ChainArgos.com

Better blockchain intelligence.

OPINION

The CLARITY Act's Decentralization Blind Spot

WHY THE CLARITY ACT'S EMERGENCY-POWER EXEMPTIONS CANNOT WORK AS DRAFTED

AUGUST 21, 2026

Executive Summary

The Digital Asset Market Clarity Act tries to solve a real problem. Developers need room to build software without triggering the full weight of financial services regulation, and platforms need a way to respond quickly when they are attacked. The current draft addresses this through exemptions that allow a security council or similar body to hold emergency powers without voiding a platform's decentralized status.

This opinion piece explains why those exemptions cannot work as drafted. The core difficulty is that no one can reliably separate security fixes from other changes in advance. So any exemption scoped to security incidents is either self-referential with a circular definition or leaves real attacks out of scope. We support the goal of clear rules. Our concern is that the current language will cover almost none of the systems that exist today and will set up years of avoidable disputes. It is also open to fairly straightforward exploitation. The closing sections offer specific recommendations that would make the exemptions workable, and explain what this analysis means for regulators and for investors.

Generally we will assume good faith here by all involved. But when discussing ways in which these kinds of exemptions can be exploited we will need to relax that a little. To illustrate these exploitation concerns, the discussion below includes several examples of seemingly inconsistent statements made by advocates for the Clarity Act and similar legislation. We are not arguing these are bad faith actors. We are merely pointing out that even if you assume these are good faith actors that would never consciously exploit a legal loophole, the positions still leave considerable room for others to engage in such exploitation.

1. Setting The Scene

The Kelp DAO incident exposed a hard design problem at the center of web3 regulation. Many products claim to be decentralized while also giving substantial control to some form of "Security Council[1]" or "Guardian[2]" or other "emergency" or "security" response mechanisms. In the Kelp DAO incident the Arbitrum Security Council and Aave Guardian worked together to seize tokens and transfer them without the private key holder's consent[3]. Both products claim to be decentralized and claim users of their platforms retain full control over their assets. Yet this incident demonstrated that those claims do not always hold true.

The latest Digital Asset Market Clarity Act draft[4] tries to carve out an exemption for these sorts of "emergency" powers. The Act says, in effect, that a platform can claim to be decentralized and be treated as decentralized under US law even when there is an identifiable small group with absolute control over assets on the platform so long as that group's control is limited to handling safety problems.

Set aside any questions you may have about how we are to assess what is a safety problem or what is a good faith response to a safety incident or even whether it is possible to codify the concept of safety vs unilateral control in a purportedly decentralized system. Here we are going to discuss how *as a technical matter* the draft is

[1] https://docs.arbitrum.foundation/concepts/security-council

[2] https://aave.com/docs/ecosystem/governance

[3] The private key holder was purportedly a North Korean hacking group. We are not advocating for their property rights so much as noting that their property was non-consensually seized.

[4] https://www.lummis.senate.gov/wp-content/uploads/Clarity-Act.pdf

unworkable in that one cannot legitimately claim the exemption and still provide anything like a real, effective, and limited-scope safety control of the type the Act envisions.

As a technical matter this is relatively straightforward to show. We do not consider this position controversial or even particularly interesting. What is interesting here is how these exemptions fly in the face of well-understood best practices on the security front developed over decades and discussed widely in the real, non-web3, computer security community.

In short, these exemptions are moot and therefore serve no practical purpose. The sections below explain why, and show that the underlying issues are well understood in the broader computer security community. This connection helps to motivate, and to identify, alternatives and compromises that can work.

For simplicity, we refer to the current draft as the "Act."

2. What The Act Says

Generally speaking the Act exempts only "decentralized" products from traditional regulation. But there are exceptions for "Emergency Measures" where a small group can have total control under a narrow set of circumstances without voiding the decentralization exemptions:

> *Emergency measures.--For the purposes of this section, a pre-defined, temporary, rules-based cybersecurity emergency measure that is exercised by an incident response or security council exclusively in response to a specific and documented cybersecurity incident or imminent threat pursuant to publicly disclosed, on-chain authorization mechanisms, that is strictly limited in scope and duration solely to address that cybersecurity incident or imminent threat, and that is exercised without unilateral control by any single person, shall not alone constitute common control or an agreement to work in concert, if those rules and mechanisms, including the procedures and operational limits governing the emergency measure, are disclosed in publicly available written documentation reasonably available to the applicable Federal agency by a decentralized autonomous organization or similar legal entity sufficiently in advance of any exercise of the emergency measure.*

Elsewhere similar language is present:

> *Emergency measures.*
> *(A) In general.--Pre-defined, temporary rules-based cybersecurity emergency measures exercised by an incident-response or security council exclusively in response to a specific and documented cybersecurity incident or imminent threat and pursuant to publicly disclosed, on-chain authorization mechanisms, strictly limited in scope and duration solely to address such specific and documented cybersecurity incident or imminent threat, and without unilateral control by any single person, shall not, by themselves, constitute common control or an agreement, arrangement, or understanding to act in concert, provided that such rules and authorities, including the procedures and operational limits governing such emergency measures, are disclosed in publicly available written documentation reasonably available to the applicable Federal regulator, by a decentralized governance system or similar legal entity sufficiently in advance of any exercise of such emergency powers.*
> *(B) Prohibition.--The emergency measures described in subparagraph (A) may not be used to implement protocol upgrades, governance decisions, or economic changes that are unrelated to the mitigation of the applicable cybersecurity incident or imminent threat, as described in that subparagraph.*

Both of these passages convey the same idea. A Security Council or other emergency response mechanism must meet two conditions so that it does not void the decentralization of a product:

- It "may not be used to implement protocol upgrades, governance decisions, or economic changes

info@chainargos.com

that are unrelated to the mitigation of the applicable cybersecurity incident or imminent threat."

