RFO and context INC0108178

FedWAN Day
Pieter Hanssens
15/10/2024

Belnet

Connecting communities

.be

Chapter 1

The incident.

The incident: parameters

Impacted Customers:

- MINFIN
- MINSOC
- FAVV
- MOBILIT

Impact: Loss of connectivity for the offices on the WAN side

Timing: 08/10/2024 13h51 – 15h19 (88’)

Belnet

Connecting communities

The incident: RFO

Cause:

- Planned Change by a Belnet network engineer when integrating a new customer (the WAN for JUST)
- Change was done on both iconnection points for JUST: our DC in BRUNOR and BRUFIN
- A faulty AS number was used on a connection between customer and Belnet
- It was the same AS number as Belnet is using between Belnet and Telenet
- This created an AS loop, that was detected by the routers
- This resulted in a withdrawal of ALL routes where this AS was present, creating the impact with other customers than just JUST

What do you mean? Isn’t our WAN separated?

Yes, it is, we use separate VRF’s, routing instances, VLANs, AS numbers etcetera to separate your FOD from other FOD’s, not only on the data-plane but also on the control plane

BUT:

We made a mistake in that we didn’t identify that this specific type of loop (AS loop) would retire ALL routes carrying this AS, regardless of the specific routing table.

-> Lesson 1: study how this can be prevented with our current hardware (without introducing other unwanted behaviours)

Belnet

Connecting communities

What do you mean? Isn't our WAN redundant?

