

RealNames Working Documents

May 7, 2002



RealNames Corporation
150 Shoreline Drive
Redwood City, CA 94065

Table of Contents

Business Proposal and Terms.....	3
RealNames Business Overview.....	8
User Experience Overview and Recommendations.....	22
RealNames Incremental Business Opportunities.....	28
White Paper.....	61
Minimum Assurances Agreement....	87

Business Proposal and Terms

Michael Arrington
650-486-5770
marrington@realnames.com

Business Proposal and Terms

1. Preamble.

In the lead up to the termination of the existing License & Marketing Agreement between Microsoft and RealNames, signed in March 2000, it has become clear that there are three key issues that must be resolved before a renewal of the agreement on its expiry on 28 June 2002 can be completed.

The issues are best labeled user experience and related resource issues, and business terms issues. This document proposes solutions to all of these issues and sets forth a framework for a long term mutually beneficial partnership.

2. User Experience.

Explanation of the Issues

The attached Keyword User Experience document describes in detail the reasoning behind the proposals set forth below, and also includes additional recommendations for product and process changes that we feel will do much to improve the overall user experience and reduce the use of Microsoft resources. The goal of each of these proposals is to set forth a workable solution that addresses these issues.

The decision last year to divide Keywords into *Basic Keywords* (only available for terms that have 10 or less searches per month on MSN) and *Keywords Plus* (which must be approved by a human being applying rules of quality control as defined in the agreement) was intended to create an automated basis for the solving of any issues around data quality and user experience. In line with these changes and to further improve the user experience, *Keyword Namespaces* were also required to pass review and approval by a human being, and Basic Keywords were removed from search results.

Following these changes it has become clear that there are additional steps we can take to further improve the user experience, and further reduce Microsoft resource costs. These solutions generally address one or more of the following: Microsoft resource use in *uploading* Keyword data; *processing* Keyword data; and *editorially monitoring* Keyword data.

- A. Uploading and Processing RealNames Keyword data. RealNames will implement a resolver lookup protocol. Implementation of this protocol will allow Microsoft to query the resolver in real time (better than 50 milliseconds) and ask the question “Do you have a Keyword for X?,” where “X” is the user input and associated context like country and language. If the answer is “yes” the query will be passed to RealNames for resolution as a Keyword. If the answer is “no” the query will not be passed to RealNames.

The advantages of utilizing the resolver lookup protocol include elimination of the data delivery (*uploading*) and indexing (*processing*) and the savings for Microsoft on both technical and staff resources that support the current processes.

RealNames already has this technology. It was previously deployed (and scaled) successfully with both Google and AltaVista in real time. Microsoft would need to add code to the search back end, invoking a lookup to the RealNames Keyword Service. This is a relatively minor change.

B. Editorially monitoring RealNames Keyword data. An unintended by-product of 2(A) above is that the editorial team at Microsoft will not have daily access to the RealNames data. To fix this, and allow Microsoft to trust RealNames data policies, RealNames will do two things.

1. Build an automatic policing agent that will throw out a daily “exception” file. This policing agent can be rules based and can be live within 30 days. As a minimum it will look at all Basic Keywords and be sure that there are less than 10 searches on MSN Search in the most recent month. Additionally it will look at all previously sold Keywords and ask if there is a spike in usage that suggests a change has occurred in the meaning of the Keyword, necessitating a human review. It could also implement any other rules Microsoft feels appropriate within the terms of the contract.
2. In addition to creating a daily exception file, RealNames proposes that they deploy a Keyword Management Services (KMS) screen exclusively for Microsoft editors, allowing access to Keywords sold and allowing (within the limits of the agreement) for the rescinding of Keywords by Microsoft editors.

This screen could be used worldwide and allow MSN worldwide to police the namespace that affects their territory. Note that RealNames already has a KMS product that is used by internal RealNames editors and the RealNames sales channel. Within 30 days a new module specifically for Microsoft’s use can be deployed. The potential for a trusted third-party to host this policing agent activity is another option for consideration.

C. If desired by Microsoft RealNames would continue to provide MSN access to its data files.

3. Microsoft Resources to Ensure Quality User Experience.

Explanation of the Issues.

Microsoft currently has to deploy resources in three ways. First, Microsoft uploads the data. Second, Microsoft processes or indexes the data in order that it is available through

MSN search. Third, Microsoft monitors the Keyword editorial quality across MSN worldwide.

The above changes in Section 2 will significantly reduce the resources Microsoft expends in all three areas. It would also allow RealNames to be “switched on” in all 244 locales that Microsoft recognizes in runonce.asp. This would create a consistent user experience across the world.

These things are all achievable within 30 days of your readiness to move forward.

4. Business Terms.

Explanation of the Issues

RealNames owes Microsoft \$25.5m which is due for payment on 2 May 2002. RealNames is unable to pay this amount and believes Microsoft’s best interests would be served by agreeing to renew the RealNames contract and growing the value of its 20% stock holding whilst continuing to enjoy a revenue share – partly guaranteed. The proposals below address this.

The proposal is as follows:

- A. Microsoft will forgive the \$25.5m owed to it by RealNames. All interest will also be forgiven. In consideration of the forgiveness of the notes and accrued interest, RealNames will pay to Microsoft \$5 m on the later of the date of execution of the renewed agreement and June 28, 2002.
- B. Microsoft will renew the RealNames contract. The term of the contract will automatically renew, subject to termination for cause, including non-payment of revenue share or minimum payments. The contract will be renewed automatically each subsequent year so long as RealNames has paid its commitments in full and has abided by the Keyword quality rules in a manner spelled out in the contract.
- C. RealNames will pay to Microsoft a percentage of total Keyword-based revenues (including Country Registry Licenses, Basic Keywords, Keywords Plus and Corporate Keyword Namespaces) for each fiscal quarter. The percentage to be paid to Microsoft for any particular fiscal quarter is dependent on total Keyword-based revenues for such quarter, as set forth in the table below.

<u>If Quarterly Revenues Are:</u>	<u>The Revenue Share is:</u>	<u>Annualized MS Payments:</u>
< \$10m	15.0%	up to \$6m
\$10m - \$25m	18.5%	\$7.4m - \$18.5m
\$25m - \$35m	22.0%	\$22m - \$30.8m
\$35m - \$45m	25.5%	\$35.7m - \$45.9m
\$45m - \$55m	29.0%	\$52.2m - \$63.8m
\$55m - \$65m	32.5%	\$71.5m - \$84.5m
\$65m - \$75m	36.0%	\$93.6m - \$108m
> \$75m	40.0%	at least \$120m

D. In the event that the revenue share for a given Microsoft fiscal year is less than the amount set forth in the table below, then in addition to the revenue share set forth above, RealNames will pay to Microsoft the difference between the Minimum Yearly Payment set forth below and the quarterly payments for such fiscal year. The payments set forth below are in addition to the \$5 m payment to Microsoft set forth in Section 4(A) above.

<u>Microsoft Fiscal Year:</u>	<u>Minimum Yearly Payment:</u>
2002 (July '02 – June '03)	\$5.0m
2003	\$6.0m
2004	\$7.0m
2005	\$8.0m
2006	\$9.0m
each successive year	lesser of (i) (previous year's guarantee * 1.1), and (ii) 40% of total Keyword revenue. Notwithstanding the foregoing, the MYP in FY 2007 and beyond will not be less than \$9m.

5. Additional Discussion Points.

- A. Discussion of marketing of Keywords by Microsoft
- B. Discussion of actions to resolve Netpia issue in Korea

RealNames Business Overview


Keith Teare
CEO
650-486-5555
keith@realnames.com

Executive Summary

After 5 years of operations RealNames Corporation has developed a comprehensive naming infrastructure for the Internet. This infrastructure consists of Name Registration Services, used by almost 100 companies worldwide to register names on behalf of almost 100,000 customers, and Name Resolution Services – used almost 500 million times in the first quarter of 2002, to deliver users to web pages.


During 2001 and 2002 RealNames has become increasingly relevant as a new layer of naming on top of the DNS. The Internet Engineering Task Force (IETF) and the Internet Corporation for Assigned Names and Numbers (ICANN) have both had discussions in which the RealNames technology has been discussed as having built a layer of naming capable of solving many problems with the old infrastructure.

In Q1 2002, with \$4.8m in revenues and \$5.6m of cash collections, the company was cash flow positive for the first time. Revenues for 2002 are set to be at record levels.

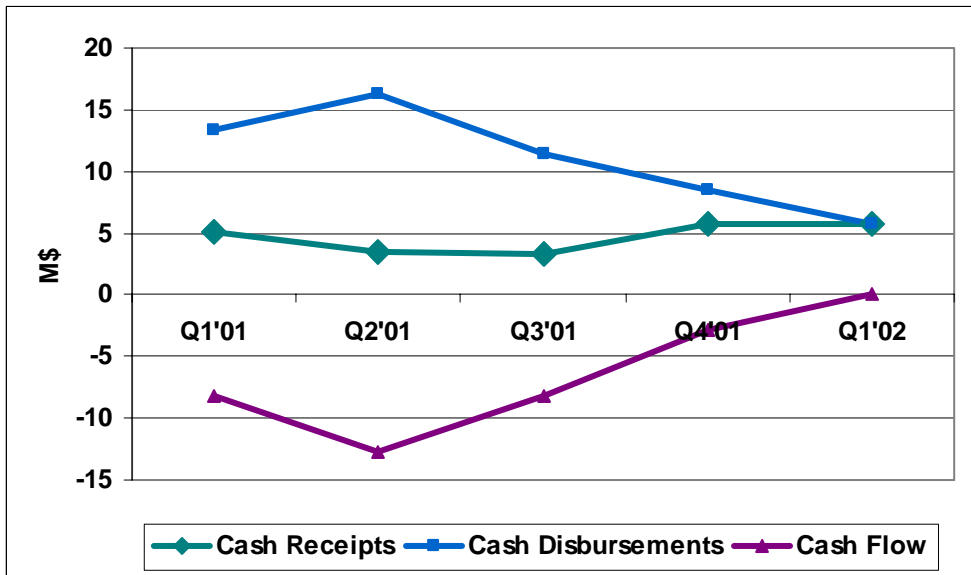


Also during 2002, VeriSign's Global Registry Services decided to add the RealNames product line to its distribution channel, adding as many as 80 additional sales outlets directly, and thousands of others indirectly, to the RealNames channel.

This has reflected itself in growing sales volumes across all products, and an increasing percentage of sales are being made by channel sales outlets. There have been 3 successive quarters of Revenue growth



matched by three consecutive quarters of decreased spending, resulting in break-even cash flows in Q1 2002 (excluding expense of Microsoft note, due May 2, 2002)



It has become clear that RealNames has a long-term, viable business model that is poised for significant growth over the coming quarters. The primary challenge facing RealNames now is to increase the growth rate of the business based on the technical and sales infrastructure already in place. It is at this point that the company has run into a crisis with its largest shareholder and most important partner – Microsoft.

Microsoft's Internet Explorer browser has almost 95% distribution to Internet users worldwide. In April 2000, Microsoft and RealNames agreed to embed the RealNames

naming technology into the browser, thus opening up the naming layer to all users worldwide. In practice it took until October 2000 to get RealNames live in the 24 countries in which Microsoft's MSN has a presence. Since then over 20 additional countries have also gone live, including China, Israel, Poland, Russia and 17 Latin American territories.

The terms of the RealNames agreement were that Microsoft would get ownership of 20% of RealNames, and additionally a \$40m cash guarantee, against a 15% revenue share of Keyword sales. \$15m would be paid between April 2000 and March 2001 and a further \$25m in the second 12 months.

In 2001 Microsoft agreed that the \$25m due in the second year of the contract would be converted into a note for Senior Debt, due on 2 May 2002. In March 2002 Microsoft further agreed to extend the termination date of the distribution contract to June 28 2002, from 12 March 2002.

The management team of RealNames believes that the business has accomplished all of the difficult tasks required of it to build a viable long-term business. We are fully committed to reap the rewards of that effort but require a long-term extension of our distribution agreement with Microsoft.

This memorandum will outline the industry, RealNames position in it, and the opportunity as we understand it. We will also provide 3 year projections based on about 15 months of real revenue experience with our channel.

RealNames Technology

There is now a consensus in both the policy and the technical communities that there needs to be a new layer of naming on top of the DNS.

The IETF has concluded that the demand for native language web addresses cannot be satisfied by building the capability into the 7-bit ASCII architecture of DNS. Instead, it will be necessary to build a new layer on top of the DNS. This has led to the formation of a group led by former Internet Architecture Board chair John Klensin, working to define such a layer. RealNames CTO, Nico Popp, and Chief Architect, Yves Arrouye, are part of this group.

ICANN has echoed the IETF findings and gave more than 2 hours of its recent conference in Ghana to a discussion of “*Higher Layer Naming Systems: Keywords and other non-DNS identifiers*”.

The innovations in the area of distributed computing – particularly the GRID project and the Object File System concept – also give rise to the need for a naming resolution protocol for files stored in “the cloud”. This points to future enterprise products which are described below.

RealNames Technology has three core components. A name *Registration Service*, a name *Resolution Service* and a name *Lookup or Directory Service*. Together these form the essential ingredients for a layer of naming. The principle building block of the RealNames technology is a *Namespace*. The registration, resolution and lookup services are all predicated on being targeted at a specific namespace.

The RealNames Keyword system is a namespace of Keywords for the entire Internet built on top of this technology. Any application can bind to this namespace through an open API for the purposes of resolution and/or lookup. Any approved reseller can sell names into this namespace through an open registration platform and its API.

The primary requirements of the technology are scalability and fault tolerance. The system must be permanently available to all applications that choose to bind to it. To this end RealNames has built a truly scalable architecture, capable of wide geographical distribution on the Internet with no obvious upper limits in terms of scale.

RealNames Products & Partnerships

On top of our Platform we have built four products.

- Basic Keywords
- Keyword Plus
- Keyword Management Services
- Keyword Namespaces

Three more are about to be announced

- Keyword Lookup Services (already live)
- Keyword Messaging Services through ContactCard
 - Email
 - SMS
 - Call
 - Page
- Web Numbers

And one is a Microsoft-specific opportunity

- MSN Keywords
 - (i) Basic Keywords look like this [click]: [Heathrow Declaration](#). As you see, there is a frame above the web page. This product has an MSRP of \$50 per Keyword per country per year. It is available over 40 countries.
 - (ii) Keywords Plus look like this [click]: [IBM NetVista](#). As you see, no frame - just an immediate visit to the site. This product has a brand protection feature. We would ONLY sell it to IBM. There is a human process of approval that ensures this protection, and this makes the product more expensive – MSRP of \$500 per Keyword, per country, per year. Most well known companies buy this product. It is available in over 40 countries.
 - (iii) Keyword Management Services is a managed version of the above 2 products plus some special services for e-commerce customers who are prepared to pay on a revenue share basis. The Keyword "IBM Thinkpad" and many others operate under this model. RealNames appoints an account manager to manage the Keywords appropriate to a customer. Some of these are paid-for Keywords and some are based on a revenue share from sales that are driven through the use of the Keyword. This is a service-based sale and is therefore variable customer by customer. Only RealNames is authorized to sell this product.
 - (iv) Keyword Namespace's work like this: "IBM anything". To test this just type "IBM buy a notebook", or anything else you want to try, so long as it starts with "IBM", into the browser. This product has an MSRP of \$70,000 per

country per year plus an initial \$80,000 implementation fee. It is available in over 40 countries.


- (v) Keyword Lookup Services are simply the directory service for Keywords. They look like this: "[Lookup IBM Thinkpad](#)" if no country code is added or like this "[Lookup UK IBM Thinkpad](#)". This beta site will go live later this year with paid for submission services allowing a customer to get prominence based on payment.
- (vi) Keyword Messaging Services through ContactCard are explained in greater detail in a separate document.
- (vii) The anticipated launch for Web Numbers is also Q4 2002. With this product, you can type "1" and it will point to the web site that has bought the number 1, or "777" for the site that has bought 777. Again pricing is yet to be determined.
- (viii) MSN Keywords are explained in greater detail in a separate document.

Country Registries: The above products are all subsumed under an architecture that is simultaneously local and global. Each of the 244 territories defined in the DNS has its own RealNames namespace. As a result RealNames has been able to delegate the authority for a country-specific namespace to a single partner. This partner runs the "country registry" for the Keyword product family. The partner pays RealNames both a license fee for the exclusive rights to the territory, and a wholesale fee for the products sold. All fees are annual and recurring. So far licenses to run the registry have been negotiated in China, Japan, France, Spain, Sweden, Denmark, Norway, Finland, Holland, Belgium, India and Israel. These deals are typically for 5 years with payments ramping through the term of the deals.

RealNames Business Metrics & Performance

User behavior

In January 2001 RealNames processed about 40m uses of a Keyword (a resolution in our language). This measures when a Keyword that is in our system is typed by a user into the browser and our system delivers the appropriate web page to them. In February 2002 we measured 163m resolutions. We delivered more than 450m resolutions of a Keyword in Q1 2002, representing 33% annualized growth.



Keywords Provide High Quality Visitors

It seems that Keyword visitors also convert to buyers at a great rate. See here: http://www.emarketer.com/analysis/marketing/cs_realnames.html?ref=ed


Average Sales Orders for RealNames Keywords and Forrester Online Retail Index, 2001

Small-ticket Items:	Forrester Online Retail Index	RealNames	% Difference RealNames/Forrester
Software	\$40.87	\$89.62	119.26%
Books	\$34.02	\$39.53	16.22%
Music	\$30.89	\$29.57	-4.28%
Videos	\$34.39	\$31.49	-8.44%
Office supplies	\$56.37	\$92.52	64.11%
Apparel	\$70.58	\$91.56	29.71%
Footwear	\$61.29	\$62.15	1.39%
Jewelry	\$67.78	\$203.73	200.57%
Flowers	\$51.51	\$50.13	-2.68%
Linens/home décor	\$62.19	\$114.84	84.67%
Health and beauty	\$43.81	\$60.23	37.48%
Small appliances	\$62.89	\$97.66	55.29%
Toys/videogames	\$47.59	\$54.42	14.34%
Sporting goods	\$80.66	\$99.96	23.92%
Tools and hardware	\$73.07	\$96.78	32.45%
Garden supplies	\$45.66	\$47.55	4.12%
Big-ticket items:			
Computer hardware	\$178.93	\$354.12	97.91%
Consumer electronics	\$141.97	\$232.56	63.81%
Appliances	\$328.16	\$359.40	9.52%
Furniture	\$200.51	\$353.98	76.54%
Food/beverages	\$80.43	\$261.72	225.42%
Air tickets	\$327.58	\$381.60	16.49%
Car rental	\$191.94	\$73.74	-61.58%
Hotel reservations	\$224.12	\$293.88	31.12%
Other	\$114.35	\$83.74	-26.77%

Source: RealNames, Forrester Research, 2002


Strong, Consistent Sales Growth

Paid Keyword sales have grown consistently over a 12-month period and are now at about 500 per day and the infrastructure is in place to drive this far higher. Cumulative registrations of paid Keywords into the RealNames System since June of 1998 are over 900,000.




The sources of Keyword sales have also become far more diversified as volumes have grown. In the three months Sept – Nov of 2001, Asia made up just 9% of sales, the Americas 45% and EMEA 46%. In the following 3 months, sales increased in China, and North America grew slightly, whilst Europe slowed. As a result, the distribution of Keyword sales were as follows: Asia 36%, the Americas 32% and EMEA 36%. This balance is likely to continue, although China looks set to continue its relatively strong growth.

