

perfs--NAR

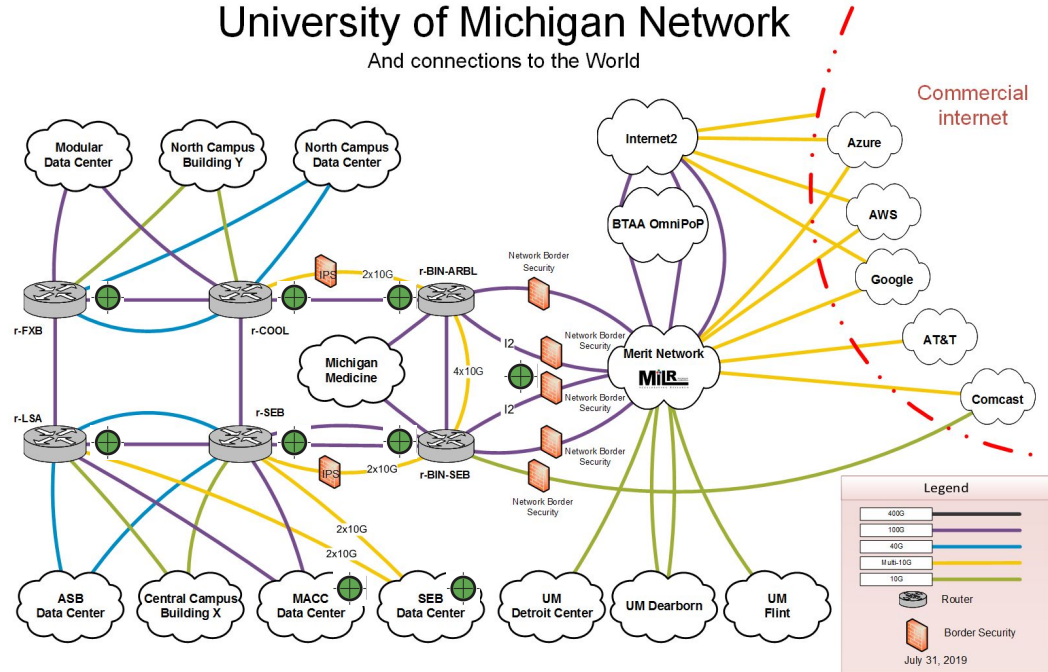


# Background and History

- Core network 100GE
- Many campus 10GE and 1GE links
- SR and LR fiber links
- 10GE testing for campus backbone and gateways
- Static perfSONAR testpoints on server-class rackmount hardware in data centers
- Inconsistent departmental perfSONAR deployments
- Service Degradation reports


## University of Michigan Network

And connections to the World



# Disclaimers

- perfSONAR **does not** endorse or promote any hardware platform or vendor
- Lessons and techniques covered in this talk are anecdotal: YMMV
- If you break it, you've bought both pieces



# History: Dell R330

## Physical Specifications:

- 1U server, 4 post rack
- 1.7" x 17" x 26.7"
- 30 lbs / 14kg
- 2 x 1GE & 1GE OOB Management
- 2 x SFP+
- Optional Dual Power Supplies
- Powerful Cooling Fans (loud)
- ~\$2,700 without optics

*Good for semi-permanent data center deployments*

## Not great:

- On a professor's desk
- Basement of a 100 year old campus building



# Need for Portability

## Portable Requirements:

- Form Factor (Size / Weight)
- Cost
- 10GE Networking
- 1GE Networking
- perfSONAR Supported OS:
  - CentOS 7 or Ubuntu 18

## Portable Niceties:

- OOB Management
- PoE
- Optional Rack Mount
- Expandability
- x86



# SuperMicro E300-8D Specs

- 10" x 9" x 1.7"
- 7.5 lbs. / 3.4 kg
- 2 x SFP+
- 6 x 1GE
- 1GE OOB Management Port
- Intel Xeon D-1518 (4 core @ 2.2GHz)
- 8GB RAM
- 500GB NVMe M.2 SSD
- Optional 2 Post Rack Ears
- ~\$900 without optics
- PCI-E Expansion options:
  - 4 x 1GE
  - 4 x SFP+
  - 2 x 10GE R45
  - 2 x 25GE




# SuperMicro Bootstrapping

- Crash cart (VGA, USB keyboard)
- Firmware settings (UEFI for NVMe)
- OOB Management (IPMI)
- Install OS
- Configure network
- Install perfSONAR with Ansible
- perfSONAR limits file (optional)



