

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/4044648>

Pushing the Internet direct to the user – security issues

Conference Paper · July 2003

DOI: 10.1049/jc:20030063 · Source: IEEE Xplore

CITATIONS

0

READS

19

1 author:



David McGovern

7 PUBLICATIONS 7 CITATIONS

SEE PROFILE

PUSHING THE INTERNET DIRECT TO THE USER – SECURITY ISSUES

D McGovern

BT Broadcast and Satellite Communications

INTRODUCTION

Recently, attention has been given to the possibility of providing Broadband service using broadcast media. Although the Internet has developed as a point to point network service, there are some limitations in its adaptation to large volumes, maintaining user performance expectations, and driving down costs. The broadcast media offer a means of tackling these problems, and also of increasing the takeup of multicast services in a way never fully exploited by the terrestrial Internet. A broadcast-based network, whether terrestrial or satellite, has the further advantages of service coverage to rural and remote areas, and a simple topology with the minimum of intermediate switching nodes between core network and user. Note that a return channel is still generally required, and this can be provided via the dial-up copper or wireless network which already supports access to the narrowband Internet.

The premise behind a commercial datacast business is that a significant proportion of Internet traffic volume is requested from a small number of sites by a large number of users. There are millions of web sites, but only a small proportion are accessed very frequently, and those exist for commercial purposes to reach the largest possible market. They have large volumes of content, for both diversity and to maintain high multimedia quality. It has been found by Nielsen

NetRatings (1), for example, that in typical developed countries some 25% of content is generated by the ten largest sites. Broadcasting of content allows the avoidance of a significant volume of repeated material, and gives corresponding cost benefits

In a broadcast-based network, multicast traffic will become the norm rather than the exception. Random browsing is supported in unicast, via the broadcast channel for large files, and via the narrowband circuit for delay-sensitive and low-volume traffic. Multicast traffic comprises the main volume, and takes two forms – streaming for all real-time multipoint content, and offline for all file transfer content. Any web site may have a combination of all three – for example browsing for obscure content or live voting, streaming for live video or audio news, and offline for video clips, MP3 audio files and file downloads of all kinds from images to software suites.

There are several challenges for the development of a large-scale broadcast-based system, although many of its components are already in place. A general system configuration is shown in Figure 1.

The net effect of pushing content to the user STB is to extend the edge of the network to the user premises. Any content which is chargeable or subject to licence restrictions must be held behind a secure interface until the correct authorisation has been entered by the user. In this paper we address the relevant security aspects, first in terms of defining the problem, then outlining the path to the solution.




Figure 1 Broadcast Broadband System Model

THE SECURITY PROBLEM

The Internet has evolved from the military (where all sites are trusted and secret security is a priority) and the academic world where information tends to be open. In the migration to the commercial world, accessible security has lagged the development of universal connectivity and open access. The Internet is essentially an open network, and it will remain so.

User expectations are generally that Internet-delivered content is free, placed there for the purpose of disseminating information. An increasing proportion of free content is there to develop user interest to purchase, but the recent bursting of the dot com bubble verified that a firm basis for revenue-based content is still required.

One part of the equation is the ability to secure point to point interactions, for example making purchases over the internet using credit cards. Existing standards operating at the transport layer, such as SSL (Secure Socket Layer) are familiar to us as https:// and are increasingly trusted by customers. Most of the fraud problems appear related to a lack of secure or honest storage of customer data by unknown vendors, rather than transmission insecurity. Another standard point to point security system in place is IPsec, applied at the IP (network) layer, allowing users the choice of very high levels of security, in exchange for more effort in creating the association. In general, standards-based point to point security is sufficiently mature to meet its purpose, and will be given a reference function in this discussion of multicast security.

Two forms of multicast traffic are defined in the pushed Internet – streaming and offline. Streamed content may be live (and usually has a limited shelf-life), or may be playout of pre-recorded content. It is intended to be used in real time, that is, at the time of reception only. If user recording of streamed content is an issue for content owners, they need to examine options to manage this. Offline multicast content comprises files for local storage and use some time after reception. Security of offline content is considered a separate activity from securing streaming content, because of the distinct problems it poses for Key Management.

