Container Identity Working Group proposal

Problem: When a Kubernetes container starts, it often needs to integrate and securely communicate with a wide ecosystem of services, including infrastructure (health checks, proxies, metrics), security systems (secret vaults, firewall systems), other end-user services in its local context, or other services in the wider world.

Goals: It must be possible for:

1. Containers to establish their *identity* via various means to other actors
2. For Kubernetes administrators to ensure that rogue containers can have their identity revoked
3. For existing security tools and solutions in wide use to integrate to link that authorization to external systems or provide assistive credentials
4. For this to be layered on top of Kubernetes instead of deeply coupled to it
5. For higher level concepts, like service identity, to also leverage or align with this approach

Mission: The Container Identity Working Group (WG) will bring together the appropriate experts inside the Kubernetes community (sig-auth, sig-node, sig-networking primarily, with sig-cluster-lifecycle as well) and outside the Kubernetes community (groups like SPIFFE or Istio) to move the discussion forward on:

1. How to allow containers to prove their identity
2. How to integrate existing process identity systems (certificate, kerberos, bearer tokens, proxies) with the goals above
3. The extensions in Kube required for this to be possible
4. How the installation and configuration of a Kubernetes cluster can result in this identity being trusted

Other topics of discussion:

1. How can service identity aggregate and integrate with container identity?
2. How does node bootstrapping (which involves establishing a chain of trust between a node and the master) integrate with container identity?
3. Interaction between identity and external secrets management: e.g. does bootstrap identity for a pod unlock a secret store for other identities, or are there multiple bootstrap identities in a keyring?
4. Identify non-goals and boundaries for this work

Organizers:

- Clayton Coleman, Red Hat

- Greg Castle, Google

This WG will follow the evolving Kubernetes working group process (example) via sig-auth.

First kickoff meeting will be in the next two weeks, please fill out doodle if you have a time preference: https://doodle.com/poll/3q6pgczzke9gwkb5

## Future agenda and meeting notes will now live in this doc

Agenda:

- Kickoff Meeting August 7th, 2017
    - Working group kubernetes-wg-container-identity@googlegroups.com
    - 10 min Introduce and discuss WG goals and mission (clayton)
        - Also discuss non-goals
        - Ist.io hopes
            - Bootstrap identity
            - Secrets and subdivision
    - 10 min Introduce some example use cases from the design discussion doc (greg)
        - External secrets solutions (vault cred)
            - Giving containers bootstrap identity they can use to obtain secrets
        - Enterprise (Kerberos/AD/etc)
            - Importing existing identity into containers
        - Multi-cloud
            - one-to-many (container gets Google and Amazon service account credentials)
        - Service to service
            - Support obtaining/injecting/rotating TLS certs, etc
        - Different priv levels (imagemagick)
            - Pod-level or container-level resolution?
            - Enable sidecar containers with more or less powerful credentials
    - 20 min Open discussion on use cases not covered, identity problems we need to solve
        - Audit/Intrusion detection?
            - Goal is to record and provide sufficient information for auditing/detection to be built on top
        - impersonation/delegation (e.g. helm doing things on users' behalf)
            - Injected credentials could be the starting point for this, but detailed use of credentials seems distinct from delivery/maintenance of credentials
        - Extensible methods for providing data to applications

            - Volume plugins, especially flex plugins (or CSI in the future) seem promising as a mechanism for injecting secret data directly into containers
        - Federation of identities
            - Is it per cluster or across clusters
            - Most organizations have central truth for identity
                - Cloud - cloud providers have their own providers
                - Enterprises - everyone has ldap
            - We can probably assume a global identity solution
                - But we recognize that we should make it possible to solve federation
        - Tenancy
            - Might be useful to discuss in more detail in future
            - Should be possible to build tenancy on top of kube, even if kube by default doesn't force it
                - Should be possible to avoid assumptions than complicate tenancy (like unified identity) by "opting-out"
    - 10 min Settle on next steps, propose:
        - Agreeing on a set of representative use cases we want to solve, using those in the doc as the starting point.
            - Topic for next meeting - discuss some of the ways that identity could be represented at the node (as described in design doc), and make sure we identify the key design points
        - Before next meeting
            - Discuss non-goals on the working group email list to eliminate / debate what we *don't* want to solve
        - Do we want to schedule a regular meeting: an hour every 2/3/4 weeks?
            - What times within the 11am EDT - 2pm EDT window (for both europe and west coast) work for people?
                - Clayton will do a doodle

Related documents shared with the thread:

- Design considerations for container identity https://docs.google.com/document/d/1no8rYJ_nzhMeXYLL6JLjVSDtrj7pd6CBBfE3cPGOv8g/edit#heading=h.xgjl2srtytjt