



THETA

NETWORK

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

WHITEPAPER



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

Last Updated: May 1, 2018
Version 1.7

Abstract

This whitepaper introduces the Theta Network, a new blockchain and token as the incentive mechanism for a decentralized video streaming and delivery network.

The Theta Network and protocol solves various challenges the video streaming industry faces today. First, Theta Tokens are used as an incentive to encourage individual users to share their redundant computing 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 4k, 8k and next generation 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 novel concepts:

- **Resource Oriented Micropayment Pool:** We have designed and implemented an off-chain "Resource Oriented Micropayment Pool" that is purpose-built for video streaming. It allows a user to create an off-chain micropayment pool that any other user can withdraw from using off-chain transactions, and is double-spend resistant. It is much more flexible compared to off-chain payment channels.
- **Proof-of-Engagement:** The micropayment records can be used as "Proof-of-Engagement (PoE)". Each payment for the video segment is associated with the ID of the video (i.e. the resource ID) and can be used to track delivery of video segments to end viewers. With PoE, viewers can earn tokens as rewards from advertisers in exchange for their attention to video streams and by providing PoE. PoE is used to prove that viewers legitimately consume the video streams, providing better transparency and a basis for viewers to earn Theta Tokens for engaging with content.

This white paper will describe these concepts and the Theta blockchain in detail. The Theta Network launched with ERC20-compliant tokens and were integrated into the SLIVER.tv platform in December 2017. 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	4
Introduction	4
Video streaming market	4
Video streaming challenges	5
Background	6
Traffic Overview	8
Opportunity	8
Theta Decentralized Streaming Network	10
Overview	10
Bootstrapping The Network	12
The Theta Blockchain	13
Overview	13
Resource Oriented Micropayment Pool	13
Double Spending Detection and Penalty Analysis	16
Applications of Micropayment Pool Beyond Video Streaming	18
Smart Streaming Contracts	19
Incentive Contract	19
Future Work	20
Founding & Advisory Team	22
Project Roadmap	24
List of Figures	26

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.

Global IP Video Traffic Growth

IP video will account for 82% of global IP traffic by 2020

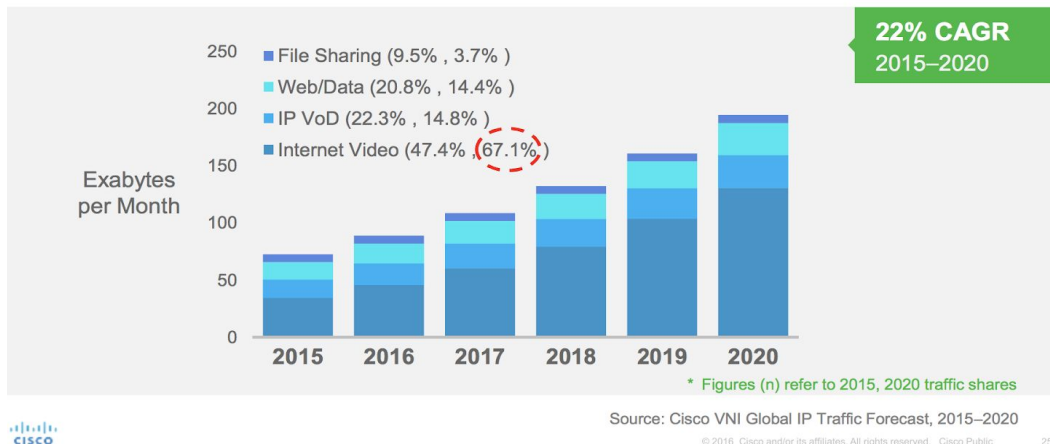


Figure 1. Global IP video traffic growth

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/>

Global Virtual Reality Traffic Growth

Virtual reality traffic quadrupled in the past year, and will increase 61-fold by 2020

