

White Paper

Solve the multiple dilemmas of UC management

Focus: Managing multiple locations and hybrid environments

In the final part of this 3-part white paper we look at managing the performance of hybrid technologies, time zones and 'tin', spread across multiple locations.

As these locations extend across servers, devices, local and wide area networks and into the cloud this creates a UC management dilemma that requires specialized performance management to resolve.

By Sue Bradshaw, Technology Writer, Integrated Research



Preface

In this final part of the ‘Solve the multiple dilemmas of UC management’ white paper series we look at the dilemma of managing the performance of UC across physical, logical and cloud locations to ensure seamless integration between UC devices, services and applications.

You can download parts 1 and 2 here:

Part 1 focuses on managing the first two dilemmas – multiple technologies and multiple vendors, and juggling performance management of multiple technologies from multiple vendors, distributed across multiple locations, supporting multiple applications. [Go to download](#)

Part 2 is looking at the next dilemma, that of managing the performance needs of multiple stakeholders. It identified where stakeholders come from in UC projects, what their needs are and how to address them. [Go to download](#)

About Integrated Research and Prognosis

Integrated Research is the creator of Prognosis® multi-vendor UC ecosystem management for Avaya Aura®, CS1000, Cisco and Microsoft Lync UC platforms. Prognosis helps you manage the infrastructure needed to implement and maintain successful UC and SIP solutions. For more information visit www.prognosis.com/uc

Dilemma 4 – Managing Multiple Locations

It's not just managing geography any more: it's about managing locations

For many years local and wide area networks have meant that people can work as if they are co-located, even though they are geographically separate. And with UC applications now deployed on these networks enabling greater real time collaboration there is less need for bricks and mortar investment. The concept of a location now includes people working from home, the mobile workforce and of course those working from the cloud.

All this makes it possible to create a borderless enterprise for business and customer services that can function independently of location and time. Businesses can leverage time zone differences to create a 24 x 7 working day if that's what they need. And as they must cater to an increasingly global client and prospect base, this may prove to be an inevitable requirement. They'll be able to evolve into 'always-on' enterprises through the availability of cloud services of all types enabling them to collaborate and make decisions in real time, regardless of what time it is or where they are located.

This kind of flexibility also enables employees to make lifestyle choices about where they live, the hours they work, and has the benefit of contributing to core HR objectives like staff attraction and retention, and employee satisfaction.

Smooth operations hide the hard work underneath

Like the familiar analogy of a swan gliding gracefully across a lake with its pedaling legs invisible beneath the water; the smooth functioning and apparent simplicity of the highly performing 'always-on' 21st century enterprise belies the heavy lifting that must occur under the surface. And up in the sky, if you include cloud as a location.

The combination of devices, technologies, vendor platforms and versions located across physical, virtual and cloud platforms are big contributors towards performance management complexity.

A way to address this complexity is to put in place a hierarchical performance management model that allows the flexibility of local monitoring autonomy that can also be rolled up into a centralized '60,000 foot' or 'bird's eye' management view.

In this way you get the best of all worlds. High level views provide the big picture. And deep drill down details enable IT staff to observe any degradation, delays in processes and communications and allow them to deal with issues immediately.

And this brings up another aspect related to location management – the way issues are identified and resolved will depend on whether you have on-premise UC infrastructure, UC as a Service (UCaaS) or a hybrid of the two.

Performance monitoring for on-premise UC

In-house UC performance management requires dedicated IT resources to integrate communication and collaboration applications, and monitor their performance and availability. Whether the staff are employees, contractors or service integrator resources, they will need to know the status, availability, utilization and performance of all the UC infrastructure to meet their responsibilities for ensuring uptime, quality and sufficient capacity.

The need to maintain some on-premise services is common for a variety of reasons, including legacy assets, customized applications, security and information protection. This includes not just where data is stored but the routes it takes and where it is backed up. Also despite the growing trend to use cloud-based UC services, many IT managers are skeptical about their maturity and are concerned about enterprise support issues, especially for remote locations.



