

ISPadmin
July, 2001
Usenet News

Introduction

In this installment, I look at many die hard Internet user's favorite application and every service provider's headache, news. Usenet news is defined by RFC977 (NNTP proposed standard) and RFC1036 (Usenet message standard).

The problem of news is a difficult one for a service provider, due to the following attributes:

- Very high volume of both posts and news data itself
- Very high (exponential?) rate of increase in both number of posts and MB of data year over year
- High number of end users
- The distributed distribution model implemented by NNTP, making reliability and "correctness" (i.e. posts containing all of the data they are supposed to) problematic
- Infrastructure costs, including bandwidth and disk space
- Personnel costs to monitor spam, illegal articles originating on their network, as well as manage the news infrastructure itself
- Legal problems associated with spam and illegal pornography and warez

Small Provider Infrastructure

A small provider (and most enterprises who offer news to their employees for that matter) will likely utilize a single machine for news. This single machine performs all three functions typically required in a news infrastructure: inbound news relay, outbound news relay, and serving news to clients.

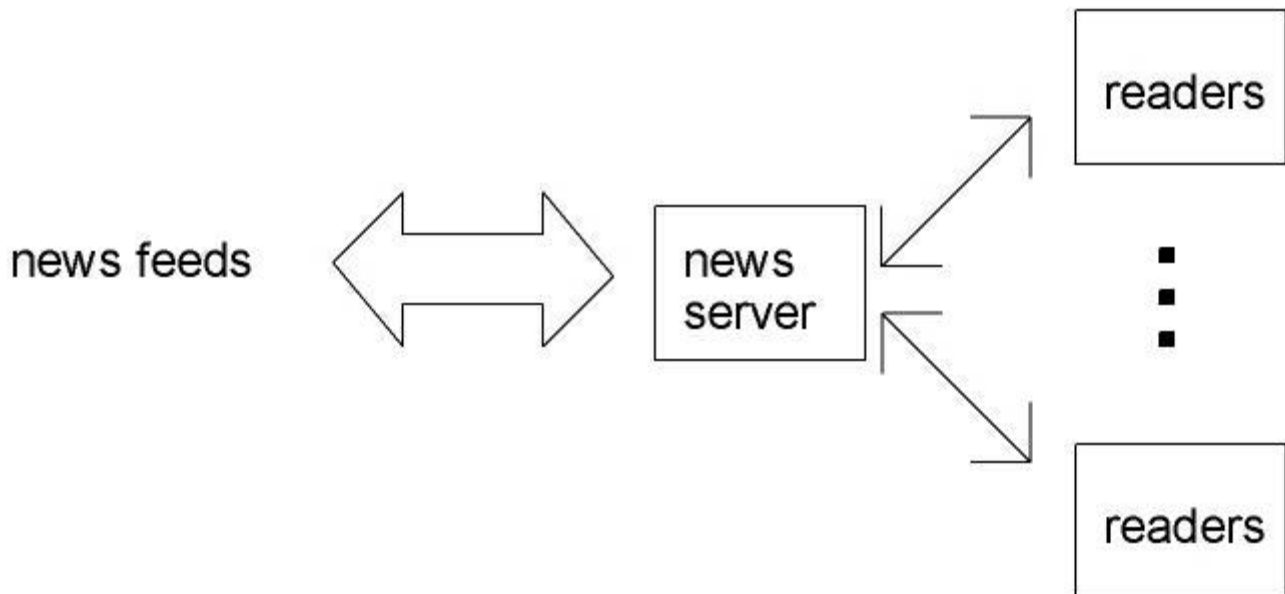


Figure 1

Figure 1 contains a diagram of how a small provider might set up their news infrastructure. The center box labeled "news server" handles all three basic news functions: inbound news relay, outbound news

relay and client news readers. The inbound articles come into the machine from the various sources of news (peers, commercial providers, upstream ISP(s), etc.). The outbound articles leave the server though the outbound news connection(s) (usually the sources as the inbound news streams). All of the customers who want news point their news clients to this same machine.

