



allegro



Technologie Solaris w Grupie Allegro



Piotr Jasiukajtis



Technologie Solaris w Grupie Allegro

- Rozwiązania oparte o systemy (Open)Solaris
- Systemy backupowe
- Wysoka dostępność w oparciu o Solaris Cluster
- Przyszłe technologie Solaris 11



Dlaczego (Open)Solaris?

- Wysoce skalowalna technologia
- Open Source / otwarte standardy
- Innowacyjne rozwiązania



Dlaczego (Open)Solaris?

- ZFS
- DTrace
- FMA
- SMF
- Solaris Zones
- Solaris Cluster / Open HA Cluster



System plików - ZFS

- Bezpieczeństwo
- Pojemność
- Swoboda w administracji
- Snapshoty



Wydajność ZFS

- ARC Cache
- L2 ARC - ZFS Intent Log ZIL - SSD
- ZFS prefetch
- Własna obsługa I/O



DTrace

- Dynamiczna instrumentacja całego systemu
- Kernel level
- User level
- Performance



FMA (Fault Management)

```
# fmstat
module          ev_rcv  ev_acpt  wait   svc_t  %w  %b  open  solve  memsz  bufisz
cpumem-retire   0        0  0.0    0.1   0   0    0     0     0     0
disk-transport  0        0  0.0    1.7   0   0    0     0     32b   0
eft             0        0  0.0    0.2   0   0    0     0     1.3M  0
ext-event-transport  0        0  0.0    0.1   0   0    0     0     0     0
fabric-xlate    0        0  0.0    0.1   0   0    0     0     0     0
fmd-self-diagnosis  76       0  0.0    0.0   0   0    0     0     0     0
io-retire       0        0  0.0    0.1   0   0    0     0     0     0
sensor-transport  0        0  0.0    0.2   0   0    0     0     32b   0
ses-log-transport  0        0  0.0    0.2   0   0    0     0     40b   0
snmp-trapgen    0        0  0.0    0.1   0   0    0     0     0     0
sysevent-transport  0        0  0.0    10.9  0   0    0     0     0     0
syslog-msgs     0        0  0.0    0.1   0   0    0     0     0     0
zfs-diagnosis   2        0  0.0    0.6   0   0    0     0     0     0
zfs-retire      2        0  0.0    0.3   0   0    0     0     28b   0
```



FMA

```
# fmadm faulty
-----
TIME          EVENT-ID          MSG-ID          SEVERITY
-----
03:22:18
Host          :
Platform     : SUN-FIRE- SERVER   Chassis_id :
Product_sn   :

Fault class  : fault.sunos.eft.unexpected_telemetry 50%
               defect.sunos.eft.unexpected_telemetry 50%
Problem in   : dev:///pci@0,0
               faulted and taken out of service

Description  : The diagnosis engine encountered telemetry from the listed
               devices for which it was unable to perform a diagnosis - no valid
               path to component was found in ereport.io.pciex.rc.nfe-msg.
               Refer to http://sun.com/msg/SUNOS-8000-J0 for more information.

Response     : Error reports have been logged for examination by Sun.

Impact       : Automated diagnosis and response for these events will not occur.

Action       : Ensure that the latest Solaris Kernel and Predictive Self-Healing
               (PSH) patches are installed.
```



FMA

```
# fmdump -e | grep ereport.cpu.intel.quickpath.mem_ce | wc -l
9316
# fmdump -e | grep ereport.cpu.intel.quickpath.mem_ce | tail
2010 16:11:41 ereport.cpu.intel.quickpath.mem_ce
2010 16:12:01 ereport.cpu.intel.quickpath.mem_ce
2010 16:22:01 ereport.cpu.intel.quickpath.mem_ce
2010 16:22:51 ereport.cpu.intel.quickpath.mem_ce
2010 16:23:21 ereport.cpu.intel.quickpath.mem_ce
2010 16:24:01 ereport.cpu.intel.quickpath.mem_ce
2010 21:44:52 ereport.cpu.intel.quickpath.mem_ce
2010 21:56:12 ereport.cpu.intel.quickpath.mem_ce
2010 22:07:32 ereport.cpu.intel.quickpath.mem_ce
2010 22:18:52 ereport.cpu.intel.quickpath.mem_ce
```



FMA

```
# fmdump -eV | head -30
TIME                               CLASS
2009 10:17:57.394191760 ereport.cpu.intel.quickpath.mem_ce
nvlst version: 0
  class = ereport.cpu.intel.quickpath.mem_ce
  ena = 0xa265b84555700401
  detector = (embedded nvlst)
  nvlst version: 0
    version = 0x0
    scheme = hc
    hc-list = (array of embedded nvlsts)
    (start hc-list[0])
    nvlst version: 0
      hc-name = motherboard
      hc-id = 0
    (end hc-list[0])
    (start hc-list[1])
    nvlst version: 0
      hc-name = chip
      hc-id = 0
    (end hc-list[1])
    (start hc-list[2])
    nvlst version: 0
      hc-name = memory-controller
      hc-id = 0
    (end hc-list[2])
  (end detector)

  compound_errorname = MC_CH_reserved_ERR
  IA32_MCG_STATUS = 0x0
```



