Opening up OpenSM

with the Subnet Monitoring Tools

OFS User Group Workshop

March 19, 2015

Timothy Meier
tameier@llnl.gov

Lawrence Livermore National Laboratory

LLNL-PRES-668552
This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

Subnet Monitor Tools

- Introduce the **smt-gui** via a million screenshots
  - part of the Subnet Monitor Tool Suite (SMT)
  - top and trace route example
- Explain the motivation for another set of tools
  - fundamentally different (pros & cons)
- Describe the OpenSM Monitoring Service (OMS)
  - design and implementation

Why do we need another set of IB Tools?

- We use;
  - ibutils
  - infiniband-diags
  - pragmatic-infiniband-utilities
  - opensm console
  - opensm.log
- None designed for continuous monitoring
  - long term comprehensive interrogation
  - benign (minimal fabric disturbance)
  - accessible
    - local/remote
    - should not require root
    - interactive and/or programmatic
    - does not require ninja skills
  - safe and secure

How should IB be monitored?

- Information Aggregation
  - One entity responsible for periodically sweeping and gathering fabric information for all to share
  - Minimal fabric disturbance (deterministic)
  - OpenSM (untapped wealth of information)
    - Subnet manager
    - Subnet administrator
    - Performance manager
- Service (1 to many sharing)
  - Authentication, authorization, encryption
  - Multiple concurrent clients
  - Application interface
  - Platform independent (sort of)

OpenSM Monitoring Service (OMS)

OMS/SMT Software Stack

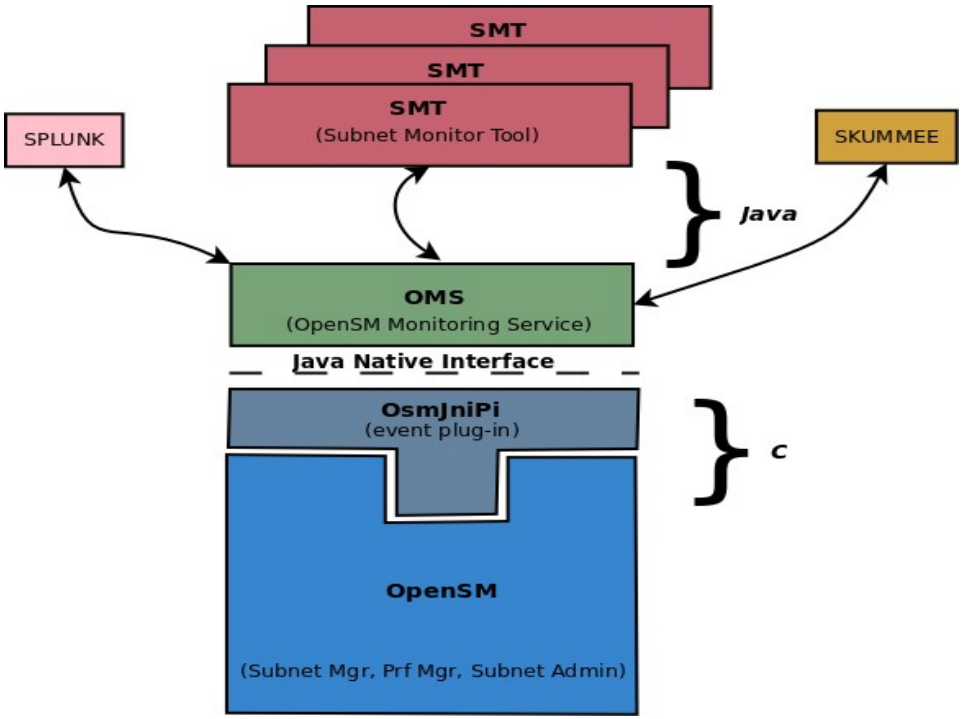

- OMS is a service that runs on the management node
- SMT commands are OMS clients
- SMT commands can be invoked on the management node or remote system

Subnet Monitor Tools (SMT)

- OMS Clients
  - Inherits all the pros and cons
- Tool set is just a collection of commands
  - New commands are relatively easy to build
  - Growing list...
- Uses a common code base
  - Raw OMS api, plus SMT convenience and analysis libraries
- Most commands operate on one or more OMS snapshots

  1. Connection to OMS on port 10011
     - Authenticated, encrypted, and compressed binary data
  2. An OMS history file
     - Compressed binary data

List of Subnet Monitor Tools

```
13:13:50 > smt -?
usage: smt <command> [-?] [?] [<command args>]

This command provides access to some of the most commonly used SMT commands.  Most commands should be
invoked directly using the form "smt-<command>", but can be invoked here for convenience.
.
 -?,--Help                      print this message
 -abt,--about                   smt-about,     - software package information
 -c,--config                    smt-config,    - checks or modifies the SMT configuration
 -con,--console                 smt-console,   - a curses application for viewing OMS information
 -e,--event                     smt-event,     - shows SM events, traps, and exceptions
 -f,--fabric                    smt-fabric,    - provides fabric level information
 -fn,--file                     smt-file,      - provides information about OMS files
 -gui,--gui                     smt-gui,       - a gui fabric exploration tool
 -h,--help                      smt-help,      - a gui help tool
 -id,--id                       smt-id,        - an identificaton tool (name resolver)
 -l,--link                      smt-link,      - provides link level information
 -lf,--logFile <file name>      the file name or pattern to use for log files
 -ll,--logLevel <log level>     the verbosity level for log files
 -m,--multicast                 smt-multicast, - a multicast group tool
 -n,--node                      smt-node,      - provides node level information
 -p,--port                      smt-port,      - provides port level information
 -part,--partition              smt-partition, - a partition tool
 -pv,--priv                     smt-priv,      - a set of privileged commands
 -r,--route                     smt-route,     - routing table tools
 -rC,--readConfig <filename>    reads the specified configuration file
 -rcd,--record                  smt-record,    - saves OMS information (flight recorder)
 -t,--top                       smt-top,       - shows top errors and traffic
 -v,--version                   print the version

examples:
> smt -?                    - provides this help
> smt --node ?              - provides help for the node command (no dash for its args)
> smt --multicast pn 10013  - multicast status for service on port 10013
.
Copyright (C) 2015, Lawrence Livermore National Security, LLC
```

smt-record: the flight recorder

- collects and saves OMS snapshots (history)
  - an OMS snapshot contains everything provided by OMS
- requires an OMS connection
- specify host, port, number to collect and file name

```
11:23:27 > smt-record -pn 10013 -nh 3 -wH hype3H.his
OMS_Collection
fabric name:                 hype355.llnl.gov
first timestamp:             Mar 04 11:23:40 2015
last timestamp:              Mar 04 14:20:40 2015
ave secs between records:    180
# secs between pfmgr sweeps: 180
# records in collection:     60
# nodes:                     164
# ports:                     759
# links:                     287
```

smt-file

- determines file type and attributes
- can manipulate or convert files
- specify file(s)

```
14:39:53 > ls -lah *3H.his
-rw-r----- 1 meier3 meier3 3.1M Mar  4 14:23 hype3H.his
-rw-r----- 1 meier3 meier3  95M Feb 25 13:55 sierra3H.his

14:40:01 > smt-file -i sierra3H.his
OMS_Collection
fabric name:                 sierra7.llnl.gov
first timestamp:             Feb 25 10:35:08 2015
last timestamp:              Feb 25 13:52:38 2015
ave secs between records:    150
# secs between pfmgr sweeps: 150
# records in collection:     72
# nodes:                     2188
# ports:                     11638
# links:                     5768

14:40:46 > smt-file -i hype3H.his
OMS_Collection
fabric name:                 hype355.llnl.gov
first timestamp:             Mar 04 11:23:40 2015
last timestamp:              Mar 04 14:20:40 2015
ave secs between records:    180
# secs between pfmgr sweeps: 180
# records in collection:     60
# nodes:                     164
# ports:                     759
# links:                     287
```

smt-gui

- exploration and visualization
- development and testing
- postmortem analysis

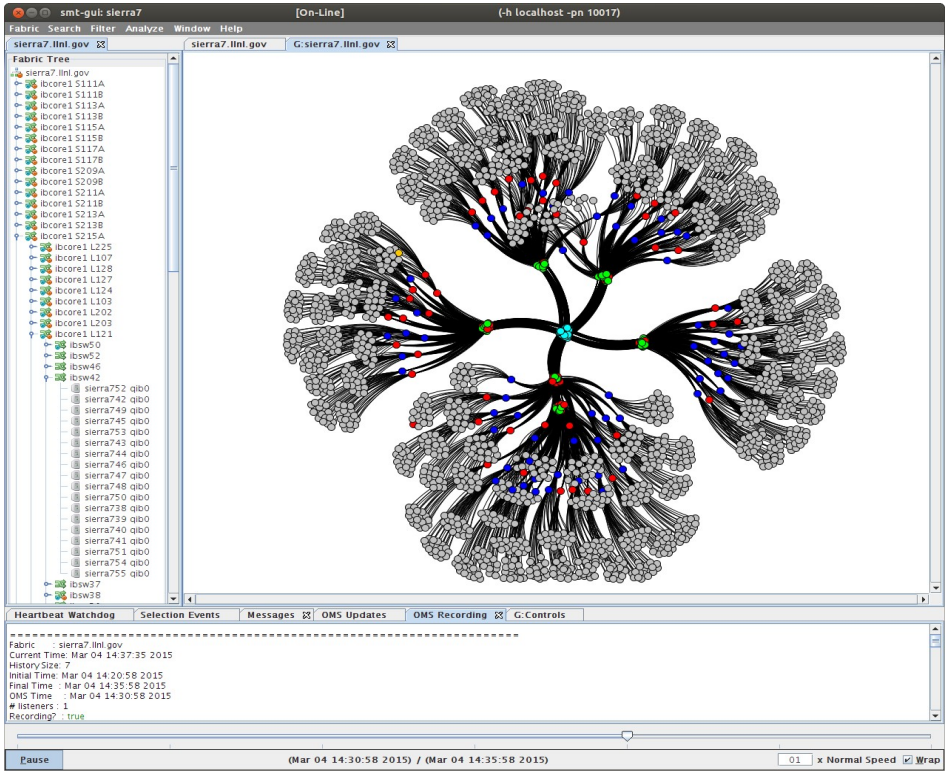

modes -

- on-line
  - connected to OMS
- off-line
  - flight recorder file
- almost identical behavior

comprehensive -

- dynamic (time based)
- includes functionality of other SMT commands
- visual analytics (charts, graphs, trees, etc.)

smt-gui

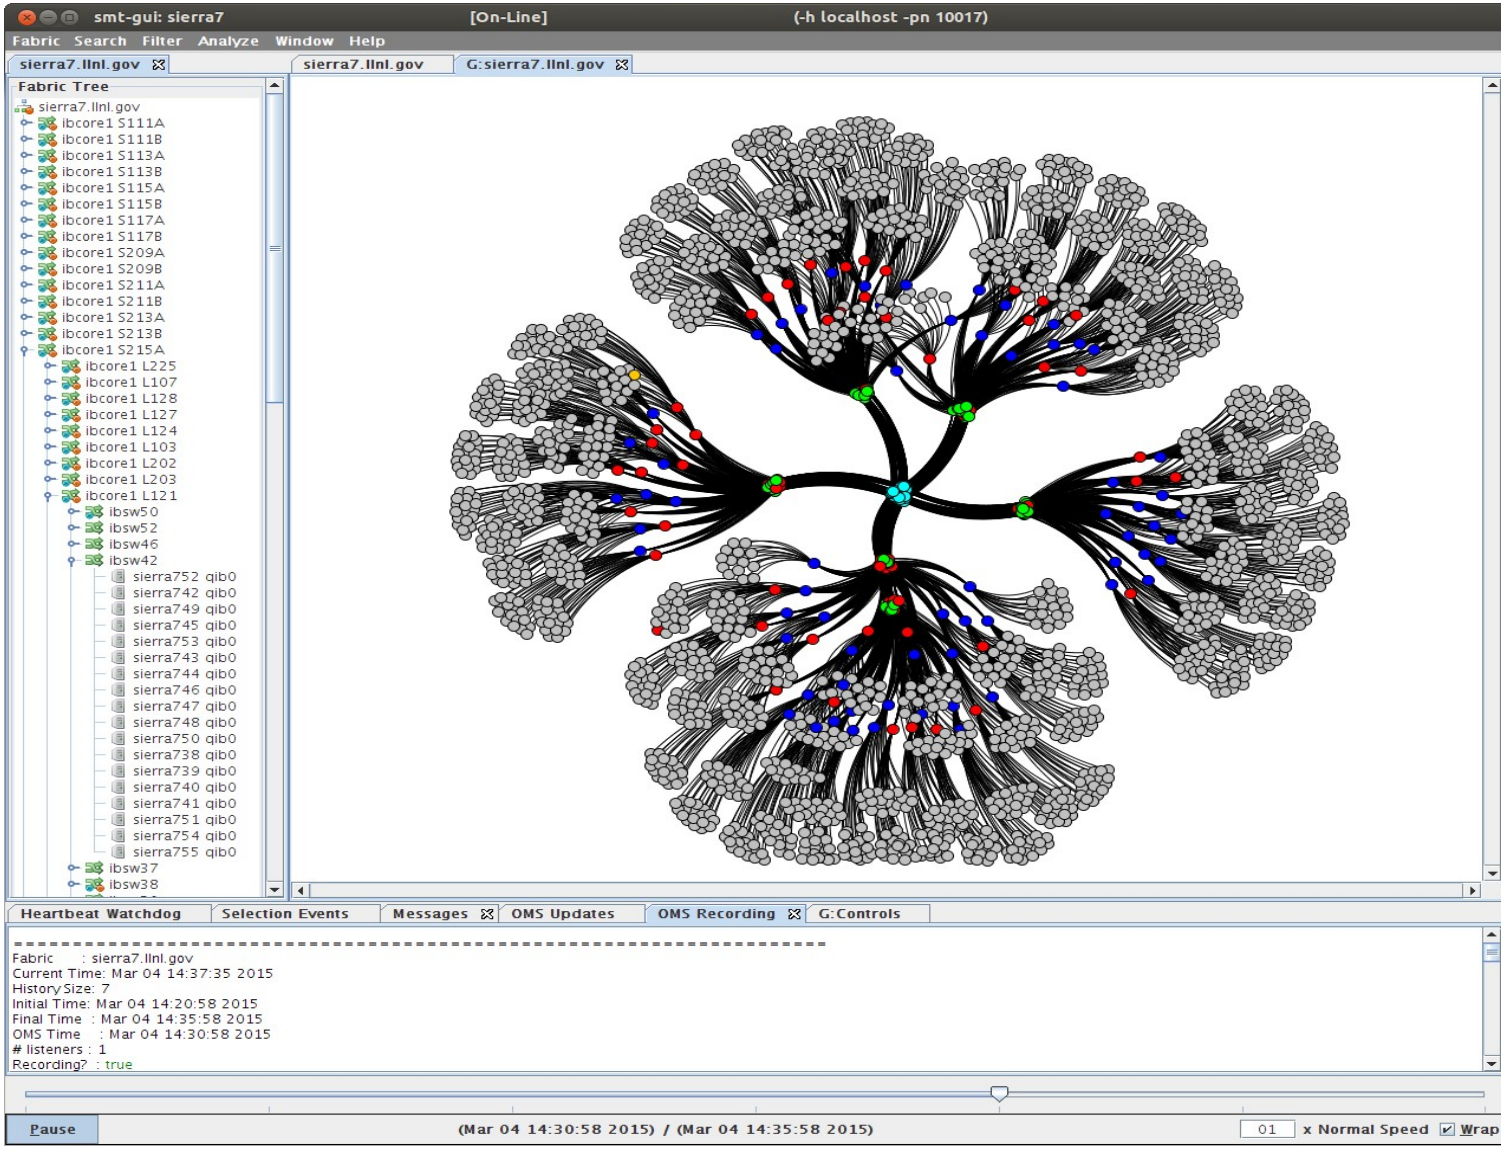
smt-gui: sierra7
[On-Line]
(-h localhost -pn 10017)
Fabric Search Filter Analyze Window Help
sierra7.llnl.gov
G:sierra7.llnl.gov
Fabric Tree
sierra7.llnl.gov
ibcore1 S111A
ibcore1 S111B
ibcore1 S113A
ibcore1 S113B
ibcore1 S115A
ibcore1 S115B
ibcore1 S117A
ibcore1 S117B
ibcore1 S209A
ibcore1 S209B
ibcore1 S211A
ibcore1 S211B
ibcore1 S213A
ibcore1 S213B
ibcore1 S215A
ibcore1 L225
ibcore1 L107
ibcore1 L128
ibcore1 L127
ibcore1 L124
ibcore1 L103
ibcore1 L202
ibcore1 L203
ibcore1 L121
ibsw50
ibsw52
ibsw46
ibsw42
sierra752 qib0
sierra742 qib0
sierra749 qib0
sierra745 qib0
sierra753 qib0
sierra743 qib0
sierra744 qib0
sierra746 qib0
sierra747 qib0
sierra748 qib0
sierra750 qib0
sierra738 qib0
sierra739 qib0
sierra740 qib0
sierra741 qib0
sierra751 qib0
sierra754 qib0
sierra755 qib0
ibsw37
ibsw38
Heartbeat Watchdog
Selection Events
Messages
OMS Updates
OMS Recording
G:Controls
Fabric : sierra7.llnl.gov
Current Time: Mar 04 14:37:35 2015
History Size: 7
Initial Time: Mar 04 14:20:58 2015
Final Time : Mar 04 14:35:58 2015
OMS Time : Mar 04 14:30:58 2015
listeners : 1
Recording? : true
Pause
(Mar 04 14:30:58 2015) / (Mar 04 14:35:58 2015)
01
x Normal Speed
Wrap

