

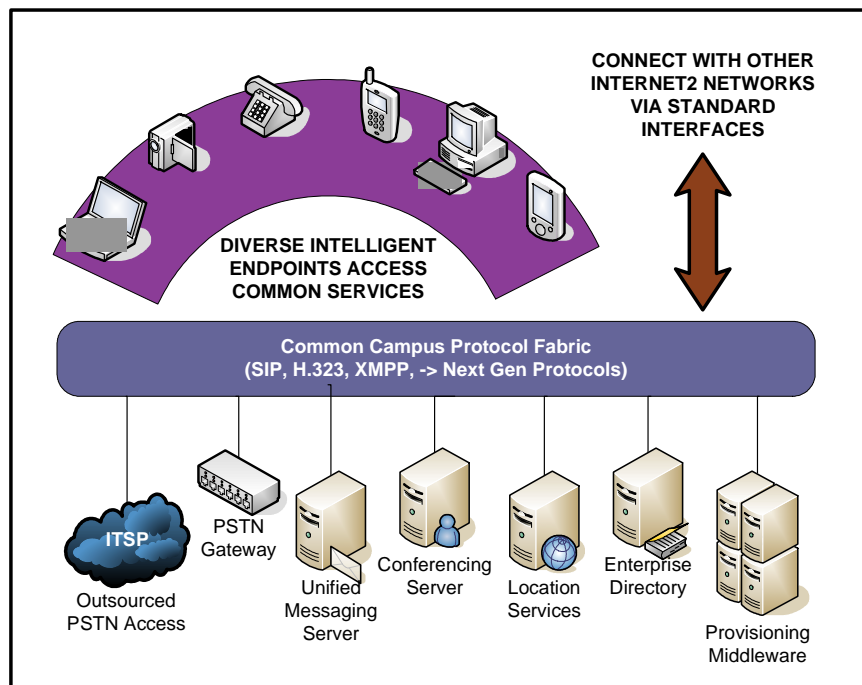
Reference Architecture



Real Time Communications (RTC) Next Generation Communications for Internet2 Campuses

The Internet2 Real Time Communications (RTC) Reference Architecture provides high level recommendations describing architectural components that Internet2 members should have in order to interoperate effectively with other Internet2 campus networks. It also identifies key technologies and components that will allow campuses to effectively develop and manage large scale RTC services and migrate those services forward as new protocols develop. This reference architecture describes near term functional specifications and long-term directions and should form a key part of a campus roadmap toward the implementation of RTC services.

The Internet2 RTC Reference Architecture enables a decomposed model so that campuses can implement best of breed components in a vendor neutral manner. Server hosted software systems integrate well into existing IT support models, but support for RTC applications requires an IT skill set that can handle system administration, software development and IP protocol analysis. Note the lack of monolithic telephony structures in the sample diagram at right.



By breaking the architecture into components, campuses may implement various features, such as VoIP, video conferencing, unified messaging, advanced call preferences, presence and location services in gradual ways that track developing user demand.

Recommendations

In the following recommendations, *must* refers to functionality that is critical for the secure operation and internetworking of RTC services between campuses, *should* refers to recommended best practices, and *may* refers to optional courses of action.

Call Signaling

Campuses *should* be prepared to operate multi-protocol RTC networks. The current strong direction is SIP and each campus *must* be capable of bidirectional SIP signaling between their campus and others as the primary communication protocol. H.323 remains important for video conferencing, but campuses *should* develop video conferencing infrastructures that support both H.323 and SIP, to capitalize on the strengths of H.323 while preparing for a migration of video services to SIP. Campuses *should* consider XMPP as well as SIP as protocols to support instant messaging. Finally, since developing RTC protocols will be more peer to peer oriented, campuses *should* already be preparing to migrate away from the protocols listed here to these more distributed protocols.

Middleware

RTC systems touch many different components within the IT infrastructure, including network management, the DNS, directories, identity management, accounting, location services, white pages, user provisioning, email and messaging, databases, monitoring and emergency services. The development and management of middleware components that manage this interconnectedness represent one of the major costs and operational considerations of RTC. Skill sets to support these activities include programmers with real time systems experience and high level systems administrators, many of whom are not traditionally found in telecommunications work groups. Campuses *should* have a strategy for developing and maintaining their middleware infrastructures.

Addressing

Campuses *must* implement DNS style addressing in the form of sip://user@university.edu. Campuses *may* implement E.164 addressing for interoperability with the PSTN but *should not* build their core architectures around numeric addresses. Campuses *should* separate their addressing architectures from their signaling infrastructure so that they can progress through changes in signaling protocols with minimal disruption to address spaces.

Authentication and Authorization

Campuses *must* implement the standard authentication methods supported by each protocol. For SIP, campuses *must* implement Digest Authentication. For H.323, campuses *must* implement one of the annexes in H.235. It should be noted that current standards and implementations generally only authenticate a user back to a central server. While this is common for IT applications, it is not a good long term solution for RTC. RTC applications will generally require that users be able to gain information about the identity of a calling party across network domains. This is important not only to minimize security risks, but also to be able to effectively offer services to individuals in other domains. Therefore, campuses *should* have a long term direction to implement products and standards that support inter-domain authorization, such as SAML from the OASIS group.

