

A Decentralized Cyber Threat Intelligence Market.

Table of Contents

PolySwarm in 60 Seconds	03
Background.....	04
Reinventing the Threat Intelligence Market.....	06
The Participants	07
Prediction Markets, Arbiters & Mediated Consensus	09
Determining Ground Truth	10
The Instruments	11
Rewarding Accuracy.....	13
Fees.....	15
Protocol Details	16
Bounty Lifecycle	17
Offer Lifecycle.....	20
Reputation.....	22
Worker Registry.....	23
Artifact Confidentiality	24
Additional Markets	25
Roadmap.....	26
Token Sale.....	29
Disclaimer.....	32

PolySwarm in 60 Seconds

Polyswarm is a decentralized threat intelligence market made possible by Ethereum smart contracts and blockchain technology.

PolySwarm incentivizes rapid innovation in the \$8.5B/yr anti-virus and automated cyber threat intelligence space with precise economic incentives that reward *timely* and *accurate* threat intelligence concerning the malintent of files, network traffic and URLs.

PolySwarm defines a real-time threat detection ecosystem involving enterprises, consumers, vendors and geographically-diverse security experts. Experts develop and hone competing “micro-engines” that autonomously investigate the latest threats, attempting to outperform their competition. PolySwarm’s “Proof of Work” is threat detection accuracy: the market rewards experts who are best able to defend enterprises and end users.

Relative to today’s ad hoc market, PolySwarm will lower the barrier to entry, provide broader coverage options, discourage duplicative effort and ensure interoperability among products and threat intelligence feeds.

Economically, PolySwarm functions as a skill-required twist on a prediction market² with thousands of micro-engines (“workers”) investigating the latest in malware evolution at machine speed – no human in the loop.

PolySwarm will be developed by **PolySwarm Pte. Ltd.** with funding derived from the sale of ERC20-compatible **Nectar** (“NCT”) utility tokens.

As a utility token, PolySwarm economically disincentivizes Nectar speculation by rewarding honest market participation through the [collection and distribution of Fees \(details on page 6\)](#) to value-adding, active participants.

¹ Jefferies Cyber Security Primer. January 18, 2017.

² https://en.wikipedia.org/wiki/Prediction_market

Background

Today's enterprises rely on an ad-hoc mixture of anti-virus subscriptions, threat intelligence feeds, and assorted dynamic analysis engines to defend against evolving adversarial cyber activity. Users must weigh the benefits and drawbacks presented by each solution and decide on the least-worst fit for their environment.

Today's market discourages solutions that provide broad threat coverage.

Today's solutions focus on a comfort zone of threats – a direct result of today's market economics.




Figure A: The black rectangle represents all the threats an Enterprise may encounter; white circles are anti-virus product #1 and #2, respectively.

It's easy (and perhaps justifiable) to ignore an anti-virus solution that doesn't detect WannaCry³, but in doing so, today's market effectively rewards overlapping coverage among vendors – a market inefficiency that causes duplicative cost. This is a classic tragedy of the commons situation.

Similarly, consider a vendor that chooses to develop expertise outside of this comfort zone. If the comfort zone is defined by the set of threats faced by most enterprises, sales for this specialized vendor will be difficult. How do you convince a potential customer that they

³ https://en.wikipedia.org/wiki/WannaCry_ransomware_attack

will have to deal with malware that you are uniquely qualified to detect / prevent / mitigate?

Finally, defenders cannot mix and match many of today's solutions, making combinatorial coverage impossible in many scenarios.

In contrast, PolySwarm will foster an ecosystem of broad coverage powered by 1000s of "micro-engine" workers, authored by geographically diverse security experts.




Figure B: PolySwarm will foster an ecosystem that produces broad threat coverage.

PolySwarm provides a legitimate revenue stream for security experts.

When local economics cannot support sufficient honest work, some security experts develop ransomware, operate bots and otherwise use their skills for evil. PolySwarm provides a region-free alternative: unencumbered reward for honest work. Experts compete to make the internet a safer place.

PolySwarm puts the user first.

The PolySwarm market directs economic incentives to where it matters most: toward accurate malintent detection ("threat intelligence"). Accuracy is determined via a novel process we call Mediated Consensus (details on page 10).

In short, PolySwarm provides users with timely access to broad, crowdsourced security expertise. Need to quickly triage a suspect file? Just Swarm It™.

Reinventing the Threat Intelligence Market

PolySwarm's Nectar ("NCT") tokens form the basis of a new market that introduces novel instruments for satisfying demand for timely and accurate assertions regarding the malintent of files, network traffic, and URLs, collectively referred to as Artifacts. These new instruments are structured to directly incentivize innovation in the threat intelligence space with a feedback loop driven by accurate results.



Enterprises



Experts



Ambassadors



Arbiters

PolySwarm employs Fees to discourage spam and incentivize honest and active market engagement. Fees are assessed on transaction types that may be abused for spam, and then distributed to active ecosystem participants; those who are introducing Artifacts and determining ground truth regarding the malintent of Artifacts. This participation is measured in a sliding window fashion, "ageing-off" older market contributions, thereby incentivizing continued market participation.

PolySwarm 1.0 (and this document) will focus exclusively on boolean (malicious / benign) determination, but the PolySwarm team has bigger plans than disrupting boolean malintent determination. During PolySwarm 1.0 development, the team will investigate methods to incentivize the production of artifact metadata such as malware family and how core PolySwarm concepts such as smart contract market design, Mediated Consensus and an information security focus can disrupt related markets such as vulnerability bug bounties⁴.

Before introducing PolySwarm's instruments, we introduce the participant classes that will utilize them.

⁴ "HackerOne but distributed, pseudonymous, and unfettered by jurisdictional encumbrances. Powered by blockchain." Not covered by this document.

