# Certifying performance of cooperative services in a digital government framework

F. Arcieri[1] F. Fioravanti[2] R. Giaccio[1] E. Nardelli[2] M. Talamo[1]

1. "NESTOR - Laboratorio Sperimentale per la Sicurezza e la Certificazione dei Servizi Telematici Multimediali", Univ. of Roma "Tor Vergata", Roma, Italia.
2. Dipartimento di Informatica, Univ. of L'Aquila, L'Aquila, Italia, & Istituto di Analisi dei Sistemi ed Informatica, C.N.R., Roma, Italia. nardelli@di.univaq.it. CONTACT AUTHOR.

## Abstract

*One of the hottest challenges in the digital government arena is the capability of providing good quality services to citizens. The critical issue is that for a given service a citizen usually interacts with a single provider, even if service supply and management requires coordination and cooperation among many autonomous organizations. This means that a single user request spreads in the underlying distributed information system and activates a number of information flows among organizations involved, with various roles and responsibility, in service provision. A main issue is how what is going on in the distributed system can be objectively monitored so that the service provider can (*i*) understand and manage problems in the overall service supply process and (*ii*) certify quality of provided service. What makes this scenario especially complex is that, beyond technical aspects, any solution that wants to be successful has to comply with requirements of independence and autonomy of the various organizations involved. In this paper we discuss how we tackled and solved this issue in real-world systems defined for the Italian Public Administration and we argue that our solution can provide a reference architecture to deal with this kind of problems.*

**Keywords**: digital government support, inter-organizational e-service certification, actual performance monitoring, application interoperability and cooperation.

## 1 Introduction

The management of digital government services has opened a new type of network service monitoring problems. In fact, these services are usually provided to citizens through a single access point, but very often a service of this kind requires the coordination of and cooperation among many autonomous organizations. Hence the architectural and technological operating scenario is very complex [14], since the various public administrations (PAs) and agencies that are involved in a single service are autonomously and independently managed [9, 20]. An IT infrastructure supporting e-government services has therefore to allow efficient monitoring of service execution without being intrusive with respect to IT solutions existing in the involved organizations. This means to be able to check and to certify the status of progress of the distributed transaction(s) activated by requests to an e-government service while treating various legacy IT components of involved organizations as black boxes.

These specific organizational characteristics of digital government services require, from the network traffic monitoring viewpoint, a new kind of application level measurement techniques. In fact, traditional approaches in this area have focused on web servers [17] or content distribution networks [22] performance measurements. For our purposes, instead, we need to measure and to certify actual performance of service flows which spread in the network in consequence of an end-user's request. To obtain precise measurements, it is then needed to record the actual behaviour in the network of IP packets corresponding to service flows. To the best of our knowledge no solution for the problem of actual performance measurement of distributed e-services has been proposed in the literature.

Emphasis on actual performance measurements versus performance estimation, where sophisticate techniques have been proposed for accounting and billing [13, 15], stems from the fact that in the digital government service framework very often a legal value is attached to information exchanged, and in these cases it is not possible to use an estimation based approach. The same motivations pre-



Proceedings of the 2003 Symposium on Applications and the Internet (SAINT'03)
0-7695-1872-9/03 $17.00 © 2003 IEEE

vent the use of flow statistics like those being provided by Cisco NetFlow [11].

Also, from the (higher level) viewpoint of the application services, only recently in the Data Base [12, 25, 26] and in the Software Engineering [18, 24, 27] research communities this problem is receiving specific attention.

In this paper we address this issue, by discussing how to monitor, measure and certify actual performance of a cooperative service provided to end-users in a network (*e-service*, for short) by means of the interaction of autonomous and independent organizations in a digital government framework.

Our solution has been tested and refined while working in the wider context of the definition of the architecture of an IT-based system allowing: (*i*) to ensure that exchanged data is kept coherent in the different PA organizations, even under updates, (*ii*) to control and certify the exchange of information between independent PA organizations, each with its own hardware and software systems, and its different organizational procedures. We addressed these two points in [2, 6], and in [7, 8], respectively, proposing solutions deriving from on-the-field experience gained while realizing real-world systems [3, 4, 5] for the italian Public Administration.

