

INVITATION TO NEGOTIATE ITN # FLRNETUP2021

Questions & Amendments

June 9, 2021

- 1. Q: Referencing Question 1 from your answers on May 13, 2021, when you reference a small and large site, are you considering a small site to consist of one core router and one hybrid switch, and a large site to consist of two core routers and two hybrid switches? If this is correct, would this design lead to a requirement of a large site DEVICE (or combination of a core and hybrid device) to provide 40x1/10G, 20x100G, and 6x400G edge connections?**

That is correct. Subtending switch option is provided to allow for 10G and lower services. 100G or higher should be provided by the core router. The combination of all devices determines the available port count at a site.
- 2. Q: Referencing Q11, can you elaborate on the ability to move the device forward several inches?**

We can obtain brackets that will move the equipment forward approximately 3.5 inches. This should also allow our fiber management to support jumpers plugged into this gear. Any further forward and the fiber management would be out of position to support jumpers to the equipment.
- 3. Q: Referencing Q21, can you elaborate how you're providing PWHE with two boxes today?**

Currently we do not have the correct linecards to support "true" PHWE, so we terminate the MPLS PW on a subtending box which exposes vlans to the core router on a bundle interface. The subtending box receives MPLS traffic from the core. This traffic flow is obviously sub-optimal for high bandwidth flows.
- 4. Q: Referencing Q26, is the maximum power draw of 80A for a large core site inclusive of both core devices and both hybrid devices? Or is that the max power draw per device or the max power draw per feed?**

That is the total draw for all devices.
- 5. Q: Are the locations for the L2 Hybrid Switch and Core router the same? If they are not, what is the typical distance between these devices?**

They would be in the same rack or in directly adjacent racks.
- 6. Q: Is there a preference between SRv6 and SRv6+/SRm6 and if so, can you elaborate on why?**

There is a potential to support SRv6 on platforms which do not support MPLS. These platforms are already deployed and will not be replaced as part of this upgrade. The decision to use SRv6 hinges on SRv6 and SR/MPLS internetworking testing that would be conducted at a later time. This is why SRv6 is noted as a "valuable" feature. Any use of SRv6+ would require demonstrated interoperability with SRv6 deployments.
- 7. Q: What is the use case for the Containers or VM plug-in module support?**

Potential for NFV, most likely around security (IDS, etc). Performance monitoring will be done "off-box" using subtending systems. No specific use cases have been committed.
- 8. Q: For Deep Buffering support, is there a target or goal for the buffer size?**

An approximate target would be 16GBytes+ per 24x400G, but that is by no means an absolute rule.

May 24, 2021

1. **Q: Is FLR expecting implementation services to be provided by the responder as part of the ITN?**

FLR is not expecting implementation services by the respondents as part of this ITN.

2. **The ITN has a brief mention around monitoring and in depth element management, what is FLR currently monitoring and what would FLR like to monitor going forward (i.e. up/down, performance monitoring etc..)?**

FLR currently monitors a wide variety of items from FLR's backbone routers:

- a. Interface up/down
- b. Protocol up/down
- c. Protocol send/received (threshold)
- d. Interface send/received/errors/pps (usage trending/threshold)
- e. Environmentals (temperature/power usage/voltage/etc.)
- f. Linecard status
- g. Power supply status
- h. Chassis alarm status
- i. Backplane status (usage/errors)

FLR also performs active measurement through the chassis using Perfsonar. Perfsonar will be expanded as part of this ITN, but it is not a part of the ITN.

FLR would continue to monitor roughly the same things, but each platform is going to have elements that are unique to it. FLR desires vendor assistance to make sure we are monitoring those unique elements important to the health of the equipment. FLR also finds it challenging to monitor BGP status "in VRF" on some platforms using standard vendor monitoring tools.

3. **Section 2.8 states training specific to FLR engineering requirements. What are the FLR requirements around training?**

Engineering training should be specific to the hardware platform proposed, and the NOS. If the NOS is IOS-XR, this training should cover the differences between 6.x 32bit and the version of IOS-XR that is running on this platform. If the NOS is not IOS-XR, it should cover the NOS in a service provider environment running the protocols that we have outlined in the ITN. A brief overview of management and monitoring strategies using industry standard protocols as well as next generation monitoring (i.e. streaming telemetry) would be useful.

