# Economic feasibility analysis of seamless multi-homing WAN solution

A. Mäkelä*, H.Warma*,K.Kilkki*, A. Decros*, J. Manner*
*Aalto University, Department of Communication and Networking, Finland
antti.makela, henna.warma, aurelien.decros, kalevi.kilkki, jukka.manner@aalto.fi

***Abstract*—High-availability wide area network (WAN) connectivity has traditionally been implemented with a single service provider building links with high quality hardware, redundancy, and guarantees, or with an economical but technically limited approach. We are proposing a new approach, Redundant Array of Independent Internet Connections (RAIIC), which could enable same technical flexibility as traditional high-availability services but lower pricing. In this paper, we analyze the current market situation in WAN market and the impact of a new stakeholder, a virtual service provider (VSP), on the market. We consider the possibilities for an existing business transforming into a VSP.**

***Index Terms*—redundancy, reliability, RAIIC, virtual service provider**

## I. Introduction

Traditionally, redundancy in the Internet and site-to-site intranet connectivity has been implemented with methods that are costly either in monetary or workload terms. A single service provider can set up a high-availability connectivity using various wide area network technologies such as MPLS VPNs [1], and customer and the service provider negotiate a service level agreement (SLA). This approach works well but can be prohibitively expensive for smaller entities who have to cope with smaller budgets. More economical alternatives, including multi-homing using routing protocols, site-to-site VPNs over Internet and network address translators (NAT), have varying degrees of limitations in flexibility.

For consumer purposes, economical broadband and even mobile broadband is starting to become more or less ubiquitous. Consumer broadband works "well enough" for most purposes, although no guarantees are ever given, and the connectivity could go down or become unusable without any kind of warning. For business-critical applications, utilizing such connectivity has traditionally been completely unacceptable due to unpredictability.

The technical limitations of economical approaches have caused the market for redundancy to be highly divided. On the high end, there are true high-availability solutions that come with full range of network services and flexibility with instant responses to outages. On the low end, there are several economical redundancy approaches, but they typically have limitations to their flexibility, recovery times and deployment. There are no feasible "mid-range" offerings.

We have previously [2] introduced a concept called RAIIC (Redundant Array of Inexpensive Internet Connections), for economical WAN connectivity for corporations or similar entities. We have researched RAIIC utilizing Mobile IP [3] and done studies with simulation [4] and a real-world implementation [5] to further evaluate the feasibility of the technology. The approach makes it possible to obtain equivalent reliability compared to traditional high-availability solutions with lower price by virtualizing a bundle of cheap, unreliable connections to a single, reliable one. Although our work uses Mobile IP, other protocols such as host identity protocol (HIP) [6] could be used instead.

In this paper we analyze whether a business model for a virtual service provider (VSP) based on this concept is feasible when RAIIC is acting as a technological enabler for VSPs. We are basing our assumptions and conclusions on published data, interviews, stakeholder identification and SWOT analysis. Similar methods have been used e.g. for analyzing business aspects of multipath TCP [7].

In Section II, we cover the current WAN market situation and different stakeholders. In Section III, we cover the differences between available access technologies. In Section IV, we discuss the current pricing levels of the technologies. In Section V, we conduct an in-depth analysis on the VSP business model. In Section VI, we conduct a SWOT analysis on how current market stakeholders could shift into VSP operations. In Section VII, we briefly observe reactions from other stakeholders. Finally, we conclude the paper in VIII.

## II. Analyzing current market stakeholders

Since we are conducting analysis on potential new player in the marketplace, the VSP, we need to analyze what are the existing roles in the market. We have identified the following roles that an entity may have. A single entity may of course have multiple roles - for example, a large telecommunications provider could act as a service provider for enterprises and consumers, provide integration services and also provide hosting services.

The roles we have identified are as follows:

- Customer - Entity, such as corporation, purchasing network services. Customer has end-users.
- End-user - A user of network services. Typically an individual, but could also be e.g. a sensor or similar automaton.
- Consumer service provider - Provides residential/consumer-grade broadband services with various technologies, including xDSL, cable modems, and various wireless technologies. The service provider

may operate their own network or lease it from another provider. Typically offers Internet access combined with own services such as broadcast television ("triple play" scenario). The services are marketed directly to end-users.

