

Communication Gateways: Friend or Foe?

Steven T. Bushby
National Institute of Standards and Technology

Published in *ASHRAE Journal* Vol.40, No. 4, April, 1998, pp. 50-53

Introduction

Since the middle 1980s the building industry has wrestled with issues related to integrating digital control devices from different manufacturers. Much has happened during that time including the development of BACnet™, an ASHRAE and ANSI standard and a European Community prestandard [1,2], and a host of proprietary integration approaches including LonTalk™ [3]. Gateways play an important role in projects involving interoperability. The strengths and weaknesses of gateways are usually not well understood by building owners and operators. This article explores the role of gateways and proposes criteria that can be used to evaluate whether or not a gateway is appropriate for a particular application.

What is a Gateway?

Computers require very precise, rigidly defined rules or protocol for successful communication. Even slight variations can render communication impossible. In order for two computers using different protocols to communicate, some kind of translation must take place. The device that performs this translation is called a gateway.

Protocol translation, like language translation, is an imperfect art. Concepts that are easily and clearly expressed in one language or protocol may be difficult or even impossible to translate into another. Consider the English word "snow." Native peoples of the arctic have many words in their languages that all translate to snow. All of the detail and nuance that was associated with the original word is lost when it is translated to snow. For the English speaking person this may not be much of a problem because the lost details may not have been useful information. However, when the word "snow" is translated in the reverse direction the problem is more difficult. Much more detail is expected. Which possibility is the best match for the current context? What problems might be caused by choosing the wrong word?

These issues must be addressed when designing gateways. How will the gateway handle differences in the way concepts such as trends, schedules, alarms, and prioritized commands are represented and manipulated? How will missing details be filled in without causing unintended consequences? These difficulties often place severe limitations on the kind of information that can be transferred. The result is often connectivity on a limited scale.

Some control system manufacturers use gateways within their own product line even though they are using their own proprietary protocols. A typical example would be field panel controllers and workstations that communicate using one protocol while the field panels communicate with unitary controllers that they supervise via a second protocol. The field panel in this example is also a gateway. Since the same company designs both protocols and they were intended to be used together, the semantic translation issues are avoided. Typically, translation problems occur only when trying to connect products made by different manufacturers or, in some cases, different versions of products made by the same manufacturer. The example shows that gateways may be stand-alone devices or be included as part of a device that has other control functionality.

Gateways are sometimes confused with routers because gateways also perform routing tasks in addition to translation. The distinction between a gateway and a router is important.

How is a Router Different from a Gateway?

Like a gateway, a router is used to connect two or more communication networks. Routers and gateways filter message traffic because messages pass through only if the source of the message is on one side and the intended destination is on the other. The difference is that for a router, the connected networks all use the same application protocol messages. Translation is not needed. A router's job is to forward the message to the next network in the path to its destination. This makes routers much simpler than gateways.

A router may be used to connect similar or dissimilar networks. If the networks are similar the router merely forwards the message, changing only the address so that it will go to the next stop. For the case of dissimilar networks the situation is slightly more complicated. This case is illustrated in figure 1 for two BACnet local area network (LAN) options. In Figure 1, a message originating with a device on a Master-Slave/Token-Passing (MS/TP) LAN passes through a router in order to reach its destination, a workstation residing on an ISO 8802-3 or "Ethernet" LAN. The content of the message and the binary representation of the message do not change as it passes through the router. However, some other changes do take place when the router forwards the message.

