Integrated Network Management

– The Key to the Efficient Management of Complex Operational Networks

K-J. Schulz & U. Christ
ESA Directorate for Technical and Operational Support,
European Space Operations Centre (ESOC), Darmstadt, Germany

L. Lechien
Aethis, Louvain-la-Neuve, Belgium

Together with industry, ESA has conducted a study of state-of-the-art network management concepts in the context of assessing its future space operations networking requirements. The validity of the principles involved and the potential for improved efficiency have been demonstrated with the help of ESA's Interconnection Ground Subnetwork (IGS) Testbed environment.

Introduction

Management is defined as the act, manner, or practice of managing: handling, supervision, or control. Recognising the importance of appropriate management whenever organisational complexity needs optimisation for efficiency or disaster prevention in an organisation or industry, management principles, techniques and strategies have been continuously developed and adapted to changing boundary conditions. The management of complex technical projects has undergone similarly expeditious development in recent years. Multimedia networks with continuously changing profiles of users and traffic flows are good examples of the complex architectures and very demanding conditions facing today's network operators. By tailoring the network management system to the particular operational system scenarios involved and careful selection of the technologies used, the stresses associated with operating such networks under extreme time pressure can be considerably reduced.

Motivation

The operational networks required to support the communications needs of space missions are a major cost element in the operation of satellite ground segments. Cost optimisation for each individual project has tended to lead to the use of different basic communication services, highly customised to the specific needs of each mission. Typical communication bearer services today are terrestrial and satellite links, Frame-Relay-based Virtual Private Networks (VPNs), Narrowband Integrated Services Digital Network (N-ISDN) on-demand services, permanent and on-demand services via satellite with Very Small Aperture Terminals (VSATs), and Broadband Integrated Services Digital Networks (B-ISDNs) based on the Asynchronous Transfer Mode (ATM). Whilst on the one hand the cost-optimised mapping of the mission requirements to the most appropriate communications service reduces service charges, on the other increased effort is needed to operate these communication systems due to the added heterogeneity. To counterbalance this increased operational involvement, sophisticated Integrated Network Management Systems are needed that will ease the routine operations and simplify troubleshooting in the event of problems.

In addition, the complexity involved in running operational mission-support networks has increased considerably in recent years because new communication techniques supporting multimedia (voice, video, data) applications make the attendant networking support quite demanding.

ESA presently runs two operational networks: OPSNET for its unmanned missions, and the Interconnection Ground Subnetwork (IGS) Testbed as a precursor to the operational IGS for the International Space Station. OPSNET typically uses leased lines enhanced by N-ISDN and Frame Relay VPN services, while the IGS Testbed uses leased lines, N-ISDN, VSAT and ATM services.

Major cost reductions have been achieved for both networks through the introduction of on-demand connectivity, which is particularly cost-efficient in providing backup capacity or additional short-term connectivity when user needs dictate. However, the reliable and controlled establishment and provision of such on-demand connectivity in the operational network environment creates an additional workload for the operators and adds considerable complexity to the routine operations, which needs to be contained in order to ensure safe long-term operation. In addition, the new multimedia applications have added totally new equipment, e.g. for video, to the operational network, which also needs to be managed and integrated into a network management scheme to ease both its routine operation and fault recovery.

In summary, therefore, the safe and reliable operation of these very dynamic and complex networks and services, requiring even greater operator involvement, calls for optimum network management support tools if the manpower required is not to increase. The goal should rather be to reduce the number of network operators required and to allow a reduction in the skill profile for 24hour/7day-shift positions.

Key features for an Integrated Network Management System

Experience in operating large space-mission support networks like OPSNET and the IGS Testbed over long periods indicates that network management does not need to be fully integrated into one system, as was the paradigm for such systems in the past. The major elements of a network, such as the Wide Area Network (WAN) switches, the routers, the voice-conferencing equipment, the video distribution equipment, etc. are equipped with dedicated, often proprietary, management systems fully optimised for the configuration and operational needs of the particular equipment. Attempts at achieving an overall integrated network management system in the past have led to tremendous software maintenance problems when new releases were deployed for the individual network elements, since this also required updates in the Integrated Network Management software. In practice, the management systems were rarely ever complete and up to date and were therefore treated with suspicion by the operations personnel.

It was therefore necessary to depart from this old paradigm and apply an integrated network management approach more suited to today's operational environment. This leads to the concept depicted in Figure 1, still keeping the dedicated Element Management Systems that are most suitable for fulfilling their particular job of configuring and managing the devices that they control. In particular, they provide very detailed Resource Views of the devices they manage. The Integrated Management System realises higher level operator and user-oriented views that resemble a certain overall service model and provide an intuitive graphical display of end-to-end service. These higher level views can only be generated by a combination of information from the different Element Management Systems. The Integrated Management System becomes more of an umbrella system, not replacing the individual Element Management Systems, but rather creating an umbrella for the entire system with added functionality.

The key features of such an umbrella Integrated Management System can be summarised as follows:

