Allen-Bradley

EtherNet/IP Performance

Application Solution

Rockwell Automation

Important User Information

Solid state equipment has operational characteristics differing from those of electromechanical equipment. *Safety Guidelines for the Application, Installation and Maintenance of Solid State Controls* (Publication SGI-1.1 available from your local Rockwell Automation sales office or online at http://www.ab.com/manuals/gi) describes some important differences between solid state equipment and hard-wired electromechanical devices. Because of this difference, and also because of the wide variety of uses for solid state equipment, all persons responsible for applying this equipment must satisfy themselves that each intended application of this equipment is acceptable.

In no event will Rockwell Automation, Inc. be responsible or liable for indirect or consequential damages resulting from the use or application of this equipment.

The examples and diagrams in this manual are included solely for illustrative purposes. Because of the many variables and requirements associated with any particular installation, Rockwell Automation, Inc. cannot assume responsibility or liability for actual use based on the examples and diagrams.

No patent liability is assumed by Rockwell Automation, Inc. with respect to use of information, circuits, equipment, or software described in this manual.

Throughout this manual we use notes to make you aware of safety considerations.

WARNING — Identifies information about practices or circumstances that can cause an explosion in a hazardous environment, which may lead to personal injury or death, property damage, or economic loss.

IMPORTANT — Identifies information that is critical for successful application and understanding of the product.

ATTENTION — Identifies information about practices or circumstances that can lead to personal injury or death, property damage, or economic loss. Attentions help you:

- identify a hazard
- avoid a hazard
- recognize the consequence

SHOCK HAZARD — Labels may be located on or inside the drive to alert people that dangerous voltage may be present.

BURN HAZARD — Labels may be located on or inside the drive to alert people that surfaces may be dangerous temperatures.

Summary of Changes

This document help you plan your EtherNet/IP network and describes considerations for improving overall network performance. Revision bars in the margin identify updated information. Changes for this version of the document include:

- information on selecting an Ethernet switch (see chapter 2)
- addition of information regarding multicast frames (see page 3-9)
- addition of information for the 1756-EWEB web server module (see chapter 3)
- addition of information for the 1769-L32E and 1769-L35E CompactLogix controllers (see chapter 3)
- addition of information for the 1734-AENT POINT I/O adapter (see chapter 3)
- additional information on how the RPI translates into the actual packet interval (see page 4-5)

Notes:

Table of Contents

Chapter 4

Predict System Performance

Appendix A

Worksheets

EtherNet/IP Overview

Introduction

This document provides an application example to help estimate the performance of an EtherNet/IP control system. This chapter provides a brief overview of an EtherNet/IP network and provides a list of reference materials that you can use to find detailed information about planning and configuring a network.

EtherNet/IP Overview

EtherNet/IP is a network suitable for use in industrial environment and time-critical applications. EtherNet/IP uses standard Ethernet and TCP/IP technologies and an open application layer protocol called the Control and Information Protocol (CIP). CIP is also the application layer used in DeviceNet and ControlNet networks. The open Application Layer protocol makes interoperability and interchangeability of industrial automation and control devices on EtherNet/IP a reality for automation and control applications.

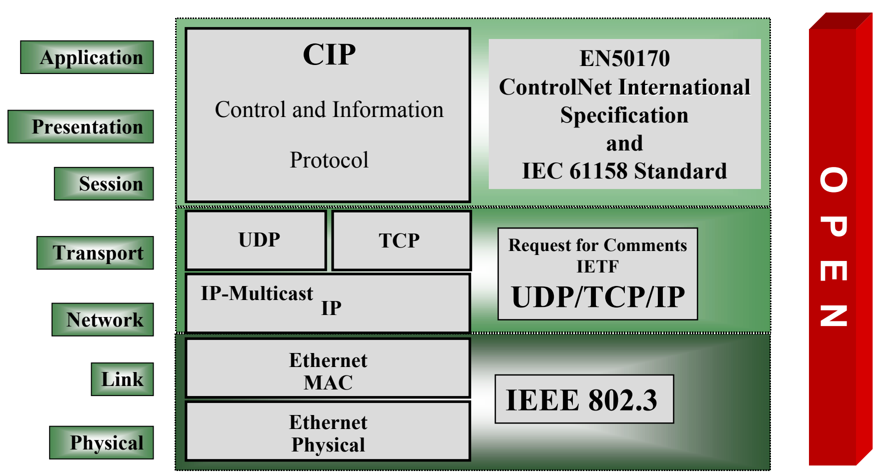

EtherNet/IP supports both time-critical (implicit) and non time-critical (explicit) message transfer services of CIP. Exchange of time-critical messages is based on the producer/consumer model where a transmitting device produces data on the network and many receiving devices can consume this data simultaneously.

EtherNet/IP supports these functions:

- Time-critical message exchange (for I/O control)
- Human Machine Interface (HMI)
- Device configuration and programming
- Device and network prognostics and diagnostics
- Compatibility with SNMP and web pages embedded in devices

Support of these functions, along with interoperability and interchangeability, positions EtherNet/IP as an open network-standard for Ethernet-based, industrial automation.

Planning an EtherNet/IP Network

When planning an EtherNet/IP network, follow these steps:

Step:	Issue to Decide:
Generate functional requirements for the network	• EtherNet/IP technology • Physical medium • Topology • Mode of operation (half-duplex or full-duplex) • Web access
Determine network infrastructure	• Network configuration/hierarchy – separate business and control traffic – use Ethernet switches instead of hubs • Layer 2 switch selection – required features include 10/100Mbps, full-duplex, IGMP snooping, and port mirroring • Layer 3 switch selection (optional) • Media selection (UTP, STP, fiber) • Any other infrastructure components • Be aware of security issues
Assign node IP addresses	See the user manual for your EtherNet/IP interface to determine the most appropriate method for assigning an IP address
Determine data requirements (implicit and explicit messages)	• Verify device capacity against load • TCP connections • CIP connections
Estimate system performance	• Determine if the system has sufficient bandwidth (packets/second) • Predict the maximum input and output times • Calculate worst-case I/O response • Estimate HMI traffic

Additional Reference Materials

The following documents provide detailed information about planning and configuring an EtherNet/IP network.

These documents:	Are available as:
A11259018 - Switch Considerations for Use with EtherNet/IP	KnowledgeBase articles See: http://support.rockwellautomation.com Select the Knowledgebase icon from the Customer Support page
EtherNet/IP standards documents	ODVA web site See: http://www.odva.org
ENET-IN001 *EtherNet/IP Media Planning and Installation Guide* ENET-UM001 *EtherNet/IP Modules in Logix5000 Control Systems User Manual*	Rockwell Automation manuals are available in the Automation Bookstore See: http://www.theautomationbookstore.com

Notes:

Select an Ethernet Switch

Introduction

Switch Functionality for EtherNet/IP Systems

When using EtherNet/IP for time-critical (implicit) messaging for I/O control, there are several features available in switches that are required and/or recommended.

Required or Recommended:	Switch Feature:
required	• full-duplex capability on all ports • IGMP snooping • port mirroring
recommended	• VLAN • autonegotiation and manually configurable speed/duplex • wire-speed switching fabric • SNMP

Full-Duplex Capability on all Ports

Full-duplex capability eliminates collisions. Combined with the speed of the switches available today, you can eliminate the delays related to collisions or traffic in the switch. The end result is that EtherNet/IP becomes a highly deterministic network that works well for I/O control.

If a device is forced to a specific duplex, (auto-negotiation is disabled), then the switch defaults to half-duplex. It is best to let the device auto-negotiate.

IGMP Snooping

Much of EtherNet/IP implicit (I/O) messaging uses IP multicast to distribute I/O control data, which is consistent with the CIP produced/consumer model. Historically, most switches have treated multicast packets the same as broadcast packets. That is, multicast packets are re-transmitted to all ports.

