RAVENNA

Draft 0.3

Network Requirements

This document describes the architectural and performance requirements of the RAVENNA technology to be provided by the underlying networking infrastructure.

ALC NetworX

ALC NetworX GmbH

Am Loferfeld 58
81249 Munich
GERMANY

Fon: +49 (89) 44236777-0
Fax: +49 (89) 44236777-1

@: info@alcnetworx.de

www.alcnetworx.de

Version information:

Draft 0.3 – 2021-04-28 – Minor additions

RAVENNA – Network Requirements

DRAFT 0.3

TOC

1 INTRODUCTION

RAVENNA is a technology for high-performance media streaming on qualifying IP networks. High-performance is defined by:

- Low latency: in most application scenarios, this translates into a latency of single-digit milliseconds; the most demanding use cases may even require sub-milliseconds latency.
- Full bit transparency: the media content is transferred without any processing, thus no coding / encoding is deployed.
- Sample-accurate synchronization: independent streams can be aligned with respect to absolute time up to a precision, which allows sample-accurate play-out / processing where desired.
- Phase alignment: regenerated media clocks can be aligned to phase accuracy according to AES11 specification for master clocks.

In order to achieve these goals, the underlying network infrastructure has to fulfill certain pre-requisites and to guarantee specific performance characteristics. Naturally, these demanding requirements can best be matched in single subnet LAN environments. But since RAVENNA is fully based in IP, with careful design / administration RAVENNA streams can potentially also be established on well-managed WAN environment.

2 GENERAL NETWORK REQUIREMENTS

2.1 General Network / Link Considerations

- In general, a structured network topology (e.g. star, tree) is recommended; daisy chaining (hop-to-hop linking) should be avoided as it adds performance uncertainties (e.g. latency, jitter, loss of timing precision, bandwidth restrictions etc.).
- The network environment needs to be managed or administrated in order to establish a qualifying operational environment.
- Sufficient link / backbone bandwidth capacity is assumed as no solution-inherent means of bandwidth management / reservation are available. Adequate amount of bandwidth headroom should be reserved (e.g. it is suggested to not assign more than 75% of the available bandwidth to RAVENNA traffic).
- Full switching capacity is assumed for active network infrastructure components (switches, routers etc.).
- The network environment needs to guarantee 100% packet delivery (zero UDP packet loss) as lost packets cannot be recovered by design, unless redundant connectivity on physically separated networks is configured.
- A single subnet LAN environment is preferred for optimal performance, any routing may add performance uncertainties (e.g. latency, jitter, loss of timing precision, bandwidth restrictions etc.).
- RAVENNA works on Fast Ethernet speed and above, but Gigabit Ethernet is recommended for node and backbone connectivity.
- When operating on Ethernet, full duplex operation is required.

2.2 IP Connectivity

- Currently, IPv4 is used as a protocol basis. IPv6 support may be added in the future.
- Native IP addressing with free port number assignment is employed; operation across NAT routers will not work.
- Multicast operation needs to be supported, as it is the main method for stream traffic transport and the basis for PTP operation.
- Packet fragmentation is not allowed for RAVENNA traffic.
- Jumbo packets should be avoided on any path carrying RAVENNA traffic as this may result in significant performance impairments (increased latency and jitter, loss of timing precision etc.)

2.3 Synchronization

- RAVENNA system synchronization is based on IEEE1588-2008 (sometimes referred to PTPv2).
- PTP packets are encapsulated in IP utilizing well-known multicast addresses (224.0.1.129…132 / 224.0.0.107).
- For adequate precision results, all paths between any PTP master and slave node need to have symmetrical operational parameters in terms of latency and jitter. In case different routes between any two nodes may be established within a network environment, measures for determinable packet routing need to be taken.
- The achievable synchronization precision usually depends on PTP packet jitter. Very low packet jitter is required for phase-accurate synchronization between nodes. While this can usually be established within non-routed environments (LAN segments), deployment of PTP-aware network equipment (switches with transparent and / or boundary clock support) may be required in critical paths, particularly in routed environments .
- PTP-aware network equipment may even be required in larger LAN environments or in heavily loaded environments to achieve sufficient synchronization performance.
- In routed or WAN environments, the deployment of independent, GPS-synchronized local GM clocks may be required to achieve satisfying synchronization performance.