Figure 1. New Integrated Management System paradigm

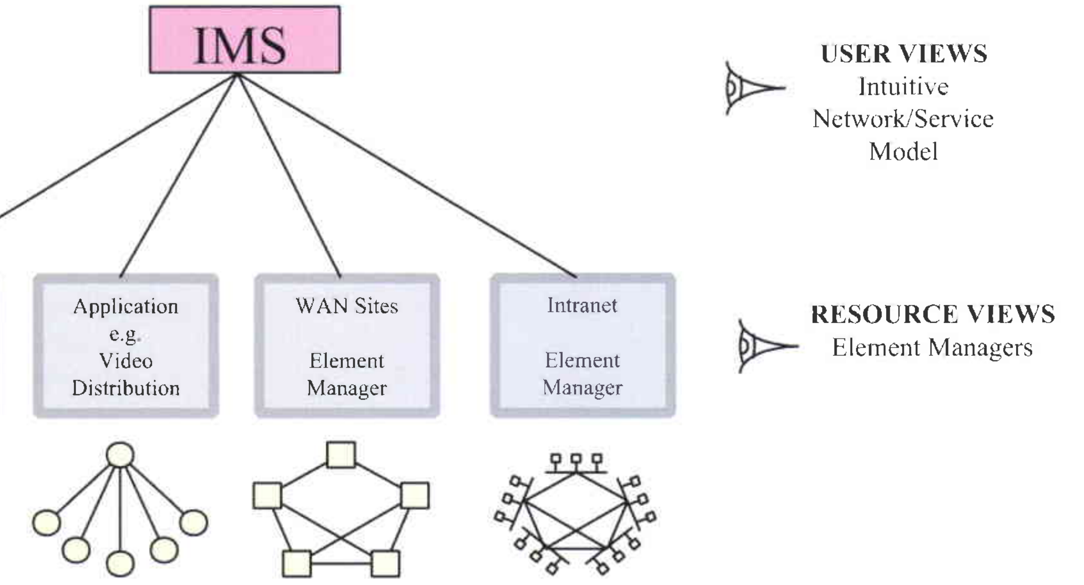

1. Online Unified View of the entire network resources via the horizontal integration of high-level management information from Element Management Systems. Generation of useful synoptics for the network operator, with the ability to export views to actual end users.
2. Offline Network History Summary, which complements the online view with long-term availability and performance information, allowing for the independent verification of Service Level Agreements (SLAs) with telecommunications operators. Such long-term reports provide the basis for future planning and allow easy tracking of failure correlations and performance evaluation.
3. Service Management Views, which present the end-to-end application view as perceived by the end user, thereby aggregating the management information of the various elements that provide the service.
4. Automated Connection Management, which ensures reliable and efficient activation and release of on-demand communication resources according to high-level communication service schedules coming from mission operations. This is also useful for automatic initiation of preplanned recovery procedures if, for instance, performance degradations or failures are detected.

A state-of-the-art Integrated Management System has to fulfil the following architectural and software requirements, for which the rationale is provided in the next section:

1. Simple Network Management Protocol (SNMP) for the transfer of information between the Integrated Management System and the Element Management Systems based on high-level information. In addition, access to low-level device-embedded SNMP agents and a migration possibility to also support legacy devices that do not yet support SNMP.
2. Commercial-off-the-Shelf Network Management Platform that provides all basic functionality for SNMP data exchange and easy generation of animated customised graphical views.
3. World Wide Web (WWW) based remote access for the export of user views.

Taking into account the above requirements for a new Integrated Management System, a prototype system was built and demonstrated in a reference network that combined on-demand ISDN and ATM connectivities for the provision of video and data services.

Design of the prototype Integrated Management System (IMS)

Past concepts for Integrated Network Management have often proved inadequate in many practical cases. Omnipotent monolithic systems, generic enough to cope both with detailed device management and the more general functions at network or even service level, required heavy development efforts for producing what often turned out to be rigid and never complete solutions.

A more pragmatic strategy is to seek simple individual solutions that can be loosely combined to form a prototype system and derive systematic techniques applicable to real operational environments. This implies the very careful selection of some key system and software requirements:

1. Standardised use of the Simple Network Management Protocol (SNMP)

SNMP is the standard network management protocol of the TCP/IP (Transmission Control Protocol/Internet Protocol) family. It is used for the exchange of management information between two conceptual entities: a manager (typically a managing application) and an agent (e.g. the management interface of a network device, or an element management system). It provides three basic communication interactions:

- 'Get-Requests', which allow a manager to acquire the current value of selected variables of a Management Information Base (MIB) instrumented by the target agent, i.e. monitoring.
- 'Set-Requests', which provide new values to be assigned by the agent to the given MIB variables, i.e. control.
- 'Traps', which are asynchronous notifications delivered by an agent to a manager upon detection of a particular condition (e.g. an alarm event).

Given the tremendous success of SNMP, this protocol has been selected for the exchange of management information between the IMS and, preferably, high-level interfaces provided by element managers or, by default, the lower level MIBs supplied by agents directly embedded in the equipment. A flexible solution, based on the LT-301, which is a small adapter from SNMP Research Inc., also allows one to add SNMP support to legacy devices that can only be controlled by messages and commands from their console port.

2. State-of-the-art off-the-shelf Network Management Platform

A network management platform provides a set of basic building blocks, e.g. for event handling or graphical network map animation, which can be customised for specific network management applications. After a comparative analysis, the HP OpenView Network Node Manager (HP-OV NNM) product has been selected as the core platform for the IMS. This platform offers a high stability for the custom-built applications and ensures compatibility with future trends in network-management technology. It is presently the market leader. It is also rather flexible and open, and thus minimises the development effort for those custom applications designed to take advantage of the platform's built-in features.

3. Use of Web-based technology for the export of high-level views

The IMS must provide a synthetic and intuitive model of the communications network. This high-level presentation is not only of major interest for the network operators, but can also be very useful for the end users. The easiest way to ensure wide accessibility is to make use of the World Wide Web, the HyperText Markup Language (HTML) and JAVA technology. To avoid unnecessary work, however, it is better to rely on the standard export mechanism inherent in the network management platform, which will benefit from future enhancements.

The resulting architecture of the IMS prototype is shown in Figure 2.

