Using Low Power Wide Area Networks to Provide out of Band Management for Data Centers

J. Redžepagić, Z. Morić, D. Regvart
Algebra University College, Zagreb, Croatia
jasmin.redzepagic@algebra.hr, zlatan.moric@algebra.hr, regvart.damir@algebra.hr

Abstract - This paper presents a study on using low-power wide area networks (LPWANs) to provide out-of-band management for data centers. The research examines the potential of LPWANs to offer a cost-effective and secure alternative to traditional out-of-band management methods or supplement them to provide additional security and fault tolerance. The study also evaluates the technical capabilities of LPWANs in terms of range, throughput, and reliability and how these factors can affect the performance and security of data centers. Additionally, the paper examines the various LPWAN technologies available and their suitability for data center out-of-band management. The research concludes by recommending the best practices for using LPWANs for out-of-band management in data centers and providing guidelines for implementation.

Keywords – LPWAN, out-of-band management, data centers, data center, data center management

I. INTRODUCTION

In managing and overseeing data centers or any equipment in general, a significant challenge that must be addressed is effectively conveying the status of the equipment, even in the absence of a primary communication network. This paper primarily concentrates on the networking aspect that links management of data centers to the Internet. Commonly, communication connections are set up to ensure at least one alternative for each link. However, referring to these as backups is inaccurate since standby links are seldom employed as just backups, typically a data center maximizes the use of every available link to connect to the broader network.

All the available connectivity is normally used both for the regular service of the data center and to provide the management interfaces from the outside. Since the design of the data center demands that it will be able to continue functioning even if the primary communication with the Internet goes down, we configure a link failover connected to a different provider, sometimes even through a different medium. For example, if our primary link is on an optical cable, our backup link will also be on an optical cable but can be wireless or copper based.

Still, it should be running through different paths and be serviced by a provider other than the primary provider. This kind of redundancy is typical and is what most data centers do.

The main reason for using optical links is that they provide enough bandwidth for the primary service of the data center, while using them for management of the data center is secondary since the management functions require an order of magnitude less bandwidth, and management is done from inside the data center.

Historically it was common to have a backup link on copper lines. Since the price difference between copper and fiber links is now nonexistent, much higher bandwidth capabilities of data centers are the primary reason to use optical links.

There are, however, situations when every available link goes down, and they are more common than we think. Although this is the worst-case scenario for any data center, and in the design phase we do everything possible to mitigate this risk, sometimes it is unavoidable, especially in smaller data centers. The reason is commonly communal, or infrastructure works outside the data center perimeter that break or disconnect communication lines. In these cases, we rely on so-called out-of-band management strategies that enable connection to the data center itself through one of the wireless networks.

In most cases this is done through the commercial 3G, 4G or 5G network. Still, as soon as the center needs to switch to this link, we lose the ability to operate the data center and can just manage it to troubleshoot and solve the problem. In most cases for this we use commercial providers. Unfortunately, this sometimes leaves us with a single point of failure since the wireless connectivity provider can depend on the same link we used as our primary. Our provider will have the same communication problems as we have, and the redundancy of the backup link is then completely gone.

A. Commonly used wireless technologies

Total data center communication interruption is the worst nightmare of any operator. Since the main risk is physical interruption of the medium we use to connect, the commonly deployed strategy is to provide as many redundant wireless links to the outside world as possible. The technology we have seen used is mostly based on the 802.11 standards since the main idea is to provide the bandwidth needed for troubleshooting and keeping the data center online.

The problem here is that most of the standards for wireless communication were not designed for long-range

communications, which means that trying to connect the data center through a typical 802.11-based network requires a lot of additional infrastructure outside of the data center. We are talking about implementing access points or point-to-point links on locations outside the data center itself and then providing the link to the Internet from that outside location through the wireless network.

In enterprise computing, everything comes down to providing redundancy. As a last resort, creating a link that is going to enable us to have minimal control over a data center and infrastructure even if the primary and secondary providers go down is something that is extremely useful in all circumstances. Ideally, this communication provider will be completely independent both on the communication level and physically from the rest of the network so that the common failure points can be isolated as much as possible. For this, there are a couple of solutions that are, as of right now, underutilized in this context.

B. Low power wide area network breakdown

As the name implies, there is a complete set of modern technologies intended for communication in the modern Internet of Things world. We use the term low power wide area network to denote all of them since these are the main characteristics we look for in this context. Although being covered by one name, when we compare their different specifications, we will see that there are many characteristics that distinguish these networks apart.

So, what are LPWAN or Low Power Wide Area Networks?

