Does Anybody Really Need Open Routers?

Jon Turner
www.arl.wustl.edu/~jst

Good News and Bad News

- Good news – good hardware is available
  - »ATCA has made world-class router hardware available
    - can buy all the parts needed for 120 Gb/s router today
    - sure, the parts cost $150K but commercial routers in this class aren't cheap either
  - »tools for NP-based forwarding engines are very good, making software development straightforward
- Bad news – it takes software too
  - »nobody wants them badly enough to invest effort
  - »system vendors definitely don't want them
  - »research community would like them, but not badly enough to learn to program NPs
  - »network operators would like them, but not if they have to pay for *open* software development

Shifting the Playing Field

- IP ossification is not a technical problem
  - »>10,000 ISPs with competing interests who must cooperate to enable change
- Overlay approach offers only credible way to overcome IP limitations
  - »solid commercial examples of successful use
  - »single organization can make change happen
- To enable rapid deployment of new overlays, need commercial overlay hosting services
  - »allows small organization with bright idea to deploy service globally without huge infrastructure investments
  - »Planetlab has shown us the potential
  - »need more capable platforms that can handle internet-scale traffic with router-like performance

Overlay Hosting Service

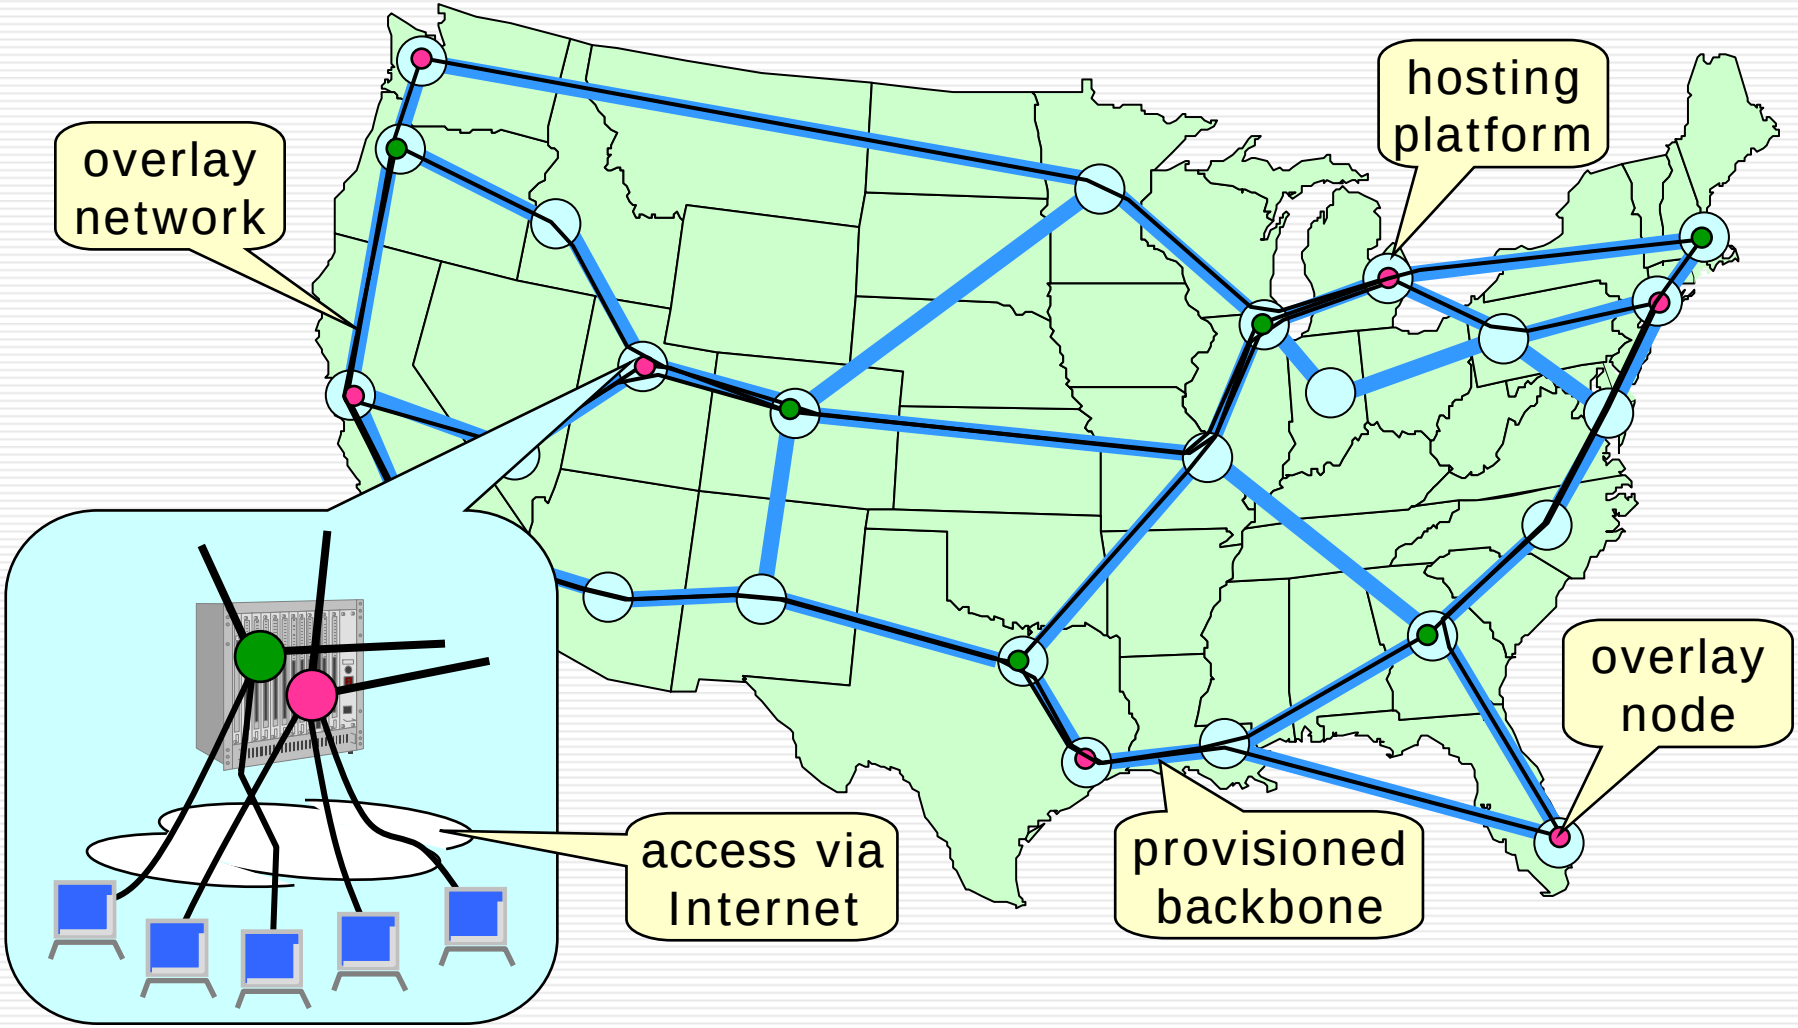

- Flexible platforms shared by multiple overlays
- Provisioned backbone, internet for access

Overlay Hosting Platform

- Processing Engines (PEs) implement overlay nodes
  - »GPE – conventional server blade
  - »NPE – network processor blade
    - nearly 4 Mp/s per NP vs 50 Kp/s
    - 100 μs latency vs. 1-300 ms
  - »shared or dedicated

