
21th NMRG Meeting in Utrecht (Netherlands)

The 21th NMRG meeting will be held in cooperation with [EMANICS](#) on Thursday October 19th and Friday October 20th in [Utrecht](#) (Netherlands), just before the [MANWEEK 2006](#) in [Dublin](#) (Ireland). The organizers of this meeting are [Aiko Pras](#), [Jürgen Schönwälder](#), and [Gabi Dreo Rodosek](#).

Title

Joint EMANICS / IRTF-NMRG Workshop
on
Future Direction of Network and Service Management Research

19-20 October 2006

Hosted by SURFnet
Utrecht, The Netherlands

Purpose

The area of network and service management research has undergone several changes. While some research topics have matured over the years into their own little areas (e.g. policy-based management or fault management), there are other areas that are more driven by advances in technologies. This includes technologies that need new management approaches but also technologies that can be utilized for more effective management.

The purpose of this jointly organized workshop is to bring together researchers, operators, vendors and technology developers to identify promising future directions of network and service management research. The outcome produced by the workshop should be a description of research directions that is felt worthwhile to explore in a timeframe of the next 5 years.

Format

The workshop will be limited to 30 participants. The organizers seek participation from researchers, operators, vendors, and technology developers.

Interested parties are invited to write up a short position statement (about one page plain ASCII text) explaining their specific background and how they plan to contribute to the workshop discussions. The deadline for sending position statements is September 15th.

If position statements of more than 30 people are received, the organizers will select participants with the goal to achieve a diverse view on the subject and to balance involvement of researchers, operators, vendors, and technology developers.

Location

The meeting is hosted by SURFnet, Utrecht, The Netherlands. [SURFnet](#) is the operator of the Dutch research network. SURFnet is located in [Utrecht](#), a city close to the [Amsterdam Schiphol international airport](#).

There is a train running from the airport to Utrecht which takes approximately 30 minutes.

The schedule of the trains running from Schiphol to Utrecht Central can be found [here](#). To find your way to the meeting place, please following these [directions](#).

Accomodation

- NH Utrecht Hotel (4 stars, 5 minutes walk)
Jaarbeursplein 24
3521AR Utrecht
Tel: +31-30-2977977
Euro: 140-240
- IBIS (2 stars, 20 minutes walk)
Bizetlaan 1
3533 KC UTRECHT
Tel: +31-30-2910366
Euro 82

For additional hotels, go to [ViaMichelin](#). Select on "Maps": City=Utrecht, Netherlands and "Nearby hotels".

Agenda (subject to changes)

The meeting will start on Thursday at 10:00 and it will close on Friday at 16:00.

Thursday (2006-10-19)

Time	Activity	Actors
10:00	Welcome and Administrivia	Chair(s)
10:15	Welcome Note of the Host	N.N. (Surfnet)
10:30	Position Statements	Participants
12:30	Lunch Break	
13:30	Position Statements	Participants
15:00	Coffee Break	
15:30	Break-out Sessions	Operators / Vendors / Researchers
18:00	End of Day #1	

Friday (2006-10-20)

Time	Activity	Actors
09:00	Break-out Sessions Results	Operators / Vendors / Researchers
10:00	Discussion	Participants
10:30	Coffee Break	
11:00	Discussion	Participants
12:30	Lunch Break	
13:30	Summary and Conclusions	Participants
15:00	Coffee Break	
15:30	Wrap-up	Chair(s)
16:00	End of Day #2	

Position Statements

In order to prepare the meeting, we collect position statements from all participants. The position statements will go online beginning of October. Any subsequent discussions before the meeting should take place on the NMRG mailing list.

1. [Claudio Bartolini](#)
2. [Mark Burgess](#)
3. [Prosper Chemouil](#)
4. [Petre Dini](#)
5. [Gabi Dreo Rodosek](#)
6. [Liam Fallon](#)
7. [Olivier Festor](#)
8. [David Harrington](#)
9. [James Won-Ki Hong](#)
10. [Simon Leinen](#)
11. [Giorgio Nunzi](#)
12. George Pavlou (**still missing**)
13. [Gregorio Martinez Perez](#)
14. [Aiko Pras](#)
15. [Danny Raz](#)
16. [Dan Romascanu](#)
17. [Jürgen Schönwälder](#)
18. [Rolf Stadler](#)
19. [Burkhard Stiller](#)
20. [Stefan Wallin](#)
21. [Bert Wijnen](#)
22. [Jae-Hyoung Yoo](#)

Participants

The list of participants below is not yet final. We still expect a few more people to join and attend the workshop.

1. Claudio Bartolini (HP)
2. Mark Burgess (University College Oslo)
3. Prosper Chemouil (Orange France Telecom)
4. Petre Dini (Cisco Systems)
5. Gabi Dreo Rodosek (University of the Federal Armed Forces)
6. Liam Fallon (Ericsson)
7. Olivier Festor (LORIA - INRIA Lorraine)
8. David Harrington (Huawei)
9. James Won-Ki Hong (Pohang University of Science and Technology)
10. Simon Leinen (Switch - Swiss Education and Research Network)
11. Giorgio Nunzi (NEC Network Labs)
12. George Pavlou (University of Surrey)
13. Gregorio Martinez Perez (University of Murcia (UMU))
14. Aiko Pras (University of Twente)
15. Danny Raz (Technion)
16. Dan Romascanu (Avaya)
17. Jürgen Schönwälder (International University Bremen)
18. Rolf Stadler (KTH)
19. Burkhard Stiller (University Zurich)
20. Jean Theunissen (Tiscali)
21. Bert Wijnen (Lucent)
22. Jae-Hyoung Yoo (Korea Telecom)

NMRG/EMANICS Workshop Position Statement

Aiko Pras <a.pras@utwente.nl>
University of Twente

In the next decade the main challenge in networking will be how to improve dependability of the IT infrastructure, including the services running on top of it. The issue will no longer be how to make under good weather conditions networks a few percent faster, but how to ensure, even in case of "bad weather" (DOS attacks, failing equipment) a predictable service on which users can rely.

Dependability is a goal, however, and not a kind of mechanism that can easily be implemented. In this sense dependability is like security, performance and scalability, which are also goals. The mechanisms (protocols, algorithms etc.) needed to make an infrastructure dependable will vary from case to case. It is essential, however, that researchers start from real problems, as identified by real managers, and not from artificial problems that look well on paper but have little relationship with reality. Examples of real problems are: how to keep viruses outside, how to fight spam, how to immediately identify system components that do not perform well, how to find inconsistencies in my configuration etc.

NMRG/EMANICS Workshop Position Statement

Name: Bert Wijnen
email: bwijnen@lucent.com

I have been active in Systems and Network Management for a pretty long time now. After having served as the lead-systems-programmer for IBM mainframe systems (supporting all of IBM Europe, Middle East, Africa) for a number of years, it seemed that my operational experience and knowledge of system internals made me a good candidate to look at the more generic Automated Operations (Dark Room operations) for IBM EMEA wide internal systems.

That "Automated Operations for IBM datacenters" work started around 1985 and it got me in touch with IBM T.J.Watson Research to work on the computing systems to automate a lot of the operational/management aspects of mainframes and the network that connected them with the end-users (mostly on dumb terminals at the time). It also got me exposed and involved in some of the IBM Mainframe NetView products, which were used for both Network (mainly SNA) and Systems management.

In 1988, IBM got involved in contributing tools/software for managing the NSFNET, controlled from the NOC at MERIT in Ann Arbor (USA). IBM first suggested and tried to manage the NSFNET (IP network) using a mainframe NetView system with interfaces to (among others) SNMP. I was instrumental in that piece of work. That exposed me to many of the tools/software used in the Unix world. With a team we first developed a new object-oriented system (DRAGONS, much like a distributed JAVA environment) that also used SNMP to speak to the managed devices and protocol stacks.

