

Workgroup: Network Working Group
Internet-Draft: draft-aw-metaverse-icn-01
Published: October 2023
Intended Status: Informational
Expires: 25 April 2024
Authors: C. Westphal H. Asaeda
Futurewei NICT

Metaverse and ICN: Challenges and Use Cases

Abstract

This document considers some challenges for ICN support of Metaverse-type applications from a networking perspective. Also, one use case is presented to promote one of our future visions.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 3 April 2024.

Copyright Notice

Copyright (c) 2023 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Introduction](#)
- [2. Definitions and Acronyms](#)
- [3. Metaverse Definition and Use Cases](#)
 - [3.1. Definition](#)
 - [3.2. Taxonomy](#)
- [4. ICN Challenges](#)
 - [4.1. Metaverse Objects](#)
 - [4.2. Centralization](#)
 - [4.3. Interoperability](#)
- [5. Use Cases](#)
 - [5.1. Moonshot Project](#)
- [6. Conclusions](#)
- [7. IANA Considerations](#)
- [8. Security Considerations](#)
- [9. Acknowledgments](#)
- [10. Informative References \(TBD\)](#)
- [Authors' Addresses](#)

1. Introduction

The experience of a virtual world, simulated online, with interactions with other people distributed over the real, physical world, is close to being supported by networking technologies in the near future. It is widely believed 6G networks will support enough bandwidth and short enough RTTs to enable such Metaverse applications.

However, there are challenges to deploy this type of application at scale, due to their distributed and shared nature. A Metaverse application will combine the properties of several existing applications: the rendering of a virtual world on a display (either a screen or a head-mounted display) draws similarity from streaming a video (albeit with additional requirements); the interactions between the users are related to social media as well as to video-conferencing.

While the Metaverse application has different requirements from social media, video streaming or video-conferencing, it inherits some of their properties. In a Metaverse, two people can isolate in a room and draw on a board to recreate a Zoom call with avatars. And indeed, this type of collaboration is one of the potential use cases brought forward by the Meta company. People would be connected with their friends, and could grant access to their virtual world based upon social connections. And a virtual world could be imagined that recreates the design look and feel, the architecture of a movie, be it that of a dark Gotham or the Paris of Amelie.

Since applications such as video streaming, video conference, or social media sharing have been considered as potential use cases for ICN architecture, it seems natural to investigate if such architecture would benefit the Metaverse use case.

This document attempts to define the framework for such an investigation, from a networking perspective. This is similar to RFC7933 which considers the interaction of ICN with adaptive video streaming.

2. Definitions and Acronyms

TBD

3. Metaverse Definition and Use Cases

First we need to define what we mean by "metaverse" and introduce some taxonomy to help us refine the architectural requirements of such an application.

3.1. Definition

We present here three definitions. We do not settle on one specific definition, as it is not our scope to offer a definitive definition of the metaverse, or to settle any debate about what is/what isn't a metaverse. Rather we see in the different definitions a different set of implications for the design of the application.

Definition #1: "a 3D virtual shared world where all activities can be carried out with the help of augmented and virtual reality services" (Damar 2021)

Definition #2: "an integrated immersive ecosystem where the barriers between the virtual and real worlds are seamless to users, allowing the use of avatars and holograms to work, interact and socialize with simulated shared experience" (Meta 2022)

Definition #3: "the next generation Internet that is always real-time and mostly 3d, mostly interactive, mostly social and mostly persistent" (John Riccitiello)

Note that the first definition is an extension of an AR/VR framework; the second definition includes an ecosystem, which assumes a set of API to integrate multiple elements into the ecosystem; the last definition views the metaverse as the replacement of the Internet, that is a global scale infrastructure that supports an unlimited range of applications and functions, with a requirement of persistence for each end user's application states.

3.2. Taxonomy

As with the definition of the metaverse, we can try to better define what a metaverse is by way of a taxonomy that differentiates according to different criteria. Depending on where the metaverse application falls within these criteria will have an impact on the architectural design, to some extent, as well as the design and user experience of the application itself.

The dimensions that we consider are listed below. This list is inspired by [[dwivedi2022](#)] but includes additional dimensions.

