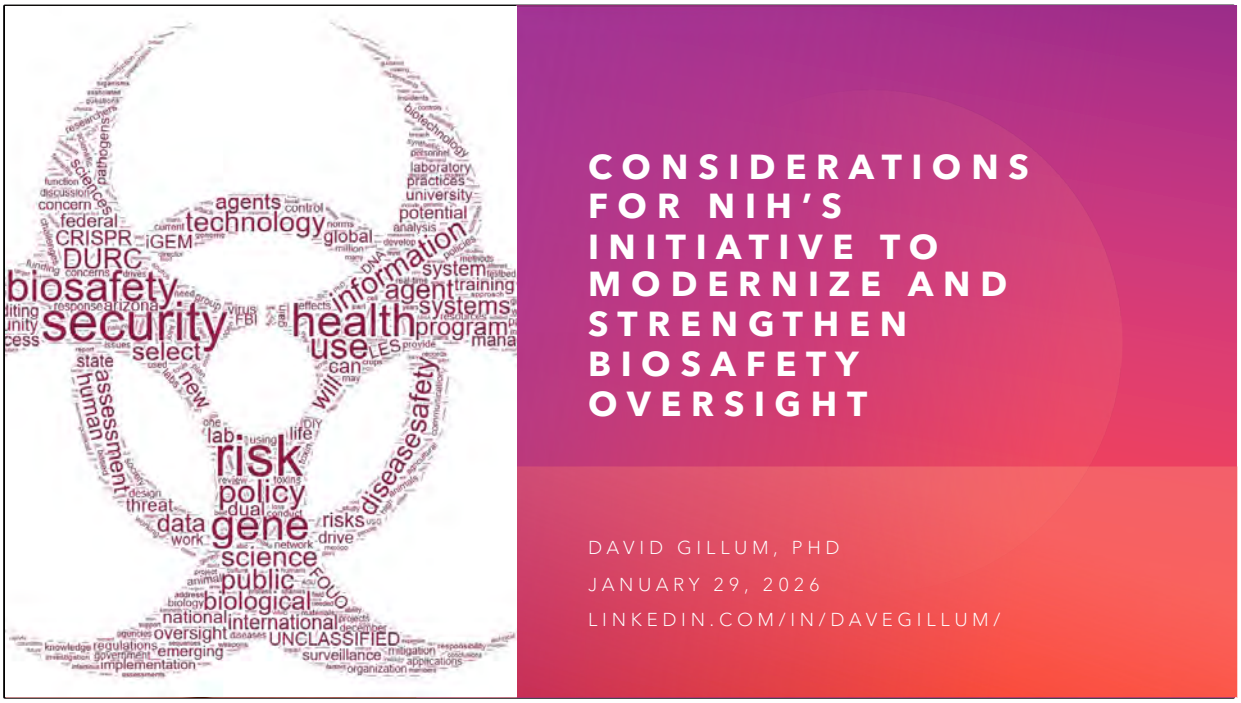


Considerations for NIH's Initiative to Modernize and Strengthen Biosafety Oversight

David Gillum, PhD
January 29, 2026
linkedin.com/in/davegillum/


Risk Groups

Why we need to modernize Risk Groups

## Human-Centric Risk Group (RG)

- Risk Groups are based on human disease severity and treatment
- Modern research spans animals, plants, insects, and ecosystems
- Non-human and environmental risks are not systematically captured
- Research systems and risk profiles change faster than oversight frameworks

| Risk Group | Basis for the Classification of Biohazardous Agents |
|---|---|
| Risk Group 1 (RG1) | Agents that are not associated with disease in healthy adult humans |
| Risk Group 2 (RG2) | Agents that are associated with human disease which is rarely serious and for which preventive or therapeutic interventions are often available |
| Risk Group 3 (RG3) | Agents that are associated with serious or lethal human disease for which preventive or therapeutic interventions may be available *(high individual risk but low community risk)* |
| Risk Group 4 (RG4) | Agents that are likely to cause serious or lethal human disease for which preventive or therapeutic interventions are not usually available *(high individual risk and high community risk)* |

Across the scientific and research compliance community, there is growing agreement that Risk Groups were designed for a much narrower research landscape. They focus on human disease severity and medical countermeasures, which makes sense historically. But modern research routinely involves animals, plants, arthropods, and environmental systems where human illness is not the dominant concern.

In those settings, Risk Groups fail to capture the primary risks. Institutions compensate by using informal interpretations or precedent, which works locally but leads to inconsistency nationally. The issue is not misuse of Risk Groups, but over-reliance on a framework that was never meant to carry this much weight.

## RG ≠ BSL (or ABSL, PBSL, ACL, etc.)

- Risk Group does not equal Biosafety Level
- Alignment weakens for animal, plant, and arthropod containment
- Risk changes with route, scale, and the possibility of escape
- Institutional Biosafety Committees bridge frameworks

Even in laboratory research, Risk Groups and Biosafety Levels are not interchangeable. Containment decisions depend on the facility, engineering controls, how work is performed, and not just which agent is used. That mismatch becomes much more pronounced in animal, plant, and arthropod research.