From that project I got involved in IETF in the Network Management area and I contributed quite a lot over the years in spec-development of SNMP and AgentX. At the same time I implemented SNMP (all versions, both agent and manager side) as a SNMP platform that could be (and has been) used on basically all systems being shipped by IBM.

Later (from 1998 onwards) I contributed to the IETF in my role as co-Area Director for the Operations and Management Area. In 2000 I switched from IBM to Lucent Technologies (unfortunately I did not write much if any code since then).

In my role of IETF co-Area Director for the Operations and Management Area I had to manage the IETF volunteers on deciding between various technologies (SNMP vs. COPS, MIBs vs. PIBs, and others). During this time I got exposed to many different, but also many overlapping and competing techniques, approaches and viewpoints. The proliferation of Internet technologies also allowed too many ideas to just ignore others and so a lot of overlap and competition resulted. Good in the sense that we learn from competing technologies, Bad in the sense that we also wasted a lot of resources. And in addition, we put a lot of burden on Operators who had to understand and use all these technologies. From what I have seen over the years, I am now convinced that if we have N technologies today, that the Operators will have to deal with N+1 technologies every time that Researchers and Standards Organizations come up with yet another technology/protocol/mechanism that will be "the best thing since sliced bread and that will obsolete all earlier technologies/protocols".

As such, my main contribution will probably be to take a devils advocate position and challenge the need for anything new that will be presented.

Bert

Pre-paid Schemes Considered Harmless for All-IP Networks

Burkhard Stiller

University of Zurich and ETH Zurich, Switzerland

Future all-IP networks will support a variety of high-value IP-based services including the upcoming integration of today's non-IP services offered by mobile telecommunication providers. Besides post-paid charging, pre-paid charging will play an important role in all-IP environments due to the financial control of providers and customers, thus, determining effective economic management means for Internet services. Therefore, the enhancement of service management functionality by such economic mechanisms determines one important area of research, especially since the support of such mechanisms within an Internet without any business-related restrictions is essential for the provisioning of open standards and technology.

Since existing pre-paid charging solutions do not meet key all-IP requirements imposed, a novel approach termed TICA (Time Interval Calculation Algorithm), including three algorithmic variants has been developed. TICA introduces the concept of service bundles which are composed out of different services. For consecutive time intervals, TICA estimates the maximum resource consumption of a vector of any pre-determined and service-related parameters, e.g., available from an accounting system. While the time is divided into intervals, the resource estimation is performed by TICA. Finally, TICA calculates the maximal charge for all services of that service bundle, yielding the total maximal charge per service bundle in that time interval under consideration.

To show TICA's optimized performance for many service mixes foreseen, a thorough evaluation has been performed, being based on different scenarios. Additionally, comparing TICA's overall performance with existing online charging solutions yields TICA's promising results in terms of a way smaller quantity of credit checks required and smaller financial risks for providers to be achieved. Finally, the three different TICA algorithmic variants reveal a clear sensitivity to input parameters defined. These findings indicate that a number of suitable parameter settings can be used by today's Internet Service Providers already. Thus, TICA determines a valuable extension to current service management approaches.

NMRG/EMANICS Workshop Position Statement

Claudio Bartolini <claudio.bartolini@hp.com>
HP Laboratories

I'm Claudio Bartolini
(http://www.hpl.hp.com/personal/Claudio_Bartolini/index.html), a senior researcher with HP Labs Palo Alto. (However, given the European spin of the meeting, I'd like to let known that my team is based in Bristol, UK). My background is on architecture and design of software systems and frameworks.

My current research is centered on business-driven IT management (BDIM), and this is the angle that I would like to bring into the discussion at the meeting. The goal of BDIM is to to enable an enterprise to manage its IT services in accordance with its business objectives. BDIM focuses on the impact of IT on business processes and business results and vice versa; besides the conventional IT metrics such as availability and response time, it looks at key performance indicators (KPIs), that is metrics that have significance from the point of view of the business supported by the IT. The BDIM approach aims at rethinking IT management from a business perspective, whether this be in an operational, tactical or strategic context.

Below is a description of my medium term research plan, which I have begun to execute on together with my colleagues at HP labs and external collaborators at various universities in Europe and South America.

Within an enterprise, the Information Technology (IT) function must deliver value by aligning itself with the objectives of the business that it supports. IT service management frameworks such as COBIT [1] and ITIL [2] define best practices and processes that support the IT function in the smooth running of their operations. However, neither these frameworks, nor any of the commercially available software tools today can help an IT manager plan and execute their courses of action by taking into account the risk of effecting the actions and their impact on the business.

The objective of my research plan is to provide an approach to business-driven IT service management based on modeling the impact on the business of actions taken on IT systems and processes. Building on the modeling activity, I aim to apply existing and/or novel decision support and automation techniques in order to suggest and help execute the actions that minimize the impact on the business.

Research plan

- * Review the state of the art of the academic and industrial research in technologies for business-IT alignment
- * Produce of a conceptual framework for business-driven IT management
- * Design and implement a software framework for business-driven IT management complete of modeling tools, decision support tools, and automation tools
- * Apply the tools and technologies for decision support and automation to at least one real-life scenario of IT service management

Research questions

The main research questions that I'm expecting to address in the course of my research will have to do with modeling, decision support and automation.

Modeling

Both the production of the conceptual framework and the design of the software artifacts will require formal models of the description of IT service management processes, and other related IT objects such as people technology and processes involved. I'm expecting that in the course of the modeling research activity I will tackle research questions similar to the following:

- * What are the key IT metrics that are of concern to IT managers?
- * What are the key performance indicators (KPIs) that measure the performance of IT service delivery and operations with respect to the value they contribute to the business supported by the IT?
- * To what degree the concerns of the business are addressed by measuring the level of alignment of observed metric to short term (e.g. service level agreements or SLAs) and mid/long-term (e.g. balanced scorecard metrics) objectives?
- * How to model the dependencies between IT metrics and business objectives, business rules and policies, using a variety of modeling techniques (logic-based, probability-based)?
- * How to link these metrics and KPIs to models of organization, objectives, IT facilities, people skills, application systems, technology, and data?
- * How to estimate risk and impact on the business of activities related to IT management processes?
- * What information models exist today that can provide a basis to start answering the questions listed above? (Possible starting points are COBIT and ITIL stated above, but also DMTF's CIM [3] - common information model, and possibly advanced modeling concepts from the semantic web, e.g. W3C's OWL [4], the web ontology language.)

Decision support

Providing decision support in order to maximize the alignment with business objectives requires determining the tradeoffs that exist between conflicting business objectives and analyze them to choose between alternative options available to an IT manager (for instance alternative service and process configurations) such that the optimum alignment with the objectives is achieved with respect to current business and system/resource conditions. The main research question to be addressed here is

- * What fundamental techniques for decision support (such as forecasting and prediction methods, mathematical optimization, simulation, etc.) can be applied to the most important decision to be made within the IT service management processes?

Automation

To close the IT management loop and act on the decisions made by the IT managers, it will be necessary to find answers to questions such as:

- * What fundamental techniques for automation (such as planning, scheduling, automated capacity planning and resource allocation), can be applied to increase efficiency in the execution of the IT service management processes?

References

- [1] COBIT - www.isaca.org/cobit
- [2] ITIL - www.itil.co.uk
- [3] DMTF CIM - www.dmtf.org/standards/cim/
- [4] W3C OWL - www.w3.org/TR/owl-features/

Management Research Position Statement

Dan Romascanu, Avaya

Background:

- Area Director, IETF Operations and Management Area
- Software Engineer, R&D Team and Project Leader, Systems Architect at LANNET, Madge Networks, Lucent and Avaya (current employer)
- Participant, editor, and working group chair in various IETF Network Management activities, author and co-author of several IETF MIB and management related documents
- Participant in management related standards development in the IEEE, TIA and MEF
- Designer and software developer for various management products including embedded device agents, management applications and management protocols implementations

Interest and contribution in the NMPR meeting dedicated to Future Direction of Network and Service Management Research:

- represent and present the position of the IETF Operations and Management area
- present proposals related to future protocols, data models and applications research as perceived in the IETF operations and management space
- present issues related to the collaboration and relationship between the work in the IETF, NMRG and other Standards Development Organizations
- present ideas and provide inputs related to management of specific areas of technology and expertise (configuration protocols and data models, performance monitoring, VoIP and real-time applications)

To manage or not to manage: Addressing the benefit
overhead tradeoff in network management

Danny Raz, Technion

The increased complexity of networking infrastructure and protocols and the desire to provide high quality services at the lowest possible cost, drive many organizations to deploy more network and system management tools in their networks. It is often argued that due to the high complexity of management, a much more cost effective way to assure performance is just to acquire more resources.