Environment: the environment can be realistic, unrealistic, fused; the more realistic (or detailed) the more bandwidth is required; conversely, some unrealistic environment can be generated and rendered from some basic models that can be distributed ahead of time. Part of the environment is also if it is generated anew or permanent. In the latter case, it can be cached at the edge or on the device.

Interface: the environment can be interacted with through an interface that ranges from a simple phone screen to a 3D head-mounted display (HMD), from a window into the virtual world into an immersive experience; other physical methods to interface the virtual environment (such as haptics) can be included as well. Multi-modal interfaces would require synchronizing multiple datastreams, such as audio-channel, video-channel, haptic feedback channel, etc. This could be an application-layer issue, or a network-layer issue. The interaction with the interface will occur at different timescales. Some timescales are below the RTT achievable by any general multi-hop technology (e.g. head tracking, eye tracking, possibly hand tracking). Others are RTT constrained by propagation latency over even modest distances (e.g. multiple milliseconds to even a local base station on a cellular network). Still others are constrained by factors like audio and video lag to 100-200ms at best. Some are constrained by user QoE expectations (e.g. transaction delay), and yet others are explicitly non-interactive and intentionally time decoupled. All of these can exist simultaneously in a rich Metaverse applications.

Interaction: the level of interaction can be specific to the virtual environment. It can be in one extreme a solitary experience (such as playing game against a computer) and extend to social networking, and/or work collaboration. The granularity of the interaction also impacts the infrastructure requirements.

Security: it is paramount to protect the security and privacy of the experience. This includes data security, privacy, software/

hardware/network security. Further, the granularity of the security may include several layers, as for instance, only a given set of participants can access a given shared metaverse; and within this metaverse, only a subset can have access to objects or rooms within. Meta-data needs to be protected, independently of any ability to participate.

Centralization: this is not a characteristic of the metaverse itself, rather a design choice on how to deploy such an application over some infrastructure. A logical administrative centralization is assumed, to ensure some global properties (say globally unique identities). However, the technical implementation will have to decide some trade-off between centralized vs distributed. This choice has an impact on the infrastructure and needs to be considered. Centralization of the metaverse, by hosting it on a specific set of servers and have clients connect to these servers, facilitates some aspects of the metaverse: for instance, it requires N connections (at the transport layer for an active session) between a client and a centralized server, where N is the number of users (clients); such centralization facilitate access controls, as per the "security" item above. It also facilitate the management of the user state. A fully distributed architecture that is fully meshed would require N^2 (potentially multicast) connections (for active sessions); further, these connections would need to be time-synchronized. However, the latency of a direct path would be faster than a triangular routing through a central server (depending on the respective positions of the participants within the network topology), and therefore the interactions could be quicker.

This list may be modified with other important dimensions based upon further discussion.

4. ICN Challenges

ICN Challenges for the Metaverse

ICN (cf [[ahlgren2012survey](#)] for a survey) is a novel network architecture that considers objects as the organizational principle for the network infrastructure. Instead of connecting to a server, ICN attempts to dissociate the objects from the server that could be hosting them, and attempts to route requests for an object directly to that object, rather than to a server that contains that object.

From the infrastructure perspective, a metaverse would be a distributed system that shares content in real time on a massive global scale, with some QoE requirements for the users, and in a secure way with complex ownership/access privileges/management of

trust. This is exactly the type of problem that ICN sets out to solve.

ICN has nice properties for fetching objects. In the case of Adaptive Video streaming, where a video stream is decomposed into chunks that correspond to a resolution and a time segment of the original video, fetching these objects directly is attractive. Popular videos will be distributed throughout the network, and therefore, can be encountered before getting it from a high level cache or an origin server.

RFC7933 considered the challenges of using ICN for video streaming and immersive streaming applications. The objective of this document is to similarly consider the impact of ICN on multiverse applications, and conversely, the requirements of multiverse applications on ICN architectures.

Some of the points discussed in RFC7933 can be straightforwardly mapped from video streaming to metaverse applications. For instance, the interactions of video streaming and ICN maps to interactions of metaverse applications on ICN; the integration of video streaming and ICN similarly maps to a possible integration of metaverse application with ICN. Encodings are also relevant in both contexts.