2.4 Streaming

- RAVENNA streams use RTP/AVT over UDP with assignable port numbers.
- IP packet sizes may vary from 40 up to 1500 bytes (including RTP / UDP / IP overhead).
- RAVENNA streams are transported with multicast by default; unicast operation is also possible for selected streams. Multicast operation in LAN environment requires support of IGMP[1] – switches need to understand IGMP, and at least one switch needs to be configured as IGMP querier. Employment of unconditional multicast forwarding will most likely result in uncontrolled network flooding!
- In routed or WAN environment, administrative measures (e.g. TE) need to be taken to enable the transport of multicast streams between the designated nodes.
- RTCP messaging is used for status monitoring, port number is automatically set to +1 with respect to the UDP port of the referring RTP stream.
- Bounded, low UDP packet jitter is desired, as any increase of packet jitter needs to be compensated with larger latency settings at the receiving end. Unbounded UDP packet jitter may lead to unrecoverable late packet arrival (which is equal to packet loss).

[1] RAVENNA requires IGMPv2 for ASM (any-source multicast). IGMPv3 is required for SSM (source-specific multicast).

2.5 QoS

- Qualities of Service measures need to be applied to ensure expedited forwarding and lowest possible jitter of RAVENNA-related traffic.
- As DiffServ (Differentiated Services) is widely supported in current network equipment, RAVENNA traffic can be tagged with configurable DSCP values to support appropriate configuration of network devices.
- PTP packets need to be tagged with the highest available priority to ensure timely transport (specifically in non PTP-aware network equipment). The recommended default value for PTP event messages is CS6, but may be adjusted to match individual constraints of certain network configurations. Note: Most existing corporate networks have assigned highest QoS priority to network management and control traffic; PTP event messages are considered to be of same class.
- RTP stream packets need to be tagged with a high priority to ensure expedited or preferred forwarding. Priority settings are assignable and may vary between streams to allow prioritization between streams or to match individual constraints of certain network configurations. Priority needs to be lower than for PTP packets. Recommended values are EF and AF41.
- Care should be taken when designing the priorities in case other prioritized traffic classes (i.e. real-time video or VoIP) traversing the same network segments. For example, most existing corporate networks have assigned EF or AFx1 tags to telephony (VoIP) traffic; this needs to be adjusted accordingly as RAVENNA real-time traffic is prone to large packet jitter or packet loss; hence, RAVENNA traffic needs to be assigned a higher priority than VoIP traffic. It is also recommended to prioritize RAVENNA traffic against any real-time video traffic, as video traffic usually is more insensitive to larger packet jitter or latency.
- DSCP assignments may be changed on-path by routing equipment to match any specific constraints; however, RAVENNA packets still need to receive highest priority compared against any other traffic on any given link (not just within an assigned VLAN or other logical traffic group). Special care needs to be taken if on-path routing equipment removes the original DSCP tags (i.e. as part of their "untrust" policy) using other mechanisms to preserve priority for RAVENNA traffic; in this case it may be necessary to reestablish meaningful DSCP tags to maintain RAVENNA traffic prioritization in the destination network segment.

2.6 Device Configuration & Connection Management

- Device configuration (device specific setup and stream configuration) is achieved via HTTP with assignable port number (default: 8080).
- Connection Management (connecting to streams) is executed via RTSP; port number may be different from the selected HTTP port.

2.7 Advertisement / Discovery

- Device advertisement and discovery is achieved through DNS-SD.
- Currently, DNS-SD is implemented using mDNS (Zeroconf), which is by default limited to a single subnet within a LAN. Through administrative means Zeroconf can be expanded to reach beyond the scope of a single subnet.
- As an alternative method for larger systems, a DNS server with UPDATE capability can be employed[2].