- And it
- can only have "pre-defined, temporary, rules-based" powers
- can only be exercised "in response to a specific and documented cybersecurity incident or imminent threat and pursuant to publicly disclosed, on-chain authorization mechanisms"
- must be "strictly limited in scope and duration solely to address such specific and documented cybersecurity incident or imminent threat"
- must not involve "unilateral control by any single person"

The conditions are long lists of things that either must be done or must be completely blocked to qualify. What we are going to explain below is that essentially no real, useful or realistic system can meet all these conditions.

The proposed language fares worse elsewhere. Lists of unrealistic conditions are at least specific. Section 604, the so-called "Blockchain Regulatory Certainty Act", provides an exceedingly vague and malleable definition relevant to a set of related issues. This section exempts parties with control over user funds from regulation if they meet this loose definition:

> *Non-controlling developer or provider.*
> *The term ``non-controlling developer or provider'' means a developer or provider of a distributed ledger service that, in the regular course of operations, does not have the legal right or the unilateral and independent ability to control, initiate upon demand, or effectuate transactions involving digital assets to which users are entitled, without the approval, consent, or direction of any third party.*

So a developer or provider can have total independent control over assets so long as it does not exercise that control "in the regular course of operations." Of course, the precise conditions of that term are not defined anywhere. If a system has a freeze or seize switch under operator control but the operator's policy is not to flip that switch, are they non-controlling?

The intent of supporters of this bill is clear. In their view the answer to that question should be "yes." For example, a recent interview[5] discusses carve-outs for "some limited control as it relates to things like cybersecurity" explicitly with respect to all-powerful security council structures. Both participants in that discussion, both lawyers with substantial US government tenure, adopt an intent-based approach endorsing the Clarity Act language with comments like "I think that is something we want to incentivize." To be clear, the "something" is total independent control over funds within a security or safety mechanism operated by humans. Functionally this is simply a bank where the computers process transactions with extremely limited human oversight or intervention.

This is a real issue as myriad projects have repeatedly stated their "policy"[6][7][8] is not to exercise control over user assets. But then we see those same projects use such control to fix problems for their own investors from time to time[9]. Whether the hacking of one's own investor falls inside or outside the regular course of operations is unclear, as is the question of who gets to make that determination. Empirical examples abound of policy decisions that are consistent only to the extent they are consistently self-serving.

Consider the single example footnoted above. Lido has consistently refused to help recover funds for scam or hack victims. But when an investor in Lido lost their governance tokens in a hack the Lido DAO and team took extreme measures to recover them. The process included a video conference with a hired lawyer and carefully planned script which relied on humans taking manual steps to verify each other's identities and then reallocating supposedly immutable assets to fix the investor's problem[10].

[5] https://podcasts.apple.com/us/podcast/has-control-replaced-decentralization-as-defis-legal/id1123922160?i=1000779033023&l=ru start at 51:40
[6] https://research.lido.fi/t/sushi-routeprocessor2-post-exploit-request-for-comment/4383
[7] https://research.lido.fi/t/defiplaza-exploit-request-for-comment/8241/2
[8] https://research.lido.fi/t/a-decision-on-an-outflow-of-40-eth-from-the-lido-dao-treasury-to-aid-in-the-sushi-recovery/4492/9
[9] https://research.lido.fi/t/proposal-to-mint-ldo-tokens-into-secured-wallet/1909
[10] https://web.archive.org/web/20250613044056/https://pastebin.com/raw/VuMTH9ZU

Surely the people on that call, whose names are given in the linked document, all believe that was in the interest of security. Does that structure deserve an exemption from regulation? Do the participants in the call get to make this determination about their own actions?

A second reading of the draft definition is instructive. A service provider would still be exempt from regulation when they had the "legal right" "to control…digital assets to which users are entitled" so long as they only exercise that right outside the "regular course of operations." A bank does not have the legal right to zero out your balance as and when it chooses. And no bank steals depositor funds in the regular course of operations. The entire point is that anyone with the mechanical ability to control funds should be regulated in the same way.

Consistency is important. But many different types of rules can be enforced consistently. The broader context also matters. A wide net is well-aligned with broad consumer- and investor-protection frameworks. A wide net that carves out a single type of database is aligned with nothing but the interests of the sellers of that type of database.

This industry once advocated for "technology-neutral" regulations. This proposal specifically favors a certain technology. It is not clarifying something that was unclear. The industry has long pointed to FinCEN guidance requiring "total independent control" over assets to attract regulatory oversight as a clear, well-defined bright line to manage these issues and protect developers. For example, Coin Center's Director of Policy took this position in recent House testimony about the Act[11]. The draft moves away from that bright line. Under its language, "total independent control" no longer attaches liability so long as it is reserved for emergency use. That is a significant shift, and it deserves open discussion.

That same testimony by Coin Center's Director of Policy decries efforts to hold "developers liable for criminal actions they have no control over." Again that is not what the draft says. The draft blocks efforts to hold developers liable for criminal actions they have the power to block but decide is not sufficiently bad to justify the inconvenience or more substantial effort required to exercise that control.

