Vol. 08, No. 2 (2026) 1347-1358, doi: 10.24874/PES08.02B.023

Proceedings on Engineering Sciences

www.pesjournal.net

PRACTICAL ASPECTS OF ROUTING WHEN BUILDING BUSINESS CONTINUITY AND DISASTER RECOVERY AT THE ENTERPRISE LEVEL

Andrian Piskozub
Mikolaj Karpinski[1]
Elena Nyemkova
Yuriy Lakh
Roman Banakh
Pawel Sawicki

Received 18.11.2025.
Revised 17.03.2026.
Accepted 02.04.2026.

Keywords:

BCP, DR, RTO, Routing, Flat networking, Clustering, MC-LAG, Juniper, VPN

A B S T R A C T

Disaster recovery and business continuity are critical components of a comprehensive risk management strategy for organizations. Minimizing downtime is one of the key aspects for a company that wants to quickly recover its IT infrastructure, applications and data after a disaster. The emphasis in this paper is placed on improving network resilience and availability of critical services, possible risks and implementation of several Recovery Time Objective improvements. Practical aspects of routing when building business continuity and disaster recovery at the enterprise level have been discussed. Mitigation steps to minimize the impact of disasters based on improving infrastructure resilience and recovery steps, focused on repairing damaged infrastructure, re-establishing business functions, are considered.

© 2026 Published by Faculty of Engineering

1. INTRODUCTION

Security has evolved from being a technical necessity to a strategic priority. In today's digital age, where data is central to nearly every aspect of our lives and businesses, ensuring that systems, data, and communications are secure is paramount. Security breaches can disrupt business operations, causing delays, downtime, or even permanent data loss. Effective security measures ensure that critical business functions remain operational.

2. PROBLEM STATEMENT

The paper describes practical experience in implementing an enterprise network infrastructure with BCP elements and considers options for improving the recovery time objective (RTO).

Business Continuity Planning (BCP) is closely connected with security. Security is essential to ensuring that critical business functions remain operational during disruptions.

[1] Corresponding author: Mikolaj Karpinski
Email: mikolaj.karpinski@uken.krakow.pl

Disaster Recovery (DR) is a subset of BCP specifically related to IT and data systems and, in our case, we will consider ensuring network resilience based on network recovery possibilities (Watters, 2013).

Many industries have strict compliance requirements related to both business continuity and security (e.g., GDPR (General Data Protection Regulation, 2018), HIPAA (Health Insurance Portability and Accountability Act, 2024), PCI-DSS (Payment Card Industry Data Security Standard, 2024)). Business continuity and security are integral components in SOC 2 (System and Organization Controls, 2024), ISO 27001 (ISO/IEC 27001, 2022), though the focus and structure can vary between the frameworks.

These regulations, standards and frameworks require organizations to ensure their data is protected, continuously available, and resilient to disruptions, with security being a foundational element of their business continuity strategies.

The main task of this paper is the deployment of an enterprise network infrastructure with improved resilience capabilities.

This task will be achieved based on the elements of Disaster Recovery Plan (DRP), aimed at overcoming and reducing risks from equipment and network failures, minimizing the impact of disasters by reducing the RTO.

The improvement of RTO is based on fast recovery of Intranet/Internet service connections in each company location as well as fast recovery of company critical services Intranet/Internet service connections. Hence, such optimized network infrastructure will improve RTO by reducing the time to recover critical systems and services and provide quick recovery of services in the company.

3. DEPLOYMENT

Network infrastructure according to the described above BCP requirements and compliance with the above-mentioned regulations, standards and frameworks has 2 data centres (DCs) to provide disaster recovery (DR) where the second DC2 site can take over in the event of a primary DC1 site failure (see Figure 1).