The Participants

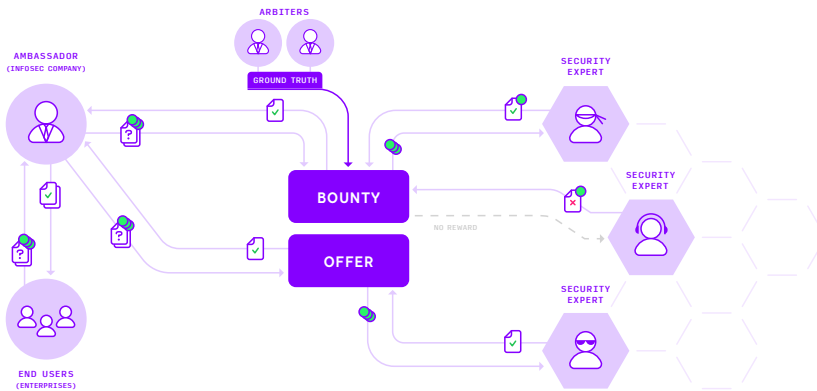


Figure C: Overview of PolySwarm’s Bounty and Offer Lifecycles.



End Users: Enterprise and home users with suspect Artifacts. End Users participate in the PolySwarm market via Bounties and Offers (more on these in a moment) and extract timely and accurate malintent classifications



Security Experts (“Experts”): Geographically diverse malware experts and reverse engineers. Experts dissect the latest suspect Artifacts and maintain PolySwarm-connected detection engines (“workers”) that determine malintent. Experts commit to “Assertions”, public statements that reflect the results of their analysis into the malintent of the Artifact. Those that have committed an accurate Assertion are rewarded (in NCT) for their efforts.

As far as core participants go, that’s actually it. Simplicity is good, but this is pretty far from the full PolySwarm story. Technically, PolySwarm could work with only these two classes of participants. Realistically, it won’t because most **End Users** will prefer to outsource the nitty gritty of interfacing with the PolySwarm market. Before going further, we must introduce **Ambassadors** (another participant) and a certain subclass of these Ambassadors.



Ambassadors: Companies that make it easy for End Users to benefit from the PolySwarm market. Ambassadors collect traditional fiat (e.g. subscription fees) and suspect Artifacts from their clients (End Users) and introduce **Bounties and Offers** into the market on their clients' behalf. It is the Ambassador's responsibility to distill the Assertions of various Experts into a simple **malicious or benign Verdict** that they deliver to their clients.

A trivial, perhaps naive approach for this distillation might be to simply average the Assertions provided by Experts. It is unlikely, however, that such an algorithm would compete favorably against a Bayesian analysis⁵, let alone expert human involvement. We expect existing Antivirus and threat intelligence firms to participate in the PolySwarm market as early Ambassadors, augmenting their in-house expertise with PolySwarm-enabled triage of suspect Artifacts.

From the User's perspective, upgrading to PolySwarm-backed protection is an easy process: choose a reputable Ambassador and pay a subscription fee.

An Ambassador's reputation is based on past Verdict performance relative to ground truth. Ambassadors are incentivized to make their Verdicts public due to their inherent desire to build reputation and attract new clients, as well as take advantage of a discount on Fees (this process is detailed later). These public Verdicts enable the creation of an Ambassador "scorecard" that rates real-world performance against non-synthetic Artifacts with data that is simply not available in today's market⁶.



Arbiters: Top-echelon Ambassadors that are responsible for determining malintent ground truth. A certain percentage of Ambassadors (in terms of Fees generated) will be considered "Arbiters".

During development, PolySwarm will assign Arbitership to existing, reputable threat intelligence vendors that are willing to maintain frequent engagement with the PolySwarm team, help identify and address platform bugs, and help build interest in the ecosystem.

Once PolySwarm 1.0 is ready for launch, these designated Arbiters will need to maintain top-echelon Ambassador volumes to maintain Arbiter status.

⁵ Kantchelian, Alex, et al. "Better malware ground truth: Techniques for weighting anti-virus vendor labels." Proceedings of the 8th ACM Workshop on Artificial Intelligence and Security. ACM, 2015.

⁶ The closest analog is AV Comparatives (<https://www.av-comparatives.org/>).

Prediction Markets, Arbiters & Mediated Consensus

In the PolySwarm ecosystem, Security Experts develop micro-engine workers that compete to quickly & accurately investigate suspect artifacts. This investigation occurs at machine speed - well before the ground truth concerning an artifact's malintent has been established.

This design shares some similarities with prediction markets, i.e. rewarding past assertions based on accuracy, but differs in two critical respects:

- 1. No future data is required to accurately classify an artifact.** The "correct" answer to whether an artifact is malicious can always be determined with certainty at the time it is made available to workers. PolySwarm is a skill-based design where nothing is left to chance⁷.
- 2. Determining ground truth will always require expertise.** This is in contrast to the "universal observability" of prediction markets (including crypto markets such as Augur⁸) that rely on unskilled participants to observe and record events as they occur.

Item #1 is PolySwarm's gambling deterrent⁹.

Item #2 introduces a technical challenge: what is the best way to incentivize authorities to continually produce ground truth? PolySwarm's answer is the Arbiter class and a process we refer to as Mediated Consensus.

Mediated Consensus is a generic design paradigm that we hope will find a home in other market design projects. In short, Mediated Consensus is market design that entrusts a critical task to a subset of participants. These participants:

1. Are qualified to complete the task (possess expertise).
2. Have their interests aligned with the overall health of the market (avoiding tragedy of the commons).

⁷ In theory. In practice, at least two events may introduce uncertainty: (1) arbiters incorrectly determine ground truth resulting in reward distribution to wrong experts, (2) the expert's analysis is accurate according to ground truth, but the expert disagrees on the boundaries of malintent - a semantic issue that is larger than PolySwarm. Is adware malware or simply a potentially unwanted application? Neither of these variables fundamentally detracts from the skill-based nature of the PolySwarm ecosystem.

⁸ <https://augur.net/>

⁹ Guessing, betting, gambling or other chance-based wagers on artifact malintent is bad for everyone and PolySwarm is specifically designed to be inhospitable to this misuse of the platform.

We've designed Arbiters to satisfy these two constraints. By definition, the Arbiter class is defined as the most active Ambassadors in the PolySwarm ecosystem at any given time.

