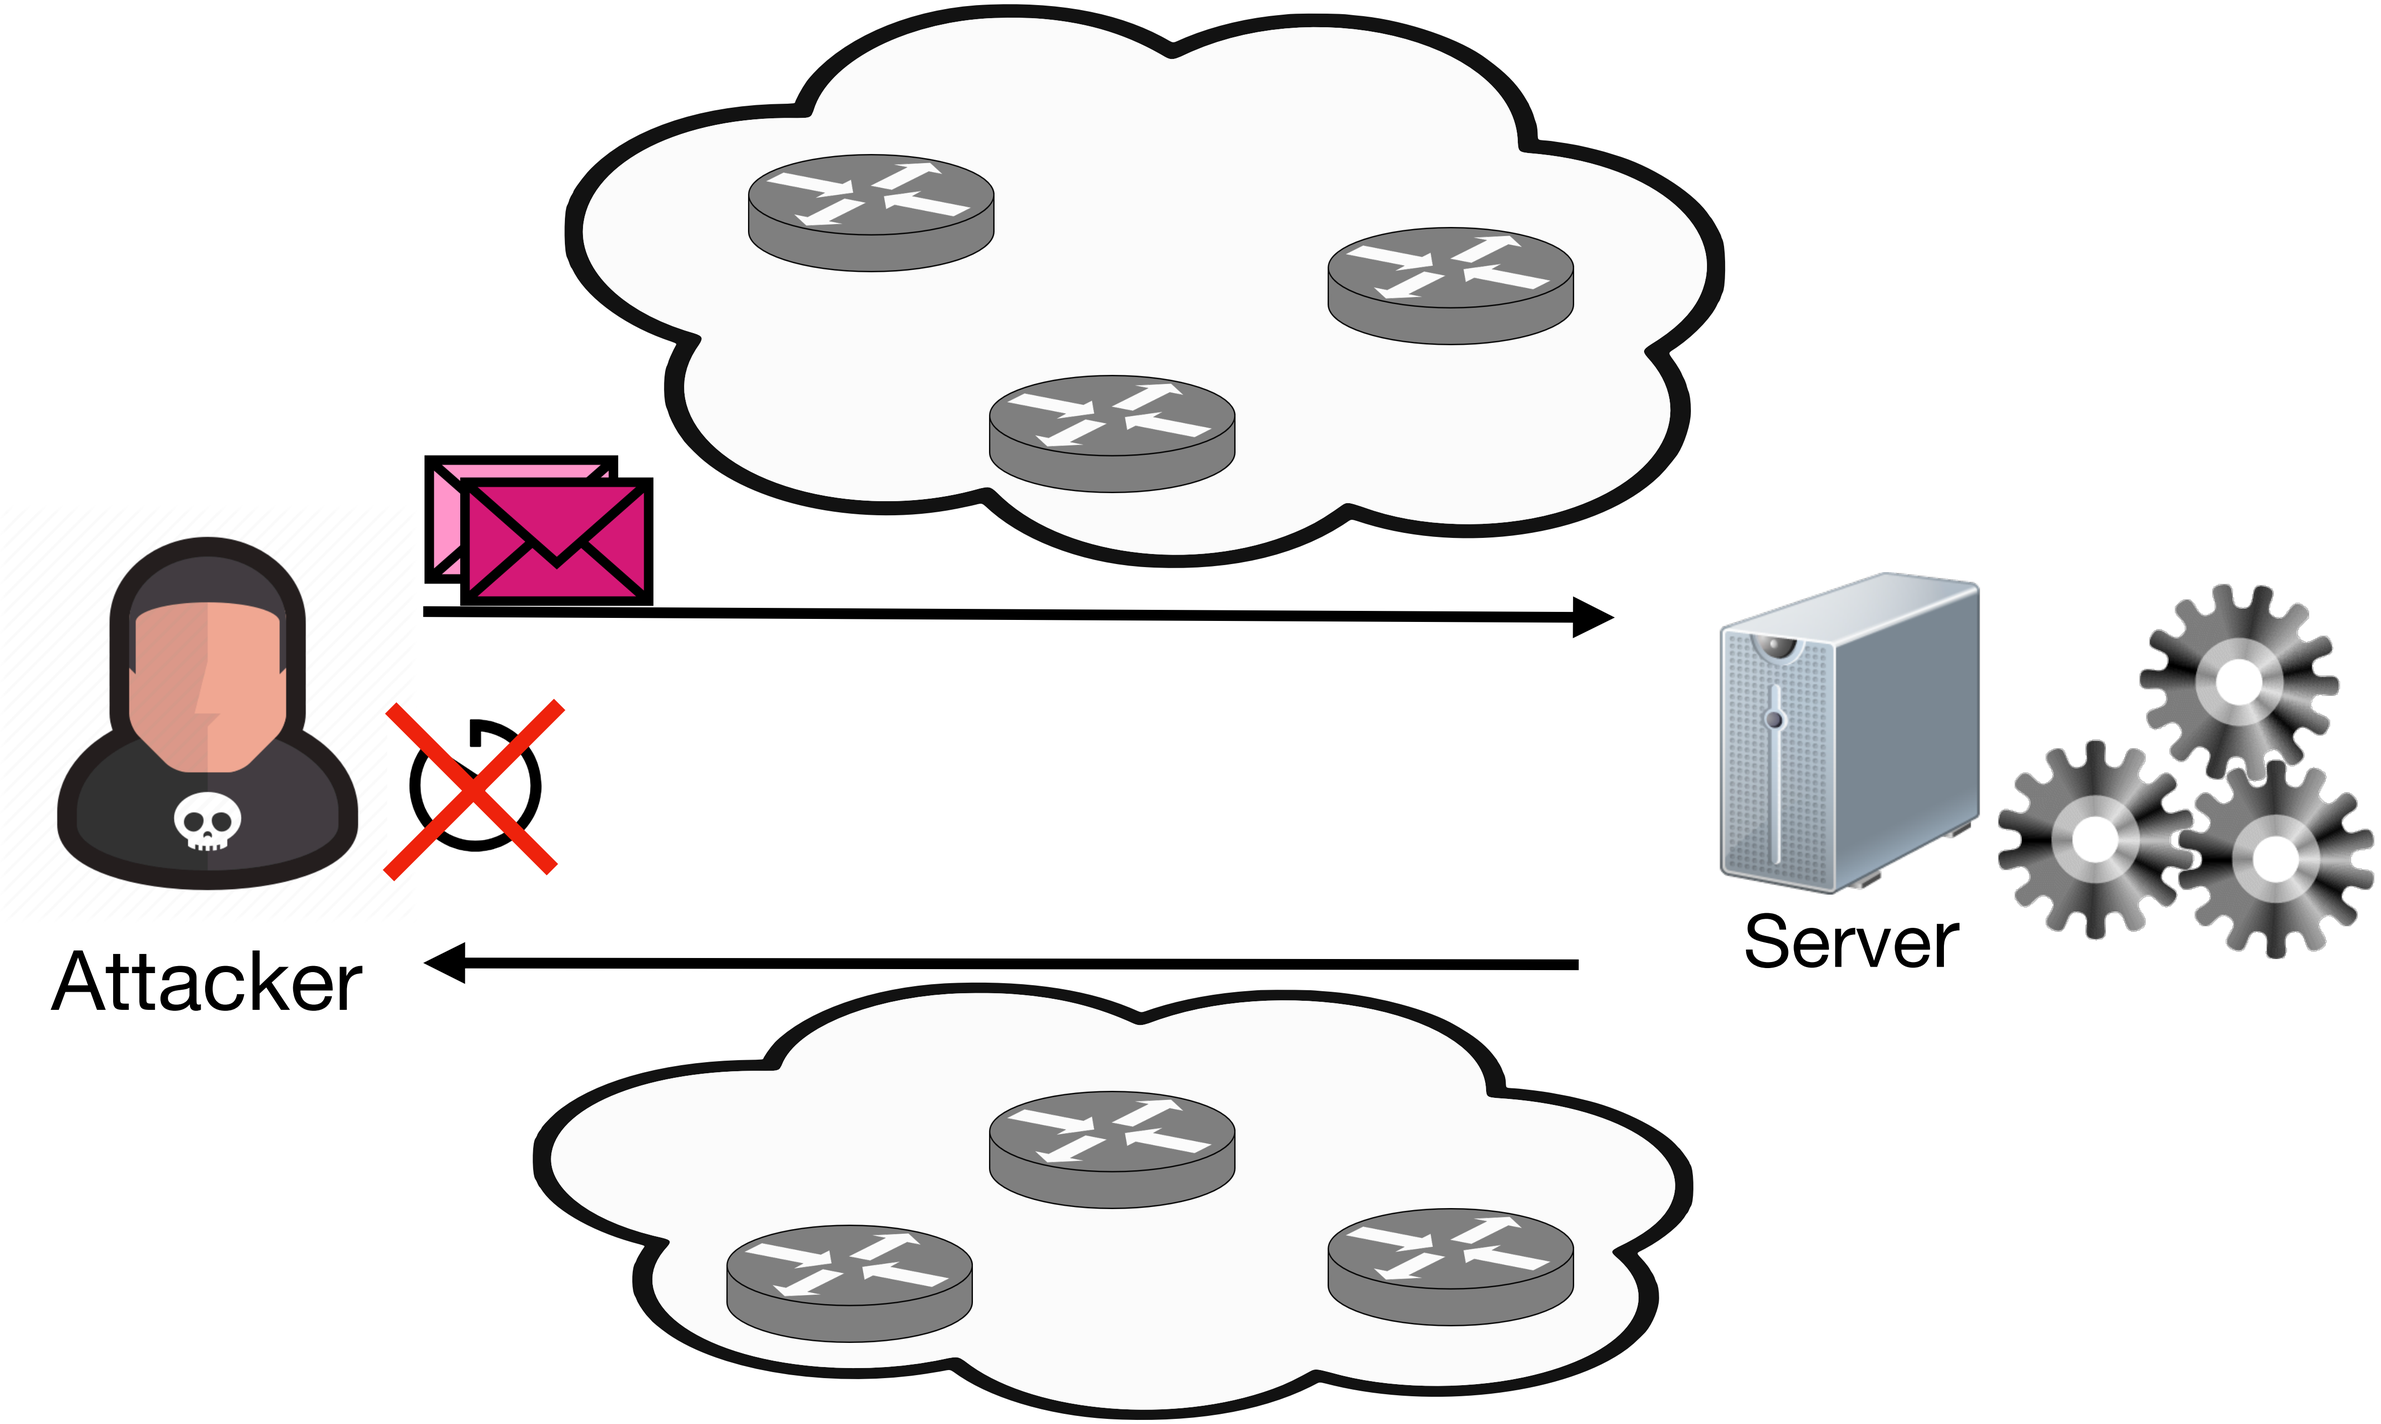
Attacker
Server

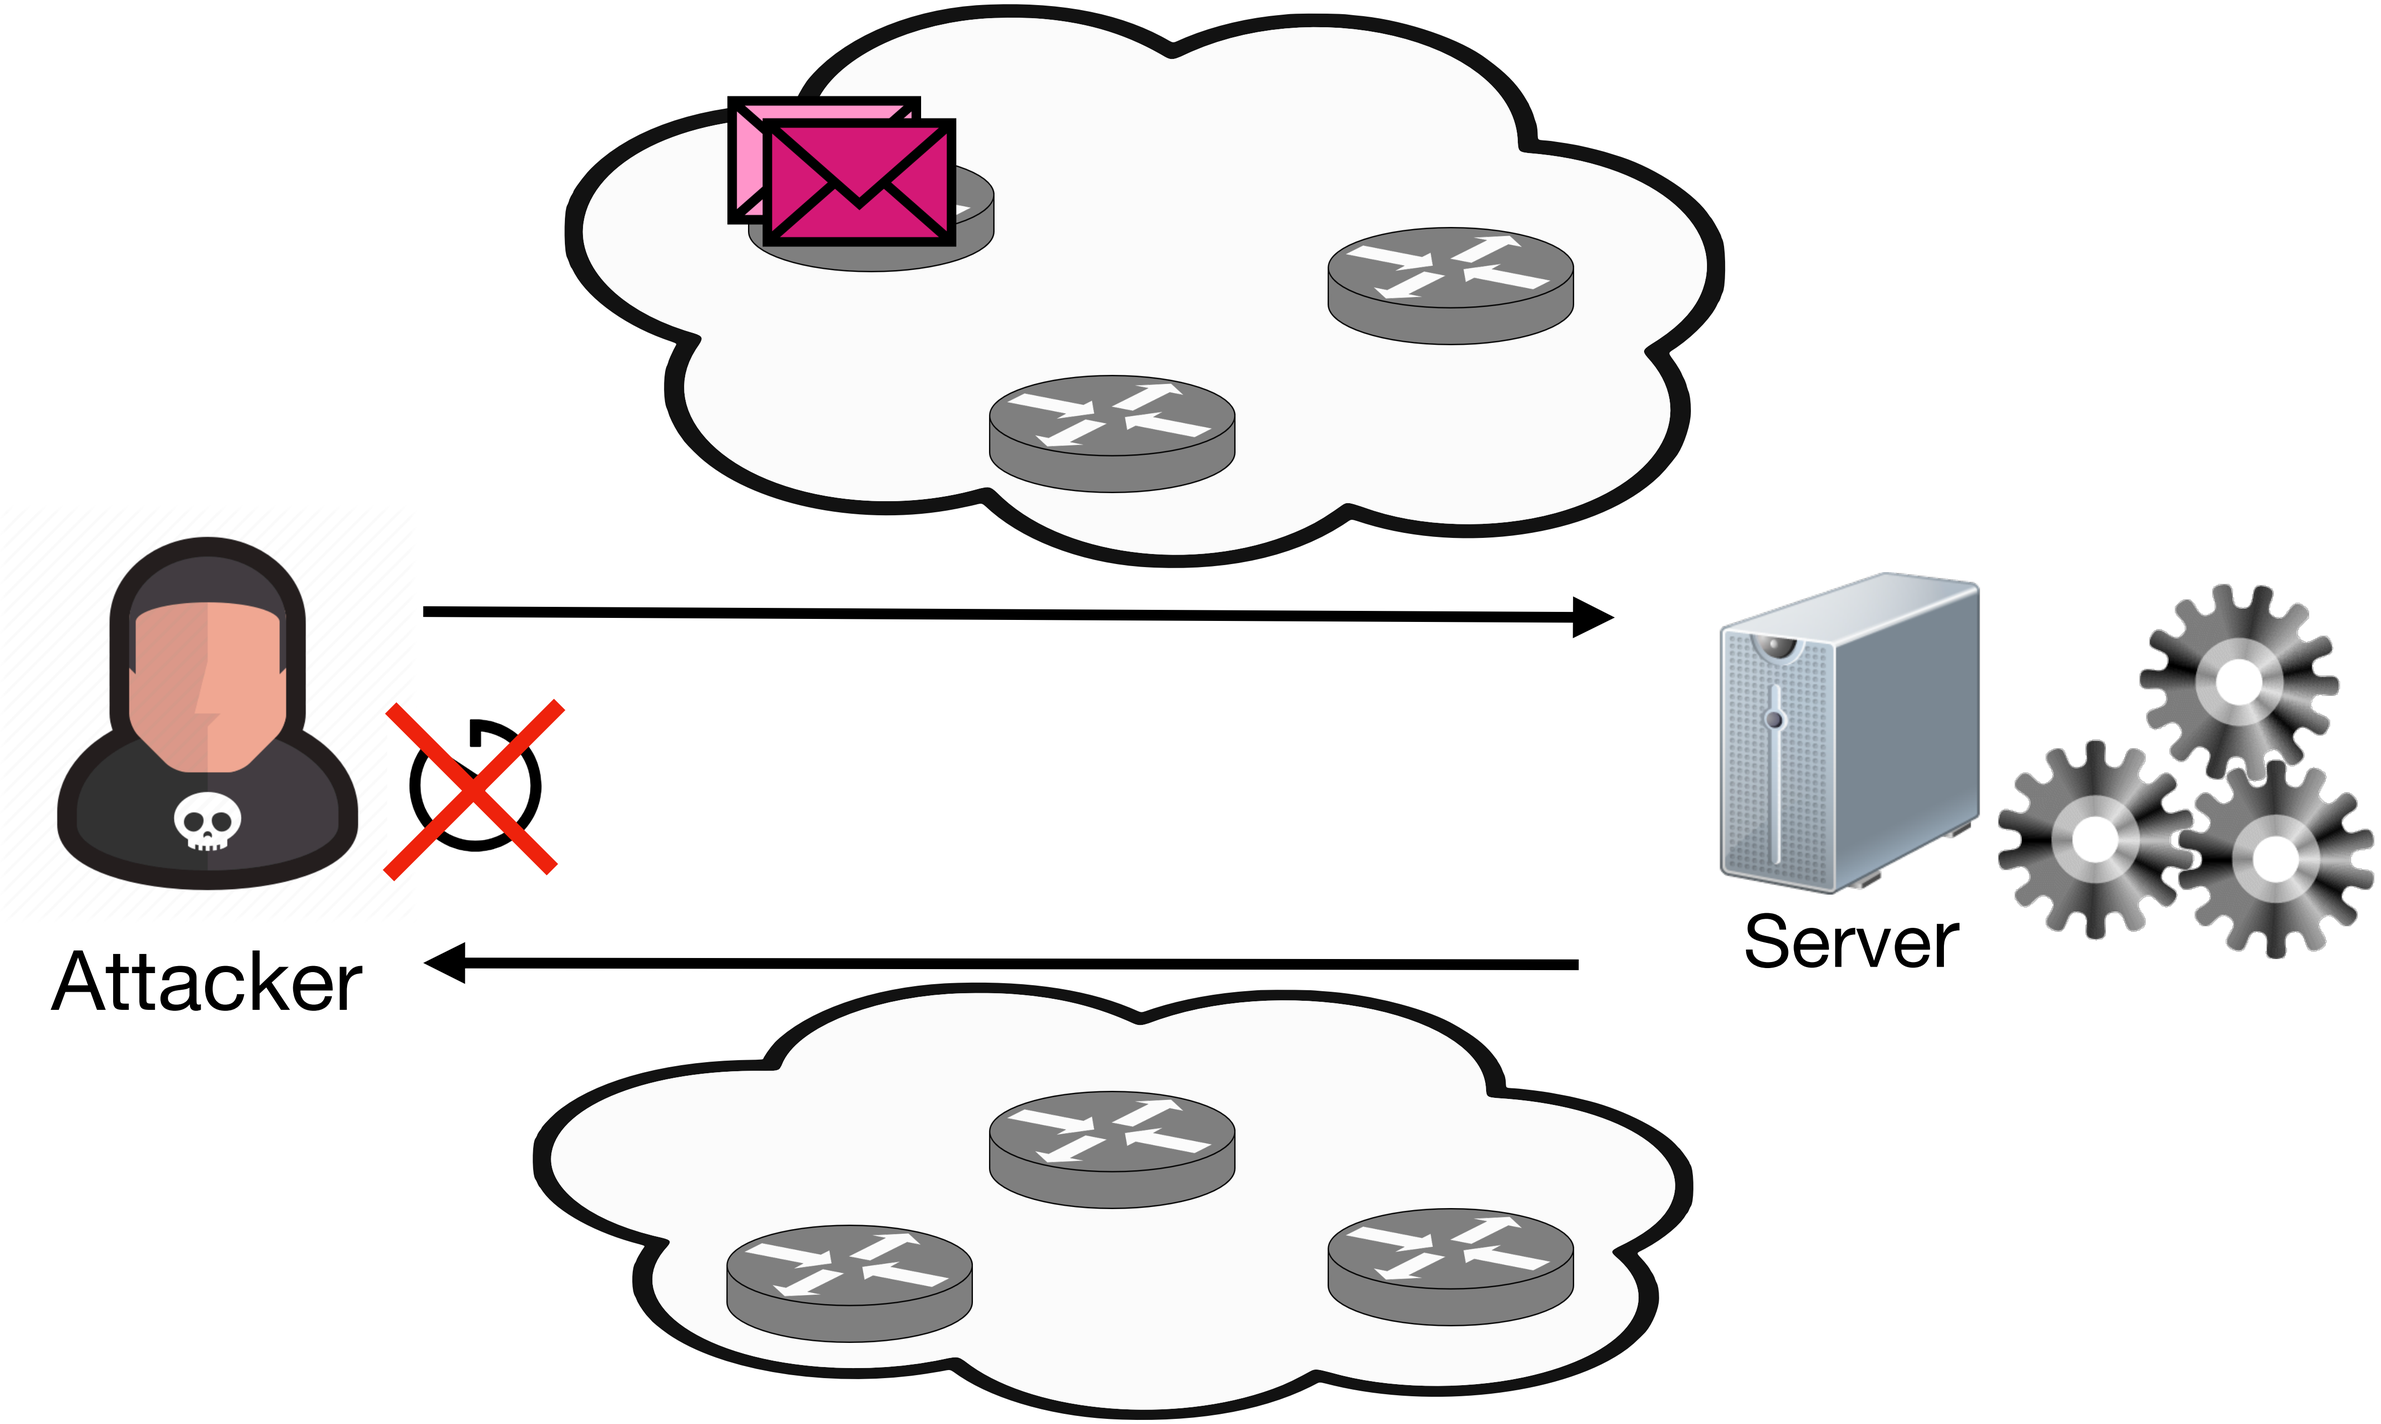
Attacker
Server

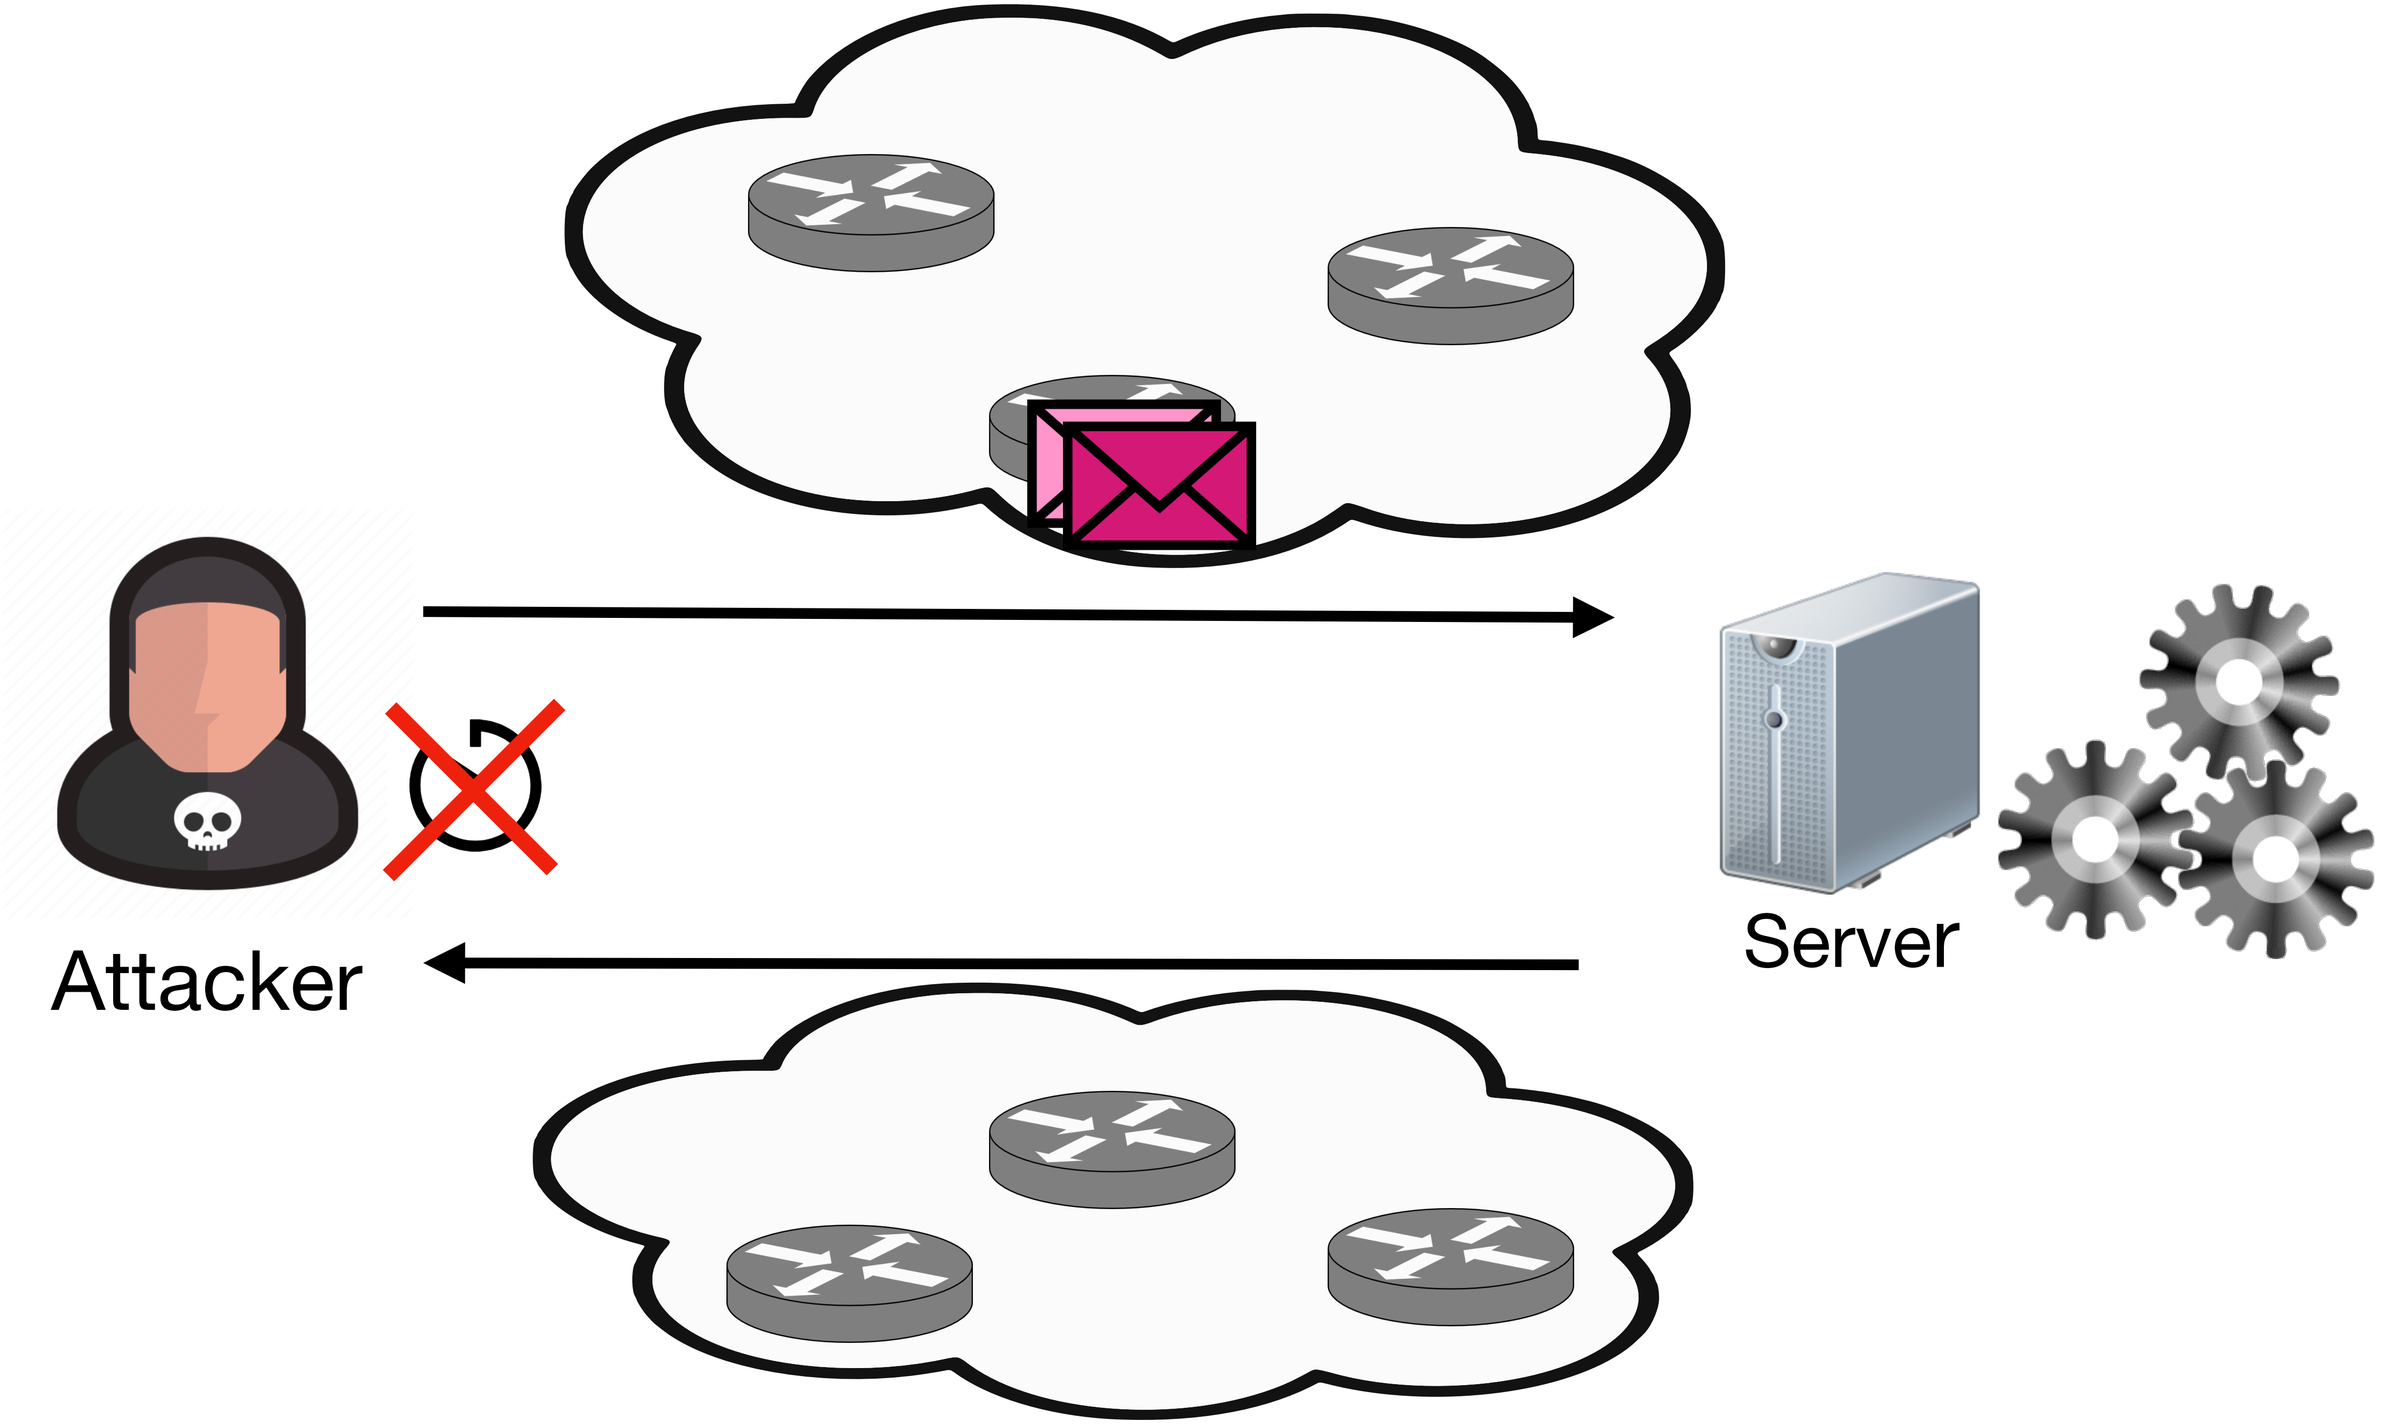
Attacker
Server

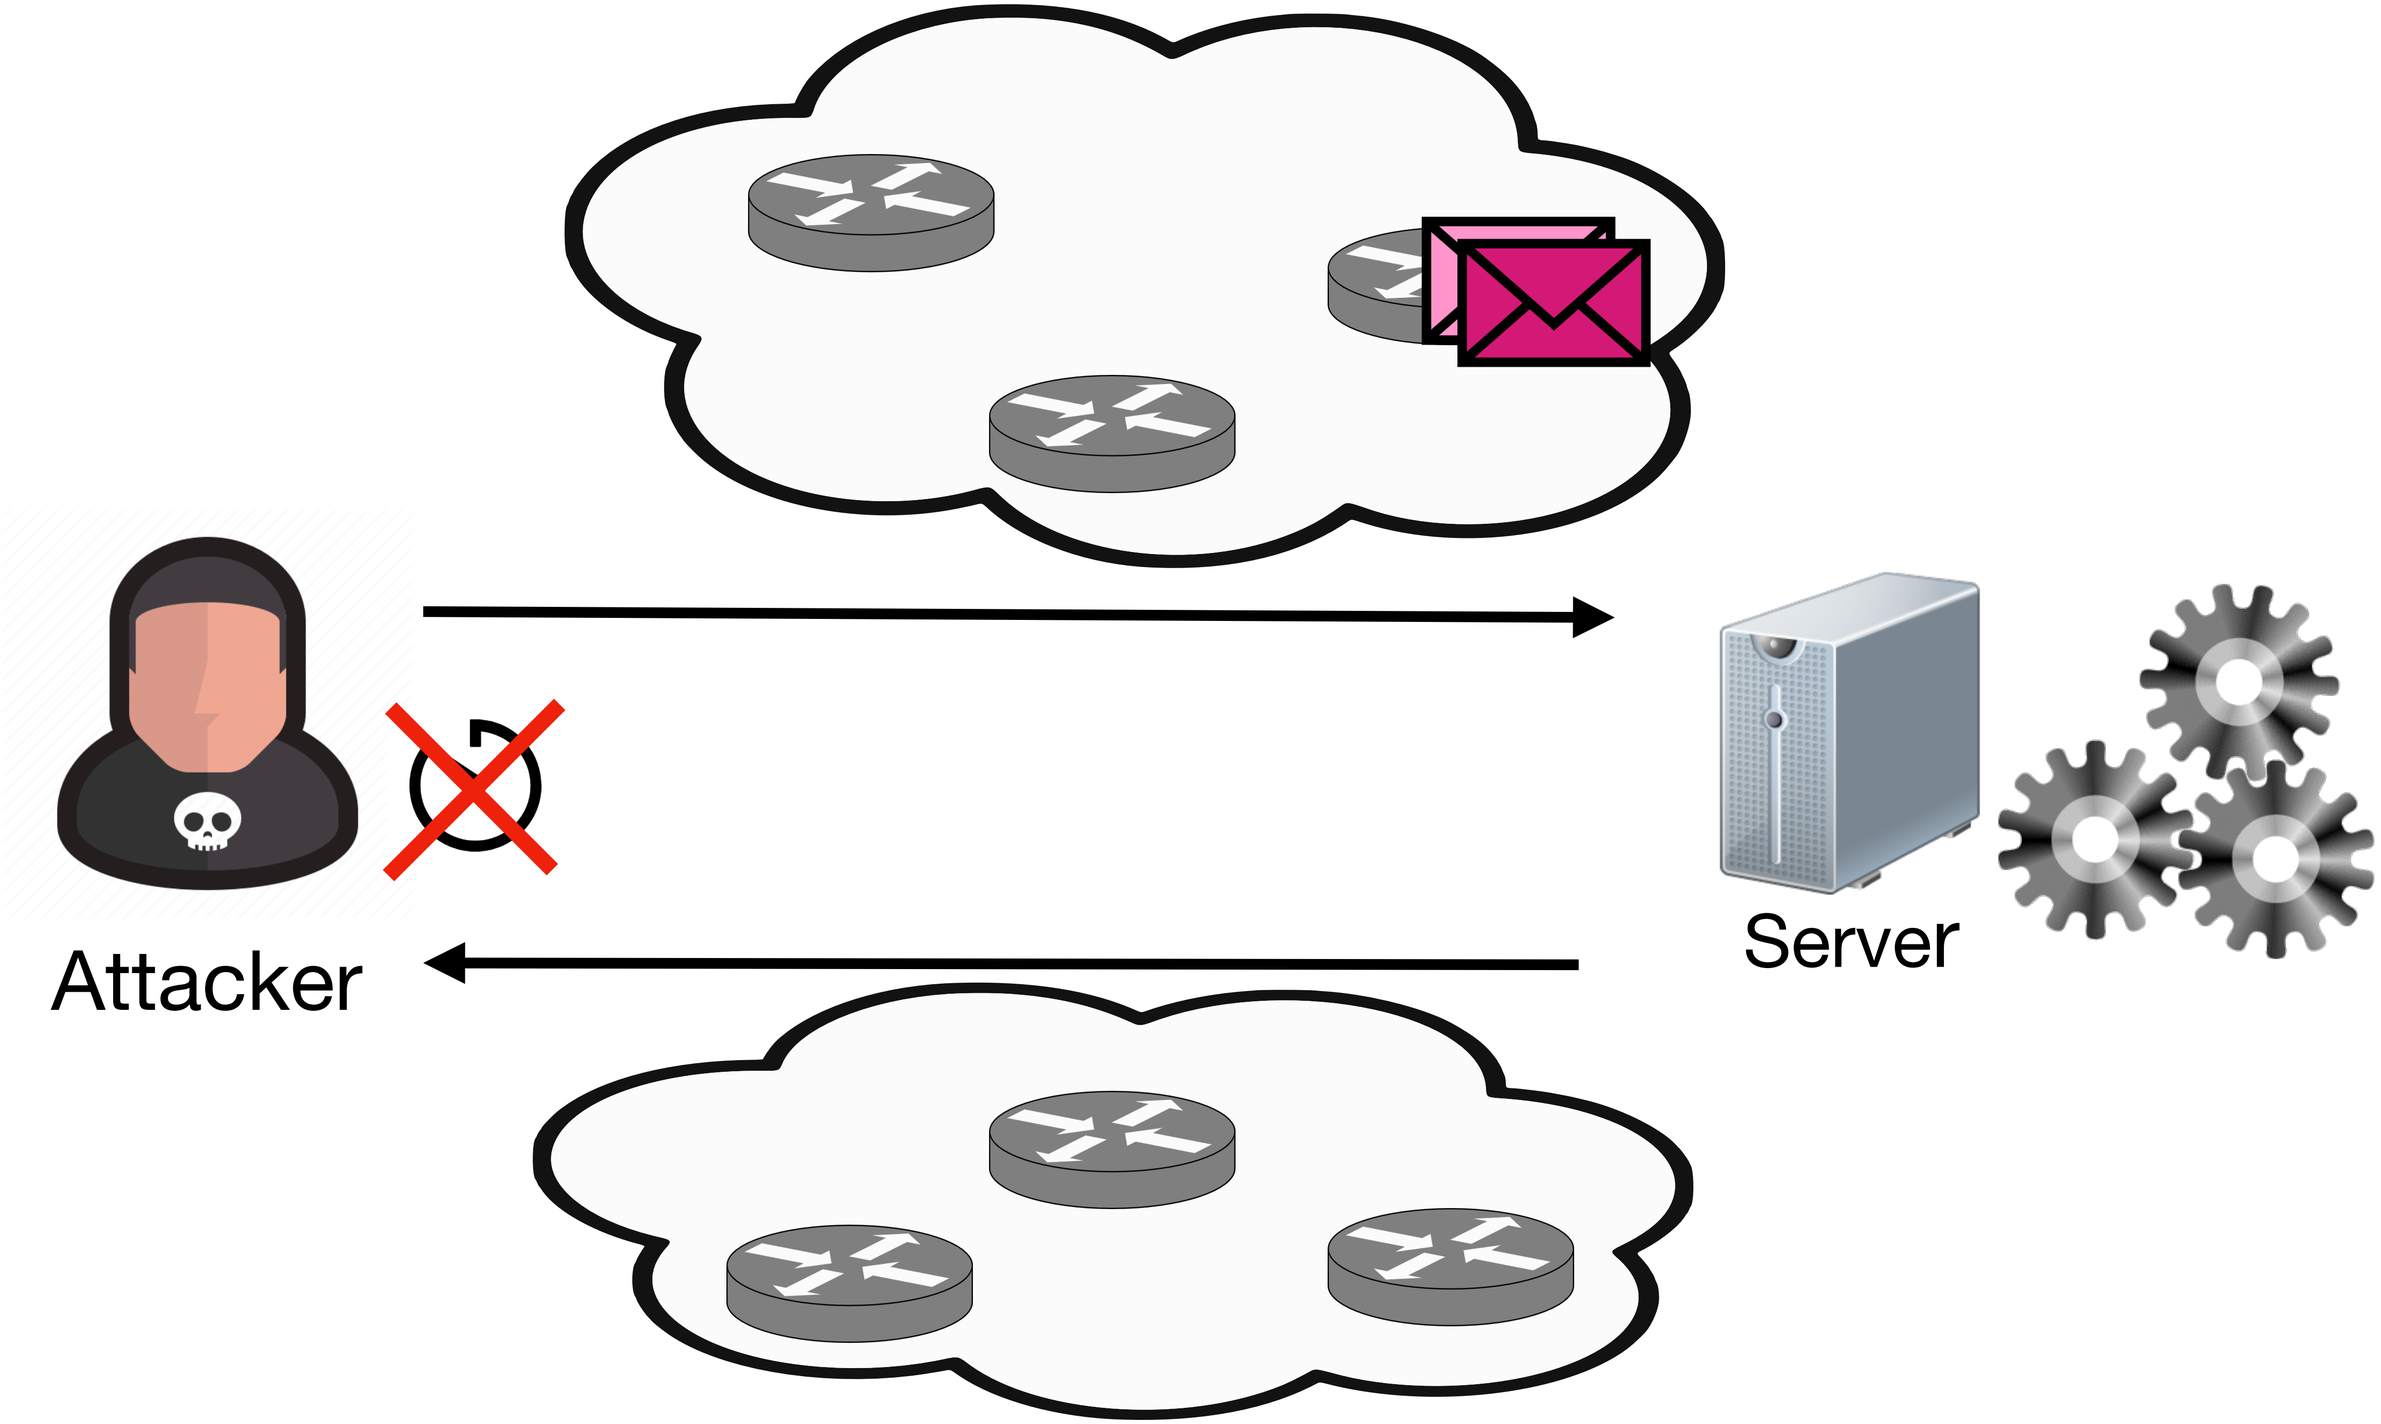
Attacker
Server

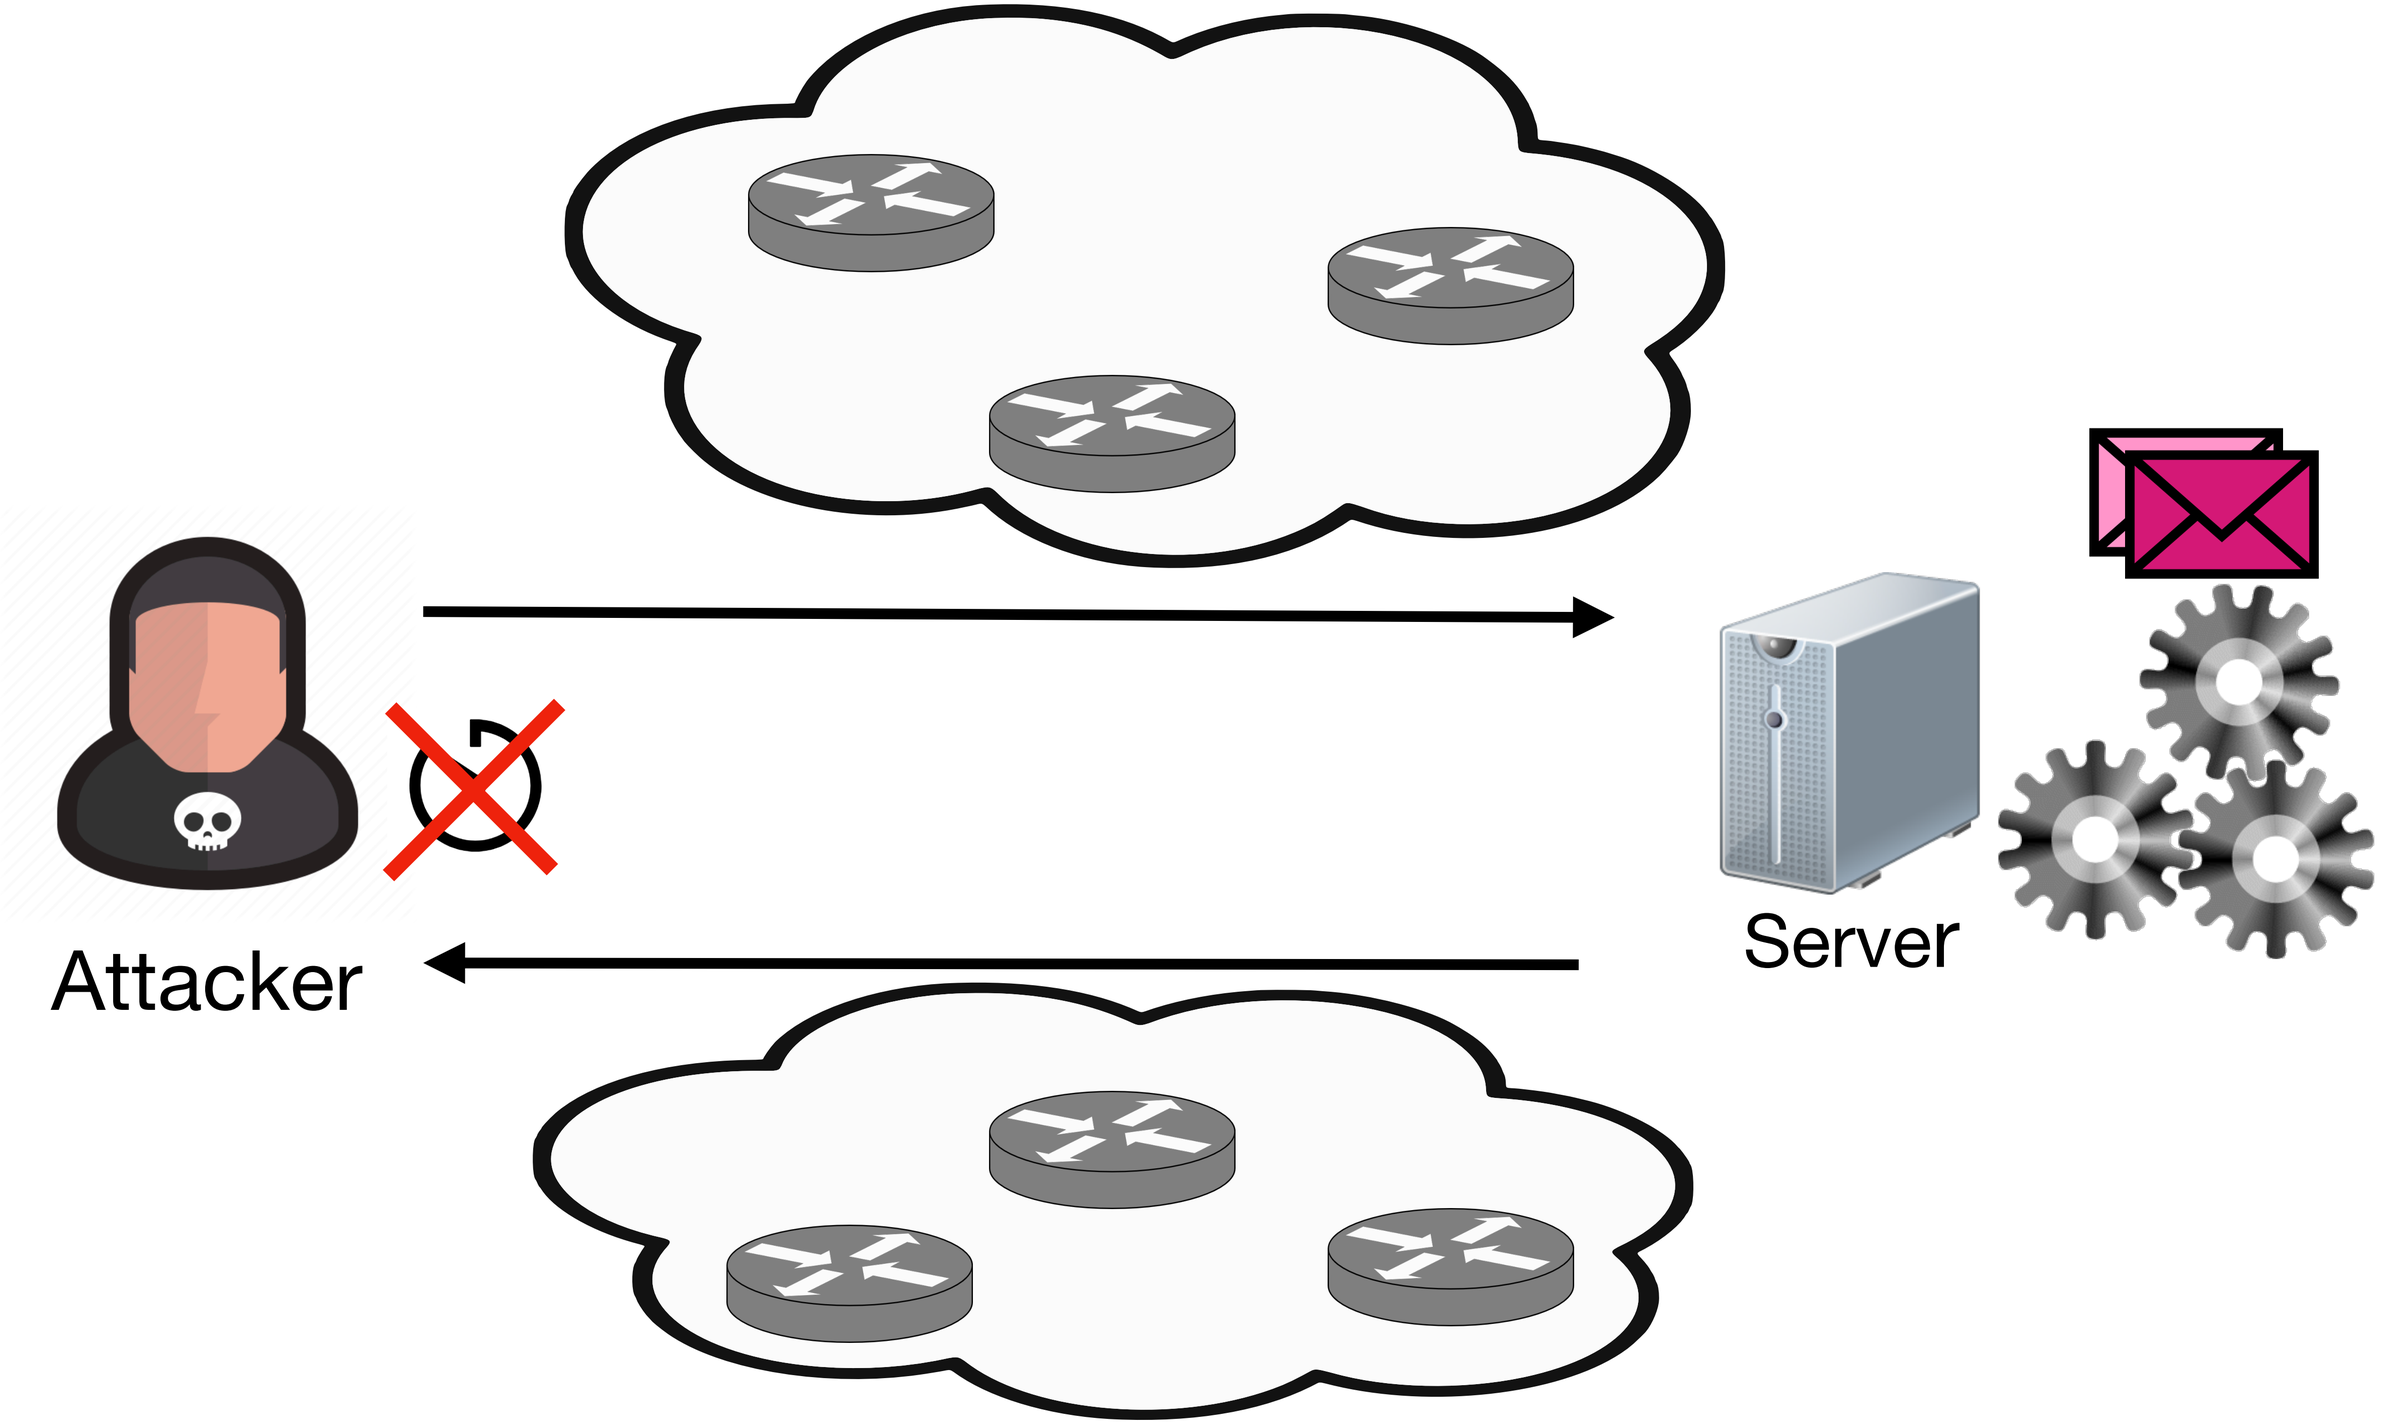
Attacker
Server

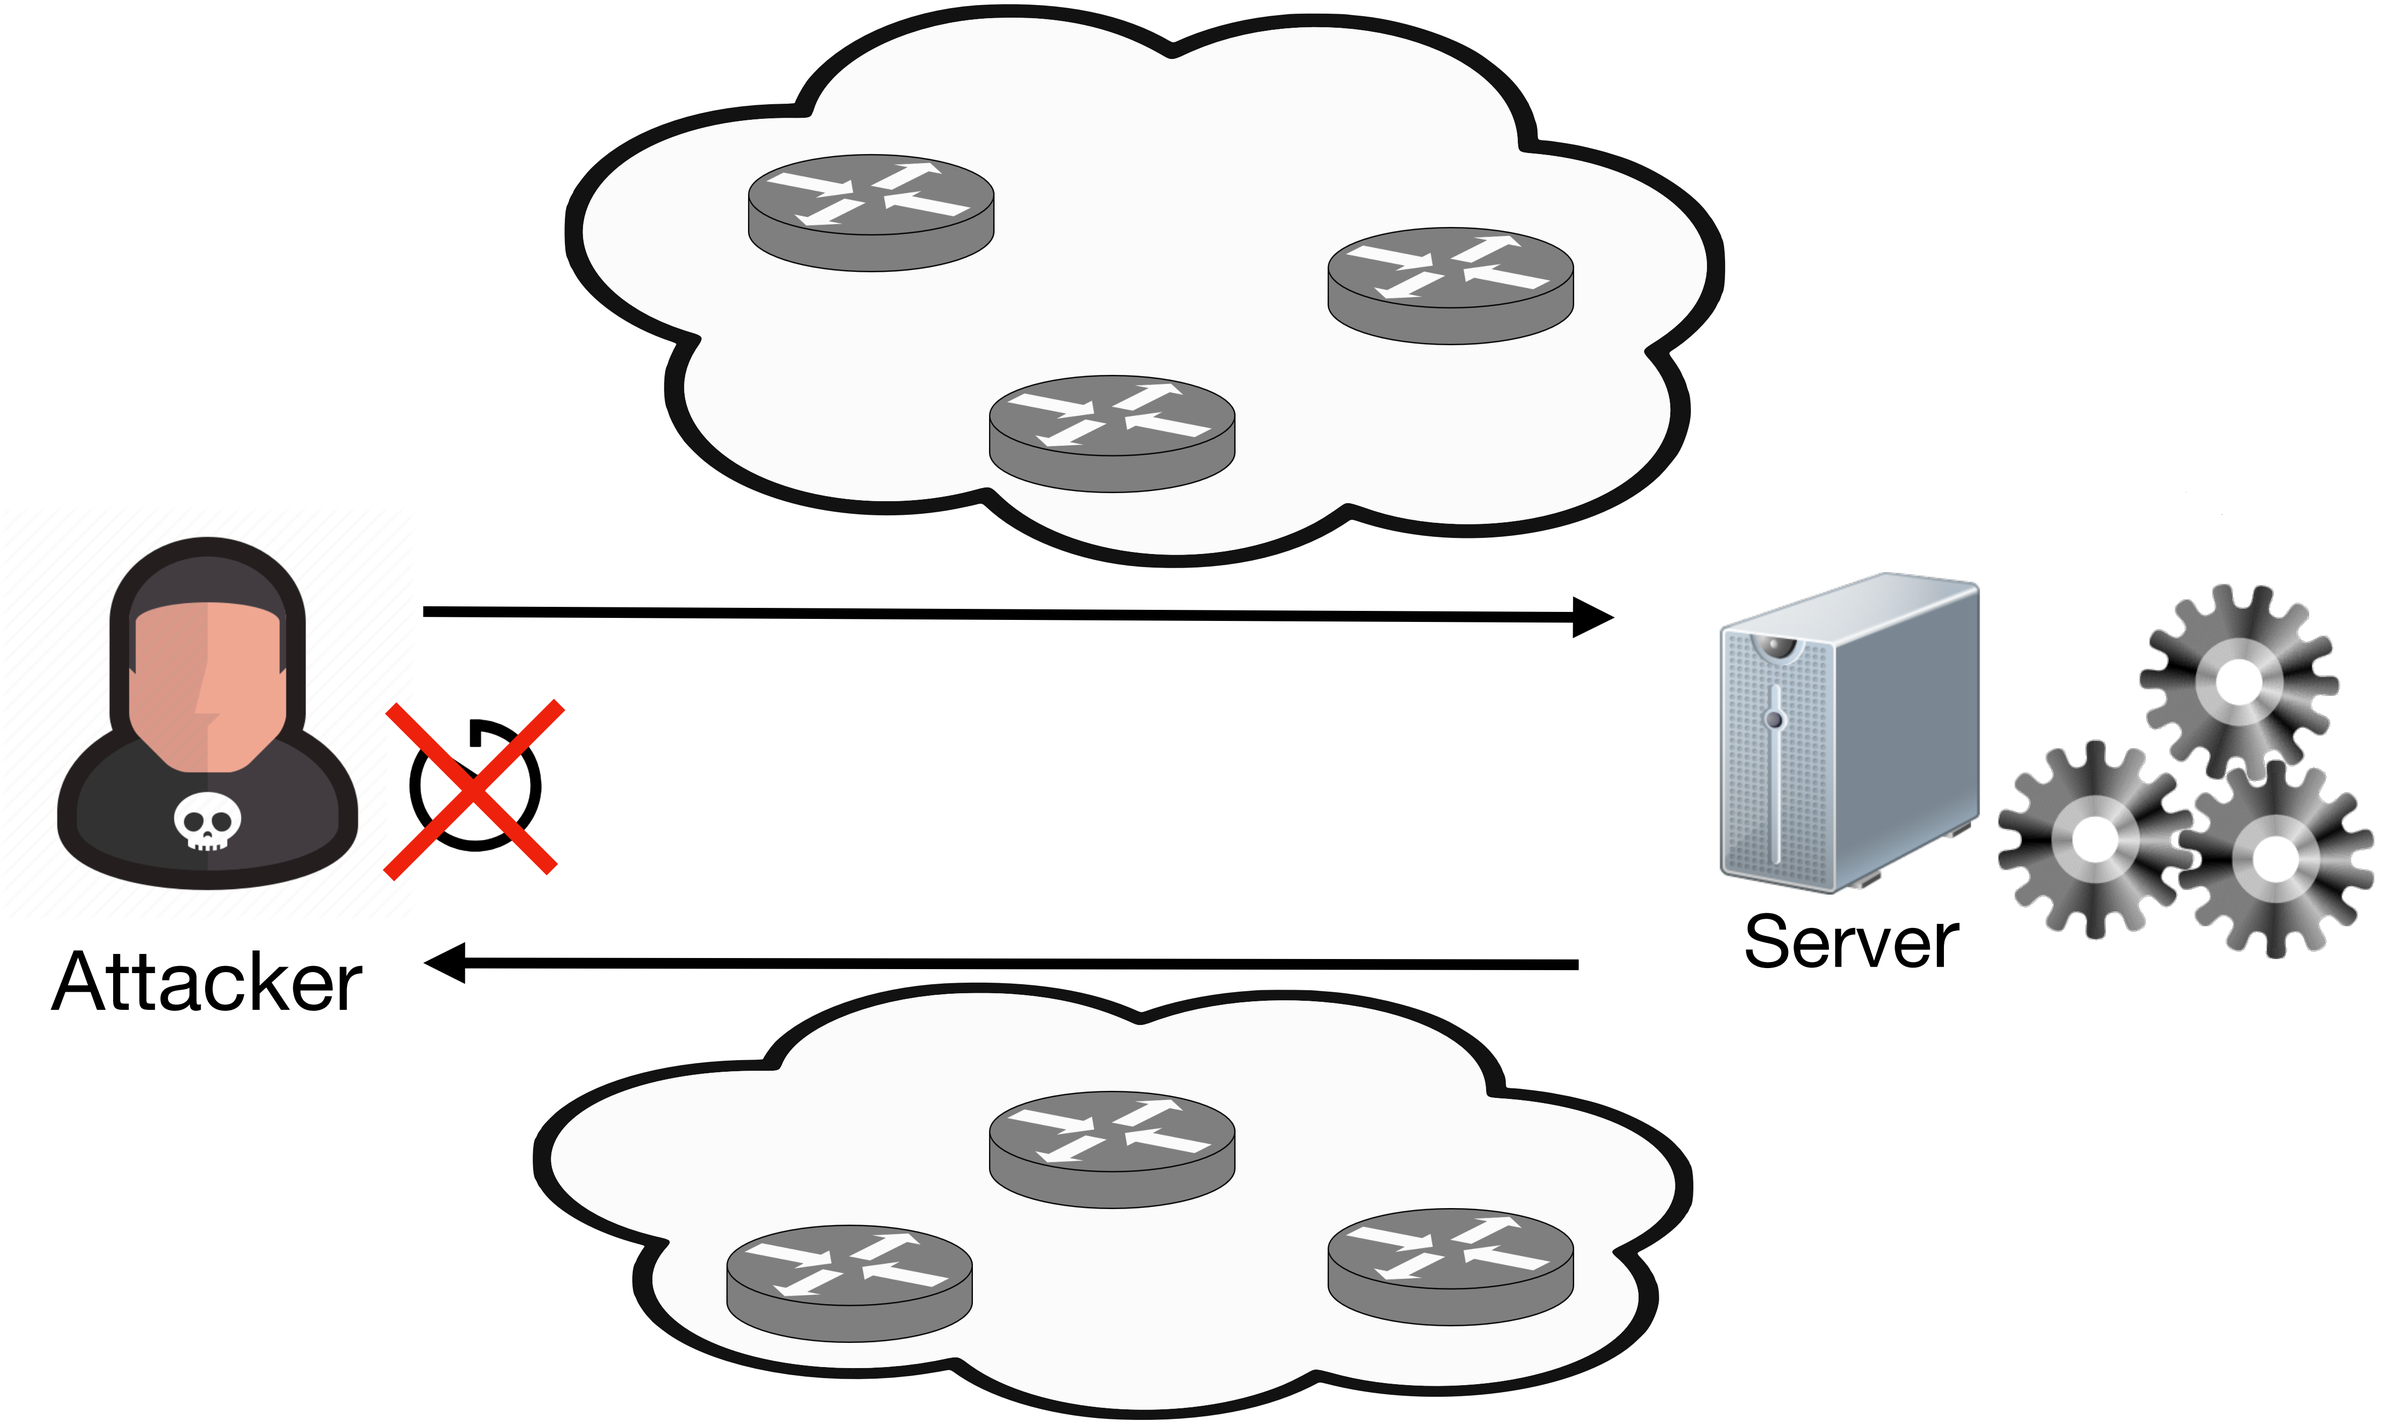
Attacker
Server

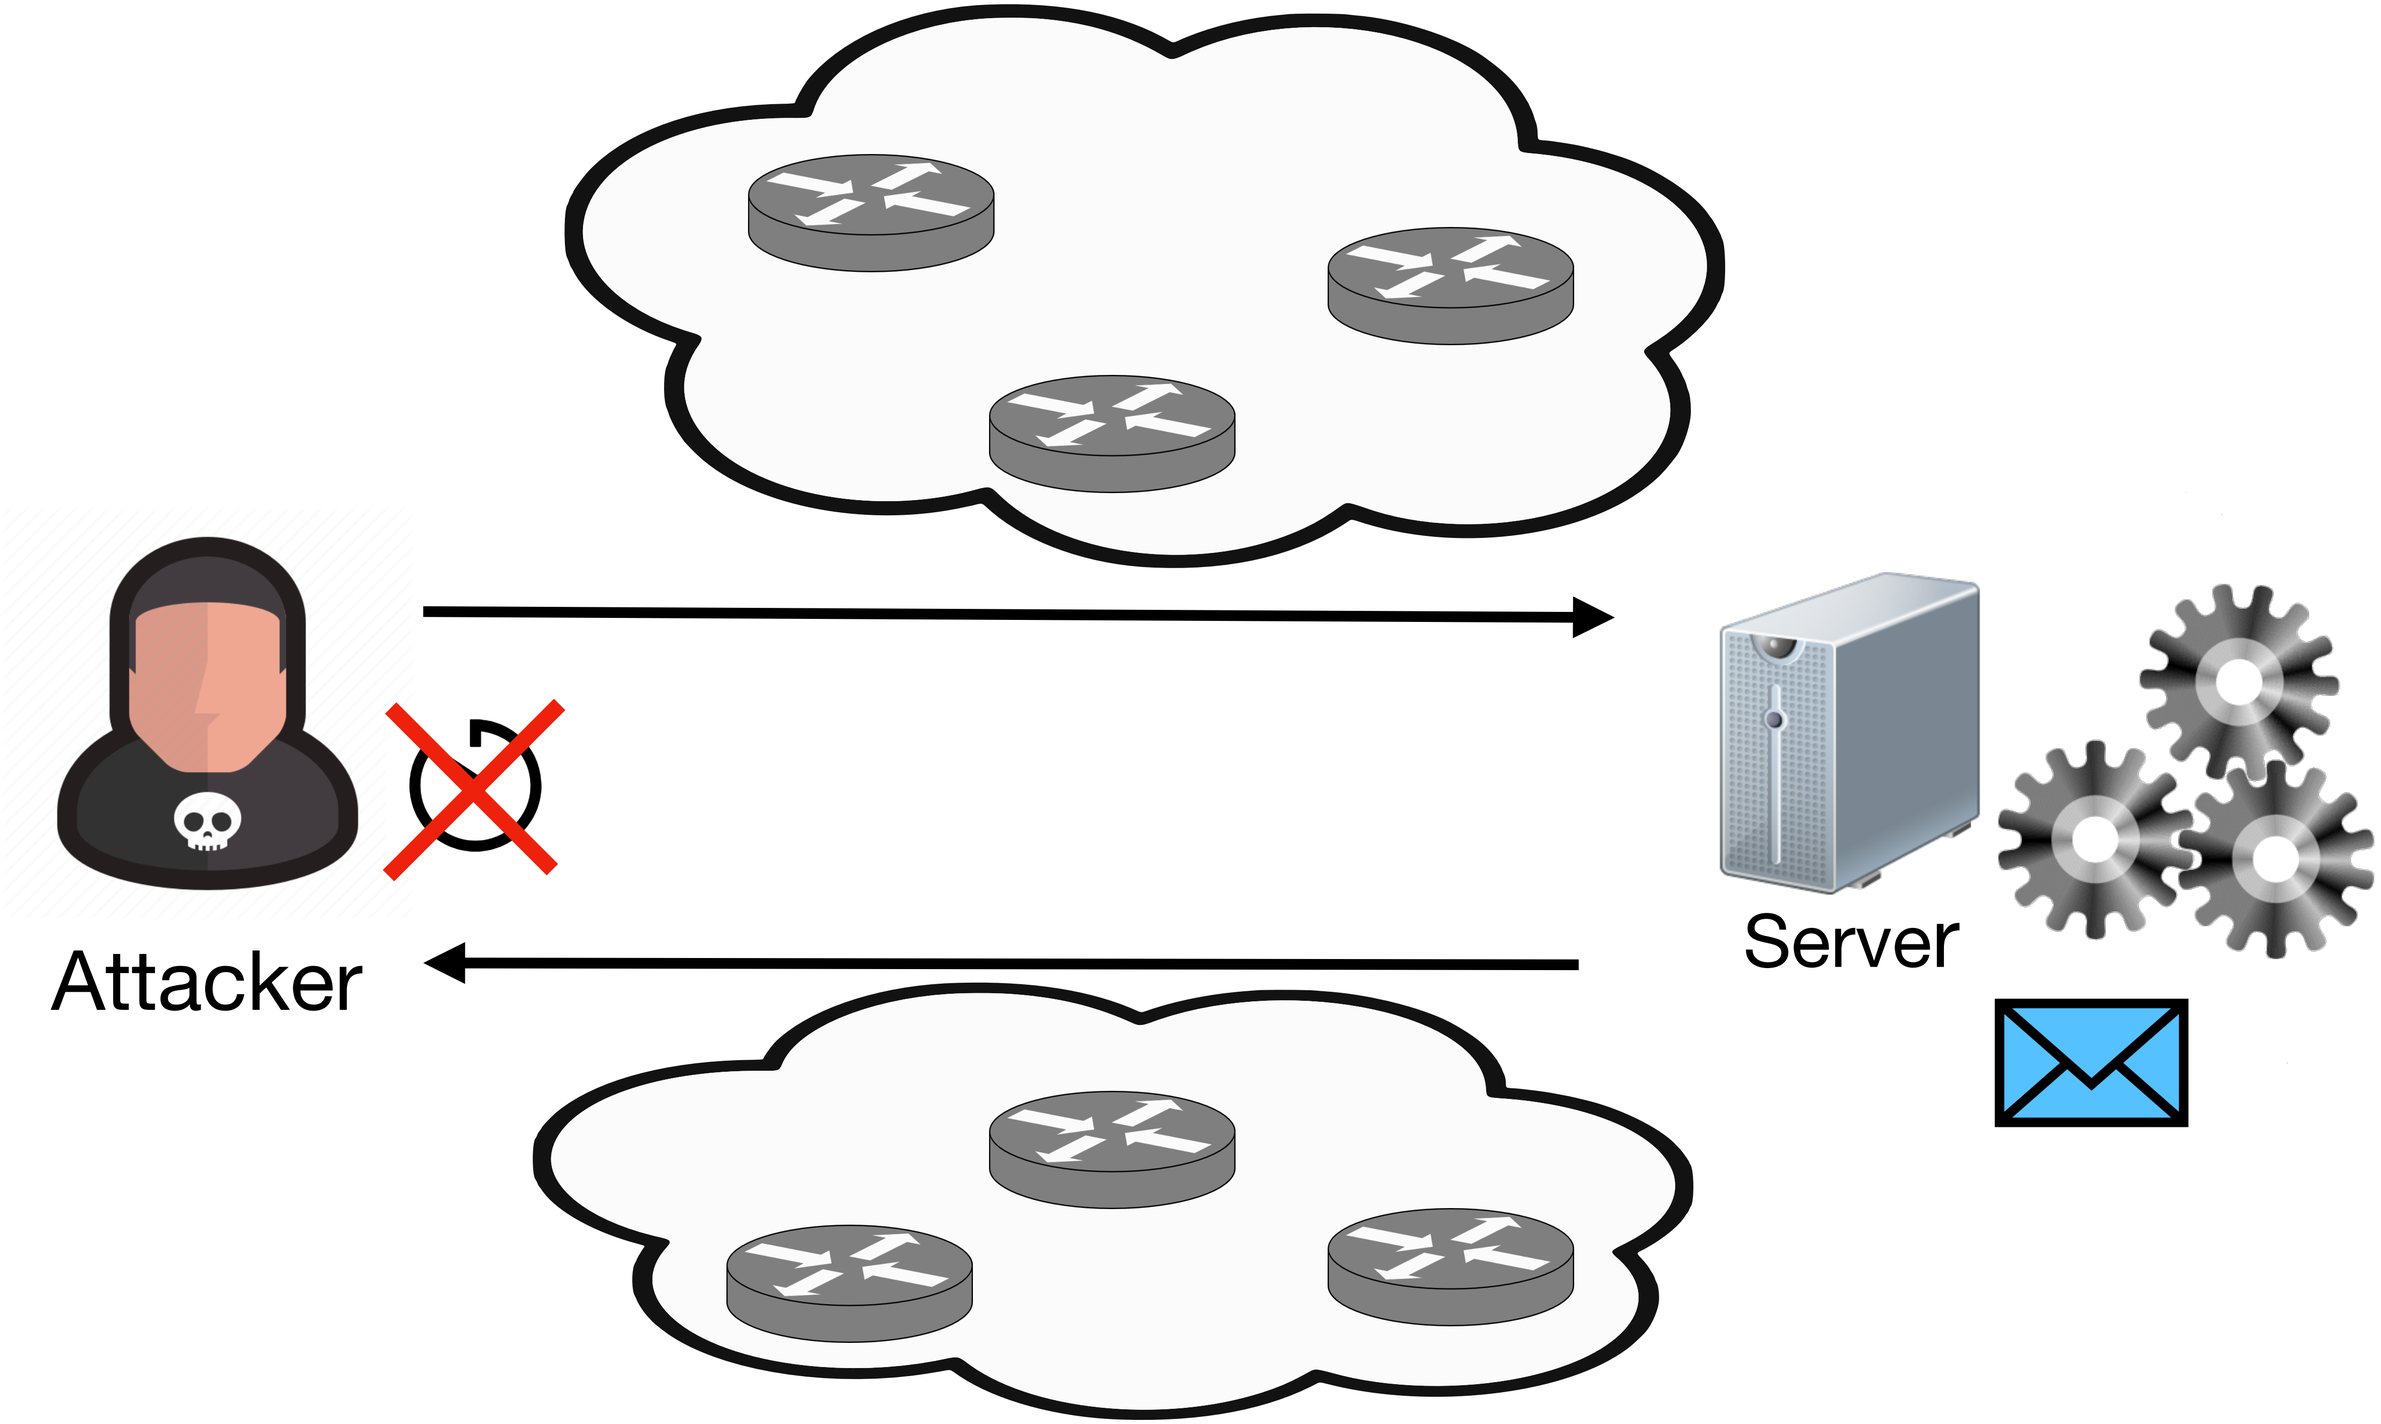
Attacker
Server

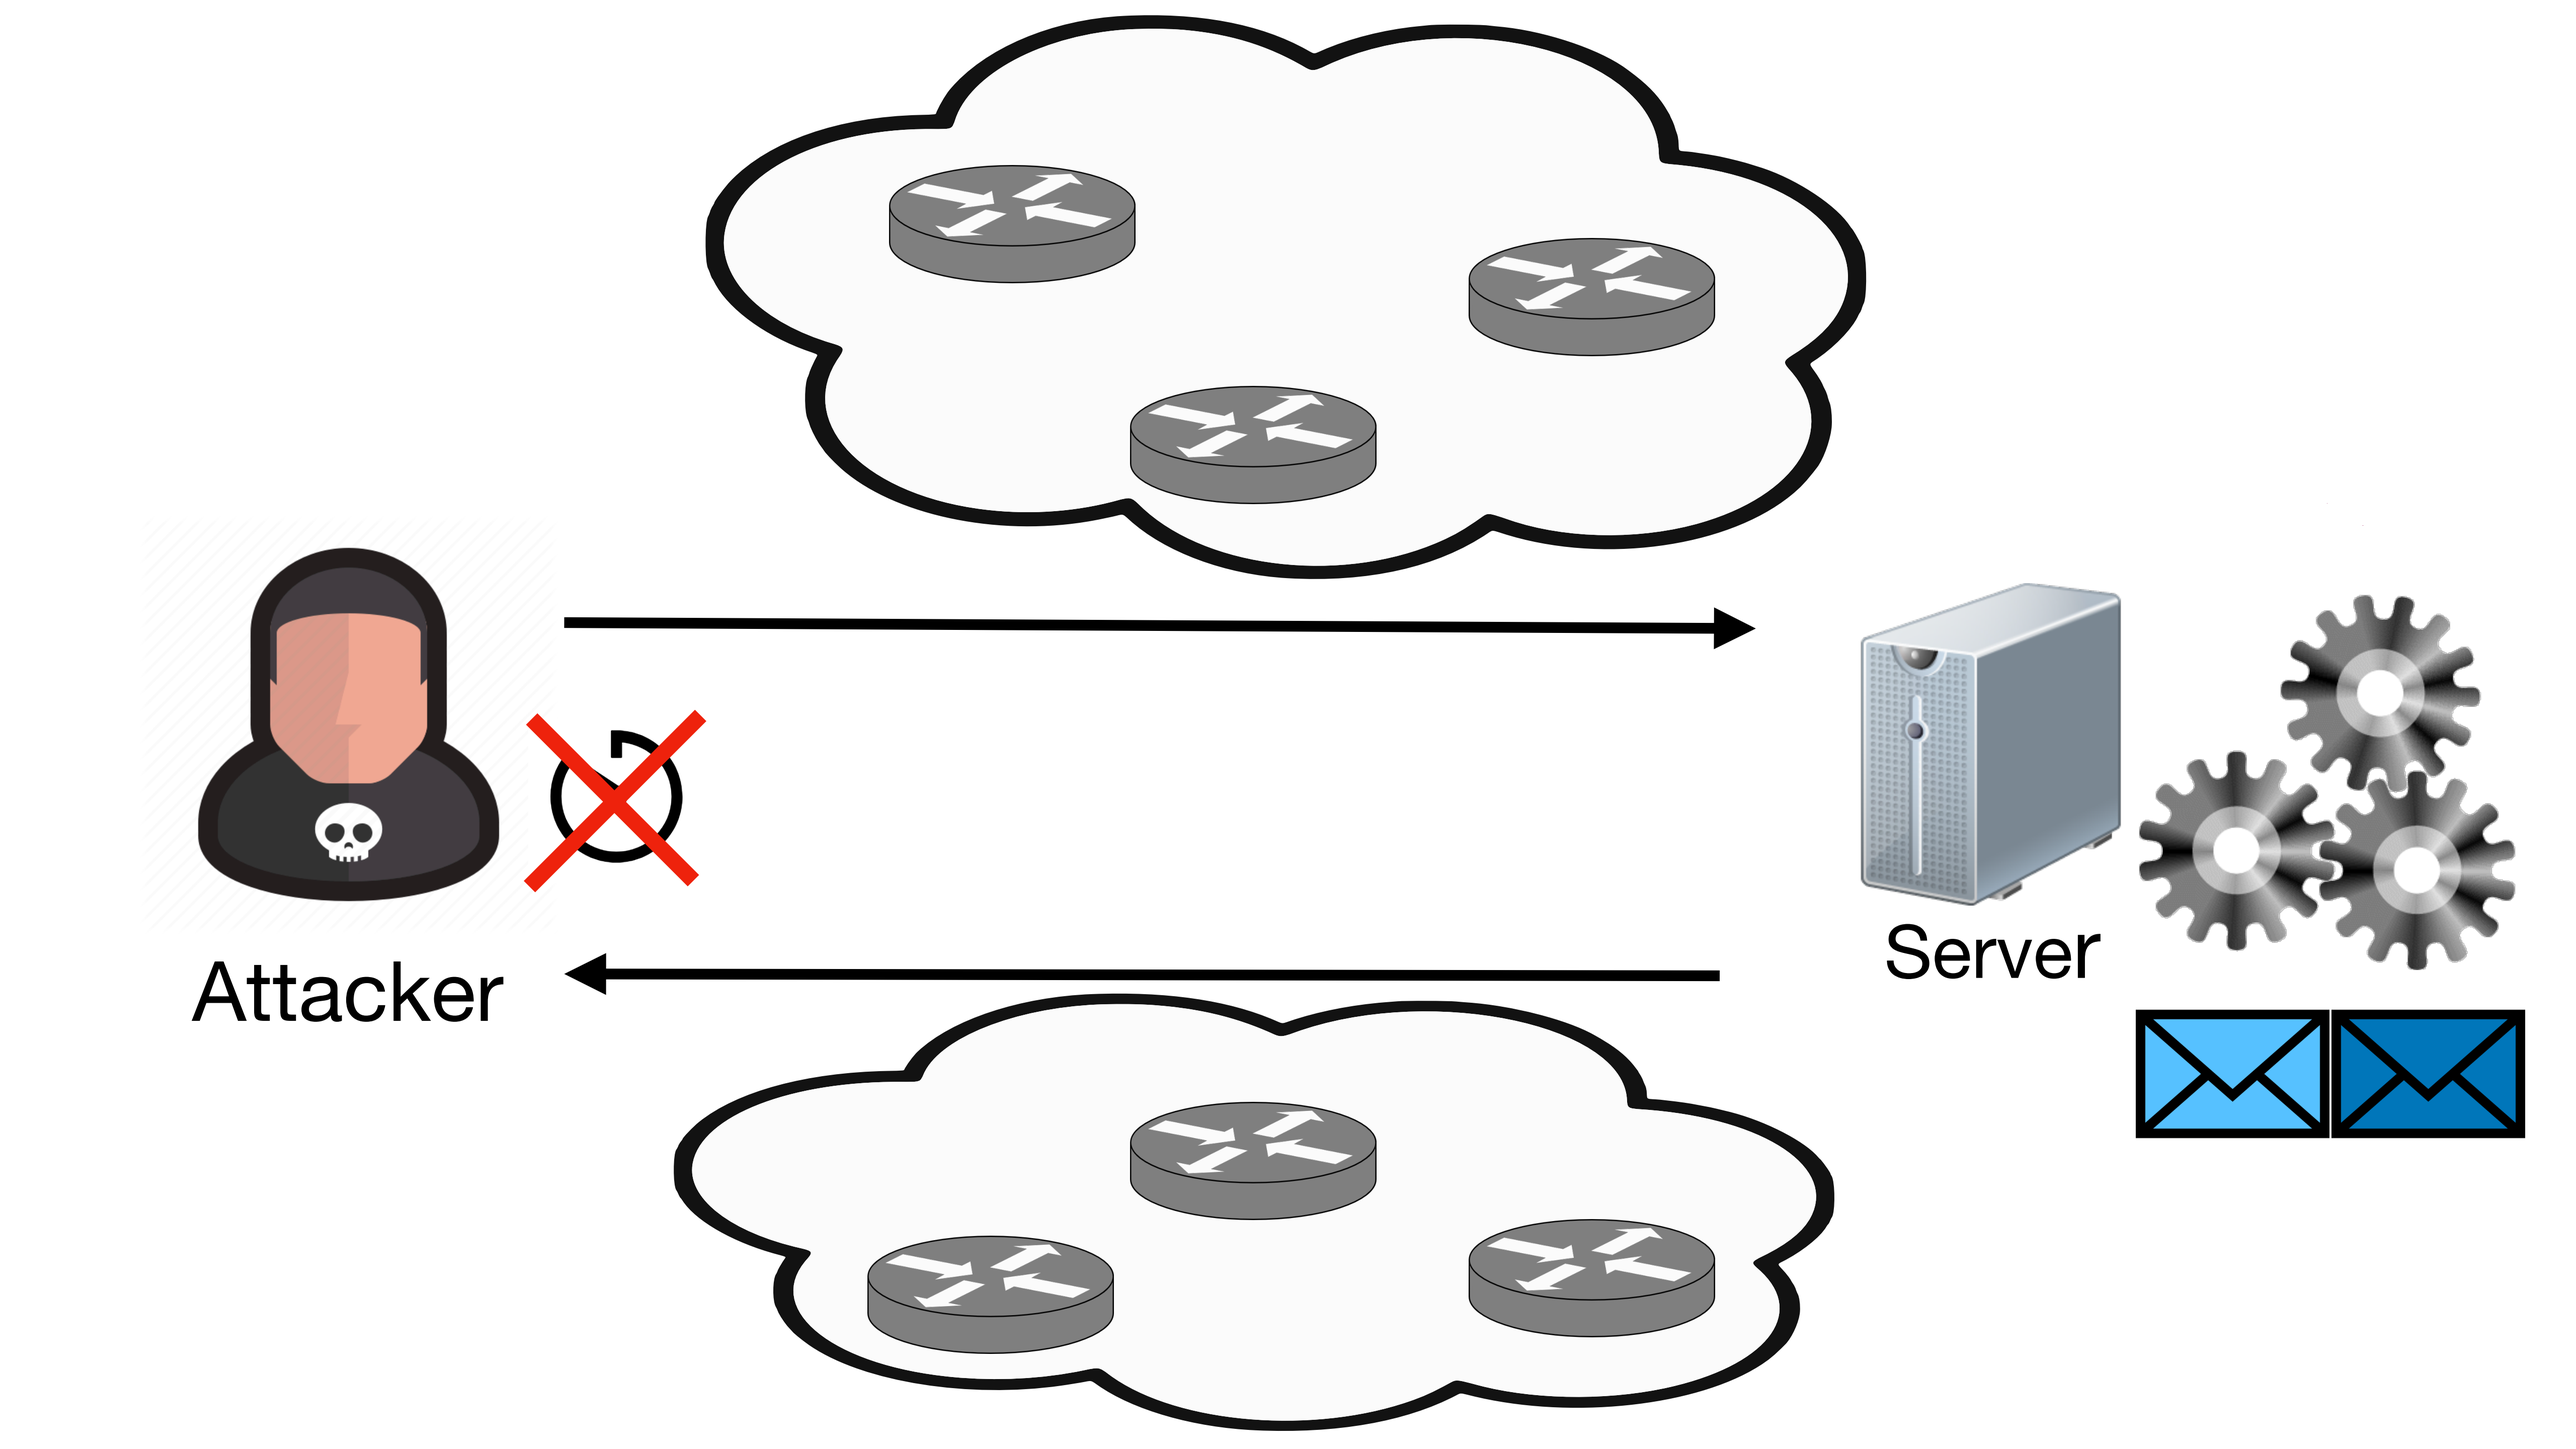
