

LOW POWER WIDE AREA NETWORK TECHNOLOGY

# **Symphony Link™ vs. LoRaWAN™**

A Guide for Engineers and Decision Makers



 **LinkLabs**



The goal of this whitepaper is to explain the Link Labs wireless technology at a fairly deep technical level. This paper will go past the hype and talk about technical facts. It will position the advantages of the system upfront, and then go under the hood to explain how the system functions to achieve its dramatic differentiation.

## Symphony Link™ is a wireless system developed by Link Labs which is used by industrial and enterprise customers who need high reliability and advanced features in their LPWA system.

Symphony Link™ is built on the LoRa® chirped spread spectrum (CSS) physical layer technology. It is an alternative specification to LoRaWAN™, which is the public network protocol developed by the LoRa Alliance™ for carrier networks. We will also explain the differences between Symphony Link and LoRaWAN, and how Symphony Link addresses the performance limitations of LoRaWAN.

### What is LoRa®?

Essentially LoRa is a waveform. It uses a chirped spread spectrum technology which encodes multiple bits per symbol, has integrated packetization and error correcting, and works with a great integrated baseband digital signal processor, called the SX1301, made by Semtech. The role LoRa plays in both LoRaWAN and Symphony Link is at layer 1, or the physical layer (PHY). It is analogous to Frequency Shift Keying (FSK) or another PHY technology in a different system.

OSI Model	Data Unit	Layers
Link Labs	Data	Application Network Process to Application
	Data	Presentation Data Representation & Encryption
	Data	Session Interhost Communications
	Segments	Transport End-to-End Connections and Reliability
LoRa®	Packets	Network Path Determination & Logical Addressing (IP)
	Frames	Data Link Physical Addressing (MAC & LLC)
	Bits	Physical Media, Signal and Binary Transmission

### What is LoRaWAN™?

LoRaWAN is a protocol developed by the LoRa Alliance™ for use by Mobile Network Operators who want to use unlicensed spectrum to communicate with IOT devices in their network.

### What is Symphony Link™?

Symphony Link is a standardized protocol developed by Link Labs and our customers who love the range of LoRa, but need a level of performance that is not available with LoRaWAN.

[www.link-labs.com](http://www.link-labs.com)

The LoRa® name, LoRa® logo and LoRaWAN™ are trademarks of Semtech Corporation



# THE REASONS CUSTOMERS CHOOSE SYMPHONY LINK™ INSTEAD OF LORAWAN™ FOR DEPLOYMENTS ARE:

## GUARANTEED MESSAGE RECEIPT.

Since LoRaWAN is an asynchronous, ALOHA-based protocol, with very limited acknowledgements, packet error (PER) rates over 50% are common. This is fine for some meter reading applications, but for industrial or enterprise sensor networks or control systems 0% PER is a requirement. We have many Symphony Link customers, from Fortune 100 enterprises to startups, who tried to build on LoRaWAN and failed. The Symphony Link MAC acknowledges every message, both up and down. LoRaWAN, SIGFOX and others utilize a “spray and pray” method of message delivery. This is not appropriate for most industrial use cases.

## FIRMWARE OVER-THE-AIR.

LoRaWAN does not offer the ability to multicast large files to multiple end devices. With Symphony Link, you can update the host firmware on your device after it has been fielded. This is huge advantage early in the IOT evolution, as it allows customers to get to market more quickly, and with less risk. For many customers, this is the biggest reason to adopt Symphony Link. LoRaWAN claims to be able to do firmware upgrades, but it does not seem plausible when you add up the time (which could be days or more), the complexity (since you have to manage the firmware upgrade completely at the application layer), and the effect on network performance (sending lots of downlink in a LoRaWAN network causes packet

error rate to skyrocket, since the gateway is transmitting often, but the nodes don't know that when they are trying to send messages).