Streamed Services

Streamed services can include news, price or sports text ticker bars, radio channels and voice commentaries, motorway or other camera views, videoconferencing, all the way up to full-quality TV. Appropriate sound or picture encoding is used to ensure that the minimum bandwidth is taken for the quality desired, and this can be selected dynamically in

the IP environment. As an example, Windows Mediaplayer 9 is capable of decoding most multimedia content with almost state-of-art decoding performance. Any streamed material with a saleable value should be encrypted to allow the owner to recover that value. Encryption is a multi faceted process, its functions defined to meet the required purposes:

Transparency – operation without restricting connectivity or access to users, and without limiting transmission options such as performance enhancers or data accelerators

Security of data and “data leakage” – it must be sufficiently difficult to attack either the encrypted datastream, or encrypted keys, that the cost of doing this is greater than the value of the attack. It depends on the key management system, and the algorithm selected for the application, but must be matched to practical means to protect keys at the content provider and in each user STB.

Security of keys – since the user premises cannot be regarded as “trusted”, it must be sufficiently difficult to obtain keys from the user STB that the effort is greater than the value. Since a multicast session is encrypted with one traffic key at a time, it follows that the discovery and quick rebroadcast of the key would compromise the service

Open-access – security must be entirely in the key generated for each session, the algorithm is open for any manufacturer, so that proprietary relationships between content owners, carriers, STB manufacturers and users are totally avoided

Join and leave – in a multicast service subscribers may join midway through, or may leave (in which case their access has to be suspended), while keeping continuous and uninterrupted service for the others

Scaleable – Multicast transmissions may presently target some thousands, but in the broadcast Internet we should plan for millions of users. The join and leave activities involving large-scale key distribution are the main issues here

Authentication – a system should guard against an “imposter” substituting his transmission for a legitimate one

Non-repudiation – a subscription manager needs to verify that a user cannot deny receiving content that he has received. It may not be sufficient to verify that he could have received it. The manager should be able to call upon a database of recipients with positive verification, or to identify exceptions.

We may observe that the characteristics of multicast security differ from those of unicast (point to point) security. Without making a line by line comparison, the multicast requirements apply for a particular purpose, but do not generally extend to the degree of confidentiality required for unicast applications like banking details, card transactions, transfer of secure documents etc. Instead, we should consider the value of the information and the opportunity to acquire it by other means.

In the residential market, much of the likely streamed content is in the public domain, in the sense that it is being made available to many subscribers, or that copies of recorded material (movies, TV programmes) are already in circulation. The value of the subscription is therefore in the immediacy of reception and in the convenience of its delivery direct to the user. The required strength of security is therefore relatively low, and is equivalent to that already employed in Conditional Access systems for digital TV.

In the business or corporate markets, stricter requirements apply. A multicast session may indeed contain information of a commercially or legally sensitive nature, and this will define the level of security capability that the system is specified to achieve. However, there is a limit to the degree of control that an enterprise can have over a number of its sites, where there may be unauthorised, disaffected or transient attendees. The system should be specified to a point where multicast sessions above a certain level of confidentiality are handled in a manner appropriate to the site conditions, for example using multiple point to point sessions, or a special-purpose proprietary system. This would place an upper limit on the strength of standard multicast security required.

Offline Services

Most Internet content is already offline, in the sense that it comprises files stored on web servers transmitted on request for storage in the PC file area known as "Temporary Internet Files". They can be selected to be available for local browsing if the option "Work Offline" is selected and the URL is bookmarked with this setting. Also, the user has the option to save any file to a nominated folder.

In the broadcast Internet this process is developed to employ various means to anticipate required content; to find, recognise and store it according to the user profile, and to enable the browser to search the local file store before searching the web. The use of high-capacity broadcast media, and user hard drives in the order of 100 – 200GB, allows storage of a large number of web pages and a large quantity of media files such as video and audio. These represent a major growth area and the driving force for this technology.