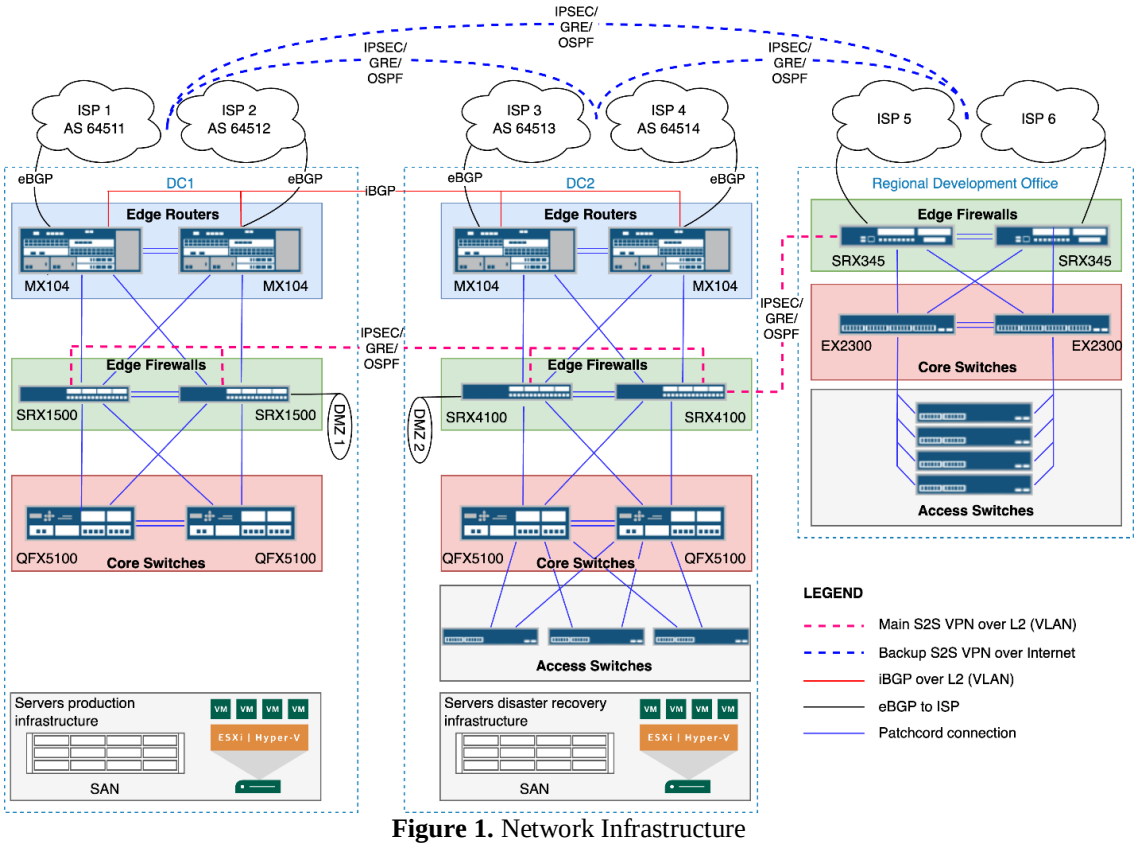


Figure 1. Network Infrastructure

The distance between DCs is above 500 km. These DCs are directly interconnected by two L2 (VLAN) channels, configured in LACP mode.

There are headquarters and several regional development offices (for simplicity, only one typical development office is shown in Figure 1), regional offices have to be connected with both DCs to get access to highly available services. All interconnections between any locations – DCs, development offices use site-to-site VPN (IPSEC) for security compliance.

Concerning disaster recovery strategy due to such a long distance between DCs, the preference was given to so-called warm site (active-passive) strategy, even though hot site (active-active) strategy has better BCP indicators - fast recovery time and minimal downtime (RTO). Warm sites have lower costs than hot sites, but here we have a worse recovery time indicator. Hence, one of our paper tasks is to improve this RTO indicator by using some routing solution, which will be described below in the paper.

Continuing to do our best in improving the RTO indicator, we have to mention the so-called flat networking network architecture where all the hosts in different locations – DCs and regional offices may have the same IP addressing.

Usually, to provide flat networking between several locations, solutions such as Layer 2 VPN (MPLS, EVPN, VPLS, VXLAN, etc.) are used. Unfortunately, in our case, due to some technical limitations, we could not use these Layer 2 VPN solutions. It means that every DC, every regional office in our case, uses its own internal IP range. The drawback of such situation is a rather high RTO for critical company services, and our task is to reduce it. The benefit is that any DC or any office network in the other location is independent in case of a disaster at another site and can continue its operation.

As Figure 1 illustrates, to satisfy high RTO requirements all the network components are duplicated – every location has several connections to independent Internet service providers; routers, firewalls and switches are configured as clusters to provide high availability network services (Hanks et al., 2016), (Chassis Cluster User Guide for SRX Series Devices, 2023).

3.1 Data centres

Both DCs have identical infrastructure which consists of two Juniper MX104 routers, two Juniper SRX firewalls (SRX1500 in DC1, SRX4100 in DC2) and a cluster of Juniper QFX5100 switches. We will not describe here in details SAN network and servers' structure since this is out of this paper scope, but we only state the usage of virtual machines infrastructure.

Instead of using Layer 2 VPN for connection between locations, we will use Layer 3 VPN (IPSEC). Since we will use dynamic routing for our tasks with RTO improvement, the OSPF protocol is used for internal routing. To provide OSPF routing specifically through site-to-site VPN between locations, we use IPSEC/GRE/OSPF (Varnum, 2014).

Two MX routers in every DC are connected in cluster solution using MC-LAG (Multi-Chassis Link Aggregation). The benefits of MC-LAG are as follows (Hanks et al., 2016):

- MC-LAG provides us with both – link-level and node-level redundancy;
- for enterprise environments, MC-LAG is a great method for multiple core routers to provide a single, logical IEEE 802.3ad interface to downstream firewalls and switches that establish IEEE 802.3ad across two physically separate MX chassis;
- MC-LAG maintains a separate control plane for each chassis that participates in MC-LAG.

We use two autonomous systems (AS) and every MX router operates with both AS having independent eBGP connections for every AS with upstream Internet service providers (ISPs) (see Figure 1):

```
ap@DC2-Router-0> show bgp summary
Threading mode: BGP I/O
Default eBGP mode: advertise - accept, receive - accept
Groups: 5 Peers: 9 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      13          1          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.255.6.202          65001     138971     139184       0       3       11:12 Establ
  inet.0: 0/4/4/0
10.255.6.203          65001        111        108       0       4       48:03 Establ
  inet.0: 0/4/4/0
10.255.6.204          65001        124        123       0       4       53:15 Establ
  inet.0: 0/4/4/0
10.255.6.222          65002     138961     787024       0       3       11:12 Establ
  VR.inet.0: 0/4/4/0
10.255.6.223          65002        109       4199       0       4       47:55 Establ
  VR.inet.0: 0/4/4/0
10.255.6.224          65002        120       4258       0       4       53:13 Establ
  VR.inet.0: 0/4/4/0
1.1.1.1               64517    1522731     223663       0       0       10:58 Establ
  VR.inet.0: 22503/22504/22503/0
2.2.2.2               64513     132567     139167       0       3       11:09 Establ
  VR.inet.0: 1/1/1/0
3.3.3.3               64513          8          7       0       6        2:16 Establ
  inet.0: 1/1/1/0
```

On every MX router in order to separate routing information between two AS, the first eBGP session, using the first AS, is done from the main instance of the MX router, the second eBGP session, which uses the second AS, is done from the virtual-router instance of the MX router. To exchange routing information between MX routers in both DCs, the corresponding iBGP sessions are set – iBGP sessions between the main instance of every MX router, and iBGP sessions between the virtual-router instance of every MX router respectively.

