

David Conrad Internet Japan. Bio

[Click here](#) for APNIC's History on Wikipedia

David Conrad is a prominent American internet pioneer, network engineer, and executive who played a foundational role in building early internet infrastructure in Japan and the broader Asia-Pacific region. He is widely recognised for co-establishing the region's primary internet registries and serving as a high-level executive at ICANN.

The Japan Years (1991–1998)

Conrad moved to Japan in 1991, during the dawn of the public internet.

- **Internet Initiative Japan (IIJ):** He was hired as **employee #7** at [Internet Initiative Japan \(IIJ\)](#), which became one of Japan's very first commercial Internet Service Providers (ISPs).
- **Birth of APNIC (1993):** Operating out of Tokyo in September 1993, Conrad spearheaded a pilot project to manage IP address allocation in Asia. This led directly to the creation of **APNIC** (Asia Pacific Network Information Centre), the second Regional Internet Registry ever established. He served as its founding **Director-General** from 1993 until 1998.
- **Founding APRICOT (1996):** While in Japan, he also co-founded the **APRICOT** (Asia Pacific Rim Internet Conference on Operational Technologies) summit in 1996, which remains the premier internet infrastructure and operations conference for the region to this day.

Global Leadership & ICANN Career

After leaving Japan in 1998, Conrad returned to the United States and transitioned into top-tier global internet governance:

- **Internet Software Consortium (ISC):** He acted as Executive Director of the ISC from 1998 to 2001, overseeing key software architectures like BIND.
- **ICANN and IANA:** Conrad joined the [Internet Corporation for Assigned Names and Numbers \(ICANN\)](#) in 2005. Across two tenures, he scaled through several prominent roles, notably serving as the **General Manager of IANA** (Internet Assigned Numbers Authority) and eventually the **Chief Technology Officer (CTO)** of ICANN. In these roles, he was responsible for overseeing the security, stability, and resiliency of the internet's core system of unique identifiers.

Current Status

He is currently the co-founder of [Layer 9 Technologies](#), a specialized technical consultancy firm focusing on internet governance, policy, and network technologies.

Would you like to look deeper into his specific contributions to **early IP address distribution**, his work with the **BIND DNS software**, or his leadership during the **IANA transition** at ICANN?

His early years

David "Randy" Conrad's earliest records showcase a deeply technical background that long predates his work in Japan. Before helping to anchor the commercial internet in Asia, his career began in the early 1980s inside American research institutions:

Early Tech Work & Education

- **The IBM/UMD PC Project (1983):** Conrad began working directly with foundational internet technologies around **1983**. He served as a team leader for a joint project between **IBM** and the **University of Maryland (UMD)**.
- **Developing TCP/IP:** His team was tasked with building and deploying a commercial version of **TCP/IP** (the foundational protocol suite of the internet) to run on the IBM PC.
- **The OSI Shift:** Following that project, he remained at the university to assist in deploying a pilot implementation of the competing Open Systems Interconnection (OSI) protocol suite, an experience he noted taught him the massive advantages of TCP/IP's multi-vendor interoperability.

The PACCOM Project (1990)

Conrad left the University of Maryland in **1990** to join the [University of Hawaii at Mānoa](#).

- He worked as a young researcher on the **PACCOM (Pacific Communications)** project, which was funded heavily by NASA and the National Science Foundation (NSF).
- His primary responsibility was bridging the gap between North America and the Pacific Rim, helping to engineer **the very first internet connectivity** for various Asia-Pacific nations, universities, and research institutions.

The Meeting That Led to Japan

While working on PACCOM, the project hosted annual summits in Hawaii. It was at these technical exchange meetings that a young Conrad first met legendary global internet architects like [Dr. Jun Murai](#) (the "father of the internet" in Japan).

Impressed by his hands-on experience establishing regional networks, Murai personally invited Conrad to move to Tokyo. Conrad accepted, relocating to Japan in **1991/1992** to help orchestrate the birth of Internet Initiative Japan (IIJ).

(Note: If you encounter biographical entries mentioning a David Conrad born in Pittsburgh who attended Brown University and Juilliard, that is a different, well-known [American actor of the same name](#).)