## NO DUTY CYCLE LIMIT.

LoRaWAN is designed around a 1% duty cycle limit driven by the European ETSI requirements. Since this limit also applies to the LoRaWAN base station, its ability to control the network, send commands, or send acknowledgements is very limited. In Europe, Symphony Link uses the Frequency Hopping Listen Before Talk (LBT) plus adaptive frequency agility band, which removes the duty cycle limit. In the 900 MHz Band, there is no duty cycle limit. The 1% prevents LoRaWAN from being used in systems that need the ability to send lots of data at a time.

## REPEATERS.

Since Symphony Link is a synchronous protocol, repeaters allow users to expand the range of the network dramatically without impacting latency. Duty Cycle limits of the LoRaWAN architecture prevent repeaters from operating with that system. Repeaters cost many times less than an outdoor access point, and thus allow Symphony Link customers to cover larger areas without spending additional thousands on infrastructure. They are also very power efficient, so repeaters can be solar or battery powered.

## Symphony Link™ Architecture







## QUALITY OF SERVICE.

With Symphony Link, the gateway is in control of the network it is creating, so we implemented a quality of service (QoS) tiering system to allow nodes with important traffic to have priority over devices with lower priority traffic. You don't want an alarm to have to compete with a water meter for channel access.

## NO PER-DEVICE CONFIGURATION.

Probably the biggest headache when using LoRaWAN and other IOT technologies is the complicated management of multiple per-device encryption keys both at the time of device production, and on the server side. With Symphony Link, the host device configuration is the same for all devices of the same type, and key exchange is handled via our advanced PKI-based Diffie Hellman AES architecture.

## REAL TIME POWER AND DATA RATE CONTROL.

One of the big weaknesses of LoRa is that its in-channel dynamic range is limited to about 20 dB. Since both LoRaWAN and Symphony Link take advantage of transmitters using the same frequency at the same time, dynamic range is important. In LoRaWAN, a transmitter that is close to the base station will prevent the node further away from being heard, due to dynamic range issues. And since there are few acknowledgements in LoRaWAN, these messages would be lost. In the future if large LoRaWAN public networks with many base stations in range are built, this effect would be somewhat mitigated. In Symphony Link, before every transmission, an end device calculates the reverse link to the gateway, and adjusts its transmit power and spreading factor or modulation rate to match.


This way nodes throughout the network have a balanced link budget. Close nodes are transmitting quietly and quickly, and far nodes are transmitting loudly and slowly. LoRaWAN has Adaptive Data Rate as well, but since it is driven by the server, if a node's link suddenly fades, the server has no way of telling it to change spreading factors to compensate. ADR for LoRaWAN is about optimizing capacity. ADR in Symphony Link is about optimizing performance and reliability. Interestingly enough, Symphony Link ADR instantly optimizes for capacity even better than LoRaWAN.

## SECURITY FLAWS.

While LoRaWAN security flaws do not pose a significant risk for most users, the use of pre-shared keys and identities does create vulnerabilities. With Symphony Link and the use of a Public Key Infrastructure (PKI), the over-the-air wireless channel is considered unbreakable by NSA standards. PKI also prevents spoofing and assures identity of infrastructure.

## MULTIPLE GATEWAY COEXISTENCE AND INTERFERENCE MITIGATION.

In LoRaWAN, all gateways, no matter who owns or operates them use the same frequency channels. That means that your LoRaWAN network sees all of my traffic and vice versa. My server just can't just decrypt your messages, but it eats into the capacity all the same. This type of layer 2 interference could be very problematic as a LoRaWAN networks scale. Symphony Link uses a dynamic channel mask that is controlled by the gateway, so it ensures as few collisions as possible. In the US, Symphony Link uses 28 times more spectrum than LoRaWAN, and in Europe 7 time more.







## HIGHER CAPACITY.

By using asynchronous features like slotting, and uplink/downlink coordination, a Symphony Link network has over **4 times the capacity of LoRaWAN**. And when you couple that with quality of service, Symphony Link is a much more robust choice for users that need it.