Directory Services and Identity Management

Directory enabled applications are common in IT, such as email, web services, and calendaring tools, but this is a new concept for RTC. "Directory-enabled RTC" means that users are provisioned and managed from a central directory (database) that provides information for all applications, thus eliminating redundant business processes for each application and leveraging the existing campus identity management infrastructure. This is a key factor in achieving the application flexibility and user self-provisioning features desired by customers, and the realization of cost savings through automated business processes. Campuses *should* use H.350 to directory enable RTC but *may* implement proprietary schemes if it can be determined that these do not impair flexibility and interoperability going forward.

Security: Encryption and Privacy

The area of encryption and privacy is a nexus of conflicting needs, policies, regulations and technologies. There are many legitimate reasons for users to encrypt their communications and ensure that only the intended recipients can access them. Some are regulatory, such as HIPAA, and others, such as financial or business information, may relate to the sensitivity of the information. In spite of this, a number of functions depend on access to the RTC communications. One such function is legal intercept, or the ability to undetectably wiretap a user. Another is the use inspection of packet contents to determine if the application is authorized (e.g. port blocking) or if the application is even RTC so that QoS can be applied. In general, once RTC communications are encrypted it is not possible for the network, law enforcement, or anyone else to understand or access that information in order to make decisions about it. Therefore, campuses *should* move away from simplistic packet inspection and signature analysis schemes for security and network policy awareness and develop more sophisticated approaches, such as relying on signaling information from call servers to identify authorized RTC streams in real time. It should be noted that current protocols do not generally address this issue well, and further that implementations in the market place address an even narrower view of privacy. Campuses should expect that shortcomings in this area may be a driver for movement away from current protocols and into next generation RTC protocols.

Security: DoS and SPAM Prevention

RTC endpoints are susceptible to DoS attacks like all other Internet hosts. Consider that in a large VoIP deployment a campus may have tens of thousands of IP telephones on the network. These are deployed and “managed” not by central IT, but individual departments who may not be aware of security threats or familiar with preventative measures like flash upgrading their phones to patched operating versions. Campuses, therefore, *should* have a strategy for managing large scale patch upgrades, and also be able to identify, segregate and protect RTC traffic in an emergency. Unwanted RTC traffic (SPAM) *should* be viewed as a serious potential threat. Campuses *should* pursue robust identity management architectures, especially those that offer end to end identity assertion as long term salves for this problem.

Disaster Recovery and Business Continuity

As technologies advance and converge the options to provide alternative or additional methods of traditional voice communications expands. Using converged communications technologies also provides better accessibility to those with disabilities and enhances the level of detail the communication can include. As new communications technologies evolve they may be interoperable with or independent of existing systems. It is critical that public safety and public communications departments are involved with the IT department in developing and supporting the overall communications architecture and communications plan. The communications plan needs to cover both business continuity events such as power outages, cut cables, and equipment failures and disaster recovery events such as natural disasters and terrorist activities. During a disaster recovery event the business continuity concerns are compounded by potentially life threatening situations and a greater chance of widespread extended outages of traditional communication services and other services including power and water. Campuses *must* ensure that this broad organizational involvement is secured.

Multipoint Conferencing

Campuses *should* plan to support media servers to mix audio, video and other media for rich conference calling environments. Campuses *should* also expect that this functionality is devolving to the endpoint, with IP phones and video phones including this capability internally for small conferencing applications. Therefore, campuses *should* consider a layered approach to multipoint conferencing in which very large conferences can be supported centrally, but routine conferences can be supported on the endpoint if the individual users value that enough to invest in that capability on their personal endpoints.

Data Collaboration Tools

Many campuses already have one or more data collaboration tool suites operational in support of individual workgroups or project areas. The current product choices available in the market are generally not interoperable with each other, thereby forcing campuses into choosing to support multiple "walled-gardens" within their own campus borders. Individuals participating in multiple collaborative communities are forced to become proficient in multiple data collaboration tool suites, potentially a different one for each and every community. Some recent attempts to select from the field of available vendor products produced mixed results and frustration with no clear "winner"; common themes in the evaluations include a field of dozens (or hundreds), with no clear selection strategy, and no consensus on standards for interoperability, i.e. any choice creates yet another proprietary collaboration space. Given this, the immediate-term selection of a data collaboration product *should* be governed, first by its appropriate match to strategic and/or project requirements, and second, by the preference for support of a majority of the campus adopted RTC architectural standards, as well as related media and storage standards that are broadly accepted in the market. As vendor, open-source or community-source products become available that provide broad-based interoperability between standards-based applications, campuses *should* look to vote with their investments in open and maximally interoperable product suites.

