NASA Goddard
High End Computer Network (HECN) Group

Force10 10GigE Eval

- Testbed/Prototyping/Collaboration
- Production Applicability & Technology Transfer
- Interoperability
- Security Considerations

HECN 10GigE Evaluation Criteria

- Jumbo Frame capable (improved performance and reduced overhead)
- Backplane/Stack (right size switch - avoid bottlenecks and waste of ports)
- Non-blocking or at least line-rate capable (minimize bottlenecks)
- Oversubscription (improved density at a cost?)
- Link Aggregation (economically increase bandwidth of link) efficiency
- VLAN capable (logical separation of networks/LANs)
- Line rate ACLs (for security reasons)

Continued…

HECN 10GigE Evaluation Criteria

- Interoperability (connectivity/compatibility with existing equipment)
- Configuration management (ease of use)
- Flexibility - options (e.g., stackable, port density, format, mix, ...)
- QoS/CoS capable (ability to favor or limit throughput)
- OSPF and BGP (for building or WAN routers)
- Latency (e.g., for clustering over GigE)
- SSH, SSL (for security)
- VLAN Stack (allow separating groups of networks/LANs on a switch)

Continued…

HECN 10GigE Evaluation Criteria

- Redundancy (various hardware and software, for failover/uptime)
- Mirroring (troubleshooting/security)
- SFLOW (monitoring/security)
- SNMP (monitoring/management)
- Price/Performance/Features
- Future: IPV6, MPLS, Multicast (MBGP/MSDP, ...)

HECN TestBench - Hardware

- PowerMAC G4 (867 MHz, Dual 1 GHz),
  G5 (Dual 2 GHz)
- PC Xeon (Dual 3.06 GHz)
- GigE Adapters: NetGear GA620T, Intel Pro/1000 Server and Quad port
- 10GigE Adapters: Intel Pro/10GbE LR Server Adapter
- Future - 64 bit processors with PCI-X 2.0 (266 MHz) bus(es).
- Future - PCI-X 2.0 10 GigE Adapters from Intel and S2IO

HECN TestBench - Software

- Standard networking tools (ping, traceroute, ...)
- nuttcp (memory to memory transfers)
- NISTnet (network delay/bandwidth/loss emulator)
- tcpdump (network sniffer)
- MRTG to graph utilization
- Apple PowerMac - Yellow Dog Linux
- PC - Red Hat Linux

HECN TestBench - Methodology

- Layout the design for the test and configure for stress testing switch
- Emphasize use of real network data flows
- Run a standard set of tests
- Verify results (nuttcp, MRTG, ...) and look for anomalies
- Rerun more specialized tests to resolve anomalies
- Layout design for testing other features, configure/test/verify/analyze

Completed Force10 E300 Testing

- Throughput testing
- Link Aggregation (two 10GigE links, similar physical and VLAN as above)
- Initial QoS testing (port based)
- VLAN Stacking