Would you like to know more about the **IBM/University of Maryland TCP/IP project**, or how his **PACCOM network routing** laid the groundwork for APNIC?

Both

Here is the closer look at how **David Conrad's** early engineering work on both projects directly shaped the modern internet.

1. The IBM / University of Maryland TCP/IP Project (1983)

In 1983, the internet was in its infancy. The US military's ARPANET had just officially switched to **TCP/IP**, but it was primarily used on large Unix mainframes. Standard corporate personal computers, like the newly introduced **IBM PC**, did not have a native, commercial way to communicate using these protocols.

- **The Mandate:** Working at the University of Maryland (UMD), Conrad led a specialized technical team in a joint venture with IBM. Their goal was to build and program **one of the very first commercial TCP/IP software packages for the IBM PC**.
- **The "Protocol Wars" Lesson:** After successfully shipping the software, UMD tasked Conrad with deploying a pilot for the *OSI (Open Systems Interconnection)* protocol suite—a competing networking standard heavily backed by European telecommunications bodies and governments at the time.
- **The Outcome:** Operating both systems side-by-side proved to Conrad that TCP/IP's open, multi-vendor interoperability was vastly superior. This firsthand technical insight turned him into a lifelong champion of the open internet model, a philosophy he took with him to the Pacific.

2. The PACCOM Project Network Routing (1990)

In 1990, Conrad moved to the [University of Hawaii at Mānoa](#) to work as a researcher on the **PACCOM (Pacific Communications)** project. Funded by NASA and the National Science Foundation (NSF), PACCOM was a massive geopolitical and technical effort to physically link the networks of the Pacific Rim to the United States.

- **The Infrastructure:** At the time, countries like Japan, Australia, New Zealand, South Korea, and Hong Kong had internal academic networks, but they were digitally isolated from North America. Conrad and the PACCOM team engineered **the very first international subsea and satellite links** to physically bridge these networks.
- **The Strategic Push:** NASA and the NSF were highly motivated to convince the scientific communities in Japan and the Pacific to adopt TCP/IP over the competing OSI standard. By routing international traffic over PACCOM using TCP/IP, they successfully pushed the entire region toward the modern internet standard.
- **The Legacy Concept ("Economy Reports"):** During the 1990 PACCOM meeting in Hawaii, Conrad sat in on what were called **"economy reports"**. Representatives from different countries would present the status of their local network infrastructure.

Hearing these reports made it glaringly obvious to Conrad, [Dr. Jun Murai](#), and others that the Asia-Pacific region desperately needed a localized system to manage its own IP addresses. Without the routing foundations laid during the PACCOM days, the administrative framework for **APNIC** would never have been established in Tokyo three years later.

Would you like to explore the **specific technical challenges** they faced running international lines in the 90s, or see how this led directly to the licensing delays when he first arrived at **Internet Initiative Japan (IIJ)**?

Yes

The transition from academic engineering in Hawaii to establishing a commercial network in Tokyo forced David Conrad to navigate intense physical, bureaucratic, and geopolitical friction.

Part 1: The Technical Challenges of PACCOM (1990)

During the early 1990s, building out the **PACCOM** infrastructure meant working at the absolute boundaries of existing technology. Conrad and his colleagues faced severe physical limitations trying to run internet connections across the Pacific Ocean:

- **The Satellite Delays:** Much of the early international connectivity relied on satellite links. The physical distance a signal had to travel up to a satellite and back down introduced immense **latency** (delay). For interactive internet applications, engineers had to extensively tune TCP/IP routing tables just to prevent packets from timing out over the vast Pacific expanses.
- **The "Thin" Subsea Cables:** Subsea fiber-optic technology was in its complete infancy. Bandwidth was incredibly narrow—often operating on 56 kbit/s or fractional T1 lines. Conrad and the team had to compress and route traffic with extreme precision.
- **The Interoperability Barrier:** Because the National Science Foundation (NSF) and NASA wanted to push the Pacific Rim toward TCP/IP, Conrad had to build custom bridges and routing pathways to convince isolated Unix systems in universities across Japan, Korea, and Australia to talk to the American backbone without crashing.