This is particularly true for performance management of Information Technology (IT), where the goal is to coordinate networked resources in such a way that the business-level objectives are met at all times, at a lowest possible cost, and with optimum capacity. As Service-Oriented Architecture (SOA) spreads as a popular way of organizing and providing distributed capabilities to solve business problems, cost-effective performance management becomes essential. Practicing IT administrators know well that committing more resources to management improves the overall quality of service up to a certain point, after which management costs start dominating the total cost of ownership and management off-sets its own advantages. Thus, although it is naturally desirable that the network would perform at the highest possible level, this may not be the best solution due to the associated cost.

Time is now mature for the research community to address this fundamental tradeoff in a rigorous way by showing exactly how much effort should be invested into management to gain the maximal benefit. In order to do that, one should accurately define both the cost associated with the management process and the expected benefit. Of course, considering the overall benefit of general management systems and all aspects of the associated overhead may be impossible due to the variety of different aspects involved and different network conditions. However, when applying to specific tasks, within the network management domain, one can rigorously define this tradeoff, and then provide a general tool to find optimal working points for such systems.

Consider for example a service that is being provided by a set of servers over the network. The goal of the service provider is to provide the best service (say, minimizing the service time) given the amount of available resources (e.g., the number of servers). The provider can add a load sharing system (for example as suggested in RFC 2391) and improve the service time. However, the same resources (budget) could be used to add additional servers to the system and thus provide better service to end customers. The dilemma here is between adding more computational power and adding management abilities, where the goal is to achieve the best improvement in the overall system performance. Note that in order to be effective, the load sharing system needs updated load information from the servers. Handling such load information requests requires small but nonzero resources (e.g., CPU) from each server. Thus, it is not easy to predict the actual amount of improvement expected from preferring a specific configuration. Yet, for this concrete example, one can formalize the cost and expected benefit and define an optimal working point.

As indicated by this example, it is important to identify just the right amount of resources that should be allocated to management tasks (such as monitoring) in order to maximize the overall system performance. In addition to being an important and interesting research direction, this approach can be proven to provide practical tool that can help in providing cost effective services to the community.

Management Research Position Statement

David Harrington, Huawei

I am an IETF-focused person, and I see a disconnect between the requirements of the telecomm industry, which is beginning to dominate the IP-based Internet, and the management protocols that exist in the IETF (e.g. SNMP) and are being developed in the IETF (e.g. Netconf).

I think research needs to be done about how the IETF datacomm-based management approach needs to change to better accommodate the management of emerging networks and services that operate over the Internet. This would include inter-domain management, peer-to-peer sharing of mgmt info, and standardization of mgmt above the router/switch/device level.

A second area of research is solution reuse (especially in the IETF). The use of reusable (standardized) components has been shown to be highly beneficial in equipment design and software development, reducing development costs and improving consistency across products and across vendors. But surprisingly, it is unclear whether component reuse is as beneficial during the process of developing standards in an organization like the IETF.

The IETF is an engineering organization, and the problems of reuse in developing standards are not well understood. It would be appropriate for the NMRG, in particular, to analyze the issues involved and provide guidance to the IETF about the benefits and costs of trying to move to reusable components in the development of IETF management standards.

Part of the research into reusable components should also consider whether moving to a component-based network management "architecture" or approach could better meet the emerging needs for services and multi-domain management.

David Harrington
dharrington@huawei.com
dbharrington@comcast.net
ietfdbh@comcast.net

Service Level Management: From Quality of Device to Quality of Service

Gabi Dreo Rodosek, University of Federal Armed Forces Munich

By moving from device-oriented to service-oriented management we have to face several new challenges. Some of the challenges refer to the formalization of Service Level Agreements (SLAs) and the question how to obtain the so-called Quality of Service (QoS) parameters such as availability out of Quality of Device (QoD) parameters such as IfOperStatus.

The formalization of SLAs is still an open issue. An approach is to formalize SLAs based upon WS-Agreement specifications that allow the specification of service functionalities as well as adequate QoS parameters to quantify the service quality. However, open issues refer to the questions what are the elements of a SLA as well as what are the relevant QoS parameters for specific services. A generic approach is still missing.

When dealing with QoS parameters the challenge is to provide customer- and service-oriented parameters such as response time if accessing a web site instead of device-oriented parameters such as InOctets. The challenge is to provide "what a customer wants to see" and not "what a provider can provide wrt. management tools and techniques". Thus, this gap can be filled by approaching from a top-down (e.g., the vision) or a bottom-up (e.g., state of the art) perspective. Issues like what are the basic data, the appropriate metrics, the relevant QoS parameters, the various reporting groups (e.g., managers, operational stuff, customers) as well as the measurement approaches need to be addressed.

21th NMRG Meeting, statement for future research directions in network management from NEC - Network Laboratories, Heidelberg

Giorgio Nunzi <Giorgio.Nunzi@netlab.nec.de>

Distributing management has been one of the dominating paradigms in the research community during the past decade. Autonomic computing showed up as a new reference architecture in the last years and self-x technologies have been investigated in different forms.

Interests are now shifting towards the technology transfer of these results to products:

- Mesh networks have clear deployment scenarios and have been realized as pilot tests or operative networks.
- The Beyond 3G effort finally defined the new architecture for mobile telecommunication networks: it is clear that it is based on a distributed paradigm and many management functions will be delegated down to the base stations.
- Plug-and-play and adaptation are technology values added to many products to attract operators' interests for costs savings in installation and maintenance.

These examples shows that self-x technology are receiving interest and certainly will be transferred into future network solutions. In this scenario, the future challenging questions to address will be: Can the transfer of self-x technology be as easy as a pure adoption? Can normal management processes be executed on top of the new architecture?

The problem arises from the conviction that many management functions will still be maintained with centralized processes. Depending of the application scenario, this will probably concern fault management, goals enforcement, accounting, etc.

Therefore, a bridge must be built between the two worlds, the population of self-managed functions and the centralized functions. The instruments required to do this can vary from correlation between different nodes, to filtering autonomous functions and aggregation of information. The questions to pose to the research community are: How can we extract information from self-managed nodes? How can we correlate differences between them? How do we present management information to human operators in an easy way?

Traditional instruments might not be adequate and might require further study. It is unfeasible to put classical management interfaces on top of autonomic functions. Data mining techniques might not reveal efficient to discover a behaviour over a distributed self-managed network. The co-existence of different management interfaces is likely to be accentuated with the use of self-x functionalities. On top, the human interface needs to be revised as well to cope with all the possible approaches used in network management.

A challenging direction is towards an integrated instrument to cope with different information elements, types of interface, correlation policies and methods to enforce decisions. Will the winning approach in network management rely on a "Google-like" approach, integrating self-x technologies with central management systems?

Position Statement submitted by Gregorio Martinez

1. Institution: University of Murcia (UMU), Spain
2. Department: Information and Communications Engineering
3. Research Group: Intelligent Systems
4. Contact people: Antonio F. Gomez Skarmeta <skarmeta@dif.um.es>
Felix J. Garcia Clemente <fgarcia@dif.um.es>
Gregorio Martinez Perez <gregorio@dif.um.es>

5. Background:

The UMU Intelligent Systems research group (ANTS, <http://ants.dif.um.es/>) is working on network and service management from early 2000; initially they started working on policy-based management applied to different networking scenarios; in this sense, the Euro6IX (European IPv6 Internet Exchanges backbone, <http://www.euro6ix.org/>) IST FP5 project was serving as a very first contact point where the PBNM paradigm was playing a key role for the management of security and routing networking services in IPv4/IPv6 networks. Later the SEINIT IST FP6 project was evolving this paradigm towards the provision of scalable management of security services in inter-domain scenarios.