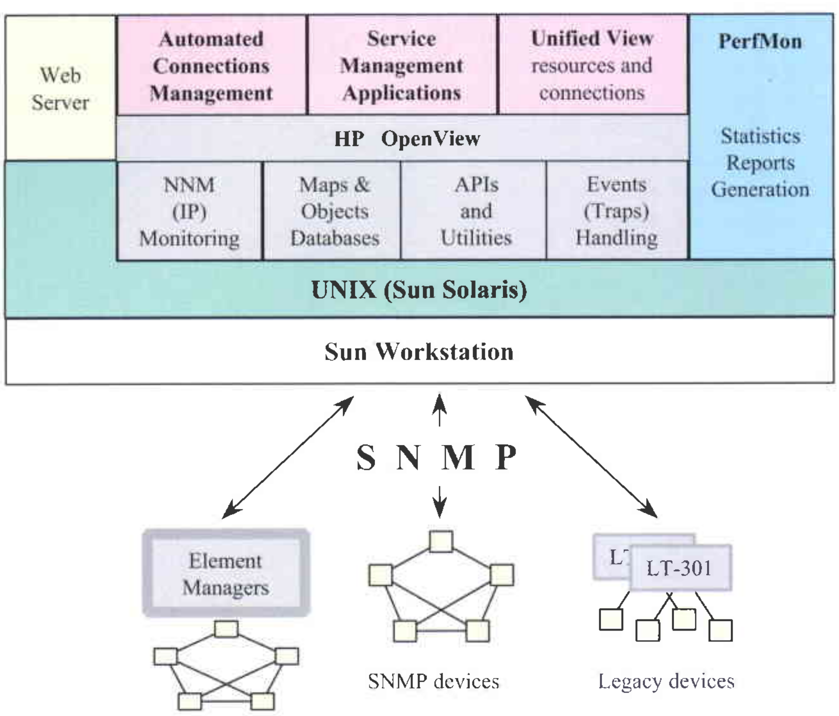

Figure 2. IMS prototype software architecture

The standard IP node management application of HP OpenView NNM ensures the detection and monitoring of the IP data network components (called the Intranet). For the majority of non-IP-related managed objects, like physical resources or logical communication services (e.g. equipment ports, protocol entities, virtual channels), a custom model is created in the OpenView objects database. Here the managed objects are represented in the graphical map as icon and connection symbols according to a layout derived from custom network operations conventions. As dynamic connections play an important role in the IMS logical views, specific techniques have been used for their graphical representation.

Status monitoring of those objects relies mainly on simple 'trap-triggered polling' mechanisms. Custom scripts are attached as automatic actions to relevant types of SNMP traps. Upon event (trap) reception, the script retrieves additional diagnostic information by querying the originating SNMP agent (polling), to identify the corresponding OpenView objects and to choose the status colour to be used for their representation in the unified view.

Some state transitions are, however, not reflected by the spontaneous generation of traps. This is typically the case for dynamic activation and release of virtual channels. For those cases, simple dedicated polling programs continuously inspect appropriate state variables from the device MIBs, thereby deriving the animation of the colour of corresponding graphical objects in the views. The OpenView Web export capability allows remote access to the views.

With the Performance Monitoring (PerfMon) tool, raw historical observations are continuously collected and stored in log files that are regularly processed to automatically update an integrated set of hyperlinked Web reports. These represent the availability and performance summaries through a mix of tabular and chart data on different time scales (day, month, 12-months).

Dedicated Service Management Applications receive service related configuration and status data by traps. They dynamically draw their own views in the OpenView graphical map. A simple menu-driven control interface permits the on-demand execution of network device configuration commands by the invocation of custom scripts of SNMP Set-Requests.

An editor (replacing the mission-planning interface) of high-level service requests generates a communications schedule that is translated into a series of UNIX at-jobs by an Automated Connections Management robot. At the specified time, those scripts are executed, and the necessary atomic commands to the managed network components are issued through SNMP Set-Requests.

Customised IGS Testbed Intranet view

The basic IP Network Node Manager

application of the HP OpenView product has been tailored to monitor the IGS Testbed Intranet. Its topology is depicted in Figure 3. As this network is still under construction, not all sites are yet represented.

It is a star-shaped network with a central node at ESOC. The Johnson Space Centre (JSC) and Marshall Space Flight Centre (MSFC) are connected through different paths (IP tunnelling over ESA/NASA IP links, Frame Relay circuits over leased lines and backup links over inverse multiplexed ISDN). An experimental connectivity with MSFC over public ATM is also foreseen.

User Support Operations Centres (USOCs) and ESA Control Centres will gradually be added, as indicated in the right part of Figure 3. At each site, a LAN segment dedicated to network management traffic (NM LAN) is separated from the local Operational LAN (OPS LAN).