Respectively, SRX firewalls in every DC use these routes from both AS and redistribute the necessary part of them into the Intranet using OSPF. Since MX routers in every DC, connected in MC-LAG, use VRRP for both - main instance and virtual-router instance, this simplifies the configuration of respective SRX clusters and provides business continuity for each data center network with access to the Internet. An important note is that each of the 4 MX routers in both data centers announces all available public networks. This allows hosts in the company network to operate even if at least one router remains fully operational with its BGP uplink. For separating routing information between two AS, SRX clusters, directly connected to respective MX router clusters using reth interface connection, also use the same idea – they also use the main instance and the virtual-router instance. Each SRX cluster instance uses as default route the corresponding VRRP IP address of the main instance or virtual-router instance of upstream MX routers respectively.

3.2 Regional offices

Regional offices have also high availability infrastructure which consists of two Juniper SRX345 firewalls and a cluster of Juniper EX2300 switches in every regional office.

Every regional location has its own DHCP, DNS and AD services, so in case there is trouble with DC1 and/or DC2 connectivity, there will be no problems with network connectivity in a regional office.

The peculiarity of regional offices – they have 2 independent ISP connections and L2 (VLAN) connection with site-to-site IPSEC VPN over it towards DC2 (see Figure 1).

In regional offices with 2 independent ISP connections, we use failover scenario – switching from main ISP to backup one in the case of Internet connectivity problems.

Since we use IPSEC/GRE/OPSF solution for interconnection between DCs and regional offices, we can simultaneously have two working sets of VPN connections between DCs and regional offices for the sake of RTO improvement. In Figure 1, we may see the main S2S VPN over L2 (VLAN) between DCs and between regional development offices and DC2, also we may see a backup S2S VPN over the Internet between locations. In case of problems with L2 (VLAN) channel between DCs or between regional development offices and DC2, due to the backup alternatives and OSPF configuration all the traffic will be rerouted to DC2 or DC1 resources through the Internet backup connection. Here is an example of IPSEC/GRE/OPSF configuration excerpt of the SRX cluster from a typical regional office towards DC2. We will skip the security zones and policies description since our task is about BCP and RTO improvement for Intranet and Internet connectivity.

Configure Phase 1 – IKE

```
set security ike proposal DC2_L2 authentication-
method pre-shared-keys
set security ike proposal DC2_L2 dh-group
group14
set security ike proposal DC2_L2 authentication-
algorithm sha-256
set security ike proposal DC2_L2 encryption-
algorithm aes-256-cbc
set security ike proposal DC2_L2 lifetime-
seconds 28800
set security ike proposal DC2 authentication-
method pre-shared-keys
set security ike proposal DC2 dh-group group14
set security ike proposal DC2 authentication-
algorithm sha-256
set security ike proposal DC2 encryption-
algorithm aes-256-cbc
set security ike proposal DC2 lifetime-seconds
28800
set security ike policy DC2_L2 mode main
set security ike policy DC2_L2 proposals DC2_L2
set security ike policy DC2_L2 pre-shared-key
ascii-text "xxxxxxxxx"
set security ike policy DC2 mode main
set security ike policy DC2 proposals DC2
set security ike policy DC2 pre-shared-key
ascii-text "xxxxxxxxx"
```

where DC2 – connection to DC2 through Internet;
DC2_L2 – connection to DC2 through direct optical link.

Configure Phase 1 – IKE Gateways

```
set security ike gateway DC2_L2 ike-policy
DC2_L2
set security ike gateway DC2_L2 address
10.255.4.27
set security ike gateway DC2_L2 local-identity
inet 10.255.4.26
set security ike gateway DC2_L2 external-
interface reth1.2591
set security ike gateway DC2_L2 version v2-only
set security ike gateway DC2 ike-policy DC2
set security ike gateway DC2 address 4.4.4.4
set security ike gateway DC2 local-identity inet
5.5.5.5
set security ike gateway DC2 external-interface
reth1.2592
set security ike gateway DC2 version v2-only
```

Configure Phase 2 – IPSec

```
set security ipsec proposal DC2_L2 protocol esp
set security ipsec proposal DC2_L2
authentication-algorithm hmac-sha-256-128
set security ipsec proposal DC2_L2 encryption-
algorithm aes-256-cbc
set security ipsec proposal DC2_L2 lifetime-
seconds 3600
set security ipsec proposal DC2 protocol esp
```

```
set security ipsec proposal DC2 authentication-algorithm hmac-sha-256-128
set security ipsec proposal DC2 encryption-algorithm aes-256-cbc
set security ipsec proposal DC2 lifetime-seconds 3600
set security ipsec policy DC2_L2 perfect-forward-secrecy keys group14
set security ipsec policy DC2_L2 proposals DC2_L2
set security ipsec policy DC2 perfect-forward-secrecy keys group14
set security ipsec policy DC2 proposals DC2
```