IGMP snooping constrains the flooding of multicast traffic by dynamically configuring switch ports so that multicast traffic is forwarded only to ports associated with a particular IP multicast group.

Switches that support IGMP snooping "learn" which ports have devices that are part of a particular multicast group and only forward the multicast packets to the ports that are part of the multicast group. \

Be careful as to what level of support of IGMP snooping a switch has. Some layer 2 switches that support IGMP snooping require a router (which could be a layer 3 switch) to send out IGMP polls in order to learn what devices are part of the multicast group. Some layer 2 switches can use IGMP snooping without a router sending polls. If your control system is a stand-alone network or is required to continue performing if the router is out of service, make sure the switch you are using supports IGMP snooping without a router present

Port Mirroring

Port mirroring is for troubleshooting. Port mirroring refers to the ability to direct the frames being transmitted on one port to another port. This lets a traffic analyzer to connected to a switch have the ability to monitor the traffic on a given port. Without port mirroring, an analyzer cannot see frames on other ports. With hubs this is not an issue, because all frames are transmitted out all ports.

Select a switch that supports port mirroring so that a traffic analyzer can function correctly on the network. Traffic analyzers are critical to supporting and maintaining Ethernet networks.

VLAN (Virtual Local Area Network)

With VLANs, you can configure a switch to share two isolated networks without the traffic from one network burdening the other. IP multicast traffic from VLAN 1 will not reach VLAN 2. A VLAN blocks broadcast traffic and adds a measure of security between networks.

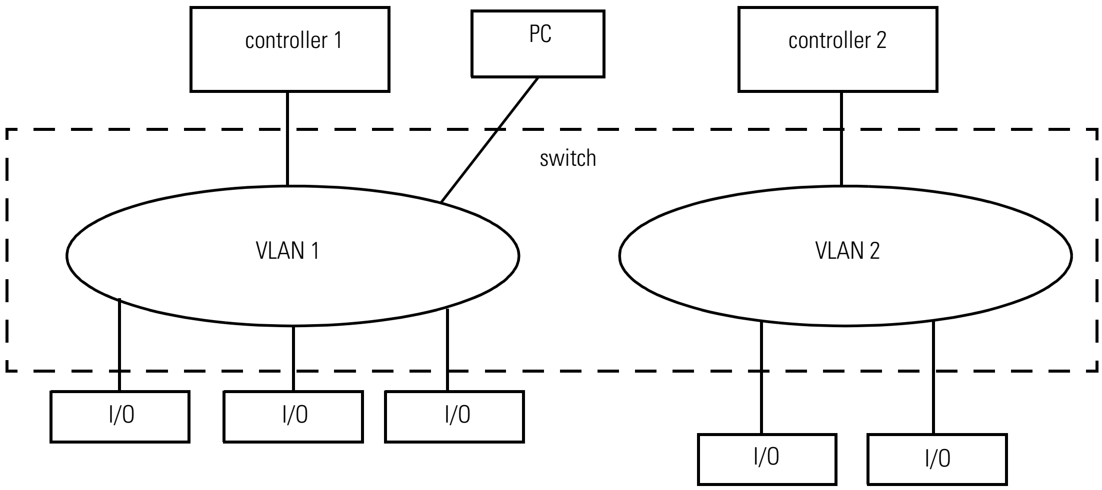

Autonegotiation

Autonegotiation lets devices select the most optimal way to communicate without the you having to configure the devices. However, if you connect a manually-configured to an autonegotiation device, there can be problems which result in a high rate of CRC errors.

All 100 Mbps devices are required to support autonegotiation, but most existing 10 Mbps devices do not. Select a switch that supports both speeds so you can connect to existing devices such as 1756-ENET modules.

Wire-Speed Switching Fabric

The switch fabric is a measure of the maximum traffic that a switch can handle without dropping a packet and without storing a packet in memory. Wire-speed switching fabric refers to a switch that can handle the maximum data rate of the network on each of its ports.

Switches are typically rated in Gbps. For a 10 port switch connected to EtherNet/IP products, the maximum data rate needed is typically in the 100-200 of Mbps range. Therefore, a 10 port switch rated at least 1 Gbps should be adequate for an EtherNet/IP implementation.

SNMP (Simple Network Management Protocol)

SNMP is a TCP/IP protocol for obtaining statistical information about a device. SNMP software lets a network manager view and modify a wide variety of network parameters and provides a common way to manage many diverse vendor products.

Notes:

How Connections and Multicast Frames Affect Data Transfer

Introduction

Connection Overview

A connection is a point-to-point communication mechanism used to transfer data between a transmitter and a receiver. Connections can be logical or physical. An EtherNet/IP CIP connection transfers data from an application running on one end-node to an application running on another end-node. A CIP connection is established over a TCP connection. A single TCP connection can support multiple CIP connections.

Example 1: I/O Connections
A Logix controller has 5 CIP I/O connections to modules in remote chassis and all of these connections are through the same local 1756-ENBT and the same remote 1756-ENBT. The following connections would exist:

- 1 TCP connection
- 5 CIP connections

Example 2: RSLinx OPC Test Client

- 1 TCP connection
- 4 CIP connections (4 is the default)

TCP connections are used for all communications on EtherNet/IP. Even with implicit connections, a TCP connection is used and remains open. A TCP connection is required for all CIP messaging, including connected and unconnected messaging.

CIP connections are almost always used. Unconnected CIP messages are used but are temporary (short lived).

Terminology

Term:	Definition:
implicit connection	An implicit connections are time critical in nature. This includes I/O and produced/consumed tags. Implicit refers to information (source address, data type, destination address, etc.) which is implied in the message but not contained in the message.
explicit connection	Explicit connections are non-time critical and are request/reply in nature. Executing a MSG instruction or executing a program upload are examples of explicit connections. Explicit refers to basic information (source address, data type, destination address, etc.) that is included in every message.
producer and consumer	Producer/consumer refers to implicit connections. With implicit connections, messages are sent cyclically (every RPI). **Example:** Assume a ControlLogix controller is controlling a single rack of Flex I/O using a rack connection. Both the ENBT module that is local to the controller and the Flex AENT module are consumers and producers of data. The AENT consumes outputs and produces inputs.
client and server	Client/server refers to explicit connections. A client creates a connection and initiates messages. A server provides a service or data. Clients can send messages continuously or intermittently. **Example:** A ControlLogix controller can send a MSG instruction to another controller.
transports	Each connection has transports. A transport is a uni-direction entity with its own numeric identifier. A implicit connection has 2 transports. A explicit connection has 1 transport. Transports are important because they help you calculate the number of packets per second for each Ethernet interface. **Example 1: I/O** For an I/O connection to a rack of Flex I/O, a connection is configured in RSLogix 5000 software by adding the Flex adapter and I/O modules in the I/O list. When the connection is created, output packets flow from the controller to the I/O rack. In addition, input packets flow from the I/O to the controller. Each direction of flow is a transport. In this example, two transports exist. One transport is from the controller to the adapter. The second transport is from the adapter to the controller. **Example 2: Produced Tag** For a produced tag connection with 2 consumers, there is a connection to each consumer. Data from the producer is produced to the wire on one transport. Each of the consumers returns a heartbeat. A total of 3 transports exist in this example. One transport is from the tag producing controller to the "wire" media. The second transport is from one consumer to the tag producer. The third transport is from the second consumer.
UCMM	In the web servers, you can see references to Unconnected Message Manager (UCMM). This type of messaging is momentary and therefore can be ignored unless you are troubleshooting. Examples of where UCMM messages are used are: • Flash update of module firmware • Some functions in RSLinx • CIP Generic MSG instruction • Opening any CIP connection (forward_open command)