As Ambassadors, Arbiters are contracted by their customers to distill Security Experts' assertions into Verdicts. Customers trust them to possess and exercise automated – and when necessary – human expertise. These companies already have a vested interest in maintaining a public record of accuracy (item #1).

As the most active Ambassadors, Arbiters have the most to gain / lose by an honest / dishonest PolySwarm ecosystem. Their large stake in the faith of the ecosystem aligns their financial interests with the health of the market (item #2).

We've designed Fees as a supplemental defense against Arbiter abuse such as Arbiter–Security Expert collusion.

Determining Ground Truth

Most of the technical details of the Arbiter voting process are intentionally left undefined at this time. The reason for this is simple: the specifics of the Arbiter decision process are predicated on the specifics of virtually all other PolySwarm processes – and we expect variations in other processes.

That being said, we expect the following high level design choices to provide value in this process:

1. Arbiters reach consensus on ground truth via a **majority vote**.
2. These votes are mined onto the **Ethereum chain en masse** (many votes per cycle), saving time and money (Ethereum gas) for the PolySwarm ecosystem.
3. To encourage participation, **Arbiters are rewarded with Fees for voting on the ground truth of artifacts**.
4. **Arbiters may abstain from voting** on any particular artifact. Arbiters may choose to do this if, for example, they feel unqualified to determine the malintent of a particular artifact.
5. As necessary, Arbiter **voting privileges are automatically delegated to additional Ambassadors (in order of volume)** to ensure quorum on all ground truth determinations.

6. As necessary, **we expect participants to challenge Arbiters' determinations as is done today:** with a blog post or technical paper describing the malintent or benign nature of an artifact that an authority such as an Arbiter had miscategorized: an external PR-driven feedback loop.

The PolySwarm team will iteratively develop the details concerning Arbiter voting and incentive structures during development. We expect early Arbiter feedback will be instrumental to this design.

Topics of research include, but are not limited to:

- Minimum Arbiter quorum percentage
- Arbiter deferral procedure (for when quorum cannot be achieved)
- Appropriate Fee reward structure
- Incentive for timely ground truth determination, e.g. only first X Arbiters to vote receive reward
- Penalty for non-participation, e.g. Arbiters must vote on ground truth for X% of artifacts in a time window, else their Arbiter status is revoked
- (As necessary) additional deterrents against perverse incentives

The Instruments

PolySwarm exposes two core instruments to End Users and Ambassadors¹⁰ that increase the efficacy of the threat intelligence market:

Polyswarm Offers: Requests made directly to reputable Security Experts for their malintent prediction. PolySwarm provides frictionless access to thousands of such researchers, enabling traditional information sharing agreements with non-traditional participants. This interaction happens inside Raiden¹¹-style Offer Channels (detailed later). Once Channels are established, Offers provide millisecond-scale latency for artifact investigation.

OFFER

¹⁰ Hereafter, End Users and Ambassadors is shortened to Ambassadors. End Users may act as their own Ambassador. We expect some large enterprises to participate in this manner.

¹¹ <http://raiden.network/>

Polyswarm Bounties: A wild-west style “Wanted” poster with accompanying Artifact contents (e.g., Wanted: This Artifact: Malicious or Benign? Reward: 1000 NCT). Security Experts make a name for themselves (build reputation) by successfully competing for Bounties. No direct analog exists in today’s market.

Offers are the closest analogue to today’s on-demand scanning market and operate with millisecond latency.

Ambassadors issue *Offers* and *Artifacts* directly to chosen Security Experts, optionally under a non-disclosure agreement. Each Expert chooses whether to accept the *Offer* based on their confidence in rendering an accurate *Assertion* for the provided *Artifact*. Experts may choose to decline if they are not confident so as to avoid adversely affecting their reputation. If the Expert accepts the *Offer*, the Expert commits to providing a malintent *Assertion* in a timely fashion. Token collection and awards for all instruments are managed and executed entirely by distributed smart contracts.

Issuing an *Offer* requires a certain familiarity with Experts best equipped to dissect the *Artifact* in question. To ease this matchmaking, Experts advertise their specialties via listings in the PolySwarm Worker Registry – PolySwarm’s analog to Ethereum dApp Registries¹² and build public reputation by successfully participating in similar Bounties.

Bounties are cheaper than *Offers* and require no upfront familiarity with specific Experts, but Bounties are not for everyone.

First, Bounties must be mined into an Ethereum block, which occurs approximately every 15 seconds – by far the dominant time cost in this arrangement. Second, when placing a Bounty, the *Artifact* must be made public¹³. After all, wild-west “Wanted” posters wouldn’t be of much help if they weren’t posted in a public place with all known information. In a similar manner, Bounties represent a public, smart (as in contract) commitment to reward providers of information that leads to the quarantine or exoneration of *Artifacts*. Bounties also form a critical component of the feedback loop that establishes ground truth.

¹² See the [Worker Registry](#) section on page 23 for more details.

¹³ True of PolySwarm 1.0, not necessarily true in future iterations. [See Artifact Confidentiality](#).

Rewarding Accuracy

In a departure from today's market, the PolySwarm market offers precise rewards based solely on the accuracy of threat intelligence, incentivizing Experts to optimize exclusively for detection accuracy (minimal false positives and false negatives), and enabling Ambassadors to extract maximal value on behalf of their clients. PolySwarm defines accuracy as agreement (or not) between Experts' Assertions and Arbiter-defined ground truth.

This accuracy feedback loop is driven by PolySwarm Bounties. Each Bounty rewards Experts who render an accurate Assertion and penalizes Experts making an inaccurate Assertion – all whilst avoiding the introduction of collusion incentives between Ambassadors and Experts¹⁴. PolySwarm Offers provide a convenient means of achieving traditional 1:1 business relationships, but do not factor into this accuracy equation.