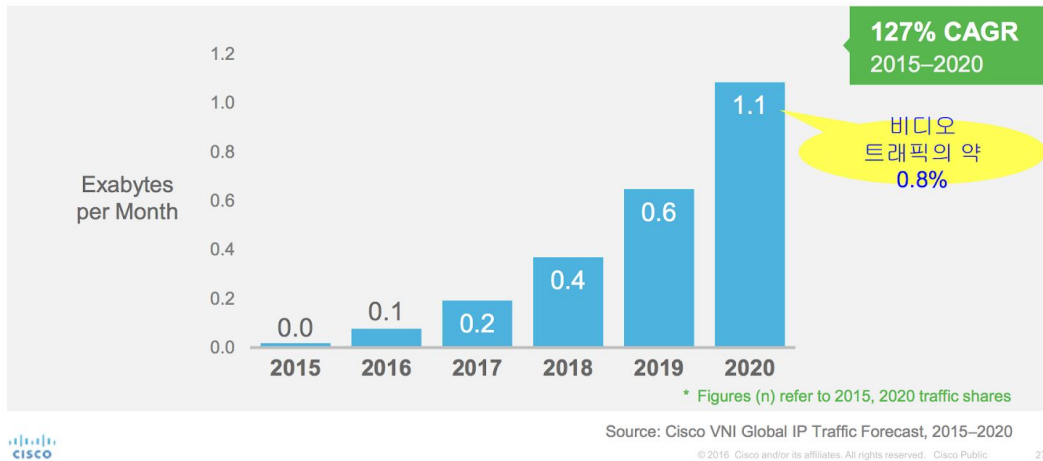


Figure 2. Global virtual reality traffic growth³

Video streaming challenges

Content Delivery Network (CDN) plays an important role in the video streaming 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 developed regions. This "last-mile" link is usually the bottleneck of the streaming delivery pipeline and often leads to less optimal user experience including choppy streams and frequent rebuffering.

To streaming sites and platforms, another major concern is the CDN bandwidth cost. For popular sites, the CDN bandwidth cost can easily reach 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, 8k, 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. Bandwidth requirement jumps quickly 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 today is the leading next-generation live esports streaming platform with over five million unique visits in March 2018, with a vision to transform the esports engagement

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 research⁴ 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 Masters⁷. 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 the first few weeks of launch. 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 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 blockchain technology to create the first **Decentralized Video Streaming and Delivery Network** whereby video viewers are incentivized to share redundant computing and bandwidth resources to address today's video streaming challenges. Using the Ethereum EVM "World Computer" metaphor, the Theta Network 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 devices as "**caching nodes**" whereby they form a video delivery mesh network that is responsible for delivering any given video stream to viewers anywhere around the world. The Theta Network can effectively address the technical challenges discussed in the previous section. First, viewers' devices 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 and next generation streaming technology practical.

To encourage viewers to contribute their computing and bandwidth resources, we introduce **the Theta protocol** as an incentive mechanism. The caching nodes can earn tokens as they relay video streams to other viewers. Not only do Theta Tokens motivate viewers to join the network as caching nodes, 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 their favorite content, and influencers earn Theta Token as gifts directly from viewers. Streaming platforms reduce CDN costs and open up new revenue opportunities with Theta Tokens.

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

- Advertisers fund ad campaigns with tokens to support influencers, streaming sites and viewers
- Caching nodes earn tokens for caching and relaying video streams to other viewers
- Viewers potentially earn tokens from advertisers as engagement rewards, and can in turn gift to favorite influencers and content creators
- Streaming sites and platforms can offload up to 80% of CDN costs and drive new revenues through Theta powered premium goods and services.

The Theta protocol builds upon the following novel concepts:

- **Resource Oriented Micropayment Pool:** We have designed and implemented an off-chain “Resource Oriented Micropayment Pool” that is purpose-built for video streaming. It allows a user to create an off-chain micropayment pool that any other user can withdraw from using off-chain transactions, and is double-spend resistant. It is much more flexible compared to off-chain payment channels. In particular, for the video streaming use case, it allows a viewer to pay for video content pulled from multiple caching nodes without on-chain transactions. By replacing on-chain transactions with off-chain payments, the built-in “Resource Oriented Micropayment Pool” significantly improves the scalability of the blockchain.
- **Proof-of-Engagement:** The Resource Oriented Micropayment Pool can be leveraged to track the delivery of the video segments to the end viewers, **since each payment for the video segment is associated with the ID of the video (i.e. the resource ID)**. This indicates that the micropayments recorded in the blockchain denote video segments delivered to end viewers. Thus the micropayment records can be used as the “Proof-of-Engagement (PoE)”. With PoE, viewers can earn tokens as rewards from advertisers in exchange for their attention to video streams and by providing PoE. 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. PoE can also be used as a basis to reward the content creators, which could be done automatically by the validators of the Theta blockchain since PoE can be inferred from the transaction records.

