Deloitte.

Deloitte Center *for*
Higher Education Excellence™

Network Engineering Trends at Research Universities

Table of contents

Overview and context

Research Institutions are at the forefront of developing new ideas and practices. However, they may face numerous challenges and constraints as they consistently push the limits of today's facilities. This has led to an unprecedented demand on the network infrastructure of these institutions to provide enhanced capabilities, speed, and security. The challenges they may face include grappling with an exponential surge in cloud adoption, evolving student usage patterns, escalating requirements for high-performance computing, an influx of IoT devices, emergence of generative artificial intelligence (GenAI), and the staggering data volumes generated by research tools and applications.

The multitude of distinct and demanding use cases within research universities could present operational and engineering challenges that transcend the capabilities of legacy architectures and conventional processes. As traditional network designs are pushed to their limits, attracting and retaining top-tier talent, such as network architects and engineers, has become an increasingly daunting challenge.

In collaboration with the University of California Berkeley, Deloitte orchestrated a series of workshops that involved participation from numerous leading R1 institutions. These workshops served as a forum to analyze pain points and exchange insights regarding solutions to evolving challenges. Drawing upon Deloitte's experience within the Higher Education sector, coupled with the collective perspectives obtained from these workshops, this article serves as a summary of prevailing networking engineering trends. It seeks to address five pivotal questions that loom in the minds of technology leaders within Higher Education:

- What network challenges are educational institutions confronted with today?
- How is the network poised for transformation in the immediate future?
- What is the long-term strategy to address research needs over the next five years?
- What implications do these changes hold for network consumers?
- What is the strategic path forward from this juncture?

In this article, Deloitte aims to shed light on the evolving strategy of network architecture within research universities, offering insights that illuminate the path forward.

What challenges are institutions facing?

Any modicum of friction perceived by consumers of the network can potentially damage an institution's reputation, not just within the campus, but online as well. Network outages will likely be broadcasted or posted across various social media platforms. Long-standing annoyances like cumbersome network access, inconsistent wireless connectivity, or lagging customer support may turn into commiserating threads on a top post in [insert your university name here] sub-Reddit. Moreover, reputation plays a critical role in recruiting and retaining top research talent. The network is an essential service, and the longer it goes neglected, the greater the potential risk an institution takes in supporting the needs of students, faculty, researchers, and staff.

Some of the network challenges that are top-of-mind for Higher Education CIOs include:

Funding constraints

Funding models vary, but often include a combination of funding from campus IT, research grants, chargebacks for select services, and planned refresh cycles, which result in **difficulty quantifying derived value** from multiple investment sources. **Funding is limited**, often with the network being perceived more as a utility rather than a service.

Talent retention

Retaining talent has been a challenge for many institutions, with the majority **regularly short on staff. Salary expectations are high**, driven by inflation and higher compensation packages offered by the private sector. Universities, like many companies, are being asked to do more with less, which can negatively impact network support and maintenance.

The continued acceptance of remote work-from-home positions in a post-COVID world, driven in part by The Great Resignation, has created hiring competition where it did not exist previously. Schools are being forced to hire from outside their state due to this hiring competition, which presents challenges when staff are needed on-site.

Wi-Fi stability

Massive growth and increasing dependency on wireless are causing **scaling and resiliency concerns.** Access point **refreshes are difficult to budget for** as wireless density has increased. Aging Wi-Fi networks need transformation to meet today's demands, but migration resources are scarce.

Cellular penetration of buildings is weakening due to carriers turning off low frequencies and newer buildings being LEED certified, which can result in a poor end user experience.

Growing IoT footprint

Security concerns are rising due to a variety of operating systems and unpatched vulnerabilities on IoT devices, including default passwords that are difficult to track without profiling. Moreover, very few devices support 802.1x, requiring **unique controls for onboarding.**

Some **IoT devices need to be on the same radio channel,** resulting in custom wireless network configurations outside of IT standard practices.

Research with robotics and drones introduces new challenges that impact and overwhelm a trifecta of services—security, Wi-Fi, and reliability.

Security

New cyber security initiatives are driving architectural changes on the network, including zero trust methodologies, increased protection at the border, and **macro- and micro-segmentation across the campus.** These functional and operational shifts are disrupting traditional methods of providing connectivity without dramatically impacting the end-user experience.

Research units often function independently, on their own terms outside of campus IT standards, which could result in **technical debt and a serious risk** of vulnerabilities, attacks, and compromises. Additionally, researchers are working with an increasing number of entities that are government regulated, **requiring unique hardware and policy requirements,** posing new challenges.

Some researchers require unfettered access with robust and mature DMZ capabilities to perform their research. The institution should attempt to protect itself from research impacting organizational workflows, while simultaneously providing adequate capabilities for researchers to perform aggressive and occasionally invasive research. Similarly, research conducted in high-risk geographies well-known for hacking poses challenges with data sovereignty and requires unique traffic pattern analysis to adequately detect and protect against attacks.

Resiliency

Planning for disaster recovery has uncovered **challenges with the network delivering services** across multiple locations and multi-cloud environments, both privately and publicly. The cloud and its consumption pose challenges based on how **connectivity to the cloud** is designed with traffic patterns no longer East-West, but now North-South as well.

Fate-sharing and large failure domains are prevalent in campus environments and stretched data centers due to legacy architectures that simple campus refresh programs do not address.

Aging buildings on campus are often housing critical infrastructure. These buildings could have several constraints, including a risk of failure, inability to deliver sufficient power, limited space, or no longer being strategically located on campus, resulting in costly foundational redesigns. The cabling infrastructure itself within buildings and across the campus is aging as well, with some institutions forced to either retrofit the cabling infrastructure or adopt workarounds like Wi-Fi-only or resource shuffling.

Speed and capacity

Researchers are **projecting massive increases** in data generation, transfer, compute, and storage in the coming years, as well as multi-institutional collaboration, driven by advancements in AI/ML. Interactive instrumentation introduces needs for long distance low-latency. Developing and maintaining the **relationship with researchers** and research departments to stay on top of these projections has proved **challenging.**

Operations

The plans for improving the network are at times **out of touch with the needs of the network consumers**. This is not necessarily the sole fault of the IT departments—rather it lies in the communication

between IT and the stakeholders and constituents that represent significant consumers and user bases. The end-users on campus often feel disconnected from IT, with cumbersome interfaces that are perceived more as a barrier than an enabler. Users have expressed a desire for more self-service or seamless "coffee shop" type of experience.

Operations still heavily involve manual interaction. Automation holds promise to improve operations and streamline deployments, but **most institutions are spending more time keeping the lights** on with their minimal staff rather than investing the time and resources required to optimize.

Artificial Intelligence (AI)

Similar to any other application which leverages the network, **AI brings its own set of unique requirements**. AI is known for performing **real-time analysis** which means any **variance in delay and jitter will impact its effectiveness**. Speed and capacity will also impact the effectiveness of AI as enormous data models are synchronized and re-calculated over the network as new data is introduced. **Modular network architectures** will **enforce predictable traffic flows,** improving the AI experience.

As these AI applications become more common, **network observability** will become key to monitor the performance of these dynamic applications. Being able to **correlate delay/jitter with network flow data** will provide the visibility necessary to **identify bottlenecks and issues** when problems arise.

Many R1 institutions are in the process of implementing next-generation network designs to meet these needs and overcome challenges. But first, gaps must be identified in current state architectures, and then addressed in the future-state design.

A glimpse at today's typical campus network

The modern R1 institution supports a wide range of users who all have unique and demanding requirements for bandwidth, latency, security, and availability. Today's network architecture must account for these requirements stemming from various schools, students, administration, and researchers. Our team found common themes and strategies in the current network design of R1 institutions, providing a glimpse into today's typical campus architecture.

The campus boundary provides the WAN connections to the internet, other internal sites, and external institutions. The boundary typically consists of multiple routers, each peered with an Internet Service Provider (ISP). The boundary routers are generally geographically separated on campus to provide additional fault tolerance. For additional resiliency, two different service providers may be used, but limited availability or technical requirements may dictate that a single service provider is used for all peers.

