



NetCache Appliances: NetNews Caching Overview

Ed Chow | Network Appliance | TR 3064

TECHNICAL REPORT

Network Appliance, a pioneer and industry leader in data storage technology, helps organizations understand and meet complex technical challenges with advanced storage solutions and global data management strategies.

Table of Contents

1. [Overview](#)
2. [The Popularity of NetNews](#)
3. [Challenge: NetNews is Expensive to Provide & Maintain?](#)
4. [The Solution: The NetCache NetNews Cache](#)
5. [NetApp - The Gold Standard for Large Scale News Services](#)
6. [NetCache - Improves News Services while Lowering Costs](#)
7. [Deploying NetCache NNTP Caches](#)
 - 7.1. [Front-ending an NNTP Server Farm](#)
 - 7.2. [Local Service at a Remote POP](#)
 - 7.3. [Outsourced NNTP Services for Smaller ISPs and Enterprises](#)
8. [Conclusion](#)

[TR3064]

1. Overview

NetNews (including Usenet, alt, regional newsgroups) is the de-facto standard for online discussions on the Internet and is therefore a required service for most Internet Service Providers (ISPs). Unfortunately, the distributed nature of NetNews and the sheer volume of NetNews traffic make it a very expensive service to provide requiring large amounts of bandwidth, server and storage resources. Network Appliance (NetApp) has combined its many years of NetNews storage experience with its industry leading NetCache® appliance caching solution to provide an NNTP (Network News Transfer Protocol) cache. This NNTP cache enables Internet Service Providers (ISP) to provide a high-quality NetNews service at a fraction of the cost of traditional NetNews servers.

2. The Popularity of NetNews

For most ISPs, NetNews is the second most popular hosted service based on traffic (next to e-mail). The popularity of NetNews is based on many factors including the broad range of categories available (currently over 30,000 newsgroups), the ability for anyone to easily post an article, and the wide availability of news readers (NetNews support is built into both Netscape® Communicator and Microsoft® Internet Explorer).

3. The Challenge: NetNews is Expensive to Provide and Maintain?

The popularity of NetNews has made it a required service for most ISPs. Unfortunately, it is also a very expensive service to offer. Unlike the Web, NetNews doesn't live on servers hosted by content providers; published articles get distributed to news servers (mainly hosted by ISPs and corporations for their own users) around the world via news feeds, and these individual servers provide news service to the client readers. The tremendous amounts of published data distributed by these feeds creates a challenge for any ISP hosting a news service: a high-quality news feed typically requires 2 full T1 lines (3 Mbps) and 200GB of storage for every week's worth of news that the ISP wants to offer. These requirements continue to grow with the popularity of NetNews. To cope with this problem, ISPs have deployed a number of different solutions, each with its own problems:

- Many large ISPs have chosen to centralize news services at data centers rather than distributing the service to Points-of-Presence (POP). The advantage of this architecture is

that only one news feed is required, thus reducing bandwidth, storage and maintenance requirements. The disadvantages are that a large cluster of news servers are required to service all the ISP's clients, additional WAN bandwidth is consumed for redundant data that is requested to the same POP and the user experiences poor response times as articles travel across an ISP's WAN. This delay becomes increasingly noticeable as client connection speeds increase.

- Many small ISPs have chosen to outsource their news services to other ISPs or news providers, thus eliminating the need to manage their own services. The disadvantages of this model are that the ISP's customers get poor response times as the news must be retrieved across the Internet and the ISPs need additional bandwidth for any redundant data requested.

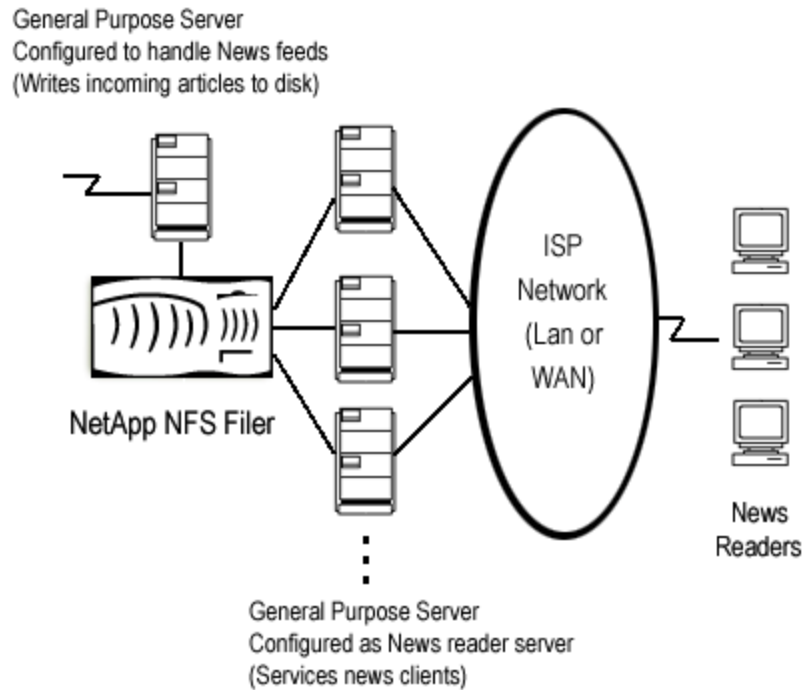
4. The Solution: The NetCache NetNews Cache