Attacker
Server

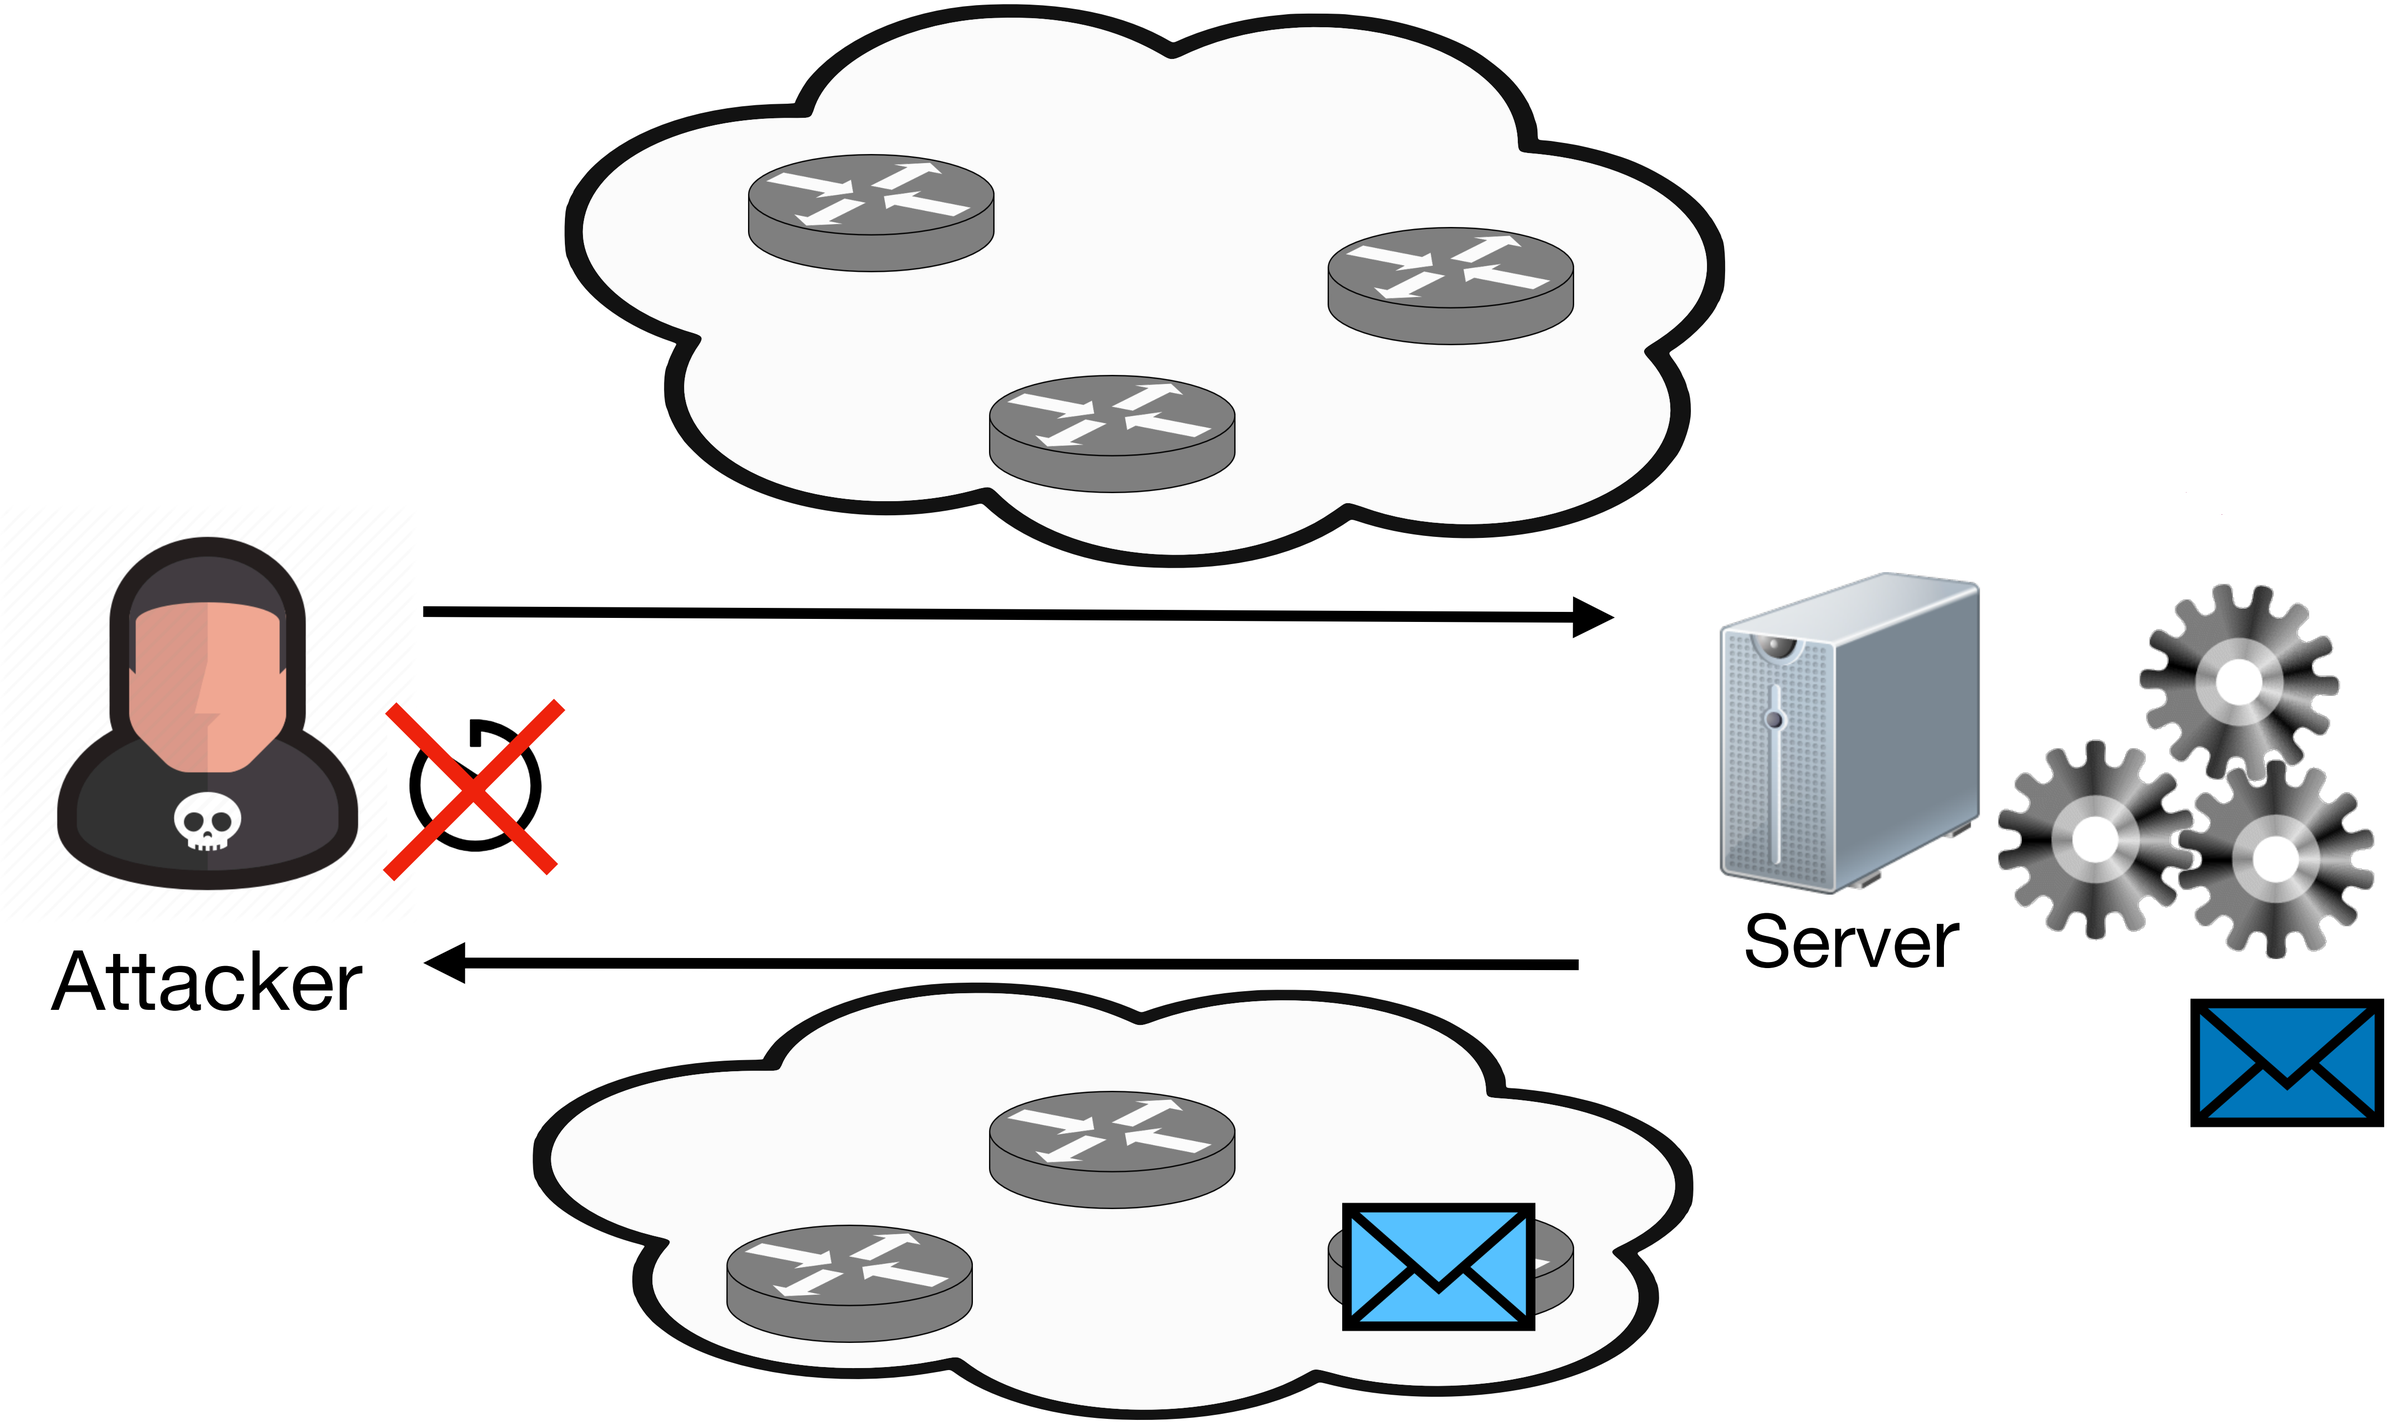
Attacker
Server

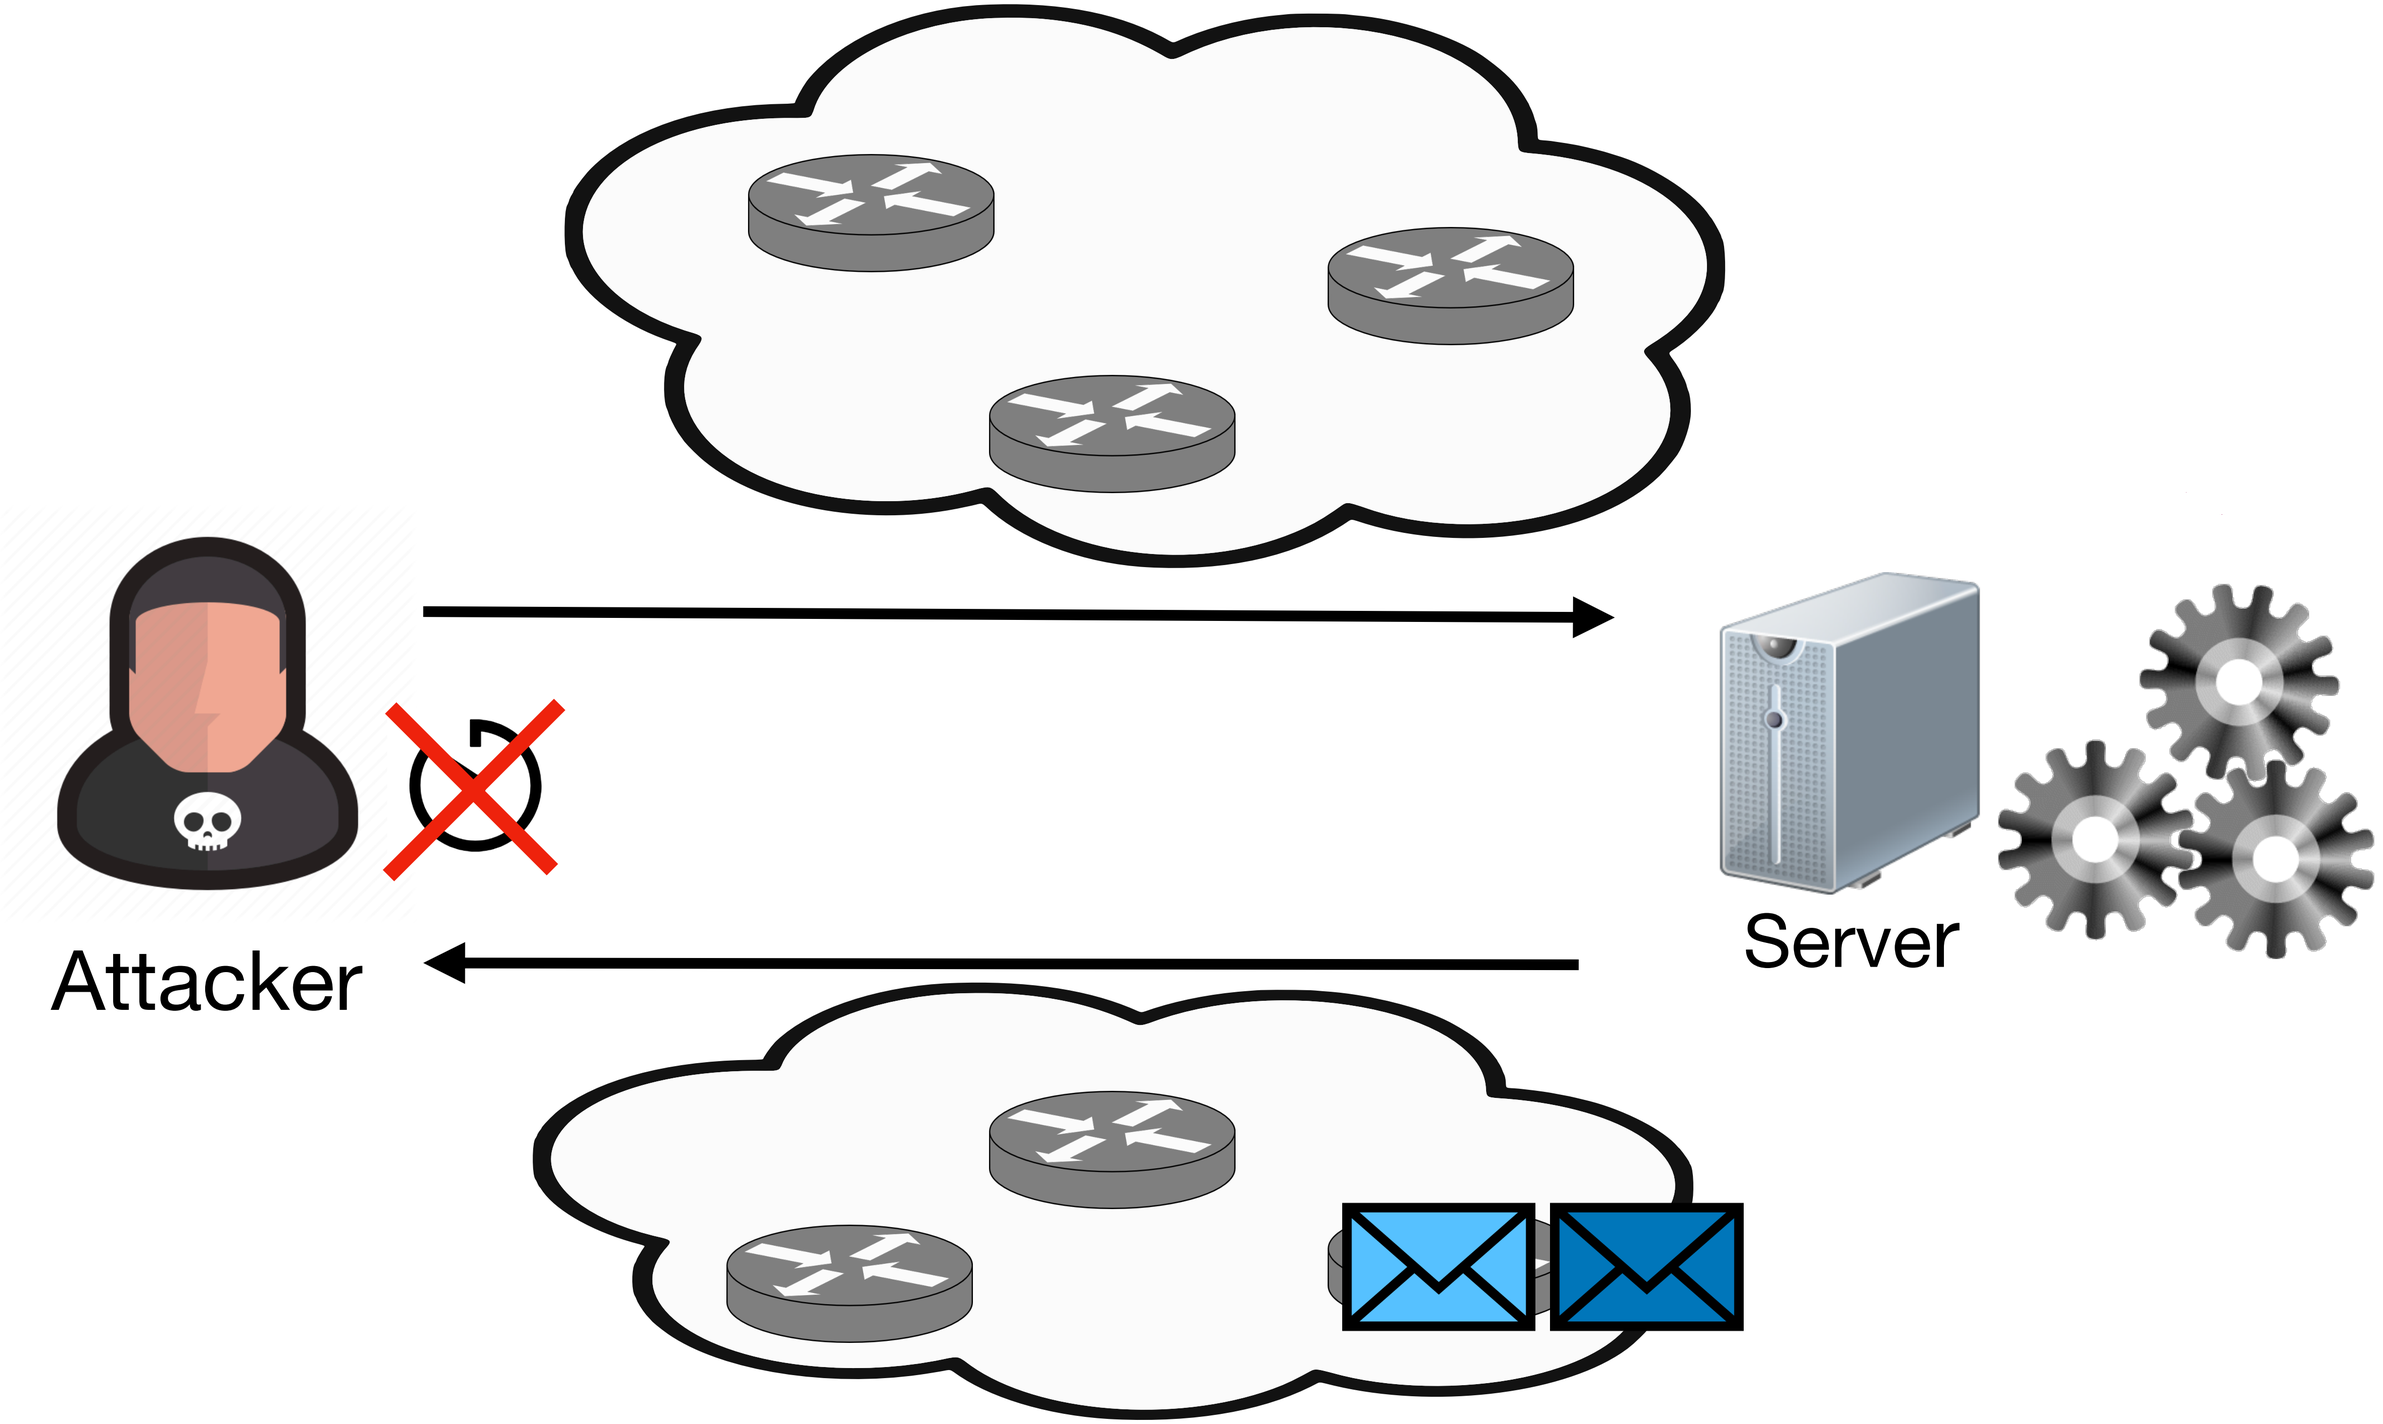
Attacker
Server

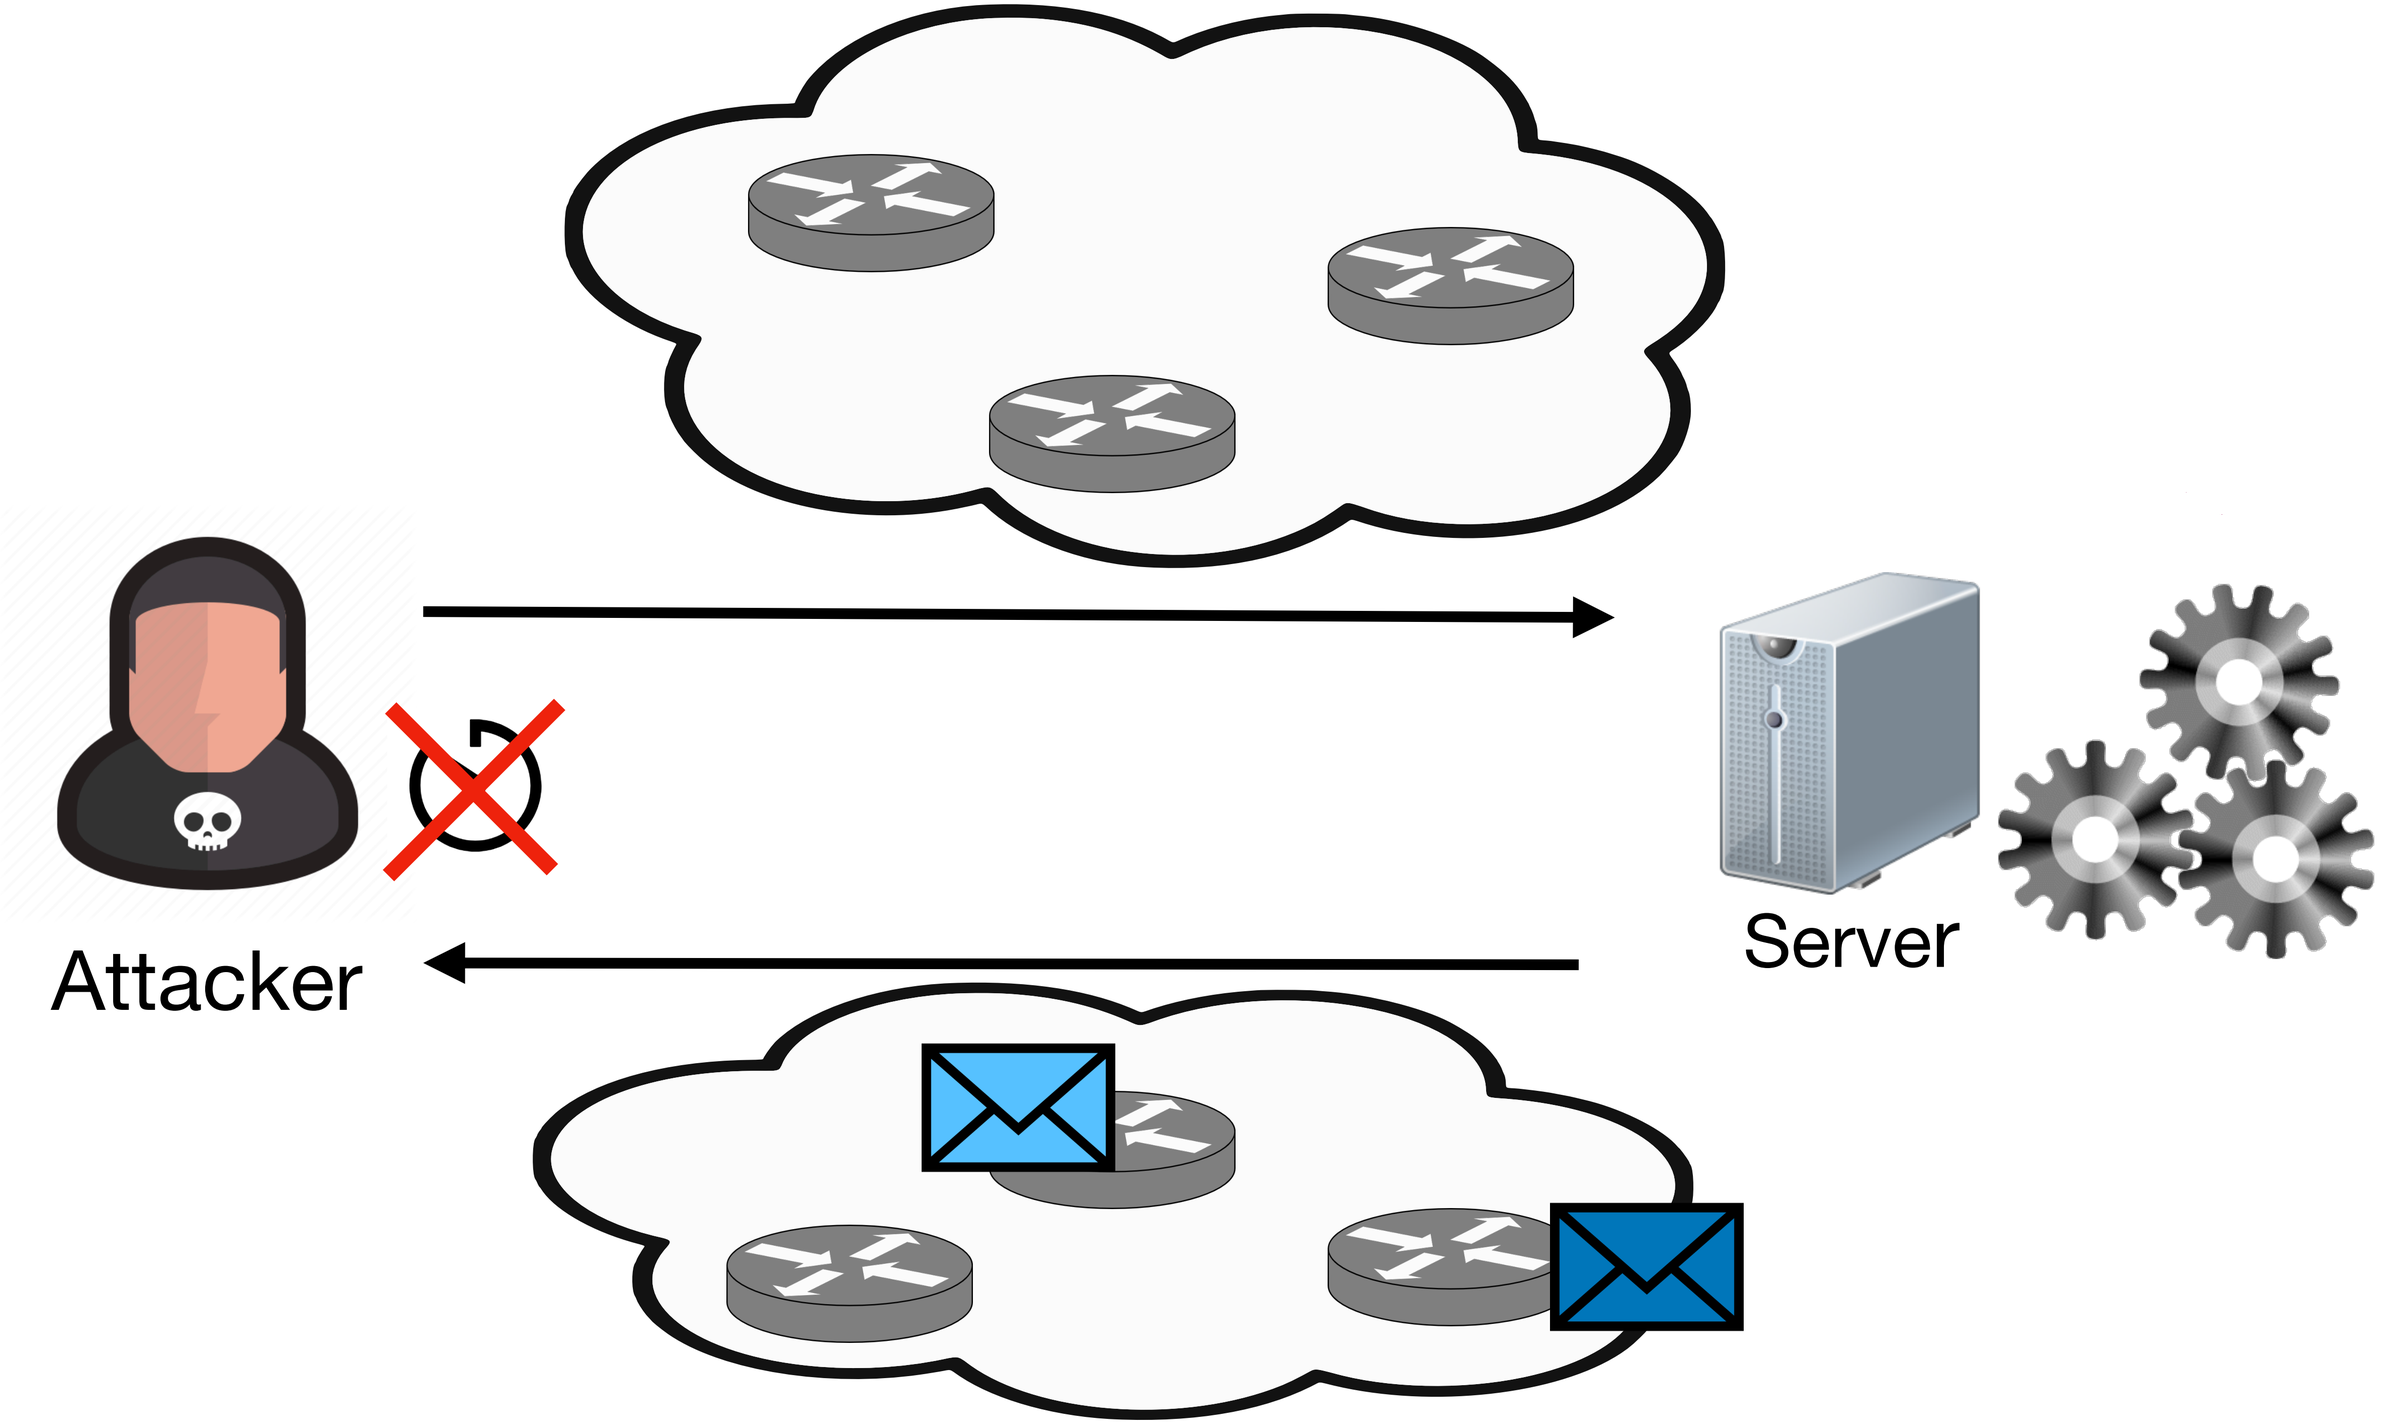
Attacker
Server

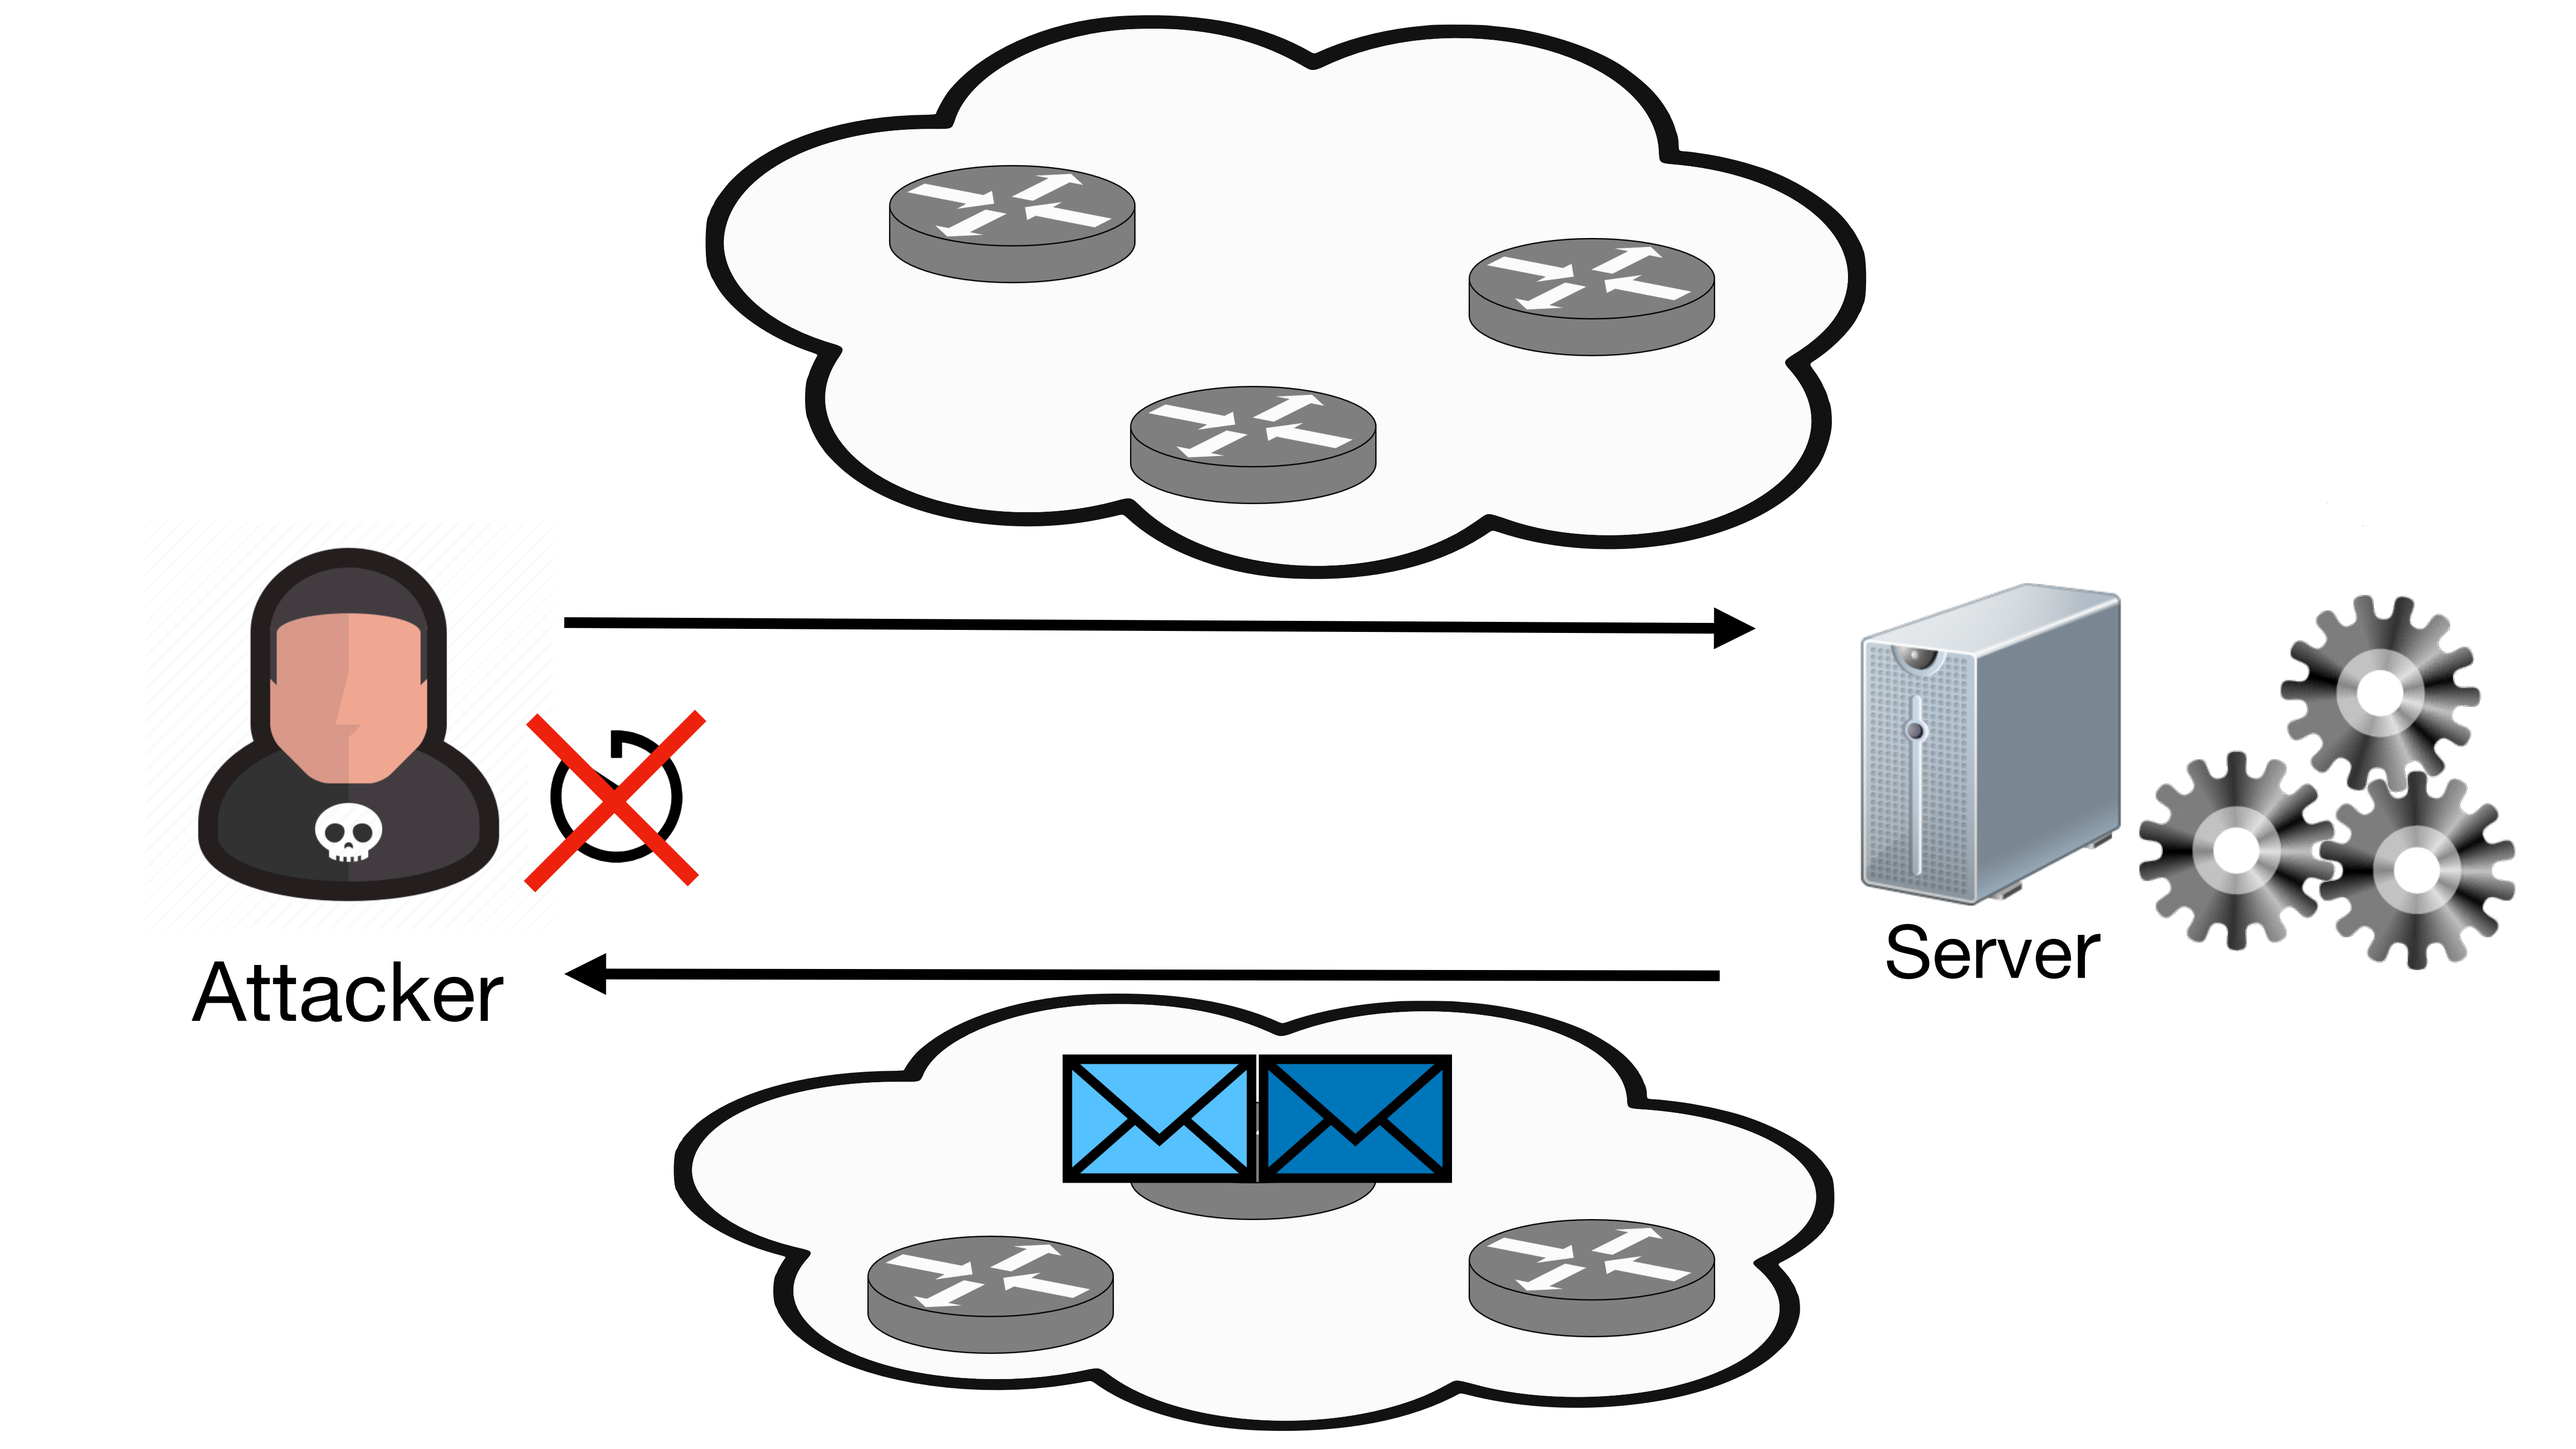
Attacker
Server

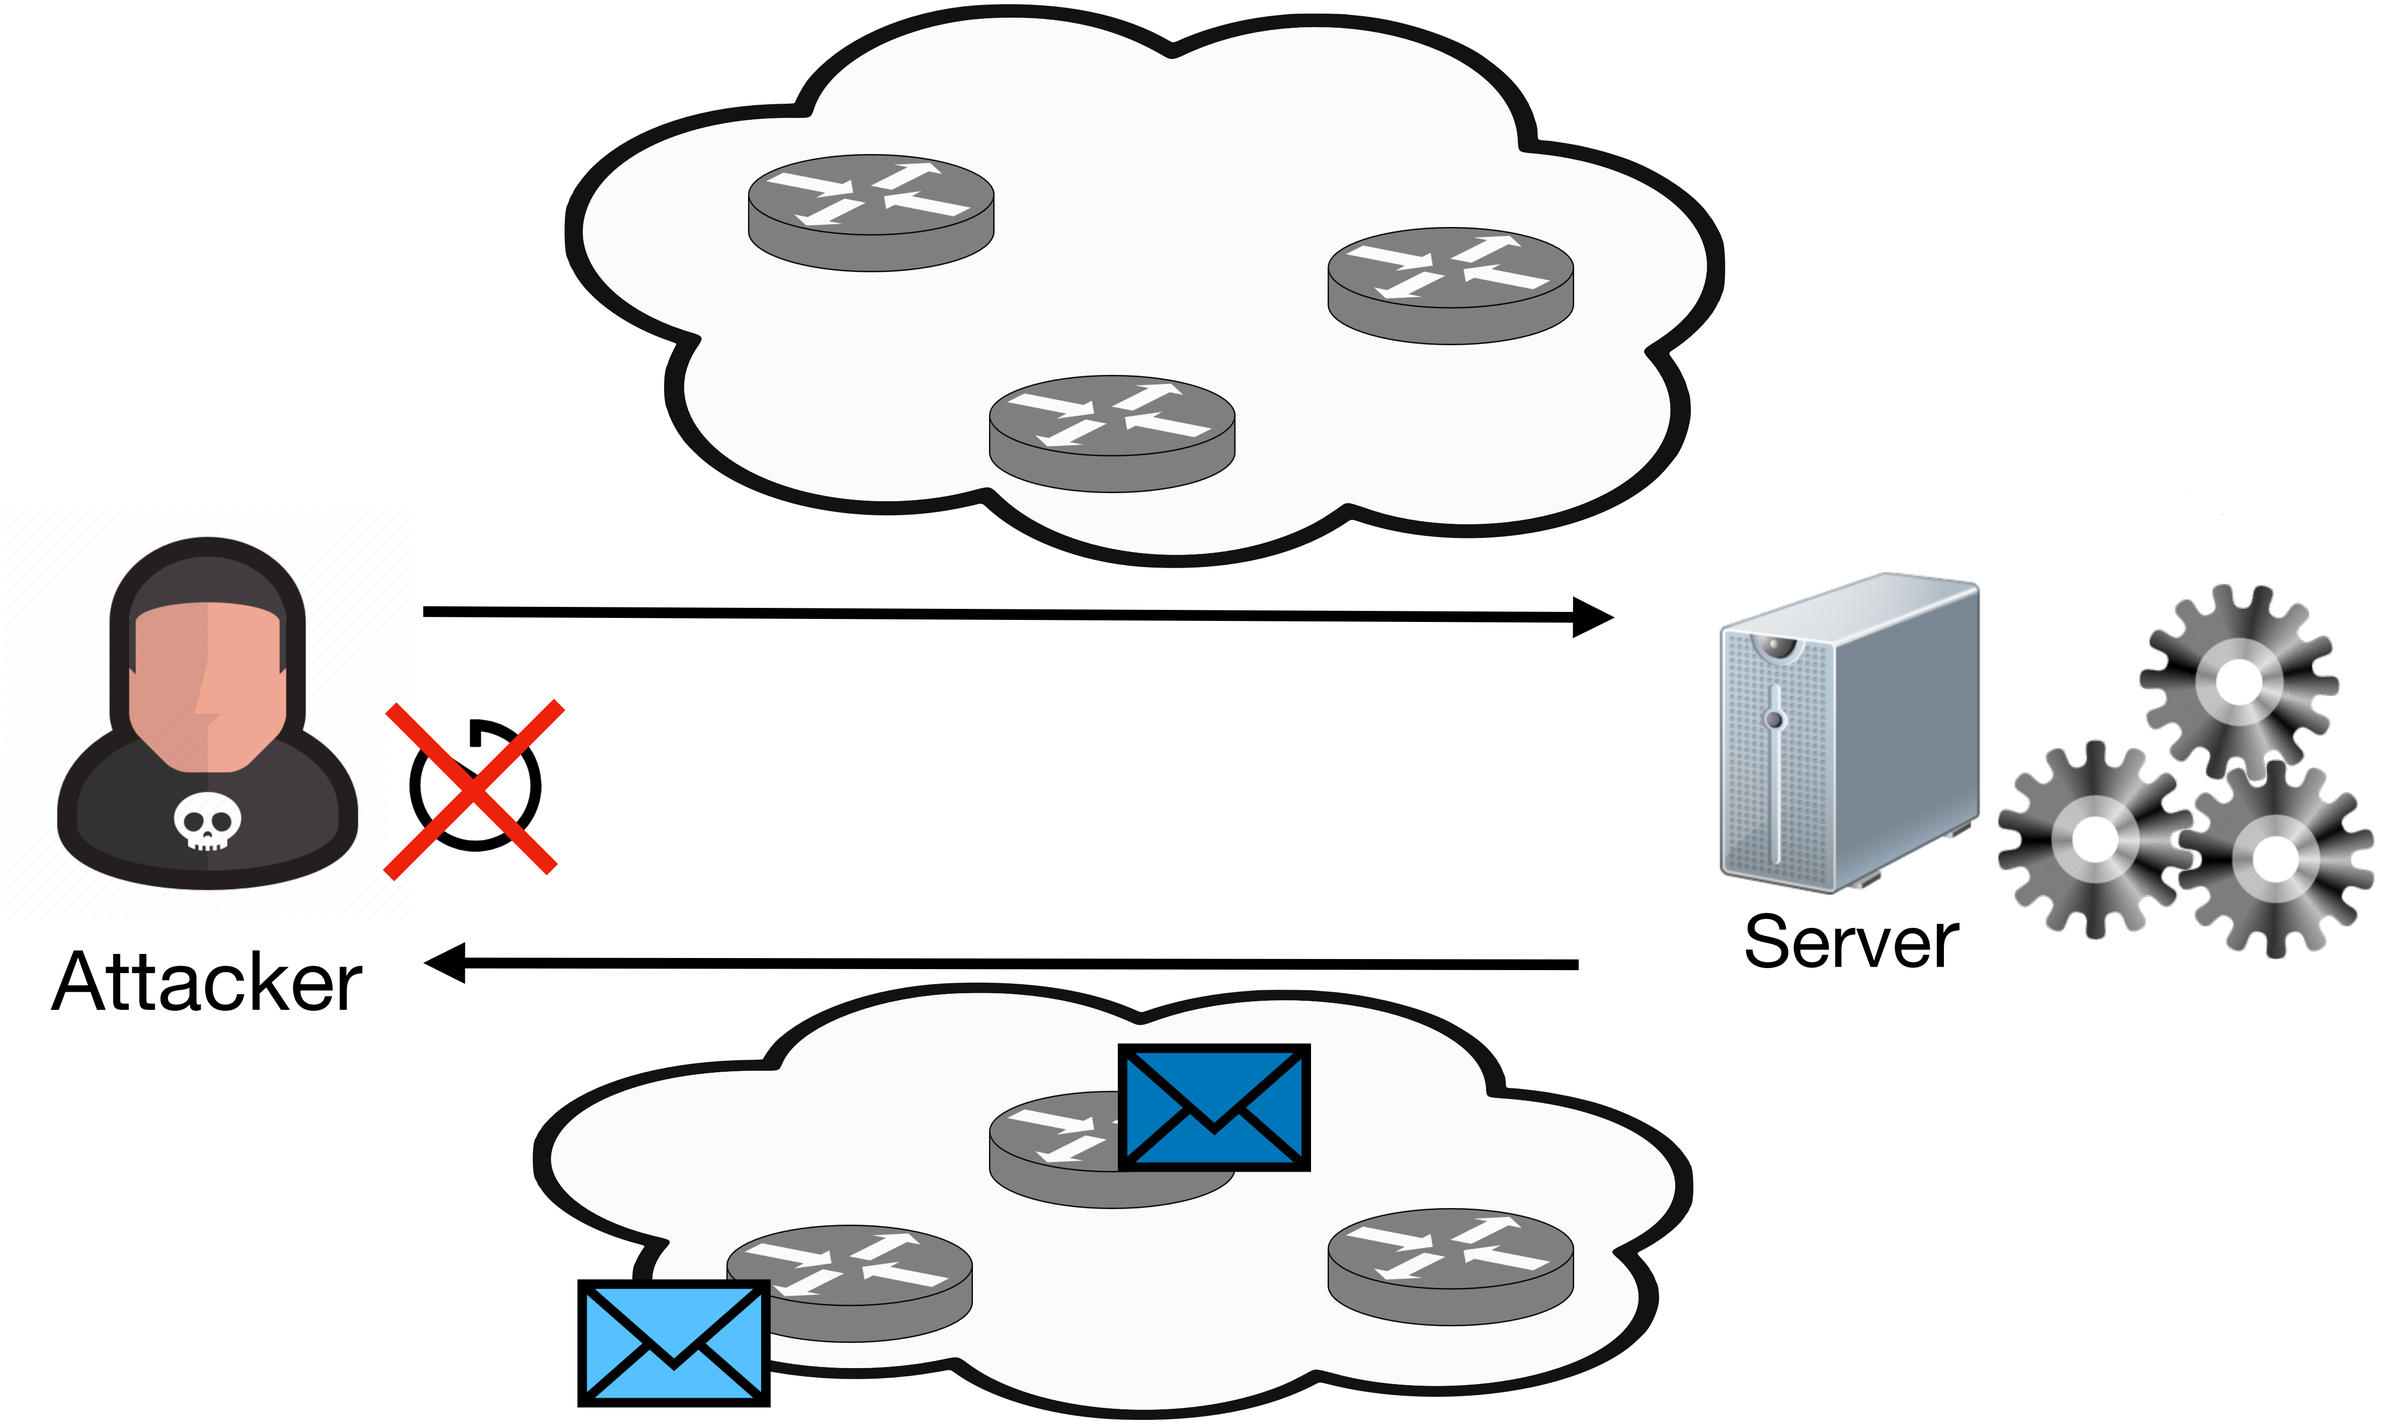
Attacker
Server

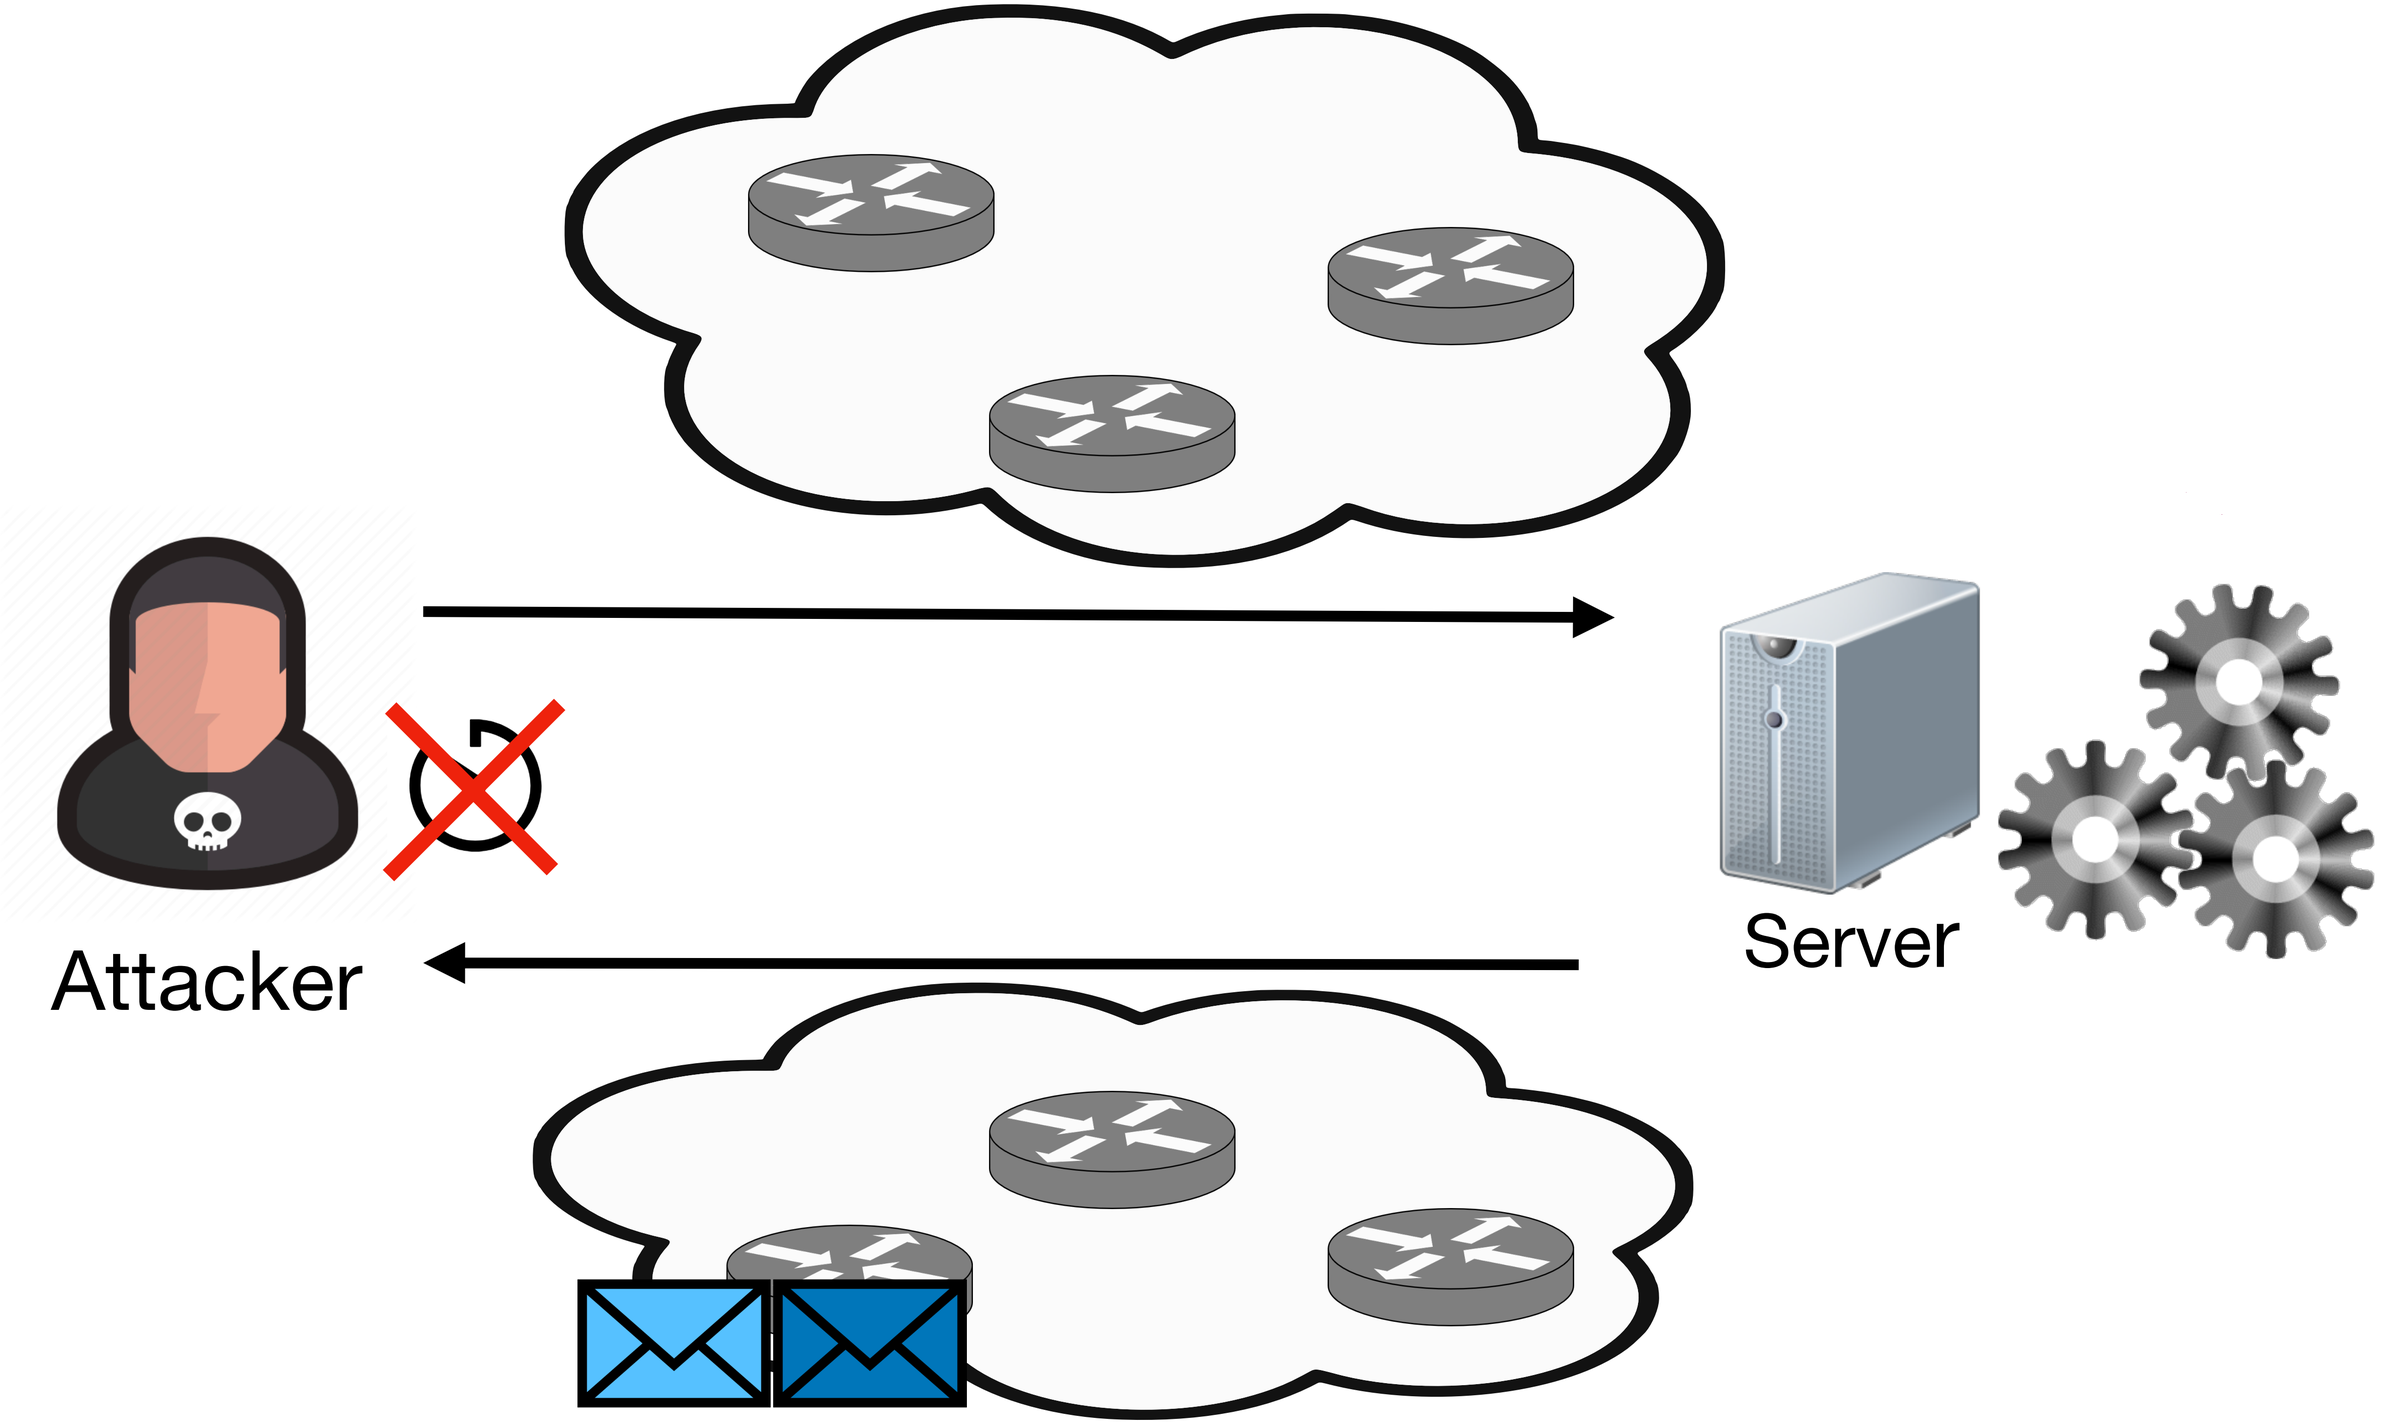
Attacker
Server

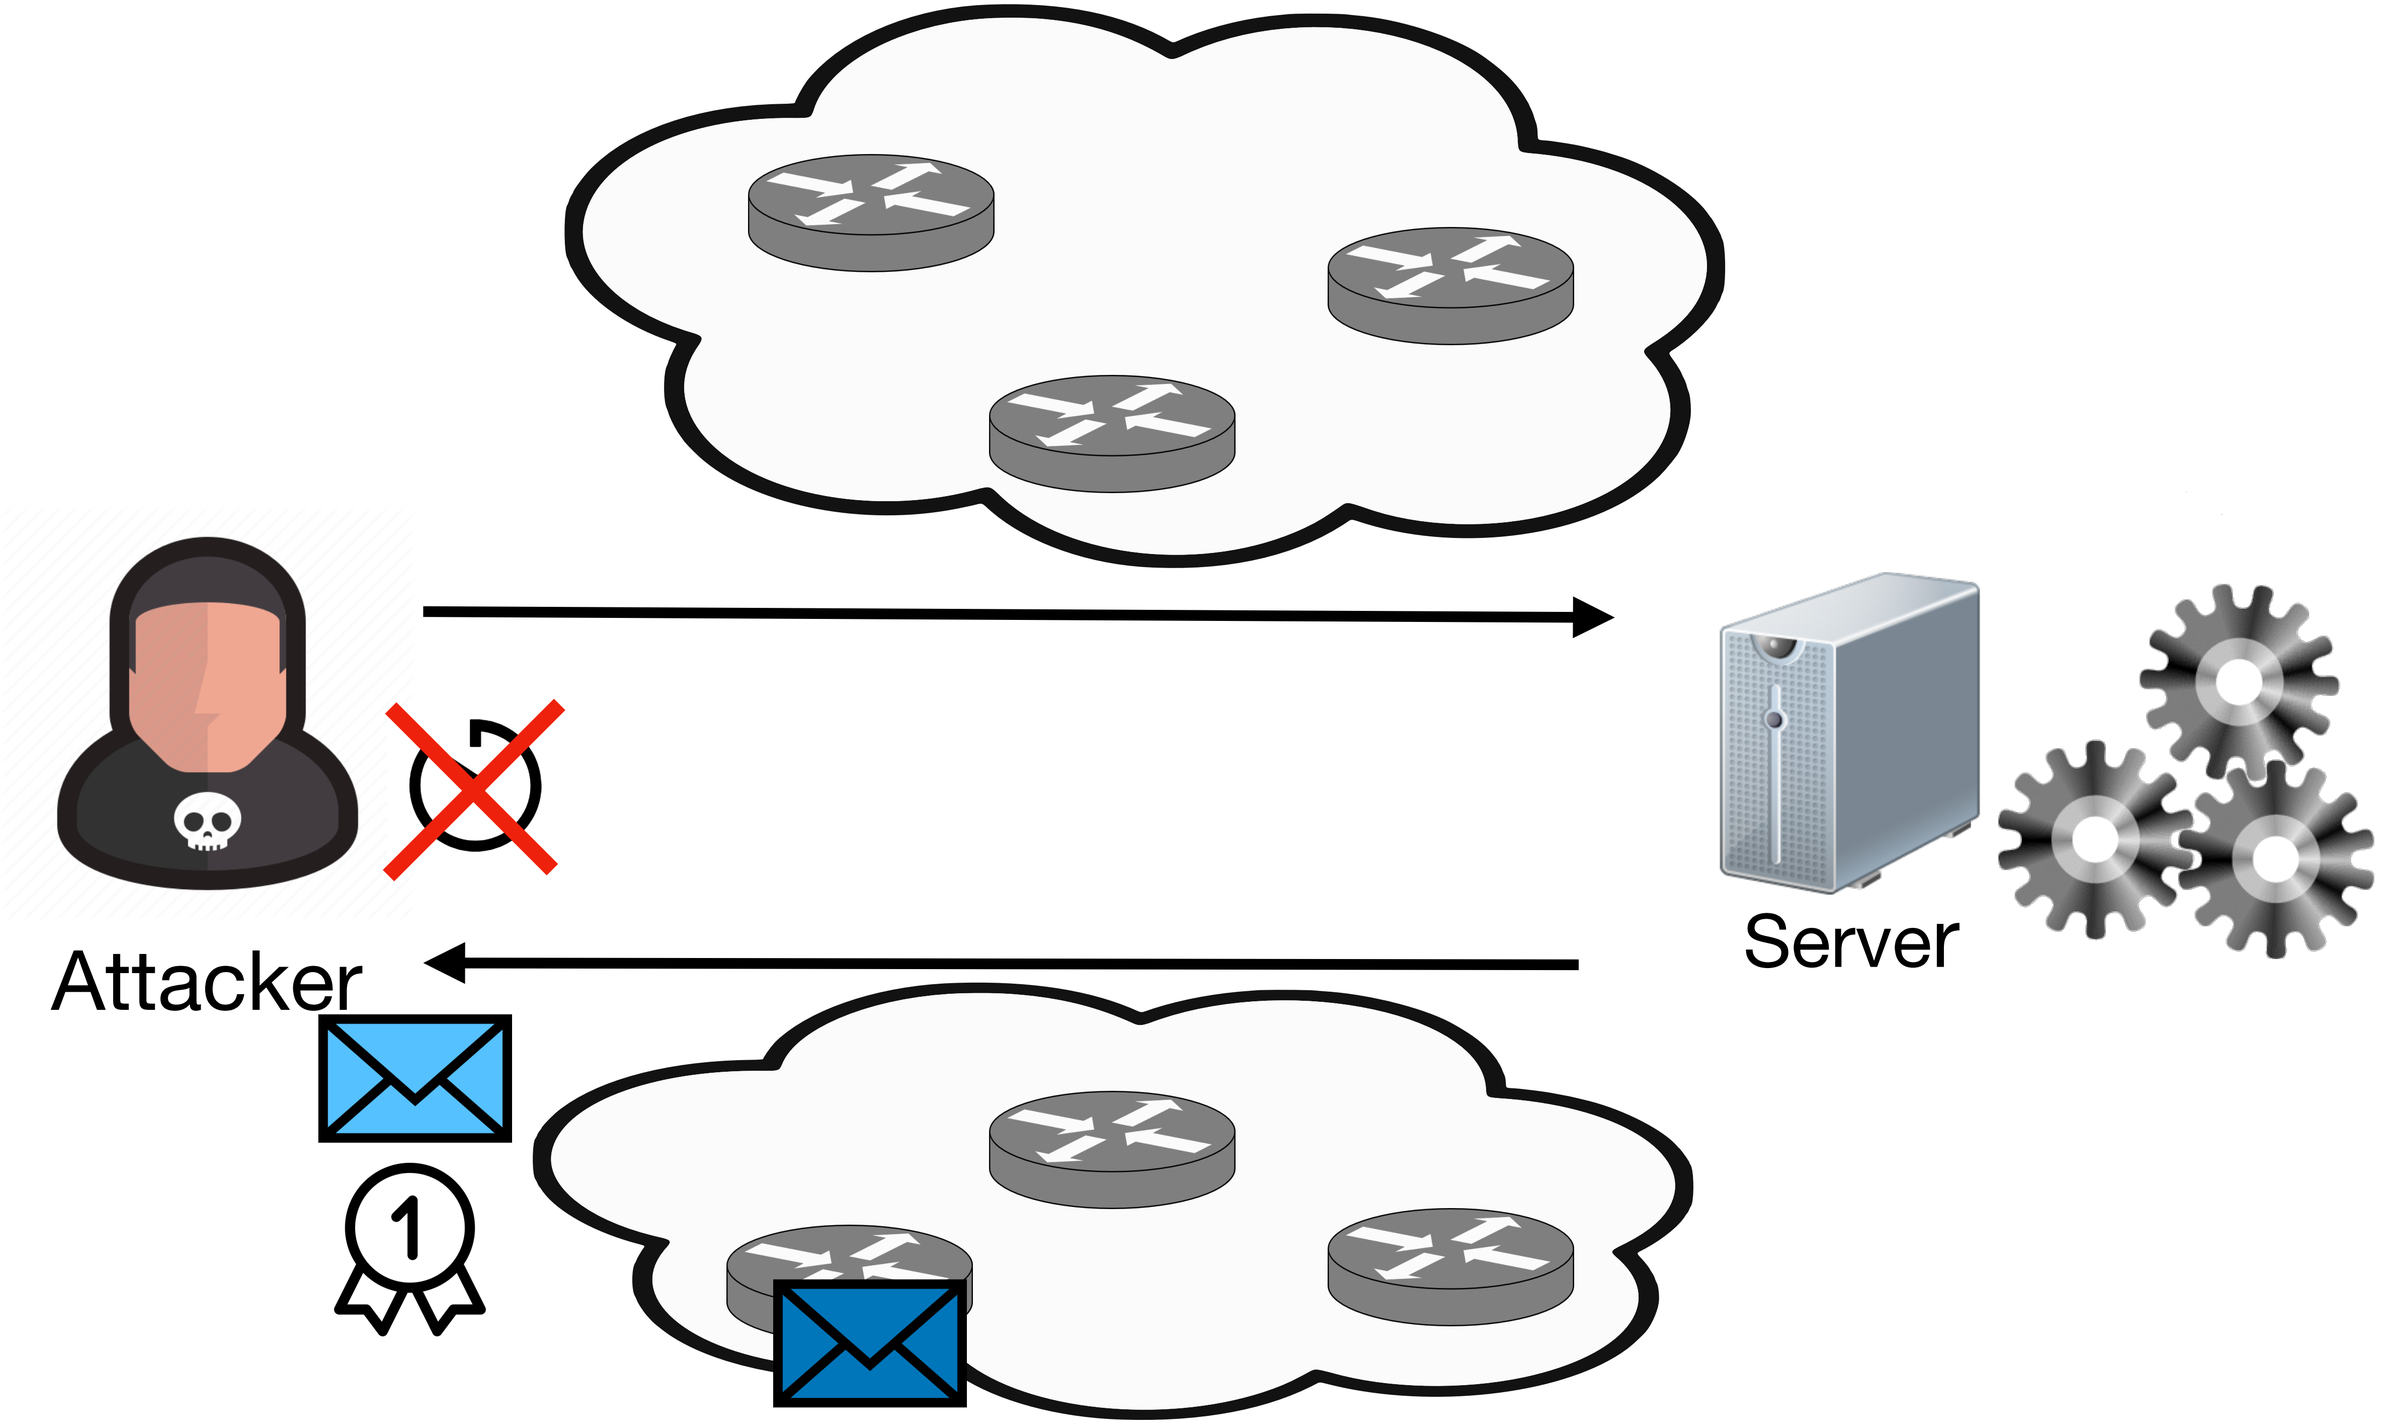
Attacker
Server
1

