802.1aq/SPB - 3rd Interop Overview

Onsite
Edgard Vargas – Alcatel Lucent
Peter Ashwood-Smith - Huawei
John Vant Erve - Avaya
Stephane Vinet - Spirent
Bis Nandy – Solana
Rup Makkar - Solana

Date
June 27-30 Ottawa/Canada

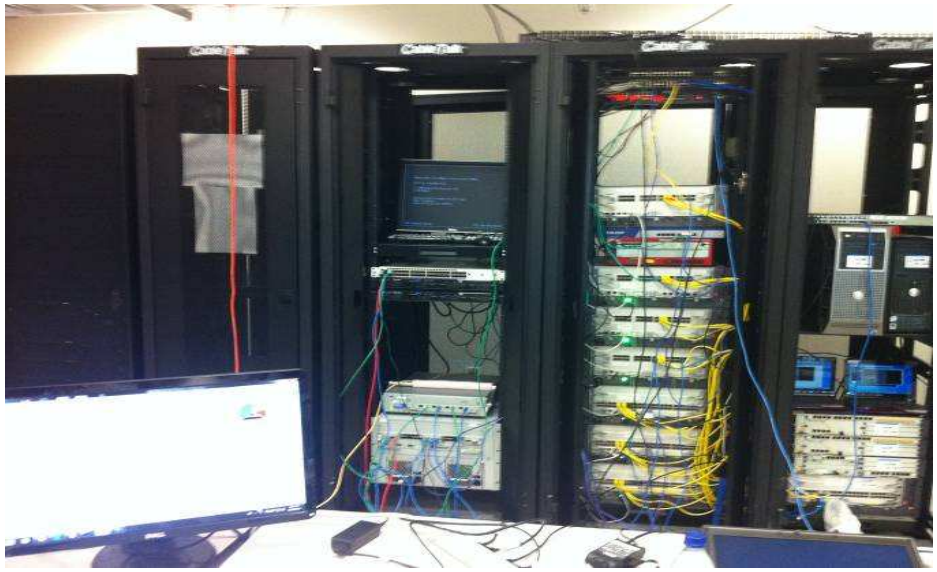

802.1aq Standardization

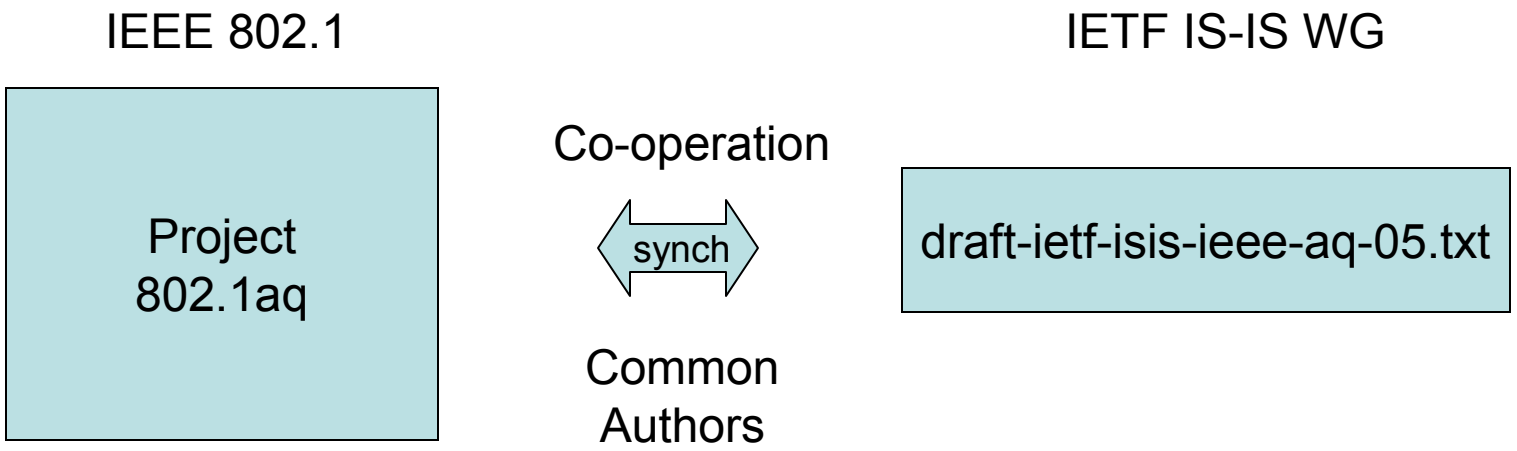

Demonstration of Implementation Interoperability

Interop Background (SPBM)

- 1st Interop was 2010 Avaya/Huawei
  - Simple control plane/data/.1ag – 5 nodes.
  - Report available via the wikipedia 802.1aq page.
- 2nd Interop was in January 2011 Avaya/Huawei
  - Full data plane/control/.1ag – 9 nodes including VM moves (report pending..)
- 3rd Interop was in June 2011 Avaya/Huawei/ALU
  - Full control plane/ some .1ag – 187 nodes and 412 links

Third Interop Summary (SPBM)

- Based the IETF draft-05 with current IANA code point allocations and the IEEE D4.0 – June 2011.
- Focus on SPBM mode
- Brought together 5 vendors products representing 6 independent implementations.
- 10 real switches
- 1 high end tester
- 1 passive network viewer (via NNI)
- 1 LINUX/Quagga emulator.