- Business service provider - Provides network services for business customers, including other service providers. Unlike in consumer case, the customer typically has at least some service assurances, such as technical support response times. While a product portfolio exists, products can be customized on a per-customer basis. This is typical especially in case of large enterprise customers - for smaller entities customization is usually limited.
- Hosting provider - Hosts internal and external services for customers, such as company website and e-mail services. The hosting provider usually has some sort of guaranteed service level. The costs for maintaining high availability are shared amongst several customers, thus making the per-customer pricing more affordable.
- Integrators - Provides customers with expertise for designing, implementing and operating their infrastructure. Typically, unless customer has highly experienced inhouse IT department, partners are used to provide at least some input in the technology strategy development.
- Network device vendors - Provides hardware and software that is used to implement network. Depending on vendor, may choose to direct their products to only one or all of the above segments. One possible division is "carrier-grade" equipment for service providers, "enterprise-grade" equipment for customers, and consumer products directly for end-users.

In addition, there are additional stakeholders, who do not significantly affect this model. Such players include e.g. resellers which arepart of vendor's logistics chain..

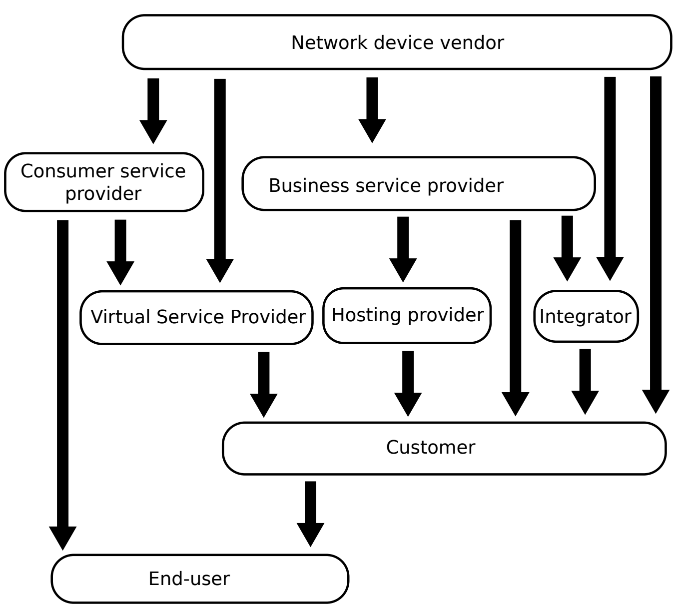


Fig. 1. Stakeholder relations diagram

In Figure 1 we are showing the customer relationships of the entities. In addition, the figure shows where our new role, the virtual service provider, would fit. VSPs would be purchasing equipment from device vendors, links from consumer service providers, and marketing the WAN service to customers. The VSP is studied fully in Section V.

## III. Differences of connectivity approaches

Since the primary driver we are introducing is a new technology, the technical and pricing differences of currently available solutions should be reviewed. To facilitate the review, we have conducted interviews with applicable parties and studied publicly available materials from service providers, such as price listings. The interviewees requested confidentiality due to contractual obligations, and thus we cannot publish the original interviews or exact monetary figures.

Since the VSP may attempt to build infrastructure over consumer-grade connections, we have to define what constitutes a consumer-grade connection and what are the differences to business-grade connectivity. Even the most basic business-grade connectivity is priced higher than technically equivalent consumer-grade connectivity. While there is some overlap in available network services and technologies, the primary differences stem from other factors.

Business grade connectivity usually always offers at least direct contacts with service provider's sales and accounting teams and dedicated technical support and incident management. This approach allows for customer-specific contracts; variable terms, discounts, and so on. From a purely technical perspective, the business-grade connectivity typically has priority in the service providers core networks and less overbooking, even without any explicit guarantees.

Possible optional network services have a much more diverse portfolio in business-grade connectivity; in consumer-grade connections you may be able to purchase some rudimentary services, but e.g. dynamic routing protocols and site-to-site WAN connectivity are completely unavailable. In contrast, with business-grade connection almost anything is possible for a certain price.

A table of the different connection technologies and their features is listed in Table I. In the following sections, we cover the approaches with more detail. Our RAIIC approach attempts to combine the benefits of traditional solutions - fast response time and flexibility - with a price level of the economical offerings. It can be considered an overlay service running over basic, non-guaranteed access technologies. However, the service is capable of switching to alternate links practically seamlessly during outages, and has no limitations on inbound connections, topologies, or similar issues. To the customer, RAIIC simply appears as another WAN technology.

