

A Survey of the Research on Future Internet Architectures

Jianli Pan, Subharthi Paul, and Raj Jain, Washington University

ABSTRACT

The current Internet, which was designed over 40 years ago, is facing unprecedented challenges in many aspects, especially in the commercial context. The emerging demands for security, mobility, content distribution, etc. are hard to be met by incremental changes through ad-hoc patches. New clean-slate architecture designs based on new design principles are expected to address these challenges. In this survey article, we investigate the key research topics in the area of future Internet architecture. Many ongoing research projects from United States, the European Union, Japan, China, and other places are introduced and discussed. We aim to draw an overall picture of the current research progress on the future Internet architecture.

INTRODUCTION

The Internet has evolved from an academic network to a broad commercial platform. It has become an integral and indispensable part of our daily life, economic operation, and society. However, many technical and non-technical challenges have emerged during this process, which call for potential new Internet architectures. Technically, the current Internet was designed over 40 years ago with certain design principles. Its continuing success has been hindered by more and more sophisticated network attacks due to the lack of security embedded in the original architecture. Also, IP's narrow waist means that the core architecture is hard to modify, and new functions have to be implemented through myopic and clumsy ad hoc patches on top of the existing architecture. Moreover, it has become extremely difficult to support the ever increasing demands for security, performance reliability, social content distribution, mobility, and so on through such incremental changes. As a result, a clean-slate architecture design paradigm has been suggested by the research community to build the future Internet. From a non-technical aspect, commercial usage requires fine-grained security enforcement as opposed to the current "perimeter-based" enforcement.

Security needs to be an inherent feature and integral part of the architecture. Also, there is a significant demand to transform the Internet from a simple "host-to-host" packet delivery paradigm into a more diverse paradigm built around the data, content, and users instead of the machines. All of the above challenges have led to the research on future Internet architectures.

Future Internet architecture is not a single improvement on a specific topic or goal. A clean-slate solution on a specific topic may assume the other parts of the architecture to be fixed and unchanged. Thus, assembling different clean-slate solutions targeting different aspects will not necessarily lead to a new Internet architecture. Instead, it has to be an overall redesign of the whole architecture, taking all the issues (security, mobility, performance reliability, etc.) into consideration. It also needs to be evolvable and flexible to accommodate future changes. Most previous clean-slate projects were focused on individual topics. Through a collaborative and comprehensive approach, the lessons learned and research results obtained from these individual efforts can be used to build a holistic Internet architecture.

Another important aspect of future Internet architecture research is the experimentation testbeds for new architectures. The current Internet is owned and controlled by multiple stakeholders who may not be willing to expose their networks to the risk of experimentation. So the other goal of future Internet architecture research is to explore open virtual large-scale testbeds without affecting existing services. New architectures can be tested, validated, and improved by running on such testbeds before they are deployed in the real world.

In summary, there are three consecutive steps leading toward a working future Internet architecture:

Step 1: Innovations in various aspects of the Internet

Step 2: Collaborative projects putting multiple innovations into an overall networking architecture

Step 3: Testbeds for real-scale experimentation

It may take a few rounds or spirals to work out a

This work was supported in part by a grant from Intel Corporation and NSF CISE Grant #1019119.

future Internet architecture that can fit all the requirements.

Future Internet research efforts may be classified based on their technical and geographical diversity. While some of the projects target at individual topics, others aim at holistic architectures by creating collaboration and synergy among individual projects. Research programs specifically aimed at the design of the future Internet have been set up in different countries around the globe, including the United States, the European Union (EU), Japan, and China. The geographical diversity of research presents different approaches and structures of these different research programs. While dividing the projects by their major topics is also possible, due to the holistic architecture goals, different projects may have some overlap.

Over the past few years' future Internet research has gathered enormous momentum as evidenced by the large number of research projects in this area. In this article, primarily based on the geographical diversity, we present a short survey limited in scope to a subset of representative projects and discuss their approaches, major features, and potential impact on the future.

We discuss the key research topics and design goals for the future Internet architectures. Research projects in the United States, European Union, and Asian countries are discussed in detail, respectively. Some of our discussions and perspectives on future Internet architectures are included later. Finally, a summary concludes the article.

KEY RESEARCH TOPICS

In this section, we discuss some key research topics that are being addressed by different research projects.

Content- or data-oriented paradigms: Today's Internet builds around the "narrow waist" of IP, which brings the elegance of diverse design above and below IP, but also makes it hard to change the IP layer to adapt for future requirements. Since the primary usage of today's Internet has changed from host-to-host communication to content distribution, it is desirable to change the architecture's narrow waist from IP to the data or content distribution. Several research projects are based on this idea. This category of new paradigms introduces challenges in data and content security and privacy, scalability of naming and aggregation, compatibility and co-working with IP, and efficiency of the new paradigm.

Mobility and ubiquitous access to networks: The Internet is experiencing a significant shift from PC-based computing to mobile computing. Mobility has become the key driver for the future Internet. Convergence demands are increasing among heterogeneous networks such as cellular, IP, and wireless ad hoc or sensor networks that have different technical standards and business models. Putting mobility as the norm instead of an exception of the architecture potentially nurtures future Internet architecture with innovative scenarios and applications. Many collaborative research projects in academia and industry are pursuing such research topics with great interest.

These projects also face challenges such as how to trade off mobility with scalability, security, and privacy protection of mobile users, mobile endpoint resource usage optimization, and so on.

Cloud-computing-centric architectures: Migrating storage and computation into the "cloud" and creating a "computing utility" is a trend that demands new Internet services and applications. It creates new ways to provide global-scale resource provisioning in a "utility-like" manner. Data centers are the key components of such new architectures. It is important to create secure, trustworthy, extensible, and robust architecture to interconnect data, control, and management planes of data centers. The cloud computing perspective has attracted considerable research effort and industry projects toward these goals. A major technical challenge is how to guarantee the trustworthiness of users while maintaining persistent service availability.