This paper is structured as follows. In section 2 we provide a general description of the operating scenario, while the system architecture of reference is presented in section 3. Subsequently, section 4 discusses how network service performance monitoring and organizational constraints interact in our scenario. Following section 5, where we present and explain our architectural solution and compare it with existing approaches, we discuss in section 6 examples of network service performance measurement and certification in a real-world system. Finally, section 7 concludes the paper.

## 2 General Description

In our framework, a single e-service request is made up by a single query from the end-user and the corresponding answer from the provider.

A single end-user request reaching the site of the provider activates, in general, a series of further requests towards other sites which contribute parts of the overall service supplied by the provider. These other sites, in turn, can send other request towards further sites and so on, and a potentially complex information thread spreads over the network.

When supplying to end-users an e-service made up by parts furnished by autonomous and independent organizations the service provider facing end-users has therefore a main problem: how to measure and to certify, with respect to requests issued for the given e-service, performance of the distributed system made up by the collection of information (sub)systems of participating suppliers and the communication subsystem interconnecting them.

Moreover, participating organizations usually run their base services using legacy systems, developed since long time according to completely independent strategies, directions, and technologies.

The difficulty in overall performance measurement and certification lies then in the fact that what is required is the capability to identify actual performance for each single information thread activated by an end-user request.

From a physical point of view such a thread is made up by a large number of IP packets travelling according to routing strategies and current traffic distribution load (both outside the control of the provider). Any measurement regarding the traffic of IP packets is completely useless in the view of providing the specific information the provider needs. Even aggregating, at each participating node, traffic measurements regarding all IP packets related to a same request does not provide useful information to the provider.

The reason is that in such a way measurements are not correlated to the information threads activated by the end-user requests. What is important to know for the service provider is how good was the performance of the distributed system for each end-user request, that is for each information flow activated by a request, and then to aggregate this information over all end-user requests to get an overall certified measurement of the provided quality of service.

Clearly the problem here is also of organizational nature, since from a purely technical point of view one could implement a single communication management system encompassing all the involved systems and give to such a system the task of carrying out such measurements. This is certainly technically feasible. The organizational problem is that such a system would give to anybody controlling it an explicit control over the involved organizations, and this clearly is not feasible in the scenario of autonomous independent organizations we are considering.

Therefore a mechanism is needed than can be reliably managed by a third party and is, at the same time, able to carry out this measure in a certified (i.e., reliable and objective) manner by reconstructing information threads of each request without being intrusive with respect to the information systems of the involved organizations.

Also, note that for an IT infrastructure supporting e-services to work efficiently, basic functions like performance measurement and certification have to be independent and abstracted from the actual e-service invoked or executed.



Proceedings of the 2003 Symposium on Applications and the Internet (SAINT'03)
0-7695-1872-9/03 $17.00 © 2003 IEEE

## 3 Reference Architecture

The reference architecture for our discussion about network service performance measurement in digital government infrastructures is clearly a distributed one. This is compliant with more recent trends and requirements in Public Administrations and e-government actions where decision capabilities are increasingly and increasingly being decentralized and put at the appropriate local level. A number of *external services* are attached to the Public Administration intranet but are not directly accessible to clients. For this purpose a client has to access the web-site of the *e-service provider* through which external services are then made available. See[8] for more details on the reference architecture.

In a typical example of a service request an end-user connects to the web server of an e-service provider through a web client, asks for a specific service and sends needed parameters.

Within the provider's site the end-user request is forwarded to an application server and here, in general, it is decomposed in a set of further requests to external services to get required pieces of information or to carry out required checks.

For example, a service provider supplying tax information with respect to a given apartment may have first to check with the Personal Data Registry Service of a given Municipality whether personal data of the end-user are correct, then to ask to the Tax Service of a (possibly different) Municipality which are the current taxation levels for the apartment, and finally to verify with the Ministry of Finance which are the end-user information rights with respect to the property of the apartment itself.