For example, Animal Biosafety Levels emphasize shedding, housing, and exposure. Plant Biosafety Levels emphasize spread and persistence. Arthropod Containment Levels emphasize escape and establishment. Risk Groups do not map cleanly to any of these systems, leaving Institutional Biosafety Committees to reconcile incompatible frameworks without consistent guidance.


There is broad recognition that Risk Groups systematically undervalue non-human consequences. An agent can pose little risk to people while still threatening livestock, crops, native species, or ecosystems. Environmental persistence and the ability to spread beyond containment may drive risk far more than human pathogenicity in these cases.

Agricultural and economic impacts are rarely captured in Risk Group thinking, even though they may represent the most significant real-world harm. This limitation becomes more visible as research increasingly intersects with food systems and the environment.

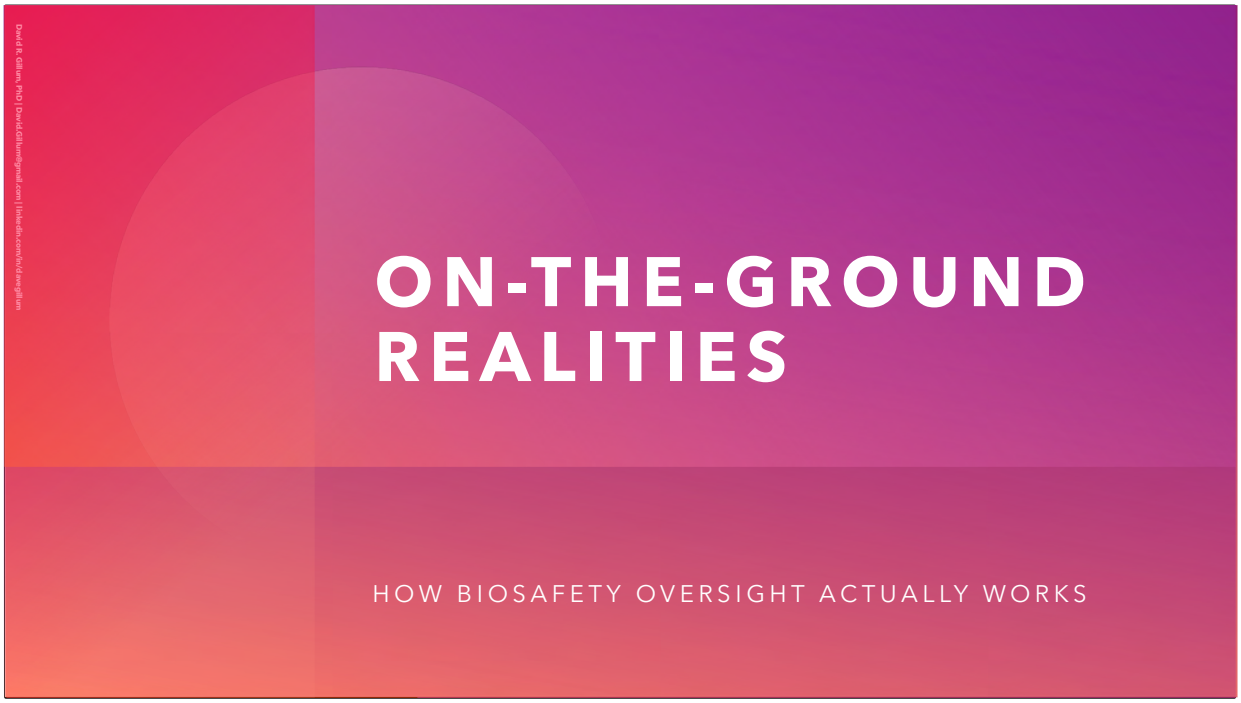


On-The-Ground Realities

How Biosafety Oversight Actually Works

## Biosafety Oversight Relies on Professional Judgment

- Oversight is not purely rule-based
- Practitioners rely on experience and judgment
- Guidance gaps require local interpretation
- Informal practices fill formal gaps

Biosafety oversight is often described as if it were a straightforward application of rules, but that is not how it works in practice. Biosafety professionals routinely rely on experience, judgment, and institutional understanding to make decisions, especially when guidance is incomplete or ambiguous.

When formal frameworks do not clearly address a situation, local committees and biosafety officers develop informal practices to bridge those gaps. This is not a failure of the system. It's evidence that oversight already depends on human judgment. Modernization efforts should recognize this reality and support it more explicitly, rather than assuming decisions can be fully automated by categories or checklists alone.

Practitioner interviews consistently have shown that biosafety decisions are interpretive, not mechanical. Professionals describe "working around" gaps in guidance and relying on peer norms, precedent, and experience. This supports the case for clearer decision support from NIH, not stricter rules. The system already runs on judgment. Policy should acknowledge and scaffold it.

