Performance Metrics for AGCS Applied to CastGate Technology and Native Multicast

Tudor Mihai Blaga, *Student Member, IEEE*, Gabriel Lazar, Bogdan Moraru and Virgil Dobrota, *Member, IEEE*

Abstract **— One of the many solutions to the present lack of native multicast is offered by CastGate. Like other Alternative Group Communication Services, CastGate makes use of UDP tunneling to relay multicast data. We determine the performance of CastGate and its impact on the network, using several metrics All the different CastGate solutions are taken into account, and the results are then compared to native multicast. Some of the metrics (stress, resource usage, stretch) can be evaluated, whilst other metrics (control overhead, join latency) can only be determined through measurements in a testbed.**

Keywords **— AGCS, CastGate, multicast.**

I. Introduction

THE issues concerning native multicast deployment or rather the lack of multicast deployment are covered in several papers [1], [2] and [3]. On one hand we have the complexity of the protocols involved and on the other hand we have the lack of customer demand.

Nowadays numerous multimedia applications make intensive use of network resources. Using unicast for these means that the bandwidth requirements increase linearly with the number of receivers, and so is the load on the servers. If multicast is used, the media content is sent only once to all the users, and thus bandwidth and server resources are spared.

Native multicast requires protocols like IGMP [4] and MLD [5] for the host management part and special routing protocols like PIM [6] for the creating of distribution trees.

Several proposals for AGCS [1] have been developed. They aim to bypass the lack of native multicast but also to go beyond the limitations of traditional multicast routing (no support for AAA).

The remainder of this paper is organized as follows. We present the characteristics of one AGCS proposal CastGate, we define the performance metrics used to evaluate CastGate, and we compare the obtained values to native multicast. Finally, we conclude this paper.

Tudor Mihai Blaga is with the Technical University of Cluj-Napoca, 400020 Constantin Daicoviciu 15, Cluj-Napoca, Romania (phone: +40-264-401816; fax: +40-264-597083; e-mail: tudor.blaga@ com.utcluj.ro).

Gabriel Lazar is with the Technical University of Cluj-Napoca, 400020 Constantin Daicoviciu 15, Cluj-Napoca, Romania (e-mail: gabriel.lazar@com.utcluj.ro).

Bogdan Moraru is with the Technical University of Cluj-Napoca, 400020 Constantin Daicoviciu 15, Cluj-Napoca, Romania (e-mail: bogdan.moraru@com.utcluj.ro).

Virgil Dobrota is with the Technical University of Cluj-Napoca, 400020 Constantin Daicoviciu 15, Cluj-Napoca, Romania (e-mail: virgil.dobrota@com.utcluj.ro).

II. Alternative Group Communication Services

These proposals can be classified according to the way multicast data is delivered. Some use tunneling and are based on UMTP [7], others use overlay multicast like Narada [3] and others use group specific routing services like Xcast [8].

The CastGate technology is the result of work by the Digital Telecommunications (TELE) research group from Vrije Universiteit Brussel. It provides seamless access to multicast content through the use of auto-tunneling [9]. It is intended as a transition technology that will lead to an increase in the number of multicast users. It uses a modified version of the UMPT called Enhanced UMTP [10].

A. CastGate

The basic CastGate architecture (Figure 1) consists of three parts: CastGate Tunnel Client (TC), CastGate Tunnel Server (TS) and CastGate Tunnel Database Server (TDS). The database contains information about all the available TSs. Multiple TDSs form what is called a Hierarchical Tunnel Database (HTD). The TS is to be found in the multicast part of the Internet, where it terminates one end of the tunnel. The TC is located at the client side, where it terminates the other end of the tunnel. It will ask the HTD for a list of TSs. The TC informs the chosen TS of the multicast group it wants to receive traffic for, and the TS tunnels the data to the client.

