Tunnel based FRR

<draft-bryant-ipfrr-tunnels-00.txt>

RTWG IETF-60 August 2004

Stewart Bryant <stbryant@cisco.com>

Mike Shand <mshand@cisco.com>

Goals

- **FRR MUST do no harm – the impact of the mechanism is never worse than if it were not used.**
- **Once a router has detected the failure, no further packets will be lost.**
- **No topology tuning required.**
- **MUST be suitable for incremental deployment**

Implications of the goals

- **Following invocation of the repair a controlled convergence is needed to avoid undoing the FRR repair, and collateral damage due to micro-looping.**
- **Controlled convergence takes time, therefore repair must be 100% to prevent extending outage for un-repaired destinations.**

Overview

- **This is a long-reach repair mechanism to complement ECMP and “downstream” routes.**
- **Works by tunnelling the packet to a router in the network, which is reachable by the repairer, and which has a natural route to the destination that avoids the failure.**
- **Simplified computation by using other side of the failure as a proxy for the packet destination.**

Basic Operation

SPF from A with all destinations reached via the failure excised

Reverse SPF from B with all Nodes that reach B via the failure excised

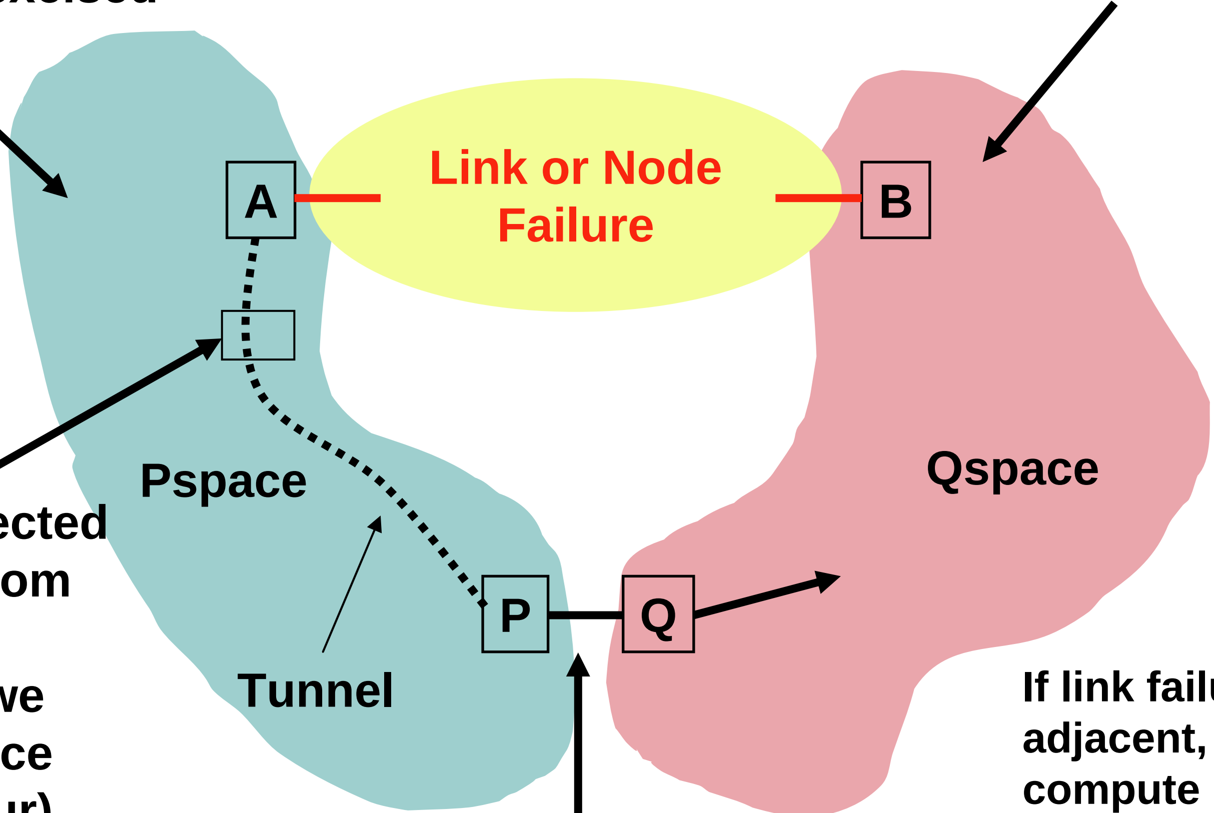

First hop directed forwarding from A can extend Pspace (i.e. we use the Pspace of a neighbour)

Max 1 hop if metrics symmetric

If link failure A & B adjacent, otherwise compute Qspace for each of B's neighbors reachable via AB

Interference

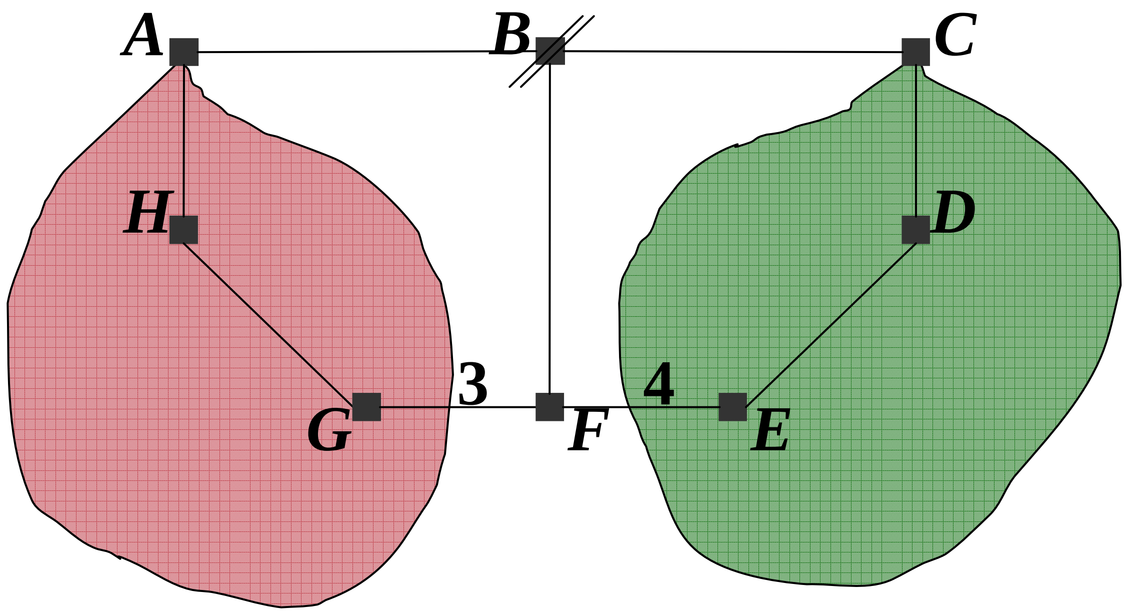

- A node repair problem that SOMETIMES arises due to the packet getting sucked back towards the failed node.
- Solved by concatenating repair paths using a selected neighbour (F) as an intermediary.
- A encaps to F, repairs to F, F decaps and repairs as normal.
- MAY need to repeat this secondary repair process to another neighbour.