Solaris Zones

- Wirtualizacja wbudowana w system
- Przejrzysta konfiguracja oraz zarządzanie
- Sklonowanie ZONY = ok 5sek
- Bezpieczeństwo usług
- Konsolidacja systemów



Rozwiązania oparte o (Open)Solaris

- Farmy obrazkowe
- Storage dla backupow
- Klastry Solaris Cluster NFS



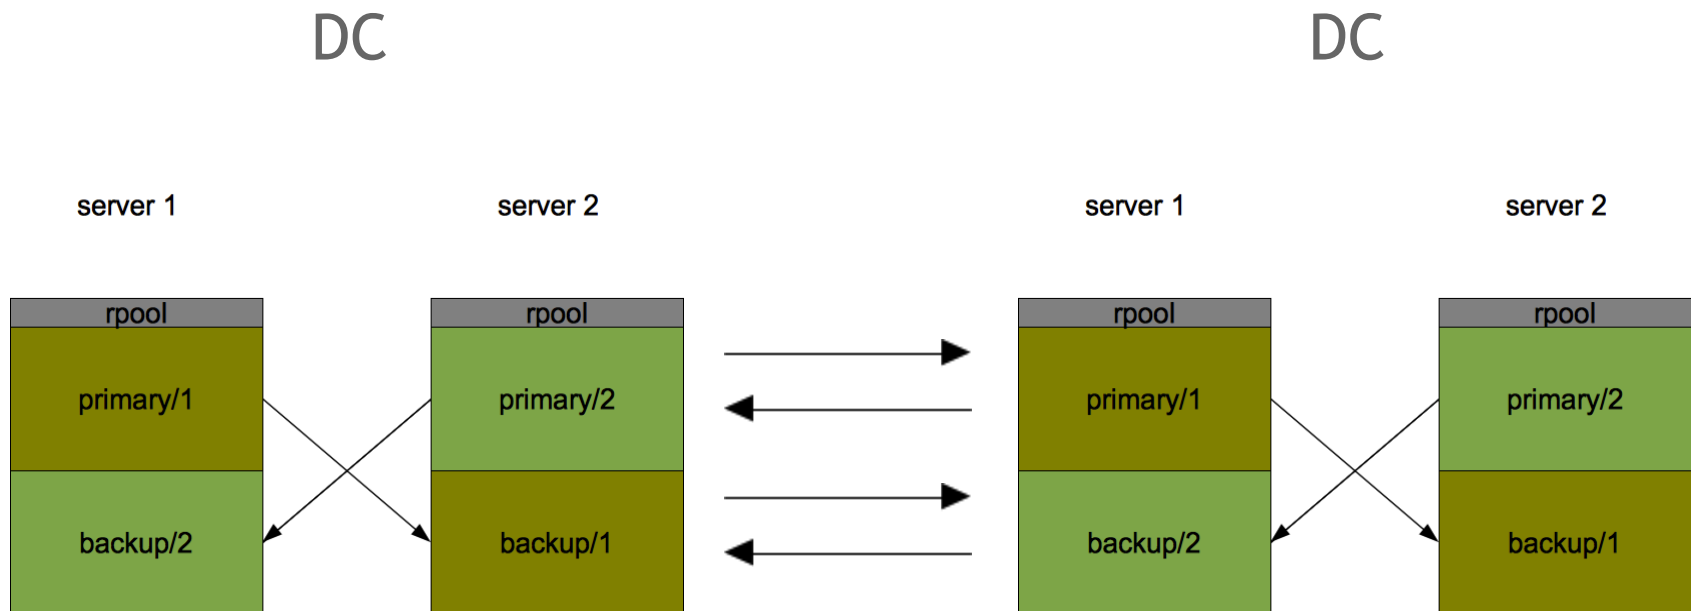
Farmy obrazkowe Allegro

422 240 838 plików

~ 10 TB miejsca




Architektura pary obrazkowej





Poznań - Frankfurt





Systemy backupowe

- Amanda Enterprise
- Solaris Zones
- ZFS snapshots
- rsync / rsyncd



Backup storage: Sun Fire X4540

```
# prtdiag | head -1
System Configuration: Sun Microsystems      Sun Fire X4540
# echo | format | tail
    /pci@3c,0/pci10de,376@f/pci1000,1000@0/sd@3,0
44. c12t4d0 <ATA-HITACHI HUA7210S-AC5A-931.51GB>
    /pci@3c,0/pci10de,376@f/pci1000,1000@0/sd@4,0
45. c12t5d0 <ATA-HITACHI HUA7210S-AC5A-931.51GB>
    /pci@3c,0/pci10de,376@f/pci1000,1000@0/sd@5,0
46. c12t6d0 <ATA-HITACHI HUA7210S-AC5A-931.51GB>
    /pci@3c,0/pci10de,376@f/pci1000,1000@0/sd@6,0
47. c12t7d0 <ATA-HITACHI HUA7210S-AC5A-931.51GB>
    /pci@3c,0/pci10de,376@f/pci1000,1000@0/sd@7,0
Specify disk (enter its number): Specify disk (enter its number):
```



