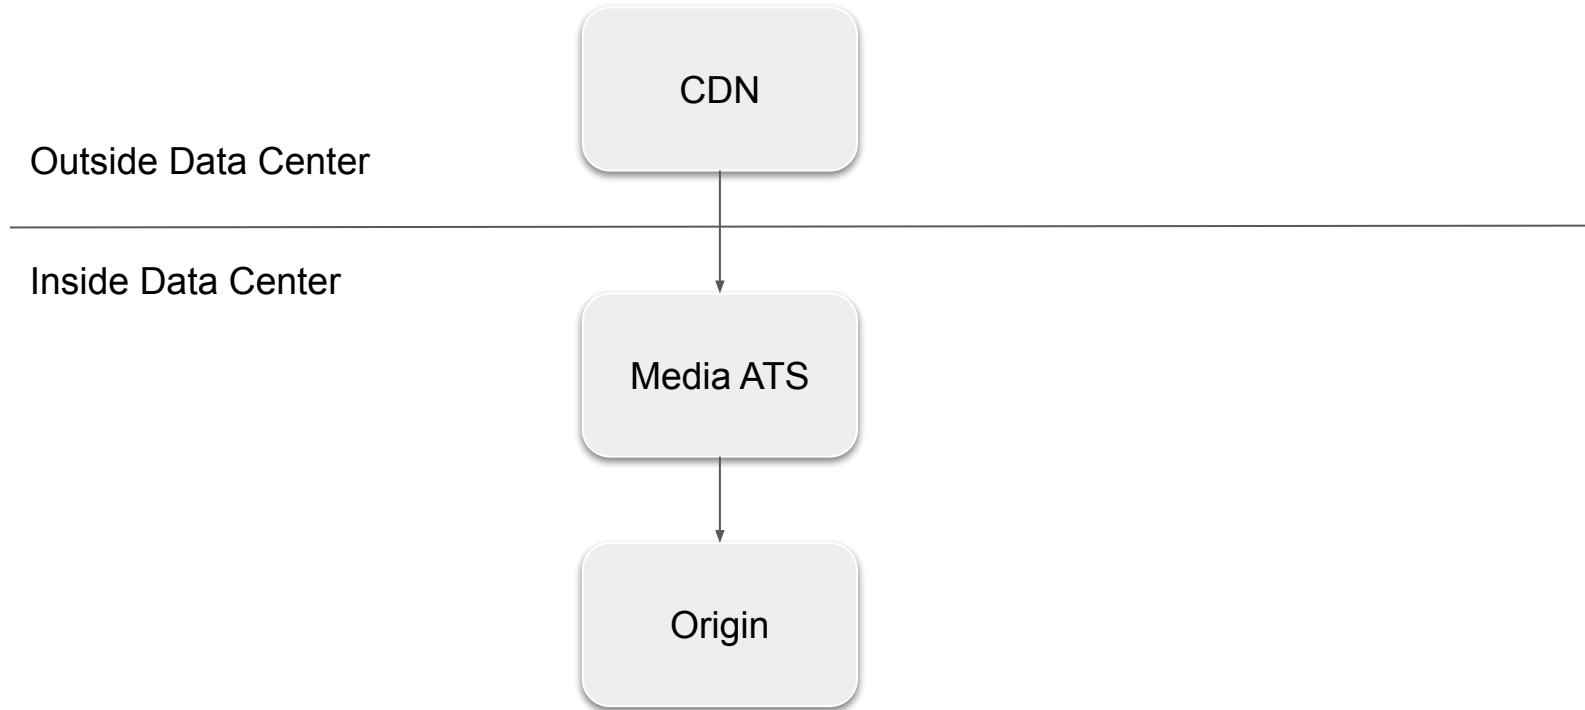


