”multiONE” with BGP communities

LHCONE meeting #51
19 September – Victoria CA
edoardo.martelli@cern.ch

Recap: multiple “LHCONEs”

Each site joins only the VPNs of the groups it is collaborating with (e.g. ATLAS-ONE, CMS-ONE, DUNE-ONE, BelleII-ONE...)

- **Major Benefit**: reduced exposure of data-centre/Science-DMZ to other sites
- **Major Challenge**: how to correctly route traffic into VPNs at sites that join several of them?

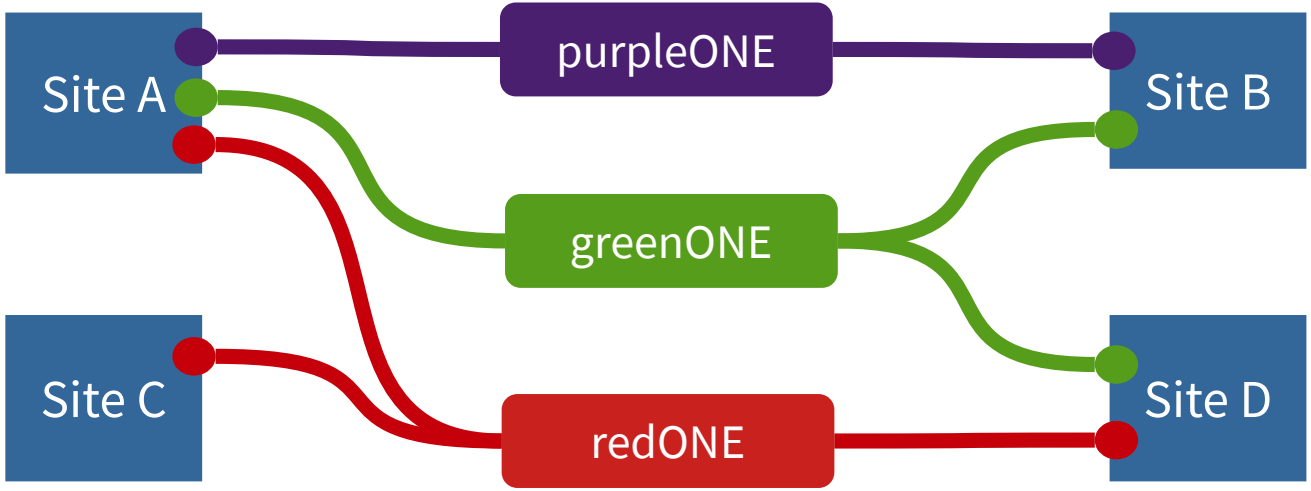

What happened so far?

Several proposals made, but all (practically and/or operationally) difficult to implement

Main problems:
- traffic separation not easy to achieve at sites
- multiple VPNs add complexity to network operations at sites and NRENs

New proposal

Don’t add any additional VPN (or maybe just one for Other Big Sciences)

Each prefix announced to LHCONE is tagged with BGP communities that identify the collaborations served by the site

The tagging is done by the sites, or by the connecting REN if they can’t do it

Sites can/should then decide to accept only the prefixes of the collaboration they are working with

In addition/alternative, RENs could announce to a given site only the prefixes of the collaborations related to the site

Practical example

Each site tags its prefixes announced to LHCONE with the BGP communities that identify the collaborations the site is participating in