major gui components

- Title bar
  - Shows the fabric name and details of the mode of operation
- Menu bar
  - Provides access to general or global functions
- Fabric Tree panel – left side
  - Hierarchical view of the nodes (navigable & selectable)
- Diagnostic Panel
  - Message area (various threads)
  - Graph controls
- Play Bar
  - Move through the OMS collection
  - Start/stop, step, and play at desired rate
- Main Panel
  - Details of selected object
  - Analysis results
  - Graphs, tables, trees
  - Etc...

composition

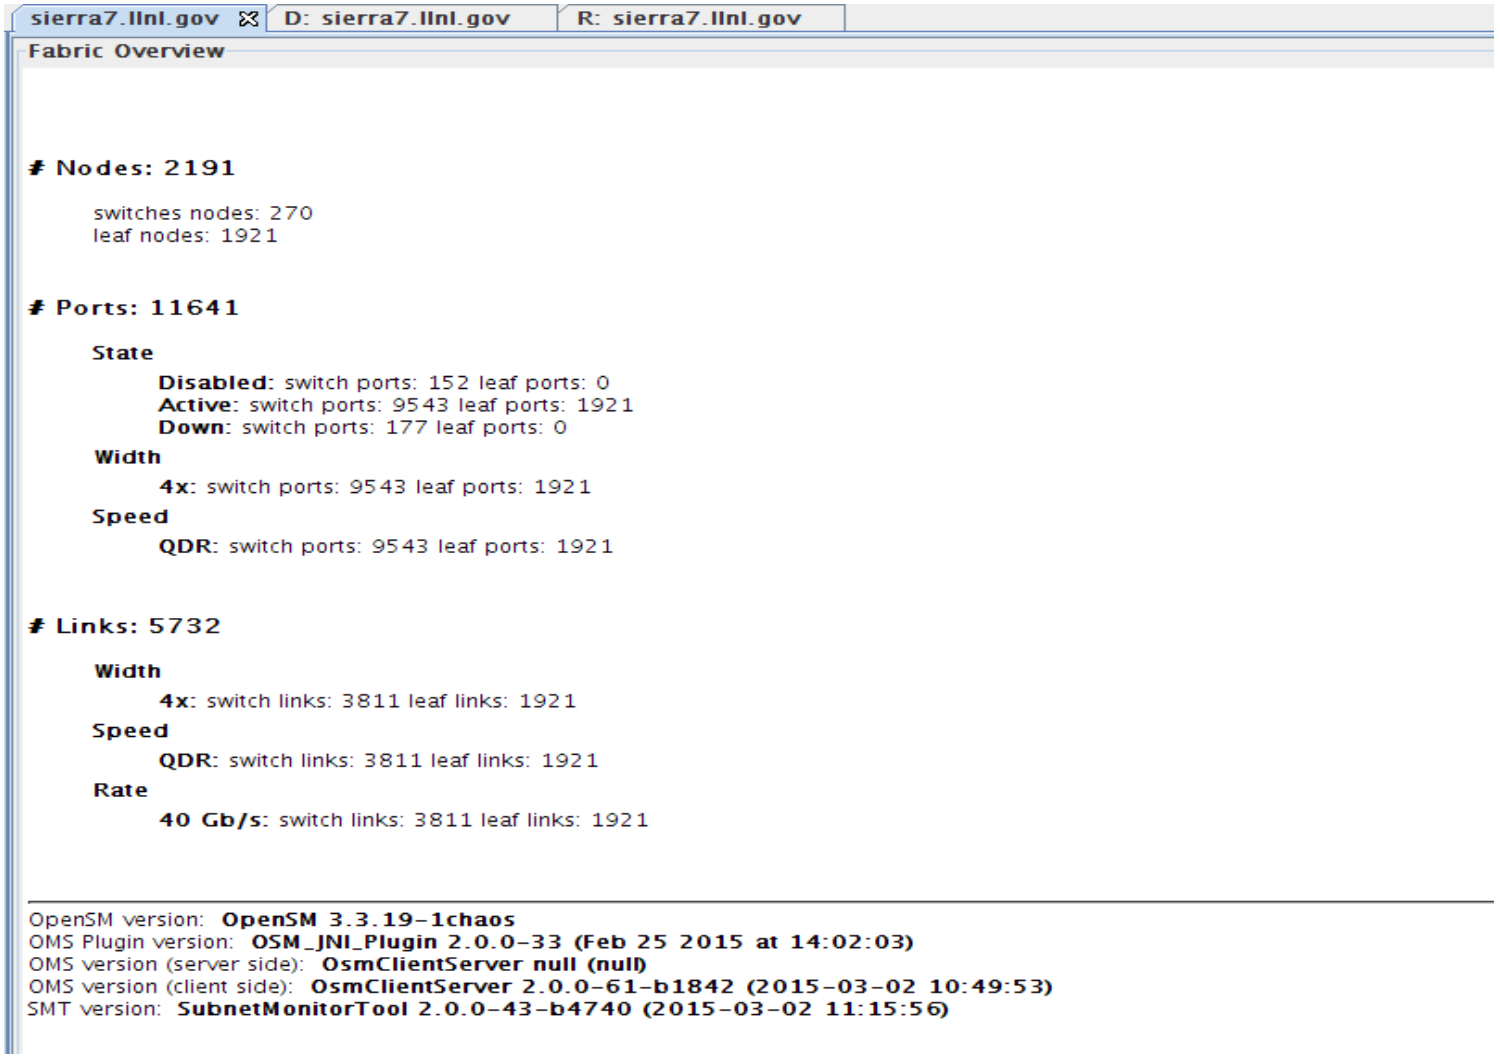

detailed information

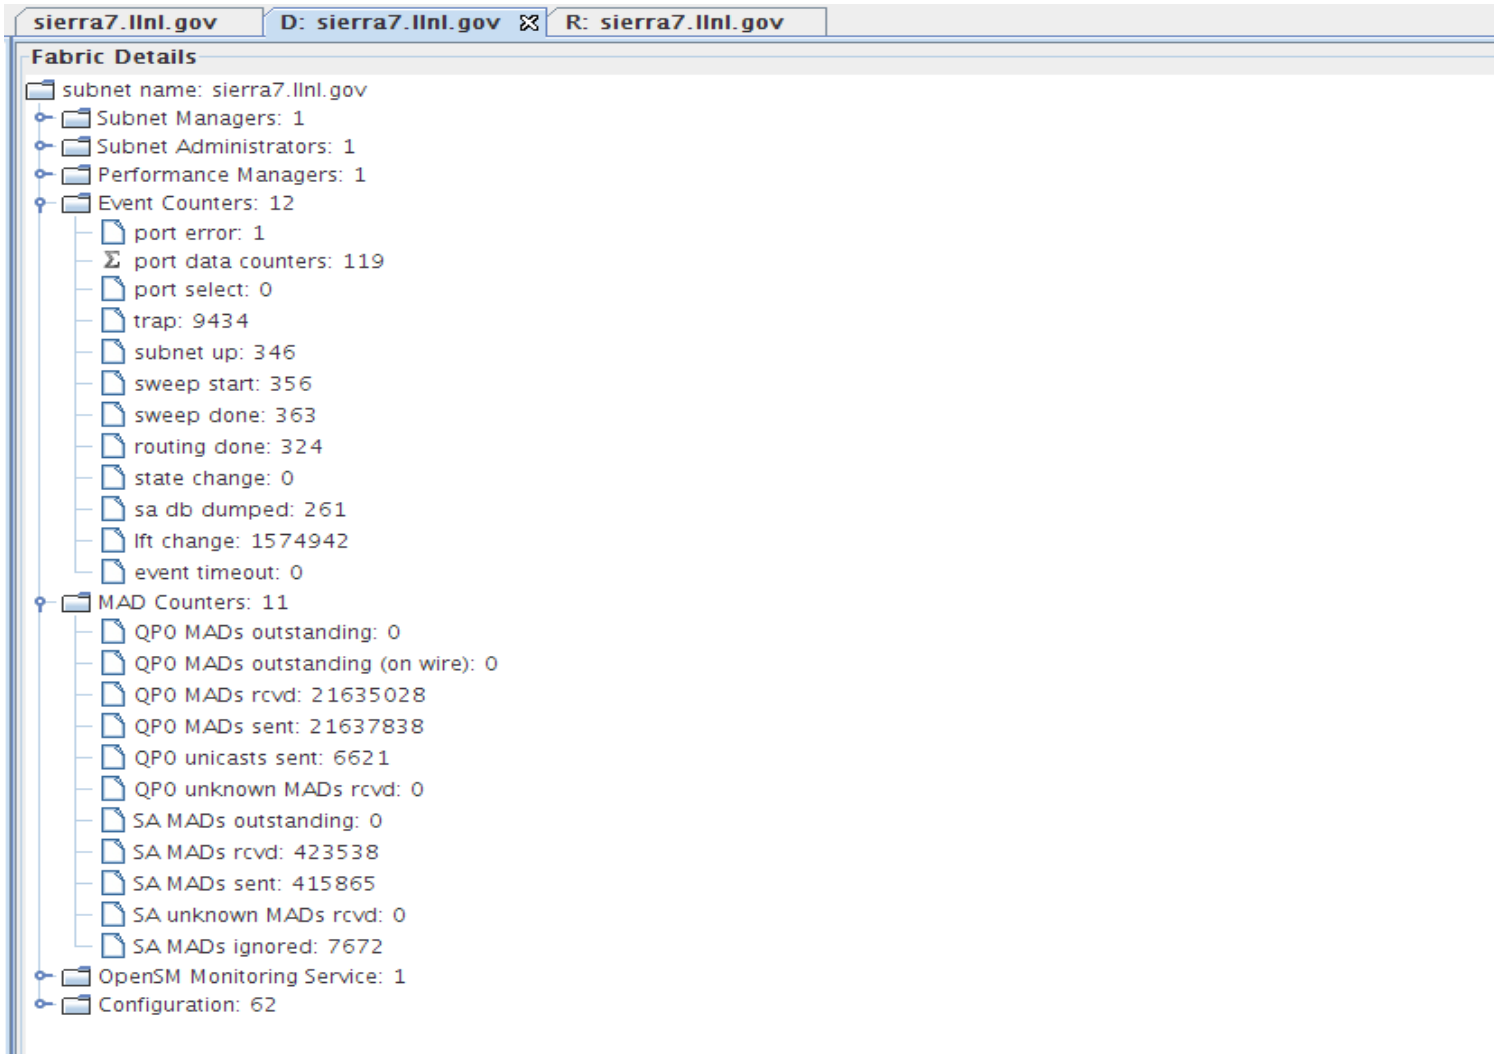

routing tables

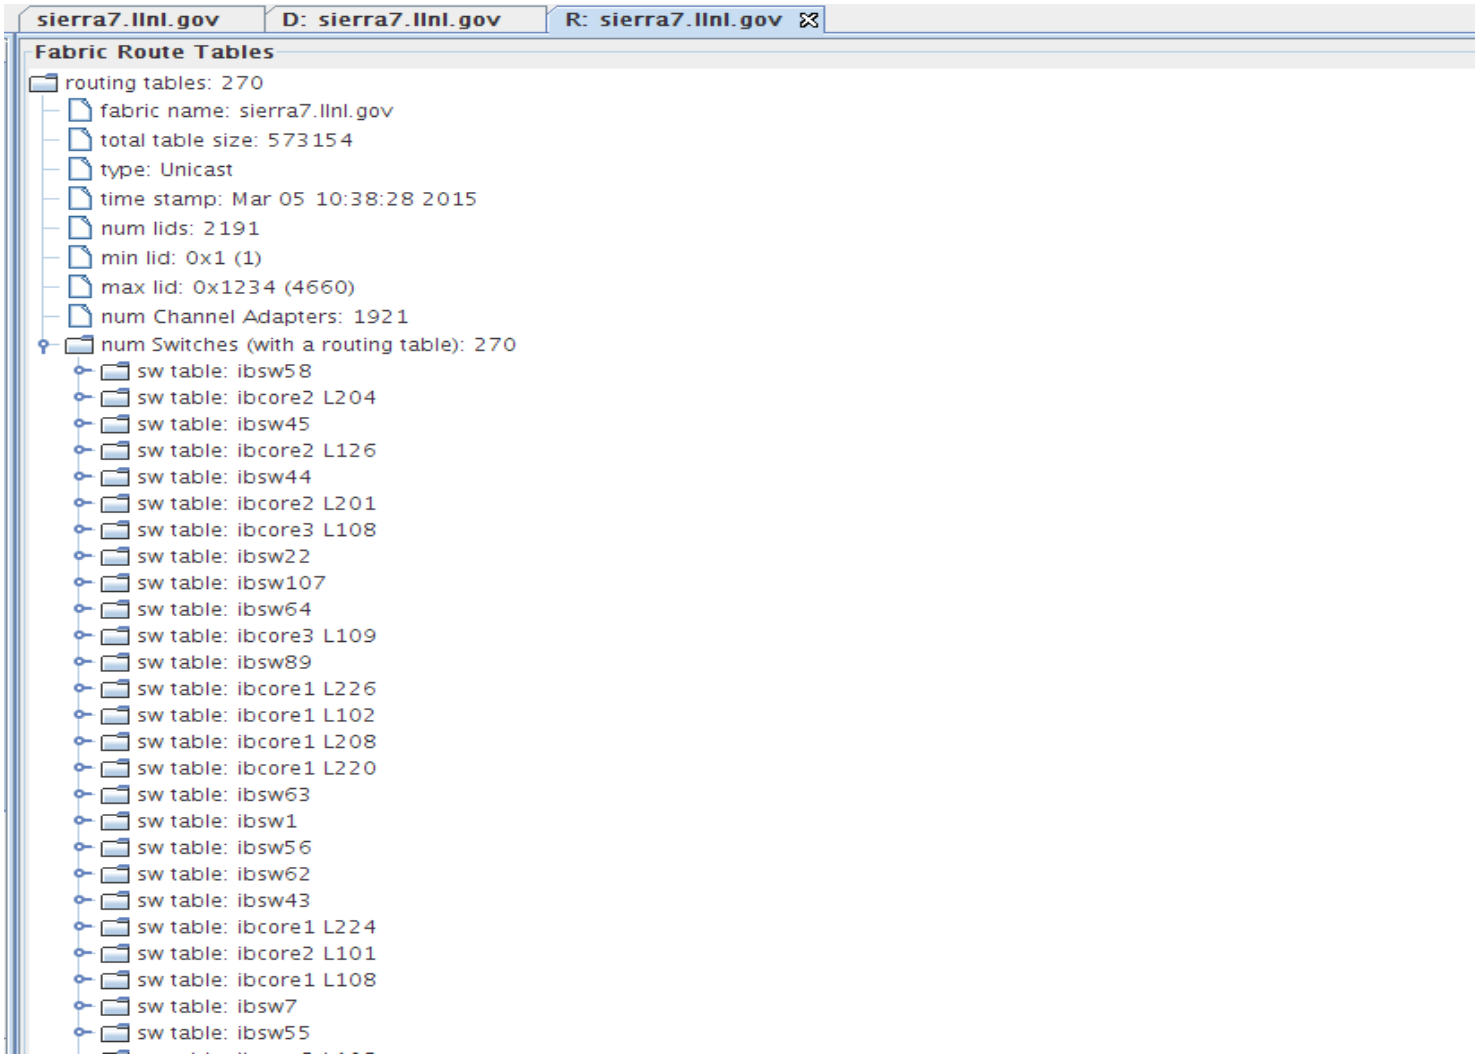

node tree

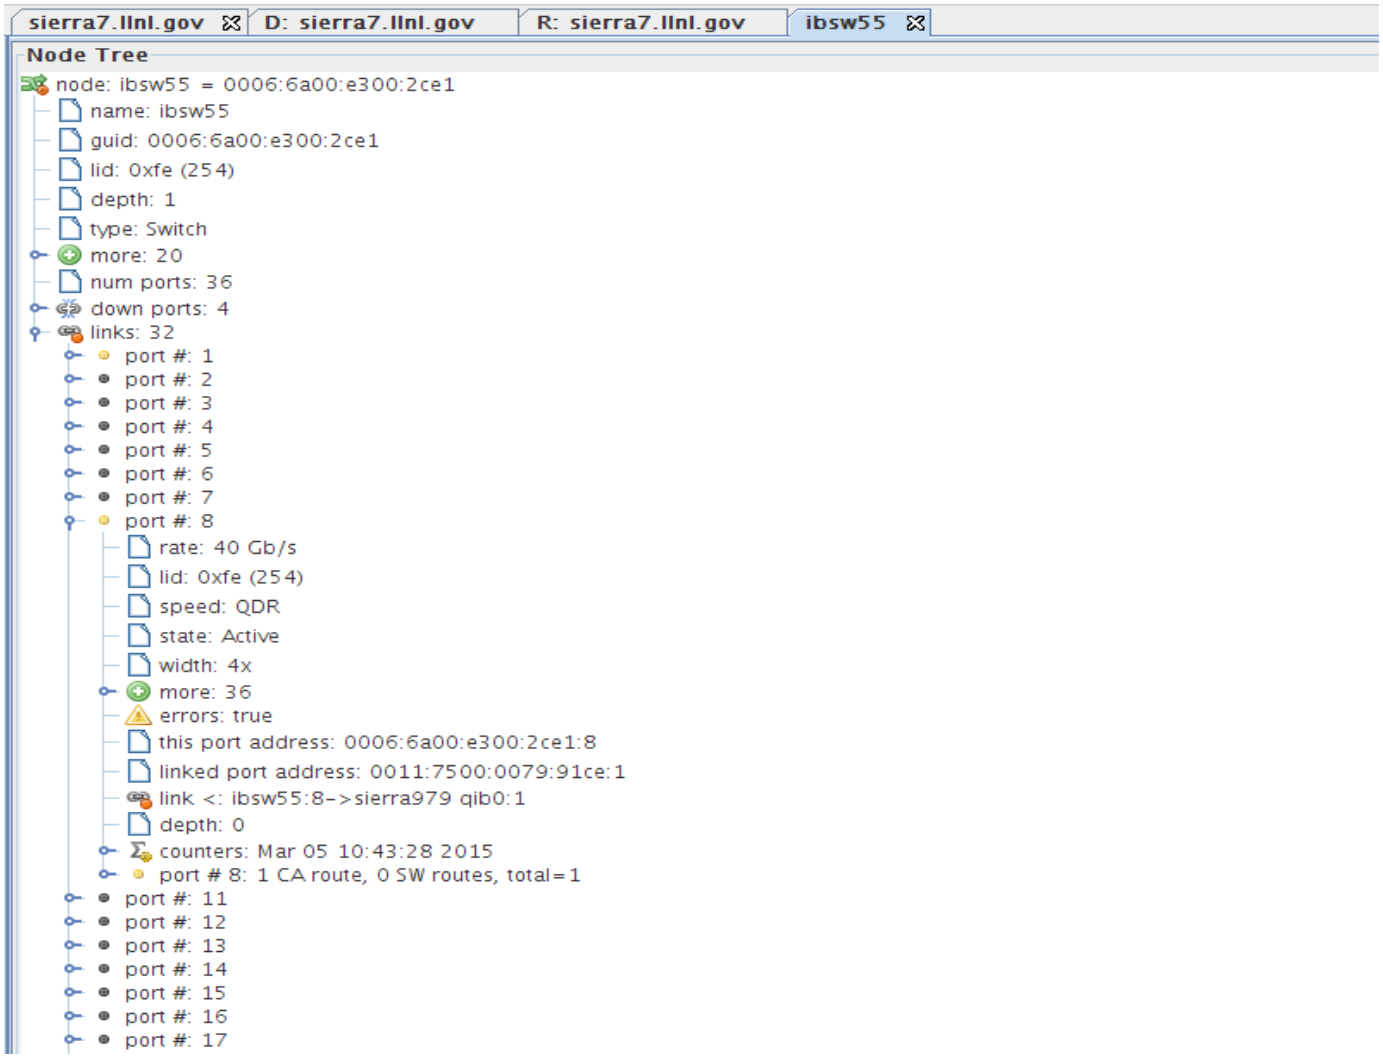
sierra7.llnl.gov
D: sierra7.llnl.gov
R: sierra7.llnl.gov
ibsw55
Node Tree
node: ibsw55 = 0006:6a00:e300:2ce1
name: ibsw55
guid: 0006:6a00:e300:2ce1
lid: 0xfe (254)
depth: 1
type: Switch
more: 20
num ports: 36
down ports: 4
links: 32
port #: 1
port #: 2
port #: 3
port #: 4
port #: 5
port #: 6
port #: 7
port #: 8
rate: 40 Gb/s
lid: 0xfe (254)
speed: QDR
state: Active
width: 4x
more: 36
