

1off.it Technical Presentation

ZeroTier Deep Dive



08 mar 2024



About me

Alessandro Campanella

I have been working in the telecommunications field for over 20 years as a producer, consultant, and trainer, both in Italy and abroad.

I was among the pioneers of wireless technology, conducting experiments on my network and gathering the experiences of the many operators and integrators with whom I have collaborated over the years.

MikroTik Evangelist
Design, analysis, and consulting for ISPs
MikroTik Training and Certification



a.campanella@1off.it



[@alessandrocampanella](https://t.me/alessandrocampanella)



MUM 2006

I was there (here)

Returning to the place of the first MUM after 18 years, reconnecting with friends and colleagues, evokes a mix of nostalgia and joy.

It's a powerful reminder of the journey undertaken and the enduring bonds formed.

This reunion is both a reflection on past experiences and a celebration of growth and enduring connections.

I was there Gandalf. I was there 3000 years ago...





ZeroTier Deep Dive

A deep dive into the capabilities of ZeroTier and MikroTik



Module 1

Introduction





ZeroTier

What is it?

“ZeroTier is a way to connect devices over your own private network anywhere in the world.

You do this by creating a network and then joining two or more devices to that network.

You can use ZeroTier to play games, connect to remote business resources or even as a cloud backplane for your enterprise.”

Adam Ierymenko

Founder, ZeroTier



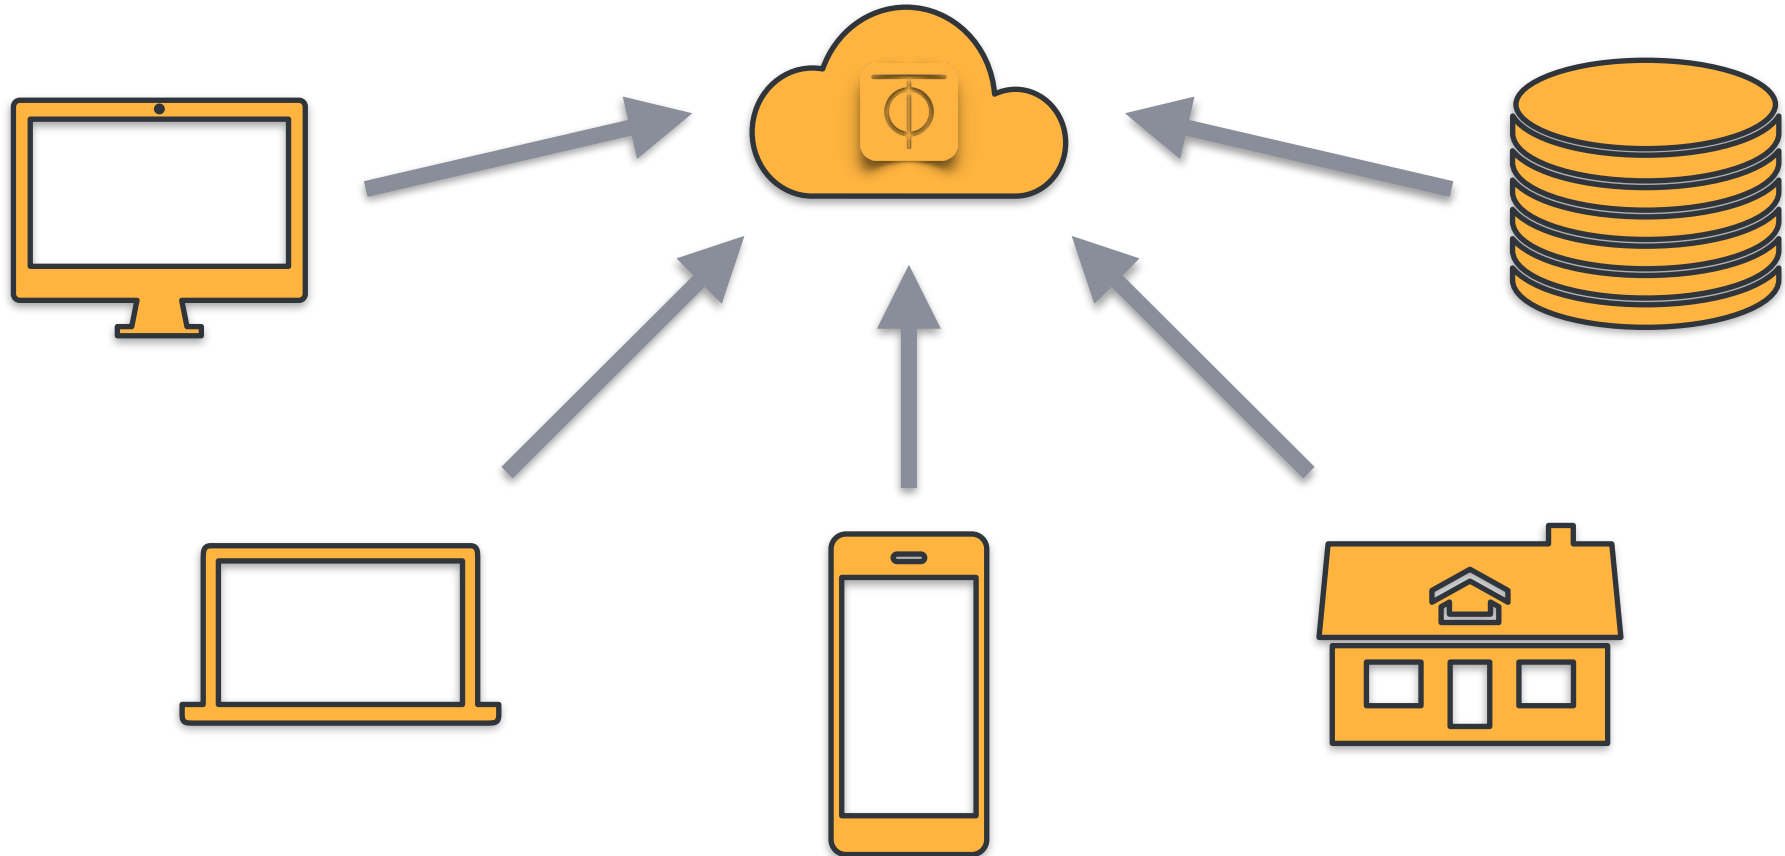
The Planetary Data Center

Manage everything everywhere in the same way

Imagine having:

- All network resources in the same room
- You have a management system with L2 and L3 administration capabilities
 - Devices can connect directly to each other as if they were adjacent at L2
 - Each device is authenticated, and traffic is subjected to strong encryption
 - You can group devices into VLANs to create group rules
 - Traffic can be filtered and monitored
- Devices can move freely within the room

The Planetary Data Center





Network hypervisor, not VPN!

