

# *y*Our Customers are deploying LISP

See how it impacts you as a Service Provider

Paul Vinciguerra

NANOG55 • June 2012

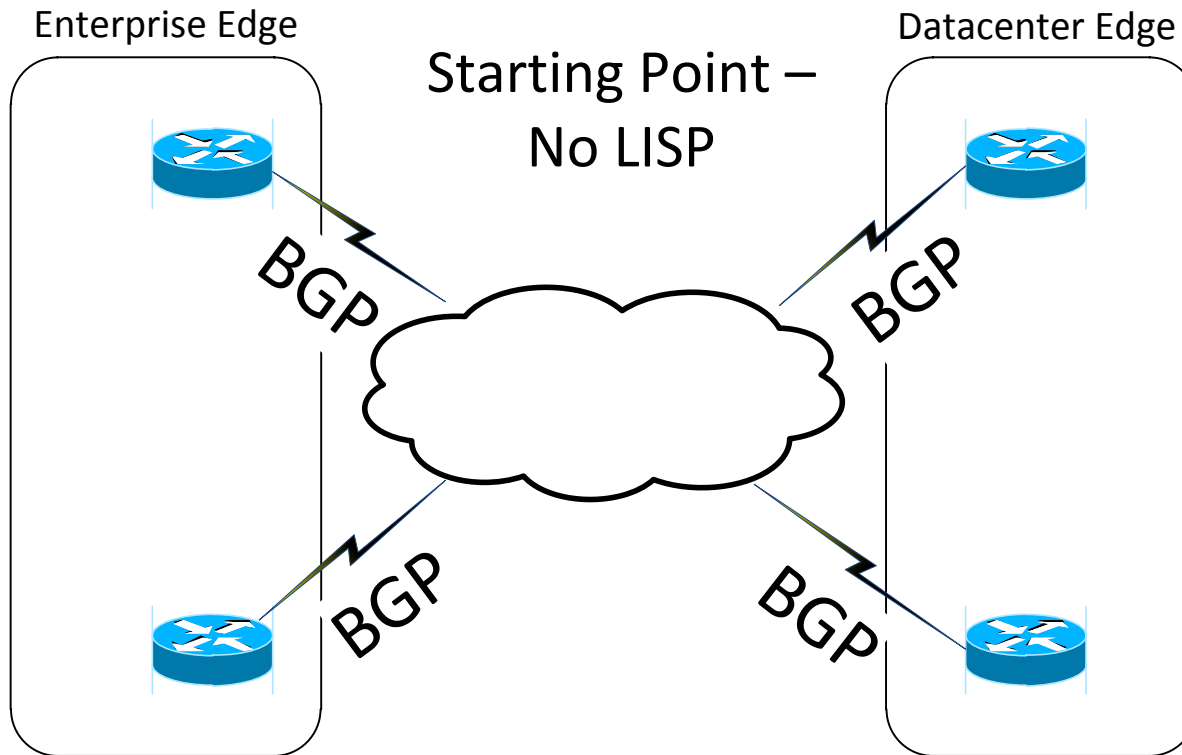
# Why would we deploy all this?

