Hálózatok építése és üzemeltetése

OSPF gyakorlat

Ismétlés

Routing protokollok

- Feladatuk
  - optimális útvonal (next hop) kiszámítása bármely csomópontok között
    - aktuális állapot információ gyűjtés a hálózatról
    - útvonalak kalkulálása
  - forwarding táblák
    - konfigurálása
    - dinamikus karbantartása, frissítése
    - bejövő routing protokoll üzenetek alapján
  - routing információk
    - feldolgozása
    - terjesztése

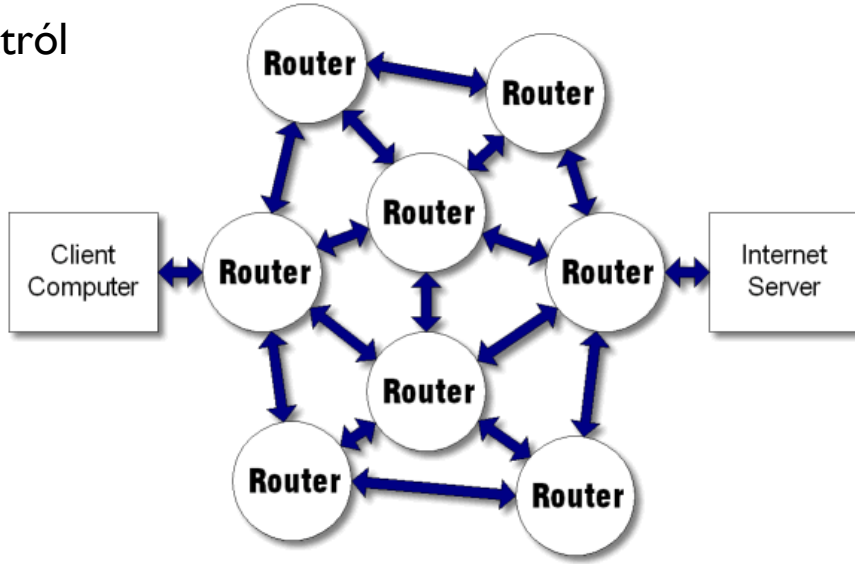

Routing protokollok

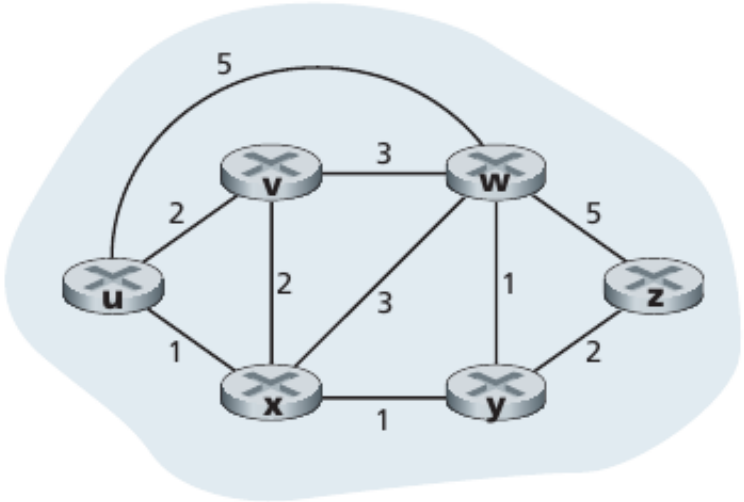

- Hálózat: abszrakt gráf
  - csomópontok: routerek
  - élek: linkek
  - élköltség: valamilyen metrika (pl. késleltetés, sávszélesség kifejezése)
- cél:
  - (valamilyen értelemben) optimális, legkisebb költségű útvonal meghatározása két csomópont között
    - pl. legrövidebb út
- Ismerős algoritmusok:
  - Dijkstra algoritmus
  - Bellman-Ford algoritmus

J. Kurose, K. Ross, "Computer Networking: A Top-Down Approach", Pearson, 2012

Csoportosításuk

- Globális vs. Elosztott
  - globális: minden router ismeri a teljes topológiát
  - elosztott: minden router csak a szomszédjait és a tőlük kapott üzeneteket ismeri
- Intra-domain vs. Inter-domain
  - intra: Interior Gateway Protocol (IGP)
    - közös adminsztratív domain
    - rugalmatlan szabályok
    - egyes esetekben nem jól skálázódik
  - inter: Exterior Gateway Protocol (EGP)
    - külön adminisztratív domainek, AS-ek (Autonomous System) között
    - jól skálázódik (internet)
- Link state vs. Distance Vector (ld. később)

Csoportosításuk

- Interior Gateway Protocol (IGP) példák
  - OSPF (OpenShortest Path First)
  - IS-IS (Intermediate System to Intermediate System)
  - RIP (Routing Information Protocol)
  - EIGRP (Enhanced Interior Gateway Routing Protocol)
- Exterior Gateway Protocol (EGP)
  - BGP (Border Gateway Protocol)
  - Id. MSc (Internet architektúra és szolgáltatások főspecializáció)

Link State alapú routing

- Működési elv
  - globális nézeten dolgozik
  - LSP: Link State Packet (id, costs, seq.no, ttl)
  - egy router
    - mindenkinek küld LSP-t (broadcast)
    - a közvetlenül kapcsolódó linkjeiről
    - periodikusan újra generálja (seq.no++)
    - legfrissebb beérkezett LSP-ket tárolja
  - mindenki ugyanazt a topológiát látja
  - azon számolja az útvonalakat
  - útvonalszámítás: Dijkstra algoritmus

Például: OSPF

- Open Shortest Path First (v2)
- nyílt, IETF szabvány
  - v2: RFC 2328
  - IP felett
- együttműködés különböző gyártók termékei között!
- korlátozott erőforrás igény
- viszonylag gyors, automatikus konvergencia topológia változásokra
- támogatja
  - különböző útvonal költségek számítását
  - hierarchikus, többszintű topológiát
  - alkalmazás típusára alapozott forgalomirányítást
  - autentikációt minden üzenetre

Hálózatemulációs környezet

Netkit, Quagga

Università degli Studi Roma Tre
Dipartimento di Informatica e Automazione
Computer Networks Research Group

Netkit

The poor man's system for experimenting computer networking

Version	2.3
Author(s)	G. Di Battista, M. Patrignani, M. Pizzonia, M. Rimondini
E-mail	contact@netkit.org
Web	http://www.netkit.org/
Description	an introduction to the architecture, setup, and usage of Netkit

netkit: a system for emulating computer networks

- based on uml (user-mode linux)
  - http://user-mode-linux.sourceforge.net/
- each emulated network device is a virtual linux box
  - a virtual linux box is one that is based on the uml kernel
- note: the linux os is shipped with software supporting most of the network protocols
  - hence, any linux machine can be configured to act as a bridge/switch or as a router

Netkit

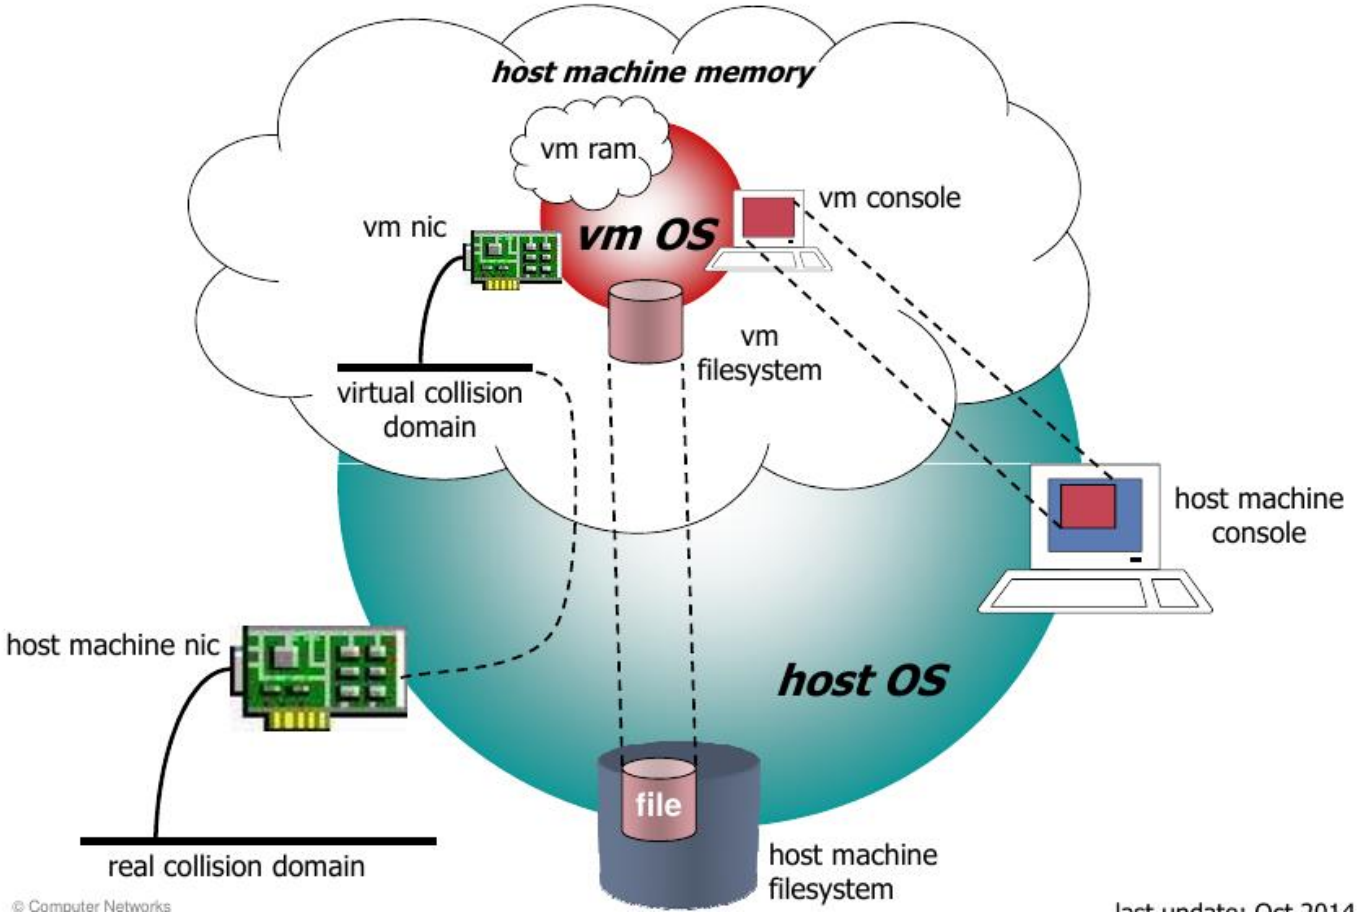
host machine memory
vm ram
vm console
vm nic
vm OS
vm
filesystem
virtual collision
domain
host machine
console
host machine nic
host OS
file
real collision domain
host machine
filesystem
© Computer Networks

Netkit

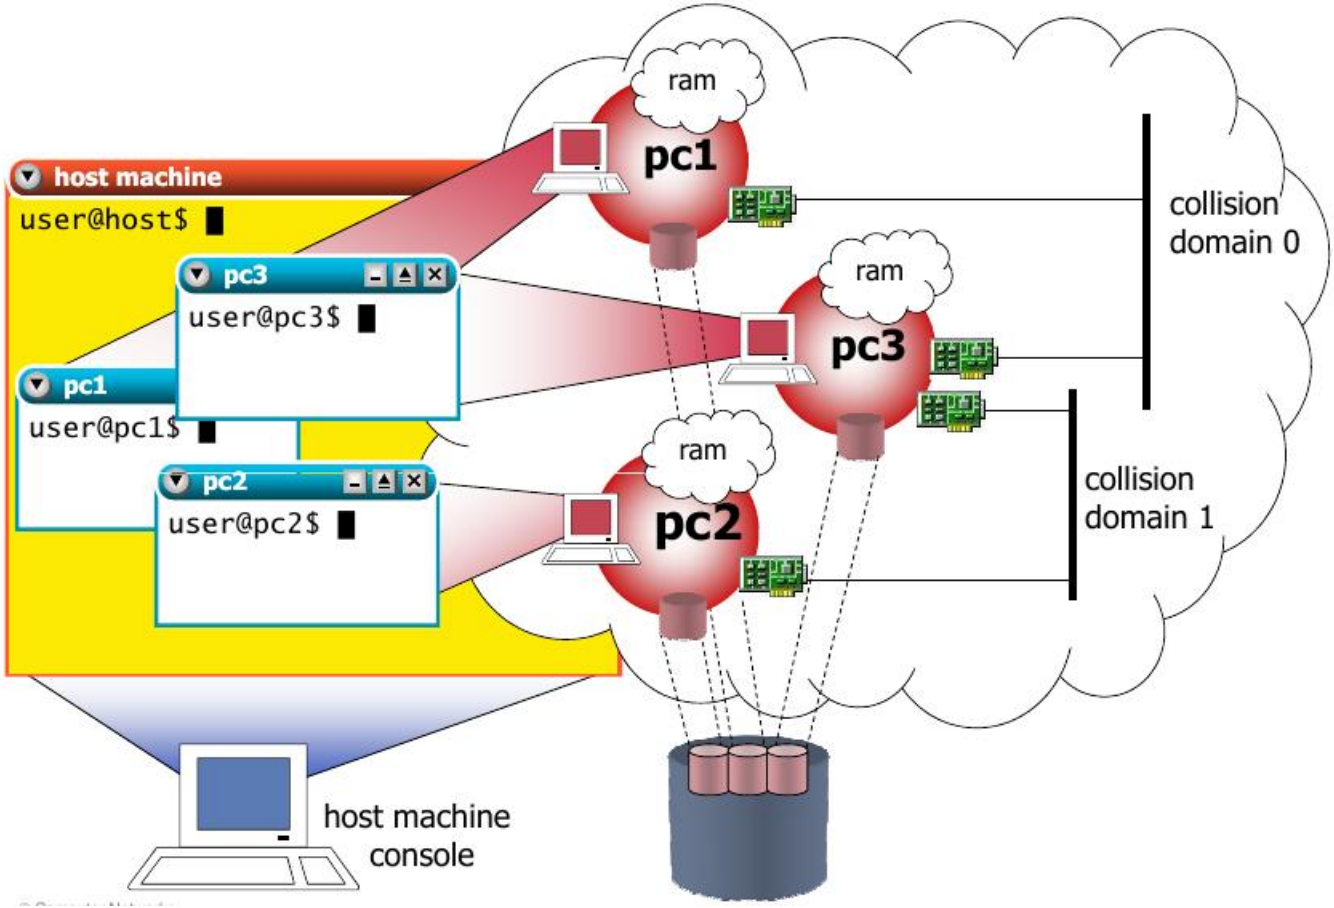
host machine
user@host$
pc3
user@pc3$
pc1
user@pc1$
pc2
user@pc2$
ram
pc1
ram
pc3
ram
pc2
collision
domain 0
collision
domain 1
host machine
console

netkit vcommands

- allow to startup virtual machines with arbitrary configurations (memory, network interfaces, etc.)
  - vstart: starts a new virtual machine
  - vlist: lists currently running virtual machines
  - vconfig: attaches network interfaces to running vms
  - vhalt: gracefully halts a virtual machine
  - vcrash: causes a virtual machine to crash
  - vclean: "panic command" to clean up all netkit processes (including vms) and configuration settings on the host machine

netkit lcommands

- ease setting up complex labs consisting of several virtual machines
  - lstart: starts a netkit lab
  - lhalt: gracefully halts all vms of a lab
  - lcrash: causes all the vms of a lab to crash
  - lclean: removes temporary files from a lab directory
  - linfo: provides information about a lab without starting it
  - ltest: allows to run tests to check that the lab is working properly