Many faces of the same coin

We give applications different names based on their main function, but all are based on **SDN (Software Defined Networking)**

- Virtual Private Network
- Multicloud
- SD-WAN



ZeroTier Network hypervisor

VL1: “Virtual Layer One”

- Uses existing communication channels to interconnect ZeroTier nodes
- Provides encryption, authentication, and key management
- Manages the construction of direct links through NAT and firewalls
- Rebuilds links when existing ones drop or when a device moves
- Offers advanced SD-WAN features like multipath and active/backup



ZeroTier Network hypervisor

VL2: “Virtual Layer Two”

- Ethernet emulator built on VL1
- L2 and L3 filtering functions and microsegmentation
- Supports multicast, broadcast, and non-IP protocols
- L3 authentication to prevent L2 attacks
- L2 Bridging and L3 Routing to connect to the rest of the world



Advantages of ZeroTier

Simple incremental implementation

- In most networks, ZeroTier just "works" without requiring changes to the existing infrastructure
- It can be deployed in isolated networks not connected to the Internet
- Can be installed on switches, routers, or directly on devices
- Adoption can be incremental, no "rip and replace"
- No dedicated hardware is needed



Advantages of ZeroTier

Security

- All traffic is encrypted end-to-end and always authenticated
- Key and permission distribution is automatic
- VL1 provides the cryptographic infrastructure used by VL2 to authenticate nodes, certificates, and access credentials
- You can independently manage your personal keys and security boundaries
 - VL1 is unaware of existing networks and network links
 - If you host your controller, you own the secret key of your network's CA



Advantages of ZeroTier

Rules and Microsegmentation

- The integrated engine allows defining rules that operate both at Layer 2 (Ethernet) and Layer 3 (IPv4/IPv6)
- The rules establish the permitted behaviors within the network
- Tags and Capabilities allow ACL-style control for each node
- "Tee" and "Redirect" rules enable the implementation of monitoring and security systems by intercepting, diverting, or duplicating traffic selectively
- Changes are generally applied within a minute



Advantages of ZeroTier

Layer2 features

- There's no overhead, and the MTU is 2800!
- ZeroTier interfaces can be bridged with physical (or virtual) interfaces for Layer 2 expansion
- There is support for multicast and broadcast
- IPv4 ARP is handled transparently in a WAN-friendly manner. For devices, it's normal ARP.
- All non-IP protocols pass through (PPPoE, VLAN, etc.)
- It can carry BGP, MPLS, OSPF, RIP, IS-IS



Advantages of ZeroTier

Layer3 features

- There are no restrictions on IP addresses; nodes can have IPv4 and/or IPv6
 - The controller can send IP configurations, routes, and DNS settings (similar to DHCP)
 - Nodes can set their own IP addresses on the ZeroTier interface just as if it were an Ethernet interface
 - IPv6 link-local addresses are set automatically
- 6PLANE, a special IPv6 address scheme
 - Each node automatically acquires an /80 IP address
 - Each node can assign up to 2^{48} local addresses without the need to announce them
 - NDP is emulated locally and can coexist with classic IPv6 addresses



Advantages of ZeroTier

Physical Network Bonding and Failover (SD-WAN)

- Nodes can use multiple connections simultaneously
- Aggregation methods:
 - Round Robin
 - Load balancing
 - Traffic duplication
 - Active / Backup with rapid failover
- More than one WAN can be used to increase performance and reliability



Advantages of ZeroTier

To conclude...

- Nodes on the same LAN automatically recognize each other and do not communicate across the WAN
- ZeroTier can recognize private WANs and use them
- Router configuration protocols such as uPnP are supported
- If nodes cannot connect directly to each other, root servers can act as a bridge and relay the traffic



Advantages of ZeroTier

Supported Platforms

- Windows
- MacOS
- iOS (Apple)
- Linux
- Android
- FreeBSD
- OpenBSD
- NetBSD
- X86 / X86_64
- ARM32 / ARM64
- MIPS / MIPS64
- PPC64
- S390x
- RISC-V



ZeroTier Use Cases

What can it be used for...

- VPN L2/L3
 - Site to Site, Roadwarrior VPN, Mesh VPN
- OOB Device Management - IoT
 - Secure access to devices directly at Layer 2
 - Bridging to network segments that are difficult to reach
- IT Services
 - Remote access to user systems for remote assistance
- Monitoraggio / Intrusion Detection
 - Mirroring or streaming traffic to IDS



Module 2

Let's create a network



Creation of a ZeroTier Network

On ZeroTier.com

- Register an account
- Create a Network
 - 1 Admin
 - Unlimited Networks
 - Up to 25 nodes
 - FREE
 - 10\$/month for additional Admin
 - 5\$/month for additional 25 nodes

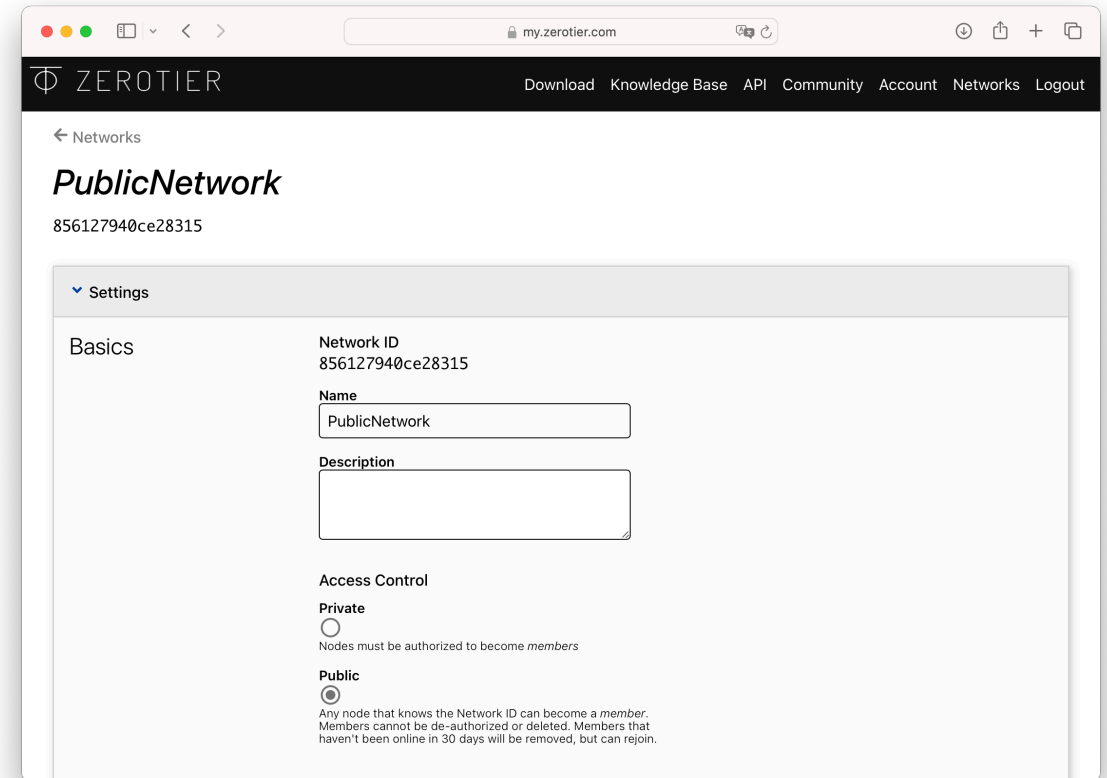
The screenshot shows the ZeroTier.com web interface. At the top, there is a navigation bar with the ZeroTier logo and links for Download, Knowledge Base, API, Community, Account, Networks, and Logout. A prominent orange button labeled "Create A Network" is centered on the page. Below this, the "Your Networks" section is displayed, showing a search bar with "3 networks..." and a table of existing networks.

