

Node Attribute Prefix Definition

Node Attribute Types

There are multiple types of node attributes

1. **Centralized**: set by yarn admin user
2. **Distributed**: reported by node managers
3. **System**: set by yarn system services, e.g scheduler
4. **User-Defined**: set by normal users

With the support of these prefixes, we are able to further define what attributes are visible for a request. E.g user A's job might be only possible to see the attributes set by user A, #1 and #2. Attributes set by other users will not be visible.

Format

- A valid DNS subdomain
- Less than 255 characters
- Only contains number/alphabets/dot/dash/underscore

Conventions:

1. **Centralized**: rm.yarn.io
2. **Distributed**: nm.yarn.io
3. **System**: *.yarn.io (yarn.io suffix is preserved for other yarn system set node attributes)
4. **User-Defined**: other DNS form of prefix except reserved ones

Default Prefix

When attribute added/queried from CLI/REST API or used by app, if prefix is not added, it will be default to be considered as centralized node attributes, yarn will automatically add "rm.yarn.io" as the prefix of the node attribute.

Type Conflict

If types have conflict, for example an attribute with type=int rm.yarn.io/attr1 added to node1, and then type=string rm.yarn.io/attr1 added to node2. The second operation will fail. Similarly in distributed node attributes, when node manager reports a attribute whose type is conflict with an existing node attribute, the request will be rejected.

Node Attribute Lifecycle

Node attributes can be added/updated/removed by admin, users, system services or node managers, RM stores all node attributes in the node attribute manager, and handles the persistent/recovery. An attribute will be automatically removed from system if no node in the cluster has the attribute. (Thus we don't need explicit add/remove operation).