Quagga/Zebra

- Routing Szoftver csomag
  - GPL
  - FreeBSD, Linux, Solaris, NetBSD
  - GNU Zebra volt előbb
  - Quagga egy fork volt
    - "The Quagga tree aims to build a more involved community around Quagga than the current centralised model of GNU Zebra."
    - az élővilágban a quagga halt ki
    - a routing világban a zebra

Quagga/Zebra

- Routing Szoftver csomag
  - zebra (core daemon)
    - kernel interface, static routes
    - zserv szerver (API) -> quagga kliensek felé
  - quagga démonok
    - routing protokollok
      - ripd, ripngd, ospfd, ospf6d, bgpd, isisd
    - mindegyikkel dedikált CLI-n (vty) keresztül kommunikálhatunk
      - hasonló interfész, mint egy HW routernél
    - speciális quagga tool: vtysh
      - közös front-end minden démonhoz

zebra: a routing daemon

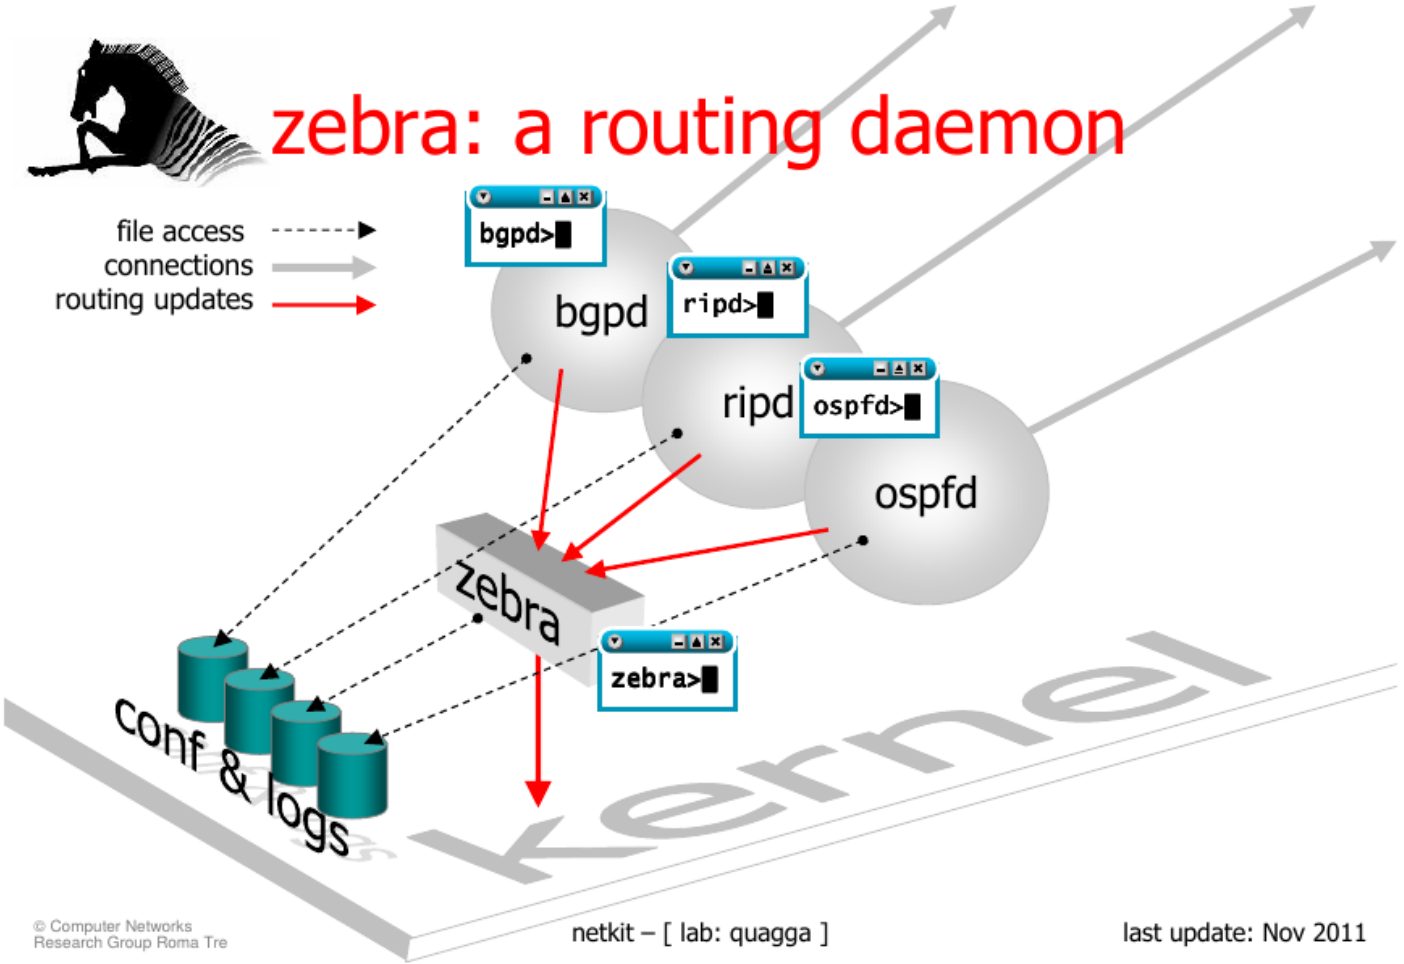

Vizsgálati környezet kialakítása

- IB213 labor: Mininet+Netkit boot image
- Saját gép:
  - Mininet+Netkit-disk1.vmdk letöltés
  - Indítás VMware player vagy importálás Virtualbox-ba
- A rendszer egy lebutított Ubuntu 64-bit linux, tartalma:
  - Mininet hálózatemuláció
  - Netkit hálózatemuláció
- OSPF lab indítása:
  - $ cd ~/netkit/labs/netkit-labs_advanced-topics/netkit-labs_ospf/netkit-lab_ospf-singlearea
  - $ lstart

OSPF lab#1

netkit-lab_ospf-singlearea

OSPF lab#1 topológia

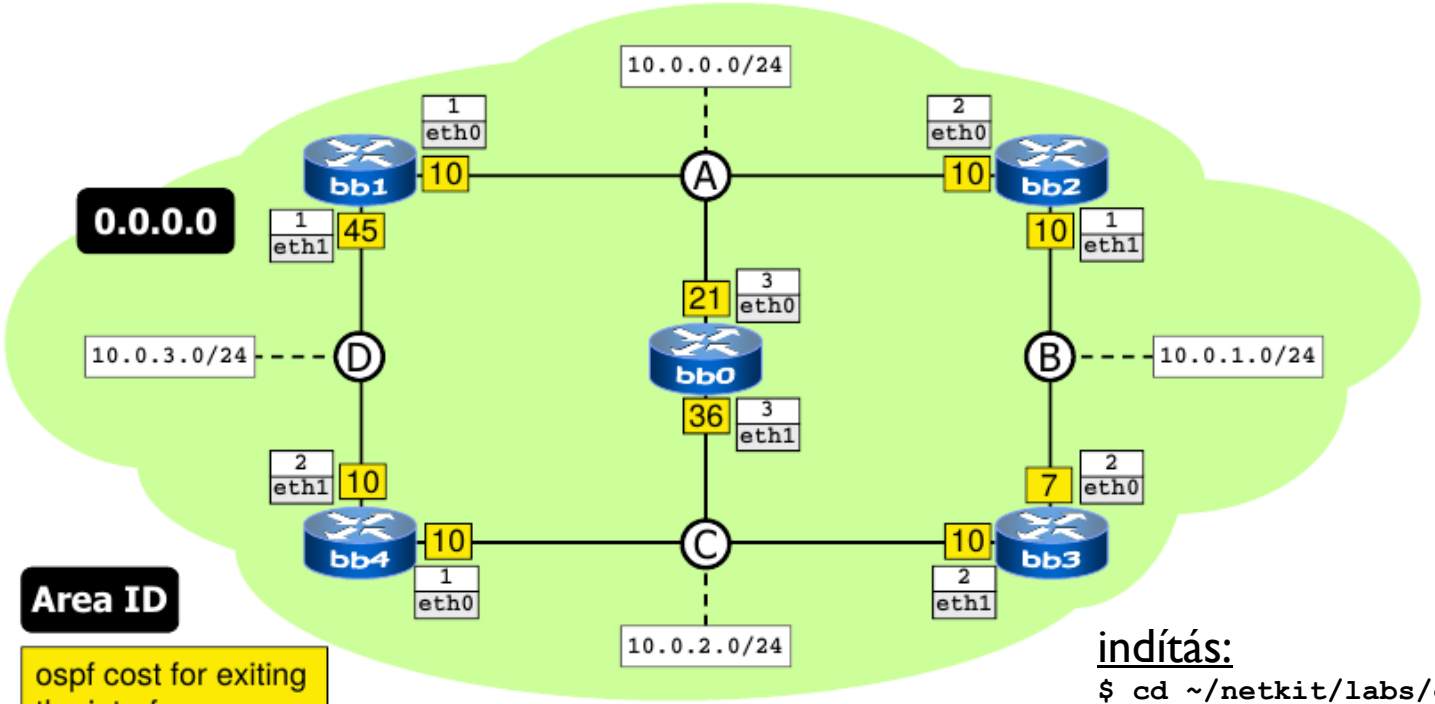

- single (backbone) area
  - 0.0.0.0
- minden interfészhez
  - ospf cost
  - default: 10
  - néha trükkösen van beállítva!

indítás:

```
$ cd ~/netkit/labs/ospf_gyak
$ lstart
```


netkit – [labs: ospf]

last update: Nov 2014

OSPF lab#1 topológia

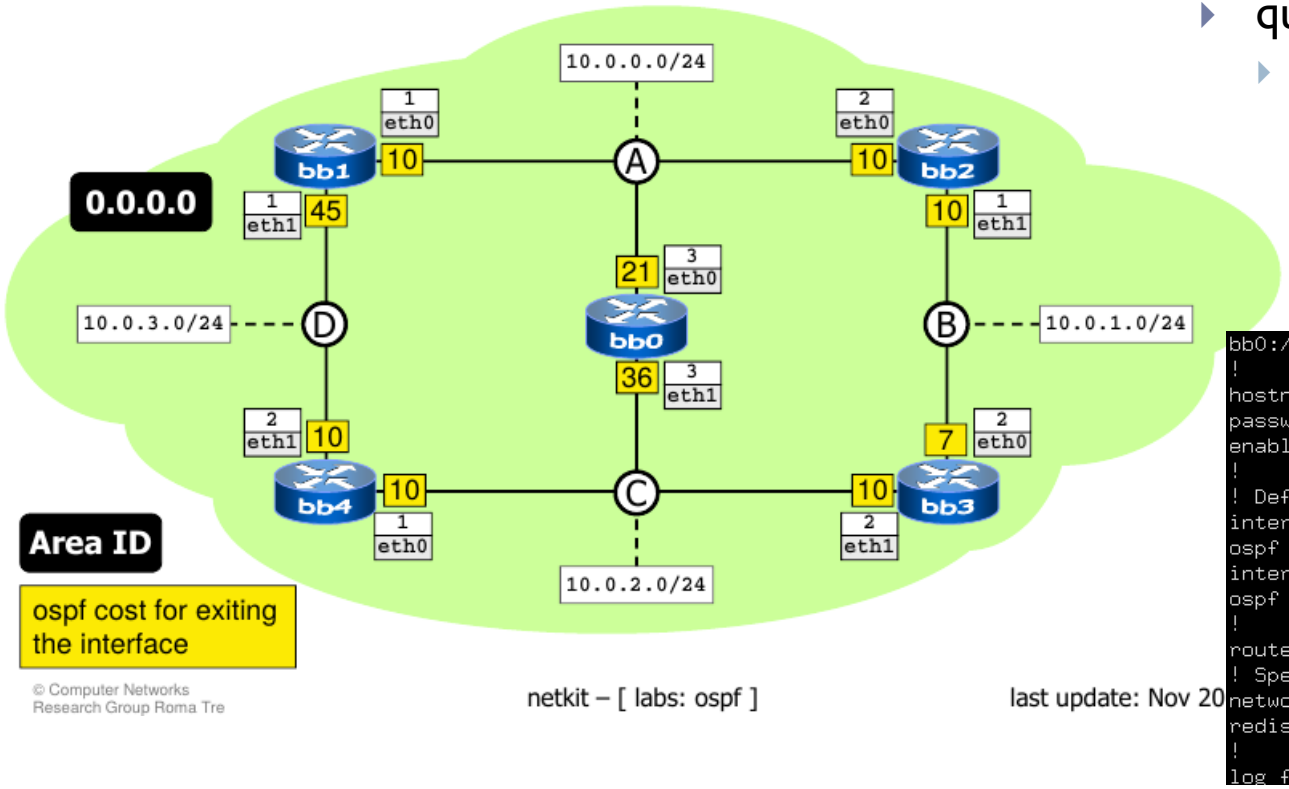

- quagga teszt, pl. bb0 routeren
  - cd /etc/zebra; ls -l
    - daemon conf fájlok
    - cat daemons
    - cat zebra.conf (passwd!)
    - cat ospfd.conf

```
bb0:/etc/zebra# cat ospfd.conf
!
hostname ospfd
password zebra
enable password zebra
!
! Default cost for exiting an interface is 10
interface eth0
ospf cost 21
interface eth1
ospf cost 36
!
router ospf
! Speak OSPF on all interfaces falling in 10.0.0.0/16
network 10.0.0.0/16 area 0.0.0.0
redistribute connected
!
log file /var/log/zebra/ospfd.log
```

