



LINX – A bit of history, and how a 10GE IXP evolves

Mike Hughes, CTO
mike@linx.net

April 2007

MENOG1

What is LINX?

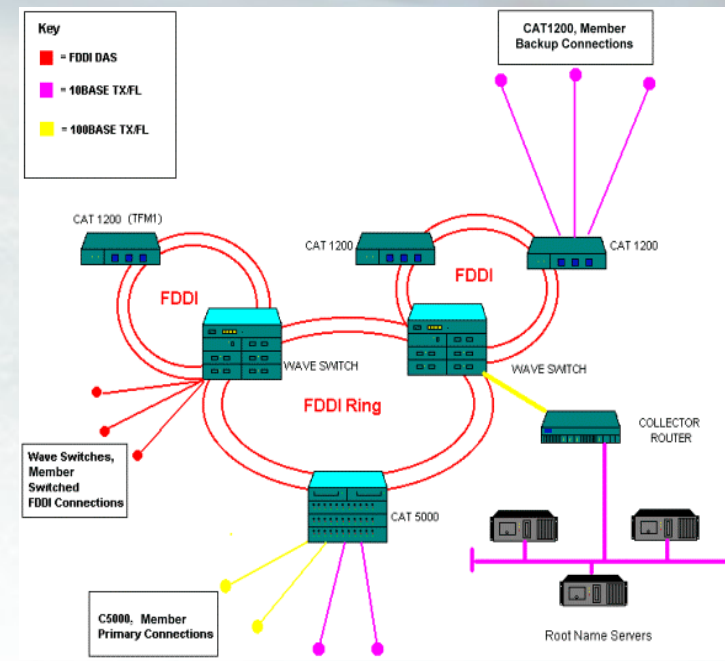
- The largest exchange in the London area
 - One of the largest exchanges globally
- A not-for-profit association of Internet companies
 - A full time staff (20 FTE) operates the exchange fabric
- We're not a co-lo, we just operate the switches – in other peoples co-location facilities
- We are “95th percentile” to quote Bill Woodcock
 - Not all exchanges (have to) look like this

Brief History of LINX

- Founded in 1994 by 5 ISPs from the UK
 - Pipex (the original “Pipex”, now MCI/Uunet)
 - Demon Internet
 - BTnet
 - UKERNA
 - EUnet GB (later PSInet, now Telstra UK)
- A switch (well 10Mb hub!) in Telehouse
 - Volunteer staff
- To stop UK traffic “elbowing” through US

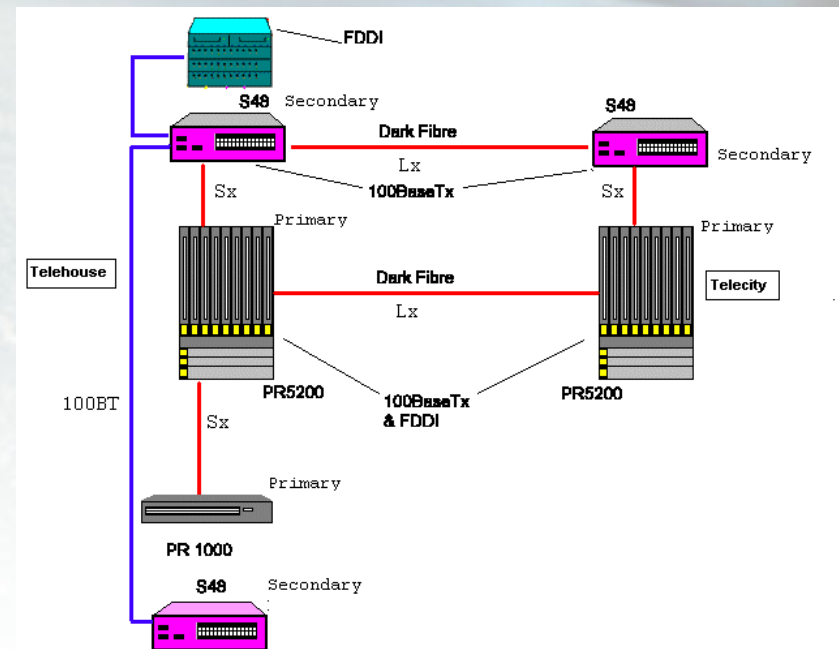
Architecture Development - 1996

- A FDDI ring based architecture
 - Cisco and Plaintree switches
 - FDDI, 100Mb TX and 10Mb connections
- 4 Full time staff