Similar views are expressed in a letter to the CFTC released by the Hyperliquid Policy Center and Phantom[12]. The letter expresses support for regulatory clarity to:

> *confirm that developing or contributing to onchain protocol software, standing alone without any ongoing control over use of that software, does not trigger registration with the Commission or constitute the solicitation, offering, or acceptance of orders or execution or confirmation of transactions under the Commodity Exchange Act ("CEA")*

This request has no carve-outs for safety or security. Separately, with respect to "Registrants," meaning regulated parties that opt in to CFTC regulation, the letter states:

> *The Commission should make clear that a registrant can, through testing, risk assessment, and other measures, have sufficient controls around the security, scalability, and continuity of the onchain protocols it uses even if they are not in a contractual relationship with a third-party vendor.*

So the position would appear to be that safety and security mechanisms are not a part of DeFi but rather come in alongside registration. This is a consistent position. At the same time, the Hyperliquid Policy Center has publicly supported passage of the Clarity Act, including the emergency measure exemptions discussed above[13]. Read together, these positions ask for an exemption from regulation even where control exists, so long as the holder does not intend to exercise it. That standard would be difficult to defend or administer, for reasons the sections below explain.

Imagine a financial institution with no compliance staff and a policy to process all funds transfer requests unless a court order arrives between the time the transfer request is initiated and finalized. Say logging into the computer which processes transfer requests requires two people to present smart

[11] https://docs.house.gov/meetings/BA/BA21/20260717/119461/HHRG-119-BA21-Wstate-SomensattoJ-20260717.pdf

[12] https://hyperliquidpolicy.org/cl/hpcphantomcftc-fintechrfi-response20260709.pdf

[13] https://x.com/jchervinsky/status/2082530128937021746?lang=bn

cards simultaneously at a terminal in a locked room to which neither of them holds the key[14]. Under the language quoted above is this financial institution entitled to immunity for obviously deficient compliance? Certainly they are not exercising these powers in the regular course of operations. Does it matter if the compliance policy manual states in no uncertain terms that in the regular course of operations they will not block any transfers or freeze any accounts?[15]

This all looks like a clear setup for some future date when a court decides control is control is control and the accused waves their hands and points at, ironically, a claimed lack of clarity to try to avoid punishment. The strategy is transparent. If the answer is that all of this can be litigated later, then the regime is one of regulation by enforcement. If that is the goal it should go on the record now.

All of these strange, vague standards and definitions are about providing carve-outs so service providers and developers can still deal with "security" and "safety" and similar problems without attracting the full weight of financial services regulations. However, one user's bug is another user's feature. These things are also hard to define.

3. All Bugs Are Security Bugs

Let's start with Linux creator Linus Torvalds' attitude on the question of what is a security bug[16]:

> *I expressly told you that security bugs should not be marked as such, because bugs are bugs.*
> ...
> *It's pointless and wrong because it makes people think that other bugs aren't potential security fixes.*

This is echoed by Linux stable kernel maintainer Greg Kroah-Hartman[17]:

> *The primary reason why the kernel does not do any security announcements is that almost any bugfix at the level of an operating system kernel can be a "security issue" given the issues involved*

Elsewhere Kroah-Hartman quotes approvingly from another well known and respected cybersecurity professional, Ben Hawkes[18][19]:

> *It's hard to capture the fact that a bug can be super serious in one type of deployment, somewhat important in another, or no big deal at all -- and that the bug can be all of this at the same time. Vulnerability remediation is hard.*

In a public 2023 talk, Kroah-Hartman is even more explicit[20]:

> *any bug at our level has the potential of being a security issue and that's true people need to realize this any bug could be a security issue not just something we call out so we don't call anything out we just say take all the fixed and move on*

Linux is an operating system and as such has access to all the data on whatever computers it is managing. The same can be said of the code which runs blockchain nodes and the smart contracts that operate platforms on blockchains. In the limit these software tools have total access to everything on their respective platforms. So there is a good chance any mistake can be bootstrapped into a security

[14] Imagine, for example, the MLRO or compliance chief holds the key.
[15] Note this sort of situation is endemic in web3 where a product or protocol has total independent control via the software and then self-publishes a policy manual indicating they will not exercise it.
[16] https://lore.kernel.org/lkml/alpine.LFD.1.10.0807151620450.2867@woody.linux-foundation.org/
[17] http://www.kroah.com/log/blog/2026/01/02/linux-kernel-security-work/
[18] https://en.wikipedia.org/wiki/Ben_Hawkes
[19] https://blog.isosceles.com/what-is-a-good-linux-kernel-bug/
[20] https://www.linuxfoundation.org/webinars/demystifying-the-linux-kernel-security-process

problem. We often find the root cause enabling a hack is only a handful of lines of code.

Now, we do not claim that every programming problem is in fact a safety or security issue. Security researchers Bruce Schneier in 2014[21] and Dan Geer in 2019[22] wrote eloquently about the question of working out what fraction of problems are security problems and what should be done about them. The key truth here is that many programming problems are in fact safety or security problems and yet we have no good automated ways of working out which are and are not safety or security problems. And even if we did, we would not know what to do about them.

The technical explanation for this is the undecidability of non-trivial semantic properties of programs in Turing-complete programming languages. The point is well summed up by yet another well known operating system developer, George Neville-Neil, who asked rhetorically in a 2019 talk[23]:

> *Anyone got a really amazing security test suite? Right, I didn't think so.*

The problem should now be clear. If we cannot tell in advance what is and is not a security problem, how is anyone supposed to set up a process to gate permission to make changes on whether or not the underlying issue is a security risk? The same applies to investor protection and general notions of safety. The entire point of a rapid response mechanism is to be able to respond rapidly. If you are going to gate that on a careful deliberative process it cannot by construction be rapid. If you are going to make it a rapid process it will make mistakes. Who is then responsible? Is whatever group intervenes supposed to be immune from liability because their intent was to help? What about all the other problems they did not intervene to fix? For example, what if their platform was used for sanctions violations? Can they have a policy to only intervene to resolve hacks and not other problems? Should courts grant immunity on the basis of those policies? Should only violations flagged by third parties designated in advance by the government count as violations, as in the framework for Wyoming's stablecoin[24]?

