Published: 15 November 2024

Authors: T. Eckert, *Futurewei Technologies USA*
M. Richardson, *Sandelman Software Works*

Resilient Remote Managability of Wide-Area Network Infrastructures

Abstract

This is a position paper for the IAB 2024 NEMOPS workshop. It it intended to promote the experimentation and standardization direction of more self-aware and managing infrastructure to support more resilient and easier to manage network infrastructures.

This position paper outlines the problem space addressed by the ANIMA-WG in the IETF, which has released initial standards recommendations. However, these alone are insufficient without a robust network device-level implementation architecture to achieve the intended benefits. The paper also broadens the problem scope, encouraging exploration of alternative and complementary solutions that involve not only IETF contributions but also collaboration with other SDOs and open or proprietary network device development communities.

Table of Contents

1. The Problem 1
2. Non-working or partial solutions 2
3. Current common practice 3
4. Solution space 3
5. Conclusion and Discussion 5
6. Informative References 5
Appendix A. ANIMA beyond remote infrastructure management 5
Authors' Addresses 6

1. The Problem

Today's core network infrastructure challenge is ensuring resilient remote management, whether by human operators or automated systems like SDN, provisioning tools, or other controllers that adjust remote equipment configurations.

Configuration changes often disrupt network services and can prevent further changes by breaking connectivity to the remote operator or configuration software during the process.

These disruption are a recurring common problem, which only sometimes gains visibility outside of the operators themselves. Big OTT operator have at times released root causes for outages, including explanations like (paraphrased) "we did change some IGP routing parameters in a remote site, lost connectivity and had to first secure personnel at the remote site to fix that problem".

In the U.S., networks providing 911 emergency telephony service are under scrutiny of the FCC (Federal Communications Commissions). It not only imposes fines up to millions of dollars for service outages, but also investigates them and publishes incident reports with analysis and future mitigation recommendations. An example report with the example paraphrased root cause is referenced in [RFC8994], whose technology ("ACP") would have avoided or minimized the service outage.

A well known incident from Canada is the July 8th 2022 outage from Rogers [ROGERS] in which we believe the ACP technology would likewise have avoided more than short-term outages. In this incident, radio towers could not be switched off remotely because they did not have routing connectivity leading to mobile phones still attempting to use them including for 911 services - instead of roaming to other operators cell towers. Likewise, the SDN cloud was thought to be set up for routing redundancy, but turned out not to be.

2. Non-working or partial solutions

One misbelief (in the option of the authors) is that these operational problems can be avoided by complete understanding of the network behavior and the interdependence of all network configurations and the impact of all external events, so that it is possible to model the right configuration as well as sequences of configuration changes to always avoid any non-remotely fixable connectivity issues.

Even when such methods are done with human intelligence, they often require a much larger sequence of intermediate configuration changes than direct configuration changes that can not be applied remotely, because they would cause connectivity interruptions in one step, which could only be restored after validation of this step being successfully executed from remote and then executing the next step.

Likewise, atomic execution of large blocks of configuration and reversal to a prior configuration upon failure of any individual command is only a partial solution to the problem, because it can at most only validate local conditions of failures.

What instead is required is what is evolving as (hopefully) common understanding: Configuration automation systems attempt to incrementally modify the observed operational configuration to match the desired / assumed to be working correctly configuration. However, this alone never guarantees that any configuration change does not create a connectivity interruption that prohibits further remote changes.

An example of such problems where multiple independently running automation programs in a large OTT provider which where reliable for different aspect of routing configuration, yet those configurations did interact, leading again to loss of connectivity that stopped the automation software to continue its work.

3. Current common practice

As a result of these problems, current practice is for network configuration to be distinguished into at least two main blocks.

1. The "services configuration" which involves any configuration elements that need to be created when creating another instance of a well defined service offered to a customer/ subscriber of the network. And accordingly deleted or changed when the service instance is deleted or changed. This is very well automated through provisioning tools, often ending up in so-called "self-service web interfaces".
2. The "infrastructure configuration" including most of the physical (non-edge) configuration, addressing, security / filtering / infra-crypto and specifically L2 / L3 routing configuration is NOT managed by the same procedures and tools as the "services configuration", but a lot more manual and with other change processes.

In support of remote infrastructure configuration, local ad hoc "remote-management" infrastructures are built into remote network equipment locations such as some "bootstrap-PC" connected via a LAN (or console serial ports) to all networking equipment, and that PC is equipped with some 4G mobile phone network remote dial-in.

The result of this current pragmatic state of affairs is that evolving of the infrastructure configuration of a network is by far not as agile as that of service configurations, making it more expensive, reducing the ability to quickly react to necessary updates for security, performance or to introduce otherwise more beneficial infrastructure options (such as better/newer protocol versions/features).

In result, the core problem leads to ossification of wide-area networks to achieve reasonable reliability - a problem not had in DC due to their ability to easily use an out-of-band management network.

4. Solution space

The known valid solution space primarily consists of technologies providing a resilient, secondary network infrastructure to remotely manage the network infrastructure. Originally, when networks where (mostly) not IP based, this was using a hodgepodge of telephone/x.25 line access to serial console ports of network equipment. In 2000, ITU-T specified G.7712/Y.1703 "Data Communications Network", in which a TCP/IP based router network replaces such infrastructure, but still may use a hodgepodge of access to the actually managed network equipment.

