Next Steps for Internet2 Real Time Communications
DRAFT

Ben Teitelbaum (editor) *et al.* *

September 30, 2004

1 Overview

1.1 Goals

This document has three primary goals: to clarify the relationships between various Internet2 activities, working groups, and initiatives working on advancing human-to-human real time communications and collaboration technologies; to lay out an emerging set of architectural and strategic principles for Internet2 real time communications (RTC); and, thirdly, to assist Internet2 advisory councils and staff in aligning and prioritizing Internet2's real time communications efforts.

*This document is the product of an *ad hoc* Internet2 real time communications strategy group, whose membership included: Larry Amiot (Northwestern University), Dennis Baron (Massachusetts Institute of Technology), Nadim El-Khoury (University of North Carolina, Chapel Hill), Mike Enyeart (Indiana University), Jeremy George (Yale University), Michael Gettes (Duke University), Ken Klingenstein (Internet2), Walt Magnussen (Texas A&M University), Kathleen McMonigal (University of Washington), Cheryl Munn-Fremon (Internet2, Deployment and Infrastructure Delivery), Chris Peabody (Georgetown University), and Ellen Vaughan (Internet2). The numerous discussions of this group and the written contributions of some of its members are reflected throughout. Special thanks also to Candace Holman (Harvard University) and David Lassner (University of Hawaii), whose feedback was invaluable. The editor, Ben Teitelbaum, may be contacted at *ben@internet2.edu*, *sip : ben@internet2.edu*, or 3025 Boardwalk, Suite 200; Ann Arbor, MI 48108; USA.

Towards this final goal, the document concludes with several recommendations for re-organizing Internet2's RTC working groups and activities for greater strategic alignment and to address several emerging challenges and opportunities. In conclusion, the document also recommends several high-priority areas, where Internet2 effort and resources should be allocated over the next year.

1.2 Process and Status

This document is the product of an *ad hoc* group of Internet2 technical leaders and working group chairs. On September 26, 2004, a draft of this document was presented to the Applications Strategy Council (ASC). The ASC recognized the area of real time communications among humans as an important area in need of strategy and coordination. Additionally, the ASC embraced the creation of a "Real Time Communications Coordinating Committee" (RTCCC) to develop and preserve architectural and strategic principles for Internet2 real time communications. The ASC embraced none of the other recommendations herein; those are pending as input to the RTCCC.

1.3 Scope

A critical subset of Internet2 applications are those that enable human-to-human real time communication (RTC) and collaboration. These applications are general-purpose, meeting the one-to-one and multi-party personal communications needs of people, whether faculty or student, poet or physicist. Internet2 currently has more than half a dozen activities in this space spread across various working groups and areas.

These activities include the projects of the Integrated Infrastructure for Instant Messaging (I2IM) [1], Presence and Integrated Communications (PIC) [2], Video Middleware for Videoconferencing (VidMid-VC) [3], and Voice over IP (VoIP) [4] working groups, as well as the Internet2 Commons [5] effort and some projects of the ResearchChannel/Internet2 Working Group. In planning for the Internet2 Real Time Communications Forum [7], held at the Spring 2004 Internet2 Member Meeting, the leaders of these efforts began a series of discussions that has produced several outcomes:

- An increased awareness of significant common ground and strategic alignment among the activities;

- An initial attempt at a "grand unified architecture" of Internet2 real time communications;
- Consensus that alignment among the groups is not perfect and that it would be beneficial to have an ongoing forum that works to foster architectural and strategic coherence among Internet2's real time communications activities.

Because real time communications pose a distinct set of architectural and performance requirements, as well as a distinct set of deployment challenges, it makes sense to focus and to coordinate Internet2's activities in this area.

1.4 Relationship to Internet2 Collaboration Efforts

Network-enabled collaboration, especially inter-institutional collaboration, is the *raison d'être* for Internet2. Consequently, nearly every Internet2 activity either directly or indirectly addresses the needs of collaborative applications. This is often accomplished by deploying general-purpose networking resources and middleware services that can be leveraged by new applications.

Consider, for example, the work of the Internet2 Middleware Initiative. Having built much of the architectural foundation for inter-realm authentication and authorization, Internet2 middleware working groups are now focusing on scaling up campus deployments, on building trust federations, and on "Shibbolizing" a large number of collaborative tools, including Wikis, WebDAVs, teaching and learning suites, and list managers.

Similarly, Internet2's RTC activities need to build a sound, general-purpose architectural foundation for new real time applications and services. Such a foundation is necessary for enabling new collaborative applications, many of which will have real time media components.

Such a foundation is not, however, sufficient to support all of the needs of emerging collaborative applications. For example, non-real time collaborative applications (*e.g.* email, distributed file systems, calendaring, "groupware", content distribution networks) represent an important class of Internet2 applications with distinct architectural and deployment challenges.

These applications are, however, outside the scope of the Internet2 RTC architecture and the strategic plan presented herein[1].