Under some circumstances the first and/or the second checks might be skipped. In this example, then, a generic end-user's service request may activate different kinds of threads. This is true in general for digital government e-services: a same kind of request to the end-user access point may activate different threads depending on values of input parameters. The main consequence of this fact is that, from the viewpoint of performance measurement, reliable values can be obtained only through actual performance measurement, since the use of estimation techniques appears to be, given the organizational constraints, highly difficult and unable to provide accurate measures. Still, on top of these difficulties, one has to take into account that, as recalled in the introduction, the legal value attached to information exchanged in digital government e-services mandates the record of actual performance.

## 4 Performance monitoring and organizational problems

The problem of e-services monitoring could either be handled by an independent trusted layer or by designing a system that directly takes into account the problem. Apparently this latter solution, designing a communication system that both handles transactions and keeps track of all transaction details, is a better one. Note, however, that this is not always possible or practical. For instance, data communication can be provided by a legacy system that is difficult to substitute or modify. More difficult problems arise in the scenario we are considering of organizations cooperating to provide an e-service:

- organizations may have different data communication and service implementation solutions, and may be unable to agree on a common standard;
- organizations may have different strategies on security and privacy;
- organizations may handle similar data in different ways, but there is the need to correlate these data.

If any problem occurs, an obvious approach is to correlate all log files and other information sources in the involved organizations, and then to understand and manage the overall problem. This clearly requires a higher organizational level with respect to those involved in the service. When an e-service is implemented through the cooperation of autonomous and independent organizations it is impossible to define a unique point responsible for the monitoring and handling of all events arising from the correlation of the log files.

In fact, such an approach would imply the involved organizations should disclose the internal details in their log files: this is not possible due to leadership and privacy problems.

Then the challenge is the definition of an architecture allowing to carry out performance measurement according to the approach described in section 3 but taking into account organizational issue here discussed.

The main technical difficulty for performance monitoring and certification purposes in our context is that all e-services involved are, from the viewpoint of the monitoring process, like black boxes and cannot be internally changed. Our approach is then based on *active tracing*. Active tracing consists in scanning in real time all outgoing and incoming inter-organizational transactions and extracting all and only the information needed for monitoring and certifying service performance.



Proceedings of the 2003 Symposium on Applications and the Internet (SAINT'03)
0-7695-1872-9/03 $17.00 © 2003 IEEE

But input and output flows of IT services taken as black boxes are nothing more than sequences of IP packets.

Then the challenge is to be able to understand, by only checking packets flowing through various points over the intranet, what is going on with respect to a specific service. Note that even if we are dealing, in these digital government cases, with the PA intranet, in general all sort of packets can be found. So the goal is to be able:

- at specific points, to collect all packets corresponding to specific application and information flows related to a specific service,
- to reverse engineer streams of packets so to reconstruct application level messages passed through monitored points,
- to correlate what has been identified at those specific points so that it can be understood which is the status of the overall distributed transaction,
- to check whether application threads corresponding to a given service request are compliant with application rules defined by the e-service provider and, if not, to understand where and when problems occurred.

This approach clearly requires algorithms for packet classification with very high efficiency, for avoiding a too large slowdown in the interaction among the basic components. We have a very good algorithm solving this problem, which is able to perform the task in real-time by using a highly efficient algorithmic technique [23], originally developed for secondary memory searching and usable as well for packet classification purposes. See[8] for further details on the reconstruction of application flow from TCP/IP packets.

## 5 The architectural solution

Our architectural solution[1] is therefore based on *network probes*. They operate at each site involved in service provision and, once properly set-up, monitor all and only that traffic, flowing between them and the communication network, for which they have been explicitly configured. Note that, of course, configuration of network probes directly derives from agreements formally established between organizations involved in the exchange of services.

Network probes are devices featuring high security levels, such as absence of terminal devices to access internal resources (e.g., keyboard, mouse, screen, ...), uninterruptable power supply, software components for automatic faults check, diagnosis and alert. Their action is very efficient and neither disturb nor slow down operations of the communication layers. They are physically placed on network links directly leading to or coming from participating nodes.

[1] Patented, 1997.

Probes are able to detect all IP packets passing on the portion of the network they are controlling, to efficiently identify all those related to application services for which they have been configured, and to select only those referring to specific kinds of transactions. From these IP packets, probes in fact reconstruct the individual messages of the application threads relative to services they have to monitor, then they extract suitable data items able to certify and to document information passed through the link under their surveillance. Reconstruction of the specific messages exchanged at the application level between the monitored service(s) and its enduser(s) is done by reverse engineering the data flow of IP packets first to the TCP level and then to the application level, as explained at the end of the previous section.