4. **By the way we read this question, we view this as a question about testing after the ITN is complete and an order has been placed. Is FLR expecting the respondent to provide a suggested equipment testing schedule, if not could FLR clarify what more information we can provide to the testing schedule.**

Yes, FLR is expecting the respondents to provide an equipment testing schedule.

May 13, 2021

1. **Q: Will FLR provide existing port density and speed requirements per location?**
Small sites should support at least 40 1/10G edge connections, 24x100G edge connections and 6400G edge connections. The solution should be configured such that subtending switches are connected to core routers by at least 2x100G and still have the necessary edge ports. Large sites should support at least 80 1/10G edge connections, 40x100G edge connections and 12 400G edge connections. The solution should be configured such that subtending switches are connected to core routers by at least 2x100G and still have the necessary edge ports. Both solutions should have a total of 6 400G ports available for core interconnections (i.e., 3 east and 3 west). For large sites, core nodes should have enough 400G ports to locally connect to each other at 3x400G. Additionally, 400G ports should support 200G
2. **Q: Will FLR provide anticipated/expected growth on port density and speed requirements per location?**
For pizza box style configurations, the growth is included in the port counts listed above. For chassis-based solutions, the solution should be able to handle the port counts listed above, and be priced in that configuration, however FLR may elect to defer the purchase of some linecards until they are necessary.
3. **Q: What are the current number of prefixes per VRF?**
Internet – 852982 v4, 117463 v6
R&E Net – 22653 v4, 2635 v6
CP – 288847 v4, 34197 v6
4. **Q: What is the anticipated lifespan of this solution?**
The solution as designed, should sustain the network with no significant additional hardware for 4 years. We would expect to add hardware to the solution in the 5–7-year timeframe as port needs demand. Hardware would be expected to have a lifetime of up to 10 years being used somewhere in the network.
5. **Q: What prefix growth is expected per VRF over the lifespan of the solution?**
Assuming a 7-year growth plan, we would expect to see 1.2M to 1.4M IPv4 prefixes in the Internet table, approximately 40K IPv4 prefixes in the R&E table, and around 700K to 800K IPv4 prefixes in the CP table. For IPv6, we see a more aggressive growth pattern of the global v6 table. We expect to see as many as 600K IPv6 prefixes on the Internet side, 150K IPv6 prefixes on the CP side and 10k prefixes on the R&E table. Due to the size predictions on the CP side, FLR may be forced to limit the CP table to higher “value” peers on a bandwidth/route basis.
6. **Q: What is the anticipated timeline for implementation for this solution?**
We would like to have the network in place by EOY 2021 or worst case, Q1 CY22.
7. **Q: With which routing protocol do the CPE service nodes connect to the Flex nodes?**
OSPF for EoMPLS services/mgmt and BGP for user traffic.
8. **Q: Based off the comment “We are interested exploring both site level routed redundancy (multiple core routers per site) as well as backhaul redundancy (layer 1 or 2 backhaul to an adjacent core node to provider layer 3 redundancy)”, is FLR expecting the solution to have two options? One with site level redundancy and one with backhaul redundancy?**
Yes.
9. **Q: Does FLR already have the server class hardware for the route reflectors specified in section 2.0, or are the servers expected to be a part of the solution?**
Vendors are free to spec this hardware as part of their proposal, but it should be separable from the network equipment.
10. **Q: How many routes do the Route Reflectors need to be able to scale to?**
We expect 4-5 of the core routers will accept full tables from upstream Tier 1 providers. We also expect 2-3 of the core routers to receive nearly full FLRnetCP tables (currently nearly 300K prefixes). Our desire

is to support optimal route reflection (based on maturity) in the future, which should not significantly increase the memory footprint of the RR, but we may wish to deploy diverse-path/shadow route reflection, add-path, or multi-rd, which do have scale implications.

