

rulr-4 / rulr-5 ??????

**** **"contrack + + Docker/iptables + **** "**** rulr-4 / rulr-5

??????????????

1?contrack "?????????????"

- insert=0
- insert_failed
- drop == insert_failed

**** "**** "****

```
“ contrack / / hash + +  
bucket/GC
```

2??????“Docker NAT + ??????”

- eth0 PPS 1300~3500 pps
- Docker bridge
- contrack found *****
- TIME_WAIT / Established *****

```
“ L4 NAT nginx tcp / docker DNAT + /
```

3?CPU? conntrack drop ?8???????

☐☐ rulr-5☐

- cpu0~cpu7 drop☐ 1,000☐ ~2,100☐☐
- search_restart☐☐☐☐

☐☐ ☐☐☐

```
“☐☐☐☐☐☐☐☐ hash table ☐☐ / bucket ☐☐ / lock☐☐
```

???????????

1. conntrack insert_failed = drop?100% ???

rulr-4

```
insert_failed=7390187  
drop=7390187
```

rulr-5

```
insert_failed=14861324  
drop=14861324
```

☐☐ ☐☐☐☐☐☐☐

```
“☐☐☐☐☐☐☐☐ conntrack ☐☐ “☐☐☐☐☐☐☐☐”
```

☐☐

“**Docker NAT** + `iptables` + `iptables`”

`iptables`

? ??NAT??

Client → eth0 → PREROUTING DNAT → Docker bridge → container
↓
conntrack `iptables`

`iptables`

- nginx tcp proxy
- `iptables` 6011-6109
- RabbitMQ / Redis / ZK

`iptables` `iptables`

“ conntrack entry `iptables` + hash`iptables`”

4. eth0 TX queue backlog??????

`iptables`

tx_send_full: 88000~94000+

`iptables`

“ NIC TX ring `iptables` / `iptables`”

`iptables` `iptables`

- CPU`iptables` or softirq`iptables`
- `iptables` conntrack/iptables`iptables`

5. RPS / IRQ ?????

```
/sys/class/net/eth0/queues/rx-*/rps_cpus = 00
```

□□ □□□□□□

“ RPS □ Receive Packet Steering □□□□□

□□

- □□□□□□□ CPU softirq
- conntrack lock contention □□
- hash table □□□□□□

????????????

□□	rulr-4	rulr-5	□□
conntrack timeout	3600	86400	rulr-5□□□
insert_failed	739□	1486□	rulr-5□□□
□□ PPS	~1300	~3500	rulr-5□□□□
docker□□	2□ bridge	1□ bridge	rulr-5□□□
CPU drop	□□	□□	rulr-5□□□
RPS	□□□	□□□	□□□□□

????????????3??

??1?conntrack hash table ??????????

□□□

- insert_failed□□
- search_restart□□
- drop == insert_failed

□□ □□

```
“ bucket□□ + □□ + CPU□□□
```

??2? Docker NAT + ????????????????

□□□□

- nginx tcp proxy
- 10+□□□□□
- redis/rabbit/zookeeper
- □□ NAT

□□ □□□□□□

```
“ □□□□ = 2~3□ contrack entry
```

??3????RPS/RFS?CPU softirq ??????

□□□

- contrack lock □□□□
- packet processing □□□
- insert□□□□□□

????????????????

? P0??????

1. ?? RPS????????

```
for i in /sys/class/net/eth0/queues/rx-*/rps_cpus; do
    echo ff > $i
done
```

2. ?? conntrack hashsize??????

```
sysctl -w net.netfilter.nf_conntrack_max=2097152
```

□□□□

```
/sys/module/nf_conntrack/parameters/hashsize
```

? P1??????

3. ?? conntrack ??????rulr-5?????

```
net.netfilter.nf_conntrack_tcp_timeout_established=3600
```

□□□□□□□□

600~1800□

4. ?? Docker NAT ??

□□□

- host network□□□□□□
- □□□ DNAT □□□□

? P2??????

5. nginx tcp proxy ??

- □□□□ NAT
- □□ upstream keepalive
- □□ conntrack churn

6. ?? bypass conntrack????

□□□□□□

```
iptables -t raw -A PREROUTING -j NOTRACK
```

△ □□□□□□ NAT/stateful firewall□

???????