The grace of a swan on water belies all the hard work going on under the surface.

Another issue that affects a business's ability to move to a cloud-based service model is customization. It has emerged as a significant factor that extensive customization becomes firmly entrenched within an organization's culture, processes and practices. As it may not be possible to continue extensive customization cost-effectively in a UCaaS environment, trade-offs may need to be considered between continuing with in-house customization versus the benefits of reduced costs achieved through UCaaS.

What do I need to know?

When managing on-premise UC infrastructure vital performance metrics will include:

- UC server health, performance and capacity
- Voice quality across vendors, locations and technologies like Cisco, Avaya, Microsoft Lync and Acme Packet
- TelePresence availability, performance and quality
- Gateway and trunk utilization and capacity
- Busy-hour activity across all PBXs
- Number of calls, voice streams and Erlangs
- UC server health like Presence, Instant and Unified Messaging
- Session Border Controller availability, performance and capacity
- Virtualized server performance

How do you manage all your sites?

You don't need to make an all or nothing choice when deciding between centralized or distributed management; you can in fact have the best of all worlds. A combination of local monitoring and central management means you can administer every site in your enterprise as you like.



You can have
the best of all
management
worlds

Manage local sites autonomously with issues and reports forwarded to the NOC

For a variety of reasons a local site may need to manage its day-to-day activities based on its available staff, resources, skills and technology. From a head office's perspective it may wish to receive that information summarized to a centralized management location, like a NOC or data center.

This information may be required on a daily, weekly or monthly basis providing visibility into trunks, server performance, call quality and any alarms that have been generated in that time. It is likely to include:

- **Moves, additions, changes and deletions (MACDs)**
Identify quantities, inventory and deployment rates of technology / products (i.e. phones, video endpoints, etc).
- **Server health and voice quality**
Provide visibility into calls, errors, performance of trunks and quickly identify all components and levels of firmware, specifically problem ones.
- **Conditions that caused alerts**
Create and fine-tune conditions to trigger thresholds that when met or breached a specified number of times, over a specified period will send a message to alert destinations like email, SNMP traps, MOM consoles and logs.
- **Summary of previous day's bad calls**
Create an overnight report for the NOC with a list of bad calls from the previous day. This enables staff to quickly identify problem areas across the enterprise before they impact users the following day. Notifying support staff and any affected users will help prevent a flood of user calls to the help desk, either in the local site or at the NOC.

Manage all sites from a central location

The advantage this brings is insight to the bigger picture. A consolidated high-level view of the entire UC ecosystem combined with deep, detailed drill down information significantly reduces the mean time to identify and resolve issues, across all locations. This insight will deliver comprehensive visibility into gateways, trunks and performance of WAN links. In a single screen NOC staff can see the health, usage and capacity of all of the organization's trunks and obtain conclusive evidence of whether problems should be addressed by the local or central staff, and which of the teams, voice or data, should address them.

Comprehensive visibility into the necessary components beginning with voice quality all the way down to the individual processes on PBX servers and supporting applications will quickly identify a faulty or failed server, application or service. For example if a central SIP trunk resource has been degrading, action can be taken or a failover plan put in place before users are impacted elsewhere in the organization.

Centralized reporting meets the requirements of IT staff who need to record and store the calling/called parties' numbers and the duration of the calls. Whether its for compliance or legal reasons, this type of information also identifies past usage and allows staff to plan for future and/or seasonal demand and either increase or decrease WAN service provider costs.

Performance monitoring for Unified Communications as a Service (UCaaS)

As corporate footprints expand, cloud-based UC gives enterprises access to strong feature sets they don't have to deploy and maintain themselves, service provider expertise, responsiveness, on-site service and reporting. They can also take advantage of redundancy, scalability, backup and disaster recovery features that are costly to maintain in-house. And they're removed from much of the detailed and time consuming work of performing system updates and enhancements.