The tagging can be done by the sites or the connecting NREN

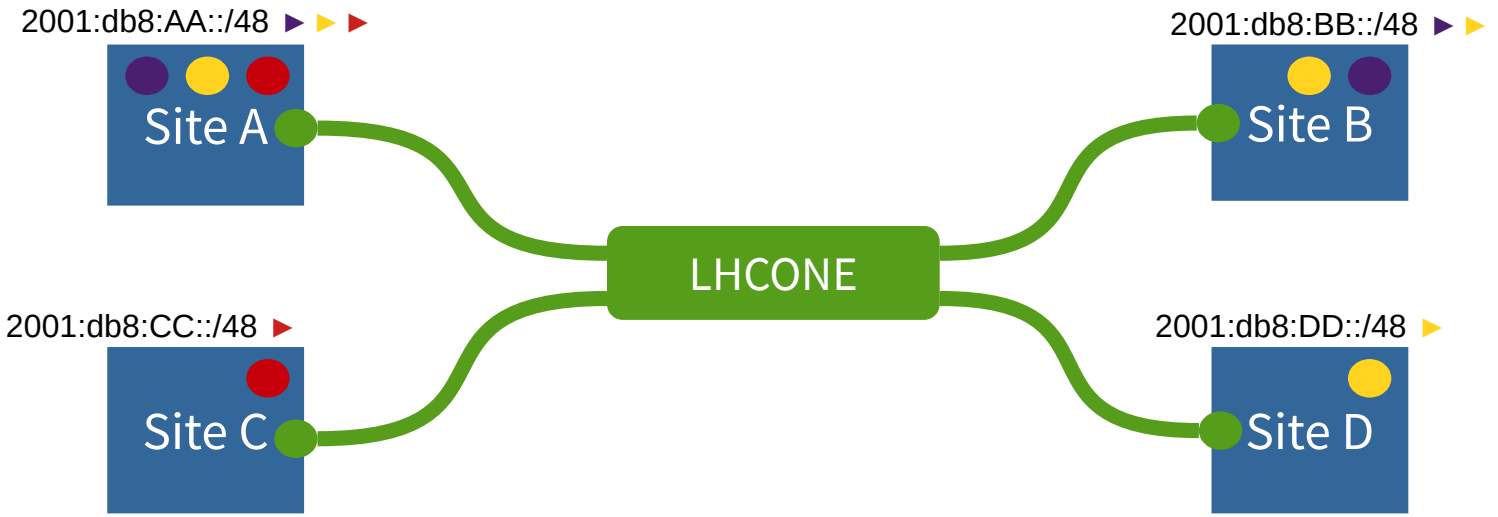

Practical example

Each site accepts only the prefixes tagged with the BGP communities of its own collaborations