## Biosafety Oversight Relies on Tacit Knowledge

- Much biosafety expertise is experience-based
- Knowledge develops through practice
- Lessons are often informal and locally retained
- Formal guidance rarely captures this expertise

A key insight from our empirical research on biosafety governance is that much of what keeps research safe is not fully written down. Biosafety professionals develop judgment through repeated exposure to complex situations, ambiguous guidance, and real-world decision-making. This expertise is often shared through mentorship, peer networks, and institutional memory rather than formal documents.

While this allows institutions to function effectively, it also creates vulnerabilities. When experienced personnel leave, when institutions change, or when new research areas emerge, this tacit knowledge may be lost or unevenly distributed. Written guidance alone cannot replace it.

Modernization efforts should recognize this reality. Policies that assume complete standardization miss how oversight actually works. Supporting mechanisms that help surface, share, and preserve practitioner knowledge would strengthen consistency without eliminating flexibility.

## Risk Assessment Is Situational, Not Formulaic


- Risk depends on how research is done and where it takes place
- Setting shapes exposure and impact
- Pathways and scale matter more than labels
- One-size frameworks fail across diverse research environments

Biosafety risk is not an intrinsic property of an agent alone. It depends on how the research is conducted, the scale of the work, and the setting in which it occurs. Practitioners consistently emphasize that exposure pathways and local environment often matter more than formal classifications.

This reinforces the idea that static, one-size frameworks struggle to scale across diverse research environments. Effective oversight requires structured flexibility that allows local committees to account for institutional dynamics explicitly, rather than forcing everything into predefined categories.

Practitioners describe risk assessment as situational and relational. They evaluate not just hazards, but who might be exposed, under what conditions, and with what downstream effects.

## Biosafety Risk Evolves Over the Research Lifecycle

- Biosafety risk changes over the life of a project
- Scale, application, and setting shift the risk profile
- Early laboratory work differs from later use or deployment
- Oversight should adjust as risk changes

Biosafety risk should not be understood as fixed at the moment a protocol is first reviewed. In practice, risk often evolves as research progresses. Changes in scale, application, delivery method, or deployment setting can significantly alter the risk profile of the same underlying system.

Early-stage laboratory work may present limited risk, particularly when it is exploratory or contained. As research advances toward scale-up, broader application, environmental interaction, or integration into operational settings, new exposure pathways and consequences may emerge.

Static frameworks struggle to capture these transitions. Modernization efforts could explicitly recognize that risk evolves over the research lifecycle and support oversight approaches that adapt as risk drivers change.

## How Oversight Moves Between Risk Tiers

- Movement between tiers follows changes in risk drivers
- Scale, exposure, setting, and biological function matter
- Transitions should be explicit, not left to assumption
- Committees need shared decision points

Much of the discussion around tiered oversight assumes that movement between levels of review is straightforward, but in practice it often is not. Institutional Biosafety Committees already make these judgments, but they typically do so implicitly, based on experience, precedent, or professional intuition rather than clearly articulated transition points.

Risk does not change simply because an agent changes categories. Risk shifts when underlying drivers change, such as scale of work, routes of exposure, methods of use, or the functional behavior of a system. These changes are often incremental, which makes them easy to miss when oversight relies on static classifications.

Modernization efforts could focus on making these transition points more explicit. This does not mean creating rigid thresholds or bright lines. Instead, it means identifying shared decision points where committees pause and ask whether the risk profile has materially changed enough to warrant a different level of review. Making those decision points explicit would improve consistency across institutions while preserving local judgment and flexibility.

## Biosafety Professionals Act as Boundary Spanners

- Practitioners bridge science, policy, and institutions
- Oversight involves negotiation and communication
- Social and institutional factors influence decisions
- Current frameworks rarely acknowledge this role

Biosafety professionals do more than evaluate protocols. They routinely act as intermediaries between scientists, administrators, compliance offices, and sometimes external stakeholders. Their work involves negotiation, explanation, and translation across domains.

This boundary-spanning role becomes especially important when research falls outside well-defined categories or raises broader institutional or societal concerns. Yet current oversight frameworks rarely acknowledge this function. Recognizing it explicitly helps explain why ethical, social, and institutional considerations naturally enter biosafety decision-making.

## Include Biosafety Practitioners in Policy Development

- Biosafety practitioners implement policy in practice
- Policy gaps often emerge during implementation
- Early practitioner input improves clarity and feasibility
- Collaboration reduces unintended consequences

Biosafety professionals are the individuals who translate policy into practice. They are responsible for interpreting definitions, aligning oversight requirements, and managing the real-world implications of regulatory decisions within institutions.

Our research makes clear that many of the challenges we see, ambiguity, inconsistency, and workarounds, emerge not because policy goals are wrong, but because policies are developed without sufficient input from those responsible for implementing them. As a result, institutions are left to resolve gaps independently, leading to national inconsistency.

Including biosafety practitioners earlier in the policy development process would improve clarity, feasibility, and alignment with institutional realities. This is not about shifting authority. It is about ensuring that policies designed to improve safety actually function as intended when applied in diverse research environments.