Theta 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.

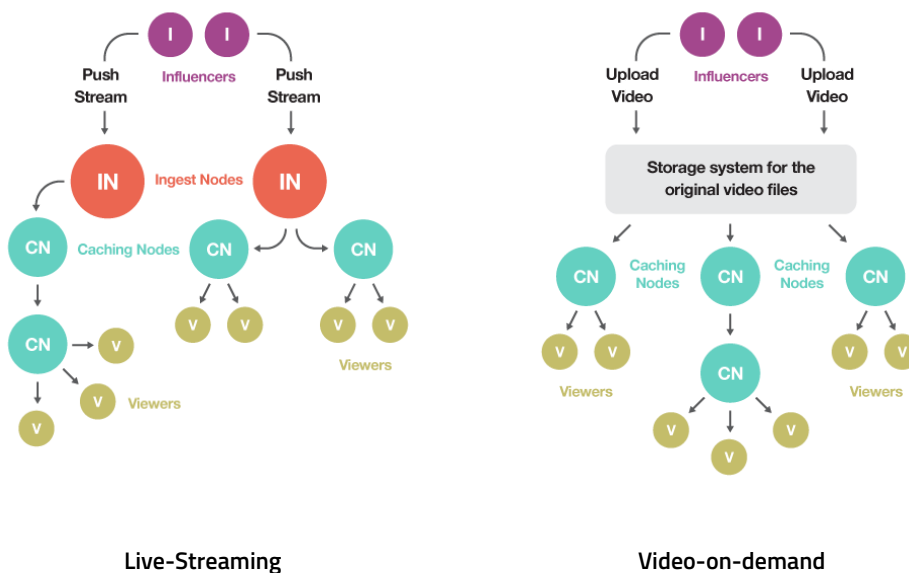


Figure 3. **Left** - Theta Network architecture diagram for live streaming
Right - Theta Network architecture diagram for video on demand

Figure 3 Left shows the Theta Network 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 real-time 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 Theta Network 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 IPFS⁹ or SWARM¹⁰, 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 Theta Network 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

$$GUID = \textit{QuadtreeRegionID} \parallel \textit{RandomBitsPostfix}$$

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.