## Our customers are moving to the “cloud”

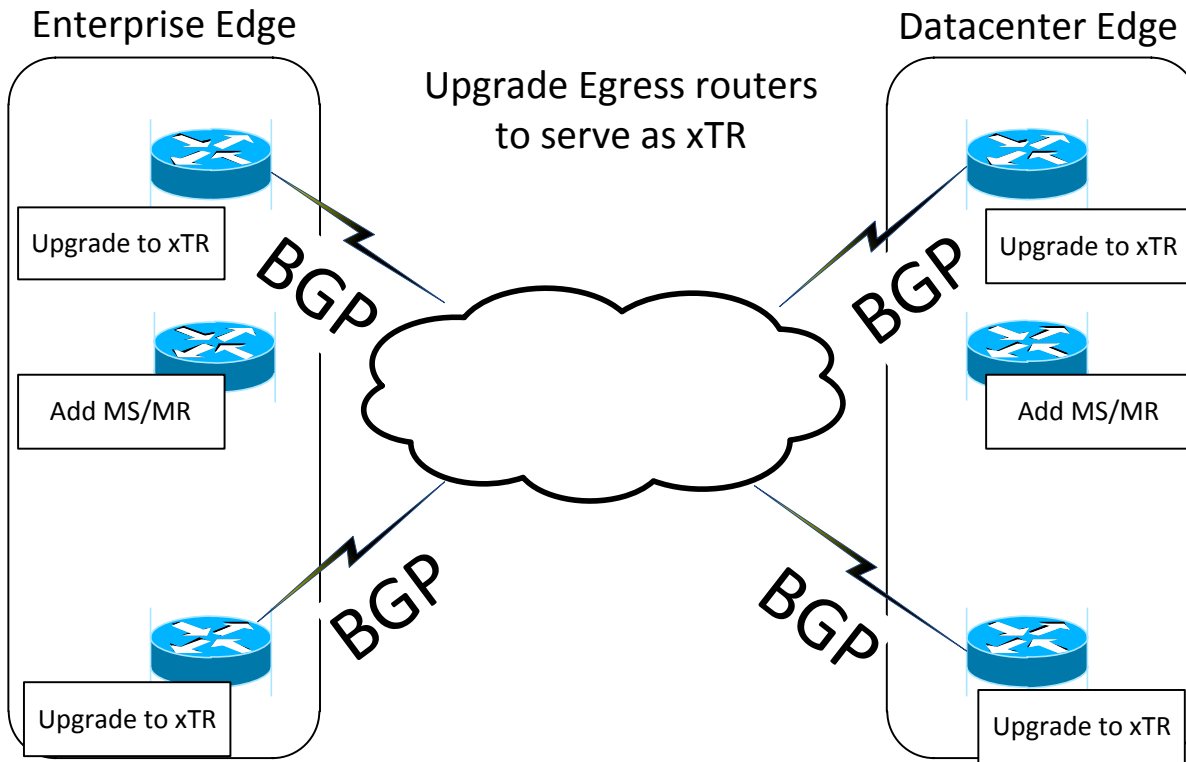
- Want high availability.
- Extra bandwidth/availability for accessing the “cloud.”
- Expecting benefits beyond BGP.