Security: Security was added into the original Internet as an additional overlay instead of an inherent part of the Internet architecture. Now security has become an important design goal for the future Internet architecture. The research is related to both the technical context and the economic and public policy context. From the technical aspect, it has to provide multiple granularities (encryption, authentication, authorization, etc.) for any potential use case. Also, it needs to be open and extensible to future new security related solutions. From the non-technical aspect, it should ensure a trustworthy interface among the participants (e.g., users, infrastructure providers, and content providers). There are many research projects and working groups related to security. The challenges on this topic are very diverse, and multiple participants make the issue complicated.

Experimental testbeds: As mentioned earlier, developing new Internet architectures requires large-scale testbeds. Currently, testbed research includes multiple testbeds with different virtualization technologies, and the federation and coordination among these testbeds. Research organizations from the United States, European Union, and Asia have initiated several programs related to the research and implementation of large-scale testbeds. These projects explore challenges related to large-scale hardware, software, distributed system test and maintenance, security and robustness, coordination, openness, and extensibility.

Besides these typical research topics, there are several others, including but not limited to networked multimedia; "smart dust," also called the "Internet of things"; and Internet services architecture. However, note that in this survey, we are not trying to enumerate all the possible topics and corresponding research projects. Instead, we focus on a representative subset and discuss a few important ongoing research projects.

Due to length limitations, we are not able to enumerate all the references for the projects discussed below. However, we do have a longer survey [18], which includes a more complete reference list for further reading.

Since the primary usage of the today's Internet has changed from host-to-host communication to content distribution, it is desirable to change the architecture's narrow waist from IP to the data or content distribution.

Categories	Project or cluster names (selected)
FIA	NDN, MobilityFirst, NEBULA, XIA, etc.
FIND	CABO, DAMS, Maestro, NetSerV, RNA, SISS, etc. (more than 47 total)
GENI	Spiral1: (5 clusters totally): DETER (1 project), PlanetLab (7 projects), ProtoGENI (5 projects), ORCA (4 projects), ORBIT (2 projects; 8 not classified; 2 analysis projects)
	Spiral2: over 60 active projects as of 2009*
	Spiral3: about 100 active projects as of 2011*
* GENI design and prototyping projects can last for more than one spiral.	

Table 1. U.S. projects and clusters on the future Internet.

RESEARCH PROJECTS FROM THE UNITED STATES

Research programs on future Internet architecture in United States are administrated by the National Science Foundation (NSF) directorate for Computer and Information Science and Engineering (CISE).

FIA AND FIND

The Future Internet Architecture (FIA) program [1] of the National Science Foundation (NSF) is built on the previous program, Future Internet Design (FIND) [2]. FIND funded about 50 research projects on all kinds of design aspects of the future Internet. FIA is the next phase to pull together the ideas into groups of overall architecture proposals. There are four such collaborative architecture groups funded under this program, and we introduce them here. Table 1 illustrates the overall research projects from the United States, including FIA and FIND.

Named Data Networking (NDN) — The Named Data Networking (NDN) [3] project is led by the University of California, Los Angeles with participation from about 10 universities and research institutes in the United States. The initial idea of the project can be traced to the concept of content-centric networks (CCNs) by Ted Nelson in the 1970s. After that, several projects such as TRIAD at Stanford and DONA from the University of California at Berkeley were carried out exploring the topic. In 2009 Xerox Palo Alto Research Center (PARC) released the CCNx project led by Van Jacobson, who is also one of the technical leaders of the NDN project.

The basic argument of the NDN project is that the primary usage of the current Internet has changed from end-to-end packet delivery to a content-centric model. The current Internet, which is a “client-server” model, is facing challenges in supporting secure content-oriented functionality. In this information dissemination model, the network is “transparent” and just forwarding data (i.e., it is “content-unaware”). Due to this unawareness, multiple copies of the same data are sent between endpoints on the

network again and again without any traffic optimization on the network’s part. The NDN uses a different model that enables the network to focus on “what” (contents) rather than “where” (addresses). The data are named instead of their location (IP addresses). Data become the first-class entities in NDN. Instead of trying to secure the transmission channel or data path through encryption, NDN tries to secure the content by naming the data through a security-enhanced method. This approach allows separating trust in data from trust between hosts and servers, which can potentially enable content caching on the network side to optimize traffic. Figure 1 is a simple illustration of the goal of NDN to build a “narrow waist” around content chunks instead of IP.

NDN has several key research issues. The first one is how to find the data, or how the data are named and organized to ensure fast data lookup and delivery. The proposed idea is to name the content by a hierarchical “name tree” which is scalable and easy to retrieve. The second research issue is data security and trustworthiness. NDN proposes to secure the data directly instead of securing the data “containers” such as files, hosts, and network connections. The contents are signed by public keys. The third issue is the scaling of NDN. NDN names are longer than IP addresses, but the hierarchical structure helps the efficiency of lookup and global accessibility of the data.

Regarding these issues, NDN tries to address them along the way to resolve the challenges in routing scalability, security and trust models, fast data forwarding and delivery, content protection and privacy, and an underlying theory supporting the design.

MobilityFirst — The MobilityFirst [4] project is led by Rutgers University with seven other universities. The basic motivation of MobilityFirst is that the current Internet is designed for interconnecting fixed endpoints. It fails to address the trend of dramatically increasing demands of mobile devices and services. The Internet usage and demand change is also a key driver for providing mobility from the architectural level for the future Internet. For the near term, MobilityFirst aims to address the cellular convergence trend motivated by the huge mobile population of 4 to 5 billion cellular devices; it also provides mobile peer-to-peer (P2P) and infostation (delay-tolerant network [DTN]) application services which offer robustness in case of link/network disconnection. For the long term, in the future, MobilityFirst has the ambition of connecting millions of cars via vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) modes, which involve capabilities such as location services, georouting, and reliable multicast. Ultimately, it will introduce a pervasive system to interface human beings with the physical world, and build a future Internet around people.