Two are the main EU IST FP6 projects who serve now as background for the research on management of security services done by our group; the first one is POSITIF (Policy based Security Tools and Framework, <http://www.positif.org/>) who proposes the design and implementation (as contribution to the open-source community) of a framework for managing security policies. The second one is DESEREC (Dependable Security by Enhanced Reconfigurability, <http://www.deserec.org/>) whose main focus is to manage critical information systems with the intention of increasing their level of dependability.

6. Current and future research lines: current work is focused on different lines; some of them are now outlined:

[Semantic-aware Security Policy Language] The work in this line is aiming at providing one step towards the objective of detecting conflicts in the specification of the security policies to be applied to a target system or network. The context for such research is the deployment of a semantically-rich checking component able to detect any inconsistency that may exist in a set of policy rules when applying them in a target system. The current approach is based on using OWL to allow the expression of elements and relationships in a distributed system. This offers two possibilities: validating the coherence and integrity of the model and also querying the model about instances which satisfy some properties of interest. On the other hand, SWRL allows specifying policies in the form of conventional if-then rules. This is an important shift in the paradigm as it opens the door to the automatic detection and fix of conflicting rules.

[Conflict analysis] Current work in this line goes in the direction of defining a formal framework for automatic conflict

resolution. We are using the powerful metaphor of intelligent agents to model the problem. As we understand it, a distributed system is managed through a distributed set of policies. We model the problem as a multi-agent system in which agents have its own and non-shared knowledge base. In this situation, conflicts may arise when, for example, an agent decides to grant some privilege to a user and a group of other agents deny the same privilege. In this case, conflicts should be solved through negotiation by means of argumentation. In this context, argumentation means trying to persuade the rest of agents that your reasoning process is more plausible than the other's reasoning process and, in consequence, reach an agreement about, in our example, privileges granted on the user.

[Design of a Web Service-oriented management framework] This work is focused on the definition of WS-based mechanisms enabling the dynamic management of security blocks; this framework is addressing those functionalities required to check, transform, distribute, enforce and monitor the security configurations that should be applied to the target systems. Asynchronous notification of relevant events (e.g., security breaches, attacks, etc.) in such kind of architecture is also being researched.

[Management of IDS/IPS systems] Current state of the art on IPS/IDS does not provide a common understanding on the knowledge being used in the detection and prevention processes; this research line is focused on defining the building blocks of IDS/IPS systems as independent services sharing a common information base. This work is also currently addressing how an attack can be modelled and managed as part of distributed hybrid IDS/IPS systems.

[Self-management of security services] This research line is attempting to analyse the need of self-management in the context of security services and proposing design paradigms based on the different theories already existing or the new ones to come. There is also an intention to relate these designs with current existing protocols and communication architectures.

[Management of Critical Information Systems] This research line is focused on defining mechanisms to respond efficiently to different kind of incidents which can occur on a critical system (attack from the outside, intrinsic failures and malicious internal use); this response will be based on a three-tiered response to exceptions and incidents.

7. Scenarios: several are the scenarios where this research work is being applied; examples of them are:

[Grid computing] Globus Toolkit 4 (GT4) provides effective resource management for the grid-computing environment; it includes security services, but lacks security policy management services. In this sense, our research is intending to serve as complement to the GT4 system, providing a wide range of security management capabilities that usually rely on platform-specific enforcement mechanisms.

[Distributed Firewalls] Our research in this area is mainly intended to make use of the concept of distributed firewalls, where firewall policies are centrally defined by an administrator (or a set of them) and enforced at each individual network endpoint, not only in a central location. This research is part of a bigger effort on distributed security management for IP networks where authors are analysing the current network-centric security model being used in Internet, identifying its limitations and proposing a new host-centric model (based on the concept of distributed

firewall) to solve some of the current issues.

8. Selected papers: a few selected recent publications are:

- * Gregorio Martínez Pérez, Gabriel López Millán, Félix J. García Clemente, Antonio F. Gómez Skarmeta, "Dynamic and Secure Management of VPNs in IPv6 Multi-Domain Scenarios", Elsevier Journal of Computer Communications, to appear in October 2006.
- * Gregorio Martínez Pérez, Antonio F. Gómez Skarmeta, Steve Zeber (1), Joe Spagnolo (2), Tim Symchych (3), "Dynamic Policy-Based Network Management for a Secure Coalition Environment", IEEE Communications Magazine, to appear in November 2006. This work is published in collaboration with (1) Defence R&D Canada, (2) NRNS Incorporated Canada and (3) Communications Research Centre (CRC) Canada.

Position Statement

A Telco's Perspective on Future Research Direction of Network and Service Management

Jae-Hyung Yoo

Assistant Vice President, Network Technology Lab., KT, Korea (styoo@kt.co.kr)

Recently, the operational environment of telcos has become more complex and heterogeneous with the continuous introduction of new technologies. Also, due to fierce competitions, telcos need to provide more advanced services and higher quality of services at lower prices to customers. To meet the needs of customers and save operational expenses, telcos are seeking more advanced network management technologies and methodologies. In this statement, we are not proposing any new specific solutions but rather introduce the field operation requirements that should be considered in future research direction of network and service management.

There are two important aspects that are often ignored or neglected in developing network management systems (NMSs). The first aspect is the requirement for new network management functions from the view point of network operators. The second aspect is the requirement for customer-oriented management functions. Below, we elaborate the required network operator oriented and customer oriented management functions from the KT's perspective.

1 Network Operator Oriented Management Functions

Network operators are requesting several important functions that are required to simplify and speed up their operations. The network operator oriented management functions are 1) integrated management, 2) automation of workflows, and 3) automation of diagnosis functions.

1.1 Integrated Management

In KT, a network operator is responsible for managing many different types of network elements, such as IP routers, transmission equipments, ATM switches, electric power systems, etc. Over the past 20 years, many NMSs and EMSs have been developed, but because the management systems are designed to manage specific domains, operators have to handle many different GUIs. Most manager GUIs have their own menu trees and specifically designed windows. Thus, it is very difficult to monitor the whole network status by accessing many different GUIs. The operators are asking for an integrated NMS which can monitor heterogeneous network elements in an integrated view.

1.2 Automation of Workflows

Many OSSs/NMSs have been developed and applied, but only about 15% of the work in network operator centers (NOCs) is automated and the rest is still done manually. The major manual job is diagnosis of root cause, preparing daily logs and statistical reports, etc. Especially, it takes a long time to retrieve various data from many OSSs and NMSs to make a report. Thus, we believe the NMS should evolve to be the core of workflows in NOCs.

1.3 Automation of Diagnosis Functions

When an alarm is reported from a network device, an operator tries to find the root cause of it through the use of many CLI commands. Since most of the NMSs were designed to monitor and control alarm and traffic flow in company-wide scale, NMSs do not have the functions to support detailed diagnosis functions. Perhaps because it is defined

that control is automatic and management involves human. However, there is much need to automate diagnosis functions, which is somewhat between control and management plane.

2 Customer Oriented Management Functions

Today's customers are demanding more information and analysis capability than ever before. The two important customer oriented management functions are 1) analysis of service interaction and 2) analysis of services affected by network and server faults.

2.1 Analysis of Service Interaction

In the near future, traffic should be managed from a customer experience perspective. Problems can occur with interacting services within media streams that have nothing to do with the transport or access network. So, service providers are introducing deep packet inspection technology to look into media streams to monitor service behavior and diagnose very specific problems a customer might experience. Let's suppose a customer is connected to the Internet with a 10 Mbps leased line. When N people are using VoIP phones and M people are using IP-TV, and the performance of both services is degraded.