errors: true
this port address: 0006:6a00:e300:2ce1:8
linked port address: 0011:7500:0079:91ce:1
link <: ibsw55:8->sierra979 qib0:1
depth: 0
counters: Mar 05 10:43:28 2015
port # 8: 1 CA route, 0 SW routes, total=1
port #: 11
port #: 12
port #: 13
port #: 14
port #: 15
port #: 16
port #: 17

port tree

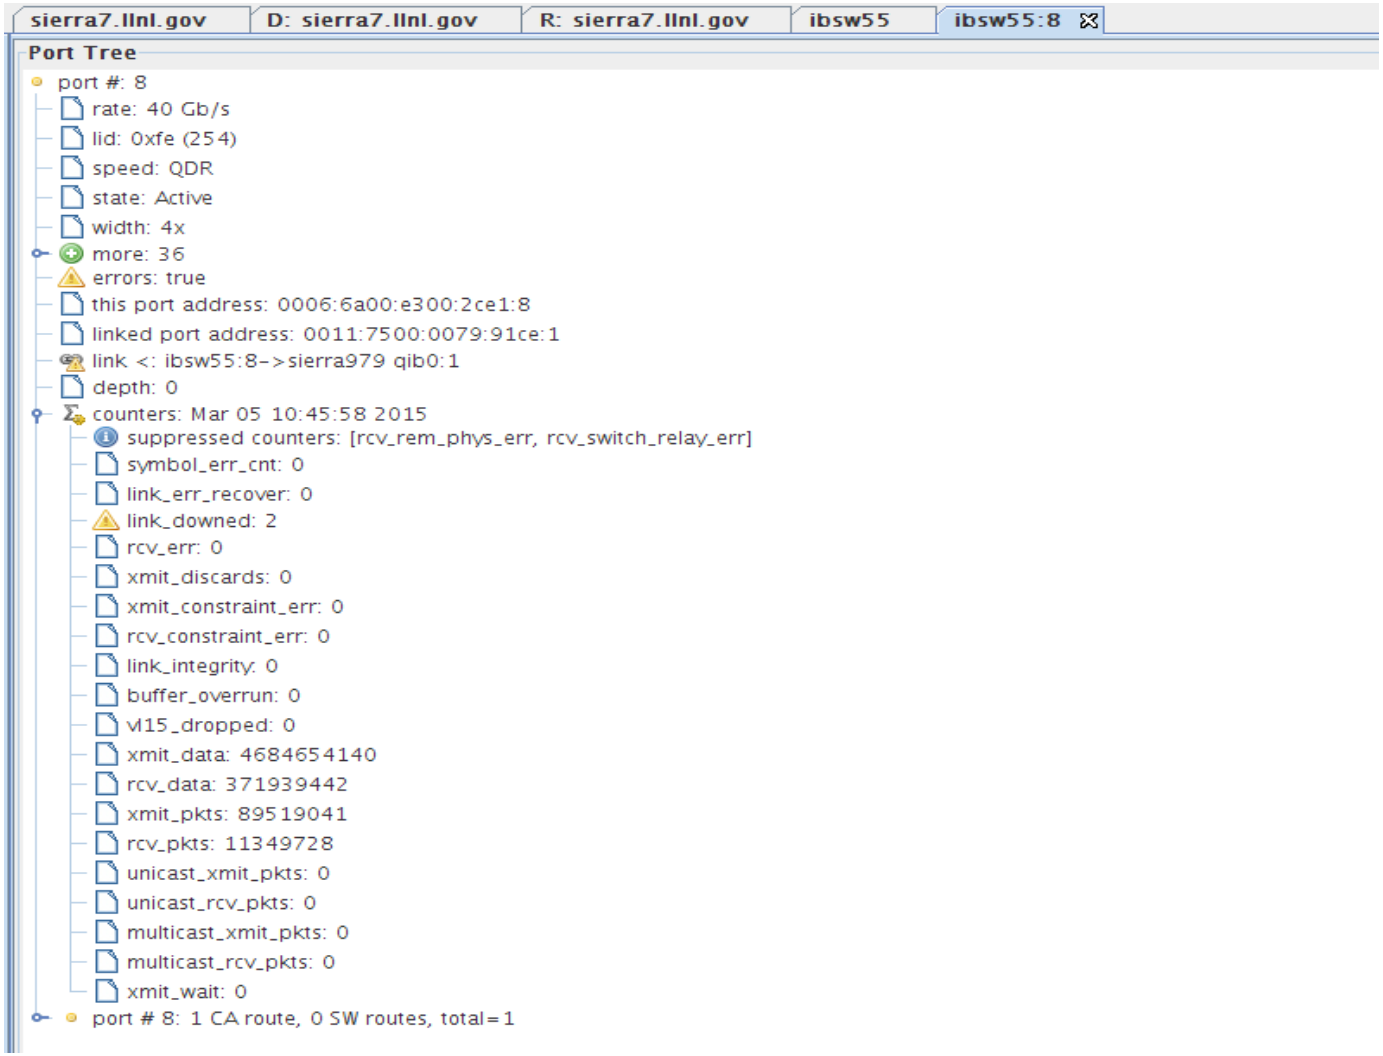
sierra7.llnl.gov
D: sierra7.llnl.gov
R: sierra7.llnl.gov
ibsw55
ibsw55:8
Port Tree
port #: 8
rate: 40 Gb/s
lid: 0xfe (254)
speed: QDR
state: Active
width: 4x
more: 36
errors: true
this port address: 0006:6a00:e300:2ce1:8
linked port address: 0011:7500:0079:91ce:1
link <: ibsw55:8->sierra979 qib0:1
depth: 0
counters: Mar 05 10:45:58 2015
suppressed counters: [rcv_rem_phys_err, rcv_switch_relay_err]
symbol_err_cnt: 0
link_err_recover: 0
link_downed: 2
rcv_err: 0
xmit_discards: 0
xmit_constraint_err: 0
rcv_constraint_err: 0
link_integrity: 0
buffer_overrun: 0
vl15_dropped: 0
xmit_data: 4684654140
rcv_data: 371939442
xmit_pkts: 89519041
rcv_pkts: 11349728
unicast_xmit_pkts: 0
unicast_rcv_pkts: 0
multicast_xmit_pkts: 0
multicast_rcv_pkts: 0
xmit_wait: 0
port # 8: 1 CA route, 0 SW routes, total=1

