infinera

# Higher Speed Ethernet Requirements

**Drew Perkins**
**dperkins@infinera.com**
**408-572-5308**

**Ted Sprague**
**tsprague@infinera.com**
**408-572-5200**

**IEEE HSSG**
**November 14, 2006**

**IEEE HSSG**
**November 14, 2006**

# Internet Backbone Growth

- Industry consensus indicates a sustainable growth rate of 75% to 100% per year in aggregate traffic demand
- Traffic increased more than 10,000x from 1990 to 2000
- Traffic projected to increase an additional 1,000x from 2000 to 2010

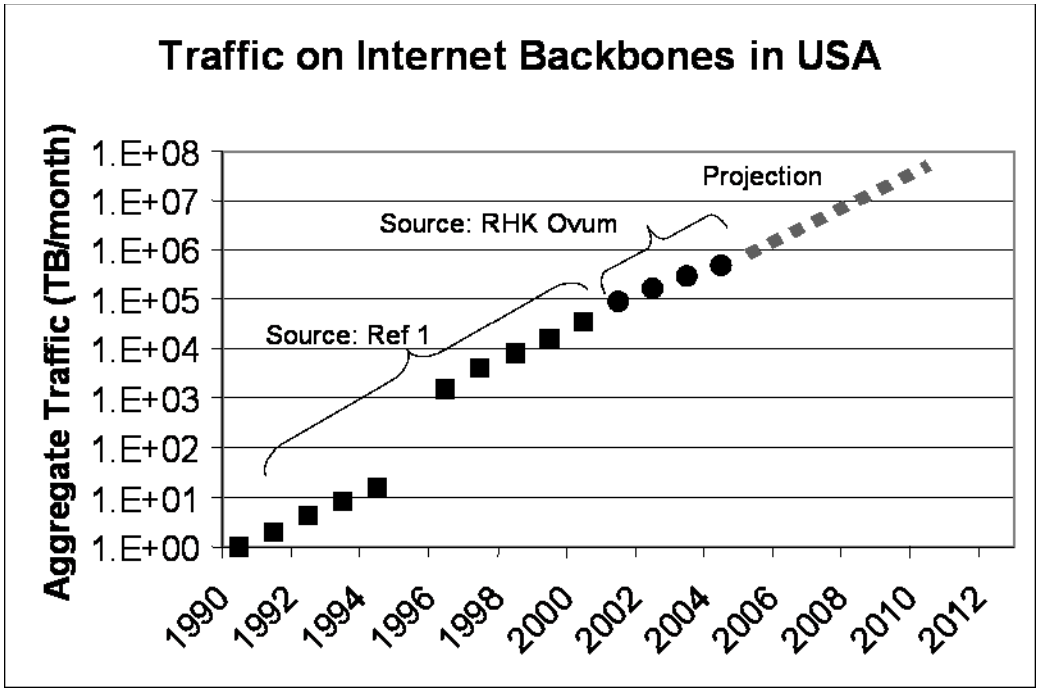


[1] K. G. Coffman and A. M. Odlyzko, 'Growth of the Internet', Optical Fiber Telecommunications IV B: Systems and Impairments, I. P. Kaminow and T. Li, eds. Academic Press, 2002, pp. 17-56.

# The Future Belongs to Tb/s Links!

- Carriers deploying Nx10 Gb/s networks
- Evaluating deployment of (Nx) 40 Gb/s router network
- This may be too little, too late
- Current Backbone growth rates, if sustained, will require IP link capacity to scale to > 1 Tb/s by 2010