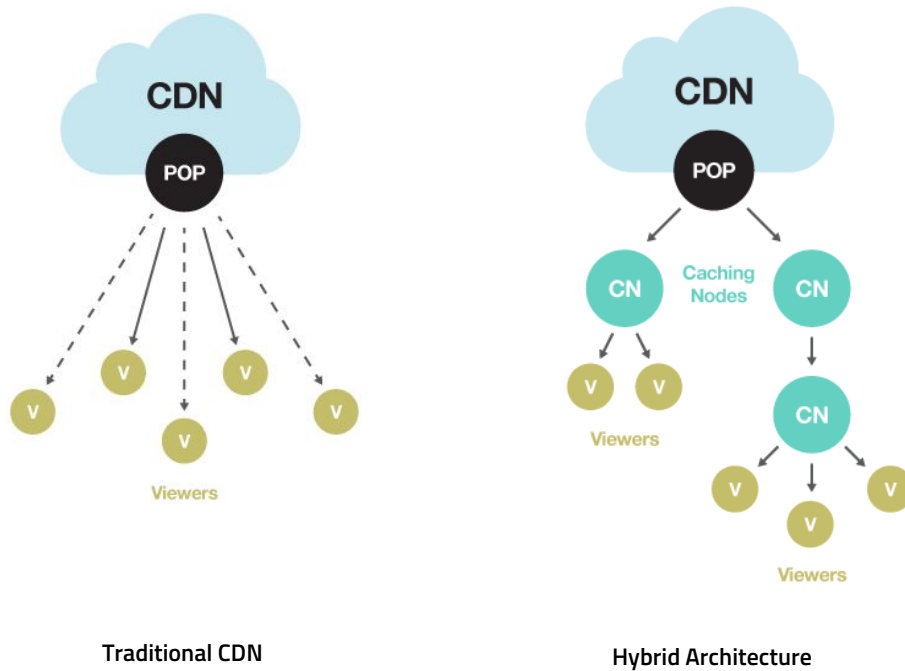


Figure 4. **Left - Traditional CDN**, where all the viewers connect directly to the POP servers. For nodes geographically far away from the POP servers, the stream quality may be lower (the dashed connections). **Right - Hybrid Architecture** where viewer nodes can pull stream from caching nodes that 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 Theta network is fully functional even with only a single caching node.** Furthermore, as will be discussed later, the caching node can earn 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 Theta Network, 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.

The Theta Blockchain

Overview

The Theta blockchain is based on a proof of stake (POS) consensus mechanism. PoS has the advantage of lower computation costs and higher transaction throughput. In this model, caching nodes and viewer clients don't need to conduct hashing computations, which lessens computational resource waste. Nonetheless, this approach has an important impact on the Theta protocol design, especially on the bandwidth sharing reward scheme.

Separation of the Validator and Cacher Role

To further reduce the computational burden of the caching nodes, the role of validator and cacher are separated. Validators are dedicated nodes for verifying transactions and assembling new blocks. With such separation, caching nodes and viewers don't need to perform extra computations aside from the usual video stream encoding/decoding tasks, as they would in a POW scenario. This enables a wide range of edge devices (e.g. personal computers, mobile devices, IoT devices) to operate as caching nodes and viewers.

Scalability

Support for high transaction throughput is a must for a video streaming focused blockchain. Imagine a scenario where a live streaming event could generate numerous token transactions for caching node rewards in a short amount of time. Further examples include virtual gifting, where a shoutout from a popular streamer can lead to a surge in token donations, possibly thousands of transactions per second. ERC-20 implementations simply cannot support this level of performance — this is our view, which has also been echoed by many other projects.

Fortunately, for many streaming use cases, the transactions are independent of each other. This unique characteristic makes techniques like sharding very appealing. Proper sharding could potentially execute all the independent transactions in parallel, thus greatly increasing network throughput. Sharding also reduces the data storage requirement for the validators since a validator only needs to work on one shard at a time.

Further, many of the token transactions are recurring (e.g. a viewer might make many micro token transactions to the same upstream caching node).

We developed a novel way to support high volumes of transactions through a Resource Oriented Micropayment Pool.

Resource Oriented Micropayment Pool

One of the biggest challenges we identified when designing the Theta protocol is how to scale our native chain for ultra high transaction throughput. Although many blockchain projects are facing scaling problems, scaling for live video streaming is different and possibly even more

complex. Existing level two scaling solutions like payment channels have various disadvantages when applied to the video streaming. We have designed and implemented an off-chain **“Resource Oriented Micropayment Pool”** that is purpose-built for video streaming. It allows a user to create an off-chain micropayment pool that any other user can withdraw from using off-chain transactions, and is double-spend resistant. It is much more flexible compared to off-chain payment channels. In particular, for the video streaming use case, it allows a viewer to pay for video content pulled from multiple caching nodes without on-chain transactions. By replacing on-chain transactions with off-chain payments, the built-in “Resource Oriented Micropayment Pool” significantly improves the scalability of the blockchain.

The following scenario and diagram provide a comprehensive walkthrough of how the Resource Oriented Micropayment Pool works in application.