link tree

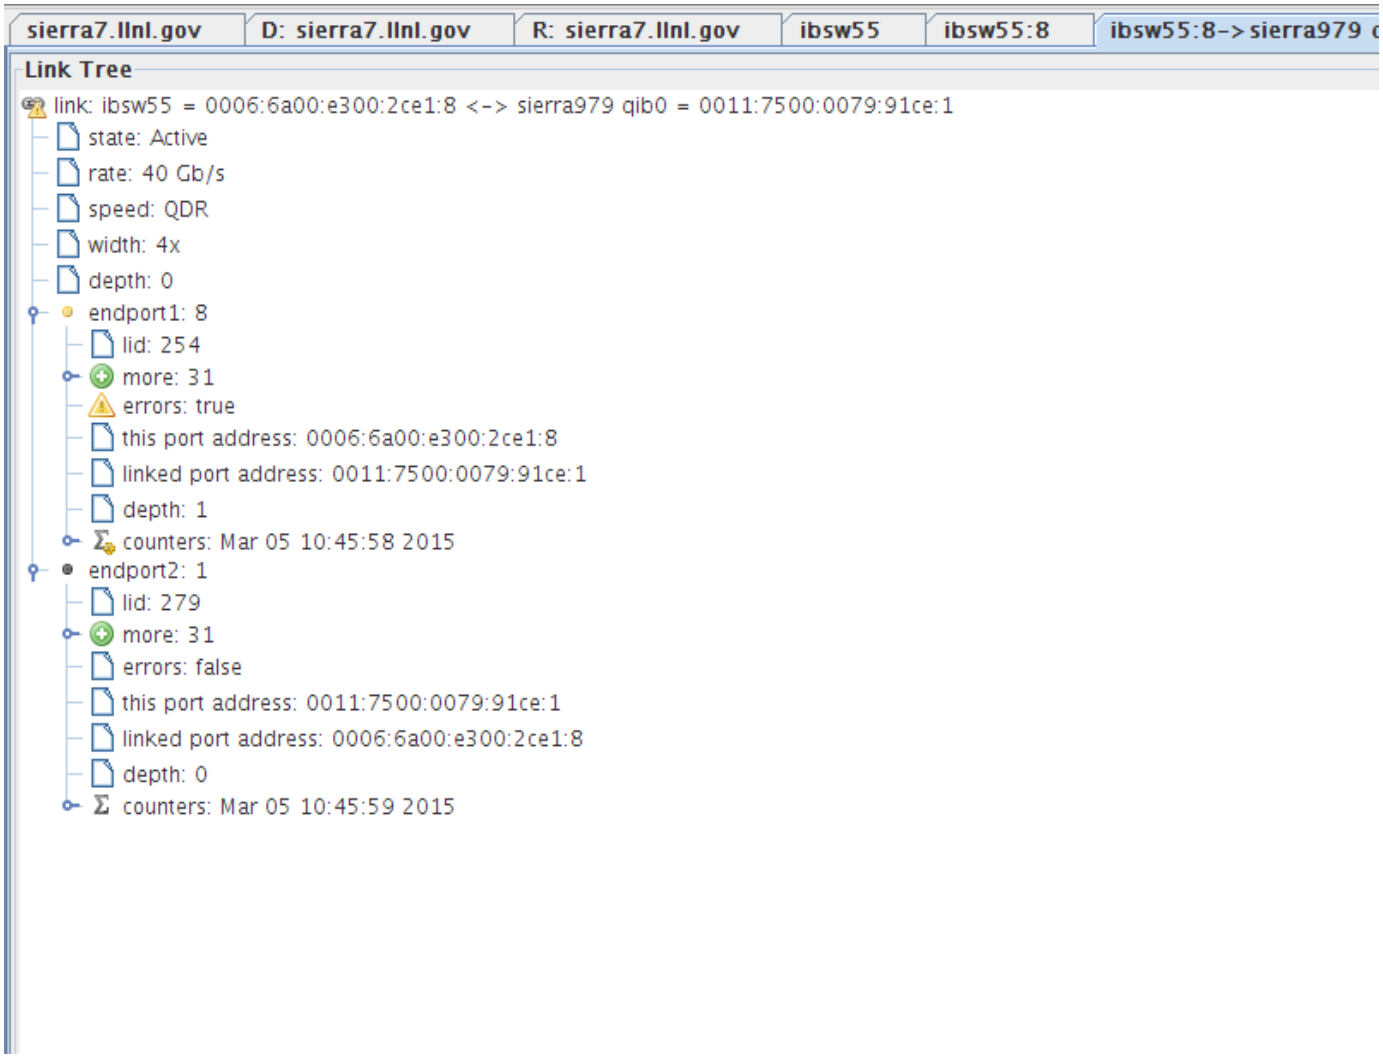
sierra7.llnl.gov
D: sierra7.llnl.gov
R: sierra7.llnl.gov
ibsw55
ibsw55:8
ibsw55:8->sierra979
Link Tree
link: ibsw55 = 0006:6a00:e300:2ce1:8 <-> sierra979 qib0 = 0011:7500:0079:91ce:1
state: Active
rate: 40 Gb/s
speed: QDR
width: 4x
depth: 0
endport1: 8
lid: 254
more: 31
errors: true
this port address: 0006:6a00:e300:2ce1:8
linked port address: 0011:7500:0079:91ce:1
depth: 1
counters: Mar 05 10:45:58 2015
endport2: 1
lid: 279
more: 31
errors: false
this port address: 0011:7500:0079:91ce:1
linked port address: 0006:6a00:e300:2ce1:8
depth: 0
counters: Mar 05 10:45:59 2015

top ports

sierra7.llnl.gov | Top Ports | G:sierra7.llnl.gov | sierra1136 qib0:1

Top Traffic Ports

time stamp: **Feb 25 10:52:38 2015**
period (secs): **300**

#	level	name	guid	port #	xmit MB/s	recv MB/s
1	0	sierra1136 qib0	0011:7500:0079:8c00	1	551	571
2	0	sierra261 qib0	0011:7500:0078:9e00	1	553	545
3	0	sierra1041 qib0	0011:7500:0079:8900	1	513	479
4	0	sierra992 qib0	0011:7500:0079:9200	1	470	488
5	0	sierra1034 qib0	0011:7500:0079:8e00	1	469	487
6	2	ibcore3 L102	0006:6a00:ec00:2a01	14	1	315
7	2	ibcore3 L102	0006:6a00:ec00:2a01	13	315	0
8	2	ibcore3 L102	0006:6a00:ec00:2a01	31	0	0
9	2	ibcore3 L102	0006:6a00:ec00:2a01	30	0	0
10	2	ibcore3 L102	0006:6a00:ec00:2a01	35	0	0
11	2	ibcore3 L102	0006:6a00:ec00:2a01	36	0	0
12	2	ibcore3 L102	0006:6a00:ec00:2a01	33	0	0
13	2	ibcore3 L102	0006:6a00:ec00:2a01	32	0	0
14	2	ibcore3 L102	0006:6a00:ec00:2a01	34	0	0
15	1	ibsw50	0006:6a00:e300:2d02	3	0	0
16	1	ibsw50	0006:6a00:e300:2d02	1	0	0
17	1	ibsw52	0006:6a00:e300:2d08	5	570	550
18	2	ibcore1 L124	0006:6a00:ec00:2b25	5	236	0
19	2	ibcore1 L124	0006:6a00:ec00:2b25	6	167	202
20	2	ibcore1 L124	0006:6a00:ec00:2b25	2	9	153

port counter

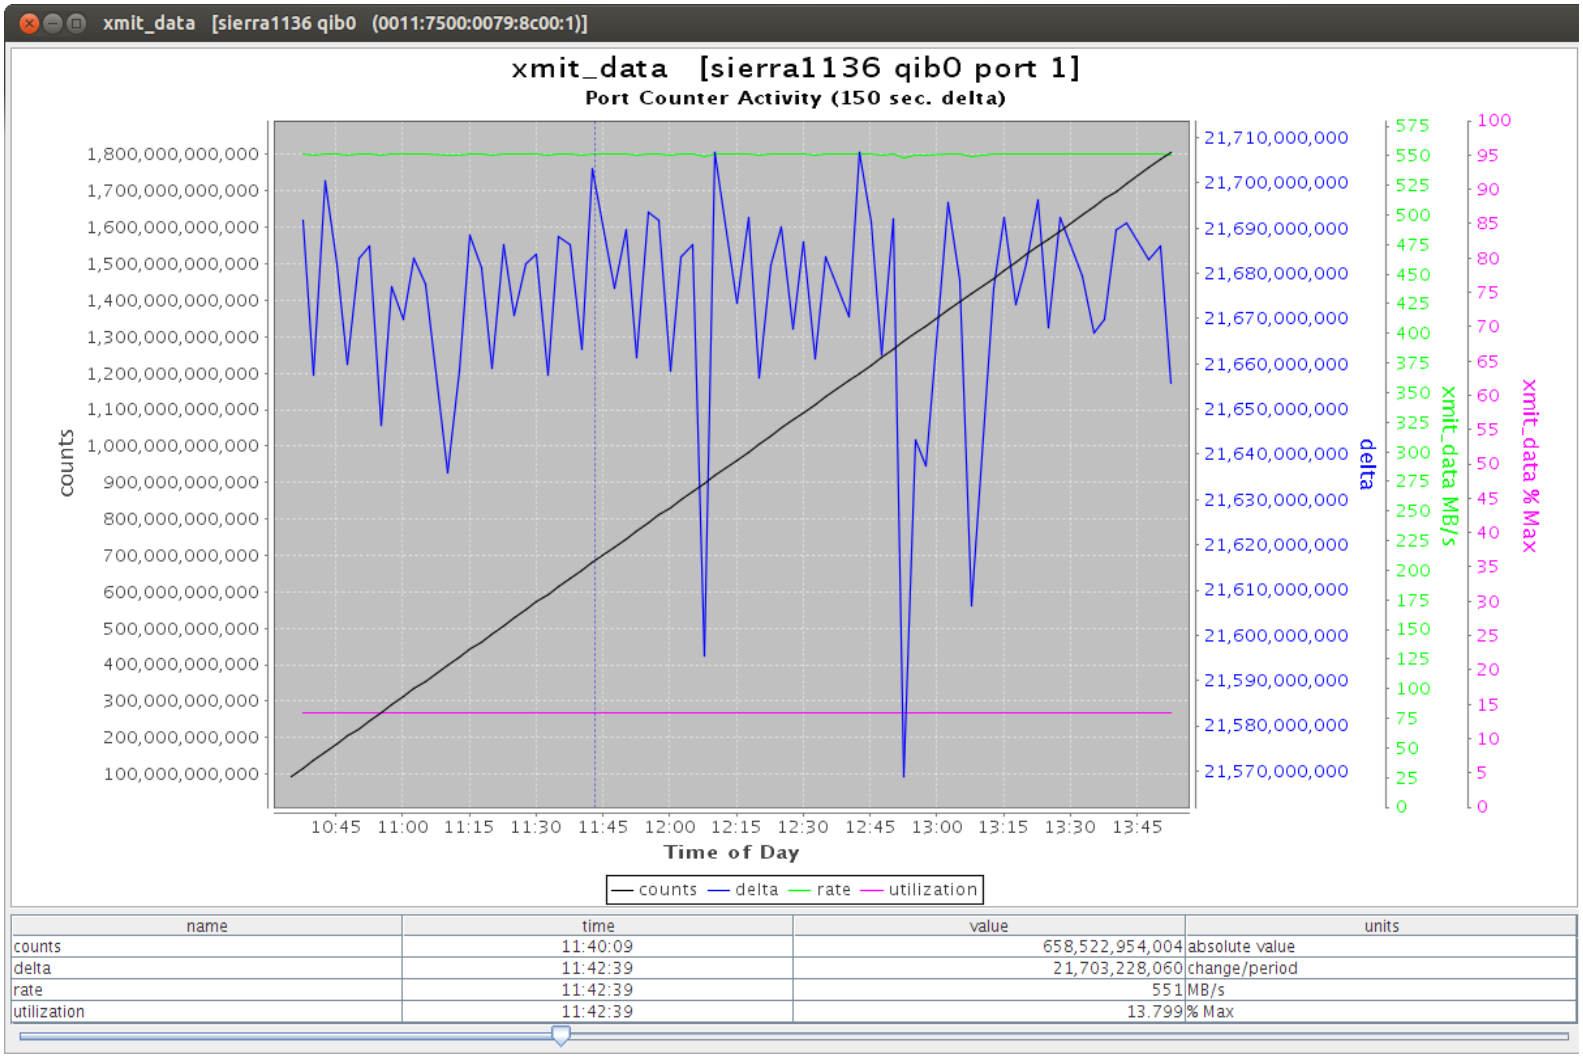

name	time	value	units
counts	11:40:09	658,522,954,004	absolute value
delta	11:42:39	21,703,228,060	change/period
rate	11:42:39	551	MB/s
utilization	11:42:39	13.799	% Max

top error links

sierra7.llnl.gov | Top Ports | G:sierra7.llnl.gov | sierra1136 qib0:1 | Error Links | ibsw58:20->ibcore1 L124:4

Top Error Links

time stamp: **Feb 25 10:57:38 201**
period (secs): **30**

Links with Errors: 138

Supressed Errors: **rcv_rem_phys_err, rcv_switch_relay_err**

#	level	link identification	delta error/period=p1,p2
1	1	ibsw2 0006:6a00:e300:2c27:32->0006:6a00:ec00:2a01: 2 ibcore3 L102	symbol_err_cnt=0,1
2	1	ibsw38 0006:6a00:e300:2bcc:21->0006:6a00:ec00:2a03: 2 ibcore1 L225	symbol_err_cnt=0,1
3	1	ibsw52 0006:6a00:e300:2d08:21->0006:6a00:ec00:2a03:16 ibcore1 L225	symbol_err_cnt=0,3
4	1	ibsw24 0006:6a00:e300:2b69:19->0006:6a00:ec00:2914: 6 ibcore1 L107	symbol_err_cnt=0,2
5	1	ibsw34 0006:6a00:e300:2c96:19->0006:6a00:ec00:2914:16 ibcore1 L107	symbol_err_cnt=0,2
6	1	ibsw61 0006:6a00:e300:307b:17->0006:6a01:e900:0122: 7 ibcore2 L128	symbol_err_cnt=0,6
7	1	ibsw71 0006:6a00:e300:2b67:16->0006:6a02:e900:0122:17 ibcore2 L127	symbol_err_cnt=0,21
8	1	ibsw72 0006:6a00:e300:2a99:16->0006:6a02:e900:0122:18 ibcore2 L127	symbol_err_cnt=0,7
9	1	ibsw66 0006:6a00:e300:2cd7: 3->0006:6a00:ec00:2923:12 ibcore2 L123	symbol_err_cnt=0,5
10	1	ibsw58 0006:6a00:e300:2ceb:20->0006:6a00:ec00:2b25: 4 ibcore1 L124	symbol_err_cnt=0,34
11	1	ibsw38 0006:6a00:e300:2bcc:32->0006:6a00:ec00:2b27: 2 ibcore3 L224	symbol_err_cnt=0,5
12	1	ibsw2 0006:6a00:e300:2c27:15->0006:6a00:ec00:2907: 2 ibcore2 L104	symbol_err_cnt=1,0
13	1	ibsw23 0006:6a00:e300:2a9b:35->0006:6a00:ec00:2b30: 5 ibcore3 L119	symbol_err_cnt=0,2
14	1	ibsw95 0006:6a00:e300:2c39:35->0006:6a00:ec00:2b31: 5 ibcore3 L221	symbol_err_cnt=0,1
15	1	ibsw99 0006:6a00:e300:2b60:35->0006:6a00:ec00:2b31: 9 ibcore3 L221	symbol_err_cnt=0,2
16	1	ibsw106 0006:6a00:e300:2cab:35->0006:6a00:ec00:2b31:16 ibcore3 L221	symbol_err_cnt=0,9
17	1	ibsw9 0006:6a00:e300:2a53:17->0006:6a00:ec00:2733: 9 ibcore2 L106	symbol_err_cnt=0,1
18	1	ibsw19 0006:6a00:e300:2be1:34->0006:6a00:ec00:2b33: 1 ibcore3 L110	symbol_err_cnt=0,1
19	1	ibsw30 0006:6a00:e300:2752:34->0006:6a00:ec00:2b33:12 ibcore3 L110	symbol_err_cnt=0,1
20	1	ibsw14 0006:6a00:e300:2c36: 2->0006:6a00:ec00:29cd:14 ibcore1 L106	rcv_rem_phys_err=1,0