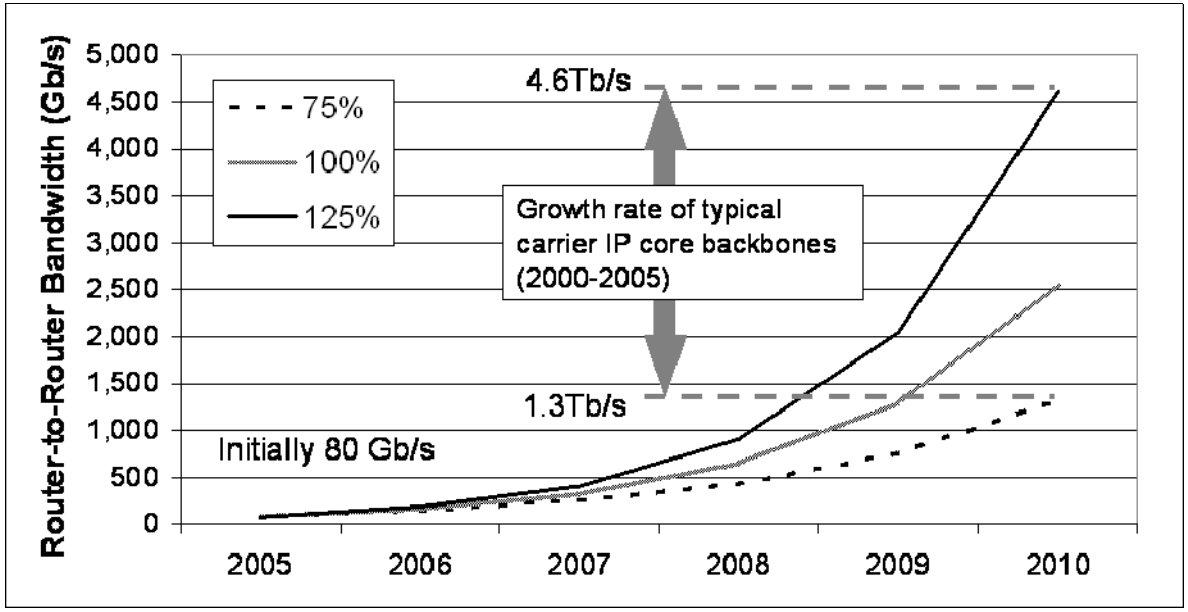

# Proposed Requirements for Higher Speed Ethernet

- Protocol Extensible for Speed
  - Ethernet tradition has been 10x scaling
  - At current growth rates, 100 Gb/s will be insufficient by 2010
  - Desirable to standardize method of extending available speed without re-engineering the protocol stack
- Incremental Growth
  - Most organizations deploy new technologies with a 4-5 yr lifetime
  - Pre-deploying based on the speed requirement 5 yrs in advance is economically burdensome
  - Assuming 5 yr window and 100% growth per year, ability to grow link speed incrementally over $2^5$ = 32x without a “forklift upgrade” desirable

# Economics of Incremental Growth (Less)

| | Cost/10G | Cost/100G | Notes |
|---|---|---|---|
| 10x10G Scalable | 10 | 100 | Relative costs |
| 1x100G Non-scalable | 8 | 80 | Given benefit of the doubt that it will be cheaper |

| | Annual | Quarterly | Notes |
|---|---|---|---|
| Growth Rate | 75% | 15.02% | |
| Discount Rate (Low) | 12.50% | 2.99% | Typical cost of capital for most incumbent carriers |
| Discount Rate (Mid) | 25.00% | 5.74% | Higher cost of capital for many competitive carriers |
| Discount Rate (High) | 50.00% | 10.67% | Very high cost of capital for startup carriers |

Total Cost

| **Month** | **0** | **3** | **6** | **9** | **12** | **15** | **18** | **21** | **24** | **27** | **30** | **33** | **36** | **39** | **42** | **45** | **48** | **51** | **54** | **57** | **60** |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Bandwidth Required (Gb/s) | 10 | 12 | 13 | 15 | 18 | 20 | 23 | 27 | 31 | 35 | 41 | 47 | 54 | 62 | 71 | 82 | 94 | 108 | 124 | 143 | 164 |
| 10x10G Scalable | 10 | 20 | 20 | 20 | 20 | 30 | 30 | 30 | 40 | 40 | 50 | 50 | 60 | 70 | 80 | 90 | 100 | 110 | 130 | 150 | 170 |
| Incremental CF | 10 | 10 | 0 | 0 | 0 | 10 | 0 | 0 | 10 | 0 | 10 | 0 | 10 | 10 | 10 | 10 | 10 | 10 | 20 | 20 | 20 |
| Net Present Cost (Low) | 10 | 20 | 20 | 20 | 19 | 27 | 28 | 28 | 36 | 36 | 43 | 43 | 50 | 57 | 64 | 70 | 76 | 82 | 93 | 105 | 115 |
| Net Present Cost (Mid) | 10 | 19 | 19 | 19 | 19 | 27 | 27 | 27 | 33 | 33 | 38 | 38 | 43 | 48 | 52 | 56 | 60 | 63 | 70 | 76 | 82 |
| Net Present Cost (High) | 10 | 19 | 19 | 19 | 19 | 24 | 24 | 24 | 28 | 28 | 31 | 31 | 34 | 36 | 38 | 40 | 41 | 43 | 45 | 47 | 49 |
| 1x100G Not-scalable | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 160 | 160 | 160 | 160 |
| Incremental CF | 80 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 80 | 0 | 0 | 0 |
| Net Present Cost (Low) | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 127 | 127 | 127 | 127 |
| Net Present Cost (Mid) | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 109 | 109 | 109 | 109 |
| Net Present Cost (High) | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 80 | 91 | 91 | 91 | 91 |
| 1x100G vs. 10x10G Ratio (Low) | 800% | 406% | 406% | 406% | 419% | 292% | 283% | 283% | 222% | 222% | 184% | 184% | 159% | 140% | 126% | 115% | 105% | 156% | 136% | 122% | 110% |
| 1x100G vs. 10x10G Ratio (Mid) | 800% | 412% | 412% | 412% | 412% | 299% | 299% | 299% | 243% | 243% | 208% | 208% | 185% | 167% | 154% | 143% | 134% | 171% | 155% | 142% | 132% |
| 1x100G vs. 10x10G Ratio (High) | 800% | 424% | 424% | 424% | 424% | 327% | 327% | 327% | 282% | 282% | 255% | 255% | 236% | 222% | 211% | 202% | 194% | 214% | 202% | 193% | 186% |

