



**OpenStack con fabric VXLAN**  
SDN

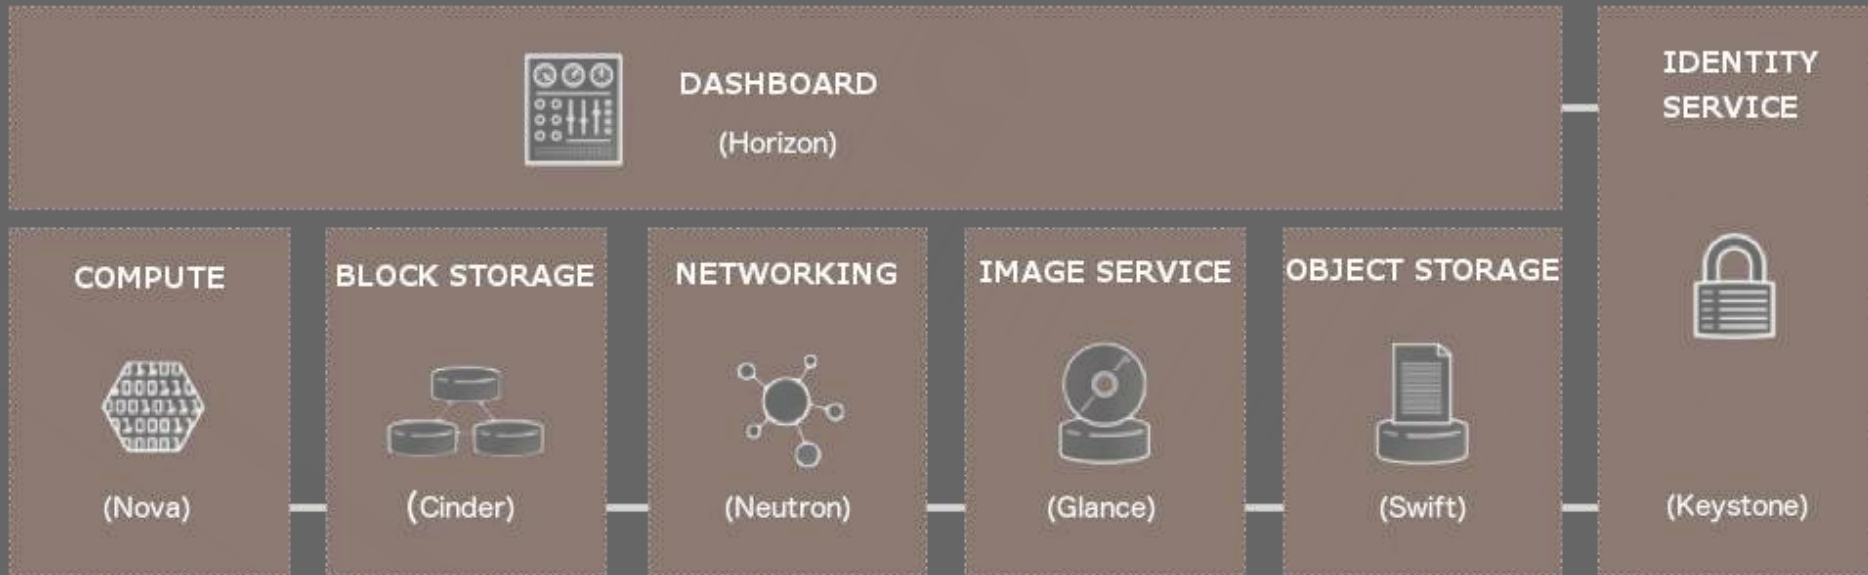
# Índice

- OpenStack
- VXLAN
- Demo integración OpenStack + VXLAN
  - Creación tenants
  - Creación redes
  - Creación routers
  - Creación VMs
  - Pruebas comunicación y revisión del underlay



# Introducción a OpenStack

# OpenStack – Componentes



## Versiones

- Austin – Octubre 2010
- Bexar – Febrero 2011
- Cactus – Abril 2011
- ...
- Havana – Octubre 2013 ← Cambios importantes en el módulo de red
- Icehouse – Abril 2014
- Juno – Octubre 2014
- Kilo – Abril 2015
- **Liberty – Octubre 2015**
- Mitaka – Abril 2016