Backup: zpool list

```
# zpool list
NAME      SIZE  ALLOC   FREE   CAP  DEDUP  HEALTH  ALTROOT
data1    20T   16.1T   3.88T   80%  1.00x  ONLINE  -
data2    20T   10.9T   9.12T   54%  1.00x  ONLINE  -
rpool    928G   255G   673G   27%  1.00x  ONLINE  -
```



zpool status

```
# zpool status data1
pool: data1
state: ONLINE
scrub: none requested
config:

    NAME                STATE                READ WRITE CKSUM
    data1
      raidz1-0           ONLINE              0     0     0
        c7t1d0           ONLINE              0     0     0
        c8t1d0           ONLINE              0     0     0
        c9t0d0           ONLINE              0     0     0
        c10t0d0          ONLINE              0     0     0
        c11t0d0          ONLINE              0     0     0
        c12t0d0          ONLINE              0     0     0
        c7t2d0           ONLINE              0     0     0
        c8t2d0           ONLINE              0     0     0
        c9t1d0           ONLINE              0     0     0
        c10t1d0          ONLINE              0     0     0
        c11t1d0          ONLINE              0     0     0
      raidz1-1           ONLINE              0     0     0
        c12t1d0          ONLINE              0     0     0
        c7t3d0           ONLINE              0     0     0
        c8t3d0           ONLINE              0     0     0
        c9t2d0           ONLINE              0     0     0
        c10t2d0          ONLINE              0     0     0
        c11t2d0          ONLINE              0     0     0
        c12t2d0          ONLINE              0     0     0
        c7t4d0           ONLINE              0     0     0
        c8t4d0           ONLINE              0     0     0
        c9t3d0           ONLINE              0     0     0
        c10t3d0          ONLINE              0     0     0
    spares
      c11t7d0            AVAIL
      c12t7d0            AVAIL
```



Backup: Solaris Zones

```
# zoneadm list -v
ID NAME          STATUS   PATH                BRAND  IP
 0 global        running /                   ipkg   shared
 1 backup2      running /zones/backup2     ipkg   shared
 2 backup1      running /zones/backup1     ipkg   shared

# zonecfg -z backup1 info fs
fs:
  dir: /backup1
  special: /data1/backup1
  raw not specified
  type: lofs
  options: []

# zonecfg -z backup1 info capped-cpu
capped-cpu:
  [ncpus: 1.00]

# zonecfg -z backup1 info capped-memory
capped-memory:
  physical: 1G
  [swap: 2G]
```