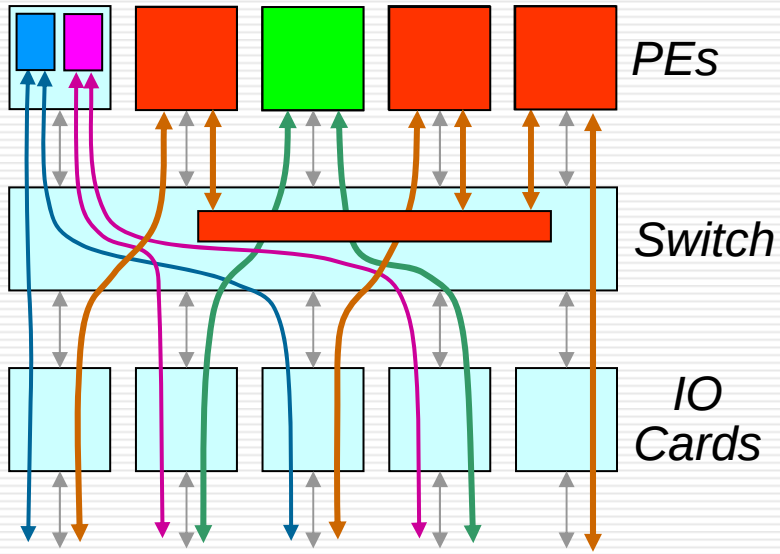

- IO Cards terminate external links, mux/demux streams
- Shared PEs managed by *substrate*
- Dedicated PEs may be fully controlled by overlay
  - »switch and IO Cards provide protection and isolation
- PEs in larger overlay nodes linked by logical switch
  - »allows scaling up for higher throughput