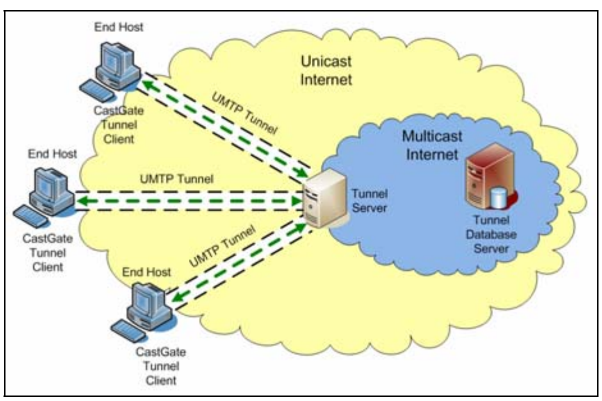

Figure 1. CastGate Client

The TC can be integrated in a multicast application or it can be a Java applet which runs in a web browser. In either situation, the operation is transparent to the end user. From the client's point of view it is as good as native multicast.

B. CastGate Router

CastGate Router is the result of further development of CastGate technology [11]. It integrates the functionality of an IGMP querier with the Tunnel Client. Thus it provides multicast access to all the hosts on the same LAN segment. The IGMP querier from the CastGate Router keeps track of the group membership for that LAN segment. Based on this information the Tunnel Client will join or leave the multicast group through the tunnel. The advantage of using a CastGate Router is that multicast traffic is tunneled only once for all the receivers on that LAN segment (Figure 2).

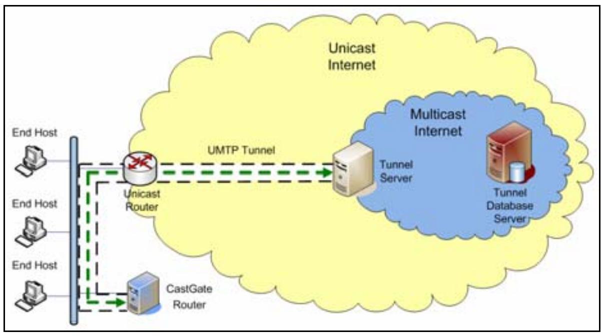

Figure 2. CastGate Router

C. CastGate with PIM-SM

The next step in the development of the technology was adding support for PIM-SM [12]. The functionality of the CastGate Router was extended in order to provide multicast access to an entire local domain. Here by domain we understand a group of networks under local administration, where any multicast protocol can be used, but without global multicast access. PIM-SM routing protocol was used because it creates multicast delivery trees with a single common root RP (Rendez-vous Point). Information about multicast activity in the domain is gathered by the RP. Placing a modified CastGate Router (Figure 3) on the same link as the RP would give us access to information regarding multicast receivers and sources in the domain.

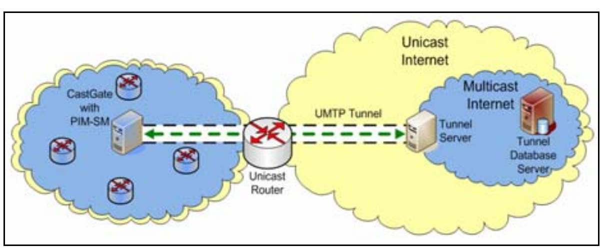

Figure 3. CastGate with PIM-SM

III. Performance Metrics for AGCS

Several performance metrics have been defined in [1] and [3] to characterize AGCS performance and impact on the network. The goal of our evaluation is to establish the downside of using the CastGate technology compared to native multicast. Some of these metrics can be determined using a simulated network architecture and some can only be determined in an operational network.

The metrics considered are:

- *Stress*: the number of identical copies of a packet carried by a physical link as the *stress of a physical link*. For example, if on a link the packet arrives tunneled from the source and then it is distributed through multicast, the stress on that link has a value of 2. In general, we would like to keep the stress on all links as low as possible.
- *Resource usage*: is defined as:

$$\sum_{i=1}^{L} d_i * s_i \qquad (1)$$

where L is the number of links active in data transmission, d_i is the delay of link i, and s_i is the stress of link i. The resource usage is a metric of the network resources consumed in the process of data delivery to all receivers. There is the assumption that links with higher delay tend to be associated with higher cost.

