THETA

NETWORK

Decentralized video streaming, powered by users and an innovative new blockchain.

WHITEPAPER

A Decentralized Video Streaming Network, Powered by a New Blockchain and Token

Last Updated: Jan 15, 2018
Version 1.6

Abstract

This whitepaper introduces a new Decentralized Streaming Network (DSN) and Theta, a new blockchain and token as the incentive mechanism for the DSN.

DSN and the Theta protocol solve various challenges the video streaming industry faces today. First, Theta tokens are used as an incentive to encourage individual users to share their redundant memory and bandwidth resources as caching nodes for video streams. This improves the quality of stream delivery and solves the "last-mile" delivery problem, the main bottleneck for traditional stream delivery pipelines, especially for high resolution high bitrate 360° virtual reality (VR) streams. Second, with sufficient amount of caching nodes, the majority of viewers will pull streams from peering caching nodes. This significantly reduces content delivery network (CDN) bandwidth costs, which is a major concern for video streaming sites. Lastly, the Theta network greatly improves the streaming market efficiency by streamlining the video delivery process. For example, advertisers can target end viewers at a lower cost and reward influencers more transparently.

The Theta blockchain introduces three novel concepts:

> **Reputation Dependent Mining:** In the Theta protocol, the caching nodes play the role of miners in the blockchain. The block reward is not a constant, but depends on the reputation score of the caching node that mined the block. To obtain more mining rewards, miners not only spend computation power to mine blocks, but also relay video streams to downstream viewers to increase their reputation scores.
>
> **Global Reputation Consensus:** We propose a mechanism for the Theta network to reach the global consensus on the reputation scores for each caching node.
>
> **Proof-of-Engagement:** We introduce a novel Proof-of-Engagement scheme to prove that viewers legitimately consume the video streams, providing better transparency to advertisers and a basis for viewers to earn Theta tokens for engaging with the content

This white paper will describe these concepts and the Theta blockchain in detail. Phase I of the Theta network will launch with the issuance of ERC20-compliant tokens during the token sale. Phase II encompassing a new blockchain and native Theta tokens is planned to launch in Q4 2018, at which time each ERC20 Theta token can be 1:1 exchanged for a native Theta token.

TABLE OF CONTENTS

Vision

Introduction

Video streaming market

Live video streaming accounts for over two-thirds of all internet traffic today, and it is expected to jump to 82% by 2020, according to Cisco's June 2016 Visual Networking Index report.[1] In the US, millennials between the ages of 18 and 34 are driving the growth of video streaming, and are heavy users of services like Instagram, Spotify and Snapchat. Streaming video among this group has jumped 256% from an average of 1.6 hours per week to 5.7 hours per week according to a SSRS Media and Technology survey, and mobile devices are leading the charge in video consumption growing 44% in 2015 and 35% in 2016.[2] The top five video streaming players in the US are Facebook, Google/Youtube, Twitter (and related properties), Live.ly and Twitch.

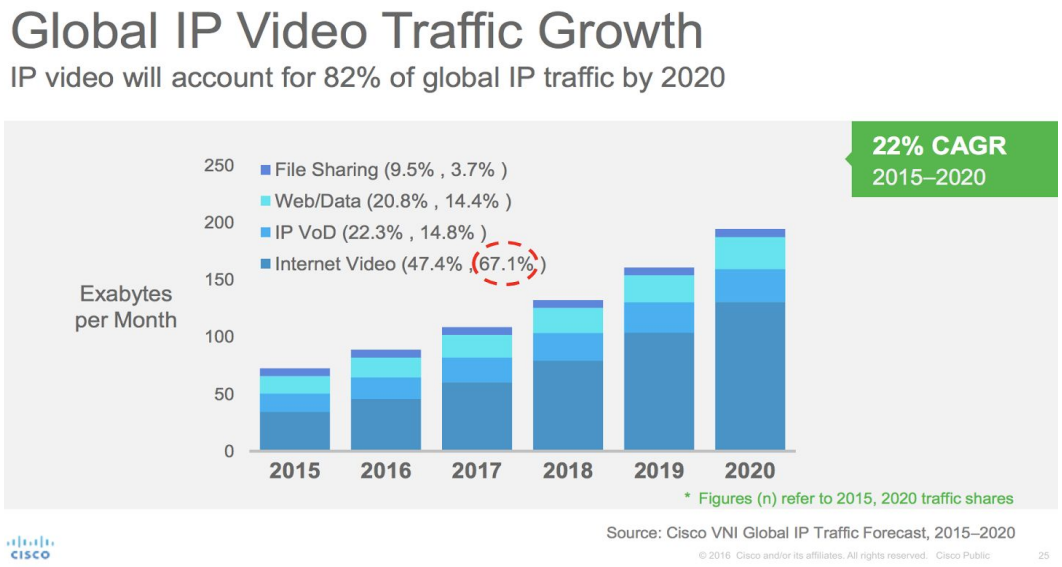

Figure 1. Global IP video traffic growth

More importantly, global virtual reality (VR) traffic including 360° video streaming content is estimated to grow 61-fold by 2020, at a staggering 127% CAGR according to the same Cisco report.

[1] https://www.cisco.com/c/dam/global/ko_kr/assets/pdf/2016-VNI-Complete-Forecast-PT.pdf
[2] http://www.zenithmedia.com/mobile-drive-19-8-increase-online-video-consumption-2016/

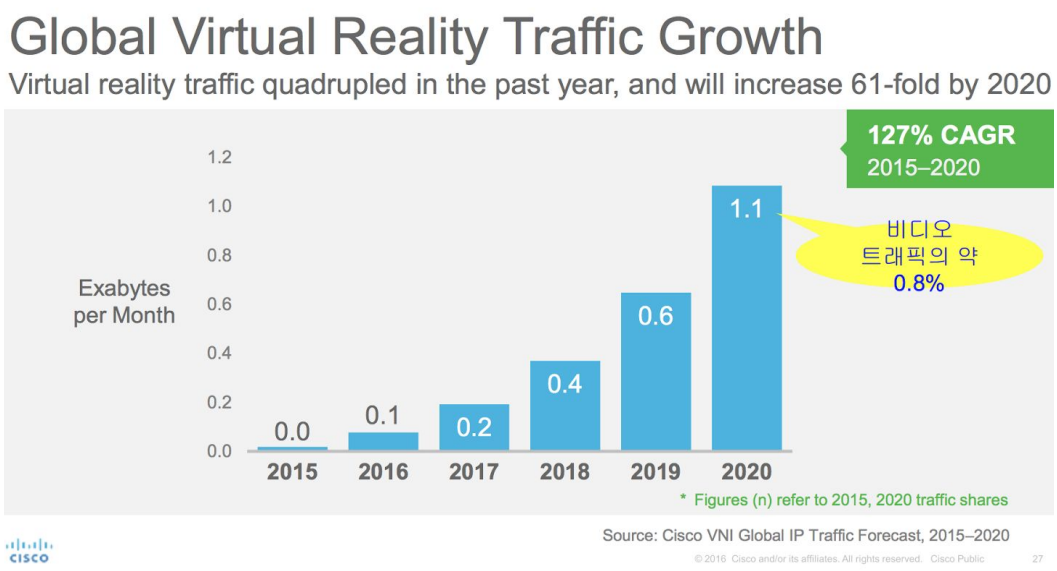

Figure 2. Global virtual reality traffic growth

Video streaming challenges

Content Delivery Network (CDN) plays an important role in the video streamin[3]g ecosystem. It provides the backbone infrastructure to deliver the video streams to the end viewers. One major limitation of today's CDN networks is the so-called **"last-mile" delivery problem**. Typically, the CDN providers build data centers called Point-of-Presences (POPs) in many locations around the world, with the hope that these POPs are geographically close to the viewers. However, the number of POPs are limited, hence cannot be close enough to many viewers, especially in less economically developed regions. This "last-mile" link is usually the bottleneck of the streaming delivery pipeline and often leads to undesirable user experience including choppy streams and frequent rebuffering.

To the streaming sites, another major concern is the CDN bandwidth cost. For popular sites, the CDN bandwidth cost can easily go up to tens of millions of dollars per year. Even if platforms own proprietary CDNs, maintenance costs are often high.

These issues are becoming even more prominent with the coming era of 4K, 360° VR streaming, and upcoming technologies such as light field streaming. Table 1 compares the bandwidth requirement of today's mainstream 720p/HD stream with 4K, 360° VR and future lightfield streams. It is clear that the bandwidth requirement increases by orders of magnitude.

[3]https://www.cisco.com/c/dam/global/ko_kr/assets/pdf/2016-VNI-Complete-Forecast-PT.pdf

Standard	Resolution	Bandwidth / Mbps	Magnitude
720p HD	1080x720	5 to 7.5	1x
1080p HD	1920x1080	8 to 12	1.6x
4K UHD	3920x2160	32 to 48	6.4x
8K 360° VR	7840x4320	128 to 192	25x
16K 360° VR	15680x8640	512 to 768	100x
Lightfield	---	5000+	1000x

Table 1. Bandwidth comparison: today's 720p/1080p video vs 4K and 360° VR streaming, vs future volumetric/lightfield streaming

To solve the VR and light field video delivery problem, the industry has started to explore "**foveated streaming**" technology. Instead of streaming the entire video in full resolution, this technology reduces the image quality in the peripheral vision (outside of the zone gazed by the fovea) in order to reduce bandwidth requirement. As the viewer turns his or her head to look at a different direction, the system adapts the spatial video resolution accordingly by fetching the high resolution video packets for the viewing direction from the server. For the foveated streaming technology to work well in practice, the round-trip time between the server and the viewer has to be small enough. For the viewers that are geographically far away for the CDN POPs, their VR stream viewing experience will be compromised even with foveated streaming technology.