OSPF lab#1 topológia

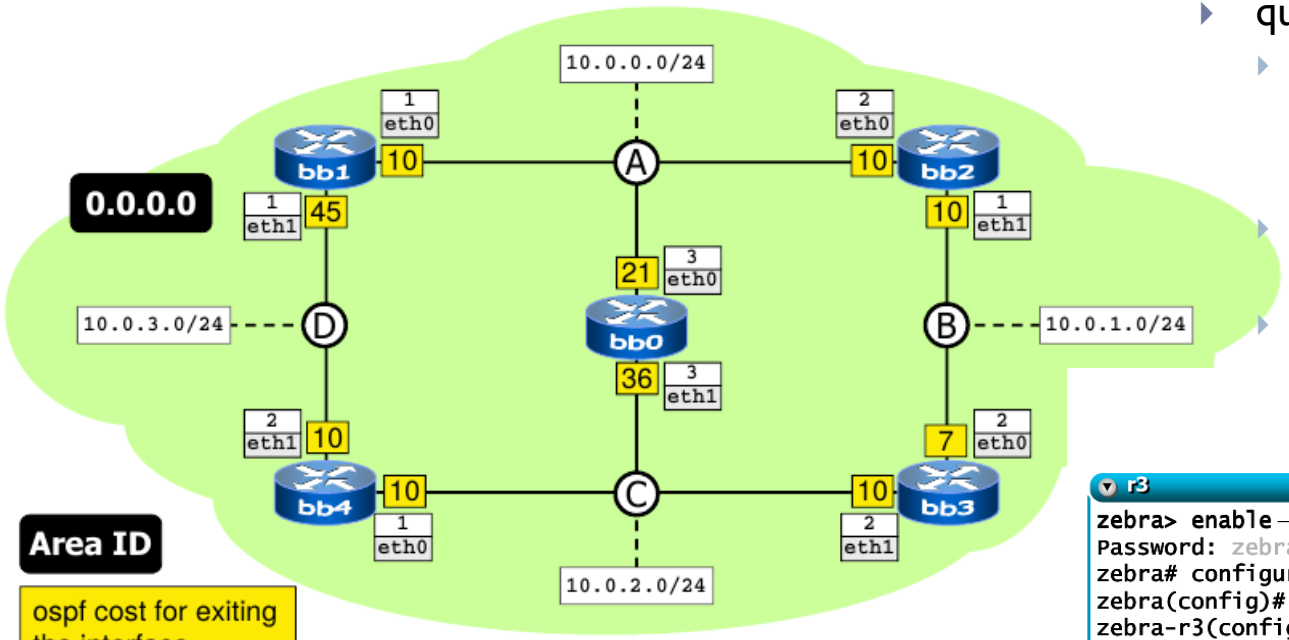

- quagga teszt, pl. bb0 routeren
  - telnet localhost zebra
    - szokásos lehetőségek
      - enable, configure terminal, ?, <tab>, show, list
  - telnet localhost ospfd
    - show ip ospf
  - vtysh (minden démonhoz)

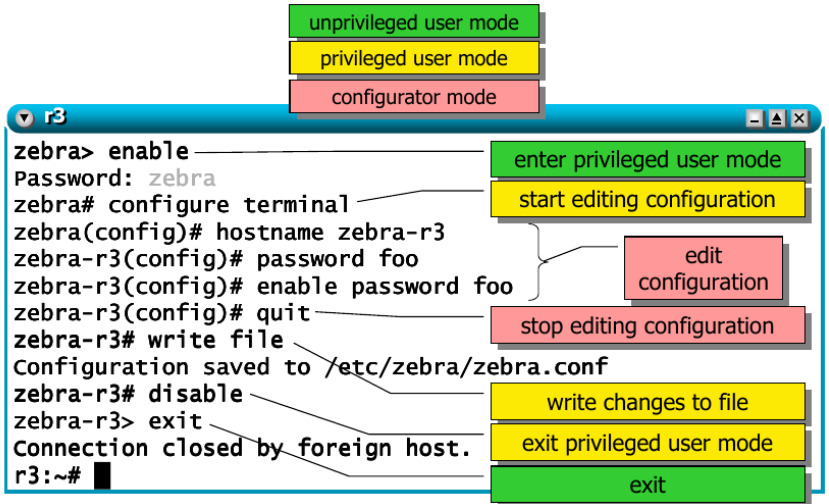

Legrövidebb utak

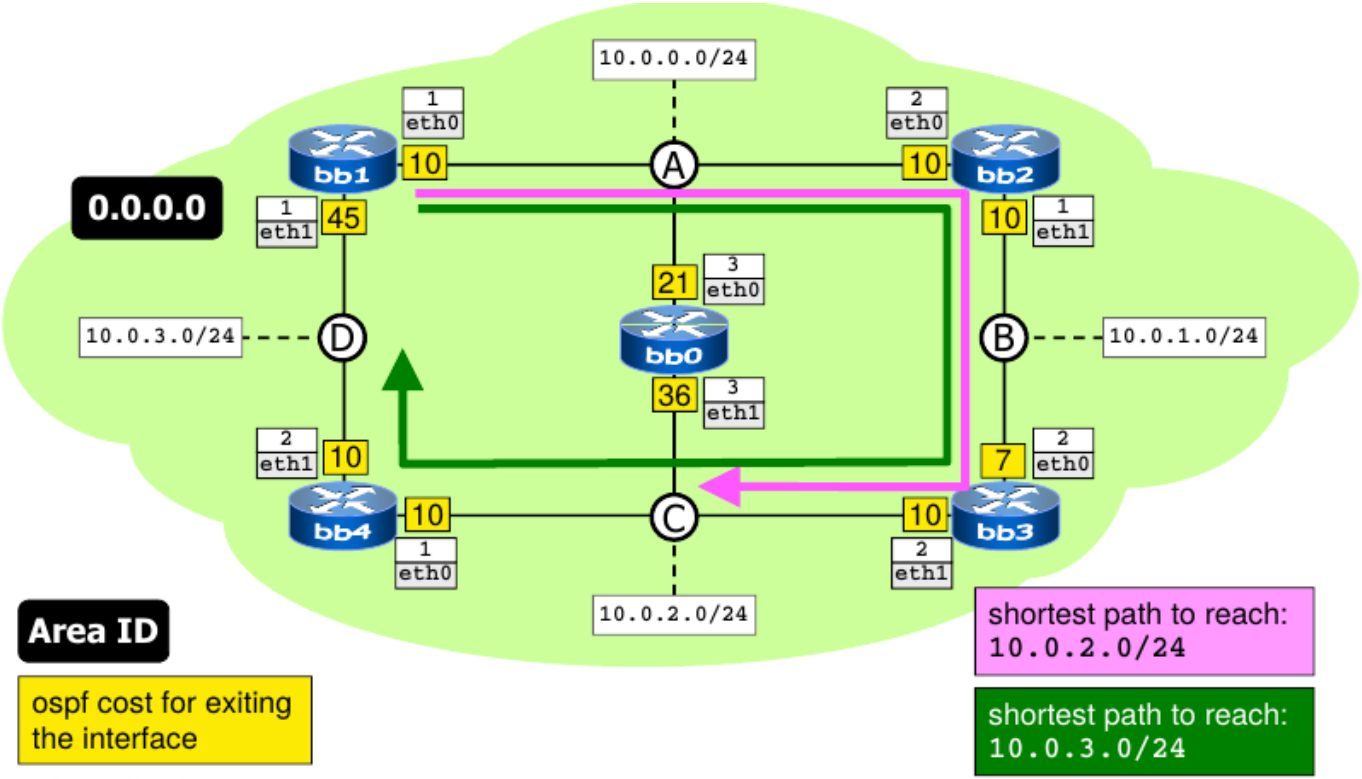

- traceroute -I icmp
  - bb1->10.0.2.1
    - melyik útvonal?
    - hogy jönnek vissza az ICMP válaszok?
  - bb1->10.0.3.2
    - melyik útvonal?

Legrövidebb utak

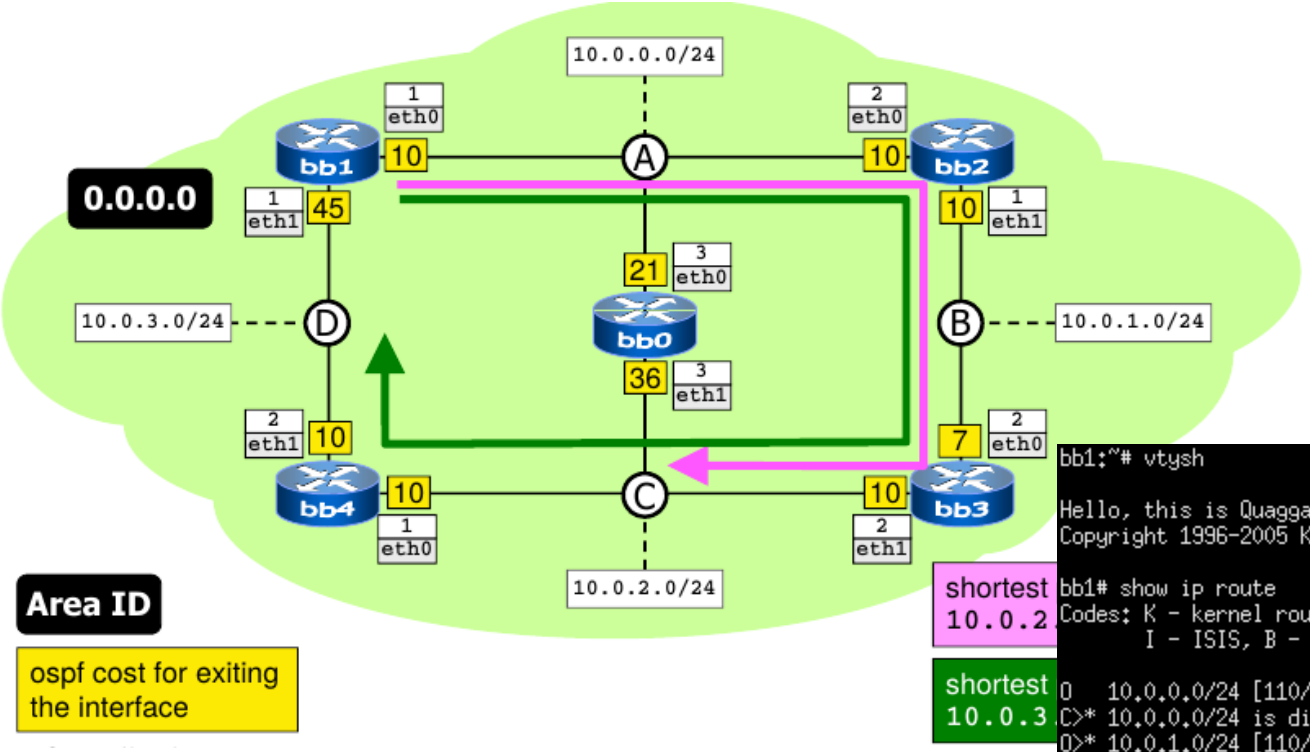

- routing táblák
  - értelmezzük minden routeren
  - vtysh
  - show ip route
    - administrative distance: 110 (default OSPF)
    - ospf metric: 10, 20, …
    - connected metric: 1

```
bb1:~# vtysh

Hello, this is Quagga (version 0.99.10).
Copyright 1996-2005 Kunihiro Ishiguro, et al.

bb1# show ip route
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       I - ISIS, B - BGP, > - selected route, * - FIB route

O   10.0.0.0/24 [110/10] is directly connected, eth0, 00:53:10
C>* 10.0.0.0/24 is directly connected, eth0
O>* 10.0.1.0/20 [110/20] via 10.0.0.2, eth0, 00:53:05
O>* 10.0.2.0/30 [110/30] via 10.0.0.2, eth0, 00:53:05
O   10.0.3.0/24 [110/40] via 10.0.0.2, eth0, 00:53:05
C>* 10.0.3.0/24 is directly connected, eth1
C>* 127.0.0.0/8 is directly connected, lo
bb1#
```