Simulate a Large Network

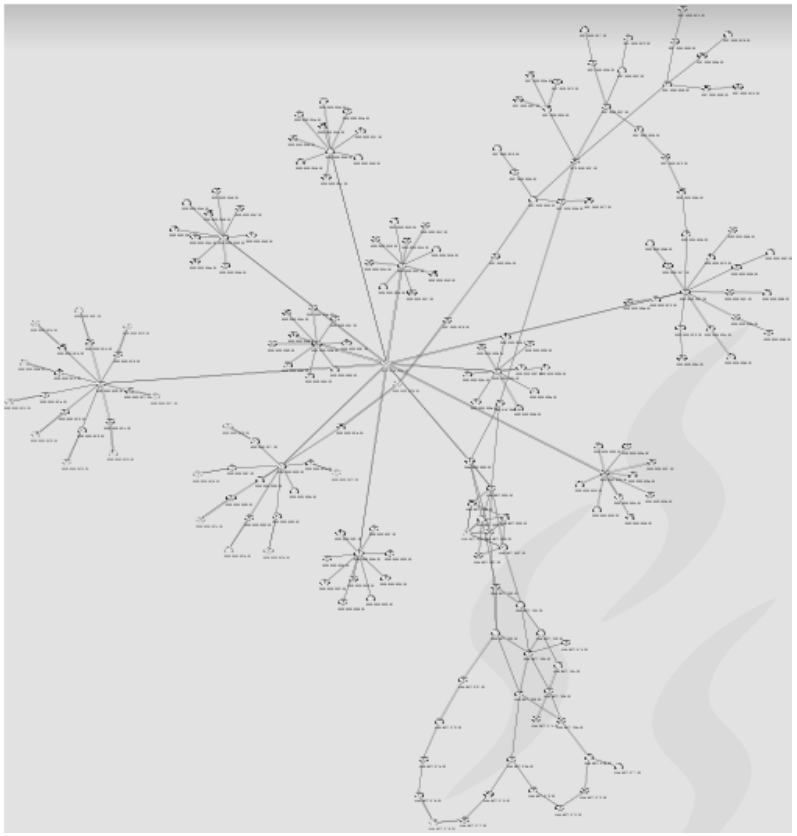

```
<ottawa-s9300-6>d spb status
SPB Status:
    mode SPBM, nodes 187, links 412, adj 6, ufib
372, mfib 45
```

Viewer –via passive NNI - Solana

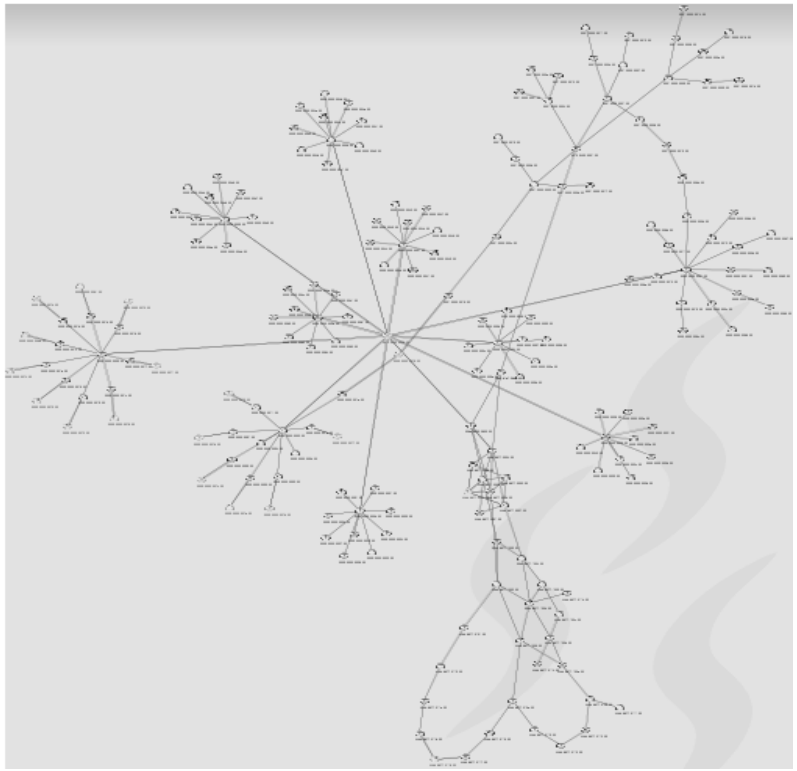

Implemented the hello protocol.

Passive Llsten to the IS-IS LSP's Draw the topology they see.

Automatic network drawing layout.