- *Stretch*: also called *Relative Delay Penalty* represents the ratio of the delay between the source and the receivers along the AGCS route to the delay of the unicast path.
- *Control overhead*: quantifies the cost of maintaining the AGCS topology, in terms of control information exchanged (number of messages and bandwidth).
- *Join latency*: also known as *Time to First Packet*, defines the time required for a newly joined member to start receiving the data flow.

In a previous paper [12] we determined through experiments the values for join latency and control overhead for IPv6 PIM-DM and PIM-SM. In the case of native multicast join latency refers to the time passed from the first Multicast Listener Report sent by the end host to the first multicast packet received by the host. Control overhead is obtained by measuring the bandwidth occupied by the PIM messages, that are used to create and maintain the multicast distribution trees.

IV. Evaluating CastGate

Using a given network topology we evaluate the following metrics: stress, resource usage and stretch. Our analysis focuses on what happens in the local domain, so the results presented refer only to these links. In order to determine join latency and control overhead measurements will be carried out in an operational network.

The test network (Figure 4) has a total of 7 links and only one connection to the Internet through router R1. Inside the domain we have three routers R2, R3 and R4. Routers R3 and R4 each have two LAN segments with multicast receivers connected. The number of end hosts is 17, from c1 to c17. We consider each LAN segment to be only one link because it is only one broadcast domain. We assume that the delay on all the links is equal and has a relative value of 1.

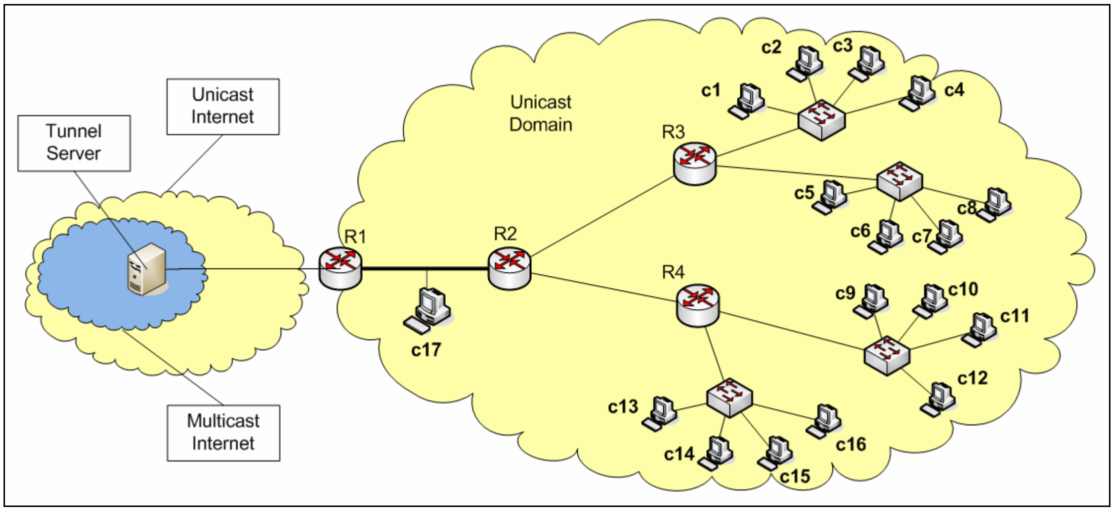

Figure 4. Network topology

Five scenarios were analyzed:

- CastGate Client
- CastGate Router
- CastGate with PIM-SM worst case
- CastGate with PIM-SM best case
- Native multicast

The first scenario taken into consideration (Figure 5) is the use of the CastGate Client. Each host in the domain runs the Java applet to receive multicast content. They are all receiving the same data. The stress values for each link can be observed in the figure. This scenario is similar to unicast from the determined metrics point of view, because the CastGate technology uses tunneling over unicast. A separate analysis for unicast was not carried out.