Multi-homed Prefixes

- **A very similar problem to interference in which nodes unaware of the failure “suck” the packet back to the failed node.**
- **Only affects node protection**
- **Solution is to encapsulate packet to alternate router with reachability to the prefix, and then repairing to that router.**

Loop-free via delayed FIB update

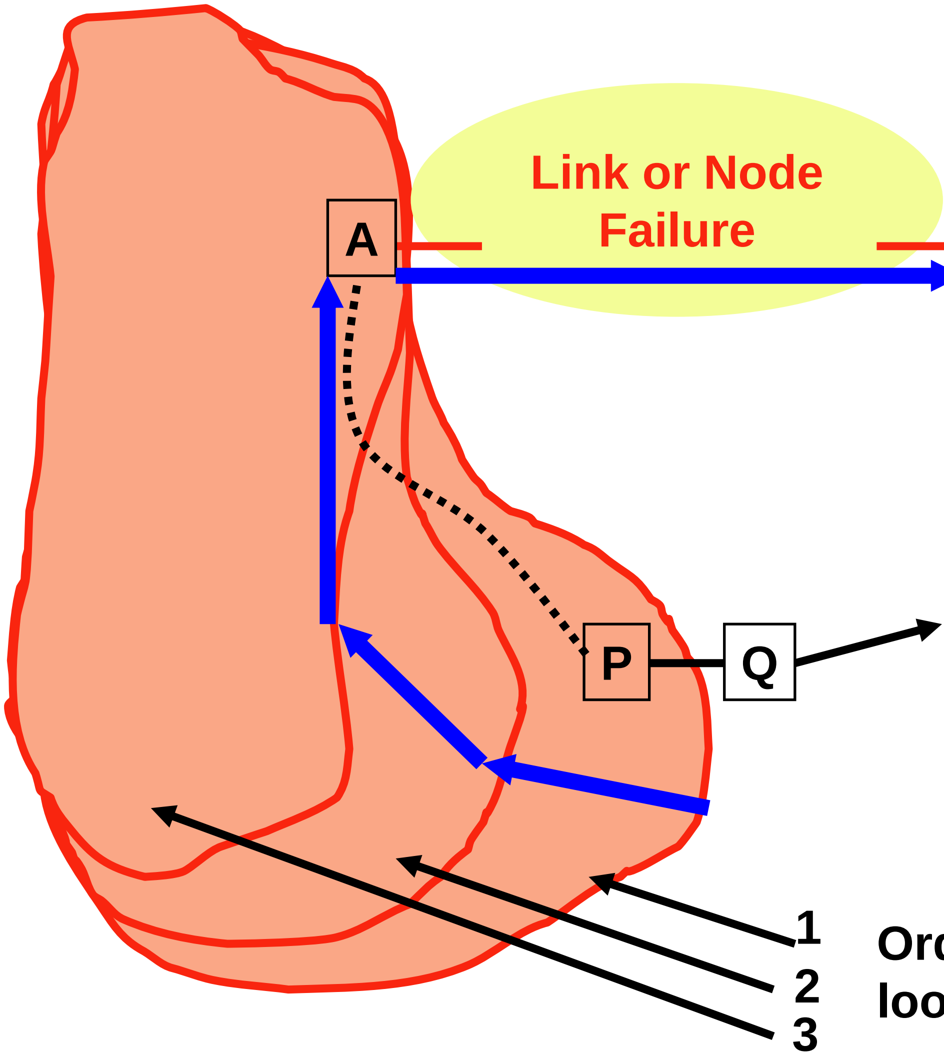

Each node computes
Its own order of transition by computing a reverse SPF as if it were rooted at B

Order of change of FIB for loop-free convergence

Data-plane modifications

- **Rapid detection mechanism and routing to alternative next-hop is common to all FRR solutions.**
- **To cover all pathological case may need three layers of tunnel encapsulation and one directed forwarding operation:**
  - **Encapsulate to MHP**
  - **Encapsulate to secondary repair**
  - **Encapsulate to P**
- **Any tunnelling mechanism may be used: IP-IP, GRE, L2TPv3**
- **The only nodes needing modification are the encapsulating routers. Tunnel decapsulation is a “standard” mechanism.**

Control Plane Modifications

- New sub-TLV to flood FRR parameters
  - Router FRR capable
  - Link protected
  - DF vector
- IPFRR routers must calculate repair strategy.
- For traffic for which node is single point of failure, repairing router must do node-link discrimination check.
- Loop-free convergence requires additional calculation and controlled execution of FIB updates.

