



# DiSC: Backpressure Mitigation In Multi-tier Applications With Distributed Shared Connection

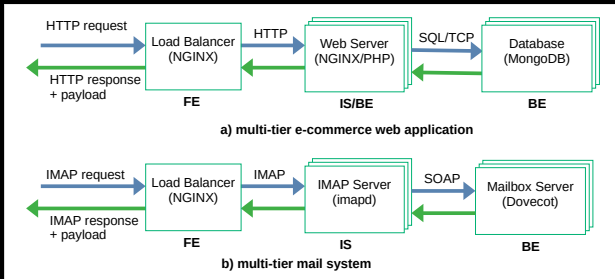
---

Alain Tchana - KrakOS/Inria/Grenoble INP-UGA

NSDI 2025 - April 29 - Philadelphia (USA)

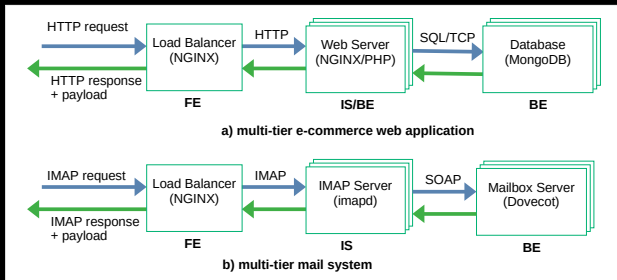
Joint work B. Ekane, D. Mvondo, R. Lachaize, D. Bromberg  
and D. Hagimont

# Motivating Examples




# Final Data

- ▶ **Final** payload (e.g., mails, attached files) are never used by intermediate servers (IS) and the frontend (FE)
- ▶ **Relaying** final data
  - ⇒ **Backpressure** on IS and FE
  - ⇒ **Scalability** issues



# Backpressure


- ▶ **Final** payload (e.g., mails, attached files) are never used by intermediate servers (IS) and the frontend (FE)
- ▶ **Relaying** final data  
⇒ **Backpressure** on IS and FE



**Higher load** on FE and IS, compared to BE

# Scalability Issues

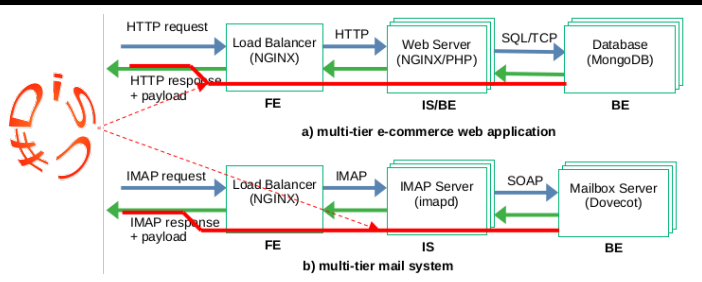
- ▶ **Final** payload (e.g., mails, attached files) are never used by intermediate servers (IS) and the frontend (FE)
- ▶ **Relaying** final data  
 ⇒ **Scalability**



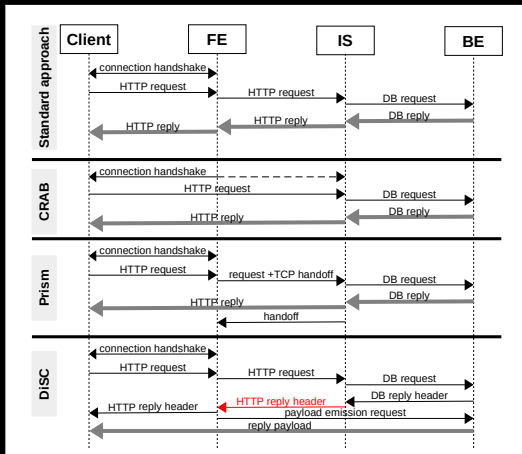
FE **sturates first**

# DiSC: Bypass + Distributed TCP

- ▶ **Final** payload (e.g., mails, attached files) are never used by intermediate servers (IS) and the frontend (FE)
- ▶ **Bypass** IS and FE for final data



# State-of-the-art



**Connection handoff:** DSR, CRAB [ATC'20] and Prism [NSDI'21], QDSR [ATC'24]

# Main limitations

---

- ▶ Only bypass a single server (LB)
- ▶ All servers must implement the same protocol (e.g., HTTP)
- ▶ Fully bypass the LB
  - ▶ LB (API gateway) collect stats, provide authentication, manage billing, implement smart load balancing algo.