# Ejemplo interfaz

The screenshot displays the OpenStack dashboard interface. At the top left is the OpenStack logo. The top navigation bar shows the current project 'gore' and the user 'admin'. A left-hand sidebar contains a menu with categories like Project, Compute, Network, Admin, and Identity, with sub-items such as Overview, Instances, Volumes, Images, Access & Security, and Identity. The main content area is titled 'Overview' and features a 'Limit Summary' section. This section contains seven circular gauges representing resource usage: Instances (Used 4 of 10), VCPUs (Used 4 of 20), RAM (Used 256 of 51,200), Floating IPs (Used 0 of 50), Security Groups (Used 1 of 10), Volumes (Used 0 of 10), and Volume Storage (Used 0 of 1,000). Below the gauges is a 'Usage Summary' section with the text 'Select a period of time to query its usage:'.

openstack

gore

admin

## Overview

### Limit Summary

Resource	Used	Limit
Instances	4	10
VCPUs	4	20
RAM	256	51,200
Floating IPs	0	50
Security Groups	1	10
Volumes	0	10
Volume Storage	0	1,000

### Usage Summary

Select a period of time to query its usage:

# Ejemplo interfaz

The screenshot shows the OpenStack dashboard interface. The top navigation bar includes the OpenStack logo, a dropdown menu for 'gore', and a user profile for 'admin'. The left sidebar contains navigation links for Project, Compute, Overview, Instances (highlighted), Volumes, Images, Access & Security, Network, Admin, and Identity. The main content area is titled 'Instances' and features a search bar with 'Instance Name' and 'Filter' options. Below the search bar are buttons for 'Launch Instance', 'Terminate Instances', and 'More Actions'. The main table lists five instances:

	Instance Name	Image Name	IP Address	Size	Key Pair	Status	Availability Zone	Task	Power State	Time since created	Actions
<input type="checkbox"/>	vm04	cirros-0.3.4-x86_64-uec	192.168.1.3	m1.nano	-	Active	nova	None	Running	6 minutes	Create Snapshot
<input type="checkbox"/>	vm03	cirros-0.3.4-x86_64-uec	172.16.202.4	m1.nano	-	Active	nova	None	Running	7 minutes	Create Snapshot
<input type="checkbox"/>	vm02	cirros-0.3.4-x86_64-uec	net01 172.16.201.4 net02 172.16.202.3	m1.nano	-	Active	nova	None	Running	7 minutes	Create Snapshot
<input type="checkbox"/>	vm01	cirros-0.3.4-x86_64-uec	172.16.201.3	m1.nano	-	Active	nova	None	Running	1 day	Create Snapshot



The image features a magnifying glass with a black handle and a clear lens. The lens is positioned over a background of binary code (0s and 1s) that is slightly blurred. The magnifying glass's frame is in sharp focus, and the binary code within the lens is also sharp, while the code outside the lens is out of focus. The overall composition suggests a focus on digital security or data analysis.