Configure Phase 2 – IPSec Peers

```
set security ipsec vpn DC2_L2 bind-interface st0.2
set security ipsec vpn DC2_L2 vpn-monitor optimized
set security ipsec vpn DC2_L2 vpn-monitor source-interface lo0.0
set security ipsec vpn DC2_L2 vpn-monitor destination-ip 10.255.6.1
set security ipsec vpn DC2_L2 ike gateway DC2_L2
set security ipsec vpn DC2_L2 ike proxy-identity local 10.255.6.26/32
set security ipsec vpn DC2_L2 ike proxy-identity remote 10.255.6.1/32
set security ipsec vpn DC2_L2 ike proxy-identity service any
set security ipsec vpn DC2_L2 ike ipsec-policy DC2_L2
set security ipsec vpn DC2_L2 establish-tunnels immediately
set security ipsec vpn DC2 bind-interface st0.0
set security ipsec vpn DC2 vpn-monitor optimized
set security ipsec vpn DC2 vpn-monitor source-interface lo0.0
set security ipsec vpn DC2 vpn-monitor destination-ip 10.255.6.1
set security ipsec vpn DC2 ike gateway DC2
set security ipsec vpn DC2 ike proxy-identity local 10.255.6.26/32
set security ipsec vpn DC2 ike proxy-identity remote 10.255.6.1/32
set security ipsec vpn DC2 ike proxy-identity service any
set security ipsec vpn DC2 ike ipsec-policy DC2
set security ipsec vpn DC2 establish-tunnels immediately
set interfaces st0 unit 0 description DC2
set interfaces st0 unit 0 family inet mtu 1460
set interfaces st0 unit 2 description DC2_L2
set interfaces st0 unit 2 family inet mtu 1460
set security zones security-zone VPN interfaces st0.0
set security zones security-zone VPN interfaces st0.2
```

Configure the Loopback (Used as the tunnel source)

```
set interfaces lo0 unit 0 family inet address 10.255.6.26/32
set security zones security-zone VPN interfaces lo0.0
```

Configure the Virtual Tunnel Interface

```
set security zones security-zone VPN interfaces gr-0/0/0.0
set interfaces gr-0/0/0 unit 0 description DC2
set interfaces gr-0/0/0 unit 0 tunnel source 10.255.6.26
set interfaces gr-0/0/0 unit 0 tunnel destination 10.255.6.1
set interfaces gr-0/0/0 unit 0 family inet mtu 1400
set interfaces gr-0/0/0 unit 0 family inet address 10.255.5.98/29
```

Configure OSPF

```
set protocols ospf area 0.0.0.0 interface gr-0/0/0.0 interface-type p2p
set protocols ospf area 0.0.0.0 interface gr-0/0/0.0 metric 30
set protocols ospf area 0.0.0.0 interface gr-0/0/0.0 retransmit-interval 3
set protocols ospf area 0.0.0.0 interface gr-0/0/0.0 hello-interval 2
set protocols ospf area 0.0.0.0 interface gr-0/0/0.0 dead-interval 8
set protocols ospf area 0.0.0.0 interface lo0.0 passive
```

Configure Routing

```
set routing-options static route 10.255.6.1/32 next-hop st0.2
set routing-options static route 10.255.6.1/32 qualified-next-hop st0.0 preference 8
```

Here you may see the general idea of switching from DC2_L2 – connection to DC2 through direct optic link (st0.2) to DC2 – connection to DC2 through Internet (st0.0) in case of optic link problem.

Verify:

```
ap@reg1-firewall> show security ike security-associations
Index   State  Initiator cookie  Responder cookie
Mode          Remote Address
90983   UP    2082972b513f52c2  b7d10f00794ed8bc
IKEv2           4.4.4.4
90984   UP    3db33e420f907baa  dbef62792071e212
IKEv2           10.255.4.27

ap@reg1-firewall> show ospf neighbor
Address            Interface              State
ID               Pri  Dead
10.255.5.97        gr-0/0/0.0              Full
10.255.6.1       128     7
```

Thus, in total, every regional office has 3 S2S VPN – one main over L2 (VLAN) channel, the second backup – with DC2 over the Internet and the third backup - with DC1 over the Internet. Thanks to OSPF metric parameter selection, we can choose the best path in terms of throughput to the final destination - as shown in Figure 1, the best path to DC1 is through DC2 using direct optical link to the latter.

Internet connectivity failover in every regional office is provided by SRX firewalls using this configuration:

```
set routing-options static route 0.0.0.0/0 next-hop 6.6.6.6
set routing-options static route 0.0.0.0/0 qualified-next-hop 7.7.7.7 preference 7
```

We also implemented the idea of providing Internet connectivity for a typical regional office through DC2 interconnections with the Internet. It is useful in case both local ISPs are inaccessible, and in case we need to provide Internet access for our employees from other AS

IP addresses that the local ISPs suggest. The routing result on a typical regional SRX cluster looks like this:

```
ap@reg1-firewall> show route 0/0 exact
inet.0: 704 destinations, 712 routes (704
active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 15w1d 06:53:15
                    >  to 6.6.6.6 via reth1.2592
                    [Static/7] 15w1d 06:53:27
                    >  to 7.7.7.7 via ge-5/0/0.0
                    [OSPF/150]  5w0d  16:08:13,
metric 20, tag 0
                    >  via gr-0/0/0.0

reg1_VR.inet.0: 699 destinations, 699 routes
(699 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0           *[OSPF/150] 5w0d 16:08:13,
metric 25, tag 0
                    >  via gr-0/0/0.0
```

The idea of providing Internet connectivity for regional office through DC2 is realized on the virtual router base.

Routing instance is a powerful tool of the JunOS operating system that allows to perform any task when used together with rib-groups, firewall filters and policy.

Among such routine-instance options as Forwarding, Nonforwarding, Virtual router (VR), VRF, VPLS, Layer 2 VPN and others (Juniper Networks. Routing Instances Overview, 2024) we chose for our task the Virtual router instance due to its possibility to support any routing protocols, GRE, IPSEC tunnels, NAT and DHCP. One of the purposes of this routing instance is to isolate the client interface(s) and its routing table. The routing instance type we are considering installs routes into a separate routing table, as in our example above – in the `reg1_VR.inet.0` table. This allows us to isolate individual segments from accessing the Internet through local ISPs, and send traffic exclusively through DC2, which, in its turn, announces the default route.

