


Internet 3.0: Performance and Productivity Issues of the Internet and a Proposal for the Next Generation



RAJ JAIN

Washington University in Saint Louis
Saint Louis, MO 63130

Jain@wustl.edu

St. Louis - Computer Measurement Group Meeting
November 18th, 2008

These slides and Audio/Video recordings of this talk are at:

http://www.cse.wustl.edu/~jain/talks/in3_cmgs.htm



http://www.cse.wustl.edu/~jain/talks/in3_cmgs.htm

©2008 Raj Jain

1



1. Internet and Productivity
2. What is Internet 3.0?
3. Why should you keep on the top of Internet 3.0?
4. What are we missing in the current Internet?
5. Our Proposed Architecture for Internet 3.0

Acknowledgement: This research is sponsored by a grant from Intel Research Council.



http://www.cse.wustl.edu/~jain/talks/in3_cmgs.htm

©2008 Raj Jain

3


Networking and Productivity

- ❑ **Networking ⇒ Enhanced productivity**
 - Faster communication, Faster access to information
- ❑ **Networking is the bottleneck**
 - Communication is more critical than computing
 - ❑ Greeting cards contain more computing power than all computers before 1950.
 - ❑ Genesis's game has more processing than 1976 Cray supercomputer.
 - Network is the bottleneck. Productivity of people, companies and countries depends upon the speed of their network
- ❑ **Networking is the reason for reduced productivity**
 - Spam, Unlimited information => Distractions, Misuse

What is Internet 3.0?

- ❑ Internet 3.0 is the architecture of the next generation of Internet
- ❑ Named by me along the lines of “Web 2.0”
- ❑ National Science Foundation is planning a \$300M+ research and infrastructure program on next generation Internet
 - Testbed: “Global Environment for Networking Innovations” (GENI)
 - Architecture: “Future Internet Design” (FIND).
- ❑ Internet 3.0 is more intuitive than GENI/FIND
- ❑ Most of the networking researchers will be working on GENI/FIND for the coming years
- ❑ Q: How would you design Internet today? Clean slate design.
- ❑ Ref: <http://www.nsf.gov/cise/cns/geni/>

Life Cycles of Technologies



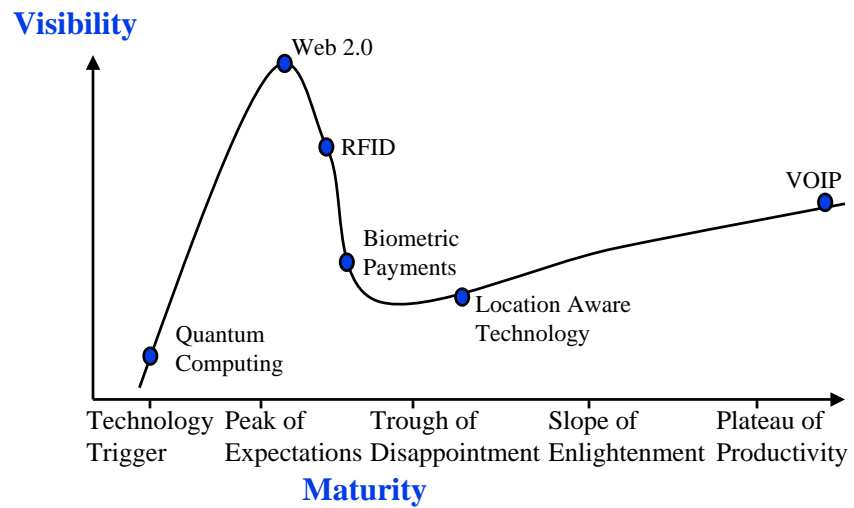
Washington University in St. Louis

http://www.cse.wustl.edu/~jain/talks/in3_cmgs.htm

©2008 Raj Jain

8

Networking Hype Cycle 2006



Washington University in St. Louis

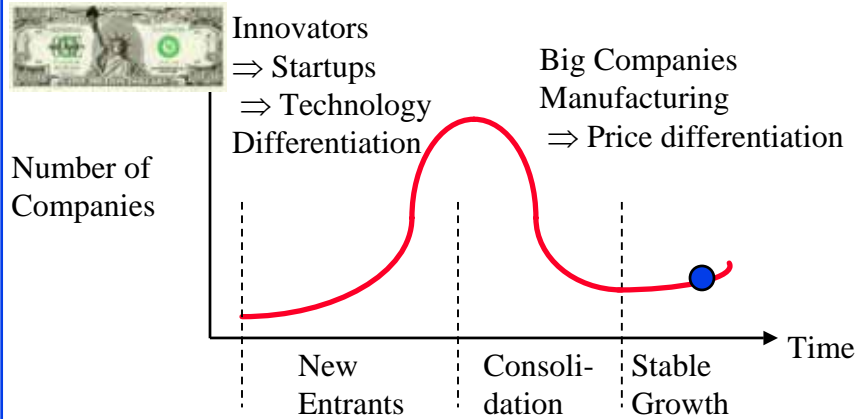
http://www.cse.wustl.edu/~jain/talks/in3_cmgs.htm