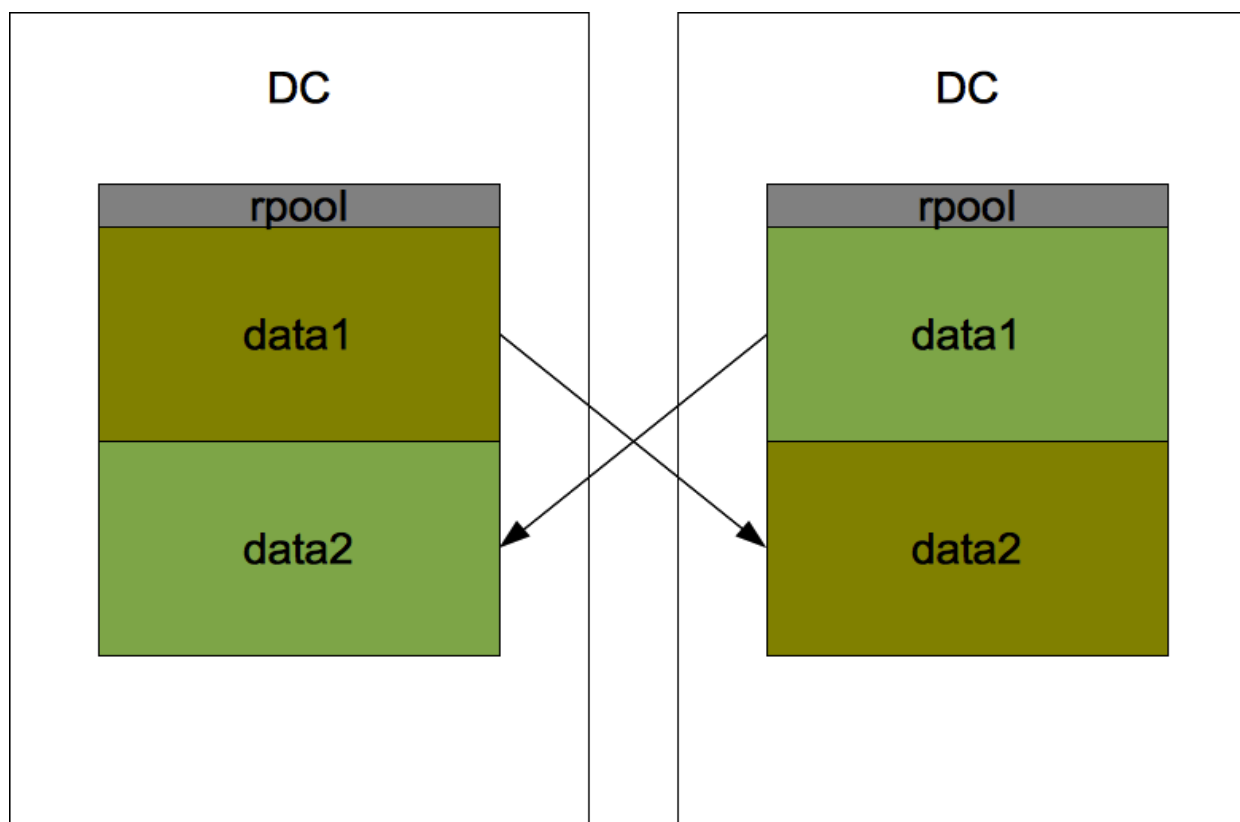


Backup: IPMP

```
# dladm show-link int0
LINK      CLASS      MTU      STATE      BRIDGE      OVER
int0      aggr       1500     up         --          nge0 nge1
# dladm show-link int1
LINK      CLASS      MTU      STATE      BRIDGE      OVER
int1      aggr       1500     up         --          nge2 nge3
# dladm show-aggr
LINK      POLICY     ADDRPOLICY          LACTIVITY    LACTIMER     FLAGS
int0      L4         auto               active       short        -----
int1      L4         auto               active       short        -----
# ipmpstat -i
INTERFACE  ACTIVE  GROUP      FLAGS      LINK      PROBE      STATE
int1       yes    ipmp0     -----   up        disabled   ok
int0       yes    ipmp0     --mb---   up        disabled   ok
# ipmpstat -g
GROUP      GROUPNAME  STATE      FDT        INTERFACES
ipmp0     ipmp0     ok         --         int1 int0
#
```

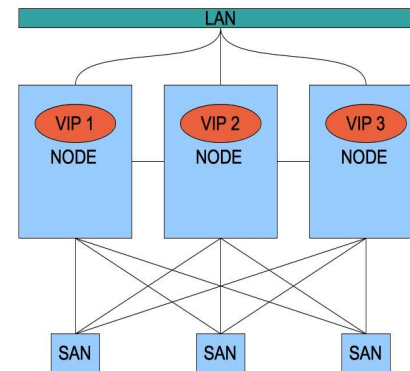



Backup: synchronizacja kopii





Wysoka dostępność Solaris Cluster






Open HA Cluster

- Solaris Cluster na platformie OpenSolaris
- Wysoka dostępność usług
- Nie wymaga shared storage
- Spójna konfiguracja



3 node NFS HA Clusters





3 node NFS HA Clusters

```
# clnode status -v

=== Cluster Nodes ===

--- Node Status ---

Node Name                               Status
-----                               -
nfs1-phys1                               Online
nfs1-phys2                               Online
nfs1-phys3                               Online

--- Node IPMP Group Status ---

Node Name      Group Name      Status      Adapter      Status
-----      -
nfs1-phys1     sc_ipmp0        Online      bnx0          Online
nfs1-phys2     sc_ipmp0        Online      bnx0          Online
nfs1-phys3     sc_ipmp0        Online      bnx0          Online
```



NFS Cluster Resource Groups

```
# clrg status -v

== Cluster Resource Groups ==

Group Name          Node Name          Suspended          Status
-----
nfs1-rg             nfs1-phys1         No                 Offline
                   nfs1-phys2         No                 Offline
                   nfs1-phys3         No                 Online

nfs1-rg2            nfs1-phys2         No                 Offline
                   nfs1-phys1         No                 Online
                   nfs1-phys3         No                 Offline

nfs1-rg3            nfs1-phys1         No                 Offline
                   nfs1-phys2         No                 Online
                   nfs1-phys3         No                 Offline
```