As the internet edge is also the first point of ingress into a campus from unsecured external networks, **a boundary security stack is typically present** behind the boundary routers, but not always. This security stack may include next-generation firewalls (NGFW), taps, IDS/IPS, proxies, and web application firewalls (WAF). There are typically requirements within an R1 institution for certain networks to be unprotected and exposed to the internet, such as cyber research networks or honeypots. An additional consideration in the boundary security strategy is throughput limitations on many security platforms that could negatively impact large data transfers from research labs. To address this requirement, we found **most institutions implement some form of firewall bypass for specific traffic flows**, typically from research labs to compute clusters and data storage. (Figure 1)

Within the LAN, most institutions are running an iteration of the **classic 3-Tier model** for the campus design (Figure 2). Although there are nuances within each, they can be categorized into the following options: 3-Tier with collapsed core, 3-Tier with L3 at Distribution Layer, and 3-Tier with L3 at the Access Layer.

For routing, there is a mix of strategies and designs. **Many universities take a traditional approach** and utilize Interior Gateway Protocols (IGPs) throughout the campus, and then redistribute into Border Gateway Protocol (BGP) at the boundary. **Other universities have adopted a more service provider approach and utilize MPLS or Ethernet VPN (EVPN)** across their environment. Although more common in the data center, EVPN is also making its way into large campus environments to deliver similar scaling and security functionalities closer to the end-users.

Within the data center, most institutions are running Leaf-Spine architectures, with a **trend towards BGP EVPN or a controller-based design** (Figure 3). This is not only providing resiliency and scale, but also allowing for unique segmentation capabilities and Layer-2 extensions to solve complex network challenges such as research computing and disaster recovery.

Figure 1

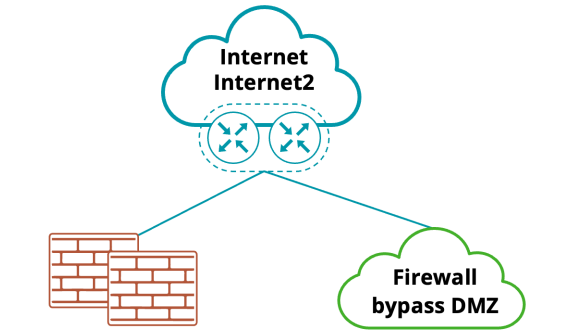

- All non-research traffic
- Research with unique data classification or compliance requirements

- Research requiring "unfettered" access
- High-speed, multi-institutional research

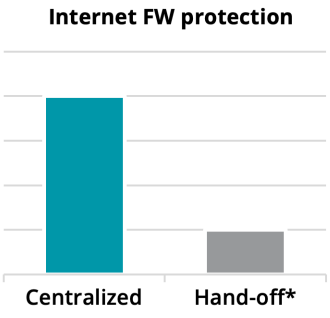

* Provides an hand-off without any centrally managed internet firewall services

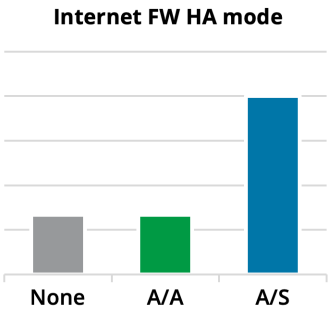

None = No default firewalls at Internet edge

A/A = Active/active

A/S = Active/standby or active/passive

Figure 2

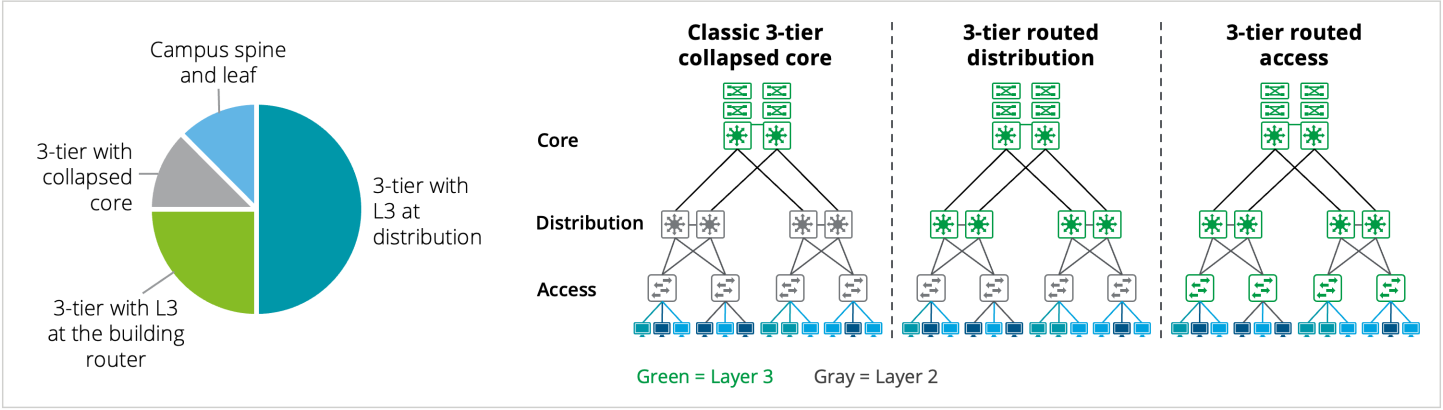

Figure 3

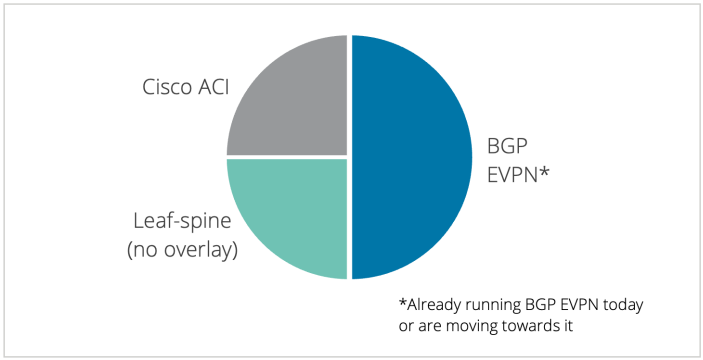

With research universities, there is often demand for higher connection speeds and bandwidth requirements, anticipated to increase significantly due to a combination of growth in wireless connectivity and research data flows across campus. **Most universities are abandoning 40G** links at the core and within the data centers for 100G connections due to a minimal cost difference, with 400G on the radar if it is not already implemented. At the access layer, edge connections are generally standardized at 1G with 10G availability in most locations. Some research labs have requirements for specific tools to connect at speeds higher than 10G, with certain labs expressing a desire for end-to-end 100G connections from their tools generating the data to the compute and storage infrastructure.

For end user connectivity, many institutions have a "**wireless first**" initiative with the intent to reduce wired connections and cabling requirements throughout campus. We found most R1 universities have over 10,000+ wireless access points (WAPs) deployed. Most wireless deployments leveraged an on-premises controller-based architecture with multiple sets of wireless controllers for redundancy and scaling purposes, utilizing a minimal number of service-set identifiers (SSIDs) to avoid overhead and reduce congestion issues.

With **massive amounts of data being generated and sent across R1 networks**, the data center strategy is a focal point of many R1 institutions. Due to the latency requirements and size of data generated by research labs, on-premises data centers for compute and storage were the primary options for most research labs. Some R1 institutions leverage third party data centers or off-campus colocation facilities to host production (non-research) computing workloads and storage. Likewise, many universities are struggling with floor space and power/cooling challenges driven by increased compute demands from researchers.

Universities are already **significantly investing or planning to invest in public cloud**. In addition to research labs, many schools, staff, and students leverage services from cloud providers such as AWS, Azure, and GCP, implying the **need for a multi-cloud strategy**. Several institutions are planning to minimize their on-premises data center footprint by leveraging cloud or colo-data centers for a majority of services, **trending towards a hybrid compute model** (Figure 4). However, most research use cases are not practical to fully operate in the cloud, usually due to cost, governance, or instrumentation gaps.