There are some important junctures between real time collaborative applications and non-real time collaborative applications that call out for architectural alignment. Example junctures include: calendaring, unified messaging, and the naming of resources, people, and groups. These are noted and discussed in section 3. An overall strategy, architecture, and work plan for collaborative applications would be useful, but is beyond the scope of this document.

1.5 Document Structure

Section 2 surveys the current set of Internet2 RTC activities, including working groups, relevant projects, and Internet2 staff activities. Section 3 sets these activities in context by proposing an emerging "grand unified architecture" for Internet2 real time communications, while Section 4 proposes a set of strategic principles for Internet2 RTC, many of which are implied by the architecture. In Section 4.3, each of the current RTC activities is evaluated with respect to the proposed strategic principles.

Section 5 proposes a set of success criteria for evaluating RTC activities. General alignment with the evolving RTC architecture and strategic principles is but one component of this. Also important are the availability of sufficient resources, member involvement and enthusiasm, lack of duplication of effort, and other criteria. Again, each of the current RTC activities is evaluated—this time with respect to these success criteria.

The future of real time communications is with enterprise-based offerings and peer-to-peer technologies. How then does a membership organization and network service provider like Internet2 work to advance this future? Section 6 lists several models for moving Internet2 RTC beyond "sandbox" trials to production services. A future document may discuss these models and their trade-offs in greater depth.

Finally, the memorandum concludes with some thoughts on re-organizing

[1] Real time applications with one or more non-human endpoint represent yet another class of Internet2 application that is outside the scope of this document. Applications such as remote instrument control, remote simulation steering, or real time distributed computing may make use of some elements of the real time communications architecture presented herein. However, these applications tend to fill niche needs, rather than general-purpose needs, and are technically different enough to be regarded as out-of-scope.

Internet2's RTC activities and proposes several top priority goals for the next year.

2 Internet2 RTC Activities

2.1 Integrated Infrastructure for Instant Messaging (I2IM)

The I2IM Working Group, chartered in Fall 2003, is chaired by Michael Gettes from Duke University and is organizationally under the Middleware area under the direction of Ken Klingenstein.

The I2IM charter calls for the group to:

1. Research the issues surrounding enterprise integration of IM services with respect to identity, authentication and authorization;
2. Investigate current and future designs of relevant technologies with respect to authorization services, and how authorization information is transmitted and interpreted, especially between domains; and how to integrate these methods with emerging campus and inter-campus authorization systems;
3. Consider and possibly instantiate a national or international federated IM service. It will look at feasibility and applicability of leveraging technical methods like SAML/Shibboleth, and the policy frameworks coming from federations such as InCommon;

Accomplishments include the development of use cases, scenarios, and requirements for authenticated instant messaging, including cases where some attributes of users are authenticated (*e.g.* registered student), but the users themselves remain anonymous. The group is currently working "Shibbolizing" the Extensible Messaging and Presence Protocol (XMPP)—the protocol behind Jabber.

2.2 Internet2 Commons

The Internet2 Commons is a production, subscription-based H.323 video conferencing service, managed by Jonathan Tyman and under the direction of Cheryl Munn-Fremon. After running for a number of years as an experiment, the Commons Subscription Service was launched in October 2003. It provides standards-based H.323 video conferencing, as well as streaming and archiving of earlier conferences. The Commons is supported 24/7 by OARNet and Ohio State University and is guided by the Commons Management Team (a member-driven advisory committee).

The goals of the Commons are more ambitious than the current H.323 offering and include: "promoting and facilitating remote collaboration by means of innovative and integrated, standards-based Internet technologies"; and "building on useful technologies to create collaboration services that are sustainable, scalable, and affordable".

Commons leaders are currently surveying member needs, studying remote collaboration tools, and looking to expand the service suite to support data collaboration and shared workspaces.

2.3 Presence and Integrated Communications Working Group (PIC WG)

The PIC Working Group was chartered in Summer 2003 and is chaired by Jeremy George of Yale University. The working group seeks to "foster the deployment of SIP-based communication that integrate multiple communications elements in the context of presence," which is the notification of user level state information to facilitate communication. The PIC working group develops "technical deployments and use cases for campus presence and integrated communications services", and seeks "to inform the emerging policy tussle".

Accomplishments include the execution of three rich presence trials that use Internet2 meetings to prototype next-generation campus communications services. These trials enhance the meeting network infrastructure with a wireless location service, SIP routing, and a set of standards-based presence services integrated with back-end meeting functions. Meeting attendees are invited to participate by downloading one of several clients that leverage the meeting middleware and provide rich presence-enabled, integrated voice, video, and IM communications.

