

# Differences between R&E Networks and the Commercial Internet

---

July 8, 2010

Research and Education networks in the US typically cite the same set of distinguishing features to separate themselves from the 'commercial' or 'commodity' Internet. R&E networks support advanced protocols like IP multicast and IPv6. They have more direct routing between endpoints, which reduces end-to-end latency, and their minimalistic design and ample capacity reduce jitter. That excess capacity, carefully maintained by a much more aggressive upgrade schedule, allows the use of experimental protocols that make unusual bandwidth demands or do not tolerate congestion. Connections to the network are restricted to members of the community, ensuring that the traffic carried is at least non-commercial, if not always strictly limited to research or education activities.

Those are only the surface differences, however. The design and operation of the R&E infrastructure is distinct as well, driven by the technical requirements for the provision of those special capabilities, and as a natural reaction to the different environment of the R&E world.

## Redundancy and outage handling

Customers have extremely low tolerance for failures of commercial Internet connections, even at home - and especially if their ISP is also their television and telephone provider. Although many single points of failure remain at the edges of most ISP networks, their designs attempt to minimize the effects of any failure, and to avoid outages entirely in the core of the network through designed-in redundancy and resiliency. Major customers can and do pay for redundant facilities, or purchase connectivity from multiple independent sources.

R&E customers typically treat their research connections as important, but not essential, and are much more tolerant of service outages because the commercial Internet offers a lower-performance but acceptable substitute during the outage. They usually have only a single connection to an R&E provider and no interest in a second one. As a result, R&E network operators invest in bandwidth capacity and advanced features rather than redundancy. Their support staff are centralized and limited in number, with on-call rotations or outsourced services taking the place of 24x7 manned network operations centers. Field services are provided by contractors, and equipment is maintained under vendor support contracts instead of spare parts inventory and redundant installations.

## Scaling

R&E networks optimize for few customers with large connections; commercial operators plan for many customers with smaller connections. An R&E customer typically connects directly to a backbone router, which minimizes latency and jitter, and makes the network simpler to administer by reducing the number of 'moving parts' in the design. Commercial networks must employ multi-level connectivity to meet their scaling demands. Home and small business customers connect to aggregation devices (appropriate to the technology - DSL, PON, or

cable), which connect to provider-edge routers; larger customers also connect there. Those routers connect to the backbone, and will frequently have multiple uplinks to diverse backbone sites, for redundancy and traffic tuning.

## Services

Commercial ISPs provide two categories of services that R&E networks characteristically do not: value-adds and combined business offerings (double and triple play services). A home customer expects value-added services: email boxes, DNS resolvers, address space assignment, a web portal with customer information. And the ISP wants to maximize revenue by cross-selling additional services: phone on broadband, broadband on cable TV, and so on. Many ISPs are aggressively targeting businesses as well as home users for these services. R&E networks focus on large customers, most of whom already provide those services for themselves and their internal users, and seldom offer value-adds or bundles.

## Transparency

The technical details of a commercial ISP network are carefully kept secrets, since knowledge of topology, equipment, configuration, connectivity and network utilization could be used against them by a competitor. The same secrecy applies when a customer tries to work with the ISP to resolve a problem, asks for information about maintenance work, or requests a post-mortem report after an outage. On the other hand, R&E networks, with their smaller customer base, strong hierarchy of customer relationships (backbone-regional-campus-user) and requirement to provide advanced services, have a history of providing much more open and responsive customer support services, including visibility into the network's operations and configuration that is unheard of in the commercial world.

## Network Reach

Though underserved areas remain across the country, 95% of Americans have access to high-speed Internet access. In addition to the scaling differences noted above, the reach of commercial ISP networks is vastly different from R&E networks. The two national backbones taken together have a total of eighteen backbone router nodes, but with overlap they cover just eleven cities. Adding transport locations, they reach a total of twenty-seven. Independent regional networks (those not funded by state government) tend to be similarly sparse; NYSERNet has five PoP locations, PNGP has four, MAGPI and NOX each operate out of a single facility.

## Regulatory Environment

R&E networks have long enjoyed freedom from most of the regulations imposed on telcos and ISPs. To the extent that they operate private networks with restricted membership and purpose, R&E providers can avoid entanglement with tariffs, USF, DMCA, CALEA and other requirements that ISPs must deal with. By the same token, when R&E providers begin to operate networks that resemble or are identical to commercial ISPs, they risk entanglement, or accusations of unfair competition.

## **Business Model**

Without exception, R&E networks in the US are run by not-for-profit or government-funded entities, and commercial Internet networks are for-profit. The differences between these three models are too many to list, and they appear at every level within the organizations: governance, financial rules, product design, service models, pricing, contracts, marketing, etc. In addition, although the distinctions are eroding at the national level, regional networks still enforce acceptable use policies that determine the makeup of their membership and the uses those members can make of the network.

## **Competition**

Helped along by a regulatory environment that recognizes the many benefits of competition and a high density of potential customers, more than three quarters of US households are served by two or more broadband ISPs. Home and small business users can choose a connection based on price, features, support, reputation or any combination of the above; large users can connect to multiple providers for redundancy. The density of R&E connections is much lower, and as noted above, though their customers typically maintain multiple ISP connections, redundancy is not required on the R&E side. A degree of competition exists at the national level, though not driven by economics.

## **Audience**

A seemingly obvious, yet sometimes overlooked distinction is the intended users of the two classes of network. Commercial ISPs may have specific markets - wholesale to major connectors and other providers, small and medium business, multi-unit dwellings, etc. However, for any particular customer, there is an ISP who will want to sell to them, since everyone is a potential Internet user (as well as a consumer of the ISP's other services). R&E networks have some variations in their intended market, but generally serve medium and large K-20 educational institutions and research labs operated by large corporations or the government. These represent a tiny slice of the market for commercial Internet service, and are sufficiently similar that the R&E networks can specialize in every way in order to meet their needs.

## **Special Service Offerings**

Like most businesses, commercial ISPs strive for efficiency and lowered costs by standardizing their service offerings. They must also protect the operational integrity of their networks by enforcing demarcation points between customer and provider, and preventing customer network problems, mis-configurations or misbehavior from affecting the provider network. R&E networks, with much smaller customer bases, can afford to customize to a much greater degree, and to offer services that commercial ISPs would consider too risky or too likely to reveal details of the ISPs network and operations.