Background

SLIVER.tv (the "company") has been at the forefront of developing next-generation video streaming technologies for VR and spherical 360° video streams since 2015, and has been integral in the founding of the Theta network. SLIVER.tv has raised over $17 Million in venture financing from notable Silicon Valley VCs including Danhua Capital, DCM, Sierra ventures, leading Hollywood media investors including Creative Artists Agency, BDMI, Advancit Capital, Greycroft Gaming Track Fund, and marquee corporate investors including GREE, Colopl, Samsung Next and Sony Innovation funds. Additionally, the company has strong Chinese investors and partners including Heuristic Capital Partners, ZP Capital, Green Pine Capital Partners, and Sparkland.

In a technology derived from "foveated streaming" SLIVER.tv's most recent technology **patent pending #62/522,505**, "METHODS AND SYSTEMS FOR NON-CONCENTRIC SPHERICAL PROJECTION FOR MULTI-RESOLUTION VIEW", *specifically addresses the problem of generating highly efficient spherical videos for virtual reality (VR) streaming, highlight, and replay. The technology performs non-concentric spherical projection to derive high resolution displays of selected important game actions concurrently with lower resolution displays of static game environments, thus optimizing tradeoff between visual fidelity and data transfer load.*

SLIVER.tv is the leading next-generation live esports streaming platform, with a vision to transform the esports spectating experience. As video games have grown in popularity to become a $40+ billion market, bigger than Hollywood and Bollywood combined, the rise of multiplayer competitive video gaming as a spectator sport has become a major new industry, dubbed **esports**. Esports is a global phenomenon with major tournaments and major pockets of fans and competitive teams in Europe, Asia and North America. The online gaming and esports ecosystems have exploded over the past five years.

A recent 2017 SuperData research[4] put the combined audience for gaming video content on YouTube and Twitch at 665 million, more than twice the US population. This surpasses the viewership of 227 million for HBO and Netflix combined. Today, esports and gaming video content account for a significant portion of all video content streamed over the Internet.

SLIVER.tv additional core patents and technology focus on various applications of cutting edge live streaming to esports content. The company's US Patent **#9,573,062** "*METHODS AND SYSTEMS FOR VIRTUAL REALITY STREAMING AND REPLAY OF COMPUTER VIDEO GAMES*"[5] and **#9,473,758** "*METHODS AND SYSTEMS FOR GAME VIDEO RECORDING AND VIRTUAL REALITY REPLAY*"[6], pioneer the capture and live rendering of the most popular PC esports games including League of Legends, Dota2 and Counter-Strike: Global Offensive in a fully immersive 360° VR spherical video stream, effectively placing the viewer and audience inside the 3D game through a live video stream.

Since launching last year, SLIVER.tv has broadcast numerous global esports tournaments in 360° VR in partnership with premier brands including ESL One, DreamHack and Intel Extreme Masters[7]. At key events in the US and Europe, SLIVER.tv has live streamed top esports games to millions of fans of *Counter-Strike: Global Offensive (CS:GO) and League of Legends (LoL).*

SLIVER.tv launched its Watch & Win esports platform in July 2017 and the first virtual token designed around esports content streaming and fan engagement. Since launch, the company has attracted millions of esports fans circulating over 1 Billion virtual tokens by actively participating and engaging with live esports matches. These users viewed over 50 million minutes of live esports streaming, nearly 100 years worth of content in a few weeks. This positions the company as one of the largest esports streaming sites built around a virtual community today.[8]

SLIVER.tv platform is continuing to expand quickly, with over 50% month-to-month growth rate in unique visits driven by word-of-mouth, referral and social channels.

[4] https://www.superdataresearch.com/market-data/gaming-video-content/
[5] https://www.google.com/patents/US9573062
[6] https://www.google.com/patents/US9473758
[7] https://www.sliver.tv/events
[8] https://www.sliver.tv/press

Traffic Overview

Total visits on desktop and mobile web, in the last 6 months

Opportunity

The company's mission is to leverage the decentralized blockchain structure to create the first **Decentralized Streaming Network (DSN)** whereby video viewers are incentivized to share redundant memory and bandwidth resources to address today's video streaming challenges. Using the Ethereum EVM "World Computer" metaphor, the DSN can be viewed as the "**World Cache**" formed by the memory and bandwidth resources contributed by viewers.

Specifically, viewers around the globe can contribute their computers as "**caching nodes**" whereby they form a video delivery infrastructure that is responsible for delivering any given video stream to viewers anywhere around the world. The DSN can effectively address the technical challenges discussed in the previous section. First, viewers' computers are geographically much closer to each other than to the CDN POPs. This reduces packet round-trip time and improves the stream delivery quality, and thus addresses the "last-mile" delivery issue. Second, with sufficient amount of caching nodes, most viewers will receive the stream from caching nodes, this will help streaming sites reduce their CDN bandwidth cost. Third, caching nodes also reduce round-trip time making foveated streaming technology practical.

To encourage viewers to contribute their memory and bandwidth resources, we introduce **the Theta protocol** as an incentive mechanism. The caching nodes can mine new tokens as they relay video streams to other viewers. Not only does Theta tokens motivate viewers to join the network as caching nodes, but it also greatly *improves the streaming market efficiency by*

streamlining the video delivery process. We will discuss more details later in the paper, but within the Theta network, advertisers can directly target viewers at a lower cost, viewers earn Theta tokens for their attention and engagement with video streams, and influencers earn Theta token as gifts directly from viewers.

The full launch of the Theta protocol in phase II introduces a **new blockchain** and a **native token** structure where:

- Advertisers spend tokens to support influencers, streaming sites and viewers
- Caching nodes mine new tokens for caching and relaying video streams
- Viewers optionally earn tokens from advertisers as engagement rewards

The Theta protocol builds upon the following novel concepts:

- **Reputation Dependent Mining:** In the Theta protocol, *the caching nodes play the role of miners in the blockchain*. Similar to other blockchains, the caching nodes validate transactions and attempt to add new transaction blocks to the blockchain based on Proof-of-Work. However, in our protocol, the block reward is not a constant. *The miners with higher **reputation score** (intuitively this means that they serve viewers more effectively) will have a higher block reward*. To obtain more mining rewards, miners not only spend computation power to mine blocks, but also relay video streams to downstream viewers to maintain and increase their reputation scores. In other words, they are essentially "functional miners".

- **Global Reputation Consensus:** We propose a scalable mechanism for the Theta network to reach the global consensus on the reputation scores for each caching node, and thus the global consensus on the amount of mining reward for each block.

- **Proof-of-Engagement:** We devised a novel Proof-of-Engagement (PoE) scheme to *ensure that viewers legitimately consume the video streams*. Viewers can earn tokens as rewards from advertisers in exchange for their attention to video streams and by providing Proof-of-Engagement. PoE not only brings benefits to viewers, but also to advertisers as it provides them with a reliable and verifiable engagement measurement of the delivered video streams.

Decentralized Streaming Network

Overview

Leveraging our experience and expertise in high resolution high bitrate video streaming, we are proposing a decentralized video stream delivery technology for both video on demand and live streaming. Figure 3 depicts the high level architecture.

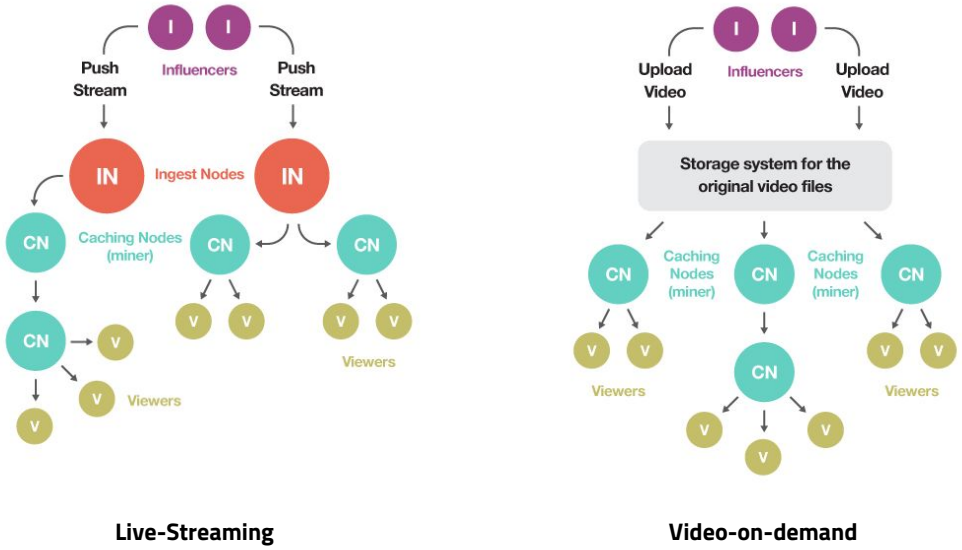

Figure 3. ***Left*** *- DSN architecture diagram for live streaming*
Right *- DSN architecture diagram for video on demand*

Figure 3 Left shows the DSN architecture for **live streaming**. The **influencers (aka streamers)** publish their video streams to the ingest nodes. **Ingest nodes** run on computers contributed by the user community. The ingest nodes are responsible for transcoding the video stream to different bitrates and resolutions. Then, the **caching nodes** pull the video streams and relay to end viewers. Any user can contribute his or her computer as a caching node by running a special video caching software client we are developing. The video caching software is responsible for choosing the upstream node, and relaying the video stream to the downstream nodes. This is similar to traditional P2P file sharing systems like BitTorrent, but with an important extra constraint on latency due to the realtime nature of live streaming. Rather than sending chunks of a file out-of-order as in P2P file sharing systems, more urgent high-priority packets will be sent first. Furthermore, to minimize re-buffering, each caching node may cache the entire video file instead of storing only a portion of a file as P2P file sharing systems do.

For ease of explanation, the viewer nodes and the caching nodes are drawn as separated nodes in Figure 3. However, it is worth pointing out that in practice, a user can run both the caching software and viewer software on the same computer.