TCP Connections

TCP connections are required for EtherNet/IP communications. TCP connections are used for unconnected CIP messages and for CIP explicit connections. Examples of TCP connections are:

- HMI (human-machine interface) to a controller that supports EtherNet/IP communications
- Logix MSG instruction to a controller or workstation
- OPC or DDE accessing a controller

TCP connection limits

Product:	TCP Connection Limits:
1756-ENBT 1756-ENET/B 1756-EWEB 1769-L32E 1769-L35E 1788-ENBT 1794-AENT	64 TCP connections These modules also have web servers which use TCP connections for non-CIP traffic (HTTP). These TCP connections do not count toward the 64-connection maximum
1734-AENT	as many TCP connections as memory is available

CIP Connections

CIP connections are required for both implicit and explicit messaging. Examples of functions supported by CIP connected messaging include:

- Logix controller message transfer to Logix controller
- I/O or produced tag
- Program upload
- RSLinx DDE/OPC client
- PanelView polling of Logix controller

There are different types of CIP connections:

CIP connection type:	Description:
bridged connection	A bridged connection is a connection that passes through the EtherNet/IP module. The end point of the connection is a module other than the EtherNet/IP module. **Example:** an explicit connection from a controller through a 1756-ENBT to another controller.
end-node connection	An end-node connection is a connection whose end point is the EtherNet/IP module itself. **Example:** an explicit connection from RSLinx to the EtherNet/IP module to set the module's IP address.
rack-optimized	A rack-optimized connections is an implicit message connection to a rack or assembly object in the EtherNet/IP module. Data from selected I/O modules is collected and produced on one connection (the rack-optimized connection) rather than on a separate direct connection for each module.
direct	An implicit message connection from a controller to an specific I/O module (as opposed to a rack-optimized connection).

CIP connected messaging limits

<table>
<tr><th>Product:</th><th>CIP Connected Messaging Limits:</th></tr>
<tr><td>1756-ENBT</td><td>Each module has a maximum of 128 CIP connections, of which:
Maximum of 128 bridged connections (any combination of implicit and explicit connections).
Maximum of 32 end-node connections.
In addition to the CIP connections, the 1756-ENBT supports:
Maximum of 16 controllers that can have a rack-optimized connection to the module.
Maximum of 16 controllers that can have a rack-optimized, listen-only connection to the module.
Maximum of 64 controllers can consume data from an implicit connection.</td></tr>
<tr><td>1756-ENET/B</td><td>Each module has a maximum of 160 CIP connections, of which:
Maximum of 128 bridged connections (any combination of implicit and explicit connections).
Maximum of 32 end-node connections.
In addition to the CIP connections, the 1756-ENET/B supports:
Maximum of 16 controllers that can have a rack-optimized connection to the module.
Maximum of 16 controllers that can have a rack-optimized, listen-only connection to the module.
Maximum of 32 controllers can consume data from an implicit connection.</td></tr>
<tr><td>1756-EWEB</td><td>Each module has a maximum of 128 CIP connections, of which:
Maximum of 128 bridged connections (explicit connections).
Maximum of 32 end-node connections.</td></tr>
<tr><td>1769-L32E
1769-L35E</td><td>Each module has a maximum of 32 CIP connections, of which:
Maximum of 32 bridged connections (any combination of implicit and explicit connections).
Maximum of 20 end-node connections.
In addition to the CIP connections, the 1769-L32E, -L35E supports:
Maximum of 32 controllers can consume data from an implicit connection.</td></tr>
</table>

Product:	CIP Connected Messaging Limits:
1788-ENBT	Each module has a maximum of 32 CIP connections, of which: • Maximum of 32 bridged connections (any combination of implicit and explicit connections). • Maximum of 20 end-node connections. In addition to the CIP connections, the 1788-ENBT supports: • Maximum of 32 controllers can consume data from an implicit connection.
1734-AENT	Each module has a maximum of 32 CIP connections, of which: • There are 0 bridged connections. • Maximum of 5 controllers with any combination of rack-optimized connections and listen-only, rack-optimized connections. • Maximum of 32 explicit end-node connections. • Maximum of 20 implicit end-node connections (this includes rack-optimized and listen-only rack-optimized connections).
1794-AENT	Each module has a maximum of 32 CIP connections, of which: • There are 0 bridged connections. • Maximum of 32 explicit end-node connections. • Maximum of 31 implicit end-node connections. In addition to the CIP connections, the 1794-AENT supports: • Maximum of 31 controllers can consume data from a direct connection.

CIP unconnected messaging limits

The following limits of unconnected messages are the maximum number of outstanding unconnected messages. These are unconnected messages that have been sent to the device and are being processed and have not yet generated a response or timeout.

Product:	CIP Unconnected Messaging Limits:
1756-ENBT 1756-ENET/B 1756-EWEB	Each module has a maximum of 256 CIP unconnected messages, of which: • Maximum of 128 unconnected messages from the EtherNet/IP port to an object on the module or to the backplane. • Maximum of 128 unconnected messages from the backplane to an object on the module or to the EtherNet/IP port.
1769-L32E 1769-L35E	Each module has a maximum of 64 CIP unconnected messages, of which: • Maximum of 32 unconnected messages from the EtherNet/IP port to the host. • Maximum of 32 unconnected messages from the host to the EtherNet/IP port.
1788-ENBT	Each module has a maximum of 64 CIP unconnected messages, of which: • Maximum of 32 unconnected messages from the EtherNet/IP port to the host. • Maximum of 32 unconnected messages from the host to the EtherNet/IP port.
1734-AENT	Each module has a: • Maximum of 20 simultaneous, unconnected messages from the EtherNet/IP port to the backplane. • Maximum of 32 simultaneous, unconnected messages from the EtherNet/IP port to an object on the module.
1794-AENT	Each module has a maximum of 256 CIP unconnected messages from the EtherNet/IP port. The 1794-AENT can receive messages from the EtherNet/IP port. Because the FLEX I/O backplane uses a polled architecture, the FLEX modules I/O modules do not initiate messages to the 1794-AENT module.

TCP and CIP Timeouts

TCP/IP Stack timeout (not user configurable)

Approximate Timeout:		Description:
keep-alive	2 minutes	Source and target stacks periodically talk to verify connectivity
inactivity	5 minutes	This timer continues to run if CIP connected or unconnected messages are not occurring. When this timer expires, the TCP stack will: • tear down a connection • not attempt to reestablish a connection (the controller must do this)
ARP	variable	The Address Resolution Protocol (ARP) timer is used when a device needs to send an IP message to another device. ARP allows the sending device to translate a device's IP address to the corresponding MAC ID so that the message can be sent over the EtherNet/IP network. ARP is used in all TCP/IP implementations (even your PC). An ARP cache temporarily stores mappings of IP addresses to MAC IDs. For example: • MAC ID (e.g. 0000bc060102) • IP address (e.g. 130.151.139.121) If you replace a Rockwell Automation EtherNet/IP module with a new module, the new module will have a different MAC ID. The ARP cache entries in other devices are now invalid because the MAC ID corresponding to the module's IP address has changed. This could cause a delay in reestablishing communications with the replacement module. The delay varies depending on the module and the network configuration in use. When a Rockwell Automation EtherNet/IP starts up, it issues a "gratuitous ARP" which causes other devices to update their ARP caches. This generally results in a quick recovery of communications with the replacement module (less than a minute). However, some switches will not forward the gratuitous ARP message onto the network, in particular, if the Spanning Tree Protocol is enabled on that port. It is recommended that you disable the Spanning Tree Protocol on those ports to which EtherNet/IP modules are directly connected (but not on ports which are linked to other switches). In the worst case, if the gratuitous ARP is not seen, an originating device could wait as long as 10 minutes for the ARP cache entry to age out and be deleted.

