



Out-of-Band Management Network Design for GPU Clusters

Plan and build private AI compute with GPU Smith. We connect workload requirements with GPU hardware, networking, deployment and operating decisions so your infrastructure fits the work it must perform.

Published by GPU Smith · 2026-10-29 · out of band management network gpu cluster · bmc redfish gpu server · dgx bmc network 1gbe · redfish gpu telemetry · bmc security best practices data center · openbmc redfish

Original article: <https://gpusmith.com/articles/en/out-of-band-management-network-bmc-redfish-gpu-clusters>



In this article

1. The four networks in a GPU cluster
 2. What DGX generations do
 3. OOB topology
 4. BMC capabilities
 5. Redfish for telemetry and provisioning
 6. BMC threat model and segmentation
 7. Firmware resiliency
 8. Provisioning flow
 9. DPU-based management
 10. Commissioning checklist
 11. Sources
-

The four networks in a GPU cluster

A GPU cluster that loses its out-of-band (OOB) network still trains models, but nobody can power-cycle a hung node or reflash a bad BMC image without walking to the rack. NVIDIA's DGX SuperPOD B300 reference architecture enumerates a compute fabric, an InfiniBand or Ethernet storage fabric, an in-band management network and an out-of-band management network, with the two management networks built on a shared Ethernet fabric "segregated with VXLAN and EVPN" ([NVIDIA DGX SuperPOD B300 XDR RA: Network Fabrics](#)).

The four planes carry different traffic and fail differently:

Network	Typical media	Carries	Must survive
1. Compute fabric	400-800 Gb/s InfiniBand or Ethernet	GPU-to-GPU collectives	Job-level failures
2. Storage fabric	100-400 Gb/s Ethernet (RoCE) or InfiniBand	Checkpoints, datasets	Storage maintenance
3. In-band management	100-400 Gb/s Ethernet	Provisioning, scheduler, NFS home, user access	Host OS failures on some nodes
4. Out-of-band (OOB)	1 GbE copper	BMC Redfish/IPMI, switch and PDU management	Host OS down, node powered off, in-band fabric down

The defining property of the OOB plane is in the last column. The NSA/CISA guidance describes BMCs as components that "operate separately from the operating system and firmware to allow for remote management and control, even when the system is shut down" ([NSA/CISA: Harden Baseboard Management Controllers](#)). Routing BMC access through the in-band fabric removes that property exactly when it is needed. The high-speed fabrics are covered in the guide on [fabric sizing, ports and rails](#).

What DGX generations do

The split has shifted across DGX generations, mainly in where the in-band ports live and how many BMCs a node carries. Values are copied from the DGX user guides fetched on 2026-10-28 (the H100 and B200 pages show "Last updated on Oct 17, 2026"; the B300 page shows "Last updated on Oct 21, 2026"; none carries a revision number).

System	System BMC	Host / in-band management	Storage and in-band card
1. DGX H100	1 GbE RJ45; Redfish, IPMI, SNMP, KVM, Web UI	Optional: dual-port 100GbE in slot 3 and 10 GbE RJ45	2 x ConnectX-7 dual-port Ethernet
2. DGX B200	1 GbE RJ45; Redfish, IPMI, SNMP, KVM, Web UI	Optional: dual-port 100GbE in slot 3 and 10 GbE RJ45	2 x BlueField-3 DPU dual-port cards
3. DGX B300	1 GbE RJ-45; Redfish, IPMI, KVM, remote storage; BMC on replaceable DC-SCM	On-board 1 GbE RJ-45 (eno1)	2 x dual-port BlueField-3, NIC mode by default

Sources: [DGX H100 user guide](#), [DGX B200 user guide](#), [DGX B300 user guide](#).

Three consequences follow. First, every generation keeps a **1 GbE copper BMC port**, so the OOB switch remains a commodity 1 GbE access switch. Second, from B200 onward the in-band and storage ports sit on BlueField-3 cards, each with its own BMC. The B300 XDR reference architecture states that "the BMC from both NVIDIA BlueField cards are connected to the out of band Network" alongside the system BMC ([NVIDIA DGX SuperPOD B300 XDR RA: Network Fabrics](#)). A DGX B300 therefore needs three OOB copper drops, not one. Third, the B300's on-board 1 GbE host port is OS-owned, and neither the user guide nor the reference architecture pages reviewed here say where it should land. Treating it as in-band (a separate VLAN or VRF from the BMCs) is the conservative choice: a host-OS interface in the BMC segment gives any compromised host a path to every BMC.