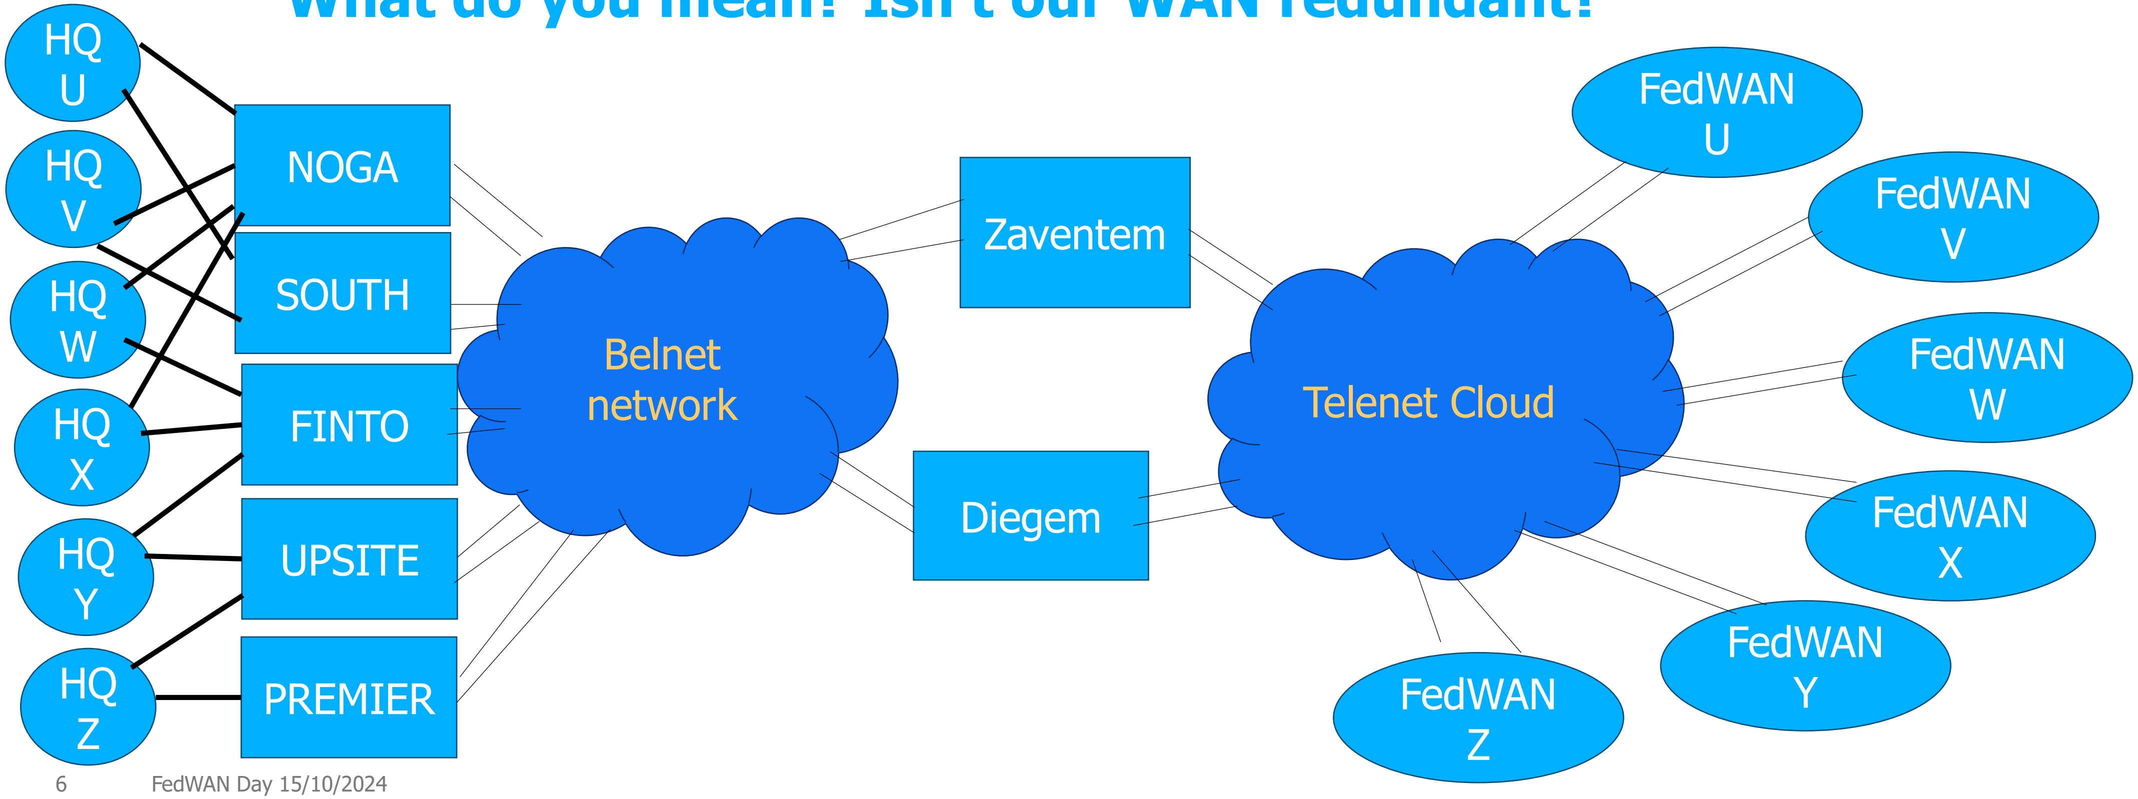

Belnet

Connecting communities

What do you mean? Isn't our WAN redundant?

