Netfilter im Großeinsatz

Maik Hentsche <maik@mm-double.de>

Alien8 <fb@alien8.de>

Chemnitzer Linux-Tage 2007

2007-03-04

Agenda

- Vorwort: Firewall-Konzepte
- Einführung in iptables/netfilter
- iptables-Development: A moving target
- Shorewall
- Logdaten
- Hochverfügbarkeit

Konzepte

Konzepte

- **Netzwerksicherheitspolicy durchsetzen**
- Paketfilter (Layer 2-7)
  - Paketeigenschaften, Paketraten, ...
- VPNs
- NAT (Network Address Translation)
- Bandbreitenmanagement
- Routing

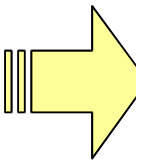

Firewalls sind ein Konzept!

Bestandteile eines Firewallsystems

- Policy
- Paketfilter oder Application Layer Gateway
- Regel Management
- Logdatenauswertung
- Alarmierung

Bestandteile

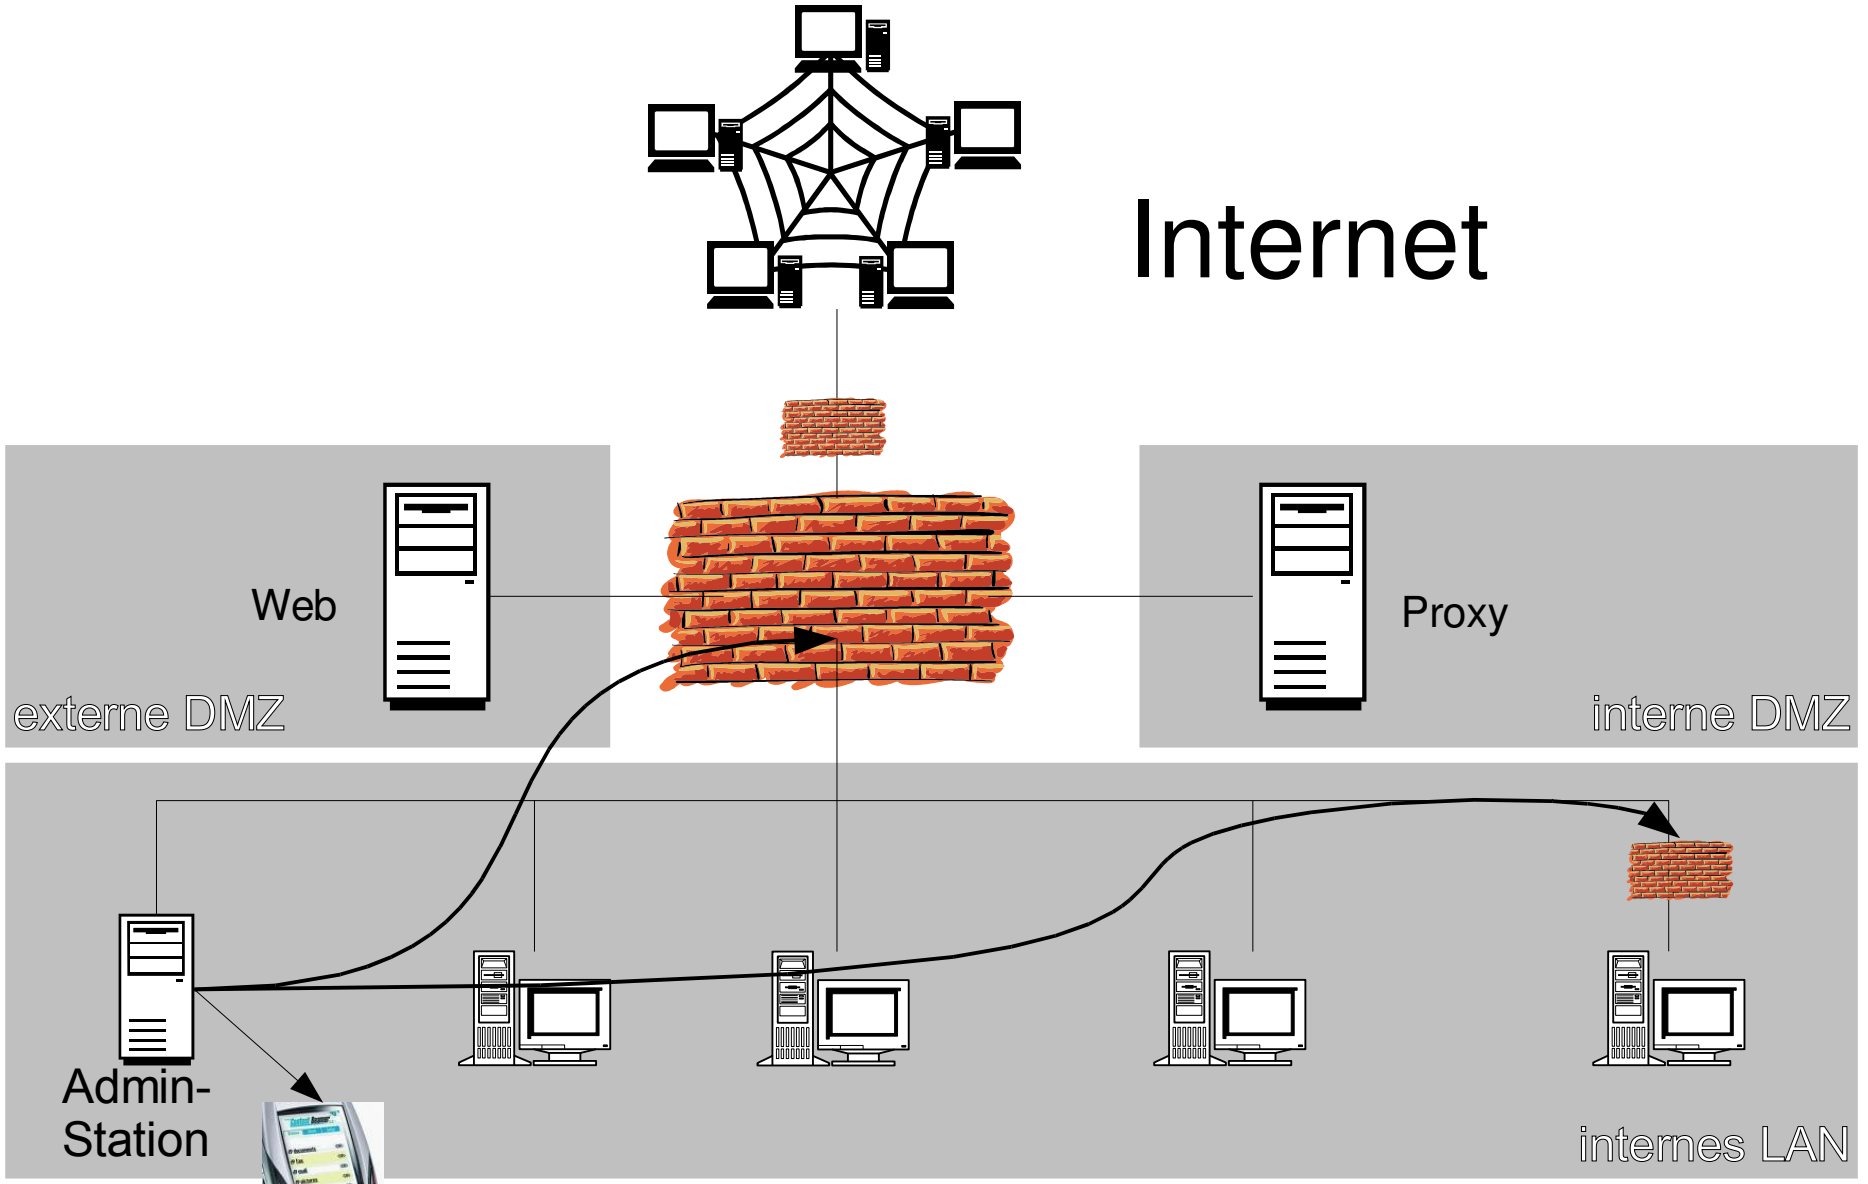

Iptables: Intro

Maiks erster Tag als iptables Commander

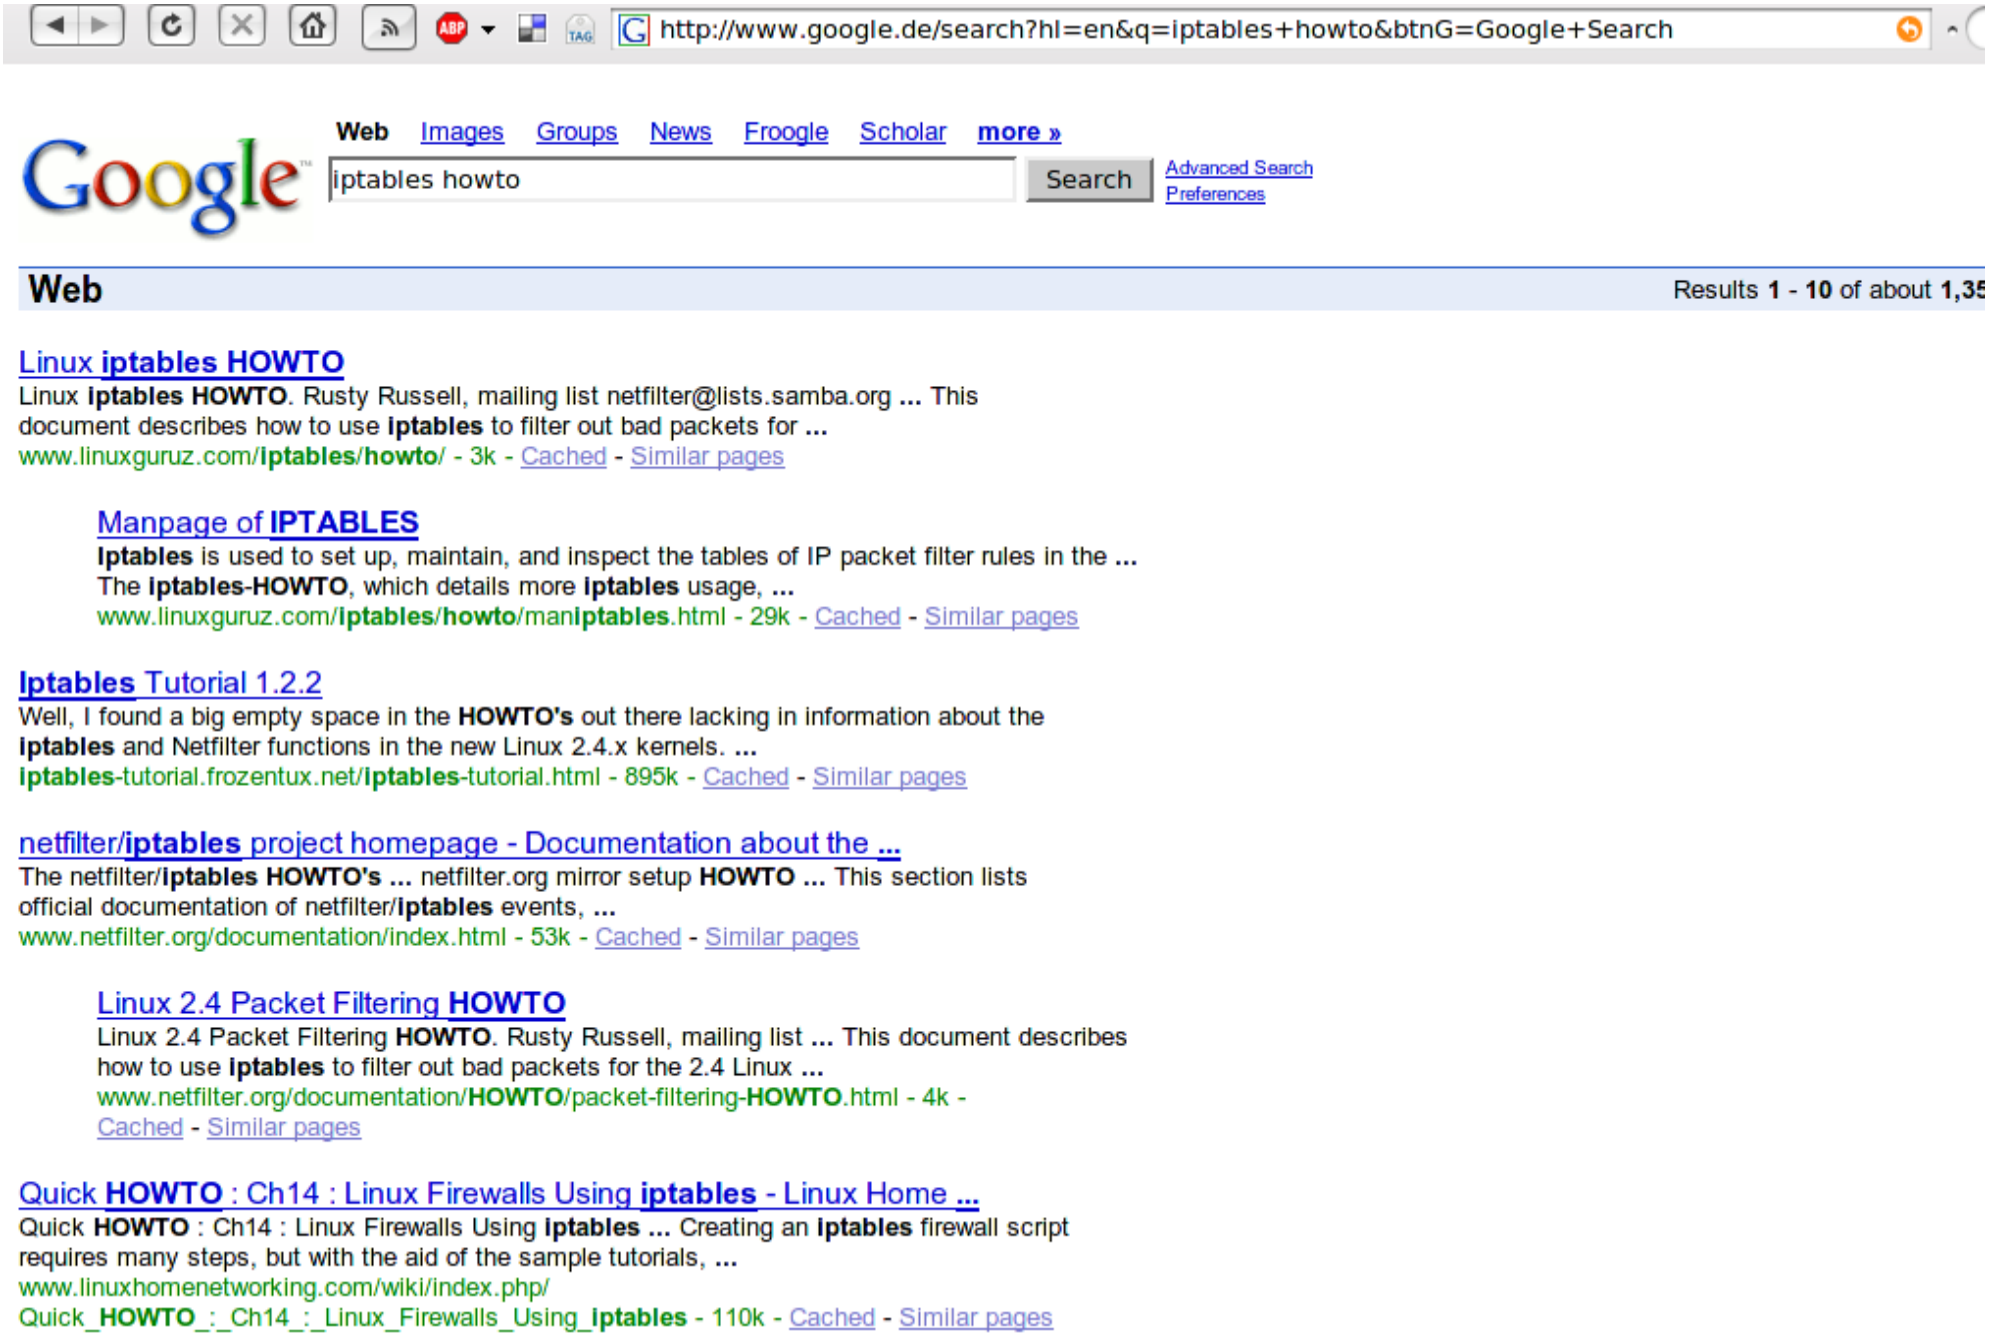
http://www.google.de/search?hl=en&q=iptables+howto&btnG=Google+Search
Google
Web Images Groups News Froogle Scholar more »
iptables howto
Search
Advanced Search
Preferences
Web
Results 1 - 10 of about 1,35
Linux iptables HOWTO
Linux iptables HOWTO. Rusty Russell, mailing list netfilter@lists.samba.org ... This document describes how to use iptables to filter out bad packets for ...
www.linuxguruz.com/iptables/howto/ - 3k - Cached - Similar pages
Manpage of IPTABLES