Note that probes neither have to collect and process all data passing on the network link they are monitoring, nor have to reconstruct exchanged flows for all servers on the network. They are configured for specific services and process just IP packets referring to them. Hence this approach is fully scalable without performance degradation.

Synthetic information items filtered out by network probes, as result of their real-time analysis of IP packets according to their configuration parameters, are sent to a central *probes controller*. This processes all received information items and establishes correlation among those referring to the same service flow. Note that with this approach network traffic overhead is very low since only synthetic elements, at the application level of the communication, and only those relative to the service for which the probe has been configured, are sent back to the probes controller. The probes controller is then able to aggregate information items referring to the execution of a given service. Probes also compute, on the payload part of packets, a suitably defined hash function, whose outcome is sent, together with other information, to the probes controller. In such a way the controller is able to check also the correct exchange of payload data between provider and end-users, hence to provide a certification of the integrity of exchanged data or to establish when and where an error, if any, has occurred.

An important aspect of this solution based on network probes is that it allows the service monitoring and certification process to be allocated, if desired, to a third-party, independent from both the servers and the clients involved in the exchange of e-services. Note that this does not mean to send externally all IP packets of PA since the third party is anyhow within the PA's intranet. But with this choice the technical leverage point to enact, control, and enforce regulation policies for e-government services is obtained. 'If', 'when', and 'what' are of course issues of competence of higher-level policy makers, but it is important to stress that



Proceedings of the 2003 Symposium on Applications and the Internet (SAINT'03)
0-7695-1872-9/03 $17.00 © 2003 IEEE

this technical solution is transparent with respect to the inner structure and operation of the involved providers. Hence it does not incur the risk of being unusable in practice due to technical mismatches and difficulties.

It is worth mentioning here the main differences between our approach and one leading system for evaluation of performances of an e-service, the Keynote [10] system, as it represents quite well the overall trend in performance evaluation over the Internet. The Keynote system provides a commercial service which makes it possible to estimate the end-user perceived performance of a web site by using up to 50 software agents which issue HTTP requests to the monitored web site at regular time intervals. Measurement agents run at different geographical locations according to a distribution which reflects the real distribution of Internet users.

The system provides the service subscriber with a statistical summary [21] of web site availability and performance.

In general, in PA transactions we must be able not only to detect the compliance of performance parameters with the contractual ones, but also to exactly determine the errors and possibly activate recovery procedures to properly close the transaction. For instance, if an error occurs during a transaction related to some update in the central Cadastral database, one *must* be able to determine the cause of the error at least to a sub-service level, so that the involved organization can correct the error cause. Hence, one has to monitor *all* critical transactions and certify their actual behaviour: this is our approach and we store in a central repository the core information of each transaction under surveillance. Capturing all these data and sending them to the centralized repository does not cause a significant penalty on the network capacity, since all the information not needed to identify and correlate the transaction are stripped out before sending the data. The Keynote system, on the contrary, monitors only a subset of the transactions over a given period of time for statistical analysis purposes. It is worth to stress again that for the certification purposes that are mandatory in a digital government framework, a statistical approach to performance monitoring is not suited.

When dealing with complex distributed applications cooperating in a global transaction, we are not only interested in certifying the performance parameters on the perceived user application entry point, but also in monitoring the same parameters with respect to the internal communications of the involved sub-services. Hence, we have network probes on all service providers access points and all interconnections among involved organizations; we also have *software network probes* (i.e., implementation of the network probe functions as software programs) on all clients. This allows us to capture all the existing transaction information with high granularity, also for statistical purposes. On the other side, the Keynote system is able to monitor a set of single end-user requests, with no knowledge about the interaction with third-party systems generated by each request. Moreover, it does not capture real transaction data; instead, it mimics real transactions for testing purposes only.

Finally, since we basically monitor all transactions, we do not need any statistical inference tool and this is clearly an advantage.