Ground truth is produced and consumed in the PolySwarm marketplace in the following manner:

1. An Ambassador places a Bounty on an Artifact, submitting a Fee to do so.
2. Various Experts render Assertions on this Artifact prior to the Bounty's Assertion deadline. Each Assertion is accompanied by an Expert-chosen NCT amount (a "Bid") that reflects the Expert's confidence in their Assertion. A Fee is assessed as a percentage of this "Bid". These Assertions are confidential up until the Assertion deadline.
3. The Ambassador produces a Verdict, taking Experts' Assertions into account however they see fit and delivers this Verdict to their client. The Ambassador is incentivized to make this Verdict public in order to build their reputation and take advantage of a Fee discount (detailed later).
4. At a later time (e.g., 7 days after the Artifact was placed on Bounty), Arbiters are offered the opportunity to vote on whether the Artifact is in fact malicious or benign. This later-determined ground truth IS NOT A BLOCKER FOR RAPIDLY RETURNING ARTIFACT VERDICTS. In other words, Ambassadors need not (and should not) wait for ground truth determination before returning a result to their customer.

¹⁴ This includes potential collusion interest between Experts and Arbiters who vote on their own Bounties. The PolySwarm Fee structure is designed such that no single Arbiter could sway ground truth sufficiently to cause a colluding Expert to receive a Bounty greater than the Arbiter's cost for introducing doubt into the market.

5. A quorum of Arbiters determines the ground truth status of the Artifact. The details of this vote to ground truth conversion process are yet to be determined. Possible arrangements are simple majority vote or proportional majority vote (e.g., based on amount of market participation). These specifics will be determined during prototype development.
6. With ground truth established, the Bounty smart contract awards NCT to Experts who rendered accurate Assertions. The amount of NCT awarded is proportional to the Expert's Bid amount relative to the total pool of accurate Bids.

On the surface, it may seem that a classic prediction market would fit well here: Experts render Assertions and, at a later date, these Assertions are compared against ground truth, ultimately triggering rewards and penalties based on their accuracy. The trouble for PolySwarm is that classic prediction markets implicitly assume a source of ground truth that is verifiable by all participants. A prediction market that asks who will win an election or whether the price of gold will exceed an amount by a certain day is easily settled because all participants can verify these data points after the prediction period has closed: Who won the election? Did gold exceed the strike price?

Determining the malintent of a suspect Artifact demands expertise that is not uniformly shared among PolySwarm participants. This conundrum has analogues to prediction markets applied to medical diagnoses¹⁵, where doctors bring varying degrees of expertise to bear.

PolySwarm entrusts ground truth determination to Arbiters. In today's market, the role of the Arbiter is effectively filled by traditional anti-virus companies. These companies "vote" (and are held publicly accountable) on "ground truth" via services like VirusTotal. In the PolySwarm market, the set of Arbiters is determined solely by market participation, offering opportunity for newer players to shake up the establishment.

In summary, we believe this process properly incentivizes Arbiters to commit resources to conduct due diligence when determining ground truth for the benefit of the market as a whole.

This concept of an elite group of ground truth "verifiers" has applicability to other sectors. We anticipate that others will explore the feasibility of using this mechanism in other token platforms that demand specialized knowledge to arrive at ground truth.

¹⁵ Kurvers, Ralf HJM, et al. "Boosting medical diagnostics by pooling independent judgments." Proceedings of the National Academy of Sciences (2016): [201601827].

Fees

Fees are assessed for various transactions in the PolySwarm market and are designed to achieve two goals:

1. **Promote efficiency in the PolySwarm market by incentivizing actions that incur minimal Ethereum gas cost.** Fixed-rate Fees are applied to PolySwarm transactions to discourage behavior that would generate superfluous transactions (including spam transactions) and therefore incur unnecessary gas cost across the market. PolySwarm's need to scale these Fees independently of Ethereum gas cost is one of several reasons the creation of NCT is essential to the existence of the PolySwarm market.
2. **Reward active PolySwarm market participants in proportion to their (honest) participation.** Fees are awarded to participants that are actively utilizing the PolySwarm ecosystem via the introduction of Artifacts (Ambassadors placing Bounties) and the determination of ground truth regarding the malintent of Artifacts (Arbiters reaching quorum on ground truth).

During network development, the PolySwarm team will red team the Fee structures presented above, iterating as necessary to best encourage a healthy market.

Bounty Placement Fees Paid by Ambassadors

When an Ambassador lists a Bounty on the PolySwarm market, the following Fees are rendered on the funds that are held in the smart contract:

1. **A fixed listing Fee.** This fixed fee encourages the Ambassador to group multiple Artifacts into a single Bounty, reducing strain on the network and gas costs for all.
2. **A Fee proportional to the Bounty amount.** Bounties with a higher initial Bounty amount will likely attract more Expert responses. This proportional Fee is assessed to scale the Fee with Expert interest (and network strain).

Bounty Assertions Fees Paid by Experts

When an Expert renders an Assertion against a Bounty, the following Fees are rendered on the funds that are held in the smart contract:

1. **A fixed Assertion Fee.** This fixed Fee disincentivizes repeated Assertions that may slow down the network.
2. **A Fee proportional to the Bid amount.**

Offer Channel Settlement: Ambassador Fees

Offers occur within Raiden-style Offer Channels, established directly between Ambassadors and Experts. When an Ambassador and/or Expert decide to settle their Offer Channel onto the blockchain, the following Fee is assessed by the Channel smart contract:

1. **A percentage of the net NCT transferred over the channel.** This Fee is assessed on the Ambassador-supplied tokens when the channel is opened. When the channel is closed, these Fees are refunded to the Ambassador in proportion to unused channel tokens.

Protocol Details

Here we detail Polyswarm Bounties, Offers, Assertions and Verdicts. Bounties and Offers are collectively referred to as Polyswarm Listings. A Polyswarm Statement is any of: Bounty, Offer, Assertion or Verdict.



Bounties




Offers



Assertions



Verdicts



Bounty Lifecycle

When an Ambassador wishes to place a Bounty on the market, the Ambassador registers a Statement via a call to the PolySwarm Bounty smart contract.