Network emulator - Spirent

Spirent SPB

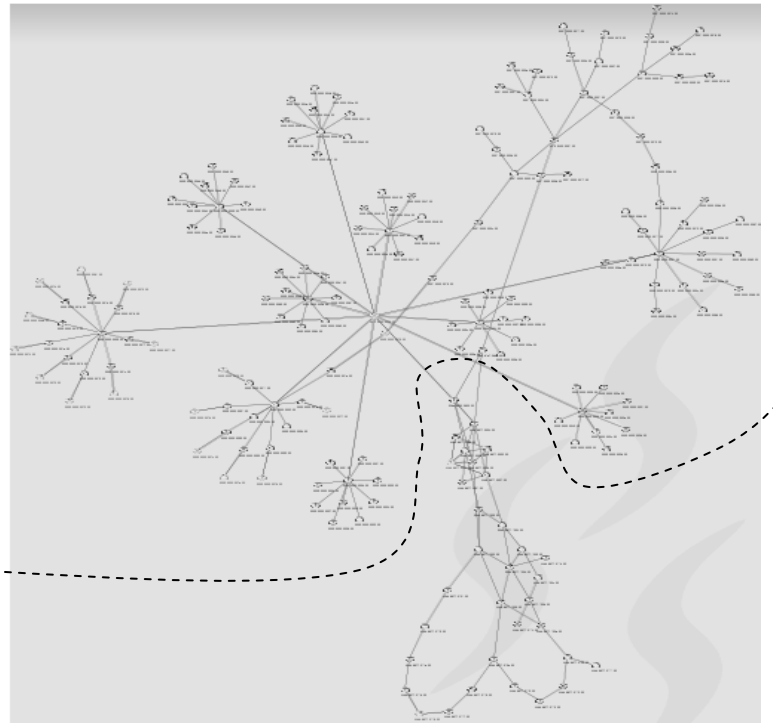

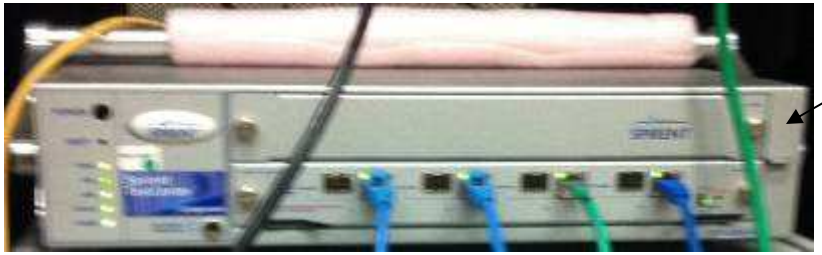

Formed 3 SPBM Adjacencies.

Inject various topologies including ISIDs info.

Injected packets with PBB encapsulation.

Also used as Customer test point for ISID.