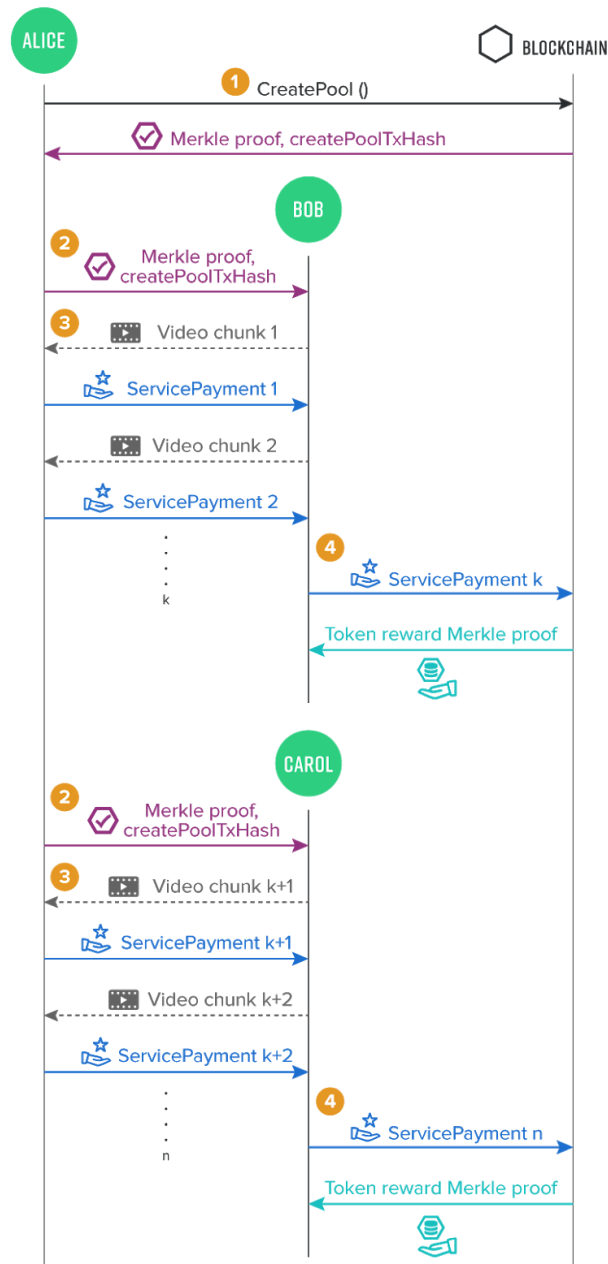


Figure 5. **Resource Oriented Micropayment Pool** shows viewer Alice making off-chain transactions to cachers Bob and Carol for video chunks

- Step 1. Micropayment pool creation: As the first step, Alice publishes an on-chain transaction to create a micropayment pool with a time-lock and a slashable collateral.

CreatePool(resourceId, deposit, collateral, duration)

A couple things to be noted. To create the pool, Alice needs to specify the "Resource ID" *resourceId* that uniquely represents the digital content she intends to retrieve. It may refer to a video file, or a live stream.

The *deposit* amount needs to be at least the total value of the resource to be retrieved. For instance, if the resource is a video file which is worth 10 tokens, then the deposit has to be at least 10 tokens.

The *collateral* is required to discourage Alice from double spending. If a double spending attempt from Alice is detected by the validators of the blockchain, the collateral will be slashed. Later in the blogpost we will show that if $collateral > deposit$, the net return of a double spend is always negative, and hence any rational user will have no incentive to double spend.

The duration is a time-lock similar to that of a standard payment channel. Any withdrawal from the payment pool has to be before the time-lock expires.

The blockchain returns Alice the Merkle proof of the *CreatePool()* transaction after it has been committed to the blockchain, as well as *createPoolTxHash*, the transaction hash of the *CreatePool()* transaction.