September'01 - November'01



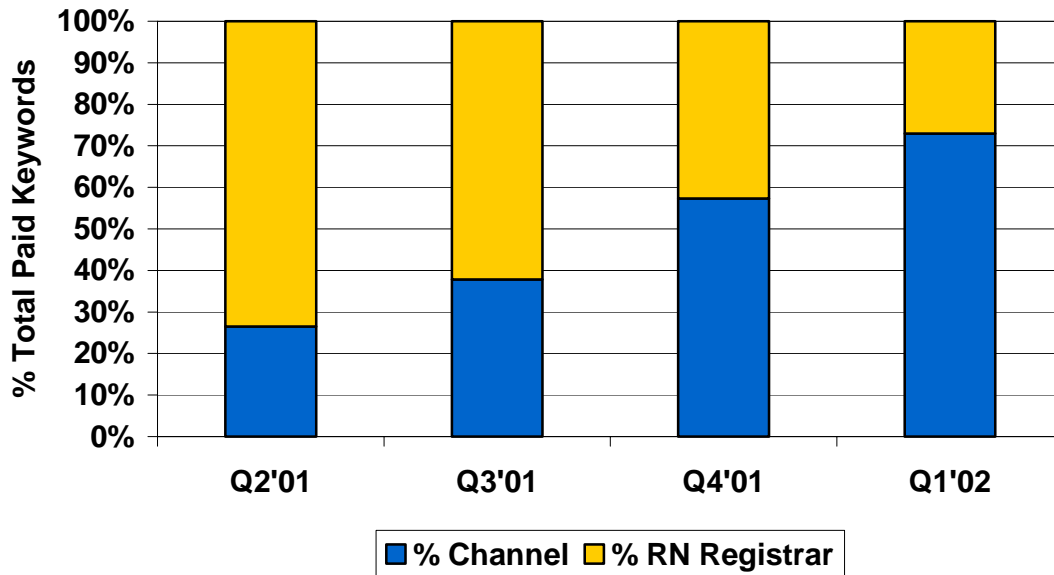
Keyword Units 31,388

December'01 - February'02



Keyword Units 48,097

The contribution of the Channel (as opposed to sales by RealNames directly through our own Registrar) grew systematically through the year. From 25% of sales in Q2 2001 to more than 70% by Q1 2002.



The company has gone through an extensive period of transformation, amid a significant "rightsizing" exercise. Headcount peaked at 330 employees during 1999 and has now settled at about 80. This is reflected in a significant reduction to burn rate, and a cash break-even quarter in Q1 2002. This does not take into account expenses associated with the previously mentioned note to Microsoft.

RealNames Corporation

Cash Flow

In Thousands

	Actual				Actual
	Q1'01	Q2'01	Q3'01	Q4'01	Q1'02
Beginning Balance	43,835	35,623	22,868	14,677	11,885
Source of Cash					
Collection	4,478	3,197	2,952	5,507	5,493
Miscellaneous Income	83	-	136	157	103
Interest Income	530	306	151	76	75
Total Source of Cash	5,091	3,503	3,239	5,740	5,671
Use of Cash					
Employee Related	5,575	4,638	4,322	3,760	3,021
Severance	-	600	100	1,085	-
Marketing	1,050	581	390	270	132
Facilities	1,453	1,070	1,250	1,146	1,018
Equipment (excl. depr)	1,319	934	1,008	624	575
Professional Services	1,803	433	373	260	291
Distribution Fees/Integration	145	6,805	2,660	551	-
Overseas	1,471	980	888	749	600
Misc	487	217	439	87	40
Total Use Of Cash	13,303	16,258	11,430	8,532	5,677
Cash Burn	\$ (8,212)	\$ (12,755)	\$ (8,191)	\$ (2,792)	\$ (6)
Ending Balance	\$ 35,623	\$ 22,868	\$ 14,677	\$ 11,885	\$ 11,879

There have also been three successive quarters of revenue growth.

RealNames

In Thousands

Revenue Detail

<i>(in thousands)</i> Channel	<u>Actual</u>		<u>Actual</u>
	<u>Q3-01</u>	<u>Q4-01</u>	<u>Q1-02</u>
Registries Rights & Services Fees	\$19	\$405	\$398
Keywords	49	122	279
VeriSign	0	0	0
Total Channel	<u>68</u>	<u>527</u>	<u>677</u>
RN Registrar			
Corporate Accounts	1,107	978	778
RN Website	466	597	719
Keyword Management Services	<u>1,160</u>	<u>1,698</u>	<u>2,574</u>
Total RealNames Registrar	<u>2,733</u>	<u>3,272</u>	<u>4,071</u>
Other/Resolution Services	<u>183</u>	<u>107</u>	<u>108</u>
Total RealNames	<u>\$ 2,984</u>	<u>\$ 3,906</u>	<u>\$ 4,856</u>

RealNames Corporation
Forecast with Eureka Included

Recognized Revenue <i>(in thousands)</i>	FY'01 Actual		FY02 Forecast					FY'03 Model					FY'04 Model				
	Q3-01	Q4-01	Q1-02	Q2-02	Q3-02	Q4-02	FY-02	Q1-03	Q2-03	Q3-03	Q4-03	FY-03	Q1-04	Q2-04	Q3-04	Q4-04	FY-04
Wholesale																	
Registries Rights & Services Fees	\$19	\$405	\$399	\$560	\$269	\$545	\$1,773	\$1,559	\$1,495	\$1,507	\$1,524	\$6,085	\$2,255	\$2,493	\$2,505	\$2,518	\$9,771
Keywords	49	122	268	561	841	1,104	2,774	1,202	1,438	1,769	2,329	6,738	2,590	2,883	3,213	3,584	12,270
Eureka Keywords	-	-	-	-	-	152	152	461	481	514	582	2,038	686	724	766	797	2,973
ccTLD's	-	-	-	-	-	50	50	153	161	174	199	687	236	251	267	283	1,037
Keyword Namespaces	-	-	-	-	-	7	7	30	63	107	150	350	195	254	330	428	1,207
	68	527	667	1,121	1,110	1,858	4,756	3,405	3,638	4,071	4,784	15,898	5,962	6,605	7,081	7,610	27,258
Retail																	
Keywords	1,573	1,577	1,497	1,365	1,450	1,580	5,892	1,868	2,241	2,689	3,227	10,025	2,463	2,817	3,227	3,705	12,212
Keyword Management Services	1,160	1,698	2,574	1,600	1,800	2,100	8,074	3,000	1,900	2,100	2,500	9,500	3,500	2,200	2,500	2,900	11,100
	2,733	3,275	4,071	2,965	3,250	3,680	13,966	4,868	4,141	4,789	5,727	19,525	5,963	5,017	5,727	6,605	23,312
Other	183	107	108	81	66	82	337	60	75	92	96	323	115	138	166	199	618
Challenge	-	-	-	900	450	-	1,350	-	-	-	-	-	-	-	-	-	-
Total RealNames	\$2,984	\$3,909	\$4,846	\$5,067	\$4,876	\$5,620	\$20,409	\$8,333	\$7,854	\$8,952	\$10,607	\$35,746	\$12,040	\$11,760	\$12,974	\$14,414	\$51,188

Cash Receipts <i>(in thousands)</i>	FY02 Forecast							FY'03 Model					FY'04 Model				
	Q3-01	Q4-01	Q1-02	Q2-02	Q3-02	Q4-02	FY-02	Q1-03	Q2-03	Q3-03	Q4-03	FY-03	Q1-04	Q2-04	Q3-04	Q4-04	FY-04
Wholesale																	
Registries Rights & Services Fees	\$39	\$1,691	\$586	\$1,118	\$865	\$1,171	\$3,740	\$1,499	\$2,586	\$2,672	\$2,672	\$9,429	\$3,002	\$4,027	\$4,117	\$4,117	\$15,263
Keywords	89	283	489	811	1,090	1,437	3,827	1,630	1,968	2,317	3,059	8,974	3,885	4,946	6,307	8,059	23,197
Eureka Keywords	-	-	-	-	-	608	608	1,235	82	133	879	2,329	1,649	234	299	1,003	3,185
ccTLD's	-	-	-	-	-	202	202	411	32	52	301	796	559	92	117	363	1,131
Keyword Namespaces	-	-	-	-	-	80	80	120	160	200	200	680	260	338	439	571	1,608
	128	1,974	1,075	1,929	1,955	3,498	8,457	4,895	4,828	5,374	7,111	22,208	9,355	9,637	11,279	14,113	44,384
Retail																	
Keywords	1,315	1,826	1,678	1,556	1,655	1,985	6,874	2,382	2,859	3,431	4,018	12,690	4,420	4,862	5,348	5,883	20,513
Keyword Management Services	1,161	1,698	2,574	1,600	1,800	2,100	8,074	3,000	1,900	2,100	2,500	9,500	3,500	2,200	2,500	2,900	11,100
	2,476	3,524	4,252	3,156	3,455	4,085	14,948	5,382	4,759	5,531	6,518	22,190	7,920	7,062	7,848	8,783	31,613
Other	480	108	336	77	106	136	655	141	152	164	178	635	214	256	308	369	1,147
Challenge	-	-	-	900	450	-	1,350	-	-	-	-	-	-	-	-	-	-
Total RealNames	\$3,084	\$5,606	\$5,663	\$6,062	\$5,966	\$7,719	\$25,410	\$10,418	\$9,739	\$11,069	\$13,807	\$45,033	\$17,489	\$16,955	\$19,435	\$23,265	\$77,144

User Experience Overview and Recommendations

Michael Arrington
650-486-5770
marrington@realnames.com

Executive Summary

The decision last year to divide Keywords into *Basic Keywords* (only available for terms that have 10 or less searches per month on MSN) and *Keywords Plus* (which must be approved by a human being applying rules of quality control as defined in the agreement) was intended to create an automated basis for the solving of any issues around data quality and user experience. In line with these changes and to further improve the user experience, *Keyword Namespaces* were also required to pass review and approval by a human being, and Basic Keywords were removed from search results.

Data obtained both before and after these changes were effected have given us excellent vision into the overall effectiveness of current Keyword policies. Ensuring a good Keyword user experience begins during the initial registration attempt, when the potential subscriber first checks for Keyword availability through the RealNames Subscription Filter Database. Additional quality assurance measures include a review of the prior day's subscription data to ensure proper filter functionality before Keyword data is uploaded to MSN. The Keywords Plus Application and Review process, a practical assessment of each term submitted for subscription as Keywords Plus based on user expectation, uniqueness, and the Registrant's asserted rights in the requested Keyword provides further assurance of quality Keyword data.

Through these efforts we have been able to minimize both the MSN requested Keyword revocations as well as RealNames initiated Keyword revocations.

Examination of the period beginning July 1, 2001 through March 31, 2002 reveals a 0.142% MSN requested revocation rate for all paid Keywords newly subscribed or renewed during this same period. Revocations during this period were primarily Standard Keywords (92%), followed by Basic Keywords and Keywords Plus – both at 4% each of total MSN requested revocations. (See product description table below). MSN requested revocation rates since the launch of the Basic Keyword and Keyword Plus products in August of 2001 demonstrate the effect of further quality improvements with revocation rates of 0.011% and 0.012%, respectively. The primary reason for revocation was due to the generic nature of a Keyword (56%), followed by contrary user expectation (35%), and not authorized ownership (9%). Asia experienced the majority of revocations at 59%, followed by EMEA with 33% and the US with 8%.

In order to continue to improve the user experience and simultaneously reduce the use of Microsoft resources required for review, we propose a number of policy and process changes in this document that we feel will move us materially towards those goals.

Current Statistics

MSN Requested Keyword Revocations - July 2001 through March 2002

	USA	EMEA	Asia	Total Revocations	Total Subscriptions 07/01 - 03/02*	Revocations as % of Subscriptions
Basic Keywords						
• <i>Contrary User Expectation</i>	0	3	0	3		0.005%
• <i>Generic</i>	0	1	3	4		0.007%
• <i>Not Owner</i>	0	0	0	0		0.000%
Total Basic Keywords:	0	4	3	7	60,911	0.011%
Keywords Plus						
• <i>Contrary User Expectation</i>	0	1	1	2		0.003%
• <i>Generic Term</i>	0	6	0	6		0.009%
• <i>Not Owner</i>	0	0	0	0		0.000%
Total Keywords Plus:	0	7	1	8	67,361	0.012%
Standard Keywords						
• <i>Contrary User Expectation</i>	5	26	28	59		0.046%
• <i>Generic</i>	7	12	72	91		0.071%
• <i>Not Owner</i>	2	11	4	17		0.013%
Total Standard Keywords:	14	49	104	167	-	0.130%
All Keywords						
• <i>Contrary User Expectation</i>	5	30	29	64		0.050%
• <i>Generic</i>	7	19	75	101		0.079%
• <i>Not Owner</i>	2	11	4	17		0.013%
Total All Keywords:	14	60	108	182	128,272	0.142%

* Includes Renewals, RN Alliance & RN Contract Subscriptions

RealNames Keyword Products Overview

	Dates Available	Annual Subscription Price	Query Level	Reviewed	IE Browser	MSN Search	Keyword Window
Basic Keyword	Beginning 08/26/01	\$50	Under 10	No	Yes	No	Yes
Keywords Plus	Beginning 08/26/01	\$498	11 and above	Yes	Yes	Yes	No
Standard Keywords	02/01/01-08/26/01	\$50	none	No	Yes	Yes	No

Keyword Quality Assurance Recommendations

RealNames believes that ongoing quality assurance development and automation will further improve and ensure an excellent Keyword user experience and simultaneously reduce the review burden on Microsoft.

A summary of RealNames current recommendations for enhancing Keyword quality include: (1) executing a resolver lookup protocol which would eliminate the Keyword data upload and indexing processes; (2) development of an automated policing agent through the Keyword Management System (KMS) which allows MSN Editors to view rules based “exception” data as well as initiate corrective action for specific data; (3) expansion of the automated policing agent to allow for Keyword review and an automated scoring process; (4) expansion of the subscription filter; (5) confirmation of Microsoft’s normalization rules; and (6) establishment of acceptable standards and

ongoing metrics review by RealNames and Microsoft Editorial teams. Each of these recommendations is described in detail below.

Resolver

RealNames will implement a resolver lookup protocol. Implementation of this protocol will allow Microsoft to query the resolver in real time (better than 50 milliseconds) and ask the question “Do you have a Keyword for X?,” where “X” is the user input and associated context like country and language. If the answer is “yes” the query will be passed to RealNames for resolution as a Keyword. If the answer is “no” the query will not be passed to RealNames. The advantages of utilizing the resolver lookup protocol include elimination of the data upload (delivery) and indexing (processing) and the savings for Microsoft on both technical and staff resources that support the current processes.

RealNames already has this technology. It was deployed (and scaled) successfully with both Google and AltaVista previously in real time. Microsoft would need to add code to the search back end, invoking a lookup to the RealNames Keyword Service. This is a relatively minor change.

Automated Policing Agent: Tools Access & Development

An unintended consequence of the solution above (implementation of a resolver lookup) is that the editorial team at Microsoft will not have daily access to RealNames data. To fix this, and allow Microsoft to embrace RealNames data policies, RealNames recommend the development of an automated policing agent that will throw out a daily “exception” file.

This policing agent can be rules based and can be live within 30 days. As a minimum it will look at all Basic Keywords and be sure that there are less than 10 searches on MSN Search in the most recent month. Additionally it will look at all previously sold Keywords and ask if there is a spike in usage that suggests a change has occurred in the meaning of the Keyword, necessitating a human review. It could also implement any other rules Microsoft feels appropriate within the terms of the contract.

In addition to creating a daily exception file, RealNames proposes that they deploy a Keyword Management Services (KMS) screen exclusively for Microsoft editors, allowing access to Keywords sold and allowing (within the limits of the agreement) for the rescinding of Keywords by Microsoft editors.

This screen could be used worldwide and allow MSN worldwide to police the namespace that affects their territory. Note that RealNames already has a KMS product that is used by internal RealNames editors and the RealNames sales channel. Within 30 days a new module specifically for Microsoft’s use can be deployed. The potential for a trusted third-party to host this policing agent activity is another option for consideration.

Development of additional tools will provide an automated process for MSN Editors to send questions, comments and requests to RealNames Editors related to specific data in

our system, as well as provide reporting capabilities for analysis of data in the filters and subscription queue.

Automated Policing Agent: Keyword Review & Scoring Process

This initiative is focused on the development of an automated scoring process that would quantify and qualify terms submitted for Keyword review on the basis of user expectation, uniqueness, and asserted rights. Results of this scoring process would be logged in KMS and used in conjunction with current processes to further ensure accurate and consistent quality review results. The Keyword Review Score would be available through KMS.

Filter Expansion

RealNames is currently expanding our subscription filters on a country-specific basis through review and categorization of additional generic terms and popular brand names. This expansion will create more robust localized functionality, further ensuring proper routing of terms through the filter to the appropriate product designation (Basic Keyword or Keywords Plus) as well as identifying terms that are not suitable for subscription as Keywords.

MSFT Normalization Rules

Close cooperation with Microsoft to maintain and update normalization rules would allow RealNames to fine-tune its subscription filters. The Unicode based nature of RealNames' Keywords allows for a vast number of special characters to be displayed that are normalized out of query strings on the Microsoft side. This creates situations in which terms proceed through the filter when they shouldn't, e.g. "sex†" - which due to normalization rules will display as "sex", a term that would not be allowed through the filter. A comprehensive dialog on Microsoft's normalization rules would enable RealNames to update filter functionality to capture and prevent this type of situation. While many normalization issues have been resolved, opportunities remain for further alignment of the RealNames filter with MSN's front and back end systems that would contribute to our quality initiatives.

Quarterly Performance Reviews of RealNames by MSN Team

Ongoing review and discussion of the quality of the Keyword user experience is a critical factor in the successful development and implementation of process improvements and initiatives. RealNames would like to propose a quarterly meeting in Redmond between our editors to work collaboratively to ensure the highest quality Keyword user experience. Through definition and communication of a clear issues resolution path and the establishment of acceptable quality standards and ongoing metrics measurement we are confident in our ability to scale and further improve our current processes.

- Data Issues Resolution Path

RealNames has demonstrated its desire and ability to respond quickly to issues presented by MSN Editorial. A clearly identified communications path for MSFT personnel to raise questions and concerns related to Keyword data would create an

added level of efficiency that would further enable RealNames to respond and resolve these types of issues quickly, minimizing their impact on the user experience.

- **Agreed Upon Performance Metrics**
By establishing acceptable quality standards and metrics for their measurement, RealNames and MSN can work together to achieve the quality standards required of the Keyword user experience.

RealNames Incremental Business Opportunities

Griffin Golamco
VP, Corporate Development
650-486-5639
ggolamco@realnames.com

Executive Summary

Project Eureka is a promotional Program (Program) designed by RealNames to provide a catalyst for exponential growth in Keywords sales, to rapidly increase consumer awareness of Keywords and to expand the channel of Keywords resellers. The project is targeted at domain name buyers and owners: For each .com, .net and .org domain name (currently totaling approximately 28 million domain names according to SnapNames data as of 03/31/02), RealNames will generate multiple, pre-approved promotional Keywords. These Keywords will be activated in the RealNames System, free of charge, for a limited period of time, during which RealNames will provide significant incentives for its channel to sell these Keywords to the domain name owners.