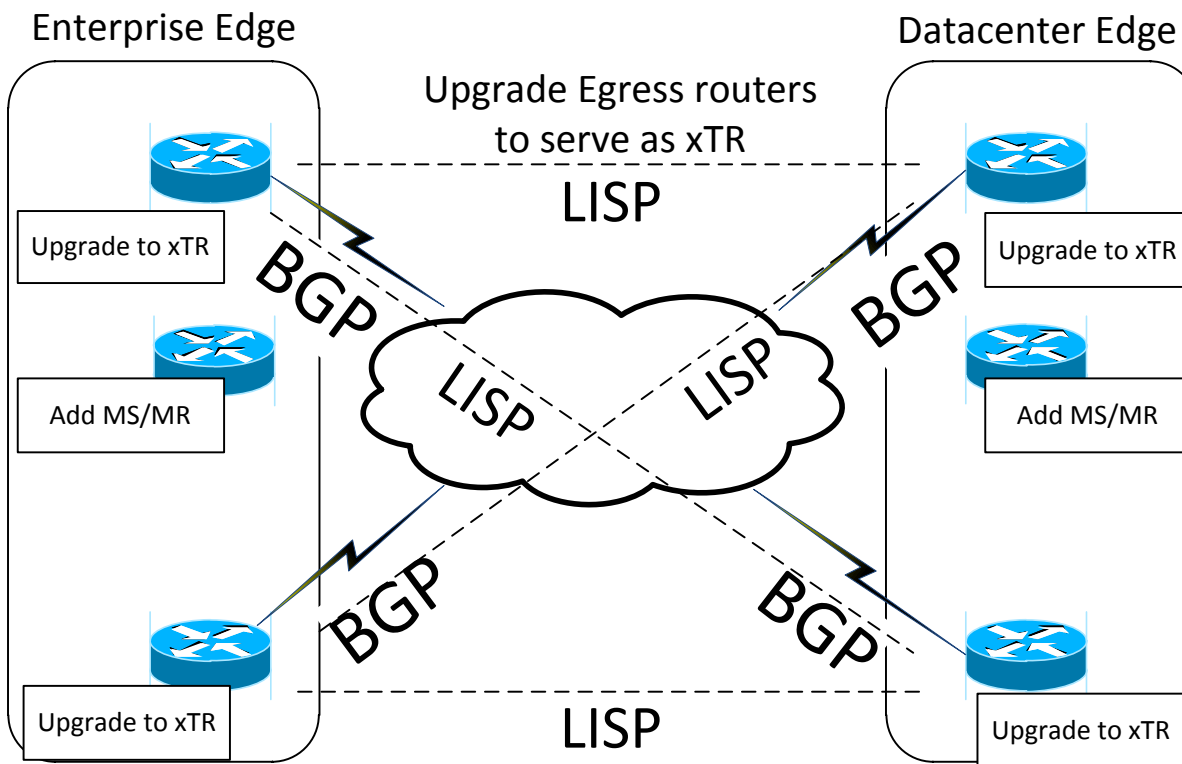
# Customer Self-Reliant Deployment



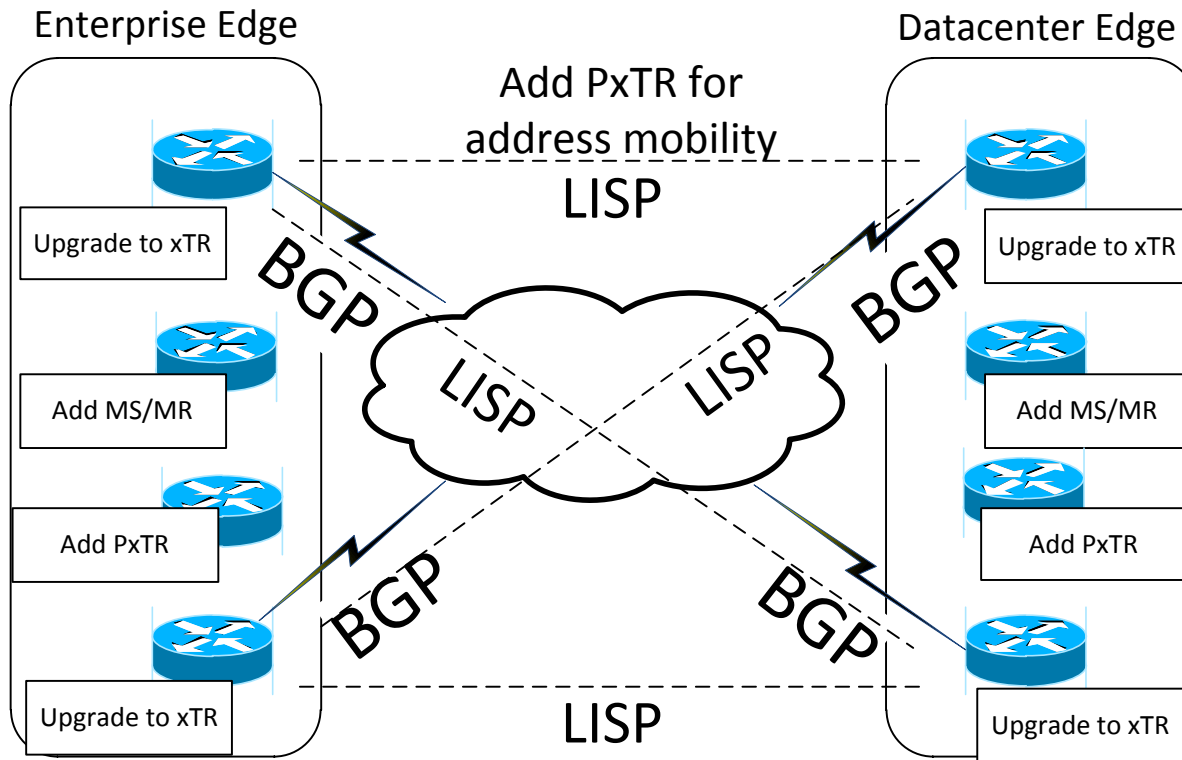
# Customer Self-Reliant Deployment



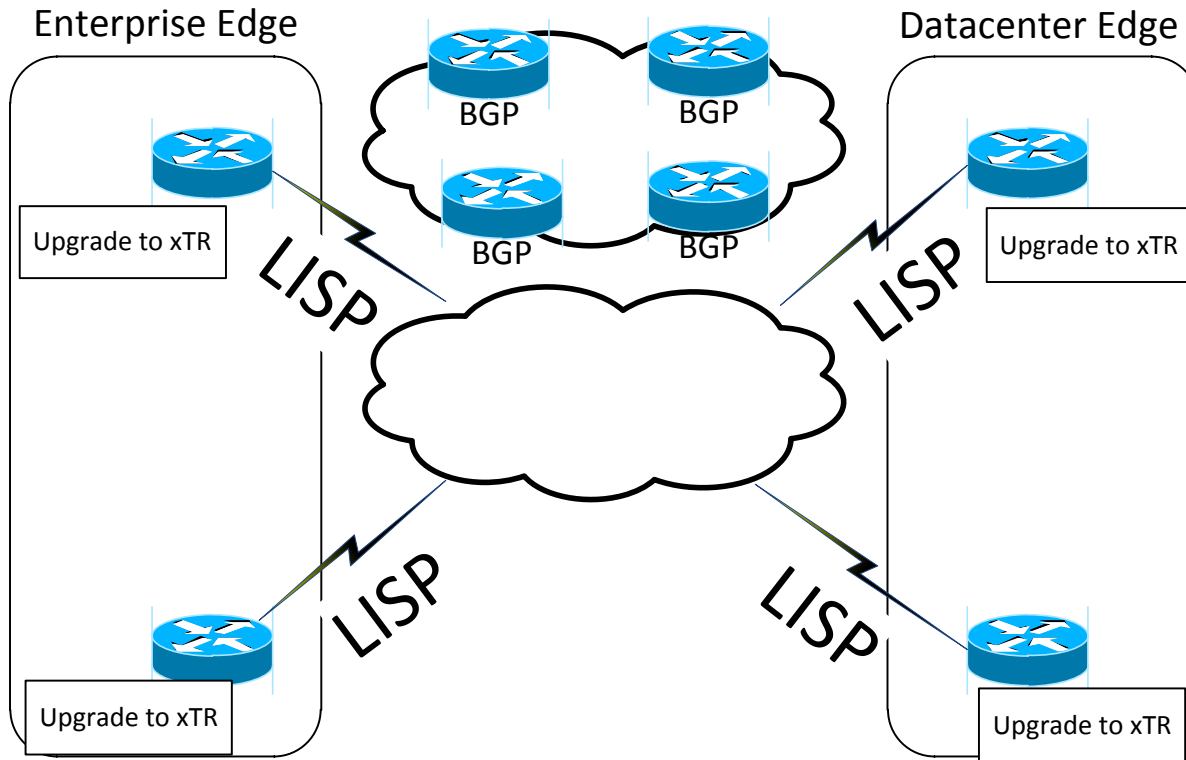
# Customer Self-Reliant Deployment



# Customer Self-Reliant Deployment



# Customer / SP Deployment



# Where do the pieces go?

Clients do everything

Or...