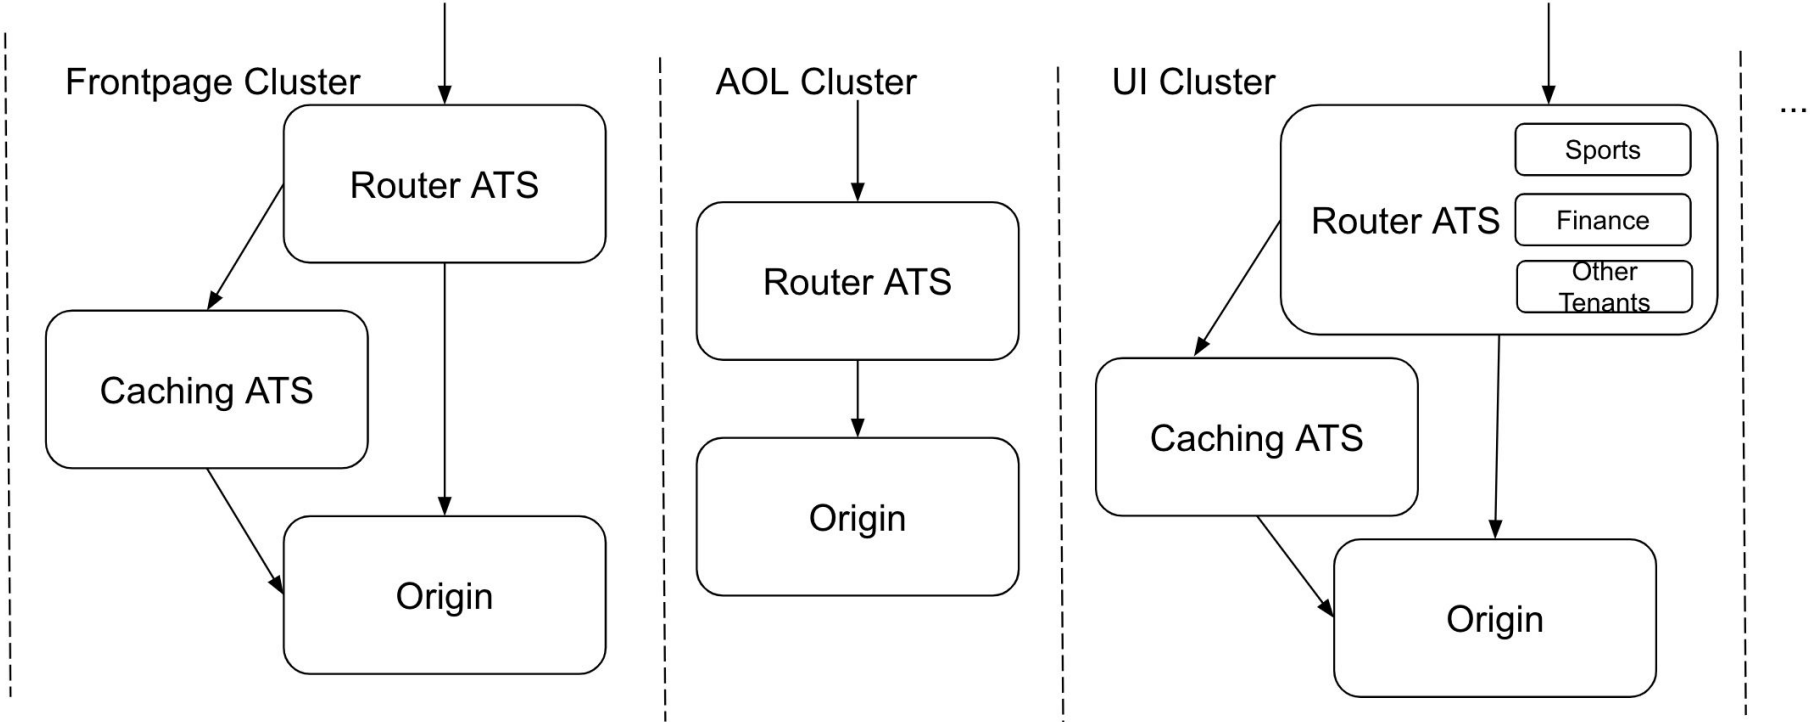
ATS + K8s

Kit Chan (kichan@apache.org)