File transfer is currently the largest reported consumer of Internet capacity, in the form of peer to peer connections for sharing of music files. The music publishing industry appears to be reluctant to commit to multicast for music distribution, because of the potential for files immediately to be "shared" with other users. These are considerations for the developing Digital Rights Management (DRM) industry, to enable copyright owners to realise revenue from each transfer, whether via central distribution or peer to peer. The broadcast model for the Internet should enable control to revert to the rights owners, so that peer to peer connection can be minimised.

DRM is a means of encrypting files before transmission and user local storage, so that the files can be decrypted by the holder of a valid key under defined commercial conditions. It therefore comprises both an encryption/decryption process, and a process of key management linked to subscription management.

The main purpose of DRM is to enforce purchase of content in the mass market, by enabling legal purchase, and making illegal copying or forwarding uneconomic or unsafe. The main enforcement role is in the consumer space where illegal practice has become a matter of routine (almost of de-facto right!) In the realm of content owners, distributors and carriers it is expected that legal processes alone will be used to protect material at source.

Each file is encrypted with a key. The key to enable users to decrypt is provided in a licence package, implemented in software, which includes the conditions of use, for example the dates of validity, the number of times it may be played, or the feature set of a computer program. The conditions can extend to stipulation of a process whereby, if a file is copied or forwarded, a fee can be collected and transferred to the rights owner. This could be a means to legitimise peer to peer forwarding and may be useful for cases where material is not available directly from a distributor.

THE SECURITY ENVIRONMENT

We can place some bounds on the multicast security environment:

Content provider model. Initially broadcast content may be aggregated and encrypted at central network nodes, but the multicast model assumes that content providers will develop their own playout facilities close to source and applying their own security. Playout sites will then be equivalent to large web sites which may be very numerous, geographically dispersed, hosted by site owners or under various hosting outsourcing arrangements. Content, ready packaged in

broadcast form, is fed to nodes via appropriate IP, Frame Relay or ATM connections.

The present model for TV content providers, with a few large broadcasters operating their own proprietary Conditional Access systems in a vertical market structure, is no longer appropriate for datacasting. For a flat, open, web-like content host network, there must be a single-family open-access security standard.

Layer in stack Arguments about the best layer in the stack to operate standard broadcast security could reach philosophical levels, and a detailed discussion is outside the scope of this paper. Instead we present a pragmatic view of some aspects.

The first candidate is Layer 2, the link layer. It is appropriate in the current DTV environment where Conditional Access is applied to traffic on the DVB bearers ie the satellite link, decrypted at the user STB. The backhaul connections to studios and playout centres are secured to a satisfactory level by other means within the Operator's control. At the user end there is still a 1:1 relationship between STB and TV set, since the problems of secure on-premise distribution via a local link have not yet been seriously addressed.

In an open IP-based broadcast service security needs the possibility to be applied end to end, and Layer 3 is the lowest layer which allows this.

Hence the second candidate for applying broadcast service security is Layer 3, the IP layer. It has been selected by IETF for the open IPsec standard, for its support for end to end security (on a point to point basis), the IP layer being the most universal of all. However, IPsec places restrictions on network routing by its required use of tunnels, and is not compatible with transmission-improving expedients such as Performance Enhancing Proxies (PEPs) which generally operate at a higher layer. In multicast the concept of tunnelling does not apply, but the problem with PEPs may still be a serious one for some satellite applications, where performance will be reduced if PEPs must be switched off to allow IP-layer security.

The third candidate is Layer 4, the transport layer. It is an appropriate layer to secure point to point systems, examples being Secure Socket Layer (SSL) and Secure Shell (SSH). In multicast Layer 4 has no particular value, since the function of TCP to "guarantee" packet delivery has no parallel in multicast.

The fourth candidate is the upper layers bundled together as applications-related. Although applications vary widely, they all have access to the forward broadcast channel, the return channel, and the security Key Management process which is the core of the

security process. By definition, applications are run on the end user devices, so are best-placed to provide the user interface for personal logon, user verification etc (which may extend to signature or iris-recognition).

Selection of the appropriate layer to encrypt multicast will be the subject of further work, but it already appears that the application layer is favourite, followed by Layer 3.

Trust of user sites Some security architectures are designed on the assumption that user devices are secure, either in their physical construction or in the environment in which they are situated. In the military, for example, this is a correct assumption, because security of decrypted data must be ensured as well as the transmission keys.