Network segmentation initiatives continue to complicate the end-user experience in campus environments. On one hand, network security is an essential requirement to protect data and the assets within the institution. On the other hand, the consumers of network services often find themselves entangled in the complexity of accessing the network or their data. Not everyone is implementing campus-level network segmentation today, but we are observing **an increasing trend where segmentation becomes a vital element of R1 network designs**. Some institutions are segmenting based on security use cases and requirements driven by their customers (e.g., school or department), and others are segmenting based on network role or function driven by a campus IT security governance model (e.g., staff, student, researcher, etc.). In general, the goal for most is an "opt-out" model for segmentation, meaning that justification and approval are required to opt out of the firewall controls in place.

Figure 4

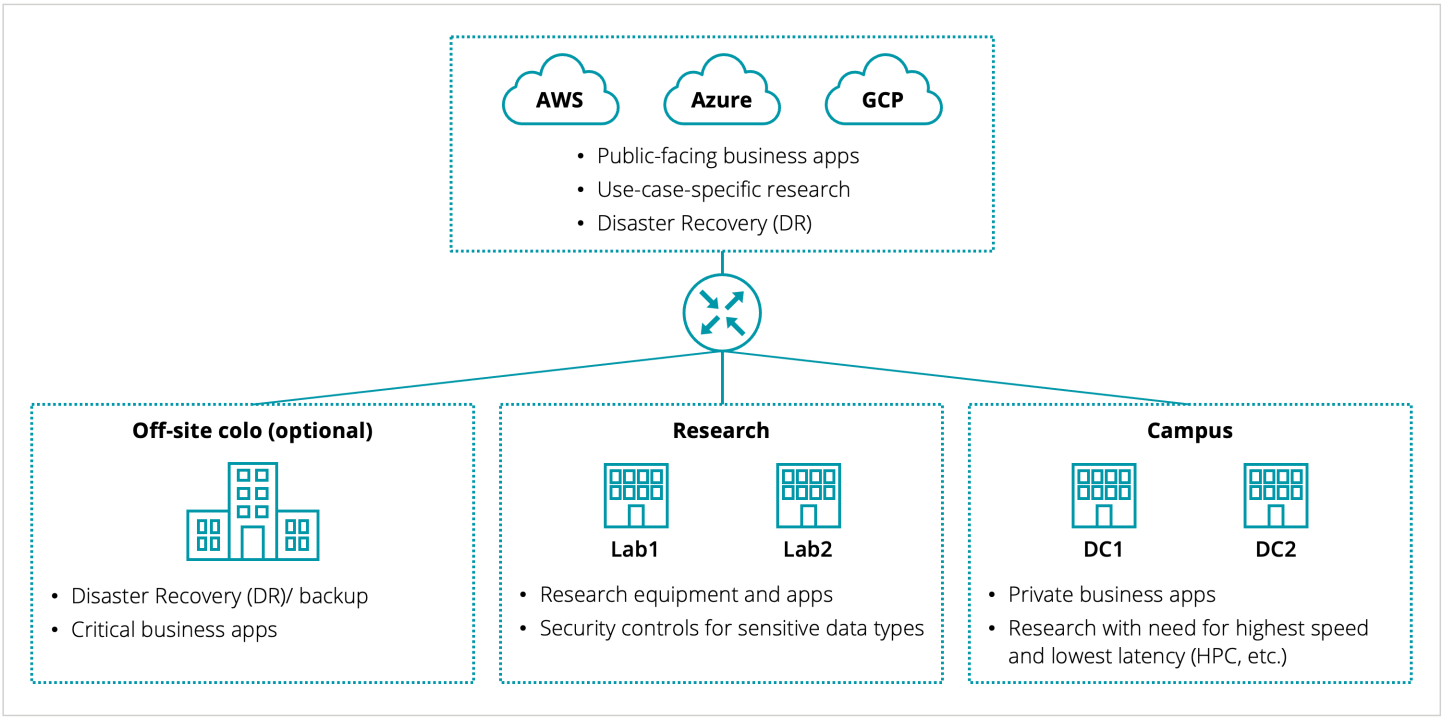

To support segmentation, many institutions are implementing some level of Network Access Control (NAC) (Figure 5). NAC is primarily used on wireless networks, generally employing a MAC authentication approach with device registration. Some institutions are implementing wired NAC, while others are in pilot phases or have select areas implementing wired NAC—primarily leveraging MAC authentication. New security concerns are prompting the adoption of zero trust initiatives that supplement macro-segmentation approaches, such as VRF and Firewalls. The usability and effectiveness of these solutions vary from barely effective to mature and robust with custom user management portals.

Most **R1 institutions are competing for funding and talent** to maintain a competitive advantage and deliver exceptional service to their customers. In many organizations, even "keeping the lights on" is challenging, with aging hardware and software that is rarely refreshed and often not delivering any new capabilities to the end-users. It is often not the lack of engineering talent, but rather the lack of organizational support that has not prioritized network services due to more pressing campus needs. The research reputation is the driving force here, and recruiting the best researchers requires cutting-edge architectures that are agile and ready for rapid transformation.

Figure 5

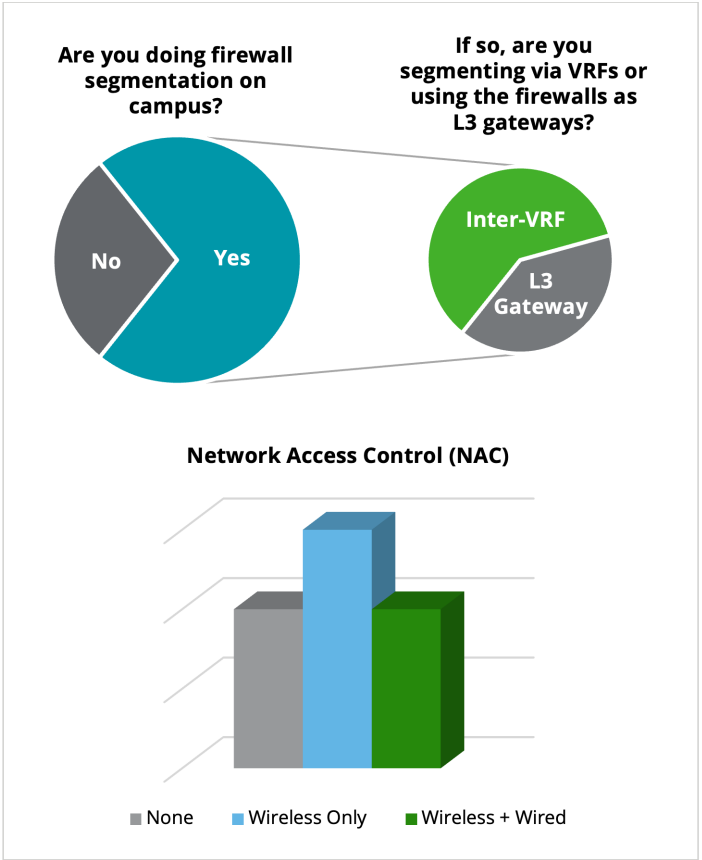

Network transformation to support the future of research

Research labs are generating more data than ever before, and projections indicate an exponential increase in the coming years. It is not only researchers, but also many schools that are utilizing tools and applications that demand significant network resources. One of the primary areas of network design and strategy that should be prioritized to address this demand is the enhancement of network connection speeds and throughput capabilities.

The demand for higher speeds is driven by increases in the bandwidth capabilities of wireless networks and a greater number of devices using wireless, as well as researchers requiring higher speeds for transferring massive data sets between on- and off-campus locations. A connectivity speed of **10G** is becoming the minimum standard for buildings, with **100G emerging** as the de-facto link speed between network equipment. The adoption of **400G is beginning to establish its place** in the core of networks, and most R1 institutions are preparing for the **future implementation of 800G+** in their long-term roadmaps, particularly those institutions with significant AI aspirations.