The Program will be implemented through RealNames' network of Registrars, including those represented through our relationship with VeriSign's Global Registry Services division (VGRS). To participate in the Program, Registrars must meet minimum marketing requirements, including, but not limited to, an e-mail marketing campaign to their entire installed base of customers. In addition, RealNames will utilize all of its marketing resources, including the significant exposure obtained by its Testbed Keywords, to promote the Program. Finally, VGRS will also dedicate significant resources to the promotion of the Program.

Initially, the Program will be implemented in the RealNames Namespace for Keywords in the United States, and will be limited to the generic top-level domains (gTLDs) currently administered by VeriSign. However, if successful, the Program will be rolled out to other gTLDs, as well as to the ccTLDs, and will also be expanded to include other Namespaces and character sets.

Basic Keywords Program Details

Keyword Generation

RealNames will generate multiple Basic Keywords for each .com, .net, .org domain name. These promotional Basic Keywords will be generated using an automated system that will include techniques such as word splitting (e.g., “Joes Pizza” for www.joespizza.com), replacing terms with symbols (e.g., “&” instead of “and”), as well as the use of thick registry data which may allow us to add location extensions (e.g., Joe’s Pizza of San Francisco).

Keyword Activation

These Keywords will be activated free of charge in the RealNames System for a period of 90 days, starting on June 15, 2002, and continuing through September 15, 2002. The promotional Keywords will resolve in Microsoft browsers set to the US-English namespace and will display a Keyword window, similar to the one shown with RealNames’ current Testbed Keywords, but with specific messaging regarding the Program.

This activation will allow potential customers to try their Keywords in the browser, risk-free. All Keywords in the Program that have not been registered by September 15, 2002, will be de-activated in the RealNames System.

Promotional Pricing

The Promotion will start on June 15, 2002, and will end on September 15, 2002. During the 90-day promotional period, the owner of each domain name will be able to purchase the promotional Keywords at a significant discount: US\$9.95 for a one-year subscription for each Keyword, as opposed to the regular price of US\$49. The term of the subscription for each Basic Keyword is one year, beginning upon registration.

Each registered Basic Keyword will operate just like any other Basic Keyword in the RealNames System.

Registrar Participants – Eligibility and Requirements

Eligibility

Any current Registrar of RealNames for U.S. Keywords and all ICANN-Accredited Registrars will be eligible to participate in the Program.

Participating Registrars - Registrars who have successfully joined the Program by June 15, 2002 (Participating Registrars), will have the exclusive right to offer the promotional

pricing for the promotional Keywords during the Program. Participating ICANN-Accredited Registrars will be given the promotional Keyword data associated with the domain names under their management. Only participating Registrars, ICANN-Accredited or otherwise, will have access to the Open Pool, as described below. Registrars may not join the Program after June 15, 2002. As a result, there is a significant incentive for Registrars to join the Program prior to June 15.

Non-Participating Registrars – Registrars who have not signed up to participate in the Program by June 15, 2002 (Non-Participating Registrars), will not have the right to offer the promotional pricing for the promotional Keywords; however, they may sell the promotional Keywords at the full, non-promotional price. Additionally, the promotional Keywords generated for domain names registered by non-participating, ICANN-Accredited Registrars will be put into an “Open Pool” of Keywords, as described below.

It is important to note that all of the Keywords not activated for the Program will be available for registration by any Registrar regardless of whether or not they participate in the Program, at the full, non-promotional registration price. Participation in the Program is required only to be able to offer the promotional pricing for Keywords.

Open Pool – Certain data, including, without limitation, promotional Keywords and the corresponding domain names under management by Non-Participating ICANN-Accredited Registrars, will be included in an “Open Pool” available to Participating Registrars in good standing.

Requirements

Initial Requirements. In order to participate in the Program, a Registrar must agree to the following:

- An e-mail campaign to the Registrar’s entire existing customer base with messaging to be provided by RealNames. There will be an initial e-mail sent by June 30, 2002, a follow-up e-mail sent by July 15, 2002 and a final e-mail sent prior to the end of the Promotion term.
- For Registrars who sell through a website, the promotional Program must receive significant visibility on the site, according to guidelines to be set forth by RealNames.

Each participating Registrar must meet certain performance metrics regarding promotional Keywords sold during the course of the Program.

Ongoing Requirements. In order to preserve the rights and benefits of participation, Participating Registrars must continue to meet specific milestones throughout the Program. In the event an ICANN-Accredited Registrar fails to meet these milestones, those promotional Keywords associated with their domain names will be added to the Open Pool.

What RealNames Provides

RealNames will generate all of the promotional Keywords and will maintain them in the RealNames System. RealNames will also administer and coordinate the Program in conjunction with VeriSign. Through the use of the Keyword Frame on both the promotional and the Testbed Keywords, RealNames will provide significant marketing exposure for the Program. For example, in Q1 of 2002, RealNames' Testbed Keywords received 331 million page views. The Keyword Frame, therefore, represents one of the most powerful marketing and promotional tools available on the Internet today. Finally, RealNames will provide all of the technical tools and partner support required for the Program. This will include full messaging suggestions, FAQs and, where necessary, direct support.

Ongoing Activity & Future Expansion

On October 1, 2002, RealNames will generate multiple, pre-approved Keywords for each new domain name registered, and on January 1, 2003, we will expand the promotion to renewing domain names. These Keywords will be provided to the Registrar of record to offer to the owner of the domain name at the promotional pricing for a 30-day trial period.

All Keywords registered during the promotion will renew at standard prices.

In Phases Three through Five, RealNames will expand the Program to cover other TLDs, languages and territories. Probable areas of expansion include:

- Providing a similar Program for the other gTLDs, as well as the ccTLDs
- Providing a similar Program for VeriSign's Internationalized Domain Names (IDNs)
- Offering multi-lingual, multi-character set promotional Keywords for the appropriate ccTLDs

Various aspects of Project Eureka will be individually tailored to accommodate the different market environments, but the overall concept will remain the same.

Keyword Plus Program

In addition to generating Basic Keywords for .com, .net and .org domain names, RealNames will also generate promotional Keywords Plus. Although the initial process for generating promotional Keywords Plus will be similar to the process described above for promotional Basic Keywords, RealNames will perform a more stringent approach to ensure quality Keywords with the appropriate user experience. Many of the Keywords Plus that will be offered through Project Eureka will be RealNames' Testbed Keywords,

all of which are approved Keywords Plus and all of which are already tied to a specific domain name.

The Keywords Plus Program will mirror the Basic Keywords Program. The promotional price for a Keywords Plus will be US\$99, as opposed to the standard price of US\$498. Participating Registrars will receive the same exclusivity regarding promotional pricing and access to the Open Pool as described above.

Domain name owners will be offered either promotional Basic Keywords or promotional Keywords Plus, but not both. Registrars that agree to participate in the Programs will be eligible to sell both Basic Keywords and Keywords Plus at promotional prices.

Expected Effect of Programs

RealNames expects that Project Eureka will generate a significant increase in awareness of Keywords. The promotional requirements for participants, in conjunction with RealNames' own marketing efforts, will greatly increase the visibility of Keywords, both to Registrars and to consumers. In addition, the increased sales of Keywords that RealNames anticipates as a result of the Program will help solidify the naming industry's impression of Keywords as the de facto standard of next generation naming.

RealNames also anticipates significant revenue to be generated from the Program, both from the promotional sales of Keywords, as well as from continued increased sales from the Program has finished as a result of the increased awareness of Keywords.

Project Eureka Technical Requirements

Project Eureka's Phase One will generate up to 140 million Keywords. These Keywords will be active and resolve in the English-United States (EN-US) namespace for a 30 – 90 day promotional period (duration TBD).

RealNames anticipates that Project Eureka will lengthen data import/export times and increase impact to the system components between MSN and RealNames.

RealNames assumes that the current implementation based on data transfer (FTP), import, and indexing from RealNames to MSN would be used for this project.

Project Eureka will be launched as follows:

Phase 1 Eureka US Kickoff (by June 15, 2002)

Duration 30 – 90 days

Total Current Keywords: 390,000

Additions: 121,800,000

Phase 2 Eureka US ongoing (October 2002)

Duration 30 – 90 days

Total Current Keywords: 122,190,000
Additions: 118,063,000

Phase 3 Eureka Worldwide Kickoff (January 2003)
Duration 30 – 90 days
Includes the addition of ccTLD's and IDN's
Total Current Keywords: 240,253,000
Additions: 48,572,000

Phase 4 Eureka Worldwide Ongoing (April 2003)
Duration 30 – 90 days
Includes the addition of ccTLD's and IDN's
Total Current Keywords: 288,825,800
Additions: 61,718,400

Phase 5 Eureka Worldwide Ongoing (August 2003)
Duration 30 – 90 days
Includes the addition of ccTLD's and IDN's
Total Current Keywords: 350,544,200
Additions: TBD

Phase 1:

Impact on System Components

- Estimated Disk space 8.5 gigabytes per day
- DP/upload time – negotiable

Impact on editorial review:

- There will be a significant impact for Microsoft on the current editorial/review process to process 140 million new Keywords.

Impact on the Resolution System

- The addition of 140 million new Keywords could increase the total re-directed traffic to RealNames from MSN by 100 million resolutions per month and elevate queries to 100 per second.

RealNames Corporation
Eureka - CNO Domains

	<u>Quarter 1</u>	<u>Quarter 2</u>	<u>Quarter 3</u>	<u>Quarter 4</u>	<u>Year 1</u>	<u>Quarter 5</u>	<u>Quarter 6</u>	<u>Quarter 7</u>	<u>Quarter 8</u>	<u>Year 2</u>	<u>Quarter 9</u>	<u>Quarter 10</u>	<u>Quarter 11</u>	<u>Quarter 12</u>	<u>Year 3</u>
Keyword Units															
Basic Keywords															
New	452,183	13,996	32,837	32,837	531,853	30,872	32,568	33,698	33,698	130,837	32,519	33,537	34,215	34,215	134,487
Renew	0	0	0	0	0	113,046	3,499	8,209	8,209	132,963	115,134	28,854	33,526	33,526	211,040
Total	452,183	13,996	32,837	32,837	531,853	143,918	36,067	41,908	41,908	263,800	147,654	62,390	67,741	67,741	345,527
As % of Total	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%
Keywords Plus															
New	1,827	57	133	133	2,149	125	132	136	136	529	131	136	138	138	543
Renew	0	0	0	0	0	457	14	33	33	537	465	117	135	135	853
Total	1,827	57	133	133	2,149	581	146	169	169	1,066	597	252	274	274	1,396
As % of Total	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
Total Keyword Units	454,010	14,053	32,970	32,970	534,002	144,499	36,213	42,077	42,077	264,866	148,250	62,643	68,015	68,015	346,923
Total Market (at Recommended Promotional Pricing)															
Bookings															
Basic Keywords	\$ 4,499,216	\$ 139,262	\$ 326,729	\$ 326,729	\$ 5,291,936	\$ 5,846,415	\$ 495,504	\$ 737,553	\$ 737,553	\$ 7,817,025	\$ 5,965,149	\$ 1,747,517	\$ 1,983,221	\$ 1,983,221	\$ 11,679,108
Keywords Plus	\$ 180,873	\$ 5,598	\$ 13,135	\$ 13,135	\$ 212,741	\$ 239,810	\$ 20,068	\$ 29,997	\$ 29,997	\$ 319,873	\$ 244,672	\$ 71,472	\$ 81,145	\$ 81,145	\$ 478,433
Total Bookings	\$ 4,680,089	\$ 144,860	\$ 339,864	\$ 339,864	\$ 5,504,677	\$ 6,086,225	\$ 515,571	\$ 767,551	\$ 767,551	\$ 8,136,897	\$ 6,209,821	\$ 1,818,988	\$ 2,064,365	\$ 2,064,365	\$ 12,157,541
VeriSign															
Bookings															
Basic Keywords	\$ 1,889,671	\$ 58,490	\$ 137,226	\$ 137,226	\$ 2,222,613	\$ 2,455,494	\$ 208,112	\$ 309,772	\$ 309,772	\$ 3,283,150	\$ 2,505,363	\$ 733,957	\$ 832,953	\$ 832,953	\$ 4,905,225
Keywords Plus	\$ 32,557	\$ 1,008	\$ 2,364	\$ 2,364	\$ 38,293	\$ 43,166	\$ 3,612	\$ 5,400	\$ 5,400	\$ 57,577	\$ 44,041	\$ 12,865	\$ 14,606	\$ 14,606	\$ 86,118
Total Bookings	\$ 1,922,228	\$ 59,498	\$ 139,590	\$ 139,590	\$ 2,260,906	\$ 2,498,660	\$ 211,724	\$ 315,172	\$ 315,172	\$ 3,340,727	\$ 2,549,404	\$ 746,822	\$ 847,559	\$ 847,559	\$ 4,991,343
Revenue															
Basic Keywords	\$ 157,473	\$ 477,292	\$ 498,476	\$ 532,782	\$ 1,666,022	\$ 602,674	\$ 709,446	\$ 748,893	\$ 792,030	\$ 2,853,044	\$ 824,479	\$ 876,611	\$ 1,008,314	\$ 1,139,110	\$ 3,848,514
Keywords Plus	\$ 2,713	\$ 8,223	\$ 8,588	\$ 9,179	\$ 28,704	\$ 10,455	\$ 12,440	\$ 13,130	\$ 13,888	\$ 49,913	\$ 14,459	\$ 15,376	\$ 17,693	\$ 19,995	\$ 67,523
Total Revenue	\$ 160,186	\$ 485,515	\$ 507,064	\$ 541,962	\$ 1,694,726	\$ 613,129	\$ 721,887	\$ 762,023	\$ 805,918	\$ 2,902,957	\$ 838,938	\$ 891,987	\$ 1,026,008	\$ 1,159,105	\$ 3,916,037
Non-VeriSign															
Bookings															
Basic Keywords	\$ 1,349,765	\$ 41,778	\$ 98,019	\$ 98,019	\$ 1,587,581	\$ 1,753,924	\$ 148,651	\$ 221,266	\$ 221,266	\$ 2,345,107	\$ 1,789,545	\$ 524,255	\$ 594,966	\$ 594,966	\$ 3,503,732
Keywords Plus	\$ 54,262	\$ 1,680	\$ 3,940	\$ 3,940	\$ 63,822	\$ 71,943	\$ 6,020	\$ 8,999	\$ 8,999	\$ 95,962	\$ 73,402	\$ 21,441	\$ 24,343	\$ 24,343	\$ 143,530
Total Bookings	\$ 1,404,027	\$ 43,458	\$ 101,959	\$ 101,959	\$ 1,651,403	\$ 1,825,868	\$ 154,671	\$ 230,265	\$ 230,265	\$ 2,441,069	\$ 1,862,946	\$ 545,697	\$ 619,310	\$ 619,310	\$ 3,647,262
Revenue															
Basic Keywords	\$ 112,480	\$ 340,923	\$ 356,054	\$ 380,559	\$ 1,190,016	\$ 430,481	\$ 506,747	\$ 534,924	\$ 565,736	\$ 2,037,888	\$ 588,913	\$ 626,150	\$ 720,225	\$ 813,650	\$ 2,748,938
Keywords Plus	\$ 4,522	\$ 13,705	\$ 14,314	\$ 15,299	\$ 47,840	\$ 17,425	\$ 20,734	\$ 21,883	\$ 23,147	\$ 83,189	\$ 24,098	\$ 25,627	\$ 29,489	\$ 33,325	\$ 112,539
Total Revenue	\$ 117,002	\$ 354,628	\$ 370,368	\$ 395,858	\$ 1,237,856	\$ 447,907	\$ 527,481	\$ 556,806	\$ 588,883	\$ 2,121,077	\$ 613,012	\$ 651,777	\$ 749,714	\$ 846,975	\$ 2,861,477
RealNames															
Bookings															
Basic Keywords	\$ 1,754,694	\$ 54,312	\$ 127,424	\$ 127,424	\$ 2,063,855	\$ 2,280,102	\$ 193,246	\$ 287,646	\$ 287,646	\$ 3,048,640	\$ 2,326,408	\$ 681,532	\$ 773,456	\$ 773,456	\$ 4,554,852
Keywords Plus	\$ 70,540	\$ 2,183	\$ 5,123	\$ 5,123	\$ 82,969	\$ 93,526	\$ 7,826	\$ 11,699	\$ 11,699	\$ 124,750	\$ 95,422	\$ 27,874	\$ 31,646	\$ 31,646	\$ 186,589
Total Bookings	\$ 1,825,235	\$ 56,495	\$ 132,547	\$ 132,547	\$ 2,146,824	\$ 2,373,628	\$ 201,073	\$ 299,345	\$ 299,345	\$ 3,173,390	\$ 2,421,830	\$ 709,406	\$ 805,103	\$ 805,103	\$ 4,741,441
Revenue															
Basic Keywords	\$ 146,225	\$ 443,200	\$ 462,870	\$ 494,726	\$ 1,547,021	\$ 559,626	\$ 658,772	\$ 695,401	\$ 735,456	\$ 2,649,255	\$ 765,587	\$ 813,996	\$ 936,292	\$ 1,057,745	\$ 3,573,620
Keywords Plus	\$ 5,878	\$ 17,817	\$ 18,608	\$ 19,888	\$ 62,192	\$ 22,653	\$ 26,954	\$ 28,447	\$ 30,092	\$ 108,146	\$ 31,328	\$ 33,314	\$ 38,336	\$ 43,323	\$ 146,301
Total Revenue	\$ 152,103	\$ 461,017	\$ 481,478	\$ 514,615	\$ 1,609,212	\$ 582,279	\$ 685,726	\$ 723,848	\$ 765,548	\$ 2,757,400	\$ 796,915	\$ 847,310	\$ 974,628	\$ 1,101,067	\$ 3,719,920
Cash	\$ 608,412	\$ 1,235,655	\$ 81,846	\$ 132,547	\$ 2,058,459	\$ 878,813	\$ 1,649,443	\$ 234,591	\$ 299,345	\$ 3,062,191	\$ 1,003,385	\$ 1,851,022	\$ 744,759	\$ 805,103	\$ 4,404,269

