Decentralized cloud WAN traffic engineering with BlastShield

Umesh Krishnaswamy

Rachee Singh

Nikolaj Bjorner

Himanshu Raj

Microsoft

Software-driven WAN

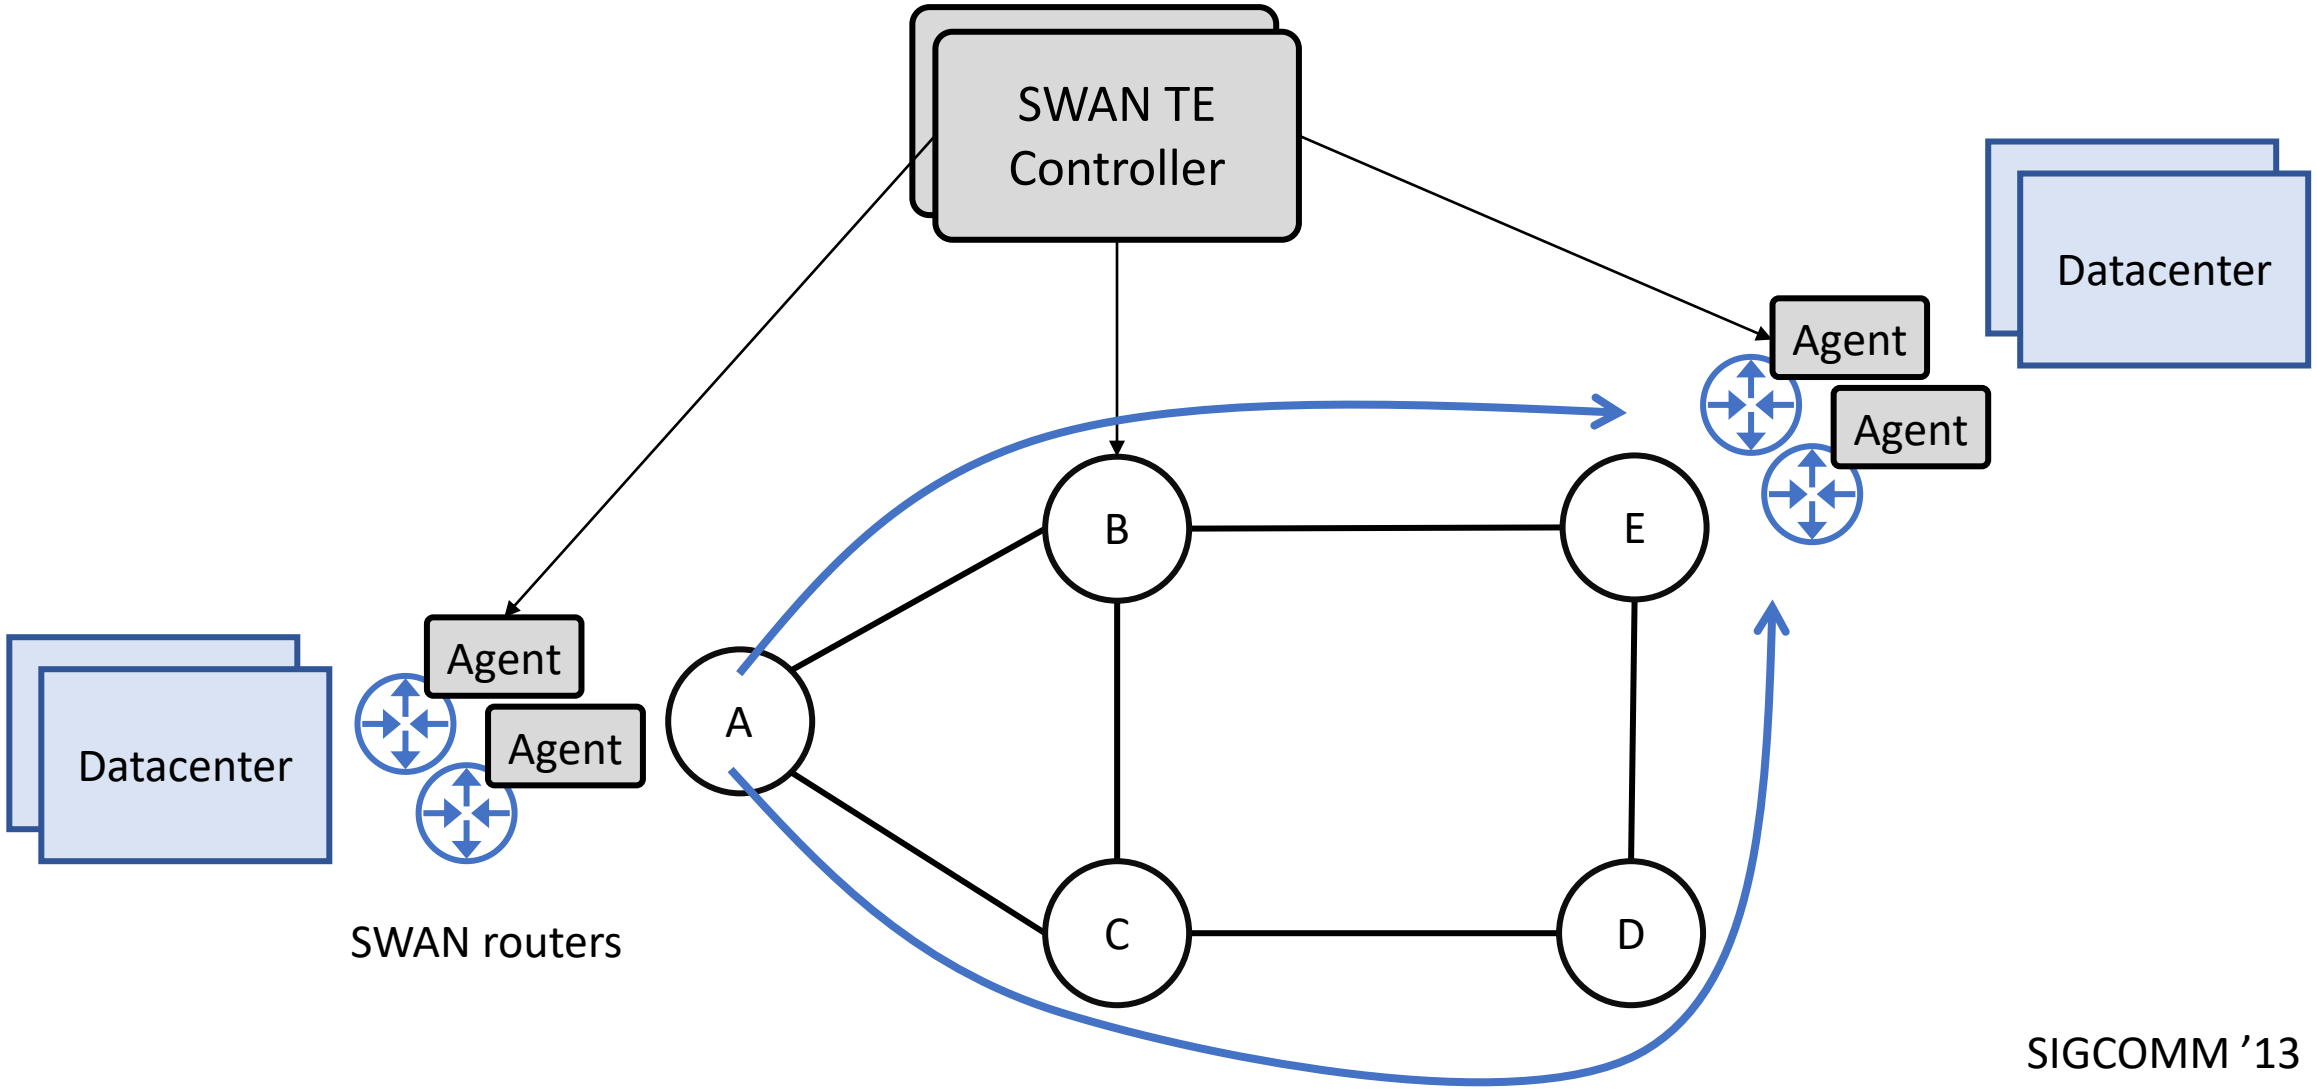

SWAN traffic

