Joint ccNSO-GAC Dublin Meeting Summary

ccNSO & GAC Joint Sessions | Key Discussions and Takeaways

Executive Overview: Five Key Discussion Areas

Station	Main Focus	Key Innovation	Notable Results
AI Detection *(.nl registry)*	Machine learning to detect patterns across large datasets with human oversight	RegCheck system for registration monitoring; continuous AI retraining; Privacy Board governance	98% reduction in fake webshops (12,000 to 241) over 5 years
Portfolio Abuse *(.US registry)*	Tracking abuse patterns across multiple domains owned by same actors	No privacy/proxy services; verified contact data; AI to distinguish compromised vs. malicious sites	5-day investigation timeline; 72-hour registrant response window; no court orders required
National Framework *(Australia)*	Multi-stakeholder responsibility across telcos, banks, platforms, and registries	National Anti-Scam Centre; fusion cells bringing sectors together; presence requirements for .au	30-day response timeline; suspend first, delete later; civil penalties framework
Bulk Registration	Understanding intent behind large-scale domain registrations	Pattern recognition over hard thresholds; behavioral analysis distinguishing burst vs. steady bulk	No consensus on definitions; focus shifting from identity to behavior and payment trails
Trusted Notifiers *(.uk registry)*	Partnering with expert organizations for specialized abuse detection	Three-tier trust system: immediate (child abuse), priority (MOU partners), law enforcement channel	2,000-3,000 takedowns/year; no lawsuits; low maintenance; cost-effective

Detailed Discussion Summaries

Station AI-Powered Abuse Detection

The Netherlands' .nl registry (SIDN) demonstrated how artificial intelligence can dramatically reduce domain abuse when properly governed. Their success with fake webshops—reducing incidents by 98% from 12,000 in 2018 to just 241 by 2023—showcases the potential of machine learning combined with human oversight.

Critical Success Factors:

- Human-in-the-loop principle ensuring AI provides recommendations, not automatic decisions
- Privacy Board governance and EU AI Act compliance preparation
- Continuous model retraining as abuse patterns evolve
- Joint development with .be registry (RegCheck system) enabling resource sharing

- Open-source foundation with multiple detection systems: RegCheck, fake webshop detection, online impersonation

The discussion emphasized that AI serves both defensive and offensive purposes—while registries use it to detect abuse, malicious actors employ it to create sophisticated attack networks. This arms race requires continuous investment in expertise and infrastructure.

Station Portfolio-Based Abuse Detection

The .US registry shared their "nexus" approach—connecting dots between domains and registrants to spot patterns of abuse. Rather than investigating individual domains in isolation, they track entire portfolios of domains owned by the same actors.

Unique Advantages:

- Prohibition of privacy/proxy services provides verified contact information for 2.5 million registrants
- AI distinguishes between compromised websites (legitimate sites that were hacked) and deliberately malicious operations
- Fast action timeline: 5-day investigations, 24-hour resolution for simple cases
- Authority to act without court orders as the authoritative .US registration source under NTIA
- Collaborative intelligence sharing within security professional community

This approach raises important questions for other TLDs that allow privacy services: Can portfolio abuse be effectively detected without knowing domain ownership? The balance between speed and due process—72-hour response windows—will inform ICANN's policy development for associated domain checks.

Station National Anti-Scam Framework

Australia's National Anti-Scam Centre exemplifies the "everyone does their bit" philosophy—recognizing that scams require coordination across banks, platforms, telcos, and domain registries. The 2022 scam spike triggered major government action, new legislation, and industry standards.

Framework Elements:

- Multi-stakeholder fusion cells bringing together experts from different sectors
- Presence requirements for .au domains significantly reduce malicious registrations
- Government-funded consumer reporting mechanism with direct registry notification
- Civil penalties framework (financial obligations without liability for actual losses)
- AUDA's proactive approach: calling hacked businesses to help them remediate rather than immediately suspending

Open challenges include information sharing versus privacy protection and defining "bulk registration." The discussion revealed tension: a brand registering 100 domains is legitimate, while a fraudster doing the same is problematic—but distinguishing them upfront remains difficult.

Bulk Registration Discussion

This conversation centered on whether large-scale domain registrations are inherently problematic. The answer: context and intent matter more than volume. APIs enable legitimate businesses to operate efficiently while also potentially enabling abuse at scale.

Key Insights:

- Identity verification doesn't reveal intent—knowing who registered domains doesn't explain why
- Pattern recognition may be more effective than hard thresholds ("burst" vs. "bulk" patterns)
- ccTLDs have more flexibility than gTLDs to experiment with local approaches
- Payment trails might be more revealing than traditional identity checks
- Friction must balance security with efficiency—too much harms legitimate business

The discussion produced shared understanding rather than concrete recommendations, with participants acknowledging that APIs themselves aren't the problem—they're essential infrastructure. The challenge is ensuring access doesn't enable abuse at scale.

Station Trusted Notifier Systems

The UK's .uk registry (Nominet) demonstrated that trusted notifier systems work—and they're surprisingly straightforward. The key insight: don't try to become experts in everything. Partner with organizations whose full-time job is identifying specific abuse types.

Three-Tier System:

- **Tier 1 - Internet Watch Foundation:** Immediate suspension for child abuse content (highest trust, highest stakes)
- **Tier 2 - MOU Partners:** Organizations like TWNIC and DotAsia get priority investigation (reports jump the queue)
- **Tier 3 - Criminal Practices Policy:** 14 UK law enforcement agencies; 48-hour registrant notice before suspension

Nominet handles 2,000-3,000 takedowns annually with no lawsuits to date, suggesting the combination of trusted sources, clear terms, and appeal processes provides adequate protection. The system is described as cost-effective and "low maintenance" compared to building in-house expertise across all abuse categories.

Cross-Cutting Themes

- **Speed vs. Due Process:** All approaches balance rapid response (immediate for child abuse, 24-72 hours for other abuse) with registrant rights and appeal mechanisms
- **Data Quality Foundation:** Accurate WHOIS data and verified identities enable portfolio detection—key question for TLDs with privacy services
- **Human Oversight:** Technology enhances but doesn't replace human judgment—AI provides recommendations, experts make final decisions
- **Collaborative Intelligence:** No registry fights abuse alone—successful approaches leverage partnerships, joint development, and community threat sharing
- **ccTLD Flexibility:** Country-code TLDs have more latitude to experiment with local approaches than gTLDs bound by ICANN contracts
- **Pattern Recognition:** Behavioral analysis and anomaly detection often more effective than rigid thresholds or definitions

* * *

These discussions inform ongoing discussions on associated domain checks, abuse reporting, and validation requirements.