RealNames Corporation
Eureka - ccTLD

	Quarter 1	Quarter 2	Quarter 3	Quarter 4	Year 1	Quarter 5	Quarter 6	Quarter 7	Quarter 8	Year 2	Quarter 9	Quarter 10	Quarter 11	Quarter 12	Year 3
Keyword Units															
Basic Keywords															
New	11,235	11,235	11,235	11,237	44,943	11,237	11,237	11,237	0	33,711	0	0	0	0	0
Renew	0	0	0	0	0	32,484	1,191	2,815	2,815	39,306	34,976	9,941	11,240	11,240	67,398
Total	11,235	11,235	11,235	11,237	44,943	43,721	12,428	14,052	2,815	73,017	34,976	9,941	11,240	11,240	67,398
As % of Total	96%	100%	100%	100%	99%	100%	100%	100%	98%	100%	99%	99%	99%	99%	99%
Keywords Plus															
New	525	19	46	46	635	45	45	45	45	182	45	45	45	45	182
Renew	0	0	0	0	0	131	5	11	11	159	141	40	45	45	272
Total	525	19	46	46	635	177	50	57	57	340	187	86	91	91	454
As % of Total	4%	0%	0%	0%	1%	0%	0%	0%	2%	0%	1%	1%	1%	1%	1%
Total Keyword Units	11,760	11,255	11,281	11,282	45,578	43,898	12,478	14,109	2,872	73,357	35,162	10,027	11,331	11,331	67,852
Total Market (at Recommended Promotional Pricing)															
Bookings															
Basic Keywords	\$ 1,292,878	\$ 47,406	\$ 112,050	\$ 112,050	\$ 1,564,383	\$ 1,703,526	\$ 170,155	\$ 249,742	\$ 249,742	\$ 2,373,164	\$ 1,825,617	\$ 598,921	\$ 662,590	\$ 662,590	\$ 3,749,718
Keywords Plus	\$ 51,975	\$ 1,906	\$ 4,505	\$ 4,505	\$ 62,890	\$ 69,857	\$ 6,891	\$ 10,159	\$ 10,159	\$ 97,065	\$ 74,870	\$ 24,497	\$ 27,112	\$ 27,112	\$ 153,591
Total Bookings	\$ 1,344,853	\$ 49,311	\$ 116,554	\$ 116,554	\$ 1,627,273	\$ 1,773,382	\$ 177,046	\$ 259,900	\$ 259,900	\$ 2,470,229	\$ 1,900,487	\$ 623,418	\$ 689,702	\$ 689,702	\$ 3,903,309
VeriSign															
Bookings															
Basic Keywords	\$ 387,863	\$ 14,222	\$ 33,615	\$ 33,615	\$ 469,315	\$ 511,058	\$ 51,046	\$ 74,922	\$ 74,922	\$ 711,949	\$ 547,685	\$ 179,676	\$ 198,777	\$ 198,777	\$ 1,124,915
Keywords Plus	\$ 15,593	\$ 572	\$ 1,351	\$ 1,351	\$ 18,867	\$ 20,957	\$ 2,067	\$ 3,048	\$ 3,048	\$ 29,120	\$ 22,461	\$ 7,349	\$ 8,134	\$ 8,134	\$ 46,077
Total Bookings	\$ 403,456	\$ 14,793	\$ 34,966	\$ 34,966	\$ 488,182	\$ 532,015	\$ 53,114	\$ 77,970	\$ 77,970	\$ 741,069	\$ 570,146	\$ 187,025	\$ 206,911	\$ 206,911	\$ 1,170,993
Revenue															
Basic Keywords	\$ 32,322	\$ 98,151	\$ 103,323	\$ 111,726	\$ 345,522	\$ 127,595	\$ 151,196	\$ 160,776	\$ 171,103	\$ 610,669	\$ 181,040	\$ 197,863	\$ 229,623	\$ 260,586	\$ 869,112
Keywords Plus	\$ 1,299	\$ 3,946	\$ 4,154	\$ 4,492	\$ 13,890	\$ 5,164	\$ 6,182	\$ 6,573	\$ 6,997	\$ 24,916	\$ 7,405	\$ 8,096	\$ 9,400	\$ 10,672	\$ 35,573
Total Revenue	\$ 33,621	\$ 102,097	\$ 107,476	\$ 116,218	\$ 359,412	\$ 132,759	\$ 157,379	\$ 167,349	\$ 178,100	\$ 635,586	\$ 188,445	\$ 205,959	\$ 239,023	\$ 271,258	\$ 904,685
Non-VeriSign															
Bookings															
Basic Keywords	\$ 646,439	\$ 23,703	\$ 56,025	\$ 56,025	\$ 782,192	\$ 851,763	\$ 85,077	\$ 124,871	\$ 124,871	\$ 1,186,582	\$ 912,809	\$ 299,460	\$ 331,295	\$ 331,295	\$ 1,874,859
Keywords Plus	\$ 25,988	\$ 953	\$ 2,252	\$ 2,252	\$ 31,445	\$ 34,928	\$ 3,445	\$ 5,079	\$ 5,079	\$ 48,533	\$ 37,435	\$ 12,249	\$ 13,556	\$ 13,556	\$ 76,796
Total Bookings	\$ 672,427	\$ 24,656	\$ 58,277	\$ 58,277	\$ 813,637	\$ 886,691	\$ 88,523	\$ 129,950	\$ 129,950	\$ 1,235,114	\$ 950,244	\$ 311,709	\$ 344,851	\$ 344,851	\$ 1,951,655
Revenue															
Basic Keywords	\$ 53,870	\$ 163,585	\$ 172,204	\$ 186,210	\$ 575,870	\$ 212,658	\$ 251,993	\$ 267,960	\$ 285,171	\$ 1,017,782	\$ 301,733	\$ 329,772	\$ 382,705	\$ 434,311	\$ 1,448,520
Keywords Plus	\$ 2,166	\$ 6,576	\$ 6,923	\$ 7,486	\$ 23,151	\$ 8,606	\$ 10,304	\$ 10,955	\$ 11,662	\$ 41,527	\$ 12,342	\$ 13,493	\$ 15,667	\$ 17,786	\$ 59,289
Total Revenue	\$ 56,036	\$ 170,161	\$ 179,127	\$ 193,696	\$ 599,020	\$ 221,265	\$ 262,298	\$ 278,915	\$ 296,833	\$ 1,059,310	\$ 314,075	\$ 343,266	\$ 398,372	\$ 452,097	\$ 1,507,809
RealNames															
Bookings															
Basic Keywords	\$ 581,795	\$ 21,333	\$ 50,422	\$ 50,422	\$ 703,972	\$ 766,587	\$ 76,570	\$ 112,384	\$ 112,384	\$ 1,067,924	\$ 821,528	\$ 269,514	\$ 298,166	\$ 298,166	\$ 1,687,373
Keywords Plus	\$ 23,389	\$ 858	\$ 2,027	\$ 2,027	\$ 28,300	\$ 31,435	\$ 3,101	\$ 4,571	\$ 4,571	\$ 43,679	\$ 33,692	\$ 11,024	\$ 12,200	\$ 12,200	\$ 69,116
Total Bookings	\$ 605,184	\$ 22,190	\$ 52,449	\$ 52,449	\$ 732,273	\$ 798,022	\$ 79,671	\$ 116,955	\$ 116,955	\$ 1,111,603	\$ 855,219	\$ 280,538	\$ 310,366	\$ 310,366	\$ 1,756,489
Revenue															
Basic Keywords	\$ 48,483	\$ 147,227	\$ 154,984	\$ 167,589	\$ 518,283	\$ 191,392	\$ 226,794	\$ 241,164	\$ 256,654	\$ 916,004	\$ 271,559	\$ 296,795	\$ 344,434	\$ 390,880	\$ 1,303,668
Keywords Plus	\$ 1,949	\$ 5,919	\$ 6,231	\$ 6,737	\$ 20,835	\$ 7,746	\$ 9,274	\$ 9,860	\$ 10,496	\$ 37,375	\$ 11,108	\$ 12,144	\$ 14,100	\$ 16,008	\$ 53,360
Total Revenue	\$ 50,432	\$ 153,145	\$ 161,214	\$ 174,327	\$ 539,118	\$ 199,138	\$ 236,068	\$ 251,023	\$ 267,150	\$ 953,379	\$ 282,667	\$ 308,939	\$ 358,535	\$ 406,887	\$ 1,357,028
Cash	\$ 201,728	\$ 410,853	\$ 32,277	\$ 52,449	\$ 697,307	\$ 300,974	\$ 558,572	\$ 92,099	\$ 116,955	\$ 1,068,599	\$ 363,043	\$ 663,659	\$ 290,481	\$ 310,366	\$ 1,627,549

Microsoft as a Registrar

Executive Summary

Despite almost three years of close business partnership between Microsoft and RealNames, Microsoft has never directly sold any of the products and services that RealNames offers. Microsoft has numerous sales channels, each of which would be well suited to selling various of RealNames' products:

- MSN direct salesforce – Direct sales to large corporations
- Microsoft bCentral – Website and direct sales to small businesses
- MSN Website – Automated website sales

If Microsoft were to become a Registrar of RealNames, they receive a much greater percentage (both in terms of cash and revenue) of Keyword sales. Given the competence and scale of Microsoft's sales channels, this could represent a significant amount of cash and revenue for Microsoft.

RealNames Corporation
MSN Registrar

Product

Pricing

Basic	20%	\$ 49
Plus	80%	\$ 498
CKN	Integration	\$ 80,000
CKN	Annual	\$ 70,000

SalesPeople	1	5	10	20	30
Account/Person	1,000	1,000	1,000	1,000	1,000
Monthly Calls	83	83	83	83	83
Average # of KW	50	50	50	50	50
Conversion	5.0%	6.0%	7.0%	8.5%	10.0%
Monthly New Accounts	4.2	25.0	58.3	141.7	250.0
Number of Keywords	208	1,250	2,917	7,083	12,500
Basic Keywords Monthly	42	250	583	1,417	2,500
Keywords Plus Monthly	167	1,208	2,875	7,042	12,458
Annual Basic Revenue	\$24,500	\$147,000	\$343,000	\$833,000	\$1,470,000
Annual Plus Revenue	\$996,000	\$7,221,000	\$17,181,000	\$42,081,000	\$74,451,000
MSFT Share (40% (excluding overall revenue share))					
Annual Basic Revenue	\$9,800	\$58,800	\$137,200	\$333,200	\$588,000
Annual Plus Revenue	\$398,400	\$2,888,400	\$6,872,400	\$16,832,400	\$29,780,400
RN Share (60%)					
Annual Basic Revenue	\$14,700	\$88,200	\$205,800	\$499,800	\$882,000
Annual Plus Revenue	\$597,600	\$4,332,600	\$10,308,600	\$25,248,600	\$44,670,600

10 Salespeople model used in summary model

MSN Keywords

Executive Summary

MSN Keywords is a product concept designed to allow Microsoft to drive significantly higher revenue from high-traffic Keywords. The Internet industry has seen recent successes of revenue models based on variable, traffic-based pricing. The amount of traffic driven by RealNames' Keywords has grown steadily, and Keywords currently deliver over 180 million resolutions (a resolution is the equivalent of a click-through) per month.

Given Microsoft's familiarity and experience with variable pricing, RealNames would look to Microsoft to be the exclusive sales channel for this product. RealNames believes that this product would be highly complimentary to the existing products and services sold by MSN's current sales force.

Pricing Model and Revenue Opportunity

MSN Search receives approximately 630 million searches per month. Through analysis of the search data, RealNames estimates that 460 million of these search queries could be viable Keywords (i.e. not properly formed domain names). 250 million of these queries are searches on 30,000 unique terms. RealNames would be willing to reserve these 30,000 terms for the MSN Keyword product.

Through further analysis of MSN's data logs, as well as its own Keyword resolution data, RealNames estimates that 125 million of the 250 million queries generated by the 30,000 terms originate from the browser command line, as opposed to from search.

Keywords provide a much higher quality visitor than any other form of traffic delivery mechanism. A Keyword user has typed in the specific name, brand or product that they are looking for into the browser command line, and have not simply clicked on a link or banner ad out of curiosity. Based on available information on current traffic-based pricing models, RealNames believes that monetized MSN Keywords can generate an average of \$0.20 per resolution.

The attached model provides a conservative estimate of the percentage of the 30,000 MSN Keywords could be monetized, how that grows over time, and what that represents in terms of revenue to both RealNames and Microsoft.

RealNames Corporation
MSN Keywords

Traffic	460 M
30,000 Top search > 2000 (Browser & Search)	250 M
% Available	54%
Browser	50%
Cost Per Click	\$ 0.20 (available industry data)

Calculation:

Browser Traffic Available for Keywords	125,000,000	125,000,000	125,000,000
Conversion Rate	1%	10%	30%
Monthly Paid Traffic	1,250,000	12,500,000	37,500,000
Annual Revenue	\$3,000,000	\$30,000,000	\$90,000,000
MSN Revenue Share 50% (excluding overall revenue share)	\$1,500,000	\$15,000,000	\$45,000,000
RN Revenue Share	\$1,500,000	\$15,000,000	\$45,000,000

10% model used in summary model

Contact Card

Executive Summary

ContactCard is a new product offering designed to offer Internet users a simple Web presence. This Web presence consists of three elements: a name (Keyword), a digital business card (ContactCard) and messaging services (e-mail, chat, SMS Messaging etc.). A ContactCard can be built within a few minutes by completing a simple Web form and fully personalized using a point and click editor integrated into Internet Explorer.

ContactCard by RealNames provides customers with a comprehensive suite of messaging and community-based functionality that can be accessed either via a Web page, or directly from the browser command line using a keyword interface based on verbs such as “e-mail”, “chat”, and “sms”. Upon purchasing a ContactCard, a customer receives an immediate Internet presence, as well as a number of Keywords to make it easy to navigate to that Internet presence.

The ContactCard application is built upon RealNames Network Naming Service architecture and provides a powerful example of the flexibility and versatility of the Keyword Platform.

ContactCard Keywords

ContactCard owners will receive at least two, and possibly more, Keywords that enable direct and simple navigation to theirContactCard. The first Keyword is a basic keyword beginning with the word “contact” (e.g. contact pepe le pue). The Contact Keyword is allocated on a first-come, first-served basis. The Contact Keyword automatically links to the customer’s ContactCard and cannot be changed. The ContactCard owner also receives a second Contact Keyword in the form of a telephone number. The telephone number is chosen by the ContactCard customer and follows a specific format. For example, a ContactCard phone number Keyword could be +14251234567, with the plus prefix denoting international direct dialing making the context of the Keyword easily recognizable around the world.

Internet Presence

ContactCard provides customers with a simple, easy to navigate web page that can contain as much, or as little information about the customer as desired. The page’s text and graphics can be customized with easy-to-use web-based tools, as shown below

Messaging Functionality

ContactCard provides a full suite of messaging applications that can be accessed not only from buttons on the ContactCard web page, but also directly from the browser command line. For each of the functions listed below, an Internet user would be able to type the action verb followed by the ContactCard Keyword to directly call the messaging function, rather than having to visit the web page and click on a button.

- *E-Mail* – Will bring up the user's default web client and automatically populate the address line with the ContactCard owner's e-mail address
- *Call* – Will launch an IP telephony client pre-populated with the ContactCard owner's phone number
- *Locate* – Will bring up a map of the area around whatever address is chosen by the ContactCard owner
- *Visit* – Will launch a daughter window displaying any other site desired by the ContactCard owner
- *SMS* – Will launch a messaging application that can send SMS messages to the ContactCard owner's wireless device
- *Chat* – Will launch an instant messaging application pre-populated with the ContactCard owner's chat handle

Community-Based Functionality

The product will enable the owner of a ContactCard to share extended information and services with the community of users visiting the card. Shared information will include friends' ContactCards, favorites, calendar and other kinds of information the card owner would want to share with friends and colleagues (movies, restaurants, etc).

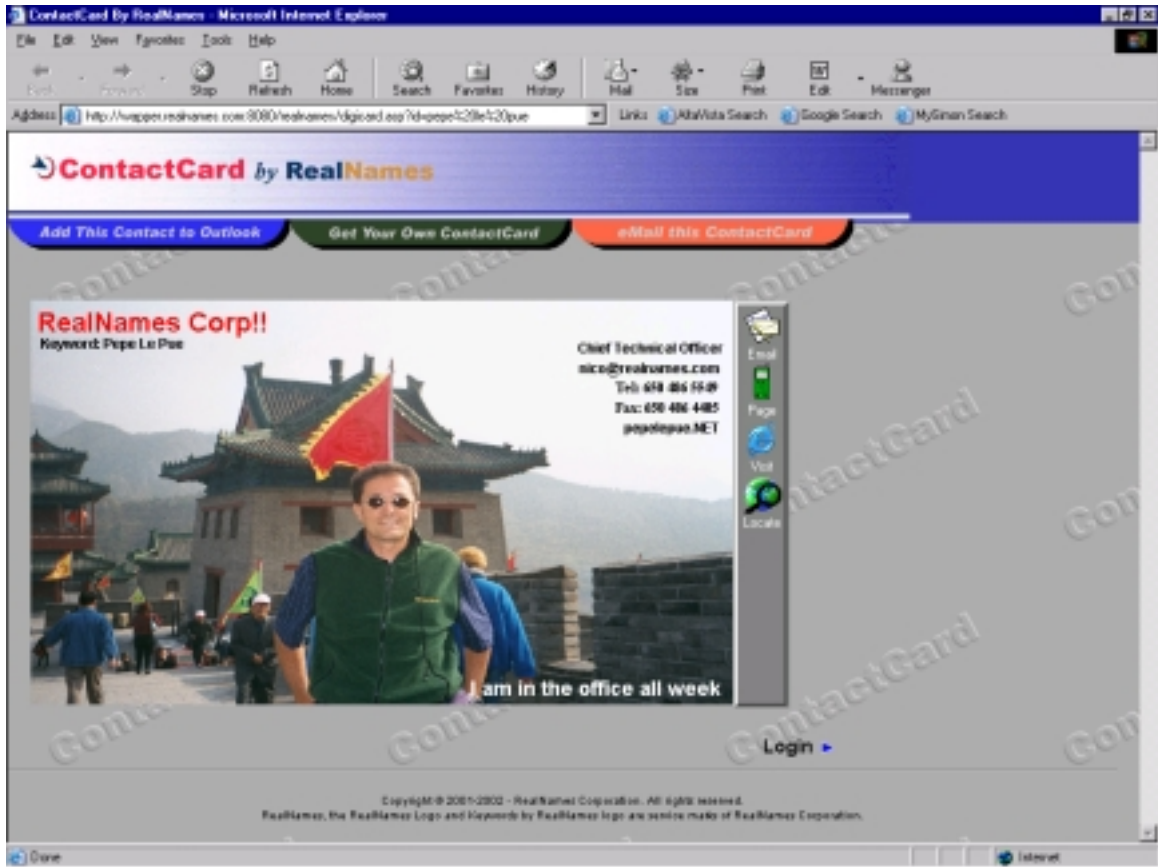
Through a login into their ContactCard, owners will be able to access more complete address book information, in the form of ContactCard links. ContactCard owners will also be able to access personalized services provided by the channel ISP or portal (MY MSN or MY Yahoo types of services). The ContactCard architecture is such that new services can be plugged-in and customized to any distribution partner. For instance, a photo plug-in could be installed that allows the publishing and sharing of photos on line.

Thus, ContactCard is not only a customer's business card on the Web, it becomes a portable address book and a hub for personalized services.