The challenges addressed by MobilityFirst include stronger security and trust requirements due to open wireless access, dynamic association, privacy concerns, and greater chance of network failure. MobilityFirst targets a clean-slate design

directly addressing mobility such that the fixed Internet will be a special case of the general design. MobilityFirst builds the “narrow waist” of the protocol stack around several protocols:

- Global name resolution and routing service
- Storage-aware (DTN-like) routing protocol
- Hop-by-hop segmented transport
- Service and management application programming interfaces (APIs)

The DTN-like routing protocol is integrated with the use of self-certifying public key addresses for inherent trustworthiness. Functionalities such as context- and location-aware services fit into the architecture naturally. An overview of the MobilityFirst architecture is shown in Fig. 2. It shows all the building blocks mentioned above and how they work together.

Some typical research challenges of MobilityFirst include:

- Trade-off between mobility and scalability
- Content caching and opportunistic data delivery
- Higher security and privacy requirements
- Robustness and fault tolerance

NEBULA — NEBULA [5] is another FIA project focused on building a cloud-computing-centric network architecture. It is led by the University of Pennsylvania with 11 other universities. NEBULA envisions the future Internet consisting of a highly available and extensible core network interconnecting data centers to provide utility-like services. Multiple cloud providers can use replication by themselves. Clouds comply with the agreement for mobile “roaming” users to connect to the nearest data center with a variety of access mechanisms such as wired and wireless links. NEBULA aims to design the cloud service embedded with security and trustworthiness, high service availability and reliability, integration of data centers and routers, evolvability, and economic and regulatory viability.

NEBULA design principles include:

- Reliable and high-speed core interconnecting data centers
- Parallel paths between data centers and core routers
- Secure in both access and transit
- A policy-based path selection mechanism
- Authentication enforced during connection establishment

With these design principles in mind, the NEBULA future Internet architecture consists of the following key parts:

- The NEBULA data plane (NDP), which establishes policy-compliant paths with flexible access control and defense mechanisms against availability attacks
- NEBULA virtual and extensible networking techniques (NVENT), which is a control plane providing access to application-selectable service and network abstractions such as redundancy, consistency, and policy routing
- The NEBULA core (NCore), which redundantly interconnects data centers with ultra-high-availability routers

NVENT offers control plane security with policy-selectable network abstraction including mul-

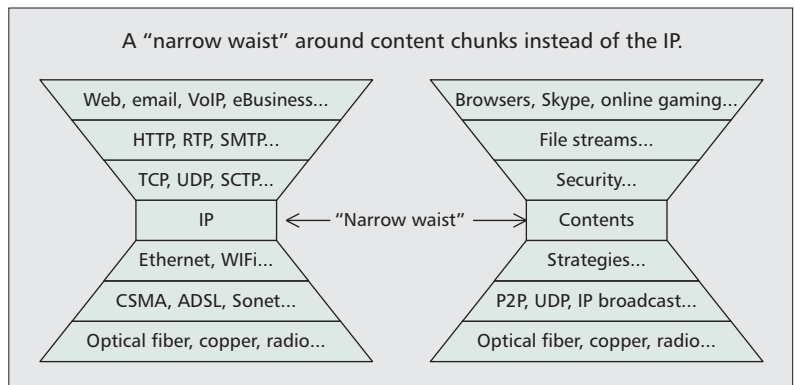


Figure 1. The new “narrow waist” of NDN (right) compared to the current Internet (left).

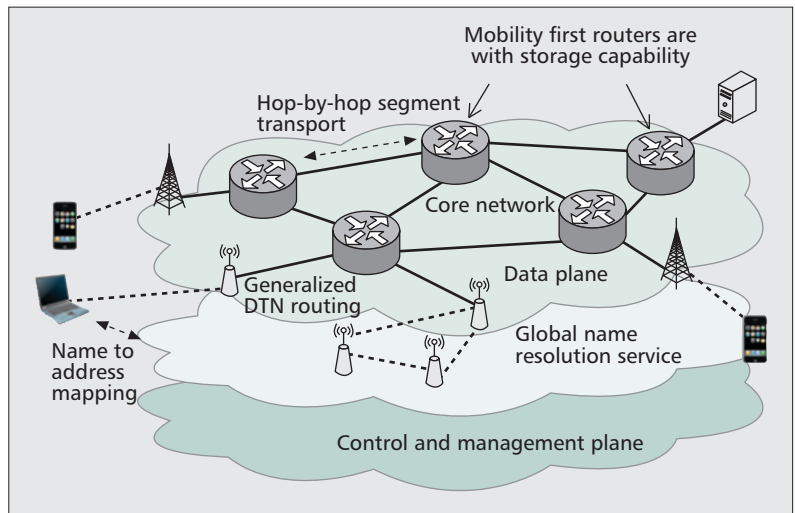


Figure 2. MobilityFirst architecture.

tipath routing and use of new networks. NDP involves a novel approach for network path establishment and policy-controlled trustworthy paths establishment among NEBULA routers. Figure 3 shows the NEBULA architecture comprising the NDP, NVENT, and NCore, and shows how they interact with each other.

eXpressive Internet Architecture (XIA) — Expressive Internet Architecture (XIA) [6] is also one of the four projects from the NSF FIA program, and was initiated by Carnegie Mellon University collaborating with two other universities. As we observe, most of the research projects on future Internet architectures realize the importance of security and consider their architecture carefully to avoid the flaws of the original Internet design. However, XIA directly and explicitly targets the security issue within its design.

There are three key ideas in the XIA architecture:

- Define a rich set of building blocks or communication entities as network principals including hosts, services, contents, and future additional entities.
- It is embedded with intrinsic security by using self-certifying identifiers for all principals for integrity and accountability properties.




Figure 3. NEBULA architecture components and their interactions.

- A pervasive “narrow waist” (not limited to the host-based communication as in the current Internet) for all key functions, including access to principals, interaction among stakeholders, and trust management; it aims to provide interoperability at all levels in the system, not just packet forwarding.