With many use cases for external data transfers through the Science DMZ via Data Transfer Nodes (DTNs), WAN circuits are expected to move to higher speeds in the coming years. Although the requirement depends on an institution's external data transfer needs, circuits should be planned for speeds exceeding 100G in the near term. Factors related to significant AI expansion suggest that the demand for higher speeds will materialize more quickly than originally projected. Education Network Service Providers such as CENIC have just begun offering native 400G services to their customers, with these higher speeds becoming more commonplace in the next 3–5 years. (Figure 6)

Research labs handle massive datasets, and some have low latency requirements for technologies such as real-time data rendering. With that in mind, researchers typically prefer to utilize on-premises **compute resources** for data analysis rather than cloud or off-site third-party data centers. Due to perceived limitations of the network environment, many labs use their own grant funding to build Do-It-Yourself (DIY) compute and storage solutions within their department or building. This approach is less than ideal as it can result in a significant increase in overall costs, especially for power and cooling in non-data center facilities. The goal of most R1 institutions is to provide GPU and CPU compute clusters in a centralized model and reduce the amount of expensive and demanding infrastructure within each lab.

Although compute and data analysis typically occur on-premises, long-term storage solutions from external third parties or off-site data centers are often utilized. The Science DMZ is purposely built to handle long-distance data transfers effectively across high-latency paths. The Science DMZ and DTNs can be leveraged to efficiently transfer large amounts of data from on-premises