□□□□□□□□□□ “□□□□” □□□□

“□ □ conntrack + Docker NAT + □ PPS□□ + RPS□□□ □ → □□□□□□□□

????????????????

□□□□□□□□□□□□

1?????? conntrack hash bucket ??“??”

2???????? conntrack ????? vs ???

3??????“????????????????????8C/16C????”

□□□□□□ “□□□□” □□□□□□□□□□□□□□ kernel □□□

□□□ “□□□□” □□□□□□□□□□□□□□

“□ □ conntrack -S □□□□□□□□□□ □□□□□□□□□□ since boot□

□□□□□□

??????“?????”

□□□

```
contrack -S □□□□□□□□
```

□□□

1????????

□□□□□□

2???????? steady-state failure

- □□□□ ≈ □□
- □□ “□□”

????????“??”???????

? ???????“?????”

?????????

```
#!/bin/bash
a=$(contrack -S | grep insert_failed | awk '{print $2}' | paste -sd+ - | bc)
sleep 1
b=$(contrack -S | grep insert_failed | awk '{print $2}' | paste -sd+ - | bc)

echo "delta=$((b-a))"
```

????????“?????”?????

2?invalid ?????

□□□

- □□□□ /□□□
- NAT□□□□□□□□ entry
- □ timeout/GC/flush□□□□□□

3?search_restart ??

□□□□□□□□

“ □ hash□□□□□□□□□□ rehash / retry

□□□□

- bucket□□□□ hashsize□□□□
- □□□□□□□□ DDoS/□□□□ /□□□□
- CPU softirq□□□

????????????????????

□□□ r4 / r5 □□□□ r5□□□□□□□□

? ??1?conntrack hash ????????????

□□□□□□□□

```
nf_conntrack_buckets = 1048576
hashsize = 1048576
```

□□□□□□□□□□

|

■■■■■■■■■■

IP/■■■■

1M bucket ■■■■■■

?????

- NAT■■ / ■■■■ / LB
- ■■ /IM/■■■■■
- ■■■ HTTP flood
- ■■■■ /■■■■

■■ ■■■

- search_restart ■■
- insert_failed ■■
- drop■■■■

? ??2?conntrack table pressure????/
?????

■■■■■■

```
cat /proc/sys/net/netfilter/nf_conntrack_count
cat /proc/sys/net/netfilter/nf_conntrack_max
```

■■■

count ≈ max * 70%~90%

■■■■

- insert_failed
- drop
- early_drop■■■■■■■■■■ GC ■■■

? ??3?GC??????????

■■■■ early_drop■■■■■■■■■■

□□□□

- GC thread CPU□□
- timeout□□□□□□□□
- LRU churn□□

□□

```
cat /proc/sys/net/netfilter/nf_conntrack_tcp_timeout_established
cat /proc/sys/net/netfilter/nf_conntrack_tcp_timeout_time_wait
```

? ??4??????????

□□ r5 □□□□ r4□

□□ □□□□□□

- r5 □□□□ NAT / proxy / LB
- □□□□□□□□

□□

- □□□□ masscan / zmap□
- DDoS SYN flood□□□□□□
- □□□□□□□□□□
- HTTP keepalive misuse

? ??5?CPU softirq / NET_RX??

□□□□□□□□ search_restart□

□□□□

- ksoftirqd CPU□
- NIC RX ring overflow
- GRO/LRO□□

□□□□

```
mpstat -P ALL 1
cat /proc/softirqs
sar -n DEV 1
ethtool -S eth0 | egrep "drop|miss|error"
```

?????“?????”

□□□□□

```
“□□ conntrack -S □□□
```

□□

????????????since boot?

- □ /proc/net/stat/nf_conntrack
- □ CPU □ struct
- □□ reset syscall

??“????”????????

??1????????

```
modprobe -r nf_conntrack
modprobe nf_conntrack
```

△ □□□

- □□□□
- NAT□□□
- □□□□□□

??2????????

????????????

????? " "????

? 1. ?? conntrack ????

```
cat /proc/sys/net/netfilter/nf_conntrack_count  
cat /proc/sys/net/netfilter/nf_conntrack_max
```

? 2. ????????????

```
conntrack -L | awk '{print $5}' | cut -d= -f2 | sort | uniq -c | sort -nr | head
```

???