2.4 ResearchChannel/Internet2 Working Group

The ResearchChannel/Internet2 Working Group, chaired by Gates Rhodes (University of Pennsylvania) and Kathleen McMonigal (University of Washington), is chartered primarily to promote on-demand educational video content. However, the working group's "Big Video" project is focused on testing and evaluating high-end interactive video applications. The Big Video project has tested unicast and multicast 30Mbps DVTS video in a number of interactive meeting settings. The group is also evaluating 270 Mbps SDI video and has future plans to investigate plasma screen resolution video systems.

2.5 SIP.edu Initiative

SIP.edu has been incubated as a project under within Voice over IP (VoIP) Working Group. Having reached a critical mass of member interest, leadership, focus, and momentum, SIP.edu is emerging as a working group in its own right. SIP.edu is led by Dennis Baron (MIT) and is chartered to grow SIP connectivity within Internet2, promote the email address as a converged electronic identifier, and to build a community of campuses that are deploying and investigating campus SIP services. The SIP.edu project has published a "cookbook"[12] that provides low cost-of-entry "recipes" for campuses wishing to implement the initial SIP.edu architecture. The project has also partnered with Cisco and Avaya to sponsor new SIP.edu sites and has recently held the first SIP.edu Implementors Workshop (June 16th, 2004). At this time, SIP.edu has scaled to provide nearly 100,000 users with SIP reachability.

2.6 Video Middleware-Videoconferencing (VidMid-VC WG)

The VidMid-VC Working Group is chaired by Nadim El-Khoury (University of North Carolina at Chapel Hill) and is chartered to further the development of middleware for digital video and related areas. The group has focused on resource discovery, authentication, and authorization for point-to-point and multi-point videoconferencing. The primary accomplishment of the group has been the standardization of the H.350 [6] standard, which defines a directory services architecture for multimedia conferencing using LDAP. The group has demonstrated endpoint self-configuration and LDAP authentication using H.350 for both H.323 endpoints and SIP user agents.

The group has also investigated a "directory of directories" that would allow trans-institutional searching of campus person directories. Next on the group's agenda: a federated architecture for authenticated video conferencing.

2.7 Voice Disaster Recovery Initiative (Voice DR)

The Voice DR project exists under the umbrella of the Voice over IP (VoIP) Working Group and is led by Chris Peabody (consultant to Georgetown University). The group is chartered to investigate and deploy voice disaster recovery solutions that exploit the complementary strengths and vulnerabilities of the Internet and the PSTN to provide a very high level of voice survivability. The project has partnered with a hosted VoIP solutions provider (Broadsoft) and a carrier (PAETEC) to build a voice disaster recovery testbed that provides inbound and outbound fall-back to VoIP in the event that PSTN connectivity is compromised or overloaded.

2.8 Voice over IP Working Group (VoIP WG)

The VoIP Working Group, chaired by Walt Magnussen (Texas A&M University) and Mike Enyeart (Indiana University) is chartered to develop and deploy advanced voice communications, study the implications of network convergence, and improve the scalability, survivability, and functional richness of voice communications over the Internet.

A key accomplishment of the group is the design and operation of an H.323 VoIP peering testbed. This project enjoyed significant support from Cisco Systems and peered more than 20 campus Call Manager deployments. Although this project has formally concluded, the peerings and gatekeepers persist.

The VoIP Working Group serves as an umbrella for the SIP.edu and Voice Disaster Recovery Initiative. The group also organizes several workshops each year that provide opportunities to exchange best practices, come up to speed on recent technical and industry developments, and to experience the technology in a "hands-on" setting.

2.9 Technology Scouting

Technology scouting is a critical component of Internet2's work that does not fit neatly within the mold of a single working group or project. The aim of technology scouting is to facilitate partnerships across the Internet2 community that share and leverage promising new technologies. Knowing how to find, qualify, and apply new, externally-developed technologies in a manner that adds member value is key. Several Internet2 staff serve as technology scouts, spotting trends and developments at an early stage and finding technologies and partners that match our technology focus.

Technology scouting efforts seek to understand which emerging technologies might have the greatest and broadest value to Internet2 members. There is a bias in favor of technologies that maintain and contribute to the accessible, standards-based character of the Internet and in favor of technologies that promote appropriate, flexible, and easily-administered open source licensing arrangements.

The approach to technology scouting has been informal—more art than science and has included:

- Making inquiries, contacts and connections by attending events such as the SCxy super-computing conferences and the annual conference of the National Association of Broadcasters, where interesting new applications are often showcased;
- Making contacts within targeted early-adopter communities, seeking to understand their needs and how Internet2 might engage;
- Brokering relationships between developers and early adopters
- Testing and working directly with the development teams;
- Raising the visibility of promising technologies through demonstrations and outreach to targeted Internet2 working groups;
- Assessing whether a community responds and begins to adopt a technology; when it does not, the technology is dropped and the technology scouts move on.