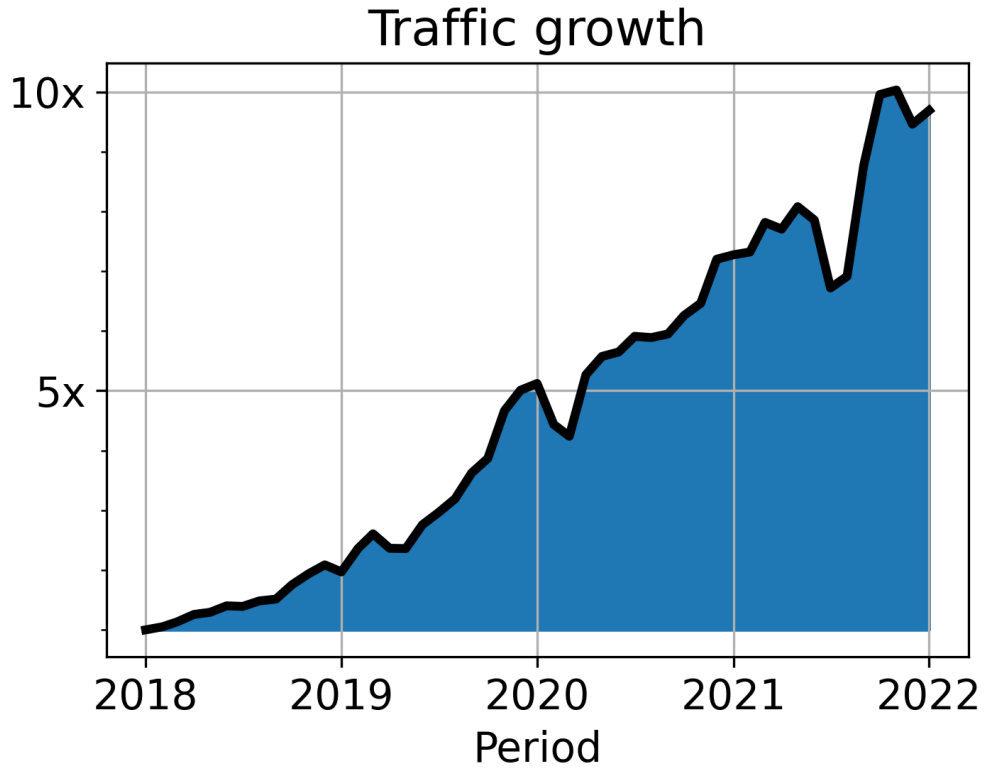

SWAN outage of global scope

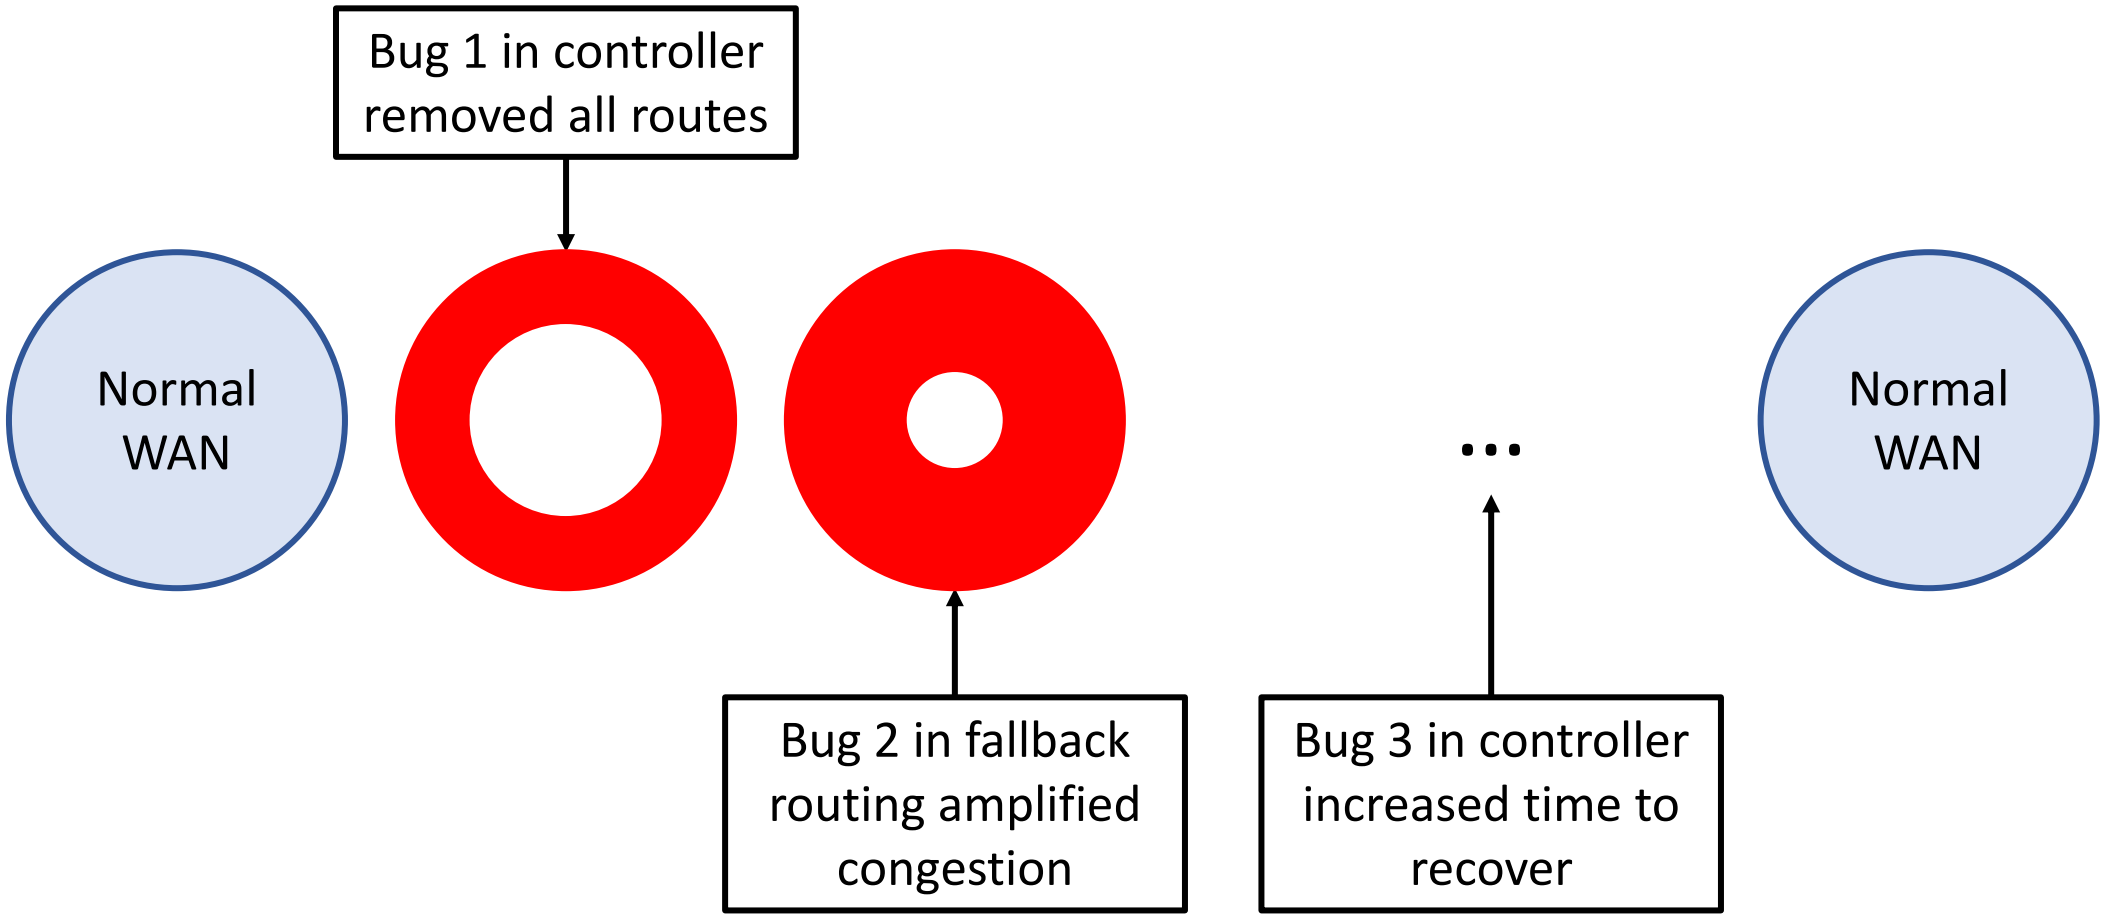

Blast radius

	Tier	Service level objective
Customer traffic	Tier-0	99.999
Discretionary traffic	Tier-1	99.9
	Tier-2	99