It is very likely that a small provider is not going to want to deal with the hassles of news and will outsource news to a provider like Critical Path (formerly Supernews/RemarQ). (The References section contains a pointer to a listing of news providers at the Open Directory project home page.) While some larger providers do utilize commercial news services providers, it is usually more cost effective for a big provider to setup their own news infrastructure.

Large Provider Infrastructure

A large provider is likely to deploy separate machines (or groups of machines) for each function: inbound news relay, outbound news relay and client news serving. Of course, functions can be combined; for example, inbound and outbound news relay can be the same machine if the inbound and outbound news volume isn't too high.

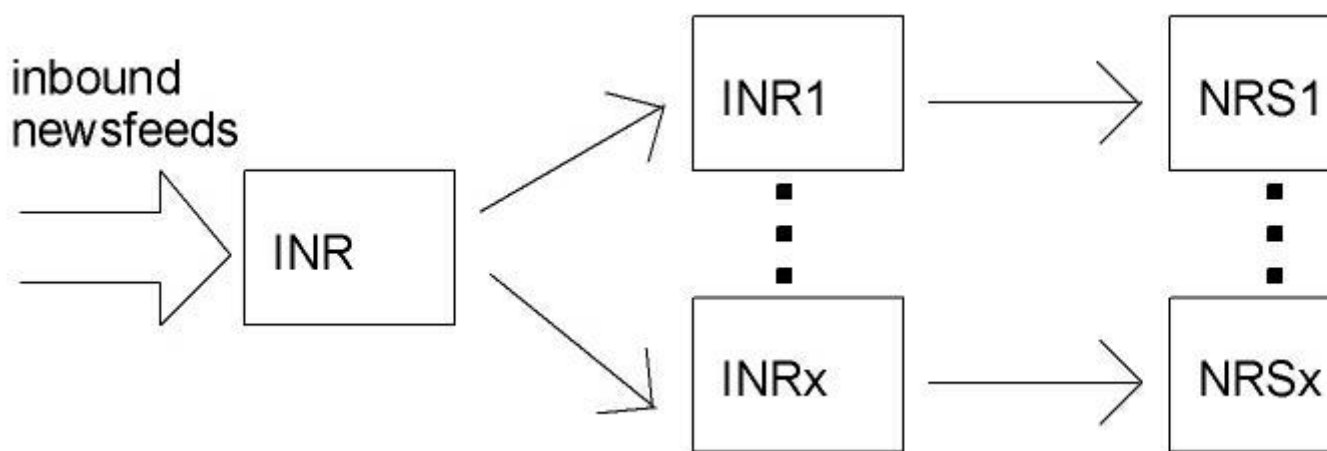


Figure 2

Figure 2 outlines how a larger provider might setup their inbound news infrastructure. The box marked "INR" illustrates an inbound news relay machine. This machine takes all off-site incoming feeds and consolidates them into a single feed. Only the very largest providers would need more than one INR machine for performance reasons. (Of course, they may have multiple INR machines for redundancy purposes.) A single machine to consolidate incoming feeds keeps transit costs to a minimum.

Multiple inbound news relay machines accept feeds from the inbound news relay machine(s). Each inbound news relay machine sends news articles to multiple news reading servers (indicated by "NRS1" and "NRSx") which news reading clients (not shown) attach to. Note that depending upon the news server hardware and software running on the news reading servers, each server can feed hundreds of news clients concurrently. The first class of machines to require scaling is usually the news reading servers, followed by the inbound news relay and finally the outbound news relay machine, which would be a distant third.




Figure 3

Figure 3 shows how a big provider could set up their outbound news (posts originating on their network) infrastructure. The boxes marked "ONR1" and "ONRx" indicate outbound news relay machines, which take articles from the news reading servers (labeled "NRS1" and "NRSx") and send them to the machine labeled "ONR". The outbound news relay might be located on the same machine that provides the inbound news relay function, depending upon the number of articles originating on the providers network. This outbound news relay machine is tasked with sending articles originating on the providers network to the Internet at large through the outbound news feeds previously configured. The outbound news relay machines (machines labeled "ONRx") are typically the last part of the news infrastructure that require scaling (after news reading servers and inbound news relays), as news clients don't usually originate many news articles.