Legrövidebb utak

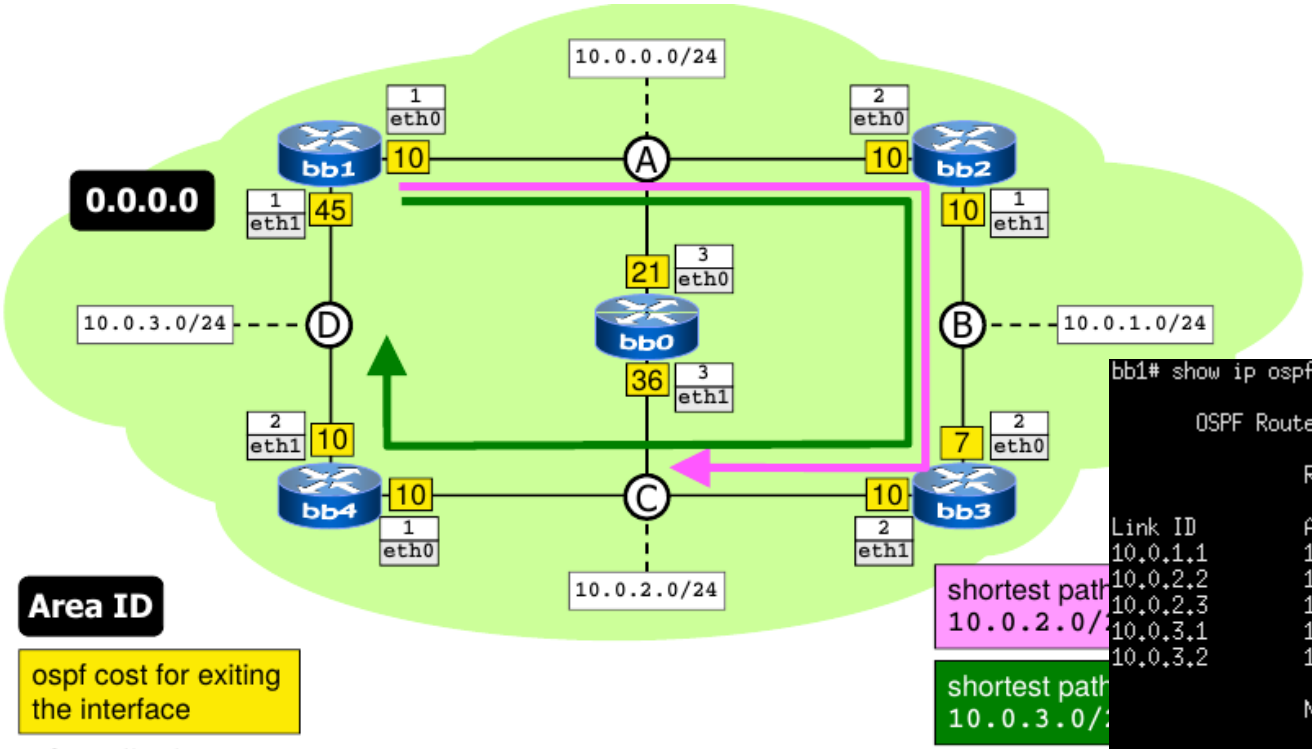

- ospf vizsgálata
  - nézzük meg minden routeren
  - vtysh
  - show ip ospf database
  - show ip ospf neighbor
  - show ip ospf route

```
bb1# show ip ospf database

       OSPF Router with ID (10.0.3.1)

                Router Link States (Area 0.0.0.0)

Link ID         ADV Router      Age  Seq#       CkSum  Link count
10.0.1.1        10.0.1.1         473 0x80000007 0xe1fe 2
10.0.2.2        10.0.2.2         474 0x80000007 0xdbfe 2
10.0.2.3        10.0.2.3         473 0x8000000a 0xd9d4 2
10.0.3.1        10.0.3.1         467 0x8000000a 0x248f 2
10.0.3.2        10.0.3.2         469 0x80000009 0x3e92 2

                Net Link States (Area 0.0.0.0)

Link ID         ADV Router      Age  Seq#       CkSum
10.0.0.1        10.0.3.1         467 0x80000006 0x61ad
10.0.1.2        10.0.2.2         474 0x80000004 0x63be
10.0.2.1        10.0.3.2         468 0x80000006 0x6a9e
10.0.3.2        10.0.3.2         468 0x80000005 0x63b7
```

Designated Router (DR) és BDR

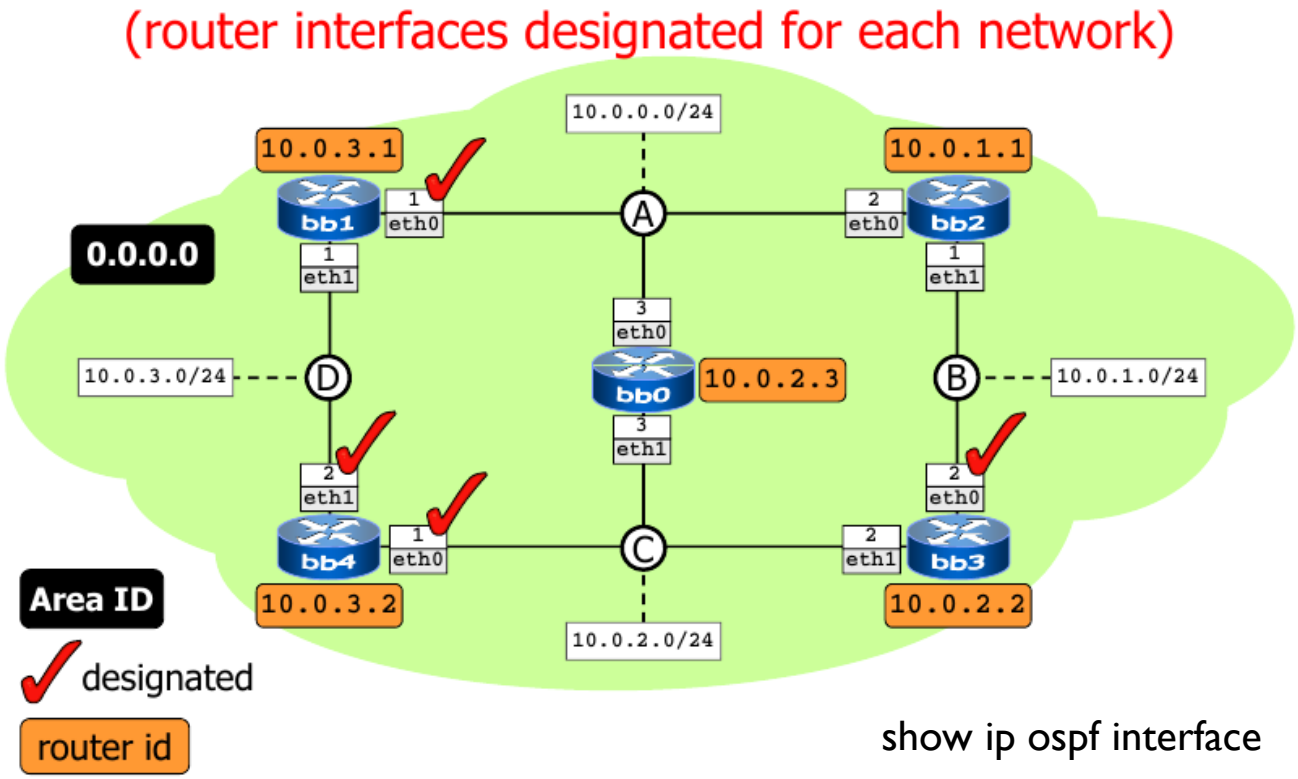

show ip ospf interface

- Broadcast hálózatoknál
  - pl. Ethernet
  - DR és Backup DR: kitüntetett routerek
    - választás alapján
    - router id alapján (max.)
    - (ami interfész id)
  - többi OSPF router csak velük van full szomszédságban
    - különben mindenki-mindenkivel kommunikálna
  - útvonalfrissítés csak DR-től
    - sok erőforrás spórolható

ospf's view of the network

- by exchanging link state update packets, every router learns about the complete network topology, that is:
  - routers
  - subnets
  - adjacencies between routers and networks

ospf's view of the network

last update: Nov 2014

forrás: http://wiki.netkit.org/netkit-labs/netkit-labs_advanced-topics/netkit-labs_ospf/netkit-labs_ospf.pdf

ospf's view of the network

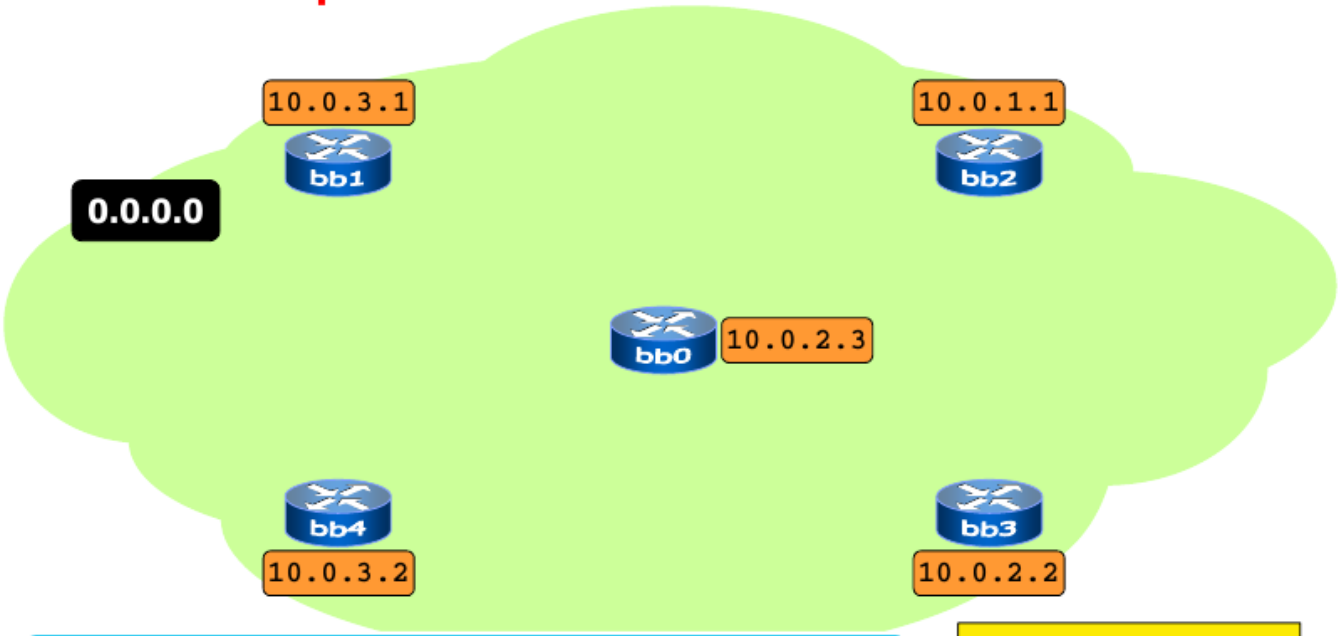


```
bb0
        OSPF Router with ID (10.0.2.3)
                Router Link States (Area 0.0.0.0)
Link ID         ADV Router      Age  Seq#       CkSum  Link count
10.0.1.1        10.0.1.1        553 0x80000003 0xe9fa 2
10.0.2.2        10.0.2.2        552 0x80000003 0xe3fa 2
10.0.2.3        10.0.2.3        552 0x80000003 0xe7cd 2
10.0.3.1        10.0.3.1        552 0x80000003 0x3288 2
10.0.3.2        10.0.3.2        548 0x80000004 0x488d 2
```

for router lsas, the `Link ID` is the router's id

router legnagyobb IP címe

last update: Nov 2014

forrás: http://wiki.netkit.org/netkit-labs/netkit-labs_advanced-topics/netkit-labs_ospf/netkit-labs_ospf.pdf

ospf's view of the network

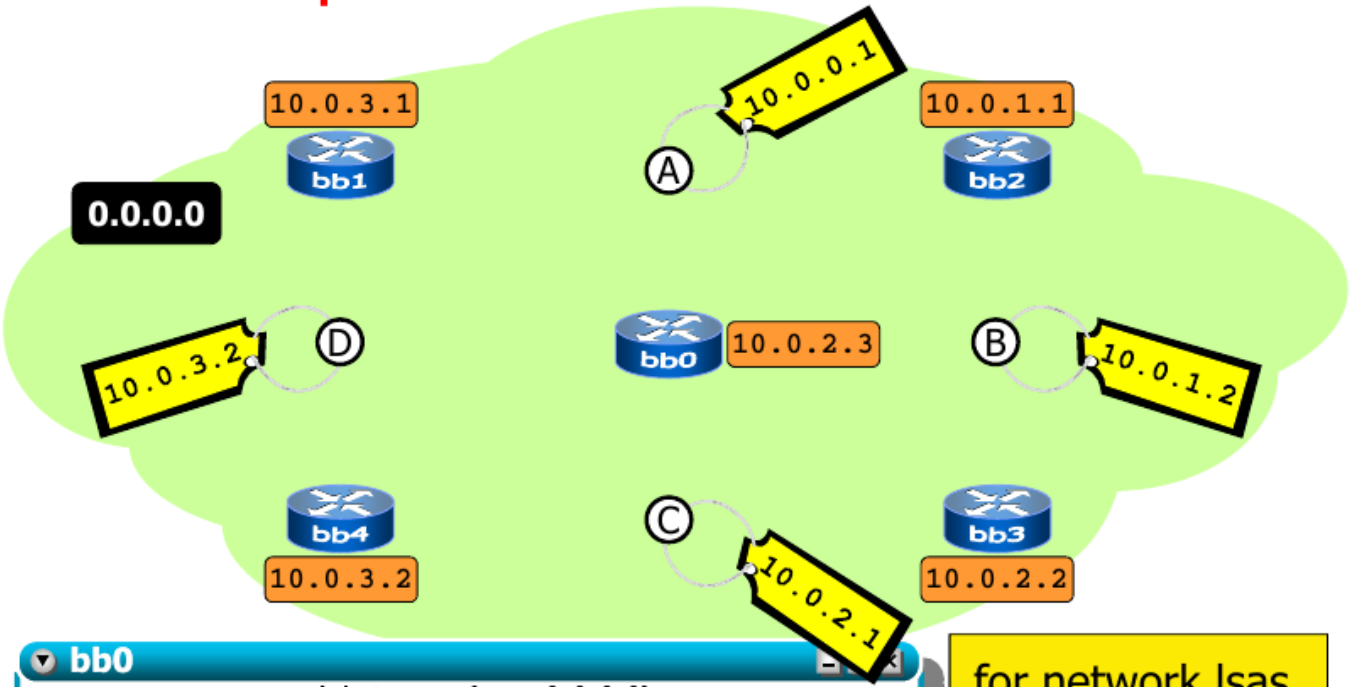


```
bb0
                Net Link States (Area 0.0.0.0)
Link ID         ADV Router      Age  Seq#       CkSum
10.0.0.1        10.0.3.1        557  0x80000001 0x6ba8
10.0.1.2        10.0.2.2        559  0x80000001 0x69bb
10.0.2.1        10.0.3.2        553  0x80000002 0x729a
10.0.3.2        10.0.3.2        553  0x80000001 0x6bb3
```