# Economics of Incremental Growth (More)

| | Cost/10G | Cost/100G | Notes |
|---|---|---|---|
| 10x10G Scalable | 10 | 100 | Relative costs |
| 1x100G Non-scalable | 12 | 120 | Model it as more expensive |

| | Annual | Quarterly | Notes |
|---|---|---|---|
| Growth Rate | 75% | 15.02% | |
| Discount Rate (Low) | 12.50% | 2.99% | Typical cost of capital for most incumbent carriers |
| Discount Rate (Mid) | 25.00% | 5.74% | Highre cost of capital for many competitive carriers |
| Discount Rate (High) | 50.00% | 10.67% | Very high cost of capital for startup carriers |

Total Cost

| **Month** | **0** | **3** | **6** | **9** | **12** | **15** | **18** | **21** | **24** | **27** | **30** | **33** | **36** | **39** | **42** | **45** | **48** | **51** | **54** | **57** | **60** |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Bandwidth Required (Gb/s) | 10 | 12 | 13 | 15 | 18 | 20 | 23 | 27 | 31 | 35 | 41 | 47 | 54 | 62 | 71 | 82 | 94 | 108 | 124 | 143 | 164 |
| 10x10G Scalable | 10 | 20 | 20 | 20 | 20 | 30 | 30 | 30 | 40 | 40 | 50 | 50 | 60 | 70 | 80 | 90 | 100 | 110 | 130 | 150 | 170 |
| Incremental CF | 10 | 10 | 0 | 0 | 0 | 10 | 0 | 0 | 10 | 0 | 10 | 0 | 10 | 10 | 10 | 10 | 10 | 10 | 20 | 20 | 20 |
| Net Present Cost (Low) | 10 | 20 | 20 | 20 | 19 | 27 | 28 | 28 | 36 | 36 | 43 | 43 | 50 | 57 | 64 | 70 | 76 | 82 | 93 | 105 | 115 |
| Net Present Cost (Mid) | 10 | 19 | 19 | 19 | 19 | 27 | 27 | 27 | 33 | 33 | 38 | 38 | 43 | 48 | 52 | 56 | 60 | 63 | 70 | 76 | 82 |
| Net Present Cost (High) | 10 | 19 | 19 | 19 | 19 | 24 | 24 | 24 | 28 | 28 | 31 | 31 | 34 | 36 | 38 | 40 | 41 | 43 | 45 | 47 | 49 |
| 1x100G Not-scalable | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 240 | 240 | 240 | 240 |
| Incremental CF | 120 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 120 | 0 | 0 | 0 |
| Net Present Cost (Low) | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 191 | 191 | 191 | 191 |
| Net Present Cost (Mid) | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 163 | 163 | 163 | 163 |
| Net Present Cost (High) | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 120 | 136 | 136 | 136 | 136 |
| 1x100G vs. 10x10G Ratio (Low) | 1200% | 609% | 609% | 609% | 628% | 438% | 424% | 424% | 333% | 333% | 276% | 276% | 238% | 210% | 189% | 172% | 158% | 233% | 205% | 183% | 166% |
| 1x100G vs. 10x10G Ratio (Mid) | 1200% | 618% | 618% | 618% | 618% | 448% | 448% | 448% | 364% | 364% | 312% | 312% | 278% | 251% | 230% | 214% | 200% | 257% | 232% | 213% | 197% |
| 1x100G vs. 10x10G Ratio (High) | 1200% | 635% | 635% | 635% | 635% | 491% | 491% | 491% | 423% | 423% | 382% | 382% | 355% | 333% | 316% | 303% | 292% | 320% | 303% | 290% | 278% |

# Proposed Requirements (cont'd)

- Hitless Growth
  - Problematic to "take down" core router links for a substantial period of time without customer service degradations
  - SLAs may be compromised or require complicated temporary workarounds if substantial down time is required for upgrade.
  - Ideally, upgrade of the link capacity should therefore be hitless, or at least only momentarily service-impacting.
- Resiliency and Graceful Degradation
  - Protocol should provide rapid recovery from failure of an individual channel or component
  - If the failure is such that full performance can not be provided, degradation should only be proportional to the failed element(s).