Dataplane complexity

Tunnel encapsulation, particularly the need to apply nested tunnels in sequence due to the need to fixup length and checksum

Control Plane Complexity – Link Protection

- **Symmetric costs**

  For each protected link, each node prunes the existing SPF and calculates 1 reverse SPF

- **Asymmetric costs**

  As above, plus up to k-1 SPF to extend Pspace <u>if needed</u>

Note – SPFs can terminate as soon as repair is found.

Control Plane Complexity – Node Protection

- **Symmetric Costs**

  If secondary repairs not needed, then for each protected neighbour we need 1 SPF prune plus k-1 reverse SPF.

  For each neighbour taking part in a secondary repair we need one additional SPF.

- **Asymmetric Costs**

  As above, plus up to k-1 SPF to extend Pspace <u>if needed</u>

Loop-free convergence

Several methods – consider ordered FIB update

Each node effected by the failure computes 1 reverse SPF (from B), and determines it's position WRT the horizon

Each node must update its FIB within a maximum time.

As an optimisation may use signalling to reduce the time needed to converge.

Comparison with other methods

- **This is a long-range method, capable of finding and using a repair point some distance from the failure.**
- **In symmetric cost networks (and non-pathological asymmetric cost networks) repair coverage is 100%, and when used with loop-free convergence, post repair packet loss is zero.**
- **Following an arbitrary number of failures, the network will recompute an equally effective repair strategy limited only by an induced single point of failure.**
- **Layered tunnelling allows us to overcome pathological topologies, and to repair multi-homed prefixes.**
- **Use of other side of failure as proxy for the destination results in a significant reduction in repair path computation.**
- **Does not require a change to forwarding behaviour of neighbours (U-turn).**

What we can take from other methods

- **Per-destination strategy may enable us to use less complex repair strategy to some destinations.**
- **IP loose source routing or multi-hop tunnels (e.g. MPLS) could enhance this solution.**