In the public network, however, the concept of protecting decrypted data does not apply. Since decryption will be done in various devices like PCs and PDAs (ie not only in user STBs), it is no longer correct to assume that the user device has a degree of physical security in its construction.

This aspect needs to be carefully factored into the security specification. Real-time transmitted data is generally decrypted by the application software at the point of use, generally a physically insecure device, where key security is normally ensured by regularly changing it. Stored file security may need keys to have a longer life, in which case a degree of physical security is required to protect the keys from being copied.

Frequency of key changes Present Conditional Access schemes use smartcards which may be changed quarterly or annually in a significant logistics effort, to be repeated early if a successful attack occurs. In the broadcast Internet model, in the general case user terminals have a return channel, available to support the regular update of keys. Given the more relaxed physical security of user devices, and the diversity of user devices, the objective is to change keys sufficiently often that it would not be worthwhile to attempt illegally to obtain and distribute keys.

The potential large scale of the service, with users for a session possibly in the millions, means that thousands of users are likely to join or leave the session part-way through. The schedule for changing keys is therefore determined by the key lifetime rather than the rate of joining and leaving. For example, if the key lifetime is determined as 30 minutes, then subscriptions will be geared to be valid for up to 30 minutes after a user declares they are leaving, with content still received during this period. In practice a dynamic process will evolve, transparent to the users, and able to accommodate various types and duration of sessions.

IMPLEMENTATION

Real-time streaming

Work is proceeding in the MSEC Group of IETF to develop standards for securing multicast streaming content. The approach taken is not specific to a particular layer, so solutions are expected to apply at the network or application layer, as appropriate.

The structure needs to define an integrated and coherent solution to the requirements. Proposals for a true open standard (security in the keys, not the algorithm or proprietary hardware) have been discussed in IETF, but only to a limited extent. RFC 2627 (2) covers matters relating to multicast service over the internet, but in the context of small closed services rather than the mass market. Recent indications are that MSEC is tackling the general issues of key dynamics, reliability of key distribution, and stream authentication, but are not yet addressing issues such as key leakage, hosting business models, or non-repudiation. An observation would be that the commercial model for multicast via the Internet has limited scope, and is not attracting sufficient interest from those working in various broadcast security fields to develop a complete, coherent family of standards.

The main requirements parameters for security of a broadcast/multicast mass-market streaming service are discussed below in the context of solutions developed by BT and other organisations.

Transparency and data security A security system needs to be transparent in the sense that it is applied with the minimum disadvantage to legitimate users of content. The user registration, authentication and billing process must be as straightforward as possible, and encryption and decryption should take place with the minimum transmission delay or susceptibility to interruption (for example during rekey synchronisation).

The core process is the generation of a new Traffic Encryption Key (TEK) at the transmission source for each "billable" session, and its secure despatch to all receivers. Data is encrypted at source by the TEK, and decrypted at each user by the same key. The multicast/broadcast channel may be used to distribute the key to each receiver, and we stipulate that there is a return channel available by which any part of the process can be selected or verified. Broadcast-only reception (ie with no return channel) is possible but is an exception for which limited solutions may apply.

The process of TEK generation is not covered here. An important output of the process is the length of the TEK, which may have an impact on the capacity and time needed to distribute it. 128 bytes is a typical TEK length, but this may vary considerably.

To send the TEK securely it must itself be encrypted, and for this an agreement is made between the source and each user on a Key Encryption Key (KEK). In order to manage the various problems involved in large multicasts, the KEK is part of a hierarchy of frequently-changed keys, but is based on unique information about each user subscription known to the source. Protection of subscriber information is important but is outside the scope of this paper.

Solutions to the various complications of large-scale secure multicast are discussed below:

Key management, non-trusted user environment If an Internet user can obtain a valid TEK, he is in a position to rebroadcast it (in the clear) and thus undermine the content subscription value. Since reception and decryption of streamed content is likely to take place in PCs (rather than an STB), there is the potential for decryption to be done in software, and for attempts to be made to discover the TEK by dedicated applications also running on the PC.