### A. Traditional high-availability WAN solution with SLA

Traditionally, high reliability is achieved by obtaining a high-availability connectivity from a service provider. In this case the provider implements redundancy by utilizing such measures as multiple physical links, failover-capable

TABLE I
PROS AND CONS OF DIFFERENT CONNECTION TECHNOLOGIES

| | | |
|---|---|---|
| Traditional high-availability solution with SLA guarantees | Pros | • Excellent service level<br>• Well-established method |
| | Cons | • Highly expensive |
| Multi-homing via several providers | Pros | • More economical than traditional high availability<br>• Service provider independence |
| | Cons | • Competence requirements<br>• Failover convergence time |
| NAT, DynDNS, VPNs | Pros | • Economical<br>• Simple to deploy |
| | Cons | • Visible to end-users in case of outages<br>• Inbound traffic reaction time |

equipment with redundant power supplies, and other elements. In addition, the service provider has processes to support high-availability SLAs, such as fast replacement of faulty equipment. Redundancy is implemented locally using such methods as virtual router redundancy protocol (VRRP) [8]. The response and recovery times to an outage is in the sub-second range, being for all intents and purposes seamless.

The pricing is heavily affected by such factors as whether the service provider has any shared infrastructure in the area, such as a metropolitan area network (MAN). While the provider's core network is usually always shared between multiple customers, at least the last-mile-links are typically dedicated for a specific customer.

### B. Multi-homing via several providers

In this case, there is necessarily no SLA from any single provider. Instead, the redundancy is created by connecting via multiple service providers concurrently. A routing protocol, typically BGP [9], takes care of switching to alternate connections in case of network or equipment failures. The routing protocol approach has significantly slower response time, since the routing changes have to propagate throughout the network, and the propagation time can grow very large. In worst cases, the propagation can last in the order of several minutes or even up to a full hour [10]. This can occur especially when the connectivity is used for accessing the Internet, as any routing changes have to propagate globally.

The monetary costs are less pronounced than with single-provider approach. However, the customer still needs business-grade connectivity with dynamic routing protocol support from multiple service providers. In addition, managing the address space within customer's own autonomous system (AS) incurs additional IT costs - regardless of whether the IT is managed internally or outsourced to an integrator.

This approach can be combined with traditional SLA-based methods. If combining approaches, the SLA guarantees are typically set lower for each individual provider than in a single-provider scenario. Due to the exponential pricing pattern in relation to reliability (see Section IV) this may allow obtaining relatively high reliability with lower costs.

### C. Economical approaches

More economical approaches for redundancy at customer sites are available. However, unlike the previous approaches they are somewhat lacking in flexibility and robustness. Furthermore, no guarantees whatsoever are offered. These kind of approaches are implemented by an overlay service over basic, non-guaranteed access technologies, including consumer-grade links and basic business-grade links.

The most rudimentary approach that simply allows redundant Internet access, NAT with multiple links [11], simply translates internal network addresses to external ones on demand and uses some sort of heuristics to determine which outgoing link to use. While simple to implement, the system can typically only accept inbound connections from a specific link, so there is no redundancy to any services offered to the Internet. Dynamic DNS [12] can be used for directing inbound traffic to a degree, but this cannot react quickly to rapid changes in conditions, such as link outages or even changing dynamic IP address assignments.

Combined with IPSec [13] or another VPN technology, the NAT approach can be used to form site-to-site WAN connectivity. However, VPN solutions typically require at least one point of a VPN connectivity to be fixed - in practice this means at least one site having a more expensive, guaranteed connectivity. Even if a single site, such as headquarters, could be set up with appropriate capabilities, this would cause intra-site traffic to traverse via the headquarters, creating a bottleneck.

Another possible approach is utilizing a hosting provider for all services the customer's network provides. However, this limits the flexibility of services available and the typical providers only implement the most common services, such as e-mail, web hosting and possibly VoIP. Business-specific, custom services implemented in this manner would require a more expensive hosted server. In addition, having confidential, business-critical data at 3rd party premises always poses a risk.

## IV. OBSERVATIONS REGARDING CURRENT PRICING LEVELS

We have conducted interviews with a few businesses and looked up pricing examples from service providers. It should be noted that obtaining accurate data from service providers is not trivial. Business-grade connectivity is almost always sold with highly customized contracts that differ in discounts, and furthermore such deals are usually confidential. However, we have been able to obtain rough figures on a few occasions. In addition, some service providers have their list prices openly available.

The pricing pattern seems rather clear to the order of magnitude. The exact differences depend on such factors as region, population density, government subsidies and other factors affecting the quality of infrastructure. Our observations are mostly made from Finland. However, while exact price points vary, the basic observations seem to apply more universally, at least to the rest of Europe and North America.

Based on our research, three basic price points seem to emerge. The lowest pricing is for the consumer-grade broadband Internet access connectivity. In Europe, this is typically in the order of 5-50 € per month, offering speeds from 1 to 100 Mbps. The service typically comes with a dynamically allocated IP address that may or may not be behind a carrier-grade NAT. Especially in fixed-term contracts, the service provider includes a modem or other network access device as a gratuity. There may be restrictive terms of service (TOS) covering issues like connection sharing and types of traffic allowed.