Blast radius = customer traffic at risk from a controller failure

BlastShield slices

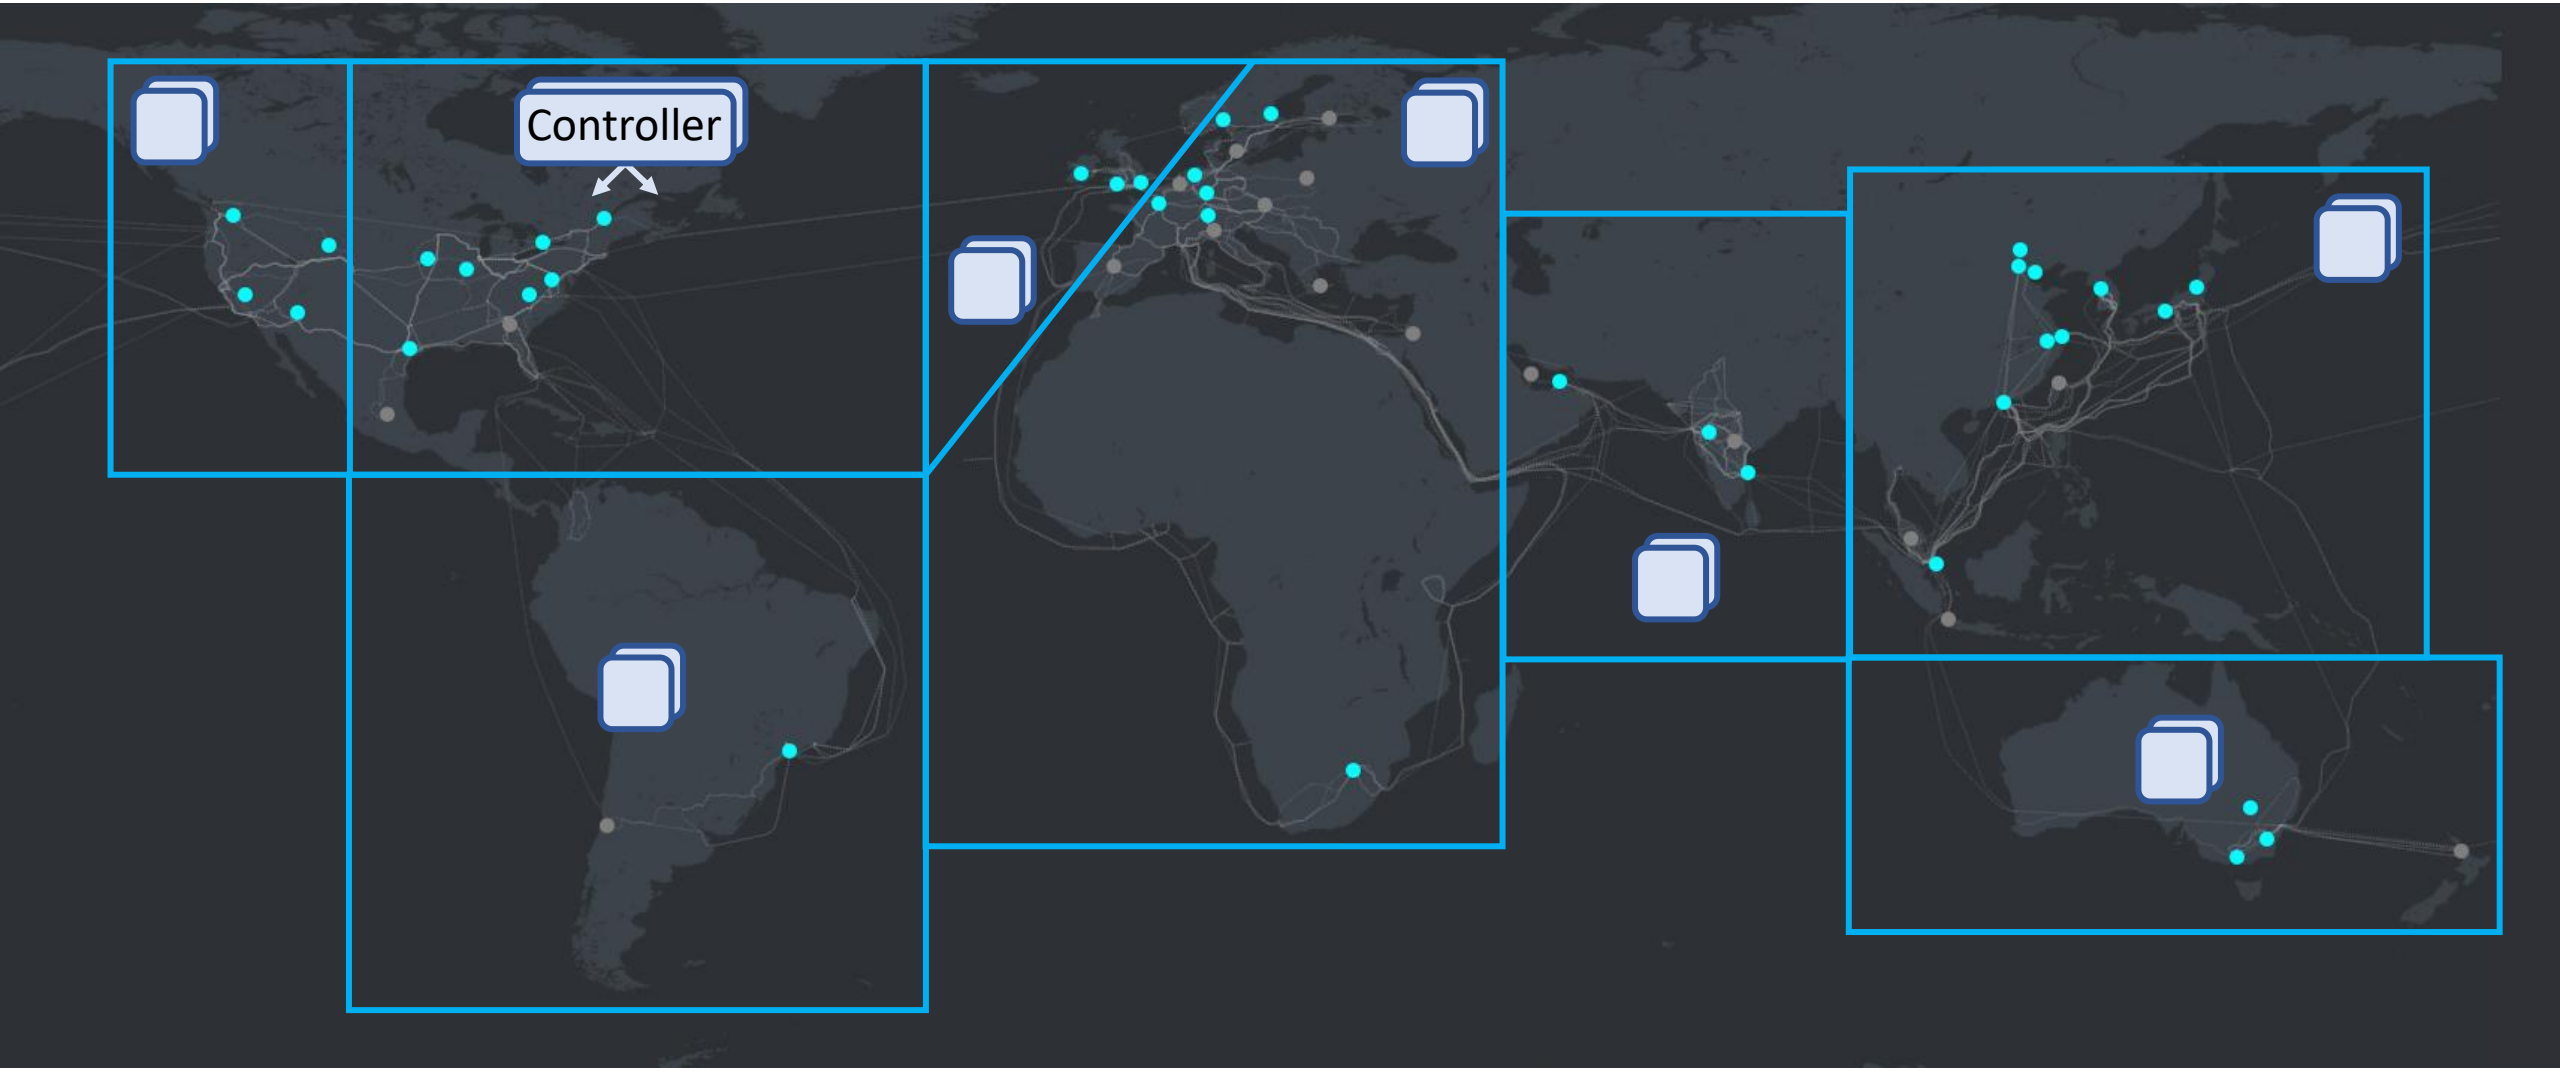

Design assumptions

Decentralized	Hierarchical
Global view	Local view
No coordination	With coordination
Inter-slice traffic < Intra-slice traffic	