link tree

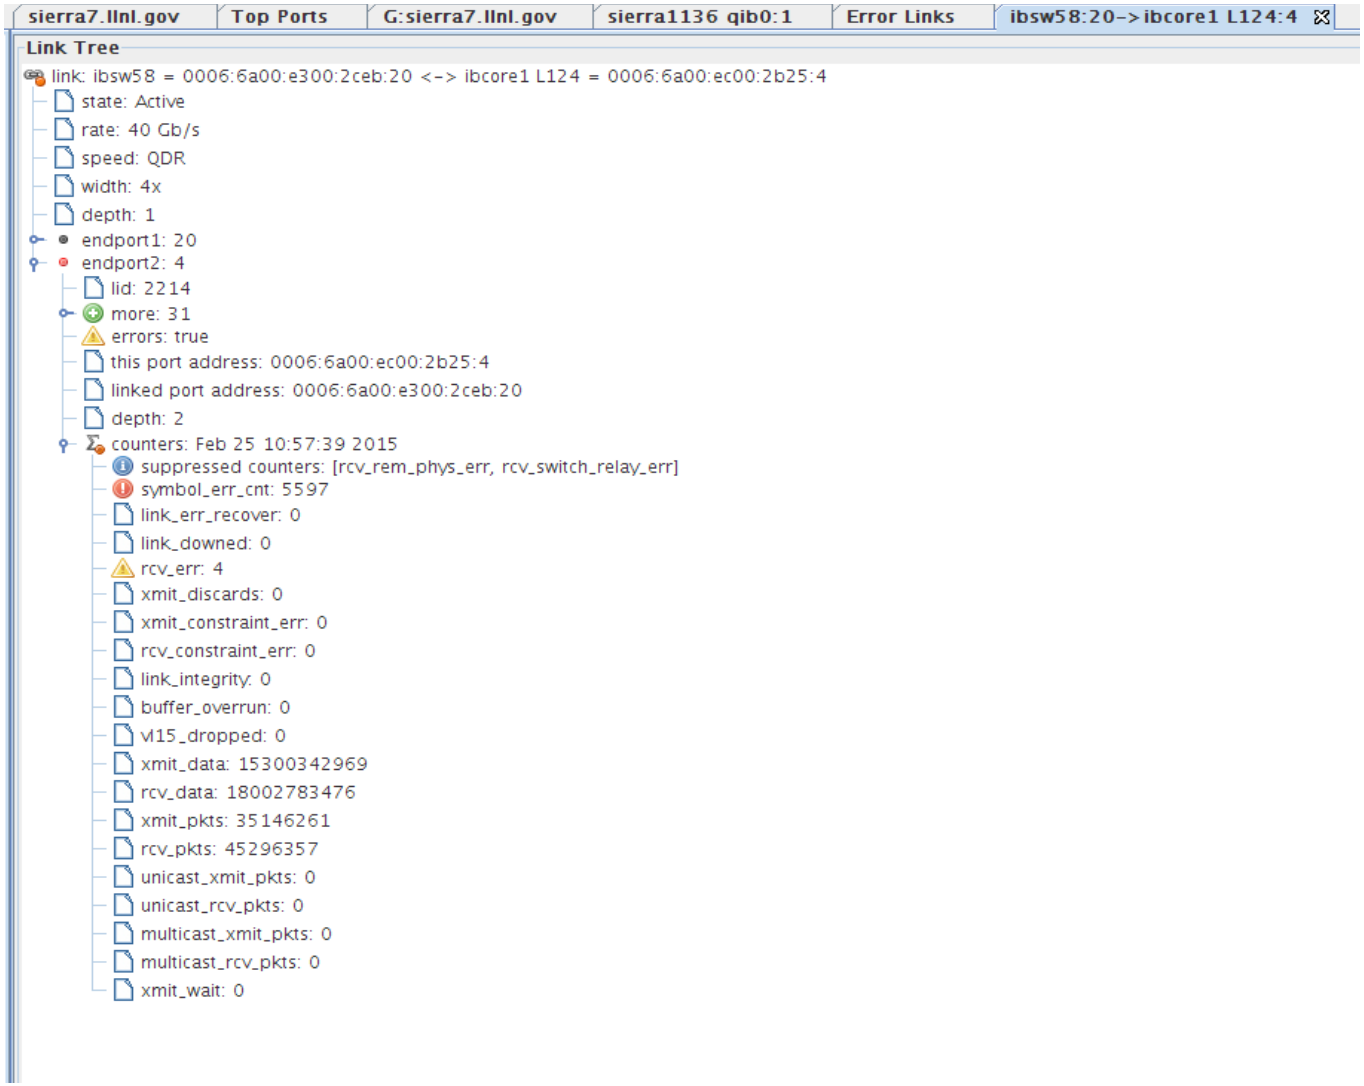
sierra7.llnl.gov
Top Ports
G:sierra7.llnl.gov
sierra1136 qib0:1
Error Links
ibsw58:20->ibcore1 L124:4
Link Tree
link: ibsw58 = 0006:6a00:e300:2ceb:20 <-> ibcore1 L124 = 0006:6a00:ec00:2b25:4
state: Active
rate: 40 Gb/s
speed: QDR
width: 4x
depth: 1
endport1: 20
endport2: 4
lid: 2214
more: 31
errors: true
this port address: 0006:6a00:ec00:2b25:4
linked port address: 0006:6a00:e300:2ceb:20
depth: 2
counters: Feb 25 10:57:39 2015
suppressed counters: [rcv_rem_phys_err, rcv_switch_relay_err]
symbol_err_cnt: 5597
link_err_recover: 0
link_downed: 0
rcv_err: 4
xmit_discards: 0
xmit_constraint_err: 0
rcv_constraint_err: 0
link_integrity: 0
buffer_overrun: 0
vl15_dropped: 0
xmit_data: 15300342969
rcv_data: 18002783476
xmit_pkts: 35146261
rcv_pkts: 45296357
unicast_xmit_pkts: 0
unicast_rcv_pkts: 0
multicast_xmit_pkts: 0
multicast_rcv_pkts: 0
xmit_wait: 0

port error

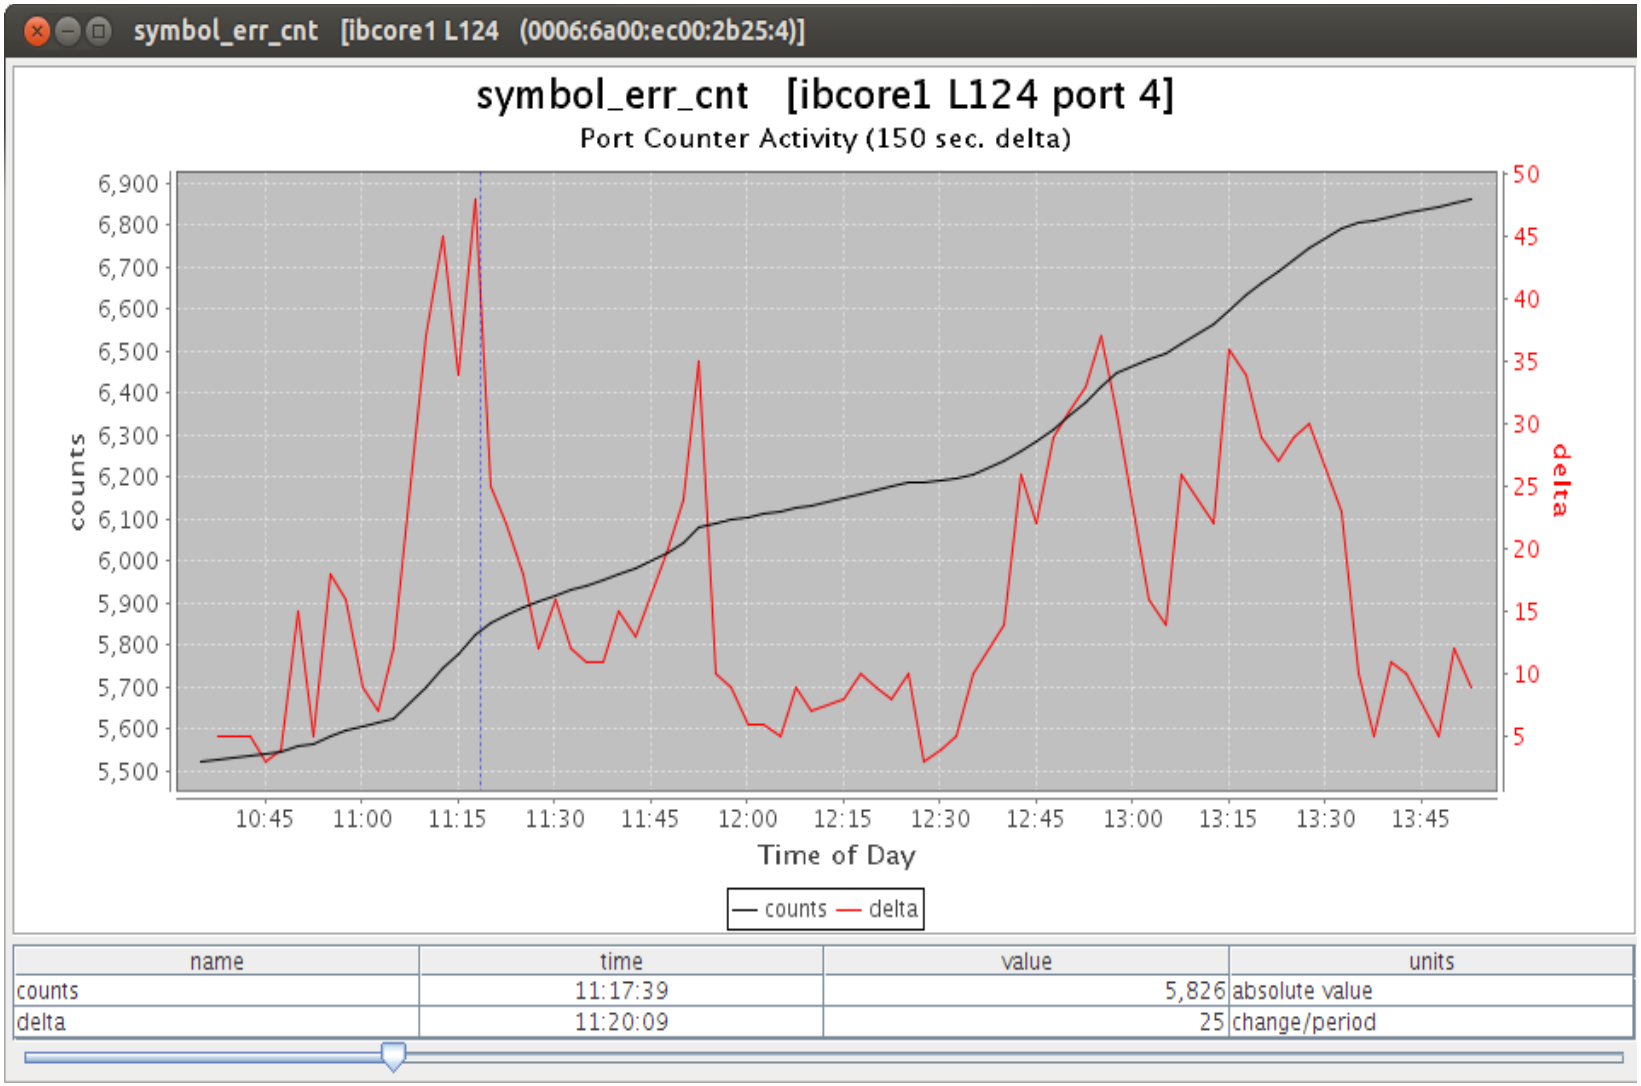

name	time	value	units
counts	11:17:39	5,826	absolute value
delta	11:20:09	25	change/period

