CLIO/UNISONO: Practical Distributed and Overlay-Wide Network Measurement

Ralph Holz, Dirk Haage
Wilhelm-Schickard-Institut für Informatik
Universität Tübingen, Germany
{holz,haage}@informatik.uni-tuebingen.de

Abstract

Building on previous work, we present an early version of our CLIO/UNISONO framework for distributed network measurements. CLIO/UNISONO is a generic measurement framework specifically aimed at overlays that need measurements for optimization purposes. In this talk, we briefly introduce the most important concepts and then focus on some more advanced mechanisms like measurements across connectivity domains and remote orders.

1 Introduction

It is generally accepted that cross-layer information can be very useful for the optimization and re-structuring of overlays. At the same time, overlays seem to be one way of bringing more functionality to a future Internet. The SpoVNet project [4], for example, uses several overlays at once to provide an underlay abstraction and advanced services (multicast, event service). In this talk, we present CLIO/UNISONO, the cross-layer measurement service used by SpoVNet services to acquire data for optimization purposes.

2 Architecture and Basic Functionality

CLIO/UNISONO is an implementation of the split architecture proposed in [1]. UNISONO encapsulates measurement methods and logic. CLIO defines an interface and component that serves as a connector for a *querier* – usually an overlay node (or a service or an application). Nodes cooperating in a distributed overlay optimization algorithm can thus acquire local and remote information. It is also possible to order measurements between two other nodes, see section 3. CLIO is specified as an XML-RPC interface, so developers can easily write a CLIO connector for their own overlays. UNISONO comes delivered with a CLIO connector [3] for the Ariba overlay that is part of SpoVNet [2]. UNISONO currently supports about 50 *data items* that can be measured and determined. We distinguish between properties of a communication endpoint (node) and link properties. Among the endpoint properties are CPU properties, RAM and HDD, but also host uptime, host interfaces (like WLAN) and battery state. Among the link properties are RTT, path MTU and path bandwidth, but also features like multicast support between nodes. Measurements can be one-shot or periodic, but also with result delivery on exceeding a threshold. Measurement methods are encapsulated with a plug-in concept. Each plug-in is associated with a cost value that represents its resource usage. This is useful for our policy engine (see below).

3 Advanced Functionality

CLIO/UNISONO offers some measurement options that advance beyond what current measurement frameworks could offer for the optimization of overlays, especially if used combined.

Measuring Across Connectivity Domains Measurements across middleboxes like NATs are usually difficult. In the case of overlays like Ariba, CLIO/UNISONO can overcome this. Ariba knows the relay nodes at the boundaries of connectivity domains. CLIO can thus determine the relay path and split a measurement into the corresponding domain-internal measurements. These are aggregated upon completion and the result computed. Without Ariba-like overlays, this method is not an option; but we can still use overlay measurements, see below.

Overlay Measurements CLIO/UNISONO can be used for message-based overlay measurements. UNISONO can ask CLIO to send messages of a certain size to another CLIO instance (in the same overlay) and wait for a reply. Messages are time-stamped upon sending and receiving.

This can be used for, e.g., latency and bandwidth measurements on overlay level.

Remote Orders With the help of CLIO, UNISONO can conduct remote measurements. A node A can request a measurement from a node B to a node C. CLIO is used to send and receive remote orders (requests for measurements) over the overlay to which it is connected. Remote Orders can be very useful for some optimization purposes. They also allow to measure a node from different points in the network and aggregate the results (cooperative measurements).

Aggregation and Low Impact Distributed measuring and aggregation is a central concept in CLIO/UNISONO. Recent results are held in a cache. If two queriers need the same measurement within a certain time frame, the measurement is only executed once. The cache is currently an in-memory SQLite database, with confidence intervals for measurement types pre-defined.

4 Security Issues and Protection

CLIO/UNISONO was designed with security in mind. Communication between CLIO instances of a SpoVNet is authenticated, encrypted and integrity-protected. UNISONO instances only communicate via CLIO, and UNISONO only allows access from pre-registered CLIOs on the local host. A Policy Engine is a central component. Policies define which measurements a UNISONO instance allows. Access is defined with role-based ACLs. Access rights are granular: it is possible to assign a maximum resource usage or a rate limit. These mechanisms can be used together to limit the impact of malicious nodes in the network trying to conduct Denial-of-Service attacks on the local host or on a remote host. Local protection can be achieved with appropriate policy settings (for outgoing measurements and incoming remote orders). However, an attacker could attempt to request remote measurements from a high number of nodes to one single node (a type of DoS). This is more difficult to defend against. One can either use very prohibitive local policies (e.g. disallowing most remote orders). Or one can use *destination limits*. These are rate limits that define the maximum number of remote measurements that may be conducted to any given destination. However, this will only work if all or most nodes use similar policies. This is the case in SpoVNets, but more difficult to achieve in other settings. The effectivity depends largely on the number of malconfigured nodes in the network and the number of nodes known to the attacker (to whom he can send remote orders). We have done some first rough estimates and believe that our protection can be sufficient for small networks (in the order of 10,000 nodes), but becomes much less effective in larger settings. For such networks, we recommend to disable remote orders entirely.

5 Benefit Examples

We give two examples how CLIO/UNISONO is already used. Multicast groups in SpoVNet can be formed according to bandwidth demands. In WLANs, nodes can additionally be grouped according to their access points. The event service in SpoVNet uses latency values for placement optimization. Remote orders and measurements across connectivity domains allow measurements that would otherwise be difficult or impossible to conduct.

6 Ongoing Work

We are currently focusing on correlation mechanisms for underlay and overlay measurements (e.g. stretch for RTT). UNISONO will be extended to raise statistics, which may help in investigating dynamically chosen caching times.

7 Conclusion

We have reported on CLIO/UNISONO, our framework for cross-layer measurements in distributed settings, which features some advanced mechanisms like measurements across connectivity domains and remote orders. We have also discussed some security aspects and described some benefits.

Acknowledgements This work was in part funded by Landesstiftung Baden-Wuerttemberg in the scope of the Spontaneous Virtual Networks (SpoVNet) project.

References

[1] D. Haage, R. Holz, H. Niedermayer, and P. Laskov. CLIO – a cross-layer information service for overlay network optimization. In *Kommunikation in Verteilten Systemen (KiVS) 2009*, Kassel, Germany, March 2009.

[2] C. Hübsch, C. P. Mayer, S. Mies, R. Bless, O. P. Waldhorst, and M. Zitterbart. Reconnecting the internet with ariba: Self-organizing provisioning of end-to-end connectivity in heterogeneous networks. In *SIGCOMM 2009 Demos*.

[3] UNISONO: UNified Information Service for Overlay Network Organization. http://projects.net.in.tum.de/projects/unisono/, 2009.

[4] O. Waldhorst, C. Blankenhorn, D. Haage, R. Holz, G. Koch, B. Koldehofe, F. Lampi, C. Mayer, and S. Mies. Spontaneous virtual networks: On the road towards the internet's next generation. *it – Information Technology Special Issue on Next Generation Internet*, 50(6):367–375, December 2008.