
US ATLAS XCache deployment with SLATE

Lincoln Bryant

US ATLAS Computing Facilities Meeting 7/18/18

Outline

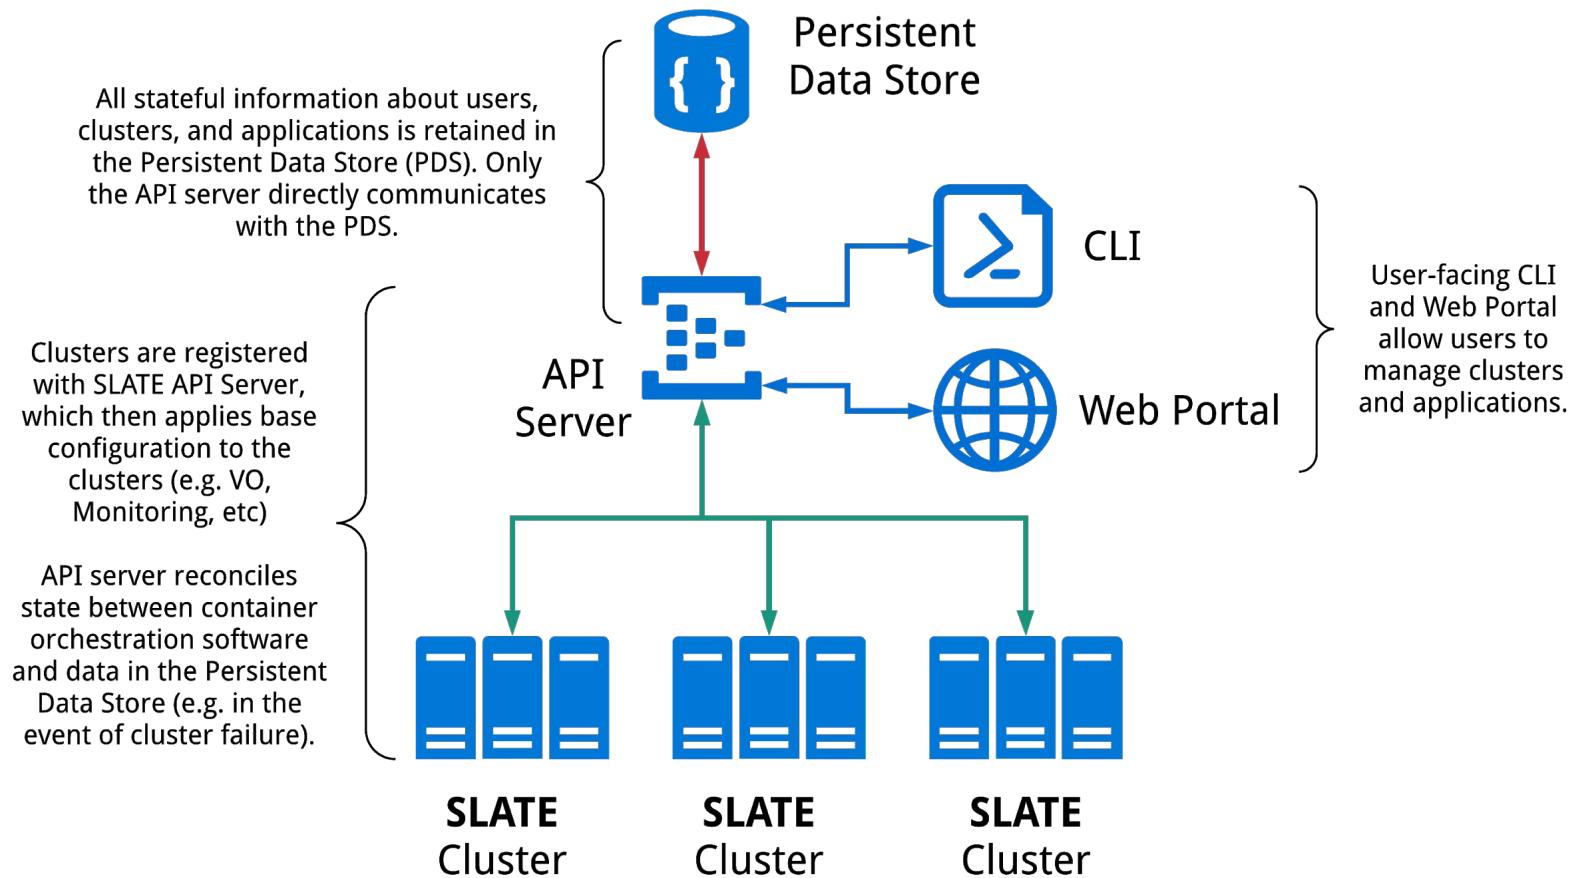


- Review of SLATE
- Site hardware
- Plans for XCache
 - (report next meeting from Ilija & Wei)