Based on Gartner Research (July 2006)

©2008 Raj Jain

9



Industry Growth: Formula for Success



- **10-20-70 Formula:** 10% of R&D on distant future, 20% on near future, 70% on today's products


Internet Generations

- **Internet 1.0** (1969 – 1989) – Research project
 - RFC1 is dated April 1969.
 - ARPA project started a few years earlier
 - IP, TCP, UDP
 - Mostly researchers
 - Industry was busy with proprietary protocols: SNA, DECnet, AppleTalk, XNS
- **Internet 2.0** (1989 – Present) – Commerce ⇒ new requirements
 - Security RFC1108 in 1989
 - NSFnet became commercial
 - Inter-domain routing: OSPF, BGP,
 - IP Multicasting
 - Address Shortage IPv6
 - Congestion Control, Quality of Service,...



Ten Problems with Current Internet

1. Designed for research
⇒ Trusted systems
Used for Commerce
⇒ Untrusted systems
2. Control, management, and Data path are intermixed ⇒ security issues
3. Difficult to represent organizational, administrative hierarchies and relationships. Perimeter based.



Problems (cont)

4. Identity and location in one (IP Address)
Makes mobility complex.
5. Location independent addressing
⇒ Most services require nearest server.
⇒ Also, Mobility requires location
6. No representation for real end system: the human.



Problems (cont)

7. Assumes live and awake end-systems
Does not allow communication while sleeping.
Many energy conscious systems today sleep.



8. Single-Computer to single-computer communication \Rightarrow Numerous patches needed for communication with globally distributed systems.



9. Symmetric Protocols
 \Rightarrow No difference between a PDA and a Google server.



Google

Problems (Cont)

10. Stateless \Rightarrow Can't remember a flow
 \Rightarrow QoS difficult.
QoS is generally for a flow and not for one packet




Our Proposed Solution: GINA

- ❑ **Generalized Inter-Networking Architecture**
- ❑ Take the best of what is already known
 - Wireless Networks, Optical networks, ...
 - Transport systems: Airplane, automobile, ...
 - Communication systems: Wired Phone networks, Cellular networks,...
- ❑ Develop a consistent general purpose, evolvable architecture that can be customized by implementers, service providers, and users



Names, IDs, Addresses




Name: John Smith

ID: 012-34-5678

Address:
1234 Main Street
Big City, MO 12345
USA

- ❑ Address changes as you move, ID and Names remain the same.
- ❑ **Examples:**
 - Names: Company names, DNS names (microsoft.com)
 - IDs: Cell phone numbers, 800-numbers, Ethernet addresses, Skype ID, VOIP Phone number
 - Addresses: Wired phone numbers, IP addresses

Realms



- ❑ Object names and Ids are defined within a realm
- ❑ A realm is a **logical** grouping of objects under an administrative domain
- ❑ The Administrative domain may be based on Trust Relationships
- ❑ A realm represents an organization
 - Realm managers set policies for communications
 - Realm members can share services.
 - Objects are generally members of multiple realms
- ❑ Realm Boundaries: Organizational, Governmental, ISP, P2P,...

Realm = Administrative Group


Physical vs Logical Connectivity

- ❑ Physically and logically connected:
All computers in my lab
= Private Network,
Firewalled Network
- ❑ Physically disconnected but logically connected:
My home and office computers
- ❑ Physically connected but logically disconnected: Passengers on a plane,
Neighbors, Conference attendees sharing a wireless network, A visitor



Physical connectivity \neq Trust

Id-Locator Split Architecture (MILSA)



- ❑ Realm managers resolve current location for a given host-ID
- ❑ Allows mobility, multi-homing
- ❑ Ref: Our Globecom 2008 paper [3]

Server and Gatekeeper Objects

- ❑ Each realm has a set of server objects, e.g., forwarding, authentication, encryption, storage, transformation, ...
- ❑ Some objects have built-in servers, e.g., an “enterprise router” may have forwarding, encryption, authentication services.
- ❑ Other objects rely on the servers in their realm
- ❑ Authentication servers (AS) add their signatures to packets and verify signatures of received packets..
- ❑ Storage servers store packets while the object may be sleeping and may optionally aggregate/compress/transform data. Could wake up objects.
- ❑ Objects can appoint proxies for any function(s)
- ❑ Gatekeepers enforce policies: Security, traffic, QoS