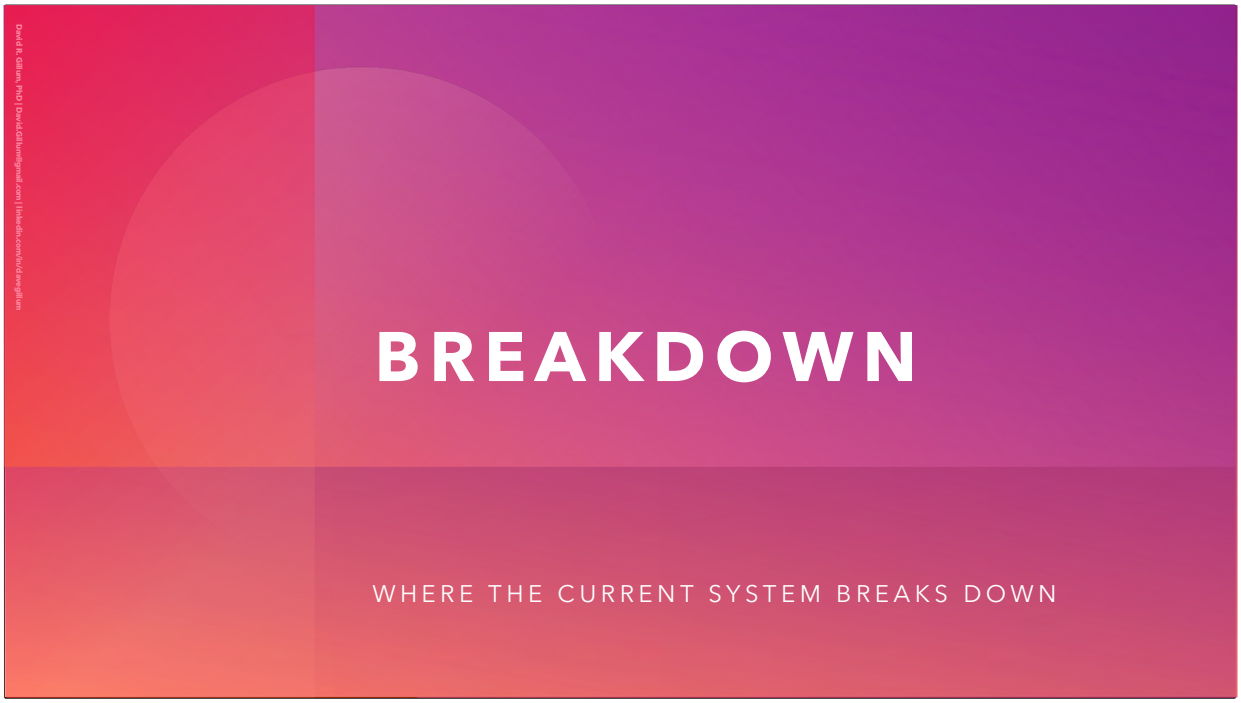


Breakdown

Where the Current System Breaks Down

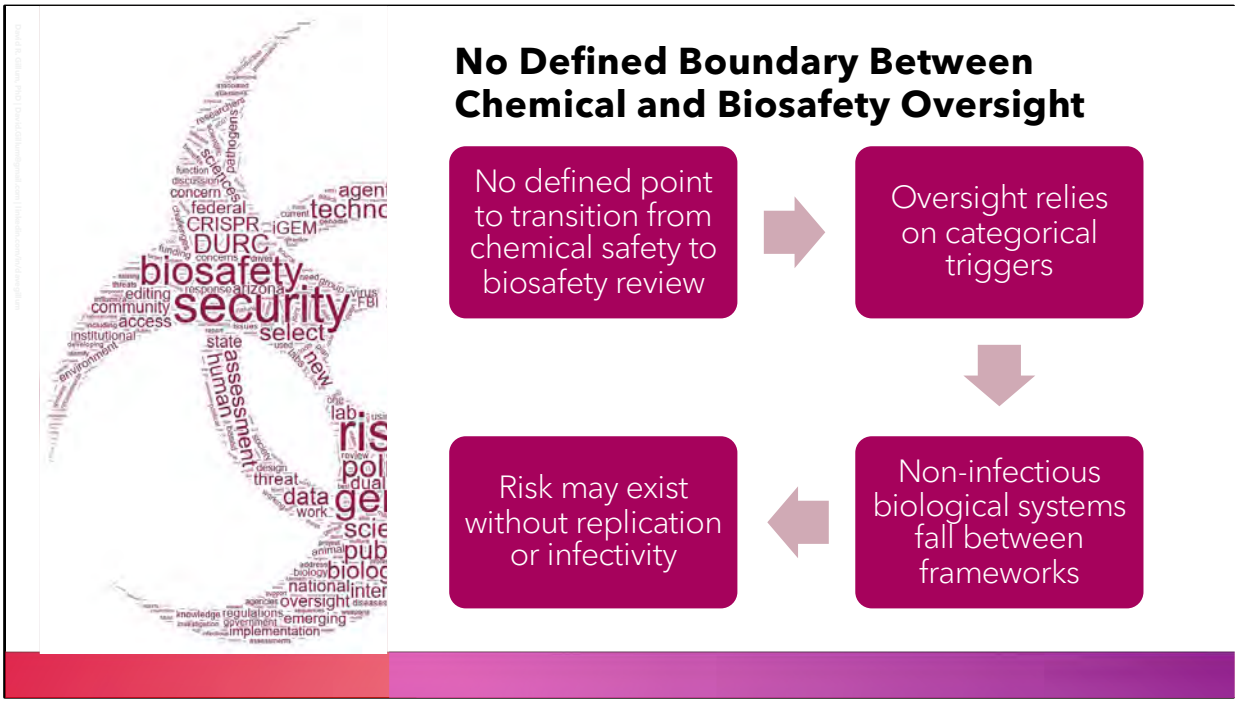


Across current scientific and compliance discussions, a recurring issue is the lack of a defined transition point between chemical safety and biosafety oversight. Existing frameworks tend to rely on categorical triggers, such as whether genetic material is present or whether a system can replicate.

That approach works when systems clearly fit one category or the other. It works much less well for biologically derived or biologically functional systems that are non-infectious and non-replicating. These systems may still pose risks driven by toxicity, persistence, scale, or interaction with living systems, yet they do not consistently trigger biosafety review.