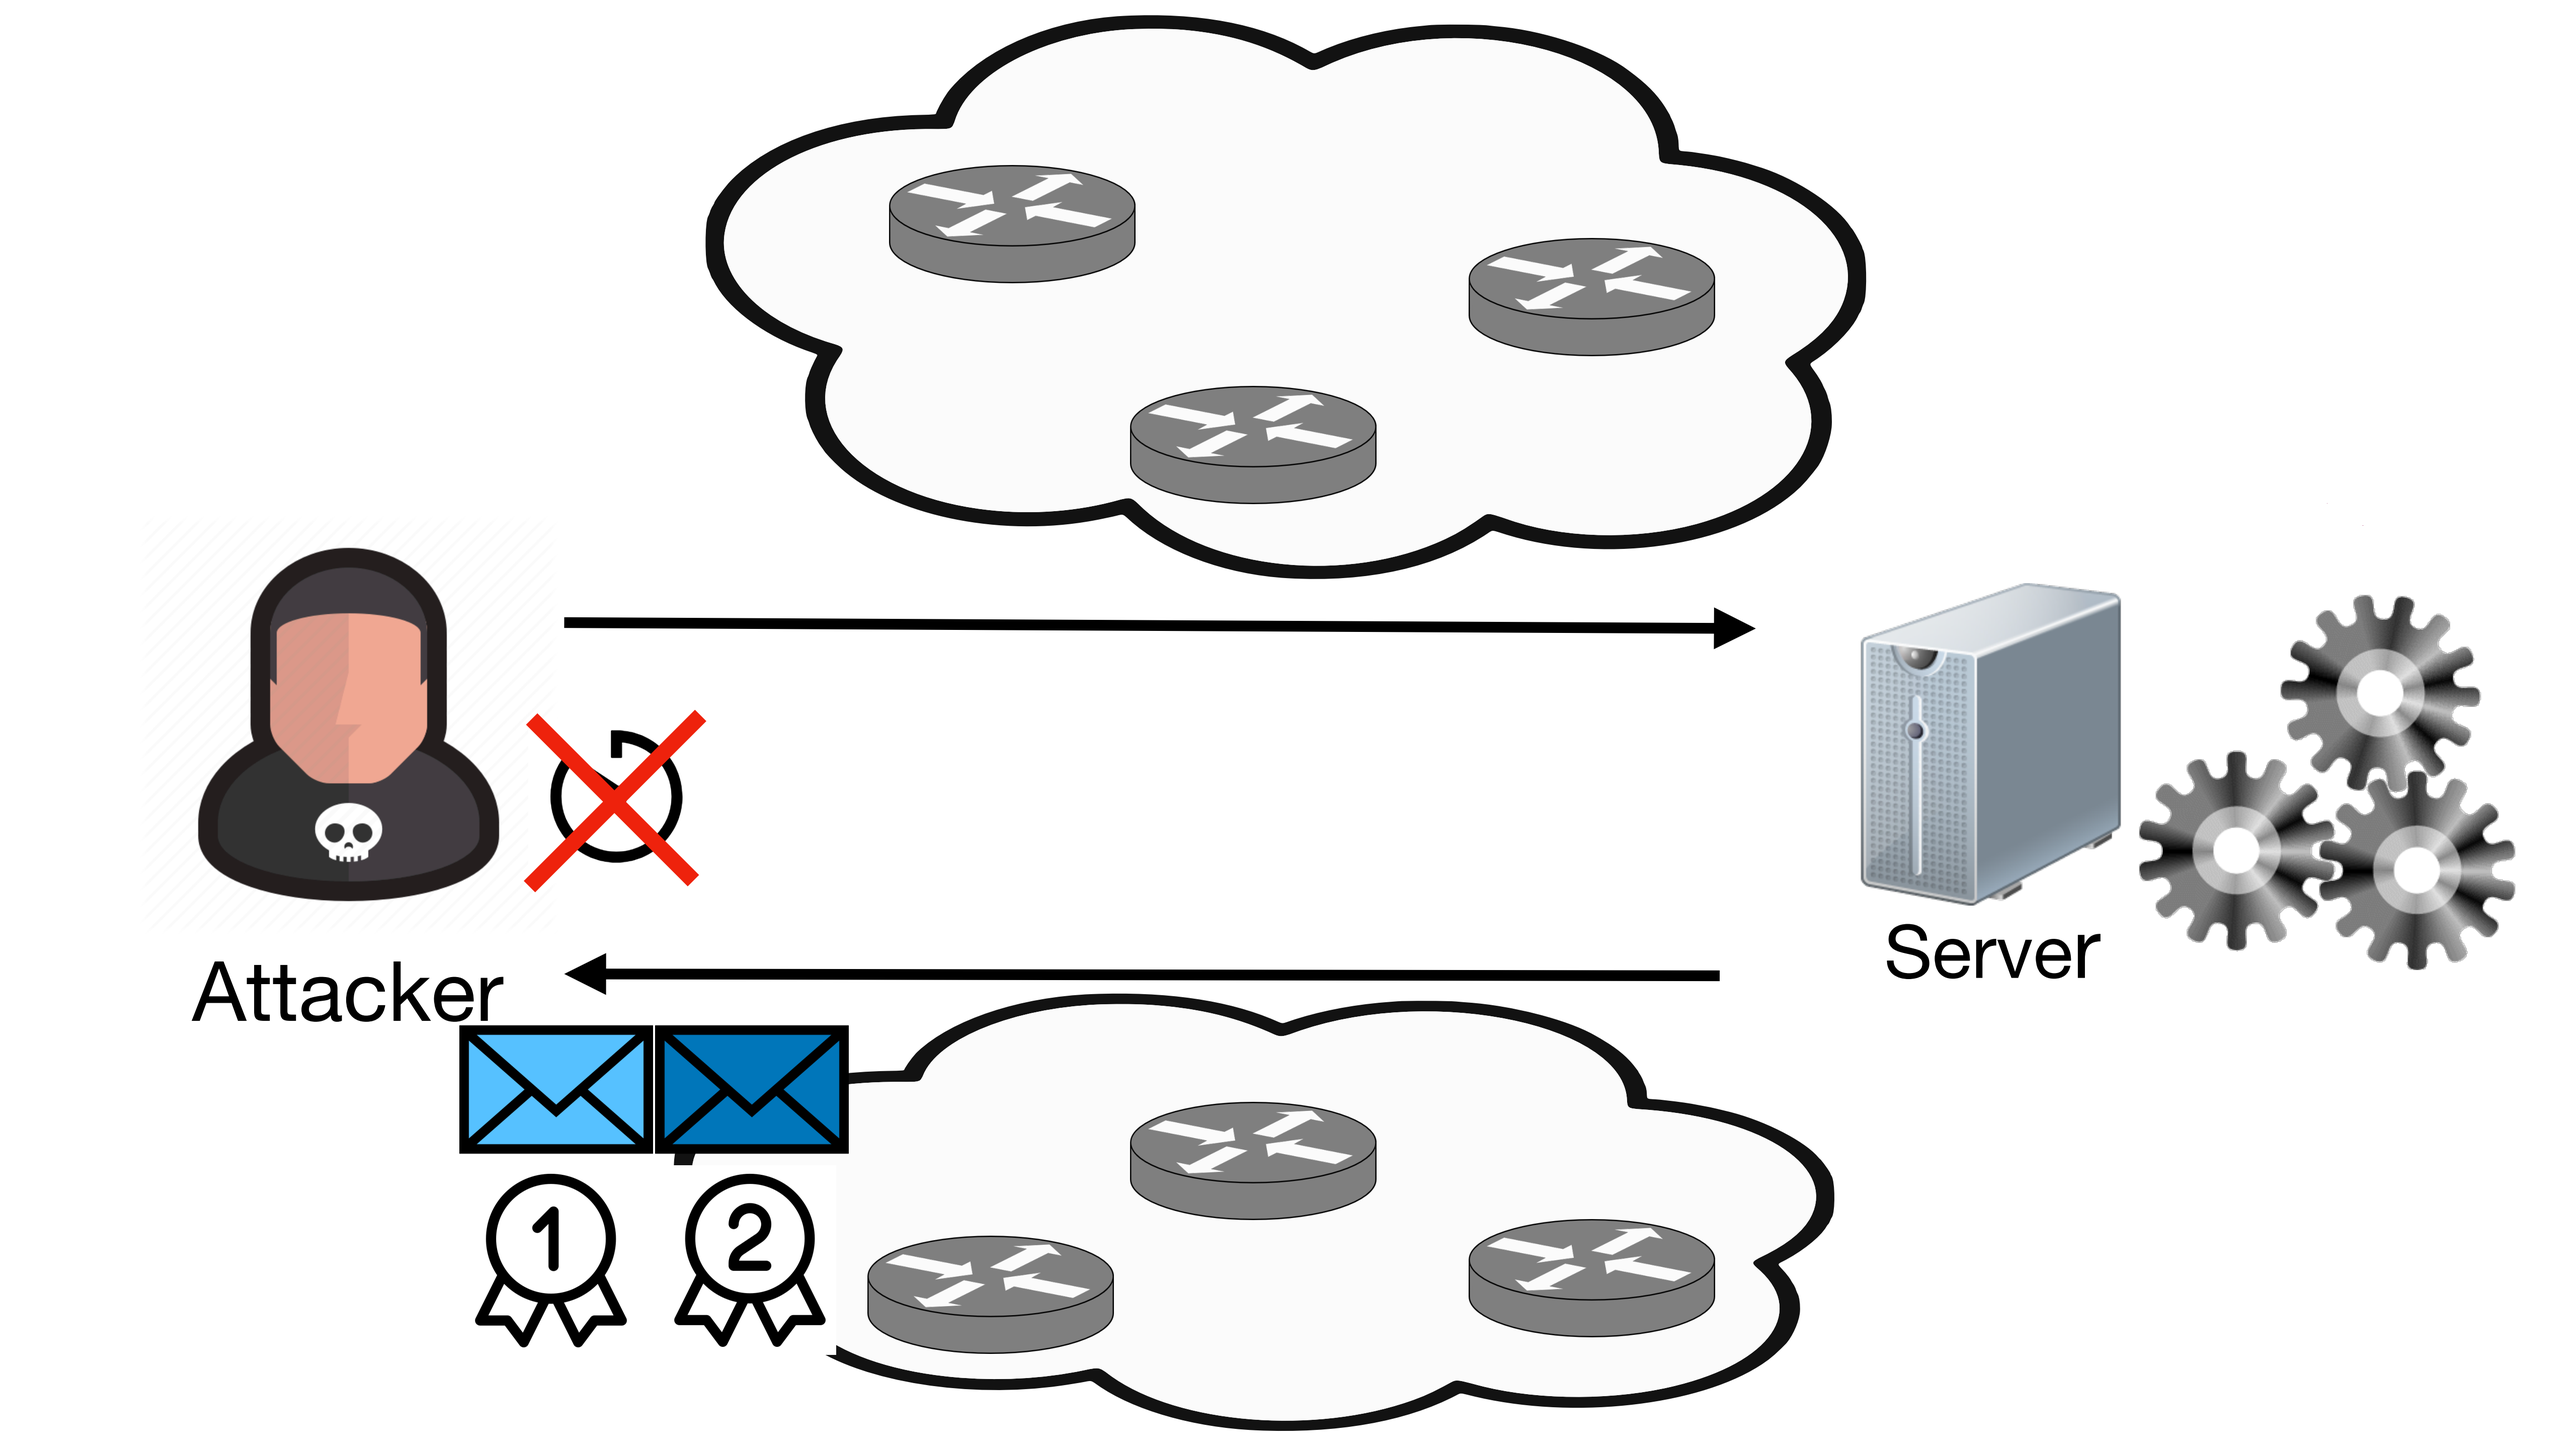
Attacker
Server
1
2

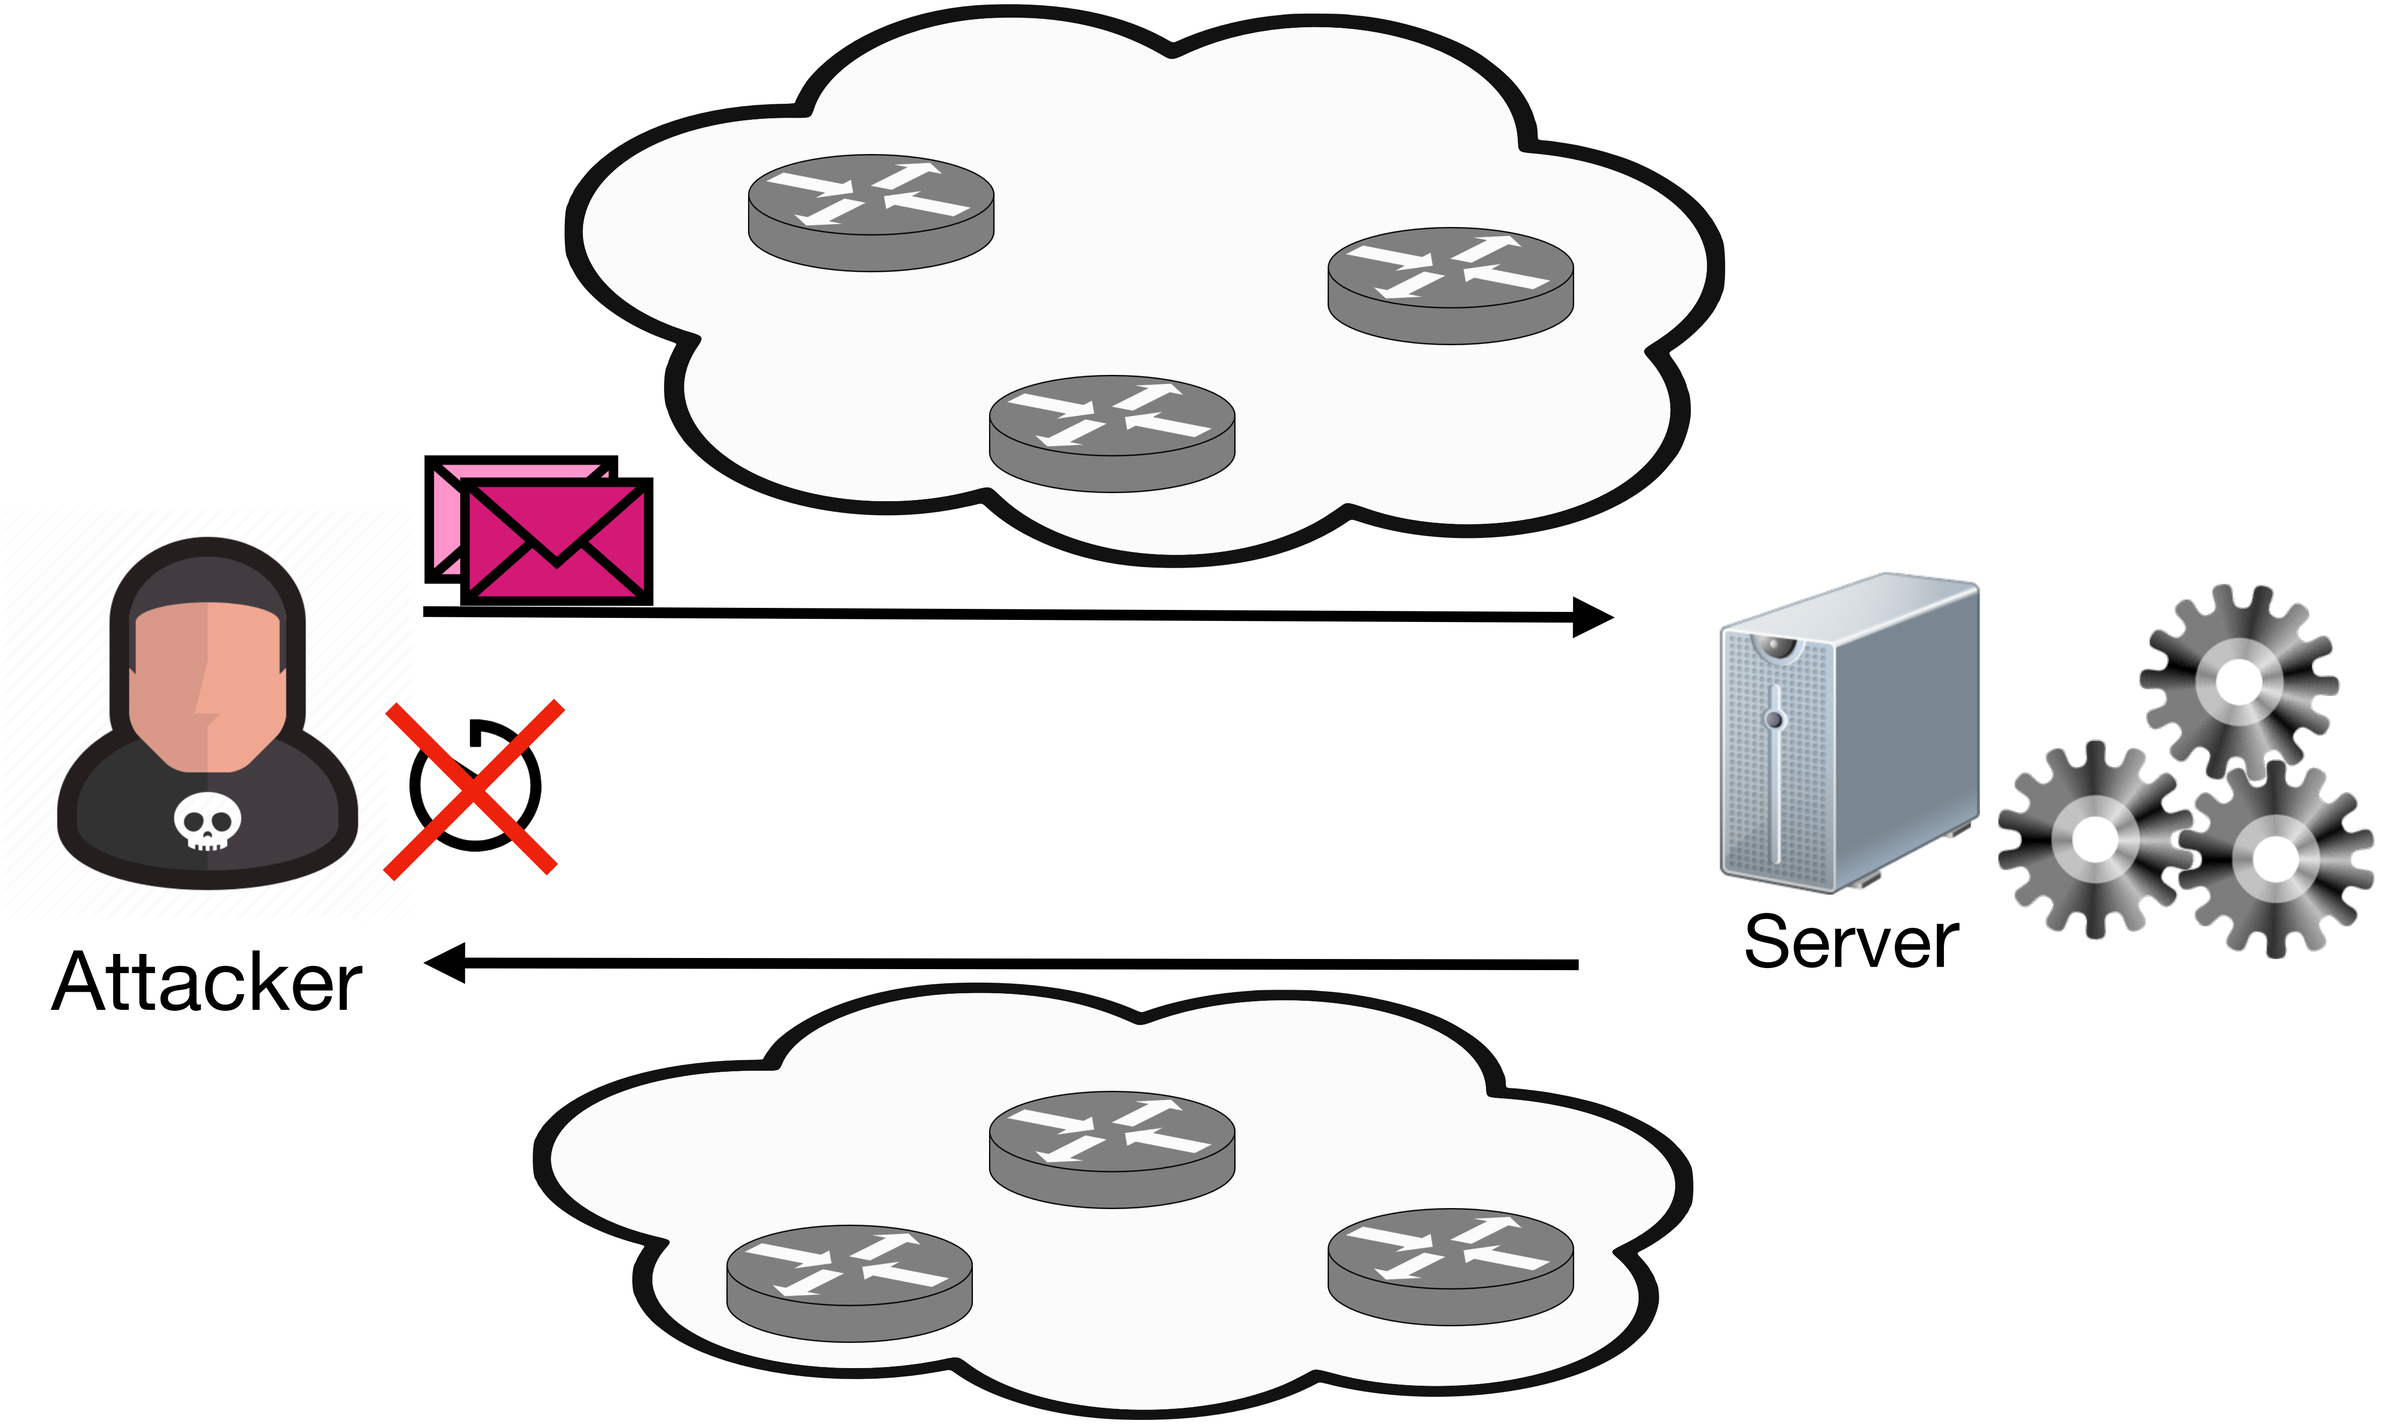
Attacker
Server

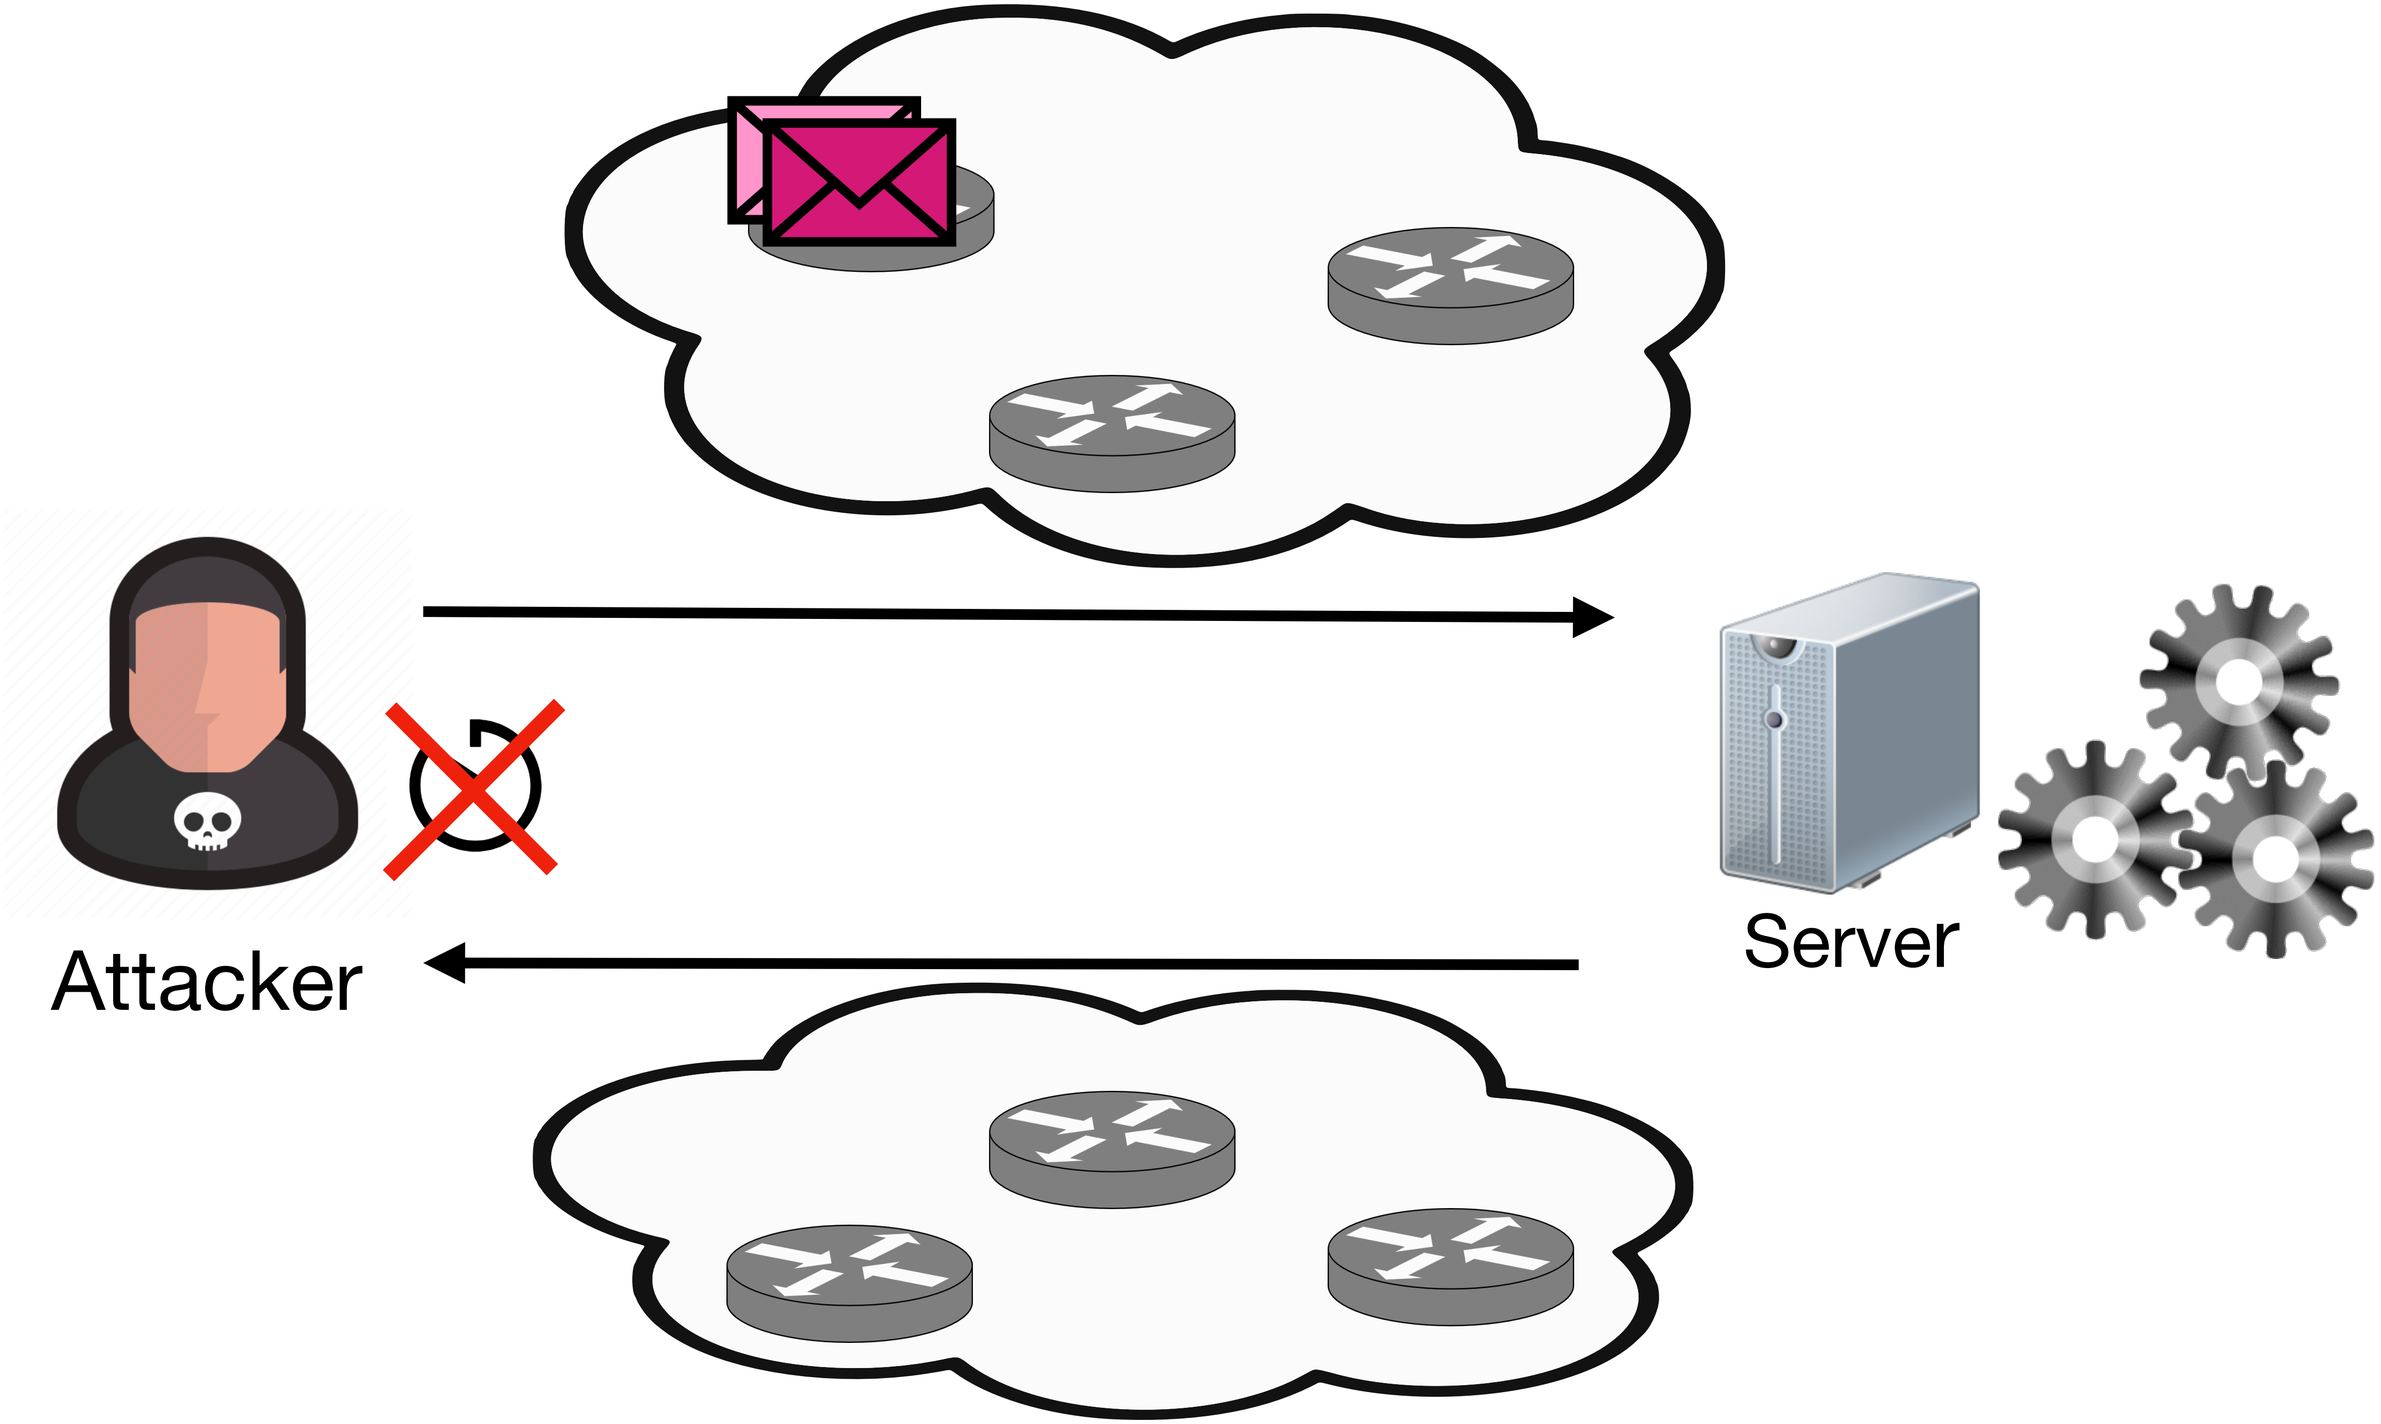
Attacker
Server

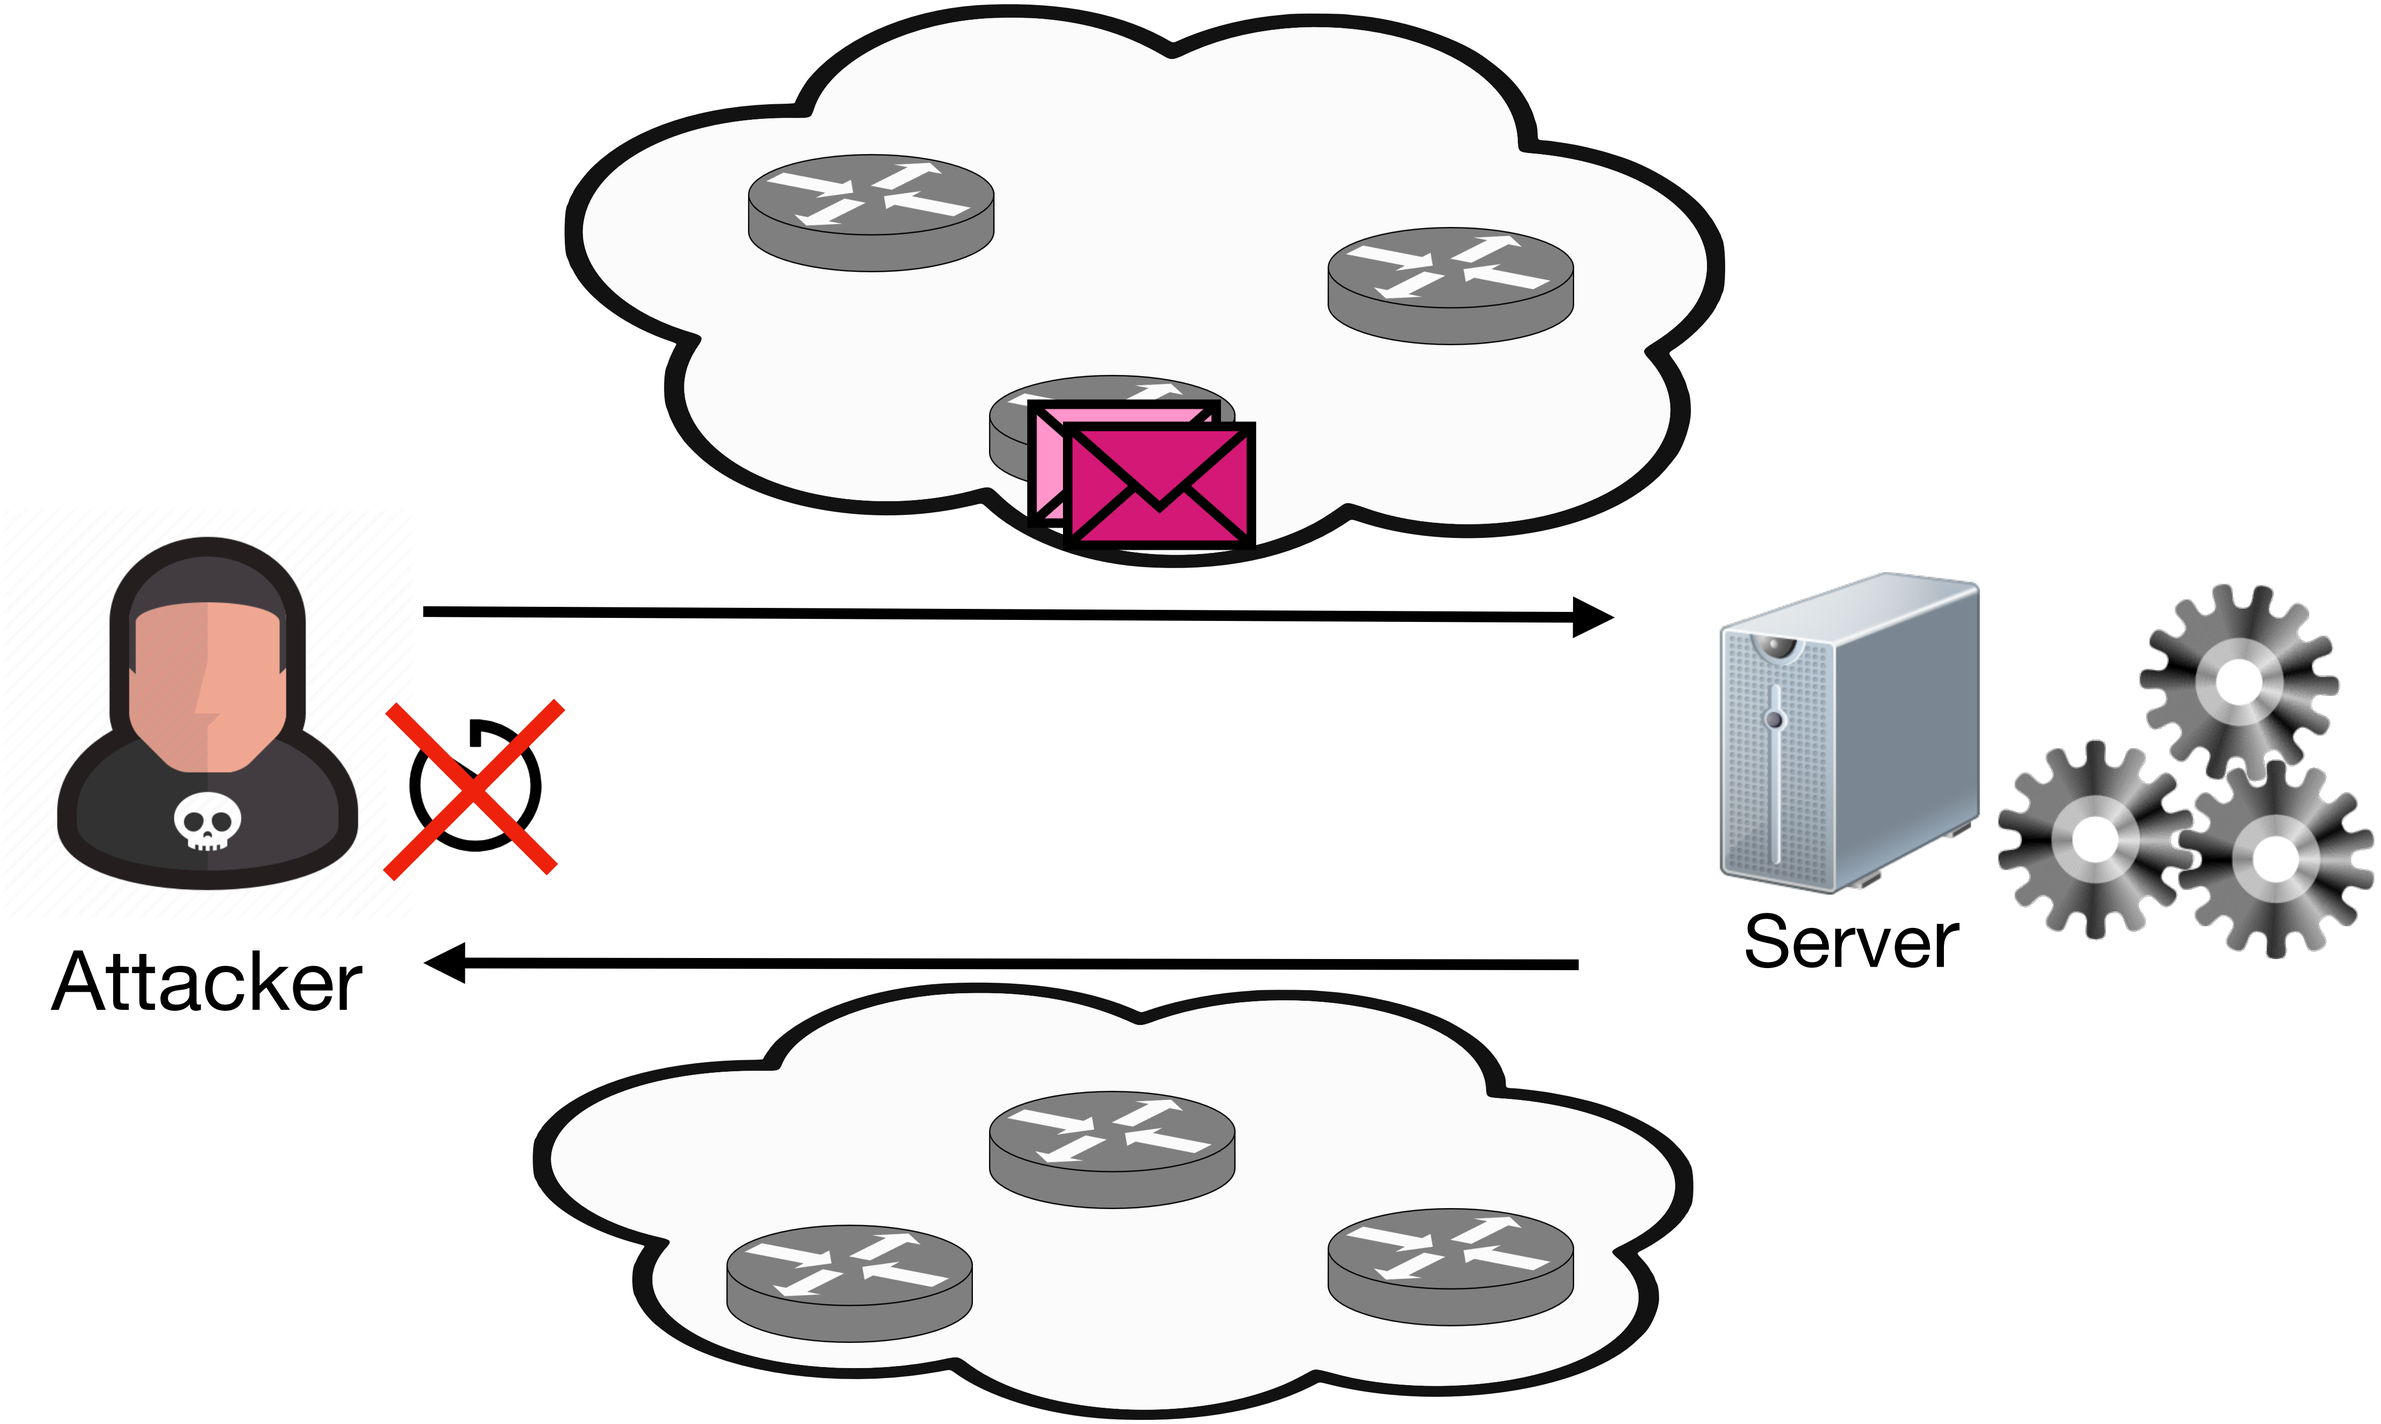
Attacker
Server

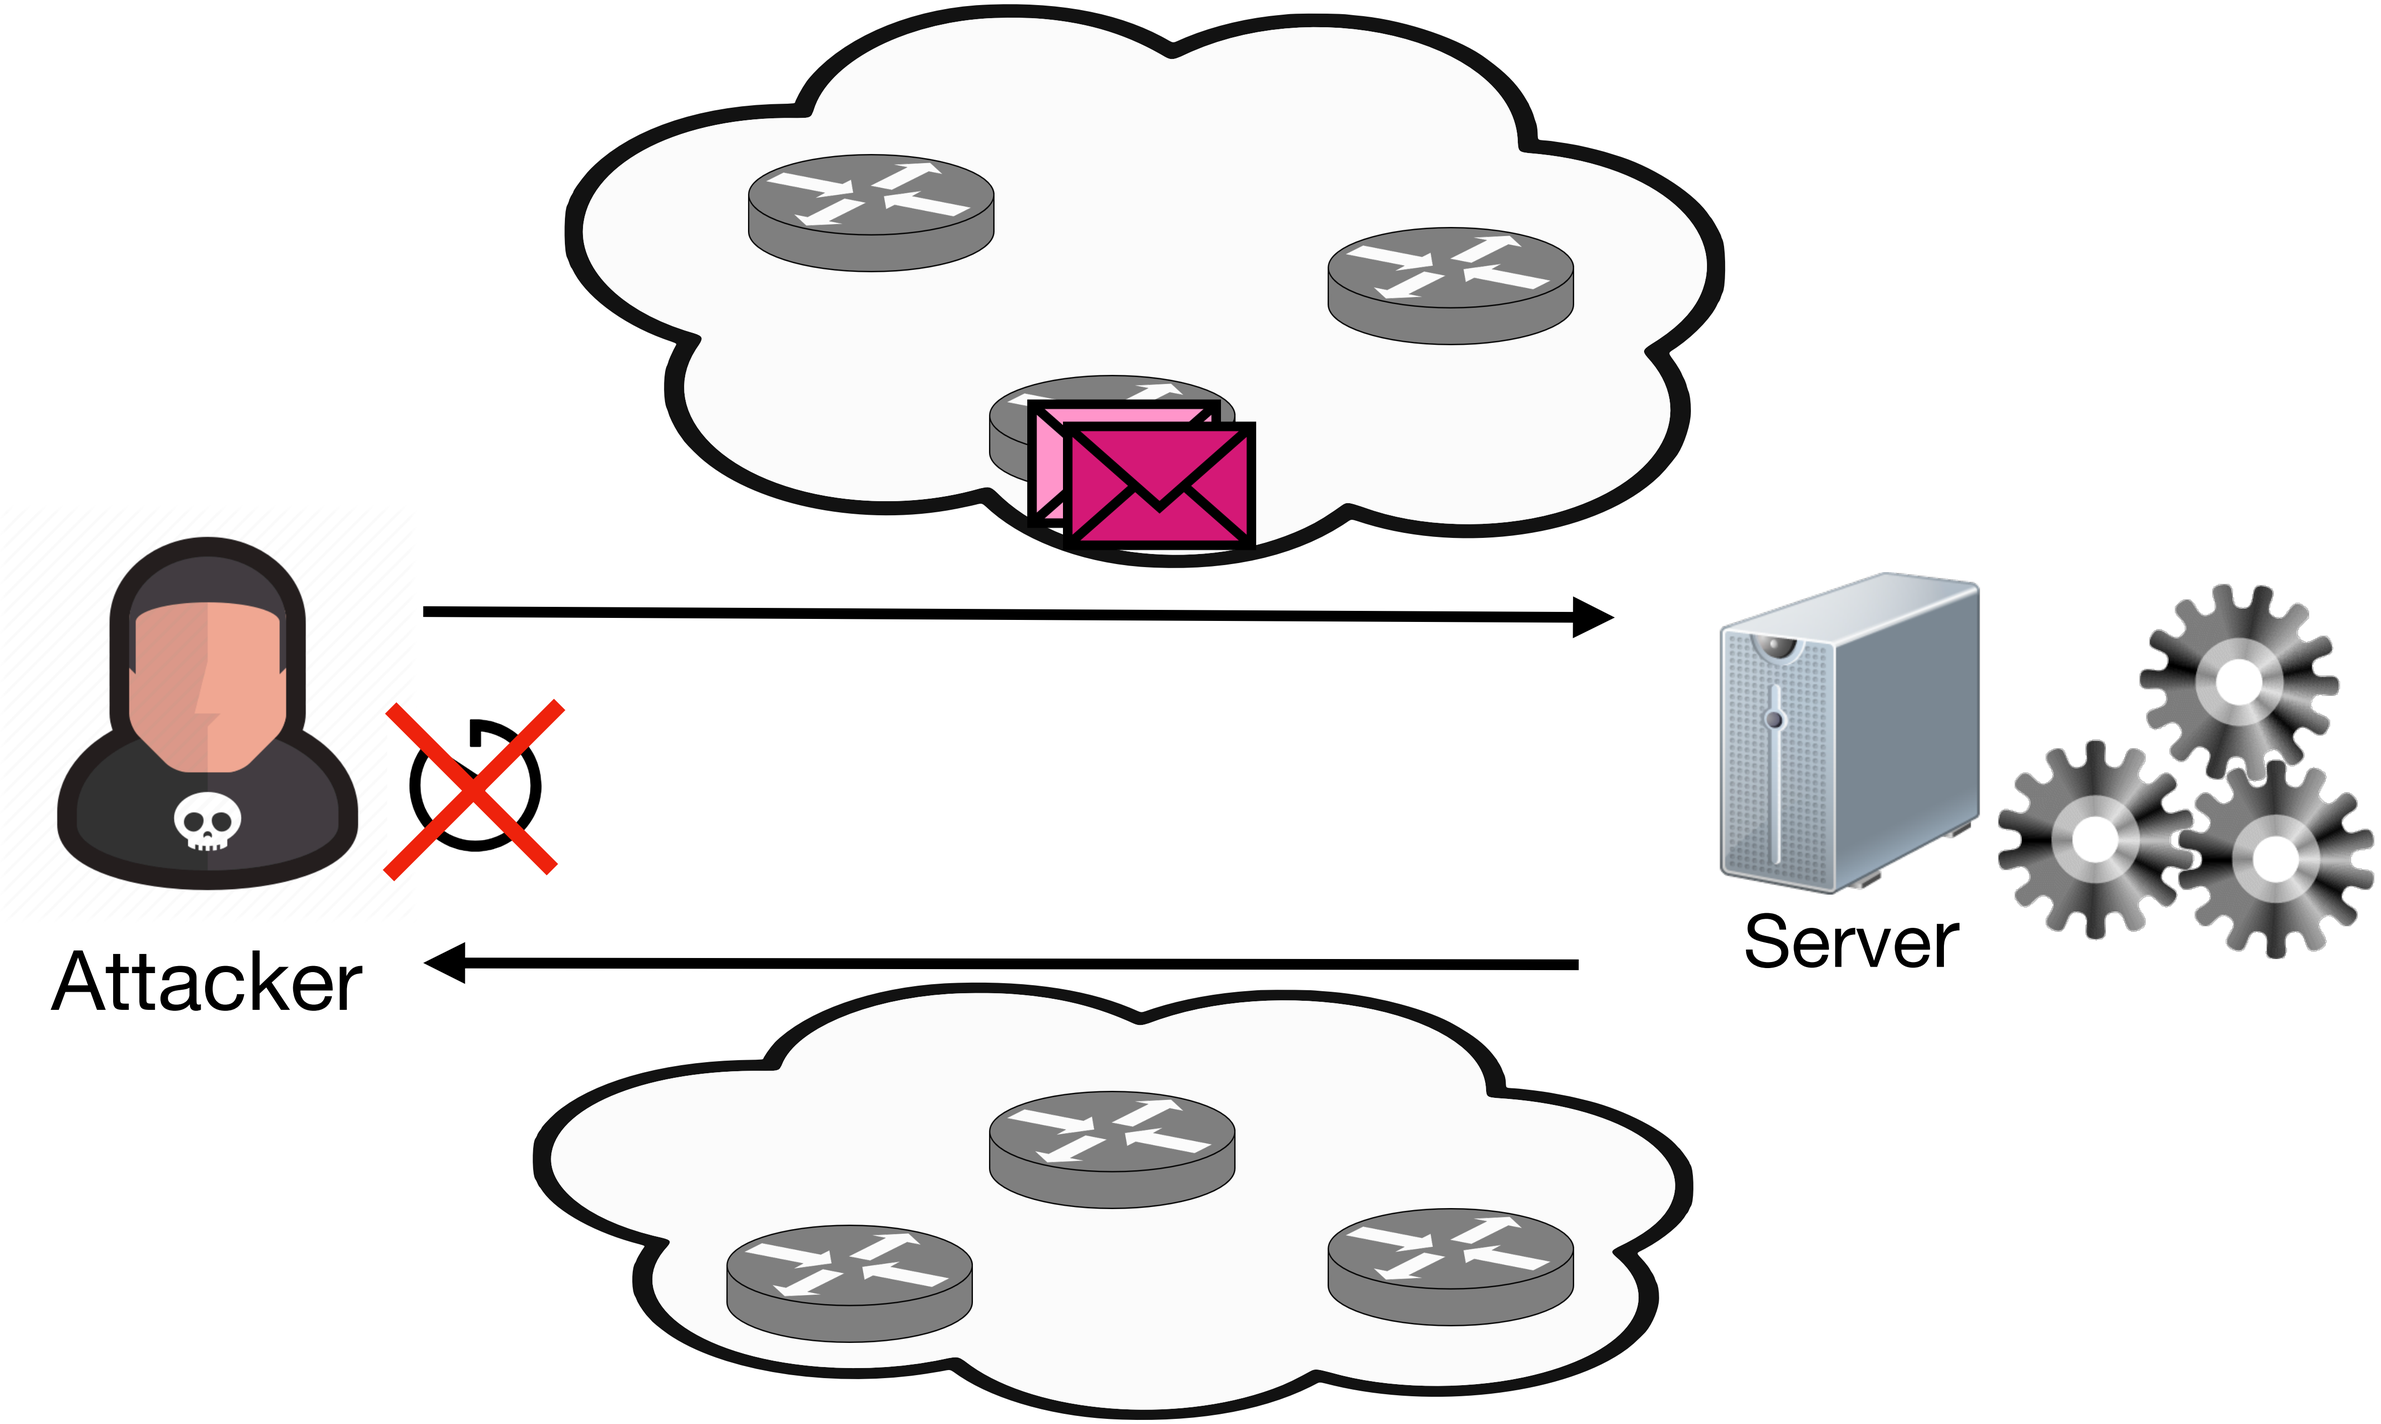
Attacker
Server

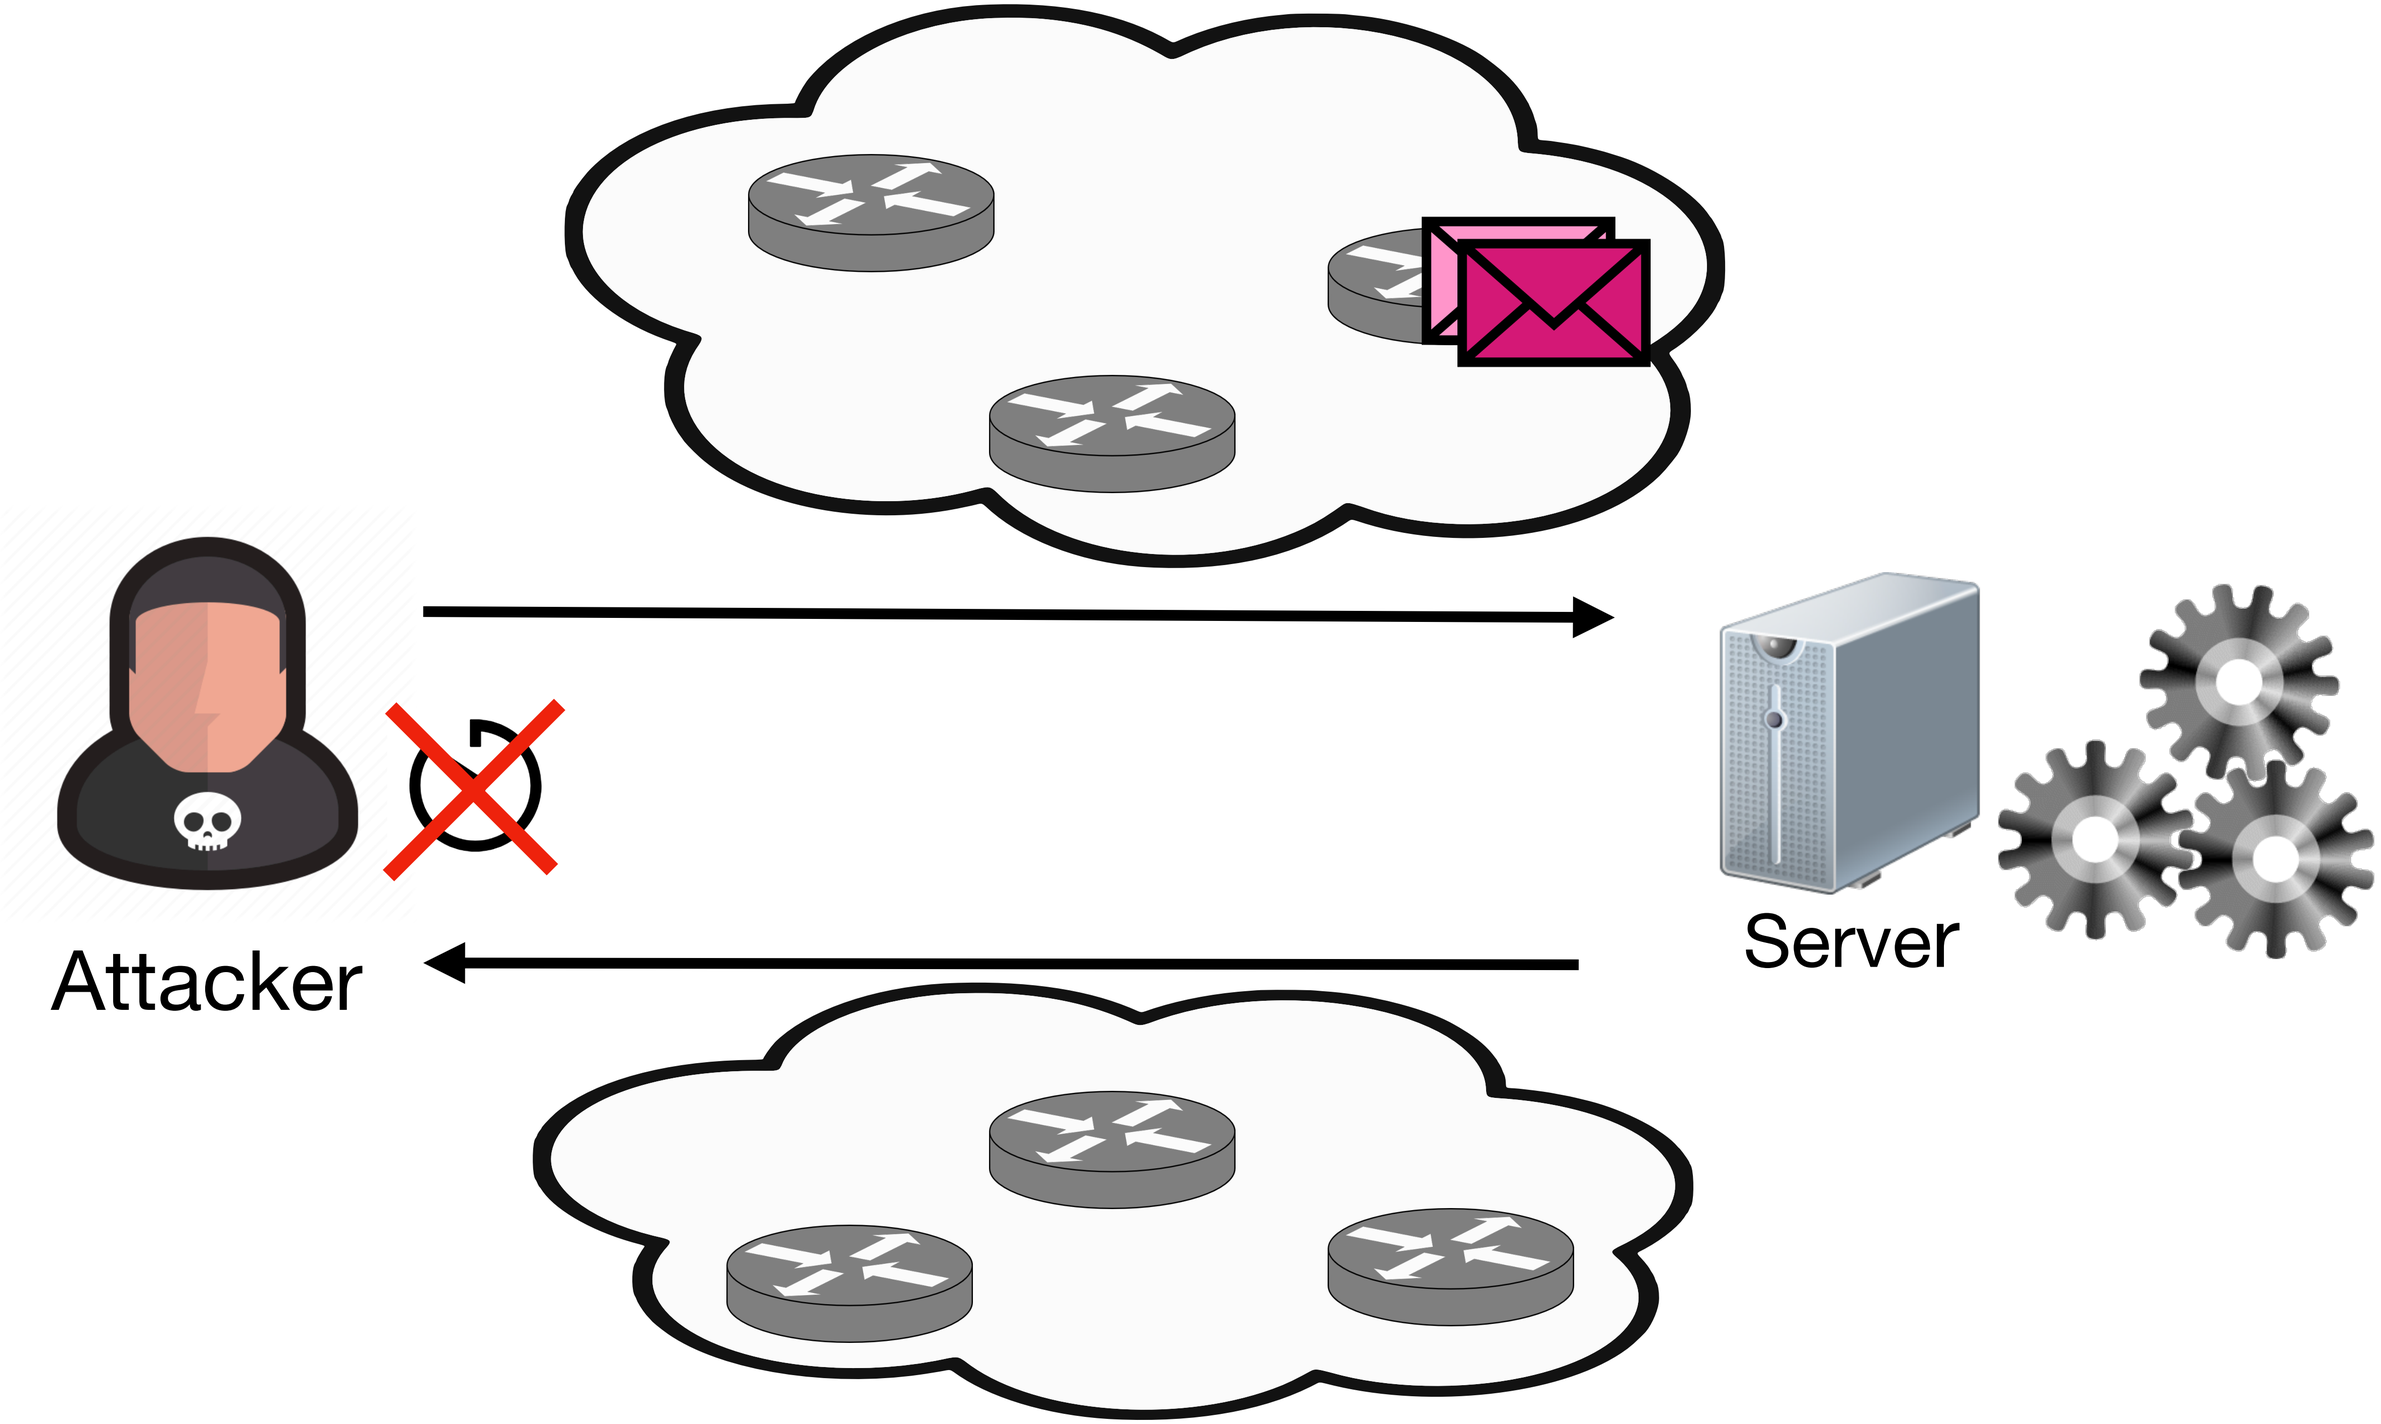
Attacker
Server

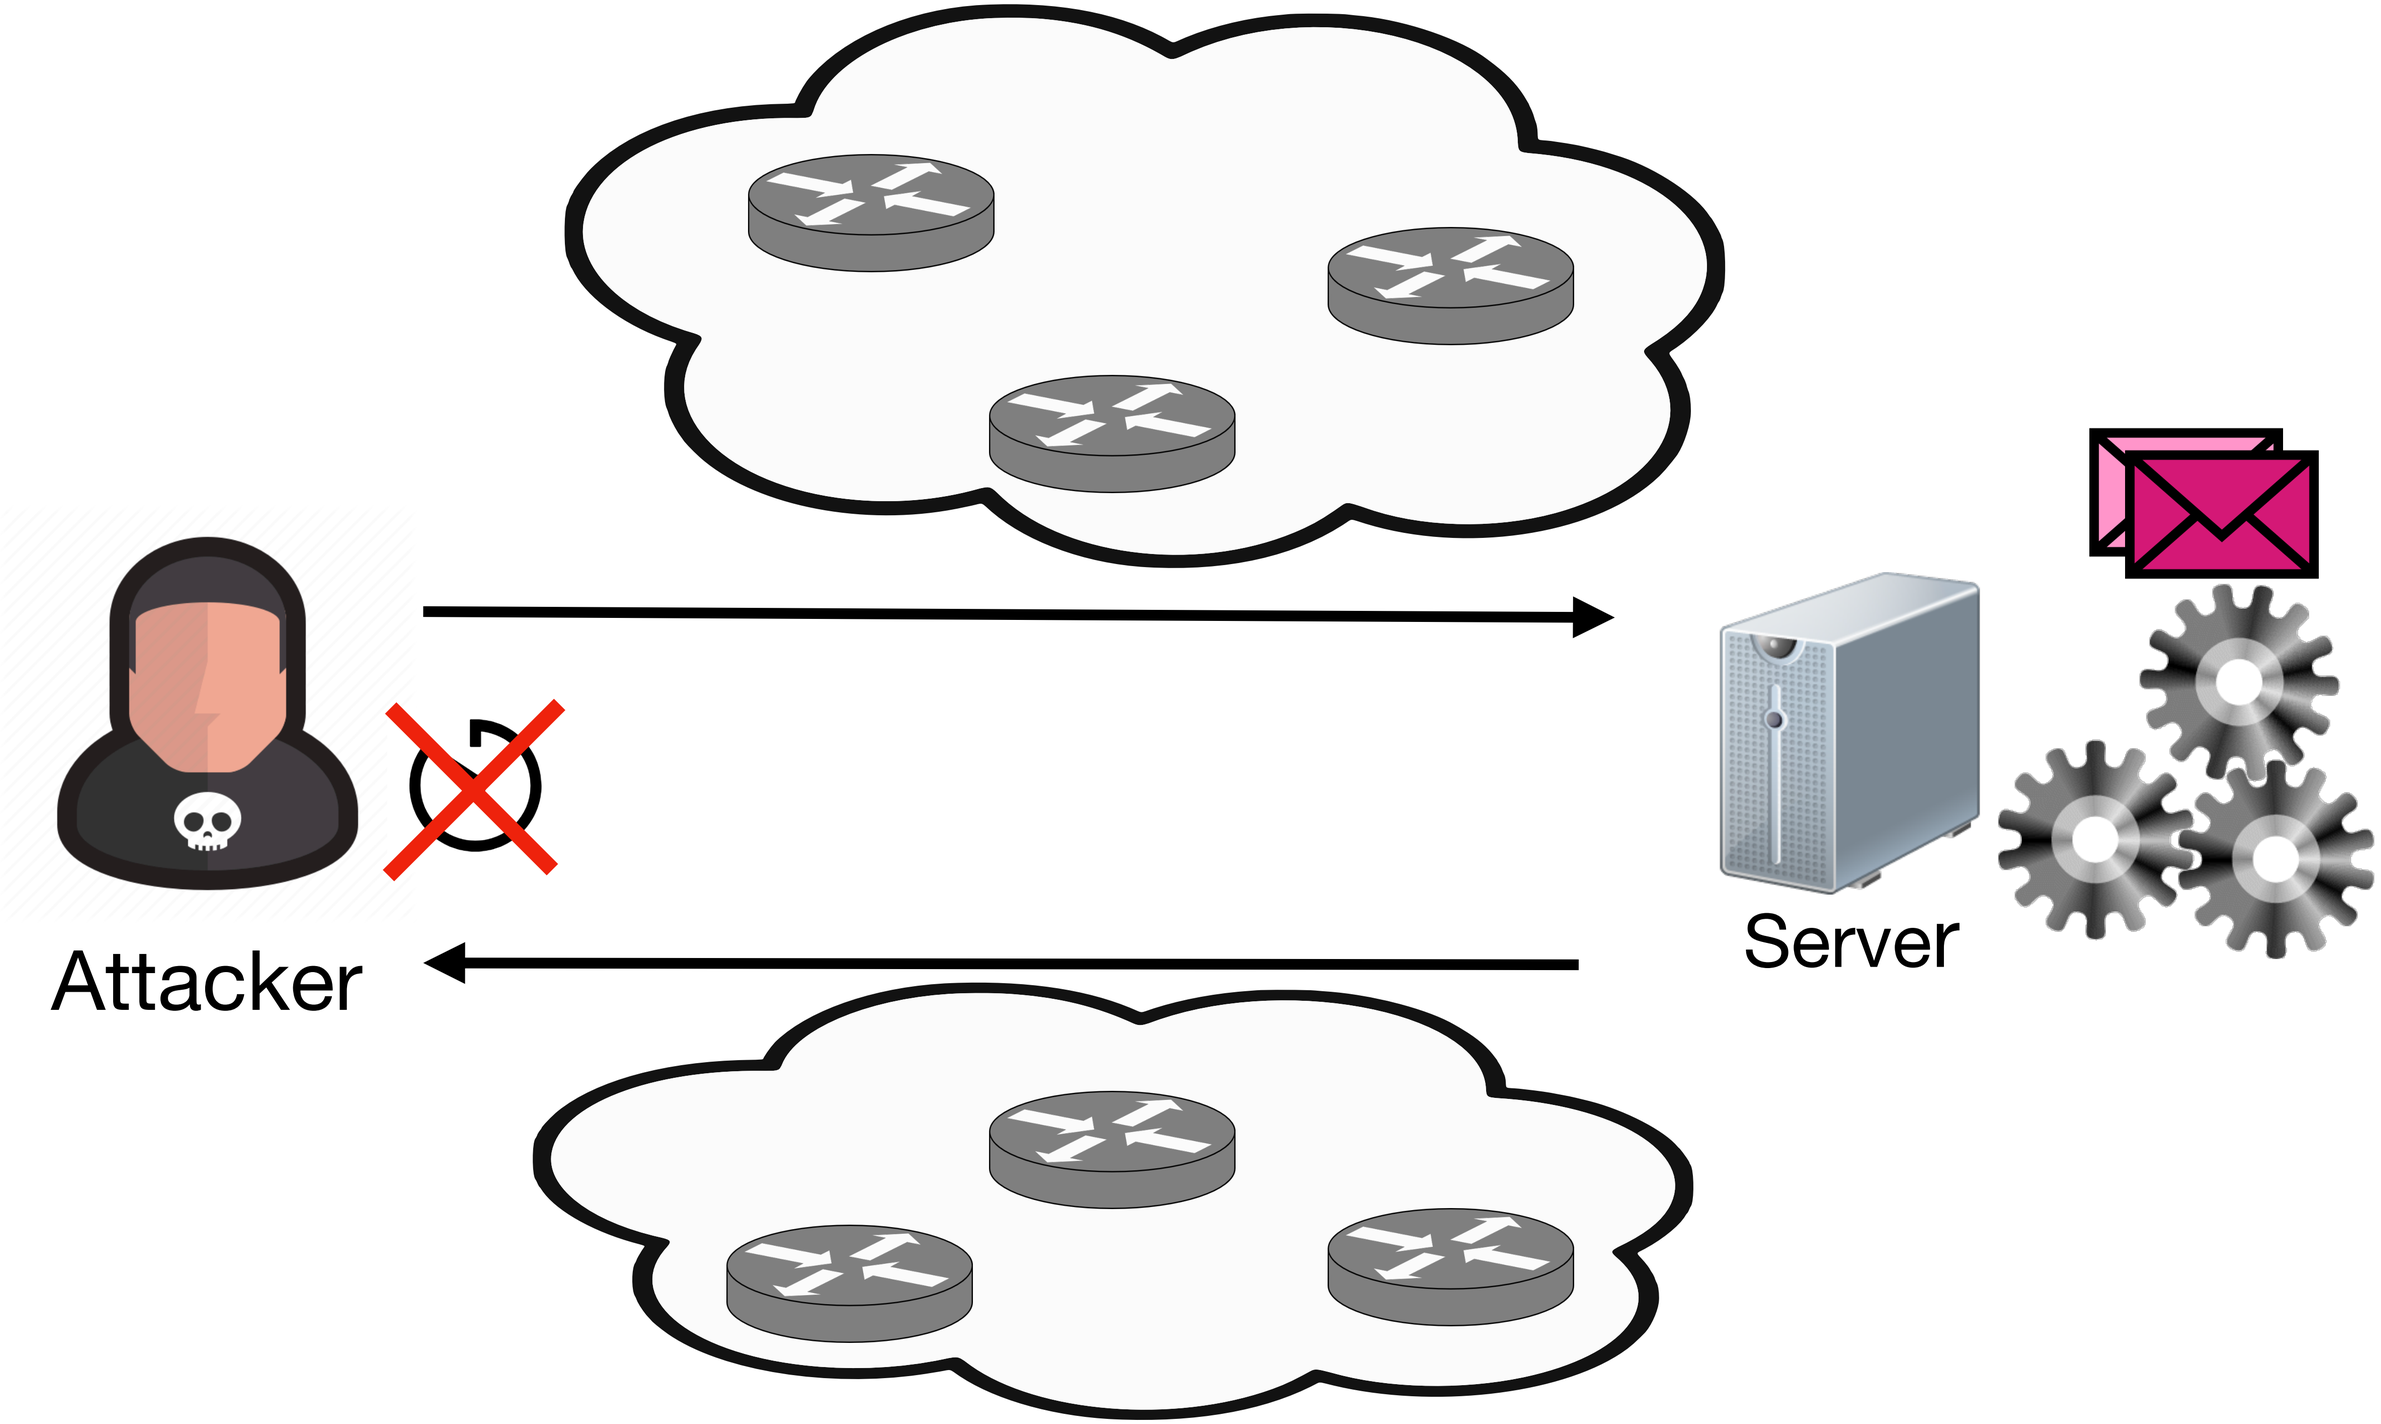
Attacker
Server

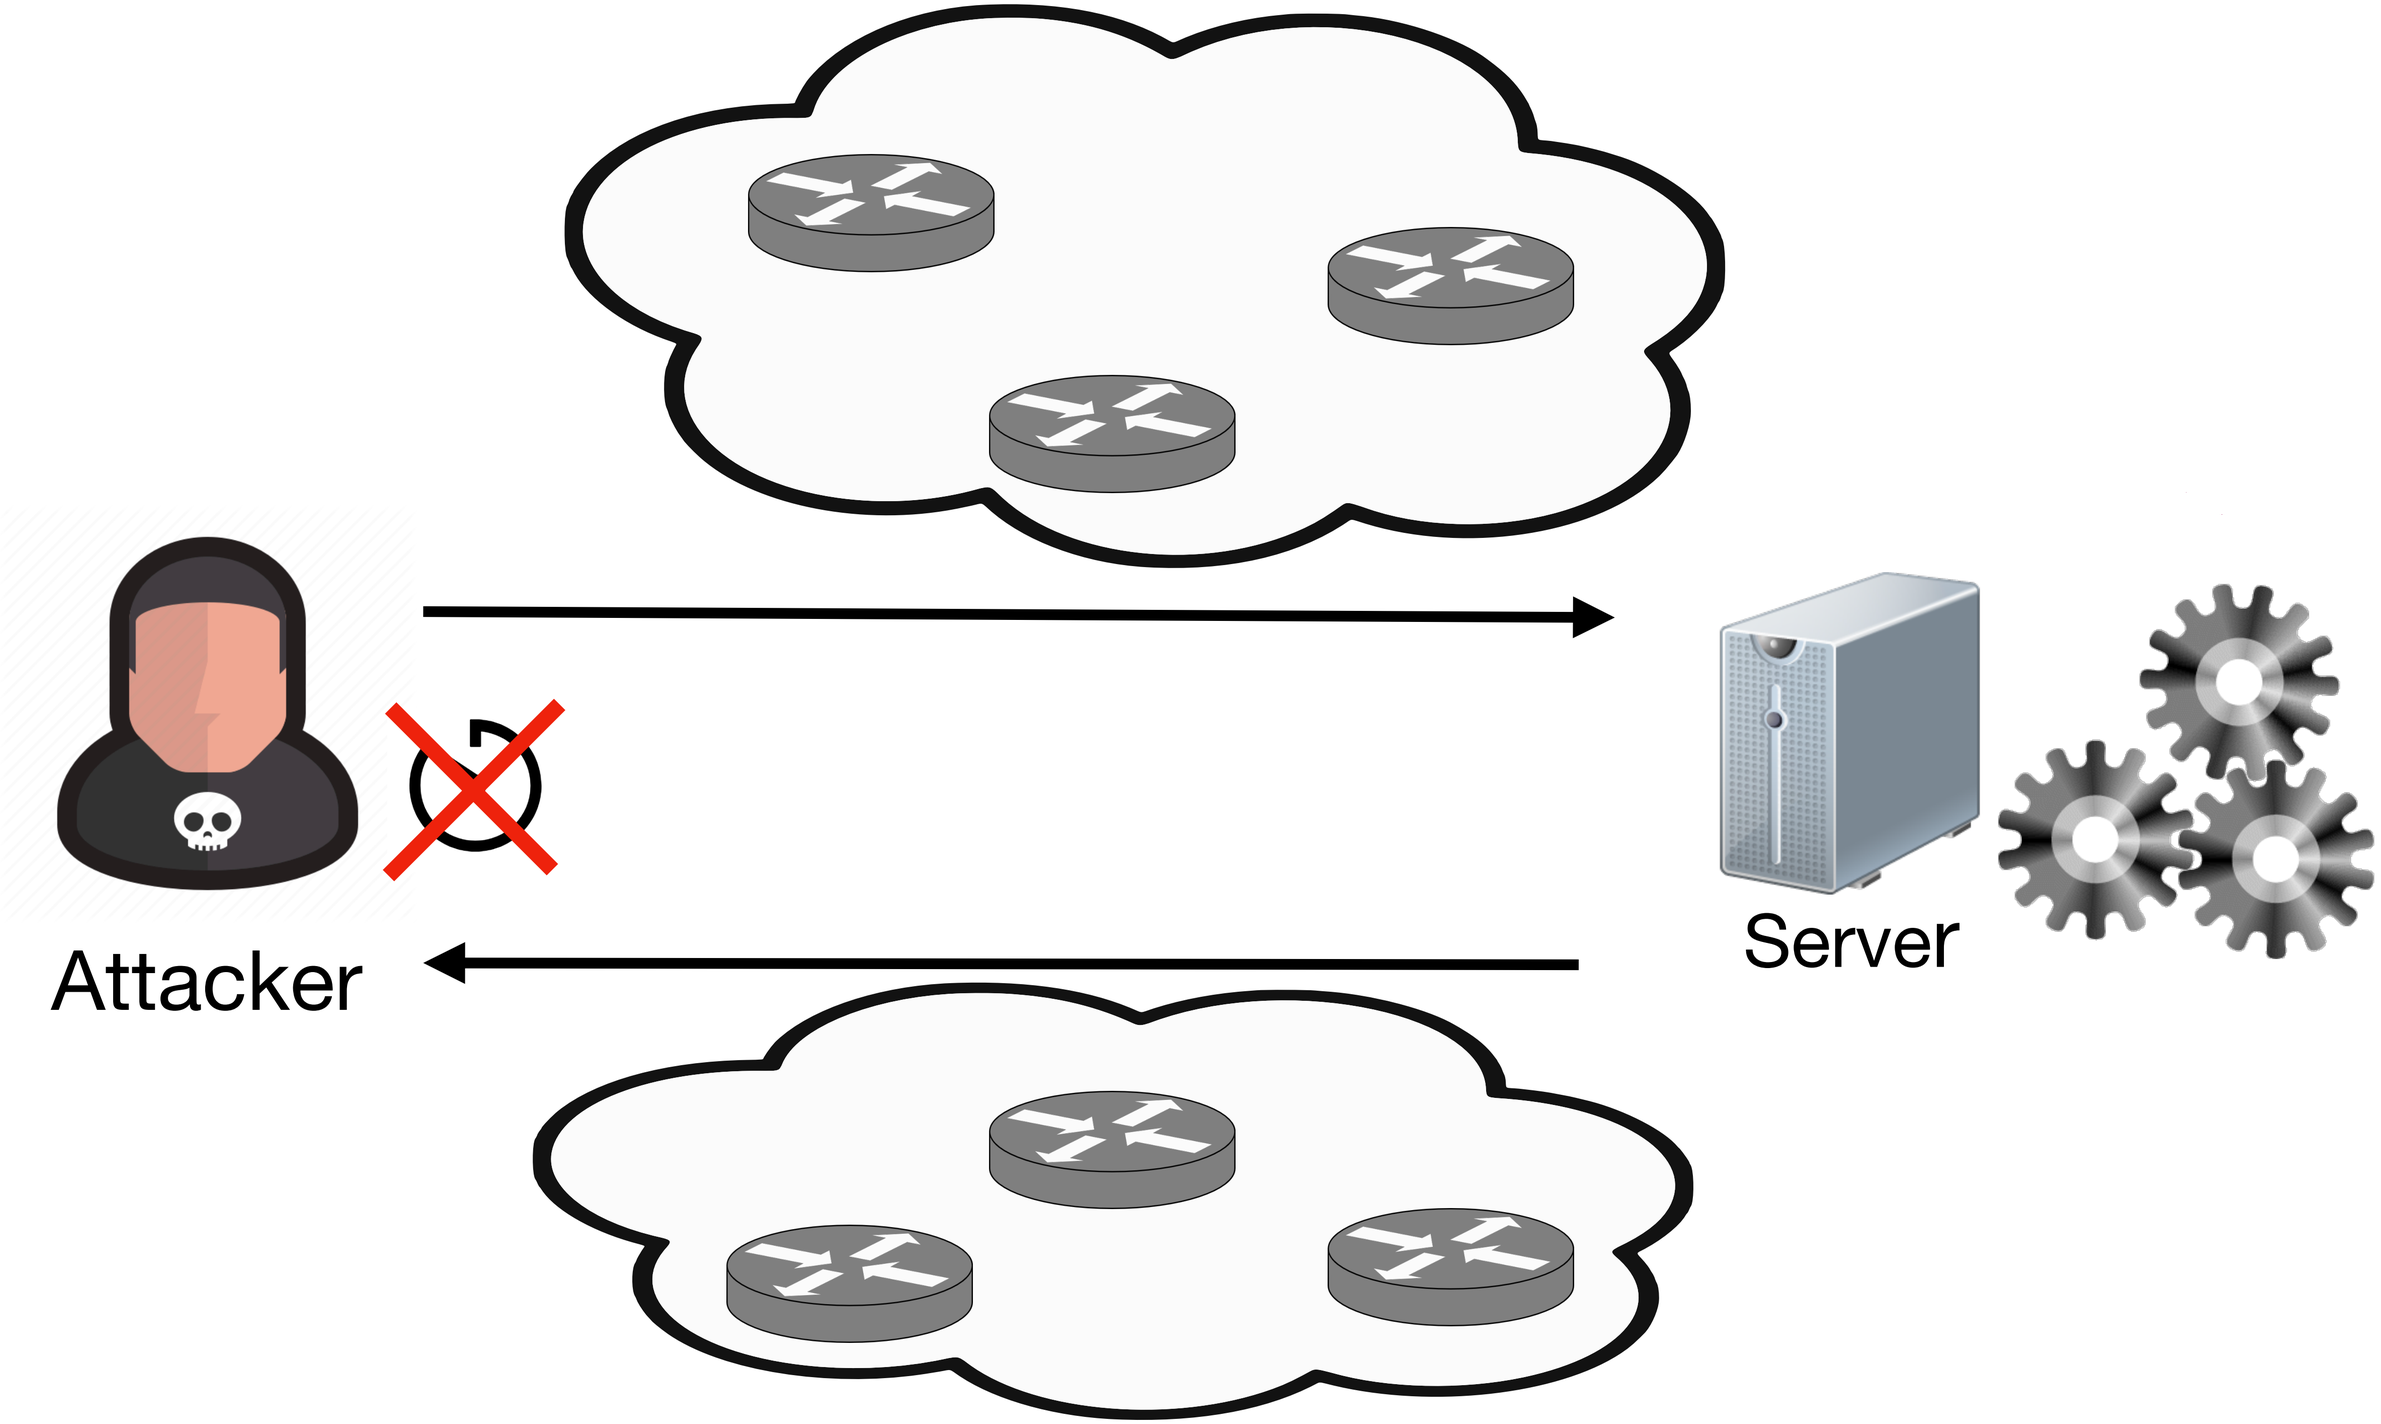