Figure 6

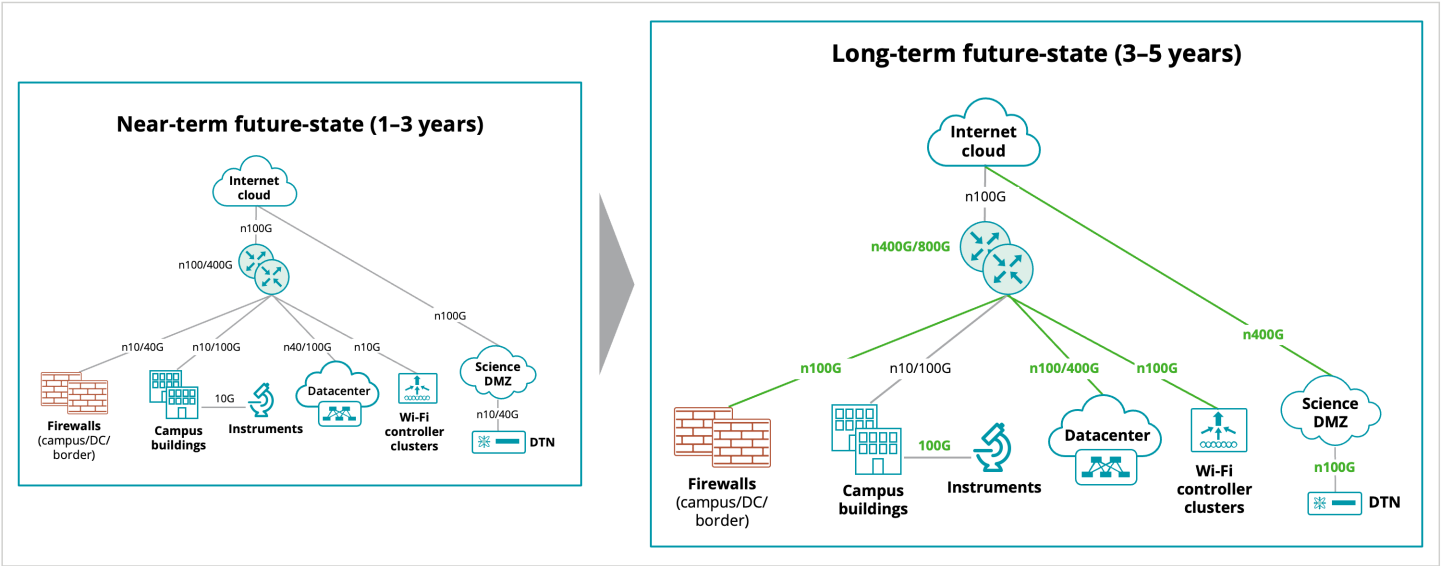

Figure 7

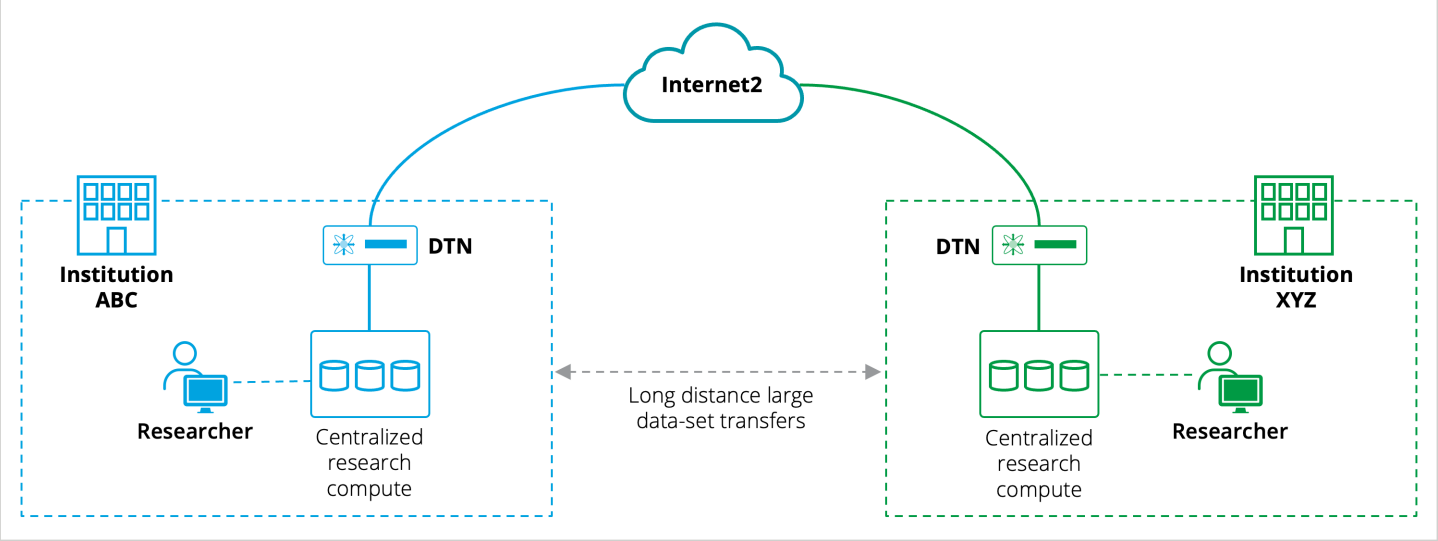

Figure 8

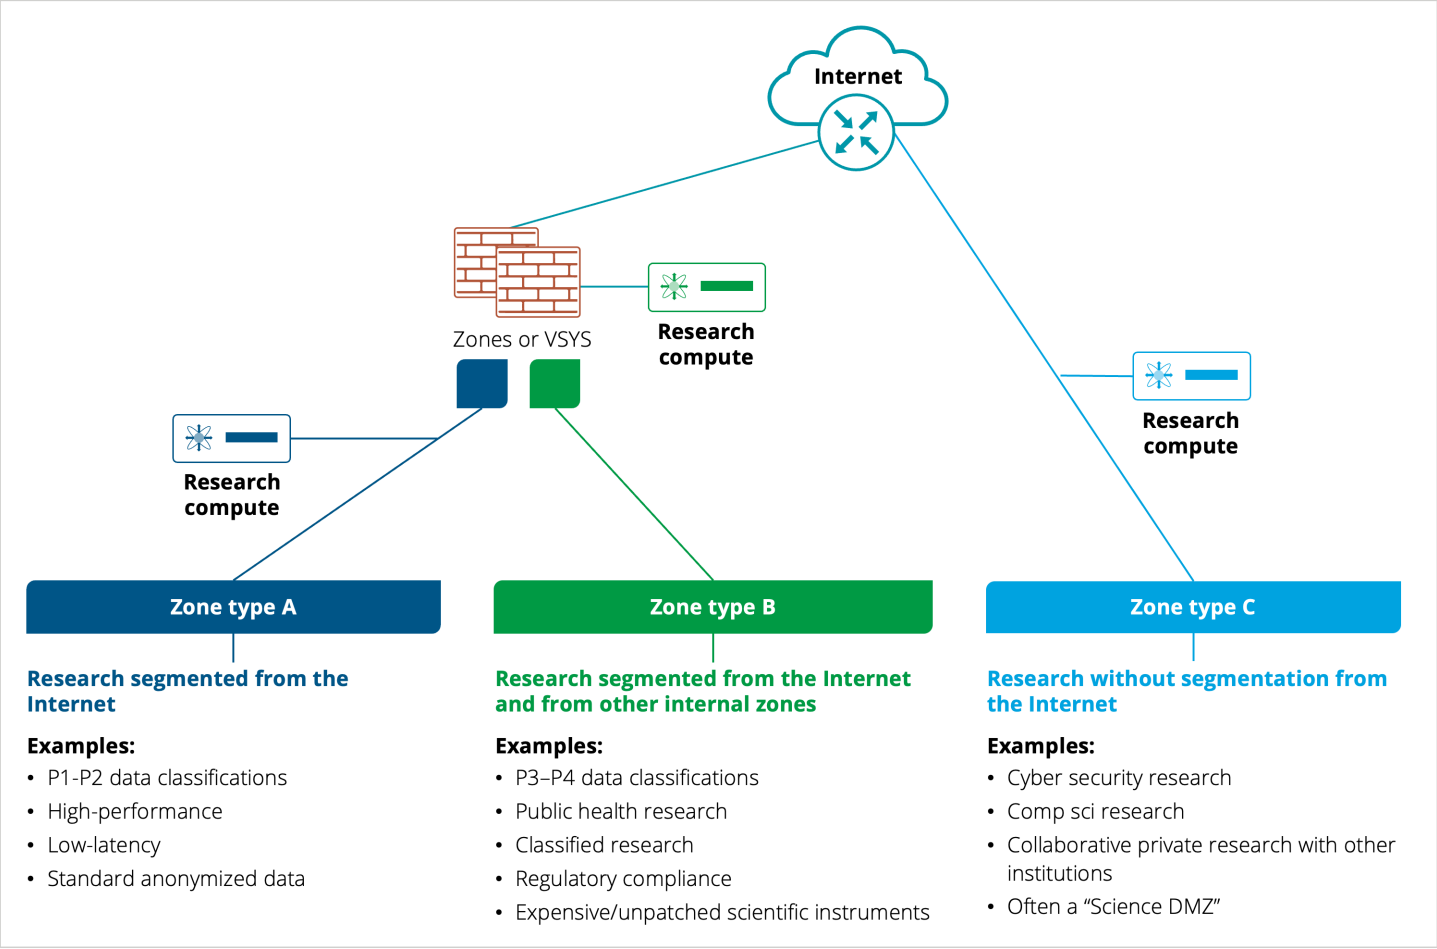

infrastructure to the cloud, off-site storage options, and other R1 institutions. The primary use case is for transferring data into and out of campus for research collaboration purposes, leveraging DTNs attached to the storage infrastructure with optimized file transfer capabilities. (Figure 7)

While network performance and usability are often primary concerns for many researchers, **security** remains an integral part of the design. Most R1 institutions are in the process of designing, implementing, or enhancing segmentation capabilities. Solutions under consideration or currently deployed include **network fabrics** that provide isolation, **Software Defined Networks (SDNs)** that address **micro-segmentation needs**, integrations between **NAC** systems and **firewalls** for tag-based or user-id based policy enforcement, and significant remodeling of network segments protected by Next-Generation firewalls (NGFWs).

Many research labs handle P1 and P2 datasets originating from tools like MRI machines (medical imaging) and genomic data from microscopes. While bypassing the firewall could be considered for large data transfers, institutions are establishing clear security boundaries to protect sensitive data where necessary. The firewall strategy commonly complies with the data protection requirements of the datasets generated and typically aligns with one of the three high-level designs: Research segmented from the internet, Research segmented from the internet and other internal zones, and Research without segmentation from the internet. (Figure 8)

Researchers often wear several hats within a university organization. They may be students, staff, dedicated researchers, or a combination of roles. As such, they are integrated into the user population, and segmenting all research users at the network level can be challenging, further complicating traditional layer-4 firewall policies. A zero-trust approach, utilizing a NAC solution paired with identity-based security policies, can reduce or eliminate some of the network segmentation requirements for researchers and simplify the firewall rule sets for R1 institutions. Many R1 institutions are actively investigating or already investing in solutions and architectures that leverage identifiable data beyond IP and MAC addresses to provide dynamic and robust network security solutions that are seamless for the end-users. This provides an opportunity to improve the security posture of the environment while concurrently enhancing the usability and flexibility of the network for researchers. (Figure 9)