The XIA components and their interactions are illustrated in Fig. 4. The core of the XIA is the Expressive Internet Protocol (XIP) supporting communication between various types of principals. Three typical XIA principal types are content, host (defined by “who”), and service (defined by what it does). They are open to future extension. Each type of principal has a narrow waist that defines the minimal functionality required for interoperability. Principles talk to each other using expressive identifiers (XIDs), which are 160 bit identifiers identifying hosts, pieces of content, or services. The XIDs are basically self-certifying identifiers taking advantage of cryptographic hash technology. By using this XID, the content retrieval no longer relies on a particular host, service or network path. XIP can then support future functions as a diverse set of services. For low-level services, it uses a path-segment-based network architecture (named Tapa in their previous work) as the basic building block; and builds services for content-transfer and caching and service for secure content provenance at a higher level. XIA also needs various trustworthy mechanisms and provides network availability even when under attack. Finally, XIA defines explicit interfaces between network actors with different roles and goals.

GLOBAL ENVIRONMENT FOR NETWORK INNOVATIONS (GENI)

GENI [7] is a collaborative program supported by NSF aimed at providing a global large-scale experimental testbed for future Internet architecture test and validation. Started in 2005, it has attracted broad interest and participation from both academia and industry. Besides its initial support from existing projects on a dedicated backbone network infrastructure, it also aims to attract other infrastructure platforms to participate in the federation — the device control framework to provide these participating networks with users and operating environments, to observe, measure, and record the resulting experimental outcomes. So generally, GENI is differ-

ent from common testbeds in that it is a general-purpose large-scale facility that puts no limits on the network architectures, services, and applications to be evaluated; it aims to allow clean-slate designs to experiment with real users under real conditions.

The key idea of GENI is to build multiple virtualized *slices* out of the substrate for resource sharing and experiments. It contains two key pieces:

- Physical network substrates that are expandable building block components
- A global control and management framework that assembles the building blocks together into a coherent facility

Thus, intuitively two kinds of activities will be involved in GENI testbeds: one is deploying a prototype testbed federating different small and medium ones together (e.g., the OpenFlow testbed for campus networks [8]); the other is to run observable, controllable, and recordable experiments on it.

There are several working groups concentrating on different areas, such as the control framework working group; GENI experiment workflow and service working group; campus/operation, management, integration, and security working group; and instrumentation and management working group.

The GENI generic control framework consists of several subsystems and corresponding basic entities:

- Aggregate and components
- Clearinghouse
- Research organizations, including researchers and experiment tools
- Experiment support service
- “Opt-in” end users
- GENI operation and management

Clearinghouses from different organizations and places (e.g., those from the United States and European Union) can be connected through federation. By doing this, GENI not only federates with identical “GENI-like” systems, but also with any other system if they comply with a clearly defined and relatively narrow set of interfaces for federation. With these entities and subsystems, “slices” can be created on top of the shared substrate for miscellaneous research-defined specific experiments, and end users can “opt in” onto the GENI testbed accordingly.

GENI’s research and implementation plan consists of multiple continuous *spirals* (currently in spiral 3). Each spiral lasts for 12 months. Spiral 1 ran from 2008 to 2009; spiral 2 ran from 2009 to 2010; spiral 3 started in 2011. In spiral 1, the primary goals were to demonstrate one or more early prototypes of the GENI control framework and end-to-end slice operation across multiple technologies; there were five competing approaches to the GENI control framework, called “clusters.”

Cluster A was the Trial Integration Environment based on DETER (TIED) control framework focusing on federation, trust, and security. It was a one-project cluster based on the Cyber-Defense Technology Experimental Research (DETER) control framework by the University of Southern California (USC)/ISI, which is an individual “mini-GENI” testbed to demonstrate

federated and coordinated network provisioning. Cluster A particularly aimed to provide usability across multiple communities through federation. The project delivered software “fedd” as the implementation of the TIED federation architecture providing dynamic and on-demand federation, and interoperability across ProtoGENI, GENI-API, and non-GENI aggregate. It included an Attribute Based Access Control (ABAC) mechanism for large-scale distributed systems. It created a federation with two other projects: StarBED in Japan and ProtoGENI in the United States.

Cluster B was a control framework based on PlanetLab implemented by Princeton University emphasizing experiments with virtualized machines over the Internet. By the end of spiral 2, it included at least 12 projects from different universities and research institutes. The results of these projects are to be integrated into the PlanetLab testbed. PlanetLab provided “GENI-wrapper” code for independent development of an aggregate manager (AM) for Internet entities. A special “lightweight” protocol was introduced to interface PlanetLab and OpenFlow equipment. Through these mechanisms, other projects in the cluster can design their own substrates and component managers with different capacities and features.

Cluster C was the ProtoGENI control framework by the University of Utah based on Emulab, emphasizing network control and management. By the end of spiral 2, it consisted of at least 20 projects. The cluster integrated these existing and under-construction systems to provide key GENI functions. The integration included four key components: a backbone based on Internet2; sliceable and programmable PCs and NetFPGA cards; and subnets of wireless and wired edge clusters. Cluster C so far is the largest set of integrated projects in GENI.

Cluster D was Open Resource Control Architecture (ORCA) from Duke University and RENCi focusing on resource allocation and integration of sensor networks. By the end of spiral 2, it consisted of five projects. ORCA tried to include optical resources from the existing Metro-Scale Optical Testbed (BEN). Different from other clusters, the ORCA implementation included the integration of wireless/sensor prototypes. It maintains a clearinghouse for the testbeds under the ORCA control framework through which it connects to the national backbone and is available to external researchers.

Cluster E was Open-Access Research Testbed for Next-Generation Wireless Networks (ORBIT) by Rutgers University focusing on mobile and wireless testbed networks. It included three projects by the end of spiral 2. The basic ORBIT did not include a full clearinghouse implementation. Cluster E tried to research how mobile and wireless work can affect and possibly be merged into the GENI architecture. WiMAX is one of the wireless network prototypes in this cluster.