Attacker
Server

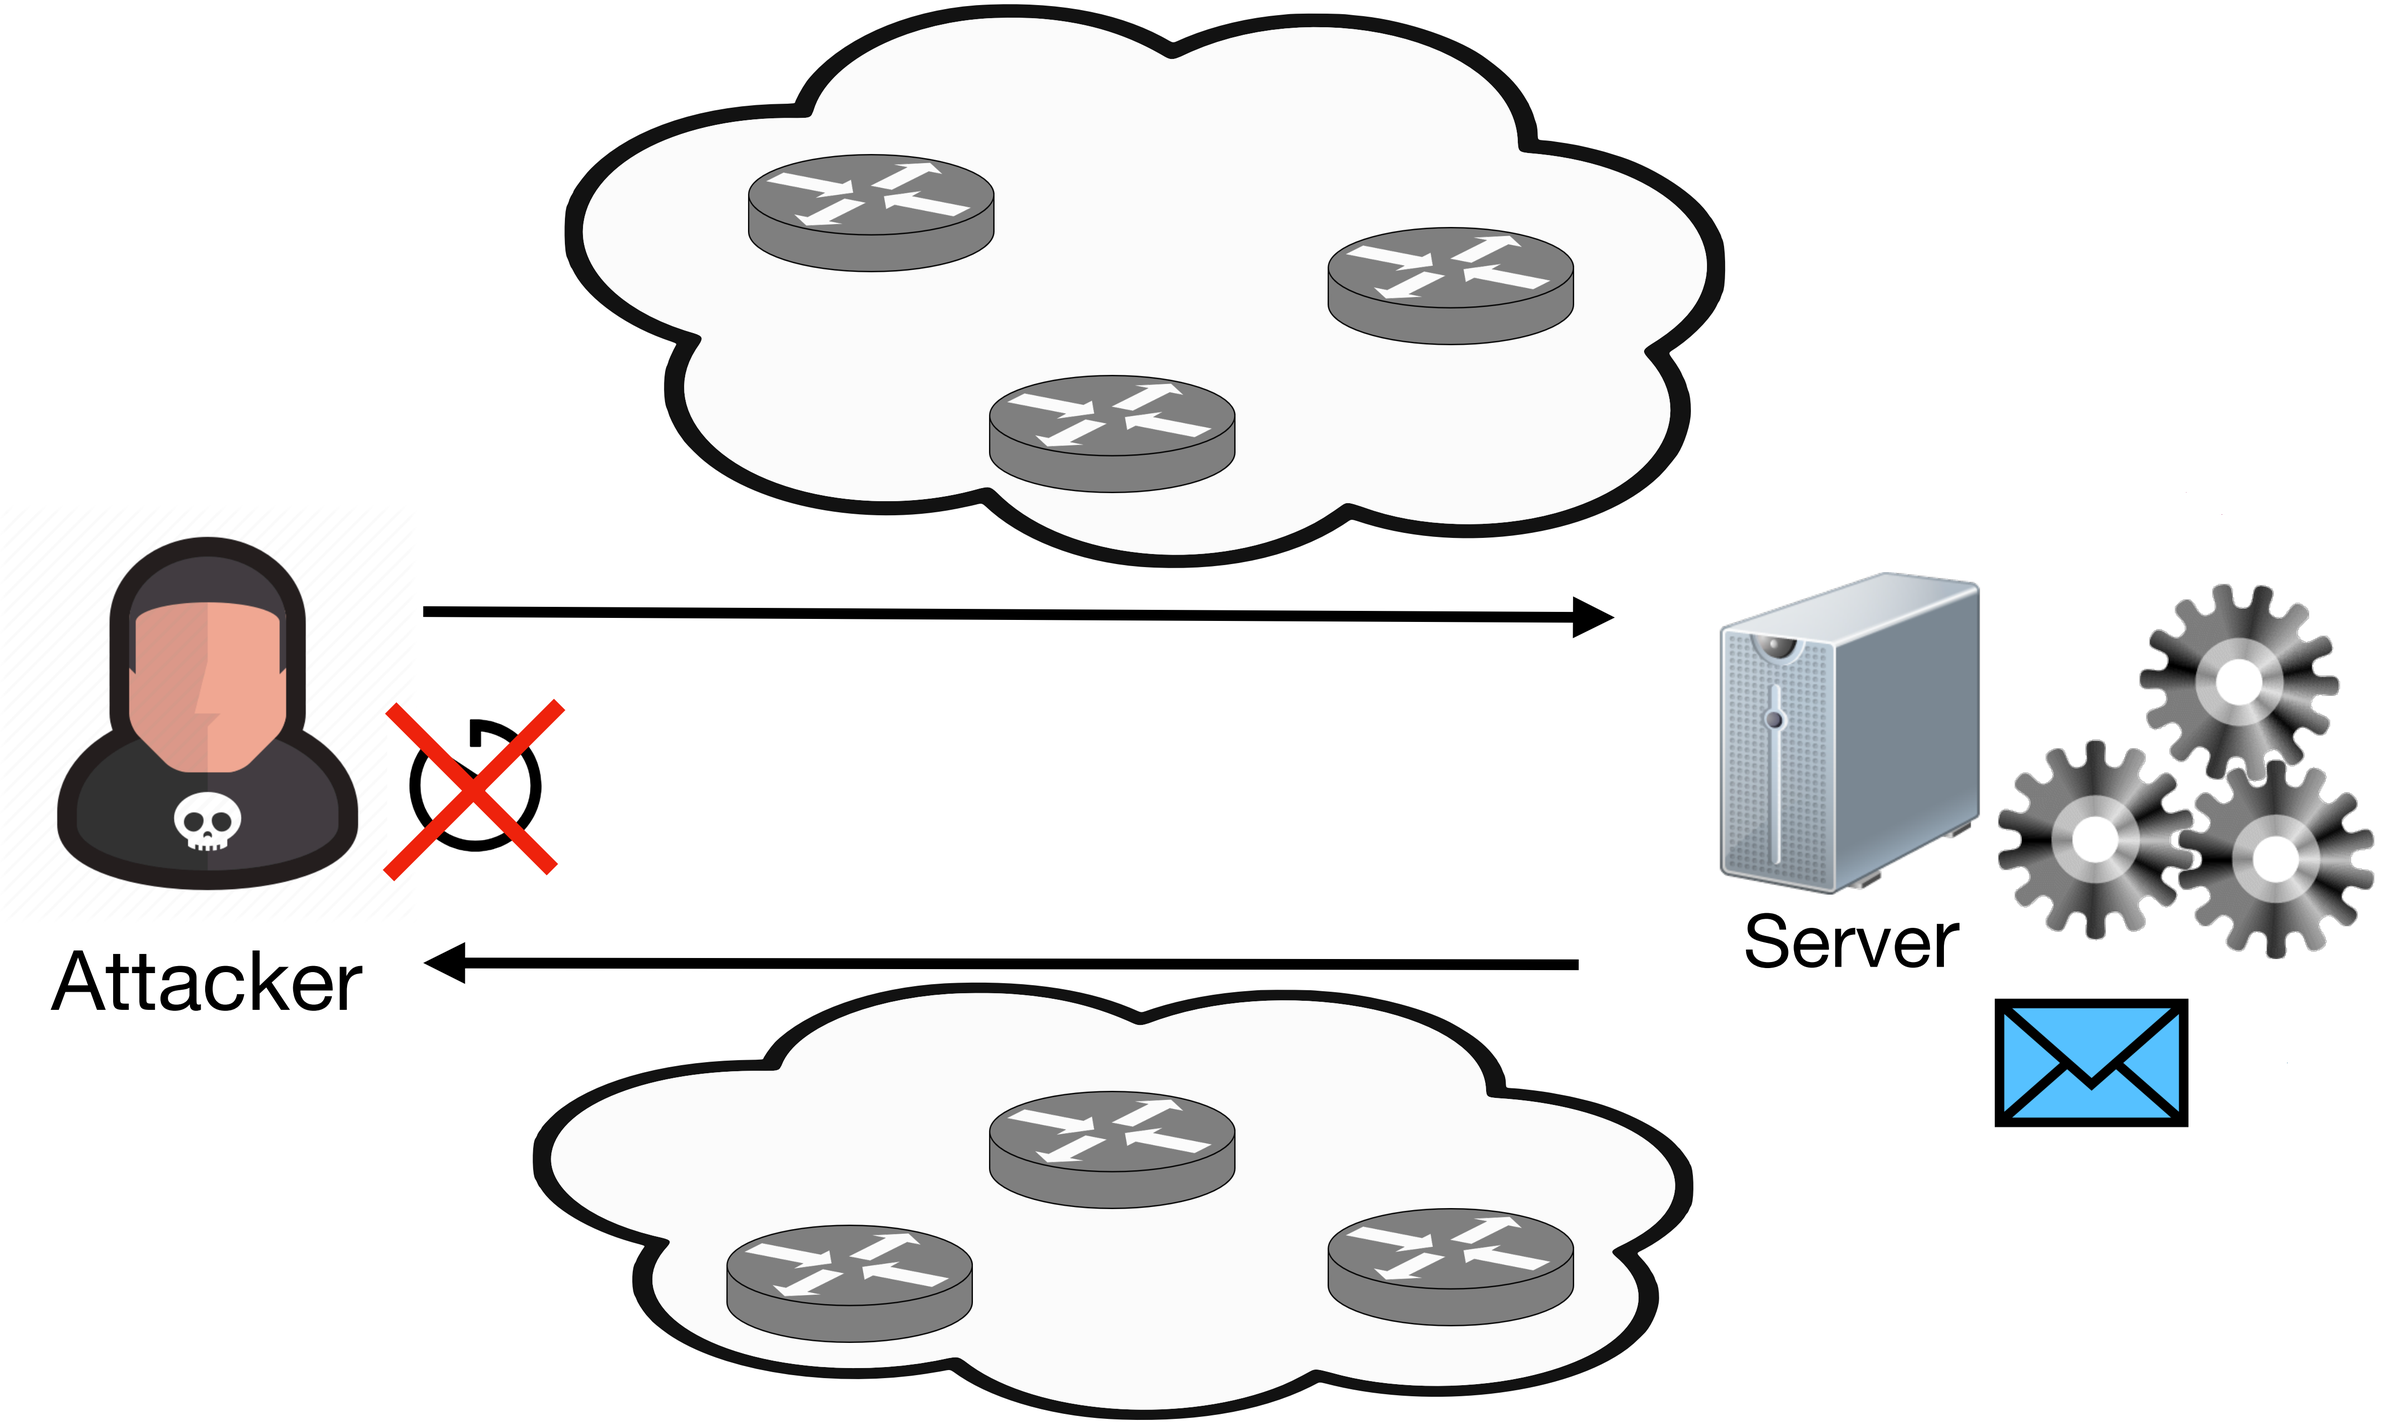
Attacker
Server

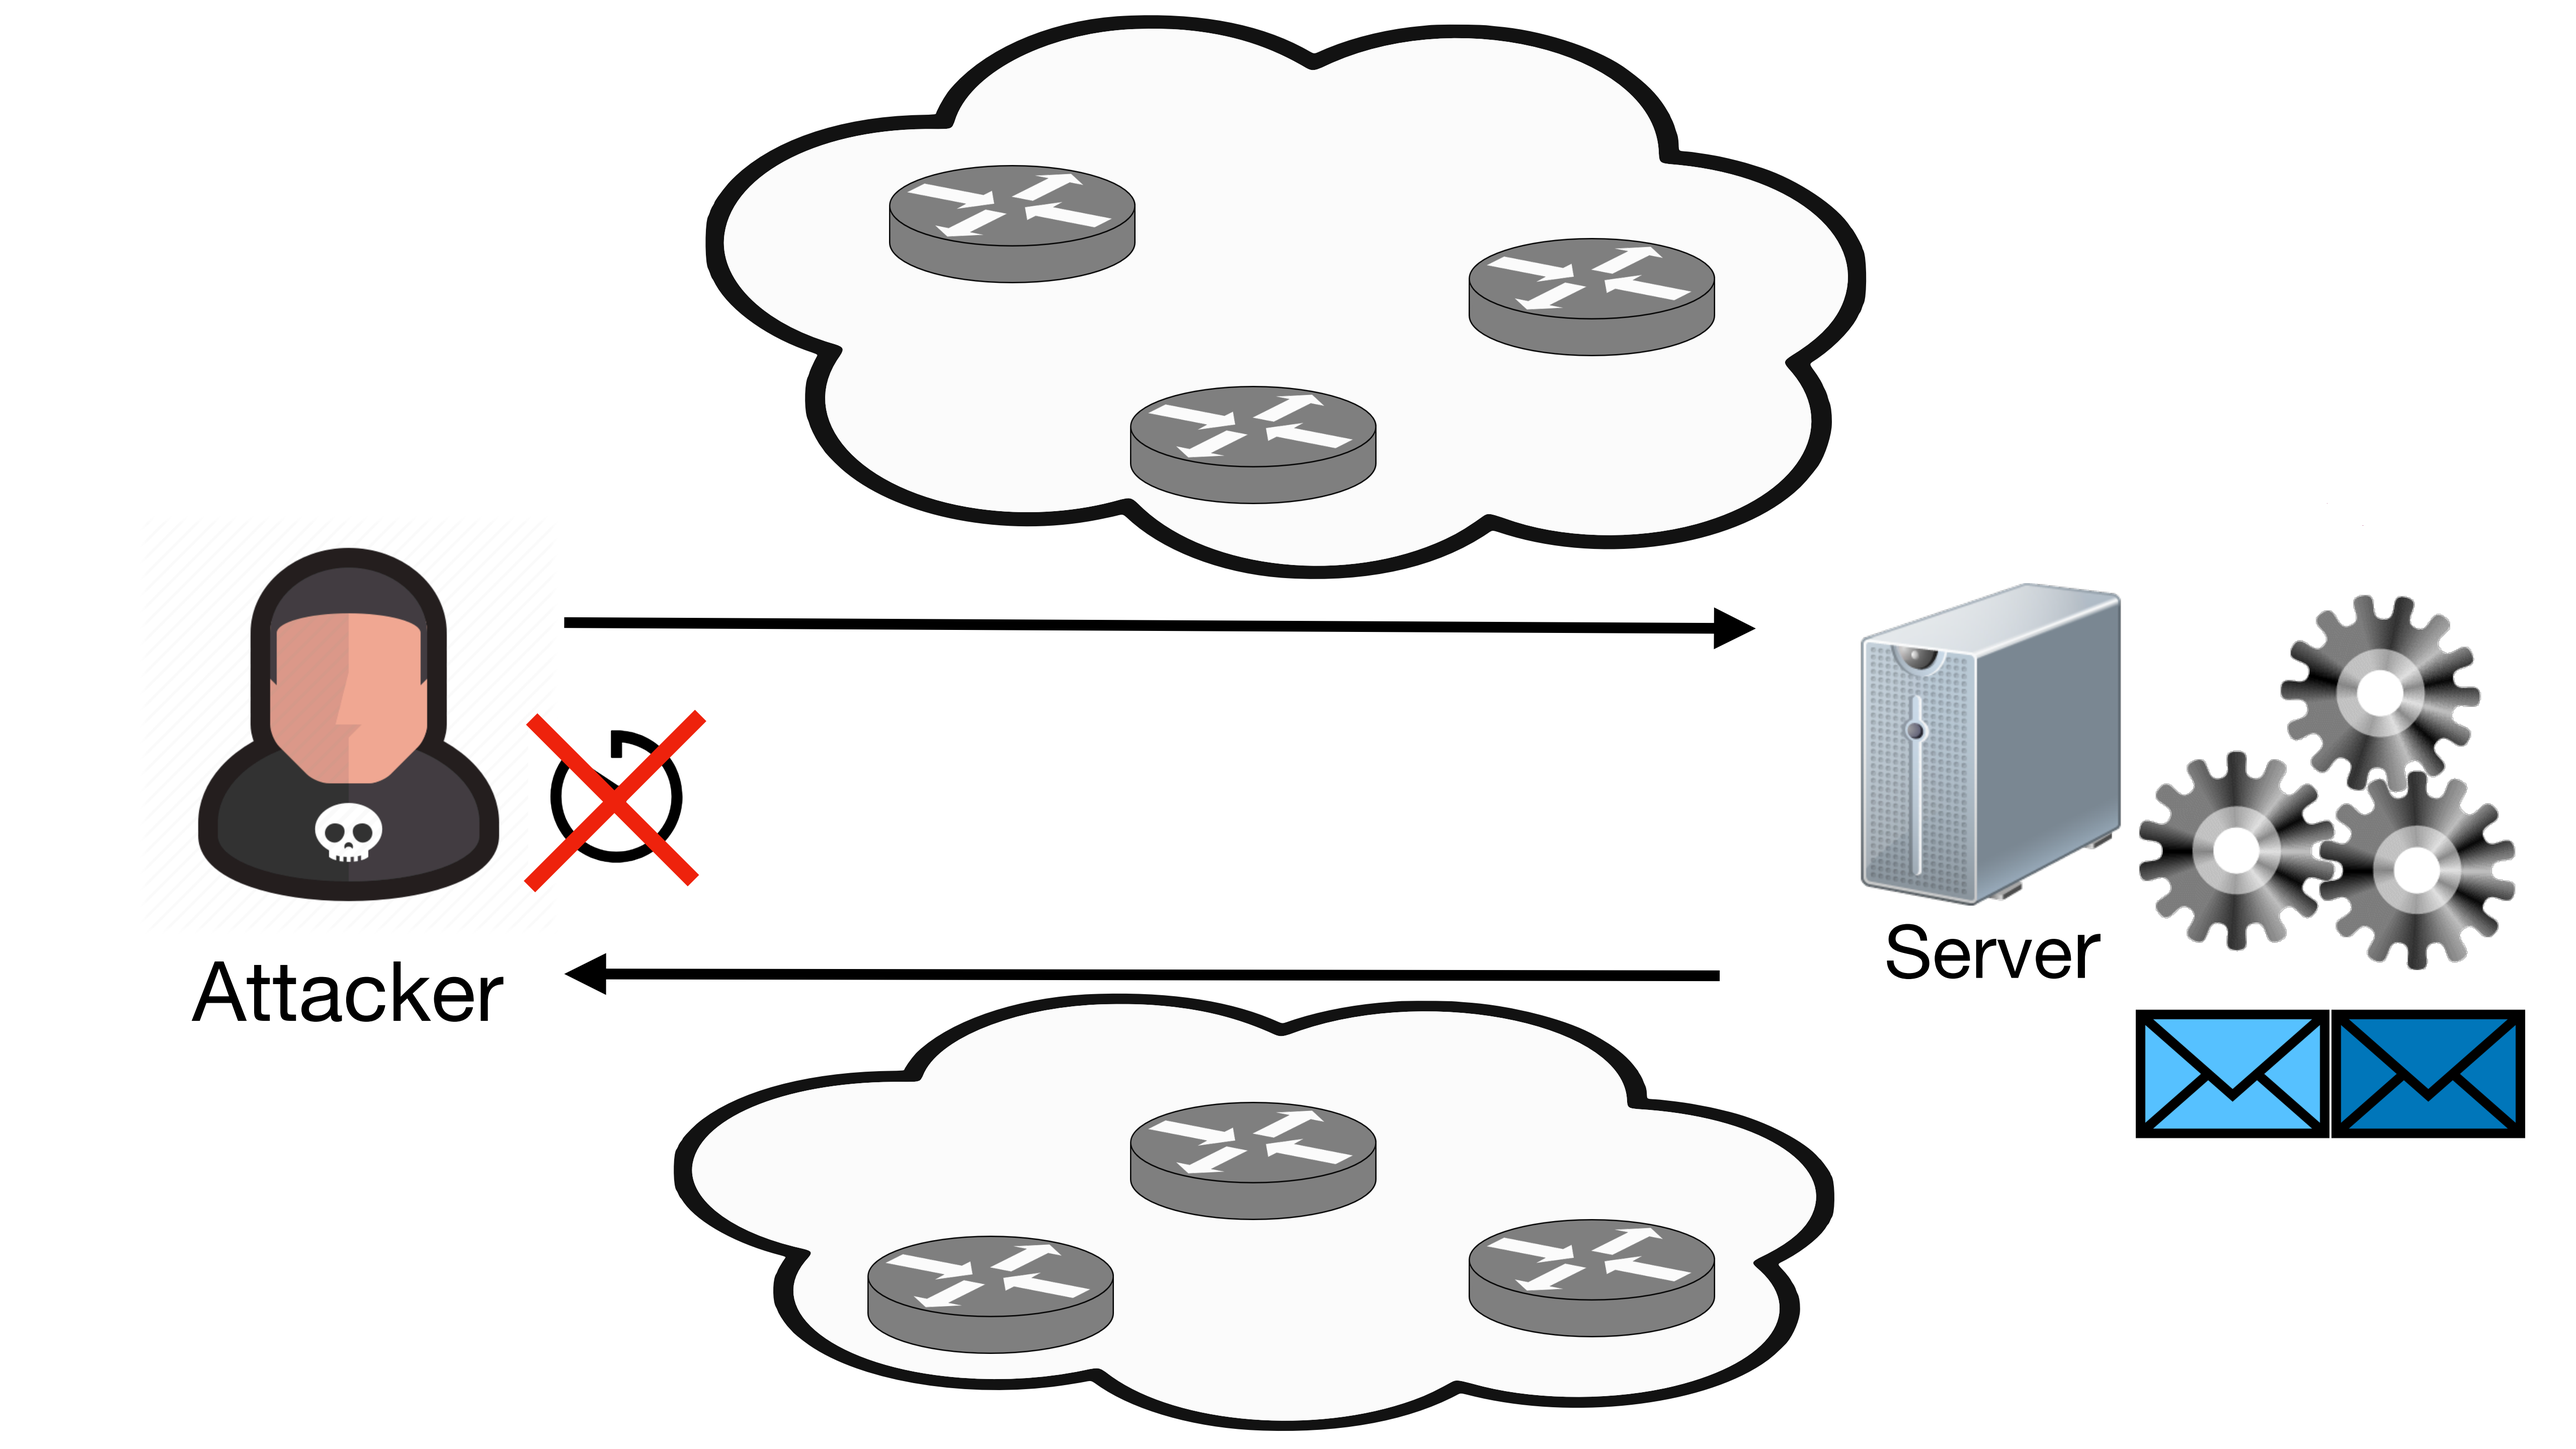
Attacker
Server

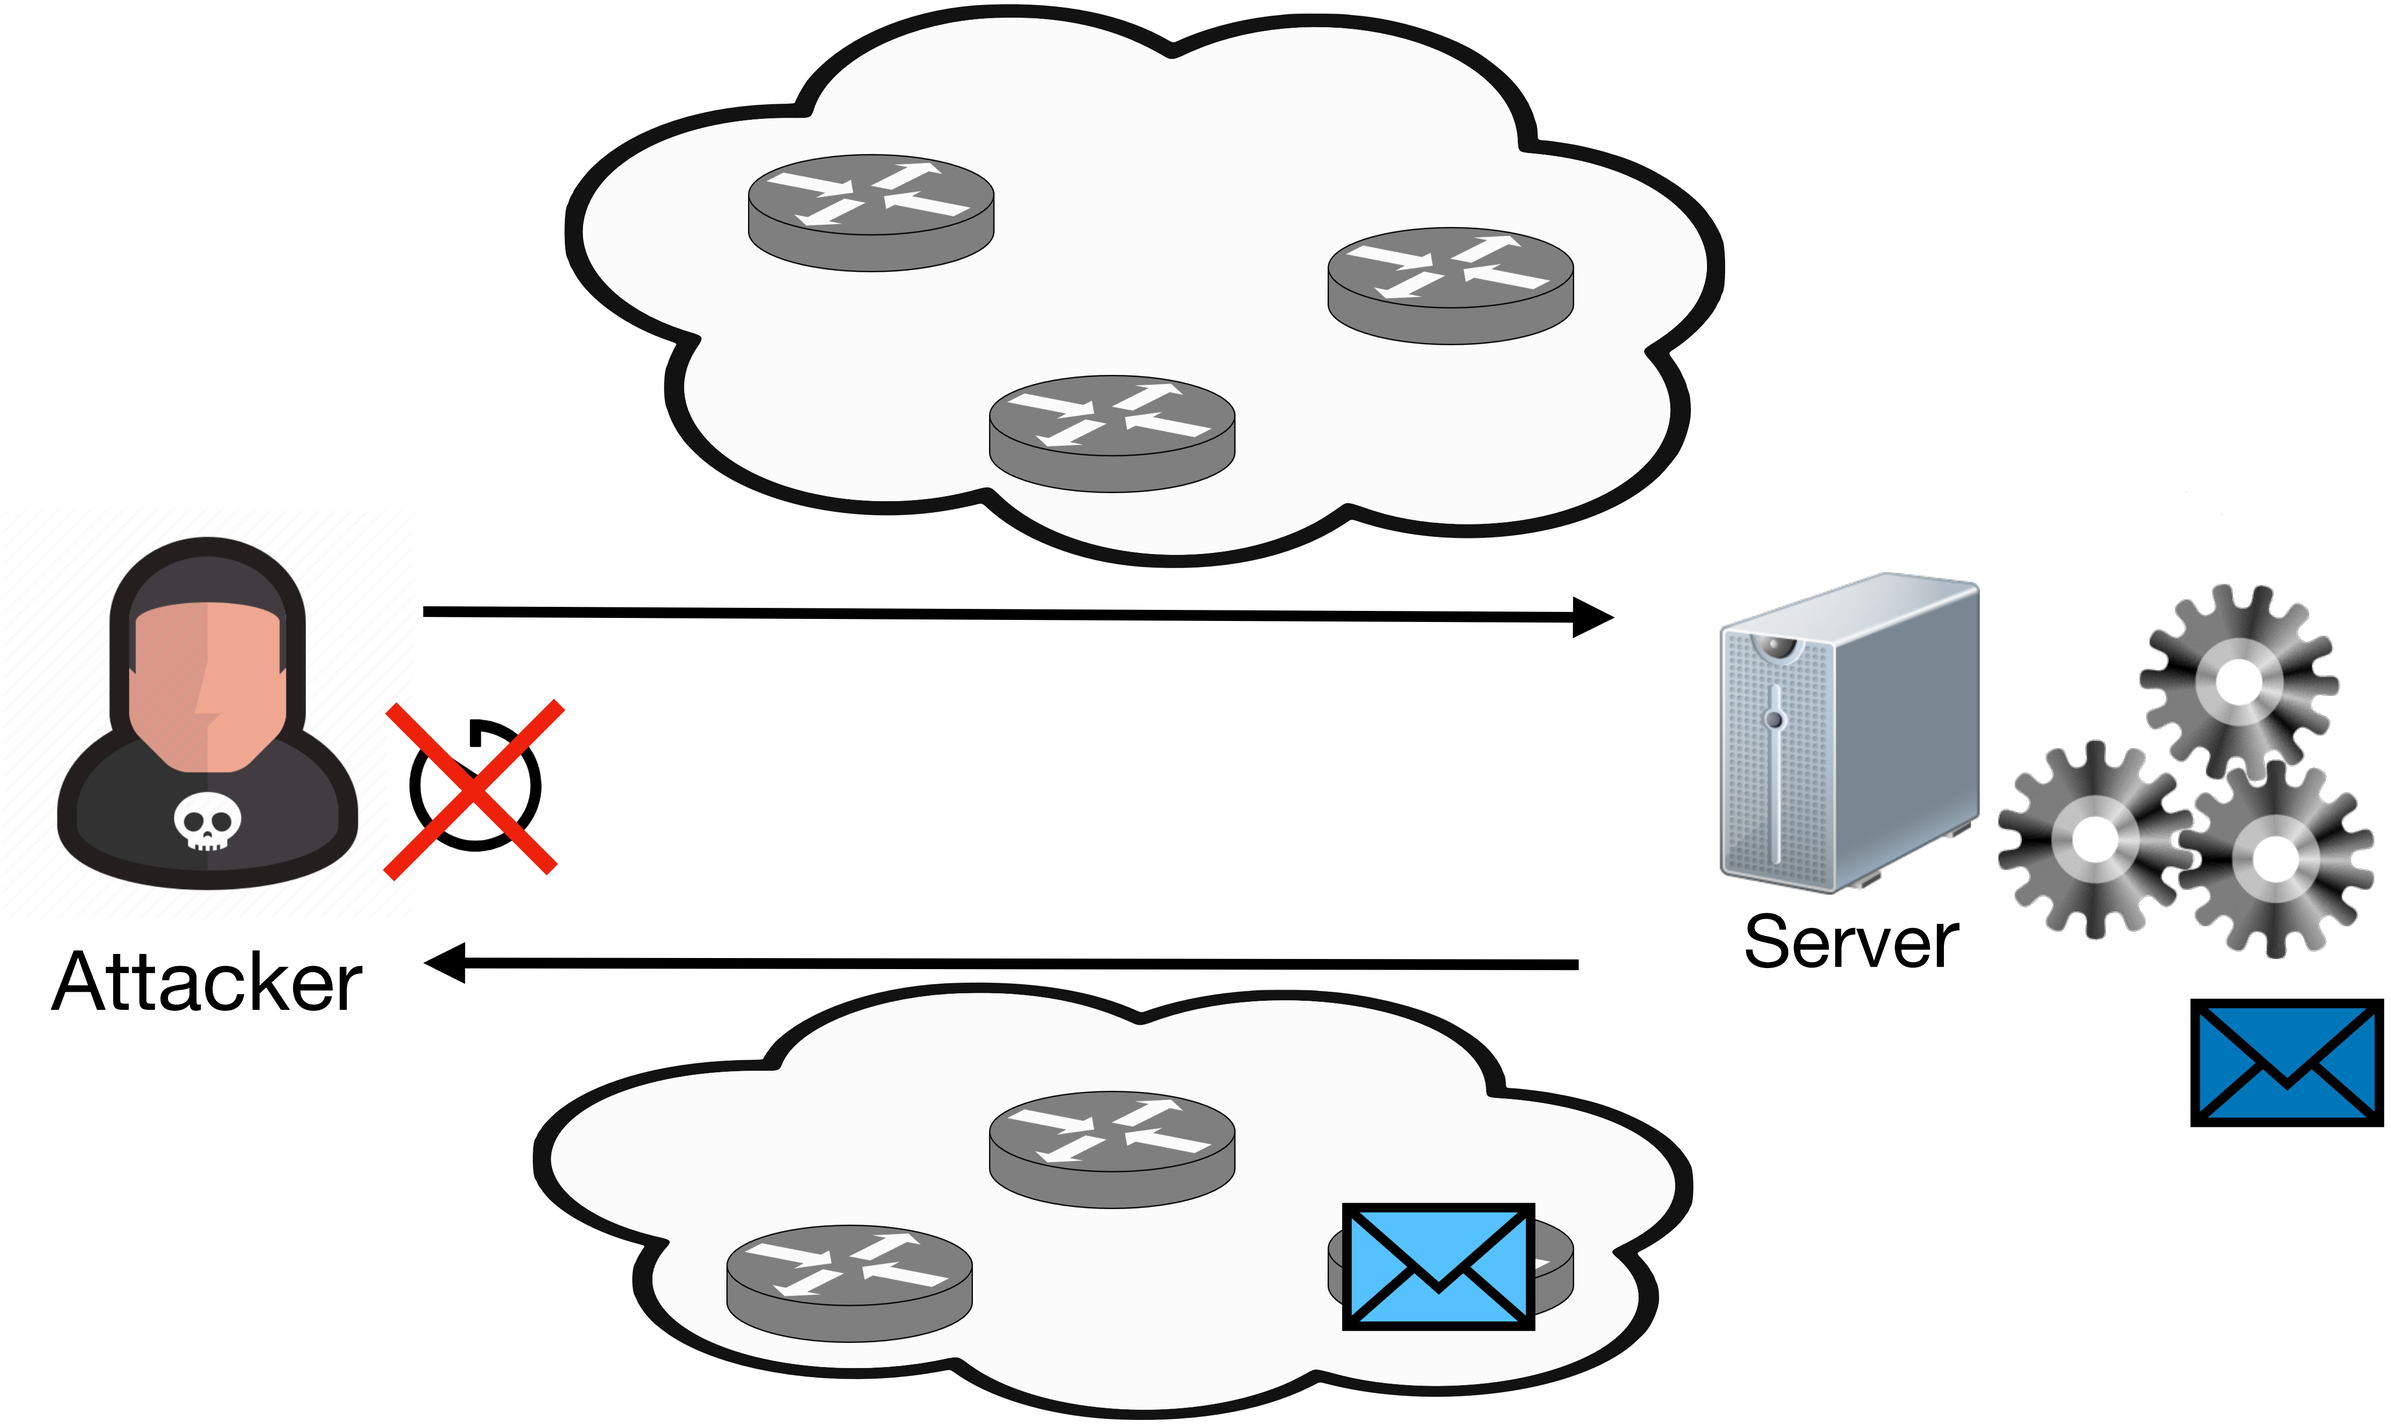
Attacker
Server

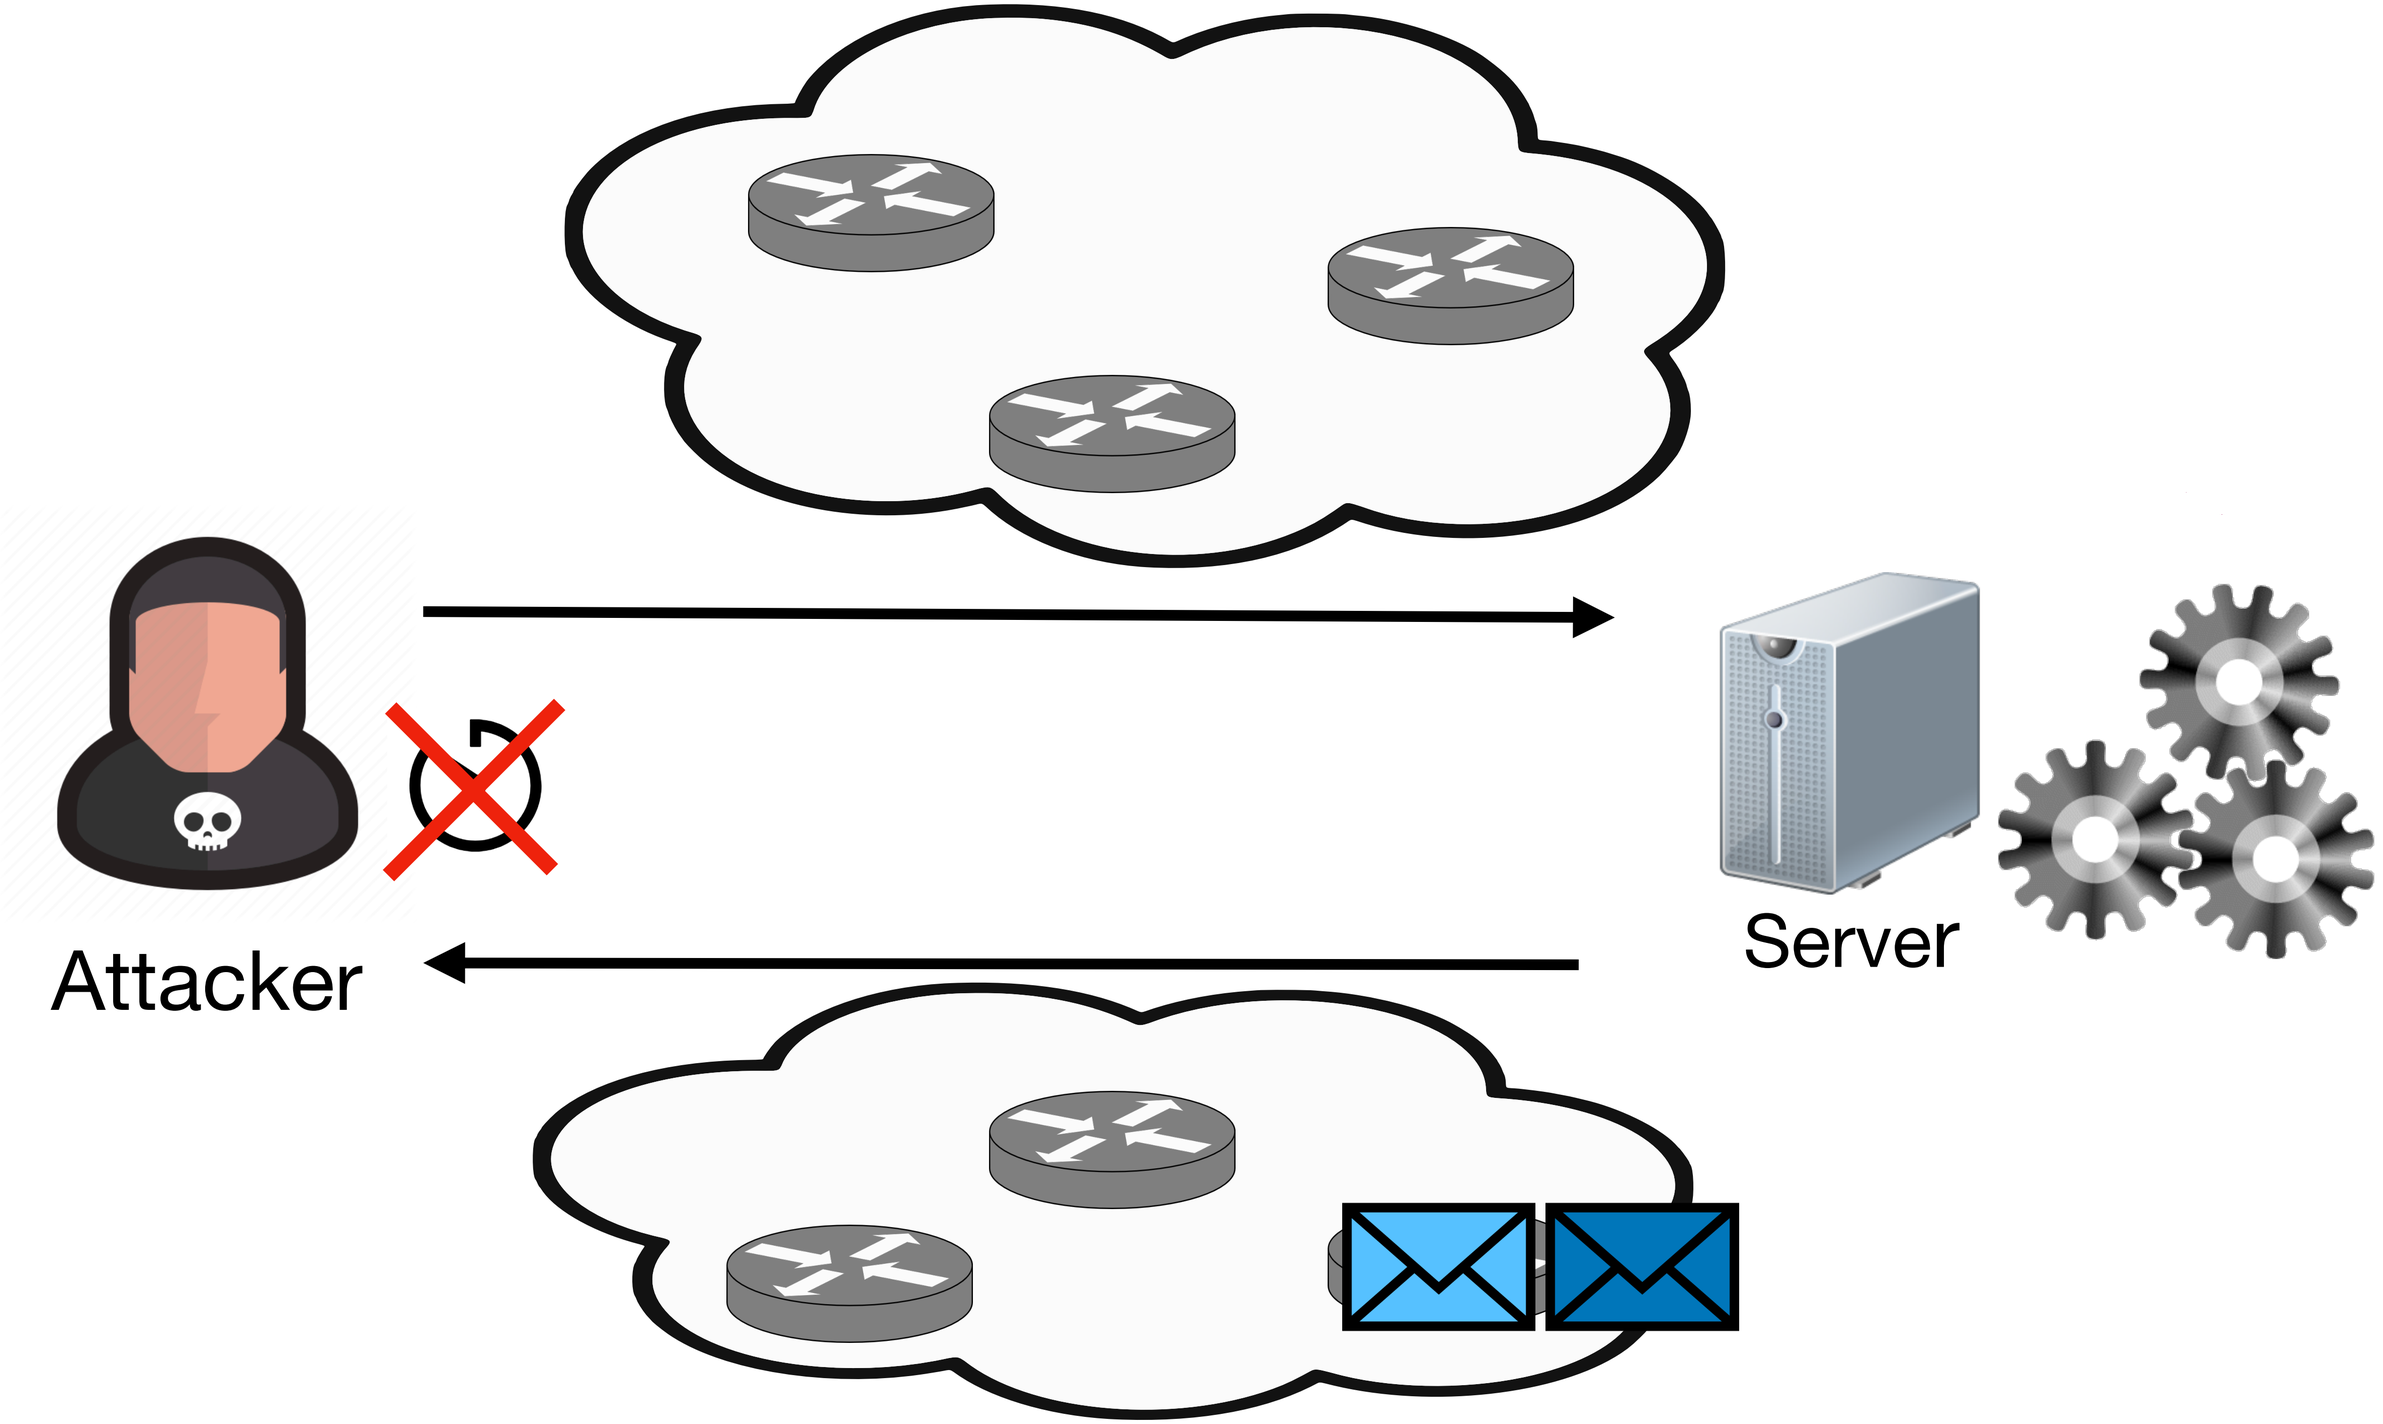
Attacker
Server

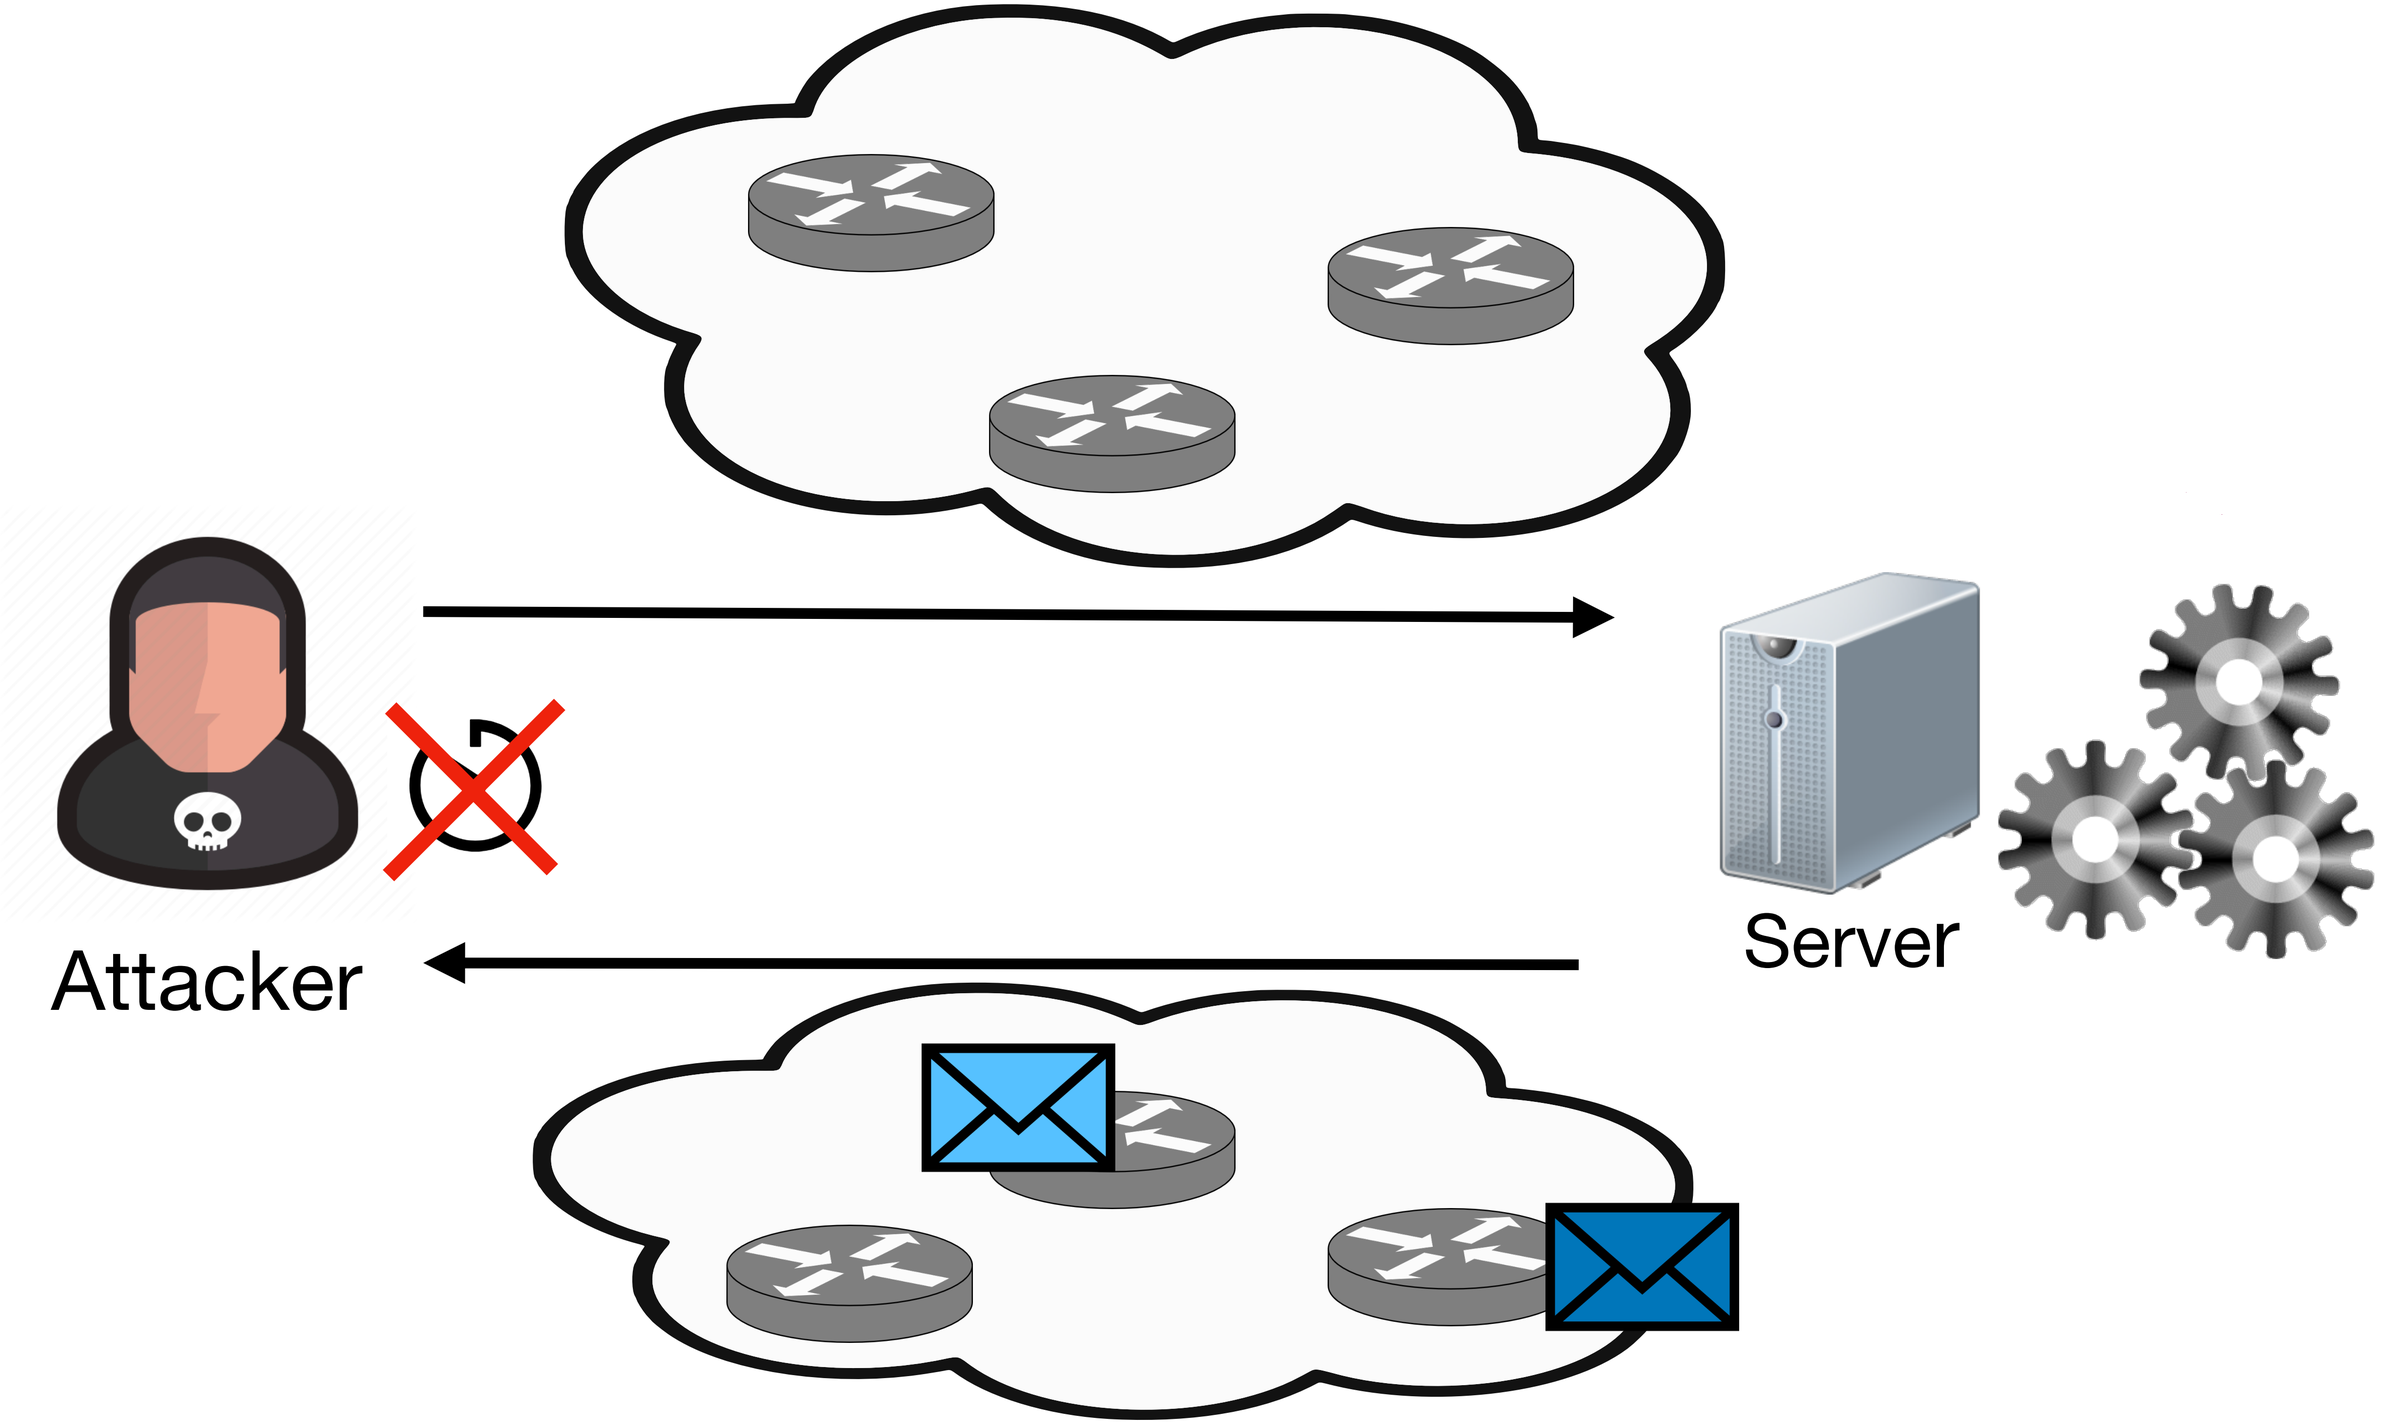
Attacker
Server

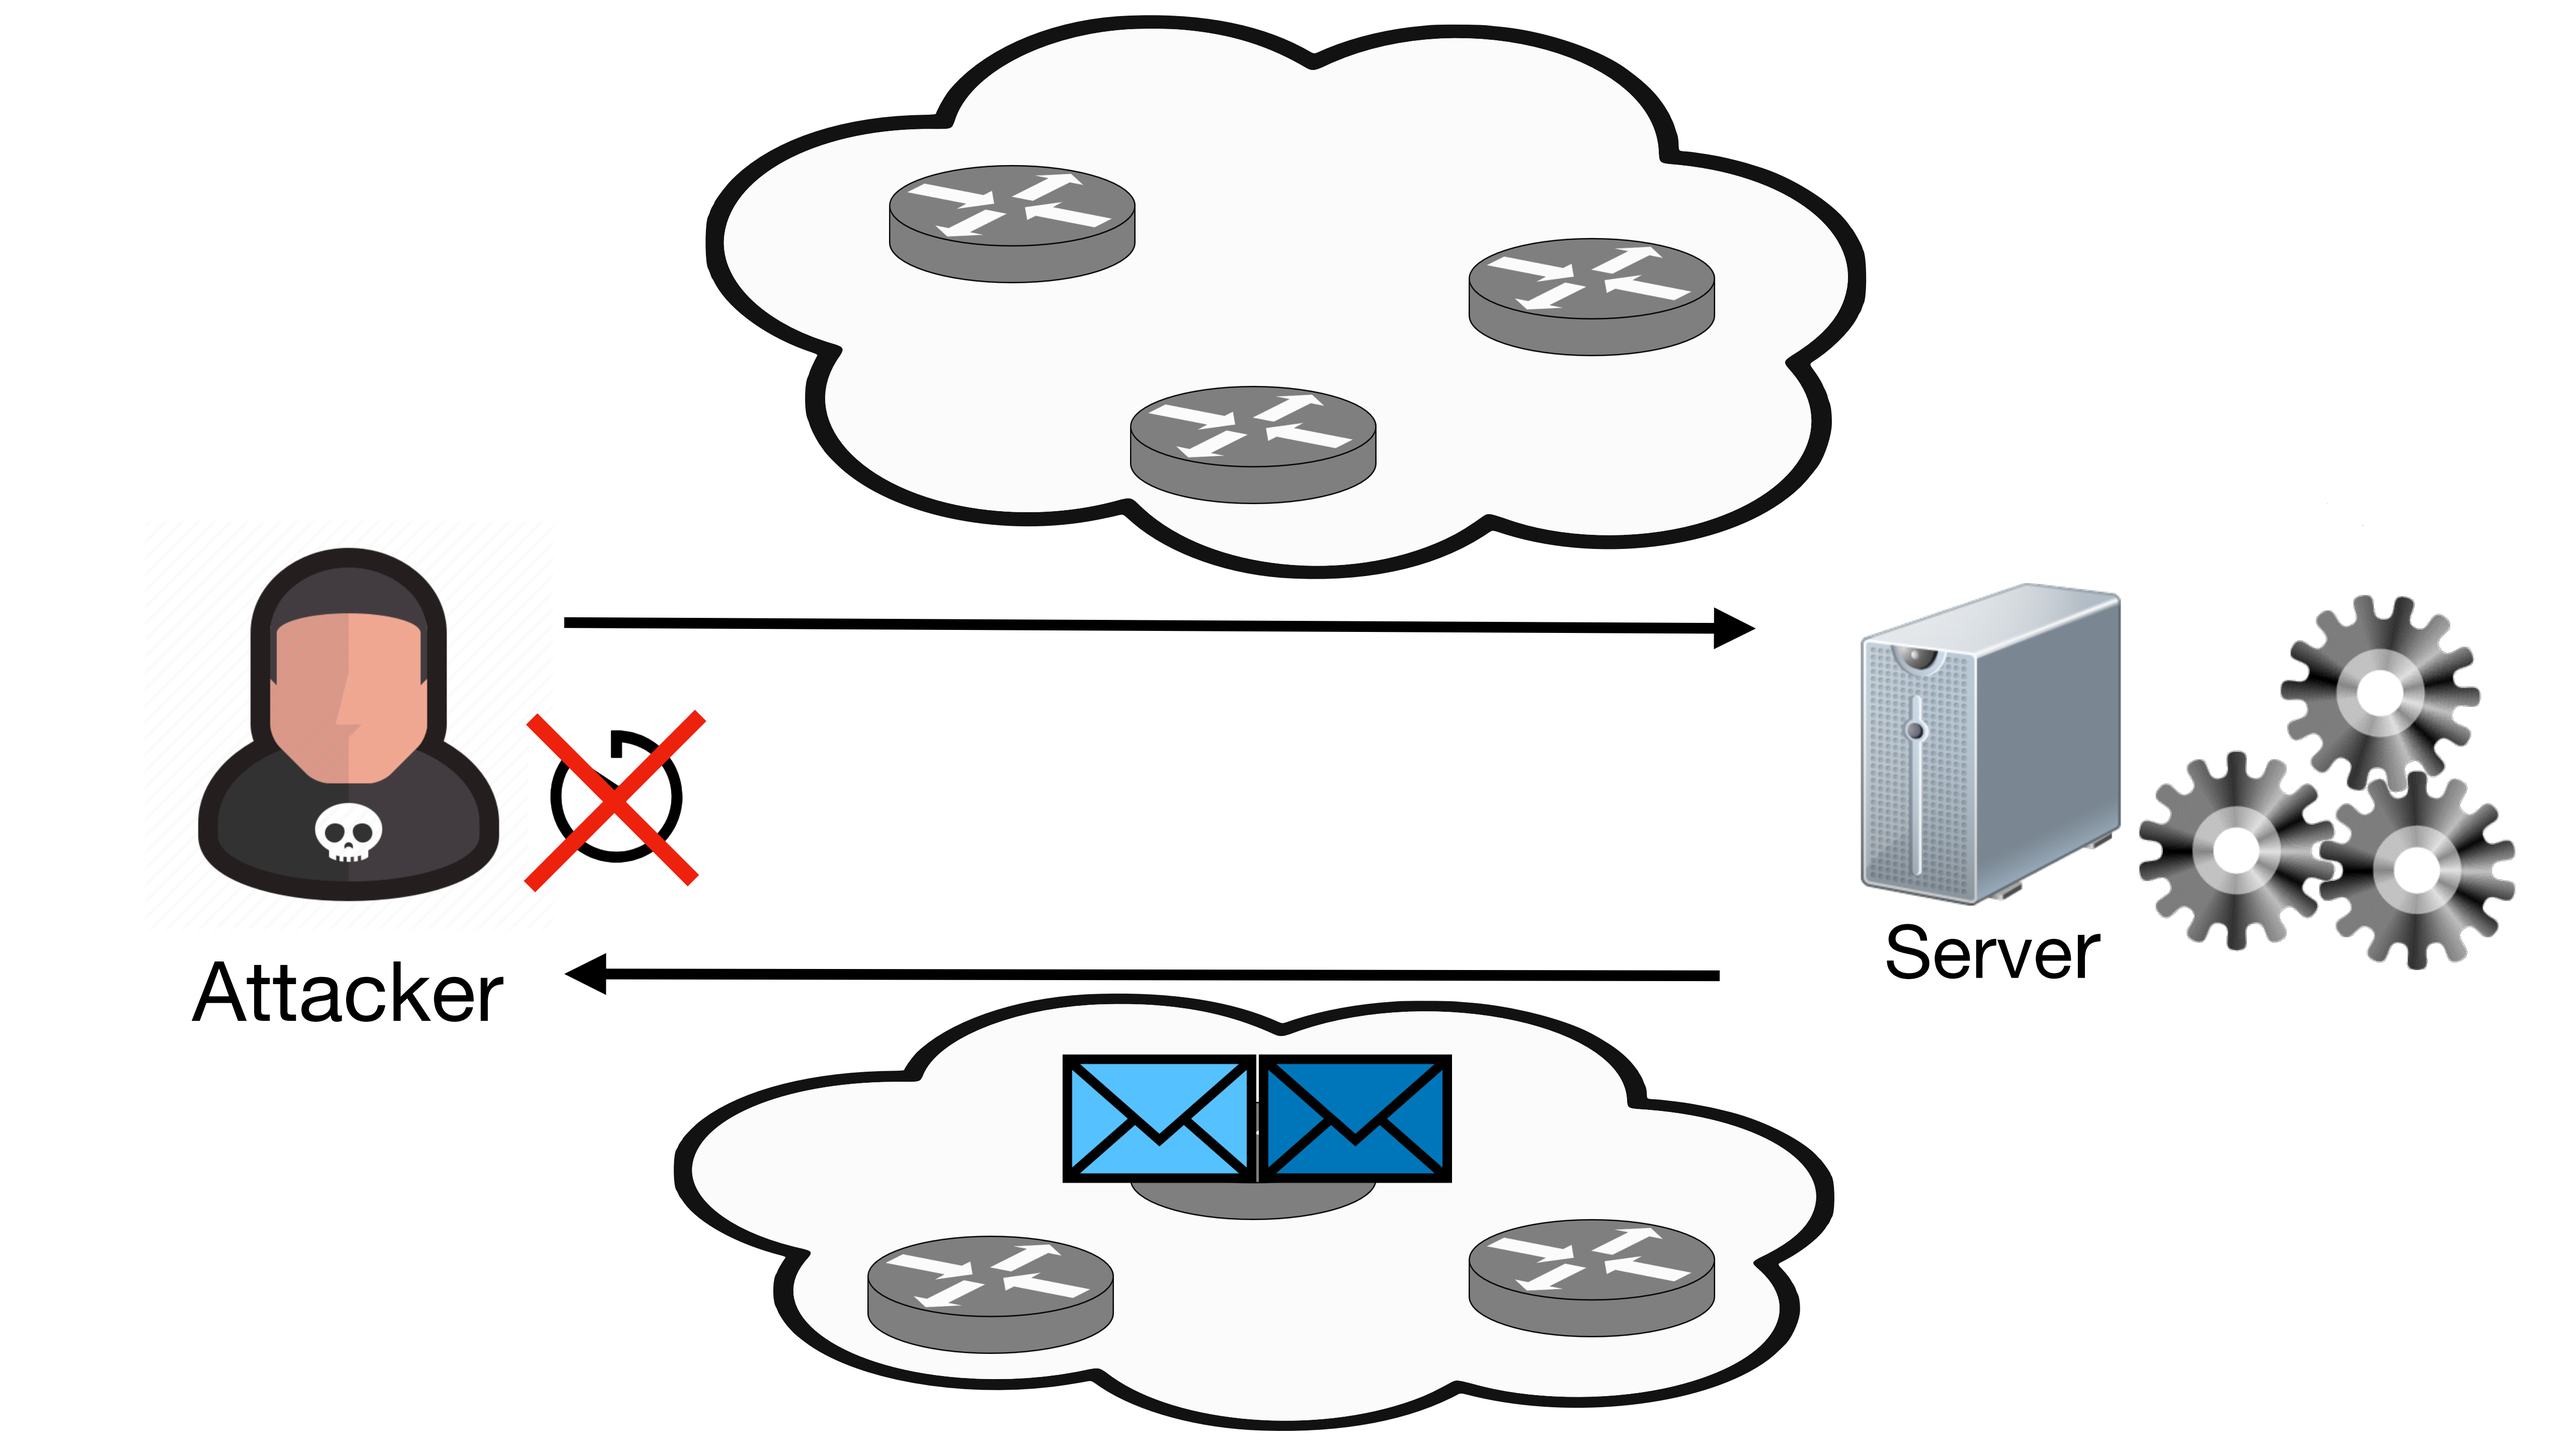
Attacker
Server

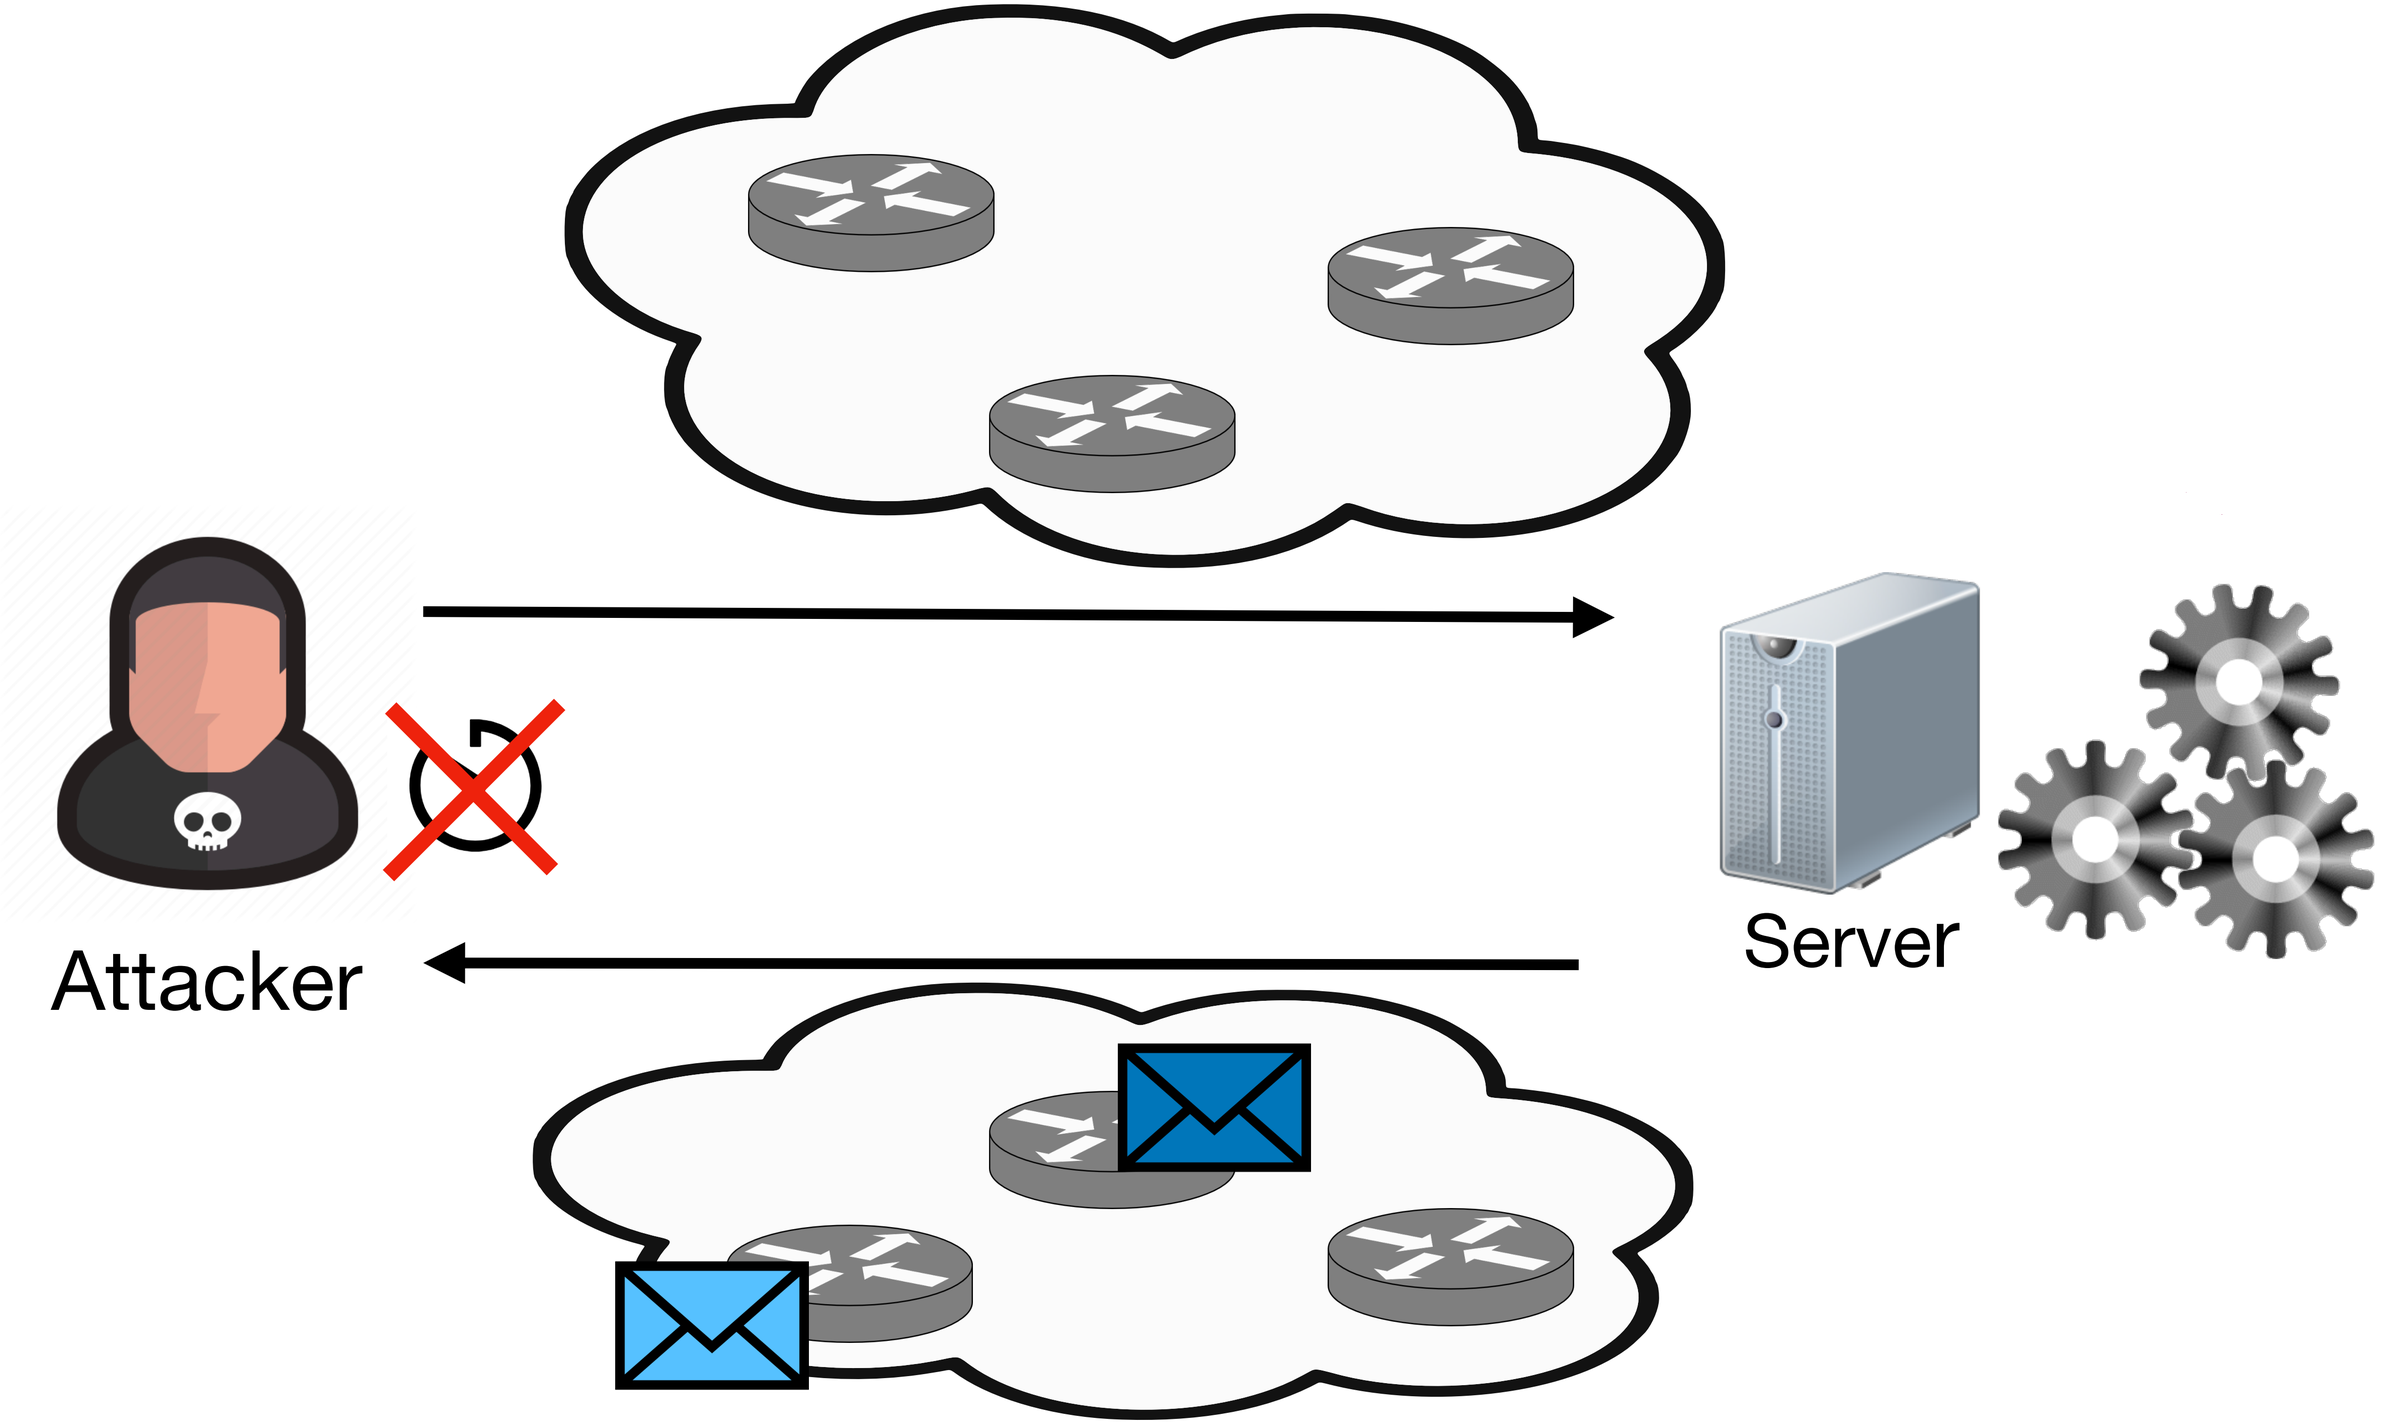
Attacker
Server