SLATE Concepts & components



- Containerized services in managed clusters
- Services curated with security in mind
- **Standalone or federated**
- Widely used open source technologies for growth and sustainability
 - Kubernetes
 - Helm

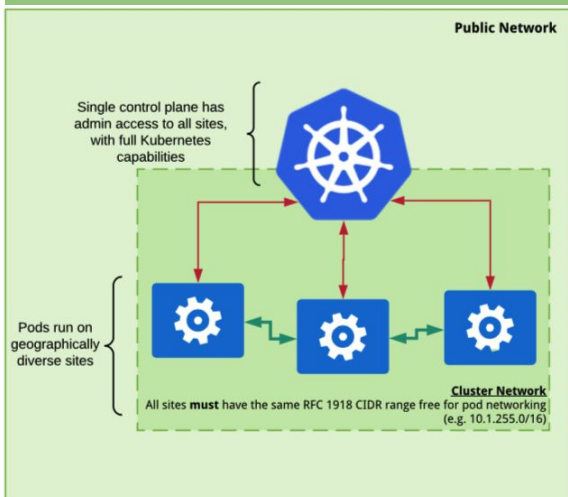


Federation Topology



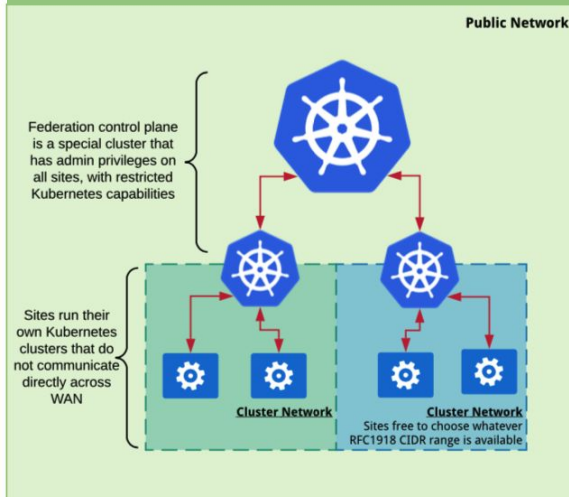
- We are currently pursuing the third approach, although this is not set in stone

Stretched K8s (PRP)



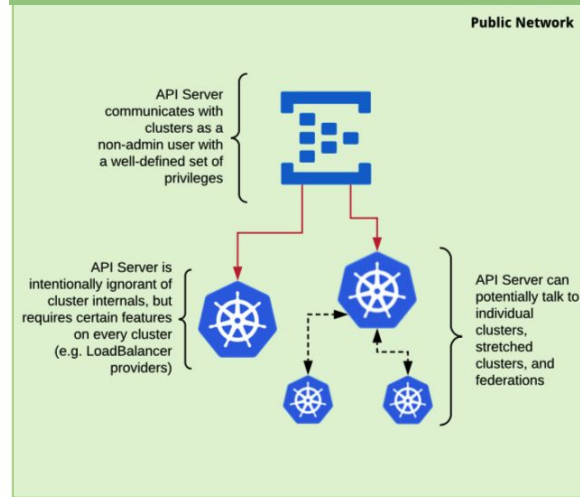
STRETCHED KUBERNETES (a la PRP)