Iptables is used to set up, maintain, and inspect the tables of IP packet filter rules in the ... The iptables-HOWTO, which details more iptables usage, ...
www.linuxguruz.com/iptables/howto/maniptables.html - 29k - Cached - Similar pages
Iptables Tutorial 1.2.2
Well, I found a big empty space in the HOWTO's out there lacking in information about the iptables and Netfilter functions in the new Linux 2.4.x kernels. ...
iptables-tutorial.frozentux.net/iptables-tutorial.html - 895k - Cached - Similar pages
netfilter/iptables project homepage - Documentation about the ...
The netfilter/iptables HOWTO's ... netfilter.org mirror setup HOWTO ... This section lists official documentation of netfilter/iptables events, ...
www.netfilter.org/documentation/index.html - 53k - Cached - Similar pages
Linux 2.4 Packet Filtering HOWTO
Linux 2.4 Packet Filtering HOWTO. Rusty Russell, mailing list ... This document describes how to use iptables to filter out bad packets for the 2.4 Linux ...
www.netfilter.org/documentation/HOWTO/packet-filtering-HOWTO.html - 4k - Cached - Similar pages
Quick HOWTO : Ch14 : Linux Firewalls Using iptables - Linux Home ...
Quick HOWTO : Ch14 : Linux Firewalls Using iptables ... Creating an iptables firewall script requires many steps, but with the aid of the sample tutorials, ...
www.linuxhomenetworking.com/wiki/index.php/Quick_HOWTO_:_Ch14_:_Linux_Firewalls_Using_iptables - 110k - Cached - Similar pages