Operating a network does not require a network ID from the LoRa Alliance. A recent policy announced by the LoRa Alliance is that anyone operating a LoRaWAN network needs to be granted a NetID by the Alliance. Since these are only available to Contributor members and above, be sure to budget **\$20,000 / year for the privilege of using LoRaWAN**. This was likely implemented to create a bit of a barrier since more uncoordinated LoRaWAN usage in an area is a very bad thing, as mentioned before. Symphony Link does not interfere with LoRaWAN and vice versa, and there is no membership fee required to use it.

There are also other features unique to Symphony Link that are critical for a subset of users. These include features like **time sync broadcast**, which allows end devices to sync to

the real time clock of the gateway, and an run time based rules or allows fine time stamping of data at the edge.


Also, by using a full frequency hopping scheme, end devices can transmit up to 1W in the 900 MHz Band. This is great for AC powered devices like electricity meters and lights. LoRaWAN achieves its fixed channel scheme by using a hybrid certification, which limits power transmission.

Symphony Link has a **fixed MTU size of 256 bytes**. Our protocol handles all of the sub packetization and retries necessary to get the message through. In LoRaWAN, the MTU changes based on what spreading factor the server assigns, and you have handle packetization, retries, and reassembly at the application layer, which is difficult.

Finally, because LoRaWAN encrypts all traffic both up and down on a one-to-one basis, implementing **multicast for control systems** like lights, if very difficult. Symphony Link implements multicast session keys which allow for groups of devices to be addressed. This is also what allows Symphony Link to transfer firmware over-the-air.

## Symphony Link™ is the only LPWAN technology that allows users with use cases like:

- Lock Control
- Demand Response
- Industrial Control Systems
- Lighting Control
- Alarm Systems
- Physical Security





# Symphony Link™ Protocol Description

Say we just turned on our Symphony Link gateway. The first thing it does is scan the band, and create an interference profile. Please note that we are explaining how the system operates in 900 MHz. For 868, the operation is slightly different, since the beacon channel is fixed, but uses a TDMA scheme.

Once the interference scan is complete, the gateway selects a 500 kHz channel for its downlink (125kHz for Europe), and then listens to that channel to make sure there is no weak LoRa traffic there, which would indicate another Symphony Link gateway further away has chosen that same channel. Also note that this is the automatic channel selection mode, the user can also set the channel manually via the network manager interface of our network management software.