Coverage In Some Operational Networks

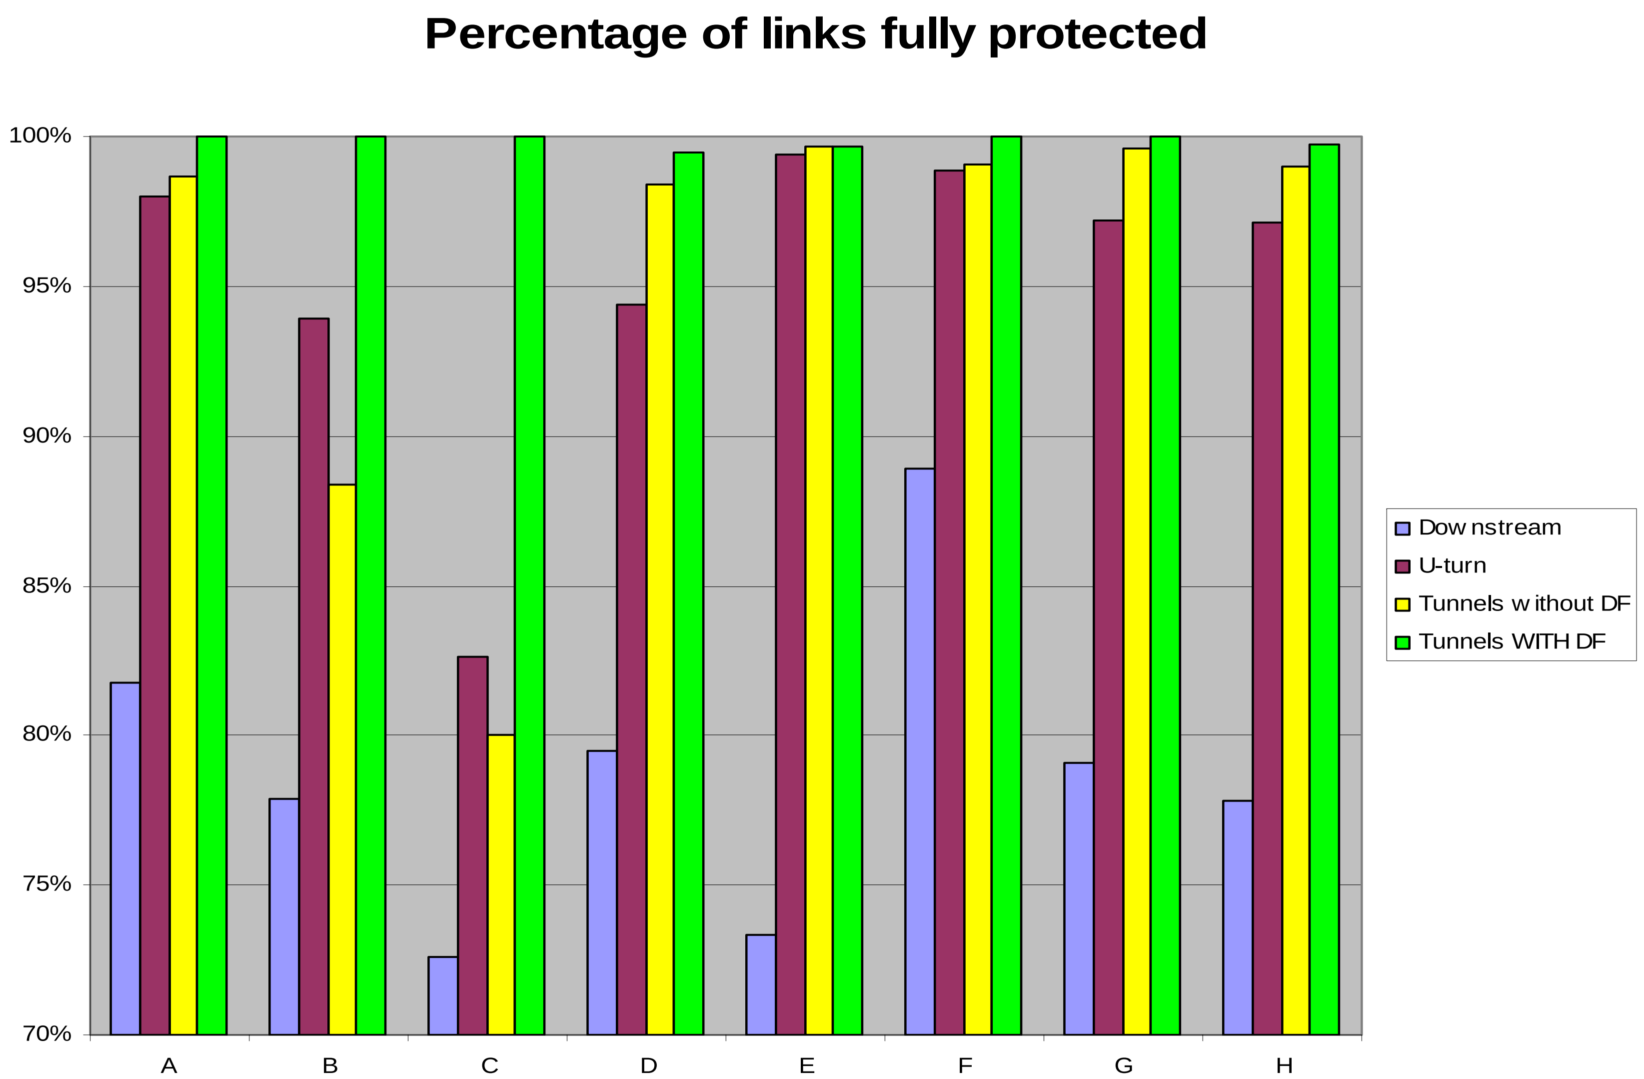

- **Thank You**