CIP timeouts

Approximate Timeout:		Description:
implicit message	network multiplier x RPI	The multiplier is selected by the controller firmware so that the timeout is greater than or equal to 100 ms. The minimum multiplier is 4. Example 1: RPI = 2ms, the controller-selected multiplier = 64. The timeout is 128ms. Example 2: RPI = 10ms, the controller-selected multiplier = 16. The timeout is 160ms.
explicit message	30 seconds	For explicit messages, connected or unconnected, the timeout is 30 seconds. This is user-changeable in the Message (MSG) instruction structure.

Multicast Frames

All input data from I/O devices is sent multicast. Therefore, each frame is broadcast throughout the system to make sure it reaches all the possible devices in the multicast group.

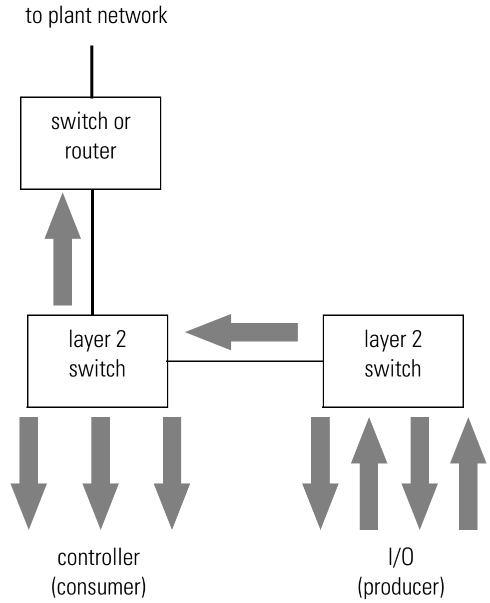

I/O devices generally produce at very fast rates (such as 10 ms), so it is easy to flood the network with multicast traffic and force each end device to spend time deciding whether to discard numerous multicast frames. If there are a lot of I/O devices, they can easily use up a significant part of a router's CPU time.

You must consider control network traffic propagating onto the plant information network, as well as, plant information network traffic propagating onto the control network. Some best practices to follow are:

- Minimize device load due to unwanted IP multicast traffic
- Minimize switch load due to unwanted IP multicast traffic
- Minimize network load due to unwanted incoming IP multicast or broadcast traffic
- Block IP multicast traffic generated within the Ethernet/IP subnet from propagating onto the plant network
- Implement standard network troubleshooting tools

Multicast address limit

Implicit connections that produce data over an EtherNet/IP network use multicast addresses. EtherNet/IP interfaces support a maximum of 32 unique multicast addresses. The actual address (such as 239.192.22.121) is determined by the EtherNet/IP interface.

Example 1: An ethernet adapter that produces data uses a unique multicast address for each I/O connection.

Example 2: A Logix controller that produces tags uses a unique multicast address for each produced tag.

The multicast address limit is independent of the connection limit for a device. Not all connections require a multicast address. And in the case of produced and consumed tags, one produced tag requires one multicast address but it also requires one connection for each consumer. If there are multiple consumers, the one multicast address would be using multiple connections.

Predict System Performance

Introduction

This chapter describes how to predict the performance of your EtherNet/IP-based control system and how to enhance that performance

System Prediction Goals

You allocate the bandwidth of your EtherNet/IP communication module between two types of messaging:

Messaging Type:	Description:
explicit messaging	Explicit messages are connections that do not use an RPI. Explicit messaging includes: MSG, PanelView, RSView, RSLogix 5000 uploads and downloads, etc.
implicit messaging	Implicit messages are connections that use an RPI. Implicit messaging is used for I/O. This includes rack optimized connections, direct connections, and messages using produced tags.

The performance predictions have two major goals:

- To determine if the system as a whole has sufficient bandwidth to fulfill the requirements of the application.
- To estimate the maximum input or output times for rack optimized connections, direct connections, and produced/consumed tags.

Part I: Determine if the system has sufficient bandwidth to fulfill the requirements of the application

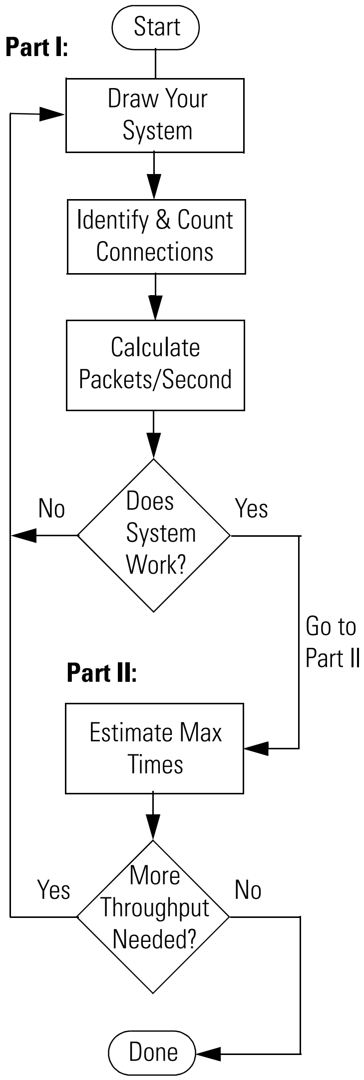

To determine if your system has sufficient bandwidth to fulfill the requirements of the application, perform the following steps:

1. Draw an overall sketch of your system that includes all processors, EtherNet/IP modules, and I/O modules, and shows all connections to the network. Include a description of what the processors are doing (e.g., messaging using produced tags), and any known RPI requirements.

2. Identify and count each type of implicit connection for the system and each EtherNet/IP module.

3. Use the formulas shown later in this chapter to calculate the packets/second loading on each EtherNet/IP module and the available bandwidth for any unspecified RPIs.

4. Based on the results of these calculations, decide if your system will work.

5. If necessary, modify your system by doing one or more of the following:
   - Increase some RPIs to allow other RPIs in the system to decrease.
   - Change connection types (e.g., direct to rack optimized).
   - Change I/O module configurations (filter times, trigger types).
   - Add EtherNet/IP modules.
   - Add Logix controllers.
   - Verify that the network infrastructure can handle the system traffic.
   - Verify that the switches support full-duplex operation and IGMP snooping. Port-mirroring is also important for switch and system diagnostic functions.

6. If you have made modifications, verify that the modified system will work by re-counting the connections and re-calculating the packets/second loading.

Part II: Predict the maximum input or output times for CIP connections

A CIP connection is an efficient communications path used for high performance. Basic types are: rack, data, produce tag. Also, a CIP connection is bi-directional which means that during every RPI interval, a packet of information is initiated from both ends of each connection. The type of data packet produced by each end of the connection depends on the connection type and is detailed below:

Connection Type:	Scanner:	Adapter:	Producer Controller:	Consumer Controller:
rack	output data	input data	n/a	n/a
input data	heartbeat	input data	n/a	n/a
output data	output data	output data echo	n/a	n/a
produce tag	n/a	n/a	tag data	heartbeat

To predict the maximum input (I/O to controller) or output (controller to I/O) times for CIP connections, continue with the following steps:

7. Estimate the maximum time intervals for:
   - rack optimized connections
   - direct connections
   - produced tags

8. If necessary, modify your system to get more throughput by performing one or more of the adjustments described under step 5.

9. If you have made further modifications, verify that the modified system will work.

Performance Calculations

The performance predictions involve three sets of simple calculations:

1. Identifying and counting the number of connections

2. Calculating the packets/second loading

3. Estimating the maximum input or output times

Identify and count connections

Use your design to identify and count the total number of rack optimized connections, direct connections, and produced/consumed tag connections for each EtherNet/IP module in your system. Performance prediction is done on a CIP connection basis.