- Step 2. Initial handshake between peers: Whenever Alice wants to retrieve the specified resource from a peer (Bob, Carol, or David, etc.). She sends the Merkle proof of the on-chain *CreatePool()* transaction to that peer. The recipient peer verifies the Merkle proof to ensure that the pool has sufficient deposit and collateral for the requested resource, and both parties can proceed to the next steps.
- Step 3. Off-chain micropayments: Alice signs *ServicePayment* transactions and sends them to the peers off-chain in exchange for parts of the specified resource (e.g. a piece of the video file, a live stream segment, etc.). The *ServicePayment* transaction contains the following data:

targetAddress, transferAmount, createPoolTxHash, targetSettlementSequence,
 $\text{Sign}(SK_A, \text{targetAddress} \parallel \text{transferAmount} \parallel \text{createPoolTxHash} \parallel \text{targetSettlementSequence})$

The *targetAddress* is the address of the peer that Alice retrieves the resource from, and the *transferAmount* is the amount of token payment Alice intends to send. The *targetSettlementSequence* is to prevent a replay attack. It is similar to the "nonce" parameter in an Ethereum transaction. If a target publishes a *ServicePayment* transaction to the blockchain (see the next step), its *targetSettlementSequence* needs to increment by one.

The recipient peer needs to verify the off-chain transactions and the signatures. Upon validation, the peer can send Alice the resource specified by the *CreatePool()* transaction.

Also, we note that the off-chain *ServicePayment* transactions are sent directly between two peers. Hence there is no scalability bottleneck for this step.

- Step 4. On-chain settlement: Any peer (i.e. Bob, Carol, or David, etc) that received the *ServicePayment* transactions from Alice can publish the signed transactions to the blockchain anytime before the timelock expires to withdraw the tokens. We call the *ServicePayment* transactions that are published the “on-chain settlement” transactions.

Note that the recipient peers needs to pay for the gas fee for the on-chain settlement transaction. To pay less transaction fees, they would have the incentive to publish on-chain settlements only when necessary, which is beneficial to the scalability of the network.

We note that no on-chain transaction is needed when Alice switches from one peer to another to retrieve the resource. In the video streaming context, this means the viewer can switch to any caching node at any time without making an on-chain transaction that could potentially block or delay the video stream delivery. As shown in the figure, in the event that Bob leaves, Alice can switch to Carol after receiving k chunks from Bob, and keep receiving video segments without an on-chain transaction.

Moreover, the total amount of tokens needed to create the micropayment pool is (*collateral + deposit*), which can be as low as twice of the value of the requested resource, no matter how many peers Alice retrieves the resource from. Using computational complexity language, the amount of reserved token reduces from $O(n)$ to $O(1)$ compared to the unidirectional payment channel approach, where n is the number of peers Alice retrieves the resource from.

Double Spending Detection and Penalty Analysis

To prevent Alice, the creator of the micropayment pool from double spending, we need to 1) be able to detect double spending, and 2) ensure that the net value Alice gains from double spending is strictly negative.

Detecting double spending is relatively straightforward. The validators of the Theta Network check every on-chain transaction. If a remaining deposit in the micropayment pool cannot cover the next consolidated payment transaction signed by both Alice and another peer, the validators will consider that Alice has conducted a double spend.

Next, we need to make Alice worse off if she double spends. This is where the collateral comes in. Earlier, we mentioned that the amount of collateral tokens has to be larger than the deposit. And here is why.