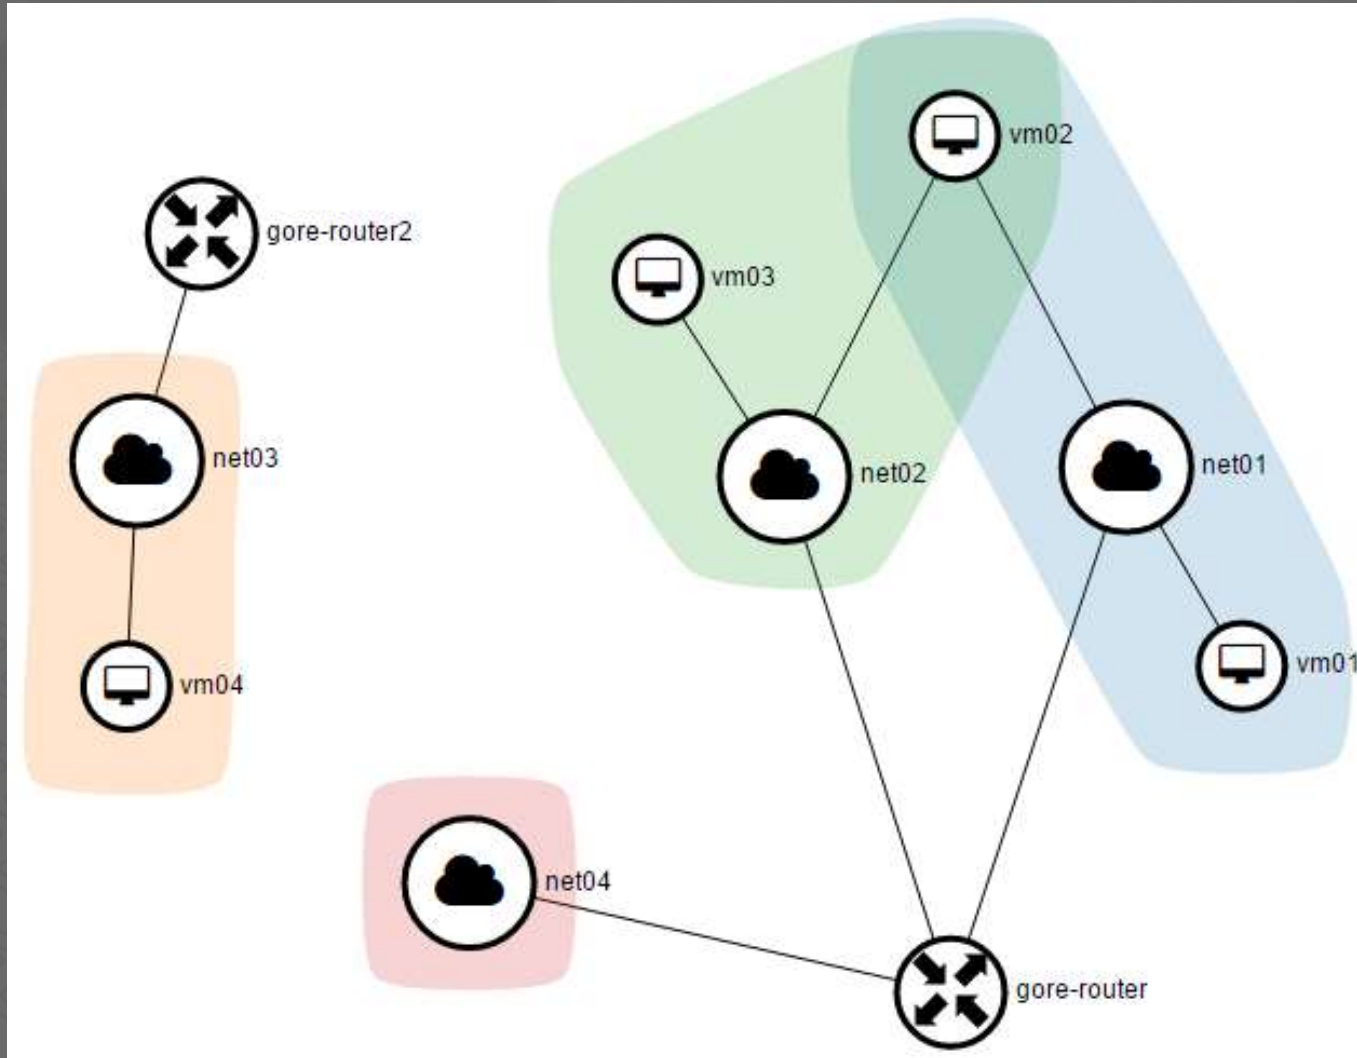
# Módulo Neutron y ML2 plugins



## Cambios en Havana (Octubre 2013)

- Deprecados los módulos monolíticos:
  - Openvswitch
  - Linuxbridge
- Aparece Modular Layer 2 plugins (ML2)
  - Distintos *type drivers*
    - Local
    - Flat
    - VLAN
    - GRE
    - VXLAN
  - Facilidad de integración de distintas tecnologías (*mechanism driver*)
    - ALE Omniswitch
    - Arista
    - OpenHardware (Cumulus, PicOS,...)
    - Brocade
    - Cisco
    - Lenovo
    - Huawei