Cyclical News File System (a.k.a. the Trash Can)

The vast majority of news implementations utilize either Internet Software Consortium's INN (originally written by Rich Salz) which is open source or Openwave's Typhoon/Cyclone series of commercial software. Before discussing the applications in particular, a short history and discussion of the circular news file system would be helpful.

Prior to the implementation of the Cyclical News File System (or CNFS) within INN, many providers (including Time Warner Cable of Maine and Ziplink) who had implemented INN 1.x switched to Typhoon and Cyclone because INN simply could not handle the load, nor expire articles automatically. Typhoon/Cyclone (and, more recently, INN) both feature an implementation of a cyclical news file system that eliminates many of the headaches when managing a news infrastructure, namely article expiration.

Article expiration is the process by which articles are "cancelled" and deleted from the list of available articles for download. As articles "age", they are expired. Historically, article expiration handled by setting parameters within an INN configuration file. Ever day at a certain time, a process ran that deleted articles that met the criteria set in the configuration file and all of the news indexes were re-indexed. This process could take hours for a large news system, frequently causing service interruptions.

With the rapid growth of the size of newsfeeds, the partition(s) articles were stored on frequently filled up if an administrator was not diligent in keeping the expiration configuration file up to date with the added news groups. Also, performing the expiration would often cause service interruptions due to the

load put on the news server while running the article expiration.

A cyclical news file system has no concept of article expiration. Therefore, there is no need to perform the CPU consuming expiration process and its associated overhead. With the advent of the 2.0 release of INN, INN includes CNFS support.

Openwave's Twister and Cyclone

Twister and Cyclone are commercial service provider grade news server implementations. Many service providers utilize these products for serving news to their customers. Some of the features of Openwave's Twister and Cyclone products include:

- Virtual server support

- Customizable anti-spam filtering

- Synchronized article numbering across entire news server infrastructure

- Real time statistics and logs capable of generating bills from

- Post filtering

- Automated moderator support

- Feeds automatically adjusted without administrator intervention for optimal throughput and efficiency

Openwave has a free version of their discussion software named Breeze. Of course, there are limits as to how many feeds and how many readers can connect to it, but it might be worth some investigation if you are in the market for Usenet news server software.

Internet Software Consortium's INN

INN is freeware and doesn't have all of the bells and whistles that a commercial application like Openwave's does. However, it is a fully functional news server and perfectly capable of serving news. The advent of the CNFS in INN makes it much more robust and usable in a service provider environment. Features of INN 2.3.1 include:

- Python, Perl and TCL authentication and filtering plugin support

- News reading over SSL

- Email gateway to news

- Exponential backoff for posting, enabling some level of anti-spam support

For many providers, large and small, INN is a fine solution to the problem of Usenet news. If the added features of a commercial news product (like article synchronization, virtual server support and real time statistics) are required, then a provider would likely utilize a commercial grade server.

News Client Software

News client software bears a brief mention. Both Microsoft and Netscape browsers contain client news reading capability. While they both can read news, I personally find them not nearly as functional as the Forte Agent news reading client. For those folks who have been around since before the GUI days, you can still read news from the Unix command line utilizing tin, trn, pine or a multitude of other character based news readers. If you are interested in finding out more about news clients, please check out the appropriate web site in the References.

Storage Considerations

When designing news infrastructure, many details must be considered. In the area of storage, single disk spindles (i.e., not RAID or other fault tolerant storage technology) are usually utilized for storage as losing articles is a tolerable event. Also, backups are almost never performed (except for those news providers who archive such things) because once again, losing articles is acceptable. Once the hardware failure is repaired, news will begin filling the disks again very rapidly!

News articles can be stored and shared via NFS mounts. Historically, many problems arose utilizing NFS for article storage, which accounts for the limited use of NFS in news implementations. Issues with utilizing NFS with news are:

- file locking
- performance

It is not recommended that NFS be utilized for news implementation, as there are much better ways (i.e. Storage Area Networks or SAN) to achieve similar functionality and performance.

Other Considerations

As you are probably aware Usenet news hosts thousands of newsgroups in a multitude of languages. For providers with networks located solely in the US, it is sufficient to carry the 50,000 or so English-only groups. International service providers would likely carry a complete feed with all non-English groups as well.