Verizon Media Properties - High Level Arch



Media ATS



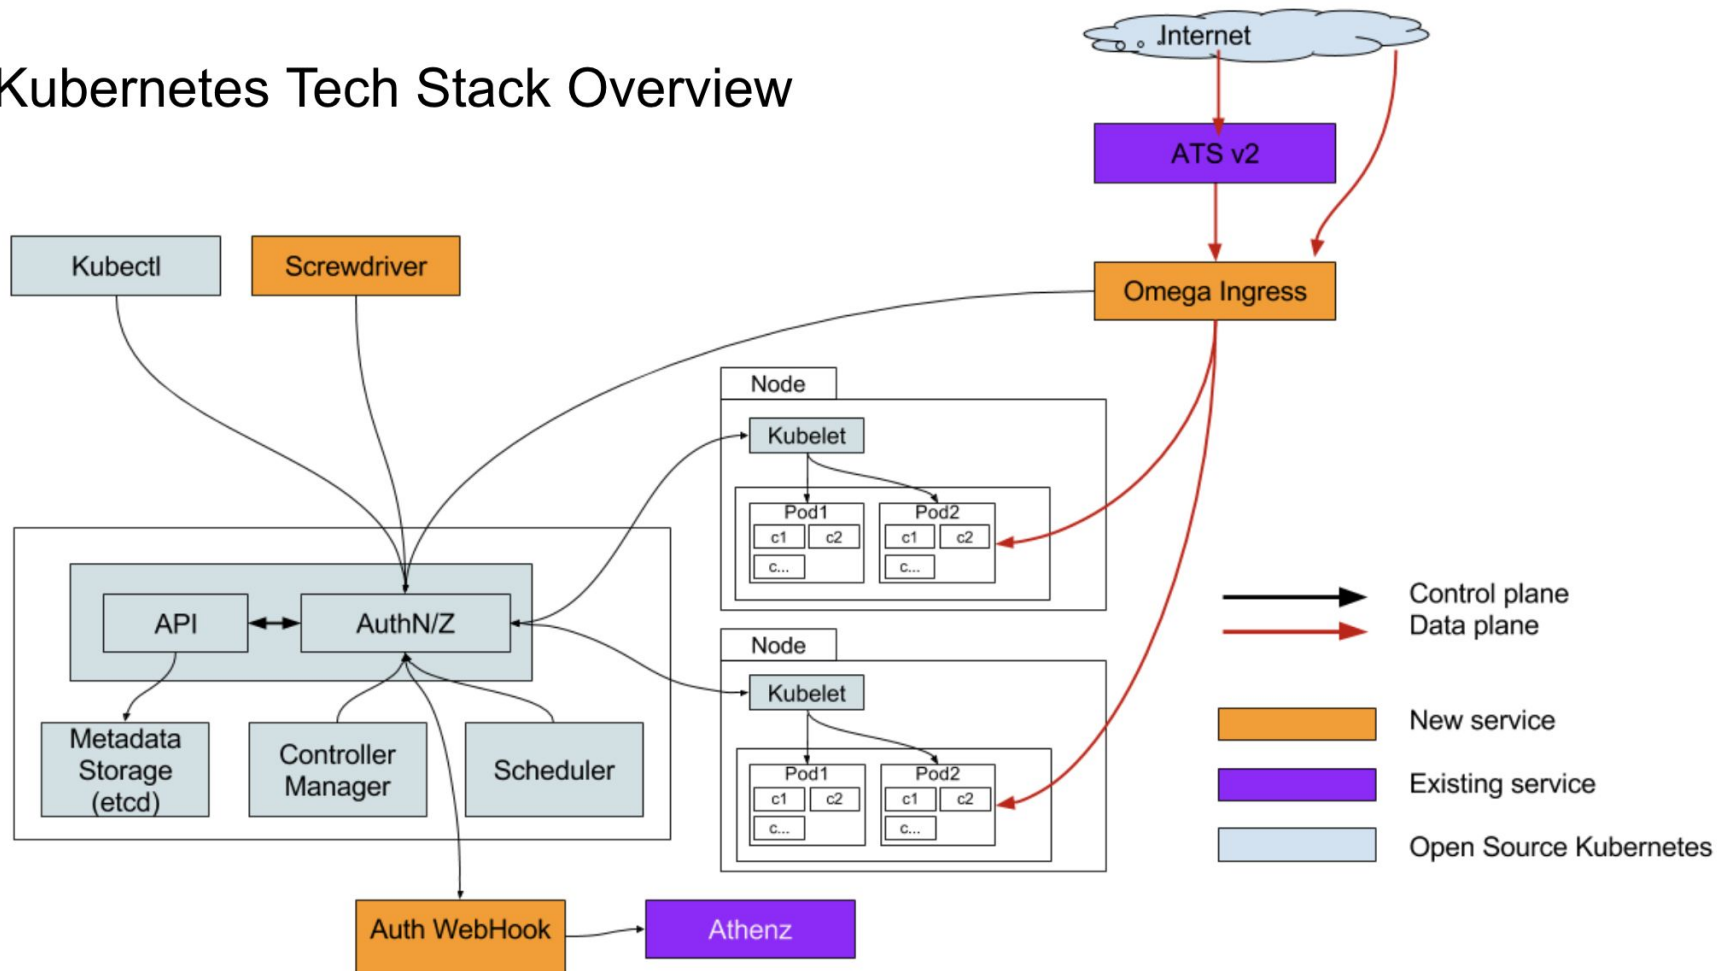
Notes on the Architecture related to this talk