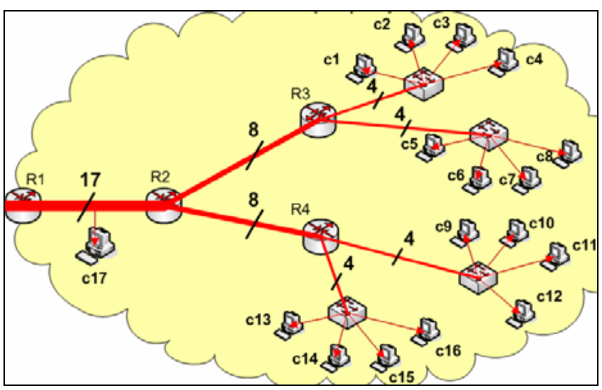

Figure 5. CastGate Client scenario

In the second case (Figure 6) one of the hosts in each LAN acts as a CastGate Router. We have four multicast enabled LAN segments, where the host use IGMP (Internet Group Management Protocol). The distribution of data through native multicast is represented with a different color than the tunneled data. Multicast traffic must be tunneled to the local domain only 5 times, for the following hosts c1, c5, c9, c16 and c17.

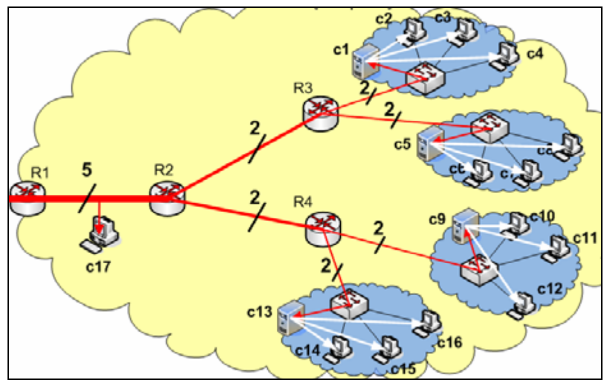

Figure 6. CastGate Router scenario

The use of CastGate with PIM-SM is considered in the third and fourth scenarios. The worst and the best cases are analyzed.

In the worst case scenario (Figure 7), the host c9, that is located in one of the LAN segments connected to router R4, serves as CastGate with PIM-SM device. All the multicast data is tunneled only once for all the receivers from the local domain.

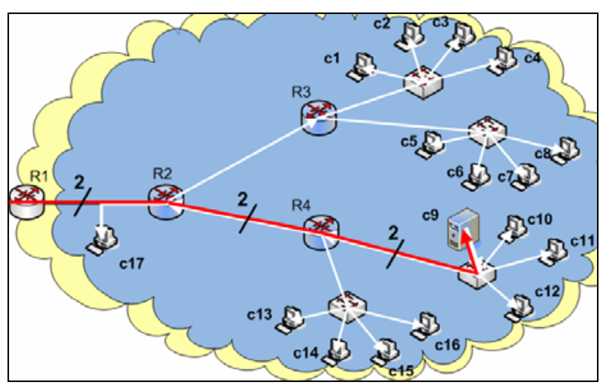

Figure 7. CastGate with PIM-SM worst case

Host c9 is the root of the multicast distribution tree; from this point on, native multicast is used. As it can be observed from Figure 7, the stress on some of the links has a value of 2, because tunneled data crosses them in one direction and native multicast crosses in the other direction.