NetApp has combined its years of experience assisting many of the world's largest ISPs to serve news with its industry-leading caching solution to create the NetCache NetNews caching solution. The NetCache NNTP cache addresses the key issues facing ISPs today by reducing news server hardware requirements, minimizing bandwidth costs, improving client response time, and reducing the administrative burden of managing a news server. The NNTP cache service can run simultaneously on the same appliance as a NetCache Web (HTTP, FTP) cache and offers all the advantages of an appliance: price/performance, reliability, security and ease of management.

5. Network Appliance - The Gold Standard for Large Scale News Services

NetApp has for years been the leading data access solution for large-scale ISP news deployments. NetApp's Data ONTAP[®] software and patented WAFL[®] file system (the microkernel OS and file system found in NetApp filers and NetCache appliances) offer many benefits in NetNews environments including:

- High performance read access to support thousands of simultaneous news readers.
- Hashed directory structures for efficient handling of millions of articles.
- Support for NVRAM to efficiently cache incoming articles.
- Built-in RAID for improved data availability.
- Appliance simplicity to reduce administrative costs.
- Highly available appliance architecture that has a measured availability of over 99.99%, ten times higher than UNIX[®] and Windows NT[®] systems.



Existing Large Scale Network News Architecture

6. NetCache - Improves News Services while Lowering Costs

NetApp has combined this extensive storage experience in providing large-scale news service with its industry leading NetCache Web caching solution to produce a high performance, easy to manage NNTP caching solution. By caching NNTP traffic, the NetCache appliance helps both large and small ISPs provide high performance news services at a fraction of the cost of deploying traditional news servers.

Key benefits of the NetCache NNTP cache include:

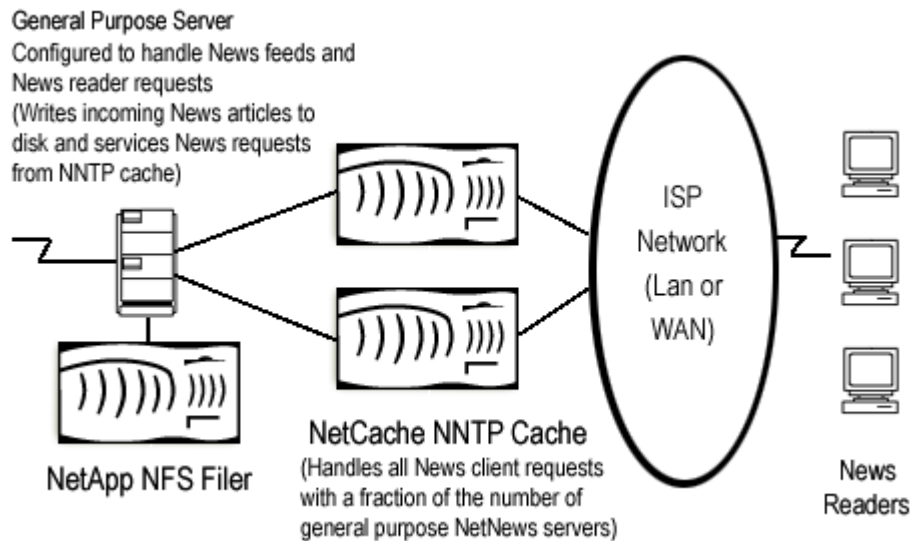
- Improved NNTP service by lowering end-user response time.
- Reduced bandwidth requirements and storage costs by only requesting and storing popular data (articles and group meta-data including overview data).
- Lower hardware costs. The NetCache appliance's high performance architecture enables ISPs to replace multiple NNTP servers with a single NetCache appliance.
- Lower administrative costs with NetCache server appliance architecture.

7. Deploying NetCache NNTP Caches

The flexibility of the NetCache NNTP cache solution enables it to be deployed in several different locations.