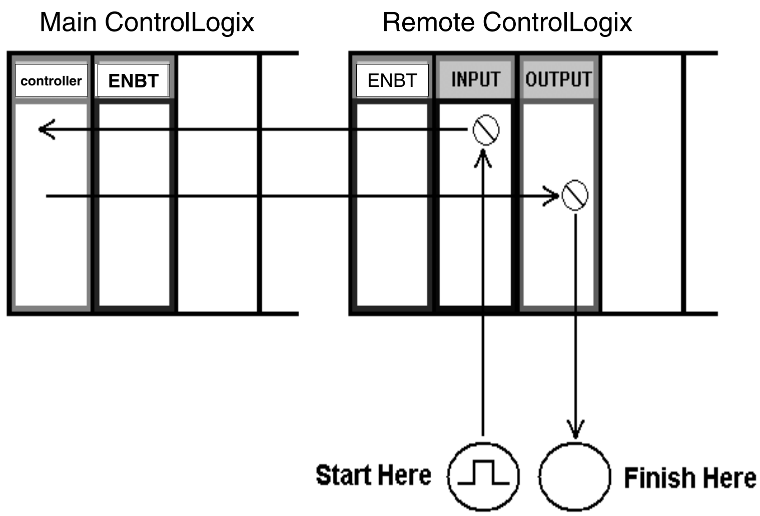

Calculate packets/second

Each EtherNet/IP module has a maximum number of packets/second. Bandwidth should be allocated as follows:

- Reserve 10% of each EtherNet/IP module's bandwidth to allow for processing of explicit messages.
- The total for implicit messaging should not exceed 90% of capacity for each EtherNet/IP module.

IMPORTANT

If you do not reserve at least 10% of each EtherNet/IP module's bandwidth, you might not be able to go online with RSLogix 5000 software or be able to access the EtherNet/IP module's embedded web server.

To remedy this situation, remove AC power on one or more racks to reduce the EtherNet/IP traffic in the overloaded EtherNet/IP modules. Then go online with RSLogix 5000 software to reconfigure the RPIs to a less frequent (slower) rate.

Each CIP connection is bi-directional, and therefore requires a minimum of 2 packets per RPI. Using 2 packets/RPI/connection, the number of packets/second to or from each EtherNet/IP module can be calculated as follows:

A. Rack Optimized: Packets/Second = (2 x connections)/RPI

B. Direct Connect: Packets/Second = (2 x connections)/RPI

C. Produced Tag (producer and all consumers are in different chassis and operating at a uniform RPI):

At Producer: Packets/Second = (1 + connections)/RPI for each produced tag
At Consumer: Packets/Second = 2/RPI for each consumed tag

For each EtherNet/IP module, the total packets/second is the sum of the above. This total should not exceed the recommended 90% packets/second limit.

How the RPI you configure translates into the actual packet interval (API) depnds on the controller initiating the communication. Both RPI and and API are in milliseconds. In general:

- ControlLogix: API = RPI
- CompactLogix: API = 2^n
- FlexLogix: API = 2^n
- SoftLogix: API = RPI

where 2^n is a value that is a power of 2, such as 2, 4, 8, 16, etc., that is equal to or faster than the RPI you configured.

So, in most cases, you can get data faster than the RPI you configured, which can increase the number of packets/second to be more than you expected based on the RPI.

Estimate the fastest RPI

The fastest RPI for an EtherNet/IP interface is:

$RPI_{(Fastest)}$ = (2 x connections)/pps

IMPORTANT It is not necessary to operate at the fastest RPI.

Example 1: The 1756-ENBT supports 5000 pps. If there are only four connections that are all at the same RPI, the fastest RPI is:

$RPI_{(Fastest)}$ = (2 x 4)/5000 = 1.6 ms

Example 2: Assume there are three connections already running at an RPI of 2 ms on a 1756-ENBT module. These connections are already using some of the communication packets:

pps = (2x3)/2 ms = 3000 pps

The fastest RPI possible for a fourth connection is:

$RPI_{(Fastest)}$ for 4^{th} connection = (2 x 1)/(5000-3000) = 1 ms

The same concept can be used for produced tags and consumed tags by replacing the (2 x connections) with (1 + connections) for produced tags or (2/RPI) for consumed tags.

Estimate maximum input or output times for CIP connections

System response is dependent on several factors. The dominant factors are RPI and the number of implicit CIP connections. To simplify, the response time of a connection can be approximated with only the RPI.

The maximum input (I/O to controller) or output (controller to I/O) times for implicit CIP connections can be estimated as follows. With this approximation, the error will be less than 10% if the RPI (in milliseconds) is at least 10 times the number of connections through the EtherNet/IP interface.

A. Rack Optimized: 1 RPI

B. Direct Connect:

Discrete: 1 RPI

Analog (non-isolated): 2 RTS

Analog (isolated): 1 RTS

C. Produced/Consumed Tag: 1 RPI

The above response times are estimates. For more accurate numbers, include system delays. See page 4-16 for more information.

Example: Predict System Performance

This example system has the following components:

- A controller that controls the I/O and produces a tag that is consumed by two other processors at an RPI of 20 ms.
- Twenty-five digital I/O modules using rack optimized connections with specified RPIs of 20 ms.
- Fifteen analog I/O modules using direct connections with specified RPIs of 50 ms.
- One PanelView with a direct connection to an array of 40 tags at 100 ms RPI and an explicit connection to an array of 100 tags sent every 300 ms.
- One workstation running RSView32, requiring explicit messaging of an array of 1500 tags every 100 ms. This workstation also connects to the company Intranet.

Part I: Determine if the system has sufficient bandwidth to fulfill the requirements of the application

Based on the system requirements, the initial network diagram is shown below:

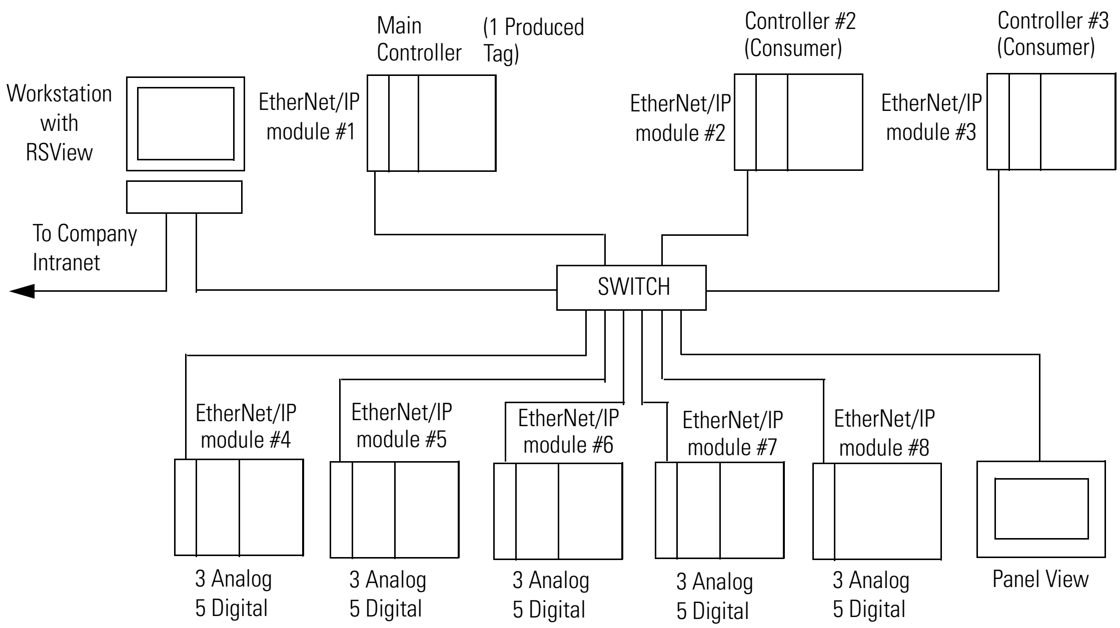

Explicit messaging

The RSView messages and the explicit messages from the PanelView are explicit messaging. Reserve 10% of the bandwidth of the EtherNet/IP module for explicit messaging:

EtherNet/IP Interface:	Total Bandwidth (packets/second):	10% Reserve for Explicit Messaging (packets/second):
1756-ENBT	5000 pps	500 pps
1756-ENET/B	900 pps	90 pps
1769-L32E 1769-L35E	4000 pps	400 pps
1788-ENBT	5000 pps	500 pps
1734-AENT	5000 pps	500 pps
1794-AENT	9500 pps	950 pps

Explicit messaging throughput is also dependent upon network availability and target availability. Therefore, reserving 10% of the total bandwidth does not guarantee throughput.

