Linux Esotico

Alberto Cammozzo
mmzz@stat.unipd.it

Mauro Luzi
mauro@pluto.it

Serate a tema PLUTO Padova
17 Dicembre 2003

Linux Esotico: Sommario

- V-server
- Filtering bridge o bridging firewall
- Hyperscsi

V-server

- Un solo server hardware
- Un solo kernel
- Diversi server indipendenti:
  - Processi
  - Networking TCP/ IP e Sys V IPC
  - Filesystem
  - Distribution
  - Servizi
  - Utenti (incluso root)

Vserver: a cosa serve?

- **Hosting**: basta un server per diversi clienti
- **Sperimentazione**: coesistenza diverse release
- **Didattica**: un server per studente
- **Security box**: isolamento servizi
- **Migrazione soffice**: da vecchio a nuovo senza reboot
- **Rightsizing**: compro solo i server che mi servono.

Vserver: uno schema

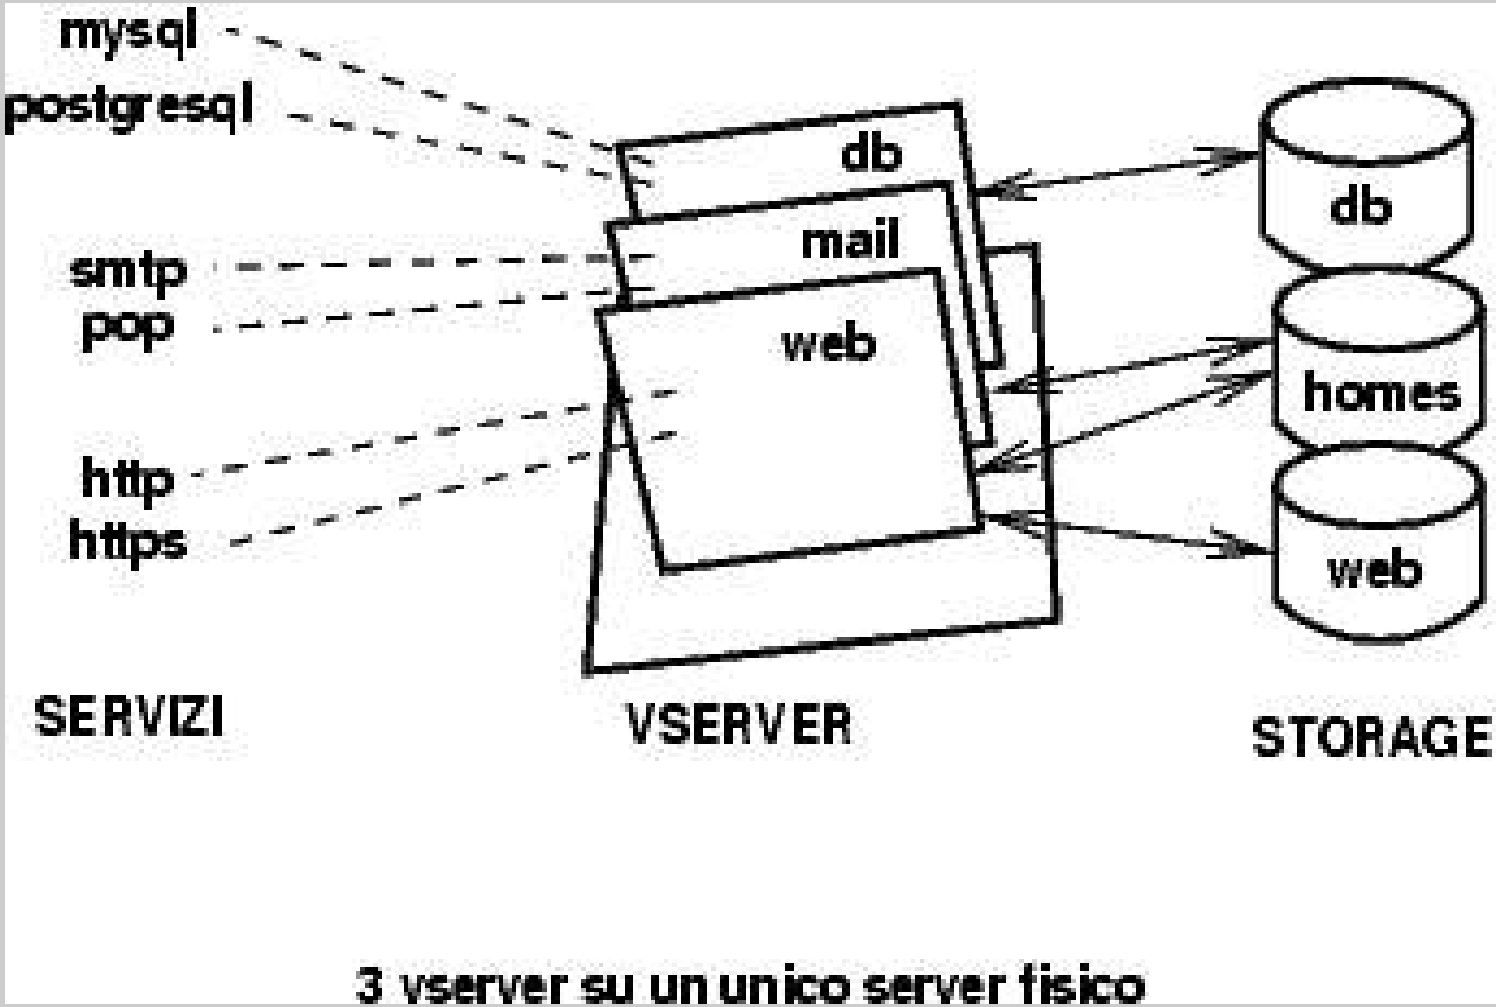

3 vserver su un unico server fisico

Vserver in pratica

```
mother :~# uname -a
Linux mother  2.4.21-ctx17 #5 SMP Fri Aug 22 08:53:25 CEST 2003 i686 unknown
mother :~# vserver-stat
CTX#   PROC QTY       VSZ        RSS  userTIME   sysTIME    UPTIME NAME
0            37      41MB        2kB   1h57m01  49m05.79  33d02h06 root server