# Network Interface Bonding

- Bonding allows each interface to share a single network configuration for IP Addr, Gateway, etc. (Active/backup)
- Simplified field deployment for 1GE, 10GE SR, and 10GE LR
- Lab testing to verify performance expectations





# Network Interface Bonding

|          |                           |
|----------|---------------------------|
| A        | 1GE OOB Management        |
| <b>D</b> | DHCP for lab provisioning |
| <b>I</b> | BONDED 1GE                |
| <b>J</b> | BONDED 10GE SR            |
| <b>K</b> | BONDED 10GE LR            |

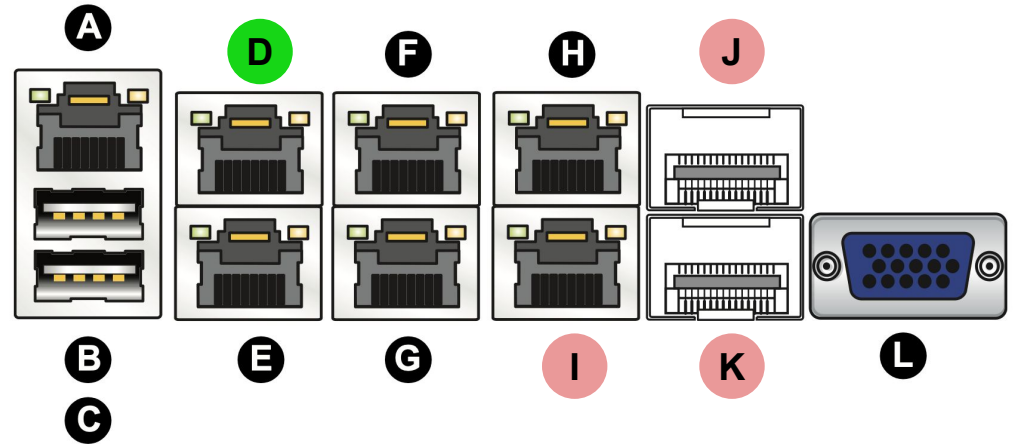


Figure 3-1. Rear Input/Output Ports

| Rear Panel I/O |               |                      |
|----------------|---------------|----------------------|
| A. IPMI LAN    | E. LAN Port 1 | I. LAN Port 5        |
| B. USB Port 1  | F. LAN Port 4 | J. LAN Port 8 (SFP+) |
| C. USB Port 0  | G. LAN Port 3 | K. LAN Port 7 (SFP+) |
| D. LAN Port 2  | H. LAN Port 6 | L. VGA Port          |

# Subnet Config /29

- Dual Homed DL support
- /29 vs /30
- Troubleshooting - verify the connectivity inside VLAN
- Support for dual testpoint field deployments
- Flexibility

| <b>/29 Subnet IP Usage</b> |                |
|----------------------------|----------------|
| Network                    | 192.168.0.0/29 |
| Gateway                    | 192.168.0.1    |
| DL1                        | 192.168.0.2    |
| DL2                        | 192.168.0.3    |
| Testpoint A                | 192.168.0.4    |
| Testpoint B                | 192.168.0.5    |
| Extra                      | 192.168.0.6    |
| Broadcast                  | 192.168.0.7    |

# Subnet Config /29

[Spreadsheet Template](#)

| Test Network   | Subnet | Net Name           | Testpoint A        | Testpoint B        |
|----------------|--------|--------------------|--------------------|--------------------|
| 141.213.137.72 | 29     | V-PERF-TEST-PAIR-1 | perf-test-pair-1-a | perf-test-pair-1-b |
| 141.213.137.80 | 29     | V-PERF-TEST-PAIR-2 | perf-test-pair-2-a | perf-test-pair-2-b |

| IPMI Network | Subnet | IPMI Net Name           | IPMI A                  | IPMI B                  |
|--------------|--------|-------------------------|-------------------------|-------------------------|
| 10.224.54.72 | 29     | V-PERF-TEST-IPMI-PAIR-1 | perf-test-ipmi-pair-1-a | perf-test-ipmi-pair-1-b |
| 10.224.54.80 | 29     | V-PERF-TEST-IPMI-PAIR-2 | perf-test-ipmi-pair-2-a | perf-test-ipmi-pair-2-b |