BlastShield controller

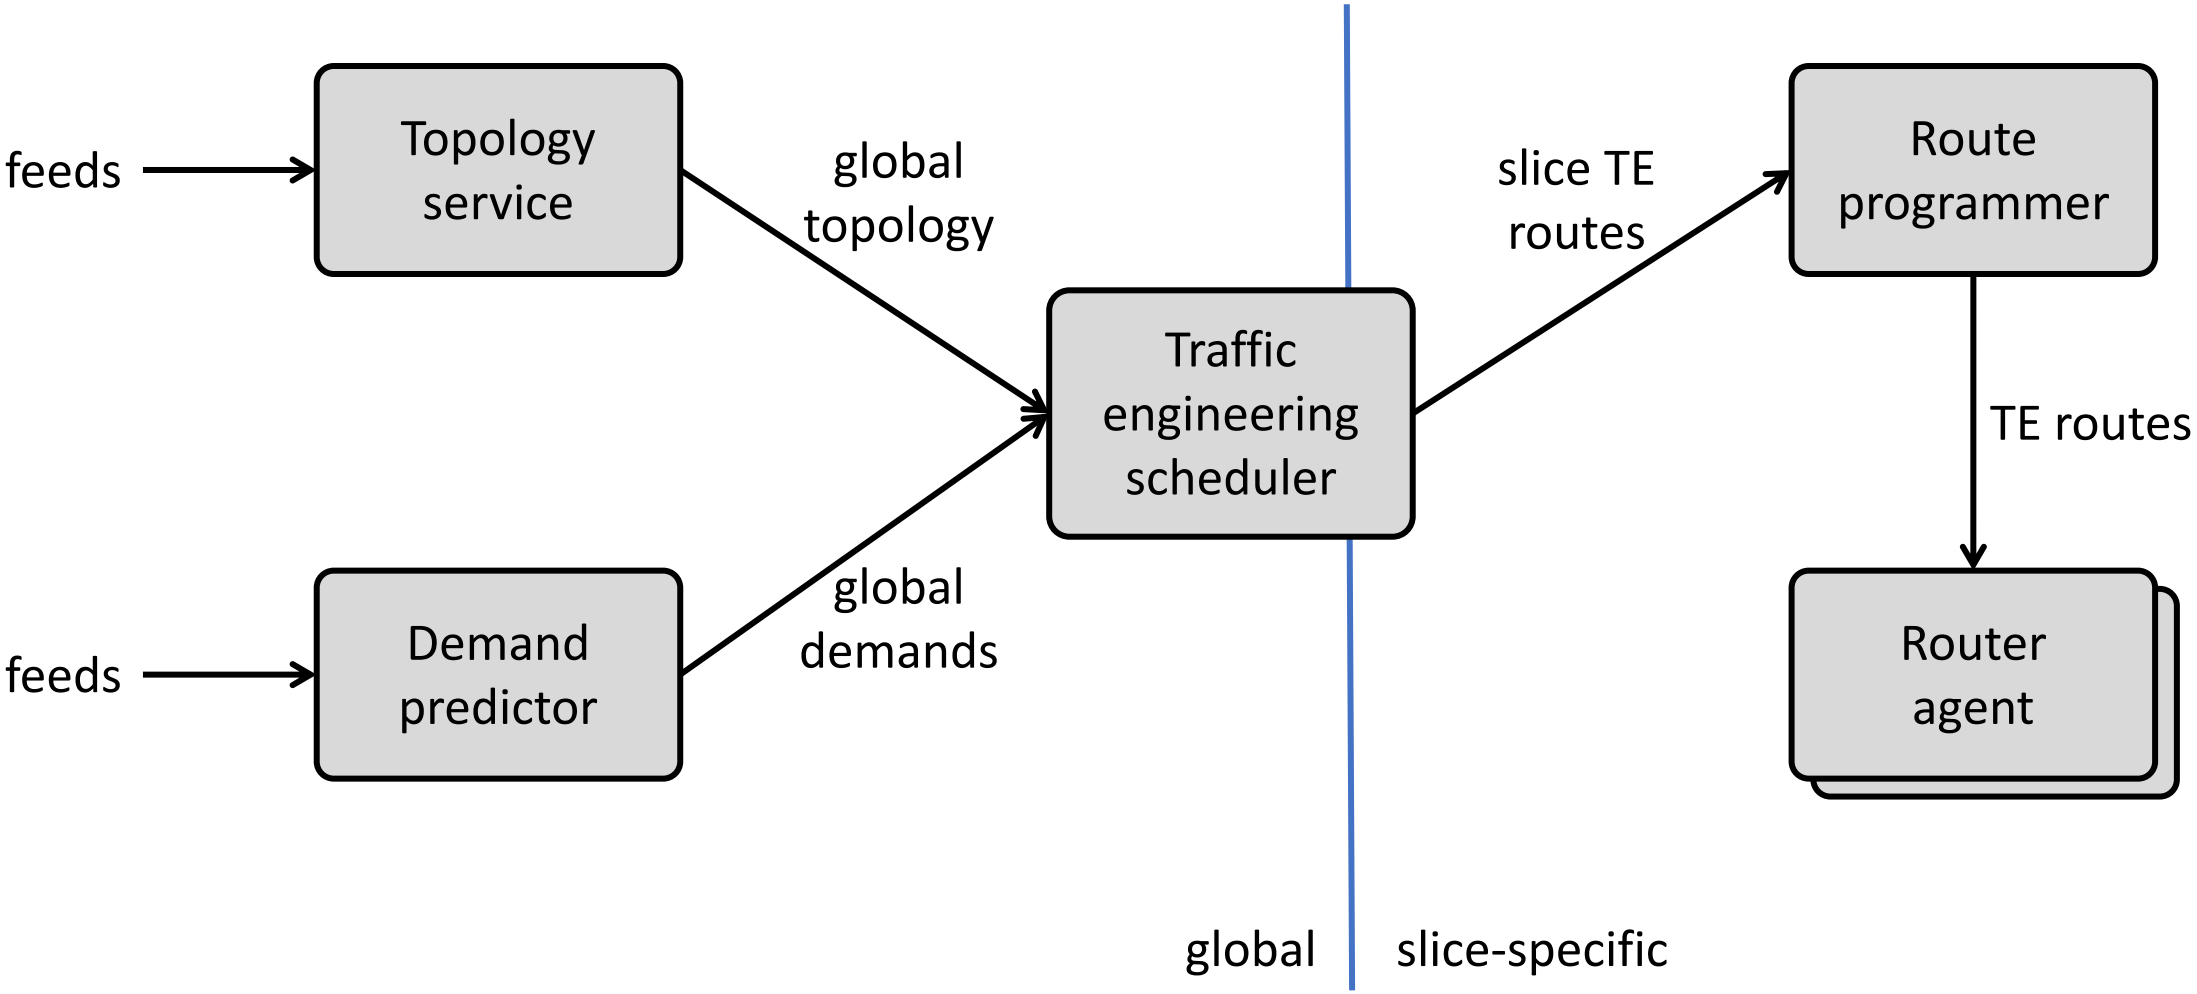

Inter-slice routing

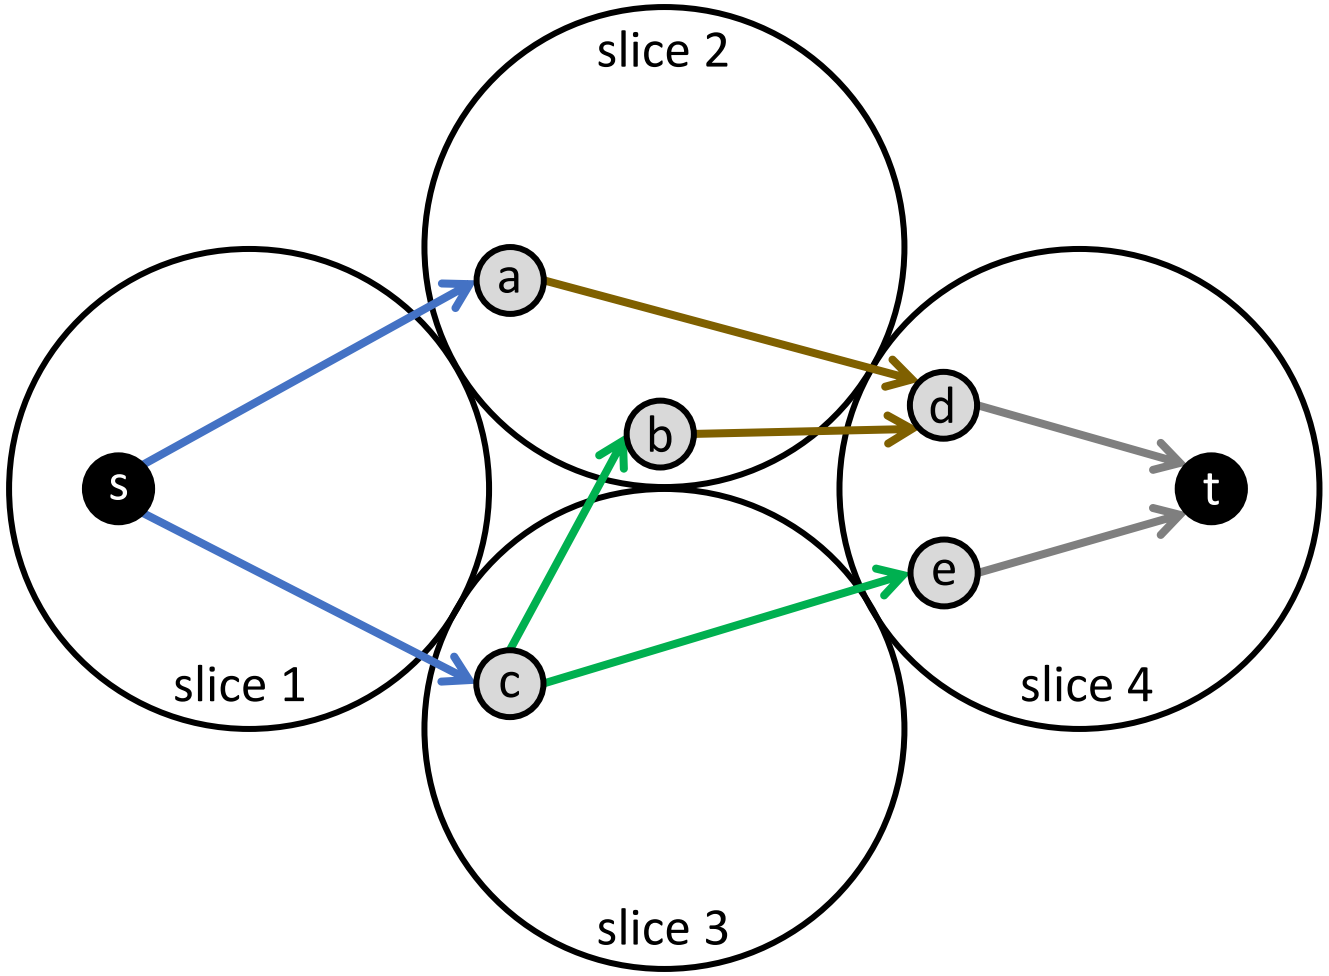

$s \rightsquigarrow a \rightsquigarrow d \rightsquigarrow t\ (w_1)$
$s \rightsquigarrow c \rightsquigarrow b \rightsquigarrow d \rightsquigarrow t\ (w_2)$
$s \rightsquigarrow c \rightsquigarrow e \rightsquigarrow t\ (w_3)$