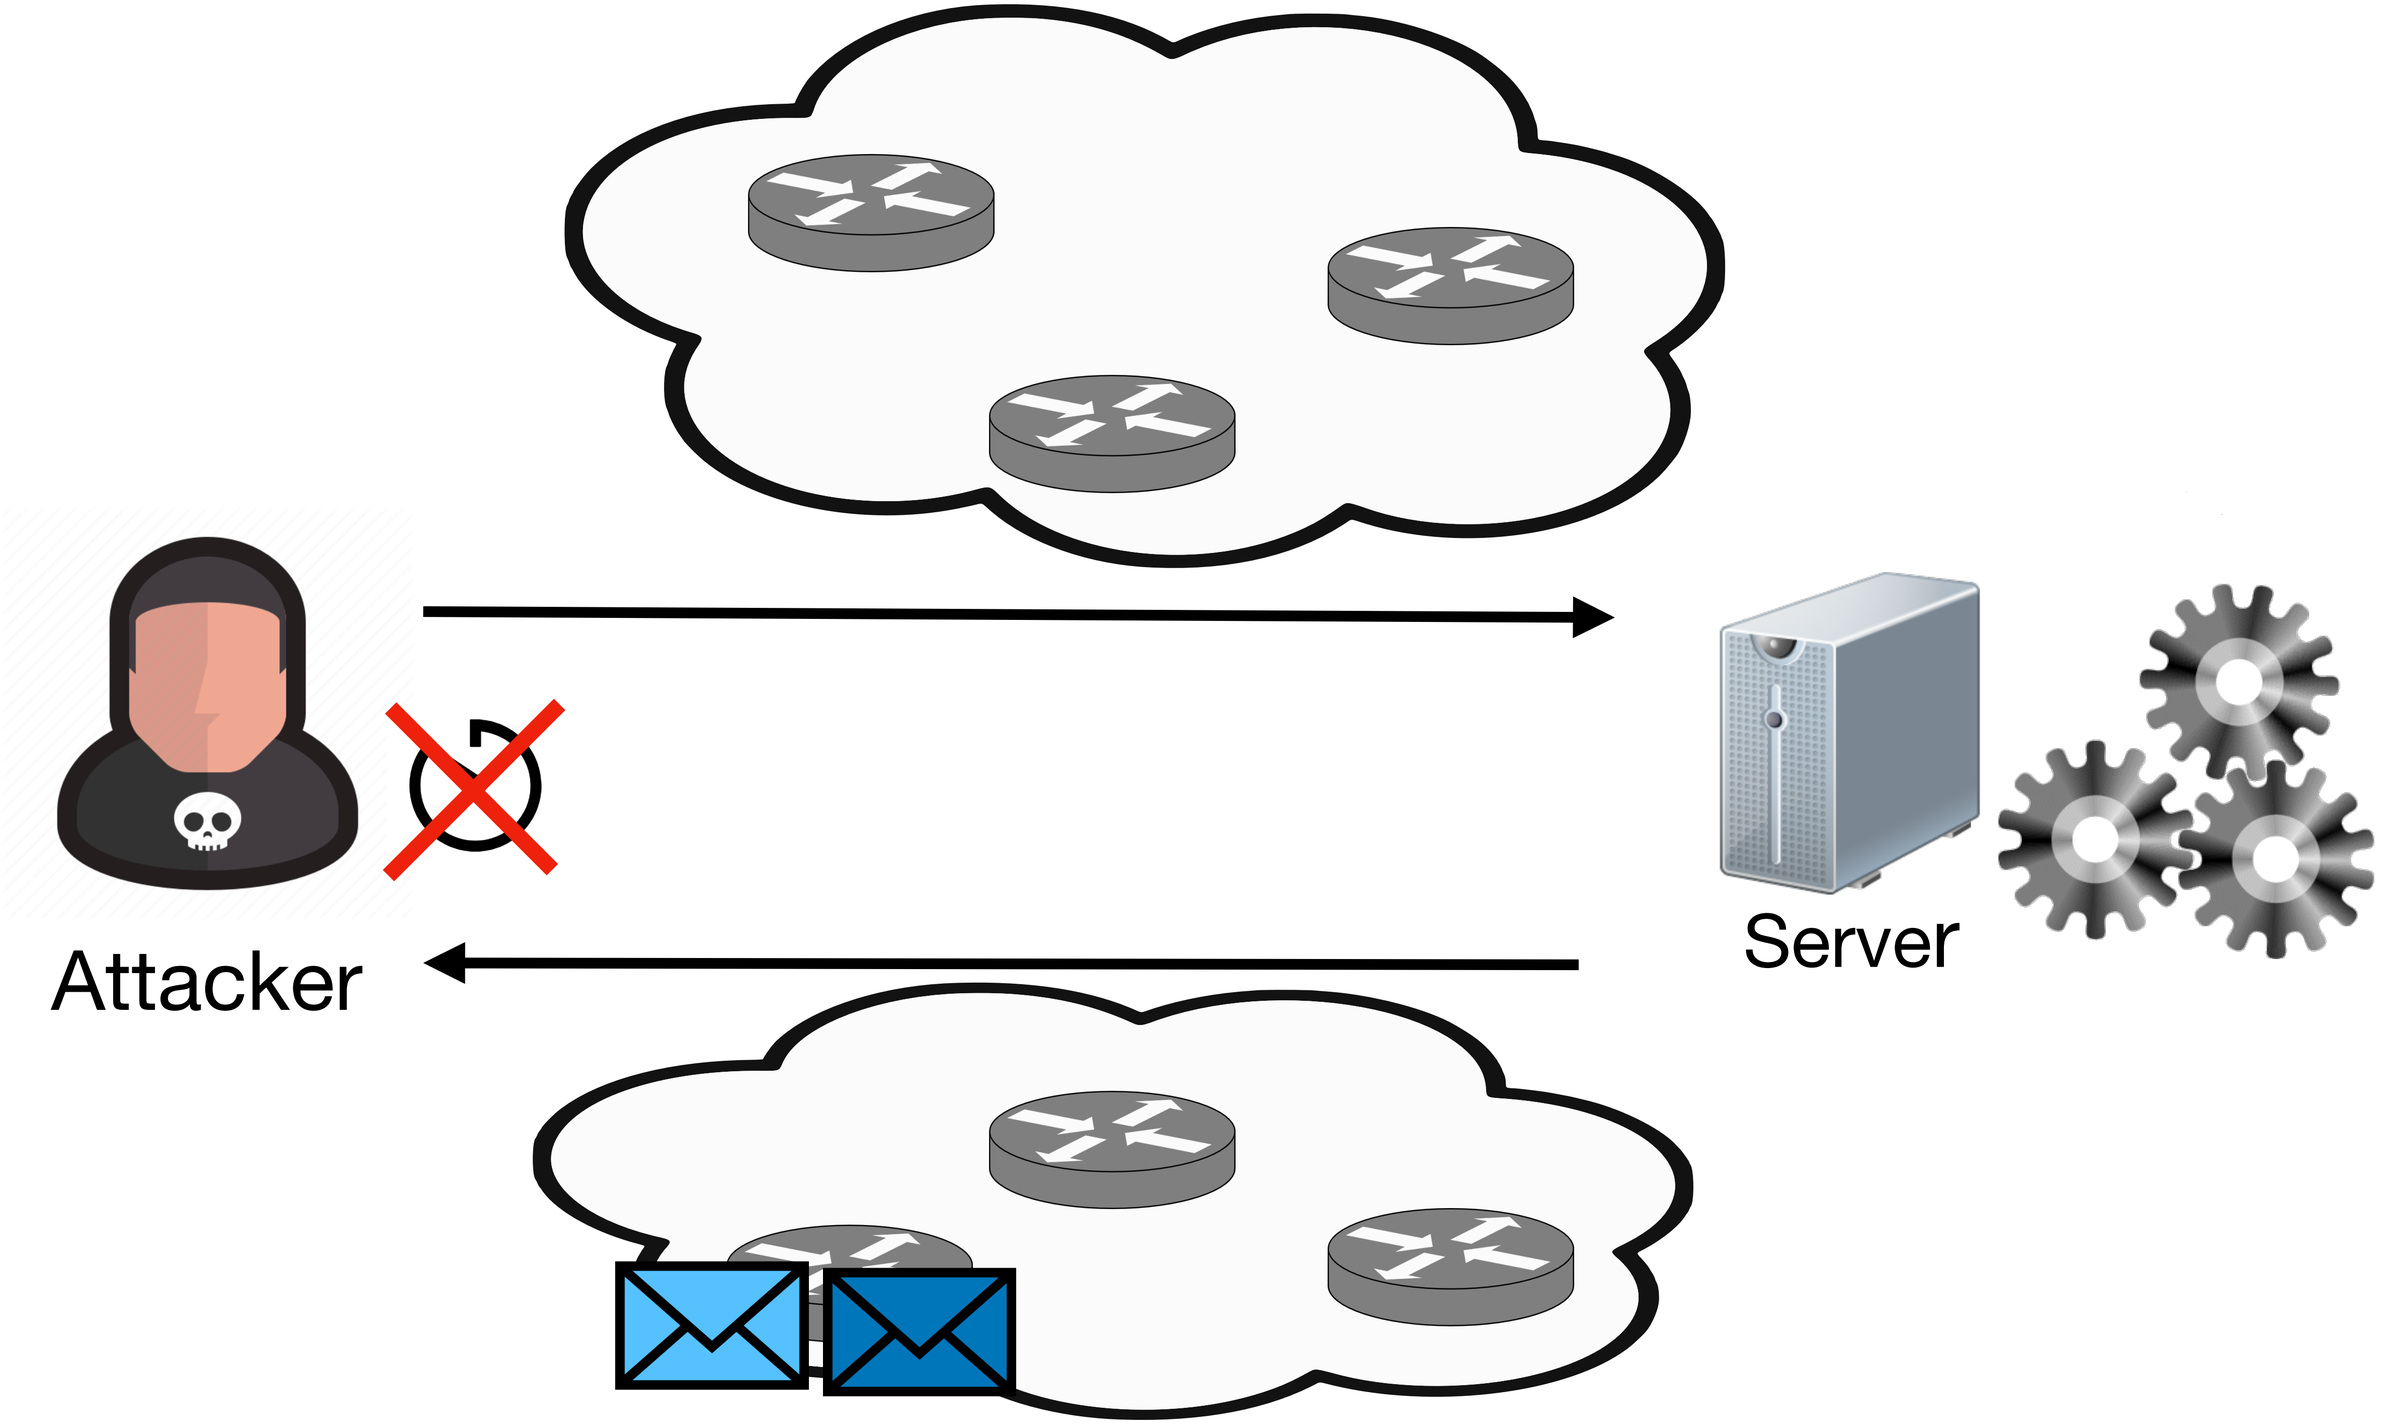
Attacker
Server

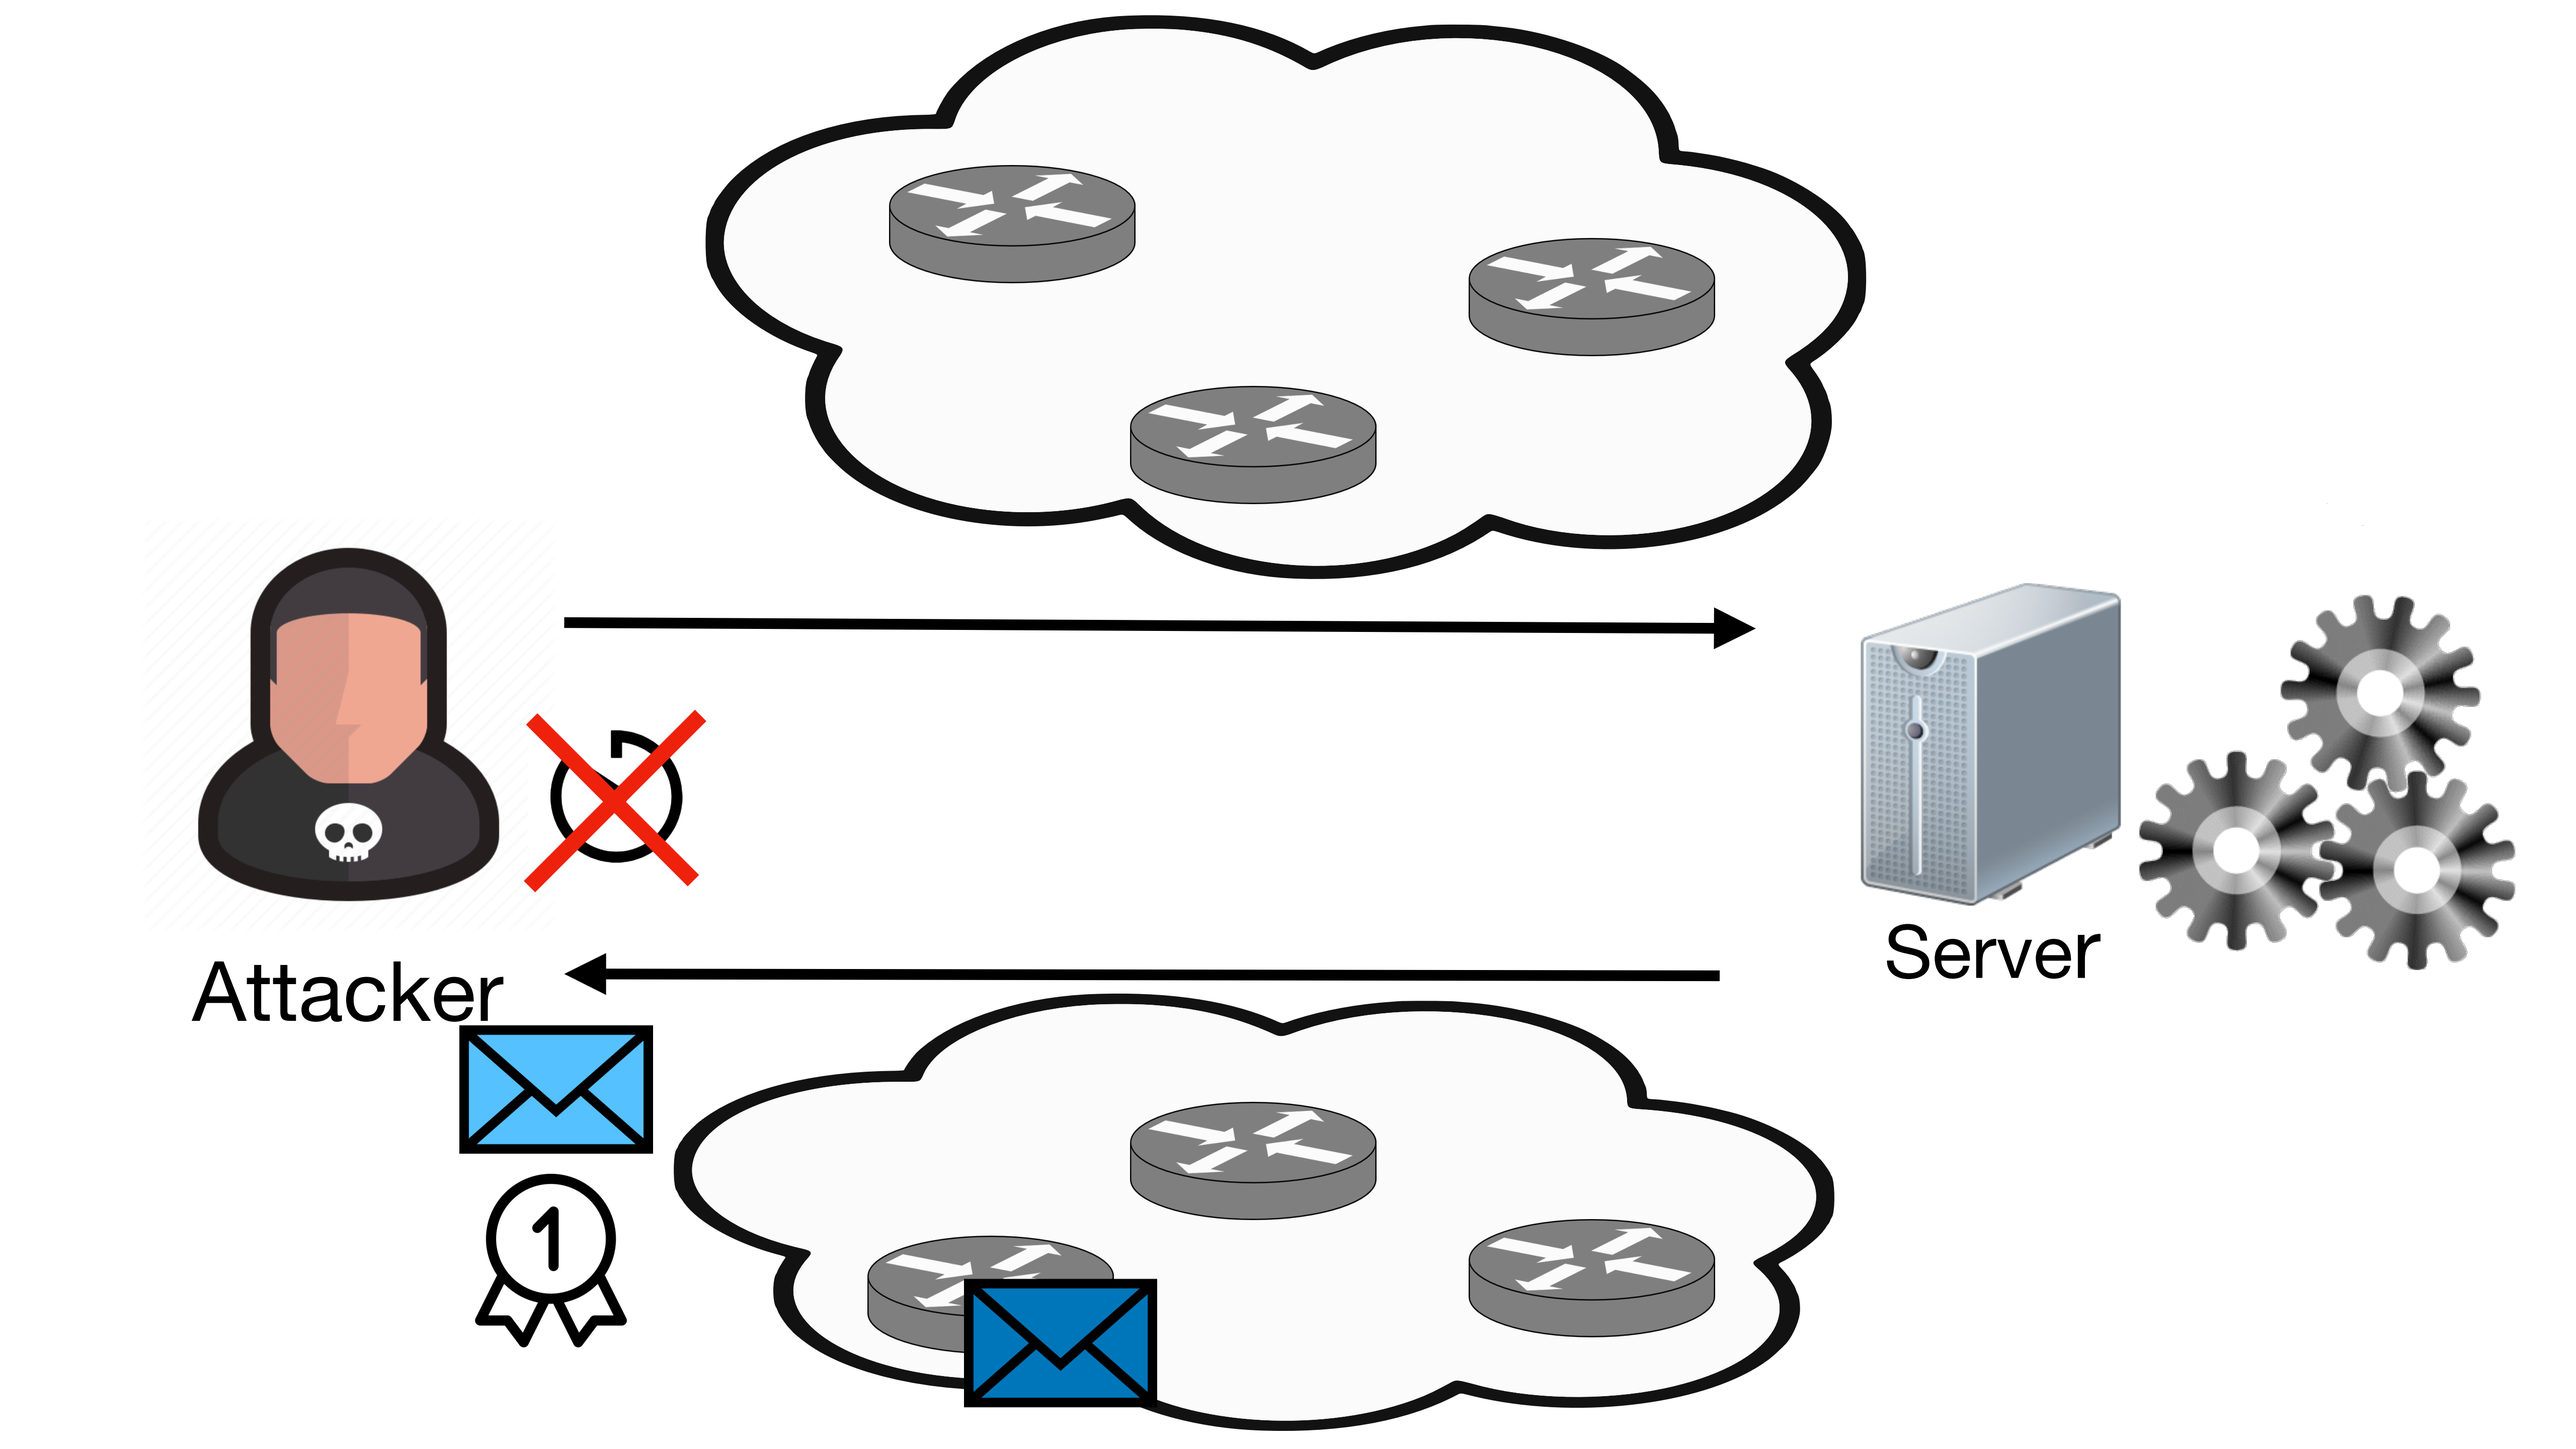
Attacker
Server
1

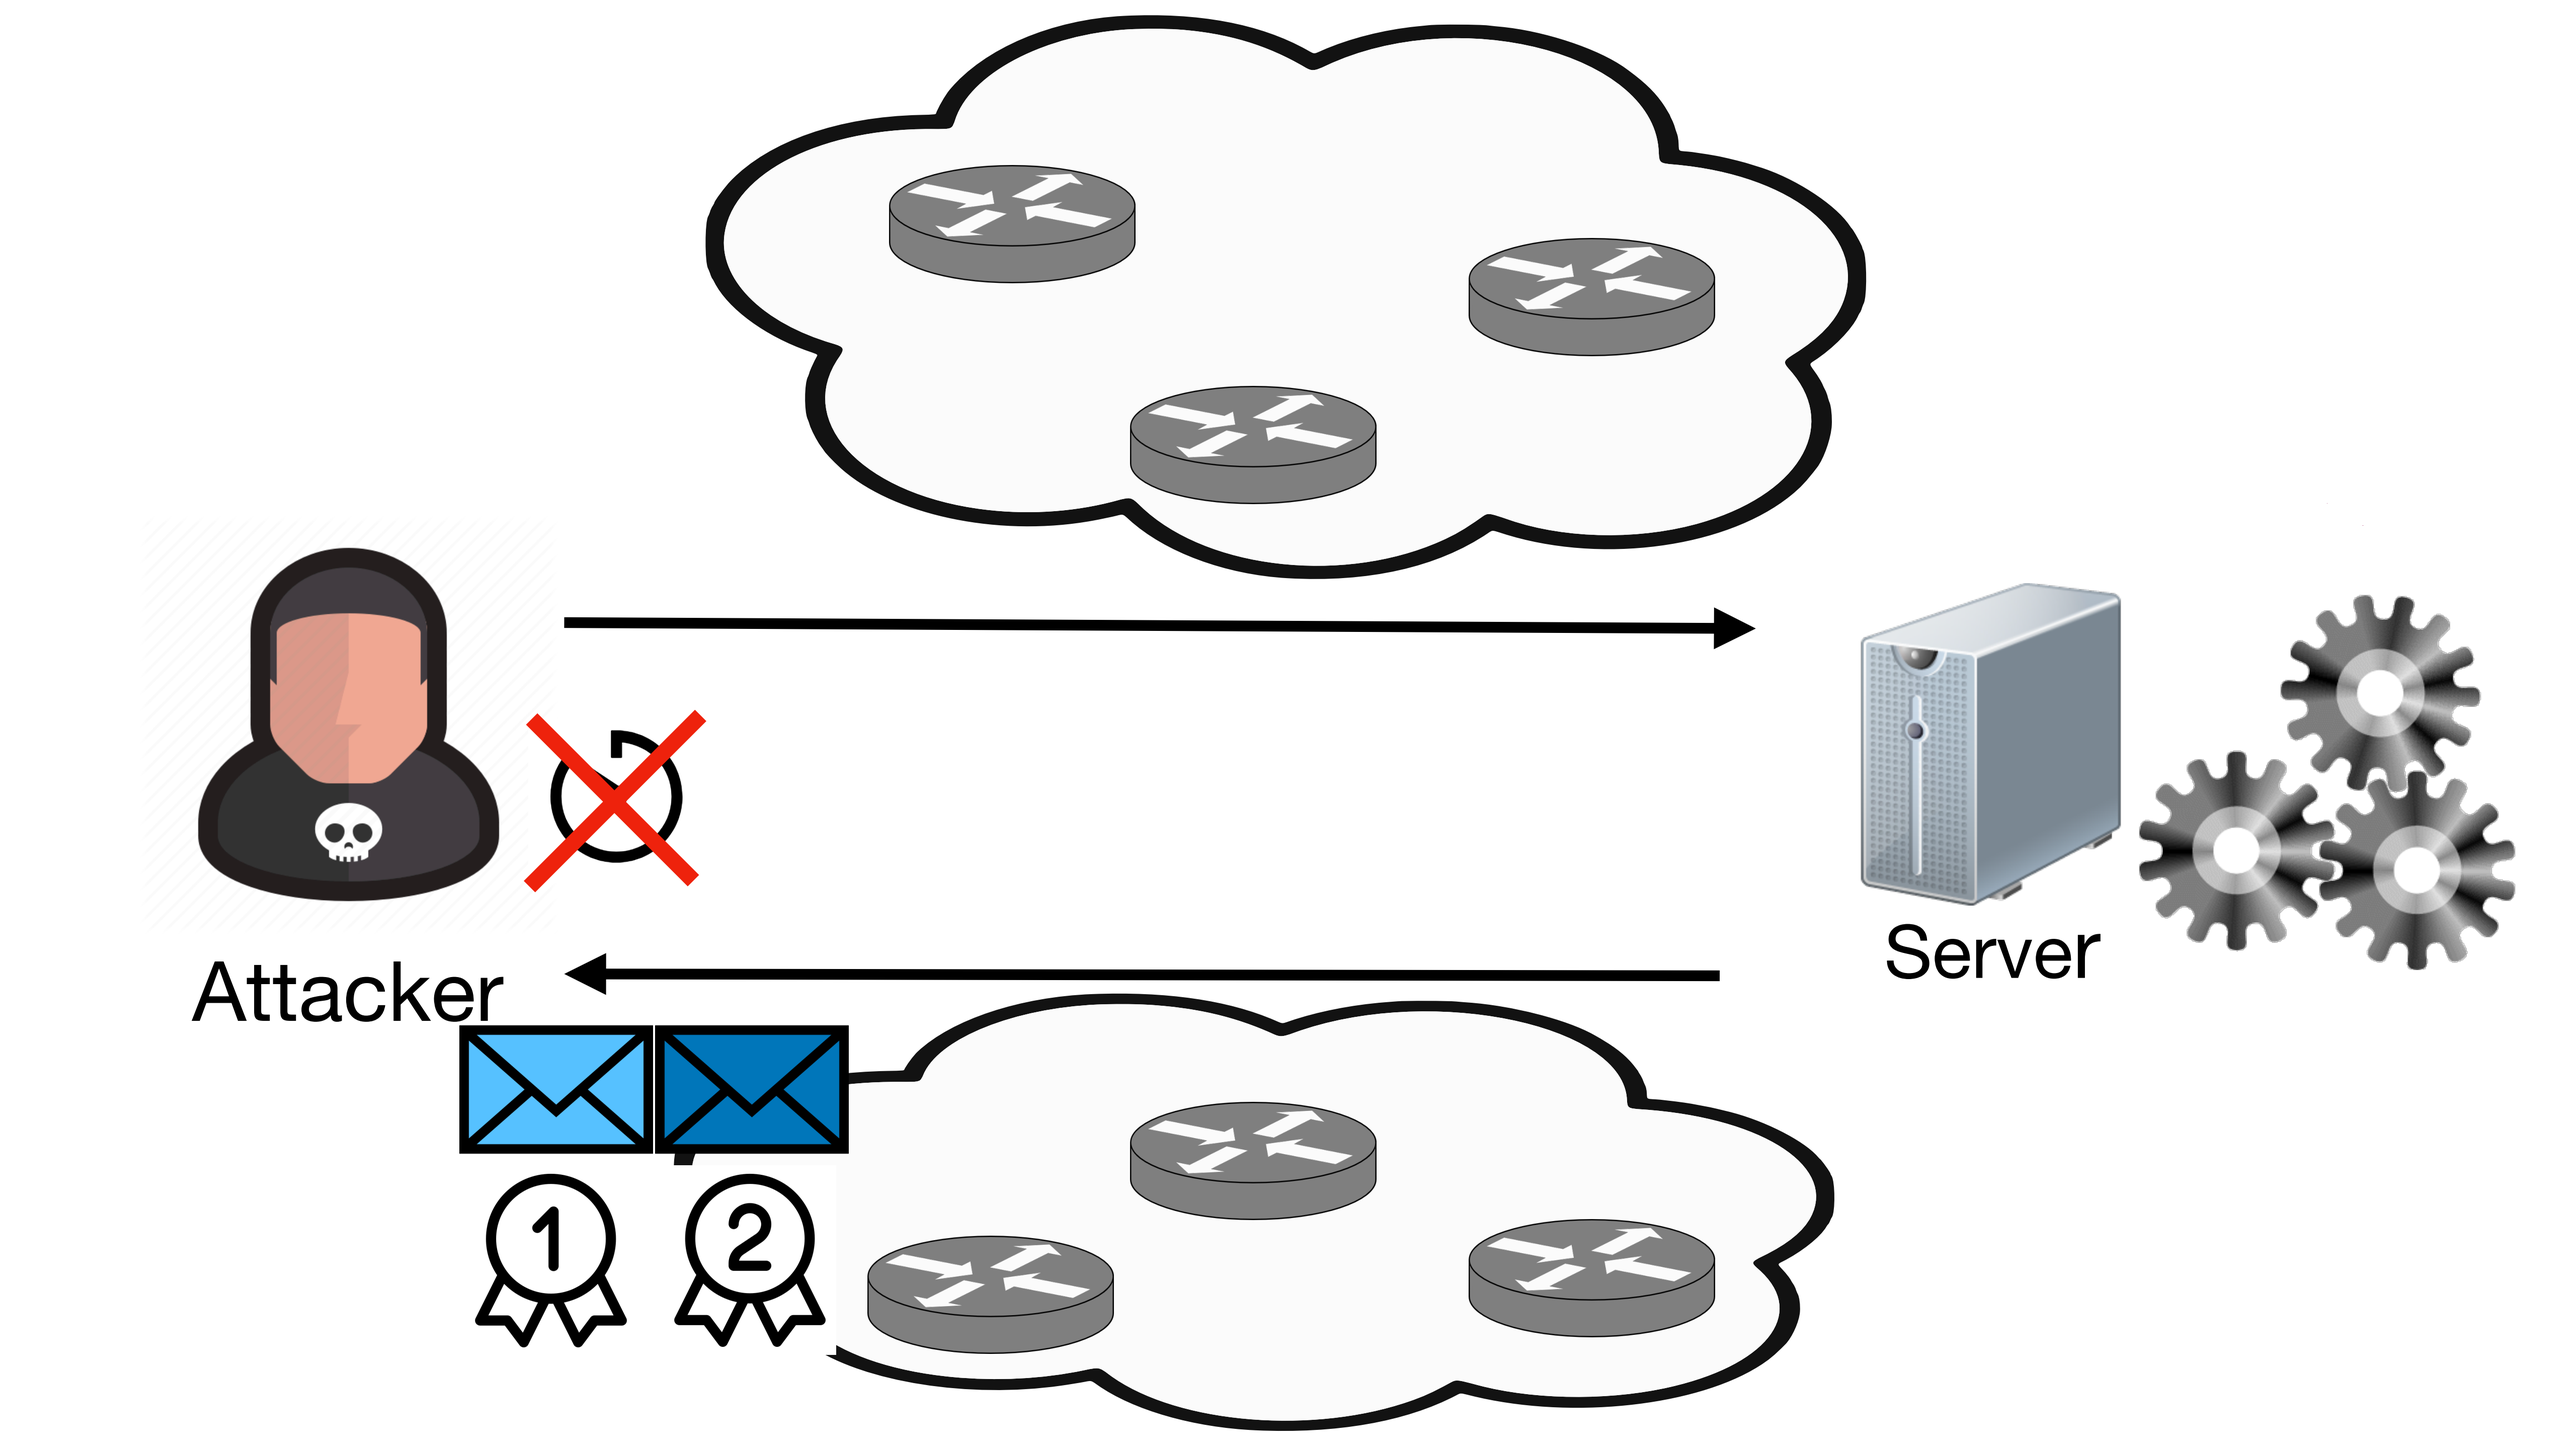
Attacker
Server
1
2

Timeless Timing Attacks: Requirements

1. Requests need to **arrive at the same time** at the server
2. Server needs to process requests **concurrently**
3. **Response order** needs to reflect difference in execution time

HTTP/2 (multiplexing)

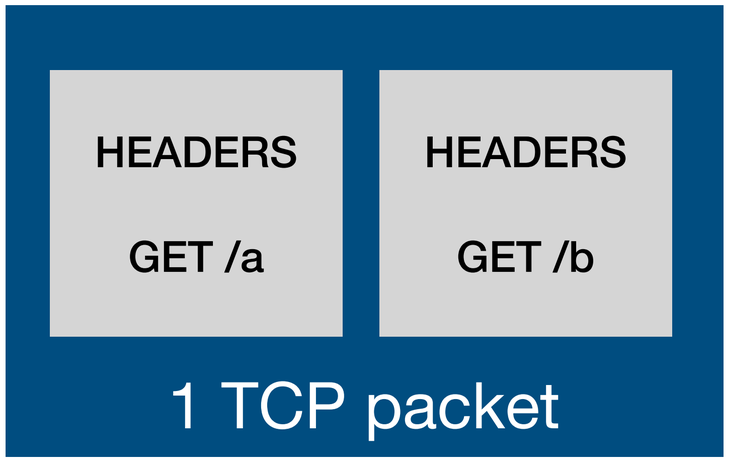