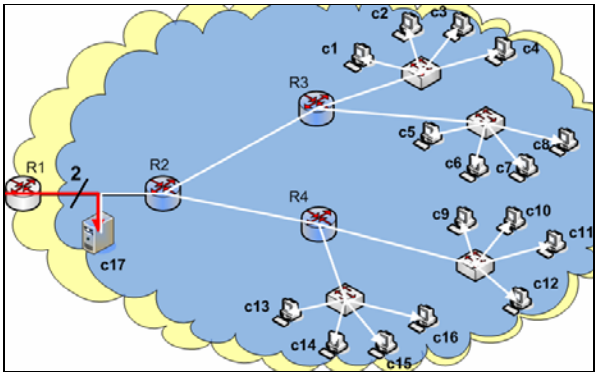

Figure 8. CastGate with PIM-SM best case

In the best case scenario (Figure 8), the host c17 that is placed on the link between R1 and R2, acts as a CastGate with PIM-SM device, thus providing multicast access for the entire domain. Its placement on the R1-R2 link offers better performance than the previous case.

The fifth scenario assumes the use of multicast in the entire Internet (Figure 9). There is no further need to tunnel the multicast traffic to the local domain. The stress has the value 1 on all the links, as a result of native multicast routing.

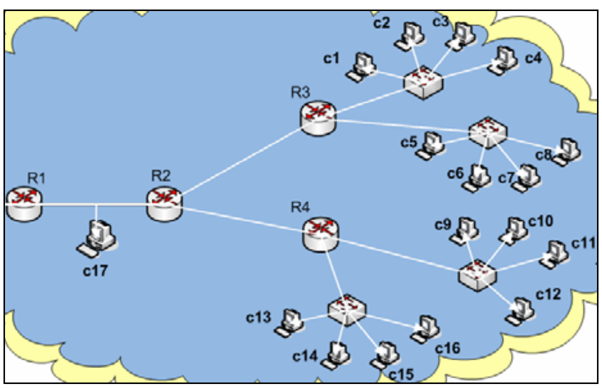

Figure 9. Native multicast scenario

Table 1 presents the values determined for the stress metric in each scenario. The results are presented in the second column on a per link basis. We observe that the access link (R1-R2) has the highest values (s1) in all of the CastGate based scenarios. This is due to the fact that all traffic for the local domain has to cross this link.

Table 1. Evaluated Stress values

	STRESS		
	s1	s2 / s3	s4 / s5 / s6 / s7
CastGate Client	17	8 / 8	4 / 4 / 4 / 4
CastGate Router	5	2 / 2	2 / 2 / 2 / 2
CastGate with PIM-SM worst case	2	1 / 2	1 / 1 / 2 / 1
CastGate with PIM-SM best case	2	1 / 1	1 / 1 / 1 / 1
Native Multicast	1	1 / 1	1 / 1 / 1 / 1

The use of CastGate with PIM-SM (best case) reduces the stress on the links by 8 to 4 times depending on the scenario. From the stress metric point of view, the two CastGate with PIM-SM scenarios are alike.

In our scenarios the resource usage, that is determined according to equation (1), because the delay has a relative value of 1, so $d_i = 1$ for $i = 1...7$, gives the following simplified formula:

$$R = \sum_{i=1}^{7} s_i \qquad (2)$$

Table 2. Evaluated Resource Usage and Stretch values

	RESOURCE USAGE	STRETCH for **c7**
CastGate Client	49	1
CastGate Router	17	1.33
CastGate with PIM-SM worst case	10	2.33
CastGate with PIM-SM best case	8	1.33
Native Multicast	7	1

Table 2 presents the values obtained for the resource usage and stretch metrics. In order to get a better measure regarding the whole local domain we must look at the resource usage metric. If we compare the CastGate with PIM-SM cases to the other CastGate based scenarios we note that it is 5 times more efficient than the CastGate Client scenario and 2 times more efficient than the CastGate Router scenario. Comparing the best case CastGate with PIM-SM with native multicast, shows that the resource usage is higher with 15%.

The stretch was determined considering R1 the source of the data. We made this assumption because of our interests in these metrics only from the point of view of the local domain. The delay from the actual source to router R1 has the same value, independent of the scenarios we considered. The reference delay for unicast in our local domain from router R1 to any of the hosts c1...c16 has the value 3.