Kernel level Cluster

```
# modinfo | grep cl_
 2 ffffffffbb82000    8c8 - 1 cl_bootstrap (DCS bootstrap module)
 3 ffffffffbb827f8    4478 - 1 cl_runtime (cluster runtime support)
171 ffffffff7f54000   b3758 - 1 cl_comm (Sun Cluster communication suppo)
172 ffffffff7feb000    2ec0 - 1 cl_load (Sun Cluster cladmin support)
173 ffffffff7fee000    d6110 - 1 cl_orb (CORBA runtime support)
174 ffffffff808f000   1b1938 - 1 cl_haci (Sun Cluster Cluster Membership )
175 ffffffff81ef000    174c0 - 1 cl_quorum (Sun Cluster Quorum support)
276 ffffffff8b87000    141f8 - 1 cl_dlpitrans (ORB TCP Transport)
280 ffffffff8b9a000   1308e0 - 1 cl_dcs (Major-Minor server code)
```

```
# echo "*cmm_dbg_buf/s" | mdb -k | tail
th ffffff02e72451e0 tm 440084950: Quorum: preempt_nodes_and_reserve called
th ffffff02e72451e0 tm 440084951: SCSI3(/dev/did/rdisk/d2s2): quorum_preempt call
ed with NULL victim key, ignoring.
th ffffff02e72451e0 tm 440084952: take_ownership: device=1: owner_id=2, preempt_
nodes=0x0
th ffffff02e72451e0 tm 440084953: cmm_reconfig_done called bring_qds_online: fin
ished
th ffffff02e9dc4c20 tm 440278268: pc_pri.pc_clpmax = 60 pc_pri.pc_clpmin -60
th ffffff02e70cfae0 tm 440383013: check_ccr_quorum(): nodes=0x7, devices=0x1, vo
tes=5, total configured votes = 5; required quorum = 3
```