We also use the rib-groups functionality for exchanging routes and connecting VRs with each other and with the router's main instance.

Here is an example of virtual router configuration excerpt of SRX cluster from typical regional office with optimized routing to Internet.

We create the `reg1_VR` virtual-router instance with the necessary parameters' settings (routing protocols, interfaces that belong to this VR):

```
set routing-instances reg1_VR instance-type
virtual-router
set routing-instances reg1_VR routing-options
interface-routes rib-group inet reg1_router
set routing-instances reg1_VR routing-options
router-id 10.255.6.126
set routing-instances reg1_VR routing-options
instance-import from_main_into_reg1
set routing-instances reg1_VR protocols ospf
area 0.0.0.4 interface lo0.1 passive
set routing-instances reg1_VR protocols ospf
area 0.0.0.4 interface reth0.436 passive
set routing-instances reg1_VR protocols ospf
export 2ospf_reg1
set routing-instances reg1_VR interface
reth0.436
set routing-instances reg1_VR interface lo0.1
set interfaces lo0 unit 1 family inet address
10.255.6.126/32
set security zones security-zone vpn2 interfaces
lo0.1
```

Then we define the necessary rib-groups:

```
set routing-options rib-groups reg1_router
export-rib reg1_VR.inet.0
set routing-options rib-groups reg1_router
import-rib inet.0
set routing-options rib-groups reg1_router
import-rib reg1_VR.inet.0
set routing-options instance-import
from_reg1_into_main
```

rib-groups allow us to import routing tables from our `reg1_router` VR to the main instance router and vice versa. Not completely, but just the parts we need. This is achieved due to a powerful `policy-statement` mechanism which allows to effectively match the necessary conditions (Juniper Networks. Policy-statement Overview, 2024). The `policy-statement from_reg1_into_main` describes what we export from `reg1_router` VR to the main instance router (OSPF routes, internal IP address ranges, directly connected networks):

```
set policy-options policy-statement from_reg1_into_main term 2 from instance reg1_VR
set policy-options policy-statement from_reg1_into_main term 2 from protocol ospf
set policy-options policy-statement from_reg1_into_main term 2 then metric 25
set policy-options policy-statement from_reg1_into_main term 2 then accept
set policy-options policy-statement from_reg1_into_main term 7 from instance reg1_VR
set policy-options policy-statement from_reg1_into_main term 7 from route-filter 10.0.0.0/8 orlonger
set policy-options policy-statement from_reg1_into_main term 7 from route-filter 172.16.0.0/12
orlonger
set policy-options policy-statement from_reg1_into_main term 7 then accept
set policy-options policy-statement from_reg1_into_main term 10 from instance reg1_VR
set policy-options policy-statement from_reg1_into_main term 10 from protocol direct
set policy-options policy-statement from_reg1_into_main term 10 then accept
set policy-options policy-statement from_reg1_into_main term 199 then reject
```

Respectively here is `policy-statement from_main_into_reg1` that describes what we export from main instance router to `reg1_router` VR (OSPF routes, internal IP addresses ranges, directly connected networks, static routes, but except 0.0.0.0/0 route):

```
set policy-options policy-statement from_main_into_reg1 term 2 from instance master
set policy-options policy-statement from_main_into_reg1 term 2 from protocol ospf
set policy-options policy-statement from_main_into_reg1 term 2 then metric 25
set policy-options policy-statement from_main_into_reg1 term 2 then accept
set policy-options policy-statement from_main_into_reg1 term 4 from instance master
set policy-options policy-statement from_main_into_reg1 term 4 from protocol static
set policy-options policy-statement from_main_into_reg1 term 4 from route-filter 0.0.0.0/0 exact
set policy-options policy-statement from_main_into_reg1 term 4 then reject
set policy-options policy-statement from_main_into_reg1 term 7 from instance master
set policy-options policy-statement from_main_into_reg1 term 7 from route-filter 10.0.0.0/8 orlonger
set policy-options policy-statement from_main_into_reg1 term 7 from route-filter 172.16.0.0/12
orlonger
set policy-options policy-statement from_main_into_reg1 term 7 then accept
set policy-options policy-statement from_main_into_reg1 term 10 from instance master
set policy-options policy-statement from_main_into_reg1 term 10 from protocol direct
set policy-options policy-statement from_main_into_reg1 term 10 then accept
set policy-options policy-statement from_main_into_reg1 term 16 from instance master
set policy-options policy-statement from_main_into_reg1 term 16 from protocol static
set policy-options policy-statement from_main_into_reg1 term 16 then accept
set policy-options policy-statement from_main_into_reg1 term 199 then reject
```

Thus, on the one hand, we isolated Internet access for individual network segments, on the other hand, due to rib-groups for route exchange and communication between VR and the main instance router we provide full connectivity with all networks in the Intranet.

4. RISK ASSESSMENT

A Disaster Recovery Plan must include a Risk Assessment to identify and evaluate potential threats and hazards that could disrupt the organization's operations, resources, and information. It should consider internal and external factors such as natural disasters, workplace accidents, etc. Since the scope of our paper is network infrastructure, we will analyze here only risks connected with network and connectivity.

The first group of possible risks consists of risks, connected with the operations of routers, firewalls and switches clusters in the case of their degradation either due to cluster node malfunction on every level or due to the cable connectivity problems. We tested all these possible problems: MX, SRX, QFX, EX node degradations, as well as link-level degradations – there were no problems with network connectivity at all. Due to this link-level and node-level redundancy we fit in our RTO – all the problems were solved in less than several seconds.