Next tier in pricing is basic the unguaranteed business-grade connectivity, providing access to better technical support, public IP address blocks, and similar services. The pricing levels are typically 2-3 times that of the equivalent consumer-grade connection, and may have additional investment costs for networking equipment. Additional services come with extra costs, but are not typically limited in any way. Furthermore, TOS is much less restrictive on the usage of the network.

Adding guarantees to the connectivity increases price considerably. Even the least expensive cases, such as SLA specifying 99% uptime, increases the pricing by an order of magnitude or more. In [14], Gartner report suggests that for every additional "nine" stacked to the reliability factor results in a 30% increase in pricing. Depending on the exact terms of the SLA, the costs quickly become very high.

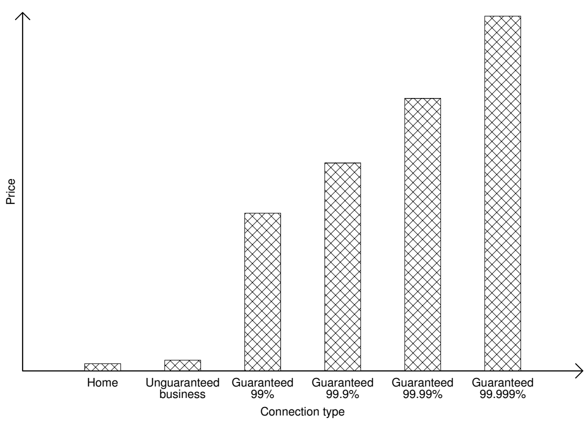


Fig. 2. Illustration on SLA's effect on connectivity price

The pricing structure is illustrated in Figure 2. There is a noticeable gap between guaranteed and unguaranteed connectivity. There are no existing "mid-range" offerings.

One of our interviewed cases is a Finnish medium-sized company with two sites; an office and a data center. Both are located in Helsinki metropolitan area, roughly 50 km apart. The company provides services over the network to external partners, and the services are hosted in the data center. As such, availability of the services and the data-center site, is business-critical. In this case, they have obtained a guaranteed, high-availability connectivity for connecting their data center to partners' extranets, and a unguaranteed site-to-site WAN connectivity for connecting their office to the data center from an European service provider.

The cost structure illustrates the effect of high-availability to the pricing profoundly. The site-to-site link is a 1 gigabit connectivity over the service providers' metropolitan area network. The infrastructure is shared, but no guarantees are given. The data center's connectivity to partners is a redundant pair of 10 Mbps links. Both of these links come with similar tech support agreements: 4 hour response time. The site-to-site link costs roughly 1500 €/month. The redundant pair of 10 Mbps links costs roughly equal amount *per link,* or 3000 €/month total.

Thus, in terms of €/Mbps, the guaranteed 10 Mbps link is more expensive by *two orders of magnitude.*

Although not exactly similar, the same orders of magnitude seem to apply in United States and Canada as well [15], [16]. Information obtained from another European service provider on the pricing structure revealed that the profit margins on enterprise connectivity are relatively high; the price points are set on customer-basis after bidding rounds. Since guaranteed connectivity is not a commodity and contracts are typically long, there is no real pressure to drive the costs down. Furthermore, the situation has apparently been similar for a considerable period. While connection technologies and thus bandwidths have improved, the price difference between guaranteed and unguaranteed connectivity has remained huge.

## V. Proposed Business Model for a New Player: Virtual Service Provider

Our approach is based on the idea that currently, guaranteed services are very expensive, either in monetary terms (redundancy from single ISP) or from management (multi-homing with routing protocols, VPNs) perspective. On the other hand, reliability based on more economical methods has limits in flexibility. The RAIIC approach allows for a connectivity that provides equivalent service compared to traditional high-availability solutions, yet utilizes economical infrastructure and thus can be offered with reduced price. Instead of exponential rise in pricing levels, there should be a linear one.

As stated in Section IV, the business-grade connectivity offerings with guarantees currently cost at least an order of magnitude more than equivalent consumer-grade connections. However, without guarantees, the technical support and similar business-grade services only cause a moderate price increase from consumer-grade connectivity.

Basic outline of the RAIIC model is illustrated in Figure 3. Virtual service provider would be an entity managing network elements providing the virtual WAN functionality with Home Agents and Mobile Routers. In addition the VSP would obtain the links for basic connectivity to customer sites to be presented as a single, virtual site-to-site connection. As shown, the virtual connectivity would automatically utilize whatever paths happen to be available. In case of a breakdown, RAIIC would transparently switch the traffic to alternate paths.