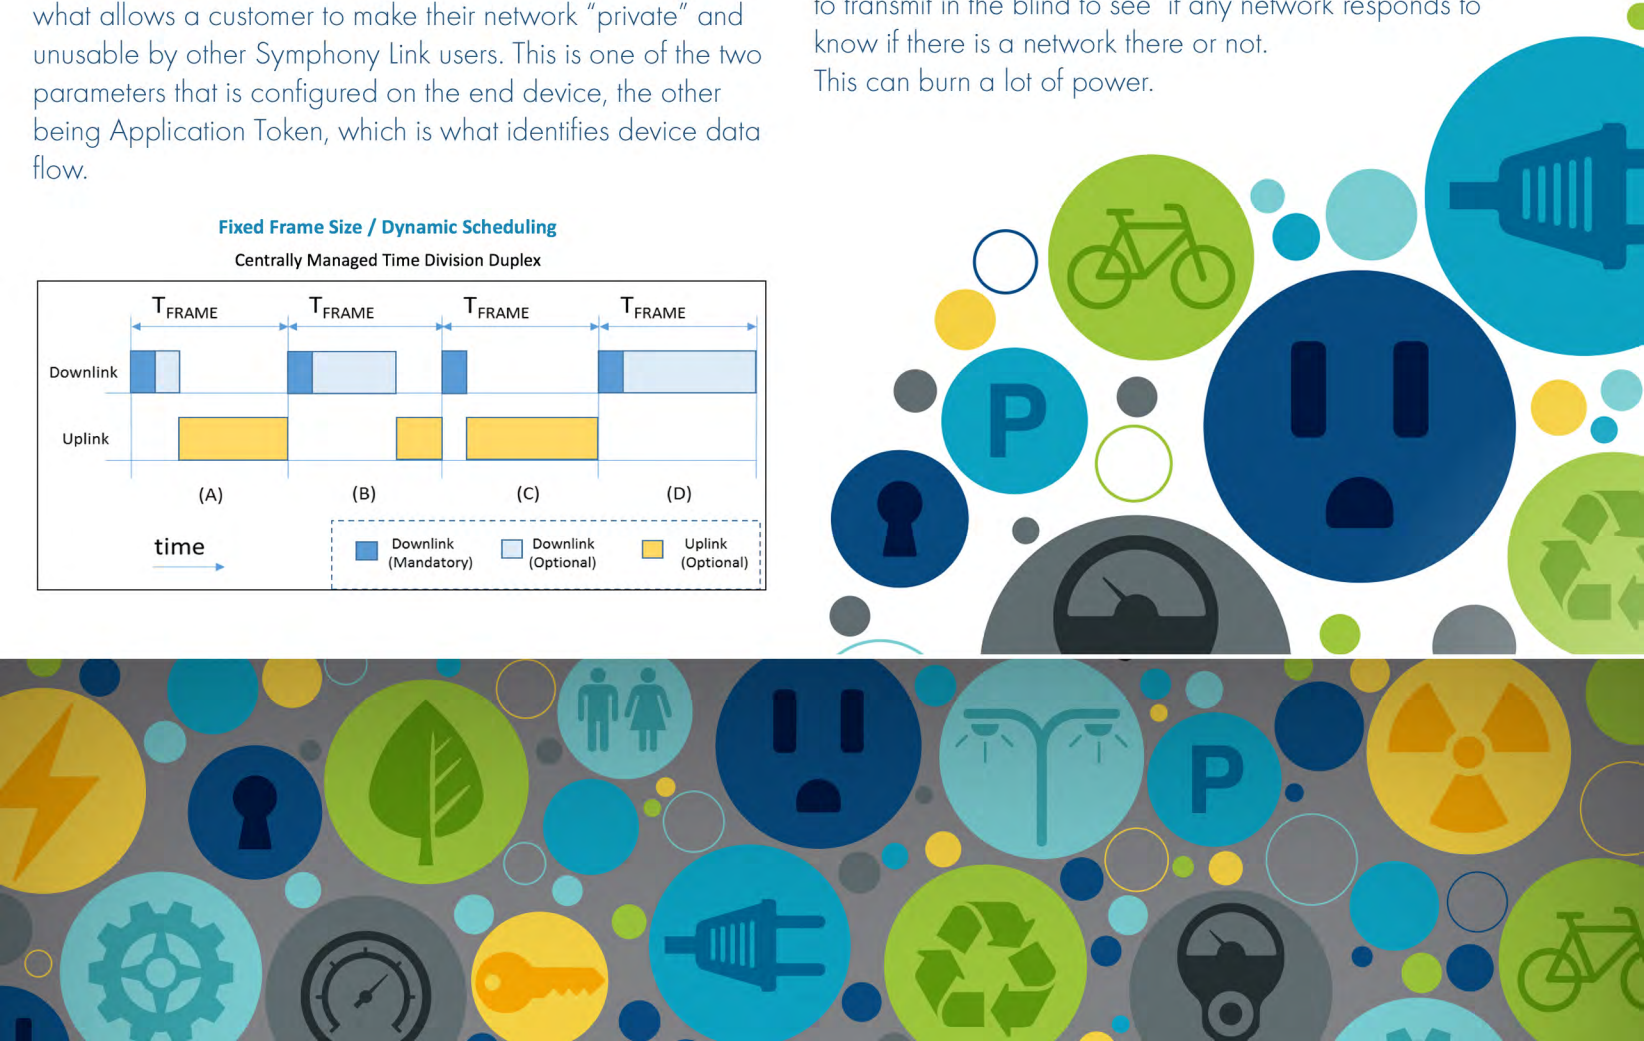
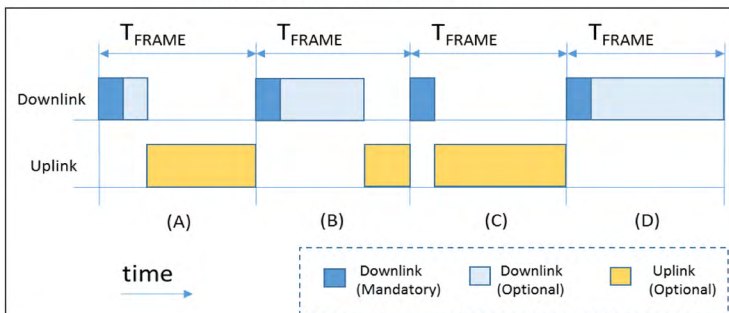
So, this channel is chosen, and now the system begins transmitting every two seconds. This message can be called a beacon or frame header. And this frame header contains several important bits of information.