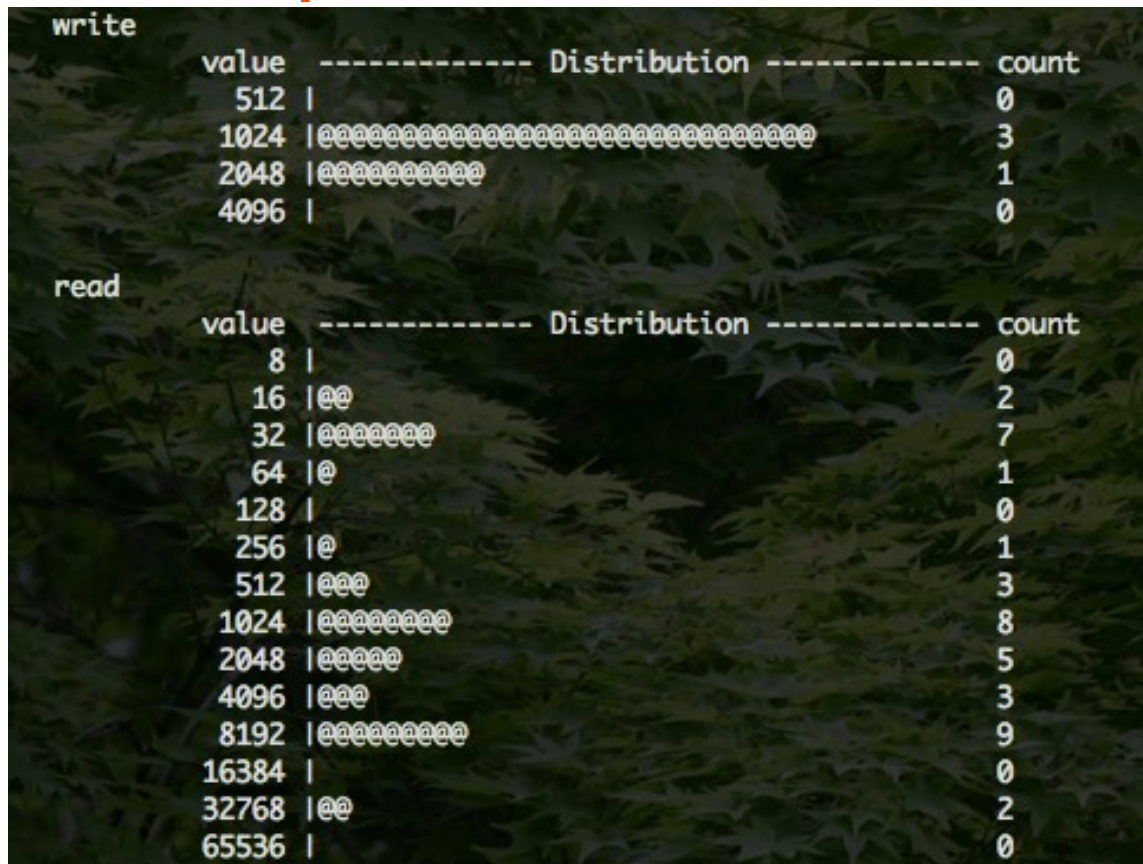


DTrace NFS provider: nfsv3rwsnoop.d

TIME(us)	CLIENT	OP	OFFSET	BYTES	PATHNAME
1751324627633	.0.28	W	0	963	/data/vol-staticpics/public/ap/209/ads/1250802837_thumb_46x46.jpg
1751324744998	.0.28	W	0	1079	/data/vol-staticpics/public/ap/209/ads/1249242228_thumb_46x46.jpg
1751324852561	.0.28	W	0	1154	/data/vol-staticpics/public/ap/209/ads/1255075665_thumb_46x46.jpg
1751324903998	.0.28	W	0	286	/data/vol-staticpics/public/ap/209/ads/131261.html
1751325091871	16.7	R	0	4096	/data/vol-istore03/istore/11366/photos/midi/778262.jpg
1751325021361	.0.28	W	0	1248	/data/vol-staticpics/public/ap/209/ads/1255234157_thumb_46x46.jpg
1751325134114	.0.28	W	0	1079	/data/vol-staticpics/public/ap/209/ads/1249242228_thumb_46x46.jpg
1751325106764	16.7	R	0	16384	/data/vol-istore02/istore/11366/photos/thumbnail/778265.jpg
1751325180791	.16.10	R	0	12288	/data/vol-staticpics/public/uploaded/209/logotip.jpg
1751325193795	16.6	R	0	16384	/data/vol-istore04/istore/11366/photos/thumbnail/778263.jpg
1751325104990	.16.7	R	4	28672	/data/vol-istore03/istore/11366/photos/midi/778262.jpg
1751325105983	.16.7	R	32	4096	/data/vol-istore03/istore/11366/photos/midi/778262.jpg
1751325247640	.0.28	W	0	986	/data/vol-staticpics/public/ap/209/ads/1274286470_thumb_46x46.jpg
1751325297007	.0.28	W	0	285	/data/vol-staticpics/public/ap/209/ads/129952.html
1751325419023	.0.28	W	0	1111	/data/vol-staticpics/public/ap/209/ads/1263001541_thumb_46x46.jpg
1751325532101	.0.28	W	0	1248	/data/vol-staticpics/public/ap/209/ads/1255234157_thumb_46x46.jpg



DTrace NFS provider: nfsv3rwtime.d





DTrace NFS provider: nfsv3rwtime.d cd

```
NFSv3 read/write by host (total us):
```

16.10	512
14.51	592
0.66	6285
3.8	7303
6.6	19784
16.9	21086
5.7	24110
14.52	46231

```
NFSv3 read/write top 10 files (total us):
```

/data/vol-co-kupic-www/cache/produkt/700/Canon-MV-700.html.new.1283242304.164634448	
/data/vol-co-kupic-www/cache/produkt/tke/Bonprix-Zakiet-w-kratke.html.new.1146874176.4	
/data/vol-staticpics/public/cz_loyalty/thumbnail_656	7294
/data/vol-staticpics/public/56/int_4455/1212342432_thumb_72x55.jpg	7303
/data/vol-staticpics/public/apnew/1/./56/428.gif	8494
/data/vol-staticpics/public/uploaded/209/mailings/komissioner-1.jpg	12828



DTrace NFS provider: nfsv3fileio.d

Rbytes	Wbytes	Pathname
0	95927	/data/vol-co-kupic-www/cache/produkt/8KM/Jaguar-S-Type-3-0-i-V
1473	0	/data/vol-staticpics/public/price_comparer/89/98/8998_1b4a4509
1533	0	/data/vol-staticpics/public/price_comparer/87/89/8789_6e932db2
1619	0	/data/vol-staticpics/public/price_comparer/87/89/8789_7b0afeb5
7096	0	/data/vol-istore02/istore/11273/photos/thumbnail/376197.jpg
7681	0	/data/vol-istore02/istore/11273/photos/thumbnail/376157.jpg
8027	0	/data/vol-istore02/istore/11273/photos/thumbnail/376073.jpg
8537	0	/data/vol-istore02/istore/11273/photos/thumbnail/376181.jpg
19860	0	/data/vol-istore03/istore/6686/photos/midi/1344614.jpg



Przyszłe technologie Solaris 11

- Crossbow
- IPS
- ipadm / dlstat / dladm ...
- ZFS crypto
- ZFS dedup
- Clearview: IPMP



Crossbow

```
# dladm show-vnic
LINK      OVER      SPEED  MACADDRESS  MACADDRTYPE  VID
ai1       bnx1      70     2:8:         random        0
dc3       bnx1      1000   2:8:         random        0
vnic2     bnx0      1000   2:8:         random        0
ds1       bnx1      1000   2:8:         random        0
clients1  bnx1      1000   2:8:         random        0
clients2  bnx1      1000   2:8:         random        0
#
```



dlstat

```
estibi@os-devel:~# dlstat -i 1
```

LINK	IPKTS	RBYTES	OPKTS	OBYTES
bnx0	202.51M	16.09G	2.04M	765.94M
bnx0	48	2.89K	1	210
bnx0	27	1.63K	1	162
bnx0	19	1.15K	1	162
bnx0	15	910	1	162
bnx0	13	790	1	162
bnx0	14	880	1	162
bnx0	45	4.25K	19	1.71K