2            10     428MB        4kB   1h10m37   4m51.83  32d08h07 zarquon
5            16     457MB        2kB  34m06.95   1m30.06  32d08h05 db
10           15     853MB        7kB  24m03.05  10m43.46  31d08h35 wonko
11           15     200MB        1kB    m55.50   1m30.75  30d04h32 agda
mother :~# vserver zarquon enter
ipv4root is now 147.162.35.4
New security context is 2
root@zarquon:/# ps ax
  PID TTY      STAT   TIME COMMAND
    1 ?        S      1:21 init
 1886 ?        S      0:03 /sbin/syslogd
 1898 ?        S      0:00 /usr/sbin/inetd
 1912 ?        S      0:02 /usr/sbin/cron
 1917 ?        S      0:45 /usr/sbin/apache-ssl
21386 ?        S      0:00 /usr/lib/apache-ssl/gcache 33 /var/run/gcache_port
21441 ?        S      0:00 /usr/sbin/apache-ssl
21442 ?        S      0:00 /usr/sbin/apache-ssl
21443 ?        S      0:00 /usr/sbin/apache-ssl
21444 ?        S      0:00 /usr/sbin/apache-ssl
21445 ?        S      0:00 /usr/sbin/apache-ssl
 9740 pts/1    S      0:00 /bin/bash -login
 9753 pts/1    R      0:00 ps ax
root@zarquon:/# logout
mother :~#
```