RFC7933 also discusses P2P video distribution, which could be map to a distributed embodiment of the metaverse. IPTV in RFC7933 brings up issues of multipath and multicast transport. Finally, digital rights management translate into how to manage access in the metaverse.

We can list some of the research challenges for the Metaverse in ICN: scalability; privacy/trust/security. Low-latency would be required for such an interactive real-time application. As a corollary, high precision transport layer, such as DetNet or any transport layer providing some latency guarantees, could be an interesting challenge.

Machine learning could help with identifying behaviors within a Metaverse to allow better operations (say by filtering out data that is unlikely to be consumed, or by pre-fetching data that is likely to be consumed; or by anticipating users' behavior). ML can be used to optimize various aspects of the system at multiple layers, and where both training and inference operations happen. Particularly for processing input video (as opposed to streaming it out), the placement of the computations strongly affect what the feasible topologies and deployment options are. There are of course also strong interactions with security and privacy.

One important question is the benefit overall of ICN for such an application in terms of sustainability. Is an ICN-supported Metaverse greener or more energy efficient than legacy applications?

Some other challenges are detailed a bit more in the following subsections.

4.1. Metaverse Objects

Objects in the metaverse have different properties than for typical ICN objects. Usually, as is the case in video streaming, an ICN object is fetched as part of a group of objects. The application layer then uses it as it chooses. A content owner provides keys for the user to access/decrypt the data.

It sees the metaverse semantics impose other requirements on the objects. In particular, the content ownership is more diffuse. There are several layers into it, with different access rules.

For instance, in a virtual world, there is a provider of a virtual world service, that would own objects pertaining to that virtual world infrastructure; this world is populated by people/avatars and objects that may want to control to whom they are visible; then they can use/create/purchase/sell virtual objects, granting new set of permissions to another layer of data.

Permissions to view, use, operate, modify, take or remove a virtual object becomes a more complex operation than just setting access rights. Then meta-data should be associated with virtual objects to keep track of events, accounting, transactions, etc, associated with that virtual object.

Data pertaining to a virtual world could be collected into a (potentially supercharged) FLIC collection, or grouped together using some form of manifest (similar to that used by DASH video streaming for instance). Other data however would need to perform several level of access controls; how to represent and organize such ownership levels, especially in a distributed manner, seems like an interesting challenge for ICN.

4.2. Centralization

Centralization vs Distributed is one of the dimensions in the taxonomy discussed above. If the metaverse is implemented as an application overlay, then it can be easily centralized. However, if the goal is to embed metaverse support into the network, then a decentralized implementation may be necessary.

An application overlay could be distributed, as with current CDNs for instance; or a network function could be centralized, as with SDN. Finding the right balance is an open challenge.

A hierarchical structure would be required to support the scale of such application. This yields questions and challenges regarding edge nodes running independently; or whether a part of the metaverse can keep running if disconnected from a central authority.

ICN decouples objects from the origin server. This is a step towards running a metaverse independently of a centralized server; however, can the whole application be decoupled from an origin server? The challenge would be to run Named Function Networking services for such an application.

4.3. Interoperability

A metaverse should interoperate along multiple dimensions. John Radoff lists as domains of interoperability: connectivity (networking, communications); persistence (identity, ownership, accounting, history); presentation (graphics models, physical properties); meaning (metadata, semantics, ontologies); behavior (rules, economies, consequences, power).

This document focuses only on the lower layer, connectivity. One key issue is to propose a common framework in ICN that would support interoperability for communications.

5. Use Cases

5.1. Moonshot Project

Japan Science and Technology Agency (JST) has launched a project, Moonshot Research and Development Program [[Moonshot](#)] (called Moonshot program, hereafter). The Moonshot program tackles important social issues, global climate change and extreme natural disasters, the Moonshot program is pursuing disruptive innovations in Japan and promoting challenging research and development based on revolutionary concepts. In one of the goals of the Moonshot program, one concrete research target is defined, which is "Reliability-ensuring cybernetic avatar infrastructure allowing interactive teleoperation [[Moonshot.TG](#)]". This research target consists of several sub-projects, such as "Development of Smart Spot Cell", "Local Network Intelligentization Algorithms", and "Reliable low-latency communications leveraged by information-centric networking technology".