for network lsas, the `Link ID` is the dr's address

last update: Nov 2014

forrás: http://wiki.netkit.org/netkit-labs/netkit-labs_advanced-topics/netkit-labs_ospf/netkit-labs_ospf.pdf

ospf's view of the network

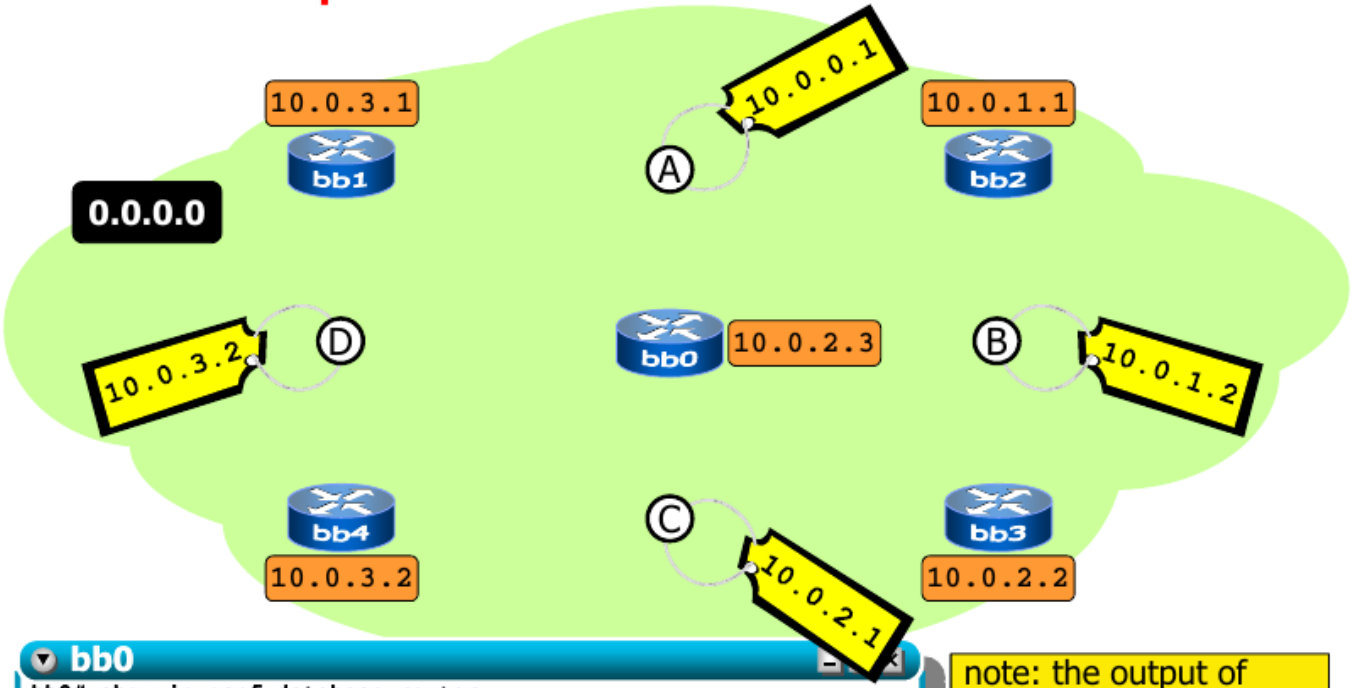


```
bb0
bb0# show ip ospf database router
 Link State ID: 10.0.1.1
   Number of Links: 2
    Link connected to: a Transit Network
     (Link ID) Designated Router address: 10.0.0.1
     (Link Data) Router Interface address: 10.0.0.2
    Link connected to: a Transit Network
     (Link ID) Designated Router address: 10.0.1.2
     (Link Data) Router Interface address: 10.0.1.1
```

note: the output of `show ip ospf database router` has been summarized

last update: Nov 2014

forrás: http://wiki.netkit.org/netkit-labs/netkit-labs_advanced-topics/netkit-labs_ospf/netkit-labs_ospf.pdf

ospf's view of the network

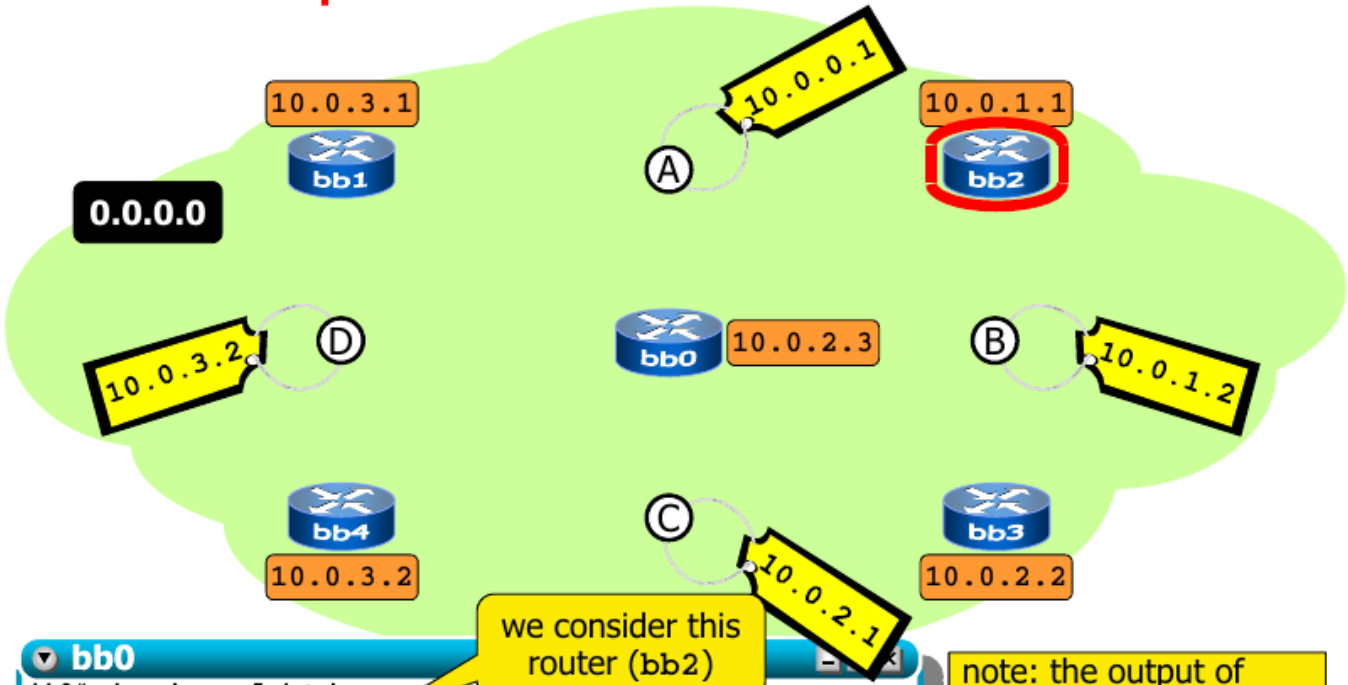

we consider this router (bb2)

```
bb0
bb0# show ip ospf database router
 Link State ID: 10.0.1.1
   Number of Links: 2
    Link connected to: a Transit Network
     (Link ID) Designated Router address: 10.0.0.1
     (Link Data) Router Interface address: 10.0.0.2
    Link connected to: a Transit Network
     (Link ID) Designated Router address: 10.0.1.2
     (Link Data) Router Interface address: 10.0.1.1
```

note: the output of `show ip ospf database router` has been summarized

last update: Nov 2014

forrás: http://wiki.netkit.org/netkit-labs/netkit-labs_advanced-topics/netkit-labs_ospf/netkit-labs_ospf.pdf

ospf's view of the network

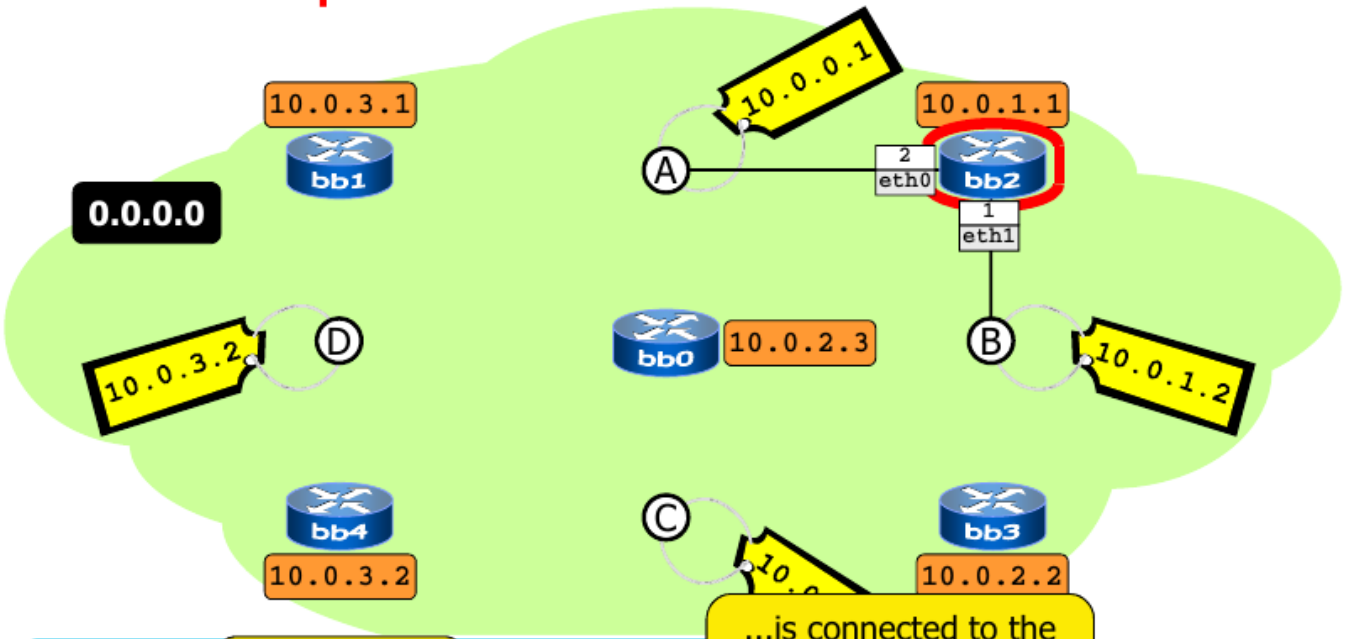

last update: Nov 2014

forrás: http://wiki.netkit.org/netkit-labs/netkit-labs_advanced-topics/netkit-labs_ospf/netkit-labs_ospf.pdf

ospf's view of the network

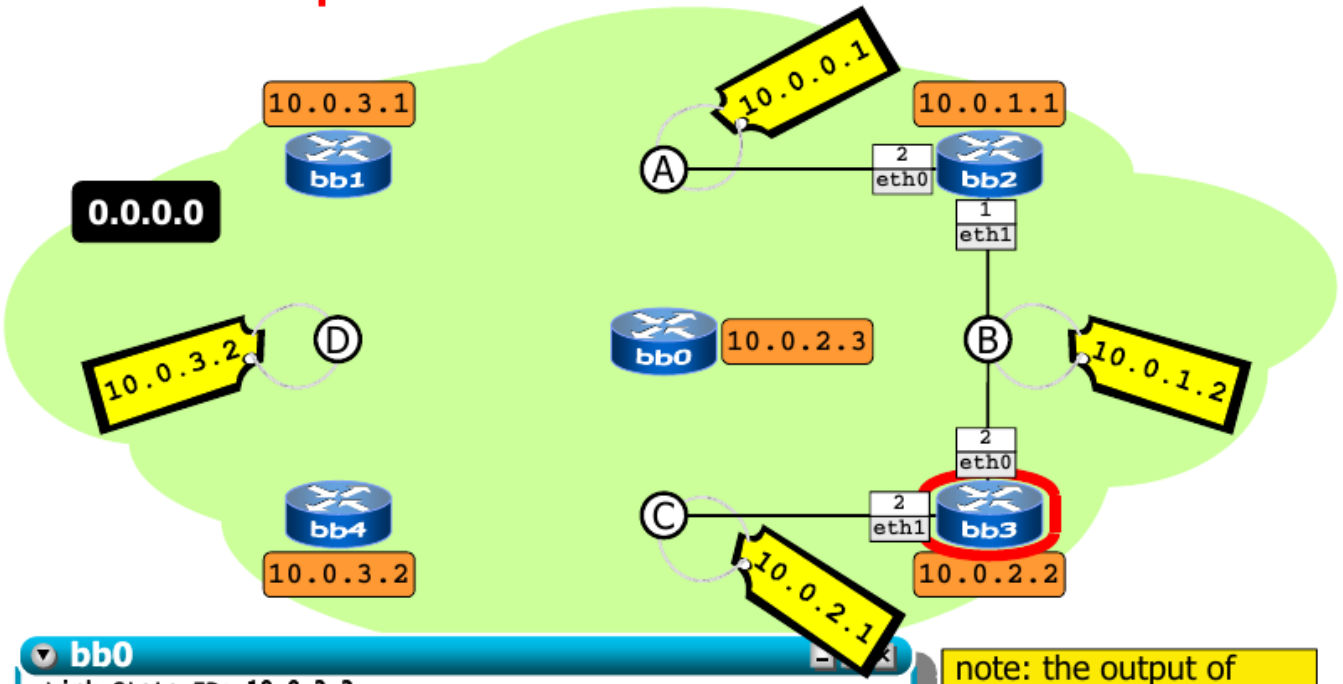


```
bb0
Link State ID: 10.0.2.2
  Number of Links: 2
   Link connected to: a Transit Network
    (Link ID) Designated Router address: 10.0.1.2
    (Link Data) Router Interface address: 10.0.1.2
   Link connected to: a Transit Network
    (Link ID) Designated Router address: 10.0.2.1
    (Link Data) Router Interface address: 10.0.2.2
```