Architecture Development - 1998

- Gigabit Ethernet switches
 - First Metro GigE deployment in EU
- Multiple site IX
- Multiple vendor
 - Packet Engines
 - Extreme
- Broke the 1G mark in Nov 1999, 10 FTE



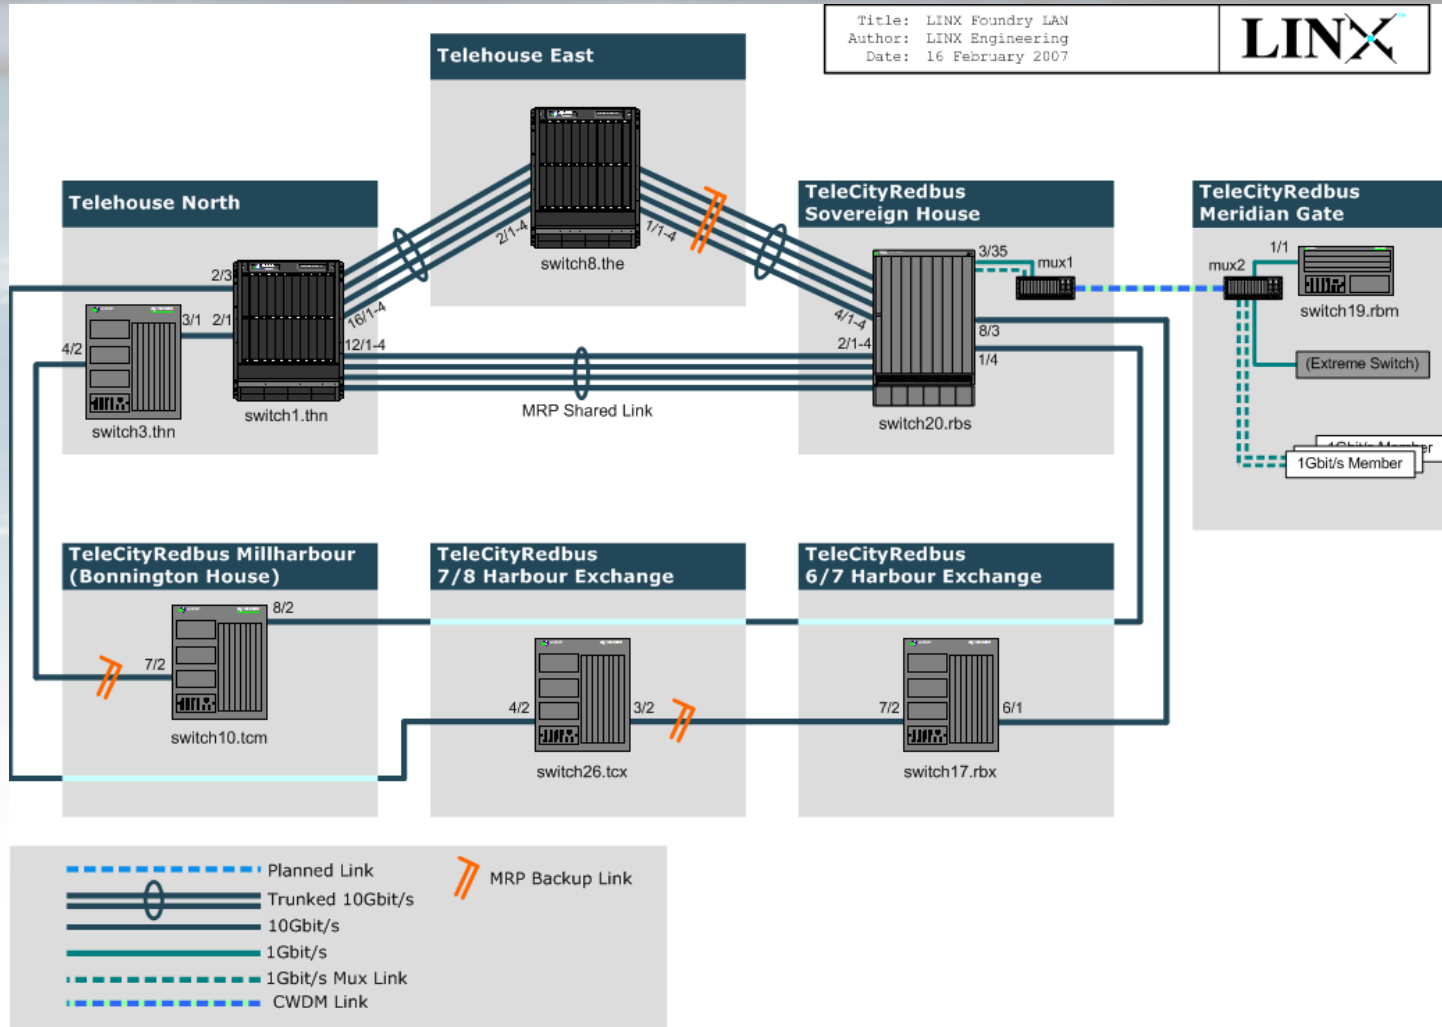
Cathartic Events in 2000

- There was an attempt to take LINX commercial in the wake of the initial boom
 - By existing LINX directors, with external backing/funding
- Member reaction – “LINX is not for sale!”
 - Concerns about LINX becoming open to capture
- Reaffirmed the mutual, not-for-profit model being the right thing for LINX