The filtering can be done by the sites or the connecting NREN

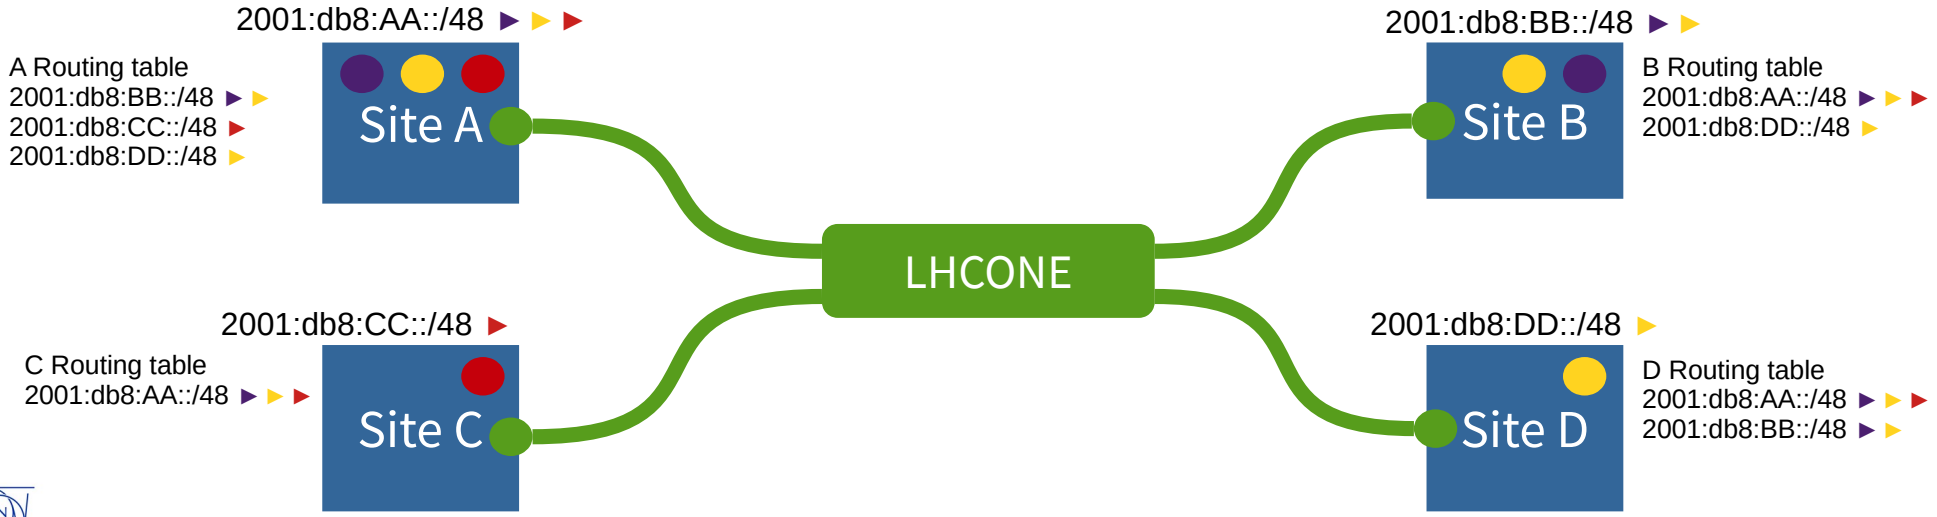

What could go wrong?

Site D doesn't filter out not-relevant prefixes

- traffic is sent to LHCONE but comes back via the Internet
- asymmetry may cause connections drop
- ..but anyway there should be no traffic between those sites, and the fix is easy to apply

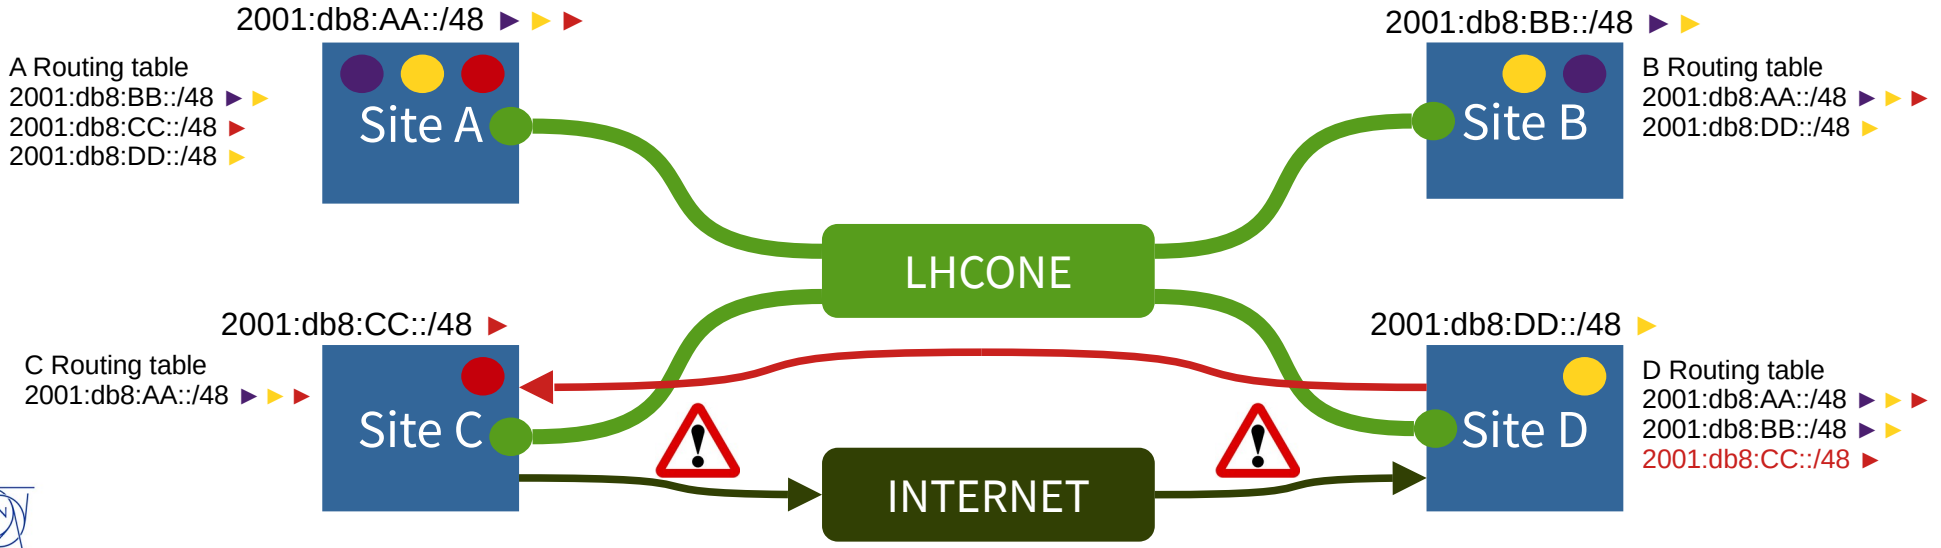

What could go wrong?