Force10 10GigE Testing

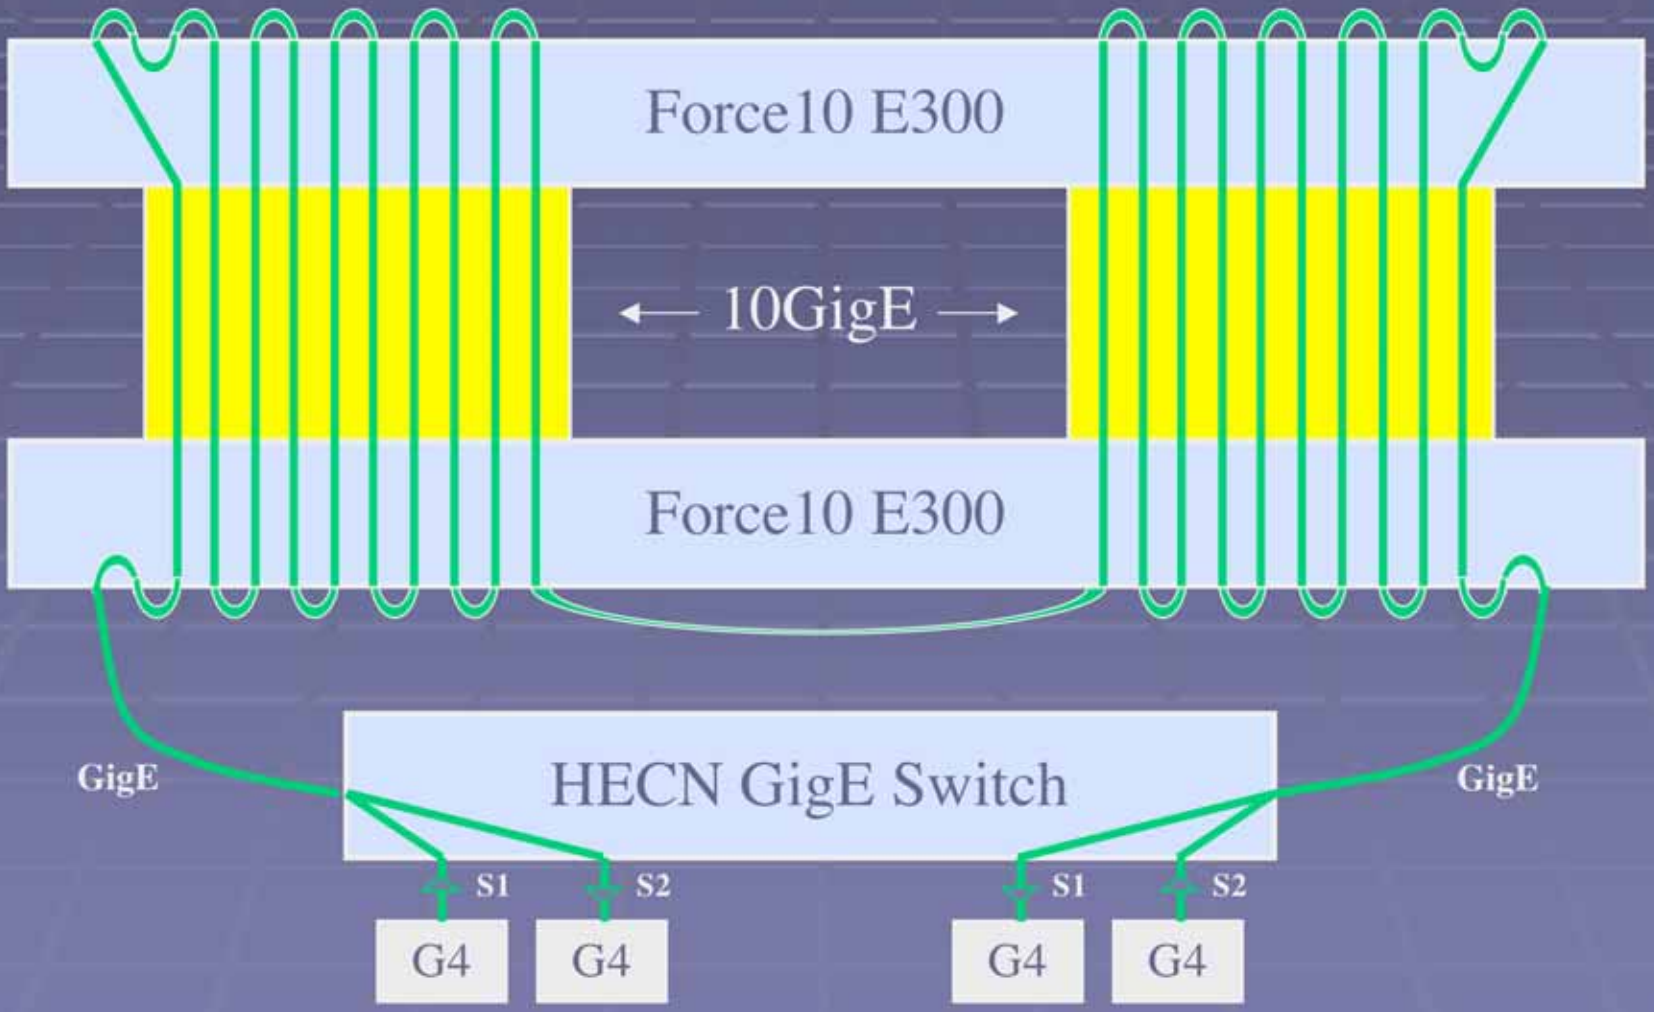

Test data generated with nuttcp and VLAN looping

NASA Goddard
BF/PL 3/21/04

Force10 10GigE nuttcp Stress Test (Over 1 Petabyte in 3 days *)