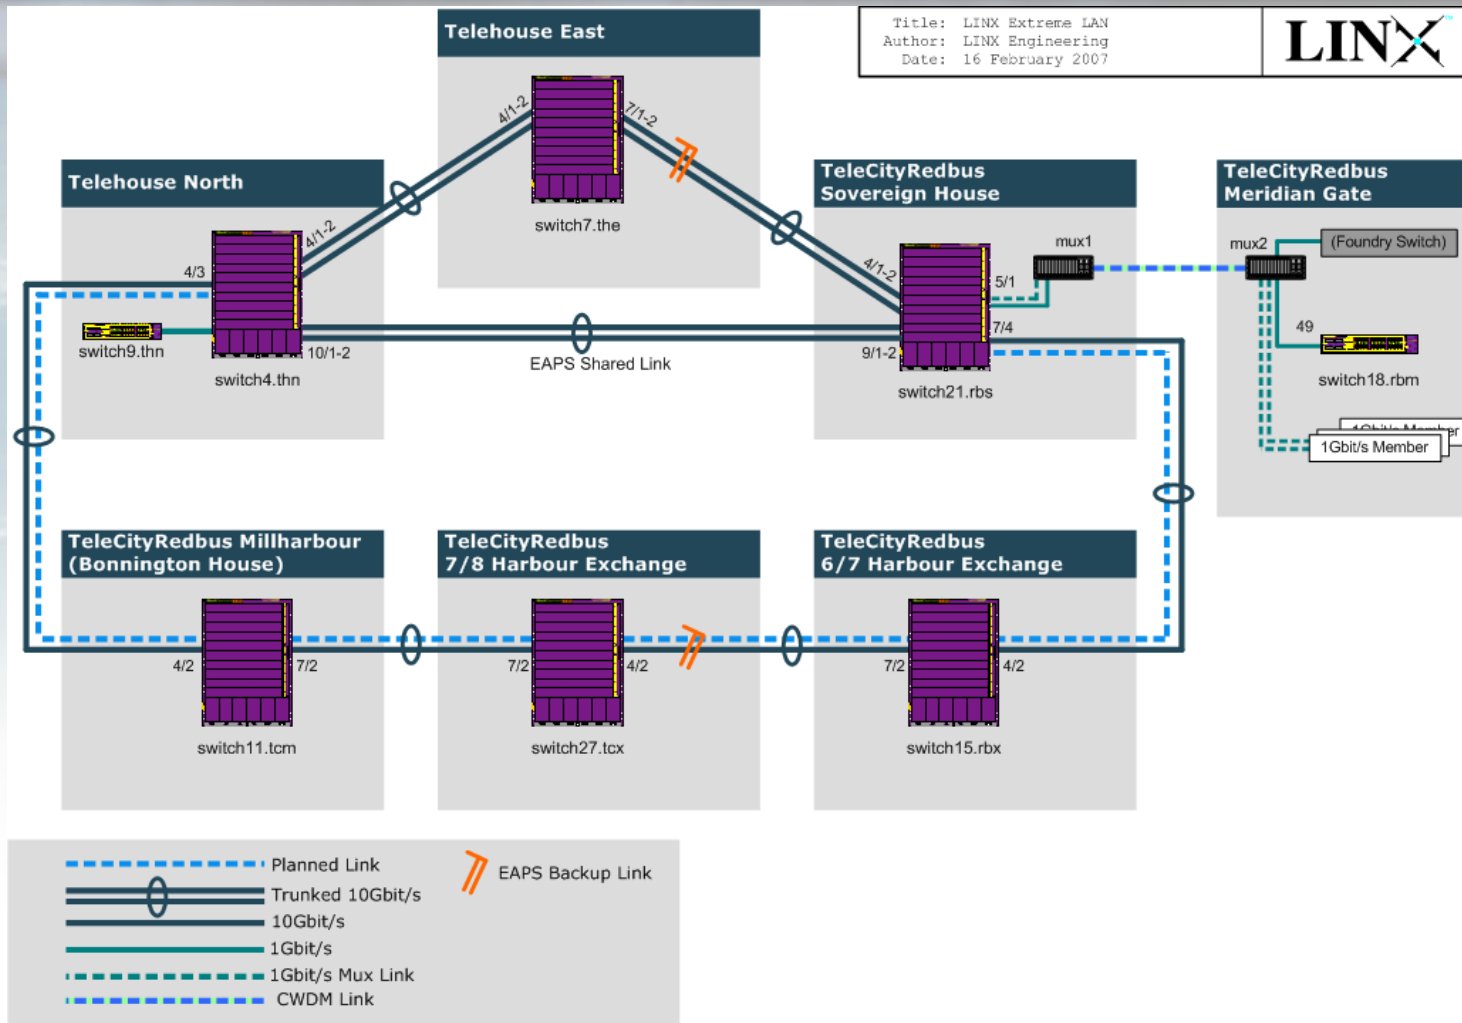
LINX Today

- Over 240 members from over 30 different countries
 - Members are legal entities with BGP clue
- Still strong UK contingent (about 50%)
- Most continents represented
- 21 Employees (20 FTE equivalent)
- 7 locations in London Docklands
- Dual LAN, Dual Vendor *nx*10G flat Ethernet network

LINX Network Diagram - Foundry




LINX Network Diagram - Extreme



Meeting the 10G Challenge

- LINX was a very early adopter of 10G
 - Foundry network first in late 2001
 - It just worked, out of the box!
 - Removed the need to buy WDM equipment
- That's been upgraded to *nx*10G in the backbone as traffic has grown
- But now networks want to attach to LINX at 10G
 - Presenting challenges for the backbone

10G Switches




April 2007

MENOG1

Interesting packet size datapoint

Packet Size Distribution at LINX



April 2007

MENOG1

Vendor Selection: What Matters?

- 10G port density
- 1G port density
- Uniform, predictable packet performance
 - Especially at smaller frame sizes!
- Important features
 - Particularly trunking/LACP
- High Availability
 - Hitless failover/upgrade, redundancy model

Challenges to come

- Scaling the network for multiple 10G connections from members
- 100G Standards process still slow
 - 100GE likely won't ship until 2010
 - Meaning $nx10G$ is best we can expect for now
- Being able to provide uniform service in multiple locations
- Potential for massive traffic growth...

