1900 L Street NW
Suite 302
Washington DC 20036, USA
contact.isbr@gmail.com
https://isbr.info

ISBR Webinar 3, 22Oct26: Synthetic virus genomes: design & biosafety considerations

AI-generated meeting notes from Microsoft Teams, please confirm accuracy.

Quick Read

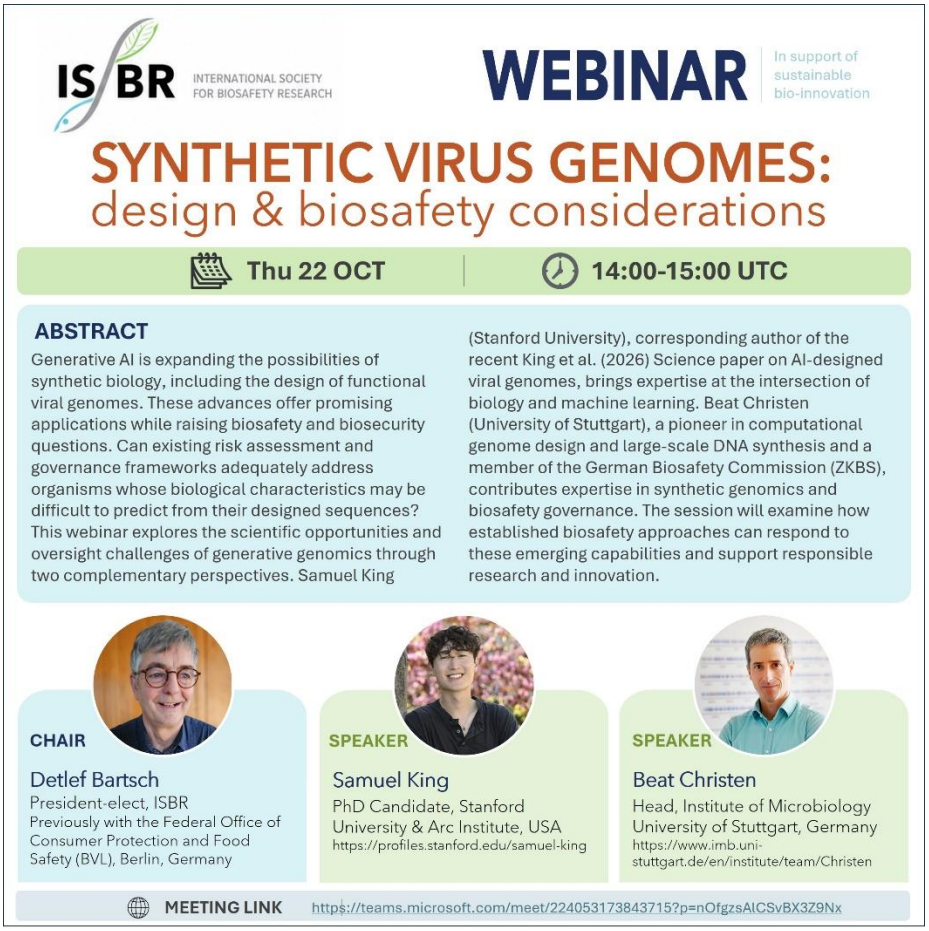
A video recording of the webinar is available HERE.

The third webinar of the International Society for Biosafety Research (22 October 2026) paired the first demonstration of AI-generated viral genomes with an assessment of what that capability does to existing biosafety control points.

Samuel King (Stanford) described the use of genome language models, trained on millions of natural genomes and prompted in DNA rather than natural language, to generate complete bacteriophage genomes, using the historically significant phiX174 as a design template. Of roughly 300 candidates synthesised and individually tested, 16 were functional, despite carrying several hundred mutations each and, in some cases, exceeding the conventional threshold for a new species. A designed-in host-tropism filter held the functional set to the intended E. coli host. A cocktail of the generated phages overcame antibiotic-resistant strains that natural phages could not, and a start-up had applied the method to a clinical case days before the webinar. Work with Google DeepMind has shown that generated phage genomes can be watermarked and detected.

Beat Christen (member, German Central Committee on Biological Safety) placed this within what he termed digital biology: the convergence of AI, DNA synthesis and the living world into a strategic platform rather than a single technology. His central argument for regulators is that the three control points currently proposed, access to models and compute, sequence-based synthesis screening, and the barrier of specialist expertise, are each being weakened by the same technology they are meant to govern, and that agentic AI linking design, synthesis and experiment compresses the cycle further. His recommended posture is anticipatory and internationally coordinated, explicitly distinguished from centralised control, combined with investment in education and in distributed manufacturing capacity so that response capability is as widely held as the capability itself.