S1: porthos (2x800 MHz PowerMac G4) -> clifford (867 MHz PowerMac G4)
S2: underdog (867 MHz PowerMac G4) -> bigdog (2x1 GHz G4 XServe)

All systems running YellowDog Linux with Intel Pro/1000 Server Adapter or NetGear GA620T NIC, and using a 9K Jumbo Frame MTU

```
clifford% nuttcp -u -r -T72h -w2048 -Iporthos2clifford porthos & \
          nuttcp -u -T72h -w2048 -Iunderdog2bigdog underdog bigdog
underdog2bigdog:  30687801.2734 MB / 259772.86 sec =  990.9732 Mbps 74 %TX 29 %RX
                  1950222 / 3929988785 drop/pkt 0.05 %loss
porthos2clifford: 30695954.5781 MB / 259773.58 sec =  991.2337 Mbps 65 %TX 17 %RX
                  6056720340 / 9985802526 drop/pkt 60.65 %loss
```

(30687801.2734 MB + 30695954.5781 MB)*10*2/1024/1024/1024 = 1.14336 PB

* Over 0.25 PB transferred bidirectionally across each 10GigE link

Force10 Petabyte Challenge

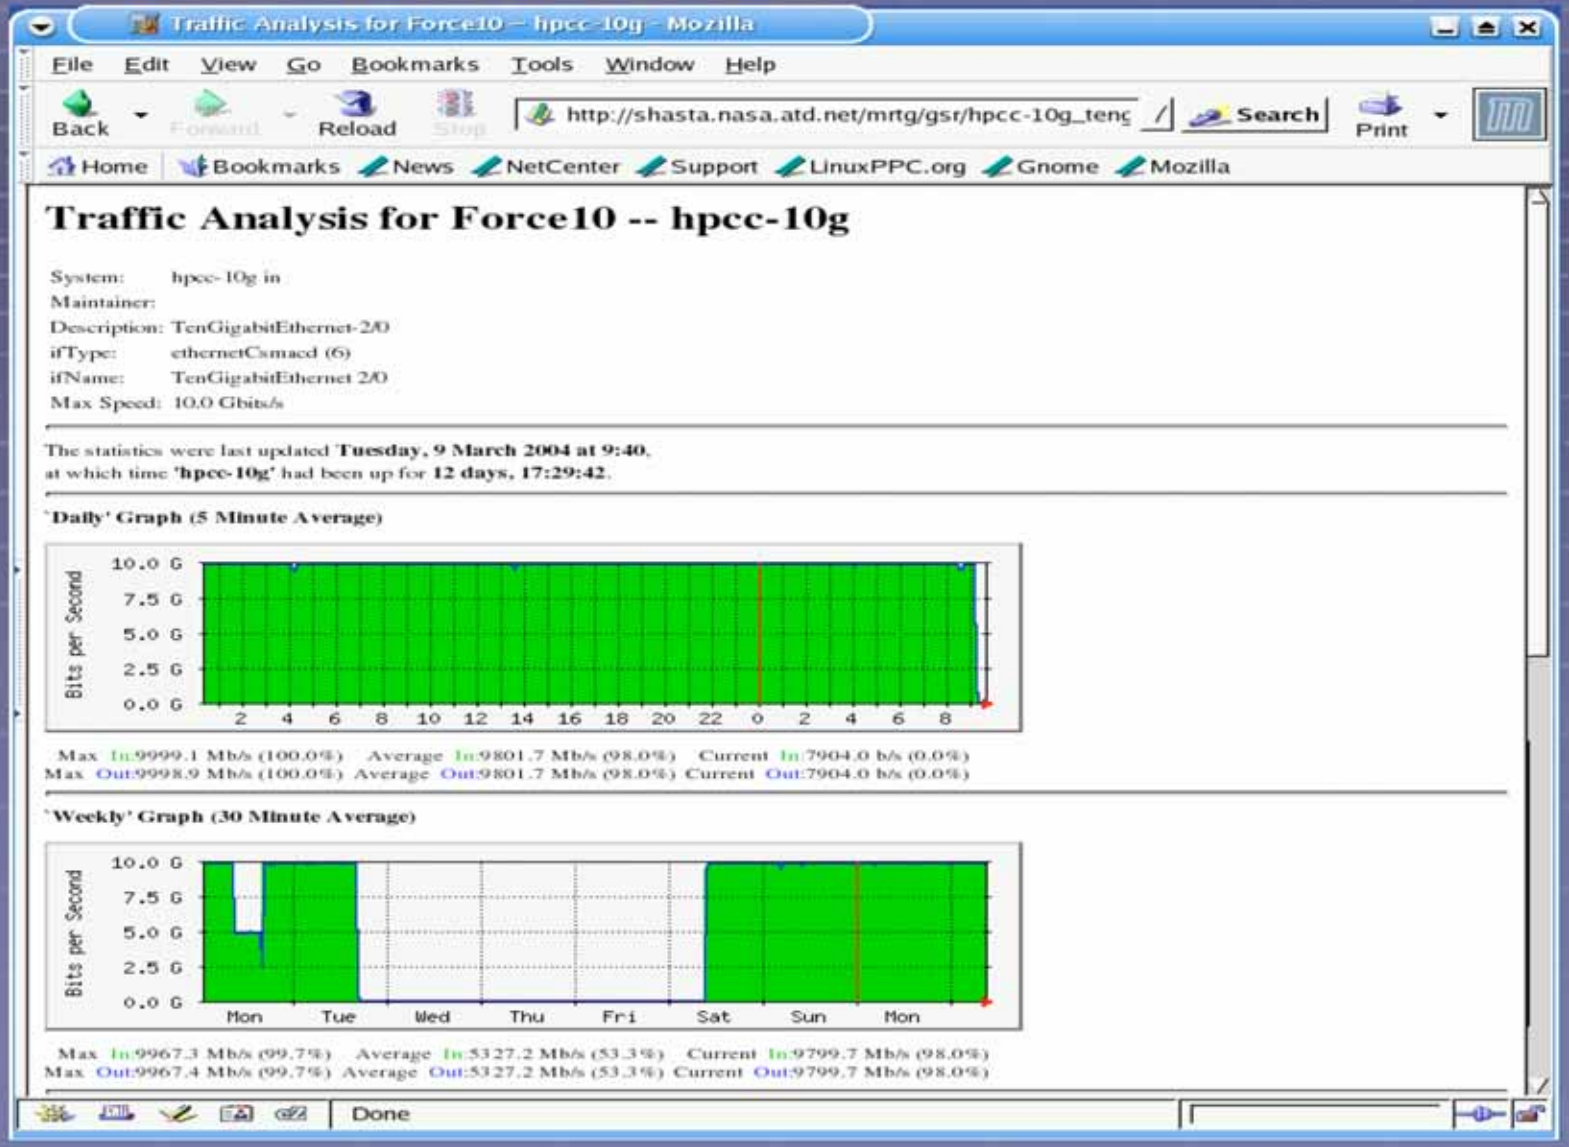

Graph for TenGigabitEthernet-3/0 basically identical

Planned Testing

Near Term Tests to complete:

- Multiple 10 GigE testing Between building
- 10GigE MUX'd over WDM link to MAX
- 10GigE Server Adapters (in our G5s) transferring through the Force10
- Compatibility/Interoperability with other switches in our network
- L3 performance
- E600
- Multicast
- Clustering over GigE
- Additional QoS Testing

Continued…

Planned Testing

Future Tests:

- IPV6
- MPLS
- Multicast protocols (MBGP/MSDP)

Additional Comments

Some other beneficial features…

- Configuration:
  - "interface range" command
  - "show vlan" shows ports and type (tagged, untagged/ member)
  - "show config" while configuring an interface
  - ease of configuring port based QOS
  - descriptive error messages
- Performance:
  - "IP 5-tuple" for Link Aggregation load balancing