atlas

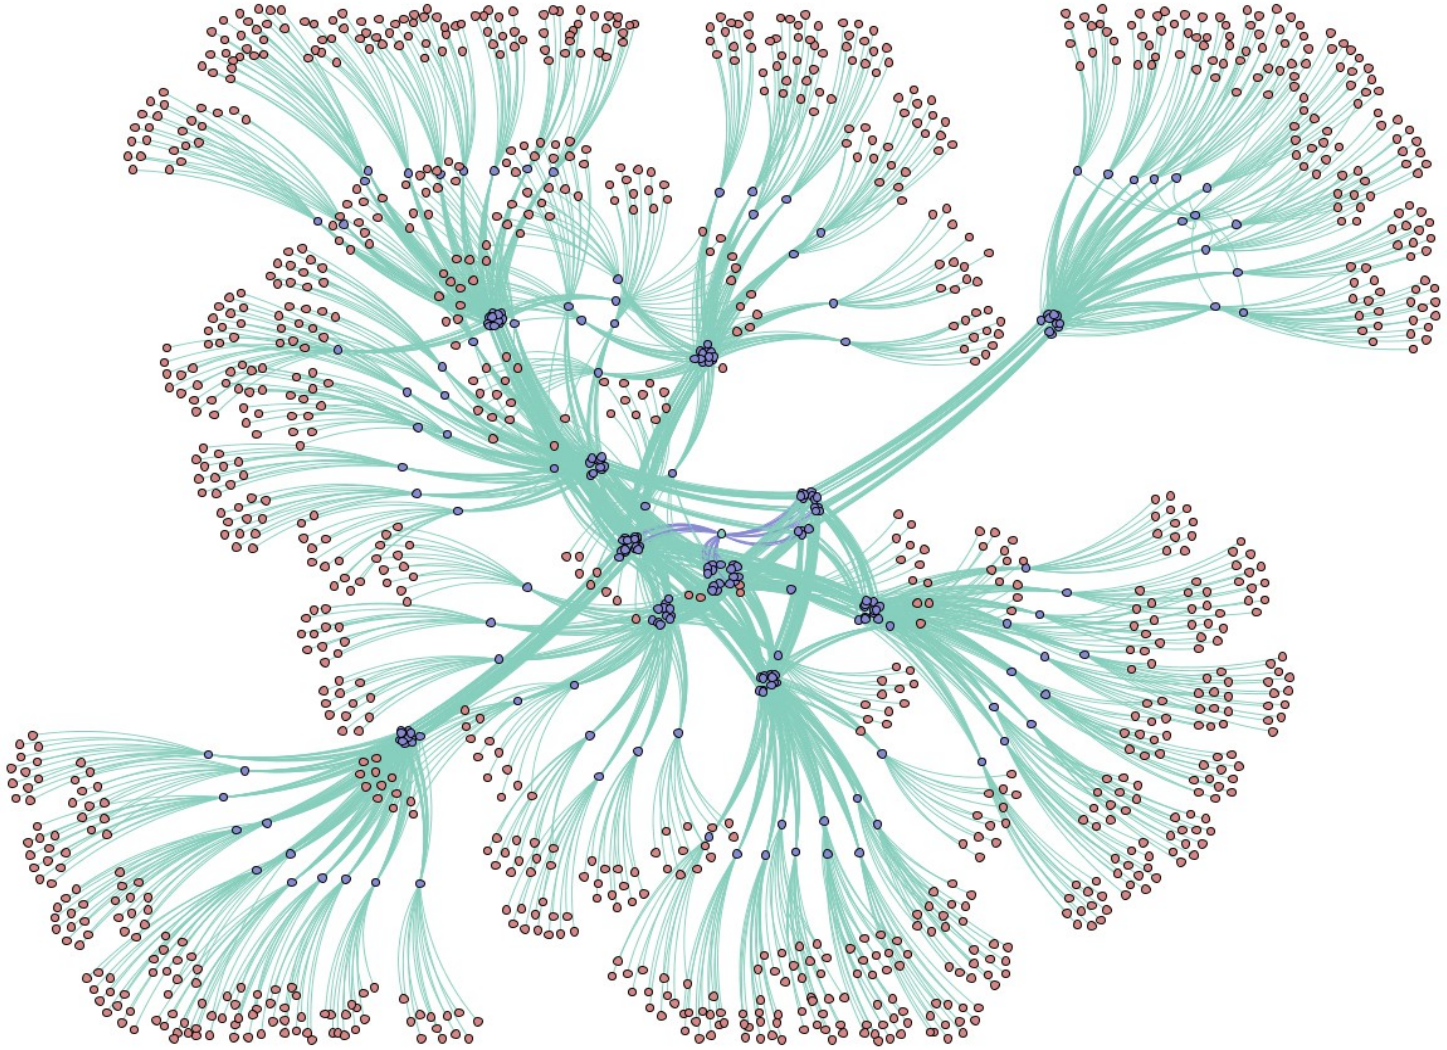

prism

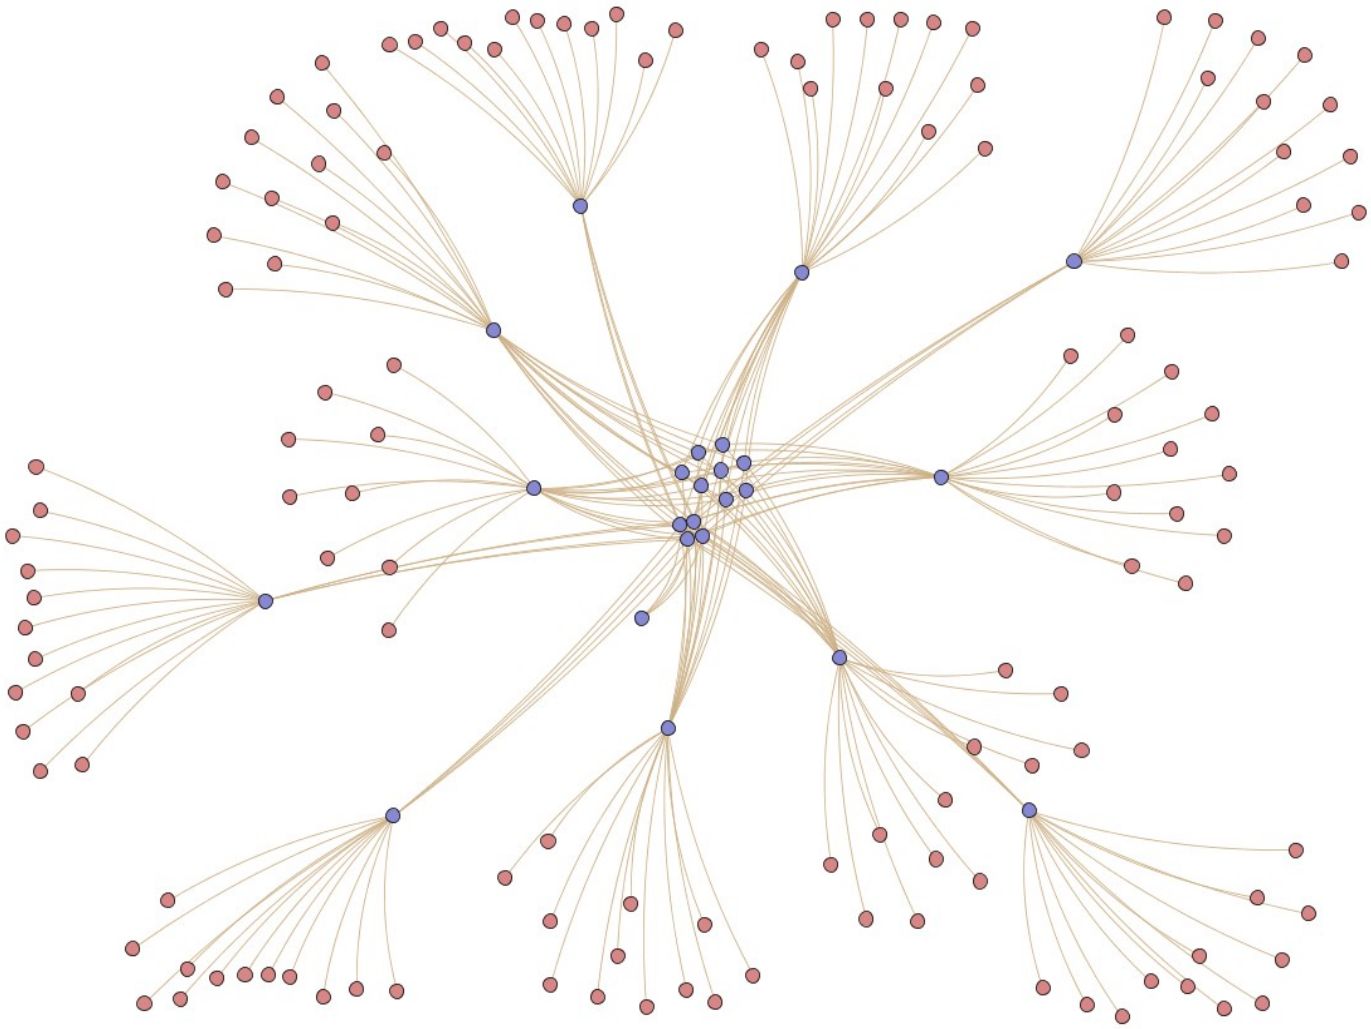

hera

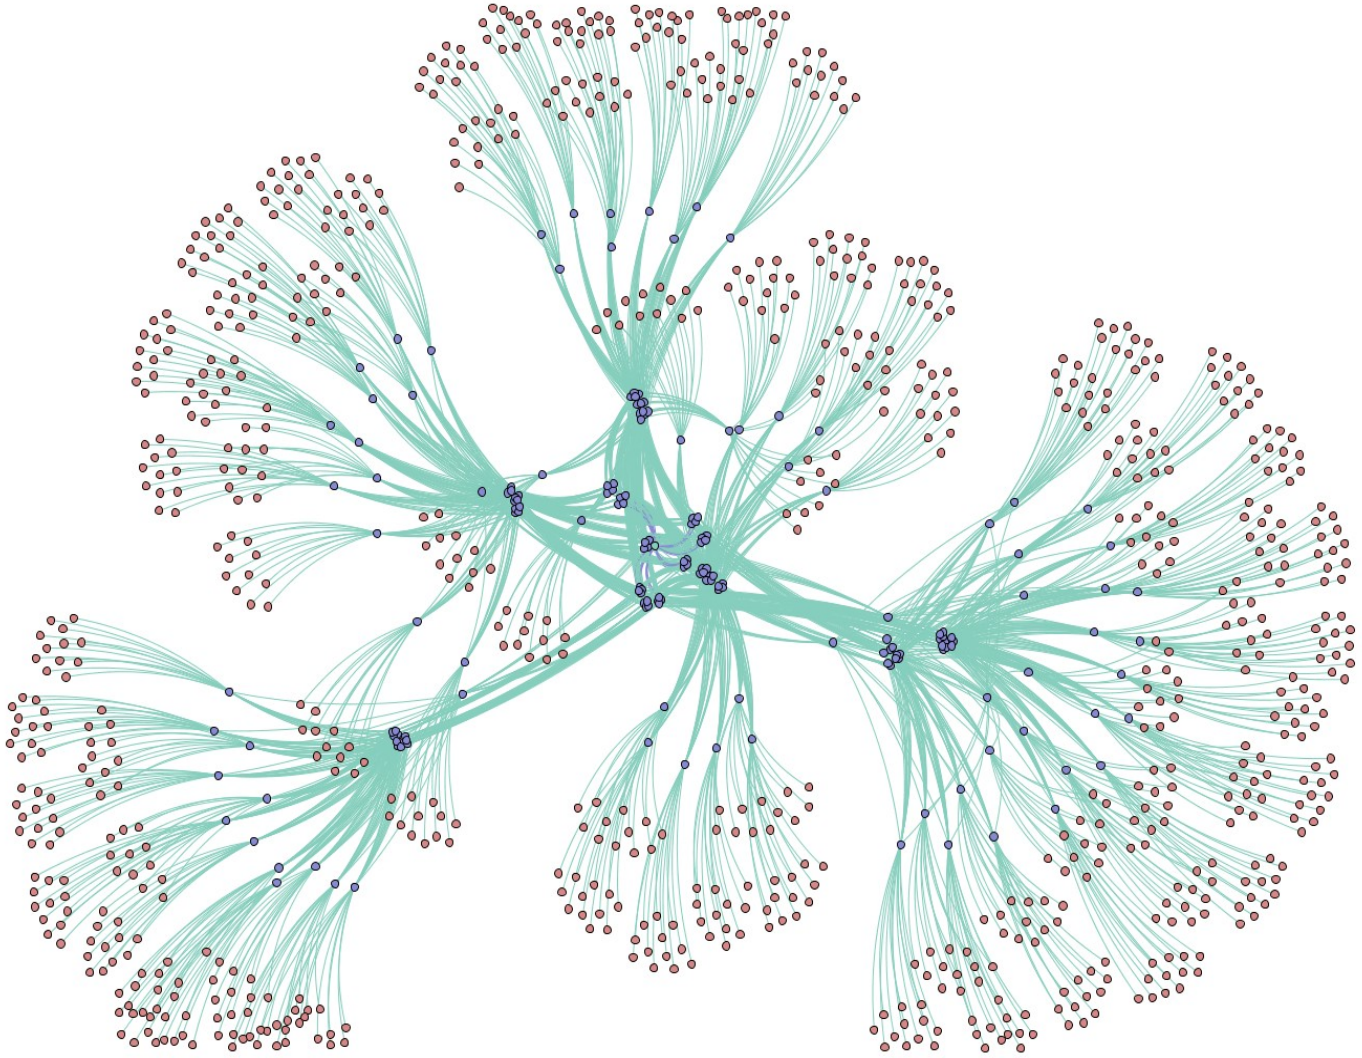

sierra

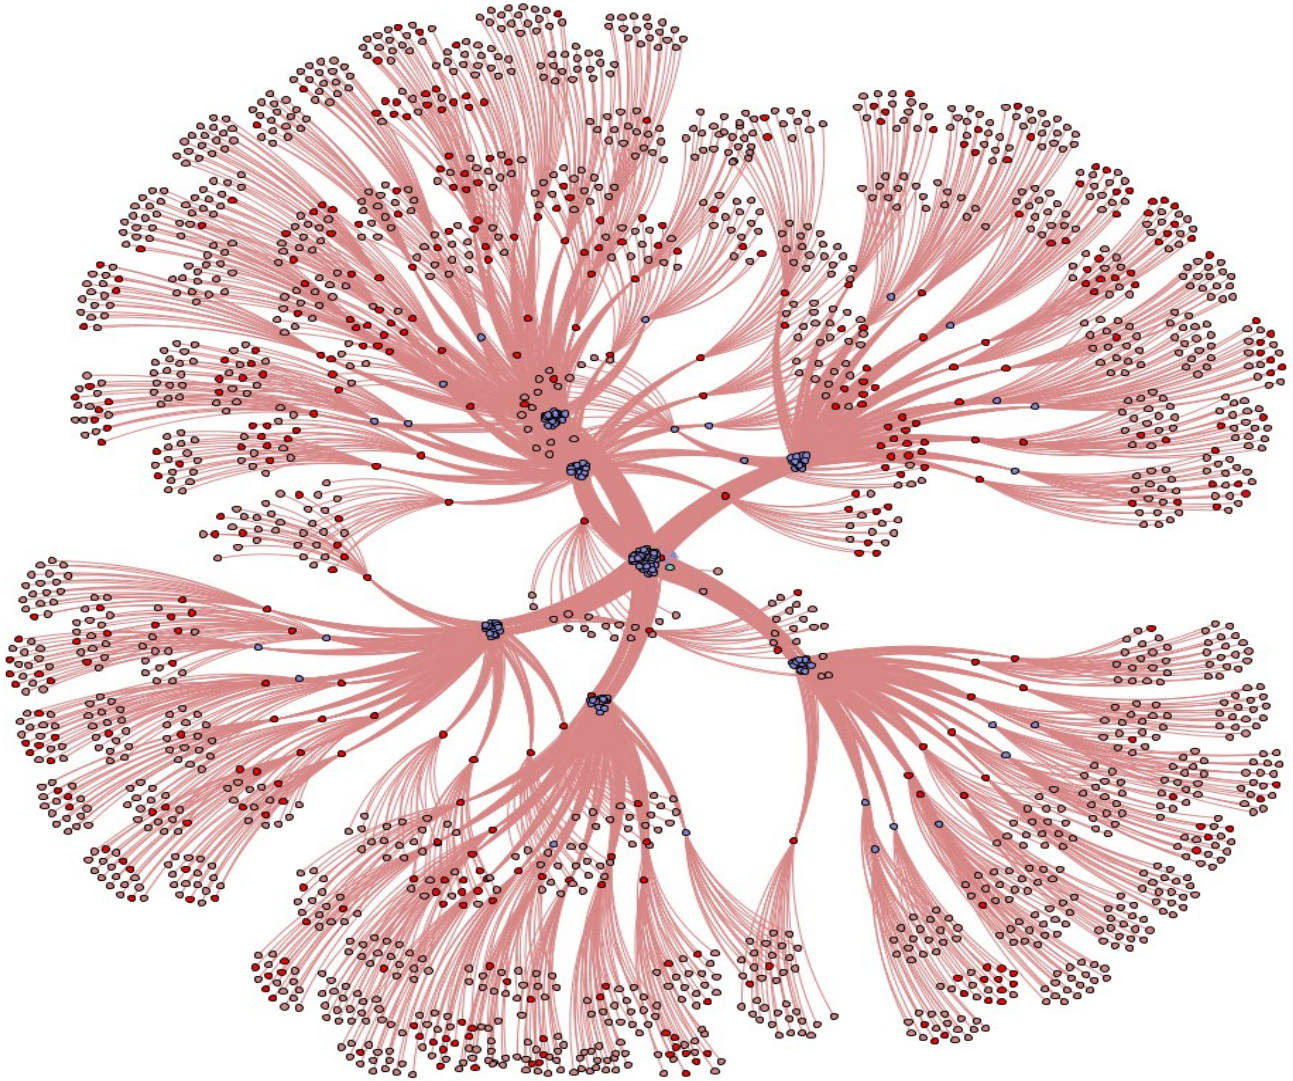

hype