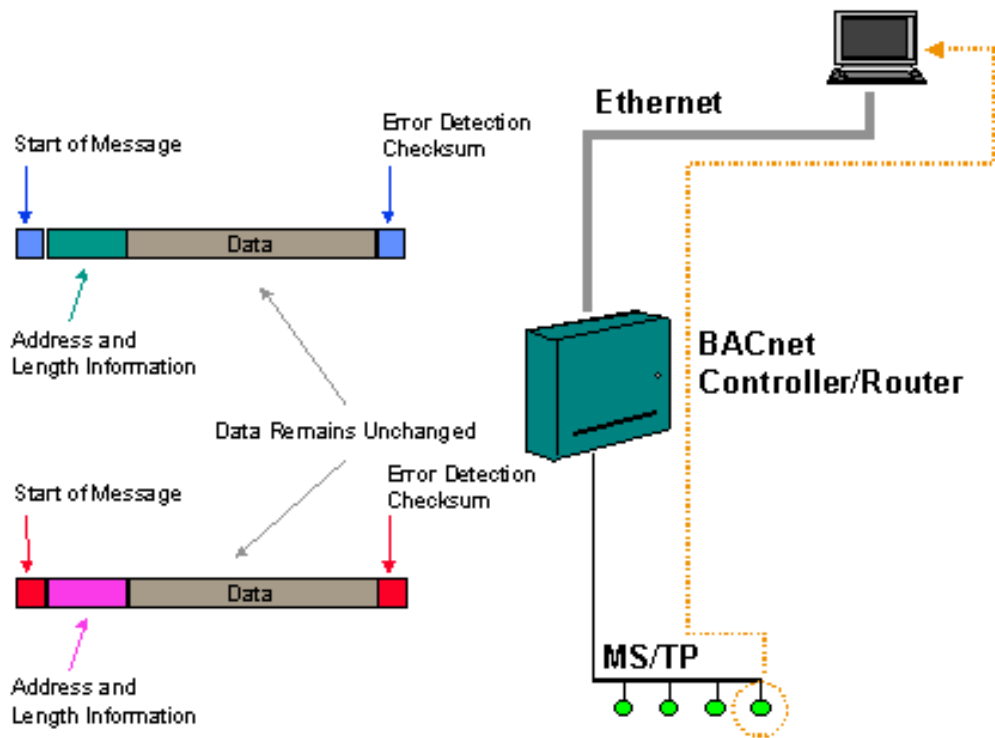


Figure 1: A BACnet router connecting an Ethernet LAN to an MS/TP LAN. The content of the message does not change even though the addressing, signaling, and framing change. In a gateway the data must also be translated.

In this example, the two LAN technologies use different electrical signaling to represent the ones and zeros that make up the message. The address and packet format also change. This is similar to sending a letter through an interoffice mail system or through the postal service. The same letter can be sent either way, although the type of envelope and the necessary addressing information is different. Using this analogy, a router removes the letter from the interoffice envelope, puts it in a postal service envelope, addresses the envelope, and sends it on its way. The content of the letter itself remains unchanged.

A gateway must perform the same addressing, formatting, and signaling changes that a router does. This in addition to translating the message.

How Can I Evaluate the Costs and Benefits of a Gateway?

Gateways can offer very important benefits to a building automation system. That is why they exist. They also have limitations and complicate the design and maintenance of the system. No universal right answer exists to the question: Is it a good idea to use a gateway in my building automation system? Careful consideration of the benefits and costs of the gateway in comparison to other options is needed before making a decision. Table 1 summarizes the benefits and limitations that should be considered when making this decision.

Table 1: Summary of gateway benefits and limitations

Gateway Benefits	Gateway Limitations
<ul style="list-style-type: none"> • provides connectivity that may otherwise be impossible • may expand options for competitive bidding • provides a point of isolation between two largely independent systems • permits interconnection of legacy systems with newer products 	<ul style="list-style-type: none"> • limited capacity and expandability • limited ability to translate dissimilar concepts • configuration and programming of devices through the gateway is generally not possible • failure results in communication loss between all devices on opposite sides of the gateway • possible time delay or the return of cached data that is old • more difficult to troubleshoot problems

The most important step in evaluating the use of gateways is to determine what information needs to be exchanged. This requires some prioritization. What information is critical? What information is important but not essential? What information is desirable but not worth significant additional expense? What kind of future expansion is anticipated? These questions must be answered before a good specification can be written and before a thorough analysis of the available options can be made.

Gateway Benefits

In some situations a gateway can provide connectivity between devices that would not be possible any other way. If this connectivity is vital to the building automation system, this factor can be so critical that it outweighs any negative factors. This is one reason why an initial critical analysis to determine exactly what information needs to be exchanged is so important. Can the gateway provide all of the information needed? If not, is the information that it provides important enough to proceed with the gateway?

In some circumstances, options are available that do not involve a gateway at all. These options may be available from only a single source or a very limited number of sources. In this case it may be beneficial to consider allowing gateways simply to create a competitive bidding situation. The financial savings resulting from competition may make any limitations of the gateway worthwhile. This approach may permit the involvement of a vendor in the process with whom you have a very good business or service relationship. This business relationship may be important enough during the life cycle of the system to make it worthwhile to accept any limitations that the gateway imposes. Similarly, a vendor's product may have some unique feature that is very desirable. It may be beneficial to permit the use of a gateway just to get this feature.