iptables/netfilter Intro

http://linux-ip.net/nf/nfk-traversal.png

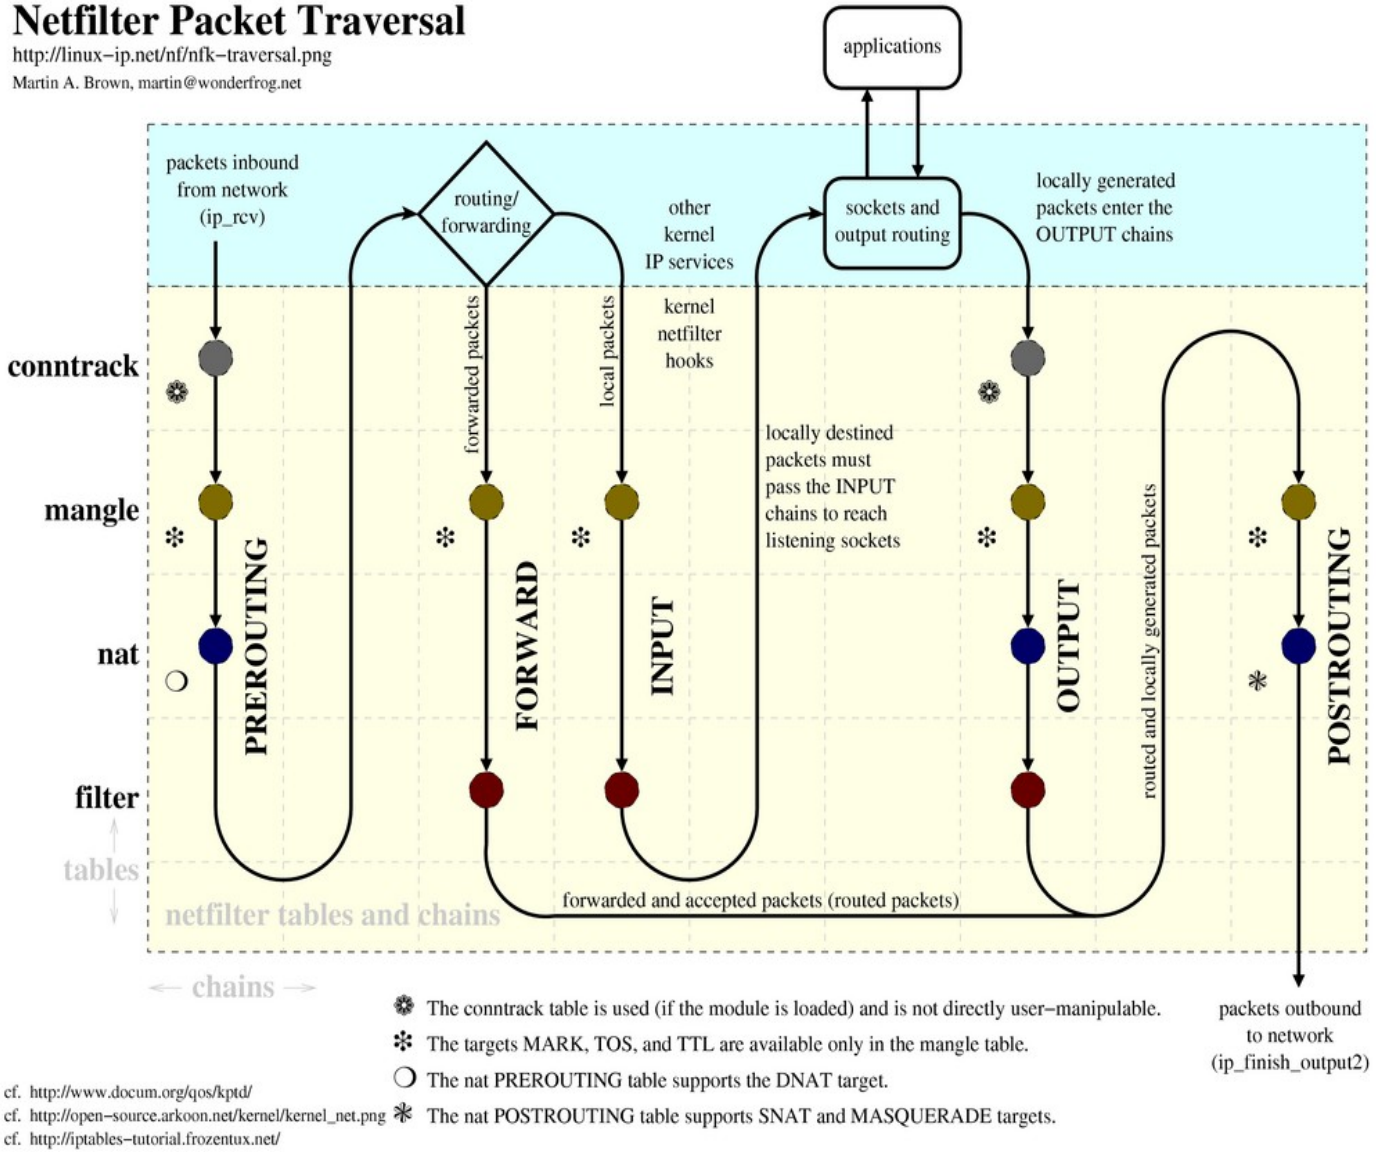

http://ebtables.sourceforge.net/br_fw_ia/PacketFlow.png

Regelgenerator:

Shorewall

Shorewall

http://www.shorewall.net

- Netfilter Administrationsstool
  - textbasiert (mehrere Konfigdateien)
  - Administrationskommandos
  - Flexible address management/routing support
    - NAT, MASQ, Proxyarp, NETMAP, mehrere ISPs
  - VPNs
  - Bandbreitenmanagement
  - Accounting
  - transparente Paketfilter

RTFM!

Shorewall Bsp Einzelplatz

- Konfiguration in /etc/shorewall
  - zones
  - interfaces
  - hosts
  - policy
  - rules