- **Client xTR** – Site
- **Datacenter xTR** – Client or SP Site
- **Service Provider** – Site and/or Infrastructure





# Is LISP on your network today?

- LISP Control – UDP/4342 (0x10F6)
- LISP Data – UDP/4341 (0x10F5)

Router#sho ip cache flow | inc 10F

SrcIf	SrcIPaddress	DstIf	DstIPaddress	Pr	SrcP	DstP	Pkts
Gi0/0	199.119.74.6	Gi0/1.2	38.122.8.142	11	A7F8	10F5	3
Gi0/1.73	199.119.73.3	Local	38.122.8.142	11	10F6	10F6	2
Gi0/0	199.119.74.6	Gi0/1.2	38.122.8.142	11	AF6C	10F5	20
Gi0/0	199.119.74.6	Gi0/1.2	38.122.8.142	11	8C6C	10F5	1
Gi0/0	199.119.74.6	Gi0/1.2	38.122.8.142	11	EC99	10F5	176K
Gi0/0	199.119.74.6	Gi0/1.2	38.122.8.142	11	DC9D	10F5	7



# Tunneling is done in UDP not GRE

– LISP Data – UDP/4341 (0x10F5)

Our traffic share was 80:1 before LISP. It is 8:1 after LISP

Why not 1:1?

IPSec Traffic: 1 source – 1 destination

SrcIif	SrcIPaddress	DstIif	DstIPaddress	Pr SrcP DstP	Pkts
LI0	38.125.5.14	Local	199.119.75.17	32 C2A6 288C	65
BV1	199.119.75.17	LI0*	38.125.5.14	32 E915 27E5	7
BV1	199.119.75.17	LI0*	38.125.5.14	32 E915 27E5	14K
LI0	38.125.5.14	BV1	199.119.75.17	32 E2DD 9274	9
LI0	38.125.5.14	BV1	199.119.75.17	32 E2DD 9274	12K
BV11	199.119.75.17	LI0*	38.125.5.14	32 9F2E 8E31	4
BV1	199.119.75.17	LI0*	38.125.5.14	32 9F2E 8E31	11K
BV1	199.119.75.17	LI0*	38.125.5.14	32 9F2E 8E31	970
LI0	38.125.5.14	BV1	199.119.75.17	32 941C 3028	20K
LI0	38.125.5.14	BV1	199.119.75.17	32 941C 3028	3




# Production LISP Volume

- It's not just the LISP ~~Beta~~ Network <sup>innovation</sup>

[LISP Beta Network MRTG Index Page](#)

Total PxTR Internet Traffic



# Production LISP Volume

- NJEDge.net is a non-profit technology consortium of academic and research institutions in New Jersey.
- NJEDge's LISP traffic volume is sustained at 2G.

<http://www.njedg.net/>



# Impact to You

You provide transit for Customer A running LISP.


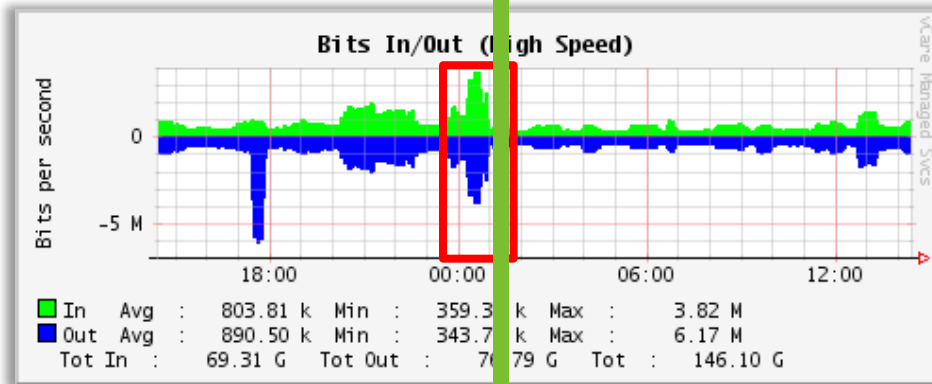
- **HQ** (LISP)
- **DC** (LISP)

You provide transit for Customer B not running LISP.