Traditionally, building automation functions such as lighting control, HVAC control, and life safety systems have been isolated from each other. Sometimes hard-wired connections between an HVAC control system and a life

safety system exist that permits monitoring of some limited status information. A growing consensus in the building industry indicates that important benefits may be obtained through more complete integration of these systems. However, complete integration may cause concern because it could possibly have an adverse impact on the ability to guarantee the integrity of critical life safety systems. A gateway provides an opportunity to meet both goals by significantly improving the ability to integrate and manage the systems at the same time retaining some isolation because the gateway limits the interaction. In some cases this may be important for meeting code requirements.

A policy decision might be made to require the use of standard communication protocols, like BACnet, for all future purchases at a facility. This can be a prohibitively expensive change if the entire building automation system must be replaced at one time. Gateways provide a way to retain installed proprietary equipment that is still functioning satisfactorily and to migrate towards standards-based products in a step-by-step manner as equipment reaches the end of its useful life.

These potential benefits can make gateways desirable. It is important to keep in mind that all gateways impose limitations. A building automation system with gateways is more complicated to design and maintain than one where all of the components communicate using the same protocol. There is a cost to using gateways that must be balanced against the potential benefits.

Gateway Limitations

All gateways have finite information storage and processing resources. The gateway must somehow map the information and concepts from one protocol to the other. This requires translation algorithms for each kind of information that needs to be translated and tables that provide mapping information. Sometimes the gateway stores a local copy of the information from one network in a form that makes it easily accessible to devices on the other network. The computer resources necessary to implement this process vary from case to case. In general, both the number of different types of information that are to be translated and the number of occurrences of each type, affect the amount of memory and processor time needed. The more information exchange your application requires, the more probable it is that a gateway will not have the capacity to do the job. Future expandability is usually limited and sometimes impossible.

Specifications for gateways must clearly state what information must be available through the gateway and what expansion capability is required. Any proposed gateway must be evaluated in the context of these requirements. Often, using a gateway will require a compromise between getting all of the desired information and keeping within budget.

As described earlier, gateways have a limited ability to translate dissimilar concepts. The more differences there are between the protocols involved, the more likely it is that the gateway will only be able to translate a portion of the available information. This will probably not be a problem for simple concepts such as temperatures and on/off status. However, problems may arise for more complex concepts such as schedules, alarms, prioritized commands, and trending data. Different approaches to handling these building automation functions make gateway design and implementation very difficult. In some cases one or more of these functions will be simply impossible to do through the gateway.

Many control systems provide a way to program or configure controllers over the network. It is usually not possible to accomplish these tasks by passing messages through a gateway. Instead, a separate connection that bypasses the gateway is usually required.

Gateways can also introduce time delays when attempting to retrieve information. In addition to the processing time required to do the translation, the connected networks may have different transmission speeds and different rules for gaining access to the media that also contribute to the delay. Sometimes gateways are designed only to pass a limited amount of very specific information. The gateway may poll devices for this information and store a local copy. When a request for this information is received the response is prepared by the gateway based on the local copy of the data. This approach can speed up the response but it does not scale easily to large systems.

It can also result in returning old and misleading data. What happens, for example, if the proprietary communication fails temporarily and the stored data cannot be updated? Does the gateway pass the previously stored data that is now obsolete? Will the recipient be able to tell the difference?

Gateways make troubleshooting network problems more difficult. Different tools are needed to see and interpret the protocols on both sides of the gateway. Any ambiguity introduced by the gateway's imperfect translation introduces the possibility of error. Even the limitation on the amount of information accessible through the gateway can make troubleshooting difficult, because it may not be possible to access all of the information needed to diagnose the problem.

It should be clear that, under the right circumstances, a gateway is very useful. However, it comes with a price. A wrench is very useful and effective tool to tighten a nut. It is not a particularly helpful tool to drive a nail, even though with persistence it may be used for that purpose. A gateway is a useful and effective tool to get limited connectivity between otherwise incompatible devices. As the amount and variety of information that needs to be exchanged grows, the viability of using a gateway to accomplish the task diminishes.

BACnet / LonMark Gateways

One example of a gateway application is integrating BACnet and LonMark devices in a single building automation system. This particular gateway application has been given a great deal of attention in the industry because some companies have adopted a marketing policy to provide a two-tiered architecture for their building control products. The idea is to provide unitary control products based on LonMark technology integrated with higher level controllers and workstations based on BACnet. What is this architecture and how does it relate to gateways?