Part 2: Moving to Japan & The IIJ Licensing Crisis (1991–1993)

When Conrad arrived in Tokyo as employee #7 of the newly formed **Internet Initiative Planning** (later [Internet Initiative Japan \(IIJ\)](#)), he slammed directly into a regulatory wall.

At the time, the **Ministry of Posts and Telecommunications (MPT)** held a rigid monopoly over Japanese communications.

- **The "Illegal" Commercial Internet:** Up until IIJ, the only networks operating in Japan were non-profit, government-sanctioned research networks like JUNET and WIDE. The MPT's framework explicitly did not allow a private company to sell "internet data" as a commercial product.
- **The License Delays:** The Ministry delayed IIJ's operator licensing for months. Bureaucrats simply did not understand the business plan of an independent Internet Service Provider (ISP) and feared it would disrupt the government-sponsored domestic network systems.
- **The "Bootleg" Router Solution:** Because commercial routing hardware was heavily regulated and astronomically expensive in Japan, the IIJ team had to get creative while waiting for official approvals. Conrad and the engineers engineered **Japan's first commercial routers out of standard DOS/V personal computers**. They loaded them with open-source BSD UNIX on bootable floppy disks, attaching modems via serial ports to run dedicated lines over PPP protocols.

Eventually, by **February 1994**, the continuous pressure from IIJ forced the Japanese government to break its rigid monopoly. IIJ was officially registered as a **Special Type II Telecommunications Carrier**, completely rewriting the rules and opening the floodgates for commercial internet in Japan.

Would you like to explore **how the DOS/V BSD router system worked**, or look into how this regulatory battle convinced him that **APNIC needed to be an independent body**?

Both

Part 1: How the DOS/V BSD Router System Worked

In the early 1990s, high-end routing hardware from companies like Cisco was strictly regulated, prohibitively expensive, and required extensive import approvals in Japan. Because **Internet Initiative Japan (IIJ)** was essentially stuck in a regulatory holding pattern, David Conrad and the early engineering team built their own infrastructure using consumer parts and open-source software.


The custom system operated on a distinct architecture:

- **The "DOS/V" Hardware Base:** The engineers utilized standard IBM-compatible PCs known as **DOS/V machines** (personal computers configured to display Japanese characters on standard VGA monitors). These off-the-shelf PCs were significantly cheaper than specialized enterprise hardware.
- **The Operating System:** Instead of DOS, the team loaded the PCs with **BSD UNIX (Berkeley Software Distribution)** via floppy disks. The [University of California, Berkeley](#) had recently integrated a highly stable, open-source **TCP/IP stack (sockets)** into BSD UNIX.
- **The Routing Functionality:** The DOS/V machine's CPU was configured to handle IP forwarding. Conrad and the team installed multi-port serial cards into the PC expansion slots. These serial ports were physically wired to modems running Point-to-Point Protocol (PPP) over leased phone lines.
- **The Result:** This engineered workaround allowed a standard, cheap PC to ingest raw internet data packages, read the IP headers, and successfully route them to the next destination. It effectively served as Japan's very first localized, commercial ISP backbone infrastructure—bypassing the need for specialized routing gear.

Part 2: How Regulatory Friction Sparked an Independent APNIC

While David Conrad was waiting for the Japanese Ministry of Posts and Telecommunications (MPT) to grant IIJ its official commercial operating licenses, he found himself with unexpected spare time. During this window, [Dr. Jun Murai](#) approached him to run a pilot project to address a growing global administrative crisis: **the allocation of IP addresses**.

The intense regulatory gridlock Conrad experienced firsthand at IIJ heavily influenced how he structured **APNIC** (Asia Pacific Network Information Centre):



- **The Problem with Government Monopolies:** Conrad's battle with the MPT proved that traditional government ministries moved far too slowly to handle the explosive growth of the internet. If a centralized government agency in Japan, Australia, or China was put in charge of assigning IP addresses, the internet would get bogged down in international geopolitics and local bureaucracy.