The **Bounty Statement** contains the following items:



1. **Listing GUID:** A globally-unique identifier for the Listing.
2. **The Ambassador's Identity.** This takes the form of a signature on the Statement, verifiable by anyone in possession of the Ambassador's address (public keys are derived from addresses).
3. **A Bounty Placement Fee.** This Fee is a fixed amount + a percentage of the Bounty Amount (described in "Fee" section above). A discount is applied if the Ambassador later publicly commits to a Verdict.
4. **The Bounty Amount.** This amount is placed into the smart contract alongside the Bounty Placement Fee and is awarded to Experts that render accurate Assertions (or refunded if no Experts respond). This amount is akin to a transaction fee in Ethereum or Bitcoin and is useful to call attention to the Bounty.
5. **Cryptographic hash of the Artifact.** Configurable, but currently SHA256. This uniquely identifies the Artifact.
6. **The URI of the Artifact.** Due to storage costs, Artifacts are not stored on-chain. We anticipate a secondary market to arise for outsourcing storage and delivery of Artifacts¹⁶ and that some Ambassadors will choose to retain a copy of all swarmed Artifacts for later competitive analysis.
7. **Assertion Deadline:** Deadline for Assertion. The Ambassador may choose to shorten or lengthen analysis periods based on their requirements for the particular Artifact. Artifacts that require a more in-depth analysis likely require more time (and also likely a higher Bounty amount to attract attention), whereas other Artifacts may require a quick turnaround.

Once the Bounty is listed, Experts will determine whether the Artifact fits their skillset and whether their expected reward warrants their attention. This expected reward is a function of their confidence interval regarding the malintent of the Artifact.

Experts that choose to participate place Bids on their Assertions that represent their level of confidence (and risk tolerance). The Bounty amount plus the sum total of all inaccurate Bids are awarded to accurate Bids in proportion to their Bid amount relative to the total accurate Bid amount.

Consider the following example (Fees are omitted here for simplicity):

1. An Ambassador places a Bounty with 2 NCT. This 2 NCT is placed into the Bounty smart contract. Thus far, this 2 NCT Bounty Amount constitutes the entire “pot” of potential winnings for accurate Experts.
2. Experts A and B study the Artifact and decide to Assert that the Bounty Artifact is *benign*. They place a 5 and 3 NCT bid on their Assertions, respectively.
3. Experts C and D study the Artifact and decide to Assert that the Bounty Artifact is *malicious*. They place a 1 and 4 NCT bid on their Assertions, respectively.
4. At some later time, the Arbiters establish ground truth and determine that the artifact was in fact *malicious*. Experts A and B therefore rendered *inaccurate* Assertions; Experts C and D rendered *accurate* Assertions.
5. The Bounty amount plus the Bids placed by Experts A and B make up the “pot” that is proportionally awarded to Experts C and D. This pot is $2 + 5 + 3 = 10$ NCT.
6. Expert C made an accurate Assertion and her Bid represented $1/5$ of the total accurate Bid amount. Expert C has her Bid returned plus her portion of the pot: $1 + 1/5 * 10 = 3$ NCT.
7. Expert D made an accurate Assertion and his Bid represented $4/5$ of the total Bid amount. Expert D gets his bid back plus his portion of the pot: $4 + 4/5 * 10 = 12$ NCT.

Bounty Assertions contain:



1. **Listing GUID + Cryptographic Hash:** A unique reference to the prompting Bounty Listing.
2. **The Expert’s Identity:** Again, as a signature on the Statement.
3. **An Assertion Fee:** This Fee is a fixed amount + a percentage of the Bid amount (see “Fees” section).

4. **Bid Amount.**
5. **The Assertion:** A boolean “Malicious” or “Benign” determination regarding the malintent of the Artifact. Assertions are kept secret until the Bounty’s Assertion deadline expires. We anticipate using blind commitments + a reveal phase to achieve this property.
6. **(Optional) Metadata:** Information derived during Assertion generation and offered as a value-add to the Ambassador to utilize in their Verdict computation (e.g., malware family). Experts may volunteer this information as a means to differentiate themselves and attract future business.

With Assertions streaming in, the Ambassador weighs Assertions and incorporates their own internal due diligence efforts, producing a “malicious” or “benign” Verdict. The Ambassador delivers this Verdict to their client as their ultimate response regarding the malintent of the Artifact. The Ambassador may also choose to make this Verdict public (and is incentivized to do so via a Fee discount) before the Verdict deadline.

Public Verdicts contains the following items:



1. **Listing GUID:** A reference to the prompting Listing.
2. **The Ambassador’s NCT Identity:** Again, as a signature on the Statement.
3. **The Verdict:** A boolean “malicious” or “benign” determination regarding the malintent of the Artifact.

At some later time, the Arbiters reach quorum on the ground truth regarding the Artifact. This ground truth is compared against each Experts’ Assertions and NCT is awarded for accurate Assertions.



Offer Lifecycle

When an Ambassador wishes to issue Offers to a chosen Expert, the Ambassador opens a Raiden-style channel with the Expert.



This is done by instantiating an **Offer Channel** smart contract with the Expert that includes the following:

1. **Identities of Ambassador and Expert.**
2. **Channel Balance:** The maximum net NCT that the Ambassador is comfortable settling with the Expert at channel close time (plus Fees).



With the Offer Channel established, the Ambassador issues zero or more Offers to the Expert. Each Offer includes the following:

1. **Listing GUID:** A globally-unique identifier for the Listing.
2. **The Offer Amount.** This Amount must not exceed the remaining balance on the Offer Channel smart contract (minus Fees).
3. **Cryptographic hash of the Artifact.**
4. **The URI of the Artifact.** The Artifact is optionally encrypted to the Offeree's identity or otherwise access protected to ensure confidentiality.
5. **Engagement Deadline:** Deadline for Engagement Response (described below).
6. **Assertion Deadline:** Deadline for Assertion on this Offer.

When presented with an Offer, the Expert (Offeree) may choose to engage with the Offer or ignore it. Each Offeree determines whether this Artifact is worth evaluating for the Offered amount in the time allotted.