note: the output of `show ip ospf database router` has been summarized

last update: Nov 2014

forrás: http://wiki.netkit.org/netkit-labs/netkit-labs_advanced-topics/netkit-labs_ospf/netkit-labs_ospf.pdf

ospf's view of the network

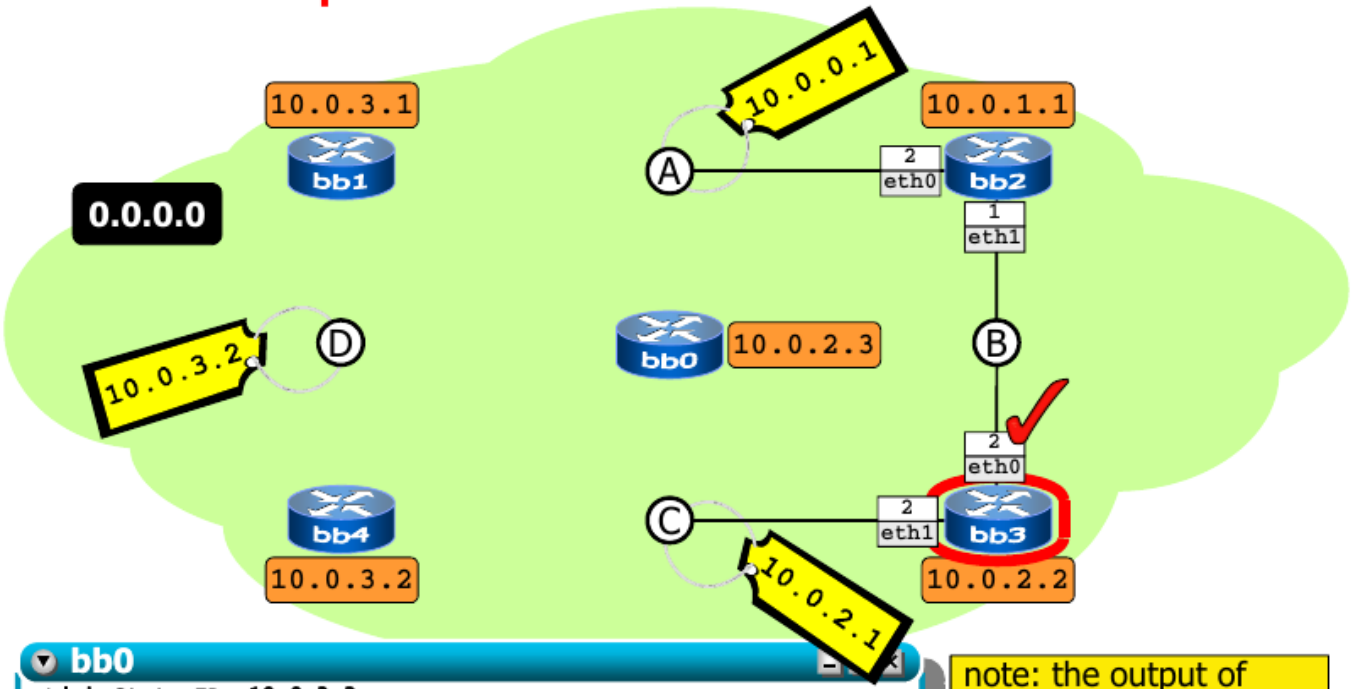


```
bb0
Link State ID: 10.0.2.2
  Number of Links: 2
   Link connected to: a Transit Network
    (Link ID) Designated Router address: 10.0.1.2
    (Link Data) Router Interface address: 10.0.1.2
   Link connected to: a Transit Network
    (Link ID) Designated Router address: 10.0.2.1
    (Link Data) Router Interface address: 10.0.2.2
```

note: the output of `show ip ospf database router` has been summarized

last update: Nov 2014

forrás: http://wiki.netkit.org/netkit-labs/netkit-labs_advanced-topics/netkit-labs_ospf/netkit-labs_ospf.pdf

ospf's view of the network

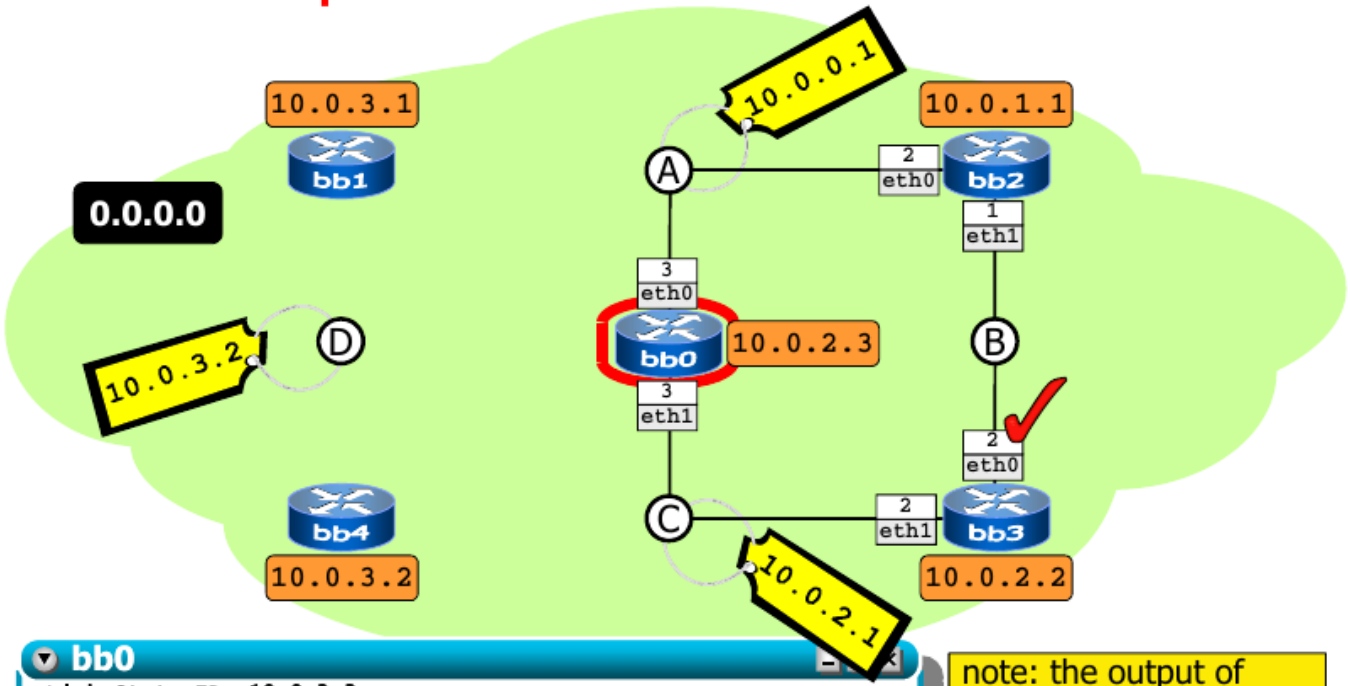


```
bb0
Link State ID: 10.0.2.3
  Number of Links: 2
   Link connected to: a Transit Network
    (Link ID) Designated Router address: 10.0.0.1
    (Link Data) Router Interface address: 10.0.0.3
   Link connected to: a Transit Network
    (Link ID) Designated Router address: 10.0.2.1
    (Link Data) Router Interface address: 10.0.2.3
```

note: the output of `show ip ospf database router` has been summarized

last update: Nov 2014

forrás: http://wiki.netkit.org/netkit-labs/netkit-labs_advanced-topics/netkit-labs_ospf/netkit-labs_ospf.pdf

ospf's view of the network

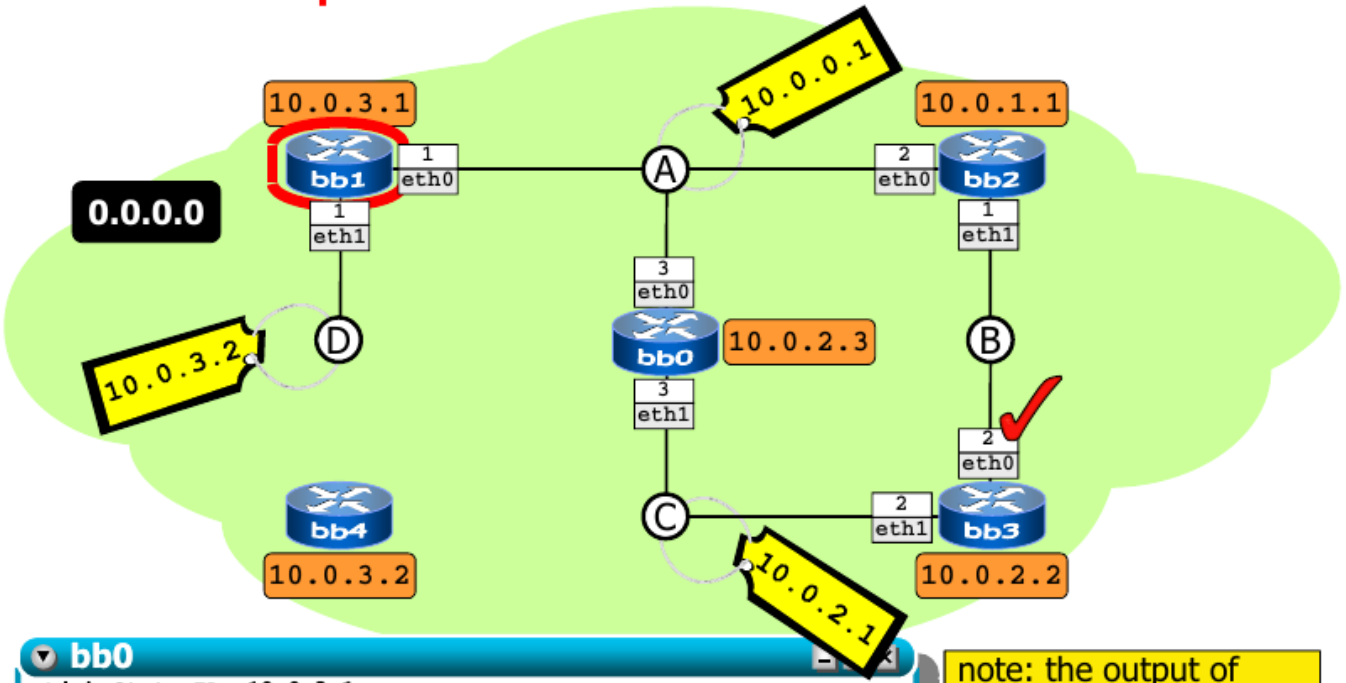


```
bb0
Link State ID: 10.0.3.1
  Number of Links: 2
   Link connected to: a Transit Network
    (Link ID) Designated Router address: 10.0.0.1
    (Link Data) Router Interface address: 10.0.0.1
   Link connected to: a Transit Network
    (Link ID) Designated Router address: 10.0.3.2
    (Link Data) Router Interface address: 10.0.3.1
```

note: the output of `show ip ospf database router` has been summarized

last update: Nov 2014

forrás: http://wiki.netkit.org/netkit-labs/netkit-labs_advanced-topics/netkit-labs_ospf/netkit-labs_ospf.pdf

ospf's view of the network

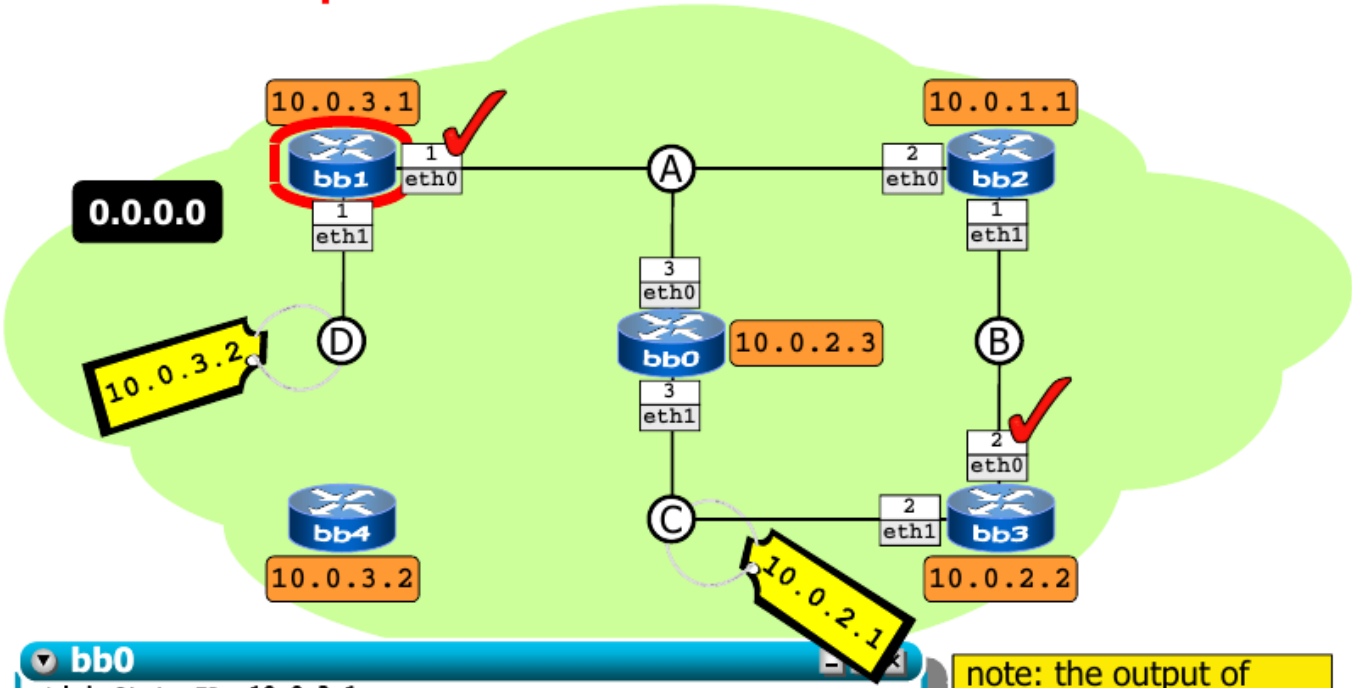


```
bb0
Link State ID: 10.0.3.1
  Number of Links: 2
   Link connected to: a Transit Network
    (Link ID) Designated Router address: 10.0.0.1
    (Link Data) Router Interface address: 10.0.0.1
   Link connected to: a Transit Network
    (Link ID) Designated Router address: 10.0.3.2
    (Link Data) Router Interface address: 10.0.3.1
```