Drivers

- General traffic growth
- Extra growth factors/drivers
 - Increase in broadband/broadband speed
 - VoIP is a red-herring, but is “sensitive” to jitter
- Faster ports
 - 10G ports, *nx*10G ports
 - 44 10G ports are have member connections
 - *nx*1G (link agg) ports

Influencers

- Cost of co-lo space
- Mergers/Acquisitions in sector
- New builds in London co-lo market
- Cost of dark fibre
- Geo-redundancy or “Telehouse effect”
- All the above may make members move their equipment and LINX connections
 - Requiring redeployment of LINX resources

Shorter Term

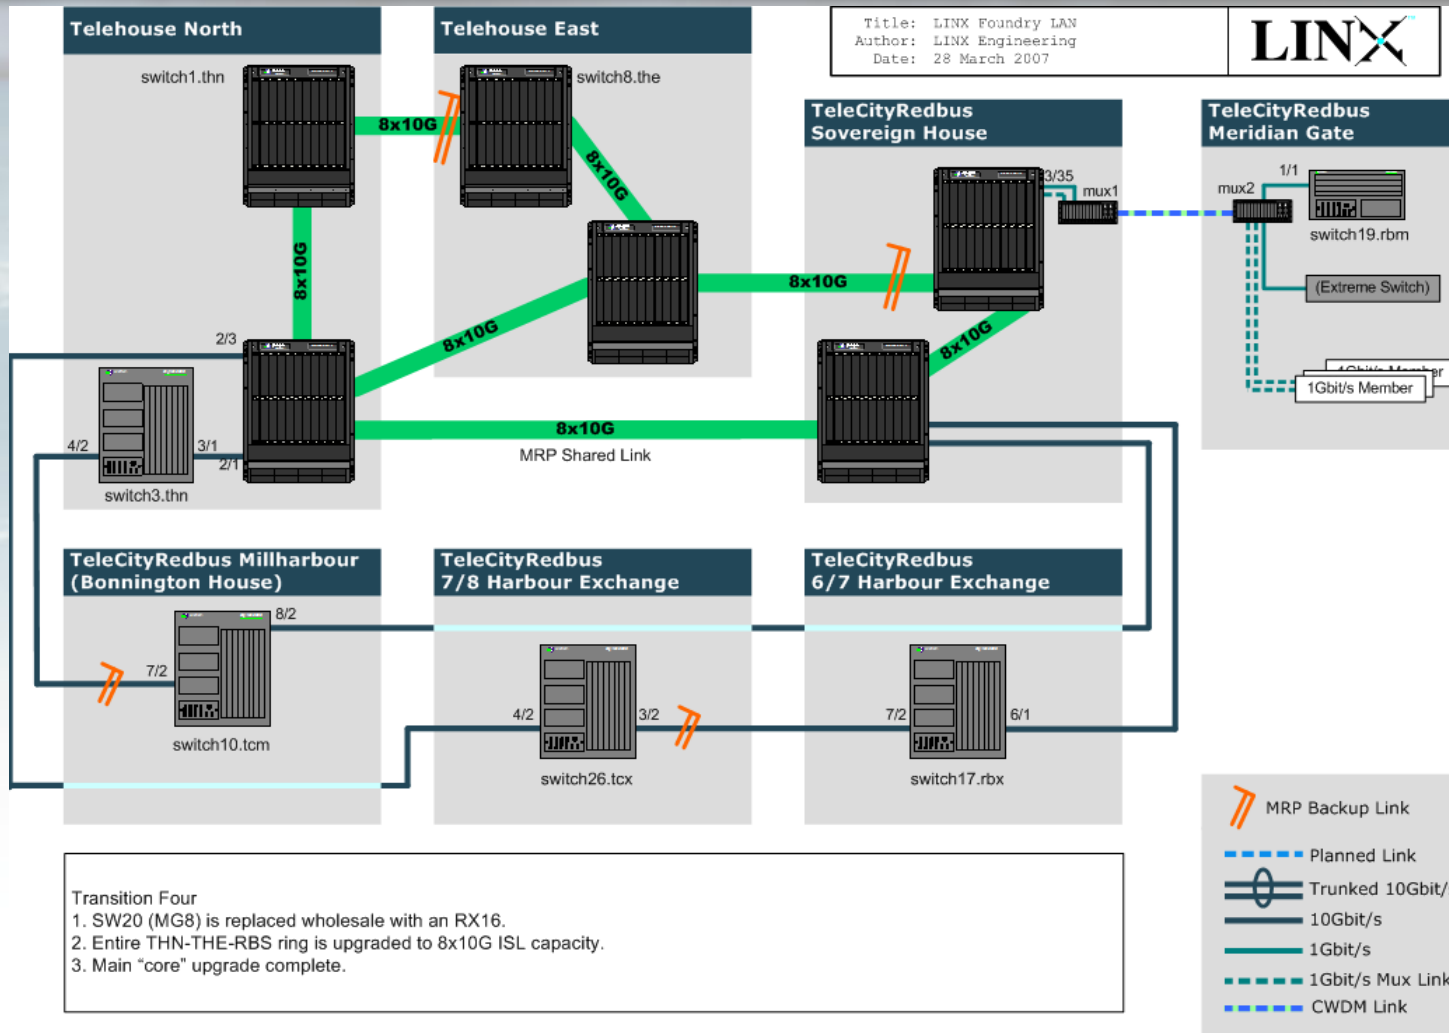
- Bigger switches and fatter Interswitch trunks can meet most needs
 - 10G connections have to be “concentrated”
 - But about 50% of a switch could easily be consumed by backbone connectivity
- Using some protocol enhancements from vendors
 - e.g. EAPSV2 and MRP phase 2
 - add multiple ring support