2.8 Redundancy

- Operational redundancy can optionally be achieved through deployment of two (physically) independent networks. Therefore, RAVENNA nodes may exhibit two independent network interfaces with different IP address assignments.
- A stream sender would transmit identical stream packages on its 2 independent network ports. Receivers would receive these identical packets independently on both network interfaces. Incoming packets are ordered and matched against their RTP time stamps. As long as at least one of any two identical packets is received in time, play-out or processing is not interrupted.
- Identical streams can be configured to different multicast address and UDP port assignments on both networks. Equal (redundant) streams are indicated by means of SDP signaling.
- SMPTE ST 2022-7 is supported.

[2] Most (if not all) RAVENNA device implementations use mDNS; DNS-SD will most likely not be implemented.

3 RAVENNA IN WAN ENVIRONMENT

RAVENNA is mainly aimed at LAN applications where the reliability of the network is high and the management of the network is under control. That does not preclude its application in a WAN setting, though. One must keep in mind the technical limitations, however.

Routing is one of the benefits of using Layer 3, the IP layer, instead of lower layers. It means that traffic is not restricted to a subnet. Since RAVENNA is fully based on IP protocols, it can potentially be used in routed environments including qualifying WANs.

Usually, there will be routers or gateways sitting at the boundaries between LAN and WAN. This is necessary for a number of reasons, including security and management. In many cases, these devices perform additional tasks like NAT (network address translation), possibly encryption and/or compression, and the lower reliability of WAN connections may require some form of error correction. Ordinary RAVENNA nodes will usually not cater for the specifics of a WAN connection, relying instead on the services of a suitable gateway.

RAVENNA can be used in WAN environments, if the following conditions are met:

- The WAN connection is reliable (zero UDP packet loss) and offers known bandwidth and latency and determinable (low) UDP packet jitter
- The WAN connection and the routers support Layer 3 QoS (Diffserv)
- Routing does not use address translation (no NAT)
- Common absolute time is provided on both ends, either by PTPv2 transparency across the WAN link, provided that satisfying precision can be achieved, or by having GPS-locked Grandmasters on both ends of the link.

3.1 FAQs

3.1.1 So if routing can be used, why not across the internet?

The internet is considered to be an unmanaged / uncontrolled environment. Non-determinable performance, poor reliability and non-availability of quality of service makes it impossible to use RAVENNA across the public internet.

3.1.2 Why would NAT not be allowed?

Having NAT (network address translation) activated in a router will typically prevent RAVENNA from working. This is because IP addresses and port numbers are being sent inside various protocols, most notably with the session description in SDP format. NAT would normally not find these occurrences, and thereby fail to translate the addresses there. The result is that the addresses communicated via SDP do not match the actual addresses that result from NAT, and communication fails. Making RAVENNA work across NAT routers would require the NAT routers to translate addresses in SDP records, too. This would require some sophisticated deep packet inspection and rewriting facilities, which are unlikely to be available in typical devices.

3.1.3 How do switches / routers deal with multicasting?

Switches use a technique called IGMP snooping to control the forwarding of multicast packets. IGMP is technically a protocol at IP layer (layer 3) while switches typically operate at the Ethernet layer (layer 2). For determining which multicast packet needs to be forwarded to which output ports, the switches need to monitor and interpret the IGMP traffic between routers and end-nodes on the network. Switches without IGMP snooping will broadcast multicast packets on all outputs, which is undesirable in media streaming applications, because it causes excess network load on all nodes who need to discard large numbers of packets they aren't interested in.

Routers usually prevent multicast traffic from being forwarded at all, as this could result in non-controllable traffic conditions. However, RAVENNA multicast stream traffic is very predictable in terms of volume and required bandwidth. The multicast addresses and related UDP ports can be assigned upon stream generation and are thus known. This would allow network administrators to statically configure forwarding of any RAVENNA-related multicast stream traffic to the desired locations.

3.1.4 Can PTP traffic be routed across WAN environment?

Basically, what has been said above for RAVENNA multicast stream traffic, applies to PTP traffic as well, as it is a standard protocol utilizing well-known multicast addresses and port numbers. However, as the achievable precision may degrade rapidly when routed across WAN connections, it may be preferable or even become mandatory to deploy local GPS-synchronized PTP Grandmaster devices to distribute accurate absolute time within the participating LAN subnets.