# Ejemplo de Network Topology

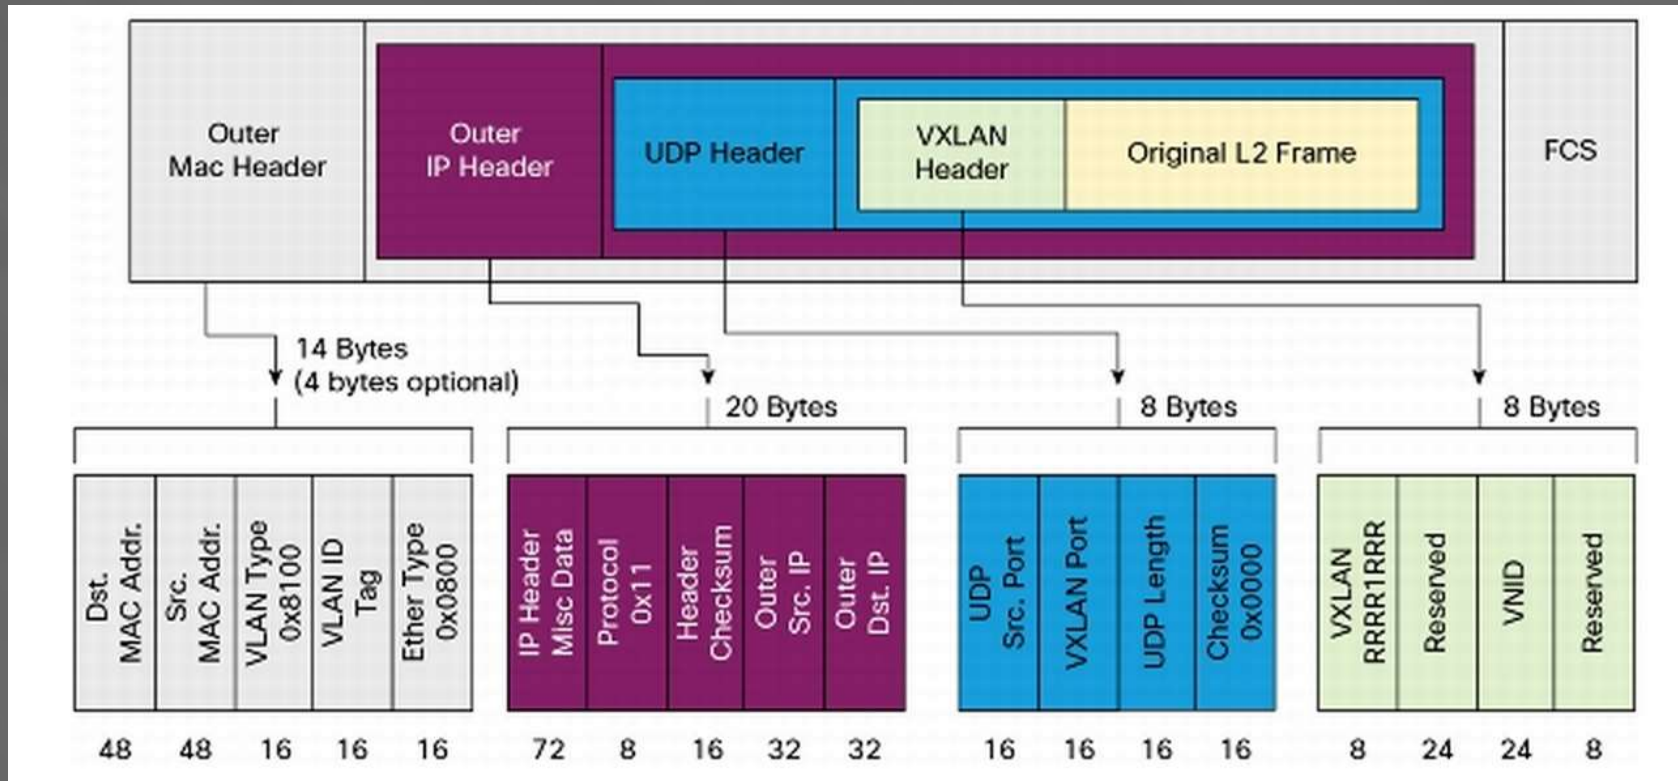


# Introducción a VXLAN

## VXLAN – Protocolo estándar redes overlay

- Protocolo de transporte para redes virtuales  
Extiende un nivel 2 sobre un *fabric IP*
- Encapsulado UDP (50 bytes adicionales)
- Funciona tanto por multicast como unicast