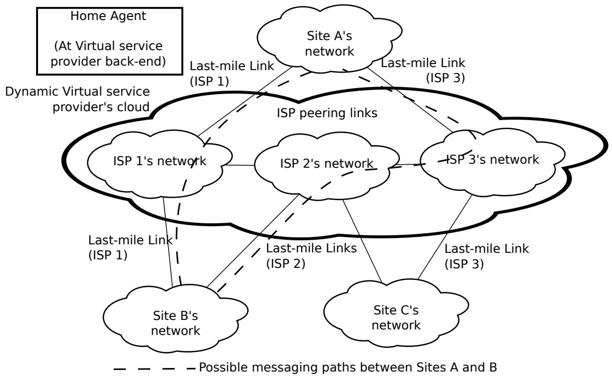


Fig. 3. RAIIC model using Mobile IP

During normal operation, all paths may be used concurrently with a load-balancing schema.

The VSP would thus implement and operate corporate WAN services to customers. The offerings would consist of planning, design, implementation and operation of high-availability WAN connectivity. The fundamental operation would resemble an union of several traditional roles.

To customers, the VSP would appear just like a traditional, business-grade service provider. The offerings would be relatively similar; connectivity with technical support at agreed-upon service levels. Internally, most of the functions, such as helpdesk, customer relations management, and marketing would be relatively identical.

Significant internal differences stemming from the technology require the VSP to be aware of possible service providers in their operating areas and statistical reliability of their unguaranteed links. The information is used for creating guidelines for quality levels in the offerings. In addition, the VSP needs to be able to respond to changing conditions proactively; e.g. if a specific provider shows signs of decreasing reliability, VSP responds by adding another link.

Although maintaining awareness of service provider quality requires additional work, the costs associated should remain relatively low - all that is required are a few consumer-grade connections for statistics-gathering purposes.

### *A. Analysis of the business environment*

In this section, we analyze several factors that may affect a business operating environment of the virtual service provider. We are considering the market environment from social, legal, economical, political and technical perspectives. These are covered below.

*1) Social considerations:* The RAIIC approach is not a consumer or a mass-market product. Social considerations stem from end-users; however, the end-users do not directly make decisions on network architecture or investments. The end-users' productivity in being able to conduct their normal, daily workflows is affected by the reliability of their network connectivity. If the network is unreliable, but budget for improving the experience is limited, the approach may become attractive as a solution.

*2) Legislative considerations:* From VSP's perspective, certain preconditions pertaining to consumer service providers terms of service are beneficial. It is typical for terms of service (TOS) and acceptable usage policies (AUP) of consumer-grade connections to disallow e.g. hosting servers or offering connectivity to third parties (e.g. outside the immediate family). If such terms are in service contract, regulations may affect whether these terms are actually contractually binding. Even if they are, the service provider may simply choose to not enforce them - even though hosting servers is typically disallowed, they are still very common.

The TOS issue can be circumvented, at cost, by utilizing the least expensive business-grade connection instead - a non-guaranteed basic Internet access. This can be slightly more expensive than consumer version, but still not the orders of magnitude that comes with guarantees.

*3) Economical considerations:* While the basic consumer-grade connections have gone through significant price erosion during the last decade, this has not significantly affected the contribution of guarantees to service pricing. The market of the high-availability connectivity seems to have been stagnated for a relatively long time and therefore the proposed market player might be able to shape the market and increase competition.

In the current situation, traditional guaranteed connectivity offerings for enterprises are perceived as highly expensive, and therefore a VSP might require extra effort to become a convincing solution. After initial momentum has been gained and the technology has established more credibility, demand will probably rise simply as a result of customers becoming informed of the approach.

For breaking into market, market research is needed for obtaining reference customers. The potential target customers are most likely medium-sized companies that are currently spending a significant amount of budget into traditional solution. Smaller entities most likely do not have such an extensive network and their references carry less weight. On the other hand, large enterprises have typically a much slower decision-making process and the costs of traditional connectivity, while expensive, do not represent a significant enough portion to attempt a new approach. Another advantage of medium-sized companies is that the network structure is typically limited to a few sites, so gathering feedback is very straightforward. The RAIIC solution can be used concurrently with traditional approaches, which allows the possibility to offer a safe migration/piloting path.

*4) Political considerations:* The local political climate and regulatory practices may either contribute prohibit or contribute to the operating environment of the VSP. We have identified two main issues from the political point of view:

- Net neutrality legislation in United States and similar approaches elsewhere could prevent service providers from affecting traffic by e.g. throttling bandwidth for lower-tiered applications. This can be considered a positive aspect for a VSP since the providers could not legally restrict the intra-site traffic from VSP.
- Competition in the service provider market. Regulation

enforcing sharing of physical infrastructure when capacity exists reduces price levels of providers, as more can enter the market using shared infrastructure. Since the VSP works by purchasing infrastructure from providers, this can be considered beneficial from both pricing and capabilities perspective.