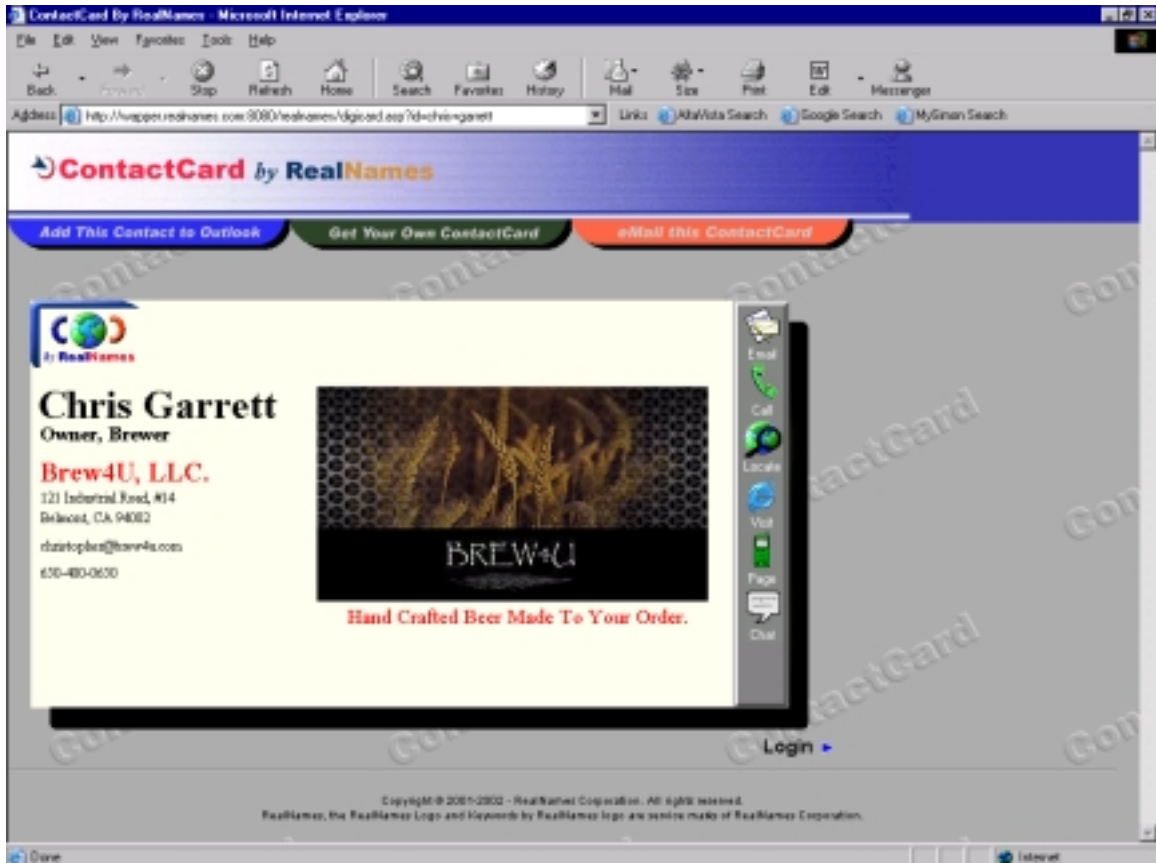
Sales Strategy

ContactCard will be sold both individually to consumers and in packages to small enterprises and large corporations. For individuals, ContactCard would most likely be sold along with other consumer-oriented web services such as e-mail, Internet service and

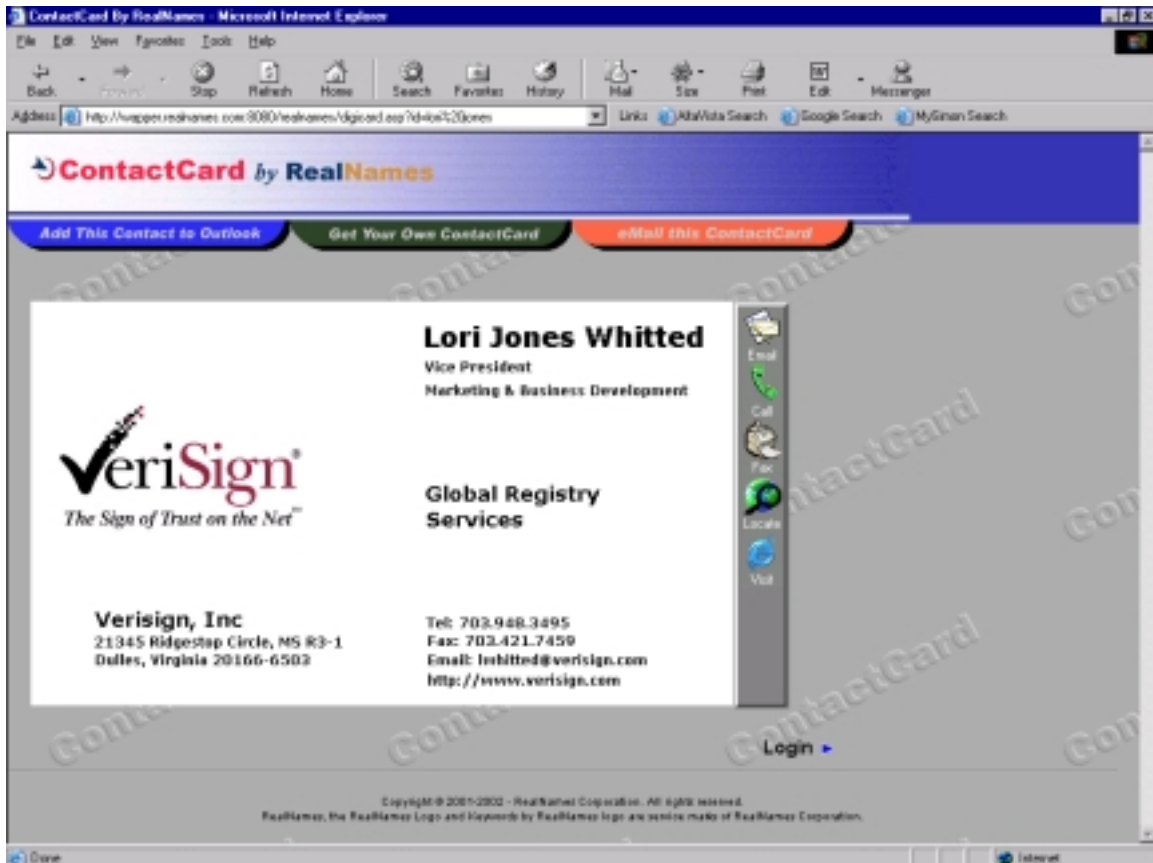
home page web-hosting services. Channel partners that would sell to consumers would include ISPs, user portals, wireless operators and domain name resellers.



Small enterprises will be able to purchase packages of customized ContactCards for their employees. These packages could be purchased separately, or along with other services such as Internet Service, domain names and web page hosting and management services.



Large corporations that want to centralize the administration of many ContactCards will be able to utilize the functionality of RealNames' Corporate Keyword NameSpace product. Corporate Keyword NameSpaces involve the purchase of a Keyword Prefix that defines a Keyword NameSpace that is carved out for the corporation. To access an employee's ContactCard, users would type in the Keyword Prefix followed by the name or phone number of the individual. For example, "VeriSign Lori Jones" would take a user to the following ContactCard:



The RealNames System recognizes the Prefix and routes anything following the prefix to the corporation's back-end system, which handles looking up the data and routing it to the appropriate ContactCard. In this way, the corporation has control over the ContactCards that are entered into its system, and can make changes without having to involve RealNames.

Revenue Opportunity

Current estimates are that ContactCards for individuals and small businesses could generate between US\$5 and US\$10 of revenue per year, either by adding value to a product or services bundle, through advertising or through a subscription fee. This represents a significant revenue opportunity for both RealNames and its distribution partners.

RealNames began releasing information about this product to selected members of its channel approximately two weeks ago. In that time, several of its largest partners (which are some of the largest distributors of any products in the world) have expressed a high level of interest in the product, and a willingness to deploy significant resources to the deployment and promotion of ContactCard.

Contact Card
RealNames Corporation
Contact Card

\$10 MSRP used in summary model

	<u>Jul-03</u>	<u>Aug-03</u>	<u>Sep-03</u>	<u>Oct-03</u>	<u>Nov-03</u>	<u>Dec-03</u>	<u>Jan-04</u>	<u>Feb-04</u>	<u>Mar-04</u>	<u>Apr-04</u>	<u>May-04</u>	<u>Jun-04</u>	<u>Jul-04</u>	<u>Aug-04</u>	<u>Sep-04</u>	<u>Oct-04</u>	<u>Nov-04</u>	<u>Dec-04</u>
Growth		2.5%	2.5%	5.0%	5.0%	7.5%	7.5%	10.0%	10.0%	12.5%	12.5%	12.5%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Cum	100,000	202,500	307,563	417,878	533,710	658,228	792,086	939,329	1,101,297	1,283,511	1,488,501	1,719,116	1,949,730	2,180,344	2,410,959	2,641,573	2,872,187	3,102,801
New Activation	100,000	102,500	105,063	110,316	115,831	124,519	133,858	147,243	161,968	182,214	204,990	230,614	230,614	230,614	230,614	230,614	230,614	230,614
Renewal Rate													50%	50%	50%	50%	50%	50%
Renewed Cards													50,000	51,250	52,531	55,158	57,916	62,259
Expired Cards													50,000	51,250	52,531	55,158	57,916	62,259
MSRP	\$ 10.00	\$ 10.00	\$ 10.00	\$ 10.00	\$ 10.00	\$ 10.00	\$ 10.00	\$ 10.00	\$ 10.00	\$ 10.00	\$ 10.00	\$ 10.00	\$ 10.00	\$ 10.00	\$ 10.00	\$ 10.00	\$ 10.00	\$ 10.00
Cash Collected	1,000,000	1,025,000	1,050,625	1,103,156	1,158,314	1,245,188	1,338,577	1,472,434	1,619,678	1,822,138	2,049,905	2,306,143	2,806,143	2,818,643	2,831,455	2,857,721	2,885,300	2,928,737
MSFT Share 40%	400,000	410,000	420,250	441,263	463,326	498,075	535,431	588,974	647,871	728,855	819,962	922,457	1,122,457	1,127,457	1,132,582	1,143,088	1,154,120	1,171,495
RN Shares 60%	600,000	615,000	630,375	661,894	694,988	747,113	803,146	883,461	971,807	1,093,283	1,229,943	1,383,686	1,683,686	1,691,186	1,698,873	1,714,633	1,731,180	1,757,242

	<u>Jul-03</u>	<u>Aug-03</u>	<u>Sep-03</u>	<u>Oct-03</u>	<u>Nov-03</u>	<u>Dec-03</u>	<u>Jan-04</u>	<u>Feb-04</u>	<u>Mar-04</u>	<u>Apr-04</u>	<u>May-04</u>	<u>Jun-04</u>	<u>Jul-04</u>	<u>Aug-04</u>	<u>Sep-04</u>	<u>Oct-04</u>	<u>Nov-04</u>	<u>Dec-04</u>
Growth		2.5%	2.5%	5.0%	5.0%	7.5%	7.5%	10.0%	10.0%	12.5%	12.5%	12.5%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Cum	100,000	202,500	307,563	417,878	533,710	658,228	792,086	939,329	1,101,297	1,283,511	1,488,501	1,719,116	1,949,730	2,180,344	2,410,959	2,641,573	2,872,187	3,102,801
Activation	100,000	102,500	105,063	110,316	115,831	124,519	133,858	147,243	161,968	182,214	204,990	230,614	230,614	230,614	230,614	230,614	230,614	230,614
Renewal Rate													50%	50%	50%	50%	50%	50%
Renewed Cards													50,000	51,250	52,531	55,158	57,916	62,259
Expired Cards													50,000	51,250	52,531	55,158	57,916	62,259
MSRP	\$ 5.00	\$ 5.00	\$ 5.00	\$ 5.00	\$ 5.00	\$ 5.00	\$ 5.00	\$ 5.00	\$ 5.00	\$ 5.00	\$ 5.00	\$ 5.00	\$ 5.00	\$ 5.00	\$ 5.00	\$ 5.00	\$ 5.00	\$ 5.00
Cash Collected	500,000	512,500	525,313	551,578	579,157	622,594	669,288	736,217	809,839	911,069	1,024,952	1,153,071	1,403,071	1,409,321	1,415,728	1,428,860	1,442,650	1,464,368
MSFT Share 40%	200,000	410,000	420,250	441,263	463,326	498,075	535,431	588,974	647,871	728,855	819,962	922,457	1,122,457	1,127,457	1,132,582	1,143,088	1,154,120	1,171,495
RN Shares 60%	300,000	102,500	105,063	110,316	115,831	124,519	133,858	147,243	161,968	182,214	204,990	230,614	280,614	281,864	283,146	285,772	288,530	292,874

Acquisition Opportunities

Executive Summary

In conjunction with a number of strategies to accelerate the internal, organic growth of its business, RealNames has developed a strategy to acquire or merge with companies where it believes that such a combination would result in one or more of the following:

- Strategic benefits, including, but not limited to:
 - Valuable technological or intellectual property additions to the RealNames platform
 - Removal of competitive pressures
 - Increased visibility and credibility with industry bodies such as ICANN and the IETF
- Significant increases in any of RealNames' current revenue streams through additional sales and distribution capabilities, which could not otherwise be achieved quickly via organic growth
- Addition of revenue streams that RealNames currently does not have, and would be complimentary to the existing business model
- Significant cost savings resulting from a combination

The primary goal of this strategy is to build a combined company with enough revenue, momentum and industry influence to be interesting to the public markets, with the ultimate goal of an initial public offering on the Nasdaq National Market in FY2003.

Advice that RealNames has received from various investment banks, including Investec, Inc. and WRHambrecht + Co, as well as from various public market investment firms, leads the management team to believe that a public offering would be possible if RealNames were to achieve an annual revenue run rate of over \$50 million while maintaining positive cash flows.

These same factors would also make the company more attractive and more valuable to any potential acquirors, and would help establish RealNames as the dominant force in next generation naming.

RealNames' current estimates for organic growth of revenue suggest that, without any acquisitions, an annualized run rate of \$50 million in revenue would be achieved in mid-FY2004. RealNames believes that a combination with two or three synergistic businesses, each with \$3 million to \$10 million in annual revenues, would enable the company to achieve a \$50 million revenue run rate as early as Q2, 2003.

The company has set the following guidelines for evaluating possible acquisition candidates:

- Additional revenue must be highly complimentary to RealNames' existing sources of revenue and overall business model

- Acquisition must be achievable through the issuance of RealNames stock – no immediate cash outlay from RealNames
- Relative valuations must be such that, at the end of all of all acquisition activity, RealNames' current shareholders maintain a majority stake in the combined entity
- All integration efforts must be completed by Q1, 2003
- The combined entity must be cash flow positive

RealNames has identified several acquisition candidates that it believes may be able to meet these criteria. Preliminary discussions and negotiations with most of these candidates have already started, and the company is confident that the goals of the acquisition strategy, as outlined above, are achievable. RealNames has also engaged the services of Investec, Inc., a full service investment bank, to assist with the acquisition process by identifying potential targets, initiating negotiations, conducting valuation analyses and closing any transactions.

Identified Potential Targets

3721

Location: Beijing, China

Employees: Approx. 200

Estimated 2001 Revenue: US\$2 million

Estimated 2002 Revenue: US\$5 million

Estimated Cash Holdings: US\$10 to US\$12 million

Cash Burn Rate: Close to cash flow positive. Should be able to achieve positive cash flows through cost synergies.

Core Business: 3721 is RealNames' primary competitor in China. They provide Keyword and search-related functionality through a plug-in that has achieved significant distribution throughout China.

Rationale for Acquisition:

- *Complimentary Revenue Streams* – 3721's revenue streams are highly complimentary to those of RealNames, with similar product functionality and identical target markets. 3721's sales efforts, both direct and through their own channel would represent a direct salesforce in Asia, similar to the RealNames Registrar that operates in the U.S. In addition, there would likely be strong revenue synergies due to the removal of competition (see below)
- *Asia Operations* – 3721 shares the same core philosophy as RealNames regarding how to build, deploy and grow a naming architecture. RealNames would anticipate keeping the majority of the 3721 organization intact, and 3721 would become the division of RealNames focused on building out the Keyword system in Asia.
- *Competition* – As noted above, 3721 is RealNames' primary competitor in Asia. It is RealNames' belief that the entire Keyword market in China is not as lucrative as it should be for either RealNames or 3721 as a result of the confusion in the market arising from the presence of competing Keyword services. If the companies were to merge, this confusion would be removed, and the entire market for Keywords would grow much more quickly.

- *Technology* – 3721 has several technologies that would augment RealNames’ naming architecture. These include (i) their plug-in which would extend the distribution of Keywords beyond the Microsoft Internet Explorer and MSN browsers; (ii) highly effective fuzzy search capabilities that have already been integrated into many Chinese search engines and (iii) better technology for handling Chinese character sets.
- *Employees* – 3721 has high quality employees, especially within their engineering and product groups, who would contribute well to RealNames as a company.

Current Status of Discussions: We have entered into fairly detailed discussions with 3721 regarding the possibility of an acquisition. They are committed to finding a way for the two companies to work together to consolidate the Keyword market in China, but are not convinced that a full acquisition is necessary.

Netpia

Location: Seoul, Republic of Korea

Employees: Approx. 50

Estimated 2001 Revenue: US\$2 million to US\$3 million

Estimated 2002 Revenue: US\$6 million

Estimated Cash Holdings: US\$500,000 to US\$1 million

Cash Burn Rate: Unknown, but likely to be negative

Core Business: Netpia is RealNames’ primary competitor in Korea. They have developed server-side software that redirects Internet traffic at the IP level. They have established agreements with the largest ISPs in Korea to install this software and take traffic away from MSN Korea and redirect it to Netpia. They use this traffic to provide Keyword navigation services very similar to those of RealNames. In this manner, they have achieved broad distribution of their Keyword service (aggressive estimates are as high as 80+%), while simultaneously decreasing RealNames’ distribution, since RealNames relies on traffic to MSN Korea for its services.

Rationale for Acquisition:

- *Complimentary Revenue Streams* – RealNames and Netpia currently sell very similar products and services, and target similar customer bases. Netpia’s sales efforts, both direct and through their own channel would represent a direct sales force in Korea, which RealNames has not had since it shut down its Korean office in late 2001. In addition, there would likely be strong revenue synergies due to the removal of competition (see below)
- *Competition* – As a result of the traffic redirection by Netpia, RealNames has been at a significant disadvantage in the Korean market. However, RealNames continues to sell more Keywords in Korea than in most other countries. The Keyword market in Korea is very robust, with high user awareness of both Netpia’s system and RealNames’. This user awareness of both systems has resulted in confusion, with some customers buying Keywords for both systems and most buying neither. If RealNames were to acquire Netpia, this confusion would disappear, and Korea could quickly become the largest Keyword market in the world.
- *Benefits to Microsoft and VeriSign* – Both Microsoft and VeriSign benefit significantly if RealNames is successful in acquiring Netpia. Microsoft benefits

because MSN Korea is currently unable to gain any traction, or make any money since the vast majority of its traffic is redirected. If RealNames were to gain control of Netpia's operations, the redirection would cease and MSN Korea would benefit greatly. VeriSign benefits because VeriSign relies on the RealNames system for resolution services for both its Internationalized Domain Names (IDNs) and for its upcoming release of ML.ml addresses. The constraints on RealNames' distribution imposed by Netpia's current operations are also constraining VeriSign in Korea. As two of RealNames' most important partners, it is important to consider the benefits to these two entities.

Current Status of Discussions: RealNames has been in ongoing discussions with Netpia for over a year, with topics ranging from cooperation to outright acquisition. Current discussions have stalled, but this may be an opportune time to re-engage. Netpia recently lost a senior executive (widely considered to have been #2 to the CEO) and it is likely that Netpia is encountering financial problems. This would make an acquisition easier and cheaper to accomplish.