# Test Procedure: Bastion Host & Testpoint

1. Build VLAN network @ test site
2. Stretch VLAN to ports you're testing through
3. Plug in the equipment, power up test hardware
4. Ping device from network
5. Log on to trusted pScheduler host
6. Run your test!

```
[epcjr@its-perfsonar-bastion mobile_demo]$ pscheduler task \  
> throughput \  
> --source 141.213.137.100 \  
> --dest 141.213.137.101  
Submitting task...  
Task URL:  
https://141.213.137.100/pscheduler/tasks/85c9f6dd-e0b2-4120-9  
Running with tool 'iperf3'  
Fetching first run...
```

```
Next scheduled run:  
https://141.213.137.100/pscheduler/tasks/85c9f6dd-e0b2-4120-9  
ns/09bd4d3f-4e6d-46b7-b874-95bdc9a86076  
Starts 2020-10-02T12:21:49-04 (~6 seconds)  
Ends 2020-10-02T12:22:08-04 (~18 seconds)  
Waiting for result...
```

```
* Stream ID 5  
Interval      Throughput      Retransmits      Current Window  
0.0 - 1.0     9.93 Gbps       2                 1.97 MBytes  
1.0 - 2.0     9.90 Gbps       0                 1.97 MBytes  
2.0 - 3.0     9.91 Gbps       0                 1.97 MBytes  
3.0 - 4.0     9.90 Gbps       0                 1.97 MBytes  
4.0 - 5.0     9.91 Gbps       0                 2.21 MBytes  
5.0 - 6.0     9.90 Gbps       0                 2.21 MBytes  
6.0 - 7.0     9.91 Gbps       1                 1.12 MBytes  
7.0 - 8.0     9.90 Gbps       1                 1.04 MBytes  
8.0 - 9.0     9.90 Gbps       0                 1.08 MBytes  
9.0 - 10.0    9.91 Gbps       1                 942.24 KBytes
```

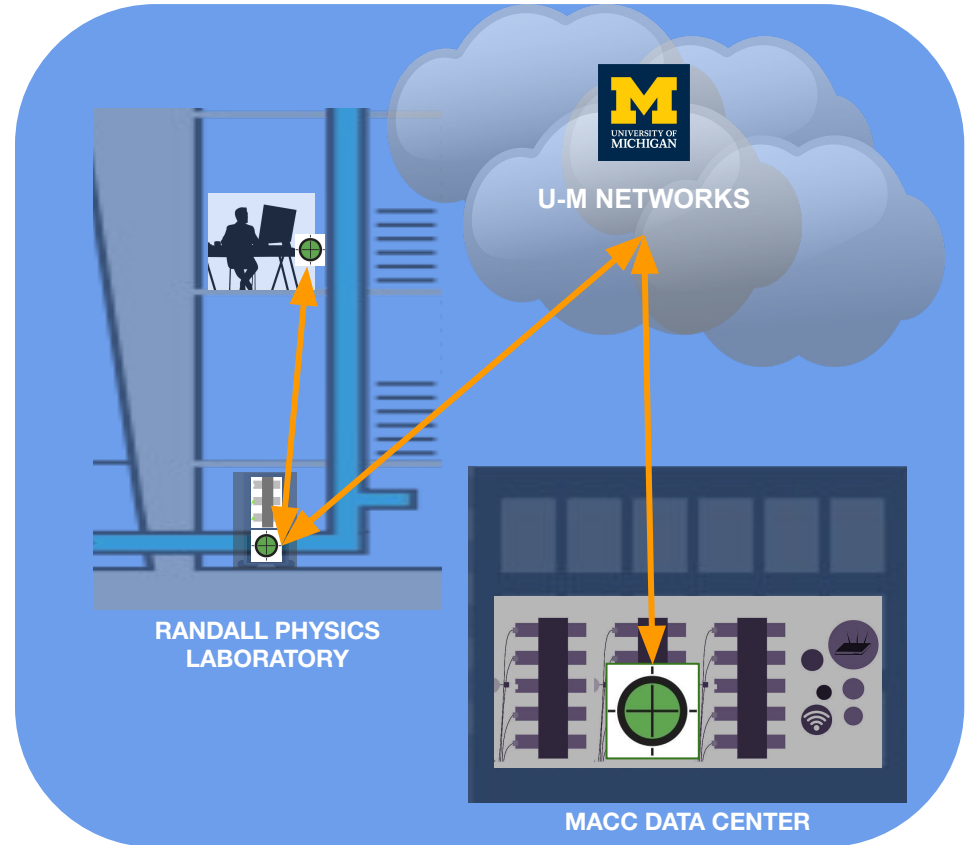
```
Summary  
Interval      Throughput      Retransmits  
0.0 - 10.0    9.91 Gbps       5
```

```
No further runs scheduled.
```

```
[epcjr@its-perfsonar-bastion mobile_demo]$ █
```