These are difficult problems, and unfortunately they cannot be entirely solved through automation.

4. Computer Science Problems

Above we provided anecdotal evidence in the form of quotes from prominent and experienced security people which strongly suggest the Act's conditions may be difficult to achieve. Now we are going to sketch why they are in fact unachievable in the real world. A more technical treatment of these issues is available in a 2025 journal article co-authored by our CEO[25].

The executive summary is short and best presented as an example. Start with some system about which you have a set of guarantees. As security issues can originate from anywhere in that system, any mechanism that can fix all security issues must be able to modify any component of that system. But once you allow even a single arbitrarily modifiable external function call, you permanently lose the ability to automatically ensure those modifications maintain whatever guarantees you started with. This is a mathematical fact about computer programming under the precise conditions covered in the referenced paper. We cannot completely excise bugs from these systems.

So now you have two choices. First, you can limit the amount of the system you can change. In practice this means there will likely be security issues you cannot remediate. And second, you can opt for more than "pre-defined, temporary, rules-based" remediation powers.

Now there is a potential loophole here, which is to define the "rules" as anything expressible in the programming language used to build the product. This sounds like a solution until you realize the "anything" in that phrase ends up meaning anything at all. As discussed in the journal article, it is possible to build systems with weaker tools such that truly restrictive rules are enforceable. Unfortunately those

[21] https://www.theatlantic.com/technology/archive/2014/05/should-hackers-fix-cybersecurity-holes-or-exploit-them/371197/

[22] https://www.usenix.org/system/files/login/articles/login_summer19_12_geer.pdf

[23] https://www.youtube.com/watch?v=7kShjboN6ek

[24] https://drive.google.com/file/d/1Z21haMRoC6mFfAJnY2FEqSqWdkx-q_p2/view ch 4 sec 3 (iii) limiting the sources of information to initiate a freeze to only parties predetermined by the commission.

[25] https://www.nature.com/articles/s41598-024-84612-9

systems are not sufficiently powerful to build many types of useful real-world services. This trade-off is fundamental to computing. It is, in a very real sense, the reason consumer devices still crash from time to time. Programming tools powerful enough to build things we want are often too malleable, too flexible, for the automated formal verification methods that could cut out all the bugs.

But this is all a bit theoretical. Let us look at the 4 conditions listed above that any Security Council or emergency response mechanism must meet to avoid running afoul of the decentralization exemptions.

Can only have "pre-defined, temporary, rules-based" powers Unless we stick with the circular and vacuous limitation of "can only be used to fix security or safety problems" this one is a problem. The only options are to accept some bugs will not be fixable, some desirable products will not be buildable, or the system will not qualify as decentralized.

Can only be exercised "in response to a specific and documented cybersecurity incident or imminent threat and pursuant to publicly disclosed, on-chain authorization mechanisms" It is of course possible to introduce a gating mechanism that only allows changes after some other process declares that a security incident occurred. But this cannot be automated for the reasons discussed above. So in practice designers must choose between missing some incidents and falling back on a circular "we only allow changes in response to things we declare to be security incidents without regard to any firm pre-defined rules." This kind of approach also entails an obvious balance between speed and correctness. A rapid response mechanism is going to misclassify problems more often than a slower mechanism. What are we to do about those mistakes?

Must be "strictly limited in scope and duration solely to address such specific and documented cybersecurity incident or imminent threat" Again the only possibilities here are circular logic and gaping holes in defenses.

Must not involve "unilateral control by any single person" This is the condition most easily satisfied. It is straightforward to build an on-chain governance mechanism such that multiple parties, maybe something like a DAO voting scheme where anyone can buy tokens and participate, must approve all changes. The rules do address this issue and achieve something. Notably this is the only point where the technology is simple and exists today and where the real open issues sit entirely in the legal domain. As several courts have found these sorts of DAOs to be general partnerships with joint and several liability they are a questionable foundation atop which to build financial services. But here, at least, we can point to concrete mechanisms that are capable of performing the required tasks.

This also puts us in an excellent position to look at the "regular course of operations" problem. Any given system having bugs, and those bugs getting exploited, are both entirely foreseeable problems. The exact details of the bug or bugs are hard to know in advance. But it is not like cybersecurity problems are some kind of black swan that nobody saw coming. A "bug" is a situation where the software behaves how the code was written but not how the people writing that code wanted it to behave. Exempting the operator of a financial services product from financial services regulations so long as their unilateral control over user funds is only exercised when the software behaves differently from the way that operator wants the software to behave is untenable. We know software is full of bugs. Managing this is the entire point. Treating bugs as rare or exceptional is a fundamental error. Modern businesses are built to withstand and survive bugs.

And at first blush it might look like imposing some kind of backward-looking liability could fix this. Maybe a regulator would be empowered to review exercises of the administrative powers and decide that the provider made a mistake. What then? If the result is a nominal penalty then there is no incentive to do any of this in good faith. And if the result is a large, painful, punishment of some kind? Then no provider is going to act to protect their users lest punishment be visited upon them. Or maybe we end up with services governed by security councils with anonymous members. So literally no liability because we do not even know who to blame. Once you think through how these rules work it is clear they are not fit for purpose. Any backward-looking regime of this kind is regulation by enforcement, a point the recommendations below address directly.