Presence

Presence is the notification of user-level state information to facilitate communication. Basic automatic presence (e.g. online, offline, idle) is part of most consumer instant messaging and presence services and services also allow richer presence information to be set manually (e.g. "at lunch", "on the phone", "in a meeting"). Going forward, location services will play an increasingly important role in rich presence. Campus-enabled integrated communications represents a unique opportunity to provide richer, automatic presence to improve productivity, enable collaboration, and enhance the online campus life experience.

Both SIP and Jabber/XMPP have mature (though not fully complete) standards for the publication and subscription of rich presence. Currently, Jabber/XMPP clients and servers are somewhat more mature and the Jabber/XMPP ecosystem is benefiting from the open federation efforts of Google Talk and others. Because of the potential benefits of this application, and because standards for presence are still emerging, campuses *should* initiate rich presence trials after establishing basic connectivity and *should* ensure that their technology roadmaps address the technical and policy implications of supporting user presence information on campus.

Location Services and Mobility

Mobility *should* be recognized as a key feature of an RTC service, so that users can freely move around campus and in their communities without provisioning changes. Campus *should* develop a location services architecture to support e911 and other developing location aware applications. Standards in this area are developing and campuses *should* follow the work of IETF's GeoPriv and ECRIT working group and activities in NENA. Campuses *should not* bind their location services architecture to a single protocol or application like SIP, but *should* instead build a generic location services architecture that can be re-used for different applications.

Accounting

Campuses *should* implement real time accounting interfaces to RTC systems to support accounting, traffic engineering, QoS signaling, location services, and CALEA and legal intercept applications. This accounting interface *should* be abstracted into a separate service from the core RTC architecture to allow the RTC architecture and the systems it touches to evolve without interdependencies.

Firewall / NAT Traversal

Campuses *should* have a strategy for utilizing standard protocols for firewall/NAT traversal. Campuses *should* deploy H.460.18/19 for H.323 services. Unfortunately, stable standards based firewall/NAT Traversal solutions have not solidified for SIP, so campuses *may* implement proprietary local firewall/NAT solutions but *should* be prepared to evolve this as better technologies emerge. A primary driver for NAT traversal is to be able to support users connecting from their home networks. However, this scenario may conflict with 911 policy if a campus public safety unit does not support off-campus service. Similarly, a campus networking department may not support devices connected to foreign networks. Therefore, firewall/NAT traversal architectures *should* be harmonized with policy.

Finding People and Services

Most campus identity management systems support regular campus faculty, staff and students. However, it is generally the case that RTC services must be provided to entities that fall outside of that narrow definition of identity. Examples such as a graduate library reference desk, a conference room, the history department, or collaborators at other institutions using campus RTC services all highlight the need for an expanded identity management service. Campuses *should* have an identity management infrastructure that can support affiliates and resources in addition to regular faculty, staff and student identities. Once these entities are represented in the directory, they can be searched and contacted, and can be managed like other RTC users. Campuses *should* also support the ability to appropriately hide identity where appropriate, such as when calling a suicide hotline, or if a caller ID should indicate "Pediatric Medicine" rather than an individual clinician.

Staffing and Operational Planning

Campuses *should* expect that the skill sets required to support RTC will lean away from traditional telephony technician skills and toward much higher level skills such as advanced systems administration, protocol analysis, security compliance, and IP networking. Thus, supporting RTC is much more like supporting email than supporting telephone service. Because of this, campuses *should* consider distributing support for RTC applications across their IT organizations, and should plan for evolving their staffing to meet the new operational requirements.

Application Clusters

RTC applications generally fall into the following application clusters:

- Business Telephony for VoIP
- Video Conferencing for collaboration and distance education.
- Data Collaboration for sharing of digital materials.

Campuses *should* recognize that each application cluster has different needs and therefore *may* identify specific products to meet specific application needs. However, campuses *must* ensure that there is a base level of interoperability among all application clusters by virtue of the fact that they share the same communications protocol infrastructure. For example, campuses *should not* implement data collaboration solutions that do not integrate with business VoIP.

Emergency and Regulatory

Emergency services infrastructures and regulatory frameworks generally lag the functionality provided by new technology, creating tensions between innovative service offerings and regulatory compliance. For example, mobility is a primary functionality that many users want from RTC, yet it breaks most existing e911 procedures. Further, the encrypted and peer to peer nature of SIP conversations make for highly scalable and inexpensive voice networks, but also renders effective court ordered wiretapping difficult or impossible. The National Emergency Number Association (NENA) is developing technical standards and guidelines for 911 services in the United States. Campuses *should* follow NENA's i2 and i3 architectures for IP based 911 services. Educause has been a leader in the development of policy around legal interception of VoIP and campuses *should* consider Educause positions when developing policy and infrastructure around wiretapping. Campuses *should* assume that policy shifts may be extreme going forward and *should* implement architectures that are policy-neutral whenever possible so that policy changes can be implemented by configuration changes and emphasis on expanded functionalities rather than changes to the core architecture.