NewTek™

NDI™ Network Device Interface

TECHNICAL BRIEF

March 24, 2016 - Version 1

TABLE OF CONTENTS

OVERVIEW

NDI (Network Device Interface) is an open protocol developed by NewTek to enable video-compatible products to share video across a local area network. We believe that the future of the video industry is one in which video is transferred easily and efficiently in IP space, and that this vision will largely supplant current industry-specific connection formats (HDMI, SDI, etc.) in the production pipeline.

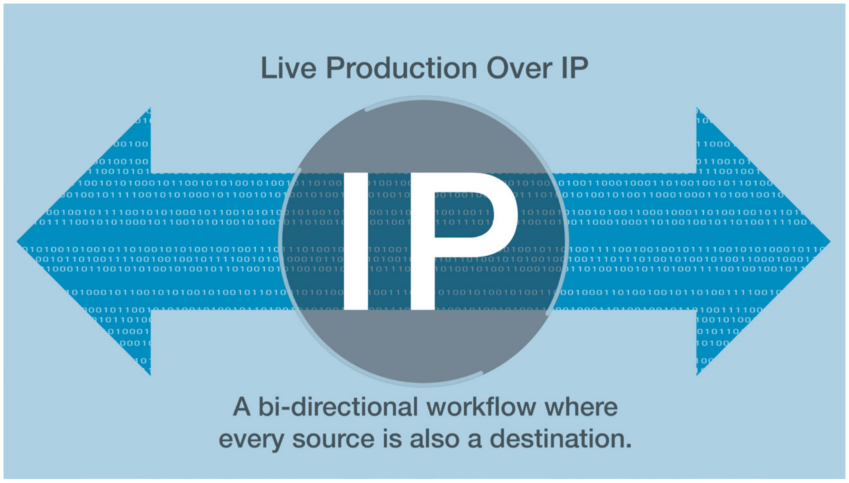
Live Production Over IP
IP
A bi-directional workflow where every source is also a destination.

NDI allows multiple video systems to identify and communicate with one another over IP, and to encode, transmit, and receive many streams of high quality, low latency, frame-accurate video and audio in real time. This new protocol can benefit any network-connected video device, including video mixers, graphics systems, capture cards, and many other production devices.

NDI can operate bi-directionally over a local area network, with many video streams on a shared connection. Its encoding algorithm is resolution and frame-rate independent, supporting 4K (and beyond) along with 16 channels (and more) of floating-point audio. The protocol also includes tools that implement video access rights, grouping, bi-directional metadata, and IP commands. And its superb performance over standard GigE networks makes it possible to transition facilities to an incredibly versatile IP video production pipeline without negating existing investments in SDI cameras and infrastructure or costly new high-speed network infrastructures.

This paper is intended to deliver the essential facts with best practices, and is intended for professionals familiar with common networking devices and concepts. What is wonderful about NDI is that it can be utilized on almost any Gigabit network. As the production grows, however, additional considerations will be required and that is what we will cover.

DISCOVERY & REGISTRATION

Sending and receiving video streams across an IP network requires applications supporting video to be able to discover receiving applications that are looking for video. NDI resolves host names to IP addresses over the local area network (LAN) and does so automatically. When you start an application that sends NDI, the devices that can receive NDI become aware instantaneously. While this is a typical function on almost all networks, there are some cases where it is important to know how this works in order to properly configure networks utilizing managed data flows and QoS protocols.

NDI utilizes mDNS (multicast Domain Name System)[1] to create the zero configuration environment for discovery. This service sends an IP multicast message that asks the host to identify itself. The target machine then multicasts a message that includes its own IP address. This multicast is seen by all NDI receiving machines on the subnet, which then use the information in that message to update their own caches. These multicast queries are sent to a multicast address and as a result, no single device is required to have global knowledge. When a service or device sees a query for any service it recognizes, it provides a DNS response with the information from its cache.

The primary benefits of using mDNS is that it requires little or no administration to set up. Unless the network is specifically configured to not allow mDNS, NDI sources will be discovered. This format works when no infrastructure is present and can span infrastructure failures.