Q1: Who is responsible for the degradation of service quality?

A1: The service provider who sold VoIP and IP-TV services.

Q2: Among the N VoIP and M IP-TV channels, which one is degraded and how is the quality is degraded?

A2: We don't know yet.

Q3: If some packets are dropped by priority control, how can the service provider explain the quality degradation to customers?

A3: It's a good question!

These questions are important but frequent questions from customers. Future NMSs should be able to provide answers to them.

2.2 Analysis of Services Affected by Network and Server Faults

In the near future, NMS should evolve to provide end-to-end visibility from the service servers to customer terminals. The fault management function should provide more detailed and useful information. That is, the alarm severity has to be defined in proportion to the level of affected services. The list of affected customer and service information should be provided to the network operator, and if a fault occurs in a service server, the number of affected customers and degradation level of service quality should also be provided.

One can think of techniques such as service-oriented architecture (SoA), policy-based network management (PBNM), packet capture and flow analysis, and artificial intelligence (AI) might be the solution to satisfy these requirements. These have been studied for many years and reasonable solutions have been developed. However, problems still exist and operators and customers asking for solutions!

Position Statement

Operations and Management of 4G Networks and Services

James Won-Ki Hong, Dept. of CSE, POSTECH, jwkhong@postech.ac.kr

Driven by the need to support a variety of mobile, personalized, context-rich multimedia services and applications, the fourth generation (4G) mobile communication systems have been under development by researchers and vendors around the world. 4G systems are expected to be deployed and operational within the next several years. With the advances in networking, multiplexing, scheduling and physical layer technologies, 4G systems are also expected to provide higher bandwidth to more users and thus provide more cost-effective services than currently existing systems.

The 4G networks are expected to co-exist and inter-work with existing 2G and 3G mobile communication systems as well as satellite, wireless LAN (WLAN), and IEEE 802.16e (WiMAX and WiBro) networks, all interconnected through the service provider IP backbone networks and the Internet. Effective, secure and efficient operations and management of the envisioned 4G network environment is a huge challenge. Traditional network operations and management methods and protocols such as CMIP, TMN, SNMP, WBEM themselves are not going to be sufficient enough to support such complex communication and service environment. New, intelligent, and self managing operations and management architectures and methods are needed to meet the challenges.

In the 4G network environment, the traditional challenges related to FCAPS (fault, configuration, accounting, performance, and security management) still exist. New challenges related to FCAPS and new functional areas such as mobility management, customer management and terminal management exist.

More research needs to be carried out on the following new challenges:

- o Fault Management - automatic, end-to-end fault detection, location and analysis
- o Configuration Management - dynamic device configuration/re-configuration, network discovery & selection, and resource provisioning & management
- o Accounting Management - authentication, authorization and accounting (AAA), usage-based, content based charging & billing management, and session management
- o Performance Management - end-to-end QoS monitoring and management, SLA
- o Security Management - information security and network security
- o Mobility Management - domestic & global roaming, horizontal & vertical handoff management
- o Customer Management - customer subscription management and profile management
- o Terminal Management - autonomic (i.e., self-x) terminal management, location management and terminal trace management

What is Network Management Research?

Network Management Research Position Statement

Juergen Schoenwaelder, IUB

I. Introduction

My position statement is split into two parts: First, I like to discuss what I believe are research questions that have (a) high practical value and (b) require do some real substantial research efforts. However, understanding that pure research is rare and money has some importance to do research, I also like to discuss what I believe the network management research community needs to do in order to enable more research on network management.

II. Research Questions

Here are my favorite research topics. They are not fancy or buzzword compliant, but I believe we really have not solved them. Lets first look at the research question for traditionally managed networks and services:

o Configuration

The way many networks are configured and operated is ad-hoc and rather chaotic. Mankind created many rather complex artifacts, but most of them are well engineered in order to make the work reliable. And engineering often means to restrict what can be done with the many knobs that you in principle have in a complex system. Many networks, especially multi-service enterprise networks, are run by turning way too many knobs which makes the resulting network fragile and subject to constant change. You would not run an airbus this way.

Researchers and operators never identified "golden" engineering rules that tell people "this is the preferred way to run your network" which then are widely adopted and simplify the operations. Instead, whenever there is a new way of approaching a problem, people find it worthwhile do follow that path (even if the price is yet unknown).

I believe it is necessary to identify "golden" engineering rules (or network design pattern) and to code them up in software tools that actually start to automate the network configuration process (much like software package management systems automate much of the software installation and maintenance process) and make network behaviors understandable and predictable (even if that comes at the price of some lost "cool" optimizations).

o Event Correlation

Event correlation or sometimes called root cause analysis has been an old topic in the network management community. Given the steady increase of events we have to deal with, it seems that more work in the area is still needed. But note that it is also required that research acknowledges prior art in this area and does not reinvent the wheels from scratch.

o Distributed Management Algorithms

Distributed management has been around as a research topic for quite some time. Most of the work so far focused on supporting

technologies. The algorithmic aspect has seen much less attention (until recently). I believe that the distributed algorithms community has developed over the years many interesting algorithms that are waiting to be applied to network management. Note that I am talking about algorithms; I am not talking about the application of the latest distributed systems technology to network management (which I find rather boring lately).

o Security and Trust Management

This is probably the hardest topic of all. Some revolutionary research projects exist where people want to design the whole Internet from scratch with security build in from day one. While this radical approach has some merits, it might not lead to deployment anytime soon. So the question for people who draw satisfaction from research that can lead to real-world deployment is how we can orchestrate the security point solutions we have got into something that actually guarantees certain security properties.

o Understanding and Modeling

For all of the research efforts mentioned above, it is crucial to develop an understanding how networks work and to derive models from operational data. Without such established models, researchers are often forced to make arbitrary assumptions and thus they can justify any idea as a great invention. In other words, research to understand network behavior in general and management interactions in particular is enabling research in order to address the points listed above.

If we look further ahead, we will see an increasing number of networks and services that are inherently "unmanaged" since management functions are dealt with by the networks or services themselves. It is very important to do exploratory research in this area. Self-organizing management overlays which provide relevant informations to applications and services to adapt themselves are an example of such exploratory research projects.

III. Enabling Research on Network Management

Money is needed to do research. There are different funding sources, like government research funding institutions, cross-government research funding institutions (e.g. the ED), foundations, and industries. The funding institutions all have different cultures and decision processes, but in general there is a higher chance to obtain money for research if one does research on a topic that falls into some big mainstream research programs.

My experience is that it is relatively difficult to obtain research funding in network management. Furthermore, it is difficult to find people who are highly qualified to do research in this area. Why?

Network management is often not seen as an established research area, even though there is an active research community which is running an established series of events. The main reason, in my view, is the failure of this community to turn research results into any standard text books which define what network management research is all about. This is a clear failure of the community.

The second problem is that it is difficult to find PhD students which are on the one hand oriented towards sound research but at the same time have enough practical experience to understand the problem domain they are working on. This leads relatively often to research

results that are practically pointless or to rather week research contributions since people focus on building a tool without doing much research on the topic.

In short, my thesis is that Network Management (and Service Management) Research does not have a "face" and you will receive different answers to the question stated in the title of this note whenever you pose this question. The lack of a "face" makes it difficult to enable research in this area.

Bio:

Juergen Schoenwaelder is working at the International University Bremen, specializing in the field of communication networks and distributed systems. His research interests are network management, distributed systems and network security. He is an active member of the Internet Engineering Task Force (IETF) where he has edited more than 20 network management related specifications and standards (and those which took most of the time were the least successful ones).

He is the initiator and chair of the Network Management Research Group (NMRG) of the Internet Research Task Force (IRTF) and serving in various roles in the organization of IEEE/IFIP workshops, conferences, and journals.