The concern here is not that oversight is absent, but that there is no shared decision point that signals when chemical safety considerations are no longer sufficient on their own.

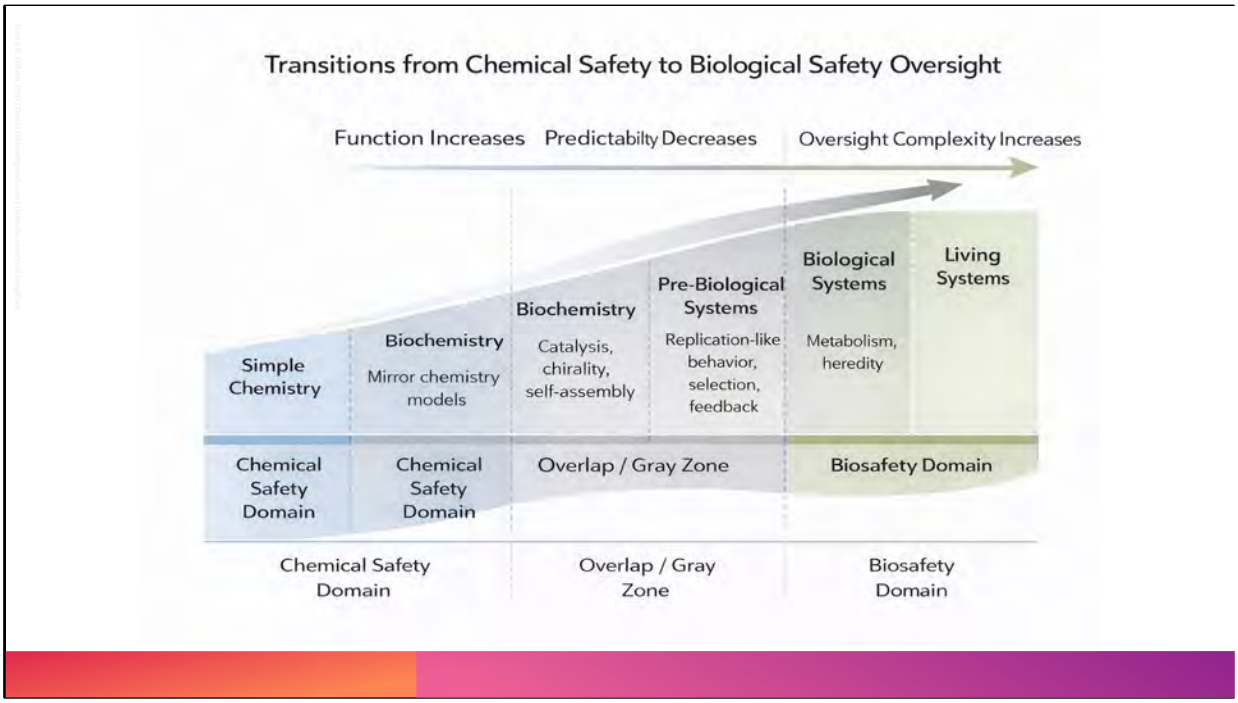


A recurring challenge in current oversight frameworks is the lack of a clearly defined transition point between chemical safety review and biosafety review. Existing approaches work well when materials are clearly chemical or clearly infectious, but they struggle with work that falls between those categories.

Oversight triggers often rely on categorical definitions, such as the presence of genetic material or the ability to replicate. These triggers are increasingly misaligned with modern research, where risk can arise without infectivity or replication.