A current challenge is to scale this informal process. Moving forward, Internet2 is working to develop sustainable processes in support of technology

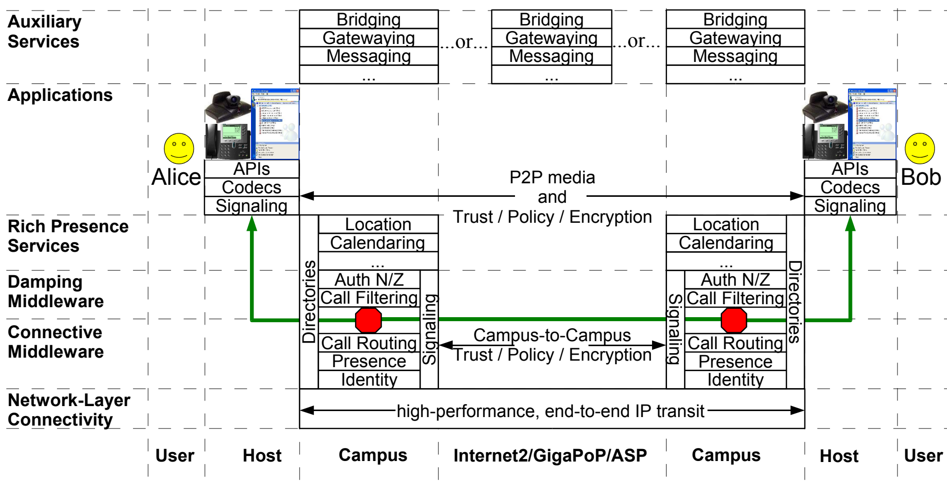

Figure 1: Internet2 RTC: Grand Unified Architecture (draft)

scouting. The desired state will include a proactive orchestrated effort and managed processes with appropriate linkage with strategic planning processes. This will include integration / coordination across areas allowing for organizational perspective. Known technology needs, principles, roadmap, and architecture will be articulated. Clear decision points and processes will be identified. Information on technical needs and findings will be shared.

3 Grand Unified Architecture

The most innovative and valuable real time person-to-person applications will leverage peer-to-peer design principles and campus / enterprise middleware. There are many parallels between RTC and email. Email is enabled through a combination of open standards (*e.g.* SMTP, MIME, RFC 822), host applications software (mail user agents like Eudora and Exchange, as well as auxiliary applications whose documents wind up as email MIME attachments), and enterprise middleware (*i.e.* mail servers and directories). Likewise, RTC will be enabled primarily by end-user agent software and campus middleware based on open standards.

An initial architecture for Internet2 real time communications emerged from the planning process for the Internet2 Real Time Communications Forum [7]. In subsequent conferences calls of the RTC group and in the preparation of this document, the architecture matured into what is described below and depicted in Figure 1. This architecture will be presented in greater detail in a forthcoming companion document.

Working from the bottom of the figure to the top, we see that RTC applications require high-performance **network-layer connectivity** to transport IP packets transparently from one host interface to another with low loss and low delay.

Investments in Internet2 backbone, regional, and campus LAN networks over the past five to seven years have resulted in exemplary network-layer connectivity. Even without network quality of service (QoS) assurances, typical end-to-end IP performance is more than adequate for most RTC applications, end-to-end transparency is high, and native IPv6 and multicast are increasing available. Even the best connectivity between host interfaces, however, is not sufficient to connect users to each other and to enable advanced real time communications. For that, additional middleware is required.

The first layer of enabling middleware is **connective middleware**. Users require electronic identifiers to which RTC calls may be placed. Campuses already manage the electronic identities and addresses of their users (*e.g.* email addresses, campus phone numbers, user names). To minimize confusion among communicating users and to provide authentication and other services to RTC applications, it is desirable for campuses to manage RTC identifiers as well. Campuses must also provide a basic location presence service, which, at a minimum, binds a user's electronic identifier to a set of applications and devices where the user may be reached. Finally, campuses must provide the ability to accept incoming call requests and to route them (according to local policy) to appropriate devices and applications.

Unchallenged much connectivity, however, can be a bad thing. Consider maximally connective middleware that always routed incoming call requests to a device or application that would interrupt the recipient and command his attention. This would be a nightmare. Recipients and call routers running on their behalf need the ability to block unwanted communications or to mitigate their impact (*e.g.* by routing such calls to a low-priority voice mailbox). Such **damping middleware** might leverage inter-realm trust federations like InCommon to authenticate and authorize incoming call requests or might employ anti-spam techniques like white lists, black lists, or any of various charging schemes to deter unwanted communications.

To enable communication at the best time and in the best manner, the networking infrastructure and directory assets of a campus may be exploited to provide **rich presence services** to users. For example, integration with an enterprise calendaring system might provide users with a convenient mechanism for automatically publishing information from their calendar (*e.g.* "in weekly XYZ project meeting") and for obscuring it (*e.g.* "in a meeting") for certain classes of watchers. A rich presence service might also leverage a campus location service to provide location presence (*e.g.* "in Michigan Conference Room; second floor; 3010 Boardwalk; Ann Arbor, Michigan") and to obscure it when appropriate (*e.g.* "in Ann Arbor"). Campus-based location services themselves might track users through the wired Ethernet infrastructure (as with Cisco's CDP), through the legacy WiFi infrastructure, or through the deployment of new beacons and sensors.