The common conclusion across both talks and the discussion: the demonstrated capability is still bounded and assessable on familiar grounds of traits and host range, but the trajectory toward more complex systems, and eventually mammalian viruses, is foreseeable. Frameworks should be written for that trajectory rather than for the present snapshot, and the window for doing so is now.



1. At a glance

- **Meeting:** ISBR Webinar 3, *Synthetic virus genomes: design and biosafety considerations*, 22 October 2026, 14:00 UTC, 73 minutes. Third in the series, following webinars on data requirements and on gene drive organisms.
- **Chair:** Detlef Bartsch, ISBR President-elect.
- **Speakers:** Samuel King, PhD candidate in bioengineering, Stanford University (Brian Hie laboratory), first author of the *Science* paper on generative design of bacteriophages with genome language models; Beat Christen, member of the German Central Committee on Biological Safety (ZKBS), which publishes monitoring reports on synthetic biology.
- **Format:** two 20-minute presentations with a moderated Q&A, questions submitted in writing.
- **Attendance:** peaked at 124 participants.
- **Attribution:** the Chatham House Rule was applied. Presentations are to be published by ISBR and are attributed here; contributions from the floor are not.

2. Presentation 1 — Generative design of bacteriophage genomes (S. King)

The problem framed. Biological design has historically succeeded at the scale of single genes or small circuits. Whole-genome design has been intractable because constraints compound: single point mutations can be lethal, elements interact, and some dependencies lie outside the genome altogether. The enabling shift is data volume — public sequence repositories now hold on the order of 10^8 sequences — combined with models able to learn design principles from it.

Method, at the level relevant to assessment. The Evo 1 and Evo 2 genome language models were trained on millions of genomes spanning eukaryotes, prokaryotes and phages. They operate only in DNA: a design intent

cannot be expressed in natural language, only as a seed sequence. Generic models did not produce plausible candidates; further training on the target family (*Microviridae*) was required before outputs resembled viable genome architecture. *phiX174* served as the template — a small genome of about 5,000 nucleotides and 11 genes, mostly overlapping and almost all essential, with any single protein mutation around 20% lethal. It was also the first DNA genome sequenced (1977) and the first phage genome synthesised (2003).

Design filters. Candidate genomes were screened against criteria calibrated to the natural family, including a host-tropism filter that selected for spike proteins resembling those of the natural phage, so that candidates would be restricted to the intended *E. coli* C host.

Results. Approximately 300 final candidates were synthesised and each tested individually; 16 were functional. Functional designs diverged substantially from natural genomes — several hundred mutations each, including insertions, deletions, recombinations and genes transferred from other known phages, with some crossing the threshold conventionally used to delimit a new phage species. Observed viability far exceeded what the known per-mutation lethality rate would predict. Tested against eight related *E. coli* strains, the functional set infected only the target strain and one further strain that the natural phage also infects — evidence that the tropism filter performed as intended.

Application and pace of translation. Phage therapy requires diverse phages to limit the evolution of resistance, and sourcing that diversity from the environment is slow. A cocktail of the 16 generated phages overcame two resistant *E. coli* strains within one and two passages respectively, where *phiX174* alone and a cocktail of 16 natural isolates never did. A start-up had used the method to produce a phage therapy applied in a clinical case of antibiotic-resistant infection days before the webinar — a point the speaker himself flagged as striking for its speed.

Traceability. Work with Google DeepMind has demonstrated that generated phage genomes can be watermarked and subsequently detected. This was offered explicitly as a contribution to the regulatory debate.

3. Presentation 2 — Digital biology and the erosion of control points (B. Christen)

The framing. Reading genomes moved biological information into the digital domain. Generative design plus synthesis reverses the flow: a genome is designed digitally, instantiated as physical DNA, and introduced into a cell. What this produces is not a single technology but a **strategic platform** formed by the convergence of three fields — artificial intelligence, DNA synthesis, and the living world — each of which is independently a focus of global strategic competition.

The trajectory. Design capability has progressed from molecules (binders, antibodies), through molecular machines and systems, to whole genomes, in a few years. The *phiX174* lineage illustrates the pattern: sequenced 1977, synthesised 2003, megabase genomes synthesised since, and now AI-generated genomes of the same phage. The speaker was explicit that no date can be attached, but that progressive extension to more complex biological systems should be anticipated — with mammalian viruses named as the next stage, where the risk profile differs fundamentally from phage work.