The end subscriber controls the groups to which they subscribe. When their news client initially connects to the news server, the news server will query the client as to what group headers to download (usually all are downloaded). Once the group headers are downloaded, the end user can subscribe (download article headers within each group) to whatever set of groups interests them, and then download individual article bodies that they wish to view.

Most providers carry local news groups (a group dedicated to restaurant reviews in the provider's city, for example). In fact, the Openwave series of news servers enables "virtual groups" to be located across servers and only visible to certain classes of clients (for example, the customers of a particular ISP in a wholesale ISP's news infrastructure). This is a very useful feature for a wholesale service provider.

One might wonder how much effort and hardware it took to run a news infrastructure at a moderate sized ISP. At Ziplink, we had a moderate size English language only (50,000 groups) news infrastructure (200 news clients). These clients were served by two Sun Ultra 5 machines, each with 512MB of RAM and approximately 36 GBs of disk space. One machine ran Cyclone and was the inbound and outbound news relay, while the other machine ran Typhoon and was the news reader machine. The load on either machine was never higher than 1. The aggregate feeds were on the order of 2 megabits/second, from a handful of UUNET news feeds and several news peers.

Occasionally a news spam complaint will arrive in the abuse mailbox of a provider. Usually, it is very easy to track down the perpetrators of news spam, as the logs and message headers themselves contain exactly when and where the message originated. Forging message headers makes this process much more difficult, but the logs again make it easy to determine positively whether or not a message

originated on a particular provider's network. Generally, Usenet news spam is much less of a problem for an ISP than junk email. In 2.5 years at Ziplink, I handled one Usenet spam complaint but hundreds of unsolicited commercial email complaints.

Legal Aspects of News

I've asked John Nicholson (the lawyer who writes in ;login: about legal issues surrounding computers) to cover ISP legal areas, as I'm not an attorney. Usenet news would definitely be an important topic for any discussion around service provider legal liability.

Most ISP's consider themselves a "common carrier". Having "common carrier" legal status would exempt ISPs from liability of what is carried over their infrastructure. How true this belief really is, I am not sure. If considered a "common carrier", a service provider cannot be held liable for pornography or illegal software (warez) originating or residing on their infrastructure. (An analogy to this would be holding the US Post Office responsible for someone sending illegal drugs through the mail.)

Most if not all providers perform no *content based* censoring (moderating) of what content flows through their network. Of course, decisions *not* based on content but strictly on technical capacities and related areas (for example, limits placed on news articles based upon the amount of disk space or network bandwidth available) is an acceptable means of controlling one's destiny as an ISP while not jeopardizing the potential for "common carrier" legal status.

Conclusion

Providing Usenet news functionality can be a difficult task for any provider. For a smaller provider, one machine can handle inbound, outbound and news reading capability. For a larger provider, the functionality is split up based on the type of functionality.

A good commercial news server is Openwave's discussion products, while INN remains the open source stalwart. Spam is not too much of an issue when it comes to news, and those who do are relatively easy to catch. Most ISP's do not filter groups for fear of losing "carrier" status, which is unclear ISPs have at this point under any circumstances.

Next time, I'll take a look at how service providers deploy their name service infrastructure. In the meantime, please send your questions and comments regarding ISP's, system administration or related topics to me.

References

NNTP proposed standard (RFC977): <ftp://ftp.isi.edu/in-notes/rfc977.txt>
Usenet message standard (RFC1036): <ftp://ftp.isi.edu/in-notes/rfc1036.txt>
RFC Editor: <http://www.rfc-editor.org/>
Critical Path Supernews/RemarQ: <http://www.supernews.net/>
Open Directory Project list of commercial news providers:
http://dmoz.org/Computers/Usenet/Feed_Services/
ISC's INN: <http://www.isc.org/products/INN/>
Openwave's discussion server software: <http://discussion.openwave.com/>
Microsoft Internet Explorer: <http://www.microsoft.com/windows/ie/>

Netscape Navigator: <http://home.netscape.com/browsers/index.html>

Forte Inc., Agent/Free Agent: <http://www.forteinc.com/>

tin: <http://www.tin.org/>

pine: <http://www.washington.edu/pine/>

trn: <http://trn.sourceforge.net/>