## 6 A real-world example of e-service performance measurement and certification

Here we discuss the actual results of e-service performance measurement and certificationfor a real-world specific service. In particular, we present how the italian Ministry of Agricultural and Forestal Policies measures and certifies e-service performances in SIM ("Sistema Informativo della Montagna"), a real-world system we have realized according to the architecture previously described. SIM is a distributed system providing e-government services to people living in mountain areas [5].

SIM acts as a mediator between citizens living in those areas and IT-based services in the fields of cadaster, labour and pension, public registry of personal data. Its design started in 1998 and the overall financial effort has been, until now, of about 52 million euros (roughly 47 million US dollars). Nowadays it is currently being used as a fully operational system in about one thousand operating centers, all over Italy, serving more than 10 million inhabitants, more than 4000 of the about 8000 municipalities and covering more than 54% of the italian territory. Its extension to the whole country is currently under implementation. It is worth stressing that since the SIM and the others Italian PA systems currently based on the network probes architecture described in the previous section are able to handle several millions of complex transactions per year in thousands of physically distinct service centers distributed all over italian territory, the architecture has obviously been proved to be fully scalable.

SIM is structured around Service Centers, providing services to end users, coordinated by various Regional Centers and a National Center. Some of the services provided by SIM are managed internally, some others (e.g., cadastral services, public registry of personal data, ...) are made available through SIM by other agencies. A national Managing Agency is in charge of the overall control, coordination and certification activities regarding SIM.

In the reported figures you can see actual performance results obtained by means of the architectural solution described above. Note that the long response times for some services are due to the fact that each transaction in-



Proceedings of the 2003 Symposium on Applications and the Internet (SAINT'03)
0-7695-1872-9/03 $17.00 © 2003 IEEE

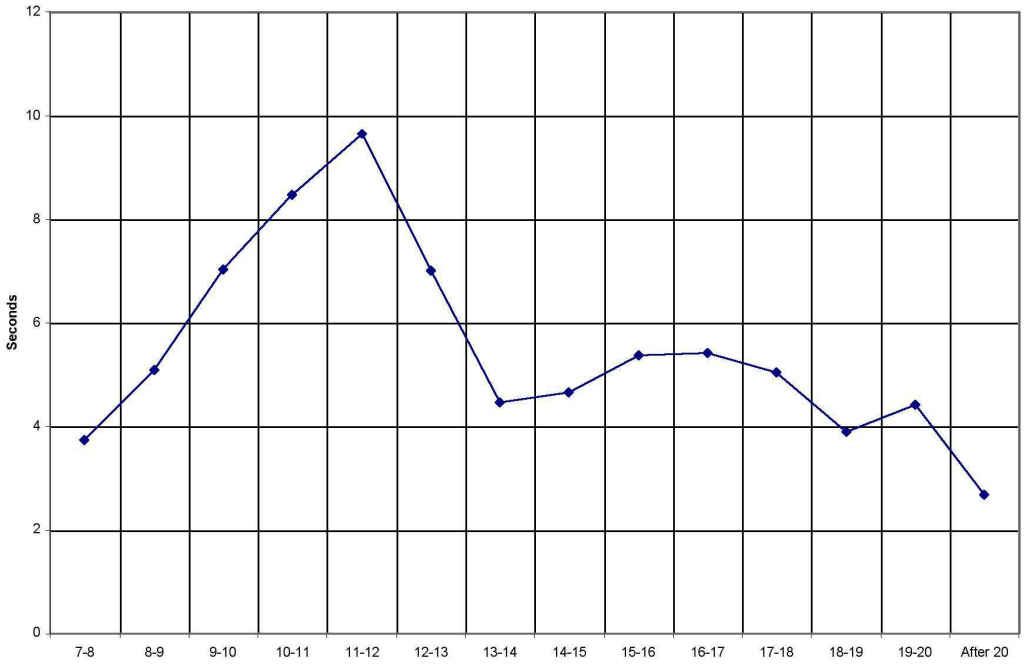


**Figure 1. Distribution of the average answer time of the cadastral services accessible through SIM.**

volves several different sub-transactions, possibly with different Public Administrations. For instance, each cadastral transaction involves an access to the cadastral web service, which then connects to the central cadastral service provider, that finally routes the request to the peripheric cadastral service provider in charge of the requested information.