Foundry RX-series



- Double the density of the older MG8
- Up to 64 line-rate 10G ports per chassis
 - Biggest on the market today
 - Keeps traffic inside a single large box
- Will be an RX32 shipping shortly
 - The size of an entire cabinet!

Upgrade at planning stage



April 2007

MENOG1

Forward Looking

- Ethernet rings can have some problems
 - All nodes in one ring have to be (roughly) equal
 - Multiple rings solves most of this
 - Still constrained by max link speed/trunk size
- Is the Swedish model - unconnected “standalone” switches - a better way?
 - Backplane bandwidth is unrestricted/cheap
 - Some redundancy/resiliency challenges

Other Scalers

- Passive Private Interconnect
 - Fibre cross-connects to shed the largest flows
 - Cheap (for the IX), easy to implement
 - Can run whatever protocol the peers choose
- More exchanges
 - Could LINX run a third platform?
 - More smaller exchanges? Influences critical mass?
- “Transmission Only”
 - e.g. WDM platforms, stub-sites (no switch)

Traffic Management

- Enable normally blocked links
- MPLS
 - The DIX-IE (Tokyo) is involved in an MPLS interconnect – using conventional routing (ISIS) to route the network and LDP to discover endpoints – “mplsASSOCIO”
- “Smarter” L2 forwarding
 - IETF – TRILL/Rbridge – ISIS for L2 networks
 - IEEE – 802.1aq Shortest Path Bridging

So, until 100G comes along...

- With nothing faster than 10G in the short term, management of 10G member connections is going to be vital for bigger IXPs
 - Keep traffic local to the switch
- 100G progress is there, but slow
- Private Interconnect is a vital complement
- Totally new or revised topologies
 - To fit traffic profiles, for traffic management

An IXP doesn't have to look like LINX

- There is no “right way” to do this
 - Though co-operative, non-profit, neutral exchanges, that start inexpensively usually work well
 - The exchange may be operated by the co-location provider, or may be separate
- LINX is huge in terms of traffic, members and locations
 - One of about 6 similar exchanges globally
 - “95th percentile”
- Many successful smaller IXPs exist

Questions



April 2007

MENOG1