



# Cosmos Network Game of Stakes - Comprehensive Overview

v1.0 draft

# What is Game of Stakes?

---

Hosted by Cosmos Network, Game of Stakes (GoS) was a first-of-its-kind competition that tested the security of Cosmos' Byzantine-Fault-Tolerance (BFT) based Proof-of-Stake (PoS) consensus algorithm prior to their mainnet launch - in short, it was a simulation. Cosmos started the game to stress-test the network, lending insight into how collusion and deception would unravel in such a system. The competition attracted participants worldwide whose objective was to compete for the most stake.

In particular, GoS tested the cryptoeconomics of the Cosmos Network Validators. Validators are critical to the success of the Cosmos Network, as they validate and relay transactions and propose, verify, and finalize blocks. In GoS, Validators were encouraged to make any software modifications they desired, collude with other validators, find bugs in the network and other validators and exploit them, and ultimately accumulate as much staked Atoms as possible.

Upon genesis of the Cosmos Hub, 300,000 Atoms would be allocated to GoS participants: the winning validator (the validator with the most stake) would receive 10,000 Atoms and the remaining validators would receive a proportion of 280,000 based on their ranking.

## Validator Requirements and GoS Rules

GoS had around 200 registered participants (validators). In order to participate as a validator, participants had to fill out paperwork and consent forms, set up appropriate hardware and software, set up a Cosmos wallet, and complete a KYC process. Many participants took GoS very seriously, constantly testing and analyzing points of failure, opportunities, and other validators. Others saw this purely as a competition of smarts and skill, never losing sight of the prize.

The Cosmos protocol measures uptime and voting record of validators and rewards (or punishes) them accordingly. GoS was no different: validators had the potential to increase their stake based on real parameters outlined by Cosmos. These include remaining online at all times, delegation of Atoms, producing blocks, completing challenges and tasks throughout GoS. Participating validators could use any services or hardware to keep their nodes online.

Slashing is a big part of the Cosmos Network and was tested during GoS. Slashing is when staked Atoms are taken away if certain behaviors are found to occur, such as double signing or significant validator downtime. Slashing equates to losing voting power. Specifically the degree to which slashing occurs in each scenario was

tested. For example, the amount of staked Atoms slashed for downtime due to conditions outside of the validator's control would be different than the amount slashed for double signing a block.

GoS was also the first test on jailing when more than one consecutive offense was made. Jailing is the lockout of a validator from the voting process for a predetermined set of time. They miss out on any rewards or inflationary tokens distributed during that time period, discouraging multiple offenses.

While bad behavior was encouraged in order to test the network, there were certain actions that would result in the immediate removal of a participant.

- Violation of Amazon Web Services and Google Cloud Policies.
- Social engineering attacks: using psychological manipulation to trick other users into making security mistakes or giving out information. Activities such as phishing, cloud account credential compromise, malware distribution, and physical security attacks on data centers were considered out of bounds.
- Attacking any testnet other than the GoS Tetnet.
- Causing long-term harm to a validator and their setup
- Exploiting application-level security vulnerabilities in Cosmos + Tendermint code, as these should be reported to Cosmos

## Potential Attack Vectors Tested in GoS

Participants were encouraged to play dirty in order to find bugs in the network and validator system. There were two behaviors in particular that GoS wanted to test: cartel formation and vote withholding. Both of these are big concerns for PoS systems because they can lead to the centralization of voting/network power and forking of the chain.

Cartels need to acquire at least  $\frac{1}{3}$  of the voting power in order to launch a successful attack, after which they censor which blocks and votes get approved. Vote withholding supports this activity by only pre-committing and pre-voting on proposed blocks from cartel members.

# GoS: The Good, The Bad, and The Cartel

---

GoS took place over 4 different networks and lasted two months: GoS, GoS 3, GoS 5, and GoS 6. PoS is incredibly challenging to get right in implementation, and therefore GoS faced some critical and noncritical chain halts and software bugs during its two month run. This was a good thing, as the point of GoS was to make the Cosmos BFT system stronger and more resilient. Validators were given the opportunity to understand and play with validator economics and devise adaptive defense strategies against malicious actors.

Along with testing for and fortifying against critical bugs, GoS was the first time that features to be implemented in Cosmos were tested. One that caused a bug was vesting. Vesting is when tokens can be staked but not distributed for a set amount of time, much like stock options at startups. So while the tokens cannot be spent they can be used to gain transaction fees for block creation and inflationary rewards. Another feature was jailing/tombstone. This happens when a validator engages in a behavior that results in them being locked out of the network for a certain amount of time. Things that can trigger being jailed include double signing blocks or being offline for too long.