Next determine if each EtherNet/IP module has enough bandwidth to handle the implicit messaging.

EtherNet/IP modules serving as adapters

Each EtherNet/IP module serving as an adapter in an I/O rack (EtherNet/IP modules 4 to 8 in the example on page 4-8) has one rack optimized connection for digital I/O and three direct connections for analog I/O. The digital I/O has a required RPI of 20 ms, and the analog I/O has a required RPI of 50 ms.

For example, use the packets/second worksheet in Appendix A for each of the five EtherNet/IP modules as follows:

	EtherNet/IP Module ID: 4 - 8	
Rack Optimized Connections 1 @ 20 ms		
Packets/Second	= (2 x connections)/RPI	
	= (2 x 1)/ 20 ms =	100
Direct Connections 3 @ 50 ms		
Packets/Second	= (2 x connections)/RPI	
	= (2 x 3)/ 50 ms =	120
Produced Tag Connections 0		
Packets/Second	= (1 + connections)/RPI for each produced tag	
	=	N/A
Consumed Tags 0		
Packets/Second	= 2/RPI for each consumed tag	
	=	N/A
Total Packets/Second		220

The total of 220 packets/second is well within the remaining bandwidth for these modules (4500 pps limit for a 1756-ENBT module; 810 pps limit for a 1756-ENET/B module).

EtherNet/IP modules 2 and 3 with consumed tags

EtherNet/IP modules 2 and 3 interface two consumer controllers to the network. Each of these controllers consumes one produced tag at an RPI of 20 ms. For either of these EtherNet/IP modules:

Produced Tag Connections 0

Packets/Second	= (1 + connections)/RPI for each produced tag	
	=	N/A

Consumed Tags 1 tag @ 20 ms

Packets/Second	= 2/RPI for each consumed tag	
	= 2/20 ms =	100
Total Packets/Second		100

The total of 100 packets/second is well within the remaining bandwidth for these modules (4500 pps limit for a 1756-ENBT module; 810 pps limit for a 1756-ENET/B module).

EtherNet/IP module serving as a scanner

EtherNet/IP module #1, in the chassis with the main controller, is the most loaded EtherNet/IP module. It must communicate with the five specified I/O racks, the PanelView, and the two other controllers (using a produced tag). It must perform explicit messaging as well.

Assume the PanelView is a direct connection at an RPI of 100 ms. There are also 5 rack optimized connections to I/O racks at RPIs of 20 ms, 15 direct connections at 50 ms RPIs for the analog modules, and 1 produced tag with 2 connections (to EtherNet/IP modules 2 and 3). Fill in the worksheet for this module as follows:

EtherNet/IP Module ID: 1

Rack Optimized Connections 5 @ 20 ms
(for I/O)

Packets/Second = (2 x connections)/RPI

= (2 x 5)/ 20 ms = 500

Direct Connections 1 @ 100 ms
(PanelView)

Packets/Second = (2 x connections)/RPI

= (2 x 1)/ 100 ms = 20

Direct Connections 15 @ 50 ms
(for I/O)

Packets/Second = (2 x connections)/RPI

= (2 x 15)/ 50 ms = 600

Produced Tag Connections 2 @ 20 ms

Packets/Second = (1 + connections)/RPI for each produced tag

= (1 + 2)/ 20 ms = 150

Consumed Tags 0

Packets/Second = 2/RPI for each consumed tag

= N/A

Total Packets/Second 1270

The total of 1270 packets/second is within the limit for a 1756-ENBT module, but it well exceeds the recommended limit for a 1756-ENET/B module. This system does not work for a chassis containing a 1756-ENET/B module.

Recommendations to achieve more throughput with an existing 1756-ENET/B module

To achieve more throughput with an existing 1756-ENET/B module, do one of the following:

- Add a 1756-ENBT or 1756/ENET/B module
- Replace the 1756-ENET/B module with a 1756-ENBT module

In this example, module 1A is an ENET/B module. Add module 1B as a 1756-ENBT module.

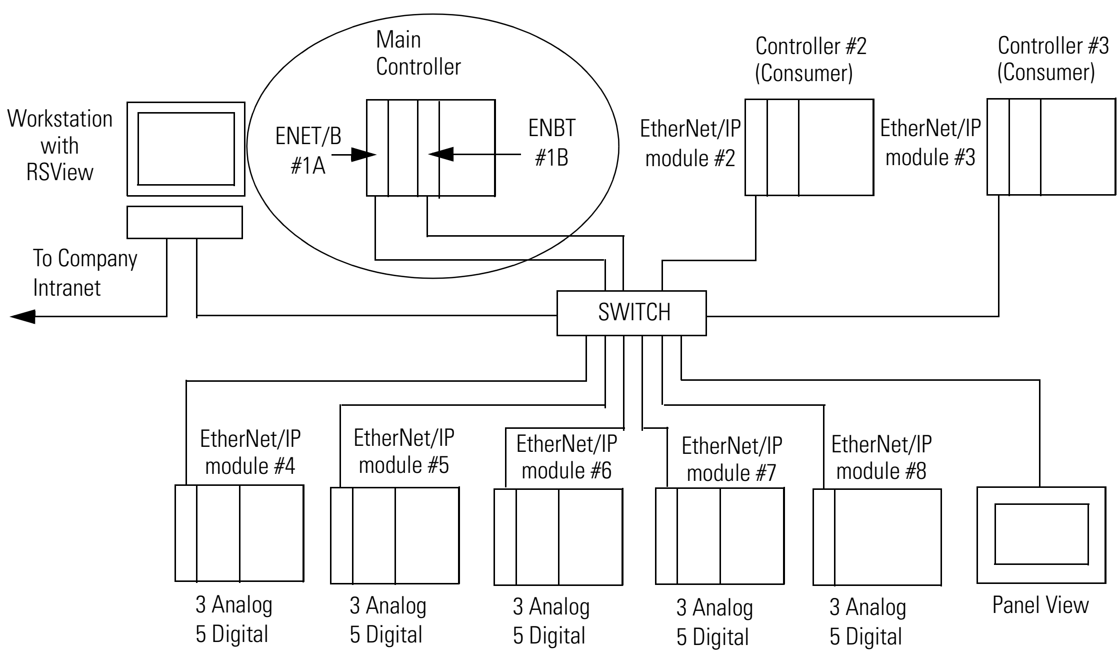

In this configuration, divide the scanning function between the 1756-ENET/B and 1756-ENBT modules. For example, connect EtherNet/IP modules 2-5 and the Panel View to the 1756-ENET/B module 1A, and connect EtherNet/IP modules 6-8 to the 1756-ENBT module 1B.

The RSLogix 5000 software I/O configuration for the rack optimized and direct connections in this system is shown below. This configuration is done in the Main Controller in rack number 1.

These example configurations show generic EtherNet/IP modules as the adapters under module 1A and module 1B. You could select any EtherNet/IP module that works for your application (1734-AENT, 1756-ENBT, 1794-AENT, PanelView terminal, etc.).