Vserver in pratica / 2

```
mother :~# vserver
vserver [ options ] server-name command ...

server-name is a directory in /vservers

The commands are:
 build   : Create a virtual server by copying the packages
           of the root server
 enter   : Enter in the virtual server context and starts a shell
           Same as "vserver name exec /bin/sh"
 exec    : Exec a command in the virtual server context
 suexec  : Exec a command in the virtual server context uid
 service : Control a service inside a vserver
           vserver name service service-name start/stop/restart/status
 start   : Starts the various services in the vserver, runlevel 3
 stop    : Ends all services and kills the remaining processes
 running : Tells if a virtual server is running
           It returns proper exit code, so you can use it as a test
 status  : Tells some information about a vserver
 chkconfig : It turns a server on or off in a vserver

--nodev  : Do not configure the IP aliases of the vserver
           Useful to enter a vserver without enabling its network
           and avoiding conflicts with another copy of this vserver
           running elsewhere
--silent : No informative messages about vserver context and IP numbers
           Useful when you want to redirect the output

mother :~# vserver wonko enter
ipv4root is now 147.162.35.81
New security context is 10
root@wonko:/# df
Filesystem           1K-blocks      Used Available Use% Mounted on
```

Vserver: come funziona / 1

- Nuove syscall:
  - **`new_s_context (int ctx)`**
  - Ctx: ID contesto, namespace univoco per i processi di quel contesto e i loro figli.
  - **`set_ipv4root(unsigned long ip)`**
  - Indirizzo IP immutabile per un dato contesto.
- Fa uso delle **capabilities** per limitare i poteri di un contesto: `/usr/include/linux/capability.h`

Vserver: come funziona / 2

- 3 nuovi comandi base:
  - `/usr/sbin/chcontext`
  - `/usr/sbin/chbind`
  - `/usr/sbin/reducecap`
- Comandi modificati:
  - `vps, vpstree, vrpm, vkill, vdu, vserver, vhalt, vtop, vservers, vserver-stat, ...`
- *Unification*:
  - Per non replicare n volte TUTTI i file

Vserver Webografia

- http:/ / www.linux-vserver.org/
- http:/ / www.paul.sladen.org/ vserver/ howto.html
- ftp:/ / ftp.solucorp.qc.ca/ pub/ vserver
- http:/ / www.solucorp.qc.ca/ miscprj/ s_context.hc
- http:/ / list.linux-vserver.org/ mailman/ listinfo/ vserver
- http:/ / www.13thfloor.at/ vserver

Filtering bridge

noto anche come

- transparent firewall
- in-line firewall
- shadow firewall
- stealth firewall
- bridging firewall
- “quella roba la”

FB: Bridge e router

- Bridge: inoltra pacchetti di livello 2 (eg ethernet)
  - Tra diversi spezzoni di rete L2
  - Instradamento semplice
  - Per evitare loop: *spanning tree*
- Router: inoltra pacchetti a livello 3
  - Tra sottoreti diverse
  - Deve conoscere la topologia della rete:
  - instradamento complesso (o statico)

FB: il firewall tradizionale

- Un firewall di solito e' un router
  - Smista tra piu' interfacce su [sub]net L3
  - Richiede un'interfaccia configurata su ogni rete.
- Accetta/ respinge pacchetti in base a determinate caratteristiche
  - Lavora prevalentemente a livello 3 e superiore (port number o anche altro)

FB: Il filtering Bridge

- Fa il bridginig, e non il routing, tra due interfacce
- Non richiede indirizzo L3 (IP)
  - Utilizzabile per filtrare porzioni della stessa net
  - Non richiede riconfigurazione router
  - Invisibile sulla rete: *Stealth*
- Permette ridondanza usando *spanning tree* (802.1)
  - Cosa difficile da ottenere con un router (a meno di usare protocolli di routing ... e fidarsi)
- Applica le stesse regole di firewall