The mDNS Ethernet frame is a multicast UDP packet that broadcasts to[2]:

- MAC address 01:00:5E:00:00:FB (for IPv4) or 33:33:00:00:00:FB (for IPv6)
- IPv4 address 224.0.0.251 or IPv6 address FF02::FB
- UDP port 5353

Because mDNS uses a link-local multicast address, its capacity is limited to a single physical or logical LAN. If the networking reach needs to be extended to multiple subnets or to an environment consisting of many different networking technologies, an mDNS gateway is implemented. An mDNS gateway provides a transport for mDNS packets across Layer 3 boundaries by filtering, caching, and redistributing services from one Layer 3 domain to another. This is a process that is managed in Layer 3 capable networking switches (refer to documentation provided from the switch manufacturer).

[1] Apple's mDNS is published as a standards track proposal (RFC 6762) https://tools.ietf.org/html/rfc6762

[2] https://en.wikipedia.org/wiki/Multicast_DNS

On Windows devices in particular, choosing the network location type is critical for the successful discovery and registration of NDI. Typically, the first time a Windows machine is connected to a network, a dialog window appears that allows the user to choose the network location type: **Home**, **Work**, **Public**, or **Domain.** By default, Windows sets a new network location to **Public.** This location is designed to keep machines from being visible and responding to broadcast pings. This location type also affects mDNS responses and, in turn, keeps NDI video streams from being discovered and registered on the network. For successful discovery and registration of NDI, network locations should be set to **Work** or **Home.**

The **Domain** network location is used for domain networks, such as those at enterprise workplaces. This type of network location is controlled by the network administrator and cannot be selected or changed. In this type of configuration, mDNS discovery must be allowed at the domain level.

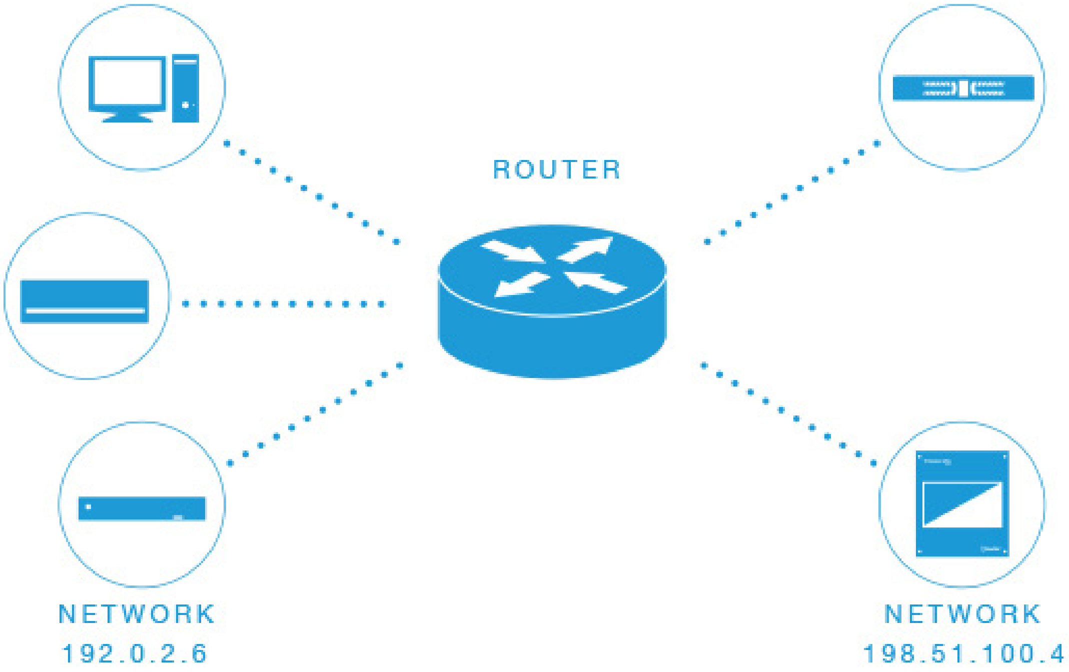

Figure 1. Sample Networking Scenario. For example, if the mDNS gateway functionality is enabled on the router in this figure, then service information can be sent from one subnet to another and vice-versa. For example, the NDI discovery information being advertised in the network with IP address 192.0.2.6 (left) is redistributed to the network with IP address 198.51.100.4 (right) and learned by the mDNS-enabled hosts and devices in that network.

VIDEO DATA FLOW

Once two NDI devices have discovered each other on the network, video can be passed from the sending device to the receiving device. After the compression of the video, the NDI sending device opens a session to the receiving NDI device. At this point, we have two endpoints that consist of an IP address and a port number.