HTTP/1 + Tor (encapsulation)

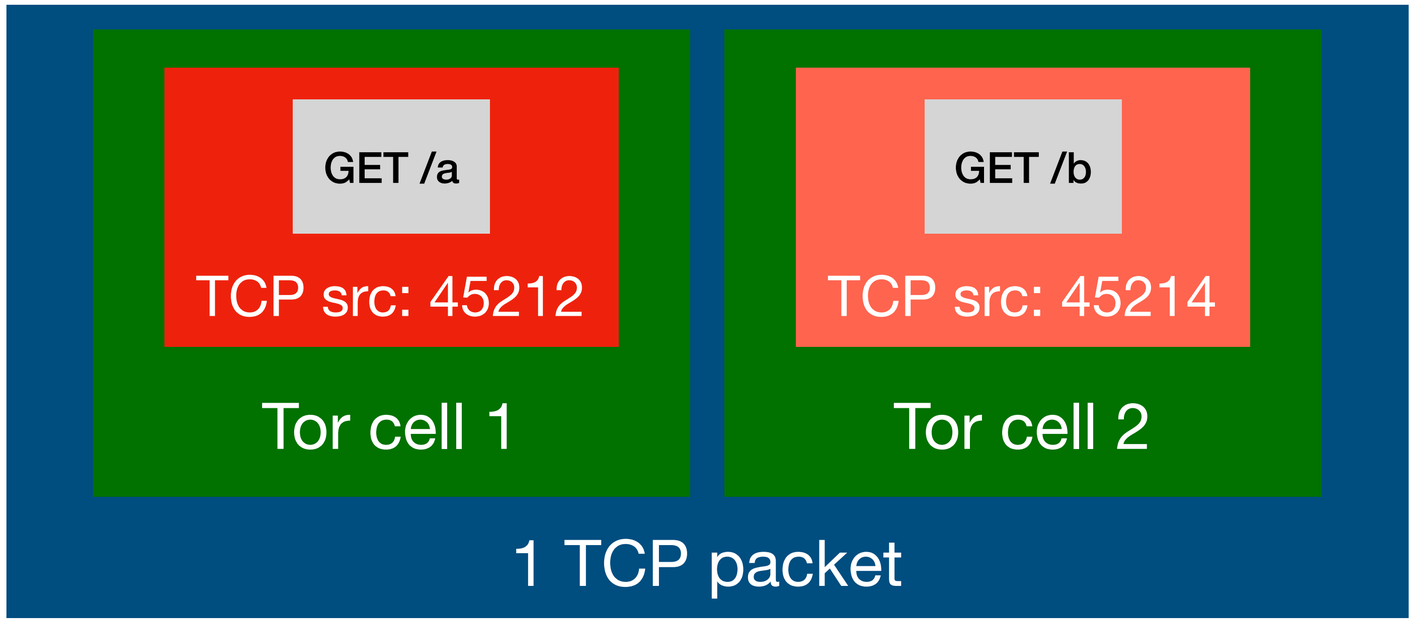

How many requests/pairs are needed?

Sequential Timing Attacks

	EU	US	Asia	LAN	localhost
50µs	333	4,492	7,386	20	14
20µs	2,926	16,820	-	41	16
10µs	23,220	-	-	126	20
5µs	-	-	-	498	42
Smallest diff	10µs	20µs	50µs	150ns	150ns

Concurrency Timing Attacks

	Internet (anywhere)
50µs	6
20µs	6
10µs	11
5µs	52
Smallest diff	100ns

Attack Scenarios