Internet Number Corporation

Location: Tokyo, Japan

Employees: Approx. 130 but should be cut to approx. 70

Estimated 2001 Revenue: US\$1 million

Estimated 2002 Revenue: US\$3 million

Estimated Cash Holdings: US\$175,000 – currently looking for financing

Cash Burn Rate: \$2.7 million

Core Business: Internet Number provides Internet navigation services through the use of number-only Keywords. They distribute access to this service through the use of a plug-in, which they have managed to get installed on a relatively large number of computers in Japan. They have distribution deals with some of the major Japanese computer companies, including NEC and Fujitsu, to install their plug-in as standard on all new computers, and they have a deal with AU (one of the three largest Japanese cell phone carriers) to embed Internet Number's technology into their cell phone software.

Rationale for Acquisition:

- *Complimentary Revenue Streams* – RealNames has recently developed its own plans for number-based Keywords, which it plans to launch in the near future. An acquisition of Internet Number would provide an installed base of such products from which to grow the market.
- *Competition* – Internet Number's products and services are competitive with RealNames' Keywords in that they both allow Internet navigation through the use of simplified addresses. Until now, the products have been different enough that this competition from Internet Number has not represented a significant barrier to RealNames' operations in Japan. However, RealNames' new number-based Keyword offerings will compete directly with Internet Number.
- *Technology* – Internet Number has provided number-based Internet navigation for several years, whereas RealNames is only now developing its strategy. Various aspects of Internet Number's technology would be valuable additions to RealNames platform, including their interface into cellular phones.

Current Status of Discussions: RealNames has approached Internet Number with a range of cooperative ideas ranging from simple business deals to a full acquisition. Internet Number has been receptive to all of the ideas and the companies are currently at the stage of exchanging relevant information.

Keylime Software

Location: Carlsbad, CA, USA

Employees: Approx. 40

Estimated 2001 Revenue: US\$3.5 million

Estimated 2002 Revenue: US\$6 million

Estimated Cash Holdings: US\$1 million

Cash Burn Rate: Approx. \$200,000 per quarter

Core Business: Keylime develops real-time customer intelligence solutions that provide enterprises with a single, unified view of the customer, turning valuable customer insights into actionable knowledge. Keylime's clients leverage this knowledge to optimize every customer interaction, increasing revenues and profitability through improved customer conversion, retention, and loyalty.

Rationale for Acquisition:

- *Diversification of Revenue Streams* – An acquisition of Keylime would RealNames with a new revenue stream that, while different from RealNames' current revenue streams, is a natural extension of RealNames' current products and business model.
- *Highly Complimentary Technology* – Keylime's technology would extend the functionality of RealNames' Keywords, creating a valuable combined package of services. Keywords provide customers with a quick and easy method of delivering Internet users to their web content. Keylime provides valuable tracking and reporting on user activity once that user has been delivered to the web content. Keylime technology would also greatly enhance RealNames' Keyword Management services by enabling RealNames to directly monitor purchasing behavior of Keyword users.

Current Status of Discussions: RealNames and Keylime are engaged in ongoing discussions. Both companies are eager to find a way to work closely together and are currently at the stage of exchanging relevant information.

Easyspace

Location: London, England

Employees: Estimated at 6

Estimated 2001 Revenue: US\$__ million

Estimated 2002 Revenue: US\$__ million

Estimated Cash Holdings: US\$__ to US\$__ million

Cash Burn Rate: Unknown, but expected to have significantly positive cash flows

Core Business: Easyspace is the largest domain name registrar in the United Kingdom and the 12th largest in the world. Easyspace provides a wide range of web hosting packages and domain names, and also offer Keywords.

Rationale for Acquisition:

- *Complimentary Revenue Streams* – Easyspace is one of RealNames’ best and most productive Registrars. By acquiring Easyspace, RealNames would greatly increase the cash and revenue that it receives for each Keyword sale. In addition, RealNames believes that it would be able to improve upon Easyspace’s Keyword sales if it had control. Easyspace’s other revenue streams would add to and diversify RealNames’ revenue, while still being complimentary to consistent with RealNames’ core business model.

Current Status of Discussions: RealNames has an excellent working relationship with Easyspace, but has not raised acquisition discussions to date.

Revenue Opportunity Summary


RealNames Corporation

Recognized Revenue 5 Year Model

Revenue based MSFT executing against all Opportunities

(in thousands)

	FY02 Forecast					FY'03 Model					FY'04 Model					FY'05 Model					FY'06 Model				
	Q1-02	Q2-02	Q3-02	Q4-02	FY-02	Q1-03	Q2-03	Q3-03	Q4-03	FY-03	Q1-04	Q2-04	Q3-04	Q4-04	FY-04	Q1-05	Q2-05	Q3-05	Q4-05	FY-05	Q1-06	Q2-06	Q3-06	Q4-06	FY-06
RealNames Revenue																									
RealNames Base Plan (B1)	\$4,846	\$5,067	\$4,876	\$5,620	\$20,409	\$8,333	\$7,854	\$8,952	\$10,607	\$35,746	\$12,040	\$11,760	\$12,974	\$14,414	\$51,188	\$15,124	\$17,933	\$19,375	\$20,961	\$73,392	\$23,498	\$24,627	\$26,691	\$28,946	\$103,762
Additional Opportunities:																									
MSFT as Registrar (B1A)	-	-	-	650	650	1,300	1,975	2,625	2,700	8,600	2,825	3,050	3,350	3,675	12,900	4,050	4,450	4,900	5,400	18,800	5,925	6,525	7,190	7,900	27,540
MSN Keywords (B2)	-	-	-	940	940	1,875	2,800	3,750	3,850	12,275	4,050	4,350	4,790	5,250	18,440	5,790	6,375	7,000	7,700	26,865	8,475	9,325	10,250	11,275	39,325
Contact Card (B3)	-	-	-	-	-	-	-	460	990	1,450	1,650	2,580	3,390	4,160	11,780	4,350	4,400	4,275	4,125	17,150	4,100	3,925	3,675	3,425	15,125
Total RealNames Rev	\$4,846	\$5,067	\$4,876	\$7,210	\$21,999	\$11,508	\$12,629	\$15,787	\$18,147	\$58,071	\$20,565	\$21,740	\$24,504	\$27,499	\$94,308	\$29,314	\$33,158	\$35,550	\$38,186	\$136,207	\$41,998	\$44,402	\$47,806	\$51,546	\$185,752
Microsoft Revenue																									
MSFT overall Rev Share	\$0	\$0	\$731	\$1,082	\$1,813	\$1,726	\$1,894	\$2,368	\$3,357	\$9,346	\$3,805	\$4,022	\$4,533	\$6,050	\$18,409	\$6,449	\$7,295	\$9,065	\$9,737	\$32,546	\$10,709	\$11,323	\$13,864	\$14,948	\$50,844
Rev Share as % of Total RN Rev			15.0%	15.0%		15.0%	15.0%	15.0%	18.5%	16.1%	18.5%	18.5%	18.5%	22.0%	19.5%	22.0%	22.0%	25.5%	25.5%	23.9%	25.5%	25.5%	29.0%	29.0%	27.4%
Initial Payment/Minimum	-	5,000	-	-	5,000	-	567	-	-	567	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
MSFT Exclusive Rev Share																									
MSFT as Registrar (B1A)	-	-	-	440	440	875	1,300	1,750	1,800	5,725	1,875	2,025	2,225	2,450	8,575	2,700	2,975	3,275	3,600	12,550	3,950	4,350	4,795	5,275	18,370
MSN Keywords (B2)	-	-	-	940	940	1,875	2,800	3,750	3,850	12,275	4,050	4,350	4,790	5,250	18,440	5,790	6,375	7,000	7,700	26,865	8,475	9,325	10,250	11,275	39,325
Contact Card (B3)	-	-	-	-	-	-	-	300	660	960	1,100	1,725	2,250	2,775	7,850	2,900	2,925	2,850	2,750	11,425	2,730	2,630	2,450	2,275	10,085
Total MSFT Revenue	\$0	\$5,000	\$732	\$2,462	\$8,193	\$4,476	\$6,561	\$8,168	\$9,667	\$28,873	\$10,830	\$12,122	\$13,798	\$16,525	\$53,275	\$17,839	\$19,570	\$22,190	\$23,788	\$83,387	\$25,865	\$27,628	\$31,359	\$33,774	\$118,624
MSFT FY Annual Revenue		\$5,000						\$14,231					\$40,787					\$67,733						\$99,471	



White Paper

Nico Popp
CTO
650-486-5549
nico@realnames.com

Table of Contents

Executive Summary

Namespace Data Model

Network Services

Network Architecture

Executive Summary

The Need for a New Naming Architecture

The Internet's naming infrastructure is creaking under the weight of increasing user expectations. Standards bodies such as the IETF and policy bodies like ICANN are awash with discussion of how to "fix" the system on which more than 500 million people rely. Clearly, a new naming architecture that can meet the needs of the public Internet and other major network types is immediately required. RealNames offers this architecture as a contribution towards meeting this requirement.

To understand the naming system that networks need today, it is extremely useful to review why the currently deployed solutions to resource naming and navigation - DNS, URIs and search engines - are lacking.

The limits of the Domain Name System are connected to its roots as a naming architecture invented in 1984, with the primary purpose of giving names to networks and to people as endpoints in email services. The network has changed dramatically since then. DNS's main weaknesses include:

- a) DNS is only able to make use of 7 bit ASCII - 26 characters in the English alphabet and the 10 in the numerical system, 0-9, plus the hyphen (37 total characters), in forming a name. 7 bit ASCII cannot handle foreign characters, creating a significant problem for languages with non-Roman scripts.
- b) DNS cannot guarantee quality of service in delivering content. A DNS resolution points a user to a physical resource and is at the mercy of bandwidth constraints and traffic peaks.
- c) DNS is a poor global naming system. A company with multiple sites worldwide has to give each of them different names [ibm.com; ibm.co.uk etc].
- d) DNS has no inbuilt reporting capabilities. In fact, reporting on DNS traffic is so complex and essential that an industry has arisen to provide the imperfect reports that are available today.

URIs and URLs have weaknesses as well:

- a) DNS gave birth to the URI. These long strings - again restricted to ASCII - allow naming of a wider set of resources. The URI can address individual web pages (with URLs), but the URI can also address people's email address - as in <mailto:person@company.com> - and even their phone number - as in <phoneto:16504865555>.
- b) The URI is a major breakthrough as a means of addressing an unlimited number and type of resources on the Internet, but it is not a naming system. Rather it is a physical addressing system. Naming systems match a physical resource with an alias. A phone number, for example, is simply a memorable (one hopes!) alias to a physical switch address. A DNS name is an alias to an IP number. Physical addresses that are also forced to play the role of names are a bad idea because an identity is then tied to a physical resource identifier. If the resource moves or changes, the name will break. No persistent naming

system for the Web was built, and the URL was adopted as the only available alternative. This is widely accepted to be a huge error.

- c) In addition, the URI is incapable of being human friendly. Home page URLs for well known things barely pass muster as human friendly, intuitive identifiers. <http://www.coke.com> is OK, but how could one expect to intuitively understand that the URL for the US Fish and Wildlife Service is <http://www.fws.gov>.
- d) URLs cannot be consistent pointers to all content across all network access devices Wireless URLs and Web URLs point to different versions of content.
- e) The URL, like DNS, cannot use non-ASCII characters, although it can use a wider set of ASCII characters than the DNS.

Limitations in DNS and URIs spawned search engines – which compensate for the lack of a manageable, human friendly naming architecture for network resources. While they solve a specific and relevant problem, search engines also have weaknesses:

- a) Search Engines can only index “static” web pages on the public network. These are pages with a physical existence on a web server. Today less than 25% of web pages are “static”. Search engines cannot provide pointers to protected content. Similarly, search engines cannot provide access to dynamic content that is refreshed frequently, or content that resides in a content management server or searchable database.
- b) Search Engines employ a “full text index” approach to content. Even with algorithms which attempt to elevate one site above another based on relevance rankings, search engines inevitably find it hard to distinguish between a home page for an entity [a company, a product, a famous person] and a reference to that entity by a third party. Search is great for research but of limited value for navigation.

Requirements for a Modern Naming Architecture

With the problems of the current naming solution in mind, RealNames technologists outlined basic requirements for a modern, ubiquitous naming architecture:

- 1) The architecture must recognize the Internet as a Platform. Microsoft’s .NET and Sun’s SunONE are each attempts to construct a holistic view of the Internet as a total end-to-end network. Naming is a key component of any end-to-end network’s platform – and an essential service in any system of Web services.
- 2) Application developers should be able to access the “naming stack” in the Internet platform to build the naming function for any and every network application. Applications would assign stack-compliant names to resources and users would follow well-understood rules to use the name assigned to retrieve resources.
- 3) The architecture should place a primary emphasis on direct navigation as opposed to search or a facilitated directory service. Research clearly indicates that end users prefer direct navigation to any other means of finding content on networks. StatMarket shows that 52% of Internet users reach sites via direct navigation – up from 46% last year -- while only 7% use search. Similarly, our own research has

shown that 75% of the top 200 EN-US queries are direct navigation queries. The naming architecture must conform to preferred user behavior.

- 4) The architecture recognizes namespaces - a namespace is an arbitrary collection of network resources. Logical organizations of resources are often shaped by the organization of the network itself. A modern naming architecture must seamlessly adopt today's network that often dictates resource organization.
- 5) A complete namespace architecture for the network must reflect the constraint of today's network topology. We have identified four large classes of namespaces in this paper, and offer an overview that demonstrates how the naming architecture would meet the needs of each class's topology.
- 6) A modern naming architecture must seamlessly adopt the current and future networks. In fact, logical organizations of resources are often shaped by the organization of the network itself. Therefore, interoperability and connectivity between namespaces is a key requirement of our namespace architecture as instances of namespaces should be able to connect to other namespaces in order to combine their data and services.

Network Naming Services

In this paper, RealNames proposes an architecture that both meets the stated requirements for a modern naming solution and solves the problems of the current systems. The cornerstone of this architecture is a name and metadata repository that can be accessed through a very small set of core services. The namespace data and service models are detailed in this document.

This architecture has been designed to become an open platform for data owners and application developers. For that reason, the architecture is primarily modeled as a Web service that we call a Network Naming Service. A Network Naming Service is a unique class of Web Service that exposes the data and service model to developers. The Web Service approach makes it easier for developers to integrate network naming within existing applications. It also facilitates the development of new applications when new data sources are made available to the network as instances of Network Naming Services.

We address the topology of namespaces, that is, the different ways in which namespaces are connected to each other on the network. We describe the four large classes of namespaces to show how this namespace architecture can be used to facilitate many different instances of namespaces. We describe how these namespaces can interact with each other for name resolution and discovery.

Because of its flexible yet simple architecture, Network Naming Services will greatly abet access to a wide variety of network resources that are not currently easily found. We also anticipate Network Naming Service instances to be run on a wide range of computer systems (including personal computers that host many of the network's "hidden" resources). Finally, we expect Network Naming Services to play a significant roles in collaborative systems such as P2P and extranets applications that have a great need for a

unified mechanism for registering, accessing and discovering distributed resources such as files, music, people, and more generally, complex network objects.

Once named resource information is exposed as a Network Naming Service, we believe that the development of new applications to manipulate these resources is quite simple. As an illustration, the last section of this document presents examples of important namespace and network applications patterns leveraging the Network Naming Services architecture. These examples include P2P file sharing, enterprise document access and World Wide Web navigation.

Namespace Data Model

The *namespace* is the central component of RealNames' architecture. Simply put, a namespace facilitates the organization of data whose primary means of being accessible is through a *name*. Since data can be clustered according to other properties too, namespaces support the association of a name and some properties as the way of building a handle used to retrieve data.

A namespace contains a set of *resource descriptors*, which are described below. A resource descriptor contains a number of *properties*, or *facets*, describing a physical resource that can be accessed by an application after its descriptor (or a given subset thereof) has been retrieved.

Names

The name is the most important property used to access data. Since many applications present names to their users, the name has a number of constraints:

- It must be multilingual. This is achieved by representing names using Unicode, which supports characters from more than 40 scripts and hundreds of languages.
- It must be usable according to local usage. This means for example that it is desirable to handle the issue of traditional and simplified Chinese characters for Chinese names, without damage to the Japanese or Korean populations.
- Names should be syntax-free, for two reasons. The first is that users should be able to pick the names they want with little constraints. The second is that applications that want to be able to enforce a specific syntax (e.g. telephone numbers, street addresses, etc...) need to be able to do so without being limited by the infrastructure it uses.

Naming Context

Many namespaces organize data not only according to a name but also to some properties. Moreover, they allow a name to appear multiple times, as long as the combination of a name and these properties is unique. The set of properties that bears this constraint is called a *naming context*. Such context is important for a number of reasons:

- It allows an application to refer to resources of identical name but different purposes, simply by adding some distinguishing properties.
- Applications can offer the most granular view of resources they can afford to without burdening the user. This is done by making most of the context *implicit*, i.e. supplied by the application, not by the user.
- Context promotes extensibility of the system, by allowing one to use a given name in contexts that were not anticipated.

The properties that make up the naming context are distinguished as such in the namespace schema.

The context can be viewed as a way of allowing sub-namespaces to exist. Each new property of the naming context further divides the sub-namespaces along a new dimension. Subsets of the context may have conflicting values in different sub-namespaces, but each full tuple of context properties is, by definition, unique within a given namespace. Sub-namespaces do not have a physical reality, they are just a convenient way of seeing the way resource descriptors are arranged in a namespace.

As an example of the use of a naming context, the RealNames Keywords application adds properties for country and service type (used to denote an application or a device). This allows one to refer to different resources using the name “IBM” in different countries and on different devices. When used in a Web browser, or a cellular phone, the name “IBM” is enough because the application is able to supply the country (from the browser settings or the cell phone carrier) and the service type (Web or mobile Web). All the user needs to manipulate is a name, yet the application is more powerful because it allows access to a bigger number of resources.

Resource Descriptors

Each entry in a namespace is a *resource descriptor*. A descriptor contains all the information available about a given resource. It can be viewed as a table associating properties names and types with values.

Properties names are associated to a type and a set of semantics. A number of names and their associated meaning are predefined by the Network Naming Service namespace schema, and can be freely reused, facilitating the creation of interoperable namespaces. The types are types available from the XML schema specifications as well as from the Network Naming Service namespace schema; those types are in the nns namespace in this document.

For example, the following is a resource descriptor:

<i>Property name</i>	<i>Property type</i>	<i>Property value</i>
Name	nns:name	清华大学
Country	nns:country	CN
Language	Language	zh-CN
AddressType	nns:addressType	web
Address	nns:address	http://www.tsinghua.edu.cn/chn/index.htm
Description	String	前身是清华学堂。为一所多科性工业大学，重点为国家培养工程技术人才。

It describes a resource named “清华大学” with a number of properties. In an application where the name is used to retrieve the resource described, the property *Address* will be used to determine the physical address of the resource.

Namespace Schema

The properties of the resource descriptors are dictated by the *namespace schema*. The schema lists the name and type of each property by defining a resource descriptor type, which contains a naming context and additional properties.

The schema for the above sample resource descriptor could be the following:

<i>Property name</i>	<i>Property type</i>
Context	namingContext
AddressType	nns:addressType
Address	nns:address
Description	xsd:string

Where the *namingContext* itself would be:

<i>Property name</i>	<i>Property type</i>
Name	nns:name
Country	nns:country
Language	nns:language

The schema for a namespace is the only authoritative description of the namespace. It can be exposed to, and discovered by, applications that need to use the namespace. Such a mechanism, using Web services, is described later in this document.