The NDI network frame is a TC packet that transmits in the Ephemeral port range. Ephemeral ports are temporary ports assigned by the NDI sending machine's IP stack, and are assigned from the designated range of ports for this purpose. When the NDI sending device terminates the transmission session, the port then becomes available for reuse, although most IP stacks will not reuse a port number until the entire pool of ephemeral ports have been used. This is important to keep in mind in the event of disconnection due to shut down of the NDI endpoints. The new connection will almost always be assigned a different port number.

The port ranges used for NDI are **49152** to **65535.** If NDI devices are used on networks with firewalls positioned on the internal network, configuration to allow for traffic between devices in this port range is required.

You can view the Ephemeral port range on a computer that is running Windows Vista or later by use of the netsh command[3]. These commands also provide a way to limit the port range, however this will impact all dynamic port assignments to that particular machine.

[3] Default dynamic port range for Windows Vista and Windows Server 2008 https://support.microsoft.com/en-us/kb/929851

GETTING VIDEO ACROSS THE NETWORK

Video, just like voice data in VoIP systems, is a very demanding data stream and will immediately expose a weakness in the network. The network must be capable of supporting multiple video, audio, and data streams in a reliable, synchronized manner, without disruption. When delay, packet loss, and jitter reach thresholds where the video is impacted visually, the usefulness of that video drops to zero. It is important to understand the complexities of video in IP data networks so that these factors can be mitigated.

Networks that are designed to move NDI video streams should be thought of as being primarily utilized for video. IP networks are by their very nature “best effort delivery” systems and were originally developed for the transport of data. Data services, by contrast to video, can function happily with packet retransmissions, lost packets, and even packets arriving out of order. Video streams, while still data, are much more rigid in their requirements.

NETWORK LAYOUT

NDI is designed for use with standard consumer off-the-shelf (COTS) networking devices. Looking closely at the network topology and configuration will help to ensure the maximum possible bandwidth is available.

When selecting a network switch, it is important to check the throughput speeds. Ensure that each port is full duplex (i.e. bi-directional communication) and that the upstream and downstream data speeds for each port are at least 1 Gigabit per second (Gbps). It is best to force the ports on managed switches to utilize 1 Gbps in contrast with Auto Negotiation. The use of Auto Negotiation can sometimes result in 100Mb connections or even lower, which does not renegotiate until the port is flooded with traffic for some time. Also, poor termination of RJ-45 connectors can impact Auto Negotiation.

When possible, it is best to use switches from the same manufacturer, or ideally, the same model of switch, throughout a single subnet. This will simplify configuration and lessen the chances of compatibility and configuration issues.

BANDWITH

NDI operates most efficiently in a dedicated network with high bandwidth and high availability. This is in contrast to unmanaged environments such as the public Internet or networks where video rides along with data without priority.

While a single stream of HD video can easily be delivered on a Fast Ethernet (100 Mbps) network, Gigabit (1000 Mbps) networks are essential in production workflows. A typical NDI stream consisting of 1080i HD video yields a data rate up to 100 Mbps per stream. This extremely efficient stream is designed to have very low latency and allows multiple streams to be stacked together on a single Gigabit network.

QUALITY OF SERVICE

Quality of Service (QoS) is a set of standards and mechanisms that provide a required level of service for network traffic over various technologies. The primary goal of QoS is to determine the priority of specific types of data, ensure sufficient bandwidth, and, in some cases, control jitter and latency. Using QoS technology can deliver control over resources and make the most efficient use of the network, especially in cases where integrated data flows are required.

It is beyond the scope of this paper to determine specific QoS configurations and tools, as NDI does not require the use of QoS to function. However, specifying devices using NDI as real-time data flow on the network can optimize transmission. For environments where the coordination of multiple types of data is required, adhering to the following industry guidelines is recommended:

- **Loss should be no more than 1%.**
- **One-way latency should be no more than 150 ms.**
- **Jitter should be no more than 30 ms.**
- **Guaranteed bandwidth of 100 Mbps per stream required.**

NETWORK INTERFACE SETTINGS

NDI is designed to enable successful video transport using the default configurations of network interface drivers, however most recent network interface drivers do support configuration of advanced properties that can help optimize NDI transmission.