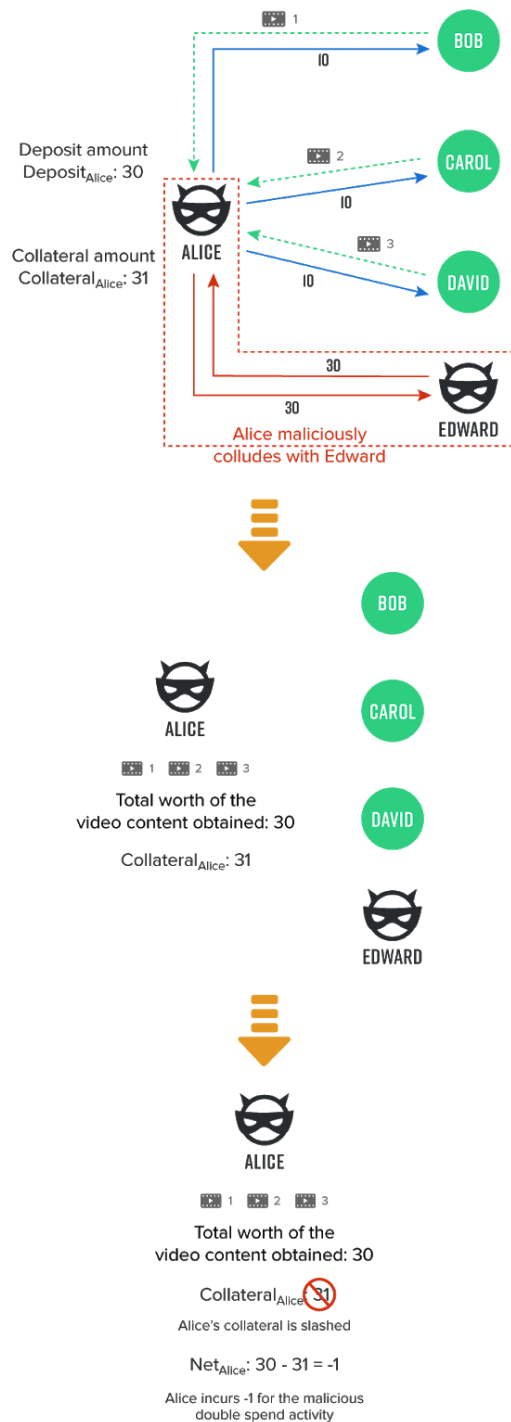


Figure 6. **Malicious Actor Detection and Penalty** shows malicious actor Alice attempting to make a double spend and the resulting penalty she receives

In the diagram above, Bob, Carol, and David are honest. Alice is malicious. Even worse, she colludes with another malicious peer Edward. Alice exchanges partially signed transactions with Bob, Carol, and David for the specified resource. Since Alice gains no extra value for the duplicated resource, the maximum value she gets from Bob, Carol, and David is at most the *deposit* amount. As Alice colludes with Edward, she can send Edward the full *deposit* amount. She then asks Edward to commit the settlement transaction before anyone else and return her the *deposit* later. In other words, Alice gets the resource which is worth at most the *deposit* amount for free, before the double spending is detected. Later when Bob, Carol, or David

commit the settlement transaction, the double spending is detected, and the full *collateral* amount will be slashed. Hence, the net return for Alice is

$$net_{Alice} = deposit - collateral$$

Therefore, we can conclude that for this scenario, as long as $collateral > deposit$, Alice's net return is negative. Hence, if Alice is rational, she would not have any incentive to double spend.

We can conduct similar analysis for other cases. The details are omitted here, but it can be shown that in all cases Alice's net return is always negative if she conducts a double spend.

Another case is that Alice is honest, but some of her peers are malicious. After Alice sends a micropayment to one of those peers, it might not return Alice the resource she wants. In this case, Alice can turn to another peer to get the resource. Since each incremental micropayment can be infinitesimally small in theory, Alice's loss can be made arbitrarily small.

Proof of Engagement

The "Resource Oriented Micropayment Pool" described can actually be leveraged to track the delivery of the video segments to the end viewers, **since each payment for the video segment is associated with the ID of the video (i.e. the resource ID)**. This indicates that the micropayments recorded in the blockchain reflect the viewing activity of the end viewers. Thus the micropayment records can be used as the "Proof-of-Engagement (PoE)".

With PoE, viewers can earn tokens as rewards from advertisers in exchange for their attention to video streams and by providing PoE. 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. PoE can also be used as a basis to reward the content creators, which could be done automatically by the validators of the Theta blockchain since PoE can be inferred from the transaction records.

Applications of Micropayment Pool Beyond Video Streaming

In the above discussion, we have presented the resource oriented micropayment pool concept in the video delivery context. It is worth pointing out that it can be used in other applications as well, as long as we can identify a resource where a user can gain no extra value when he obtains multiple copies of the resource from different peers. Here are a few examples:

- 1) A resource can be a generic file. So the resource oriented micropayment pool can facilitate generic file sharing/storage.
- 2) In the context of app distribution (e.g. Apple AppStore, GooglePlay), an app can be considered as a resource.
- 3) A resource can be a certain type of service. E.g. file compression service, video transcoding service, solving a set of linear equations, etc. For these services, the requester gains no extra value by getting the same service from two vendors.