TE path segments programmed by slice controllers

Blast ripple and routing loops

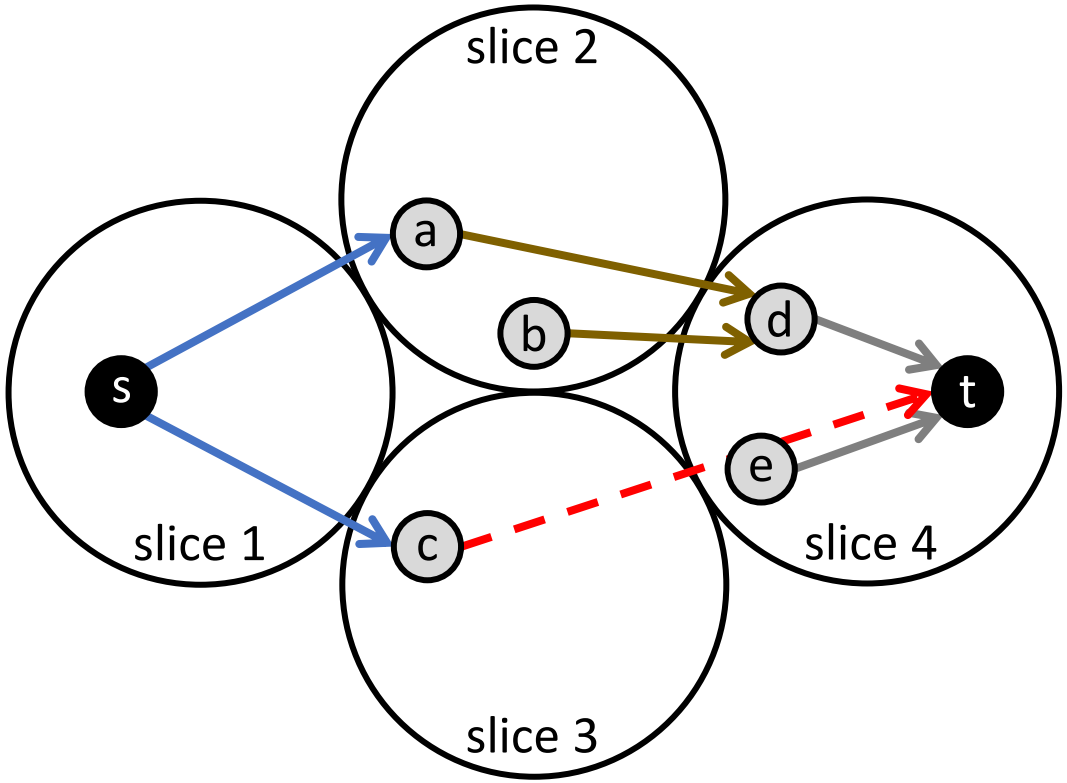

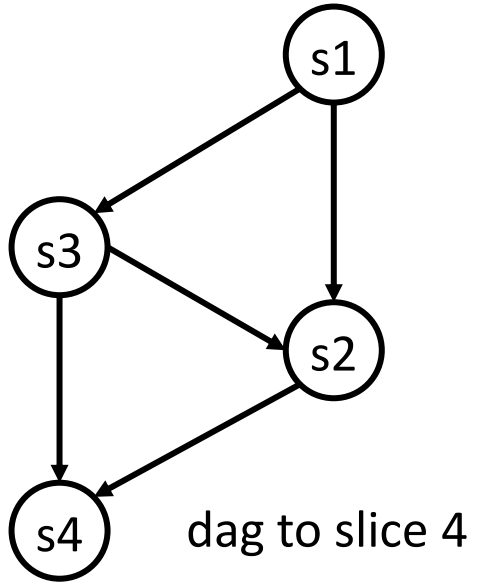

Traffic uses fallback routes to destination when downstream controller fails.