Figure 3 Right shows the DSN architecture for **video on demand**. Compared to live streaming, the only difference is that the **influencers** (aka **content creators**) upload the source video file to

a **storage system** for later retrieval, rather than publishing the stream to the ingest nodes. The storage system can be a decentralized file system such as IPFS[9] or SWARM[10], or cloud-based storage like AWS S3. The rest of the network is essentially the same as in the live streaming scenario.

For any viewer and caching node to receive a stream with high quality and low latency, its physical proximity to their peer caching nodes is essential. Thus, when a new node joins the DSN network, it goes through a bootstrapping process to discover the caching nodes that are physically nearby. P2P file sharing systems typically employ Distributed Hash Table (DHT) such as Kademlia to find peer nodes. In Kademlia, each node is assigned a random GUID, and the distance between the two nodes is calculated as the XOR of their GUIDs. While this distance measure is sufficient for non-realtime applications such as file sharing, in the context of video streaming, it is important for the nodes to identify peering nodes that are geographically close. To achieve this, we propose to encode the geographic information into the GUID as shown below

$$GUID = QuadtreeRegionID \,||\, RandomBitsPostfix$$

The GUID is the concatenation of two parts. The first part is the $QuadtreeRegionID$, a 64 bit long string that encodes the geographical region of the caching node. The Quadtree data structure is employed as the spatial index. The second part is a 32 bit random bit postfix to differentiate caching nodes in the same quadtree region. This way the XOR of two GUIDs is a rough estimation of the geographical distance between two nodes. Thus a new node can first collect a list of candidate peer nodes with small GUID XOR distances. Then, it can send a ping request to the candidates to measure the actual round trip time.

[9] https://ipfs.io/
[10] https://github.com/ethereum/go-ethereum/wiki/Swarm---distributed-preimage-archive

Bootstrapping The Network

The above discussion assumes a fully decentralized architecture where all the ingest and caching nodes come from users. To realize this long term vision, we need a sufficient number of users to participate, which may take some time. To **bootstrap** the network, we propose a **hybrid architecture** that works with the existing CDN networks and the decentralized caching nodes.

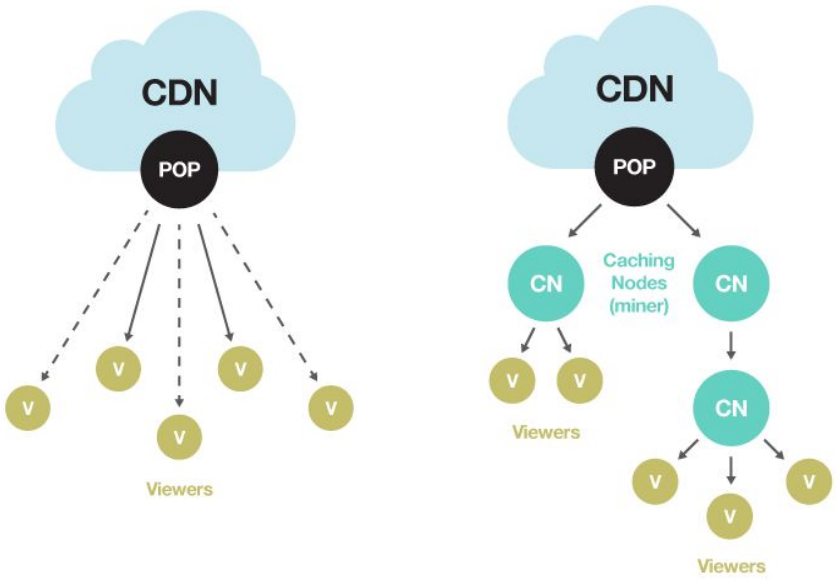

Traditional CDN **Hybrid Architecture**

Figure 4. ***Left - Traditional CDN****, where all the viewers connect directly to the POP servers. For nodes geographically far away from the POP servers, the stream quality may be lower (the dashed connections).* ***Right - Hybrid Architecture*** *where viewer nodes can pull stream from caching nodes that are geographically closer than the POP servers, resulting in better streaming quality.*

Figure 4 illustrates the hybrid architecture and compares it with traditional CDN networks. Whereas in the traditional CDN, every node pulls the stream directly from POP servers, in the hybrid architecture, only a subset of nodes pull the stream from POP servers. Other nodes simply pull the stream from their peer caching nodes which provide better and more effective connection. The caching nodes thus augment the traditional CDN backbones with more caching layers for end viewers geographically far away from the POPs. We note that this hybrid architecture applies to both video on demand and live streaming scenarios since traditional CDN networks work similarly.

The advantage of bootstrapping with the hybrid architecture is important. **The DSN is fully functional even with only a single caching node**. Furthermore, as will be discussed later, the caching node can mine new Theta tokens while providing the caching and video relaying services. **This provides strong economic incentives for more users to join the caching network over time**. As more user contributed nodes join the DSN, the stream delivery capability of the network improves. At the point when there are sufficient amount of user contributed nodes, the network can run on its own without the traditional CDN backbone.

How Theta Tokens Work in the DSN

Theta tokens work as a long-term sustainable incentive mechanism to motivate various stakeholders to participate in the DSN. It is tempting to design the incentive mechanism based on the traditional peer-to-peer model where the end viewers send Theta tokens to caching nodes in exchange for the stream relay service. However, such a scheme is impractical since ad-supported viewers today on Twitch or YouTube, for example, are used to free video streams. **We need an incentive mechanism to motivate the caching nodes to relay the video stream in absence of Theta tokens collected from viewers.** For this purpose, we propose to have the caching nodes play the roles of Theta miners. Similar to miners in other blockchains, they validate Theta token transactions and assemble new blocks to obtain mining rewards. However, as detailed later in the whitepaper, we design the mechanism such that for a caching node, **the more video stream relaying work it conducts, the higher mining rewards it gets**. This motivates the caching nodes to relay video streams in addition to mining new blocks.

In this section, we will start by listing the stakeholders in the streaming ecosystem. We then describe how the Theta tokens flow among them and tie them together. Finally, we discuss how Theta improves the efficiency of the streaming market in the DSN.

Video Streaming Stakeholders

The decentralized video streaming ecosystem has the following key stakeholders:

- **Caching Nodes:** They provide caching services to improve the video stream delivery. They are rewarded for the caching service they provide to the end viewers.

- **Viewers:** They are the end consumers of the video content. Viewer attention is the most important asset in the video streaming ecosystem. Viewers can be optionally rewarded for their attention and engagement with the video stream and with advertisements[11].

- **Ingest Nodes:** They provide the ingest service for the live streams, generating various stream resolutions, bitrates, etc. They are rewarded for providing this service to the caching nodes.

- **Influencers/Streaming Sites:** They are the producers of the video content. In the live streaming context, these influencers are also known as streamers. In the video on demand context, influencers are also referred to as content creators. They will be rewarded with Theta tokens for producing the content.

- **Advertisers:** The goal of advertisers is to promote their products and services to end viewers. To achieve this, advertisers will spend Theta tokens as part of their advertising budget to sponsor influencers and video streams through video ads including pre-rolls, mid-rolls and other formats.

[11] *Brave Software*, Basic Attention Token (BAT) - Blockchain Based Digital Advertising

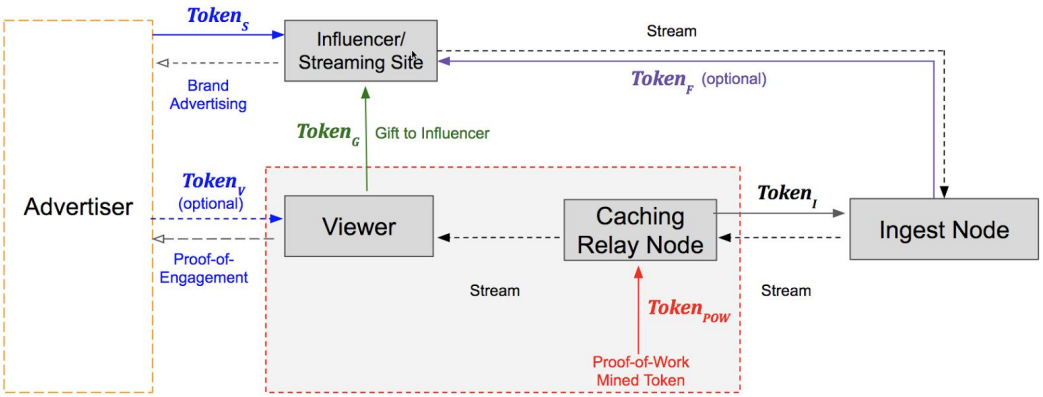

Figure 5. How Theta tokens works in the DSN

Token Flow Among Stakeholders

Caching nodes play the roles of miners in the blockchain. As they validate the transactions and add new blocks to the chain, tokens $\boldsymbol{Token_{POW}}$ are awarded to them. In the Theta blockchain, the block reward is not a constant. Miners who provide better services to viewers will have higher block reward when they mine a new block. This motivates the caching nodes to relay video streams to more viewers and more efficiently. We will dive into more detail later in the paper.

Note that in practice, a user might watch a video stream while running the caching node software on his or her computer. In such a case, **the user's computer is both a viewer and a caching node**. This allows the viewers to earn mining rewards as a caching node, and then optionally gift to influencers. This is represented by the dashed red rectangular box in Figure 5.

In the live streaming scenario, caching nodes pull the stream from the ingest nodes and in return they reward $\boldsymbol{Token_I}$ to the ingest nodes. And optionally the ingest nodes tip $\boldsymbol{Token_F}$ to influencers.