The connective middleware, damping middleware, and rich presence services just described all make heavy use of existing and new enterprise directories. Internet2's work in this space (*e.g.* eduPerson, H.350) may be leveraged to ensure that RTC's directory needs are met within the overall enterprise directory framework.

Applications are, of course, also essential. Applications may run on users' personal computers or, as is becoming increasingly common, on dedicated devices (*e.g.* portable communications appliances, IP phones, videoconferencing stations). To reach their full potential, RTC applications should exploit the connective, damping, and presence middleware described above. RTC applications should also exploit the capabilities of the network, which, in the context of Internet2, could mean using bandwidth-intensive media types, native IPv6, or native multicast. Applications should strive to excel with innovative user interfaces, media codecs, form factors, and P2P architectures (*e.g.* for multi-party conferencing and collaboration). Finally, it is essential that applications give users the means to manage both their "information privacy" (*i.e.* who sees what types of private information) and their "seclusive privacy" (*i.e.* when is a user available for real-time communication and who is authorized to cause an interruption).

A final "layer" in the architecture is a catch-all for **auxiliary services**. Central conference control and bridging services may be desired for multipoint conferences[2]; calls may need to be gatewayed to other networks (*e.g.*

[2]Note that multi-party conferencing does not necessarily require a central bridge or

to the PSTN); and, to handle cases when real time communications is impossible or deferred, there will need to be some juncture with messaging systems (*e.g.* voice mail, paging, and unified messaging). These services may be provided by the campus / enterprise, a traditional service provider, or a third-party applications service provider.

Finally, we note that the emerging RTC architecture makes heavy use of both **peer-to-peer trust** and **campus-to-campus trust**. Such trust fabrics are required to secure RTC through the authentication and authorization of call and presence routing and the encryption of media and signaling. Campus-to-campus trust relationships are required to enforce institutional policies and to leverage the role of the campus as a service provider and a manager of users' identities. Peer-to-peer trust relationships are necessary to enforce finer-grained, user-defined policies, such as those that are needed to protect users' information privacy and seclusive privacy.

4 Strategic Principles

As a membership organization and a network service provider, Internet2 faces a fundamental challenge in finding viable models for developing and deploying new campus-based services. This section proposes some strategic principles and guidelines for Internet2 RTC that advance the architecture described in section 3, making the most of Internet2's organizational strengths and principles.

These are strategic principles for organizing our RTC efforts for maximal effect and should be regarded as supplementary to Internet2-wide principles stated elsewhere. Although, some of the principles stated below are redundant with Internet2-wide principles (*e*.g. promote open standards), other important Internet2-wide principles are not reiterated because they are clearly stated elsewhere (*e*.g. promote a liberal intellectual property climate [8]).

4.1 Affirmative Principles ("Do"s)

Enable General-Purpose Communication The most successful applications are general-purpose—flexible enough to meet a variety of needs,

MCU. Successful multi-party conferencing systems have been designed that leverage native IP multicast or peer-to-peer overlay networks.

from the frivolous to the profound. Email, telephony, and word processing all have this quality. RTC applications have the ability to improve communications among people, whether they are gossiping freshmen or collaborating Nobel prize winning physicists.

Leverage Campus Directories Our members already manage the identities of their users and have rich directory assets that should be exploited to provide new value to RTC applications; when possible, new directories to support RTC applications should be integrated with campus directory infrastructure to avoid duplication and disintegration.

Leverage Campus Networks Because they run over high-performance campus LANs, all Internet2 RTC applications naturally leverage campus investments in network infrastructure. This principle, however, is about developing and exploiting campus networks in new ways. Examples of such strategic leverage might include running production voice over a converged LAN infrastructure, using the wired or wireless LAN to derive geospatial location, or using the LAN to power phones, network appliances, or other small devices.

Link and Converge Identities Future RTC applications must be user-centric and context-sensitive. Calls should be addressed not to devices, but to users. The addresses of users and of communication endpoints must be linked to provide seamless reachability, preferably at a single, converged electronic address. User identity may also be linked to roles or other attributes to provide anonymous, but authenticated, communications[3] or other semantically rich addressing (*e.g.* call a particular user by voice and video when she is next available for 30 minutes).

Grow Connectivity As "Metcalfe's Law" states: the value of a network is proportional to the square of the number of users. Consequently, growing the number of users reachable by new RTC applications is critical. Directory and call routing middleware must be deployed to give every Internet2 user an electronic identity, to maintain a binding between this identity and a set of application-level addresses where the

[3] Example uses include: a student seeking clarification on a professor's plagiarism policy who wishes to authenticate as a registered student for a particular course without revealing his specific identity; or anonymous suicide prevention hot line that needs to verify that repeated calls are indeed from the same caller without revealing the caller's specific identity.

user may be reached[4], and to provide the signaling and call routing for RTC applications and users to establish connectivity with each other. Connective middleware must also facilitate many-to-many group communication, by connecting individuals to “channels” or “rooms” where multi-party real time collaboration may occur.