Native K8s Federation



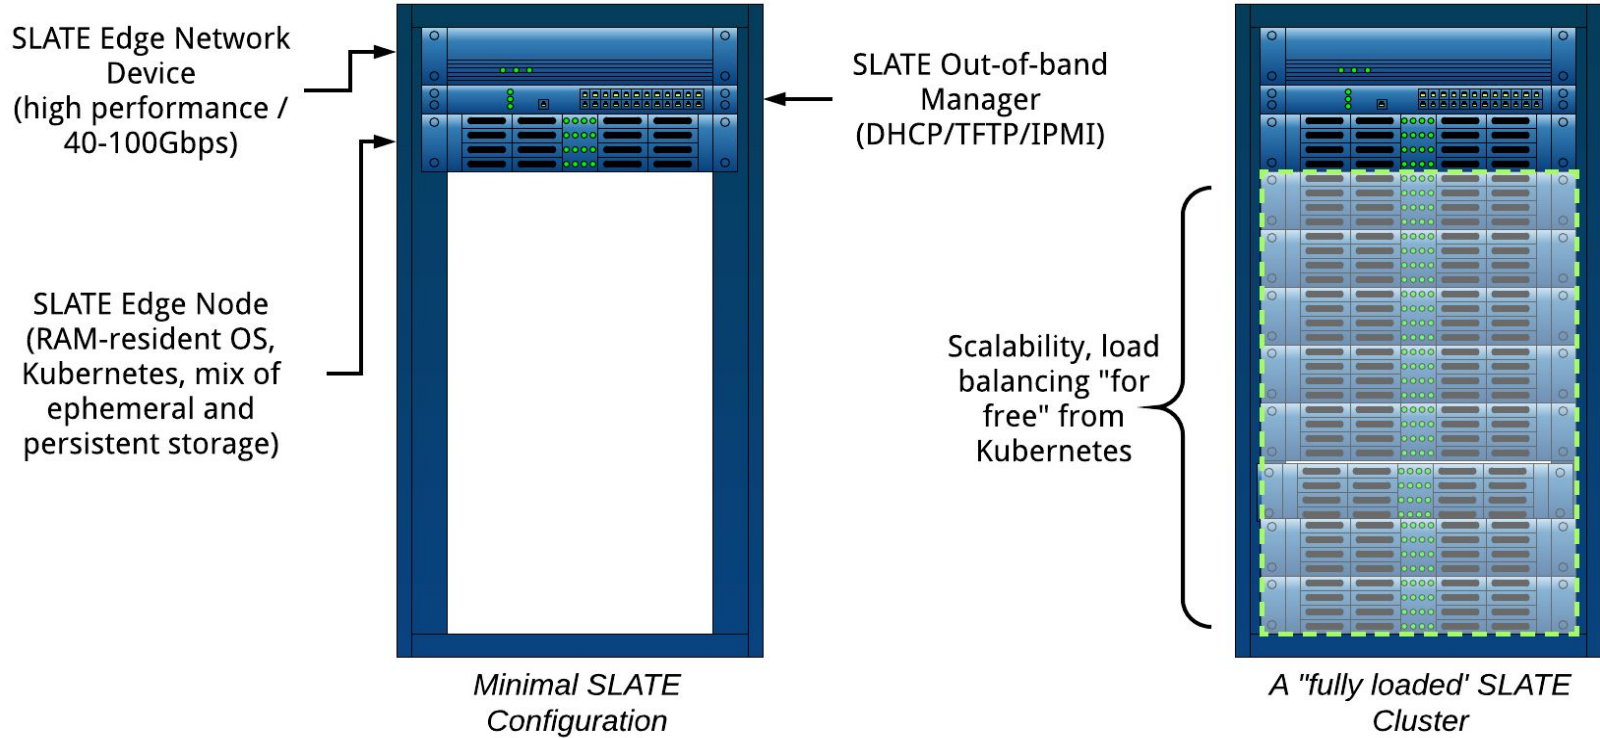
FEDERATION CONTROL PLANE

SLATE Federation



SLATE PSEUDO-FEDERATION

SLATE Hardware Architecture



SLATE Hardware - OOB Manager



Out-of-band Manager

- R440
- single CPU
- 32GB RAM
- 480GB SSD

VENDOR PART NO.	ITEM	QUANTITY
210-ALZE	PowerEdge R440 Server	1
384-BBQW	Motherboard	1
461-AADZ	No Trusted Platform Module	1
321-BCUV	2.5" Chassis with up to 8 Hot Plug Hard Drives	1
405-AAOM	Internal PERC	1
338-BLUS	Intel Xeon Silver 4114 2.2G, 10C/20T, 9.6GT/s , 14M Cache, Turbo, HT (85W) DDR4-2400	1
374-BBBX	No Additional Processor	1
330-BBHL	Riser Config 1, 1 x 16 FH	1
370-ADNU	2666MT/s RDIMMs	1
370-AAIP	Performance Optimized	1
780-BCDI	No RAID	1
405-AANP	PERC H330+ RAID Controller, Adapter, Low Profile	1
619-ABVR	No Operating System	1
385-BBKT	iDRAC9,Enterprise	1
379-BCSG	iDRAC,Legacy Password	1
542-BBBP	On-Board LOM	1
429-ABBF	No Internal Optical Drive for x4 and x8 HDD Chassis	1
450-AGOY	Dual, Hot Plug, Redundant Power Supply (1+1), 550W	1
350-BBKT	Dell EMC Luggage Tag	1
770-BCJI	ReadyRails Sliding Rails Without Cable Management Arm	1
815-3441	Dell Hardware Limited Warranty Plus Onsite Service	1
815-3443	Basic Hardware Services: Business Hours (5x10) Next Business Day On Site Hardware Warranty Repair 5 Years	1
370-ADND	16GB RDIMM, 2666MT/s, Dual Rank	2
400-ARXG	480GB SSD SATA Mix Use 6Gbps 512n 2.5in Hot-plug Drive, SM863a,3 DWPD,2628 TBW	1
540-BBHP	Intel X710 Dual Port 10Gb Direct Attach, SFP+, Converged Network Adapter	1

SLATE GEN 1 HARDWARE CONFIG:

<http://slateci.io/docs/slate-hardware/management-node.html>

SLATE Hardware - Container Edge Node



Container Edge Node:

- R740XD
- (8+) 8TB disks
- 1.6TB NVMe
- (2) 12C/24T CPUs
- 384GB RAM
- 2x10Gbps or better

VENDOR PART NO.	ITEM	QUANTITY
210-AKZR	PowerEdge R740XD Server	1
329-BDKH	PowerEdge R740/R740XD Motherboard	1
461-AAEM	Trusted Platform Module 2.0	1