Figure 9

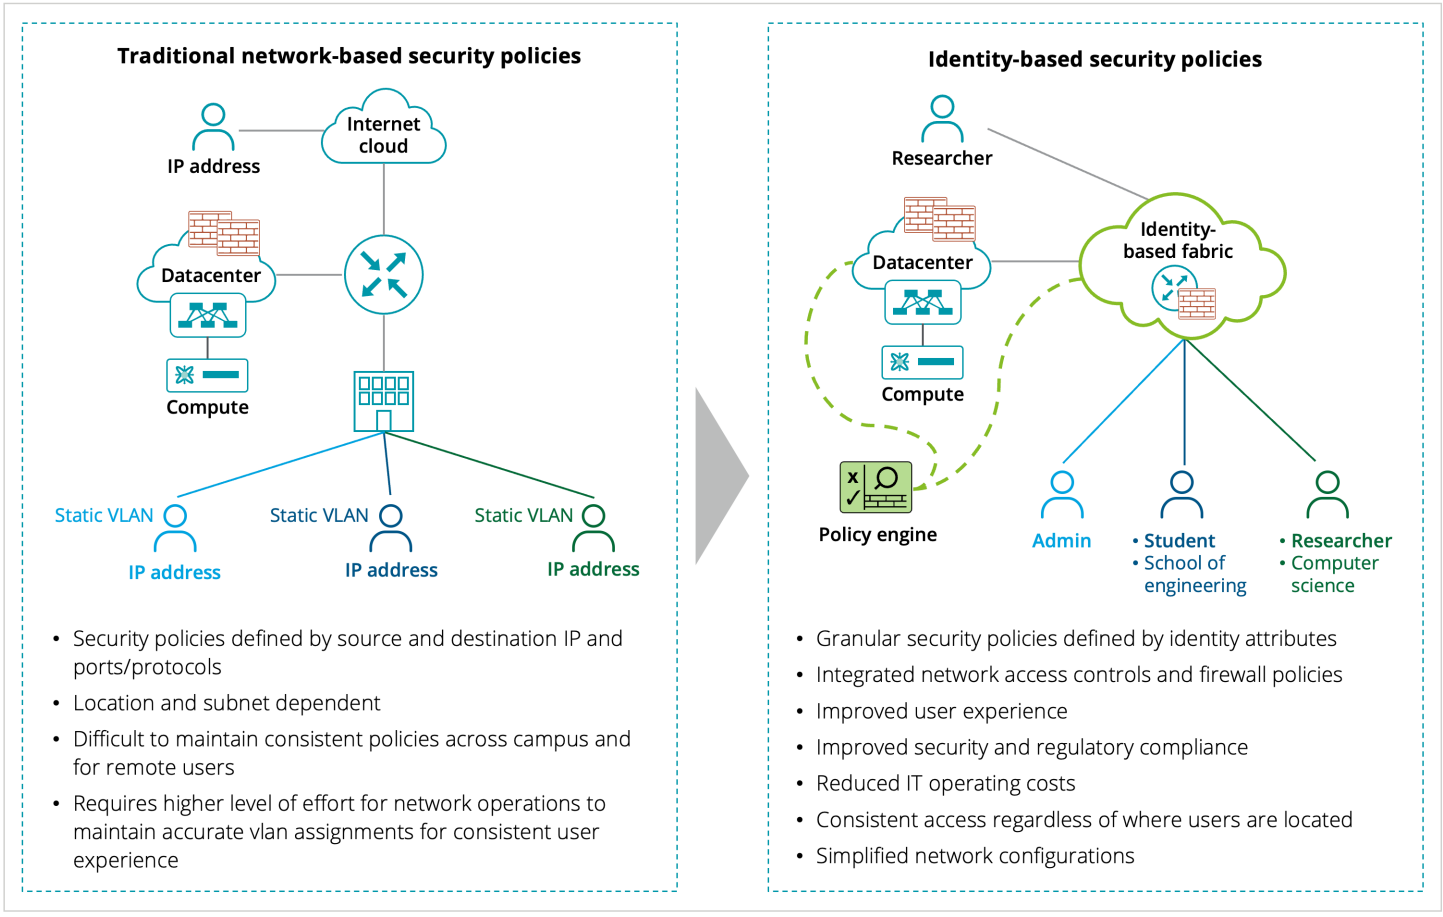

- Security policies defined by source and destination IP and ports/protocols
- Location and subnet dependent
- Difficult to maintain consistent policies across campus and for remote users
- Requires higher level of effort for network operations to maintain accurate vlan assignments for consistent user experience

- Granular security policies defined by identity attributes
- Integrated network access controls and firewall policies
- Improved user experience
- Improved security and regulatory compliance
- Reduced IT operating costs
- Consistent access regardless of where users are located
- Simplified network configurations

Observability is a critical aspect of every network enterprise, and this is particularly true when addressing the data size and latency requirements of researchers. Most network teams utilize a combination of **performance monitoring** tools to track metrics such as link utilization, network device or link status, and syslog monitoring. However, many organizations lack the capabilities to validate true end-to-end network performance. Performance monitoring captures actual performance data from the production network. This can be achieved through **telemetry** collected from various sources to provide a comprehensive view of the network experience by evaluating transactions, NetFlow, and device statistics. Additionally, employing toolsets that **generate synthetic transactions** to simulate user transfers through the network and record actual throughput performance provides tremendous value to research labs.

One of the primary drivers of network automation at R1 institutions is the need to do more with less. New technologies are being evaluated for their **automation capabilities** and **built-in analytics** or telemetry features, with the aim of **streamlining operations** as much as possible and reducing the expertise required by junior engineers for support. SDN architectures are being evaluated for their **orchestration** capabilities and robust **integration** with various systems such as ticketing and chat. The trend for many is to purchase commercial solutions over developing them in-house.

Having the right observability tools within an environment is imperative, but their value is limited by how effectively the information they generate is viewed and reported. Institutions that have implemented custom dashboards to quickly identify anomalies, trends, and performance metrics can provide a clearer and more concise narrative for their critical research stakeholders. (Figure 10)

For researchers to be successful, they require a **research organization and support model** that takes into account their unique technical requirements, funding constraints, and stringent timelines. Research IT and campus IT must maintain a cohesive and collaborative relationship. Research IT works closely with the research labs, helping to identify and communicate specific requirements to the broader university community. Without an integrated support model, understanding and engineering solutions for the exact and specific requirements of each lab becomes challenging, frequently leading research labs to seek their own solutions outside of centrally offered services.

Research labs evolve based on grants, which makes budgeting difficult to forecast as Research grants typically provide limited funding, requiring cost-effective solutions. These grants also often come with aggressive timelines for completing the associated research, requiring connectivity solutions that can be delivered quickly and support agile and responsive solutions. Some organizations are addressing these challenges by implementing **a support model that enhances collaboration between research IT and the labs for requirements gathering, procurement, and ongoing support**.

To reduce DIY solutions in research labs and improve supportability and standardization, some research institutions are leveraging a central procurement and management solution through campus IT. In this model, researchers collaborate directly with IT teams to define their needs, evaluate standardized solution offerings, and procure through campus IT (Figure 11). The network infrastructure is then deployed, managed, and monitored by campus IT, potentially reducing both the overall cost and the level of effort required to establish lab network connectivity.

Figure 10

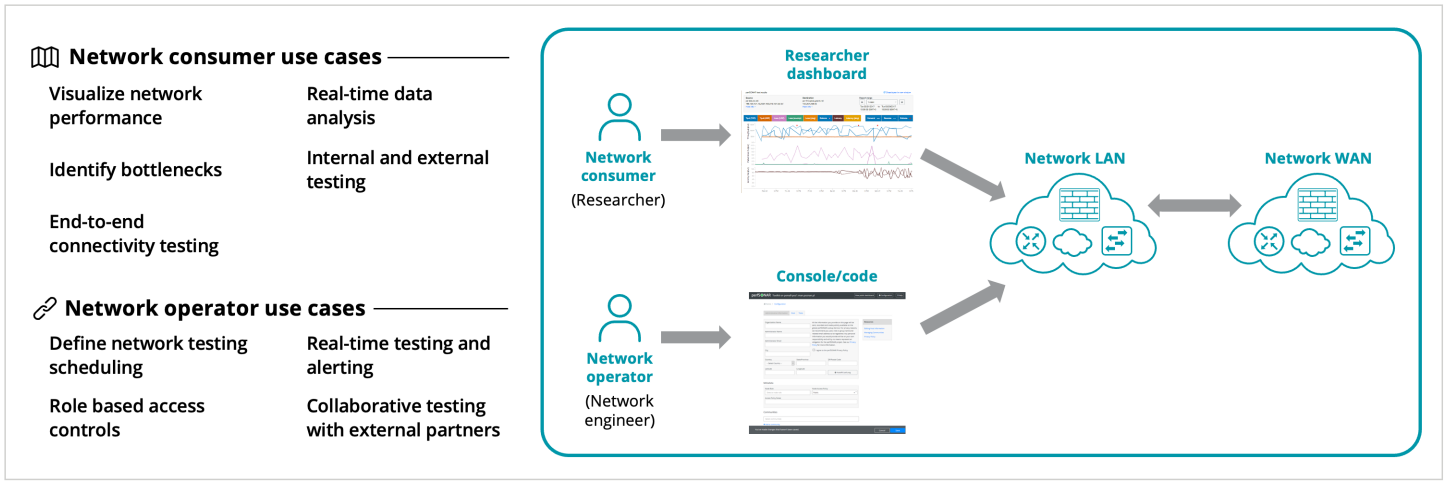

Keeping up with this shift in capabilities requires the appropriate talent and operational mindset. Institutions are adopting a hybrid (remote + in-office) model to **attract and retain talent**, which allows for a wider geographical reach in recruiting. Other institutions strategically prefer an insourcing model, bringing in the right talent for specific projects or initiatives.

With the right skills in-house, institutions are increasingly focused on automation as one of their foundational principles for future-state network initiatives. Ad-hoc scripts are insufficient for the scale and agility required in research universities, where changes are frequent, and the stakes are high. Network emulation and digital twin technologies are providing engineers with confidence in network design validation to meet unique research network requirements. Automation platforms, Infrastructure-as-Code, Generative AI and AIOps are at the forefront of strategic planning for the next 2–3 years.

Figure 11

IT support model for research

Decentralized

- Increased flexibility and freedom for researchers
- Potential cost savings for hardware and software
- Inconsistent technology and designs
- Difficult to integrate into campus architecture
- Lack of adherence to security best practice and policy
- Increased dependency of research lab staff for IT configurations and support

Centralized

- Agile and responsive support for research IT needs
- Improved security posture
- Decreased dependency on research staff
- Infrastructure is easily repeatable and expandable
- Increased resiliency and performance in research IT infrastructure
- End-to-end observability of research infrastructure and traffic flows

What does this all mean for the consumers of the network?

Future-state designs showcase a promising future for network consumers, successfully overcoming many of the challenges faced by research universities today.