For the control processor:

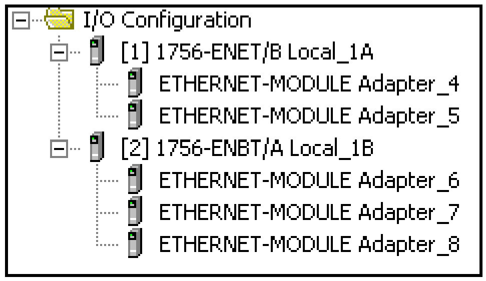

The RSLogix 5000 software configurations for produced tags in this system is shown below. These configurations are done in the consumers (i.e., the processors in racks 2 and 3).

For the consumer in rack 2:

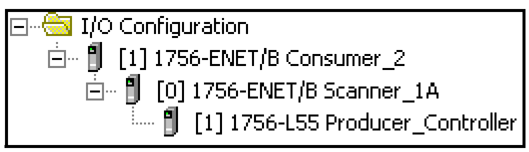

For the consumer in rack 3:

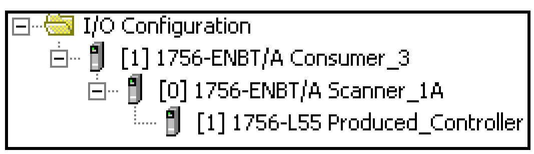

In this new configuration, 1756-ENET/B module 1A has a produced tag with connections to the Consumer_2 and Consumer_3 modules at an RPI of 20 ms, two rack optimized connections to the Adapter_4 and Adapter_5 modules at RPIs of 20 ms, one direct connection to the PanelView at an RPI of 100 ms, and six direct connections to the analog modules in racks 4 and 5 at RPIs of 50 ms.

For 1756-ENET/B module 1A, calculate the total packets/second as follows:

EtherNet/IP Module ID: 1A

Rack Optimized Connections 2 @ 20 ms
(for I/O)

Packets/Second = (2 x connections)/RPI

= (2 x 2)/ 20 ms = 200

Direct Connections 1 @ 100 ms
(PanelView)

Packets/Second = (2 x connections)/RPI

= (2 x 1)/ 100 ms = 20

Direct Connections 6 @ 50 ms
(for I/O)

Packets/Second = (2 x connections)/RPI

= (2 x 6)/ 50 ms = 240

Produced Tag Connections 2 @ 20 ms

Packets/Second = (1 + connections)/RPI for each produced tag

= (1 + 2) / 20 ms = 150

Consumed Tags 0

Packets/Second = 2/RPI for each consumed tag

= N/A

Total Packets/Second 610

EtherNet/IP interface 1A has a total of 11 implicit connections.

The total of 610 packets/second is well within the bandwidth for 1756-ENET/B module 1A (810 pps limit for a 1756-ENET/B module).

1756-ENBT module 1B has three rack optimized connections to the Adapter_6, Adapter_7, and Adapter_8 modules at RPIs of 20 ms. It also has nine direct connections to analog modules in racks 6-8 at RPIs of 50 ms. It has no produced or consumed tags.

For 1756-ENBT module 1B, calculate the following:

EtherNet/IP Module ID: 1B

Rack Optimized Connections 3 @ 20 ms
(for I/O)

Packets/Second	= (2 x connections)/RPI	
	= (2 x 3)/ 20 ms =	300

Direct Connections 9 @ 50 ms
(for I/O)

Packets/Second	= (2 x connections)/RPI	
	= (2 x 9)/ 50 ms =	360

Produced Tag Connections 0

Packets/Second	= (1 + connections)/RPI for each produced tag	
	=	N/A

Consumed Tags 0

Packets/Second	= 2/RPI for each consumed tag	
	=	N/A

EtherNet/IP interface 1B has a total of 12 implicit connections.

The total of 660 packets/second is well within the bandwidth for 1756-ENBT module 1B (4500 pps limit for a 1756-ENBT module).

Conclusion

The redesigned system now has sufficient bandwidth for the application. This concludes the first goal of performance predictions, determining if the system can fulfill the application's requirements.

Next, determine the maximum input, output, and produced tag times for the system. Then you can determine if these are acceptable for the application.

Part II: Estimate the maximum input or output times for CIP connections

Calculate the worst-case, maximum input (I/O to controller) or output (controller to I/O) times for CIP connections in our system.

For a **Rack Optimized Connection** the maximum input or output time for a CIP connections is estimated as:

$$T_{MAX\ (Rack\ Optimized)} = RPI = 20ms$$

For the **Direct Connect Analog Modules**, assume non-isolated modules with the real-time sampling (RTS) rate equal to the RPI (i.e., 50 ms). Therefore, the maximum input or output time is estimated as:

$$T_{MAX\ (Analog\ Non\text{-}Isolated)} = 2RTS$$

$$= 2 \times 50 \text{ ms} = 100 \text{ ms}$$

Isolated modules have an RTS rate of 1RTS.

For a **Produced or Consumed Tag**, the maximum input or output time is estimated as:

$$T_{MAX\ (Produced/Consumed\ Tag)} = 1RPI$$

$$= 1 \times 20 \text{ ms} = 20 \text{ ms}$$

Assume that these times are acceptable for the example application. If you find that the times for your system are too slow, you can make adjustments to operate at faster RPIs. This may entail selecting I/O modules that operate at faster data rates, adding more EtherNet/IP modules, and/or other changes as outlined in step 5 on page 4-2.

No further modifications have been made, so the system is complete.

Refine estimates

To further increase the accuracy of these times, include considerations for system delays:

For a **Rack Optimized Connection**, to the RPI, add these delays:

$$T_{MAX\ (Rack\ Optimized)} = \text{input filter} + \text{RPI} + \text{transmission} + \text{switch} + \text{queue}$$

where:

Type of delay:	Description:
input filter	Discrete input modules have filters. The default for a 1756 discrete I/O module is 1 ms. For The default for a FLEX discrete I/O module is 0.25 ms. input delay = 1 ms There are no filters for outputs, so there is no additional delay for outputs. Outputs are always sent using an RPI timer.
transmission	The transmission delay is the interval of time that it takes a packet to be transmitted at a specific bit rate (e.g. 100Mbps). For example, in a 7-slot ControlLogix chassis, the size of the entire packet is approximately 122 bytes (including header, all protocols, all data, and CRC). At 100Mbps, this packet takes approximately 10 microseconds. (0.01 ms) on the wire. transmission delay = 0.01 ms x (number of CIP connections)
switch	Switch latency is the delay between reception of the first bit and transmission of the first bit. This delay depends on the type of switch. It is typically 0.1 ms. switch delay = 0.1 ms
queue	Input data is sent from the remote rack (adapter), through a switch, through a communication module in the controller rack, and finally to a controller (scanner). If two or more input CIP connections are simultaneously ready to be transmitted, they must be transmitted sequentially. It takes 0.2 ms for a 1756-ENBT to process 1 implicit packet. Note that 0.2 ms is equal to the reciprocal of 5000 (pps). The total queue delay is 0.2ms times the number of CIP connections through the module. queue delay = 0.2 ms X (number of CIP connections)

For the **Direct Connect Analog Modules**, you add the same transmission, switch, and queue delays as for rack-optimized data:

$$T_{MAX\ (Analog\ Non\text{-}Isolated)} = 2RTS + transmission + switch + queue$$

For a **Produced or Consumed Tag**, you add the same transmission, switch, and queue delays as for rack-optimized data:

$$T_{MAX\ (Produced/Consumed\ Tag)} = 1RPI + transmission + switch + queue$$

Refine the example times for EtherNet/IP interface 1A

For the example system with 1756-ENET/B interface 1A (page 4-12), the refined calculation including delays is:

T = Flex input filter + RPI + transmission + switch + queue
= 0.25ms + 20ms + [0.01ms x (connections through this module)] + 0.1ms + [1.1ms x (connections through this module)]
= 0.25ms + 20ms + [0.01 x 11]ms + 0.1ms + [1.1ms x 11]
= 0.25ms + 20ms + 0.11ms + 0.1ms + 12.1ms
= 32.56ms

The queuing delay uses a multiplier of 1.1ms because the 1756-ENET/B module supports only 900 pps. This means that a single packet takes 1/900 second (1.1ms) to process. If the example used a 1756-ENBT module, which supports 5000 pps, the multiplier would be 1/5000 seconds (0.2ms).