- Ramdisk as cache storage
- Hierarchical caching in use
- A Hardware VIP load-balances multiple router ATS

Phase 1: The Beginning

(The Omega)

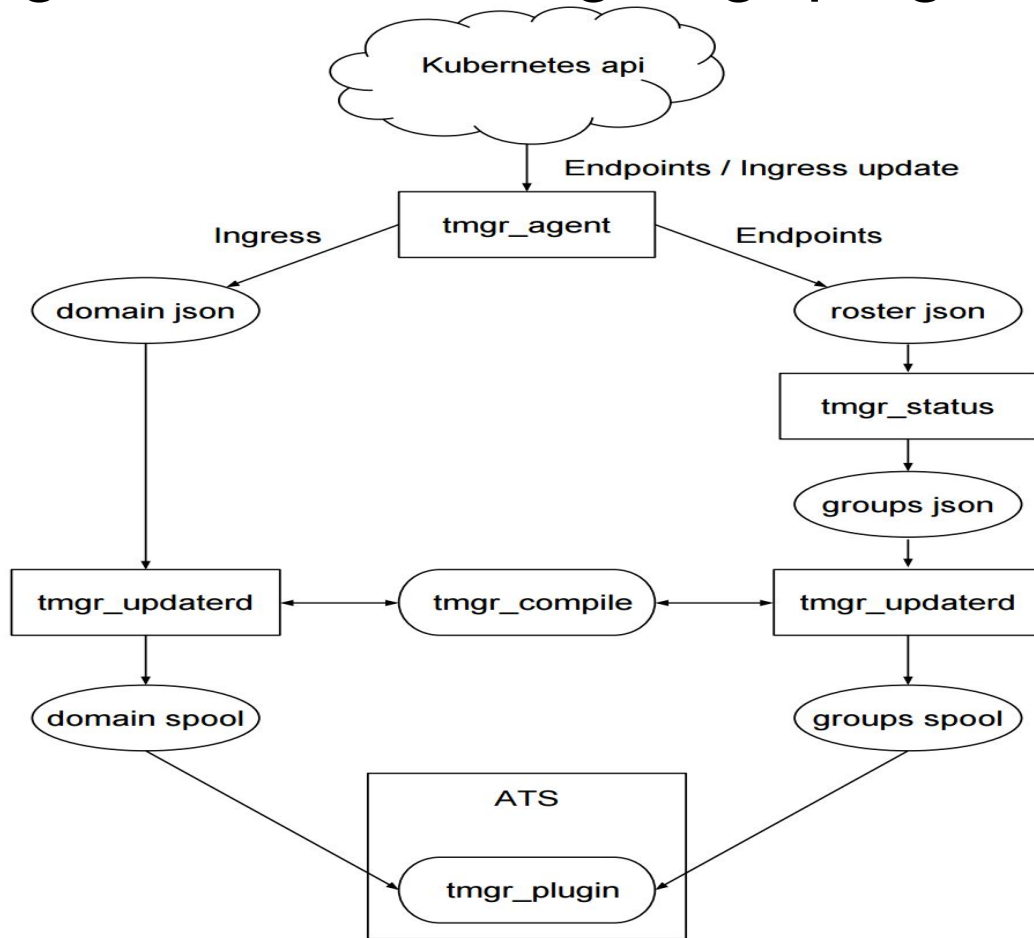
Kubernetes Tech Stack Overview



Omega

- Media has built a *lightweight* set of tools and templates that make it straightforward to deploy Kubernetes workloads onto Yahoo's infrastructure.
- Takes care of common needs such as monitoring, authentications and authorizations
- Handles per-environment and per-colo differences, namespace & access control
- Centralizes base container responsibility for Node.js and Jetty to PE
- Full Kubernetes interface is still available to developers; Omega just makes things **much** easier

Omega Ingress - ATS Using tmgr plugin



Results

- Majority of Media workloads migrated to K8s
 - Savings on machines
 - Higher Utilization
 - Unify deployments and management
-
- Post Phase 1 optimization - eliminating the omega ingress

Phase 2: Caching ATS

Dockerized Caching ATS Manager by K8s

- Base image - from internal RHEL image with ATS + Plugins + Tools/scripts installed
- Extended image - add records.config, storage.config, remap.config, plugin.config + other configurations per environment
- Images hosted in private repo
- Images also contain simple “self-test”
- ENTRYPOINT
 - Create the ramdisk
 - Setup liveness and readiness probes' endpoints
 - Start the server

Details on Deployment

- Initially with Dedicated nodes and hostNetwork
- Side-cars: splunk and other monitoring purposes
- Integrate with authentication and authorization system
- Post deployment tests
- Limiting the capabilities
 - Privilege acquisition in ATS - <https://github.com/apache/trafficserver/pull/4538/files>
 - Or you have to run your docker instance in privileged mode (bad security practice)

Changes to Router ATS

- Similar to using tmgr plugin
- Updates to parent.config
- traffic_ctl config reload
 - Quite slow
 - Affect live traffic