OOB topology

The B300 reference architecture uses NVIDIA SN2201 switches for OOB, described as "48 port 1 Gbps Ethernet and 4 x 100 Gbps" ([NVIDIA DGX SuperPOD B300 RA: Key Components](#)). The OOB network "is physically rolled up into the aggregation layer (spine layer) of each SU as a dedicated VXLAN" and "carries all IPMI related control traffic" ([NVIDIA DGX SuperPOD B300 RA: Network Fabrics](#), last updated Nov 19, 2025). The XDR variant adds that the OOB network also "serves as the network for fabric management of the compute InfiniBand fabric" ([NVIDIA DGX SuperPOD B300 XDR RA: Network Fabrics](#), last updated Sep 02, 2026). The arithmetic below applies to any 48-port 1 GbE switch.

Port-count formula

For one scalable unit or pod:

- $P_{\text{node}} = N \times (b + d + h)$, where N is the number of GPU nodes, b system BMC ports per node, d DPU or SuperNIC BMC ports per node, and h host management ports placed on the OOB switch (0 if they go in-band).

- P_{infra} = PDU, CDU/cooling controller, fabric switch, storage controller, UFM/management server BMC and console server ports.
- $P_{\text{total}} = \text{ceil}((P_{\text{node}} + P_{\text{infra}}) \times (1 + s))$, where s is the spare-port policy.
- $\text{Switches} = \max(\text{ceil}(P_{\text{total}} / 48), \text{placement minimum})$, since cabling policy can force more switches than port count.

Worked example (illustrative)

Assumptions, all illustrative: 64 DGX B300 systems (the SU size in the B300 Spectrum-4 reference architecture, per [NVIDIA DGX SuperPOD B300 RA: Key Components](#)); $b = 1$, $d = 2$ (two BlueField-3 BMCs), $h = 0$; 16 racks of 4 systems; 2 metered PDUs per rack; 40 infrastructure management ports for fabric switches, storage, management servers and cooling controllers (a placeholder, not a vendor figure); $s = 20\%$.

1. Node ports: $64 \times (1 + 2 + 0) = \mathbf{192}$.
2. PDU ports: $16 \times 2 = 32$.
3. Infrastructure: 40.
4. Subtotal: $192 + 32 + 40 = 264$.
5. With 20% spare: $264 \times 1.2 = 316.8$, rounded up to **317**.
6. Port-count minimum: $\text{ceil}(317 / 48) = \text{ceil}(6.6) = \mathbf{7}$ switches (336 ports).

The placement policy changes the answer. A switch in every compute rack needs 16 switches for $16 \times (12 + 2) = 224$ rack-local ports, or 14 of 48 ports used per switch (29%), plus separate switches for the infrastructure rows. One switch per pair of racks needs 8 switches at 28 ports each (58%), with copper runs crossing one rack boundary. Seven central switches minimize hardware but need long copper runs.

Uplink bandwidth is rarely the constraint: $48 \times 1 \text{ Gb/s} = 48 \text{ Gb/s}$ of access ports against two 100 Gb/s uplinks is 48:200, or 0.24:1. Steady-state telemetry is small: if, illustratively, 192 BMCs each return a 50 KB sensor payload every 10 seconds, the aggregate is $192 \times 50 \text{ KB} / 10 \text{ s} = 960 \text{ KB/s}$, about 7.7 Mb/s. The event that matters is mass reimaging via virtual media: illustratively, a 2 GB image to 64 nodes is $64 \times 16 \text{ Gb} = 1,024 \text{ Gb}$, at least $1,024 / 25 = 41$ seconds from an image server on one 25 Gb/s link, before BMC-side limits that should be measured rather than assumed.

Addressing follows from the same count: 317 addresses exceed a /24 (254 usable hosts), so either a /23 per SU or separate /24s for node BMCs and infrastructure.