Current Development Platform

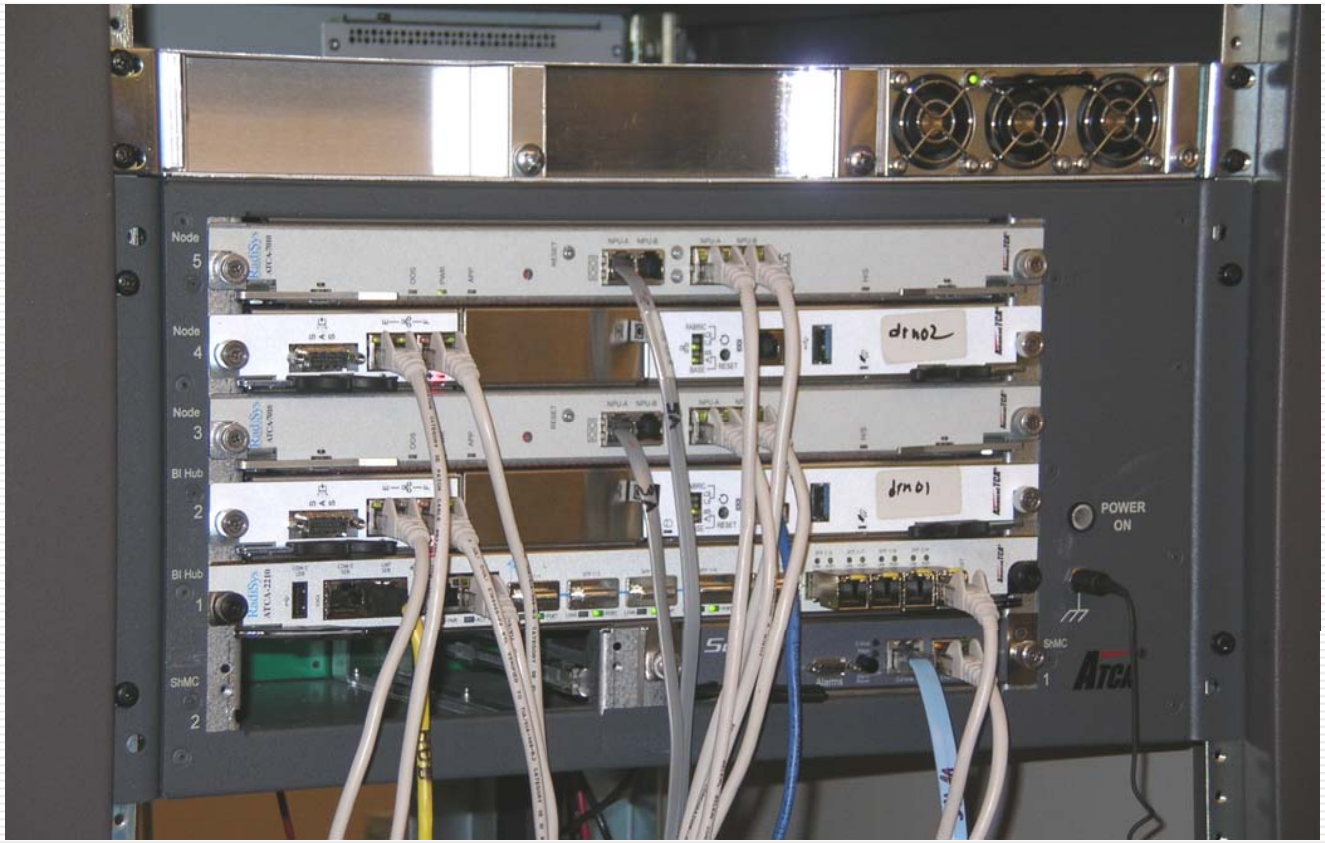

- Scale up using 14 slot chassis
- Multi-chassis configurations also possible

ATCA Boards

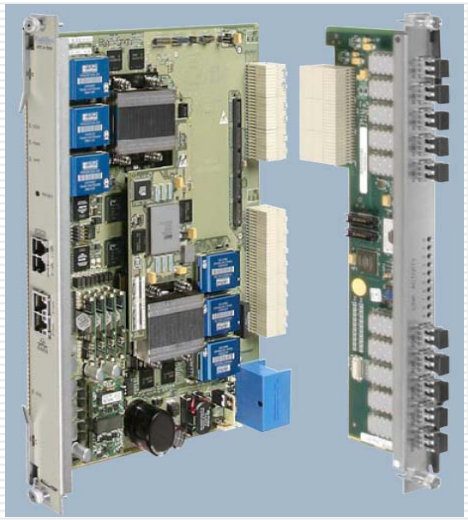

- Intel server blades
  - for CP and GPE
  - dual Xeons (2 GHz)
  - 4x1GbE
  - on-board disk
  - Advanced Mezzanine Card slot
- Radisys NP blades
  - for LC and NPE
  - dual IXP 2850 NPs
    - 3xRDRAM
    - 4xSRAM
    - shared TCAM
  - 2x10GbE to backplane
  - 10x1GbE external IO (or 1x10GbE)
- Radisys switch blade
  - up to 16 slot chassis
  - 10 GbE fabric switch
  - 1 GbE control switch
  - full VLAN support
- Scaling up
  - 5x10 GbE to front
  - 2 more to back

What You Need to Build Your Own

Qty	Description	Supplier	Model
1	Dual Network Processor Module with IO	Radisys	A7K-PPM10-CFG002
2	Dual Network Processor Module		A7010-BASE-2855
2	18 MB IDT TCAM Module		A7010-TCAM-01-R
3	10 Gb/s Fabric Interface Card		A7010-FIC-2X10G
1	10 GE/1GE Switch & Control Module		A2210-SWH-CFG-01
1	RTM with extra IO ports		A5010-SPM-01
5	1GE plugin modules (4 per kit)		A2K-SFP-C
2	Server blade with 2 dual-core Xeon processors	Intel	MPCBL004N01Q
1	Zephyr 6 Slot ATCA Shelf	Schroff	ZR5ATC6TMDPEM2N
1	Shelf Manager		21593-375
1	Alarm Board		ISAP2
1	1U Power Supply Shelf	Unipower	TPCPR1U3B
1	48 Vdc/25A Power Supply		TPCP7000
1	115 Vac/15A Power Cord		364-1409-0000