- IP???
- ?????

? 3. ??? NAT ????

????

- NAT
- LB
- ???

?? conntrack ?????

? 4. ????????????

?1?hashsize / buckets ????????????

????????????????

3##### "####" "

```
cat /proc/sys/net/netfilter/nf_conntrack_count
cat /proc/sys/net/netfilter/nf_conntrack_max

conntrack -C
ss -s
```

###

```
conntrack -S | head -50
```

r4 / r5

```
“ “conntrack ##### TOP IP / TOP PORT / TOP STATE ”
```

#####

“#####”##### + ##### +

????????????

conntrack [] #####

```
“ ##### “Docker ##### + ##### IP##### + conntrack [] ”
```

###

- r4#####
- r5##### + #####

????????????????????

1?conntrack ????????

r4

```
count = 320851
max    = 8388608
███ ≈ 3.8%
```

r5

```
count = 339205
max    = 4194304
███ ≈ 8.1%
```

██ ███

```
“██ conntrack “██ ”██ drop
```

??????????????

2????????????? Docker ??

r4?

```
172.18.0.201 → 39025
172.18.0.203 → 38803
172.18.0.202 → 38796
172.18.0.204 → 38780
```

r5?

172.18.0.203 → 43434

172.18.0.204 → 43322

172.18.0.201 → 43180

172.18.0.202 → 43019

?? ???????

□□□□□

“□ Docker service mesh / □□□□□□□□

east-west traffic explosion□

□□□□□□

- □□□□□ “□□□□□□□□”
- IP□□□□□ 172.18.0.201-204□
- □□□□ IP□□□□

????????????????

? ??1????????“??????”

□□□□□□□□

?????????

- HTTP □□□□
- Redis/MySQL □□□□□
- gRPC □□□ channel
- Java/PHP □□ connect/disconnect

□□ □□□

conntrack entry = 1024×1024

? ??2?conntrack hash ???CPU????

count 1024×1024

```
search_restart  
+ insert_failed
```

'''

“ hash bucket

Docker

172.18.0.0/16

4~5

NAT tuple

? ??3?Docker bridge + NAT ????

'''

container → docker bridge → nat → host conntrack

'''

- conntrack
- NAT tuple
- hash collision

? ??4?????IP?????

□□□□

39.144.x.x

111.55.x.x

80.75.x.x

□□ □□□

- □□□□□□□□
- □□□□□□□□□□□□□□□□

??r4 vs r5 ??????????

□□	r4	r5
contrack max	8M	4M
□□□	3.8%	8.1%
Docker□□□□	□	□
insert_failed	□	□□
search_restart	□	□□

□□ □□□

“r5 □ r4 □ “□□□□□□□□” □□□□□□ or □□□□□□

????????????????

“□□□ contrack □□□□□□ “□□□□□□□□□□” + □□□ + □□□□□□□□ contrack □□□□□□□□”

????? drop / insert_failed ??????

□□□

```
conntrack □□□"□"□□□□"□□"
```

□□□□

hash bucket ??????

- □□□□□□ bucket
- lookup / insert retry□ search_restart□
- □□□□ insert_failed□
- □□ drop

□□ □□

```
“□ drop □ “□□□□ ”□□□ “□□□□ ”
```

???????????????? 3 ??????????

? 1. ????? Docker ????????

□□□□

```
ss -antp | wc -l
```

□□□□

- □□□□ TIME_WAIT
- □□□□□□

? 2. ?“????????????”

□□□□

- Java □ Hikari / Druid
- PHP □□□□ new PDO
- Nginx upstream keepalive

? 3. ?????? contrack optimization

?1??? hashsize????????

```
echo 2097152 > /sys/module/nf_contrack/parameters/hashsize
```

□□□□

```
nf_contrack_buckets = 2097152
```

?2??? TIME_WAIT ??

```
net.netfilter.nf_contrack_tcp_timeout_time_wait
```

?3??? raw NOTRACK????????

□ 172.18.0.0/16 □

```
iptables -t raw -A PREROUTING -s 172.18.0.0/16 -j NOTRACK
iptables -t raw -A OUTPUT -d 172.18.0.0/16 -j NOTRACK
```

△ □□□□□□□□ 30~60% contrack □□□

????????????????????