5. Evidentiary Problems

The July 22, 2026 draft[26] includes a safe harbor for certain designated private entities when they share information they think is related to illicit finance. Specifically it states:

> *A designated private sector entity that transmits, receives, or shares information for the purposes of identifying and reporting activities that may constitute illicit finance violations, or threats and emerging risks relating to illicit finance violations, shall not be liable to any person for such disclosure or for any failure to provide notice of such disclosure to the person who is the subject of such disclosure or any other person identified in such disclosure.*

This is far too broad a grant of immunity. We have concrete examples of evidence fabrication and falsification in courts, and examples of egregious errors and incompetence, from exactly the sorts of parties that are likely to end up on that designation list. No broad and strong presumption of good faith is reasonable here. Any grant of immunity must be narrow and subject to challenge.

Even if it were possible to build perfectly reliable automated systems to solve these problems, and we know it is not, the web3 industry has a long history of presenting tools as flawless which it turns out the developers knew at that time had significant holes. Similarly, there are myriad cases of proposed "solutions" which fall down on first contact with a real attacker.

[26] https://www.lummis.senate.gov/wp-content/uploads/Clarity-Act.pdf

This designated entity safe harbor concerns entities with status akin to licensure providing non-public information via non-public channels to law enforcement with the expectation they will action the information. Strong grants of immunity are not justifiable given the history of actors in this sector and given the known impossibility of performing the relevant tasks flawlessly. We are not suggesting that governments can only act when given perfect information. We are simply suggesting it is impossible to properly manage imperfect information when you grant immunity, broadly and up front, to those providing the information. The act can add some kind of rebuttable presumption or other conditions on the immunity. As it stands the phrase "shall not be liable" is just too broad.

6. Realities

The reality is that it is impossible to build many useful, desirable web3 services in ways that allow for addressing security issues quickly and comply with the decentralization conditions in the Act. This does not mean every service is impossible to build. And certainly it is theoretically possible to build immutable services which are both secure and qualify as decentralized.

The problem is that it is exceedingly difficult to build complex services which are secure out of the box. So in practice the only kinds of Act-compliant decentralized services anyone can really build are going to be extremely simple or extremely risky. There are simple immutable systems which have proven to be secure and which look in fact to be secure. But these tend to be small units of functionality and make up an even smaller fraction of the token values of the kinds of services that are likely to want the decentralization or non-controlling exemptions. To see this, consult L2BEAT's layer-2 dashboard[27]. Only a vanishingly small fraction of the Ethereum Layer 2 blockchains even claim to have sufficiently restricted Security Councils.

Yes there are a few larger DeFi protocols that would qualify. Uniswap almost surely qualifies. But Aave certainly does not. Hyperliquid does not. Lido does not. A review of the largest protocols listed on DefiLlama again shows that only a small fraction are eligible[28]. All the guardians and security councils we see in the real world have unfettered power to change their systems. Most are able to effect changes immediately with no opportunity for users to opt out of potentially disruptive modifications. This is very much a security-first mindset built atop an assumption the security council itself is not a risk. Financial services regulations, wisely, do not assume management is always trustworthy, mainly because we have centuries of evidence that management is often the biggest security problem an institution faces.

The Act as written today provides exemptions for some products that actually exist and some products that could exist. But it covers at most a small fraction of the products claiming to be decentralized that exist in the real world today. What we need here is an open discussion of this problem so everyone involved is on the record about what will and will not qualify under these rules. If anyone wants to accept vague, circular, promises like "we will only use our unfettered power to fix XYZ designated security problems" they should be prepared to defend that position publicly. They should also explain what they want done if the group making such a commitment then violates that very commitment. Otherwise this entire exercise is just setting up future regulatory battles over what the language means. And the entire point of having a "Clarity" act was to fix that problem today.

[27] https://l2beat.com/scaling/summary
[28] https://defillama.com/

info@chainargos.com

7. Recommendations

The problems described above are structural, but they are not fatal to the Act's goals. The analysis points to six specific changes.

First, retain a bright-line control test. If any party holds the mechanical ability to move, freeze, or reassign user assets, that party should be treated as a controlling party regardless of its stated policy. This is the standard the industry itself long defended, and it is the only test that does not depend on reading intentions.

Second, exempt genuinely immutable systems cleanly. Systems with no administrative control mechanism at all pose none of the problems described above. The law should say so plainly, which would give builders a clear and achievable safe design target.

Third, keep and strengthen the requirement that no single person hold unilateral control. This is the one condition in the current draft that is both technically achievable and verifiable on-chain today. It should be paired with a requirement that council members be identified, so that accountability is possible.

Fourth, replace advance scoping with review after the fact. Since no one can define in advance which changes are security fixes, the law should not pretend otherwise. A workable structure would grant time-limited emergency powers, require public disclosure of any intervention within a fixed period, and apply a rebuttable presumption of good faith rather than automatic immunity or automatic liability.

Fifth, narrow the information-sharing safe harbor the same way. A rebuttable presumption subject to challenge protects honest actors without granting blanket immunity to a sector with a documented history of error and fabrication.

Sixth, put qualification on the record before disputes arise. Regulators should publish which existing designs qualify under the final language. Every major protocol operator should be able to know, in advance, which side of the line it stands on. That is what a clarity act is for. We are not advocating for a civil law system but rather for interpretive guidance codified within the law. The law should also acknowledge explicitly that the backward-looking portions of the regime amount to regulation by enforcement, and state plainly that this is intended.