For the ad-supported streams, advertisers typically allocate a marketing budget, $\boldsymbol{Token_A}$, as shown in Figure 5. A majority of this budget, $\boldsymbol{Token_S}$, is spent to sponsor influencers or streaming sites. The remaining portion, $\boldsymbol{Token_V}$, is optionally used as an engagement reward for viewers of the stream, namely:

$$Token_A = Token_S + Token_V \quad (1)$$

Influencers and streaming sites in turn earn $\boldsymbol{Token_S}$, and optionally $\boldsymbol{Token_G}$ which are gifts from their fans and viewers. Viewers are rewarded with $\boldsymbol{Token_V}$ once they prove they have watched and engaged with the live stream content through Proof-of-Engagement.

Streaming Market Efficiency Improvement

As shown in Figure 5, Theta token ties the stakeholders of the streaming ecosystem together to bring more transparency and better market efficiencies. Below we list several examples illustrating the benefits Theta brings to the ecosystem.

- Most importantly, the Theta network increase influencers' rewards. Advertisers can sponsor influencers directly, and viewers and ingest nodes can additionally gift Theta tokens to influencers.

- Second, the Theta network increases transparency to advertisers. Specifically, the "Proof-of-Engagement" to be introduced shortly provides a probabilistic proof that a certain number of viewers have watched a given video stream. The proof is publicly auditable, unlike today's streaming ecosystem where such data is only made available by the streaming sites.

- Finally, the Theta network allows advertisers to optionally reward viewers ($\boldsymbol{Token_V}$ in Figure 5) for their attention to the stream and advertisements. Rewarding viewers for their attention could result in an improvement in viewer engagement and retention.

The Theta Blockchain

Overview

The Theta blockchain is similar to the Bitcoin[12] blockchain in many ways. Caching nodes are the miners for the Theta blockchain. The Theta network implements a transaction ledger via a blockchain. Each block consists of a set of transactions of Theta tokens. To add a new block to the blockchain, a miner needs to validate the transactions and solve a Proof-of-Work based cryptographic hash puzzle. The Proof-of-Work parameter is dynamically adjusted so that one block is added to the chain in a roughly fixed amount of time, similar to Bitcoin. Furthermore, as many other blockchains, a caching miner node should consider a transaction as committed after a sequence of valid blocks have been added after his block.

The Theta blockchain also introduces three novel concepts: **1) Reputation Dependent Mining, 2) Global Reputation Consensus, and 3) Proof-of-Engagement.** Below we will discuss each of these concepts in detail.

Reputation Dependent Mining

The caching nodes in the DSN are the miners for the Theta blockchain. They validate transactions and attempt to add new blocks to the blockchain by solving a Proof-of-Work (PoW) hash puzzle. However, for the DSN to function properly, in addition to puzzle solving, the caching nodes need to relay video streams to end viewers.

In order to incentivize them to relay the video stream, in the Theta protocol, the mining reward of a block depends on the **reputation score** of the caching node that mined the block:

$$reward = (1 + r) \cdot reward_{base} \quad (2)$$

Here $reward_{base}$ is a predefined constant shared by all the caching nodes. And r is the reputation score, which is a real number between 0 and 1. **Hence, a caching node with a higher reputation score would have a higher block reward**. The reputation score of a caching node characterizes the contribution of the caching node to the streaming ecosystem. Essentially, a caching node that *relays more video streams **more recently*** will have a higher reputation score.

The reputation score of a caching node is calculated based on the service certificates it collects from its downstream nodes. The concept of "service certificate" is inspired by the "pooled mining" approach[13] [14]. In pooled mining, the miners submit "near-miss solutions" to the master on a regular basis to prove that they are constantly trying to solve the PoW mining puzzle for the master. In the Theta network, as a caching node relays the video stream to its downstream viewers, it will ask its viewers to solve a "service certificate PoW puzzle" and return their puzzle solutions. These puzzle solutions are the service certificates which prove that the caching

[12] *Nakamoto*, Bitcoin: A Peer-to-Peer Electronic Cash System
[13] *Luu et al.* SMARTPOOL: Practical Decentralized Pooled Mining
[14] https://en.bitcoin.it/wiki/P2Pool

node has provided service to the viewers for a certain amount of time. If a viewer fails to provide the puzzle solutions, the caching node can choose to terminate the stream. When the streaming session ends, the viewer simply stops puzzle solving. In other words, the viewers offer their computational power in exchange for the video stream relay service.

The **service certificate PoW puzzle** is formally defined as the following minimization problem:

$$nonce^* = argmin\ HASH(pk_i \,||\, pk_j \,||\, block\_hash_k \,||\, nonce) \tag{3}$$

Here || is the string concatenation operator. pk_i and pk_j are the public key of caching node *i* and viewer node *j*, respectively. $block\_hash_k$ is the hash of block k, the most recently confirmed block in the current longest chain. The viewer node *j* needs to return the $nonce^*$ that attains the **minimal hash value** it found before the next block is added to the longest chain. Whenever a new block is confirmed, the caching node sends a new puzzle to its downstream viewers, where in the new puzzle, the $block\_hash_k$ parameter in Formula (3) to point to the newly added block. The caching node also asks the viewers to submit their solutions $nonce^*$ for the previous puzzle. The HASH function is an ASIC resistant hash function to reduce the advantage of special-purpose mining hardware. Candidates includes the memory hard hash functions that were employed by Ethereum or ZCash.

Node *i* will need to keep the record of the puzzles and the nonces, so it can justify its reputation score to the other caching miner nodes when it mines a new block. For this purpose, it is sufficient to store the **service certificates**, which is formally defined as the following tuple

$$sc = (pk_i,\ pk_j,\ block\_hash_k,\ nonce^*) \tag{4}$$

Intuitively, a service certificate is a probabilistic proof that the caching node *i* has served a downstream node *j* for a certain amount of time. Below we provide the analysis.

Denote the minimal hash value the viewer *j* found as

$$H^* = HASH(pk_i \,||\, pk_j \,||\, block\_hash_k \,||\, nonce^*) \tag{5}$$

We can estimate the amount of time that caching node *i* served viewer *j* during the time block b was the most recent block using the following formula

$$T_{SC} = min(E[n \mid H^*] / s,\ T_b) \tag{6}$$

Here s is the number of hashes that the caching node requires the viewer node to perform per unit time. $\mathbf{E}[n \mid H^*]$ is the expected number of hashes to get a hash value smaller than or equal to H^*, which can be calculated using the Bernoulli trial model

$$E[n \mid H^*] = \sum_{k=1}^{\infty} k \cdot (H^*/M) \cdot (1 - H^*/M)^{k-1} = M / H^* \tag{7}$$

where M is the max possible hash value. T_b is amount of time that block b is at the tail of the longest chain. It can be estimated using the average amount of time needed to mine a new block, which can be regarded as a constant.

Now we can define of the reputation score *r* as follows

$$r = f(T_{total}) = 2 / (1 + e^{-\beta \cdot T_{total}}) - 1 \quad (8)$$

where T_{total} is the total amount of time node i served its downstream nodes in a **recent time window**. This time window is a constant shared among all caching miner nodes, e.g. the past 24 hours. β is a constant scaling factor shared by all the miners. T_{total} can then be calculated as the sum of T_{sc} given the service certificates the node collected within the time window, as shown in Formula (9).

$$T_{total} = \sum_{sc} T_{sc} \quad (9)$$

Function $f(T_{total})$ is a monotonically increasing function which maps T_{total} to a real number in the [0, 1] range as shown in Figure 6. When T_{total} is 0, $f(T_{total})$ is 0. As T_{total} approaches infinity, $f(T_{total})$ approaches 1. Thus, the more video relay service the caching node delivers within the recent time window, the higher reputation score and hence higher block reward it can obtain.

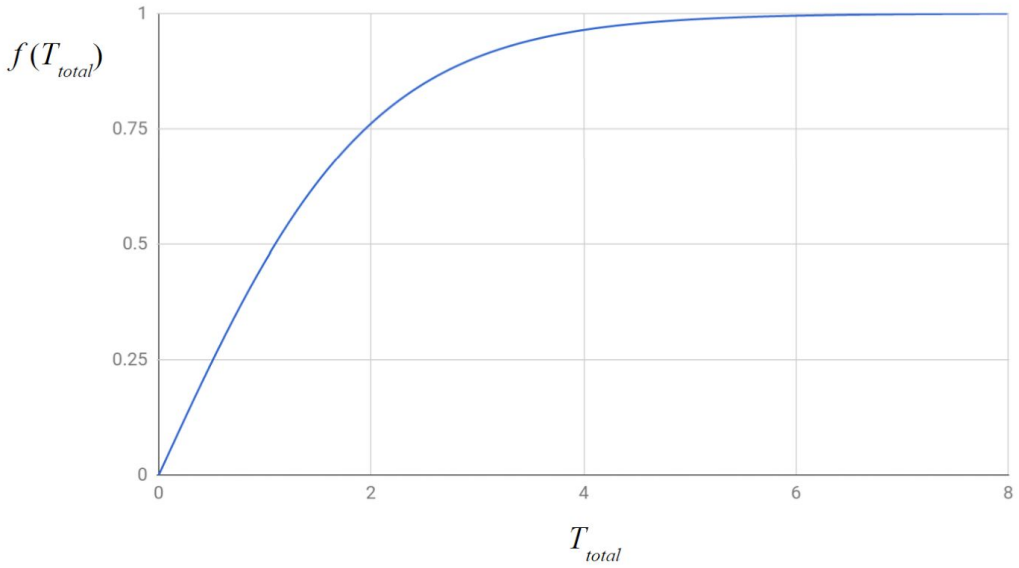

Figure 6. Function $f(T_{total})$ maps T_{total} to a real number in the [0, 1] *range*

One interesting property of the function $f(T_{total})$ is that if the parameter β is chosen properly **it actually discourages malicious miners from making fake service certificates.** In the Appendix, we will show that if β is proportional to the total network hashing power (similar to Bitcoin where the difficulty level of the mining puzzle scales with the total hashing power), forging service certificates actually **reduces** the expected mining reward. Intuitively, to forge service certificates, a malicious miner needs to spend hashing power. As a result, its chance of mining the next block reduces. Even though the fake service certificates it generated can increase block reward, but with properly chosen β, the expected mining reward per unit time **strictly** reduces. Thus, a rational miner would not try to forge service certificates.