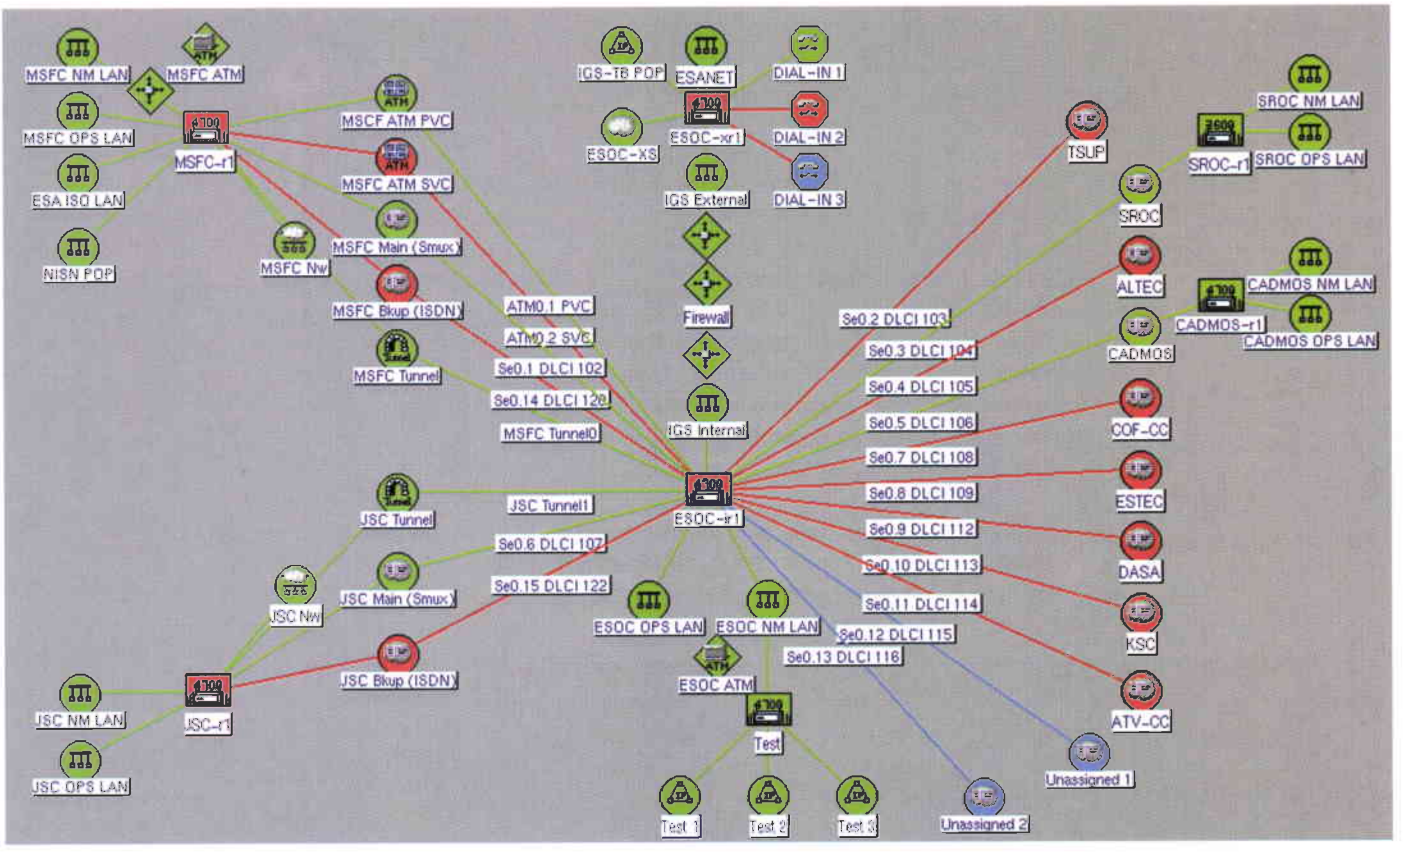

Figure 3. IGS-Testbed Intranet view

Reference network for the IMS prototype

The reference network for the development of the IMS prototype involves public ISDN-based connectivity, as already used in actual mission-support scenarios, and .the use of an experimental ATM infrastructure. It offers combined data and video-conferencing services. The demonstration in the context of the IGS Testbed has been conducted in two steps. First, a unified model of a representative heterogeneous network configuration has been integrated as an overall synoptic of the status of resources and connections. Then a dedicated Video Service Management facility has been built-in for the easy setting up and control of video conferences. In addition, an Automated Connection Management Application has been developed for orchestrating dynamic service provisioning according to high-level communication service requests coming from mission planning.

Figure 4 shows the reference network configuration used for the final demonstration of the combined exploitation of the different IMS custom applications in a scenario of automatic setting up of end-to-end video conferences over ISDN and ATM.

A video Multipoint Control Unit (MCU) allows one to configure two-party and multi-party conferences between video channels. Five of the eight available user ports were used in the demonstration. One remote video codec was connected through inverse multiplexed dynamic ISDN calls (6 individual ISDN B-channels of 64 kbps each need to be aggregated to obtain the 384 kbps required for the video transport). Another video codec could be connected either through ISDN calls or through dynamically enabled ATM Permanent Virtual Circuits (PVCs). A complete scenario involving the combined use of the Automated Connections Management Application, of the specific Video Service Management Application, and of the Unified