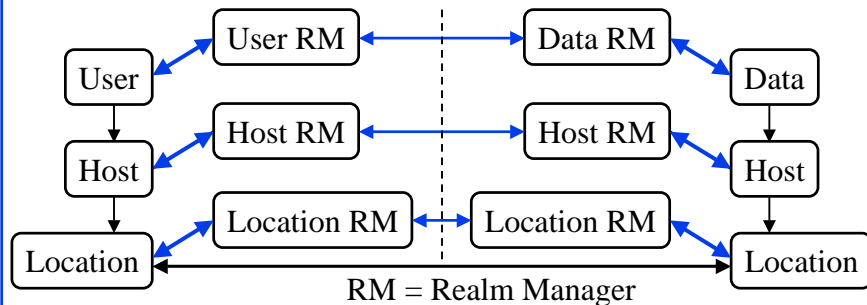
Servers allow simple energy efficient end devices

User- Host- and Data Centric Models

- ❑ All discussion so far assumed host-centric communication
 - Host mobility and multihoming
 - Policies, services, and trust are related to hosts
- ❑ User Centric View:
 - Bob wants to watch a movie
 - Starts it on his media server
 - Continues on his iPod during commute to work
 - Movie exists on many servers
 - Bob may get it from different servers at different times or multiple servers at the same time
- ❑ Can we just give addresses to users and treat them as hosts?
 - No! ⇒ Policy Oriented Naming Architecture (PONA)



Policy Oriented Naming Architecture




- ❑ Both Users and data need hosts for communication
- ❑ Data is easily replicable. All copies are equally good.
- ❑ Users, Hosts, Infrastructure, Data belong to different realms (organizations).
- ❑ Each object has to follow its organizational policies.

PONA (Cont)


- ❑ User and data realms are higher level than host realms
- ❑ Most communication is user-data communication
- ❑ User, Host, and Data can move independently
 - Hosts move from one location to next
 - Users and data can move from one host to the next
- ❑ User ID \Rightarrow Host ID \Rightarrow Host Location = Address
- ❑ User realm managers provide User ID to Host ID translation
- ❑ Realm managers enforce organizational policies
- ❑ Realm managers setup trust relationships between organizations

Virtualizable Network Concept



Ref: T. Anderson, L. Peterson, S. Shenker, J. Turner, "Overcoming the Internet Impasse through Virtualization," *Computer*, April 2005, pp. 34 – 41.


Realm Virtualization



- ❑ Old: Virtual networks on a common infrastructure
- ❑ New: Virtual user realms on virtual host realms on a group of infrastructure realms. 3-level hierarchy not 2-level. Multiple organizations at each level.

Internet 1.0 vs. Internet 3.0

	Feature	Internet 1.0	Internet 3.0
1.	Energy Efficiency	Always-on	Green ⇒ Mostly Off
2.	Mobility	Mostly stationary computers	Mostly mobile <i>objects</i>
3.	Computer-Human Relationship	Multi-user systems ⇒ Machine to machine comm.	Multi-systems user ⇒ Personal comm. systems
4.	End Systems	Single computers	Globally distributed systems
5.	Protocol Symmetry	Communication between equals ⇒ Symmetric	Unequal: PDA vs. big server ⇒ Asymmetric
6.	Design Goal	Research ⇒ Trusted Systems	Commerce ⇒ No Trust Map to organizational structure
7.	Ownership	No concept of ownership	Hierarchy of ownerships, administrations, communities
8.	Sharing	Sharing ⇒ Interference, QoS Issues	Sharing <i>and</i> Isolation ⇒ Critical infrastructure
9.	Switching units	Packets	Packets, Circuits, Wavelengths, Electrical Power Lines, ...
10.	Applications	Email and Telnet	Information Retrieval, Distributed Computing, Distributed Storage, Data diffusion



Summary

1. Networking is the key enabler for productivity and also the bottleneck.
2. Internet 3.0 is the next generation of Internet. It must be secure, allow mobility and energy efficiency.
3. Active industry involvement in the design essential. Leading networking companies must actively participate.
4. Must be designed for commerce
⇒ Must represent multi-organizational structure and policies
5. Users, Hosts, and infrastructures belong to different realms (organizations). Users/data/hosts should be able to move freely without interrupting a network connection.

References

1. Jain, R., "Internet 3.0: Ten Problems with Current Internet Architecture and Solutions for the Next Generation," in Proceedings of Military Communications Conference (MILCOM 2006), Washington, DC, October 23-25, 2006, <http://www.cse.wustl.edu/~jain/papers/gina.htm>
2. Subharthi Paul, Raj Jain, Jianli Pan, and Mic Bowman, "A Vision of the Next Generation Internet: A Policy Oriented View," British Computer Society Conference on Visions of Computer Science, Sep 2008, <http://www.cse.wustl.edu/~jain/papers/pona.htm>
3. Jianli Pan, Subharthi Paul, Raj Jain, and Mic Bowman, "MILSA: A Mobility and Multihoming Supporting Identifier-Locator Split Architecture for Naming in the Next Generation Internet,," Globecom 2008, Nov 2008, <http://www.cse.wustl.edu/~jain/papers/milsa.htm>