321-BCPU	Chassis with Up to 12 x 3.5 Hard Drives for 2CPU Configuration	1
338-BLUT	Intel Xeon Silver 4116 2.1G, 12C/24T, 9.6GT/s , 16.5M Cache, Turbo, HT (85W) DDR4-2400	1
374-BBPQ	Intel Xeon Silver 4116 2.1G, 12C/24T, 9.6GT/s , 16.5M Cache, Turbo, HT (85W) DDR4-2400	1
370-ADNU	2666MT/s RDIMMs	1
370-AAIP	Performance Optimized	1
405-AAMR	PERC H730P+ RAID Controller, 2GB NV Cache, Adapter, Full Height	1
370-ADNF	32GB RDIMM 2666MT/s Dual Rank	12
400-ASIB	8TB 7.2K RPM NLSAS 12Gbps 512e 3.5in Hot-plug Hard Drive	8
403-BBQL	Dell 1.6TB, NVMe, Mixed Use Express Flash, HHHH AIC, PM1725a, DIB	1
619-ABVR	No Operating System	1
385-BBKT	iDRAC9,Enterprise	1
379-BCSG	iDRAC,Legacy Password	1
330-BBHH	Riser Config 4, 3x8, 4 x16 slots	1
540-BBBB	Intel X520 DP 10Gb DA/SFP+, + I350 DP 1Gb Ethernet, Network Daughter Card	1
385-BBLE	IDSDM and Combo Card Reader	1
385-BBCF	Redundant SD Cards Enabled	1
385-BBKG	16GB microSDHC/SDXC Card	2
384-BBPZ	6 Performance Fans forR740/740XD	1
450-ADWS	Dual, Hot-plug, Redundant Power Supply (1+1), 750W	1
389-BTTO	PE R740XD Luggage Tag	1
770-BCIP	2U Combo Drop-In/Stab-In Rail	1
813-6070	Basic Hardware Services: Business Hours (5x10) Next Business Day On-Site Hardware Warranty Repair, 5 Years	1

SLATE GEN 1 HARDWARE CONFIG:

<http://slateci.io/docs/slate-hardware/container-host.html>

Joining SLATE



- Typical SLATE site:
 - PerfSONAR node, Management node, Container Edge node(s)
- Installed via USB flash drive (by on-site staff) or DRAC (remotely by SLATE staff)
- Cluster registered with SLATE Central Services
- Kubernetes configured "the SLATE way" and joined to the SLATE federation

SLATE Kubernetes



- Latest version of Kubernetes (v1.11)
- Single-node "seeder", horizontally scalable
- SLATE Federation with share-nothing infrastructure
- Many additional tools and features configured out-of-the-box
 - e.g., Load Balancer, CVMFS driver, Monitoring, external DNS services, ...