As a result, chemically-produced and biologically derived or biologically functional systems can fall into gaps between oversight frameworks, leading to inconsistent review of similar risk profiles across institutions.

Clearer transition criteria, focused on risk drivers such as toxicity, persistence, scale, and exposure pathways, would help determine when chemical safety alone is no longer sufficient and when biosafety review should be integrated. This would improve consistency while preserving institutional flexibility.

## Modernization Aligns Oversight With Emerging Biotech

- Research increasingly spans chemical and biological domains
- Synthetic and hybrid systems challenge legacy definitions
- Mirror-biology may push current systems
- Guidelines and rules should anticipate, not react

The current NIH Guidelines were developed for a different research landscape. Today's work increasingly spans chemical, biochemical, and biological domains, often without clear boundaries. Synthetic systems, cell-free platforms, and novel nucleic acid constructs challenge legacy definitions.

Modernization is necessary not because existing oversight failed, but because research has evolved. Modernization efforts position NIH to anticipate emerging risk rather than respond after gaps appear.

## Inconsistent Oversight for Similar Risk Profiles

- Similar systems receive different oversight
- Review depends on institutional interpretation
- Fragmentation creates gaps and overlaps
- Committees lack shared decision support and transition points


Because there is no defined transition point, oversight outcomes often depend on how an institution interprets the boundary between chemical and biosafety domains. Two systems with similar risk profiles may be reviewed differently based on organizational structure, historical practice, or which committee encounters the work first.

As hybrid chemical-biological systems become more common, this fragmentation creates gaps in some cases and overlapping review in others. It also creates uncertainty for review committees and investigators, who may struggle to explain or defend why similar systems are treated differently.

What is missing is not expertise, but a shared, risk-based framework that helps determine when biosafety considerations should be integrated, regardless of whether a system meets traditional biological definitions.

## Modernization Must Be Implementable

- Policy assumptions often exceed institutional capacity
- Biosafety offices operate with limited staff and resources
- Broad triggers dilute focus on highest-risk work
- Effective oversight must match real-world operations

Inconsistency across institutions is often framed as a failure of oversight, but it is frequently a signal of misalignment between policy design and operational reality. Biosafety programs are typically small and responsible for multiple functions, including review, training, inspections, incident response, and coordination with other oversight bodies.

When oversight triggers are overly broad or insufficiently targeted, attention shifts away from work that presents the greatest risk. Committees and biosafety officers are forced to prioritize informally, which leads to variability across institutions.

Modernization should aim to improve decision quality and risk prioritization, not simply expand the scope of review. Aligning guidance with real-world operations is essential if changes are going to strengthen safety rather than add procedural burden.

## Add Structured Decision Support

- Institutional Biosafety Committees already assess complex risk
- Similar risks are reviewed differently across institutions
- NIH could provide structured decision support and learning mechanisms
- Improves consistency without limiting local judgment

Institutional Biosafety Committees already perform nuanced, risk-based assessments. They routinely evaluate agent characteristics, experimental design, scale, exposure pathways, and setting. The issue is not a lack of expertise. It is the lack of consistent national decision support.

Because current policies and guidelines provide limited guidance on how to weigh these factors, similar risk profiles may be handled differently across institutions. This variability reflects gaps in guidance rather than differences in safety culture or diligence.

Modernization could include structured decision support tools such as risk-tier matrices, decision checklists, and illustrative examples. These tools would not replace local judgment or flexibility. Instead, they would support committees in making decisions that are more consistent, transparent, and defensible across institutions.

## Baseline IBC Forms Would Improve Consistency

- IBCs use institution-specific registration forms
- No common baseline questions exist across institutions
- Forms must be updated as policies and guidance change
- A baseline template would support consistency and learning

Currently, every institution's Institutional Biosafety Committee uses its own protocol registration forms. While this allows local flexibility, it also means there is no shared baseline for the questions investigators are asked or the risk factors committees consider. Similar research may be framed and reviewed very differently depending on the institution.

IBCs and biosafety offices are also responsible for keeping these forms up to date as policies, guidelines, and regulations evolve. This requires continuous monitoring and revision, and changes may be implemented unevenly across institutions.

Providing a baseline template of questions would not eliminate local customization. Instead, it would give institutions a consistent foundation to build on, supporting more uniform risk assessment, reducing administrative burden, and making it easier to incorporate future policy changes in a timely and consistent way.

## Learning From Incidents and Near Misses

- Incidents and near misses provide critical insight
- Informal reporting already occurs at many institutions
- Learning-focused systems improve consistency and trust
- Reporting should emphasize improvement, not blame

A consistent theme from practitioner experience is the importance of learning from incidents and near misses. Many biosafety professionals already track these events informally to refine practices, adjust oversight, and reduce the chance of recurrence.

Near misses, in particular, often reveal weaknesses in procedures, training, or assumptions about risk before harm occurs. However, when reporting systems are unclear or perceived as punitive, valuable information can be lost.