The central governance argument. The three control points currently proposed are each being weakened by the very technology they are intended to govern:

Proposed control point	Why it holds today	Why it weakens
Restrict access to models and compute	Frontier models require concentrated infrastructure held by a few actors	Distributed and local models, cloud services, and open weights that can be further trained spread the capability to many actors

Screen DNA synthesis orders and customers	A limited number of providers ship globally, making a practical chokepoint	Function is increasingly difficult to infer from sequence alone for novel designs; distributed and benchtop synthesis dilutes the chokepoint
Rely on the barrier of specialist expertise	Molecular biology competence has required years of training and tacit laboratory experience	AI-assisted workflows provide skills uplift — synthesis of literature, contextual guidance, problem solving — compressing what took decades to acquire

He noted that certification or evaluation of AI models before release, and strengthened synthesis screening, remain worth implementing now — but should be understood as measures for the present, not durable solutions.

The inflection point ahead. The more significant shift is agentic AI that connects all three domains rather than operating within one: specifying a design, optimising it for synthesis, directing the experiment, and feeding results back into further design cycles. This would accelerate beneficial discovery substantially while unsettling assumptions about where technical expertise, human oversight and institutional control sit.

Proposed response. If capability becomes distributed, the capacity to respond must be distributed and internationally connected too. Two strategies are needed simultaneously: **global intelligence** — education in these technologies and international scientific networks, bringing the next generation in early — and **local resilience** — a global net of DNA manufacturing capability rather than a few centres, so that a collective response to a threat can be mounted quickly. His closing proposition: technological capability consistently outpaces institutions, and responding with new restrictions after each breakthrough that attracts public attention, while sometimes necessary, is not a substitute for building long-term responsibility now.

4. Discussion

Scalability to larger genomes. Transfer to other bacteriophages was considered quite feasible now that the method exists, given adequate training data and the ability to test in the laboratory. Bacterial or animal genomes remain distant, held back by three simultaneous bottlenecks: computational evaluation (knowing which candidates are worth testing), synthesis cost at that scale, and the absence of experimental methods for introducing very large DNA constructs into living cells. Advances would be needed in all three.

Novelty versus reorganisation. A question from the floor pressed this point specifically to anchor practical biosafety assessment, since assessment rests on potential new traits. The answer: the functional genomes were relatively close to natural genomes in overall architecture, while genuinely novel within their family. Because the models learn only from natural genomes, radically unfamiliar designs are both unlikely to be produced and effectively un-evaluable — there is no basis for judging which such candidate would be worth testing. Design of a wholly novel virus would require combining many methods at the level of individual proteins and assemblies.

Mirror life. Both speakers declined to speculate. No clear route was identified from current models, trained on natural genomes, to mirror genomes. Both welcomed that the biosecurity community is already anticipating the question.

Societal readiness and international governance. A reasonably clear picture of the present is available, but blind spots are likely because the fields are interconnected and each development triggers secondary ones. International coordination was held to be increasingly important, and was expressly distinguished from centralised control: use will be global and the resulting problems global, so coordination and shared resilience are the realistic instruments.

Enthusiasm versus caution. Both postures are present on the same globe simultaneously — some nations, institutions and companies will proceed where they see advantage, while others pause to deliberate. Since no

institution can be compelled to open that deliberation, the practical contribution of bodies such as ISBR is to convene a discussion that bridges jurisdictions.

Where risk is most likely to arise. Immediate deliberate misuse was viewed as less probable than risk emerging through legitimate development — the example offered was engineering a stealth virus to deliver therapeutic genes to the brain — where the risks are not yet characterised and the development cycle will run ahead of understanding.

5. Take-home messages

1. **AI-designed viral genomes have moved from prospect to practice.** Complete functional phage genomes have been generated, validated in the laboratory, and already applied in a clinical case. The interval between demonstration and use was measured in days.
2. **Novelty is still bounded — and that is the opportunity.** Functional designs remain within the architecture of their natural family, so assessment can still proceed on the familiar grounds of traits and host range. This is the window in which durable frameworks can be built, before designs outrun the basis for evaluating them.
3. **Containment and traceability can be engineered at the design stage.** The host-tropism filter restricted host range as intended, and watermarking of generated genomes was demonstrated. Both are concrete candidate hooks for oversight rather than aspirations — arguably the most actionable regulatory finding of the webinar.
4. **The conventional control points are eroding together.** Compute and model access, sequence-based synthesis screening, and the expertise barrier each weaken as the technology advances. A framework resting primarily on any one of them will not hold for long.
5. **Govern the trajectory, not the snapshot.** The next steps — other phages, then progressively more complex systems, and eventually mammalian viruses — are foreseeable in direction even where they cannot be dated. A framework addressing only today's concerns will be obsolete on arrival.
6. **Coordination is not centralisation.** The posture advocated was internationally coordinated anticipation combined with distributed capability, education and response capacity — not a single controlling authority, which the global distribution of the technology would in any case defeat.
7. **The near-term risk pathway runs through legitimate work.** Deliberate misuse was judged less likely in the immediate term than uncharacterised risk emerging inside beneficial development cycles — which places the burden on risk assessment capability rather than on prohibition.