To understand this scenario it is important to understand the different pieces that make up the total picture. LonTalk is a proprietary communication protocol controlled by a private corporation. It is not a *de jure* standard, meaning that no nationally or internationally sanctioned standards body has approved it. Also, there is not a significant installed base to justify a claim of *de facto* standard. The LonTalk protocol is open for anyone to use and several building control companies have publicly indicated an intention to do so. The LonTalk protocol does not define how application specific information is to be represented and exchanged except to provide some general purpose tools. To make LonTalk products interoperable, external constraints that address this issue must be added.

The LonMark Association is a fee-based organization made up of manufacturers using LonTalk technology. The purpose of the organization is to create implementer's agreements, called LonMark profiles. These LonMark profiles provide the external constraints needed to assure interoperability at some level. Devices that conform to the same LonMark profile will interoperate with respect to the functionality defined by that profile.

Since LonMark profiles of the LonTalk protocol are fundamentally different from BACnet, LonMark products cannot be integrated with BACnet products unless a gateway is used. Companies who offer two-tiered architectures based on LonMark and BACnet have a gateway in-between. A gateway must be present or no communication between the two parts of the system can take place.

A BACnet/LonMark gateway is fundamentally no different than a BACnet gateway to any other proprietary system. It has the same strengths and weaknesses. Whether it makes sense for your building can be decided by considering the pros and cons described earlier. It is necessary to weigh the potential gains against the constraints and limitations imposed by the presence of the gateway. The answer will differ for each type of situation.

LonMark has one advantage over other proprietary protocols that might be connected to BACnet by a gateway. An owner can add new LonMark devices to the LonMark portion of the system and realistically expect them to work with the existing LonMark devices as long as they support the same profiles and transmission medium. However, there is a danger from this expansion at the gateway. The gateway was probably designed to pass a specific set of information. Adding new devices to either side requires changes to the gateway. In some cases the gateway will have sufficient resources to handle the additional workload and only reprogramming will be

necessary. Every gateway has a limited capacity. If the new controllers exceed this capacity a new or additional gateway would be required to achieve connectivity.

BACnet provides two standard ways to connect cost-sensitive devices that compete in the marketplace with LonMark proprietary products. One option is the Master-Slave/Token-Passing (MS/TP) protocol that is based on EIA-485 [4] technology. The other option makes use of the LonTalk protocol itself. BACnet permits the use of LonTalk but only as another LAN technology used to convey BACnet messages. Recall that LonTalk products are only interoperable if manufacturers impose external application-specific constraints. BACnet does this by requiring the use of BACnet objects and messages. This approach provides a way for manufacturers to take advantage of some of the desirable features of LonTalk in such a way that these products could integrate with other BACnet products without the need for a gateway. Integrating BACnet products that use LonTalk as the LAN technology with other BACnet products requires only a router

Conclusion

Is a communication gateway a friend or a foe? There is no general answer. In some circumstances, a gateway is the only option for a critical integration and is definitely worthwhile. In other cases, a gateway is acceptable because the advantages and disadvantages roughly balance each other. However, sometimes a gateway is an unnecessary liability and should not be used. The key point is that no two jobs are the same and the costs and benefits must be weighed on a case-by-case basis.

The most important part of the process is to know where to look for the benefits, as well as the trouble, so that an informed decision can be made. First, determine what information needs to be exchanged and prioritize the list. Then, request detailed information about the gateway design from the vendor so that the potential benefits and weaknesses described in this article can be evaluated.

Disclaimer

Certain trade names are mentioned or identified in an illustration in order to specify adequately their relationship to the BACnet standard. In no case does such identification imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that the products are necessarily the best available for the purpose.

References

1. ANSI/ASHRAE Standard 135-1995: *BACnet - A Data Communication Protocol for Building Automation and Control Networks*, Atlanta Georgia: American Society of Heating Refrigerating, and Air-Conditioning Engineers. 1995.
2. CEN ENV 1805-1 *Data Communication for HVAC Application, Management Net, Part 1. BACnet*. European Committee for Standardisation, Central Secretariat, Rue de Stassart 36 B-1000, Brussels, Belgium.
3. Echelon, *LonTalk Protocol Specification*, Echelon Corporation, 4015 Miranda Ave., Palo Alto, CA 94304, USA.
4. EIA-485 (1993), *Standard for Electrical Characteristics of Generators and Receivers for Use in Balanced Digital Multipoint Systems*. Electronics Industries Association, 200 Eye St., N. W., Washington, D. C. 20006.