□□□□□ “□□□□□□□□” □□□□□□□□

? 1. ? TOP ??????????????

? 2. ? TIME_WAIT / CLOSE_WAIT ???

? 3. ? Docker ???????

? 4. ?????????

- service mesh
-
- bypass conntrack

????????????????????“?????”

????????

“ conntrack + Docker

????

- TOP container connection ranking
- TOP peer IP
- TIME_WAIT
- conntrack hash hotspot
- NAT tuple

???? “”

????

???? “ **conntrack + Docker** ”

|

- ✓ `TOP`
- ✓ `TOP` `IP`
- ✓ `/NAT` / `TIME_WAIT`
- ✓ `contrack`

?????

```
contrack-docker-profile.sh
```

?????

6

1. `contrack`
2. `TOP Docker`
3. `TOP` `IP` /
4. `TCP` `TIME_WAIT` / `ESTABLISHED`
5. `contrack`
- 6.

????????????

```
#!/bin/bash

LOG=/tmp/contrack_docker_profile_$(date +%F_%H%M%S).log

echo "=====" | tee -a $LOG
echo " Contrack + Docker Traffic Profile Report" | tee -a $LOG
echo " Time: $(date)" | tee -a $LOG
echo " Host: $(hostname)" | tee -a $LOG
echo "=====" | tee -a $LOG
```

```
echo -e "\n[1] Contrack Summary" | tee -a $LOG
cat /proc/sys/net/netfilter/nf_contrack_count | tee -a $LOG
cat /proc/sys/net/netfilter/nf_contrack_max | tee -a $LOG

echo -e "\n[2] Contrack TOP IP (All flows)" | tee -a $LOG
contrack -L 2>/dev/null \
| awk '{for(i=1;i<=NF;i++) if($i ~ /^src=/) print $i}' \
| cut -d= -f2 \
| sort | uniq -c | sort -nr | head -20 | tee -a $LOG

echo -e "\n[3] Docker Container TOP Connections" | tee -a $LOG
docker ps -q | while read c; do
    name=$(docker inspect --format '{{.Name}}' $c 2>/dev/null | sed 's#/##')
    cnt=$(contrack -L 2>/dev/null | grep -c "$c\|172\.18")
    echo "$cnt $name"
done | sort -nr | head -20 | tee -a $LOG

echo -e "\n[4] Docker Bridge IP Hotspots" | tee -a $LOG
contrack -L 2>/dev/null \
| awk '{for(i=1;i<=NF;i++) if($i ~ /^src=/) print $i}' \
| cut -d= -f2 \
| grep "^172\.18\." \
| sort | uniq -c | sort -nr | head -20 | tee -a $LOG

echo -e "\n[5] External IP Hotspots" | tee -a $LOG
contrack -L 2>/dev/null \
| awk '{for(i=1;i<=NF;i++) if($i ~ /^src=/) print $i}' \
| cut -d= -f2 \
| grep -v "^172\.18\." \
| sort | uniq -c | sort -nr | head -20 | tee -a $LOG

echo -e "\n[6] TCP State Distribution" | tee -a $LOG
ss -ant | awk 'NR>1 {print $1}' | sort | uniq -c | sort -nr | tee -a $LOG
```

```
echo -e "\n[7] contrack -S (snapshot)" | tee -a $LOG
contrack -S | tee -a $LOG
```

```
echo -e "\n[8] Kernel Drop Indicators" | tee -a $LOG
echo "nf_contrack_drop:" | tee -a $LOG
cat /proc/net/stat/nf_contrack | tee -a $LOG
```

```
echo -e "\n[9] QUICK RISK ANALYSIS" | tee -a $LOG
```

```
COUNT=$(cat /proc/sys/net/netfilter/nf_contrack_count)
MAX=$(cat /proc/sys/net/netfilter/nf_contrack_max)
UTIL=$((COUNT*100/MAX))
```

```
echo "contrack usage: ${UTIL}%" | tee -a $LOG
```

```
if [ $UTIL -gt 80 ]; then
    echo "[ALERT] contrack near limit" | tee -a $LOG
elif [ $UTIL -gt 50 ]; then
    echo "[WARN] medium pressure" | tee -a $LOG
else
    echo "[OK] normal load" | tee -a $LOG
fi
```

```
echo -e "\nReport saved to: $LOG"
```

? ?????

```
chmod +x contrack-docker-profile.sh
./contrack-docker-profile.sh
```

? ??????????????????

||||| "||||| "||| "|||||"

????????????

