

perfs--NAR

Small, Low-cost Nodes “the Flock”

JET Meeting
Scott Chevalier (schevali@iu.edu)
Network Systems Analyst, IN@IU
August 16, 2016



Who is this guy?

- Was once an Assistant Network Administrator (1995)
 - This included monitoring labs, teaching HS students, summer re-wiring projects, other duties as required...
- With IU since 2001
 - Started off as an Asst. Webmaster for the IUPUI campus, to Webmaster, to GlobalNOC

What are we talking about today?

- Problem statements
 - Limited budgets for perfSONAR deployments
 - Need many boxes for mobile, adhoc test-points
- A possible solution: Low-cost, Small-form Nodes
- Primary Benefits and Use Cases
 - Low-cost nodes for multiple, meshed testpoints
 - Small-form for easy transport as adhoc testers
- Installation and Configuration
- Meshes in the wild
- Next steps...

The Beginning


- Many institutions/groups experimenting with and deploying perfSONAR on low cost / small form factor nodes
- perfSONAR project committed to officially test/support small number of configurations as part of version 3.5 roadmap
- Community survey undertaken to determine what node characteristics were most important to the community.

Initial Assumptions

- Initially shooting for “very low price” nodes
- ARM based solutions investigated
 - Cubox/perfCube (Brian Tierney, etc.)
 - Beaglebone / Raspberry Pi (Alan Whinery, etc.)
- Great for some use cases (OWAMP)
- Can not support 1Gbps testing, Challenges in supporting ARM
- Cost profile could allow very large deployment

Survey Surprise

- Survey respondents generally preferred slightly higher cost and the ability to test at 1Gbps



A Course Change

- Began investigating slightly higher cost/performance nodes
- Targeting around \$200 per node
- Capable of achieving 1Gbps BWCTL throughput
- Easy to install
- Intel CPUs to avoid complications with supporting ARM builds

“Problem” Statement



- **First world problem: there are now dozens and dozens of small nodes on the market**
- **Many groups/institutions have very specific needs or goals which may require odd/interesting choices of gear.**

Possible Solutions



...

Possible Solutions – LIVA by ECS

- Price-point: \$100-150
- Build requirements: Debian (Ubuntu 12.4.5 desktop; perfSONAR endpoint)
- 1GE tested
- 2G memory onboard
- 32 or 64G onboard flash drive
- All parts included
- **May be discontinued/replaced by newer model(s); Liva X, Liva X2**



Possible Solutions – LIVA by ECS

Pros:

- Cost
- Sold as complete kit

Cons:

- Reports of CPU bottlenecking in testing
- EMMC drive requires Debian/Ubuntu
- “Headless” operation issues



LIVA/Debian Installation

- Ubuntu 12.04.5 **Desktop**
 - Server doesn't have proper EMMC driver
- Debian (Ubuntu) perfSONAR packages:
http://docs.perfsonar.net/install_debian#configuration
- Using EXAMPLE script log found at:
<http://download-test.odi.iu.edu/liva-ps/go>
<http://download-test.odi.iu.edu/liva-ps/foo>
 - And associated files on the USB drive (ntp10 and conf10)
- More instructions for LIVA installation:
http://docs.perfsonar.net/low_cost_nodes_configuration.html

Possible Solutions – GigaByte BRIX

- Price-point: \$150-200 (all pieces)
- Supported Build: CentOS Toolkit
- 1GE tested
- 2/4/8G memory
- Recommend 32G SSD
- Requires memory/drive purchase separately (price at top includes)



Possible Solutions – GigaByte BRIX

Pros:

- Supported Build: CentOS Toolkit
- 1GE tested

Cons:

- Power button on top
- Bulky power brick
- Single NIC



Possible Solutions – Intel NUC

- Price-point: \$150-200 (all pieces)
- Supported Build: CentOS Toolkit
- 1GE tested
- 2/4/8G memory
- Recommend 32G SSD (Requires mSATA)
- Requires memory, drive, and (sometimes) power cord purchase separately (price at top includes)