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

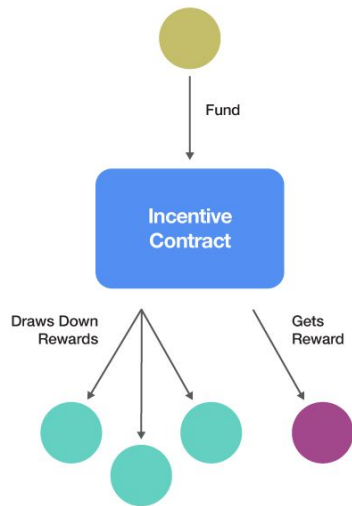
The Theta blockchain supports a specialized smart contract called **incentive contract**. An incentive contract is designed for several use cases which may involve a large number of different parties. Rather than require complex application logic, this smart-contract simplifies the process of collecting 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.** Incentive contracts can enable 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 cacher can share caching rewards with viewers and with content streamers.** A caching node can put a percentage of their rewards, their earned tokens, into an

incentive contract and these can be distributed to the streamers who own the content they are streaming or share caching rewards with viewers.

These smart contracts are executed by the validators. 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 7. 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.

Future Work

In this whitepaper, we introduced the Theta protocol, a new blockchain and token as the incentive mechanism for a decentralized video streaming network. The Theta Network encourages viewers to share their computing 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 Caching Rewards.** Smart Contracts in the form of incentive contracts can be used to distribute caching 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 cacher.
- **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 Theta 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 Contract¹¹. This could enable DApps built on top the Theta blockchain to issue their own tokens.

¹¹ <https://github.com/ethereum/wiki/wiki/White-Paper>

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
- **Steve Dakh**, CTO of SmartWallet and a founding member of the Ethereum project
- **Ma Haobo**, CEO of aelf
- **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

Project 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 private sale begins

Q4 2017 - Theta ERC20 Tokens integrated on SLIVER.tv platform

Q1 2018 - Theta private sale completed, tokens issued, full integration on SLIVER.tv

Q4 2018 - Planned launch of native blockchain and protocol tokens

THETA BLOCKCHAIN DEVELOPMENT ROADMAP

(ONGOING) SANDBOX TESTING ENVIRONMENT

- Complete initial implementation of PoS blockchain (based on Tendermint)
- Support multiple validators in different geographies
 - Conduct scalability tests
- Implement off-chain micro transaction protocol

(LATE Q2) TESTNET DEPLOYMENT

TESTNET LAUNCH

- Deploy testnet clones for early development partners
- Conduct security review: protocol design, cryptography design, 3rd party integrations

MESH STREAMING

- Develop packet routing algorithms, peer group formation methods, stream pulling methods
- Develop usage tracking & rate limiting mechanisms
- Establish dedicated testing channel on SLIVER.tv platform
- Implement blockchain-streaming client interaction protocol

COMPLETED

- Introduced Theta Tokens on SLIVER.tv
- Implemented PoS blockchain infrastructure
- Implemented block explorer v0.1
- Implemented initial off-chain micro transaction protocol
- Built mesh streaming prototype
- Architected blockchain-streaming client interaction protocol

(Q2) TESTNET PREPARATION

PUBLIC INTERACTIVE DEMO

- Community users simulate how peer nodes in Theta protocol cache and relay streams

BLOCKCHAIN EXPLORER

- Functionality: browse block history, transactions history, users' addresses

(Q3) WALLET CLIENT LAUNCH

- Secure multi-platform wallet for transacting Theta protocol-related currency

List of Figures

Figure 1. Global IP video traffic growth

Figure 2. Global virtual reality traffic growth

Figure 3.

Left - Theta Network architecture diagram for live streaming

Right - Theta Network architecture diagram for video on demand

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 - The 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 5. Resource Oriented Micropayment Pool shows viewer Alice making off-chain transactions to cashers Bob and Carol for video chunks

Figure 6. Malicious Actor Detection and Penalty shows malicious actor Alice attempting to make a double spend and the resulting penalty she receives

Figure 7. Incentive contract examples