A modernization effort could consider how existing policies and guidelines support learning from these events in a way that emphasizes improvement rather than blame. This is not about enforcement or surveillance. It is about creating feedback loops that strengthen decision-making, improve consistency across institutions, and reinforce a culture of safety and transparency.

## Modern Biosafety Oversight Requires Feedback Loops

- Oversight decisions shape practice over time
- Practice reveals gaps in guidance
- Learning should inform future policy updates
- Modernization should be iterative, not static

Modern biosafety oversight does not operate as a one-time decision. Oversight decisions influence how research is conducted over time, shaping laboratory practices, training norms, and institutional expectations. In turn, those practices reveal where guidance works well and where it falls short.

Biosafety practitioners routinely encounter situations where existing guidance does not fully address emerging research approaches, hybrid systems, or evolving use settings. In these cases, institutions develop local solutions, workarounds, or informal norms. While this allows research to proceed safely, it also means that lessons learned often remain localized.

A modernized oversight framework should support feedback loops that allow learning from practice to inform future guidance and policy updates. This includes insights gained from incidents, near misses, and routine oversight decisions, not just adverse events.

Framing modernization as an iterative process acknowledges that bioscience will continue to evolve. Rather than relying on static classifications or periodic overhauls, oversight systems should be designed to adapt over time,

incorporating real-world experience to improve clarity, consistency, and effectiveness.


Societal Considerations

Dealing with significant societal consequences

## The Need for Adding Ethical Expertise to IBCs

- Focused on biosafety governance, not human subjects review
- Some biosafety decisions involve more than technical biosafety and biosecurity risk
- Risk mitigation does not resolve all oversight questions
- Institutional Biosafety Committees currently lack formal ethics capacity
- Assessment of significant societal consequences must occur at the local level

As biosafety oversight evolves, some decisions extend beyond technical questions of containment, exposure pathways, or procedural controls. In certain cases, even when risks can be mitigated, there remains a question about whether the remaining level of risk is acceptable given the intended benefits of the research.

Increasingly, those determinations involve assessing the potential for significant societal consequences. These may include impacts beyond the laboratory, such as effects on surrounding communities, environmental tolerance, public confidence, or broader institutional risk. Importantly, these assessments cannot be made generically or centrally. They must occur at the local level, where institutional setting, infrastructure, and community setting are understood.

Institutional Biosafety Committees are the bodies best positioned to make these local assessments. However, while IBCs are well equipped to evaluate technical biosafety risk, they are not typically structured to explicitly assess ethical tradeoffs or societal impact. Adding ethical expertise would help ensure that evaluations of significant societal consequences are deliberate, transparent, and well supported.

## How Ethical Expertise Strengthens Biosafety Oversight

- Supports evaluation of societal risk–benefit tradeoffs
- Improves consistency and transparency in decision-making
- Helps document how societal consequences were considered
- Reinforces public confidence in oversight decisions

When oversight decisions involve significant societal consequences, they inherently involve value-based tradeoffs. Ethical expertise can help committees surface those tradeoffs, evaluate them explicitly, and document how societal considerations informed the final decision.

This does not require turning Institutional Biosafety Committees into ethics boards or duplicating other review mechanisms. Instead, ethical expertise can function as targeted decision support, strengthening the quality, clarity, and defensibility of biosafety oversight, particularly in cases where public trust and community impact are at stake.

## Modernization Should Integrate Existing Oversight Systems

- Biosafety oversight does not exist in isolation
- Chemical safety, ethics, and compliance already play roles
- Modernization should clarify coordination, not add layers
- Clear roles reduce duplication and confusion

As biosafety oversight becomes more risk-based and more attentive to societal consequences, it increasingly intersects with other oversight systems that institutions already have in place. These include chemical safety programs, environmental health and safety offices, compliance functions, and, in some cases, ethics review mechanisms.

One concern that often arises in modernization discussions is the risk of creating overlapping or duplicative oversight. That is not the goal. The goal is clarity. Modernization should help institutions understand how biosafety oversight coordinates with existing systems, rather than competing with them or operating in parallel.

Clear articulation of roles and coordination points would reduce confusion for investigators and committees alike. It would also help ensure that risks are addressed holistically without unnecessary duplication of review. This kind of integration strengthens oversight while respecting the structures institutions already rely on.

## Coordination Across Biosafety Policies and Agencies

- Biosafety oversight spans multiple agencies and policies
- Current requirements are fragmented and evolve unevenly
- Institutions must reconcile overlapping guidance independently
- Clear coordination improves consistency and implementation

Biosafety oversight does not operate under a single policy or agency. Institutions must navigate guidance and requirements from multiple federal entities, often addressing related risks through separate frameworks that evolve on different timelines. In practice, alignment across these policies is not always clear, and institutions are left to reconcile overlaps, gaps, or inconsistencies on their own.

This creates variability in how similar research is reviewed and managed across institutions. It also places a significant burden on local biosafety programs, which must track changes across multiple policies while ensuring day-to-day compliance.