Power and redundancy

The OOB switches should not share fate with the nodes they manage: feed them from a separate or UPS-backed circuit, give each two uplinks to different aggregation switches, and keep a console-server path to the aggregation layer. The DGX guides reviewed here list one 1 GbE BMC interface per system, so a failed OOB access switch isolates every BMC behind it; the per-pair placement above bounds that blast radius to 8 nodes in the worked example.

BMC capabilities

DGX H100 and B200 BMCs list "Redfish, IPMI, SNMP, KVM, and Web user interface" ([DGX H100 user guide](#); [DGX B200 user guide](#)). The DGX B300 BMC page lists the ports to open on the firewall, which doubles as an inventory of the BMC's attack surface ([DGX B300 user guide: Using the BMC](#)):

Port	Protocol (as listed)	Function	Recommended exposure
1. 443	HTTPS / Redfish	Web UI; Redfish with authentication	Automation and jump hosts only
2. 80	Listed as HTTPS	Redfish service root	Redirect or block
3. 623	RMCP+	IPMI	Disable or restrict to legacy tools
4. 7582	KVM	Secure KVM redirection	Jump hosts only
5. 5124	CD media redirection	Secure CD media	Provisioning hosts only
6. 161, 199	SNMP	SNMP requests, SMUX	Monitoring hosts, read-only
7. 123	NTP	Time sync	Internal NTP only
8. 1900, 50000, 427	UPnP, SLP	Discovery	Block

The exposure column is this guide's suggestion, not NVIDIA's. Preferring Redfish over IPMI is partly protocol security: in IPMI 2.0 RMCP+ RAKP authentication, a remote attacker can obtain password hashes from the BMC's RAKP message 2 response for offline guessing, a specification-level weakness ([NVD: CVE-2013-4786](#)). NVIDIA's BlueField BMC documentation (v26.07, updated September 1, 2026) puts it plainly: "IPMI is supported for backward compatibility," and "it is recommended to start new deployments with Redfish only" ([NVIDIA BlueField BMC Software v26.07: BlueField Management](#)).

Redfish for telemetry and provisioning

DMTF's Redfish Specification (DSP0266 version 1.25.0, dated 2026-08-07) defines Redfish as "a standard that uses RESTful interface semantics to access a schema-based data model to conduct management operations," with JSON payloads and OData conventions ([DMTF DSP0266 1.25.0](#)). DMTF lists Release 2026.2 as current ([DMTF: Redfish](#)). Normative details that matter for tooling:

- The service root is `/redfish/v1/`.
- Implementations "shall support" TLS 1.2 or later, and only cipher suites IANA marks "Recommended."
- HTTP Basic authentication requires HTTPS; session login is a POST to `/redfish/v1/SessionService/Sessions`, HTTPS only, and suits pollers better.
- Eventing comes in two styles: push (the service POSTs to a subscriber) and server-sent events (SSE) over a long-lived GET.

The Redfish Resource and Schema Guide (DSP2046 2026.2) maps the cluster operations onto standard resources ([DMTF DSP2046 2026.2](#)):

Task	Redfish resource	Note
1. Firmware inventory and update	UpdateService (SimpleUpdate, MultipartHttpPushUri)	Inventory at /redfish/v1/UpdateService/FirmwareInventory
2. Power and thermal sensors	Sensor, PowerSubsystem, ThermalSubsystem	Older Power and Thermal schemas are deprecated
3. Metric streams	TelemetryService, MetricReportDefinition	Collects and reports metric data
4. OS install media	VirtualMedia InsertMedia action	Attaches remote media
5. Alerts	EventService subscriptions or SSE	Push or SSE

DGX B300 implements these concretely. Redfish "is enabled by default in the DGX B300 BMC and the BIOS," and NVIDIA documents platform sensors at /redfish/v1/Chassis/DGX/Sensors, GPU tray metrics at /redfish/v1/TelemetryService/MetricReportDefinitions/HGX_PlatformEnvironmentMetrics_0, and virtual media via /redfish/v1/Managers/Self/VirtualMedia/CD_1/Actions/VirtualMedia.InsertMedia (DGX B300 user guide: [Redfish APIs Support](#)). NVIDIA's examples use `curl -k`, which skips certificate validation; production tooling should instead enroll BMC certificates from an internal CA so that `-k` is never needed.