7.1. Front-ending an NNTP Server Farm

The NetCache appliance's high performance and reliable architecture make it ideal for front-ending large-scale NNTP server farms to significantly reduce the number of servers required and the administrative burden of managing those servers.



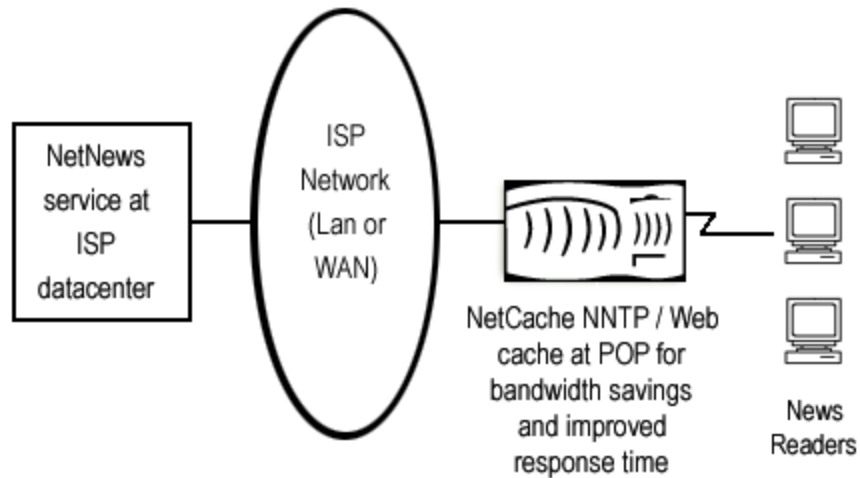
Large Scale NetNews Architecture Simplified with NetCache NNTP Caches

Key Benefits:

- Reduced hardware and software costs. The NetCache appliance's high performance NetNews caching architecture enables the replacement of multiple NNTP servers with a single NetCache appliance.
- Reduced administrative costs. The NetCache appliance simplifies maintenance by providing a self-managing service. NetCache appliances only process the data that users want and it automatically adjusts how long data is kept.
- Faster response times for better end-user performance. The NetCache appliance's microkernel architecture provides excellent response times even under heavy load.
- Can be combined with existing filer-based NetNews architecture for additional scalability and high availability.

7.2. Local Service at a Remote POP

The NetCache solution's highly reliable and remotely manageable appliance architecture allows global and national ISPs to deploy NetCache appliances at larger POPs to improve end-user performance and reduce bandwidth costs.




NetCache NNTP/Web Cache POP Deployment

Key Benefits:

- Improved quality of service by locating news service closer to the end-user.
- Reduced bandwidth costs by not having to retrieve popular news articles multiple times or send a news feed to every POP. The NetCache NNTP cache has demonstrated byte-hit rates of up to 80% and the majority of NNTP requests do not require expensive WAN bandwidth.
- Lower cost of ownership via the NetCache solution's highly deployable, remotely serviceable appliance architecture.
- Improved Web service (better response times and lower bandwidth costs) via the NetCache appliance's high performance HTTP cache. NetCache appliances can simultaneously cache HTTP and NNTP traffic.

7.3. Outsourced NNTP Services for Smaller ISPs and Enterprises

The NetCache appliance's simplicity, secure architecture, and industry leading price/performance enable smaller ISPs and enterprises to provide their users with a high performance news service at a fraction of the cost of maintaining their own NNTP server.



NetCache NNTP/Web Cache-Outsourced NetNews Deployment

Key Benefits:

- Private labeled NNTP service.
- No need to purchase or maintain news servers.
- Higher performance with lower bandwidth requirements than traditional, outsourced services as popular news articles are serviced directly from the cache.
- Reduced administrative costs via self-tuning appliance architecture.
- Improved Web service (better response times and lower bandwidth costs) via the NetCache appliance's high performance HTTP cache. NetCache appliances can simultaneously cache HTTP and NNTP traffic.

8. Conclusion

As the volume of NetNews continues to grow, ISPs are searching for a cost-effective way to provide this popular service. The NetCache NNTP cache combines many years of NetNews storage and caching experience to provide ISPs with a high performance and easy to maintain NetNews service at a fraction of the cost of traditional solutions.



Network Appliance, Inc.
 495 East Java Drive
 Sunnyvale, CA 94089
www.netapp.com

© 2005 Network Appliance, Inc. All rights reserved. Specifications subject to change without notice. NetApp, NetCache, and the Network Appliance logo are registered trademarks and Network Appliance, DataFabric, and The evolution of storage are trademarks of Network Appliance, Inc., in the U.S. and other countries. Oracle is a registered trademark of Oracle Corporation. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.