- **The Need for Neutrality:** The Asia-Pacific region was incredibly fractured by different languages, cultures, and levels of technological development. Conrad realized that a regional registry could not be beholden to any single country's telecom regulations.
- **The Birth of the Bottom-Up Model:** Operating initially as an informal experiment out of Tokyo in **September 1993**, Conrad established APNIC under a fundamentally different philosophy. It was built as a non-profit, neutral, **bottom-up organization** governed directly by the network operators who used it, rather than by politicians or states.

By keeping APNIC isolated from state-level regulatory interference, Conrad successfully scaled the registry into an independent entity. It seamlessly distributed IP space across multiple nations throughout the 1990s, completely unhindered by the type of regulatory friction that had nearly stopped IJF from launching.

How APNIC transitioned from a part-time pilot project to a legally recognized entity

The transition of **APNIC** from an informal, part-time experiment into a legally recognized international organization is a masterclass in early internet governance. Driven by David Conrad and regional internet pioneers, the evolution occurred in three distinct stages over five years.

1. The "Interim Project" and Authority Crisis (1993–1994)

When the APNIC pilot project officially concluded its experimental phase between **September 1993 and June 1994**, its creators faced a massive roadblock.

- **No Legal Umbrella:** The pilot had been birthed under an informal regional collective called the *Asia Pacific Networking Group (APNG)*. However, APNG was just a loose committee; it had **no legal authority** or status to sign contracts, open bank accounts, or employ staff.
- **The IANA Safety Net:** To keep the registry from collapsing, the consultative committee decided that APNIC's operational authority must bypass regional governments entirely and derive directly from the **Internet Assigned Numbers Authority (IANA)**, led by Jon Postel.
- **The Bootstrapped Office:** IANA publicly legitimized APNIC by officially delegating early IP blocks to it. Still, the entity had no formal corporate structure, no membership base, and no fee model. It survived strictly as a guest in Tokyo, heavily subsidized by **JPNIC** (the Japanese Network Information Centre), which generously provided 10% of its own budget and shared its University of Tokyo computer centre space.

2. Formal Incorporation and Global Split (1995–1996)

By 1995, the workload was scaling uncontrollably as commercial internet exploded across Asia. David Conrad, operating as the first Director-General, knew APNIC could no longer run on hand-outs and volunteer hours.

- **The Shell Legal Entity (1996):** In **1996**, APNIC officially achieved independence from APNG and established itself as a formal corporate entity. Due to severe tax and legal complexities for foreign non-profits operating inside Japan, APNIC's legal advisors actually recommended incorporating the initial business entity—**APNIC Ltd**—all the way across the globe in the **Seychelles**.

- **The New Membership Business Model:** With a legal entity finally secured, Conrad established a formal **membership and tier-based fee structure**. For the first time, ISPs and enterprises across Asia paid direct membership dues to APNIC to secure their IP address allocations, making the organization entirely self-funded and financially sustainable.

3. The Move to Brisbane and Institutional Maturation (1997–1998)

Even with a legal corporate structure and steady funding, APNIC's physical operational environment in Tokyo was highly restrictive. High operating costs and strict local labor laws meant the registry was suffocating—trapped with a tiny staff of only 4 to 5 people trying to service the entire Asia-Pacific continent.

Tokyo (1993–1997)	Brisbane, Australia (1998–Present)
• Restrictive local environment	• Stable, high-tech infrastructure
• High cost of operation	• Low cost of living & operations
• Maxed out at 4–5 staff members	• Massive tax advantages for non-profits

- **The KPMG Search:** In 1997, APNIC hired global consulting firm **KPMG** to audit the entire region and find a permanent, optimal home for the registry. After surveying several economies, **Brisbane, Australia** won the bid. Brisbane offered stable infrastructure, lower operational costs, and immense tax advantages specifically tailored for international membership organizations.
- **The Final Relocation (1998):** Between **April and August 1998**, the physical assets and registry data were systematically migrated from Tokyo to Australia with zero downtime or service interruption. The organization re-registered locally as an Australian company: **APNIC Pty Ltd**.