11. **Q: In this statement “Cost effectiveness and compatibility of the proposal with the colocation facilities that will host this equipment,” what compatibility criteria will be used in the evaluation process?**

Primarily we are concerned about rack space, device dept, and power consumption. Our colo facilities use single frame telco-style racks and many locations are not extremely deep. Depth of more than about 29” may require a customized rack solution that gets more difficult the heavier the device is. Mid mounting is probably not a realistic solution, but the device can be moved forward several inches to make some room.

12. **Q: Is FLR looking for a service where the vendor provides on-site support with parts replacement, or is FLR just looking for vendor part replacement with FLR doing the hardware swap?**

In all cases, escort would be required by smarthands to reach the equipment anyway, so a part-only solution is fine.

13. **Q: Is FLR planning on managing the pool of spares and expecting to get them to the site of failure?**

Yes, when spares are called for rather than a direct swap from a vendor parts depot, FLR would manage that process.

14. **Q: Can you provide additional details on the make up of the ITN evaluation team.**

The committee will be made of up Executive, Operational, Financial, and Architectural/Engineering leadership

- a. **Q: Can you provide additional information on the scoring/weighting model that will drive the evaluation criteria?**

Due to the fact that this is an ITN, the evaluation process is iterative and largely will follow the criteria in section 2.6

15. **Q: Please confirm that "...involved in..." as used in Section 23.0 on page 14 does not preclude an OEM from providing technical assistance to its resellers.**

This is confirmed.

16. **Q: What is the Add/drop filter bandwidth? Are these colored 50GHz?**

The current filters are 40 channel, odd, fixed channel panels. The panels themselves are expected to pass a channel wider than 50Ghz, however we have no hard information on the full band-pass of a channel across the worst-case path which would be 3 intermediate WSS/ROADM nodes. Empirical testing may be required to understand the optical performance, which FLR is would support. We can also support initial lab testing in the FLR optical lab using DUT and an optical signal analyzer.

17. **Q: What is the WSS pass band filter shape at 3dB and 0.5 dB bandwidth?**

See the above answer

18. **Q: How many ROADMs are the signals transiting in the network?**

Worst case there would be 3 intermediate ROADM nodes in any 400G path.

19. **Q: Regarding Section 25.0 ADDITIONAL TERMS AND CONDITIONS of the Florida LambdaRail ITN, is there a separate T&Cs document that contains the general and special conditions referenced in this section or is the FLR_ITN_Network_Upgrade_2021-v220210412(3).pdf document the only document provided for this RFP?**

This is not a RFP. It is an ITN. Any T&C's that result from this solicitation would be included in any awarded contract subsequent to negotiation.

20. **Q: Are all day 1 features needed on every device (core, peering, hybrid L2)**

The appropriate features are marked with a C=Core, P=Peering, H=Hybrid L2

- 4-6 million IPv4 prefixes in FIB, 1-2 million IPv6 prefixes in a Forwarding Information Base (FIB) [C/P]

- Fast Routing Information Base (RIB) to FIB/TCAM (Ternary Content Addressable Memory) programming (forwarding plane) [C/P]
- Fast BGP prefix learning (control plane) [C/P]
- MPLS with VRF support (RFC4364), v4 and v6 (6vPE) [C/H]
- EoMPLS/Draft Martini/RFC4447 [C/H]
- MPLS Traffic Engineering with RSVP [C/H]
- Flexible mapping of Vlans into L2VPNs (individual, range, multiple range) [C/P/H]
- Segment Routing over MPLS with SR-TE and TI-LFA [C/H]
- Lossless/Line-rate performance when running with a rich set of features and
- minimal oversubscription on line-card and switch fabric elements [C/P]
- BGP Flowspec in hardware [C/P]
- IPv6 support with “full width” IPv6 ACL support [C/P]
- uRPF with loose and strict mode [C/P]
- Support for traffic ACLs at scale (at least 64K entries) [C/P]
- Hitless ACL TCAM programming [C/P/H]
- Dynamic Deep buffering support [C/P/H]
- Support for 400G including DCO ZR+ [C]
- Egress rate limiting (policing/shaping) on a per-sub interface bases [C/H]
- IPv4 and v6 Multicast (ASM/SSM) in VRF, Draft Rosen [C]
- Multi-core route-processor design with sufficient memory for future growth [C/P]
- Pre-emptive multitasking OS with protected memory [C/P]
- Two stage config to commit with pre-commit syntax and commit-time semantic checks [C/P/H]
- Flexible configuration templating and manipulation from the CLI [C/H]
- Configuration change tracking (userid/time-date/change made) [C/P/H]
- Configuration reversion to point in time [C/P/H]
- Rich set of APIs. Must include at least two of the following [C/P/H]
 - o XML
 - o REST
 - o RESTconf
 - o Netconf
 - o gRPC
- Hardware based Control Plane Policing/Protection [C/P/H]
- Port mirroring without performance impact to traffic (SPAN/SPAN to MPLS PW/ERSPAN). This should include both ACL based SPAN and full interface/sub-interface based SPAN. [C], [P and H can use dedicated SPAN port]
- Full Netflow v9/IPFIX support with sampling with MPLS support [C/P]
- Route leaking across all VRFs and the default table [C]
- Service Provider oriented feature set [C/P/H]
- BGP PIC (protocol independent convergence) or other fast BGP failover [C/P]
- Support for disabling any vendor locking of optics [C/P/H]
- Rich route policy language with configuration re-use as a central element [C/P/H]