NETWORK ID	NAME	DESCRIPTION	SUBNET	NODES	CREATED
1c33c1ced02e5752	lab		172.29.0.0/16	2	2023-11-25
856127940ce28315	PublicNetwork		172.25.0.0/16	4	2023-12-15
52b337794f1b7bd6	zt1		10.147.19.0/24	5	2023-12-14

Creation of a ZeroTier Network

Private or Public?

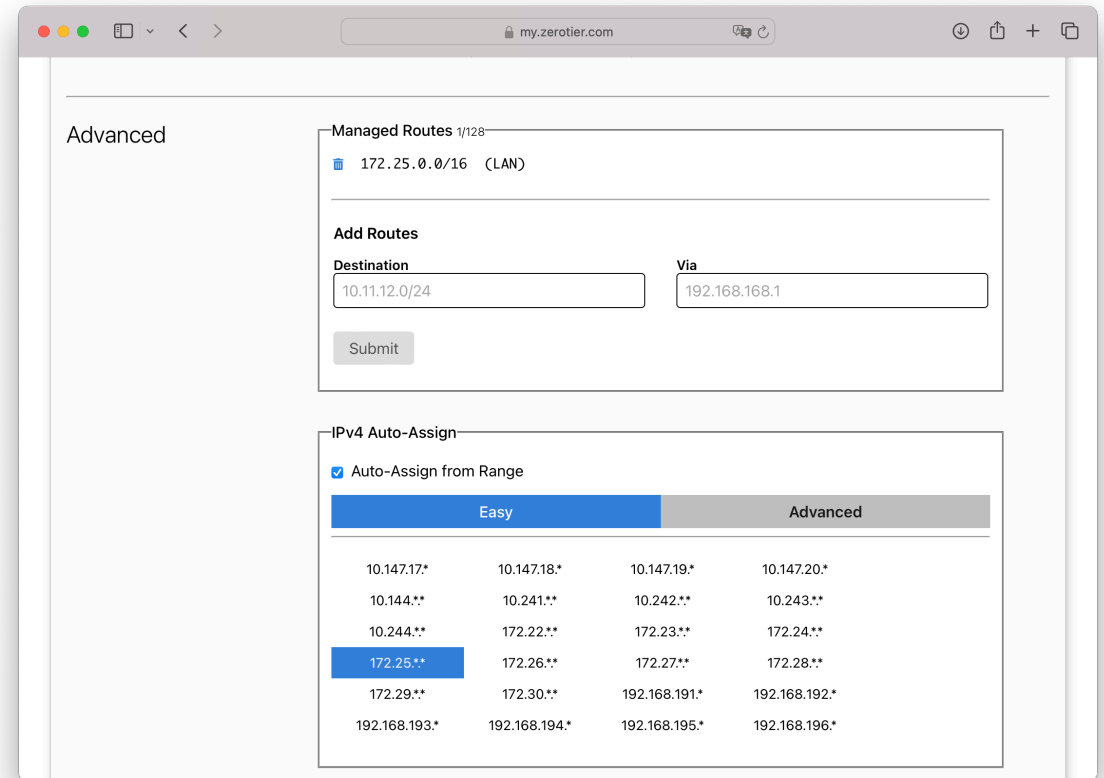
- The Network ID is the unique identifier of our network, used as the address for connection
- A Private network requires authentication for each node
- A Public network does not require authentication



Creation of a ZeroTier Network

Assign IP and routes

- We can choose a subnet from those proposed or we can customize it
- By default, we will send clients a route to reach the other nodes
- We can add routes and gateways



The screenshot shows the ZeroTier web interface for configuring routes. The browser address bar shows `my.zerotier.com`. The page is titled "Advanced" and displays "Managed Routes 1/128" with a list containing `172.25.0.0/16 (LAN)`. Below this is the "Add Routes" section with a "Destination" field containing `10.11.12.0/24` and a "Via" field containing `192.168.168.1`, with a "Submit" button. The "IPv4 Auto-Assign" section has a checked "Auto-Assign from Range" option and a table with two tabs: "Easy" (selected) and "Advanced". The table lists various IP ranges for auto-assignment.

Easy		Advanced	
10.147.17*	10.147.18*	10.147.19*	10.147.20*
10.144.**	10.241.**	10.242.**	10.243.**
10.244.**	172.22.**	172.23.**	172.24.**
172.25.**	172.26.**	172.27.**	172.28.**
172.29.**	172.30.**	192.168.191.*	192.168.192.*
192.168.193.*	192.168.194.*	192.168.195.*	192.168.196.*

Creation of a ZeroTier Network

Manage the members

- Assign names
- Customize the IPs

my.zerotier.com

Members

Switch to beta list

Search (Address / Name)

Display Filter

- Authorized
- Not Authorized
- Bridges
- Inactive 2
- Active 2
- Hidden 0

Sort By

- Address
- Name

< 1-4 / 4 >

	Address	Name/Description	Managed IPs	Last Seen	Version	Physical IP
	65035509e0 16:e6:e1:59:9d:c7	(short-name) (description)	172.25.0.2 + 172.25.0.x	1 MINUTE	1.10.3	93.70.123.99
	a542b17534 16:26:a0:bd:e1:13	(short-name) (description)	172.25.0.1 + 172.25.0.x	1 MINUTE	1.10.3	79.9.94.233
	c0825d87cb 16:43:60:51:13:ec	(short-name) (description)	+ 172.25.0.x	3 DAYS	1.12.0	5.91.6.229
	f927c7234f 16:7a:c5:cb:b7:68	(short-name) (description)	172.25.154.165 + 172.25.0.x	ABOUT 6 HOURS	1.12.2	79.9.94.233