## Chain Malfunctions

The development team had coded in detection systems for various things that would halt or disrupt the network. Some of these halts were key to understanding how the network would respond automatically to these mechanisms in the mainnet. These malfunctions also started discussions around which upgrades needed to happen pre-mainnet launch (critical) and which could wait until a later date (noncritical). These malfunctions also provided an opportunity to improve the QA processes, for example creating a more robust detection system for potential slow downs, flaws in the import/export system, etc.

The network halted four times throughout GoS for various reasons:

1. Pre-launch liveness failure, December 11, 2018: An incorrect block was proposed because the validator was running an incorrect version of the Cosmos SDK. This caused all validators to freeze and prevented the next voting round from starting. This issue was fixed by restarting validator and sentry nodes.
2. Fee distribution logic and block size limit, December 19, 2018: The GoS chain was halted at block height 11443 due to a bug in the unbonding and fee distribution logic of the SDK consensus mechanism. Block size limits were also slowing down the network. The block size was increased from 50 kb, a size that was basically taken up by validator signatures and made it difficult to claim and delegate rewards. The network was upgraded to reflect these changes.

3. Double sign and export logic, December 21, 2018: A validator had double signed a block, causing inconsistency in the state of the chain. A slashing period is generated at the genesis of the chain, which failed to happen and caused the chain to halt. As a result GoS was suspended for the holidays in order to fix the bug.
4. Prevote failure, January 4, 2019: Version 3 of GoS halted just before reaching 10,000 blocks when a significant amount of validators failed to pre-vote and a large percentage of validators were also knocked offline.

## Bugs in the Network

In addition to the chain halts, the network experienced bugs that the Cosmos/Tendermint team was quick to address. These potentially critical bugs include:

1. Token Printing Bug: a simple code fix that would have proven to be critical during mainnet launch.
2. Excessive Gas Fees: A flaw in the Cosmos SDK resulted in transactions consuming an excessive amount of gas then ultimately failing. This flaw was limiting the number of transactions to about 7 per block and required a simple upgrade to fix.
3. Simulation code bug: at the end of January a software bug was found in the simulation code that unveiled other errors not previously detected. The code was found to not have been running properly for the past 18 days, the testnet was suspended until the bugs were fixed.
4. Vesting Errors: GoS was the first time that vesting was tested, but the code for vesting caused transactions with fees to fail. A consensus breaking update was required to fix it.

## Cartel Formation and Forking Them Out

At the end of January 2019 Certus One (a GoS participant and winner) published an article on Medium outlining evidence that a cartel had formed. They initially caught this back in December, and brought it up to the Cosmos team and other trusted validators. It was determined that while this behavior would be disastrous during the Cosmos Mainnet Launch, it would be allowed to play out until malicious behavior had actually occurred in GoS. Formation of cartels is a valid concern in PoS systems as they can lead to network censorship, and the standard response has been to fork them out.

The first indication that a cartel had formed was when the Certus One team analyzed registration data from GoS. The IP addresses and number of registered participants provided the first clue: the evening of December

7th saw a significant spike in the number of registrations and many (second highest) registrations were coming from a Digital Ocean (hosting provider) virtual machines. Not only were these registrations happening from the same IP Address and time frame, their GitHub accounts were also less than a week old. 73 accounts to be exact. With this amount of accounts it is entirely possible that, if these were from one group, they could have controlled ~40% of the voting power from the get-go.

When GoS launched these validator nodes all came online at the same time, confirming the suspicion that they were being controlled by the same person. Certus One also attempted to knock the validators (not only done to the cartel) offline with spam transaction attacks and were successful. All of the cartels validators went offline at the same time, which halted the chain as they represented a significant amount of validators. By the time Certus One posted their blog, the cartel had a suspected 53% voting control and were in the lead based on the uptime and staking metrics.

In response to the blog post bitfish labs, the cartel organizer posted a Medium blog outlining their strategy and intent. Bitfish has a history in mining, and they wanted to use GoS to expand their knowledge and expertise to include PoS networks. According to them, they used friends and family to register for GoS and never intended to get even close to the 33% voting power required for BFT consensus failure. They attempted small censorship attacks using a node nicknamed “communist” but chose not to deploy an attack of this size at scale because they did not want to harm the network.

The cartel was forked out of GoS, validators committed to not running GoS 4 (that still had the cartel) and instead opted to run GoS 5 the next day (which cut out the cartel). Despite being forked out, bitfish views their run as successful. Not only were their numbers excellent, other validators had reached out to get help on their own setups. Bitfish has since become an official Cosmos validator that has consistently been in the top 4 since mainnet launch.

# GoS Results

---

GoS was immensely successful. Not only was it the first test of its kind but it proved that the Tendermint BFT could run a large-scale distributed network of 300+ validator nodes while maintaining BFT consensus. Along with testing the technical aspects of the Cosmos Network, the community aspect was also tested. Validators worked together to knock out bad players, helped each other with infrastructure problems, found errors in the code, and ultimately proved how strong the community is.

