January 2023

INTELLIGENT CLUSTER PROVISIONING

Shivram Salem Arunachalam

Raji Vinayagam

Rahul Joshi

Roli Kumar

Sajeev Ramachandran

Follow this and additional works at: https://www.tdcommons.org/dpubs_series

INTELLIGENT CLUSTER PROVISIONING

AUTHORS:
Shivram Salem Arunachalam
Raji Vinayagam
Abhisek Talasu
Rahul Joshi
Roli Kumar
Sajeev Ramachandran

ABSTRACT

Currently, cluster formation through a network management system (NMS) within a software-defined wide area network (SD-WAN) is a manual, multi-step process that is error-prone and time-consuming and which requires multiple interventions (along with validations) by a user at each and every intermediate step. Moreover, that process is extremely difficult to scale with the increased cluster sizes that are required in larger deployments. To address the types of challenges that were described above, techniques are presented herein that support an end-to-end automation of the above-described cluster provisioning process. Aspects of the presented techniques encompass the scanning of an inter-cluster network and the automatic discovery of peer nodes that are viable for cluster formation; the automatic checking of prerequisites for cluster formation (which may include checks regarding a central processing unit (CPU), memory, disk resources, a persona, and a current software version); and support for a 'one touch' provisioning tool that may be employed by a network equipment vendor's support staff or a customer to manage, expand, discover, form, etc. a cluster.

DETAILED DESCRIPTION

As an initial matter, it will be helpful to confirm an element of nomenclature. The narrative that is presented below makes reference to a network management system (NMS) within a software-defined wide area network (SD-WAN) and is referred to herein as a network management system controller. Among other things, the network management system controller may offer a highly customizable graphical user interface (GUI)-based dashboard that simplifies and automates the deployment, configuration, management, and operation of a network equipment vendor's SD-WAN. The network management system

6827

controller may further offer an application programming interface (API) through which it may be possible to, for example, control, configure, and monitor individual devices. One possible API may conform to the constraints of a representational state transfer (REST) architectural style.

Currently, cluster formation involves a manual, multi-step process. Additionally, the process is error-prone and time-consuming and requires multiple interventions (along with validations) by a user at each and every intermediate step. The user must also understand all of the applicable limitations and must check to ensure that each prerequisite that is needed to form a cluster (with respect to compute, memory, storage, and persona requirements) is satisfied. Moreover, the process that was described above is extremely difficult to scale with the increased cluster sizes that are required in larger deployments.

Figures 1 through 3, below, present elements of an exemplary arrangement that helps to illustrate various of the challenges that were described above.
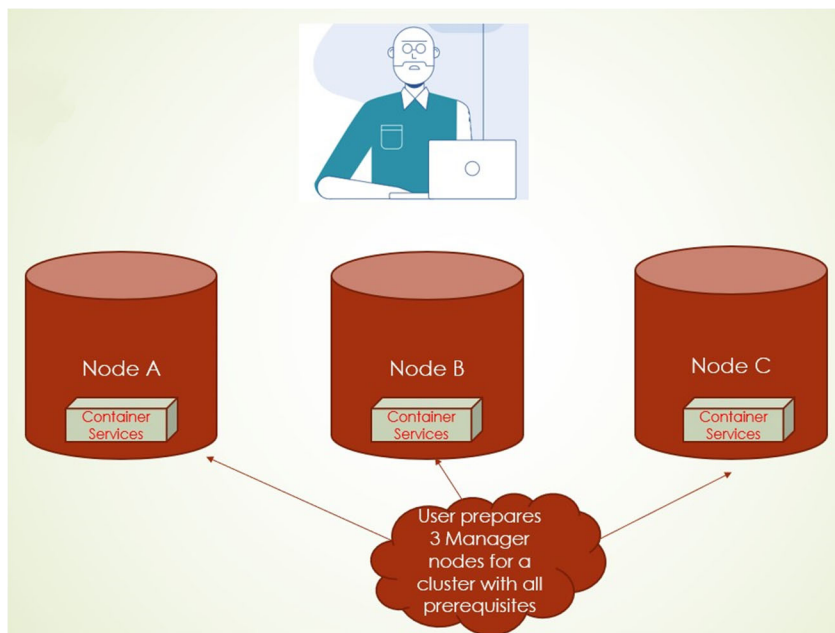

Figure 1: Exemplary Arrangement

In Figure 1, above, a user prepares to perform cluster formation using the current (manual, multi-step) approach as described above. Figure 2, below, illustrates various of the steps that are involved in converting a first node (Node A) into a cluster.