Site D doesn't tag its prefixes

- all sites will drop those untagged prefixes
- traffic will go via LHCONE but come back via the Internet
- asymmetry may cause connections drop
- Site D will immediately realize the mistake because LHCONE doesn't work for it

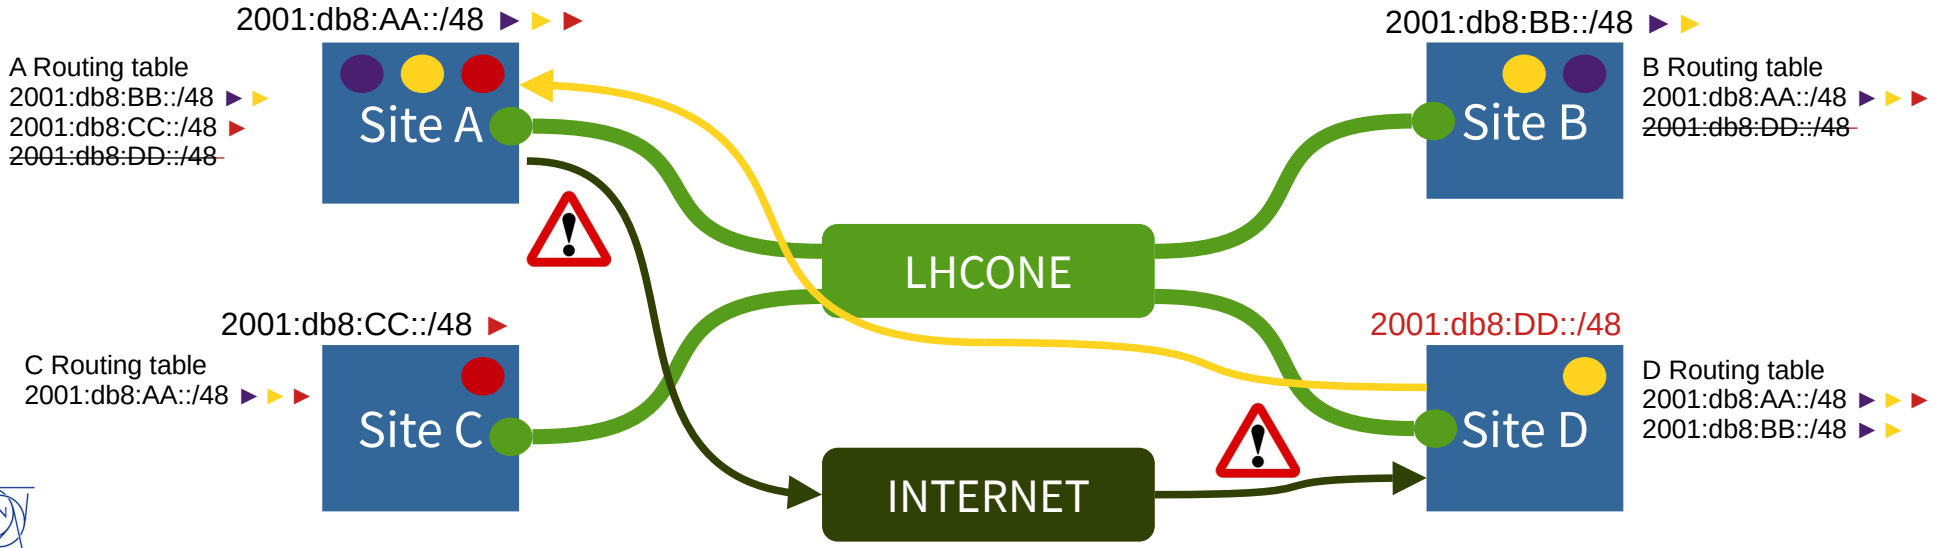

Implementation proposal

- Define BGP communities for the different LHCONE collaborations

- Implement prefix tagging at sites
    - if sites can’t do it, RENs will do for them

- Gradually implement filtering at sites and/or specific announcements at RENs

Benefits

- Reduced exposures of sites
- No additional VPNs to configure
- No changes at sites when a new site connects to LHCONE
    - only when a new collaboration joins, if they are interested in it
- Communication Errors will be an incentive to adhere
- Communication Errors will highlight already existing implementation errors and weaknesses

Opinions?

edoardo.martelli@cern.ch