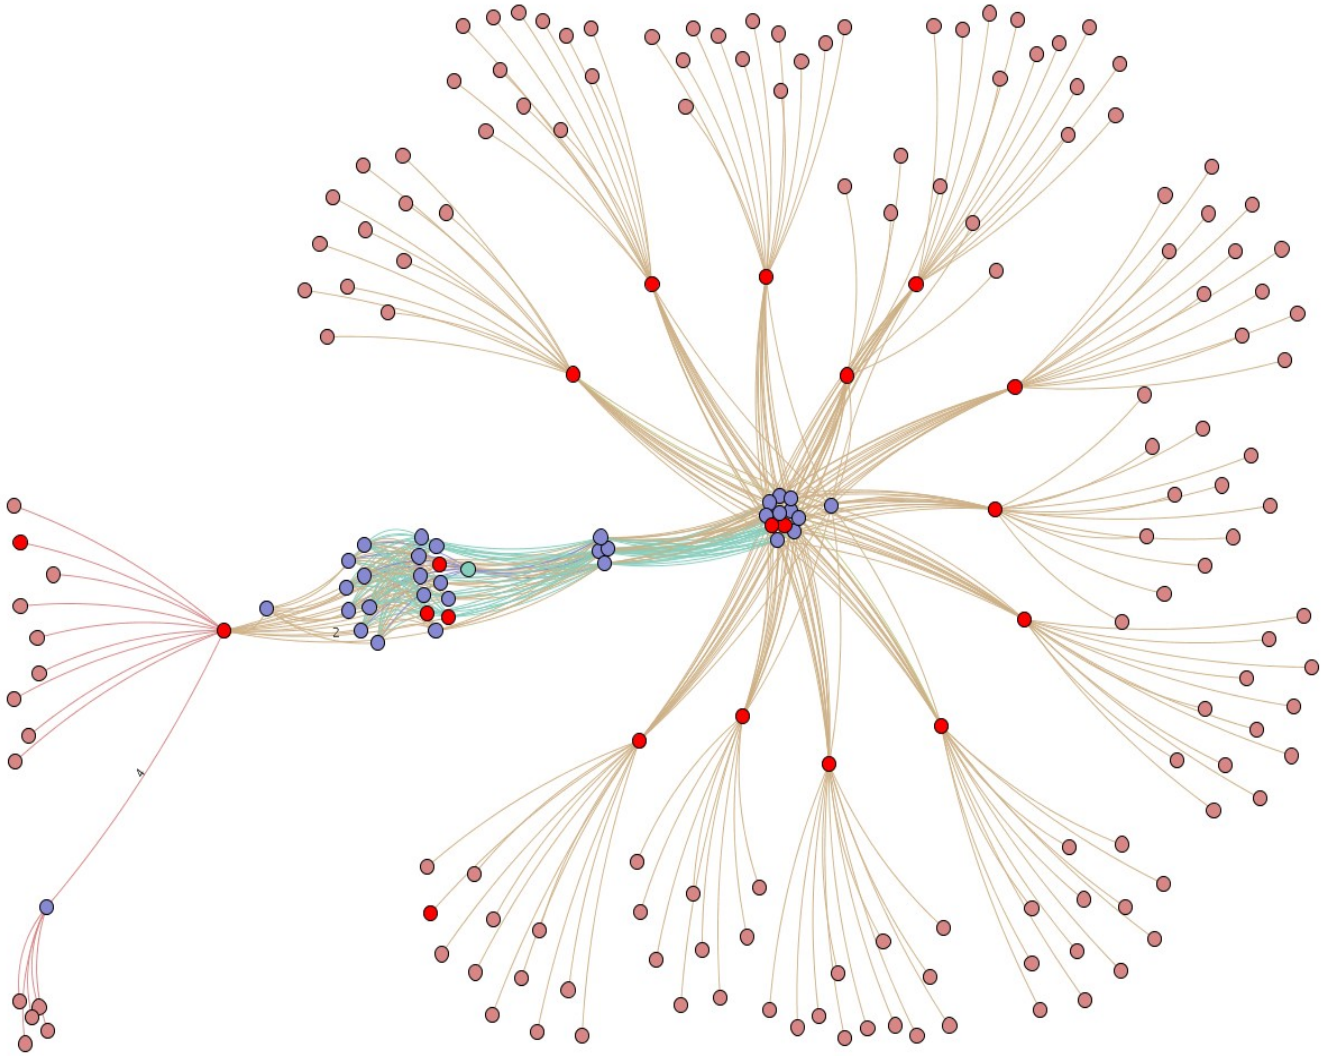

grove

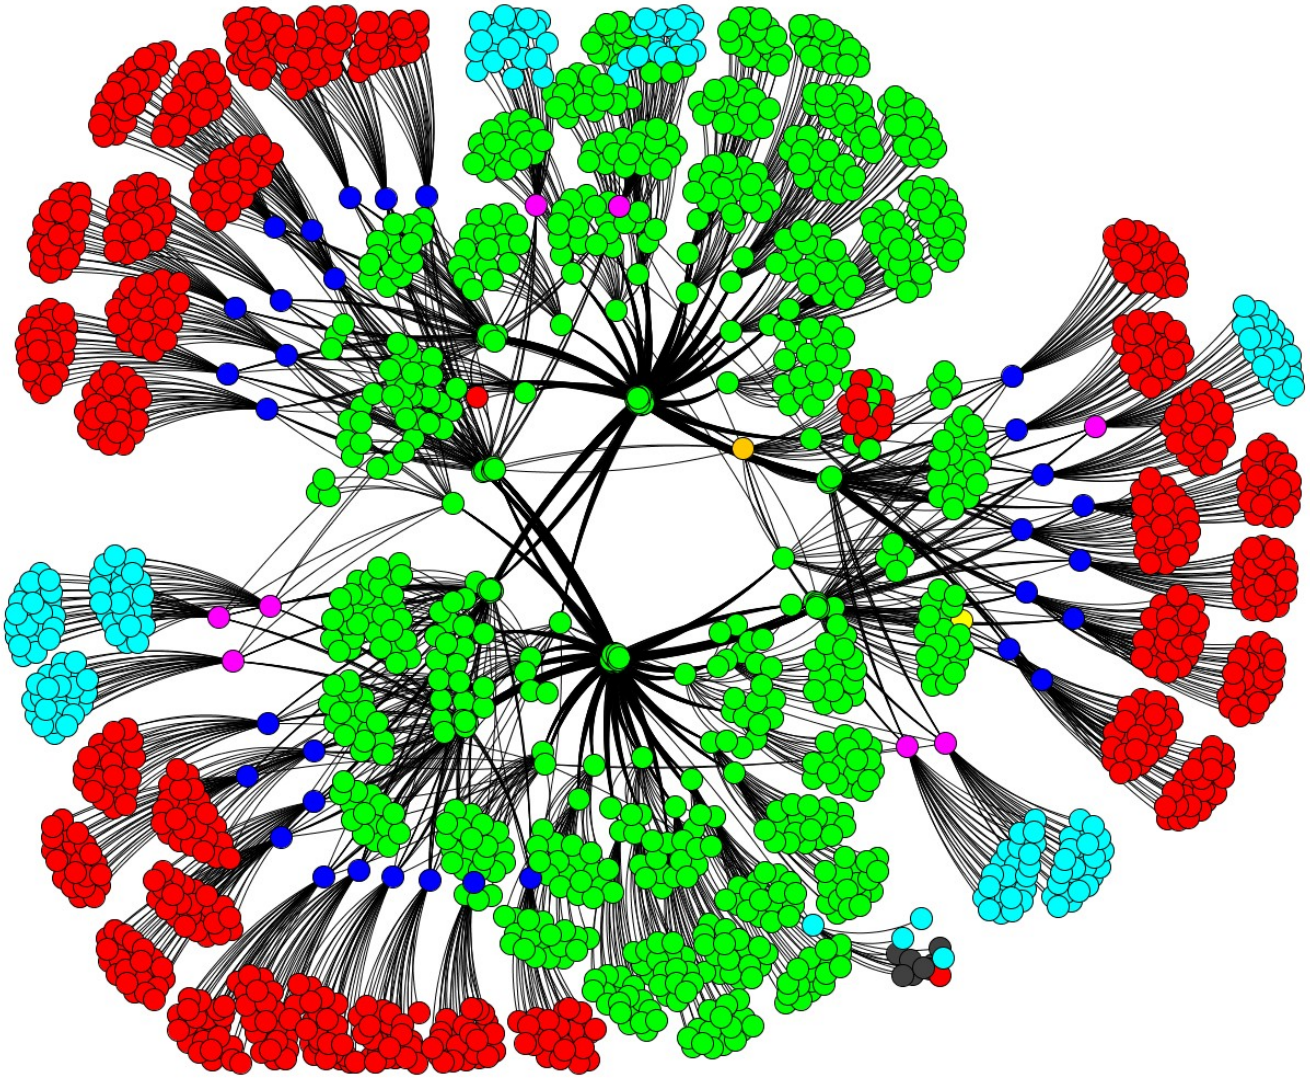

dynamic fabric graph

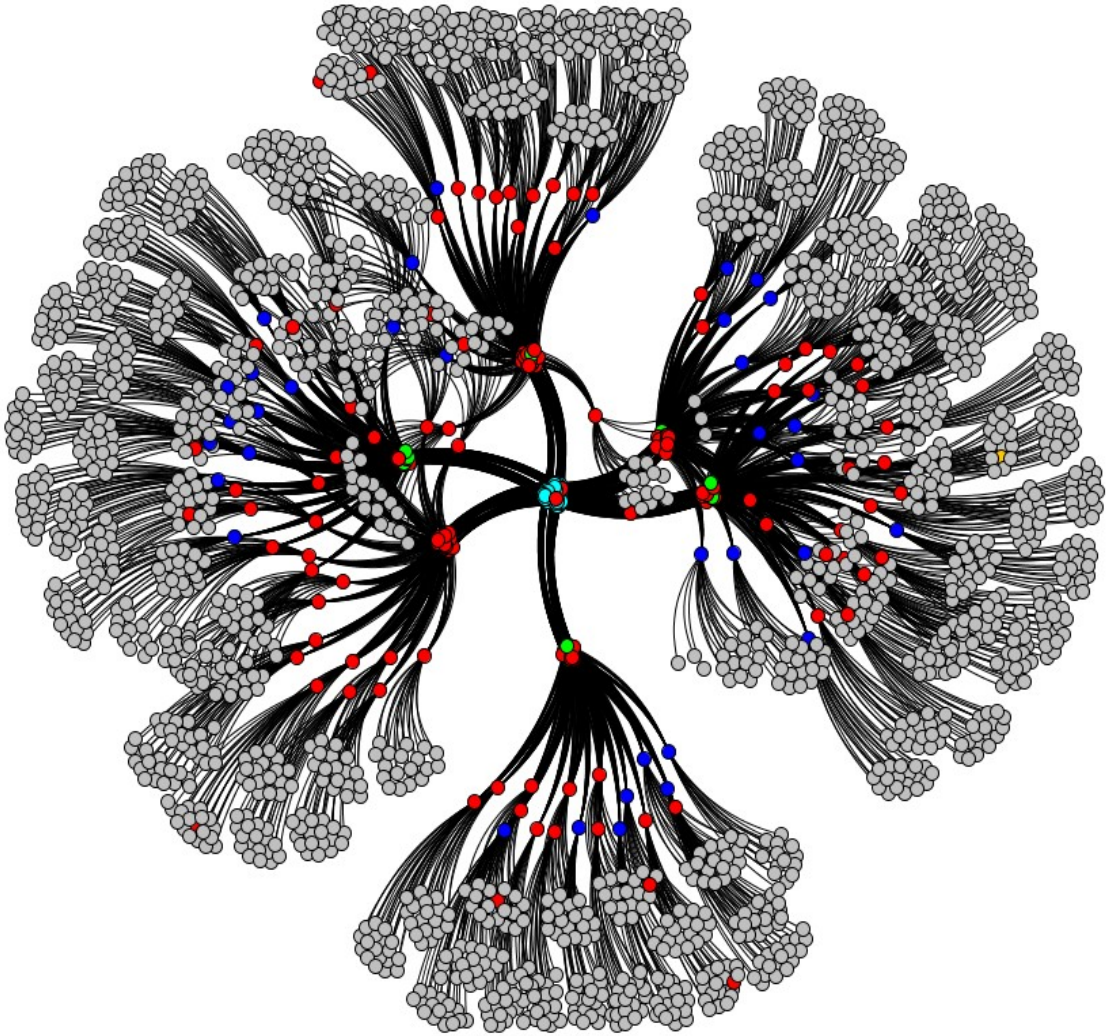

path selection

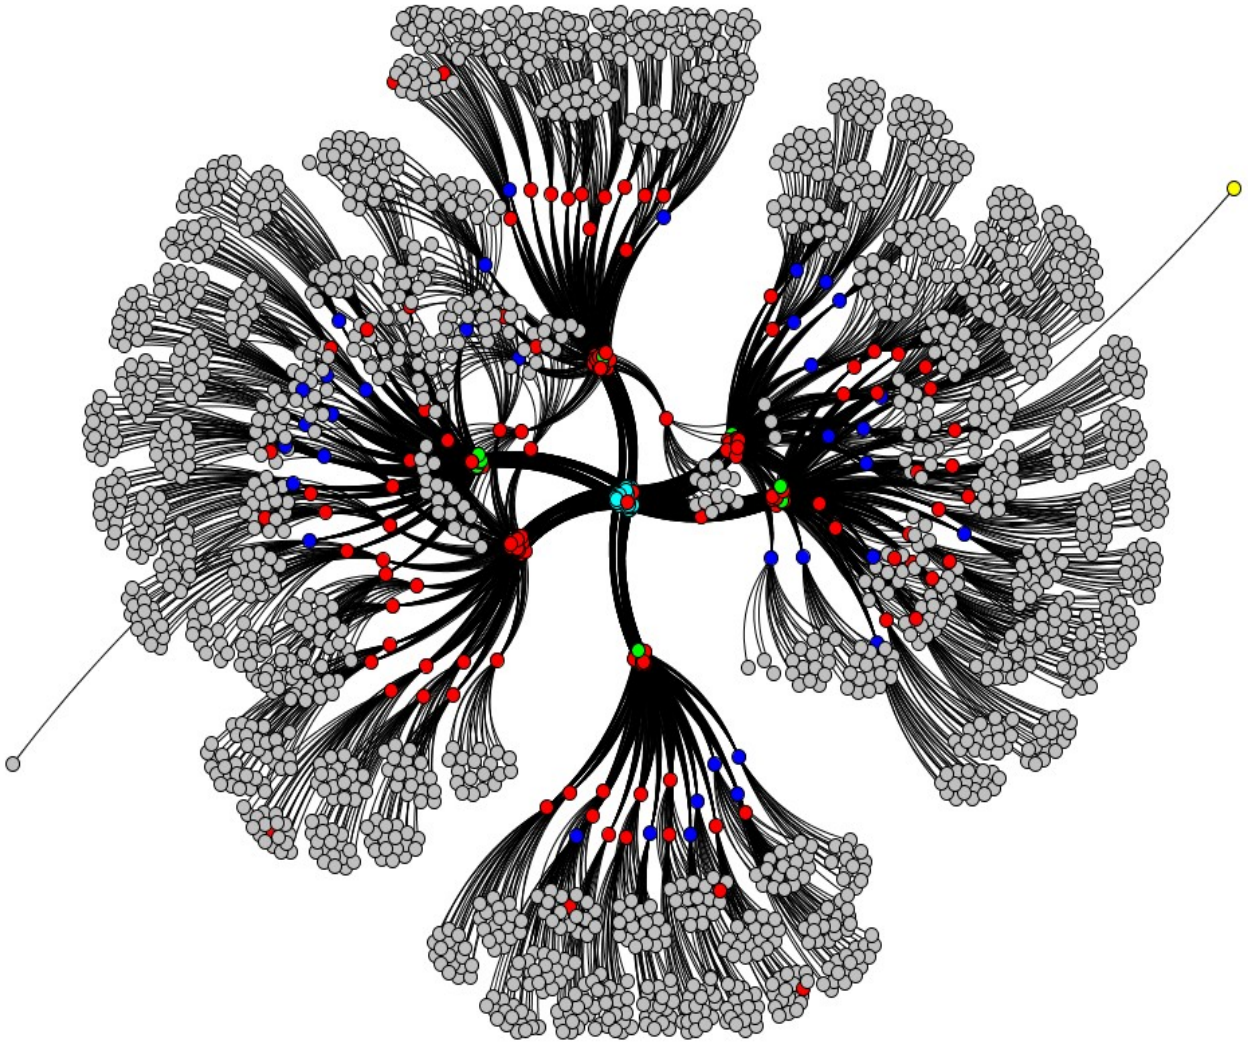

path selection (decorated)