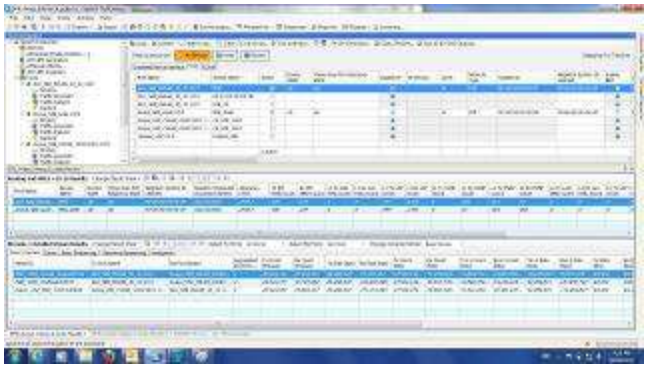

Network emulator - Quagga

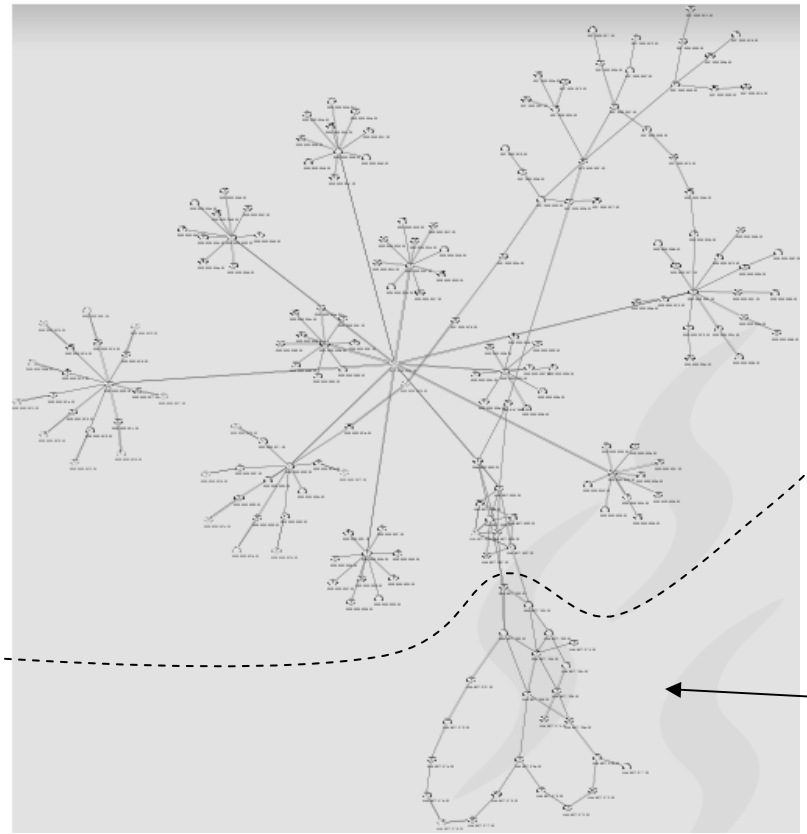

Quagga

Linux/Laptop

Modified Quagga ISIS+Java

Draw picture & inject/change

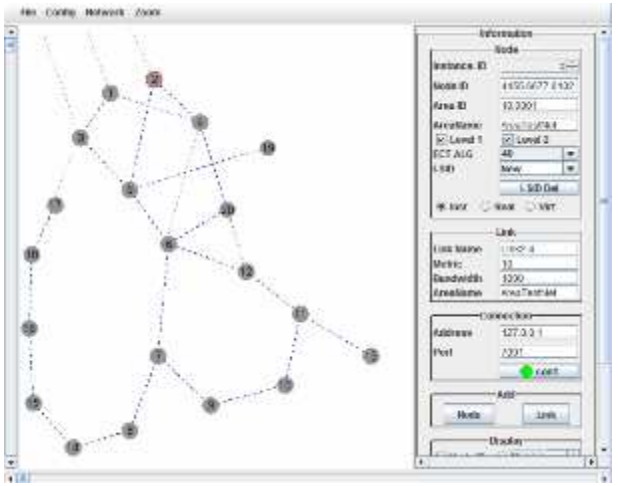

Physical Network

Physical topology

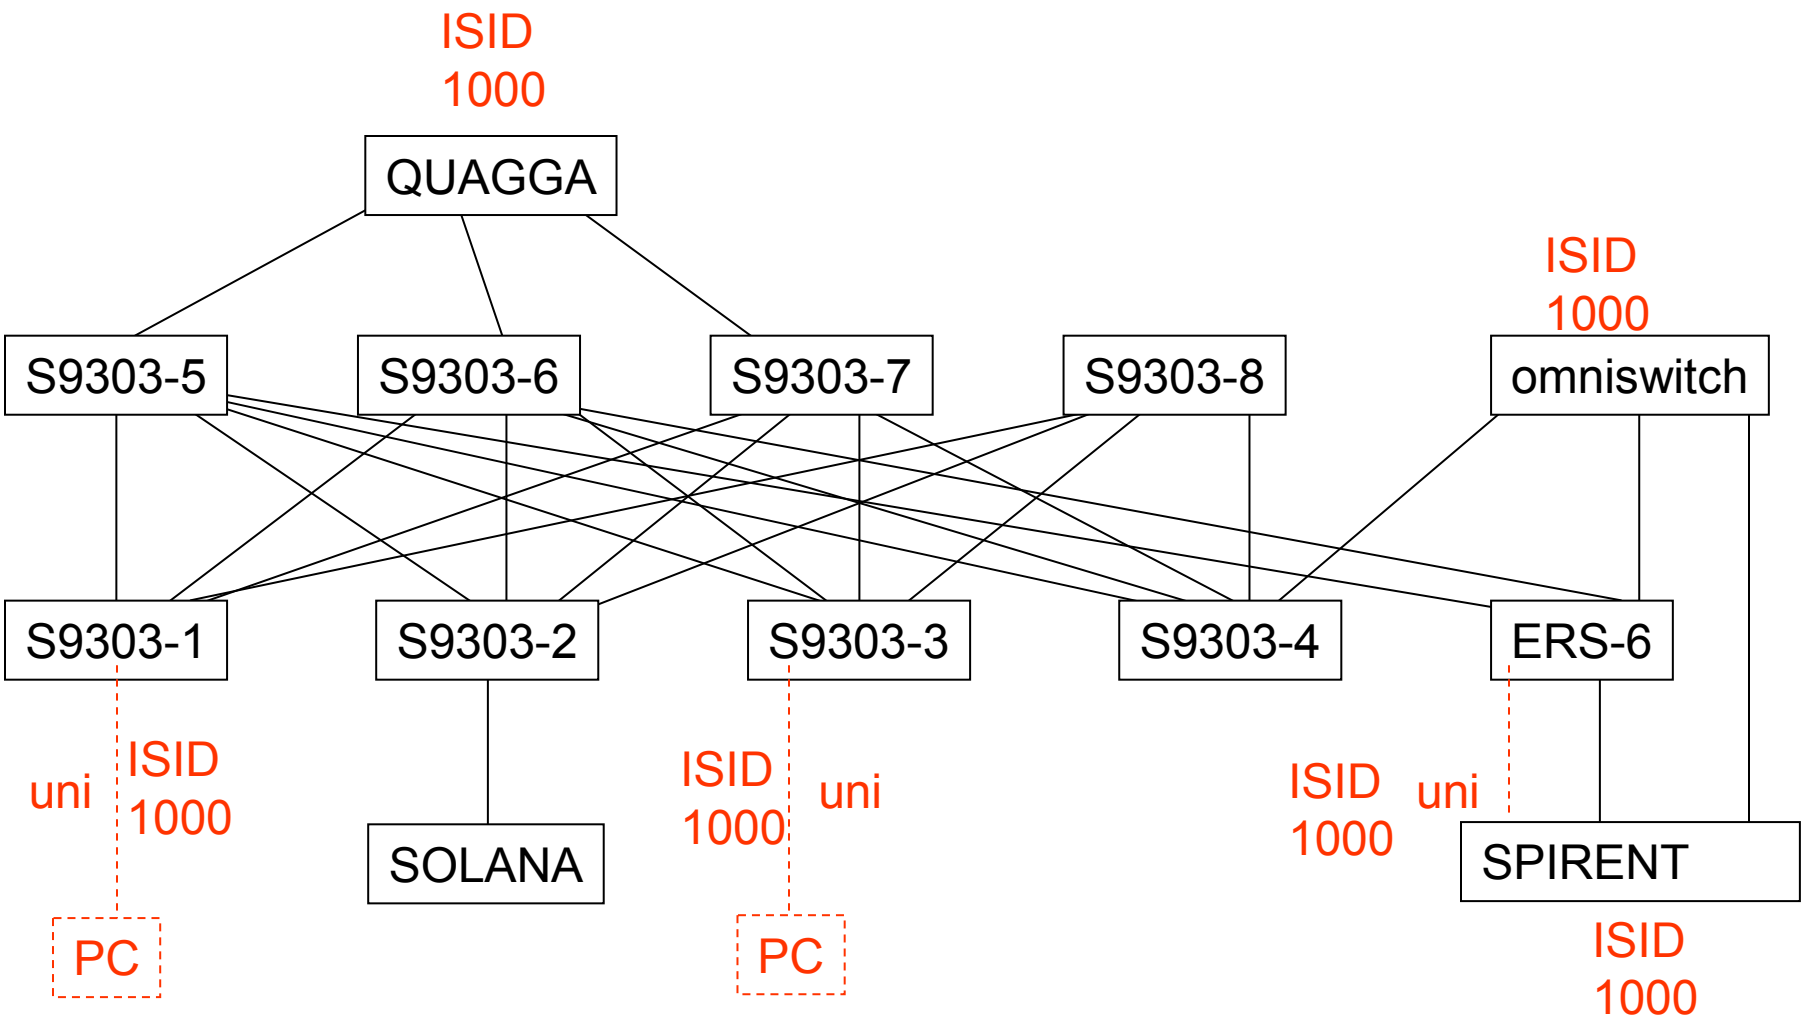

Actual TLV's Tested

PDU	TLV	sub-TLV	Type	Tested	Comments
IIH					
	NLPID		129	Y	RFC 1195
	MT-Port-Cap		143	Y	Tested MTID#0 - single topology
		SPB-MCID	4	Y	Digest of all VID/FID mappings
		SPB-Digest	5	N	Topology Diges - Optional
		SPB-B-VID	6	Y	Part of establishing an adjacency, necessary to a common view of the network.
LSP					
	MT-Capability		144	Y	A TLV level identifying the topology (MTID) for all sub TLVs.
		SPB-Inst	1	Y	SPB Instance, Identifies the bridge uniquely, along with the Base VIDs and ECT Algorithms. Must be carried in fragment ZERO LSP.
		SPB-I-OALG	2	N	Instance Opaque ECT, for new ECT algorithms.
		SPBM-SI	3	Y	SPBM Service Identifier, Maps I-Sids in the context of a B-VLAN to a B-MAC.
	Extended-IS-Reachability/MT-Intermediate-System		22/222	Y	Contains the System ID of a neighbor.
		SPB-Metric	29	Y	Link Metric, necessary for distributing cost of links (if not present for an ISIS adjacency, then adjacency may not carry SPB traffic).
		SPB-A-OALG	30	N	Adjacency Opaque ECT, for new ECT algorithms.

These are the proper IANA code points

```
<ottawa-s9300-6>d spb status
SPB Status:
    mode SPBM, nodes 187, links 412, adj 6, ufib 372, mfib 45
    Collision:
<ottawa-s9300-6>
<ottawa-s9300-6>d isis peer

                          Peer information for ISIS(1)

  System Id     Interface          Circuit Id       State HoldTime Type     PRI
------------------------------------------------------------------------------
s9303-1         GE1/0/10           0000000002        Up   23s      L1       --
s9303-2         GE1/0/11           0000000004        Up   25s      L1       --
s9303-3         GE1/0/12           0000000002        Up   29s      L1       --
s9303-4         GE1/0/13           0000000003        Up   27s      L1       --
Instance_2      GE1/0/14           005               Up   26s      L1       --