Redfish telemetry complements in-band GPU telemetry: the BMC sees power, temperatures and firmware state even when the host is down, while DCGM sees utilization, ECC and XID events. The [DCGM exporter monitoring and triage guide](#) covers the in-band side.

OpenBMC context

Outside DGX, BMC firmware may be a proprietary vendor stack or derived from OpenBMC, which describes itself as "an open source BMC firmware stack that will work across heterogeneous systems" ([OpenBMC](#)). Its web server, `bmcweb`, implements "a protocol compliant, DBus to Redfish translator" alongside KVM and serial-console websockets ([openbmc/bmcweb on GitHub](#)). Ask the vendor which stack it ships; it sets patch cadence.

BMC threat model and segmentation

NSA and CISA note that BMC firmware is "highly privileged, executes outside the scope of operating system (OS) controls," and that a compromised BMC can be used to disable TPM or UEFI Secure Boot and manipulate attached storage; EDR and TPM attestation are described as "ineffective at mitigating a compromised BMC" (NSA/CISA: [Harden Baseboard Management Controllers](#)). The sheet (June 2023, version 1.0) recommends protecting credentials, VLAN separation, hardened configurations, monthly or quarterly update checks, integrity monitoring, periodic firmware scanning, and "do not ignore BMCs."

Mapped onto a GPU cluster, the controls look like this:

1. **No routable exposure.** No default route to the internet; if BMCs must fetch updates, permit only update traffic during the window, as NSA/CISA suggest.
2. **Allow-listed sources.** ACLs on the OOB segment admit only the provisioning servers, the telemetry collector, the fabric manager and a hardened jump host. Tenant and user networks never reach it.
3. **Unique credentials.** Replace factory credentials before first connection, use unique vaulted per-node passwords, and rotate them on schedule and after staff changes.
4. **Protocol minimization.** Disable IPMI-over-LAN, UPnP and SLP where tooling allows; restrict SNMP to read-only from monitoring hosts.
5. **DPU BMCs count.** BlueField-3 BMCs are BMCs: same credential policy, same segment, same patch cadence.

Firmware resiliency

NIST SP 800-193, *Platform Firmware Resiliency Guidelines* (final, May 2018), is organized around protecting firmware against unauthorized changes, detecting unauthorized changes and recovering "rapidly and securely" ([NIST SP 800-193](#)). The full text lists the BMC as a platform component and assigns these functions to logical roots of trust: a Root of Trust for Update (RTU) "responsible for authenticating firmware updates," a Root of Trust for Detection (RTD) for corruption detection, and a Root of Trust for Recovery (RTRec) for restoring firmware ([NIST SP 800-193 PDF](#)). Its authenticated update mechanism "employs digital signatures."

Operationally, this means procurement questions and update policy:

- Ask the vendor what acts as RTU, RTD and RTRec for BMC and BIOS images, and whether a failed BMC flash recovers automatically. The DGX B300's BMC sits on "a new, replaceable DC-SCM" module ([DGX B300 user guide](#)), which shortens physical recovery but does not replace an automatic one.
- Stage BMC, BIOS and DPU BMC updates as a canary rack first, then by failure domain, never the whole SU at once.
- Record the pre- and post-update firmware inventory from Redfish for every node, and alert on any version change outside a change window.

Provisioning flow

A repeatable bring-up sequence uses the OOB network for everything that happens before the host OS exists:

1. **Inventory as code.** Record serial number, rack position, and system and DPU BMC MACs in a version-controlled inventory.
2. **DHCP reservations by BMC MAC.** Map each BMC MAC to a fixed address and hostname. The DGX B300 guide documents static BMC addressing for networks without DHCP ([DGX B300 user guide: Using the BMC](#)), but reservations keep the source of truth in one place.
3. **Credential and certificate enrollment.** Rotate factory credentials and install internal-CA certificates over Redfish.
4. **Firmware baseline.** Compare `FirmwareInventory` with the pinned baseline and update via `UpdateService` until every node matches.

5. **OS install.** Boot via PXE on the in-band network or via Redfish VirtualMedia.InsertMedia; PXE usually suits mass installs, virtual media single-node recovery.
6. **Hand-off to acceptance.** Burn-in begins only after inventory, firmware and OOB reachability match the record.