More precisely in figure 1 the hourly distribution of the average answer time (in seconds) of cadastral services provided through SIM is shown. This is the average time elapsed between a request issued by the SIM National Service Center to the cadastral system and its answer.

In figure 2 the hourly distribution of the average time (in seconds) spent by end-users to obtain answers from SIM (upper graph), and its partition among SIM processing time (center graph) and the overall time needed for all network communication (bottom graph) is shown.

In figure 3 the average processing time (in seconds) needed to SIM to process queries for nine different subservices provided by SIM is shown. Histograms in white refers to the 11:00-12:00 time band, while those in black refers to the 12:00-13:00 one.

## 7 Conclusions

In this paper we have considered network service performance measurement and certification issues arising in the Information Technology infrastructure supporting digital government services.

Performance measurement and certificationof e-services is a critical activity for any integrated network service, but is of the utmost importance for e-government since usually these services are realized through the cooperation of autonomous and independent organizations.

We have presented a solution allowing to efficiently and effectively deal with these issues. It also provides the technical platform for constructing control policies, if these are deemed necessary. Our solution directly derives from successful experiences in the development of a number of real-world systems [2, 4, 5]. We have shown, by presenting examples of performance measurements and certification of a real world large scale systems (namely, SIM) that our solution is both effective and scalable.

Our approach requires a careful modification to be used in the case of public network, like Internet, and in the newest distributed computing infrastructures, like the GRID [16]. In fact, very often in this case services are provided on top of a secure communication layer (e.g., SSL).



Proceedings of the 2003 Symposium on Applications and the Internet (SAINT'03)
0-7695-1872-9/03 $17.00 © 2003 IEEE

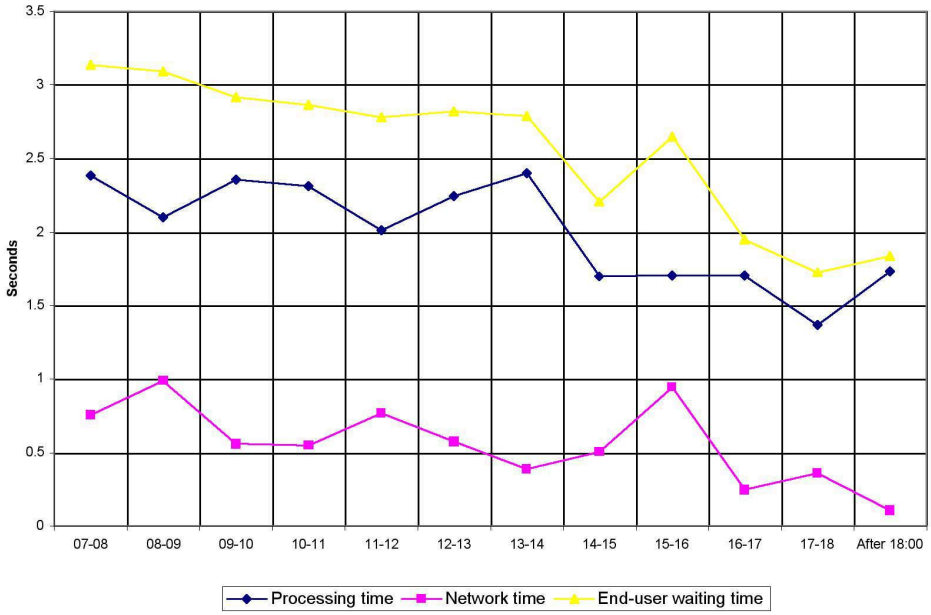


**Figure 2. Distribution of the average enduser waiting time (upper) for SIM services and its partition between processing (center) and network (lower) time.**

The architecture here described can still be a conceptual reference, but different technical solutions have to be implemented for the probes.

## References


[1] B. Aboba, J. Arkko, D. Harrington, Introduction to Accounting Management, RFC 2975, October 2000.

[2] F.Arcieri, E.Cappadozzi, P.Naggar, E.Nardelli, M.Talamo: Coherence Maintainance in Cooperative Information Systems: the Access Key Warehouse Approach, accepted for publication in the *Int. J. of Cooperative Information Systems*, 11(1-2):175–200, 2002.