Enter-leave constraints restrict paths to achieve loop-free routing

Source routing

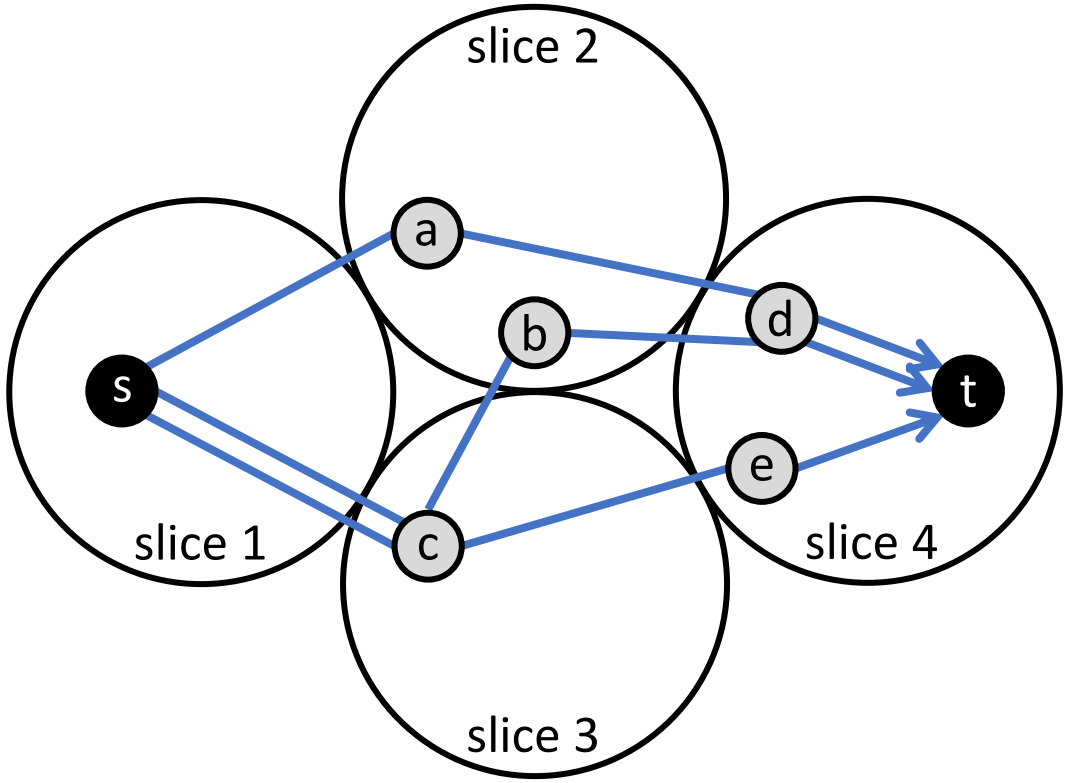

TE paths programmed by source slice controller

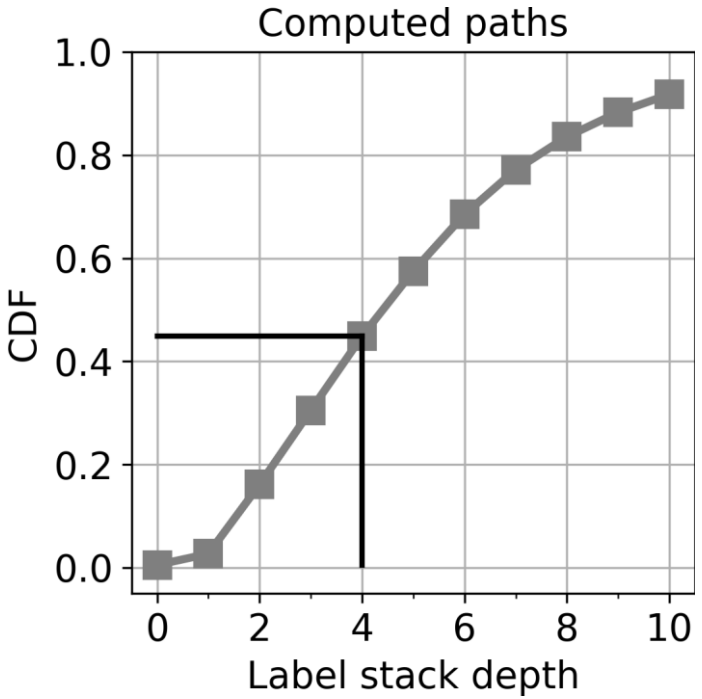

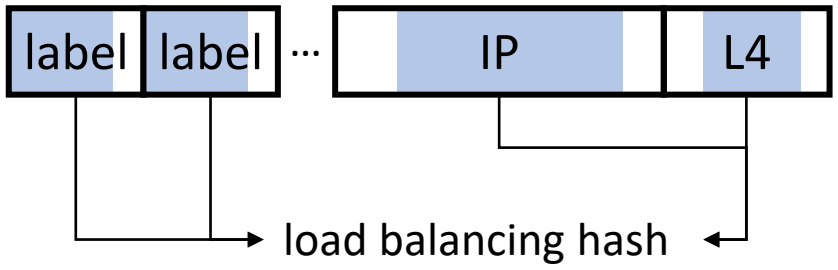

Traffic engineering scheduler