< 1-4 / 4 >

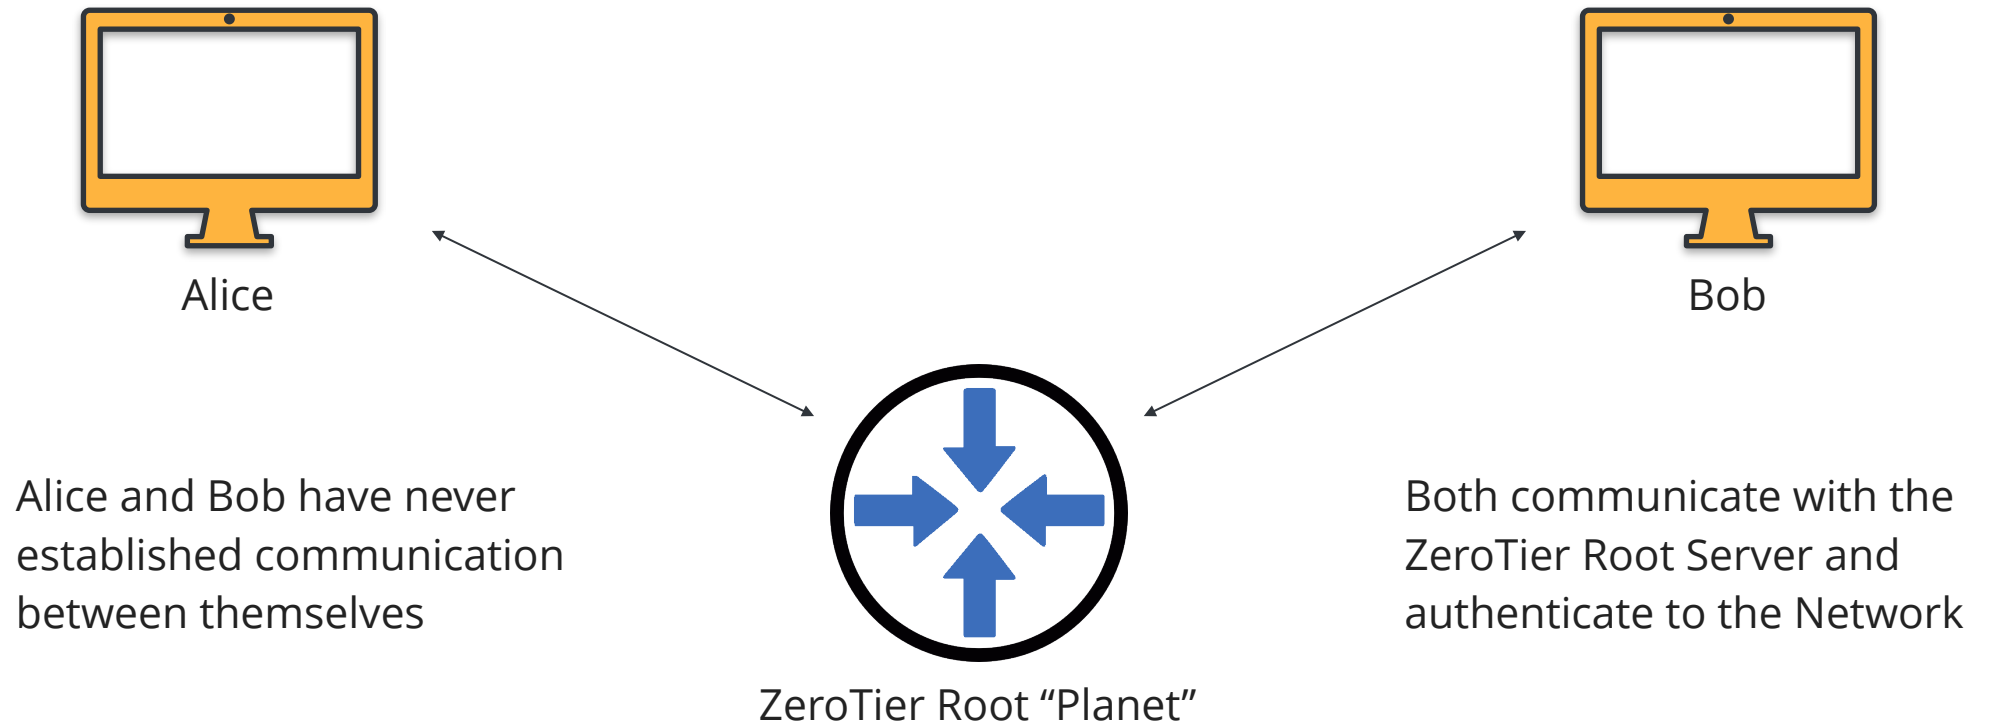


Module 3

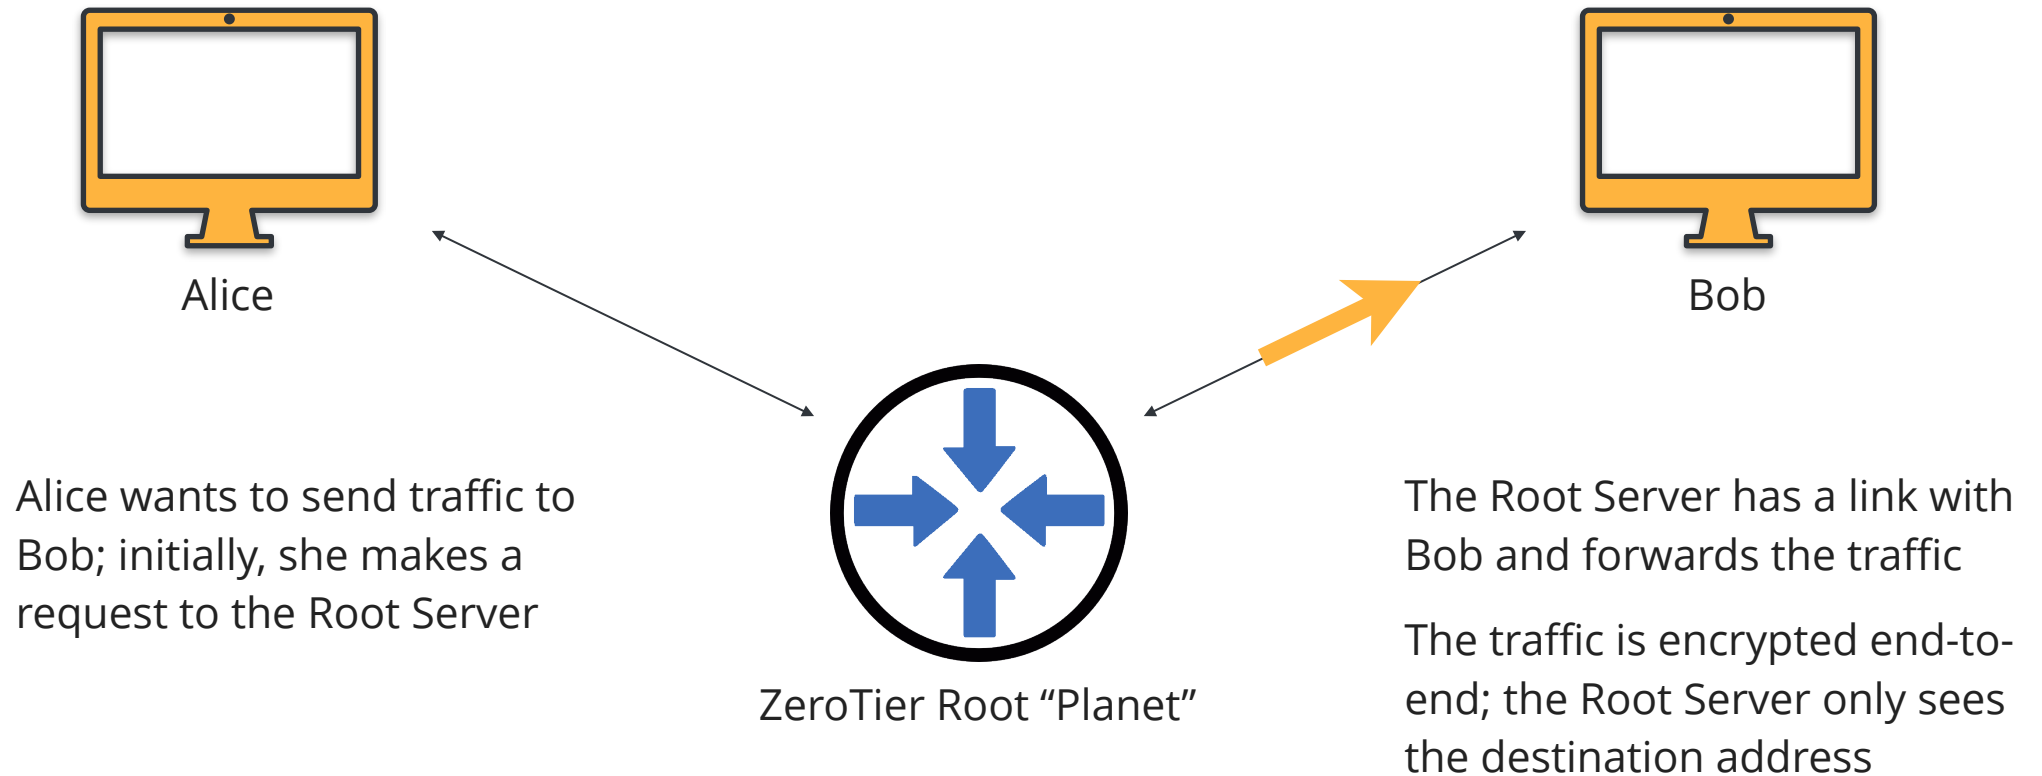
Connection!



VL1 Peer to Peer Link



VL1 Peer to Peer Link



VL1 Peer to Peer Link




The Root Server understands that Alice and Bob want to communicate