note: the output of `show ip ospf database router` has been summarized

last update: Nov 2014

forrás: http://wiki.netkit.org/netkit-labs/netkit-labs_advanced-topics/netkit-labs_ospf/netkit-labs_ospf.pdf

ospf's view of the network

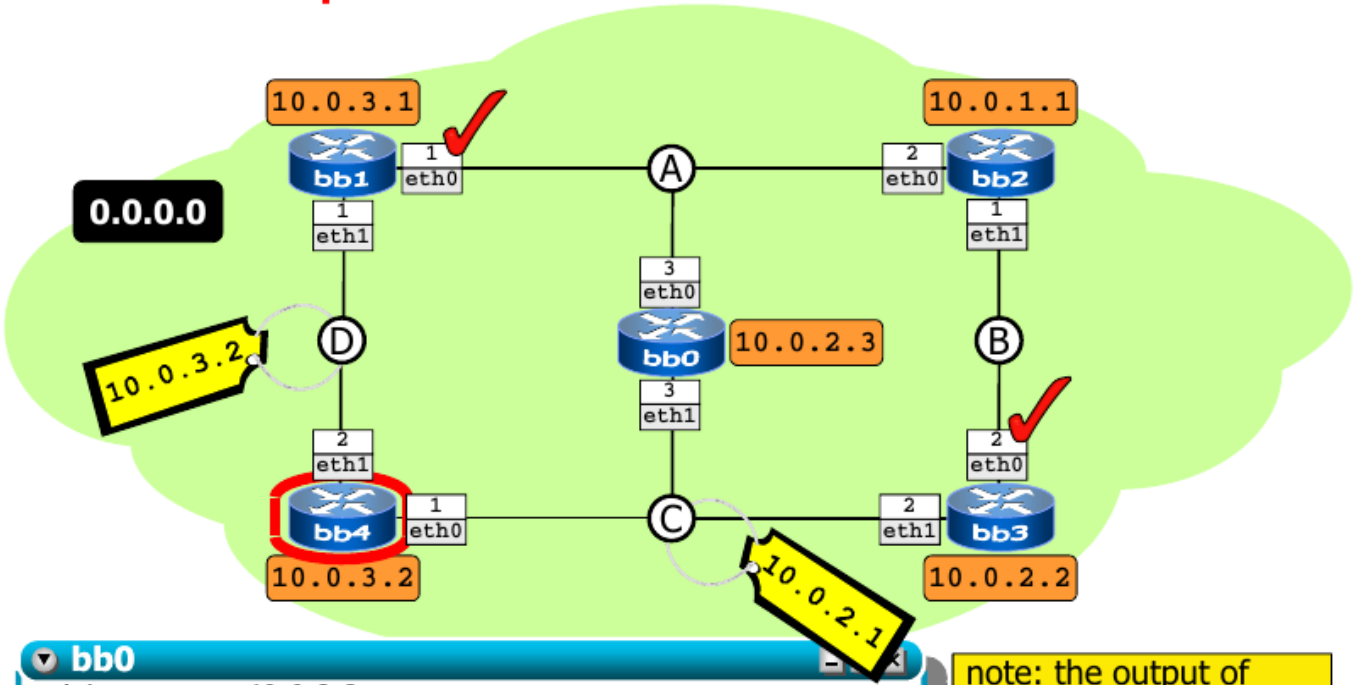


```
bb0
Link State ID: 10.0.3.2
  Number of Links: 2
   Link connected to: a Transit Network
    (Link ID) Designated Router address: 10.0.2.1
    (Link Data) Router Interface address: 10.0.2.1
   Link connected to: a Transit Network
    (Link ID) Designated Router address: 10.0.3.2
    (Link Data) Router Interface address: 10.0.3.2
```

note: the output of `show ip ospf database router` has been summarized

last update: Nov 2014

forrás: http://wiki.netkit.org/netkit-labs/netkit-labs_advanced-topics/netkit-labs_ospf/netkit-labs_ospf.pdf

ospf's view of the network

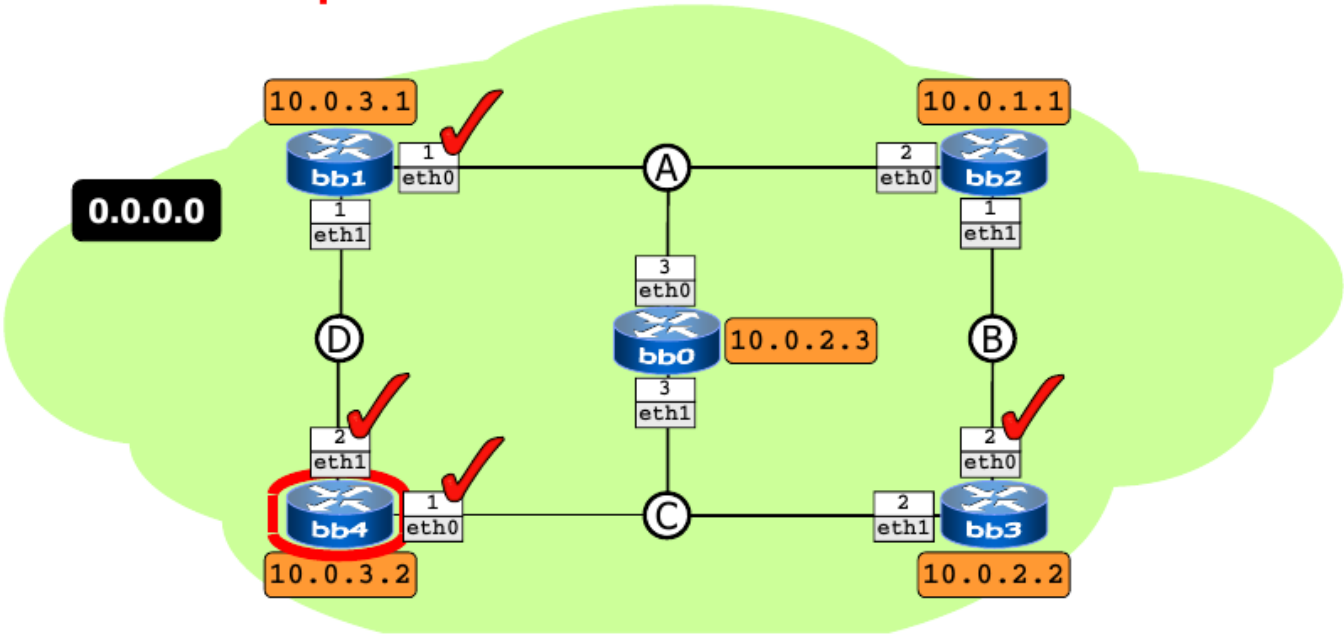


```
bb0
Link State ID: 10.0.3.2
  Number of Links: 2
   Link connected to: a Transit Network
    (Link ID) Designated Router address: 10.0.2.1
    (Link Data) Router Interface address: 10.0.2.1
   Link connected to: a Transit Network
    (Link ID) Designated Router address: 10.0.3.2
    (Link Data) Router Interface address: 10.0.3.2
```

note: the output of `show ip ospf database router` has been summarized

last update: Nov 2014

forrás: http://wiki.netkit.org/netkit-labs/netkit-labs_advanced-topics/netkit-labs_ospf/netkit-labs_ospf.pdf

ospf's view of the network

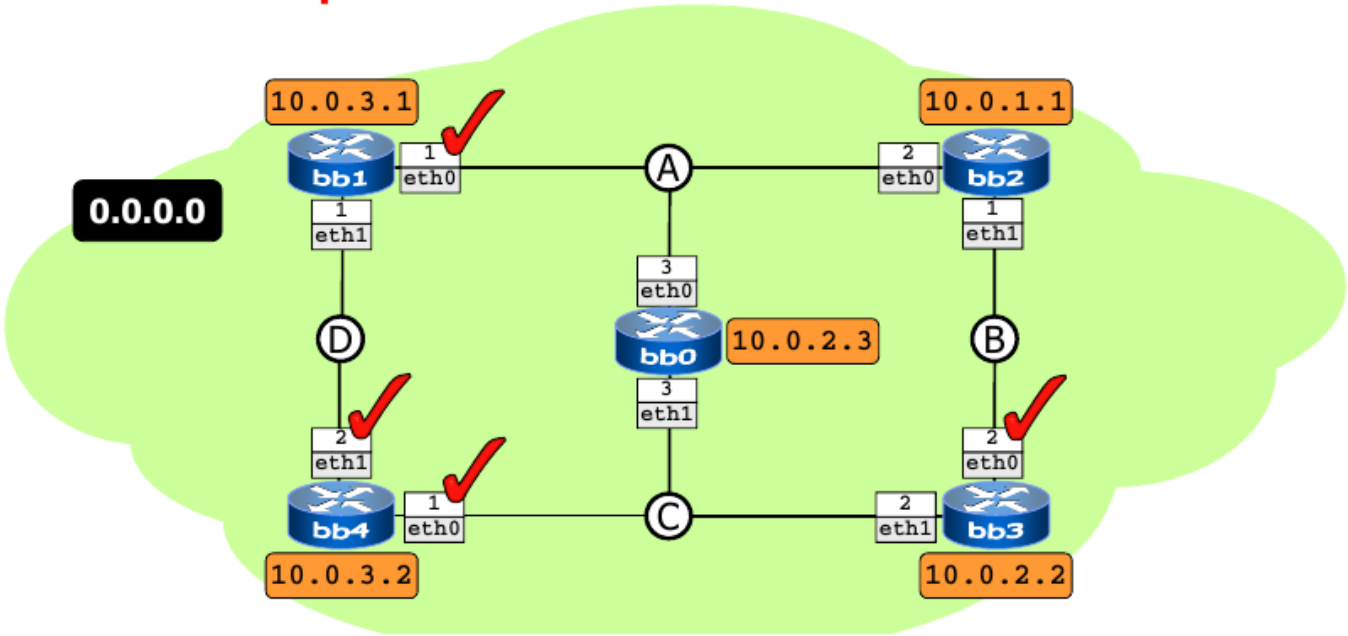

last update: Nov 2014

forrás: http://wiki.netkit.org/netkit-labs/netkit-labs_advanced-topics/netkit-labs_ospf/netkit-labs_ospf.pdf

ospf's view of the network

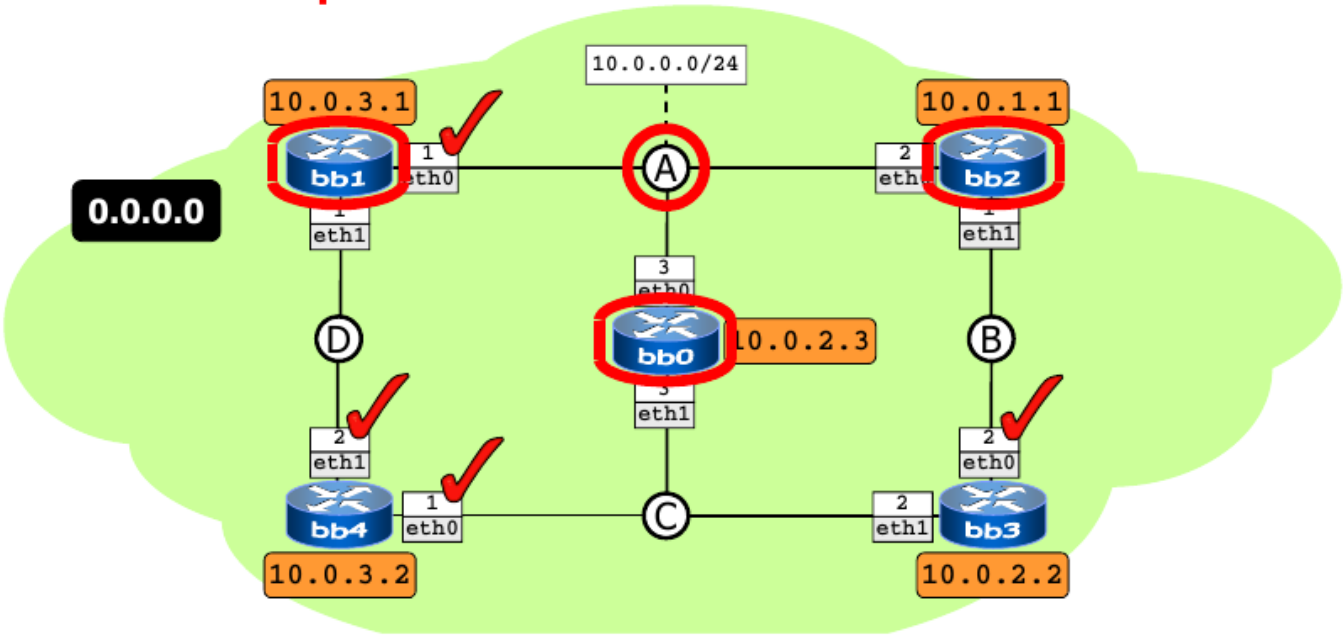


```
bb0
Link State ID: 10.0.0.1 (address of Designated Router)
 Advertising Router: 10.0.3.1
 Network Mask: /24
       Attached Router: 10.0.3.1
       Attached Router: 10.0.1.1
       Attached Router: 10.0.2.3
```