3.1.5 Are device and stream management or advertising and discovery required to be visible / accessible at all participating sites?

While these services can be made available for remote access (the underlying protocols just need to be allowed to travel across the WAN connection), it is not necessary for establishing synchronized RAVENNA stream exchange. All relevant connection parameters of a certain stream are stored in a related SDP file generated by the sending node. This file needs to be accessible by the designated receiving node. If remote access through HTTP is impossible or not desired, the file can be moved (copied) by any available means into the vicinity of the designated receiving node for local access. After reading the file, the receiving node can connect to the desired stream without any further HTTP or RTSP request.

3.1.6 How can the QoS requirements be met?

While QoS based on DiffServ works well in administrated LAN environment, things get difficult when several independent services are routed across a common WAN connection. Usually, DSCP will be ignored by edge routers or gateways in order to allow QoS assignments specifically required for shaping / balancing the consolidated traffic for the given connection. The problem here is that these edge routers do not now any specific requirements of individual applications. However, to ensure best possible transfer of RAVENNA traffic across the link, RAVENNA packets need to receive the same relative priorities and transport guarantees as within the LAN environment, but this time with respect

to *any* other traffic on that given link. So, remapping of available QoS means must be applied with some knowledge about the requirements mandated for unimpaired RAVENNA operation. Applicable means may include bandwidth reservation (IntServ), appropriate MPLS administration or other means of Traffic Engineering (TE).

However, it needs to be ensured, that upon leaving the WAN environment and re-entering the destination LAN segment, the original QoS settings need to be reestablished – at least to a point where the required performance parameters can be met. Thus, transparent packet routing by applicable means of encapsulation (MPLS / VPLS or similar) may be required.

5 APPENDIX

5.1 Abbreviations

AES	Audio Engineering Society
AES11	a standard covering the synchronization of digital audio signals
DiffServ	Differentiated Services, a mechanism for providing Quality of Service (QoS) guarantees on IP networks (RFC 2474)
DNS	Domain Name System (RFC 1034)
DNS-SD	DNS-based Service Discovery (IETF draft-cheshire-dnsext-dns-sd-10)
DSCP	Differentiated Services Code Point (RFC 2474)
GPS	Global Positioning System
HTTP	Hypertext Transfer Protocol (RFC 2616)
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IGMP	Internet Group Management Protocol (RFC 2236)
IP	Internet Protocol (RFC 791)
IPv4	Internet Protocol version 4 (RFC 791)
IPv6	Internet Protocol version 6 (RFC 2460)
LAN	Local Area Network
mDNS	multicast-DNS (part of Zeroconf specification)
multicast	simultaneous transmission of messages to a group of network destinations identified by a virtual multicast group address (one-to-many transmission)
Node	a device acting as connector or link between two domains or layers
PTP	Precision Time Protocol (an acronym for IEEE 1588)
QoS	Quality of Service
RAVENNA	Real-time Audio Video Enhanced Next-generation Network Architecture
RFC	Request for Comments, an IETF memorandum on Internet standards and protocols
RTCP	Real-time Transport Control Protocol (RFC 3550)

RTP	Real-time Transport Protocol (RFC 3550)
RTP/AVP	RTP Audio Video Profile (RFC 3551)
RTSP	Real-Time Streaming Protocol (RFC 2326)
SDP	Session Description Protocol (RFC 4566)
SMPTE	Society of Motion Picture and Television Engineers
SMPTE ST 2022-7	Standard for Seamless Protection Switching of RTP Datagrams
UDP	User Datagram Protocol (RFC 786)
unicast	transmission of messages to a single network destination identified by a unique address (one-to-one transmission)
WAN	Wide Area Network
Zeroconf	Zero configuration networking (RFC 3927)

Document information:

Authors:

Stefan Heinzmann

Andreas Hildebrand

Editor: Andreas Hildebrand

Version: Draft 0.3

Last revision: 2021-04-28 by Andreas Hildebrand

File: RAVENNA Network Requirements - Draft 0.3.docx