Predefined Namespace Data Types

The Network Naming Service defines a number of data types to be used in the definition of namespace schemas in addition to the XML schema data types. These types, or derivation of those types, must be used when defining new schemas. The functionality of the namespace Web services is described in term of these basic types, and the extensibility of the service is achieved by deriving types from these basic types

I. nns:name

A *string*. It is the name by which a resource descriptor is known. Note that different namespaces may treat the name differently by applying namespace-specific normalizations to the name. Since these processes happen within the namespace, an application does not have to know what they are in order to be able to use the namespace.

II. nns:addressType

A *string* identifying the application domain of the address used to get the resource. The address type conditions which applications can manipulate the address. The following strings are reserved and their semantics fixed; if they are used, they must be used according to their specified used.

- **web:** identifies traditional applications that can handle the Web as used on a personal computer or a personal digital assistant able to render unprocessed Web pages.
- **mobile web:** identifies applications such as mobile phone browsers that handle a specialized version of Web pages destined to work with their particular limitations.
- **e-mail:** identifies e-mail applications.

Namespaces can use other strings for their own specific usages. It is expected that the list of reserved strings will be expanded when application categories take advantage of the namespace infrastructure. For that reason, applications should use names starting with an *x-* prefix in order to avoid future conflicts.

III. **nns:address**

An *anyURI* which contains a resource URI as defined by RFC 2396. It denotes the physical address of the resource. This address can be used to access the resource itself. Since not all addresses can be expressed in terms of a URI, especially for applications that use their own protocols, the namespace architecture provides a MIME type, *application/x-rnns-data* which can be used in a data URL, as specified by RFC 2397, to encode application-specific addresses as a URI. Applications are also free to use a different property for their private addresses.

IV. **nns:country**

A *string* which contains an ISO 3166-1 alpha-2 country code. It is used to identify the country to which a name is scoped, if country-based scoping is used in a namespace (i.e. if the naming context contains a country code). It is not used to identify the country of the resource itself.

V. **rnns:language**

The property *rnns:language* is of type *rnnsd:languageCode*. It is used to identify the language to which a name is scoped, if language-based scoping is used in a namespace (i.e. if the naming context contains a language code). Typically this language matches the language of the resource itself, as one aims to present users of a given language with resources in that language.

The following is the XML Schema defining the basic types used for describing namespaces.

```
<?xml version="1.0" encoding="utf-8" ?>
<xs:schema id="NNSSchema"
  targetNamespace="http://localhost/WebService1/NNSSchema.xsd"
  elementFormDefault="qualified"
  xmlns="http://localhost/WebService1/NNSSchema.xsd"
  xmlns:mstns="http://localhost/WebService1/NNSSchema.xsd"
  xmlns:xs="http://www.w3.org/2001/XMLSchema" version="1.0">
  <xs:simpleType name="name">
    <xs:restriction base="xs:string"></xs:restriction>
  </xs:simpleType>
  <xs:simpleType name="country">
```

```
        <xs:restriction base="xs:string">
            <xs:length value="2" />
        </xs:restriction>
    </xs:simpleType>
    <xs:complexType name="namingContext">
        <xs:sequence>
            <xs:element name="Name" type="xs:string" />
        </xs:sequence>
    </xs:complexType>
    <xs:complexType name="resourceDescriptor">
        <xs:sequence>
            <xs:element name="Context" type="namingContext" />
            <xs:element name="AddressType" type="xs:string" />
            <xs:element name="Address" type="xs:anyURI" />
        </xs:sequence>
    </xs:complexType>
    <xs:element name="ResourceDescriptor"
        type="resourceDescriptor"></xs:element>
    <xs:simpleType name="addressType">
        <xs:restriction base="xs:string" />
    </xs:simpleType>
    <xs:simpleType name="address">
        <xs:restriction base="xs:string" />
    </xs:simpleType>
</xs:schema>
```

Namespace Service Model

The previous section defines the namespace data model: a metadata repository for named resource descriptors. This namespace data model naturally leads to the definition of a service model. The service model specifies the primitive services and the interfaces that they make available to a namespace client to provide interaction with the namespace data.

Abstract Service Model

The namespace service model includes three core services. These three core services are the registration service, the lookup service and the search service. Each core service is itself defined by one or many primitive interfaces that are introduced below.

Registration Service

The namespace registration service allows the registration and management (update, deletion) of all names and resource descriptors. The registration service defines three abstract interfaces.

- **Create:** This interface creates a new named resource descriptor. The “Create” interface takes a resource descriptor as input and returns zero or one resource descriptor based on success. Concrete implementation of this interface must implement the main integrity rule of a namespace that enforces the uniqueness of a name within the namespace-defined naming context.
- **Update:** This interface updates an existing resource descriptor. The “Update” interface takes a resource descriptor as input and returns zero or one resource descriptor based on success.
- **Delete:** This interface deletes an existing resource descriptor. The “Delete” interface takes a resource descriptor as input and returns zero or one resource descriptor based on success.

Lookup Service

The namespace lookup service allows the resolution of a name and its naming context into a complete resource descriptor. The lookup service defines one unique abstract interface called “Lookup”.

- **Lookup:** The parameterized input of the “Lookup” interface is name and naming context. The output is zero or a single resource descriptor. Concrete implementations of this interface require that the name and naming context be fully specified in order for the namespace to retrieve the unique resource descriptor associated with them (otherwise an empty result set is returned).

Search Service

The namespace search service allows the discovery of names and resource descriptors. The search service defines one unique abstract interface called “Search”.

- Search: The parameterized input of the “Search” interface is a query expressed using the Keyword Query Language (KLS). The Keyword Query Language is described in detail in the next section. The parameterized output for the Search interface returns zero, one or many resource descriptors matching the search query.

The Keyword Query Language

The Keyword Query Language enables search queries in the form of logical conjuncts expressed as a sequence of key-value pairs. The keys of the search expression are the property names of the resource descriptors as defined by the namespace schema. Therefore, a search query can be expressed in terms of the name, naming context and any other properties of the resource descriptors for that namespace. A search expression can be a conjunct of one, many or all the properties of a resource descriptor:

The Keyword Query Language has been designed to become a command-line like search interface that can be used both by client applications and Web services as well as end users. Therefore, KLS allows both simple and advanced search queries to be specified as a structured search keyword. The syntax of the Keyword Query Language goes like this:

property1:value1 property2:value2...propertyN:valueN

“Property” is a property of the namespace resource descriptor, and “value” is a static value of a field or a search filter for that field value (see next section on search filters). A key benefit of the KLS is to allow user to input elaborate search queries even though the search user interface may be reduced to a single text field (e.g. browser address bar). The KLS syntax is extensible to all namespace schemas (specific resource descriptor properties simply map into a new search field).

The following KLS query string expresses a search for all resource descriptors whose name is like the string ‘sony’, expressed in the ‘en-us’ language for the country ‘us’ and whose description contains the word ‘electronics’. Note that the “name” property is the default search field when omitted and that these fields are space-separated.

“sony language:en-us country:us description=electronics”

The constrained form of a KLS search expression reduces all search queries to a logical AND operation across the stipulated properties. Since the search service interprets the input as a logical AND, only OR operations across resource descriptor properties are supported. However, note that OR operations on value of a single property are possible through search filters as explained below.

Search Filters

In the search query, field values are typically expressed as static values. However, advanced queries are also possible using search filters. In fact, search filters allow the construction of search expressions on property values. Search filters are kept simple by design and therefore restricted to the main logical operators.

- Logical OR. This logical operator is represented as ‘|’ in a search filter.
- Logical AND. This logical operator is expressed as ‘&’ in a search filter.
- Logical NOT. This logical operator is expressed as ‘!’ in a search filter.

Search filters also implement the wild card character to search for presence (any value) or perform elaborate substring matches. The wild card is represented by the character ‘%’ in a search filter.

For example, the following query searches for all names starting with the prefix “sony” or ending with the suffix “vaio”.

```
name:'sony% | %vaio'
```

Concrete Service Model - Namespaces as a Unique Class of Web Service

In this section, the abstract namespace service model outlined above is presented using the Web service formalism. This formalization leads to the representation of our namespace architecture as a new and unique class of Web service that we call Network Naming Service.

A Network Naming Service instance is fully defined through its Web Service Description Language (WSDL) document. As a Web service, a concrete instance of a Network Naming Service must define and implement all three concrete ports corresponding to the registration, lookup and search services. Each interface of each core service is mapped into a port operation in the WSDL. Operations have input and output messages. These messages are resource descriptors expressed as XML messages using the namespace schema and XML namespace definition derived from the namespace data model. As explained before, a specific instance of a namespace can extend the core data schema. If so, these extensions will appear within the primitive operation messages for the WSDL of the Network Naming Service.

However, the WSDL representation of a Network Naming Service also allows us to go one step further by moving from the purely abstract model of core service and interface towards the concrete definitions of a namespace implementation. In fact, using the WSDL binding element, namespace implementers can fully specify the transport protocols, message representation, and physical port address that they have chosen for developing and deploying their namespace. In other words, the WSDL is both a formalism to express the namespace architecture presented in this document as well as the authors preferred mechanism for exposing its concrete implementation to a client application.

Namespace Operation Composition

One fundamental component of our namespace architecture is the ability to compose namespace operations across multiple namespaces. Composition is particularly relevant to lookup and search operations across distributed namespace instances as illustrated in this paper.

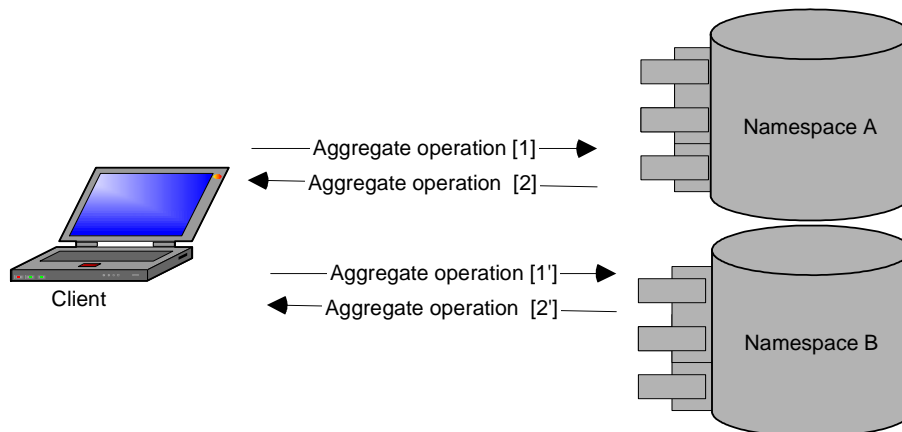
The next section identifies and defines three main composition patterns for namespaces: the aggregation, routing, and referral patterns.

Aggregation

The aggregation pattern occurs when a client application sends an operation input message to multiple namespaces at the same time and aggregates their respective results set into a larger one representing the sum of each individual namespace output.

Aggregation requires that every namespace operated upon have a similar data model or that the client is able to map all the different namespace descriptors into a common schema.

The aggregation pattern is particularly relevant in a network topology where resource descriptors cannot be found in one unique centralized namespace, but instead are distributed across multiple namespace nodes, which is generally the case for P2P networks.




Chaining

The chaining pattern occurs when a namespace decides to forward a request to another namespace with or without processing that request against its own data repository.

All namespace operations can be composed through chaining. Note that the chaining pattern does not require all namespaces in the chain to have the same data model. However, in the case of distinct data models, any namespace node in the chain must be able to map the query resource descriptor that it receives into a resource descriptor compatible with the data model of the node that it sends to. Note that a namespace node

is only expected to return a response resource descriptor in accordance with its own schema.

The chaining pattern is especially useful to implement a namespace proxy. Therefore, one application of the chaining pattern is query caching geared toward improved performance (i.e. local caching server). Another use of a namespace proxy is to provide namespace adapters. The proxy namespace “adapts” the query to the data model of another namespace without requiring the client application to have to understand and handle another namespace schema.

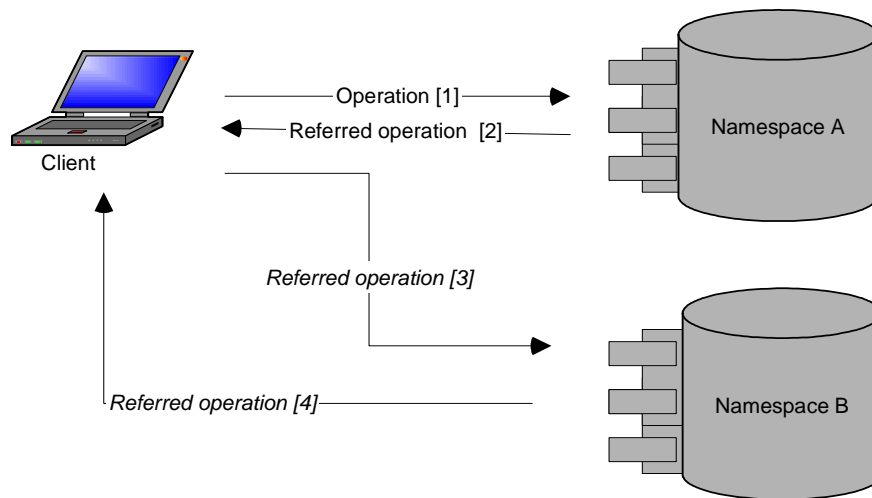


Referral

The referral pattern allows a namespace to redirect an operation to another namespace. Unlike the chaining pattern, in the referral pattern, the forwarding of the operation occurs through the client and not the namespace service. Since a namespace operation is implemented through a concrete Web service port, the referral pattern requires that the Web service port be implemented through a protocol capable of message redirection. This places a strong requirement on the transport protocol. Note that such capabilities are already an integral part of the HTTP and SMTP protocols and under development for SOAP.

When a referral occurs across two namespaces whose data model differs, the referrer namespace must translate the input message into a message compatible with the schema (resource descriptors) of the target namespace. This is not an unrealistic requirement, as one would expect a namespace to only refer queries to a namespace with similar scope and, at a minimum, to a namespace whose data model it understands. This constraint has the advantage to remove the application from the burden of mapping namespace schemas.

Referral is expected to be the most common composition pattern for namespaces. For example, using the referral pattern, a local namespace running on the user local machine as a Web service can delegate the user query to an enterprise namespace serving the corporate Intranet. In turn, the corporate Intranet namespace can cascade the query to a larger namespace serving the global Internet.



Simple Security Model

Confidentiality, Integrity and Authentication

Message confidentiality and integrity are handled using the existing security mechanisms from the transport layer. When needed, the Secure Socket Layer (SSL) along the Transport Security Layer (TLS) is used to provide confidentiality for all namespace operation messages.

For authentication, a namespace may implement token-based authentication such as username and password. The Network Naming Service namespace defines a username and password element so that these authentication tokens can be passed within the SOAP header of a namespace operation message. More elaborate authentication schemes such as XML signatures and X509 should also be possible. For those, we refer the reader to the emerging WS-Security standard specification that describes how to include and encode signature and encryption headers as part of SOAP messages.

Access Control Roles

A namespace defines two types of resource descriptor and properties: public and private. Public resource descriptors and properties can be viewed by all users. Therefore, namespace content viewing will generally not require any user authentication. However, a namespace can also store private resource descriptors and private properties. In that case, content viewing will require user authentication.

A namespace defines four distinct user roles. These roles are: NamespaceAdministrator, NamespaceWriter, PublicNamespaceReader and PrivateNamespaceReader. Note that since the registration operations are stateless, when required, credentials must be passed with each operation request.

The NamespaceAdministrator is a super user with access rights to all namespace operations across all resource descriptors (private and public). This role has the highest level of privileges and requires authentication.

A NamespaceWriter is an authenticated role. NamespaceWriters can create new resource descriptors. NamespaceWriters can read, update and delete all the resource descriptors and properties that they own (but only the one they own).

The PublicNamespaceReader role describes an unauthenticated use of a namespace. All public resource descriptors and public properties can be viewed by any user registered as a PublicNamespaceReader. Private resource descriptors and properties are not accessible to PublicNamespaceReaders. The PublicNamespaceReader role provides a simple mechanism to allow content viewing without requiring any authentication.

The PrivateNamespaceReader role is an authenticated role that permits read access to all private resource. This role enables the creation of private namespaces where content viewing requires user authentication.

Namespaces and Network Architecture

Logical organizations of resources are often shaped by the organization of the network itself. A modern naming architecture must seamlessly adopt the current and future networks, which often dictates resource organization. As mentioned previously, a namespace contains a set of resource descriptors that contain many properties that describe physical resources. This collection of resource descriptors can be organized based on many different facets or properties that make the namespace interesting.

One such property is the name for the physical resource. Names in our system generally map to addresses for resources on the network. These are addresses for resources such as objects on the Web, documents or files on the network and so on. However, we leave the underlying network-layer addressing and routing of the IP architecture unchanged. We layer our services as an overlay network over the IP-based network. Name resolvers can make use of many namespace properties in order to locate resources on the network. Since the organization of the network often influences organization of resources, interoperability between namespaces becomes an important feature of the system.

This section discusses the topology of namespaces, that is, the different ways in which namespaces are connected to each other on the network. We shall first describe four large classes of namespaces that show how this namespace architecture can be used to facilitate many different instances of namespaces. We shall then describe how these namespaces can interact with each other for name resolution and discovery.

Types of Namespaces

Personal Namespace

Each user can name and reference a set of resources specific to her/him. This namespace typically resides on the user's device, giving the user complete control over resources of interest to him. Instances of resources include document files, frequently-accessed web pages, e-mail addresses, images and so on. The personal namespace is very similar to the Favorites list in the IE browser in that it helps the user create short cuts for resources on the network or his own hard drive. For example, the user can name a file and map the location of the file on his hard drive to the name, as shown below. The file is then be referenced by its user-assigned name.

<i>Property name</i>	<i>Property type</i>	<i>Property value</i>
Name	nns:name	Expense Report
AddressType	nns:addressType	Document
Address	nns:address	C:\My documents\expense.xls
Description	String	This is the expense report for my business trip

Resolution of names is done at the local level. The name resolver checks the user's machine for a matching entry in the user's personal namespace and returns the address of the resource if successful. At the same time, if a user wants to know which names are

stored in the repository, he can query the repository for a list all the resources and names that are assigned to each of those resources. All services mentioned above are available to this namespace.

Enterprise Namespace

Such a namespace typically deals with resources accessible to a limited community of users such as a corporate Intranet. The Enterprise Namespace is a collection of resources that can be shared amongst groups of workers within a department or between employees across departments. Resources are named and referenced by their assigned names. The enterprise namespace then stores resources that are sharable but not in a user's personal namespace. The enterprise namespace repository is usually hosted on corporate Intranet servers. For example, there might be shared documents between different design groups or a repository of employee contact information that is accessible by user name. The IT manager or system administrator generally manages the repository. Registration of documents within the enterprise namespace is not automatic. Resources must be submitted to the namespace administrator or IT manager for inclusion in the namespace along with access permissions.