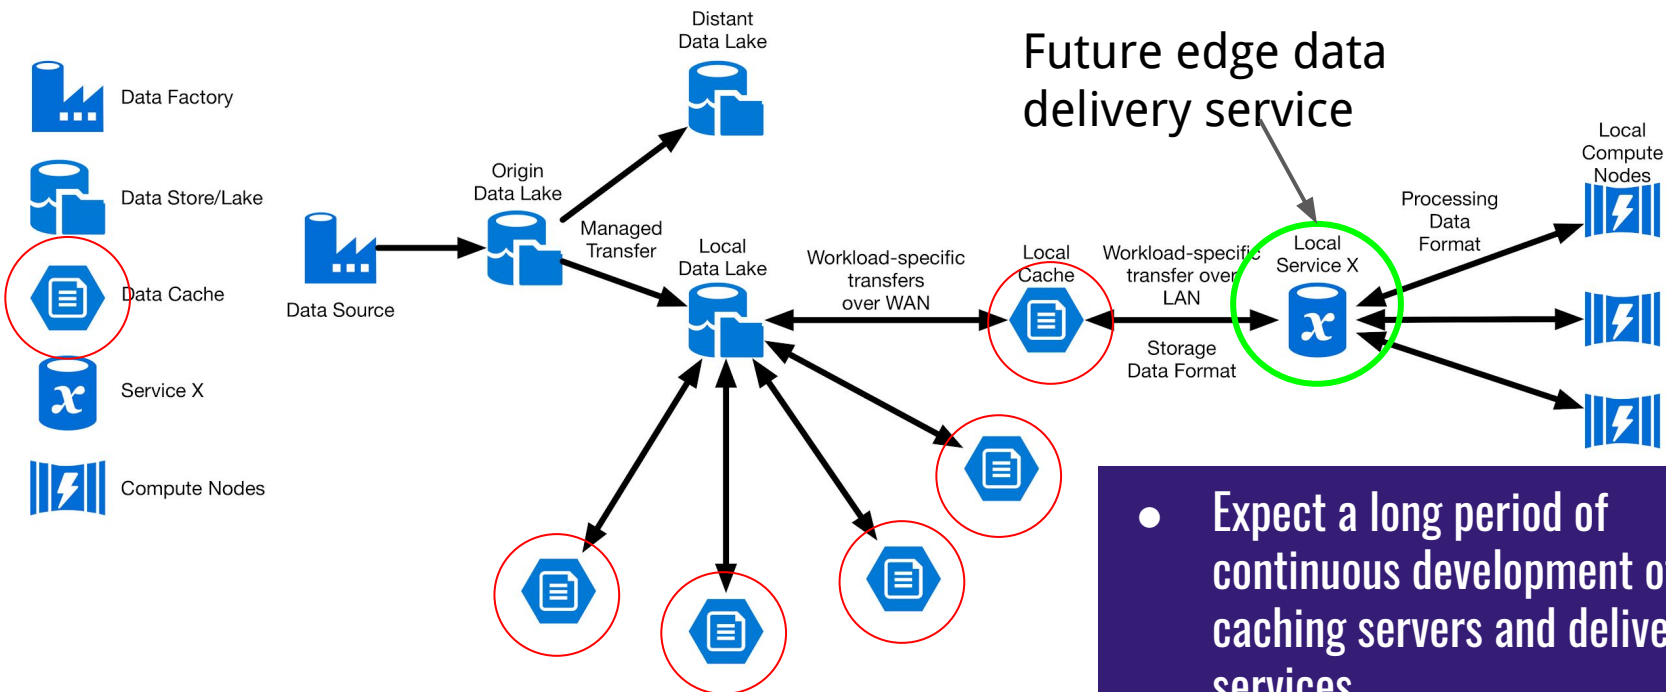
Applications

US ATLAS XCache



- XRootD-based caching service for sites
- Installed, configured, operated and monitored:
 - **by** the XCache Team
 - **on** SLATE hardware
 - **for** ATLAS sites

XCache Service Topology in Data Lakes



- Expect a long period of continuous development of caching servers and delivery services
- Roll out updates centrally
- Configure & control centrally



Extra technical slides

SLATE Bootstrap



- iPXE image provided by SLATE
 - Can either be mounted over NFS via DRAC or via USB stick
- Admin configures network settings
 - for now - static or DHCP over untagged 1G interface best supported.
- iPXE chainloads into a SLATE image hosted on the web
- SLATE image is downloaded and installed to disk