ERS-6           GE1/0/24           0825505536        Up   23s      L1       --

Total Peer(s): 6
```

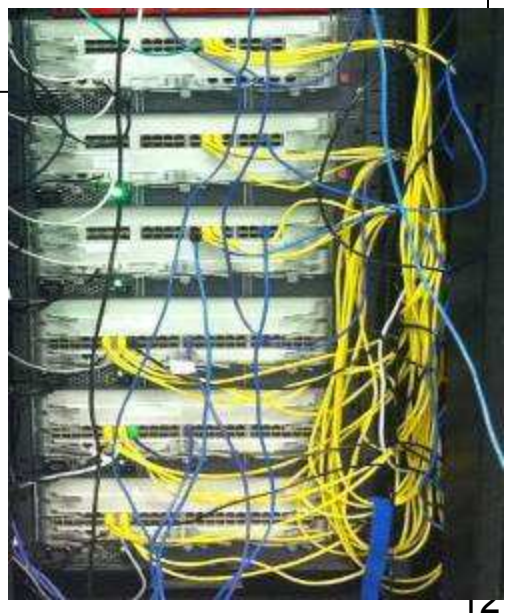

Huawei “S9303”

```
ERS-6:3# show isis adj
********************************************************************************
Command Execution Time: THU JUN 30 11:36:15 2011 UTC
********************************************************************************

================================================================================
                                ISIS Adjacencies
================================================================================
INTERFACE L STATE        UPTIME PRI HOLDTIME SYSID             HOST-NAME
--------------------------------------------------------------------------------
Port2/18  1 UP         18:24:49 127       25 0010.9400.0008    (null)
Port2/19  1 UP         04:27:27 127       29 4455.6677.0005    s9303-5
Port2/20  1 UP         04:25:41 127       28 4455.6677.0006    s9303-6
Port2/4   1 UP         14:39:30 127       21 00e0.b1e7.0bd3    omniswitch

--------------------------------------------------------------------------------
 4 out of 4 interfaces have formed an adjacency
--------------------------------------------------------------------------------
```

Avaya “ERS-6”

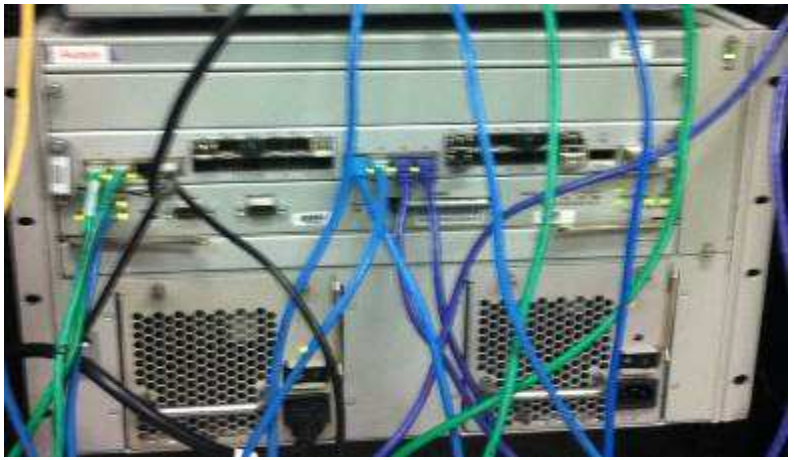

```
6900-X40-A2_14_2_3-> show ip isis adjacency
================================================================================
ISIS Adjacency
================================================================================
System ID           Type      State     Hold      Interface             Hostname
--------------------------------------------------------------------------------
4455.6677.0004      L1        UP        27        1/1                   s9303-4
00be.b000.0600      L1        UP        25        1/2                   ERS-6
0010.9400.0001      L1        UP        28        1/3                   None
--------------------------------------------------------------------------------
Adjacency : 3
================================================================================
```