We would like to emphasize again that Formula (9) should only account for the service certificates generated in the recent time window. The reason that we account for **recency** is two-fold: Firstly, this encourages the caching nodes to keep relaying the video streams. It cannot simply rely on the video relaying work done in the distant past to obtain high reputation score. Since a caching miner node cannot predict when it will mine the next block, it will be always motivated to relay the video streams and collect the service certificates. Secondly, it reduces the amount of service certificates that need to be validated by other caching miner nodes. This, in turn, reduces the bandwidth overhead of the Theta protocol.

To determine whether a service certificate is generated in the defined time window, one can leverage the $block\_hash_k$ parameter in the service certificate. With the $block\_hash_k$ one can look up the corresponding block from the blockchain and obtain its index in the longest chain. Since the time required to mine each new block is relatively a constant, one can determine whether the corresponding block, and hence the service certificate were generated within the given time window.

Once a caching node mines a new block, it propagates the block to other caching nodes in the network for validation. For a caching node to validate a new block, in addition to verifying the transactions and the PoW puzzle solution, it needs to verify that the reputation score is computed correctly based on the service certificates. In the following section, we will discuss how the global consensus on the reputation score can be formed.

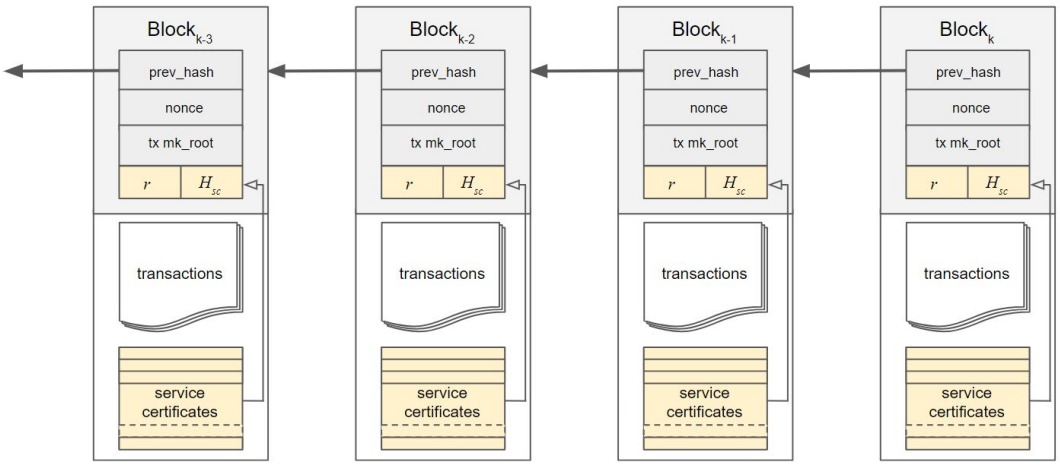

Figure 7. Extending the blockchain to facilitate global reputation consensus

Global Reputation Consensus

For ease of discussion, below we will refer to the caching miner node that proposed the new block as the "proposer miner", and the nodes that validate the new block as the "**validator miners**".

In order to facilitate the validator miners to validate the reputation score, we extend the blockchain structure with the following additional fields.

First, we add 1) the reputation score *r* and 2) the following hash to the block header

$$HASH(T_{total} \parallel sc_1 \parallel sc_2 \parallel ... \parallel sc_r) \quad (10)$$

where sc_1, sc_2, ..., sc_r are the service certificates the proposer miner collected within the recent time window. Second, we include the list of service certificates in the block body.

Figure 7 depicts the extended blockchain structure with the additional fields.

When a validator miner receives a newly proposed block, it needs to perform the following verifications:

- The following PoW mining puzzle inequality holds

$$HASH(block\_header) \leq M / difficulty$$

  Similar to Bitcoin, the $block\_header$ includes the hash of the previous block, the nonce, the public key of the proposer miner, and other metadata. In addition, as mentioned above, the $block\_header$ also includes the reputation score r of the proposer miner and the hash value defined by Formula (10). This prevents potential tampering of the reputation score and the hash value.

- All the Theta token transactions are valid. The validity conditions are the same as Bitcoin.

- The validator miner verifies the reputation score r is computed correctly based on the service certificates included in the block body. More specifically, it needs to verify the following conditions
  - The $block\_hash_k$ parameter in each certificate must point to a block that was generated within the defined recent time window.
  - There is no duplicated certificates, i.e. two certificates should not have the same (pk_i, pk_j, $block\_hash_k$).
  - The hash value calculated by Formula (10) matches with the corresponding hash value in header of the proposed block.
  - The reputation score r is computed correctly given the service certificates using Formula (4) to (9).
  - The Theta token amount specified in the coinbase transaction in the new block equals $(1 + r) \cdot reward_{base}$, as specified by Formula (2).

If the newly proposed block passed the above checks, the validator miner can confirm the block to the blockchain.

Support For Pooled Mining

Pooled mining[15] is a mining approach where multiple miners contribute to the generation of a block, and then split the block reward according to the contributed hashing power. Pooled mining effectively reduces the variance of mining reward for the participant miners.

On the Theta blockchain, miners can also form pools. Similar to Bitcoin mining pools, the miners solve the PoW mining puzzle on behalf of the master. In addition, their downstream viewers generate the service certificates for the master, by using the public key of the master as pk_i in the service certificate PoW puzzle given by Formula (3).

[15] https://en.bitcoin.it/wiki/Pooled_mining

However, pooled mining presents two unique challenges to the Theta blockchain:

- First, the number of service certificates for each block could potentially be large, since a pool can serve many viewers. Including all the certificates in the block body may result in considerable bandwidth and storage overhead.

- Second, it is worth noting that on the Theta blockchain, pooling not only reduces reward variance, but also increases the expected mining reward, since the pool can collectively serves more viewers than an individual caching miner node. Hence, compared to other cryptocurrencies like Bitcoin, this would provide extra incentive for the miners to form big pools and potentially lead to centralization of power.

To address these challenges, we propose the following techniques.

TECHNIQUE 1:

Limit the total number of service certificates per block

The protocol can limit the total number of service certificates per block. This has two advantages.

Firstly, this places an upper bound on the additional storage space required for the service certificates per block. The storage and bandwidth overhead of the protocol are limited and more predictable.

Secondly, after the limit is reached, the expected block mining reward would not increase with more mining nodes joining the pool. Hence there is no extra incentive for the miners to form big pools compared to other cryptocurrencies.

TECHNIQUE 2:

Less frequent service certificate PoW puzzle changes

Earlier in the whitepaper we proposed to have the caching node send a new service certificate PoW puzzle to its viewers each time a new block is confirmed, where the $block\_hash_k$ parameter in the new puzzle points to the newly confirmed block. Consider a scenario where viewer node j pulls video stream from caching node i continuously for a couple hours straight. Each time a new block is added, viewer node j needs to generate a different service certificate. To reduce the certificate storage overhead, node i could instead only change the puzzle for every $d > 1$ blocks. Figure 8 illustrates the case where $d = 3$. For a given d, the expected service time for a given service certificate should be calculated using the following Formula

$$T_{sc} = min(E[n \mid H^*] / s,\ d \cdot T_b) \quad (11)$$

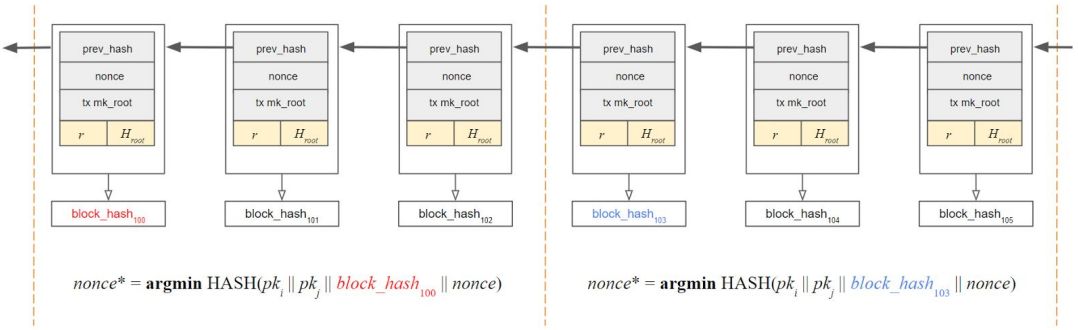

Figure 8. Service certificate PoW puzzle change every d blocks

TECHNIQUE 3:

Randomized reputation score validation protocol

The technique we propose here is to store the service certificates off-chain instead of inside the block body. They can be stored either by the caching miner node that mined the new block, or on decentralized storage systems such as IPFS/SWARM.