On the technical side, he can be considered an "SNMP veteran" and coder of network management toolkits and applications. He has always been involved in the operation of University computer networks and he tries to understand not only the theory of network management, but also the practice of it.

Position Statement for the Joint EMANICS / IRTF-NMRG Workshop on
Future Direction of Network and Service Management Research

Network Management Research Centre, NMRC
Ericsson Ireland

BACKGROUND

The Ericsson portfolio of Telecom Management systems provide advanced tools for efficient management of GSM, WCDMA, & CDMA radio networks, core networks both circuit switched and packet switched, service networks, transport networks and so on. The network management functionality is developed based on our expertise in network technologies and in network operations. Ericsson is the largest provider of Telecom Management systems worldwide. Ericsson is a large-scale user of network management systems as part of its managed services business.

The role of the NMRC is to provide Ericsson with network management architectures, concepts, applications and competencies of strategic importance to secure long term product and service offerings in the Network Management domain.

MANAGED DOMAIN

Managed networks of all types are increasingly characterised by their inherent and growing complexity.

First, there is growing complexity introduced by the move to a packet based (IP) infrastructure and the proliferation of different access types and technologies. The increasing number and types of nodes and information, the high rate of change and diversity of network technology types, and continuing fixed-mobile convergence provide a challenge to today's management systems. The scale of managed networks is set to increase significantly as customer premises equipment and even terminals comes into the scope of management.

Second, there is service complexity today that did not exist in the past. Rapid introduction of new services aided by the advent of IMS puts high demands on the service delivery and service assurance organizations of operators as they strive to ensure a high quality of service. As service delivery often extends over several different networks, the demands of provisioning, management and quality control become more complex and difficult to manage.

Third, there is more complexity in the business conditions and environment, as competition hardens and price pressure increases.

Ericsson considers that, as the networks, services, and business needs become more fragmented and more complex, competitive advantage will accrue through advanced and flexible management systems and applications which are scalable, robust, and largely autonomous.

FUTURE DIRECTIONS OF NETWORK MANAGEMENT RESEARCH

The Ericsson NMRC believes that network management systems must evolve to become more scalable, flexible, open, self-configurable (on network level), autonomous, reliable, robust, and easier/more intuitive to operate. In order to achieve this goal, we identify the need for distributed management as being critical. We further identify the following research domains as having priority for investigation and exploration in the next 5 years:

- o Distributed Data Management Systems
- o Discovery

- o Openness and Flexibility
- o Scalability
- o Reasoning and Learning
- o Automated Deployment
- o Network and Data Visualization
- o Self Managing Systems
- o Information Modelling
- o Security
- o Real Time Distributed Data Access
- o Information Categorization
- o Closed Loop Management

Position paper (draft) for NMRG/EMANICS Meeting

Mark Burgess
Oslo University College

What are the challenges for the next 5 years in system and network management? I suggest three things worth discussing.

1) Decentralization

There is still a persistent belief in the need to centralize management solutions. We have to demonstrate with a usable technology, however imperfect, that decentralized management can be made to work with no greater uncertainties or unreliabilities than centralized models. This is key to solving the challenges of both scalability and commercialization.

2) Economics and Environment (with and without money)

The force driving the development of the network is commerce. Economics is such a fundamental part of civilization that the effect of communications technology on commerce must be understood. Today we tend to think in the traditional terms of payment for services. But informal trading of services has long been practised by service providers, esp. within BGP peering relationships. Money is only one of several currencies by which parties can trade online. The economics of service provision is also a part of this picture. There is much to be gained by studying energy efficiency and heat wastage in data centres and desktops also.

3) Convergent change management for general languages

We end with a technical problem. In configuration management, there is the immunity model for reliable autonomic change management. But this model works only for primitive strings of single symbols. To make autonomic, policy based management work as a self-regulating enterprise, we need also to be able to implement the same techniques for general coded languages of the Chomsky hierarchy. Alternatively, we need custom developed operating systems with properties unlike those that we have today.

Position Statement

=====

Michael Alexander <malexand@wu-wien.ac.at>

Network Management (NM) research has traditionally been focused on protocols, such as pertaining to configuration management and discovery. Yet, all areas of Network Management, ranging from Element Management (EM) to Network Management Systems (NMS) and Service Management Systems have in common that they critically rely on underlying data structures describing devices and services being managed. Network elements (NE) frequently spawn ten thousands of managed objects, such as for a multiservice switch.

In a common fallacy, programmatic access methods towards an NE such as operations on attributes via SNMP, CLI or XML-Netconf etc. are frequently taken as equating to 'network management'. Yet, in order for a Craft Terminal/EMS/NMS/Service Management System to provide sensible services to either a GUI operator or a mechanized interface, the device data of an NE has to be in a structured model. Managed objects alone - but in rare cases - do not allow to provision a subscriber, retrieve an alarm, examine a performance metric etc. From e.g. an NMS viewpoint, they tend to be 'flat' data structures which need to be put into a defined structure the NMS can understand so to perform any logical flow of actions on them. That is, a "flat" device attribute list does not enable an operator or mechanized interface to manage an NE. Unfortunately, there is little commonality in data data models, respectively managed objects exposed by NEs to their northbound interfaces. Even worse, the plethora of NMS/Service Managers each have their own models, leading to the present-day situation in which a large set of management systems with different models map to a large set of NEs with - differing models, both of which ever changing.

The following comparison highlights the problem from an implementation-effort perspective: It takes an experienced NM person about a year to understand a medium-complexity carrier device's data model in its entirety. This familiarity is necessary to build EMSs/NMSs/Service Management Systems with human-operatable GUIs, or machine-to-machine interfaces for high-level operations (enable a service/disable a service etc.). To adapt a given NMS's model for only one NM constituent, such as configuration management, the effort to initially map and then maintain the NE's to the NMS's data model is vast. For example, only to cover configuration management of a single service of a multi-service switch man-years of effort may be necessary to match the two models.

Yet in contrast, the effort to effect an attribute value change with SNMP, CLI or XML- Netconf etc. even with little prior knowledge of a device's functionality, model and behavior is comparatively miniscule. Setting a given NE attribute via a prompted CLI may take an NM-person with only rudimentary knowledge of a given device minutes; SNMP, depending on the design of the MIB takes longer, similarly for XML or CORBA etc. Comparing the discrepancy in the distribution of efforts again: the implementation of one leg of NM tends to be measured in man-years, the second leg of device access may range from man-hours to man-weeks. Logically, one might suppose that research would match this distribution.

NM research, however, is by a large margin focused on access-methods, discovery etc. while the more heterogeneous data models and data modeling per se have received comparatively little attention. As there are real possibilities to find underlying and universal commonalities in EMS/NMS/Service Management Systems on the one side

and NEs, protocols, transmission facilities, lines and services on the other, the two distributions should be closer aligned. Despite the complexity, it is possible to identify and define frameworks of each constituent and their relation to each other that would substantially ease the management of present and future devices, networks and services.

Dr. Alexander is a lecturer at the Department of Information Systems, Wirtschaftsuniversit?t Wien (WU). His professional experience includes education and product management at IBM, Siemens, Nortel Networks and Alcatel. Prior to WU, he was Product Line Manager for Alcatel ADSL (Craft/Network Management) and Optical Access Networks (Craft/EMS/NMS and Voice Communication). He is author of a text book on networks and network security as well several articles.

NMRG/EMANICS Workshop Position Statement

Olivier Festor <olivier.festor@loria.fr>
LORIA - INRIA Lorraine

Autonomous/mic Management

Automating as much as possible the management functions is not new in the device, network and service management community. In fact, automation is a goal behind almost every evolution in the control and management plane. A good illustration of this is, as presented in the DSOM'2003 panel on self-management by Alexander Clemm, the evolution of address configuration procedures which evolved from manual configuration, to fully distributed stateless configuration, through semi-automated allocation with DHCP. Many other such examples of evolutions exist in the Internet.