21. **Q: What is the use case for the PWHE? Are you delivering BNG services? If so, can you elaborate the details of your use case today and how you want to change these for the future?**

The use case is for MPLS based CPEs to transport Layer 3 services through intermediate MPLS fabrics that

would then terminate on a core router as a virtual interface. The member would have a direct single hop BGP relationship with the core, but the layer 2 transport would provide automatic redundancy to reach the core router. We are not currently able to do this directly, currently we use two different boxes to do this, and we would like to reduce our box count.

22. **Q: How important is 25G support?**

25G support would be seen as a nice extra, but not required. It is currently unclear whether carrier partners will widely deploy 25G such that we could take advantage of it for member backhaul.

23. **Q: What are the densities needed for 1, 10, 25, 100, and 400G subscriber ports?**

See above.

24. **Q: You request DC power for the optical equipment. Is DC the preferred power source for all equipment?**

We would like to see both DC and AC options as there are several sites where AC is cheaper and several sites where DC is cheaper. Once we analyze the proposals and look at the space and power requirements, we will refine our request.

25. **Q: Will the subscribers peer directly with core or L2 Hybrid devices or with regional RRs?**

Subscribers will peer directly with the core, or they may also peer directly with a local CPE device. Subscribers will not peer with regional RRs or the L2 Hybrid devices.

26. **Q: You mention the desire to utilize pizza-box style hardware. What specific requirements or goals do you want to obtain for power and space?**

We are looking to preserve rack space and minimize our power footprint as much as possible. There are also some constraints at some of our locations which limit us to about 29" total depth. Some alternate rack configurations are possible, but we would like to stay under this depth if possible. Our absolute maximum power draw in a large core site would be around 80A @ 48vDC, but we hope to be significantly under that. Our absolute maximum rack space footprint for switch/router equipment would be around 32U (we have separate space for optical equipment), however we would, again like to be under that, which greatly assists our migration to the new platform and may allow us to reduce our footprint in the colo which would save on OpEx. TCO including space and power is an important factor in the decision of which equipment to move forward with.

27. **Q: Is there interest in a SDN-based controller system for provisioning and service assurance?**

As a value-add, yes.

28. **Q: What features of traffic engineering do you expect to implement today or in the future?**

Today we do simple TE/FRR tunnels for our layer 2 services using MPLS/TE and RSVP. Bandwidth management is not a primary concern. Rather it's just providing fast failover in the case of a topology change and eliminate shared fate for member services that are redundant on each side of our transport. We are interested in moving our Layer 2 services over to SR/MPLS. Due to the size of our network, we feel that we can perform "segment management" using internal tools, however solutions for a PCE are welcome (see above).

29. **Q: Do you need TE features such as dynamic bandwidth management?**

We don't anticipate the need for such a feature set.

30. **Q: You've requested the rate or scale of Dynamic Deep buffering support. What measurements are you requesting?**

Primarily maximum buffer depth (in xBytes or timescale of queueing at line rate), however we would also like to understand the processes used in queue management, how egress queueing is managed and how the system (i.e., the series of routers/switches) might perform in "downshifting" traffic from a 400G core to a 1G edge given the tendency for "flow sync" across high-speed backbones.