ALU “omniswitch”

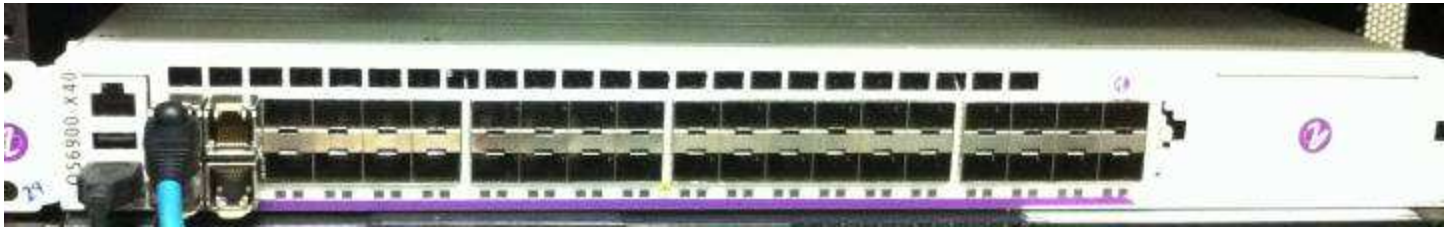

Link State Database

```
<ottawa-s9300-6>d isis lsdb

                        Database information for ISIS(1)
                        --------------------------------

                          Level-1 Link State Database

LSPID                 Seq Num      Checksum      Holdtime      Length  ATT/P/OL
---------------------------------------------------------------------------
0001.0000.0001.00-00  0x00000066   0xd16         648           202     0/0/0
...
Instance_2.00-00      0x000000b6   0xa06e        850           209     0/0/0
Instance_3.00-00      0x000000b6   0x85c4        313           269     0/0/0
Instance_4.00-00      0x000000b4   0xcd4f        811           239     0/0/0
Instance_5.00-00      0x000000b6   0xe9e9        1002          275     0/0/0
ERS-6.00-00           0x000001e9   0x28d5        694           324     0/0/0
omniswitch.00-00      0x00000085   0xb9bf        839           176     0/0/0
s9303-1.00-00         0x00000143   0x8961        717           162     0/0/0
s9303-2.00-00         0x0000036f   0x1714        715           219     0/0/0
s9303-3.00-00         0x00000311   0xb5ba        727           200     0/0/0
s9303-4.00-00         0x0000035a   0xbbba        676           219     0/0/0
s9303-5.00-00         0x000003d0   0x7cde        432           183     0/0/0
s9303-6.00-00*        0x000003ce   0x81d8        503           202     0/0/0
s9303-7.00-00         0x000003d3   0xff18        889           183     0/0/0
s9303-8.00-00         0x0000016d   0x15cc        1015          145     0/0/0
```

Spirent

Quagga

Avaya

Alu

Huaweii

Solana is passive so no LSP

Filtering Database

```
ERS-6:3/show/isis/spbm# unicast-fib
********************************************************************************
Command Execution Time: THU JUN 30 12:04:29 2011 UTC
********************************************************************************

================================================================================
                          SPBM UNICAST FIB ENTRY INFO
================================================================================
DESTINATION        BVLAN  SYSID           HOST-NAME         OUTGOING     COST
ADDRESS                                                     INTERFACE
--------------------------------------------------------------------------------
00:01:00:00:00:03  40     0001.0000.0003  NULL              2/18         70
..
00:55:66:77:10:05  41     0055.6677.1005  Instance_5        2/20         40
00:e0:b1:e7:0b:d3  40     00e0.b1e7.0bd3  omniswitch        2/4          100
00:e0:b1:e7:0b:d3  41     00e0.b1e7.0bd3  omniswitch        2/4          100
44:55:66:77:00:01  40     4455.6677.0001  s9303-1           2/20         20

44:55:66:77:00:01  41     4455.6677.0001  s9303-1           2/20         20
44:55:66:77:00:02  40     4455.6677.0002  s9303-2           2/20         20
44:55:66:77:00:02  41     4455.6677.0002  s9303-2           2/20         20
44:55:66:77:00:03  40     4455.6677.0003  s9303-3           2/20         20
44:55:66:77:00:03  41     4455.6677.0003  s9303-3           2/20         20
```

Shortest Path Next hop

```
6900-X40-A2_14_2_3-> show ip isis spf
=================================================================
ISIS Path Table
=================================================================
Node                          Interface          Nexthop
-----------------------------------------------------------------
0001.0000.0001.00             1/3                0010.9400.0001
0001.0000.0002.00             1/3                0010.9400.0001
0001.0000.0003.00             1/3                0010.9400.0001
....
4455.6677.011a.00             1/2                00be.b000.0600
4455.6677.011b.00             1/2                00be.b000.0600
4455.6677.1008.00             1/2                00be.b000.0600
4455.6677.1009.00             1/1                4455.6677.0004
4455.6677.100a.00             1/2                00be.b000.0600
4455.6677.100b.00             1/2                00be.b000.0600
4455.6677.100c.00             1/1                4455.6677.0004
-----------------------------------------------------------------
SPF count: 185
=================================================================
```