<i>Property name</i>	<i>Property type</i>	<i>Property value</i>
Name	nns:name	401k Forms
Country	nns:country	US
Language	Language	en-US
AddressType	nns:addressType	Document
Address	nns:address	\\HR\enroll\Documents\401k.pdf
Description	String	This is the company 401K enrollment plan

Access to the namespace is generally limited to company employees and is determined by the roles assigned to each user. This paper defines the different access roles that could be assigned to users. For example, the system administrator will be assigned a NamespaceAdministrator role for obvious reasons and all users in a dept could be assigned a PrivateNamespaceReader role. However, this does not prevent other implementations of this scheme where a separate authentication policy might allow finer control over access to some or all of the namespace.

Enterprise namespace resolvers handle resolution of names in the enterprise namespace. The enterprise namespace resolution services behave similarly to other namespace resolution services, except the scope of name resolution is limited to names within the enterprise namespace. We do envision having many different namespaces and resolvers within any enterprise. These namespaces may choose to communicate with each other with simple rules for query resolution. If a resolver cannot fulfill a resolution request, it may forward resolution requests to other resolvers (namespaces) that it might know of either through the process of referral or chaining.

Global Namespace

These namespaces have the scope of the Internet. RealNames has built and commercialized an example of a public global namespace. The global namespace repository handles many different namespaces (for example, digital business cards, Instant Messaging handles, Mail clients), but generally the scope of resolution is much broader than that of the enterprise namespace, although commercial interests may govern it. This namespace is different from the enterprise namespaces in that all Internet users can discover and use names. However, registration of names is not automatic and commercial entities could charge for registering a name.

<i>Property name</i>	<i>Property type</i>	<i>Property value</i>
Name	nns:name	Hotmail
Country	nns:country	US
Language	Language	en-US
AddressType	nns:addressType	web
Address	nns:address	http://www.hotmail.com
Description	String	This is a web based e-mail service

Records in this namespace repository embody many or all of the properties described previously. All namespaces services described are implemented in the global namespace resolvers. This allows users to discover namespaces, discover resource name-to-address mappings and user-transparent location of resources or direct navigation based on a single matching record in a particular namespace.

P2P Namespace

A user machine can host a set of resources that the user can name and publish to other peers on the same P2P domain. The types of resources that are hosted in this namespace are determined by the P2P application. For example, file-sharing applications allow naming of files, while instant messaging systems allow naming of instant messaging handles for users. The implementation of the P2P system will determine how the resources are discovered and accessed. We will not get into the details of the implementation of the P2P system since the focus of this paper is architectural support for naming and resolution of named resources. Registration of resources into the namespace repository is similar to that of the personal namespace. The following table shows an example of a resource descriptor to an audio file, “ray of light.avi” available on the user’s public folder.

<i>Property name</i>	<i>Property type</i>	<i>Property value</i>
Name	nns:name	Ray of Light
AddressType	nns:addressType	audio file
Address	nns:address	http://www.myserver.com/Public/Madonna/ray of light.avi
Description	String	Madonna's CD

This namespace offers all core namespace services: registration, resolution and discovery of shared resources. This namespace is a bit different from other namespaces described above in that it participates in a particular application where resources are shared amongst many hosts on an open network. The namespace administration is not under the control of a central administrator; rather it is under the control of an individual, much like a personal namespace.

Resolution Services Across Namespaces

Our namespace architecture allows instances of namespaces to connect with each other and provide a rich resolution service to the users. Resolutions across namespaces are influenced by resolution policies defined by each application. The resolution policy involves use of other namespaces in order to successfully fulfill name resolution requests. The order in which these namespaces are queried is called a resolution cascade. The resolution cascade is initially set up at application installation time. As noted previously, the process of forwarding a query from one namespace resolver on the resolution cascade to the next one, through a client, is called referral. The decision of whether or not to refer resolution requests is a matter of resolution policy. The following table shows a resolution cascade listing. The resolution cascade is the local namespace followed by the enterprise namespace and then the RealNames namespace.

<i>Namespace</i>	<i>Address</i>
Local	C:, D:
Enterprise	http://www.ens.com
RealNames	http://resolver.realnames.com

It is possible for a user to maintain a list of namespace resolvers to which the client sends a request in order to resolve a name. For example, as shown in figure 1, the user, would like to retrieve a project document and so forms a resolution query named ("Project Eureka").

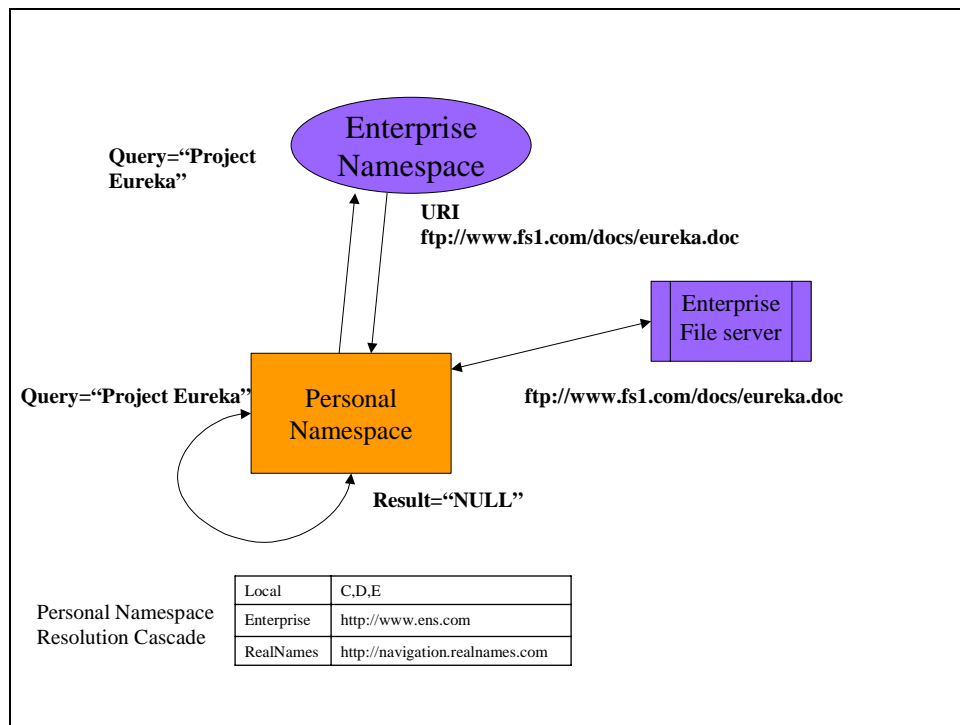


Figure 1

The application first checks the user's resolution cascade, selects the first entry on that list and sends a query to that namespace in order to find a matching record. In this case, the first entry on the resolution cascade happens to be the user's personal namespace. If a matching entry is found in the personal namespace, then the resolution request completes successfully (i.e., the file is available). If the personal namespace does not contain a matching entry, the client sends a query to the enterprise resolver, which is next in the resolution cascade and then to the RealNames global repository if required. Figure 1 shows that a matching resource descriptor was indeed available in the enterprise namespace.

The user can add other namespace resolvers to the resolution cascade as he becomes aware of them. By default, each resolver checks its own namespace before forwarding the request to the next resolver in the cascade. So a user probably will want to register a name in the local namespace once it has been retrieved from the network. Applications that implement this scheme will provide such facilities to the user and the application's resolution process will automatically follow the order of resolution that the user defines.

By manipulating the resolution cascade, the user has control over how he/she wants to implement the resolution policy. In fact, the user may decide to always refer resolution requests to a known resolver and let that resolver refer further in its own resolution cascade. Since all namespaces have a common query interface, the namespaces are interoperable. Given a resource name, each namespace resolver returns the location of the resource or URI if successful; otherwise, it may chose to forward the resolution request. This is shown in figure 2 below. In the figure, the client makes a query "contact Keith" and first checks the local namespace for a match. This is based on the resolution cascade on the client. Since it does not find a matching value in the resource descriptor, the client

refers to the next resolver in the resolution cascade, resolver 1. Resolver 1 then follows the same process and refers to Resolver 2 and so on.

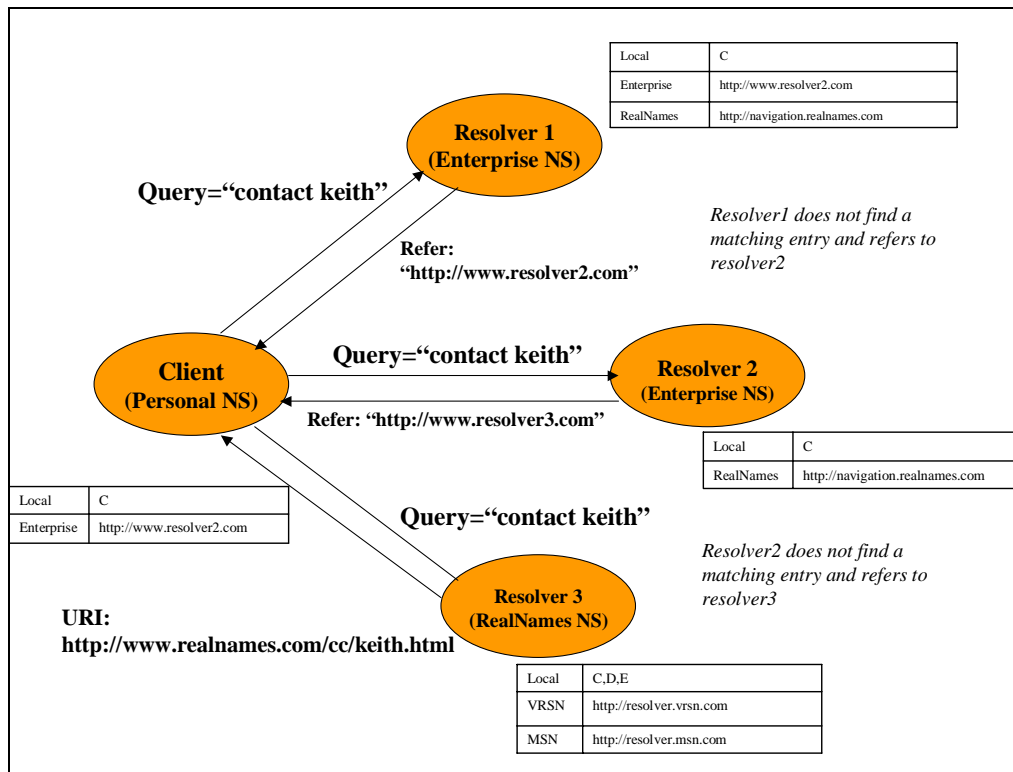


Figure 2

Finally, the RealNames' resolver finds a match and retrieves the address and sends the result back to the client.

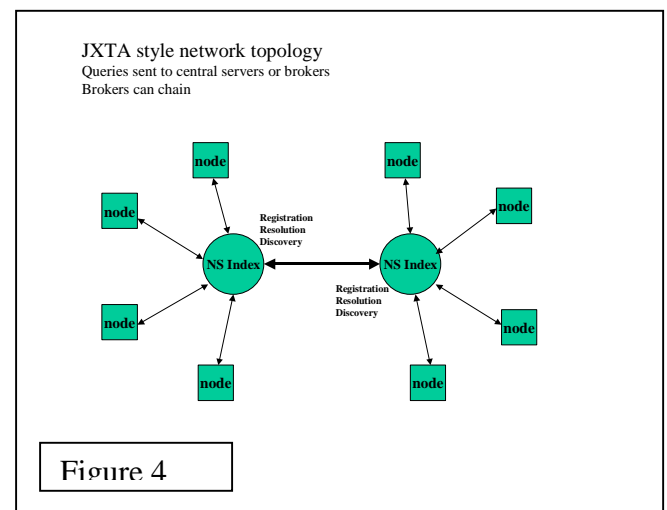
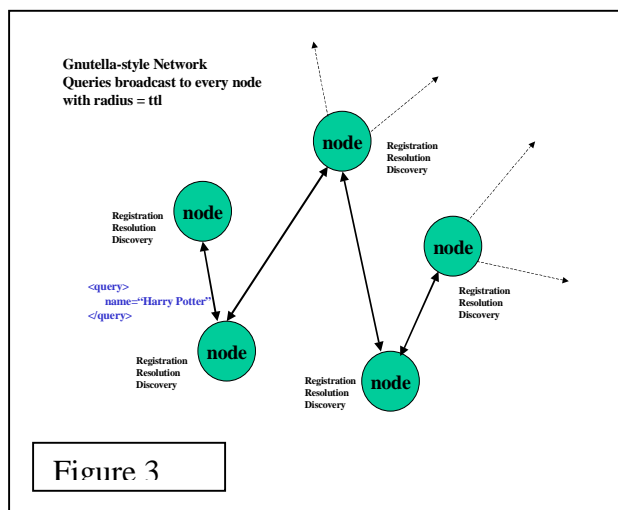
User-defined policy decisions such as the one mentioned above are not binding on applications. For example, the IE browser may decide to just refer to the RealNames global namespaces in order to resolve names and completely ignore the user defined resolution cascade. Of course, the RealNames global namespace resolver may decide to chain resolution requests following its own resolution cascade and still provide the user with a resolution service that works across namespaces. Alternatively, other network-aware applications, may implement a policy whereby they always follow the user-defined resolution cascade first and then query a known resolution service if that is not part of the user's resolution cascade.

P2P Systems

We shall now describe how P2P applications can use the namespace architecture in order to name shared resources and discover them. P2P systems facilitate resource sharing across a large number of hosts in the network. Interoperability between namespaces

distributed on the peer-to-peer network can be used to enable new and different applications.

At present, there are many peer-to-peer network topologies over which different services have been built. The namespace architecture described here is agnostic to the network topology used in peer-to-peer applications. We have provided two diagrams below in order to illustrate our point. Figure 3 shows a Gnutella style decentralized network topology. All nodes in such a network run core namespace services. A query originating from a node propagates through the network based on each node's connectivity and is limited by its TTL (or time-to-live). As each node receives the query, it checks the resolution cascade and either returns a result or forwards the request to other nodes (through the process of chaining) on the resolution cascade. Figure 4 shows a network topology where central servers host namespaces and each node on the network registers its names with the central server.



Nodes in the network only talk to the central servers to which they are allowed to talk. Central servers run core namespace services and query other central servers (or hubs or brokers). This style of network topology is used in the JXTA architecture.

Nodes or peers on the network in a P2P system have a local repository, which hosts resources that the user would like to share with other members of the P2P network. All nodes in a P2P network use the same search/discovery semantics as described previously. When a user wants to search for a resource (a file, for example) on the network he creates a search query first using the Keyword Query Language. The client examines the user defined resolution cascade on the user's system and sends the query to the first namespace on that list. That host in turn will forward the request to the list of hosts that are on its resolution cascade and so on. Search results are returned to the originating host from every host that is queried. The P2P application running on the original client can then aggregate the results and display them.

The search semantics that we have described are primitives for search and do not address issues of request routing, replication etc. This is left to the application, which can determine whether or not to send the search request to all the hosts currently available on the network, or to a subset of such nodes in its vicinity. Peers in a P2P network enter and leave the network at random. The application also needs to keep track of which hosts on

the resolution cascade are available at a particular instance and then chain requests appropriately. The namespace architecture does not address such issues but it provides a common set of name resolution and discovery services upon which this functionality can be built.

In order to design an effective peer-to-peer application, many different issues need to be addressed. Amongst them are issues such as request routing, authentication, efficient data location and storage, redundancy, search and so on. We have not described implementation of all these services. But we have shown how our naming and name resolution architecture can support location of distributed resources, a fundamental requirement for any peer-to-peer application. Other functionality of a peer-to-peer system can be layered on top of this.

Minimum Assurances Agreement

Jenny Cox
Interim General Counsel
650-486-5696
jcox@realnames.com

May 7, 2002

RealNames Corporation
150 Shoreline Drive
Redwood City, 94065

Microsoft Corporation
One Microsoft Way
Redmond, WA 98052

Re: Business Relationship

Ladies and Gentlemen:

RealNames Corporation (“RealNames”) and Microsoft Corporation (“Microsoft”) have entered into the License and Marketing Agreement dated March 13, 2000 as amended on April 13, 2001, September 28, 2001, and March 13, 2002 (the “Agreement”) and, in connection therewith, RealNames has issued four Senior Promissory Notes to Microsoft dated July 15, 2001, October 15, 2001, January 15, 2002 and April 15, 2002 (collectively, the “Notes”). The Agreement expires June 28, 2002. The parties are currently discussing mutually agreeable terms for a continuation of the business relationship. In the interests of permitting sufficient time for such discussions the parties wish to enter into this binding letter of intent.

[Insert outline of specific terms, if any, during meeting]

The parties agree to continue negotiating in good faith towards a definitive agreement for a renewal of the Agreement until the effective date of a renewal of the Agreement (the “Discussion Period”).

The parties further agree that, during the Discussion Period, they will refrain from initiating any litigation or other legal action or proceeding against each other on account of any claims relating in any way to the Agreement or the Notes.

The terms of the Nondisclosure Agreement between the parties dated June 2, 1999 are incorporated herein by reference. In no event shall either party be liable to the other party in connection with this letter of intent or the definitive agreement for any consequential, incidental, indirect, special, or reliance damages (or any loss of revenue or profits), however caused, whether for breach of contract, negligence or under any other legal theory, whether foreseeable or not and whether or not the other party has been advised of the possibility of such damages. Any notice required to be delivered hereunder shall be delivered via federal express to the addressee or signatory hereto, as the case may be.

If this letter is acceptable to Microsoft, please indicate by signing below.

RealNames Corporation

By: _____

Title: _____

Accepted and Agreed:

Microsoft Corporation

By: _____

Title: _____

May 7, 2002

RealNames Corporation
150 Shoreline Drive
Redwood City, CA 94065

Microsoft Corporation
One Microsoft Way
Redmond, WA 98052

Re: Business Relationship

Ladies and Gentlemen:

RealNames Corporation (“RealNames”) and Microsoft Corporation (“Microsoft”) have entered into the License and Marketing Agreement dated March 13, 2000 as amended on April 13, 2001, September 28, 2001, and March 13, 2002 (the “Agreement”) and, in connection therewith, RealNames has issued four Senior Promissory Notes to Microsoft dated July 15, 2001, October 15, 2001, January 15, 2002 and April 15, 2002 (collectively, the “Notes”). The Agreement expires June 28, 2002. The parties have discussed various outcomes to their current business relationship including mutually agreeable terms for a termination of the business relationship.

In order for RealNames to explore its opportunities independent of the business relationship, the parties have agreed that Microsoft will not seek payment on the Notes. The parties expect to execute the final documents including a release of the Notes by June 28, 2002.

Until the execution of final documents as contemplated by this Letter Agreement, the parties agree to refrain from initiating any litigation or other legal action or proceeding against each other on account of any claims relating in any way to the Agreement or the Notes.

The terms of the Nondisclosure Agreement between the parties dated June 2, 1999 are incorporated herein by reference. In no event shall either party be liable to the other party in connection with this letter of intent or the definitive agreement for any consequential, incidental, indirect, special, or reliance damages (or any loss of revenue or profits), however caused, whether for breach of contract, negligence or under any other legal theory, whether foreseeable or not and whether or not the other party has been advised of the possibility of such damages. Any notice required to be delivered hereunder shall be delivered via federal express to the addressee or signatory hereto, as the case may be.

If this letter is acceptable to Microsoft, please indicate by signing below.

RealNames Corporation

By: _____

Title: _____

Accepted and Agreed:

Microsoft Corporation

By: _____

Title: _____