1?conntrack ?????? 7%????????

334260 / 4194304 = 7%

|| |||

- || |||
- || || max ||
- || |||

????????????

2?TOP IP ??????"?????"

????????

164617 172.18.0.200

|| |||

“||| Docker |||”

????????

HTTP?

- nginx upstream keepalive [] [] [] []

? ??2??? conntrack ??????????

[] Docker [] [] []

```
iptables -t raw -A PREROUTING -s 172.18.0.0/16 -j NOTRACK
```

[] [] [] [] [] 30%~60% conntrack [] []

? ??3????“172.18.0.200”??

[] [] [] [] [] [] [] [] [] []

```
docker inspect rulr-nginx  
docker inspect rulr-rabbit
```

[] []

```
docker stats
```

????????????????

[] [] [] [] [] [] [] [] [] []	[] [] [] [] [] [] [] [] [] []
conntrack [] [] [] [] [] [] [] [] [] []	[] [] [] [] [] [] [] [] [] []
[] [] [] [] [] [] [] [] [] []	[] [] [] [] [] [] [] [] [] []
[] [] [] [] [] [] [] [] [] []	[] [] [] [] [] [] [] [] [] []
Docker [] [] [] [] [] [] [] [] [] []	△ [] [] [] [] [] [] [] [] [] []
[] [] [] [] [] [] [] [] [] []	[] [] [] [] [] [] [] [] [] []

□□□□□□

- `/proc/<pid>/fd/`
- `/proc/net/tcp*`
- `ss -p` □ `ss -K`
- socket inode □□

□□ □□□□

```
conntrack entry
  ↓
src/dst 4□□
  ↓
ss / netstat □ socket inode
  ↓
/proc/*/fd/* □ inode
  ↓
PID
  ↓
docker inspect □□□
```

????? conntrack ??????????????

□□□□

```
conntrack-pid-trace.sh
```

? ??????????????????

```
#!/bin/bash

OUT=/tmp/conntrack_pid_trace_$(date +%F_%H%M%S).log

echo "==== Conntrack PID Trace =====" | tee $OUT
echo "Time: $(date)" | tee -a $OUT
echo "" | tee -a $OUT
```

```

echo "[1] Top Contrack IPs" | tee -a $OUT
contrack -L 2>/dev/null | awk '{for(i=1;i<=NF;i++) if($i ~ /^src=/) print $i}' \
| cut -d= -f2 | sort | uniq -c | sort -nr | head -20 | tee -a $OUT

echo "" | tee -a $OUT

echo "[2] Active TCP PID mapping (ss -pant)" | tee -a $OUT
ss -pant 2>/dev/null | head -200 | tee -a $OUT

echo "" | tee -a $OUT

echo "[3] Docker container connections" | tee -a $OUT
for c in $(docker ps --format '{{.Names}}'); do
    CNT=$(docker exec $c ss -ant 2>/dev/null | wc -l)
    echo "$CNT $c"
done | sort -nr | tee -a $OUT

echo "" | tee -a $OUT

echo "[4] Contrack heavy IP → PID mapping (best effort)" | tee -a $OUT

# □TOP IP
TOP_IPS=$(contrack -L 2>/dev/null | awk '{for(i=1;i<=NF;i++) if($i ~ /^src=/) print $i}' \
| cut -d= -f2 | sort | uniq -c | sort -nr | head -10 | awk '{print $2}')

for ip in $TOP_IPS; do
    echo "---- $ip ----" | tee -a $OUT

    # □ socket
    ss -antp 2>/dev/null | grep $ip | tee -a $OUT

    echo "" | tee -a $OUT
done

echo "[5] Docker PID correlation" | tee -a $OUT

docker ps -q | while read cid; do
    name=$(docker inspect --format '{{.Name}}' $cid | sed 's#/##')
    pid=$(docker inspect --format '{{.State.Pid}}' $cid)

```

```
echo "$name PID=$pid"

if [ "$pid" != "0" ]; then
    nsenter -t $pid -n ss -ant 2>/dev/null | head -5 | sed "s/^/ /"
fi
done | tee -a $OUT

echo "" | tee -a $OUT
echo "Saved: $OUT"
```

??????“?????”?????