DPU-based management

BlueField devices complicate the plane boundaries. NVIDIA's DOCA documentation (3.5.0) defines DPU mode, in which "the NIC resources and functionality are owned and controlled by the embedded Arm subsystem," and NIC mode, in which "the Arm cores of BlueField are inactive, and the device functions as an NVIDIA ConnectX network adapter" ([NVIDIA DOCA: BlueField Modes of Operation](#)). DGX B300 ships BlueField-3 in NIC mode by default ([DGX B300 user guide](#)). In NIC mode the card is an in-band data path plus a BMC on the OOB network. In DPU mode the Arm subsystem is another managed operating system on each card, with its own provisioning, patching and access control. Qualification questions for newer DPUs are covered in the [BlueField DPU qualification guide](#).

Commissioning checklist

The OOB plane should be accepted before the GPU nodes, and its checks belong in the [cluster acceptance testing runbook](#):

Check	Pass criterion
1. Reachability	Every system and DPU BMC in inventory answers Redfish on 443 from the automation host
2. Isolation	BMCs unreachable from tenant, user and in-band host networks; no internet egress
3. Credentials	No factory credentials accepted; each BMC has a unique vaulted password
4. Protocols	IPMI-over-LAN, UPnP and SLP disabled or blocked where tooling permits
5. Time	BMC clocks synchronized to internal NTP
6. Firmware	Redfish firmware inventory matches the pinned baseline on every node
7. Telemetry	Sensor and metric-report data arrive at the collector for every node
8. Failure drill: power	Powering off a node leaves its BMC reachable and reporting
9. Failure drill: switch	Pulling an OOB uplink keeps every BMC reachable via the second uplink
10. Failure drill: recovery	One node recovers from a deliberately aborted BMC or BIOS update per the vendor procedure

The failure drills prove the OOB network stays reachable when everything else is down. GPU Smith can help review an OOB and BMC design against these checks.

Sources

1. [DGX H100 User Guide: Introduction to the NVIDIA DGX H100 System](#) - NVIDIA, BMC port and protocols, ConnectX-7 storage/in-band cards, optional management interfaces (accessed 2026-10-28)
2. [DGX B200 User Guide: Introduction to the NVIDIA DGX B200 System](#) - NVIDIA, BMC port and protocols, BlueField-3 storage/in-band cards (accessed 2026-10-28)
3. [DGX B300 User Guide: Introduction to the NVIDIA DGX B300 System](#) - NVIDIA, BMC on DC-SCM, on-board 1 GbE host port, BlueField-3 NIC mode default (accessed 2026-10-28)
4. [DGX B300 User Guide: Using the Baseboard Management Controller](#) - NVIDIA, BMC firewall port table, static BMC addressing (accessed 2026-10-28)
5. [DGX B300 User Guide: Redfish APIs Support](#) - NVIDIA, Redfish enabled by default, firmware, telemetry and virtual media URIs (accessed 2026-10-28)
6. [DGX SuperPOD B300 Reference Architecture: Network Fabrics](#) - NVIDIA, OOB on SN2201, VXLAN roll-up, IPMI traffic (accessed 2026-10-28)
7. [DGX SuperPOD B300 Reference Architecture: Key Components](#) - NVIDIA, SN2201 port configuration, 64-system SU (accessed 2026-10-28)
8. [DGX SuperPOD B300 XDR Reference Architecture: Network Fabrics](#) - NVIDIA, network enumeration, BlueField-3 BMCs on OOB, fabric management over OOB (accessed 2026-10-28)
9. [Redfish](#) - DMTF, current Redfish release listing (accessed 2026-10-28)
10. [DSP0266 Redfish Specification 1.25.0](#) - DMTF, RESTful/JSON/OData definition, service root, TLS, authentication, eventing (accessed 2026-10-28)
11. [DSP2046 Redfish Resource and Schema Guide 2026.2](#) - DMTF, UpdateService, TelemetryService, VirtualMedia, deprecated Power/Thermal (accessed 2026-10-28)
12. [Harden Baseboard Management Controllers](#) - NSA and CISA, BMC threat model and eight recommended actions (accessed 2026-10-28)
13. [CVE-2013-4786](#) - NIST NVD, IPMI 2.0 RAKP password-hash disclosure (accessed 2026-10-28)
14. [NVIDIA BlueField BMC Software v26.07: BlueField Management](#) - NVIDIA, last updated September 1, 2026, 1GbE OOB port, Redfish-only recommendation for new deployments (accessed 2026-10-28)
15. [OpenBMC](#) - OpenBMC project, self-description (accessed 2026-10-28)
16. [openbmc/bmcweb](#) - OpenBMC on GitHub, Redfish translator in bmcweb (accessed 2026-10-28)
17. [NIST SP 800-193: Platform Firmware Resiliency Guidelines](#) - NIST, publication status and date, protect/detect/recover (accessed 2026-10-28)
18. [NIST SP 800-193 \(PDF\)](#) - NIST, roots of trust for update, detection and recovery, BMC scope, authenticated updates (accessed 2026-10-28)
19. [BlueField Modes of Operation](#) - NVIDIA DOCA 3.5.0, DPU mode, NIC mode, zero-trust mode (accessed 2026-10-28)