and sends them information on how to reach each other (blue arrow)

ZeroTier Root "Planet"

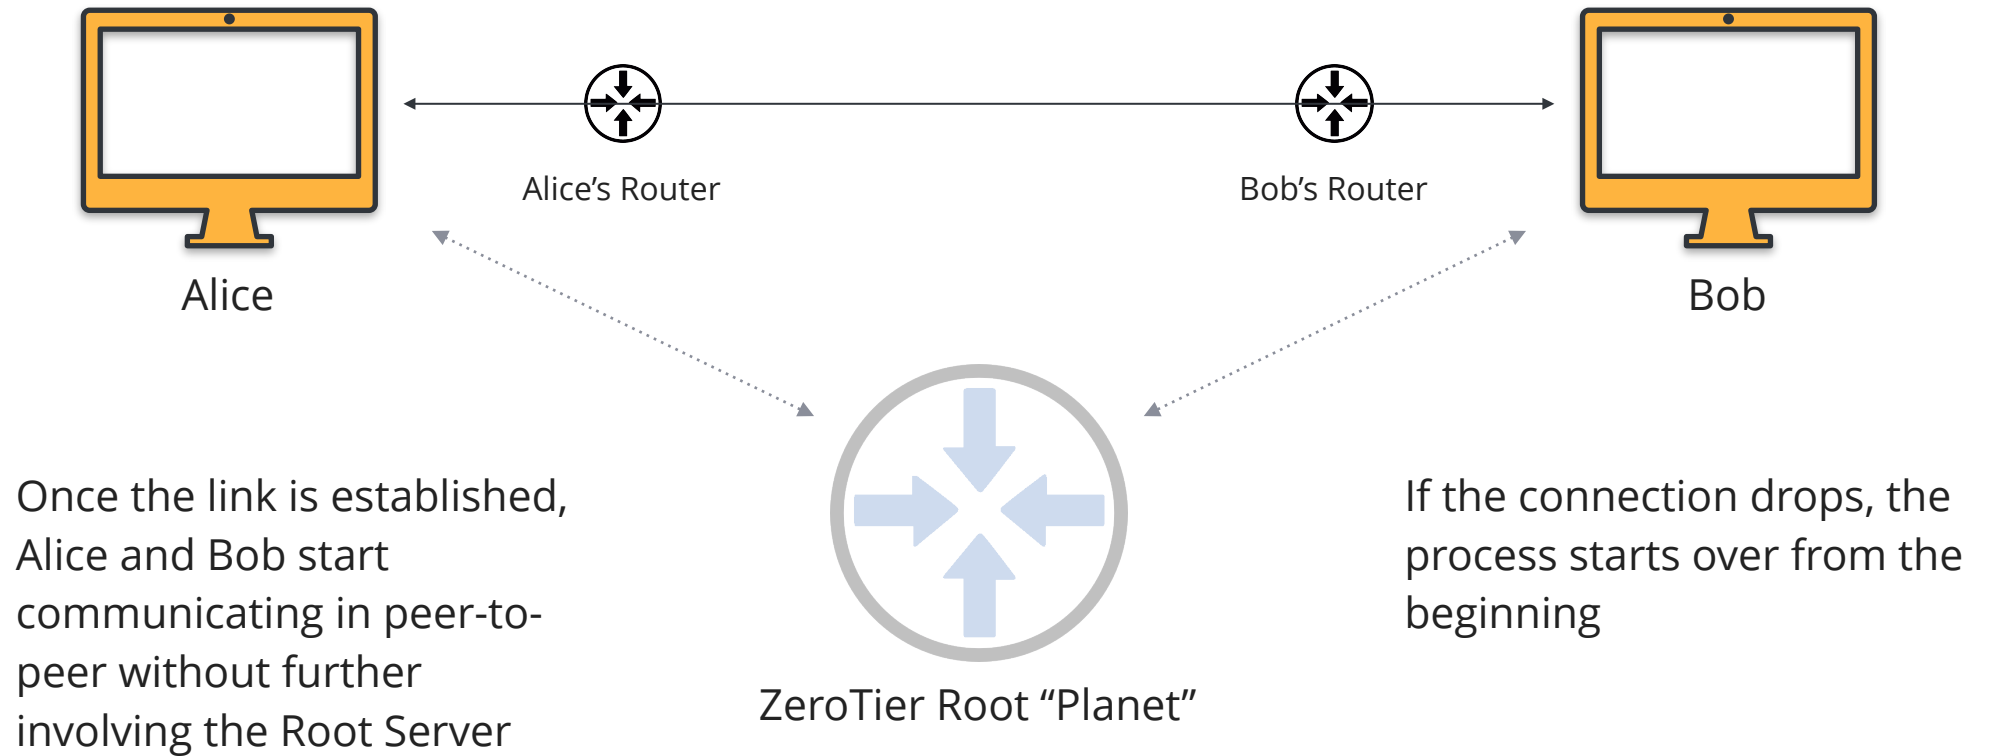
Alice and Bob use the transit offered by the Root Server to send their own information: IP and presence of NAT (yellow arrow)