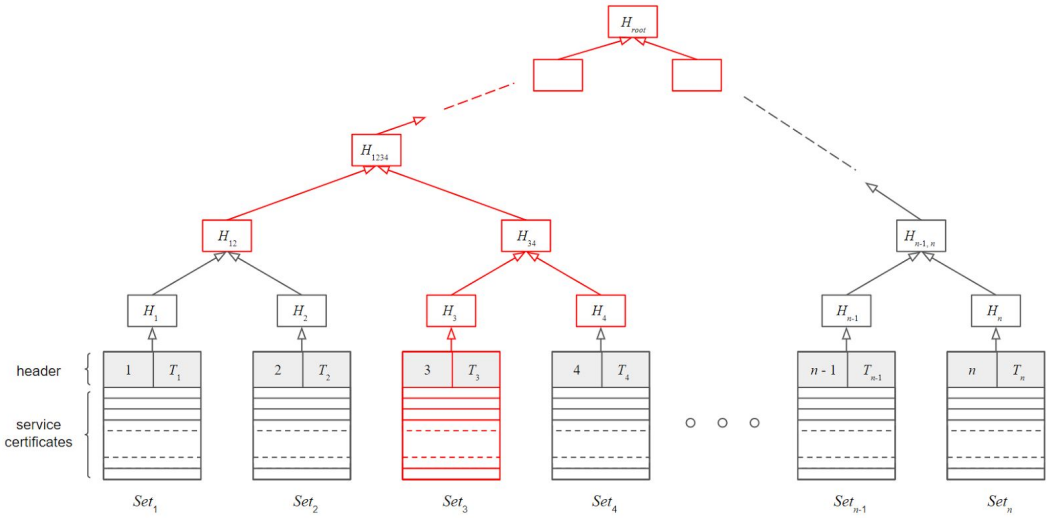

Figure 9. Organizing the service certificates into a Merkle tree for better scalability. The Merkle tree branch for Set_3 is highlighted in red.

On the off-chain storage, the service certificates can be organized into a Merkle tree[16] as shown in Figure 9. Here the certificates are split into n sets, with each set contains similar number of certificates. The maximum number of certificates in each set is a pre-defined constant. Each set also has a header containing two numbers, the set index i, and T_i, the sum of the serving time for each certificate (calculated using Equation (6)) in the set.

The root of the merkle tree is stored on-chain in the header of the newly mined block to prevent potential tampering of the service certificates. In addition, the accumulated service time of each set of certificates T_1, T_2, ..., T_n are also added to the block header.

[16] https://en.wikipedia.org/wiki/Merkle_tree

The Merkle tree structure enables the following randomized reputation score validation protocol:

- **Step 1**. Upon receiving a newly proposed block, the validator miner extract the accumulated service time of each set of certificates T_1, T_2, ..., T_n from the block header. It then verifies the reputation score stored in the block header is computed correctly with the following formulas

$$T_{total} = \sum_{i=1}^{n} T_i$$

$$r = 2 / (1 + e^{-T_{total}}) - 1$$

  If the reputation score is computed correctly, continue to Step 2. Otherwise, reject the new block.

- **Step 2**. The validator miner randomly picks k set of certificates ($k > 1$), and request for these sets and the Merkle tree branch connecting the root and these sets. These data can be downloaded from either the proposer miner, or decentralized storage system, depending on however the service certificates are stored off-chain. The validator miner then verifies that each set of certificates meets the following conditions.
  - The $block\_hash_k$ parameter in each certificate must point to a block that was generated within the defined recent time window.
  - There is no duplicated certificates, i.e. two certificates should not have the same (pk_i, pk_j, $block\_hash_k$)
  - The set index i stored in the set header matches with the requested set index.
  - The sum of the serving time of all the certificates in set i equals to the T_i.
  - Moreover, the validator miner needs to verify the validity of the Merkle tree branch to ensure the certificates included in the set have not been tampered. Figure 9 highlights the Merkle tree branch for the 3rd set of certificates in red.

  Note that it is important have $k > 1$, otherwise the proposer miner can store duplicated certificates in different sets without being detected. If the k set of certificates all pass the verification process, the validator miner will add the new block to the currently longest chain.

- **Step 3**. If any validator miner identifies a set of certificates that does not pass the above verification, it propagates 1) this set of certificates and 2) the Merkle branch from this set to the root through the network. Other validator miners can then verify this set on their own to decide the validity of the claimed reputation score.

It can be shown that since each validator miner picks the k set of certificates at random, as more caching miner nodes join the network, the probability that the proposer miner can give an incorrect reputation score and yet get block confirmations from all honest validator miners approaches zero. The formal analysis is omitted here.

Furthermore, since each validator miner now only needs to validate a small subset of certificates, the amount of data the validator miners need to download is greatly reduced. If a validator miner identified an invalid set, it just need to propagate that set of certificates through the networks, which should not add much overhead since each set only contains a small number of certificates.

Theta Token Issuance Model

Phase I of the Theta network will issue a fixed number of ERC20-compliant application tokens with utility and functionality on the SLIVER.tv and other video platforms including:

- Enabling end-user viewers to gift Theta tokens to influencers directly or purchase virtual items and goods which can then be gifted to influencers.

- Allowing advertisers and brand sponsors to fund their advertising campaigns with Theta tokens. These Theta tokens will automatically be used to reward influencers whose video streams these ads will be displayed in, and optionally, advertisers can reward viewers for their attention to the stream content and advertisements.

- Providing end-users with the ability to purchase premium content, virtual goods and other paid products and services.

With the full launch of the native Theta network and tokens in phase II, **miners will have a fixed amount of Theta tokens they can mine each year**[17], similar to Ethereum, for two primary reasons:

- First, this ensures that there is always an incentive for the caching nodes to relay the video streams to improve their reputation and thus mining rewards. If the total Theta token supply has an upper limit, the caching nodes will not be motivated to relay the video streams after all the Theta tokens are mined.

- Second, it is expected that a portion of the Theta tokens will be lost each year due to transmission to addresses that are no longer accessible. Moreover, Theta tokens might be lost due to loss of private keys, or token destruction by purposely sending to an address that never had an associated private key generated. It is expected that eventually the expected rate of annual loss will balance the rate of issuance.

Anti-fraud Considerations

- **Fake Service Certificates**: As discussed earlier, if a malicious miner try to forge service certificates, its expected mining reward actually decreases. Thus, a rational miner would not try to forge service certificates. We will provide a more detailed analysis in the Appendix.

- **Service Certificate Withholding Attack**: A malicious viewer might withhold the best nonce it finds. However, we note that the difficulty level of finding a nonce that reflects the service time is reasonable for a typical PC, unlike the extremely difficult Bitcoin mining puzzle as of today. Hence, it is relatively easy to detect such attacks statistically. If a viewer is identified to withhold the best nounce it finds, the upstream caching node can simply terminates the video relay service.

[17] https://blog.ethereum.org/2014/04/10/the-issuance-model-in-ethereum/

- **Stream Relay Denial Attack**: A malicious caching node might deny service request from viewers and just mine for the tokens. We show here that this would actually reduce its mining reward. As discussed earlier, the reputation score of a caching node depends on the amount of video streams to the end viewers in a recent time window. Therefore, if a caching node is inactive for a certain amount of time, its reputation would decrease, which in turn, would reduce the mining rewards it can obtain. To maximize mining reward in the end, the dominant strategy for a caching node is to serve viewers consistently.

Smart Streaming Contracts

The Theta blockchain supports a set of specialized smart contracts which are novel and can be utilized to implement several scenarios which will help facilitate reward collection and distribution between various players in a distributed streaming and entertainment ecosystem.

Incentive Contract

The Theta blockchain supports a specialized smart contract called **incentive contract**. An incentive contract is designed for several use cases which may involve a large number of different parties. Rather than require complex application logic, this smart-contract simplifies the process of collecting and distributing tokens based on certain pre-defined criteria and allocations across different parties, some of whom may not be identified until later.

We believe the Incentive contract is a novel concept which will have many uses, including but not limited to:

- **Advertisers rewarding viewers and streamers**. In this case, the incentive contract serves as a way to deposit tokens in a decentralized advertising network, which are drawn upon by various parties as the advertisements are shown and watched. The tokens get spread across multiple parties, including one or more streamers, and any number of viewers (ranging from a few thousand to as high as a few million), based upon proof of engagement. The streamers may be identified at the start, or they may be added later. The allocation may be specified based on percentages (for example, 50% to streamers and 50% to viewers) or other ways, and the contract is depleted over time based on viewership, etc.

- **Viewers can gift rewards to multiple parts of the streaming chain**. An incentive contract filled by a viewer for a specific streamer can also benefit the ingest nodes and caching nodes that brought the stream to a specific viewer.

- **Gift contract for multiple streamers**. A viewer may want to gift tokens as they are earned to multiple streamers. An incentive contract can be used to set up these gifts, which can be allocated across multiple parties.

- **Paid/Premium video content**. The concept of an incentive contract can be used to support paid/premium content in the network in the future. See “Future Work” section. To watch paid content, a viewer must put tokens into an incentive contract, which gets

split based upon pre-defined allocations with the content owner (who may or may not be streaming it).

- **Subscriptions and Decentralized Entertainment Networks (DEN)**. Incentive contracts can enabled premium services like a Netflix, Amazon Prime or Hulu to be built on top of our decentralized video streaming network. Viewers can subscribe by allocating tokens into incentive contracts which get parsed out between the many content owners in the network and the streaming/bandwidth providers who enable the network.

- **A miner can share mining rewards with viewers and with content streamers**. A miner can put a percentage of their rewards, their mined tokens, into an incentive contract and these can be distributed to the streamers who own the content they are streaming or share mining rewards with viewers.

These smart contracts are executed by the caching miner nodes. Typically, in order for a recipient to receive the payment, it needs to send a certain proof to the smart contract. Upon proof validation, the smart contract initiates the token reward, and the original funders of the incentive contract do not need to be involved.