Route query for two ECT-ALGORITHMS.

```
<ottawa-s9300-6>d spb uroute s9303-1 omniswitch vlan 40
PATH:
  s9303-1 -> s9303-6 -> ERS-6 -> omniswitch

<ottawa-s9300-6>d spb uroute s9303-1 omniswitch vlan 41
PATH:
  s9303-1 -> s9303-7 -> s9303-4 -> omniswitch

<ottawa-s9300-6>
```

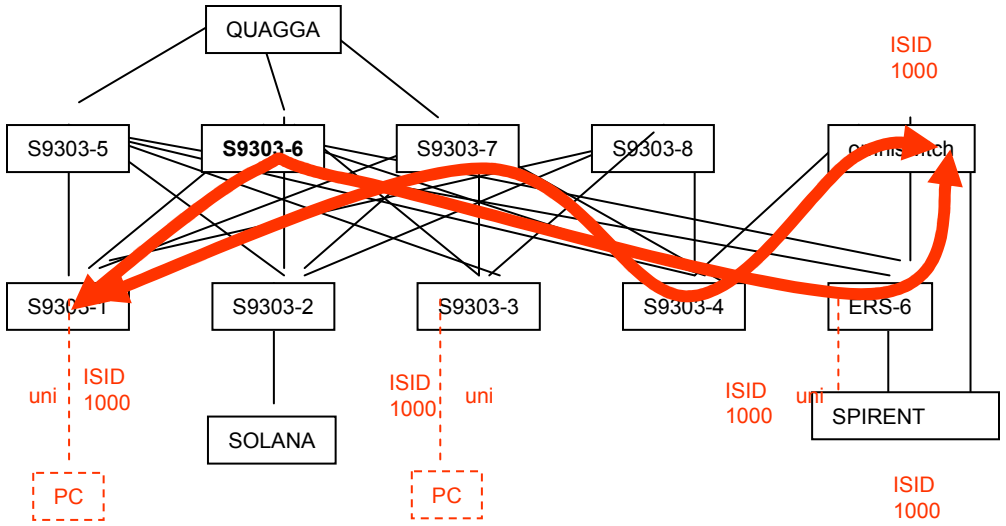

Route query for distant nodes

```
<ottawa-s9300-6>d spb uroute s9303-1 Instance_22 vlan 40
PATH:
  s9303-1 -> s9303-6 -> Instance_2 -> Instance_3 -> Instance_8 -> Instance_14 ->
Instance_24 -> Instance_23 -> Instance_22

<ottawa-s9300-6>d spb uroute s9303-1 Instance_22 vlan 41
PATH:
  s9303-1 -> s9303-7 -> Instance_4 -> Instance_5 -> Instance_8 -> Instance_14 ->
Instance_24 -> Instance_23 -> Instance_22
```