VL1 Peer to Peer Link



Alice and Bob attempt to establish a peer-to-peer connection using various connection strategies, including UDP hole punching

VL1 Peer to Peer Link





Module 4

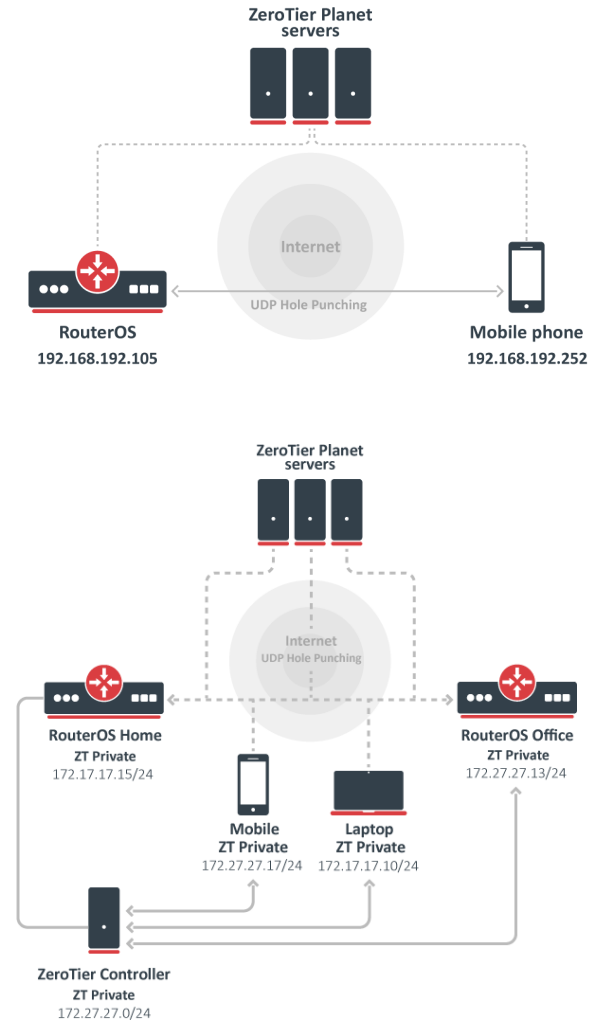
ZeroTier and MikroTik



ZeroTier and MikroTik

zerotier.npk

- Installable in RouterOS starting from version 7.1rc2
- Requires ARM/ARM64 architecture
- Available in both **client** and **controller** mode



ZeroTier and MikroTik

Client Configuration

- Enable ZeroTier Instance
 - Default port 9993
- Router connects to Root Server (PLANET)
- Create the ZeroTier interface specifying the Network ID
- Check firewall

The screenshot shows the MikroTik WinBox interface. The 'ZeroTier' window displays a table of instances:

Name	Port	State
R_zt1	9993	running

The 'New Interface' dialog is open, showing configuration for a ZeroTier interface named 'zerotier'. The 'General' tab is active, and the 'Instance' dropdown is set to 'zt1'. The 'Allow Managed' checkbox is checked.

The terminal window shows the following output:

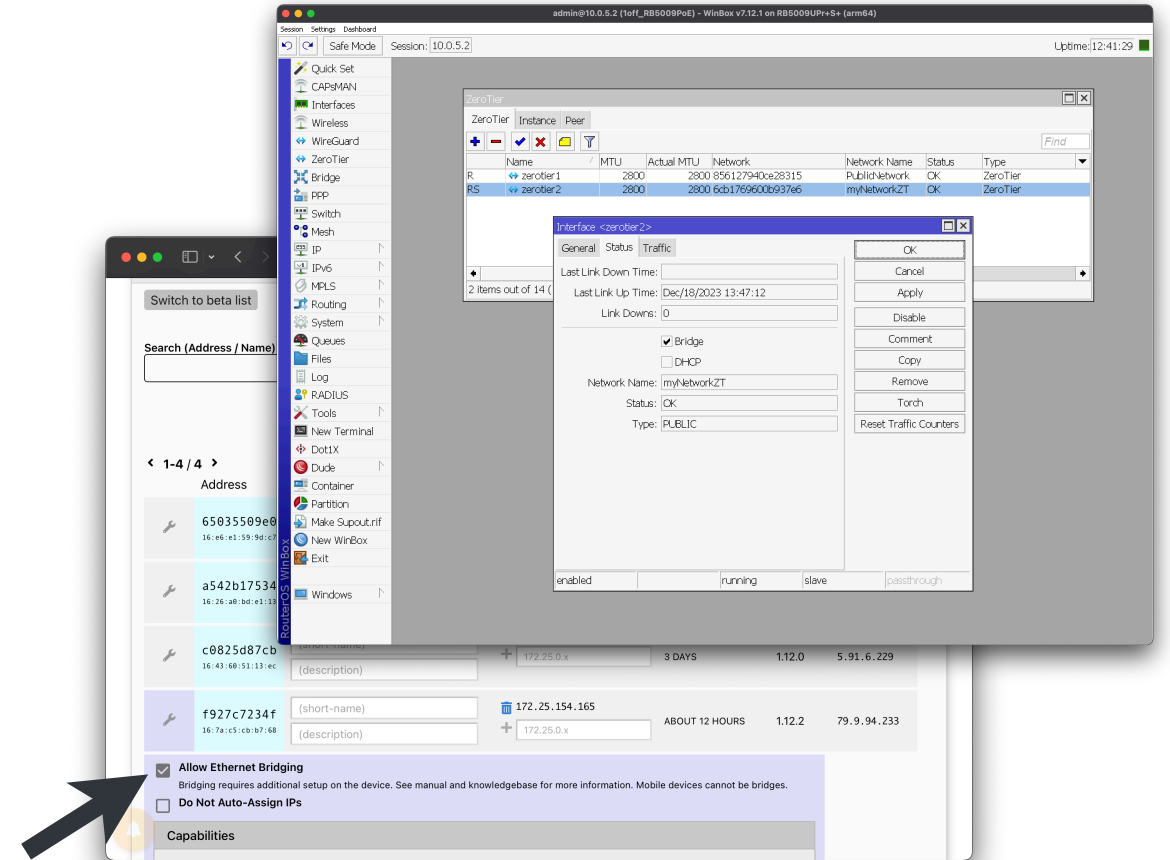
```
[admin@10ff_RB5009PoE] > zerotier/peer/
[admin@10ff_RB5009PoE] /zerotier/peer> pr
Columns: INSTANCE, ZT-ADDRESS, LATENCY, ROLE
# INSTANCE ZT-ADDRESS LATENCY ROLE
0 zt1 62f865ae71 180ms PLANET
1 zt1 778cde7190 123ms PLANET
2 zt1 cafe04eba9 21ms PLANET
3 zt1 cafe9efeb9 173ms PLANET
4 zt1 856127940c 138ms LEAF
5 zt1 65035509e0 35ms LEAF
6 zt1 6cb1769600 LEAF
[admin@10ff_RB5009PoE] /zerotier/peer>
```

```
/ip firewall filter add action=accept chain=forward in-interface=zerotier1 place-before=0
/ip firewall filter add action=accept chain=input in-interface=zerotier1 place-before=0
```

ZeroTier and MikroTik

ZeroTier L2 bridging

- We can enable the Ethernet Bridging function for each of the Members
- In RouterOS, the flag for the additional functionality appears immediately
- Let's bridge the ZeroTier interface





Module 5

ZeroTier Flow Rules





ZeroTier Flow Rules

Flow rules, not just firewall

- Traffic within the ZeroTier Network can be observed and controlled with a system of globally applied rules.
- The rules are applied to both ends of the communication.
- To bypass a rule, an attacker would need to compromise the entire connection system.
- The system is stateless, meaning it does not implement connection tracking.
- What makes this rules engine unique is the ability to set **Capabilities** and **Tags** and to analyze traffic remotely (**tee**).

ZeroTier Flow Rules

Firewall

- By default, all traffic is discarded.
- The "**accept**" action allows the traffic to pass.
- The "**drop**" action cannot be overridden by Capabilities.
- The "**break**" action can be overridden by Capabilities.

```
# Whitelist only IPv4 (/ARP) and IPv6 traffic and allow only ZeroTier-assigned IP addresses

drop          # drop cannot be overridden by capabilities
not ethertype ipv4 # frame is not ipv4
and not ethertype arp # AND is not ARP
and not ethertype ipv6 # AND is not ipv6
or not chr ipauth # OR IP addresses are not authenticated (1.2.0+ only!)
;

# Allow SSH, HTTP, and HTTPS by allowing all TCP packets (including SYN!/ACK) to these
ports

accept
dport 22 or dport 80 or dport 443
and ipprotocol tcp
;

# Drop TCP SYN,!ACK packets (new connections) not explicitly whitelisted above

break        # break can be overridden by a capability
chr tcp_syn  # TCP SYN (TCP flags will never match non-TCP packets)
and not chr tcp_ack # AND not TCP ACK
;

# Accept other packets
accept;
```

ZeroTier Flow Rules

Capabilities

- Roles can be created that entirely or partially override the "**break**" commands.
- It is necessary to specify a unique numerical ID.
- The Capability can be set from the web interface.

```
# Create a capability called "superuser" that lets its holders override all but the initial "drop"

cap superuser
  id 1000 # arbitrary, but must be unique
  accept; # allow with no match conditions means allow anything and everything
;
```

The screenshot shows the ZeroTier web interface for managing capabilities. At the top, there are filters for Authorized, Not Authorized, Bridges, Inactive, Active, Hidden, Address, and Name. Below the filters is a table of capabilities:

Address	Name/Description	Managed IPs	Last Seen	Version	Physical IP
65035509e0 <small>16:e6:e1:59:9d:c7</small>	(short-name) (description)	172.25.0.2 + 172.25.0.x	1 MINUTE	1.10.3	93.70.123.99
a542b17534 <small>16:26:a0:b0:e1:13</small>	(short-name) (description)	172.25.0.1 + 172.25.0.x	LESS THAN A MINUTE	1.10.3	79.9.94.233
c0825d87cb <small>16:43:60:51:13:ec</small>	(short-name) (description)	+ 172.25.0.x	3 DAYS	1.12.0	5.91.6.229
f927c7234f <small>16:7a:c5:cb:b7:68</small>	(short-name) (description)	172.25.154.165 + 172.25.0.x	ABOUT 13 HOURS	1.12.2	79.9.94.233

Below the table, there are checkboxes for "Allow Ethernet Bridging" (checked), "Do Not Auto-Assign IPs" (unchecked), and a "Capabilities" section with a checkbox for "superuser" (unchecked). A black arrow points to the "superuser" checkbox.