Possible Solutions – Intel NUC

Pros:

- Supported Build: CentOS Toolkit
- 1GE tested

Cons:

- Some early versions still sold without power cord, and require mSATA
- Bulky power brick



Possible Solutions – Zotac

- Price-point: \$150-200 (all pieces)
- Supported Build: CentOS Toolkit
- 1GE tested
- 2/4/8G memory
- Recommend 32G SSD
- Requires memory, drive purchase separately (price at top includes)



Possible Solutions – Zotac

Pros:

- Quad-core CPU
- Front-facing power indicator/button
 - And more small touches like this
- Ports, ports, ports!
- Tool-less case dis/assembly

Cons:

- Single NIC



Possible Solutions – SuperMicro by Servers Direct

- Price-point: \$600-1000 (all pieces)
- Supported Build: CentOS Toolkit
- Rack-mounted
- 1GE/10GE tested
- Customizable builds



Passive Cooling (possible issue)

- Most of the above include No fan in the box for cooling
- No direct issues with this were found during testing, but it is definitely worth considering when placing in warmer spaces.



Things to Keep in Mind

- You get what you pay for... mostly
- Prices vary widely and often
- Most of these options are Single NIC
 - pS PTB recommend Dual NIC systems to test OWAMP and throughput on unique interfaces
- Read the fine print
 - Most of these come as barebones systems, but the descriptions can be misleading
- Newer, better always coming Soon™

Speaking of Soon™ – 4 NIC Node

- Price-point: \$170-200 (all pieces)
- Supported Build: CentOS/Debian
- Should allow for individual NIC per test; OWAMP and Throughput
- Quad-core CPU
- “Industrial”
- 4 NICs!



So this replaces large servers?

- No, not really.
 - There is still a need to have larger boxes with more memory, higher speeds, warranties, stability, etc.
 - BUT there are niches where these are really useful and the cost makes them easy “starter” boxes
- Limited to 1GE (currently)
 - Fragile and possibly “Leggy” (not easily mounted into existing rack structures)

Low-cost, Small Nodes in the Wild

- Raspberry Pi deployment (Alan Whinery - University of Hawaii)
 - ultra-low-cost (30 USD), ultra-small size, USB hub power capable
 - Low-speed Ethernet, not as useful for bwctl (throughput); fine for OWAMP
- LIVA deployment beginning (IN@IU/APAN/Esnet/IU GlobalINOC)
 - 100-125 USD, small form factor, 15 watts usb micro
 - No fan, cannot support CentOS 6/Toolkit build on native drive (Ubuntu 12.04/Debian build)
- NUC/GigaByte/Zotac (in testing IN@IU/SCInet15 Flock)
- Servers Direct deployment (KENET - NSRC and IN@IU)
 - 1 RU rack-mounted server, 600-1000 USD
- Know of Another?
 - Strengths?
 - Limitations?



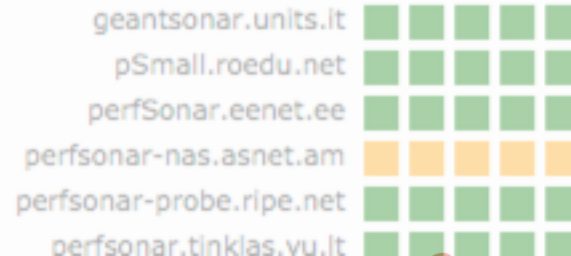
The Swarm (Whinery, U of Hawaii)

- Raspberry Pi deployment (Alan Whinery - University of Hawaii)
 - ultra-low-cost (~35 USD), ultra-small size, USB hub power capable
 - Low-speed Ethernet, not as useful for bwctl (throughput); fine for OWAMP
- Link: <http://docslide.us/documents/slide-1-uh-swarm-dense-perfsonar-deployment-with-small-inexpensive-devices.html>