# Proposed Requirements (cont'd)

- WAN Transportability
  - The protocol should be compatible with WAN optical transmission technology and deployed topologies in a manner that is:
    - Cost-effective
    - Reliable
    - Operationally Manageable
  - These characteristics should be sustained across networks having regional, national and inter-continental reach
  - The protocol should be resilient to intra-channel/intra-wavelength propagation delay differences (skew)
  - Support for AIS, FDI, BDI and TCM mechanisms

# Proposed Requirements (cont'd)

- Technology Reuse
  - Highly desirable to leverage existing 10G PHYs, including 10GBASE-R, W, X, S, L, E and LRM in order to foster ubiquity and avoid duplication of standards efforts
- Deterministic Performance
  - Latency/Delay Variation should be low for support of real-time packet based services
    - Streaming video
    - VOIP
    - Gaming

# Drivers for a Super-λ (Multi-wavelength) Protocol

- Per-channel bit rate growth historically and dramatically out-paced by Core Router interconnection demand growth

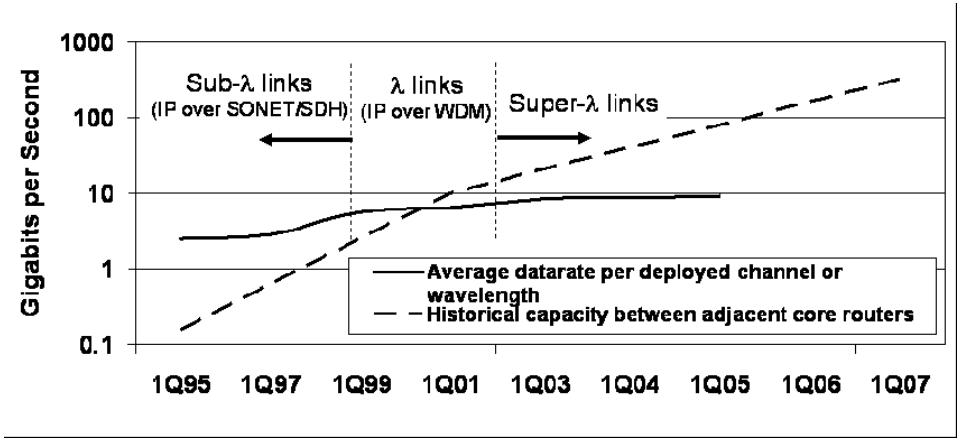


- Requirement for compatibility with WAN transport strongly favors approach leveraging multiple wavelengths (Super-λ)

# Possible Channel Bonding Techniques

- Limitations of LAG/ECMP derived from fixed assignment of “conversations” to channels
- Traffic may be distributed over multiple links by a variety of techniques
  - Bit/Octet/Word Distribution
    - In a manner similar to 10GBASE-X, increments of the serial stream are assigned sequentially to lanes
    - Minimal additional overhead required to allow re-alignment at the receiver
  - Packet Distribution
    - Sequence numbers added to packets to enable re-ordering at the receiver
    - Large packets within the stream may induce excessive delay/delay variation to smaller, latency-sensitive packets
  - Packet Distribution with Fragmentation
    - Fragmentation bounds buffering requirements and delay associated with packet size and packet size variation
    - Overhead/link inefficiency is a function of the maximum fragment size chosen
    - At 100 Gb/s and above, a fragment size can be chosen such that an effective compromise between link efficiency and the QoS of individual, time-sensitive flows can be readily achieved

# Summary

- HSE Requirements
  - Incremental Growth
  - Protocol Extensible for Speed
  - Hitless Growth
  - Resiliency and Graceful Degradation
  - WAN Transportability
  - Technology Reuse
  - Deterministic Performance
  - Multi-channel Operation

infinera

# Thanks!

**Drew Perkins**
**dperkins@infinera.com**
**408-572-5308**

**Ted Sprague**
**tsprague@infinera.com**
**408-572-5200**