"Reliable low-latency communications leveraged by information-centric networking technology" will be conducted research and development of ICNx, which is an extension/enhancement of ICN. ICNx

realizes highly reliable, low-latency, and highly efficient communications between humans (i.e., operators) and Cybernetic Avatars (CAs) over the wired and wireless networks. ICNx uses content identifiers (e.g., content names) for communication, exchanges data without relying on the locations of servers or clouds, and enables multicast that copies and transfers data within the network. It also provides a connectionless transport that improves throughput while suppressing congestion and a user-driven many-to-many multicast ($M \times N$ communication) that performs data transfer and sharing between multiple users and multiple CAs.

ICNx is a communication network architecture that provides interactive communications with the data transfer throughput of 100 Mbps and latency of 100 ms or lower, which are the requirements for a smooth communication between humans and CAs, in the condition that underlying wired and wireless networks have enough bandwidth. In order to achieve highly reliable communication, they will develop a mechanism to compensate for any data loss by transferring redundant data according to the network conditions and potentially considering advanced methods using network coding. To realize and verify the proposed technology, they will develop an ICNx communication platform using open-source software, called Cefore [[Cefore](#)]. Cefore complies with CCNx version 1 protocol messages specified in RFCs 8569 [[RFC8569](#)] and 8609 [[RFC8609](#)] published by the IRTF ICN Research Group, and runs on Linux (Ubuntu), macOS, and Raspberry Pi OS. In this project, they will develop an ICNx communication platform using Cefore, as well as Application Programming Interfaces (APIs) for ICNx applications.

6. Conclusions

This document attempts to present some of the challenges of supporting a Metaverse application within an ICN architecture. We presented a taxonomy and listed some of the challenges. This is an initial draft to initiate a discussion with the ICNRG.

7. IANA Considerations

This document does not have any IANA requests.

8. Security Considerations

No particular security considerations at this point.

9. Acknowledgments

We thank Dave Oran for providing valuable comments on an earlier version of this document.

10. Informative References (TBD)

[ahlgren2012survey]

Ahlgren, B., Dannewitz, C., Imbrenda, C., Kutscher, D., and B. Ohlman, "A survey of information-centric networking", IEEE Communications Magazine Vol.50 No.7, 2012.

[Cefore] "Cefore Home Page", <<https://cefore.net/>>.

[dwivedi2022] al, Y. K. D. E., "Metaverse beyond the hype: Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy", (Elsevier) International Journal of Information Management Vol. 66, Oct 2022, 2022.

[Moonshot] "Realization of a society in which human beings can be free from limitations of body, brain, space, and time by 2050.", <<https://www.jst.go.jp/moonshot/en/index.html>>.

[Moonshot.TG] "Reliability-ensuring cybernetic avatar infrastructure allowing interactive teleoperation", <<https://ca-platform.nict.go.jp/en/index.html>>.

[RFC7933]

Westphal, C., Ed., Lederer, S., Posch, D., Timmerer, C., Azgin, A., Liu, W., Mueller, C., Detti, A., Corujo, D., Wang, J., Montpetit, M., and N. Murray, "Adaptive Video Streaming over Information-Centric Networking (ICN)", RFC 7933, DOI 10.17487/RFC7933, August 2016, <<https://www.rfc-editor.org/info/rfc7933>>.

[RFC8569] Mosko, M., Solis, I., and C. Wood, "CCNx Semantics", RFC 8569, July 2019, <<https://www.rfc-editor.org/rfc/rfc8569>>.

[RFC8609] Mosko, M., Solis, I., and C. Wood, "CCNx Messages in TLV Format", RFC 8609, July 2019, <<https://www.rfc-editor.org/rfc/rfc8609>>.

Authors' Addresses

Cedric Westphal
Futurewei

Email: cedric.westphal@futurewei.com

Hitoshi Asaeda
National Institute of Information and Communications Technology
4-2-1 Nukui-Kitamachi, Koganei,
Tokyo 184-8795
Japan

Email: asaeda@nict.go.jp