A more detailed description of the clusters and their specific approaches and corresponding features can be found in our previous survey [18]. Even more details can be found from GENI project websites and wikis [7].

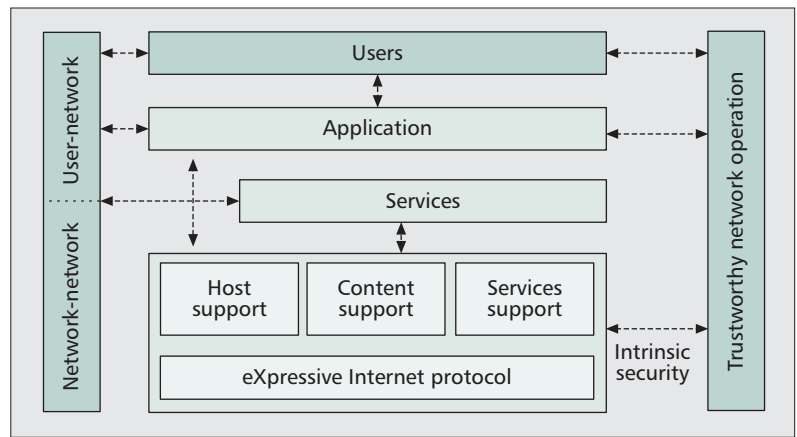


Figure 4. XIA components and interactions.

We can see that spirals 1 and 2 integrated a very wide variety of testbeds into its control framework. Spiral 2 was the second phase aiming to move toward continuous experimentation. Key developments include improved *integration* of GENI prototypes; architecture, tools, and services enabling experiment *instrumentation*; *interoperability* across GENI prototypes; and researcher *identity management*. In spiral 3, the goal is to coordinate the design and deployment of a first GENI Instrumentation and Measurement Architecture. Supporting experimental use of GENI and making it easier to use are also key goals. Also, more backbone services and participants are expected to join in the GENI framework for this spiral.

Another notable and unique characteristic offered by GENI is that instrumentation and measurement support have been designed into the system from the beginning since the ultimate goal of GENI is to provide an open and extensible testbed for experimentation with various new Internet architectures.

RESEARCH PROJECTS FROM THE EUROPEAN UNION AND ASIA

The European Union has also initiated a bundle of research projects on future Internet architectures. In this section, we introduce the research organized under the European Seventh Framework Programme (FP7) along with that in Japan and China.

EUROPEAN UNION

The European Future Internet Assembly [19] (abbreviated FIA as in the United States) is a collaboration between projects under FP7 on future Internet research. Currently, the FIA brings together about 150 projects that are part of FP7. These projects have a wide coverage, including the network of the future, cloud computing, Internet of service, trustworthy information and communication technology (ICT), networked media and search systems, socio-economic aspects of the future Internet, application domain, and Future Internet Research and Experimentation (FIRE) [10]. The FIA maintains a European Future Internet Portal [20],

A significant trait of the “Network of the Future” is that the research projects cover a very wide range of topics and a number of commercial organizations, including traditional telecommunication companies, participate in the research consortiums.

Categories	Project names (selected)
Future Architectures and Technologies	4AWARD, TRILOGY, EIFFEL, SPARC, SENSEI, Socrates, CHANGE, PSIRP, etc.
Services, Software, and Virtualization	ALERT, FAST, PLAY, S-Cube, SLA@SOI, VISION Cloud, etc.
Network Media	3DLife, COAST, COMET, FutureNEM, nextMEDIA, P2P-Next, etc.
Internet of Things	ASPIRE, COIN, CuteLoop, SYNERGY, etc.
Trustworthiness	ABC4Trust, AVANTSSAR, ECRYPT II, MASTER, uTRUSTit, etc.
Testbeds	FIRE, N4C, OPNEX, OneLAB2, PII, WISEBED, G-Lab, etc.
Others	HYDRA, INSPIRE, SOCIALNETS, etc.

Table 2. EU research projects on future Internet.

which is an important web portal for sharing information and interaction among the participating projects. Multiple FIA working groups have been formed to encourage collaboration among projects.

Of these projects, around 90 of them were launched following the calls of FP7 under the “Network of the Future” Objective 1.1. They can be divided into three clusters: “Future Internet Technologies (FI),” “Converged and Optical Networks (CaON),” and “Radio Access and Spectrum (RAS).” The total research funding since 2008 is over €390 million. A subset of the projects is shown in Table 2.

A significant trait of the “Network of the Future” [17] is that the research projects cover a very wide range of topics and a number of commercial organizations, including traditional telecommunication companies, participate in the research consortiums. Since there are a large number of projects, we selected a few representative ones and explain them in some detail. They are all under FP7 based on a series of design objectives categorized by the ICT challenge #1 of building “Pervasive and Trusted Network and Service Infrastructure.”

Due to the large number of projects, for the architecture research in this article, we selected a project named 4WARD (Architecture and Design for the Future Internet), and for the testbed we selected FIRE. We selected them due to the fact that FIRE is often deemed the European counterpart project to GENI, and the 4WARD project aims at a general architectural level of redesign of the Internet, and we feel that it is representative of the rest. It also involves a large number of institutions’ participation and cooperation.

In the following, we discuss these two projects briefly.

4WARD — 4WARD [9] is an EU FP7 project on designing a future Internet architecture led primarily by an industry consortium. The funding is over 45 million dollars for a 2-year period.

The key 4WARD design goals are:

- To create a new “network of information” paradigm in which information objects have their own identity and do not need to be

bound to hosts (somewhat similar to the goal of the NDN project)

- To design the network path to be an active unit that can control itself and provide resilience and failover, mobility, and secure data transmission
- To devise “default-on” management capability that is an intrinsic part of the network itself
- To provide dependable instantiation and interoperation of different networks on a single infrastructure.

