# Bulk Domain Registrations

*ccNSO & GAC Joint Discussion Summary & Insights*

## What's the Big Picture?

Starting point for the discussion was the INFERMAL study and the correlation identified in this study between unrestricted APIs and DNS abuse. The conversation at ICANN84 centered on a surprisingly nuanced question: when someone registers hundreds or thousands of domain names at once, is that a problem? The short answer turned out to be: it depends. Think of an API as a fast lane at the registry—it lets people register many domains quickly through automated systems. But just like a fast lane isn't inherently good or bad, bulk registration isn't automatically suspicious.

What emerged was a tension between trying to identify who's doing the registering versus trying to understand why they're doing it. Both government and technical communities realized they're looking at the same phenomenon from different angles, and neither has all the answers.

## Key Themes That Emerged

### The Identity vs. Intent Puzzle

Here's where things get interesting. Just because you can identify someone doesn't mean you know what they're planning to do. A brand owner might legitimately register 15,000 domain variations to protect their trademark. Meanwhile, a domain investor in a country might snap up every domain name containing the word "flower" to resell to florists later—perfectly legal but looks identical to bulk registration from the outside.

The challenge is that knowing someone's identity doesn't automatically reveal their intentions. Switzerland's .CH registry has an interesting approach: when something looks suspicious, they pause the registration and ask for identification before proceeding. It's a simple friction point that gives them time to assess the situation.

### Burst vs. Bulk: Patterns Matter

The discussion revealed an important distinction between "burst" and "bulk" registrations. A burst might be 500 domains registered in 10 minutes. Bulk might be 10,000 domains registered steadily over a week. Both involve large numbers, but the patterns tell different stories about intent and risk.

Instead of focusing solely on identity verification, several participants suggested watching for unusual behavioral patterns. If a registrar suddenly shows a spike in daily registrations, that's a red flag worth investigating—regardless of who's behind it.

### ccTLDs vs. gTLDs: Different Worlds

Country code top-level domains (ccTLDs like .de or .ch) operate completely differently from generic TLDs (like .com). Each ccTLD sets its own rules, has its own accreditation process for registrars, and serves its local community. There's no universal playbook.

This means some ccTLDs might cap the number of registrations per day, while others have no restrictions at all. They often have closer relationships with their domain holders and can spot unusual activity more easily. The smaller scale and local focus gives them flexibility that ICANN-accredited registrars operating globally don't have.

## Real-World Examples

**Legitimate bulk registration scenarios that came up:**

- Brand protection: A company registers thousands of variations of their brand name to prevent cybersquatting and typosquatting
- Domain investors: Portfolio holders who collect themed domain names (like all variations with "flower") to resell as packages to relevant businesses
- Resellers: Intermediaries who register domains on behalf of multiple clients in a single batch transaction

The discussion made clear that having 10,000 domains isn't unusual or automatically problematic. The sheer existence of large portfolios is normal in the industry.

## Emerging Issues & Open Questions

### The Role of Friction

From the GAC perspective, the goal is to create just enough friction to make malicious actors think twice—without making it impossible for legitimate users to operate efficiently. It's a delicate balance. Too much friction frustrates legitimate businesses. Too little makes abuse trivially easy.

The question remains: what kind of frictions are effective? Requiring documentation? Verification delays? Graduated access based on reputation? Nobody has definitively answered these questions to date.

### Who's Responsible?

The conversation kept circling back to a fundamental question: is bulk registration primarily a registrar problem or a registry problem? Registrars interact directly with customers and process the registrations. Registries see the aggregate data and patterns. Both have different tools, incentives, and responsibilities

One theme emerged: registries earn more revenue from bulk registrations, which creates a potential conflict of interest. How do you balance business incentives with security concerns?

### The Payment Trail

When someone runs a phishing campaign and claims it was a mistake, who gets the money back? This question highlighted how payment verification might be a effective identification mechanism. Registrars typically pay registries in bulk, which obscures the individual registrant's payment trail.

Following the money might be more effective than following the identity, but it also raises privacy concerns and practical implementation challenges.

### Threshold Questions

Nobody could agree on where to draw the line. What number makes a registration "bulk"? 100 domains? 1,000? 10,000? The consensus was that hard numbers aren't as useful as pattern recognition. A ccTLD might not have a formal definition of bulk registration but can still recognize unusual activity when they see it.

## What's Next?

The discussion didn't produce concrete recommendations—that wasn't the goal. Instead, it created shared understanding between the technical community (ccNSO) and government representatives (GAC) about the complexity of bulk registrations.

Several participants emphasized that APIs themselves aren't the problem—they're essential infrastructure that enables efficient domain registration. The challenge is ensuring that access to these powerful tools doesn't enable abuse at scale.

Moving forward, the focus will likely shift from simple identity verification to behavioral analysis and pattern recognition. The key insight: the need to understand what registrants are doing, not just who they are. Whether that's achievable without compromising privacy and efficiency remains an open question.

## Key Takeaways

- Bulk registration isn't inherently problematic—context and intent matter more than volume
- ccTLDs have more flexibility than gTLDs to experiment with different approaches and local solutions
- Behavioral patterns (burst vs. steady bulk, unusual spikes) may be more revealing than identity verification alone
- The industry needs to balance security with efficiency—too much friction harms legitimate business
- Payment trails and verification might be more effective than traditional identity checks
- No consensus exists on threshold definitions—pattern recognition may be more practical than hard numbers
- APIs are essential tools, not problems—the focus should be on who has access and how it's monitored