The second group of possible risks consists of risks, connected with external factors like problems with upstream ISPs, problems with channels to ISPs or between our locations (see Figure 1). The results of BCP tests are as follows:

- The simulation of failure of eBGP sessions to upstream providers in DCs showed that if at least one eBGP upstream session to one of four routers in DCs stays operational, then all company network infrastructure is operational. Due to BGP protocol nature the rebuilding of routing information between routers takes up to 1 minute.
- The simulation of failure of one of two L2 (VLAN) channels (configured in LACP mode) between two DCs – there is no impact at all, since it is LACP.
- The simulation of failure of both two L2 (VLAN) channels (configured in LACP mode) between two DCs:
  - there is little to no risk of DCs interconnectivity problems since rerouting through backup S2S VPN over Internet between DCs occurs in 8 seconds;
  - there is a risk of iBGP malfunction between both DCs MX routers since they all announce the IP address ranges of both DCs. The solution is provided to ensure that network infrastructure will continue to operate. The idea is to make changes in BGP announcements so that every DC will announce only its own IP address ranges and thus the problems with iBGP sessions between MX routers will not have any influence on this situation. Such architecture provides BGP disaster recovery solution for both DCs. Due to BGP protocol nature the rebuilding on routing information on routers takes up to 1 minute after changes in BGP announcements.
- The failure of ISP channels in regional offices – switching to backup ISP channel as soon as link on the first ISP channel goes down. Besides, as it is shown above, there is an option to reach Internet through DC2 due to OSPF announcement.
- The failure of L2 (VLAN) channel between regional development office and DC2 - there is little to no risk with connectivity problems to

DCs since rerouting through backup S2S VPN over Internet between regional development offices and DCs occurs in 8 seconds.

5. FLAT NETWORKING DEPLOYMENT

From the point of view of RTO improvement using flat networking means that we have the same IP address ranges for internal (in Intranet) and external addressing (in DMZ) respectively in several locations. In case we use virtual machine (VM) infrastructure with flat networking, e.g., ESXi or Hyper-V, VMs can be quickly moved to different physical hypervisors, in different locations, reducing downtime. The critical services restoration scenario looks like this:

- critical services, situated in one DC, are replicated to the other DC;
- in the case of problem with any critical service in one DC its replica is powered on in the other DC with minimal RTO since there is no need to reconfigure TCP/IP parameters of this replica VM.

On the contrary, when we have no flat networking between DCs, even using replication of VMs between DCs, after VM is replicated, it has to be powered on and reconfigured according to IP addresses in that another location. Often it is not only about changing IP addresses, but also about changes in application services configurations and DNS records of this service. It means that at best the recovery of such critical service, based on replicated VM, takes at least one hour. If the company accepts such a scenario, then that's ok. Otherwise, we must find a solution to minimize critical service RTO.

As mentioned earlier, since we cannot use Layer 2 VPN solutions, in order to reduce RTO, we will provide flat networking scenario without L2 VPN. The idea is shown in Figure 2.

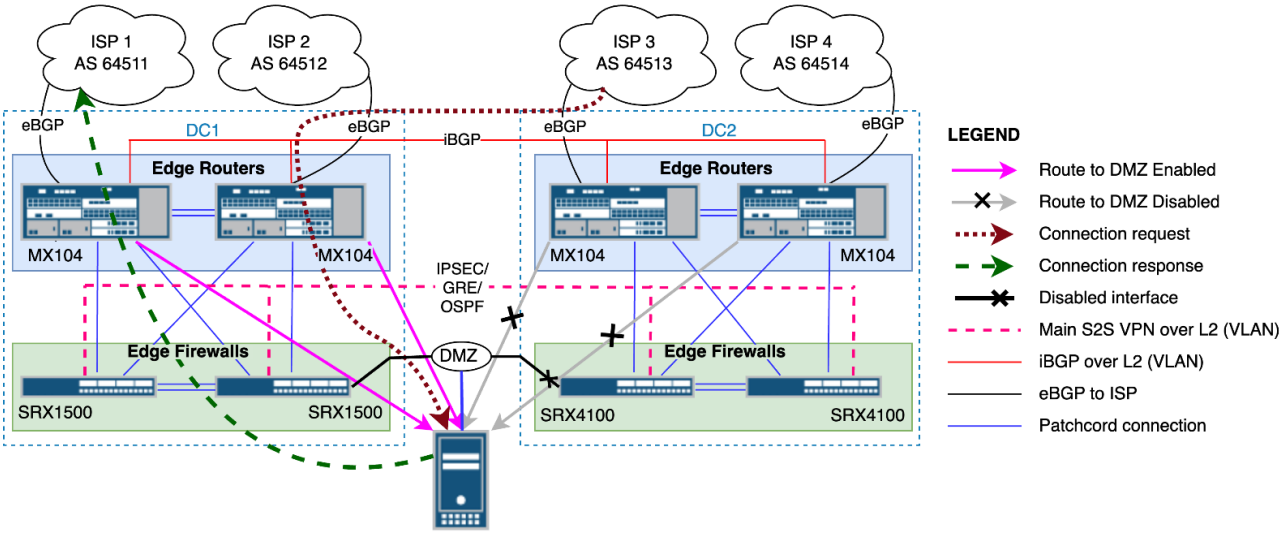


Figure 2. Flat networking diagram

Suppose we have some critical service, this service may be internal for the company, or maybe external, situated in DMZ. We need to extend those VLANs, in which these highly critical services are situated, to both DCs. We may extend as many VLANs with high critical services, as we need. To demonstrate the idea in Figure 2 we will extend the external DMZ VLAN to both DCs. Since extended VLAN is needed to ensure the same IP addressing for highly critical services between DCs, there is one common DMZ with critical service, connected to it, depicted in Figure 2.