Regulatory requirements help the technology to maintain better connectivity. Without regulatory intervention, the operating environment is more problematic, and may affect offerings.

*5) Technological considerations:* As the service is based on providing redundant connectivity over an array of unreliable ones, the implementation will have specific requirements.

The most obvious requirement is that there should be several providers for last-mile connectivity. However, the technology approaches should avoid any common elements; e.g. obtaining xDSL access technologies from multiple providers may not bring true redundancy, as the connectivity may use same physical cabling. xDSL, cable modems, and possible wireless technology, however, do not have overlapping elements.

Lack of available service providers is a larger issue. If there is only one choice for a provider, redundancy by utilizing multiple providers is obviously not possible. Influx of wireless technologies may alleviate this somewhat, as there are less physical infrastructure requirements.

Beyond the last mile, the service provider cores should be independent as well. However, several providers may share a common core of a higher-level network operator. If the network operator's core network becomes congested, all service providers may be affected. Thus, the VSP needs to be aware of such issues.

The implementation on site does not have any special requirements except the ones stemming from traditional high availability. The Mobile Routers should be provided with failover pairs, and individual routers should have redundant ports.

### *B. Differentiating from existing providers*

To the customers and end-users, the RAIIC approach would appear just like any other high-availability WAN connection. The primary difference between existing business providers' high-availability WAN offerings is the pricing level. Compared to the traditional high-availability service, the VSP would be capable of offering this at significantly lower price.

Additional differentiation is flexibility and service provider agnosticism. The VSP does not require the physical network access providers to be the same at each customer site. The traditional high-availability offerings require all sites to be connected to the same provider's network. Stemming from this flexibility is the possibility that the RAIIC-based offerings could be even more reliable than the traditional high-availability offerings, especially in areas with multiple providers.

### *C. Cost and revenue analysis*

The expenditures specific to the VSP, are shown in Figure 4. Expenses that are usual to business, such as office space and support functions are not shown. There are relatively few capital costs; the initial investments consists primarily of setting up the infrastructure required for the Home Agent. The costs for statistics-gathering infrastructure are relatively low. Once customer base grows, the largest capital costs stem from obtaining access connections. Besides the initial customer installation, the VSP may also choose to act proactively e.g. due to prevailing weather conditions and install additional connectivity later.

Operational expenditures that can be easily estimated are maintaining own infrastructure and personnel costs. Additionally, the VSP should be prepared to pay sanction fees for breaking an SLA in case the approach fails despite all precautions due to all access links being inoperative.

| OPEX |
|---|
| • Back-end infrastructure maintenance<br>• Senior level helpdesk functions<br>• Customer site installation work<br>• Sanction fees |
| CAPEX |
| • Purchasing access connections<br>• Gathering statistics on providers' network quality<br>• Deploying back-end infrastructure |

Fig. 4. VSP-specific expenditures

Revenue for the VSP would be coming from implementation fee and periodic fees for the virtual connection. The pricing level for the virtual, bundled connectivity should be distinctly competitive compared to traditional high-availability approaches, but still provide enough revenue to offset the costs. Since the traditional high-availability methods cause such a huge increase in price levels, this is probably achievable. The revenue obviously needs to offset all the costs and offer a satisfactory profit margin.

The model differs from traditional high-availability offerings by basing it on a more economical technology platform. However, business-grade connectivity typically offers more than just the connection; the customers are also purchasing a more dedicated helpdesk, account manager, and similar services.

In the traditional model, the implementation requires less planning and design work. The connectivity offerings are usually well-established products with a relatively smooth workflow but the actual installation may be very costly. Physical installations of primary and backup cabling in separate locations, obtaining high-quality hardware with seamless failover capabilities and fast on-site maintenance contracts, and so on. The RAIIC requires far more detailed design work to verify that access connections are truly independent, but physical installations are performed by the providers.

However, based on our data, the helpdesk functions and similar services do not appear to be significant drivers in the pricing model. The functions are shared across several customers; single helpdesk may serve requests from all customers, and the costs are divided amongst them. Thus, bulk of the cost does indeed come from installing and maintaining the network operations.

In cases where a customer would not want high-availability to a particular sites, the VSP may choose to simply resell a regular connection from the lowest bidder at each location.

### D. Offering SLAs without owning the infrastructure

The largest single issue, is how can the VSP offer services over infrastructure it does not own, or even control. Furthermore, the parties who own the infrastructure - existing service providers - may choose to alter their network services. An SLA typically stipulates some sort of sanction fees in case agreed-upon service levels are not met.