Secure Communications RTC must be secured. While the pioneers of email may be forgiven for not having foreseen the problem of spam, we now have the foresight and technology to deploy RTC applications to prevent spam before it becomes a problem. Likewise with denial-of-service (DOS) attacks. RTC infrastructure is a high-value target that must be designed and operated to survive attacks from determined and intelligent adversaries. Finally, users expect a high degree of privacy for voice communications and, increasingly, for other real time communications. RTC must provide means to protect the privacy of users by encrypting communications, by scoping the dissemination of sensitive information, and by controlling who can interrupt whom and when.

Promote Advanced Media RTC should promote an exploration of advanced media types (*e.g.* high-bandwidth and multicast media streams) that leverage Internet2’s advanced network infrastructure. However, resource-intensive media types should never be promoted simply to justify the capacity of Internet2 networks; there must be a value proposition. Advanced media codecs should support compression and loss-tolerance. Most importantly, RTC applications must support congestion-responsive adaptation, as users will move through a variety of networking environments (*e.g.* wired Internet2 connectivity, campus wireless, NATed residential broadband) even within a single communications session.

Promote Open Standards To promote a competitive market of interoperable products, solutions, and services, it is a principle of Internet2 to promote open standards; likewise, it is a principle for Internet2 RTC.

Engage the Membership Internet2 represents a unique partnership of academia, industry, and government, committed to advancing internet applications. It is essential that Internet2 RTC initiatives engage not just the core membership (universities), but the membership

[4]Think of an application-level address as a triplet consisting of IP address, port number, and protocol.

broadly. Internet2's corporate membership has recently expanded beyond vendors, to include large enterprises and IT consulting firms. As IT consumers and integrators, these organizations have much in common with university members.

Displace New communications technologies often do not replace or fully inter-operate with earlier technologies—the FAX did not replace the postal system, email replaced neither FAX nor the postal system, and television did not replace radio. Instead, new communications technologies tend to displace the use of earlier technologies, existing side-by-side with them for a prolonged period.

It is essential, especially with voice applications where expectations are deeply entrenched, that we not limit the potential value of new applications either by forcing ourselves to match every feature and attribute of the earlier technology or by requiring full interoperability with legacy protocols or full access to legacy communications networks.

Unfortunately, it is often easier to fund the replacement of an existing system that has reached the end of its life-cycle or to fund incremental improvements on an existing system than it is to fund a wholly new service, system, or capability.

Study the Paths in the Snow Technology, policy, and business elites have a poor historical track record of predicting how users will use communications networks. Meanwhile, users are highly motivated to communicate with each other and have a history of finding value in unexpected ways (*e.g.* SMS, email, Napster, residential telephony)[9]. We should focus on building application-neutral and media-neutral middleware and then, rather than predicting what applications, media types, or services users will want to run on top of this, study what users (and the applications they use) are actually doing. This is colloquially referred to as "paths in the snow engineering"[5].

Develop Diagnostics As with all emerging technologies, good diagnostic algorithms, practices, and tools are required. Users and systems administrators will experience buggy code, interoperability problems, security compromises, poor performance, and pathological interactions with NATs and firewalls. It is especially important that we develop

[5] The reference is to a maxim about paving paths on a campus quadrangle: don't guess where the students will walk; wait until it snows, observe where they walk, and then pave your paths.

standards, best practices, and tools for diagnosing interdomain RTC performance problems. It is one of the unique strengths of Internet2 that we are able to cooperate across multiple domains to solve end-to-end performance problems.

4.2 Negative Principles ("Don't"s)

Don't Create Applications (Except When Necessary) Rather than create applications software, Internet2's RTC efforts should seek to develop and deploy enabling middleware. Host application software development in particular is difficult, expensive, and not our forté[6]. It is essential, of course, that new applications be developed that leverage Internet2 network and RTC middleware deployments. Internet2 should encourage third parties—including our corporate partners and open-source development foundations like the Open Source Applications Foundation (OSAF) and the Mozilla Foundation—to develop these applications . It also may make sense in some situations to partner with applications developers to understand and document the requirements of the higher-education market or to develop integrated proof-of-concept demonstrations.

Don't Centralize (Unless Necessary) As has been argued elsewhere in this document, the strongest opportunities to advance real time communications are at the edge (in host software and enterprise middleware). Although the centralization of some functions may be unavoidable (*e.g.* root name servers or certificate authorities), unnecessarily centralizing components of the RTC architecture will limit the potential value and scalability of advanced RTC and should be avoided whenever possible.