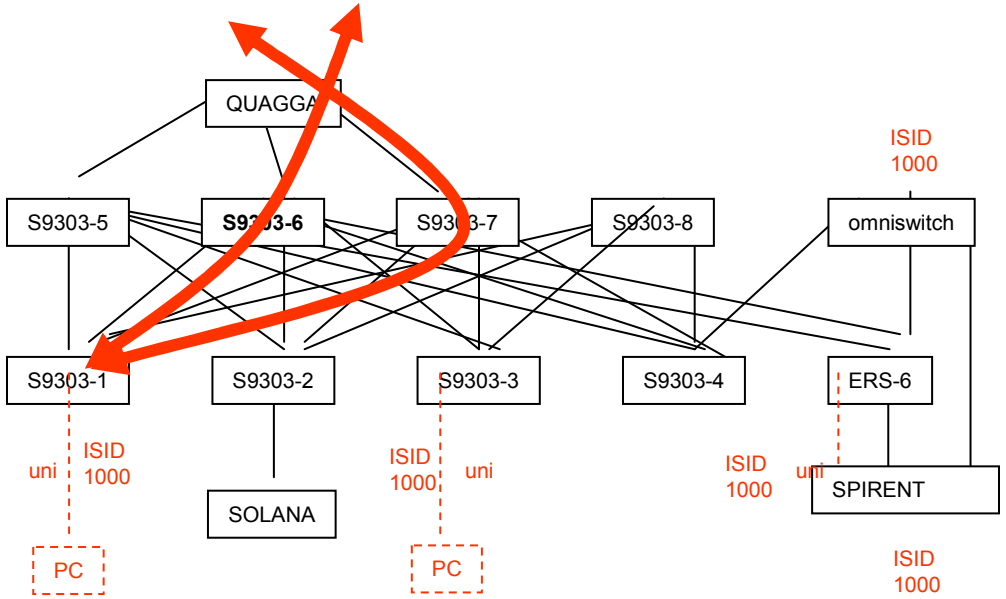

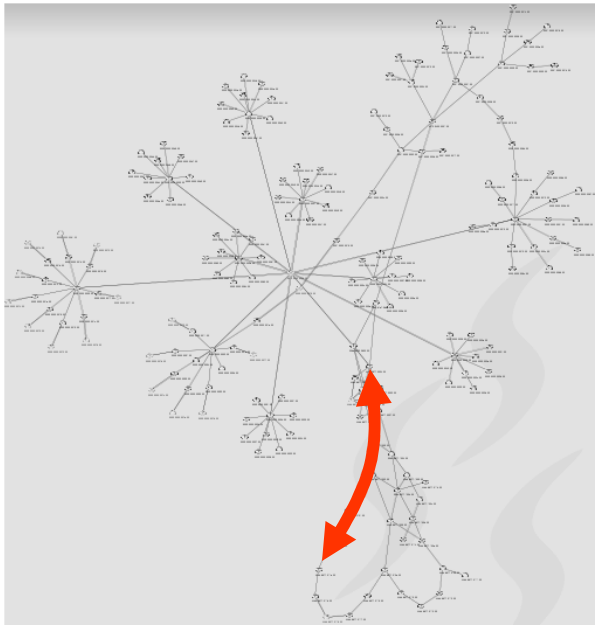

Multicast FDB

```
ERS-6:3/show/isis/spbm# multicast-fib
************************************************************************
Command Execution Time: THU JUN 30 12:08:33 2011 UTC
************************************************************************

========================================================================
                        SPBM MULTICAST FIB ENTRY INFO
========================================================================
MCAST DA           ISID   BVLAN  SYSID               HOST-NAME         OUTGOING-
INTERFACES
------------------------------------------------------------------------
23:00:0e:00:03:e8  1000   40     0001.0000.000e      NULL              2/20,2/2
...
03:00:06:00:03:e8  1000   40     00be.b000.0600      ERS-6             2/18,2/20,2/4,2/2
03:00:06:00:03:e9  1001   41     00be.b000.0600      ERS-6             2/18,2/20,2/4,2/2
73:0b:d3:00:03:e8  1000   40     00e0.b1e7.0bd3      omniswitch        2/18,2/20,2/2
73:0b:d3:00:03:e9  1001   41     00e0.b1e7.0bd3      omniswitch        2/18,2/2
73:00:01:00:03:e8  1000   40     4455.6677.0001      s9303-1           2/18,2/4,2/2
73:00:01:00:03:e9  1001   41     4455.6677.0001      s9303-1           2/18,2/2
```

Compliance points

- Bridge Priority
- Test > 2 ECT (but should be easy)
- Do more adjacency rejection tests
- MIBs?
- More .1ag tests

Recommended but not required for compliance

- Be nice to have Default B-VIDs/ECT mappings. I.e Best Common Practice eg: 4000-4015 as default set?
- Align SpSouceID's, MEPID's.
- Hostname very useful, if not set by user suggest (ASCI of your OUI) + instance.

Next Steps

- Tata has offered to host 4th Interop in NJ and offered use of a transatlantic link.
- More in depth.1ag tests
- IETF interop informational but want to use also to recommend good practices (BCP).
- Perhaps test the MIBs?

Thank-you

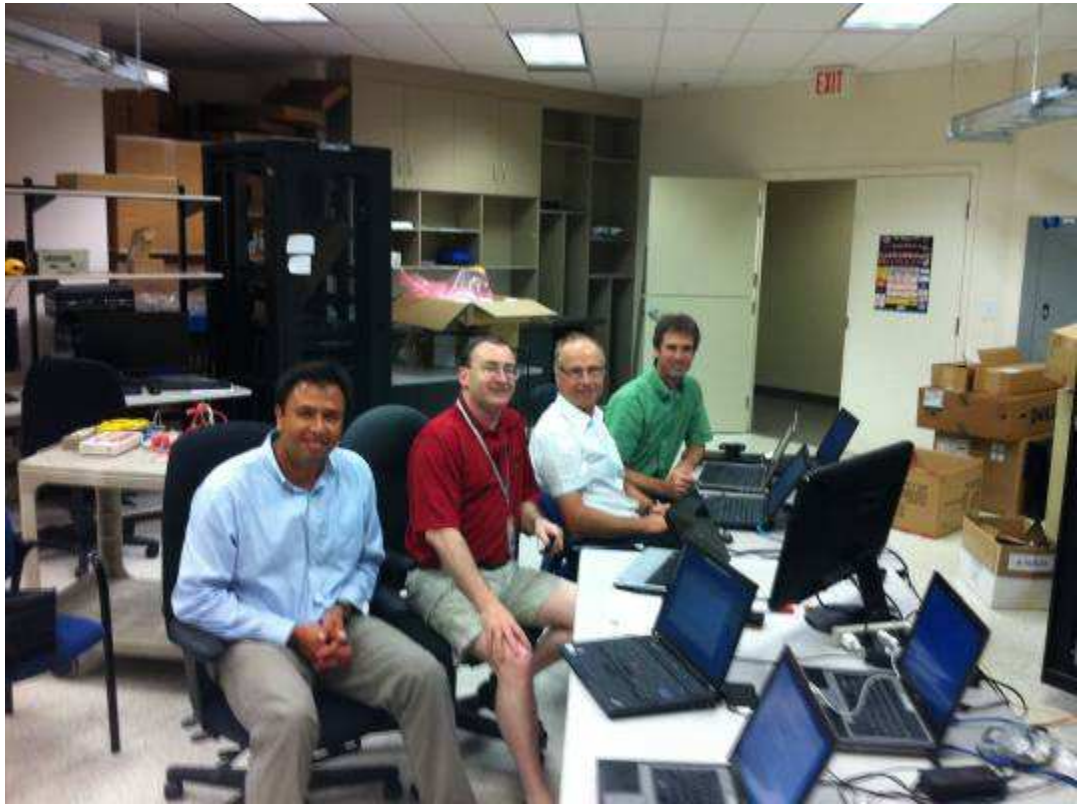

Edgard, Peter, John, Stephane, Bis & Rup (not shown)...
and many others of course in 3 different countries.