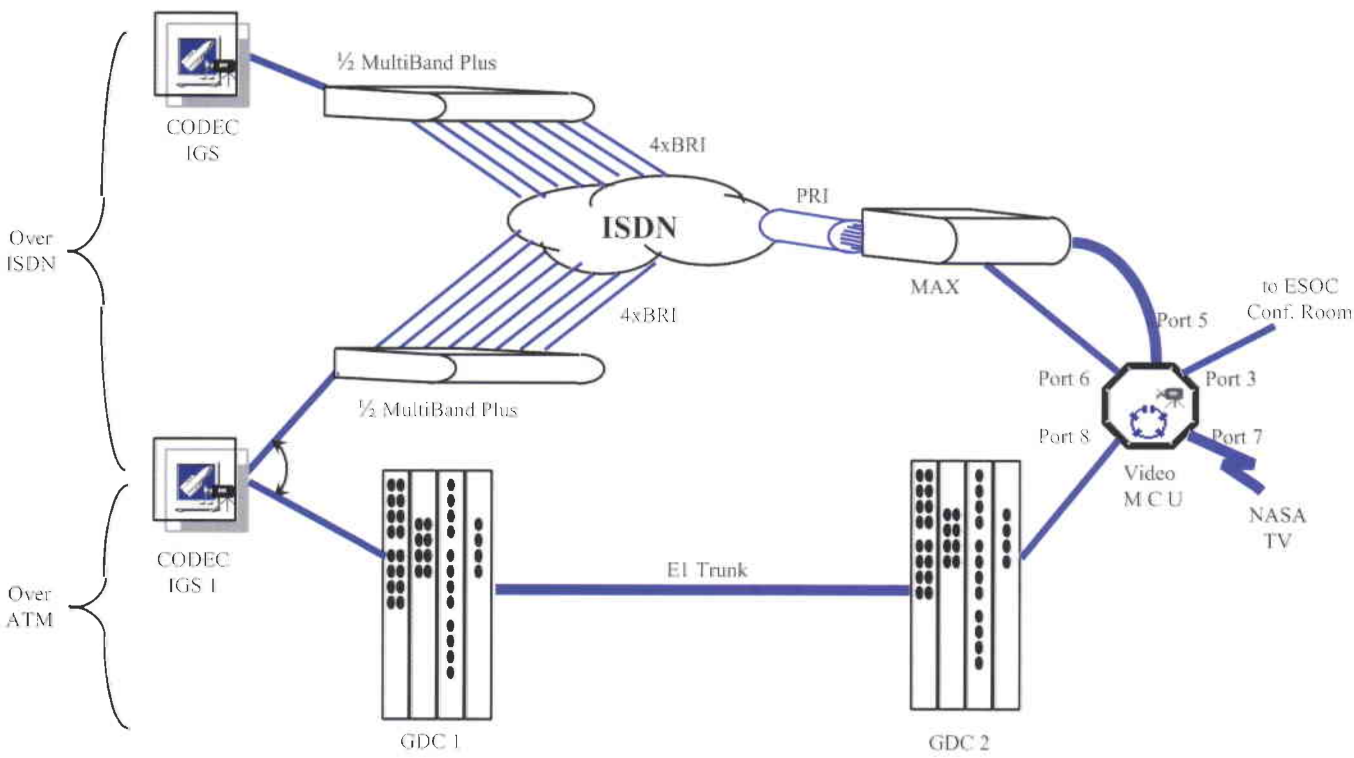

Figure 4. Reference network for the IMS prototype

View has been demonstrated with this reference network infrastructure.

The real-time unified view

Although the unified view is manually custom designed, it has been possible to model the network objects in a systematic fashion by combining the network element MIBs information and by adapting IGS Testbed-specific network operation conventions.

Particular emphasis was put on the intuitive visualisation of communication relationships. For instance, coloured graphical links synthesise the state of all active components in a path, as well as the state of the logical end-to-end service channels. Different techniques have been used for an animated display of dynamic connections, with minimal effort. The methods rely on an a priori knowledge of the conventional configuration of the communication resources.

The example in Figure 5 illustrates the technique used for inverse multiplexed trunks over the public ISDN, the left-hand OpenView window being the synoptic of the reference network shown in Figure 4.

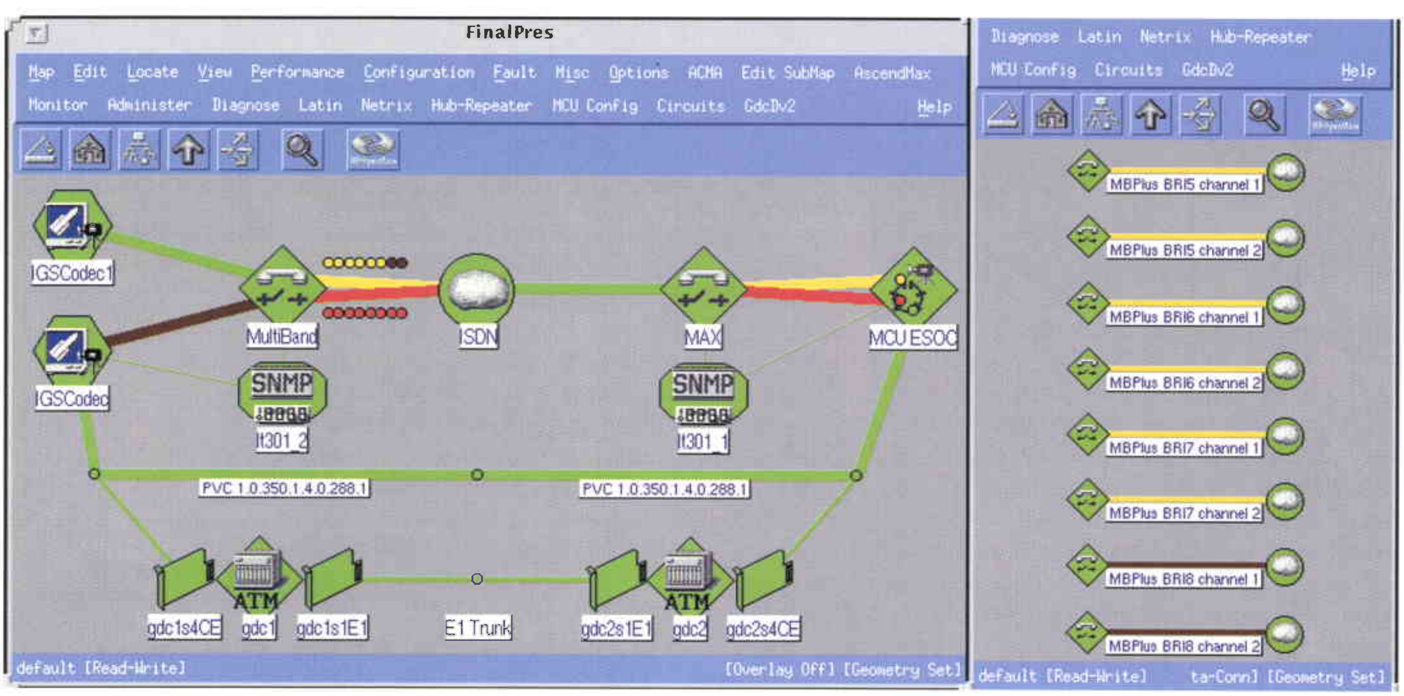

Figure 5. Real-time unified view of the reference network

Table 1. Cost-efficient dynamic ISDN trunk bandwidth allocations for multiplexed services

Combined services	**Number of 64 kbps ISDN B-channels supporting the multiplexed trunk**
Voice + low data rate	2xB = 128 kbps
Voice + low data rate + high data rate	4xB = 256 kbps
Video	7xB = 448 kbps
Video + voice + low data rate	8xB = 512 kbps
Video + voice + low data rate + high data rate	10xB = 640 kbps

Cost-efficient multiplexing of on-demand multimedia services over long-distance ISDN trunks involves the dynamic aggregation of a minimum number of 64 kbps ISDN B-channels to achieve the bandwidth needed at a given moment. Typical assignments are presented at Table 1.

In our example, a trunk could amount to a maximum of 4 dedicated Basic Rate Interfaces (BRIs), where each BRI may support up to two B-channels. The direct (nonmultiplexed) transmission of a video signal necessitates 6 B-channels. In the upper view, the trunk is represented by a single connection symbol (e.g. the upper link between the MultiBand icon and the ISDN network icon). On the right-hand side of Figure 5, a zoom view displays the possible 8 B-channels of the 4 BRIs, of which 6 are in the process of being established.