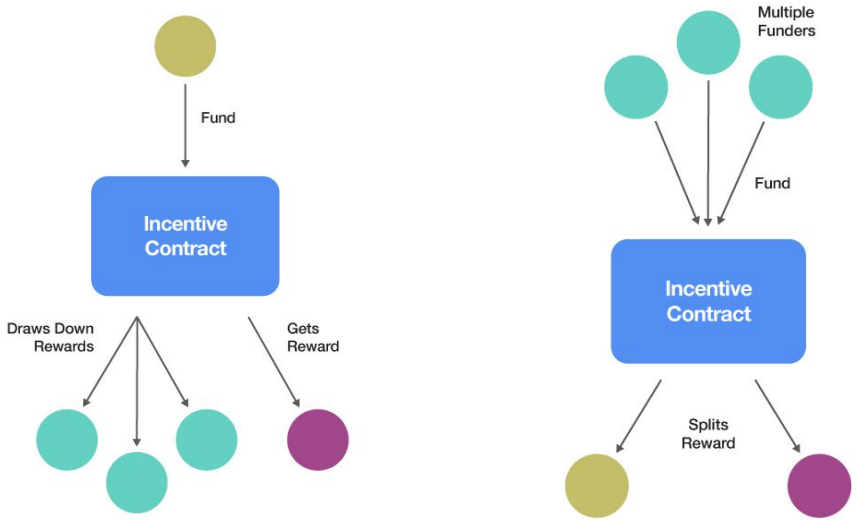

Incentive Contract Example 1:
Single Funder, multiple recipients

Incentive Contract Example 2:
Multiple funders, multiple recipients

Figure 10. Incentive contract examples

One of these smart-contracts, the live stream watching reward, which allows the advertisers to reach the end viewers directly, is expanded below. To claim live stream watching reward, an end viewer needs to supply the smart contract with the Proof-of-Engagement which will be discussed below in detail.

Live Stream Watching Reward

Proof-of-Engagement

Engagement reward is proven to be a very effective way to increase user retention in the live streaming industry. Examples include daily logins rewarded with virtual currency common in many mobile games and apps.

In the game streaming context, Twitch recently introduced the "Twitch Drop" feature which ***randomly*** rewards the viewers with in-game loot[18] as they watch streams on Twitch. Engagement reward is also effective for an advertiser to directly target end viewers. In the "Twitch Drop" example, the rewarded in-game loot provides incentives for the viewers to come back and play the game, thus effectively retargeting the viewer on behalf of the game publisher, which can also be viewed as the advertiser sponsoring the stream.

SLIVER.tv's Watch & Win platform is another example of a system that rewards viewers for watching the eSports streams. Analytical data shows that the watching incentive increases viewer engagement and retention.

However, these types of engagement reward schemes are usually vulnerable to attacks. For example, a malicious user can perform the Sybil attack where he or she creates a large number of accounts to claim engagement rewards. Below we will propose an engagement incentive protocol based on **Proof-of-Work** that increases the robustness of the system against these attacks.

In our proposed protocol, for viewers to claim Theta token rewards from advertisers, they will need to solve cryptographic puzzles during the live stream similar to the Bitcoin mining puzzles. However, in this case, the Theta rewards is not mined. Instead, the advertisers prepare the puzzles and assign the amount of Theta token rewards to each puzzle. The puzzles are **embedded** in the live stream itself and sent to viewers continuously during the live broadcast session. The first viewer that solves a puzzle will write the solution to the blockchain and then receive the corresponding reward after the solution is verified.

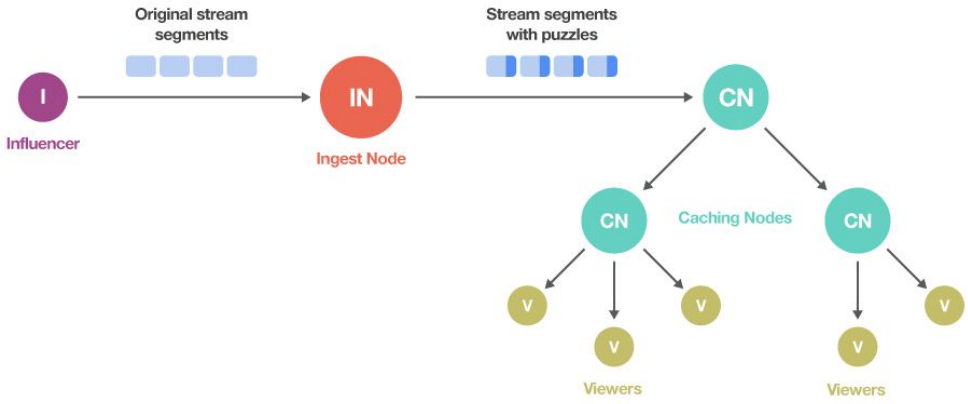

Figure 11. The ingest server provided by the streaming site embeds cryptographic puzzles into the video stream and sends the modified stream to the CDN.

[18] https://dev.twitch.tv/drops

The node at the bottom-left corner of the Figure 11 represents the influencer. He or she simply publishes the stream to the ingest node as usual. As discussed earlier, the ingest node is responsible for transcoding the source stream into **video segments** and then sending them to the CDN network.

Besides transcoding, the ingest server will also create a **cryptographic puzzle** similar to the Bitcoin mining puzzle **for each video segment**. The puzzle is "embedded" in the corresponding video segment. To create the puzzle, the ingest server appends a short string to to the end of each video segment as shown in Figure 8. The short string for $segment_m$ has the following components:

$$(prize_m,\ difficulty_m,\ randstr_m) \tag{12}$$

$prize_m$ specifies the amount of Theta tokens for solving the puzzle. $difficulty_m$ defines the difficulty target. $randstr_m$ is a short random string (e.g. 128 bits) generated by the ingest server on the fly. The puzzle here is to find a ***nonce*** such that the following hash value is smaller than the maximum possible hash value M divided by the difficulty target:

$$HASH(segment_m \,||\, randstr_m \,||\, pk \,||\, nonce) \ \leq\ M\ /\ difficulty_m \tag{13}$$

In the above formula pk is the public key of the viewer node.

For each puzzle, only the first node that provides the solution can win the prize. As mentioned earlier, the issuance of the puzzle prizes are carried out using a smart incentive contract. If a node solves a puzzle, it sends the solution to the incentive contract immediately in order to claim the prize. For the incentive contract to verify the solutions of the puzzles, the advertisers and/or the viewers can save the puzzles (segment plus the short string postfix) to decentralized storage system such as IPFS/SWARM. Upon validation of each puzzle solution, the incentive contract generates the transaction to send the prize to the viewer.

The parameter $difficulty_m$ controls the expected amount of time a viewer node needs to solve the puzzle similar to Bitcoin mining. It should not be too easy so a viewer node can solve it in milliseconds -- this will create an unfair advantage for the nodes that has lower streaming latency. Likewise, it should not be too hard otherwise it takes too much computing time to find the solution, which may discourage viewers from participating.

The prize of each puzzle is determined by the streaming site, so **the streaming site has full control of the total budget**. The streaming site can choose prizes wisely so it encourages the viewers to watch the stream but keeps the total cost under budget.

It is worth noting that in Formula (13), the video segment $segment_m$ is required in order for the viewer nodes to compute the hash. This means the viewer has to "watch" the video segment to win the prize. To encourage viewers to watch the entire stream, **the prizes for the segments can increase monotonically with time**. Then, it can be proven that for viewers, the optimal strategy to maximize its expected total reward is to pull each video segment and try to solve each of the puzzles.

The puzzles might create overhead for the stream delivery. However, the total number of bits of the additional string ($prize_m$, $difficulty_m$, $randstr_m$) is small enough so the overhead for stream delivery is negligible. Typically a video segment is roughly 1 Mbits, whereas the additional string is less than 512 bits. The overhead is only around 0.05%.

In the discussion above, we only attach a single puzzle to one segment. In practice, we can attach multiple puzzles to each segment to increase the number of rewarded viewers for the live streams. The overhead should still be negligible.

In short, it is worth noting that Proof-of-Engagement is a probabilistic proof that a certain number of viewers have watched the video stream. This proof is publicly auditable, unlike today's streaming ecosystem, such data is only available to streaming sites. In this case, it greatly improves transparency to advertisers.

Extensions

The incentive protocol proposed above embeds the cryptographic puzzles into the video streams to encourage viewers to watch the streams. This idea can be extended to provide other types of user engagement rewards. For example, many website and apps use "daily login" reward to encourage users to come back to the site. We can assign a cryptographic puzzle each day a user logs in, and the user gets the reward if he or she can provide the solution. Following the same line of reasoning above, the cryptographic puzzles can discourage the user from creating multiple accounts to loot daily rewards.

Furthermore, introduce smart contracts among the leaf nodes to solve a specific puzzle together. If any of the leaf nodes in the contract solves the puzzle, the prize will be divided among the participants. This will further increase the number of rewarded viewers for the live streams, and shorten the expected amount of time to get the first reward for the participating viewers.

One shortcoming with the Proof-of-Work based Proof-of-Engagement scheme is that it consumes computing resource on the viewer's computer to solve the cryptographic puzzles. In practice, we can allow the viewers to tune the amount of CPU/GPU usage for puzzle solving, so the puzzle solving will not overtax their computer resources.

An alternative approach to address the computation overhead concern is to devise stake based Proof-of-Engagement instead of work based. In the Appendix we will give details of a stake based Proof-of-Engagement scheme and discuss the pros and cons compared to work based Proof-of-Engagement.

Anti-Fraud Considerations

- For the live stream watching reward, the cryptographic Proof-of-Work prevents false claims for prizes from the viewers. On the other hand, if the advertiser cheats (e.g. save incorrect puzzles to storage), its credibility quickly and transparently collapses, and other honest competing advertisers will earn the trust from viewers.

- **Sybil Attack**: It is worth noting that our proposed protocol greatly improves the resistance to Sybil attack. A malicious user might create thousands of nodes on a computer, but *since these nodes share the same computing resource, the probability that the user finds the puzzle solution doesn't increase*. A professional attacker might gather a large amount of computing resources to solve the puzzles, however, as we mentioned earlier, **the prizes for the puzzles are determined by the streaming site**. It is possible to choose the prizes such that the payout for the professional attacker is below his or her investment in computing resources therefore making it economically unprofitable. Nonetheless, the prizes are still attractive to all legitimate viewers -- after all, they are

simply using redundant computing resources to solve the puzzles while they view the live-stream of interest, so end viewers lose nothing and only stand to gain.