GEANT Small Nodes Meshes

■ Loss rate is ≤ 0
■ Loss rate is ≥ 0
■ Loss rate is ≥ 0.01

- LIVA deployment beginning (IN@IU/APAN/Esnet/IU GlobalNOC)
 - 100-125 USD, small form factor, 15 watts usb micro
 - No fan, cannot support CentOS 6/Toolkit build on native drive (Ubuntu 12.04/Debian build)
- NUC/GigaByte/Zotac (in testing IN@IU/SCInet15 Flock)
- <http://perfsonar-smallnodes.geant.org/maddash-webui/>




KENET Low-cost Server-class

- Servers Direct deployment (KENET - NSRC and IN@IU)

- 1 RU rack-mounted server, 600-1000 USD

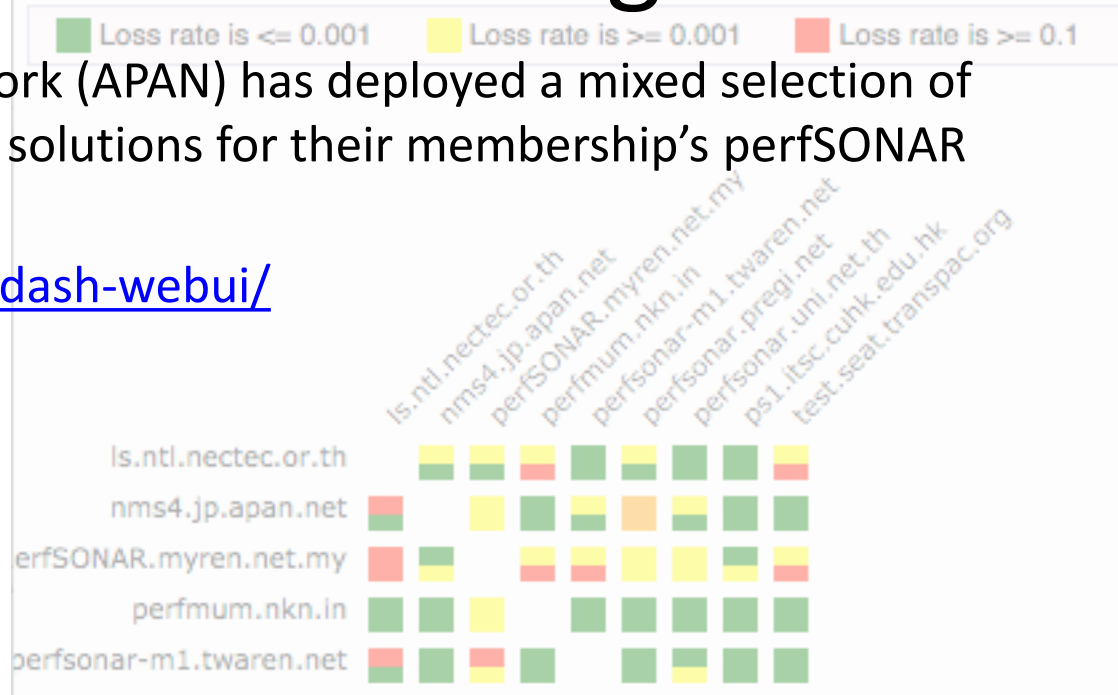
■ Loss rate is ≤ 0
■ Loss rate is ≥ 0
■ Loss rate is ≥ 0.01