FB: un confronto

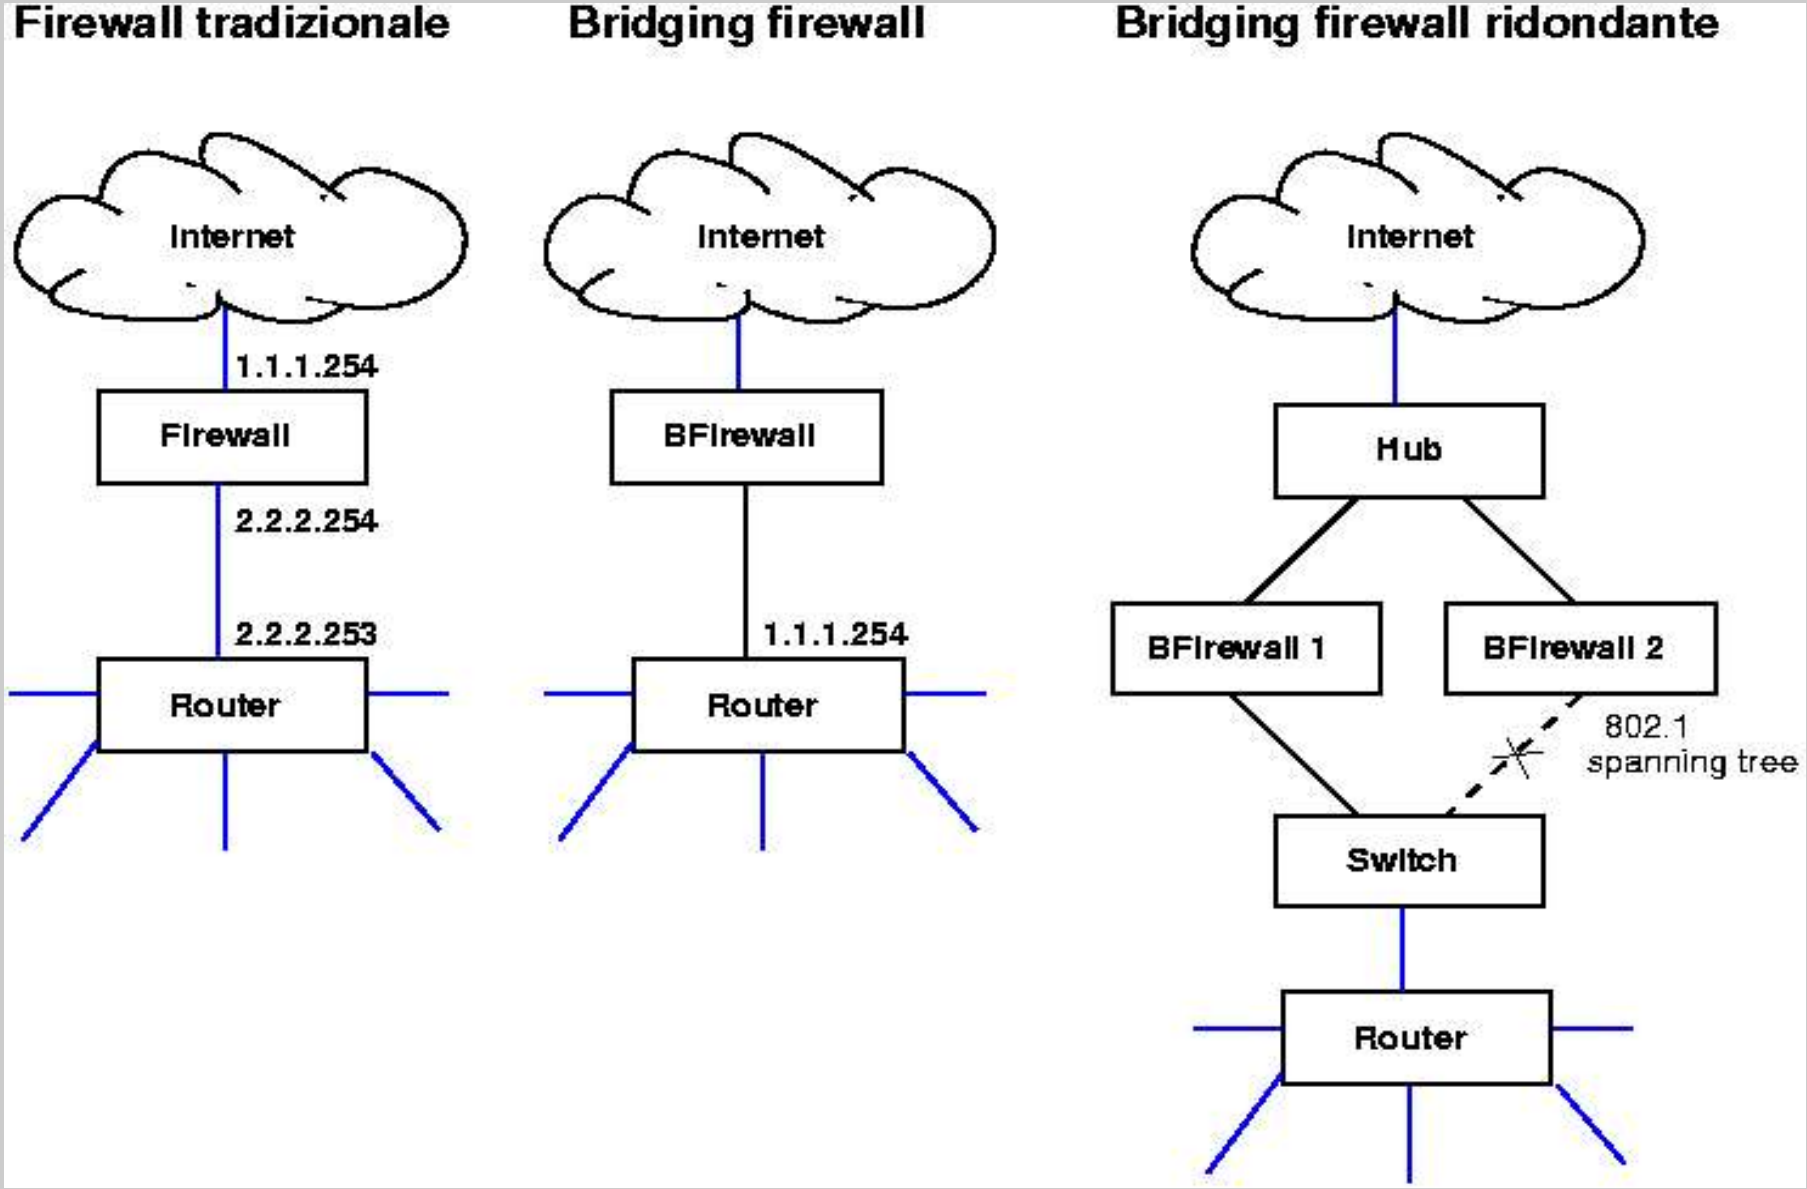

FB: altro uso furbo

Come tenere sulla stessa [sub] net rete buoni e cattivi

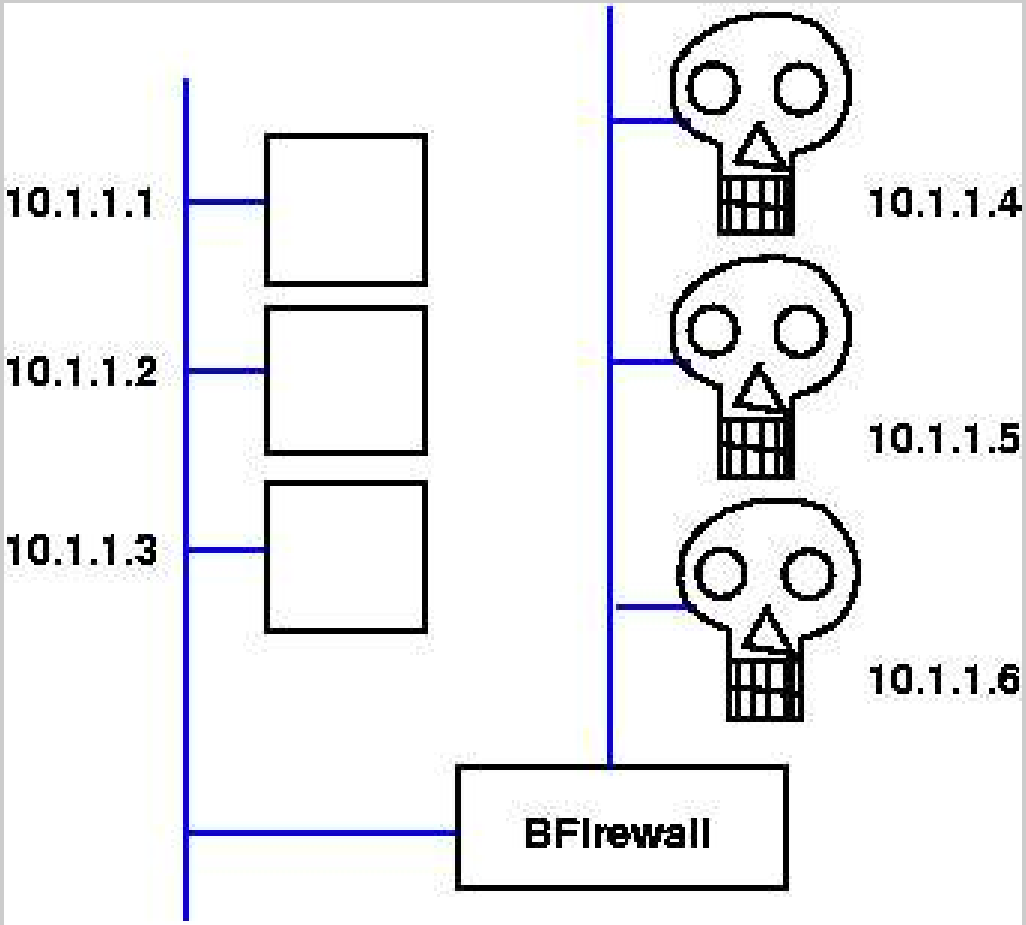

Filtering Bridge: Webografia

- http://www.tldp.org/HOWTO/Bridge+Firewall.html
- http://ezine.daemonnews.org/200211/ipfilter-bridge.html
- http://www.freebsd.org/doc/en_US.ISO8859-1/articles/filtering-bridges/
- http://www.linuxjournal.com/article.php?sid=4478

Hyper SCSI

HyperSCSI: il problema

Vari modi per montare i dischi:

- DAS: Direct attached storage: SCSI & co.
- NAS: network attached storage:

  nuovo nome per NFS, SMB & co.
- SAN: Storage area network

  Fiber Channel: caro impestato!

  iFCP: Internet Fiber channel protocol

  iSCSI: internet SCSI

HS: HyperSCSI

- Data Storage Institute, Singapore.
- Protocollo SCSI standard
  - Su frame ethernet (fast o giga)
  - Su frame IP
- Hs-server esporta un device / dev/ sd*
  - Anche usb, ide, ...
- hs-client importa un device / dev/ sd*
  - Puo' montarlo come gli pare, fare mkfs, fsck, software RAID...
  - Compaiono in / proc/ scsi

HS in pratica

- Linux (free), Windows ($$), Solaris
- Due moduli (server, client)
- Quattro programmi:
  - / sbin/ hs-configure
  - / sbin/ hs-wrapper
  - / sbin/ hs-server
  - / sbin/ hs-client

HyperSCSI Webografia

- http://nst.dsi.a-star.edu.sg/mcsa/hyperscsi/docs.html
- http://nst.dsi.a-star.edu.sg/mcsa/hyperscsi

Fine