These benefits mean that if you're accessing UC from the cloud as a service, you are freed from the majority of time-consuming on-premise responsibilities. However your new responsibilities include ensuring that service level agreements are met and that you receive the services you're paying for.

This might include:

- Access to up-time statistics
- Minimal down time and mean time to repair
- Acceptable voice quality
- Presence updates synchronized quickly
- Emails arriving within an acceptable time, without spam or virus-laden attachments
- Response times for trouble tickets
- After hours support and service
- Service health information
- Integration and working with other providers
- Reports and metrics regarding applications used, and maintenance updates completed



Performance monitoring for hybrid on-premise and UCaaS

Most businesses are at least considering, or if not already using cloud-based UC services to reap benefits like cost, energy and resource savings. However, certainly in larger organizations, few find that the decision to use a cloud-based service is an all-or-nothing decision. The need to maintain some on-premises services is common for a variety of reasons, including legacy asset support, customized applications, and security and information protection.

However as time passes the flexibility of being able to deploy hybrid UC gives system architects a choice of embracing the cloud where it makes sense while continuing to use on-premise platforms as needed. If the decision is made to embrace UCaaS, staff may find other integration dilemmas present themselves. One of these is the multi vendor, multi location dilemma that you and your provider will need to solve to ensure facilities like Presence work successfully across the enterprise's staff and business processes, irrespective of location.

You'll also need to ensure that monitoring and support tools and procedures allow you to troubleshoot problems holistically and avoid finger-pointing. One of these issues is ensuring the currentness of applications and firmware as you and your UCaaS provider may upgrade at different times.

This makes it vital to know what versions of applications, hardware, software and firmware you have. Small version and firmware inconsistencies can create those type of knotty problems that can be very difficult to resolve.

The good news is that whether your infrastructure is in your data center, in the cloud or a hybrid of the two (or more) you can manage physical, virtual and cloud service delivery from a central location on one screen. The choices, permutations and complexity that are now available can be quite overwhelming so the following example of managing UC performance for a mining company's Asia Pacific (APAC) region might be helpful.



Case Study:

Multiple technology, vendor and location management through a single pane of glass

A mining company with its head office in Hong Kong has a global network comprising multiple international sites in Australia, Singapore, Japan and Indonesia. The local sites require a degree of autonomy but global IT needs a 'bird's eye' view of the entire enterprise to ensure cost efficient routing, fail over capabilities and data sharing.

Background, geography and a bit of technology

The company is managed from Hong Kong, which is the APAC hub. Four secure high-performance MPLS links providing routing, security and redundancy connect the company's international sites.

All IP traffic runs over these connections including virtual private network services back to the APAC hub. Voice and video packets are marked for expedited forwarding.

Growth has created a hybrid environment

As the company has grown in part via merger and acquisition, it has a mix of VoIP and UC technology that it needs to manage from the APAC hub. As each international location needs a degree of autonomy for efficiency and authority, they manage their own IT needs during the business day, including configuration and administrative tasks. The global IT staff in Hong Kong receive daily summaries of MACDs, server health, voice quality and capacity information; and always retain the ability to manage each site directly.

To see how this works let's take one country as an example – Australia. We'll then see how this data rolls up to the Hong Kong office and how control is managed from Hong Kong when needed.


Australian technology snapshot

The Prognosis server in Hong Kong is configured as a managing server at the top of the hierarchy for each of the international locations in Sydney, Jakarta, Tokyo and Singapore. This gives it the ability to manage any other Prognosis server in the network.

A Prognosis server in Sydney manages monitoring servers for four interstate and provincial locations. Prognosis servers are located in Melbourne and Adelaide, while the smaller sites in Broome and Perth are managed directly from Sydney. Sydney has deployed Avaya Aura® throughout with Microsoft Lync clients and the Melbourne office has a Cisco Unified Communications Manager cluster, and uses Microsoft Exchange and Unity connection.