[3] F.Arcieri, C.Cammino, E.Nardelli, M.Talamo, A.Venza: The Italian Cadastral Information System: a Real-Life Spatio-Temporal DBMS, *Workshop on Spatio-Temporal Database Management (STDBM'99)*, Edinburgh, Scotland, U.K., Sep.99, Lecture Notes in Computer Science vol.1678, 79–99, Springer-Verlag.

[4] F.Arcieri, E.Cappadozzi, E.Nardelli, M.Talamo: Distributed territorial data management and exchange for public organizations, *3rd International Workshop on Advanced Issues of E-Commerce and Web-Based Information Systems (WECWIS'01)*, San Jose, Ca., USA, Jun.01, IEEE Computer Society Press, 2001.

[5] F.Arcieri, E.Cappadozzi, E.Nardelli, M.Talamo: SIM: a working example of an e-government service infrastructure for mountain communities, *Workshop on Electronic Government (DEXA-eGov'01)*, Conf. on Databases and Expert System Applications (DEXA'01), Sep.01, Munich, Germany, IEEE Computer Society Press, 2001.

[6] F.Arcieri, E.Cappadozzi, G.Melideo, E.Nardelli, P.Naggar, M.Talamo: A formal model for data coherence maintenance. *Int. Workshop on Foundations of Models for Information Integration (FMII'01)*, 10th Workshop in the series Foundation of Models and Languages for Data and Objects (FMLDO), Viterbo, Italy, Sep.01. Lecture Notes in Computer Science Vol., Springer-Verlag, 2001.

[7] F.Arcieri,G.Melideo, E.Nardelli, M.Talamo: Experiences and issues in the realization of e-government services. *Int. Workshop on Research Issues in Data Engineering (RIDE'02)*, San Jose, Ca., USA, Feb.02, IEEE Computer Society Press, 2002.

[8] F.Arcieri, G.Melideo, E.Nardelli, M.Talamo. A reference architecture for the certification of e-services in a digital government infrastructure. *Distributed and Parallel Databases*, 12:217–234, 2002.

[9] A.Bouguettaya, M.Ouzzani, B.Medjahed, J.Cameron: Managing government databases. *Computer*, 34(2):56–64, Feb.01.

[10] N.Brownlee, C.Loosley, Fundamentals of Internet Measurement: A Tutorial, *CMG Journal of Computer Resource Management*, 102, Spring 2001.

[11] Cisco NetFlow, http://www.cisco.com/warp/public/732/Tech/nmp/netflow/

[12] F.Casati, U.Dayal, M.-C.Shan: Report on the VLDB Workshop on Technologies for E-Services (TES), *SIGMOD Record*, 29(4):11–15, 2000.

[13] N.Duffield, C.Lund, M.Thorup. Charging from sampled network usage. ACM-SIGCOMM Internet Measurement Workshop (IMW'01), San Francisco, Ca., USA, Nov.01.




Proceedings of the 2003 Symposium on Applications and the Internet (SAINT'03)
0-7695-1872-9/03 $17.00 © 2003 **IEEE**