With the legal and geographic transition successfully finalized, David Conrad stepped down as Director-General, passing the leadership baton to **Paul Wilson (as Director General from 1998 to 2024)**. Having evolved from a single, shared computer in a borrowed Tokyo university room, APNIC was now a fully recognized, financially independent international institution anchoring global internet governance.

Would you like to explore how **Paul Wilson expanded APNIC** once it landed in Brisbane, or look into the specifics of how **IANA first divided global IP address blocks** among the early registries?

How IANA first divided global IP address blocks among the early registries

In the early days of the internet, there was no complex global system for distributing IP addresses. **Jon Postel**, the founder and head of the **Internet Assigned Numbers Authority (IANA)**, initially managed the entire global pool of IPv4 addresses using a simple, handwritten physical notebook.

As the internet scaled from an American research network into a global commercial phenomenon, this centralized, single-person system became unsustainable. The transition to dividing the global IP pool among Regional Internet Registries (RIRs) occurred in a structured, phases-based rollout:

1. The Classful Network Era (Pre-1993)

Before RIRs existed, IANA handed out IP addresses in rigid blocks based on a system called **Classful Routing**.

- **Class A Blocks:** These massive chunks contained **16.7 million addresses** each (known as a /8 block). They were handed out directly by IANA to early pioneers, massive corporations (like IBM, Apple, and Xerox), and US government agencies.
- **The InterNIC Monopoly:** For everyone else, IANA deputized a single American entity called **InterNIC** (operated by Network Solutions) to hand out smaller Class B and Class C blocks to the rest of the world on a first-come, first-served basis.

2. The Great Efficiency Crisis & CIDR (1993)

By 1993, the Classful system was causing two catastrophic problems: the global routing tables were overloading, and Class B addresses were rapidly running out.


To save the internet from collapsing, engineers invented **CIDR** (Classless Inter-Domain Routing). This allowed IANA to break the address pool down into highly customizable, flexible fragments rather than rigid classes. This breakthrough coincided perfectly with David Conrad establishing **APNIC** in Asia and the formalization of **RIPE NCC** in Europe.

3. How IANA Divided the Map

With CIDR in place, Jon Postel and IANA officially began carving up the global IPv4 address space to distribute chunks to the newly forming regional registries.

IANA used a **geographically ring-fenced block allocation** strategy to keep global routing efficient:

- **The Regional Delegations:** IANA block-allocated massive /8 chunks (each containing 16.7 million addresses) directly to the specific registries based on their geographic boundaries:
 - **RIPE NCC (Europe & Middle East):** Received blocks like 62.0.0.0/8, 80.0.0.0/8, and 193.0.0.0/8.
 - **APNIC (Asia-Pacific):** David Conrad's registry was allocated blocks like 202.0.0.0/8 and 203.0.0.0/8.
 - **ARIN / InterNIC (North America):** Retained historical blocks like 24.0.0.0/8 and 63.0.0.0/8.
- **The "Just-In-Time" Supply Rule:** To prevent hoarding, IANA did not give the registries all their space at once. A strict inventory rule was established: an RIR like APNIC had to prove to IANA that its current /8 blocks were **80% utilized** by local ISPs. Only when they hit that threshold would IANA release the next block from the global pool.



4. The Global Policy Integration

By the late 1990s and early 2000s, this division method was formalized into an international framework known as the **ICP-2 (Internet Coordination Policy 2)**. This baseline explicitly dictated that IANA was the ultimate root authority of the pool, but it must remain a neutral "wholesale" supplier.

This elegant system ensured that while IANA held the master keys to the global ledger, local registries like APNIC had full autonomy to divide their allocated blocks among their own domestic internet service providers without needing American approval for every transaction.

Which specific IP blocks were originally handed over to APNIC during David Conrad's tenure

During David Conrad's tenure as the founding Director-General of [APNIC](#) (1993–1998), the internet registry was handed its very first chunks of the global IPv4 address pool directly by Jon Postel at IANA.

Because this was the dawn of **CIDR** (Classless Inter-Domain Routing), IANA block-allocated massive, geographically ring-fenced /8 blocks (each containing **16.7 million addresses**) specifically to be distributed across the Asia-Pacific region.