31. **Q: Can you elaborate on the full requirements to meet a Service Provider oriented feature set?**
Many of the specific requirements have been enumerated in this ITN. Other than that, we like to make sure that the development happening on the platform is targeted at problems service providers are likely to see, both now and in the future, and not so much enterprise networks. Active development on new and emerging standards for service provider networks is of high importance to FLR.
32. **Q: What is your migration strategy from MPLS to SR?**
So just to be clear, we have no intention of moving away from MPLS per se. We are interested in moving to SR/MPLS as a layer 2 transport technology. If we do elect to do that, we would first like to fully implement the new network using existing "legacy" MPLS technology based on Martini draft circuits. Once all is up and running, we would then migrate tunnels on a tunnel-by-tunnel basis over to SR.
33. **Q: The RFP states requirements for a legacy LDP environment and a segment routing environment, can you explain where one system would be used vs the other?**
The legacy LDP environment would be used to interoperate with legacy CPE gear that does not support SR, until its retirement. SR would be used on all new deployments and on legacy gear which supports SR.
34. **Q: Are the route scale requirements the same for core router, subtending switch, peering router, and route reflector devices? If not, what are the FIB/RIB requirements for each class of device?**
See above for most of the detail. In short, Core and Peering router will have the same route scale requirements, subtending switch would have a very low scale as it would only be using SR/MPLS or LDP/RSVP Martini tunnels, so mostly IGP only routes. RR would need high scale because of the possible need for add-path and the multiple Tier 1 full feeds we would be accepting.
35. **Q: What IGP's are to be used within the core?**
Currently the core runs OSPFv2 and OSPFv3. We anticipate a migration to ISIS, however there is testing to be done with some of our legacy platforms that are not being upgraded before the final decision is made.
36. **Q: What is the transport for the OOB management network between the core nodes?**
The OBM network is a dedicated 100Mbps optical ring which rides over our DWDM platform. Each OBM node is routed and provides a routed network to the local equipment. We anticipate that we will upgrade the core-to-core OBM links to 10G as a function of this upgrade, or shortly after.
37. **Q: What are the class of service requirements?**
Currently FLR does not provide QoS. If we decided to do so, our CoS requirements would be extremely simple. The most likely deployment model would be "scavenger" services for either bulk data, or "DDoS avoidance"
38. **Q: Can a full set of port requirements be provided for each site?**
That has been covered in questions above.
39. **Q: Will technically equivalent capabilities be considered in lieu of the Cisco specific requirements?**
We don't believe there are many Cisco specific requirements in the ITN. Any Cisco specific ones are there to facilitate migration from the old network to the new network, as we cannot do a "flag data" cutover, or to support legacy equipment which will not be swapped during this upgrade. All other technologies should be based on open standards and backed by an international standards body. As we progress through the ITN, we can get into specifics about the capabilities deployed and how we interoperate with existing hardware, and perform a migration to new hardware.
40. **Q: Will there be an opportunity to ask follow-up questions that will arise from the answers to the above questions?** Yes
41. **Q: Can you please elaborate on your MACSec requirements? What is your target encryption level? Where in the network do you want to implement MACSec?**
We would be targeting GCM-AES-128. Primary application would be to encrypt traffic over third party



Florida's 100 Gigabit Research and Education Network

(carrier) circuits. We may elect to also run it over our in-house DWDM system but that has not been decided.

April 30, 2021

1. **Q: From Section 2.4, Page 8 Section 2.4, Page 8, First paragraph- "Currently FLR maintains a Cisco NCS2006 based platform. This is configured as a 40 channel 50GHz ITU grid system using odd channel spacing." What is the model number of the existing ROADMs within the NCS2006 platform?**
We use SMR1 and SMR2 roadms: 15454-40-SMR1-C, 15454-40-SMR2-C

2. **Q: Section 2.4, Page 8, Option 3 - What is the model number of the existing ROADMs within the NCS2006 platform?**
The expected handoff for a 400G solution (option 1/2) is 400G, the expected handoff to the router for option 3 would be 200G.