The network part is self-explanatory, we need to concentrate on the rest of the words in the definition. The “wide area” denotes the capability of the network to cover not only points that are close together but instead to enable the communication to cover a wider area, measured in tens of kilometers or more. Inevitably this means that we are using some sort of hub and spoke configuration since we are trying to connect many devices in a single network.

In this paper, we are not interested in the network's low-power capabilities because we will presume that our equipment will run on “normal power” instead of batteries. When it comes to communicating with different sensors used in the Internet of Things applications, communicating using low power means that we can stretch a device’s battery life from hours to years.

Like all technologies, low power wide area networks face their own difficulties, and one of these challenges significantly impacts our intended goal. Networks typically facilitate direct bidirectional communication via a "channel" or a "circuit," contingent on the network type. Most of the solutions discussed in this paper will supply messages as distinct communication units. Consequently, our capacity to "connect" to the data center will be considerably constrained.

This, combined with bandwidth limitations of all the different standards, means that this proposed type of out-of-band communication will enable us only to monitor and have the minimal capability to issue commands to the data center.

II. LPWAN STANDARDS

LPWAN standards [1] on the market right now try to completely conform to the original idea of low power and wide area network. This means that different standards try to solve the problem of communicating over a large area using low power in diverse ways. This contrasts with the much more used standards like 802.11x, which have changed from their initial idea to something that tries to solve multiple different problems at once, and because of it fails its most common usage.

A. SigFox

Sigfox is one of the well-known standards for the Internet of Things created with the idea of connecting not only low-power devices but also devices in places unreachable by regular telecommunication networks, such as underground sensors (for example, water meters in urban areas). This comes at a cost since lowering the power levels required for communication also means reducing the bandwidth to a point where we are talking about 100 bps downlink speed on Sigfox. This standard was controlled by a French-based firm that filed for bankruptcy in January 2022 and is now acquired by another company that continues to run the French network and the rest of its operations.

From our point of view, SigFox is not a network we can use to communicate to data centers since the bandwidth is extremely low and the number of messages a single device can use is limited. We mention it because the network itself has a lot of available implementations, so many people first think about SigFox when LPWANs are mentioned.

B. LORA

LORA is an acronym coming from “LOng RAnge.” The name itself is proprietary name for a proprietary radio communication technique based on CSS technology and it is patented by another company from France.

LORA standard covers the physical protocol while LoRaWAN defines the communications and the architecture of the system[4]. LoRaWan is an official standard under ITU-T named Y.4480. The development is managed by the LORA alliance which has hundreds of members and is an open nonprofit organization.

The standard uses ISM license free part of the spectrum and is dependent on the geographic region. In Europe it works in the 868-megahertz part of the spectrum, in South America its 915 to 928 megahertz and in United States is 902-to-928-megahertz part of spectrum. For some applications 2.5 gigahertz spectrum is also used worldwide. Rates that can be achieved are slow - between 300 bps to 27 kbps. The exciting part of the protocol is that different gateways can receive data from one device simultaneously and then the data packets are forwarded to a centralized network server which means that reliability is high.

C. NB-IoT

Narrow-Band Internet of Things is LPWAN technology that is developed by 3GPP alliance and is mainly targeted towards cellular devices and services. This technology focuses on the low cost, long battery life and high density of the connections. Interestingly, this standard uses guard bands of the normal LTE standard to achieve speeds between 26 kilobits per second to 127 kilobits per second in downlink, or 16.9 to 159 kilobits per second uplink.

Together with LORA this standard is the one that is most heavily implemented in the field and is available across the world. The reason why it is interesting as a LPWAN solution is normally low cost and the low power requirements but the main disadvantage is that in most places it is directly connected to the mobile operators and if we require actual out-of-band communication this means that we are basically unable to use this standard since it is directly connected to being able to use the infrastructure of a given operator – if we decide that we are going to rely on something that is a de facto derivative of LTE, why not use LTE itself to establish a fully functional direct connection in the first place.

III. MESSAGING PROTOCOLS VS CIRCUITS

With all these different network technologies the main thing that is stopping us from directly connecting is their ability to provide us with a circuit or a channel of communication. All of them are based on the idea of the gateway or some other service providing a messaging queue the devices will use to deposit or send messages. Clients can connect and read messages and respond to them. Messages themselves are going to be either distributed or are going to wait in the queue to be read and the reason for all of this is that devices themselves need to be able to decide when to send data to save power .

Our proposed solution for this is to create a management network that is not only going to consist of backup communication loop using one of the standards but also that is going to be based entirely on the messaging running something like RabbitMQ or MQTT. Both of these are different standards of messaging protocols common in the Internet of Things space. When we compare these two, we will probably use MQTT since it is much closer to the idea of using the smallest number of messages to exchange data.