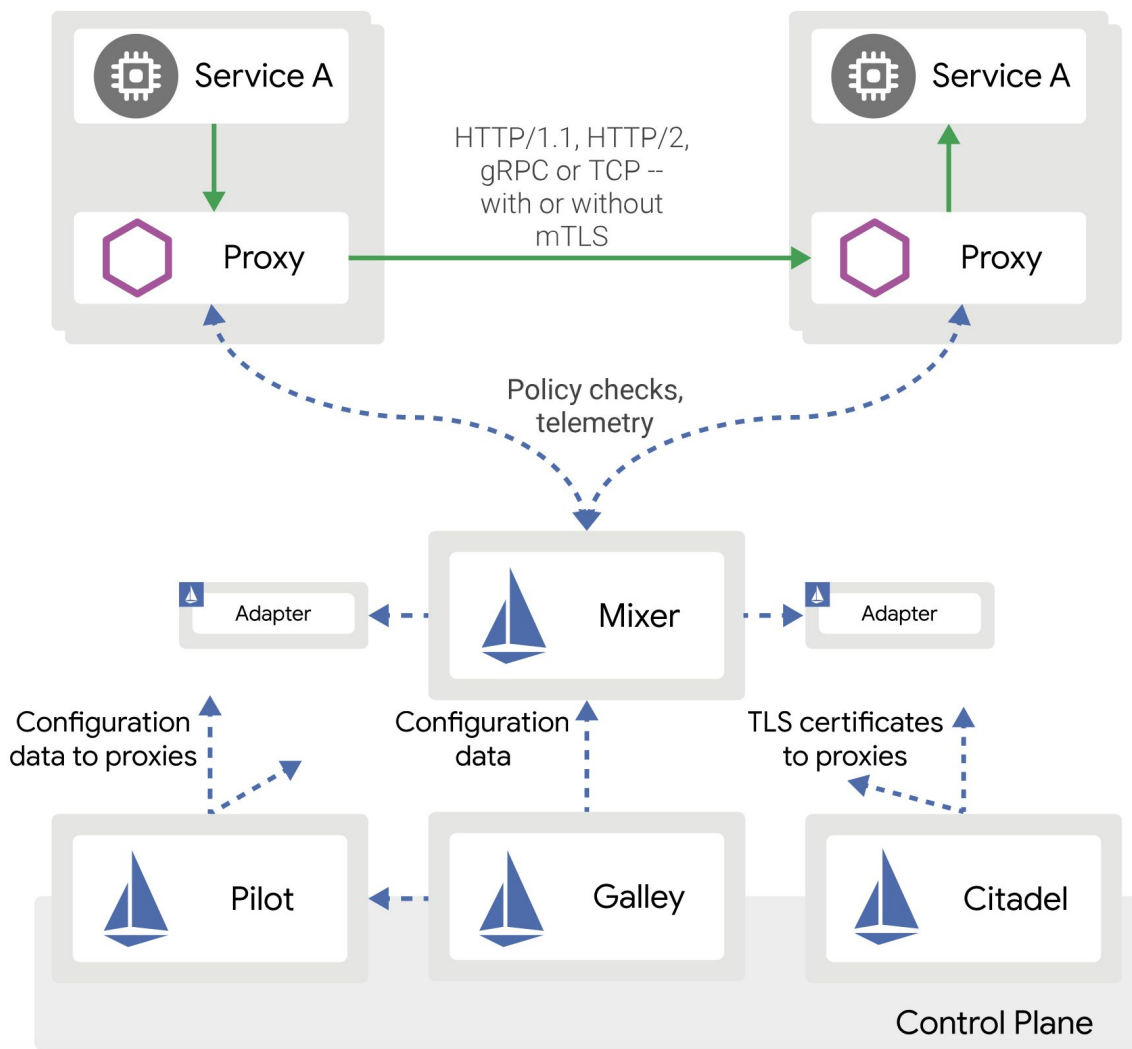
Phase 3: Everything Else

Challenges

- Use cpuset for calculating thread
 - Changes needed in `ink_number_of_processors()`
- More performance tests needed
- Dedicated Hosts and `hostNetwork` initially
 - Thus can use an external VIP for load balancing
- Create K8s Ingress controller for ATS
 - Example with nginx - <https://akomljen.com/kubernetes-nginx-ingress-controller/>
 - Solution already existed? - <https://github.com/torchbox/k8s-ts-ingress/blob/master/example-deployment.yaml>

Future in k8s

- Istio / Envoy
 - E2E mTLS
 - Fault injection
 - Circuit breaking
 - Tracing



Future in K8s - Knative