First off, this message is encrypted with the network ID. This is what allows a customer to make their network “private” and unusable by other Symphony Link users. This is one of the two parameters that is configured on the end device, the other being Application Token, which is what identifies device data flow.

The second piece of information is the uplink/downlink time boundary. This tells the nodes that are awake during this frame, when the gateway will be done transmitting. Since LoRa is a half-duplex technology, it is important to prevent up/down collisions, which are very prevalent in LoRaWAN. A LoRaWAN gateway must respond to a node’s request for Acknowledgement or downlink within a fix time frame. Any LoRaWAN message sent during this time will not be received by the gateway. Requesting more acknowledgements in LoRaWAN just compounds the problem exponentially.


A third piece of information is the uplink channel frequencies of the forthcoming uplink frame. Since Symphony Link utilizes a “block hopping” uplink method, where the bank of receivers hop in every frame, the nodes need to be told where these channels are. This is also how Symphony Link can have so many more networks on the air at a given time without interference. Also, since the gateway is telling the nodes what channels are available, Symphony Link can have a 1 channel, 8 channel, or 64 Channel gateway, and the endnode doesn’t care. A repeater or inexpensive single channel gateway can bring the same features to the network as a larger gateway, with the only difference being uplink capacity. Plus there is the added benefit of the endnode being able to detect the presence of a network passively, even across both 868 and 915. A LoRaWAN endnode has to transmit in the blind to see if any network responds to know if there is a network there or not. This can burn a lot of power.

Fixed Frame Size / Dynamic Scheduling  
Centrally Managed Time Division Duplex





## Symphony Link™ Spectrum Diagram



The gateway also sends a quality of service level for that frame, so that if the network is congested, less important nodes can wait before transmitting. Finally, the gateway transmits the compressed acknowledgement packet, which contains an acknowledgement for all messages in the previous frame. In LoRaWAN, the acknowledgements that do happen are always one to one, and when you add the LoRa preamble to those messages, they are a huge bandwidth hog. By compressing ACKs into one message, we save significant time on air over LoRaWAN.

There is additional information that nodes need to transact with the network, and this information is contained in the info block message, which is transmitted by the gateway every 8 frames. This contains information about regulatory power limits, transmit power of the gateway, which is important for nodes as they calculate their power and spreading factor for each transmission. If a network needs to trade reliability for capacity, the info block can also tell nodes to apply additional signal margin to their adaptive power and data rate calculation. The software version of the gateway is also transmitted, to prevent a mismatch in capabilities between the end node and the gateway. The info block message can also turn on or off the

listen before talk mode in the nodes, which is only required for operation under certain regulatory regimes. (The European LBT+AFA band and Japan are major examples.) Also, the info block tells the nodes whether or not they are connected to the network management cloud, and thus whether the PKI server can be used for the public key source, and whether the node needs to register before joining the network. This allows some networks to operate completely disconnected from the internet.

Okay, so let's say that a node decides it wants to transmit an uplink payload to the gateway. Of course, since this is a low power network, it has been completely asleep and idle for some long length of time. It will wake up and tune its receiver to the frequency where it last heard the gateway. It assumes some worst case crystal offset, and thus starts listening some milliseconds prior to the beginning of the gateway frame header preamble.

It then should detect and process the frame header message. It learns when the uplink window starts, and what frequencies are available.

Then it sleeps for the rest of the downlink period, unless it is a "downlink always on" node, which we will explain in a moment.