Two metrics were signaled as valuable by the Tendermint team:

1. Cumulative bonded stake: the amount of staked Atoms accumulated during GoS, the more the better.
2. Number of missed pre-commits: a pre-commit is a vote on a block signed by a validator. If a pre-commit is not included in a block, the validator is considered to have missed the block. This can happen due to a variety of factors, both in and out of the validators control: poor connection, being offline, censorship, etc. The more pre-commits the better.

Analysis conducted by Castenode (a GoS participant) also signals some metrics that express validator performance. One of the most descriptive was average misses per 200 blocks, which shows how reliable a validator was. If their missed blocks were concentrated in a small time frame then there could have been a malfunction on their end that they fixed. Missed blocks over long periods of time showed that the validator was unreliable.

Over the course of the GoS 494,582 blocks were created. Below is a table of the top 10 validators by stake, excluding GoS 5 because there was a bug and some validators did not receive their vesting stake at genesis (see bug 4 above).

Validator	Cumulative Stake
Certus One	2,654,380
bharvest	2,653,679
ATEAM	2,648,313
compass	2,643,182
J	2,635,193
Castlenode	2,630,429

Validator	Cumulative Stake
JASON	2,630,393
loco.block3.community	2,624,590
node	2,612,182
bflabs	2,608,753

Many validators deployed an auto-bonding script that maximized their accumulated stake. These results are not that interesting, as the top 10 all had over 2,000,000 stake at the conclusion and the top 50 did not have a drastic leap in the amount of stake. Where the results start to become interesting is in missed blocks. Below is a table of the top 10 validators by missed blocks, those that missed the least amount of blocks.

Validator	Missed Blocks
Certus One	154
validator.network	567
ATEAM	719
Castlenode	869
J	967
JASON	1,021
kochacolaj	1,429
SF	1,710
DokiaCapital-ION	1,726
layover_run	1,914



With this metric there is a significant difference between the top and even second place validator. Certus One missed only a third of the blocks as the next validator. One of the best strategies to improve this metric was having multiple nodes running for redundancy, both physical and through various cloud providers. Just like in the Cosmos mainnet, validators need to have redundancy and protocols in place for when they get attacked to prevent slashing and jailing. Security is and always should be a validators main concern.

Finally, looking at pre-commits missed per 200 block window Certus One once again takes a significant lead over other validators. On average per 200 blocks they only missed 0.062 precommits, nearly six times less than the second place validator that missed 0.351 precommits per 200 blocks. While GoS rankings is not a guarantee of how well validators will perform in the mainnet, it is an excellent exercise in working out kinks in their operations and security.

## Certus One: Winners of GoS

Certus One objectify demonstrated exemplary technical skills and extensive Cosmos knowledge during GoS by placing first in every category. The metrics participants were ranked off of were:

1. Cumulative Bonded Stake
2. Missed Blocks
3. Average Ranking Per 200 Blocks Window
4. Proposer of Most Missed Blocks
5. Missed Blocks (Without Censorship)
6. Average Ranking Per 200 Blocks Window (Without Censorship)

The team secured their rankings through creating tools, such as an autobonding script in addition to gathering and analyzing game data to adapt their strategy.

Certus One also performed several attacks of which included transaction spam on validator nodes suspected to be controlled by the cartel. When spammed, the nodes all went offline, halting the chain (see chain malfunctions number 4). In addition to their forking out of the cartel, they helped with genesis transactions and found a fix for the state machine bug in GoS 1. According to Zaki Manian, Director of Tendermint Labs, "Certus One's accomplishments are incredibly impressive. They had the most precommits, flawless uptime, led the GoS 6 hard fork. They helped out in validating the genesis transactions and provided a fix for the state machine bug in the first Game of Stakes network. They clearly dominated Game of Stakes and cemented their reputation as one of the best validators."

As a first experiment of its time, Cosmos succeeded in its objective of testing various adversarial conditions as

well as the abilities of node operators. What emerged was a clear distinction of knowledge, skill, and creativity among the pack of participants. The community thought of GoS as a major success and other blockchain networks are looking to launch their own version.

# References

---

1. <https://contribute.cosmos.network/b/xvatkq/view>
2. <https://blog.cosmos.network/the-game-of-stakes-is-open-for-registration-83a404746ee6>
3. <https://medium.com/certus-one/uncovering-a-game-of-stakes-cartel-f895d9591da1>
4. <https://medium.com/bitfishlabs/bitfish-response-to-gos-5-24eaeaea9876>
5. <https://medium.com/tendermint/the-game-of-stakes-its-time-for-some-game-theory-278608c94e6e>
6. <https://medium.com/@castlenode/a-retrospective-on-game-of-stakes-44ab8bd90659>
7. <https://github.com/castlenode/gos-sql/#unofficial-gos-results>
8. <https://medium.com/coinmonks/breaking-down-the-cosmos-game-of-stakes-5cbc538bcd8b>