The Offeree's future business is based on their ability to produce accurate results. If the Offeree is not confident in her ability to produce an accurate result in the time allotted (e.g., the Artifact is a filetype that the Offeree is ill-equipped to analyze), she may choose to reject the Offer.

If the Offeree decides to reject the Offer, she either actively issues a “decline” Engagement Statement or simply waits for the Engagement Deadline to expire. It likely behooves the Offeree to actively reject, rather than wait for the Engagement Deadline, as Ambassadors will likely favor Experts who do not force them to wait until the Engagement Deadline to determine next steps.

If the Offeree accepts the Offer, she issues an “accept” Engagement Statement prior to the Engagement deadline.

Engagement Statements contain the following items:



1. **Listing GUID + Hash:** A reference to the prompting Listing.
2. **Engagement Commitment:** A boolean “accept” or “decline” that states whether the Offeree will engage with this Offer and commit an Assertion prior to the Offer’s Assertion Deadline.

Once the Offeree accepts an Offer, the Offeree has until the Assertion Deadline to render an Assertion against the Offer.

These Assertions contain the following:



1. **Listing GUID + Hash:** Unique reference to the prompting Listing.
2. **The Assertion:** A boolean “malicious” or “benign” determination regarding the malintent of the Artifact.
3. **(Optional) Metadata:** Information derived during Assertion generation and offered as a value-add to the Ambassador to utilize in their Verdict computation (e.g., malware family). Experts may volunteer this information as a means to differentiate themselves and attract future engagements.

After the Expert delivers her Assertion to the Ambassador, the Ambassador signs a Raiden-style message that commits to also deliver the Offer amount, so at no time is the Expert risking more than the output of a single analysis. The sum of all such delivered amounts are settled to the blockchain when the Offer Channel is closed.

Unlike Bounties, Offers do not incorporate an on-chain ground truth feedback mechanism. Instead, Offerees who either fail to render an Assertion before the deadline or render an inaccurate Assertion (as determined by later off-chain analyses) may not attract future engagements. Reputation maintenance therefore becomes akin to an off-chain feedback loop that uses market pressure to keep Offerees honest.

When an Offer Channel is closed, a Fee is collected on the net NCT transferred. As with all Fees, this Fee is awarded to active market participants. The Fee is enforced by the master Offer Channel smart contract which refuses to establish non-conforming Offer Channels.

Reputation

In contrast with other token platforms (e.g., Augur), PolySwarm does not attempt to formalize the notion of reputation and build it into market transactions.

The reason for this is simple: it is incredibly difficult to anticipate which data points various participants value – and how these data points are weighed. The best fit for Ambassador A may not be the best fit for Ambassador B and it would be futile to try to get them to both operate under the assumption that it was. Instead, PolySwarm mechanics simply mandate or incentivize disclosure of the kinds of data that likely represents useful input into a reputation function.

Reputation algorithms will likely be regarded as a trade secret in the same manner as Assertion to Verdict distillation algorithms. Ambassadors will differentiate themselves by engaging with optimal Experts for their Artifact workload. Experts will differentiate themselves by maintaining a favorable track record of correct Assertions when compared against ground truth. And, optionally, by providing metadata, for example, the malware family they believe a particular artifact belongs to.

The prototype Ambassador implementation to be developed will distill various data points into a simple reputation “score” and an easy-to-understand color scheme. This reference distillation will likely be helpful to many, but may be replaced or modified as participants see fit. Experts will likely also want to implement a reputation system for evaluating Ambassadors, but the business logic of such a system are expected to be more complex and differ more substantially among participants.

The reference PolySwarm daemon will “decay” old data, maintaining a “sliding window” into participants’ reputation. We believe a sliding window approach is necessary to avoid lockout of new market participants.

We anticipate making use of the following data points in a reference Ambassador implementation:

1. **Assertion accuracy** (relative to ground truth and/or off-chain analyses).
2. **Frequency and average response times for both Bounties and Offers.** In the case of Offers, Ambassadors may choose to publish aggregate data that is normally only observable within an Offer Channel. Ambassadors may favor Experts who demonstrate rapid and consistent response times.
3. **Ability to maintain confidentiality of Artifacts delivered under Offers.** This is likely a hybrid on-market, off-market analysis.

Ambassadors may choose to review their competitors' past Verdicts and attempt to uncover errors (e.g., Ambassador A stated that the sample was benign when in fact it was malicious). Experts may seek to avoid doing business with Ambassadors whose due diligence fails to hold up to such scrutiny in general or is more frequently incorrect under specific circumstances (e.g. Ambassador A frequently miscategorizes Microsoft Word documents resulting in certain Experts avoiding analyzing Word documents on behalf of Ambassador A). In other words, among other inputs, Experts may be interested in Ambassadors' Verdict accuracy.

Consumers will likely rely on aggregate Consumer Reports or AV Comparatives-style websites and publications that report on each Ambassador's track record, perhaps with sample inputs with known malintent and verification that an Ambassador's Verdicts always match what it represents back to its clients.

Worker Registry

Experts in the PolySwarm ecosystem will, in many cases, choose to encapsulate their expertise into automated Workers. Workers are systems that automatically respond to Bounties and Offers, perhaps with expertise for certain filetypes or network traffic. The PolySwarm marketplace will provide a distributed Worker Registry to ease discovery of Experts' offerings.

We will develop a common interface and description language for publicly registering Workers. These components will interact with a PolySwarm Registry that is modeled after initial Ethereum dApp registry smart contracts. PolySwarm's Worker Registry will contain descriptions that advertise particular skillsets (e.g., "I'm good with Mac OS Mach-O binaries" or "I handle Microsoft Word Documents").

Artifact Confidentiality

Confidential disclosure of Artifacts is a first class feature in the PolySwarm marketplace via the *Offer* instrument. Nevertheless, without taking additional precautions, Ambassadors utilizing Offers must trust Security Experts' motivations and technical ability to ward off attackers, possibly in perpetuity.

We expect this trust requirement to give rise to secondary markets and novel applications of technology with the goal of minimizing or consolidating trust. The following are several examples of instruments built on top of PolySwarm to further address confidentiality requirements.