1. direct timing attack

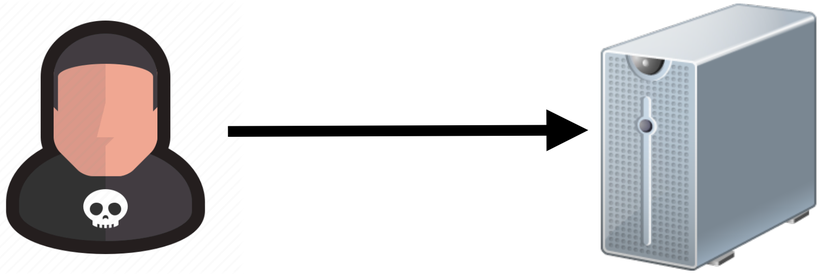

3. Wi-Fi authentication

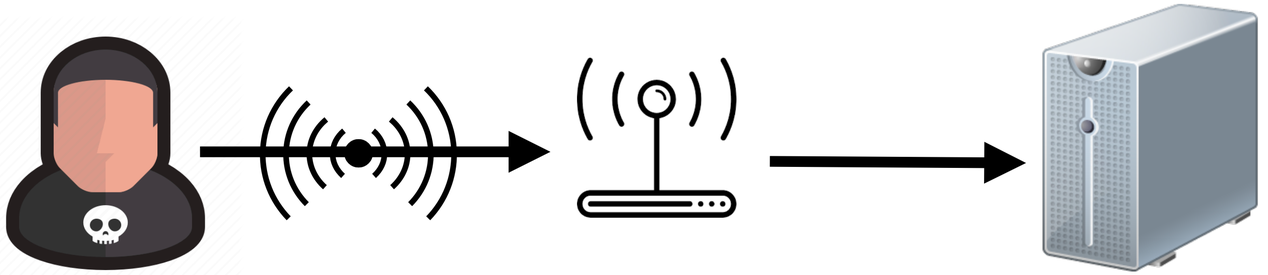

2. cross-site timing attack

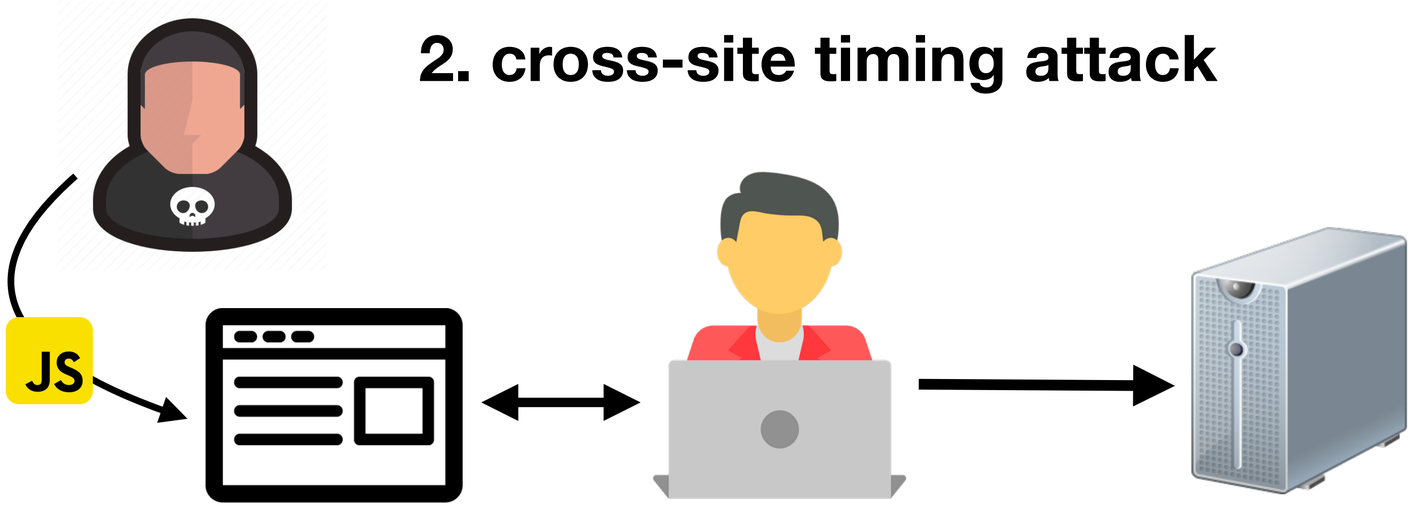

Conclusion

- Concurrency-based timing attacks are **not affected by network jitter** at all
- Perform **remote** timing attacks with an **accuracy similar to** an attack against the **local system**
- Attacks can be launched against protocols that feature **multiplexing** or by leveraging a transport protocol that enables **encapsulation**
- All **protocols that meet the criteria** can be **susceptible to concurrency-based** timing attacks: we created practical attacks against **HTTP/2**, **EAP-pwd** (Wi-Fi), HTTP/1.1 over **Tor**
- Future work: extensive evaluation of network protocols on susceptibility of attacks


Questions?
@tomvangoethem
tom.vangoethem@cs.kuleuven.be