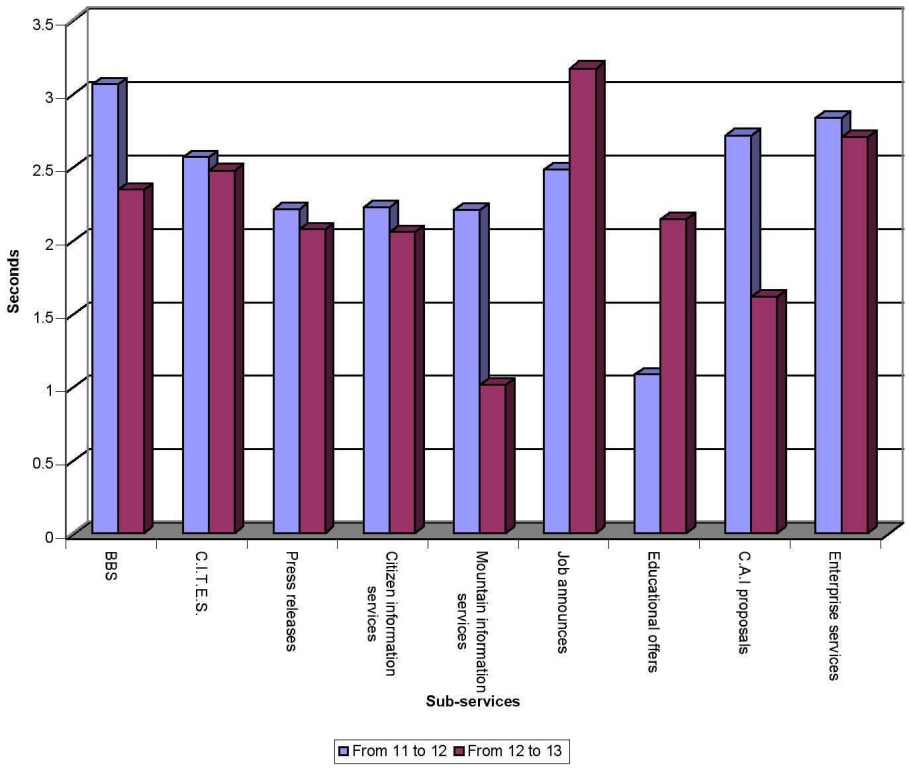


**Figure 3. Average SIM processing time for nine different subservices, during 11:12 (lighter) and 12/13 (darker) time bands.**


[14] A.K.Elmagarmid, W.J.McIver: The Ongoing March Toward Digital Government, Guest Editors' Introduction to the special section on Digital Government, *IEEE Computer*, 34(2):32–38, Feb.01.

[15] C.Estan, G.Varghese. New directions in traffic measurement and accounting. ACM-SIGCOMM Internet Measurement Workshop (IMW'01), San Francisco, Ca., USA, Nov.01.

[16] I.Foster, C.Kesselman (editors), The GRID: Blueprint for a New Computing Infrastructure, Morgan Kauffman Publishers, 1999.

[17] S.Gigandet, A.Sudarsanam, A.Aggarwal. The Inktomi Climate Lab: an integrated environment for analyzing and simulating customer network traffic. ACM-SIGCOMM Internet Measurement Workshop (IMW'01), San Francisco, Ca., USA, Nov.01.

[18] W.Hasselbring: Information System Integration: introduction to the special section, *Communications of the ACM*, 43(6):33–38, June 2000.

[19] T. Hacker, W. Thigpen, Distributed Accounting on the GRID, Grid Forum Working Draft.

[20] J.Joshi, A.Ghafoor, W.G.Aref, E.H.Spafford. Digital government security infrastructure design challenges. *Computer*, 34(2):66–72, Feb.01.

[21] Keynote Data Accuracy and Statistical Analysis for Performance Trending and Service Level Management. A Keynote White Paper http://www.keynote.com .

[22] B.Krishnamurthy, C.Wills, Y.Zhang. On the use and performance of content distribution networks. ACM-SIGCOMM Internet Measurement Workshop (IMW'01), San Francisco, Ca., USA, Nov.01.

[23] E.Nardelli, M.Talamo, and P.Vocca. Efficient searching for multidimensional data made simple. 7th Annual European Symposium on Algorithms (ESA'99), Prague, Czech Republic, Jul.99, Lecture Notes in Computer Science vol.1643, pp. 339–353, Springer-Verlag.

[24] C.U.Smith and L.G.Williams, Performance Evaluation of Distributed Software Architectures, em. Proc. Computer Measurement Group, Anaheim, CA, Dec. 1998.

[25] Bulletin of the Technical Committee on Data Engineering, *Special Issue on Interoperability*, D.Kossmann (ed.), 21(3), September 1998.

[26] Bulletin of the Technical Committee on Data Engineering, *Special Issue on Infrastructures for Advanced E-Services*, G.Weikum (ed.), 24(1), March 2001.

[27] J.Yang, M.P.Papazaglou: Interoperation support for electronic business, *Communications of the ACM*, 43(6):39–47, June 2000.




Proceedings of the 2003 Symposium on Applications and the Internet (SAINT'03)
0-7695-1872-9/03 $17.00 © 2003 IEEE