The primary blocks handed over during his tenure were:

The Core 202/7 Block Allocation (1993)

When David Conrad first met with Jon Postel at a California Pizza Kitchen to outline the launch of APNIC, Postel formally agreed to delegate the **202/7 range**. In networking terms, a /7 block is massive because it combines two consecutive /8 blocks:

- 202.0.0.0/8
- 203.0.0.0/8

These two blocks became the primary, original "fresh" registry space that Conrad's team used to bootstrap the commercial internet across Japan, Australia, China, South Korea, and the rest of the Pacific Rim throughout the mid-1990s.

Early Multi-Registry /8 Blocks (May 1993)

Right at the transition point when APNIC was transitioning into its pilot phase, IANA also assigned management authority for sections of various older /8 blocks that already contained mixed global registrations. APNIC was given administrative authority over regional networks sitting inside:

- 150.0.0.0/8
- 153.0.0.0/8

The Migration of Legacy Space (1995–1998)

As APNIC grew into a legally recognized entity, Conrad's team was handed responsibility for managing "historical" IP address records. Before APNIC existed, entities like the Australian

Academic and Research Network (AARNet) or InterNIC handed out Class B and Class C blocks to Asian organizations directly.

During Conrad's tenure, IANA systematically began shifting the registry records for those old allocations into the APNIC database. This included fragments of historical Class B and Class C blocks spanning ranges such as 128.0.0.0/8 **through** 143.0.0.0/8, as well as the famous 192.168.0.0/8 space.

Would you like to know how Conrad's team **rationed these early blocks** to prevent different countries from hoarding them, or explore the history of the famous 202 and 203 blocks in early web routing?

Both

The allocation of the 202.0.0.0/8 and 203.0.0.0/8 blocks by IANA to APNIC in **April 1994** marked a major turning point in internet history. It shifted address management away from an American monopoly and established a localized, well-regulated digital economy for the Asia-Pacific region.

Part 1: The History of the 202 and 203 Blocks in Early Routing

Before these blocks were delegated, if an entity in Tokyo, Sydney, or Seoul wanted to connect to the internet, they had to apply directly to **InterNIC** in the United States to secure an IP address block.

- **The Geopolitical Pivot:** In April 1994, IANA officially shifted authority over the 202.0.0.0/8 and 203.0.0.0/8 blocks to APNIC. Overnight, these two numbers became the digital prefix for an entire continent. If a data packet on the global network started with 202 or 203, routers from California to London instantly knew to direct that traffic across the Pacific subsea cables toward Asia.
- **The Backbone of Early Web Infrastructure:** Because these blocks were the very first fresh spaces managed natively within Asia, they were used to build the foundational architecture of the Asian internet. The region's earliest commercial websites, root DNS servers, and ISP exchange points were entirely built within this numeric real estate. Even today, a traceroute to some of Asia's oldest telcos or academic institutions will still reveal routing paths grounded in the 202 and 203 ranges.

Part 2: How David Conrad's Team Rationed the Space to Prevent Hoarding

When David Conrad took over the pilot project, he was running APNIC out of a borrowed computer laboratory at the University of Tokyo with a staff of just himself and an email mailing list of volunteers.

He faced an immediate crisis: the commercial internet was exploding, and a few technologically advanced countries could easily hoard the entire 202/7 address space, leaving developing nations with nothing. To prevent this, Conrad's team engineered strict rationing and conservation mechanisms:

```
[Requestor Asks for IPs] —> Must Show Local Need —> Subnet Size Hard-Limited —>
Assigned /24 Prefix Max
```