Here are flat networking configuration excerpts of corresponding SRX and MX clusters in both DCs.

Let's suppose that DC1 is the primary DC and DC2 – respectively the backup DC (DR DC) (see Figure 2).
In our example we will use:

- VLAN 1080 for DMZ, VLAN 802 – for internal IP addressing and extend both VLANs to both DCs;
- Flat DMZ IP addressing is 9.9.9.128/26, flat internal segment – 10.224.0.0/24;
- DC1 SRX1500 edge firewall public IP address on external interface – 11.11.11.11;
- DC2 SRX4100 edge firewall public IP address on external interface – 12.12.12.12.

One of the elements of ensuring the operation of highly critical services in both the primary DC and the DR DC without changing TCP/IP parameters, namely without changing the default gateway parameter, is VRRP launching in the above-mentioned VLANs between the firewall clusters in both data centers.

One problem we came across here is asymmetric routing (see Figure 2) – when, for example, a request to a critical service comes from the one data center, and the response is sent to the default gateway, which is the VRRP master on firewall cluster of the other data center. Asymmetric

traffic flow in this case is a problem since we are working with next generation firewalls which, in their turn, want to control each session from the beginning to the end. Thus, we need to avoid asymmetric traffic flow.

Thanks to the main DC /DR DC approach, in each DC on the corresponding firewall cluster we configure interfaces of the critical services VLANs with VRRP support. After that, we disable the corresponding interface on the DR DC side, thereby leaving the only available interface - through the firewall on the main DC side – DC1 in our case (see Figure 2).

Checks to be done on DC1 SRX1500 edge firewall

For flat internal segment:

```
set interfaces reth0 unit 802 vlan-id 802
set interfaces reth0 unit 802 family inet address 10.224.0.3/24 vrrp-group 1 virtual-address 10.224.0.1
set interfaces reth0 unit 802 family inet address 10.224.0.3/24 vrrp-group 1 priority 200
set interfaces reth0 unit 802 family inet address 10.224.0.3/24 vrrp-group 1 accept-data
set protocols ospf area 0.0.0.0 interface reth0.802 passive
```

For flat DMZ:

```
set interfaces reth0 unit 1080 vlan-id 1080
set interfaces reth0 unit 1080 family inet address 9.9.9.130/26 vrrp-group 1 virtual-address 9.9.9.131
set interfaces reth0 unit 1080 family inet address 9.9.9.130/26 vrrp-group 1 priority 200
set interfaces reth0 unit 1080 family inet address 9.9.9.130/26 vrrp-group 1 accept-data
```

Checks to be done on MX routers DC1-Router-0, DC1-Router-1.

There should be present the following rule:

```
set routing-options static route 9.9.9.128/26 next-hop 11.11.11.11
```

The result to check is:

```
ap@DC1-Firewall> show route 10.224.0.0
inet.0: 1320 destinations, 1336 routes (1320 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.224.0.0/24      *[Direct/0] 1w6d 01:36:49
                    >  via reth0.802

ap@DC1-Firewall> show route 9.9.9.128
inet.0: 1320 destinations, 1336 routes (1320 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
9.9.9.128/26 *[Direct/0] 1w6d 01:39:06
                    >  via reth0.1080

ap@DC1-Firewall > show vrrp
Interface      State        Group    VR state VR Mode   Timer    Type   Address
reth0.802      up             1      master   Active A  0.156 lcl    10.224.0.3

vip    10.224.0.1
reth0.1080     up             1      master   Active A  0.066 lcl    9.9.9.130

vip    9.9.9.131
```

Checks to be done on DC2 SRX4100 edge firewall, considered in this scenario as member of backup DC

For flat internal segment:

```
set interfaces reth0 unit 802 disable
set interfaces reth0 unit 802 vlan-id 802
set interfaces reth0 unit 802 family inet address 10.224.0.2/24 vrrp-group 1 virtual-address 10.224.0.1
set interfaces reth0 unit 802 family inet address 10.224.0.2/24 vrrp-group 1 priority 100
set interfaces reth0 unit 802 family inet address 10.224.0.2/24 vrrp-group 1 accept-data
set protocols ospf area 0.0.0.0 interface reth0.802 passive
```

For flat DMZ:

```
set interfaces reth0 unit 1080 disable
set interfaces reth0 unit 1080 vlan-id 1080
set interfaces reth0 unit 1080 family inet address 9.9.9.129/26 vrrp-group 1 virtual-address 9.9.9.131
set interfaces reth0 unit 1080 family inet address 9.9.9.129/26 vrrp-group 1 priority 100
set interfaces reth0 unit 1080 family inet address 9.9.9.129/26 vrrp-group 1 accept-data
```

The result to check is:

```
ap@DC2-Firewall> show route 10.224.0.0
inet.0: 625 destinations, 674 routes (625 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.224.0.0/24        *[OSPF/10] 1w0d 10:44:49, metric 2
                     >  to 10.255.4.42 via reth0.804
```

This result means there is no direct connection to flat internal segment.

```
ap@DC2-Firewall > show route 9.9.9.128
inet.0: 625 destinations, 674 routes (625 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
0.0.0.0/0          *[Static/5] 1w6d 01:17:59
                    >  to 12.12.12.13 via reth1.557
```

This result means there is no direct connection to flat DMZ

```
ap@DC2-Firewall > show vrrp
VRRP is not running
```

Due to disabling reth0.802 and reth0.1080 there is no VRRP on DR DC SRX firewall.