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. <https://docs.nvidia.com/dgx-superpod/reference-architecture/scalable-infrastructure-b300-xdr/latest/network-fabrics.html>
2. https://media.defense.gov/2023/Jun/14/2003241405/-1/-1/0/CSI_HARDEN_BMCS.PDF

3. <https://docs.nvidia.com/dgx/dgxm100-user-guide/introduction-to-dgxm100.html>
4. <https://docs.nvidia.com/dgx/dgxm200-user-guide/introduction-to-dgxm200.html>
5. <https://docs.nvidia.com/dgx/dgxm300-user-guide/introduction-to-dgxm300.html>
6. <https://docs.nvidia.com/dgx-superpod/reference-architecture/scalable-infrastructure-b300/latest/dgx-superpod-components.html>
7. <https://docs.nvidia.com/dgx-superpod/reference-architecture/scalable-infrastructure-b300/latest/network-fabrics.html>
8. <https://docs.nvidia.com/dgx/dgxm300-user-guide/bmc.html>
9. <https://nvd.nist.gov/vuln/detail/CVE-2013-4786>
10. <https://networking-docs.nvidia.com/bluefieldbmc/26.07/bluefield-management>
11. https://www.dmtf.org/sites/default/files/standards/documents/DSP0266_1.25.0.pdf
12. <https://www.dmtf.org/standards/redfish>
13. https://www.dmtf.org/sites/default/files/standards/documents/DSP2046_2026.2.pdf
14. <https://docs.nvidia.com/dgx/dgxm300-user-guide/redfish-api-supp.html>
15. <https://openbmc.org>
16. <https://github.com/openbmc/bmcweb>
17. <https://csrc.nist.gov/pubs/sp/800/193/final>
18. <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-193.pdf>
19. <https://networking-docs.nvidia.com/doca/sdk/bluefield-modes-of-operation>

THE PUBLISHER

About GPU Smith

GPU Smith provides independent engineering and research for private AI infrastructure. We help technical buyers and operators reason about GPU workload sizing, hardware selection, procurement, deployment and inference operations, from the compute system to the networking, power and cooling requirements around it.

Start with the workload

Useful infrastructure decisions begin with the models, data, throughput, latency and operating constraints a team actually has. GPU Smith helps connect those requirements with system design and deployment choices. The goal is a privately controlled AI environment whose capacity and operational demands are understood before a purchasing decision.

Hardware and supplier research

Our public reference library covers GPUs, complete systems and networking components. The server-vendor directory and manufacturing research help teams investigate suppliers, compare options and follow sources behind technical claims. We distinguish manufacturer specifications from measured benchmarks and advertised capabilities from tested configurations. Reference pages are research resources, not a live inventory listing or a binding equipment quote.

Deployment and operations

GPU Smith's engineering scope includes infrastructure integration, networking, power, cooling and the practical operation of inference workloads. Our published articles explain the tradeoffs behind deployment, procurement and ongoing operations so teams can ask better questions and document decisions.

Work with GPU Smith

Explore the [hardware library](#), [GPU references](#), [networking references](#) and [vendor research](#). [Contact GPU Smith](#) with your workloads and deployment constraints to discuss a project and confirm current scope, pricing and availability.

Inclusion of a manufacturer or product in our research does not imply a partnership, certification or endorsement.

Website: <https://gpusmith.com>

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.