Conclusion:

For this example, interface 1A has an RPI (20ms) that is only 2 times larger than the queuing delay (12.1ms). This means that the RPI value does not dominate the equation and the RPI value will be a poor estimate of the total input delay.

The rough estimate for a rack optimized connection (see page 4-16) was 20ms (RPI). But the refined calculation is significantly larger, mainly because of the queuing delay. You could reduce the queueing delay by using a 1756-ENBT module instead of the 1756-ENET/B module.

Worksheets

EtherNet/IP Module Packets/Second

EtherNet/IP Module ID: ______

Rack Optimized Connections _____

Packets/Second	= (2 x connections)/RPI	
	=	_____

Direct Connections _____

Packets/Second	= (2 x connections)/RPI	
	=	_____

Produced Tag Connections _____

Packets/Second	= (1 + connections)/RPI for each produced tag	
	=	_____

Consumed Tags _____

Packets/Second	= 2/RPI for each consumed tag	
	=	_____

Total Packets/Second _____

EtherNet/IP Module ID: ____

Rack Optimized Connections ____

Packets/Second = (2 x connections)/RPI

= ____

Direct Connections ____

Packets/Second = (2 x connections)/RPI

= ____

Produced Tag Connections ____

Packets/Second = (1 + connections)/RPI for each produced tag

= ____

Consumed Tags ____

Packets/Second = 2/RPI for each consumed tag

= ____

Total Packets/Second ____

EtherNet/IP Module ID: ____

Rack Optimized Connections ____

Packets/Second = (2 x connections)/RPI

= ____

Direct Connections ____

Packets/Second = (2 x connections)/RPI

= ____

Produced Tag Connections ____

Packets/Second = (1 + connections)/RPI for each produced tag

= ____

Consumed Tags ____

Packets/Second = 2/RPI for each consumed tag

= ____

Total Packets/Second ____

Maximum Input (I/O to Controller) or Output (Controller to I/O) Times for CIP Connections

For 2 Sigma (95%) Accuracy:

Rack Optimized Connection:

$T_{MAX\ (Rack\ Optimized)}$ = 1 RPI = ______ ms

Direct Connection:

$T_{MAX\ (Discrete)}$ = 1 RPI = ______ ms

$T_{MAX\ (Analog\ Non\text{-}Isolated)}$ = 2 RTS + RPI

= = ______ ms

$T_{MAX\ (Analog\ Isolated)}$ = RTS + RPI

= = ______ ms

Produced or Consumed Tag:

$T_{MAX\ (Produced/Consumed)}$ = 2RPI = ______ ms

For 3 Sigma (100%) Accuracy:

Rack Optimized Connection:

$T_{MAX\ (Rack\ Optimized)}$ = 2 RPI = ______ ms

Direct Connection:

$T_{MAX\ (Discrete)}$ = 2 RPI = ______ ms

$T_{MAX\ (Analog\ Non\text{-}Isolated)}$ = 2 RTS + 2RPI

= = ______ ms

$T_{MAX\ (Analog\ Isolated)}$ = RTS + 2RPI

= = ______ ms

Produced or Consumed Tag:

$T_{MAX\ (Produced/Consumed)}$ = 3RPI = ______ ms

Notes:

How Are We Doing?

Your comments on our technical publications will help us serve you better in the future.
Thank you for taking the time to provide us feedback.

You can complete this form and mail it back to us, visit us online at www.ab.com/manuals, or email us at RADocumentComments@ra.rockwell.com

Pub. Title/Type EtherNet/IP Performance and Application

Cat. No. EtherNet/IP products Pub. No. ENET-AP001D-EN-P Pub. Date October 2004 Part No. 957899-85

Please complete the sections below. Where applicable, rank the feature (1=needs improvement, 2=satisfactory, and 3=outstanding).

Overall Usefulness	1 2 3	How can we make this publication more useful for you?
Completeness (all necessary information is provided)	1 2 3	Can we add more information to help you? ☐ procedure/step ☐ illustration ☐ feature ☐ example ☐ guideline ☐ other ☐ explanation ☐ definition
Technical Accuracy (all provided information is correct)	1 2 3	Can we be more accurate? ☐ text ☐ illustration
Clarity (all provided information is easy to understand)	1 2 3	How can we make things clearer?
Other Comments		You can add additional comments on the back of this form.

Your Name ______________ Location/Phone ______________

Your Title/Function ______________ Would you like us to contact you regarding your comments?

___No, there is no need to contact me

___Yes, please call me

___Yes, please email me at ______________

___Yes, please contact me via ______________

Return this form to: Allen-Bradley Marketing Communications, 1 Allen-Bradley Dr., Mayfield Hts., OH 44124-9705

Phone: 440-646-3176 Fax: 440-646-3525 Email: RADocumentComments@ra.rockwell.com

PLEASE FASTEN HERE (DO NOT STAPLE)

Other Comments

PLEASE FOLD HERE

PLEASE REMOVE

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 18235 CLEVELAND OH

POSTAGE WILL BE PAID BY THE ADDRESSEE

Allen-Bradley
RELIANCE ELECTRIC
DODGE
ROCKWELL SOFTWARE
Rockwell Automation

1 ALLEN-BRADLEY DR
MAYFIELD HEIGHTS OH 44124-9705

Rockwell Automation Support

Rockwell Automation provides technical information on the web to assist you in using our products. At http://support.rockwellautomation.com, you can find technical manuals, a knowledge base of FAQs, technical and application notes, sample code and links to software service packs, and a MySupport feature that you can customize to make the best use of these tools.

For an additional level of technical phone support for installation, configuration and troubleshooting, we offer TechConnect Support programs. For more information, contact your local distributor or Rockwell Automation representative, or visit http://support.rockwellautomation.com.

Installation Assistance

If you experience a problem with a hardware module within the first 24 hours of installation, please review the information that's contained in this manual. You can also contact a special Customer Support number for initial help in getting your module up and running:

United States	1.440.646.3223 Monday – Friday, 8am – 5pm EST
Outside United States	Please contact your local Rockwell Automation representative for any technical support issues.

New Product Satisfaction Return

Rockwell tests all of our products to ensure that they are fully operational when shipped from the manufacturing facility. However, if your product is not functioning and needs to be returned:

United States	Contact your distributor. You must provide a Customer Support case number (see phone number above to obtain one) to your distributor in order to complete the return process.
Outside United States	Please contact your local Rockwell Automation representative for return procedure.

www.rockwellautomation.com

Power, Control and Information Solutions Headquarters

Americas: Rockwell Automation, 1201 South Second Street, Milwaukee, WI 53204-2496 USA, Tel: (1) 414.382.2000, Fax: (1) 414.382.4444
Europe/Middle East/Africa: Rockwell Automation, Vorstlaan/Boulevard du Souverain 36, 1170 Brussels, Belgium, Tel: (32) 2 663 0600, Fax: (32) 2 663 0640
Asia Pacific: Rockwell Automation, Level 14, Core F, Cyberport 3, 100 Cyberport Road, Hong Kong, Tel: (852) 2887 4788, Fax: (852) 2508 1846

Publication ENET-AP001D-EN-P - October 2004 PN 957899-85
Supersedes Publication ENET-AP001C-EN-P - January 2003