Thus, on one hand, 4WARD promotes the innovations needed to improve a single network architecture; on the other hand, it enables multiple specialized network architectures to work together in an overall framework. There are five task components in the 4WARD research:

- A general architecture and framework
- Dynamic mechanisms for securely sharing resources in virtual networks
- “Default-on” network management system; a communication path architecture with multipath and mobility support
- Architecture for information-oriented networks

Note that 4WARD is one of many projects under the FP7 framework on future Internet architecture research. Readers can find more information on a complete list of the projects from [19]. Some typical projects focusing on different aspects of future architecture are listed in Table 2.

Future Internet Research and Experimentation (FIRE)

FIRE [10] is one of the European Union’s research projects on testbeds and is like a counterpart of GENI in the United States. FIRE was started in 2006 in FP6 and has continued through several consecutive cycles of funding. FIRE involves efforts from both industry and academia. It is currently in its “third wave” focusing on providing federation and sustainability between 2011 and 2012. Note that the FIRE project’s research is built on the previous work on the GEANT2 (Gigabit European Academic Networking Technology) project [11], which is the infrastructure testbed connecting over 3000 research organizations in Europe.

FIRE has two interrelated dimensions:

- To support long-term experimentally driven research on new paradigms and concepts and architectures for the future Internet
- To build a large-scale experimentation facility by gradually federating existing and future emerging testbeds

FIRE also expects not only to change the Internet in technical aspects but also in socio-economic terms by treating socio-economic requirements in parallel with technical requirements.

A major goal of FIRE is federation, which by definition is to unify different self-governing testbeds by a central control entity under a common set of objectives. With this goal in mind, the FIRE project can be clustered in a layered way as depicted in Fig. 5. As shown in the figure, it contains three basic clusters. The top-level cluster consists of a bundle of novel individual architectures for routing and transferring data. The bottom cluster consists of projects providing support for federation. In the middle is the federation cluster, which consists of the existing testbeds to be federated. These existing small and medium-sized testbeds can be federated gradually to meet the requirements for emerging future Internet technologies. Documents describing these sub-testbeds can be found on the FIRE project website [10].

ASIA

Asian countries such as Japan and China also have projects on future Internet architectures.

Japan — Japan has broad collaborations with both the United States and European Union regarding future Internet research. It participates in PlanetLab in the United States, and the testbed in Japan is also federated with the German G-Lab facility. The Japanese research program on future Internet architecture is called New Generation Network (NWGN) sponsored by the Japan National Institute of Information and Communications Technology (NICT). The Japanese research community defines the clean-slate architecture design as “new generation” and the general IP-based converged design as “next generation” (NXGN) design. NWGN started in June 2010 and expects to change the network technologies and Internet community with broad impact in both the short term (to 2015) and long term (to 2050). Like the projects in the United States and European Union, NWGN consists of a series of sub-projects collaborated on by academia and industry. The sub-projects range from architecture designs, testbed designs, virtualization laboratories, and wireless testbeds to data-centric networking, service-oriented networks, advanced mobility management over network virtualization, and green computing. Rather than enumerating all projects, we briefly discuss the architecture project AKARI [12] and the testbed projects JGN2plus [13] and JGN-X (JGN stands for Japan Gigabit Network). The reason we selected these projects is similar to the reason we selected FIA and GENI. AKARI is so far the biggest architectural research project in Japan; JGN2plus and JGN-X are the testbed research counterparts to GENI and FIRE.

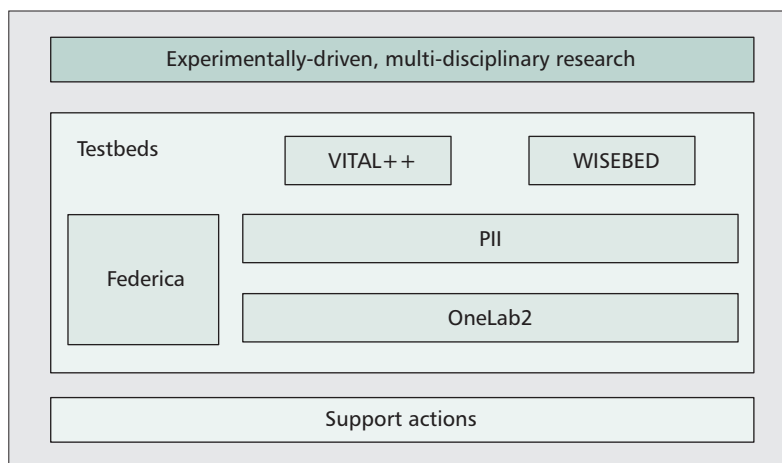


Figure 5. FIRE clustering of projects.

AKARI: AKARI means “a small light in the darkness.” The goal of AKARI is a clean-slate approach to design a network architecture of the future based on three key design principles:

- “Crystal synthesis,” which means to keep the architecture design simple even when integrating different functions
- “Reality connected,” which separates the physical and logical structures
- “Sustainable and evolutionary,” which means it should embed the “self-*” properties (self-organizing, self-distributed, self-emergent, etc.), and be flexible and open to the future changes

AKARI is supposed to assemble five subarchitecture models to become a blueprint NWGN:

- An integrated subarchitecture based on a layered model with cross-layer collaboration; logical identity separate from the data plane (a kind of ID/locator split structure)
- A subarchitecture that simplifies the layered model by reducing duplicated functions in lower layers
- A subarchitecture for quality of service (QoS) guarantee and multicast
- A subarchitecture to connect heterogeneous networks through virtualization
- A mobile access subarchitecture for sensor information distribution and regional adaptive services

AKARI is currently in the process of a proof-of-concept design and expects to get a blueprint in 2011. Through systematic testbed construction and experimentations, it aims to establish a new architecture ready for public deployment by 2016.

JGN2plus and JGN-X: JGN2plus is the nationwide testbed for applications and networks in Japan, and also the testbed for international federation. It includes broad collaboration from both industry and academia. It evolved as JGN, migrated to JGN II in 2004, and then to JGN2plus in 2008. From 2011, the testbed is under JGN-X, which targets to be the real NWGN testbed to deploy and validate AKARI research results. JGN2plus provides four kinds of services:

Any architecture that requires investment without immediate payoff is bound to fail. Of course, the payoff will increase as the deployment of the new technology increases, economies of scale reduce the cost, and eventually the old architecture deployed base will diminish and disappear.