^C



ipadm

```
estibi@os-devel:~# ipadm show-if
IFNAME      STATE      CURRENT      PERSISTENT
lo0         ok         -m-v-----46 ---
bnx0        ok         bm-----4-  ---
estibi@os-devel:~# ipadm show-addr
ADDROBJ      TYPE      STATE      ADDR
lo0/v4       static   ok         127.0.0.1/8
bnx0/_a      static   ok         .6/16
lo0/v6       static   ok         :::1/128
```



ipadm

```
estibi@os-devel:~# ipadm
usage: ipadm <subcommand> <args> ...
create-if      [-t] <interface>
disable-if    -t <interface>
enable-if     -t <interface>
delete-if     <interface>
show-if       [[-p] -o <field>,...] [<interface>]

set-ifprop    [-t] -p <prop>=<value[,...]> -m <protocol> <interface>
reset-ifprop  [-t] -p <prop> -m <protocol> <interface>
show-ifprop   [[-c] -o <field>,...] [-p <prop>,...]
              [-m <protocol>] [<interface>]

create-addr   [-t] {-T static <static_args> | -T dhcp <dhcp_args> |
               -T addrconf <addrconf_args>} <addrobj>
static_args = <[-d] -a {local|remote}=addr[/prefixlen]>
dhcp_args = <[-w <seconds> | forever]>
addrconf_args = <[-i interface-id]
               [-p {stateful|stateless}={yes|no}]>

down-addr    [-t] <addrobj>
up-addr      [-t] <addrobj>
disable-addr -t <addrobj>
enable-addr  -t <addrobj>
refresh-addr [-i] <addrobj>
delete-addr  [-r] <addrobj>
show-addr    [[-p] -o <field>,...] [<addrobj>]

set-addrprop [-t] -p <prop>=<value[,...]> <addrobj>
reset-addrprop [-t] -p <prop> <addrobj>
show-addrprop [[-c] -o <field>,...] [-p <prop>,...] <addrobj>

set-prop     [-t] -p <prop>[+|-]=<value[,...]> <protocol>
reset-prop   [-t] -p <prop> <protocol>
show-prop    [[-c] -o <field>,...] [-p <prop>,...] [<protocol>]
```



ipadm show-addrprop

```

estibi@os-devel:~# ipadm show-addrprop
ADDROBJ      PROPERTY    PERM  CURRENT          PERSISTENT    DEFAULT    POSSIBLE
lo0/v4       broadcast   r-    --              --            --         --
lo0/v4       deprecated  rw    off             --            off        on,off
lo0/v4       prefixlen   rw    8               --            8          1-30,32
lo0/v4       private     rw    off             --            off        on,off
lo0/v4       transmit    rw    on              --            on         on,off
lo0/v4       zone        rw    global          --            global     --
bnx0/_a      broadcast   r-    255             --            255       --
bnx0/_a      deprecated  rw    off             --            off        on,off
bnx0/_a      prefixlen   rw    16              --            8          1-30,32
bnx0/_a      private     rw    off             --            off        on,off
bnx0/_a      transmit    rw    on              --            on         on,off
bnx0/_a      zone        rw    global          --            global     --
lo0/v6       broadcast   r-    --              --            --         --
lo0/v6       deprecated  rw    off             --            off        on,off
lo0/v6       prefixlen   rw    128             --            64         1-126,128
lo0/v6       private     rw    off             --            off        on,off
lo0/v6       transmit    rw    on              --            on         on,off
lo0/v6       zone        rw    global          --            global     --
  
```



ipadm show-ifprop

```

estibi@os-devel:~# ipadm show-ifprop
IFNAME      PROPERTY      PROTO PERM  CURRENT  PERSISTENT  DEFAULT  POSSIBLE
lo0         arp           ipv4  rw   on       --          on       on,off
lo0         forwarding    ipv4  rw   off      --          off      on,off
lo0         metric        ipv4  rw   0        --          0        --
lo0         mtu           ipv4  rw   8232     --          8232     68-8232
lo0         exchange_routes ipv4  rw   on       --          on       on,off
lo0         usesrc        ipv4  rw   none     --          none     --
lo0         forwarding    ipv6  rw   off      --          off      on,off
lo0         metric        ipv6  rw   0        --          0        --
lo0         mtu           ipv6  rw   8252     --          8252     1280-8252
lo0         nud           ipv6  rw   on       --          on       on,off
lo0         exchange_routes ipv6  rw   on       --          on       on,off
lo0         usesrc        ipv6  rw   none     --          none     --
bnx0       arp           ipv4  rw   on       --          on       on,off
bnx0       forwarding    ipv4  rw   off      --          off      on,off
bnx0       metric        ipv4  rw   0        --          0        --
bnx0       mtu           ipv4  rw   1500     --          1500     68-1500
bnx0       exchange_routes ipv4  rw   on       --          on       on,off
bnx0       usesrc        ipv4  rw   none     --          none     --
bnx0       forwarding    ipv6  rw   --       --          off      on,off
bnx0       metric        ipv6  rw   --       --          0        --
bnx0       mtu           ipv6  rw   --       --          --       --
bnx0       nud           ipv6  rw   --       --          on       on,off
bnx0       exchange_routes ipv6  rw   --       --          on       on,off
bnx0       usesrc        ipv6  rw   --       --          none     --
  