- <http://maddash-uon.kenet.or.ke/maddash-webui/>



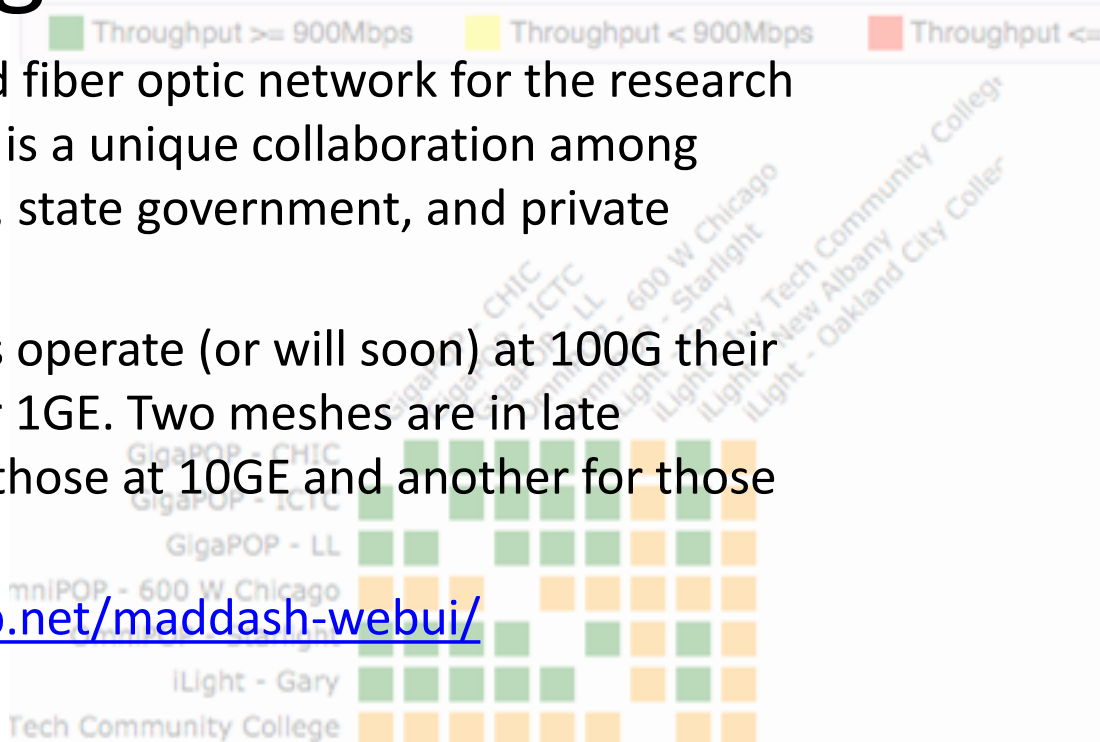
APAN Mixed Nodes Testing Mesh

- Asia Pacific Advanced Network (APAN) has deployed a mixed selection of server-class and small node solutions for their membership's perfSONAR mesh.
- <http://ps2.jp.apan.net/maddash-webui/>