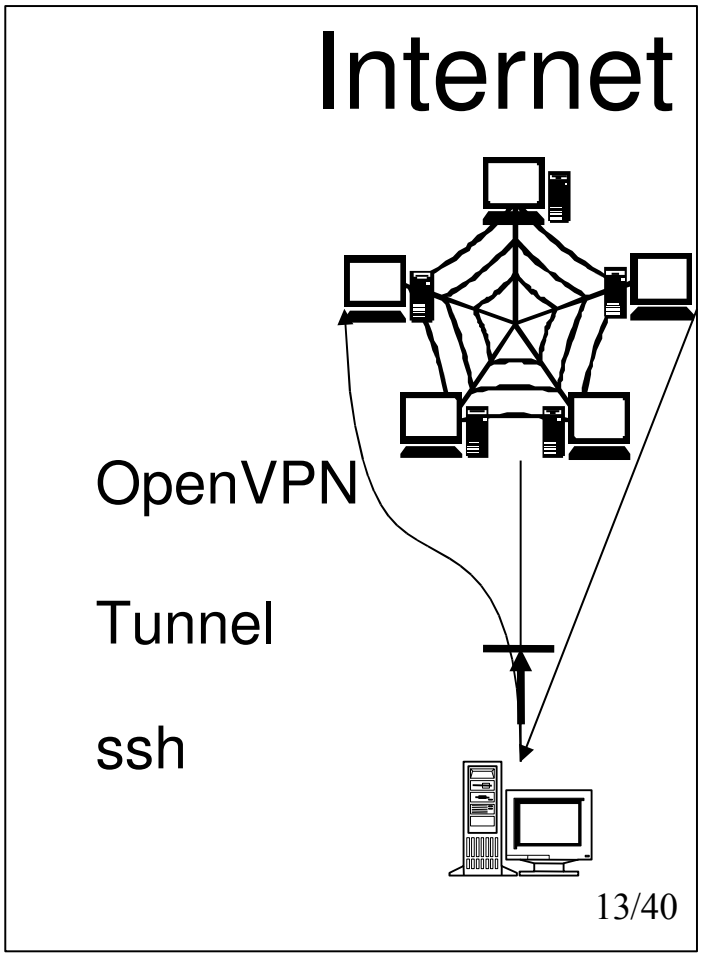

Shorewall Zonen-Konzept

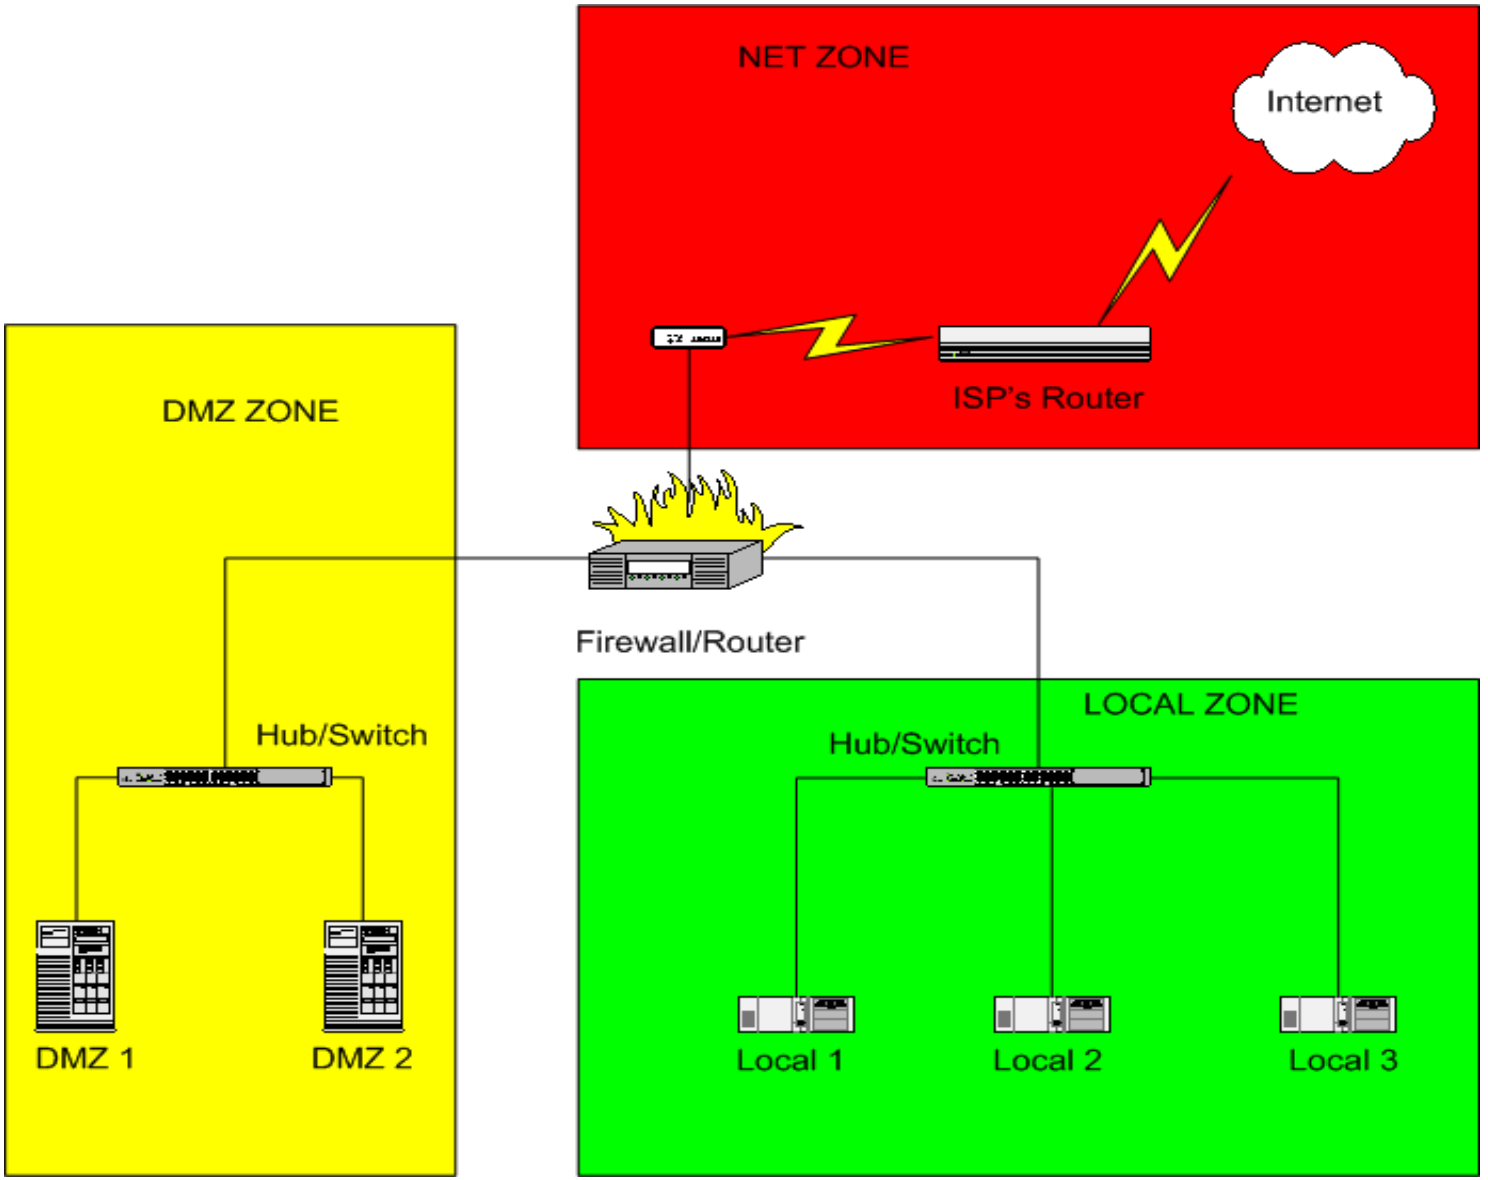

Shorewall Bsp: Server

- Konfiguration in /etc/shorewall
  - zones
  - interfaces
  - hosts
  - policy
  - rules
  - accounting
  - masq
  - tcdevices, tcclasses, tcrules

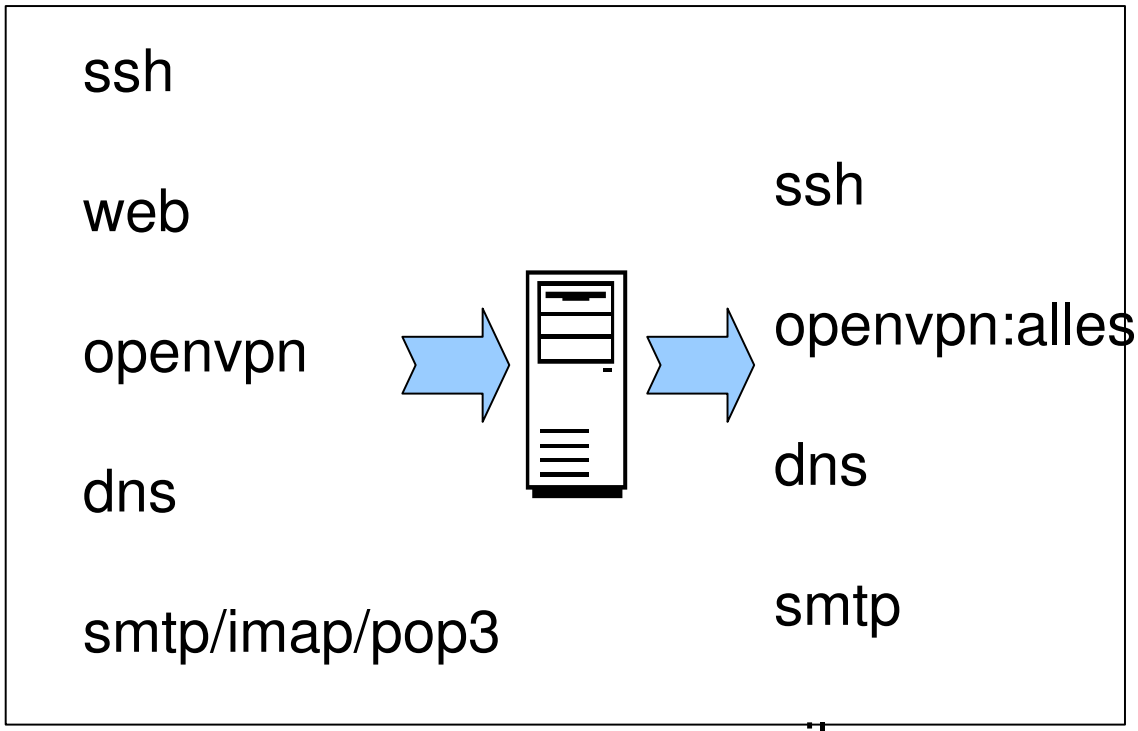

Shorewall Logging

- via ULOG
- via Syslog facility *kern*
  - abgestuft nach Level/Priority

```
destination df_shorewall_info {
  file ("/var/log/shorewall/info.log"
        owner(root) group(adm) );
};
filter f_shorewall_info {
  level (info) and match ("Shorewall");
};
log{
  source (s_all);
  filter (f_shorewall_info);
  destination (df_shorewall_info);
};
```

Shorewall Light

- Firewall Hosts
  - `/usr/share/shorewall-lite/shorecap > \ capabilities`
- Administrativer Host
  - neues Verzeichnis wie in /etc/shorewall
  - `shorewall -e <neues verzeichnis> \ /path/to/fw-script`
  - `shorewall load firewall` (auto via scp/ssh)
- sonst:
  - Makefiles

Netfilter

Development

iptables

- iptables, ip6tables, arptables, ebtables ->viel gemeinsamer Code
- Zusammenlegung ->Ziel pkttables
- langsam bei großer Anzahl Regeln -> nfHipac
  - existiert, ist aber noch kein offizieller Patch

Userspace Kommunikation

- sei 2.6.14 Netlink Device
- Abstraktion über Bibliotheken
- Subsysteme
  - conntrack, ulog, queue, helper, count
- neue ABI für Regelübergabe
  - XML, proc, Filesystem, Netlink

Logging mit ULOG

ULOG - Kernelspace

- Logging mit syslog – schlechte automatische Auswertung, unflexibel
- ULOG – neues Logziel für iptables
- benutzt Netlink Device – Unterscheidung anhand Netlink Gruppe
- "Burst Modus" – mehrere Lognachrichten zusammen versenden

ULOG - Userspace

- ulogd – Versionen 1.24 und 2.00-alpha
- specter – Version 1.4, Fork von ulogd 1.02
- Plugin basiert – hohe Flexibilität
- unterschiedliche Behandlung der Netlink Gruppen

Zentrale Logauswertung

Zentrale Logdatenauswertung

- Logdaten in zentraler Datenbank
- Auswertetools über diese Datenbank
  - Filter (fertige/eigne)
  - Korrelation
- Reports
- Alarmierung
- Sicherheit (Verschlüsselung, Signierung)
- Caching wenn Server nicht erreichbar