When a call is in progress, the trap-triggered polling mechanism detects how many channels are already dialled, and assigns them a yellow colour. Once the call is in service, the actual number of channels involved is coloured in green. Upon call release, all channels will be shown in red. The status colours of the compound trunks in the upper-layer view are painted accordingly. A variant of this techniques shows every channel as a bullet near the trunk representation in the upper submap; the visual effect is that of an animated sliding bar of LEDs.

Long-term statistics

The PerfMon tool, initially developed by Aethis for ESA to monitor routers and Local Area Networks (LANs) of the administrative network in the frame of the CNMS project, has been enhanced and has been integrated with the IGS Testbed. As a spin-off, it was also installed in the recently established OPSNET Intranet.

This tool consists of a series of dedicated poller programs that efficiently collect samples of status and performance indicators via SNMP Get-Requests. These observations are regularly processed by flexible scripts, written in PERL, to update summary reports showing statistics for the short, medium and long term. The resulting reports present multiple facets of the network history: availability, traffic utilisation, Quality of Service (QoS) indicators, and delay measurements for different protocols supported by the routers.

As the set of reports generated is presented on a web site, presenting hyperlinked intuitive views at different levels of granularity, operators and users can obtain a complete picture of past network behaviour by navigating through the web site, and thereby avoiding printing large piles of paper.

The sample web report in Figure 6 shows average and peak traffic over an X.25 link between ESOC in Darmstadt (D) and Redu in Belgium (B).

Figure 6. Long-term statistics

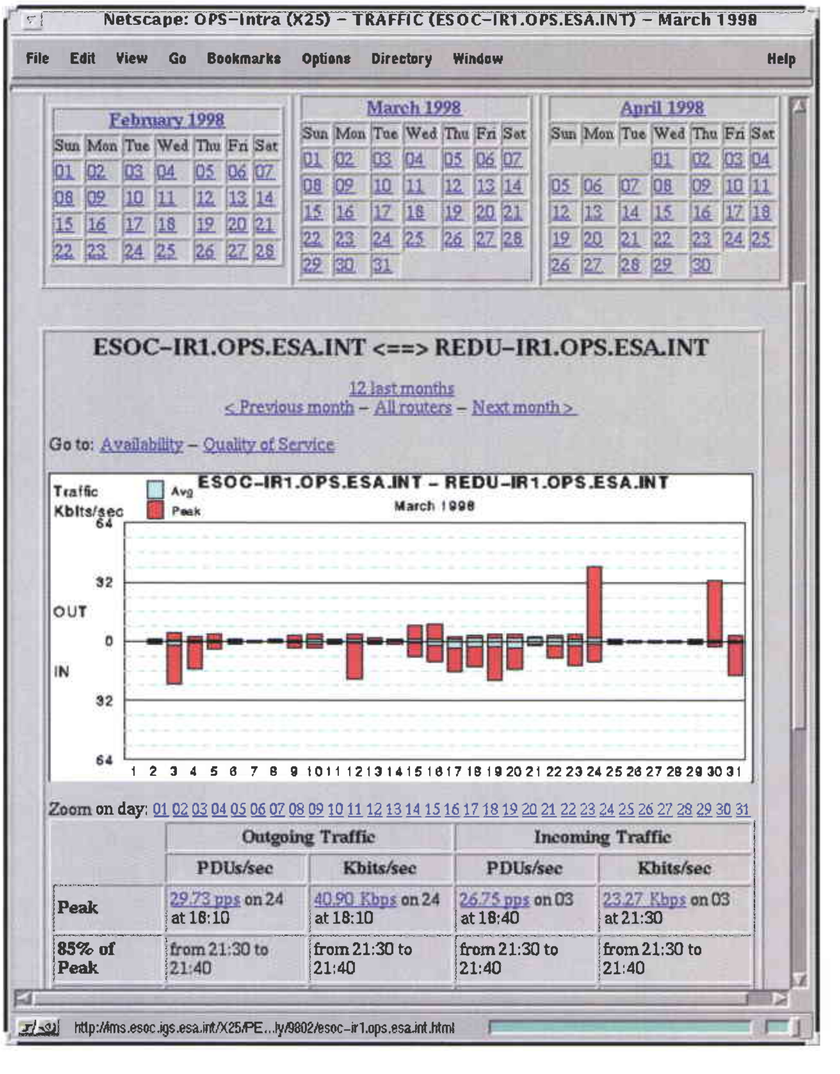

Video service management

The video Multipoint Control Unit (MCU) is a console-controlled device, the management interface of which has been mapped to be 'SNMP-manageable' by using a LATIN (Legacy Adapter To INternet) agent. A specific Video Service Management application has been developed for the IMS to make the MCU's management more intuitive. This application generates its own OpenView submap according to layout directives in a configuration file. Typically, the icons of the 8 MCU ports are placed in a circle (Fig. 7).