With new funding models established to foster and sustain relationships between IT, Academics, and Research—each party is **aware of and invested** in the success of delivering an exceptional experience and capabilities. Moreover, optimized support models facilitate **expedited timelines** for grants, including agile and responsive solutions that enable quick setup while adhering to organizational standards and support structures.

Consumers can benefit from **higher speeds and increased capacity**, which eliminate bottlenecks and avoid roadblocks. Edge connectivity from 1G to 100G is now available, providing a comprehensive range of speeds tailored to end-user needs. The implementation of 400G is underway, with 800G+ on the horizon for backbone connectivity and specialized research use cases.

Wi-Fi capabilities continue to advance, leveraging expanded radio usage and massive scaling via cloud-delivered architectures that provide seamless roaming, even during periods of heavy mobility such as class transitions or events.

Frictionless network access is safeguarded by security policies—offering consumers a secure environment that is segmented appropriately based on data confidentiality yet remains discreet and invisible to those using the network—utilizing a combination of profiling, identity, NGFW-backed segmentation, and zero trust architectures.

Highly scalable architectures that transcend the limitations of traditional 3-Tier architectures will introduce new capabilities to consumers, such as "DMZ Anywhere" approach or identity-driven network policies leveraging automated overlays.

Resilient architectures with geo-redundancy offer peace of mind to network consumers by enabling deployment across multiple data centers or clouds, utilizing common tools and support structures.

The acceptance of multi-vendor architectures **addresses supply chain constraints and impacted timelines**, achieved in part by employing automated methods to operate networks agnostic of vendor platform, more seamlessly integrated with standards-based protocols.

Automation platforms will introduce highly sought-after **self-service capabilities** for consumers, featuring a combination of AI-driven interfaces and intelligent infrastructure and services management. End-users are empowered to self-register, deploy, troubleshoot, and even design, with minimal interaction with IT.

A high-level network architecture framework in support of research may look like the following (Figure 12):

Figure 12

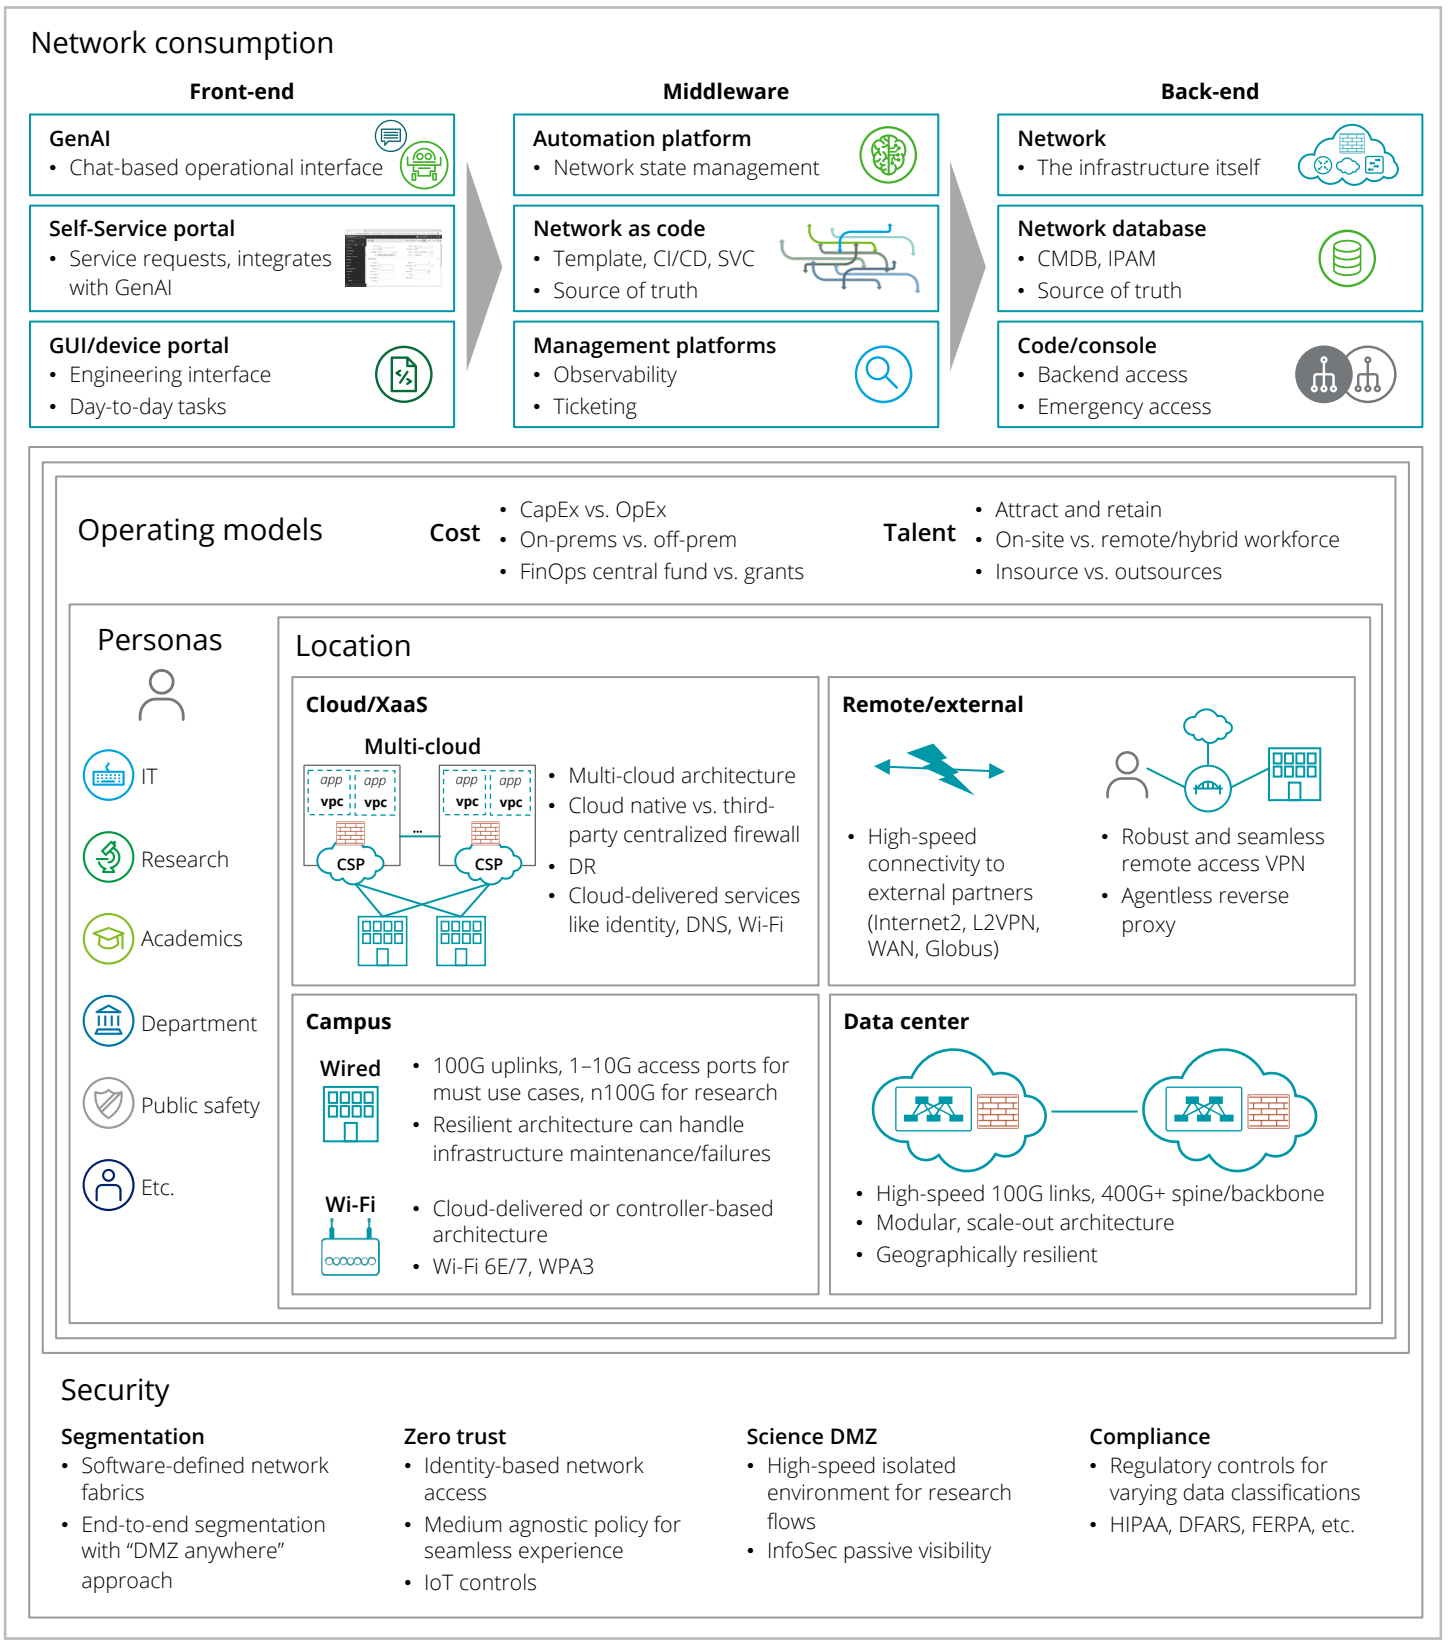
Network consumption
Front-end
GenAI
• Chat-based operational interface
Self-Service portal
• Service requests, integrates with GenAI
GUI/device portal
• Engineering interface
• Day-to-day tasks
Middleware
Automation platform
• Network state management
Network as code
• Template, CI/CD, SVC
• Source of truth
Management platforms
• Observability
• Ticketing
Back-end
Network
• The infrastructure itself
Network database
• CMDB, IPAM
• Source of truth
Code/console
• Backend access
• Emergency access
Operating models
Cost
• CapEx vs. OpEx
• On-prems vs. off-prem
• FinOps central fund vs. grants
Talent
• Attract and retain
• On-site vs. remote/hybrid workforce
• Insource vs. outsources
Personas
IT
Research
Academics
Department
Public safety
Etc.
Location
Cloud/XaaS
Multi-cloud
app
vpc
CSP
• Multi-cloud architecture
• Cloud native vs. third-party centralized firewall
• DR
• Cloud-delivered services like identity, DNS, Wi-Fi
Remote/external
• High-speed connectivity to external partners (Internet2, L2VPN, WAN, Globus)
• Robust and seamless remote access VPN
• Agentless reverse proxy
Campus
Wired
• 100G uplinks, 1–10G access ports for must use cases, n100G for research
• Resilient architecture can handle infrastructure maintenance/failures
Wi-Fi
• Cloud-delivered or controller-based architecture
• Wi-Fi 6E/7, WPA3
Data center
• High-speed 100G links, 400G+ spine/backbone
• Modular, scale-out architecture
• Geographically resilient
Security
Segmentation
• Software-defined network fabrics
• End-to-end segmentation with “DMZ anywhere” approach
Zero trust
• Identity-based network access
• Medium agnostic policy for seamless experience
• IoT controls
Science DMZ
• High-speed isolated environment for research flows
• InfoSec passive visibility
Compliance
• Regulatory controls for varying data classifications
• HIPAA, DFARS, FERPA, etc.