Trust Minimization via Intel SGX

Intel SGX is a new technology initially shipped with the Skylake family of processors. At a high level, SGX provides opaque computation environments inside "enclaves". Memory contents of SGX enclaves are unavailable to other software and users on the same system, irrespective of permission level. In the context of PolySwarm, SGX enclaves would allow Experts to scan Artifacts without revealing the Artifacts to the Researcher and without revealing the Experts' intelligence to anyone else. SGX therefore promises mutual trust reduction between Ambassadors and Experts.

Experts may choose to run accredited (and attested) SGX enclaves that ingest the Artifact and their intelligence and produce an Assertion that may be returned to the Ambassador. There are technical hurdles to overcome, but we are confident SGX can be integrated and scaled with enclaves bootstrapping one another in a distributed fashion. Future iterations of PolySwarm may allow Ambassadors to author future "SGX Bounty" statements into the PolySwarm market, revealing the Artifacts only to bootstrapped enclaves.

Trust Consolidation via Scan-as-a-Service (SaaS) Ambassadors

Experts may choose to outsource computation tasks to third parties for a number of reasons that may or may not be PolySwarm-specific. For example, we expect many Experts to make use of AWS, Google Compute Engine, Azure or other such on-demand computational services to scale their Assertion pipeline in response to PolySwarm market demand. This is not an intrinsically PolySwarm-motivated compute decision.

We expect PolySwarm-specific incentives for SaaS entities to arise naturally. In the context of confidentiality, participants will create SaaS companies that serve as trusted third parties sitting between Ambassadors and Experts and maintaining nondisclosure agreements with both. These SaaS companies would serve the same purpose as SGX enclaves – protecting Artifacts from Experts and intelligence details from Ambassadors.

Ambassadors may find it easier to trust third party SaaS entities with public reputation rather than a large number of individual Experts whose identity may not be known. Such Ambassadors will author more Offers targeted at SaaS entities rather than directly at Experts.

Additional Markets

The PolySwarm market will create demand for a number of additional markets. One such market (Scan-as-a-Service) was previously discussed in the context enabling additional confidentiality controls. Below are some additional expected markets.

Ambassadors

This is an obvious market and one heavily discussed throughout this document. Many Consumers will choose to do business with the market through an Ambassador, but working through an Ambassador is not a strict requirement.

Artifact Publishing Services

When placing a Bounty or Offer, the Ambassador must make their Artifact available to Experts. Artifact hosting entities will offer this service, making such sharing seamless. Their clients will be Consumers and Ambassadors. We expect many Ambassadors to also offer Artifact Publishing services.

Reputation Tracking Services

In the United States, three credit bureaus are trusted by banks to maintain a ledger of “facts” regarding individual trustworthiness. The PolySwarm market will present a similar demand; market participants will want to quantify the trustworthiness of other participants for a variety of reasons already discussed. It is likely that reputation services will arise that track trustworthiness of participants and offer this information to mitigate counterparty risk.

Roadmap

Here we present preliminary PolySwarm development milestones. Each milestone will: (1) offer new functionality and (2) provide documentation and tests that enable new Ambassador, Expert, and End User transactions.

PolySwarm is at a distinct developmental advantage relative to networks like Bitcoin and Ethereum that must retain all transactions in perpetuity. Once ground truth is determined or an Offer Channel is settled to the chain, Ambassadors and Experts have received what they need out of that transaction: intelligence regarding the Artifact and NCT, respectively. This atomicity provides leeway and allows us to make backwards-incompatible changes, up until the 1.0 release, without cost and consensus challenges experienced in other networks.

Vo.1 Alpha (30 April 2018)

PolySwarm Alpha will focus on delivering a prototype end-to-end workflow for Bounties after successful token sale. This initial release will allow the PolySwarm team to begin the bootstrapping process in earnest. Ambassadors and Experts will be able to transact via a PolySwarm test network (“testnet”). The basic Bounty smart contracts will be available on the PolySwarm testnet, allowing Ambassadors to place Bounties, Experts to render Assertions, and Ambassadors to render Verdicts.

We anticipate that PolySwarm Alpha will accomplish the following:

- A “Bounty Manager” smart contract which holds Bounty Amounts and Fees, accepts Assertions, publishes Verdicts, and rewards accurate Experts
- Artifact posting and retrieval to / from off-PolySwarm storage (e.g. S3 and/or IPFS)
- Prototype algorithm for selecting Arbiters at Bounty expiration
- Reference implementations for Bounty publishing, Assertion response, and Arbiter ground truth determination

Vo.2 Beta (31 May 2018)

PolySwarm Beta will focus on delivering PolySwarm Offer support over Raiden-style Channels. Offers are designed to facilitate higher-throughput Artifact evaluation relative to Bounties and as such demand more efficient blockchain settlement for each transaction. PolySwarm Beta will implement Offer smart contracts. We anticipate implementing this functionality as direct, peer to peer, Raiden-style Channels established between Ambassadors and Experts which utilize the testnet.

We anticipate PolySwarm Beta will accomplish the following:

- Offer Channels
 - Distributed Channel establishment
 - multiple Offers via a single Channel
 - NCT settlement at Channel Close
- Reference implementation for Offer production and response
- Reference implementation for Arbiter ground truth notification and response
- Fee collection and distribution smart contracts
- Research and quantify trust (in NCT) required for a specific Artifact volume
- Offer/Bounty development toolkit: a test and implementation framework for Experts and Ambassadors to automate servicing of Bounties and Offers

Vo.3 Gamma (31 July 2018)

PolySwarm Gamma will focus on providing matchmaking features that facilitate collaboration between Ambassadors, Experts, and Arbiters. Experts will be able to encapsulate their expertise into Workers tuned to specific types of Artifact. We expect them to be able to advertise these Workers in a PolySwarm Registry. This Worker Registry should allow Ambassadors to engage with Workers in a risk-free manner prior to the Stable release (Gamma will continue to utilize the PolySwarm testnet).