How does your traffic flow?



# Site vs. PxTR Traffic Flows



# Traffic Flow

Source	Dest	EIDs in MS?	PxTR	Flow
Non-EID	Non-EID	NA	No	SRC - DST global routing
Non-EID	EID	Yes	Via PiTR	SRC to PiTR (IP) PiTR to EID (LISP)
EID	Non-EID	Yes	Via PeTR*	EID to PeTR (LISP) PeTR to DST (IP)
EID	EID	Yes	No	LOC-LOC (LISP) global routing
EID	EID	No	Via PiTR/PeTR*	EID to PeTR (LISP) PeTR to PiTR (IP) PiTR to EID (LISP)

\* PeTR is optional but needed for uRPF and IPv6 over IPv4 environments



# DDT The Global Mapping System for LISP

<http://tools.ietf.org/id/draft-fuller-lisp-ddt-01.txt>

- EID – EID in separate islands is the worst case for traffic.
- DDT allows for lookup across disparate islands.

Source	Dest	EIDs in MS?	PxTR	Flow
EID	EID	No	Via PiTR/PeTR*	EID to PeTR (LISP) PeTR to PiTR (IP) PiTR to EID (LISP)
EID	EID	Yes (via DDT)	No	LOC-LOC (LISP) global routing

\* PeTR is optional but needed for uRPF and IPv6 over IPv4 environments





# Use RIG to see DDT Referrals

## VXNet EID Space

rig 199.119.75.0 to 199.119.73.8

Send Map-Request to DDT-node 199.119.73.8 ... node referral, rtt: 4 ms

EID-prefix: [0] 199.119.75.0/24, ttl: 1440

referrals: 199.119.73.6, 2607:1F00:0:4::6

Send Map-Request to DDT-node 199.119.73.6 ... delegation hole, rtt: 0 ms

EID-prefix: [0] 199.119.75.0/28, ttl: 15



# Use RIG to see DDT Referrals

## LISP Beta Network

VXNet-01#rig 153.16.0.0 to 199.119.73.8  
Send Map-Request to DDT-node 199.119.73.8 ... node referral, rtt: 0 ms  
EID-prefix: [0] 153.16.0.0/16, ttl: 1440  
referrals: 158.38.1.91, 173.36.254.167, 208.217.187.20

Send Map-Request to DDT-node 158.38.1.91 ... node referral, rtt: 112 ms  
EID-prefix: [0] 153.16.0.0/19, ttl: 1440  
referrals: 173.36.254.164, 206.223.132.89, 198.6.255.37, 149.20.48.61

Send Map-Request to DDT-node 173.36.254.164 ... map-server not registered, rtt: 80 ms  
EID-prefix: [0] 153.16.0.0/24, ttl: 1  
referrals: 149.20.48.61, 198.6.255.37, 173.36.254.164, 206.223.132.89



# Operational Issues

## MTU

**LISPAadj:** IP midchain out of LISPO, addr 96.56.153.210  
Change MTU from 9180 to 1464 due to PMTU 1500

**LISPAadj:** IP midchain out of LISPO, addr 96.56.153.210  
pick source RLOC 199.119.74.6 MTU 1464 (pMTU 1500)

*(Our MTU problems were L2 at IX)*

**Latency to PxTR** *(Less of an issue than expected)*

**RLOC = Routing by Uplink** *(Are all uplinks routable?)*



# Resources

- [lisp.cisco.com](http://lisp.cisco.com)
- [lisp4.net](http://lisp4.net)
- [ddt-root.org](http://ddt-root.org)
- [openlisp.org](http://openlisp.org)
- [lispmob.org](http://lispmob.org)



Questions?