Where do we go from here?

Leading R1 research Institutions are at the forefront of transforming expectations in network connectivity services. Knowledge sharing and collaboration are paramount in this era of rapid technological evolution. As institutions recognize the impact and potential of emerging technologies such as AI, IoT, and immersive experiences, business strategies must evolve to embrace this new and disruptive paradigm, or risk being left behind. Cloud consumption models, along with billing insights, should provide researchers with a pathway to store and process data, independent of the underlying platform, with robust support structures to avoid the need for "home-grown" solutions. The network technology stack is evolving, characterized by increased security requirements, enhanced performance and reliability expectations, and a growing demand for a seamless experience irrespective of physical location. The journey is daunting, but with the appropriate strategy in place, securing a competitive advantage that attracts the best student and resea

The process typically begins with defining requirements and business objectives. We conduct workshops with stakeholders such as researchers and other significant network users from academics to understand use cases, pain points, and upcoming initiatives that may impact infrastructure and services. Deloitte assesses the current state environment and identifies gaps that must be addressed to satisfy the business objectives. We work with our clients to develop a strategic vision and roadmap, often partnering to create a next-generation network architecture.

At Deloitte, we can help guide our clients through this transformational journey.

To learn more about our services and capabilities (Figure 13), and how we can add value to your organization, please reach out to our Higher Education or Edge & Connectivity practice leaders or visit us on deloitte.com.

Figure 13

Capabilities

Differentiated outcome focused services and industry knowledge brings together network consulting capabilities to deliver accelerated customer value

Multi-Cloud Networking/ SD WAN	Network Security	Edge Mobility	Campus Wired & Wireless	Data Center Networking	SDN-Network Virtualization	Networking Sustainability/ Green	Network as a Service (Naas)
Design and implementation of large-scale modern network fabric with help of native tooling, third-party overlays, integrated security functions, or a combination of all	Securing assets across the landscape. Encompassing local on-premise assets, external cloud space, and mobile users. Next-generation security for today's Internet of Things.	Integrating E2E solutions across Network/Edge Design and Planning, Network/Edge Deployment, Operations Enablement, Assurance, Tooling & Automation (5G/Edge)	Design and implementation of campus wired and wireless solutions, empowering enterprise network modernization and deploying zero trust frameworks to deliver a unified user experience	Storing, processing, and moving your data and applications safely and anywhere in the world	Planning, design, implementation and operations for virtualized network functions and services	Climate-smart considerations in every decision, recommendations and actions across the portfolio of services, including Network Consulting	Device and incorporate cost optimization strategies to achieve operational or business goals leveraging NaaS or Managed Services that are at the optimum cost point and right sized
Network Operations Automation	Network Segmentation	Network/Edge Design and Planning	Wiring and Wi-Fi Approach & Design	Datacenter, Co-Lo, and CNF Optimization	Network Function Virtualization Design	Strategy & Roadmap (Sustainable Targets)	Cost Strategy
Firewall & WAN Optimization	Identity Management	Network/Edge Deployment	Wi-Fi Coverage Surveys	HA & DR Strategy	VNF Onboarding and Provisioning	Operations	Managed Services
Requirement Planning (SD-WAN Hardware & Software)	Secure Access Services Edge	Operations Optimization	Deployment, Placement & Management	Data Center Fabrics Implementation	Service Chaining	Finance	Asset Leasing
SD-WAN Integration with Cloud Providers	Security Orchestration and Telemetry	Network Assurance, Tooling & Automation	Site and cabling assessments	Network Design and Strategy	Next-generation Protection & Security	Reporting & Transparency	Operational planning
Managed SD-WAN Delivery		Infra Planning, Configuration, Mgmt.	Device onboarding and provisioning	Network Assurance, Tooling & Automation			

Thank you!

Deloitte extends our deepest gratitude to the University of California Berkeley for collaborating with us on the R1 Network Peer Institution workshops that ultimately led to the development of this publication. We appreciate your dedication and the valuable insight you provided that made this publication possible.

Contact us

Roy Mathew
Principal | US Higher Education Practice Lead
Deloitte Consulting LLP
rmathew@deloitte.com

Richard Johnson
Managing Director | Edge & Connectivity Practice Lead
Deloitte Consulting LLP
richarjohnson@deloitte.com

James Qua
Principal | Higher Education Cloud Engineering Practice Lead
Deloitte Consulting LLP
jqua@deloitte.com

David Varnum
Specialist Leader | Edge & Connectivity Practice
Deloitte Consulting LLP
dvarnum@deloitte.com

Authors:

- David Varnum, Chief Network Architect
- Bryce Koepke, Network Architect
- Stephen Occhiogrosso, Network Architect
- Richard Johnson, Edge & Connectivity Leader

Deloitte.