The VSP requires at least some information on reliability levels of consumer-grade service providers in the area as a basis for drafting service contracts. In a simple example, having a basic assumption that each individual link has 2% downtime, this becomes an uptime of 1 - $0,02^n$ with n links to each site. In case of 3 links, this is already 99,9992% uptime. However, uptime is not all that counts - individual service providers may suffer from a periodically congested network, which cause jitter and packet loss.

Mitigating issues stemming from utilizing third party (and in some case, competitor's) infrastructure is possible, however, these cause extra costs or lost business opportunities for the VSP. Mitigation methods include at least:

- Offer the service only in areas where you can assume smooth operation, e.g. metropolitan areas with multiple ISPs and relatively quick recovery times.
- Overbuy redundancy; in the above example of 99,9992% uptime, obtain a fourth connection for even more reliability while only guaranteeing the first 99,9992% - leaving a high margin for error.
- Implement the service with separate statistics-gathering and normal operation phases. During statistics-gathering obtain data on true behavior of the service provider networks and add or reduce links based on the results.
- Have network vendors implement a smart heuristics mechanisms on the Mobile Routers for making the handover decisions.

## VI. Analysis on becoming VSP

In this section, we analyze how different kinds of established players could expand and extend their operations to start operating as a VSP. We use SWOT analysis as our tool.

### A. Integrators

Integrators may participate either in their traditional role as consultants, who mostly plan, design and implement the network, but leave the operations to other parties, or also take care of the operations. In the former case, the technology does not create new opportunities as such. It simply becomes an additional choice in their approach portfolio. If the integrator also takes care of the operations, they become, in essence, the Virtual Service Provider.

Traditionally, integrators subcontract network connectivity from a single service provider based on quotations. With the VSP model it is possible to implement single customer network with a multi-provider deployment, which could reduce expenses even further. However, the established relations to service providers may suffer which can affect pricing for the offerings implemented with traditional methods. The SWOT analysis for integrators to become a VSP is shown in Table II.

TABLE II
SWOT Analysis for Integrator as VSP

| | |
|---|---|
| Strengths<br>• Technical expertise for implementation<br>• Established customer relations | Weaknesses<br>• Helpdesk processes typically optimized for incident responses, not continuous, proactive processes<br>• No experience on service provider business |
| Opportunities<br>• Transforming network operations to service provider-agnostic model | Threats<br>• Financial sustainability in case of high sanction fees<br>• Backlash from established service providers especially in pricing for existing, non-VSP connections |

### B. Business service provider

Traditional business service providers will most likely perceive the new technology as a threat. Providing high-availability connectivity currently has very high profit margins. Thus, they will likely attempt to marginalize the approach. This may include marketing, aggressive discounts, or similar methods to prevent the adoption.

However, a business service provider may also see the VSP model as a possible way to supplement their existing product portfolio, by offering a moderately priced high-availability. In addition, the VSP model allows the service provider to extend their reach by making high-availability possible outside their own infrastructure. However, obtaining access lines from other, competing, service providers may not prove straightforward. The competitors may utilize different kinds of methods to hinder the emerging VSP business, including regulatory methods. The SWOT analysis for service providers to become a VSP is shown in Table III.

### C. Hosting provider

A hosting provider already has high-availability infrastructure and proactive maintenance processes for hosting services to customers. The Home Agent which acts as a signaling anchor for the RAIIC model is simply another service. The hosting provider could use the RAIIC approach to allow businesses that have shunned hosting providers due to confidentiality considerations to keep their data inside their premises, but still make them economically available. However, helpdesk functions are usually centered on maintaining shared components, not customer-specific issues,

TABLE III
SWOT ANALYSIS FOR BUSINESS SERVICE PROVIDER AS VSP

| Strengths | Weaknesses |
|---|---|
| • Processes already optimized for service provider business<br>• Direct ownership of at least one access line | • No experiences on quality of other providers' networks<br>• Additional access lines may be difficult to obtain without losing transparency to customer |
| **Opportunities** | **Threats** |
| • Extending infrastructure reach virtually<br>• Creation of new market for moderately-priced high-availability | • Aggressive reaction by existing business service providers (competitors)<br>• Customers migrating away from traditional high- availability offering |

and thus resources are relatively limited. The SWOT analysis for hosting providers to become a VSP is shown in Table IV.