Future Work

In this whitepaper, we introduced a new Decentralized Streaming Network (DSN) and the Theta protocol, a new blockchain and token as the incentive mechanism for the DSN. The Theta network encourages viewers to share their memory and bandwidth resources and solves a number of technical and business challenges. The launch of the native Theta network is planned for late 2018.

In the initial phase, we have assumed the video and livestream to be ad-supported free content. The next phase of our research will provide support for premium paid content, piracy issues and anti-piracy mechanisms in a trustless environment.

There are many other technical aspects of the protocol and network which we classify as future work, beyond the initial launch of the native Theta network:

- **Shared Mining Rewards**. Smart Contracts in the form of incentive contracts can be used to distribute mining rewards between different parties. This would allow certain servers to incentivize end viewers to watch their stream, separate from advertiser's rewards for viewers, and to reward streamers for providing their content to the miner. This may be coordinated with the reputation system such that as more rewards are distributed, the better the chances of miners getting the rewards.

- **Anti-Piracy**. The network can be expanded to include anti-piracy - since tokens may be used to stream and cache certain content, the tokens serve as a "dis-incentive" within the network as the content can be tagged as required tokens or "premium content".

- **General Purpose Service Platform**. The reputation dependent mining protocol is in fact independent of streaming. It can be extended to handle other types of service to allow end users to receive service for free. Another interesting extension is to support more general types of smart contracts in addition to the smart streaming contracts. In the extreme case, the smart contracts supported could be Turing complete similar to the Ethereum Smart Contract[19]. This could enable DApps built on top the Theta blockchain to issue their own tokens.

- **Combining Proof-of-Stake and Proof-of-Work for Reputation Dependent Mining**. Proof-of-Stake has the advantage of less electrical power waste and more eco-friendly. A possible improvement for the reputation dependent mining framework is to incorporate Proof-of-Stake. For instance, a viewer can send certain amount of tokens to the upstream caching node during the streaming session to increase the chance that the upstream caching node generates the next block. The caching node returns the same amount of tokens to the viewer after the streaming session

[19] https://github.com/ethereum/wiki/wiki/White-Paper

(enforced by a smart contract). This way the viewer still gets the video stream for free, and yet the expected minting reward of the upstream caching node increases. One issue with such a scheme is the cold start problem where a new viewer might not have enough tokens. However, this can be solved combining Proof-of-Work style mining. After the viewer obtains enough initial tokens through Proof-of-Work mining, he can get the video streams for free without spending computational power.

Appendix

Service Certificate Analysis

In this section, we show that by selecting constant β in Formula (8) properly, an attacker cannot benefit from generating fake service certificates. More specifically, if the attacker allocates a portion of his hash power to make fake service certificates, his expected mining reward strictly reduces. Hence, a rational attacker would not try to forge certificates.

To prove this, let us assume that the attacker's computing resource can perform n hashes per unit time. If he uses all his hashing power to mine new blocks. His expected mining reward per unit time can be calculated by

$$E[reward] = n \cdot p \cdot reward_{base} \tag{14}$$

where p is the probability of solving the mining puzzle in one hash.

If he allocates a portion $\alpha \bullet n$ of hashing power to generate service certificates (where $0 < \alpha < 1$), and use the remaining $(1 - \alpha) \bullet n$ hashing power to solve the mining puzzle, his expected mining reward per unit time can be calculated by

$$E[reward_{\alpha}] = E[k \cdot (1 + f(T_{total})) \cdot reward_{base}] \tag{15}$$

where k is a random variable representing the number of times the attacker solves the mining puzzle in one unit time with hashing power $(1 - \alpha) \bullet n$. On the other hand, T_{total} can also be viewed as a random variable representing the estimated amount of service time given the service certificated generated with hashing power $\alpha \bullet n$. These two random variables are **statistically independent** since they are derived from two independent random processes. Hence, we can simplify Equation (15) to

$$\begin{aligned} E[reward_{\alpha}] &= E[k] \cdot E[1 + f(t_{total})] \cdot reward_{base} \\ &= (1 - \alpha) \cdot n \cdot p \cdot (1 + E[f(T_{total})]) \cdot reward_{base} \\ &= (1 - \alpha) \cdot (1 + E[f(T_{total})]) \cdot n \cdot p \cdot reward_{base} \end{aligned}$$

We will omit the details here, but it can be shown that if we choose the constant β to be proportional to n_{ent} the total amount of hashes the entire network can perform in a unit time, we can make $\mathbf{E}[f(T_{total})]$ to be a **sublinear** function of $\alpha \bullet n$

$$E[f(T_{total})] < 2 \cdot (\alpha \cdot n) / n_{ent}$$

Hence, we have the following

$$E[reward_{\alpha}] < (1 - \alpha) \cdot (1 + 2 \cdot (\alpha \cdot n) / n_{ent}) \cdot n \cdot p \cdot reward_{base} \quad (16)$$

If the attacker has no more than 50% of the total network hashing power, the following inequality holds

$$2 \cdot (\alpha \cdot n) / n_{ent} \leq \alpha$$

Plug it back into Inequality (16), we have

$$\begin{aligned} E[reward_{\alpha}] &< (1 - \alpha) \cdot (1 + \alpha) \cdot n \cdot p \cdot reward_{base} \\ &= (1 - \alpha^2) \cdot E[reward] \\ &< E[reward] \end{aligned}$$

Thus, splitting any hashing power to forge service certificate results in strict decrease of mining reward.

On the other hand, it is easy to show that with service certificates generated by real viewers, the expected mining reward of a caching miner node strictly increase. Therefore, we can conclude that **our reputation dependent mining scheme rewards honest miners and punishes dishonest miners.**

Founding & Advisory Team

The founding members of the Theta network include:

Mitch Liu - Mr. Liu is the co-founder and CEO of SLIVER.tv, the leading esports entertainment platform with patented technology to live stream top esports events in fully immersive 360° VR in partnership with Intel Extreme Masters, Turner ELEAGUE, ESL One and Dreamhack among other global tournament operators. Along with his co-founder Mr. Long, they currently hold two patents and two additional pending patents for virtual reality 360° video streaming, and new algorithms for generating highly efficient live spherical video streams.

In 2010, Mr. Liu co-founded Gameview Studios best known for its Tap Fish mobile game franchise with nearly 100 Million downloads. The company was acquired by DeNA, a leading Japanese mobile gaming company within 6 months of launch. Prior to that, he co-founded Tapjoy in 2007, a pioneer of rewarded social and mobile video advertising, and grew that company to $100MM in revenues. He received a BS in Computer Science & Engineering from MIT, completed his thesis research at MIT Media Lab "Interactive Cinema" video group and received a MBA from Stanford Graduate School of Business.

Jieyi Long - Mr. Long is the co-founder and Chief Technology Officer of SLIVER.tv. He leads the technical team and developed multiple patented technologies including VR live streaming and instant replay for video games. He received a B.S. degree in Microelectronics from Peking University in Beijing, China. He also received a Ph.D. degree in Computer Engineering from Northwestern University in Evanston, IL where he conducted research in mathematical modeling and algorithms to optimize large scale electronics systems, and a cryptography enthusiast.

Ryan Nichols - Mr. Nichols is the Head of Product and Platform for SLIVER.tv. He leads the company's eSports entertainment platform built around one of the largest esports virtual economies with 1B+ virtual tokens circulated within two months of launch. Leading previous startups, he's designed and launched virtual currency systems for a variety of multiplayer games, including a cross-game virtual currency API used by hundreds of third-party game developers and tens of millions of players worldwide. Mr. Nichols was a director for Tencent on the globally popular WeChat app, and a co-founder of a live video streaming app for foodies.

Rizwan Virk - Mr. Virk is an advisor, investor and the interim Head of Corporate Development at SLIVER.tv. Mr. Virk also serves as the current director of Play Labs @ MIT, and did his research at the MIT Media Lab. Mr. Virk is an early investor in cryptocurrency and blockchain companies, including *Ripio/BitPagos*, *CoinMkt*, *Bex.io*, and has been active with *BitAngels* since 2013. Mr. Virk is the co-author of several cryptocurrency related papers including *Online Automatic Auctions for Bitcoin Over-The-Counter Trading (2015)* and *Creating a Peer to Peer System for Buying and Selling Bitcoin Online (2013)* and was the designer of *Bitcoin Bazaar*, one of the first peer-to-peer mobile applications for in-person trading of bitcoin. Mr. Virk received his BS in Computer Science & Engineering from MIT and his Master's in Management from Stanford Graduate School of Business.

The advisory team to Theta includes:

- ***Justin Kan,*** *Co-founder of Twitch*

- ***Steve Chen,*** *Co-founder of YouTube*

- ***Fan Zhang,*** *Founding member, Sequoia Capital China*

- ***Travis Skweres,*** *Founder CoinMkt, one of the first US bitcoin exchanges*

- ***Rajeev Surati,*** *MIT PhD, video compression and streaming expert*

- ***Prof Shoucheng Zhang,*** *Founder Danhua Capital*

- ***Sebastian Serrano,*** *Founder Ripio, first global lending network on blockchain*

- ***Cliff Morgan,*** *CEO, GFUEL energy drink*

- ***Sam Wick,*** *Head of UTA Ventures, United Talent Agency Hollywood*

- ***Dennis Fong,*** *CEO, Plays.tv aka "Thresh" the world champion of Quake/Doom*

Roadmap

Q4 2015 - SLIVER.tv founded

Q2 2016 - $7.2MM Seed financing

Q3 2016 - esports video streaming platform launched, first patent granted

Q2 2017 - $9.8MM Series A financing, Theta project launched

Q3 2017 - Theta token pre-sale

Q4 2017 - Phase I network launch ERC20-compliant tokens

Q4 2018 - Planned phase II launch of native blockchain and protocol tokens

Token Sale & Anticipated Use of Proceeds

Anticipated use of proceeds will be as follows:

- 60% software development R&D
- 20% marketing, community development
- 20% operational, legal expenses and overhead

List of Figures