Modernization presents an opportunity to improve coordination across this policy landscape. Even modest alignment in definitions, expectations, and assumptions can improve consistency, reduce confusion, and support more effective implementation without creating new regulatory layers.


Modernization path

Steps to Creating a Better Governance System

## Risk Groups as One Input, Not the Backbone

- Risk Groups remain useful baseline information
- They should not determine oversight alone
- Decisions should integrate agent, system, and use
- Exposure pathways and setting matter

In current scientific and compliance discussions, there is broad agreement that Risk Groups still have value. They provide a shared starting point and a common language around hazard severity, particularly for human laboratory research. The concern is not that Risk Groups are wrong, but that they are being asked to carry more responsibility than they were designed for.

As research expands across animals, plants, arthropods, and environmental systems, decisions increasingly depend on how an agent is used, the system it is introduced into, and the exposure pathways that result. In those cases, Risk Groups alone are not sufficient to determine appropriate oversight or containment.

A modernization effort could explicitly reposition Risk Groups as one input into a broader risk evaluation, rather than the backbone of oversight decisions. That shift would better reflect how Institutional Biosafety Committees already approach complex protocols.

## Develop Non-Human Risk Classification Counterparts

- Human Risk Groups cannot represent all risk domains
- Parallel risk concepts needed for:
  - Animals
  - Plants
  - Arthropods
  - Environmental systems
- Risk language should align with containment systems

A recurring theme in current discussions is that a single, human-centric risk classification framework cannot reasonably support all areas of bioscience. Animal, plant, arthropod, and environmental research each involve distinct types of harm, exposure, and spread that are not captured by human disease-focused Risk Groups.

As a result, committees often try to translate human Risk Groups into Animal Biosafety Levels, Plant Biosafety Levels, or Arthropod Containment Levels, even though those systems were not designed to align directly. This creates friction and inconsistency.

One potential path forward is the development of parallel or complementary risk classification concepts that speak directly to non-human systems and align more naturally with their respective containment frameworks. The goal would be shared understanding and consistency, not additional layers of review.


Many current oversight triggers rely on categorical definitions, such as whether genetic material is present, whether an organism can replicate, or whether something meets a traditional definition of a biological agent. These triggers work well for classic use cases, but they increasingly fall short for modern research.

There are many systems where risk is driven not by infectivity or replication, but by factors such as toxicity, environmental persistence, scale of production, or interaction with living systems. Examples include cell-free systems, bioactive molecules, and hybrid biological-chemical constructs.

A modernization effort could shift emphasis away from categorical triggers and toward the underlying risk drivers that actually determine potential harm. This would better align oversight decisions with how risk manifests in practice.

## Clarify When Expedited Review Is Appropriate

- Expedited review works when risk profile is unchanged
- Appropriate for administrative or low-risk modifications
- Not appropriate when risk drivers change
- Risk stability should guide review level

Expedited review is an essential tool for efficient biosafety oversight, and there is broad agreement that it should be used when risk is well understood and stable. The key determinant should not be familiarity with an agent, but whether the overall risk profile has changed.

Administrative updates, renewals without substantive modifications, and work involving well-characterized low-risk systems are appropriate candidates for expedited review. In contrast, changes that affect scale, exposure pathways, host range, environmental interaction, or intended use should trigger full committee review.

Clarifying these principles would support consistent and defensible use of expedited review across institutions while preserving appropriate safeguards and flexibility.

## What NIH Modernization Should Prioritize

- Align risk frameworks with modern research
- Support consistent local decision-making
- Clarify transitions and tiered oversight
- Preserve flexibility while strengthening trust

These discussions point toward a set of clear priorities for modernization. First, risk frameworks need to align with the realities of modern research, including non-human systems, emerging technologies, and evolving use applications. Second, Institutional Biosafety Committees need support in making consistent, defensible decisions without losing the flexibility that allows them to respond to local conditions.

Clarifying how oversight transitions between tiers, including when expedited review is appropriate, is a key part of that effort. Just as important is maintaining public and institutional trust by ensuring that oversight decisions are transparent, well reasoned, and grounded in both technical and societal considerations.

Modernization does not require reinventing the system. It requires refining it so that it better reflects how research is conducted today and how oversight already functions in practice.

## Why Modernization Improves Safety Outcomes

- Focuses oversight on real risk drivers
- Reduces attention spent on low-risk work
- Improves early identification of emerging risk
- Strengthens consistency and transparency

Modernization is not about expanding oversight or adding new layers. It is about focusing attention where it matters most. By shifting away from static categories and toward risk drivers, oversight resources can be directed toward work with the greatest potential for harm.

Clearer transitions, decision support, and learning mechanisms also improve early identification of emerging risks. Consistency across institutions strengthens trust, both within the research community and with the public. Together, these changes support better safety outcomes without constraining innovation.


Questions?

David Gillum, PhD
David.Gillum@gmail.com
linkedin.com/in/davegillum/