8. What This Means for Regulators and Investors

For regulators and legislative staff, the practical lesson is that no amount of drafting skill can rescue an exemption scoped to security incidents, because the boundary it depends on cannot be defined. The choice is between a bright-line control test with review after the fact, or a discretionary standard that will be litigated for years. The recommendations above describe the first path. The second path is, roughly, the approach the United States took from the mid-2010s until early 2025, resolving these questions case by case through enforcement. The entire point of a clarity act is to avoid taking that road again.

For investors in tokenized assets, the practical lesson is to ask who holds the mechanical ability to move or freeze an asset, not what the project's policy says. Public dashboards make this checkable today. A decentralization label without that check is a marketing claim, and the record shows that emergency powers tend to be used when the operator's own interests are at stake. Investors can begin to approach this problem by asking a simple question. If something goes wrong with this tokenized asset, who, if anyone, can we sue? When the answer is "nobody" you likely have a genuinely decentralized product and comparatively little regulatory risk. For everything else, there is a strong positive correlation between the regulatory risk of each identified party and the strength of your potential case against that party.

info@chainargos.com

Who are we?

ChainArgos is the blockchain intelligence firm best known for uncovering crypto-asset exchange Binance's $1.4bn BUSD stablecoin undercollateralization, forcing the New York Department of Financial Services to take action.

We provide unparalleled blockchain intelligence by focusing on the financial drivers of transactions, facilitate investigations and analysis centered on the economic value of transfers, and provide insight into the motivation behind specific flows.

ChainArgos is recognized globally as a leader in blockchain intelligence.

We've tracked illicit flows funding terrorism and sanctions evasion, analyzed transaction patterns connecting global scams, and uncovered crypto-asset trading opportunities before the market.