All voice, video, application and internet traffic from Australia is routed via Sydney's Acme Packet SBC to its locally cloud-based service provider.

In Adelaide an office has been acquired as a result of a merger, and like Sydney, runs Avaya Aura and Microsoft Lync. The Perth and Broome offices run Cisco Unified Communications Express for a small number of users, as well as Cisco Unified Unity Express for voice mail. There is a high-speed and high-capacity link between Melbourne and Adelaide.





The requirement for the company's Hong Kong office is to manage this hybrid of technologies and vendors from its central location.

Staff receive reports at the end of each day detailing any devices that have been re configured, added or deleted as well as an overall picture of voice quality, infrastructure availability, performance and alerts.

Once the business day is over, Hong Kong takes over direct control of the Australian offices through its central management and data aggregation server, and receives alerts if processes, or performance benchmarks are breached.

This management server enables head office staff to view all vendor platforms and technologies through a single pane of glass. There is no need to log into each location and a backup route is configured for performance data to flow to Hong Kong if the primary route fails.

At no point do the international offices have access to Hong Kong's Prognosis resources, but the flexibility is there to allow it if the firewall is configured and administrator's permissions are added.

Alerts

Alerts are configured for core process thresholds at all locations together with CPU, memory and disk usage and message queues.

The results of these alerts are sent via SNMP traps to an enterprise Manager of Managers in Hong Kong's NOC only if they are classed as critical, with the alert duplicated via email to the Prognosis administrators in Sydney. All informational and warning alerts are managed out of Sydney.

Reports

Each night a summary of MACDs, resource issues, mail queues, number of calls are summarized and sent to Hong Kong. These reports ensure that although the NOC doesn't manage day to day performance in Australia, it can ensure that operations and resources are optimized and expensive outages are minimized.

Summary

A flexible management hierarchy that encompasses multiple vendors' platforms, technologies and versions is successfully managed centrally via a minimal number of Prognosis servers, which can be either physical or virtual machines. It's also possible to install Prognosis as an application on any virtual machine to provide deep drill down process-level information for that guest, as well as all other guests and the host machines they reside on.




In this way irrespective of vendor, version, technology, stakeholder or location Prognosis ensures:

- Local and central IT staff have proactive insight to their sphere of operations
- Forensic analysis enables reduced mean time to identify, convince and repair
- Operations and resources are optimized
- User perception and satisfaction are achieved
- Expensive outages are minimized
- Capacity planners have the information they need to ensure successful day-to-day operations, future growth and adoption of new technology like Microsoft Lync is successful

For more information please visit our website at <http://www.prognosis.com/uc> or contact us at info@ir.com.

Contact us

Americas **t:** +1 (303) 390 8700 **f:** +1 (303) 390 8777 **e:** info.usa@ir.com
Asia Pacific/Middle East/Africa **t:** +61 (2) 9966 1066 **f:** +61 (2) 9966 1042 **e:** info.ap@ir.com
United Kingdom **t:** +44 (0) 1344 894 200 **f:** +44 (0) 1344 890 851 **e:** info.europe@ir.com
Germany **t:** +49 (69) 710 455 255 **f:** +49 (69) 710 455 450 **e:** info.germany@ir.com

 twitter.com/IRPrognosis
 facebook.com/IntegratedResearch
 linkedin.com/companies/integrated-research

For more information visit our website at www.ir.com

UC0192-01



About Integrated Research (IR) - IR is a leading global provider of high-definition performance monitoring, diagnostics and reporting software. The creator of award-winning Prognosis® – IR builds scalable, customizable solutions designed to manage continuity-critical IT infrastructure, payments and communications ecosystems. IR provides real time solutions that give customers the insight they need to run their business, reduce their operating costs, minimize risk, and gain competitive advantage.