ZeroTier Flow Rules

Tag

- Tags can be created for micro-segmentation.
- We can create custom rules for those who possess a specific tag.
- Tag operations allow for limiting communication between Members.

```
# Create a tag for which department someone is in
tag department
  id 1000          # arbitrary, but must be unique
  enum 100 sales  # has no meaning to filter, but used in UI to offer a selection
  enum 200 engineering
  enum 300 support
  enum 400 manufacturing
  default 0
;

# Allow Windows CIFS and netbios between computers in the same department using a
tag
accept
  dport 139 or dport 445
  and ipprotocol tcp
  and tdiff department 0 # difference between department tags is 0, meaning they match
;

# Create a tag for user and routers
tag devtype
  id 1000          # arbitrary, but must be unique
  enum 100 user    # has no meaning to filter, but used in UI to offer a selection
  enum 200 router
  default 100
;
break
  tdiff devtype 0 # obtain client isolation
;
```

ZeroTier Flow Rules

Tee

- Traffic can be duplicated entirely or in part and sent to an analyzer.
- The recipient must be a member of the Network.

```
# Send a copy of EVERY packet on both sender and receiver side to ZeroTier address
"deadbeef11".
tee -1 deadbeef11;

# Send only first 128 bytes of every inbound TCP SYN, RST or FIN packet to deadbeef11
tee 128 deadbeef11
  chr inbound
  and chr tcp_syn or chr tcp_rst or chr tcp_fin
;

# Send only first 128 bytes of every outbound TCP SYN, RST or FIN packet to deadbeef22
tee 128 deadbeef22
  not chr inbound
  and chr tcp_syn or chr tcp_rst or chr tcp_fin
;
```



Module 6

MikroTik ZT Controller





MikroTik ZeroTier Controller

No user limit, but no flow rules

- RouterOS implements ZeroTier Controller functions.
- We can create and manage a Network directly from the router without having to enter the ZeroTier portal.
- The Network Controller is not a Root Server!
 - The Root Server is used to build connections at the VL1 layer.
 - The Network Controller manages the Certificate Authority to which the VL2 layer belongs.
- The controller is responsible for authorizing and customizing the parameters of the Members.
- In RouterOS, the roles of Controller and Member can coexist.

Network Controller in RouterOS

Let's create the Network Controller (CLI only):

- Add a Controller

```
/zerotier/controller/add name=ZT-private instance=zt1 ip-range=172.27.27.10-172.27.27.20  
private=yes routes=172.27.27.0/24
```

- Add an interface to the network we created

```
/zerotier/interface/add network=879c0b5265a99e4b name=myZeroTier instance=zt1
```

- Authorize new member

```
/zerotier/controller/member/set 0 authorized=yes
```

- Check received parameters (IP address / IP route)



Modulo 7

ZeroTier Docker Controller





MikroTik Docker Container

ZeroTier Controller in a MikroTik Docker Container

- RouterOS supports Docker Container functions starting from RouterOS v7.4beta4.
- The Container.npk package is only compatible with **arm**, **arm64**, and **x86** architectures.
- It is possible to virtualize ZeroTier Controller, ZeroTier Client, and ZeroTier Root Server (Moon) as containers.
- The use of an external disk is highly recommended.

Docker Container in RouterOS

Let's configure the router to host a Container:

- After installing the package, enable the container mode

```
/system/device-mode/update container=yes
```

- Create a virtual interface and configure its IP address

```
/interface/veth/add name=veth1 address=172.17.0.2/24 gateway=172.17.0.1  
/ip/address/add address=172.17.0.1/24 interface=veth1
```

- Set an external library

```
/container/config/set registry-url=https://registry-1.docker.io tmpdir=usb1/pull
```

Docker Container in RouterOS

Install the ZeroTier Controller Docker:

- Set up a mount point to customize the configuration files for Networks and Members

```
/container/mounts/add dst=/var/lib/zerotier-one/controller.d/network/ name=nets src=/usb1/ZT/network
```

- Install Container

```
/container/add interface=veth1 logging=yes mounts=networks root-dir=usb1/ZeroTier workdir= remote-image=namestars/ztncai-arm64:latest
```

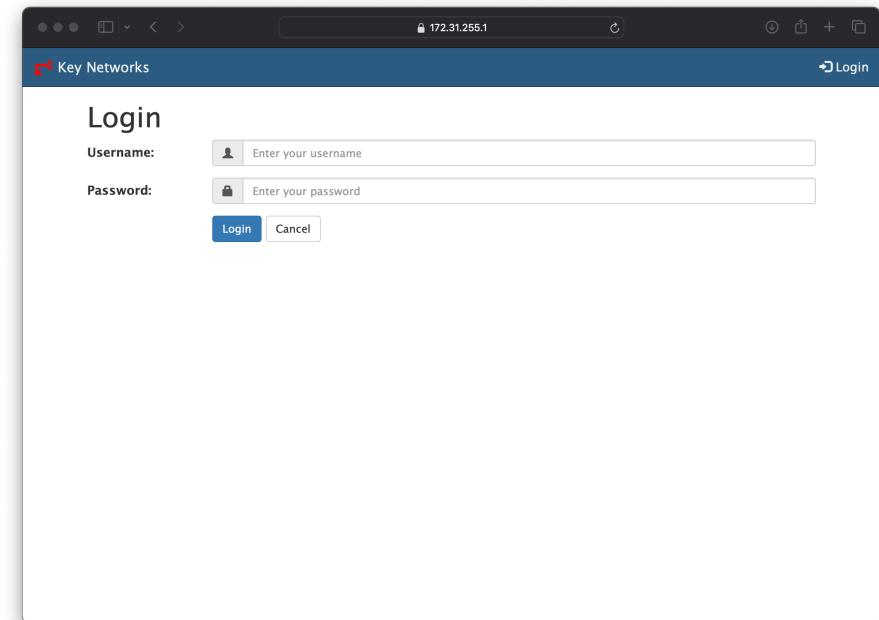
- The password is available in /var/log/docker-ztncai.log

```
/container/shell 0  
cat /var/log/docker-ztncai.log
```

Docker Container in RouterOS

Access via web interface:

- Default port for https: 3443
- Login: admin
- Set password
- Create a Network
- Connect Members





Docker Container in RouterOS

Configure ZeroTier Container on X86:

- Let's execute the steps previously discussed to prepare a CHR to host Containers
- Install the Docker Container zerotier/zerotier:latest
- Connect Docker to our Controller

```
/container/shell 0  
zerotier-cli peers  
zerotier-cli join [networkid]  
zerotier-cli listnetworks
```

- We can now route traffic through the CHR. (VETH are only L3)



Module 8

Missing in RouterOS





Missing in RouterOS

it would be perfect if:

- it were possible to install the ZeroTier client on other architectures and not only on arm.
- the Flow Rules were implemented for the controller configured on RouterOS.
- it were possible to use ZeroTier's native load balancing features.
- the VETHs were Layer 2 or it were possible to pass more than one Virtual Ethernet to a Docker.

Toff.it Technical Presentation

Thank you Q&A

Alessandro Campanella