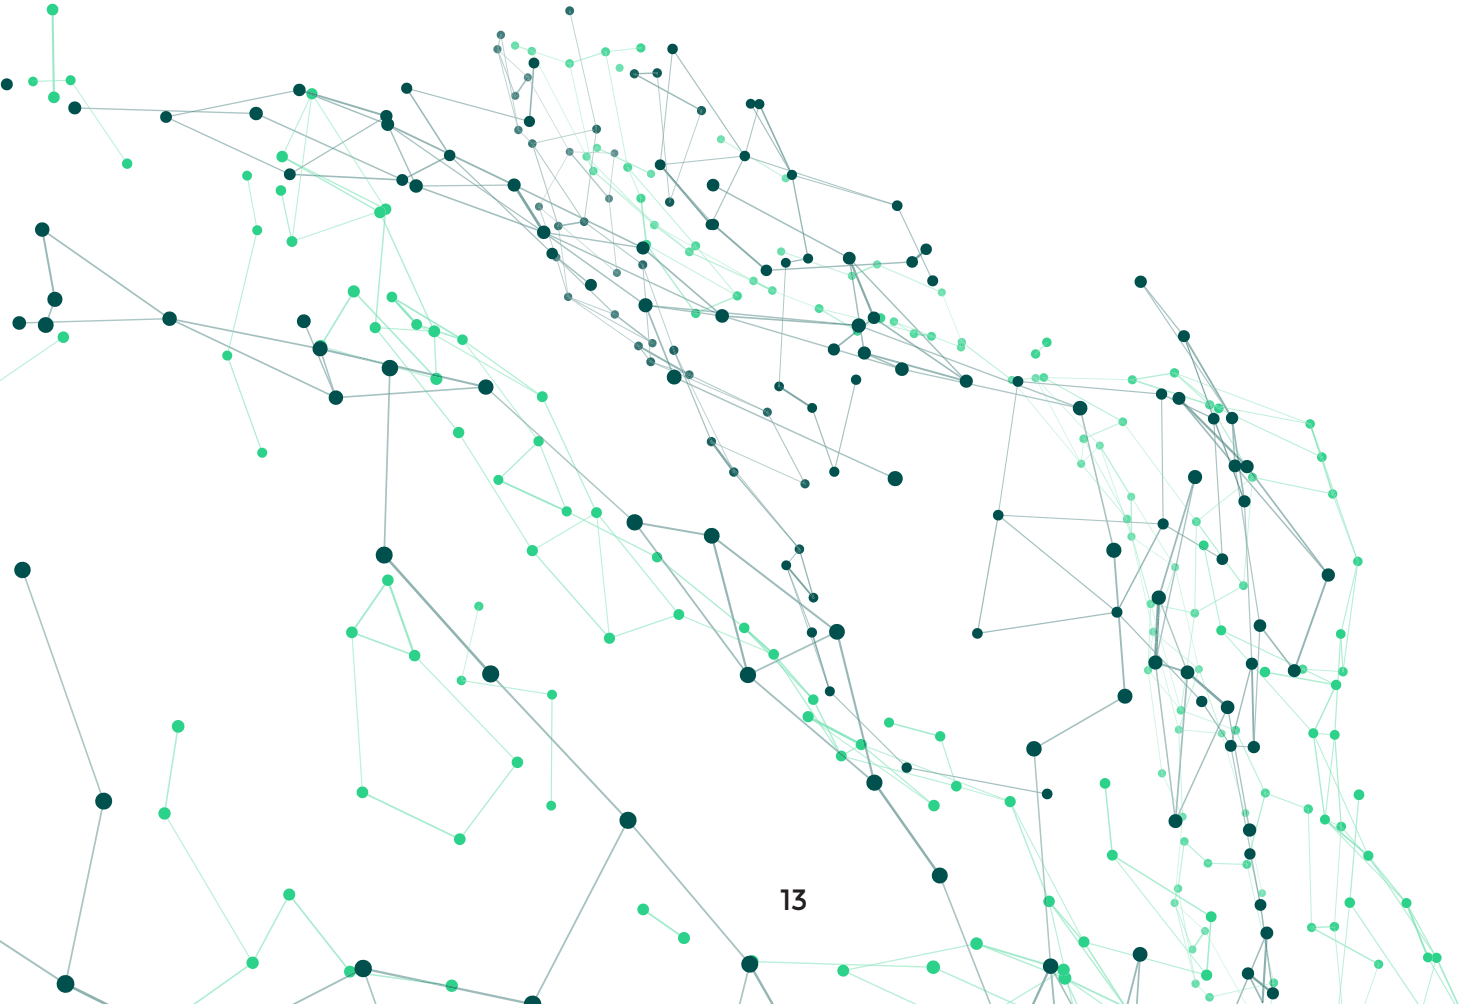

Where else have you seen us?

ChainArgos works with the United Nations, governments, central banks, financial institutions, hedge funds, proprietary trading firms, regulators, law enforcement and intelligence agencies, research institutes, universities, and crypto-asset service providers globally.

We're trusted by top news outlets including the Wall Street Journal, Bloomberg, Forbes, Fortune, Thomson Reuters, and the South China Morning Post, for unimpeachable blockchain intelligence.

Here's just a selection of our blockchain intelligence that created news:

Bloomberg

Binance Acknowledges Past Flaws in Maintaining Stablecoin Backing

- Blockchain analyst Reiter had flagged gaps in Binance-peg BUSD
- Binance says earlier 'operational delays' have now been fixed

Forbes

Did Digital Currency Group Profit From $60 million In North Korea Crypto Money Laundering?

THE WALL STREET JOURNAL.

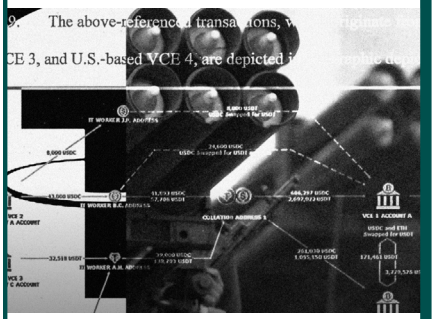

From Hamas to North Korean Nukes, Cryptocurrency Tether Keeps Showing Up

Tether has allegedly been used by Hamas, drug dealers, North Korea and sanctioned Russians

THE WALL STREET JOURNAL.

The Shadow Dollar That's Fueling the Financial Underworld

Cryptocurrency Tether enables a parallel economy that operates beyond the reach of U.S. law enforcement

Bloomberg

Stablecoin Operator Moves $1 Billion in Reserves to Bahamas

- Move reflects worsening US banking conditions for crypto firms
- TrueUSD's circulation has more than doubled in the last month

South China Morning Post

How crypto investigators uncover scammers' blockchain billions, scale of money laundering in Asia

info@chainargos.com

Who uses blockchain intelligence?

Finance and Banking

Assess the risks and opportunities in crypto-assets, stablecoins, and decentralized finance. Develop innovative products, explore tokenization opportunities, and generate new revenue streams.

Compliance

Fight money laundering, expand know-your-customer tools, and combat the financing of terrorism while expanding your customer base. Manage risk from customer crypto-assets and confidently verify sources of crypto-asset wealth.

Law Enforcement

Terrorists and criminals are using blockchain technology to avoid the banking system, launder money, and fund operations. Blockchain wallet analysis and transaction tracing fights crime, prosecutes criminals, and tracks illicit fund flows.

Regulators and Policymakers

Develop and implement effective crypto-asset and stablecoin supervisory, licensing tax, compliance, and regulatory frameworks to foster innovation, while managing threats to national security and the financial system.

How are we different?

We deliver actionable blockchain intelligence.

Say “no” to pseudo-science and “yes” to blockchain intelligence you can count on for commerce, compliance, and crime-fighting.

ChainArgos is built by finance, legal, and technology professionals to deliver actionable blockchain intelligence focused on financially-relevant analysis.

Whether you’re looking to on-board a customer, determine source of wealth, or ensure your evidence isn’t rejected on appeal, our blockchain intelligence is based on established principles of statistics, math, and forensic science.

Extreme Versatility

Create compliance and commercially-driven analysis in a single place and arrive at better business decisions faster.

No-Code Customization

Build any query or analysis without programming skills or coding.

Financially-Relevant

Standard financial measures combined with blockchain intelligence for actionable insight.

Data Integrity

ChainArgos runs its own blockchain nodes, and we never enrich our data with yours, so you can be sure of data integrity.

API Ready

Robust and resilient APIs with 99.99% uptime. Minimal code required for easy integration.

Automated Alerts

Schedule automated alerts and reports via Email, Webhook, Amazon S3 and SFTP so you’re always in the know when something happens.

info@chainargos.com

How do we do it?

Blockchain intelligence is a relatively new industry, and it's not uncommon to hear of methods which have little basis in finance, let alone forensic science.

Let's look at one example to understand the limitations of blockchain tracing.

