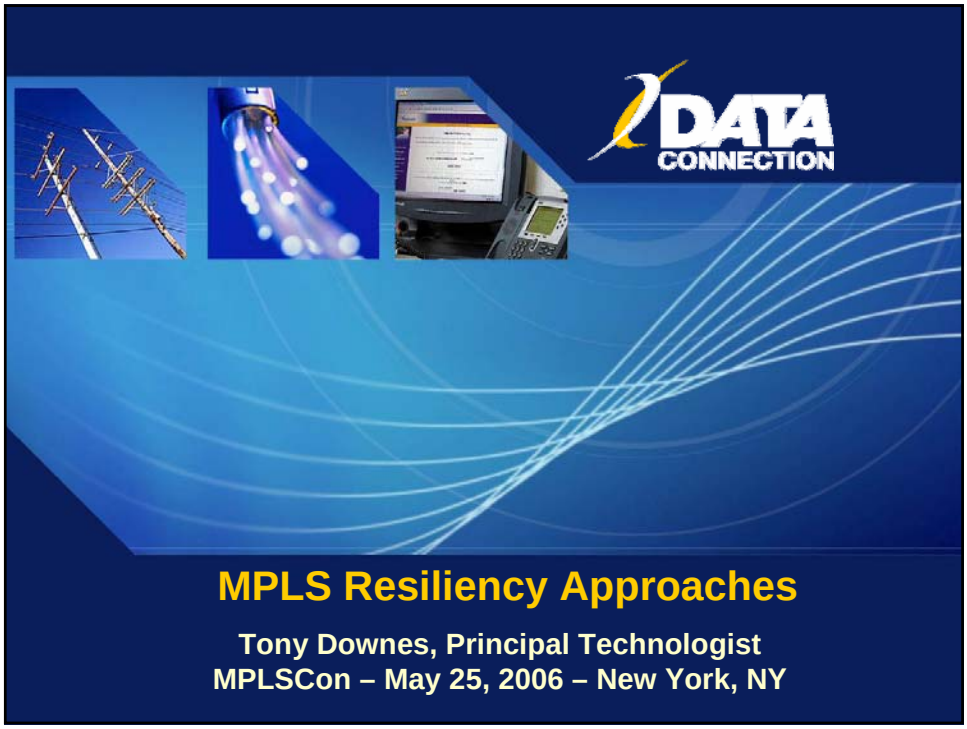

MPLS Resiliency Approaches

Tony Downes, Principal Technologist
MPLSCon – May 25, 2006 – New York, NY

Data Connection Overview

- Background
  - Founded in 1981
  - Headquarters in Enfield, UK
  - 300 employees across 7 locations
  - Network protocols
    - VoIP – SIP, MGCP/Megaco, Session Border Controller (SBC)
    - IP Routing – OSPF, BGP, IS-IS, PIM
    - MPLS – RSVP, LDP, VPLS, VPWS, …
    - ATM
  - Internet applications
  - MetaSwitch
- Independence and Stability
  - Steady, profitable growth
  - Privately held & self-funded by Employee Benefit Trust

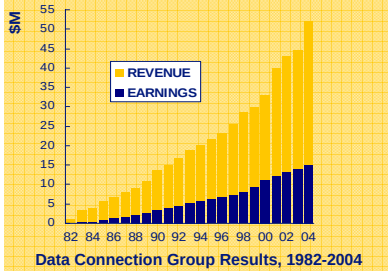

Data Connection Group Results, 1982-2004

Copyright © 2005 Data Connection Limited • Proprietary and Confidential • All Rights Reserved.

MPLS Resiliency - Requirements

- Fundamental requirement
  - No interruption to traffic
  - Typically requires less than 30-50msec of traffic loss – this is the minimum loss that does not (seriously) affect voice
- Classes of “failure” within a system
  - Software failure
  - Hardware failure
  - Controlled software upgrade (and downgrade)
  - Controlled hardware replacement
  - Mis-configuration / “operator error” – often the major cause!
- Classes of “failure” within the network
  - Device failures
  - Link failures
  - Mis-configuration / “operator error” – often the major cause!
  - Poor network design – not link / node disjoint

Copyright © 2005 Data Connection Limited • Proprietary and Confidential • All Rights Reserved.

MPLS Resiliency – Graceful Restart

- RFCs 3473 (RSVP-TE), 3478 (LDP)
- Resynchronization of path state with a neighbor following
  - Failure and restart of local node’s control plane software
  - Failure and restart of neighbor’s control plane software
  - Temporary failure of a link (RVSP-TE only)
- Can handle many software and/or link failures, but
  - Assumes forwarding is maintained separately (eg on line card)
  - Resynchronization takes time (~secs) – which can be an issue in a larger network or where failures are not rare
  - Resources (eg LSPs) can be “stale” – eg where LSPs are bought down by other nodes during the failure
- Can be used for
  - software upgrade / downgrade, hardware replacement
  - protection of out-of-band signaling

Copyright © 2005 Data Connection Limited • Proprietary and Confidential • All Rights Reserved.