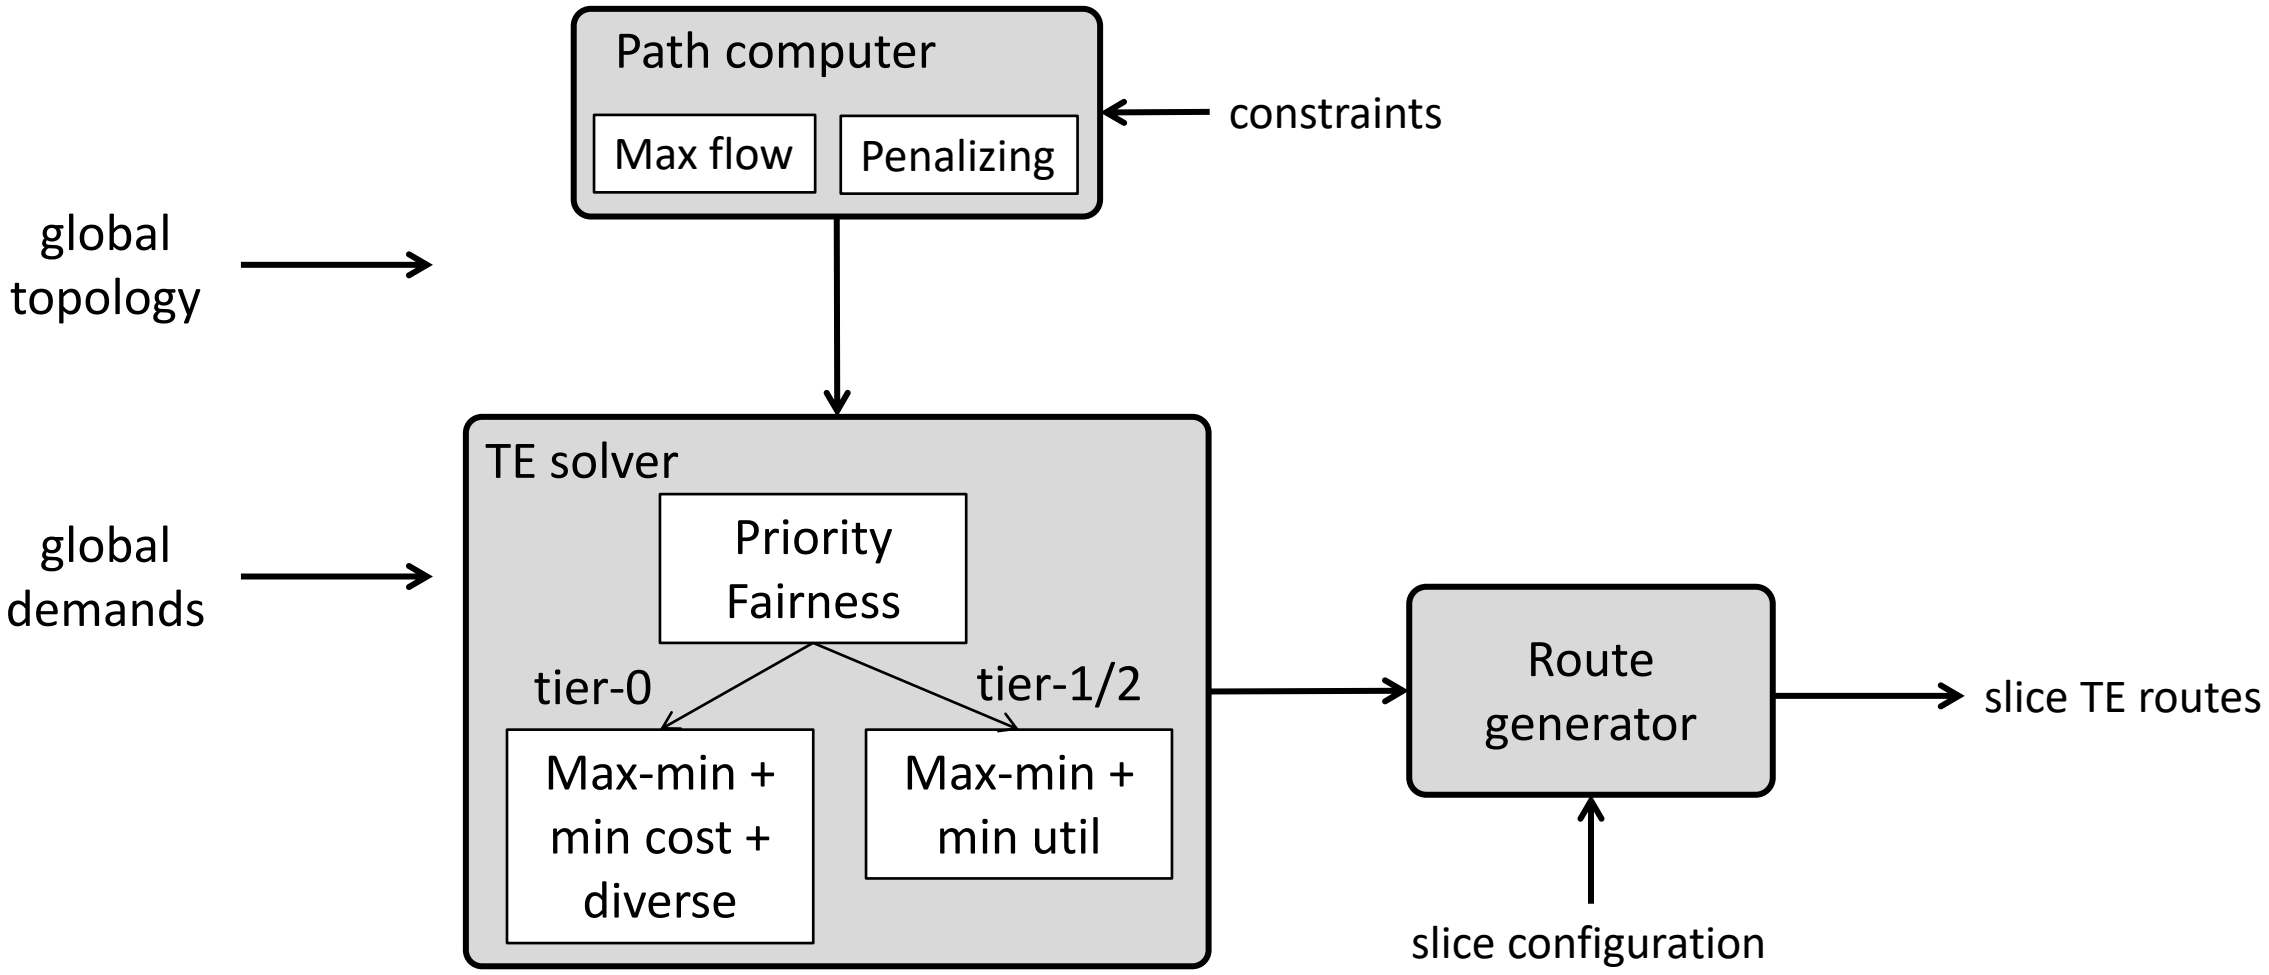

Symphony or cacophony

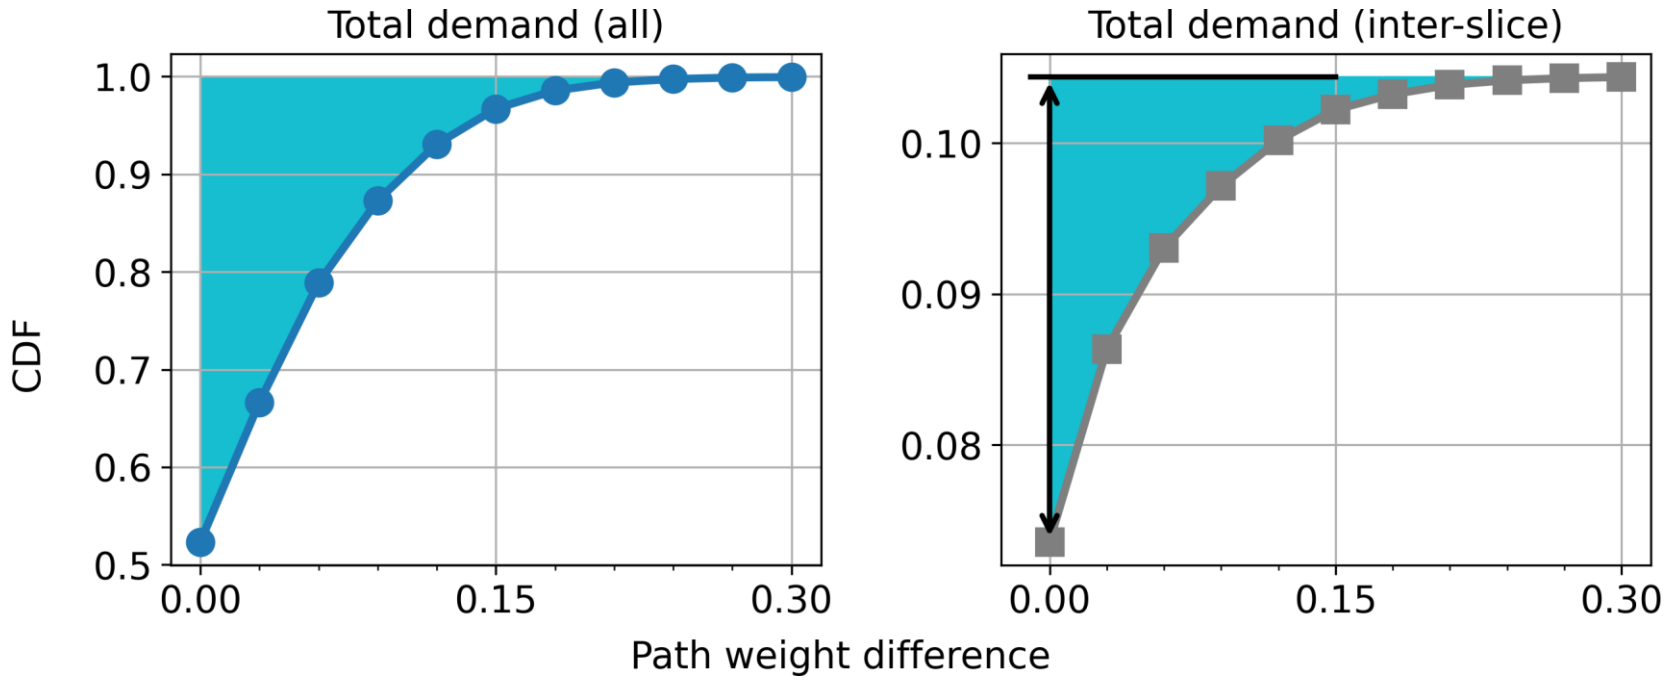

TE inefficiency 3%

Blast radius reduction

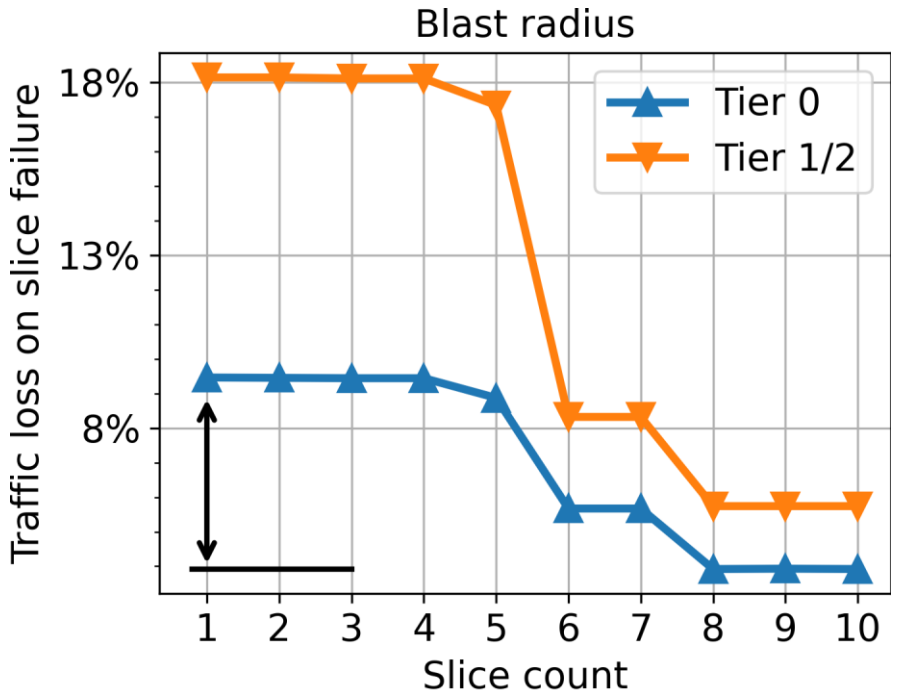

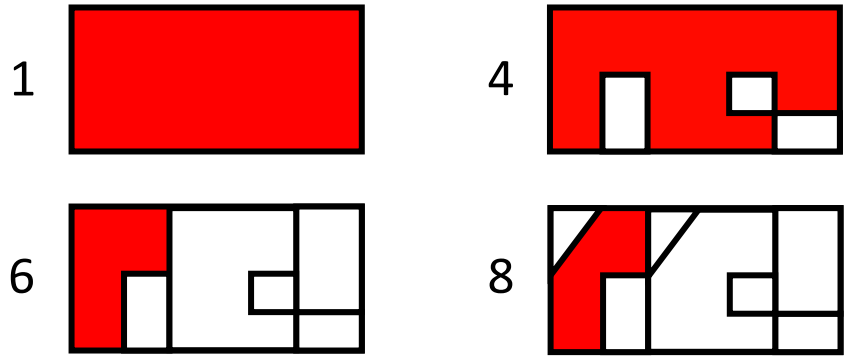

Blast radius reduction 6%

Blast radius reduction

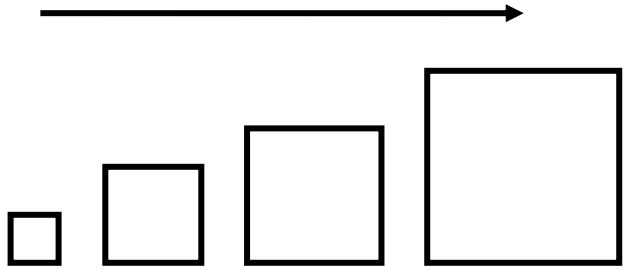