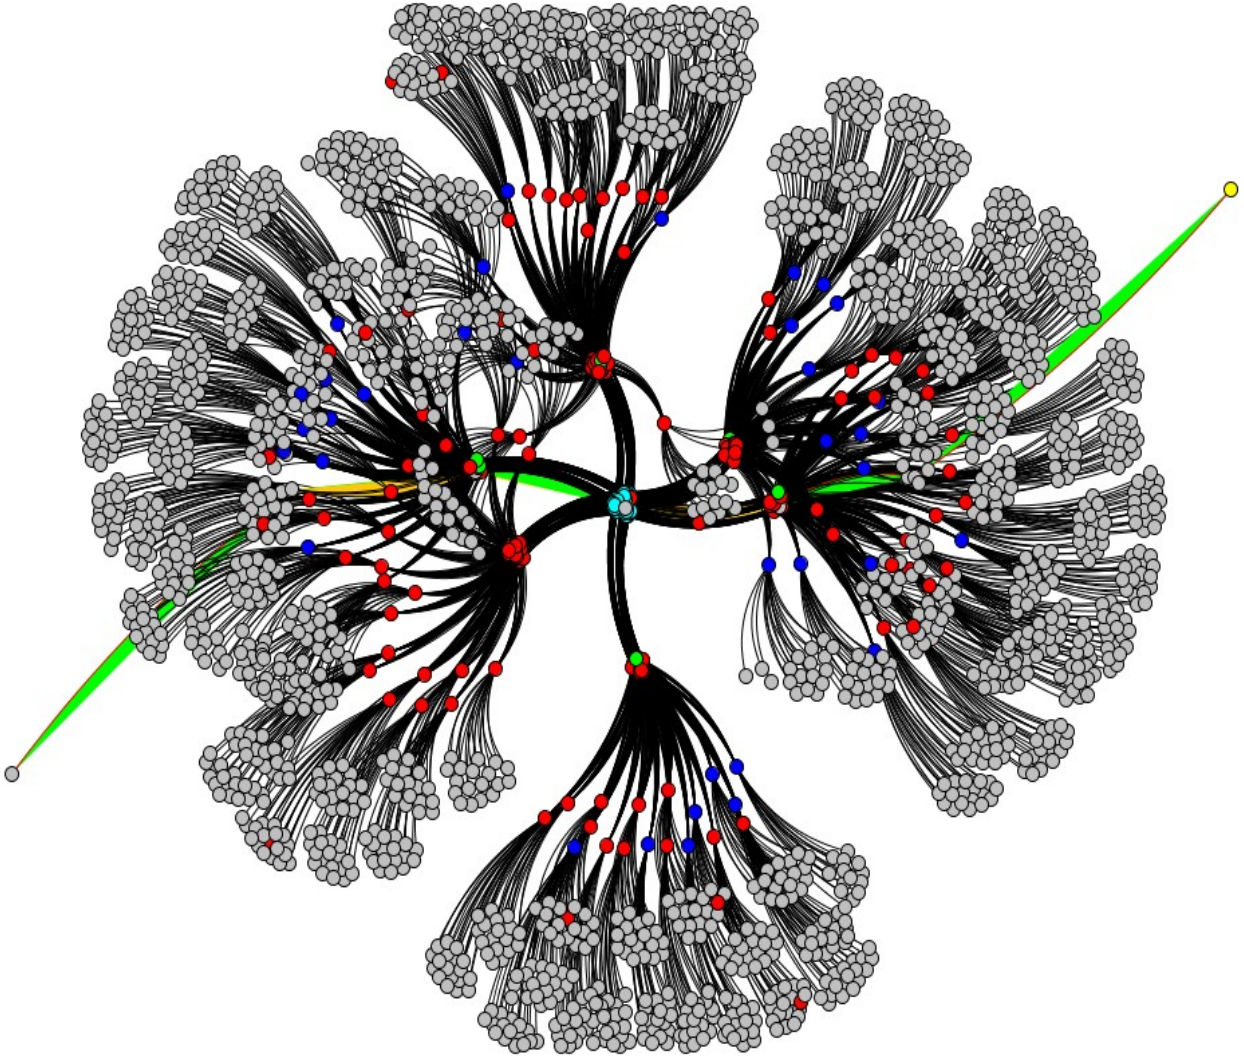

path selection (revealed)

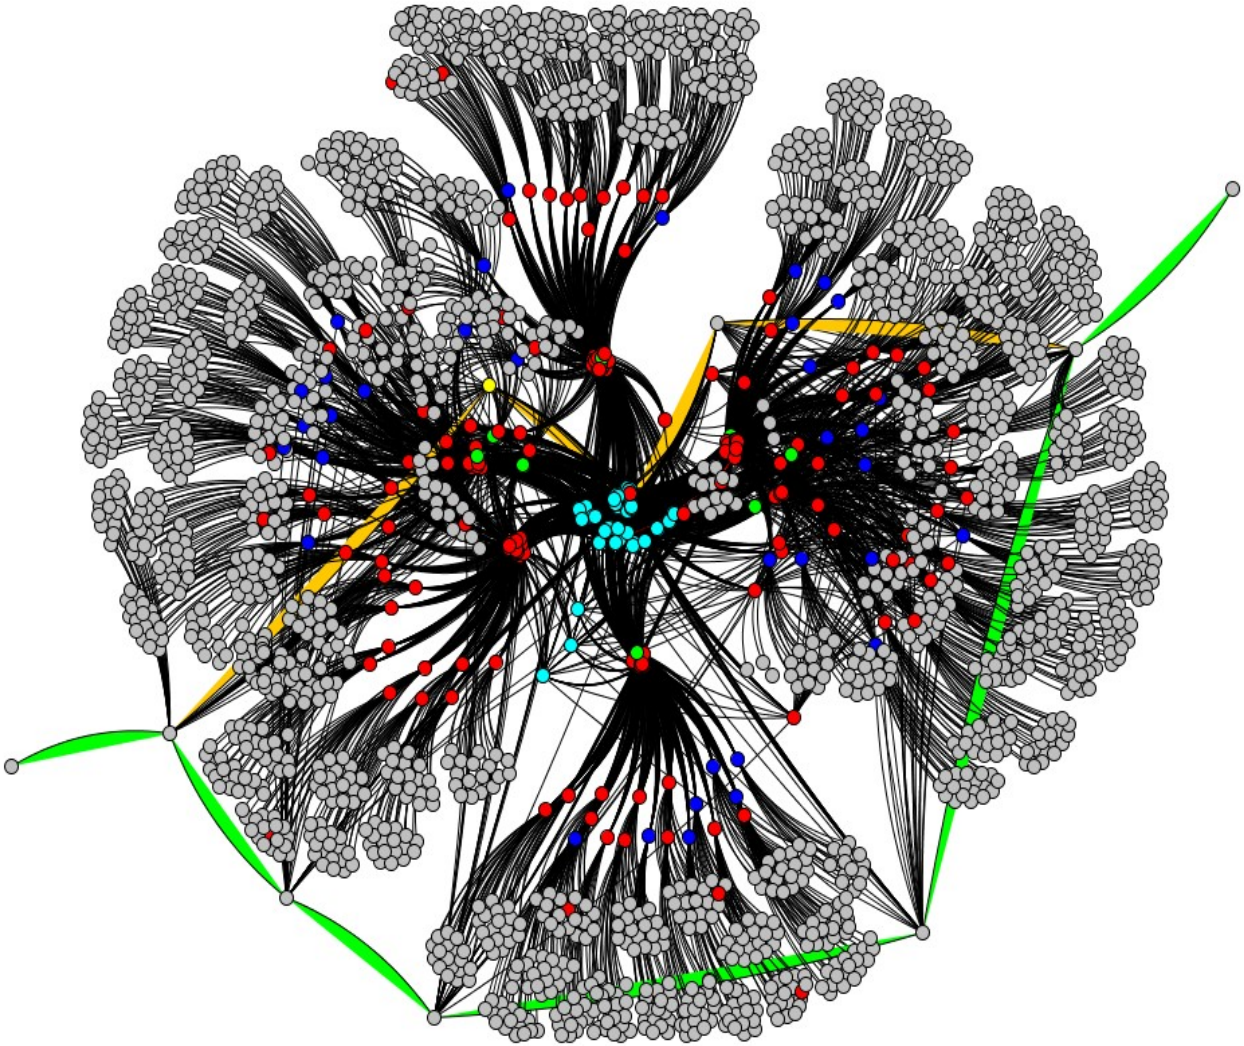

path tree (trace route)

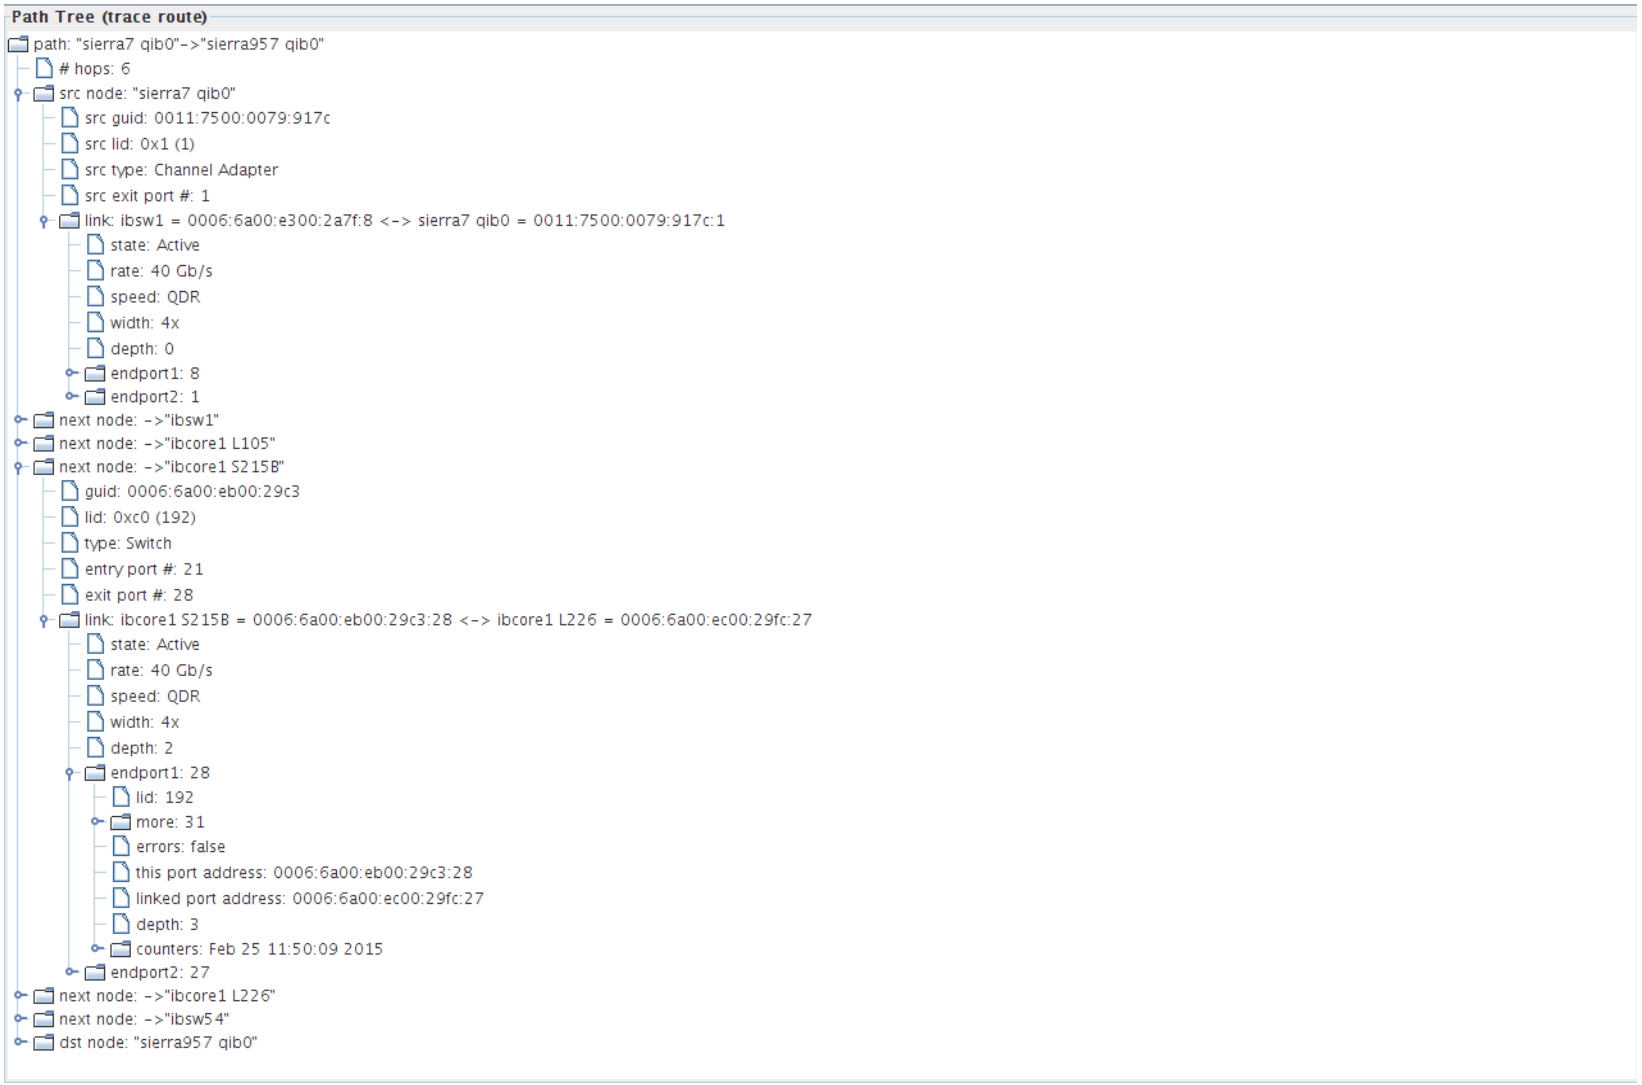

path utilization

Path Utilization

time stamp: **Feb 25 11:50:08 2015**
period (secs): **1050**

Transmit Path (to dst):

path id: **"sierra7 qib0"->"sierra957 qib0"**

hop	node	output port	xmit delta (counts)	xmit rate	units	% rate
src-0	sierra7 qib0	0011:7500:0079:917c:1	12104328	0	MB/s	0.00 %
1	ibsw1	0006:6a00:e300:2a7f:1	217792717	0	MB/s	0.02 %
2	ibcore1 L105	0006:6a00:ec00:29cc:21	11176330	0	MB/s	0.00 %
3	ibcore1 S215B	0006:6a00:eb00:29c3:28	557086	0	MB/s	0.00 %
4	ibcore1 L226	0006:6a00:ec00:29fc:18	1211663	0	MB/s	0.00 %
5	ibsw54	0006:6a00:e300:2cd6:26	132015656081	479	MB/s	11.99 %
dst-6	sierra957 qib0	0011:7500:0079:91aa				

Receive Path (from dst):

path id: **"sierra957 qib0"->"sierra7 qib0"**

hop	node	output port	xmit delta (counts)	xmit rate	units	% rate
src-0	sierra957 qib0	0011:7500:0079:91aa:1	141345839713	513	MB/s	12.84 %
1	ibsw54	0006:6a00:e300:2cd6:1	1843181	0	MB/s	0.00 %
2	ibcore1 L121	0006:6a00:ec00:2959:19	2790935	0	MB/s	0.00 %
3	ibcore1 S113A	0006:6a00:eb00:28a7:11	6019342	0	MB/s	0.00 %
4	ibcore1 L104	0006:6a00:ec00:29b4:1	7609414	0	MB/s	0.00 %
5	ibsw1	0006:6a00:e300:2a7f:8	12101012	0	MB/s	0.00 %
dst-6	sierra7 qib0	0011:7500:0079:917c				

xmit_data: 11.99%

Concluding Remarks

- OpenSM maintains a substantial amount of fabric information
- smt-gui is only one of many SMT commands
- most commands are NOT gui based
- most commands have dual operating modes
- all rely on the OpenSM Monitor Service (OMS)
- future plans?
  - smt-agents for other monitoring, analysis and visualization tools (such as SPLUNK and other internal LLNL systems)
  - enhanced support for congestion management, partitions, multicast groups, etc.
  - open to other ideas
- availability?
  - included in the TOSS distribution
  - expected to be on GitHub this year

Questions?

Thank You

Lawrence Livermore National Laboratory