Consider the following adjustments, but note that making adjustments on individual adapters can significantly affect performance and reliability, both positively and negatively. It is important to consider testing performance with a network analyzer before and after each setting change. The following adjustments are intended to help, however performance will depend on particular network and usage (names and available settings vary between vendors, adapter models, and even between different driver versions):

Speed and Duplex: This setting allows for selection of the desired speed and duplex of the network adapter. Usually this is set to Auto Negotiation. To ensure the maximum available throughput, this setting should be set to 1 Gbps Full Duplex.

Energy Efficient Ethernet: When enabled, this allows the adapter to engage power saving features while keeping connections active. This technology uses the standard IEEE 802.3az to allow for less power during periods of low data activity. Adapters that utilize the IEEE 802.3az standard should have no impact on performance of NDI, however some integrated circuits exist that were developed before the standard was finalized or do not adhere to the standard at all. In these cases, it is best to disable the energy efficiency while determining best network optimization.

ENCODING/DECODING

COMPRESSION

NDI uses compression to enable transmission of a large number of video streams across existing infrastructure, specifically discrete cosine transform (DCT), which converts video signals into elementary frequency components. This method of compression is commonly used in encoding formats and mezzanine codecs within the industry.

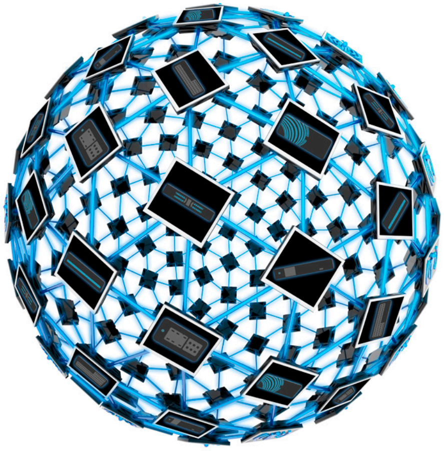

One of the most efficient codecs in existence, NDI achieves significantly better compression than the majority of codecs that have been accepted for professional broadcast use. On a typical, modern Intel-based i7 processor, the codec is able to compress a 1920x1080 video signal at 250 frames per second using a single core.

The peak signal-to-noise ratio (PSNR) of the NDI codec exceeds 70dB for typical video content. Uniquely, and importantly, NDI is the first ever codec to provide multi-generational stability. This means that once a video signal is compressed, there is no further loss. As a practical example, generation 2 and generation 1000 of a decode-to-encode sequence would be identical. Examples of this concatenation are provided in NewTek's NDI SDK[4].

The NDI codec is designed to run very fast and is largely implemented in hand-written assembly to ensure that the process of compressing video frames occurs as quickly as possible. Latency is both a factor of the network connection and the endpoint products. NDI has a technical latency of 16 video scan lines, although in practice, most implementations would be one field of latency[5]. Hardware implementations can provide full end-to-end latency of within 8 scan lines.

FORMATS

NDI fully supports all resolutions, frame rates, and video streams, with and without alpha channel. In practical terms, resolution and frame rates will be determined by the capabilities of endpoint devices.

The most common implementations are expected to utilize 8-bit UYVY and RGBA video, however support for 10-bit and 16-bit is available. The internal pipeline of the codec is maintained entirely at 16-bit or better.

[4] For more information, please visit http://ndi.newtek.com.

[5] The implementations of NDI using the SDK prior to March 3, 2016 typically provide frame-at-a-time delivery for reasons of compatibility with older systems. In these cases, the minimum latency is likely to be one frame.

GLOSSARY

Cache

Cache refers to a reserved section of computer memory or an independent high-speed storage device used to accelerate access and retrieval of commonly used data.

Domain

A domain refers to a LAN subnetwork of users, systems, devices and servers. Domain can also refer to the IP address of a website on the Internet.

DNS

DNS (Domain Name System) is a system used by the Internet and private networks to translate domain names into IP addresses.

mDNS

mDNS (Multicast DNS) refers to the use of IP multicast with DNS to translate domain names into IP addresses and provide service discovery in a network that does not have access to a DNS server.

Ethernet

Ethernet, standardized as IEEE 802.3, refers to a series of LAN (Local Area Network) technologies used to connect computers and other devices to a home or business network. Ethernet is a physical and data link layer networking protocol that supports data transfer rates starting at 10 Mbps, typically over twisted pair cabling, but also fiber optic and coaxial cabling.