Multicast service parameters In point to point sessions all data including key exchange is bilateral and data keys are usually valid for the duration of the session (though can be changed if desired). In multicast sessions users who leave a session and are no longer entitled to take part force the changing of the TEK, and a multicast with a flat Key Management system and a large number of subscribers would have a huge rekey problem. The solution is described in a number of IETF MSEC Internet Drafts.

Of particular interest here is the Group Secure Association Key Management Protocol (GSAKMP) (3) which includes hierarchical key dissemination, authentication of users, group key generation, and access control. (NB GSAKMP Light has superseded GSAKMP.) It defines a number of virtual, logical nodes in a binary tree structure, such that each user has a unique set of nodes and corresponding KEKs with which to encrypt the multicast session TEK. When a new session is initiated the number of KEKs is proportional to the natural logarithm of the number of users, so that typically a session with 1 million users, instead of 1 million KEKs, requires only around 40. Compared to the traffic volume of a valuable multicast this is a relatively small overhead. GSAKMP has been implemented by LogicaCMG, Sparta Inc and others.

Authentication The IETF MSEC group has addressed the subject of secure source authentication for large-scale multicast and broadcast data streams. TESLA has been described in an Internet Draft (4) as providing multicast source authentication of individual data packets, regardless of the packet loss rate, and with low overhead.

Charging procedures Work has been done in BT and other organisations to a detailed level on charging procedures suitable for an open standard. This does not yet appear to have been incorporated into the MSEC or other IETF mandates.

Offline files

Security of offline content has been treated separately because it presents some particular issues, mainly around the delay between encryption and decryption, and because the MSEC work has covered mainly streaming services. The Digital Rights Management (DRM) forum in IETF has been dormant for nearly 18 months while delegates reconcile the need for open standards with the proprietary nature of the content licensing process.

DRM is not itself restricted to protect offline material. In the case of video, for example, the same player can be used to show real-time streamed or locally stored content. An open-standards path is apparently available via MSEC for streaming security, whereas the proprietary route to a small number of de-facto standards is the likely present route to offline security. The content industry must decide which way to go, and it may be stimulated by the development of the data broadcast channel.

DRM products are available on the market with the following main components; encryption of content, management of rights, and a licence procedure which covers the delivery and enforcement of rights. Some solutions are designed as an open structure, able to integrate with existing front end services such as those provided by Microsoft, Real, Adobe and others.

CONCLUSION

A broadcast-based Internet is architecturally quite simple, and it offers a number of advantages in terms of coverage, and user value in terms of content quality and cost. Its primary use of multicast in a very efficient mode means that new opportunities are presented for marketing all forms of multimedia content. To enable this, a family of open security standards is required.

The Internet Engineering Task Force is probably the best-qualified body to develop these standards, and is making good progress in its MSEC Forum on aspects

of protecting streamed (or real-time) data. The commercial drivers for extending the scope fully to include mass-market issues are not yet in place. These include very large multicasts, the problems of large number of joins and leaves, and efficient source authentication and user non-repudiation.

Similarly, protection of offline content is not yet the focus of the standards bodies. The content industry has to some extent been caught out by the rapid growth of unauthorised file sharing among Internet users, and has a number of proprietary solutions for Digital Rights Management. For reasons of transparency and universality it is hoped that the main *de facto* standards will be developed into open standards.

There is no fundamental reason why streaming and offline services should be secured by separate means. The commonality of platforms, applications and operating systems should encourage the integration of interactive, streaming and offline service under one standard security umbrella.

REFERENCES

1. Nielsen NetRatings Internet usage data found at http://www.nielsen-netratings/hot_off_the_net.i.jsp
2. RFC 2627, 1999, Key Management for Multicast: Issues and Architectures, <http://www.ietf.org/rfc/rfc2627.txt>
3. Harney et al, 2002, GSAKMP Light, IETF MSEC Internet Draft, <http://www.securemulticast.org/msec-drafts.htm>
4. Perrig et al, 2001, TESLA: Multicast Source Authentication Platform, IETF Internet Draft, <http://www.ietf.org/internet-drafts/draft-ietf-msec-tesla-spec-00.txt>