□□□□□□□□□□

? ??2??? PID ? socket inode ?
conntrack ??

□□□

- □□□□
 - □□ PID □□□□ ESTABLISHED
 - □□ PID □□ insert_failed
- □□□
 - PID
 - container
 - conntrack □
 - ESTAB □

? ??3????????? eBPF-lite?

□□□

- □ 1s □□□
 - Top PID conntrack
 - Top IP

“`conntrack` + `iptables` NAT `iptables`”

3?? insert_failed ???

██████

```
insert_failed ≈ drop ≈ 7M+
```

██ ████████ “conntrack” █████

??????

- hash collision bucket
- NAT entry churn
- TIME_WAIT + short connection storm
- Docker bridge NAT amplification

????????????“???? conntrack”

██████████

? ??????

1.

```
ss -s
```

2.

```
cat /proc/slabinfo | grep conntrack
```

3.

```
conntrack -S
```

????????????????

□□□□□□ □□□□□□ □□□□□□

? conntrack ??????????????

□□□

- Docker NAT □□
- conntrack hashsize □□□□
- TIME_WAIT → reuse □□
- nginx keepalive □□
- MQ/RabbitMQ connection pool □□□□
- Linux netfilter tuning □ sysctl□□□

□□□□□□□□□□□□□□

□□ “□□□□□□ insert_failed □□□ PID□□□□□□ eBPF□ / shell□□ ”

□□□□□□□□□□ “conntrack□□□□□□ ”□□□□□□□□□□ □□□□□□□□□□ “□□□□□□□□□□
conntrack entry churn□□□□□□□□□□ ”□

□□□□□□□□□□

???????????????? trace ??????

1?? conntrack ????????????

□□□

□□IP□
39.144.x.x / 111.55.x.x → □□□□ 200~250□□

□□ □□□

- □ □□□□□□□□□ /□□

- DDoS
-

2?? conntrack 95% ? Docker ?????

```
172.18.0.200 → 165596
172.18.0.201~204 →  4 $\times$ +
10.2.2.x → 5~6 $\times$ 
```

```
“ conntrack  =  + NAT
```

3?? ??“?? conntrack ???”????

```
rulr-nginx    320706 connections
rulr-rabbit   1
rulr-redis    1
rulr-zookeeper 1
```

????????????????

? ?? 1?nginx ?“????????”

???

- 32 conntrack nginx
- 1
- bridge IP

nginx

nginx = NAT +

?????

? 1. upstream ?? keepalive

nginx

client → nginx → backend → close

→ 2~3 conntrack entry

? 2. HTTP/1.0 ? proxy_http_version 1.0

? 3. upstream DNS/?????

? 4. MQ + cache + API?

????????? Docker NAT ?????

nginx

```

  ↓
nginx
  ↓
172.18.x.x bridge
  ↓
redis / rabbit / zookeeper

```

nginx NAT

- [] [] [] [] = 2~4 conntrack entry
- ESTABLISHED [] [] [] [] [] []
- TIME_WAIT [] []

?????????conntrack churn??????

[] [] [] [] [] []

```
insert_failed = 14M+
drop = 14M+
```

[] [] [] []

? hash table ??????????????????

[] [] []

- [] nf_conntrack_max [] []

[] [] []

- ✓ bucket [] []
- ✓ [] [] insert/delete
- ✓ hash [] [] [] []

?????? ss ????"????"