IGMP

IGMP (Internet Group Management Protocol) is the protocol used in IP multicasting that allows a host to report its multicast group membership to networked routers in order to receive data, messages, or content addressed to the designated multicast group.

IP

IP (Internet Protocol) is the communications protocol for the Internet, many wide area networks (WANs), and most local area networks (LANs) that defines the rules, formats, and address scheme for exchanging datagrams or packets between a source computer or device and a destination computer or device.

IPv4

IPv4 (Internet Protocol Version 4) is the fourth and most commonly used version of the Internet Protocol. IPv4 uses a 32-bit IP address scheme for network identification and communication, with each unique IP address expressed as four numbers (between 0 and 255) separated by decimal points.

IPv6

IPv6 (Internet Protocol Version 6) is the latest version of the Internet Protocol, developed to eventually replace IPv4 (Internet Protocol Version 4). IPv6 uses a 128-bit IP address scheme for network identification and communication, with each unique IP address expressed as eight groups of four hexadecimal digits (numbers from 0-9 or letters from A-F) separated by colons. In addition to increasing the number of available IP addresses exponentially, IPv6 simplifies and streamlines network communication, while increasing security, compatibility, and efficiency.

LAN

LAN (Local Area Network) is a network that connects computers and devices in a room, building, or group of buildings. LANs are typically deployed in homes, offices, and schools, where users share access to the same server, resources, and data storage. A system of LANs can also be connected to form a WAN (Wide Area Network).

Layer 2

Layer 2 refers to the second layer, or Data Link layer, of the OSI networking model. A layer 2 switch uses hardware-based switching to transmit data between connected devices based on their MAC (Media Access Control) layer addresses.

Layer 3

Layer 3 refers to the third layer, or Network layer, of the OSI networking model. A layer 3 switch uses hardware-based switching to transmit data between connected devices based on their IP (Internet Protocol) addresses. A layer 3 switch can support packet inspection and routing protocols to prioritize and forward traffic.

MAC Address

MAC (Media Access Control) address refers to a unique physical address that identifies a network node.

Mbps

Mbps (Megabits per second) is a unit of measurement for data transfer speed, with one megabit equal to one million bits. Network transmissions are commonly measured in Mbps.

NDI

NDI (Network Device Interface) is an open protocol developed by NewTek for IP transmission and live production using standard LAN networking. NDI allows networked video systems to identify and communicate with each other over IP and encode, transmit, and receive multiple streams of broadcast-quality, low-latency, frame-accurate video and audio in real time.

OSI

The OSI (Open System Interconnection) reference model is a standard that defines worldwide network communication, developed by ISO (International Organization for Standardization). The OSI reference model divides network communication into seven layers: 1) Physical, 2) Data Link, 3) Network, 4) Transport, 5) Session, 6) Presentation, and 7) Application.

Packet (Frame)

A packet, also known as a frame or datagram, is a unit of data transmitted over a packet-switched network, such as a LAN, WAN, or the Internet.

Port

A port is a communications channel for data transmission to and from a computer on a network. Each port is identified by a 16-bit number between 0 and 65535, with each process, application, or service using a specific port (or multiple ports) for data transmission. Port can also refer to a hardware socket used to physically connect a device or device cable to your computer or network.

QoS

Qos (Quality of Service) is the measure of performance for system or network, with considerations that include availability, bandwidth, latency, and reliability. QoS can also refer to the prioritization of network traffic to ensure a minimum or required level of service, predictability, and/or control.

Subnet

Subnet (short for subnetwork) refers to a distinct subdivision of an IP network, usually created for performance or security purposes. Subnets typically include the computers, systems, and devices in one location, office, or building, with all nodes sharing the same IP address prefix.

TCP

TCP (Transmission Control Protocol) is a network communications protocol, which enables two host systems to establish a connection and exchange data packets, and ensures data is delivered, intact, to the correct destination. TCP is typically grouped with IP (Internet Protocol) and known collectively as TCP/IP.

UDP

UDP (User Datagram Protocol) is an alternative protocol to TCP that is used when reliable delivery of data packets in not required. UDP is typically used for applications where timeliness is of higher priority than accuracy, such as streaming media, teleconferencing, and voice over IP (VoIP).

WAN

WAN (Wide Area Network) is a network that spans a relatively broad geographical area, such as a state, region, or nation. WANs typically connect multiple smaller networks, such as LANs (Local Area Network) and MANs (Metropolitan Area Network). The Internet is an example of a WAN.