Since we initially assumed that MX routers in both DCs work, announcing the corresponding IP address range of critical services network from both DCs simultaneously, we must disable static routes on MX routers to the downstream firewall cluster in the DC that is currently considered DR DC (see arrows named “Route to DMZ disabled” in Figure 2).

Checks to be done on MX routers DC2-Router-0, DC2-Router-1

There should NOT be present the following rule:

```
set routing-options static route 9.9.9.128/26
next-hop 12.12.12.12
```

In our example the request from Internet to the mentioned critical service will follow the arrow "Connection request" (see Figure 2) and the response will follow the arrow "Connection response" respectively. In Figure 2 we depicted incoming traffic to our critical service going through ISP3 in DC2 and the corresponding outgoing traffic going through ISP1 in DC1 which is normal situation from the point of view of BGP functionality. What is important to us is that both the connection request and the connection response must go through firewall cluster on main DC.

Due to the presence of critical services VMs replicas on the DR DC side, in the event of a main DC network failure we switch to the DR DC network. To do this, we need to swap the roles of the critical services VLANs interfaces on the firewalls – to enable the mentioned interfaces on the DR DC firewall and disable the corresponding interface on the main DC firewall. Thus, now the firewall in the DR DC has become the VRRP master.

In case of DC1 disaster scenario DR DC (DC2) becomes main and we need to swap the roles of MX routers and SRX firewalls in DCs - to enable the previously disabled static routes on DR DC MX routers to the downstream DR DC firewall cluster. And vice versa – to disable the previously enabled static routes on main DC MX routers (see arrows named "Route to DMZ enabled" in Figure 2) to the downstream main DC firewall cluster.

This approach of course has the disadvantage of having to run a script to reconfigure the corresponding network changes in main DC and DR DC, but we still have the benefit of improving the RTO from about an hour to a few minutes required to run this script. Next step, it's time to start the critical service VM replica in DR DC with no changes in network configurations.

6. CONCLUSION

The paper describes practical experience in implementing an enterprise network infrastructure with BCP elements and considers options for improving RTO.

The main task of this paper is the deployment of an enterprise network infrastructure with improved resilience capabilities. To provide disaster recovery possibilities there have been chosen so-called active-passive strategy with one main DC and one DR DC.

We considered various options for improving network survivability and availability of critical services, considered possible risks and implemented several RTO improvements by implementing solutions at the routing level. There has been suggested the flat networking solution, the peculiarity of which is that L2VPN solutions were not used.

The suggested solution for reconfiguring network in main DC and DR DC to provide flat networking can be automated, but this may already be the subject of further research.

References:

Chassis Cluster User Guide for SRX Series Devices. (2023). *Juniper Networks*. Retrieved November 15, 2024, from https://www.juniper.net/documentation/us/en/software/junos/chassis-cluster-security-devices/topics/topic-map/security-chassis-cluster-verification.html

General Data Protection Regulation. (2018). *European Union*. Retrieved November 12, 2024, from https://gdpr-info.eu/

Hanks, D. R., Reynolds, H., & Roy, D. (2016). *Juniper MX series* (2nd ed.). O'Reilly Media, Inc.

Health Insurance Portability and Accountability Act. (2024). *U.S. Department of Health & Human Services*. Retrieved November 11, 2024, from https://www.hhs.gov/hipaa/index.html

ISO/IEC 27001. (2022). *International Organization for Standardization*. Retrieved November 12, 2024, from https://www.itgovernance.co.uk/iso27001

Juniper Networks. (2024). *Policy-statement overview*. Retrieved February 16, 2025, from https://www.juniper.net/documentation/us/en/software/junos/cli-reference/topics/ref/statement/policy-statement-edit-policy-options.html

Juniper Networks. (2024). *Routing instances overview*. Retrieved February 16, 2025, from https://www.juniper.net/documentation/us/en/software/junos/routing-overview/topics/concept/routing-instances-overview.html

Payment Card Industry Data Security Standard. (2024). *PCI Security Standards Council*. Retrieved November 14, 2024, from https://www.pcisecuritystandards.org/standards/pci-dss/

System and Organization Controls. (2024). *SOC2 UK*. Retrieved November 14, 2024, from https://soc2.co.uk/

Varnum, D. (2014). Juniper–Cisco–GRE IPSec with OSPF. *Overlaid.net*. Retrieved November 13, 2024, from https://overlaid.net/2014/02/01/juniper-cisco-gre-ipsec-with-ospf/

Watters, J. (2013). *Disaster recovery, crisis response, and business continuity: A management desk reference*. https://doi.org/10.1007/978-1-4302-6407-1

Andrian Piskozub
Lviv Polytechnic National University
Lviv,
Ukraine
andriian.z.piskozub@lpnu.ua
ORCID 0000-0002-3582-2835

Mikolaj Karpinski
University of the National Education Commission,
Krakow,
Poland,
Ternopil Ivan Puluj National Technical University,
Ternopil,
Ukraine
mikolaj.karpinski@uken.krakow.pl
ORCID 0000-0002-8846-332X

Elena Nyemkova
Lviv Polytechnic National University
Lviv,
Ukraine
olena.a.niemkova@lpnu.ua
ORCID 0000-0003-0690-2657

Yuriy Lakh
Lviv Polytechnic National University
Lviv,
Ukraine
yurii.v.lakh@lpnu.ua
ORCID 0000-0003-4153-8125

Roman Banakh
Lviv Polytechnic National University
Lviv,
Ukraine
roman.i.banakh@lpnu.ua
ORCID 0000-0001-6897-8206

Pawel Sawicki
SunsetPicnic UG Ltd.
Berlin,
Germany
pwlsawicki@gmail.com
ORCID 0000-0002-8798-7279