These DCNs are still standard in most large Internet Service Providers that evolved over more than 2 decades, but to the extend that they do not provide full secure and fast TCP/IP connectivity to the managed equipment they are often only used in emergencies. Even dedicated TCP/IP management ports as found in many of today's network equipment may not have the same speed as the managed networks own connectivity, so it is not preferred for management operations such as firmware download or diagnostics upload and streaming that require ongoing higher throughput.

As aforementioned, in better designed DC, such an "Out of Band" Management network is common practice and access speed to servers and network equipment for management operations typically fast enough (1/10 Gbps), that all management operations can actually be run across that out of band management network, including any telemetry or other large data transfers. This is of course the case because the additional cost for this out of band network is negligible.

In most wide-area networks on the other hand, management happens in-band across the IP network itself to save on the high cost of the additional wide-area network cost for an out-of-band-network, leading to the aforementioned problems (Section 1) or partial workarounds (local management LAN in remote sites).

In specific type of networks, such as optical networks, even standards specify physically in-band, but logically out-of-band management channels. This allows to provision changes to optical network infrastructures without the aforementioned problems in IP routed networks.

The IETF ANIMA WG has specified in [RFC8994] one universal, but implementation wise challenging approach for the problem, bringing the approach of optical networks to that of routed layer 3 networks. It also includes a high degree of security to avoid the problem of additional attacks when the network is also deployed with less human expert oversight/control, but the core of the solution could equally exist without protection against attacks: A virtual in-band network automatically created by the network equipments, when connected to each other without any operator input, and for which the configuration is also non-configurable by any network management interface, yet providing full routed IPv6 connectivity between some central management site and all transitively connected network equipment supporting the technology. This so-called "Autonomic Control Plane" virtual network operates logically completely independent of the so-called "Data Plane" network, e.g.: the normally operator/provisioning-software managed part of the infraastructure devices.

In mission centric designed networks such as most constrained IoT network technologies, solutions are simpler, yet achieving the same goals, simply by not providing any flexible (and hence risk of being mis-configured) infrastructure functionality, but instead hard-coding all the "Data Plane" necessary to provide connectivity. Arguably also any L2 or L2+L3 switches which out-of-the-box enable all interfaces to operate in a Spanning Tree Domain do provide such an "autonomous" network connectivity, but they do not provide a safe configuration path to change that behavior to a more desirable one (such as involving routing) - and they do of course provide no security or scalability of deployment across diverse type of network links (beyond LAN ethernet).

5. Conclusion and Discussion

The authors suggest to put more focus on this problem area increase reliability/availability of networks, reduce operational complexity of remote infrastructure provisioning/management and hence make wide-area network infrastructures more agile - similar to DC.

During the workshop, it would be very welcome to learn about the understanding of more participants about this problem space. Q&A could include questions such as.

o How do you manage the infrastructure configuration part of your wide-area network ?

o Are driving towards specific designs in your management infrastructure, especially the back end to allow agile re-configuration of your network infrastructures core configuration parts such as routing, security, filtering, hardware, etc. If so, what do you do ?

o Do you employ any of the aforementioned or other specific functionality or additional hardware in the infrastructure itself (out-of-band links/devices) to support this management operations processes ?

o What do you see as the biggest gaps ? What could the IETF and/or other organizations help to solve them ?

o What could we do do make the solution components that we have already defined in the IETF more viable to you ?

6. Informative References

[RFC8990] Bormann, C., Carpenter, B., Ed., and B. Liu, Ed., "GeneRic Autonomic Signaling Protocol (GRASP)", RFC 8990, DOI 10.17487/RFC8990, May 2021, <https://www.rfc-editor.org/rfc/rfc8990>.

[RFC8994] Eckert, T., Ed., Behringer, M., Ed., and S. Bjarnason, "An Autonomic Control Plane (ACP)", RFC 8994, DOI 10.17487/RFC8994, May 2021, <https://www.rfc-editor.org/rfc/rfc8994>.

[ROGERS] "2022 Rogers Communications outage", n.d., <https://en.wikipedia.org/wiki/2022_Rogers_Communications_outage/>.

Appendix A. ANIMA beyond remote infrastructure management

It should be noted that the original goal of the ANIMA ACP solution predates the rise of remote "SDN" (controller) based network management, but was instead targeting to ease the more fully decentralized self-configuration of any type of networks.

For example, IGP routing protocols such as those standardized in the IETF where originally designed to provide fully automatic resiliency and failover under impairment of any subset of nodes and replacement thereof. Alas, this concept was never expanded to the actual self-configuration of the nodes and their IGP. As it turns out, any such self-configuration would already require some form of transitive connectivity or signaling that in today's IP networks only the IGPs enable. Hence the idea of an ACP that provides a fully automatic but minimalistic instance for such secure connectivity, so that distributed automation agents could automatically configure all the required aspects of the network infrastructure services especially IP addressing and routing.

The ANIMA architecture supports the development of such decentralized self-configuring agents through the GRASP protocol [RFC8990], which is built into the ACP and provides automatic, secure network wide self-configuration primitives.

One example type of self-configuration agents that have in the past been demonstrated to show the benefits of the ACP are self-configuration of the various security options for routing protocols including IGPs and BGP, as well as other network services protocols.

Authors' Addresses

Toerless Eckert
Futurewei Technologies USA
2220 Central Expressway
Santa Clara, CA 95050
United States of America
Email: tte@cs.fau.de

Michael Richardson
Sandelman Software Works
Email: mcr+ietf@sandelman.ca
URI: http://www.sandelman.ca/