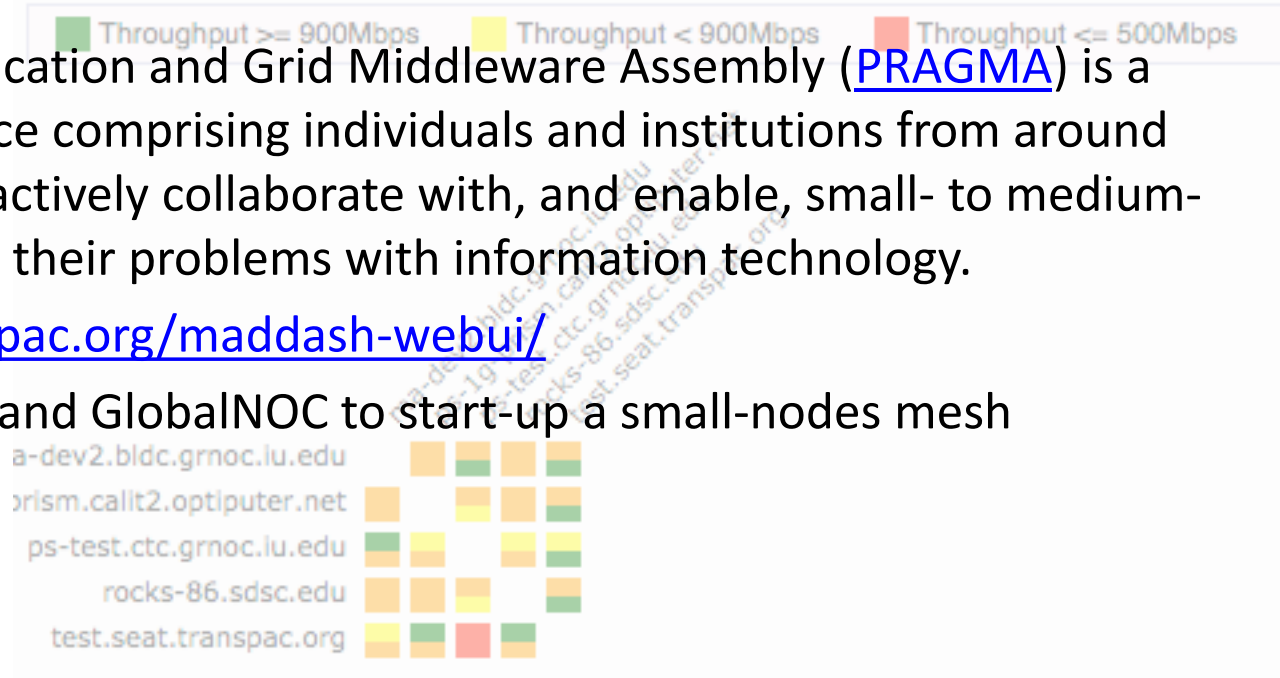
I-Light Mesh

- [I-Light](#) is Indiana's high-speed fiber optic network for the research and education community. It is a unique collaboration among Indiana colleges, universities, state government, and private sector broadband providers.
- While I-Light fiber backbones operate (or will soon) at 100G their members connect at 10GE or 1GE. Two meshes are in late planning/early deployment; those at 10GE and another for those at 1GE.
- <http://dc.ctc.Indiana.gigapop.net/maddash-webui/>



PRAGMA Mesh

- The Pacific Rim Application and Grid Middleware Assembly ([PRAGMA](#)) is a community of practice comprising individuals and institutions from around the Pacific Rim that actively collaborate with, and enable, small- to medium-sized groups to solve their problems with information technology.
- <http://data.ctc.transpac.org/maddash-webui/>
- Worked with IN@IU and GlobalNOC to start-up a small-nodes mesh
- Very Early stages



So is that it?

- No.
 - We as a community should continue to investigate and test new technology/options as they become available.
 - Try new ways to use older ideas
 - Research: Faster, smaller, cheaper, better. When we can get all of these we can stop. (no, not really)
- Small node prices and models are constantly fluctuating; next time we talk there will be something new to evaluate.
- Trials and test meshes exist in several institutions to continue evaluating boxes in the \$100-\$200 range for both CentOS and Debian distributions.

So is that it?

- As small nodes improve, and costs decrease, even more options will become viable
- While the initial tests can be good reference, the perfSONAR Dev group isn't in the full time business of defining all variables on all devices. Which is where the community comes in!
- “We” are working to make it easier for the pS Community to share their reviews of devices, experiments in the works, and use cases from the field.

What's Next?

- Develop flashable images
- More HOWTOs for setting up nodes
- Scripts to make setup easier
- Post results and information to perfSONAR website: <https://github.com/perfsonar/project/wiki/perfSONAR-Endpoint-Node-Project>

In Closing...

- More nodes equals more data for troubleshooting
 - Lower cost = more possible nodes
- Small size means they can ride along in backpacks for adhoc testing
- Other possible benefits:
 - Power over Ethernet (POE), Fit in places where space is at a premium

Resources

Low-cost Nodes Primer -

<https://www.perfsonar.net/deploy/hardware-selection/low-cost-hardware/>

Finding the right Installation -

<http://docs.perfsonar.net/install/options.html>



perfs--NAR

Small, Low-cost Nodes “the Flock”

JET Meeting
Scott Chevalier (scheval@iu.edu)
Network Systems Analyst, IN@IU
August 16, 2016