Safe deployment lowers failure probability

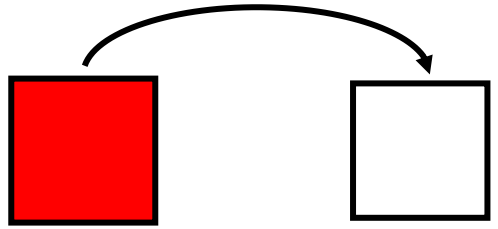

Applications have option to fail out

Summary

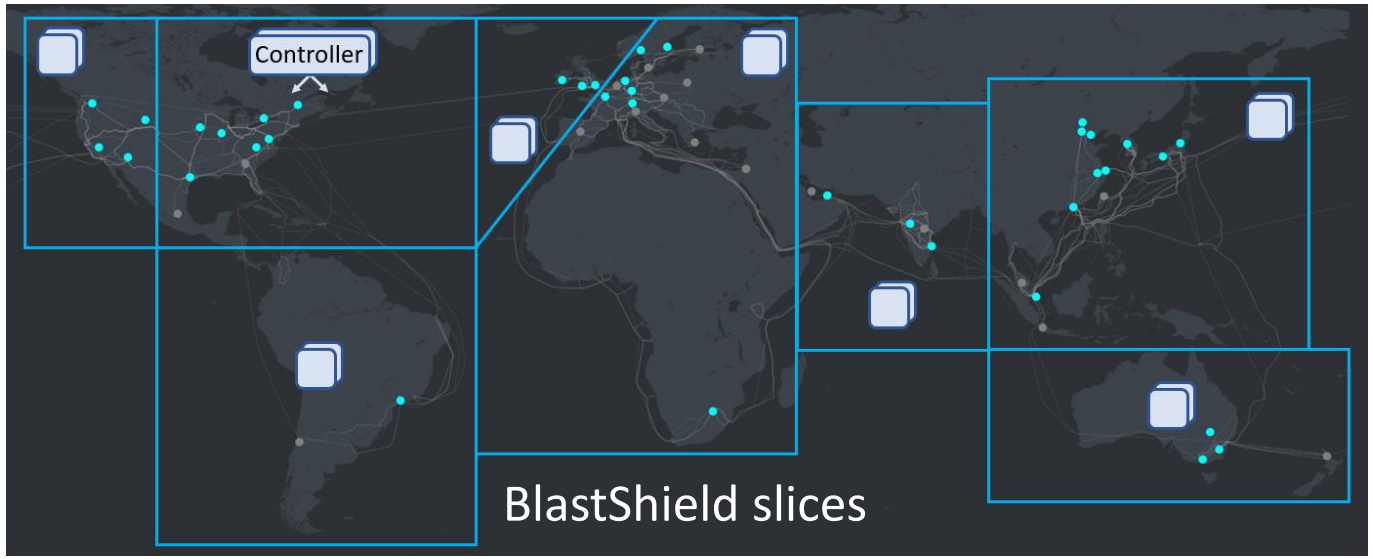

Decentralized	Hierarchical
Global view	Local view
No coordination	With coordination
Inter-slice traffic < Intra-slice traffic