- Layer 3 (L3) IP connection
 - L2 Ethernet connection
 - Optical testbed
 - Overlay service platform provisioning
- There are also five subtopics in the research of JGN2plus:
- NWGN service platform fundamental technologies
 - NWGN service testbed federation technology
 - Middleware and application of lightpath NWGN
 - Component establishment technologies for NWGN operation
 - Verification of new technologies for international operation

JGN2plus expects to create collaboration among industry, academia, and government for NWGN experiments. It also aims to contribute to human resource development in the ICT area via these experiments.

China — The research projects on future Internet in China are mostly under the 863 Program, 973 Program, and “12th Five-Year Plan Projects” administrated by the Ministry of Science and Technology (MOST) of China. Currently there are several ongoing research projects, which include:

- New Generation Trustworthy Networks (from 2007 to 2010)
- New Generation Network Architectures (from 2009 to 2013)
- Future Internet Architectures (from 2011 to 2015)

Project 1 was still IP - network research instead of a clean-slate future Internet. It consists of research sub-projects on new network architecture, next generation broadcasting (NGB), new network services, a national testbed for new generation networks and services, new routing/switching technology, a new optical transmission network, and low-cost hybrid access equipment.

Besides the research projects on future Internet architecture, there are also ongoing research projects for building a China Next Generation Internet (CNGI) testbed. It is based on the previous infrastructure network testbed of the China Education and Research Network (CERNET [14] and CERNET2 [15]) and the China Science and Technology Network (CSTNET). A terabit optical, terabit WDM, terabit router plus IPTV testbed called (3T-NET) was also announced on July 2009 as NGB. The testbed projects are mostly industry oriented with specific interest in IPv6 related protocols and applications.

We observed that the current future Internet architecture research in China leans heavily toward IPv6 related testbed, which is relatively short-term. To some extent, it reveals the pain China felt due to the collision between the extreme shortage of IPv4 address space in China and the ever expanding demands from increasing customers and novel services. Longer-term research projects on innovative architectural research are still in the cradle compared to those of the United States and European Union.

DISCUSSIONS AND PERSPECTIVES

Having presented a variety of research projects, we find that there are several issues worth discussing. In this section, we give our perspective regarding these issues. Of course, there is no agreement among researchers regarding these perspectives, and none is implied.

Clean-slate vs. evolutionary: Clean-slate designs impose no restriction and assumption on the architectural design. The key idea is not to be subjected to the limitations of the existing Internet architecture. It is also called “new generation” by Japanese and Chinese researchers. While the architectures can be revolutionary, their implementation has to be evolutionary. Today, the Internet connects billions of nodes and has millions of applications that have been developed over the last 40 years. We believe any new architecture should be designed with this reality in mind; otherwise, it is bound to fail. Legacy nodes and applications should be able to communicate over the new architecture without change (with adapter nodes at the boundary), and new nodes and applications should similarly be able to communicate over the existing Internet architecture. Of course, the services available to such users will be an intersection of those offered by both architectures. Also, the new architecture may provide adaptation facilities for legacy devices at their boundary points. Various versions of Ethernet are good examples of such backward compatibility. Some variations of IP are potential examples of missing this principle.

New architecture deployment will start in a very small scale compared to the current Internet. These early adopters should have economic incentives for change. Any architecture that requires investment without immediate payoff is bound to fail. Of course, the payoff will increase as the deployment of the new technology increases, economies of scale reduce the cost and eventually the old architecture deployed base will diminish and disappear.

Integration of security, mobility, and other functionalities: It is well understood and agreed that security, mobility, self-organization, disruption tolerance, and so on are some of the key required features for the future Internet. However, most of the projects, even for those collaborative ones like in FIA program, put more emphasis on a specific attribute or a specific set of problems. It seems to be a tough problem to handle many challenges in a single architecture design. Currently, for the collaborative projects such as FIA, they are trying to integrate miscellaneous previous research results into a coherent one trying to balance some of the issues. Although different projects have different emphases, it is beneficial to create such diversity and allow a bunch of integrated architectures to potentially compete in the future. However, we believe that there is still a long way to go before there is a next-generation architecture unifying these different lines of designs. For example, we observe that the four U.S. FIA projects concentrate on four different specific issues. Self-certifying and hash-based addresses are effective tools for security. However, securi-

ty needs much more consideration on both micro and macro scopes. Content- and information-centric features are also important trends, but how to integrate these differing requirements and resulting architectures is still a pending problem. We expect that more integration research will be required when such issues emerge in the future. It is therefore desirable for different projects to create some synergy for the integration process.

Architectures built around people instead of machines: It has been widely realized that the usage pattern of the Internet has changed, and the trend of building future Internet architecture around the contents, data, and users seems to be justifiable and promising in the future. Design goal changes naturally lead to design principle changes. Different patterns may emerge without any further synthesis. Current existing projects on future Internet architectures sort out different principles according to their own design emphases. From our perspective, it is essential and important to form a systematic and comprehensive theory in the research process rather than designing based only on experiences. It may take several continuous spirals between theoretical improvement and practical experience to achieve a sound architecture. We believe more research in this area may be desirable and meaningful for future Internet research.

Interfaces among stakeholders: Future Internet architectures are required to provide extensible and flexible explicit interfaces among multiple stakeholders (users, Internet service providers, application service providers, data owners, and governments) to allow interaction, and enforce policies and even laws. A typical example is Facebook, which creates a complex situation for data, privacy, and social relationships. Societal and economic components have become indispensable factors in the future Internet. The transition from the academic Internet to a multifunctional business-involved future Internet puts much higher requirements on the architectural supports to regulate and balance the interests of all stakeholders. In both technical and non-technical aspects, the future Internet architectures are required to provide extensible and flexible explicit interfaces among multiple actors to allow interaction, and enforce policies and even laws. The deep merging of the Internet into everyone's daily life has made such endeavors and efforts more and more urgent and important. From our perspective, significant research efforts are still needed in aspects such as economics, society, and laws.