RabbitMQ, although also a similar messaging protocol, is primarily targeted to large systems that are online and exchange many messages. This is in stark contrast to what LPWAN networks are trying to accomplish.

This idea also has one unintended consequence that is extremely useful to us. Since a lot of the monitoring in the data center is based on monitoring different parameters of the data center itself - for example temperature, fire detection, monitoring current and power or monitoring access we can completely decentralize this part of the monitoring network. Instead of using a dedicated monitoring network inside the data center and a separate one to monitor devices from outside, we can use the low-power network to directly get the data into the local MQTT server and then provide synchronization to the servers outside of the data center. Add to that the ability to issue commands remotely and we have a working solution.

IV. INFRASTRUCTURE DESIGN

What is our proposed design? We will divide it into a few separate logical units, starting with the local infrastructure.

The idea is to install all the sensors that will monitor our local environment as we would normally do but to add to all the sensors the capability to send messages. Of course, sensors that are LORA capable from the start will be prevalent. Still, there are things we need to monitor that are incapable of messaging or connecting to anything other than a normal network.

Since we need messages to have a common platform for exchanging information, every sensor should be able to send them.

Ideally, this would be through a local LoRaWan Gateway. Still, it is also possible to use the local network since we are under assumption that we are designing a system to enable troubleshooting our data center from anywhere and that right now we are planning to mitigate the risk of all the links to the internet being disconnected. If we can create a local network of sensors that will be completely independent of the local LAN, this is just a nice bonus.

Next step is to create a local MQTT server that will collect all the information from the sensors. This server will be capable of forwarding all the messages to an external MQTT server outside our monitored network.

These steps cover the idea of monitoring our data center. To add functionality that will enable us to troubleshoot and possibly solve problems remotely we also need to create some sort of way to issue command and reconfigure equipment in the datacenter.

Our proposed way would be to create premade triggers triggered by messages containing commands we will issue from our control client. This part of the solution will have to be custom made, there are no solutions for this that are made for this particular usage case, but this capability exists in the protocol, messaging supports not only messages but also different ways of forwarding complex data.

Idea of premade commands would be to be able to issue specific commands with variable parameters to units in the data center, we could for example create a bridge to our virtualization environment to establish control over virtual machines or create a bridge to internal management network to reconfigure switches and routers.

The last part of the system would be the command-and-control client that will connect to the internal or external messaging server. Its main purpose would be to be able to display run-time parameters of our data center and to issue commands.

Unfortunately, for this paper we can only give a general idea how to design such a system. Further research and development is required to implement our ideas.

V. Limitations of MQTT and RabbitMQ

Main problem that we are going to have with this solution is security. MQTT as a server and protocol are inherently unsecured, and we would have to implement encryption and security for all the data in transit. Using RabbitMQ would solve some problems since it has built-in security for communication, but RabbitMQ has other difficulties when dealing with large numbers of IoT sensors. This dilemma will have to be resolved in the implementation itself.

VI. Conclusion

Implementing LPWAN solution to both monitor and control data centers is interesting from two standpoints: it provides a way of monitoring the data center utilizing a completely separate network, and thanks to a large number of inexpensive sensors it gives us the ability to monitor a lot of metrics that would otherwise require more complex solutions to implement, things like power monitoring, access monitoring and such.

We propose a straightforward design comprising of two separate servers that will enable both local and remote monitoring of data centers, with limited command capabilities. This solution can be created to be completely independent of existing networks, both in the data center and the Internet. The hardware required is already available and implementation can be done using existing software with slight modifications.

Result is a monitoring solution that enables both on-site and off-site monitoring, is resilient to network problems and is scalable to even the largest installations.

Reference

[1] Ayoub, W., Samhat, A. E., Nouvel, F., Mroue, M., & Prévotet, J.-C. (2019). Internet of Mobile Things: Overview of LoRaWAN, DASH7, and NB-IoT in LPWANs standards and Supported Mobility. *Communications Surveys and Tutorials*, *21*(2), 1561–1581. https://doi.org/10.1109/comst.2018.2877382ï

[2] Gill, P., Jain, N., Nagappan N., Understanding Network Failures in data centers: measurement, Analysis, and implications, Sigcomm 11 Conference Committee. (2011). *Proceedings of Sigcomm 2011 and Best Papers of the Co Located Workshops.* Association for Computing Machinery.

[3] Gall, J., & Gall, J. (2002). The systems bible is the beginner's guide to systems large and small , the third edition of Systemantics. General Systemantics Press. ISBN 0961825170

[4] https://lora-alliance.org/about-lorawan/ accesed on 20.2.2023.