- **The Death of Large Static Blocks:** Under the old "Classful" system, if a company asked for a network, they were blindly handed a Class B block containing 65,536 addresses, even if they only needed 500. Conrad strictly enforced **CIDR (Classless Inter-Domain Routing)**. APNIC mathematically sliced the 202 and 203 blocks into precise, variable-length subnets matching exact, immediate infrastructure needs.
- **Strict Documentation of "Demonstrated Need":** Under APNIC's early framework, an internet service provider could not simply buy IP addresses to hoard for future use. ISPs had to submit detailed architectural plans proving they had the physical hardware and active customer demand to justify the allocation.
- **The /24 Hard Minimums:** APNIC tightly restricted initial allocations to tiny fractional blocks, frequently defaulting to small /24 **allocations** (just 256 addresses at a time) for new members. An ISP had to prove they had utilized at least 80% of their existing pool before APNIC would unlock the next consecutive block.
- **The Birth of National Registries (NIRs):** To balance regional equity, Conrad worked with local pioneers to establish national sub-registries, such as **JPNIC** in Japan and **KRNIC** in South Korea. APNIC allocated bounded chunks of the 202 and 203 pools to these national bodies. This ensured local internet authorities could distribute addresses fairly to their domestic businesses without a single country draining the shared regional pool.

Through these rationing protocols, a tiny, bootstrapped organization in Tokyo successfully managed 33.5 million addresses so efficiently that they sustained the explosive, multi-nation growth of the early Asian internet for years.

Specific criteria Conrad used to decide when an ISP qualified for more IP space

David Conrad and the early APNIC team determined ISP eligibility for more IPv4 space using a strict set of quantitative utilisation metrics, slow-start scaling, and documented proof of immediate technical need*.

Conrad co-authored these foundational rules in [RFC 2050](#) (*Internet Registry IP Allocation Guidelines*), alongside Jon Postel and Daniel Karrenberg. The framework replaced arbitrary assignments with measurable thresholds:

1. The 80% Utilization Rule

Before an ISP could receive a single additional block from APNIC, it had to demonstrate exhaustive usage of its current inventory:

- **The Threshold:** The ISP had to prove it had globally assigned or routed at least **80% of its previously allocated address space**.
- **Whois Verification:** This usage could not merely exist on internal spreadsheets. The ISP had to demonstrate that its customer assignments were actively registered in the public **APNIC Whois database**, allowing staff to audit downstream assignments.

2. Time-Phased Usage Milestones (25% / 50%)

To qualify for any subnet size, an ISP's architectural deployment schedule had to meet strict timeline ratios:

- **Immediate Need:** The ISP had to document that at least **25% of the requested address space** would be deployed immediately upon receipt.
- **One-Year Target:** The ISP had to provide a credible business plan proving it would achieve at least **50% utilisation within 12 months**. Requests based on speculative two-to-five-year projections were rejected to prevent space from sitting idle.

3. The "Slow-Start" Procedure & Assignment Windows

Rather than handing ISPs whatever block size they requested, Conrad implemented a phased trust process known as **slow-start**:

- **The Baseline Window:** A brand-new ISP began with an "**Assignment Window**" of zero. This meant the ISP could not delegate sub-blocks to its own customers without first sending the documentation to APNIC staff for review and manual approval.
- **Doubling the Window:** Once an ISP proved it was following best conservation practices and consistently hitting the 80% threshold, APNIC expanded its Assignment Window. The ISP earned the autonomy to independently assign larger chunks (e.g., advancing from a /24 to a /23, then a /22).
- **Subsequent Allocation Sizing:** Subsequent blocks were capped at what the ISP was projected to consume over a **short operational window (typically 3 to 12 months)**, preventing the stockpiling of multi-year inventories.

4. Technical Justification & Efficient Topologies

ISPs were required to submit engineering schematics verifying that their network topology maximized address conservation:

- **Variable-Length Subnet Masking (VLSM):** The applicant had to confirm that its routing hardware actively supported CIDR and VLSM rather than wasting address space on fixed legacy subnet masks.
- **No Static Dial-Up IPs:** An ISP could not qualify for large blocks if it intended to assign static IP addresses to dial-up modem users. APNIC mandated dynamic pooling (using protocols like DHCP and RADIUS) so that single IP addresses were shared among multiple transient users.
- **Hardware Proof:** ISPs had to supply proof of physical network buildouts—such as receipts for terminal servers, modems, or leased line contracts—to prove the infrastructure was real and ready to be addressed.

**** End of Page**