5 Success Criteria

The following criteria are proposed for evaluating current and future Internet2 RTC activities.

Architectural Alignment Is the activity aligned with the Internet2 RTC

[6] Server-side software development, which doesn't require slick GUIs and careful attention to human-factors, is another matter.

architecture?

Strategic Alignment Is the activity aligned with the strategic principles described above?

Viable Business Model Does the activity have sufficient funding or a viable business model for generating revenue or savings?

Adequate Leadership Does the project have a leader(s) with strong vision and project management skills?

Sufficient Support Does the activity enjoy adequate Internet2 support (*e.g.* scribes, flywheels, technical writers)?

Work Plan Does the project have a plan with specific deliverables and a time line?

Member Enthusiasm Does the project have strong "mind share" among members and a "critical mass" of member enthusiasm at the right levels (*e.g.* technologists, CIOs, faculty)? Is there demand for the project's deliverables?

Not Redundant Does the effort fill a niche within the spectrum Internet2 activities and is it not redundant with emerging commercial products and services or other higher-education initiatives?

Probability of Success Overall, how high is the probability of success for this effort?

6 Execution Strategy: from Sandboxes to Services

A number of possible execution strategies exist for Internet2 to organize to move new RTC technologies and applications from leading-edge technology trials, evaluations, and testbeds to production campus services. Models that might be employed include:

- To serve as a forum for developing, sharing, disseminating, and advocating best practices
- To help certain campuses to become deployment showcases

- To develop and promote open standards that leverage the strengths of our enterprises and to promote these new standards to vendors (*e.g.* H.350)
- To prototype new services (à la PIC), hoping to inspire campus deployments
- To build distributed consulting services that concentrate expertise and provide a for-fee consulting service to Internet2 members looking to deploy emerging RTC applications and middleware on campus
- To contradict ourselves and deploy some services centrally, acting like a benevolent carrier that provides services and guidance in the short-to-intermediate term, but that relinquishes control later on when campuses step up to the challenge themselves
- To organize as a community to develop requirements for campus-enabled applications and to partner with corporate or open source development efforts (*e.g.* Mozilla Foundation or Chandler) to develop and support these applications

7 Possible Next Steps

In conclusion, we recommend several high-priority "next steps" that will help align Internet2's RTC activities with the emerging RTC architecture presented in Section 3 and the set of strategic principles presented in Section 4. We believe that by focusing on these goals Internet2's RTC activities will be organized and better positioned for success.

7.1 Reorganize RTC Working Groups

To develop and preserve the Internet2 RTC architecture and strategic principles, to provide guidance to existing RTC activities, and to evaluate new ones, an umbrella advisory committee should be chartered, chaired, and staffed. This committee should be modeled after the Middleware Architecture Committee for Education (MACE) and might be called the Internet2 "Real Time Communications Coordinating Committee", abbreviated "RTCCC" or, for the truly geeky among us, "RT$(C)^3$".

The RTCCC would require a mixture of expertise. Master architects would be needed, as would senior IT managers, analysts, and engineers with deep practical experience deploying and operating real-time campus communications services. Like MACE, RTCCC should include the chairs of the various RTC working groups and should itself be chaired by someone who is deeply respected in the field and able to serve as an effective and fair-minded leader.

One of the first actions of the RTCCC should be to review the space of Internet2 RTC activities, looking not only at working group charters, but also at a set of success criteria (such as those proposed in section 5). This review may result in rechartering, realignment, and streamlining to ensure that Internet2 resources are allocated effectively and that Internet2's RTC activities are in-line with the broad strategic objectives for RTC.

7.2 Understand Campus Deployment Barriers

It has been argued earlier in this document that campuses are uniquely positioned to offer new, high-value, general-purpose integrated communications services. Justifying IT investments in these new services, however, can be a challenge. It is often easier to focus on replacing legacy systems and services (*e.g.* a legacy PBX) to reduce costs and make incremental improvements, than it is to embrace qualitatively different service models. This is especially true when disruptive technologies threaten organizational structures and demand new job skills.

It is essential that we begin a set of campus trials that engage real users and, by observing the "paths in the snow" and comparing notes with each other, work together to develop creative new business models for new RTC services. Launching such participatory campus trials has proven difficult[7].

We need to understand fully the deployment barriers faced by new RTC services in the campus environment. The barriers are clearly not with users, many of whom have adopted radically new communications patterns and defected from traditional campus communications offerings. Over the last 10-15 years, use of wireline campus telephony has been displaced by the

[7] For example, it has been the experience of the SIP.edu project that many campuses are willing to deploy an unsupported, low-visibility service (*e.g.* inbound SIP reachability through a legacy TDM switch), but are far more reluctant to support a trial that would allow users to register their own phones and soft-clients, despite the much greater value proposition.

use of cellular telephony[8]; over the last five years, many users have become avid users of proprietary instant messaging and presence services[9]; and, more recently, with the arrival of Skype, Free World Dialup, and others, use of wireline and wireless telephony has begun to be displaced by internet-based, voice-centric integrated communications. Rather than sit idly by while our users defect to third-party service providers, we need to test the conjecture that our users would be better served (and that research and education would be better served) if these services were offered by campuses themselves.