The results presented in Table 2 are determined for host c7. With the exception of the CastGate with PIM-SM worst case scenario, we obtain the same values for any of the hosts. Note the increase for the stretch metric in the CastGate Router and the two CastGate with PIM-SM scenarios. This increase in delay compared to unicast delay is due to the fact that in these scenarios traffic is first tunneled to a device in the network and then this devices forwards it through native multicast to the receivers. This means that data has to pass the same link twice. The highest value 2.33 (7/3) is obtained in the CastGate with PIM-SM worst case for host c7. In this case the values differ depending on the considered host. For any of the host located in the LAN segment connected to R3, we have the same value as for c7. If the stretch is evaluated taking into consideration one of the hosts in the same LAN segment as the CastGate with PIM-SM device we get the same value as in the best case scenario. If the host is located in fourth LAN segment the value obtain is 1.66 (5/3).

In order to determine the join latency and control overhead measurements are under progress. These results will be compared to existing data for IPv6 PIM-DM and IPv6 PIM-SM.

I. Conclusion

The focus of this paper is the evaluation of the different CastGate solutions: CastGate Client, CastGate Router and CastGate with PIM-SM. The following metrics: stress, resource usage and stretch were determined. As results show, stress and resource usage decreases significantly when CastGate with PIM-SM is used. The values are very close to the native multicast scenario. The resource usage for CastGate with PIM-SM is only 15% higher that for native multicast, while it is 2 to 5 times less than the other CastGate scenarios. We must also notice the increase in stretch for the CastGate Router and CastGate with PIM-SM scenarios. The results prove also that the location of the CastGate with PIM-SM device is very important. The two cases considered (worst – best) show that the stretch metric is the most affected.

However these results must be confirmed by practical experiments by determining metrics like control overhead and join latency. At this point we can state that CastGate is a possible solution until native multicast is fully available, and we recommend the use of CastGate with PIM-SM where possible, due to better efficiency.

References

[1] Ayman El-Sayed, V. Roca, and L. Mathy, "A Survey of Proposals for an Alternative Group Communication Service", *IEEE Network*, January-February 2003, pp. 46-51.

[2] C. Diot, et al., "Deployment Issues for the IP Multicast Service and Architecture", *IEEE Network*, 2000, pp. 78–88.

[3] Y. Chu, S. Rao, and H. Zhang, "A Case for End System Multicast", *Proceedings of the ACM SIGMETRICS*, 2000.

[4] B. Cain and S. Deering, "Internet Group Management Protocol, Version 3", *RFC 3376*, October 2002.

[5] R. Vida and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", *RFC 3810*, June 2004.

[6] B. Fenner, M. Handley, H. Holbrook, and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", *draft-ietf-pim-sm-v2-new-11.txt*, 25 October 2004.

[7] Ross Finlayson, "The UDP Multicast Tunneling Protocol ", *draft-finlayson-umtp-09.txt*, November 2003.

[8] R. Boivie, et al., "Explicit Multicast (Xcast) Basic Specification", *draft-ooms-xcast-spec-09.txt*, December 2005.

[9] Pieter Liefooghe, "An Architecture for Seamless Access to Multicast Content", PhD Thesis, *Vrije Universiteit Brussel*, 2002.

[10] Pieter Liefooghe, "CastGate: An Auto-Tunneling Architecture for IP Multicast", *draft-liefooghe-castgate-02.txt*, October 2004.

[11] Pieter Liefooghe, M. Goossens, A. Swinnen, and B. Haagdorens, "The VUB Internet Multicast "CastGate" Project", Technical Report 10/2004 v1.8, *Vrije Universiteit Brussel*.

[12] Tudor Mihai Blaga, et al., "Steps towards Native IPv6 Multicast: CastGate Router with PIM-SM Support", *Proceedings of the 14th IEEE Workshop on Local and Metropolitan Area Networks - LANMAN 2005*, Chania, Greece, 2005.