```



ipadm show-prop

```

estibi@os-devel:~# ipadm show-prop
PROTO PROPERTY      PERM CURRENT  PERSISTENT  DEFAULT  POSSIBLE
ipv4 forwarding      rw  off   --         off      on,off
ipv4 ttl           rw  255  --         255     1-255
ipv6 forwarding    rw  off   --         off      on,off
ipv6 hoplimit      rw  255  --         255     1-255
ipv6 hostmodel     rw  weak  --         weak    strong,
src-priority,
weak
ipv4 hostmodel     rw  weak  --         weak    strong,
src-priority,
weak
icmp recv_maxbuf   rw  8192  --         8192    4096-65536
icmp send_maxbuf   rw  8192  --         8192    4096-65536
tcp  ecn           rw  passive --         passive never,passive,
active
tcp  extra_priv_ports rw  2049,4045 --         2049,4045 1-65535
tcp  largest_anon_port rw  65535  --         65535    1024-65535
tcp  recv_maxbuf     rw  128000 --         128000   2048-1073741824
tcp  sack           rw  active  --         active   never,passive,
active
tcp  send_maxbuf     rw  49152  --         49152    4096-1073741824
tcp  smallest_anon_port rw  32768  --         32768    1024-65535
tcp  smallest_nonpriv_port rw  1024  --         1024     1024-32768
udp  extra_priv_ports rw  2049,4045 --         2049,4045 1-65535
udp  largest_anon_port rw  65535  --         65535    1024-65535
udp  recv_maxbuf     rw  57344  --         57344    128-1073741824
udp  send_maxbuf     rw  57344  --         57344    1024-1073741824
udp  smallest_anon_port rw  32768  --         32768    1024-65535
udp  smallest_nonpriv_port rw  1024  --         1024     1024-32768
sctp extra_priv_ports rw  2049,4045 --         2049,4045 1-65535
sctp largest_anon_port rw  65535  --         65535    1024-65535
sctp recv_maxbuf     rw  102400 --         102400   8192-1073741824
sctp send_maxbuf     rw  102400 --         102400   8192-1073741824
sctp smallest_anon_port rw  32768  --         32768    1024-65535
sctp smallest_nonpriv_port rw  1024  --         1024     1024-32768

```



IPMP Rearchitecture

```
# ipmpstat -g
GROUP          GROUPNAME     STATE        FDT          INTERFACES
ipmp0          ipmp0         ok           10.00s      proxy1 proxy0
int0           int0          degraded    10.00s      igb1 [igb2]
daemons0      daemons0     ok           --          daemon1 daemon0
# ipmpstat -i
INTERFACE      ACTIVE        GROUP        FLAGS        LINK        PROBE        STATE
proxy1        yes          ipmp0        --mb---     up          ok           ok
proxy0        yes          ipmp0        -----     up          ok           ok
igb2          no           int0         -----     up          failed      failed
igb1          yes          int0         --mb---     up          ok           ok
daemon1       yes          daemons0    --mb---     up          disabled    ok
daemon0       yes          daemons0    -----     up          disabled    ok
```



IPMP: ipmpstat -p

```
# ipmpstat -p
TIME      INTERFACE  PROBE  NETRTT   RTT      RTTAVG   TARGET
0.13s    igb2       56268  0.19ms   0.24ms   0.29ms   108
1.09s    igb1       56241  0.27ms   0.38ms   0.31ms   110
1.12s    proxy1     56831  0.76ms   0.79ms   1.18ms   1
1.52s    proxy0     56800  0.82ms   0.93ms   1.28ms   1
1.78s    igb2       56269  0.28ms   0.32ms   0.31ms   109
2.49s    proxy0     56801  0.28ms   0.32ms   1.16ms   1
2.86s    igb1       56242  0.15ms   0.20ms   0.30ms   117
2.91s    proxy1     56832  0.47ms   0.55ms   1.10ms   1
3.11s    igb2       56270  0.20ms   0.29ms   0.27ms   118
3.96s    igb1       56243  0.11ms   0.21ms   0.28ms   103
4.28s    igb2       56271  0.23ms   0.39ms   0.31ms   119
4.44s    proxy1     56833  0.41ms   0.52ms   1.03ms   1
^C
```



Automated Installer

- Sieciowa instalacja systemów
- Następca Jumpstart
- Automatyczny deployment farm serwerów



AI: installadm

```
# installadm list
Service Name Status      Arch  Port  Image Path
-----
ai-srv1_b111 on      x86   46503 /export/ai/os-ai-b111
os-ai-b117  on      x86   46501 /export/ai/os-ai-b117/
os-ai-b123  on      x86   46502 /export/ai/os-ai-b123/
os-ai-b126  on      x86   46504 /export/ai/os-ai-b126/
# █
```



Pytania?



Piotr Jasiukajtis