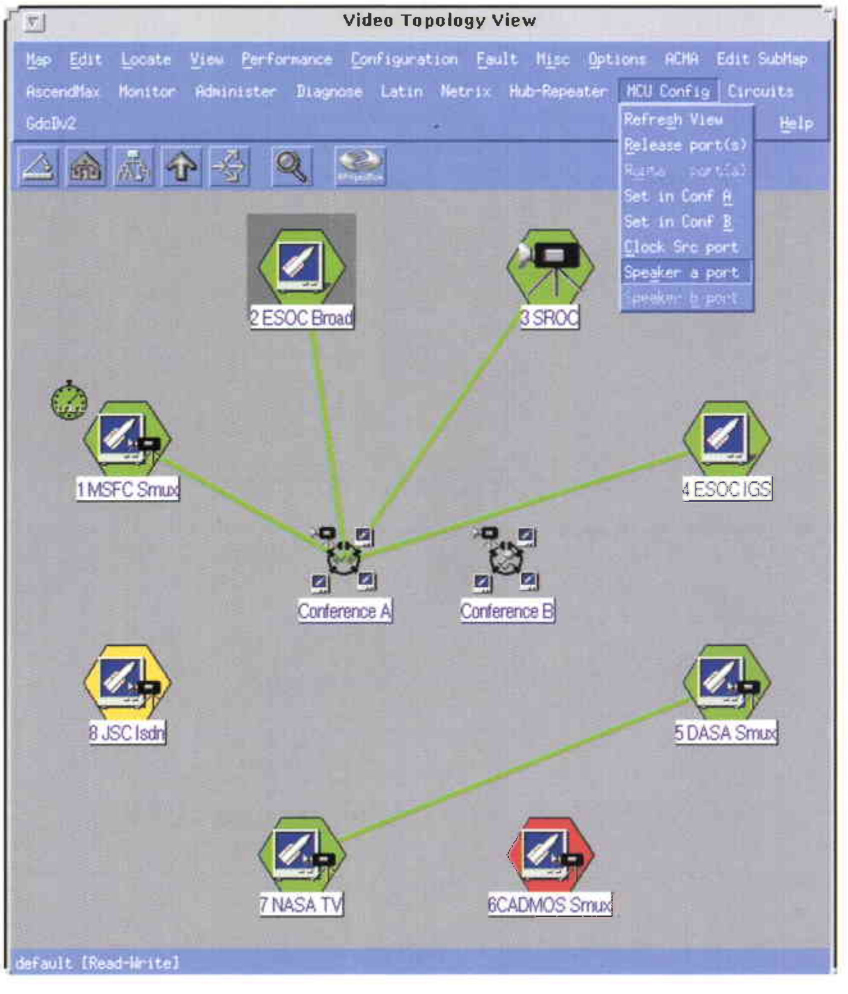

Figure 7. Video Service Management view

The device can establish two-party conferences in which both connected codec devices receive signals from each other. It also supports up to two multi-party conferences ('A' and 'B'), represented by the two central icons.

In a multi-party conference, the 'Current Speaker' signal is sent to all others, while the signal of the 'Previous Speaker' continues to be sent to the now Current Speaker. In order to distinguish the three roles, a simple 'monitor' icon is used for passive participants, while a 'camera' icon identifies the Current Speaker, and the Previous Speaker icon combines a monitor with a small camera. The icon of each conference is also dynamically selected to identify its operating mode. When it is 'Voice Switched', the louder sound source is automatically assigned as Current Speaker, while in 'Local Control' mode this selection is manually controlled.

In Figure 7, a multi-party conference 'A' has been established between the ports 1, 2, 3 and 4. Port 3 is the Current Speaker and port 1 is the Previous Speaker, whilst ports 2 and 4 are passive participants. A two-party conference is also established between ports 5 and 7. Port 6 is experiencing a bad signal, e.g. communication with the corresponding remote codec device is not established. The signal at port 8 is not synchronised. The signal at port 1 is used as the reference clock for the MCU device.

In addition, a context-sensitive menu allows simple control of the MCU behaviour. It is possible to refresh the view (to get the initial configuration before the occurrence of the first poll). Selected port(s) can be released from their two- or multi-party conference. Two selected ports can be connected in a two-party conference. Selected port(s) can be assigned to conference A or B. A single member port can be configured as Current Speaker of its conference A or B. A single selected port can be used as Clock Source.

Automated Connection Management

The scope of a prototype Automated Connection Management Application (ACMA) has been defined following an analysis of three types of information:

- a prototype database of high-level Communications Service Requests has already been built by the Mission Planning Team
- all communications services anticipated for the real IGS have been identified
- a survey of the logical communication resources of the IGS Testbed controllable by management commands has been performed.

The principal hypothesis in mission-related operational environments is that the detailed characteristics of possible connections are pre-configured in the network elements. The definition of those service instances associated with a description of their control is specified in a 'Network Resources Description' file. A formal high-level definition of the Communication Service Schedule, normally derived from mission planning, forms the second input for the Automated Connections Management Application (ACMA). In the prototype implementation this is generated with a local graphical editor.

The ACMA robot identifies all services that have to be activated or released at each transition of the communication service schedule (Fig. 8). According to the network resource definitions, the robot translates the end-to-end service definitions into sets of necessary connections, and maps their definition into elementary actions to be executed on various components of the network infrastructure. This results in a series of command scripts (typically SNMP Set-Requests) that are scheduled as UNIX at-jobs for automated execution at the requested points in time.

The feedback from the ACMA is displayed in the unified view by its status monitoring mechanisms. As only real feedback is obtained directly from the network, a command that does not succeed will not erroneously be taken into account in the display.

Lessons learned

More and more of the network devices on the market provide an SNMP management interface. However, the experience gained in integrating their monitoring and control into the IMS unified view and ACMA applications highlights the importance of the completeness of their SNMP MIBs. To allow the correct integration of a complete IMS model and a further automation of its construction, the resource MIBs have to provide access to:

- state information about the hardware interfaces
- state information about the logical connections
- complete configuration to be automatically learned by the management applications
- sufficient performance indicators for providing relevant on-line diagnostics and for feeding statistics reports with complete sets of useful observations
- full control of the activation and release of logical communications resources prepared in the network elements.

This strong requirement is an important recommendation for future procurements of network equipment to ensure smooth integration with high-level management functions. Ideally, this comprehensive SNMP management interface has to be offered by the dedicated element managers performing a first level of consolidation of the management view of individual network resources.

For better automation of the creation of the unified service model, however, the network element configuration MIBs represent just one of the two information sources needed. Knowledge about conventions of network operations and usage practices needs to be taken into account in designing a formal high-level representation of the overall communications model.