TABLE IV
SWOT ANALYSIS FOR HOSTING PROVIDER AS VSP

| Strengths | Weaknesses |
|---|---|
| • Infrastructure already in place<br>• Proactive processes in maintaining infrastructure | • No expertise on service provider business or operations<br>• Helpdesk processes not tuned for customer-specific issues |
| **Opportunities** | **Threats** |
| • New market segment for confidentiality-aware customers | • Scalability of processes with high demand |

## VII. IMPACTS ON OTHER STAKEHOLDERS

Network device vendors will have to conduct evaluation on whether the new technology is attractive enough to implement. Since the feature is about high availability, the implementation may require stringent programming practices and thus there is an increased cost in designing, implementing an maintaining the product. On the benefit side, the vendors will have a new item to market. The technology may become a feature on pre-existing products, or a new, dedicated product altogether.

Consumer service providers do not necessarily have any sort of reaction to the technology at all. However, they may become partners to integrators or VSPs, by providing detailed data on their connectivity implementations, advance warnings about planned downtimes, and other such issues. On the threat side, they may choose to place limitations on such services, in similar manner to limitations on e.g. P2P traffic.

Customers should benefit regardless of what happens to the technology adoptions. Even if the technology is marginalized by traditional business service providers, the increased discounts still provide incentives to obtain high availability via traditional methods.

## VIII. CONCLUSIONS AND FURTHER RESEARCH

After analyzing the potential business model of the virtual service provider, we conclude that there can be several ultimate outcomes, assuming the technology works as expected. In the most positive outcome for all parties, the new offering becomes a new moderately priced option for high availability. The VSP provides for a new market segment in co-operation with existing players, allowing entities which until now have not been able to purchase guaranteed connectivity to obtain one.

In a more competitive outcome, the new approach is seen as a competitor to the established high-availability solution providers, and starts making inroads into the traditional high-availability market. In this case, the reactions from business service operators, such as discounts, may supersede the approach.

In all cases, the customer should still see the benefits - either the new approach offers more economical high-availability or the pricing level of traditional solutions drops heavily.

This is the first effort on analyzing the economic feasibility of the RAIIC solution but before setting up a VSP business, further research is needed. Besides refining the technology itself, a quantitative analysis of the costs and revenues of the VSP should be carried out and potential markets to establish the business should be evaluated.

## IX. ACKNOWLEDGMENTS

This work was supported by TEKES as part of the Future Internet program of TIVIT (Finnish Strategic Centre for Science, Technology and Innovation in the field of ICT).

## REFERENCES

[1] I. Pepelnjak, J. Guichard, "MPLS and VPN Architectures", *Cisco press*, 2002

[2] A. Mäkelä, "Concept for providing guaranteed service level over an array of unguaranteed commodity connections", *in proceedings of The 25th Symposium On Applied Computing (ACM SAC 2010)*, March 2010.

[3] C. Perkins et al, "IP Mobility Support for IPv4, Revised", *IETF RFC* 5944, November 2010

[4] A. Mäkelä, J. Manner, "Bundling consumer connections - a performance analysis", *in proceedings of The 1st International Workshop on Protocols and Applications with Multi-Homing Support (PAMS 2011)*, March 2011

[5] A. Mäkelä, J. Manner, "Performance of a economical, redundant system for Intranet connectivity", *in proceedings of 16th IEEE symposium on Computers and Communications (ISCC'11)*, June 2011

[6] R. Moskowitz et al, "Host Identity Protocol", *IETF RFC 5202*, April 2008

[7] T. Levä, H. Warma, A. Ford, A. Kostopoulos, B. Heinrich, R. Wider, P. Eardley, "Business Aspects of Multipath TCP Adoption", *in proceedings of 6th EURO-NF Conference on Next Generation Internet (NGI)*, 2010.

[8] S. Nadas et al, "Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6", *IETF RFC 5798*, March 2010

[9] Y. Rekhter, T. Li, S. Hares et al, "A Border Gateway Protocol 4 (BGP-4)", *IETF RFC 4271*, January 2006

[10] R. Teixeira, S. Uhlig, C. Diot, "BGP Route Propagation between Neighboring Domains", *in Passive and Active Network Measurement, Lecture Notes in Computer Science*, 2007

[11] Cisco, "IOS NAT Load-Balancing for Two ISP Connections", *Cisco Document ID 100658*, February 2008

[12] P. Vixie et al, "Dynamic Updates in the Domain Name System (DNS UPDATE)", *IETF RFC 2136*, April 1997

[13] S. Kent, P. Seo, "Security Architecture for the Internet Protocol", *IETF RFC 4301*, December 2005

[14] N. Rickard, "Cost Cutting by Rightsizing Network Reliability", *Gartner research report G00155940*, April 2008

[15] Verizon Business, "Service Publication and Pricing Guide" [Online]. Available: http://www.verizonbusiness.com/us/publications/service_guide/

[16] TELUS Communications Company, "The TELUS Tariffs list" [Online]. Available: http://about.telus.com/publicpolicy/tariffs/