Experimental facilities: Most of the current testbeds for future Internet architecture research in different countries are results of previous research projects not related to future Internet architectures. The networks use different technologies and have different capabilities. Although the federation efforts are meaningful, they may be restricted in both manageability and capability by such diversity. Testbeds from different countries are also generally tailored or specialized for the architectural design projects of those countries, with different features and emphases. Federation and creating synergy among such testbeds may be challenging. From

our perspective, such challenges also mean a valuable opportunity for research on sharing and virtualization over diverse platforms.

Service delivery networks: The key trend driving the growth of the Internet over the last decade is the profusion of services over the Internet. Google, Facebook, YouTube, and similar services form the bulk of Internet traffic. Cloud computing and the proliferation of mobile devices have led to further growth in services over the Internet. Therefore, Internet 3.0 [16], which is a project in which the authors of this article are involved, includes developing an open and secure service delivery network (SDN) architecture. This will allow telecommunication carriers to offer SDN services that can be used by many application service providers (ASPs). For example, an ASP wanting to use multiple cloud computing centers could use it to set up its own worldwide application-specific network and customize it by a rule-based delegation mechanism. These rules will allow ASPs to share an SDN and achieve the features required for widely distributed services, such as load balancing, fault tolerance, replication, multihoming, mobility, and strong security, customized for their application. One way to summarize this point is that service delivery should form the narrow waist of the Internet (Fig. 1), and content and IP are special cases of service delivery.

Internet 3.0, which is a project in which the authors are involved, includes developing an open and secure service delivery network architecture. This will allow telecommunication carriers to offer SDN services that can be used by many application service providers.

SUMMARY

In this article, we present a survey of the current research efforts on future Internet architectures. It is not meant to be a complete enumeration of all such projects. Instead, we focus on a series of representative research projects. Research programs and efforts from the United States, European Union, and Asia are discussed. By doing this, we hope to draw an approximate overall picture of the up-to-date status in this area.

REFERENCES

- [1] NSF Future Internet Architecture Project, <http://www.nets-fia.net/>.
- [2] NSF NeTS FIND Initiative, <http://www.nets-find.net>.
- [3] Named Data Networking Project, <http://www.named-data.net>.
- [4] MobilityFirst Future Internet Architecture Project, <http://mobilityfirst.winlab.rutgers.edu/>.
- [5] NEBULA Project, <http://nebula.cis.upenn.edu>.
- [6] eXpressive Internet Architecture Project, <http://www.cs.cmu.edu/~xia/>.
- [7] Global Environment for Network Innovations (GENI) Project, <http://www.geni.net/>.
- [8] OpenFlow Switch Consortium, <http://www.openflowswitch.org/>.
- [9] The FP7 4WARD Project, <http://www.4ward-project.eu/>.
- [10] FIRE: Future Internet Research and Experimentation, <http://cordis.europa.eu/fp7/ict/fire/>
- [11] GEANT2 Project, <http://www.geant2.net/>.
- [12] AKARI Architecture Design Project, <http://akari-project.nict.go.jp/eng/index2.htm>
- [13] JGN2plus- Advanced Testbed Network for R&D, <http://www.jgn.nict.go.jp/english/index.html>.
- [14] China Education and Research Network, <http://www.edu.cn/english/>.
- [15] CERNET2 Project, http://www.cernet2.edu.cn/index_en.htm.
- [16] Internet 3.0 project, <http://www1.cse.wustl.edu/~jain/research/index.html>.
- [17] The Network of the Future Projects of EU FP7, http://cordis.europa.eu/fp7/ict/future-networks/home_en.html.

-
- [18] S. Paul, J. Pan, and R. Jain, "Architectures for the Future Networks and the Next Generation Internet: A Survey," *Comp. Commun.*, U.K., vol. 34, issue 1, 15 Jan. 2011, pp. 2–42.
- [19] Future Internet Assembly, <http://www.future-internet.eu/home/future-internet-assembly.html>.

BIOGRAPHIES

JIANLI PAN [S] (jp10@cse.wustl.edu) received his B.E. in 2001 from Nanjing University of Posts and Telecommunications (NUPT), China, and his M.S. in 2004 from Beijing University of Posts and Telecommunications (BUPT), China. He is currently a Ph.D. student in the Department of Computer Science and Engineering at Washington University in Saint Louis, Missouri. His current research is on future Internet architecture and related topics such as routing scalability, mobility, multihoming, and Internet evolution. His recent research interests also include green building in the networking context.

SUBHARTHI PAUL [S] (pauls@cse.wustl.edu) received his B.S. degree from the University of Delhi, India, and his Master's degree in software engineering from Jadavpur University, Kolkata, India. He is presently a doctoral

student in the Department of Computer Science and Engineering at Washington University. His primary research interests are in the area of future Internet architectures.

RAJ JAIN [F] (jain@cse.wustl.edu) is a Fellow of ACM, a winner of the ACM SIGCOMM Test of Time award and CDAC-ACCS Foundation Award 2009, and ranks among the top 50 in Citeseer's list of Most Cited Authors in Computer Science. He is currently a professor in the Department of Computer Science and Engineering at Washington University. Previously, he was one of the co-founders of Nayna Networks, Inc., a next-generation telecommunications systems company in San Jose, California. He was a senior consulting eEngineer at Digital Equipment Corporation in Littleton, Massachusetts, and then a professor of computer and information sciences at Ohio State University, Columbus. He is the author of *Art of Computer Systems Performance Analysis*, which won the 1991 Best-Advanced How-to Book, Systems award from Computer Press Association. His fourth book, *High-Performance TCP/IP: Concepts, Issues, and Solutions*, was published by Prentice Hall in November 2003. Recently, he co-edited *Quality of Service Architectures for Wireless Networks: Performance Metrics and Management*, published in April 2010.