Figure 2: First Node Converted into a Cluster

Figure 3, below, depicts various of the steps that are involved in adding a second node (Node B) and a third node (Node C) to the cluster.


Figure 3: Second and Third Node Added to Cluster

To address the types of challenges that were described and illustrated above, techniques are presented herein that support an end-to-end automation of a network management system controller-based cluster provisioning solution. Aspects of the

3   6827

presented techniques encompass the automatic discovery of network management system controller cluster nodes using an inter-Manager network. Further aspects of the presented techniques support the automatic validation and provisioning of a cluster without user intervention.

Figure 4, below, presents elements of an overall architecture that is possible according to aspects of the techniques presented herein and which is reflective of the above discussion.
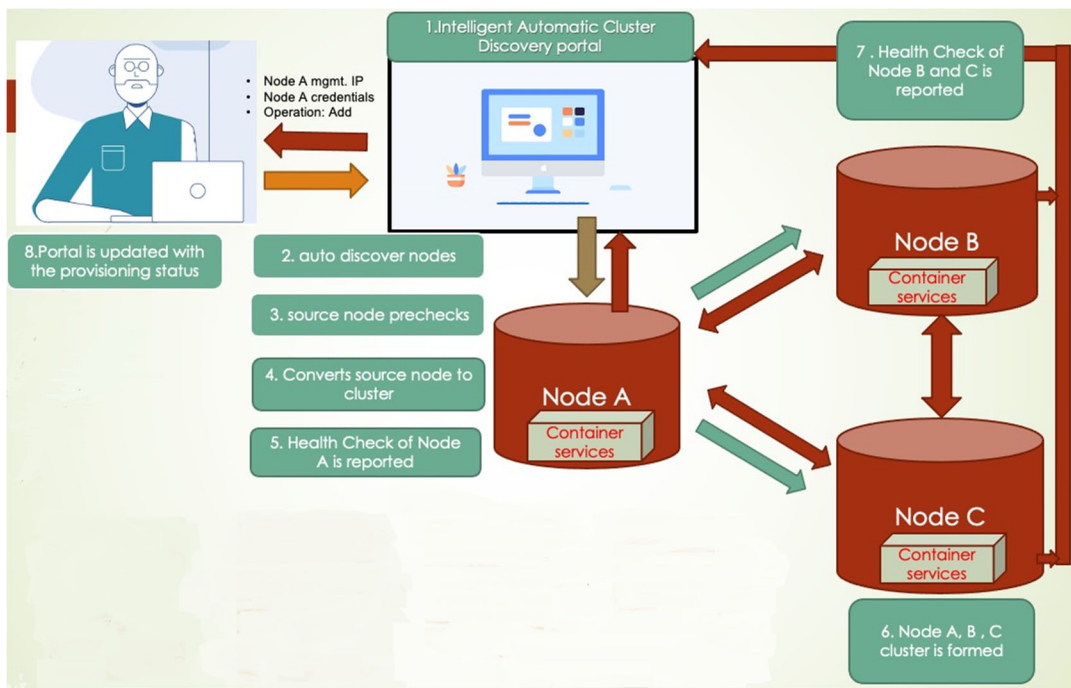

Figure 4: Exemplary Architecture

Figure 4, above, identifies a series of steps, which are labeled 1 through 8 in the figure, various of which will be described below.