Continuing to improve automation in the management plane is and will remain a major challenge in the next decade. Open challenges in this area are: (1) ensuring traceability of the actions taken by the autonomous management plane, (2) ensure safety of the autonomous management plane (provide guarantees, even statistical ones, that the system brings/maintains the managed environment in a stable and better state after management actions have been taken).

Scale

The number of communicating devices, the networks that interconnect them and the services that run on top of them are increasing at a very fast rate often in constrained environments. This growth in size, often combined with increasing constraints on the underlying infrastructure (power, connectivity, CPU limitations) puts strong requirements on the management plane. While some performance evaluations were reported in a couple of contributions published in the last decade, better understanding, modelling and evaluation of scale in the management plane is still necessary.

Recent work on this topic is very promising (management trace analysis, behavior modelling, and benchmarking activities in the EMANICS NoE, large scale systems update analysis by Gkantsidis & al. in SIGCOMM 2006). Such efforts need to be continued and extended. Some open challenges in this area are : (1) the definition of common metrics and benchmarks like they exist in the database community, (2) the design and operation of large testbeds to enable reproducible experiments, (3) in depth study of large scale configuration as well as fault management activities.

Managing in Hostile Environments

Networks are getting more and more hostile to traditional management approaches. Entities like firewalls, NAT devices, problems like security flaws in some management approaches or the way they are used, dynamics of managed devices, changing connectivity conditions or simply the absence of support (agent) for a standard management approach make the life of traditional management very hard.

Management data is still available in the devices, networks and services, but the way it is accessed, processed and used is currently changing drastically. It is a fact that management functions get more and more embedded and take various new forms (trust, reputation schemes, incentives, ?) . This mutation of management functions is very promising. It also represents a real challenge for the next

decade. Open challenges in this area are: (1) design of new ?opportunistic? management models and algorithms for specific management functions on specific problems that use alternative ways to collect/process management data and act on the managed environment, (2) increase investigations on ?probabilistic management?, well adapted to networks which exhibit strong dynamic behaviours.

Promising Network Management Research Areas: A Network Vendor's Perspective

Abstract:

The Cisco Network Management Research Council (NMRC) is a task force within Cisco that is organized out of the CTO office of Cisco's Network Management Technology Group. The mission of the NMRC is to facilitate interaction between Cisco product development groups and researchers outside of Cisco who investigate areas of potential interest to those development groups. This has resulted in a number of joint research projects between Cisco and research organizations, in several cases funded by Cisco's University Research Program (URP) or as Cisco Applied Research and Development (CARD) projects.

As part of participating in the workshop, we would like to present an overview of research areas in network management that the NMRC has identified as particularly promising and relevant to Cisco's business. These are research areas which we believe have significant potential to address existing and/or future management needs of Cisco customers such as enterprise IT organizations and service providers. By presenting the overview, we hope to gather feedback from the research community on their perspective on those areas that we have identified, as well as contribute towards guiding the research community towards areas that promise to have substantial practical relevance and commercial impact.

=====

Petre's bio (if needed at all; if too long cut it):

Prof. Dr. Petre Dini [SM] (pdini@cisco.com) is now with Cisco Systems, as a senior technical leader and principal architect, being responsible for policy-based strategic architectures and protocols for network management, QoS, SLA, and performance, programmable networks and services, provisioning under QoS constraints, wireless networks and protocols, and consistent service manageability. He's applied industrial research interests include instrumentation software agents, performance, scalability, autonomic computing, wireless and mobile networks, constraints in wireless networks, adaptive networks, sensor networks, and policy-related issues in adaptable networks. Until 1990 he worked as a project director on the development of various industrial applications including CAD/CAM, nuclear plant monitoring, and real-time embedded software.

From 1991 he led various Canadian projects related to object-oriented management applications for distributed systems, and to broadband services in multimedia applications, until early 1996. In 1996 he joined Computer Science Research Institute of Montreal and coordinated many projects on distributed software and management architectures. In this period he was an Adjunct Professor with McGill University, Montreal, Canada, and a Canadian representative in the European projects. Since 1998 he was with AT&T Labs, as a senior technical manager, focusing on distributed QoS, SLA, and performance in content delivery services. He was the Co-Chair of Policy-Based Management Work Group in Telemangement Forum, is a Rapporteur for ITU-T/SG4, and actively involved in IEEE industrial initiatives. He has been an invited speaker to many international conferences, a tutorial lecturer, and chaired several international conferences. He published more than 100 papers in prestigious conferences and journals, and has more than 20 granted or pending patents.

Petre received his M.Eng. from Polytechnic Institute of Timisoara, Romania, in Computer Engineering, and a Ph.D. in Computer Science from University of Montreal, Canada. He is currently an Adjunct Professor at Concordia University, Montreal, Canada, a Senior IEEE member, and an ACM

member.

Joint Emanics / IRTF-NMRG Workshop on Future Direction of Network and Service Management Research

19-20 October 2006, Utrecht, The Netherlands

Being an Integrated Operator gives France Telecom with opportunities to offer new services in a seamless manner and thus to be an important player in an ubiquitous world. In contrast, we have to face challenging issues to guarantee interoperability between our various networks and services, as well as to interwork with other operators/providers, being able to ensure Quality of Service as well as trust to our customers.

Networks and Service Management is thus an important domain that must be mastered by the Network Operator/Service provider. Hereafter are some overall issues that we might wish to consider in the joint Meeting about Future Directions of NSM:

1. The various areas that are concerned (home networks, access networks, Backbone networks and the many service platforms) that we are managing include a number of various technologies that have a huge impact on our OPEX. One major business objective is to take advantage of this many resources to use them as best as can and to be develop agile management.
2. The heterogeneity of our networks and technologies is an opportunity to deliver services through different ways but requires additional features that may increase CAPEX. Therefore a trade-off must be found between flexibility between the various technologies available and the cost of implementation.
3. Scalability of procedures and mechanisms become of utmost importance as the number of terminals, devices and network elements increases.
4. Complexity might highly be increased and limit the role of human operators for management tasks usually carried out by NMgers, thus calling for increased automation.
5. The need for automation signifies that a kind of autonomic and self-management tools and mechanisms must be developed in order to make the systems more intelligent regarding the constraints and the environments. Learning and decision-making architecture have to be developed in order to account for change of context and Applications requirements on QoS.
6. Since networks and services have to interwork via different service providers, it is requested that architecture, models and procedures, protocols be standardized through ad hoc SDOs.
7. Finally, QoS and Trust/Immunity/Security must be considered as key objectives and methods and related protocols need to be developed accordingly.

Prosper Chemouil
France Telecom R&D

New Approaches to Distributed Management

Rolf Stadler <stadler@s3.kth.se>

The need for network management in large-scale and dynamic network environments calls for solutions beyond traditional management paradigms. Approaches based on management overlays, p2p computing, as well as distributed aggregation and control schemes, have been recently proposed to engineer management systems that scale beyond 1000s of nodes and are robust regarding topology changes and failures. The design space is huge and still largely unexplored, and the challenges are tremendous.

EMANICS Workshop, 19-20 October 2006, Utrecht
Position Statement for Simon Leinen <simon@limmat.switch.ch>

Personal Background

I graduated in Computer Science at TU Berlin, then worked for several years as a system administrator at a university, and did some consulting and development jobs both in the research and commercial world. For the past ten years, I have been working for SWITCH as a network engineer.

SWITCH

We run a national backbone for the Swiss higher education and research communities. Our network group is small and performs design, engineering, and operational duties. So I help decide how we build the network, but also have to help fix things when it breaks.

Current Operational/Network Management Issues at SWITCH

In former generations of our backbone, we rented "services" such as PDH or SDH circuits, ATM and MPLS connections, from carriers, and built an IP network over them. Today however, our network runs on leased dark fibers, and we are operationally responsible for the optical transmission system in addition to the IP network. This makes things both simpler and more complex: More layers must be monitored, but it's easier to get an integrated picture now.