MPLS Resiliency – High Availability

System Manager
- Creates backup process if required
- Initiates replication procedures
- Handles failovers

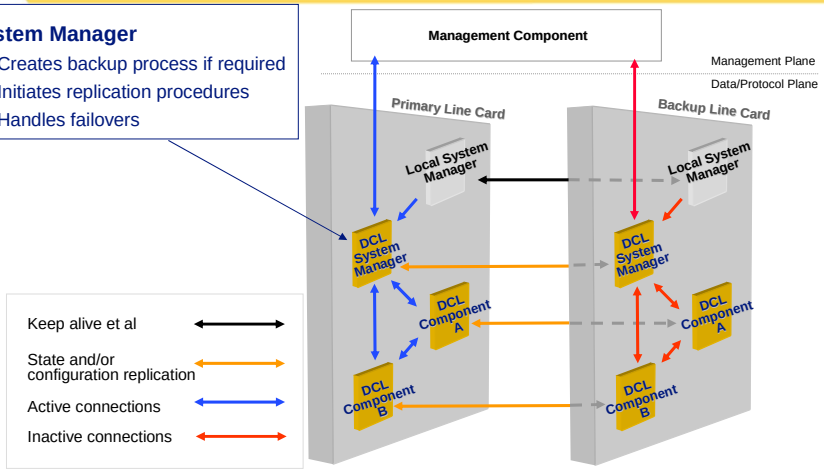

- Handles software and hardware failures, hardware replacement and software upgrade/downgrade
- Requires extra hardware – and very careful software design

Copyright © 2005 Data Connection Limited • Proprietary and Confidential • All Rights Reserved. 5

MPLS Resiliency – Make-Before-Break

- draft-ietf-mpls-rsvplsp-tunnel
- Set up new LSP route (tunnel instance) and switch data to it only when it is established
- Can handle node and/or link failures
  - But requires that all LSPs have backup tunnels
  - But requires ingress to detect and switch to backup
  - Is useful for operator cleanup after failures / recovery

Copyright © 2005 Data Connection Limited • Proprietary and Confidential • All Rights Reserved. 6

MPLS Resiliency – Protection Switching (1)

- Data switched from failed LSP to backup LSP at repair point (usually ingress)
- Backup LSP may be pre-provisioned or signaled upon failure (although backup route may have been pre-computed)

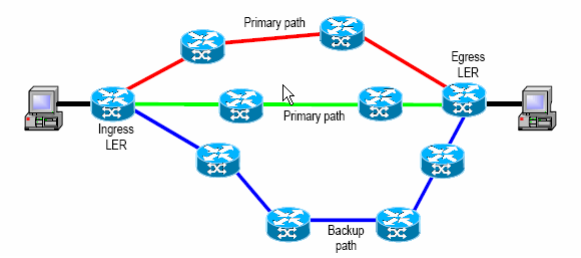

Copyright © 2005 Data Connection Limited • Proprietary and Confidential • All Rights Reserved.

MPLS Resiliency – Protection Switching (2)

- If the backup is pre-provisioned, the backup LSP
  - may already be carrying duplicate date (1+1 protection)
  - is idle and ready for immediate use (1:1 protection)
  - is in use carrying low-priority traffic which can be discarded (1:1 with extra traffic)
  - is pre-signaled and resources have been reserved; the resources are in use for low-priority traffic which can be discarded (1:1 without extra traffic)
  - Backs up multiple primary LSPs and is used for the first to fail (1:n protection)
- Main concern is speed of repair
  - All options require signaling – from point of detection to point of repair, or for the full LSP
  - For example – 10,000 LSPs failed over in 10 secs at a PLR

Copyright © 2005 Data Connection Limited • Proprietary and Confidential • All Rights Reserved.

MPLS Resiliency – Fast Reroute (FRR) (1)

- Draft-ietf-mpls-rsvp-lsp-fastreroute
- Fast Reroute establishes backup LSP tunnels for local repair

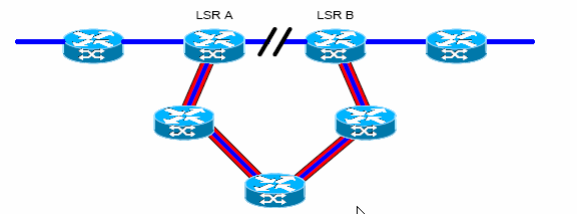

Copyright © 2005 Data Connection Limited • Proprietary and Confidential • All Rights Reserved.

MPLS Resiliency – Fast Reroute (FRR) (2)

- Detour
  - For each protected LSP, RSVP signals a detour LSP from each PLR (point of local repair) to a downstream MP (merge point)
- Facility (Bypass)
  - An independent bypass LSP is pre-provisioned by the management system between PLRs and MPs such that many LSPs can be switched onto the bypass LSP
- Fast Reroute
  - is designed for in-band signaling
  - requires label-stacking
  - is uni-directional only
  - graceful restart can be used to recover FRR LSPS

Copyright © 2005 Data Connection Limited • Proprietary and Confidential • All Rights Reserved.

MPLS Resiliency – Optical Recovery

- Protection/Segment Recovery (draft-ietf-ccamp-gmpls-segment-recovery)
  - Detour: Similar to FRR and only for in-band signaling
  - Dynamic Control (Bypass): Similar to FRR
  - Explicit
- End-to-end Recovery (draft-ietf-ccamp-gmpls-recovery-e2e-signaling)
  - Uni-directional 1+1 protection
  - Bi-directional 1+1 protection
  - 1:1 Dedicated Protection (with Extra Traffic)
  - Shared Mesh restoration
  - Full LSP restoration
- Both are still "works-in-progress" – e2e has been in progress since 2003

Copyright © 2005 Data Connection Limited • Proprietary and Confidential • All Rights Reserved.

MPLS Resiliency Approaches – Summary

- Multiple schemes
  - Some local (eg HA), some global (eg FRR)
  - Pro's and Con's for each
  - Suited to different requirements
- In practice
  - Many service providers ask for them all
  - Most equipment vendors have to provide them all
  - Protocol vendors have to do them all!
- Check out Data Connection's white paper…

  "Protection and Restoration in MPLS Networks" white paper at http://www.dataconnection.com/products/whitepapers.htm

Copyright © 2005 Data Connection Limited • Proprietary and Confidential • All Rights Reserved.