Prelude

http://www.prelude-ids.org

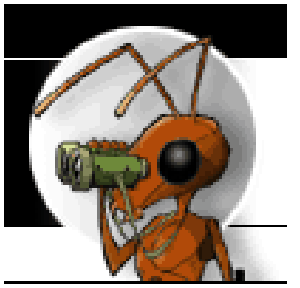

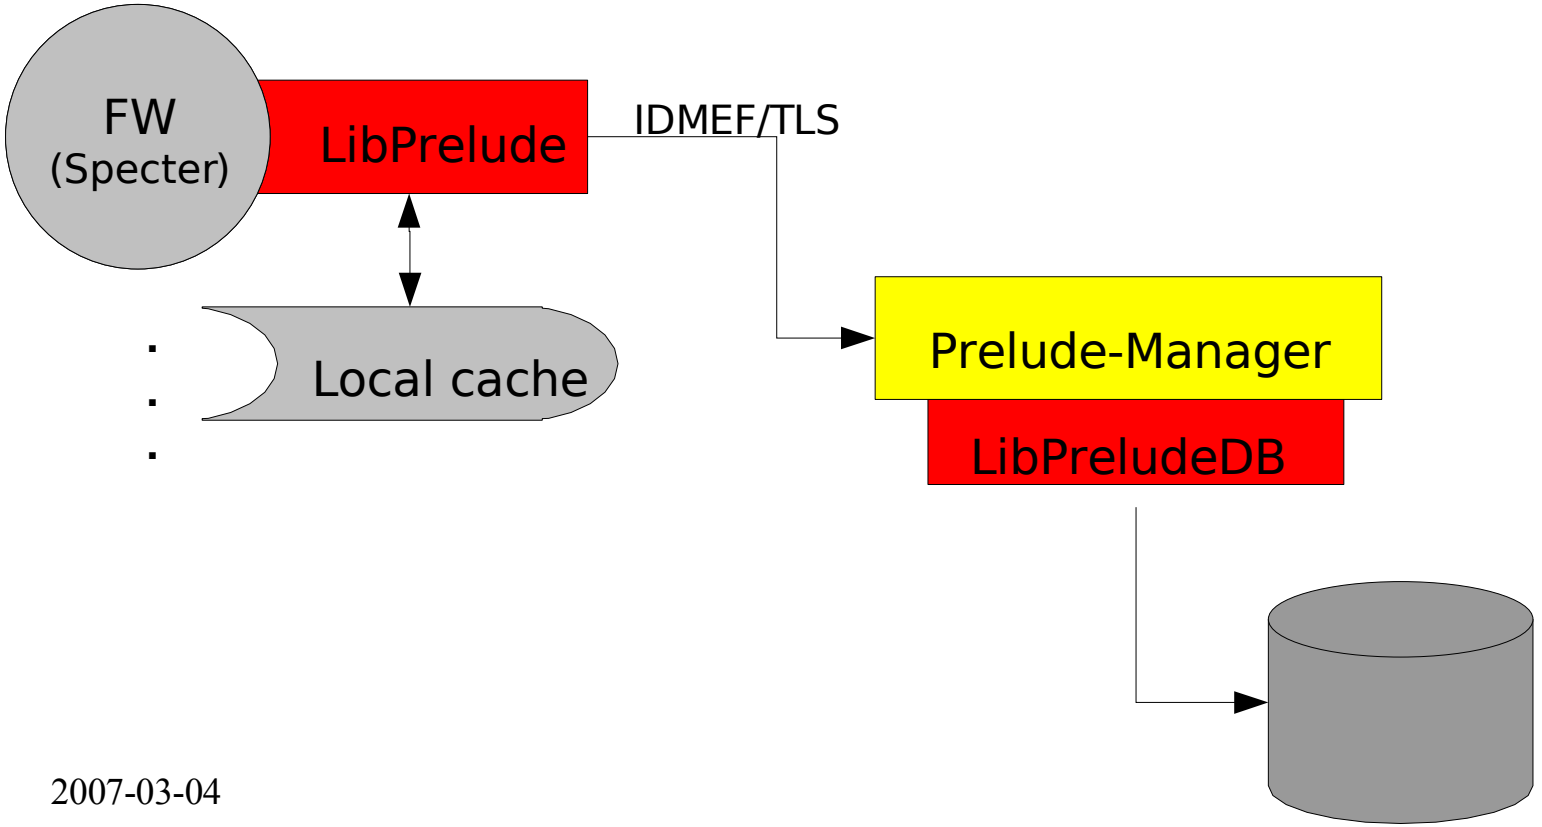

Prelude

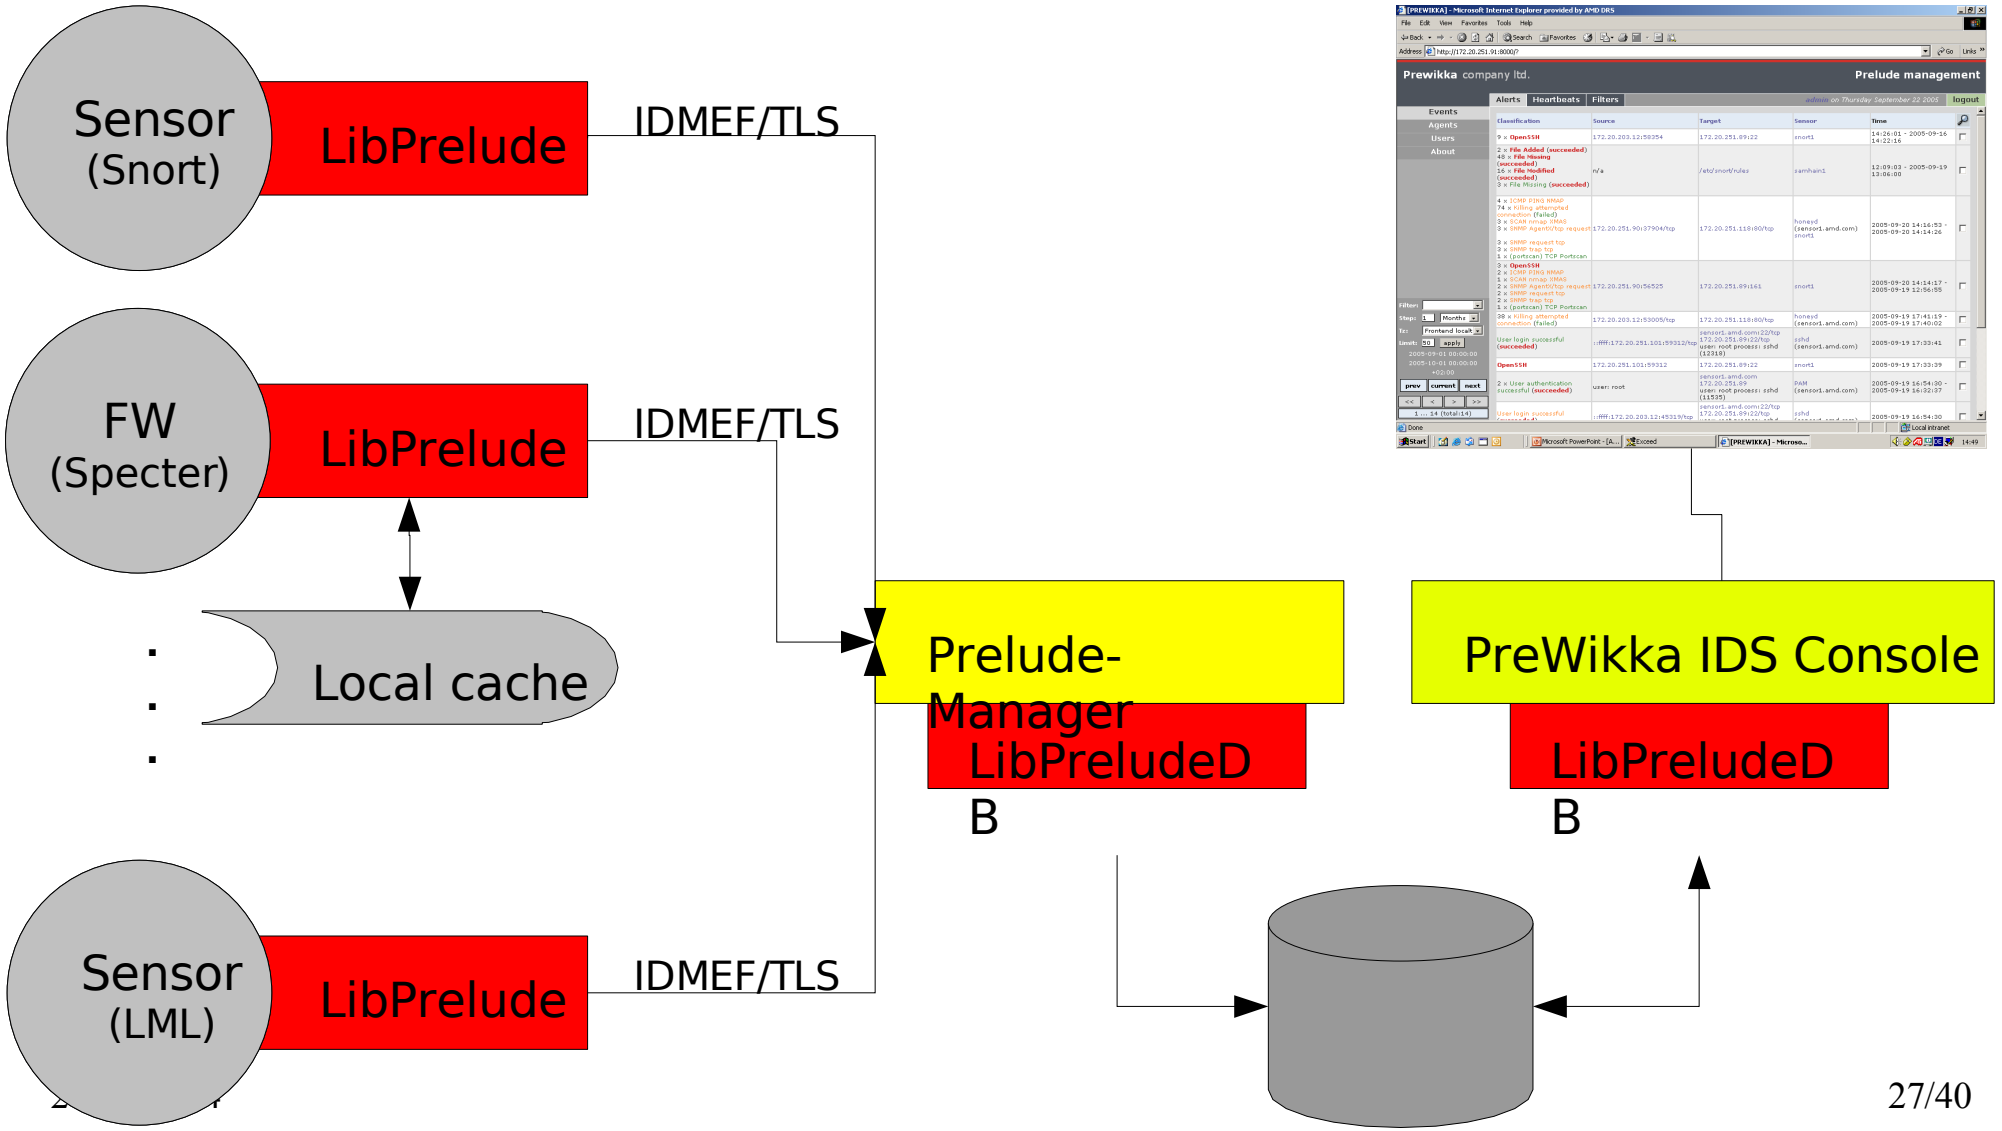
Sensor
(Snort)
LibPrelude
IDMEF/TLS
FW
(Specter)
LibPrelude
IDMEF/TLS
Local cache
Sensor
(LML)
LibPrelude
IDMEF/TLS
Prelude-
Manager
LibPreludeD
B
PreWikka IDS Console
LibPreludeD
B