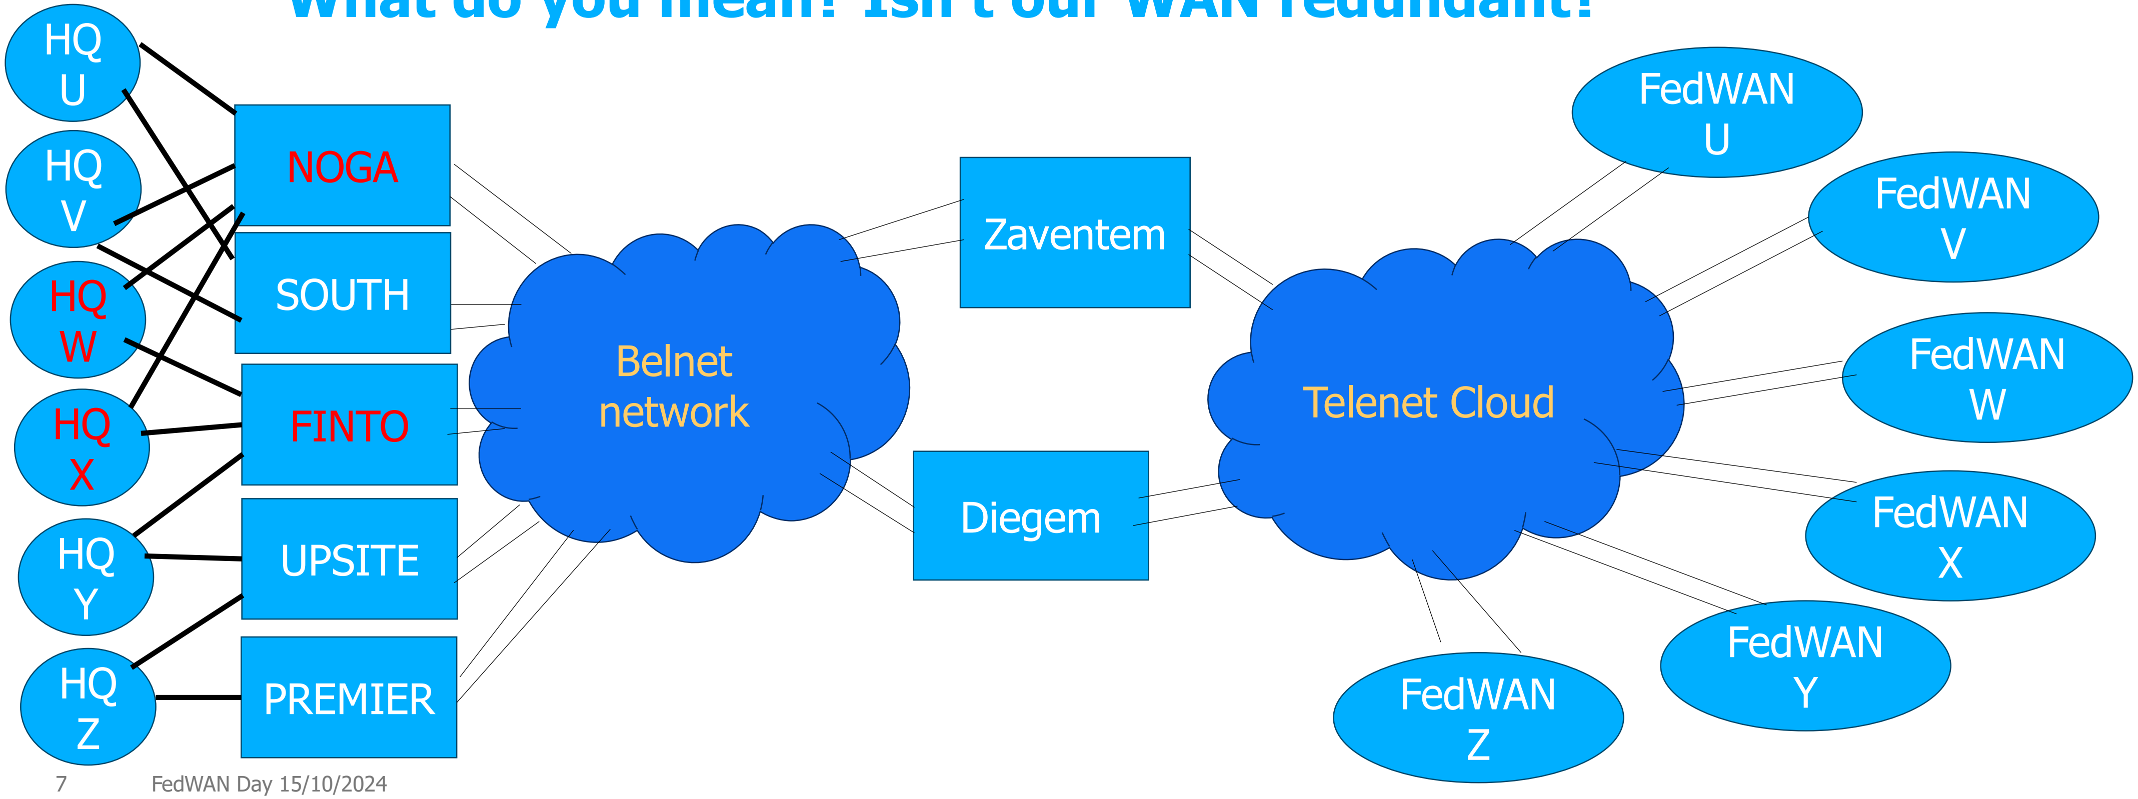

Belnet

Connecting communities

What do you mean? Isn’t our WAN redundant?

Even though very specific parameters had to be met:

- Customer already makes use of FedWANs
- Customer’s HQ is connected BOTH on North Galaxy AND Financial Tower

Impact was substantial, isolating the 4 impacted FedWAN’s

-> Lesson 2: even if the IA doesn’t identify any risk, use a stoptimer between configuring different nodes, allow simultaneous manipulations as a last resort

Belnet

Connecting communities

Why did it took 88 minutes to resolve?

No major alert was raised due to:

- no monitored elements were faulty
- i.e.: ports, links, bgp sessions, ospf states, lsp swaps, hardware issues all remained in the same state

Detection in this incident came through customer calls, some latency cropped up

- From Customer Servicedesk to ITSM team to Networks team or from Customer to Account Manager to ITSM team to Networks team lost us precious minutes

Once outage was known, identification and resolution took another 15 minutes

-> Lesson 3: better monitoring methods, encompassing more than just the current elements

Chapter 1

The aftermath.

No excuses

We were in the wrong as a Service Provider to you and we assume our responsability.

We will work on these three lessons:

- Software improvements to prevent these “multiple customer impacting” loops
- No simultaneous implementations on more than 1 location, however low or non-existant the impact may have deemed to be,
- Better monitoring on more parameters to detect these events quicker

Can I answer your questions?