# Initial Field Deployment: End-to-End Testing - Verify Problem Exists

- Testpoint A using same fiber cabling as user
- Testpoint B in same data center rack as server
- Tests ran from user test to various permanent, temporary PerfSonar boxes
- All had issues!




# Test Procedure: Benchmark Test Lab Equipment

- Lessons: done in a lab before the field deployment verify that SM can actually hit 1G and 10G throughput goals connected back to back
- Configuration (IP address, cabling)
- Results / score




# Test Procedure: UMnet Testing - Pilot Shortest Loop

- Testing equipment always on the same network
- Allows to stay on switch, test a fresh, new, unboxed switch



# Test Procedure: Test Larger Segments

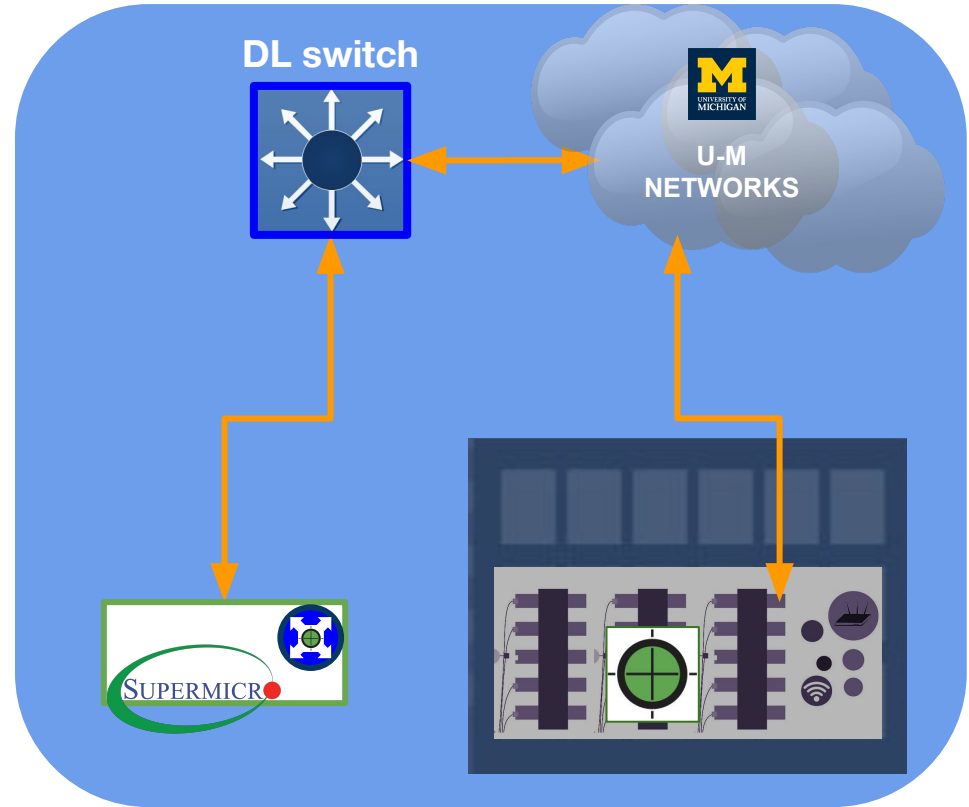
- Allows testing cabling, hardware, without leaving to core.
- Allows permanent install that can be tested against
- Same or different network testing options





# Test Procedure: Entire Network

- From DL to permanent PerfSonar installations
- Allows core testing
- Limitations: speeds availability - core is 100gb



# Credits and Resource Links

## Credits:

- Jeff Hagley
- Ajay Nautiyal
- Katarina Thomas
- Marlin Whitaker

## Resources:

- [pScheduler command and CentOS Interface Bonding](#)
- [perfSONAR Ansible Playbook](#)
- [perfSONAR Limits Tutorial](#)
- [SuperMicro E300-8D](#)
- [Deployment Template](#)