Prewikka

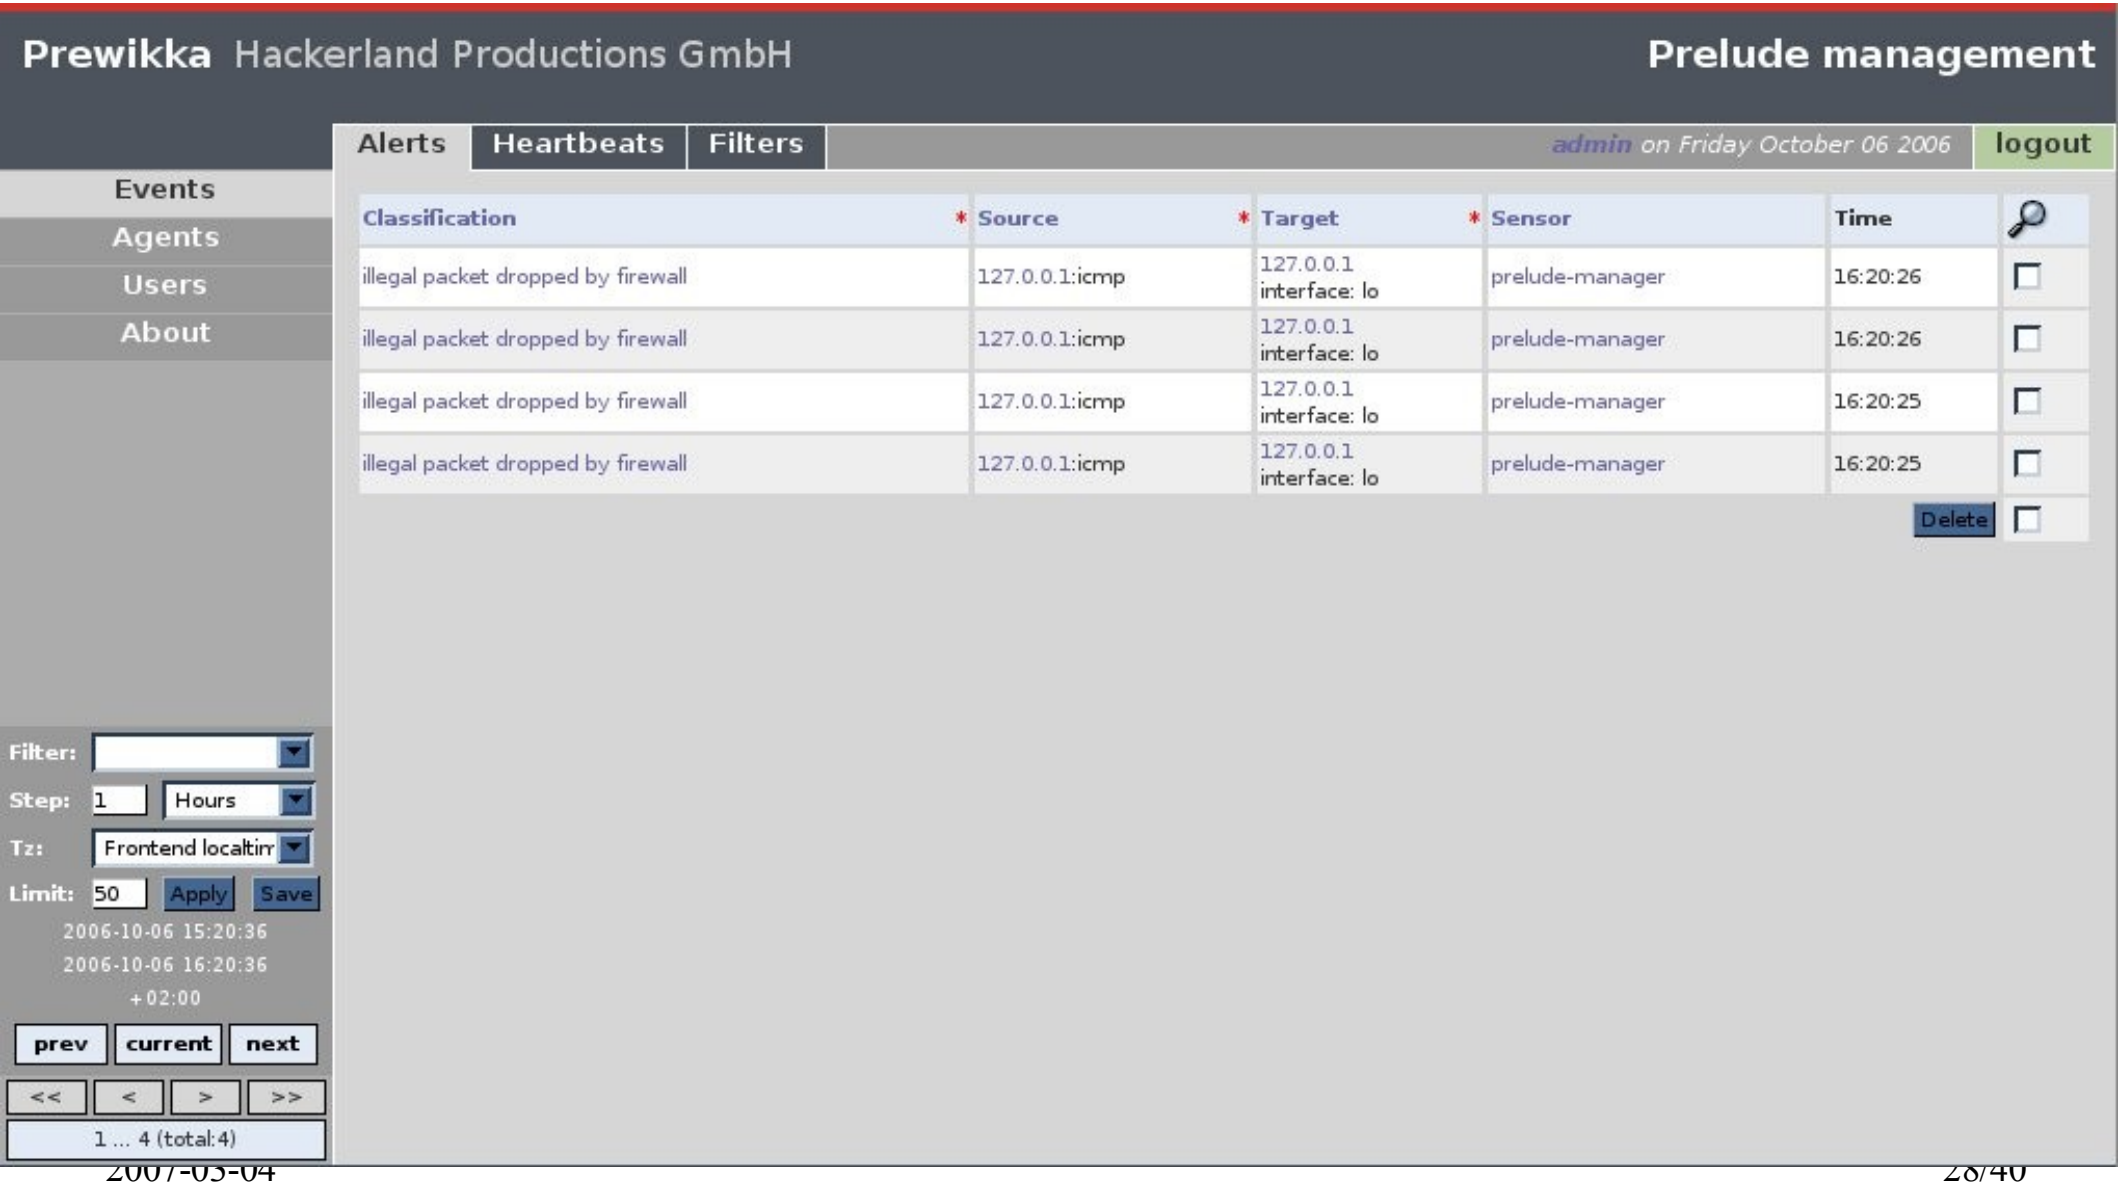
Prewikka Hackerland Productions GmbH
Prelude management
Alerts
Heartbeats
Filters
admin on Friday October 06 2006
logout
Events
Agents
Users
About
Classification
Source
Target
Sensor
Time
illegal packet dropped by firewall
127.0.0.1:icmp
127.0.0.1
interface: lo
prelude-manager
16:20:26
16:20:26
16:20:25
16:20:25
Delete
Filter:
Step:
1
Hours
Tz:
Frontend localtim
Limit:
50
Apply
Save
2006-10-06 15:20:36
2006-10-06 16:20:36
+02:00
prev
current
next
<<
<
>
>>
1 ... 4 (total:4)

Hochverfügbarkeit

Hochverfügbarkeit: Warum?