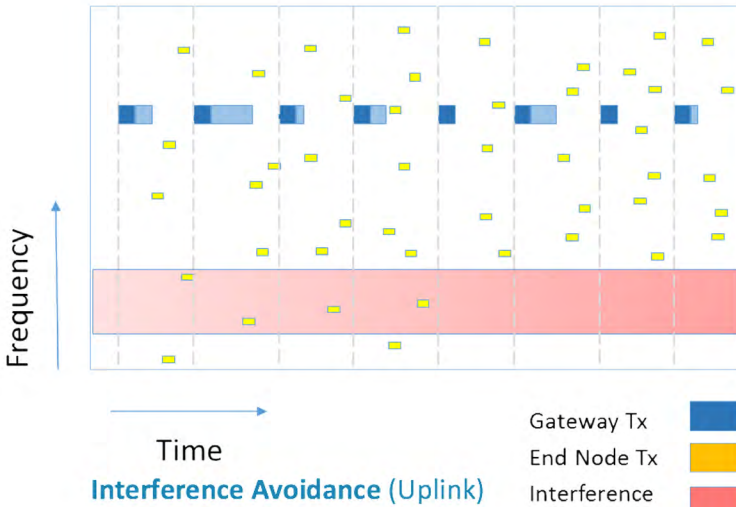
Once the downlink portion of the frame is over, it tunes to a random frequency from the set that the gateway just advertised. Inside the remaining frame are a series of time slots, which each make up 10 bytes of message payload length, plus a Listen before talk period.

The node at this point calculates the power and spreading factor needed to close the link back to the gateway. Let's say the signal is relatively strong, so it chooses spreading factor 7, and 0 dBm transmit power.





This node has a 37 byte message to transmit, so it will need 4 sub-frame slots.



Now, the node's quality of service plays an important role at this point. Assuming the gateway isn't suppressing nodes with that QOS, the QOS is what dictates what percentage of time slots a node can take. A high quality of service, it will pick 4 sub frame slots every time.

At the lowest QOS, a node is only allowed to use one sub frame slot each frame, so it would take 4 frames to get 37 bytes out the door.

But, let's assume the QOS is in the middle, like 8, in the QOS range of 0-15. So it picks 4 slots. It then goes back to sleep,

and wakes back up the frame after next to receive the acknowledgement. If there was a collision with one of the 4 sub-packets, then only 3 would be ACK'd and the node would retransmit the missing sub-packet in that frame.

Since the Symphony Link MTU or Maximum transmit payload length is 256 bytes, it can take up to 26 slots to uplink a large message. The nice thing is that any slots that are missed will be resent automatically by the module. In LoRaWAN, if you want to send more than about 12 bytes, you have to handle retries and packetization at the application layer. Just think about how hard that would be to implement for more than a few nodes.

A repeater in Symphony Link works by fitting a beacon, downlink, uplink, and transfer message all within the uplink portion of a normal Symphony frame. This is because the modulation rate, or spreading factor of repeaters are 2x faster than the gateway. This gives up 3dB of the link, but this is not really noticeable for repeaters, which can add huge coverage to the network. While repeaters have much less capacity than the gateway, a single gateway can host dozens of repeaters. This is a great architecture to cover giant areas cost effectively.

These are just the basics of how Symphony Link works. There are many more features like firmware transfers, multicast, and key exchanges, but the above should give you a good idea of the basic system architecture.

## ABOUT THE AUTHORS

Our team of engineers and scientist mostly hail from the Johns Hopkins Applied Physics Lab where we worked together for many years developing wireless systems for the U.S. Government. We have decades of experience solving problems throughout the stack from ASIC design to satellites.

**Want to learn more about Link Labs?**

Visit our website

[www.Link-Labs.com](http://www.Link-Labs.com)

**LinkLabs**

**For additional questions:  
[info.link-labs.com/contact](http://info.link-labs.com/contact)  
+1 (202) 524-1390**

**Learn about our other products:  
[link-labs.com](http://link-labs.com)**