# VXLAN – Encapsulado UDP






# VXLAN – Protocolo estándar redes overlay

Permite encapsular trafico de nivel 2 sobre un *fabric IP*



## VXLAN – VTEP: Virtual Tunnel End Point


- Punto de entrada/salida de una red virtual
- Componente software o hardware
- Necesario conocer qué hay detrás de cada VTEP, 4 métodos
  - HER: Head-End Replication, necesario configurar manualmente todos los VTEP
  - Push: Los VTEP son actualizados mediante un controlador SDN
  - Poll: Los VTEP actualizan una base de datos con toda la información (OVSDB)
  - Dinámico: Mediante autodiscovery y uso de multicast



Middle of Row




Layer 2 / MLAG




Layer 3 / ECMP


L2 over Layer 3  
VXLAN




Servers Middle of Row Servers



Servers



Servers



Servers

Server Scale: 100 to 2,000

100 to 10,000

100 to 100,000+

100 to 100,000+

## VXLAN – La realidad

- Evitar IP Multicast
  - IP Multicast es un protocolo eficiente para la gestión del tráfico BUM
  - Pero nadie quiere multicast en la red
- VTEP: Hardware vs Software
  - Tráfico Norte-Sur, es necesario salir del mundo virtual!
  - Equipos físicos en redes VXLAN (almacenamiento, servidores no virtualizados, etc.)
  - Rendimiento de los VTEP software no siempre es suficiente
- CUIDADO CON LA MTU!!! Casi nadie fragmenta y no hay gestión de frag. needed

## VXLAN – Decisiones clave

- VTEP: Software VS Hardware
  - Flexibilidad VS Rendimiento
- Replicación de nodos VS replicación head-end distribuida
- Controlador SDN externo VS Neutron standalone
  - Funcionalidad VS Coste