- Service funktioniert für die meisten Nutzer.
- Availability (A) definiert als: A=(actual life time / end of life time)
- High Availability: A>= 99.9 % (8.5 h downtime pro Jahr)
- Geplante Downtime zählt auch!
- Zuverlässigkeit:
  - Seriell verbundene Systeme: $R(t) = R_{S1}(t) * R_{S2}(t)$,

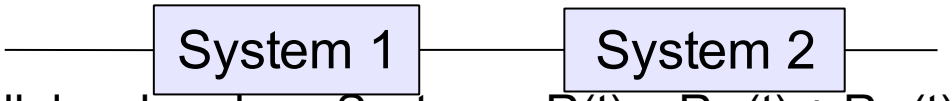

Parallel verbundene Systeme: $R(t) = R_{S1}(t) + R_{S2}(t) - R_{S1}(t) * R_{S2}(t)$

Beispiel
0.95
0.9025
0.9975

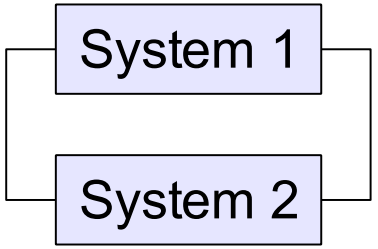

-> Schaffung redundanter Systeme!

Hochverfügbare Firewalls

- Hot/Standby
  - ein System aktiv, zweites scharf aber inaktiv
- Load-Sharing
  - Traffic wird zwischen mehreren Systemen aufgeteilt
- virtuelle IP/MAC Adresse
- Abgleich der Zustandstabellen

heartbeat

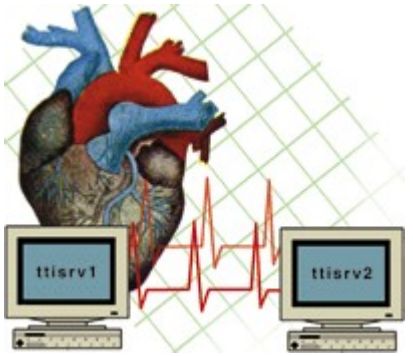

- Failover bei Node- oder Service Ausfall
- <=16 node Clusters
- heartbeat über serial, UDP bcast, mcast, ucast
- Active/Passive oder Active/Active (mehrere IPs)
- eingebautes Resource Monitoring
- STONITH
- mehr: Vortrag auf LinuxKongress 2006

heartbeat Demo

- ha.cf: generelle Konfigs
- haressorces: Virtuelle Dienste
- authkeys: Cluster-Authentifizierung
- Services über Shellscripts in ressource.d/

ucarp

- Entwickelt von OpenBSD
- Hot-Standby Cluster
- Antwort auf Cisco patentbelastetes VRRP

```
ucarp --interface=eth0 --srcip=10.1.1.1 --vhid=1 \
                      --pass=mypassword \
                     --addr=10.1.1.5 \
                     --upscript=/etc/vip-up.sh \
                      --downscript=/etc/vip-down.sh
```

```
14:44:25.154618 arp who-has 10.1.1.5 (ff:ff:ff:ff:ff:ff) tell 255.255.255.255
14:44:26.157500 arp reply 10.1.1.5 is-at 52:54:00:12:c3:d2
14:44:26.158614 IP 10.1.1.2 > 224.0.0.18: VRRPv2, Advertisement, vrid 1, prio 0, \
                authtype none, intvl 1s, length 36
14:44:26.158798 arp who-has 10.1.1.5 (ff:ff:ff:ff:ff:ff) tell 255.255.255.255
```

State Synchronisation

ct_sync

- Connection tracking seit kernel 2.4
- 2002 erste Anstrengungen für State Sync -> ct_sync
- funktionierender Patch bis 2.6.10
- Verwendet NACK für Synchronisation
- Active-Active im Development Tree
- Nachteile: keine Unterstützung für NAT und Conntrack Helper Module
- "pretty much dead"

conntrackd

- Userspace Implementierung mit netlink Bibliotheken
- erste Veröffentlichung Mai 2006, ready to use (experimental)
- Protokoll – Alarm ohne Rückmeldung, Multicast
- interner und externer Cache
- Active/Active (nur) mit symmetrischem Routing

conntrackd

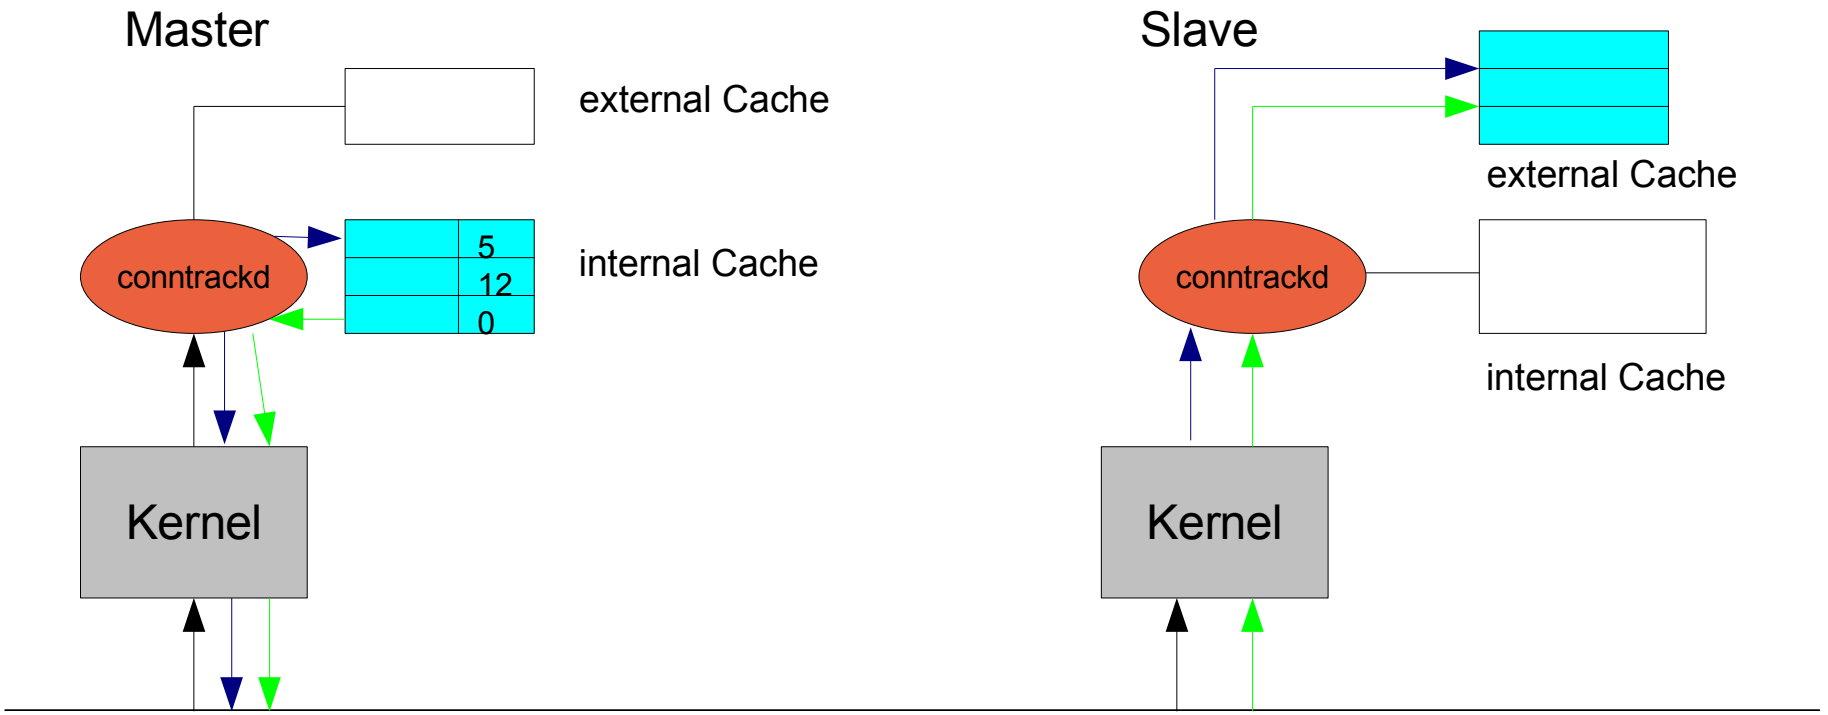
Master
external Cache
conntrackd
5
12
0
internal Cache
Kernel
Slave
external Cache
conntrackd
internal Cache
Kernel