# Behind DiSC

---

- ▶ A response is composed of
  - ▶ header+payload+footer
- ▶ Bypass only applied to the payload
- ▶ Header and footer follow the normal communication path


# Design goals

---

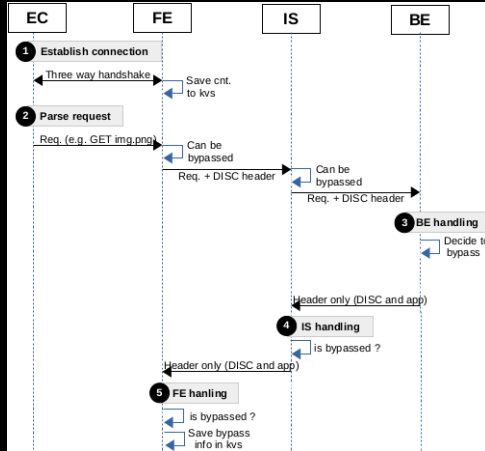
- ▶ Transparent to clients
- ▶ Fine-grained bypass
- ▶ Modest modification for the application
- ▶ Support all application-level protocols
- ▶ Keep the storage at the same security level
- ▶ Support TLS

# DiSC

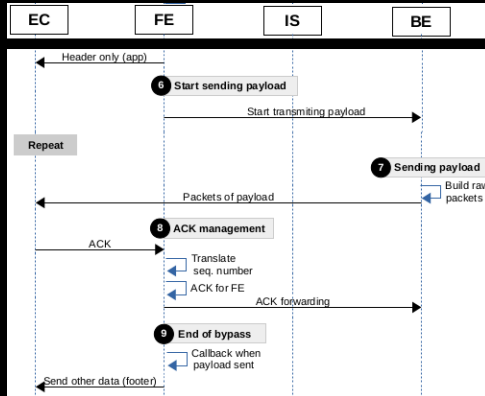
- ▶ DiSC-PROT, a new protocol
  - ▶ A node can indicate that it can be bypassed
  - ▶ Lies between TCP and the application-level protocol
- ▶ Main techniques
  - ▶ Transparent (BPF) Address spoofing, ACK redirect, Seq number rewriting, TLS session serialization-export



# End-to-end coordination with DiSC



# End-to-end coordination with DiSC




# Evaluation

---



- ▶ Testbed: CloudLab
- ▶ Micro-benchmarks: Nginx-based n-tier web app that returns static contents
- ▶ Macro-benchmarks: SpecWeb, SpecMail, Train Ticket, and Social Media (DeathStarBench suite)

# Load offloads from FE to BE




**DiSC** efficiently offloads load from FE to BE, almost as **DSR**

# Good scalability



# Better tail latency



# More results



## DISC: Backpressure Mitigation in Multi-tier Applications with Distributed Shared Connection

Brice Ekam, Dyob Myonko, Renaud Lachaise,  
Yérou-David Bemborg, Alain Tchana, Daniel Haymann,  
Eric Armes, Inria, CNRS, DISA, France  
INRT, Université de Toulouse, CNRS, Toulouse INP, UT2 Toulouse, France

### Abstract

Most data-center applications are based on a multi-tier architecture, involving either custom-provided software components (e.g., web servers) or web applications (e.g., cloud services). Such applications are prone to resource contention. In this paper, we propose a novel backpressure mitigation, DISC, that addresses resource contention. This problem is due to the fact that, in a multi-tier architecture, the requests are not processed in a FIFO order. This leads to a significant increase in the response time of the requests. In addition, this leads to a significant increase in the response time of the requests. In addition, this leads to a significant increase in the response time of the requests.

We have recently completed and communicated through preprint using protocols such as HTTP, SOAP, REST, etc. In the remainder of this paper, we use the terms "user" and "service" interchangeably, as our construction is agnostic to the number and granularity of these components.




Figure 1: A schematic view of a chain of tiers. Tier  $n$  requests data which is relayed by tiers  $1, \dots, n-1$  to the client.

...involves the generation and execution of application-level policies. DISC is (1) flexible — it allows the user to specify the number of tiers, (2) the number of tiers, (3) and the number of tiers. In addition, DISC can be used to mitigate resource contention in a multi-tier architecture. It is designed to be used in a multi-tier architecture. It is designed to be used in a multi-tier architecture. It is designed to be used in a multi-tier architecture.

### 1 Introduction

The reasons of modularity and, above all, scalability, Internet services are built in the form of an assembly of components in an architecture formerly referred to as multi-tier or tiered architecture. In such an architecture, components

The independence is a fundamental property that makes the design of multi-tier architectures. However, this independence is a challenge to address each tier independently. One of the most important challenges is performance degradation. One of the most important challenges is performance degradation. One of the most important challenges is performance degradation.

# This work

---

- ▶ Introduced DiSC
- ▶ Intermediate server bypass to directly transfer final data
- ▶ Reduces CPU utilization on intermediate servers
- ▶ Reduces tail latency
- ▶ Better scalability
- ▶ Meta-data are not bypassed
- ▶ Support TLS, no client adaptation, modest server modification, based on BPF
- ▶ Experiments on several apps (C, Java-based, traditional n-tier and microservices-based)

# Contact me!

---

- ▶ [alain.tchana@gmail.com](mailto:alain.tchana@gmail.com)
- ▶ Postdoc in Systems at `lig-krakos.imag.fr`
- ▶ Grenoble (Alpes, Ski) - France (+250 types of cheese)