# Laboratorio

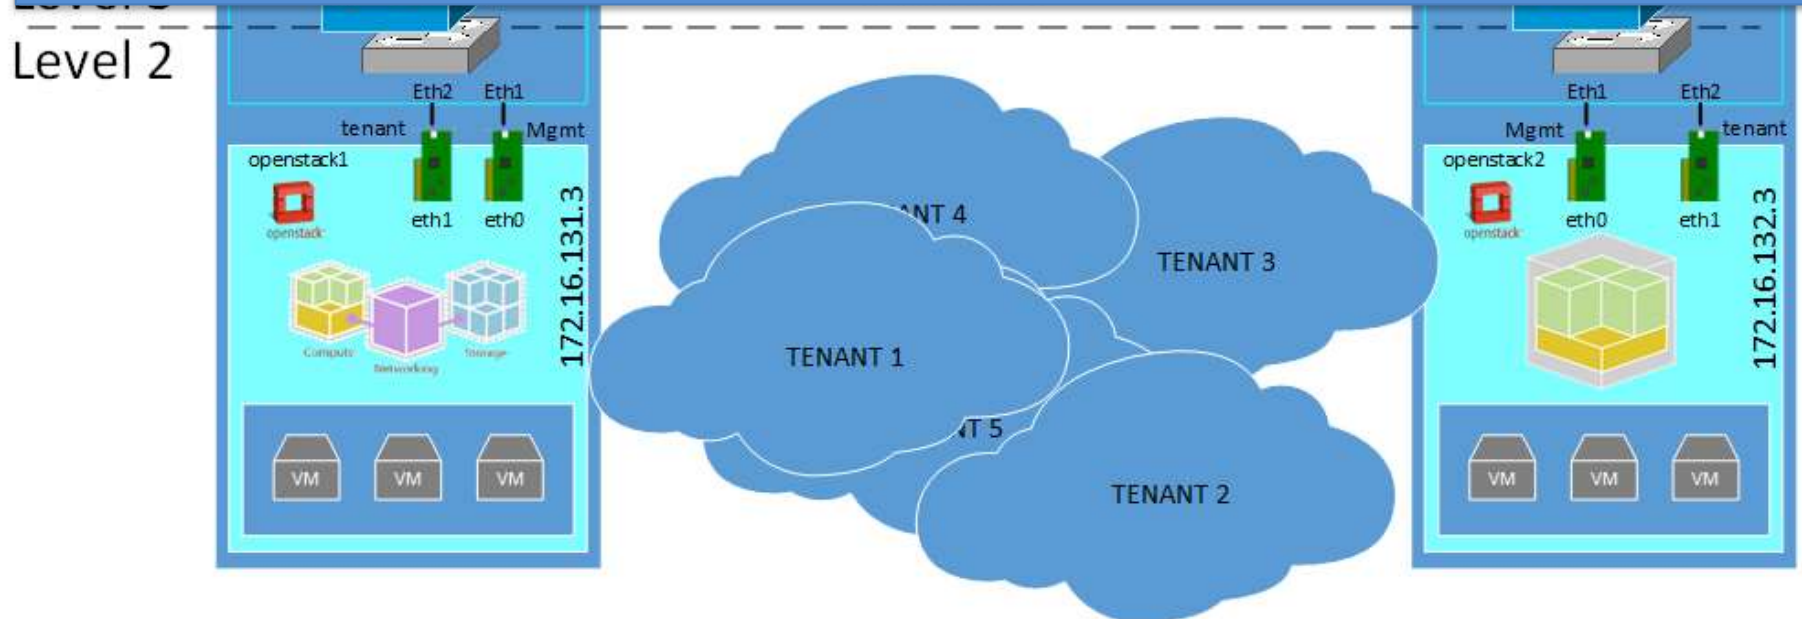


# Laboratorio

# Dispositivos

- CORE de red, interfaces:
  - 172.16.130.1/24
  - 172.16.130.1/24
  - 172.16.130.1/24
- 2 x Switch físico, switch ToR con funcionalidad VTEP
  - sw-openstack1: 172.16.131.2/24, Lo1: 10.0.0.1
  - sw-openstack2: 172.16.133.2/24, Lo1: 10.0.0.2
- 1 x Switch / VXLAN Router
  - sw-openstackR: 172.16.130.2/24, Lo1: 10.0.0.3
- 2 x Nodos OpenStack, uno en cada Rack:
  - Openstack1: 172.16.131.3
  - Openstack2: 172.16.132.3

# Extensión redes L2 mediante VXLAN Enrutado mediante VXLAN routing en hardware





DEMO





¿Preguntas?



**open3s**  
Open Source & Security Services

[www.open3s.com](http://www.open3s.com)

## Barcelona

Carrer dels Almogàvers, 107-115, Bajos , 08018

E: [info@open3s.com](mailto:info@open3s.com)