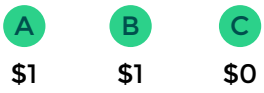

Fig. 1

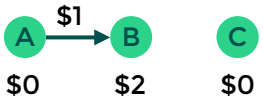

Fig. 2

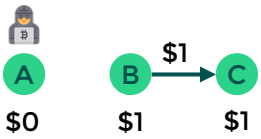

Fig. 3

In Fig. 1, A and B start with $1, while C starts with $0. In Fig. 2, A transfers their $1 to B who now has $2. Finally, in Fig. 3, B transfers $1 to C, who now has $1.

If it turns out A is an illicit actor, with what degree of confidence can we say that C has received $1 from illicit sources? 50-50?

Would you accept a "risk score" of 50%?

Follow the money.

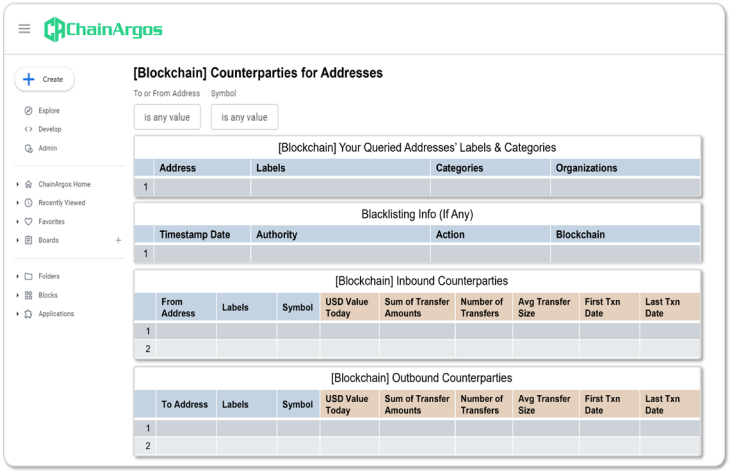

Instead of passing off "risk scores" as "risk management" ChainArgos helps you follow the money.

Most blockchain transactions don't derive from a single source, and believing they do is what leads to poor outcomes.

Make better decisions by focusing on what matters - where the money went, where it came from, and where does it look like it's headed to?

How much does one address deal with another? What's the average transaction size? What's the frequency? What's the crypto-asset or stablecoin of choice? What's the transaction behavior? When did the transaction size change?

And so much more.

info@chainargos.com

Better attribution.

Don't risk critical legal, trading, and compliance decisions to questionable or subjective attribution methods. Trust math and science.

ChainArgos is the only blockchain intelligence firm that delivers programmatic address labels and wallet tags that are unassailable whether you're making business decisions or preparing to sue someone.

Blockchain addresses are automatically ranked and labeled based on a variety of factors including:

- **Transaction Count**: the number of transactions by an address. Sending $100,000 in one transaction may have very different implications from sending 10 transactions of $10,000 each. Either way, you'll know the difference.
- **Lifetime Sent/Received**: lists the biggest sender and/or receiver of any given crypto-asset or stablecoin currently. Markets are extremely dynamic. The biggest movers today may not be the same tomorrow.
- **Max. Historical / Current Balances**: helps you decide whether an address is participating in affiliated crypto-assets and/or stablecoins based on their maximum historical balance and who's stocking the highest current balances.
- **Recipient Number**: gives you a sense of whether they were an early adopter, or even possibly an insider of a crypto-asset or stablecoin. Recipients are ranked according to the date and time they received a crypto-asset or stablecoin.

Say "no" to dodgy wallet tagging and "yes" to attribution you can trust.

info@chainargos.com

Legal Disclaimers.

THE INFORMATION CONTAINED IN THESE MATERIALS IS FOR INFORMATION PURPOSES ONLY AND NOT INTENDED TO BE RELIED UPON.

The information contained herein is information regarding research and analysis performed by ChainArgos Pte. Ltd., a company incorporated with limited liability under the laws of the Republic of Singapore with registration number 202303560W ("the Company"). The information herein has not been independently verified or audited and is subject to change, and neither the Company or any other person, is under any duty to update or inform you of any changes to such information. No reliance may be placed for any purposes whatsoever on the information contained in this communication or its completeness. No representation or warranty, express or implied, is given by, or on behalf of the Company or any of their members, directors, officers, advisers, agents or employees or any other person as to the accuracy or completeness of the information or opinions contained in this communication and, to the fullest extent permitted by law, no liability whatsoever is accepted by the Company or any of their members, directors, officers, advisers, agents or employees nor any other person for any loss howsoever arising, directly or indirectly, from any use of such information or opinions or otherwise arising in connection therewith. In particular, no representation or warranty is given as to the reasonableness of, and no reliance should be placed on, any forecasts or proposals contained in this communication and nothing in this communication is or should be relied on as a promise or representation as to the future or any outcome in the future.

This document may contain opinions, which reflect current views with respect to, among other things, the information available when the document was prepared. Readers can identify these statements by the use of words such as "believes", "expects", "potential", "continues", "may", "will", "should", "could", "approximately", "assumed", "anticipates", or the negative version of those words or other comparable words. Any statements contained in this document are based, in part, upon historical data, estimates and expectations. The inclusion of any opinion should not be regarded as a representation by the Company or any other person. Such opinion statements are subject to various risks, uncertainties and assumptions and if one or more of these or other risks or uncertainties materialize, or if the underlying assumptions of the Company prove to be incorrect, projections, analysis, and forecasts may vary materially from those indicated in these statements. Accordingly, you should not place undue reliance on any opinion statements included in this document.

By accepting this communication you represent, warrant and undertake that you have read and agree to comply with the contents of this notice.

info@chainargos.com