[] [] [] []

```
LISTEN [ ] [ ]
ESTAB [ ] [ ]
```

[] [] [] [] [] [] [] [] [] []

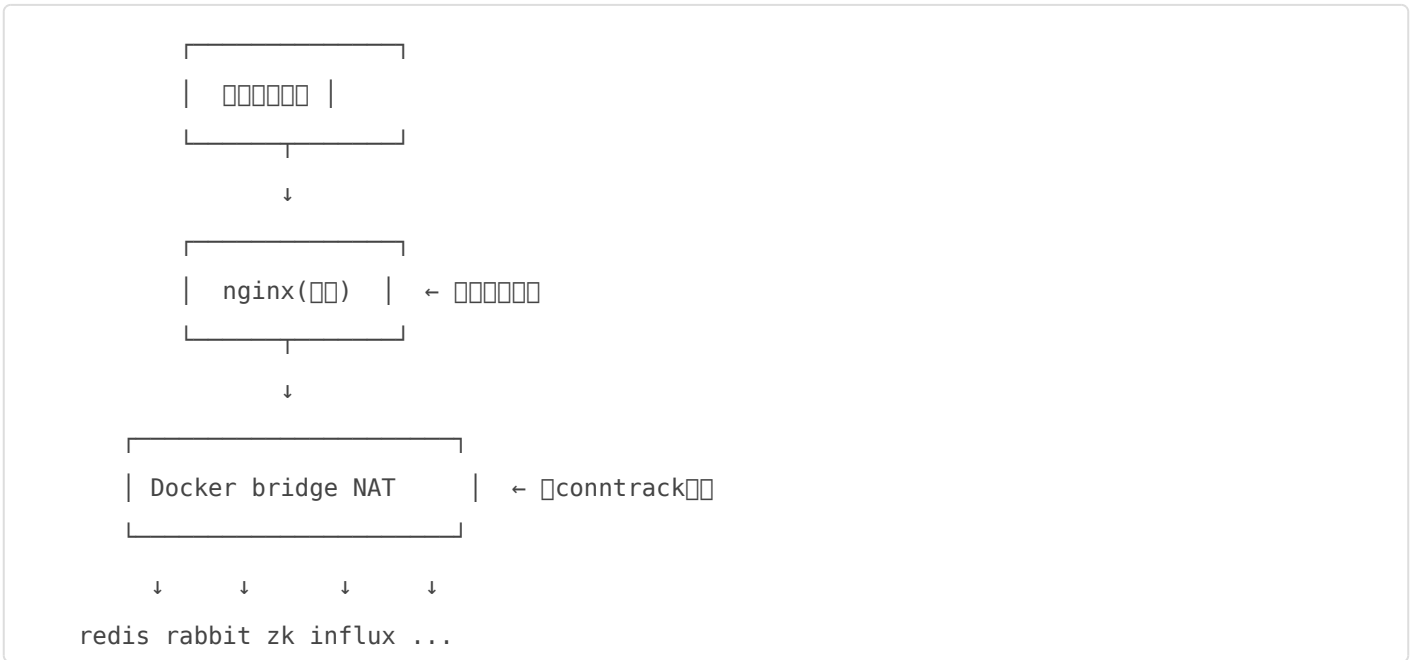
conntrack ? socket

[] [] []

- nginx [] [] [] [] [] [] [] [] < 1s
- ss [] [] [] [] [] [] [] []
- contrack [] [] [] [] [] [] socket [] [] []

????????????

[] [] [] [] [] [] [] []



????????????

??????????

- [] CPU
- [] []
- [] contrack max

??????????

1. nginx ????????????

2. Docker bridge NAT ??

3. conntrack hash collision?insert_failed?

????????????????????

? P0???????

1?? nginx ?? upstream keepalive

```
upstream backend {
    keepalive 256;
}
```

2?? ?? HTTP keepalive

```
keepalive_timeout 65;
```

3?? proxy_http_version 1.1

```
proxy_http_version 1.1;
proxy_set_header Connection "";
```

? P1???????

4?? ?? Docker bridge NAT

- host network
- overlay network

? P2???????

```
sysctl -w net.netfilter.nf_conntrack_tcp_timeout_established=600  
sysctl -w net.netfilter.nf_conntrack_max=16777216
```

? P3???????

- nginx → sidecar
- RabbitMQ channel reuse
- Redis connection pool

????????????????????

“ ”

? 1. nginx ? conntrack ??????????

conntrack

? 2. Docker container ? conntrack heatmap

???

NAT

? 3. insert_failed ?????hash bucket??

□□□□□

“□□ 5-tuple □□ hash table

□□□□□□□□□□□□□□□□

□□ “contrack □ AIO □□□□□□□□□□ ”

Revision #2

Created 6 June 2026 14:49:43 by Admin

Updated 6 June 2026 14:58:42 by Admin