7.3 Complete Desktop Video AuthN/Z Work

It should be a strategic priority to confirm the VidMid-VC working group's functional specifications and use cases for desktop video authentication and authorization. We must also partner with Internet2 corporate partners or invest in open-source development efforts to grow the number of H.350-enabled clients and to verify interoperability. More work must also be done to enable media-neutral, SIP-based authentication and authorization leveraging Internet2's work on federated authentication.

7.4 Execute a Thorough Study of Voice Services

Internet2 members are increasingly expressing a desire to scale up their use of Internet2 network infrastructure for plain-old telephone service (POTS), including: inbound and outbound calling between member campuses (dialed as E.164 phone numbers); inbound and outbound calling from / to the PSTN; voice disaster recovery; and international long-distance. A priority going forward should be to convene a "Voice Services Advisory Committee" (VSAC) that will be charged with evaluating alternatives for Internet2 to provide voice services to its membership. The group will survey the needs and capabilities of the university membership, survey voice solutions and products offered by Internet2 corporate members, and propose one or more service bundles with a thorough evaluation of the business case for each.

An ideal outcome from this study group would be a voice service offering that coupled immediate or near-term savings to Internet2 members with a

[8] Despite poorer fidelity, reliability, and support for emergency communications!

[9] A 2002 survey by the Pew Internet and American Life Project found that 29% of college students use instant messaging as their "primary internet communication tool"[17].

longer-term strategic plan for bootstrapping campus-enabled advanced RTC services. This might require Internet2 to act as an "enlightened carrier", that strives to put itself out of business by helping its member customers take on increasing responsibility for running campus-enabled advanced RTC services.

7.5 Align the Commons

The Internet2 Commons Subscription Service, which provides H.323 video bridging, is in the midst of a big marketing push. The RTCCC should begin to work hand-in-hand with the Commons Management Team to ensure that current and future Commons offerings (*e.g.* collaboration services) are in alignment with the strategic and architectural directions of Internet2 real time communications.

A general process must also be developed for migrating experimental RTC services to production Internet2 services. This is not to say that all experimental RTC services should become production Internet2 services—indeed, most should become production campus services—but that when Internet2 does assume responsibility for new production RTC services, there be a well-thought-out process in place for service migration.

References

[1] Integrated Infrastructure for Instant Messaging Working Group (I2IM), http://middleware.internet2.edu/i2im/.

[2] Internet2 Presence and Integrated (PIC) Working Group, http://pic.internet2.edu/.

[3] Video Middleware for Videoconferencing Working Group (VidMid-VC), http://middleware.internet2.edu/video/.

[4] Internet2 Voice over IP (VoIP) Working Group, http://voip.internet2.edu/.

[5] Internet2 Commons, http://commons.internet2.edu/.

[6] H.350 Directory Services for Multimedia, International Telecommunications Union, Study Group 16, http://www.itu.int/ITU-T/studygroups/com16/h350/.

[7] Internet2 Real Time Communications Forum, Spring 2004 Internet2 Member Meeting, http://events.internet2.edu/2004/spring-mm/sessionDetails.cfm?session=1313&event=203, April 21, 2004,

[8] Internet2 Intellectual Property Framework, http://www.internet2.edu/resources/IntellectualPropertyFrameworkv2.1.pdf, April 9th, 2003.

[9] A. M. Odlyzko. The history of communications and its implications for the Internet, http://www.dtc.umn.edu/~odlyzko/doc/history.communications0.pdf, June, 2000.

[10] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. Johnston, J. Peterson, R. Sparks, M. Handley, and E. Schooler, "SIP: Session Initiation Protocol," RFC 3261, Internet Engineering Task Force (IETF), June 2002.

[11] Internet2 SIP.edu Initiative, home page, http://voip.internet2.edu/SIP.edu/.

[12] SIP.edu Cookbook, D. Baron (editor), http://mit.edu/sip/sip.edu/.

[13] EduPerson Object Class Specification, Internet2 Middleware Architecture Committee Directory Working Group (MACE-Dir), http://www.nmi-edit.org/eduPerson/internet2-mace-dir-eduPerson-200210.pdf, October, 2002.

[14] Shibboleth Project Home Page, http://shibboleth.internet2.edu/.

[15] InCommon Federation Home Page, http://incommon.internet2.edu/.

[16] SIP for Instant Messaging and Presence Leveraging Extensions (simple), IETF working group, http://www.ietf.org/html.charters/simple-charter.html.

[17] The Internet Goes to College: How students are living in the future with today's technology, Steve Jones, Pew Internet and American Life Report, September 15th, 2002.