Messungen

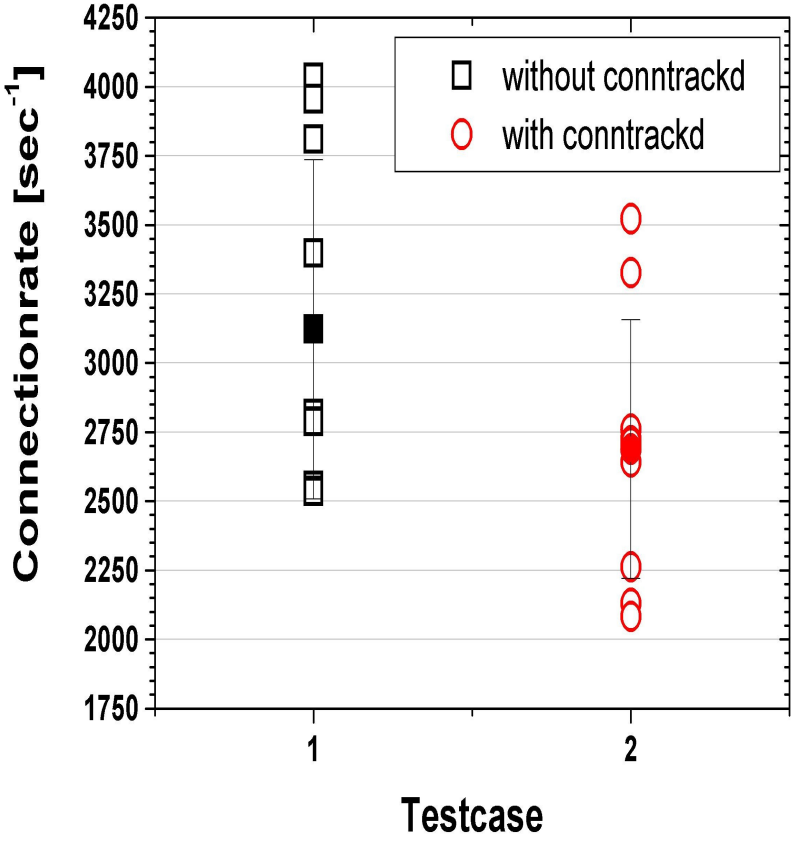
Connectionrate [sec^{-1}]
4250
4000
3750
3500
3250
3000
2750
2500
2250
2000
1750
without conntrackd
with conntrackd
1
2
Testcase

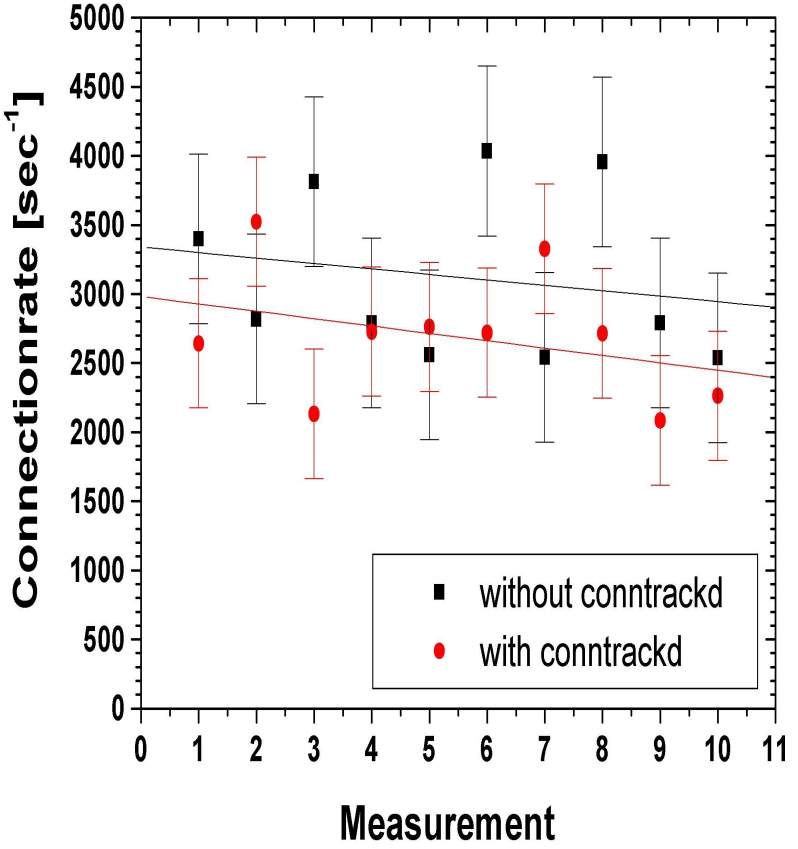
Connectionrate [sec^{-1}]
5000
4500
4000
3500
3000
2500
2000
1500
1000
500
0
without conntrackd
with conntrackd
0
1
2
3
4
5
6
7
8
9
10
11
Measurement

NACK

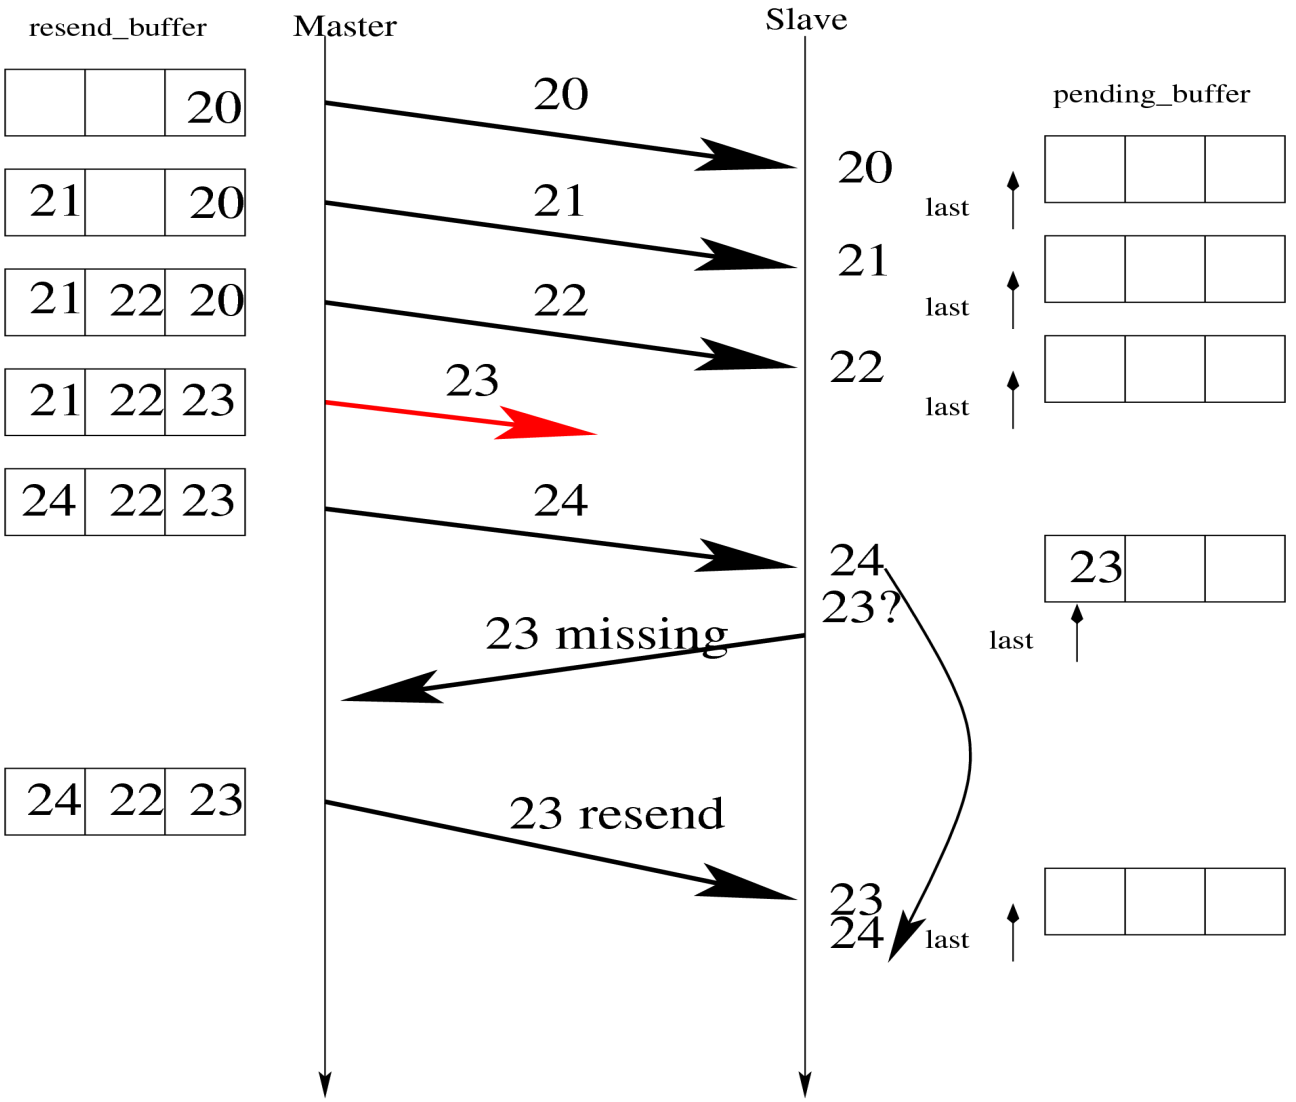
resend_buffer
Master
Slave
pending_buffer
20
21
22
23
24
23 missing
23 resend
23?
last

Happy Firewalling!

Demonstration

Versuchsaufbau