Finally, the pragmatic approach followed in the incremental set-up of the IMS prototype has proved its worth. In can be considered as a rapid prototyping of the required functionality, which allows immediate exposure of the proposed solution to operations personnel and immediate incorporation of their valuable comments. The mapping to the greatest extent possible of the individual features required to built-in mechanisms in the commercial-off-the-shelf network management platform allows

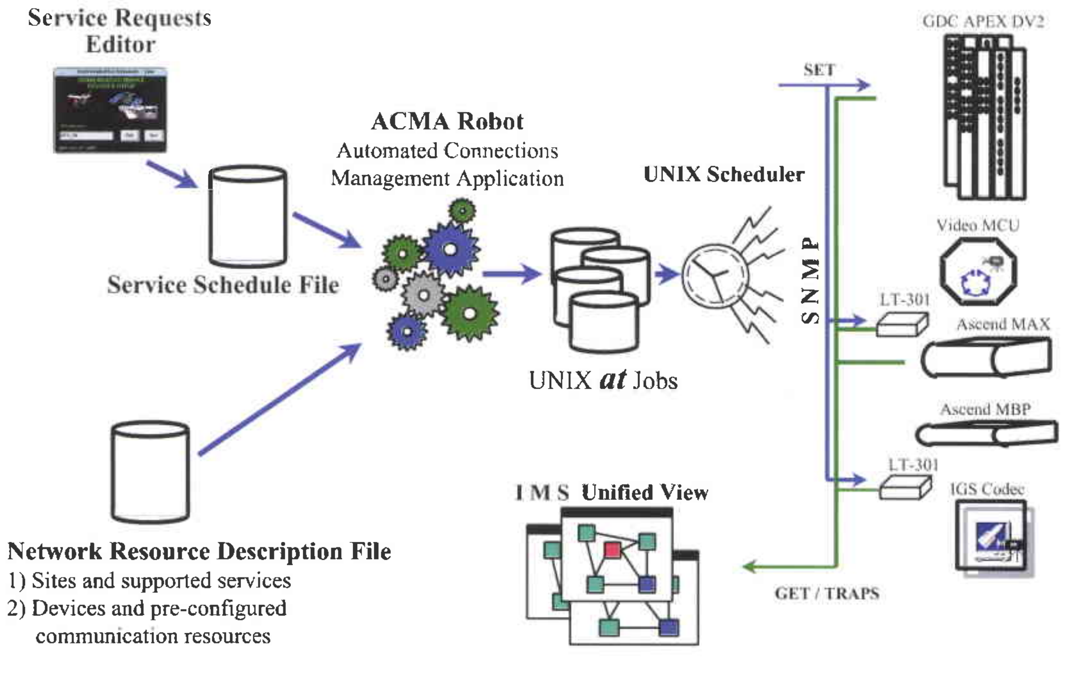

Figure 8. An Automated Connections Management Application (ACMA)

practical solutions suitable for real network operations to be achieved quickly. Those individual solutions can then be loosely combined to form a complete IMS system.

Next steps

The manual process of building up the unified view network model and the associated intuitive views will become too heavy when scaled to really large networks. A formalism has therefore to be inferred from this initial work for designing tools that can help in further automating the generation of the IMS custom configuration. Additional sophisticated Service Management Applications can be developed to tackle special management requirements of voice, video and data services. For instance, an appropriate visualisation of any degradation in quality of service, e.g. temporary loss of redundancy, would be valuable in supporting proactive troubleshooting.

The Automated Connections Management still needs to be extended to cover the complete set of communications characteristics, e.g. asymmetric/multipoint connections and specification of requested quality of service. Additional validity checks and the detection of potential conflicts in the communications service schedules are also necessary to allow the ACMA to be interfaced to the actual mission planning.

Finally, the intuitive presentation of overall communications state and its distribution to a large community of concerned users can be enhanced by the exploitation of new features provided by the recent version 6.0 of the HP OpenView platform. In particular, custom Alarm Correlation rules will help in integrating interrelated elements of information obtained from different management interfaces, and the new JAVA-based application programming interface will permit a better graphical presentation of web-based remotely accessible logical service oriented views.

Further developments of the IMS technology foresee adaptations to the operational conventions and constraints of OPSNET. Possible uses in other network environments such as the IGS are also envisaged.

Conclusion

The requirements for operational networks to support such mission scenarios as the Columbus science operations include integrated multimedia support and the integration of several different technologies (ATM, VSAT, ISDN, etc.), resulting in heterogeneous networks. The usage of the various network resources and services can be rather dynamic. The on-board experiment profile changes over time, resulting in changed geographical data flows on the ground, changed data rates or data volumes to be transported, and the involvement of different users in voice and video conferences. To minimise traffic charges, the network resources allocated have to match the actual needs and the appropriate connection technology has to be selected. Since Columbus will be operational for at least a decade, the connectivity and data-transport charges that will be accrued are major cost drivers.

The study and final prototype demonstration that have been reported here have shown that an Integrated Network Management approach can have worthwhile cost and labour-saving benefits in such situations:

- The integration of subnetworks and services into a single network management platform allows a unified umbrella view of the entire system to be maintained by a single operator.
- The automated connection and service configuration management allows complex configuration changes can be preconfigured and then executed at the scheduled time for a predefined period. Operator intervention is required only for real-time changes or trouble-shooting purposes. Configuration changes to circumvent performance degradations or equipment failures can also be automated according to preselected routines.
- Reports summarising the network performance history and trend indications are generated automatically.
- The implementation demonstrated a very effective way to match the human intuitive perception. This enables the operator to think in terms of more abstract logical network services including end-to end services.

All in all, the Integrated Network Management system approach allows one to optimise the cost of providing and operating network resources for highly complex operational systems with minimum staff requirements.

Acknowledgement

The authors would like to thank the other Aethis team members – T. Grootaers, L. Steenput and J. Maes – who helped to turn this project into a successfully demonstrated prototype. The demonstration would not have been possible without the active support of the Cap-Gemini team in charge of the IGS Testbed infrastructure at ESOC, namely G. Buscemi, A. Boccanera, F. Sintoni and A. Bernal.