Initially, a user may log into an Intelligent Automatic Cluster Discovery portal and provide the Internet Protocol (IP) address and login credentials of a primary or first network management system controller node. Upon providing such information, a request may then be passed to a backend. By using the provided information regarding the first network management system controller node, the system may complete a Secure Shell Protocol (SSH) login to the first network management system controller node and identify a possible

service interface which may be used for cluster formation. Further, a ping sweep may be completed to locate all of the reachable IP addresses in the inter-cluster network. Additionally, an asynchronous client Uniform Resource Locator (cURL) request may be made to each of the reachable IP addresses to identify all of the network management system controller nodes in the network.

At this point, all of the critical information for the first network management system controller node (such as a persona, a software version, and system specifications such as a central processing unit (CPU), memory, and disk space) which are needed for cluster formation may be gathered and stored in memory. The backend may then use the provided administration credentials of the first network management system controller node and try to authenticate all of the other network management system controller nodes through a successful login attempt. Further, all of the network management system controller nodes for which a login attempt failed may be discarded as they do not share the same administration credentials of the first network management system controller node.

Next, the system specifications may be collected from all of the network management system controller nodes that had a successful login attempt. Such a collected system specification may then be compared with the first network management system controller node specification and may be validated in terms of meeting the prerequisites for cluster formation. Once such a validation is complete and successful, the node may then be added to the final list of network management system controller nodes (which are now ready to become part of a cluster).

Prior to cluster formation, the first network management system controller node may be checked to ensure that it is in cluster mode and that it is ready and healthy to form a cluster. After this crucial step, a cluster add operation may be triggered on the first network management system controller node by providing information for a second network management system controller node (from the list of final network management system controller nodes) using a REST-based API. At periodic intervals, a polling may be completed on all of the participating cluster nodes and all of the necessary tasks (such as a cluster add node operation and a restarting and rebalancing of services) may be monitored to completion.

5   6827

Upon completion of all of the above-described mandatory tasks, a health check may be performed on all of the cluster members using a REST-based API to ensure that all of the distributed services are running and healthy.

Figure 5, below, presents elements of an illustrative process flow that is possible according to aspects of the techniques presented herein and which is reflective of the preceding discussion of Figure 4, above.
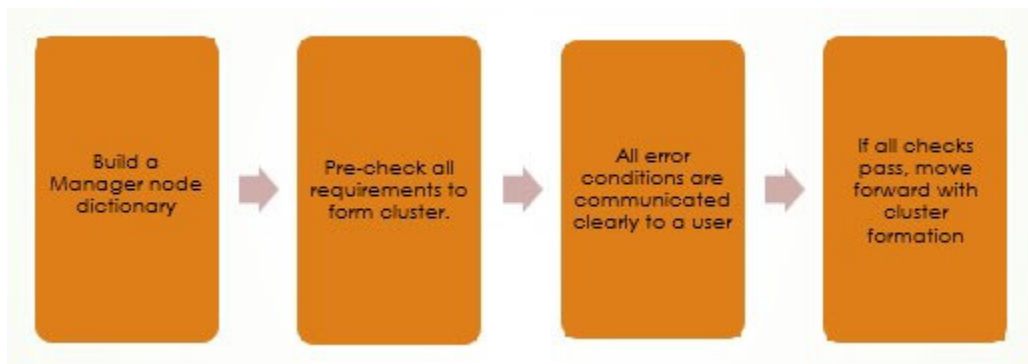

Figure 5: Illustrative Process Flow

According to further aspects of the techniques presented herein, a live log link may be provided in a portal which may be used to view the current state of a cluster formation.

In summary, techniques have been presented herein that support an end-to-end automation of the current (manual, multi-step) cluster provisioning approach that was described above. Aspects of the presented techniques encompass the scanning of an inter-cluster network and the automatic discovery of peer nodes that are viable for cluster formation; the automatic checking of prerequisites for cluster formation (which may include checks regarding a central processing unit (CPU), memory, disk resources, a persona, and a current software version); and support for a 'one touch' provisioning tool that may be employed by a network equipment vendor's support staff or a customer to manage, expand, discover, form, etc. a cluster.

6   6827