note: the output of `show ip ospf database network` has been summarized

last update: Nov 2014

forrás: http://wiki.netkit.org/netkit-labs/netkit-labs_advanced-topics/netkit-labs_ospf/netkit-labs_ospf.pdf

ospf's view of the network

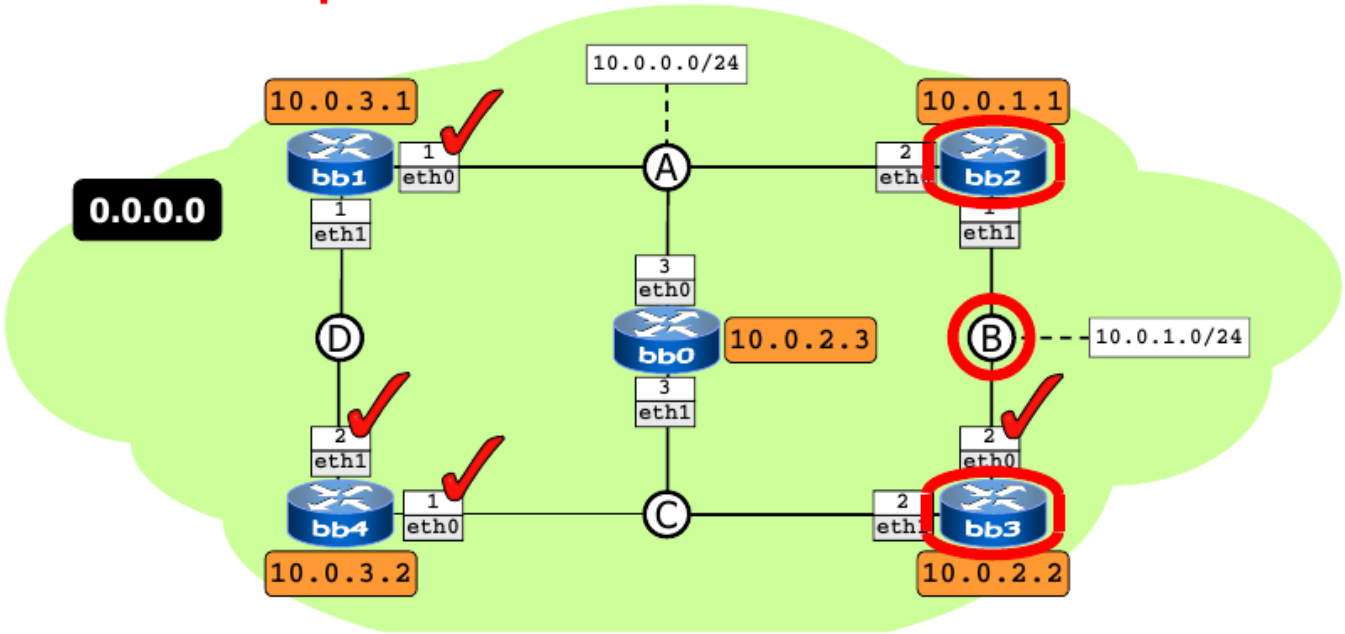


```
bb0
Link State ID: 10.0.1.2 (address of Designated Router)
 Advertising Router: 10.0.2.2
 Network Mask: /24
       Attached Router: 10.0.1.1
       Attached Router: 10.0.2.2
```

note: the output of `show ip ospf database network` has been summarized

last update: Nov 2014

forrás: http://wiki.netkit.org/netkit-labs/netkit-labs_advanced-topics/netkit-labs_ospf/netkit-labs_ospf.pdf

ospf's view of the network

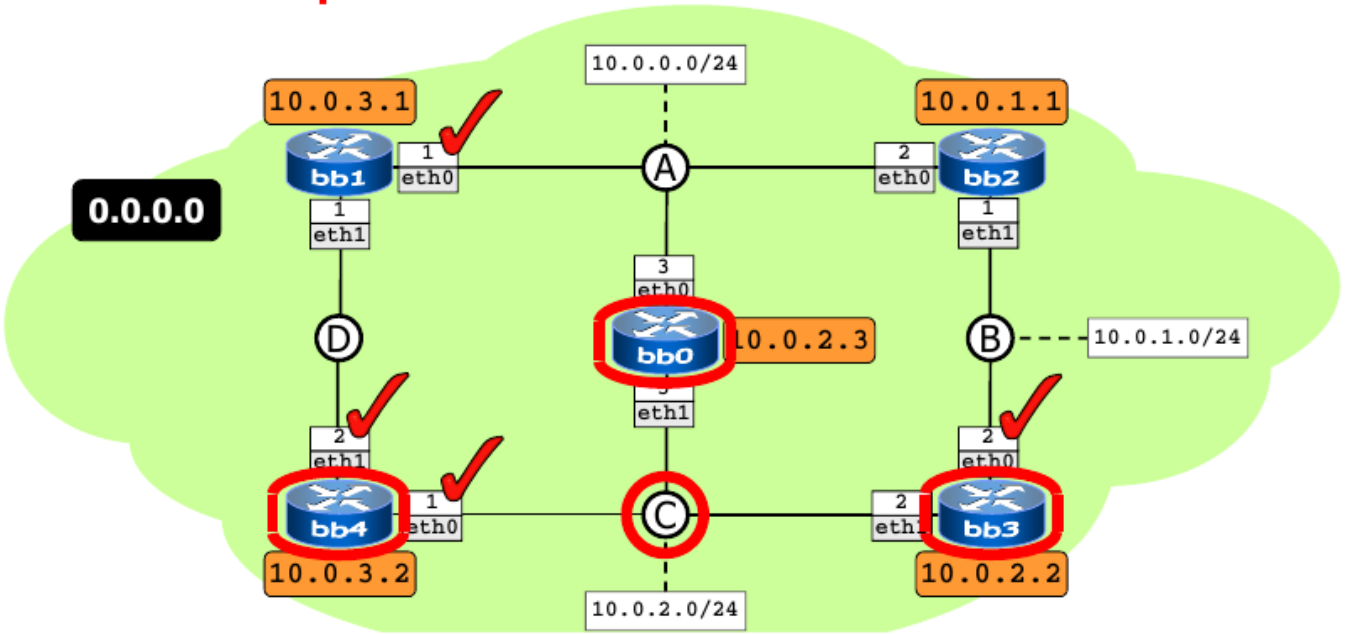

bb0

```
Link State ID: 10.0.2.1 (address of Designated Router)
 Advertising Router: 10.0.3.2
 Network Mask: /24
       Attached Router: 10.0.3.2
       Attached Router: 10.0.2.2
       Attached Router: 10.0.2.3
```

note: the output of `show ip ospf database network` has been summarized

last update: Nov 2014

forrás: http://wiki.netkit.org/netkit-labs/netkit-labs_advanced-topics/netkit-labs_ospf/netkit-labs_ospf.pdf

ospf's view of the network

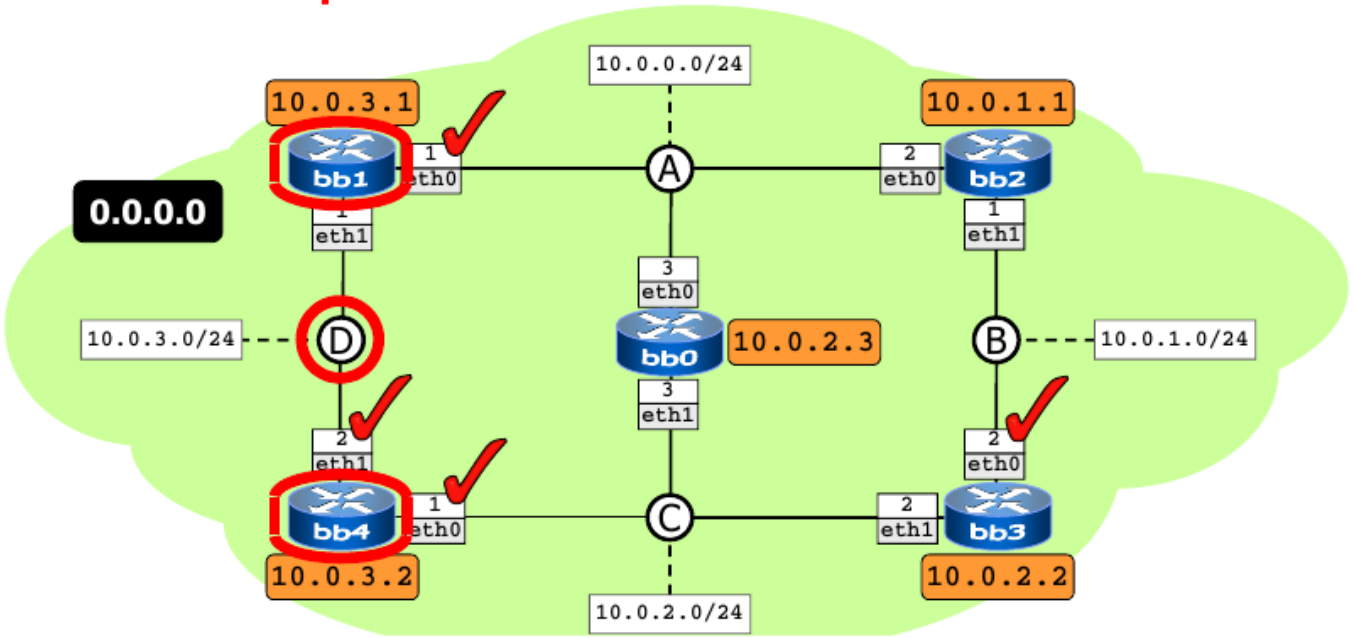


```
bb0
Link State ID: 10.0.3.2 (address of Designated Router)
 Advertising Router: 10.0.3.2
 Network Mask: /24
       Attached Router: 10.0.3.1
       Attached Router: 10.0.3.2
```

note: the output of `show ip ospf database network` has been summarized

last update: Nov 2014

forrás: http://wiki.netkit.org/netkit-labs/netkit-labs_advanced-topics/netkit-labs_ospf/netkit-labs_ospf.pdf

ospf's view of the network

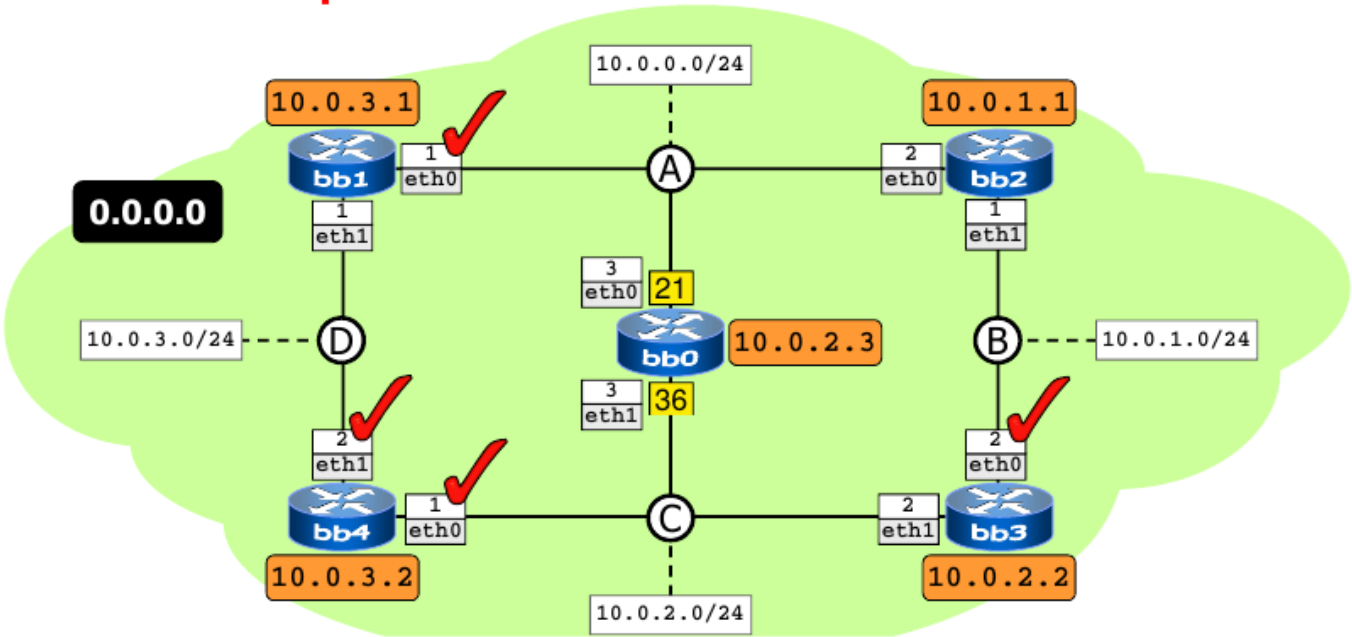

bb0

```
bb0:~# vtysh -e "show ip ospf interface" | egrep "eth|Cost"
eth0 is up
  Router ID 10.0.2.3, Network Type BROADCAST, Cost: 21
eth1 is up
  Router ID 10.0.2.3, Network Type BROADCAST, Cost: 36
```

a shortcut to quickly get the cost

ospf interface costs can be queried on all routers

last update: Nov 2014

forrás: http://wiki.netkit.org/netkit-labs/netkit-labs_advanced-topics/netkit-labs_ospf/netkit-labs_ospf.pdf

Dinamikus működés

- OSPF Hello üzenetek
  - tcpdump -ne -vv
  - vizsgáljuk meg részletesebben
- Mi történik, ha lemegy egy link?
  - ifconfig vagy ip parancs használható
  - hogy változnak az útvonalak?
    - show ip ospf route
  - ha DR megy le ezzel?
    - show ip ospf database network
- Mi történik, ha lemegy egy router?
  - pl. minden interfészét down állapotba kapcsoljuk