Additionally, this release will consider enhancing and extending the Bounty instrument in light of Alpha and Beta feedback and experience. As one example, we anticipate that Bounties will benefit from the ability to solicit Assertions that are kept confidential from other Experts and revealed only after the Assertion deadline.

We anticipate PolySwarm Gamma will accomplish the following:

- Worker Description Language (WDL) describing artifact analysis capabilities and Expert authorship attribution (for reputation)
- Distributed, searchable, registry of Workers and their capabilities.
- Confidentiality for pre-deadline Bounty Assertions
- Finalize Arbiter selection, response, and end to end Bounty feedback loop

We anticipate that Gamma will be feature complete.

Polyswarm 1.0 (Q4 2018)

PolySwarm Stable will focus on eliminating bugs in PolySwarm Gamma prior to migration from the PolySwarm testnet (test tokens) to the real PolySwarm market.

As a stretch goal and as time / funding allows, Stable may introduce first-class Artifact confidentiality via a limited-audience Bounty.

Polyswarm 2.0 (Q2 2019)

PolySwarm's second stable release currently has two main goals: expanding Offers and Bounties to handle a wider range of Artifact types (e.g. URLs and Network streams) and easing PolySwarm utilization for the Enterprise and Home User.


Artifact type expansion in future PolySwarm releases should make it possible to handle, for example, automated phishing attempt detection of end user URLs by placing Bounties and Offers directly on hosted content.

Extending PolySwarm's protections directly to the endpoint will be a priority for post-Stable development. We anticipate developing an Open Source reference implementation of an Endpoint protection suite which automatically blocks malicious artifacts based on PolySwarm's intelligence. We see this open agent framework as a key to widespread end user adoption.

Future Features

PolySwarm's future releases will mirror the evolving threat landscape and real-world usage of the network. We see several opportunities, contingent on observed PolySwarm usage, for the following features to benefit the network and the ultimate goal of end user protection:

- Flexible Artifact confidentiality
 - Basic: Artifacts selectively disclosed to Experts or SaaS providers
 - Advanced: Artifacts never revealed to anyone outside of SGX enclaves of fully homomorphic engines
- New Instruments (e.g. subscription-ready threat feeds)
- Deployment support for endpoint-resident Worker containers that supplant today's endpoint protection suites.



Token Sale

Technical development, community engagement and participant outreach for the PolySwarm ecosystem will be conducted by PolySwarm Pte. Ltd. with funding derived from the sale of PolySwarm's Nectar utility tokens (NCT).

These tokens will be exchanged by participants for threat intelligence. NCT powers all the major PolySwarm interactions: Bounty placement, Offer channel establishment / teardown / exchange, assertion registration and ground truth determination. By selling utility-specific NCT tokens, exchange rates for NCT to threat intelligence should enjoy isolation from the volatility in the larger ETH market, ensuring active participation in PolySwarm without concern for fluctuations in ETH.

This section is intended to give a high level overview of PolySwarm's tokenomics. Should any conflict between this document and the Token Sale Agreement ("Agreement") arise, the Agreement shall take precedence. [The Token Sale Agreement is available online.](#)

Schedule

PolySwarm Pte. Ltd., Inc is currently operating a token PreSale for invited participants. The public NCT sale is scheduled to open to qualified buyers on February 20, 2018 at 19:00 UTC. It is scheduled to close March 22, 2018 at 19:00 UTC.

Purchase Requirements

NCT will be made available exclusively for purchase with Ether (ETH). Potential buyers are advised to secure sufficient ETH in advance of the sale.

The token sale will employ Know Your Customer (KYC) controls with a unique technical mechanism to tie specific purchases to KYC data points including, but not limited to: Country of Origin, email address, full name, IP address and acceptance of the terms in the Token Sale Agreement.

Caps / Limits

No new tokens will be created after the close of the token sale period.

The total number of NCT tokens will be determined during the token sale.

The maximum number of tokens is capped. This cap is based on the maximum funding USD cap, the price of NCT (in ETH) and the current ETH to USD exchange rate. Refer to the token sale web page for specifics on NCT caps, price, funding tranches and more.

Allocation

70% of NCT tokens will be sold during the Token Sale.

15% of the remaining 30% will be distributed to enterprises, vendors and Security Experts interested in helping bootstrap the PolySwarm ecosystem. Recipients will be chosen at PolySwarm Pte. Ltd.. sole discretion.

The final 15% of tokens will be used by PolySwarm Pte. Ltd. in a manner that best accelerates adoption of the PolySwarm ecosystem. As an illustrative example, tokens may be offered as incentive to participate in PolySwarm "hackathons". This allocation will be at PolySwarm Pte. Ltd.. sole discretion.

Distribution

We envision that ETH derived from the sale of PolySwarm's Nectar utility tokens ("NCT") will be allocated in the following manner:



54.51% – Protocol and Software Development
3.00% – Office Expenses
4.06% – Operational Overhead (travel, hosting, etc)
14.49% – Legal
13.62% – Marketing
10.32% – Tax

The PolySwarm team expects Legal spending to decrease in FY19/20. The surplus will be redistributed into Marketing, Overhead and Developer Salaries.

Our 10.32% estimated Tax is low, but not unreasonable. This is good for participants: more funding will be allocated directly to the development and promotion of PolySwarm.

Disclaimer

This whitepaper is intended to provide technical background to would-be token purchasers. The PolySwarm token sale will fund initial prototype development and testing of the economic instruments presented here. We expect, as would any prudent participant, that the details presented in this document may change during pre-Stable development and testing periods. PolySwarm Pte. Ltd. and token purchasers' interests are aligned to make PolySwarm a viable marketplace that truly disrupts the threat intelligence industry.

Conclusion

PolySwarm is a novel market and associated ERC20 token, NCT, that is posed to disrupt the threat intelligence industry by facilitating new methods of interaction between security experts, vendors and enterprises. The PolySwarm team will be conducting a token launch to fund software development and to cultivate a community around this new marketplace.

We hope you'll join us in protecting users from **All Angles™**.