Our main operational challenges don't concern device management technologies, but rather support for communication in the broadest sense: internal coordination and documentation, customer information such as tickets, and alerting and information exchange with providers of network or outsourced monitoring services. We would welcome tools that could use network management information to facilitate such communication tasks.

Cross/Inter-Domain (Performance) Monitoring

Multi-domain monitoring systems such as Route-Views, Traceroute or Looking Glass gateways, or RIPE's RIS and TTM services, are extremely useful in the Internet model of service provision. They often have deployment and/or usability issues because of the tension between the wish for information hiding in a competitive marketplace and the need to combine information from different domains to help customers.

We are involved in research and development activities of the GN2 EU project. GN2 includes an important activity in multi-domain performance monitoring (JRA1), where a Web Services-based system ("PerfSONAR") is being developed with transatlantic cooperation.

Parts of the system have already been deployed, and allow unified access to operational information (topology, link status, load etc.) over many independent (research) networks. Finding good ways to use such information to support network operations looks like an interesting research topic.

NETCONF

In the IETF, I have co-chaired the NETCONF working group, where a basic protocol was defined for the exchange of (parts of) device configuration between network managers and devices.

Ongoing and proposed future work for the Working Group tends towards generalizing the NETCONF approach to cover management tasks not strictly related to configuration. Some of these tasks, e.g. polling

and notifications, are already covered by existing protocols. I'm trying to understand the reasons why these existing tools are rejected by (slightly bell-shaped) parts of the community.

NMRG/EMANICS Workshop Position Statement

Stefan Wallin <stefan.wallin@dataductus.se>
Data Ductus

1 Service Quality

1.1 Who we are

We work as network management integrators for Telco Operators world-wide. The HP OpenView products are the main solution components. We have been working with classical alarm management and lately service management and SLA management. We do part-time research work at CDT Luleå University.

1.2 Introduction

We will look at the concept of service quality within the telecom service provider context.

The subject of service quality is very complex and broad in its nature. The two terms service and quality are inherently lacking solid definitions. Adding the two into a new concept does not make the picture more clear. We will use the two following definitions:

- * Service: Something that the service consumer pays a service provider for.
- * Quality: A measure correlated to customer satisfaction with the service. An underlying complexity is the different contexts of a service; one as seen by the service provider and on the other hand as seen by the service user:
- * Delivered Service Quality: often an objective measurement of the service quality using technical performance indicators. The time window is often short if not even momentarily
- * Perceived Service Quality: a subjective judgement of the service quality. Covers a broad aspect of service quality indicators over a long period of time; Installation of the service, contacts with helpdesk, etc.

1.3 Different ways of measuring service quality

There are several different approaches to measure service quality, each of them with focus on one, or both, of the above contexts.

- * Traffic shaping: this is a technology which tries to manage the network service in order to actually deliver the defined QoS. This is mainly the efforts of the IETF work on QoS. Different techniques are applied to manage throughput, packet-loss, latency and jitter. Well-known IETF work is Diff-Serv and MPLS.
- * Probes: simulates end-users behaviour using the network as a black-box in order to measure the actual delivered network service quality. Cisco Service Assurance Agent, SAA, falls into this category. Mobile networks have probes for probing telecom network services like Voice, MMS, etc..
- * Modelling, calculation: a system where a formal model ties different measurements into a service model. The service quality is actually calculated by a calculation engine. These are fairly complex tools where a formal language is used to define the service model. The model captures the overall structure of the services, and mapping in several layers down to individual network resources. It calculates so called Key Performance/Quality Indicators, KPIs/PQIs, using

expressions based on input events and polls. Service state is propagated in the service model which is typically tree-structured. Since there is a model the system and user can perform 'reasoning', a certain service state can be reduced to the original events/polls. What-if scenarios can also be applied, showing the affected services as effects of network events and changes.

- * User feed-back systems: informal methods like web forms, telephone interviews etc in order to catch users perception of the service

Apart from covering different contexts they also have different capabilities:

- * Monitor: only reporting on real-time and historical QoS
- * Control: can actually control network resources in order to deliver a defined QoS
- * Root-cause: can drill down from a service failure to the root-cause of the failure

1.4 Organisational aspects

It is impossible to talk about service quality without looking into the involved parties.

- * Service user: context of perceived service quality.
- * Network provider: network management center, network operations centre, etc: context of delivered service quality. The NMC needs to have an understanding from network events/state into affected services and customers in order to prioritize work.
- * Customer care: needs to map the two contexts to each other:
 - o Customers are complaining, why?
 - o Which services and customers will be affected by a network failure?

1.5 Pros and cons of different techniques

Traffic shaping

- + Actually controls the network to deliver service quality
- + Embedded in some protocols
- Only available for some network services

Probes

- + Easy to deploy
- + Measures the service usage
- No understanding of why?, root-cause etc.

Modelling, calculation

- + Reasoning: service-impact, root-cause, what-if
- + Creates a model of the network
- Complex to deploy
- QoS defined by calculations, they may be wrong

User feed-back

- + Reports on true user perception
- + Easy to capture
- No mapping to why?
- Subjective

1.6 What is service quality, really?

Although tools and standards for QoS have been available for a while. We are not sure that users of services perceives a big

improvement. The number of available services have increased but the service quality have rather decreased. What is the reason? We think that one of the problems is that the QoS concepts have been focusing on the provider side and the technical/characteristics side of quality. We need to add the perceived quality as defined by subjective measurements. Furthermore, delivered and perceived service quality needs to be integrated.

```
User
|
User Terminal
|
Service Access Point
|
Media
|
Service Delivery Point
|
Provider infrastructure
|
Service producer
```

The above ASCII art illustrates the chain involved in the total service usage chain. Our proposed overall architecture is following:

- * Service model which models the total chain above
 - o BUT, the service model needs to be dynamic and easy to change
 - o The model needs to support pivoting in order to present different views
- * Collection of quality parameters from all components in the chain, both technical provider and user subjective indicators.

Position statement for workshop on Future Direction of Network and Service Management Research

Thomas Delaet <thomas@delaet.org>

The scale of IT infrastructures is increasing rapidly, as is the heterogeneity of devices and their software. Since businesses depend more and more on their IT infrastructures, system administrators have to guarantee functionality, security, availability, performance, etc. Although many separate system administration tasks can be automated, manual editing still prevails in managing the interactions between different subsystems. In a large and heterogeneous IT infrastructure, this approach is very error-prone. Moreover, manually editing configuration files does not guarantee correct behaviour of functional and non-functional characteristics. My research is positioned in the area of policy-based management and entails creating an environment that manages the growing complexity of IT infrastructures. I call my approach to system administration PoDIM, short for Policy-Driven Infrastructure Management.

My environment consists of a language and a platform implementation.

- * The language is composed of a core language and can be extended with multiple subsystem languages. A subsystem is defined as a functional building-block of a system. The configuration of devices can be considered the sum of partition subsystems.
- * The platform is a distributed platform which allows to delegate system administration responsibilities to third parties and enforce those delegations. My platform interacts with different subsystem components to generate the configurations for all individual systems in an infrastructure.

Subsystem can be categorized in low-level and high-level subsystems. Examples of low-level subsystems are the scheduler and network configuration. An e-commerce application running on a webserver farm or an ERP-client application running on most desktop computers are examples of high-level subsystems.

The most recent description of PoDIM can be found in a paper I submitted to the IM 2007 conference (currently under review) at http://purl.org/podim/papers/im_2007.pdf. More information about the PoDIM prototype can be found on <http://purl.org/podim/devel>.

I think my contributions to the workshop discussions can be twofold.

- * My research entails an extensive (still incomplete) study of existing environments and research approaches for infrastructure management. This allows me to compare different approaches with each other and discuss them with other participants.
- * One of the core subjects of my research is to improve understanding of the complex web of interactions and dependencies between different subsystems in an infrastructure. It would be interesting to discuss these interactions with interested people.

In short, I think I have the required research background on current approaches to system administration and this workshop gives me a unique opportunity to discuss my (and other) approaches to system management with other stakeholders.