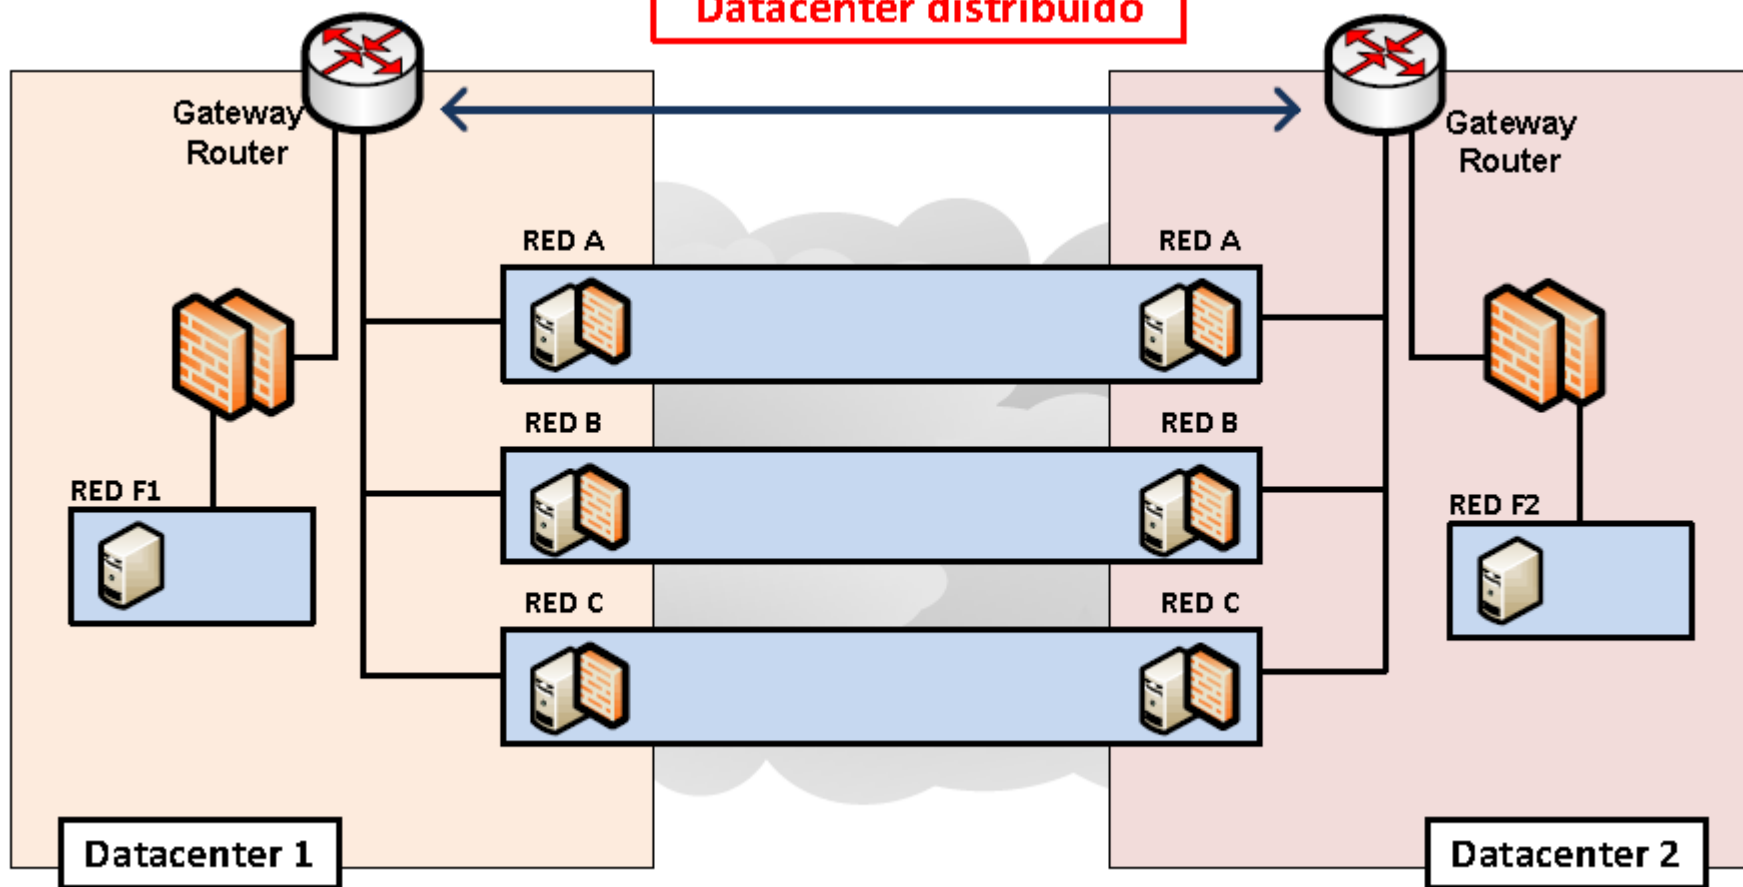
T: 93 268 73 20

F: 93 268 73 20



# Extensión VLANs

**Datacenter distribuido**



**Datacenter distribuido**



# Arquitectura – Diferentes diseños de red

